

**CENTRE FOR ADVANCED
SPATIAL ANALYSIS
Working Paper Series**



Paper 66

**INTERPRETING
INTERPOLATION:
THE PATTERN OF
INTERPOLATION
ERRORS IN
DIGITAL SURFACE
MODELS DERIVED
FROM LASER
SCANNING DATA**

**Smith, S. L.
Holland, D.
and Longley, P.**



Centre for Advanced Spatial Analysis
University College London
1-19 Torrington Place
Gower Street
London WC1E 6BT

[t] +44 (0) 20 7679 1782
[f] +44 (0) 20 7813 2843
[e] casa@ucl.ac.uk
[w] www.casa.ucl.ac.uk

http://www.casa.ucl.ac.uk/working_papers/paper66.pdf

Date: September 2003

ISSN: 1467-1298

© Copyright CASA, UCL

Interpreting Interpolation: The Pattern of Interpolation Errors in Digital Surface
Models Derived from Laser Scanning Data

Smith, S. L.[†], Holland, D.[†], and Longley, P.[‡]

Ordnance Survey, Rm C530, Romsey Road, SOUTHAMPTON, S016 4GU
slsmith@ordsvy.gov.uk

[†] Ordnance Survey
Research and Innovation

[‡] Centre for Advanced Spatial Analysis,
University College London

Centre for Advanced Spatial Analysis, University College London,
1-19 Torrington Place, London WC1E 6BT, UK

Interpreting Interpolation

The Pattern of Interpolation Errors in
Digital Surface Models Derived from Laser Scanning Data

Smith, S. L¹, Holland, D¹., and Longley, P²

¹Ordnance Survey

Research and Innovation

² Centre for Advanced Spatial Analysis,

University College London

Ordnance Survey, Rm C530, Romsey Road, SOUTHAMPTON, S016 4GU

slsmith@ordsvy.gov.uk

Errors within height models have, in the past, been communicated in terms of global measures of accuracy for the model. Such quantification ignores the spatial structure of errors across the surface, hindering subsequent analysis. This paper demonstrates the importance of understanding the spatial structure of error using, as an example, the creation of a Digital Surface Model (DSM) from laser scanner data.

1 Introduction

Laser scanner data are obtained as an irregularly spaced set of points. Many software packages require that these points are interpolated onto a regular grid for analysis and visualisation of a height model. Despite the fact that laser scanner points are sampled at very small separation distances, the interpolation from points onto a grid can introduce a degree of uncertainty into the model. The level of this uncertainty can vary greatly with different interpolation methods. The aim of this paper is to identify these differences, and to suggest how they may cause the propagation of error in later stages of the modelling process. The interpolation techniques used within this investigation are: bilinear, bicubic, nearest neighbour, and biharmonic splining.

Each of the methods produces slightly different height values across the surface. Several global characteristics of the various surfaces have been noted previously in the literature. Zinger et al (2002) commented that linear interpolation will tend to overly smooth and deform building edges. However, such general characteristics reveal little about the exact spatial pattern of error within a surface model. Lloyd and Atkinson (2002) further investigated the quantification of error within interpolated surfaces. The authors focused on a comparison of Inverse Distance Weighting (IDW) and kriging interpolation, and quantified the inaccuracy in each surface. Such measures are useful general indicators of error within surface models, again however they do not reveal anything about the spatial pattern of errors across the surface. It is argued here that understanding this pattern is important for many data users who wish to use the interpolated surface in subsequent analysis.

The errors in four DSMs are identified and quantified, and the spatial pattern explored. The effect of changing grid size on the patterns of error is also commented upon. Particular emphasis is placed on the significance of the spatial pattern of the error for 3D feature extraction and reconstruction.

2 Methodology

Four DSMs were created from a subset of a first return laser scanning dataset, supplied by the Environment Agency. The data were captured from an airborne sensor, at a point density of ~2m. The raw points (figure 1) were first resampled onto a regular 1m grid¹, using four interpolation methods. Spatial interpolation may be defined as the procedure of estimating the value of a field variable at unsampled sites within the area covered by sample locations (Zhang and Goodchild, 2002). The basic assumption underlying any interpolation procedure is that points that are close in space are more likely to be similar than points further apart. The four methods used for interpolation in this research were bilinear, bicubic, nearest neighbour and biharmonic splining. All interpolation was conducted in the Matlab environment, and the resultant surface forms produced are shown below in figures 2a and 2b.

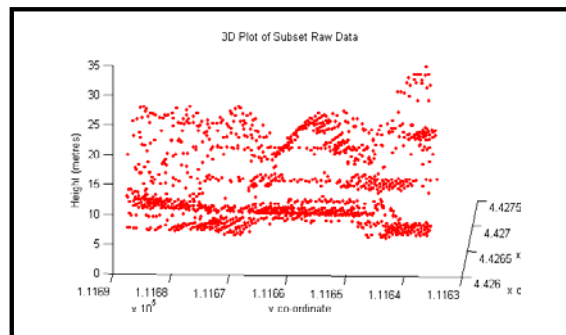


Figure 1: The raw lidar point cloud. For many software packages and applications, these points need to be resampled onto a regular grid in order to create a surface model.

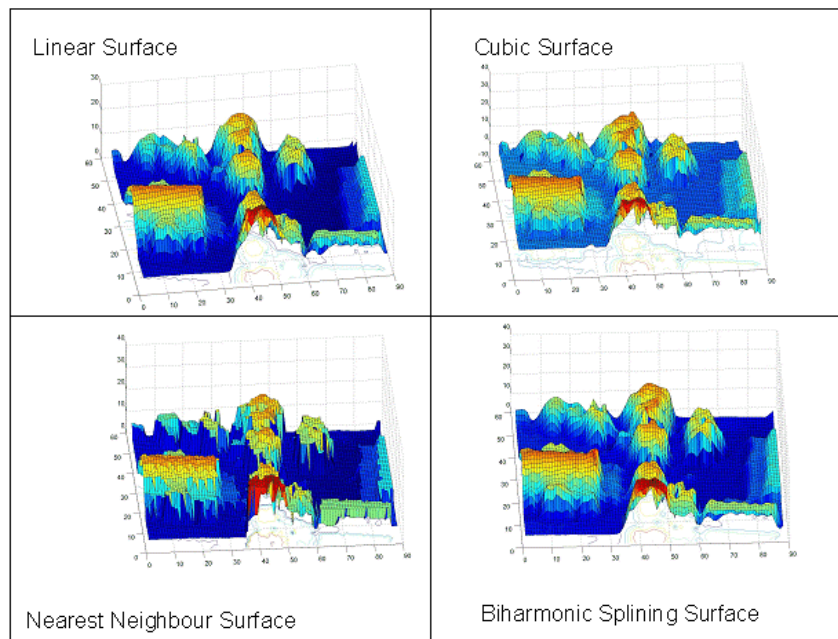


Figure 2a: Showing the different surfaces created using the interpolation methods. Surfaces created from raw first return laser scanner data kindly supplied by the Environment Agency for England and Wales.

¹ The effect of using alternative grid sizes was also investigated – see section 3

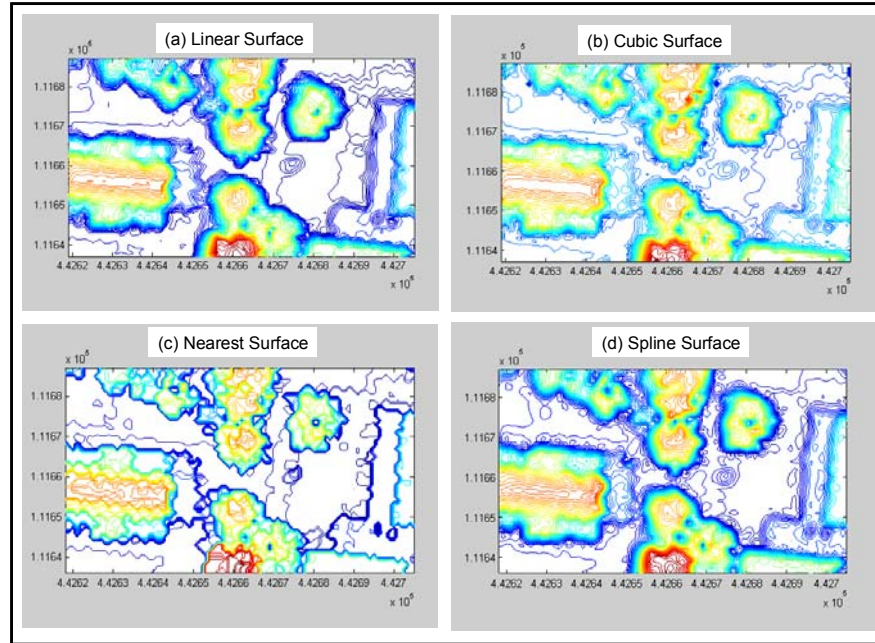


Figure 2b: Showing the contour maps for the four DSMs. Note the blocky nature of the surface contours created using the nearest neighbour methodology

It can be seen from the above figures that there are differences between the forms of the surfaces. In particular, the nearest neighbour surface is noticeably blocky in appearance, whereas the linear and the splined surfaces are extremely smooth. The form of the cubic surface appears to be somewhere between the two. However, comparisons based on qualitative characteristics of surface form reveal nothing about the amount, or the distribution, of errors across the surface. In order that the characteristics of the errors introduced by interpolation are fully understood, some quantification is required. Error prediction may provide important information about the deficiencies of a certain method and so may be an important input when designing a methodology for a particular application. In this investigation, the methodology for measuring error was designed to identify the following:

- Where the values on the interpolated surface differ from the raw data points
- A quantification of these differences
- Why these differences occur
- Why the errors are significant

For the purposes of this investigation, the error (ϵ) at each investigated point within the surfaces was considered to be the difference between the raw data point ($Z(x)$) and the interpolated value ($Z_i(x)$) for that location (see eq.1 below).

$$\epsilon(x) = Z(x) - Z_i(x) \quad (1)$$

Where: ϵ = predicted error, x = location of point, Z = height value, Z_i = interpolated height value

In accordance with this definition, the error for each point on the surface was calculated using a combination of two methodologies: cross-validation and jack-knifing. Different characteristics of error are highlighted by these two methods.

The cross validation procedure involved using all of the data points to create the DSM. The values on the interpolated surface were then compared to each of the raw data values. This technique can expose general characteristics of the surface, and can provide some initial indication of differences between the surfaces. Preliminary statistical analysis was conducted on the error predictions. Techniques used included standard deviation, root mean squared error (RMSE), maximum positive error, maximum negative error, and mean. The results showed that the Nearest Neighbour interpolator produced the highest mean error, and the highest RMSE. In contrast, the splining interpolator was found to produce the least error across the surface. Whilst useful indicators of trends, these global statistics cannot be used to assess the spatial variation of error across the surface. Often it is this variation which is of most significance, particularly where the DSM is to be used for 3D modelling. The pattern of individual errors was examined by plotting the locations of the interpolated points and assigning them a colour in accordance with the error calculated for that point (figure 3). It was observed that there was a strong spatial dependence of the highest magnitude errors. Visual inspection of the corresponding orthorectified photograph of this area showed that the patterns of highest error occurred where there was vegetation in the scene.

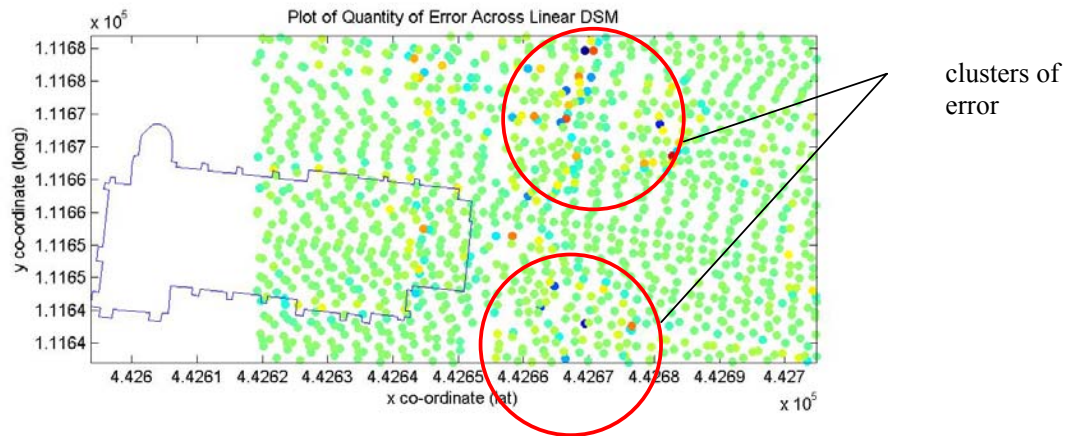


Figure 3: Point plot showing the spatial variation in error across the surface. Low/no errors are represented in green, positive errors in red, and negative errors in blue. The blue line represents the boundary of a church – and highlights the relationship between building roof edge and magnitude of error. Building outline vector reproduced with permission of Ordnance Survey © Cc Ordnance Survey. All rights reserved. The error shown in the diagram above is the difference between the interpolated point height value and the raw data height for that point.




Figure 4: Orthorectified photograph of the corresponding area. Aerial photography reproduced with permission of Ordnance Survey © Cc Ordnance Survey. All rights reserved.

The analysis was repeated to identify whether there was a difference in point errors between building and vegetation points. It was found that there was significantly more error over vegetation. This is likely to be the result of the characteristics of the surface. Laser scanner pulses cannot penetrate building surfaces (except for windows), as a result all of the captured points will lie on one continuous plane. Vegetation, however, will permit some penetration of the laser pulse between branches and so data points captured in vegetation will be much more scattered in 3D – the computation of an interpolation surface is accordingly more difficult in these areas. The pattern of higher magnitude errors may be used in subsequent feature extraction to both differentiate between buildings and vegetation, and possibly to identify the approximate boundaries of above ground objects. Whilst it is no surprise that the greatest errors occur at feature edges (where there is the greatest change in slope) the patterns and quantification of this error has not been investigated to date. This is potentially significant as it indicates that at least some surface texture information can be obtained from first return laser scanner datasets – thus reducing the need for additional information or data such as intensity returns. The potential for the spatial pattern of error to also reveal this information is therefore significant.

Cross validation is a useful indicator of the general characteristics of an interpolation method, however it cannot be used as a measure of the robustness of a particular technique. As an alternative, Lloyd and Atkinson (2002) suggest that the jack-knifing technique can be used to identify which interpolation algorithm is the most accurate. The jack-knifing technique (previously used by Lloyd and Atkinson, 2002; Priyakant et al 2003; and Deutsch and Journel, 1998) involves extracting random data points, and using the remaining data to create a surface. The interpolated surface is then used to predict the values of the extracted data, and the differences between the extracted and predicted values considered to be a measure of the error. The central idea underpinning this technique is that more robust algorithms will perform well despite a reduction in the number of raw data points used to form the surface. This technique was repeated using 95%, 50%, and 25% of the raw data points. The linear surfaces resulting from the jack-knifing are shown below.

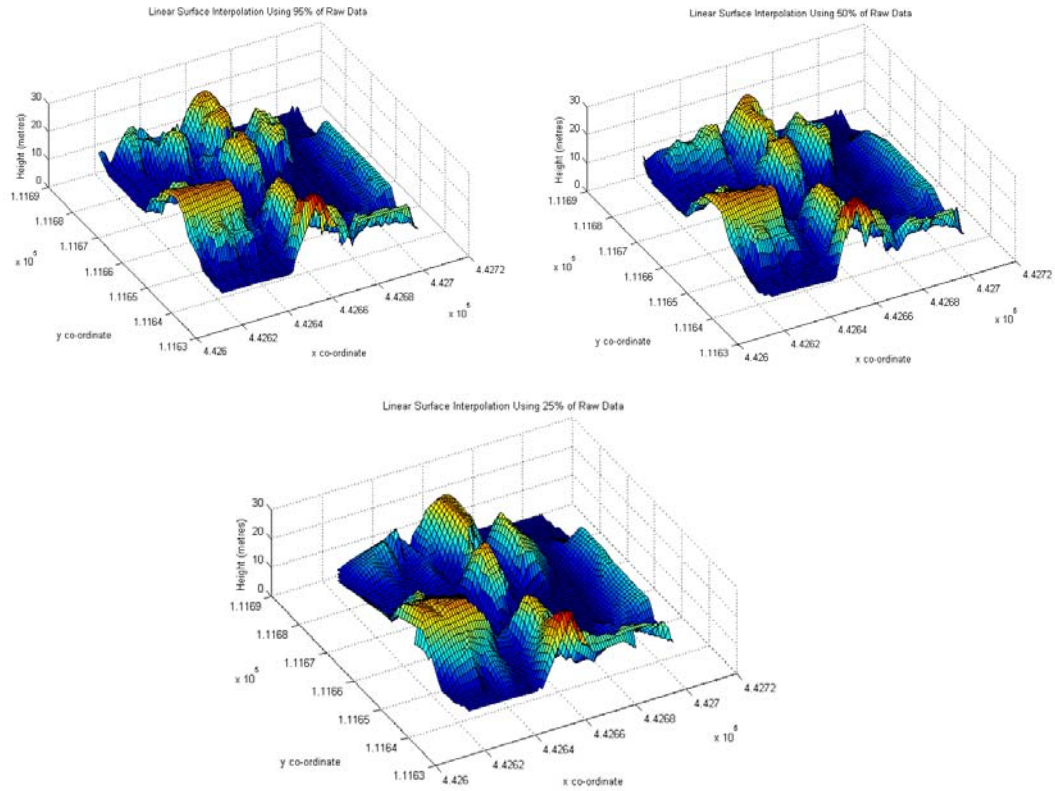


Figure 5: The Linear Interpolation Surfaces using 95%, 50%, and 25% of the raw data respectively. Note the loss of surface form as the number of raw data points decreases.

Figure 6 below, shows similar surfaces, but created using the nearest neighbour algorithm. The surface form using this algorithm is much blockier and there is a marked decline in surface form for the lower levels of jack-knifing.

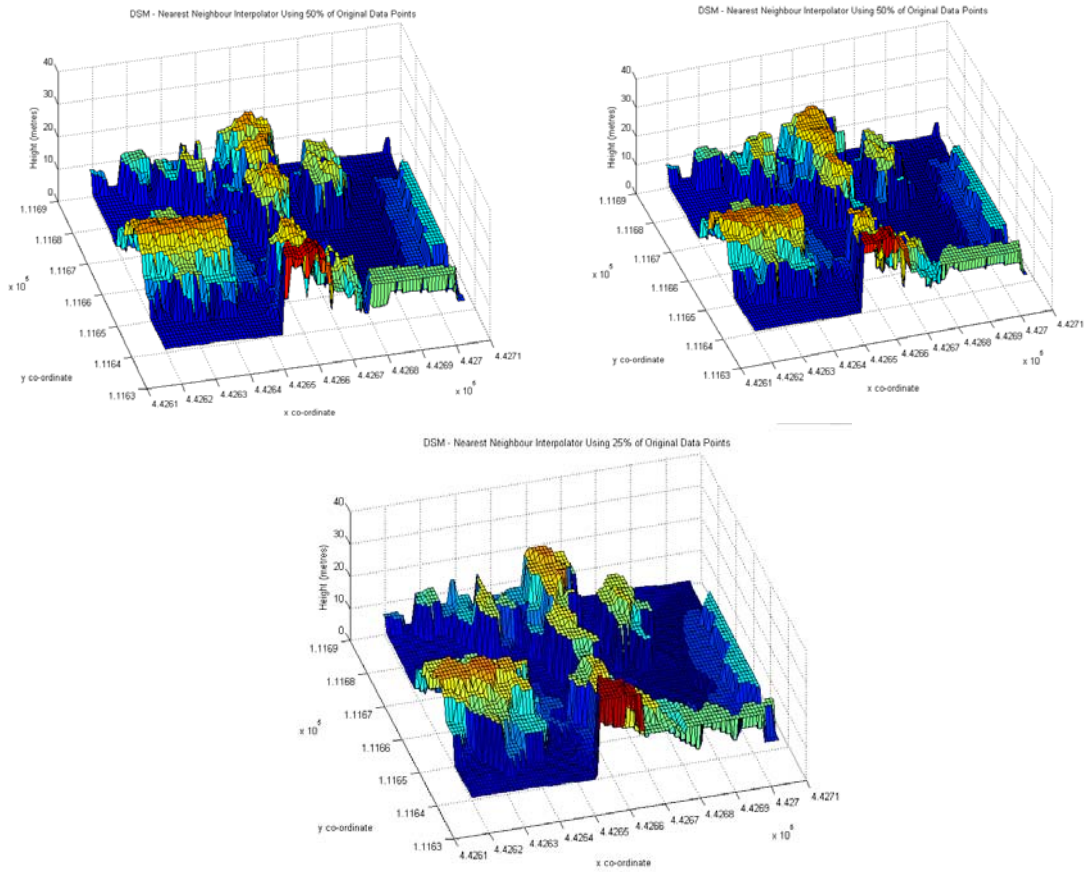


Figure 6 The three DSMs created using the nearest neighbour algorithm

Using the jack-knifing results, overall model accuracy and stability was assessed in terms of absolute errors for each point. The results from the jack-knifing (table 1) were more informative than the results obtained using the cross-validation methodology, as they showed greater difference in the characteristics of the surfaces.

Table 1: Accuracy statistics from the jack-knifing investigation

	Mean error (m)	Maximum(+) error (m)	Maximum (-) Error (m)	Std. dev (m)	RMSE (m)
Linear 95%	0.034066	9.1443	-10.658	2.3505	2.3314
Cubic 95%	0.050776	9.6616	-10.562	2.4045	2.3852
Nearest 95%	0.050776	9.6616	-10.562	2.4045	2.3436
Spline 95%	-0.023134	8.9793	-8.737	2.0028	1.9872
Linear 50%	-0.10566	9.8812	-14.549	2.7361	2.736
Cubic 50%	-0.037267	10.431	-14.549	2.8114	2.8093
Nearest 50%	-0.051381	11.776	-14.749	3.0516	3.0496
Spline 50%	-0.014377	11.561	-15.032	2.7741	2.772
Linear 25%	-0.034536	12.652	-14.295	3.3768	3.3752
Cubic 25%	-0.067997	13.589	-15.019	3.3481	3.3469
Nearest 25%	-0.01017	13.3	-15.32	3.6435	3.6416
Spline 25%	0.052757	15.276	-15.397	3.1782	3.1769

The results show that, in general the nearest neighbour interpolator produces the highest RMSE results, meaning that it repeatedly introduced the highest amount of error across the surfaces. On

the whole the splining method was the most accurate (the mean error and RMSE were lower than other methods). In order to check the reliability of these observations, the test was run three times (each time with a different random selection of points). The repeat tests confirmed these observations.

As with the cross validation method, the pattern of error was mapped to ascertain whether the jack-knifing procedure predicted the same spatial clustering of error as had been previously identified. The plot below (figure 7) shows more clearly the full distribution of the spatial pattern of point error across one of the linear surfaces. In figure 7 the points represent the locations of the interpolated points, whilst the size of the points represents the magnitude of the error at these locations. The correspondence of the larger errors with the ‘above ground’ features (ie. the buildings and vegetation) can be seen clearly. Within these regions, the errors are larger at the edges of the features.

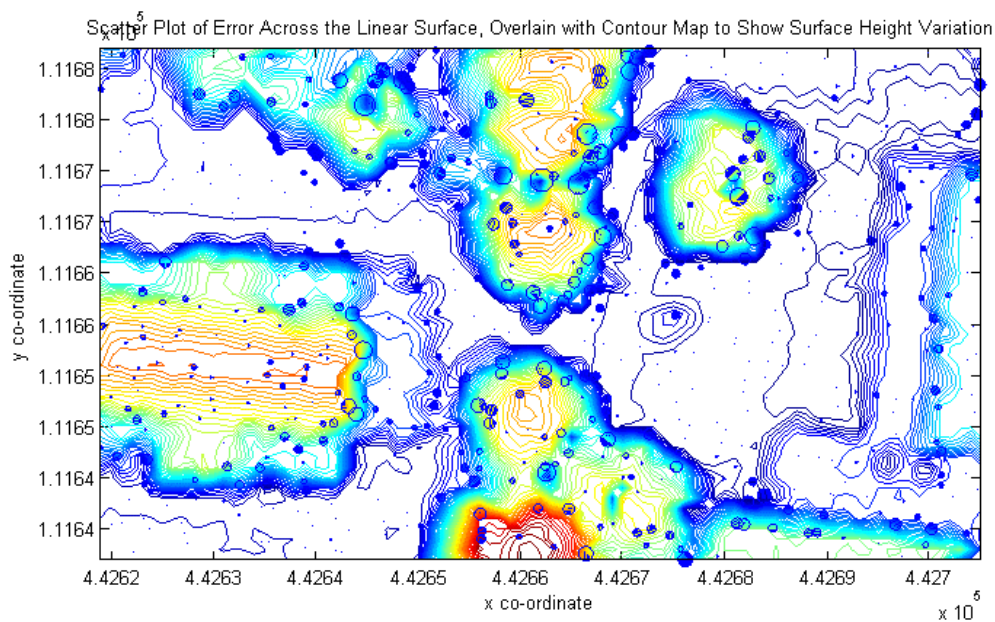


Figure 7 Scatter plot showing the spatial distribution and the magnitude of error across the linear surface

Regression analysis confirmed that the pattern of the error magnitude was correlated with surface roughness.

Alterations in the size of the grid used for interpolation were also found to have an effect on the pattern of errors across the interpolated surface. It was concluded that the optimal sample spacing for minimising error is that which is as close as possible to the original point spacing. Also, the nearest neighbour interpolator produces much greater errors than alternative methods at lower resolution grid spacings. The significance of such variations in the pattern of errors, caused by grid size changes and different interpolation methods, is discussed below.

3 The Significance of the Results: A Worked Example

The best way to demonstrate the importance of quantifying the spatial distribution of error, is with reference to a practical example application of the use of such interpolated surface models.

For the creation of many 3D urban models DSMs are required as an initial input, and the accuracy of the reconstructed features – such as buildings, vegetation and street furniture – is

reliant on the accuracy of the underlying DSM. Understanding spatial variations in this accuracy is of significant importance where the 3D models are to be used for viewshed analysis or for wave propagation modelling. Small differences in surface form may have a considerable effect on the subsequent model. This is highlighted using the example below.

The diagram below (figure 8) shows, in plan, four surface models that were produced in this investigation using the same subset of laser scanner data as above. In these images, surface height is represented by colour – the lighter the pixel the higher the surface – the basic form of a building in the left of the image can be noted, as can the more irregular areas of lighter pixels which represent vegetation. Such surface information is often used to classify images for use in city modelling. The first stage of this modelling process is identifying the features that are above ground, the product of which is a mask of ‘above ground features’. In line with such methods, the pixels in the surfaces shown in figure 8 were classified simply as being either bare earth or ‘above ground’ features (figure 9a) based on the height value of the pixel. In this classification pixels were considered to represent objects above ground if they had a value greater than 8m above datum. In the subsequent masks (figure 9b) the white pixels represent above ground features, and the black pixels are bare earth. It can be seen from these images that there is a difference in the shape of the identified buildings. This difference is also represented here as the percentage of ‘above ground’ pixels per image. There is a variation between the surfaces in both the area of classified above ground features, and in the distribution of these. Where such building masks are to be used as the basic planimetric footprints in the subsequent creation of 3D city models, the accuracy of these is of primary importance. It is clear from this example that any buildings rendered from these different footprints will all have different volumes and form, this has an impact on any proceeding analysis done using the built model.

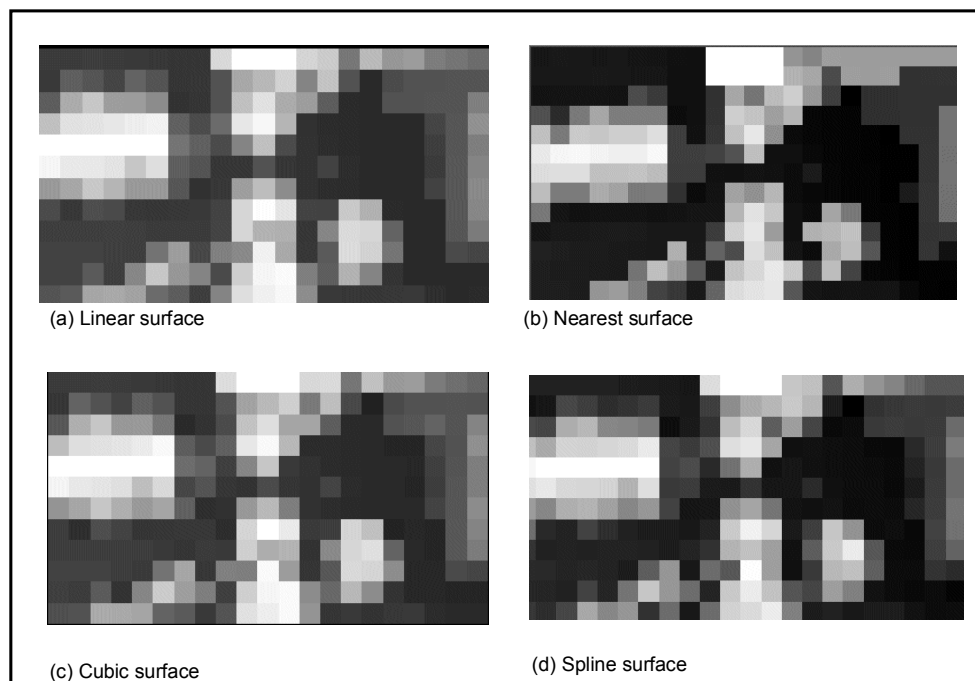


Figure 8: Four surface models in plan view

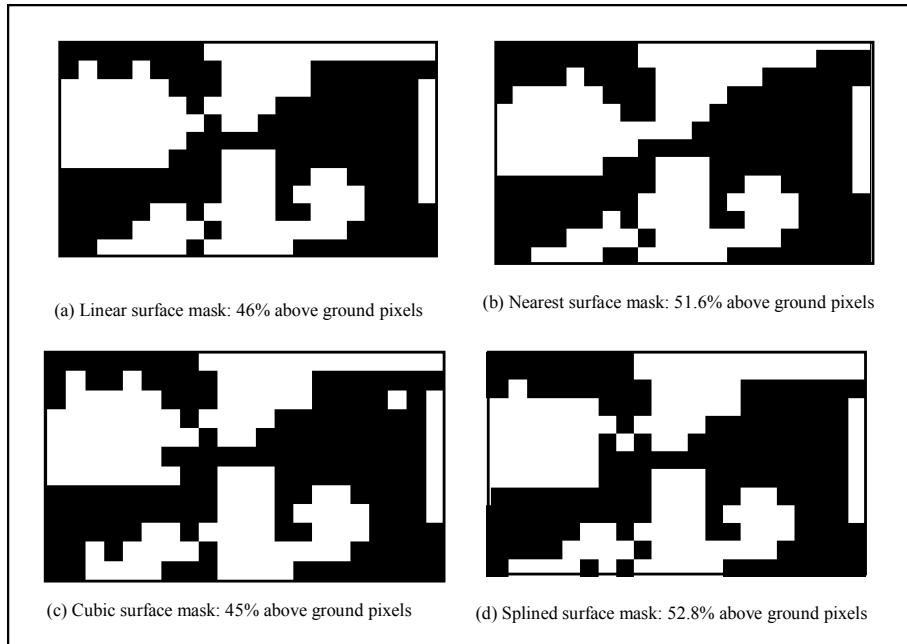


Figure 9a Binary mask for the four surface models. Note the difference in the shape of the features classified as above ground

Figure 9b shows the differences between three of the surface classifications, the pixels where there are contrasting classifications are depicted in RGB, the pixels where the classifications are the same are shown in black (bare earth) and white (above ground). The coloured pixels therefore represent the regions of greatest uncertainty for any urban modeller. An understanding of where the greatest errors occur can therefore aid in the identification of the optimal interpolation algorithm, as well as being used in any accuracy statement for the final model.




Figure 9b The difference in planimetric building form between the nearest neighbour, splined, and cubic surfaces. In this image the coloured pixels represent the areas in which there is a difference between the predictions of the three surfaces.

Conclusion

This paper has shown that there is significant variation between the forms of DSMs created using different interpolation algorithms on laser scanning data. It was found is suggested that the most error is introduced by the nearest neighbour algorithm, and the least error is introduced by the biharmonic splining method. Of the methods investigated, the splining method is thus considered to be the optimal method for minimising error caused during the gridding of points. Despite being shown to produce the least overall error, splining and cubic interpolation were observed to oversmooth building edges and may therefore be unsuitable for some applications. Ultimately the choice of optimal algorithm for a particular application must be decided by the user – understanding spatial variations in accuracy can therefore promote better informed decision making.

It is argued in this paper, that an understanding of the spatial pattern of error in digital surfaces is as important as the provision of the surface model itself. It is suggested that global accuracy statements alone are of little use for many applications, and that as such there is a real requirement for an understanding of the spatial variation in error across the surface. Two reasons are offered in support of these claims:

- For reliable modelling and feature extraction based on DSMs, it is of paramount importance to determine where errors are, particularly in regions of interest such as in the vicinity of buildings and vegetation.
- Secondly, it has been shown in this paper that the error pattern itself may be used to improve the classification in the feature extraction stage. In this situation, an understanding of the pattern of error can actually be used to enhance the model itself, and can certainly be used as an input for subsequent modelling and classification from laser scanner data.

Note to reader: this paper was originally submitted for the GISRUK 2003 conference

References

- Deutsch, C.V and Journel, A.G** (1998) *GSLIB: Geostatistical Software and User's Guide*, 2nd ed. New York, Oxford University Press
- Lloyd, C.D and Atkinson, P.M** (2002) Deriving DSMs from LiDAR data with kriging. *International Journal of Remote Sensing* (23), 12. Pg. 2519 – 2524
- Priyakant, N.K.V., Rao, L.I.M., and Singh, A.N.** (2003) Surface approximation of Point Data using different Interpolation Techniques – A GIS approach.
<http://gisdevelopment.net/technology/survey/techgp0009pf.htm>
- Zhang, J and Goodchild, M** (2002) *Uncertainty in Geographical Information*. Taylor and Francis, London.
- Zinger, S, Nikolova, M., Roux, M, and Maître, H** (2002). 3D Resampling for Airborne Laser Data of Urban Areas. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXIV, 3B, Commission III