

Affine registration of multispectral images of historical documents for optimized feature recovery

Cerys Jones

UCL Department of Medical Physics and Biomedical Engineering,
London, UK

William A. Christens-Barry

Equipoise Imaging LLC, Ellicott City, MD, USA

Melissa Terras

College of Arts, Humanities and Social Sciences, University of
Edinburgh, Edinburgh, UK

Michael B. Toth

R.B. Toth Associates, Oakton, VA, USA

Adam Gibson 

UCL Department of Medical Physics and Biomedical Engineering,
London, UK

Abstract

Multispectral (MSI) imaging of historical documents can recover lost features, such as text or drawings. This technique involves capturing multiple images of a document illuminated using different wavelengths of light. The images created must be registered in order to ensure optimal results are produced from any subsequent image processing techniques. However, the images may be misaligned due to the presence of optical elements such as filters, or because they were acquired at different times or because the images were captured from different copies of the documents. There is little prior work or information available about which image registration techniques are most appropriate. Image registration of multispectral images is challenging as the illumination changes for each image and the features visible in images captured at different wavelengths may not appear consistently throughout the image sequence. Here, we compare three image registration techniques: two based on similarity measures and a method based on phase correlation. These methods are characterized by applying them to realistic surrogate images and then assessed on three different sets of real multispectral images. Mutual information is recommended as a measure for affine image registration when working with multispectral images of documentary material as it was proven to be more robust than the other techniques tested.

Correspondence:

Adam Gibson, UCL
Department of Medical
Physics and Biomedical
Engineering, London, UK.

E-mail:

adam.gibson@ucl.ac.uk

1 Introduction

Multispectral imaging of historical artefacts can recover features that are invisible to the human eye (Easton, 2003; Bearman and Christens-Barry, 2009; Marengo *et al.*, 2011; Liang, 2012). The images are then processed to enhance and combine the information contained in the image sequence. These image processing techniques assume that a pixel in the same position in each image corresponds to a unique point on the object. However, this is not the case if any misalignments were produced during the capture process, for example due to movement of the camera or artefact, the filters introduced into the optical path, or if images are captured on different occasions or setups. Any of these will compromise the accuracy of the subsequent image processing if uncorrected.

In order to process and compare multispectral images, the misalignments must be corrected using image registration (Gottesfeld Brown, 1992; Flusser and Zitova, 2003) which is applied in many fields such as medical imaging (Oliveira and Tavares, 2014), remote sensing (Ma *et al.*, 2015), and computer vision (Wang, 2014). However, registration of multispectral images is challenging as the brightness and contrast vary with different illumination conditions, and features which are present in one wavelength may be absent in another (see examples in (Bearman and Spiro 1996); Easton *et al.*, 2003; Bearman and Christens-Barry (2009); Knox *et al.* (2011)). Furthermore, cameras often have imaging sensors comprising 30 megapixels or more (Bearman and Christens-Barry, 2009; Easton *et al.*, 2010; Knox *et al.*, 2011; Janke and Macdonald, 2014; Bennett, 2015;) that capture high resolution images for which other applications may not be suitable if computer memory, processing power or time is limited. Consequently, a registration method that is effective at aligning multispectral images of heritage objects must be determined.

The use of multispectral image registration for heritage artefacts is reported inconsistently in the literature. Some researchers claim that registration is not required to correct for external filters (Easton, 2003; Marengo *et al.*, 2011; Liang, 2012), many fail to refer to registration at all (Bacci *et al.*,

2005; Agathi-Anthoula and Alexopoulou, 2013; Samadelli *et al.*, 2015) and others describe the process in some detail (Pelagotti *et al.*, 2008; Remondino *et al.*, 2011; Hollaus *et al.*, 2012; Giacometti *et al.*, 2017). There is no consensus on when registration may be needed or on which may be the most appropriate methods to use during investigations of textual materials, unlike other areas of science where registration has been studied intensely (Oliveira and Tavares, 2014).

A broad range of different registration techniques have been applied to heritage imaging. For example, feature-based methods, such as those that use SIFT (Scale Invariant Feature Transform) (Remondino *et al.*, 2011), have been used to register multispectral images of artwork. These techniques rely on identifying features that appear consistently throughout the sequence of images. However, in multispectral images of documents, a feature might not be visible in certain wavelengths and heavily damaged objects such as the Herculaneum scrolls (Ware *et al.*, 2000) may not have any identifiable features. Cross-correlation and related methods (Gottesfeld Brown, 1992; Flusser and Zitova 2003; Fei *et al.*, 2001; Lettner *et al.*, 2007) assume the illumination does not change or only varies by a constant between images, which is not true for multispectral images. Other methods based on statistical properties of the intensities, such as mutual information, are commonly used in medical imaging (Oliveira and Tavares, 2014) and have been successfully applied to multispectral images (Cappellini *et al.*, 2005; Pelagotti *et al.*, 2008; Pronti *et al.*, 2015). Fourier transform based methods which align images in the frequency domain have also been used in the literature (Tonazzini *et al.*, 2009; Bianco *et al.*, 2013). However, papers tend to focus solely on a single registration method and there is no real agreement over the best technique to use.

The search space of transformations for the registration technique must first be determined. For example, the technique can be chosen to search for solely translations or affine distortions. An affine transformation is the combination of a linear transformation and a translation (shift). This includes translations, rotations, scale changes and shear distortions (i.e. transforming a rectangle into a

parallelogram). An affine transformation is uniform across the whole image whereas other non-linear transformations (such as a local stretch) can vary between different areas. Certain distortions, such as translations and rotations (which will be expected if the same object is imaged at different times) or refraction effects (due to light passing through a filter) are described by an affine transformation while others, such as cockling or bending of an object due to changing environmental conditions, require non-linear, local registration. However, affine transformation can provide a close approximation to non-linear transformations and it is common practice for the first stage of a non-linear registration to be an initial affine transformation. In this research, two techniques that register images in the spatial domain using similarity measures (mean-squared differences and mutual information Chen *et al.* (2000)) and one technique that registers images in the frequency domain (the phase correlation technique Reddy and Chatterji (1996)) were compared. The first method was selected due to its computational efficiency and the latter two were chosen as they are invariant to changes in illumination, a necessity when working with multispectral images. The three registration methods were applied to artificially created and distorted multispectral images, and to multispectral images of heritage documents. The performances of the three methods were then compared and assessed.

2 Method

2.1 Multispectral imaging

The UCL Multispectral Imaging System (R B Toth Associates, USA) contains a PhaseOne IQ260 camera (PhaseOne, Denmark) with an 8964×6716 pixel, 16-bit, monochrome digital back and a 120 mm apochromatic lens. LED lighting panels (Equipoise Imaging LLC, USA) illuminate in 12 different wavelengths from 370 nm to 940 nm. The aperture and ISO were set to $f/8.0$ and 200, respectively, ensuring adequate depth of field and detector noise, while the shutter speed varied from $1/6$ s to 30 s. A motorised filter wheel enables the fluorescence to be captured by excluding the illumination wavelengths using violet (400 nm), green (515 nm),

and red (590 nm) longpass filters (ThorLabs, USA). The acquisition was controlled using Spectral XV (Equipoise Imaging LLC, USA) software, integrated with the Capture One (PhaseOne, Denmark) camera software. The images underwent flat-field correction to remove any non-uniformities in the lighting using the Paleo Prep Bar Toolbox (Equipoise Imaging LLC, USA) in ImageJ. The image registration was carried out using MATLAB 2016a (The Mathworks Inc, USA). The computer used for the image processing in this research had an Intel[®] Core[™] i76560U 2.20 GHz CPU and 16 GB RAM.

2.2 Surrogate sequences of images

A legal contract from the year 1869 was chosen as a test object. It is a parchment document with printed text and signatures in iron gall ink and pencil providing a range of distinct features. The document was imaged under green light and segmented into four regions according to the intensities. The four regions were given values for the reflectance spectra for print (Klein *et al.*, 2008), iron gall ink, red ink, and parchment (Knight, n.d.) at 12 wavelengths, which were obtained from the literature. These were assigned to the segmented regions of the original image, creating 12 synthetic multispectral images, all of which were spatially identical but with different, known intensities.

The surrogate sequence was distorted with a range of translations (from zero to 22 pixels in x and y) and rotations (from zero to 1.5 degrees) to mimic the effect of a manuscript or camera moving and rotating across the x - y plane. Three additional images were created by scaling three of the images with a scaling factor between 1.05 and 1.15 to simulate the transformations caused by the filters. Arbitrarily large translations and rotations may occur, due to variable placement of the subject in the camera field of view during image capture, if images are captured on different days (or even on different systems). We have assumed that such distortions are initially corrected by the user, for example by manually aligning the images in image processing software by eye.

The first sequence, consisting of co-registered synthetic images, was used to assess the registration

methods when registering images which are spatially identical but differ in intensity. The second sequence, of spatially distorted images, was used to assess the spatial registration performance. In both cases, the ‘gold standard’ of the initial undistorted images was available, allowing objective, quantitative analysis of realistic images and distortions. Throughout, the green image (519 nm) was taken as the reference image as green is the central wavelength of the lights and so the differences in intensities between the reference and target images should be minimized.

2.3 Images of historical documents

The surrogate documents were used to provide objective and quantitative criteria of registration quality as the correct ‘target’ solution is known. This is not the case in imaging of historical documents; consequently, in order to inform the applicability of the technique to real objects, three documents were chosen to test different applications of image registration on historical documents. These were chosen to demonstrate a variety of contexts where image registration may be valuable, ranging from minor effects (refraction due to the presence of filters in the optical path) through an object that may show non-linear distortion (due to humidification) to a more challenging application of image registration of different objects imaged at different times.

(1) ‘PEARSON’ is a printed letter dated 1929 held by UCL Special Collections (accession number GALTON/LAB/3/1/3 FOLDER 5). The letter is from Sir James Purves Stewart to Professor Karl Pearson, a well-known mathematician who established the world’s first Department of Statistics at UCL (Norton, 1978). The document had previously been exposed to water, resulting in mould growth that rendered much of the text illegible. This test showed the importance of image registration to compensate for the misalignments that are present due to the filters used during image acquisition and the effect on subsequent image processing.

(2) ‘CHADWICK’ is a page from a manuscript written in 1853 using iron gall ink which was a report to the Metropolitan Commission of Sewers, held by UCL Special Collections (accession number

CHADWICK/45-66/55-60/56). It is almost illegible under room lighting but can be read clearly under ultraviolet. It was imaged before and after an experiment in which the paper was humidified (at 90%, 70%, and 50% humidification) and then dried. This tested the use of registration on a single object whose shape, appearance, and position could not be assumed to be constant.

(3) ‘ALDERMEN’ is a collection of 15th Century ink and watercolour drawings of Aldermen of the city of London, held by London Metropolitan Archives (accession number SC/GL/ALD/001). They all have similar designs and it is believed they were drawn from a template (“Wards - City of London.”). Multispectral imaging was used to enhance the pen outlines and image registration was performed to compare the outlines for the different drawings (Payne and Smith, 2014). This tested the ability of the techniques to register images of different objects acquired at different times for comparison.

2.4 Image registration

2.4.1 Mean squared differences

Mean squared differences is one of the simplest similarity measures used for image registration but is rarely used to register multispectral images. It relies on the assumption that two images are identical other than a spatial transformation and Gaussian noise, but this assumption does not hold for multispectral images because different substances respond differently to the range of wavelengths used and thus may not appear consistently. However, due to its simplicity and computational speed, the mean squared differences method was implemented and tested.

The mean squared differences measure involves finding the difference between the intensity values at each pixel in the target and reference images, then squaring and averaging them.

For two images I and J , each made up of $m \times n$ pixels, the mean squared differences measure, MSD , is

$$MSD(I, J) = \frac{1}{nm} \sum_{j=1}^n \sum_{i=1}^m (I(i, j) - J(i, j))^2$$

where i, j are the row and column indices.

When the mean squared differences measure is minimized, the images are assumed to be registered. The algorithm used to perform the minimization was gradient descent (Eastman and Le Moigne, 1998), which finds the minimum of a function by iteratively taking steps in the direction of the negative gradient. Gradient descent is susceptible to finding local minima, however, as the mean squared differences measure is a convex function, any local minimum must also be the global minimum. The algorithm was implemented using MATLAB's 'imregister' function (MathWorks, 2017, p. 1246) with the 'monomodal' input and 'affine' parameter. A multiscale approach was used to increase the computational speed by first registering a low-resolution image and then increasing the resolution in steps (Chen *et al.*, 2000). 100 iterations were performed at each of the five steps.

2.4.2 Mutual information

Mutual information has been used to register multispectral images in cultural heritage (Pelagotti *et al.*, 2008; Pronti *et al.*, 2015). Instead of dealing with the intensity values directly, mutual information uses statistical properties known as the marginal and joint entropies of the intensity values to compare the images. The marginal entropy of an image I is defined as $H(I) = -\sum_{x \in X} p(x) \log p(x)$, where $p(x)$ is the probability that intensity value x occurs in image I , and X is the set of intensity values in image I . The joint entropy of images I and J , $H(I, J) = -\sum_{x \in X} \sum_{y \in Y} p(x, y) \log p(x, y)$, measures the average uncertainty in I and J simultaneously, where $p(x, y)$ is the joint probability of intensity values x and y both occurring in images I and J , and X, Y are the set of intensity values in images I and J , respectively. The mutual information (Shannon, 1948), MI , between two images I and J is then

$$MI(I, J) = H(I) + H(J) - H(I, J).$$

As two images become more similar to each other, they share more information, and thus the mutual information increases.

Mutual information can have local maxima, so an evolutionary optimiser was applied instead of the gradient descent method used previously (Maier *et al.*, 2006). An evolutionary optimiser adjusts the

initial parameters by randomly selecting values that are within a search radius. If this new transformation increases the mutual information by providing a better alignment of the two images, then its parameters are retained and it becomes the new centre of the search area for the optimiser. If this transformation does not improve the alignment, the optimiser continues from the previous search area. The specific optimiser used in this research to minimize the negative of the mutual information was the (1+1) evolutionary algorithm (Styner *et al.*, 2000).

As mutual information depends on the number of grey levels in the images, computing the mutual information of a 16-bit image is computationally expensive. Therefore, the registration algorithm was first performed on an 8-bit image to provide an initial registration transformation. This was then applied to the 16-bit image and the registration algorithm was performed again on the transformed 16-bit image. To further increase the speed, multiscale registration was performed on the 8-bit image with 50 iterations for three steps, then the registration was performed on the 16-bit image with 150 iterations at each of five steps. This multiscale approach took only 14 min to register the sequence of images rather than 37 min when registering only the 16-bit images with the same number of iterations.

Mutual information was implemented using MATLAB's 'imregister' function (MathWorks, 2017, p. 1246) which uses Mattes' method (Mattes *et al.*, 2003), with the 'multimodal' and 'affine' parameters. Imregister only terminates the registration when the number of iterations is completed and does not provide for any alternative stopping criteria. The mutual information was plotted against number of iterations and attained its maximum after 600–750 iterations, confirming that 150 iterations at each of five stages provided a reliable compromise between registration accuracy and computation time.

2.4.3 Phase correlation

The phase correlation method registers images in the frequency domain, instead of working in the spatial domain as in the previous two methods.

The phase correlation method uses the property that multiplications in the frequency domain are equivalent to convolutions in the spatial domain to calculate the cross-correlation of the images and thus determine the extent of the translation. This involves calculating the cross-power spectrum of the Fourier transform of both images and then taking the inverse Fourier transform. The location of the peak of the inverse Fourier transform gives the translation parameters. The cross-power spectrum for two images I and J with Fourier transforms F and G , respectively, is defined as follows:

$$\frac{F(u, v)G^*(u, v)}{|F(u, v)G(u, v)|} = e^{i2\pi(ut_x + vt_y)}$$

where image J differs from image I by a translation of (t_x, t_y) , i is the complex number, and G^* is the complex conjugate of G .

The images must first be corrected for any changes in rotation and scale before the translation can be determined. Prior processing with a log-polar transform in the frequency domain provides the rotation and scale changes (Reddy and Chatterji, 1996).

Registration methods using Fourier transforms can register images captured under different conditions, such as multispectral images, and are resilient to noise (Flusser and Zitova, 2003). The method was implemented using the MATLAB function ‘*imregcorr*’ (MathWorks, 2017, p. 1230) with input ‘*similarity*’.

2.5 Analysis of registration results

The registration algorithms were tested objectively and subjectively on the surrogate and distorted sequences, and subjectively on the three sets of real multispectral images.

The subjective analysis of the registration was completed by creating a false-colour image between the original, undistorted image and the registered image. If the images were misaligned, the misalignments would appear in colour, whereas perfectly aligned images would appear in grey.

Quantitative analysis was performed by analysing the transformation matrices determined by the registration techniques. As the surrogate images were distorted by specified transformations, the

inverse transformation that was required to register each image is known. The transformations were compared using the Frobenius norm.¹ This is done by summing the differences squared between each entry in the correct transformation and the transformation calculated by the registration technique and then taking the square root. If this value is 0, the transformations are equal and thus the registration technique perfectly aligned the images. The norm increases as the accuracy of the registration decreases. As the norm of the difference between the transformations only depends on the translation, rotation and scale parameters, it is independent of the image size, bit-depth, interpolation error, and changes in intensity value.

3. Results

3.1 Surrogate images with differing intensities

The 12 surrogate images with no spatial distortions were registered to determine the performance of the registration algorithms when registering a series of images with different intensities. The Frobenius norm of the transformations calculated by the mean squared differences measure were large and ranged between 1030 and 1090 (average 1050) and thus represents a failure of the mean squared differences measure to register every surrogate image. Mutual information and phase correlation both gave perfect registration in all cases. The mean squared differences measure was therefore excluded from further analysis.

3.2 Surrogate images with differing intensities and spatial distortion

The 15 surrogate images with translations, rotations, and scale were registered using mutual information and phase correlation. The Frobenius norm for the images registered with phase correlation ranged between 0 and 322.5 (average 23.4), however, when excluding the single image with the largest value (322.5) which corresponded to the image with the largest scale factor, the average reduced to 0.4, suggesting that the remaining images were accurately registered. For the images registered with mutual



Fig. 1 Results of image registration showing the surrogate images superimposed on the original, unregistered image (left), registration with mutual information (centre) and registration with phase correlation (right). Top row shows the surrogate image with translation (3 pixels in x and 5 pixels in y); second row shows the surrogate image with translation (15 pixels in x and 13 pixels in y) and rotation (0.6°); third row shows the surrogate image with translation (3 pixels in x and 5 pixels in y) and scale ($\times 1.15$); and bottom row shows the surrogate image with translation (22 pixels in x and 21 pixels in y), rotation (1.5°) and scale ($\times 1.05$)

information, the Frobenius norm ranged between 0.1 and 0.6 (average 0.3).

Figure 1 shows false-colour images of four of the surrogate images with the corresponding target images (left), registered images by mutual information (centre), and registered images by phase correlation (right). Images registered with mutual information are visibly indistinguishable from the target images whereas slight misalignments remain

for the third example (translation and scale) after registration with phase correlation. This agrees with the Frobenius norm values, which were similar for mutual information and phase correlation, except for the image that was significantly higher for the registration with phase correlation. Registration with phase correlation was faster than with mutual information, taking 1.5 min compared to 14 min.

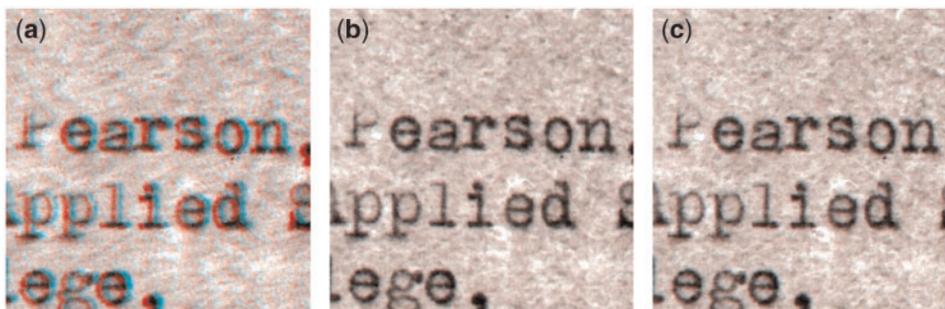


Fig. 2 Crop of image from the Pearson letter showing (a) unregistered images; (b) images registered using mutual information; (c) images registered using phase correlation

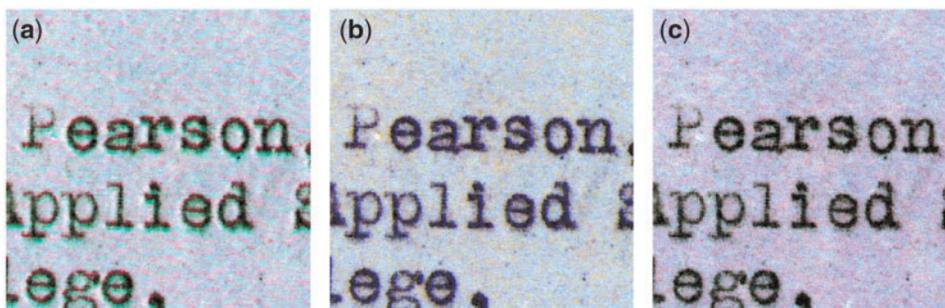


Fig. 3 Crop of the principal components from the Pearson letter superimposed on each other showing the effect of (a) no registration; (b) registration using mutual information; (c) registration using phase correlation

3.3 Images of historical documents

3.3.1 *Pearson*

The Pearson letter was imaged illuminated in the twelve wavelengths from 370 nm to 940 nm and three different filters were used to capture the fluorescence from the ultraviolet and deep blue (448 nm) illumination, giving a total of seventeen images. For this study, 400 x 400 pixel crops of the images taken in 370 nm with and without a violet 400 nm long-pass filter were analysed. The image taken at 370 nm without a filter was assigned to the red channel and the equivalent image with a filter was assigned to the green and blue channels of the images shown in Fig. 2. Lines that are well registered appear once, whereas multiple edges are due to the misalignments in the images.

An area of the document containing the printed text was chosen so that the edges of the letters

would clearly show the extent of any misalignment. Figure 2a shows the unregistered images in which the image with a filter appears as a shadow due to its misalignment. Registration by mutual information (Fig. 2b) and phase correlation (Fig. 2c) successfully corrected the misalignments.

The images were analysed using Principal Component Analysis (PCA). Three principal components were chosen and assigned to the red, green and blue channels of the images shown in Fig. 3. The misalignments can be seen in Fig. 3a. Fig. 3b shows the result of PCA applied to images registered using mutual information and Fig. 3c shows PCA applied to images registered using phase correlation. Fig. 3b–c show clearer edges with no ‘shadow’, demonstrating that misalignments in the initial images will also be present in the processed images.

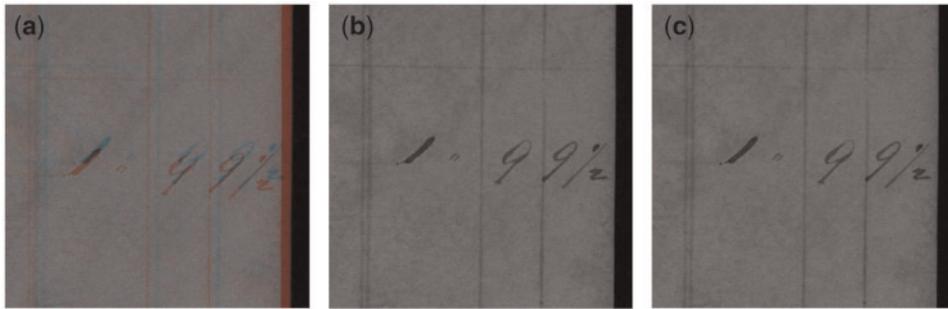


Fig. 4 Crop of handwritten figures from CHADWICK showing the effect of (a) no registration; (b) registration using mutual information; (c) registration using phase correlation



Fig. 5 Drawings of Alderman John Norman (left) and Alderman Simon Eyre (right)

3.3.2 CHADWICK

Multispectral images of the Chadwick document were captured illuminated in the multiple wavelengths. Those of the document after the humidity treatment were registered to the image taken at 519 nm before treatment. False colour images were created from crops of the images taken at 519 nm before and after humidifying. The image taken

before treatment was placed in the red channel and the image captured after was placed in the green and blue channels of the images shown in Fig. 4. Registration for scale, translation and rotation was implemented. The mutual information and phase correlation techniques both successfully registered the images before and after treatment.

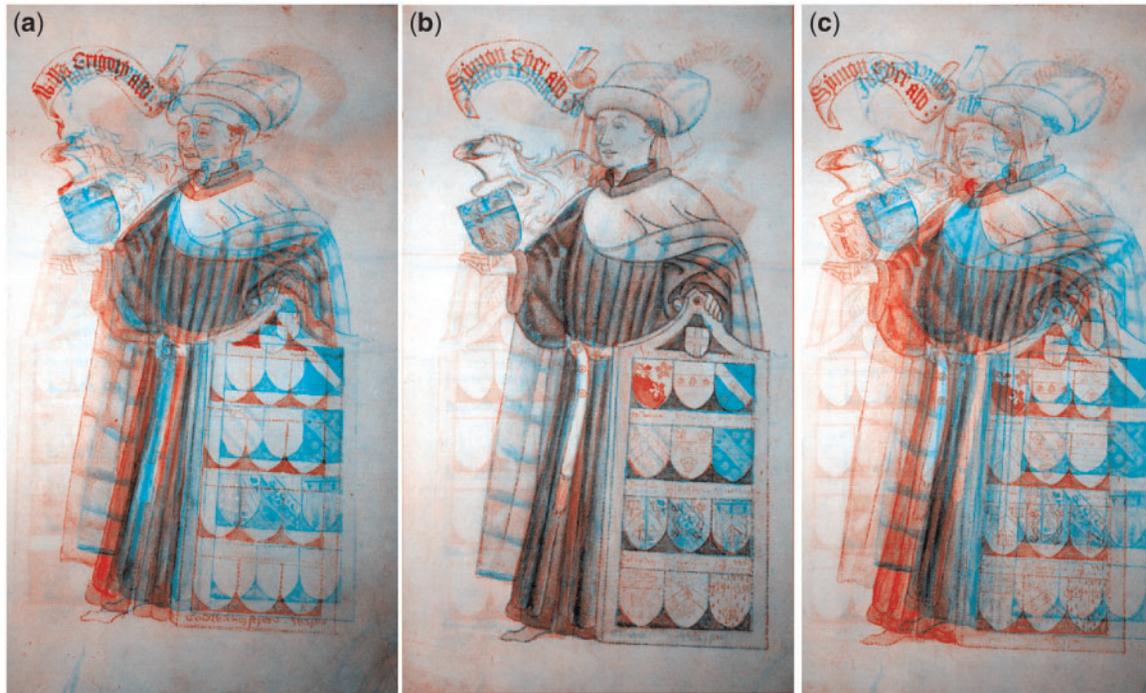


Fig. 6 Registration of the infrared images of Alderman Simon Eyre to Alderman John Norman, showing from left (a) unregistered images; (b) images registered using mutual information; (c) images registered using phase correlation

3.3.3 ALDERMEN

Figure 5 shows two examples of the Aldermen watercolour drawings. Corresponding images acquired under infrared light (940 nm) clearly showed the outlines of the figures, whereas the images captured under visible light also contained other features from the paint. Images of the drawing shown in Fig. 5b under infrared light with wavelengths centered at 940 nm were registered to the infrared image of the drawing in Fig. 5a using mutual information and phase correlation. To test the hypothesis that a template was used to draw the outlines, the registration was limited to rotations and translations only as a template would not lead to scale or shear transformations.

Images acquired under infrared illumination were most sensitive to underdrawings and therefore to the template. False colour images created by placing the infrared image of Alderman John Norman (shown in Fig. 5a) into the red channel and the equivalent image of Alderman Simon Eyre (in

Fig. 5b) into the green and blue channels. These are shown as Fig. 6.

Mutual information successfully registered the images of the two Aldermen, however, phase correlation failed to accurately align them. Examination of Fig. 6b revealed that there is a good alignment of the outlines of the Aldermen's clothing and shields, but poor alignment of their hats and scrolls. The outlines of their faces are aligned, but their features are less so (see crops in Fig. 7b). This indicates that the template included the clothing, shields, and face outlines, but the hats and scrolls were drawn freehand.

The timings for the four sequences of images are contained in Table 1. Although the length of time taken for both algorithms to complete depended on the size of the images, the type and size of the distortions, the processing capacity of our system and implementation of our code, the mutual information algorithm is invariably slower than the phase correlation algorithm due to the computational



Fig. 7 Crops of fig. 6 showing the faces of the Aldermen. In fig. 7b, the collars, outlines of the faces and lower parts of the hats are aligned but the top of the hats, features of the faces and the scrolls in the top left are not, showing that the drawings differ in those areas

Table 1. The timings for the mutual information and phase correlation methods on the surrogate images and the three case studies (in minutes)

	Surrogate images	PEARSON	CHADWICK	ALDERMAN
Mutual information	14 min	9 min	10 min	6.6 min
Phase correlation	1.5 min	3 min	3 min	1.5* min (*failed to register)

complexity of the mutual information measure. Furthermore, the timing of the mutual information method also depends on the number of iterations and the size of the search radius for the evolutionary optimiser. Therefore, although the exact timings in Table 1 are dependent on the case studies and implementation of the algorithms, the mutual information technique can be assumed to be considerably slower.

4 Discussion

Registration using mean squared differences as a measure failed for the surrogate images due to the change in intensity between the reference image and target images. It was expected that the performance of this method would be poor but the extent of the failure meant its use in multispectral imaging cannot be recommended.

Both mutual information and phase correlation successfully registered the surrogate images. When these images were spatially distorted, the mutual information method correctly aligned all of the spatial distortions with an average Frobenius norm of 0.4. The phase correlation method also correctly registered most of the spatial distortions, however failed for the image with the largest scale factor, 1.15. The surrogate images with and without spatial distortions enabled quantitative analysis of the registration accuracy to be completed.

Mutual information and phase correlation were applied to several multispectral image sequences for different image registration tests. The PEARSON test (Figs 2 and 3) showed that even during static imaging of an object, image registration is necessary if the optical path is distorted in some images, for example by filters. Both the mutual information method and the phase correlation method were able to accurately register the Pearson images, however the phase correlation method registered the

images considerably faster, taking only 3 min to register each image as opposed to 9 minutes for the mutual information method. Applying principal component analysis to the images showed that any spatial distortions present in the original images are also present in the subsequent processed images and thus the images must be registered before any processing techniques are applied.

The CHADWICK test (Fig. 4) demonstrated the use of image registration on a single document imaged on different occasions. Both techniques successfully registered the Chadwick images and, as before, phase correlation was faster taking only 3 min instead of the 10 min that the mutual information method required.

The ALDERMEN test showed that registration can be used even with different objects, provided they have some points in common, however, in this case only the mutual information method was able to register the images and the phase correlation method failed for every image. The phase correlation method is known to fail for large rotations and scale changes due to aliasing in the low frequencies, which produces false peaks and reduces the peak at the correct transformation, resulting in an inaccurate alignment (Stone *et al.*, 2003). This may explain why the phase correlation method failed to accurately register the spatially distorted surrogate image with the largest scale factor. However, the degree of rotation for the Aldermen images was small and there was no change in scale and thus this is unlikely to be the reason the registration failed in this case. As the Aldermen images contain several periodic structures, such as the shield containing the coat of arms, several peaks were present in the cross-power spectrum. Additionally, significant differences in the content of images (e.g. the features of the coat of arms and styles of hats vary) reduces the presence of a single, narrow peak in the cross-power spectrum that the Fourier shift property (Reddy and Chatterji, 1996) would predict for images that differed only in their registration. Phase correlation was substantially faster than mutual information as it does not require multiple iterations and is computationally cheaper. It has been proposed as an initial step prior to final registration by mutual information (MathWorks, n.d.).

5 Conclusion

Three different approaches to affine image registration have been assessed for multispectral images of documents. The methods were first tested on surrogate images, for which an absolute, objective test was available, and on images acquired of real historical documents. It was found that the mutual information method was most successful. The simplest method tested, which minimised the mean squared differences measure, failed due to the changes in intensity. The phase correlation method successfully registered the PEARSON and CHADWICK images but failed to correct the misalignments in the ALDERMAN images and the surrogate image with large scale distortions. The phase correlation method was also faster than the mutual information method as it is not an iterative method and is computationally cheaper.

Mutual information has been shown to be robust at registering multispectral images. However, if time is of concern, a pre-processing step using phase correlation could be implemented when registering images of the same object. We also recommend that image registration processes are adequately documented when reporting on analysis of multispectral images of historical documents.

This work only considered affine registration, i.e. translations, rotations, scale and shear transformations applied to the entire image. The range of such transformations was limited in this study to realistic cases likely to be encountered by investigators in practice was assessed; applicability of the results and conclusions of this study to material distortions of greater extent or that represent particular types of distortions or materials awaits further analysis. In the case of non-affine distortions, such as local distortions that might occur if paper or parchment is heated, torn or otherwise distorted, other non-linear image registration methods may yield improved results, although mutual information methods could still provide an initial first step.

Acknowledgements

We are very grateful to the conservators and archivists who provided access to samples and advised

