

Investigating the use of 3D digitisation for public facing applications in cultural heritage institutions



A thesis submitted to University College London (UCL) for the
degree of Doctor of Engineering (EngD)

John Hindmarch

UCL Civil, Environmental and Geomatic Engineering

London 2015

Declaration of ownership

'I, John Hindmarch, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.'

Signed _____

Date _____

Abstract

This thesis contains research into the use of 3D digitisation technologies by cultural heritage institutions in public facing applications. It is particularly interested in those technologies that can be adopted by institutions with limited budget and no previous experience in 3D digitisation.

Whilst there has been research in the area of 3D imaging by museums and cultural heritage institutions, the majority is concerned with the use of the technology for academic or professional, curatorial purposes and on technical comparisons of the various technologies used for capture. Similarly, research conducted on the use of 3D models for public facing and public engagement applications has tended to focus on the various capture technologies, while little has been published on processing raw data for public facing applications – a time-consuming and potentially costly procedure. This research will investigate the issues encountered through the entire 3D digitisation workflow, from capture through processing to dissemination, focusing on the specific problems inherent in public facing projects and the heterogeneous and often problematic nature of museum objects.

There has been little research published on the efficacy of 3D models both as providers of informational content and as public engagement tools used to fulfil a cultural heritage institution's public facing remit. This research assesses the utility of interactive 3D models, as well as rendered animations of 3D content used as in-gallery exhibits and disseminated online. It finds that there is a *prima facie* case for believing that 3D models may be used to further a museum's engagement and educational aims, and that there is an appetite among the general public for the use of this type of content in cultural heritage applications. The research will also compare a variety of methods for assessing the success of models.

Acknowledgements

I would like to thank my two supervisors, Professor Melissa Terras and Professor Stuart Robson from for their advice, support and unbounded patience throughout the course of this research.

I would also like to thank my many colleagues and collaborators in the 3DImpact Group in UCL's department of Civil Environmental and Geomatic Engineering for their help over the last five years, with special mentions for Ali Hosseiniinaveh, Lindsay MacDonald and Mona Hess. I would particularly like to thank Anita Soni for her moral support and least squares advice.

I must also thank Matthew Shaw and William Trossell from ScanLAB Projects, and the staff at the Science Museum and the Courtauld Institute of Art, particularly Dr Alexandra Gerstein, without whom the projects in this thesis would not have been possible.

Table of Contents

1.	Introduction	18
1.1	This thesis.....	18
1.2	Cultural heritage institutions' dual remit.....	19
1.3	The role of 3D digitisation.....	22
1.3.1	In professional applications.....	22
1.3.2	In public facing applications	25
1.4	This research.....	29
1.4.1	Structure of this thesis.....	32
2	Review of 3D digitising technologies	35
2.1	Introduction.....	35
2.2	3D Digitisation Technologies	38
2.2.1	Photogrammetry	38
2.2.2	Laser scanning – time of flight scanners.....	46
2.2.3	Close range triangulation scanners	49
2.2.4	Projected/structured light systems	54
2.3	Strengths and limitations of digitisation methods in public facing cultural heritage applications.....	60
2.3.1	Common weaknesses	60

2.3.2 Strengths and weaknesses	67
2.4 Overall conclusion.....	75
2.5 Software	76
2.5.1 Processing.....	76
2.5.2 Dissemination	78
2.5.3 Virtual reality	79
2.5.4 Point clouds or meshes?	80
2.6 Cultural heritage objects online – the state of the art	87
2.7 The future: measuring BRDFs.....	94
2.8 Conclusions.....	100
3 Digital Surrogacy, Aura and the role of 3D models	102
3.1 Introduction.....	102
3.2 The Digital Surrogate.....	102
3.2.1 The public facing digital surrogate.....	108
3.2.2 The museum's dual remit.....	110
3.2.3 Objectivity	111
3.2.4 Authenticity	115
3.2.5 Interaction	121
3.3 Museology and its history	125

3.3.1	Aura, mechanical reproduction and the digital model.....	131
3.3.2	Three objects with 'aura'	134
3.4	The cultural heritage object in the age of digital reproduction	138
3.5	Conclusions	141
4	Scanning the Science Museum's Shipping Gallery	143
4.1	Introduction and Research Aims	143
4.2	The Shipping Gallery	143
4.3	The project	147
4.3.1	The commission.....	147
4.4	Capture.....	149
4.4.1	Targets.....	149
4.4.2	Control survey	151
4.5	Scanning	152
4.5.1	Scanning: Methodological Issues	154
4.6	Processing	155
4.6.1	Colourisation & Registration	156
4.6.2	Cleaning.....	156
4.7	Output.....	158
4.8	Further work.....	160

4.8.1	Euclidean	161
4.8.2	Amazon.....	162
4.8.3	WebGL.....	162
4.9	Survey & Online Analysis.....	165
4.9.1	Methodology	165
4.9.2	Demographics.....	170
4.9.3	Reception.....	174
4.9.4	Potential uses	176
4.9.5	Further development.....	177
4.9.6	Style.....	179
4.10	Methodological Issues	181
4.11	Conclusions	182
5	Illuminating objects	188
5.1	Introduction and research aims.....	188
5.2	Illuminating Objects	188
5.3	Object 1: Spanish lustreware bowl	192
5.3.1	Capture.....	195
5.3.2	Processing.....	200
5.3.3	Rendering	207

5.4	Object 2: Miniature Bibles	212
5.5	Capture.....	213
5.5.1	Methodology.....	215
5.5.2	Issues.....	216
5.6	Processing	219
5.6.1	Filtering	219
5.6.2	<i>Registration</i>	220
5.6.3	Overlapping scans.....	222
5.6.4	Hole filling	223
5.7	Rendering	225
5.8	Analysis	226
5.8.1	The bowl.....	227
5.8.2	The bible.....	231
5.9	User Survey	233
5.9.1	Methodology.....	233
5.9.2	Analysis	234
5.9.3	Demographics.....	235
5.9.4	Technical/experiential.....	237
5.9.5	Understanding.....	238

5.9.6	Interest in 3D models in general.....	239
5.10	Correlations.....	240
5.11	Comment analysis	242
5.12	Methodological issues	244
5.13	Solutions?.....	246
5.14	Alternatives to online surveys.....	248
5.15	Conclusions	250
6	The Courtauld Bag	253
6.1	Introduction and research aims.....	253
6.2	The Courtauld Bag.....	253
6.3	The commission	256
6.4	Imaging	257
6.4.1	Reflectance Transformation Imaging	257
6.4.2	Laser scanning	260
6.4.3	Photography	263
6.5	Processing.....	269
6.5.1	Structure from motion model	269
6.6	Point cloud processing:	279
6.6.1	Exterior.....	279

6.6.2	Point cloud processing: Interior	286
6.7	Output.....	290
6.7.1	The 3D model	290
6.7.2	RTI Imagery.....	294
6.7.3	Outputting video	300
6.7.4	The final output.....	302
6.8	Analysis	304
6.8.1	Aims	304
6.8.2	Methodology.....	305
6.8.3	Results.....	309
6.8.4	Caveats	318
6.9	Conclusions	322
7	Conclusions and further work	325
7.1	Motivations	325
7.2	Thesis structure and chapter conclusions.....	325
7.3	Overall conclusions and reflections	328
7.3.1	Reflections.....	329
7.3.2	Findings of use to cultural heritage institutions.....	330
7.4	Further work.....	335

7.5	Summary of Conclusions	338
8	References and Bibliography	341
9	Appendix A: Science Museum Shipping Gallery Survey	355
9.1	Online survey results.....	355
9.2	Online comments for the Shipping Gallery Video	362
9.2.1	Metafilter.....	362
9.2.2	B3ta.com	366
9.2.3	Alphr.....	368
9.2.4	Gizmag.....	368
9.2.5	YouTube	369
9.2.6	Oculus.com.....	372
9.2.7	Gizmodo	373
9.2.8	Selection of Tweets referencing the Shipping Gallery scan	375
9.3	Media links	381
10	Appendix B: Illuminating Objects survey results	382
11	Appendix C: A possible method for processing highly specular objects.....	393
11.1.1	Our method	394
11.1.2	Matlab code.....	399
12	Appendix D: Suggested method for enhancing user interaction with point cloud models	402

12.1	Motivation	402
12.2	Method	402
13	Appendix E: A selection of scanning projects undertaken during the course of this research	405
14	Appendix F: Publications and presentations	410
14.1	Museum Computer Network, Dallas, November	410
14.2	Digital Humanities Congress, Sheffield, September 2014.....	411
14.3	ISPRS Technical Commission V Symposium, Italy, June 2014	412

Table of Figures

Figure 1.1: Example of deformation monitoring	24
Figure 2.1: The evolution of photogrammetry	42
Figure 2.2: Arius Foundation scanner and Nikon Metris K-Scan	51
Figure 2.3: Michelangelo's David being scanned.....	53
Figure 2.4: Example of a fringe projection	55
Figure 2.5: A 3D model of Jeremy Bentham's auto icon	57
Figure 2.6: Example of an occlusion due to small scale geometry.....	62
Figure 2.7: Laser scan of an elephant's tooth.....	64
Figure 2.8: objects with complex surface properties.....	65
Figure 2.9: Images of meshed models from the Smithsonian 3D web viewer	84
Figure 2.10: objects from the museum of Sheffield's metalwork collection.	88
Figure 2.11: WebGL viewers	89
Figure 2.12: Example of the Smithsonian's material editor.	90
Figure 2.13: Material samples used to measure BRDFs for the MERL database.....	95
Figure 2.14: Two methods for recording BRDFs	97
Figure 2.15: Screenshot of Bonn's object viewer.....	99
Figure 3.1: CG rendering of a glass of milk.....	119
Figure 3.2: Petrie clothing exhibit	136
Figure 3.3: Two examples of moon rock exhibits	137
Figure 4.1: Images of the Shipping Gallery, then and now.....	144
Figure 4.2: Label text in the Shipping Gallery.....	146
Figure 4.3: View of gallery showing paper targets	149

Figure 4.4: Spherical targets	151
Figure 4.5: Faro 120 scanner with DSLR camera mounted	153
Figure 4.6: Images of raw scan data.....	157
Figure 4.7: Turbine model stills from video.....	159
Figure 4.8: Images from the animation.....	160
Figure 4.9: Close up of HMS Monarch rendered via WebGL.	163
Figure 4.10: Figurehead	164
Figure 4.11: Marine turbine	164
Figure 4.12: Science Museum Survey questions	168
Figure 4.13: Age vs familiarity	173
Figure 4.14: Q4 "Did you enjoy the video?"	174
Figure 4.15: age (Q12) vs enjoyment (Q4)	176
Figure 5.1: The lustreware bowl in the Illuminating Objects display case.....	190
Figure 5.2: A typical photo from the imaging session	193
Figure 5.3: Photographs showing iridescence caused by the thin glaze	193
Figure 5.4: The Nikon scanner in action.....	196
Figure 5.5: Scanning errors.....	197
Figure 5.6: One of the images from the photography session.....	182
Figure 5.7: Level 1 and level 0 models	202
Figure 5.8: A cloud of points can be seen detached from the surface of the bowl	203
Figure 5.9: Area of high specularity in showing noisy data	204
Figure 5.10: Hole filling algorithms.....	205
Figure 5.11: Detail of the bowl's rim.....	206
Figure 5.12: Photographs of one of the bibles	212
Figure 5.13: The bible underneath the Arius scanner	214
Figure 5.14: Scan taken at 45° and 0°, and the result when the two are registered... ..	216

Figure 5.15: An early model of the bible.....	217
Figure 5.16: Detail of the bible's spine.....	217
Figure 5.17: Comparison of the two covers	223
Figure 5.18: Less than perfect hole filling	225
Figure 5.19: Splitting the bible by manually painting points.....	209
Figure 5.20: Survey questions.....	235
Figure 5.21: Interest vs understanding	222
Figure 5.22: Bible landing page.....	248
Figure 6.1: Image of the bag (from the first photogrammetry imaging session)	254
Figure 6.2: the bag before and after cleaning	255
Figure 6.3: The bag under the RTI dome.....	258
Figure 6.4; The four areas selected for imaging under the dome.	260
Figure 6.5: Holes in the laser scan caused by excessive reflection.....	261
Figure 6.6: Clouds of false points corresponding to holes in the scan	261
Figure 6.7: Three separate laser scans showing colour discrepancies	262
Figure 6.8: Schematic diagram showing angle of laser and surface	263
Figure 6.9: Camera setup for capturing photogrammetry images	265
Figure 6.10: Camera setup, wide view, including the two lights with soft boxes.....	265
Figure 6.11: First attempt at modelling the interior of the bag	267
Figure 6.12: The UCL engineering department's camera calibration object	271
Figure 6.13: Screenshot from VSFM showing the full 'virtual' network	277
Figure 6.14 Top and base models before registration	278
Figure 6.15: Missing data and model with missing data replaced	280
Figure 6.16: Close-up image of the braiding on the bag's edge	281
Figure 6.17: Hole filling on one of the bag's roundels	282
Figure 6.18: Interior of the bag, second attempt	287

Figure 6.19: Image of the texturing process.	288
Figure 6.20: Model with interior.	289
Figure 6.21: Nine input images for MacDonald's RTI rendering algorithm.	296
Figure 6.22: Comparison of RTI rendering techniques	300
Figure 6.23: rough plan of the exhibition (not to scale)	312
Figure 6.24: View of the bag on display in Court and Craft	317
Figure 11.1: The user selects an area containing a specular highlight	395
Figure 11.2: Histograms for the area selected in fig 11.1	395
Figure 11.3: The whole process repeated with a 'good' area	396
Figure 11.4: Input photo with corresponding output.....	397
Figure 12.1: Model rendered normally and with coloured point cloud	404

1. Introduction

1.1 This thesis

This research is concerned with the use of 3D digitisation and subsequent dissemination of 3D models in public facing applications by cultural heritage (CH) institutions¹. The different technologies available for 3D digitisation will be assessed in terms of their suitability for use in cultural heritage, with particular emphasis placed on recent solutions that can be adopted at little cost and by institutions with no previous experience in the field.

It should be noted how 3D digitisation differs from 3D *reconstruction*. Digitisation is a recording or measuring process which aims to create an objective representation of the thing being digitised². This is in contrast to reconstructions made for visualisation purposes which could be considered ‘artist’s impressions’.

During the course of this research, several digitisation projects have been undertaken with cultural heritage institutions. The Science Museum Shipping Gallery project was the first to capture and preserve an entire gallery using terrestrial laser scanning, and the resulting video was used to test the public appetite for this type of resource. Two projects conducted for the Courtauld Institute of Art involved digitising individual

¹ Whilst the act of digitising an object in 3D, and subsequently the creation of a visualisation of the object can be considered two separate processes, it will be seen in this research that each process inform the other. The requirements of the visualisation (for example, the necessary level of detail, whether the resulting file needs to be served on or offline, etc.) will place restrictions on the capture method whilst the facts of digitisation (what detail can be captured, the amount of processing required) will help to determine what sort of visualisation can be created.

² As will be discussed in section 3, how far genuine objectivity can, in fact, be achieved, is a debatable point.

objects. The first used close range laser scanning and photogrammetry to create interactive 3D models which were placed online using a WebGL renderer. The second used photogrammetry and RTI imaging to produce models which were then used in a pre-rendered animation which was shown in the gallery alongside the original object. User research was conducted on both projects to assess the utility of these techniques in public facing applications, and how they might be used to further the aims of cultural heritage institutions (see section 1.2 for a discussion of these aims).

The combination of these two research questions will hopefully allow cultural heritage institutions to make more informed decisions in their use of digitised 3D content

1.2 Cultural heritage institutions' dual remit

Museums³ and other cultural heritage institutions, like the objects they contain, form an extremely heterogeneous class. However, despite their widely varying nature, however, they can all be considered to operate under a dual remit which could be simply characterised as *preservation* and *display*. In 2007, the International Council of Museums (ICOM) defined 'museum' as:

"a non-profit, permanent institution in the service of society and its development, open to the public, which acquires, conserves, researches,

³ Throughout this thesis I will be using 'museum' and 'cultural heritage institution' interchangeably, so in future instances 'museum' should be read to include art galleries as well as those organisations responsible for public monuments, art installations etc. This is not to imply that, for example, museums and art galleries are identical, or that these institutions have identical roles and functions as the owner of a publicly accessible historic palace. However, as custodians of cultural heritage, with similar responsibilities for preservation and the provision of public access, in regards to using 3D digitisation in public facing applications their aims are close enough that for the purposes of this research they can be considered synonymous.

communicates and exhibits the tangible and intangible heritage of humanity and its environment for the purposes of education, study and enjoyment.”⁴

These two strands of a museum’s function could further be characterised as their *professional* and *public-facing* remits. The former would include all aspects of a museum related to the acquisition, conservation and preservation of cultural heritage and the academic research conducted on it; the latter would include aspects concerned with displaying and providing access to cultural heritage and the dissemination of knowledge.

Part of the professional remit includes documenting and archiving cultural heritage and providing access for research purposes. This function is undeniably important and central to their role as cultural heritage institutions: the preservation of cultural heritage to protect against the inevitable effects of entropy, the detailed description and recording of objects for when preservation is not sufficient, and the increase in our knowledge and understanding of our cultural heritage and its context. All these roles can be seen to benefit not just ourselves but future generations.

The second part of the remit, ‘for the purposes of education, study and enjoyment’ would appear to be secondary to the first. Simply put, preservation must be a primary concern, otherwise there will be nothing to display in the future. Without research there will be no knowledge to disseminate. Whilst this is broadly true, the primacy of

⁴ ICOM Statutes, 21st General Conference in Vienna, Austria, in 2007, <http://icom.museum/the-vision/museum-definition/>. This is not a completely uncontroversial definition, and will be looked at in more detail in 3.

the museum's professional role does not diminish the importance of the second.

Unesco, (United Nations Educational, Scientific and Cultural Organisation) in their 1978

*Recommendation for the protection of movable cultural property*⁵ state that "cultural

property representing the different cultures *forms part of the common heritage of*

mankind"⁶, ICOM's code of ethics state "Museums that maintain collections hold them

in trust for the benefit of society and its development"⁷ while the Museums

Association's Code of Ethics⁸ says museums should:

1: Hold collections in trust for the benefit of society

2: Focus on public service

3: Encourage people to explore collections for inspiration, learning and
enjoyment

4: Consult and involve communities, users and supporters

These ideas, that CH institutions hold their contents in trust for 'mankind' lends

considerable importance to their public facing remit. One could argue that, in fact, the

museum's professional remit is there to support the public one; that the preservation

and understanding of cultural heritage is meaningless if what is preserved and what is

understood is inaccessible to the people for whom it is held in trust.⁹

⁵ Whilst this document was specifically concerned with 'movable' property, the statements are equally applicable to other items of cultural heritage

⁶ UNESCO General Conference, 1978. Recommendation for the protection of movable cultural property. http://portal.unesco.org/en/ev.php-URL_ID=13137&URL_DO=DO_TOPIC&URL_SECTION=201.html

⁷ ICOM Code of ethics, section 2: <http://icom.museum/the-vision/code-of-ethics/2-museums-that-maintain-collections-hold-them-in-trust-for-the-benefit-of-society-and-its-developme/#sommairecontent>

⁸ <http://www.museumsassociation.org/ethics/code-of-ethics>

⁹ And, it should be remembered, the people who paid, through taxes, entrance fees etc., for the preservation and research.

1.3 The role of 3D digitisation¹⁰

1.3.1 In professional applications

Documentation of objects in a museum's collection is an essential part of their professional remit¹¹, and the digitisation^{12,13} of such records and subsequent provision of access (for instance, through a searchable database accessed via the web) has been shown to not only aid professional research but also benefit the general public^{14,15}.

Compared to 2D, 3D digitisation is a relatively new field, but the potential benefits to the cultural heritage professional have been widely discussed and researched^{16,17} for example by Taylor et al:

"3D imaging of works of art offers a significant new analytical tool to curators, historians and conservators, which provides some new and unique types of information which otherwise is not obtainable using traditional techniques. The high-resolution 3D image data contain a wealth of information that can be used

¹⁰ 'Digitisation' can be defined as both a recording and conversion process. An analogue object is converted into a digital form by a sampling process. In the case of 3D digitisation, the object's surface is sampled at discrete points to build up a 3D model.

¹¹ ICOM code of ethics, Section 2.1, Collections Policy: <http://icom.museum/the-vision/code-of-ethics/2-museums-that-maintain-collections-hold-them-in-trust-for-the-benefit-of-society-and-its-developme/#sommairecontent>

¹² Terras, M. (2015). "Cultural Heritage Information: Artefacts and Digitization Technologies" In Chowdhury, G. and Ruthven, I. (2015). "Cultural Heritage information", Facet. p. 63-88.

¹³ Terras, M. (2010). "The Rise of Digitization: An Overview". In Rukowski, R. (ed). (2010). "Digital Libraries". Sense Publishers, The Netherlands. p. 3-20.

¹⁴ Hughes, L. M. (2012). Introduction: the Value, Use and Impact of Digital Collections. Evaluating and Measuring the Value, Use and Impact of Digital Collections. London: Facet Publishing.

¹⁵ Terras, M., Ross, C. & Motyckova, V., 2012. Measuring impact and use: scholarly information-seeking behaviour. In Hughes (2012), pp85-102

¹⁶ La Pensée, A., Cooper, M. I., & Parsons, J. B. (2006). Applications in the field of cultural heritage using "off-the-shelf" 3d laser scanning technology in novel ways. In Proc. 7th Int. Symp. On Virtual Reality, Archaeology and Cultural Heritage,(VAST) (Vol. 215).

¹⁷ Payne, E. M. (2013). Imaging techniques in conservation. Journal of Conservation and Museum Studies, 10(2), pp17-29.

for modelling, display, comparison, measurement and analysis applications.”¹⁸

An example of an application is remote object assessment¹⁹, where researchers or curators can access and virtually examine potentially fragile and friable objects without the need to handle (or transport) the original. A 3D model is a particularly rich data set²⁰ for monitoring an object over time for conservation purposes (‘4D imaging’), and 3D point cloud software (such as GeoMagic²¹ or the open source CloudCompare²²) make detecting changes in an object over time simple, intuitive and above all accurate (Figure 1.1).

3D documentation has also found considerable utility in the field of archaeology, and non-contact 3d documentation techniques have been in use for several decades²³. Traditional methods of documenting an archaeological site (ie, surveying techniques and measurements made with tapes etc.) can be time consuming and only measure a set of discrete points. 3D technologies, including photogrammetry^{24,25,26}, terrestrial

¹⁸ Taylor, J., Beraldin, J_A, Godin, G., Cournoyer, L., Rioux, M., and Domey, J. (2002) 3D imaging technology for museums and heritage published in *Proceedings of The First International Workshop on 3D Virtual Heritage*, October 2-3, 2002, Geneva, Switzerland, pp. 70-75.

¹⁹ Hess, M., Millar, F. S., Ong, Y. H., MacDonald, S., Robson, S., Brown, I., & Were, G. (2008). E-Curator: 3D Colour Scans for Object Assessment.

²⁰ Certainly compared to traditional methods, such as making individual measurements over time

²¹ <http://www.geomagic.com/>

²² <http://www.danielgm.net/cc/>

²³ Jiang Yu Zheng, Zhong Li Zhang and N. Abe, "Virtual recovery of excavated archaeological finds," *Multimedia Computing and Systems, 1998. Proceedings. IEEE International Conference on*, Austin, TX, 1998, pp. 348-357. doi: 10.1109/MMCS.1998.693663

²⁴ Susie Green, Andrew Bevan, Michael Shapland, A comparative assessment of structure from motion methods for archaeological research, *Journal of Archaeological Science*, Volume 46, June 2014, Pages 173-181, ISSN 0305-4403, <http://dx.doi.org/10.1016/j.jas.2014.02.030>.

²⁵ Henrique Lorenzo and Pedro Arias, A methodology for rapid archaeological site documentation using ground-penetrating radar and terrestrial photogrammetry, *Geoarchaeology, Special Issue: Archaeological Geophysics: A Global Perspective*, Volume 20, Issue 5, pages 521–535, May 2005

²⁶ Andrew Bevan, Xiuzhen Li, Marcos Martinón-Torres, Susan Green, Yin Xia, Kun Zhao, Zhen Zhao, Shengtao Ma, Wei Cao, Thilo Rehren, Computer vision, archaeological classification and China's terracotta warriors,

laser scanning²⁷ and airborne (LIDAR)²⁸ scanning allow an entire site to be mapped relatively quickly, in detail and with a high degree of accuracy.




Figure 1.1: Example of deformation monitoring over time using GeoMagic Qualify software. Image: Anita Soni, UCL

3D digitisation can also help preserve, at least virtually, cultural heritage that is lost through decay, accident or deliberate destruction. Cultural vandalism, whether by an individual or state sponsored is in no way a new phenomenon²⁹, and recent events in

Journal of Archaeological Science, Volume 49, September 2014, Pages 249-254, ISSN 0305-4403, <http://dx.doi.org/10.1016/j.jas.2014.05.014>.

27 José Luis Lerma, Santiago Navarro, Miriam Cabrelles, Valentín Villaverde, Terrestrial laser scanning and close range photogrammetry for 3D archaeological documentation: the Upper Palaeolithic Cave of Parpalló as a case study, Journal of Archaeological Science, Volume 37, Issue 3, March 2010, Pages 499-507, ISSN 0305-4403, <http://dx.doi.org/10.1016/j.jas.2009.10.011>.

28 Michael Doneus, Christian Briese, Martin Fera, Martin Janner, Archaeological prospection of forested areas using full-waveform airborne laser scanning, Journal of Archaeological Science, Volume 35, Issue 4, April 2008, Pages 882-893, ISSN 0305-4403, <http://dx.doi.org/10.1016/j.jas.2007.06.013>

²⁹ Cf. Henry VIII's dissolution of the monasteries

the middle-east clearly demonstrate the fragility of much of the world's cultural heritage. Whilst it would be naive to suggest that a digital model is in any way a replacement (or digital surrogate – a concept which will be discussed in some detail in section 3.2) for a lost or destroyed object, a digitised model may ameliorate the loss³⁰.

1.3.2 In public facing applications

Compared to the professional sphere, until recently there was a dearth of public facing applications involving 3D digitisation. There are several reasons for this. Historically, 3D digitisation required expensive, specialist equipment (for example laser scanners or specialist photogrammetry equipment³¹), and once the object was digitised, there was no easy way to disseminate good quality models to the general public. An early example of a public facing CH application was the 90's Virtual Stonehenge project³² which provided a 3D model³³ derived from photogrammetry over the internet. The model required the user to download both the data for the model and a software plug-in to render it, and the quality of the model was hugely compromised by limits on the amount of data that could be transferred over the dial-up connections of the time.

In recent years, the barriers to both digitisation and dissemination have fallen

³⁰ For example, see the work of Project Mosul, "a response to the destruction of cultural heritage by the Islamic States" which uses "crowd-sourced imagery to digitally reconstruct the heritage that has been destroyed". <http://projectmosul.org/>

³¹ I will examine the costs of various technologies, including recent low cost alternatives in Chapter 2

³² Burton, N. R., Hitchen, M. E., & Bryan, P. G. (1999). Virtual Stonehenge: a Fall from Disgrace?, BAR INTERNATIONAL SERIES, 750, 265-265.

³³ The virtual Stonehenge project refers to their model as a 'virtual reality' model; however, the output is a 3D model explorable on a 2D display (a computer monitor) via the internet. To clarify, and reflecting contemporary usage, this research will restrict the term 'virtual reality' to immersive experiences delivered by 3D VR technologies such as the Oculus Rift, HTC Vive, and to a lesser extent Google Glass. The models used may be identical, the difference is in the method of delivery.

dramatically. The advent of low-cost sensors and practically no-cost photogrammetry has made 3D digitisation an affordable option, whilst the emergence and rapid adoption of broadband web access and new technologies such as WebGL³⁴ have presented a clear a path to the audience. Meanwhile, the growth of 3D printing, whilst not directly related to this research, has led to an increase in demand for free, user friendly 3D capture and processing software, increasing the size of the 3d ecosystem. One can imagine this process becoming a virtuous circle, where increase in demand leads to better software, leading in turn to more demand, etc.

However, whilst the barriers to entry have fallen a long way, when compared to 2D digitisation, a 3D project is still a very time consuming and non-trivial thing to embark on, and thus there need to be compelling reasons to undertake one. I would argue that in terms of fulfilling a museum's remit, the potential benefits of using 3D models in public facing applications are possibly greater than those in the professional sphere. As we saw above, providing public access to its collections is a key component of a museum's remit, and also an ethical obligation. An interactive 3D model placed on a website is immediately accessible by an audience many orders of magnitude larger than could ever visit even the world's most popular museum. Similarly, a museum will usually only have a small fraction of its collection on display at any one time, purely due to the physical limitations of space. Percentages vary wildly from institution to institution³⁵; the Victoria and Albert Museum claims to have 24% of their 'display

³⁴ <https://www.khronos.org/webgl/> - WebGL will be covered in more detail in chapter 2

³⁵ <http://www.bbc.com/culture/story/20150123-7-masterpieces-you-cant-see>

collections' on view at any one time³⁶ though the figure can be as low as 2% (the Smithsonian³⁷) or even 0.5% (the British Museum³⁸). Thus there is an immediate conflict between the museum's mission to collect and its mission to allow access. Again, a digitised object or even an entire digitised exhibition takes up no physical space³⁹ and thus could help resolve some of this conflict.

There are other conflicts; as is the case in the professional sphere, the mere fact of fulfilling one mission (access) can entail risk to another (preservation). As UNESCO explain,

“the growing desire of the public to know and appreciate the wealth of the cultural heritage, of whatever origin, has nevertheless led to an increase in all the dangers to which cultural property is exposed as a result of particularly easy access or inadequate protection”⁴⁰

Simply put, the more you allow (or encourage) people to access an object or the more the object has to travel to be accessed, the greater the risk of loss or damage to the object. Take, for example, the constant tension between allowing visitors better access to monuments such as Stonehenge, and the potential of irreversible damage to the site.

There are also the more ineffable potential benefits that 3D models may provide. Do

³⁶ <http://www.vam.ac.uk/content/articles/s/size-of-the-v-and-a-collections/>

³⁷ <http://newsdesk.si.edu/factsheets/fact-sheet-smithsonian-collections>

³⁸ Gardner, L. (2009). The Uses of Stored Collections in some London Museums. *Papers from the Institute of Archaeology*, 18(S1).

³⁹ There are, of course, other types of space issues involved in storing digital objects.

⁴⁰ Unesco (1978) *ibid*

we learn about objects better when presented with a 3D model (as opposed to say, a set of 2D images)? Is the information we glean from a digitised 3D model of an object of better quality? Retained for longer? All these questions pertain directly to the museum's educational remit. We could also ask if 3D objects entertain and engage us better than 2D content. In fact, Metallo and Rossi argue that the ability to digitise museum objects with "incredible accuracy and realism" marks the point where "research begins to blur with entertainment"⁴¹. There are practical and financial considerations for museums too. Are we more likely to visit a website containing 3D models, stay longer, view more? And crucially, are we more – or indeed less – likely to visit a physical museum after seeing 3D models of its exhibits?

Whilst there has been much research into the 3D digitisation of cultural heritage, this tends to be either reviews of technology, novel methodologies or individual case studies of digitisation projects. There has been very little research into the outputs of these technologies and projects. Until this research is carried out, institutions adopting 3D technologies in public facing applications risk what Pallud⁴², building on the work of Monod and Klein, calls technological determinism, where the "implementation of [digital technologies] in museums is assumed will positively impact visitor satisfaction even while there is little verification of whether these technologies really achieve their goal."

⁴¹ Metallo A & Rossi V (2011) *The Future of Three-Dimensional Imaging and Museum Applications*, Curator, The Museum Journal, Vol 54. No 1, Jan 2011, pp 63-69

⁴² Pallud, J. (2009). The application of a phenomenological framework to assess user experience with museum technologies, ECIS 2009 Proceedings. Paper 395.

Whilst there may be good, intuitive reasons to believe that the provision of 3D models by museums may help them fulfill their public facing remit (for example, as Metallo and Rossi say: "our experience with museum websites so far has shown that putting high resolution images of collections online just increases audience engagement and familiarity with collections ... We expect 3D to do the same."⁴³), there is also the possibility that the use of 3D models (or the incorrect use) may have a negative impact. We will examine Benjamin's concept of 'aura'⁴⁴ and how it pertains to both cultural heritage and 3D digitised objects in section 3.3.1, but we must accept the possibility that the provision of 3D models may have a negative effect on the public's appreciation of cultural heritage, may discourage visits to physical exhibits and somehow cheapen the role of the museum. It is the intention of this research to at least begin to answer some of these questions.

1.4 This research

There are two main strands to the research in this thesis. One is to ascertain how feasible it would be for a cultural heritage institution with a limited budget⁴⁵ and no previous experience in the area to carry out a 3D digitisation project successfully, from capture through processing to dissemination. In other words, can a small institution use 3D digitisation techniques to create a virtual model of sufficient quality that it can be used to fulfill one or more of the institution's remits.

⁴³ Metallo & Rossi (2011)

⁴⁴ Benjamin, Walter. "1939. The Work of Art in the Age of Mechanical Reproduction." *Media and Cultural Studies. Key-Works* (1936): 18-40.

⁴⁵ Particularly in the current climate where cultural heritage institutions are facing drastic reductions in their public funding: <http://www.museumsassociation.org/campaigns/funding-cuts/fighting-the-cuts>

By small institution, here we refer to a museum's capacity for digitisation rather than, for example, physical size, size of collection or visitor numbers. A 'small' institution in this research would be one that has no dedicated digitisation team (or, more specifically, no dedicated 3D digitisation team), nor the budget to outsource a 3D digitisation programme to a commercial provider. In this research, the Courtauld Institute would be considered a 'small' institution: while they have conducted several 2D digitisation projects, these are often funded through institutional partnerships or via, for example, lottery funding⁴⁶. Whilst a limited amount of money was available from the Courtauld to support the digitisation projects in this research (see chapters 5 & 6), there certainly weren't the funds to have paid to outsource the entire project to a commercial provider.

To contrast, in this research the Science Museum would be considered a 'large' institution, with the budget to engage external providers (in this case Scanlab) to carry out large scale, expensive 3D digitisation projects⁴⁷. All cultural heritage institutions will exist somewhere on the continuum from 'small' to 'large', with some considerably smaller and less well funded than the Courtauld (a local museum, for example), and some larger and better resourced than the Science Museum (the Smithsonian Institute in the US, for example).

Of course, the low cost solutions examined in this research are applicable to

⁴⁶ The Courtauld Institute of Art Institutional Audit, Annex to the report, Feb 2011. Accessed 2/3/16 at <http://www.qaa.ac.uk/en/ReviewsAndReports/Documents/The%20Courtauld%20Institute%20of%20Art/Co urtauld-Institute-of-Art-IA-annex-11.pdf>

⁴⁷ Though, as we will see in chapter 4, the cost as charged was far below the actual commercial cost of the project.

institutions of all sizes, and as we shall see, 'low cost' does not necessarily imply lower quality results. These methods may be the best choice for a particular digitisation project regardless of the budget and equipment available, though they, and therefore the results of this research, may potentially be of more relevance to smaller institutions where alternative options are limited.

The workflows required to create models using a variety of technologies (close and long-range laser scanning, photogrammetry) and to make them available either via pre-rendered video or as fully interactive online models will be examined, and potential problems identified.

The second strand is to conduct user-testing on the outputs produced as part of the research in order to ascertain the utility of the techniques. Public response to the models will be gauged in terms of a general appetite⁴⁸ for the use of 3D in cultural heritage contexts, and the utility produced by the models in terms of both their informational content and capacity to engage the audience. These findings will be used to measure how well the models perform in the furtherance of the museums' remits. Previous research in this area has concentrated on the mechanics of capturing, processing and disseminating 3D content; its utility has often been assumed⁴⁹ or the question ignored. Whilst there are strong circumstantial and intuitive reasons for

⁴⁸ Whilst a public appetite for something – in this case 3D digitised models – does not necessarily entail that the 'something' is worth pursuing, one would assume that for the public to be engaged by the resources, and therefore able or more likely to extract informational content, there must be an appetite *for*, or at least the absence of an aversion *to*, 3D content.

⁴⁹ as CH 3D digitisation becomes a 'solved problem' attention will turn to the users, see for example: Alelis, Genevieve, Ania Bobrowicz, and Chee Siang Ang. "Comparison of engagement and emotional responses of older and younger adults interacting with 3D cultural heritage artefacts on personal devices." *Behaviour & Information Technology* (2015).

believing 3D content will be well received (from previous experience in the commercial sector, the wide adoption of 3D in virtual worlds and gaming), few experiments have been conducted that provide either quantitative or robust qualitative evidence as to how cultural heritage content has been received. Better knowledge of both the processes involved in producing digitised objects, and the utility (or otherwise) of 3D models, can be used by cultural heritage institutions to allocate their resources in such a way as to maximise the return on their investment. The ability to make more informed decisions will help them avoid the trap of technological determinism and concentrate on those projects that will best fulfill their specific requirements.

1.4.1 Structure of this thesis

Chapter 2 is an introduction to, and brief history of, the various technologies used in 3D digitisation and their suitability for use in public facing cultural heritage applications. It also examines in some detail the issues faced by all digitising technologies when imaging ‘difficult’ objects, and looks at some of the newest methods for capturing complex optical properties. In particular, it looks in at developments in the last five years (the period of this research) and their implications for the future of 3D digitisation. It also looks at technologies relevant to disseminating 3D models via the web and looks at some of the current state-of-the-art, public facing cultural heritage applications.

Chapter 3 discusses some of the fundamental philosophical issues of 3D digitisation, in particular the concept of a ‘digital surrogate’, a common, if perhaps ill-defined term used in regards to digitised cultural heritage. It looks in more detail at some of the

museological issues raised in this chapter, and finally discusses Benjamin's concept of aura and what ramifications this may have on virtual cultural heritage and how it is presented.

Chapters 4, 5 and 6 consist of case studies of 3D digitisation projects conducted as part of this research. All the chapters cover the workflow required to create the model plus required output, and detail the specific issues encountered in each project. Each chapter finishes with an analysis of the research conducted on each output.

Chapter 4 deals with a project conducted between myself, a professional scanning company (ScanLAB) and the Science Museum to digitise an entire gallery. The output in this case was a pre-rendered video which was placed online, and research on the video was conducted using a mixture of online surveys and comment analysis.

The digitisation projects in chapters 5 & 6 both concern projects undertaken with the Courtauld Institute. In **chapter 5**, two objects from the gallery's Illuminating Objects were digitised, and interactive models placed online on the Courtauld's website. User testing was conducted via online surveys. **Chapter 6** details a major project which involved digitising a very problematic object and creating a pre-rendered video which would appear in-gallery in a major exhibition at the institute. Research was conducted using interviews with visitors to the gallery.

Chapter 7 contains conclusions, reflections and suggestions for further work, and addresses some of the methodological problems with the survey methods used in this research.

The Appendices include details of the survey results, media coverage and online comments for chapter 4, survey results from chapter 5, two pieces of code, for masking specular reflections in 2D images and a novel method for labelling 3D point clouds, a selection of scanning projects undertaken during the period of this research, and details of publications and conference presentations given by the author.

2 Review of 3D digitising technologies

2.1 Introduction

The use of three dimensional (3D) models or replicas for cultural heritage applications is by no means a new phenomenon. The practice stretches back centuries and long before the advent of computer graphics and virtual reality, museums would create and display physical casts of objects, allowing visitors to “experience architectural monuments and sculptures as if they were physically present at the original site”⁵⁰. The motivation has been around for a long time, but only in the last 100 years or so has technology allowed the relatively easy digitisation of heritage objects through non-contact recording. This chapter examines the various technologies available for acquiring and disseminating accurate 3D digitised models, and their relevance and application regarding public facing cultural heritage applications.

It will also highlight the most important changes in technology over the last five years – the period of this research – and the resulting ‘democratisation’ of 3D digitisation.

When this thesis was started, the ability to create, share and disseminate realistic 3D models was restricted largely to experts using specialist – and expensive – equipment and/or software. With the emergence or maturation of several key technologies – free, user-friendly photogrammetric software coupled with ubiquitous digital photography;

⁵⁰ Bearman, D (2011) 3D Representations in Museums, Curator: The Museum Journal, Vol 54, Iss 1, pp55-61, Jan 2011

the arrival of Microsoft's Kinect and other low cost 3D scanners, and the huge growth in 3D printing driving the need for cheap, user-friendly 3D software, the ability to digitise the real world in three dimensions is now in everybody's hands. Literally, given that a 3D model can be created using nothing more than a smartphone. And with the emergence of WebGL and HTML5 in the last five years⁵¹, it is now possible to display 3D models natively in browsers, allowing anyone with a web connection to immediately view and interact with the models online. This in turn has lowered the barriers preventing cultural heritage institutions from digitising their collections for public facing applications. While the process of creating virtual models required specialist equipment and skills and thus potentially considerable expense, making the case for the use of these models in public facing applications – as opposed to professional curatorial applications where the cost/benefit may have been easier to calculate – was difficult. Whilst, as we will see in the next chapter, there is good prima facie evidence to believe that 3D models may help cultural institutions further both their educational and entertainment remits, until recently it would still have been difficult for all but the largest (and wealthiest) institutions⁵² to commit to the expense of digitisation.

However, with the barriers lowered, it is now possible for small institutions to create and disseminate their own 3D models, if not easily, then at least cheaply. It is one of the aims of this thesis to show that digitised models of sufficiently high quality that

⁵¹ The first browsers (Chrome and Firefox) to feature WebGL appeared in 2011

⁵² For example, the Smithsonian in the US, whose work in this area will be examined later in this chapter.

they are engaging and can aid understanding, are within the reach of users with no previous skill or training⁵³ and with no expensive hardware. Coming into this research, I myself had no previous experience in this field, and was unaware of most of the technologies involved. Whilst I have had access to both hardware and software that may not be available to all, and more importantly, many helpful experts in the area, I hope this research still shows that the skills required are learnable and that, from a position of ignorance, a small CH institution could indeed embark on its own 3D digitisation programme. To clarify: certain skills will still need to be learned, but these do not require any previous knowledge or training or a background in a specific area (engineering, computing, photography etc.). The experiments in later chapters show that the creation of 3D models is a craft process and as such results will tend to improve with practice; I am not implying that a user with zero experience, a smart phone and some free software can immediately start creating 3D models of sufficient quality to be used in public facing applications (though equally, there is no specific reason why they can't), merely that the potential is there.

The rest of this chapter looks at the main technologies used for 3D digitisation and their potential applicability for public facing cultural heritage purposes. It will also discuss the problems shared by all of these methods when capturing objects with complex surface properties. Finally it will take a brief look at current research into new

technologies that may go some way to solving these problems. I shall also look at the technologies involved in processing and displaying the models, and offer some examples of current projects that are ‘state of the art’ in public facing 3D visualisation.

2.2 3D Digitisation Technologies

2.2.1 Photogrammetry

Photogrammetry, literally the “science of making measurements from photographs”⁵⁴, and its application to cultural heritage recording, has a history dating back almost two centuries. Sir Charles Wheatstone discovered the principles of stereoscopy in 1838⁵⁵ and the principles of photogrammetry – the process of making accurate 3D measurements from 2D images – were well established by the latter half of the nineteenth century^{56,57}. Architect Albrecht Meydenbauer was one of the pioneers of both photogrammetry and its use in cultural heritage, developing its methods through the 1860s. Meydenbauer used photogrammetry to accurately record and document important buildings and his ultimate aim was to create a cultural heritage archive, establishing the Prussian Photogrammetric Institute in Berlin in 1885. According to Albertz, “he was convinced that the most important cultural heritage objects should be recorded in such a way that they could even be reconstructed in cases of

⁵⁴ www.photogrammetry.com

⁵⁵ Wheatstone, C (1838) Contributions to the Physiology of Vision.--Part the First. On Some Remarkable, and Hitherto Unobserved, Phenomena of Binocular Vision, Philosophical Transactions of the Royal Society of London, Vol. 128, pp. 371-394

⁵⁶ See, for example, Meade Bache, R, (1892) Civil and Military Photogrammetry, Proceedings of the American Philosophical Society, Vol. 30, No. 138 (April), pp. 229-240, or,

⁵⁷ Adams, CB, (1893), Method of Photogrammetry, US patent US 510758 A, Dec 12 1893

destruction”⁵⁸. This was both a prescient attitude with two destructive world wars just over the horizon, and a motivation that still resonates today - see for example Project Mosul⁵⁹, “a response to the destruction of cultural heritage by the Islamic States” and the reconstruction of the Bamiyan Buddhas, destroyed by the Taliban in 2001⁶⁰.

The end of the 20th/beginning of the 21st century saw many uses of photogrammetry in both cultural heritage documentation and visualisation. Notable examples include the campaign conducted in the Tomb of Christ in Jerusalem in 1989, where photogrammetric techniques were combined with traditional surveying methods to produce both accurate architectural plans and elevations as well as textured 3D models of the tomb⁶¹.

In the UK, English Heritage have been active users of photogrammetry (and other 3D technologies): “Since 1983, numerous historic buildings, monuments and landscapes ... have all benefitted from photogrammetric application in some form”⁶². While the digitisation of sites such as Hadrian’s Wall, Stonehenge and Whitby Abbey, have been used predominantly for the creation of orthographic photo reproduction and 3D survey data, textured 3D models have also been produced for public facing applications⁶³.

⁵⁸ Albertz J, 2001. Albrecht Meydenbauer - pioneer of photogrammetric documentation of the cultural heritage. Proceedings of the XVIIth CIPA Symposium (eds. J. Albertz), September 18 - 21, Potsdam, Germany

⁵⁹ <http://projectmosul.org/>, the reconstruction of the Bamiyan Buddhas is covered later in this chapter

⁶⁰ Grün, A., Remondino, F. and Zhang, L. (2004), Photogrammetric Reconstruction of the Great Buddha of Bamiyan, Afghanistan. *The Photogrammetric Record*, 19: 177–199

⁶¹ Cooper, M; A, R; Robson, S; Littleworth, R; M, ; (1992) The Tomb of Christ, Jerusalem; analytical photogrammetry and 3D computer-modelling for archaeology and restoration. In: Archives for Photogrammetry and Remote Sensing. (pp. 774 - 785). International Society for Photogrammetry and Remote Sensing: Washington DC.

⁶² Bryan, P (2005), The role of photogrammetry in understanding enhancing and enjoying England’s historic environment, in Recording, Modelling and Visualization of Cultural Heritage, Baltavista et al (Eds), Taylor and Francis Group, London, pp 78

⁶³ *Ibid*, p80

These virtual reconstructions further English Heritage’s “corporate aims of promoting the historic environment”⁶⁴, for example their 1996 Virtual Stonehenge project, where data from a 1993 photogrammetric survey⁶⁵ was used to create a virtual reality (VR) model of the monument.

In recent years, certain technologies have emerged (or matured sufficiently) that have made photogrammetry a much simpler and more viable process. Early photogrammetry, from about the 1960s onwards was used to measure individual points on an object, and analytical plotters used to create contour models from these points⁶⁶. These line drawings could be converted to Computer Aided Design (CAD) models which could then be rendered as solid objects⁶⁷. In the early and mid-90s, dedicated photogrammetric workstations were introduced, allowing many more points to be processed, generating recognisable (if sparse by today’s standards) point clouds⁶⁸. Today, however, the continuing improvement in personal computing power coupled with contributions from the fields of machine vision and computer graphics, has made processing dense point clouds containing millions of points feasible even on a desktop PC.


⁶⁴ Ibid, p80

⁶⁵ Ibid, p81

⁶⁶ Ebert, J. I. (1984). Remote sensing applications in archaeology. *Advances in Archaeological Method and Theory*, 293-362.

⁶⁷ Cooper et al (1992)

⁶⁸ Bryan, P. G., & Clowes, M. (1997). Surveying Stonehenge by photogrammetry. *The Photogrammetric Record*, 15(89), 739-751.






Left: A 'highly accurate photogrammetric plot' from 1979. The image was produced from metric photographs (one shown), showing the outlines of a wall as well as spot elevations.

Image from: Ebert (1984)

Below: (left) a linestring drawing derived from photogrammetric measurements of the Tomb of Christ. (right) a CAD model of the Edicule (cupola and roof) derived from the linestrings

Images from: Cooper et al (1992)





Left: an image of a DEM (Digital elevation model) of one of the standing stones from the Virtual Stonehenge project, 1994. The largest stones' models contained up to 60,000 points

Image from: Bryan & Clowes (1997)



Left: a point cloud generated for the Courtauld Institute in 2014, containing approx. eight million coloured points (see chapter 6)

Figure 2.1: The evolution of photogrammetry

The emergence of digital photography has changed the photogrammetric workflow to an 'all digital' model – it is now a trivial matter to take hundreds, potentially thousands, of digital images which can be instantly viewed, uploaded and processed. There is no longer a cost/time issue involved in processing plates or film and no practical restriction on the number of images that can be captured at one time. Digital

photography, and specifically the incorporation of (relatively) high quality cameras into mobile phones has led to an exponential increase in the number of photographs being taken⁶⁹, and phones' ubiquity means many people are carrying the necessary tools to conduct photogrammetry in the palm of their hands. It should be noted, of course, that the quality of the model produced will depend to a large extent on the quality of the photographs used. Apart from normal photographic issues like depth of field and consistent lighting, all cameras and lens system will distort⁷⁰ images to a certain extent. However, even mobile phone cameras have been shown to perform as well as consumer grade digital (compact, non SLR) cameras^{71,72} in photogrammetric applications, and photogrammetry software (ie, 123D Catch, VSFM, PhotoScan) typically includes distortion correction as part of its processes. Nevertheless, there are certain techniques that can be applied in order to ensure better accuracy and less noise in the photogrammetric model, for instance metric calibration of the camera and un-distortion of the images, though these may require specialist equipment and software and be beyond most 'amateur' photogrammetrists. See section 6.5.1.1 for

⁶⁹ There are a variety of wildly differing estimates as to how many images we are taking today, and to how many have been taken in total, though all agree that digital photography has led to a vast number of images being captured compared to the pre-digital era

⁷⁰ Akca, D., & Gruen, A. (2009). Comparative geometric and radiometric evaluation of mobile phone and still video cameras. *The Photogrammetric Record*, 24(127), 217-245.

⁷¹ Chikatsu, H., & Takahashi, Y. (2009, August). Comparative evaluation of consumer grade cameras and mobile phone cameras for close range photogrammetry. In *SPIE Optical Engineering+ Applications* (pp. 74470H-74470H). International Society for Optics and Photonics.

⁷² I was recently fortunate enough to see the results of an as yet unpublished research project using photogrammetry to capture a series of at-risk Etruscan tombs. Partly due to the difficulty in getting electricity and bulky equipment to the site, the team used mobile phone cameras and their built in flashes to record the scene. The results certainly look impressive, and the researchers say that the models using the mobile phone images, whilst lacking some of the texture quality of models made using DSLR cameras, captured the geometry of the tombs with the same degree of accuracy.

more detail on the calibration process.

The increase in the number of photos taken, and their availability via the internet has also allowed 'crowd-sourced photogrammetry', notably in the 'Build Rome in a Day' project, which aims to reconstruct a virtual model of the city of Rome using the two million photographs tagged with 'Rome' on Flickr^{73,74}. Tourist images sourced from the internet were also used as one set of inputs in the attempt to virtually reconstruct the Bamiyan Buddhas⁷⁵.

The growth of the web and broadband penetration^{76,77} has, as well as allowing access to billions, if not trillions, of photographs, also placed free online photogrammetry software within reach of anyone. Autodesk's 123D Catch, (based on Project Photofly, originally released May 2011⁷⁸) is a photogrammetry app that can be downloaded to PC or smartphone⁷⁹ and which will process a set of images into a 3D (mesh) model which can be shared, edited or printed (possibly using other apps in the 123D suite). As all processing is done in the cloud, the power of the local machine is not a bottleneck on performance. Although free⁸⁰, and an entirely 'black box' solution⁸¹, 123D Catch still

⁷³ Agarwal, S., Furukawa, Y., Snavely, N., Simon, I., Curless, B., Seitz, S. M., & Szeliski, R. (2011). Building rome in a day. *Communications of the ACM*, 54(10), 105-112.

⁷⁴ <http://grail.cs.washington.edu/rome/>

⁷⁵ Grün, (2004), *ibid*.

⁷⁶ <http://www.statista.com/statistics/272235/fixed-broadband-penetration-in-the-united-kingdom-uk/>

⁷⁷ Seybert, H. (2011). Internet use in households and by individuals in 2011. *Eurostat. Statistics in Focus*, 66.

⁷⁸ http://autodesk.blogs.com/between_the_lines/2011/05/project-photofly-v2-released-for-download.html

⁷⁹ <http://www.123dapp.com/catch> - nb, the original 123DCatch did not require a download, photos could simply be uploaded to a website for processing.

⁸⁰ For non-commercial use; if you wish to create 3D models for commercial use, a subscription is required

⁸¹ Ie, there are no parameters that can be changed and there is no feedback on the processes, error estimations etc.

provides output which may still be sufficient for some public-facing CH applications⁸².

Visual Structure From Motion (VisualSfM or VSFM)⁸³, the software used in this thesis (see chapters 5 and 6) is a free, open source package with a relatively user-friendly graphical user interface (GUI), and while it may need some knowledge to install (particularly on Linux based machines), and its performance is directly related to the power of the machine it is running on⁸⁴, it is still well within the reach of most ‘amateur’ photogrammetrists.

As a result of these new technologies, any museum or small cultural heritage organisation has the potential to create its own high quality 3D models. Some investment may be required— while the necessary hardware and software can be acquired for little or no expense (ie, some form of digital camera is required for photogrammetry, though many institutions will already have access for 2D documentation purposes), it will still need to be acquired, set up, installed and maintained. The digitising staff will require time to train, and the creation of 3D content - both capturing, and particularly processing, is time consuming. Depending on the required output, the learning curve involved in producing models of sufficiently high quality can be both long and steep, however there are many tutorials and guides

⁸² Kersten, T. P., & Lindstaedt, M. (2012). Image-based low-cost systems for automatic 3D recording and modelling of archaeological finds and objects. In *Progress in cultural heritage preservation* (pp. 1-10). Springer, Berlin & Heidelberg.

⁸³ <http://ccwu.me/vsfm/>

⁸⁴ See 6 for more details of VSFM’s performance on various hardware configurations

online⁸⁵, as well as helpful and responsive communities⁸⁶. Chapters 5 and 6 show typical⁸⁷ workflows for photogrammetric projects (albeit for difficult objects) from capture through processing to dissemination.

2.2.2 Laser scanning – time of flight scanners

Whilst photogrammetry can be considered a ‘passive’ recording technology, with a sensor recording electromagnetic radiation (usually, but not exclusively, visible light) from a source such as a camera flash or the sun, long-range terrestrial laser scanning is an active recording technology whose fundamental principle is similar to that of radar⁸⁸, developed in the early part of the 20th century and building on the experiments of Hertz in the 1880s: electromagnetic radiation is projected onto the target and the reflection captured by a sensor. Lasers began to be used in the 1960s⁸⁹, and while both technologies involve measuring the time it takes for radiation to bounce off a target, the laser’s coherent beam and short wavelength (1500-2000nm compared to the 1cm – 1m radio waves used in radar) allows for higher resolution in both angular and range measurements⁹⁰. Long range laser scanners are also referred to as LIDAR⁹¹ devices which includes both airborne and terrestrial scanners.

⁸⁵ For example: <https://dinosaurpalaeo.wordpress.com/2013/10/31/photogrammetry-tutorial-1-equipment/>, <http://www.tested.com/art/makers/460142-art-photogrammetry-how-take-your-photos/>, http://guides.archaeologydataservice.ac.uk/g2gp/Photogram_Toc,

<http://blog.sketchfab.com/post/121838008009/how-to-set-up-a-successful-photogrammetry-project>

⁸⁶ https://groups.google.com/forum/?utm_source=digest&utm_medium=email#!forum/vsfm/topics

⁸⁷ If there is a ‘typical’ workflow – due to the heterogeneity of cultural heritage objects,

⁸⁸ Beraldin et al, (2010) Laser Scanning Technology, in Airborne and Terrestrial Laser Scanning, Vosselman and Maas (eds), Whittles Publ., Dunbeath, p3

⁸⁹ Goyer, G. G. & R. Watson (September 1963). The Laser and its Application to Meteorology in Bulletin of the American Meteorological Society. 44 (9): pp564–575

⁹⁰ Vosselman & Maas (2010), p3

⁹¹ The word is variously described as an acronym of ‘Light detection and ranging’ or simply a portmanteau of ‘light’ and ‘radar’

Terrestrial time of flight laser scanners⁹² are used for scanning objects at a distance of between one and several hundred metres⁹³, with a range accuracy that is usually in the one mm to one cm⁹⁴ range. Therefore they are generally used for recording buildings or large interior spaces (see Science Museum Shipping Gallery scan, chapter 4), though they can also be used for recording large objects (Figure 4.8). Long-range scanners will be either Time of Flight (ToF) or phase based⁹⁵. The former simply measures the time it takes for the laser to be reflected back to the scanner, thus – knowing the speed of light – giving a distance measurement to the object, while the latter uses the phase difference in a reflected sin wave beam (or beams) to calculate the distance. Phase based recorders tend to be quicker, measuring on the order of millions of points per second, whilst ToF scanners are slower (100s of thousands of points per second), have a slightly shorter effective range, but are more accurate⁹⁶. The Faro Focus3D x130 terrestrial laser scanner, a newer version of the scanners used in the Shipping Gallery project (section 4) has a range of .6 – 130m, a ranging error of +/-2mm and captures approximately one million points per second⁹⁷. The operation and output of the two types of long-range scanner are identical, however.

⁹² As opposed to, for example, an instrument mounted on a plane which may have a range of 1000s of metres and an accuracy of 10s of centimetres

⁹³ Aloysius Wehr, LIDAR: Airborne and terrestrial sensors (2008), in Advances in Photogrammetry, Remote Sensing and Spatial Information Sciences: 2008 ISPRS Congress Book, Li, Chen and Baltsavias (Eds), Taylor & Francis Group, London, pp80-81

⁹⁴ Boehler, W., & Marbs, A. (2002, September). 3D scanning instruments. In Proceedings of the CIPA WG 6 International Workshop on Scanning for Cultural Heritage Recording, Ziti, Thessaloniki (pp. 9-18).

⁹⁵ Phase based scanners still use ‘time of flight’ to calculate distance, the difference is in the method they use to calculate the time.

⁹⁶ Beraldin et al, (2010) Laser Scanning Technology, in *Airborne and Terrestrial Laser Scanning*, Vosselman and Maas (eds), Whittles Publ., Dunbeath, pp1-39

⁹⁷ <http://www.faro.com/products/3d-surveying/laser-scanner-faro-focus-3d/features#main>

The terrestrial laser scanners used in cultural heritage recording are typically hemispherical scanners⁹⁸ - a rotating mirror sweeps a laser beam in the vertical direction while the instrument revolves, giving the scanner a 360° field of view in the horizontal plane and a 300° field of view in the vertical plane, leaving a small 'blind spot' underneath the scanner itself.

Whilst these scanners record the 3D position and intensity value (of the reflected light) for each point, colour can be added to the point cloud using a camera built in to the scanner or alternatively a standard DSLR which can be mounted to the scanner. After the scan is complete, the camera takes a series of images which are stitched together into a panorama and projected on to the point cloud.

Long range scanners are commonly used for surveying, forensic and metrology purposes, but also have a long history in cultural heritage applications. They have been particularly useful in archaeological applications⁹⁹ where large and potentially transient areas (ie, a dig site) need to be digitised quickly and accurately, and in recording monumental architecture and natural environments¹⁰⁰. A prime example of their use in recording areas of cultural importance is by CyArk¹⁰¹, a non-profit organisation dedicated to digitising and virtually preserving some of the world's most at risk heritage sites:

⁹⁸ As opposed to 'window' scanners, which will scan a particular area. Hemispherical scanners such as the Faro models can also be used as window scanners, the user able to specify what area to scan.

⁹⁹ Vozikis, G., Haring, A., Vozikis, E., & Kraus, K. (2004). Laser scanning: a new method for recording and documentation in archaeology.

¹⁰⁰ Remondino, F., & Rizzi, A. (2010). Reality-based 3D documentation of natural and cultural heritage sites—techniques, problems, and examples. *Applied Geomatics*, 2(3), 85-100.

¹⁰¹ <http://www.cyark.org/about/>

“CyArk was founded in 2003 to ensure heritage sites are available to future generations, while making them uniquely accessible today. CyArk operates internationally as a 501(c)3 non-profit organization with the mission of using new technologies to create a free, 3D online library of the world's cultural heritage sites before they are lost to natural disasters, destroyed by human aggression or ravaged by the passage of time.”

Terrestrial laser scanners are also becoming increasingly common in public facing applications. A National Geographic TV series, Time Scanners¹⁰² (also shown in the UK on Channel 5 as the ‘Secrets of...’) was based round terrestrial laser scans of famous monuments (the pyramids, Rome’s Coliseum etc.), while Channel 4’s Time Team also featured laser scanning¹⁰³ in its Jersey programme. An impression of the breadth of use of terrestrial scanners, in cultural heritage, entertainment and art can be seen in the work of ScanLAB¹⁰⁴, who were commissioned to scan the Science Museum’s Shipping Gallery for subsequent evaluation by this EngD thesis (Chapter 4).

2.2.3 Close range triangulation scanners

Close range scanners, used for imaging objects at ranges usually less than 5m and often considerably shorter (the Arius Foundation 150 scanner used in chapters 5 and 6 has a recording volume of approx. 1m³), operate on the triangulation principle. A laser spot is projected onto the surface to be measured, and the reflection recorded on a sensor at

¹⁰² <http://natgeotv.com/uk/time-scanners/about>, <http://www.channel5.com/shows/secrets-of-romes-colosseum/episodes/secrets-of-romes-colosseum>

¹⁰³ <http://www.digitalsurveys.co.uk/case-study/mont-orgueil-castle-time-team>

¹⁰⁴ <http://scanlabprojects.co.uk/portfolio>

some known distance from the emitter. Knowing the relationship (baseline) between the emitter and sensor, the position of the reflection on the sensor, and the angle of both projection and collection, simple trigonometry can be used to calculate the reflected spot's position in space relative to the scan head¹⁰⁵. If the scan head's position is also known, either through a physical connection to a coordinate measuring machine (CMM) as in the Arius scanner, or an optical CMM as in the Nikon K Series, the objects surface can be measured in a coordinate space independent from the laser scanner itself. In practice this means that multiple scans can be made of an object (ie, from a variety of positions or angles), with each individual scan pre-registered in an absolute coordinate system. This is particularly important for scanners such as the Arius which builds up complete scans in a series of stripes. Without the CMM, each stripe would need to be registered individually.

Scanning volume, and thus the maximum sized object that can be captured by the scanner, depends to a large extent on the individual machine. 'Fixed' scanners such as the Arius which sits within a mechanical CMM has a scanning volume of approximately 1m³. At the other end of the cost scale, the NextEngine¹⁰⁶, uses a turntable to scan an object in the round, and the scanning volume is limited to what can fit on the turntable – approximately .5m³. Handheld scanners such as the Nikon have larger recording volumes, approximately 5x3m³, though multiple scans can be registered together so larger objects can be captured.

¹⁰⁵ Beraldin (2010) p8

¹⁰⁶ <http://www.nextengine.com/>




Figure 2.2:

Left: The Arius Foundation scanner, showing the scan head mounted within the highly accurate CMM. Above: the Nikon Metris K-Scan showing the handheld laser scan head and the camera bar CMM. The three cameras on the camera bar (one per axis) recognises patterns of LEDs on the scanhead, fixing its position in space to a high degree of accuracy. (image from <http://www.nikonmetrology.com/>)

Typically, a laser scanner will only record the geometry of the object being scanned, and if a colour model is the desired output some other method of colourising the point cloud is required, for example, draping textures taken from colour photographs. However, if instead of a single monochromatic laser, the scanner emits a number of coloured lasers (ie, red green and blue, as used by the Arius scanner) the relative amounts of each colour reflected allow the scanner to record a value for the colour of the surface. Measuring colour and position at the same time allows an accurately coloured model of the object to be built up, point by point. The ability to record colour and geometry is clearly useful in cultural heritage applications, though the vast majority of close range laser scanners only capture geometry.

High end close range laser scanners can achieve geometrical accuracies of the order of 25 microns and record objects with extremely dense point spacings of 50 or 100

microns¹⁰⁷, making them useful for recording archival quality models for use in professional cultural heritage purposes¹⁰⁸. For example, the Digital Michelangelo project¹⁰⁹ scanned the 5m statue of David (amongst others) with enough accuracy to record and analyse individual chisel marks (see Figure 2.3), while NRC in Canada conducted a detailed examination of the Mona Lisa and Da Vinci's painting techniques using a detailed scan of the painting¹¹⁰. Whilst some degree of accuracy will always be required, it may be of secondary importance in some public facing applications.

¹⁰⁷ Hess, M., & Robson, S. (2010). 3D colour imaging for cultural heritage artefacts. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 38(Part 5), 288-292.

¹⁰⁸ Hess M et al (2008), Final Report: E-Curator: 3D Colour Scans For Remote Object Identification and Assessment, UCL Museums and Collections. Available at: <http://www.ahessc.ac.uk/e-curator>

¹⁰⁹ Levoy, M., Pulli, K., Curless, B., Rusinkiewicz, S., Koller, D., Pereira, L. & Fulk, D. (2000, July). The digital Michelangelo project: 3D scanning of large statues. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (pp. 131-144). ACM Press/Addison-Wesley Publishing Co.

¹¹⁰ Blais, F., Taylor, J., Cournoyer, L., Picard, M., Borgeat, L., Godin, G., ... & Lahanier, C. (2007). Ultra high-resolution 3D laser color imaging of paintings: the Mona Lisa by Leonardo da Vinci.





Figure 2.3: Top: Michelangelo's David in the process of being scanned.
Middle: (left) a colour photo of the statue, (right) an artificial rendering using the scan data
Bottom: Close up view of the statue's right eye. (Left) a photograph, (right), an untextured scan

All images: Marc Levoy/Digital Michelangelo Project

2.2.4 Projected/structured light systems

Projected or structured light scanners lie somewhere between passive and active systems. Like photogrammetry, they capture 2D images of the light reflected from the surface to be measured; like laser scanners they project the light themselves, though in this case the light is a two dimensional image rather than a single spot. They work on the principle that a stripe (or other known pattern) of light projected onto a three dimensional surface will appear distorted when viewed from a viewpoint other than the projector's¹¹² (Figure 2.4). Knowing the relationship between projector and sensor, and the distortion of the stripe, the underlying geometry can be recovered. If the entire surface (or a particular section of the surface) is covered in a pattern of stripes, the geometry of the entire area can be calculated. More complicated patterns, such as fringe projection and changing patterns like a Gray code¹¹³ can improve accuracy.

Structured light scanning can be a cheap and fast option, and can be performed using a standard projector and digital camera, or an integrated system such as those from Breukmann¹¹⁴. The method has proved useful for many cultural heritage applications^{115,116}. Specifications vary very much depending on the camera and lens used and the resolution of the projected pattern. A telephoto lens, for example on a Breukmann triTOS scanner only captures an area 80x60mm per scan, but has a

¹¹² Devrim, A, Grun, A, Breuckmann, B and Lahanier, C (2007) High definition 3D-scanning of arts objects and paintings. Institute of Geodesy and Photogrammetry, ETH Zurich.

¹¹³ <http://mathworld.wolfram.com/GrayCode.html>

¹¹⁴ <http://aicon3d.com/products/breuckmann-scanner.html>

¹¹⁵ Acka, D, Grun, A, Breuckmann, B & Lahanier, C (2007) High definition 3D-scanning of arts objects and paintings. Institute of Geodesy and Photogrammetry, ETH Zurich.

¹¹⁶ Alkis, Z. (2006). 3D modeling of the Weary Herakles statue with a coded structured light system. Institute of Geodesy and Photogrammetry, Swiss Federal Institute of Technology.

horizontal resolution of approx. 60µm.

Colour is not automatically captured by structured light scanners (as the object needs to have a pattern projected on it), but if a colour camera is used to capture the fringe projection, photographs can be taken from the same position as the scan, allowing images to be draped over the resulting point cloud relatively accurately and easily.




Figure 2.4: a) example of a fringe projection b) how an object distorts the pattern. Image from people.stud.edu.sg/~chenlujie/doc/FringeProj.pdf

2.2.4.1 *Kinect and low cost scanners*

The Microsoft Kinect was a gaming accessory for use with (and often packaged with) XBox360 console, and was launched in autumn 2010, during the period of this research. Based on a sensor from Israeli company PrimeSense, it consists of a projector which shines a pattern of infra-red dots, an infra-red camera and an RGB camera. Like any structured light scanner, by measuring the distortion of the projected pattern, the Kinect can reconstruct a depth map of its environment. While the raw output of the Kinect is low resolution (VGA, or 640x480 pixels) and extremely noisy, it's importance to 3D scanning, or more generally, the 3D industry, has been enormous. Because of the

reach of the Xbox, economies of scale enabled Microsoft to manufacture the Kinect for a relatively low cost (it retailed for well under £100), placing millions of projected light depth cameras in the hands of users and potential hackers. Within months of its appearance, users had reverse engineered the drivers and released them to the open source community¹¹⁷, Microsoft released the official Software Development Kit (SDK)¹¹⁸ in winter 2011.

Within months of the SDK being released, researchers had turned the Kinect into a 3D scanner using SLAM (Simultaneous Localisation and Mapping) algorithms¹¹⁹. Features are extracted from each frame taken by the RGB camera, and using the depth map, assigned a position in 3D space. The features are matched in adjacent frames, and as each subsequent frame is registered with the previous one, a coloured 3D point cloud is gradually built up.

Thus, soon after the Kinect's release, software became available¹²⁰ that turned the cheap depth sensor into a 3D scanner, and whilst aimed at the burgeoning 3D printing market, the ease with which a user can create convincing, if low quality, 3D models, (Figure 2.5) was a huge shift in 3D scanning.

¹¹⁷ http://openkinect.org/wiki/Main_Page

¹¹⁸ <http://research.microsoft.com/en-us/news/features/kinectforwindowssdk-022111.aspx>

¹¹⁹ Engelhard, N., Endres, F., Hess, J., Sturm, J., & Burgard, W. (2011, April). Real-time 3D visual SLAM with a hand-held RGB-D camera. In Proc. of the RGB-D Workshop on 3D Perception in Robotics at the European Robotics Forum, Västerås, Sweden (Vol. 180).

¹²⁰ Ie, ReconstructMe (<http://reconstructme.net/reconstructme-ui/#features>) which is a paid for product, or Skanect (<http://skanect.occipital.com/>) which is free for non-commercial use.



Figure 2.5: A 3D model of Jeremy Bentham's auto icon created by UCL's 3DImpact group using a Kinect scanner and ReconstructMe software. Further work was carried out on the model in Meshlab. See <http://www.youtube.com/watch?v=Gy7muFzA1e0>, and more information found at <http://uclgeomatics.com/2012/11/09/jeremy-bentham-in-3d/>

Whilst the models produced by the Kinect or similar low cost sensors are too rough to be used in many cultural heritage applications, they are a low cost and quick way of producing quick models for visualisations (or printing) and could potentially have some use in public facing applications. Perhaps a more important effect of the sudden availability of low cost 3D scanners is the increase in awareness of 3D scanning and corresponding demand for user-friendly software for both processing and sharing 3D models online. The ease of use has also seen Kinect and Kinect-like sensors used in children's workshops at CH institutions like the British Museum¹²¹.

¹²¹ Pers. Communication with Katherine Biggs, former Education manager for Samsung Digital Discovery Centre at the British Museum

2.2.4.2 Hand held scanners

2015 has seen the launch of two hand held scanners, the Faro Freestyle¹²² and Artec Spider¹²³. Both scanners operate on the same principle as the Kinect – using structured light to create a depth map image, and SLAM algorithms to register each frame into a coloured 3D point cloud. In contrast to the Kinect, however, both scanners produce accurate and dense point clouds, offering resolutions of 1-200µm at .5m range. At £9,000 (The Faro¹²⁴) or £11,000 (the Artec¹²⁵), the scanners are not cheap, but perhaps within the budget of some CH institutions. The Artec scanner in particular is being marketed as suitable for CH applications¹²⁶, both public facing and professional. This type of scanner does have many advantages that may make them attractive to people involved in cultural heritage; they create full colour point clouds, are portable (ie, they can capture objects in situ), easy to use¹²⁷, versatile (they can be used to scan objects at a variety of scales, as well as small rooms) and they are cheaper than many of the alternatives.

2.2.5 Reflectance Transmission Imaging

RTI (Reflectance Transmission Imaging) is a technique designed to produce computer-generated imagery of surfaces with a high degree of photo-realism¹²⁸. It could be said

¹²² <http://www.faro.com/products/3d-documentation/handheld-3d-freestyle-3d/features#main>

¹²³ <http://www.artec3d.com/hardware/artec-spider/specifications/>

¹²⁴ <http://surveyequipment.com/faro-scanner-freestyle-3d/>

¹²⁵ <http://www.artec3d.com/hardware/artec-spider/specifications/>

¹²⁶ <http://www.artec3d.com/applications/> , http://www.artec3d.com/case_studies/The+Digital+Soane-+How+Artec+3D+scanners+bring+together+new+technologies+and+contemporary+art_31833

¹²⁷ I have had a demo with the Faro Freestyle, and while it is perhaps not as easy to pick up and use as some of the videos may suggest, with a little practice capturing objects becomes relatively simple.

¹²⁸ Malzbender et al, Polynomial Texture Maps, SIGGRAPH 01, Proceedings of the 28th annual conference on Computer graphics and interactive techniques, pp 519-528, NY, 2001

to be an improvement and refinement of bump-mapping¹²⁹, in which for a given pixel on a 2D texture, an accurate surface normal is recorded as well as the RGB colour, allowing for a more realistic rendering under arbitrary illumination.¹³⁰

For RTI, the object is placed under a hemispherical dome with the surface to be captured pointed upwards towards the dome's apex, where a camera or equivalent imaging device is mounted. In the case of the dome constructed by Lindsay MacDonald¹³¹ a member of UCL's 3DImpact group, there are 64 LED flashes arranged around the dome in three rings, and whose geometrical centroids have been measured to within three millimeters. By photographing the object 64 times, each illuminated from a different, but known angle, a reflectance map of the object's surface is built up, allowing the object to be accurately rendered as if illuminated by an artificial light source from any direction.¹³²

The ability to simulate a raking light effect means that in cultural heritage, the technique is often used to examine two dimensional surfaces with very shallow three dimensional features; for example faint inscriptions or carvings¹³³, or painted

¹²⁹ ibid

¹³⁰ Blinn, JF, (1978) Simulation of Wrinkled Surfaces, SIGGRAPH '78 Proceedings of the 5th annual conference on Computer graphics and interactive techniques, pp 286-292

¹³¹ Macdonald, L & Robson, S (2010), Polynomial Texture Mapping and 3D Representations, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII, Part 5 Commission V Symposium, Newcastle upon Tyne

¹³² In fact, the dome is not required. By placing a reflective sphere next to the object being imaged, a hand held or movable flash can be used and the light's position calculated from the position of the specular highlight on the sphere. See Mudge et al, New Reflection Transformation Imaging Methods for Rock Art and Multiple-Viewpoint Display, The 7th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST (2006), Ed. Ioannides, M et al

¹³³ Graeme Earl, Kirk Martinez, Tom Malzbender, (2010) Archaeological applications of polynomial texture mapping: analysis, conservation and representation, Journal of Archaeological Science, Volume 37, Issue 8, August 2010, pp 2040-2050

surfaces¹³⁴. It should be noted that while the technique produces a highly accurate rendering of the surface, and a strong three dimensional effect when the object is rendered in real time with a moving (and often user-controlled) light source, the images produced are strictly two dimensional, and thus the technique is often referred to as '2.5D'¹³⁵. Thus while the viewpoint can be altered, the illusion of three dimensionality disappears if the reconstructed viewpoint moves too far away from orthogonal to the surface (ie, too far from the original camera position)¹³⁶.

2.3 Strengths and limitations of digitisation methods in public facing cultural heritage applications

2.3.1 Common weaknesses

Having listed various 3D digitising technologies, it is important to point out the major weaknesses that, unfortunately, are shared by all these methods. All the methods discussed above rely on measuring some form of electromagnetic radiation reflected from the object. Therefore, any property of the object, such as the micro-geometry of its surface or the optical properties of its material(s) that affect or somehow modify the reflected radiation will inevitable effect the ability of the scanning method to measure accurately. Thus certain objects or classes of objects will prove problematic

¹³⁴ Padfield, J., Saunders, D., & Malzbender, T. (2005). Polynomial texture mapping: a new tool for examining the surface of paintings. ICOM Committee for Conservation, 1, 504-510.

¹³⁵ García Fernández, J. (2013) Interpretation of topographic data from 'shape from shading' method, application in Villagarcia De Campos Castle, International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XL-5/W1, 2013 3D-ARCH 2013 - 3D Virtual Reconstruction and Visualization of Complex Architectures, 25 – 26 February 2013, Trento, Italy

¹³⁶ This, and some of the potential of ptm imaging, can be seen using the University of Leuven's online interactive viewer at <http://homes.esat.kuleuven.be/~mproesma/mptmp/vandyck/WebPlayer.html> acc. 4/6/14 (requires the Unity plug-in)

for any common 3D digitising technology; in fact, it is reasonable to say that, with very few – if any – exceptions, *every* cultural heritage object will exhibit one or more properties that make it difficult to capture.

Perhaps the most obvious of these properties is complex geometry. Trivially, a scanning method cannot capture what it cannot see: if parts of the object occlude other parts, and the scanner, camera, sensor etc. cannot be positioned in such a way that the area is visible (in photogrammetry, from more than one angle), there will necessarily be missing data and a ‘hole’ in the model. This can be caused by the actual small scale geometry of the object so that occlusion is unavoidable, for example there may be cracks or small ‘trenches’ in the object that are too dark to be photographed, or too narrow for a triangulation scanner to image¹³⁷ (Figure 2.6). There is a trade-off between accuracy and a scanner’s ability to access an entire surface; a longer baseline between laser and sensor can provide greater accuracy, but necessarily limits the complex geometry that can be accessed. Alternatively, occlusions may be due in part to geometry and part due to practicalities; if for example a statue is being imaged in situ, there may simply be areas that cannot be reached due to practical considerations of lighting, stepladders and the care required when working with particularly rare or fragile objects.

¹³⁷ A trap that I fell into in a lot of early laser scanning is that just because you can see the laser spot on a surface, this does not mean that surface is being recorded; the sensor, necessarily offset and viewing at some other angle than the laser beam, must also be able to see the spot – so even if the laser can ‘reach’ the bottom of a crack, the sensor will be unlikely to see it.




Figure 2.6: Example of an occlusion due to small scale geometry. Even though the laser can 'see' the crack in the object, and the laser spot will be visible on the surface, the sensor cannot image the spot and therefore no data will be collected

As well as small details, objects with large-scale concave features, such as an interior, can be difficult to capture due to the difficulties of lighting and imaging the insides (see section 6.4.3.2 for an example).

Non rigid objects (such as those made from fabric¹³⁸) also present

considerable difficulties. For most purposes, objects will need to be imaged from a variety of viewpoints to capture a full 360° model, and if the object moves or changes shape between scans, it will be impossible to register the scans together. Very thin objects (for example, the brim of a hat) can also be problematic as it may be impossible to image both sides of the object's surface at one time (for example, one cannot capture both sides of the object in one photograph), and since the two surfaces will necessarily share no features (or perhaps very few at their interface), registering the two scans can be difficult – see the lusterware bowl in section 5.3 for an example.

Whilst problems with geometry are common, it is often the materials of the object being imaged that can cause even greater problems. All the digitising methods discussed above measure the amount of electromagnetic radiation reflected from a

¹³⁸ For example, the Punch puppet scanned for the 3DCoform project: <http://exhibition.3d-coform.eu/?q=node/65>

surface, and, in general, measure the reflectance for just one angle of reflection. Thus a single intensity (or colour) is recorded for each spot on the surface. This is not a problem if the surface exhibits perfect diffuse reflectance, where the same amount of radiation is reflected irrespective of the angle at which it is measured. This is known as a ‘Lambertian’, perfectly diffuse, or perfectly matte surface. These materials are very rare and are found in expensive calibration objects, for example colour checker charts used to colour correct imagery (see 6.5.1.2) or Spectralon¹³⁹, as used in the Arius scanner’s colour correcting white cube.

In the real world, object surfaces will never be *perfectly* diffuse and will almost always exhibit one or more material properties that cannot be captured by the techniques mentioned above¹⁴⁰. Almost all objects and materials will exhibit some amount of shininess or gloss. On these surfaces, the amount of light reflected depends to a greater or larger extent (depending on the degree of gloss) on the relationship between the angle of incidence of the ray and the viewing angle. For very shiny objects, the magnitude of the reflected light peaks very sharply around a certain angle, causing bright spots known as specular highlights. Specularity is defined by various parameters representing the strength of the reflection and its sharpness (how quickly the peak reflection tails off away from the specular angle), and also the colour – whilst in general specular reflections take the colour of the light source, different metals

¹³⁹ <https://www.labsphere.com/products/diffuse-reflectance-coatings-materials/spectralon-reflectance-material/>

¹⁴⁰ Some complex surface properties could potentially be extracted from the raw data gathered by normal digitisation methodology, but this is not part of their normal workflow and is a complex, decidedly non-trivial operation.

reflect different coloured highlights due to the photoelectric effect. It should also be noted that due to the fact that most materials will exhibit *some* level of gloss, the measurement of diffuse colour by the techniques above may give different results depending on the particular arbitrary angle chosen for the measurement. The outcome of this is that two scans taken of the same object from even slightly different angles will often record subtly different colour values, and when the two scans are registered, the colour discontinuity is clearly visible. An extreme example of this can be seen in Figure 5.17, but it is a problem present in objects that exhibit very little shininess, such as the elephant's tooth in Figure 2.7.




Figure 2.7: Laser scan of an elephant's tooth made using the Arius scanner. Colour discontinuity between two scans can be seen to the right of point 1, with darker points from one scan overlay lighter points from another, and at point 3, where the discontinuity can be seen as an abrupt colour difference along a straight line

Other materials demonstrate anisotropy, where the surface appears different depending on the angle at which it is viewed. Examples include cloth with a 'nap', like velvet which will appear very different when viewed from different directions. There are many other complex properties of surfaces not captured by the techniques above. These include (but are not limited to) translucency and transparency (where some or all of the radiation is transmitted rather than reflected by the surface), subsurface scattering (where some of the radiation penetrates the material, reflects internally and is emitted from a point some distance from where it originally hit the surface), iridescence (colour changes due to interference effects caused by microscopic layers, as seen on soap bubbles, oil films and shells) and fluorescence (where light is absorbed and then emitted at a different wavelength).



Some of these properties can be approximated and reproduced when rendering the model. For example Phong shading¹⁴¹, introduced in the 1970s and still very popular in computer graphics, includes values for a material's diffuse and specular properties.

With a digitised object, having measured an object's diffuse properties, the specularity can be approximated in the renderer by trial and error (and aesthetic judgement), as in the case of the lustreware bowl in section 5.3.3. Alternatively, one could use previously measured values for the relevant material, though things become more complicated for objects made up of a variety of materials, and/or materials in various states of wear.

Simple models such as Phong shading do not capture the other properties mentioned above, and whilst these can be approximated in rendering, the more complex properties require equally complex rendering techniques (see, for example sub-surface scattering in Jensen et al¹⁴²). Also, despite being able to render these effects realistically, it does not solve the problem of measuring them in the first place.

These issues are of special relevance to cultural heritage applications – particularly public facing ones – where a convincing visual simulation of the object is required, and where the 3D model is intended to act as a digital surrogate with respect to visual

¹⁴¹ Phong, B. T. (1975). Illumination for computer generated pictures. *Communications of the ACM*, 18(6), 311-317.

¹⁴² Jensen, H. W., Marschner, S. R., Levoy, M., & Hanrahan, P. (2001, August). A practical model for subsurface light transport. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (pp. 511-518). ACM

inspection. Indeed, for many CH objects, it is the complex interaction with light that makes them interesting. Referring again to the objects digitised in this research, certainly for the Lustreware bowl and the Courtauld bag, it is, in the first case the interplay of metallic specular highlights in the glaze that gives the object both its name and its aesthetic appeal, and the bag's highly polished silver contrasting with the black ground contributes much to its beauty. Cultural heritage objects also come in a wide variety of materials; metals, gems, glass, marble, jade, organic materials such as feathers or fur, polished wood and stone are just some of the materials which exhibit complex surface characteristics and which cannot, therefore, be captured perfectly by traditional methods¹⁴³. New techniques are currently in development which aim to solve these problems in a cultural heritage context, and I shall look at them in section 2.9.

2.3.2 Strengths and weaknesses

Whilst it is possible to examine the strengths and limitations of individual digitisation methods, it must be stressed that a real evaluation of the different techniques can only be done on a case by case basis, and while it is possible to generalise over certain classes of objects or certain types of digitisation projects, the decision on which technology to use will depend to a great extent on the particular nature of the object in question and the specific desired output.

¹⁴³ I refer to this as the Goldilocks problem – there is only one way an object can be ‘just right’ for scanning, but an infinite number of ways it can be problematic... Unfortunately, not many cultural heritage objects are perfectly diffuse spheres!

2.3.2.1 Accuracy and resolution

For applications that require high levels of accuracy, for example deformation monitoring, a close range triangulating laser scanner would traditionally have been the best choice. Whilst photogrammetry can yield results as accurate as the best laser scanners¹⁴⁴, this involves more advanced methods using stereo pairs¹⁴⁵, properly calibrated cameras etc.

Resolution is equivalent to point density, or how many points are measured per a specific area. The resolution required for a project depends largely on the size of the object to be captured, how close an inspection of the model you wish to allow the user and the dimensions of the smallest detail you wish to capture. Laser scanners and some structured light systems can achieve point spacings of up to 50µm, equivalent to 400 points per mm², time of flight scanners will achieve point spacings of several mm at a range of 10m.

2.3.2.2 Scale

Related to accuracy, scale, or the ability to make 'real world' measurements from the virtual model is something that comes free with a laser scan. Laser scanners record absolute measurements when they record an object, so that measurements made between corresponding points on the real and virtual objects should yield the same

¹⁴⁴ Ahmadabadian, A. H., Robson, S., Boehm, J., Shortis, M., Wenzel, K., & Fritsch, D. (2013). A comparison of dense matching algorithms for scaled surface reconstruction using stereo camera rigs. *ISPRS Journal of Photogrammetry and Remote Sensing*, 78, 157-167.

¹⁴⁵ Using stereo pairs involves taking every image twice, using a pair of cameras rigidly attached to each other, with a known distance between the two principle points. This added information allows for more accurate measurement and can generate a scale for the model. See: Ahmadabadian, A. H., Robson, S., Boehm, J., Shortis, M., Wenzel, K., & Fritsch, D. (2013). A comparison of dense matching algorithms for scaled surface reconstruction using stereo camera rigs. *ISPRS Journal of Photogrammetry and Remote Sensing*, 78, 157-167.

results. Photogrammetrical models are generally scale-free and any measurements made on the virtual model, whilst hopefully proportional to the real world, are purely arbitrary. To give an example, a photogrammetrical model of a car could have been generated from a set of photos of a toy or a real vehicle, there would be no way of telling simply from examining the model.

Whilst accuracy (and scale) may not be an important consideration in many public-facing applications, it could be an issue in a model created for multiple purposes. Scale can be introduced to a virtual model either by using stereo pairs (the known baseline between the cameras providing the necessary information) or by simply including something with a known measurement in the imaging, like a scale bar or simply a ruler. Again, scale may not be of primary importance in a public facing application, though it is advisable to provide some information on object size (even if it is just a description of the dimensions of the original object).

2.3.2.3 Colour

Ignoring the many issues described in 2.3.1, this discussion will consider only diffuse colour measurement. Most laser scanners do not automatically record colour, and those that do, such as the Arius, can be very expensive. However many scanners, such as the lower-end Next Engine, include RGB cameras that can be used to project texture on to the raw geometry, and long range ToF scanners can feature either built in cameras, or brackets to which cameras can be attached, to record colour. These methods are not quite as accurate (in terms of mapping colour to the surface, rather than the accuracy of the colour information itself) as a multi-wavelength laser scanner

as there may be an offset between the camera and laser position - the quality of built in optics can be poor as well. Projected light scanners will also use a camera to record colour, which is then projected onto raw geometry though here, as the camera that captures colour is often the same one that records geometry, there is no offset problem. Photogrammetry, of course, provides colour automatically, though in all cases where photography is used to add colour to an object, the images should be properly colour calibrated to ensure accuracy. (See 6.5.1.2 for an example of the colour calibration workflow for a photogrammetry project)

2.3.2.4 Versatility

Often the size of the object to be recorded will govern the choice of technology. Close-range scanners' recording volumes range from under 1m³ for 'fixed' scanners to several m³. Movable systems such as those with optical CMM systems such as the Nikon Metris can be used to capture much larger objects by stitching multiple scans together.¹⁴⁶ In this case the maximum size is determined by time available and the amount of data that can be captured and processed.

For buildings (interior or exterior), or other large volumes such as archaeological digs, time of flight scanners, with ranges of 100s of metres are the logical choice. For objects in between, the decision to use close-range or ToF scanners should probably be made on the desired output; a close range scanner will provide a much denser and more detailed point cloud than a ToF scanner, allowing close inspection of the 3D model, but

¹⁴⁶ Hess, M., Robson, S., Millar, F. S., Were, G., Hviding, E., & Berg, A. C. (2009). Niabara-the western solomon islands war canoe at the british museum-3D documentation, virtual reconstruction and digital repatriation. InVirtual Systems and Multimedia, 2009. VSMM'09. 15th International Conference on (pp. 41-46). IEEE.

it will take considerably longer. For example, scanning an 11m war canoe with a handheld triangulation scanner¹⁴⁷ took many days; the same object could potentially be captured in an hour or less using a ToF scanner, though the resulting point cloud would be less dense and less accurate.

Projected light scanners generally share similar specifications to close range scanners, depending on the power of the projector and the camera used. As the projected pattern is shone over a larger area, it will necessarily lose both intensity and sharpness, leading to a trade-off between scanning volume and accuracy.

Photogrammetry is the most versatile method, as the size of the object to be captured is limited only by the available camera lenses. If you can clearly image an object, whether it is something tiny requiring macro lenses, or an entire building requiring a wide angle¹⁴⁸ lens, from multiple angles, you should be able to process a 3D model.

2.3.2.5 Geometry and texture

As discussed above, all the methods under investigation will struggle with complex small scale geometry, where parts of the object shadow or occlude other parts. For larger scale geometry, the portability of the scanning technology becomes important.

As geometry becomes more complex, the number of scanning angles required to capture an entire surface increases and with fixed scanners, such as the Arius, this can cause issues with fragile objects, as it would need placing in a variety of positions, some of which may not be feasible. A hand-held scanner, on the other hand, can

¹⁴⁷ Hess (2009) *ibid*

¹⁴⁸ Though images taken using wide angle lenses will need to be calibrated to remove distortions first

capture an effectively infinite number of scan angles without moving or handling the object. Similarly, photogrammetry is fairly versatile when it comes to geometry as the camera, depending on the particular set-up, available tripod and ability to focus, can move freely - though with complex geometry care must be taken that every point on the surface is captured in, preferably, a minimum of three images.

Objects with little or no surface texture (blank walls, ceilings, areas of uniform colour etc.) will cause problems for photogrammetry, as it needs to match visible features in multiple images to calculate 3D points. Laser scanners and projected light systems¹⁴⁹ do not have this issue.

2.3.2.6 Portability

In cultural heritage, many of the objects that need imaging are rare or fragile and there may be issues and cost implications around transporting the object. Larger objects such as statuary or architectural features obviously cannot be moved to the scanner, so the ability to capture an object in situ can prove key. The amount of clearance around the object and the ability to access it from multiple angles can govern what methods are suitable. Again, photogrammetry would appear to be the most portable technology and therefore the most versatile, followed by handheld scanners.

2.3.2.7 Necessary skills

Whilst laser scanners are complicated pieces of technology, their operation is relatively

¹⁴⁹ Some projected light systems, or more accurately some software using projected light cameras, may have trouble with textureless objects if they use the output from an rgb camera to do feature matching between frames. Skanect for the Kinect uses this method, for example.

simple. From personal experience, the majority can be operated after one or two hours of training¹⁵⁰. However, it should be noted that scanning objects, using whichever method, is as much art as science, and experience can help both achieve the best results and potentially speed up the recording process. Each technology, individual scanner – and indeed object – will have its own idiosyncrasies, however, and experience with one may not necessarily help with another. Photogrammetry requires photographic skills that may already be present within an institution, and thus has a low barrier to entry, though as mentioned above there are more advanced techniques and methodologies (that will be covered in later chapters) that can improve results.

It should also be noted that capturing the object is merely the first step, once acquired the raw data must be processed (to remove unwanted data or noise, correct scanning errors, interpolate missing data etc.), and then prepared for dissemination (rendered as video, inserted into an interactive viewer etc.). The processing stage in particular is a time consuming, and, from personal experience, difficult process.¹⁵¹ The same rules apply; it is a craft rather than a science, with a steep learning curve, and again, there is no real substitute for practice and experience.

2.3.2.8 Cost

It is difficult to place a definitive cost on any particular technology, as hardware prices can vary hugely even within a particular class – a triangulation laser scanner can cost

¹⁵⁰ This is the time it takes to train to use the machine; as ever, to become truly skilled takes practise and experience.

¹⁵¹ For examples of complete workflows for creating 3D content, from capture through processing to dissemination, see chapters 4, 5 and 6

anything from a few thousand pounds (the NextEngine), to many hundreds of thousands (the Arius). Time of flight scanners will generally cost between £30k and £100k, whilst the new generation of handheld scanners, at around £10k are more realistic options for small CH institutions (see 1.4). There is no single solution that fits all potential capture targets, however, and thus photogrammetry, being both the cheapest option (under £1000 for a perfectly adequate camera and lenses, if the institution does not already possess them), and the most versatile, is certainly the most cost-effective.

This does not take into account the cost of software, however, and most scanners will come with their own point cloud processing applications. Photogrammetrists have plenty of options to choose from when it comes to creating their model, from free (VSFM, 123D Catch) to commercial (Agisoft Photoscan), and there are also free options when it comes to processing 3D models (such as CloudCompare or MeshLab).

It should also be noted that time (and related personnel costs) is potentially the largest expense in a digitisation project. To give two examples from this research: the Science Museum project with ScanLAB took five nights or approximately 100 man-hours for capture, and 3200 man-hours for processing. The Courtauld Bag took approximately 10 hours in total for capture but around 100 hours to process the model(s), not including rendering. These may be excessive (the Science Museum was a huge project and the Courtauld bag a very recalcitrant subject), but are indicative of the investment that may be required beyond hardware and software.

2.4 Assessing 3D digitisation technologies: Overall conclusions

For applications, where visual fidelity is perhaps more important than geometric accuracy, and certainly those involving object-sized capture, it is hard to look beyond photogrammetry. Its low cost, versatility, availability, and the suitability of its outputs for public facing applications means that there would need to be a compelling reason not to recommend this method to cultural heritage institutions interested in performing 3D capture. It should be noted that the experiments in this thesis (chapters 4, 5 & 6) are presented in chronological order, and the capabilities that allow low cost photogrammetry emerged during the course of the research. With the benefit of hindsight, if the research was starting today then photogrammetry would have been the first choice for all the object-scanning projects. This would not be due to the cost of the alternatives (photogrammetry vs the Arius laser scanner), but due to the ease of use and suitability and quality of output.

But, once again, it is worth reiterating that there is no ‘one size fits all’ solution and that every digitisation project needs to be assessed individually and a suitable methodology chosen that is determined by both digitisation target and desired output. These conclusions are also predicated on the fact that there is a single desired output for the digitisation, though there are also good arguments for adopting a ‘SOAP’, or “scan once for all purposes” attitude¹⁵², where a single model is made from which others can be created (ie, models for archiving, for research, for on- and offline dissemination etc.). In this case, photogrammetry may still be the preferred option,

¹⁵² See UKOLN good practise guide: <http://www.ukoln.ac.uk/interop-focus/gpg/DigitisationProcess/>

though the arguments for it over, say, high quality laser scanning are perhaps less convincing.

2.5 Software

2.5.1 Processing

Whilst the majority of existing software for processing 3D point cloud models is still very much rooted in a surveying or geomatics (Faro Scene¹⁵³, Leica Geosystems Cyclone¹⁵⁴) or engineering and metrology (GOM Inspect¹⁵⁵ and CloudCompare¹⁵⁶) background, there are indications that as 3D scanning is becoming more accessible and the scanners themselves are being marketed more towards both hobbyist and cultural heritage applications, the software is also evolving with a greater focus on aesthetic functionality and the ability to manipulate colour and texture. See for example Artec Studio 10¹⁵⁷, supplied alongside the Artec handheld scanners which includes some automated features for improving texture. However, as mentioned above, from personal experience processing 3D models is considerably more time consuming (and potentially frustrating) than the capturing procedure, and this must be taken into account when costing a project. Whilst speed does improve with operator experience, different objects may require different techniques and often compromise is required both due to time constraints and the limits of the software being used. Whilst some functionality has been automated to a lesser or greater degree, the ultimate aim of

¹⁵³ <http://www.faro.com/en-us/products/faro-software/scene/overview>

¹⁵⁴ http://hds.leica-geosystems.com/en/Leica-Cyclone_6515.htm

¹⁵⁵ <http://www.gom.com/3d-software/gom-inspect.html>

¹⁵⁶ <http://www.danielgm.net/cc/>

¹⁵⁷ <http://www.artec3d.com/software/studio/>

producing an authentic and aesthetically consistent 3D model is still very much a manual process. A basic workflow indicating the level of automation of each step is shown in Table 1, much more detailed workflows are covered in chapters 4, 5 and 6.

Process:	Registration	Hole filling	Texture processing	Filtering
Description	Taking multiple scans and unifying them into a single unified model	Filling in areas with missing or bad data with interpolated points	Correcting colour discrepancies between scans, correcting textures on hole filling areas, correcting for specularity	Preparing the model for rendering by reducing the number of points (and file size) and evening point density across the model
Level of automation	Highly automated, though may need the user to select several common points on two or more scans. Close visual inspection is required to check alignment, and process may need to be repeated	Small holes can be filled automatically, larger holes or holes in noisy data may require an iterative process	A largely manual process, and requires tools not present in many software options.	Largely automatic, but usually an iterative process depending on the rendering method.
Availability	Present in the vast majority of point cloud processing software	Present in the majority of software, though different algorithms may produce different results	Still quite rare, though hopefully becoming more common	A common function in point cloud software

Table 1: Point cloud processing workflow

2.5.2 Dissemination

As well as low cost scanning and easily available photogrammetry, one of the key technologies to emerge in the last five years is WebGL, a JavaScript API for rendering

graphics – including 3D graphics – via a web browser. Released in 2011¹⁵⁸, it allows 3D models to be displayed natively in most modern browsers¹⁵⁹, ie, without the necessity for plug-ins or downloads. For public facing applications, the ramifications are enormous: whereas before WebGL, users who wished to access 3D models via the internet needed to download extra software, and/or were restricted to relatively primitive pseudo 3D apps such as QuickTimeVR¹⁶⁰, today it is a relatively simple matter to embed complex 3D models in web pages. Section 2.7 examines some state of the art applications available today which use WebGL to display 3D models of cultural heritage objects.

2.5.3 Virtual reality

It is beyond the scope of this research, but nevertheless, the imminent arrival of affordable, functional and consumer-friendly virtual reality (VR) devices such as Google Cardboard¹⁶¹, Oculus Rift¹⁶² and HTC's Vive¹⁶³ could dramatically change how 3D models are consumed in the future. Whilst one would assume capture methods for, and potential uses of, 3D cultural heritage models will not be dramatically altered by the change in delivery method, VR does have the potential to profoundly alter the 3D landscape in the next few years.

¹⁵⁸ <https://www.khronos.org/news/press/khronos-releases-final-webgl-1.0-specification>

¹⁵⁹ The number of browsers supporting WebGL applications has grown steadily over the last five years with newer versions of all the major desktop and most mobile browsers allowing interactive 3D applications: <http://caniuse.com/#feat=webgl>

¹⁶⁰ F. Gabellone, M.T.Giannotta (2005) Realtime 3D multimedia system for the distance visiting of cultural heritage. A case study on the chamber tombs in Via Cristi, Taranto. CIPA 2005 XX International Symposium, 26 September 2005, Torino, Italy

¹⁶¹ <https://www.google.co.uk/get/cardboard/>

¹⁶² <https://www.oculus.com/en-us/>

¹⁶³ <http://www.htcvr.com/>

2.5.4 Point clouds or meshes?

As all the capture methods discussed above measure discrete points, their raw output will be a ‘point cloud’ – literally, a cloud of measured points. All the projects undertaken in this thesis use point clouds as the final output as well. However, the majority of projects using 3D models (including all those mentioned in the next section) use a *meshed* model as the output, and it is worth discussing the pros and cons of both methods.

A mesh model is constructed of polygons (almost always triangles, and, certainly in the context of digitised objects, usually textured) and offer several advantages over point clouds. A simple way of generating a mesh from a point cloud would be to treat every point as a vertex, however, this would be inefficient; whilst the point density required to represent a large flat area is, from a visualisation standpoint the same as that required to represent a curved or bumpy surface, a single triangle (or small number of triangles) can cover a large flat area with very little reduction in geometrical detail. Thus, an efficiently meshed model created from a point cloud can contain significantly less triangle vertices than the original points, resulting in a smaller file size and a more efficient rendering. However, it should be noted that depending on the particular geometry of the object (the more complex, the less efficient the mesh) and the target file size, some geometrical information will inevitably be lost in the meshing process.

The other key point about meshed models is they have been the standard format for

rendering 3D computer graphics for many decades¹⁶⁴, and therefore not only are software methods for processing and rendering meshed models far more common and well developed than the point cloud equivalents, computer hardware in the form of specialised graphics chips (graphical processing units or GPUs) are specifically designed to render triangles quickly and efficiently¹⁶⁵.

These three reasons; the more efficient modelling and smaller file sizes of meshes¹⁶⁶, the mature nature of the technology leading to more and better software solutions for mesh models, and the design of hardware to facilitate rendering meshes make a compelling reason to choose meshes over point clouds.

Why, then, are point clouds used in this thesis? The first reason is simplicity; using point clouds as the output format removes one stage from the workflow – that of meshing the point cloud. Many point cloud processing software applications (GeoMagic, Pointstream, Meshlab etc.) will include some automated process to generate a textured mesh from a point cloud. However this automated process is not always entirely successful and the mesh may need considerable work to make it a coherent model. Some point clouds, such as the bag in 6 are too noisy to make an accurate mesh and in some cases where the point density is too low relative to the

¹⁶⁴ Foley, J. D., & Van Dam, A. (1982). Fundamentals of interactive computer graphics (Vol. 2). Reading, MA: Addison-Wesley, pp. 505-511

¹⁶⁵ Luebje, D & Humphreys, G, How GPU's Work: mowgli.hadassah.ac.il/mod/resource/view.php?id=38019

¹⁶⁶ Note that whilst the triangulation of a point cloud can lead to more efficient file sizes, it will usually -but not necessarily - result in loss of detail/information. The acceptability of this loss must be assessed on a model by model – and visualisation by visualisation basis; in the same way that the level of jpeg compression for an image can be chosen such that the file size is smaller but there is no *discernible* loss of information, similarly a model can be meshed such that there is no *visible* loss of detail.

complexity of the geometry¹⁶⁷ (such as the Science Museum Shipping Gallery scan) it may be impossible to generate any sort of mesh.

A second reason was I had early access to a piece of software designed to render point clouds efficiently and which was integrated into the workflow of the Pointstream software I was using to process the data. Thus, to go from point cloud to a model that could be rendered in a browser was an extremely simple process, again removing much of the incentive to use meshed models.

The third reason is perhaps a more subjective one. As we will see in section 4, 3D models presented in the form of point clouds have a particular aesthetic, and, in my experience and opinion, look better at higher levels of zoom than 3D models presented as meshes. The process of meshing a point cloud, where the aim is to create an efficient model, will necessarily require some ‘averaging’ of the data. Whilst we said above that a small number of triangles can represent an arbitrarily large flat surface, in reality, the triangles will probably represent a *relatively* flat area and depending on the level of efficiency required, this can result in an averaging out of very small scale geometry. Representing curved surfaces require more, smaller triangles, so as geometrical complexity increases, the relative efficiency of a meshed model falls off¹⁶⁸. Whilst some meshes record ‘per vertex’ colour, with colour information recorded at

¹⁶⁷ In the case of the shipping gallery, the internal geometry of the gallery with its many display cases and detailed models made meshing impossible. However, if the same long range scanning methods had been used to capture, say, the outside of a building constructed mainly of planes and large scale geometrical features, even with the same point spacing a meshed model would have been possible.

¹⁶⁸ Kobbelt, L., & Botsch, M. (2004). A survey of point-based techniques in computer graphics. *Computers & Graphics*, 28(6), 801-814.

each triangle vertex (inherited, for example, from the original point cloud, they are usually coloured by storing a single image or several separate images that can then be draped over the mesh, each region on the image corresponding to a particular triangle on the model¹⁶⁹. In this way, colour information can be stored efficiently as well, perhaps in a compressed jpeg image. However, this can cause a noticeable blurring effect when zooming in on a model, breaking the illusion of reality and giving the model an air of artificiality. This can be avoided by generating and storing higher resolution textures or using more complex techniques like bump¹⁷⁰ or normal¹⁷¹ mapping, though this goes against one of the incentives for creating a mesh in the first place; having a more efficient model and smaller file size.

¹⁶⁹ Colour information can be inherited from 'averaged-out' points as well. If a single triangle represents a flattish area defined by several hundred points, the colour information from those points can still be captured in the texture for that triangle.

¹⁷⁰ Kilgard, M. J. (2000). A practical and robust bump-mapping technique for today's GPUs.

¹⁷¹ Heckbert, P. S. (1986). Survey of texture mapping. *Computer Graphics and Applications, IEEE*, 6(11), 56-67.




Figure 2.9: Images of meshed models from the Smithsonian 3D web viewer (see section 872.7 below). Already some texture blurring can be seen on the image on the left, it is more pronounced in the close up view on the right

Point clouds, on the other hand keep their detail at higher levels of zoom, though point size needs to be controlled carefully to avoid transparency. From my experiences showing people point cloud models, I would venture the possibility that the blurring effect one gets with meshes is a familiar effect, whereas the breaking up of a point cloud at high magnifications is, for most, a novel experience. Meshes, on the other hand are common to anyone with any experience of computer graphics, and in fact, the meshes experienced on line are of a much lower quality than one has come to expect from offline applications such as gaming. Point clouds, having no frame of reference, are treated by the user as something other than artificially created mesh models and this may help to reinforce the notion that point clouds are somehow

objective records and that the points are all measured. As far as I am aware, no research has been carried out comparing users' responses to models presented either as meshes or point clouds, and would appear to be a valid line of enquiry.

2.6 Other 3D technologies in cultural heritage: Procedural modelling, constructive solid geometry and voxels

Procedural modelling creates 3D models from an algorithm or set of rules, for example, architectural models could be built up from rules governing parts (windows, doors, walls etc.) and their relations to each other.¹⁷² Procedurally generated models can be created quickly and can be far more efficient than a traditional 3D model where every part of the model must be described in full detail 'from scratch'. The models produced may lack the one-to-one correspondence with reality one might get with a scanned model, and, generally, can only represent regular (presumably man-made) objects that possess a fairly well-defined visual grammar, but the technique can still find use in cultural heritage. Haegler et al¹⁷³ find the lack of objective realism an advantage in archaeological reconstruction. Where the finished result will necessarily involve a certain amount of uncertainty, the ability to generate many alternative, hypothetical, models quickly and easily is a benefit of procedural generation.

Constructive solid geometry, like procedural generation, aims to simplify the creation of 3D content by building up complex models from a set of primitive objects such as

¹⁷² Pascal Müller, Peter Wonka, Simon Haegler, Andreas Ulmer, and Luc Van Gool. 2006. Procedural modelling of buildings. *ACM Trans. Graph.* 25, 3 (July 2006), 614-623. DOI=<http://dx.doi.org/10.1145/1141911.1141931>

¹⁷³ Simon Haegler, Pascal Müller, and Luc Van Gool. 2009. Procedural modeling for digital cultural heritage. *J. Image Video Process.* 2009, Article 7 (February 2009), 1 pages. DOI=<http://dx.doi.org/10.1155/2009/852392>

spheres and cubes. As in the case of procedural generation, the objective one-to-one correspondence of a surface scan is sacrificed for speed and compact file size, though Vilbrandt et al¹⁷⁴ also emphasise the interoperability and sustainability of this method compared to surface scanning.

Another method of 3D reconstruction uses voxels. A voxel is a 3D pixel ('volume element' as opposed to 'picture element'), and any volume can be divided into a 3D grid of voxels in an analogous way to a picture containing a 2D grid of pixels. Voxels are the 'native' format of volumetric (as opposed to surface-) scanners, such as CT, MRI and ultrasound scanners. Voxel rendering is a useful technique in cultural heritage where the interior, as well as the exterior, of an object is to be rendered, for example in the reconstruction of Egyptian Mummies from the British Museum's collection¹⁷⁵

¹⁷⁴ Vilbrandt, C., Pasko, G., Pasko, A., Fayolle, P.-A., Vilbrandt, T., Goodwin, J. R., Goodwin, J. M. and Kunii, T. L. (2004), Cultural Heritage Preservation Using Constructive Shape Modeling. Computer Graphics Forum, 23: 25–41. doi:10.1111/j.1467-8659.2004.00003.x

¹⁷⁵ C Baldock, S W Hughes, D K Whittaker, J Taylor, R Davis, A J Spencer, K Tonge, and A Sofat (Dec 1994), 3-D Reconstruction of an Ancient Egyptian Mummy Using X-Ray Computer Tomography, J R Soc Med 87: 806–808, doi:10.1177/014107689408701229

2.7 Cultural heritage objects online – the state of the art

To provide some context and show how the art has developed in the four or so years since WebGL appeared, one of the first projects to use the new technology was a Jisc sponsored Sheffield Museum project to digitise some of its metalwork collection. Today, just a few years later, the results already appear quite primitive.¹⁷⁶ The models are, by today's standards, very low resolution ('low poly'), with numerous errors and areas like the interior of the cream jug (Figure 2.10) are actually solid 2D textures. The viewer itself allows the user to zoom past both near and far clipping planes, occasionally causing the entire object to disappear.

¹⁷⁶ Battle axe: <http://collections.museums-sheffield.org.uk/view/objects/asitem/360/15/>, Cream jug: <http://collections.museums-sheffield.org.uk/view/objects/asitem/360/15/>



First Previous Next Last
 1 | 2 | - | 99 | 100 | 101 | 102 | 103 | - | 672 | 673



Carving set
 Date Made/Found: 1870-1899
 Material and Medium: steel

Quick Search Go

switch views:



Figure 2.10: One of the objects from the museum of Sheffield's metalwork collection. Above: the full webpage

Right, top: A close up from the full screen view. Degenerate triangles – errors in the mesh – can be seen along the handle, whilst the end of the blade is invisible due to clipping. It is possible to zoom in so far that the entire object becomes invisible.

Right, below: A cream jug, the obvious colour discontinuity clearly showing the problems encountered when digitising a reflective material such as metal.



Compare this to more recent efforts, such as the 3D Petrie Museum¹⁷⁷, the Google Cultural Institute Art Project¹⁷⁸, and the Smithsonian's X3D project¹⁷⁹ (Figure 2.11).

¹⁷⁷ <http://www.ucl.ac.uk/3dpetriemuseum/3dobjects>

¹⁷⁸ <https://www.google.com/culturalinstitute/browse/3d?projectId=art-project&hl=en>

¹⁷⁹ <http://3d.si.edu/>




Figure 2.11: From top, 3D Petrie Museum homepage and viewer; Google Cultural Institute 3D page and viewer and Smithsonian X3D homepage and viewer

The Petrie's objects are all linked by a theme: Egypt, as the Petrie museum is a museum of Egyptology. Google's models are similarly grouped according to the institution they came from. However, there is still little context and the overriding impression is that the models are there, displayed as they are *because these are the*

objects which we have digitised in 3D. A museum is very unlikely to create an exhibition of objects with no thematic link and no explanation as to why these particular objects are there, but this is the impression one gets from these collections. This is even more apparent in the Smithsonian collection, where, partly due, one presumes, to the extremely heterogeneous nature of the Smithsonian's own 19 million objects, the models range from complete burial sites to fossils to the Wright Brother's plane to flowers and animals.

The Smithsonian also provides the user with tools to affect and alter the rendering of the object, for example, they can arbitrarily change the object's material properties such as specularity and reflection to an extreme degree (Figure 2.12). This must inevitably cause tension with the concept that these are reliable, somehow objective records of the objects themselves.




Figure 2.12: Example of the Smithsonian's material editor. Left, a ritual ewer displayed with default values, and right, with maximum reflectance

Both these things, the mode of presentation and the ability to alter the models themselves, reinforces the impression that these are experiments in 3D as opposed to

applications with a museological or cultural heritage purpose. This is not meant as a criticism of any of the projects mentioned above, it is a natural and necessary step at this stage in the development of 3D digitisation and its use in public facing applications – while it is still exciting and novel and where truly mass digitisation of collections is still some way off. The Smithsonian is explicit¹⁸⁰ about their 3D project – developed with the help of Autodesk and still officially in ‘beta’ – being an ‘experimental lab’, designed to investigate the various potentials of public facing 3D content. However, as I shall explore in more detail in the next chapter, much (if not all) of a museum object’s aura is due to the context in which it is presented, and there are reasons to believe the same will be true of virtual objects. Until 3D models are treated, not as experiments, but as tools to achieve particular museological purposes, it is hard to judge their utility as such.

With its ‘tours’, the Smithsonian does, in fact, show how 3D models can be placed in a larger context¹⁸¹. The Repatriation and Replication of The Kéet S'aaxw (Killer Whale Hat)¹⁸² is a prime example. During the tour, the model is visible and the user can interact with it in the normal model viewer, while a multi-media presentation involving text, images and video is shown on the right. Occasionally the tour manipulated the 3D object to draw attention to certain features. It is a novel and engaging experience and, I believe, one that points to the future for public facing CH applications. For the moment however, while the user accesses the tour from the Smithsonian X3d portal, and while

¹⁸⁰ <http://3d.si.edu/about>

¹⁸¹ <http://3d.si.edu/tour-browser>

¹⁸² <http://3d.si.edu/tour/repatriation-and-replication-k%C3%A9et-s%E2%80%99aaxw>

they have the ability to arbitrarily alter the properties of the object, these are still very much experiments in the use of 3D.

2.8 Other uses of 3D content for public engagement

2.8.1 Immersive environments and 'serious' games

While this research primarily deals with 3D digitised museum objects and exhibits *qua* museum objects and exhibits, as opposed to 3D models used in other public engagement applications, for instance as assets in other media such as 'serious games'¹⁸³ or immersive environments (such as virtual museums embedded in worlds such as Second Life¹⁸⁴). Whilst there are obvious differences between a 3D game and an immersive environment, there are similarities both in their method of creation and presentation, for example, using game engines such as Unity or Unreal^{185,186}) The idea of using 3D environments for public engagement has a long history in cultural heritage (for example, see Miller et al, 1992 for an early virtual museum using QuickTime technology¹⁸⁷). However, whilst some 3D content in both games and virtual environments uses digitised assets, much of the content consists of visualisations or artists' recreations. Even when digitised assets are used, often the models must be so severely decimated to work within the game engine that they could be considered

¹⁸³ Herminia Wei-Hsin Din. 2006. Play to learn: exploring online educational games in museums. In ACM SIGGRAPH 2006 Educators program (SIGGRAPH '06). ACM, New York, NY, USA, , Article 13 .

¹⁸⁴ See for example: <http://secondlife.com/destinations/arts>

¹⁸⁵ George Lepouras, Costas Vassilakis (2004), Virtual museums for all: employing game technology for edutainment, Virtual Reality, June 2004, Volume 8, Issue 2, pp 96-106

¹⁸⁶ Eike Falk Anderson , Leigh McLoughlin, Fotis Liarokapis, Christopher Peters, Panagiotis Petridis, Sara de Freitas, (2010) Developing serious games for cultural heritage: a state-of-the-art review, December 2010, Volume 14, Issue 4, pp 255-275

¹⁸⁷ Miller, Gavin, et al. "The virtual museum: Interactive 3d navigation of a multimedia database." The Journal of Visualization and Computer Animation3.3 (1992): 183-197.

well-made visualisations rather than objective recordings of reality. However, it should also be noted that the ability to re-use lower quality versions of digitised objects and environments in games or virtual environments may provide further incentive for their creation in the first place.

2.8.2 Crowd Sourcing

Due to its ability to both engage users, and create multi-user collaborative environments, 3D content also appears in several crowd sourcing projects in the cultural heritage area.^{190,192}.

2.8.3 Kiosk applications

Kiosk applications, or applications running on hardware within museums and galleries is a common method for using 3D content. Indeed, before disseminating 3D content over the web became feasible, apart from distributing physical media such as compact discs (as in the virtual Stonehenge project), kiosk-type installations were the only way for CH institutions to use 3D content in public facing applications.¹⁹³

Today, many kiosk applications can also be accessed online¹⁹⁴, though the kiosk still has advantages. For example, hardware configuration such as screen size and touch screen

¹⁹⁰ Chih-Hao Yu, Tudor Groza, Jane Hunter, (2013), Reasoning on Crowd-Sourced Semantic Annotations to Facilitate Cataloguing of 3D Artefacts in the Cultural Heritage Domain, The Semantic Web – ISWC 2013, Volume 8219 of the series Lecture Notes in Computer Science pp 228-243

¹⁹² Daniel Biella, Daniel Sacher, Benjamin Weyers, Wolfram Luther , Nelson Baloian, Tobias Schreck, (2015) Crowdsourcing and Knowledge Co-creation in Virtual Museums, Collaboration and Technology, Volume 9334 of the series Lecture Notes in Computer Science pp 1-18

¹⁹³ M. Carrozzino, C. Evangelista, A. Scucces, F. Tecchia, G. Tennirelli, and M. Bergamasco. 2008. The virtual museum of sculpture. In Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts (DIMEA '08). ACM, New York, NY, USA, 100-106.

¹⁹⁴ Rafal Wojciechowski, Krzysztof Walczak, Martin White, and Wojciech Cellary. 2004. Building Virtual and Augmented Reality museum exhibitions. In Proceedings of the ninth international conference on 3D Web technology (Web3D '04). ACM, New York, NY, USA, 135-144.

interfaces are known, and extra hardware allowing, for example, haptic interaction¹⁹⁵ can be provided. Higher quality models can be provided than would be possible over the web (see, for example, chapter 6). There is still, however, a lack of published research into user experience with such installations.

2.9 The future: measuring BRDFs

All of the digitising techniques discussed above measure the reflectance of a surface for one viewing angle and one angle for the light source. Thus, as we have seen they all have difficulty measuring consistent values for non Lambertian surfaces, and for representing the true appearance of an object and its interaction with light.

The solution to this problem is to measure the reflectance of a surface for many combinations of light source and viewing angle. What is measured in this case is a BRDF, or Bidirectional Reflectance Distribution Function. The BRDF is a four dimensional function that for any given incident angle of light and viewing angle gives the ratio of reflected to incident radiance.¹⁹⁶

There are in fact, many types of BRDF, increasing in complexity and dimensionality depending on which and how many parameters are measured¹⁹⁷. Non-spatially varying BRDFs are generally used to measure homogenous samples of a particular material, giving a set of measurements that can then be used to render that material. For

¹⁹⁵ Butler, M., & Neave, P. (2008). Object appreciation through haptic interaction. Hello! Where are you in the landscape of educational technology? Proceedings asilite Melbourne 2008, 133-141.

¹⁹⁶ Nicodemus, Fred (1965). ["Directional reflectance and emissivity of an opaque surface"](#) Applied Optics Vol. 4, Iss. 7, pp767-775

¹⁹⁷ They are, in fact, all subsets of the 16 dimensional GRF, or General Reflectance Function, Haindl, M., & Filip, J. (2013). Visual Texture: Accurate Material Appearance Measurement, Representation and Modelling. Springer Science & Business Media., p9

example, the MERL BRDF database includes BRDFs for 100 materials¹⁹⁸ (Figure 2.13).

Spatially Varying Bidirectional Reflectance Distribution Functions (SVBRDFs) measure BRDFs over a two dimensional surface, whilst the Bidirectional Texture Function (BTF) and Bidirectional Surface Scattering Reflectance Distribution Function (BSSRDF) also measure non-localised effects such as subsurface scattering¹⁹⁹. To measure iridescence and/or fluorescence, a new dimension must be added – all the measurements taken must be repeated for different wavelengths of light.




Figure 2.13: Material samples used to measure BRDFs for the MERL database. Image: <http://www.merl.com/brdf/>

¹⁹⁸ See: <http://www.merl.com/brdf/> and Matusik, W. (2003). A data-driven reflectance model (Doctoral dissertation, University of North Carolina, Chapel Hill).

¹⁹⁹ Filip (2013) pp12-19

There are many methods for capturing BRDFs, though all operate on the same basic principle, moving one or more of the object, camera and light source in order to measure multiple values for reflectance.

Gonioreflectometers measure four dimensional BRDFs and feature a fixed sample with a moving sensor and light source. The Kaleidoscope method uses an arrangement of mirrors and beam splitters to capture multiple reflections in a single image²⁰⁰. With the advent of (relatively) cheap digital photography, the dome or camera array, where multiple cameras and/or light sources are used to is becoming more prevalent, particularly in cultural heritage recording²⁰¹

Two good examples of BRDF recording specifically for cultural heritage have emerged in the last few years, CultLab3D²⁰² in Darmstadt and the Dome II project at the University of Bonn²⁰³. CultLab3D involves an automated ‘conveyor belt’ method for capturing the geometry and optical properties (BRDF), aimed at mass digitisation projects, whilst Bonn’s method features a modular dome that can be dismantled, moved and reassembled where required. Both examples use projected light to measure 3D geometry, CultLab3D at another ‘station’ on the conveyor belt, whilst Bonn includes projectors inside the dome.

²⁰⁰ Schwartz, C., Sarlette, R., Weinmann, M., & Klein, R. (2013, June). DOME II: a parallelized BTF acquisition system. In Proceedings of the Eurographics 2013 Workshop on Material Appearance Modeling: Issues and Acquisition (pp. 25-31). Eurographics Association.

²⁰¹ Methods can be hybrids of the above, for example in the CultLab3D project mentioned below, multiple cameras and light sources are places on a moving arch, giving the same effect as a dome

²⁰² Digitizing Cultural Heritage. About the Cover CultLab3D, IEEE Computer Graphics and Applications (CultLab3D have yet to publish any results from their process for commercial reasons)

²⁰³ Schwartz et al (2013)




Figure 2.14: Two methods for recording BRDFs. Top, CultLab3D's conveyor belt system with a moving arch of light sources and cameras. Bottom, the University of Bonn's 'Dome II' Images: CultLab3D and Schwarz et al (2013)

Other examples include the 3DCoform project's mini-dome, which has an array of fixed cameras and light sources and which uses photogrammetry to capture geometry whilst simultaneously measuring a BTF. Another example of a 'low cost' solution is a dome featuring 50 compact digital cameras²⁰⁴. By controlling each shutter and flash independently, a series of 50 images can be captured from each camera, illuminated by each of the other cameras' flashes. This process is repeated for every camera, generating a stack of 2500 images, from which both geometry and BTF can be reconstructed.

It should be noted that rendering an object with a BTF is considerably more complex than a simple diffuse or artificial diffuse + specular model such as Phong shading. Since much more information must be recorded per pixel, the amount of data, and the amount of processing, is necessarily far greater. However, outputs of models rendered with realistic optical properties can already be viewed on the web, for example using Bonn's impressive WebGL BTF-Object viewer²⁰⁵. CultLab3D also has models to view online²⁰⁶, though as all the available objects appear to be made of single, mostly diffuse, materials, and are (virtually) lit with similarly diffuse lighting, it is difficult to ascertain how successful their technique is.

²⁰⁴ Pers communications, unfortunately further details are protected by an NDA

²⁰⁵ <http://btf.cs.uni-bonn.de/viewer/buddha.html>

²⁰⁶ <http://www.cultlab3d.de/results.html>




Figure 2.15: Screenshot of Bonn's object viewer. Whilst there are more traditional CH objects available to view, this shoe beautifully illustrates the contrast between the glossy metallic and dull matte fabrics

With the speed at which technologies such as photogrammetry have moved from highly technical specialist-only techniques to something that can be carried out by anyone using just a mobile phone, it is easy to envision a time when recording and rendering cultural heritage objects with their full optical properties becomes equally

cheap and routine. Whether these particular approaches eventually become common place in cultural heritage (for example, 3DCultLab's mass-digitisation approach may prove problematic for objects with complex geometry, and for particularly fragile or valuable objects that can't be placed on a conveyor belt), some form of more realistic recording and rendering of objects would seem to be inevitable in the coming decades, and would certainly revolutionise the use of 3D models in public facing applications.

2.10 Conclusions

Many 3D digitising technologies are mature techniques with a long history of use in cultural heritage. Within the last five years, however, a suite of new technologies (including certain photogrammetric methods, WebGL and digital photography) have reached a point where 3D digitisation has become a low cost operation within the reach of most cultural heritage institutions. For the reasons discussed above, for example cost, ease of use and desired outcome, photogrammetry would certainly seem to be the method of choice for *most* public facing CH digitisation projects, though there will inevitably be occasions when another technology will be the better choice.

However, all the technologies discussed will struggle with certain objects due to their specific geometric and particularly material properties, and in fact the class of CH artefacts that can be imaged with something approaching 100% success is vanishingly small. Nevertheless, there are very good models from various sources already available online that demonstrate the potential of 3D digitisation in public facing applications.

Technologies do exist that can capture an object's optical properties in the form of a BRDF or BTF, but these are as yet unavailable to non-specialists. However, with the

speed at which other 3D capture techniques have become accessible, it is certainly possible that simple BTF recording may become viable within the next decade, and that the enhanced ability they offer in capturing and rendering authentic and accurate 3D models of cultural heritage objects will have profound effects on the utility of public facing 3D content.

3 Digital Surrogacy, Aura and the role of 3D models

3.1 Introduction

This chapter will examine the concept of the 3D virtual model as a digital surrogate. It will look at the key properties required by a 3D model and specifically how these properties differ between professional and public facing applications in cultural heritage. It will discuss some key concepts in museology and how 3D models might help fulfil a cultural heritage institution's remit, and finally it will look at Benjamin's concept of aura, and its relevance to 3D digitisation.

3.2 The Digital Surrogate

'Digital surrogate' is a common term used when talking about 3D models in cultural heritage (CH) applications^{207,208,209,210} and it is worth spending some time picking apart the meaning. Given that the standard dictionary definition of 'surrogate' is a person or thing acting as a substitute"²¹¹, and the definition of 'substitute' is 'a person or thing that serves in place of another'²¹², we can say that a model is a digital surrogate if it can act as a substitute or 'stand in' for the real object²¹³ for any particular purpose:

²⁰⁷ Mudge, M., Ashley, M., & Schroer, C. (2007). A digital future for cultural heritage. Available in: http://culturalheritageimaging.org/What_We_Do/Publications/cipa2007/CIPA_2007.pdf.

²⁰⁸ Cameron, F. (2007). Beyond the cult of the replicant: Museums and historical digital objects—traditional concerns, new discourses. Theorizing digital cultural heritage: A critical discourse, 49-75.

²⁰⁹ Hess, M., Robson, S., Millar, F. S., Were, G., Hvding, E., & Berg, A. C. (2009, September). Niabara—the western solomon islands war canoe at the british museum—3D documentation, virtual reconstruction and digital repatriation. InVirtual Systems and Multimedia, 2009. VSMM'09. 15th International Conference on (pp. 41-46). IEEE.

²¹⁰ David Arnold & Guntram Geser (2008), EPOCH Research Agenda for the Applications of ICT to Cultural Heritage Full Report, Archaeolingua

²¹¹ <http://www.collinsdictionary.com/dictionary/english/surrogate>

²¹² <http://www.collinsdictionary.com/dictionary/english/substitute>

²¹³ As in the previous chapter, for simplicity's sake I will use the term 'object' though we could be referring to an entire exhibit, gallery, building, archaeological site etc.

Interacting with the model will provide the same results as interacting with the object. For example, a measurement performed between two points on the digital model will return the same answer as a measurement made between the two corresponding points on the real object. But when a 3D model is described as being a digital surrogate for an object, there is clearly a missing clause; a digital model cannot serve as a substitute for a physical object for *all* purposes. Trivially, you cannot pick up a digital surrogate, weigh it in your hands or ascertain its material properties through the sense of touch, nor can you extract physical samples from it. There is clearly more to the concept of digital surrogate than it being simply a substitute or stand in. Different authors in the field of cultural heritage and digital humanities have attempted to define digital surrogate more rigorously:

“[the digital surrogates’] goal is to reliably represent real world content in a digital form. Their purpose is to enable scientific study and personal enjoyment without the need for direct physical experience of the object or place. Their essential scientific nature distinguishes them from speculative digital representations.” (Mudge 2007)

The ‘essential scientific nature’ of the digital surrogate refers in part to the method of acquisition; these are digitised models created via some repeatable methodology – there is a traceable connection between each point of data in the model and a corresponding point on the subject – and the distinction is drawn between these and ‘speculative digital representations’ – what we might characterise as ‘artist’s impressions’. Note that Mudge’s description is a functional one, it makes no claims as

to the model's quality or any of its properties, but talks about 'goals' and 'purposes'.

This functional definition is a useful one, and Arnold (2008) makes it more explicit:

"The capture of digital representations of cultural artefacts and environments is almost always related to a specific use. It is therefore appropriate to try to distinguish categories of use that give rise to different requirements and obligations. Three might be distinguished as digital surrogate, visualization, and representations captured for illustration or entertainment. The digital surrogate is the closest fidelity to the actual object that can be achieved digitally and theoretical representations for other purposes might be extracted from the surrogate. However, in practice it is unlikely that such levels of detail will be justified or achievable over all cultural artefacts and other categories may well be sufficient for identification or to get an impression in a Web page."²¹⁴

Here, he first uses a functional distinction, distinguishing between digital surrogates and lower quality, less accurate 'visualizations', or 'representations captured for illustration or entertainment'. This can be interpreted as a distinction between 'professional' and 'public facing' applications. He goes on to define the digital surrogate as the 'closest fidelity to the actual object that can be achieved digitally' and from which other models – presumably the visualisations and other representations – can be extracted. The digital surrogate, the best possible digitisation of the object, serves as a master record.

²¹⁴ Arnold & Geser, 2008, p63

This definition begs several questions though. As technology evolves, the ‘closest fidelity’ that can be achieved is constantly changing. Does this entail that a model that is a digital surrogate today is tomorrow’s ‘representation’? Different technologies may be chosen for different purposes as well; for example a model used for assessing damage or monitoring cracks in an object over time may need to measure accurate geometry but no surface texture while other types of research may need perfect colour recording but are more ambivalent towards geometric accuracy. Arnold seems to suggest that models for these two purposes may be extracted from the one ‘digital surrogate’, but at the capture stage one must choose what technology to use – the one that records geometry to the highest level of accuracy, the one that captures the best colour, or a compromise solution where colour and geometry are both captured with an acceptable degree of accuracy? The alternative would be to create two (or more) digitisations, each capturing a different aspect of reality. In this case we would not only have to abandon the notion of a single ‘digital surrogate’, but would also, presumably, drastically increase the cost and difficulty of digitisation.²¹⁵

So, in creating a single model with the ‘highest fidelity’ to reality, we must choose which particular aspect of reality we are interested in. Thus a clearer definition for digital surrogate might be: the model is a digital surrogate if it can substitute for the object *for the purpose of x*. This definition has the advantage that the success or

²¹⁵ Whilst these are philosophical discussions, there are clear practical questions. Is it better (more efficient/cost effective) to scan once for all purposes – even if this means spending more time/money on the process than is needed for current purposes (ie, a public facing visualisation). This is a complex and important question, and ultimately beyond the scope of this research. Intuition says that this is something that must be decided on a case by case, object by object and institution by institution basis.

otherwise of 'x' can be measured or evaluated.

What if 'x' is 'professional curatorial purposes' or 'academic research'? Even limiting the interaction to non-contact inspection of an object's surface properties, it is still difficult to see how any 3D model, certainly those feasible given today's technology, could ever substitute for the real object in *all* situations. As we saw in the previous chapter, the very best laser scan available today will fail to measure many surface properties, and thus fail to accurately capture texture for a wide variety of cultural heritage objects:

"ones that exhibit complex reflectance properties such as anisotropy or iridescence, ones that exhibit significant self-shadowing or mutual illumination, ones that exhibit significant subsurface reflection, objects that are highly specular or translucent, and objects with intricate surface geometry."²¹⁶

The digital model will not be useful for multi-spectral imaging unless non-visible wavelengths of light are captured during scanning, and the best scanning resolutions will still fail to capture details that might be revealed by a camera with a telephoto lens²¹⁷ or a curator with a powerful magnifying glass. Even the best geometrical measurement of an object will still capture discrete points and thus to create a complete surface from which we could conceivably make any arbitrary measurement will necessarily involve some interpolation of the surface between measured points.

²¹⁶ Hawkins, T., Cohen, J., &Debevec, P. (2001, November). A photometric approach to digitizing cultural artefacts. In Proceedings of the 2001 conference on Virtual reality, archaeology, and cultural heritage (pp. 333-342). ACM

²¹⁷ For example, the Arius scanner at UCL has recently been upgraded to a 50 micron resolution, a camera with a good lens may have a resolving power in the region of 10-20 microns

And whilst some improvement can be expected, there are both physical and practical limits to the accuracy that can be achieved using current methods (of both laser scanning and photogrammetry), including diffraction (in camera optics and of the laser itself) and speckle effects^{218,219,220}.

So even the best digital surrogate is a substitute for the real object in a very limited domain: visual, non-contact inspection in a restricted spectrum, and with restricted resolution (that of the capturing technology). For a CH professional with access to the object, the digital surrogate is almost always going to be a poor substitute for the real object.

That isn't to say that in certain professional applications the digital model doesn't have some obvious advantages over the physical object. The E-Curator project^{221,222} and subsequent work by Mona Hess²²³ has shown the potential for 3D scanning technology for professional CH purposes and there are clear advantages to working in the digital realm: Interrogating the model does not require handling of potentially fragile objects or objects which may be in hard to access storage or on public display. The model can be examined remotely from (potentially) anywhere with an internet connection, it can

²¹⁸ Beraldin, J. A. (2004). Integration of laser scanning and close-range photogrammetry-the last decade and beyond. In International Society for Photogrammetry and Remote Sensing.

²¹⁹ Beraldin, J. A., Picard, M., El-Hakim, S., Godin, G., Borgeat, L., Blais, F & Bandiera, A. (2005). Virtual reconstruction of heritage sites: opportunities and challenges created by 3D technologies

²²⁰ Marani, R., Nitti, M., Cicirelli, G., D'Orazio, T., & Stella, E. (2015). Design of High-Resolution Optical Systems for Fast and Accurate Surface Reconstruction. In Sensing Technology: Current Status and Future Trends III(pp. 47-65). Springer International Publishing.

²²¹ Hess, M et al (2008) "3D Colour Scans for Object Assessment", EVA 2008 London Conference, July 22-24, 2008

²²² Hess, M et al (2009) "E-Curator: A 3D Web-based Archive for Conservators and Curators", Ariadne, Iss 60, July 2009

²²³ Hess, M, (2015) A metric test object informed by user requirements for better 3D recording of cultural heritage artefacts, thesis

be simultaneously accessed by an unlimited number of people, and objects situated in different collections on different continents can be compared side by side.

Digital surrogacy is therefore no longer tied to ambiguous criteria such as 'closest fidelity', or 'essential scientific nature', but is defined by the ability of the model to substitute for the object *for a specific purpose*²²⁴. The creation of a model that can fulfil the particular purpose required depends, then, on communication between the cultural heritage professional and the digitisation expert. The CH professional must be able to specify the properties they need recording, and the digitisation expert must be able to communicate the limits of the capturing technology.

3.2.1 The public facing digital surrogate

We have seen that the idea of digital surrogacy is potentially problematic in professional cultural heritage applications as the standards a digital surrogate must reach are necessarily high and, when compared to accessing the real object, the digital surrogate is only a viable substitute for the original in certain specific circumstances. So what is a digital surrogate when talking about public facing applications? Visitors – as in members of the public with no privileged access to the exhibits – to cultural heritage institutions will almost always have a very different experience of an object to that of a CH professional. Often the object will be behind glass, almost certainly placed at a

²²⁴ In the case where the object is not available, it is tempting to say that the digital model is automatically a surrogate, as any interrogation of the model will necessarily be better than the alternative, ie, nothing. However, this is a problematic stance to take, particularly in the professional sphere as wrong conclusions are potentially more damaging than no conclusions at all. The idea of a model acting as a surrogate iff similar interrogations provide similar results is still important, even if, without the object, there is no way of immediately verifying this. The confidence in the model-as-surrogate could instead come from previous comparisons between model and object, or alternatively, thought a traceable and repeatable capture methodology.

minimum distance from the visitor, is only viewable from certain angles and with uncontrollable lighting. It would appear that for practically any purpose the visitor might have in regards to the object, the digital surrogate has a lot less work to do to become a viable substitute²²⁵ than in the professional case.

In fact, one could argue that there may be certain circumstances where the digital surrogate may be not just an adequate substitute, but a superior one; it may provide a better experience for the visitor than their interaction with the real object (see, for example, users' responses to the digital model of the Courtauld Bag in 6.8.3.4). On top of this, the digital model provides the same ancillary benefits as in the professional case; it increases the object's accessibility by many orders of magnitude; the surrogate can be 'on display' permanently (ie, not constrained by the limited space in museums), can be viewed from anywhere, by any number of people, in any number of contexts.

Of course, we are falling into the same trap: we have talked about a digital surrogate for public facing applications without talking about the specific purpose of the surrogate, as if both 'cultural heritage objects' and 'the public' are well defined homogenous sets and we can easily generalise over both. Instead, both are exceptionally heterogeneous²²⁶, and to find a single (non-trivial) property shared by all cultural heritage objects is an impossible task. Rather, the set of CH objects is a class

²²⁵ It should be noted that, given these circumstances, there is nothing to prevent a 2D digitisation – a photograph – performing as a digital surrogate for the object. Interestingly, there is nothing in the previous definitions, Mudge or Arnold's, that does not apply as equally to digital photography as it does to 3D digitisation.

²²⁶ Neidhardt, F. (1993). The public as a communication system. *Public Understanding of Science*, 2(4), 339-350.

defined by a series of polythetic properties or family resemblances. Similarly, to find a single purpose for which people interact with objects in cultural heritage institutions is an equally difficult task. To achieve a functional definition of a public facing digital surrogate we must approach it from another direction, the CH institutions themselves. What is their role, or remit – or more probably, their *roles* or *remit*s – and how may digital surrogates promote or fulfil those purposes?

3.2.2 The museum's dual remit

In the introduction, we examined the two complementary roles of the museum: to collect, research and preserve for the future (the professional, cultural heritage role); and to display, disseminate, educate and entertain in the present (the public-facing role). There is no necessary connection between the two, and no reason to believe the two roles will share aims and requirements, though there are obviously good reasons based on economy and convenience why the two should be fulfilled in the same location²²⁷. These differing aims and requirements will be reflected in how 3D digitisation projects are evaluated.

As discussed, the professional case (collect, conserve, research) depends upon a certain objectivity in the model (a one-to-one correspondence between model and object), in the public arena, however, I shall argue that authenticity is the key quality.

²²⁷ In parallel, whilst there is no necessary connection between creating digital models for academic and public facing purposes, there are still clear economic benefits if a model acquired for one purpose can be repurposed for another.

3.2.3 Objectivity

It is possible to characterise one of the main aims of 3D digitisation in a cultural heritage context as the creation of an objective record of a particular object. This is simply an extension or evolution of the documentation process every museum is required to undertake²²⁸, whether through physical measurements of the object, textual descriptions, photography or some other medium. There are extremely detailed requirements for the 2D capture of cultural heritage objects^{229,230}, but while efforts to generate an equivalent set of guidelines for 3D imaging are ongoing^{231,232,233}, the variety of available technologies (ie, terrestrial vs close range scanning, photogrammetry etc.) and subject matter (archaeological sites, complete buildings, objects and collections etc.) have complicated this process.

To generalise, in 3D digitisation for cultural heritage, the intention is to create a dataset which records surface properties of an object (predominantly, geometry and diffuse colour) in as much detail and with as much accuracy as possible (ie, to within the error ranges of the technology being used) in order to conduct research on that object²³⁴. For a record to be considered objective, it would need to fulfil criteria such that if one

²²⁸ According to the ICOM code of ethics: <http://icom.museum/the-vision/code-of-ethics/#intro>

²²⁹ For example, those of the Federal Agencies Digitization Guidelines Initiative, see Technical Guidelines for Digitizing Cultural Heritage Materials: Creation of Raster Image Master Files at <http://www.digitizationguidelines.gov/guidelines/digitize-technical.html>

²³⁰ <http://www.jiscdigitalmedia.ac.uk/guide/photographic-guidelines#pg4>

²³¹ Bryan P et al (2009) Metric Survey Specifications for Cultural heritage (English Heritage) 2nd ed, Swindon, UK

²³² CARARE 2013, Europeana CARARE project – bringing content for archaeology and historic building to Europeana users (2010-2013)

²³³ JISC Digital Media (2014). Digital 3D content infokit, <http://www.jiscdigitalmedia.ac.uk/infokit/3d/3d-home>

²³⁴ Fontana, R., Gambino, M. C., Greco, M., Pampaloni, E., Pezzati, L., & Scopigno, R. (2003, October). High-resolution 3D digital models of artworks. InOptical Metrology (pp. 34-43). International Society for Optics and Photonics.

were to carry out measurements on the dataset and repeat those measurements on the original object, one would obtain the same result; for example, measuring the distance between two points in the dataset would produce the same results as physically measuring the distance between the equivalent points on the real object. These virtual measurements could then be used for a number of purposes, such as measuring the propagation of cracks over time²³⁵, or for more abstract purposes than simply determining distances on a surface. Two projects have used detailed 3D models plus historical data to map astronomical events such as sunrises at Stonehenge²³⁶ or star positions as viewed from Roman temples in antiquity²³⁷.

The potential uses of such an objective dataset have been iterated many times, for example:

“3D imaging of works of art offers a significant new analytical tool to curators, historians and conservators, which provides some new and unique types of information which otherwise is not obtainable using traditional techniques. The high-resolution 3D image data contain a wealth of information that can be used for modelling, display, comparison, measurement and analysis applications.”²³⁸

The advantage of a 3D dataset or 3D model over a series of physical measurements is

²³⁵ Armesto, J., Arias, P., Ordóñez, C., Lorenzo, H., & Caparrini, N. (2007, September). Damage quantification and monitoring in masonry monuments through digital photogrammetry. In Key Engineering Materials (Vol. 347, pp. 291-296).

²³⁶ Burton, N. R., Hitchen, M. E., & Bryan, P. G. (1999). Virtual Stonehenge: a Fall from Disgrace?.BAR INTERNATIONAL SERIES, 750, 265-265.

²³⁷ Frischer, B., & Fillwalk, J. (2012, September). The Digital Hadrian's Villa Project: Using virtual worlds to control suspected solar alignments. In 2012 18th International Conference on Virtual Systems and Multimedia.

²³⁸ Taylor J, Beraldin JA, Godin G, Cournoyer L, Rioux M, and Domey J (2002) “3D imaging technology for museums and heritage” Proceedings of The First International Workshop on 3D Virtual Heritage, October 2-3, 2002, Geneva, Switzerland, pp. 70-75.

that the dataset can be analysed at any time without reference to the original object (which may be fragile, unavailable or even lost or destroyed), and measurements can be made between any arbitrary points on the object – if we want to compare an object's changes over time we are not restricted to comparing the particular physical measurements made at the time the object was documented²³⁹. For example, a project at UCL involved creating highly accurate scans of masonry blocks using the Arius Foundation scanner. The blocks were then stored in various conditions (in a field, indoors, etc.) and rescanned at regular intervals. By comparing the point clouds captured over time, the exact changes to the blocks could be noted and recorded.

Similarly, if we want to compare measurements conducted on two related objects, we are not relying on the same measurement having been made on both. In terms of recording surface properties, a single photo can only show one aspect of the object and a finite series of photos can only show a finite number of views, while a 3D dataset – despite itself being constructed from a finite number of views – has the potential to show the object from an effectively infinite number of aspects.

As well as being able to perform repeatable and verifiable measurements, another key aspect for our dataset to be truly objective is that the data should be as 'raw as possible' and mediated as little as possible through human agency. Where subjective decisions have to be made, the reasoning behind and ramifications of the decisions should be recorded to allow future researchers and users of the dataset to make

²³⁹ Uccheddu, F., Pelagotti, A., Cappellini, V., & Massa, E. (2014, July). 3D Technologies for Measurement of Painting Surface Deformations: A Case Study. In Proceedings of the EVA London 2014 on Electronic Visualisation and the Arts (pp. 56-62). BCS.

informed decisions.²⁴⁰ One must recognise and acknowledge the limitations of the device(s) used to capture the data; for example measurements can only be given within the error range of the capturing device, and even when properly calibrated, there may be minor differences in colour captured in individual scans due to changing ambient lighting conditions or differing amounts of surface reflectivity (see previous chapter).

The Arius3D Foundation Model 150 scanner used in the E-Curator project captures data with “a level of geometric and colour standardisation that easily surpasses any other available recording process”²⁴¹ and yet even with this level of fidelity, subjectivity enters the process: “when editing the colours, 'artistic' decisions will be made by the technician in charge.”²⁴² Indeed, these subjective decisions were raised as issues by the museum professionals participating in the E-Curator project, and reinforce the need for the provision of paradata when creating 3D models for professional purposes.

As well as ambiguous data, many models will contain holes (missing data) where the geometry occludes certain areas from the sensor, where they are made of a material that the sensor cannot resolve, or areas have simply been missed in the scanning process. Interpolating existing data to fill these holes, or approximating the surface with new data is another process that concerned the users in the E-Curator report, and can lead to tension between the requirements of objectivity and authenticity²⁴³.

²⁴⁰ See for example the paradata entries in the London Charter for the computer-based visualisation of cultural heritage: <http://www.londoncharter.org/downloads.html>

²⁴¹ Brown et al (2008) E-Curator: 3D colour scans for remote object identification and assessment – Final Report

²⁴² Brown (2008) *ibid*

²⁴³ Brown (2008) *ibid*

3.2.4 Authenticity

Algharabat and Dennis²⁴⁴ define authenticity in a computer mediated environment as: "a psychological state in which virtual objects presented in 3D in a computer-mediated environment are perceived as actual objects in a sensory way."²⁴⁵ This definition is particularly useful as it can be applied to NIVRs, or 'non-immersive virtual realities' as opposed to other definitions which were devised for 'traditional' immersive VRs which attempt to embody the user in a virtual environment and involve concepts like telepresence and feelings of transportation; concepts not relevant when viewing virtual museum objects in isolation.

Authenticity in this form is a subjective, psychological concept and therefore a property of the viewer as much as the object itself, and authenticity in this sense has been shown to increase both engagement and learning potential^{246,247}, and thus would appear to be an important concept regarding museums' public facing remit.

How does a 3D model acquire authenticity? It is perhaps easier to provide a list of features that negatively affect it, for example impossible geometry, edges that don't

²⁴⁴ Algharabat R & Dennis C (2010) "3D Product Authenticity Model for Online Retail: An Invariance Analysis", Int. Journal of Business Science and Applied Management, Vol 5, Iss 3

²⁴⁵ I take it from this definition that the authors do not equate this psychological state with being 'fooled' by a virtual object, ie, believing they are looking at a solid object rather than an image on a computer monitor. Rather, the viewer can reconstruct a consistent internal representation of a 3 dimensional object from the visual inputs from a virtual object, in the same way they would construct an internal model from the visual inputs from a 'real' object. Any inferences they then draw as to the 3D properties of the object will be the same whether their visual inputs came from the real or virtual object. One could restate this as the virtual object is acting as a digital surrogate for the real object with respect to its 3D properties (shape etc)

²⁴⁶ Algharabat, R. , Dennis, C. , Morschett, D. , Rudolph, Th. , Schnedlitz, P. , Schramm-Klein, H. , Swoboda, B.(2010c). 'Modelling the impact of 3D authenticity and 3D telepresence on behavioural intention for an online retailer'. European Retail Research. 24, 2, Gabler Verlag, 93-109

²⁴⁷ Kakuta T, Oishi T, Ikeuchi K, (2008) "Development and Evaluation of Asuka-Kyo MR Contents with Fast Shading and Shadowing", Proc. Int. Society on Virtual Systems and MultiMedia (VSMM 2008), Oct. 2008, pp.254-260.

meet up, holes that reveal the lack of solidity of the model and discontinuities in colour where we expect the smooth gradients that characterise the world around us. Table 2²⁴⁸ lists the various cues which *do* allow us to create a 3D scene from our visual inputs. Some of these cues, specifically those involving physical changes in the eyes (the first three items in Table 1) are simply unavailable to our virtual model²⁴⁹, or not relevant to viewing objects individually, such as atmospheric effects and relative size.

Cue	Explanation	Availability in desktop 3D
accommodation	The adjustment to the lens required to bring the object into focus	no
vergence	The convergence or divergence of the eyes required to produce an apparently single image	no
binocular disparity	The difference between the image as viewed by the two eyes	no
occlusion	The hiding of parts of an object by other objects	yes
relative size	The proportion of the view taken up by an object	yes
relative density	How close together objects appear	yes
height in the visual field	The up-down position within the visual field	yes
aerial perspective	The degree of atmospheric colour distortion (normally making objects appear more blue)	yes
perspective	The convergence of parallel lines going away from the viewer	yes
Shading	The differences in apparent colour of surfaces depending on their angle from the light source	yes
texture gradients	The density of object textures (objects further away will have more dense textures)	yes
motion parallax	The change in occlusion of objects as the view position changes (especially moving left-right)	yes
motion perspective	Changes in object size and density as the view position changes (especially moving nearer-further)	yes

Table 2: Visual Depth Cues - from Hedberg et al (2002)

²⁴⁸ Hedberg J, Harper B, Dalgarno B. (2002) "The contribution of 3D environments to conceptual understanding" O. J. McKerrow (Eds.), Winds of Change in the Sea of Learning: Proceedings of the 19th Annual Conference of the Australasian Society for Computers in Learning in Tertiary Education Vol 1 (pp. 149-158)

²⁴⁹ assuming we are not displaying it in stereoscopic 3D, though even with this 'pseudo-3D' display, there are conflicts between vergence and accommodation [Reichelt 2010]

Of those we *can* simulate in our 3D model, clearly properly rendered geometry will play a large part, taking care of occlusion and perspective, and, when combined with a virtual camera, motion-related changes such as parallax and perspective. However, the 'shading' entry is an interesting one, as, while we can simulate some changes in apparent colour due to the relationship between light source and object (ie, using an approximation to reality such as the Phong Model²⁵⁰), current technology only measures an object's diffuse colour – modelling, for example, the specularity of the object involves manually adjusting parameters according to a subjective judgement (see for example, the lustreware bowl in section 5.3.3). Our rendered model will only be a completely accurate and objective representation of the interaction between light and surface for a perfectly lambertian material. As we have seen, (section 2.3.1), there are many other components that make up the complete lighting or reflectance model for an object, properties we simply cannot easily extract from the data provided by the common 3D digitisation technologies (for example laser scanning, structured light and photogrammetry) currently used in capturing cultural heritage objects.

There are ways of simulating the look of more complex, non-Lambertian materials in the rendering process, for example Arius' Pointstream.js html5 renderer (as used in the Illuminating Objects project, chapter 5) allows you to adjust the global shininess and specularity of the model, or to define different parts of the point cloud and adjust their

²⁵⁰ Phong BT, (1975) Illumination for computer generated pictures, Communications of the ACM, Vol 18 Iss 6, June 1975, pp 311-317


values individually. Another method, used by the V&A's digitisation project²⁵¹ in cases where a higher quality model is required, is to use a renderer (in this case via the open source JavaScript library X3Dom) which can assign different material properties to different areas of a meshed model. Different material properties can then be simulated via shaders, either custom written or drawn from existing databases of materials. However, these are somewhat subjective processes (similar to the 'artistic' decisions taken in the ECurator project) and there is no way to infer material properties from the underlying, objective dataset.

These common rendering techniques and lighting models only simulate relatively simple surface properties such as specularity and shininess, not the more complex physical properties such as subsurface reflection or iridescence. It is difficult to demonstrate the gap between a 'diffuse' model and a realistic one without seeing the 3d model 'in action' - the effects talked about are due to the changing interactions between the light and model, and thus difficult to illustrate in a static image. Figure 3.1 shows a CG image of a glass of milk rendered with and without sub-surface scattering²⁵², a material property where, whilst the majority of light hitting an object is reflected, some penetrates the surface, is reflected internally and is finally transmitted from a point some distance from where the light ray originally hit the object. Whilst the glass of liquid is an extreme example, many materials commonly found in cultural

²⁵¹ Pers. comms with the digitising team from the V&A

²⁵² The image is from Jensen, H. W., Marschner, S. R., Levoy, M., & Hanrahan, P. (2001, August). A practical model for subsurface light transport. In Proceedings of the 28th annual conference on Computer graphics and interactive techniques (pp. 511-518). ACM. Note, this is simply a rendering technique, the scattering shown is simulated and not a measured property of the object.

heritage artefacts, including, but not limited to, marble, jade, shell, skin and other organic materials²⁵³, are subject to some degree of sub-surface scattering.



*Figure 3.1: CG rendering of a glass of milk, with (right) and without (left) sub-surface scattering.
Image from Jensen et al (2001)*

It is fairly clear from Figure 3.1 which of the models is a more ‘authentic’ model of a real object. As discussed in (Chapter 2), technology exists to capture some of the more difficult material properties, but is not necessarily accessible to small institutions in the same way that cheap laser scanning or photogrammetry is. And with traditional techniques, even where the exact optical properties of an object are not required, the task of simply recording, processing and rendering accurate geometry and diffuse

²⁵³ 96% of the light reflected by skin, for example, is subject to sub-surface scattering, with only 4% reflected directly from the surface: Krishnaswamy, A. and Baranoski, G. V.G. (2004), A Biophysically-Based Spectral Model of Light Interaction with Human Skin. *Computer Graphics Forum*, 23: 331–340. doi: 10.1111/j.1467-8659.2004.00764.x

colour for an object with complex optical properties is considerably more complex than for something well-behaved and diffuse²⁵⁴.

It would seem from the previous arguments that creating a public facing digital surrogate is an easier task than creating one for professional purposes, but this is not necessarily the case. While the creation of a professional surrogate requires certain levels of accuracy and a rigorous workflow, the purpose may be a very narrow one – we may only need to record a certain area of the object, we may only require geometry or colour, not both etc. The public facing model may require less accuracy (it has bigger ‘error bars’) but it must present a holistically authentic object with an illusion of solidity and reality not required in the professional sphere.

For many cultural heritage objects, it is its aesthetic properties that give the object much of its value, properties that are not wholly captured in a set of objective measurements of shape or diffuse colour. We live our lives surrounded by 3D objects with complex surface properties, and we find it easy, in the vast majority of cases and apart from rare ‘optical illusions’, to successfully parse the interaction between light and surface. If this interaction is missing or ‘off’ in a 3D model, we will either reject the object as an authentic representation of reality, or misinterpret the information from our senses – for example assuming a shiny object is made of some matte material. In both cases, the 3D model will fail as a digital surrogate, even when that surrogacy is restricted to purely visual inspection.

²⁵⁴ See chapters 5 and 6 on some of the challenges of digitising shiny and specular objects

From these arguments, it can be seen that the domain of objects from which we can create digital surrogates for public facing applications is necessarily restricted. In fact, counter-intuitively, it may be more restricted than the domain of objects we could digitise for professional purposes, which may not require the whole object to be digitised, or all of its properties to be recorded.²⁵⁵

3.2.5 Interaction

According to Algharabat²⁵⁶, there is a strong link between perceived authenticity and a user's interactions with, and specifically their apparent control over, the virtual model. Smooth animation with an acceptable frame-rate, as well as an intuitive and responsive method of manipulating the virtual object can go a long way to creating a strong feeling of authenticity. Again, there is nothing in the capture process, or any properties inherent in the dataset that entails a particular interaction model. There may in fact be a negative correlation between the size and density (and therefore detail/quality) of the dataset and the quality of the interaction model. However, this means care must be taken when providing a system for interacting with a 3D model. As my research shows (section 4.9), while there isn't yet a single standardised system for manipulating 3D models via mouse and keyboard, interaction is one of the most desired features when presented with a virtual object.

²⁵⁵ It should be noted however, that many of the problems that prevent accurate texture information being recorded (subsurface reflection, specularity etc.) also make it difficult to accurately record geometry

²⁵⁶ Algharabat (2010)

By interaction we are primarily talking about manipulating the 3D model in space (or, more accurately, manipulating the camera or user's viewpoint relative to the object). There are other forms of interaction, of course. For example, we can allow the user to interact by labelling or highlighting particular areas on the object, or by having certain parts of the object ('hotspots') trigger actions such as playing media when clicked on with the mouse (or via touchscreen etc.). Appendix D shows a simple method for making 3D point clouds interactive via WebGL.

3.2.6 Sustainability and intellectual property

We have talked about the museum's dual role which can be crudely defined as preservation and dissemination. Whilst we have talked predominantly about the public facing, dissemination role, the role of sustainability in any digitisation programme is an important one; whilst a fragile physical object can be preserved in a digital form (insofar as the limits of the digital surrogate allow, as discussed earlier in this chapter) there is a responsibility on behalf of the digitising institution to ensure that, like the original, the digital model is both preserved and accessible in the future. Similarly, if a resource is provided online, care must be taken to ensure that the resource is available to as many people for as long as possible²⁵⁷.

Issues involved in sustainability include choice of 3D file format, and choice of software

²⁵⁷ There may be occasions where a particular online resource is created to coincide with a particular event, exhibition, anniversary etc., and that resource is intended to be available only during a specified timeframe. However, online resources do not suffer the same constraints (physical space, limited resources etc.) as 'real-world' objects, and thus the online component of a real world exhibition (etc.) can live on long after the primary reason for its creation is gone. Even if the original 3D resource is no longer required, for example, a kiosk app in an exhibition that is now closed, given the effort taken in creating 3D assets, it would also make sense to preserve them for reuse or repurposing in the future.

used to render the model. Proprietary file formats can offer advantages such as bespoke compression techniques which can greatly reduce the file size (particularly important for online applications), however relying on continual support. Choosing a correct file format for preservation is a common question in the digital humanities^{258, 259, 260} though the advice is far more comprehensive when dealing with 2D digitisation than 3D file formats. However, whilst .obj²⁶¹ is a common and open format, most software will allow a 3D file to be exported in an ascii format, which whilst inefficient and can lead to extremely large file sizes (and thus unsuitable for use in many applications, particularly online), is extremely interoperable, open and is suitable for archival purposes. Software used to render 3D content must also be carefully chosen. Proprietary, paid for solutions can offer more flexibility and rendering options, but as well as expense you are again reliant on the providing business. Creating your own viewer, or using an existing open source solution gives the user much greater control over rendering options but will require some investment in time and/or money both in the initial coding or customisation stage, and also in maintenance. Browser technology is constantly changing and online code may need to be updated to ensure continued functionality (see section 5.3.3).

Placing 3D models online, particularly using open source software and file formats

²⁵⁸ Arms, Caroline, and Carl Fleischhauer. "Digital formats: Factors for sustainability, functionality, and quality." Archiving Conference. Vol. 2005. No. 1. Society for Imaging Science and Technology, 2005.

²⁵⁹ <http://bentley.umich.edu/giving/donate-your-archives/guidelines-for-the-selection-of-sustainable-preservation-quality-file-formats/>

²⁶⁰ http://www.jiscdigitalmedia.ac.uk/infokit/file_formats/digital-file-formats

²⁶¹ <http://www.dpconline.org/advice/preservationhandbook/technical-solutions-and-tools/file-formats-and-standards>

does entail releasing the 3D models 'into the wild'. A 3D model served over the web can be downloaded and the model then be altered, copied or 3D printed by users. This has obvious ramifications for the owning institution; prints or physical models of the object may be a potential revenue stream and there may be issues over who owns the rights to the 3D model. Copyright issues around 3D digitisation are complex²⁶² and still, to a large extent, undefined. Ownership should be defined at creation, through a contractual agreement between the object owners and digitiser. Permissions can be defined via a license attached to the 3D model; Jisc recommends the use of creative commons licenses²⁶³.

These issues go beyond simply providing 3D models online. As is shown in this thesis, as the ability to create digitised models is democratised it will become increasingly feasible for a member of the public to create models simply by photographing an object on display in a museum (or even using photographs sourced via the web). Whilst the quality of these models may not be equal to those created in controlled circumstances, it is already possible for someone to create a model of sufficient quality for 3D printing. In this case, the rights issues become complex and, in a sense, the law has yet to catch up with the technology. An example of the legal and ethical considerations around this sort of modelling were illustrated by the recent 'covert' scan of the Nefertiti Bust in the Neues Museum in Berlin²⁶⁴. Whilst it is debatable whether

²⁶² <https://www.jisc.ac.uk/guides/3d-digitisation-and-intellectual-property-rights>

²⁶³ <https://creativecommons.org/licenses/>

²⁶⁴ <http://www.smithsonianmag.com/smart-news/thanks-sneaky-scanners-anyone-can-3d-print-copy-nefertiti-bust-180958213/?no-ist>

this 3D model can be considered a true digitisation or artist's impression/ visualisation²⁶⁵, the potential to create 3D models from museum objects without the permission or knowledge of the institution is already there and will only become more prevalent as the technology improves. Whilst these complicated questions are beyond the scope of this research, anyone involved in 3D digitisation projects within cultural heritage institutions must consider them as part of their planning process.

3.3 Museology and its history

Following Casey²⁶⁶, it is helpful to divide museum practise into three typologies: The legislative, interpretive and performing museum. These three types of museum do represent a chronological sequence, but there is much overlap and shouldn't be considered as exclusive categories; it could be argued that nearly all museums partake of all of the three types to varying degrees, and it is simply the relative proportions which have changed over time.

Casey describes the oldest type of museum, the legislative, as a pre-19th century conception. Legislative museums aim to be "paragons of the aesthetic and intellectual pursuit, to create a venue for display not debate", while the museum itself becomes merely "a container for collections of objects." In this incarnation, according to Casey, the museum gained its authority through its collections, unique and special objects intended to be viewed with a passive, awe-struck demeanour. This type of institution

²⁶⁵ <http://3dprintingindustry.com/2016/02/26/3d-scanned-nefertiti-real-fake-or-real-fake/>

²⁶⁶ Casey (2003) "The museum effect: gazing from object to performance in the contemporary cultural-history museum", Archives & museum informatics [1042-1467]

survives today in the form some traditional art galleries, where old masters hang in quiet, church-like spaces and appreciated in reverential silence.

The modern, interpreting museum is perhaps unsurprisingly closer to our conception of 'museum', and arose from the realisation that whilst we should, as Gurian says, "acknowledge the power of some objects to speak directly to the visitor, for example, in the sensual pleasure brought about by viewing unique original objects of spectacular beauty," – as in a legislative institution – "the notion that objects, *per se*, can communicate directly and meaningfully is under much scrutiny."²⁶⁷ This shift in thinking also recognises that meaning is not fixed and inherent in the object, it is shifting and contingent and can change over time as discoveries are made and society evolves. As Messham-Muir says, "Since the emergence of the New Museology in the 1980s, it is an axiom of museum interpretation that an artefact's meaning and significance is contingent upon its social contexts."²⁶⁸ Indeed, changes in an object's meaning aren't restricted to re-interpretations of the past; as well as preserving cultural heritage, museums also act as transmitters of the current prevailing culture by interpreting the past through its prism:

"Through label text, docent tours, and multimedia tools, the museum provides a framework for how objects should be viewed and understood ... Rather than having objects speak for themselves, museum professionals interpret cultural

²⁶⁷ Gurian, Elaine Heumann (1999) "What Is the Object of This Exercise? A Meandering Exploration of the Many Meanings of Objects in Museums", *Daedalus*, Vol. 128, No. 3, America's Museums (Summer, 1999), pp. 163-183

²⁶⁸ Messham-Muir 2006 *ibid*

significance for visitors by structuring art and artefacts around easily identifiable chronologies, geographies, formal themes, and narratives.”

The third type, the performing museum, is as much a reaction to changes in society and the commercial pressures faced by museums as it is a 'better way to do museology'. The museum has assimilated “commercial strategies to entertain audiences ... the contemporary museum privileges the processes of display over the particularity of objects to convey information.”²⁶⁹ The performing museum sees as much, or more, emphasis placed on how an object is displayed than on the object itself; the museum object is further removed from its legislative pedestal and becomes just another (albeit important) feature of the exhibit as a whole. We can see this 'performing' in practise through the use of multimedia installations, interactivity and 'event-driven' exhibits. Chakrabarty²⁷⁰ sees the shift in museum philosophy in the 20th century as a reflection of the evolution of western democracies away from a pedagogical model to a performative one. One could restate this as a societal change from authoritarian to participative; the change in museums reflecting the idea that we are no longer so receptive to information (or interpretations) handed-down to us from an unimpeachable source, instead we place more emphasis on constructivist learning. Messham-Muir sees in this evolution a move away from “cognitive forms of interpretation” to “affective forms” and the “general trend towards 'experiential' display practices within museums in general”²⁷¹, ideas have relevance to the use of

²⁶⁹ Casey 2003 *ibid*

²⁷⁰ Chakrabarty, Dipesh (2002), “Museums in Late Democracies”, *Humanities Research*. X: 1

²⁷¹ Messham-Muir 2006 *ibid*

virtual models.

As we have mentioned, the chronological move through the categories is not an inevitable or well defined flow and contemporary institutions will partake of the different aspects of museology in differing amounts. Even within the same building, different galleries or exhibitions may emphasise different aspects and sometimes these differences may be expressed in the way individual objects are displayed in the same exhibition within the same gallery. It is an evolution, and even the most performing museum will retain genes from the legislative days while the most legislative of institutions may have picked up some tricks from the new museums. This evolution can be seen in microcosm in one particular gallery in the Science Museum. The Shipping Gallery is – or was – the museum's oldest existing exhibition, having been opened in the 1960s and closed earlier this year. In 2014 the space was filled by a brand new exhibition, The Making of Modern Communication (MMC) and the difference in style of the two exhibitions illustrates the change from legislative/interpretative to interpretative/performing.

In the Shipping Gallery, objects were arranged in a rough chronological order, with a general flow of time from one end of the large hall to the other. The hall was divided into two main halves, one containing model ships and the other models of various parts of ships related to propulsion, steering and other details of ship building. The objects were densely packed with no single path through indicated and very little context for the visitor, the only information provided on small typed labels which are

also densely packed, full of facts, figures and esoteric technical jargon²⁷². Little or no attempt is made to explain or embed the details in a larger context, and there is no clear attempt to delineate a hierarchy of the models or focus attention. Each object is as important as any other and is to be appreciated in isolation, or only very loosely as part of a wider narrative. The history of advances in steering, engines and propeller design can be inferred from the exhibition, but it is up to the visitor to extract the signal from the noise; the viewer is responsible for creating their own narratives from the objects on display. We can contrast this with its replacement, MMC. From a Science Museum presentation:

“This gallery will immerse visitors in the experience of change since the introduction of the electric telegraph in the 1830s. It will tell parallel stories through the eyes of those that invented, operated, and were affected by each new wave of communications technology. Within each Network, visitors will be invited to explore stories about people and technology. Our audience research shows that *‘Visitors are seeking to understand the impact of objects on people’s lives at the time’*. They want an insight into the historical context in order to have an engaging experience with objects. To address this, each Network features transforming events that illustrate the significance of communications technologies to people’s lives.” [my emphasis]

Objects are barely mentioned, and only in the context of their impact on people’s lives.

²⁷² It is of course possible that visitors in the 60s were more knowledgeable about, and familiar with, nautical matters and the technical aspects of sailing, though I doubt to an extent that would make much difference to this argument

They are reduced to a secondary role, no longer to be appreciated for their own sake as in the legislative museum. Stories, events and experiences seem to be the primary components of this exhibition; the objects have become subservient to narrative. As Gurian says, “in the late 1960s and early 1970s, the definition of museums always contained reference to the object as the pivot around which we justified our other activities” [Gurian '99], whereas today, as MacDonald points out, museums primary role has become dissemination of information rather than merely collections of objects [MacDonald '92]. Today it could be argued that a museum's other activities (storytelling, education, entertainment) justify the objects: If exhibiting an unremarkable mobile phone is the best way of telling a story or imparting information, then that justifies the phone's exhibition rather than any particular characteristics of the object itself. We accept the mobile phone as a museum exhibit; instead of the museum gaining its authority from its objects, the objects gain authority from the museum.

This evolution would, on the surface, make an argument for the use of digital models more compelling. The object has changed from an end-in-itself to a means-to-an-end, and if the value of a museum object is its ability to impart information or sustain a narrative, and nothing to do with any intrinsic worth, then a digital model which could impart the information or support the narrative equally well could be substituted for the object with the exhibit suffering no detrimental effects. It would follow, then, that we could replace all the objects in our museum with digital copies – or indeed any other information containing object, such as a textbook – so long as they impart the

same information: the medium is irrelevant, only the message is important.

And yet museums *do* contain objects, not textbooks, and they continue to be visited.

The reason for the visit may be varied – it may be for recreation, entertainment, a social experience or learning²⁷³, but museums' enduring popularity is also something to do with their unique objects – as Falk and Dierking put it, “in some cases the museums themselves are considered unique or national treasures ... in other cases they present unique or special objects”. Building on research by Graburn and Yellis, Falk and Dierking ultimately put museums continuing popularity down to the public's need for “reverential experiences”. Perhaps we haven't moved as far from the legislative museum after all.

3.3.1 Aura, mechanical reproduction and the digital model

This idea of ‘reverential experiences’ is an important one. Not only is it somehow an integral part of what we might think of as the archetypal museum experience, it is also a clear distinction between cultural heritage’s professional and public worlds. Whilst a museum curator may feel reverence for the objects they handle, it is irrelevant to their professional role, and their interactions with objects or their digital surrogates. If a reverential experience is a fundamental, or even an important or common, part of a museum visitor’s interaction with an object, and the experience is part of what

²⁷³ Falk, JH & Dierking LD (2000), “Learning From Museums: Visitor Experiences and the Making of Meaning”, AltaMira Press

allows/encourages people to engage with museum objects, then what import does that have for a public facing digital surrogate?

While the requirements for digital surrogacy for both professional and public facing models are the same – that the model can substitute for the object for a particular purpose – it is the purposes which differ between the two applications. In both cases the model can be considered successful if the viewer's purpose is fulfilled; the professional acquires a measurement for the model which is as good as or better (within some error bars) as that which they would have got from the real object, the member of the public learns something about, is entertained by, or engages with the model – probably a mix of all three. However, if the ability for members of the public to engage with a museum object is, even in some part, related to the 'reverential experience', we must examine how, or if, the property or properties of the object that are responsible for the experience are captured or capturable in a digitised model. And if they aren't, can a public facing 3D model ever be considered a digital surrogate?

The property in question is very similar, if not identical to the idea of an object's 'aura'²⁷⁴. This concept, introduced by Benjamin in the 1930s represents an ineffable quality inherent in the object, and which is not captured by any form of mechanical reproduction²⁷⁵. Whatever you think of Benjamin's very much of-its-time Marxist

²⁷⁴ While some museum objects don't appear to have the property (for example, the cutaway model of a working toilet in the Science Museum, complete with plastic 'contents'), it could be argued that *all* objects displayed in cultural heritage institutions have a certain amount of aura inherited from the institution itself by virtue of their status as 'an item worthy of display', but there are clear differences between, say, the Science Museum's toilet and its Apollo space capsule. The toilet model is an object *within* a CH institution which no doubt tells us something *about* our cultural heritage (technology and innovation being part of that heritage), but is not, perhaps, itself a CH object.

²⁷⁵ For Benjamin, mechanical reproduction has a history as long as art itself; from copies of ancient Greek

arguments, his concept of aura as applicable to CH objects is certainly useful, and would seem to have an intuitive validity; with many museum objects or gallery artworks it is easy to feel something over and above the mere physicality of the object itself. We often feel some connection to the past through the object, visitors use words like ‘transported’ to describe the effect and cultural heritage objects seem to have a particular power to evoke emotional responses in viewers.

This power of ‘aura’ to elicit an emotional response can be described *affectual power* of an object. Witcomb, building on the work of Ross Gibson, describes an object's power to affect alteration, this being an emotional response to an object that through imagination and empathy allows us to “experience what it is to be other” and thereby come to a greater understanding²⁷⁶. Many other authors^{277,278,279,280} refer to a similar power, possessed by museum objects, that allows them to evoke an emotional response. This power is often framed in terms of the potential of an object to engender feelings of empathy in the viewer.

Some objects may have aura and affectual power independent of any knowledge in the viewer, due to its materials, craftsmanship or aesthetic qualities, though these still require some pre-existing knowledge in the viewer: the value of gold and other

statuary through woodcuts and engravings and up to the new technology of the day, photography and film. One can only imagine what he would have thought of 3D digitisation.

²⁷⁶ Witcomb A (2007), “The Materiality of Virtual Technologies”, *Theorizing Digital Cultural Heritage* Ed. Cameron F & Kenderdine S, MIT Press, Cam, MA

²⁷⁷ Messham-Muir, Kit. (2006) “Affect, Interpretation and Technology.” *Open Museum Journal* 7

²⁷⁸ Hooper-Greenhill, E (2000) “Museums and the Interpretation of Visual Culture”, Routledge New York

²⁷⁹ Bennett J, 1997, 'Kama and Eroticism: The Five Senses in the Work of Francesco Clemente and Pierre Klossowski', in Body, Bookman Schwartz, Melbourne, Australia, pp. 140 - 150

²⁸⁰ Muller K (2002) “Museums and virtuality” *Curator, the museums journal*. Vol 45, No 1, Jan 2002

precious materials, the skills required to create something of beauty as well, perhaps, as aesthetic values inherited from their culture. But for many objects the aura is derived not from the object's material and physical properties but its history, its unique 'biography', to borrow another of Benjamin's words. An otherwise mundane or common object becomes part of our heritage due to its connection with a historical figure or event and looking at aura in this way, we can see it something not inherent in the object itself, but a relationship between the viewer and object and depends on the viewer's knowledge of the object and/or on cultural values.

3.3.2 Three objects with 'aura'

For an example of this affectual power in a real museum object, Messham-Muir²⁸¹ recollects viewing an exhibit of shoes which was part of a larger holocaust exhibition at the Imperial War Museum. Focussing on one single shoe, the experience profoundly affected him, the "raw stark materiality" of the object providing a link between the observer and the shoe's owner, 60 years in the past. He talks of our relationships with objects and how they allow us to "enter into powerful empathic relationships that seem to transcend place and time"; the author here having a reverential experience that is clearly due to more than the stark physical facts of the exhibit. The experience is a result of three things: the context of the exhibit itself, including both the extra information imparted by the facts of the particular exhibit and the way in which the shoe is presented, as well as the larger context of both the holocaust exhibition and the Imperial War Museum itself; the sum total of the author's knowledge and

²⁸¹ Messham-Muir, 2006

experience as it pertains to the holocaust (and indeed, shoes); and ultimately the shoe itself, the material object with a biography that intersects the events of the holocaust. It is these three things working together that induces the empathic response and leads the viewer to a greater understanding of the horrors of the holocaust itself.

Another, perhaps less emotive example comes from my own experience and involves an exhibit in the Petrie museum of an old (several millennia old!) item of clothing (Figure 3.2). Whilst damaged and decayed it is still instantly recognisable as some sort of shirt, and, despite its condition, could quite easily pass for a piece of modern clothing. The item is impressive enough, as any sufficiently old and fragile object might be, but on its own it invoked no particular emotional response in this viewer. However the context of the exhibit, in this case the label text, explains how the object was found inside out on the floor of a dwelling, as if it had been taken off and thrown down²⁸², and it was only after reading the label that I experienced an involuntary and powerful emotional response, as if the millennia had been compressed: the simple action of taking off a piece of clothing and throwing it, inside out, on the floor, to be picked up later is such a basic human experience that I felt an immediate sense of empathy with the past. This feeling did not exist until I had read the information label, and so clearly is not a product of the object alone. On the other hand, it is obvious that the label alone would not have had the same effect without the object to which it refers. Again, we see the affectual power of an object as a product of three things; the context, in

²⁸² The 'as if' is important here, demonstrating that this is, of course, an interpretation of the facts – or simply speculation on behalf of the curatorial staff; an example of the importance of context and the malleability of an object's 'true' biography.

this case the label information, my own experiences (of throwing clothing on the floor), and the object itself.



Figure 3.2: Petrie clothing exhibit. Picture: Petrie Museum

The third example is that of a moon rock, specifically the sample displayed at the Science Museum, though a quick google search reveals there are many examples of lunar rock on display at museums around the globe (Figure 3.3). The moon rock in the exhibit in the Science Museum has a definite 'aura' due to its truly extraordinary biography, but the object itself is a nondescript lump of greyish rock which – to

practically any observer – could have been picked off the side of a road anywhere in Britain. Again, the aura is provided by the context in which the object is displayed: the entire exhibit calls attention to the small rock at the centre, and screams 'this is important'. It also serves the not unimportant task of signalling to the visitor that this rock is, in fact, from the moon. This information would be insufficient to create affectual content unless the visitor had prior knowledge of the effort that went into bringing this lump of moon back to Earth and thus the rock's rarity and value²⁸³. The particular aura we experience is a subjective one and thus a difficult concept to elucidate.




Figure 3.3: two examples of moon rock exhibits. Left: The National Mining Hall of Fame and Museum, Colorado right: Tellus Science Museum, Atlanta. Pictures: NMHFM & Charles Atkieson/examiner.com

²⁸³ Of course, the knowledge doesn't have to be 'prior', but could be learned 'simultaneously', if, as is fairly common, the moon rock is part of a larger exhibit on lunar exploration. This leads to the interesting interpretation of some museum exhibits, where a virtuous circle can arise from the synergy between object and information. The new knowledge acquired by the visitor from the exhibit increases the object's aura which increases the viewer's engagement which increases the amount of knowledge they acquire...

3.4 The cultural heritage object in the age of digital reproduction

So how does this concept of aura or an object's affectual power relate to digital, virtual copies? On the surface, it would seem to argue against the utility of public facing 3D models: if what makes a museum object 'special' is its affectual power, and its ability to affect is a product of an aura which is in turn a product of the object, somehow allowing us to experience and partake of its history, there doesn't seem much hope. As Messham-Muir puts it, it was the object's "raw materiality" which induced the response. How could our digital model of the shoe from the holocaust – even if it were our idealised perfect digital surrogate – evoke the same emotional response? Unless we are dealing with a purely aesthetic object, the aura seems to have little to do with the surface properties of the object, and it is the surface, after all, that is the sole domain of the digitised object.

One can argue that intangible things can, and do, evoke emotional responses – images can certainly have affectual power. However, emotionally resonant photographs tend to be pictures of events rather than objects, and thus the photographer's presence at an event creates a physical connection to a particular place and time – the photo has its own biography. A digitised model is just another kind of photograph, but they tend to be of objects, not events. It is conceivable that a 3D model of, for instance, a crime scene or aftermath of a disaster may have the same or more affectual power as a photograph, so there is nothing inherent in virtual models as a medium that precludes them from having affectual power, but in these cases it is unclear as to whether the model is of an object (or space) or an event.

Taking our digital model in isolation, the prognosis would be poor: it is hard to see where its aura, and subsequently its affectual power, will come from. However, as we have seen, museum objects are rarely viewed in isolation, instead they are presented in context and often used as tools for leveraging a viewer's existing knowledge. The shoe, shirt or rock, on their own, are mundane and unremarkable. In the contexts of their respective exhibits, they exhibit aura. So the question is not does our virtual object have an aura, but can an entire exhibit, in context, have affectual power when one part of it, the object itself, is replaced with a digital copy. From my own research, context is a vital part of a 3D model. The model of the shipping gallery is impressive in itself, and, for reasons discussed in (section 4.9) may have some affectual power even in isolation, but it is clear from the comments of users that, when combined with the background music, and more importantly the narration, which provides context for both the digitisation process and the object itself, the whole 'virtual exhibit' is capable of eliciting a strong emotional response, again based on three things: the digital model, the context, and the 'emotional baggage' people bring to the experience themselves due to previous experience with the original object or its subject matter.

Cameron²⁸⁴ argues, that simply due to the fact that the museum has chosen this object, and expended effort in making the model, that the virtual object inherits a portion of the museum's authority and therefore, presumably, some of the museum's aura-giving ability may apply. It is debatable whether the simple fact of a museum

²⁸⁴ Cameron, F (2007), "Beyond The Cult of the Replicant", Theorizing Digital Cultural Heritage Ed. Cameron F & Kenderdine S, MIT Press, Cam, MA

making a model can bestow aura, as surely this would depend to a certain extent on the viewer's knowledge of the process of selection and digitisation, but the idea of a virtual model inheriting some aura from existing in a museum environment may be a useful one: a digital model of a moon rock displayed in isolation will be a digital model of a grey lump; a digital model of the entire exhibit, with the rock initially displayed in its altar-like case, perhaps embedded in a larger virtual museum, may look like something else entirely. Similarly, a description of the process involved in 3D digitisation²⁸⁵ displayed alongside the model, certainly while 3D digitisation is still a novel and unfamiliar technology, may reinforce this idea.

Another reason for cautious optimism is that while we don't have a physical connection – there is none of Messham-Muir's 'stark materiality'²⁸⁶ – with the virtual object, we don't have a real physical relationship with the vast majority of museum objects anyway; we are usually separated by glass and/or distance. Pallud advances an argument based in phenomenology that simply seeing an object (the relationship we generally have with museum objects) is a flawed way of experiencing an object and that "this argument leads to the conclusion that being able to touch things *or to manipulate* them contributes to a better experience and to better interpretation"²⁸⁷ (my emphasis). Obviously without the use of haptic technology (which, while a valid

²⁸⁵ For an example, see: <http://www.courtauld.ac.uk/gallery/exhibitions/2013/illuminating/dish/3d.shtml>

²⁸⁶ Messham-Muir (2006)

²⁸⁷ Pallud J (2009) "The application of a phenomenological framework to assess user experience with museum technologies" ECIS 2009: 2375-2386

area of research^{288,289} is beyond the scope of this research) we cannot 'touch' virtual objects, we can manipulate them, and it is possible that this experience may lead to a richer engagement with an object than simply looking at the real thing. Whether this engagement is enough for the digital object to provoke an affectual response is an open question.

3.5 Conclusions

We have seen that 'digital surrogate' is a useful term to describe 3D models, and that the requirements for a model to be considered a digital surrogate are very different for professional and public facing purposes (and indeed, within those two classifications). While it may appear easier to create digital surrogates for the latter, the requirements for authenticity, and the capture of an object's aura in a virtual model are problematic. In the first case, the authenticity requirement, and thus the need to create a visually consistent model that successfully captures the object's unique properties drastically limits the types of objects (or materials) which are amenable to digitisation. However, as we saw in chapter 2, technology is constantly evolving and as BRDF or BTF measuring techniques become more accessible, it would be hoped that these limits will become less restrictive.

In the case of an object's aura, we have seen that its aura – or affectual power - is not

²⁸⁸ Zimmer R, Jefferies J, "Accessing material art through technologies of mediation and immediation", *Futures*, Volume 39, Issue 10, December 2007, Pages 1178-1190

²⁸⁹ Figueroa Pablo, Mauricio Coral, Pierre Boulanger, Juan Borda , Eduardo Londoño, Felipe Vega, Flavio Prieto, Diego Restrepo (2009) "Multi-modal exploration of small artefacts: an exhibition at the Gold Museum in Bogota", Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology, November 18-20, 2009, Kyoto, Japan

necessarily inherent in the object itself, but rather an emergent property of the object, its context and the user. In this case, it may be possible for the digital surrogate to inherit some of an object's aura if its context is also recreated; in short, a virtual museum object should be treated as a museum object. There are some indications in the next chapter that this may in fact be the case.

4 Scanning the Science Museum's Shipping Gallery

4.1 Introduction and Research Aims

This chapter is concerned with a project undertaken in May 2012, a collaboration between UCL, ScanLAB²⁹⁰, and the Science Museum which used terrestrial laser scanning to create a full point cloud model of the museum's Shipping Gallery prior to its decommission.

Research is conducted to evaluate the public response to such a project, the utility of recording galleries in this way and possible further uses for terrestrial scanning in museums and cultural heritage institutions. Initial findings, based on responses to an online survey as well as analysis of comments left on a variety of websites, suggest that there is a large appetite among the public for this kind of project, and that this is a potentially successful method for preserving, and indeed commemorating, large museum exhibitions.

However, when analysing responses to a digital resource such as the Shipping Gallery video, it is very difficult separating the audience's reactions to the resource from their feelings towards the source; in this case the gallery itself. This suggests that some of the aura of the original gallery is somehow preserved in the digitised copy and conveyed by the context in which it is presented.

4.2 The Shipping Gallery

Originally opened in 1963, by the time of its closure almost half a century later in 2012

²⁹⁰ <http://scanlabprojects.co.uk/>

the Shipping Gallery was the Science Museum's largest single space and its oldest surviving exhibition²⁹¹. Containing over 1800 individual exhibits, the gallery featured a mixture of models of important pieces of maritime technology, dioramas depicting the evolution of shipping around the world, historical objects and, perhaps most famously, a large array of model ships and boats representing important landmarks in the history of shipping²⁹².




Figure 4.1: Images of the Shipping Gallery, then and now. Top: Images from a 1963 New Scientist article. Bottom: recent visitor images of the gallery (Credit: Dave Patten, Flickr)

²⁹¹ Information from: <http://www.gizmag.com/shipping-gallery-3d/28844/>, http://www.sciencemuseum.org.uk/about_us/history/shipping.aspx (text and narration)

²⁹² One of the models, a 1:64 scale, four metre long recreation of the Mauretania, sister ship of the Lusitania, was recently sold at auction for £135,000: <http://www.antiquestradegazette.com/news/2015/may/20/ship-model-speeds-to-135000-record-at-charles-miller-auction/> (acc. 27/5/15)

The gallery itself remained largely untouched since its opening, and in some ways could be considered a museum exhibit in itself. While the rest of the Science Museum evolved over the last half a century, introducing interactive technology and a wider socio-historical perspective^{293,294} to its exhibitions, the Shipping Gallery remained almost exactly as it was at its opening, as if preserved in amber. Its floorplan felt very dense compared to other galleries in the museum (for comparison, the Information Age gallery which replaced the Shipping Gallery contains just 800 exhibits²⁹⁵), and whilst the exhibits were arranged with a rough chronological flow from one end of the gallery to the other, there was no clear path or obvious signposting. Labels on individual exhibits often featured long, manually typed chunks of text which, especially in regard to the technical models and items explaining marine technologies, contained esoteric language and assumed a certain level of knowledge and understanding.²⁹⁶

²⁹³ http://www.sciencemuseum.org.uk/about_us/history.aspx?page=4

²⁹⁴ MacDonald, S, 2001, *Behind the Scenes at the Science Museum: Knowing, Making and Using*, in Academic Anthropology and the Museum: Back to the Future, ed. Bouquet, M, 2001, Bergahn Books

²⁹⁵ <http://eandt.theiet.org/magazine/2014/10/analysis-science-museum-information-age.cfm>

²⁹⁶ Of course, this is all relative. It is interesting to note the following from a 1963 New Scientist article on the occasion of the gallery's opening: "The gallery devoted to sailing ships ... is a second, excellent demonstration of what can be achieved when a collection is mounted with up-to-the-date design and display. Alcoves are devoted to groups of ship models from different parts of the world and from different periods of history. In each case the background sets the models into their social context. A complete re-writing of all the old labels in the gallery has cut them down from 160,000 words - two complete novels - into brief and readable information for the non-specialist" Michealis, A R, New Scientist, No. 347, July 1963




Figure 4.2: Label text in the Shipping Gallery (credit: Dave Patten, Flickr)

There was little interactivity or attempt to place exhibits in a wider context, and any overarching narrative to the gallery would have to be provided by the visitor themselves.

The ‘old fashioned’ nature of the exhibition and its incongruity within the museum as a whole was reflected in attendance. Even during the museum’s busiest times the gallery was often deserted. On some of my own visits, during, for instance, school holidays when much of the museum was hard to navigate due to the volume of people, the Shipping Gallery may have had one or two – or, quite often, no – visitors²⁹⁷.

Despite – or, perhaps, because of – its anachronisms, the Gallery’s closure appeared to evoke strong emotional responses in those for whom it held special significance.

Evidence of, and reasons for this will be examined in more detail in the analysis of survey results and online comments in section 4.9, but it appears to be a result of both the gallery’s longevity –and consequent nostalgic associations – and its subject matter. Shipping and maritime technology is deeply entwined with the national image of Britain’s golden age as a maritime power and a global trading empire. Nevertheless, one cannot help but notice the disconnect between the strong responses to the gallery’s closure and the lack of interest when it was still there.

4.3 The project

4.3.1 The commission

The Science Museum decided to replace the Shipping Gallery with a new exhibition, the Making of Modern Communication, in 2013. I had previously shown the museum examples of the work of ScanLAB²⁹⁸ to demonstrate the potential of capturing large

²⁹⁷ From the comments on metafilter (<http://www.metafilter.com/130281/Scrapped-but-not-forgotten>): “This is sad also because the Shipping Gallery had a few benches where you could eat your lunch in complete peace, even during the school holidays Always deserted.”

²⁹⁸ ScanLAB is a terrestrial laser scanning company set up by two former students of UCL’s Bartlett School of Architecture. They create point cloud models for a diverse set of clients including architecture firms, television

spaces and environments with time of flight laser scanners, and the museum made the decision to record the gallery using this technology. Early in 2012 they approached ScanLAB with the possibility of scanning the entire space prior to its decanting.

The museum's initial commission for ScanLAB was for a complete scan of the gallery and the production of a two minute video from the data. The gallery's curator, David Rooney, would record a narration to be placed over the video and the full data set would then be made available to the public. The contract was purely between the museum and ScanLAB; however after introducing the two parties I was able to assist ScanLAB with the scanning project, document the processes involved and conduct research on creating interactive experiences with the data (section 4.8.3 & 4.7) and on the primary pre-rendered video output (section 4.9).

The amount charged for the project was approximately £4,000 which included both the scanning and the video output, though ScanLAB estimate the actual cost, and the figure they would quote if taking on a similar sized job today, to be between £40,000 and £50,000²⁹⁹. The low cost was partly due to an under-estimation of the size of the project – this was, at the time, the largest single job undertaken by ScanLAB – but mainly due to knowledge of the museum's budget and the treatment of the project as a marketing opportunity and loss-leader; the high profile nature of the campaign and the media attention surrounding it leading to further work for the company.

productions, artists, Greenpeace and cultural heritage organisations. See scanlabprojects.co.uk for examples of their work.

²⁹⁹ Figure quoted by ScanLAB in personal communications

4.4 Capture

4.4.1 Targets

Prior to scanning, approximately 200 A4-sized paper checkerboard³⁰⁰ targets were evenly distributed throughout the gallery space. The numbered targets were placed predominantly in high positions that could be seen from as wide an area as possible (Figure 4.3). These targets, visible in many scans, were used as common points during registration. The targets were left in situ throughout the scanning period.



Figure 4.3: View of gallery showing paper targets

As well as paper targets, spherical targets were used to aid the registration of clusters of scans as well³⁰¹. Spheres are useful for registering individual scans as the capture of

³⁰⁰ Both the checkerboard targets and the spheres are automatically identified by the Faro software used for registration, see 4.6.1

³⁰¹ Becerik-Gerber et al, Assessment of target types and layouts in 3D laser scanning for registration accuracy, Automation in Construction, Volume 20, Issue 5, August 2011, Pages 649-658

even a portion of the sphere's surface, from any direction, is enough to determine the sphere's centre. Therefore a common point (the target's centroid) can be determined for two scans which capture a part of the sphere even if the scans are from opposite directions with limited overlap and which share no common points on the target's surface. The spherical targets are also automatically recognised in the scan data by most terrestrial laser scanning software (including Faro Scene, used in this project) and thus simplify the processing of multiple scans.

The spheres, attached to magnets, could be placed around the gallery (Figure 4.4) – for example, attached to the display cases themselves – again in positions visible from multiple scans. In this case however, the targets were left in situ for three or four scans before being moved for the next cluster. The targets, both paper and spheres, are visible in both the scan data and video.



Figure 4.4: Spherical targets

4.4.2 Control survey

Before scanning, and with the aid of Anita Soni, a colleague from UCL CEGE, the gallery space was surveyed and a control network established using a Leica TS15i³⁰² total station. 11 control points were established from which a selection of paper targets, spaced evenly around the gallery, were measured to within mm accuracy. A least squares network adjustment was made using Starnet³⁰³ software to minimise residual errors.

The survey provides a 'ground truth', ensuring that physical measurements of the gallery's dimensions taken from the scan data could be made with a high level of

³⁰² http://www.leica-geosystems.co.uk/en/Leica-Viva-TS15_86198.htm

³⁰³ <http://www.microsurvey.com/products/starnet/> Version 7 was used.

accuracy and precision³⁰⁴, and also constituted a framework within which the laser scans could be registered (see 4.6). Without an externally referenced and verified coordinate system, simply registering individual scans or clusters of scans to each other could lead to errors propagating through the data and it being impossible to ‘close the loop’³⁰⁵. This was the first time ScanLAB had used a survey before a scanning project, and the extra data provided would prove to be invaluable when registering the scans and processing the data later (see section 4.6).

4.5 Scanning

Two Faro Photon 120³⁰⁶ laser scanners were used for the project. The Photon is a phase based³⁰⁷ time of flight laser scanner which uses a rapidly rotating mirror to capture up to one million points per second, with a maximum range of 120m and an accuracy of +-2mm at 10m. As well as an xyz position, the points captured have an intensity value depending on the strength of the laser reflection; to add colour, a DSLR camera is attached to the scanner via a bracket (see Figure 4.5), and the tripod is wound down so that the camera’s optical centre is aligned with the point where the laser is emitted. The scanner revolves while the camera takes a series of wide-angled photos, which are then stitched into a single panoramic image. By projecting this panoramic image onto the scan, individual points are given an RGB value resulting in a coloured point cloud.

³⁰⁴ Important if the museum wished to use the data for other purposes, such as BIM or exhibition planning

³⁰⁵ Ie, if for example the scans were registered in a clockwise loop from one corner of the gallery, even if errors between individual scans were small, the sum of these errors may grow to a point where the registration of the final scan with the first would be impossible.

³⁰⁶ http://www.dirdim.com/pdfs/DDI_FARO_Laser_Scanner_Photon.pdf

³⁰⁷ Pfeifer, N., & Briese, C. (2007). Laser scanning-Principles and applications. p6

The two scanners were used simultaneously, starting in opposite corners of the gallery and moving down alternate 'legs' in such a way that they would not interfere with, or appear in, each other's scans. Scanning was conducted in separate clusters of 10-15 scans, and where possible the scanners were placed directly in front of glass cases containing the exhibits in order to capture the models or objects in as much detail as possible and to minimise noise caused by reflections and refractions by the glass (see section 4.6). In Figure 4.8, the black or dark circles visible in the rendered images are areas directly below the scanner which cannot be captured and thus give an impression of the spacing between scans and their positions.



Figure 4.5: Faro 120 scanner with DSLR camera mounted

Clusters were planned in advance and small round stickers placed on the floor to indicate scan positions; once the scan was completed the stickers were marked with an 'x'. In this way a record was kept of which positions had been scanned, ensuring that positions were neither missed nor duplicated. As mentioned above, spherical targets were positioned and repositioned for small groups of 3-5 scans within these clusters. A total of 275 scans, capturing approximately ten billion points, were made over five nights.

4.5.1 Scanning: Methodological Issues

Scanning could only be conducted whilst the museum was closed to the public, roughly between the hours of 7-8pm and 6-7am. This had several ramifications, not least of which, according to ScanLAB, was fatigue³⁰⁸. Working five consecutive night shifts (whilst still working on other projects, or, for example, sourcing replacements for malfunctioning hardware during the daytime) caused tiredness which in turn could lead to mistakes. For example, the process of taking the colour photos after the scan involves several steps: attaching the camera to the bracket, sliding the bracket to a central position and winding the tripod to the correct position. Missing one of these steps would mean the process would have to be repeated; if the scanner was moved before the mistake was realised, the entire scan would have to be repeated.

Similarly, if the reverse was forgotten, ie, the tripod wasn't wound back up (meaning the camera could not be used from the correct position) or the bracket not slid back to its normal operating position (thereby occluding part of the scan), the next scan would be useless and would have to be repeated. Small details like this, when combined with fatigue were potential sources of frustration and loss of valuable time.

Working at night introduced problems as well. For instance, the museum's lifts were switched off afterhours and thus a large amount of bulky, heavy equipment (the batteries alone weigh 10 kilos each) had to be carried a considerable distance through the museum and up several flights of stairs at both the beginning and end of each

³⁰⁸ This aligns with personal experience; I spent three and a half nights with ScanLAB at the museum (whilst doing very little during the days!) and by the end of the project the fatigue was certainly noticeable.

scanning session. This time must be factored into the project.

The control of lighting in the museum is outsourced to an external company, and whilst they had been informed of the project and its requirements, at midnight of the first night of scanning, all the lights in the gallery were switched off, and could only be switched back on remotely. It proved impossible to contact the company and so most of that night was lost.

Battery life for the scanners is limited to about four hours and therefore each one needed to be recharged once or twice per night. To avoid downtime, a temporary power source consisting of a case containing car batteries was used to power one scanner whilst its battery was recharging. If necessary, the scanners could be run off the mains, though this entailed trailing long extension cords around the gallery which would subsequently appear in the scans. A close eye had to be kept on charges and charging times to avoid this scenario as much as possible.

At the end of each scanning session, data needed to be downloaded from the scanners' on board 80gb hard drives onto laptops. This could take up to an hour, more time which needs to be factored into the process.

4.6 Processing

Whilst the scanning, spread over five nights, took approximately 120 man hours, processing the captured data took approximately four months, or 1350 hours. The processing took place on fairly modest hardware: two 8Gb MacBook Pros using Faro

Scene³⁰⁹ (version 4.1) software.

4.6.1 Colourisation & Registration

After the point clouds were colourised with the panoramic photographs, the scans were registered in clusters of 10-15 individual scans using a combination of the paper targets and spheres. Due to limitations on the amount of data the software could process at any one time, the clusters were not registered with each other. Instead the locations of the checkerboard targets, as measured in the control survey, were imported into Scene, and each cluster was placed in the coordinate system by registering it with these targets.

Whilst the registration process was partly automated, there was considerable craft involved in working within and around the software's limitations and idiosyncrasies, and the entire process took three weeks (approx. 240 hours). ScanLAB point out³¹⁰ that if the project were to be repeated today, improvements in Faro Scene software coupled with more powerful hardware would speed up the process, but only to a certain extent; there is still a considerable time cost involved.

4.6.2 Cleaning

Once the scans were registered, the data was cleaned of noise, a large amount of which was caused by the glass in the display cases. Whilst the glass was largely invisible where the laser hit orthogonally, glancing angles caused false points through both

³⁰⁹ Scene is point cloud software developed specifically for use with Faro scanners: <http://www.faro.com/en-us/products/faro-software/scene/overview>

³¹⁰ pers. comms

reflection and refraction. The majority of these unwanted points showed up as halos of bad points around the edges of the glass cases, and all had to be removed manually. Whilst automatic filters could be – and were – used up to a point, the filtering algorithms could not distinguish between isolated noise and the fine details on the exhibits (for example, masts or rigging on the ship models). The laborious cleaning process took two people at ScanLAB a total of three months (approx. 1000 hours), by far the longest (and thus most expensive) part of the entire project.



Figure 4.6: Images of the raw scan data showing some of the sources of noise connected with the glass cases

4.7 Output

Once the data was registered and cleaned, it was imported into Pointools³¹¹ to create the animation. However, as Pointools could only handle approximately 2 billion points, the amount of data used had to be reduced. To get the data down to a usable size, areas of little interest (ie, floor, walls and ceiling) were sampled down to roughly 10%, whilst for the exhibits and models, approximately 20% was used. Overall, only about 17% of the total data – still two billion points or half a terabyte – was used to produce the video.

The curator, David Rooney, had identified ten models or ‘points of interest’ in the gallery, and thus the initial process involved finding a camera path through the gallery which would identify each of these exhibits. An iterative process was used, which involved plotting a path, rendering several frames, examining the result and if necessary tweaking the camera’s trajectory. An initial complete path sticking to the original commission of a two minute video was unsatisfactory, and instead a longer, more ‘cinematic’ option was created. For the individual models, such as the marine turbine (see Figure 4.7), separate animations were created using the entire (ie, not down-sampled) data set, and these were spliced into the video at the appropriate locations.

³¹¹ Pointools is point cloud editing software capable of rendering large data sets at high quality: http://www.bentley.com/en-US/Promo/Pointools/pointools.htm?skid=CT_PRT_POINTOOLS_B

The final animation was seven minutes long, which the museum were happy to accept.

The entire process of composing the video – planning and refining the camera path through the virtual space – took another three weeks (approx. 240 hours), and the completed animation took a full 48 hours to render³¹². The finished clip was sent to



Figure 4.7: Turbine model stills from video. Left - fully detailed model, Right – using the down-sampled data

David Rooney who recorded his narration to fit the video. ScanLAB also commissioned Box Of Toys Audio³¹³ to compose incidental music and add some ‘marine’ sound effects, and finally to composite this audio with the narration. The finished video was published on July 22, 2013 and can be viewed on the Science Museum website³¹⁴ and on the Science Museum’s YouTube channel³¹⁵.

³¹² Again, using a MacBook Pro

³¹³ <http://www.boxoftoysaudio.com/> like ScanLAB, Box of Toys charged a fraction of the market rate for their work

³¹⁴ http://www.sciencemuseum.org.uk/about_us/history/shipping.aspx

³¹⁵ <https://www.YouTube.com/watch?v=gDTbFhFZl9I>, see appendix C for some other websites where the video can be viewed




Figure 4.8: Images from the animation; top left - a screenshot from the YouTube video, others - images rendered directly from Pointools

4.8 Further work

The video provided to the Science Museum was created by rendering a predetermined path through the point cloud. An interactive model was not one of the outcomes required by the museum, and would be, due to the size of the dataset, a formidable challenge. This section details some attempts to turn the Shipping Gallery scan into a

more interactive experience, something that research shows would be extremely popular with users (see 4.9.5).

4.8.1 Euclideon

Euclideon³¹⁶, an Australian company creating a point cloud rendering solution predominantly for the games industry, created an online demo using a portion (about half) of the cleaned and registered point cloud. Their ‘solidscan’ technology is described as a ‘voxel rasteriser’ and (presumably) consists of a proprietary data structure and search algorithm which only renders those points which will be visible on the display, so rendering speed is limited by screen resolution and not the size of the point cloud.

The resulting demo required a plug in to be downloaded, and whilst it did allow smooth navigation through the virtual space, it streamed the point cloud as voxels using an LOD (Level Of Detail) process which gradually (and slowly) increased the resolution whilst the user was stationary. The user experience was not particularly good and suffered from the same issues as the WebGL experiments (see 4.8.3); the scan resolution and point density acquired from the terrestrial laser scanners is not high enough to allow interactivity and close up inspection of objects in the point cloud (see Figure 4.9, Figure 4.10 & Figure 4.11).

Ultimately, the results of ScanLAB’s experiments with Euclideon were disappointing and further partnership with them proved too expensive to pursue. With current

³¹⁶ See: <http://www.euclideon.com/>. Though the company is still active, this website does not appear to have been updated since 2013.

technology it is unlikely that a fully navigable, online version of the Shipping Gallery could be created, certainly without considerable expense³¹⁷.

4.8.2 Amazon

Amazon have also taken a copy of the data³¹⁸, with the intention of creating an augmented reality app that would allow visitors to the current gallery to look ‘through’ their mobile devices and see the space as it was. The status of their experiment is unknown at present, despite prompting.

4.8.3 WebGL

Experiments were conducted by myself using WebGL to render interactive models of some of the individual exhibits taken from the animation. The website³¹⁹ was developed using XBPS (Cross Browser Point Stream), an open source WebGL point cloud viewer³²⁰. Figure 4.11 show screenshots taken from the viewer.

The files for the models were provided by ScanLAB in ascii format (for interoperability) and which were converted to .psi files in Pointstream. The original models consisted of approximately 10 million points each, and a sampling of around one third, or 3 – 3.5 million points, was found to be a suitable compromise between detail and usability

³¹⁷ It has been suggested that it could be achieved using a system similar to the now defunct OnLive gaming service (onlive.com), where the ‘heavy lifting’ of rendering the point cloud is carried out on powerful servers and each frame is then streamed to the client. However, you would still need powerful technology running in the back end, as well as expensive infrastructure, bandwidth etc.

³¹⁸ From pers. comms with ScanLAB

³¹⁹ http://www.homepages.ucl.ac.uk/~uczcjhi/Sci_mus_viewer/index.html (Works best in Chrome) Click on the image to see a 3D model; only ‘figurehead’, ‘turbine’ and ‘Vanguard’ have been implemented. Due to changes in browser technology since the code was written, the ‘free cam’ mode which allowed mouse control is no longer supported. Controls are: WASD to move in the x and z planes, arrow keys to look around, insert + delete move the camera up and down along the y axis whilst + and – increase and decrease the point size.

³²⁰ The project was funded by Pointstream and the technology formed the basis of Pointstream.js.
http://zenit.senecac.on.ca/wiki/index.php/XB_PointStream

(including both loading times and framerate). No other modifications were made to the original data.

Whilst these were rough 'proof of concept' experiments, and more work could have been done, for example, cleaning up the models to remove noise, it is apparent that with the level of detail available, even using the full, unsampled data, the models (with the possible exception of the figurehead due to its simple 'large scale' geometry) would not stand up to close inspection or give a satisfying user experience.



Figure 4.9: Close up of HMS Monarch rendered via WebGL. Whilst the hulls and large scale features of the ship models were captured fairly well, the limited resolution of the scanners failed to pick out details such as masts and rigging.



Figure 4.10: Figurehead, Top right with large point size, bottom right with minimum point size. This model, with its relatively simple geometry, proved to be the most successful object




Figure 4.11: Marine turbine, bottom left with large point size, bottom right, minimum. Due to the small size of the detail on this object (a 'spaghetti'-like profusion of small pipes), and the occlusion due to the limited number of scan angles, this model is unsatisfactory under close inspection

4.9 Survey & Online Analysis³²¹

Following guidelines from a variety of resources and institutions including TISDR^{322,323,324}, and methodologies used in a number of other analyses of digital resources in the cultural heritage field³²⁵, a mixture of quantitative and qualitative³²⁶ research was conducted around the online video using a combination of survey results and analysis of online comments.

4.9.1 Methodology

4.9.1.1 Quantitative (survey) methodology

As the main output was an online resource (the video on YouTube/the Science Museum website), an online survey was created using the Opinio platform (via UCL services³²⁷) which was accessed via a link placed below the video on the Science Museum website³²⁸.

Since the survey is concerned with evaluating a resource only available via the web, one of the main issues with online-only surveys³²⁹ doesn't apply as our target population is limited to people with internet connections/web users. However, another

³²¹ For full details of all responses see Appendix C

³²² The Toolkit for the Impact of Digitised Scholarly Resources, a Jisc funded project at the Oxford Internet Institute: <http://microsites.oi.ox.ac.uk/tidsr/>

³²³ <http://www.lse.ac.uk/media@lse/research/EUKidsOnline/BestPracticeGuide/Home.aspx>

³²⁴ http://www.postgrad.com/editorial/advice/phd/research_methods/_...qualitative_vs_quantitative_research/

³²⁵ For an example of a mixed quant/qual approach, see: Grincheva, N. (2015). 'The World Beach Project' Going Viral: Measuring Online Influence—Case Study of the Victoria and Albert Online Museum Project; Journal of Creative Communications, Vol 10 iss. 1, pp 39-55.

³²⁶ In the case of this research, the line between quantitative and qualitative is not drawn sharply

³²⁷ <http://www.ucl.ac.uk/isd/services/learning-teaching/elearning-staff/core-tools/opinio>

³²⁸ http://www.sciencemuseum.org.uk/about_us/history/shipping.aspx

³²⁹ See: Couper, M, 2000, Web Surveys: A review of Issues and Approaches; Public Opinion Quarterly, Vol 64. Iss 4 pp 464-494

issues brought up by Couper – ‘nonresponse error’, or the target audience not responding (in sufficient numbers) was an issue. The resulting small sample size can lead to sampling errors (and reduces the statistical significance of the results), though as we shall see the survey respondents exhibit a certain amount of self-selection anyway. While this does mean our sample is not, perhaps, entirely random or a valid reflection of the frame population as a whole, there are other reasons to believe this is not necessarily a fatal problem for this particular analysis. These will be discussed during the analysis and in the conclusion.

Given the target audience was the general public, the survey was kept short – under five minutes – (to ensure completed responses) and was written in a conversational tone using fairly informal language³³⁰. The survey consisted of 12 multiple choice questions (Figure 4.12) each with either four or five possible responses, plus one open-ended question with a free-text input asking for any further comments. In total there were 35 responses, 34 of which completed all questions with one respondent (776753)³³¹ answering none of the multiple choice questions but leaving a long free text comment.

These numbers, and the nature of the survey (with only five discrete answers possible per question) are not conducive to rigorous statistical analysis, and while statistical

³³⁰ Compare and contrast with TIDSR’s ‘example’ survey, Meyer, E.T., Madsen, C., Eccles, K. (2009). TIDSR Survey on the Use of Digitised Resources. Available online: <http://microsites.oi.ox.ac.uk/tidsr/kb/27/tidsr-survey-use-digitised-resources>. This is a 20 minute survey and is aimed at a much more limited domain - professionals in the field of cultural heritage. There are 47 questions, some of which have multiple sub-questions each with up to ten potential responses; it was judged that this sort of comprehensive questionnaire would be too daunting and involved for a non-professional general audience.

³³¹ Survey respondents are identified by a unique number assigned by Opinio

tests do suggest correlations between certain questions, with the number of respondents it is impossible to claim statistical significance for these results. Therefore any suggestions of correlations in the following analysis should be considered with this in mind.

Self-selection among survey respondents is a problem with any online community³³², and in examining both responses to the survey (particularly the free text comments), and responses to the video on the web in general, it is clear that the respondents are not drawn randomly from a cross section of the general public. If one assumes that respondents to the survey have watched the entire video, then the fact that they are willing to sit through a seven minute video suggests they either have some prior interest in laser scanning, or, more likely (judging from the responses) the Shipping Gallery itself. This is not necessarily a problem for this research, however, and in itself throws up some interesting results (See section 4.10).

³³² Wright, K. B. (2005), Researching Internet-Based Populations: Advantages and Disadvantages of Online Survey Research, Online Questionnaire Authoring Software Packages, and Web Survey Services. *Journal of Computer-Mediated Communication*, Vol 10, iss 3

Question 1: How did you originally find this video?

Question 2: How familiar are you with the Science Museum?

Question 3: How familiar are/were you with the Shipping Gallery?

Question 4: Did you enjoy the video?

Question 5: Do you think this is a good way of preserving old exhibitions?

Question 6: Would you like to see this sort of thing done for existing exhibitions? For example to help plan a visit to the museum?

Question 7: Imagining for a moment that Shipping Gallery exhibition still existed, would this video have made you more or less likely to visit the gallery?

Question 8: Would you like to have been told more about the technology used to create the video?

Question 9: Would you like to be able to explore the model yourself (ie, 'walk around' the gallery) rather than follow a video?

Question 10: This video showcases a very distinctive style, with the objects in the gallery appearing translucent and almost 'ghostly'. Did you like this style of presentation, or would you have preferred to see something more solid?

Question 11: Do you have much experience with virtual worlds and environments on computers? Ie, perhaps through playing games?

Question 12 Please tell us how old you are...

Figure 4.12: Science Museum Survey questions

4.9.1.2 Qualitative (comment analysis) methodology

As well as the survey responses, comments³³³ have been gathered for analysis from a variety of sites which hosted, or linked to, the video, predominantly the Science Museum's YouTube channel and metafilter³³⁴. Relevant sites were found through a

³³³ Approximately 100 comments were accessed from six sites. Spam, empty and entirely irrelevant comments were not recorded.

³³⁴ A full list of comments and the websites they came from can be found in appendixA, which also contains a

combination of standard Google searches (looking for combinations of the keywords ‘Science Museum’, ‘Shipping Gallery’, ‘scan’ ‘3d’; etc.) and image searches, identifying sites which were using the stills released as publicity images, via the Tineye³³⁵ engine. A list of websites covering the release of the video, and any comments taken from those sites, can be found in appendix 9.2. Included is a collection of tweets regarding the video (again, found through searching Twitter with similar terms to those listed previously), for simplicity’s sake these will also be referred to as ‘comments’.

The majority of comments, particularly on the most popular sites (for instance, YouTube and Twitter) were empty of (for our purposes) useful content – for instance, notifications that the content had been shared, spam, or personal communications between users; these comments were ignored or removed (though, obviously, still visible on the original websites) The remaining comments were analysed using a form of grounded theory³³⁶, in that there were no preconceptions or hypotheses in mind when the analysis was conducted. Rather, themes and patterns of common or shared concepts (for example, the repeated use of ‘spooky’ or ‘ghostly’ language to refer to the video) were allowed to emerge naturally from the data.

Where appropriate, the comments have been analysed and included alongside the survey results. Whilst little quantitative conclusions can be drawn from them (apart, perhaps, from a general liking for the video), qualitatively they often reflect and

list of other media sources (both on and offline) which featured the project or embedded the video. Where quoted in the text, comments are identified by website and username.

³³⁵ <https://www.tineye.com/>

³³⁶ Glaser, B. G., & Strauss, A. L. (2009). *The discovery of grounded theory: Strategies for qualitative research*. Transaction Publishers.

reinforce the survey findings. The similarity of sentiments expressed via unsolicited comments on a variety of websites and those in the survey itself also provide confidence that the survey, with certain caveats, is in fact representative of a wider audience.

4.9.2 Demographics

The majority (18/34, 55%) of respondents found the video through friend recommendations. As there was no specific social media category, it is unclear how many of these recommendations were through specific personal correspondence and how many were through, for example, Twitter, Facebook etc. The other respondents found the video through a mixture of other websites, (print) news articles and from browsing the Science Museum website.

Divided by age (from Q12), the largest single group of respondents were aged between 34 and 50 (16/34, 47%), and the second largest (10/34, 29%) were over 50. Only six respondents (17%) were from the 25-34 group and just two from 16-24. There were no respondents under 16.

From Q3, “How familiar are/were you with the Shipping Gallery?” respondents to the survey exhibited a greater familiarity with the gallery than one would expect from a random sample of the general public. Whilst approximately one third (11/34, 32%) had never visited the gallery (all respondents were, however, aware of the Science Museum), nearly half (15/34, 44%) had either visited in the last three years (7/34) or were regular visitors who ‘loved the gallery and [had] visited often’ (8/34). This implies a certain amount of self-selection among the respondents – people more familiar with

the gallery were either more likely to follow a link and watch the video, and/or were more likely to go on to fill in the survey having watched the video. Examining both the quantitative and qualitative results of the survey, particularly the comments left in Q13, it is clear that many of the survey respondents have both strong feelings about, and personal experience with, the Shipping Gallery. This can be seen from the use of emotive language in comments like, “it is a frustrating tantalising reminder of what an asset we used to have ... I would forgive you for the loss of the gallery...” (783283), “I was so disappointed to discover that one of my favourite galleries has gone.” (773459), or “it was a crime that it was destroyed” (776753). These comments are indicative of the depth of feeling the Shipping Gallery inspired in some visitors, and these same sentiments are reflected in comments left elsewhere on the web (“I’m devastated that the gallery is gone”- Anthony Cooper, YouTube³³⁷). This emotional connection could be due to both the longevity of the gallery itself and its connection with the museum’s (and by extension, its visitors’) own history. The gallery represents a particular past and is – or was – a relic of a certain type of museology, which taking into account the ages of the survey respondents, may resonate with powerful memories and personal experiences (For example: “The Shipping Gallery has been with me since I was a child, I visited with my Grandad and my father” (777484)). The gallery’s demise, therefore, comes to represent a disappearing past and is symptomatic of a wider trend in museums: “In the digital age what we need more and more are galleries of real objects, not images or screens ... Give me real objects and day, not more screens.”

³³⁷ All YouTube comments are taken from below the video on the Science Museum’s own channel, <https://www.YouTube.com/watch?v=gDTbFhFZl9I>

(776753), “It makes the whole thing look a bit fake which is a shame for such a real place like the science museum full of REAL things such as history and weight and material.” (795195), “A shame so many iconic galleries are vanishing, too much emphasis is placed on visiting school children, and one of the world’s great museums is becoming a Disney play area,” (Lawrence Windrush, YouTube), “[the end of the Shipping Gallery is] everything I hate about modern museums. Remember when museums used to have old stuff? And lots of it?” (EnterTheStory, Metafilter)³³⁸. It is interesting to note, however, that negative feelings towards the closing of the gallery are not necessarily reflected in feelings towards the video itself (see section 4.9.3). Apart from personal feelings towards the gallery and museums in general, there is some indication that the particular subject matter of the gallery is relevant, with shipping and naval technology inextricably linked with national self-image and nostalgia for the days of empire – “a poignant reminder of kinder days when governments cared about our heritage” (776696), “it is as if the Nation is turning its back on its maritime tradition” (778798), “The Science Museum now has almost nothing about the marine technology that had a huge impact on Britain’s history.” (David Shirres, YouTube).

Thus when examining responses to the survey, it is difficult to separate thoughts on the video, and laser scanning in general, from thoughts on this particular gallery. While this makes generalising on the utility or efficacy of digital resources problematic, it may also have wider ramifications for digitisation; the fact that people’s experiences with the

³³⁸ <http://www.metafilter.com/130281/Scrapped-but-not-forgotten> 10/3/15

original object influence their interpretation of the digital is evidence that 'aura'³³⁹ and other metaphysical baggage of physical objects can be inherited by their digital copies. (For further discussion on this point see chapter 3)

As an aside, the observation that people more familiar and more emotionally invested in the gallery are more likely to have completed the survey may also go some way to explaining the age of responders. As could perhaps be expected from both the subject matter and the age of the exhibition, familiarity with the gallery (Q3) and age appear to be slightly correlated (Figure 4.13), with the older age groups being more familiar. This would then tend to skew the ages of respondents towards the older end.




Figure 4.13: Age vs familiarity

³³⁹ In the Benjamin sense: Benjamin, W (1939)

4.9.3 Reception

Results from Q4 (Figure 4.14) demonstrate that, certainly amongst the survey respondents, the video was well received. The average response is 2.23 (with 1 indicating the highest level of enjoyment, 5 the lowest and 3 being average), and the largest single response was '1', (12/34, 35%) judging the video as 'excellent'. Only 15% (5/34) rated their enjoyment of the video as below average. Positive comments left on the survey ("wonderful use of technology, well done!" (793924), "Thank you for undertaking and achieving something so valuable" (777484), "Ingenious and what a lovely film" (794252), "Very nice!" (772762) outnumber, and are generally less ambivalent and more enthusiastic, than the negative ("but the fly through tour is far too superficial to be useful, except to give a very limited first impression" (794123), "the level of detail presented was frustratingly low" (776740), "This video is no substitute whatsoever for one's being able to examine the characteristics of individual ship models in the gallery" (1128793)).




Figure 4.14: Q4 "Did you enjoy the video?"

This impression that the video was generally well received is reinforced elsewhere on

the web, in the case of The Science Museum's YouTube channel³⁴⁰, both quantitatively and qualitatively. According to total number of views, the Shipping Gallery video is the 18th most watched (out of 192) with just over 27000 views. However, looking at the 30 most popular (most watched) videos on the channel, and taking the ratio of 'likes' to views, the Shipping Gallery (134 likes) is clearly in first place with 4.96 likes per 1000 views. To compare, in second place is the "James Watt: First Hero Of The Industrial Age" video with 37,000 views and 141 likes at a ratio of 3.81 likes per 1000 views, while the average ratio for the 30 most watched videos is 1.55 likes per 1000 views.

Comments on the YouTube page also reflect this, with the positive³⁴¹ (eight, including "Wow, 3D laser scanning works better than I imagined it would", "Amazing use of technology..." "Fascinating... and a beautiful video too", "Stunning :)", "probably the coolest virtual gallery of all time", "absolutely stunning use of technology! ... [I] never imagined quite how stunning the results could be... marvelous.") outnumbering the specifically negative (two, "Very disappointing...", "a little unsatisfactory as shown"). Similarly positive comments were left elsewhere on the web ("I actually applauded at the end of that." (Spesh, B3ta.com), "beautiful" (Jeribus, metafilter)).

In terms of who enjoyed the video, there is an apparent correlation³⁴² (Figure 4.15) between the age of the respondent and their enjoyment, with younger age groups

³⁴⁰ <https://www.YouTube.com/user/sciencemuseum/videos?flow=grid&view=0&sort=p>, accessed 4/4/15. See Appendix x for details

³⁴¹ With respect solely to the video; there are also negative comments about the decommissioning which make no reference to the digitisation

³⁴² In this case, despite the low numbers and the type of data, the correlation is statistically significant with $r(32) = .39$ and $p = <0.05$

showing a greater level of enjoyment. There also appears to be a correlation (though less strongly indicated) between familiarity with the gallery (Q3) and how much the video was enjoyed, with those more familiar with the gallery enjoying the video the most. (Strangely, this would seem to be contrary to what we would expect, considering the previous finding that age is inversely correlated with familiarity (Figure 4.13)) This result implies that familiarity with, and, extrapolating further – emotional attachment to – the gallery is reflected in an increased level of enjoyment of the video. Indeed, just taking the responses from those who used particularly emotive language in their free text comments³⁴³, the average enjoyment score is 1.83, ie, significantly better than the score for all respondents.




Figure 4.15: age (Q12) vs enjoyment (Q4)

4.9.4 Potential uses

Questions 5, 6 and 7 show that 3D scanning as a method is viewed as a positive

³⁴³ See Appendix C.

technique, and could be useful for institutions both as a means of preserving galleries and as marketing material for new and existing exhibitions. To Q5 ("Do you think this is a good way of preserving old exhibitions?") 24/35 or 69% of responses were positive (18/34, 53% very positive) with two non-committal and 8/34, or 24% indicating they did not think this was a good method. But according to the survey, the technique has potential beyond just retrospectively recording old exhibits; Q6 ("Would you like to see this sort of thing done for existing exhibitions? For example to plan a visit to the museum?") and Q7 ("Imagining for a moment that the Shipping Gallery exhibition still existed, would this video have made you more or less likely to visit?") show that there may be utility for museums in capturing current spaces. These questions show that providing similar videos for current exhibitions could be useful for users in planning visits (26/34, 76% responded positively to Q6), and may also actively encourage visits to the museum (again, 26/34 responding positively). Perhaps unsurprisingly, the people most likely to visit the gallery after seeing the video were those most familiar with the gallery and had visited often (with an average score of 1 – ie, all those who answered '1' to Q3 said they would definitely have visited after watching the video), however, the group that demonstrated the second highest likelihood of visiting were those who answered '4' to Q3, ie, those who had never visited the gallery before, indicating that a resource such as this may help encourage visitors to try out new exhibitions.

4.9.5 Further development

By far the strongest response to any of the survey questions was for Q9 ("Would you like to be able to explore the model yourself (ie, walk around the gallery) rather than

follow a video?”). All 34 respondents answered either ‘1’ (“Yes, that would be amazing”) (26/34, 76%) or ‘2’ (“It would be nice to have some control”). (8/34, 24%) None of the respondents either “weren’t bothered” or expressed a negative view. This view is reflected in both the text comments left on this survey and elsewhere on the web, with the ability to navigate the space often linked to a wish to be able to more closely examine individual models and exhibits. For example, “The gallery and its models were an essential research tool. Access to 'solid' virtual models would allow it to continue ... if the fidelity can be improved and I control my on movement around the space I would forgive you for the loss of the gallery.” (783283), “but if more exhibits are going to be scanned it would be far better to have something more solid to fully do it justice” (787921), “I hope someone makes a Oculus rift viewable version of the Shipping Gallery exhibition”³⁴⁴ (793924) “The idea is good, but the fly through tour is far too superficial to be useful, except to give a very limited first impression. A fully explorable model, with the ability to view the exhibits (and labels) close-up, could be a very valuable tool” (794123), “I wanted to see the detail of the exhibits, the level of detail presented was frustratingly low. A good way to get taster but I don't live near London I would like the option to explore each item in detail.” (776740). Again, these comments were echoed many times on other websites (“will the finished version allow one to look closely at an individual item and read its caption?” (monotreme, Metafilter), “surely the artefacts are the important bits? If they had been scanned in

³⁴⁴ The Oculus Rift was mentioned in a surprising number of comments across various websites. This is probably due to the large amount of media surrounding the 1080p HD version's debut around the E3 conference in June 2013, less than a month before the Shipping Gallery was released.

high detail in isolation, they could be rendered into a virtual gallery” (ChrisH, Pc Pro), “I hope a full display of the data points will look better than these interesting but more ghost-like images” (Rhadiem, YouTube)). Though many of the wishes expressed in the comments are unrealistic given the data acquired and current rendering technologies (see 4.8.3), and conflate the two separate issues of having control and being able to see more detail, it does demonstrate an appetite for further use and provides a case for the development of 3D capture and display technologies. It also suggests that combining techniques may be a solution, for example, using terrestrial scanning to capture overall spaces and close range scanning or photogrammetric techniques to create more detailed models of individual objects.

4.9.6 Style

The style of the video – the translucent ‘point cloud aesthetic’ was quite divisive. Whilst the average score for Q10 (“This video showcases a very distinctive style, with the objects in the gallery appearing translucent and almost ‘ghostly’. Did you like this style of presentation, or would you have preferred something more solid?”) was 3 – or almost exactly ‘average’, the largest single response was ‘4’ (12/34, 35% “I found the style confusing and unappealing”), while 6 people (18%) chose ‘1’ (“I loved the style, it made the video”). Three people specifically mentioned in comments that they would have liked more ‘solid’ models, including two who had previously indicated that they liked the translucent style. Whether it is actually the *style* of the video, or the amount of data available that people are unhappy about is unclear; as we have seen in section 4.8 simply increasing point size reduces the model’s translucency and increases the

impression of a solid object, but it does not necessarily improve the appearance.

Terrestrial laser scanning with its multi-millimetre resolution is not conducive to generating solid models of objects with high frequency details.

There was a fairly even spread of familiarity with virtual worlds (Q11: “Do you have much experience with virtual worlds and environments on computers? ie, perhaps through playing games?”), with an average score of 2.6, or somewhere between “Some, but not a lot of experience” and “I have had a fair amount of experience”.

There was no correlation between age and experience with virtual worlds, and a possible, though small negative correlation between experience with virtual worlds and opinion on the style of the video. Those people who were more familiar with virtual worlds had a slightly lower opinion of the video’s style, possibly because they were used to the more solid and ‘real’ 3D meshes of videogames.

Q10 does specifically refer to a ‘ghostly’ style, and therefore could be construed as a leading question. However, on other websites commenters also used this and synonymous terms unprompted (“That is at the same time beautiful and very very ghostly.” (Jeribus, metafilter), “ghost-like images” Rhadiem, YouTube), “The 3D model is very spooky” (Damienmce, metafilter)).

Again, it is important to consider the video and its style in context. It is, after all, a memorial to something that is no longer with us, and the elegiac tones of both the incidental music and narration fit well with the ‘ghostly’ nature of the translucent point cloud aesthetic. The power of the three aspects – visuals, narration and music – to come together and evoke emotion in the viewer is summed up by one comment:

“Great. Now I’m feeling wistfulness over the non-existence of a place I never knew existed.” (Bonobothegreat, metafilter). Whether a similar style would be as appropriate for, for example, virtual tours of current exhibitions is uncertain.

4.10 Methodological Issues

One of the advantages of creating a YouTube video is that is easily embeddable on any webpage, and thus can be quickly disseminated over the web and shared via social media³⁴⁵. The downside of this, for research purposes, is that the resource is then viewable in a variety of places³⁴⁶ and thus reactions and comments are also spread amongst many sources. The link to the survey, of course, remains in only the one place. In future, it would be helpful to have the survey linked from multiple locations; particularly on the Science Museum’s YouTube channel which had a large amount of traffic, or even better, via a clickable link at the end of the video, so that the survey would be accessible no matter where the video was viewed.

We have discussed above the self-selection of the audience, and thus any conclusions drawn from the survey and to a lesser extent comments elsewhere must take into consideration that the responses are very probably not those of a random sampling of the general public, nor a random sampling of Science Museum visitors or potential visitors.

³⁴⁵ For example, as well as the SM website and YouTube channel, the video can currently be found embedded on a wide variety of sites, amongst others, the Huffington Post, Digital Arts Online, 3D Blog, LiDAR News and the BBC (URLs in appendix xxx)

³⁴⁶ Ie, YouTube, the Science Museum’s own website, b3ta.com, gizmag.com, etc.

4.11 Conclusions

We have established that terrestrial laser scanning is a viable method for recording and preserving museum spaces, although the required hardware and skillsets are unlikely to be present in most cultural heritage institutions and the cost of employing a professional scanning company may be beyond the reach of most museum budgets.

There is as yet no easy workflow between capture and dissemination – as we have seen, collecting the data is a small part of the entire process, and there are further barriers, both technological and skill-based, to producing a usable output, whether that is pre-rendered video or something more interactive. ScanLAB are professionals with much experience in producing aesthetically pleasing and dramatic outputs from raw point cloud data, but this sort of expertise is still both rare and costly. It is, however, presumed that the cost will continue to come down due to advances in both scanning hardware and software, and greater competition in the marketplace.

User research shows that interactivity is much desired by the audience, but that there will be issues concerning level of detail of the models and overall scan resolution that may make this particular capture method unsuitable for interactive applications. There are also potentially ongoing costs and support requirements incurred in providing interactive applications, and sustainability of any sort of digital resource is a very real issue^{347,348}. This can be seen in the next chapter where the user experience degraded

³⁴⁷ Maron, N., Smith, K. K., & Loy, M. (2009). Sustaining digital resources: an on-the-ground view of projects today: Ithaka case studies in sustainability. Ithaka S+ R.

³⁴⁸ Guthrie, K., Griffiths, R., & Maron, N. (2008). Sustainability and revenue models for online academic resources. An Ithaka Report.

extremely quickly once the vendor stopped supporting the online renderer.

From section 4.9.3 we see that whilst the video was generally well received, it was most popular amongst those who already had familiarity with, and a positive feeling towards, the Shipping Gallery itself. It would therefore appear that, at least in this case, the digital resource itself is not the only factor to be considered when evaluating its success; users' feelings toward the source are also relevant. This further implies the possibility that some of what Benjamin³⁴⁹ calls the physical object's 'aura' can, in fact, be transferred to or inherited by the digital object. This would support the argument from chapter 3, that the context of a digital resource is as important as the context of the original object. The comments demonstrate that the Shipping Gallery itself cannot be examined without reference to its context, over and above the bare facts of its physical properties and its contents – the individual models and exhibits it contains. The gallery's wider context includes people's memories, and not simply memories of the gallery itself; the exhibitions longevity means that it becomes a focus for emotional recollections of family and childhood. Few galleries will be quite as old as the Shipping Gallery and have had time to build up an aura of their own; be able to provide some of the powerful cross-generational memories referred to in the comments. But it is interesting to note that an exhibition, or even a physical space can accrue aura almost independently of the objects it contains, and that that aura is perhaps something worth preserving in and of itself. From the comments it is clear that the digitised version as presented has inherited some of that aura – if aura is taken to mean

³⁴⁹ Benjamin, W (ibid.)

affectual power. Whether that aura is present in the scan data – the point cloud – itself, or is a product of the visuals (including all the contingent aesthetic choices made in rendering) plus the music and narration is a valid question, but as we saw in chapter 3, it is neither uncommon nor problematic for an object's aura to be a product of the object-plus-context³⁵⁰.

Other comments refer to the gallery's context in terms of the history of shipping and its particular connection with British history, empire and the recollection (real or imagined) of a golden age entangled with notions of sea power. These, for some viewers of the video, are emotive subjects and thus the Shipping Gallery itself, and to a certain extent, the scan and then the video, become representations or carriers of a much larger significance, in much the same way that Messham-Muir's shoe (see section 3.3.2) carries the context of the holocaust. Given that we, the viewer, bring our own biographies and contexts to the exhibit, it is reasonable to assert then, that the aura of an exhibit, digital or otherwise, will be different for different people.

From this we can see that it is important not to treat the output of the laser scanning process in isolation, stripped of the context provided by the surrounding material. In the case of the Shipping Gallery output, this would include not only any accompanying explanatory text and audio content – both the music and, of course, the narration – but also the audiences' pre-existing knowledge and feelings towards the source. We have discussed the heterogeneity of cultural heritage objects, but should not disregard the

³⁵⁰ It should be noted that throughout the research we have assumed that viewers watched the video with sound, but in fact there is no way of knowing whether individual users did, in fact, experience the audio along with the visuals. This would suggest a potentially fruitful avenue for further research.

heterogeneity of the audience as well. If the purpose of the creation of the digital asset is to create something with affectual power, the question of *who* we want to ‘affect’ must be addressed.

From 4.9.4, we see that the video could potentially serve the purpose of encouraging visitors to visit the physical museum. Those already familiar with the gallery and who described themselves as regular visitors were most likely to visit after watching, with the video perhaps serving as a simple reminder of the gallery’s existence. The next most likely to visit after viewing the video were those who had never visited the gallery before, and whilst it may simply be the case that the video made them aware of something they weren’t aware of before, it could also imply what Cameron³⁵¹ calls ‘selective canonization’ - “the value of the ‘real’ increases when digitized ... owing to the resources required in the compilation of a 3D rendering”. The video’s narration, the many articles published on the project both in print and online, and, to a lesser extent, the surrounding text on the Science Museum’s website make it clear that not only is laser scanning a new and technologically sophisticated technique, but that it involved considerable effort as well; the fact that the Shipping Gallery was chosen as the recipient of this process serves to increase the importance of the gallery in the eyes of viewers, granting it a significance it may not have had previously. The creation of the copy does not diminish the original object (in the eyes of the viewer), it enhances it. The irony is that if this sort of technique becomes common, this effect will be reduced.

³⁵¹ Cameron, F, 2007, Beyond the Cult of the Replicant: Museums and Historical Digital Objects—Traditional Concerns, New Discourses, in Theorising Digital Cultural Heritage (pp49-75), Ed Cameron F & Kenderdine, S, MIT 2007

In terms of the video's presentational style, the aim is clearly not to produce a photo-realistic rendering of the gallery. While this is, in part, due to the fact that it would in any event be impossible to render a convincingly photo-realistic view from the available data, it is also a conscious decision taken by ScanLAB to use the pure point cloud aesthetic rather than attempting to 'solidify' the model or find a compromise between the fully translucent pure point cloud technique and a more realistic mode, for example a splatting technique³⁵².

It has been argued that the development of computer graphics is in some way a quest for photo-realism³⁵³, and indeed Gillings describes virtual archaeology (though it could be applied to the wider field of digitising cultural heritage) as a "relentless questing for the elusive grail of photorealism"³⁵⁴. The Shipping Gallery video shows that there is another way, and that photorealism is not the only choice when representing reality. 3D digitisation can pursue other ends and utilise other artistic means: whilst the ultimate aim – or at least one of the aims – of the project was to preserve the gallery (and indeed, something approaching an objective reality is preserved in the digital, numerical information recorded by the laser scan), this does not necessarily entail we must strive for an objectively realistic rendering. A painting may preserve a view as effectively as a photograph, and there are other valid alternatives to photo-realistic representations, as argued by Roussou and Drettakis³⁵⁵. Again, this view ties in with the

³⁵² Zwicker M et al. 2001. Surface splatting. In Proceedings of the 28th annual conference on Computer graphics and interactive techniques (SIGGRAPH '01). ACM, New York, NY, USA, 371-378.

³⁵³ F. Durand. 2002. An Invitation to Discuss Computer Depiction. ACM/Eurographics Symp. NPAR'02.

³⁵⁴ M. Gillings. 2000. Plans, Elevations and Virtual worlds: the development of techniques for the routine construction of hyperreal simulations. *Virtual Reality in Archaeology*, Archaeopress, Oxford, 2000.

³⁵⁵ Maria Roussou, George Drettakis. Photorealism and Non-Photorealism in Virtual Heritage Representation.

idea of aura and affectual power. The history of visual art has shown that a painting's (or other artwork's) ability to evoke an emotional response is not connected with the 'realism' or its objective recording of a scene or event I would argue that whilst the laser scan data of the gallery is primarily a means of preserving the original, the video output itself is a celebration, an elegy, for something lost, and as such, valued primarily for its affectual power. As such the aesthetic choices made for the video, whilst sacrificing some of the 'realism' of the scene (objects' translucency etc.), may have contributed to the ultimate success of the output.

Roussou and Drekkatis also argue that interaction can be as important, or more important, than the visual fidelity of a digital resource in creating a sense of place and physicality, and it is clear from this research that there is certainly an appetite for more interactive and user-controlled experiences. While the pre-rendered video was popular, all respondents to the survey, and many other online commenters, wanted to be able to use the data to navigate the space themselves, and even in 2013 there were already users thinking ahead to VR and Oculus Rift-type environments. As we have seen in this, and also the Courtauld Bag (Chapter 6) projects, pre-rendered outputs give the creator control over what the user sees, and therefore can compensate for issues with the available data. As capturing and rendering technologies improve the provision of fully interactive experiences on the scale of the Shipping Gallery may become possible.

A.Chalmers and D.Arnold and F. Niccolucci. First Eurographics Workshop on Graphics and Cultural Heritage (2003), 2003, Brighton, United Kingdom. Eurographics, pp.10, 2003, Proceedings of the International Symposium on Virtual Reality, Archeology and Cultural Heritage.

5 Illuminating objects

5.1 Introduction and research aims

The following chapter describes a collaboration with the Courtauld Gallery, resulting in the creation and online display of two 3D models of objects participating in the Gallery's ongoing *Illuminating Objects* programme.

The aim of the digitisation portion of the IO project, as envisaged by the Courtauld, was to offer something over and above their existing web content. The model would be considered successful if it adds utility to the website and enhances the user's experience, understanding (of the object) and/or improves engagement. These outcomes were measured using surveys attached to the website.

At the same time, this chapter examines different methodologies for creating 3D models of potentially problematic objects with challenging material properties, and the feasibility of providing interactive models via webGL that are large and detailed enough to be useful and convincing and yet small enough to provide a smooth and acceptable user experience. In particular, it compares a low cost photogrammetric method with one using an extremely expensive (and rare) laser scanner.

5.2 Illuminating Objects

This pilot programme 'aims to shine new light on unexpected objects from the decorative arts and sculpture collections, through partnerships with SOAS, King's College London, the University of Kent in Canterbury, Imperial College and University

College London.³⁵⁶ For the programme's director, Dr Alexandra Gerstein, Illuminating

Objects serves two main purposes;

"The project enables us to take objects from our collection that are rarely, or never on display, and bring them out to the public. These are objects for which we don't have a proper context, or that are usually shown in a group, but which we wanted to look at them as individual works. Equally as important, if not more important, the programme enables us to form connections and forge partnerships with other academic communities outside of our own institution. Because we chose research students, mostly PhDs and one MA, who are working in fields other than art history, this provides fresh perspectives on our objects."³⁵⁷

Post-graduate students from a variety of institutions and disciplines serve a three month internship at the Courtauld, and are invited to select an object from the gallery's stores which falls within their area of study and fits their research interests. The objects are removed from storage and placed on display in the gallery in a custom made display case³⁵⁸ designed specifically for the Illuminating Objects programme. The students spend their internships researching the object, and that research would then be disseminated to the public via two labels on the display case³⁵⁹ and on the Courtauld's website. Thus the programme aims to achieve a variety of ends: to shine a

³⁵⁶ Courtauld 2013 <http://courtauld.ac.uk/gallery/exhibitions/2013/illuminating/dish/index.shtml>

³⁵⁷ Interview with Dr Gerstein, Oct 2014

³⁵⁸ A single large case was designed with custom lighting that would be flexible enough to accommodate any object that may be chosen.

³⁵⁹ Approx. 150 words each, less if the labels included an image.

spotlight on a hitherto unseen object, whilst giving the intern experience in public engagement and communicating their research to a wide non-specialist audience.

“The students find this aspect very interesting, how to communicate ideas and structure their text in a web friendly way. They work closely with our web manager, who provides direction and gives them an idea of why it is different communicating ideas through text and images on the web as opposed to a printed medium or text in a gallery. The audience is different, and their modes of reading are different.



Figure 5.1: the lustreware bowl in the Illuminating Objects display case. Photo courtesy Court Inst. & SOAS

“Each object is approached individually – they are all treated the same insofar as they are interesting objects – but the object itself, and the amount of

published literature, more or less determines the investigation, the level of detail. More importantly is the student's discipline and how well that marries with the object. It's not an entirely academic approach, but nor is it about conjecture or creative writing; what it *isn't*, is art history, and that's where there is a learning curve, for us as much as the students. If you look at the glass, [the fourth object], compared to the dish [the second], there is a lot less detail. The website is in the process of change, in flux as well. It doesn't necessarily reflect the audience it is targeting. At the moment it is very academic, but we are a museum, and a gallery as well as a university."³⁶⁰

The creation and provision of 3d virtual models was not part of the original Illuminating Objects specification, rather it is an extension to the programme, providing an extra piece of content intended to enhance the objects' web presence, and hopefully improve engagement and understanding of the objects for web visitors.

Each object in the programme was considered on a case by case basis, and the decision whether to digitise or not made primarily on the suitability of the object for digitisation – for example, the capture of a collection of 16th -18th century filigree drinking glasses was deemed infeasible due to the inherent difficulties in scanning or otherwise modelling glass objects³⁶¹.

³⁶⁰ ibid

³⁶¹ The utility of the 3D model, and how much it might contribute to the overall web experience was also considered, though as this is a research project into the question of 3d models' potential utility, this considered to be begging the question, and indeed, the provision of questionnaires alongside the virtual models is aimed at assessing the utility or otherwise.

5.3 Object 1: Spanish lustreware bowl

The first object to be considered for digitisation was selected by Tanja Tolar, a PhD candidate at SOAS, whose research focusses on the connections between the Christian and Islamic worlds in material culture. The object is a large (approx. 47cm diameter) Spanish lustreware bowl dating from the early 16th century, and was on display at the Courtauld as part of the Illuminating Objects programme from February 6 to April 29, 2013³⁶². Unfortunately, what makes the object special, and which was the focus of Tanja's research, is also what proved to be most problematic in creating the digital model: the lustre. As Tanja explains on the website, lustreware is a technology developed in the Middle East around the 9th century,

"It is a painterly technique involving the application of compounds of copper and silver, mixed with clay or ochre, to an already fired and glazed ceramic object. After the lustre has been applied, the vessel is then fired again at lower temperatures in another, smaller kiln. Firing in this special, oxygen-reduced kiln enables the metallic compounds to bond to the surface of the glaze. Recent scientific analysis has shown this metallic layer to be extremely thin, at 0.2 microns (or 0.0002 mm). The result is a metallic sheen resembling precious objects of gold and silver, which only reveals its full iridescence in optimum lighting conditions."³⁶³

The dictionary defines lustre as 'the state or quality of shining by reflecting light' or 'a

³⁶² <http://www.courtauld.ac.uk/gallery/exhibitions/2013/illuminating/dish/>

³⁶³ Tolar, T, 2013, <http://courtauld.ac.uk/gallery/exhibitions/2013/illuminating/dish/lustre.shtml>

substance, as a coating or polish used to impart sheen or gloss³⁶⁴. Both meanings can certainly be applied to this object. What makes the dish so visually interesting, is the complexity of its interaction with light, in particular the quality and intensity of the specular highlights³⁶⁵. On the darker, reddish areas of this object, the glaze contains a high amount of copper, producing vivid magenta highlights [Figure 5.2]. There is also an iridescence reminiscent of oil on water, caused by the thin layer of glaze [Figure 5.3].



Figure 5.2: A fairly typical photo from the imaging session, showing the excessive specularity, particularly on the darker areas



Figure 5.3: photographs showing the iridescence caused by the thin glaze

The nature of this lustre presents the 3d modeller with two problems. The difficulties of capturing a shiny object with many and intense specular highlights are well

³⁶⁴ Definition from dictionary.com, <http://dictionary.reference.com/browse/luster>

³⁶⁵ Specular highlights are the mirror-like reflections experienced when light is shined on a glossy surface; the light is reflected in a single direction leading to bright spots.

known^{366, 367,368}, but as well as complicating the capture process, objects with intense visual effects create challenges in rendering as well. In the capturing case, the aim is to *minimise* the object's specularity, however, in order to at least convey the object's unique properties these highlights must be 'reintroduced' to the final rendered model. Whilst this has been a 'solved' problem for many years, with animators using ever-more sophisticated approximations to reality such as the Phong model³⁶⁹, the parameters that must necessarily be introduced into the rendering algorithm to recreate a particular material are not measured as part of the normal capturing process³⁷⁰ and thus their introduction at the rendering stage must lead to a degree of subjectivity in the finished model. Also, as a comment on current technology (a contingent, as opposed to a necessary issue), lighting models and shader techniques

³⁶⁶ Light reflecting off a shiny (non lambertian) object's surface has a diffuse component plus a specular component (a completely lambertian object only has the diffuse component, the opposite would be a mirror, which only has a specular component). The diffuse component is independent of the viewing angle, and thus the appearance of the object appears the same as the viewer moves around it (or the object moves); ie the value of the reflected colour at any point remains the same independent of viewing angle. The Specular component depends upon the angle between light source and viewer; and thus the appearance of the object will change as the light source, viewer or object moves. Since a laser scan or a photo records both the diffuse and specular components of an object, the recorded colour of a single point will differ depending on the angle at which the recording is made. With a laser scan this can cause issues recording the true diffuse colour at a particular point, and subsequently, when registering two scans taken from different angles, colour discontinuities. With photogrammetry, as proved to be the case with this bowl, pixels representing the same point in two images may show two very different colour values, interfering with the algorithms identifying them as the same physical spot.

Capturing both diffuse and specular components (and any other components that might go to specifying a complete reflectance model involves capturing a texture *function* rather than a simple texture and will be examined in more detail in another chapter. However, an object's specularity can be added to a diffuse model during the rendering process by selecting certain values representing the object's material; therefore we are interested only in capturing the purely diffuse component. I shall look in more detail at rendering specularity later in this chapter.

³⁶⁷ Mallick, SP et al, Beyond Lambert: Reconstructing Specular Surfaces Using Color, Computer Vision and Pattern Recognition, 2005, Vol 2 pp619-626

³⁶⁸ Ihrke, I et al, State of the Art in Transparent and Specular Object Reconstruction, Eurographics 2008

³⁶⁹ Strauss, PS, A realistic lighting model for computer graphics, Computer Graphics and Applications, IEEE Volume 10, Issue 6, pp 56-64 , Nov 1990

³⁷⁰ These parameters can be measured when capturing an object's BRDF (Bidirectional reflectance distribution function) or BTF (Bidirectional texture function), but this is considerably more complex than normal scanning, and I will examine these techniques in another chapter.

are both far more mature and far more common for mesh models than for point clouds.

5.3.1 Capture

The bowl was brought to UCL by Dr Gerstein and Tanja Tolar, and myself and fellow UCL PhD candidate Ali Hosseininaveh had access to the object for most of the day (approximately six hours). Our intent was to capture the bowl using two methods, laser scanning and photogrammetry. First, we scanned the object using the Nikon Metris K-Scan portable scanner, which combines a handheld laser scanning head with an optical CMM (Co-ordinate Measuring Machine) system [Figure 5.4]³⁷¹. The scanner captures geometry in high detail but doesn't provide any texture information. For the scanning it was decided to use the Nikon rather than the more accurate Arius scanner³⁷² – which also records colour – as we were unsure how either of the scanners would cope with the object's shininess and doubtful as to the Arius' ability to capture useful colour information from such a specular surface. As the virtual model was designed to be displayed online (requiring down-sampling of the scan data anyway)³⁷³, any advantages


³⁷¹ The scan head has several clusters of LEDs which are recognised by the cameras on the stationary camera bar. There are three individual cameras, each equipped with a cylindrical lens, which measure the position of the scan head along one axis. The scan head shines a laser spot on to the object, and via triangulation, measures its position relative to itself. By combining this measurement with the information from the cameras, the laser spot's position within the overall coordinate system (defined relative to the camera bar) can be determined with high accuracy. For more information see

[http://www.nikonmetrology.com/en_EU/Products/Laser-Scanning/Handheld-scanning/K-Scan-MMDx-walkaround-scanning/\(key_features\)](http://www.nikonmetrology.com/en_EU/Products/Laser-Scanning/Handheld-scanning/K-Scan-MMDx-walkaround-scanning/(key_features))

³⁷² See 5.4.1, on capturing the Miniature Bibles.

³⁷³ Using larger data sets, in the order of many millions of points, increases file sizes and causes usability issues. Particularly large data sets can also cause crashes due to hardware limitations (particularly in available gpu memory), software issues (out of memory issues in the browser) and the maximum size of arrays allowed by webGL. Therefore, our original datasets, which can often consist of 20 million points or more must be sampled.

gained from the Arius' higher accuracy and point density would be lost, while the Nikon's mobile scan head made it a considerably quicker option – about an hour for a complete scan of both sides of the bowl, compared to what would probably have been a whole day's scanning with the Arius. After the scanning, the bowl was moved to another room in which we took a series of photos which could be used to construct a photogrammetric model, or as a possible alternative, to texture the model³⁷⁴ provided by the Nikon scanner.



*Figure 5.4: Left: The Nikon scanner in action;
Above: the model during capture*

5.3.1.1 *Laser scanning*

The scanner used, a Nikon Metris K-600, measures points with an accuracy of approximately $60\mu\text{m}$ ³⁷⁵ in a volume between 1.5 and 3.5 metres from the camera bar.

³⁷⁴ By draping or projecting colour photos onto a non-coloured 3d model

³⁷⁵ Nikon Metris manual, http://nees.umn.edu/facilities/files/Krypton/Kx00_manual_EN.PDF

The scanning procedure took approximately one and a half hours, though prior to the bowl's arrival, a similar amount of time was spent setting up and calibrating the equipment³⁷⁶. The upper surface of the bowl was scanned, then the object reversed and the underside captured. In both scans we also attempted to capture portions of the rim including a section of the opposite side to ensure that both models shared some common points. These common points would then allow us to align and register the two surfaces.




Figure 5.5: The scanning errors can be seen clearly on the inner rim of the bowl

The procedure appeared to go smoothly with the scanner coping well with the object's shine; there did not appear to be a large amount of noise, holes or other unwanted artefacts in the model. However, on subsequent close inspection of the data, it was clear that there were, in fact, significant errors. Many of the scan stripes were grossly misaligned (in the order of 5mm), as can be seen in a close up of the bowl's inner rim [Figure 5.5], clearly showing three separate surfaces. As the individual scan stripes were not preserved in the data, it was impossible to correct this misalignment in post-processing. As this sort of scanning artefact later appeared in subsequent scans conducted with the Nikon, it is likely that this was a hardware fault in

³⁷⁶ For the Nikon scanner, the calibration process consists of making a series of measurements of calibration objects with known dimensions (a long bar and a smaller cube) in various orientations.

the scanning equipment, rather than a result of the object's particular properties (ie, shininess).

5.3.1.2 Photogrammetry

The photography took place in a room with a partial blackout. Lighting was provided by two professional quality flashes in softboxes, synced to the camera and aimed at the ceiling in order to provide as diffuse a light as possible and to minimise the specular highlights on the object. The camera used was a Nikon D700 with a 35mm lens³⁷⁷, and once an optimal field of view was chosen and the camera focussed, the focus ring was fixed with tape to ensure identical settings were used in every photo. After the shoot the camera was calibrated using these same settings; I shall talk more about camera calibration (and photogrammetry in general) in 6.

The distance from lens to object was measured with string, which was then used to help position the camera for subsequent shots; for each camera position the object was brought into focus by physically moving the camera. The first image included a ColorChecker Colour Rendition Chart³⁷⁸ to allow for later colour calibration of the images.

The bowl was placed on a black board, and scale bars were placed either side. As the bowl is fairly symmetrical and has a broadly repeating pattern, extra random background texture³⁷⁹ was provided by placing printed materials around the bowl

³⁷⁷ Shutter speed was set to 1/125s, aperture F18 and film speed ISO200

³⁷⁸ C. S. McCamy, H. Marcus, and J. G. Davidson (1976). "A Color-Rendition Chart". Journal of Applied Photographic Engineering 2(3). 95–99.

³⁷⁹ See: http://www.artec3d.com/news/Scanning+objects+with+repetitive+geometry+and+no+texture_31022

[Figure 5.6]. With hindsight we would have replaced the black board with more random textures as the lack of detail led to ambiguity in the finished model between the edge (rim) of the bowl and the black background.



Figure 5.6: One of the images from the photography session. Note the use of printed matter as texture, photogrammetry targets on the backing board, and the specularity of the object

The bowl was then imaged in two separate rings, the first fairly oblique, at approximately 30° to the horizontal, the second with as much elevation as the tripod could provide, at approx. 70° . These parameters were chosen as a compromise between photogrammetric best practises (for example, those published by ISPRS³⁸⁰),

³⁸⁰ Tips for the effective use of close range digital photogrammetry for the Earth sciences, ISPRS - Commission V - Close-Range Sensing: Analysis and Applications Working Group V / 6 - Close range morphological measurement for the earth sciences, 2008-2012

and what is practical in the circumstances, taking into account the location, equipment (tripods etc.) and time available. About 10 photos per ring were taken at equal intervals around the object. The images' overlap was largely irrelevant as all photos included nearly the whole bowl, whilst ten photos per ring was enough to ensure that, with the limited occlusions on the object, each point on the bowl appeared in multiple images. The bowl was then reversed and the whole process repeated. Again, with hindsight, we could have made more effort to image the rim of the bowl in order to provide some sort of common features in both models (the front and reverse of the object). This would have allowed us to align the two models - although the difficulty of capturing both sides of the bowl in a single image may have made this impossible anyway³⁸¹. As well as the photogrammetric images, a set of photos were taken as orthogonally to the bowl as possible in order to provide images suitable for texturing a scanned model.

5.3.2 Processing

After the laser scan data was examined and found to be faulty, due to what we assume were hardware problems (see Figure 5.5), the decision was made not to attempt to rescan the bowl owing to the difficulties of accessing object. Instead, we opted to concentrate on creating a photogrammetrical model. The methodology was as follows:

- After capture, the camera settings used for the imaging were left unchanged

³⁸¹ As the bowl had to remain flat (horizontal), lighting both sides simultaneously would have been an issue, as would the depth of field, not to mention the difficulty of capturing detail when the surface was practically perpendicular to the focal plane.

(with tape securing the focussing ring to prevent accidental movement) and the camera was then calibrated using VMS³⁸² software. I shall talk more about camera calibration in section 6.5.1.1.

- Using the photo with the ColorChecker Card and X-Rite colour management software, a colour profile for the images was acquired. All the images were then colour corrected using Adobe Lightroom³⁸³.
- Using the camera calibration file from (1) and LDC³⁸⁴ software, the images were undistorted.³⁸⁵
- We used Bundler³⁸⁶, a free Structure-from-Motion software package to reconstruct the camera positions and create a sparse point cloud.
- The output from bundler was then used as an input for PMVS2³⁸⁷, to create two dense point clouds per object (the two objects being the front and reverse of the bowl).
- Apart from 'level', we used default settings in PMVS2. The first model was created with the images at full resolution (level 0), the second (level 1) with the images sampled at half resolution, ie, $\frac{1}{4}$ number of pixels³⁸⁸. The second model,

³⁸² Vision Measurement System from Geometric Software, available from <http://www.geomsoft.com/>

³⁸³ The Macbeth board contains patches with known colour values, and made of a material that reflects light consistently under a wide range of lighting conditions. By adjusting the colour profile of the photos so that the colour patches in the image match the expected values, any colour distortion due to the camera and the lighting conditions can be corrected. I shall deal with this process in more detail in a later chapter.

³⁸⁴ Lens Distortion Correction from UCL Geomatics Software, <http://eprints.ucl.ac.uk/43957/>

³⁸⁵ These distortions can be caused by defects in the camera lenses or other physical systems in the camera (such as the sensor), and I will explain this procedure in more detail 6.

³⁸⁶ For more information on Bundler, see: <http://www.cs.cornell.edu/~snavely/bundler/>

³⁸⁷ For more information on PMVS2, see: <http://www.di.ens.fr/pmvs/documentation.html>

³⁸⁸ This was accomplished by changing the 'level' parameter in the PMVS option file before processing. Level '0' specifies that the full resolution images should be used for computation, while level '1' indicates that the resolution should be halved.

using the subsampled images, produces a sparser point cloud, but in some respects this was a more robust model, suffering from a significantly smaller amount of noise than the denser model []. It is unclear why this should be the case, it may be that the averaging of pixels carried out when the images are down-sampled reduces the amount of noise in the photographs themselves, the averaging improves feature matching between pairs of images or improves the generation of the dense point cloud. This sparser model proved very useful for processing later on.



Figure 5.7: Left, the level 1 model; right, level 0. The increased point density in the second model is clear.

Once the point clouds had been created, the front and reverse of the bowl were processed separately. Although the same issues were found on both sides, and the same techniques used to fix them, in this description I shall concentrate on the top of the bowl, as this was – due to its shininess and the predominance of the particularly troublesome red copper glaze – far more problematic.




Figure 5.8: A cloud of points can be seen detached from the surface of the bowl

The two point clouds were imported into Pointstream 3DImageSuite software for processing; shows the differences between the level 0 and level 1 models. As can be seen, the level 0 model provides a dense point cloud³⁸⁹, while the level 1 model, on its own, is too sparse to be used as an effective model. On closer examination, it can be seen that the dense model suffers from far more noise, holes and unwanted artefacts than the sparser model. It is difficult to clearly illustrate the extent of the noise in a two dimensional image, but an impression may be gained from . This shows a detail of the inner rim of the bowl, with a cloud of points both above and below the nominal surface. There were similar issues in all the areas which experienced high levels of specularity in the photos, predominantly those areas with the copper coloured glaze. While Pointstream has hole-filling capabilities, and can interpolate both colour and

³⁸⁹ The full level 0 model, including the background, and before any processing, contained 4.2M points. The same level 1 model approx 1.2M.

geometry information to fill gaps in the data³⁹⁰, in some areas of the model the noise was so problematic that once the 'bad' points were erased, the resulting holes were too large for the automatic hole-filling algorithms to work. However, whilst sparse, the level 1 model suffered from far less noise (Figure 5.10), and thus provided enough data to act as a 'scaffold' for the algorithms to operate on. Specifically, whilst there were fewer points than the dense model, and point density dropped even further in the problematic areas, the points that were there exhibited less 'noisy behaviour', lying closer to the plane of the surface, thus allowing the reconstruction of that surface.




Figure 5.9: The left image shows an area of high specularity in the level 1 model, showing sparse but relatively clean data. The second image shows the same area in the level 0 model, showing dense but highly noisy data (again, the level of noise is hard to discern in a 2D image)

Large holes were filled by first using a low grid setting, and gradually increasing the density until the hole fill could add a surface with the same density as the rest of the model. The underlying grids could then be erased leaving a single hole-filling object³⁹¹.

³⁹⁰ UCL Museums and Collections, 2008. E-Curator Project: 3D colour scans for remote object identification and assessment. www.museums.ucl.ac.uk/research/ ecurator/

³⁹¹ Pointstream has since replaced this grid-filling option with a newer algorithm

Once the holes were filled, subsequent corrections were made on a case-by-case basis, depending on the appearance of the interpolated points. Often the new data inherited the colour of the specular highlights from the 'bad' points and had to be corrected by sampling a texture from a 'good' area and painting it over the new areas. Smoothing filters – point, colour and normal – were used where appropriate, for example where the hole-filling had resulted in an excessively bumpy surface.




Figure 5.10: Top left, a particularly noisy area in the level 0 model; bottom left, the same area in the level 1 model. Top right, the area after bad data has been erased, leaving a hole too large for the hole filling algorithms. Bottom right shows the two models combined, the points from the level 1 model allowing the hole to be filled

Once the two models of the bowl, front and back, had been processed, they needed to be scaled and registered. A measurement was made using the scale bars in the original photogrammetric models and a scale ratio of the two models calculated (note this is a

dimensionless ratio and has no relation to any actual physical measurement of the bowl) and used to scale the model of the reverse to the front.

Registering the front and back models proved difficult, as the models had no shared points. In fact, the edge of the bowl was a problematic area in general [Figure 5.11] the



Figure 5.11: Detail of the bowl's rim, even after cleaning shows a ragged and noisy edge

lack of texture on the surrounding card led to the bowl merging with the black background leaving a 'ragged edge' with no clearly discernible rim. As mentioned earlier, photos which included both surfaces would have helped, though may have been impractical; alternatively we may have been able to take a 'clean' section of the laser scan data which

included points on both sides of the bowl, and then register each of the point clouds to that. Time constraints prevented investigating this method however. Instead, the two models were aligned by hand. An attempt at registering using Pointstream's ICP (Iterative Closest Point)³⁹² algorithm was made, by choosing equivalent points on both sides (for example, the two holes through the bowl near the rim, and the centre point) resulted in the two bowls being aligned 'back to back'. This at least meant that the models were aligned along a central (z-)axis, and thus one of the models could be

³⁹² see Bi,Z.M. and Wang,L, Advances in 3D data acquisition and processing for industrial applications, Robotics and Computer-Integrated Manufacturing, Volume 26, Issue 5, October 2010, pp403-413 and Paul J. Besl, Neil D. McKay, "A Method for Registration of 3-D Shapes,"IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 14, No. 2

flipped 180° around one of the other axes and moved along z until the thickness of the bowl was reproduced. This made the process a deal simpler than trying to move the objects 'freehand' with all six degrees of freedom. Again, this methodology involved a compromise between accuracy and time, and is not what one would consider 'best practice'. However, bearing in mind the model's intended purpose, the results were judged acceptable, in that the model looked like a convincing three dimensional solid object under the sort of viewing conditions (primarily zoom level) available on the online viewer.

5.3.3 Rendering

To render the model online we used Pointstream's web³⁹³ viewer, a proprietary piece of software developed from the open source XBPS (Cross-Browser Pointstream viewer)³⁹⁴. The viewer had many useful features, some of which are inherent in any html5/WebGL application, notably that it allows the model to be viewed natively in any WebGL enabled browser³⁹⁵ without requiring a download or plug-in. This particular viewer also enables easy export of the model from the Pointstream software (both the data file and the html required to view it), the ability to easily preview the model in the browser and has options to adjust the material properties of the object for rendering. The use of a proprietary (.pjs, Pointstream JavaScript) file format allows good compression of the data (the .pjs file for the bowl comes in at 16Mb vs 110Mb for the

³⁹³ See http://www.arius3d.com/pointstream/ps_webgl.html

³⁹⁴ Available from <https://asalga.wordpress.com/category/xb-pointstream/>

³⁹⁵ At the time of this project, just over 50% of users (<http://caniuse.com/webgl>) had browsers which either support or partially support WebGL applications, the figure today is closer to 80%

same point cloud exported in a generic ascii format) requiring smaller bandwidth and creating a smoother user experience. The model is also streamed by the viewer, so the object begins to appear in the browser – and is interactive – as soon as the page loads, giving the user instant visual feedback that something is happening. Another feature of Pointstream.js is that the model is constantly rotating. This is to make clear to the user that what they are seeing on screen is an actual 3D model and not simply a 2D image³⁹⁶.

The disadvantages in using this product arise from the non-open source nature of the code, which means the creator has no control over the viewer's camera (camera here referring to the user's viewpoint) behaviour³⁹⁷, and no access to the shaders³⁹⁸ (they can be viewed using a browser plug in such as WebGL Inspector, but not altered).

Dealing with proprietary software also leaves you reliant on the fortunes of the providing company; if – as was the case with the makers of Pointstream – the company stops supporting their product, changes in browsers can quickly render the software inoperable. In the case of Pointstream.js, several months after the models were initially uploaded, the viewer stopped working in Firefox, then newer versions of Chrome – by which time the company had stopped updating it. Not long after that, they ceased hosting the JavaScript files, stopping the player completely.³⁹⁹

³⁹⁶ From personal communication with Anu Rostogi, one of the developers

³⁹⁷ For example, it would be useful to control the level of zoom given to the user, to prevent them zooming too far into a model that doesn't have the point density or detail to accommodate it. Similarly, objects in the viewer are perpetually revolving, the justification being that the movement reveals the 3d nature of the model – which may not have been apparent with a static image - without any interaction from the user, but it can be distracting.

³⁹⁸ Shaders are programmes which convert the raw data, for example the points' xyz coordinates and rgb values, into pixels displayed on the screen. They also define how lighting and other effects are applied.

³⁹⁹ See section 3.2.6 for some thoughts on sustainability for digital resources

The method for preparing a model for online viewing is as follows:

- The original model, after processing, contained approximately 5 million points. This was sampled down using Pointstream's unique filter option to a model with just over .5 million points. When exported to .pjs, this produces a file of 16.5 Mb; a good compromise between the amount of detail in the model and the time it takes for the model to stream⁴⁰⁰.
- A point size (the size of the splat⁴⁰¹ rendered for each point in the model) must be selected. The online rendering is very different to the view in the Pointstream software, so this must be an iterative process of choosing a point size, exporting the model and testing it in the browser. Whilst the shader does appear to use some form of attenuation⁴⁰², this is not perfect, and thus a balance must be struck between the object's appearance at the initial 'zoomed out' view and that when the user zooms in (again, we have no control over how far the user can zoom in the Pointstream viewer). Selecting too small a point size will lead to individual points becoming visible, and the model becoming translucent under zooming; too large and the model will have a blurred, pixelated look when fully zoomed out.

⁴⁰⁰In effect, this loss of data is only really be apparent to the user on high zoom levels. Generally, at the sort of distances allowed by the online viewer, the model's resolution is greater than that of the computer display.

⁴⁰¹A 'splat' is simply the artefact rendered onscreen for each point in the pointcloud; it can be a single pixel a square, circle etc. of any size or a more complicated entity; ie a surfel, see Pfister et al; Surfels: Surface elements as rendering primitives, Proc. ACM Siggraph 2000, 2000, pp 335-342

⁴⁰²The Pointstream.js shaders use an attenuation value so that splats closer to the viewer are rendered larger than those far away. This ensures that, as far as possible, the user sees a continuous surface as they zoom in on an object, even though, effectively, they are seeing less points per a given area on their screen.

- Within Pointstream, the following material properties can be adjusted for the model; shininess, which affects the sharpness of specular highlights (a high value produces smaller, focussed highlights), specular intensity, which affects the brightness of the highlights, and specular colour, which changes the colour of the highlights⁴⁰³. As above, the effects of the settings in Pointstream and in the online viewer are different, so an iterative process is again required, until a suitable aesthetic effect is obtained. These values⁴⁰⁴ can, however, be edited in the raw html exported by the software, which can save a significant amount of time.
- To create a more accurate rendering, a second version of the model was created. Pointstream's colour selection tools were used to split the bowl into two separate point clouds, one containing all the points for the yellow 'ground' and another containing just the points corresponding to the areas containing the metallic copper pigments. Since the viewer can render multiple point clouds simultaneously, with different material properties, the model featuring the darker areas with the copper glaze could be rendered with magenta highlights while the yellow areas, without the metallic glaze, would show simple white highlights. The result is fairly subtle, and did not make a huge difference to the final appearance, though the technique would prove useful for the miniature

⁴⁰³ For more on material properties see: <http://www.glprogramming.com/red/chapter05.html#name6>. Also from that chapter: "Because the interaction between an object's material surface and incident light is complex, specifying material properties so that an object has a certain desired appearance is an art." Indeed.

⁴⁰⁴ Only two values are editable in the html file, shininess and spec. colour (as an rgb triplet). This value is simply the specified specular colour multiplied by the specular intensity as a percentage, so a white specular highlight (255,255,255) plus a 50% intensity results in a final specular colour of (128,128,128).

bible model, the next object in the Illuminating Objects programme.

5.4 Object 2: Miniature Bibles

The second object(s) considered for 3d capture, a pair of miniature German picture bibles, ('Dess Alten Testaments Mittler' and 'Dess Neuen Testaments Mittler'⁴⁰⁵), were selected by Josephine Neil, a PhD candidate in Theology and the Arts, at King's College London. The books were produced in Augsburg in the late seventeenth century, likely intended as portable objects for private contemplation⁴⁰⁶. The books are bound in leather with detailing in silver on both front and back covers, and a small amount of gold leaf on the spines. Inside, each bible contains more than 300 pages, each containing an engraving illustrating a biblical story.




Figure 5.12: Photographs of one of the bibles; Left, the front cover, Above, the frontispiece showing an example of the engravings present throughout the volume. Images courtesy Josephine Neil and the Courtauld Inst.

The books were presented on the website with an interactive "pseudo-3d" page-turning mechanism⁴⁰⁷ which allows visitors to flick through a selection of the bible's

⁴⁰⁵ Authors Johanna Christiana Kusel and Maria Magdalena Kusel; published Augsburg c. 1700; <http://www.worldcat.org/title/dess-alten-testaments-mittler-dess-neuen-testaments-mittler/oclc/604301726>

⁴⁰⁶ <http://courtauld.ac.uk/gallery/exhibitions/2013/illuminating/bible/furtherinfo.shtml>

⁴⁰⁷ <http://www.courtauld.ac.uk/gallery/exhibitions/2013/illuminating/bible/pageturning/index.shtml> The page turning mechanism requires Flash

pages with the keyboard or by dragging pages with the mouse. There was some concern that having a 3d model as well as the page-turning feature may be 'a gimmick too many', however, according to feedback from both users and Dr Gerstein, the combination of the 'exterior' 3d view and the ability to view inside the books complimented each other well.

5.4.1 Capture

The decision was made to only capture one of the pair of books, as whilst there are subtle manufacturing differences to the books' bindings, the design, subject matter and engravings on the covers, the two are close enough that for our purposes they can be considered identical. The decision was made to use the Arius3D Foundation Model 150⁴⁰⁸ colour scanner for capture [Figure 5.13]. The books' small size (approx. 35x40x20mm) and relatively simple geometry, coupled with the Arius' ability to capture accurate geometry and texture information simultaneously meant that the scanners relatively slow speed⁴⁰⁹ would not be an issue.

The Arius scanner uses a coordinate measuring machine to move the scan head in pre-programmed paths, whilst a rotating mirror sweeps a laser point in lines across an object's surface. A sensor in the scan head measures the spot's position and triangulation is used to calculate the point's coordinates in three dimensional space. At

⁴⁰⁸ Hess, M., Millar, F. S., Ong, Y. H., MacDonald, S., Robson, S., Brown, I., & Were, G. (2008). E-Curator: 3D Colour Scans for Object Assessment.

⁴⁰⁹ The Arius captures approximately 3000 points per second, which implies that an object with points in the order of 2,000,000 could be completely scanned in just over 11 minutes. The actual time taken depends on the geometry of the object – ie, how many individual scan stripes must be made, the amount of overlap between individual scans, and of course the absolute size of the object. Needless to say, 11 minutes underestimates the required time by at least an order of magnitude.

its maximum resolution the scanner records one line every 100 microns, producing a point density of 100 points per square millimetre. The accuracy as measured along the z (depth) axis is better than 25 microns.⁴¹⁰




Figure 5.13: The bible underneath the Arrius scanner

The laser is made up of three individual beams with wavelengths in the red, green and blue parts of the spectrum. By measuring the reflected strength of each beam an accurate measure of the object's colour can be recorded. A calibrated white cube made of a 100% diffuse material, Spectralon, is included in each scanning sweep, ensuring that the colour information is measured accurately independent of ambient conditions.⁴¹¹

The books were brought to the scanner

at UCL whilst *en route* from the Courtauld to The British Library, and thus we had access to the objects for just two and a half hours. While this was enough time to scan each surface and capture enough data to create a complete model – with certain

⁴¹⁰ ECurator: Aims and objectives/ 3D colour laser scanning:
<https://www.ucl.ac.uk/museums/research/ecurator/aims>

⁴¹¹ Hess M & Robson S, 3D Colour Imaging For Cultural Heritage Artefacts, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol XXXVIII, Part 5, Commission V Symposium, Newcastle Upon Tyne, UK, 2010

caveats which I shall consider later – it is probably the absolute minimum needed.

5.4.1.1 Methodology

We used the scanner's highest resolution, recording one scan line every 100 microns.

The book was propped up at a slight angle for scanning to help minimise specular reflection; having the book's surface – particularly the shiny silver on the cover – at 90 degrees to the laser would increase specularity⁴¹². However, considering the curved nature of the scan line due to the rotating mirror used to sweep the laser, it is inevitable that some point on the surface will be tangential to the laser. Setting the object up at a steeper angle will lead to the opposite problem – a lack of reflected light due to the laser hitting the surface at grazing angles. As ever, compromise is key, though in the limited time available, we were unable to experiment to achieve the best results. And again, as ever, lack of time spent in the capturing process led to problems, and subsequently increased time spent, in the processing stage.

The book was scanned in 12 positions, which can be thought of as roughly equivalent to scanning each of the six faces of a cube from two directions. This was considered the minimum number of scans required to get full coverage of all six faces whilst minimising occlusions; more would have been preferable but not possible in the time available. A total of 33 scans were made resulting in a point cloud (pre-processing) of 2.8 million points.

⁴¹² When the laser hits a shiny surface at a right angle, the amount of light reflected can overload the sensor, and the model suffers from 'burn outs'. These appear as holes in the model, often with a cloud of bright/wrongly coloured points floating some way from the surface.

5.4.1.2 Issues

After the scans were completed and roughly aligned, the following issues were noted.

The silverwork on the covers showed some specularity, with the scanner recording

different colour information depending on the angle at which the book was scanned.

This was most evident on the front cover, which was scanned in multiple stripes, and

led to extreme colour discontinuities. (Figure 5.14). This effect was much less

pronounced on the back cover, which was captured in a single scan. These scanning

artefacts led to noticeable issues with the final model, and the implication is that given

more time, we could indeed have achieved better results.



Figure 5.14: From left, a scan taken at 45°, one at 0°, and the result when the two are registered.

The leather areas, on the covers but particularly on the spine, suffered the opposite problem to the shiny metal areas. Whilst, in places where it had been rubbed smooth, the leather also suffered from slight specularity, in many areas the dark, matte surface did not reflect enough light for the sensor to capture, leading to a large number of holes in the scans (Figure 5.16: A detail of the bible's spine, showing some specularity but also large areas where no data was recorded). These weren't noticed at the time of

scanning for the simple reason that the scanning software displays the scans on a black background, so holes in what was effectively a black surface didn't show up. Errors like this could be easily avoided in future by simply examining the model on a contrasting background during the scanning process.



Figure 5.15: An early model of the bible showing the holes caused by the clasps' occlusions



Figure 5.16: A detail of the bible's spine, showing some specularity but also large areas where no data was recorded

An expected – but potentially avoidable – problem were the occlusions around and behind the book's clasps. These were mitigated as far as possible by scanning the area from a variety of angles, but holes were nevertheless inevitable (Figure 5.15). With hindsight, it would have been a fairly simple matter to have opened the clasps to scan the area behind, though it is also highly likely that releasing the clasps keeping the book 'squeezed' shut would change the object's geometry enough to cause problems. The clasps caused another avoidable issue with the geometry. When closed, they

protrude slightly beyond the cover, so that when the book is resting on one side, it is actually resting on the edges of the clasps. This led to small changes in shape depending on whether the book was resting on the front or back, and this in turn caused some problems registering the scans, particularly the clasps themselves. This issue could perhaps have been avoided by resting the book's cover on something so that the clasps were not in contact with any surface.

All these problems could have been solved, or at least reduced, with more time. One lesson learned from this (and other) scanning processes is that, where possible, the object should be scanned over more than one session. This would allow for an initial rough alignment of scans, hopefully highlighting any issues and problematic areas which can then be addressed in a subsequent scanning session. Obviously this is not always possible; if not, then carrying out rough alignments as you go along (again, not always possible), and carefully examining the scans as they are made may catch problems as they occur.

It has also been pointed out that test scans could have been carried out on a similar object with similar material properties, hopefully highlighting potential problems and informing the subsequent scanning sessions with the real object. Whilst in the case of the miniature bibles, sourcing a similar object would have been difficult, a combination of objects could potentially be used, each representing a different aspect. For example, tests could be conducted on a book of any size, but with similar clasps, or books with dark leather spines, metal covers or gilded pages. These tests may have helped save time in capturing the actual bibles, and produce better final results.

5.4.2 Processing

Like the bowl, the scans were processed in a beta version of the Pointstream 3DImageSuite software. Whilst the particular filters and associated settings are necessarily specific to this software, they have been included as they are similar to, and representative of, the types of process that may be carried out in other point cloud processing applications.

5.4.2.1 Filtering

The following filters, all using default settings, were applied to each of the raw scans:

Select high incident angle points (incident angle threshold 50.0°): this removes points captured when the angle between the laser and the surface to be measured is too high. The high angle results in less light being reflected back to the sensor⁴¹³ and an elliptical distortion in the shape of the laser spot and therefore less reliable measurements for these points.

Select dark edge points (Edge distance tolerance 0.4; RGB value 100,100,100). When a spot is measured on the very edge of an object, ie, where the spot is overlapping the edge and is ‘half-on, half-off’ the measured surface, the smaller effective area of the scan spot results in a correspondingly reduced amount of light returned to the sensor, and subsequently the spot will be measured as darker than its actual value. This filter removes those points.

Select isolated points (distance 0.4, cluster size 10): This filter aims to remove noise,

⁴¹³ When the laser is perpendicular to the surface, the angle of incidence would be 0° and the reflection would be strongest, 90° would be when the laser is tangential to the surface normal and would result in no light reflected back to the sensor

and assumes that individual (or small clusters of) points disconnected from the surrounding surfaces must be scan artefacts.

All these filters create a selection of points, and while you have the option of deleting them immediately, it is preferable to create a new scan object containing them, and only then delete them from the original scans. If on completion of the model, there are areas missing data, then proceeding on the assumption that 'bad' points are better than no points at all it is sometimes possible to fill or partially fill the holes with this filtered data by selectively 'un-erasing' points.

5.4.2.2 Registration

The software measures points' positions within the coordinate system provided by the CMM, so scans conducted without moving the object – ie, individual stripes along the object, or scans conducted with the scan head at different angles – are already aligned. Therefore, registration is generally a process of registering clusters of scans together.

Before this it can be useful to use Pointstream's slow 'A1' 'align with every point' ICP algorithm to register the scans within clusters. This ensures an extra degree of accuracy in the registration and reduces the potential for errors introduced by incorrect calibration of the scanner⁴¹⁴ or for any subtle (unintended) movement of the object between scans. This step is not essential, but the time it takes is trivial compared to the total processing time.⁴¹⁵

⁴¹⁴ The scan head can be aligned at a variety of angles, and must be calibrated separately for each one. For example, the UCL Arius scanner has been calibrated to be used at 30, 45, 60 and 90 degrees to the vertical.

⁴¹⁵ In the case of the bibles the process took approximately 2-3 minutes per scan (or cluster of scans), though this can be considerably longer depending on the particular scans being registered and the hardware being used.

Initial registration for non-pre-registered scans involves manually picking common points on the two scans. At least four pairs of points is recommended, and preferably they should not all lie on or close to a single plane and be fairly well spread over the scan's area. Once common points have been selected, the software uses ICP (iterative closest point) algorithms to align the two point clouds⁴¹⁶ by globally minimising the total RMS⁴¹⁷ distance between points as it tests rotations and translations (in all six degrees of freedom).

On completion of the registration, statistics are presented including an RMS for the alignment. A large error here indicates an issue with the registration, and the process should be repeated with more and/or different marker points selected. While the RMS value is useful, a visual inspection of the registration is essential, as a 'good' RMS, in the order of .003 or so, does not automatically entail a good registration. Similarly, small errors can be propagated so that even though no single RMS between pairs of scans indicates an error, overall, scans can still end up with very poor alignments. For example, imagine a sphere with scans starting at one longitudinal stripe and proceeding around the circumference. Each scan may be aligned well with the previous one, and show a very low RMS, but if these small errors are all in the same direction, they can propagate around the sphere so that when the final scan is aligned, there may be a large gap between it and the first – despite none of the individual registrations

⁴¹⁶ Paul J. Besl, Neil D. McKay, "A Method for Registration of 3-D Shapes," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 14, No. 2

⁴¹⁷ Root Mean Square – in this case a measure of the average distance between 'the same' points in the two scans being aligned

showing an unacceptable error. With a simple shape and relatively small number of scans like this bible, this is not a major problem, but in more complex objects with complex geometry and a large number of scans, this effect can lead to large discontinuities between scans. Every registration should be followed by a detailed and close visual inspection, it is not enough simply to rely on the statistics.

5.4.2.3 Overlapping scans

Once the scans have been registered, overlapping areas must be processed. Many, (if not all, in the case of a small object like the bible) areas of the model will be covered by more than one scan, and this is undesirable for a variety of reasons. One simple reason to avoid overlaps is that redundant points lead to a larger model, and a larger model entails more data, slower processing and a larger file size. A more important issue, especially on a model with shiny areas like the bible's metal cover, is colour consistency. As mentioned above, an area which exhibits specularity will give different colour values depending on the angle of the scan. An extreme version of this is shown in Figure 5.14, where the same area has been scanned at an angle of 0° and 45° .⁴¹⁸ In the situation where two scans have measured different colours, the simplest case involves deciding (subjectively, but with reference to the real object) which scan has the most accurate colour, or at least most consistent, and deleting the points belonging

⁴¹⁸ This of course brings into the question the whole concept that the scanner is, in fact, recording an objective colour value for a point on an object. If a small change in scanning angle can entail a large change in colour value, then what basis do we have for asserting that this is a real measurement? This issue applies more to documentation and digitisation, where the model is held to represent or record some underlying objective 'truth'; the creation of models for visual effect, as in this case, can afford some subjectivity, as the recorded colour of surfaces with complex reflectance properties is merely the starting point in a rendering pipeline designed to recreate an impression of the original object.

to the other scans. If for any reason this is not feasible (ie, areas where none of the overlapping scans have ‘good’ colour, or where there are unavoidable discontinuities between individual scan stripes etc.) colour can be sampled from one scan and painted on to the points from another, though this works best with areas of relatively flat colour and/or random textures. Along the edges of and between scans, colour can be blended using a smoothing filter. These techniques can cause a certain amount of texture blurring, however, as can be seen in Figure 5.17.



Figure 5.17: Comparison of the two covers. On the left, the cover was captured in a single scan, the cover on the right was a result of combining the two scans in Figure 5.14. The resulting blurring of the texture can clearly be seen.

5.4.2.4 Hole filling

As well as colour issues caused by specularity, there were three types of hole⁴¹⁹ in the scan data. The first were simply occlusions, often inevitable in any scanning campaign,

⁴¹⁹ Chalmoviansky and Juttler refer to two types of holes; those caused by occlusions (the geometrical properties of the object) and those caused by material properties of the object. I treat those caused by specularity (too much light reflected) and dark surfaces (too little light) as separate cases, as each can cause different problems in processing. See: Chalmoviansky P and Juttler B, Filling Holes In Point Clouds, in Wilson M & Martin R, Mathematics of Surfaces, Lecture Notes in Computer Science, Vol 2768, 2003, pp 196-212

but exacerbated here by the limited time and subsequent limited number of scans and scan angles conducted. These holes range from the very small, for instance those found at the interface between the leather and silverwork on the covers, and which were easily filled using the basic hole-filling algorithms, to more problematic holes caused by the clasps occluding the pages (Figure 5.15). These larger areas were filled using the same method described for the bowl earlier, though as can be seen, (Figure 5.18) the end result was not entirely successful due to the particularly high frequency and regular detail of the page edges: interpolating texture data is more successful when the textures are random.

Holes can also arise due to the object's material properties. Where the laser was (or was close to) perpendicular to the metalwork on the cover, the amount of light reflected can overpower the sensor leading to 'burn outs', areas whose appearance is similar to an over-exposed spot in a photograph. These areas were relatively small and easy to fill. More problematic, and caused by exactly the opposite effect, were areas of the leather work that were too dark and consequently did not reflect enough light (Figure 5.16).

In contrast to the burn-out problem, the leather gave better results when the laser was



Figure 5.18: Less than perfect hole filling due to the high frequency & regular texture of the bible's pages. It should be noted of course, that this is less obvious in the finished model, as the clasps also tend to occlude the user's view of the area.

orthogonal to the surface (maximising the amount of light reflected back to the sensor), however, the rough, bumpy texture (combined with the gentle curvature of the spine) meant that a single scan angle could only capture limited areas of the spine, and, as mentioned above, the extent of these holes wasn't discovered until the post-processing stage. In this case, the fairly homogenous, randomly bumpy texture of

the leather was conducive to 'invisible mending' with interpolated data.

In terms of processing the holes and filling in missing data, those resulting from specularity are often more problematic, as the holes are often accompanied by clouds of 'noise' – clusters of points above or below the plane of the object, often with a false colour much brighter than the surrounding points. These require careful cleaning before application of the automatic hole filling algorithms as points in the wrong place/with the wrong colour can throw off the hole filling procedure.

5.4.3 Rendering

Again, the Pointstream html5 viewer was used. The bible model was split into two point clouds, one containing the points representing the leather portions, the other the

metal and gold leaf edges of the pages^{420 421}. The clouds were selected by manually selecting (painting) the leather areas and creating a new scan object from the selected points. As with the bowl, values⁴²² for the bible's spot size and material properties were selected through an iterative process.




Figure 5.19: Left, splitting the bible by manually painting the leather points; Right, a portion of the resulting 'metal' point cloud

5.5 Analysis

An initial assessment of the contribution of the 3D models to the Illuminating Objects programme is one of cautious optimism. Certainly the feedback from users is positive [Section 5.6]. However, as already discussed, there are clear issues with the models,

⁴²⁰ The metal and gold leaf surfaces were treated as having the same material properties for reasons of convenience and speed; whilst the gold could perhaps have been treated as less shiny than the metal, they were close enough to combine.

⁴²¹ The clouds contain approximately 400k and 1.2M points respectively; the combined filesizes for the two .pjs files is approximately 18Mb

⁴²² The metal was rendered with a shininess value of 80 and specular colour (67,67,67), the leather with 38 and (26,26,26). The colour values show a simple white colour for both sets of specular highlights, but with an intensity roughly 2.5 times greater for the metal.

and it is interesting to note how and why these issues affect the objects' reception and utility. Both models have inconsistencies and imperfections that are readily noticeable to anything more than a cursory inspection, but it would appear that these imperfections are less important to users than the overall effect and the wider context of both the object and the website.

5.5.1 The bowl

As discussed at some length, the fundamental problem encountered with the process of capturing and displaying the bowl is, unfortunately, an unavoidable one: the nature of the object itself. The shininess and specularity were – and were expected to be – problems from the start. Whilst there have been many approaches to identifying and processing specular highlights in images for the purpose of photogrammetry, some apply to limited domains such as laparoscopy^{423 424}, and most concentrate on identifying – and in some cases inpainting – small, focussed highlights⁴²⁵ and may not be applicable to the extensive highlights such as those seen on the bowl (Figure 5.2). Other approaches are aimed solely at reconstructing geometry from specular objects and are not concerned with reproducing texture. Our less sophisticated (though reasonably successful) method for identifying highlights is detailed in Appendix A. At the current time, there does not appear to be a method for creating a complete and fully textured model from a highly specular source, and considerable processing of the

⁴²³ Stoyanov, D, et al, A practical approach towards accurate dense 3D depth recovery for robotic laparoscopic surgery, Computer Aided Surgery 2005, Vol. 10, No. 4 , Pages 199-208

⁴²⁴ Arnold, M, et al, Automatic segmentation and inpainting of specular highlights for endoscopic imaging, EURASIP Journal on Image and Video Processing, Vol 2010

⁴²⁵ Gershon, R, et al, The use of color in highlight identification in Proceedings of the 10th International Joint Conference on Artificial Intelligence, 1987 vol. 2, pp. 752–754.

incomplete and noisy data is inevitable. Whilst, in my experience, and when aiming for an aesthetically pleasing result, all models require extensive processing (estimated, at least in the case of laser scanning, at five times the time it takes for the initial capture⁴²⁶ - though this estimate is still on the low side and applies to a 'best case scenario' object as opposed to a difficult one such as the bible or particularly the bowl), but the particular problems in this object meant the whole process of erasing noise, filling holes and then cleaning up the new points was a fairly arduous and time consuming process⁴²⁷. In total, somewhere between 60 and 80 hours were spent on processing the model and preparing it for online display. Even then it is impossible to say that at any stage the model was 'finished'. However, as the Courtauld wanted the 3d model to launch at the same time as the rest of the bowl's web content went live (or as close to it as possible), the amount of processing was at least partly defined by this deadline. Certainly more time could have been spent cleaning up and repairing data, achieving a better alignment etc. We can state, with reasonable confidence, that we produced the 'best possible' model - *with the resources available*. These resources include the time available and the hardware and software used, but also, importantly, the skills and experience of the operator, in this case, myself. As such the amount of processing was determined by the time available⁴²⁸; the quality of the processing by

⁴²⁶ Ecurator final report

⁴²⁷ At least some of the issues with processing are due to the choice of software used. Whilst there were some bugs in the Pointstream beta version, the software is designed to process data from the Arius scanner and many of the difficulties were exacerbated by the fact that our data was generated from other sources.

⁴²⁸ It should be noted that even if we had spent more time on the model, subsequent processing would probably have suffered from the law of diminishing returns, with the time taken for a given increase in quality rising steadily.

the user's skill⁴²⁹. There is a steep learning curve both in the general principles of processing point clouds, and with the specific software used, and it should be remembered that when we say we produced the 'best possible' model, we are not claiming that someone else with more experience, using different software etc. may have been able to create a *better* model within the same timeframe.

Ultimately, the creation of a model like this, for online display to a wide audience, is an artistic process aimed at achieving an acceptable aesthetic output. This criteria was decided upon between myself and the Courtauld before the commencement of the digitisation as an aim that was both achievable and useful within the concept of the Illuminating Objects Programme, one of the aims of which was to expose interesting and hitherto largely un-seen objects to a wider audience. Thus by an 'acceptable aesthetic output' we mean a model that would be engaging and would benefit the user's understanding of the object.

One problem with this approach is that there is no easy metric that says when this aim is achieved and the object is 'finished'. A subjective judgement must be made, but ultimately the success or otherwise can only be determined by gauging the reaction of the intended audience.

Looking more closely at the bowl model, it is, as always, important to bear in mind the *purpose* of the model – or, perhaps more helpfully, to consider what it is *not* for. This

⁴²⁹ This is, I believe, an important point: whilst the process of preparing a model can be – depending on the particular characteristics of the model – be automated to some extent, it is ultimately a craft process, and as such depends to a large extent on the operator's skill, experience and familiarity with the tools at their disposal.

object has no real detail that is revealed by zooming in closely, nor is the model intended to be used for research or as an archival record of the object. Thus the tolerances for imperfections and for interpolated data (ie the extensive hole-fills) are far larger than in an equivalent model created for use by curators, conservators and academics. Many of the noisiest areas do not stand up to close inspection, but close inspection is not what this model was created for. Rather it is to give the viewer an impression of the physical shape of the bowl and its particular visual qualities. I believe the model was successful in achieving the former⁴³⁰, perhaps less so the latter.⁴³¹ This judgement is shared by Dr Gerstein. She is satisfied with the overall look of the model – though with the caveats that she understands both the limits of the technology and the problems inherent in this particular object (specifically the way it reflects light) – and believes that, nevertheless, it “does the job”⁴³². At the same time she admits to being uncomfortable with the addition of the lighting effects to the model. Because the lighting effects were 'invented' and not part of the objectively⁴³³ measured properties

⁴³⁰ I believe it is suggestive that the Courtauld's web designer was genuinely surprised by the bowl's geometry when she saw the 3d model for the first time (for example, the presence of the raised 'boss' in the center); this despite having designed all of the bowl's online presence and presumably having been fairly involved in examining the photography. This suggests, albeit anecdotally, that the 3d model does offer something over and above the two dimensional imagery available elsewhere on the site.

⁴³¹ It could also be noted that the model also offers the user the opportunity to view the design and decoration of the entire bowl, front and back, including the detail of the inscription on the inner rim in a way that is not possible – ironically due to the object's specularity! - from any of the photographic imagery elsewhere on the site.

⁴³² All quotes from Dr Gerstein taken from an interview conducted Oct 2013, at the Courtauld Institute.

⁴³³ I accept the difficulty of using terms like 'objective' with regard to any virtual model. Subjectivity is introduced at all stages of the process, from the choices made by the designers and operator of the equipment used to capture and the software used to process the data, all the way through to the authors of the viewing software and even the creators of the browser engine used to ultimately process the commands and render the images. Here I use 'objective' in a fairly weak sense, in that we can potentially follow all the steps in the process back to some physical interaction with the object itself, as opposed to the process of adding lighting effects meant to evoke properties of the object, but which are ultimately based on what are purely subjective choices.

of the model (as opposed to the geometry and diffuse colour), this prevents her from “bringing [the model] into discussion with a colleague ... it wasn't a record.”

Ultimately, what makes the object both beautiful and unusual is the complex way that light interacts with its surface. For instance, there is an iridescence to some of the highlights [Figure 5.3] not captured by the virtual rendering. The highlights on the real object are also 'more interesting' in terms of shape, colour and intensity than those on the virtual model. On the actual bowl they appear as random pools of light playing across the bowl's surface, whereas the highlights on the model spread uniformly and evenly. The complexity of the real highlights are no doubt due to the particular physical features of the bowl; both the detailed geometry which was not captured by the photogrammetric process, but also natural imperfections or variations in the thickness or properties of the glaze which of course cannot be replicated simply by collecting texture and geometry (no matter how detailed).⁴³⁴ We must ask ourselves the question, particularly in regard to Dr Gerstein's comments, whether it would have been better to simply present the model 'as is', rather than to attempt to recreate the subjective experience of seeing the bowl.

5.5.2 The bible

In contrast to the bowl, the bible's appeal resides less in particular visual effects and more in the object itself. Whilst there were issues with the shininess of the metalwork on the bible's cover, resulting in small holes in the data and problems with colour

⁴³⁴ The differences may also be due in part to the details of the lighting model used by the Pointstream viewer, though I would guess that these would be minor compared to the differences caused by physical effects

consistency, from a processing point of view, these were relatively minor when compared to the bowl. On an initial inspection, there are more obvious flaws in the bible model than the bowl; there are many holes, especially where the inside covers meet the pages and on and around the clasps, and these are noticeable as you rotate the object. A section of the very top of the spine, where the scanner didn't pick up the dark leather, is missing. The pages behind the clasps look 'messy' and the texture on the front cover is blurred and lacking in detail compared to the back of the book. These flaws are clearly visible without zooming in on the model; the bowl's major flaws, in contrast, only become obvious when the user zooms in. And yet from (early) user response, and comments from Dr Gerstein, the bible is the more engaging and successful model.

It is possible that the flaws in the bible are only glaring if you are aware of the problems with the model beforehand. It could also be that users are simply more tolerant of these errors (one user comment: "Sometimes you could partly see through the object to the back ... although it didn't particularly detract from the object") if the overall experience is engaging.

Why should the bible be more engaging? It is of course possible that a miniature bible is simply a more interesting object than the bowl, indeed, according to the survey attached to the online model 75% of respondents already had a 'limited' or better interest in this type of object before visiting the site, compared to the bowl where 80% of respondents had no previous interest. I would hypothesize that there are certain ineffable qualities that the book model has; in some way the book's compact shape

rewards interaction better than the large ‘flat’ more two dimensional bowl. The lighting is simpler and perhaps more realistic on the book, we are more familiar with the simple contrast between shiny metal and matte leather compared to the more complex properties of the bowl (again, the very qualities that make the bowl both unique and beautiful work against it). As Dr Gerstein says, “The books really worked well, it wasn't perfect but I really liked the weight of it, you felt it was a volume.” The synergy with the other features on the website may also explain the bible's overall appeal: “Combined with page turning mechanism it worked extremely well. As a package, it worked. I was extremely pleased.”

5.6 User Survey

5.6.1 Methodology

As in the previous chapter (4.9), since the resources on which the research is to be conducted are digital, web based resources, an online survey was created. Again, the Opinio platform was used and the survey linked from the web pages containing the 3D models.

The survey is similar to that created for the Science Museum project, as the research aims and target audience were broadly similar. However, due to both the subject matter and audience – the Illuminating Objects programme by its nature deals with interesting, but fairly obscure and ‘niche’ objects, as opposed to a much-loved and well known public gallery – the survey is a similar length but slightly less informal in tone. Also, as the resources were interactive and required fairly novel technologies (ie,

webGL), the survey contains more questions regarding technical and experiential aspects of the resource, for instance download times and smoothness of interaction. The surveys for both bowl and bible are identical, and unless specified, the combined results will be used in this analysis.

Whilst issues with the methodology will be examined at the end of this chapter, the problem of response – or lack of it – is too important not to bring up here. The volume of responses (24 for the bible, eight for the bowl, or 32 combined⁴³⁵) makes it difficult, as in the Science Museum survey, to make strong claims for the data. Unlike the Science Museum survey however, there is not a corresponding volume of qualitative data to be found elsewhere which could be analysed on its own or used to support the purely quantitative findings of the survey. Reasons for the low number of responses will be examined in section 5.9 but the small sample size should be borne in mind throughout the following analysis.

5.6.2 Analysis

The questions in the survey are printed below in Figure 5.20, the full responses can be found in appendix C.

⁴³⁵ Not every respondent answered every question, so in the following analysis, where total responses are shown to be out of a number less than 32, ie, 9/31 or 7/30, this indicates a missed response.

Question 1: Downloading: How long did the model take to download fully?

Question 2: Once the model was downloaded, how did it appear to move?

Question 3: Interaction: How did you find your interaction with the model (moving it, zooming etc.)

Question 4: Do you have any other comments on the technical aspects of this site and your experience with it?

Question 5: Do you think that viewing the 3D model improved your understanding of the object?

Question 6: Did the model look like a convincing, 'real', three-dimensional object?

Question 7: Would you like to see more 3D models on the Courtauld's website?

Question 8: Would seeing 3D models online make you more or less likely to visit the gallery? For example, having seen the model of this miniature bible, are you more or less likely to visit the gallery to see the real thing?

Question 9: Do you have any other comments or thoughts on your experience with the model?

Question 10: How would you rate your level of online experience?

Question 11: How much experience have you had with 3D models and 3D technology? (This could be, for example, through accessing models like this one, from playing videogames or using other virtual environments and technologies)

Question 12: Before visiting this site, did you have a particular interest in this bible, or this type of object?

Question 13: Please add any comments you may have on the previous questions, and especially on your particular interest (if any) in this, or this type of, object.

Question 14: Thanks for completing the survey! Feel free to add any further comments below:

Figure 5.20: Survey questions

5.6.3 Demographics

The respondents identified (Q10) as a group with a lot of online experience, all professing to be an average web user or better. 18/31 (56%) are 'familiar with most online technologies' while 9/31 (28%) are 'regular web users, comfortable online'. This may be a representative sample of the general public, or at least the subsection of the general public with the ability and motivation to access a museum's website, though it is likely that there is further self-selection occurring that biases this group towards the

technical/web savvy as well. It may be assumed that visitors to the site would be less likely to attempt the survey if they had not been able to access the resource at all, and this would disqualify, for example, all Internet Explorer users as at the time of the survey this browser did not include a WebGL renderer⁴³⁶. According to NetMarketShare⁴³⁷, IE represented over 50% of all desktop browsers in 2013, and its users tend towards the least technical and 'web savvy' end of the spectrum of web users⁴³⁸. Thus, the experiential level of respondents may be skewed towards the more technical simply by their ability to access the resource.

Among the group experience with 3D technologies (Q11) also appears quite high, with 20 out of the 30 respondents declaring an above average experience with 3D models and just four claiming little or no experience. This may be due to the same selection criteria described above, a reflection of the ubiquity of 3D technology in games and the media, or it may be that people with interest in 3D technologies were more likely to find and access the digital resource. The truth is probably a combination of all three.

Finally, the respondents include a mix of people with pre-existing interest in the objects and those with none. Overall, 12/30 or 40% expressed a previous interest with 18 (60%) having limited or no interest. People generally had less of an interest in the bowl than in the bibles, a finding corroborated by the different amounts of publicity the two

⁴³⁶ WebGL was not usable in IE until IE 11, released in 2014

⁴³⁷ <https://www.netmarketshare.com/browser-market-share.aspx?qprid=2&qpcustomd=0&qpssp=2013&qpnp=1&qpftimeframe=Y>

⁴³⁸ For comparison, W3Schools, the tutorial website run by the world wide web consortium includes browser statistics for people visiting its sites- http://www.w3schools.com/browsers/browsers_stats.asp; presumably people with interests in html and other web technologies. Their statistics show that IE users made up only approximately 10-15% of their visitors, demonstrating that technical web-savvy users would opt for Chrome or Firefox (both WebGL enabled) over Microsoft's browser.

objects received⁴³⁹. This interest was reflected in visitor numbers to the two web pages, and ultimately, in the number of survey responses for the two objects.

5.6.4 Technical/experiential

For the majority of users, the technical experience was smooth and problem-free in all three areas – speed of loading, smoothness of movement and ease of interaction⁴⁴⁰.

Again, the results are likely to be skewed towards a positive result; users who couldn't view the resource at all, or were put off by a particularly poor experience (excessive loading times etc.) would, presumably, be less motivated to complete the survey^{441,442}.

There is little correlation between the first two questions, which is unsurprising as download speed (Q1) is primarily a reflection of the user's online connection while movement/animation (Q2) and (to a lesser extent) interaction (Q3) are all performed locally and are dependent on the power of the user's computer.

The average response for all three technical questions (Q1-3) were better than average (2.5 in each case). The mean response for Q1 (loading times) is 2.3, Q2 (animation), 1.9

⁴³⁹ The bowl received no publicity outside of the Courtauld and SOAS websites, the bible slightly more, though very little compared to the Shipping Gallery: <http://www.indielondon.co.uk/Events-Review/german-miniature-picture-bibles-under-the-spotlight>, <https://www.fieldandrurallife.com/e-magazines/field-and-rural-life-e-magazine/item/16614-illuminating-objects-german-miniature-picture-bibles-come-under-the-spotlight-at-the-courtauld-gallery-.html> & <http://www.visitmuseums.com/exhibition/illuminating-objects-german-miniature-picture-bibl-3653> - note that only the first two of these pages mentioned (and linked to) the website.

⁴⁴⁰ The smoothness of interaction is a combination of technical performance and the control scheme; if it was purely down to technical factors then one would expect a very strong correlation between Q2 and Q3. While there is (remembering the small sample size) a strong correlation ($P = .55$), the fact that a small number of individual responses do show large disparities between the two questions indicates that there is more to interaction than just technical performance, reflected in the text comments about control schemes.

⁴⁴¹ In fact, one respondent couldn't access the resource at all (in Firefox) but did take the survey. They didn't answer any of the questions but did leave a comment.

⁴⁴² There are many reasons why a WebGL app won't run in a browser that ostensibly should; WebGL may be switched off in the browser's config files (quite common on managed PCs) or the gpu may be blacklisted by the browser vendor for security reasons

and Q3 (interaction), 2.0. It is difficult to generalise from these results, but we can say, at least, that it is possible to render a 3D point cloud model within the browser that is detailed and large enough to provide a user experience that is smooth, convincing and useful (see Q5 and Q6 below). A file of around 40-60 Mb appears to be a reasonable size to download for most people⁴⁴³.

5.6.5 Understanding

One of the major research aims in this chapter is to assess the utility of 3D models and see how they can aid in understanding a three dimensional object (Q5). In terms of these objects, the results indicate that the models were indeed successful, with the majority of users indicating that their understanding was improved⁴⁴⁴ (8/28 or 29% 'very much' and 15/28 or 54% 'a little'), with four users not reporting an increase in understanding and one user saying their understanding was made worse by the 3D model⁴⁴⁵. Overall, the average users' score was 1.92 (2.5 being average).

In examining the different responses to Q5 for the bowl and the bible, it must be reiterated that, particularly when it comes to the bowl, there is a very small sample size, so these findings are extremely tentative. Given that, the results for the bowl

⁴⁴³ Note that the .psi files being used are a proprietary format and quite efficiently compressed (approx. one fifth of the size of an equivalent ascii file). Also the Pointstream.js viewer starts rendering the object before the file is fully downloaded, which may increase users' tolerance for download times.

⁴⁴⁴ Of course, the caveat should be noted that this was a purely self-assessed metric, and there is no way of telling from this survey whether the users' understanding of the objects' three dimensional properties was, in fact, improved.

⁴⁴⁵ The user who claimed their understanding was worsened by the model had a very good technical experience with the model, so their response cannot be blamed on, for example, poor interaction. It is interesting to note that they expressed a strong pre-existing interest in the object (the bible, in this case), and thus one can assume had a better understanding of the object to start with. An interesting area of future research would examine the ways a 3D model could inform users familiar with the object compared to those with no prior understanding.

seem significantly better, with viewers of the bowl model registering an average score of 1.42 for Q5 compared to 2.1 for the bible. It is possible to speculate on one possible reason for this disparity: the bowl has a slightly less ‘predictable’ or obvious shape. The unusual and distinctive ‘boss’ in the centre, which is clearly visible on both sides of the bowl, is not immediately evident from the 2D images of the bowl elsewhere on the site⁴⁴⁶. In contrast, the bible is a fairly predictable ‘book’ shape, albeit on an unusual scale. There are fewer surprises to its geometry and thus perhaps it is harder to improve the understanding of its 3D shape. This would support the (one would think fairly uncontroversial) hypothesis that 3D models of objects have greater utility when the object in question has interesting three dimensional geometry; especially when that geometry is hard to represent in two dimensional imagery.

The majority of users found the models convincing as 3D objects. 18/30, or 60%, found it ‘not perfect, but quite real’ while six people were ‘completely convinced’ and six thought it ‘looked a bit fake’. A mean score of 2.0, (with 2.5 being average) confirms the models were successful as convincing 3D representations of the objects; the average was identical for the bible and bowl models.

5.6.6 Interest in 3D models in general

As in the Science Museum survey, there appears to be an appetite for this type of resource amongst the public, with 11 out of 27 respondents (40%) expressing a desire to see ‘lots more models’ on the Courtauld’s website, and seven wishing to see ‘some

⁴⁴⁶ See footnote 74 re. the Courtauld’s web editor’s reaction to the 3D model.

more'. Eight remain ambivalent. Only one respondent expressed a positive (negative?) desire to see fewer models.

As a marketing exercise to drive people to exhibitions, it appears that the further use of 3D models could be a successful tactic. Eight of 29 (27%) respondents said they would 'definitely' be more likely to visit the gallery to see the real object, while 13 (45%) said they would possibly be more likely, depending on the object. Another eight (28%) were ambivalent, but none of the users indicated that the 3D model would put them off visiting the physical gallery⁴⁴⁷.

5.7 Correlations

Some correlations have been discovered in the data, and whilst due to the small sample sizes and nature of the data, it is impossible to ascribe statistical significance to them, they might at least suggest areas for future research.

One of the possible correlations visible in the data is a negative one between how much interest the visitor had in the object prior to using the resource (Q12), and the increase in their understanding of the object (Q5) (Figure 5.21). This is perhaps unsurprising; the less knowledge the visitor has of the object before using the resource, the more potential there is for their understanding to be increased; those familiar with the object have 'less to learn'. This has implications for 3D digitisation projects, for both audience and object selection. Is it better to digitise popular but well known objects, or more obscure artefacts? Should the resources be aimed at people

⁴⁴⁷ Again, we are relying on 'self-assessed' responses; a more useful metric would be to survey visitors to the gallery itself and ascertain how many had previously visited the 3D resource.

with pre-existing knowledge of the object or a more general audience?




Figure 5.21: Interest vs understanding

There are also possible correlations between the smoothness of both animation and interaction and how convincing the model was as a 3D object, and the understanding gained of the object. In both cases the numbers are small, and so any implied correlation must necessarily be weak. However, in those visitors whose interactions with the digital object were 'intuitive and natural' or 'fairly easy', the average score for Q5 was 1.8 (with 1 being 'best' and 2.5 being average), whereas for those whose interactions were just 'ok', or 'frustrating', the average score is 2.3. Again, a small result but one that supports the intuitive view that a better interactive experience will aid learning. Similarly, for those who found the object wholly convincing, their average score for Q5 (understanding) was 1.7; this drops to 1.9 for those who considered the model 'quite real' and to 2.3 for people who found the model 'a bit fake'.

5.8 Comment analysis

Analysing the free text responses does not elicit too much new information. However, by far the most common sentiment refers to the model's automatic rotation (2075.3.3) ("I would have preferred it if the model wasn't spinning all the time" (796843), "The bible needs to stop spinning when you're no longer rotating it." (716808), "The model never stops moving which could be a good thing unless you want to look in detail at one part of it." (780923), "I found the fact that the model kept rotating the whole time slightly annoying." (780928)). As discussed earlier, the rationale for the constant rotation is to ensure that visitors to the web page realise that the resource is a 3D model and not simply a still, 2D image. However, it would seem from these comments that a compromise may be necessary – for example, the object could rotate when the visitor comes to the page, but stops after the first time the user interacts with the model.

There were several comments on the controls. One user said "I did double click to try to zoom in, but it just reset it. I think I am used to google maps controls" (780925) which points to the fact that, as yet, there is no accepted 'default' mode of interaction with 3D objects⁴⁴⁸, and that conventions can differ between platforms even in the

⁴⁴⁸ Apart from the games industry, where the 'WASD' for translating movement and the mouse or arrow keys for rotation are fairly standard. However, this set up is designed for navigating 3d spaces as opposed to examining 3D objects, and is not applicable to touch screen/mobile devices. The fact that the games industry is struggling to find a default keyboard/mouse alternative in the touchscreen space is indicative of the inherent difficulties. See, for example:

<https://tampub.uta.fi/bitstream/handle/10024/84053/gradu06268.pdf?sequence=1>

comparatively mature world of 2D manipulation. Similarly, some functionality differs on Mac and Windows machines⁴⁴⁹ (“ctrl and left mouse button zooms in on the whole display for the mac rather than the bible, so a bit difficult” (716798)). As the market splinters even more, especially in the mobile space (and VR is likely to become a relevant option in the next few years), either the adoption of a standardised control scheme, or the provision of alternatives, will become more important.

In general, users appeared to be satisfied with the models’ appearance, (“Astounding level of detail.” (716844)), though two were disappointed that a full screen mode (and presumably, therefore, a higher level of zoom) wasn’t available (“It doesn’t make sense to have such a small viewing window, users expect to be able to use full screen feature” (716872), “A full screen version would be better. Even though I could zoom in I found it a bit small.” (715880)). A full screen/higher level of zoom would necessarily require a larger/denser model, to avoid either the muddy low resolution effect that increasing the point size creates, or alternatively the translucent ‘point cloud’ aesthetic as seen in the Shipping Gallery project. From some angles, at the highest levels of zoom available, there is already some translucency, particularly in the bible model, though interestingly, some users can look past (excuse the pun) this effect: “Sometimes you could partly see through the object to the back of the object although it didn’t particularly detract from the object” (780928). This user rated the ‘convincingness’ of the object as 2, or “It wasn’t perfect, but looked quite real”, indicating that, for at least

⁴⁴⁹ It is interesting to note that one user accessed the models from an ‘Android tablet with Firefox’ (718204), so even in 2013, at least some mobile technology was capable of rendering the models.

some people, the 3D rendering does not need to be absolutely perfect to convince.

As an aside, several visitors to the ‘bowl’ page commented on the ‘hollow face illusion’⁴⁵⁰, where the bowl model seems to snap between concave and convex as it revolves (“I think the geometry of the plate confused me when turning from one face to another. It looked as if there was a dent in the centre but I’m sure it was an optical illusion!” (781650), “I’m not certain which pattern should be the inside of the bowl and which should be on the bottom as sometimes they would swap over and what was the inside became the outside.” (780923)). This illusion would appear to be an artefact of this particular object, combined, perhaps, with the player’s automatic rotation (the effect is lessened, but does not entirely disappear, when the user is rotating the model manually), so may be of limited importance.

5.9 Methodological issues

The single most important methodological issue with this survey was the low response number. This does not necessarily equate to a low *response rate*, ie, the proportion of users who viewed the object and then undertook the survey, it is instead (or as well as) an artefact of the low numbers of visitors to the pages. The visits to the bowl page, for example, were measured in ‘hundreds rather than thousands’⁴⁵¹, as may be expected for a niche website which received little publicity outside of the Courtauld institute itself (see footnote 84). Bearing this in mind, the conversion rate from visits to survey

⁴⁵⁰ https://en.wikipedia.org/wiki/Hollow-Face_illusion

⁴⁵¹ From personal communications with the Courtauld’s web team – note this is a ‘raw’ figure as well, and so includes bounces etc., and thus the number of deliberate or purposeful visitors is probably a fraction of the total.

response is probably in line with expectation. The miniature bibles represent perhaps a more interesting subject for the general public, and accordingly received more publicity (though still very little compared to the Shipping Gallery video). As such it is not surprising that the bible survey should have received approximately four times as many responses as the bowl; this is roughly in line with the Courtauld's estimation of visitors to the two pages.

Combined with the low number of visitors over all, there is also the issue of the technologies used. At the time of the survey, as we have discussed, Internet Explorer users (amongst others) would not have been able to access the resource at all, and indeed some users of ostensibly webGL enabled browsers would have found it difficult as well, due to security policies implemented by IT departments, insufficiently powerful hardware or simply the wrong GPU (Graphical Processing Unit). Today, this situation is greatly improved – almost all of the latest iterations of the major browsers are webGL enabled, and mobile devices, becoming more powerful with each generation, are capable of rendering complex three dimensional scenes. However many users are still using older versions of browsers and according to some sources, the market penetration for webGL is still only around 80%⁴⁵².

This leads to another general accessibility question: is it ok for an institution to host a resource that is only accessible to a section of their audience? This was of particular relevance to the Courtauld at the time of this project, as only around half of their

⁴⁵² <http://caniuse.com/#feat=webgl> & <http://webglstats.com/>

visitors would be able to see the 3D model. The Courtauld acknowledged this, but were happy to go ahead in the knowledge that this was an experiment, at the cutting edge of several new technologies. However, some consideration should perhaps be given to providing alternative or fall-back content, for instance a pre-rendered video as in the Shipping Gallery project.

5.10 Solutions?

The simple answer to the low numbers of survey respondents would be to drive more traffic to the site, though it is difficult to see how this could be achieved without distorting the research. The Illuminating Objects programme was initially publicised by the Courtauld via their PR team, website and Twitter, and also via the UCLDH Twitter account. As we saw earlier, there was limited uptake in the media, particularly for the bowl. One option to create more traffic would be to extend the PR campaign beyond those interested in the objects as objects, for example, to explicitly involve people in the 3D scanning community. However, the illuminating objects programme, by its very nature deals with niche objects. Publicising it beyond these areas, for example, aggressively marketing the site to other people who might be interested risks changing the visitors' demographics and thus skewing the results of the survey. The aim was, after all, to see if the resource is useful for the 'normal' visitors to the site; people who are interested in learning about the objects, not, primarily, the 3D technology. Whilst it would no doubt be interesting to hear what people involved in 3D digitising thought of the models, the research is specifically aimed at the general public and in measuring their interest and the utility of 3D models for them. By the (albeit self-assessed) level of

3D experience in the survey respondents (10/32, 31% said “I have had a lot of experience with 3D”, 10 said “I have had a fair amount of experience”), one could guess that the demographics are already skewed towards people working, if not in 3D per se, then at least in the digital sphere, and some comments support this view (“As a general rule, online digitization increases visitorship. This has been long established by libraries with digitized special collections holdings online.” (716872), “interested in seeing an example of a 3d model used in digital library collection” (719162)).

Another issue involves driving people directly to the page with the model (and survey link) rather than to the object’s home, or landing, page. Whilst this no doubt improves the number of users that see the model (and therefore complete the survey), this is not the normal user experience, and serves to show the model outside of its usual context. This can be seen from the comments: “It's best to use 3D to offer the in-person experience as closely as possible. In this case, I would expect to be able to open the book and page through it. Content + purpose should drive form + function.” (716872) and “For an art gallery, it's odd to focus on a closed manuscript as 3D model, the art is inside...” (716872). If the users had accessed the model in the ‘normal’ way, via the object’s landing page (Figure 5.22), they would have seen the ‘Turn the pages’ link which would have allowed them to page through the book and see the artwork inside.



Figure 5.22: Bible landing page

5.11 Alternatives to online surveys

There are two main questions to answer in this section; is there an alternative way of recruiting research subjects, and what are the alternative methods for conducting the research.

One solution to a lack of subjects, and therefore a lack of data to analyse, would be to extract more information from each subject. This could be in the form of a more in-depth survey, trading off the potential of losing respondents due to survey length with more information per survey⁴⁵³. Alternatively, interviews could be conducted (possibly

⁴⁵³ Edwards et al, 2009. Methods to increase response to postal and electronic questionnaires. Cochrane

via email) with users recruited via the website. Again, one would expect fewer recruits, but more information could be extracted from each one via an interview than a survey. Another option would be to conduct face to face user testing, using a mix of structured (ie, task based), semi-structured (guided testing) or unstructured testing (using, for instance, think-aloud protocols)⁴⁵⁴. The advantages of this method are clear, with, presumably, a much richer data set per user. The disadvantages would include similar difficulties in recruiting users, and the lack of anonymity in responses. This sort of focussed user-testing is also more suitable for measuring usability and particular task-based performance, and introduces certain artificialities into the user experience.

The other path to take would be to increase the proportion of people who see the survey link and follow it, and one method would be to incentivise users to participate. A common technique in online surveys is to offer a monetary incentive to users, or, as is more common, to offer to enter respondents into a draw to win a small prize. However, as well as the intuitive feeling that this is not perhaps suitable for academic research (and would somehow ‘cheapen’ the survey in the eyes of users), it is also unclear what effect this actually has on recruitment⁴⁵⁵ and may affect the quality of data collected⁴⁵⁶. An alternative would be simply to ‘market’ the survey more aggressively; for instance by making the link and call to action more prominent on the

Database of Systematic Reviews 2009, Issue 3. Art. No.: MR000008

⁴⁵⁴ <http://www.nngroup.com/articles/which-ux-research-methods/>

⁴⁵⁵ Edwards et al (ibid) records an increase of about a third in respondents when offered monetary incentives, while Kalantat & Tallet (below) report very little effect.

⁴⁵⁶ Kalantar, JS & Talley NJ, 1999, The Effects of Lottery Incentive and Length of Questionnaire on Health Survey Response Rates: A Randomized Study. *Journal of Clinical Epidemiology*, Vol 52, Iss 11, pp 1117-1122, Nov 1999

web page. However, it must be borne in mind that in a collaborative project such as this, the website is part of a large institution and there are other branding, content and design considerations that will take precedence over the needs of one researcher: the webpage is not there to drive visitors to a survey. With the benefit of hindsight, I am confident that, had we known about the low response rate in advance, the Courtauld's (very helpful) web team would have been amenable to giving the survey link more prominence. However, as is the case with the majority of 'event' websites such as the Illuminating Objects pages, interest, and online traffic, peaks at the launch when the publicity appears, and quickly tails off⁴⁵⁷; changing the design after this initial peak would have little effect on overall numbers⁴⁵⁸.

5.12 Conclusions

Certain tentative conclusions may be drawn from this project; for example, it is possible for a museum or cultural heritage organisation to create a 3D point cloud model and serve it to the public online. At current (or indeed, three year old) levels of technology, the model used can be detailed enough that users find it a valuable resource, and yet small enough (in terms of file size) to still provide an acceptable user experience.

However, one of the themes of this thesis is how accessible this technology is for cultural heritage institutions. Whilst few, if any, have access to a £500k laser scanner, as was used for the bibles, many institutions can be expected to have access to

⁴⁵⁷ For example, 15 of the 24 responses to the bible survey were completed within a week of the launch

⁴⁵⁸ It should be remembered as well that Arius stopped supporting the WebGL viewer not long after the launch of the bible page, and functionality (for instance, Firefox support) began falling off soon after.

photographic equipment (and photographic experience) that is in many cases superior to those used for the bowl. The outputs for both methods are, in fact, comparable, and both exhibited similar weaknesses when capturing shiny objects. For these reasons, photogrammetry would appear to be a more suitable method for cultural heritage institutions.

The computer hardware used in this project was a reasonably powerful, but affordable (under £1,000) desktop computer. In terms of software, the creation of the bowl model used a free, open source package (VisualSFM) with a relatively user friendly GUI. No special skills, training or prior experience are required to use it⁴⁵⁹. On the other hand, the Pointstream software used to process both models (and indeed, all the object models mentioned in this thesis) was a commercial product which is no longer available, and whilst there is free software available for processing point clouds, it does not have the same functionality as Pointstream. Similarly, the rendering software (Pointstream.js) was provided free for this project, though it was intended to eventually become a commercial project; it too is no longer available. Open source and free alternatives exist⁴⁶⁰ which have much of the same functionality as Pointstream.js, though they lack integration with point cloud processing software and some of its ease of use⁴⁶¹. Sustainability was a major issue with this project, as, not long after the

⁴⁵⁹ Not everyone is fortunate enough to be embedded in a department with considerable experience in photogrammetry and photogrammetric software, VSFM is nevertheless a fairly simple piece of software that, with a small amount of trial and error, can be learned fairly quickly. And while the documentation is fairly sparse, there is a helpful online community that will answer many questions.

⁴⁶⁰ For example XBPS used elsewhere in this thesis and upon which the Pointstream.js software was based, or potree – see <http://potree.org/> for more details

⁴⁶¹ For example, it was easy to set up lighting models in the Pointstream software, which were then output as simple chunks of html which could be pasted into any web page.

second (bible) model went online, the company providing both the Pointstream software and the WebGL viewer went out of business and ceased supporting either application. Due to the natural evolution of browser technology, as individual browsers were updated, functionality of the online viewer began to degrade. Within a few months of the support ending, the viewer stopped working at all in Firefox, and Chrome followed not long after. For a CH institution who has invested time and expense in digitisation and wishes to extract maximum utility from their 3D content, the issue of sustainability and the longevity of digital resources must be a real concern, and something addressed at the very beginning of a project.

To conclude, the majority of the photogrammetric workflow for capturing and disseminating 3D models via the web is within the budgets and capabilities of many cultural heritage institutions, requiring the acquisition and installation of potentially free software and a small investment in staff training and development. While there is still a 'missing piece' with the lack of good, free software designed for producing aesthetically pleasing point cloud models, there are indications that at least in the commercial sphere, this type of software is beginning to appear (section 2.5).

It is harder to draw concrete conclusions from the survey results, but I believe they make a *prima facie* case that the models succeeded in their aim; to provide an engaging resource (ie, that improves engagement with both website and object) that can also contribute to understanding. However, it is clear that further research is needed.

6 The Courtauld Bag

6.1 Introduction and research aims

This chapter describes a major project undertaken for The Courtauld Gallery, in which both 3D modelling and RTI imaging was used to create an animated video which was displayed in-gallery as part of a major exhibition. It introduces both the project and its subject, a 14th century silver and brass bag, before detailing the capture methodology and processing workflow. It details the difficulties involved in capturing an object of this type using both laser scanning and photogrammetry and the ramifications of this for the rest of the workflow, including point cloud processing and rendering. Section 6.7.2 describes Lindsay MacDonald's novel technique for rendering RTI imagery which was used to create a photo-realistic animation of some of the bag's details.

Finally, it describes and analyses the research conducted on the output video, consisting of informal interviews conducted with visitors to the gallery.

6.2 The Courtauld Bag

The Courtauld Bag was manufactured in Mosul in northern Iraq sometime between 1300 and 1330, approximately half a century after the invading Mongols had ended the 500 year reign of the Moslem Abassids, and just a few years after the rulers of the Mongol's Il-Khanate dynasty had converted to Islam⁴⁶². The bag is a unique object and one of the most important examples of metalwork from the Islamic world⁴⁶³. The bag is

⁴⁶² Melville, C, Northern Iraq: Historical and Political Context, in Court and Craft: A Masterpiece of Northern Iraq (Catalogue), The Courtauld Gallery in Assoc. with Paul Holberton Publishing, London, 2014, p16

⁴⁶³ Vegelin Van Claerbergen, E, Foreword,Catalogue, pp8 ibid

constructed from sheets of brass, soldered or hinged together⁴⁶⁴, and decorated with a black 'ground' (an organic substance, probably bitumen or conifer resin⁴⁶⁵) and silver and gold inlays, and the overall shape is reminiscent of 'the rounded shape of a leather or textile bag'⁴⁶⁶. The geometric patterning, picked out in silver wire, as well as other details of construction, also imply that the object was designed to resemble other leather or textile bags of the time⁴⁶⁷, though unsurprisingly there are few existing examples of contemporary bags made from perishable materials.



Figure 6.1: Image of the bag (from the first photogrammetry imaging session)

⁴⁶⁴ Ward, R, Catalogue, p76, *ibid*

⁴⁶⁵ Ward, R, Catalogue, p76, *ibid*

⁴⁶⁶ Ward, R, The Courtauld Bag: What's in a Name, pp11, *ibid*

⁴⁶⁷ Ward, R, The Courtauld Bag: What's in a Name, pp11, *ibid*

The bag came to the Courtauld in 1966 as part of the bequest of Thomas Gambier Parry⁴⁶⁸. After repair work was carried out upon the bag at acquisition, it remained in the gallery's stores until the current exhibition. In 2012 a workshop was convened at the Courtauld Gallery to discuss the bag⁴⁶⁹. Many scholars contributed their expertise, and from this initial collaboration the idea of the Court and Craft exhibition was born.



Figure 6.2: the bag before and after cleaning. Note the increased specularity in the second image; the initial assessment of the bag's 'scannability' was conducted before cleaning

⁴⁶⁸ Ward, R, The Courtauld Bag: What's in a Name, pp11, *ibid*

⁴⁶⁹ Ward, R, Author's Acknowledgements, pp 8, *ibid*

After a considerable amount of restoration and cleaning, conducted by the Victoria & Albert Museum's metals conservator, Diana Heath, the bag became the focus of the 'Court and Craft: A Masterpiece from Northern Iraq' exhibition⁴⁷⁰. Whilst strongly focussed on the bag itself, the exhibition also included many contemporaneous objects from collections all over the world which helped provide context and illustrate and expand upon the subject matter found on the bag's decoration. Alongside the exhibition, which ran from February to May 2014, a series of academic symposiums⁴⁷¹ and workshops were held as well as public lecture events and a late-night opening featuring Islamic food, music and dance⁴⁷².

6.3 The commission

Having previously collaborated on the Illuminating Objects project, the possibility of modelling the bag was discussed with Dr Gerstein in 2013. The commission was to produce a 3D model of the bag from which a non-interactive video could be produced, to be shown in the gallery as part of the exhibition. Non-interactivity was specified due to the difficulties of supplying an adequate interface within the gallery space, and also the feeling that offering interactivity would entail that only one user (whoever was interacting at that moment) could properly view the exhibit at any one time.

The options for creating the required output were a laser scanned model, using the Arius laser scanner (as used for the miniature bibles in 5), a photogrammetrical model

⁴⁷⁰ <http://www.courtauld.ac.uk/gallery/exhibitions/2014/Court-and-Craft/> (accessed 28/1/15)

⁴⁷¹ <http://blog.courtauld.ac.uk/researchforum/converging-on-the-object-the-courtauld-metal-bag/>

⁴⁷² 6 March 2014, <http://www.artfund.org/news/2014/02/27/museum-lates-in-march>

(using the same techniques as the lusterware bowl, also in 5), a combination of the two, using photographic imagery to add texture to accurate geometry derived from the laser scanner, and imaging in the RTI (Reflectance Transmission Imaging) dome, for close up modelling of some details on the bag.

The initial assessment of the object was conducted before the major cleaning and restoration project, and even then, the shininess of the metal suggested there may be problems with capture. After the difficulties caused by the specularity of the bowl and bible (see chapter 5), it was expected that the reflectivity of the highly polished metal would be a major issue, but again, we adopted a 'see what we can get' attitude to capture – by attempting two methods, photogrammetry as well as laser scanning, we increased our chances of achieving a good result without accruing much extra expense to the gallery.

Due to the bag's rarity and value, it could only be transported by professional art movers at a cost of £500 per day, and the object needed to be accompanied by a member of staff from the Courtauld qualified to handle it. Therefore, to keep costs down the wallet was brought to UCL for one day (approximately 10am to 4pm) for both a test scan with the Arius laser scanner and for imaging under the RTI dome. It was decided that any photography could take place at a later date at the gallery.

6.4 Imaging

6.4.1 Reflectance Transformation Imaging

The bag was imaged by Lindsay MacDonald under UCL's dome (see section 2.2.5).




Figure 6.3: The bag being placed under the RTI dome for imaging (L-R) Dr Alexandra Gerstein (Courtauld), Dr Lindsay Macdonald (UCL), Jack Kettlewell (Courtauld).

With many imaging methods, the bag's shininess presents major difficulties. For example, in photography, it is impossible to capture a 'definitive' view of an object whose appearance changes depending on the relationship between the camera's viewpoint and the direction of light. Similarly a laser scan can be 'confused' by specularity; the colour recorded by the reflected laser differing depending on the exact angle between the laser and surface. However, using RTI these difficulties become strengths; the technique is in fact designed to record the way reflections change with lighting angle, and to render the specular highlights themselves. A new processing and

rendering technique created by Lindsay MacDonald (section 6.7.2.2) allows for an even greater degree of photo-realism than previous RTI rendering methods.

Finally, a note on the areas selected for imaging. There were many areas on the bag which would have been both amenable to the RTI technique, and worthwhile imaging. However, time constraints meant the decision was taken to image four separate areas, chosen by Dr Gerstein (Figure 6.4). The areas were photographed at one of two levels of zoom, governed by the size of the area to be imaged. The selected areas were:

- With a 55mm zoom lens:
 - The entire top of the bag, showing the banqueting scene in full
- With a 105mm lens:
 - The horseman roundel on the rear of the bag, showing the various metals used in the bag as well as areas of extensive damage which reveal some of the techniques used in the bag's creation;
 - One of the roundels on the front of the bag, which is almost entirely obscured by the lid when the bag is closed;
 - Another small detail of a roundel, but focussing mainly on the geometric pattern of the silver wire background, showing minor damage which helps reveal techniques used in the bag's creation.



Figure 6.4; the four areas selected for imaging under the dome. Note, these images are actually examples of the final (rendered) output of the RTI process

6.4.2 Laser scanning

A series of test scans were carried out using UCL's Arius Foundation 150 laser scanner.

The Arius has a scan line point spacing of 100 microns⁴⁸² and a range measurement

resolution of 25 microns⁴⁸³, whilst three coloured lasers⁴⁸⁴ combine to record accurate colour information at every point.

⁴⁸² Since upgraded to 50 microns

⁴⁸³ Arius3D Foundation Scanner Site Specifications Document, issued 17/10/2005

⁴⁸⁴ Red, green and blue at 473, 532 and 638 nanometers



Figure 6.5: Holes in the laser scan caused by excessive reflection




Figure 6.6: Clouds of false points corresponding to holes in the scan

Even before the wallet was thoroughly cleaned and restored prior to the exhibition (Figure 6.2), the predominantly silver object exhibited a certain degree of specularity, and it was unsure how well the scanner would cope with the shiny surface.

With very little occlusion, the object's geometry was captured well. However, there was a significant amount of noise and many areas of particularly high reflectance overloaded the sensor causing erroneous readings. These areas were easy to identify, consisting of a hole or holes in the data accompanied by a cloud of 'false points' (Figure 6.5 & Figure 6.6).

Capturing accurate and consistent colour was, as expected, an issue. The recorded colour at any point was highly dependent upon the angle of the scan, with both the whereabouts of the measured point on the curved scan line, and the overall angle of the scan head to the object surface, affecting the measured value. In Figure 6.7, three separate scans are shown. The first, 'top' scan was of a fairly flat area of the bag, and was conducted with the scan head approximately perpendicular to the surface. While there are certain colour discrepancies, and a few holes caused by excessive specular

reflection, the overall colour of the scan is reasonably consistent. The second 'middle' scan is of an area of the bag with a high amount of curvature, and the scan was conducted with the scan head at 30°, or approximately perpendicular to the bag at its point of greatest curvature.




Figure 6.7: Three separate laser scans showing colour discrepancies caused by the scanning angle

As can be seen in the scan, there is a large amount of specularity along the centre of the scan where the laser hits the surface at (close to) a right angle, whilst towards the upper and lower edges of the scan the colour becomes noticeably darker. This is due to the angle the laser makes with the surface moving further from the perpendicular towards the edges of the scan, partly due to the curved nature of the laser's path, but

mainly because of the shape of the bag at this point (Figure 2.6).

With a more diffuse material, this would not have been such an issue, but the colour

discrepancies caused by the bag's glossy surface would have made it impossible to combine the individual scans into one coherent coloured object. Despite the colour issues, the geometry was captured well enough for us to consider creating a




Figure 6.8: schematic diagram showing angle of laser and surface when scanning curved sections of the bag

colourless model that could then be textured with photos or combined with a photogrammetrical model. Given the size and geometry of the bag, and the time taken to complete the scans during the test, it was estimated that a complete scan of the object would take a minimum of two full days. Due to the costs which would have been incurred for both transport and insurance, and the necessity of having a curator qualified to handle the object present at all times, it was decided to proceed with just the photogrammetry – which could be done on site at the gallery – with the option of returning for further scanning in the future if necessary.

6.4.3 Photography

6.4.3.1 Imaging the exterior

Having imaged the bag under the dome and conducted the test scan, a day was spent

photographing the bag at the Courtauld Institute. The imaging of the object was carried out in an empty classroom using a Nikon D3200 DSLR camera fitted with a 40mm lens⁴⁸⁵ and mounted on a tripod. Images were taken at the highest resolution, 6080x4012, and stored as raw .nef and jpeg files. Lighting was provided by two Balcar studio flashlights fitted with 'soft-box' diffusers.

Once the camera was focussed, the lens was taped in position to prevent any accidental change and all settings were kept constant until after the calibration process (see section 6.5.1.1) was complete. As far as possible the lights were arranged to minimise strong specular effects and provide an even global illumination, with light bouncing off the (neutrally coloured) walls and ceiling. The position of the lights could be adjusted for each image so as to reduce any unwanted reflection; whilst it is a general rule in photogrammetry that the light source should remain constant throughout capture, in this case we were rotating the object rather than moving the camera, so the lighting was moving in relation to the object for each image anyway. A wireless remote control was used to trigger the camera to avoid unnecessary movement caused by manually operating the shutter.

A small cove was created with a piece of thin black card (Figure 6.9), and the object placed on another piece of card marked with targets⁴⁸⁶. The camera was kept in the

⁴⁸⁵ Using ISO 400, f/16 and a 1/60s exposure, shooting in raw. A high f stop was chosen to provide a large depth of field, ensuring as much of the bag as possible (though not all, see Figure 6.1) was in focus. A higher f stop would have risked losing sharpness due to diffraction; light rays diffract due to edge effects as they pass through the aperture and as the aperture gets smaller a larger proportion of the light rays passing through will experience these edge effects, see Rolls, Peter (1968) Photographic Optics, in Photography For The Sciences Ed. Engel, Charles, pp 75; and <http://www.cambridgeincolour.com/tutorials/diffraction-photography.htm>

⁴⁸⁶ The targets are photogrammetry targets which can be automatically recognised by some photogrammetry software. However, in this case they were not used as targets, but merely as added background texture to

same position except for small movements for framing and focussing purposes, and the card with the targets was used as a makeshift turntable, eliminating the need to touch the object and ensuring the targets were stationary relative to the object. Each time the wallet was placed in a new position, it was imaged with an Xrite ColorChecker colour rendition chart⁴⁸⁷.



Figure 6.9: Setup for capturing photogrammetry images – note the colour chart in front of the bag



Figure 6.10: Camera setup, wide view, including the two lights with soft boxes

The wallet was first photographed upright with the lid closed, in three rings of 17-22 images each. The lowest ring was as close to level with the object as possible, the highest at the steepest angle that could be achieved given the tripod and set-up, the middle ring approximately half way between the two. The two upper rings were then repeated with the lid open, and finally the wallet was placed upside down and another two rings of images captured the base. A further set of images, taken perpendicularly

help the photogrammetry software orient the images. As will be seen later, the targets were not necessary, and in fact had to be removed from the images prior to processing

⁴⁸⁷ C. S. McCamy, H. Marcus, and J. G. Davidson (1976). "A Color-Rendition Chart". *Journal of Applied Photographic Engineering* 2(3). 95–99.

to the object's main surfaces and at a 45° angle at the corners were taken for potential texturing purposes. In total, 59 photos were taken of the upright object, 37 of the bag's base, and 30 with the lid open.

Apart from the specularity, the other noticeable problem was a depth of field effect; even on a small aperture setting (f16), keeping the entire bag in sharp focus proved impossible, particularly when taking images along its longer axis.

6.4.3.2 Imaging the interior

As will be explained in the processing section, the pictures taken proved insufficient to properly model the interior of the bag (see Figure 6.11). This was due to a combination of two factors: the angle of the highest ring was too shallow to properly capture the lower section (ie, inside base) of the bag and due to the shape of the bag and the narrowness of the opening, images of the upper parts of the interior were too oblique to capture a lot of detail. Secondly, again due to the concave shape, it was impossible to light the bag's interior consistently without strong shadowing. Considerable time was spent trying to extract a usable model from the photographs available. The photogrammetric process – ie, the generation of the point cloud by the software, would take in the order of five or six hours, and several attempts were made using different sets of images and different settings in the software (Section 6.6.2).



Figure 6.11: First attempt at modelling the interior of the bag through photogrammetry

With the benefit of hindsight and experience, it is fairly apparent from examining the available images that they would never have been sufficient to create a usable model due to the amount of shadowing, the oblique angle of many of the surfaces in the images, depth of field effects causing large parts of the images to be out of focus and the difficulty of capturing all areas of the interior due to the bag's shape. This is certainly a case where someone with more training and experience in (2D) photography and lighting may have achieved better results, though it is also true that the nature of this particular object would still cause difficulties.

Due to these difficulties and the lack of success in producing a model, a second photography session was organised, again at the Courtauld, during which the sole aim would be to image the bag's interior. Since there would be no need to include the open

lid, etc., in this particular model, there was greater flexibility when positioning the bag.

The bag was placed on its side so that the opening was in the vertical plane, meaning that with the same basic tripod, the camera could point directly into the bag. The lights could be arranged to consistently illuminate the interior, and while shadowing could not be eliminated entirely, the amount of sharp shadows was reduced, and since in this case the bag remained stationary while the camera moved, the lighting was at least constant through all the images. However, even with the new arrangement, due to the geometry of the bag (particularly, again, the concave nature of the sides – with a narrow opening and much wider body), the position of the lid when open, and depth of field issues, it was still not a trivial matter to capture all the surfaces of the interior, in particular the upper sides close to the opening, which at best could only be imaged at a very oblique angle.

In total, 31 pictures were taken, but even with the new images, it proved impossible to create a useful photogrammetrical model of the interior of the bag. However, the additional images did prove useful for texturing our 'fake' interior model (see section 6.6.2).

6.5 Processing

6.5.1 Structure from motion model

6.5.1.1 *Camera calibration*

Immediately after capturing the images, the camera was calibrated in order to “determine the geometric camera model described by the parameters of interior orientation”⁴⁸⁸. Photographs taken by any camera fitted with a lens system are subject to a variety of distortions and systematic errors, compared to the image that would be captured if the lens-based camera was replaced by a ‘perfect’ pin-hole model. These errors will cause inaccuracies when the 2D coordinate system of the images is translated into the 3D coordinate system, leading to inevitable inaccuracies in the 3D model. Therefore performing some sort of calibration is important to the photogrammetric process where accurate reconstructions are required⁴⁸⁹, and whilst most photogrammetrical software (including VSFM)⁴⁹⁰ contains automatic calibration algorithms, calibrating the camera and undistorting the images prior to the photogrammetric process allows for more control and greater accuracy.

The calibration process measures six parameters:

- An accurate value for principal distance, as opposed to the lens' nominal focal length

⁴⁸⁸ Luhmann T, Robson S, Kyle S, Boehm J, “Strategies for camera calibration”, Close-Range Photogrammetry and 3D Imaging, 2nd Ed, De Gruyter Berlin/Boston 2014

⁴⁸⁹ Remondindo, F & Fraser, C, “Digital Camera Calibration Methods: Considerations and Comparisons”, ISPRS Commission V Symposium 'Image Engineering and Vision Metrology', IAPRS Volume XXXVI, Part 5, Dresden 25-27 September 2006

⁴⁹⁰ Pierre Grussenmeyer, O. Al Khalil. A comparison of photogrammetry software packages for the documentation of buildings. 2008.

- The principal point, found where the optical axis of the lens system intersects the imaging sensor; it should lie at the centre of the sensor, ie, $x\text{-pixels}/2, y\text{-pixels}/2$.
- Orthogonality – the camera's CCD sensor may not be mounted exactly parallel to the focal plane.
- An affine transformation to correct for non-square pixels in the sensor
- Radial distortion, which occurs when the lens or lens system is not rectilinear; straight lines in the object do not appear as straight lines in the image. Points in the image are shifted either outwards towards the edge of the image (barrel distortion) or towards the centre (pincushion distortion).
- Tangential (or decentering) distortion; a distortion due to the misalignment of the optical centres of the various lens elements.⁴⁹¹

An object covered in 300 numbered, retro-reflective targets in known positions (Figure 6.12) was used for the calibration process. The camera set-up was unchanged from the object imaging session. 14 photographs were taken, eight in landscape format taken at an azimuth of approximately 45° and from the eight cardinal directions (N, NE, E, SE, S etc.). Four more images were taken at the same angle and from the four corners (NE, SE, SW, NW) but with the camera rotated 90° (ie, in portrait format) and finally two images taken from directly above the object, looking down, with the camera again

⁴⁹¹ Beauchemin S & Bajcsy R, "Modelling and Removing Radial and Tangential Distortions in Spherical Lenses", Multi-image Analysis, Lecture Notes in Computer Science Volume 2032, Ed. Goos et al, 2001, pp 1-21

rotated 90° between shots.




Figure 6.12: The UCL engineering department's camera calibration object

It should be noted that the camera was set-up to image an object whose bounding box is 152x220x135mm, while the calibration object's box is approximately five times larger. Thus it is impossible to take photos of the calibration object that both capture enough targets *and* are in sharp focus. However, the calibration software can still calculate the centroid of the targets from fairly fuzzy images, so in this case it wasn't a major issue.

Using Vision Measurement System (VMS)⁴⁹² and a pre-existing file containing the known positions of the targets on the calibration object, the targets in the images are compared to the known locations, and the discrepancies used to calculate the intrinsic camera parameters. The VMS output was then used as an input by Lens Distortion

⁴⁹² <http://www.geomsoft.com/VMS/index.html>

Correction (LDC)⁴⁹³ software to undistort the images⁴⁹⁴.

6.5.1.2 Colour Correction

The raw images containing the ColorChecker chart were converted to a digital negative (.dng) and opened with ColorChecker Passport. This automatically detects the colour patches on the ColorChecker chart and creates a digital camera profile (.dcp) file. The entire set of raw images along with the profile was then imported into Adobe Lightroom 5.2 and the colour correction applied. Manual adjustments to the tone curve were made using the grey patches on the ColorChecker chart⁴⁹⁵, and a small amount of sharpening – just enough to remove some blur, but not enough to introduce noise – applied to the images. Finally, the adjustments were batch-applied to the entire set of images, and the corrected photos exported as tiffs.

6.5.1.3 Creating the models

Initially, three separate models were created, using as their inputs the images of the upright bag with the lid closed, the upright bag with the lid open, and the bag's base (taken with the bag in an upside down position). The plan was to combine the three models – the base with both the open and closed upright models – to create two models of the complete object, one with the lid open, one with the lid closed.

VisualSFM⁴⁹⁶ was used for feature matching and sparse reconstruction. VSFM uses an

⁴⁹³ <http://eprints.ucl.ac.uk/43957/>

⁴⁹⁴ the undistortion was actually carried out after the colour correction; LDC works with tiff files, whilst the colour correction was carried out on the raw .nef files; carrying out the distortion correction as a final step reduces the number of times the images had to be converted from one format to another.

⁴⁹⁵ The grey values used were provided by Ivor Pridden, 2D and 3D imaging technician on the 3D Petrie project.

⁴⁹⁶ <http://ccwu.me/vsfm/>

implementation of David Lowe's SIFT (Scale Invariant Feature Transform) algorithm⁴⁹⁷ to identify and match features pairs in the images, and Bundler⁴⁹⁸, a modified version of Lourakis and Argyros' Sparse Bundle Adjustment^{499,500} to reconstruct the image network and produce a sparse model of the object. VSFM also includes implementations of PMVS2⁵⁰¹ (Patch-based Multi-view Stereo Software) and CMVS⁵⁰² (Clustering Views for Multi-view Stereo) to create the dense point cloud. Attempting to run the entire workflow (SIFT -> bundler -> CMVS-> PMVS2) on a desktop pc with 16Gb ram entailed out-of-memory crashes at the dense point cloud creation stage.⁵⁰³

The output from the VSFM process after the bundler stage but before CMVS was then uploaded to a UCL server with 64Gb ram on which (CMVS) and (PMVS2) were installed. The 'Csize' and 'level' parameters were adjusted in order to achieve the best possible (highest) density for the point cloud given the available resources. 'Csize' (cell size) controls the density of reconstructions, as "the software tries to reconstruct at least

⁴⁹⁷ Lowe, D.G., "Object recognition from local scale-invariant features," Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on , vol.2, no., pp.1150,1157 vol.2, 1999

⁴⁹⁸ <http://www.cs.cornell.edu/~snavely/bundler/bundler-v0.3-manual.html>

⁴⁹⁹ Lourakis, M & Argyros, A, "SBA: a software package for generic sparse bundle Adjustment" ACM transactions on Mathematical Software, 2009, pp1-30

⁵⁰⁰ The bundle adjustment treats the input images as bundles of light rays; the adjustment uses least squares to generate a network where all "corresponding (homologous) image rays ... intersect in their corresponding object point with minimal inconsistency" - Luhman et al, 2014 "Multi-image processing and bundle adjustment" pp323

⁵⁰¹ <http://www.di.ens.fr/pmvs/>

⁵⁰² <http://www.di.ens.fr/cmvs/>

⁵⁰³ Other options selected for VSFM were: -> enable GPU -> set max Dim -> 6400, which ensures the input images are not sampled and used at their full size (4000x6400); -> enable GPU -> customised param -> -nomc, which removes the GPU memory cap for the SIFT process; SfM -> more functions -> Set Fixed Calibration -> (fx, cx, fy, cy) where fx and fy = focal length in pixels and cx and cy are the principle points in pixels; these figures are all derived from the calibration process and override the default values which are taken from the exif information or estimated from the image size. Finally SfM -> more functions -> use Radial distortion (VSFM's inbuilt undistorting algorithm_ was turned off, as the images had already been corrected.

one patch in every $csize \times csize$ pixel square region in all the target images)⁵⁰⁴. Lower numbers result in denser models; the default value is 2. 'Level' specifies the 'level' in the image pyramid used for computation⁵⁰⁵, the image pyramid being a stack of images with each image up the stack being half the dimensions of the one below. Level = 0 ensures that all images are used at their full resolution, while the default value, 1, means all the images are down-sampled to half resolution (2 = one quarter etc.). Clearly, using the highest possible settings, ($csize = 1$ and $level = 0$) would generate the densest possible point cloud. However, even on a server with 64 Gb of ram, attempting to use these settings caused an inevitable crash with out-of-memory errors. The best results came using both $csize$ and $level = 1$. Table 3 below shows the point clouds generated for the bag's base using various settings.

CSize	Level	Points	
2	1	1.7 million	(default)
2	0	5 m	
1	1	7.9 m	(used)
1	0	n/a	(couldn't generate)

Table 3: PMVS settings and subsequent point cloud sizes

⁵⁰⁴ PMVS documentation: <http://www.di.ens.fr/pmvs/documentation.html>

⁵⁰⁵ ibid

6.5.1.4 Point cloud processing: registration

All point cloud processing was carried out using a beta version of Pointstream 3DImageSuite⁵⁰⁶.

6.5.1.5 Scaling

The first task was to register the two point clouds (base and top), but because the two models had been created as individual photogrammetry projects, they each had their own arbitrary scales and coordinate system. Scaling the models proved to be particularly troublesome. Seven pairs of features visible in both models were chosen, and which covered the bag's three major axes. Six to eight measurements were made between each pair of features, with any obvious outliers discarded. A scaling factor for each feature pair (or, more accurately, pair of feature pairs) was calculated by averaging the measurements for each model and calculating the ratio. An overall scale factor was calculated by averaging these results.

The noisiness of the point cloud made making consistent and accurate measurements between features difficult; when zoomed in to make the measurement, discrete and easily identifiable features on the macro scale tended to dissolve into messy clouds of points. Therefore, whilst the calculated scaling factors from each feature pair ranged from 3.595 to 3.560, a range of just .1%, the averaged value still caused an obvious visual inconsistency where the models were scaled and registered in Pointstream. Even after much manual 'tweaking' of the scaling (an iterative process of registering the two

⁵⁰⁶ Beta versions were provided on a regular basis by the developers

models using Pointstream's ICP algorithm, manually adjusting the scale factor in small increments and repeating the registration process), any value used caused noticeable artefacting where the two models overlapped. This is possibly due to the highly regular and high-frequency patterning on the wallet, which accentuated any small discrepancy; no doubt this would have been less of an issue if the wallet featured a natural or random texture.

6.5.1.6 Creating a single model

With the difficulty in scaling the two models, another approach was tried. Each of the colour corrected, undistorted images were edited manually (using free image editing software GIMP), and all of the coded targets in the image backgrounds were painted out. Because of the coving, this resulted in the object appearing to sit on a completely featureless black background.

Initially the five original rings of images, three with the object upright and two of the base were used as an input to VSFM, but the images were still processed as two separate disconnected models. However, when all the available images (the five rings *and* the images intended for texturing) were used, VSFM managed to create a physically impossible 'virtual' network, as if the camera had been able to orbit the model 360° in any direction (Figure 6.13). The output from this network was still processed in PMVS2 as two clusters due to memory constraints, but the resulting models were both scaled and in the same coordinate system, so no more work was needed to register them.




Figure 6.13: Screenshot from VSFM showing the full 'virtual' network. The three rings of images from the 'upright' imaging and the two from the 'upside-down' set can clearly be seen

There were still issues caused by having two models; in both cases the point density was reduced towards the edge (bottom edge in the case of the 'top' model, top edge in the case of the base), and there were colour discrepancies due to the differing lighting of the bag in its two orientations⁵⁰⁷. Some of the colour issues were ameliorated by

⁵⁰⁷ A later model was processed with PMVS installed on UCL's Legion supercomputer, using 100 cores and with 1tb ram available (unfortunately memory reporting was unavailable at the time, so it is unknown exactly how much was used). It used the same VSFM output (ie, the same 'virtual' network), but forced PMVS to use just one cluster. The results were, as expected, largely similar, but the colour discrepancies and reduced point density are not present, or at least not as obvious. Unfortunately, it was too late to use this model in the actual project but certainly worthwhile noting for the future. Whilst this kind of processing power is out of reach of most heritage institutions, the possibility of utilising scalable cloud computing might make it feasible in the future.

simply adjusting the brightness and contrast levels on the base scan so it more closely matched the top. The lower point density around the middle of the bag where the scans joined created a dark stripe, though it should be noted that there is a band around the middle of the real object that is slightly duller and darker than the rest of the bag, presumably due to more wear and tear on the bag's 'extremities' – there is noticeably more silver wire missing in this area compared to the rest of the bag.




Figure 6.14 Top (left) and base (right) models before registration

6.5.1.7 The 'open bag' and interior

The final model created was the top of the wallet with the lid open. Unfortunately, the solution to the scaling problem described above could not be used again, as the base of the model could only be photographed with the lid closed; therefore there would be discrepancies between the two sets of images, those with the bag upright with the lid open and upside down with the lid closed. This issue was less significant, however, as there were only two areas of interest in this model; the underside of the lid itself, which only connected with the other model along the hinge-line, and the area under the front flap of the lid.

To combine the open and closed models, the lid was erased from the closed model, leaving a gap at the top and a large area of missing data where the lid flap had obscured the bag beneath. Conversely, on the open model, the lid and the front of the bag were isolated and the rest of the data erased. A process of iterative registration and scaling was undertaken to combine the two new areas with the closed model. The scaling process was simpler in this case as we were dealing with just the front area of the bag obscured by the lid flap in the closed model, which was effectively a single two dimensional surface⁵⁰⁸, and any inconsistencies were mitigated by merging the two models along the edges of the roundels and other macro features on the front, minimising the areas where the high-frequency patterns on the two models met.

The result was two models, one open and one closed, though at this point the open model did not have a usable interior.

6.6 Point cloud processing:

6.6.1 Exterior

Before any processing, Pointstream's algorithms were used to select and erase 'dark edge points' – those points on the edges of the object where only a portion of the laser spot hits the object and is reflected back, leading to anomalously low RGB values at that point; 'isolated points' – those points situated above a threshold distance from any other points and which are therefore likely to be noise, and 'high incident' points,

⁵⁰⁸ As well as the lid itself, but as that joined the bag along one line only, and the actual point of contact was obscured by a row of hinges, the registration could be considerably 'rougher' than that of the front panel. Note also that the lid was in fact itself a purely two dimensional texture, with no thickness or underside, and thus could only be viewed from above – not, in this case, an issue, as will be covered in the 'output' section

those points which the laser has measured at a raking angle, and which therefore reflect only weakly leading to lower quality data.

6.6.1.1 Occlusion

Occlusion was really only an issue around the rings on either side of the bag, and the underside of the lid flap on the closed model. The flap naturally rests about 5mm clear of the surface of the bag, and so casts a 'shadow' on the surface below (Figure 6.15). This was fixed by using data from the open bag model to fill the gap using a similar iterative scale and register process to that described above. The new data was further blended with the existing data by manually painting in a shadow (simply drawn by hand using a point-darkening paint brush).



Figure 6.15: missing data (left) and model with missing data replaced and manually drawn shadow added

6.6.1.2 Noise & holes

The point cloud was noisy, though this was only obvious when zoomed in to the model. Highly specular areas again caused problems; producing distinct clouds of points above the bag's surface and coincident with holes. While these errors were

similar in both appearance and location to those on the laser scan (Figure 6.5), they were less serious (Figure 6.17).

Even in otherwise 'successful' areas of the model, the point cloud was noisy. Again, this noise was only really noticeable when zoomed in, and would have been an issue if the user had control of the camera. Particularly problematic were the areas of simulated braiding that edge the bag. A combination of specularity (these areas were harder to light consistently than the flatter areas), lack of texture (due to excessive wear on the corners) and occlusion caused considerable noise and some large holes/low point density (Figure 6.16).




Figure 6.16: Close-up image of the braiding on the bag's edge, showing the noise and low & uneven point density. Red points are those that have been selected for erasure

There were numerous holes in the data, though how we define a 'hole' is open to debate⁵⁰⁹. Some were obvious – for example, those on the roundels caused by large

⁵⁰⁹ A hole is a well defined object when dealing with a surface such as in a meshed model, but the description

smooth and shiny areas of polished silver. Here, the lack of texture and possibly specular effects led to obvious gaps in the data (Figure 6.17). Other 'holes' could be characterised simply as areas with a low point density; these areas are dependent almost entirely on the point size selected for the model, and appear and disappear as the point size is raised and lowered: the lower the point size, the more 'holes'.



Figure 6.17: L-R; One of the bag's roundels showing holes; holes selected; roundel after holes were filled by Pointstream's view-based hole-filling algorithm

This can be an issue when different software applications, with different point size options, are used to view the model (for example, section 6.7.1.1). Choosing a point size is a case of finding an appropriate balance; small point sizes lead to a 'cleaner' rendering, with less blurring of textures, but the visible gaps between points can lead to holes and a translucent appearance. Large points ensure there are no gaps, but the shape of splats (usually a square) become obvious and textures get a blurred, pixelated look⁵¹⁰. Both problems are particular noticeable when the model is moved or revolved,

is necessarily arbitrary when talking about a point cloud, which has no real surface. However, as can be seen in figure 6.17, the features in question certainly appear as 'holes' to a visual inspection, and are fixed using Pointstream's 'hole-filling' options, so we will continue to use the term.

⁵¹⁰ Here I am referring specifically to point clouds used to represent 'solid' 3D objects. Other options can create other aesthetic effects which may be more desirable than an illusion of solidity; for example, the

as in the latter case the splats are always drawn orthogonal to the view and thus keep their orientation as the model's changes, and in the former, the translucency and holes mean the back of the object can be seen (moving) through the front, both effects undermining the illusion of a solid 3D object.

Data from, for example, the Arius laser scanner, is ordered, and the points are spaced evenly across the model in a two dimensional grid. Apart from areas where scans overlap, one can be confident that the spaces between points are similar to the machine's resolution, commonly 100 microns. Thus, on a flat area, a point size of approximately 100 microns would provide optimal coverage. The data from photogrammetry does not have a single density, and it is impossible to give a meaningful average point spacing. Thus choosing a point size is a judgement process. We were helped in this instance by the fact that users would not be able to zoom into the model; a point size that works at one zoom factor may be inappropriate at another.

As mentioned above, the point density was also variable, both on a micro scale, with small areas of particularly low density, but also on a macro scale with the area around the middle of the wallet (where the two clusters met) showing a markedly less dense point distribution than other areas.

There were also issues with colour matching, due to the object's shininess and the differing lighting conditions both as the object was rotated, and as the object was placed in its different orientations. Both these effects were most noticeable around the

science museum rendering deliberately uses the smallest possible point size producing a particular translucent 'point-cloud' aesthetic.

middle of the object, where the two clusters met. The point density of each cluster tailed off towards the ‘extremities’ (Figure 6.14), and also the colour in this area was dark, as the lighting direction in both models was predominantly from above. These two effects combined to create a slightly darker band around the bag’s middle – though as noted earlier, there is a similar darker band around the middle of the actual object, probably due to abrasion removing more of the silver inlay and darkening what remained at the point where the bag is widest. This was mitigated by altering the contrast and brightness of the two clusters to achieve a better match, but it is still noticeable. Again, we are helped that the user does not have control, as the darkened areas could legitimately be explained by the particular lighting model applied in the video; if the user had control of the model’s orientation, the lighting artefacts caused during the imaging session might appear more pronounced.

Other problematic areas were the two rings on either side of the wallet; this was the only area where occlusion was an issue (the rings themselves obscuring the bag behind them), and also the rings (added to the bag relatively recently⁵¹¹) are smooth and highly polished, and thus were not captured well. They also suffered some unavoidable movement when the bag was placed upside down, so in effect there were no images of their underside. The area was cleaned up as much as possible using the hole filling procedures, and combining the best data from the various models, but even after processing a close inspection of the area would reveal missing data.

⁵¹¹ Ward, R, “The Courtauld Bag”, Catalogue, *ibid*

6.6.1.3 Pointstream Issues

One of the issues with processing the data was the fact that Pointstream is designed to work with structured data from the Arius scanner, imported via Arius' proprietary file format, .a3d. Whilst models in excess of 30 million points can be worked on with minimal problems in Pointstream, the data imported via .ply files from PMVS2, though possessing a similar number of points, pushed the software to the limit of usability. The problems included frequent crashes (the autosave functionality was used extensively, though the act of saving is one of the operations that could cause crashes!), rendering errors (random or erroneous clipping planes and problems (possibly with z-buffering) where objects would occlude other objects closer to the camera) and last but by no means least, excessive lag. Whilst lag is a nuisance in any software, when dealing with 3D objects it can be a serious issue as it is often only possible to understand a 3D model by manipulating it; severe lag causes a disconnect between the manipulation and the visual feedback causing frustration. These problems were mitigated by isolating small areas of the model, but this does not eliminate the need to view the entire model at times. Performance was improved by moving to a machine with a more powerful graphics card (though less RAM), and while this eliminated most of the rendering errors there was still lag. In short, a model with around 30 million points is the maximum that could be worked on in Pointstream given the available hardware.⁵¹²

⁵¹² A powerful, but not excessively so desktop PC with an i7-4770k 3.5GHz CPU, 8GB ram and a 4 GB GeForce 690 GPU. The PC was also equipped with an SD card for storage, which may have improved performance if a lot of use was being made of the hard drive.

6.6.2 Point cloud processing: Interior

Of all the potential outputs required by the Courtauld Institute, the view of the wallet with the lid open was the most important, as this would reveal areas of the wallet (the underside of the lid, the area on the front of the wallet covered by the lid flap, and the interior) which were not accessible in the gallery (the bag was displayed closed). These areas were also of particular interest as they revealed some details of the manufacturing techniques (for example the fine scoring on the underside of the lid that was used to help attach the bag's original fabric interior), as well as some of the damage done to the object (like that caused by the fitting of a clasp to the front flap) - and subsequent repairs made - over the last 700 years⁵¹³.

A model of the open wallet showing the front and lid had been created, but only small areas of the interior had been successfully – if sparsely – reconstructed from the images (Figure 6.11). After a second attempt at photographing the wallet, concentrating on capturing the bag's interior (see section 2.3.2), the resultant model had better coverage than the first attempt, though the point cloud still covered no more than 30% of the bag's inside surface. Relaxing the constraints in VSFM, for example reducing the minimum number of images in which a point must be visible from 3 to 2 increased the coverage of the model somewhat, but drastically increased the amount of noise (Figure 6.18).

⁵¹³ Court and Craft: A Masterpiece of Northern Iraq, Exhibition Catalogue, Rachel Ward (editor), The Courtauld Gallery in assoc. with Paul Holberton Publishing, London, 2014, pp95-97



Figure 6.18: Interior of the bag, second attempt

The lack of success can be attributed to two major factors, both related to the geometry of the object itself. The bag has a convex, bulbous shape with only a narrow opening, which makes it difficult to both light and image all areas of the interior. A delicate lid which must be placed flat in front of the bag when it is on its side, also cuts down the number of angles from which the interior can be photographed. This particularly affects the sides of the bag near the opening where the point cloud was sparse-to-non-existent (Figure 6.18). These are, unfortunately, the areas most visible in the required video output.

After many attempts at producing a point cloud of the interior, another approach was used. Instead of aiming for an accurate one-to-one representation of the bag's interior, it was decided to use the images captured to texture a 'fake' interior. The point cloud

was imported into GeoMagic⁵¹⁴ software and converted into a colourless mesh. This mesh was automatically ‘healed’ using Geomatics’ built in tools (filling in holes, removing degenerate triangles etc.), then decimated and simplified several times until we had a roughly 10,000 triangle approximation of the wallet’s overall shape.




Figure 6.19: Image of the texturing process. Top left, the current model; Top right, the (cropped) photograph to use as texture; bottom, the model in the process of being aligned with the image

Above: ‘Work in progress’ model showing strips of texture applied from photographs

Images of the interior were then projected onto this mesh (*Figure 6.19*), which was finally imported back into Pointstream and converted into points. This was done partly to ensure a consistency in rendering between all areas, and because some of Pointstream’s tools are designed for working on points rather than meshes. Areas with missing imagery, and obvious discontinuities between areas textured with different images were healed by sampling nearby textures and painting the colourless areas, or blending textures along the joins.

⁵¹⁴ Geomagic Studio: <http://www.geomagic.com/en/products-landing-pages/re-designx-wrap>



Figure 6.20: Model with interior. Despite some obvious defects, areas such as the large patch of new solder at the back, and the nails used to secure the pieces holding the rings can be seen.

This interior point cloud was then shrunk by a small amount (in the order of 1-2%) and, after an initial ICP registration, was transformed manually so it sat inside the wallet (Figure 6.20). The results of this process were not perfect by any means, and the textured interior does not have the same photo-realistic quality as the exterior. However, given the required output, it is good enough at least to give an impression of the inside, and allows the display of the open bag model, showing off the underside of the lid and otherwise-hidden front of the object.

6.7 Output

The output required by the Courtauld was a non-interactive, pre-rendered animation to be displayed in the gallery during the exhibition. Having seen the outputs from both the 3D modelling and the RTI imaging, it was decided the video would consist of a rendering of the open bag, as well as animations showing the details imaged in the dome.

6.7.1 The 3D model

The non-interactive nature of the video meant we would have complete control over what the viewer could see, and therefore could take certain 'short-cuts' with the model that would have been impossible had the user had control and been able to examine the bag from any angle. For example, in the open bag model, only one surface of the lid – the underside, or 'inside' lid – is modelled, but since the viewer only ever sees this surface, this is not an issue and the control of zoom level allowed for the selection of a single best spot size for the rendered video. Some areas, for example the base of the bag, which were never visible in the chosen rendering, could be erased from the model, reducing the size of the point cloud and improving the performance of the software.

6.7.1.1 Attempt 1: ScanLAB and Pointools

The first attempt at rendering a video using the original, closed-bag model was originally outsourced to ScanLAB, who had produced such good results with the Shipping Gallery project (see Chapter 4). There were two reasons for this; one, there

were problems rendering video from Pointstream both in terms of software stability⁵¹⁵ and functionality⁵¹⁶; secondly, outsourcing the rendering would enable work to continue in parallel on the open-bag model. The Courtauld agreed to pay ScanLAB for the production of the video, but it soon became apparent that there were many issues with this outsourcing approach.

Pointools, the software used by ScanLAB and which I thought would be the best option for this project is very good at rendering point clouds 'as point clouds', using a particular aesthetic. The very small points and subsequent 'ghostly', translucent objects, while appropriate for large scale projects such as the Science Museum Shipping Gallery, is less successful for small objects which are subject to a closer inspection. The software has very limited options in terms of point size and point attenuation and thus it is very hard to achieve a balance whereby the object appears fully solid, but not blurry and pixelated.

As the model was processed in one software package and rendered in another there was a disconnect between what was seen during processing and what was delivered to ScanLAB. Areas which would appear solid in Pointstream would have holes and translucent areas when viewed in Pointools, similarly, areas which had been repaired in Pointstream would appear obvious in Pointools due to differences in point density between original and interpolated data. As an added difficulty, these often subtle

⁵¹⁵ The software invariably crashed when rendering video, though it was discovered later that this may have been due to the compression codec used only accepting certain frame sizes (both axes must be divisible by four)

⁵¹⁶ For example, camera paths are limited to simple 'turntable' motions

effects were very hard to detect in a static model, and only when the bag is seen moving (revolving) in the rendered video, do they become noticeable.

The only interoperable format between Pointools and Pointstream was ascii⁵¹⁷, which more than quadrupled the file size of the model to over 2gigabytes. Therefore the only practicable way to transfer data to ScanLAB was to hand deliver a copy on a flash drive. This, combined with the problems mentioned above, meant that any changes required by ScanLAB would have to be carried out and a new model physically delivered to their office. On top of this, the actual rendering time for the video could be anything up to 12 hours, and that problems and issues often only appeared once the rendered video was viewed meant that the time between iterations made editing impracticable.

Due in large part to these issues, the video as originally delivered to the Courtauld was unacceptable; at certain angles the bag still appeared slightly translucent and certain other areas showed small holes. As this was very close to the press launch of the exhibition, the video was re-edited using Windows Movie Maker to include only parts the Courtauld were happy with, and (partly due to Lindsay MacDonald's success with his new technique – see 6.7.2.2) it was agreed that more emphasis would be placed on the videos from the RTI imagery. The final video was up in time for the exhibition launch and was well received by both staff at the Courtauld and early visitors⁵¹⁸, and featured in a Reuters package⁵¹⁹ on the exhibition distributed widely in the Middle

⁵¹⁷ American Standard Code For Information Interchange is a 50 year old standard for encoding character data, using one byte per character. Its advantage is its ubiquity, and therefore interoperability, the disadvantage is the complete lack of compression, hence the large file size.

⁵¹⁸ Pers. comms with various members of Courtauld staff

⁵¹⁹ <http://www.itnsource.com/en/shotlist//RTV/2014/03/12/RTV120314113/>

East. Having the initial video up gave us some breathing space to continue working on the model of the open bag.

6.7.1.2 Attempt 2: Pointstream

Due to the issues with using Pointools, particularly the time between iterations but also the quality of the final output, for the second video (using the open bag model) it was decided to use Pointstream to render the video. As obtaining full video output was still problematic, instead we simply output a sequence of still frames and turned them into video using free software VirtualDub⁵²⁰.

There were several benefits to this method; using VirtualDub to create video from uncompressed TIFF images gave fine control over frame rate, compression and final file size, and using Pointstream to output the original images allowed a ‘wysiwyg’ approach, with each output frame appearing identically as it did in the software.

On the other hand, Pointstream does not have Pointools’ ability to plot arbitrary camera paths, so the animation was limited to revolving the object on a virtual turntable. The animation chosen was a 150° rotation around a vertical axis, with the object angled so as to provide the best view of the interior and lid whilst hiding the missing or incomplete areas. 900 frames (for 30 seconds of video at 30fps) were rendered as TIFF files at hd (1920x1080) resolution, and the subsequent image sequence turned into an avi movie using VirtualDub with no additional compression.

⁵²⁰ <http://www.virtualdub.org/>

6.7.2 RTI Imagery

6.7.2.1 RTI version 1: 'traditional' method

The first output from the RTI imaging was created using PTMfitter⁵²¹ and PTMViewer - free software from Hewlett Packard and Cultural Heritage Imaging⁵²². The fitter creates a .ptm file from the input images (ie, the stack of 64 photos taken with the dome), whilst the viewer provides an interactive application that allows the user to view the object whilst manipulating an arbitrary light source. The image can be cropped by entering values for zoom and x & y panning, while the virtual light can be moved with the mouse or by entering individual values for its x & y position.

The viewer does allow for the export of the current view as a .jpg or .png file, but there is no facility to output video – to create an animation, each frame must be output as an individual image. This process involves first choosing a path for the light source and a combination of zooming and panning to create a smooth movement, selecting the starting and finishing values for the five parameters and calculating the interpolated values for each frame. This entails naming and outputting an image, manually updating five text fields, and repeating the process up to 900 times per video. Each 30 second animation, barring mistakes and errors, taking a minimum of several hours to output.

⁵²¹ PTM (Polynomial Texture Mapping) is the original term for RTI, though now it is considered a specific subset of the RTI technique

⁵²² <http://culturalheritageimaging.org/Technologies/RTI/>

6.7.2.2 RTI version 2: MacDonald's method

The second version to be generated from the RTI images uses a novel method devised by Lindsay MacDonald^{523,524}, an improvement on previous rendering techniques aiming at a more photo-realistic effect, such that frames rendered would be indistinguishable from photographs of the object.

The technique involves creating nine input images from the original stack of 64 images. These record – per pixel – the normals, albedo (both monochrome and colour), specular angle and specular colour and four further parameters governing a Lorentzian function that control in detail how specular highlights are rendered.⁵²⁵ (Figure 6.21)

The normals and albedo are calculated using a photometric stereo technique (MacDonald 2014), while errors which would be introduced by specular reflections and self-shadowing are avoided by sorting the intensity values at each pixel in each of the 64 images and discarding those values above and below certain thresholds, ie, those pixels which are too dark (implying shadowing) or too bright (implying a specular reflection).

⁵²³ MacDonald, L.W, 2014 Colour and Directionality in Surface Reflectance, Proc. Artificial Intelligence and the Simulation of Behaviour (AISB), Goldsmiths College, London, April

⁵²⁴ MacDonald, L., Hindmarch, J., Robson, S., and Terras, M.: Modelling the appearance of heritage metallic surfaces, Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci., XL-5, 371-377, doi:10.5194/isprsarchives-XL-5-371-2014, 2014.

⁵²⁵ MacDonald (2014)




Figure 6.21: Nine input images for MacDonald's RTI rendering algorithm. 1 & 2: Albedo (mono & colour); 3: Specular colour; 4,5,6,7: Specular parameters; 8: Normals, 9: Specular angle

The RGB values for the remaining pixels can be assumed to represent the diffuse colour and intensity at that point, and the normal vector and albedo can be acquired by solving equation (1) for each lamp which falls within the diffuse thresholds (where L_i = a vector representing the intensity and direction of the incident light, L_r = a vector representing the intensity and direction of the reflected light, ρ = the albedo at that point, \mathbf{N} = the normal vector and α is the angle between the viewer and the normal).

$$L_r = \rho L_i \bullet \mathbf{N} = \rho |L_i| \cos\alpha \quad (1)$$

A specular quotient (the ratio of the actual intensity of light reflected for each lamp compared to that which a perfectly Lambertian surface would reflect) is derived for each lamp by comparing the actual intensity values per pixel with what would be expected from a perfectly diffuse material, ie, the value obtained from equation (1). The specular angle is 'almost universally assumed'⁵²⁶ to be double the angle subtended between the viewer and the surface normal. However, this assumes a perfectly smooth surface, whereas real world surfaces are, at the microscopic level, rough and made up of microfacets⁵²⁷, each with its own normal direction. To calculate the *actual* specular angle for each pixel, a sum of all lamp vectors similar to (within a 20° cone of) the nominal specular angle, and weighted according to their specular quotient (ie, the brightest reflections contribute more to the specular angle), is taken.

Each pixel in an image straddles many microfacets, any of which could be facing (have

⁵²⁶ MacDonald (2014)

⁵²⁷ Walter B et al (2007) Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques* (EGSR'07), Jan Kautz and Sumanta Pattanaik (Eds.). Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 195-206

normals) in any direction, and the effect is to 'smear' the specular angle so that light is reflected at twice the normal angle \pm a certain value which depends on the actual roughness of the surface at that point. Examining the actual specular quotient for each pixel at various angles around the nominal specular angle shows the extent of the 'smear'. MacDonald has modelled this smear using a Lorentzian function (similar in shape to a normal distribution, with a relatively strong peak and long tail on either side), which, compared to four previous models using a variety of alternate functions⁵²⁸, has much broader flanks, meaning that the specular component will still make a small but relevant contribution even at angles a long way ($45^\circ+$) from the actual specular angle. The Lorentzian has four parameters governing the sharpness of the peak and slope of the flanks, calculated for each pixel in the image, meaning the actual roughness at each point of the surface is estimated. The colour of the specular component (usually the colour of the incident light, but in the case of metals, the colour of the metal itself) is calculated by measuring the RGB values at each pixel in the image stack and, as above, weighting the values according to the specular quotient of that image at that point; ie, pixels in those images where the illuminating lamp is close to the specular angle, and thus the intensity of the reflected light is higher, contribute more to the RGB value of the specular reflection.⁵²⁹

⁵²⁸ MacDonald (2014)

⁵²⁹ MacDonald (2014)

6.7.2.3 Comparison of RTI rendering techniques

In Figure 6.22, the column on the left shows a sequence of images rendered using PTMFitter and PTMViewer, and on the right, images rendered using Matlab and MacDonald's new method. Both sequences show the same area of the bag, with a light source moving from approximately 30° to the horizontal (from the 'west' direction) to 90° (vertical) in 30° increments. The increased 'dynamism' of the renderings created using MacDonald's approach can clearly be seen; the dramatic change in intensity of the specular highlights as the light moves creates an effect which is both more visually interesting, and more realistic than previous methods.

6.7.2.4 Application to 3D imaging

This rendering technique could potentially be applied to a full 3D model. The extra information could be encoded in textures applied to a 3D mesh, and the required calculations, which after all only require lighting and viewing (camera) angle as inputs, carried out in a shader. However, as the technique works best when imaging a 2D surface, any significant curvature would cause problems. For an example such as the bag, where the surface can potentially be split into individual flat surfaces (details with high curvature, such as the braiding around the edges could be modelled with a more traditional shader model) this may be feasible; objects with more complex geometry would be more problematic, if possible at all.



Figure 6.22: Comparison of RTI rendering techniques. Images in the left column are from PTMViewer, images on the right show the same scene with the same lighting environment, rendered using MacDonald's method

6.7.3 Outputting video

Along with the nine input layers for each of the four areas imaged, MacDonald provided Matlab code to render the output images. By inputting starting and end values for the light path (azimuth and elevation), as well as the number of frames

required, the Matlab code could output a set of sequential images showing a static image with a moving light source. Simple zoom functionality was added by simply modifying the size of the output images – from full frame (3900x2616px) down to HD size (1920x1080). As all images were resized to HD resolution before creating the video, cropping the image from full size down to HD has the effect of zooming in on smaller and smaller areas⁵³⁰. Panning was added by including an offset value for the cropped area. The individual frames were then converted to a video using VirtualDub and output as mp4 videos with minimal compression.

The production of the final RTI videos was an iterative process with Dr Gerstein at the Courtauld. The initial videos produced made the most of the ability to render images with an arbitrary virtual light source. Dramatic effects could be achieved by equally dramatic movements of the virtual light source; starting the light off with an elevation of 0° (ie, with the object in complete darkness) and sweeping it up to the zenith at 90° (with the light perpendicular to the object surface and thus most brightly lit) provides a striking effect of moving shadows and shifting contrasts that evokes a sense of movement that will be familiar from raking light applications.

For the Courtauld, however, the purpose of the video was to allow people to examine areas of the bag in detail, and overly dramatic lighting effects that distracted the viewer, and times when the bag was in complete or near-darkness, defeated that purpose. In other words, the bag should be the star of the video, not the rendering

⁵³⁰ We could, of course, have zoomed in closer by cropping to smaller than hd size; this would, however, entail upscaling the image for hd video output.

technique. Ultimately, we settled on a video format which was used for each area (with minor alterations). We would begin by slowly zooming in from the full frame image to the particular area of interest with a static light source. Once the video was zoomed in, the light source's azimuth would orbit 180° (from east to west) but at no point would the light source's elevation be less than 75° or greater than 105° (ie, it would always remain within 15° of the zenith, ensuring that the object would always be well lit). The effect, though subtle, was enough to provide some movement to the image and to throw into relief some of the details on the bag.

6.7.4 The final output

The final video, as displayed in the Court and Craft exhibition⁵³¹ was 2 minutes 10 seconds long (see table 2) and displayed at 1080p resolution at 29fps on a dedicated media-player PC with a 24" screen. The video was played on an infinite loop.

⁵³¹ The video can be viewed at <http://www.courtauld.ac.uk/gallery/exhibitions/2014/Court-and-Craft/model.shtml>; full screen and 1080hd is recommended

Time	Content
0-30 seconds	The open bag model revolving slowly, rendered from a slightly elevated angle, giving a view of the interior and underside of the lid.
	<fade>
30-40 seconds	A slow zoom into the left hand side of the wedding scene on the top of the bag, with a stable light source
40-60 seconds	A pan across the top of the bag with a subtly revolving light source
	<fade>
60-71 seconds	A zoom into the roundel on the front left of the bag (normally obscured by the lid flap) with a static light source
71-81	Static shot of the roundel with subtly revolving light source
	<fade>
81-92	Zoom into the hunting horseman figure on the front of the bag
92-104	Static shot of horseman with subtly revolving light source
	<fade>
104-115	Zoom into silver wire-work, top right of the bag's reverse
115-130	Static shot of wire work with revolving light source
	<fade>

Table 2: 'storyboard' of final video

6.8 Analysis

6.8.1 Aims

With the finished video playing in the gallery, we evaluated the success of the project.

The research aims were twofold: to examine if the video was used (engaged with) by visitors to the gallery, and to see how successful the video was in providing additional information to the user, by allowing views of the interior (the 3D model) and allowing visitors to examine exterior features in more detail (the RTI renderings). A mixed method approach was taken, involving elements of both quantitative and qualitative research; though tending towards the qualitative end of the continuum⁵³². The quantitative side of the research was aimed at establishing a minimal baseline success metric for the display; how many visitors to the gallery engaged with the video, compared to other multimedia installations and other exhibits in general. A quantitative method, such as an online or anonymous questionnaire (as used elsewhere in this research; chapters 4 & 5) was also considered for ascertaining a level of satisfaction with the video. However, ultimately we are interested in subjective judgements, and capturing phenomenological and experiential responses to the video and so a qualitative methodology was preferred⁵³³. Conducting interviews immediately outside the gallery (as opposed to directing people to an online survey) also meant that the users' experiences were still fresh.

⁵³² Newmam, I & Benz, C R (1998) "Qualitative-quantitative research methodology: exploring the interactive continuum", Southern Illinois University, Chapter 2, "Qualitative and Quantitative Research Methods: An Interactive Continuum"

⁵³³ See: King, N, Using Interviews in Qualitative Research (2004), In C. Cassell & G. Symon, eds. Qualitative Methods in Organizational Research, A Practical Guide. London: Sage, pp. 11–23

6.8.2 Methodology

Before the research began, the methodology was discussed with the Courtauld and I was granted permission to conduct interviews in the gallery. Before each session I would identify myself at the security office, and also with the member of staff on duty in the exhibition itself. At all times I was identified by my UCL Identification Card worn around the neck on a Courtauld lanyard. In order not to skew responses I did not specifically introduce myself or explain who I was to interviewees, apart from saying I was conducting research on the exhibition⁵³⁴.

Research was conducted on five separate occasions, spread over a mixture of mornings or afternoons and between weekdays and weekends (table 4).

	Visitors	Saw & ignored	Saw & Engaged	Total	Engaged with audio exhibit	Interviews conducted
Mon 5 May (pm)	30	6	12	18	2	5
Thurs 8 May (am)	31	10	13	23	5	7
Sat 10 May (am)	65	31	15	46	10	6
Wed 14 May (pm)	27	8	9	17	2	7
Sun 18 May (pm)	41	11	13	24	3	8
Totals	194	66 (34%)	62 (32%)	128 (66%)	22 (11%)	33

Table 4: quantitative research diary

⁵³⁴ Only one interviewee made an explicit connection between my UCL ID and the label next to the installation, asking if I was responsible for the video, though it is unknown how many, if any, made an implicit connection which may have influenced their responses.

6.8.2.1 Quantitative

Over the course of the five sessions (when not conducting interviews⁵³⁵), a record (a simple tally) was kept of the total number of visitors to the Court and Craft exhibition; what proportion of those visitors actually noticed/were aware of the video installation, and of those, how many engaged with it and stayed to watch for a 'reasonable'⁵³⁶ amount of time. Numbers of visitors interacting with the other multimedia installation (an audio recording of the inscription on the top of the bag, played through the two pairs of headphones provided) were also noted.

Certain assumptions and simplifications were made. Groups of two or three people who entered, and stayed together throughout their visit to the gallery were treated as one visitor for the purposes of this research. People or groups who entered the gallery and either walked straight through, or turned and left immediately without engaging with any aspect of the exhibition, were disregarded entirely.

Visitor numbers were sufficiently small that it was possible to identify each of these groups with a reasonable accuracy, even when some time was spent conducting the

⁵³⁵ This will inevitably have introduced some inaccuracies to the results, however, the total time spent interviewing was small compared to the total time spent in the gallery, and as mentioned in the text, visitor numbers were sufficiently small to make it possible to keep track of all the people in the gallery at any one time. Also, any information missed (ie x number of new visitors, y number of people watching the video etc.) would in all likelihood have the same ratios as the information recorded, and therefore would not affect the overall findings which are, after all, concerned with ratios and not absolute numbers.

⁵³⁶ 'reasonable' is not, of course, a rigorously defined period. By reasonable I certainly do not mean that they watched the whole video (very few did remain for the full 2.5 minute running time), and indeed the definition covers a wide variety of actual viewing times; rather, it implies that once the user had ascertained what the video was, they made a decision to remain and watch at least some of it. This is to be contrasted with those whose attention was caught by the video, stopped to see what it was, and having done so, left almost immediately, with total viewing time on the

interviews. During the week there were rarely more than five or six visitors (or groups of visitors) in the gallery at any one time, though this figure was slightly higher at the weekends.

6.8.2.2 Qualitative

Interviews were conducted by approaching visitors as they were leaving the gallery, explaining I was conducting research on the exhibition. Initially a completely random selection of visitors was used, but after the first session, when it became clear that only a fraction of interviewees had watched the video, and little if any useful information was gleaned from those who hadn't, people were only approached if they had been observed engaging with the video. The interviews themselves took place in a hallway just outside the gallery itself or seated in an adjoining exhibition space. No one who was asked declined to be interviewed, though on two occasions the interviews were abandoned when it was discovered that the interviewees spoke no English. No deliberate attempt was made to spread interviews over demographics, rather a random sampling was taken to reflect the exhibition's visitors. Interviews were recorded on a Dictaphone/audio cassette recorder.⁵³⁷

A total of 33 interviews were conducted (Table 4). The interviews were semi-structured⁵³⁸, for a variety of reasons. The audience was extremely heterogeneous (see Results, section 6.8.3) with a variety of backgrounds and motivations for visiting the

⁵³⁷ This method was chosen as whilst it is an old technology, I have considerable experience using the device for interviews. In retrospect it may have been a mistake not using a more modern digital technology; a small number of interviews were lost when one of the tapes – probably due to its age – was destroyed by the machine.

⁵³⁸ King (2004) *ibid*

gallery. Also, I was conscious that people were giving up their time voluntarily, and different people had varying amounts of time they could spare. Semi-structured interviews were considered preferable in that they could be tailored to the interviewee⁵³⁹ with different questioning strategies used for different people; responses could be explored on a person-by-person basis. Also, interviews could be conducted according to the perceived patience of the interviewee.

More importantly, the information we were looking for was experiential and involved subjective judgements. Therefore an open-ended interview format was considered the most effective way of eliciting these responses and avoiding 'leading' the interviewer and steering them towards a particular answer.

Therefore all the interviews began with the questions:

- 1) What did you think of the exhibition?
- 2) Are you visiting the gallery specifically to see this exhibition?

Depending on the responses, for example, if the visitor had come to see Court & Craft specifically, further questions were asked to elicit more information on the visitor's background and reasons for their interest in the bag and exhibition.

After some general conversation about the exhibition, aimed at developing a rapport with, and relaxing, the interviewee, they were asked

- 3) What was the highlight of the exhibition?

⁵³⁹ Zhang, Y. , & Wildemuth, B. M. (2009). Unstructured interviews. In B. Wildemuth (Ed.), Applications of Social Research Methods to Questions in Information and Library Science (pp.222-231). Westport, CT: Libraries Unlimited.

Only then were they asked specifically about the video:

- 4) Did you see the video?
- 5) Did you enjoy it?
- 6) What was it about it that you liked/didn't like?

Depending on responses and time available, the following (or similar) questions were asked:

- 7) Did you think the video belonged/was incongruous in the exhibition?
- 8) Did you realise what you were watching? (ie, an animation as opposed to live action video)

As stated above, in this semi-structured interview format, questions may have been asked in a different order, follow-up questions may have been interspersed, and the questions may have been asked with different wordings.

6.8.3 Results

6.8.3.1 Quantitative

As per the results shown in Table 4, the quantitative survey demonstrates that whilst only a minority of gallery visitors (approximately one third) engaged with the video, this compares favourably with the other multi-media exhibit, is roughly in line with some of the other exhibits in the gallery and commensurate with the 'browsing' behaviour exhibited by the majority of gallery visitors.

To examine the numbers in Table 4 in more detail, during every session apart from Saturday, between one half and two thirds of visitors noticed the video (and, conversely, between one third and one half of visitors to the exhibition gave no indication that they were aware of the video). Of those, approximately half engaged. In

other words, during the week, between one quarter and one third of visitors to the exhibition actively engaged with the installation.

Apart from a small difference in engagement between weekend and weekday visitors (discussed below in the results section), there was no discernible difference between responses from different times and days. Ie, no difference in responses were noticeable between morning and afternoon visitors, or those on Saturday or Sunday or weekdays.

Of those that engaged, only about a quarter watched the entire video (ie, both the model and the RTI output), most staying between 15 and 60 seconds, a period comparable with the time visitors were observed engaging with other individual objects. A very small minority watched the entire video more than once. By means of a comparison, the number of visitors engaging with the only other multi-media exhibit in the gallery, an audio installation featuring a reading, in Arabic, of the inscription on the top of the bag, had far fewer engaged users (11% compared to 32% , - though it should be noted that the audio recording required more pro-active behaviour from the user - they actually had to pick up and wear one of two pairs of headphones to experience the audio; by comparison, simply watching the video could be considered a more passive experience.

While these figures of only about one third of visitors engaging with the video sounds like a small proportion, it fits with the observed behaviour of the audience. The numbers of visitors engaging with the video was, in fact, comparable to any of the individual objects in the gallery. More informal observations suggest that for most of the exhibits, particularly those on the walls (ie, not part of the composite displays in

the large cases in the middle of the gallery), approximately a third of visitors engaged with each one (for example, spent an extended time looking, reading label text etc.). As will be seen from the interviews, the majority of visitors to the exhibition (again, approximately two thirds), had no particular interest in the exhibition or its contents, and were visitors to the Courtauld who were unaware of the exhibition's existence until they arrived and discovered it. As such, they exhibited a definite 'browsing' behaviour in the gallery⁵⁴⁰. The majority of visitors did not view the exhibition in a systematic way and made no particular effort to engage with every object and exhibit in the space. Instead they tended to bounce from object to object exhibiting a sort of random, Brownian motion, following a path that depended on what caught their eye and interest at any particular moment. This would, perhaps, explain the smaller percentage of people engaging with the video on Saturday. One could speculate that at the weekend the Courtauld Gallery experiences a higher proportion of tourists and general visitors. These groups, one would assume, would be less engaged and therefore their sampling behaviour would be even more pronounced.

From observing some individual visitors in more detail⁵⁴¹, it was clear that engagement with the video was proportional to overall engagement. Ie, people who were more engaged with the exhibition overall (measured by total time spent in the gallery, time spent with each object/exhibit, number of exhibits engaged with, time spent reading

⁵⁴⁰ As defined by Graff in: Graff, B. (1994) Visitor studies in Germany: methods and examples. In *Towards the Museums of the Future*, eds. R. Mile and L. Zavala, pp. 75–80. London, Routledge.

⁵⁴¹ During quieter times it was possible to follow individuals around the gallery and observe their behaviour in more detail.

label text etc.) were more likely to engage with the video, and the time spent watching the video was also proportional to their overall engagement. This would support the argument that video displays such as this can from an integral part of a ‘traditional’ exhibition such as Court and Craft, and that this, and other multi-media displays are not just an alternative to traditional exhibits that appeal to non-traditional audience. Indeed, two visitors, who were both extremely engaged with the exhibition as a whole, and spent the longest time in the gallery of all those observed (approx. 20 minutes), both made repeated trips between the object itself and the video, observing something in the video and then going to confirm it or examine it in closer detail in the Bag itself (or possibly vice-versa).




Figure 6.23: rough plan of the exhibition (not to scale)

One further thing to note re. the total numbers of engaged visitors was the video’s position within the gallery (Figure 6.23: rough plan of the exhibition

(not to scale). It was placed near a corner and, from some angles, partially obscured by a wall, and certainly was not easily noticeable from either of the two entrances to the gallery. On the other hand, it was placed quite close to one of the most popular objects in the exhibition, the large and sumptuously illuminated and gilded Koran, and of course, it was the only installation in the gallery which exhibited any movement whatsoever. Without detailed tracking of visitors’ movements, eye lines etc., it is

impossible to say what effect the particular layout of the gallery might have had on users' engagement with the video, though for the reasons stated above, and due to the general 'random walks' and browsing/sampling behaviour of the visitors, it is unlikely that the particular positioning had a huge effect on numbers, though this is speculation and would need further research.

Before moving on to the more strictly qualitative aspects of the research, there are a couple of qualitative observations to be made. Firstly, the effect of what could be called a type of peer pressure, in that, if one person was stood watching the video, it would seem to encourage others, so that often 'clumps' of three or more people would form. This is a well-known effect in museum studies⁵⁴², where, for example, studies have shown that standing a museum employee 'in disguise' next to an exhibit, as if they were examining it, would tend to increase the number of visitors choosing to engage with it. It is unclear whether the same effect was occurring in this particular exhibition with the other exhibits, but a similar clumping was not observed elsewhere. Though this may simply be due to the longer times people spent observing the video compared to some of the other exhibits.

The second observation is again, slightly anecdotal, and possibly too small an effect to be significant, but it certainly appeared during the observation periods that the video prompted more interaction in visitors, both within groups and between individuals and different groups. By and large, the attitude of visitors to Court and Craft was one of

⁵⁴² Falk, J. H., & Dierking, L. D. (1992). The Social Context: Groups in Museums, in *The museum experience*. Howells House., pp41-67

‘silent appreciation’. The gallery space was quiet, and, barring the odd hushed whisper there was very little communication or discussion between visitors. The video was one exhibit, however, that did seem to encourage discussion and communication among the audience, both within and between groups. I would hesitate to infer anything from this anecdotal evidence, but this could prove an area for further research.

6.8.3.2 Qualitative

The audience for the exhibition was extremely heterogeneous, though definitely skewed towards the older end. Visitors interviewed included fashion students, professional artists and historians of the period, all of whom had a specific interest in some aspect of the exhibition. However, the audience was predominantly tourists, general visitors to the Courtauld and ‘passing trade’: people with no particular interest in or previous awareness of the Court and Craft exhibition or indeed, the bag itself.

6.8.3.3 Engagement

Of those who did engage with the video, the reaction was overwhelmingly positive. 93% of interviewees (30/33) expressed a favourable opinion of the video (“there was not quite enough of it”(1)⁵⁴³, “it was sort of shaky, not very smooth, otherwise I think it was perfect”(9)), with the remaining 7% ambivalent rather than specifically negative (“if I can see the original, I don’t need the video – I watch films all day long”(6), “I like to look at the piece. Let’s not deconstruct too much”(8) & “I wouldn’t normally want to watch a video... with a video you can’t control it... you can’t come back to a particular

⁵⁴³ Numbers in brackets indicate respondents

moment. Much better looking, if you can, at the original”(15)).

15% (6/33) of the visitors interviewed were very positive, making unsolicited favourable comments about the video before they were asked any specific questions regarding it (“I wanted to look at the video for longer, if I’m honest, because I found that really interesting”(7)). Of these, a smaller number (7%, 3/33) named the video as a highlight, or the highlight of the exhibition (“The 3D images were a highlight, you can look at the bag closer, look at the detail”(16)) (for comparison, 50% named the bag itself while the majority of the remainder were split between the illuminated Qu’ran and the brass ewer⁵⁴⁴).

Interestingly, only one interviewee was disappointed that the video was non-interactive, whilst one other (a Courtauld graduate student who had attended one of the symposia where the possibility was mentioned) asked if an interactive version would be provided online. The only concrete reason for *not* having interactivity that was given did, in fact, echo the Courtauld’s own reasoning; “you get the problem only one person can do that at a time”(1).

People were not, generally, interested in the technology or methodology behind the video’s creation, unless they had some previous interest in the area. For example, one visitor (another student at the Courtauld, studying fashion) who was “into digital humanities and the use of technology in museums”(7) wished there was more

⁵⁴⁴ The Ewer, a 37cm tall jug from the same location and period as the bag, and showing similar techniques and decoration, and a sumptuously illuminated and gilded Qur’an, 57cm x 80cm (40cm per folio) also from 14thC Mosul – see Ward (2014), pp 142-5 and 161-5

information, but also said that the label probably wasn't the appropriate place for it and assumed that she could find out more online⁵⁴⁵. Another visitor, who had just supervised the installation of a 3D printer in his library in Canada was interested in the potentials of that new technology and how it might relate to the modelling processes in the video. Two interviewees wished there had been more informational content *in* the video, some explanation of what they were watching, either in the form of a voice over or text.

6.8.3.4 Informational content

The two aspects of the video that were most appreciated and most commented on were the ability to see inside the bag (15%) ("it helped to see the inside. You could see the lock wasn't there originally."(10), "We felt when we were looking at the (real) bag we wanted to look inside – we like looking inside!"(13) & "The revolving bag was a good idea so you could actually see the interior"(18)), and the enhanced view of some of the fine details of the bag provided by the RTI animations (33%) ("[the video was...] very helpful, because [the bag] is actually quite difficult to see, with the lighting and the glass..."(14), "you can't really see (the bag) and the video really helps, because so much is in the detail.", "[you can] really see the particulars, see the different shades where the light moves... really a top notch thing to do"(9), "it helped to see the detail"(12), "it was clearer, it did bring out things we couldn't see"(13)).

⁵⁴⁵ Whilst there was a brief explanation accompanying the video on the Courtauld website, there is very little information on the creation of the model.



Figure 6.24: View of the bag on display in Court and Craft (picture courtesy The Courtauld Gallery)

As can be seen from the above comments, a common complaint was that the detail on the real object was hard to make out due to the manner in which it was displayed (Figure 6.24), with visitors mentioning both the glass case (27%) and the lighting (15%) (“it was so hard to see that small [bag] and so hard to get the lighting right ... you need other things that allow the public to see better.” (17)) as reasons. The first of these issues could be applied to any object – particularly a small object like the bag – kept behind glass. As Dr Sussan Babaie, the Courtauld’s Islamic art expert who was involved in the exhibition says, “what you can see of the grooves [in the video] would be impossible without taking [the bag] out of the case”. The second issue, the lighting making it difficult to pick out detail on the bag is, arguably, an unavoidable result of this

particular object's nature. The lighting in the case was subtle and well-planned, and one may, perhaps, draw a parallel between the difficulties in capturing the bag for 3D modelling and the visitors' experiences: What makes the bag such an attractive object – the highly polished and glossy surface and its constantly shifting interaction with light, also makes it hard for an observer to get a clear look at the details. Any slight movement in the observer will cause dramatic shifts in the object's appearance. The RTI renderings, with the subtle lighting effects applied, appear to have addressed this problem, allowing for a close up view of a comparatively stable image, whilst keeping enough of the object's 'real world' specularity to provide a compelling experience.

6.8.4 Caveats

In discussing the results, certain caveats should be noted. Whilst I was aware of the potential of some of these issues before the interviews were started, the overwhelming positivity of the responses, and conversely the lack of explicitly negative comments, was surprising, and implies the results should be treated with some suspicion, and that the unconscious biases outlined below may have had a greater effect than predicted.

The exhibition itself, certainly among the sample interviewed, was universally appreciated, and responses were overwhelmingly positive. Visitors liked the subject matter (the confluence of Mongol and Islamic arts and craft), the selection of objects (including, of course, the bag itself) and the way the exhibition was presented. It is not obvious how much of this general 'good feeling' towards the exhibition would affect the reception of the video, and whether a similar instillation in a less well received

exhibition would have the same positive feedback.

Whilst anonymous in the sense that personal details were not taken, the interviews were clearly not anonymous in the same way that an online survey, or written questionnaire is anonymous. Whilst efforts were taken to elicit unbiased and honest responses from interviewees, the fact that they were asked 'face-to-face' by an interviewer who they would no doubt assume had something to do with the Courtauld and or the exhibition may have had a biasing effect on responses.

Although efforts were made to disguise the real purpose of the interview (ie, that it was, in fact, about the video as opposed to the exhibition in general), once specific questions were asked it presumably became clear fairly quickly. One interviewee, in fact, made the connection between my UCL ID and the label text next to the video, whilst several others asked if I was responsible for the video (though only towards the end of, or after the interview was finished) – fortunately they were complimentary.

Having outlined in some detail the difficulties presented by the object with respect to the creation of a 3D model, and the success of the video in terms of both providing an alternate, inside view of the bag and clarity on some of the details, one question must be asked: Could the same effect not have been achieved far more simply with either photography or video? When interactivity is taken out of the question, what does the laborious creation of the 3D model, or specialist RTI imaging and processing give you that photos or film of an open bag etc. doesn't? In fact, photography and film could give images at a much higher resolution than was achieved via either '3D' method. There are, however, certain advantages, both concrete and ineffable that our process

has over traditional imaging methods.

When time (specifically my own and that of Lindsay McDonald) is taken out of the equation, the cost to the Courtauld of the video was low. However, if capture and processing time *were* factored in, the cost would have been prohibitive⁵⁴⁶. In fact the only expense was the cost of transporting the bag (accompanied by a curator) to UCL for one day, and a payment to ScanLAB for their work on the first video. Commissioning a professional video would have incurred extra expense, but a photographic campaign, with high quality results, *was* carried out for the catalogue; could these photos, perhaps displayed in a loop on a screen similar to the video not have achieved the same ends? The creation of the model does allow an infinite number of animations to be produced (with the caveat that the model was created in the knowledge that certain angles would not be used); once photographs are taken, or a film is made, you are stuck with those angles and with editing that footage. Similarly, once the RTI imagery has been captured, one is free to render an infinite number of scenarios. This flexibility is certainly an advantage of our method over traditional video or photographic campaigns.

As mentioned before, interactivity was specifically not part of the requirements for the gallery installation, though the possibility of an interactive model, similar to those used

⁵⁴⁶ It is difficult to put an exact figure on the total time spent processing the 3D model and the RTI imagery and the rendering of the two videos, but the processing of the 3D model alone took well over 250 man-hours. That is not to say that, with more experience and a certain amount of hindsight the same results could not have been achieved quicker, but it is hard to see that it could be done in much less than 100. Whilst most of the video creation time was simply spent waiting for the images to render, the time spent was nevertheless in the order of 40 hours. It is even harder to put a time on MacDonald's contribution as his method was being used for the first time, but an optimistic estimate for a project of this nature would be somewhere in the region of 150-200 hours in total. At a UCL costing of £20/hr, the total bill would be of the order of £3-4000.

in the Illuminating Objects project, was always there. Unfortunately, in reducing the bag model to a size that could be displayed on the web proved impossible without degrading the quality to a point where the model was unusable⁵⁴⁷. However, this is with today's limitations, and there is nothing to say that in the future both browser technology and web infrastructure will have reached a point where 30 million point models can be served online. It is unlikely that this model will ever go online, nevertheless the potential for interactive applications that 3D modelling provides is still there.

A more ineffable advantage of our processes lies in its novelty; although very little information was provided on the label next to the video, it made it clear that this was a new 3D modelling technique, and it was surprising that every person interviewed realised that this was a virtual rendering and not simply a video – though some believed they were in fact watching video until they read the label. My impression is that people were, if not particularly interested in finding out more, nevertheless impressed in some way that this was something new, and perhaps unique. Dr Gerstein has confirmed that, while the Courtauld Gallery has no particular remit, or incentive to experiment with new techniques and technology in the same way that the Science Museum has, they were pleased to be involved in something new and potentially transformative. There is a feeling that this was a worthwhile project in part because of its novelty. As Dr Babaie says: “everyone loves that video, it takes the breath away. I

⁵⁴⁷ This was attempted by myself, by simply sampling the point cloud (in a similar fashion to that used on the bowl and bible in chapter 5), and also by Anu Rostogi and his colleagues, the developers of the Pointstream.js online point cloud viewer used in previous projects.

think there are people who took note of the reference, because I think it will get people to think of these techniques."

6.9 Conclusions

Again, this chapter demonstrates the potential for cultural heritage institutions to create their own 3D content using predominantly low cost or free resources. In this case, however, it should be noted that the workflow used certain techniques, such as RTI imaging using a custom made dome that may not be available outside of a specialised department⁵⁴⁸. Similarly, cultural heritage institutions would not be expected to have camera calibration objects available, and would therefore be unlikely to undistort images before use in photogrammetric applications. Whether this would be an issue in applications where metric accuracy is not an overriding concern, and how much this method improves on the calibration routines built into photogrammetric software is debatable however, and this step could conceivably be ignored for most, if not all, public facing applications. To reiterate, however, the creation of usable 3D content via photogrammetry is, with the caveat that it can be a frustrating and time-consuming process, well within the reach of small institutions.

The research demonstrates that there is utility in using 3D content inside exhibitions, even when that object is available to view in the same gallery. In this instance, users were able to extract informational content from the 3D model that they could not from

⁵⁴⁸ 'domeless' RTI is possible (see Mudge, Mark, et al. "New Reflection Transformation Imaging Methods for Rock Art and Multiple-Viewpoint Display." VAST. Vol. 6. 2006.), but could be considered an 'advanced technique'. While the equipment may be within reach of small CH institutions, the skills needed may not be.

the original object. It has also shown the usefulness of RTI imaging when dealing with shiny and highly specular objects, and particularly the ability of MacDonald's new method to render photo-realistic animations. We can conclude that the exhibit was successful both as a piece of engaging content and as something providing specific informational content. In the first case, visitors engaged with the video much as they did the other exhibits in the gallery, the multi-media installation was not seen as incongruous within the context of the exhibition and the majority of visitors found the content compelling. In the latter case, the video adequately fulfilled its dual roles of allowing visitors to see the bag's interior and enabling them to inspect some of the bag's features in greater detail.

Examining the concept of aura in relation to this exhibition, the bag stands in contrast to the Shipping Gallery in that one could argue that the bag's particular aura, the properties that make it unique and valuable are very much its aesthetic qualities. Whereas we saw that much of the Shipping Gallery's affectual power is derived from its context and the particular relationship users had with it, the bag is very much a beautiful object in and of itself. That is not to say it does not receive some aura from context; it is the centrepiece of an exhibition based entirely around it; the exhibition makes it clear that it is extremely rare in terms of its physical properties, its cultural origins and, despite its immense age, its condition. Most comments, however, referred to the object's intrinsic beauty rather than its context (though many visitors praised the look at early Islamic art, this praise was aimed at the entire exhibition rather than the bag itself). The beauty is almost entirely derived from the complex interplay between

light and surface, and in these circumstances, as we saw in section 3.4, the 3D model has a particularly difficult task acting as a digital surrogate. However whilst there was no real indication that the model inherited aura from the original, or was capable of affectual power, this was not its purpose. The original object can carry that burden; the model, as indicated above, serves to increase the informational content users could extract from the Bag.

However, it is also possible that the rendered video could succeed in its stated aims and yet the project itself still be considered unsuccessful. We have discussed the possibility that the animation could potentially have been replaced by a simple filmed video which may have had the same effects for less expense (if expense had in fact been incurred in this project). For particularly difficult objects, and outputs that require a very high standard, more research needs to be conducted into the cost-effectiveness of 3D digitisation.

7 Conclusions and further work

7.1 Motivations

This research has several motivations. Firstly, the increased accessibility of low cost 3D digitisation solutions and the ability to easily serve 3D content over the web has created the situation where any cultural heritage institution with limited budget and no pre-existing skills or experience has the potential to create and share their own 3D digitised content. Part of this research intends to ascertain how realistic a proposition this actually is, and what issues may arise along the way.

The second major strand of this thesis is a response to the lack of existing research into the utility of 3D digitised content in public facing cultural heritage applications. As we saw in the introductory chapter, there are good intuitive reasons and circumstantial evidence to support the belief that 3D content may help to further a museum's public facing remits of access, education and entertainment. However, without research into the potential benefits of 3D digitisation we risk falling foul of 'technological determinism' and simply adopting new technology because it is possible.

7.2 Thesis structure and chapter conclusions

Chapter 1 introduced the thesis and outlined the major research questions. **Chapter 2** reviewed existing technologies, their historical use in cultural heritage applications and in particular the most recent developments in 3D digitisation and associated technologies. It concluded that photogrammetry has many advantages over other digitising techniques; in cost, requisite hardware and experience, and its ability to produce fully textured models that are suitable for use in public facing CH applications.

But whilst photogrammetry will probably be the method of choice for cultural heritage institutions, each project, and indeed each object must be treated on a case by case basis and there is no single technology that will be fit for all potential purposes. In fact, issues which preclude one particular scanning technology may well preclude all of them as there are many surface and material properties that cannot be captured by any of the traditional methods. However, recent developments in BRDF and BTF capture have produced impressive results and whilst not accessible to non-specialists at the current time, these technologies may become a viable alternative in the near future.

Chapter 3 examined the concept of a digital surrogate, and what it means for a 3D model to be ‘successful’. It determined that authenticity was the concept most applicable to public facing applications, and that a convincing 3D object was more important than an attempt at objective accuracy. It also looked at the ‘aura’, or affectual power of museum exhibits: concepts which at first glance appear to be inimical to virtual objects. When examined closer, however, aura is seen to be a property that could, in fact, be shared by a digital model and is much a product of context as of the object itself. This chapter also made the argument that, therefore, virtual objects or virtual exhibits should be treated in a similar way to the ‘real thing’ in terms of the context in which they are presented to the user or visitor.

Chapter 4 established that long-range terrestrial laser scanning was a viable method for recording entire galleries or exhibitions and that there is a public appetite for this type of digitisation. Unlike the two subsequent projects, however, the capture and

processing of this model required professional expertise and this type of ambitious project may be beyond the resources of most CH institutions.

It also revealed that, whilst the video was clearly appreciated by the majority of the audience, an interactive application would be preferable. However, interactivity may not be feasible with the output from this sort of capture, and it may be necessary to manage the expectations of an audience not familiar with this kind of technology. The research found that whilst divisive, the point cloud aesthetic may be a suitable (and in some cases, preferable) alternative to photo-realistic rendering, and that, when presented in the proper context, virtual models can exhibit affectual power and elicit emotional responses. There is an argument that some of the positive response to this approach may be down to the novelty and 'exotic' nature of the procedure (see the tweets and comments in appendix A, and the number of people who excitedly mention 'lasers!'). It will be useful to observe how users' reactions to this type of scanning change, should it become more commonplace.

Chapter 5 shows that it is possible to render a point cloud model using webGL that is small enough to offer a good user experience yet detailed enough to be useful. This chapter demonstrates workflows for laser scanning and photogrammetry, demonstrating that low or no-cost solutions are both possible and comparable to traditional methods. This chapter also highlights potential issues with objects that have features that make capturing them problematic, for example shininess and other complex material properties.

The survey results, whilst too small a sample to draw any completely concrete

conclusions, do allow us to make some tentative suppositions. There is some evidence that 3Dmodels do aid understanding of complex and unusual or unexpected geometry better than 2D images, and that the models are engaging and considered useful by the audience. Photogrammetric and laser-scanning methods were compared, and the outputs from the relatively low-cost photogrammetric process shown to be comparable to those from the expensive scanning.

Chapter 6 showed a more comprehensive and involved workflow for photogrammetry projects, and further builds on the previous chapter in its coverage of capturing difficult objects, in this case a highly specular polished metal bag. The workflow for the Courtauld Bag demonstrates the attention that needs to be paid to all aspects of a project, from capture through processing and through to dissemination. In this case the desired output was the animation of a pre-rendered video, a process which proved to be far more complex and time consuming than had initially been envisioned.

The use of 3D models in-gallery, even in the form of pre-rendered video has been shown to have utility for visitors, both in understanding the exhibit and revealing views unavailable in the normal gallery display.

7.3 Overall conclusions and reflections

This research demonstrates that someone with little or no experience in any 3D (or, indeed, 2D) imaging technology can, using free software and relatively cheap or common hardware, create models of sufficient quality for public facing applications. These models, whether interactive or in video form, used on a website or inside an exhibition are at least acceptable and potentially useful to visitors. It has also

demonstrated at least a *prima facie* case that there is considerable appetite for the use of digitised 3D content among the public and that the use of such models may be a cost-effective way of fulfilling an institution's public facing remits.

7.3.1 Reflections

In retrospect, the thesis shows a certain amount of naivety on my part as to what is and isn't feasible in 3D capture and processing. This led to difficulties processing some of the models, primarily the Courtauld Bag, and whilst this was not necessarily a problem for this particular project – as my time was effectively free and the final output was acceptable – the inability to successfully predict how long a project will take would have had cost ramifications in a live project. Ultimately, if personnel costs had to be taken into account, it may be that the time and expense for the delivered output was not cost effective. This stresses the importance of expectation management, but also that the process of creating authentic 3D models is as much a craft or artistic process as a scientific one, and that whilst I have stressed that it is indeed *possible* for someone with little or no experience to create a 3D model, experience can make the whole process simpler, faster and more effective.

Unfortunately, the heterogeneous nature of CH objects and each one's own particular quirks and the difficulties they present mean that familiarity with one object or class of object does not automatically entail familiarity with all. The same can be said for the different imaging technologies, and indeed, the required output.

During the course of this thesis, further attention could have been paid to the user-research required on each model. Once one digitising project was complete, the

research moved on to the next believing that the user-research would, in some sense, take care of itself. Again, this was naïve and more attention should have been paid to obtaining user feedback, particularly on the Illuminating Objects projects. With hindsight, I would have conducted more structured research to supplement the online surveys and used combinations of research methods such as guided user-testing, interviews and observational studies to ensure a range of data was collected on these projects. However, the data presented here – extracted from online surveys, and online comment, are similar to the types of feedback sought by those delivering cultural heritage online, and so do reflect the types of analysis institutions undertaking similar projects may pursue.

7.3.2 Findings of use to cultural heritage institutions

1. Many, if not most, cultural heritage objects will exhibit some properties which make them difficult to capture and/or render.
2. While it may still be possible to create a useful 3D model from difficult objects, the degree of difficulty will impact on both the time required to process the model and the final quality of the model.
3. Processing data from raw point cloud to a finished model is the most difficult and time consuming part of the workflow.
4. Whilst huge steps have been made in automating the capturing process, software used to process the raw data is still hard to use and does not have the necessary features required to easily create authentic and aesthetically pleasing models.

5. Deficiencies in the final model can be hidden by controlling the interaction the user has with the object, for example by limiting the level of zoom available or creating a pre-rendered animation that avoids the worst areas. However, it has been shown that users are willing to forgive or ignore minor imperfections in 3D models.
6. Some decisions, particularly those made in the rendering phase of the project, will come down to subjective, aesthetic choices. It is important that in these cases the digitiser works closely with the curator or exhibition organiser.
7. Whilst the ease with which models can be created has increased dramatically in the last five years, the quality of the outputs has increased less dramatically; whilst continuing improvements computer technology allow for ever bigger models with more points and therefore more detail, the ability to easily capture complex surface properties is still some way off.
8. However, as recently as 2011 when this research began the original title referred only to laser scanning as a means of digitising in 3D. Just four years later, I believe it is fair to say that, at least in the field of public facing applications, photogrammetry has superseded laser scanning as the technology of choice. Futurology is an inexact science, but methods do already exist to capture complex optical properties and these are being refined all the time; there is a distinct possibility that the ability to record and render shiny and other 'difficult' objects with the ease with which today we can record diffuse colour may be just a few years away.

9. 3D digitisation, particularly for public facing applications, is a craft process.

Whilst some steps in the workflow, such as point cloud registration or hole-filling, can be automated to a certain extent, a lot of detailed manual work is still required to create authentic 3D models. As a craft process, performance inevitably improves with practice.

10. Further to the previous point, I have stated repeatedly through this thesis that

people with no previous experience can use some of the techniques discussed to create successful models. However, it should be noted that while I had zero previous experience coming into this research, I was nevertheless embedded within a research group made up of many individuals with huge amounts of experience and who were always generous with their help and advice. Not every potential 3D digitiser will have access to that sort of expertise, but there are a large number of online resources to aid them, including user groups, discussion boards and video tutorials.

11. Having said that, there will be some projects that, due to particularly difficult

objects, scale, or hardware requirements, will require the input of experts or professionals. The Science Museum project in this research is a case in point. Whilst there is nothing to stop anyone with access to a terrestrial scanner (or who can hire one) conducting such a project, the sheer scale and complexity of the task, the amount of data collected and the degree of processing required for the quality of the desired output may necessitate potentially expensive professional assistance. As the technology matures and the appetite for this

sort of project increases, one can imagine that the price will fall, however.

12. In 2D digitisation, institutions are encouraged to scan “once for all purposes”⁵⁴⁹.

This is for a variety of reasons, for example, to keep the costs of digitisation down or to minimise the handling of the original object. While this should also be a *theoretical* aim for 3D digitisation, in practice, due to the choices and compromises that need to be made in the capturing process this will be unachievable for the vast majority of objects. That is not to say that, for example, an archival model cannot be repurposed for use in a public facing application, simply that it may not be practical for a single digitised model to be the ‘best possible version’ for multiple purposes.

13. Methods used to assess the success of a model must pay attention to, and reflect, the specific purposes of that model. Whilst online surveys are useful up to a point, it may be possible to obtain more objective data. For example, for models aimed at increasing engagement, time spent on the website and other common web analytics could be used (perhaps compared to similar 2D content), while for models aimed at conveying information or imparting knowledge, specific tests to measure the efficacy (again, compared to similar information presented in 2D) could be conducted. Thinking about how success metrics could be measured at the first planning stages of a 3D digitisation project may in fact help to define the model’s purpose and thereby inform the

⁵⁴⁹ Terras, M. Encyclopedia of Library and Information Sciences.

digitisation process itself.

14. Expectation management is an important part of any 3D digitisation project. It is easy to only show the most successful examples when demonstrating potential outputs to those interested in using the technology. These models, for example the outputs from CultLab3D or Bonn's Dome II technologies, or even the 3D models on the Smithsonian website may have been created by teams with far greater resources than are actually available, and quality of this level may be beyond the capabilities of the low budget, in-house digitisation teams we have focussed on in this research. This does not mean useful models cannot still be produced, but particularly with difficult objects, it is important to have – and communicate – realistic expectations for the final output. Otherwise, a disconnect between what is expected (or promised) and what is delivered can cause issues which could potentially damage the uptake of 3D digitisation in the future.

15. Sustainability must be considered in any public facing application, particularly those relying on relatively new technologies such as webGL. As detailed in chapter 5, a proprietary renderer (Pointstream.js) was chosen over an open source alternative (XBPS) for the Illuminating Objects project for several reasons. It was offered free of charge, it was fully integrated into the workflow of the 3D processing software being used, and it offered greater functionality than the open source version. However, the company providing the software went out of business soon after the launch of the project, and, due to minor

technical changes in web browsers, functionality began to degrade soon after.

Within a few months of support ending, the renderer stopped working altogether.

16. For the cultural heritage institution, a certain amount of programming knowledge, both in JavaScript and the WebGL API would be required to integrate an open source solution into their website. Whilst this would provide some defence against obsolescence, it would still require someone (presumably the open source community) to constantly update the code in order to keep it fully functional as browsers evolve and change. Open source and proprietary solutions both have advantages and weaknesses, and again, the decision is something that needs to be made on a case by case basis depending on the particular functionality required, budgetary constraints and the reliability of both the particular open source community or company providing the solution.

7.4 Further work

Cultural heritage institutions are extremely diverse, and thus the difficulty of making broad generalisations from just a few digitisation projects should be acknowledged. What works at a particular scale may not work at another, one class of objects may be more amenable to digitisation than another, what suits one institution might not suit another. Until research is conducted on a wide variety of projects, and the results aggregated, we should be wary of making definitive pronouncements on the utility of 3D digitisation in general.

One of the defining features of much cultural heritage is its affectual content and ability to evoke powerful emotional responses. Whilst the fact that this same emotional response can be evoked by virtual representations is an important finding and deserving of further research on its own, it is also something that needs to be taken into account when researching other aspects. For instance, it is clear that in the case of the Science Museum, people's feelings towards the Shipping Gallery and its demise coloured their responses to the video (though in which direction is less clear). Similarly, the video of the bag was placed in a popular exhibition which, with very few exceptions, visitors were enthusiastic and appreciative of. This general 'feeling of goodwill' towards the exhibition as a whole may have affected people's opinions of and responses to the video. Whilst we have made the case that it is impossible, and also pointless, to divorce the digitisation from the context in which it is presented, only more research on a variety of projects in a variety of situations will help us reach objective conclusions on 3D digitisation as a technique.

As well as the heterogeneity of cultural heritage, we should also address the heterogeneity of our audience. We have talked about 'public facing' applications and 'the general public' without really defining who or what is being talked about. We saw in the case of the Shipping Gallery, and of the bowl and bible, that prior interest in and knowledge of the original object will affect the reception of the digitised version. Some stratification of the audience, via broad demographics or by interest is necessary to determine who or what will gain the most from 3D digitisation. This in turn will help CH institutions target their digitisation projects to extract the maximum benefit and fulfil their own particular aims.

A more functional approach to research would also be helpful in this respect. Rigorous testing of the informational content of 3D models, and their effectiveness in conveying that information could be tested and compared to traditional 2D methods. How 3D models are used by different groups – school pupils conducting a project compared to academic researchers or merely interested amateurs – could prove a fruitful line of enquiry.

Finally, there are clear differences between the projects in terms of the use the 3D content can be put to, specifically *when* it is used. The Shipping Gallery project, by its very nature, could only be used once the gallery was gone (or at least going); the Illuminating Objects content was designed to coincide with the Illuminating Objects programme, and more specifically to appear on line while the object was on display in the gallery (approx. three months per object); and the Courtauld bag content was created specifically to appear in the gallery for the course of the exhibition. The purposes behind the three different presentations was therefore necessarily different. The Shipping Gallery project's aim was to preserve and commemorate a now dead gallery space, the Illuminating Objects to provide information to and engage online visitors, whilst the Courtauld bag video was designed to enhance an exhibition and provide informational content (and also, perhaps, engagement) to gallery visitors. Success criteria for the different projects will therefore differ, and further research needs to take place to determine when and where 3D content should be used to extract the most utility. However, an area that will be of particular interest to cultural heritage institutions is the effect digitisation projects have on the physical museums.

Does the provision of 3D models encourage or discourage visits to see the originals?

Can a point cloud model of a gallery persuade people to visit or help plan their museum trip, either increasing visitor numbers or improving the visit itself (and therefore the potential for repeat custom)? Users of the 3D models in the Illuminating Objects programme said in the survey that seeing the models would make them more likely to visit the physical exhibition, but whether, in fact, this is the case would need a more in depth study with greater numbers, and preferably some solid quantitative data rather than vague assertions of intent. Obviously, if it transpires that users of online 3D content are less likely to then visit the physical object will have ramifications for the institution.

In a more abstract sense, does the physical embodiment of the museum assume less importance if it can fulfil some of its public facing role online? Can the institution fulfil its remits, presumably more cheaply, to a wider audience and with less risk to the physical objects, using purely virtual content? And if the physical instantiation of the museum assumes less importance, does this then have a circular effect on its perceived authority, and thus on the amount of aura it can impart to its objects, both real and digital? These are deep and complex questions in museology, and beyond the scope of this thesis, but they may have important implications for the future of museums and cultural heritage.

7.5 Summary of Conclusions

This research has successfully investigated a variety of 3D digitisation projects used in public facing applications by cultural heritage institutions. The Science Museum's

Shipping Gallery project demonstrated a novel method for recording, preserving, and potentially publicising, exhibitions at the gallery scale. Research on the video output, including surveys and comment analysis, reveal a real public appetite for this type of presentation.

The Illuminating Objects programme showed the workflow required to digitise objects and put them on the web as interactive models, and demonstrated that low cost photogrammetric solutions are equally as effective as traditional laser scanning methods. User research, though drawn from a small sample size, does draw tentative conclusions that interactive models such as those provided for both bowl and bible can aid understanding and provide an engaging experience for museums' online visitors, helping to fulfil their public facing remits.

The Courtauld Bag project again shows the potential of low cost digitisation solutions, though it also demonstrates the issues encountered when attempting to capture and process particularly difficult objects. It also details the first use of a novel RTI technique developed by Lindsay MacDonald which is shown to be a viable alternative for capturing and rendering objects with high degrees of specularity. The interviews conducted with gallery visitors to the Courtauld Institute show that 3D models created in this way can have utility even when displayed non-interactively, and alongside the original object as part of an exhibition.

The use of 3D content, virtual or otherwise, by cultural heritage institutions has been shown to have a long history. But in the last three or four years, and for the first time in history, it has become possible for an individual or small team to create and share their

own digitised 3D content to a huge audience. The models that can be created are of sufficient quality that, for visual inspections in public facing applications, they can rival or in some cases even surpass interaction with the original object. Institutions that hold items of cultural heritage in trust for society have an ethical obligation to make their collections accessible to the public, and to both educate and engage their audience. 3D digitisation has the potential to dramatically change how these obligations can be fulfilled, and revolutionise the way people access and interact with their own heritage. More research is clearly needed in order to confirm these potentialities, but the evidence collected so far is enough to suggest that 3D content will prove to be an innovative and cost-effective means for fulfilling public facing remits. The research in this thesis will hopefully encourage cultural heritage institutions to embark on their own 3D digitisation projects, whilst allowing them to make informed decisions regarding all aspects of the technology and workflow. It will also, just as importantly, enable them to make realistic assessments of the potentials of their 3D content, thus extracting maximum utility from their outputs.

If 3D digitisation of our heritage can be shown unequivocally to have utility, the benefits to professionals and institutions are potentially enormous. More importantly, perhaps, are the benefits to society as a whole.

8 References and Bibliography

Devrim, A, Grun, A, Breuckmann, B and Lahanier, C (2007) High definition 3D-scanning of arts objects and paintings. Institute of Geodesy and Photogrammetry, ETH Zurich.

Acka, D, Grun, A, Breukmann, B & Lahanier, C (2007) *High definition 3D-scanning of arts objects and paintings*. Institute of Geodesy and Photogrammetry, ETH Zurich.

Adams, CB, (1893), Method of Photogrammetry, US patent US 510758 A, Dec 12 1893

Agarwal, S., Furukawa, Y., Snavely, N., Simon, I., Curless, B., Seitz, S. M., & Szeliski, R. (2011). Building rome in a day. *Communications of the ACM*, 54(10), 105-112.

Ahmadabadian, A. H., Robson, S., Boehm, J., Shortis, M., Wenzel, K., & Fritsch, D. (2013). A comparison of dense matching algorithms for scaled surface reconstruction using stereo camera rigs. *ISPRS Journal of Photogrammetry and Remote Sensing*, 78, 157-167.

Akca, D., & Gruen, A. (2009). Comparative geometric and radiometric evaluation of mobile phone and still video cameras. *The Photogrammetric Record*, 24(127), 217-245.

Albertz J, 2001. Albrecht Meydenbauer - pioneer of photogrammetric documentation of the cultural heritage. Proceedings of the XVIIIth CIPA Symposium (eds. J. Albertz), September 18 - 21, Potsdam, Germany

Algharabat R & Dennis C (2010) "3D Product Authenticity Model for Online Retail: An Invariance Analysis", *Int. Journal of Business Science and Applied Management*, Vol 5, Iss 3

Algharabat, R. , Dennis, C. , Morschett, D. , Rudolph, Th. , Schnedlitz, P. , Schramm-Klein, H. , Swoboda, B.(2010c). 'Modelling the impact of 3D authenticity and 3D telepresence on behavioural intention for an online retailer'. *European Retail Research*. 24, 2, Gabler Verlag, 93-109

Alelis, Genevieve, Ania Bobrowicz, and Chee Siang Ang. "Comparison of engagement and emotional responses of older and younger adults interacting with 3D cultural heritage artefacts on personal devices." *Behaviour & Information Technology* (2015)

Alkis, Z. (2006). *3D modeling of the Weary Herakles statue with a coded structured light system*. Institute of Geodesy and Photogrammetry, Swiss Federal Institute of Technology.

Aloysius Wehr, LIDAR: Airborne and terrestrial sensors (2008), in Advances in Photogrammetry, Remote Sensing and Spatial Information Sciences: 2008 ISPRS Congress Book, Li, Chen and Baltsavias (Eds), Taylor & Francis Group, London, pp80-81

Eike Falk Anderson , Leigh McLoughlin, Fotis Liarokapis, Christopher Peters, Panagiotis Petridis, Sara de Freitas, (2010) Developing serious games for cultural heritage: a state-of-the-art review, December 2010, Volume 14, Issue 4, pp 255-275

Armesto, J., Arias, P., Ordóñez, C., Lorenzo, H., & Caparrini, N. (2007, September). Damage quantification and monitoring in masonry monuments through digital photogrammetry. In Key Engineering Materials (Vol. 347, pp. 291-296).

Arms, Caroline, and Carl Fleischhauer. "Digital formats: Factors for sustainability, functionality, and quality." Archiving Conference. Vol. 2005. No. 1. Society for Imaging Science and Technology, 2005.

Arnold, M., Ghosh, A., Ameling, S., & Lacey, G. (2010). Automatic segmentation and inpainting of specular highlights for endoscopic imaging. *Journal on Image and Video Processing*, 2010, 9.

Baldock, C, Hughes, S W, Whittaker, D K, Taylor, J, Davis, R, Spencer, A J, Tonge K, & Sofat, A (Dec 1994), 3-D Reconstruction of an Ancient Egyptian Mummy Using X-Ray Computer Tomography, *J R Soc Med* 87: 806-808, doi:10.1177/014107689408701229 Bearman, D (2011) 3D Representations in Museums, *Curator: The Museum Journal*, Vol 54, Iss 1, pp55-61, Jan 2011

Beauchemin S & Bajcsy R, "Modelling and Removing Radial and Tangential Distortions in Spherical Lenses", Multi-image Analysis, Lecture Notes in Computer Science Volume 2032, Ed. Goos et al, 2001, pp 1-21

Becerik-Gerber et al, Assessment of target types and layouts in 3D laser scanning for registration accuracy, *Automation in Construction*, Volume 20, Issue 5, August 2011, Pages 649-658

Benjamin, Walter. "1939. The Work of Art in the Age of Mechanical Reproduction." *Media and Cultural Studies. Key-Works* (1936): 18-40.

Bennett J, 1997, 'Kama and Eroticism: The Five Senses in the Work of Francesco Clemente and Pierre Klossowski', in *Body*, Bookman Schwartz, Melbourne, Australia, pp. 140 - 150

Beraldin et al, (2010) Laser Scanning Technology, in Airborne and Terrestrial Laser Scanning, Vosselman and Maas (eds), Whittles Publ., Dunbeath, p3

Beraldin, J. A. (2004). Integration of laser scanning and close-range photogrammetry-the last decade and beyond. In International Society for Photogrammetry and Remote Sensing.

Beraldin, J. A., Picard, M., El-Hakim, S., Godin, G., Borgeat, L., Blais, F & Bandiera, A.

(2005). Virtual reconstruction of heritage sites: opportunities and challenges created by 3D technologies

Andrew Bevan, Xiuzhen Li, Marcos Martinón-Torres, Susan Green, Yin Xia, Kun Zhao, Zhen Zhao, Shengtao Ma, Wei Cao, Thilo Rehren, Computer vision, archaeological classification and China's terracotta warriors, *Journal of Archaeological Science*, Volume 49, September 2014, Pages 249-254, ISSN 0305-4403, <http://dx.doi.org/10.1016/j.jas.2014.05.014>

Bi, Z.M. and Wang, L, Advances in 3D data acquisition and processing for industrial applications, *Robotics and Computer-Integrated Manufacturing*, Volume 26, Issue 5, October 2010, pp403-413 and Paul J. Besl, Neil D. McKay, "A Method for Registration of 3-D Shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 14, No. 2

Daniel Biella, Daniel Sacher, Benjamin Weyers, Wolfram Luther , Nelson Baloian, Tobias Schreck, (2015) Crowdsourcing and Knowledge Co-creation in Virtual Museums, Collaboration and Technology, Volume 9334 of the series Lecture Notes in Computer Science pp 1-18

Blais, F., Taylor, J., Cournoyer, L., Picard, M., Borgeat, L., Godin, G., ... & Lahanier, C. (2007). Ultra high-resolution 3D laser color imaging of paintings: the Mona Lisa by Leonardo da Vinci.

Blinn, JF, (1978) Simulation of Wrinkled Surfaces, *SIGGRAPH '78 Proceedings of the 5th annual conference on Computer graphics and interactive techniques*, pp 286-292

Boehler, W., & Marbs, A. (2002, September). 3D scanning instruments. In *Proceedings of the CIPA WG 6 International Workshop on Scanning for Cultural Heritage Recording*, Ziti, Thessaloniki (pp. 9-18).

Brown et al (2008) E-Curator: 3D colour scans for remote object identification and assessment – Final Report

Bryan, P., Blake, B., & Bedford, J. (2013). Metric survey specifications for cultural heritage. English Heritage.

Bryan, P (2005), The role of photogrammetry in understanding enhancing and enjoying England's historic environment, in *Recording, Modelling and Visualization of Cultural Heritage*, Baltavista et al (Eds), Taylor and Francis Group, London, pp 78

Bryan, P. G., & Clowes, M. (1997). Surveying Stonehenge by photogrammetry. *The Photogrammetric Record*, 15(89), 739-751.

Burton, N. R., Hitchen, M. E., & Bryan, P. G. (1999). Virtual Stonehenge: a Fall from

Disgrace?. BAR INTERNATIONAL SERIES, 750, 265-265.

Butler, M., & Neave, P. (2008). Object appreciation through haptic interaction. Hello! Where are you in the landscape of educational technology? Proceedings asclite Melbourne 2008, 133-141.

C. S. McCamy, H. Marcus, and J. G. Davidson (1976). "A Color-Rendition Chart". *Journal of Applied Photographic Engineering* 2(3). 95–99.

Cameron, F, (2007), Beyond the Cult of the Replicant: Museums and Historical Digital Objects—Traditional Concerns, New Discourses, in *Theorising Digital Cultural Heritage* (pp49-75), Ed Cameron F & Kenderdine, S, MIT 2007

CARARE 2013, Europeana CARARE project – bringing content for archaeology and historic building to Europeana users (2010-2013)

M. Carrozzino, C. Evangelista, A. Scucces, F. Tecchia, G. Tennirelli, and M. Bergamasco. 2008. The virtual museum of sculpture. In Proceedings of the 3rd international conference on Digital Interactive Media in Entertainment and Arts (DIMEA '08). ACM, New York, NY, USA, 100-106

Casey (2003) “The museum effect: gazing from object to performance in the contemporary cultural-history museum”, Archives & museum informatics [1042-1467]

Chakrabarty, Dipesh (2002), “Museums in Late Democracies”, Humanities Research. X: 1

Chalmoviansky P and Juttler B, Filling Holes In Point Clouds, in Wilson M & Martin R, Mathematics of Surfaces, Lecture Notes in Computer Science, Vol 2768, 2003, pp 196-212

Chikatsu, H., & Takahashi, Y. (2009, August). Comparative evaluation of consumer grade cameras and mobile phone cameras for close range photogrammetry. In SPIE Optical Engineering+ Applications (pp. 74470H-74470H). International Society for Optics and Photonics.

Chih-Hao Yu, Tudor Groza, Jane Hunter, (2013), Reasoning on Crowd-Sourced Semantic Annotations to Facilitate Cataloguing of 3D Artefacts in the Cultural Heritage Domain, The Semantic Web – ISWC 2013, Volume 8219 of the series Lecture Notes in Computer Science pp 228-243

Cooper, M; A, R; Robson, S; Littleworth, R; M, ; (1992) The Tomb of Christ, Jerusalem;

analytical photogrammetry and 3D computer-modelling for archaeology and restoration. In: Archives for Photogrammetry and Remote Sensing. (pp. 774 - 785). International Society for Photogrammetry and Remote Sensing: Washington DC.

Couper, M, 2000, Web Surveys: A review of Issues and Approaches; Public Opinion Quarterly, Vol 64. Iss 4 pp 464-494

David Arnold & Guntram Geser (2008), EPOCH Research Agenda for the Applications of ICT to Cultural Heritage Full Report, Archaeolingua

Digitizing Cultural Heritage. About the Cover CultLab3D, IEEE Computer Graphics and Applications

Ebert, J. I. (1984). Remote sensing applications in archaeology. *Advances in Archaeological Method and Theory*, 293-362.

Edwards, P. J., Roberts, I., Clarke, M. J., Diguiseppi, C., Wentz, R., Kwan, I., ... & Pratap, S. (2009). Methods to increase response to postal and electronic questionnaires (Review). *Cochrane Database Syst Rev*, 3, MR000008.

Engelhard, N., Endres, F., Hess, J., Sturm, J., & Burgard, W. (2011, April). Real-time 3D visual SLAM with a hand-held RGB-D camera. In Proc. of the RGB-D Workshop on 3D Perception in Robotics at the European Robotics Forum, Västerås, Sweden (Vol. 180).

F. Durand. 2002. An Invitation to Discuss Computer Depiction. ACM/Eurographics Symp. NPAR'02.

F. Gabellone, M.T.Giannotta (2005) Realtime 3D multimedia system for the distance visiting of cultural heritage. A case study on the chamber tombs in Via Cristi, Taranto. CIPA 2005 XX International Symposium, 26 September 2005, Torino, Italy

Falk, JH & Dierking LD (2000), "Learning From Museums: Visitor Experiences and the Making of Meaning", AltaMira Press

Falk, J. H., & Dierking, L. D. (1992). The museum experience. Howells House.

Federal Agencies Digitization Guidelines Initiative, Technical Guidelines for Digitizing Cultural Heritage Materials: Creation of Raster Image Master Files at <http://www.digitizationguidelines.gov/guidelines/digitize-technical.html>

Figueroa Pablo, Mauricio Coral, Pierre Boulanger, Juan Borda , Eduardo Londoño, Felipe Vega, Flavio Prieto, Diego Restrepo (2009) "Multi-modal exploration of small artifacts: an exhibition at the Gold Museum in Bogota", Proceedings of the 16th ACM Symposium on Virtual Reality Software and Technology, November 18-20, 2009, Kyoto, Japan

Foley, J. D., & Van Dam, A. (1982). Fundamentals of interactive computer graphics (Vol. 2). Reading, MA: Addison-Wesley, pp. 505-511

Fontana, R., Gambino, M. C., Greco, M., Pampaloni, E., Pezzati, L., & Scopigno, R. (2003, October). High-resolution 3D digital models of artworks. In Optical Metrology (pp. 34-43). International Society for Optics and Photonics.

Frischer, B., & Fillwalk, J. (2012, September). The Digital Hadrian's Villa Project: Using virtual worlds to control suspected solar alignments. In 2012 18th International Conference on Virtual Systems and Multimedia.

García Fernández, J. (2013) Interpretation of topographic data from 'shape from shading' method, application in Villagarcia De Campos Castle, International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XL-5/W1, 2013 3D-ARCH 2013 - 3D Virtual Reconstruction and Visualization of Complex Architectures, 25 – 26 February 2013, Trento, Italy

Gardner, L. (2009). The Uses of Stored Collections in some London Museums. Papers from the Institute of Archaeology, 18(S1).

Gershon, R., Jepson, A. D., & Tsotsos, J. K. (1987, August). The Use of Color in Highlight Identification. In *IJCAI* (pp. 752-754).

Glaser, B. G., & Strauss, A. L. (2009). The discovery of grounded theory: Strategies for qualitative research. Transaction Publishers.

Goyer, G. G. & R. Watson (September 1963). The Laser and its Application to Meteorology in Bulletin of the American Meteorological Society. 44 (9): pp564–575

Graeme Earl, Kirk Martinez, Tom Malzbender, (2010) Archaeological applications of polynomial texture mapping: analysis, conservation and representation, Journal of Archaeological Science, Volume 37, Issue 8, August 2010, pp 2040-2050

Susie Green, Andrew Bevan, Michael Shapland, A comparative assessment of structure from motion methods for archaeological research, Journal of Archaeological Science, Volume 46, June 2014, Pages 173-181, ISSN 0305-4403,
<http://dx.doi.org/10.1016/j.jas.2014.02.030>.

Grincheva, N. (2015). 'The World Beach Project' Going Viral: Measuring Online Influence—Case Study of the Victoria and Albert Online Museum Project; Journal of Creative Communications, Vol 10 iss. 1, pp 39-55.

Grün, A., Remondino, F. and Zhang, L. (2004), Photogrammetric Reconstruction of the Great Buddha of Bamiyan, Afghanistan. The Photogrammetric Record, 19: 177–199

Gurian, Elaine Heumann (1999) "What Is the Object of This Exercise? A Meandering Exploration of the Many Meanings of Objects in Museums", *Daedalus*, Vol. 128, No. 3, America's Museums (Summer, 1999), pp. 163-183

Guthrie, K., Griffiths, R., & Maron, N. (2008). Sustainability and revenue models for online academic resources. An Ithaka Report.

Haegler, S, Müller, P & Van Gool, L.. 2009. Procedural modeling for digital cultural heritage. *J. Image Video Process.* 2009, Article 7 (February 2009), 1 pages.
DOI=<http://dx.doi.org/10.1155/2009/852392>
Haindl, M., & Filip, J. (2013). Visual Texture: Accurate Material Appearance Measurement, Representation and Modeling. Springer Science & Business Media., p9

Hawkins, T., Cohen, J., & Debevec, P. (2001, November). A photometric approach to digitizing cultural artifacts. In *Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage* (pp. 333-342). ACM

Heckbert, P. S. (1986). Survey of texture mapping. *Computer Graphics and Applications*, IEEE, 6(11), 56-67.

Hedberg J, Harper B,Dalgarno B. (2002) "The contribution of 3D environments to conceptual understanding" O. J. McKerrow (Eds.), *Winds of Change in the Sea of Learning: Proceedings of the 19th Annual Conference of the Australasian Society for Computers in Learning in Tertiary Education* Vol 1 (pp. 149-158)

Hess M et al (2008), Final Report: E-Curator: 3D Colour Scans For Remote Object Identification and Assessment, UCL Museums and Collections. Available at: <http://www.ahessc.ac.uk/e-curator>

Hess, M., Millar, F. S., Ong, Y. H., MacDonald, S., Robson, S., Brown, I., & Were, G. (2008). E-Curator: 3D Colour Scans for Object Assessment. *EVA 2008 London Conference*, July 22-24, 2008

Hess, M., Were, G., Brown, I., MacDonald, S., Robson, S., & Simon Millar, F. (2009). E-Curator: a 3D Web-based archive for conservators and curators. *Ariadne* online magazine for information professionals in archives, libraries and museums in all sectors., (60).

Hess, M, (2015) A metric test object informed by user requirements for better 3D recording of cultural heritage artefacts, thesis

Hess, M., & Robson, S. (2010). 3D colour imaging for cultural heritage artefacts. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 38(Part 5), 288-292.

Hess, M., Robson, S., Millar, F. S., Were, G., Hviding, E., & Berg, A. C. (2009, September). Niabara - the Western Solomon Islands war canoe at the British Museum-3D documentation, virtual reconstruction and digital repatriation. In *Virtual Systems and Multimedia, 2009. VSMM'09. 15th International Conference on* (pp. 41-46). IEEE.

Hooper-Greenhill, E (2000) "Museums and the Interpretation of Visual Culture", Routledge New York

Hughes, L. M. (2012). Introduction: the Value, Use and Impact of Digital Collections. Evaluating and Measuring the Value, Use and Impact of Digital Collections. London: Facet Publishing.

ICOM Code of ethics, section 2: <http://icom.museum/the-vision/code-of-ethics/2-museums-that-maintain-collections-hold-them-in-trust-for-the-benefit-of-society-and-its-developme/#sommairecontent>

ICOM Statutes, 21st General Conference in Vienna, Austria, in 2007,

Ihrke, I et al, State of the Art in Transparent and Specular Object Reconstruction, Eurographics 2008

Jensen, H. W., Marschner, S. R., Levoy, M., & Hanrahan, P. (2001, August). A practical model for subsurface light transport. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (pp. 511-518). ACM

JISC Digital Media (2014). Digital 3D content infokit, <http://www.jiscdigitalmedia.ac.uk/infokit/3d/3d-home>

Kakuta T, Oishi T, Ikeuchi K, (2008) "Development and Evaluation of Asuka-Kyo MR Contents with Fast Shading and Shadowing", Proc. Int. Society on Virtual Systems and MultiMedia (VSMM 2008), Oct. 2008, pp.254-260.

Kalantar, JS & Talley NJ, 1999, The Effects of Lottery Incentive and Length of Questionnaire on Health Survey Response Rates: A Randomized Study. *Journal of Clinical Epidemiology*, Vol 52, Iss 11, pp 1117-1122, Nov 1999

Kersten, T. P., & Lindstaedt, M. (2012). Image-based low-cost systems for automatic 3D recording and modelling of archaeological finds and objects. In *Progress in cultural heritage preservation* (pp. 1-10). Springer, Berlin & Heidelberg.

Kilgard, M. J. (2000). A practical and robust bump-mapping technique for today's GPUs.

King, N, Using Interviews in Qualitative Research (2004), In C. Cassell & G. Symon, eds. *Qualitative Methods in Organizational Research, A Practical Guide*. London: Sage, pp. 11-23

Kobbelt, L., & Botsch, M. (2004). A survey of point-based techniques in computer graphics. *Computers & Graphics*, 28(6), 801-814.

Krishnaswamy, A. and Baranowski, G. V.G. (2004), A Biophysically-Based Spectral Model of Light Interaction with Human Skin. *Computer Graphics Forum*, 23: 331–340. doi: 10.1111/j.1467-8659.2004.00764.x

La Pensée, A., Cooper, M. I., & Parsons, J. B. (2006). Applications in the field of cultural heritage using “off-the-shelf” 3d laser scanning technology in novel ways. In Proc. 7th Int. Symp. On Virtual Reality, Archaeology and Cultural Heritage,(VAST) (Vol. 215).

Lens Distortion Correction (LDC) from UCL Geomatics Software,
<http://eprints.ucl.ac.uk/43957/>

George Lepouras, Costas Vassilakis (2004), Virtual museums for all: employing game technology for edutainment, *Virtual Reality*, June 2004, Volume 8, Issue 2, pp 96-106
Levoy, M., Pulli, K., Curless, B., Rusinkiewicz, S., Koller, D., Pereira, L. & Fulk, D. (2000, July). The digital Michelangelo project: 3D scanning of large statues. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (pp. 131-144). ACM Press/Addison-Wesley Publishing Co.

London Charter for the computer-based visualisation of cultural heritage:
<http://www.londoncharter.org/downloads.html>

Henrique Lorenzo and Pedro Arias, A methodology for rapid archaeological site documentation using ground-penetrating radar and terrestrial photogrammetry, *Geoarchaeology, Special Issue: Archaeological Geophysics: A Global Perspective*, Volume 20, Issue 5, pages 521–535, May 2005

Lourakis, M & Argyros, A, “SBA: a software package for generic sparse bundle Adjustment” *ACM transactions on Mathematical Software*, 2009, pp1-30

Lowe, D.G., "Object recognition from local scale-invariant features," *Computer Vision*, 1999. The Proceedings of the Seventh IEEE International Conference on , vol.2, no., pp.1150,1157 vol.2, 1999

Luebje, D & Humphreys, G, How GPU's Work:
mowgli.hadassah.ac.il/mod/resource/view.php?id=38019

Luhmann T, Robson S, Kyle S, Boehm J, “Strategies for camera calibration”, *Close-Range Photogrammetry and 3D Imaging*, 2nd Ed, De Gruyter Berlin/Boston 2014

M. Gillings. 2000. Plans, Elevations and Virtual worlds: the development of techniques for the routine construction of hyperreal simulations. *Virtual Reality in Archaeology*, Archaeopress, Oxford, 2000.

Macdonald, L & Robson, S (2010), Polynomial Texture Mapping and 3D Representations, International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII, Part 5 Commission V Symposium, Newcastle upon Tyne

MacDonald, L.W, 2014 Colour and Directionality in Surface Reflectance, Proc. Artificial Intelligence and the Simulation of Behaviour (AISB), Goldsmiths College, London, April

MacDonald, S, 2001, *Behind the Scenes at the Science Museum: Knowing, Making and Using*, in Academic Anthropology and the Museum: Back to the Future, ed. Bouquet, M, 2001, Bergahn Books

MacDonald, L., Hindmarch, J., Robson, S., and Terras, M.: Modelling the appearance of heritage metallic surfaces, Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci., XL-5, 371-377, doi:10.5194/isprsarchives-XL-5-371-2014, 2014.

Mallick, SP et al, Beyond Lambert: Reconstructing Specular Surfaces Using Color, Computer Vision and Pattern Recognition, 2005, Vol 2 pp619-626

Malzbender et al, Polynomial Texture Maps, SIGGRAPH 01, Proceedings of the 28th annual conference on Computer graphics and interactive techniques, pp 519-528, NY, 2001

Marani, R., Nitti, M., Cicirelli, G., D’Orazio, T., & Stella, E. (2015). Design of High-Resolution Optical Systems for Fast and Accurate Surface Reconstruction. In *Sensing Technology: Current Status and Future Trends III*(pp. 47-65). Springer International Publishing.

Maria Roussou, George Drettakis. Photorealism and Non-Photorealism in Virtual Heritage Representation. A.Chalmers and D.Arnold and F. Niccolucci. First Eurographics Workshop on Graphics and Cultural Heritage (2003), 2003, Brighton, United Kingdom. Eurographics, pp.10, 2003, Proceedings of the International Symposium on Virtual Reality, Archeology and Cultural Heritage.

Maron, N., Smith, K. K., & Loy, M. (2009). Sustaining digital resources: an on-the-ground view of projects today: Ithaka case studies in sustainability. Ithaka S+ R.

Matusik, W. (2003). A data-driven reflectance model (Doctoral dissertation, University of North Carolina, Chapel Hill).

Meade Bache, R, (1892) Civil and Military Photogrammetry, Proceedings of the American Philosophical Society, Vol. 30, No. 138 (April), pp. 229-240, or,

Melville, C, Northern Iraq: Historical and Political Context, Court and Craft Catalogue, pp16, The Courtauld Gallery in Assoc. with Paul Holberton Publishing, London, 2014

Messham-Muir, Kit. (2006) "Affect, Interpretation and Technology." Open Museum Journal 7

Metallo A & Rossi V (2011) *The Future of Three-Dimensional Imaging and Museum Applications*, Curator, The Museum Journal, Vol 54. No 1, Jan 2011, pp 63-69

Meyer, E.T., Madsen, C., Eccles, K. (2009). TIDSR Survey on the Use of Digitised Resources. Available online: <http://microsites.oi.ox.ac.uk/tidsr/kb/27/tidsr-survey-use-digitised-resources>.

Mudgea, M., Malzbender, T., Schroer, C., & Lum, M. (2006, October). New Reflection Transformation Imaging Methods for Rock Art and Multiple-Viewpoint Display. In *VAST* (Vol. 6, pp. 195-202).

Mudgea, M., Ashleby, M., & Schroer, C. (2007). A digital future for cultural heritage. Available in: http://culturalheritageimaging.org/What_We_Do/Publications/cipa2007/CIPA_2007.pdf.

Müller, P, Wonka, P, Haegler, S, Ulmer, A & Van Gool, L, 2006. Procedural modelling of buildings. *ACM Trans. Graph.* 25, 3 (July 2006), 614-623.
DOI=<http://dx.doi.org/10.1145/1141911.1141931> Muller K (2002) "Museums and virtuality" *Curator, the museums journal*. Vol 45, No 1, Jan 2002

Museum Association Code of Ethics:
<http://www.museumsassociation.org/ethics/code-of-ethics>

Neidhardt, F. (1993). The public as a communication system. *Public Understanding of Science*, 2(4), 339-350.

Newmam, I & Benz, C R (1998) "Qualitative-quantitative research methodology: exploring the interactive continuum", Southern Illinois University, Chapter 2, "Qualitative and Quantitative Research Methods: An Interactive Continuum"

Nicodemus, Fred (1965). ["Directional reflectance and emissivity of an opaque surface"](#) *Applied Optics* Vol. 4, Iss. 7, pp767-775

Nikon K-Scan [http://www.nikonmetrology.com/en_EU/Products/Laser-Scanning/Handheld-scanning/K-Scan-MMDx-walkaround-scanning/\(key_features\)](http://www.nikonmetrology.com/en_EU/Products/Laser-Scanning/Handheld-scanning/K-Scan-MMDx-walkaround-scanning/(key_features))

Padfield, J., Saunders, D., & Malzbender, T. (2005). Polynomial texture mapping: a new tool for examining the surface of paintings. *ICOM Committee for Conservation*, 1, 504-510.

Pallud, J. (2009). The application of a phenomenological framework to assess user experience with museum technologies, *ECIS 2009 Proceedings*. Paper 395.

Paul J. Besl, Neil D. McKay, "A Method for Registration of 3-D Shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 14, No. 2

Payne, E. M. (2013). Imaging techniques in conservation. *Journal of Conservation and Museum Studies*, 10(2), pp17-29.

Pfeifer, N., & Briese, C. (2007). Laser scanning-Principles and applications. p6

Pfister, H., Zwicker, M., Van Baar, J., & Gross, M. (2000, July). Surfels: Surface elements as rendering primitives. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (pp. 335-342). ACM Press/Addison-Wesley Publishing Co..

Phong BT, (1975) Illumination for computer generated pictures, *Communications of the ACM*, Vol 18 Iss 6, June 1975, pp 311-317

Pierre Grussenmeyer, O. Al Khalil. A comparison of photogrammetry software packages for the documentation of buildings. 2008.

Pointools Pro: http://www.bentley.com/en-US/Promo/Pointools/pointools.htm?skid=CT_PRT_POINTOOLS_B

Project Mosul, <http://projectmosul.org/>

Remondino, F & Fraser, C, "Digital Camera Calibration Methods: Considerations and Comparisons", ISPRS Commission V Symposium 'Image Engineering and Vision Metrology', IAPRS Volume XXXVI, Part 5, Dresden 25-27 September 2006

Remondino, F., & Rizzi, A. (2010). Reality-based 3D documentation of natural and cultural heritage sites—techniques, problems, and examples. *Applied Geomatics*, 2(3), 85-100.

ScanLAB Projects: <http://scanlabprojects.co.uk/>

Schwartz, C., Sarlette, R., Weinmann, M., & Klein, R. (2013, June). DOME II: a parallelized BTF acquisition system. In *Proceedings of the Eurographics 2013 Workshop on Material Appearance Modeling: Issues and Acquisition* (pp. 25-31). Eurographics Association.

Seybert, H. (2011). Internet use in households and by individuals in 2011. *Eurostat. Statistics in Focus*, 66.

Stoyanov, D., Darzi, A., & Yang, G. Z. (2005). A practical approach towards accurate dense 3D depth recovery for robotic laparoscopic surgery. *Computer Aided Surgery*, 10(4), 199-208.

Strauss, PS, A realistic lighting model for computer graphics, Computer Graphics and Applications, IEEE Volume 10, Issue 6, pp 56-64 , Nov 1990

Taylor J, Beraldin JA, Godin G, Cournoyer L, Rioux M, and Domey J (2002) "3D imaging technology for museums and heritage" Proceedings of The First International Workshop on 3D Virtual Heritage, October 2-3, 2002, Geneva, Switzerland, pp. 70-75.

Terras, M. (2010). "The Rise of Digitization: An Overview". In Rukowski, R. (ed). (2010). "Digital Libraries". Sense Publishers, The Netherlands.

Terras, M. (2015). "Cultural Heritage Information: Artefacts and Digitization Technologies" In Chowdhury, G. and Ruthven, I. (2015). "Cultural Heritage information", Facet. p. 63-88.

Terras, M., Ross, C. & Motyckova, V., 2012. Measuring impact and use: scholarly information-seeking behaviour. In Hughes (2012), pp85-102

Toolkit for the Impact of Digitised Scholarly Resources, a Jisc funded project at the Oxford Internet Institute: <http://microsites.oi.ox.ac.uk/tidsr/>

Tips for the effective use of close range digital photogrammetry for the Earth sciences, ISPRS - Commission V - Close-Range Sensing: Analysis and Applications Working Group V / 6 - Close range morphological measurement for the earth sciences, 2008-2012

Uccheddu, F., Pelagotti, A., Cappellini, V., & Massa, E. (2014, July). 3D Technologies for Measurement of Painting Surface Deformations: A Case Study. In Proceedings of the EVA London 2014 on Electronic Visualisation and the Arts (pp. 56-62). BCS.

UNESCO General Conference, 1978. Recommendation for the protection of movable cultural property. http://portal.unesco.org/en/ev.php-URL_ID=13137&URL_DO=DO_TOPIC&URL_SECTION=201.html

Vegelin Van Claerbergen, E, Foreward, pp8, ibid

Vilbrandt, C., Pasko, G., Pasko, A., Fayolle, P.-A., Vilbrandt, T., Goodwin, J. R., Goodwin, J. M. and Kunii, T. L. (2004), Cultural Heritage Preservation Using Constructive Shape Modeling. Computer Graphics Forum, 23: 25–41. doi:10.1111/j.1467-8659.2004.00003.x

Vozikis, G., Haring, A., Vozikis, E., & Kraus, K. (2004). *Laser scanning: a new method for recording and documentation in archaeology*. na.

Walter, B., Marschner, S. R., Li, H., & Torrance, K. E. (2007, June). Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques* (pp. 195-206). Eurographics Association.

Ward, R (Ed) Court and Craft: A Masterpiece from Northern Iraq, 2014, The Courtauld Gallery in asoc. W. Paul Holberton Publ., London

WebGL: <https://www.khronos.org/webgl/>

Wheatstone, C (1838) Contributions to the Physiology of Vision.--Part the First. On Some Remarkable, and Hitherto Unobserved, Phenomena of Binocular Vision, *Philosophical Transactions of the Royal Society of London*, Vol. 128, pp. 371-394

Witcomb A (2007), "The Materiality of Virtual Technologies", *Theorizing Digital Cultural Heritage* Ed. Cameron F & Kenderdine S, MIT Press, Cam, MA

Rafal Wojciechowski, Krzysztof Walczak, Martin White, and Wojciech Cellary. 2004. Building Virtual and Augmented Reality museum exhibitions. In *Proceedings of the ninth international conference on 3D Web technology (Web3D '04)*. ACM, New York, NY, USA, 135-144. Wright, K. B. (2005), Researching Internet-Based Populations: Advantages and Disadvantages of Online Survey Research, Online Questionnaire Authoring Software Packages, and Web Survey Services. *Journal of Computer-Mediated Communication*, Vol 10, iss 3

XBPointstream: http://zenit.senecac.on.ca/wiki/index.php/XB_PointStream

Zhang, Y., & Wildemuth, B. M. (2009). Unstructured interviews. In B. Wildemuth (Ed.), *Applications of Social Research Methods to Questions in Information and Library Science* (pp.222-231). Westport, CT: Libraries Unlimited.

Zimmer R, Jefferies J, "Accessing material art through technologies of mediation and immediation", *Futures*, Volume 39, Issue 10, December 2007, Pages 1178-1190


Zwicker M et al. 2001. Surface splatting. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques (SIGGRAPH '01)*. ACM, New York, NY, USA, 371-378.

9 Appendix A: Science Museum Shipping Gallery Survey

9.1 Online survey results


Question 1

How did you originally find this video?




Question 2

How familiar are you with the Science Museum?




Question 3

How familiar are/were you with the Shipping Gallery




Question 4

Did you enjoy the video?




Question 5

Do you think this is a good way of preserving old exhibitions?




Question 6

Would you like to see this sort of thing done for existing exhibitions? For example to help plan a visit to the museum?




Question 7

Imagining for a moment that Shipping Gallery exhibition still existed, would this video have made you more or less likely to visit the gallery?




Question 8

Would you like to have been told more about the technology used to create the video?




Question 9

Would you like to be able to explore the model yourself (ie, 'walk around' the gallery) rather than follow a video)?




Question 10

This video showcases a very distinctive style, with the objects in the gallery appearing translucent and almost 'ghostly'. Did you like this style of presentation, or would you have preferred to see something more solid?




Question 11

Do you have much experience with virtual worlds and environments on computers? ie, perhaps through playing games?



Question 12

Please tell us how old you are...



Question 13

Do you have any other comments? Please feel free to expand on your previous answers, tell us what you think about the video, the technology behind it or, well, anything!

Respondent	Comment
783283	The gallery and its models were an essential research tool. Access to 'solid' virtual models would allow it to continue. As it stands it is a frustrating tantalising reminder of what an asset we used to have. if the fidelity can be improved and I control my on movement around the space I would forgive you for the loss of the gallery.
787921	I liked the video and understand that due to the amount of data it would have been difficult to use it all but if more exhibits are going to be scanned it would be far better to have something more solid to fully do it justice. especially if the exhibit is going to be dismantled.
793924	I wonderful use of technology, well done ! I hope someone makes a Oculus rift viewable version of the Shipping Gallery exhibition, I think a lot of people around the world will then have a chance to visit :-) All the Best Saul Wynne London UK
777484	Thank you for undertaking and achieving something so valuable. When I heard that the gallery was closing, I felt as sad as when the land transport gallery closed all those years ago; my collection of science museum books (from the 1930s onwards)has a few photographs to remind me but nothing as incredible as these immersive scans. The Shipping Gallery has been with me since I was a child, I visited with my Grandad and my father. I remember the Blue Peter lifeboat and

	<p>revolving radars on the roof of the bridge (long gone!); and the inside of the ships' bridge when all worked and the equipment hadn't been stripped by tourists for souvenirs! My Dad loved the south pointing chariot! We visited together the day after the Gallery closed and only found about the its demise when we wanted to visit. To still be able to 'look around' is wonderful! I pray that you don't fill the space with those very odd exhibits such as in the welcome wing which rely too much on their presentation style to deliver their message; the confused 'Who Am I' exhibition was dreadful. Too flashy, too much effort and expense in overblowing minor points of interest better expressed in other ways. Stop showing off - there is no point in making a name for yourself to be forgotten in 2 years time! The reason the Shipping Gallery endured was that the presentation was clean, accessible to all and (aside only from the style and fonts used in the signs) didn't ever age. If any exhibit installed today still looks modern in 50 years time, I'll be amazed. Please don't allow yourselves to become bewitched by all the clever-clever touch screens, moody lighting, learning 'experience' rubbish; there is still much to be said for displaying an historically important object next to a piece of card explaining what it is - the Turing exhibit got the balance just right. But the scanning project is quite brilliant; whoever thought to record the past in this way is a visionary. Future generations will thank you for it. Perhaps the project will be complete once detailed scans of the models are made, allowing people to walk round and drill down to the detail of the ships and engines, perhaps even those round-nose - almost sharp - brass buttons on the square boxes just below the glass that used for almost 40 years to bring the reciprocating engines to life! In the way that a photograph is a 2D rendering of the 3D world, you have made a 3D rendering of our 4D world - you miss only the time dimension. I wish that you could go the last step and bring my Grandad back so we can walk round it together. Long live the Science Museum! Mike Greene rafmike@gmail.com</p>
783420	You need to include information on how one can see the items that were once part of this exhibition now that they are no longer in this gallery.
794123	The idea is good, but the fly through tour is far too superficial to be useful, except to give a very limited first impression. A fully explorable model, with the ability to view the exhibits (and labels) close-up, could be a very valuable tool.
777476	Would like to do a virtual tour and not just see a video of what you can do , if an exhibition has been closed and people want to see it a virtual tour is a good way to preserve it

795195	It makes the whole thing look a bit fake which is a shame for such a real place like the science museum full of REAL things such as history and weight and material.
780280	I hoped that it would be possible to really see the models that the museum for some reason or another put away far from the visitors.
776696	The video was for me a poignant reminder of kinder days when governments cared about our heritage and museums cared about their curators. Sadly lost forever. One day I imagine all museums will simply be like this - ghosts caught in the machine.
778798	I loved the shipping gallery and would rather you kept it, since it shut I have not visited the Science museum. With the NMM moving its models to Chatham it is as if the Nation is turning its back on its maritime tradition
773459	Impressed by video, much potential for showing exhibits, but I came here because I was so disappointed to discover that one of my favourite galleries has gone. Virtual is not Real!
774873	Whilst it is fantastic that technology can be put to work preserving exhibits like this it still looks like a stuffy old museum but with a modern twist with a very dull narrative basically it is missing the wow factor even though the production is cutting edge
776740	I wanted to see the detail of the exhibits, the level of detail presented was frustratingly low. A good way to get taster but I don't live near London I would like the option to explore each item in detail.
777474	How can we actually see the model ships now that the Shipping Section is closed?
1128793	This video is no substitute whatsoever for one's being able to examine the characteristics of individual ship models in the gallery, as one used to be able to do when it was open to the public. The technology used in this video does not even allow one to identify a particular ship model of interest in the gallery, let alone study it. Current data download times preclude any detailed inspections of the models. I am very disappointed to see what the Science Museum has done with its once famous Shipping Gallery. 2D Computer images will never replace 'being there and seeing the real thing,' in my opinion.
776753	The shipping gallery was one of the best museum galleries in the world - it was a crime that it was destroyed. I have not heard any good argument for why it needed to be replaced, especially when the science museum has huge new gallery space in the Wellcome wing. In the digital age what we need more and more are galleries of real

objects, not images or screens. Increasingly objects are going into storage and being digitised, and gallery space filled with installations and open space. Give me real objects and day, not more screens.

9.2 Online comments for the Shipping Gallery Video

9.2.1 Metafilter

<http://www.metafilter.com/130281/Scrapped-but-not-forgotten>

July 23, 2013 1:21 PM

The Science Museum in London closed their Shipping Galleries in 2012, having been open for almost 50 years. But in case you missed it, here's a narrated [short virtual tour](#), as it looked then.

Opened in 1963, the Shipping Galleries were home to 1800 maritime exhibits, including many incredibly large and detailed ship models; such as the original builders' model of Brunel's infamous [SS Great Eastern](#), the biggest ship in the world by far when she was built in 1858. There was also the first marine gas turbine, working engine models and many other unique exhibits of maritime history.

Before it was closed to make room for new exhibits, 275 laser scans collected 2 billion precise measurements, and now the data is being used to reconstruct an incredibly accurate virtual model of the Shipping Galleries. This tour is just an early taster, and uses a mere 10% of the collected data. [[via](#)]

posted by [ArkhanJG](#) (15 comments total) [18 users marked this as a favorite](#)

That is at the same time beautiful and very very ghostly.

posted by [jeribus](#) at [1:43 PM](#) on July 23, 2013 [[1 favorite](#)]

That was an amazing place to visit.

posted by [unSane](#) at [1:48 PM](#) on July 23, 2013

I'm really down about that, I love that gallery. Lived at the Science Museum with my little ones before they went to school, spend many an hour in that hall. Preserving a virtual reality version while they put those models in the attic is unsatisfactory

posted by [C.A.S.](#) at [1:48 PM](#) on July 23, 2013 [[1 favorite](#)]

One of the things I loved about the Science Museum when I went there was the various galleries that had been untouched for decades. The mathematics hall, circa the 1960s, had a bunch of models showing geometrical solids, with plain typed notes alongside them. There were exhibits that seemed 1980s, with touch screen computers being the big thing; others had the post-Exploratorium interactive components. There was a postmodern gallery where everything was web-based and interactive. It was like visiting a museum exhibit about the history of science museums, and that meta-aspect was fascinating. Not to say that some of those old galleries weren't boring as shit.

posted by [Homeboy Trouble](#) at [2:01 PM](#) on July 23, 2013 [[7 favorites](#)]

Soon, Google will laser scan and index the interior of every structure like that.

posted by [planetesimal](#) at [2:39 PM](#) on July 23, 2013

Great. Now I'm feeling wistfulness over the non-existence of a place I never knew existed.

For people on the other side of the pond, there's an impressive (though much smaller) [collection of model ships](#) in the basement of the Art Gallery of Ontario in Toronto.

posted by [bonobothegreat](#) at [3:32 PM](#) on July 23, 2013 [[3 favorites](#)]

The Shipping Gallery is going to be replaced by [InformationAge](#), six collections of related objects that will be a "celebration of information and communication technologies".

To be honest, this is probably overdue. The boat collection would be best served at the [maritime museum](#), or various other sites around the UK. After all, the models aren't being disposed off, they are being put into storage.

posted by [The River Ivel](#) at [3:45 PM](#) on July 23, 2013

> The Shipping Gallery is going to be replaced by InformationAge, six collections of related objects that will be a "celebration of information and communication technologies".

> the models aren't being disposed off, they are being put into storage

That sums up everything I hate about modern museums. Remember when museums used to have old stuff? And lots of it?

The last time I visited a museum it was nine tenths packaging, and only a tiny fraction of its stuff was on display. I would happily scrap all the display boards and architectural space and white surfaces and interactive exhibits for rooms and rooms piled high with OLD STUFF.

posted by [EnterTheStory](#) at [4:18 PM](#) on July 23, 2013 [[5 favorites](#)]

Oh man, I remember this. It was right in at the back and ridiculous. A diorama of the Port of London and even better a display of things which get delivered to the port: bottles of bleach, rope, other things!

posted by [Damienmce](#) at [6:07 PM](#) on July 23, 2013

[Things we get from ports!](#) The 3D model is very spooky. Sort of worrying when places you know from real life cease to exist and become entirely virtual. Not to self, start weird nostalgia internet business

scanning 3D models of schools to sell to people today when they're in their dotage, sitting in retirement homes connected to Oculus Rift.

posted by [Damienmce](#) at [6:19 PM](#) on July 23, 2013 [[1 favorite](#)]

I'm glad I saw the gallery before its demise. It seemed like it was designed to teach marine architecture students the history of ship design. The virtual fly through was pretty, but will the finished version allow one to look closely at an individual item and read its caption?

posted by [monotreme](#) at [8:31 PM](#) on July 23, 2013

This is sad also because the Shipping Gallery had a few benches where you could eat your lunch in complete peace, even during the school holidays. Lovely boat models *and* I've napped in there on a bench with my daughter asleep in her stroller, in the days when I never got any sleep. Always deserted.

'Little detailed models of really big things inside glass cases' is something that children don't even bother to compute any more.

At least there are some fantastic galleries still open in the Science Museum that are also usually deserted. 'Glimpses of Medical History' is a huge collection of life-size mannequins doing stuff like biting on leather gags while being held down by other mannequins as their diseased legs are sawn off.

posted by [colie](#) at [2:24 AM](#) on July 24, 2013

I hadn't heard this until right now, and I'm rather sad about it. I'm actually in London for the summer and was looking forward to visiting this gallery - it was a favorite of a friend of mine mostly because of the quiet, the loving attention put into the gallery, and the fact that it really did feel untouched by time. It was very obviously a labor of love, and that's why he and I love(d) it so much.

Seconding the recommendation of Glimpses of Medical History, which is

delightful if only because of the utter disjointedness of the displays. They definitely aren't arranged according to any sort of reasonable timeline or categorization...

posted by [naturalog](#) at [6:30 AM](#) on July 24, 2013

A continuing annoyance of mine is the way that so many museums make no effort to enable virtual visitations. It's not the same as being there, but why can't I virtually travel through the Smithsonian, for example, taking my time to view all the artifacts, read all the placards, and watch all the motion pictures? I would pay to do that.

posted by [LastOfHisKind](#) at [4:54 PM](#) on July 24, 2013

9.2.2 B3ta.com

That's properly awesome.

I'm intrigued as to how it's set up, how colours are scanned, whether you manage to retain any internal detail (i.e. how much can I see inside the cutaway models), etc.

Also, in which skip should I go diving for the WW2 battleships? I'd like a Prince of Wales or Hood if you have them still, but I'll settle for anything really.

In order to embed, change the link to: www.youtube.com/watch?v=qDTbFhFZl9I

( [wheresthefish](#) This div will be replaced, Wed 24 Jul 2013, 13:22, [Ignore](#), [I like this!](#), [Reply](#))

Voxels!

Yay! make it interactive :)

wonderful

( [zacherynuk](#) @echo off, Wed 24 Jul 2013, 13:38, [Ignore](#), [I like this!](#), [Reply](#))

Oooh the figurehead one is wonderful

Couldn't you use something like cloudcasterlite atop a beefy box with loads of RAM and striped SSD ?

With a decent GPU - I am sure there is GPGPU code for accelerating it all nicely :)

Maybe a project for later...

(.sendStatus 200, Wed 24 Jul 2013, 15:31, [Ignore](#), [I like this!](#), [Reply](#))

 **That was really, really great.**

Must have been amazing to be involved.

Click!

( **Fork** Has lost his bash virginity!, Wed 24 Jul 2013, 13:43, [Ignore](#), [I like this!](#), [Reply](#))

 **I actually applauded at the end of that.**

Really wonderful, but also tinged with nostalgia for me as the Sci Museum is somewhere I have been going all my life, from when I was at school to taking my own kids there - even last year my 20 yr old wanted to go back there as it was somewhere her and I used to go to together when she was at primary school.

It's changed a lot over the years - lots of the slightly outdated and dusty exhibits gone to be replaced with shiny interactive exhibits for schoolkids. In many ways I think that's a shame, maybe it's just nostalgia, I don't know.

I'll miss the shipping galleries though, the models were so beautifully crafted. Any ideas what's happening to them?

( **spesh**, Wed 24 Jul 2013, 14:09, [Ignore](#), [I like this!](#), [Reply](#))

9.2.3 Alphr

<http://www.alphr.com/features/385582/3d-museums-that-never-close>

Dairs • [a year ago](#)

... in our provincial galleries!

vjosullivan • [a year ago](#)

@ChrisH The artifacts are indeed the important bit and they can still be scanned and displayed in the way you describe, at any time in the future. However, this was the only opportunity to record the gallery and re-create it.

ChrisH • [a year ago](#)

We seem to have spent a lot of effort scanning and storing huge detail about a building, but surely the artefacts are the important bits? If they had been scanned in high detail in isolation, they could be rendered into a virtual gallery and most people wouldn't have cared. I'd have thought that would probably make it much easier to share the data with the public too.

•

Jaberwocky • [a year ago](#)

I'm going to hazard a guess here that the black circles all around the floor of the 3d point cloud video must be where they set up the scanning lasers. They must be the blind spots, under the rotating lasers themselves, where they can't scan.

9.2.4 Gizmag

<http://www.gizmag.com/shipping-gallery-3d/28844/>

While this is better than nothing, I still kind of hope that this will all be recreated in some way at the National Maritime Museum.

scc970

17th September, 2013 @ 11:11 am PDT

9.2.5 YouTube

https://www.youtube.com/all_comments?v=gDTbFhFZl9I

Chip Spencer 4 months ago

Wonderful and stunning. Were the images created by laser scanning and structured light?

rhadiem 5 months ago

Looks great, but I hope a full display of the data points will look better than these interesting but more ghost-like images that only give a sense of the gallery. Great idea to do this though, rather than have it lost forever.

aarocka11 8 months ago

When will the data be released? I want to see the gallery through my oculus rift.

Chris Nikolajsen Shared on Google+ · 11 months ago (edited)

Laser scanned building displayed as a point cloud.

Mozzie 11 months ago

make an oculus level for this

David Shirres 1 year ago

Very sad that the museum has closed its shipping gallery. The Science Museum now has almost nothing about the marine technology that had a huge impact on Britain's history. Digitisation of the galleries is a poor substitute but better than nothing.

Tudor Cook Shared on Google+ · 1 year ago

absolutely stunning video!

I'd read about the digitisation project some months back and the record that it would provide, but this shows that the end result is so much better than the description would lead you to believe.

Simply a fantastic use of technology

Read more

Tudor Cook 1 year ago

absolutely stunning use of technology!

I'd read about the digitisation project some months ago but never imagined quite how stunning the results could be...marvelous.

Abdullah Waseem Shared on Google+ · 1 year ago

great work

Spanna Hanz 1 year ago

Wow!

robi b Shared on Google+ · 1 year ago

Laser mapping eh? Looks good enough to get a general gist. Though I'm surprised the narrator only mentioned a billion or so points. I would have thought it would be a few more orders of magnitude larger to capture everything in that much detail.

Jeremy Bell Shared on Google+ · 1 year ago

Wow, 3D laser scanning works better than I imagined it would.

IdeasForTheKids Shared on Google+ · 1 year ago

Amazing use of technology to preserve exhibitions and open them up for even more varied use

Manny Coulon Shared on Google+ · 1 year ago

Brilliant use of the latest 3D scanning technology to preserve and open up the +Science Museum's Shipping Gallery, which, after nearly 50 years, has been dismantled and put in to storage. Awesome computer processing power will make the gallery and its exhibits accessible in many, many new ways for future generations. Exciting potential for museums as this technology becomes more mainstream...

lawrence windrush 1 year ago

A shame so many iconic galleries are vanishing, too much emphasis is placed on visiting schoolchildren, and one of the worlds great museums is just becoming a Disney play area.

Colman Carpenter Shared on Google+ · 1 year ago

Just read an article about how, and why, this project came about in PC Pro magazine. Fascinating...and a beautiful video too!

I remember walking through this gallery a few years ago...but didn't realise that even then it had been around for 45 years or so.

Stephen Whitelaw 1 year ago

Stunning :)

Graitec Ltd Shared on Google+ · 1 year ago

Stunning and productive use of point cloud data - A very good video.

arthur brogard 1 year ago

Very disappointing, what we've got here. But apparently there's something much better somewhere - two billion scans worth or something. Where is it? How do we see it?

lara Ferguson 1 year ago

brilliant...can't wait to actually explore the exhibition in 3d myself.....

Tio Rams 1 year ago

If they use Eclideon's technology, Geaverse, they can stream this on the web. Via browser..

DomainRider 1 year ago

A lot of potential there - so when will this virtual gallery be made available for individual walk-through at decent resolution?

It needs an option to turn transparency on and off, and a future enhancement could include dynamic displays (rotating engines, etc).

johnabdn 1 year ago

Looking at your other videos, the Science Museum seems to be only catering for primary school children now.

johnabdn 1 year ago

Who decided this was outdated? I still enjoyed it.. Damn shame..

Aaro Sahari 1 year ago

The gallery was outdated as a means of telling this particular story. This is an interesting opening to preserving our way of telling stories through exhibitions though. Evidently the Science Museum has higher resolution material in story though (see New Scientist 24 August 2013).

Anthony Cooper 1 year ago

I agree, the technology is impressive, but I'm devastated that the gallery is gone. Does the SM have plans to relocate the display items, eg. to Greenwich or some other theme-relevant museum? If not, then it's a tragedy for ship buffs.

Scientist-Online 1 year ago

+Anthony Cooper I was under the impression that it was all going into storage, which does seem a shame.

MaximusNYC 1 year ago

If these exhibits were so wonderful, why were they taken down? A digital fly-thru is not an adequate replacement for the actual historic models the narrator describes.

Juncus Bufonius 1 year ago

It seems a little unsatisfactory as shown. See through objects and just too little resolution to make out any detail. If that's all there is then some photos would have been just as good. If there is more where is it?

William H.G. Johnson Shared on Google+ · 1 year ago

This is probably the coolest virtual gallery of all time.

CommunistHamster 1 year ago

This is pretty impressive.

It would be very interesting to see this converted to a polygon mesh, with commentary boxes a lá HL2 Lost Coast, then made into a barebones Unity level with Oculus Rift support.

Chris Williams 1 year ago

this needs to be packaged as a program oculus rift support

9.2.6 Oculus.com

<https://forums.oculus.com/viewtopic.php?t=3796&p=48216>

The UK Science Museum has had a display on shipping, aka the Shipping Gallery, since 1963. They've just closed it and are replacing it with something on IT. However, before doing so, they 3D scanned the entire gallery. They plan to release the data publicly at some point this year.

I just thought this was cool, could be something for the future, and would be a brilliant use for the Oculus Rift. Imagine if, in the future, you could go and visit previous displays, specials, etc, even after they've been closed. Going there in person would be best, but what if you can't travel? Thought it may interest some people on here.

See http://www.sciencemuseum.org.uk/about_us/history/shipping.aspx for more details.

Ian

kingtut

Re: Science Museum: 3D scan of Shipping Gallery

Postby Calanar » Fri Aug 23, 2013 10:15 am

This is really amazing. I hope in the future to take virtual tours of museums and sites I'd have to pay a fortune to visit otherwise. Also the ability to preserve these experiences means a museum or site can have many multiples of tours at once. So they can change the grounds and still show the old exhibits or mix and match for a theme. I don't think society yet realizes what a boom this will be. Imagine you are a teacher in one country teaching history from another. You could go to every museum in the world that has a related exhibit virtually or perhaps a reenactment of the event or site involved in the history itself. The future is unlimited. I grew up as personal computers were taking off and it was a liberating time. Imagine all of the ways this can change society. So many opportunities will now come to us if we only dream big enough.

9.2.7 Gizmodo

<http://www.gizmodo.co.uk/2013/07/take-an-amazing-virtual-tour-of-one-of-the-most-important-science-museum-exhibitions-ever-scraped/#comment-1393691142>

goodfondue • [2 years ago](#)

I owe this exhibition a debt of gratitude. Saw it when I was a kid, joined the Merchant Navy after school and now get to sail icebreakers for a living.

A really great video for an inspiring place!

clipper • [2 years ago](#)

As a professional Marine Engineer - the last time I saw the Science Museum Maritime Galleries, they were vastly superior to anything in the Maritime Museum at Greenwich - who overdosed big time on Naval, Sail & art, etc.

The working models illustrating the developments of propulsion systems, for example in the Science Museum were wonderful, so very sorry to hear the display have been all broken up. But very well done with the fly through

Typical of this bloody country, not realising the importance the Marine Industries contributed - and contributes - to the world.

Ah well, another reason not to visit London.

[Someone Else](#) • [2 years ago](#)

That explains why it's dumbed down then. Should aim it at the Android generation.

[Sam Gibbs](#) • [2 years ago](#)

Yeah, I know what you mean. But it's one of those things. Gotta aim it at the PlayStation generation. Or, really, these days it's the iPhone generation, I guess.

[Someone Else](#) • [2 years ago](#)

No, it's become too much about "making science fun and accessible" by which I mean dumbing it down and having it explained by cartoon characters. Just seeing the giant machinery was enough to inspire my interest in science, engineering and technology.

[Sam Gibbs](#) • [2 years ago](#)

Maybe that's just because you're remembering it through the eyes of a child? All the wonder etc. I used to have the odd lecture in there, was a good place for science. That and the Natural History Museum to be fair.

[Mr. T](#) • [2 years ago](#)

Fear not Sam, soon you can spend time in a lot better way on PS4TW!!!

[Someone Else](#) • [2 years ago](#)

I too have fond memories of the shipping Galleries (and the science museum in general) went back there recently as was appalled at how dull it is now.

9.2.8 Selection of Tweets referencing the Shipping Gallery scan

Chris Alexander @cmalexander Jun 19

So cool! Cost mentioned? MT @MarDixon: This is the @sciemcemuseum Shipping gallery 3D model <http://bit.ly/1m0ccg6> @5easypieces #museumnext

Visido Imaging @Visido 6 Feb 2014

Beautiful #3D #preservation of now closed Shipping Gallery exhibit in Science Museum. <http://youtu.be/gDTbFhFZl9I> via @youtube

BT Archives @BTArchives 29 Jan 2014

Fantastic @sciemcemuseum fly-through of old Shipping Gallery, space to be new comms #sminfoage exhib open in Sep 2014
http://m.youtube.com/watch?v=gDTbFhFZl9I&desktop_uri=/watch?v=gDTbFhFZl9I ...

Steve Baines @sjbaines 28 Nov 2013

Just found @sciemcemuseum scanned entire Shipping Gallery before closing it - fantastic 'virtual preservation' idea! <http://tinyurl.com/qy78k9d>

Manny Coulon @mannyc 28 Nov 2013

Love how 3D scanning technology has been used to preserve and open up the @sciemcemuseum shipping gallery. Brilliant <http://bit.ly/188PdWb>

Joris Schets @jorisschets 14 Nov 2013

Science Museum Creates Stunning 2 Billion-Point 3D Model Of Shipping Galleries With 275 Lasers <http://huff.to/1bEOe5q> via @HuffPostUKTech

National Museum news @nmdcnews 27 Oct 2013

Endangered site that you'd like laser recorded (like @sciemcemuseum did with their shipping gallery) - read on here: <http://archive.cyark.org/submit-site>

James Lyon Fenner @dr_small_craft 22 Oct 2013

Absolutely incredible! What a fitting tribute to the Shipping Gallery.
http://www.sciencemuseum.org.uk/about_us/history/shipping.aspx ...

Matt Mccarter @oatfedgoat 12 Sep 2013

Its just so beautiful it's worth watching again. Point cloud of shipping gallery at science museum.
http://m.youtube.com/watch?v=gDTbFhFZl9I&desktop_uri=%2Fwatch%3Fv%3DgDTbFhFZl9I
...

Manuel Dornbusch @loki1978de 6 Sep 2013

Science Museum preserves Shipping Gallery as virtual exhibit <http://feedly.com/k/1305qgP> I was in the museum in 2010, but not in that Gallery

dennyhardiana @dennyhardiana 3 Sep 2013

#wow Shipping Gallery as virtual exhibit by [sciencemuseum\(.org.uk\)](http://sciencemuseum.org.uk)
<http://www.youtube.com/watch?v=gDTbFhFZl9I> ...

Uber Nemo @ubernemo 2 Sep 2013

The Science future of museums? Museum preserves Shipping Gallery as virtual exhibit
<http://www.gizmag.com/shipping-gallery-3d/28844/> ... via @gizmag

Alice Lighton @alicelighton 28 Aug 2013

A gorgeous virtual tour of the (now-ex) shipping gallery at the @sciencemuseum, narrated by @rooneyvision <http://www.youtube.com/watch?v=gDTbFhFZl9I> ...

Lyn Jeffery @LynJ 26 Aug 2013

3d virtual reality of closed London Museum of Science shipping exhibition made from lidar "echolocation"ish process [http://ispr.info/2013/08/26/londons-science-museum-uses-laser-scanning-to-create-faithful-virtual-tour-of-closed-gallery](http://ispr.info/2013/08/26/londons-science-museum-uses-laser-scanning-to-create-faithful-virtual-tour-of-closed-gallery/?utm_source=rss&utm_medium=rss&utm_campaign=londons-science-museum-uses-laser-scanning-to-create-faithful-virtual-tour-of-closed-gallery) ...

Aube Lebel @LebelAube 25 Aug 2013

[Virtual exhibition] The Science museum scanne en 3d la shipping gallery avant de la faire disparaître <http://bbc.in/14nxbB3> cc@muzeonum

Liam O'Neill @LiamONeill34 23 Aug 2013

@sciemcemuseum <http://www.bbc.co.uk/news/magazine-23733895> ... The shipping gallery scan with an @oculus 3D VR headset would be an amazing experience.

James Poskett @jamesposkett 23 Aug 2013

3D copy of closed shipping gallery preserved at @sciemcemuseum. Interesting future for digital heritage #histsci <http://www.bbc.co.uk/news/magazine-23733895> ...

Julia Murray @juliamurray22 23 Aug 2013

The Shipping Gallery was archived to make way for Information Age, but it lives on in time & space @sciemcemuseum <http://www.bbc.co.uk/news/magazine-23733895> ...

Claire Allan @clairenothelen 23 Aug 2013

This is a great video showing changes @sciemcemuseum from Shipping gallery to Information Age <http://m.bbc.co.uk/news/magazine-23733895> ... #thwb #museums

Rebekah Higgitt @beckyfh 23 Aug 2013

@tillyblyth Of course there are some @NMMGreenwich who deeply mourn the passing of the Shipping Gallery! @rooneyvision @sciemcemuseum

Steve Bowbrick @bowbrick 6 Aug 2013

Oh my: opened the year I was born, @sciemcemuseum's shipping gallery is now an amazing, rather melancholy 3D model <http://bit.ly/12WO5aQ>

Roger Highfield @RogerHighfield 1 Aug 2013

Our ghost gallery
<http://www.flickr.com/photos/sciemcemuseum/9339479911/in/photostream/> ... Video:
http://www.sciemcemuseum.org.uk/about_us/history/shipping.aspx ...

Amy Oliver @MuseumMogul 26 Jul 2013

Beautiful 3D rendering of the Science Museum in the UK. Perhaps this tool can be used for all types of preservation? http://www.huffingtonpost.co.uk/2013/07/22/science-museum-shipping-gallery-model_n_3633725.html ...

George Mokhtar @GeorgeMokhtar 26 Jul 2013

"@oatfedgoat: Here's the stunning point cloud video shown at @UCLGeomatics and @3DLaserMapping conference last week
http://m.huffpost.com/uk/entry/3633725?utm_hp_ref=tw ... "

Penny Edwell @PennyEdwell 25 Jul 2013

Check out @scencemuseum's AMAZING #3D model of its shipping gallery. Narrated by curator David Rooney : <http://bit.ly/15gyZL7> #musetech

Museum Studies Leics @LeicsMusStud 24 Jul 2013

Brilliant digital fly through of the Shipping Gallery at @scencemuseum...
<http://ow.ly/ngEnr> #digitalheritage

Rob @rotster 24 Jul 2013

Forgive the #NerdyMuseumTweet, but I just LOVE this digital fly-through of the Shipping Gallery at @scencemuseum <http://www.gizmodo.co.uk/2013/07/take-an-amazing-virtual-tour-of-one-of-the-most-important-science-museum-exhibitions-ever-scrappe>/ ...

Official W3DC News @Web3DCommunity 24 Jul 2013

Science Museum Creates Stunning 2 Billion-Point 3D Model Of Shipping Galleries With 275 Lasers <http://huff.to/1bEOe5q> via @HuffPostUKTech

Dr Mariann Hardey ^{□2} @thatdrmaz 24 Jul 2013

The Science Museum... The Shipping Gallery... Dave Rooney... Lasers... What's not to like? ht @ukmcg <http://www.gizmodo.co.uk/2013/07/take-an-amazing-virtual-tour-of-one-of-the-most-important-science-museum-exhibitions-ever-scrappe/> ...

Mike Ellis @m1ke_ellis 24 Jul 2013

Absolutely love this: @scencemuseum + @rooneyvision + the shipping gallery + lasers.
<http://bit.ly/14D4b9z>

Roger Highfield @RogerHighfield 24 Jul 2013

MT@SciencePunk: Here's what the @ScienceMuseum's Shipping Gallery looks like captured on 3D scan <http://www.youtube.com/watch?v=gDTbFhFZl9I> ... <ghost ships

Tori Herridge @ToriHerridge 24 Jul 2013

It lives on (virtually)! MT @trueanomalies @melissaterras crazy detailed 3D tour of @sciencemuseum's shipping gallery
http://www.sciencemuseum.org.uk/about_us/history/shipping.aspx ...

Jon Voss @jonvoss 24 Jul 2013

Agreed--amazing! Archiving entire museum exhibits:
http://www.huffingtonpost.co.uk/2013/07/22/science-museum-shipping-gallery-model_n_3633725.html?utm_hp_ref=tw ... via @JessicaKausen

Fay Curtis @fay_fay_fay 23 Jul 2013

Cool. MT @MuseumMinute: Science Museum Creates Stunning 2 Billion-Point 3D Model Of Shipping Galleries w/ 275 Lasers <http://huff.to/1bEOe5q>

Jessica Kausen @JessicaKausen 23 Jul 2013

Is it weird that this gave me chills? http://www.huffingtonpost.co.uk/2013/07/22/science-museum-shipping-gallery-model_n_3633725.html?utm_hp_ref=tw ... cc @nickstanhope @jonvoss @Historypin bc I think you'll all enjoy it!

Krista Steele @SciVizKrista 23 Jul 2013

What's possible w/ #laserscanning? Museums can scan their collections & promote them online as a mkting tool http://www.huffingtonpost.co.uk/2013/07/22/science-museum-shipping-gallery-model_n_3633725.html?utm_hp_ref=tw ... @FARO_HQ

Roger Highfield @RogerHighfield 23 Jul 2013

MT@melissaterras: 3d model made of @sciencemuseum gallery before it was dismantled
http://www.sciencemuseum.org.uk/about_us/history/shipping.aspx ... <it's fab!

Hayley M. Birch @gingerbreadlady 23 Jul 2013

Right up @bonny_jennett's street RT @melissaterras: 3D fly through of @sciencemuseum's shipping gallery http://www.sciencemuseum.org.uk/about_us/history/shipping.aspx ...

sumant singh bhatia @sumantbhatia 23 Jul 2013

Really nice, even emotional, video on the closing of Science Museum's shipping gallery, but now preserved in 3D <http://bit.ly/1352fjE>

Paige Dansinger @museumpaige 22 Jul 2013

@nealstimler @museums365 @dklevan Or what this could do for significant sites, reconstruction or immersive exhibits: <http://m.huffpost.com/uk/entry/3633725> ...

Alli Burness @alli_burnie 22 Jul 2013

Beautifully narrated. #musetech MT @rjstein: Ghostly scan of galleries. Museum Preserved In Time With 275 Lasers <http://zite.to/12Yc2Na>

Krista Steele @SciVizKrista 22 Jul 2013

Wanted: Laser Scanning experts to capture the world- growing career- digital preservation http://m.huffpost.com/uk/entry/3633725?utm_hp_ref=tw ... @CyArk @FARO_HQ @Sci_Vis

Matt Mccarter @oatfedgoat 22 Jul 2013

Here's the stunning point cloud video shown at @UCLGeomatics and @3DLaserMapping conference last week http://m.huffpost.com/uk/entry/3633725?utm_hp_ref=tw ... #UKBIMCrew

Krista Steele @SciVizKrista 22 Jul 2013

Museum educators, teachers and professors- laser scanning will revolutionize how you present information http://m.huffpost.com/uk/entry/3633725?utm_hp_ref=tw ... @FARO_HQ

9.3 Media links

A selection of links to online media reports of the Shipping Gallery scanning project.

http://www.huffingtonpost.co.uk/2013/07/22/science-museum-shipping-gallery-model_n_3633725.html

<http://www.digitalartsonline.co.uk/news/motion-graphics/science-museum-reveals-3d-model-of-shuttered-gallery/>

<http://new-aesthetic.tumblr.com/post/56428171065/shipping-galleries-3d-model-science-museum>

<http://3dblog.org/london-science-museum-creates-a-realistic-3d-model-of-a-shipping-gallery/>

<http://blog.lidarnews.com/preserving-the-shipping-gallery>

<http://www.bbc.co.uk/news/magazine-23733895>

<http://www.modelboats.co.uk/forums/postings.asp?th=55278>


<http://www.newscientist.com/article/mg21929311.700-virtual-reality-resurrects-a-defunct-exhibition.html#.VPcrJvmsV8E>

10 Appendix B: Illuminating Objects survey results


Question 1

Downloading: How long did the model take to download fully?

Bowl:




Bible:




Question 3

Interaction: How did you find your interaction with the model (moving it, zooming etc.)

Bowl:



Bible:



Question 4


Do you have any other comments on the technical aspects of this site and your experience with it?

Bowl	
728131	The image wouldn't load in Internet Explorer, but it worked fine using Firefox, as suggested.
780923	The model never stops moving which could be a good thing unless you want to look in detail at one part of it.
780925	I did double click to try to zoom in, but it just reset it. I think I am used to google maps controls
781650	Very impressive
792837	No - all was well
796843	I would have preferred it if the model wasn't spinning all the time
Bible	
715880	A full screen version would be better. Even though I could zoom in I found it a bit small.
716798	Didn't load in safari, ctrl and left mouse button zooms in on the whole display for the mac rather than the bible, so a bit difficult
716804	It did not load at all and I use Firefox already.
716808	The bible needs to stop spinning when you're no longer rotating it.
716844	Astounding level of detail.
716872	It doesn't make sense to have such a small viewing window, users expect to be able to use full screen feature. Presume this is created with actual photography but the quality is such that it's not 100% clear; it looks like a synthetic model.
716959	I couldn't zoom (using Chrome)
718204	using Android tablet with Firefox... :-)
757640	Can I open the book?
780928	Sometimes you could partly see through the object to the back of the object although it didn't particularly detract from the object. I found the fact that the model kept rotating the whole time slightly annoying.


Question 5

Do you think that viewing the 3D model improved your understanding of the object?

Bowl:




Bible:




Question 6

Did the model look like a convincing, 'real', three-dimensional object?

Bowl:




Bible:




Question 7

Would you like to see more 3D models on the Courtauld's website?

Bowl:




Bible:




Question 8

Would seeing 3D models online make you more or less likely to visit the gallery? For example, having seen the model of this Spanish bowl, are you more or less likely to visit the gallery to see the real thing?

Bowl:



Bible:



Question 9

Do you have any other comments or thoughts on your experience with the model?

Bowl

I'm not certain which pattern should be the inside of the bowl and which should be on the bottom as sometimes sometimes they would swap over
780923 and what was the inside became the outside.

It's hard to judge depth unless you have it side on. Normal photos have some light and shadow whereas this doesn't so you can't tell what depth the
780925 various elements are at in relation to each other.

I think the geometry of the plate confused me when turning from one face to another. It looked as if there was a dent in the centre but I'm sure it was
781650 an optical illusion!

Bible

I remember trying to do similar work (in the late nineties). Incredible how
716844 primitive that was compared to this.

For an art gallery, it's odd to focus on a closed manuscript as 3D model, the
716872 art is inside...


757640 I look forward to seeing where this approach goes in future.

This model was possibly larger on screen than it would be in real life. While having the option to zoom in on a very small object is great it was difficult to
780928 get a sense of scale.


Question 10

How would you rate your level of online experience?

Bowl:




Bible:




Question 11

How much experience have you had with 3D models and 3D technology? (This could be, for example, through accessing models like this one, from playing videogames or using other virtual environments and technologies)

Bowl:




Bible:




Question 12

Before visiting this site, did you have a particular interest in this bowl, or this type of object (ceramics, Spanish pottery, lustreware etc.)?

Bowl:



Bible:



Question 13

Please add any comments you may have on the previous questions, and especially on your particular interest (if any) in this, or this type of, object.

Bible

As a general rule, online digitization increases visitorship. This has been long established by libraries with digitized special collections holdings

716872 online.

interested in seeing an example of a 3d model used in digital library

719162 collection

Question 14

Thanks for completing the survey! Feel free to add any further comments below:

Bowl

792837 Great survey - very necessary I look forward to seeing the results.

Bible

716844 Really superb job, folks. Thanks for showing it off!

It's best to use 3D to offer the in-person experience as closely as possible. In this case, I would expect to be able to open the book and page through it. Content + purpose should drive form +

716872 function.

11 Appendix C: A possible method for processing highly specular objects

I shall now describe a process that unfortunately couldn't be implemented, but may point the way to future methods of reducing holes and noise in photogrammetric models of shiny objects. One of the frustrating aspects of dealing with an object with specular highlights is that the area in question is present in many images, and only shows the specular highlight in some, with other images revealing the 'true' diffuse colour. Thus, there may be enough 'clean' information on a particular area to ascertain a good 3D point. Unfortunately there is no way at present to tell the photogrammetric software which data – or pixels in the image – are 'clean' and which are specular highlights and therefore should be ignored.

Our method was to use masking, in which we would process the images and mask the areas that showed specularity – replacing them with black pixels. This involved two processes; coming up with a method of identifying highlighted pixels in the image, and then getting the algorithms to ignore these pixels when processing the model. The first process was successful, the second, less so.

Some attempts at identifying specular highlights image are based on the assumption that the 'specular' pixel is saturated in at least one (RGB) colour channel⁵⁵¹. Others use the HSV (hue, saturation, value/luminance) colour space, reasoning that highlights would have a different profile than diffuse reflection – ie, the highlight would be less

551 Brelstaff and, G & Blake, A, Detecting specular reflections using lambertian constraints, Proceedings of the 2nd International Conference on Computer Vision, 1988, pp. 297–302

saturated (S) but have a higher luminance (V)⁵⁵². Our method allows both approaches, the user selecting the most appropriate for the object and images in question, though with this particular project we used the HSV channel almost exclusively.

However, once we had created the masks, it proved impossible to insert them into the photogrammetric workflow. While PMVS does include a masking function, it is for segmenting images, removing the background to create models of foreground objects.

(From the documentation: "The software tries to reconstruct 3D points until image projections of these points cover all the target images (*only foreground pixels if segmentation masks are given*)"⁵⁵³). Thus our resultant model, after inputting our masks, showed large holes wherever a specular highlight had occurred in any of the images.

11.1.1 Our method

The masks were created in Matlab [see later for code] using the following steps:

- Show the user a zoomed in image of a photo containing specular highlights, and let the user select an area with 'bad' pixels, ie, pixels representing a highlight

⁵⁵² Demaagd, K, et al, Practical Computer Vision with SimpleCV: The Simple Way to Make Technology See, 2012, O'Reilly, pp93

⁵⁵³ <http://www.di.ens.fr/pmvs/documentation.html>




Figure 11.1: The user selects an area containing a specular highlight

- Matlab then calculates a colour profile for the selected pixels, in both RGB and HSV colourspace, and displays these two profiles as histograms.




Figure 11.2: the histograms for the area selected in fig 13; hue = cyan; saturation = magenta, value(luminance) = yellow

- Repeat 1-3 but selecting a 'good' area of the image




Figure 11.3: The whole process repeated with a ‘good’ area. Notice the difference in the histograms, particularly in the hsv space.

- Once you have the two sets of histograms, the colour profiles can be compared and values selected which distinguish the bad from the good pixels. In the example shown in *fig. 21 and 22*, and using the hsv histograms the following code is used to identify pixels with specular highlights.

```
if ((hue > .6 && hue < .92) || (sat < .19 && val > .6));
```

- At the moment, those values must be hard-coded, and once entered creates a new image with the specular highlights replaced with black pixels [*fig 23*]. (For the actual masks, non-specular pixels are replaced with white pixels to create a two-tone image, I have skipped this step to better show the results of the masking process.
- The values can be tweaked until a satisfactory result is obtained, and then be tested on multiple images. Once the appropriate values have been selected, the set of images can be batch processed, outputting a mask for each image.




Figure 11.4: Left, the input photo, right, the corresponding output with the specular highlights replaced by black pixels.

There are caveats to go with this approach. While it appeared to be successful with our image set, it is probable that this particular object is naturally amenable to this method. The object's colour is fairly uniform, and consists of either the yellowish

ground or the reddish copper glaze. This makes the object's colour profile fairly narrow, and makes discriminating between good and bad pixels simpler. With an object that has a much broader range of colours, it may be harder to find settings that don't remove good pixels along with the bad. Of course, the same settings do not need to be used for all of the images, and they could be split into smaller sets of similar images. It may also be the case that for some points on the object, once certain images have been masked, there are no longer enough images containing that point to create a good 3d point. If this resulted in small holes in the model, that may not be an issue, as they are in some ways preferable to noise (the first step in hole filling is usually the erasure of noisy areas anyway).

11.1.2 Matlab code

```
%#ok<*NOPTS>
%This code is used to determine values then used to create the masks
im = imread('C:\Users\Phd-uczajhi\Desktop\Bowl-
right_cam\DSC_2075.JPG', 'jpg');

%This returns the pixels selected by the user
%large_image_display code from
http://www.mathworks.co.uk/help/images/creating-%the-modular-tools.html

[select, position] = large_image_display(im)

position
newImg =
im(position(2):position(2)+position(4),position(1):position(1)+positio
n(3),1:3);
figure, imshow(newImg);

%createColorHistograms(newImg);
hold on;
x = 0:1:255;
figure ('name', 'rgb') ;
red = (newImg(:,:,1));
r = red(:)';
r = cast(r, 'double');
[graph1,graph2] = hist (r,x);
bar(graph2,graph1, 'FaceColor', 'r','EdgeColor', 'r')
alpha(0.3);
hold on;

green = (newImg(:,:,2));
g = green(:)';
g = cast(g, 'double');
[graph1,graph2] = hist (g,x);
bar(graph2,graph1, 'FaceColor', 'g','EdgeColor', 'g');
alpha(0.3);
hold on;

blue = (newImg(:,:,3));
b = blue(:)';
b = cast(b, 'double');
[graph1,graph2] = hist (b,x);
bar(graph2,graph1, 'FaceColor', 'b','EdgeColor', 'b');
alpha(0.3);

x = (0:.01:1);

hsv_newImg = rgb2hsv(newImg);

figure ('name', 'hsv') ;
```

```

hue = (hsv_newImg(:,:,1));
h = hue(:)';
[graph1,graph2] = hist (h,x);
bar(graph2,graph1, 'FaceColor', 'c', 'EdgeColor', 'c');
hold on;
sat = (hsv_newImg(:,:,2));
s = sat(:)';
[graph1,graph2] = hist (s,x);
bar(graph2,graph1, 'FaceColor', 'm', 'EdgeColor', 'm');
hold on;
val = (hsv_newImg(:,:,3));
v = val(:)';
[graph1,graph2] = hist (v,x);

bar(graph2,graph1, 'FaceColor', 'y', 'EdgeColor', 'y');

hold off;

figure ('name', 'mask image');

hsv_newFullImg = rgb2hsv(im);
maskImage = im;
width = size(im,2);
height = size(im,1);
for i = 1 : width;
    for j = 1 : height;
        hue = hsv_newFullImg(j, i, 1);
        sat = hsv_newFullImg(j, i, 2);
        val = hsv_newFullImg(j, i, 3);
        if ((hue > .6 && hue < .92) || (sat < .19 && val > .6)) ;
            maskImage(j, i, 1:3) = [0,0,0];
        end
    end
end
imshow(maskImage);

```

create_mask.m

```

%This code uses the values determined via masking.m and applies them
to a batch of images, creating a mask for each one

numbImages = 41;
imgType = '.jpg';
imgPath = 'F:\BOWL_IMAGES\Bowl_front\pmvs\visualize\';
imgName = '000000';
outputPath = 'F:\BOWL_IMAGES\Bowl_front\pmvs\mask\';
outputName = '000000';

%main loop
for mainCounter = 0 : numbImages-1;
    filename = strcat(imgPath, imgName);
    if mainCounter < 10;
        filename = strcat(filename, '0');
    end
    filename = strcat(filename, num2str(mainCounter), imgType);

```

```

img = imread(filename);
%read image

%create mask
hsv_newFullImg = rgb2hsv(img);
maskImage = img;
width = size(img,2);
height = size(img,1);
for i = 1 : width;
    for j = 1 : height;
        hue = hsv_newFullImg(j, i, 1);
        sat = hsv_newFullImg(j, i, 2);
        val = hsv_newFullImg(j, i, 3);
        if ((hue > .65 && hue < .9)|| (sat <.9 && val > .6)) ;
            maskImage(j, i, 1:3) = [0,0,0];
        %
        else
            maskImage(j, i, 1:3) = [255,255,255];
        end
    end
end

%output mask

filename = strcat(outputPath, outputName);
if mainCounter < 10;
    filename = strcat(filename, '0');
end
filename = strcat(filename, num2str(mainCounter), '.jpg');
imwrite(maskImage, filename, 'jpg');
end

```

12 Appendix D: Suggested method for enhancing user interaction with point cloud models

12.1 Motivation

From the research conducted during this thesis, it is clear that interaction is something desired by users of 3D content. Point cloud viewers such as the open source XBPS used throughout this project allow simple interaction, in the form of customisable cameras that allow users to rotate objects and zoom in and out. To provide richer experiences, it would be useful to be able to label areas of the point cloud, or make areas of the point cloud 'clickable'. For instance, clicking certain areas of the model could start video or other media playing, or reveal some extra textual descriptions of the object. One could even allow users themselves to label and attach media to areas of the point cloud.

Methods do exist to allow identification of particular points in a point cloud model, using, for example raycasting⁵⁵⁴. These methods are fairly complex to implement, however, and tend to return the index of a point. The method suggested below is very simple, involving just a few lines of code added to the point cloud renderer and shader, and thus could be implemented by those with little programming experience. It also returns a 3D coordinate for the point, potentially making it more useful than an index.

12.2 Method

Assign each axis (X,Y,Z) a colour (R,G,B), and divide each one into 256 equal lengths, giving a 3D volume divided into 256x256x256 boxes. It is a simple matter to render the

⁵⁵⁴ For example: http://threejs.org/examples/#webgl_interactive_raycasting_pointcloud

point cloud with each point coloured according to the box it is in using shader code

similar to the following:

```
R = (ps_Vertex.x - minX / maxX-minX);  
G = (ps_Vertex.y - minY / maxY-minY);  
B = (ps_Vertex.z - minZ / maxZ-minZ);  
  
frontColor = vec4(R,G,B,1.0);
```

When the user clicks on the point cloud, read the colour under the mouse cursor. By converting the RGB value back into XYZ coordinates, the position of the clicked point in 3D space can be calculated. The key to this method is to render the point cloud twice; once, as normal, and the second time, using the above colouring method but rendered to a buffer rather than the screen. When the user clicks on the normally coloured model, we can read the colour from the buffered image.

The key issue is the fact that we have to render the point cloud twice, potentially causing a large performance hit. However, since the coloured point cloud only needs a resolution of 256x256x256, a heavily decimated version of the original cloud containing only a few thousand points can be used, so long as the points are rendered large enough that the model is ‘solid’ and there are no obvious gaps.

Rough ‘proof of concept’ examples of this method, both based on XBPS, can be found at:

<http://www.homepages.ucl.ac.uk/~uczcjhi/cloudLabel/files/index.html>

and

<http://www.homepages.ucl.ac.uk/~uczcjhi/cloudPicker/files/index.html>

In the first case, click anywhere on the model to attach a 'label' (in this case just a short line), in the second, click on the model's nose to start a video playing.

To control the camera use 'WASD' to move and the cursor keys to look. The '+' and '0' keys on the num pad will toggle views of the original point cloud and coloured point cloud respectively.




Figure 12.1: Left, the model rendered normally; right, showing the coloured point cloud

13 Appendix E: A selection of scanning projects undertaken during the course of this research

Below are a small selection of the scanning projects I have been involved with during the course of this research, along with descriptions of the issues encountered (if any) with each one. Many are representative of types of object that may be encountered in cultural heritage.

Target & Method	Issues
Blaschka squid model ⁵⁵⁵ Arius laser scanner	<p>The object is made of glass, and unpainted areas (for example, parts of the tentacles) could not be captured at all. Similarly, the eyes are made of black glass balls, were too dark to reflect the laser and thus were also invisible.</p> <p>The tentacles and area around the mantle exhibited extremely complex geometry and many occlusions, while the fragile nature of the object made capture from many angles impossible.</p> 

⁵⁵⁵ https://en.wikipedia.org/wiki/Leopold_and_Rudolf_Blaschka

UCL North Lodge Gallery space scan	Virtually none; the gallery consists of a single, roughly cubic room decorated with flat paintings and exhibiting no complex geometry. A single scan captured the entire room, so no registration was required. Good quality textures were obtained, almost good enough to read the label text on the walls.
------------------------------------	--




Elephant's tooth from the Grant Museum	Apart from some complex geometry in one area which was impossible to capture, and minimal amounts of gloss and some specularity on the exposed tooth surface, this was one of the most successful models I have produced.
--	---



Horus, Petrie Museum loan object Arius scanner	The varnished object exhibited some specularity, but with multiple scans it was possible to get 'clean' data for most areas. Some complex geometry, but largely successful.
	
Scarab, Petrie Museum loan collection Arius scanner	The scarab's very dark, very matte material meant little laser light was reflected, and only scans made orthogonal to the surface captured much data. This led to problems scanning 'round corners' and made registering, for example, the base with the sides almost impossible.
	

<p>Centaur statues, Courtauld Institute Photogrammetry</p>	<p>With UCL colleague Mona Hess, we photographed two centaur statues in the foyer of the Courtauld Institute of Art. The photography was hampered by the position of the statues, placed close to a low wall and metal railing. The photogrammetric processing was performed by Bernie Frischer's team and the models placed inside the Virtual Hadrian's Villa project (http://idialab.org/virtual-hadrians-villa/)</p>
	

Copper figurine Arius scanner	<p>The figure has areas which are too dark to capture, as well as shiny areas exhibiting a high degree of specularity. Complex geometry leads to many occlusions (for example under the arms).</p> 
Shipping Gallery model Nikon Metris	<p>This model from the Science Museum's Shipping Gallery was scanned during the decanting process. The extremely complex geometry of the model (for example, the lifeboats on deck, masts etc.) led to much occlusion and meant the scan could not be completed in the time available (a full day). The geometry also made it virtually impossible to texture the model using photographs.</p> 

14 Appendix F: Publications and presentations

14.1 Museum Computer Network, Dallas, November

Abstract for presentation:

3D SCANNING OF THE SCIENCE MUSEUM'S SHIPPING GALLERY

The (London) Science Museum's Shipping Gallery, largely unchanged since installation in the 1960s, was a vast space with hundreds of ship models and large pieces of historical maritime equipment. When the gallery was decommissioned in 2012, the museum, in a combined project with University College London and Scanlab, created a 3D record of the exhibition. The resulting model documents and preserves the gallery, allowing users to experience the virtual exhibition long after the physical space was dismantled. The gallery was first surveyed, then scanned over five nights. Two FARO Photon 120 terrestrial laser scanners captured 275 scans. The resulting 256 GB dataset contained two billion colored points measured to sub-millimeter accuracy. This presentation will demonstrate the model and some planned outputs, as well as discussing issues of feasibility, user experience, and the potential benefits of large-scale 3D capture of museum spaces. This project was supported by UCL's VEIV EngD Centre.

14.2 Digital Humanities Congress, Sheffield, September 2014

Abstract for presentation:

3D Modelling of Islamic Metalwork: Processes and Potentials⁵⁵⁶

The 'Courtauld Bag', a brass bag inlaid with silver and gold and manufactured in Mosul in the early 14th century, is a unique object recognised by specialists as one of the most important examples of Islamic metalwork in the world. A major exhibition, 'Court and Craft: a masterpiece from Northern Iraq', was created around this beautiful object, and ran at the Courtauld institute from February to May of this year.

As part of the exhibition, UCL's 3D imaging group were commissioned to create an animation to be displayed in the gallery alongside the object. The bag was scanned with an Arius Foundation laser scanner, imaged under a PTM¹ dome and finally photographed for photogrammetric reconstruction. Despite the shiny, metallic nature of the object, a detailed 3D model was created using structure-from-motion², while a brand new technique was used for specular reconstruction from the PTM images, creating stunning photo-realistic renderings of small details of the bag. These renderings were combined to create a two and a half minute video which was shown in the exhibition. (The video can be viewed at <http://www.courtauld.ac.uk/gallery/exhibitions/2014/Court-and-Craft/model.shtml>)

Research is ongoing, the juxtaposition of rendered 'cgi' video and the real object affording a unique opportunity to examine and evaluate the use of modern technology and imaging techniques in a traditional exhibition environment. A prominent artist and senior research fellow at the University of the Arts, London, Jananne Al-Ani, observed the imaging, again affording a unique opportunity to explore the intersection of three disparate disciplines, art, technology and cultural heritage.

Our research now focusses on use and usage of the model, as we investigate the potential of using these techniques within the cultural and heritage sector. This paper will present both the building and the user testing of the model, highlighting best practice and public engagement aspects of using 3D within museums and galleries.

¹ Hammer, Øyvind, et al. "Imaging fossils using reflectance transformation and interactive manipulation of virtual light sources." *Palaeontologia Electronica* 5.4 (2002): 9.

² Kersten, Thomas P., and Maren Lindstaedt. "Image-based low-cost systems for automatic 3D recording and modelling of archaeological finds and objects." *Progress in Cultural Heritage Preservation*. Springer Berlin Heidelberg, 2012. 1-10.

⁵⁵⁶ <http://www.hrionline.ac.uk/dhc/paper/18>