

CELEBRATING THE ENGINEERING DOCTORATE IN

MRTUAL ENVIRONMENTS MAGNGAND/ISUALSATION

AT UNIVERSITY COLLEGE LONDON

CELEBRATING THE ENGINEERING DOCTORATE IN

VIRTUAL ENVIRONMENTS WAGNGAND VISUALISATION

AT UNIVERSITY COLLEGE LONDON

Book editor: Dr Jamie O'Brien

Editorial assistant: Gwyneth Bradbury

Production manager: Charmaine Francis, CPR LDN

Designer: Britta Koch, CPR LDN

Print: Creative Print Run London LTD

www.cprlondon.com

With special thanks to all the contributors.

All enquires about EngD VEIV should be addressed to Dr Jamie O'Brien jamie.o'brien@ucl.ac.uk

Published by EngD Centre in VEIV Department of Computer Science University College London **Gower Street** London WC1E 6BT

First edition, 2013

All rights reserved.

No part of this publication may be reproduced or used in any form or by any means- graphic, electronic or mechanical, including photocopying, recording, taping, or information storage and retrieval systems- without permission of the publisher

© 2013 UCL, VEIV

Printed in the UK









CONTENTS

A key strength of the EngD VEIV Centre is the diversity of our portfolio. Many projects fall loosely into one of five categories: system interactions, animated bodies, dynamic environments, enhanced vision and intelligent materials. Even within these categories we find a rich variety of innovative technologies, engineering advances and professional interests. What binds EngD VEIV projects together is their common interest in enhancing our experience of the world, whether for entertainment, acuity or interaction. Other EngD projects relate to hardware architectures and robotics. In these project spaces we also find that the focus is the responsiveness and adaptability of computing technologies to challenging requirements.

INTRODUCTION

- 1 FOREWORD
- **3 THE ENGD VEIV**
- **4 INDUSTRY PARTNERS**

SYSTEMINTERACTIONS

A graduate of the EngD, Dr Nick 'Sheep' Dalton, gave the title for a ground-breaking study of the built environment, "Space is the Machine". The spaces of our interactions are configurable to our infinite set of requirements, and several EngD projects address the evolving interplay of humans and machines in a variety of settings.

- **10 Expressive robots** Dave di Ducca examines the ways in which the built environment might respond to human inhabitants with expressive, facial characteristics
- **12** Face-to-face encounters Ye Pan seeks to optimise the realism of facially orientated social encounters in teleconference systems
- **14** Affective masking Jason Drummond uses virtual environments to experiment with the cognitive effects of interacting with masks
- **15** Enhancing neuro-therapies Dr Jamie O'Brien used experiential mapping and VR simulations to enhance activity-based therapies

- **16 Making virtual music** Rebecca Kleinberger (a graduate of the MRes component) examined visual feedback for musical ensembles
- **17** Representing energy David Hawkins, Ian Henderson, Greig Paterson and Joe Williams developed a system to effectively represent household energy use
- **18** Learning in museums Tiphaine Bardon, Andrew Barnes, Ciro Bevilacqua, Gianfranco Gliozzo, Julian Hodgson and Dane Virk developed a project to enhance museum visitors' learning experiences using visual and auditory micro-augmentations
- **19** The Digital museum John Hindmarch uses scanning techniques to create and test virtual museums
- **21** Mobile games Adam Davis, Yotam Doron, Kazim Pal and Ye Pan built and deployed a mobile game in which players capture buildings in the city landscape
- **22 Positioning fidelity** Dr Mojtaba Bahrami studied the improvement in accuracy of global positioning systems
- **23 Public trust with GIS** Dr Artemis Skarlatidou investigated trust-building using GIS in public participation
- **24** Mapping healthcare Jess Wardlaw, an 1851 Industrial Fellow, is investigating the use of health mapping for the efficient regional delivery of care services
- **24** Tracking wildlife Gianfranco Gliozzo is deploying mobile GIS platform to help non-professional participants track wildlife

ANIVIATED BODIES

Audiences for graphics expect convincing experiences, which VEIV's scientists and engineers seek to achieve by balancing hardware capabilities with computer-driven, naturalistic effects. Several EngD projects advance computing technologies to achieve realistic effects for animation in a diverse set of applications.

- **26** Facial emotions Timothy Scully, 1851 Industrial Fellow, models the flow of blood into the face to simulate the realistic effects of emotional responses
- **28 Virtual clothing** Dr Bernard Spanlang developed ground-breaking methods to simulate clothing on virtual humans
- **29** Layering body parts Dr Simon Pilgrim advanced a technique for realistic character animation based on layered modelling of body parts
- **30** Animating figures Dr Anthony Ruto examined novel approaches to animating human figures from still models
- **31** Body scans Dr Avy Tahan pioneered techniques in point distribution modelling of human figures
- **31** Reconstructing faces Dr Howard Towner developed new methods for filling in missing data from faces
- **32** Articulated joints Jacques Cali, Dan Calian and Rebecca Kleinberger developed a breakthrough technique for printing 3D figures featuring articulated joints

DYNAMIC ENURONMENTS

Understanding dynamic environments requires tools that help us capture and analyse their complexities. VEIV engineers have developed numerous technologies to model changing environments with high confidence, and to transfer knowledge effectively for professional applications

- **34** Creating Content Gwyneth Bradbury is investigating semi-automated content-creation tools, which allow a games artist to rapidly create virtual landscapes
- **35 Mixing data streams** Dan Calian is advancing a technique to incorporate data from the environment into a rendering platform with the aim of reducing the time taken by a mobile graphics platform to render a scene
- **36** Fluid simulations Julian Hodgson, 1851 Industrial Fellow, works on the technically challenging problem of user-friendly fluid simulations for film animations
- **37 Mapping investment** Dr Patrick Weber developed geographic information systems for mapping London's inward-investment landscape, which helped planning and investment decision-making in the capital
- **38** Virtual neighbourhoods Dr Nick 'Sheep' Dalton advanced spatial analysis methods to help understand how a person can achieve a sense of place within a virtual world
- **41 Managing 3D assets** Jozef Doboš is developing techniques to manage complex collections of 3D models of built-environment objects for large-scale engineering projects
- **43** Automated buildings Fabian Wanner is advancing techniques to construct 3D models of urban scenes
- **44 Urban design feedback** Greig Paterson is developing a real-time feedback system to enhance the efficient design of buildings
- **47 Visualisation energy performance** David Hawkins is building and testing a visualization tool to support the refurbishment of buildings for a low-carbon future

- **48 Wind flow** David Andreen, Gwyneth Bradbury and Samuel Wilkinson advanced a system to model the flow of wind against computer-generated architectural forms
- Aero-dynamic design Samuel Wilkinson is advancing specialist tools for the simulation of wind over built-environment surfaces
- **51** Mitigating climate change Gurdane Virk is building decision-support tools to help planners and architects design urban environments to address the effects of climate change
- Fuel monitoring Ciro Bevilacqua is advancing methods to monitor and represent fuel consumption among communities
- Social arrangements Leto Peel is examining how the analysis of arrangements of people and objects can reveal the nature of their activities
- Understanding cost Dr Sarah Graham deployed an interactive system for representing the whole-life cost of construction
- Advanced design Francis Aish develops state-of-the-art visual tools for representing major architectural designs
- **Design geometries** Abel Maciel tests and deploys new techniques for geometric modelling of urban designs
- Agent simulations Erica Calogero advances techniques for high-resolution, agent-based modelling of people in urban landscapes

ENHANCED VISION

Digital photography, and other imaging technologies, can help us see features of the environment that are not available to the naked eye. The computing algorithms that process the imaging data can determine patterns at microscopic or highly complex levels allowing, for example, for potential structural problems to be anticipated, or moving elements such as human organs to be monitored during clinical procedures

- **58** Rail track monitoring Anita Soni is testing and deploying imaging technologies that monitor the signs of stress on railway tracks
- **60** Monitoring moving organs James Martin, 1851 Industrial Fellow, is advancing medical imaging technologies that track cancerous cells on a moving organ in real time
- **Document imaging** Tiphaine Bardon is testing the potential of molecular-level imaging for the analysis and preservation of rare archival documents
- **Document reconstruction** Kazim Pal is developing a range of techniques to digitally unfold fragile documents that have been damaged by water and fire
- Working with content Dr James Tompkin advanced a range of computational tools to support people in capturing, editing, and discovering media content
- **66** Analysing mood Yotam Doron is developing new methods for analysing critical factors in digital films that affect the scene's mood
- Time-of-flight confidence Malcolm Reynolds, Jozef Doboš and Leto Peel developed an innovative technique for establishing per-pixel confidence in time-of-flight cameras
- **Multi-view display** Tim Scully, Sarat Babu, Christos Sakellariou and James Martin built an innovative multi-view display system using an array of off-the-shelf projectors and a lenticular screen
- Thermal leak imaging Maciej Gryka, Fabian Wanner and Jessica Wardlaw adapted a thermalimaging system to detect heat leakage from a public building
- **69** Aesthetic judgement Cristina Amati is developing a range of techniques to help determine how a viewer regards an image as being beautiful, with applications in photograph manipulation
- **Segmenting Images**, Dr Alastair Moore enhanced object recognition by sub-dividing the digital image

INTELLIGENT MATERIALS

Since its inception, the EngD VEIV Centre has anticipated the application of advanced computing to the production of built-environment materials. Fabrication technologies now allow us to realise innovative forms that serve to enhance the responsiveness and efficiency of urban structures

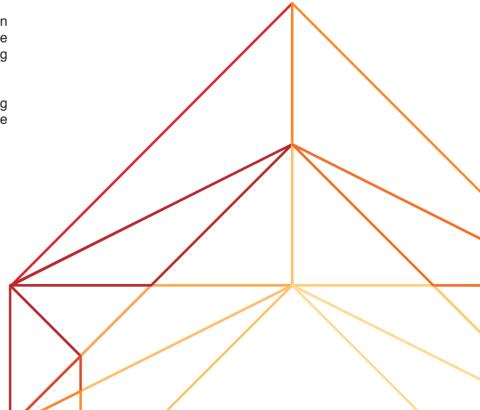
- **72** Responsive architecture Christopher Leung is designing and deploying responsive architectural components that allow a building to adjust to changing weather conditions
- **74** Generative forms Katrin Jonas is implementing innovative methods for generating complex architectural forms from simple components
- **76 Hybrid materials** Sarat Babu, 1851 Industrial Fellow, is creating hybrid materials that help adapt structures to the specific pressures of their environment
- **77** Articulated models Jacques Cali has pioneered a printing technique that creates human and animal forms with articulable joints integrated into the model
- **78** Intelligent materials Dr Sean Hanna built on a project with Dr Siavash Haroun Mahdavi and Dr Jamie O'Brien to design and deploy a technique for growing materials against certain environmental pressures
- **80** Functional complexity David Andreen is advancing the application of complex fabrication to enhance the performance of materials

MISC: HARDWARE AND DESIGN

- **82 High-end fabrication** Dr Siavash Haroun Mahdavi built on his EngD research in adaptive robotics to launch among the earliest companies to specialise in high-end 3D fabrication
- **83** Systemic architecture Christos Sakellariou is designing and building innovative computer architectures based on adaptable forms in nature
- **84 Dynamic metrology** Andrew Barnes is building and testing robotics architectures for the efficient measuring of manufactured components
- **85** In-between contexts: Karen Martin has critically explored ways of enhancing people's experience while they wait for something

POSTSCRIPT

87 POSTSCRIPT



INTRODUCTION

FOREWORD

Prof Anthony Steed, Former Director and Co-founder

When we wrote the original proposal for the EngD VEIV in 2000, we knew that visualisation, imaging and virtual environments technologies would become more vital to science and engineering over the next decade, but we under-estimated the range and diversity of application areas. With the constant improvement and accessibility of visualisation systems, and the associated increase in computing power from graphic processing units, many more application areas have been enabled that would otherwise remained the preserve of dedicated visualisation facilities. Coupled with accelerating change in other hardware areas such as new types of camera, mobile devices, tracking and sensing platforms, and additive manufacturing, whole new areas of research would become practical, such as automated design engineering, crowd-sourcing data sets, immersive installations, game-inspired spatial analysis, and so on.

This book reflects this range and diversity of application areas. Over the lifetime of the Centre we have worked with companies both large and small in a range of industry sectors from construction and manufacturing to video games and film production. The experience of our graduates is that they may work across a range of different industry sectors employing knowledge gained in one sector in others. This requires a diverse skill-set and we have set out the programme of the Centre to prepare students for a variety of working practices to be found in the different industries. This involves working across academic disciplines, but also developing a sensitivity to the nature of the industry, be it one with carefully managed multiyear projects or one where short-term entrepreneurial projects are the norm. As can be seen from the projects in this book: some were realised within large corporations as new services or products and others led to the foundation of new companies.

Another aspect of diversity in the EngD VEIV has been the relationship with the sponsoring company. With some of our SMEs, EngD students have effectively been the main blue-skies R&D personnel, looking at potential new technologies for application in the market or the impact of upcoming technology changes. In other cases, the student has been part of a large R&D team, but focussed on longer-term projects. We have had excellent feedback about the impact of the EngD projects with R&D departments and beyond. The students' research

itself has been valuable, but very significant value has also been found in having someone around who is up to date in a field, and has access to other state-of-the-art research being done by other students (PhD and EngD). Indeed, informal side projects between students have led to new projects within and between sponsor companies.

An EngD project is a serious commitment from all parties. We have found than many of our EngDs are part of a growing involvement between the sponsor company and university research group. I would highlight, as a personal example, my relationship with Electronic Arts (EA). The early relationship was recruitment: EA came to UCL to present to the undergraduate and taught postgraduate students with an aim of attracting students into the games industry. Over a few years, and based on the growing awareness of research in graphics and interactive technologies, this developed into collaborative student projects. As a specific challenging technical need emerged, (the migration of game production pipelines to the Sony PlayStation 3/Microsoft XBox 360 generation), EA sponsored two EngD students. They were looking at how content could be developed for the significantly more power platforms without simply expanding budgets. This very productive engagement then led to me taking an Industrial Secondment to EA for six months sponsored by the Royal Academy of Engineering. I worked on a game team to get a better understanding of the engineering process and to work on some game production tools. This experience and the results of the EngD research meant that we were able to establish a significant presence in research in this area, which this subsequently led to further collaborations with EA and other games companies.

The growth of the EngD has been paralleled by a significant investment by UCL in growing the associated groups. In 2001, when the centre was founded, there was a single research group covering vision, imaging and virtual environments, with six academic staff. Since then we have split VIVE into two groups (Vision and Imaging Science and Virtual Environments and Computer Graphics), both of which are now larger than the original VIVE was. Our partners in the Bartlett have similarly grown, as have associated groups such as UCL Interaction Centre and UCL Civil, Environmental and Geomatic Engineering.

Finally, I hope that the EngD VEIV Centre will be able to continue to support such a compelling portfolio of projects. It is a pleasure to see the wide range of excellent research in one place. I hope this volume inspires new research directions using virtual environments, imaging and visualisation technologies.

A Brief Introduction to the EngD VEIV

The EngD Programme

The Engineering Doctorate Centre in Virtual Environments, Imaging and Visualisation (EngD VEIV) was founded in 2001 at University College London in the Department of Computer Science and the Bartlett School of Architecture. Funded mainly by EPSRC the EngD VEIV supports research engineers (REs) studying in areas related to the use of visualisation and imaging in engineering and design. An EngD is a doctoral degree similar to a PhD, DPhil or DSci. However, unlike those degrees, and more like degree programmes in Europe and North America, the EngD is four years in duration and contains some taught elements. The EngD is thus a relatively new type of degree for the UK. The EngD is targeted at leading students and engineers who want to direct research in industry, or want to maintain significant links with industry in an academic career.

Every EPSRC-funded RE has an external sponsor, and normally the RE will work significant proportions of their time at that organisation. It is also possible to do an EngD as a secondee from an existing position in a company. This makes a doctoral degree much more accessible for senior engineers with established careers and positions. The taught programme, leading to the qualification of MRes, covers a broad range of topics, from basic engineering skills (design process, software engineering) through to business development courses. At UCL, REs have the opportunity to apply for courses at the London Business School.

Sponsors and Students

As of April 2013, 75 REs have entered the course. Their backgrounds range from recent graduates, through to qualified architects and engineers with significant career experience. The REs are sponsored by a wide variety of collaborative partners including SMEs, charitable organsiations and global businesses; over 65 companies have sponsored EngD projects since the Centre's inception, including some that have entered novel partnerships with other sponsors.

Financial Details

The EngD VEIV was founded in response to an EPSRC call to found new centres for doctoral training as a consequence of successful pilots in other centres. The EngD Centre in VEIV has repeatedly been successful in securing new funding and has been rated as excellent in reviews. Unlike a traditional studentship, in the EngD model the industrial sponsor provides significant financial input to the project and provides resources to the student. In 2009 the EngD VEIV centre secured funding for fifty new projects in the years 2009 through 2013. As of the writing of this book, we are bidding to continue funding through to 2018.

INDUSTRY PARTNERS

The EngD VEIV programme helps to enhance the design benefits of Arup Foresight by supporting a doctoral student in evaluating design tools for complex engineering challenges. Our office seeks to advance Arup's performance by driving radical innovations in major infrastructural projects. The success of our innovations depends on their adoption or acceptance among the public, so capturing support is, in this regard, a key driver for success.

Arup Foresight's EngD student has dedicated his efforts to deploying a social-media public engagement platform, upon which Arup's 3D assets are provided as enriched models of structural innovations. This web-enabled platform allows Foresight to engage a massive and diverse population of participants, which represents a qualitative shift in the scope and validity of our design consultations.

A key advantage for Arup of sponsoring an EngD VEIV student is the inclusion in our team of a researcher who is dedicated to evaluating virtual design tools. Our success as engineers depends in part on the effectiveness of our tools. With so many available to us, it is essential that we make clear, informed judgements about their appropriateness to the tasks at hand. Testing new tools involves prototyping, which is a valuable task for an EngD student to undertake. In this way, our EngD student helps Foresight to think in new ways, and to imagine solutions to the problems of the next generation.

The UCL VEIV EngD on Non-Contact Monitoring on Railway Infrastructure is a collaboration between Thameslink, Crossrail Ltd, LUL and Rail Industry Partners on assessing the application of laser scanning technology to structural movement monitoring on railway projects. The key research goal is to analyse the current movement monitoring systems, which are based on physically attached target monitoring of structures at discrete locations, and assess the advantages and capability of laser scanning to overcome the logistical limitations of this approach while achieving the high level of accuracy required.

The secondary aim of the study is to demonstrate how the vast quantities of data generated by laser scanning can be managed to provide fast and reliable monitoring solutions which can be presented to engineers and stakeholders in a manner which allows them to understand and interpret the data more readily.

The study focuses on a small number of specific live-site monitoring scenarios on the Thameslink Programme, a £5 Billion Upgrade to Major Stations and Signalling Systems which runs through central London (from St Pancras to London Bridge).

Dr Alvise Simondetti, Global Leader of Virtual Design, Arup Ltd



Barry Gleeson, Senior Programme Engineer (Survey Systems), Thameslink It is expected to provide qualitative and quantitative analysis to inform Network Rail Thameslink, and the UK rail industry generally, on the practicalities and possibilities of adopting such technologies in the future.

Dr Daniel Hulme, graduate EngD VEIV, CEO and Founder Satalia Ltd



I couldn't have chosen a better career path than when I decided to do an EngD at UCL. One of the main objectives of the EngD programme is to create the next generation of academic and industrial leaders, and the programme does exactly that. I was lucky enough to do what was essentially a Masters and PhD, plus developed invaluable knowledge and skills to understand technology transfer and innovation that you wouldn't typically get from doing a traditional PhD.

My research focus was in the area of optimisation. I designed and implemented optimisation algorithms that could solve real-world problems faster and better than state of the art, which had wide-ranging applications in both academia and industry. Subsequently I span out a company loosely inspired by this research that was funded and supported by UCL. As CEO of this company (Satalia) I can testify to how well the EngD prepared me for this role.

The EngD is an essential degree for those interested in science and the pathways to impact of research, which are both critical in securing academic funding or being a successful entrepreneur, engineer or manager. The EngD gave me the option to have a career in academia and equipped me to be a successful influencer in business and industry.

"Investing in Research Engineers" Dr Duncan Wilson Director, Intel Collaborative Research Institute for Sustainable Cities



The Intel Collaborative Research Institute for Sustainable Connected Cities thrives on the collaboration between world-class researchers in academia and those at the cutting edge of developing computing technology to connect and enrich the lives of every person on Earth. The partnership between Intel, UCL and Imperial is founded on the principle that bringing together a diverse cohort of researchers will lead to novel ways of thinking about how we can enhance the social, economic and environmental well-being of our cities. At the heart of that research are the PhD's and EngD's who are developing insights into problems that are the foundation for the projects we deliver.

In our experience the Engineering Doctorate attracts researchers who enjoy the parallel worlds of developing thought leadership while also wanting to develop and deploy their ideas at scale and in a commercial context. From the Intel perspective this technology focus aligns directly with our strategic collaboration with the academic community. The commercial training offered through the programme complements the exposure the EngD research engineers obtain through working closely with an industrial partner. The training is not just theoretically taught but also becomes practically tested through working on real industrial projects.

At UCL we have the benefit of having London as our living laboratory. We are making a significant investment to deploy a city platform to enable researchers and engineers to test their ideas for a sustainable connected city. Achieving this vision will entail investigating, developing and deploying adaptive technologies that can optimise resource efficiency, and enable new services that support and enhance the quality of life of urban inhabitants and city visitors. There are many fundamental technical, social and urban challenges and opportunities that need to be addressed to accomplish this vision. The Institute embraces an interdisciplinary research approach, combining methodologies from computer science, the social sciences, interaction design and architecture to improve how cities are managed and maintained in order to ensure and enhance citizen well-being. It is the intersection of these disciplines, approaches and methods that are the bread and butter of the corporate research laboratory and as such the fertile ground for the industrial research of the EngD.

An Engineering Doctorate is different. It produces a different kind of researcher. Standing with one foot in the corporate lab, the other in academia, and communicating with supervisors who exist in both. That can be tough, since different environments have different pressures which often have to be managed by the person in the middle. But that is where the strength of training is most important. The product of such an environment is a researcher with academic rigour who can also appreciate the practical constraints of a commercial environment. It is quite timely that I am typing this note while also trying to find the next member of our team in London, and guess what, there are a couple of candidates who have been through the EngD process.

As somebody who has been involved in the digital industry on the architectural side from the 1980's it gives me great pleasure to contribute to the discussion of what might happen in the future. Clearly a new generation of designers – both of systems and system users – need educating and allowed exploration into the only technology that offers a survival strategy for mankind.

Virtual environments have grown from one-off rendering shots with ray-tracing in the 80's (I used one of the first commercially available systems then) to total environment immersion in today's computer games. The focus has been to replicate reality, but increasingly the virtual world is acquiring its own 'mores and morals', some with unexpected consequences – the Arab spring comes to mind. Social media has brought its own successes and startling dangers.

The virtual environment mimicking the world is now expanding into simulating proper physics – from environmental control to structural analysis. Many animation tools and techniques exist to map this data to the 'dumb' 2D screen of your PC or tablet. New techniques and interaction media will need to be developed and shared to make better communication between man and machine.

"The Next Generation
Virtual Environment
and Interaction
Engineering"
Lars Hesselgren,
Director of Research PLP
Architecture



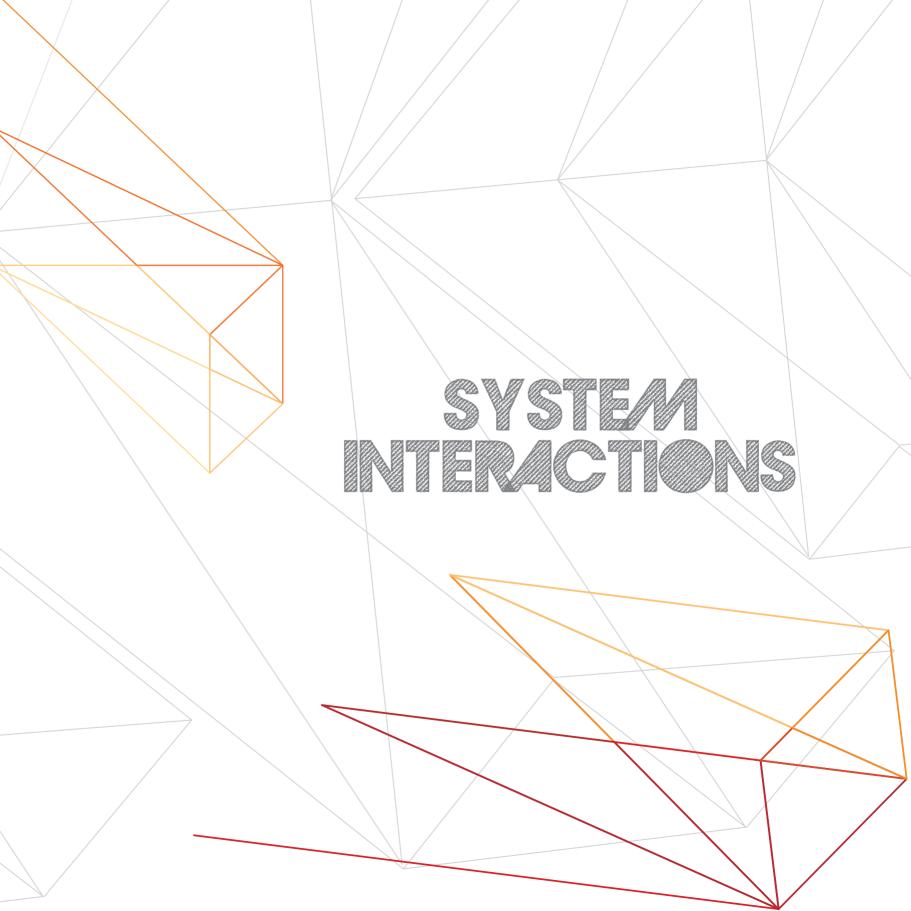
Steve Jobs did the world a huge favour by moving the interaction paradigm from the mouse to the swipe. The new gestural interfaces are coming on, and clearly have a long way to go. Ever tighter integration to the personal sphere will come with systems such as Google glasses, predictive system responses, and direct mind reading.

However it is worth reflecting that much of the success of the modern world is based not on direct interaction but by abstraction – particularly mathematics as a tool which also underlies all digital systems. Easier access into this abstract world and its manipulative tools underlie the resurgence of interest in parametric and generative systems.

Many other systems will adapt to the wide and cheap availability of computing power. We will move from easily digitisable systems such as books, music, TV, movies to systems with intelligent behaviour. Cars are currently developing into new transport systems potentially allying the strengths of public systems into those hitherto regarded as private systems. One implication of this is regaining urban space back from the car to the humans; increasingly cities compete in human attractiveness as much as in location (location being irrelevant in a digital world).

Tall buildings are immensely successful, increasing numbers are being built. They are an essential tool in building dense and sustainable cities and megacities. The design tools available to the construction industry professions leverage the vastly greater computational resources available in the cloud and elsewhere. Traditional design approaches will examine a range of alternatives – typically in the hundreds. New technologies such as parametric and generative design tools allow orders of magnitude (millions rather than hundreds) larger solution sets to be both to be generated and assessed. Evaluations are also being automated (using tools such as genetic algorithms) – the fitness criteria of any optimisation are critical to the success of such an approach.

The role of education is vital in the area both of forming young minds and allowing research into the future. The EngD Centre in VEIV is such an institution and I await the future with bated breath.





Robotic faces respond expressively to a human participant

Expressive robots

Dave di Ducca & Jason Bruges Studio

Supervised by Prof Stephen Gage, UCL Bartlett School of Architecture

There is a common assumption that we all experience and inhabit the same environment. This project investigates design affordances, human/non-human interaction, and cybernetics, developing and building interactive test installations. This normalising of perception serves many purposes. In everyday life, it allows us to navigate and socialise, refer to, and share our environment with others. Researchers interested in perception, however, have long suspected this is not the case. Instead suggesting that we each form a unique construction of our environment. Broadly, the question from a designer's perspective is: if we each find unique experiences in our world, why is the world around us so often designed with the assumption of a common experience?

This research focuses on the installation of person-responsive objects in art and commercial environments. Experience in these environments is spatial and for such projects to be commercially and curatorially desirable, objects must seek attention. We can then begin to determine whether different classes of audience react differently.

Our human/non-human inter-reaction perspective differs somewhat from that of human-computer interaction which operates with an assumption that people want to communicate in an aspatial, goal-directed environment. Observers approach interactive objects and installations in real-life conditions, in different ways and with varying expectations based on individual prior experiences. Increasingly, these

observers assume that objects in the environment which change state are doing so because they are observing people's body and facial movement. Therefore, the channel of communication is open. We notice also that the observers are inclined to invent complex interpretations of what are often simple reactions.

Since we now have the ability to engineer much more sophisticated installations, how do we do this and will it allow us to more precisely target the effects of these installations?

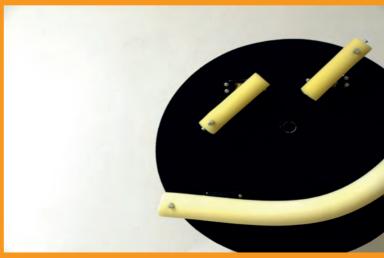
The focus of this research project is the development of spatial installations which evolve their output through interaction with observers. Through this they develop predictions of how to behave towards new observers, forming a low level of artificial intelligence. The project is an investigation into the possibilities of a multi-modal architecture, capable of evolving an autonomous method for choosing how to present itself to observers.

Personal statement

While studying on the MArch/DipArch programme at the Bartlett School of Architecture, UCL (2008-2010), I began to focus on embedding technology in spatial situations; to explore how interaction can be facilitated by integrating sensing and computational technology.

My recent research installation, "I Have these Two Faces"*, combined embedded computing, computer vision and genetic programming techniques. The installation comprised two animatronic objects, resembling simplistic faces. Each face has a web-cam, used to detect and track observers. Using common computer vision techniques, the installation can infer simple characteristics of observers. The behaviour of the installation evolves over time, though a genetic algorithm in the software. By inferring differences and responses in observers' characteristics, the installation evolves to present itself to different observers in a manner informed by interactions with previous observers, with similar characteristics. For example, at one exhibition the installation learned to smile at children and frown at adults.





Credit: Paul Clarke

Face-to-face encounters

Ye Pan, support by UCL Overseas Research Scholarship

Supervised by Prof Anthony Steed, UCL Computer Science

The movement of human gaze is very important in face-to-face conversation. Some of the quality of that movement is lost in videoconferencing because the participants look at a single planar image of the remote person. We use an array of cameras to capture a remote user, and then display video of that person on a spherical display.

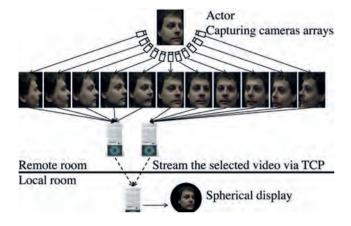


Figure 1

The goal of the system is to allow local users to perceive the eye gaze of a remote user accurately. *Figure 1* depicts the system design. A remote user, the actor in the remote room, is captured by eleven cameras controlled by two PCs. In the local room, a single PC renders video on a spherical display which is seen by a local user, the observer. Depending on the observer's position, the most appropriate camera feed is streamed from one of the two camera controller PCs to the renderer PC. Streaming is done using TCP.

In the remote room, eleven low-cost PlayStation Eye USB digital cameras are mounted on a half-annular table (see Figures 2a and b). The cameras capture at 30 Hz at 320 x 240 pixel resolution, which provides qualified videos while saving bandwidth for real-time streaming. The cameras are manually adjusted to look at the point above the centre of the table. Calibrating the cameras gives their positions and orientations accurately which are used in the rendering process, described below.

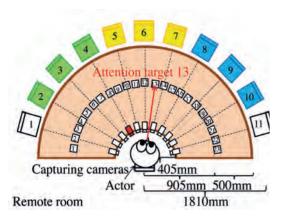


Figure 2a

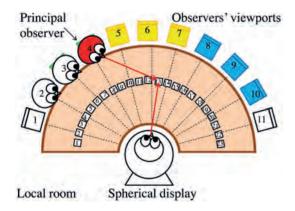


Figure 2b

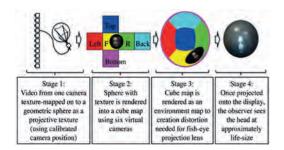


figure 3

In the local room, a spherical display is located at the centre of another half-annular of the same size. Eleven observer viewports set around the table correspond to each camera in the remote room. The spherical display is a projection device with a 16" sphere-shaped surface and an internal fisheye lens to project imagery on to the inside.

The presentation of the remote participant onto the sphere is done in four main stages shown in Figure 3. First, a sphere acts as a proxy geometry of a human head, onto which the video images are displayed using projected texture mapping (PTM). According to the observer's viewport, the video captured by corresponding capturing camera is selected. This video is projected onto the polyhedron. which is approximately human head-size. This ensures that the capturing camera and the observer's eye are in close alignment. Next, we render the scene into an environment map using six cameras positioned outside the cube at the level of the observer's eye. We then draw a 3D sphere, using an environment map, which simulates the reflectance of a surface by using the eve vector as a lookup into the texture. rather than a simple texture co-ordinate. This spherical environment map is suitable for projection through the fish-eye lens, allowing us to acheive 360 degree horizontal visibility, and the observer sees the life-size, projected head. The next stage is to compare the spherical display to a faceto-face setting and a planar display to test the effectiveness of the camera array and spherical display system. This allows observers to accurately judge where the remote user is placing their gaze.

Jason Drummond & Anthropia

Supervised by Prof Anthony Steed, UCL Computer Science

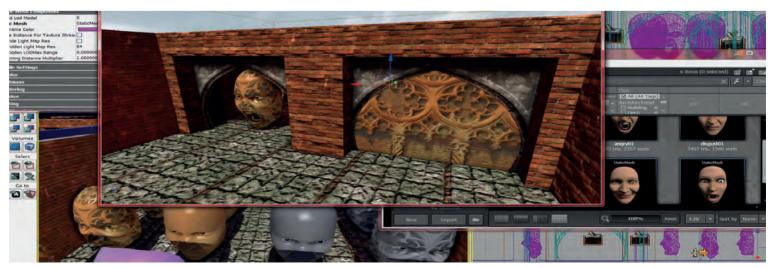
This project makes use of virtual environment (VE) properties that go beyond real-world simulation. The core technique used is based on research found within perception science, an effect known as backwards visual masking.

Here, a rapidly displayed target image is rendered explicitly imperceptible via the subsequent display of a masking image. The aim of this project was to investigate the potential of visual masking within VEs to induce changes in affective cognition. Of particular importance would be changes in a positive direction as this could form the foundation of a psychotherapeutic tool to treat affect disorders and other conditions with an affective component.

The initial pair of experiments looked at whether visual masking was possible within VEs, whether any measurable behavioural influence could be found and whether there was any evidence that affective cognitions could be influenced. It was found that the technique

worked and could influence both behaviour and affective cognition. Following this, two experiments looked further at parameter manipulation of visual masking within VEs with the aim of better specifying the parameter values. Results indicated that the form of visual masking used worked better in a VE when the target and mask were both highly textured and that affective effects were modulated by the number of exposures of the target. The final pair of experiments attempted to induce an affect contagion effect and an affect cognition-modification effect. An affect cognition-modification effect was found whereas an affect contagion effect was not.

Overall, the results show that using visual masking techniques within VEs to induce affect cognition changes has merit. The study's finding lays the foundation for further work and supports the use of this technique as basis of an intervention tool.



A virtual laboratory for testing backwards visual masking

Enhancing neuro-therapies

Dr Jamie O'Brien & NHS

Supervised by Prof Alan Penn, UCL Bartlett School of Graduate Studies

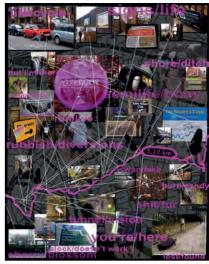
A patient experiencing neurological impairments must re-train his or her everyday skills in environments that reflect their real-world requirements. Clinical environments do not usually offer the level of functional complexity and stimulation that is required to optimise 'occupational' training. The research involved working with neuro-rehabilitation patients to find out which sorts of tasks they found meaningful or purposeful. Tasks were then simulated using virtual reality, and patients were invited to undertake their preferred tasks in the VR environment while their state of emotional arousal was measured as a gauge of immersion in the task. This level of immersion was assumed to correspond to the effectiveness of the VR and provide realistic simulations of therapeutic task environments. Innovation assessments also found that clinicians are very open to new technologies and methods that might help their patients, but also that the institutional pressures of healthcare can limit innovation. A later stage of the research used cognitive and experiential mapping to gain a deeper undertanding of everyday requirements of people with neurological impairments (pictured here).



The EngD programme enhanced my technical and research skills to help advance VR as a state-of-the-art technology at the forefront of clinical practice. The business training component of the programme included an Enterprise Scholarship at the London Business School's Centre for Scientific Enterprise, which helped me to understand the importance of applying an innovation strategy to the technology's practical implementation.

I graduated in 2008 and, since that time, have developed my research as a Visiting Fellow at Lancaster University's Centre for Mobilities Research, where I am working on a book about the system dynamics of innovative landscapes, which involves at the complex interplay of people, knowledge and space.

I returned to the VEIV in 2009 as a member of staff to manage the EngD Centre. I continue to catalyse relationships with industry and collaborations within our rich industrial-academic network, and to support students in benefitting to the greatest advantage from the EngD programme.



Credit: Kai Syng Tan





Making virtual music

Rebecca Kleinberger (Masters in Research graduate)

Supervised by Prof Anthony Steed, UCL Computer Science

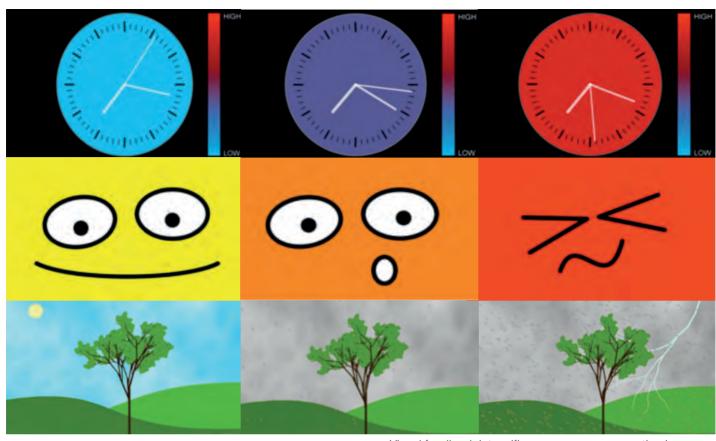
The experience of sound is incredibly important for an immersive experience: since the beginning of cinema, not only music but also sound effects were often produced in real time in cinema halls at each showing. While the graphic technologies have reached a very high level of realism during the past few decades, the quality of the audio rendering has only recently become of interest for the enhanced realism of animations. We see that the technologies commonly used can be surpassed by modelling the different sound phenomena from a new perspective. By using physically based sound modelling, we achieve a more accurate result and we also have better control over the system of the parameters have a strong semantic interpretation.

Concurrently, it is well known that non-musical communication between musicians is important during performances. Among all the meta-data produced in concert, gesturing (from eye contact to beat tapping to anticipative motions) is a privileged channel conveying intentions, enabling everything to synchronise, and adapt in real time. But keeping an eye on other musicians while focussing on one's own playing can be a perilous exercise. We believe that processing the information with an automatic visual system would improve the ensemble experience. In the same way, we investigate the optical treatment of gestural information by a leader with its transmission to other members.

Rebecca Kleinberger is currently a PhD candidate at MIT's Media Lab on the project Opera of the Future



Optical feedback enhances the musicians' performance



Visual feedback intensifies as energy consumption increases

Representing energy

Group project with David Hawkins, Ian Henderson, Greig Paterson and Joe Williams

A 'user-friendly energy visualisation' study was carried out in a London house to test three different methods of displaying real-time electricity use. The visualisation themes – functional, emotional and alert - were selected to test research theories on home energy monitoring systems. Each visualisation was assessed in terms of its impact on the occupants' energy awareness and energy-using behaviour. A monitoring display facility, accessed by tapping the visualisation on a touchscreen monitor, displayed daily budgets, current electricity use and electricity use history graphs.

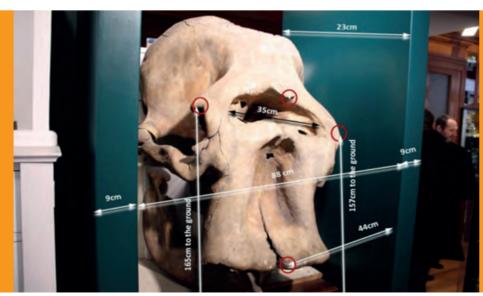
It was found that user interrogation of the system increased during periods of high electricity use. The family stated

That the system had helped to improve awareness of their electricity use, and that their household energy efficiency had improved. They expressed a preference for the alert-themed visualisation – a weather scene – and found the daily budget facility particularly useful. The exercise was successful in delivering a technical framework for testing energy use visualisation methods and the findings provided useful recommendations for similar future investigations eased during periods of high electricity use. The exercise was successful in delivering a technical framework for testing energy use; visualisation methods and the findings provided useful recommendations for similar future investigations.

Learning in museums

Group project with Tiphaine Bardon Andrew Barnes Ciro Bevilacqua Gianfranco Gliozzo Julian Hodgson and Dane Virk

Micro-Augmentations in museums have been proposed as a novel method to enhance a visitor's experience. These would be in the form of a subtle and short auditory or visual stimulus, related to an exhibition provided to the visitor. The group was given the task of designing a technological system to produce micro-augmentations in UCL's Grant Museum of Zoology. Two types of augmentations were proposed, an audio and a visual. The two groups independently investigated the most appropriate technologies, considering the time limitations, and the contextual constraints of the Grant Museum. The audio group decided to implement a directional audio speaker, which was triggered by an ultrasonic sensor. The visual group designed a mobile application that would augment features of an exhibit. Both groups initially installed their technologies separately, observing their effectiveness and consequently improving the concepts and design. The final installation was combined in order to test the use of both stimuli on one exhibit. The success of the combined. compared to the separate installation, highlights some key questions concerning micro-augmentations. The success of an installation will depend on how it enhances a visitor's experience, but the balance between this enhancement and the subtlety of the stimuli is a major design and curating challenge.



The UCL Grant Museum's elephant's skull is the focus of an AR application

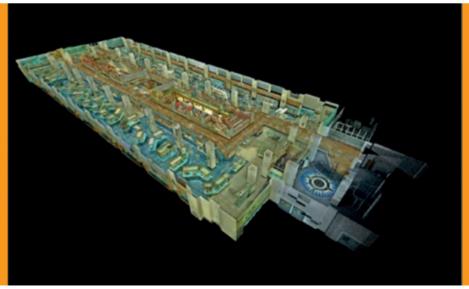
The digital museum

John Hindmarch & The Science Museum

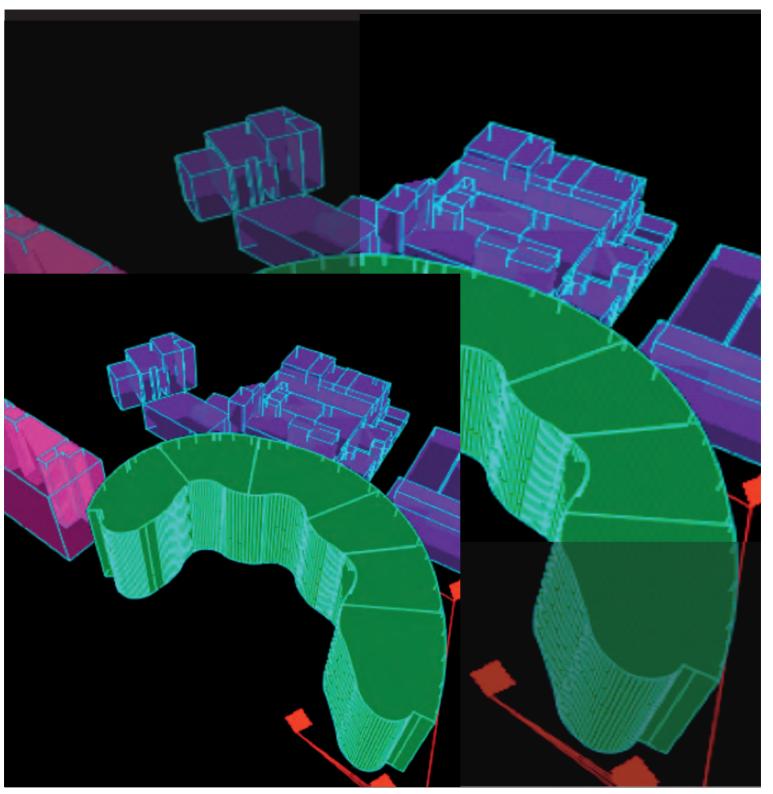
Supervised by Dr Melissa Terras, UCL Centre for Digital Humanities and Prof Stuart Robson, Department of Civil, Environmental and Geomatic Engineering

This project examines the impact of digitisation on the experience of learning in museums. One of the main aims of 3D digitisation in a cultural heritage context is 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 undertake, 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. Specifically, in 3D digitisation, the intention is to create a dataset which records the surface properties of an object (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). For a record to be considered objective. it would need to fulfil criteria such that if a curator were to carry out measurements on the dataset and repeat those measurements on the original object, he or she 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.

3D digitisation is a complicated and expensive process, and one could argue that the closer we approach to maximal objectivity the more expensive and time consuming it becomes; similarly, the more authentic we want our model, the more complex and resource hungry our rendering process becomes. However, it is possible that the use of 3D models may have the potential to add value to museums' visitor-experience remit. Yet without quantifiable metrics, it is impossible for institutions to make educated decisions as to how virtual objects can further their aims, while at the same time, if unrealistic expectations are placed on the models it is quite possible that a single unsuitable use of the technology could turn an institution or individual against other uses in the future. If we can identify the possible uses that virtual objects may have for museums, the areas where 3D models may add the most value, and provide experimental evidence to support this and measure the utility of such objects, we can prevent institutions wasting time and money on illadvised projects and properly advise curators on getting the most value from the digitisation.



A virtual model of the Science Museum's Shipping Gallery was created over five nights and from 275 individual scans



A mobile games involves the capture of buildings in an urban space

Mobile games

Group project with Adam Davis, Yotam Doron, Kazim Pal and Ye Pan

The purpose of the project was to design and implement a pervasive location-based, mobile-phone game, which makes use of three-dimensional map data of London. Location-based mobile applications typically do not have access to such data, and only use a two-dimensional map. This limits the usefulness of the application and the types of interaction that are possible. For example, augmented reality applications such as Layar, which overlay geo-located points of interest onto the phone's camera feed, only make use of the user's position and the position of the point of interest. They don't know anything about the geometry of the environment, and thus will often display information that is of limited use, and can sometimes be misleading. For example, with no knowledge of the environment geometry, it is impossible to know what the user can actually see from their position, and what is obscured from their view. One of the aims of this project is to use knowledge of the three-dimensional geometry of London, to perform line-of-sight computations between player locations, as well as provide three-dimensional visualisations of the game.

The framework developed in this project was only possible due to then -recent advances in mobile technology, such as the widespread adoption of 3G internet and the development of smartphones powerful enough to render 3D graphics and perform 3D mathematical operations. Prior to the development of these technologies, an application such as ours would have required a specialist piece of hardware to serve as the mobile client. A few such platforms had been developed, such as the Gizmondo in 2005, but were not well received and would not offer a very large user base. However, modern mobile phones provided us with the perfect platform for our application, since they are powerful enough to perform all the necessary computation, widespread enough to provide a substantial user base, and come with a ready-made infrastructure such as 3G internet, and App Stores for distributing the application to users.

Positioning fidelity

Dr Mojtaba Bahrami & Ordnance Survey

Supervised by Prof Marek Ziebart, UCL Civil, Environmental and Geomatic Engineering

This project provided techniques to enable ordinary end-users to instantly improve the accuracy and precision of their raw GPS/GNSS positioning in urban environments. In October 2009 Mojtaba won an international Student Paper Award from the U.S. Institution of Navigation (ION) to present his work at the ION GNSS 2009 conference, the 22nd International Technical Meeting of the Satellite Division of the U.S. Institute of Navigation. Mojtaba presented his technical paper "Getting Back on the Sidewalk: Doppler-Aided Autonomous Positioning with Single-Frequency Mass Market Receivers in Urban Areas" at the conference, in the session on Urban and Indoor Navigation Technology.



Global positioning systems achieve greater fidelity

Public Trust with GIS Artemis Skarlatidou & Arup Prof Tao Cheng, UCL Civil, **Environmental and Geomatic Engineering**



GIS promotes public trust in the decision to locate a nuclear power plant

Maps have a long history in the communication of spatial information, yet Web Geographical Information Systems (Web GIS) expanded map use to a wide variety of contexts and to include people who do not have profesional expertise with spatial and GIS issues. This non-expert interaction with Web GIS generates Human-Computer Interaction (HCI) implications. While such HCI elements as usability have been the focus of GIS research. additional HCI aspects, such as trust, were overlooked. The significance of trust in non-expert interaction with Web GIS becomes more apparent as these tools are used to engage the public at different levels of public participation. The public participation literature suggests that when Information and Communication Technology (ICT) mediums such as Web GIS are used to engage the public, it is essential that they improve public knowledge and trust. This project examined how this can be achieved, using the case of the site selection of a nuclear waste repository in the UK.

Firstly, the EngD research focused on an HCI-based investigation of existing Web GIS applications to understand the functional and perceptual attributes that influence non-experts' trust perceptions and introduced a set of trust guidelines. These guidelines informed the

development and design of the PE-Nuclear tool, a Web GIS to inform lay people about the site selection process of a nuclear waste repository in the UK. Secondly, the Mental Models approach was used to support the development of the PE-Nuclear tool's information content based on lay people's mental models, needs and expectations. Finally, the tool was evaluated to investigate separately whether the trust guidelines and the information content improved public trust and knowledge.

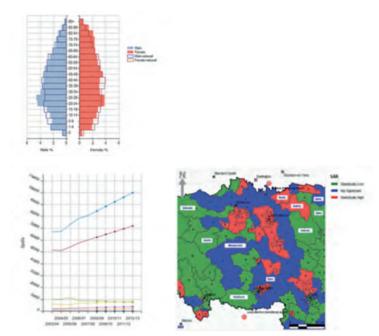
The research findings and methodological framework provided a holistic approach for the development of Web GIS applications, which have the potential to enhance public knowledge and help non-experts develop rational trust perceptions, protecting them from unethical and inappropriate use of the technology. Furthermore, this research supported the identification of critical gaps and methodological implications that should inform future GIS research, especially of an HCI interest. The scientific knowledge gained through multidisciplinary nature of this research contributed to other fields such as Risk Communication, Public Participation, but also provided important lessons to inform the current Nuclear Waste Management Programme in the UK.

Mapping healthcare

Jessica Wardlaw & Dr Foster Intelligence 1851 Industrial Fellow

Supervised by Prof Muki Haklay, UCL Civil, Environmental & Geomatic Engineering

The aim of the project is to introduce GIS techniques into the analysis and visualisation tools used in the processing of healthcare statistics. The project will focus on a range of geographical visualisations to improve use and accessibility of information across the health and care system, in support of the overall aim of giving people more choice and control over health and social care. This will be done by adding usable and accessible GIS as an integral part of management information software solutions to support better commissioning, choice, quality and efficiency: all areas of the NHS reform programme where better use of information by patients, users, professionals and managers is vital.



Detailed health mapping enhances delivery of care services



Wildlife monitoring is improved through mobile technologies

Tracking wildlife

Gianfranco Gliozzo & Zoological Society of London Supervised by Prof Muki Haklay, UCL Civil, Environmental and Geomatic Engineering

The research is framed in the general ExCites 'Extreme Citizen Science' research group which is aimed to extend the range of applications, participation and engagement of people in citizen science projects. The specific research project focuses on developing novel applications of geographical mobile gaming that will be used in the Zoological Society of London activities in monitoring biodiversity and promoting conservation.

The main focus of the research is the engagement of new types of participants in these monitoring activities, from different age and interest groups. The games can be either mobile- or webbased with the target of raising awareness and participation on the conservation and biodiversity initiatives of the ZSL.



Facial emotions

Timothy Scully & Procter & Gamble 1851 Industrial Fellow

Supervised by Dr Tim Weyrich, UCL Computer Science

Humans are very adept at facial recognition, with a significant portion of our brains dedicated to analysing both facial features and skin colour. Infants, for example, are able to mimic facial expressions from two days old. It is this adeptness that demands high realism from a computer-generated rendering of a human face.

Realistic computation of skin colour has been tackled previously by the computer graphics community. Skin colour is dominated by two chemical colorants, or chromophores. The first of these chromophores is melanin which the skin produces in response to UV light, and is commonly known as our 'tan'. The second is haemoglobin, which is a constituent of our blood. Combining information about the concentrations of these two chromophores with a simulation of skin interaction with light, allows computer graphics software to compute its colour.

Using these models it is possible to describe completely skin colour by measuring the concentrations of these two chemicals and making assumptions about the skin's structure. We can then compute skin colour for a wide range of both blood and melanin concentrations. Using these computed colours we can "invert" digital photographs to split the recorded skin colours into melanin and haemoglobin concentrations.



Realism in facial expressions is achieved by simulating blood-flow

This project also measures and models dynamic changes in our haemoglobin concentration. Melanin concentration changes slowly over time (it is only in response to prolonged exposure to sunlight that we get a darker tan). However, blood (and therefore haemoglobin) concentration changes rapidly in response to the muscle activation and compression/decompression of our skin. This project studies the changes in blood flow and its dynamic on skin colour.

To simplify the data capture process, the focus is on mechanical changes in facial expression affect the underlying blood concentration. While it is common knowledge that emotions change the blood flow in humans – blushing when we are embarrassed, blood draining away when we are afraid, it is difficult to quantify these properties experimentally. However, as an initial starting point, we are discovering a link between the changes in facial expression, and underlying blood flow changes on a purely mechanical level. This simplifies our model, as any emotional changes in blood flow are ignored.

To study these correlations, a novel device has been built to capture both changes in facial geometry and changes in blood flow. Changes in facial geometry are captured using binocular stereo to capture moving facial meshes. While changes in blood flow are measured using the video and the inversion mapping described above. We then use state-of-the-art machine-learning techniques, to study this relationship.

Virtual clothing

Dr Bernhard Spanlang & Body Metrics

Supervised by Prof Bernard Buxton, UCL Computer Science

This project involved the design and evaluation of virtual clothing, with evaluation by computer models simulating the photo-realistic appearance of real garments. With the aim of developing technology to seamlessly and automatically implant images of virtual clothing into photographs of real environments, so called mass-spring particle systems were harnessed for the simulation of cloth drape behaviour.

To mimic the drape behaviour of real fabric physical properties, such as resistance to stretching, shearing and bending were mapped to the computer model. Simulations of light interactions with virtual fabrics are carried out using global illumination methods. A good approximation to the complex light reflectance and transmittance behaviour of real fabric's microstructure is important for realism. In order to measure the accuracy and quality of the simulation image metrics, some of which take the human visual system into consideration, the difference between photographs of real garments, and simulated garments blended into real scenes were investigated.

Successfully completing his EngD in 2005, Bernhard was awarded a scholarship from the Centre for Scientific Enterprises Limited at London Business School, in order to evaluate the commercial potential of his research.



Realistic modelling of clothing is achieved through analysis of the drape's structure

Layering body parts

Simon Pilgrim & Electronic Arts

Supervised by Prof Anthony Steed, UCL Computer Science

This research focused on the animation of 3D characters for video. While previous techniques for authoring and playback on games hardware suffer from a range of computation limitations – from storage to rendering – making it hard to maximise the visual quality of the scene, this work presents a level of detail mechanism based upon the animation of a character and demonstrates the benefits compared to previous (geometric) methods.

The research area concerns the development of more realistic methods of character animation using layered models, investigating the possibility of a 5-layer model suitable for both straightforward editing by artists, and then conversion into a real-time responsive 3D model for use in a game or animation.

The work contributes very strongly in progressive skinning, a system which allows computational time allocated to

animate each character, to be easily controlled based on the complexity of the scene. Progressive skinning is achieved by reinterpreting existing skeletal animation systems, so that individual 'bone' matrices of a character can be progressively added, or removed, in real time. This avoids the computational costs of preparing and animating bones, which have a limited contribution to the animation of the character. Further skinning systems are also developed that focus on anatomical, clothing and facial animations, and are compared to other methods, including blend shapes, in terms of resource cost and hardware scalability.

Simon now works as a senior programmer and research and development engineer with Sony Computer Entertainment Europe, responsible for high-performance character animation tools and runtime systems for gaming platforms such as Sony Playstation.



Realistic body modelling is achieved by using layers

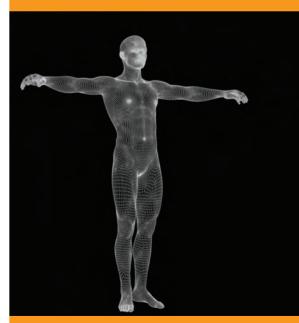
Animating figures

Dr Anthony Ruto & London College of Fashion with BBC R&D

Supervised by Prof Bernard Buxton, UCL Computer Science

This project developed static (limited pose) and dynamic (varying pose) human body models and registration algorithms. These made it possible to animate static human body scans captured using current body scanning technology. The registration algorithms enabled the static and dynamic models to match accurately the shape and pose of captured body scans and allowed the models to simulate the shape and pose of the captured scans.

In fashion, the dynamic models may make it possible to assess garment designs "in action" as new garment designs could be draped on scanned individuals being animated. In entertainment, the dynamic models can make it easier and cheaper to use captured body scans in virtual production for TV, films and documentaries, as they reduce the amount of work needed to animate the captured body. Further advanced techniques enable the dynamic animation of static human body scans for both fashion and entertainment purposes through the use of dynamic statistical human body models. Static human body models will be built from a large set of body scans (some of which will be provided by the SizeUK body database) and techniques for robustly fitting these static scans to newly captured body scans were evaluated. These techniques provided a means of registering captured body scans with dynamic human body models. The static body models were then updated with dynamic properties for both motion and deformation during motion in order to create the dynamic human body models. Existing work on virtual clothing and character animation was then incorporated into the work to provide applications for both the fashion and entertainment industries.



Human body scans were animated using a library of dynamic movements

Body scans

Dr Avy & Body Metrics

Supervised by Prof Bernard Buxton, UCL Computer Science

Point Distribution Models (PDM's) of the human body, created from a set of body surface scans (point clouds) can be used to represent variations in 3D body shape through a small number of shape parameters called Modes of Variation. This project addressed the difficulty of creating accurate PDM's by representing an equal number of anatomically corresponding points.

Reconstructing faces

Dr Howard Towner & BT

Supervised by Prof Mel Slater, UCL Computer Science

This research examined occluded facial expressions to investigate how missing or occluded data can be accommodated by other sources of information and statistical techniques. One strand of this work used a statistical model of a facial expression, based on the positions of certain feature points, such as locations on the lips, eyes, or brows. This model, combined with the positions of a few feature points, can be used to estimate the complete facial expression. The other strand of work investigated the extent to which facial expression is related to other physiological modalities, such as breathing and heart.



Missing facial data can be reconstructed through inferences

Articulated joints

Group project with Cristina Amati, Jacques Cali, Dan Calian and Rebecca Kleinberger

Hobbyists and enthusiasts now have the ability to produce their own physical models and items. However, these models are often printed as static figurines. If an articulated model were to be created, it would often involve the input of skilled artists in the design process. In this paper, we propose a method for converting 3D models of characters of arbitrary form to one-pass printable models with functional articulations, requiring no assembly, while still allowing friction within the joints to let the models maintain their poses. Our process requires only the mesh to be printed, and a simple rigging that describes the desired articulations. Our results show the effectiveness of our method on input models of varying geometrical complexity, while also allowing the user to apply some rotational constraints.





Creating content

Gwyneth Bradbury & Disney Interactive

Supervised by Dr Tim Weyrich, UCL Computer Science

In the drive for more realistic virtual worlds and richer game levels, temporal investment is becoming a huge bottleneck in the process of content creation. The creation of a world relies on a long and complex pipeline which requires artists and designers to carry out increasingly time-consuming and detailed tasks. Such investment, both financially, and in terms of committed work hours, is not always possible. So it is important for tools to be available which reduce the artists' workload, facilitating a faster work-flow and creative pipeline, as well as good control over the final visual result.

Current solutions can be categorised as follows: either they require a significant amount of user-interaction, which is both detailed and tedious. They are procedural or automated, giving the user too little control, and the output cannot be easily edited. Or they are based on real-world data or simulations which are also time-consuming. Data-driven approaches to virtual scene content creation are, however, becoming more widespread. Set production for the recent production Prometheus saw real satellite imagery used to construct

a preliminary scene to which geographical changes could be made and additional scenery inserted.

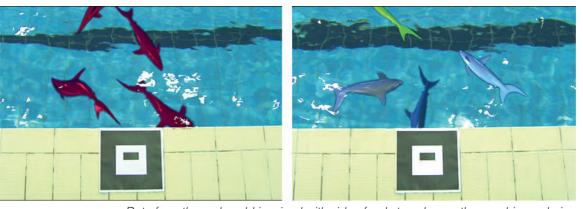
This project investigates the creation of a semi-automated, data-driven, interactive content-creation tool which is intuitive to use, while allowing control over complex parameters. The proposed system, would enable artists to make use of abstract landscape representations which incorporate real world data, allowing them to create virtual landscapes quickly, and with an appropriate amount of control.

The project looks at a range of approaches to address the problems involved in partially automated reconstruction of virtual environments, with particular focus on landscapes. Multi-spectral stereo image capture is investigated, with the aim of extracting useful information about a natural scene from a small database of images. Aerial footage is investigated in order to create natural-looking environments, which incorporate the statistical distributions of different types of vegetation, which can be seen in real-world landscapes.





Recreating a landscape for a game environments using an aerial image



Data from the real world is mixed with video feeds to enhance the graphic rendering

Mixing data streams

Dan Andrei Calian & Disney Interactive

Supervised by Prof Jan Kautz, UCL Computer Science

Any scenario in which virtual content is integrated with a video feed from the real world is defined as 'augmented reality'. This virtual content can be produced, for example, by the graphical output of a rendering algorithm, and can be integrated through the use of 3D registration technologies. With the continuing advancement of mobile hardware such as smartphones and tablets, augmented reality is becoming widespread. However, the general quality of graphics in augmented-reality applications on mobile devices remains low, in comparison to a desktop implementation, with the quality gap being attributed to the difference between overall performances.

The goal of this project is to use scenario-specific, contextual information to improve the fidelity and believability of virtual augmentations. The premise being that shifting most of the computation of rendering algorithms (traditionally performed by a computer) to the environment, will enable the creation of faster algorithms. An example of this approach, essentially a pre-computation phase, would be the capture of light probes. Light probes form the basis of many popular rendering approaches, or the capture of other features of the environment, for the purpose of accelerating graphics rendering.

With relation to the gaming and toy manufacturing industries, the importance of the success of this

research, cannot be understated. The project's output is directly applicable to the development of augmented reality games for mobile and console-class hardware. The application of this research to children's toys relies on the fact that known environment structures can be inserted within a toy's construction. Later on these can be exploited by the environment data-dependent rendering algorithms, and improving efficiency of the technology.

Personal statement

During the first year of the doctorate, which led to a Master in Research qualification I have performed initial exploratory research into the field of augmented reality and realistic rendering. My individual dissertation focused on building an underwater augmented reality rendering system, with real-time single-bounce global illumination. Another important part of this first year was a group research project on 3D printing models with articulations.

We believe that this doctorate will advance research in realistic rendering algorithms by leveraging the use of contextual information. By formalising the theory behind this reasoning, we can propose novel rendering algorithms with this theoretical basis.

Fluid simulations

Julian Hodgson & Passion Pictures 1851 Industrial Fellow

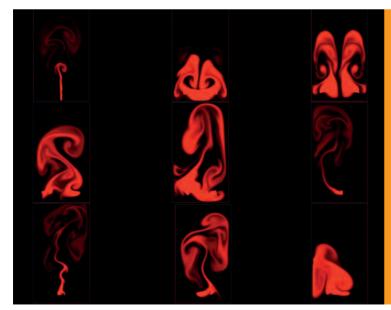
Supervised by Prof Anthony Steed UCL Computer Science

Fluid simulations are computationally (and financially) expensive and also sensitive to initial conditions. The time to produce them and the difficulty to achieve the artistic vision frustrates everyone in the creative line from the initial artist, director and producer, to the audience. Improvements have been achieved, but there are still tools to be developed to make that process more efficient and user-friendly, as well as being accessible to the less technical artist. The aim is to develop new tools for creation and manipulation of fluid simulations by collaborating between artists working in production and academic researchers.

Personal statement

My time in post-production has shown me that there is a need to create realistic fluid simulations in computer graphics. I have worked in the computer graphics and animation industry for 15 years both in software development as a programmer, and in production as a technical director and supervisor. After completing an MSc in mathematics, I went on to work in graphics hardware and software companies, and then post-production studios in London and San Francisco. I am currently working at Passion Pictures, managing software development and writing character rigging and simulation tools.

For me the EngD is about finding the link between research and its application in industry. This is not easy, as the requirements differ for each side, yet these constraints can create a unique challenge. Artists can benefit from new solutions to problems that have not been implemented in any existing commercial packages. Conversely, the workplace is an inspirational source of problems to solve that may not have yet been seen in academia.



Improving the process in animation of realistic fluid dynamics

Mapping investment

Dr Patrick Weber & ThinkLondon

Supervised by Dr Dave Chapman, UCL Civil, Environmental and Geomatic Engineering

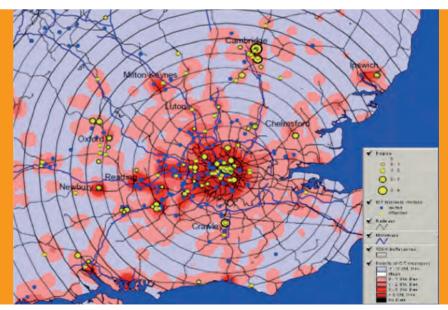
This research set out to develop a decision support system to iteratively explore, compare and rank London's business neighbourhoods. This was achieved through the development, integration and evaluation of spatial data and its manipulation to create an interactive framework to model business location decisions. The effectiveness of the resulting framework was assessed using a scenario-based user evaluation, based on the case study of foreign investors in London.

As a joint venture between ThinkLondon and UCL, the project assisted with the development of a Geographic Information System, which is used widely to engage and consult with clients. The GIS allows ThinkLondon to market itself in a more intelligent and proactive manner, helping the organisation to focus its resources more efficiently. Information delivery to inward investors has considerably improved through the use of static and

interactive maps. The GIS also gives ThinkLondon's project managers the power to produce custom maps tailored to their customers' needs.

ThinkLondon sponsored this research project with the aim of quantifying and qualifying London's diverse business neighbourhoods, which lead to the implementation of a Spatial Decision Support System for inward investors into London.

After completing his EngD, Patrick continued his research as a Research Associate with UCL, working on the connection between using GIS for day-to-day activities and as part of research or studies. This lead to an open-source, web based eLearning environment for teaching GIS skills to non-expert users.



The potential for business investment is mapped for London and the South East

Virtual neighbourhoods

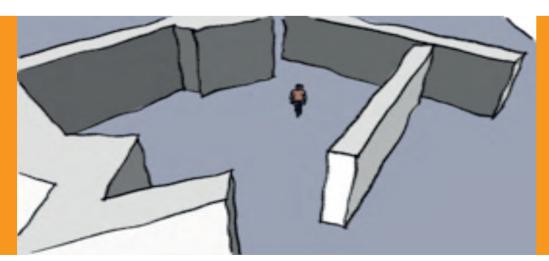
Dr Nick 'Sheep' Dalton & Ovinity

Supervised by Prof Alan Penn, UCL Bartlett School of Graduate Studies

This project began with a simple question: when we talk about 'place', your place, my place, my chair, my home, my country, my neighborhood, is place a thing or just a metaphor? Nick was interested in our attachment to place. The primary question addressed the theme of place existing with a logical syntax of space as used by specialist architects. Can we become as attached to our virtual game worlds as we are to reality? Many theorists from Aristotle onwards have separated place from space. Does this mean we can have 'virtual place' which is as emotionally strong as a real place?

Personal statement

I wanted to explore what place was in the virtual realm and to understand it I decided to look at place in the real world. To do this I wanted to take tools from architecture with the notion of recycling them for virtual worlds. At that point I discovered the theories of real places were under developed and untested. I began by developing new software to compute large-scale exotic graph theory called 'Intelligibly' and 'Synergy' in the process of visualising them I noticed that areas of the same colour (constant synergy) looked to be like the areas of familiar neighborhoods. Yet finding suggestive coloured patches was not enough. I began developing new methods to capture people's subjective notions of where their neighborhood was. For years, I walked over many square miles of London collecting data. After creating new methods which allowed me to synthesise



A virtual lab helps find the spatial factors by which people identify a neighbourhood

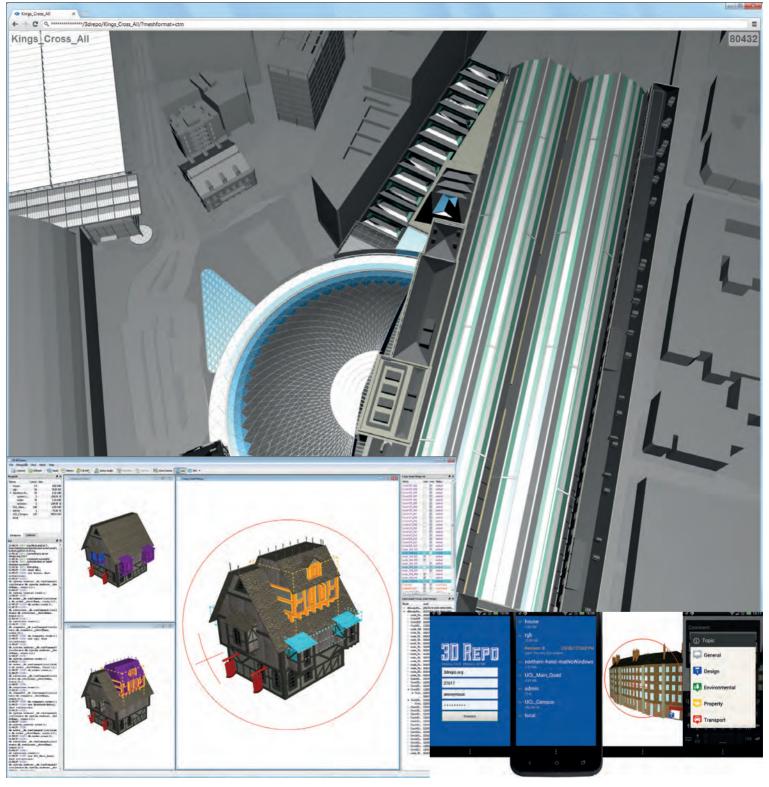


people's subjective impressions of where their place was, I was able to develop new statistical methods to test if my computation theory of neighborhood was related to the real world feelings about place. At this time found I was in a race to test my theory before others did, I processed a lot of data I found that my results were positive and were stronger than any other current theory.

I came a long way from my roots looking at games but I discovered something more significant, that the seeds of our neighborhoods hide beneath our feet and in the walls, both in our space and our hearts so to speak. I saw some wonderful things I never expected along the route, I even discovered a London neighborhood being born from nothingness. I discovered a lot about research and myself through my EngD. That you should be open to discovery whenever you find it, no matter that I had a reading age of an eight year-old child - my dyslexia was worse than I had suspected. Despite this I was able to complete my thesis.

The notion of designing for neighborhoods is a strong recurring theme in architecture and urban planning. My work has contributed to those who need evidence that design can have an impact on people's lives. As more cities are built around the world my theories and tools could be used by those who want urban design to connect people not isolate them.

I am now a lecturer in Computing at the Open University researching ubiquitous computing something I don't think I could have, as a dyslexic, dreamt of doing before.



Developing interactive tools for the management of 3D assets for major engineering projects

Managing 3D assets

Jozef Doboš & Arup Foresight

Supervised by Prof Anthony Steed, UCL Computer Science

The maintenance of assets in a large 3D visualisation can involve many authors with potentially different skills and tools. This presents problems, not limited to maintaining consistency of the models and dealing with concurrent edits in the same part of a 3D scene. To address these issues, a unified and integrated framework, 3D Repo* has been developed that supports collaborative editing and distribution of 3D assets. This framework tracks multiple revisions of assets so that they can be integrated at a later date. It further supports distributed editing over the Internet and additional lightweight clients, in webbrowsers and mobile devices.

With the recent progress in the capabilities of graphics processing units, there is a desire for even larger and more detailed models to be used in engineering projects. The standard paradigm of collaborative editing is the sharing of files between various instances of applications. Each user loads a particular 3D scene into a modelling tool, modifies it, and then re-saves the whole file again, making version tracking unnecessarily obfuscated. By the time any simulation based on the model is completed, the design might progress so rapidly that the results are often inapplicable. What's more, when developing land or transportation systems, it is a legal requirement to host public consultations. Such an event can incur significant costs, and, for very large projects, it risks not reaching the key stakeholders.

Spatial databases originate from the Geographic Information Systems domain. They store features such as points, lines, areas and their associated metadata. A few commercial and open-source databases, support spatial data directly, e.g. Oracle Spatial and PostGIS for PostgreSQL. However, they focus on 2D geometry, not 3D models and their properties. They also do not support the type of revision history that is commonly required in 3D asset management. In contrast, NoSQL databases avoid rigid table structures, and tend to be optimised for

large read-write operations. It is thus possible to store models in a database and, due to its flexibility, also track semantic relationships, and even individual revisions as done in 3D Repo.

This project's framework uses the open-source MongoDB as its 3D repository where it assigns two collections (tables) per 3D scene. One for all the scene graph constituents and one for all the documents that belong to a revision history. Several database front-ends were developed, most notably a C++ read-write application with a 3D differencing interface so that only delta changes of any revision are stored in the database.

3D Repo automatically serves models from a centralised repository, and records annotations alongside the 3D assets, and their individual revisions. In order to offer a scalable visualisation platform, it streams decomposed scene graph components onto client devices, reconstructs their 3D representation, and displays them for viewing. The interaction between modelling software, and the repository takes place through the import and export of files. This successfully de-couples the modelling from its long-term storage as files are now considered to be only a temporary representation of 3D assets. Such an approach is of great importance, as the range of editing tools used across the industry is vast.

Given the direct database connection, the scalability of MongoDB and its in-built support for geospatial indexing, this approach offers significant benefits over alternative techniques such as geographic information systems (GIS). This NoSQL-based solution can be extended to support various data types, including the Building Information Modelling (BIM) and other representations.



Artists can create a virtual building using digital capture techniques

Automatic buildings

Fabian Wanner & Framestore

Supervised by Prof Jan Kautz UCL Computer Science

The quality of the graphics and effects used in films and games is continuously getting higher, as a result of increases in detail and accuracy of the models used. These advances in quality not only improve the visual satisfaction of the audience and users, but also lead to increases in production time and cost. Nowadays, large numbers of 3D artists are needed to create the visual content to the necessary quality standards. In recent years, many techniques have been developed to simplify this process for the artists by providing them with more support. Different tools, algorithms and techniques are supposed to make the work of the artists easier by simplifying individual tasks, by increasing quality or by automating tedious elements of the reconstruction process.

One way to increase model quality and realism is to use actual objects as a base. Doing so results in simplifications in several areas. The models automatically make physical sense, textures can be obtained relatively easily and depending on the reconstruction system used a lot of details can be gathered quickly. Most existing reconstruction systems use images, a point cloud, camera information or a combination of these to allow artists to quickly and accurately obtain a model of the scene of interest. This project tries to add other sources of information to make the task of achieving a model either simpler or to improve accuracy. In particular, edges, repetitions and shadows in the scene have been used. Each of them has a different purpose to help the reconstruction artists.

We can use repetitions to reduce the amount of interaction necessary by the user. This can be done by recognising areas in images that are the same or very similar and apply changes applied to one of them to all of the similar parts of an object.

The detection of edges in the scene allows the correction of user input. In most cases when the user makes an interaction, for example to outline a plane, it will be near a textural or structural edge in the scene. This information can be used to reduce the accuracy of interaction required by the user.

Shadows are a very useful tool for testing the correctness of a model and, if necessary, fixing it. This can be done by calculating the shadow location from the scene into the image and comparing it to the shadow location in the image itself. In an ideal situation, these should be equivalent, but this is rarely the case due to the many sources of error. Then the model can be manipulated in such a way to satisfy shadows, point clouds and images as closely as possible using an optimisation procedure. The combination of these improvements enables artists to increase reconstruction accuracy while decreasing reconstruction time.

Urban design feedback

Greig Paterson & AEDAS

Supervised by Dr Dejan Mumovic, UCL Bartlett School of Graduate Studies

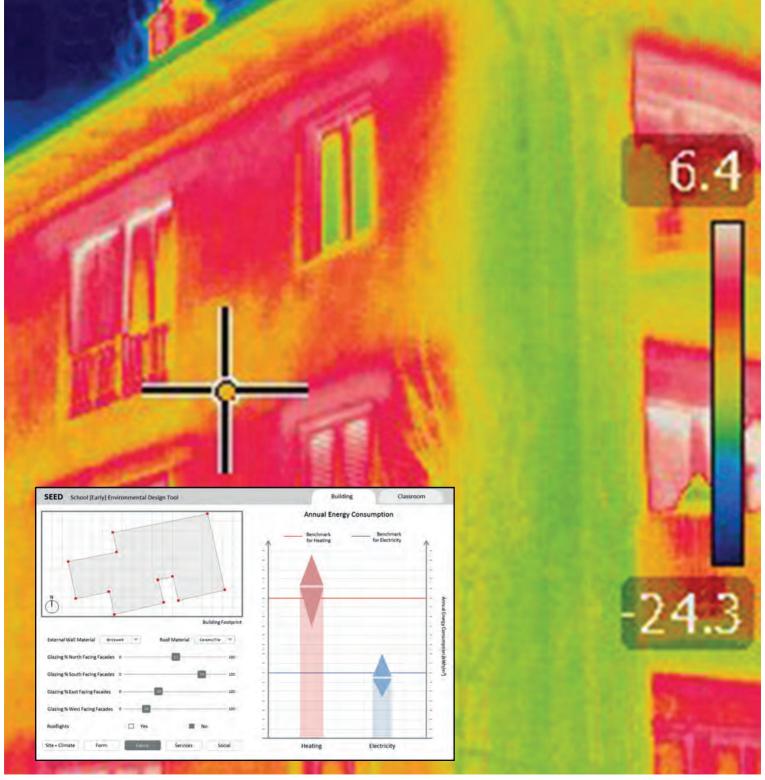
Many design aids are currently available to architects, helping them make environmental design decisions throughout the design process. Building simulation, whereby a virtual building is simulated within a virtual climate, is one of the most flexible design aids available today. However, due to its complexity and time needed to execute, building simulation is often rejected by architects at the early design stages – a time at which some of the major decisions affecting the sustainability of a project are made. Furthermore, studies have shown that the predicted energy performance of buildings is often far lower, than the actual performance of buildings once built.

The aim of this project is to develop a user-friendly design tool, which allows environmental performance indicators to be communicated to the user in real-time, while early design parameters are altered interactively. This would allow the architect to sketch performance as well as form. The development of this tool is not meant to replace traditional building simulation – instead, it aims to act as a quick plausibility check for non-technical designers, as they exploit the design space during the early stages of design.

Personal statement

My research explores machine-learning techniques, utilising monitored energy data, to form a platform for performance prediction. Using a selection of schools in England as test cases, my research has involved training an artificial neural network to predict the energy consumption of new schools by linking actual heating and electrical energy consumption data from the existing building stock, to a range of design and building use parameters.

The VEIV EngD programme is a fantastic way to carry out this research facilitating important links between academia and industry. Working closely within the R&D team of the international architecture firm, Aedas, I am able to gain feedback as to the important issues facing the design community locally and globally. One of the main objectives of the EngD scheme is to ensure a tangible 'product' for industry an objective of the VEIV Centre is the communication of complex data. The marriage of these two objectives works particularly well within an architectural design studio as the outcomes can be directly integrated into the design process, helping to create the next generation of Aedas's buildings.



Novel interfaces for design tools enhance the building simulation process

BUILDING LIFETIME

Visualising the carbon impacts over a building's lifetime

Visualising energy performance

David Hawkins & Feilden Clegg Bradley

Supervised by Dr Dejan Mumovic, UCL Bartlett School of Graduate Studies

The lifecycle carbon impact of buildings includes both the carbon impact associated with their operation and the embodied carbon impact related to their constituent materials. This developing area of interest in building design is particularly important during the refurbishment or replacement of an existing building, where expected reductions in operational carbon impact may be offset by the addition of new materials.

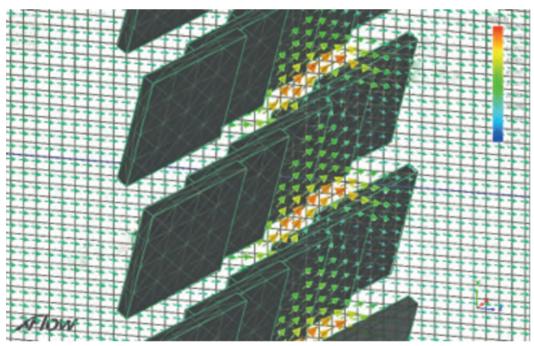
This project centres on lifecycle carbon considerations, specifically in the re-development of higher education buildings. This includes the following key research areas: the extent to which building re-development reduces operational carbon impact; how embodied carbon impacts vary between different design scenarios, particularly when allowing for future lifecycle impacts. The significance of lifecycle carbon impact as a factor in re-development decisions. The research method includes development of a database of higher education buildings to analyse key energy-use determinants. Studies of current and projected lifecycle carbon impacts for case study buildings, and a survey of industry professionals to understand the significance of carbon impact in building redevelopment decision-making.

A target outcome is to consolidate research findings in a demonstration visualisation tool, that allows the impact of carbon management decisions across a whole estate to be simulated. The tool will also serve an educational function to highlight how lifecycle carbon impacts vary over the life of a building. The tool makes use of 3D mapping data, combined with a database of carbon performance characteristics for building archetypes based on the research. An interface will be provided to allow the user to adjust characteristics for each building. The tool will then display lifecycle carbon impacts developing over time.

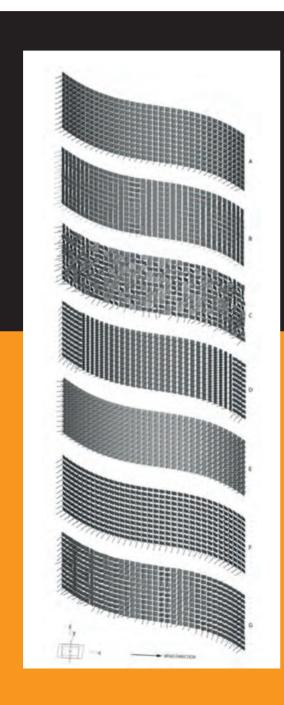
Wind flow

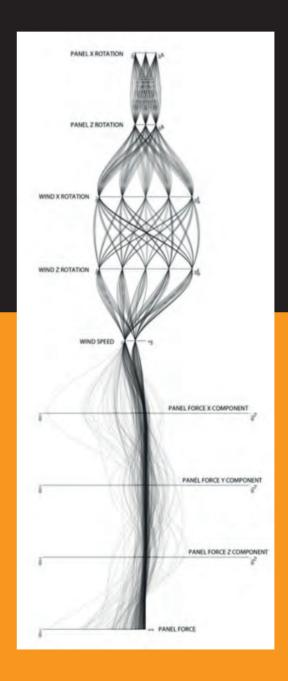
Group project with David Andreen, Gwyneth Bradbury and Samuel Wilkinson

The use of Computational Fluid Dynamics (CFD) for wind simulation and subsequent façade component optimisation can be a computationally inefficient process, requiring lengthy simulations of many iterations of the entire building under various wind conditions. This project proposed a method for learning the behaviour of individual components at a small scale to then apply to large-scale structures. Bentley Generative Components was used to model the component parametrically, NextLimit XFlow was used to simulate the wind environment and Matlab was used to correlate the training set and implement a k-nearest neighbour search for the application test. A number of application tests were conducted in order to explore the inherent issues and possibilities of using such an approach. One of the potential benefits of this method is reduced batch-simulation time, for a later generative design stage detailed optimisation. The possibility of learning local component behaviour and mapping this to larger structures was shown to have more effective success for certain applications over others, and demonstrated a strong need for further research.



Modelling the flow of air as wind helps to shape architectural forms





Aero-dynamic design

Samuel Wilkinson & PLP architecture/ Bentley

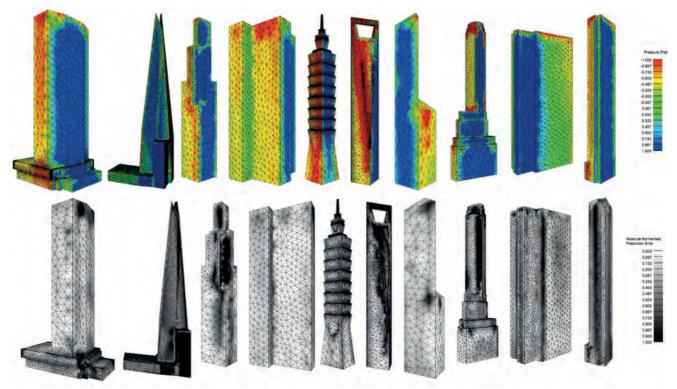
Supervised by Dr Sean Hanna UCL Bartlett School of Graduate Studies

The trend towards creating ever taller buildings continues with dramatic technological improvements in materials, structures, and modelling. However, as height increases, so do gravitational, wind, and seismic forces. Wind forces on tall buildings must be mitigated to avoid occupant discomfort from swaying, reduce risk of façade damage, and improve structure efficiency.

Although these problems can generally be avoided through early decision form-finding. Current fluid simulations are costly, and do not provide the rapid performance feedback necessary for such explorative or optimisation studies. This project develops an approach that uses machine learning to make predictions of the wind. This occurs by it recognising shape features

through the use of procedural geometry generation and regression analysis. The end result will be a prototypical tool, or a method whereby wind loads can be immediately predicted, and visualised on models of arbitrary complexity.

The research project is co-sponsored by Bentley Systems (the developer of Generative Components, a parametric CAD software for the construction industry) and PLP Architecture, a young architectural practice working on a wide range of international projects. The opportunity to develop software in conjunction with end users presents many benefits to both companies, especially when they share the common goal of integrating computational and environmental design methods in practice.

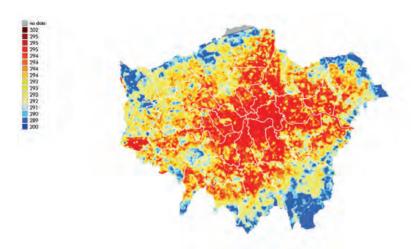


Machine learning helps predict the surface wind pressures on architectural forms

Mitigating climate change

Gurdane Virk & The Chartered Institute of Building Service Engineers

Supervised by Prof Mike Davies, UCL Bartlett School of Graduate Studies



Modelling urban land use to help address climate change

This project evaluates the way in which urban climates affect the built environment and how changes to land use and building design affect it in turn. Due to increasing urbanisation and projected climate change, the effects of urban warming will be exacerbated. The relationship between developments and urban climates, therefore, becomes increasingly relevant. The physical properties of urban environments contribute to urban warming depending on the spatial and temporal scale at which they are observed. The spatial scale of these effects range from the building to street, on both neighbourhood and city scales.

Urban climates are modified by the physical properties of cities, leading to the creation of Urban Heat Islands (UHI), as well as creating microclimates within them. Interventions and mitigation strategies can abate urban warming, but their impact depends on their spatial scale. Consequently, there are various strategies to mitigate the negative effects of a warming climate. These allow cities to be planned, so that the environment surrounding a building reduces the impact of the external microclimate. Alternatively, buildings can be designed to perform better, under future climate scenarios.

Building regulations and standards are progressively more performance-related – regarding the assessment of overheating risk within a building, as well as its sustainability credentials within the context of the surroundings. If policy makers are to master-plan cities, they need to understand the effect of possible abatement measures at different scales. Similarly, designers, planners and architects need decision support tools to identify and prioritise any adaptation strategies. And to be able to effectively comply with stricter regulations, through sustainable design.

The goal of is to further develop current models, incorporating a larger range of spatial and temporal scales. The project also aims to produce design guidance and knowledge resources in the form of modelling tools for members of CIBSE, encouraging sustainable design in urban environments. There is currently limited guidance available to the industry on the effect of building design, within the context of the urban climate. Previous CIBSE guides have provided a useful treatment of UHI effects, but lack a wide range of microclimatic effects. Ultimately, this project will contribute to new Technical Memorandums and Application Manual guides on Urban Climates.

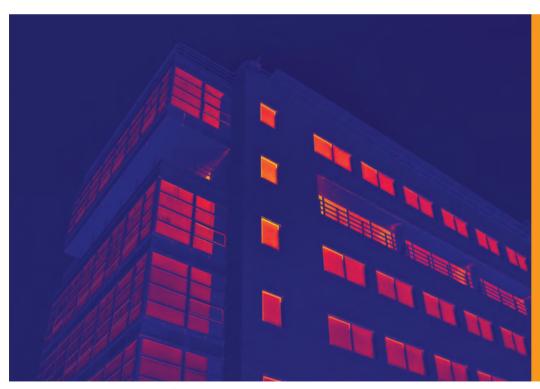
Fuel monitoring

Ciro Bevilacqua & BRE

Supervised by Dr Ben Croxford UCL Bartlett School of Graduate Studies

The depletion of fuel reserves, and the concerns on the energy security that stem from it, have increased the focus on the efficiency of the generation and delivery of the energy to the communities. In this scenario, renewable energy technologies and district heating are considered to be a very efficient form of generation and distribution of energy, offering the potential for considerable savings of carbon.

The aim of the project is to produce a tool that is to assist in the appraisal of the community energy systems at concept stage on the basis of environmental and energy metrics. The tool would have the capability to capture the implications of the integration of different energy sources and assist in the engineering optimisation of supply and demand through thermal storages, low temperature heat supply and improved building performance.



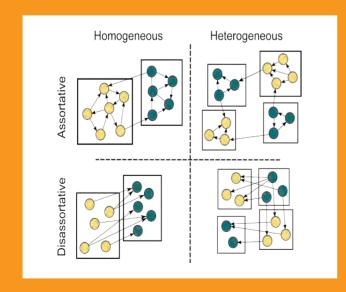
Fuel monitoring with visual tools helps indicate energy consumption

Social arrangements

Leto Peel & BAE Systems

Supervised by Dr Simon Julier UCL Computer Science

Interaction networks are made up of nodes, representing people, and links representing observed interactions between them. Within such networks community structures have been observed where nodes are more densely connected within communities than across communities. People tend to interact more frequently with similar people and so people within communities will share common attributes or purposes. This research analyses the interactions of network communities. In particular I am interested in investigating the detection of communities, discovering roles within communities and community evolution.







Visual tools help understand the cost of construction

Understanding cost

Dr Sarah Graham & Atkins

Supervised by Prof Alan Penn, UCL Bartlett School of Graduate Studies

The outcome of this research was a rapid, efficient means of producing and interacting with the information required for strategic design at the earliest stages of projects, enabling Whole Life Cost issues to be investigated, conveyed to and understood by all relevant parties. By combining emerging technologies in human-computer interface design with best practice in environmental design and Whole Life Costing. Within the context of real building projects the research developed a pragmatic and practice-driven approach to overcome a major obstacle to achieving sustainability in construction.

Advanced design

Francis Aish & Foster + Partners

Supervised by Prof Alan Penn, UCL Bartlett School of Graduate Studies

This study comprises the development of integrated software and methodologies that tighten the design optimisation loop. The research has three main foci. Firstly, the development and integration of advanced parametric modelling tools. These will enable designers to capture and manipulate design intent, and then programmatically generate dependant representations. Secondly, the development and integration of analysis tools. These will assist the qualitative assessment of building performance and allow the semi-automatic optimisation of the parametric model. Finally, the aim is to investigate enhanced user interfaces as an enabler for first two foci.

Design geometries

Abel Maciel & 3DReid

Supervised by Prof Alan Penn, UCL Bartlett School of Graduate Studies

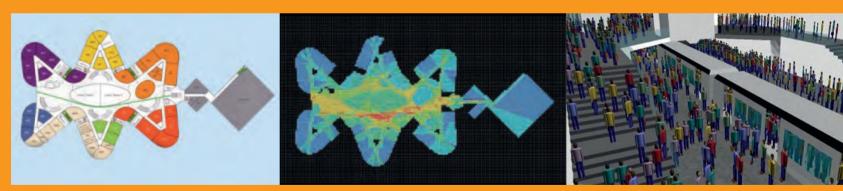
This project explores the possibilties for the graphical modelling of design as an integration strategy for the management of building construction projects. It is being developed in two stages, the first of which is to investigate geometrical modelling as the pivot for determinants of architectural design, such as sustainability and planning principles.

Agent simulations

Erica Calogero & Buro Happold

Supervised by Prof Alan Penn, UCL Bartlett School of Graduate Studies

This project uses agent-based modelling techniques to design and implement a high-resolution simulation of people movement in large urban areas for the purpose of risk reduction and mitigation in the event of system shocks. The research analyses both existing population and statistical data along with high resolution spatial network data that interfaces with existing GIS technologies in order to enable a baseline model of people movement that can dynamically update and respond to changes in the transport network. The results from the research are aimed to help emergency managers carry out risk assessments before, during and after such shocks take place. Bayesian belief networks will be used to model the choices that individuals make in the face of risk.



The school building modelled in Yamamoto sematic map modellig software showing a sample shortest route from the entrance on the left of the building to the foot of the dining room stairs on the opposite side of the building

Agents model pedestrian flow to help shape architectural designs



Rail track monitoring

Anita Soni & ThamesLink

Supervised by Prof Stuart Robson, UCL Civil, Environmental and Geomatic Engineering

The protection of existing buildings and activities, as well as major transport infrastructure hubs (e.g. Kings Cross/St Pancras, Blackfriars, Farringdon and London Bridge), is essential to making such developments sustainable as transport solutions. Monitoring is a fundamental requirement for safety, protection and operational efficiency. Current railway monitoring activities involve placing targets by drilling, clamping or gluing directly onto the structures which can be intrusive, create safety issues, and be time consuming and expensive.

Advances in survey techniques, particularly with laser scanning and photogrammetry, offer the possibility to remotely capture vast amounts of data quickly without the need of direct surface contact. This could significantly reduce the costs and time expenses, as well as safety issues currently involved. In parallel new visualisation and

mixed-reality computing technologies offer the potential for rapid on-site confirmation of changes rather than the numeric feedback of current techniques.

This project looks at key issues for the application of new spatial measurement techniques including the accuracy and precision achievable, handling the data volumes generated and the ability to rapidly analyse and communicate results. This will be done by investigating a selection of current railway monitoring situations and to segment their existing instrumentation with new technologies. Careful design using numerical simulation in combination with real data will allow quantitative and qualitative comparisons of the capabilities of different techniques and lead to the development of best practice for their adoption, including in cost/benefit analysis.

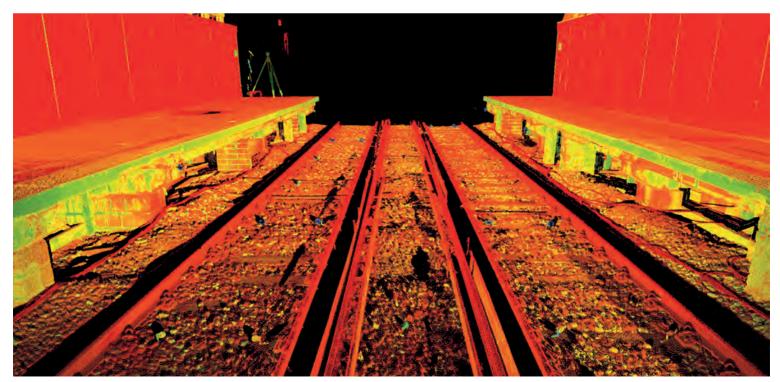


Figure 1

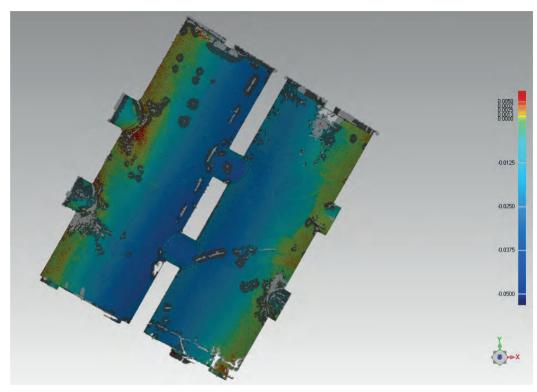


Figure 2

Figure 1 Monitoring railway track. Target-based monitoring provides a discrete set of measurements to detect movement around the track – these targets can be seen in the point cloud as blue patches near the track and platform. Common problems with the targets include a restricted line of sight when taking a measurement or the target itself getting knocked while engineering works take place. These issues highlight the potential of a laser-scanning solution to railway track monitoring. The point cloud generated can provide continuous geometrical information of the track and platform areas without the use of targets.

Figure 2 Structural monitoring of masonry arches. A displacement map, in plan view, showing the level of movement detected (in metres) while construction work takes place nearby. This has been captured by laser scanning the area over a period of time without placing any targets onto the structure. These results validate the target-based monitoring and provides the engineers with an enhanced 3D visual representation of the movement to carry out further analysis.

Monitoring moving organs

James Martin & VisionRT 1851 Industrial Fellow

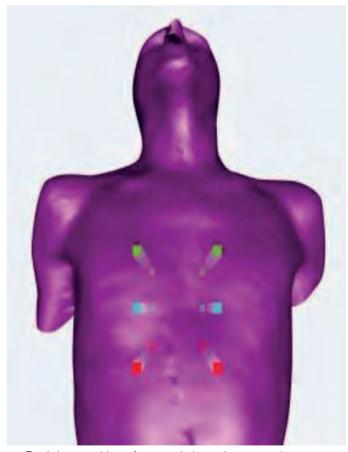
Supervised by Prof David Hawkes, UCL Centre for Medical Imaging

This research attempts to advance a cutting-edge treatment for lung cancer by incorporating the effects of respiratory motion into radiotherapy. A recent development in stereotactic ablative radiotherapy (SABR) has been the use of cone-beam CT (CBCT) scans to accurately position the tumour immediately before each session of treatment. The project's novel approach uses this scan to evaluate and model the motions of key sites involved in treatment. The sponsor, VisionRT, produces a proprietary camera system to visualise the patient in real time. This technology can then be used to monitor tumour and other motion during treatment, and deliver the dose more effectively.

Personal statement

Lung cancer is the most common cancer in the world, with less than 10% of those diagnosed surviving a further five years. By improving the effectiveness of current radiotherapy treatments, my aim is to make a real impact on patients' lives for the better.

In recognition of the potential impact this work could have, I have been awarded an Industrial Fellowship from the 1851 Royal Commission. This prestigious award has allowed me to accelerate the pace of my work, and also explore the potential of our research, to be incorporated into future medical device products.



Real-time tracking of organs is based on cone-beam scans

Document imaging

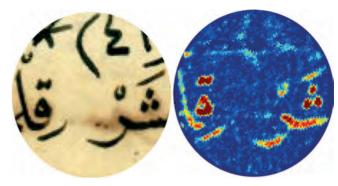
Tiphaine Bardon & TerraView in collaboration with Dutch National Archives

Supervised by Dr Matija Strlič, UCL Centre for Sustainable Heritage

This cross-disciplinary project explores the potential of terahertz time-domain spectroscopic imaging in the field of cultural heritage with a particular interest in archival documents that are difficult to read or access – such as unopened letters, scrolls or brittle documents.

Terahertz (THz) is highly penetrative, non-ionizing radiation between the infra-red and microwave regions (40 GHz – 4 THz) of electromagnetic radiation. It interacts with matter by changing the amplitude of inter-molecular vibrations and the rotations of molecules within a material. Terahertz spectroscopy can therefore provide information on the molecular structure of a sample. Combined with a scanning system, it can also produce depth-resolved images of an object based on its differences in composition, going from one layer to the other, up to a few centimetres in depth with a resolution of ~80 μm .

As the condition or shape of some archival documents may not allow for access to information without a significant risk of damage (due to degradation or structure of a manuscript), Terahertz techniques may enhance the readability of such inscriptions in a non-invasive way. Besides, new chemical information from historic documents obtained with terahertz technology can lead to a better understanding of their degradation chemistry. Therefore, THz-pulsed spectroscopy and imaging, has the potential to significantly improve the current practice of structural and chemical monitoring of archival heritage.



Ancient and fragile document are preserved using micro-level analysis

The project involves a collaboration with UCL Centre for Sustainable Heritage, in partnership with the company Teraview, and supported by the Dutch National Archives at The Hague. Collaboration with the latter has proven particularly valuable, providing crucial input and guidance, with respect to stakeholder needs and interests from the curatorial point of view.

Personal statement

My EngD at the VEIV Centre offers real advantages for all involved. On the one hand, I benefit from the access to cutting-edge analytical set-ups and expertise of the company involved. This has enabled me to focus my research on concrete applications while working closely with stakeholders. On the other hand, the partners benefit from the developed expertise, from IT to data management, as well as from access to infrastructure and knowledge at UCL. At collaborating laboratories such as the Centre for Research and Restoration of the Museums of France, at Le Louvre in Paris.

Completing the MRes Heritage Science and developing a group project with students at the EngD VEIV Centre has enabled me to develop technical and inter-disciplinary skills. I consider myself very lucky to be involved with the EngD Centre. It allows me to grow both as an engineer and a researcher.

Document reconstruction

Kazim Pal & London Metropolitan Archives

Supervised by Dr Tim Weyrich, UCL Computer Science and Dr Melissa Terras, UCL Centre for Digital Humanities

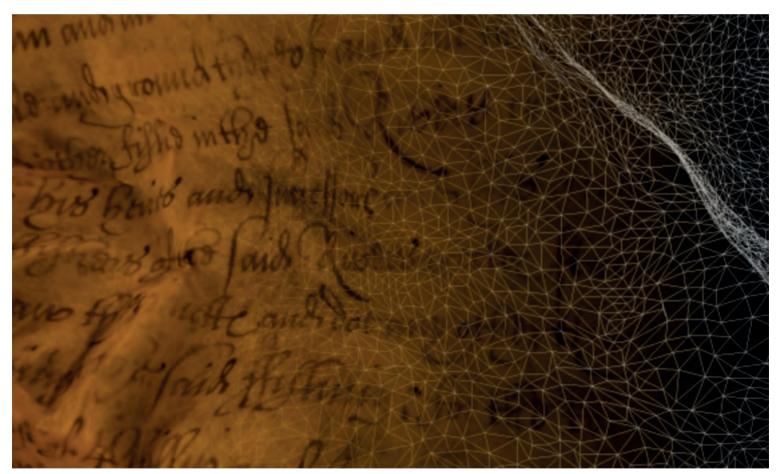
The Great Parchment Book of The Honourable The Irish Society is a 17th-century property survey of the Irish counties of Londonderry and Coleraine, was commissioned by Charles I. Its contents are of great interest to historians, studying the history of this region. The book fell victim to both fire and moisture damage, causing such strong distortions which conservators tell us cannot be restored using current methods. The study involves experiments with ways to 'virtually restore' the book employing a mix of imaging, computer vision, and computer graphic techniques. The aim is to expose the content of the pages in a legible form, and make the book accessible historians and other researchers. Creating detailed models of the pages requires a careful imaging process that



Credit: London Metropolitan Archives, City of London

gets inside every crease and fold and captures every letter at as high a resolution as possible. This results in a set of around 50 high-resolution images (for each side of each page). From these image sets can be generated high-detail and high-resolution textured 3D reconstructions of the pages, which use a pipeline of reconstruction and texturing algorithms.

Given these reconstructions, the investigator has developed a visualisation application that allows a user to browse and interactively flatten the pages. In such a way, further techniques are being developed to make contents accessible and aid the transcription process.



CREDIT: http://reclamationandrepresentation.blogspot.co.uk

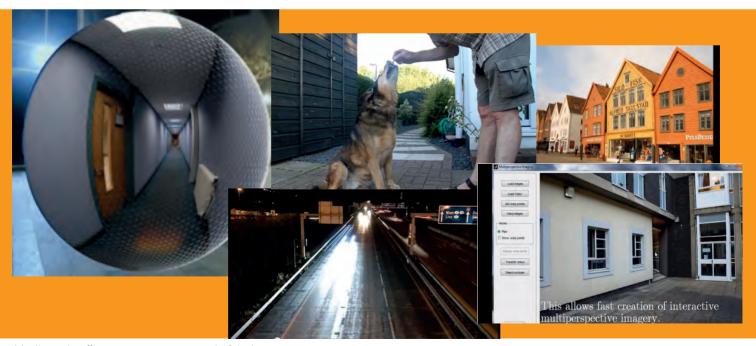
Working with content

James Tompkin, EngD & BBC R&D

Supervised by Prof Jan Kautz, UCL Computer Science

Visual computing has made great strides in helping people to capture, edit, share, and explore media collections. In capturing media, new computational cameras and software techniques have improved the quality and flexibility of photography. Now that the camera is ubiquitous, how can we make capturing beautiful images simple for novice users? How can we remove the hours of training necessary to make content edits? Perhaps most significantly, how can we make searching and exploring vast media collections simple, accurate, and semantically meaningful?

State-of-the-art research has made some progress in answering these questions for photographs, but relatively little work concerns video. This research aims to tackle some of these problems. Specifically, the investigation develops tools that exploit creative capture for new interactive media applications. The outcome of the project has been automated sophisticated video content edits, such as removing unwanted people from crowded scenes in holiday videos, and in the process removed skill barriers for novices. Finally, the project has advanced video collection, browsing and exploring techniques, to allow video collections of places and events to be navigated in an immersive and semantically meaningful way.



Creative Capture:

Towards Moment Images: Automatic Cinemagraphs

http://www.youtube.com/watch?v=-8O6C6EncLk

Visual computing advances were exploited to create new experiences from traditional and creative capture, to find content matches between different images and videos.

Interactive Multi-perspective Imagery from Photos and Videos

http://www.youtube.com/watch?v=LaQNP9e-8Bq

Artists regularly manipulate perspective in paintings, but photography commonly has only perspective and wide-angle lenses. The project introduces new creative tools for everyone to be able to play with multiple perspective imagery. This has been achieved by automatically computing the geometric relationship between different photographs of a scene, so that anyone can easily composite multi-perspective photographs.

Interactive Viewpoint Video Textures

http://www.youtube.com/watch?v=cnvVtB4I3OE

Interactive free-viewpoint object displays on websites allow museums to show ancient artefacts from any angle, and for shoppers to see the form and appearance of their particular object of desire before purchase, be it a shoe, a phone, or a new car. A method was developed to simply, and automatically, create interactive free-viewpoint object displays for objects which show stochastic or periodic motion, or any action that can be easily repeated.

Video Editing: How Not to Be Seen - Object Removal from Videos of Crowded Scenes

http://www.youtube.com/watch?v=3YJsNEzpmLc

Removing dynamic objects from videos and reconstructing them is an extremely useful and powerful tool – say, to remove the people who walk in front of your camera at the crowded plaza in your holiday videos. Also, to reconstruct what was never captured momentarily - the actions of your family and friends - and make the unwanted person disappear seamlessly. Or, for a movie production, for a period film on location, to remove anachronistic objects which would never appear at the time, when the movie is set.

Exploring Media:

Videoscapes: Exploring Sparse, Unstructured Video Collections

http://www.youtube.com/watch?v=_yE2BRTC4iE

The abundance of mobile devices and digital cameras with video capture makes it easy to obtain large collections of video clips that contain the same location, environment, or event. However, such an unstructured collection of videos is difficult to comprehend and explore. Current interfaces to these video collections are often simply lists of text-ranked videos which do not exploit the visual content relationships between videos, nor other implicit relationships such as spatial or geographical relationships. The project sought to deploy a system that analyses collections of unstructured but related video data to automatically find semantically meaningful links between videos. This allows interactive exploration of video collections by visually navigating – spatially and/or temporally – between different clips.

Shots taken at different times of day are analysed and matched for mood

Analysing mood

Yotam Doron & The Foundry

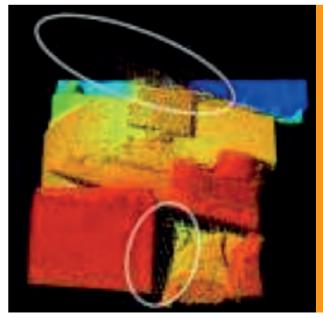
Supervised by Prof Jan Kautz, UCL Computer Science

This project aims to develop efficient techniques for automatically editing and analysing digital film. The main areas of interest are colour grading, disparity mapping and focus detection. In colour grading, the colours, brightness and contrast properties of a sequence are altered both for technical reasons, for example to match sequences shot at different times and in different lighting conditions, and in order to achieve a visual style and mood. Disparity maps are central to stereo vision, and in film production allow for combining 3D graphics with stereo footage.

Time-of-flight confidence

Group project with Malcolm Reynolds, Jozef Doboš and Leto Peel

Time-of-Flight (ToF) cameras provide high framerate depth measurements, within a limited range. However these readings can be extremely noisy, and display unusual phenomena; certain types of scenes, such as edges or materials with low reflectivities. The magnitude of these errors is much larger than in similar technologies such as laser scanners. The CamCube from PMD provides an amplitude value for each pixel (and therefore each 3D point), which can be used as a measure of confidence. This project demonstrated some deficiencies of using this as a measure of pixel reliability, and propose an improved per-pixel confidence assignment based on a Random Forest classifier. The results of this pipeline can be used in a variety of processing pipelines. Throughout several experiments with a PMD [vision] R CamCube 2.0 camera it was shown shown that the confidence measure offers superior results to classical amplitude filtering approaches, and that it has the potential to be incorporated into a real-time, hand-held system without the need for additional inertial sensors. or high-precision robotic arms as done previously. At the same time, it improves the reconstruction accuracy, and removes the need for point-cloud smoothing and median filtering.

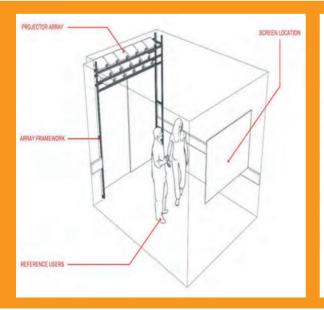


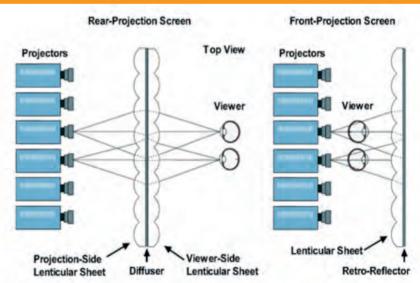
Confidence in the image is represented by colour

Multi-view display

Group project with Tim Scully, Sarat Babu, Christos Sakellariou and James Martin

The goal of this project was to build a multi-view display using arrays of projectors. The design of the display was based on a commercial example, the HoloVisio display by Holografika. This involved using a multi-view display with arrays of projectors. Each projector was positioned at a slightly different angle relative to a lenticular screen, which only allowed certain projections to be seen as the viewer moved around the screen. The approximate arrangement of the projectors and screen can be seen in the figure, below. By orienting the projection angle for each projector, the incident rays on the screen are unique. The lenticular screen is shaped in such a way that it refracts the light back so it is best viewed at a particular angle. In order to recreate the 3D effect, each eye of the observer needs to see a slightly different image. When viewed from different positions, two images of a 3D scene will differ due to the parallax effect. Objects closer to the viewer will have a greater change between each eye's image, and those further away will look more similar to each other. A good quality lenticular screen will allow each projection, from each projector, to be more distinguishable from each other, hence leading to a more impressive 3D effect. There are two methods of construction of this multi-view system: back-and-front-projection. The rear-projection method requires two lenticular sheets with a non-diffuse layer sandwiched in between, whereas the front-projection method requires one lenticular sheet, with a retro-reflective material behind it. Due to the simpler screen construction and also reduced room.





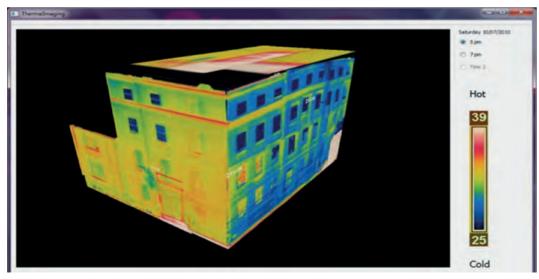
An array of projectors focused on a lenticular screen creates a stereoscopic projection

Thermal leak imaging

Group project with Jessica Wardlaw, Maciej Gryka and Fabian Wanner

Heating and cooling efficiency is very important in today's world because heating is expensive, due to the high oil and energy costs. Furthermore, global warming has to be stopped as quickly as possible. Some of the main causes of energy wastage are small gaps in the outside of buildings, bad insulation or inefficient windows. All of these can be detected with the help of thermal imaging cameras. Currently, there exist a number of companies that specialise in such services. However, we are not aware of any that perform 3D reconstruction of obtained models. This project analyzes the effect of a new representation for thermal images, on the ability of users to locate heat leakages on a building. This new visualisation is a rough 3D model of a building. The project implemented an algorithm, which created a simple geometric representation of a scene. This was then textured with visual, as well as thermal images, to allow users to switch between two views. It also allowed switching between images taken at different times.

The study involved a user evaluation that compared users' abilities to detect heat losses with normal thermal images and our 3D model. Also, obtaining feedback on the graphical user interface (GUI), that we had designed. Non-expert users were able to interpret otherwise complex thermal data and to start explaining the patterns they saw. Because the model facilitated a global overview of the building and put each thermal image into context. The model allowed for a much richer and more intelligent interpretation of the thermal data, with users becoming more confident of their assessments.



Thermal imaging of a building reveals the angle of the sun on the near side, and the leakage of heat on the cooler far side

Aesthetic judgement

Cristina Amati & Anthropics

Supervised by Dr Tim Weyrich, UCL Computer Science

Since its dawn, mankind has always been fascinated with imagery. Starting from imitating and recording their surroundings to creating original works of art, people have cultivated a sense for the beautiful and the sublime. Nevertheless, beauty has always been regarded as a highly subjective attribute, that acts pleasurably upon our senses. Thus, it was never formally described, because it encompases so many dimensions of human perception and interpretation, coupled with personal and social context. The question that we are facing is whether beauty is just a combination of lines, shapes and forms as dictated by the formalist theories of the 20th century, the meaning assigned to the objects in the image, or an intricate combination of both.



Comparing images adjusted for colour helps in understanding how we judge what is beautiful

Recently, computer science has tapped into the quantitative aspect of aesthetics by analysing vast photo collections that have been peer-rated by online users. The general approach is to train a classifier or regressor using image features based on low-level descriptors of colour, shapes, photographic rules and other observations which can then classify images as pleasant/unpleasant, or assign them an aesthetics rating.

The success of these methods are highly encouraging, by hinting that there might be an underlying, generally accepted, model of beauty. However, these models are inherently cryptic. They do not actually explain what a beautiful image is, and how it is composed. Moreover, we cannot assert whether the extracted models can really be extrapolated to quantify the preferences of the majority of people. Or if they represent a biased view of one photographic community, or trend.

An aim of this work is to seek to answer the question: "Why is an image beautiful?" To address this, the research decouples the visual qualities of images from their evocative function, analyses them separately, and then combines them into a holistic aesthetics model. Recent work in neurology (and particularly neuro-aesthetics) has shown that there is an underlying physiological mechanism that makes us respond positively to certain types of imagery.

Psychology, art theory, and aesthetics philosophy can give us an insight into the power of images, and the evolution of human thought in art. Thus, we can discover recurring patterns between images, and attribute emotions and meanings.

Once the aesthetics model has been acquired this, the goal is to use it for improving existing pictures. This can assist the creation of new imagery, aesthetics-based searches and image retrieval. Ideally, allowing us to incorporate ratings and suggestions directly onto digital cameras, helping photographers obtain better photographs, or to improve their skills.

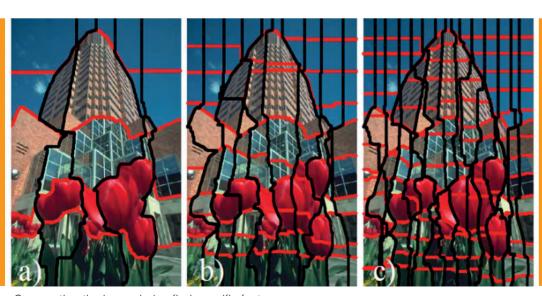
Segmenting images

Dr Alastair Moore & Sharp Electronics

Supervised by Prof Simon Prince, UCL Computer Science

Image segmentation in computer vision refers to the process of dividing an image into multi-pixel, often irregular, contiguous regions. The resulting groups of pixels are then referred to as 'superpixels' and can be used to give a different, more salient representation of an image. One typical use for superpixel image segmentation is to find lines and boundaries in an image (as these will likely also be on the boundaries between the superpixel regions) or to assign a single, often more relevant, meaning to each region. In medical imaging, where a stack of images is frequently available, image segmentation leads to a similar stack of contours which can be used for 3D, volumetric reconstruction.

Segmentation is often an important first step in many image-processing algorithms and some practical applications for superpixel image segmentation are in object-location and object-class recognition in an image. Methods for image segmentation range from the simplest methods of thresholding to more 'intelligent,' learning-based methods. This project offered a solution for globally-optimal image-segmentation which performs better than previous methods by using information about the boundaries of a superpixel as well as its contents.



Segmenting the image helps find specific features



Responsive architecture

Christopher Leung & Haque Design + Research

Supervised by Prof Stephen Gage, UCL Bartlett School of Architecture

This research interest surrounds the ways that buildings and, in particular, façades could be dynamic and physically articulated to respond to changes in synoptic weather conditions across seasons. A practical demonstration of this approach was the Deployable External Insulation (DEI) developed at the Bartlett, and realised with a full-size pavilion in 2008, using the Bartlett's manufacturing workshop. This showed how a kinetic façade with movable shutters could be operated by passive thermo-hydraulic actuators. These actuators exploit the liquid-to-solid transitions of Phase-Change Materials (PCMs) such as paraffin waxes, when they absorb sufficient thermal energy above ambient they expand to generate a large hydrostatic force; a property that can be used to directly move building elements.

A novel mechanism was invented to demonstrate how an arrangement of this actuator could respond to different outdoor conditions by moving the façade's shutters into three different positions (see figure, right). When it is cold or during night-time, the shutters are closed to provide insulation and reduce heat-loss. During daytime, they are fully opened to admit daylight. And during hot weather, they are half-closed to provide shade and limit the risk of overheating.

Two linked actuators, each powered by a different paraffin wax, respond individually to the amount of ambient thermal energy that is available. The climate-controlled chambers at the Bartlett were used during the extensive development and testing of this mechanism with a series of full-size prototypes. This invention is the subject of a patent application to the UK intellectual property office, through UCL's business department.

Visible and thermal-band time-lapse imaging was used to produce photogrammetric solutions. These allow the actual 'as-operated' observations to be compared with the 'as-simulated' by virtual models. The linking of practical, computer-based, thermal laboratory and field observations has given the insights necessary to account for the passive flow of thermal energy through these unique actuators. This informed the development of a thermal engineering analysis, using the finite element method.

The approach taken by this research, challenges the architect with the task of achieving a balance in the design of the building envelope. One in which the passive flow of energy into and out of this type of actuator system can be marshalled as the organising principle for the variety of its states, while matched to the prevailing seasonal variation in ambient conditions.

Personal statement

As a practising architect, I am motivated to conduct applied research that adapts and transfers technology from other fields. My doctoral studies on the EngD have explored the feasibility of a thermo-hydraulic actuator system that embodies both the actuation method, and the control logic entirely passively, as an integral part of the architectural design for the building's envelope. The EngD Centre sits amongst the engineering, prototyping, climate testing and manufacturing facilities that are needed to explore this research topic. I am currently completing a thesis as a candidate for the EngD and regularly give lectures on the topic of climate-responsive architecture and his findings on the subject.



Generative forms

Katrin Jonas & Buro Happold

Supervised by Prof Alan Penn, UCL Bartlett School of Graduate Studies

Recent advances in digital technologies, and improved accessibility to them, have led to rapid changes in architectural design. There has been a shift in the aesthetics from the modern orthogonal building envelope proposal to elaborate curved and folded forms. There are two dominant cost factors in such building designs, firstly geometric design changes during the engineering of the project, and secondly the manufacture of customised building parts.

The current paradigm in the design and engineering of these elaborate building structures is such that the conceptual model is provided by the architect, and the engineer must then break it down into buildable parts, in an iterative design process together with the architect. A substantial amount of time goes into this process of reverse engineering a design. Every time the design changes, and a new iteration cycle is entered, the information has to go through all previous stages again, and the cost for the project rises. It is often here that a project has to suddenly stop, or drastic rationalisation has to take place, in order to keep the project alive and within budget. Reverse engineering the building this way, and applying a rationale to the geometry retrospectively by fabrication and construction, is a process of post-rationalisation. The alternative is pre-rationalisation. This means incorporating fabrication and building rationale in the initial design conception. In such a bottom-up or generative process, the components are defined first, and the global form derives from them. Pre-defining the building components allows them to be prefabricated, improving their cost and performance over time.

The existing gap between elaborate form proposals and the procedure to engineer, manufacture and build them has been widely recognised and there is interest in pre-rationalising the approach and implementing rules which go beyond the mere aesthetic 'Gestalt'. To address this problem, the project prototyped a software plug-in. This implements a bottom-up process which allows architects to design complicated forms from simple components and connections, thereby turning the current paradigm on its head. The existing tool creates and optimises digital models of articulated envelope structures, such as roofs and building façades. The focus of the plug-in prototype is to utilise a discrete set of local components that are assembled to form different global envelope geometries, and to optimise them using an evolutionary optimisation algorithm. It can be imagined as akin to Lego, where the user can specify how many different types of Lego bricks are employed to build the next structure. Being able to control the number of different elements from which a complex structure is built, has a significant impact on both the cost of a building, the envelope, and its sustainability, through the re-use of formwork/moulds.

Personal statement

My work at Buro Happold involved the actual design, comparative form analysis, structural performance and manufacturing optimisation, as well as the material definition and reverse engineering for production of geometrically complex structures. I was working as part of a small team called the Generative Geometry Group, later the Integrated Geometry Unit. An important part of the work was the analysis, communication and preservation, or development of the design intend throughout the design and realisation process.

Being an EngD candidate allowed for in-depth analysis of projects, and the subsequent recognition of problem pattern. I also received a postdoctoral knowledge transfer grant to further the potential application of the technology I developed during my studies.



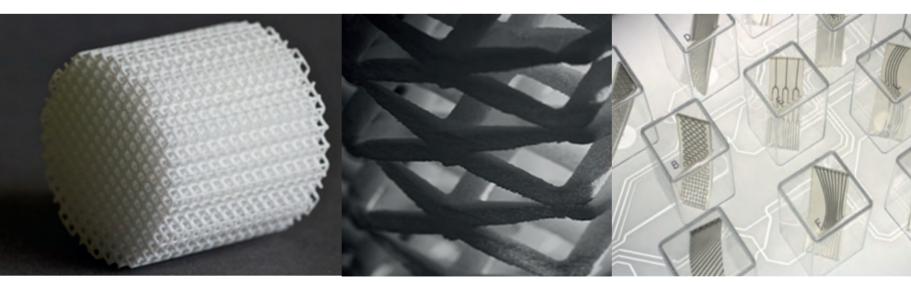
Architectural designs are built from pre-defined components

Hybrid materials

Sarat Babu & Bread 1851 Industrial Fellow

Supervised by Dr Sean Hanna, UCL Bartlett School of Graduate Studies

Innovations in fabrication technologies and material engineering are developing new methods and levels of adaptability in the materials we apply to objects and structures. This project looks at understanding how to incorporate these advances into the holistic design and architectural process to fully exploit the advantages they present. Based on a proprietary process of manufacturing hybrid functionally graded materials developed by BREAD, the project is investigating this fabrication method with new models and design systems by which it can be applied to the real world.



Innovative materials adapt to their specific environments

Articulated models

Jacques Cali & ArkaSoft

Supervised by Dr Tim Weyrich, UCL Computer Science

Over the past two decades a lot of emphasis has been placed on capturing the world around us and displaying it on screen. Now, with the advent of 3D printing and improved fabrication methods, this project is looking at ways in which we can use traditional graphic methodologies to fully utilise this technology. In this way, real, physical models with a controlled appearance and mechanical properties can be created. By adding a physical element to the pipeline, a new set of non-trivial challenges is being introduced

The nature of this project means that topics such as geometry processing, physical appearance modelling and measurements, and calibration issues are all relevant.

The next stage of the project is to look at other ways in which we can control physical aspects, in terms of mechanics and visualisation. New fabrication techniques have enabled us to create seemingly impossible objects. However, some of these objects are too complex, or time-consuming, to be generated manually. As such, the study examines ways in which computer graphics techniques could be applied, as well as machine-learning algorithms, to aid the creation process.

The main contributions from the work include: a derivation of a single versatile physical joint design incorporating a wide range of movement, and friction to control static poses. An interactive joint-fitting pipeline, and underlying algorithms to support it. The result means that a user can, with minimal interactions, generate physical, 3D printed, friction enabled, articulated models that work directly out of the 3D printer.



Animal forms can be printed to include moveable joints

Intelligent materials

Dr Sean Hanna & Foster + Partners / RCUK Academic Fellow Collaborative project (2003) with Dr Siavash Haroun Mahdavi and Dr Jamie O'Brien

Supervised by Prof Alan Penn, UCL Bartlett School of Graduate Studies

The structural properties of organic materials such as wood or bone are dependent on the shape taken by the living cells, as they develop within an environment subjecting them to load and stress. Material thereby becomes distributed where it is needed, and in the appropriate geometry, resulting in an efficiency that is hard to match in artificial materials. Then-new technologies of digital manufacturing, such as laser sintering and stereolithography came to allow the manipulation of geometry at scales measured in nanometres, by which many of the properties of such natural geometries may be replicated. A sample of structure only a centimetre or two in length might contain hundreds, or even thousands, of structural members, each less than a millimetre long, individually placed to best carry loads, minimise weight, or take impact in particular ways (see figures, right).

Designing such structures, might require the placement of millions of members in an object of reasonable size. This work developed a method, by which the intense computation required to simulate and optimise the geometry of objects, designed for a particular purpose can be made feasible. For a chair, loads from the seat must be distributed appropriately through the legs, requiring the density of structure to increase around crucial connections and their orientation to follow the lines of force, even though the geometry is a seamless and continuous whole. For a crash helmet, impact must be distributed and absorbed throughout the object, minimising the amount transferred to the wearer. The design of millions of members is made possible by taking advantage of the cellular nature of the structures. As in nature, a uniform topology of cell is produced by evolution in the form of a genetic algorithm.

This provides an efficient structural basis for changes in geometry and density to be made more easily. In later work, these geometrical changes are made by analogy, to the living cell's intelligent response to stimulus, and allowing the cell to learn appropriate geometries with machine intelligence, such as artificial neural networks.

Beginning as a collaborative project between Sean Hanna, Siavash Haroun Mahdavi and Jamie O'Brien in 2003, this research has continued in a number of directions. Hanna has focused on the application of machine learning to allow greater efficiency by virtually embedding the required 'intelligence' within each artificial cell, and extending the method to better accommodate the freeform geometry required of many designed objects. Haroun Mahdavi applied the technique to the manufacture of robots, and successfully developed its commercial potential, (see page 82). Both have collaborated with industrial designers in the design of products incorporating the technology, such as Assa Ashuach's 'Ostion Chair' and Timothy Schreiber's 'PAN_07'.

The VEIV Centre continues research in this area, with several research engineers currently investigating the control of materials in digital manufacturing under Hanna's supervision.



Innovative structures are built organically to meet the pressures of their environment

Functional complexity

David Andreen & Laing O'Rourke

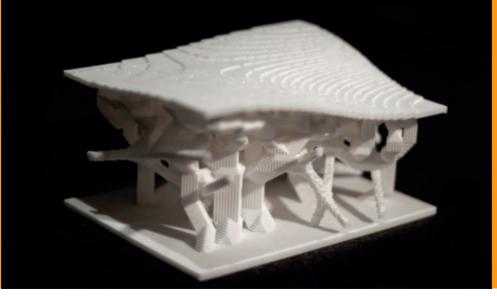
Supervised by Dr Sean Hanna, UCL Bartlett School of Graduate Studies

Emerging digital fabrication technologies are posed to radically change the way we design and build buildings in the near future; computer-numerically controlled processes, such as 3D printing, enable complex and custom manufacturing at a previously unattainable level. Within this context, conventional design strategies are inadequate and so new algorithmic tools are emerging. In this process, the design of the physical environment is becoming increasingly similar to the processes normally associated with virtual environments.

Traditionally, the architect or engineer is used to being in complete control of the design process, predicting all outcomes and possible failure states. The downside of this paradigm is the inability to deal with rapidly escalating complexity, following the manufacturing digitalisation processes. And also the catastrophic effects that often follow when failure states do occur, due to errors in the predictions, or unusual load conditions. In order to break through this complexity ceiling, self-organising

algorithms can be embedded in the design process, releasing the designer to focus on the wider picture.

This work investigates how complex geometries are generated in nature, and the performative role this geometry plays for the organism that creates them. Particularly, this deals with the mound structures of sub-Saharan termite species. These structures act as a respiratory organ for the mound, bringing in oxygen and ventilating metabolic waste gasses while maintaining a delicate internal homoeostasis. It is speculated that this happens through the interaction of the internal mound geometry and the ambient wind which is turbulent and low-energy. Combined with computerised fabrication processes, an understanding of these phenomena would allow engineers to incorporate similarly functional structures into building envelopes, potentially increasing comfort, while reducing the energy consumption, and carbon footprint of the building concerned.



Structures in nature, such as termite mounds, inspire new methods in constructing materials



High-end fabrication

Siavash Haroun Mahdavi & BAe Systems

Supervised by Dr Peter Bentley, UCL Computer Science

This work deals with a robot where no domain knowledge is given at all. The genetic algorithms used to control these robots are deeply embodied within their environment. This embodiment means that a very simple genetic algorithm is able to control smart materials like shape memory alloys without any previous domain knowledge about them. The control and exploitation of smart materials is also extended to stereo-lithographic models. The aim of the work was to create a robot made entirely out of smart materials, and able to perform a variety of tasks from locomotion, to damage recovery, even optimising its shape to get the best reception.

Personal statement

Since completing my EngD I have continued to publish papers on 'Microstructure Optimisation' and give lectures at numerous international conferences including one at NASA. With support from UCL Business, I started Complex Matters Ltd in 2005, a design consultancy that works with designers and engineers in creating objects with optimised structures that enable them to exhibited unparalleled material properties. More recently, I moved on to work at Within Technologies Ltd, specialising in higher-end engineering applications and, in 2008, I launched Digital Forming, a company that specialises in the development of software for the mass customization of 3D Printed consumer products. I have also worked on high-profile works in design, including the Osteon Chair in collaboration with world renowned product designer Assa Ashuach.



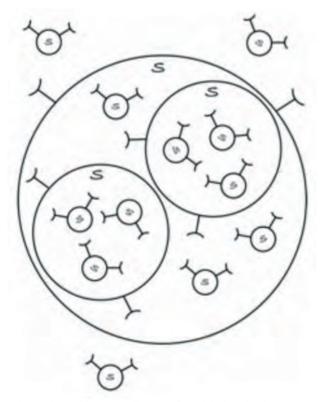
Credit: Assa Ashuach

Systemic architecture

Christos Sakellariou & Toumaz Technology

Supervised by Dr Peter Bentley, UCL Computer Science

It is widely accepted that natural computation, such as that which happens in our brains, is far superior to machine computation for tasks such as learning and parallel processing. As conventional silicon-based technologies are about to reach their limits, researchers have drawn inspiration from nature, founding new, computational paradigms. One example is 'Systemic Computation' (SC), which is designed to be a model of natural behaviour as well as a model of computation. Incorporating natural characteristics, it defines a massively parallel non-von Neumann computer architecture that can model natural systems efficiently.



A schematic architecture diagram for a bio-Inspired computer

There are three SC implementations to date. The first two attempts simulate a systemic computer, using conventional sequential hardware, and provide a satisfactory proof-of-concept, but suffer from poor performance. The latest attempt successfully maps a part of the model on the parallel resources of a GPU, and achieves performance gains in the order of one hundred. Clearly, the full potential of SC cannot be exploited using conventional hardware.

This work provides a thorough review of Natural Computation and Systemic Computation, with focus on the various hardware-based approaches to perform. The investigation involves the compatibility and suitability of the available hardware implementation platforms, with respect to the SC paradigm, identifying FPGAs as the most appropriate among them, and describe the first prototype Hardware Architecture of Systemic computation (HAoS).

HAoS is a novel, custom, digital design, which addresses the SC architecture parallelism requirement by exploiting the inbuilt parallelism of a Field Programmable Gate Array (FPGA), and by using the highly efficient matching capability of a Ternary Content Addressable Memory (TCAM). Basic processing capabilities are embedded in HAoS, in order to minimise time-demanding data transfers, while the optional use of a CPU provides high-level processing support. It is a functional standalone prototype, which takes into consideration programmability and scalability. Analysis shows that the proposed architecture provides an effective solution, in terms of efficiency, versus flexibility trade-off.

The base HAoS architecture is optimised and extended to a practical hardware platform accompanied by a software framework to provide an efficient SC programming solution. This overcomes the limited performance or capabilities of previous implementations.

Dynamic metrology

Andrew Barnes & Renishaw

Supervised by Prof Steve Hailes, UCL Computer Science

This research project is focused in the development and application of modern systems where the combination of scalability, cost and performance of core components are key to the sponsor's product line. The goal is to provide innovative solutions that allow their customers to better control the quality of their products. Each of these solutions currently requires a local, real-time controller, and a local user interface (usually another computer). The motivation for this research lies in finding better ways to manage the computational requirements, while reducing the number of components and thus the cost per installation.

Heterogeneous computing solutions may allow for an amount of shared processing over non-core requirements, while also providing direct access to hardware such as a GPU, or dedicated real-time processor for core requirements; effectively running a hybrid virtual server. One of the main problems that this work aims to solve is the accessibility and management of core hardware, through the architecture present in modern computers. We anticipate that this may require complex routing/mapping of hardware address, and context switches.

Another solution may also be found in combining multiple servo clocks (generating the precise ticks used in real-time processing) into one physical execution unit. This work would contribute to reducing the system's hardware complexity, and thus its cost. Wherever high performance and low latencies are required, software-based solutions will struggle to provide a like-for-like result against a hardware implementation. This can be seen clearly with the comparison between a real GPU and an emulated GPU, used in most virtual machines. To a point, software-based solutions provide scalability, along with the benefits in virtual machines of being independent of any hardware, which allows for features such as migration.



Enhancing the performance of metrology robotics

In-between contexts

Karen Martin

Supervised by Prof Alan Penn, UCL Bartlett School of Graduate Studies

The built environment is a complex synthesis of social, spatial and technological systems that intertwine at various points in space and time. More than ever, we are encountering hybrid electronic and physical spaces where computing capabilities are either embedded into the environment or are portable, personal and connected. As a result previously well-defined and distinct locations and activities are becoming increasingly indeterminate. These are the in-between places and activities of everyday experience: the wait at the bus stop; the queue at the supermarket; the daily commute; picnicing in the park.

Urban computing is acknowledged to have the potential for transform our relationship with the built environment yet uncovering the link between spatial form and social action is challenging. Developing strategies for the design of urban computing tools and services for inbetween spaces assists designers for urban computing to address the complexity of urban space and make the most of the benefits offered by urban computing.

Achieving this requires a thorough understanding of how in-between spaces are constructed. My work addresses this by examining the components of in-between space and aspects of design collaboration. Based on an understanding of these spatial and social properties strategies for the design of urban computing tools and services can be described.

An observation study of in-between spaces and a series of design workshops identified four properties that mediate the social-spatial relationship in situations of in-between space. These findings suggests that interaction can be thought of as a material, capable of being shaped and moulded in a similar way to wood, concrete or light. Manipulating the properties of interaction in this way enables designers to create specific spatial, social or technological experiences, for example, building tools, environments or applications to encourage communication, enhance privacy or control crowdedness.



Novel architectures form around transient spaces

POSTSCRIPT

Dr Jamie O'Brien, Editor & Manager, **EngD Centre in VEIV**

As Editor, my aim in preparing this book was to showcase VEIV research engineers' outstanding contributions to the state of the art. The contributors' articles each outline their highly inventive, ground-breaking and often thoughtprovoking projects across the breadth of science, engineering and design disciplines. In summation, I trust that I speak for our 75 students and over 100 staff members in saying that VEIV has now secured a leading position at the forefront of academic-industrial collaborative research in this technology field.

The Centre's highly inventive output has been enriched by the flows of knowledge through its external partnerships. I have been involved with the engagements strengthen VEIV as an academic centre, and also to offer novel models for industrial R&D. A key advantage of working with an academic institution is the protection of R&D from the pressures of the competitive business environment. In this way, VEIV provides the time and resources for creative researchers to build, test and deploy effective solutions in challenging domains.

VEIV's success has been measured in our students' outstanding academic achievements, including several publications at some of the field's highest rated destinations, such as SIGRAPH, ACM Transactions in Graphics and Institute of Navigation GNSS+. Several EngD graduates are now working in academic roles, or in advancing the productive relationship between academic and industry. The Centre has also benefitted greatly from the award of five Industrial Fellowships by the Royal Commission for the Exhibition of 1851. These valuable Fellowships reward the potential impact of the students' research projects, as well as the excellent quality of their academic-industrial collaborations. As Manager of VEIV, I believe we would not have achieved such robust endorsements of the Centre's endeavor and performance if it were not for the active participation of its academic, industrial and non-profit collaborators. I hope this volume has served to portray something of this dynamic and enterprising spirit.

The Engineering Doctorate Centre in Virtual Environments, Imaging and Visualisation is University College London's leading centre for science and engineering research in this field. EngD projects have led to technology advances in several areas, including body scanning, medical imaging, 3D fabrication, and generative design. A range of other projects have enhanced graphics and interactions in healthcare, urban planning, robotics and communications.

Founded in 2001, the EngD Centre in VEIV broke new ground for academic and industrial collaboration in doctoral training. Since then the Centre has achieved novel advances through dynamic and creative collaborations with a range of external partners. Innovation in science and engineering continues to be channeled via world-class training and substantial support for enterprise.

This volume celebrates over a decade of outstanding research and academic achievement. The portfolio of current and complete projects is a legacy to VEIV's position at the forefront of collaborative research in graphics, interactions and complex design. VEIV now looks ahead to furthering the advancement and innovation of visual and interactive technologies.

