The Effect of Changing Grid Size in the Creation of Laser Scanner Digital Surface Models

Smith, S.L[1], Holland, D.A[1], and Longley, P.A[2]

Research and Innovation,
Ordnance Survey, Romsey Road
Southampton, SO16 4GU
Telephone: +44 (0) 2380305009
E-mail: sarah.smith@ordnancesurvey.co.uk

[1] Research & Innovation, Ordnance Survey
[2] Centre for Advance Spatial Analysis, University College London, 1-19 Torrington Place, Gower Street London, WC1E 6BT

Abstract

Raw laser scanning data are captured and supplied as 3D points. These point data require an appropriate interpolation function to be applied to them in order to represent a continuous height surface. Many software packages require that points be interpolated onto a regular grid, in order that the resultant surfaces can be displayed as a Digital Surface Model (DSM). Such interpolation procedures introduce error across the surface, and these errors will be accommodated by any subsequent analysis on the DSM, such as image segmentation or feature extraction. Whilst it is generally understood that interpolation introduces error, there is a paucity of information regarding the characteristics of the errors created using different methodologies. Out limited understanding of the magnitudes and the spatial structure of errors creates uncertainty for users of derived products, such as the Digital Elevation Model (DEM), normalised Digital Surface Models (nDSM), and also uncertainty in processes such as object reconstruction, and viewshed modelling. For object reconstruction and feature extraction in particular, an understanding of the magnitude, pattern and characteristics of the errors introduced by different interpolation methods is vital.

This paper presents the results of a sensitivity analysis of the effects of varying the resolution of resampled locations upon the magnitude of errors. Three grid sizes are investigated, and the magnitudes and spatial pattern of errors within each are identified. Forty five Digital Surface Models (DSMs) were created using five interpolation algorithms and 3 different grid spacings. The errors in each are quantified, and the results presented here. Conclusions regarding optimal grid spacings for different applications are proposed.

1 Introduction

1.1 Laser Scanning

Laser scanning is an active remote sensing technique, in which pulses of laser light are directed towards the ground. The time taken for these pulses to return to the sensor is measured and processed in order to determine the distance between sensor and the object or surface. These data are combined with information about the known sensor position (using GPS and INS), atmospheric conditions, hardware characteristics and other relevant parameters, to generate an XYZ coordinate triplet for ground points. During data capture millions of irregularly spaced points are captured, and these can be interpolated to create a continuous Digital Surface Model (DSM).

Today, laserscanner data and the DSMs derived from them, are used in an increasingly broad range of applications, including urban landscape analysis. Here, accurate DSMs are required for modelling telecommunications, urban microclimates, flooding, and wireless communications. For many of these applications not only must the surface be modelled, but also the above ground features, such as buildings and vegetation, need to be accurately reconstructed in 3 dimensions. This requires that such objects be identified and then extracted from the DSM. Laser scanning has become the data source of choice for feature extraction in recent years, principally because of the accuracy characteristics of both the height and the range measurements. However, the accuracy of the 3-D reconstructions of objects extracted from the DSM is a function not only of the accuracy of the raw data, but also of the interpolated DSMs from which they are extracted.

In the process of feature extraction, objects are extracted from the dataset, which may be either an unconnected point cloud, a connected triangular irregular network (TIN) of original data points, or a grid in which the raw irregular data points have been interpolated onto a regular grid. Object extraction from the point clouds and TINs may entail both a complex methodology (for example, Roggero, 2002; Maas and Vosselman, 1999; Maas 1999; Vosselman, 2000) and intensive data processing. However, as the TIN interpolation retains the raw values at sample locations, it is generally considered that it introduces substantially less error into the surface model than the grided approach, and that this retention of accuracy justifies the greater complexity and processing required for feature extraction using this approach. However, this irregular TIN format is unsuitable for many applications such measurement of planimetric and vertical shifts between overlapping strips of laser data (Behan, 2000), and more commonly, the creation of Digital Surface Models (DSMs) for analysis and visualisation of height data in many commercial software packages. Extraction from the grid is comparatively straightforward compared to similar operations using a TIN or point cloud. While the known disadvantage of extraction from the grid is that the interpolation process introduces error into the surface model, the nature and magnitude of the errors that are introduced has not hitherto been the subject of intensive research. Thus, the relative merits of extraction from point clouds, TINs or grid remains unknown.

There are a number of different ways in which irregularly spaced points may be interpolated onto grids. This paper investigates the effects of using different interpolation methods at differing resolutions upon such errors.

1.2 Spatial Interpolation

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. All interpolation algorithms therefore aim to estimate values at unsampled locations and, in so doing, may change the values of the observed sample points in order to create a plausible continuous representation of the field.

Interpolation, or resampling as it is often termed, has two components; the interpolation method and the resampled locations – which refers to the pixel size or the grid intervals chosen for resampling onto a regular grid.

1.2.1 The Interpolation Method

There are two principal methodological approaches for interpolation **-** deterministic and geostatistical interpolation – and both are explored in this investigation. Deterministic interpolators use mathematical functions to drive the interpolation process, whilst geostatistical approaches use both mathematical and statistical functions for the interpolation.

There are different types of deterministic interpolation, firstly there are those which interpolate on the basis of similarities between neighbouring points - Inverse Distance Weighting is a widely used example (Longley et al, 2001: 106-8). In this method a small moving window is used to identify a set number of points. Such interpolation is used to provide a locally weighted average. The second type of deterministic interpolators is defined by the degree of smoothing in the surfacing. Here *all* the points in a subset are used to derive a polynomial equation which is then used to predict the values at unsampled locations, and the interpolator is said to be global. These latter methods are also known as fitted function techniques. In summary, weighted average methods emphasise local detail, whereas fitted functions emphasise global behaviour. Watson (1992) notes that fitted functions have a tendency to overshoot in situations where a tighter, sharper curve established by weighted averages would suggest more conservative changes in relief.

Geostatistical methods are founded on statistical concepts, which represent the arithmetical relationships, or spatial autocorrelation, between raw data. In addition to the surface prediction, geostatistical methods provide a measure of the accuracy of the prediction (in the form of a variance function). Geostatistical methods, known as Kriging, are more complex than the deterministic approach as they require an understanding of the principles of statistical spatial autocorrelation. However, this complexity is justified for interpolation where the variation of an attribute is very irregular, and the density of observations is such that the simple methods of interpolation may yield unsatisfactory results (Burrough and McDonnell, 1998). Kriging is an interpolation technique which models spatial variation across a surface by looking at general trends/patterns **and** also by modelling the residuals from this (ie. the fluctuations which are not explained by the general trend). In this way Kriging is more robust than other methods, which can tend to oversimplify the spatial pattern of variables. Because of this it is often claimed that Kriging provides the most accurate predictions of all interpolation procedures (Armstrong, 1998). This paper aims to quantify any increases in accuracy offered by Kriging, and to assess whether the higher costs incurred by Kriging interpolation for urban surfaces are justified, or whether a simpler deterministic approach offers a viable alternative.

A variety of interpolators was chosen for investigation in this study, including both deterministic and geostatistical methods. Only local deterministic interpolators were chosen as it was considered that the global approach would be inadequate for modelling the discontinuities in the urban environment. The methods chosen were those most commonly used, namely: nearest neighbour, bilinear, bicubic, biharmonic splining, and Ordinary Kriging (OK). Each of these interpolation methods produces slightly different height values across the surface. Such are the magnitudes of errors involved in the interpolation that Watson (1992) noted that some surface representations may have at best a tenuous link with the raw data values used to create them. This paper investigates the variations in such error in the resultant surfaces derived from different interpolation methods at different grid spacings.

1.2.2 The Grid Spacing

The size of the grid spacing has a strong influence on the errors introduced during interpolation. It has been suggested previously (Behan, 2000) that the optimal spacing should be as close as possible to the original point spacing – ideally this means that each pixel in the raster should contain one and only one raw data point. More points in each cell will promote information loss – as there can be only one value per cell. Similarly, if the pixel size is too small and there are a large number of pixels which contain no laser points, then the redundancy increases, as do storage requirements. This is an important consideration for many users of applications in which analysis is conducted over a large geographical extent.

2 Previous Literature

Whilst the effect of different interpolation methods on the form of the surface has been investigated in the past (eg. Zinger et al, 2002; Morgan and Habib, 2002; Lloyd and Atkinson, 2002; Smith et al, 2003) there has been little research into the effect of changing grid size in the interpolation stage save for that of Behan (2000). Behan (2000) quantified error within models produced from different interpolation algorithms. It was found that the most accurate surfaces were created using grids which had a similar spacing to the original points. Behan's (2000) study looked at global or average error differences between two interpolation methods. This paper aims to extend Behan's (2000) work, by comparing *five* interpolation methods (both deterministic *and* geostatistical), and investigating the magnitudes of error created by each method at *three* different grid spacings.

Rees (2000) investigated the interpolation of gridded DEMs to higher resolutions – whilst this study did not look at the interpolation of irregularly spaced data onto a regular grid (the subject of this paper) many of Rees' (2000) conclusions are nevertheless relevant. Rees (2000) concluded that simple bilinear and bicubic interpolations are adequate for most elevation model requirements in non-urban areas, and whilst he acknowledged a slight improvement in accuracy offered by Kriging he suggested that these improvements were outweighed by the much higher computational demand imposed by the geostatistical approach. This hypothesis will be tested with reference to the interpolation of points over urban surfaces within this investigation.

Kidner (2003) compared a variety of polynomial equations (up to 36 term) to establish the most accurate deterministic algorithm for interpolation of regular grid data onto higher resolution grids. He reported that in all cases the higher order (up to biquintic) deterministic interpolation techniques produced the most accurate surfaces. This conclusion is not surprising given that higher order interpolation techniques will, by definition, attempt to model the surface by reducing the amount of curvature permitted within the reconstruction. In other words, the higher the order of the polynomial the smoother the surface and the less *overall* error across the surface. However, Kidner (2003) did not compare the success of the algorithms across different terrain, and it remains likely that for areas with frequent discontinuities (such as urban surfaces) there will be little difference in the magnitudes of the errors for any deterministic polynomial interpolator. The basis for this hypothesis is that higher order polynomials will attempt to model the discontinuous surface with more complexity than occurs in reality and that this will potentially introduce error of a similar magnitude to that introduced by the much simpler classical interpolation methods such as the bilinear. This hypothesis is investigated in this paper.

3 Methodology

Two adjacent tiles of laser scanner data covering Southampton city centre were provided by the Environment Agency for England and Wales. The data were captured from an airborne sensor, at a point density of ~2m. The complete raw data-set contained some 600,000 points, over an area of approximately 6km^2. To speed up computation and analysis a small areal subset of this dataset was used for this investigation. The subset was chosen to represent the typical surfaces in the urban environment, incorporating buildings, an area of bare earth and some vegetation. The subset contained some 1315 points. The original data and the location of the subset are shown in Figure 1.

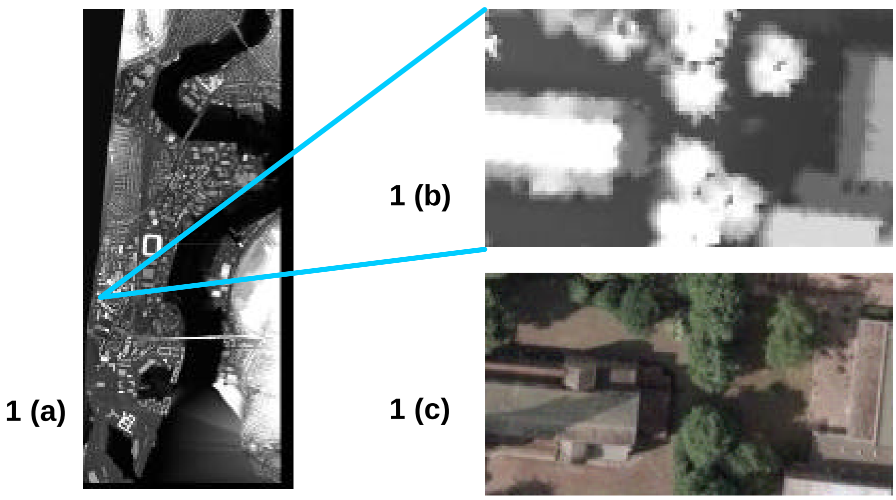

Figure 1(a) The extent of the laser scanning data – raw data supplied by the Environment Agency, Figure 1(b) the subset (88 by 52m), Figure 1(c) Orthorectified aerial photograph of the subset, reproduced with kind permission of Ordnance Survey CC©. All rights reserved.

The church, on the left, and the flat roofed industrial building on the right are visible. The two main clusters of vegetation can also be seen lying between the two buildings.

The 1315 points within the subset were first interpolated onto a regular 1m grid using each of the five chosen methods in turn. Interpolation using the deterministic methods (bilinear, bicubic, biharmonic splining and nearest neighbour) was conducted in MatLab (The MathWorks). The Kriging was performed external to MatLab. A number of authors advocate the use of specialised geostatistical software for Kriging. Burrough and McDonnell (1998), for example, consider that it is more satisfactory to export the data to a specialised geostatistical package. The Kriging for this investigation was performed in GSTAT, and GSLIB. These packages have both been successfully used in previous geostatistical research (Lloyd and Atkinson, 2002). The first stage of Kriging requires a preliminary analysis of the raw data to determine which type of Kriging should be employed. The histogram (Figure2(a)) shows the distribution of the height values used in this investigation. The distribution was strongly negatively skewed. In such cases, it is advised (Burrough and McDonnell, 1998) that either the skew is ignored and ordinary Kriging is conducted as it is relatively stable, or that the raw data should be transformed into a log-normal distribution. Log-normal Kriging entails interpolation of log-normally distributed (rather than Normally distributed) data Deutsch and Journal (1992). The predictions resulting from the Kriging must then be back-transformed. The extreme sensitivity of the errors for back transformation have been previously noted (Deutsch and Journel, 1992), this renders lognormal Kriging very difficult to use in practice.

The study data for this investigation were logged in accordance with the above methodology. However this return a multi-modal rather than a Normal distribution (figure 2(b). Given this and the known problems associated with this technique it was decided that the log-normal Kriging (with a non-normal transformed data set) would be unlikely to yield more accurate results that ordinary Kriging. Thus, the decision was taken to conduct the investigation based on ordinary Kriging.

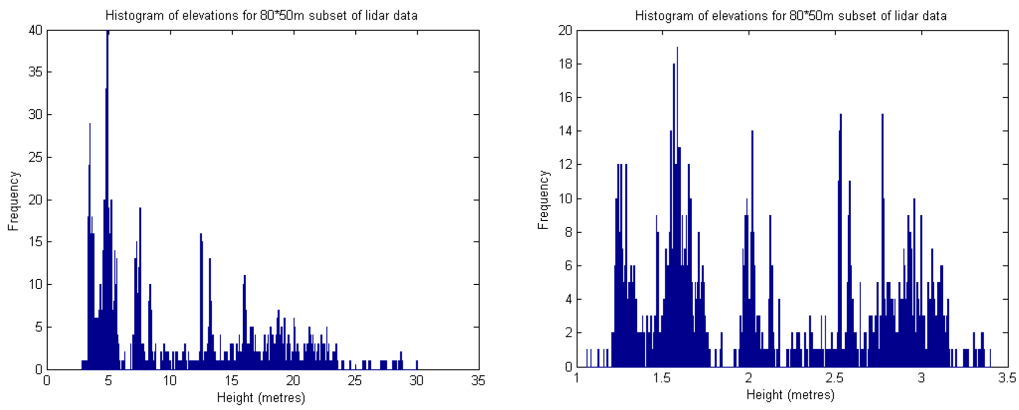

Figure 2(a) Frequency distribution of the raw data – note the strong negative skewing
Figure 2(b) Frequency distribution of the logged data – note the bimodal distribution

Kriging also requires the construction of a semi-variogram – the structure of which provides details which feed into the Kriging itself. The raw variograms were plotted and the closest model fitted to them, in all cases a spherical model was chosen, and the line fitted using either unweighted ordinary least squares (OLS), or weighted least squares (WLS). The results from the fitting of the experimental variogram models were fed into a parameter file to be used in the Kriging of the surface.

In total forty five DSMs were created using different interpolation algorithms and varying grid spacings. Each surface was created using 1m, 2m, and 4m grids. The surfaces were created using a random selection of 95% of the raw data, the remaining 5% of values were retained. The success of the surface reconstruction was then assessed by using the interpolated surfaces to predict the values at the retained locations. This technique is known as jack-knifing (Deutsch and Journel, 1998) and has been used as a test of surface accuracy in a number of papers (Lloyd and Atkinson, 2002). The jack-knifing technique was used to calculate the error in the surfaces, statistics were calculated for these error measurements in order to assess the relative accuracy of the different methods at different grid resolutions. Here it is important to distinguish between error and accuracy. Atkinson and Foody (2002) state that error relates to a single value and is data-based, whereas accuracy relates to the average of data values and is model-based. For the purposes of this investigation the error within each of the surfaces was calculated for each point in the jack-knifing set. The error (e) at each point 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).

$$\mathbf{e\ (x) = Z(x) - Z_i(x)} \qquad (1)$$

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

In addition to absolute error, the range (maximum and minimum of error values), and mean error were calculated. Standard Deviation and Root Mean Square Error (RMSE) were used to predict the accuracy – and provided an expectation of overall error.

Results for all methods and grid sizes were repeated three times, using different randomly selected subsets, to ensure consistency.

4 Results

The resultant surfaces for each of the three grid spacings for the bilinear and the Kriging interpolation are shown in figure 3.

Surfaces created using Bilinear Interpolation

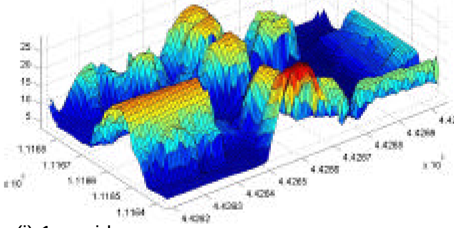

(i) 1m grid

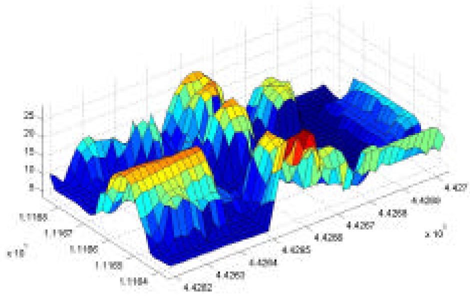

(ii) 2m grid

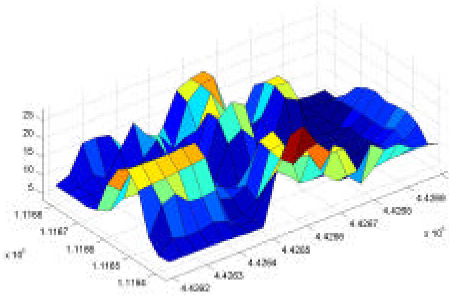

(iii) 4m grid

Figure 3a: Showing the surfaces created from three different grid sizes using a bilinear interpolator

Surfaces created using Ordinary Kriging (OK)

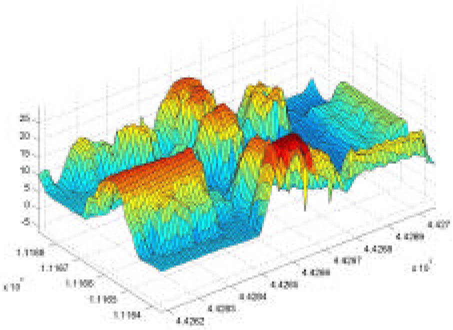

(i) 1m grid

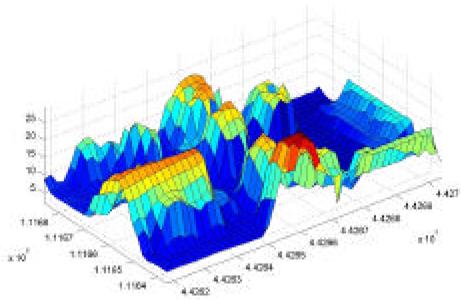

(ii) 2m grid

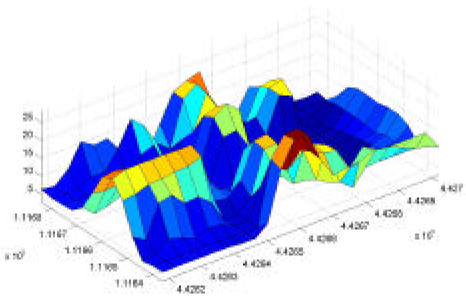

(iii) 4m grid

Figure 3b: Showing the surfaces created from three different grid sizes using ordinary Kriging

As anticipated, Figures 3a and 3b show that there is a loss of detail in surface form at lower resolutions with both the bilinear and the Kriging methodology. The kriged surfaces exhibit smoother edges than the bilinear surface at all resolutions, however there is little real qualitative difference between the surface forms. Indeed, the effect of the decrease in resolution on the accuracy of the modelled surface requires quantification. This investigation was designed to test whether there was a significant increase in accuracy with increasing resolution for all of the interpolation methods. If it is shown that there is no significant increase in error with lower resolutions then it follows that the greater computation and storage implications for higher resolution grids are not be justified. The jack-knifing methodology was also designed to assess the stability of the interpolation methods. It is suggested here that a more stable, or robust method, will exhibit only very slight increase in error at lower resolutions. The results of run 1 are shown below.

Table 1: Results of Run #1

		Max error (m)	**Min error(m)**	**Std Dev (m)**	**RMSE (m)**	**Mean (m)**
KRIGE	1m	6.645	-6.587	1.8829	1.8686	0.011515
	2m	6.522	-6.281	1.8371	1.8231	-0.0002
	4m	7.549	-5.115	1.9705	1.9583	0.10376
BILINEAR	1m	5.5569	-7.1404	1.9754	1.9635	-0.01888
	2m	5.2434	-7.2639	1.9271	1.9167	-0.16308
	4m	4.2895	-7.4366	2.1219	2.2212	-0.73055
BICUBIC	1m	6.1286	-6.3013	2.6447	2.0081	0.023724
	2m	5.8282	-6.5699	1.9565	1.9427	-0.15169
	4m	2.9361	-7.4991	2.0209	2.0515	-0.49475
SPLINE	1m	7.2971	-6.2743	2.1804	2.168	0.13414
	2m	9.2836	-6.1538	2.4014	2.3885	0.15993
	4m	13.085	-22.733	3.9339	3.9139	-0.27943
NEAREST	1m	8.4	-7.14	2.5792	2.5759	0.29453
	2m	8.4	-6.94	2.4145	2.4227	0.36381
	4m	8.33	-14.72	2.6679	2.6467	-0.06254

The quantification of error analysis revealed significant differences between the interpolation methods in terms of the amounts of error they introduced to the DSM. It was found that Ordinary Kriging produced the lowest error per point at all resolutions for all three experiments. Ordinary Kriging was also considered to be one of the most stable of all methods across all three experiments, as it exhibited only a small increase in error (4cm RMSE) with lower resolution grids. The greatest errors were introduced by the nearest neighbour method which produced errors in the region of between 50 and 60cm greater than the kriged surfaces for the same resolution. All three experiments showed that there was very little difference between the errors introduced by the bilinear and the bicubic methods. However, in two out of the three experiments, the bilinear method produced lower error than the bicubic method. This was surprising, given that higher order polynomials generally produce better results than lower order counterparts. Similarly, the biharmonic splining method generated comparatively poor results . For the 1m and 2m grid spacings the splining produced errors of a similar magnitude to the bilinear and bicubic methods. However, for the 4m grid the splining method produced very large errors – predicting much higher results for the surface than reality. This was considered to be a direct result of the highly discontinuous urban surface being modelled here – in such areas the splining interpolator will attempt to overly smooth the surface, overshoot, and attempt to extrapolate at the edge of the data – all contributing to slightly higher errors over the surface.

Almost all methods produced least error when interpolated onto the 2m grid. This finding concurs with the conclusions of Behan (2000), who suggested that the optimal grid sizing should be as close a possible to the original point spacing. The original point spacing of the raw data used in this investigation was slightly more than 2m. It was considered here that the interpolation onto smaller grid spacings (1m resolution) caused the introduction of noise into the surface and this accounts for the higher errors which were evident in the 1m grid. The highest errors were found in the four metre grid, where information was lost owing to there being more than one raw value per pixel. Despite this, it should be noted that the increase of error in the 4m grid was in the region of only 20cm for four of the methods. For many applications this small increase in error would be outweighed by the decrease in file sizes for the lower resolution grids.

Finally, there has recently been a call for not only more information about the magnitude of errors within remotely sensed data and its derived products, but also for the spatial patterns of such error. In this sense, the spatial output from any analysis should be two fold (Atkinson and Foody, 2002): (i) a map of the variable of interest, and (ii) some assessment of the uncertainty of that map. With this in mind, the errors calculated in this investigation were mapped, and the size of the magnitude at each location represented by the size of the point. The error maps for the kriged and the biharmonic splining surfaces are shown in Figure 4.

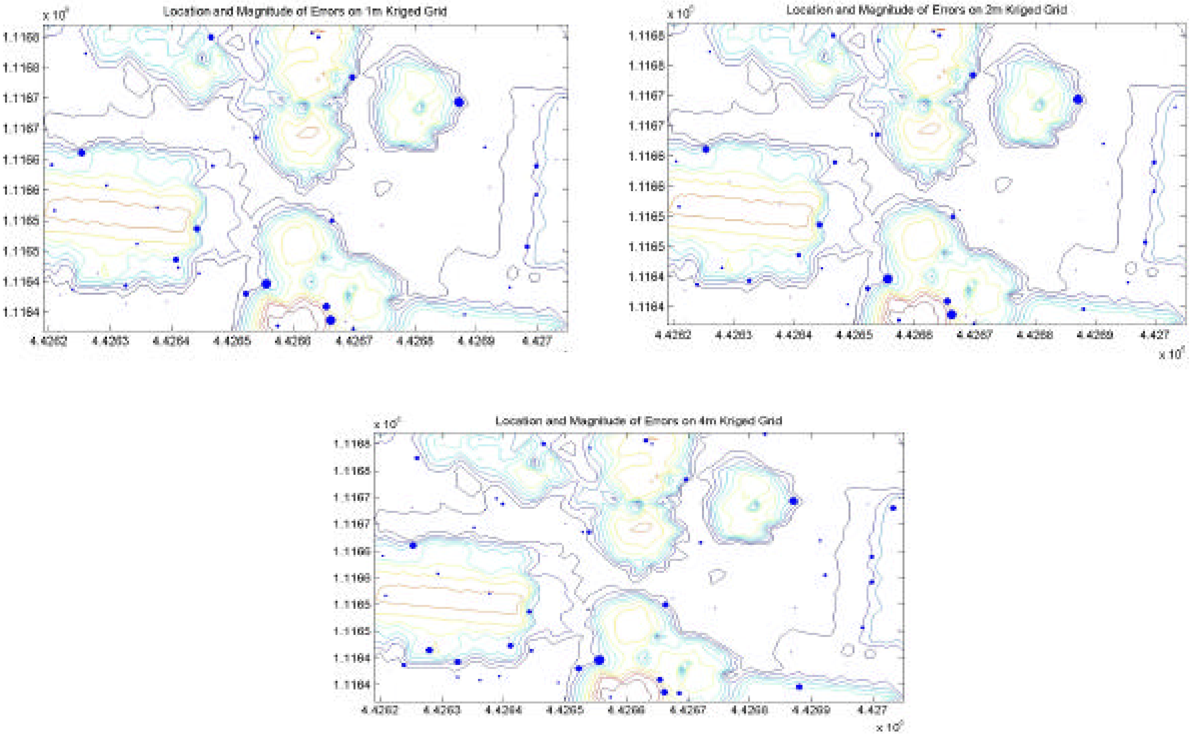

Figure 4(a) Location and Magnitude of Errors Created by the Kriging Interpolator at 3 Different Resolutions. Point size relates to magnitude of error. Contours of the surface are overlain for context

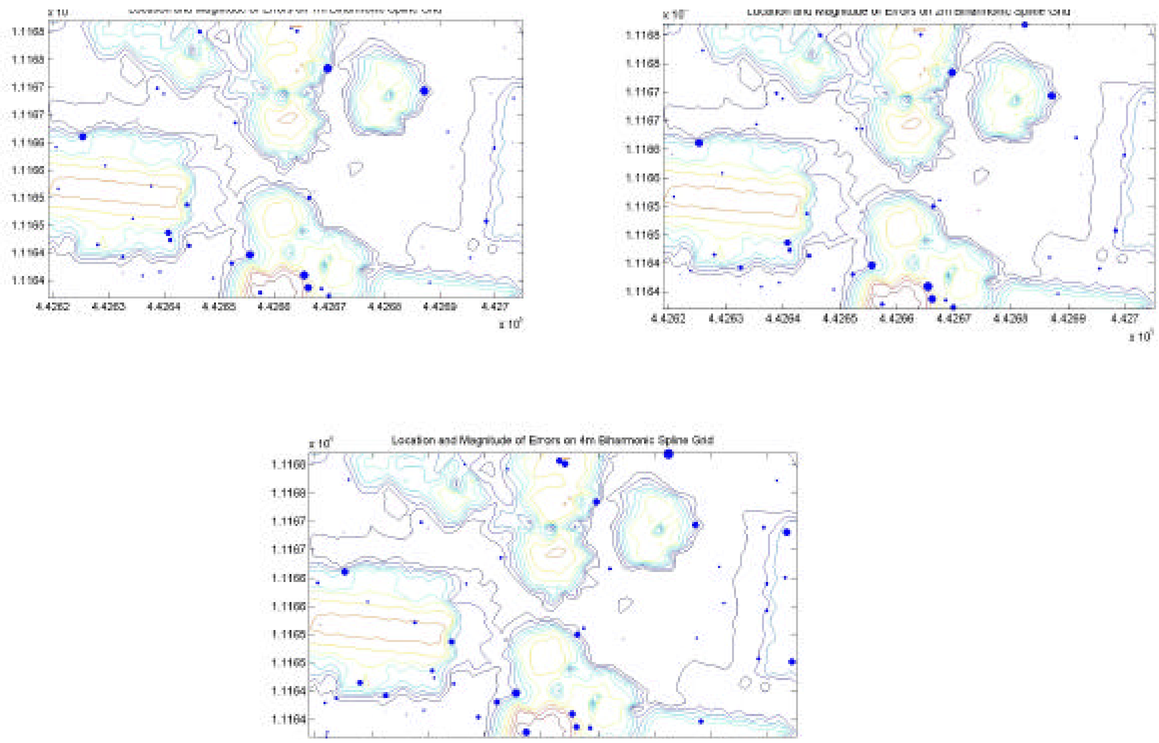

Figure 4(b) Location and Magnitude of Errors Created by the Biharmonic Splining Interpolator at 3 Different Resolutions.

Figures 4(a) and 4(b) clearly show that the pattern of errors are largely coincident with the occurrence of breaklines in the subset, and that these errors are larger in the lowest resolution (4m) grids for both methods. However, it can be noted that in the biharmonic splining 4m grid the pattern of errors is slightly different. Here, the largest errors occur around the maximum heights in the raw dataset, namely over the peaks of trees. This is consistent with the earlier suggestion that the larger errors in the splined surface had been caused by overshoots. Such patterns of error are potentially useful for aiding the choice as to which of the interpolation methods to use for a particular application.

5 Conclusions

This investigation has shown that changes in grid sizes have very different effects on the magnitude of error introduced by different interpolation algorithms. On the basis of the findings of this investigation, the following recommendations are made for the modelling of DSMs in the highly discontinuous mix of artificial and natural structures that characterise urban environment:

- Where accuracy is the most important factor, optimal grid spacing for any interpolation method should be as close as possible to the original point spacing. This corroborates the findings of Behan (2000).

- Where grid spacings are close to original point spacings both Kriging and the bilinear interpolation offer the most accurate surface representations. Given the greater computational demands of Kriging, coupled with the insignificant increases in accuracy (~1cm RMSE), it is recommended that in such cases the bilinear approach is adequate. The fact that in this investigation the bilinear interpolation algorithm produced better results than the bicubic method was surprising, given that the bilinear method is effectively just a simpler version of the bicubic method. It was considered that this may be a function of the subset used in the investigation. Further work is underway to test this.

- Where file size and/or computation times are the most important factor it is suggested that the bilinear interpolator be employed. In terms of computation time, the fastest method is the nearest neighbour algorithm, which works up to twice as fast as the bilinear method. However, it was found that this method produced errors of between 50-150cm greater than the bilinear method in urban areas. As such it is recommended that the bilinear algorithm be employed for the combination of optimal computation time and accuracy offered. It should also be noted that where a low resolution is required (to minimise file storage), Kriging produces the most accurate surfaces. However the increases in accuracy are in the order of 10-15cms, and it is considered here that such small increases in accuracy are outweighed by the computational demands of this method, and it is suggested that a deterministic interpolation be employed.

- When interpolating onto a lower resolution grid it is recommended that biharmonic splining be avoided – owing to the large overshoots created by this method in areas of sparse data. This investigation suggests that a more accurate surface may, for city centre surfaces, be produced from a lower order polynomial (such as a bilinear interpolator). This conclusion is at variance with that of Kidner (2003) who proposed that higher order polynomials will always produce a more accurate surface.

- In terms of identifying which interpolation methods are most appropriate for modelling different elements of the urban environment, it was noted that the Ordinary Kriging surface appeared to represent vegetation more accurately than other methods. This is shown in figure 4, where there a fewer errors over the vegetation regions in the 4m Kriged surface than there are in the splined surface. Ordinary Kriging was also found to create less error on breaklines of buildings than other methods, and as such may be more useful for any subsequent feature extraction using object geometry. This will be investigated further in future work.

- Other future work by the authors will assess the potential for altering the grid spacing across the scene to minimise error created during interpolation. Additional work looking at a two stage interpolation process is also underway. In this approach, the breaklines within the study area are first identified, and these used to constrain the Delaunay triangle creation for the bilinear and bicubic methods. The results of this new method will then be compared to Kriging. Additional work may also look at the effects of using different types of kriging for interpolation in urban areas.

Ultimately the choice of optimal grid size, and interpolation method depend entirely on the application for which the surface is to be used. Studies such as the one presented in this paper are merely designed to aid with this decision making and to allow the user access to more information about both the magnitude and the spatial patterns of errors created by different methods.

Acknowledgements
The authors thank the Ordnance Survey for funding the research presented in this paper, and for providing the aerial photography show in Figure 1. Nick Holden of the Environment Agency is thanked for providing the laser scanning data used in this investigation. Finally, Glen Hart, Principal Consultant of Research and Innovation, is thanked for his constructive comments on an earlier version of this paper.

References

Armstrong, M. (1998) *Basic Linear Geostatistics*. Springer. Berlin.

Atkinson, P. and Foody, G. (2002) Uncertainty in Remote Sensing and GIS: Fundamentals. In Foody, G and Atkinson (eds) (2002) *Uncertainty in Remote Sensing and GIS.* Wiley. Chichester.

Behan, A (2000) On the Matching Accuracy of Rasterised Scanning Laser Altimeter Data. *IAPRS, Vol. XXXIII, Amsterdam*

Burrough, P. and McDonnell, R. (1998) *Principles of Geographical Information Systems*. Oxford University Press. Oxford.

Deutsch, C.V. and Journel, A.G. (1992) *GSLIB: Geostatistical Software and User's Guide*, 2nd edn, New York: Oxford University Press

Deutsch, C.V. and Journel, A.G. (1998) *GSLIB: Geostatistical Software and User's Guide*, 3rd edn, New York: Oxford University Press

Kidner, D. (2003) Higher-order Interpolation of regular grid digital elevation models. *International Journal of Remote Sensing*, vol.24, no.14. 2981-2987

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

Longley, P.A, Goodchild, M.F., Maguire, D., and Rhind, D.W. (2001) Geographical Information Systems and Science. Wiley. Chichester.

Maas, H-G. (1999) Fast Determination of Parametric House Models from Dense Airborne Laserscanner Data. *Proceedings of ISPRS Workshop on Mobile Mapping Technolog*y. Bangkok.

Maas, H-G, and Vosselman, G. (1999) Two algorithms for extracting building models from raw laser altimetry data. *ISPRS Journal of Photogrammetry and Remote Sensing 54*, 153 - 163

Morgan, M and Habib, A (2002) Interpolation of LiDAR Data and Automative Building Extraction. *ACSM-ASPRS 2002 Annual Conference Proceedings*

Rees, W.G. (2000) The accuracy of Digital Elevation Models interpolated to higher resolutions. *International Journal of Remote Sensing*, vol. 21, no 1, 7-20

Roggero, M. (2002) Object segmentation with region growing and principal component analysis. *Proceedings of ISPRS Commission III Symposium, Part 3A* pg. 289 – 294

Smith, S.L, Holland, D.A, and Longley, P.A (2003) Interpreting Interpolation: The Pattern of Interpolation Errors in Digital Surface Models Derived from Laser Scanning Data. *Proceedings of GISRUK 2003.*

Vosselman, G. (2000) Slope Based Filtering of Laser Altimeter Data. *IAPRS, Vol. XXXIII,* Amsterdam.

Watson, D.F. (1992) Contouring. A Guide to the Analysis and Display of Spatial Data. Pergamon. Oxford.

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

Biography of Principal Author

Sarah Smith is currently in the second year of her Ph.D. researching 3D modelling from laser scanning data, with particular focus on the spatial variation of error within each stage of the modelling process. She works as a Research Scientist in Research and Innovation at Ordnance Survey, and is studying at University College London.