Research Article

Gary R. Nobles*, Christopher H. Roosevelt

Filling the Void in Archaeological Excavations: 2D Point Clouds to 3D Volumes

<https://doi.org/10.1515/opar-2020-0149>

received December 7, 2020; accepted May 31, 2021

Abstract: 3D data captured from archaeological excavations are frequently left to speak for themselves. 3D models of objects are uploaded to online viewing platforms, the tops or bottoms of surfaces are visualised in 2.5D, or both are reduced to 2D representations. Representations of excavation units, in particular, often remain incompletely processed as raw surface outputs, unable to be considered individual entities that represent the individual, volumetric units of excavation. Visualisations of such surfaces, whether as point clouds or meshes, are commonly viewed as an end result in and of themselves, when they could be considered the beginning of a fully volumetric way of recording and understanding the 3D archaeological record. In describing the creation of an archaeologically focused recording routine and a 3D-focused data processing workflow, this article provides the means to fill the void between excavation-unit surfaces, thereby producing an individual volumetric entity that corresponds to each excavation unit. Drawing on datasets from the Kaymakçı Archaeological Project (KAP) in western Turkey, the article shows the potential for programmatic creation of volumetric contextual units from 2D point cloud datasets, opening a world of possibilities and challenges for the development of a truly 3D archaeological practice.

Keywords: Kaymakçı, 3D GIS, 3D visualisation, volumetric analysis, photogrammetry

1 Introduction

Recent innovations in 3D recording and the display of archaeological entities have focused primarily on the capture and visualisation of landscapes, excavations, and objects as point clouds and derived textured meshes. Innovations external to archaeology in optical laser scanning, structured light scanning, and Structure from Motion (SfM, colloquially referred to as Photogrammetry) and broader image-based data capture technologies, coupled with their increasing cost-effectiveness, have seen the spread of such technologies across the archaeological discipline (Bagi, 2018; Bevan et al., 2014; Campana, 2017; De Reu et al., 2013; De Reu et al., 2014; Doneus et al., 2011; Douglass, Lin, & Chodoronek, 2015; Kjellman, 2012; McPherron, Gernat, & Hublin, 2009; Richards-Rissetto et al., 2012; Roosevelt, Cobb, Moss, Olson, & Ünlüsoy, 2015; Willis, Koenig, & Black, 2016; Zaitceva, Vavulin, Pushkarev, & Vodyasov, 2016; Zollhöfer

Article note: This article is a part of the Special Issue on Art, Creativity and Automation. Sharing 3D Visualization Practices in Archaeology, edited by Loes Opgenhaffen, Martina Revello Lami, Hayley Mickleburgh.

* **Corresponding author: Gary R. Nobles**, Department of Geomatics, Oxford Archaeology, Janus House, Osney Mead, Oxford OX2 OES, Oxfordshire, United Kingdom; Research Center for Anatolian Civilizations, Koç University, Rumelifeneri Yolu, Sarıyer, Istanbul, 34450, Turkey, e-mail: garynobles20@gmail.com

Christopher H. Roosevelt: Research Center for Anatolian Civilizations, Archaeology and History of Art Department, Koç University, Rumelifeneri Yolu, Sarıyer, Istanbul, 34450, Turkey, e-mail: chroosevelt@ku.edu.tr
ORCID: Gary R. Nobles 0000-0002-9451-5413; Christopher H. Roosevelt 0000-0002-4302-4788

et al., 2015; amongst many others). Despite external innovations in spatial data capture, however, technological innovations internal to archaeological fieldwork methods, post-fieldwork data processing, and analytical workflows have been limited (cf. Avern & Franssens, 2012; Gavryushkina, 2018; ISAAK, 2020; Merlo & Shell, 2005). Obstacles to innovation derive in part from qualitative and quantitative differences in datasets available for experimentation and from limitations in available software.

Challenges also include the differing goals of capturing 3D data from archaeological excavations. 3D data frequently must speak for themselves, with object models uploaded to online viewing platforms, surfaces visualised in 2.5D, or both reduced to 2D representations. Models of excavation units, in particular, often remain as surfaces that are unable to represent volumetric units of excavation or stratigraphy. Surface visualisations as point clouds or meshes are frequently viewed as the end of the process. For those working towards an end goal of a truly 3D Geographic Information System (GIS), however—one anchored in 3D space and intended to enable 3D spatial analysis (Merlo, 2016; van Leusen & Nobles, 2018)—the results of such visualisations could be considered just the beginning of a fully volumetric way of recording, visualising, analysing, and thinking about the 3D archaeological record.

In describing the creation of an archaeologically focused recording routine and a 3D-focused data processing workflow, this article provides the means to fill the void between the surfaces of an excavation unit, thereby producing an individual volumetric entity that represents the excavation unit in 3D. Drawing on the large and robust datasets of the Kaymakçı Archaeological Project (KAP) in western Turkey, the article demonstrates the potential of a programmatic approach to the creation of volumetric contextual units from 2D point cloud datasets. Resultant datasets are 3D GIS compatible, even if truly 3D GIS software remains currently unavailable. Before describing the point-cloud-to-volume workflow itself, we review distinctions in multi-dimensional data, the importance of thinking spatially, and the methods used to produce the KAP dataset, returning in the end to a discussion of some foreseeable challenges and opportunities for the development of a truly 3D archaeological practice.

2 3D Technologies, 3D Data, and Spatial Thinking

Workflows that leverage new 3D data capture methods do so adhering largely to existing and traditional ways of archaeological working and thinking. That is, the integration of developing geospatial technologies has done little to challenge traditional approaches to archaeology. Very few archaeologists use these technologies to disrupt the craft or innovate significantly to create new ways of working, thinking, interacting, and interpreting archaeological information. Instead, new technologies are often adopted following the development of a new tool with push button functionality and set defaults, often yielding aesthetically pleasing results that can be interpreted visually with little to no knowledge of underlying processes. This trend has a history longer than its association with 3D technologies (Gillings, 2012; Huggett, 2004; Kvamme, 1999; Lock, 2009) and is considered to serve as a divide between archaeologists who supply a digital service and digital archaeologists (Huggett, 2015a, p. 80).

Thorough engagement with 3D technologies, however, requires at least the conceptual embrace of spatial thinking (see also Goodchild & Janelle, 2010; Lock & Pouncett, 2017; National Research Council, 2006, p. 12; Sinton, 2011). This requires re-evaluating the concept of space and its representations in archaeology and, in particular, how 3D geometries move beyond traditional 2D abstractions to provide a framework for new digital surrogates or analogues that represent the entirety of the archaeological record. It requires also the adoption of new tools of management, representation, and analysis that build on existing spatial database and GIS-like functionality and enable the answering of spatially derived research questions that depend on the visual manipulation and analysis of 3D entities (Lock & Pouncett, 2017; van Leusen & Nobles, 2018). Nevertheless, the field is not quite “there” yet.

Current approaches to recording 3D space and subsequent workflows often, if not nearly always, produce only 2.5D geometries that are frequently reduced to 2D outputs (Ledoux & Meijers, 2011; Penninga, 2008, p. 15). The typical deliverables of a 3D workflow include orthophotos that enable 2D digitisation and conformance to conventional archaeological reporting. This represents a regression from 3D data capture to

2.5D data and ultimately preferences 2D spatial thinking. 3D spatial thinking can only develop if the underlying data remain truly 3D. This call to develop a theoretical framework for 3D archaeology is echoed more broadly by those calling for Digital Archaeology to be viewed as more than a technical service and actively develop its own body of theory (Gillings, 2012; Huggett, 2012, 2015a, 2015b; Perry & Taylor, 2018; Verhagen, 2018; van Leusen & Nobles, 2018; Zubrow, 2006, p. 9).

Getting the most out of these data capture methods requires us to move beyond 2.5D geometric spatial representation to create true 3D geometric abstractions of space. The latest wave of remote sensing technologies captures real-world surfaces as highly detailed point clouds. Workflows developed for landscape and excavation applications usually lead to the production of 2.5D meshes and DEMs, as discussed by Verhoeven (2017). Some approaches extrude overlying meshes into 3D geometries, presenting the illusion of a fully recorded 3D space. The point clouds that result from archaeological data capture routines record surfaces and are thus classed only as 2D, regardless of whether the point clouds reside in 3D space or the geometries themselves reflect 2D areal surfaces. Only once a point cloud fully bounds a volumetric space should it be referred to as a 3D point cloud. 3D point clouds thus often encompass objects or meaningful units of space, the full volumes of which can be recorded: an excavation unit, for example. When representing such volumetric entities, 3D point clouds can also be referred to as semantic point clouds. Semantic point clouds are more commonly associated with Deep Learning approaches (Poux & Billen, 2019; Tostberg, 2017) and are classified according to the object they represent. In the case of an archaeological deposit, for example, semantic value is attributed only to those points classified as defining the volumetric bounds of the deposit; these are therefore retained, defining the deposit's semantic point cloud. All points not forming part of the 3D topology are thus considered extraneous.

Despite this strict definition of true 3D entities, the majority of archaeologists purporting to work with 3D GIS actually engage only with 2.5D GIS or 3D modelling. Nonetheless, significant effort has been dedicated to moving from 2D surfaces towards volumes in 2.5D or 2.5D+ GIS workflows (e.g. Barceló, De Castro, Travet, & Vicente, 2003; Berggren et al., 2015; Campanaro, Landeschi, Dell'Unto, & Leander Touati, 2016; Dell'Unto et al., 2016; Dell'Unto, Landeschi, Apel, & Poggi, 2017; Landeschi, Dell'Unto, Ferdani, Leander Touati, & Lindgren, 2015; Landeschi et al., 2019; Poggi & Buono, 2018; Schubert, Predoi, & Jeffery, 2018; van Riel, 2016). Such efforts usually rely on extrusions of 2D geometries through vertical offsets based on attribute values (referred to as 2.5D; Tse & Gold, 2004) or extrusions between under- and overlying 2D surfaces (referred to as 2.5D+; Ledoux & Meijers, 2011; Penninga, 2008, p. 15). Both methods result in the fabrication of falsely vertical feature 'walls', however. Ultimately, these methods do not create accurate and semantically meaningful 3D geometries. With rare exceptions (e.g. Gavryushkina, 2018; Landeschi et al., 2016, p. 11), their uses are increasingly geared towards visualisation with only limited analytical application. Such 2.5D and 2.5D+ methods are very different from attempts to record, visualise, manipulate, and analyse the accurate and precise 3D records of archaeological contexts in true 3D (Jensen, 2018).

Only a few studies have successfully worked with true 3D geometries (e.g. polyhedra, voxels; Crema, 2011; Lieberwirth, 2008; Losier, Pouliot, & Fortin, 2007; Merlo, 2016; Orengo, 2013). One recent innovation follows a 3D modelling approach that divides the space between two surfaces into tetrahedral cubes that can be used to approximate volumetric space for analytical purposes (Dekker, 2020a, 2020b). As highlighted by this example, archaeologists are trying to move from meshes to volumes. In contrast, this paper suggests that moving from 2D point clouds to 2D/2.5D meshes in order to reach 3D is an unnecessary step. Instead, we need to take a more direct conceptual approach, working directly from point clouds to generate entities with true 3D geometries. Before demonstrating a workflow developed for just this purpose, we briefly introduce the archaeological field project that is the source of the data used to test the workflow, along with the methods used to collect them.

3 Kaymakçı and its Born-Digital Approach

Excavations at Kaymakçı, a second millennium BCE site located in western Turkey, began in 2014 (Roosevelt et al., 2018). The site had been discovered and studied non-invasively beginning much earlier

(e.g. Luke & Roosevelt, 2017; Luke, Roosevelt, Cobb, & Çilingiroğlu, 2015; Roosevelt & Luke, 2017), revealing its citadel form with architectural remains near or on the surface, a complex organisation into several definable sectors or neighbourhoods, and relatively shallow stratigraphy. A born-digital recording system was developed for the project before excavations began and has been used to document and analyse the results of work in seven excavation areas over the course of five excavation seasons to date (Roosevelt et al., 2015; Scott, Roosevelt, Nobles, & Luke, 2021, this Special Issue).

Here, we focus on excavation area 99.526, a 9 by 9 m area in the western part of a relatively flat section of the citadel known as the Southern Terrace (Figure 1). Four seasons of work here (2014–2016, 2018) concluded upon exposure of bedrock in limited sondages. Field seasons in June and July each year included excavations undertaken by a combination of specialists, students, and hired workers, each receiving varying levels of training in the digital capture protocols, the workflows of which are under constant evaluation and development (see below; Scott et al., 2021).

Excavated to a depth of nearly 3 m from the original ground surface, the area is stratigraphically complex, feature rich, and defined by several phases of use in its several-hundred-year history spanning the mid- to late second millennium BCE (ca. eighteenth–thirteenth century BCE) (Roosevelt et al., 2018). In the earliest pre-architectural phase dating to the transition from the local Middle Bronze Age (MBA) to the Late Bronze Age (LBA), a densely packed series of surfaces, hearths/ovens, crucible-like features, pits, secondary deposits, and thin fills attest to the open-air use of the area (Roosevelt, Kaner, & Luke, 2020). At the transition from the local Late Bronze (LB) 1 to the LB 2 phase, the Southern Terrace appears to have been transformed into an urban-like complex of streets and alleys between architectural complexes of mixed domestic and household industrial function. The construction of a series of parallel dry stone wall foundations in 99.526 defined two elongated interior spaces associated with yet more surfaces, hearths/ovens, and pits. Additional fills containing fragments of mudbrick and melted mudbrick debris separate this second period of use in the area from a third period of use occurring subsequently in the LB 2 phase and attest to superstructures of mudbrick. This third period of use saw stone and mudbrick additions to the walls, slight reconfigurations of interior spaces, associated with a series of surfaces, hearths/ovens, and pits, and the addition of a semi-subterranean circular feature interpreted to be a grain silo built abutting an exterior wall.

Figure 1: Aerial view of area 99.526 following the completion of excavations, looking N. The excavation area measures 9 by 9 m.

The complex appears to have been abandoned thoroughly by the end of the 13th century BCE, after which time its mudbrick walls decayed and a series of fills accumulated to cover the stone wall foundations (Roosevelt, Ünlüsoy, & Luke, 2019).

The methods employed by the KAP recording system in the excavation of area 99.526 are the same as those employed across the site. While these have been described fully elsewhere (Roosevelt et al., 2015; Scott et al., 2021), a review of salient points is necessary here. First, all recording is done via SfM with digital photography, and the whole of an excavation area is recorded only at the start and close of a field season. In the course of excavation, the bottom of every spatial context, that is, a stratigraphic or excavational unit, is recorded, with previously recorded “bottom” data serving as “top” data for subsequently deeper and/or contiguous spatial contexts. In this way, initial recording actually documents the interfaces between the volumetric entities that define the actual spatial contexts. Coded targets for georeferencing are usually positioned outside each spatial context’s extent, with coordinate data provided through RTK-GNSS.

The data capture workflow is divided between the field teams, responsible for the field recording described above, and the so-called ‘3D Spatial’ team. Because we record via SfM, each collection of photographs and the coordinate data associated with a single spatial context is referred to as a “photobatch,” named with a dated timestamp. Photobatches are collected by the field team and transferred via an onsite network to a centralised server at the nearby research centre. The 3D Spatial team generates point clouds at the research lab in real time, parallel with excavation work. Once the point cloud for a particular photobatch has been processed, the field team is informed about its coverage as well as its photographic quality and geospatial accuracy. If sufficient, excavation may resume in immediately contiguous spaces; if insufficient, partial or full recapture of the photobatch may be necessary to provide the required high-quality point cloud data.

To provide an idea about the quantity of spatial data produced by KAP, we again refer to excavation area 99.526, described above. Over the duration of the multi-year excavations in this area, 404 contexts were documented by 449 individual point clouds. Altogether, the point clouds include a total of 2,618,058,092 individual points; that is, just in excess of 2.6 billion points. With a total excavated volume of 106.36 m³, the average recorded point density is 24615062.92 points per m³, or 246150.63 points per cm³. This equates to an average of 6480341.8 points per spatial context.

With such statistics, it might be easy to boast that these point densities represent one of the most detailed recordings of an excavation ever, yet (even if this remains likely) these numbers derive from raw point clouds, many of which overlap in capturing the same surfaces multiple times. This point redundancy results from the fact that each photobatch includes not only the points specific to the bottom of the context being recorded, but also points outside the extent of the target context, hence the inflation of the average point per cubic meter and point per spatial context values provided above. Nonetheless, the result of these field data capturing and initial processing workflows is a discrete set of point clouds that, between them, define the volumetric entity of each and every excavated spatial context. Extracting from this set of relatively messy and overlapping 2D point clouds a semantically meaningful 3D point cloud that defines one particular spatial context volume is the aim of the point-cloud-to-volume workflow described in full in the remainder of this article.

4 A Point-Cloud-to-Volume Workflow

Moving beyond the display of overlapping point clouds or meshes to the creation and manipulation of digital entities that define only the volume between them is far more challenging than it seems, with no current software providing an all-in-one solution. Programmatic approaches have been experimented with previously, yet only lengthy, complicated solutions involving manual, subjective editing of each point cloud have demonstrated satisfactory results to date (e.g. Roosevelt et al., 2015). The methodology described here still requires human input, yet moves beyond manual editing towards a smooth, programmatic workflow. Key aims in developing the workflow included database storage of (A) pre- and post-processed point cloud data (and associated metadata), along with (B) their projected coordinate system locational data, adhering to a Geographic Information Science (GISc) ethos rather than 3D modelling or the broader geodesign paradigm (Flaxman, 2010). The latter was achieved only by means of a workaround that

truncates and later restores the geographic coordinate data by a fixed, consistent value rather than through on-the-fly geographic projection. The primary application – a custom-built point cloud processing program – transforms multiple overlapping point clouds into semantically valid 3D point clouds that represent homogeneous volumetric entities. Before describing the program itself as well as the creation of closed volumetric entities, we first describe the data management and documentation necessities that provide the foundation for the full workflow.

4.1 Data Management and Documentation

To facilitate the isolation of points from several overlapping 2D point clouds and their combination into a singular semantic 3D point cloud, thoughtful data management and documentation of both stratigraphic and photobatch sequences are essential. With respect to data management, both file format and data storage procedures are key. Point cloud data can be stored in a number of formats. Flat file formats (PLY, STL, OBJ, etc.) are typically used in 3D modelling and 3D printing. Web visualisation utilises these formats as well or depends on their conversion to multi-resolution formats (e.g. NEXUS (Ponchio & Dellepiane, 2015); iTowns (Picavet, Brédif, Konini, & Devaux, 2016); Nobles, Çakırlar, & Svetachov, 2019; Potree, 2020). Point cloud data resulting from SfM recording protocols used here can be stored in flat file formats and imported into GIS software just like any other point dataset. When dealing with related sets of millions or billions of points, however, flat file storage becomes vastly inefficient and eventually unworkable.

Creative methods of storing large sets of point cloud data while avoiding inefficiencies were originally developed for data gathered using airborne laser scanning or LiDAR datasets on landscape scales. Even if such LiDAR point cloud datasets contain data of far more dimensions (Bell, Chambers, Butler, & Gerlek, 2020, pp. 290–297) than the SfM-derived datasets treated here (containing only xyz coordinate and rgb colour data), both are subject to the same types of storage, processing, and visualisation workflows. KAP thus initially stores each point cloud dataset as a LAS file, typical of LiDAR datasets, adopting the file name of its associated photobatch. This is only the first step, however, as storing each LAS dataset in a folder hierarchy requires manual searching and opening.

More efficient management of point clouds for GIS requires the use of geospatial databases and software libraries specifically developed for the purpose. The data management solution adopted here includes a backend PostgreSQL database with PostGIS and pgPointCloud extensions and Point Data Abstraction Library (PDAL) integration. PDAL works on the model of the Geospatial Data Abstraction Library (GDAL); while GDAL engages with geospatial data, PDAL focuses on point and especially point cloud data, enabling read, write, and filter operations (PDAL Contributors, 2018). A basic PDAL workflow, known as a pipeline, reads in a point cloud dataset, filters it into collections of points, and stores these points as patches within the geospatial database. Patches of 600 points each are the suggested optimal patch storage size (Bell et al., 2020, p. 138), enabling a GIS to access and visualise the patches effectively. The real power of this storage method is in its efficient division of large datasets into component patches, enabling engagement with the relevant patch(es) for specific processes, rather than the entire point cloud dataset. An additional key benefit of storing point cloud data in a geospatial database, and especially semantic 3D point clouds representing volumetric entities, is the well-known relational power of such databases: point cloud datasets can easily be related to other spatial and aspatial (or attribute) data for querying, display, and analytical purposes.

In addition to the efficiency and potential querying, display, and analytical benefits of storing point clouds in a geospatial database, such storage allows them to be accessed by a variety of platforms, including GIS software, as mentioned above, but also programming and statistical languages like R and Python, and even the Computational Geometry Algorithms Library (CGAL). Depending on the infrastructure, point clouds stored in this way can be visualised and interacted with through the web (De Kleijn, De Hond, & Martinez-Rubi, 2016), 3D modelling platforms, and bespoke interfaces, opening up various virtual reality possibilities and presenting the maximum possible potential for technological accessibility and visualisation.

Beyond efficient and accessible data management, meticulous documentation is necessary to facilitate the generation of semantic 3D point clouds. First, the top and bottom point cloud files associated with each spatial context were originally stored in a spreadsheet, an easy format for documenting and tracking such information. When it comes to understanding the relationships between contiguous point clouds, however, spreadsheet documentation is too restrictive and insufficiently illustrative. For a more intuitive and graphical means of presentation, spreadsheet information is transformed into directed graphs using open source yEd software in GraphML format (Figure 2). Each graph of related point clouds runs parallel to a typical stratigraphic matrix, documenting which point clouds are required for the creation of a target 3D point cloud that represents a spatial context. In other words, this graphical documentation makes clear how each point cloud interrelates to others. In doing so, it makes clear which point clouds need to be used in the first steps of the workflow involving the processing of component point clouds into the volumetric spatial context they enclose.

4.2 The Point Cloud Processor (PCPro)

The data management and documentation protocols described above allow the user to quickly identify and access those point clouds that form the tops and bottom of each spatial context, thus laying out a clear processing starting point. The custom-built Point Cloud Processor (PCPro) then guides the user through a relatively linear workflow including (A) importing point clouds from the database; (B) translation (via coordinate truncation); (C) sampling; and (D) removal of redundant points. The final step of (E) converting the point cloud to a solid 3D geometry continues the workflow outside the PCPro (Figure 3). The PCPro was compiled from diverse open source components written primarily in Python, but including also C++ bindings, SQL, and numerous libraries (e.g. os, sys, PyQt5, pycpg2, datetime, scipy.spatial, numpy, pandas, pptk, open3D). Its graphical user interface (GUI) was designed to allow members of the team without programming knowledge, including the field team, the potential to process the 3D data they previously captured into volumetric, semantic point clouds.

4.2.1 Importing Point Clouds from the Geospatial Database

By means of the pycpg2 library, Python can interact directly with data stored in a PostGIS geospatial database (Varazzo, 2011). Pulling from the database, the GUI presents the user with a list of available point clouds (Figure 4). With reference to the graphical documentation described above, the user chooses the relevant top(s) and bottom of the target spatial context. Behind the scenes, the program creates a working copy, preserving the original data in the database. It also translates the working copy into a local coordinate system for further processing because, as in many 3D modelling environments, many Python libraries work with data in a 32-bit floating-point format that requires the truncation of coordinates to seven figures. While not ideal, this is achieved simply with subtraction of a constant value from the x and y coordinates of the projected system; in this case 580,000 from the x and 4,275,000 from the y Universal Transverse Mercator (UTM) coordinate system used by KAP. The elevational z coordinate is preserved in its original state because of its typically low value. Once the processing is completed, the data can be returned to the projected coordinate system with the addition of these same values, enabling their display “in digital situ,” so to speak, among other associated and geospatially referenced data.

4.2.2 Sampling and Spatial Filtering

The point clouds we record with SfM methods are often incredibly dense and processing them in their entirety would take hours (see below). More importantly, experiments with the KAP dataset reveal that

Figure 2: An example of the reconstruction graph logic, with purple boxes representing photobatches (point clouds) and green boxes the spatial contexts which they define. Percent values refer to sampling values (see text).

relatively low-density point clouds of 20,000 to 50,000 points each are more than sufficient to represent the top and bottom surfaces of contexts typically recovered with KAP's modified single-context recording

Figure 3: Flow diagram of the point-cloud-to-volume workflow, with followed (thick black) and ideal (thin brown) processing paths, showing data processes and results.

system. Flexibility is required, however, as the surface of a context extending over a large area – a house floor, for instance – requires more points than the surface of a small pit. A 1% sample of a surface point

Figure 4: Examples of the PCPro GUI, showing import, cleaning, and filtering tabs.

cloud is typically sufficient, yet a 10% sample is occasionally required depending on the extent and complexity of the target surface. The sampling ratio is recorded on a per point cloud basis during processing. Sampling thresholds can be evaluated visually using the integrated Point Processing Toolkit (HERE Europe B.V., 2018), here used to display both top and bottom point clouds together (Figure 5).

Figure 5: Visualisation of the sampling process. Red points represent the top point cloud, green points the bottom point cloud.

After sampling, initial cleaning of the two point cloud datasets is achieved through a spatial filtering calculation using the concave and/or convex hulls of each of the two selected point clouds. Points located within the intersection or overlap of these two hulls are retained, while those located outside are discarded, in some cases vastly reducing the size of each point cloud. In this spatial calculation, the hulls are calculated ignoring z coordinates by default, and so essentially in only two dimensions. This entire step is optional, however, and may be omitted in particular situations (e.g. where the bottom point cloud extends significantly past the extent of the top point cloud); such situations are anticipated to be rare for most excavations.

4.2.3 Calculating Between Point Distances (BPD)

The next step in data cleaning involves a minimum “Between Point Distance” (BPD) calculation performed using the `cdist()` tool from Python’s Scipy Library (see Virtanen et al., 2020). The process calculates the minimum distance from a given point in the top point cloud to the nearest point in the bottom point cloud, records that distance as an attribute of the target point, and then iteratively continues to the next point in the top point cloud until the minimum BPD of all points in the top point cloud have been recorded. The same process is then applied from the other direction, starting from the points in the bottom point cloud.

The results of BPD calculations can be assessed visually again using the `pptk` package, assigning the same colour values to BPD point attributes (selected in the GUI using the ‘[’ and ‘]’ keys) in both point clouds. In the example given here (Figure 6), blue is assigned to points with a low BPD and green to a higher BPD. Conceptually, the blue, low value BPD points represent close fits between top and bottom point clouds, and thus, archaeologically, areas where little to no volumetric change was made during excavation. In other words, low value BPD points represent points captured in SfM recording located outside the target spatial context, are redundant, and should not be included in the construction of the volume representing the target spatial context. Conversely, the green, higher value BPD points give a clear indication of the volumetric bounds of the target spatial context.

Figure 6: Visualisation of the minimum between point distance (BPD) calculation. Blue points show similarity between the two point clouds, with low BPD values; green points show differences between the two point clouds, with high BPD values.

Removing blue, low value BPD points are conducted via filtering with a selectable distance tolerance or threshold. Higher or lower thresholds can be set for the differing levels of precision associated with different in-field recording methods (e.g. total station versus RTK GNSS coordinate recording). Following extensive testing with the KAP dataset and its RTK GNSS recording method, a 3 cm threshold was adopted for most contexts, whereby points with BPD values less than 3 cm were removed. This filtering removes most of the core noise in the point clouds, yet the minor background scatter or outlier noise associated with most point cloud creation workflows – data artefacts or points missed in earlier cleaning steps – still needs isolation and removal.

Before considering outlier removal, however, it is instructive to review this most intensive stage of processing *vis-à-vis* a common concern in processing 3D data, particularly those derived from SfM

workflows: processing time. The time efficiency of the `cdist()` function was evaluated using sets of random point clouds containing between 10,000 and 1,000,000 points on the KAP server (Intel® Xeon® Silver 4114 CPU @ 2.20 GHz; 64.0 GB RAM; 64-bit Windows Server 2016 Standard OS; NVIDIA Quadro RTX 4000 GPU), without GPU acceleration. Figure 7 demonstrates the time taken to calculate the BPD of two point clouds of various sizes. BPD calculation for a pair of point clouds each containing 1 million points takes around seven hours per point cloud, or around 14 h to process both. A pair of point clouds containing 50,000 points each takes 56 s per calculation, while a pair of point clouds containing 100,000 points each takes 3 min and 39 s per calculation. Any pair of point clouds with fewer than 200,000 points each can take up to 30 min to calculate the BPD for the set. Sampling the raw point clouds before BPD calculation, therefore, is a crucial step for ensuring an efficient workflow, and experimentation is required to ascertain the level of sampling required to achieve acceptable end results. If the convex and concave hull filtering is applied, the size of the point cloud can be further reduced saving even more time in this step. Concerning the KAP dataset, point clouds of between 20,000 to 50,000 points each were generally the accepted quantity depending on the complexity of the surface.

Figure 7: Graphs showing BPD processing times relative to the size of the primary and secondary point clouds. The top graph represents the entire testing range, whereas the bottom represents the indicated subset.

4.2.4 Removing Outlying Points

Background scatter and other point cloud outlier data are filtered with the help of the Open3D Library (Zhou, Park, & Koltin, 2018), an extremely powerful Python and C++ library designed specifically for working with 3D data and containing numerous filtering options. Because our workflow begins with

sampling the point cloud data and aims to preserve the original variability of the 3D extents of spatial contexts as much as possible, the Open3D Library's recommended first downsampling routine using a voxel grid is omitted from the PCPro workflow. Similarly, the library's radius outlier removal tool, which applies a spherical search at set distances and removes points that do not have a sufficient number of neighbouring points within a set distance, is omitted because it failed to produce desired results. The statistical outlier removal tool, however, provided desired functionality. In assessing a point cloud, this tool considers two parameters, the number of neighbouring points to assess in each iteration and a cut off distance based on a ratio value. The ratio is a threshold based on the standard deviation of the global average. The higher the ratio, the less aggressive the function. The PCPro uses both this tool and Open3D's visualisation library to display those points flagged for removal by outlier filtering, so users can learn which values are applicable in which situations and see the resulting point cloud (Figure 8). Typical outlier removal filtering parameters adopted for the KAP dataset, obtained after extensive testing, are displayed in Table 1.

Figure 8: Visualisation of the outlier point removal process, with outlier points to be removed shown in red. Left: top point cloud. Right: bottom point cloud.

The removal of outliers brings us to the end of the PCPro workflow for the time being (see below), with resultant outputs including sets of cleaned point clouds that accurately define the tops and bottoms of spatial contexts with no redundancy. In addition to its smooth workflow and time efficiencies, the PCPro's direct engagement with the spatial database enables iterative experimentation until satisfactory results are achieved. Original data remain safely stored in the database throughout all sampling and filtering steps, and fresh copies can be retrieved repeatedly at will, resetting the process to earlier steps. Once typical filtering parameters are assessed for a particular dataset, however, it is our experience that continued experimentation is rarely needed, allowing regular progress through large datasets.

4.3 Producing Solid Geometry Volumes

For final processing of top and bottom point clouds into semantic 3D point clouds and closed volumetric entities, current limitations in spatial databases and related libraries forced the workflow to take a detour from its GIScience-oriented approach to the more established 3D Modelling or Geodesign paradigm. Each cleaned point cloud, now stored as xyz point geometries in PostGIS, was exported to a PCD (point cloud) file format for use in the open source CloudCompare software. In CloudCompare, the normals of each top and bottom point cloud were calculated using the triangulation method before saving the point clouds as separate files. These were then merged to create the semantic 3D point cloud representing the excavated spatial context (Figure 9).

Table 1: Typical values in outlier filtering via open3D library statistical outlier removal tool adopted for use with the KAP dataset resulting effects

Effect	Neighbours	Ratio
Fine (low number of points removed)	6	10
Fine (low number of points removed)	50	10
Medium fine	50	5
Coarse (high number of points removed)	500	1
Coarse (high number of points removed)	1000	1

The semantic 3D point cloud is then converted into a solid geometry in ways similar to previously established protocols (e.g. Roosevelt et al., 2015), using Kazhdan’s (2005) Poisson Surface Reconstruction method with an octree of 8 (Figure 9). We hope and expect that solid geometry construction can eventually be worked into the PCPro workflow using Molero’s (2020) Python bindings, but this and other integrations currently await programmatic developments that will allow the entire workflow (including point cloud normal alignment and merging) to be subsumed within the database-linked GIScience approach. At present, the ready availability of other open source software that provides suitable stop-gap functionality enables the completion of the processing.

That said, to bring the process full circle, the solid geometry volumes would ideally be imported back into PostGIS as 3D (polyhedral) geometries. A working 3D mesh importer for PostGIS can no longer be found, however, despite previously available examples (e.g. Rolland’s (2014) pg3Dimporter, previously

Figure 9: Visualisation of the result of cleaning and merging the top and bottom point clouds of KAP spatial context 99.526.43. Top: the 3D semantic point cloud. Bottom: the same, converted to a solid geometry volume.

available by Dropbox link). Critical here are differences in the ways in which geometry normals are recorded according to 3D modelling (with direct recording as attributes) versus geodatabase (inferred by direction of construction) paradigms. Attempts at writing a parser to make this conversion were successful, yet the complexity of the polyhedral geometries associated with excavated context volumes resulted in repeated program instability and thus no further experimentation. Similar attempts to import these complex meshes into other GIS databases also failed; for example, COLLADA files imported into ArcGIS Pro (on the suggestion of P. Gerrits, personal communications). Unfortunately, it seems that no GIS database can yet cope with the complexity of the real-world geometry of excavated contexts, and this is likely because none have been designed as true 3D GIS systems for 3D geometries rather than for 2D and 2.5D+ geometries (ESRI, 2012; Khuan, Abdul-Rahman, & Zlatanova, 2008, p. 290).

Nonetheless, other non-geodatabase-linked options exist for displaying, manipulating, and analysing the end products of the workflow described here: solid geometry volumes representing archaeological contexts. Visualisation was not a core focus in the development of this workflow, which instead prioritised retention of accurately representative geometries suitable for quantitative analysis. The quantitative potential is demonstrated through the Python trimesh library (Dawson-Haggerty, 2019), for instance, enabling the calculation of the volume of each context in cubic units. Further visualisation is possible via a number of other simple to complex open source software solutions (e.g. MeshLab, CloudCompare, Blender, ParaView), and their selection depends on a user's particular needs and personal preference for software and display types. In what follows, for instance, we use Blender to produce most figures, yet a variety of both open source and commercial software packages could do the same.

5 Filling the Void: Reflections on Approach, Visualisation, and Spatial Thinking

The immediate result of the development of this workflow was the creation of hundreds of solid geometry volumes representing the excavated spatial contexts at Kaymakçı. These results allow us to reflect on the importance of a programmatic approach and to explore both the potential of context volumes for analysis and interpretation and the challenges of visualisation as we move towards the future and truly 3D GIS.

5.1 The Importance of a Programmatic Approach

The “redundant plentitude” of data (quoting the call to the symposium for which the first steps of this article were developed) produced by SfM recording shouldn't encourage us to find ways to reduce data in the collection phase. Yet it does require a particular (set of) approach(es) that enables us to move beyond sheer masses of messy, raw data to meaningful, semantic entities. Only small fractions of collected data are actually necessary for the production and display of meaningful solid geometry volumes. In our case, a programmatic approach was instrumental in allowing us to reduce data gargantua to streamlined and usable sizes objectively, while preserving all data for subsequent processing for yet unrealised purposes. That is, the current programmatic approach and its associated data management protocols allow both the achievement of our immediate goal – the creation of volumetric entities representing archaeological stratigraphy and the process of its removal – and the preservation of the data in full detail for future, yet-to-be posed research questions, with yet-to-be determined data requirements. These protocols depend on accurate recording and full documentation of photobatches, best achieved not later but simultaneously with in-field recording. Furthermore, the data management protocols required for this programmatic approach prepare the data for reusability and the potential reproducibility of our methods following FAIR principles (FAIRPort Mobile Data Open Movement, 2020; Open Data Charter, 2019; Wilkinson, Dumontier, & Aalbersberg, 2016).

An additional benefit of the programmatic approach is that it allows the versioning of entire workflows for tracking developmental iterations and resultant outputs. While the current features and optimal configurations of the PCPro presented here might represent version *alpha* 0.1.0, we foresee the incremental inclusion of multiple improvements in future releases concerning both processing and meaningful outputs. For example, integrating OpenCL and CUDA would enable GPU acceleration, reducing the time required for individual point cloud processing to perhaps only a few seconds. Taking advantage of appropriate PDAL libraries, automated dimensional and orientation calculations could be written directly to the associated spatial database, making such information available alongside the semantic 3D point cloud. Important next steps include these and other improvements to the general PCPro workflow, generalising its required inputs for use broader than the KAP dataset and thus enabling others to begin creating their own truly 3D volumetric datasets. Built-in capacity for experimentation with sampling, filtering, and other parameter thresholds enables others also to cater the PCPro to the specific characteristics of the volumes they wish to model as well as to desired balances of processing time and model quality.

5.2 Visualising Individual Features, Deposits, and Stratigraphic Sequences

In the first publication of the KAP recording system (Roosevelt et al., 2015), the solid geometry volumetric fills of a stone-lined, semi-subterranean pit (99.526.80), interpreted to be a grain silo, were generated using a manual point cloud cleaning procedure. To test the new programmatic approach, the same dataset of point clouds was passed through the point-cloud-to-volume workflow described above. Comparison of the products shows expectedly similar results (Figure 10), with the point cloud processor removing redundant and outlier points programmatically and thus more consistently. The semi-automated workflow reduces inconsistencies with formalised filtering and use of distance tolerances rather than manual and subjective judgements that can differ between researchers. It also reduces the effort required and inevitably speeds up

Figure 10: Comparison of individual and grouped solid geometry volumes produced manually (at left, Roosevelt et al., 2015) and via the programmatic workflow presented in this article (at right). Numbers in parentheses at right indicate context numbers within excavation area 99.526. Arrows at left point to the wispy volumes of sediment resulting from cleaning the interstices between wall stones.

the process from hours to minutes. It bears mention, however, that the workflow also removes some details that a manual and thus more nuanced and subjective approach might intend to preserve: the vertical wisps associated especially with contexts 43 and 47 in Figure 10, for instance, which represent the volumes resulting from subsequent cleaning of the interstices of a wall. Special attention must be paid, then, to preserve such features if desired. The workflow performs particularly smoothly with straightforward stratigraphic sequences of pits and fills. To test other features and stratigraphic scenarios, other datasets were selected from excavation area 99.526.

At a site like Kaymakçı, stone wall foundations are commonly exposed features. Yet most walls remain *in situ* once exposed, posing challenges for their isolation as 3D entities in the KAP recording system. Solid geometry volumes can be created only for spatial contexts that have been fully removed and thus have recorded top and bottom point clouds. Hence, it is only those contexts that are removed, or destroyed, that become fully digitised in 3D (à la the title of Roosevelt and colleagues' (2015) article [for more discussion, see Scott et al., 2021]). Features like walls that remain “in real situ” must be displayed as surfaces, around which solid geometry volumes can be displayed “in digital situ.” Walls or wall elements that are removed – secondary phase wall additions, for instance – can be fully rendered as solid geometry volumes (Figure 11).

More complex examples of feature and deposit stratigraphy are quite common at the site. By way of example, we can refer to oven 99.526.79 (see also Scott et al., 2021). The excavation of this oven began by sectioning it and removing thin fired surfaces, layer by layer, on one side, while the other section was excavated in a single contextual unit (Figures 12 and 13). In other cases, multiple point clouds formed the compound top of a single context, requiring the point clouds to be merged before continuing through the workflow. All in all, the workflow performed well with most target contexts. Even so, the PCPro has difficulty with thin or tapering layers like partially preserved floor surfaces and subsurfaces, and particularly with contexts whose thickness is less than the recording margin of error and related filtering tolerance of 3 cm. In such cases, regardless of their archaeological significance, it was necessary to group multiple thin contexts together to create a solid geometry volume that represented their combined totality. Improvements to these situations might be achievable through the use of TotalStation and EDM devices rather than RTK GNSS. This would increase relative accuracy between top and bottom point clouds, yet, like all workflow modifications, would involve other less positive trade-offs as well, particularly concerning duration of time for field recording.

Figure 11: Visualisation of the post-excavation surface model of area 99.526 with solid geometry volumes. The orange elements are secondary wall additions 99.526.14 (0.497 m^3) and 15 (0.252 m^3), which were removed during the course of excavation. They are situated “in digital situ” atop wall 99.526.51, which is displayed as a surface rather than a volume because it was never removed and remains standing (“in real situ”). Middle: top-down view with north at the top. Left and right: oblique views to north and east, respectively.

Figure 12: Visualisation of the components of oven 99.526.79, presented in an alternative excavation sequence (a–f), so as to better visualise the stratigraphy of the feature. Oblique view to the southwest.

While (re)viewing individual features and deposits at various stages of their excavation can be instructive (e.g. Figures 10–13)—demonstrating their positionality, configuration, and extents—larger stratigraphic sequences of features and/or deposits can be (re)viewed similarly, from a variety of angles, and “in digital situ.” These sequences can be presented with exploded views that preserve stratigraphic relationships while enabling evaluation and testing of interpretations (Figure 14). Juxtaposed with Harris Matrix-like stratigraphic sequences, such “exploded” views make visually clear the complexity of physical overlaps and the positions of specific context volumes in a sequence, thereby enabling more intuitive understandings of particular contexts as well as site formation processes (see also De Roo et al., 2016). These types of interactions with the digital archaeological record can greatly aid the reconstruction of phasing for entire excavation areas. Yet their effectiveness in achieving such purposes depends greatly on how they are visualised.

Figure 13: Visualisation of oven 99.526.79, with all component volumes grouped into a single volume. Oblique view to the southwest.

Figure 14: Visualisations of selected contexts in area 99.526 to help communicate stratigraphic relationships. Bottom: 1:1 volumetric. Top: Exploded. Right: traditional abstraction of the stratigraphic sequence.

To be sure, the benefits of such visualisations depend on exploration and experimentation with display and symbology parameters in 3D modelling software, and this will likely remain valid even when 3D GIS software is available. How can a full sequence of context volumes be visualised at excavation-area scales without significantly obscuring at least some of the data? Simply put, it's impossible. Processes of reasoning in spatial thinking thus need to be engaged in displaying volumetric entities in ways that enable interpreting the digital archaeological record in effective ways. Transformation of the digital archaeological record in relative space (including rotating, tilting, and the “exploded” views mentioned above) and adjustment of relative transparencies to highlight context volumes of interest (e.g. Figure 15) are just two methods of subtly modifying virtual realities to avoid visual obstruction of relevant data and aid processes of reasoning and interpretation.

Figure 15: Visualisations of selected contexts in area 99.526 to demonstrate display challenges and the use of relative transparencies. Top: top-down view with north at the top. Middle: oblique view to north. Bottom: horizontal view to north.

Similarly, because the SfM recording workflow captures and can apply RGB colour values to point clouds, photorealistic visualisation of particular contexts is possible. Owing to the subsampling of data for efficient processing, however, RGB-rendered point clouds and resulting colour meshes appear somewhat blurred and distract more than inform (Nobles et al., 2019). Full detail texture baking over subsampled point clouds is an option for future development. For the time being, however, simple false colour visualisation seems to be the most effective solution for distinguishing particular context volumes.

5.3 3D GIS, Technological Innovation, and Spatial Thinking

Visualisation of 3D data and its current challenges discussed above involve primarily aesthetic manipulations of representation, whether conducted in 3D modelling or 3D GIS environments. As important as they are in a thinking, spatial, approach to understanding spaces and places, more archaeologically meaningful (and not just spatial) interactions must await the development of 3D GIS. The power of displaying and manipulating solid geometry volumes “in digital situ,” then, has yet to be realised fully owing to the previously bemoaned lack of appropriate software solutions. Individual features, deposits, and stratigraphic sequences can be displayed in varying configurations to enhance understanding, but these understandings remain aesthetically rather than computationally derived for the time being. Were 3D GIS software available, we would be able to query by attribute or location, select, group, and order the context volumes along with other typical GIS and spatial database functions (van Leusen & Nobles, 2018, p. 476). We would also be able to section and visualise them in any configuration desired to understand and analyse stratigraphic relationships more fully. In the meantime, as with the final processing steps described above, similar functions can be conducted without connection to a spatial database and yet with clear interpretive benefits in 3D modelling environments. One potential 3D GIS future may be reached soon with robust point-cloud integration. Following successful recent fundraising, the process of storing point clouds in spatial databases and interacting with them natively within QGIS is in development and may appear as early as 2022 (Dobias & Razmjooei, 2020). Until then, however, we should consider carefully what digital tools should be developed to best enhance archaeological decision making and interpretative practices, rather than letting the development of digital tools dictate how those archaeological practices evolve.

This is not to say that technological innovation is not an important driver; it is, but the technoscape is only one element of broader culture flows (Appadurai, 1996) and how societies, in this case disciplines, change. Technological change indeed does drive significant “societal change” in archaeology, from broad research designs to specific methodologies, whether related to archaeological field practice or archaeological GIScience. In the case of 3D GIS, we have mentioned that the technology is not readily available to proceed; we lack a truly 3D GIS, one capable of working with multiple types of 3D geometry (e.g. polyhedrals, voxels) in projected 3D space (see Goodchild & Janelle, 2010). In one sense, we are being led by technological determinism, applying methods which are being afforded and not reflecting critically enough on the requirements of 3D archaeology.

A truly 3D archaeology seems to be limited by two fundamental factors, then: technology and spatial thinking. The capabilities of software tools developed by large companies too often guide approaches, shape thinking, and thus shape knowledge production, whether or not a particular tool’s inner workings are understood or the relevance of its outputs to archaeological questions has been demonstrated. Archaeologists thus need to ensure with careful inspection that the tool fits the problem. Tools and technological directions do not usually enhance archaeological thinking all on their own (e.g. Lock & Pouncett, 2017), but the ways in which archaeologists employ tools and use and interpret their resultant data make them relevant to the field.

Presenting data in 3D space (with 3D or 2.5D+ geometries) requires a major rethink of how we conceptualise that space, regardless of the technology involved. It goes beyond geometries towards a more holistic perspective, questioning how we structure 3D data and relate the rest of the archaeological record to them. How 3D entities and associated data are presented or visualised and analysed affects communication,

comprehension, and thus the ability to use spatial reasoning to explain the archaeological record. So how should archaeologists embrace working in digital 3D environments in ways that best suit archaeological questions and explanatory models? Lock and Pouncett (2017) suggest that archaeologists engaging with spatial technologies should develop clearer lines of thinking about concepts of space, tools of representation, and processes of reasoning, encouraging the development of the deeper understandings that digital archaeologists require (see Huggett, 2015b).

Technologies relevant to archaeology will continue to emerge regularly. As digital archaeologists, it falls to us to delve into their “nuts and bolts,” to dive beneath the surface of any particular technology, to evaluate its application, weed out its gremlins, and report back to the wider discipline. With respect to 3D environments, we cannot afford to be satisfied only by visually pleasing aesthetics, even if they enhance communication; geometric robustness is an enabling factor for more analytical directions and thus must come first. A productively critical stance is required now more than ever as the discipline of archaeology continues to navigate the digital and spatial turns seen across broader humanities and social science fields.

6 Conclusion

From its inception, the born-digital and born-3D work at Kaymakçı was founded on a firm theoretical principle: a robust protocol and technological toolkit could create digital and volumetric entities representing excavation and stratigraphic units by filling the void between recorded archaeological surfaces (Roosevelt et al., 2015). While in-field SfM recording had ensured the capture of appropriate data from the beginning, the realisation of the holistic ambition for a fully 3D archaeological record had made only slow progress until recently. The workflow this article describes, and in particular the PCPro, introduces a new programmatic approach to the challenge, beginning with the 2D, surface point clouds that result from SfM recording.

The process of archaeologists developing a programmatic approach to resolve an archaeologically derived yet computationally dependent challenge has required a focus on the quality of the data and associated excavation, recording, and documentation methods at a forensic level. The workflow described here was custom-built based on working norms of one particular project, but can be generalised to handle the point clouds derived from any excavation system. Resultant outputs are fully 3D, semantic point clouds that can be visualised using a variety of applications and approaches. While forced to depart from a GIScience environment built on the prioritisation of geospatial database integration, the point clouds can easily be transformed into solid geometry volumes using open-source software, completing the point-cloud-to-volume workflow. To this end, the recording methods used to create the KAP dataset and the development of this workflow have together enabled the programmatic generation of truly 3D volumetric geometries. Additionally, the process has highlighted challenges and opportunities.

The voids addressed in this paper, then, refer not just to those between 2D surfaces and 3D volumes. They refer also to technological and conceptual voids concerning analytical software environments and spatial thinking. We have now improved the ease with which digital entities can be produced that are “3D GIS-ready,” yet 3D GIS is not yet itself ready. While current software solutions offer aesthetic and communicative opportunities for displaying this 3D data, what continues to be missing is a connection to its substance, that is, the archaeological information with which it is associated. Missing also is a holistic approach to understanding what can and should be accomplished archaeologically with 3D data, in thinking spatial analyses. Only once such related archaeological data are attached to entities within a 3D GIS, available for GIS-like spatial and non-spatial manipulation and analysis, can we begin to consider truly 3D archaeological work.

Acknowledgements: For Kaymakçı excavation permissions and assistance, we thank the General Directorate of Cultural Heritage and Museums of Turkey’s Ministry of Culture and Tourism, and its annual representatives, as well as the director and staff of the Manisa Museum Directorate. We are grateful to A. Crowe, C. B. Scott, and E. Kaner for their supervision of the meticulous excavation and recording of area 99.526 at

Kaymakçı, as well as to the entire KAP team who have contributed to the development of its 3D dataset and evolving approaches to its analysis, especially E. Moss. For their support and helpful comments on earlier drafts, we also thank C. B. Scott and C. Luke.

Funding information: Research for this paper was conducted with fellowship support from Koç University's Research Center for Anatolian Civilizations (ANAMED) and excavation support from the National Endowment for the Humanities (Award RZ5155613), National Science Foundation (Award BCS-1261363), Institute for Aegean Prehistory, Loeb Classical Library Foundation, Merops Foundation, Boston University Vecchiotti Archaeology Fund, Koç University, and private donors. Any opinions, findings, conclusions, or recommendations are those of the authors and do not necessarily reflect the views of the supporting institutions.

Conflict of interest: Authors state no conflict of interest.

References

- Appadurai, A. (1996). *Modernity at large cultural dimensions of globalization*, Minneapolis: University of Minnesota Press.
- Avern, G., & Franssens, W. (2012). A digital drawing tool for recording excavations: The Nikon iSpace System. In M. Zhou, I. Romanowska, Z. Wu, P. Xu, & P. Verhagen (Eds.), *Revive the past. Computer applications and quantitative methods in archaeology (CAA). Proceedings of the 39th international conference, Beijing, April 12–16* (pp. 21–29). Amsterdam: Pallas Publications.
- Bagi, O. (2018). The process of 3D documentation in archaeological fieldwork: A case study from the archaeological site of Metsamor. *Polish Archaeology in the Mediterranean*, 26(1), 795–808. doi: 10.5604/01.3001.0012.1807.
- Barceló, J. A., De Castro, O., Travet, D., & Vicente, O. (2003). A 3d model of an archaeological excavation. In M. Doerr & A. Sarris (Eds.), *The digital heritage of archaeology. Computer applications and quantitative methods in archaeology*. Athens: Hellenic Ministry of Culture. Archive of Monuments and Publications.
- Berggren, Å., Dell'Unto, N., Forte, M., Haddow, S., Hodder, I., Issavi, J., ... Taylor, J. S. (2015). Revisiting reflexive archaeology at Çatalhöyük: Integrating digital and 3D technologies at the trowel's edge. *Antiquity*, 89(344), 433–448. doi: 10.15184/aqy.2014.43.
- Bell, A., Chambers, B., Butler, H., & Gerlek, M. (2020, November 5). PDAL: Point cloud data abstraction library. Release 2.2.0. <https://pdl.io/PDAL.pdf>
- Bevan, A., Li, X., Martínón-Torres, M., Green, S., Xia, Y., Zhao, K., ... Rehren, T. (2014). Computer vision, archaeological classification and China's terracotta warriors. *Journal of Archaeological Science*, 49(1), 249–254. doi: 10.1016/j.jas.2014.05.014.
- Campana, S. (2017). Drones in archaeology. State-of-the-art and future perspectives. *Archaeological Prospection*, 24(4), 275–296. doi: 10.1002/arp.1569.
- Campanaro, D. M., Landeschi, G., Dell'Unto, N., & Leander Touati, A. M. (2016). 3D GIS for cultural heritage restoration: A “white box” workflow. *Journal of Cultural Heritage*, 18, 321–332. doi: 10.1016/j.culher.2015.09.006.
- Contributors, P. D. A. L. (2018). *PDAL point data abstraction library*.
- Crema, E. R. (2011). Aoristic approaches and voxel models for spatial analysis. In E. Jerem, F. Redő, & V. Szevényi (Eds.), *On the road to reconstructing the past. Computer applications and quantitative methods in archaeology (CAA). Proceedings of the 36th international conference. Budapest, April 2–6, 2008* (pp. 179–186). Budapest: Archeaeolingua.
- Dawson-Haggerty, M. (2019). Trimesh. Version 3.2.0. <https://trimsh.org/>
- Dekker, B. (2020a). *Een realistisch digitaal model voor archeologische opgravingen* [Unpublished BA honors thesis]. Amsterdam: University of Amsterdam.
- Dekker, B. (2020b, November 11). 3D volume generation. Archaeological-surfaces-to-volume. <https://github.com/TJRL/Archaeological-surfaces-to-volume>
- De Kleijn, M., De Hond, R., & Martinez-Rubi, O. (2016). A 3D spatial data infrastructure for mapping the Via Appia. *Digital Applications in Archaeology and Cultural Heritage*, 3(2), 23–32. doi: 10.1016/j.daach.2016.03.001.
- De Reu, J., De Clercq, W., Sergeant, J., Deconynck, J., & Laloo, P. (2013). Orthophoto mapping and digital surface modeling for archaeological excavations: an image-based 3D modeling approach. *Digital Heritage International Congress (DigitalHeritage)* (pp. 205–208). New York: IEEE. doi: 10.1109/DigitalHeritage.2013.6743734.
- De Reu, J., De Smedt, P., Herremans, D., Van Meirvenne, M., Laloo, P., & De Clercq, W. (2014). On introducing an image-based 3D reconstruction method in archaeological excavation practice. *Journal of Archaeological Science*, 41, 251–262. doi: 10.1016/j.jas.2013.08.020.

- De Roo, B., Stal, C., Lonville, B., De Wulf, A., Bourgeois, J., & De Maeyer, P. (2016). Spatiotemporal data as the foundation of an archaeological stratigraphy extraction and management system. *Journal of Cultural Heritage*, 19, 522–530. doi: 10.1016/j.culher.2015.12.001.
- Dell'Unto, N., Landeschi, G., Leander Touati, A.-M., Ferdani, D., Dellepiane, M., Callieri, M., & Lindgren, S. (2016). Experiencing ancient buildings from a 3D GIS perspective: A case drawn from the Swedish Pompeii project. *Journal of Archaeological Method and Theory*, 23, 73–94. doi: 10.1007/s10816-014-9226-7.
- Dell'Unto, N., Landeschi, G., Apel, J., & Poggi, G. (2017). 4D recording at the trowel's edge: Using three-dimensional simulation platforms to support field interpretation. *Journal of Archaeological Science: Reports*, 12, 632–645. doi: 10.1016/j.jasrep.2017.03.011.
- Dobias, M., & Razmjooei, S. (2020, August 24). Crowdfunding: Support for point cloud data in QGIS. *Lutra Consulting*. <https://www.lutraconsulting.co.uk/blog/2020/08/24/pointcloud-qgis/>
- Doneus, M., Verhoeven, G., Fera, M., Briese, C., Kucera, M., & Neubauer, W. (2011). From deposit to point cloud – a study of low-cost computer vision approaches for the straightforward documentation of archaeological excavations. *Geoinformatics FCE CTU*, 6, 81–88. doi: 10.14311/gi.6.11.
- Doneus, M., Mandlbürger, G., & Doneus, N. (2020). Archaeological ground point filtering of airborne laser scan derived point-clouds in a difficult mediterranean environment. *Journal of Computer Applications in Archaeology*, 3(1), 92–108. doi: 10.5334/jcaa.44.
- Douglass, M., Lin, S., & Chodoronek, M. (2015). The application of 3D photogrammetry for in-field documentation of archaeological features. *Advances in Archaeological Practice*, 3(2), 136–152. doi: 10.7183/2326-3768.3.2.136.
- ESRI. (2012). *The multipatch geometry type – An ESRI white paper December 2008*. Redlands, CA: ESRI. <http://www.esri.com/library/whitepapers/pdfs/multipatch-geometry-type.pdf>
- FAIRPort Mobile Data Open Movement. (2020, December 6). Data FAIRport. Find, access, interoperate & re-use data. <https://www.datafairport.org/>
- Flaxman, M. (2010). Fundamentals of geodesign. In E. Buhmann, M. Pietsch, & E. Kretzler (Eds.), *Proceedings of digital landscape architecture* (pp. 28–41). Anhalt: Anhalt University of Applied Science.
- Gavryushkina, M. (2018). *Layer by (3D) Layer 3D GIS stratigraphic analysis of chlorakas-palloures, cyprus* [MA thesis]. Leiden: University of Leiden.
- Gillings, M. (2012). Landscape phenomenology, GIS and the role of affordance. *Journal of Archaeological Method and Theory*, 19(4), 601–611. doi: 10.1007/s10816-012-9137-4.
- Goodchild, M. F., & Janelle, D. G. (2010). Toward critical spatial thinking in the social sciences and humanities. *GeoJournal*, 75(1), 3–13. doi: 10.1007/s10708-010-9340-3.
- HERE Europe B. V. (2018). The point processing toolkit (pptk). <https://github.com/heremaps/pptk>
- Huggett, J. (2004). Archaeology and the new technological fetishism. *Archeologia e Calcolatori*, 15, 81–92.
- Huggett, J. (2012). Core or periphery? Digital humanities from an archaeological perspective. *Historical Social Research*, 37(3), 86–105.
- Huggett, J. (2015a). Challenging digital archaeology. *Open Archaeology*, 1(1), 79–85. doi: 10.1515/opar-2015-0003.
- Huggett, J. (2015b). A manifesto for an introspective digital archaeology. *Open Archaeology*, 1(1), 86–95. doi: 10.1515/opar-2015-0002.
- ISAAK. (2020). Initiative for statistical analysis in archaeology Kiel. <https://github.com/ISAAKiel>
- Jensen, P. (2018). Semantically enhanced 3D: A web-based platform for spatial integration of excavation documentation at Alken Enge, Denmark. *Journal of Field Archaeology*, 43(sup1), S31–S44. doi: 10.1080/00934690.2018.1510299.
- Kazhdan, M. (2005). Reconstruction of solid models from oriented point sets. In *Proceedings of the third Eurographics symposium on Geometry processing*, (SGP '05) (pp. 73–82). Goslar, Germany: Eurographics Association.
- Khuan, C. T., Abdul-Rahman, A., & Zlatanova, S. (2008). 3D solids and their management in DBMS. In P. van Oosterom, S. Zlatanova, F. Penninga, & E. Fendel (Eds.), *Advances in 3D geoinformation systems* (pp. 279–311). New York: Springer.
- Kjellman, E. (2012). *From 2D to 3D A photogrammetric revolution in archaeology?* [MA thesis]. Tromsø: University of Tromsø.
- Kvamme, K. L. (1999). Recent directions and developments in geographical information systems. *Journal of Archaeological Research*, 7(2), 153–201.
- Landeschi, G., Dell'Unto, N., Ferdani, D., Leander Touati, A.-M., & Lindgren, S. (2015). Enhanced 3D-GIS: Documenting Insula V 1 in Pompeii. In F. Gillign, F. Djindjian, L. Costa, P. Moscati, & S. Robert (Eds.), *CAA2014 21st century Archaeology: Concepts, methods and tools. Proceedings of the 42nd annual conference on computer applications and quantitative methods in archaeology* (pp. 349–360). Oxford: Archaeopress.
- Landeschi, G., Dell'Unto, N., Lundqvist, K., Ferdani, D., Marco, D., & Touati, A. L. (2016). 3D-GIS as a platform for visual analysis: Investigating a Pompeian house. *Journal of Archaeological Research*, 65, 103–113. doi: 10.1016/j.jas.2015.11.002.
- Landeschi, G., Apel, J., Lundström, V., Storå, J., Lindgren, S., & Dell'Unto, N. (2019). Re-enacting the sequence: Combined digital methods to study a prehistoric cave. *Archaeological and Anthropological Sciences*, 11(6), 2805–2819. doi: 10.1007/s12520-018-0724-5.
- Ledoux, H., & Meijers, M. (2011). Topologically consistent 3D city models obtained by extrusion. *International Journal of Geographical Information Science*, 25(4), 557–74. doi: 10.1080/13658811003623277.

- Lieberwirth, U. (2008). 3D GIS voxel-based model building in archaeology. In A. Posluschny, K. Lambers, & I. Herzog (Eds.), *Layers of perception. Proceedings of the 35th international conference on computer applications and quantitative methods in archaeology (CAA), Berlin, Germany, April 2–6, 2007* (pp. 1–8). Bonn: Dr. Rudolf Habelt GmbH.
- Lock, G., & Pouncett, J. (2017). Spatial thinking in archaeology: Is GIS the answer? *Journal of Archaeological Science*, 84, 129–135. doi: 10.1016/j.jas.2017.06.002.
- Lock, G. (2009). Archaeological computing then and now: Theory and practice, intentions and tensions. *Archeologia e Calcolatori*, 20, 75–84.
- Losier, L.-M., Pouliot, J., & Fortin, M. (2007). 3D geometrical modeling of excavation units at the archaeological site of Tell ‘Acharneh (Syria). *Journal of Archaeological Science*, 34(2), 272–288. doi: 10.1016/j.jas.2006.05.008.
- Luke, C., Roosevelt, C. H., Cobb, P., & Çilingiroğlu, Ç. (2015). Composing communities: Chalcolithic through Iron Age survey ceramics in the Marmara lake basin, western Turkey. *Journal of Field Archaeology*, 40(4), 428–449. doi: 10.1179/2042458215Y.0000000009.
- Luke, C., & Roosevelt, C. H. (2017). Cup-marks and citadels: Evidence for libation in the second-millennium BCE Marmara lake basin, western Anatolia. *Bulletin of the American Schools of Oriental Research*, 378, 1–23. doi: 10.5615/bullamerschoorie.378.0001.
- McPherron, S. J. P., Gernat, T., & Hublin, J. J. (2009). Structured light scanning for high-resolution documentation of in situ archaeological finds. *Journal of Archaeological Science*, 36(1), 19–24. doi: 10.1016/j.jas.2008.06.028.
- Merlo, S., & Shell, C. A. (2005). Developing a multidimensional GIS framework for archaeological excavations. In S. Dequal (Ed.), *Proceedings CIPA 2005 XX International Symposium: International cooperation to save the world’s cultural heritage: Torino, Italy, 26 September–1 October 2005* (pp. 1–5). Torino, Italy: CIPA Organising Committee.
- Merlo, S. (2016). *Making visible: Three-dimensional GIS in archaeological excavation*. Oxford: British Archaeological Reports.
- Molero, M. (2020). Poisson surface reconstruction python binding. <https://github.com/mmolero/pypoisson>
- National Research Council. (2006). *Learning to think spatially*. Washington, DC: The National Academies Press. doi: 10.17226/11019.
- Nobles, G., Çakırlar, C., & Svetachov, P. (2019). Bonify 1.0: Evaluating virtual reference collections in teaching and research. *Archaeological and Anthropological Sciences*, 11(8), 5705–5716. doi: 10.1007/s12520-019-00898-1.
- Open Data Charter. (2019). Who we are. <https://opendatacharter.net/>
- Orengo, H. A. (2013). Combining terrestrial stereophotogrammetry, DGPS and GIS-based 3D voxel modelling in the volumetric recording of archaeological features. *ISPRS J. Photogrammetry Remote Sens*, 76, 49–55. doi: 10.1016/j.isprsjprs.2012.07.005.
- Penninga, F. (2008). *3D Topography: A Simplicial Complex-Based Solution in a Spatial DBMS* [PhD thesis]. Delft University of Technology, Delft.
- Perry, S., & Taylor, J. S. (2018). Theorising the digital: A call to action for the archaeological community. In M. Matsumoto & E. Uleberg (Eds.), *CAA2016: Oceans of data: Proceedings of the 44th conference on computer applications and quantitative methods in archaeology* (pp. 11–22). Oxford: Archaeopress.
- Picavet, V., Brédif, M., Konini, M., & Devaux, A. (2016). iTowns, framework web pour la donnée géographique 3D. *Revue XYZ* (Association Française de Topographie), 147, 49–52.
- Poggi, G., & Buono, M. (2018). Enhancing archaeological interpretation with volume calculations. An integrated method of 3D recording and modeling. In M. Matsumoto & E. Uleberg (Eds.), *CAA2016: Oceans of data proceedings of the 44th conference on computer applications and quantitative methods in archaeology* (pp. 457–470). Oxford: Archaeopress.
- Ponchio, F., & Dellepiane, M. (2015). Fast decompression for web-based view-dependent 3D rendering. In J. Jia (Ed.), *Web3D 2015. proceedings of the 20th international conference on 3D web technology* (pp. 199–207). New York: ACM Special Interest Group on Computer Graphics and Interactive Techniques.
- Potree. (2020). WebGL point cloud viewer for large datasets. <https://github.com/potree/potree/>
- Poux, F., & Billen, R. (2019). Voxel-based 3D point cloud semantic segmentation: Unsupervised geometric and relationship featuring vs deep learning methods. *ISPRS International Journal of Geo-Information*, 8(5), 213. doi: 10.3390/ijgi8050213.
- Richards-Rissetto, H., Remondino, F., Agugiaro, G., Von Schwerin, J., Robertsson, J., & Girardi, G. (2012). Kinect and 3D GIS in archaeology. In G. Guidi (Ed.), *2012 18th international conference on virtual systems and Multimedia, Milan, 2012* (pp. 331–337). Piscataway, NJ: IEEE. doi: 10.1109/VSM.2012.6365942.
- Rolland, J. (2014). *A GeoSpatial World*. <http://ageoguy.blogspot.com/>
- Roosevelt, C. H., Cobb, P., Moss, E., Olson, B. R., & Ünlüsoy, S. (2015). Excavation is destruction digitization: Advances in archaeological practice. *Journal of Field Archaeology*, 40, 325–246.
- Roosevelt, C. H., & Luke, C. (2017). The story of a forgotten kingdom? Survey archaeology and the historical geography of central western Anatolia in the second millennium BC. *European Journal of Archaeology*, 20(1), 120–147. doi: 10.1017/ea.2016.2.
- Roosevelt, C. H., Luke, C., Ünlüsoy, S., Çakırlar, C., Marston, J. M., O’Grady, C. R., ... Slim, F. (2018). Exploring space, economy, and interregional interaction at a second-millennium B.C.E. citadel in central western Anatolia: 2014–2017 research at Kaymakçı. *American Journal of Archaeology*, 122(4), 645–688. doi: 10.3764/aja.122.4.0645.
- Roosevelt, C.H., Kaner, T., & Luke, C. (2020). Kaymakçı Arkeoloji Projesi: 2018 yılı kazı ve araştırma sonuçları. *Kazı Sonuçları Toplantısı*, 41(1), 437–459.

- Roosevelt, C. H., Ünlüsoy, S., & Luke, C. (2019). Kaymakçı Arkeoloji Projesi: 2016–2017 yılı kazı ve araştırma sonuçları. *Kazı Sonuçları Toplantısı*, 40(1), 487–504.
- Scott, C. B., Roosevelt, C. H., Nobles, G. R., & Luke, C. (2021). Born-digital logistics: Impacts of 3D recording on archaeological workflow, training, and interpretation. *Open Archaeology*, 7(1), 574–588. doi: 10.1515/opar-2020-0150.
- Schubert, L., Predoi, A., & Jeffery, K. (2018). Interpolating from stratigraphy from indirect information. In M. Matsumoto & E. Uleberg (Eds.), *CAA2016: Oceans of data: Proceedings of the 44th conference on computer applications and quantitative methods in archaeology* (pp. 185–196). Oxford: Archaeopress.
- Sinton, D. S. (2011). Spatial thinking. In J. Stoltman (Ed.), *21st century geography: A reference handbook* (pp. 733–744). Thousand Oaks, CA: SAGE Publications.
- Tosteberg, P. (2017). *Semantic segmentation of point clouds using deep learning* [MA thesis]. Linköping University, Sweden.
- Tse, R. O. C., & Gold, C. M. (2004). TIN meets CAD: Extending the TIN concept in GIS. *Future Generation Computer Systems*, 20(7), 1171–1184. doi: 10.1016/j.future.2003.11.007.
- van Leusen, P. M., & Nobles, G. (2018). 3D spatial analysis: The road ahead. In M. Matsumoto & E. Uleberg (Eds.), *CAA2016: Oceans of data: Proceedings of the 44th conference on computer applications and quantitative methods in archaeology* (pp. 471–478). Oxford: Archaeopress.
- van Riel, S. (2016). *Exploring the use of 3D GIS as an analytical tool in archaeological excavation practice* [MA thesis]. Lund University, Lund. doi: 10.13140/RG.2.1.4738.2643.
- Varazzo, D. (2011). *psycopg2*. <http://initd.org/psycopg/>
- Verhagen, P. (2018). Spatial analysis in archaeology: Moving into new territories. In C. Siart, M. Forbriger, & O. Bubenzer (Eds.), *Digital geoarchaeology. New techniques for interdisciplinary human-environmental research* (pp. 11–25). New York: Springer. doi: 10.1007/978-3-319-25316-9.
- Verhoeven, G. J. (2017). Mesh is more – Using all geometric dimensions for the archaeological analysis and interpretative mapping of 3D surfaces. *Journal of Archaeological Method and Theory*, 24(4), 999–1033. doi: 10.1007/s10816-016-9305-z.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., ... SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nature Methods*, 17(3), 261–272.
- Wilkinson, M. D., Dumontier, M., Aalbersberg, IJ. J., Appleton, G., Axton, M., Baak, A., ... Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*, 3, 160018. doi: 10.1038/sdata.2016.18.
- Willis, M. D., Koenig, C. W., & Black, S. L. (2016). Archaeological 3D mapping: The structure from motion revolution. *Journal of Texas Archaeology and History*, 3, 1–36. doi: 10.21112/ita.2016.1.110.
- Zaitceva, O., Vavulin, M., Pushkarev, A., & Vodyasov, E. (2016). Photogrammetry: From field recording to museum presentation (Timiryazev Burial Site, Western Siberia). *Mediterranean Archaeology and Archaeometry*, 16(5 Special Issue), 97–103. doi: 10.5281/zenodo.204982.
- Zhou, Q., Park, J., & Koltin, V. (2018). Open 3D: A modern library for (3D) data processing. *arXiv*, 1–6. <https://arxiv.org/abs/1801.09847v1>
- Zollhöfer, M., Siegl, C., Vetter, M., Dreyer, B., Stamminger, M., Aybek, S., & Bauer, F. (2015). Low-cost real-time 3D reconstruction of large-scale excavation sites. *Journal on Computing and Cultural Heritage November*, 9(1), 1–10. doi: 10.1145/2770877.
- Zubrow, E. B. W. (2006). Digital archaeology: “A historical context.” In T. L. Evans & P. Daly (Eds.), *Digital Archaeology: Bridging method and theory* (pp. 8–26). United Kingdom: Routledge.