

CASTLES IN THE CLOUDS: LIDAR FOR HISTORICAL STUDY AND TERRAIN ANALYSIS

Peter L Guth

*Department of Oceanography
United States Naval Academy*

Abstract

Examples of castles and fortresses seen in publicly available LiDAR data by the national mapping agencies in Finland, the Netherlands, Norway, Slovenia, the United Kingdom, and United States demonstrate the value of LiDAR data for understanding military history. The data sets, which include derived grids and the original point clouds with densities from 2–24 points/m², always identify the ground points, generally include LiDAR return intensity values, and sometimes include point classifications that discriminate vegetation and buildings. Grids with 0,5- or one-metre resolution could be created from the highest density point clouds in this study, while grids with one- or two-metre resolution can be created from the lowest density clouds. The digital surface model, which includes everything seen by the sensor, notably ground, buildings, and vegetation, can be created with higher resolution than the bare earth grids. The surface model provides the best visual representation of the castle and its surroundings. Viewed in interactive 3D, the data allows familiarisation with the landscape. Optimal displays depend on the desired scale and the terrain characteristics, but hillshade or shaded reflectance maps, reverse greyscale slope maps, and openness maps all work effectively. Further analysis may include functions such as viewsheds, which enhance the understanding of key terrain.

Introduction

LiDAR provides an unparalleled view of the terrain, and has revolutionised how we create digital topography. LiDAR reveals the natural landscape, vegetation, and man-made features with a random array of millions of points located in three-dimensional space in a point cloud. Despite the random locations of individual points, the point clouds collectively depict topography with unparalleled accuracy and detail. Geographical information systems (GISs) have traditionally used two data models, raster and vector, but LiDAR data has established the point cloud as a distinct third data type in terms of both data storage and analysis. We will consider castles to be fortified structures for defensive military use, in a broad sense, to include fortified residences for nobles, purely military structures, and private creations of wealthy

individuals. As large structures with a critical relationship to the terrain, castles provide a fertile locale for analysis with LiDAR.

Laser scanning of the environment can employ three different technologies: airborne, terrestrial, and mobile. Airborne LiDAR uses a plane or helicopter, and covers very large areas, approaching national coverage in the United States and Western Europe. Terrestrial LiDAR puts the laser on a tripod, and covers a small region with a discrete number of individual scans, which can be stitched into seamless coverage. Mobile LiDAR mounts the scanner on a backpack, vehicle or drone and collects data along a travel path. Fundamental differences among these include the density of points, the differential coverage of the tops and sides of structures, and the trade-offs between point density and area covered. Point density can range from 1–100 points/m² for airborne LiDAR, and 100s to 1000s of points/m² for mobile and terrestrial LiDARs. The airborne LiDAR will cover much larger areas in much less time, but with significantly less detail. Because the active LiDAR sensor must have line of sight to a point to collect data, regions could be blocked from view unless the sensor moves. Airborne LiDAR collects most of its points on the top of man-made structures, whereas mobile and terrestrial collect mostly on the sides. LiDAR systems generally have a global navigation satellite system (GNSS) system, and may also have an inertial navigation unit (INU), so that the collected point clouds can be rapidly registered to ground coordinates.

While the cost of LiDAR sensors has been dropping precipitously, they are still much more expensive than cameras. A low-cost alternative to LiDAR has emerged with structure from motion (SfM), which uses digital photogrammetry with many digital photographs.^{1,2} In contrast to traditional photogrammetry with very precise manipulation of carefully acquired images from expensive mapping cameras, SfM uses redundancy from many more photos from common consumer-grade cameras to derive a model that can be registered to the ground, and one of the possible outputs is a 3D point cloud. With the camera mounted on a small drone, SfM brings 3D mapping to farmers or quarry operators, surveyors, scientists or even hobbyists. SfM has one significant drawback compared to LiDAR, in that photography does not penetrate forest canopy and thus SfM can only produce a digital surface model (DSM).

After collection and integration with the navigation system of the collection aircraft, the raw GIS data consists of a point cloud with 3D coordinates for each return and the intensity of the return, which relates to the characteristics of the reflecting surface. The point cloud supports an interactive 3D display with colouring either by return intensity or point elevation (Figures 1 and 2). Additional information on collection information, such as scan angle or return number, will not be useful for most users. Principles of remote sensing apply to the LiDAR image interpretation, with shape and texture revealing vegetation and man-made features. In addition to the 3D point cloud displays, which show up much better interactively on screen compared to the flat page view, the GIS program can draw 2D slices through the point cloud, which allow easier measurement of features (Figure 3).




Figure 1. 3D view of the point cloud for Bled castle in Slovenia, showing three different colouring methods.




Figure 2. Point cloud display from LiDAR survey of Mulderslot, the Netherlands, showing return intensity and elevation in nearly overhead 3D views.




Figure 3. 2D slices through the point cloud of Bergen, Norway, which shows the easily identifiable planar roof geometries and tree canopies with colour from the point cloud classification.

The advent of airborne LiDAR allows creation of more detailed gridded digital elevation models (DEMs, used here in the generic sense of any terrain grid) than were previously available, at higher resolutions and with a corresponding increase in data volumes for the resulting data sets. Rather than earlier grid resolutions of 10 to 100 m, one metre is the most common grid resolution for LiDAR DEMs, and 0,25 or 0,5 metre is also possible. The size of the grid goes up with the square of the change in grid spacing: a one-metre grid has 100 times more points than a 10-metre grid, and 10 000 times more points than a 100-metre grid. LiDAR data is collected from an aircraft with a scanning laser that collects data with sampling rates in the hundreds of kHz, with several hundred thousand pulses emitted per second. For areas covered by vegetation or power lines, each pulse can have multiple returns from within the vegetation canopy as well as a possible ground return.

Previous LiDAR analysis of castles

Airborne LiDAR provides a more complementary view of castles than mobile or terrestrial LiDAR, with lower resolution but wider context because the larger areal coverage places the castle within its regional framework. Historic England produced a handbook on using airborne LiDAR for archaeological survey,³ which provides a good introduction, but used LiDAR with much lower resolution than currently available. Their detailed handbook on using laser scanning, with terrestrial and mobile LiDAR, has 16 case studies, including Belsay Castle.⁴ Terrestrial LiDAR and SfM have been used to map Chillon Castle in Switzerland, which showed that both could produce detailed and accurate 3D models.⁵ A backpack-mounted LiDAR produced a point cloud of Newark Castle in 15 minutes.⁶ These point clouds, with extremely large point densities, can create views with photographic realism, and help create architectural quality drawings. These use many of the same techniques and software as the building inventory modelling (BIM) commercial market, and result from a dedicated survey covering a restricted area. In contrast, the study reported here focused on the GIS or terrain analysis approach, using airborne LiDAR, which is becoming routinely available for large portions of many countries as a standard product of the National Mapping Program. Off-the-shelf data covers any desired castle in any country.

LiDAR data examined

This study collected samples of airborne LiDAR data showing castles from seven countries, which had freely available data from the national mapping agency in May 2017 (Table 1). Appendix 1 has notes on the data sets in Table 1 and their processing, which should be of interest to readers who are more technically inclined. Initially, the choice of castles was only meant to appeal to military geoscientists, but the focus on military castles and fortresses turns out to highlight many of the issues that affect point clouds and the grids created from them, which may not be as obvious in other locations. The sample illustrates the range of decisions made by the mapping agencies in the production of LiDAR, and the kinds of data that can be routinely collected and processed.

Table 1. LiDAR data sets examined for castle depiction.

Castle	Country	Cloud Density (pts/m ²)	Grids	Projection	Grid Spacing (m)	Classification Categories	Survey Year
Bergenhus Fortress	Norway	23.4	DSM, DTM	UTM	1	Ground, vegetation, buildings	2016
Mulderslot	Netherlands	19.7	DSM, DTM	Amersfoort	0.5	Ground, buildings	2015
Kronborg (Helsingor)	Denmark	9.5	DSM, DTM	UTM	0.4	Ground, vegetation, buildings	2014-2015

Castle	Country	Cloud Density (pts/m ²)	Grids	Projection	Grid Spacing (m)	Classification Categories	Survey Year
York	UK	5	DSM, DTM	UK OS	1	Ground, vegetation	2008-2015
Bled	Slovenia	4.5	DTM	Slovenia 96	1	Ground, vegetation, buildings	2014-2015
Turku	Finland	2.6	DTM	UTM	2	Ground, vegetation	2012 (?)
Fort Monroe	VA, US	2.2	DTM	UTM	1	Ground	2013

LIDAR DEMs

Additional processing of the LiDAR point cloud greatly enhances its value. Anomalous isolated points, both low and high from things like bird strikes, can be filtered out. Point classification within the point cloud, using the geometry of the surrounding points, such as flat roofs and the multiple stacked returns from vegetation (Figure 3), can identify features such as power lines, buildings and trees. In addition, to generate some of the derived grids, the classification algorithms can identify returns from the ground surface. Evans and Hudak describe an algorithm to identify vegetation and the ground, which requires parameters for scale and a curvature threshold, which must be tuned to the characteristics of both the LiDAR data and the landscape. This precludes fully automatic processing.⁷ In general, the point classification algorithms will be a black box applied to the data during production, often with proprietary software, but users should ensure the producer adds the desired level of classification. All seven countries listed in Table 1 classify the ground points, because those are required for DTM generation to be discussed below. The United States generally includes no other classification; other countries also classify buildings, vegetation, or both. In the case of LiDAR users in the military, or other users talking to military members not familiar with GIS, great care must be taken when talking about ‘classified’ LiDAR data. The GIS meaning of classified, ‘assigned to classes or categories’, does not mean the same thing as the military meaning of ‘secret, or restricted to authorized users’. Classified ‘secret’ LiDAR data sets doubtless exist, and they might also have GIS classifications, but that use of ‘classification’ applies only to a subset of military users who can access that data.

Four different grids can be created from the LiDAR point cloud (Table 2). The generic term ‘digital elevation model’ (DEM) refers to any terrain grid. The digital surface model (DSM) was produced by only about half of the national mapping agency in the sample data (Table 1), and is actually the easiest to generate from the point cloud. The non-vegetated surface (NVS) has not been formally defined previously, but is also simple to create. Use of the NVS presupposes that actual building shapes will prove more useful or less disconcerting than the unnatural interpolated surface

after building point removal, but will be a subjective decision by the user of the grid. The digital terrain model (DTM) is the most common grid distributed by mapping agencies, available from all the mapping agencies surveyed, but has obvious visual artefacts from buildings and roads. With the DSM and DTM, users can easily compute the canopy height model (CHM) with a simple map algebra grid subtraction. The CHM would be virtually identical if the NVS were subtracted from the DSM, varying only with the definition of canopy and whether it should reflect just the thickness above buildings (using the NVS), or the height above a hypothetical ground surface in the absence of buildings (using the DTM). It would not be a major issue in most areas.

Table 2. Grids created from LiDAR point clouds. See examples in Figure 4.

Grid	Synonyms	Comments
DSM – digital surface model	First return, ceiling	Easiest to create, from the highest point in each cell. Only need to identify and remove obvious anomalies like bird strikes.
NVS – non-vegetated surface	Last return, floor	Lowest point in the grid cell, but tree stumps can still be an issue in dense vegetation with low pulse density. Buildings affect most military users of the grids, so leaving them in may create a more useful data set, which shows buildings as obstacles for line of sight and mobility.
DTM – digital terrain model	Base or bare earth, last return, floor	Requires identification of ground points, buildings, and power lines. Challenged by tree stumps and large buildings. Building ghosts look unrealistic. TIN (triangulated irregular network) gridding can create anomalies at data voids.
CHM – canopy height model	HAG (height above ground), nDSM (normalised DSM)	Difference between DSM and DTM, on DSM and NVS. Has pits in the canopy where tree stumps create spikes in the DTM.

Maps of Fort Monroe, Virginia, show one-metre grids created in a GIS program from a LiDAR survey with about 2.3 points/m² (Figure 4). The LiDAR intensity map and United States geological Survey (USGS) aerial photo help with interpretation of the grids. The DSM, NVS and DTM from the ground points all use a drop in the bucket algorithm. Each LiDAR return is assigned to a grid cell and compared to any existing value. The DSM keeps only the highest point, and the DTM the lowest. Small holes with no data have been interpolated from the neighbouring points. The white points show cells with no LiDAR returns ('holes' or voids not automatically filled), which in this case come from still water with specular reflections as is clear from the ortho image. These LiDAR pulses will only return to the sensor when it looks directly down, which is rare with the scanning patterns employed. The DTM from ground points uses only those points assigned a 'ground' classification, and ignores all other points. This point cloud only classifies them as 'other', but the imagery shows them to be trees and buildings. Most of the trees do not produce holes, because with multiple returns in every grid cell, some pulses hit the ground, but the buildings show

clear rectangular holes. One method for dealing with these voids uses a triangulated irregular network (TIN) to interpolate the DTM. This works reasonably well inside the fort, but the large triangles in the moat show up clearly, because the first points encountered by the LiDAR beams are at different elevations on each side of the moat.

The USGS DTM, the only grid available from the national mapping agency, has had significant clean-up compared to the DTM from the classified point cloud: the moat has been filled in, one of the three bridges over the moat remains but two were removed, and gaps where no points were identified on the ground have also been filled in. The massive masonry walls remain in the DTM, because the classification algorithm considered them ground points. Tweaking the point classification parameters to detect large buildings leads to other compromises. The metadata shows that the contractor making the DTM used a proprietary classification system and algorithms to create the hydro-flattened grid, which allows automated algorithms to create a realistic stream network




Figure 4. The top two rows show the LiDAR return intensities and 1-metre grids created in GIS software from the LiDAR point cloud covering Fort Monroe, Virginia. The bottom row shows the USGS standard mapping products, an aerial photograph and DTM.

The NVS has some remnants of tree stumps, because of the one-metre grid size and the relatively low sampling density. Using a larger grid size for the NVS surface removes the tree stumps, because increasing the number of ground pulses in each grid cell increases the likelihood that one pulse hits the ground (Figure 4). The two-metre NVS retains most of the detail in the buildings. If the DTM or NVS is supposed to capture the natural landscape, perhaps it should not capture every fallen log or small rock, while taking only the lowest points on a coarser grid may produce a grid that matches expectations more closely by ignoring the finest scale features and introducing a degree of smoothing.




Figure 5. NVS grids created with 1- and 2-metre resolution grids using the lowest point in the grid cell. The spikes in the NVS would correspond with pits in the CHM.

Many algorithms used to create the grids create a TIN from the LiDAR points and then interpolate the grid nodes from the triangles. When there are insufficient points, the triangles become large and the resulting planar surfaces appear prominently on the grids, especially when displayed with shaded reflectance (Figure 6). The algorithm can automatically skip interpolating in large triangles, as may have happened with the proprietary algorithm at Fort Monroe (Figure 4), where the producers filled in the moat with a uniform elevation. A simpler, more automated solution just leaves those areas blank, and requires users to know that water areas are frequently voids in LiDAR-produced grids. TINs are also a likely solution for filling the voids left from buildings, which can be recognised in most DTMs as large planar surfaces but which tend to be flatter and not as obvious as the water voids.



Figure 6. TIN anomalies (circled) in the DSM grid created for Denmark's Kronborg castle.

LiDAR DEM display

LiDAR DEMs provide the most value in open terrain, with some vegetation and man-made structures (Figure 7), which are very common around castles. For situational awareness while navigating over the terrain or performing terrain analysis of a castle, the DSM provides the best display. When using the DEM alone, colour can depict elevation combined with hillshading or shaded reflectance to show local relief. If colour overlays for additional layers will be on the map, making the DEM greyscale could improve the user interaction. In addition to shaded reflectance or hillshade maps, openness maps and a reverse slope map could also provide useful displays; simple elevation shading provides little value. The openness map, originally designed for a larger region, averages the angle to the nearest blocking terrain in eight directions to highlight ridges and valleys.⁸ Openness takes time to create for the larger region sizes originally proposed, but the display works very well for LiDAR DSMs with a small computational region and very fast response. A normal slope map would




Figure 7. Display options for the castle in Turku, Finland, which is the highest local elevation in this region.

show steep slopes in white and flat areas in black, but by reversing the scale, the steep slopes in black represent obstacles and the light areas highlight easy mobility. In the map below, the trees are the darkest black, while the castle is slightly less steep.

Part of the power of the GIS comes from its ability to combine layers and improve situational awareness. Figure 8 shows the classification from the LiDAR point cloud overlaid on a reverse greyscale slope map. Users must understand two things about such displays. Firstly, the standard for point classification allows three vegetation categories (low, medium and high), but leaves the class boundaries, and whether to use all three, to the person doing the classification. In this case, no metadata, which might explain the classification decision, accompanied the data, and the national web portal does not make it easy to establish whether metadata exists there. The user would have the choice of just regarding the vegetation categories as relative, or draw slices through the data to see where the breaks occurred (Figure 3 shows this, although from a different country, which might have chosen different boundaries). Secondly, when multiple GIS layers are combined with transparency in a single map, the resulting colour comes from multiple layers. Each of the four categories in Figure 8 has a range of colours: the buildings are all variations of light blue, the high vegetation is darker green, but much of it appears almost black because the trees are so steep and thus black in the slope layer dominates the merged image. The lack of uniform colour for the categories does not hamper interpretation.

Downloaded to a smartphone with its integrated GPS, this map allows users to follow their progress on the ground, and anticipate the next terrain to be encountered. It provides a good map while visiting castles and investigating their relationship with the local terrain. If maps are downloaded before going to the field, the work can be done without a network connection. The map is also simplified from an aerial




Figure 8. DSM shading combined with point cloud classification colouring for Egeskov castle. The colouring includes buildings, two height classes of vegetation, and 'bare earth'. The castle sits in a small lake, with the TIN facets clearly visible.

photograph, potentially requiring less training from the user to understand what is shown. Software can offer a choice of the imagery of the LiDAR derived layers, and mobile phones can now perform computations like viewshed analysis and 3D terrain visualisation.

Discussion

LiDAR topography, both as the collected point clouds and as derived grids, offers a tremendous tool for understanding castles and their location. The data can be displayed on personal computers and smartphones. Neither the storage requirements nor the processing power required limits the usability for moderate regions covering 100s of km², ideal for most historians and the study of historical battles involving

castles. The grid for a one-metre DEM covering one km² requires 4 MB of computer storage, as does a four-band (RGB and NIR) air photo. For 100 km² (a region 10 km x 10 km), an image, DSM, and DTM would require about 1.2 GB of storage, which fits easily on modern smartphones, which start with base storage of 64 GB. Using just base maps, such as that in Figure 8, 100 km² requires only 400 MB.

Acquiring data can be as simple as a download from the national mapping agency. For all of the countries listed in Table 1, both the LiDAR point cloud and the DTM are standard products currently covering a large fraction of the country and with ambitious plans for improved coverage. If the DSM represents the preferred grid for analysis, for three of the seven countries, GIS software would have to be used to create the grid. Freeware, open source, and commercial programs can all perform this step. The data will have to be collected, prepared for field users, and the users will have to be trained. Base map tiles, using a display like that shown in Figure 7, can be served via a local intranet or downloaded from a personal computer.

Data should be in the UTM format, and if required, converted to UTM long before getting to end users. Depending on the software used, some national projections may not be supported, and conversion can be complicated since it might require a datum shift as well as reprojection. In the sample in Table 1, about half the national mapping agencies used UTM coordinates on the WGS84 data or a functional equivalent such as NAD83 or ETRS89. The projection information must also be in the LAS files, which is not always the case for some mapping agencies. Inserting it should ideally be done when the files are created. If not, the projection should be inserted as far as possible upstream in the data distribution process, since individual end users should not have to deal with the issues.

The classification of the point clouds should include the bare earth, vegetation and buildings. Inclusion of power lines, perhaps not required for historical analysis but important for other terrain analysis, will require the use of the LAS 1.4 format.⁹ This version of the LAS format has not been used extensively to date (none of the national data sets in Table 1 uses it), but the importance of obstructions for helicopters means the LiDAR data should include it. Only three of the seven national mapping agencies in Table 1 currently include both vegetation and buildings, and the United States includes neither.

Users must learn about the characteristics of LiDAR and point classification, so they can interpret the display. They need to know how water and buildings appear in the data, because the examples shown here demonstrate the lack of consensus on how that should work (Figure 9). They will also have to learn about horizontal and vertical datums. The high resolution of LiDAR reveals discrepancies in tidal datums and ellipsoidal heights, and military coastal operations will have charts which do not use the terrestrial datum. The United States will adopt new data in 2022, requiring users to understand the data shifts, and will have four different reference frames for the four tectonic plates. As the tectonic plates of the earth continue to move and higher resolution data sets proliferate, earth coordinates will become truly 4D with a time component.




Figure 9. Water treatment in the LiDAR-derived DEMs for three castles. In addition to the TIN artefacts at Elsinore pointed out in Figure 6, the NE corner of the map shows waves captured by the sensor.

Within the range of point density represented in this comparison, 2–23 points/m², the differences in resolution have subtle effect. All can create DSMs with one-metre resolution, and the denser clouds could support 0,5-metre resolution. Because the DTM or NVS requires returns that penetrated the vegetation canopy, they require a higher density and one- to two-metre resolution is probably most appropriate for these data sets. The higher resolutions will be most obvious in slices through the data (Figure 3), with more points in the canopy.

The NVS was not been produced by any mapping agency in the sample, but might reflect obstacles better for mobility and intervisibility. The buildings will block people and line of sight, so leaving them in the DEM might serve current users better. Vegetation will also generally block line of sight, so intervisibility could use the stored DSM (Figure 10), or add the CHM to the NVS on the fly. A user trying to understand the historical context of a castle must understand which buildings were present at the time of interest, as well as determine the status of vegetation because over time, demand for construction and firewood has kept many areas with scant tree cover.




Figure 10. Cumulative viewshed in green from the four corner towers of the Harlech Castle in Wales, on the current DSM.

Conclusion

Airborne LiDAR data offers enhanced capabilities for understanding the role of castles on the landscape. Its display on handheld devices could enhance scientific study and general visits. For onscreen visualisation, 3D display of the point cloud and slices through it provide the greatest aid to interpretation. For base maps, either in the lab or in the field, or for reports, the DSM provides the most realistic and easiest view. Because only about half of the national mapping agencies providing LiDAR data currently provide DSMs, users who want them will have to download the point cloud and create their own, which is a relatively painless procedure. While the current DSM provides a great deal of context for the study of castles, fully understanding a castle's context will require knowledge of both the buildings present in the past and the vegetation at the time. Both of these can be created as layers in GIS, further improving the analysis.

LiDAR topography will also prove useful in urban warfare, where apartment buildings supply the equivalent of castles to small units. These battlefields might cover only a few blocks, and frequent data collection via drones or aircraft could provide up-to-the-minute information about destroyed buildings and current obstacles. The battles in Fallujah, Iraq in 2004 were the bloodiest in the US war effort in Iraq, and the largest urban conflicts since the Vietnam war. The efforts to retake Mosul, Iraq and Raqqa, Syria, in 2016–2017 featured urban destruction comparable to that in cities such as Stalingrad in World War 2, and current knowledge of the terrain just one or two blocks away would greatly benefit small unit commanders.

Acknowledgments

The views expressed are those of the author and do not reflect the official policy or position of the US Naval Academy, the Department of Defense, or the US Government. I thank the two reviewers, who correctly noted that the original version of the article spliced together two distinct topics; their suggestions greatly improved the article.

Appendix 1. Technical notes on the LiDAR data sets:

1. Point cloud densities are approximate and depend on choices to deal with water voids, irregular data set outlines, flight line overlap and multiple returns.
2. All countries supply DTMs, and some also supply DSMs for free download.
3. Downloads for the point clouds use the LAS format, with either regular ZIP or LAZ compression;¹⁰ none uses the proprietary zLAS format from ESRI.¹¹
4. Most grids are in Geotiff format.¹² At least one country supplies ASCII X-Y-Z triplets without metadata like the projection or number of columns and rows. When the data is not in standard Geotiff format, GIS professionals will have to reformat it for regular users.
5. Most use a UTM projection, but several use a national projection, some of which do not have straightforward conversions to UTM or support

in all software. If reprojection is required for the chosen software, GIS professionals should handle it.

6. All include the ground classification, but the level of additional classification categories varies. GIS professionals could enhance the usability for regular users by performing a more involved classification.
7. All these data sets include the greyscale LiDAR intensity; none includes RGB channels allowed in the LAS specifications. Spanish national LiDAR, acquired after this study had been complete, includes RGB channels.
8. None uses the newer LAS 1.4 format, which would allow the classification to include power lines.¹³

Endnotes

- ¹ N Micheletti, JH Chandler & SN Lane. “Structure from motion (SFM) photogrammetry”. In LE Clarke & JM Nield (eds), *Geomorphological techniques* (online edition). London: British Society for Geomorphology. 2015. <http://www.geomorphology.org.uk/sites/default/files/geom_tech_chapters/2.2.2_sfm.pdf> Accessed on 15 January 2018, Chap. 2, Sec. 2.2.
- ² MJ Westoby, J Brasington, NF Glasser, MJ Hambrey & JM Reynolds. “‘Structure-from-motion’ photogrammetry: A low-cost, effective tool for geoscience applications”. *Geomorphology*. 179. 2012. 300–314. doi.org/10.1016/j.geomorph.2012.08.021.
- ³ S Crutchley & P Crow. “The light fantastic: Using airborne lidar in archaeological survey”. *Historic England*. 2010. <<https://historicengland.org.uk/images-books/publications/light-fantastic/>> Accessed on 8 March 2018.
- ⁴ Historic England. “3D laser scanning for heritage: Advice and guidance on the use of laser scanning in archaeology and architecture”. 8 February 2018. <<https://historicengland.org.uk/images-books/publications/3d-laser-scanning-heritage/>> Accessed on 8 March 2018.
- ⁵ C Strecha & K Montgomery. “Mapping Chillon: Comparing image reconstruction accuracy to LiDAR”. *LiDAR Magazine* 5/5. 2015. 20–22. <http://www.lidarmag.com/PDF/LiDARMagazine_StrechaMontgomery-MappingChillon_Vol5No5.pdf> Accessed on 9 March 2018.
- ⁶ N Dupont. “The king, the castle and the laser beam: LiDAR mapping for heritage applications”. *3dlasermapping.com*. 26 November 2016 <<https://www.3dlasermapping.com/blog-post/LiDAR-mapping-for-heritage-newark-castle/>> Accessed on 9 March 2018.
- ⁷ JS Evans & AT Hudak. “A multiscale curvature algorithm for classifying discrete return LiDAR in forested environments”. *Geoscience and Remote Sensing* 45/4. 2007. 1029–1038.
- ⁸ R Yokoyama, M Sirasawa & RJ Pike. “Visualizing topography by openness: A new application of image processing to digital elevation models”. *Photogrammetric Engineering and Remote Sensing*. 68/3. 2002. 257–265.

- ⁹American Society for Photogrammetry & Remote Sensing. “LAS specification version 1.4 – R13”. 15 July 2013. <https://www.asprs.org/a/society/committees/standards/LAS_1_4_r13.pdf> Accessed on 20 October 2017.
- ¹⁰M Isenberg. “LASzip: Lossless compression of LiDAR data”. *Photogrammetric Engineering and Remote Sensing* 79/2. 2013. 209–217.
- ¹¹ArgGIS. “LAS Optimizer 1.2”. 2014. <<https://www.arcgis.com/home/item.html?id=787794cbbd384261bc9bf99a860a374f>> Accessed on 25 October 2017.
- ¹²N Ritter & M Ruth. “Geotiff specification 1.8.2”. *Wayback Machine*. <<http://web.archive.org/web/20160403164508/http://www.remotesensing.org/geotiff/spec/geotiffhome.html>> Accessed on 20 October 2017; see also <<http://trac.osgeo.org/geotiff/>>.
- ¹³American Society for Photogrammetry & Remote Sensing *op. cit.*