

New RTI Technology for Palæography

Lindsay MacDonald
Engineering Sciences
University College London,
London, UK
lindsay.macdonald
@ucl.ac.uk

Taylor Bennett
Archaeological Science
Oxford University,
Oxford, UK
taylor.bennett
@wolfson.ox.ac.uk

Christopher Ramsey
Archaeological Science
Oxford University,
Oxford, UK
christopher.ramsey
@arch.ox.ac.uk

Charles Crowther
Faculty of Classics
Oxford University,
Oxford, UK
charles.crowther
@classics.ox.ac.uk

A new illumination dome has been designed for palæographic use. It has the innovation that the majority of the 128 LED lights are concentrated at low angles of elevation, to increase angular resolution and enable finer discrimination of surface features. Image processing techniques applied to subsets of the images enable the enhancement of surface features. A new web-based OxRTIViewer provides options for multiple annotation and drawing layers, enabling palæographers anywhere easily to view each other's interpretations and collaborate across the internet. With this new system, we are creating a platform for crowd-sourcing of epigraphy, opening a new path for collaborative digital scholarship and direct public engagement with Roman writing tablets.

Polynomial texture mapping (PTM). Palæography. Archaeology, Roman writing tablets.

1. ROMAN WRITING TABLETS IN LONDON

During the excavation 2010–14 for the new Bloomberg headquarters in London, over 400 Roman wooden stylus-writing tablets were recovered from mud deposits along the historical course of the Walbrook stream, a tributary of the Thames. Located between the Bank of England and St Paul's Cathedral, the site occupies 3.2 acres on the west side of the original Roman settlement. Bloomberg reconstructed the Temple of Mithraeum which had been destroyed in 1954 to build Bucklersbury House. The architects have incorporated an exhibition of the Roman settlement and Temple within the basement, and the stories of both the Temple and the tablets are featured on the company's London home page.

The tablets are the earliest written documents found in Britain, dating from 50-80 AD in the early Roman settlement of Londinium, the predecessor of modern London. An individual waxed tablet would have been vulnerable, so they were bound together in pairs, with the wax text on the inside, plain wood on the outside, with an address scratched into it. Uses included commercial, administrative and legal purposes, such as contracts and inventories, and sometimes they served as the text messages of their day, for ephemera such as shopping lists, invitations and receipts. Only about 90 of the Bloomberg tablets have so far been partially deciphered, revealing financial accounts, legal documents, educational material and first mention of the name Londinium.

The Bloomberg tablets were recovered in a waterlogged condition, and cleaned under running water with a soft brush to remove surface soil. They were then conserved by the scientific department of Museum of London Archaeology MOLA by immersion in polyethylene glycol (PEG) followed by freeze-drying, leaving cellulose fibres and solid components of the wood intact with PEG filling the pores. The resulting objects are stable and easy to handle for inspection and photography, while retaining all the fine marks and surface features.

A typical Roman writing tablet consisted of a thin panel of wood, in most cases silver fir (*Abies alba*), of approximate dimensions 14x11 cm. Because silver fir was not available in tree form in Roman Britain, some archaeologists have speculated that writing tablets were recycled from barrel staves. Often two pieces were fastened as a diptych. A rectangular depression was filled with blackened beeswax, then inscribed with a metal stylus, leaving small scratches and depressions in the underlying wood (Bowman & Tomlin 2005).

Although the original wax surfaces containing messages have been lost, the stylus marks left on the wood enable traces of the original inscriptions to be observed. Revealing the inscriptions, presents significant challenges. There is little tonal contrast between the stylus scratches and the wooden substrate. The wood is of uniform hue, with a limited variation in lightness. The most salient visual features are cracks and the woodgrain, and chisel marks from the original tablet fabrication. In

some cases, the tablets appear to have been reused, resulting in two or more overlaid sets of inscriptions, in effect a palimpsest. The task of the palaeographer to identify the letter forms and thereby to interpret the Latin text is exceedingly difficult. Texts can be disentangled only when corrections have been made, or the lines of writing are aligned differently (Terras 2005).

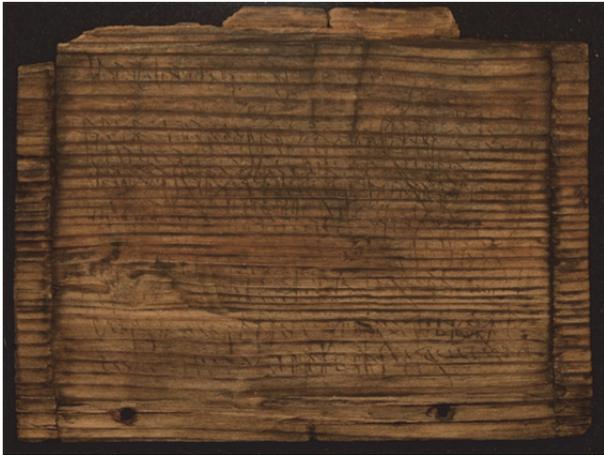


Figure 1: Roman stylus tablet BZY10_6371 <WT28>, dimensions 135x102 mm, showing superimposed texts.

Ninety of the Bloomberg tablets have been partially transcribed and translated (Tomlin 2016). The texts include the oldest financial document from the city of London (dating to 8 Jan 57 AD), legal documents including a judge making a 'pre-judgement', and educational material. One tablet shows the alphabet written out, perhaps indicating the presence of the first school in Britain.

The visibility of a scratch on the wood surface is enhanced under light cast at a raking angle perpendicular to the direction of the scratch. Digital photography with a series of directional light sources, with processing by the well-established methods of Polynomial Texture Mapping (PTM) and Reflectance Transformation Imaging (RTI), enables interactive visualisation of each tablet under a movable virtual light source, and facilitates their reading (Earl et al. 2011).

2. DESIGN OF NEW ILLUMINATION DOME

By the nature of dome geometry, the viewpoint (camera position) remains fixed while the incident rays from the lights sample the reflectance distribution function of the object surface over the enclosing hemisphere. The ability of the system to resolve angular variations in reflectance is therefore determined by the spacing of the lights. In previous dome designs, the lights are usually distributed rather uniformly. With 64 lights, for example, the mean angle between pairs of neighbouring lights is approximately 20°, so by the

Nyquist criterion the array would resolve a maximum angular frequency of one cycle in 40°, or 4.5 cycles over any great semicircle passing through the pole of the hemisphere. This limits the fineness of the specular reflection peak that can be resolved by the system, and also sets an upper limit to the order of basis functions needed to represent the intensity distribution.

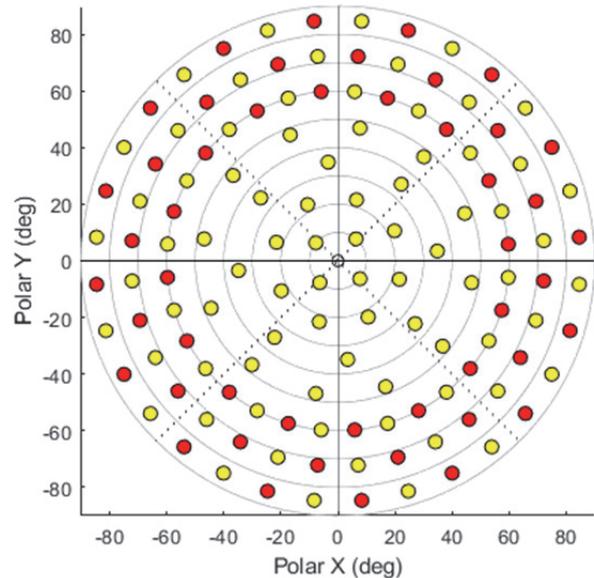


Figure 2: Concept layout of basic set of 80 lights (yellow) interleaved with an extended set of 48 lights (red).

The design of the new dome was based on the observation that light incident at lower angles of incidence ('raking light') is more effective in enhancing the contrast of local cuts and discontinuities on an object surface. Hence, for making tablet incisions visible, it should be more efficacious to position the majority of lights at low angles of elevation. The concept was, in addition to a basic set of 80 lights distributed rather uniformly, to have an extended set of 48 lights at elevation angles below 30°. As constructed, the three lower tiers of 32 lights each are at mean elevation angles of 5.3°, 18.0°, 31.6° (Fig. 3). Around the circumference, the 32 lights are at intervals of 11.25°. The four upper tiers contain 12, 8, 8, 4 lights at mean elevation angles of 47.3°, 58.6°, 69.4°, 78.2°.

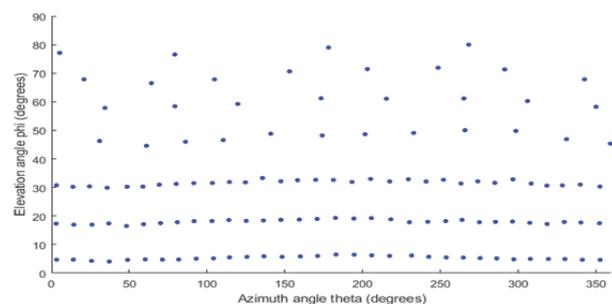


Figure 3: Measured angles of all 128 lights in dome.

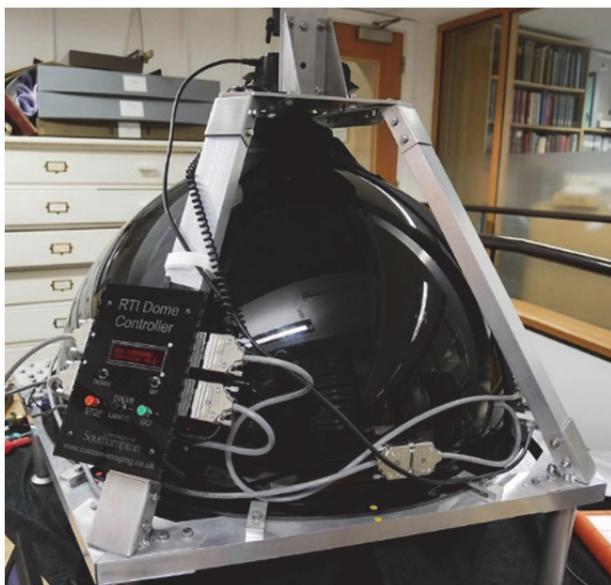


Figure 4: The Oxford RTI 'superdome' (Kirk Martinez).

The new dome is based on the principles and fabrication techniques of the previous dome at CSAD in Oxford, but with a number of refinements to improve mechanical stability and image quality that make it more suitable for the decipherment of worn inscriptions. It was designed and constructed at the University of Southampton under the supervision of Dr Kirk Martinez. The lights chosen were Bridgelux Vero 10 LED arrays, surface-mounted on inside of the shell, with specification:

CCT	CRI	Flux	Power	Efficacy
K	%	lm	W	lm/W
4000	90	1076	9.3	116

The spectrum of the light output has the bimodal distribution typical of LEDs, with a trough around 480 nm (Fig. 5), but nevertheless a high colour rendering index of 90.

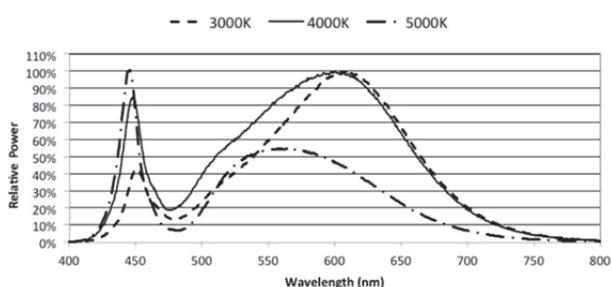


Figure 5: Spectral power distribution of Bridgelux LED.

The dome is a hemisphere of 80 cm diameter, made of moulded acrylic sheet of thickness 6 mm. It has been cut into four equal quadrants, held together by velcro straps, for ease of transport and reassembly on site. It is intended for table-top operation and is raised above the surface by 22 cm on posts to facilitate the insertion and positioning of

objects beneath camera. A heavy-duty aluminium frame with a tripod support for the camera mounting plate (Fig. 4) prevents movement of the camera during operation. The height of the sensor above the dome equator is 590 mm. The largest object that can be imaged has a nominal diameter of one-third that of the dome, i.e. 27 cm. At the dome radius of 400 mm, the 9.9 mm diameter of an LED's illuminated area subtends an angle of 1.4°.

The lights and camera are synchronised by a local micro-controller. In normal operation, for capture of an image set, the lights are switched individually in an automatic sequence. The duration of the on-time for each LED is programmable, typically of the order of one second, just sufficient for the synchronised camera to take a single image frame. Images are stored on the camera's memory card during the shooting sequence, then delivered to the host computer via a USB3 interface directly from the camera in raw format (e.g. NEF file for Nikon camera). They are subsequently converted to 16-bit linear TIFF via the DCRAW utility for processing via Matlab.

3. IMAGE PROCESSING TECHNIQUES

Previous practice for PTM has been to fit the biquadratic intensity distribution for each pixel over all images in the set, corresponding to all of the lights in the dome. The RTI fitting follows the same approach, but using hemispherical harmonics as basis functions. The effect, as with all regression procedures, is to approximate the intensity values while distributing the errors over the whole set.

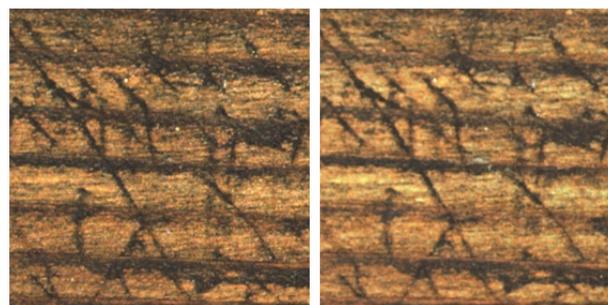


Figure 6: (left) Detail of tablet from camera image; (right) screen shot from PTM viewer.

In the viewer software the intensity of each pixel changes as the angle of the virtual light source is adjusted, but the change follows the fitted function, not the original image data. The effect is to smooth out the variations over illumination angle. Although it is a pleasing effect, and believable in terms of the illusion of moving a 'virtual torch' over a static surface, much of the local contrast of the surface is lost. To show the effect, an image detail was extracted, of dimensions 600x600 pixels, above and to the right of tablet centre. This corresponds

to a region of 14x14 mm on the surface. All of the individual images exhibit greater local contrast than the screen grab from the PTM viewer, at corresponding directions of illumination (Fig. 6). The PTM fitting technique (Malzbender et al. 2001) involves a singular value decomposition (SVD) over the image set to extract the six components of the biquadratic function ($X^2, Y^2, XY, X, Y, 1$). These are encoded as coefficients of each pixel in PTM representation, but can be extracted for visualisation as monochrome images. The fourth (X) component (Fig. 7) almost eliminates the horizontal ridges because it detects left-right gradients.

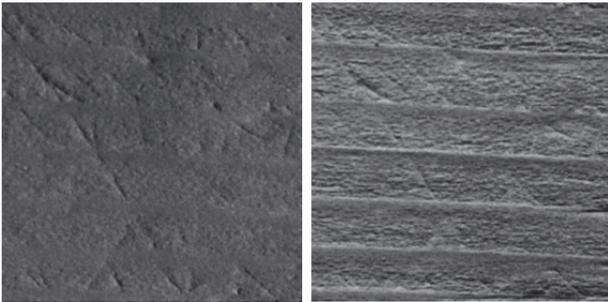


Figure 7: Two of the six PTM components extracted as images: (left) X; (right) Y.

The PTM fitting procedure can be applied to any subset of the 128 lights in the dome. This enables localised regions of the dome hemisphere to be mapped to the virtual hemisphere of the PTM viewer. For example, Fig. 8 shows a range of angles of radius 35° , centred on $\theta = 270^\circ$, $\varphi = 50^\circ$, within which there are 14 dome lights. Fitting the PTM to this subset of lights effectively maps the centre of the region to the north pole and the boundary (red line) to the equator of the coordinate system in the viewer. This has two advantages for the user: first, it magnifies the coordinate system, enabling finer adjustment of the virtual illumination angle; and second, by restricting the number of images, features visible at those illumination angles become more prominent in the PTM display.

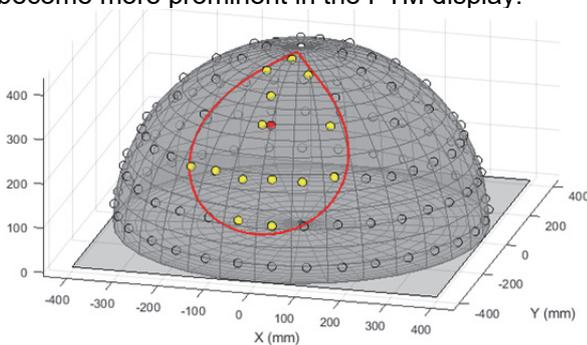


Figure 8: Dome lamps within a restricted range of angles of illumination around a given centre point (red);

Around the circumference of the dome, each of the three lowest tiers has 32 lights arranged on an almost square grid in θ and φ . Plotting the intensity

values for the green channel of a single pixel vs azimuth angle θ gives a family of three graphs, as shown in Fig. 9. The intensity varies with angle in a systematic way for this point on the slope of an incision, and increases with angle of elevation φ , as would be expected from Lambert's Law.

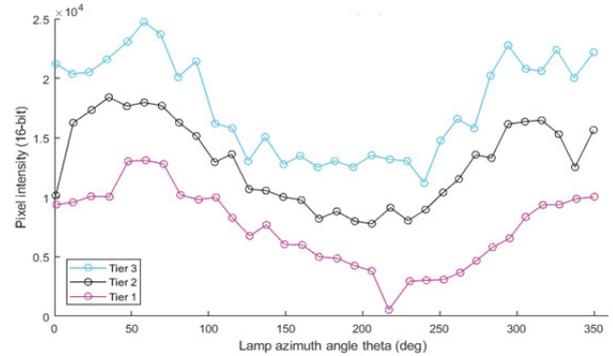


Figure 9: Pixel intensity vs azimuth angle of illumination for the three lowest tiers of lights in dome.

The images corresponding to the lowest three tiers of lights may be treated as a circular band of 32x3 intensity values, to which classical image processing techniques may be applied for finding edges and enhancing contrast. For example, a Sobel filter, which usually operates on a 3x3 block of neighbouring pixels, may be applied to the intensities of 3x3 neighbouring light positions for a single pixel. Fig. 10 (left) shows the computed gradients for all 600x600 pixels over the whole image area for the 3x3 array of lights centred on Light 21 of Tier 2, using a Sobel kernel with diagonal axis from upper-left to lower-right.

0	-1	-2
1	0	-1
2	1	0

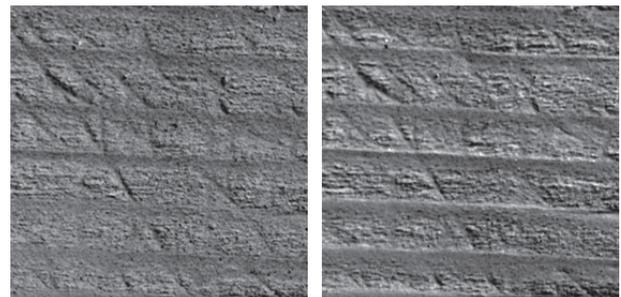


Figure 10: (left) Sobel gradients for 3x3 light array centred on light 21; (right) Difference between images for diametrically opposite lights 5 and 21 in Tier 1.

Another way of highlighting surface relief is to take the difference between images from diametrically opposite lights in the dome. Fig. 10 (right) shows the result for lights 5 and 21 in Tier 1. For this and Sobel case, the green channel of each individual image is first normalised in intensity and the result (which contains both positive and negative values) is mapped to a tone scale where mid-grey (value 127 in an 8-bit image) represents zero.

Access to the full set of images enables other kinds of image enhancement to be undertaken. For example, the images corresponding to the 32 lights in the upper part of the dome (see Fig. 3) can be summed to give the mean intensity of reflection to semi-diffuse illumination from angles within 45° of the camera axis. This could be considered to represent the ambient light within the dome. Subtracting a portion of this from an image corresponding to a single light has the effect of removing the unwanted 'flare light' and increasing the contrast of the image (Fig. 11).

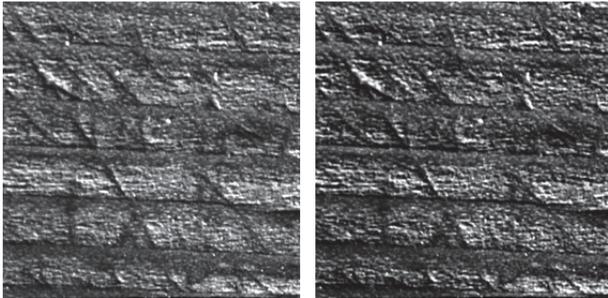


Figure 11: (left) Unmodified green channel from light 21; (right) After subtraction of $0.4x$ average of upper images.

The ridge structure of the wood is clear from the surface normal vectors (Fig. 12), generated by the bounded regression technique, in which only a subset of pixel intensity values is used, omitting shadowed and specular values (MacDonald 2015). The ridges stand above the mean surface plane, whereas the incisions cut below it. Because the ridges are nearly parallel and rather evenly spaced, a spatial frequency (Fourier domain) filtering technique could be used to remove them. The albedo (diffuse surface reflectance, independent of illumination geometry) shows that stylus incisions are intrinsically dark, i.e. not just in shadow. From the surface normals the surface gradients in the X and Y directions are easily calculated (Fig. 13). From these, the height at each point may be determined, giving a digital surface model (DSM).

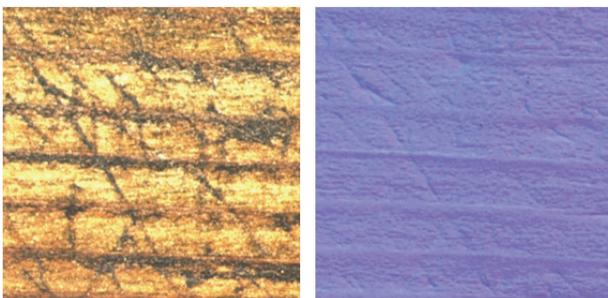


Figure 12: (left) Albedo and (right) surface normals, in conventional false-colour encoding.

A different visualisation technique is to amplify the angles of the surface normals, as if increasing the local surface gradients. After the surface normal

vectors have been calculated, the local image contrast can be enhanced by multiplying the angle subtended with the view vector by a constant, typically in the range 1.2 to 1.5, which in effect increases the surface gradients. A new set of 128 images is calculated from the modified normals by assuming a Lambertian surface with the same RGB colour as in the original PTM. Then a PTM is generated from the new image set, assuming the same spatial coordinates for the dome lights. This method is effective in increasing visibility of very shallow surface relief, and helpful for palaeography.

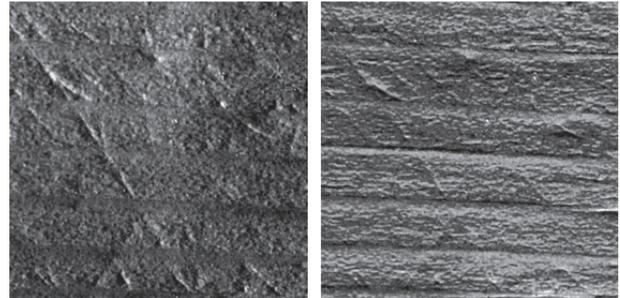


Figure 13: (left) Horizontal (X) and (right) Vertical (Y) surface gradients, derived from surface normals.

4. THE OXFORD RTI VIEWER SOFTWARE

A new interactive web-based viewer (OxRTIViewer) for PTM files has been written in Scalable Vector Graphics (SVG) for viewing and sharing image sets in web browsers. It has been developed as a test-bed for new applications, with a number of features that make it different from other web-based viewers. In particular, we have kept the software dependencies and requirements very limited:

- The software is implemented in a single html file, with no use of libraries or add-ons, and has been tested across a wide range of different web-browsers (Firefox, Chrome, Safari, Opera, Internet Explorer) and on a range of different platforms (iOS, Mac OSX, Windows). The viewer makes use of the standard html, javascript and css methodologies for general operation.
- All of the image reconstruction from the PTM components is implemented using SVG. Each PTM image is encoded in three jpg or png files, two of which store the six parameters of the PTM biquadratic function and the third stores the underlying RGB colour. These images are referred to using `<feImage>` elements, and are manipulated using `<feColorMatrix>`, `<feComposite>`, `<feComponentTransfer>`, `<feFuncR>`, `<feFuncG>` and `<feFuncB>` elements within an SVG `<filter>`. The parameters can be dynamically adjusted to produce the active PTM view. The browser uses the graphics processing unit (GPU) to process these commands,

providing the necessary speed, but no direct WebGL programming is required.

- In addition to traditional PTM views, the OxRTIViewer implements a novel calculated texture view, based on the surface normal (estimated from the PTM parameters) and using the SVG filter <feConvolveMatrix> element to perform a local kernel mask calculation. This texture map can be viewed either in false colour or as a surface using the Phong lighting model, as implemented in SVG.
- To simplify deployment, the viewer currently has no server-side communications and limited file handling. Under cross-origin resource sharing mechanisms, precomputed images can be served from any valid URL. The only requirement for description of a PTM is a single JSON file, which gives the location of the images and some basic information about the object size and parameter scaling values (Fig. 15). This file can be either in the same domain as the viewer, or uploaded by the user from his or her own computer. It can also be downloaded by the user to recall settings.



Figure 14:
OxRTIViewer GUI

- Various viewing options are provided. One view is simplified with limited controls, which might be used for a public-facing site. A more complex view provides most of the operations found in standalone viewers (Fig. 14).
- There is an annotation layer for the addition of line drawing annotations and text annotations. These are all stored in the JSON object which can be downloaded or uploaded by the user.
- The viewer could in principle be embedded in a more complex software environment with the data stored to a database or to a server file, and most of the settings can be adjusted using arguments of the URL. This enables multiple views of an object to be displayed in different frames of a web-page, which might be useful in museum applications.

The approach taken, although technically simple, turns out to be very effective for the viewing of PTM images. The main limitation is on viewing size and resolution. The viewer becomes slow when the

magnification of the image is much larger than the typical computer screen size. We assume that above this limit the GPU is unable to perform the calculations efficiently and so when moving the image or the lighting the view becomes less smooth. For many applications this may not be an issue, but clearly a WebGL approach, which specifically calculates the area of an object being viewed, would be more efficient and faster.

```
{ "json_application": "OxRTIViewer.LRGB_PTM",
  "application_url": "C:/OxRTIViewer.html",
  "url": "file:///C:/Test/Tablet.json",
  "notes": "Roman writing tablet BZY10_6371",
  "createdBy": "Lindsay MacDonald",
  "createdAt": "01-Apr-2019 23:48:03",
  "size": [1200,1000], "scale": [2,2,1,1,1,2],
  "bias": [129,116,112,91,112,0],
  "files": ["file:///C:/Test/Tablet_pt1.png",
    "file:///C:/Test/Tablet_pt2.png",
    "file:///C:/Test/Tablet_rgb.png"] }
```

Figure 15: Example of JSON file for OxRTIViewer

Another limitation of the SVG implementation is the imposition of 8-bit integer arithmetic. This is well suited to PTM, because the encoding was originally developed for 8-bit arithmetic, but is problematic for the hemispherical harmonic basis functions used in RTI (Gautron et al. 2006) because of the complex calculations involving square roots and trigonometric functions. Again this would be handled much better by floating-point arithmetic in WebGL.

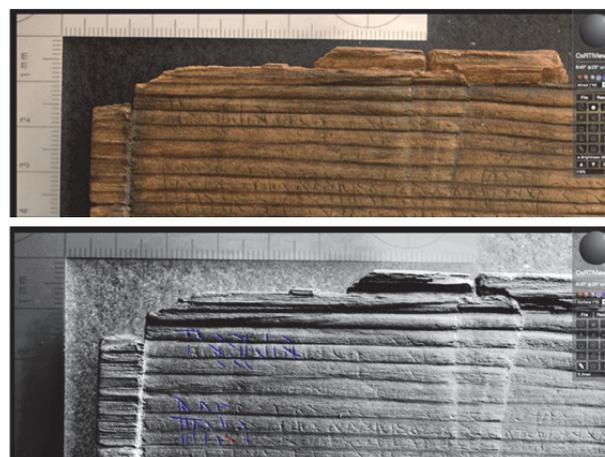


Figure 16: OxRTIViewer user interface showing tablet in standard mode (top) and annotation tool (bottom).

It is our objective to implement in the viewer some or all of the image processing operations described in the previous section. The problematic issue is providing access to the full image set. A conventional PTM of the full image frame from the Nikon D800 camera occupies 318 Mbyte, which has to be transferred to the host computer and stored locally. If the user is to select an arbitrary subset of the original images then the full set of 128 images needs to be available, which occupies

about 14 Gbyte in an uncompressed 24-bit TIFF format. Even if smart data handling techniques were employed (limiting to viewport, progressive resolution, image compression, multiple servers, etc.) the demands on storage, network bandwidth and computation would be substantial.

5. PHOTOGRAPHY AND PALÆOGRAPHY

The new illumination dome was moved to MOLA (Museum of London Archaeology) in February 2018 to capture image sets of the collection of 405 Bloomberg tablets (Fig. 17). The dome is calibrated after initial setup or whenever there are any changes to the geometry of the setup, such as the position of the camera or stage, or the lens and its aperture. The calibration process includes: (a) capturing of flat-field images for all 128 lights to enable correction for light intensity fall-off because of the dome geometry; (b) colour profiling using X-Rite ColorChecker targets and software; and (c) recording the dome's LED coordinate positions with a reflective sphere (standard highlight-RTI method).



Figure 17: Taylor Bennett with dome at MOLA lab.

We are capturing both sides of each of the 405 Bloomberg stylus tablets. Each image set consists of 128 raw images (Nikon NEF format), each of 45-48 Mbytes, hence about 6 Gbyte for storage of the full set (12 Gbyte for both sides of the tablet). We therefore anticipate that the complete archive will require approximately 4 Tbyte of storage. Each raw image from the Nikon D800E is 7360x4912 pixels in size, which when converted to an RGB 16-bit per channel (48-bit per pixel) TIFF file occupies 209 Mbyte, hence 27 Gbyte per set (i.e. over five times that of the NEF format). An archive containing both NEF and TIFF versions of each image would therefore require about 26 Tbyte of storage.

The stylus tablet to be photographed is placed on a height-adjustable lab-jack stand (Fig. 18), mounted on ball-bearing slider rails that allow the stage to be inserted beneath the centre of the dome and retracted when the capture sequence has been completed. The height of the equator of the dome above the slider-board is 22 cm (fixed by mounting posts), and the height of the top of the lab-jack stand is approximately 17 cm (this may vary according to the thickness of the object), so the distance between the top of the stage and the dome equator is 5 cm. Most of the Bloomberg tablets are captured using a Nikkor 105-mm macro lens, giving a spatial resolution of approximately 43 pixels/mm on the tablet surface. In a small number of cases where a tablet is too large to fit within the field of view of the 105-mm lens (approximately 15x10 cm), overlapping image sets of the tablet are captured, followed by an image set of the entire tablet using a 60-mm macro lens. The typical exposure setting is $f/10$ at 1/10 second.



Figure 18: Roman writing tablet on lab-jack inside dome.

In addition to documenting the stylus tablets using RTI photography, a goal of this project is to crowd-source the epigraphy of the tablets (Bennett et al. 2018). A prerequisite of this crowd-sourcing effort has been the development of a web-based viewer that provides rendering options (specular enhancement, texture kernel filter, and false-colour surface normals) with drawing and annotation tools, and the ability to record all settings, drawings and annotations. The OxRTIViewer was designed to provide the features needed for crowd-sourced collaboration that are not available in the existing web-based WebRTIViewer (ISTI-CNR 2012). The graphical user interface (GUI) for OxRTIViewer is illustrated in Figures 14, 16 and 19.

Initially, a range of interested users from the academic and student communities at the University of Oxford and a cohort of experts from outside Oxford were invited to test the OxRTIViewer using a small number of pre-processed PTM files. We have also conducted

preliminary user testing with the managers of collections of PTMs at the University of Oxford, including the Cuneiform Digital Library Initiative, Minoan seal stones and Linear B tablets at the Ashmolean Museum, and the Bodleian Library. We have also shared OxRTIVIEWER with selected international users, including Cultural Heritage Imaging (CHI) of San Francisco and the PIXEL+ project at the University of Leuven. User feedback to date includes:

- provide preset illumination positions and rendering options for reproducibility of views;
- allow rapid comparisons of different epigraphic interpretations of the scripts;
- enable recording and retrieval of drawings, annotations and rendering parameters;
- support use of a hand-held stylus and/or touch screen for drawing and annotating;
- include an 'eraser' function for revising drawings;
- zoom, pan and rotation should respond in real-time to touch-screen or tablet inputs;
- provide better on-screen documentation and help facilities;
- consider integrating the viewer with a crowd-sourcing platform such as Zooniverse.

An additional rendering option that could be useful is Diffuse Gain, a feature of the stand-alone RTIVIEWER (CHI 2013). Algorithmic methods for selecting the most favourable settings of light position and other rendering parameters could be developed, such as the maximum entropy method (Malzbender and Ordentlich 2006). Another option might be to automate the motion of light positions and other rendering parameters to take advantage of movement as an aid to perception of fine details by the human visual system.

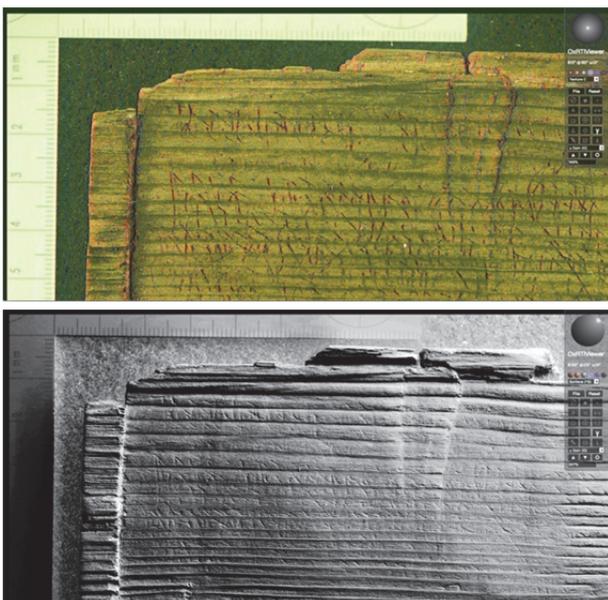


Figure 19: Display by OxRTIVIEWER in two rendering modes: texture filter and specular enhancement.

In conclusion, we have developed new technology for RTI, including an illumination dome with 128 lights, novel image processing algorithms, and an innovative web-viewer. We believe that these tools will enable greater levels of achievement in the palaeographic interpretation of Roman writing tablets and will also open the way to crowd-sourced collaboration on the decipherment of ancient texts.

6. REFERENCES

- Bennett, T., Crowther, C., MacDonald, L.W. and Ramsey, C. (2018) The Bloomberg Stylus Tablets. *Proc. 4th Intl. SEAHA Conf.*, London, pp. 41-42. <http://www.seaha-cdt.ac.uk/wordpress/wp-content/uploads/2018/07/Abstract-Book-.pdf>
- Bowman, A.K. and Tomlin, R.S.O. (2005) Wooden Stylus Tablets from Roman Britain. In Bowman, A.K. & Brady, M. (eds) *Images and Artefacts of the Ancient World*, British Academy Paper #4, OUP.
- CHI (2013) *Reflectance Transformation Imaging: Guide to RTIVIEWER*. http://culturalheritageimaging.org/What_We_Offer/Downloads/rtiviewer/RTIVIEWER_Guide_v1_1.pdf
- Earl, G., Basford, P.J., Bischoff, A.S., Bowman, A., Crowther, C., Dahl, J., Hodgson, M., Martínez, K., Isaksen, L., Pagi, H. and Piquette, K.E. (2011) Reflectance transformation imaging systems for ancient documentary artefacts. *Proc. BCS Conf. on Electronic Visualisation & the Arts (EVA)*, London.
- Gautron P., Krivanek J., Pattanaik S.N. and Bouatouch K. (2004). A Novel Hemispherical Basis for Accurate and Efficient Rendering. *Rendering Techniques, Proc. 15th Eurographics Conf. on Rendering Techniques (EGSR'04)*, pp. 321-330.
- ISTI-CNR (2012) WebRTIVIEWER and WebGLRTI Maker. In Palma, G. (ed) *Reflectance Transformation Imaging (RTI)*. Pisa, Italy.
- MacDonald, L.W. (2015) 'Representation of Cultural Objects by Image Sets with Directional Illumination', *Proc. Computational Color Imaging Workshop (CCIW'15)*. St. Etienne, France. Springer LNCS 9016, pp.43-56.
- Malzbender, T., Gelb, D. and Wolters, H. (2001) Polynomial Texture Maps, *Proc. ACM Siggraph* 28:519-528. <http://www.hpl.hp.com/research/ptm/>
- Malzbender, T. and Ordentlich, E. (2006) Maximum entropy lighting for physical objects. *Proc. Conf. on Visualization and Data Analysis*, SPIE Vol. 6060.
- Terras, M. (2005) Reading the readers: Modelling complex humanities processes to build cognitive systems. *Literary and Linguistic Computing*, 20(1), pp.41-59, Oxford.
- Tomlin, R.S.O. (2016) *Roman London's First Voices: Writing Tablets from the Bloomberg Excavations, 2010-14*. MOLA Monograph Series 72, ISBN 978-1-907586-40-8, Museum of London.