

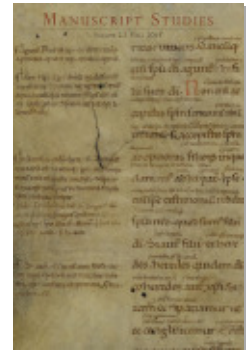


PROJECT MUSE®

The St. Chad Gospels: Diachronic Manuscript Registration and Visualization

Stephen Parsons, C. Seth Parker, W. Brent Seales

Manuscript Studies: A Journal of the Schoenberg Institute for
Manuscript Studies, Volume 2, Number 2, Fall 2017, pp. 483-498
(Article)



Published by University of Pennsylvania Press

DOI: <https://doi.org/10.1353/mns.2017.0022>

➔ *For additional information about this article*

<https://muse.jhu.edu/article/676177>

The St. Chad Gospels

Diachronic Manuscript Registration and Visualization

STEPHEN PARSONS
University of Kentucky

C. SETH PARKER
University of Kentucky

W. BRENT SEALES
University of Kentucky

WELL-MAINTAINED MANUSCRIPTS CAN APPEAR to have always existed in their current, “perfectly preserved” archival state. The reality, however, is that no manuscript can be completely and permanently protected from alteration. All documents change with the passage of time. Left untouched, documents will suffer physical degradation—discoloration and fading, breakdown of fibers, flaking of pigment. Yet, conservation efforts designed to combat such degradation introduce physical changes of their own, such as crumbling, tearing, or staining.

A better understanding of the changes introduced by the conservation and degradation of manuscripts can provide more accurate knowledge of the documents’ histories as well as help improve conservation efforts. Yet, such understanding necessitates concrete identification of the physical changes that have occurred to manuscripts. This task proves quite challenging since

these alterations are usually small or subtle and occur over extended periods of time. Recent digital methods like high-resolution imaging make it possible to capture the changes, as manuscripts are imaged multiple times and in multiple modalities. However, it remains difficult to discover and directly visualize those changes throughout the collections of images. Post-processing and visualization methods applied to sets of images over many instances are needed to allow scholars to “see” what exactly happens to manuscripts over time and to benefit from this knowledge.

The 2010 digitization of the St. Chad Gospels by the Seales research team offers a perfect opportunity to study how best to align and visualize layered image sets. The St. Chad Gospels is an eighth-century manuscript held in the Lichfield Cathedral’s library in Lichfield, England. It is historically and culturally significant and contains a wealth of information. Among its 236 surviving folios are four pages of framed text, eight illuminations, and marginalia that include some of the oldest surviving Old Welsh writing.¹ The research team led by Seales in 2010 collected spectral data and 3D shape models of every page,² as well as documentary photography, video, and historical image sets from earlier photographic sessions: 1912, 1929, 1962, and 2003.³ Although the number of pages imaged each time in the historical image sets varies (1912: 9; 1929: 240; 1962: 238; 2003: 181; and 2010: 264), the composite of all these acquisitions makes a rich collection for seeing change over time.

Viewing these images in the visual context of each other can add information beyond what a simple facsimile would provide.⁴ Each time the

1 John Davies, ed., *Encyclopaedia Wales* (Cardiff: University of Wales Press, 2008), 577.

2 See Daniel Staley, “Multi-Spectral and 3D Data Acquisition of Antiquated Manuscripts,” *MS Project Report, University of Kentucky Department of Computer Science*, 3 November 2010, and Jennifer Howard, “Cutting-Edge Imaging Helps Scholar Reveal 8th-Century Manuscript,” *Chronicle of Higher Education*, 5 December 2010, <http://chronicle.com/article/Cutting-Edge-Imaging-Helps/125616>, accessed 23 August 2016.

3 1912: Images acquired under unknown circumstances. 1929: Images acquired by the National Library of Wales, Aberystwyth. 1962: Roger Powell, “The Lichfield St. Chad’s Gospels: Repair and Rebinding, 1961–1962,” *The Library* 5, no. 4 (1965): 259–65. 2010: Images acquired by the British Library.

4 W. Brent Seales and Steve Crossan, “Asset Digitization: Moving Beyond Facsimile,” *SIG-GRAPH Asia 2012 Technical Briefs*.

manuscript is imaged, unique information is captured, but the value of that information increases when it is viewed in the context of other years' photographs, because differences can then be observed. Since imaging conditions (for example, field of view, focal length, and image resolution) can vary dramatically across acquisitions, providing this context requires both an alignment of the images as well as a visualization tool that can enable meaningful comparison. The known problem of image registration is well-suited to aligning the images, and a viewer can be built to help make sense of the registered results.

Our objective, therefore, is to combine the data sources, creating a complete version of the manuscript that encompasses all known images. Specifically, we focus on the diachronic axis of the image sets (that is, their change over time) and attempt to organize them in a way that allows meticulous observation of the manuscript's evolution over time. We also present a registered image viewer tool that enables direct comparison of the original images while maintaining their visual context.

Related Work

IMAGE REGISTRATION

Image registration is a well-studied problem with a wide array of applications, largely in medical imaging. Image registration is the process of aligning two images by mapping a *sensed* or moving image into the coordinate system of a *reference* or fixed image R , yielding a registered image S_R .⁵ The images are of the same scene or object, but they may differ by time, modality, viewpoint, illumination, or a number of other parameters.

5 See Lisa Gottesfeld Brown, "A Survey of Image Registration Techniques," *ACM Computing Surveys* 24, no. 4 (1992): 325–76; Siddharth Saxena and Rajeev Kumar Singh, "A Survey of Recent and Classical Image Registration Methods," *International Journal of Signal Processing, Image Processing and Pattern Recognition* 7, no. 4 (2014): 167–76; and Barbara Zitová and Jan Flusser, "Image Registration Methods: A Survey," *Image and Vision Computing* 21, no. 11 (2003): 977–1000.

In the case of the St. Chad Gospels, the 2010 photograph of each page is treated as the reference, and each additional image of that page is individually registered to the reference. The 2010 pages are chosen as the reference set because they are the most recent, most complete, and highest resolution of the image sets. The 2010 set is also multispectral and includes multimodal images captured simultaneously with the visible-spectrum photographs. Since these multimodal images were captured without moving the manuscript or the camera, they are already aligned or registered to the base 2010 images.

REGISTRATION OF MANUSCRIPT IMAGES

In 2008, the Archimedes Palimpsest project took a number of image sets, including Heiberg's 1906 photography of the palimpsest, and registered them to high-resolution multispectral photos taken in 2007.⁶ The study similarly addressed a registration problem over a diachronic axis. This manual registration process involved scaling and warping, and presumably relied upon an affine transform, meaning it would not have accounted for any nonlinear warping that might have occurred between images.

In 2009, Baumann and Seales applied registration tools to manuscript pages in an effort to overcome a problematic camera sensor.⁷ While creating a set that included a photo and a 3D scan of each page (taken simultaneously), they discovered that a dirty sensor on the camera had produced some images of unacceptable quality. They re-photographed those pages, but, due to time constraints, they were unable to perform new 3D scans of the pages. Instead, they registered the high-quality sensed images to the "dirty" refer-

6 See Reviel Netz and William Noel, *The Archimedes Codex: Revealing the Secrets of the World's Greatest Palimpsest* (London: W&N, 2008); and M. B. Toth and Doug Emery, "Infinite Possibilities: Archimedes on the Web," *Google TechTalks*, 8 September 2008, <http://www.youtube.com/watch?v=z3fZdIw-s-E>.

7 Ryan Baumann and W. Brent Seales, "Robust Registration of Manuscript Images," *Proceedings of the 9th ACM/IEEE-CS joint conference on Digital libraries (JCDL 2009)*: (New York: ACM, 2009).

ence photos taken earlier. Doing so allowed them to texture the 3D scans with the high-quality images and then perform digital flattening techniques to remove imperfections in the images due to warped pages. They used warping with a triangulated mesh, which was an improvement over the affine transform used with the Archimedes Palimpsest.

These projects show a marked improvement in the preservation and visualization of ancient manuscripts, but leave room for further work. To be most practical, a diachronic edition of a manuscript should be complete, well aligned with high-quality registration, and accessible via an interactive and useful interface.

VISUALIZATION OF REGISTERED IMAGES

Making use of registered images requires a method for viewing the reference image R and the registered image S_R so that their differences are readily apparent. In many cases, an observer wants to see how a specific feature of the manuscript has changed over time. Such discoveries can be made by identifying a point on R and asking how the corresponding point appears on S_R .

The most obvious way to accomplish such comparison is to examine R and S_R side by side—for example, with each occupying half of the view area. While this method places both images directly in front of the observer at the same time, it remains difficult to directly compare the specific points of interest. In switching attention from one image to the other, the observer loses their visual context and position. This method does not take advantage of the images being registered to the same coordinate system, and an observer could just as well use R and the original sensed image S .

This limitation can be overcome by “flickering” R and S_R . The observer can view one image and then switch the display to show the other in its place. Since S_R is registered to R , a point of interest will occupy the same spot on the screen regardless of which image is being viewed. By focusing on one point and switching back and forth repeatedly, an observer is able to see how that area changes across the images. However, entirely switching out the images remains overwhelming to the observer, who is tasked with mentally tracking the points and trying to remember what they looked like

on the other image. The result is a confusing visualization that does not clearly illustrate differences.

Another similar approach is to somewhat combine the side-by-side and flicker methods. Here, two images are stacked, and instead of switching completely between them, the user uses a sliding bar over the images to choose how much of each image to view. However, this process, like the others, allows comparison along only one axis at a time.

Composite images are one of the most common methods for visualizing registered image data, and they have been proven to overcome some of the above-mentioned challenges. This method combines the two images into a third image that contains information from both. One type of simple composite image is the difference $|R - S_R|$, or “diff” image, which is simply the absolute difference between the images at each pixel. Where the images are identical, meaning no change is revealed across registered images, the composite image is typically dark. Bright spots indicate a difference between the images that could then be examined on either original. More complex composite images exist that offer an enhanced visualization when compared with that of the absolute diff.⁸ However, no composite image can ever fully display the original data, since the combining process reduces the data of multiple images into one impression.

Approach

IMAGE REGISTRATION

The registration task for the Chad Gospels dataset was to produce a set $P = \{R, S_{R1}, S_{R2}, \dots, S_{Rn}\}$ of aligned images for each page of the manuscript,

8 D. P. Wallace, Y.-P. Hou, Z. L. Huang, E. Nivens, L. Savinkova, T. Yamaguchi, and M. Bilgen, “Tracking Kidney Volume in Mice with Polycystic Kidney Disease by Magnetic Resonance Imaging,” *Kidney International* 73, no. 6 (2008): 778–81. Kelly Rehm, Stephen C. Strother, Jon R. Anderson, Kirt A. Schaper, and David A. Rottenberg, “Display of Merged Multimodality Brain Images Using Interleaved Pixels with Independent Color Scales,” *Journal of Nuclear Medicine* 35, no. 11 (1994): 1815–21.

where the 2010 image and $0 \leq n \leq 4$ is the number of sensed images for that page, with each S_{Rn} from 1912, 1929, 1962, or 2003. This set was produced by beginning with the 2010 image R for each page, and generating a pair (S, R) for each non-2010 image of the same page. Each pair (S, R) was then passed through the registration pipeline to produce an S_R , and these results were combined to produce P for each page.

The generalized registration process can be broken down into the following steps, each of which can be implemented in various ways: feature detection, feature matching, transform model estimation, and image re-sampling and transformation.⁹ In transform modeling, the transformation is typically either *rigid-body* (uniform affine transformations such as rotation, translation, and scaling) or *non-rigid* (also known as nonlinear, local, or deformable). The appropriate method depends on the data and how it has changed between the sensed and reference images.

Feature Detection and Matching

The nature of the Chad Gospels dataset makes it possible to manually select matching feature points for each pair of sensed and reference images. The data set is small enough to make this tractable, and the consistent image structure allows for a straightforward process.

The registration of each image pair (S, R) starts with the identification of at least five matching points between the two images. Using a visual interface makes it straightforward for a person to select five points on each image: each of the four corners, plus a more central point. In practice, we tried to select the corners to be the actual corners of the physical manuscript page, while the central point was chosen to be some identifiable feature, such as the tip of the serif on a central letter. This process resulted in a set of landmarks $L_{S,R}$, mapping pixel indices S on to their corresponding indices on R :

$$L_{S,R} = \{((xi_S, yi_S), (xi_R, yi_R)) \mid 0 \leq i < l\}, l = \text{no. of landmarks for pair } (S, R)$$

9 Zitová and Flusser, “Image Registration Methods.”

Occasionally after registration, particular images exhibited regions that did not converge to alignment. In these cases, it was possible to manually insert additional feature points in the problematic region, which proved adequate for improving the resulting registration.

The use of automated feature detection and matching, using an algorithm such as scale-invariant feature transform (SIFT),¹⁰ was explored but ultimately proved unnecessary for this particular set of images. However, should the same methods be applied to an image set of a larger scale or with more variation in content, an automated approach would be appropriate.

Registration

Registration methods vary considerably, and their utilization depends on the underlying data. For example, in some cases it may not be appropriate to introduce nonlinear transformations to a sensed image. In our study, however, such a transformation was considered acceptable due to the value of registering letterforms on the page to their counterparts in images of the same page taken at different times. The Chad Gospels photographs were taken with different camera sensors and lenses, from different distances and angles. In addition, each folio is physically warped from the effects of time and age, creating page topologies that are inconsistent across imaging sessions. These subtle differences in the resulting images combine to create a registration problem that not only allows, but actually requires, nonlinear transformations in order to achieve the desired effect.

Our registration framework uses the Insight Segmentation and Registration Toolkit (ITK)¹¹ and consists of landmark and deformable registration stages. The process on the Chad Gospels began with the input images

10 David G. Lowe, "Object Recognition from Local Scale-Invariant Features," *Proceedings of the International Conference on Computer Vision*, vol. 2, 1999, 1150–57. doi:10.1109/ICCV.1999.790410.

11 Terry S. Yoo, Michael J. Ackerman, William E. Lorensen, Will Schroeder, Vikram Chakraborty, Stephen Aylward, Dimitris Metaxas, and Ross Whitaker, "Engineering and Algorithm Design for an Image Processing API: A Technical Report on ITK—the Insight Toolkit," *Studies in Health Technology and Informatics*, 85 (2002): 586–92. Toolkit available at www.itk.org.

S and R and the landmarks $L_{S,R}$. The pixel indices in $L_{S,R}$ were converted to points in the physical space used for registration in ITK.

The images were converted to grayscale for the registration process, which made it easier to compute a similarity metric at each iteration:

$$S' = \text{grayscale}(S), R' = \text{grayscale}(R)$$

For the landmark registration stage, an initial transformation was applied to S' that aligned each of its landmark points with its corresponding point on R . The remaining points on the images were interpolated using a displacement field built by a kernel-based spline function. This transformation T_1 creates S_{LW}' , which is roughly aligned with the reference image R in size and space. While the landmark feature points for the Chad Gospels were all registered perfectly at this point, the rest of the page remained poorly aligned. Due to the nature of the deformation, different regions of the page were misaligned independently of each other.

$$\text{Landmark registration stage: } T_1(S') \rightarrow S_{LW}'$$

The deformable registration stage was then performed using a regular step gradient descent optimization with a relaxation factor of 0.85. The deformation was computed using a basis spline or B-spline transform with spline order 3. The alignment was evaluated using the Mattes mutual information metric¹² with 50 histogram bins and the number of samples at each iteration being 1/80 the number of pixels in the reference image. This sampling volume can be increased for more accurate registration at the cost of registration program runtime.

There are multiple possible stop conditions in this configuration:

- The maximum number of iterations (default 100) is reached.
- The optimizer step size falls below the threshold value.

12 D. Mattes, D. R. Haynor, H. Vesselle, T. K. Lewellen, and W. Eubank, "PET-CT Image Registration in the Chest Using Free-Form Deformations," *IEEE Transactions on Medical Imaging* 22, no. 1 (2003): 12028.

- The change in metric between successive iterations falls below the threshold value.

The deformable registration stage yielded the transformation T_2 , which gave a grayscale registered image S_R' from S_{LW}' :

$$\text{Deformable registration stage: } T_2(S_{LW}') \rightarrow S_R'$$

The full-color registered image was generated by applying the transformations to the original sensed image:

$$\text{Registration process: } T_2(T_1(S)) \rightarrow S_R$$

The composite transformation $T = T_2(T_1)$ was also recorded in a file. This eased development and experimentation, as one can reapply the resulting transformation to a source image without the computational expense of calculating the transformation again. Additionally, the transformations may themselves hold data worth investigating in future works.

Results

REGISTRATION

Results of the registration process for the alignment of Chad Gospels pages are shown in figure 1, with page 219 used as an example. Diachronic registration was successful for pages with text as well as for those containing illuminations. These registered images along with the layered image viewer allow researchers to quickly move through the document to identify and examine points of interest that reveal diachronic changes in the artifact. For example, figure 2 reveals pigment that chipped off between the years 1929 and 1962, and again from 1962 to 2003. The small spot of dark pigment in the middle can be seen reducing in size as it chips away from the lighter page.

In some cases, the registration process failed to yield perfect alignment because the original images were so severely out of alignment. Where a



FIGURE 1. Page 219. From left: 1929 image, 2010 image, Diff after registration.




FIGURE 2. Pigment chipping on page 218. From left: 1912, 1929, 1962, 2003, 2010.



FIGURE 3. Successive diffs of results for page 234, as landmarks are manually added for correction. Areas with poor alignment appear embossed.

misalignment was deemed unacceptable, it was corrected by manually adding landmark points in the offending region and then re-registering that image pair, such as with page 234 in figure 3. The text initially appears bright where it is not aligned, but as features are added to help correct this, the brightness decreases or disappears.

VISUALIZATION

A layered image viewer was built for the visualization of registered image sets. This tool represents an improvement over previous methods by allowing direct comparison of original images while maintaining their visual context. The viewer consists of a background image and a foreground image. The background image is featured prominently in the interface. A circular “flashlight” then allows a user to “shine” the foreground image onto the background, which is aligned with the foreground image. In figure 4, the “background” image is from 2010 and the “foreground” image seen with the flashlight is from 1962.

The viewer understands a data set as a set of pages, each page being composed of the aforementioned mutually registered set *P*. The data set is described in a JavaScript Object Notation (JSON) file passed to the viewer.

Keyboard shortcuts allow quick switching of the foreground and background images. Other commands enable rotating through the various pages of the manuscript. The flashlight can also be resized and reshaped. These functions allow a user to quickly navigate the pages of the manuscript and pinpoint regions of interest. This tool is an improvement over the aforementioned methods of viewing registered images, which are not well suited for more than two layers.

The viewer uses the Deep Zoom Image (DZI)¹³ tiled image format so that the user can zoom in and achieve high resolution without needing to load the entire full-resolution images in the beginning. OpenSeadragon¹⁴ is used to display the DZIs. The JSON file describing the data set simply points the viewer to the appropriate DZI files and directories, which can be hosted locally or elsewhere.

Figure 5 demonstrates how the viewer can be used in tandem with registered imagery to identify and explore diachronic changes at a fine-grained level. The viewer clearly shows that sometime between 1912 and 1962, a

13 Deep Zoom Image format, developed by Microsoft as part of the Silverlight framework.

14 OpenSeadragon: An open-source, web-based viewer for high-resolution zoomable images, implemented in pure JavaScript, for desktop and mobile, <http://openseadragon.github.io>.

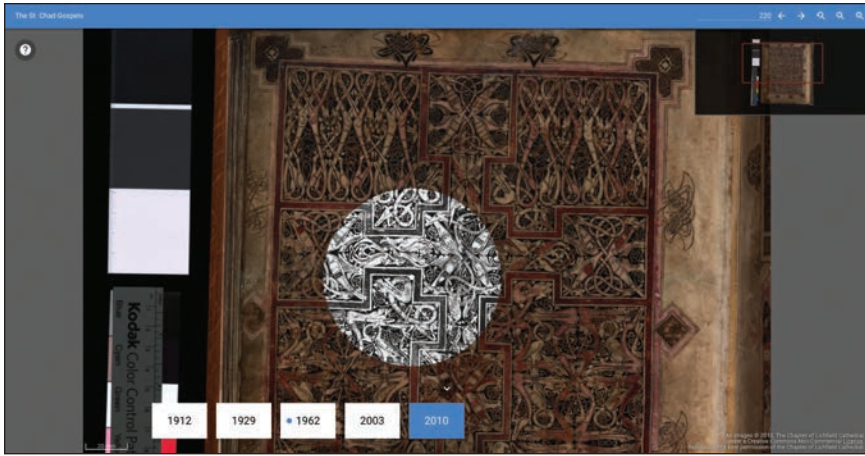


FIGURE 4. Registered image viewer.



FIGURE 5. Changes to the patched region of page 143, as shown in the layered image viewer. From left: 1912 image, 1962 image, 2010 image.

small patch applied to the edge of page 143 was removed and repatched with thread. This repair exactly matches the description of the repair process performed by Powell prior to his 1962 imaging.¹⁵

PRESERVATION

The Internet Archive¹⁶ was chosen as the repository for our diachronically registered images of the St. Chad Gospels. As the name implies, the Internet

¹⁵ Powell, “The Lichfield St. Chad’s Gospels.”

¹⁶ Internet Archive, www.archive.org.

Archive is designed for long-term preservation. It also provides a data ingestion pipeline that processes the images and makes them publicly available for download in various formats. Uploads were performed with a command line tool from archive.org that allows for batch transfer specified by a comma-separated values (CSV) file.

Conclusions and Future Work

High-resolution imaging of manuscript pages is an emerging method for improving document longevity and accessibility. In addition to providing facsimile resources, scans of these pages offer new opportunities for discovery about the original ancient manuscripts. Multimodal and diachronic image sets, for example, can reveal patterns or changes that were not visible to the naked eye working only with the physical manuscript or a digital facsimile.

We have presented a method for the organization, registration, and visualization of diachronic manuscript image sets. Registered images examined with the viewer allow researchers to quickly yet precisely navigate a manuscript, search for changes that have occurred over time, and then pinpoint those warranting a closer inspection. We believe these registration and visualization methods will prove to be powerful tools not only for the St. Chad Gospels, but for other manuscripts and, more generally, for data sets containing images showing some change, whether in time, modality, or other elements.

Several opportunities exist to enhance our methods. Fine-tuning the registration process can completely eliminate problematic areas that are not perfectly aligned. In this work, the alignment was considered successful in the context of comparing images of text over time, but for other applications, the registration may need improvement. Additionally, the feature detection and matching process can be automated, a step that was not necessary for this data but is worth pursuing with a larger or more varied set of images.

As 3D data exists for these pages, digital flattening techniques¹⁷ can also be applied to the 2010 images before using them as reference images in the

17 Michael S. Brown, Mingxuan Sun, Ruigang Yang, Lin Yun, and W. Brent Seales, "Restoring 2D Content from Distorted Documents," *IEEE Transactions on Pattern Analysis and*

registration process. This step can remove any imperfections or warps due to the physical deformation of the manuscript pages under the camera. By subsequently registering sensed images to a flattened reference image, the sensed images will have these deformations corrected as well.

The current implementation of the viewer allows the user to view two images at once, but future versions can have multiple flashlights, allowing any number of layers to be shown in the same view. Finally, the registration process and viewer can be extended to a variety of other datasets beyond the St. Chad Gospels or even manuscripts in general.

This project builds on and anticipates other work that goes beyond the simple advancement of the digital techniques used. Ultimately, the methods we describe lead to the improved preservation and understanding of ancient documents, adding to an old body of knowledge in ways that were not before possible.

Acknowledgments

All images of the St. Chad Gospels copyright The Chapter of Lichfield Cathedral, under a Creative Commons Non-Commercial License. Reproduced by kind permission of the Chapter of Lichfield Cathedral.

University of Kentucky, Center for Visualization and Virtual Environments; Furman University, Department of Classics.

This material is based upon work supported by the National Science Foundation under grant nos. IIS-0535003, 0916148, 0916421, and EAGER-1041949.

Supplemental Materials

The source code used to register the St. Chad Gospels datasets is archived and available on GitHub: <https://github.com/viscenter/registration-toolkit>.

Machine Intelligence 29, no. 11 (2007): 1904–16. Michael S. Brown and W. Brent Seales, “Image Restoration of Arbitrarily Warped Documents,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, no. 10 (2004): 1295–1306.

The datasets presented in this paper are available for download through the Internet Archive: <https://archive.org/details/@viscenter>. Data will also be made available by request.

The registered image viewer is available at <http://infoforest.vis.uky.edu/>. The source code for the viewer is also available on GitHub: <https://github.com/viscenter/layered-viewer>.