

Image Capture and Processing for Enhancing the Legibility of Incised Texts

Ségolène Tarte¹, Mike Brady², Alan Bowman³, Melissa Terras⁴

¹*Oxford e-Research Centre, University of Oxford, UK, e-mail: segolene.tarte@oerc.ox.ac.uk*

²*Department of Engineering Science, University of Oxford, UK*

³*Centre for the Study of Ancient Documents, Classics Faculty, University of Oxford, UK*

⁴*Department of Information Studies, University College London, UK*

Keywords: incised tablets, shadow-stereo, feature extraction, region labelling, texture mapping.

Abstract: This paper presents methodologies that have been adopted to enhance the legibility of incised texts, in particular Roman wooden stylus tablets, for which the texts consist of the incisions left in the wood through a now-perished coat of wax originally covering the wood. Digitization of such artefacts is the first step in the development of an interpretation of the document. At this stage, mimesis of the real-world strategy of the classicists is a guiding principle. Taking into account the 3D nature of the document, shadow-stereo and reflectance transformation imaging allow us to capture and to encode multiple images of the text under varying illumination conditions for further processing and visualization. Image processing algorithms were developed to isolate the text features. Background correction is first performed; then ways to achieve text feature extraction have been explored: phase congruency, which exploits the fact that visual features are detected for some properties of local phase of the image in the Fourier domain; and Markov Random Fields, which take a statistical approach to region labelling for image segmentation. Most techniques used here were inspired by approaches adopted by medical image processing. These methods had however to be largely adapted to our specific application; in general, the images of artefacts and their features of interest are not only different from medical images, also the type of visual expertise required to detect the text features differs greatly from that of radiologists. We conclude by observing that by better understanding the nature of the classicists visual expertise, we will further be able to integrate prior knowledge into a model of visual perception adapted to the classicists needs, hence supporting them in building meaning out of a pure signal, in building an interpretation of an artefact.

INTRODUCTION

Incised texts such as those on Roman stylus tablets or clay cuneiform tablets present the particularity of displaying script in volume, and, unlike epigraphic material, of being rather small. Deciphering such texts draws on a wide range of expertise that papyrologists develop throughout their career. These skills are as much related to visual perception of the artefact and its text as they are to palaeographical, linguistic and historical knowledge. Within the

scope of the e-Science and Ancient Documents project (eSAD¹), we are developing a software tool to support papyrologists in their interpretation task. In order to build such an Interpretation Support System (ISS), a deeper understanding of the tasks involved in the interpretation practice is necessary. One particular trait of the hermeneutic process is that, while it unfolds, the mobilized knowledge is multifarious and often implicit (Youtie, 1963); interpreting an artefact has recourse to intricately intertwined visual, palaeographical, linguistic and historical skills (Tarte, 2010). Building on previous work that elaborated a model of papyrological reading² (Terras, 2006), we were able to identify levels at which digital support can be provided. This digital assistance aims to support experts in their decipherment of text-bearing artefacts by making explicit some of the mechanisms involved in reading an ancient and damaged text. The focus of this paper is on the stages of digitization of the artefact and processing of the acquired digital data; these stages correspond to the levels of papyrological reading concerned with the physical aspect of the document, the features of a character and the identification of a character. They are at the low end of the spectrum in terms of intrinsic meaning of the levels of reading, and they heavily condition the progressive build-up of meaning. We show here how we have attempted to digitally imitate the incredibly specialized expertise of the papyrologists in order to facilitate interpretation. Roman wooden stylus tablets from Vindolanda and Tolsum, for which the texts consist of the incisions left in the wood through a now-perished coat of wax originally covering the wood, constitute the main case study of the paper.

IMAGE CAPTURE

Creating a digital representation of a text-bearing artefact is never neutral, and care must be taken that the intended use of the digital avatar of the object is clear when proceeding with digitization. In this context, it is the observation of the experts setting to interpret a text-bearing artefact that informed us on how the tablets could be digitized whilst retaining in particular one major piece of information that papyrologists exploit when deciphering it, namely, the volumetric nature of the text, i.e. one crucial aspect of its materiality.

Shadow-stereo

Experts who have access to the actual tablet they intend to transcribe and interpret have developed a very specific and intuitive strategy to enhance the visibility of the incisions that the stylus left in the wood. They lay the tablet flat on their hand, lift it up at eye level against a light source and apply pitch-and-yaw motions to the artefact. What they effectively do is enhance the visibility of the incisions by accentuating the highlights and shadows that the raking light generates; the lower the light the longer the shadows projected by the text in (inverted) relief created by the incisions carved on the surface of the tablet. The principle that is put into application here is the shadow-stereo principle by which concave shapes are revealed from shading, and the motion of the shadows exposes the location of the incisions (Brady et al., 2005). This process can be digitally imitated through a set-up where a digital

¹ Project website: <http://esad.classics.ox.ac.uk>

² There are ten levels of papyrological reading defined in (Terras, 2006). They were identified through the study of the subjects around which discussions revolved during the interpretation of Roman wooden tablets from Vindolanda; they are ordered according to increasing intrinsic meaning: (0) Archaeological or historical context; (1) Physical attributes of the document; (2) Features of a character; (3) Possible character; (4) Possible sequence of characters; (5) Possible word or morphemic unit; (6) Grammar; (7) Meaning or sense of a word; (8) Meaning or sense of a phrase or group of words; (9) Meaning of the document.

camera is affixed above the tablet laying flat and a high-resolution digital picture is taken for each one of a set of pre-established positions of a light source around the tablet. Each light position is described by an elevation angle and an azimuth angle, where the azimuth angle corresponds to an angle deviation from the horizontal in the plane of the tablet and the elevation angle describes the height of the light with respect to the plane of the tablet (Fig. 1). Both angles are measured from the centre of the tablet.

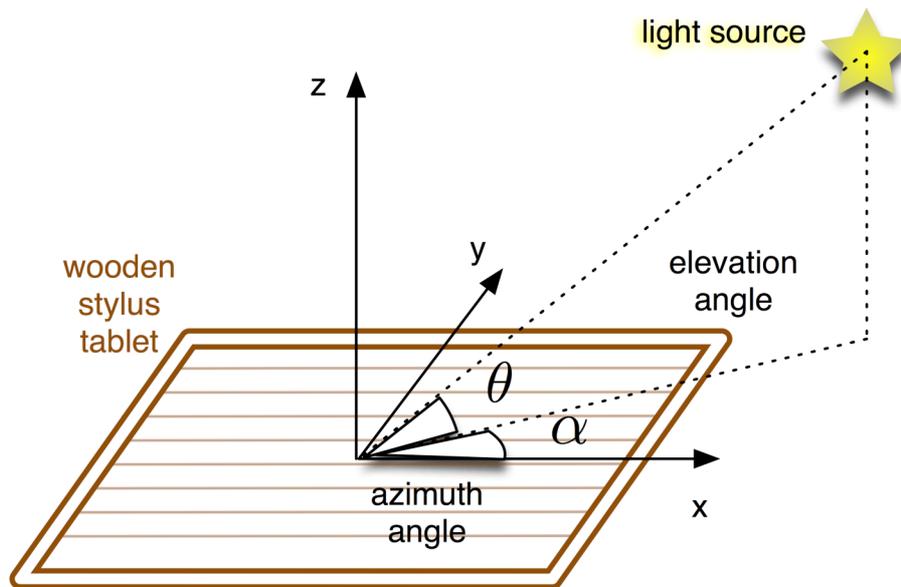


Figure 1: Definition of the azimuth angle α and elevation angle θ to describe the lighting conditions with respect to the tablet.

Reflectance Transformation Imaging through Polynomial Texture Maps

Taking into account the volumetric aspect of the materiality of the text by taking a set of high-resolution digital pictures with varying lighting directions is the first step of the attempt to mirror the shadow-stereo principle in the digital world. Typically, a set of between $N=28$ and $N=56$ images is collected. Each image is 4288×2848 pixels for artefacts of approximately the size of a modern postcard ($8 \times 12 \text{ cm}$), making each image approximately 13 MB . The collected data for each tablet was further used to recreate digitally the shadows and their motion. By adopting an appropriate model of image formation not only can one store the set of images efficiently without having to store each image, one can also interpolate between light positions, thus simulating lighting conditions that were not actually captured. The image model that is adopted relies on reflectance transformation (Malzbender et al., 2001; Goskar and Earl, 2010). Consequently, each image can be expressed in the RGB space as:

$$R_{(x,y)}(u,v) = L_{(x,y)}(u,v)R^0_{(x,y)}$$

$$G_{(x,y)}(u,v) = L_{(x,y)}(u,v)G^0_{(x,y)}$$

$$B_{(x,y)}(u,v) = L_{(x,y)}(u,v)B^0_{(x,y)}$$

$L_{(x,y)}(u,v)$ is the luminance at pixel (x,y) for the lighting condition (u,v) ; (u,v) parameterizes the light position so that $u = \cos\alpha \cos\theta$ and $v = \sin\alpha \cos\theta$ for azimuth angle α and elevation angle θ (Fig. 1). An appropriate model for L is a biquadratic in (u,v) of the form:

$$L_{(x,y)}(u,v) = a_0^{(x,y)}u^2 + a_1^{(x,y)}v^2 + a_2^{(x,y)}uv + a_3^{(x,y)}u + a_4^{(x,y)}v + a_5^{(x,y)}$$

It is this function $L_{(x,y)}$ that needs to be determined on the basis of the acquired images, for each pixel (x,y) . Determining the coefficients $a_0^{(x,y)}, \dots, a_5^{(x,y)}$ of $L_{(x,y)}$ and the light independent coefficients ($R^0_{(x,y)}, G^0_{(x,y)}, B^0_{(x,y)}$) means solving an over-constrained system of equations for each pixel (x,y) , which can be done by Singular Value Decomposition (SVD). So that instead of storing the $3N$ values for each pixel as we would if we were to store each of the N acquired images, all that needs to be stored for each pixel is the set of: the basic ($R^0_{(x,y)}, G^0_{(x,y)}, B^0_{(x,y)}$) values and the 6 values $a_0^{(x,y)}, \dots, a_5^{(x,y)}$, so only the equivalent of up to 7 times the size of one colour image, rather than $N \in \{28, 56\}$ colour images. This new image is called a Polynomial Texture Map (PTM), as it describes each pixel based on a biquadratic polynomial ($L_{(x,y)}(u,v)$) which enables to characterize the variation in pixel colour (or texture) according to the light. It is then also possible to simulate any (α, θ) lighting condition on the artefact, and, with an adequate piece of software³, to interactively simulate light motion.

IMAGE PROCESSING

Here again, mimesis is our guiding principle. The visual system is a powerful and complex system that can only be emulated if we understand how it functions. Its complexity prompts to set as a first aim to emulate human performances, before attempting to surpass it. In particular, we need to grasp how the experts' visual system proceeds to discriminate background from text. It appears that concurrently to enhancing the image, as presented in the Image Capture section, the visual system adapts to ignore distracting noise as well as it looks for text elements with some expected shapes or occurrences of characteristic features. In image processing terms, this can be assimilated to concurrent noise removal and image segmentation. On one hand, noise removal addresses the need to homogenize the background by removing patterns or behaviours that are irrelevant to the text. And on the other hand, image segmentation is concerned with extracting the meaningful areas or features of the image and thereby intends to identify where the text is located.

Background Correction

From the PTM files, it is possible to extract a set of images that have the optimal lighting conditions for the visibility of the stylus strokes. Those conditions correspond to low raking light positions, where the direction of lighting is approximately aligned with the grooves of the wood-grain, thus minimizing the shadow cast by the wood-grain itself, yet producing the most visible shadows for strokes not aligned with the wood-grain. We sample 18 images from the PTM, corresponding to: $-15^\circ \leq \alpha \leq 15^\circ$ or $165^\circ \leq \alpha \leq 195^\circ$ (alignment with the wood-grain); and $0^\circ \leq \theta \leq 6^\circ$ (raking light). These images however are quite noisy, and before exploiting the multi-view nature of the PTM, we need to remove the distracting elements from each of these images, namely: the uneven illumination, and the stripes of the wood-grain⁴. To perform illumination correction, homomorphic filtering is the method of choice. It

³ e.g. developed by HP labs: <http://www.hpl.hp.com/research/ptm/> –last checked: September 29, 2010

⁴ From this point on, all images are assumed to be grey value images. To obtain a grey value image I from a RGB image, we eliminate hue and saturation, and retain luminance: $I = 0.2989 \times R + 0.5870 \times G + 0.1140 \times B$

was implemented with a moving average algorithm, which approximates a low-pass filtering, and achieves a levelling of the brightness throughout the image (Pan, 2004). Removing the wood-grain is further achieved by taking advantage of the surface properties of the tablet and of the expected geometrical behaviour around the grooves of the wood-grain. Assuming Lambertian reflectance of the surface and that the local colour is constant, assuming also that the wood-grain is aligned with the lighting direction ($\alpha \approx 0^\circ$ or $\alpha \approx 180^\circ$), the pixel values obey: $\forall y, \sum_x I(x,y) = c \cdot \cos \gamma(y)$, where $I(x,y)$ is the intensity value at pixel (x,y) ; c is the constant colour and $\gamma(y)$ is the angle between the incident light (the incident angle is denoted θ in Fig. 1) and the normal vector to the surface of the tablet $\vec{n}(x,y)$. For a given x (a column of the image), as the coordinate system is aligned with the wood-grain, the normal to the tablet at position (x,y) only depends on the row (y), so that in effect, $\vec{n}(x,y) \approx \vec{n}_x(y)$. So that, by artificially forcing $\gamma(y) = \pi/2 - \theta$ (where θ is the incident light angle) one can level up the grooves of the wood-grain, thereby removing them (Fig. 2).

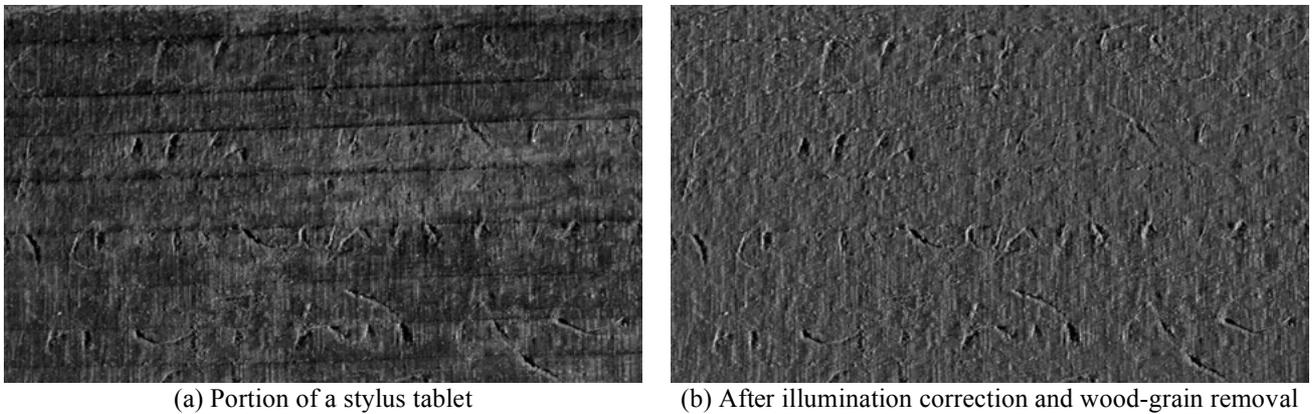


Figure 2: Before and after background correction

Image Segmentation for Text Feature Detection

In the following three subsections, we present methods that were applied to perform feature detection.

▪ Combining the PTM Information for Feature Detection

After the individual images have been cleaned up from the obvious noisy signal (as described above), the images need to be combined in such a way that the visibility of the strokes constituting the text is further enhanced. To that end, a single image is produced; it is the absolute difference image between the maximum of all images (the image that keeps all the brightest pixels in the set of images) and the minimum of all images (the image that keeps all the darkest pixels in the set of images):

$$\forall(x,y), I_{diff}(x,y) = \left\| \max_{k=1,\dots,18} I_k(x,y) - \min_{k=1,\dots,18} I_k(x,y) \right\|$$

Because the strokes are illuminated from opposite directions, the areas where highlight and shadow occur overlap and the difference image accentuates their location (Fig. 3 and 4-a). The next step is to identify these locations by segmenting this difference image. “Segmentation subdivides an image into its constituent regions or objects” (Gonzalez and Woods, 2008). As such, image segmentation is a very generic problem for which there are only domain-specific solutions. Specific domains develop specialized ways of looking and of seeing that impose implicit expectations as to what is to be seen. Each domain-specific

solution to the image segmentation problem integrates (to a certain extent) some of the domain-specific knowledge. Strategies to carry out image segmentation can however generally be seen as region labelling methods or as feature detection approaches, or as a combination of both. Region labelling methods set out to determine patches of images that share a given property, whereas feature detection procedures implement a search for breaks in the uniformity of the image. Ultimately in ideal cases, one can see the borders of the patches produced by region labelling as features; and the objects detected by feature detection algorithms can be seen as delimiters for regions. In short, region labelling concentrates on extracting uniform regions (for a domain-specific definition of “uniform”), and feature detection works to extract breaks in uniformity (for a domain-specific definition of “break in uniformity”).

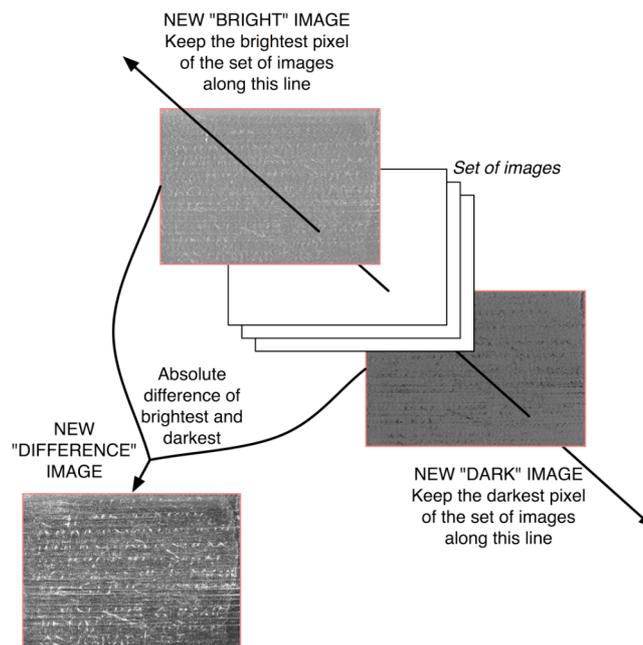


Figure 3: Obtaining the difference image I_{diff} (lower left image)

▪ Markov Random Fields

Markov Random Fields (MRFs) belong to the region labelling category of segmentation. MRFs take a statistical approach to vision and encode contextual constraints in their strategy (Li, 2009). In particular, MRFs consider the probability for a pixel to belong to one or another region of the image. Based on neighbourhood considerations, on an estimated prior model, and on assumptions about the interactions between the various labelled regions, it computes the likelihood for a pixel to belong to a given region, and allows to correct pixel values accordingly. An optimization scheme finds the equilibrium point, where not only is the labelling determined, also the prior model is progressively adapted to fit the data better. To implement this optimization scheme, we have adopted the Expectation Maximization (EM) algorithm, as described in (Van Leemput et al., 1999), with a local behaviour following an isotropic Gaussian distribution. As a result, the pixel values are “corrected” and a smoother, more consistent image is produced. Although no visual difference is noticeable between the input and output images of the MRFs, the subsequent image processing algorithms seem to give better result if this MRF image correction step is taken. Interestingly however, this algorithm performs poorly in terms of pure region labelling, due to the amount and nature of the noise present in the image. We shall consider this noise problem further in the Discussion.

▪ **Phase Congruency**

After correcting the image by means of MRFs, the main aim is to detect text features. Text features are here characterized by ridges, or adjacent steps patterns. It has been shown that the feature information is predominantly contained in the local phase information of an image. Local phase is a quantity that appears in all signals and is defined as follows. Locally, every (1D) signal can be modelled as a periodic signal of period T and thus expressed as a Fourier series:

$$I(t) = A_0 + \sum_{n=1}^{\infty} A_n \cos\left(\frac{2\pi}{T}nt + \varphi_n\right)$$

where φ_n is the local phase at the n th harmonic. Remarkably, it was hypothesized (and further demonstrated) by teams of psychologists, computer scientists and neuroscientists (Oppenheim and Lim, 1981; Morrone and Owens, 1987; Henriksson et al., 2009), that when a feature is present in a signal, the value of φ_n is constant with value Φ , and doesn't depend on the rank of the harmonic; the phase is then said to be congruent throughout harmonics. So that a strategy to detect features is to look for those locations where φ_n does not depend on n . Further, the actual value of Φ , when phase congruency occurs, characterizes the nature of the feature. In particular: $\Phi = \pi/2$ indicates a valley; $\Phi = \pi/2$ denotes a ridge; $\Phi = 0$ represents a step up; $\Phi = \pi$ designates a step down (Venkatesh and Owens, 1990). Computing the phase value however (be it for 1D signals or images), and finding where it is congruent throughout harmonics is a complex task. It requires to first evaluate phase values that correspond to ranges of harmonics, which is done through band-pass filtering in the Fourier domain (Kovesi, 2000; Felsberg and Sommer, 2004). The choice of the band-pass quadrature filters used to compute phase at various scales (harmonics) depends heavily on the type of features that are searched, and in particular their scale coverage and sharpness (Boukerroui et al., 2004). We have opted for a difference of Gaussians. The next step is to assess if congruency

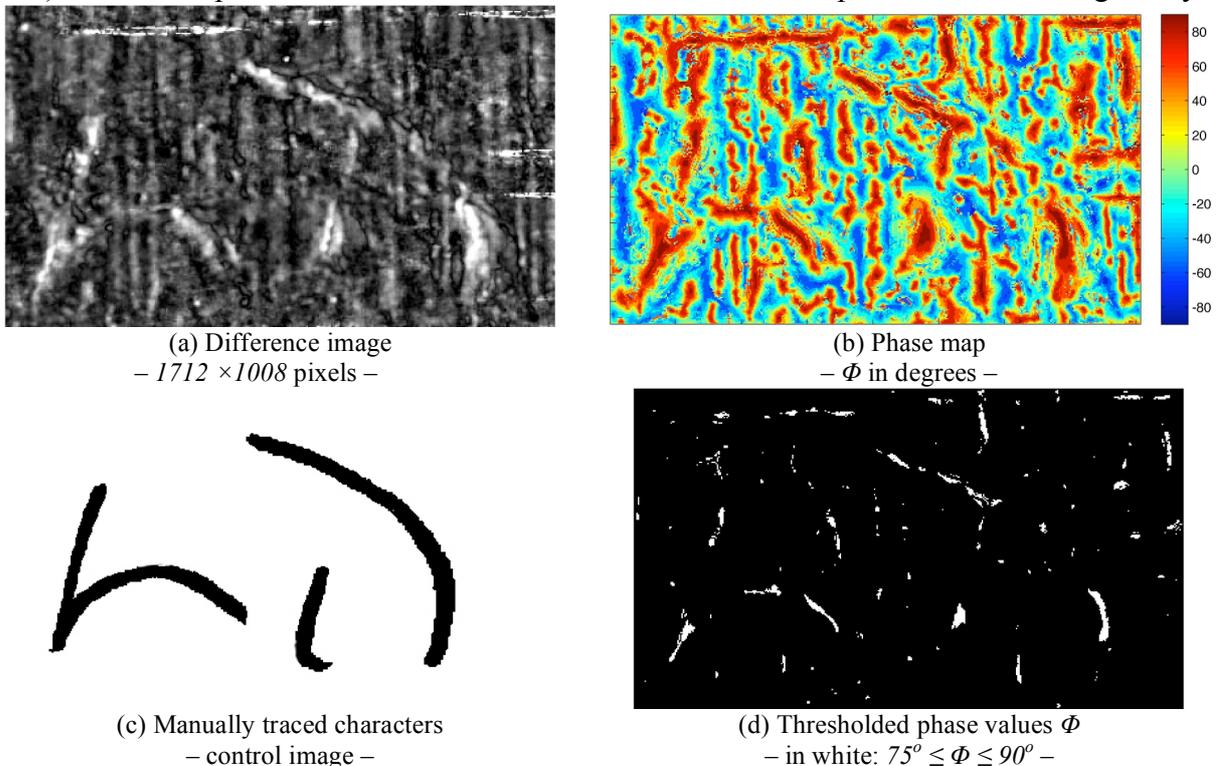


Figure 4: Portion of a stylus tablet with two characters.

occurs throughout scale-space⁵. To that effect, we use NP-windows (Kadir and Brady, 2005) to evaluate the probability distribution of phase values throughout scales. The result is a phase map, where each pixel is associated to its most probable phase value (Fig. 4-b). In the context of the text features in the difference image (as defined in the section ‘Combining the PTM Information for Feature Detection’), we are looking for ridges. So, we can threshold the phase map to keep all pixels with local phase value between $5\pi/6$ (75°) and $\pi/2$ (90°) (Fig. 4-d). The thresholded phase map shows the locations where a ridge feature was detected. Subsequent processing needs then to be applied to determine which are the locations that correspond to text features.

DISCUSSION

Image segmentation is an arduous task that presents numerous hurdles. The methods used here are methods that have largely been successful for medical images. Their performance on images of ancient texts, despite being very promising, are not as straightforwardly good as we had hoped. There are several reasons for that. One reason is that the images of ancient documents are noisy. Naturally, medical images are noisy too (Pham et al. 1999), but the nature of the noise in medical images is preponderantly due to the imaging modality⁶. In general, these imaging modalities are based on the understanding of a physical phenomenon, such as the attenuation of X-rays (in X-ray imaging) or the re-orientation of water molecules (in Magnetic Resonance Imaging). The noise that affects the images in those cases is then only due to the measurement device that produces the images, as the physical model of interaction between the emitted “waves” (X-rays, MR) and the body they traverse is a complete one. Similarly, photographic imaging is the transposition of an optical effect, and is also affected by noise. However, this noise due to the photographic procedure in images of artefacts isn’t the only type of noise present in the images. Noise, defined as “random fluctuations that obscure or do not contain meaningful data or other information”⁷, is also in our context anything that obscures the reading (our sought “information”). As such, stains and the general state of degradation of the artefact where writing is located constitutes noise, and we do not have any model for this type of noise. The state of the artefact itself, after a long sojourn in various (sometimes extreme!) conditions, is a major factor in the legibility of its text. So compensating for this kind of noise, albeit possible in some cases (as we attempted here with MRFs and wood-grain removal), is usually the main difficulty. Phase congruency is a powerful tool, but it expects the features to follow an almost perfect model, and the complexity of the algorithms used makes them all the more sensitive to noise.

Another interesting consideration, and point of comparison with medical imaging, is the way the images are interpreted. The algorithms used here were originally developed to support the interpretation and understanding of medical images. Medical images are ‘figurative’ images, for which the information relates to geometrical, topological or functional representations of a physiological or anatomical phenomenon (Cohn, 2007). They are laden with meaning of a very different nature than the meaning carried by ‘textual’ images. The assumption in using the same type of algorithms in both cases is that the primitives that allow further interpretation are of the same kind and follow perfect models of edges, ridges, valleys and

⁵ We here use the terms scale-space and Fourier domain interchangeably.

⁶ It is worth stressing that although we evoke medical images as a unified set, they do vary widely in nature and quality. We do not intend to present a review of medical imaging here, and so we only focus on the commonality between those various medical images from an image processing point of view.

⁷ From the New Oxford American Dictionary, Apple Inc.’s Dictionary V.2.0.3.

steps. Yet, even if the primitives on which interpretation is based are similar, how much of the implicit knowledge and assumptions have been encoded in the algorithms is unclear. It is important here to be conscious that the knowledge attached to ‘figurative’ medical image is likely to follow very different semantic rules from those attached to ‘textual’ images. To evaluate if and how much implicit knowledge has been encoded in the algorithms, it would be useful to understand if the mental organization (and summoning) of textual information and that of figurative information differ.

The cognitive processes involved in the act of interpretation and the fact that they are mostly implicit are one of the reasons why it is so tricky to emulate what the visual system does. Seeing that and seeing as (Neer, 2005) are intricately intertwined, and despite the fact that levels of reading have been clearly identified in the model of papyrological reading we have adopted form (Terras, 2006), oscillations between these levels constantly occur and are triggered mostly by expectations, of the visual, historical or philological type (Tarte, 2010). It is because meaning seeps in at all stages of the interpretation through the oscillations between the levels of reading that providing digital support for such a task is an exciting challenge. How can a digital tool keep allowing and even facilitate the oscillations, while providing useful tools at each level of reading? One element of answer is certainly: by not constraining the experts to a rigid pre-determined workflow.

CONCLUSION

The challenge of image processing and capture for ancient incised documents is multiple and specific to both the imaging method and the artefact. We have presented how we propose to capture the 3D nature of the artefact in a way that mimics the experts’ strategy to enhance the visibility of the text. By exploiting the play of light with the 3D nature of the text, we are able to enhance its legibility. Background correction is used to neutralize some of the distracting noise present in the images through homomorphic filtering and wood-grain removal; and a powerful feature detection technique, based on phase congruency –which was shown to be a crucial mechanism of visual perception in feature detection– is further applied to isolate text features. Each of these individual techniques were ported from medical imaging but had to be combined in a specific way and adjusted to the needs of the classicists. To further support the experts in their interpretation enterprise, we need to understand how to combine visually and meaningfully the amount of information extracted by image processing, and complement it with access to prior palaeographical knowledge. Throughout this work, the understanding of the cognitive processes involved in the hermeneutic task and attempts to emulate them are serving as guides to inform the implementation of our tool. It is important to stress here that the ultimate aim is not to replace the experts, but rather to provide them with a useful tool that can help them produce explicit evidence for their findings. Future work will concentrate on designing a system that offers these functionalities to support the papyrologists. We will strive to allow them to work through their natural workflow, in particular to allow the all-important meaning-generating oscillations between the levels of reading. This system will be integrated within a rationale recording tool that will allow to make hypotheses and their supporting evidence explicit (Roued-Cunliffe, 2010). The results of the algorithms presented here may be used (if desired!) as evidence in such a chain of reasoning, all in an effort to reveal the intricacies of the transition from signal to meaning, from the image of the text-bearing artefact to the interpretation of the text.

ACKNOWLEDGEMENTS

This work was supported by the AHRC-EPSRC-JISC Arts and Humanities e-Science Initiative.

The authors wish to thank Dr V. Grau for his help with MRFs, as well as colleagues at the OeRC, for fruitful discussions around images and their understanding.

The authors are also most grateful to the members of the eSAD team for their help and support H. Roued-Cunliffe.

Finally, we are grateful to the anonymous reviewers of this paper who provided us with insightful and encouraging feedback.

REFERENCES

- Boukerroui, D., Noble, J., and Brady, M. (2004), On the Choice of Band-Pass Quadrature Filters, *J. Math. Imaging Vis.*, volume 21(1):pp. 53–80.
- Brady, M., Pan, X., Schenck, V., Terras, M., Robertson, P., and Molton, N. (2005), Shadow stereo, image filtering and constraint propagation, in A.K. Bowman and M. Brady, editors, *Images and Artefacts of the Ancient World*, volume 4 of British Academy Occasional Paper, pp. 15–30, Oxford University Press/British Academy, Oxford.
- Cohn, S. (2007), Seeing and drawing: the role of play in medical imaging, in C. Grasseni, editor, *Skilled visions: between apprenticeship and standards*, volume 6 of EASA - Learning Fields, pp. 91–105, Berghahn books.
- Felsberg, M. and Sommer, G. (2004), The Monogenic Scale-Space: A Unifying Approach to Phase-Based Image Processing in Scale-Space, *J. Math. Imaging Vis.*, volume 21(1):pp. 5– 26.
- Gonzalez, R.C. and Woods, R.E. (2008), *Digital Image Processing*, Pearson Education, 3rd edition.
- Goskar, T.A. and Earl, G.P. (2010), Polynomial texture mapping for archaeologists, *British Archaeology*, volume 111:pp. 28–31.
- Henriksson, L., Hyvarinen, A., and Vanni, S. (2009), Representation of Cross-Frequency Spatial Phase Relationships in Human Visual Cortex, *J. Neurosci.*, volume 29(45):pp. 14342–14351.
- Kadir, T. and Brady, M. (2005), Non-parametric Estimation of Probability Distributions from Sampled Signals, OUEL 2283, Robotics Research Laboratory, Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, U.K.
- Kovesi, P. (2000), Phase congruency: a low-level image invariant, *Psychological Research*, (64):pp. 136–48.
- Li, S.Z. (2009), *Markov Random Field Modeling in Image Analysis*, Advances in Pattern Recognition, Springer, 3rd edition.

- Malzbender, T., Gelb, D., and Wolters, H. (2001), Polynomial texture maps, in *SIGGRAPH'01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pp. 519–528, ACM, New York, NY, USA.
- Morrone, M. and Owens, R. (1987), Feature detection from local energy, *Pattern recognition letters*, volume 6(5):pp. 303–13.
- Neer, R. (2005), Connoisseurship and the Stakes of Style, *Critical Inquiry*, volume 32(1):pp. 1–26.
- Oppenheim, A. and Lim, J. (1981), The importance of phase in signals, *IEEE Proceedings, Special Issue on Digital Image Processing*, volume 69(5):pp. 529–41.
- Pan, X., Brady, M., Bowman, A., Crowther, C., and Tomlin, R. (2004), Enhancement and feature extraction for images of incised and ink texts, *Image and Vision Computing*, volume 22:pp. 443–51.
- Pham, D.L., Xu, C., Prince J.L. (1999), A Survey of current methods in medical image segmentation. Technical report JHU/ECE 99-01, Johns Hopkins University, Baltimore, MD.
- Roued-Cunliffe, H. (2010), Towards a Decision Support System for Reading Ancient Documents, *Lit Linguist Computing*, —forthcoming (pre-print retrievable online on the eSAD website—last checked: September 29, 2010).
- Tarte, S.M. (2010), Papyrological Investigations: Transferring Perception and Interpretation into the Digital World, *Lit Linguist Computing*,—forthcoming (pre-print retrievable online on the eSAD website—last checked: September 29, 2010).
- Terras, M. (2006), *Image to Interpretation. An Intelligent System to Aid Historians in Reading the Vindolanda Texts*, Oxford University Press, Oxford.
- Van Leemput, K., Maes, F., Vandermeulen, D., and Suetens, P. (1999), Automated Model-Based Tissue Classification of MR Images of the Brain, *IEEE Transactions on Medical Imaging*, volume 18(10):pp. 897–908.
- Venkatesh, S. and Owens, R. (1990), On the classification of image features, *Pattern recognition letters*, volume 11(5).
- Youtie, H.C. (1963), The papyrologist: artificer of fact, *Greek, Roman and Byzantine Studies*, volume 4(1):pp. 19–33.