Stratigraphic Modeling and 3D Spatial Analysis Using Photogrammetry and Octree Spatial Decomposition

Hartmut Tschauner[1] and Viviana Siveroni Salinas[2]

[1]Department of Anthropology
Binghamton University
Binghamton, NY, USA
[2]Department of Anthropology
University of Pittsburgh
Pittsburgh, PA, USA
hartmut.tschauner@gmail.com

Abstract

The destructive nature of archaeological excavation makes the complete documentation of excavated evidence critically important to our ability to evaluate archaeological knowledge claims. Long demanded by textbooks, such complete documentation and publication have only recently become practical through the use of digital spatial technologies. These technologies open an entirely new dimension by making published evidence susceptible to independent re-analyses by the archaeological community. However, conventional stratigraphic recording methods are highly selective, and 2.5D GIS data structures do not support modeling and true 3D spatial analysis of stratigraphic volume solids. This paper presents a stratigraphic recording strategy based on oblique digital photogrammetry that is simple and inexpensive, yet comprehensive and accurate, and discusses approaches to solid modeling and 3D spatial analysis that meet the needs of archaeological stratigraphic analysis. These approaches are implemented in a custom software component under development and will be made available to the archaeological community at nominal cost.

1 Introduction

Spatial data management technologies have been available for decades and, while still far from universally adopted in archaeology, have been in use by individual practitioners for a long time. Archaeology's traditional emphasis on graphical documentation has always implicitly recognized the primacy of the spatial dimension of the archaeological record: time is derived from space (stratigraphy) and our understanding of past cultural systems is based on the spatial relations between features and artifacts (context). However, archaeologists have been slow in adapting the general-purpose spatial data management tools offered by GIS to the core needs of their field. Along with these tools, they have often imported analytical frameworks and approaches from ancillary fields that were earlier adopters of GIS, thereby giving archaeological GIS the status of a specialty subfield closely allied with environmental studies rather than a basic tool of use to every professional engaged in collecting, managing, and interpreting primary archaeological data.

One core area of archaeology—arguably its defining feature in the public eye—that cries out for customized spatial technologies is the field collection, analysis, storage, and dissemination of primary excavation data. Most archaeological uses of GIS have skipped to "sexier" analytical applications. The availability of data appropriate for such analyses has been taken for granted. Moreover, borrowing models from other fields and reflecting the capabilities of commercially available GIS products, the overwhelming majority of applications are flatly two-dimensional (pun intended), and regional analyses far outnumber intra-site studies (for an arbitrary recent selection, see Anaya Hernández, et al. 2003; Beck, et al. 2001; Cummings and Whittle 2003; De Silva and Pizziolo 2001; Goings 2003; Grau Mira 2003;

Holcomb 2001; Jennings and Craig 2001; Ladefoged et al. 2003; Llobera 2001; Stancic and Veljanovski 2000).

However, when existing paper-based records are digitized, the ability of GIS to integrate layers of information from many sources brings previously ignored accuracy issues to the fore. Strings and line levels are wholly inadequate recording "technologies" for the digital age in which graphical documentation for the first time can be much more than pretty pictures. Building spatial databases for formal analysis from even the seemingly most thorough, exhaustive, and lavishly illustrated archaeological reports tends to be a sobering experience. Moreover, traditional, paper-based excavation records offer notoriously selective and inaccessible collections of spatial data that do not support formal analyses of the spatial dimension of the archaeological record. In particular, profile drawings of stratigraphic sections at the edges of excavation units provide a spotty and arbitrary sample of a site's stratification that may or may not be representative of an entire unit. Save a few spot elevations on top of "significant" features, between sections the vertical dimension goes essentially unrecorded. Thus, although probably no other activity consumes more field time than the graphical documentation of spatial information, that information is notoriously incomplete, and it is buried in a paper medium that deprives it of the essential quality of "data," that is, being susceptible to analysis, both by its collector and other researchers, linked to a particular theoretical question or framework. These spatial pseudo-data, along with numerous domains of descriptive (attribute) and image information, are collected separately in multiple places and formats—a hodgepodge of lists, forms, descriptive texts, drawings, and photographs—which hinders data integration

and holistic cross-examination of all classes of evidence. This poor quality and inaccessibility of spatial data is quite disturbing in a field in which the answer to just about every research question has at least a spatial component.

To make matters worse, in many parts of the world an ever smaller fraction of this severely limited record of archaeological spatial information collected at the source is eventually published, let alone at a usable scale, given the cost of printing oversized plans. This publication bias sends a clear message that cannot fail to have a feedback on the standards of field data collection: the academy rewards creative interpretation rather than "mere" collection of evidence, your conclusions are more important than the quality of the evidence that supports them. In a discipline that inevitably destroys its evidence in the process of studying it, this creates a real and present danger of archaeology devolving into a pre-scientific state of knowledge claims largely being justified by an author's personal credentials, as well as the arguments' logical consistency, and hardly being backed by solid evidence *that others can challenge* (cf. Barker 1993:13-14). As Barker puts it, "[i]f we misread our documents as we destroy them, the primary evidence we offer to those interested in the past will be wrong and those following us will be misled but will have no way of knowing it." All too often, field research has become a rite of passage, a procedural step required to lend credence to conclusions that the research proposal has drawn beforehand, without the primary evidence supporting those conclusions ever being laid out for public scrutiny. Before this backdrop, the universal use of digital spatial technologies in the field, both for more complete collection of more accurate spatial data and as a prerequisite for their electronic publication in full and in a readily analyzable format, takes on enormous urgency, not as a mere technical improvement, but as a matter of significant theoretical repercussions.

As a class of spatial data, units of archaeological stratification are sufficiently different from those of other disciplines to require specialized modeling approaches that are not part of the standard repertoire of commercial GIS and need to be adapted from existing methodologies, mostly from geology, or developed from scratch. Like geological strata, units of archaeological stratification are volume solids and ought to be modeled as such. Unlike geological strata, archaeological units are typically not interpolated between sparse sampling locations (boreholes), but fully exposed by horizontal excavation. Current GIS data structures cannot model solid volume entities; they are at best 2.5 dimensional, with an elevation attribute tagged onto two-dimensional features and all spatial analysis taking place in a horizontal plane. To these models, the superimposed strata in a typical archaeological excavation unit are spatially identical, as they all cover the same horizontal extent. The few geological and groundwater-modeling packages that have pioneered a voxel approach to solid modeling are geared towards three-dimensional interpolation from sparse borehole samples. They have trouble creating accurate stratigraphic units from fully exposed and mapped interfaces and handling the complexity of all but the simplest cases of archaeological stratification. Moreover, they do not facilitate attaching unlimited attribute data to spatial units,

which is critical to archaeology with its myriad classes of artifact, geological, and biological data tied to excavation units. Finally, geological and groundwater-modeling packages tend to be exceedingly expensive, aimed at deep-pocket target markets in the oil and gas exploration and utilities industries.

Given this situation on the software market, it is not surprising that published applications of GIS to archaeological intra-site analysis and excavation mostly deal with architecture or horizontal distributions of artifacts susceptible to spatial analysis in two dimensions (e.g., Cutting 2002; Buck, et al. 2003; Craig 2000; Dawson 2003; Fronza et al. 2001; Green et al. 2002; Levy et al. 2002; Peretto et al. 2001; O'Halloran and Spennemann 2002; Pugh 2003; Valenti 1998). Vertical, stratigraphic, and truly three-dimensional (3D) analyses are rare (e.g., Nigro et al. 2003; Spikins et al. 2002).

The present paper addresses this gap in recording strategies and spatial data management software, presenting an effective, affordable, integrated solution for archaeological stratigraphic modeling and 3D spatial analysis. This solution is composed of a field recording strategy based on off-the-shelf oblique photogrammetry software and a custom modeling and 3D spatial analysis component under development that uses an octree spatial decomposition approach to solid modeling of stratigraphic units and is integrated with relational database management systems for optimal joining of spatial and attribute queries. The bulk of the following pages will be dedicated to describing the field recording method, using real-world examples from our excavations at the site of Huayuri on the south coast of Peru (Siveroni et al. 2004). This will be followed by a discussion of the requirements for a specifically archaeological software package for the management and 3D spatial analysis of stratigraphic volume solids, as well as the novel queries that such a system will facilitate. I will close by offering a sketch of the architecture and capabilities of the open-source stratigraphic solid modeling software that is under development by the author and slated for initial release in 2007.

2 Stratigraphic Recording: Contact Topography using Digital Photogrammetry

Citing the destructive nature of excavation, archaeological textbooks have always admonished excavators to record and publish the full 3D structure of excavated sites. In the paper-and-pencil era, this was actually an unrealistic reporting standard. Prior to the advent of digital mapping technologies, stratigraphic recording through section drawings at ultimately arbitrary locations (edges of excavation units, balks) and spot elevations was by necessity selective and simply did not produce the data required for full 3D models. Moreover, the communicative power of such models, when printed on paper, would have been quite limited anyway. Today's digital spatial technologies, however, make it fairly painless to collect, store, and publish the necessary information. These technologies finally enable us to live up to what has always been the textbook standard of archaeological reporting.

To qualify as "painless," a data collection technology must not only be simple, reliably accurate, and efficient to use under field conditions, but also affordable within the budgetary constraints faced by the average archaeological project, including projects based in developing countries and graduate student dissertation research. Otherwise, technology will not be of much help in counteracting the dangerous tendencies discussed in the introduction. On purely technical grounds, 3D laser scanning clearly has great potential as a stratigraphic recording technology, but its sticker price puts it out of reach of all but the wealthiest archaeological projects and institutions. Equipment prices will have to come down by at least 90% before 3D laser scanning is ready to go mainstream in archaeology, and that is not likely to happen any time soon. In the meantime, oblique photogrammetry offers a productive and cost-effective alternative with a spatial resolution that, while significantly lower than that of 3D laser scanning, is sufficient for most archaeological purposes and far superior to traditional profile drawings and spot elevations. While photogrammetry has been used in archaeology for quite some time (e.g., Anderson 1982; Fussell 1982; Grün et al. 2003; Meyera et al. 2006; R$uuther 1998; Sauerbier and Grün 2003; and numerous CIPA papers (see http://cipa.icomos.org/)), to our knowledge it has not been previously employed as a systematic and comprehensive stratigraphic recording device.

2.1 Contact Topography

Although we ultimately want to build solid volume models, the most economical way to do so, both from the fieldwork and software perspectives, is surface-based contact topography, i.e., microtopographic maps of the upper interfaces of all units of stratification. Standard 2.5D GIS software is able to create and display the interfaces; specialty software subsequently creates solid volumes sandwiched in between. Cut features, such as pits, postholes, etc., are treated the same way as layers. Their top interfaces are outlined and mapped on the surface from which they were cut; their bottom interfaces are mapped after the features have been excavated. Deep and narrow features may require multiple outlines on each surface they penetrate. A 3D model is subsequently

composed from these slices, using a "tomographic" approach not unlike computer tomography in medical imaging.

As a spatial data collection task, stratigraphic recording has high accuracy requirements, particularly in the vertical dimension, which is the hardest to measure with accuracy. Layers of less than 1 cm in thickness may represent meaningful stratigraphic events and therefore need to be distinguished. At the same time, systematic and comprehensive interface mapping of entire sites produces a substantial volume of spatial data, typically in the hundreds of interfaces and hundreds or thousands of polygon features each season. For example, the modest-sized excavations (ca. 270 m²) at the Late Intermediate Period site of Huayuri on the south coast of Peru (Siveroni et al. 2004) so far have produced over 1,300 interface maps (Figure 1). Huayuri is a dense residential site with a complex sequence of superimposed occupations that have left a multitude of often intricately nested features posing a significant challenge to comprehensive three-dimensional recording (Figure 2).

2.2 Oblique Digital Photogrammetry

Oblique digital photogrammetry has proven itself capable of coping with this challenge, both in terms of accuracy and fieldwork efficiency, generating 2.5D surface models, polygon features, and orthophotos in one quick and simple process. 2.5D surface models are extracted from sets of
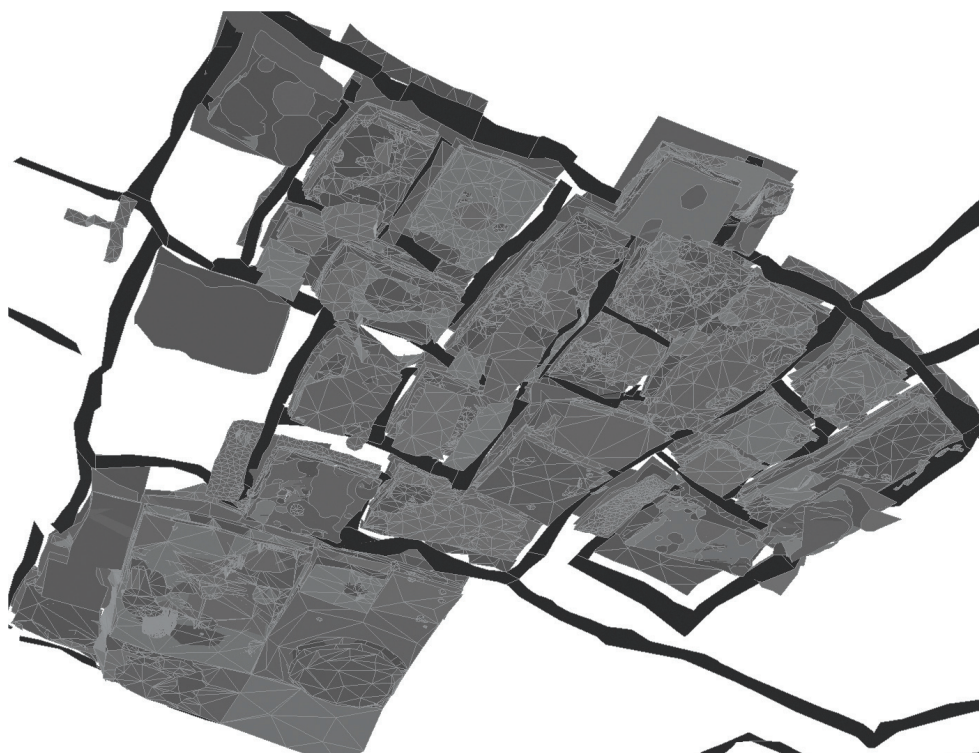

Figure 1. Stratigraphic features at the Late Intermediate Period site of Huayuri, Santa Cruz Valley, south coast of Peru. Some 1,300 stratigraphic units have been recorded in an excavated area of ca. 270 m².

Figure 2. Abundance of cut features in C3R19 at Huayuri. Photogrammetric recording has proven quite capable of handling situations as complex as this one.
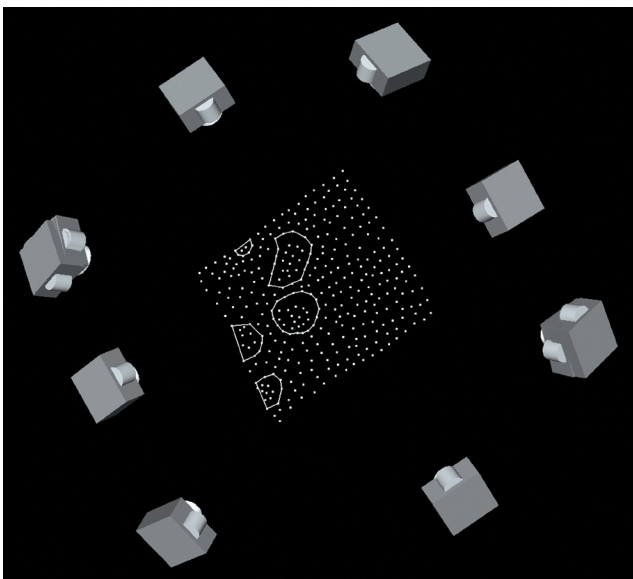

Figure 3. Typical camera positions used in recording a stratigraphic interface. The eight positions virtually assure that all significant points will show up on at least two, usually on three or more, images with strong (i.e., near right) angles between them.

ground-referenced, oblique digital photographs taken from several angles. Photogrammetry software designed for reverse engineering, forensics and accident-scene reconstruction, and architectural applications is relatively inexpensive and appropriate for the scale of most excavation units. The following discussion is based on my experiences with PhotoModeler Pro 5.2 from EOS Systems (cf. Green et al. 2002).

Photographs to be processed in PhotoModeler may be taken with standard digital cameras. These need to be individually calibrated, but calibration for small to medium-sized target areas is a simple process that can be performed in-house. A standardized dot pattern known to the software is photographed from various angles and by resolving the photogrammetric equations in reverse, any deviations from the known measurements of the pattern are used to determine the relevant parameters of the camera-lens system. This process needs to be repeated for every camera, lens, and zoom level/focal distance. It is also advisable to create separate calibrations for different sizes of areas to be measured in the field, as the size of the calibration pattern ought to be roughly the same as that of the objects to be measured. This imposes a practical limit on the size of area that can be measured. Although the calibration dot pattern may be projected at any size against a screen or clean wall, reasonably controlled lighting conditions will only be available indoors and rich collections of graffiti may make it hard to find a suitably clean wall on a university campus. However, this problem is more apparent than real. Given available camera resolutions, the level of detail desired of field photography, and the limited distances from the target attainable through ladders or photo towers, it is preferable to cover larger excavation units with multiple, partially overlapping sets of photographs. These sets are easily stitched together when processing the imagery, and the extra time for field photography is minimal.

In the field, photographs for photogrammetric stratigraphic recording are best taken at the end of a standard field photography session for a newly exposed interface, adding just a couple of extra minutes to the session. Additional special photogrammetry sessions may be required as individual features are exposed. Taking these pictures hardly differs from standard archaeological field photography. The camera may be hand held or mounted on a tripod or pole. Photographs are taken at a vertical angle of approximately 45º; precise control of this angle is not required. All photographs in a set need to be taken with the same lens, focal distance/zoom level, and from roughly the same distance. For this reason, fixed lenses are safer than zoom lenses, and autofocus must be turned off. For maximum accuracy, the area of interest should cover as much of the camera's field of vision as possible. To accomplish optimal coverage, it may be necessary to switch between landscape and portrait modes (or anything in between) within the same set of photographs. This will not cause any problems in processing the image sets.

Every significant point needs to show up on at least two ideally epipolar images, i.e., photographs taken from positions at right angles to each other with respect to the target, which provide the strongest basis for measuring 3D positions. Two photographs are the minimum to determine a point's 3D position. A third image will add a redundant measurement that is helpful in detecting blunders. Field experience has shown that taking pictures from eight positions

around a unit is just the right level of redundancy (Figure 3). It takes only a couple of minutes and virtually guarantees that even at the bottom of an excavation unit all relevant points will show up on at least two near-epipolar images.

2.3 Targets

The key to a successfully employing photogrammetry as a stratigraphic recording tool is the placement of standardized, high-contrast targets on the surfaces to be recorded. Artificial targets are crucial because archaeological interfaces tend to be of irregular shapes and seldom offer any sharply defined points that might be easily and precisely matched on several photographs. Moreover, assuming a reasonable amount of background noise, PhotoModeler can automatically recognize and reference such targets (particularly circular ones) on sets of photographs of the same scene with an accuracy of less than one pixel, which is not achievable by a human operator marking "natural" points. Since the identification of circular targets is based on a sphericity index, the targets need to have a minimum diameter of 3-5 pixels on each image, thus requiring different target sizes for different excavation areas.

The density of target placement determines the spatial resolution of the resulting model. At Huayuri, we place about 500 targets on each interface of 6 to 10 m² in 2D surface area (Figure 4). Targets are arranged to form a dense, more or less regular grid, with additional ones placed wherever surface detail needs to be worked out. This density is evidently a far cry from point clouds produced by 3D laser scanning, but the resulting models are quite detailed and even aesthetically pleasing without any manual touch-up. Extremely accidented surfaces covered with rubble or wall fall, particularly if there are vertical rock faces, may be the one exception from this rule. However, it is possible to stick targets onto the vertical faces using reusable adhesive or pins (for soft materials).

In addition to providing surface mass points, targets also serve to outline features on a surface, such as the mouth of a pit, a lens, or an artifact concentration (Figure 4). The resulting points are later easily connected—not unlike children's connect-the-dots coloring books—to form 2.5D polygons that represent the features in the excavation's spatial database. A volume model may subsequently be constructed from several such slices through the same feature photographed on successive surfaces cut by that feature. This approach is particularly useful for narrow features such as

Figure 4. Circular targets placed on a stratigraphic interface at Huayuri. Here, 516 targets have been placed in an area of about 6.5 m². The density, and hence the resolution of the model, could be increased with little extra effort. Feature edges have been outlined with circular targets. The resulting 3D points are easily connected to form 2.5D polygons in a GIS database. A full 3D model may be composed from multiple such polygon slices though a feature that may be difficult to capture otherwise, for example a narrow posthole. Some images in a set may be taken with circular targets removed and only coded targets left on the interface. These images are used to produce orthophotos for publication (cf. Figure 9).

postholes, as it may be impossible to place and photograph targets on their bottom interfaces.

The targets placed on the interfaces fall into two categories: (1) at least 6-10, relatively large, coded targets, evenly distributed across the image area and used for orienting the photographs, i.e., determining the camera positions from which the pictures were taken (Figure 4); and (2) targets for automatic marking that produce the mass points defining the surface models (Figure 4). Coded targets for orienting the images need to show up on every image in a set. Referencing them is either a somewhat tedious manual process of marking matching targets on multiple images or, if standardized targets are used, may be automatically performed by the software. For automatically marked mass points, we use circular, retro-reflective targets mounted on chips of heavy plastic material or—at sites with strong winds—heavy, magnetic rubber. Coded targets are mounted on similar, if larger chips. These may, of course, also be used for mass points, but this is not recommended since coded targets are substantially more expensive than the circular ones and their larger area effectively limits the density of target placement. We produce the target chips ourselves from commercially supplied rolls of adhesive target tape. Depending on field conditions and care in handling, these chips may survive multiple field seasons. A more elegant and 100% wind-proof alternative to plastic targets is a target projector, but these devices are costly and do not operate on battery power; thus, a power outlet or generator is needed in the field.

2.4 Georeferencing

Photogrammetric measurements are entirely relative; thus, it is necessary to shoot in at least three coded targets with a total station, preferably more for backup and accuracy checks. PhotoModeler will use these points to perform a least-squares adjustment, properly scaling and rotating the model. The total-station is the most expensive piece of equipment required for this stratigraphic recording method. However, total stations are in common use by archaeologists today. Moreover, since the number of points to be shot is fairly low and no linework is required, even a simple and cheap instrument without fancy data collection, graphical map display, and geometry editing capabilities will be adequate to do the job.

2.5 Problem Areas

At Huayuri, this method of stratigraphic recording has proven robust, productive, and accurate. A few minor problems are caused by excessively bright sunlight, heat, and wind, but there are simple solutions to all of them. Bright sunlight may result in extreme contrasts and the black plastic background of the targets reflecting as much light as the reflective dot (Figure 5). Under these circumstances, automatic target marking will not be effective since target recognition is based on brightness differences (as well as target shape) and it will be impossible to find a contrast setting that

will not blur the target boundaries in some part of an image. Noisy backgrounds with objects whose shapes mimic that of the targets (e.g., pebbles) may have a similar effect.

In these cases, automatic target marking may be performed in separate runs for sections of similar contrast within an image, but sometimes manual target marking may be inevitable. This can be somewhat tedious, but at least it does not add to the field time required for stratigraphic recording, and manual marking unfailingly works. Evidently, this problem is easily prevented by using an awning (Figure 6) or roof, which many excavations in hot areas will have anyway. Excessive heat may also deform or even melt the plastic material of the targets. In hot climates, target exposure to the sun should therefore be kept to a minimum, and the targets must be stored in the shade.

Wind may either move the targets or cover them with dust or sand. At the bottom of an excavation pit, however, this problem is much less acute than might be expected. Where the wind is too strong, targets mounted on magnetic rubber or even heavier materials will be the answer. Despite these minor issues, not a single image set—including our very first experiments—has ever failed to process.

Figure 5. Excessively bright sunlight may make the black plastic background of some targets reflect as much light as the central reflective dot and cause extreme brightness differences between different sections of the image. Multiple automatic marking runs or even manual target marking may be required. This problem is best avoided by using an awning or roof (see Figure 6).

2.6 Productivity

Photogrammetry is not only a robust method of stratigraphic recording but also exceptionally productive. The entire recording procedure—placing targets, photography, and shooting reference points—will take about 20-30 minutes per surface. Cleaning the unit prior to photography is only necessary if orthophotos for publication are to be produced, but this will cause no extra work if photogrammetry images are taken as part of regular photo sessions. Additional detail comes virtually for free; placing another 100 or so targets will take no more than a couple of minutes. Around 15-60 minutes of office time are required to post-process each surface, depending on how much manual labor is required. In most cases, the actual time will be closer to the lower end of this range.

In return, we simultaneously obtain detailed 2.5D surface models of layers and features (Figures 7, 1), 2.5D polygons representing features or slices through features (Figure 8), and orthophotos (Figure 9). Thus, this recording method

replaces profile drafting, plan view drafting, and part of field photography—the whole range of graphical documentation, eliminating the need for essentially all hand drafting. The resulting records are far more complete than traditional paper records and immediately available in digital format, ready to enter a GIS database.

2.7 Accuracy Under Field Conditions

The accuracy attainable with photogrammetry under actual field conditions is quite acceptable. From the Huayuri site, we have 81 coded targets whose position was determined both photogrammetrically and by total station and that were not used in the PhotoModeler least-squares adjustments. Thus, these measurements are redundant and may be used to check the accuracy of the photogrammetric against total-station measurements. Photographs were taken with a Canon EOS 20D 8.2 megapixel camera using Canon EF-S 18-55 mm and EF 28-105 mm lenses. The photogrammetrically determined positions of these 81 points differ by 11.6
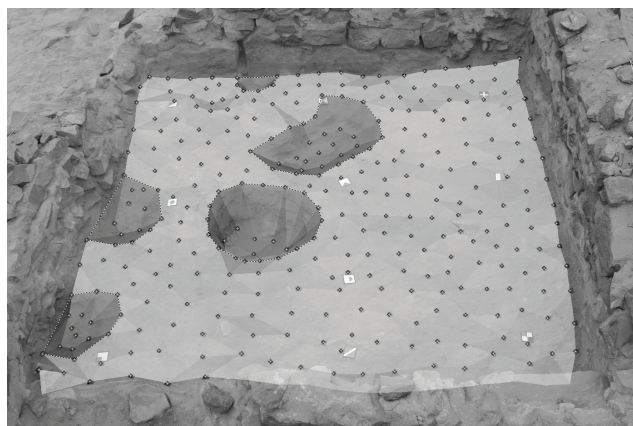

Figure 6. A simple, improvised awning and photo tower go a long way toward avoiding potential problems with photogrammetric recording in the field.

Figure 7. Outputs of photogrammetric stratigraphic recording (1): 2.5D surface projected onto a digital photograph.

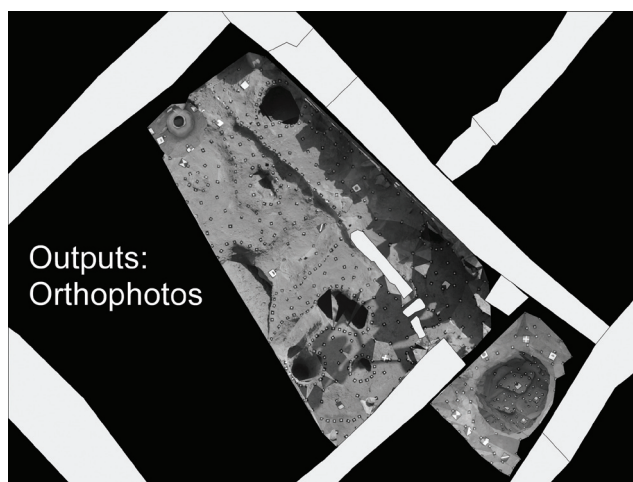*Figure 8. Outputs of photogrammetric stratigraphic recording (2): 2.5D polygons.*

Figure 9. Outputs of photogrammetric stratigraphic recording (3): orthophotos. Circular targets for automatic marking may be removed and only coded targets left in the scene if orthophotos are to be published. See Figure 6 for an example.

± 7.5 mm (median 9.6 mm) in two dimensions and 13.0 ± 7.7 mm (median 11.2 mm) in 3D. If we compare the 339 distances measured between each of the 81 points and the remaining coded targets in the same image set (i.e., on the same surface), the photogrammetrically determined distances on average are off by about 7.7 mm (median 6.4 mm) relative to the total-station measurements (Table 1; Figure 10). This corresponds to an average of about 0.9% of the distance between each pair of points.

While these tolerances are not perfect from a surveyor's standpoint, they are significantly better than what can be expected from even a skilled draftsperson preparing a profile drawing or plan view with strings, line levels, and measuring sticks—the more so the larger the area covered by the drawing. Errors are also much less dependent on individual operator skill, concentration, and patience than with hand drafting. The greatest accuracy gain over profile drafting, however, is at a more fundamental level, insofar as the complete coverage facilitated by this method introduces

a meaningful concept of accuracy of stratigraphic recording at the level of the entire site. Even if a profile drawing were a perfectly accurate rendition of the stratification encountered at the particular unit edge shown, there is simply no way to tell how accurately that profile describes the rest of the unit.

• As a rule, field researchers face a trade-off between the quality and quantity of information collected by a recording method and the amount of field time to be invested in collecting that information. Judging from my experience in coastal Peru, photogrammetric stratigraphic recording is a rare exception from that rule. When leveraged as a replacement of essentially all paper-based spatial recording methods traditionally employed by field archaeologists, photogrammetry is a robust, time-saving device that simultaneously produces vastly more and better-quality information in an appropriate digital format for

Table 1. Difference of distances between redundant coded targets not used in least-squares adjustment of photogrammetry projects, as determined by photogrammetry and by total station (Leica TCR 703). Photogrammetric measurements made in PhotoModeler Pro 5.2 software, based on 8.2-megapixel images taken with a Canon EOS 20D camera and Canon EF-S 18–55 mm and EF 28–105 mm lenses.

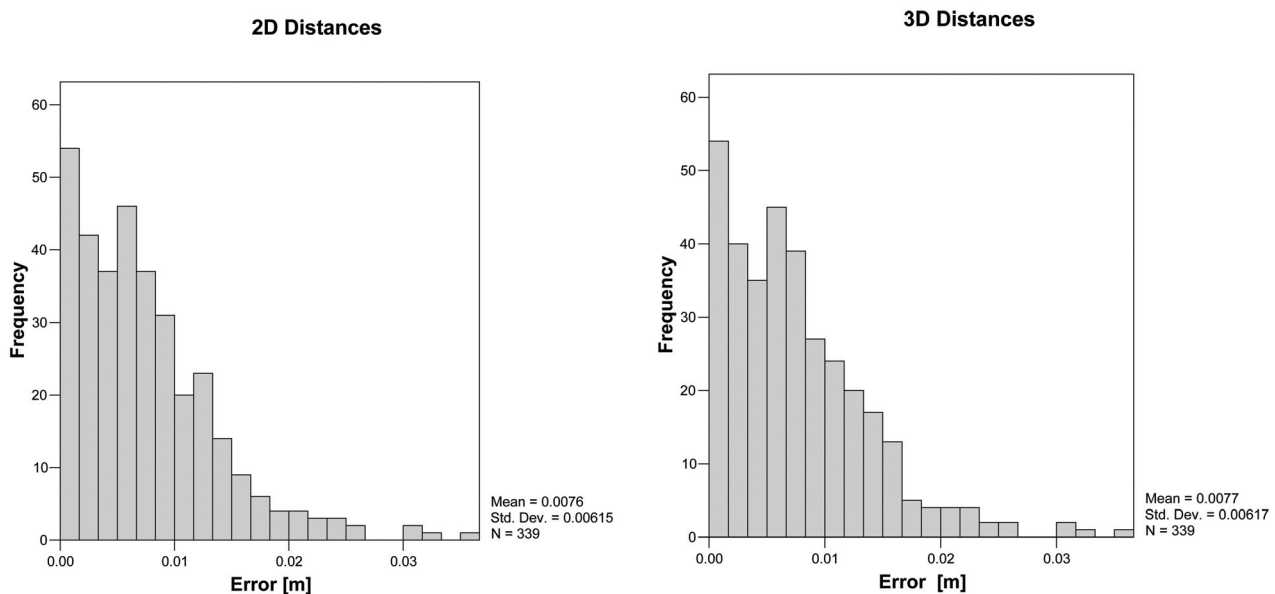| N = 339		2D		3D	
		Statistic	Std. Error	Statistic	Std. Error
Mean		0.0076 m	0.00033 m	0.0077 m	0.00034 m
Mean percent of distance		0.9094 %	0.05772 %	0.9128 %	0.05846 %
95% Confidence	Lower Bound	0.0069 m	0.0070 m	0.0070 m	
Interval for Mean	Upper Bound	0.0082 m	0.0083 m	0.0083 m	

Figure 10. Difference of (a) 2D (left) and (b) 3D distances (right) between redundant coded targets not used in least-squares adjustment of photogrammetry projects, as determined by photogrammetry and by total station (Leica TCR 703). Photogrammetric measurements made in PhotoModeler Pro 5.2 software, based on 8.2-megapixel images taken with a Canon EOS 20D camera and Canon EF-S 18–55 mm and EF 28–105 mm lenses.

storage, analysis, and publication in GIS databases. As a bonus, it is a fairly inexpensive and simple technology that should be within reach of most archaeological projects and institutions.

3 Analyzing Digital Stratigraphic Data: 3D Spatial Analysis

The stacked 2.5D surfaces produced by photogrammetric contact topography make a productive recording strategy and handy visualization device using existing GIS software (rotating in 3D, peeling off layers, etc.). These advances in communicating stratigraphic information are not to be belittled, but the analytical potential of surface models is limited (section generation, volume calculations). While surface models implicitly contain all the necessary information, 2.5D data structures do not support truly 3D spatial analysis, and therefore do not allow us to fully leverage digital stratigraphic data. Just as in 2D GIS, we want to be able to explore the spatial distributions of cultural, biological, geological, geophysical, or geochemical evidences across stratigraphic units through qualitative and quantitative thematic maps of 3D volume solids and their associated attributes. We want to be able to query our stratigraphic units based on their spatial relations and attributes, for example: "Find all the excavated volumes whose density of pottery Type A per unit volume is within the upper quartile of the distribution of densities of Type A for the whole site and that have produced charred wood remains; then produce a table of their 3D distances to the nearest hearth." This query amounts to an attribute-driven intersection of 3D solids accompanied by the corresponding set-theoretic operation on the attributes attached to those solids. We want to be able to use stratigraphic information to test our ideas about the distribution of cultural features. Suppose, for example, that Layer C contains the highest concentration of pottery Type A per unit volume and we suspect that Type A is associated with Layer C, but appreciable amounts of it are also found in other layers. To substantiate our idea, we could demonstrate a statistically significant difference in the per-volume density of Type A between those stratigraphic units that are adjacent to Layer C and those that are not or show a fall-off of this density with increasing distance from Layer C.

The above are just a few examples of questions that 3D spatial analysis may answer. Similar analyses altering attribute sets based on the spatial relations and Boolean operations between the geometry objects described by those attributes are the bread-and-butter issues of two-dimensional GIS, but current 2.5D data structures do not support them in 3D. Therefore, stratigraphic applications, particularly formation-process research, require the introduction of 3D solid modeling into GIS analysis.

3.1 Approaches to Solid Modeling and 3D Spatial Analysis

Archaeological stratigraphic analysis requires a system that will perform the same basic spatial analysis operations on 3D solids as current GIS software does on 2D geometries. Such a system ought to be integrated into a relational database management system (RDBMS) for the utmost flexibility in cross-querying all types of data—spatial, temporal, quantitative, and categorical.

There are two basic approaches to solid modeling in computer graphics. Constructive solid models are built up from geometric primitives such as cubes, spheres, cones, or cylinders, which are combined into more complex shapes using Boolean operations (union, intersection, subtraction), or by rotating 2D vectors. The resulting models are either true solids with mass properties or boundary representations describing only the surface of a solid. This approach, taken by CAD and 3D design software, is geared to the symmetrical shapes of industrial engineering applications. Therefore, it is utterly impractical for modeling archaeological deposits of arbitrary, asymmetrical, and often degenerate shapes. However, the creation of such solids is easily automated in a CAD environment by extruding grid cells between a stratum's upper and lower interfaces into vertical, prismoidal columns and merging the columns into a single complex solid through a union operation (Figure 11). This approach facilitates attaching unlimited attribute data by linking external database records to the solid features and supports spatial analysis through Boolean operations. However, these operations are computationally expensive and in a CAD environment, attributes can only be used as criteria for selecting features for analysis; not only will attributes not be subjected to set-theoretic operations corresponding to the spatial operations performed, but they will simply be dropped.

While constructive solid modeling builds up a model from geometric primitives, spatial decomposition starts from a spatial universe that covers the entire 3D extent of a model and breaks that universe down into smaller volumes. The simplest approach to such partitioning is a voxel (volume pixel) model (Figure 12), the 3D analog of a raster surface. This approach is common in 3D medical imaging (e.g., Ashburner and Friston 2000; Good et al. 2001; Studholme et al. 1997). The simplicity of the data structure is an appealing feature of voxel models, as they map onto a straightforward 3D array of fixed-size cubes that provide a ready unit of analysis for topological overlays through 3D map algebra. However, since the voxel approach divides space into uniformly sized blocks, even in regions of homogeneity, it imposes severe computational and memory overheads. The small voxel size required for accurate shape representations results in the memory requirements of even moderately complex models—on the order of a single excavation unit—quickly outgrowing currently available computer resources.

Clearly, then, it would be desirable to have variably-sized voxels, modeling homogeneous regions by large blocks and complex regions by small blocks, thereby maintaining the essential simplicity of the voxel approach while dramatically reducing its computational overhead. This may be accomplished by combining the spatial partitioning approach with a tree-like data structure. Algorithms such as binary space partitioning or octrees (e.g., Carlbom et al. 1985; Chan and Kwok 1993; Jones 1989; Keim 1999; Navazo et al. 1986)
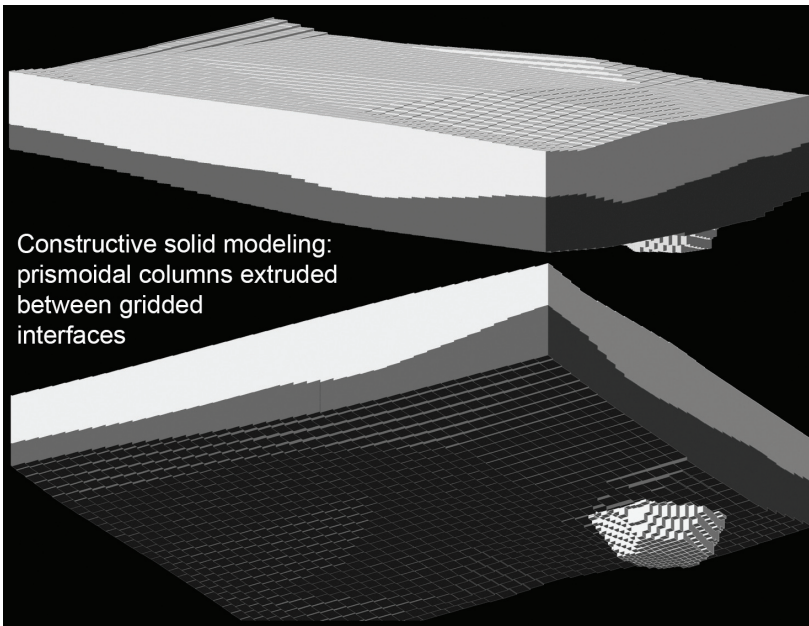
Constructive solid modeling:
prismoidal columns extruded
between gridded
interfaces

Figure 11. Constructive solid modeling of irregular-shaped stratigraphic units: pris-moidal columns extruded between gridded interfaces. This example was created by a VBA macro in AutoCAD®. Higher resolutions of even such simple models will keep the most powerful PC busy for quite some time.
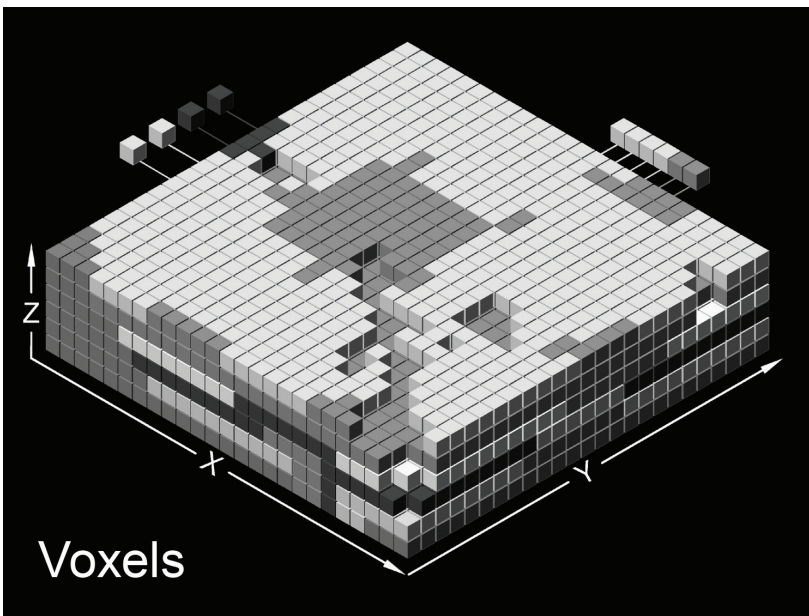

Voxels

Figure 12. Voxel models are the simplest form of spatial decomposition. The model universe is broken down into equal-sized cubes. The simplicity of voxel models is appealing, but since all voxels are of the same size, homogeneous regions must be modeled at the same resolution, resulting in enormous computational and memory overheads.

recursively partition space into nested, convex subspaces, stopping when no further detail is required. This allows them to store abundant detail in regions that need it—primarily in boundary regions—without wasting storage in regions that do not need it.

Octrees (Figure 13) initially partition the model universe into eight cubic subspaces. Where more structure is needed, each cube may be recursively decomposed into another set of eight cubes. The resulting data structure is a tree, each of whose nodes has either eight children or none. The tree is usually held in a doubly linked list. This recursive decomposition of space allows octrees to represent homogeneous and heterogeneous entities at different resolutions in the same model. The more homogeneous the modeled entities, the greater the storage compression relative to voxel models will be. At the same time, much of the topological simplicity of the voxel approach is maintained. In so-called pointer-less octree implementations, each octant is represented by a unique integer locational key or octal code that not only indicates the octant's size and location but also its spatial relations with other octants in the model. The model holds a sorted list of these locational keys (known as a linear octree) that facilitates spatial searches based on computationally cheap integer arithmetic and bit manipulations. Since stratigraphic units are by definition homogeneous regions and our goal is to enable true 3D spatial analysis of complex stratigraphic models, linear octrees are a most attractive approach to stratigraphic modeling.

Besides the sheer numbers of units involved, archaeological stratigraphic modeling is further complicated by the non-convex or even degenerate shapes of many archaeological deposits and volume features. Such shapes may not be adequately captured by an automated tri-angulation based on the 3D convex hull of a cloud of points measured on the units' interfaces. However, using 3D α-shapes, we may still largely automate the construction of solid models, avoiding tedious user inputs of breaklines or boundary polygons and manual additions or deletions of tri-angles. α-shapes (Bernardini and Bajaj 1997; Edelsbrunner 1992; Edelsbrunner and Mucke 1990, 1994) are a generalization of the convex hull that formalizes the intuitive notion of the "shape" described by a point cloud. In contrast to the convex hull, they may be concave or disconnected or even include points and lines. Edelsbrunner and Mucke (1994) compare the process of creating α-shapes for a point cloud to scooping ice cream from a bowl without touching the chocolate chips embedded in the ice cream. The parameter α is the squared radius of the scoop. If the radius is small, we will be able to scoop up most of the ice cream, eventually leaving only the chocolate chips. Thus, for small values of α (α→0), the α-shape will be the point set itself. Conversely, a large scoop will not fit between any pair of chocolate chips, preventing

us from reaching the ice cream in the space enclosed by the chips. Thus, the α-shape for large values of α (α→∞) is the convex hull of the point set. While there is a whole family of α-shapes for any point set, capturing its shape at different levels of detail, the α-shape for a given (user-provided) value of α is mathematically well defined and non-arbitrary.

3.2 Software Development: Solid 3D GIS for Stratigraphic Analysis

The senior author is currently developing a stratigraphic solid modeling and 3D spatial analysis application, using a linear octree approach and 3D α-shapes. The application builds on two software libraries that handle the complex underlying computational geometry and data storage issues and allow him to focus on custom-tailored archaeological functionality. The OctSolid library (by J&L Associates) covers the basic operations of creating, compressing, maintaining, and analyzing (spatial searches, Boolean operations) linear octrees; the Computational Geometry Algorithms Library (CGAL), an open-source software library written and maintained by a consortium of European and Israeli institutions, includes an implementation of 3D α-shapes.

Since the target audience is expected to use commercial GIS and CAD-based civil engineering programs for collecting and maintaining their data, the application is written as an extension to commercial software (Figure 14). The initially targeted GIS platform is ESRI's ArcGIS, not only because it is the most widespread, but also because it has an excellent 3D viewer (ArcScene) that is capitalized on for displaying and thematically symbolizing boundary representations of solid models, eliminating the need of writing a display system.

The following features are planned to be included in the application:

1. Round-trip conversion between linear octrees and 2.5D GIS features (points, polygons, multipatches, TINs) and raster surfaces.
2. Import of voxel models from remote sensing and geological applications.
3. 3D interpolation from sparse sample points (e.g., auger, geochemical sampling).
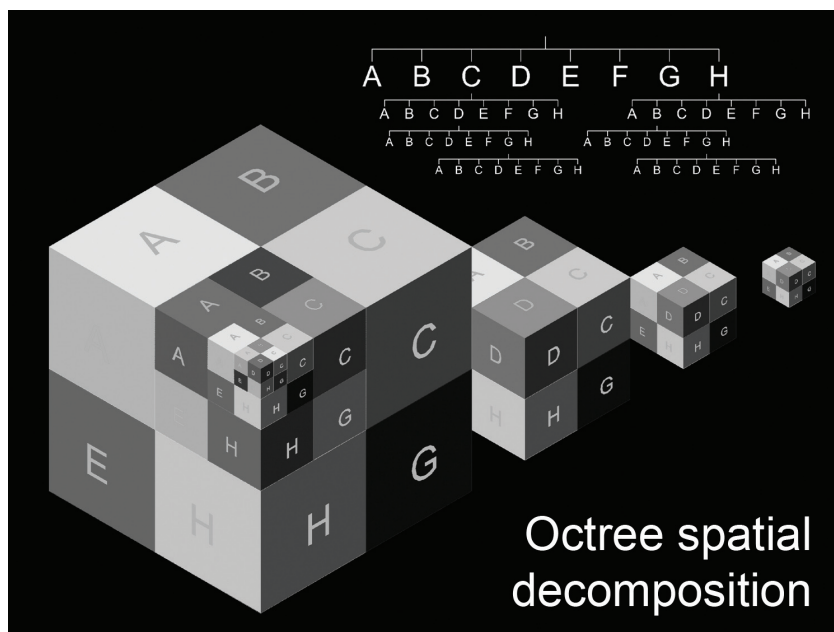4. Associating attribute data with octants and storing

Figure 13. Octree spatial decomposition. The model universe is broken down into 8 equal-sized blocks each of which may be recursively partitioned into 8 children if additional detail is required. This allows an octree model to have multiple resolutions, while preserving much of the simplicity of the voxel approach.
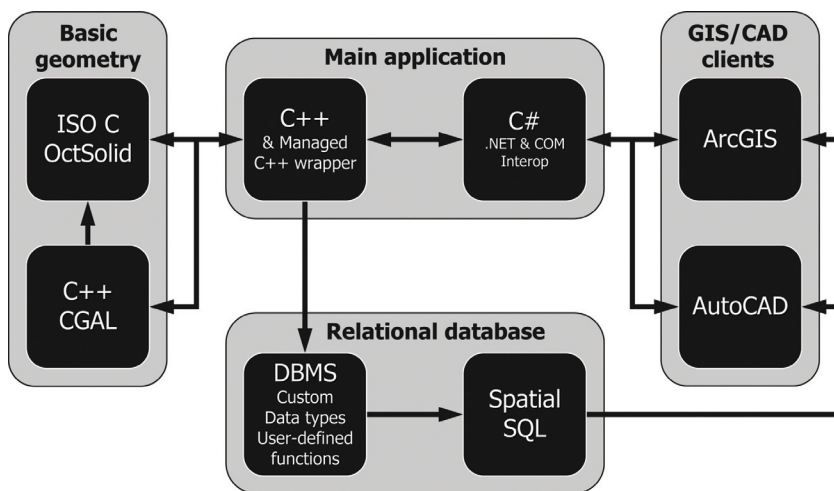

Figure 14. Architecture of the archaeological solid modeling and 3D spatial analysis application under development. Two software libraries handle the complex underlying computational geometry and data storage issues. The OctSolid library (by J&L Associates) covers the basic operations of creating, compressing, maintaining, and analyzing linear octrees; the Computational Geometry Algorithms Library (CGAL), an open-source software library written and maintained by a consortium of European and Israeli institutions, includes an implementation of 3D α-shapes.

octants as a custom data type in at least one major RDMBS. The obvious choices are Oracle, DB2, or PostgreSQL/PostGIS, which already have native spatial data types and 2D spatial analysis capabilities.

5. Spatial indexing of octant data is accomplished by indexing on the integer octal code, which encapsulates an octant's position in the tree (cf. Kunii et al.

1986; Mao et al. 1992).

6. Spatial querying and topological overlays (Boolean operations): union, intersection, pair-wise difference, adjacency, distance. Transfer of attribute data to output solids using transfer rules (e.g., copy, equal or proportional division, sum, mean, maximum, minimum).

7. Implementation of spatial search and overlay operations as spatial SQL functions supporting mixed spatial and aspatial queries in at least one major RDBMS.

8. Display of boundary representations of models and analysis outputs in commercial 3D viewers (initially ESRI ArcScene and possibly AutoCAD).

9. Volume and surface area calculations.

10. Sections and fence diagrams at arbitrary locations in the site model, output as 2.5D GIS features.

A first release that may not include all of the above functionality is planned for 2008. It will be made available to the archaeological community at nominal or no charge.

4 Conclusion

Stratigraphic information has been an underreported and underutilized class of archaeological data in the paper-and-pencil era. Digital recording, storage, and publication enable us to dramatically raise the standards of excavation reporting, without investing more time and money in data collection and dissemination. This means much more than an increase in the amount and accuracy of information recorded and published; it represents a qualitative change in the nature and use of spatial information in archaeology. For the first time, the spatial dimension of the archaeological record may be made available in a format that is open to formal analysis and puts our colleagues in a position to challenge the evidence on which our knowledge claims are based. We will only reap these benefits, however, if we go digital all the way, from data collection in the field to analysis, storage, and publication, and if we start producing our own, specifically archaeological analysis tools. This paper has taken some steps in this direction. Little would be gained from a massive improvement of field data collection if the data end up languishing in "dead" paper records or go entirely unpublished. Little would be gained from digital publishing if we waist our resources digitizing incomplete and less than accurate paper records. Given the rapid depletion of archaeological resources in the ground, archaeologists have an obligation to use all available means to provide future generations with a basis on which to form their own views of the human past. Today, these means emphatically are technological. No excavation ought to be regarded as properly published unless it is documented in a comprehensive, publicly available, three-dimensional spatial database.

References Cited

Anaya Hernández, Armando, Guenter, S. P., Zender, M. U. 2003. Sak Tz'i', a Classic Maya center: a locational model based on GIS and epigraphy. *Latin American Antiquity* 14(2):179-191.

Anderson, Richard C. 1982. Photogrammetry: the pros and cons for archaeology. *World Archaeology* 14(2):200-205.

Ashburner, John, Friston, Karl J. 2000. Voxel-based morphometry—the methods. *NeuroImage* 11:805-821.

Barker, Philip. 1993. *Techniques of Archaeological Excavation*. 3rd ed. London: Batsford.

Beck, C., Taylor, A. K., Jones, G. T., Fadem, C. M., Cook, C. R., Millward, S. A. 2001. Rocks are heavy: transport costs and Paleoarchaic quarry behavior in the Great Basin. *Journal of Anthropological Archaeology* 21(4):481-507.

Bernardini, F., Bajaj, C. 1997. *Sampling and Reconstructing Manifolds Using Alpha-Shapes*. Technical Report CSD-TR-97-013. West Lafayette, Indiana: Department of Computer Science, Purdue University.

Buck, P. E., Sabol, D. E., Gillespie, A. R. 2003. Sub-pixel artifact detection using remote sensing. *Journal of Archaeological Science* 30(8):973-989.

Carlbom, I., Chakravarty, I., Vanderschel, D. 1985. Hierarchical data structure for representing the spatial decomposition of 3-D objects. *IEEE Computer Graphics and Applications* 5(4):24-31.

Chan, Hok-Wah, Kwok, Wong-Chuen. 1993. An enhanced octree structure for representing the spatial/boundary decomposition of 3-D objects. In, *TENCON '93. Proceedings*. IEEE Region 10 Conference on Computer, Communication, Control and Power Engineering, Control and Power Engineering, vol. 2, pp. 1017-1020.

Craig, N. 2000. Real-Time GIS construction and digital data recording of the Jiskairumoko excavation, Peru. *SAA Archaeological Record* 18(1):24-28.

Cummings, V., Whittle, A. 2003. Tombs with a view: landscape, monuments and trees. *Antiquity* 77(296):255-266.

Cutting, M. 2002. The use of spatial analysis to study prehistoric settlement architecture. *Oxford Journal of Archaeology* 22(1):1-21.

Dawson, P. C. 2003. Space syntax analysis of central Inuit snow houses. *Journal of Anthropological Archaeology* 21(4):464-480.

De Silva, M., Pizziolo, G. 2001. Setting up a "human calibrated" anisotropic cost surface for archaeological landscape investigation. In, *Computing Archaeology for Understanding the Past: Computer Applications and Quantitative Methods in Archaeology. Proceedings of the 28th Conference, Ljubljana, April 2000*. Z. Stancic and T. Veljanovski, eds., pp. 279-286. BAR International Series 931. Oxford: Archaeopress.

Edelsbrunner, H. 1992. *Weighted Alpha Shapes*. Technical Report UIUCDCS-R-92-1760. Urbana, Illinois: Department of Computer Science, University of Illinois.

Edelsbrunner, H., Mucke, E. P. 1990. Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Transactions on Graphics* 9(1):66-104.

Edelsbrunner, H., Mucke, E. P. 1994. Three-dimensional alpha shapes. *ACM Transactions on Graphics* 13(1):43-72.

Fronza, V., Nardini, A., Salzotti, F., Valenti, Marco. 2001. A GIS solution for excavations: experience of the Siena University LIAAM. Proceedings of the 28th Conference, Ljubljana, April 2000. In, *Computing Archaeology for Understanding the Past: Computer Applications and Quantitative Methods in Archaeology*. Z. Stancic and T. Veljanovski, eds., pp. 173-177. BAR International Series 931. Oxford: Archaeopress.

Fussell, Angela. 1982. Terrestrial photogrammetry in archaeology. *World Archaeology* 14(2):157-172.

Goings, C. A. 2003. A predictive model for lithic resources in Iowa. *Plains Anthropologist* 48(184):53-67.

Good, Catriona D., Johnsrude, Ingrid S., Ashburner, John, Henson, Richard N. A., Friston, Karl J., Frackowiak, Richard S. J. 2001. A voxel-based morphometric study of ageing in 465 normal adult human brains. *NeuroImage* 14:21-36.

Grau Mira, I. 2003. Settlement dynamics and social organization in eastern Iberia during the Iron Age (eighth-second centuries BC). *Oxford Journal of Archaeology* 22(3):261-279.

Green, J., Matthews, S., Turanli, T. 2002. Underwater archaeological surveying using photomodeler, virtual mapper: Different applications for different problems. *International Journal of Nautical Archaeology* 31(2):283-292.

Grün, Armin, Sauerbier, Martin, Lambers, Karsten. 2003. Visualisation and GIS-based analysis of the Nasca geoglyphs. In *The digital heritage of archaeology – Proceedings of the 30th CAA Conference, Heraklion, Crete, April 2002*, edited by M. Doerrand and A. Sarris, pp. 161-167. Athens: Hellenic Ministry of Culture, Archive of Monuments and Publications.

Holcomb, D. W. 2001. Imaging radar and archaeological survey: an example from the Gobi Desert of southern Mongolia. *Journal of Field Archaeology* 28(1-2):131-141.

Jennings, Justin, Craig, Nathan. 2001. Politywide analysis and imperial political economy: the relationship between valley political complexity and administrative centers in the Wari Empire of the Central Andes. *Journal of Anthropological Archaeology* 20(4):479-502.

Jones, C. B. 1989. Data structures for three-dimensional spatial information systems in geology. *International Journal of Geographic Information Systems* 3(1):15-31.

Keim, Daniel A. 1999. Efficient geometry-based similarity search of 3D spatial databases. In, *Proceedings of the 1999 ACM SIGMOD International Conference on Management of Data*, pp. 419-430. New York: Association for Computing Machinery.

Kunii, T. L., Fujishiro, I., Mao, X. 1986. G-quadtree: a hierarchical representation of gray-scale digital images. *The Visual Computer* 2(4):219-226.

Ladefoged, T. N., Graves, M. W., McCoy, M. D. 2003. Archaeological evidence for agricultural development in Kohala, island of Hawai'i. *Journal of Archaeological Science* 30(7):923-940.

Levy, T., Adams, R., Hauptmann, Aandreas, Prange, M., Schmitt-Strecker, S., Najjar, M. 2002. Early Bronze Age metallurgy: a newly discovered copper manufactory in southern Jordan. *Antiquity* 76(292):425-437.

Llobera, Marcos. 2001. Building past landscape perception with GIS: understanding topographic prominence. *Journal of Archaeological Science* 28(9):1005-1014.

Mao, X., Fujishiro, I., Kunii, T. L., Shimizu, A. 1992. GIDBS: a database system for hierarchically represented gray-scale images. In, *Three-Dimensional Modeling with Geoscientific Information Systems*. A. K. Turner, ed., pp. 259-281. Dordrecht, Boston, and London: Kluwer Academic.

Meyer, Élise, Parisel, Claude, Grussenmeyera, Pierre, Revez, Jean, Tidafi, Temy. 2006. A computerized solution for epigraphic surveys of Egyptian temples. *Journal of Archaeological Science* 33(11):1605-1616.

Navazo, I., Ayala, D., Brunet, P. 1986. A geometric modeller based on the exact octtree representation of polyhedra. *Computer Graphics Forum* 5:91-104.

Nigro, J. D., Ungar, Peter S., Ruiter, D. J. de , Berger, L. R. 2003. Developing a geographic information system (GIS) for mapping and analysing fossil deposits at Swartkrans, Gauteng Province, South Africa. *Journal of Archaeological Science* 30(3):314-324.

O'Halloran, C., Spennemann, D. H. R. 2002. Wave action impact on archaeological sites in a freshwater reservoir: the case of Lake Hume, New South Wales. *Australian Archaeology* 54:6-12.

Peretto, C., Arzarello, M., Fontana, F., Gallotti, R., Gruppioni, G., Lembo, G., Pavia, M., Thun Hohenstein, U., Minelli, A. 2001. Excavation in progress: the Palaeolithic site of Isernia La Pineta, Italy. *Prehistoria* 1:138-149.

Pugh, T. W. 2003. A cluster and spatial analysis of ceremonial architecture at Late Postclassic Mayapán. *Journal of Archaeological Science* 30(8):941-953.

R$uuther, H. 1998. Photogrammetry in architecture, archaeology and urban conservation. *The Photogrammetric Record* 16(92):347-348.

Sauerbier, Martin, Lambers, Karsten. 2003. A 3D model of the Nasca lines at Palpa (Peru). *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XXXIV-5/W10: on CD-ROM

Siveroni, Viviana, Tschauner, Hartmut, Rospigliosi, Cristina, Falcón, Henry, Terukina, Ana. 2004. *Informe del Proyecto Arqueológico Huayurí (PAHUAY), temporada 2002*. Instituto Nacional de Cultura.

Spikins, Penny, Conneller, C., Ayestaran, H., Scaife, B. 2002. GIS based interpolation applied to distinguishing occupation phases of early prehistoric sites. *Journal of Archaeological Science* 29(11):1235-1245.

Stancic, Z., Veljanovski, T. 2000. Understanding Roman settlement patterns through multivariate statistics and predictive modelling. In, *Beyond the Map: Archaeology and Spatial Technologies*. G. R. Lock, ed., pp. 147-156. NATO Science Series. Series A, Life Sciences 321. Amsterdam etc.: IOS Press.

Studholme, Colin, Hill, Derek L. G., Hawkes, David J. 1997. Automated three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. *Medical Physics* 24(1):25-35.

Valenti, Marco. 1998. La gestione informatica del dato: percorsi ed evoluzioni nell'attivita della Catedra di Archeologia Medievale del Dipartimento di Archeologia e Storia delle Arti - Sezione Archeologica dell'universita di Siena. *Archeologia e Calcolatori* 9:305-329.