The influence of Digital Surface Model choice on visibility-based Mobile Geospatial Applications

Sam Meek[1], James Goulding[2], Gary Priestnall[3]

[1]Horizon Doctoral Training Centre, University of Nottingham Innovation Park, NG7 2TU
Tel. (+44) 0115 8232554
Psxsm6@nottingham.ac.uk

[2]Horizon Digital Economy Institute, University of Nottingham Innovation Park, NG7 2TU
Tel. (+44) 0115 8232557
James.Goulding@nottingham.ac.uk

[3]School of Geography, Sir Clive Granger Building, University of Nottingham, NG7 2RD
Tel. (+44) 0115 9515443
Gary.Priestnall@nottingham.ac.uk

Summary: : The following paper outlines the methodology and preliminary results for an experiment designed to understand the accuracy of visibility models when used in the field by a mobile media consumption app called *Zapp*. Levels of accuracy are determined in relation to points of interest that can be seen from random sites within the University of Nottingham's University Park campus, the study area of this experiment. Testing was carried out on three different surface models derived from 0.5m LiDAR data by visiting physical sites on each surface model with 14 random POI masks being viewed from between 10 and 16 different locations, totalling 190 data points. Each site was ground truthed by determining whether a given POI could be seen by the user be and also be identified by the mobile device.

KEYWORDS: Mobile GIS, Viewsheds, Visibility, Digital Surface Models

1. Introduction and purpose

In this paper we examine the effectiveness of using different digital surface models to underpin mobile geospatial applications. Our experiments show that choice of surface model has important consequences on the efficacy of visibility-based geospatial software. The test-bed for the experiments undertaken was *Zapp*, a mobile geospatial application that allows users to query, from a distance, points of interest (POIs) via use of the device's on-board sensors, Meek and Priestnall, (2011).

Zapp functions by allowing users to aim a crosshair (overlaid on the device's camera preview) at some point within the visible landscape. The application dynamically ascertains the area that the user is targeting via a line-of-sight algorithm, combining device sensor information with height data from an underlying surface model in order to calculate the exact grid cell being selected. Finally the application cross-checks that grid selection with a POI database, and returns corresponding information if a match is found.

Our in-the-field experiments harnessed three implementations of the application; each compiled using a different surface model, and assessed against physical ground truth readings. Base data for these models originated from 0.5m LiDAR (Light Detection And Ranging) of the canopy digital surface model (DSM), a digital terrain model (a version of the DSM with all surface features such as trees and buildings removed) and a POI database. This study makes preliminary investigations into which underlying surface data models best correspond to what can be seen on the ground, and therefore would be most effective in underpinning future iterations of visibility-based mobile applications.

2. Related work

The application used in our experiments is built on a previous iteration of the Zapp software, which was designed to allow for POI *capture* rather than selection. In this latest version, the software again uses the device's on board sensors in combination with Fisher's line-of-sight algorithm Fisher,(1996), to calculate what the app is "looking" at. However, instead of data collection in the field, the application now allows identification of POIs in order to enable relevant media consumption. In this sense Zapp has commonalities with software such as *MediaScape* Stenton, Hull *et al.* (2007), both being centred around the concept of location-triggered media. The main difference is that, whereas in MediaScapes the media is activated when devices enter a pre-defined trigger area, Zapp activates media when the user points the device at an object in the landscape which has media associated with it.

There are several different methods of interacting with the landscape from a mobile device, but one technology that has strong links with Zapp's "point-to-discover" strategy is the *Geowand,* which describes a device that the user physically points at a point of interest in order to select it. Studies have examined different methods of reporting back to the user from a geowand: Robinson, Eslambolchilar *et al.* (2009) investigated haptic feedback which gave the user an idea of the amount of data available to them through the level of vibration; Lei and Coulton (2009) explored a map interface where the user had opportunity to take contextually relevant photos; and Wilson and Pham (2003) tested Geowand control of devices within in a smart home setting.

Zapp differs from prior applications in that, although it also requires the user to physically align the device with a POI, it employs a surface model to determine intervisibility rather than querying a spatial database and feeds back to the user with a light AR interface.

3. Experimental Methodology

The aim of our experiments was to test the effectiveness of three different surface models within a visibility-based mobile application. The models were loaded onto multiple devices to allow simultaneous testing (thus minimizing GPS signal variation), each models being generated from various alterations to the LiDAR data captured at 0.5m resolution in summer 2009 (re-sampled to 2m due to memory pressures on the mobile devices being utilized). The models tested were as follows:

1. **DSM:** Full LiDAR surface model
2. **DSM-Trees:** The LiDAR surface model with trees removed
3. **DSM+Extrusions:** The Full LiDAR surface model (including the buildings and foliage), with POIs additionally extruded 100m above the surface.

The three different models that are illustrated diagrammatically in figure 1 below:

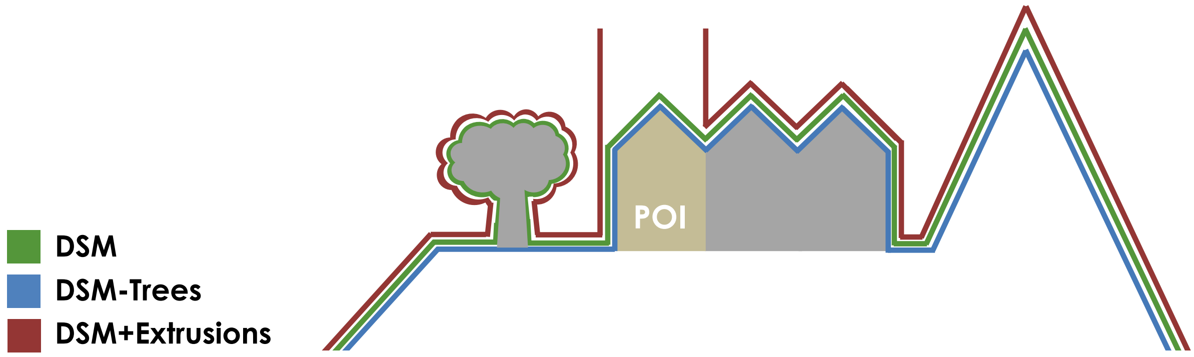

Figure 1. The DSM includes buildings and foliage (green). DSM + POI augments this by extruding the points of interest (red), while DSM – Trees removes foliage (blue).

Surface models were converted into respective rasters for use in line of site algorithms, with the *DSM raster* (which corresponds to the original LiDAR data) acting as a first attempt at modelling the real world as well as a control surface model (see Figure 3a). This also represents the theoretical maximum level of obstruction to visibility as vegetation is modelled as a solid canopy.

The rationale behind the second surface model, the *DSM–Trees* raster (illustrated in Figure 3b), was that lines of trees are semi-permeable and the LiDAR only contains a model of the canopy thus creating barriers to visibility within the model, by removing these barriers we are removing artificial assumptions within the underlying data which were created in the data collection process. The DSM + Extrusions raster created was to ameliorate the problem of foliage walls by extruding the POI buildings above the tree line. Thus, foliage would be maintained but give the sensors on the devices a better chance to "hit" the POI.

3.1 POI selection

A set of 79 possible POIs was created, spread across the University of Nottingham's main campus (see Figure 2). Although our experimental conclusions are necessarily limited to topographies similar to this study area, in order to ensure our results were not biased to a specific set of buildings and features, we generated 14 random subsets of POIs, giving 14 distinct experimental runs from which to test our results.

While the same DSM and DSM-Trees raster could be used across all experiments, a separate *DSM+Extrusions* raster was also to be generated for each experimental run. The generation of this third model type is dependent on the particular subset of POIs being used in a given run. This meant that unlike rasters 1 and 2 (figure 2), a separate version of raster 3 had to be created for each corresponding POI mask (figure 3).

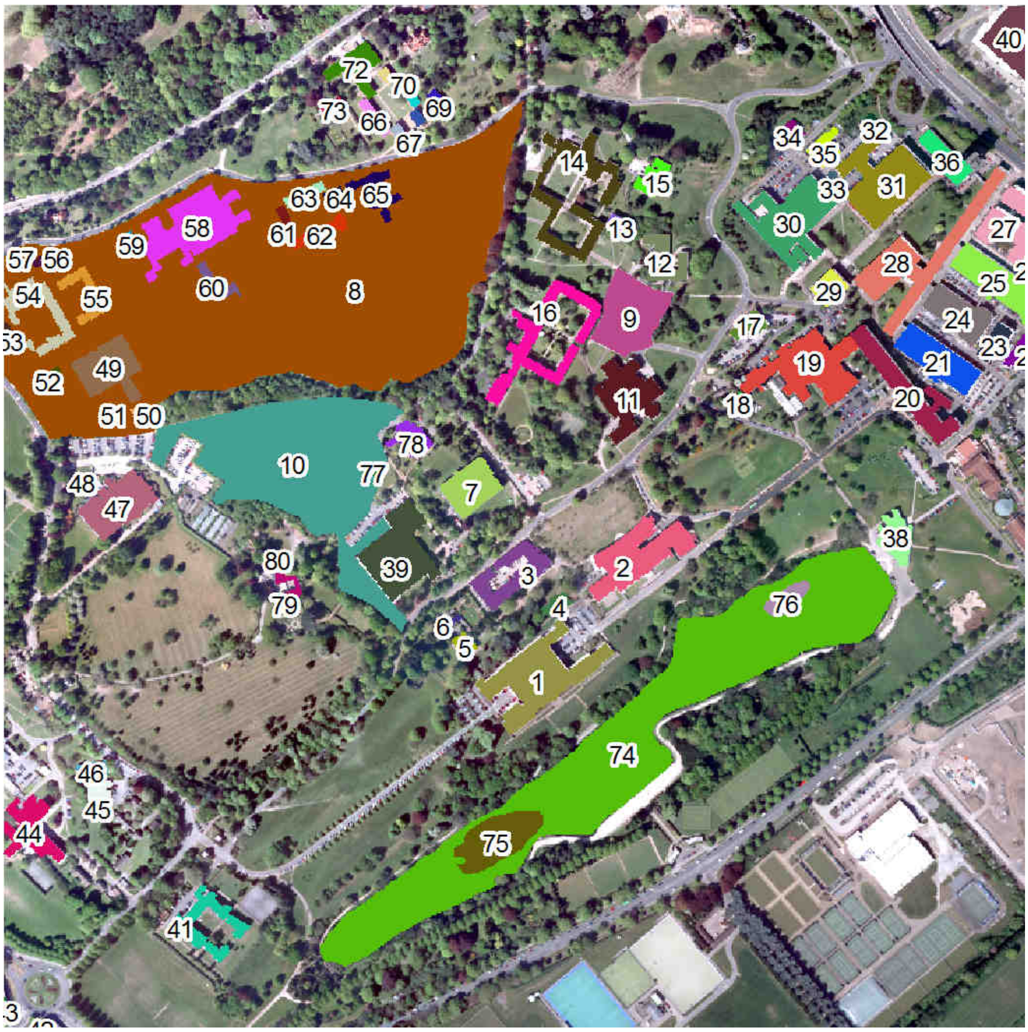

Figure 2. All possible Points of Interest (POIs)

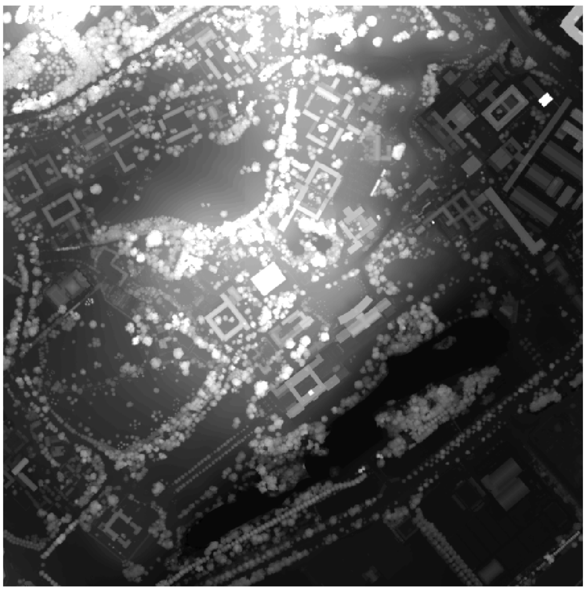

Figure 3a. The DSM Raster - unaltered LiDAR data.

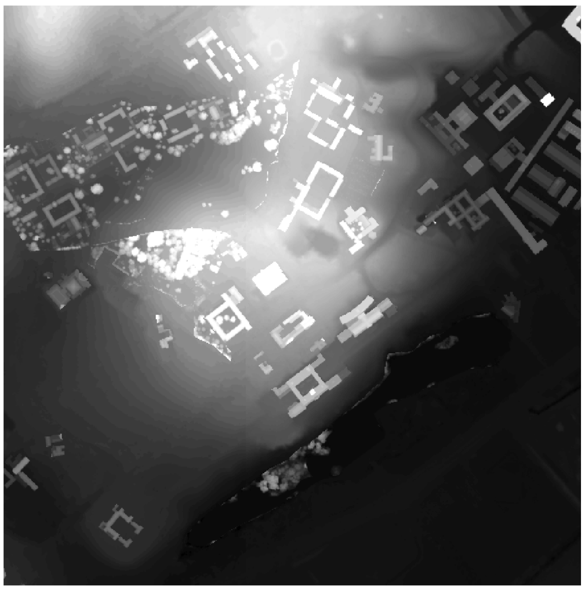

Figure 3b. The DSM - Trees – LiDAR data with trees removed

Figure 3c. An example POI mask from one of the 14 experimental runs.

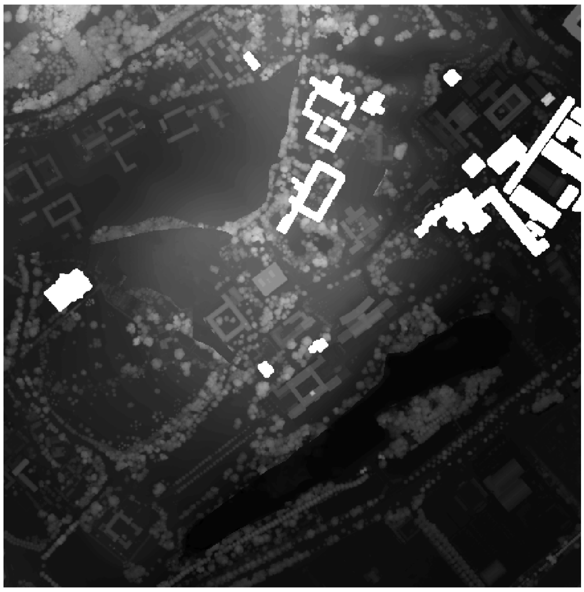

Figure 3d. A DSM+Extrusions raster that extrudes upwards the LiDAR data corresponding to the POIs masked in Figure

3.2 Viewpoint selection

20 distinct viewpoint sites were randomly selected for each experimental run (all viewpoints were constrained to being physically accessible to fall inside of campus, to avoid any issues with private land surrounding the campus and for convenience).

For each of the 14 experimental runs, a team of two people physically visited each of the viewpoint sites with three devices (installed with DSM, DSM-Trees and DSM+Extrusions rasters respectively) attached to a pole. The researcher in control of the devices then generated a ground truth, by conferring with person responsible for the recording of devices as to which POIs (if any) could be seen from that point and therefore what the devices should be able to "see". This resulted in a theoretical collection of 280 multivariate data points to be collected for each of the 3 surface models.

At each viewpoint the devices themselves were pointed in the direction of possible POIs and response results recorded. For each POI in the mask, there were four possible outcomes from the field results:

1. **True negative** – the POI cannot be physically seen and nor can the device see it
2. **False negative** – the POI can be physically seen but the device cannot see it
3. **True positive** – the POI can be seen physically and by the device
4. **False positive** – the POI cannot be seen physically but the device can see it

To determine whether the POI could be seen from the point of view of the user, we employed a rule which said that a POI was deemed visible if it was distinguishable from the landscape, trees or other buildings around it. In other words the POI had to be identifiable as a separate entity in order to be considered as "seen".

The number of sites visited using each POI set can be found in *table 1* and the locations of the sites found in figure 4. At each site the three devices with the different rasters implemented were test by attaching the devices to a pole and attempting to pick out the POIs which were included in the particular set which was being tested at the time.

Table 1. No. sites visited per POI mask

POI Mask	No. Sites Visited
1	15
2	16
3	10
4	14
5	13
6	10
7	16
8	14
9	15
10	10
11	14
12	14
13	14
14	15
Total	190

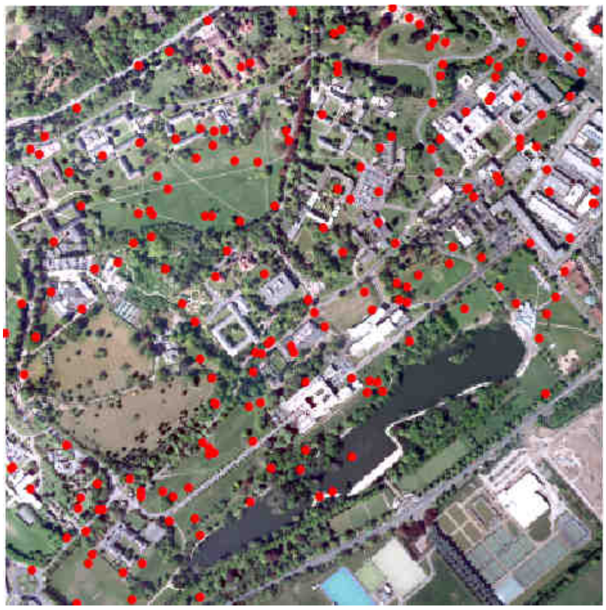

Figure 4. Visited site locations

4. Experimental Results

After carrying out the field testing, the identification of POIs in the field is summarised in Table 2.

Table 2. Experimental Results

Result	Result	DSM	DSM-Trees	DSM+Extrusions
0	True Negative	6589	6508	6559
1	False Negative	126	59	108
2	False Positive	11	91	40
3	True Positive	93	161	112

The table shows that the *DSM-Trees* raster, with no vegetation, produces by far the highest instances of true positives, but at the cost of also producing the most false positives. Without the influence of the tree line, there are fewer barriers to Zapp hitting its target. The raw *DSM* Raster is far too conservative as it suffers with the problems associated with the tree line. When designing the *DSM+Extrusions* raster, it was thought that increasing the building POI size would help account for more salient features by increasing the target size proportional to the footprint of the building, albeit only in the vertical direction. This was a successful approach in a few situations, most notably where the top of the building would poke over the top of a tree line, however these situations proved to be few and far between to make a significant difference for this raster's ability to model the real world.

5. Conclusion

These preliminary results show that none of the rasters provide a perfect fit with the real world – all contain both false positives and false negatives. The *DSM-Trees* raster seems to allow for the identification of POIs far more easily than either of the others. The raw *DSM* Raster is far too conservative and using this raster in visibility-based software likely to cause frustration for the user, with the applications rarely responding to what is being pointed at. The *DSM+Extrusions* raster does have differences to the *DSM-Trees* raster. It provides a more forgiving user experience but at the cost of a lot of false negatives, which is likely to result in even greater frustration for the user. The next steps for our work are to attempt to account for the vagaries of sensor errors on the devices and to deal probabilistically with non-POI barriers such as tree lines in a way which reflects reality on the ground.

6. Acknowledgements

This work was partly funded by the RCUK's Horizon Digital Economy Research Hub grant, EP/G065802/1. Digital surface model was supplied by Infoterra Ltd.

7. References

Fisher, P. (1996). Extending the applicability of viewsheds in landscape planning. *Photogrammetric Engineering & Remote Sensing* 62(11): 1297-1302.

Lei, Z. and P. Coulton (2009). A mobile geo-wand enabling gesture based POI search and user generated directional POI photography. *International Conference on Advances in Computer Entertainment Technology*, New York, NY, USA, ACM.

Meek, S. and G. Priestnall (2011). Towards a field toolkit for in-field reconstruction of 3D surface models. *Geographic Information Science Research UK 2010*, Portsmouth, UK.

Robinson, S., P. Eslambolchilar, et al. (2009). Sweep-Shake: finding digital resources in physical environments. *11th International Conference on Human-Computer Interaction with Mobile Devices and Services*, New York, NY, USA, ACM.

Stenton, S. P., R. Hull, et al. (2007). Mediascapes: Context-Aware Multimedia Experiences. *IEEE Multimedia* 14(3): 98-105.

Wilson, R. W. and H. Pham (2003). Pointing in intelligent environments with the World Cursor. *Interact 2003*, Ohmsha, Japan.

8. Biography

Sam Meek is a PhD student in the Horizon Doctoral Training Centre funded by EPSRC and is based both in Geography and the Learning Sciences Research Institute at the University of Nottingham . His research interests include mobile learning, usability in mobile GIS and mobile computing.

Dr. James Goulding is a Research Fellow at the EPSRC funded Horizon Digital Economy Institute at the University of Nottingham. His background is was in Econometrics before progressing to Computer Science. He is internationally published in the fields of Data Theory, Information Retrieval, Collective Intelligence and Geospatial Science.

Gary Priestnall is an Associate Professor within the School of Geography at the University of Nottingham, whose research interests lie in geographic representation, landscape visualisation, and the application of spatially-aware mobile computing in geography fieldwork.