

Interactive Visualization of Hyperspectral Images of Historical Documents

Seon Joo Kim, *Member, IEEE*, Shaojie Zhuo, *Student Member, IEEE*, Fanbo Deng, *Student Member, IEEE*, Chi-Wing Fu, *Member, IEEE*, and Michael S. Brown, *Member, IEEE*

Abstract—This paper presents an interactive visualization tool to study and analyze hyperspectral images (HSI) of historical documents. This work is part of a collaborative effort with the Nationaal Archief of the Netherlands (NAN) and Art Innovation, a manufacturer of hyperspectral imaging hardware designed for old and fragile documents. The NAN is actively capturing HSI of historical documents for use in a variety of tasks related to the analysis and management of archival collections, from ink and paper analysis to monitoring the effects of environmental aging. To assist their work, we have developed a comprehensive visualization tool that offers an assortment of visualization and analysis methods, including interactive spectral selection, spectral similarity analysis, time-varying data analysis and visualization, and selective spectral band fusion. This paper describes our visualization software and how it is used to facilitate the tasks needed by our collaborators. Evaluation feedback from our collaborators on how this tool benefits their work is included.

Index Terms—Hyperspectral visualization, data exploration, image fusion, document processing and analysis

1 INTRODUCTION

Archives serve as the custodial record keepers of our nations' collective memories. The role of these institutions is to amass collections of historically significant records and maintain them for future generations. In recent decades, digital imaging has become a mainstay tool used by archives and related institutions such as libraries and museums in their repertoire of management and preservation methods. Processing and analyzing digital images offer a non-invasive approach to study and disseminate historical documents without the risk of damaging the primary source. In addition, digital processing allows the document's contents to be enhanced and analyzed in ways that are not possible with the actual documents themselves.

This work is done in partnership with the Nationaal Archief of the Netherlands (NAN), one of Europe's leading research archives, and Art Innovation, a manufacturer of hyperspectral imaging hardware designed for historical documents. The NAN has recently started capturing HSI of historical documents and exploring the data for a variety of tasks related to document analysis, management, and conservation. In contrast to the standard 3-channel imaging (RGB) which combines response of all visible electro-magnetic (EM) radiation into three bands, HSI captures a dense set of spectral measurements over a much broader spectral range, including the invisible spectral bands in ultra-violet (UV) and near-infrared (NIR) (Fig. 1). This enables the HSI to provide a high-quality quantitative recording of the historical document that is decoupled from the environmental lighting. Such quantitative measurements offer new opportunities for a wide range of tasks in document management and analysis, including, but not limited to, the analysis of document elements (e.g., ink pigments, paper media, biological damage and foxing, etc.), the monitoring of spectral changes due to environmental aging, and the ability to emphasize faded or hard to see document content through spectral information.

This project is led by NAN conservator, Roberto Padoan, and senior engineer, Marvin Klein, from Art Innovation. Our role in this collaboration is the development of software to visualize and manipulate the

- Seon Joo Kim, Shaojie Zhuo, Fanbo Deng, and Michael S. Brown are with the National University of Singapore, E-mails: kimsj.zhuoshao,dfanbo,brown@comp.nus.edu.sg
- Chi-Wing Fu is with Nanyang Technological University, E-mail: cwf@ntu.edu.sg

Manuscript received 31 March 2010; accepted 1 August 2010; posted online 24 October 2010; mailed on 16 October 2010.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

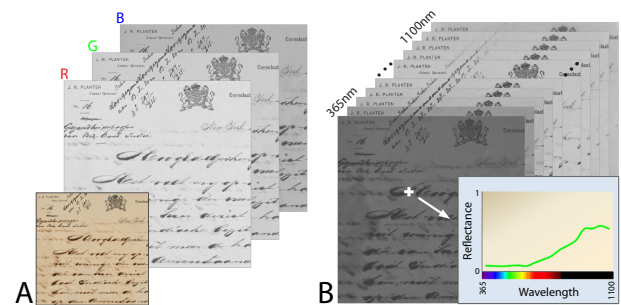


Fig. 1. (A) Standard three channel RGB image, (B) 70+ band hyperspectral image.

HSI data. The capture and analysis of the highly calibrated HSI data offered by Art Innovation's hardware represents a pioneering effort in the archival community. While there is an open agenda to the potential applications of this device in the archival setting, several key tasks for using the HSI data have been established by our collaborators.

The first task involves using the HSI data to analyze deterioration artifacts commonly found in archival materials, such as biological and physical damages, metal gall inks corrosion, paper oxidation (foxing), and pigment fading. This task is important to help decide on an adequate conservation treatment to avoid further damages [25]. While the variations in environmental lighting, humidity, and temperature are known to induce damage to the document over the long term, the exact process of document aging is still unknown. Understanding how HSI data is correlated to the aging process will play a significant role in helping to manage valuable historical documents by measuring subtle changes that are difficult, if not impossible, to distinguish by the unaided eye. To this end, the NAN has devised an elaborate controlled experiment to induce artificial aging on a wide assortment of documents of various materials and inks. The goal is to systematically monitor and quantify changes in the HSI of these documents to track the aging process.

Another important task of our collaborators is the enhancement of deteriorated documents for improved legibility and interpretation of the materials. In particular, document content often appears more salient at different spectral bands. The goal in this task is to use the varying contextual information at different wavelength to selectively enhance the desired content in an interactive fashion.

To assist our collaborators, we have designed and implemented a comprehensive visualization tool that offers an assortment of visualization and analysis features including:

- Basic features for exploring and visualizing the HSI data to identify and monitor damages or specific region-of-interest (ROI) on archival materials. The features include interactive wavelength selection, interactive region-of-interest exploration, data similarity analysis (with user selected bands), and RGB image rendering with tunable illumination;
- Analysis and visualization tools of time-varying reflectance change for monitoring and quantifying the effects of external aging factors including lighting, and changes in humidity and temperature;
- Interactive visualization of HSI data based on gradient domain fusion for enhancing the legibility of HSI data based on dynamic selection of spectral bands. There are two major contributions within this feature. First, we apply illustrative visualization to the fusion of HSI data to enhance the contextual details in the data. Second, the fusion technique makes use of the focus+context visualization by enabling the user to explore the fused result with focus on a local region while showing the regions outside with original appearance.

The remainder of the paper is organized as follows: we begin by reviewing related work in Section 2 and describe the imaging process in Section 3. Basic visualization features are explained in Section 4. Visualizing changes in the documents due to aging factors is described in Section 5, and HSI enhancement and visualization based on gradient domain fusion is presented in Section 6. We conclude with a discussion and summary that includes evaluation from our collaborators in Sections 7 and 8.

2 RELATED WORK

Visualization is highly effective in enhancing data exploration and data understanding in a wide range of applications, ranging from biology and medicine to climate research. For example, Kruger et al. [19] applied GPU volume rendering for virtual sinus endoscopy, while Miller et al. [24] employed interactive visualization to study the structure of intercluster galaxies with a collection of visual and analytical tools. Johnson et al. [14] studied the data management and visualization aspects of large transitional flow on desktop PCs and Kehrer et al. [15] demonstrated how interactive visual exploration helps scientists in climate research to evaluate their hypothesis. In this work, we deal with practical problems faced by conservators in libraries, museums, and archives, with the goal of helping them analyze and manage historical documents by deploying related analysis and visualization techniques.

Document Processing.

Hyperspectral imaging is a relatively new procedure in libraries and archives. Most existing work in document processing deal with standard RGB images with the focus on correcting the artifacts such as ink-bleed and corrosion. Without additional benefits of spectral bands, artifacts are detected by adaptive thresholding [27], source separation [29], or user-assisted training [22]. The goal of these approaches is the detection of the artifact with little emphasis placed on the final output of the image, i.e., there is no mechanism to enhance the visualization of the data to preserve the look and feel of the original document.

There does exist prior work addressing historical documents and paintings captured using multispectral or hyperspectral imaging. Historical books and archival documents were imaged by a multispectral imaging system in [23] and a hyperspectral imaging system in [25] which is our collaborator's early work. In [8], a 13-channel spectral imaging system was used to capture Leonardo Da Vinci's Mona Lisa. These prior works only discussed that enhancement or restoration could be achieved using multispectral or hyperspectral imaging, however, the realization of such enhancement was ad hoc in nature. A comprehensive framework has yet to be demonstrated.

Our work presents one of the first fully functional software for analyzing and manipulating hyperspectral images of old documents. In contrast to the methods mentioned above, our approach is user-centric with tools that allow the users to navigate the hyperspectral data and perform analytical analysis using various visualization techniques. We can also utilize the additional data in the invisible wavelength to enrich the image with details or reduce artifacts in the image while keeping it visually as close to the original RGB image as possible.

Hyperspectral/Multispectral Data Visualization.

Hyperspectral and multispectral imaging¹ has been applied to data ranging from those in very large scale to objects at atomic scales. Examples include spectral imaging of astrophysical data [12, 21], earth science data [9, 13, 10], and biological and medical data [30, 31]. Among these works, most relevant works to this paper are those that focus on the visualization of hyperspectral (multispectral) geo-science data (remote sensing). In remote sensing, hyperspectral images at different spatial, temporal, and spectral resolution are fused to increase the visual saliency of the area being imaged. To merge multispectral satellite images, Choi et al. [7] proposed a curvelet-based image fusion method which represents edges better than wavelets resulting in richer information in both the spatial and spectral domains. Socolinsky and Wolff [28] presented an optimal grayscale visualization method by taking multispectral contrast into account to increase the quality of image fusion. Cui et al. [9] introduced an interactive scheme for visualizing hyperspectral images using convex optimization. In [10, 13], design goals and solutions for the display of fused hyperspectral aerial images on tristimulus displays were presented. There are also softwares available for visualizing hyperspectral data for remote sensing, e.g., [1, 2].

The focus of visualization in remote sensing is to increase the visual saliency and assist with the interpretation of the data, in many cases simply by false-coloring. In contrast, we focus on visualization that suits the HSI of historical documents and related tasks. Specifically, we present an image fusion technique for enhancing the contextual details with an emphasis on maintaining the natural look and the feel of the document. In addition, our fusion technique enables the focus+context visualization through user-interaction by introducing a special lens feature.

3 HYPERSPECTRAL IMAGING : DATA ACQUISITION

Before going into the details of our visualization software, we first describe the imaging process used to acquire the HSI data. As previously mentioned, this work is part of a collaboration with the NAN where documents are imaged using the SEPIA Quantitative Hyper-Spectral Imager (QHSI) device developed by Art Innovation [16]. The device performs hyperspectral imaging by capturing a very narrow spectral band of EM radiation one at a time by placing a bandpass filter in front of the light source to block out all but a selected band of the EM spectrum. A monochromatic camera is then used to capture the amount of light that is reflected by the document at that selected band. The filter is changed for each image, thus capturing different parts of EM spectrum to build up the HSI. Fig. 2 (A) illustrates this imaging process. A picture of the device is shown in Fig. 2 (B).

The QHSI device captures images at different wavelength bands from 365 nm (UV) to 1100 nm (near infrared (NIR)) with the step size of 10nm in most cases except the bands in 300nm's and 1000nm's. The images have the resolution of 4 mega pixels (2048×2048) for a physical surface area of 125 mm × 125 mm and are captured at 16 bit per pixel. Such high-resolution (approximately 256 pixels per mm²) provides a reliable spatial measurement suitable for even thin lines of handwriting and printed text. Before imaging archival materials, the QHSI device is calibrated via a spectral pattern with known spectral response.

¹The difference between multispectral vs. hyperspectral imaging is related to the number of bands as well as the manner in which the data is collected. In this paper, we refer to our data as hyperspectral as it provides a dense and highly calibrated spectral response from a single sensor.

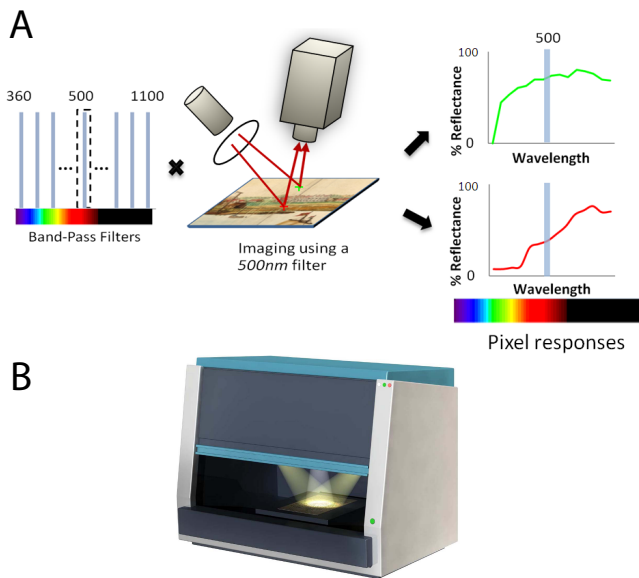


Fig. 2. (A) The hyperspectral imaging process: at each scan, a monochrome camera measures the reflected light from the document surface. The document reflects a very narrow band of EM radiation due to the bandpass filter positioned in front of the light source (500nm in this example). This process is repeated using multiple bandpass filters (70 in the paper) to build the HSI. (B) SEPIA quantitative hyperspectral imager from Art Innovation [16].

The following notations will be used to represent the acquired HSI data. The term I_λ indicates the image of the data at a spectrum λ , \mathbf{x} indicates a pixel location, and $\mathbf{s}_\mathbf{x}$ indicates the spectral response (spectrum) at point \mathbf{x} :

$$\mathbf{s}_\mathbf{x} = [I_{\lambda_1}(\mathbf{x}), I_{\lambda_2}(\mathbf{x}), \dots, I_{\lambda_m}(\mathbf{x})]^T. \quad (1)$$

4 BASIC FEATURES FOR HSI VISUALIZATION

This section describes the elementary features offered by our visualization tool. Using the basic functions of the software, the user can load an HSI datacube and interactively navigate through the spectral bands using a slider widget. Each spectral image is displayed as a monochromatic image representing the percentage of spectral reflectance at each pixel for this band. Fig. 1 shows an example of images at different bands.

Note, that the HSI data does not explicitly contain an RGB representation of the imaged document, however, it does have the complete spectral response of each pixel for the visible range. Thus we are able to provide a function that can render the RGB images of the HSI data under different illuminations (e.g., daylight at different color temperatures, white illumination, fluorescent lighting, or user-specified color temperature). This is a useful feature as the RGB image represents the most natural way to display the document as it would appear to our unaided eyes under the various illumination (see accompanying video). While pure white light was useful for the majority of visualization, the ability to change the type of lighting was welcome by our collaborators because it provides the ability to see how the document appears under different exhibition lighting environments.

Another basic operation necessary to analyze the hyperspectral data is to compare spectrums of different image points to see how similar (or dissimilar) the spectrums are. This operation is useful for tasks such as separating the foreground contents from the background of the document and measuring the corrosion and ink-bleed severity in the document. Our software provides this similarity computation feature by user mark-up interaction. The user marks on a region or regions of interest, then the similarity measure (S) between every pixel in the image and the mean spectrum of the marked-up area ($\bar{\mathbf{p}}$) is computed

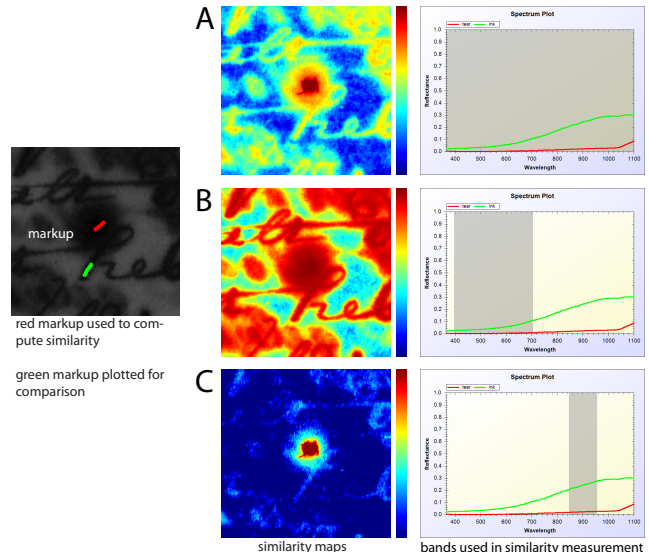


Fig. 3. Similarity computation. Similarity measure (Eq. 2) between the mean spectrum of the marked area (red) and other points in the data are computed by using the entire spectral bands (A), visible bands (B), and selected bands in NIR (C). Different bands can be selected for different tasks providing more flexibility in similarity computation.

as follows :

$$S(\mathbf{x}) = 1 - \frac{\|\bar{\mathbf{p}} - \mathbf{s}_\mathbf{x}\|^2}{m}, \quad (2)$$

where m is the number of spectral samples. An example of the similarity computation is shown in Fig. 3(A).

Our colleagues also wanted to perform this similar metric using selective bands, and not always the entire HSI spectrum. This allows the evaluation of results obtained by using only a few bands for similarity analysis, as prior research in the archival domain has established that certain bands are more suitable for various tasks and materials being observed (e.g., animal skin parchment versus paper-based substrates). Such experimentation is also useful for managing future data collection in which only the useful bands may need to be captured. Fig. 3(B-C) shows an example where various band combinations are being used, in particular the visible range (B) and a few NIR bands (C).

Another basic feature included in our software enables the user to focus on a certain region of interest on the document and visualize how the content in the selected region changes over different wavelength. To provide this context-focus feature, a lens interface termed the hyperspectral lens (inspired by [4]) was designed to view either a selected rectangle, circle, or free-form region, while maintaining the current spectral context. Fig. 4 illustrates an example of this feature. This basic visualization metaphor is also available for use with the band fusion feature that will be discussed in Section 6.

5 MONITORING CHANGES IN TIME VARYING HSI DATA

As discussed in Section 1, one of the tasks targeted by the NAN is to utilize hyperspectral imaging to detect changes in documents due to environmental factors affecting archival documents, namely, exposure to light, changes in humidity and temperature, and combination of both. While the effects of these factors are known to induce damage to the document over the long term, a systematic way to monitor and visualize changes to track the exact process of aging is not currently available in the archival community. This is particularly important for highly valued documents which are subject to more exposure due to exhibition. Currently, the majority of archives monitor the conditions of documents by visual inspections, a method that is highly unreliable and often catches changes only after a major deterioration has occurred, such as age-related spots.




Fig. 4. Context-focus through hyperspectral lens. The user can explore the HSI data with the focus on a local region of interest by using the hyperspectral lens feature in our software.




Fig. 5. Artificial aging applied periodically to a document. Each document is divided into four strips with the bottom strip serving as a control with no aging induced. The remaining strips from the bottom are artificially aged using 1) lighting; 2) temperature and humidity change; and 3) both lighting and temperature/humidity change.

The major drawback to this effort is that the aging of documents is a very slow process and could potentially take decades under natural conditions to collect data with meaningful changes. To expedite data collection for testing, the NAN has devised an elaborate and on-going controlled experiment to induce artificial aging on a wide assortment of documents of various materials and inks. In this experiment, a series of documents have been cut into four sections (or strips), resulting in four pieces from the same source. One of these pieces is used as the “control” and placed in standard storage conditions, i.e., boxed with no light exposure and with constant temperature and humidity. The remaining strips are periodically artificially aged using 1) lighting; 2) temperature and humidity change; and 3) both lighting and temperature/humidity change. Specifically, artificial aging for lighting replicates the spectral characteristics equivalent to four months of light exposure in a matter of days. Temperature and humidity change are performed following the procedures practiced in the conservation field [5]. After each aging cycle, the HSI datacube of the document is captured and then the material is placed back into storage to settle. Fig. 5 shows an exemplar document from this experiment after four applications of this cumulative artificial aging. The visible change is difficult to observe by unaided eyes.

Our software offers several ways to analyze the effects of environmental aging as recorded in the HSI data. One of the most obvious ways is to compute the spectral difference between the same strip as it is aged (method 1). The strips are geometrically aligned and their per pixel spectral difference is computed based on Eq. (2). An advantage of this type of visualization is that it shows the effect of the aging spatially. Another way to visualize aging is to show the changes due to different environmental factors (method 2). The intent is to see the changes due to different methods of aging rather than over time. The mean spectrums of the background substrate (which is interactively selected by the user) of different strips are compared for this visualization. Finally, we provide histograms plot features to help visualize the aging process (method 3). By jointly displaying histograms of the document strips with increasing amounts of cumulative aging, we can elucidate the overall trend of how the HSI data changes due to different environmental stress. Since HSI is a 3D data (cube), the histogram of a HSI data is 3D. Our tool visualizes the histograms of HSI in two ways. First, the software provides a slider widget that allows the user to see




Fig. 6. (A) Two strips artificially aged at the first aging exposure and then after the fourth cumulative exposure based on changes in the temperature/humidity. (B) Spectral plots of the background paper (top) and the ink (bottom) in the strips. The ink shows more spectral change than the paper substrate. (C) Spectral difference map shows that the response of the paper substrate to the aging process is spatially varying.

the 2D histogram at each spectral band. We also provide a 3D spectral histogram with additional axis for the wavelength to help reveal the overall trend between two aged strips at the same time.

Case Examples

Two datasets are shown as examples of how our collaborators can use the technique described in this section to analyze and quantify changes in the documents due to external aging factors. The first data is a handwritten document circa 1800s and the second example is a blank paper that has been constructed based on a historical paper making method from the same era (Fig. 7).

Fig. 6 shows an example where the per pixel spectral difference between the same strip of HSI data is compared at the first aging exposure and then after the fourth cumulative exposure (method 1). The difference map in Fig. 6 reveals non-uniform response of the paper substrate to the aging process. In this example, we have also shown the spectral plots of ink-regions and background regions (i.e., the substrate). The ink’s spectral response has changed more than the paper, which is reflected in the difference map. The spectral change for the ink is also more distributed over the entire spectrum, while the paper is affected in wavelengths below 700nm.

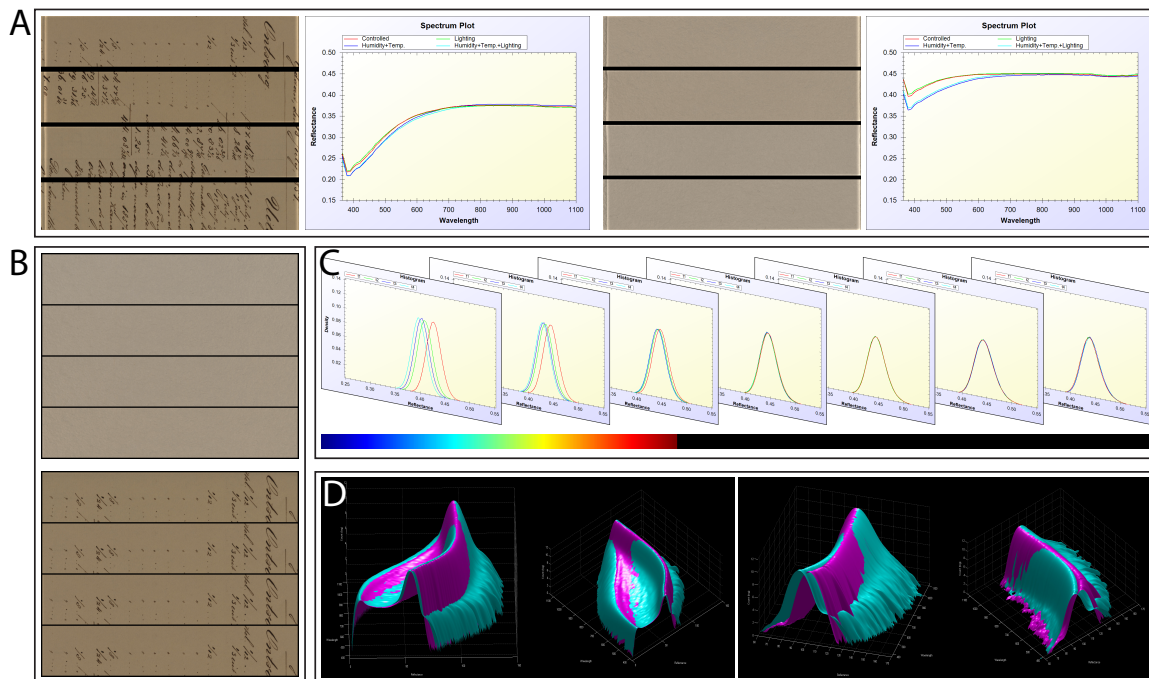


Fig. 7. (A) Effects of different aging methods. As can be seen, changes in humidity/temperature has the biggest effect on aging. (B) Strips are aged cumulatively more in the vertical direction. (C) 2D histograms for each aging step are jointly displayed at each spectral band. (D) 3D spectral histograms of two time instances reveal the overall trend between two aged strips at the same time.

Fig. 7(A) shows the mean spectrums of the background which indicate the difference in aging due to different aging factors (method 2). The plots for both cases reveal that the change in humidity/temperature affects the document aging the most. Notice that computing the difference map as in Fig. 6 is not a viable option since the content of each strip is different. Visualization of aging through histogram (method 3) is illustrated in Fig. 7(C). Four histograms at each spectral band in Fig. 7(C) represent the histograms of the data as it is progressively aged (Fig. 7(B)). We also provide a 3D spectral histogram with additional axis for the wavelength to help reveal the overall trend between two aged strips at the same time (Fig. 7(D)). This again reveals that the spectral response to the aging is most notable in bands below 700nm. Furthermore, it is possible to verify that we are actually tracking aging process rather than errors from the imaging system through visualization schemes shown in this section. This is important because aging is a slow process and the effect of aging in these experiments are very small.

6 VISUALIZATION THROUGH FUSION

Although displaying individual image per wavelength provides a direct view on the raw HSI data, contextual details in the entire data volume are rarely distributed evenly through different spectral ranges. Our NAN collaborator expressed the desire to visualize the HSI data in its natural appearance, but at the same time emphasizing various contents in the invisible range to enhance the legibility of the data. A natural way to do this is an illustrative visualization approach where the aim is to maximize the amount of contextual details while maintaining the appearance of the document in the visible range. In addition, our collaborator preferred to interactively select the bands that provide contextual details.

As discussed in Section 4, an RGB image of the HSI data under a user-specified lighting condition can be generated and displayed. However, such an image cannot represent the entire HSI data since it integrates only the visible spectrums (400nm to 700nm) and excludes potentially valuable information in the invisible range (UV, NIR).

A new scheme for visualizing the information in the invisible bands while keeping user’s perceptual context with views on the visible

portion (RGB) is necessary. For such visualization, we introduce a gradient-based fusion technique which has been shown to be effective for computer vision and graphics tasks such as image editing [26], contrast adjustment [11], image stitching [20], and image fusion [17]. Specifically, we want to reconstruct a new image $I'_i(\mathbf{x})$ by fusing the contextually rich gradients of the image at a wavelength λ ($F_\lambda(\mathbf{x}) = \nabla I_\lambda(\mathbf{x})$) with the color of the original RGB ($I_i(\mathbf{x})$), where index i indicates color channel, \mathbf{x} represents pixels, and ∇ is the gradient operator. To do this, we solve for each color channel of the new image I' by minimizing the following cost function :

$$\operatorname{argmin}_{I'_i} \sum_{\mathbf{x}} \gamma (I'_i(\mathbf{x}) - I_i(\mathbf{x}))^2 + |\nabla I'_i(\mathbf{x}) - \beta F_\lambda(\mathbf{x})|^\alpha. \quad (3)$$

The first term in above equation makes the color values of the reconstructed image I' to be close to the color of the original image I under an ℓ^2 norm. The second term forces the gradient of the reconstructed image to be close to the gradient of an image at a wavelength λ under a sparse norm ($\alpha \leq 1$). Note that it is the gradient structures that contain the relevant information, not the gradient magnitudes. Hence, the reason for using the sparse norm is because the sparse norm on the gradient terms encourages the edge structures in I' to align spatially with those in F_λ compared to the ℓ^2 norm where $\nabla I'$ will be matched closely with F_λ in magnitudes [17]. The parameter γ controls the balance between the color term and the gradient term. The parameter β scales the strength of the gradient. To optimize Eq. 3, we modified a fast optimization scheme used for image deconvolution recently introduced in [18].

This gradient fusion method for transferring information from the invisible range to the RGB image is incorporated into the hyperspectral lens feature (Section 4). Not all areas have to be visualized using the new fusion technique since the invisible bands do not contain any additional information in many regions. To this end, we add additional feature to the existing user-interface so that the user can efficiently select the area for fusion and the band to transfer the data from by looking at the fused result. The user selects a region-of-interest within the image with the hyperspectral lens feature. When the user scrolls through the wavelength, the hyperspectral lens can now display the

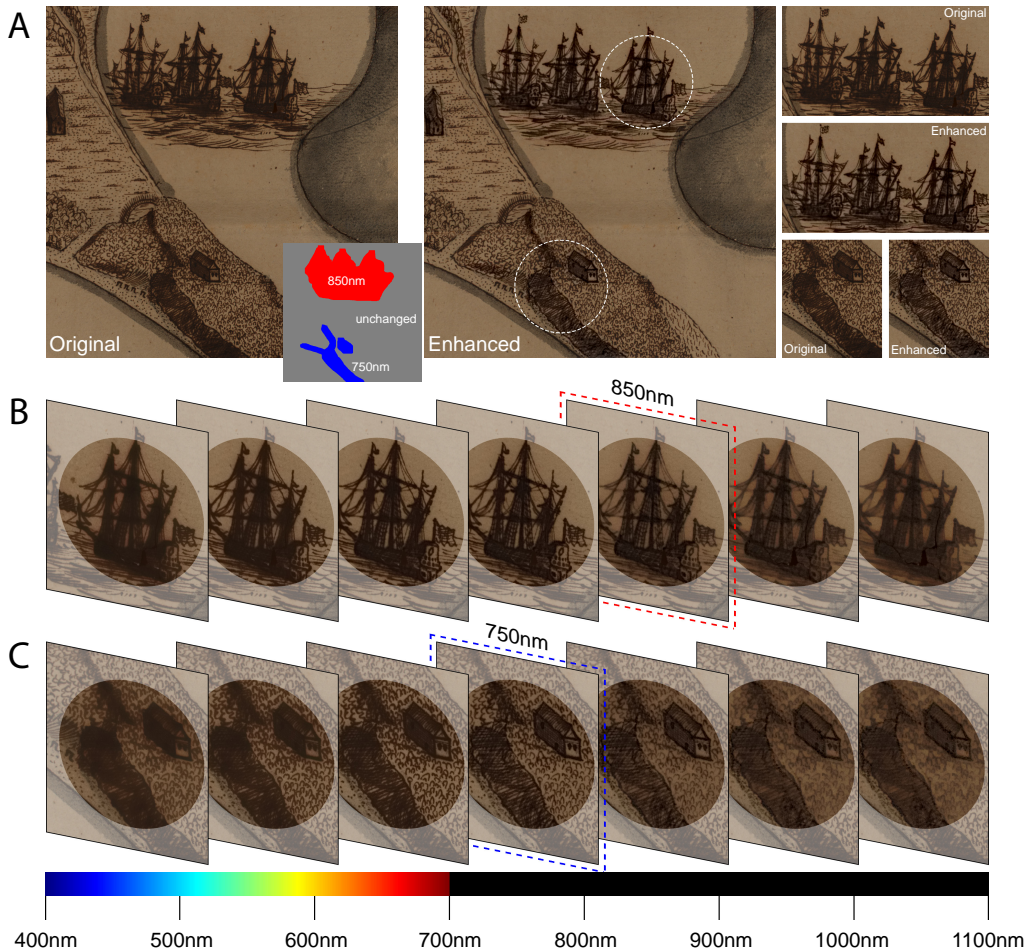


Fig. 8. (A) Contrast of some parts in the original RGB image (left) is enhanced with our gradient fusion method (middle). Side by side comparison of enhanced regions are shown on right. (B)(C) Enhancement using each band is shown locally through the hyperspectral lens. User can interactively choose the band that suits the application (highlighted with dashed line), where in this case is the contrast enhancement.

newly fused color image using the gradient of the image at the current wavelength instead of showing the gray-scale image at each band. This enables the user to browse through the HSI data for analysis with color perception.

We can also use the fusion technique to visualize the HSI data with an artifact-free RGB image in which the artifacts on the document such as ink-bleed, ink corrosion, and foxing (age spots) are removed. For this task, the user can use the similarity computation feature introduced in Section 4. After the user marks-up on a foreground text area, the HSI data can be segmented into the foreground and the background by using the computed similarity scores. A new gradient map (F_i) is constructed by combining the gradients of the original RGB (∇I_i) as the foreground and the gradients of a band (∇I_λ) chosen by the user with less artifacts as the background (Fig. 10 (B)) :

$$F_i = M \cdot \nabla I_i + (1 - M) \cdot \nabla I_\lambda, \quad (4)$$

where M is a binary mask with 1's in the foreground and 0's in the background. With the new gradient map F_i , a new RGB image I' is reconstructed by optimizing the cost function in Eq. 3.

Case Examples

Fig. 8 shows part of a late 17th century map that belonged to Michiel de Ruyter, a famous Dutch admiral. The original RGB image in Fig. 8 contains several regions with low contrast and it is difficult to clearly see the textures within those regions. There is even a tear inside one of the ships that is difficult to recognize in the original RGB image. With our fusion tool, our NAN collaborator enhanced the local areas with low contrast by selectively choosing spectral bands that have more

salient content. The salient content in this example represents crisper edges. In this case, the spectral band at 850nm is selected to enhance the boat region, while 750nm is selected for the lower part of the map. A new gradient map is composed from these selections (small color map in Fig. 8 (A)), and the enhanced image is reconstructed by optimizing Eq. 3. Another case is shown in Fig. 9 where our tool was used to reveal a tear in a document suffering from significant ink-corrosion while maintaining the original look of the document. For both cases, the use of the hyperspectral lens is used to provide a focus+context visualization of the desired regions.

In addition to enhancing information, the fusion technique can be used to reduce artifacts. Fig. 10 shows an example where the age spots due to foxing can be reduced while the look and the feel of the document are preserved. The ability to preserve the look and the feel of the original RGB image holds a significant advantage over most existing document processing techniques in which the output of the artifact removal is a binary image with a uniform color background such as the mean of the input background color or simply white. While the binarization enhances the ability to interpret the data, the texture and the look of the original document is completely lost (Fig. 10 (D)).

7 DISCUSSION

Before our involvement, our NAN collaborator only had a set of menu-driven Matlab routines at his disposal to process the data. Analysis was performed in a serial fashion with results being saved to images. Not surprisingly, the interactive nature of our software was welcomed. Our NAN collaborator noted that our interactive software allowed the

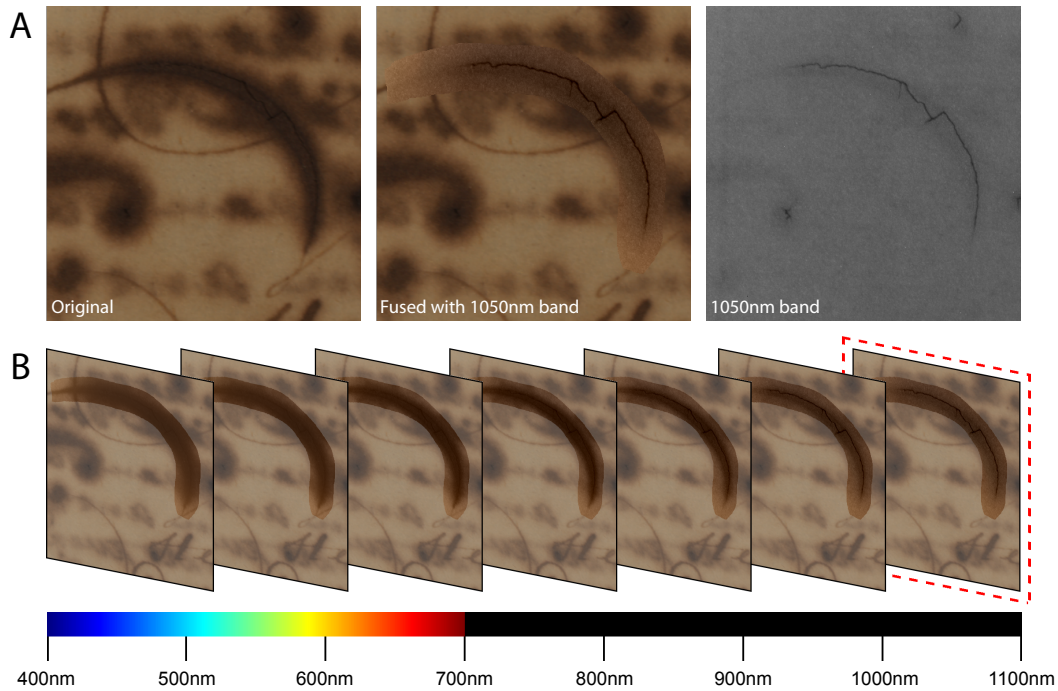


Fig. 9. (A) The image at 1050nm (right) reveals a tear in the document that is extremely difficult to catch in the original image (left). Using our fusion technique for visualization, we can show the condition of the document while keeping the color in the image (middle). (B) Fusion results using different bands with our user interface.

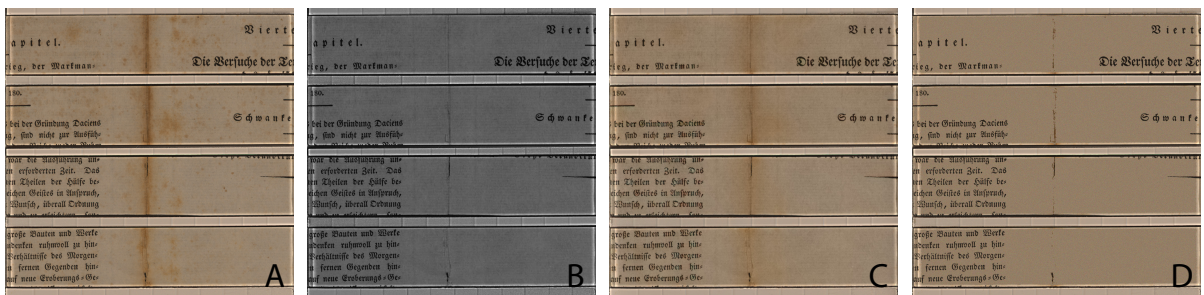


Fig. 10. (A) Original RGB image of the HSI data. (B) Source for providing gradients for the background (900nm). (C) New RGB image with artifacts removed using our technique. (D) RGB image with the background filled with the mean value of the background. The resulting image using our technique maintains the look and the feel of the original in contrast to the result in (D).

basic tasks of computing similarity measurements and producing false color images to be performed in a matter of minutes, while the previous menu-driven software could take up to hours to produce the desired output. It was from this early uses of our tool that our NAN collaborator suggested that a region of interest tool (hyperspectral lens) be introduced to enable focus+context visualization concept to allow focus on a particular region of the document while the rest of the image remained in natural appearance.

Towards the effort in using the HSI data to monitor the aging of documents under environmental exposures, our NAN collaborator pointed out that computing the spatial variation of aging on the whole document was a first of its kind in the archival community, where it was previously limited to comparing only a small number of spots on the document using single-point spectrum measurements. Our visualization of spatial variations of aging helps to reveal non-uniform aging that are common to papers, which is due to non-uniform light exposures as well as the paper's own physical construction. Multi-dimensional histogram plots which help to reveal the trends in aging in different materials due to different factors was also welcomed by our collaborators. The ability of our software to provide instant feedback with different types of histograms provides more in-depth analysis of the data

to our collaborators. The 3D histogram was also useful as it helped to overview the global trend in the spectrums. We note that measuring the effects of environmental aging is part of an on-going research at the NAN (and also at other archives). Currently, even with the assistance of our visualization tool, no quantitative assignment of damage can be applied to the documents. To be able to do this will require a long-term experiment where artificially aged documents begin to exhibit strong artifacts such as foxing. This is an evolving experiment of which our visualization tool will play a vital role.

While selective band fusion has been demonstrated in other domains, within the archival community it is a first of its kind. Our collaborators informed us that they had previously attempted to produce examples similar to ours, but did so manually using Adobe Photoshop to composite contrast enhanced monochrome images by stitching different parts of the spectral images together, which proved unsatisfactory. Our approach allows this to be done interactively and within the context of an RGB image. Our NAN collaborator explained that there are two main uses of this feature in his domain. First, the fusion procedure can be used to produce an enhanced version of the object that can be displayed next to the original in the exhibition case. Producing a result that is perceptually similar helps to maintain the look and

feel of the document, while revealing salient information. Second, this feature is useful to the conservator for producing reports to share with colleagues to help reveal regions that require repair, as demonstrated in Fig. 8 and Fig. 9. By having the damaged region shown in the context of the whole document, the conservator in charge of physical repair can easily identify the corresponding region on the physical document. Our collaborators (both from NAN and Art Innovation) also suggested to maintain the control over which bands to use for fusion, even though we indicated that this could be potentially automated based on some objective criteria (such as local contrast or local gradient). The reason given is that the use of this particular feature can vary greatly between tasks and documents, and it is difficult to predict exactly what they may want to highlight.

Finally, one theme that recurred in the discussions with our colleagues was that while the NAN is fortunate enough to be able to perform research on the HSI, the hardware cost is currently prohibitive for most archives. Instead, highly controlled RGB imaging together with a set of filters and illumination for selective invisible spectrums may be a preferred option for many tasks. Thus, the ability of our software to explore the data and perform analysis in band-selective manner is useful in providing the insight for the construction of future multispectral or single band devices, custom tailored for particular tasks.

8 CONCLUSION AND FUTURE WORK

We have presented a visualization tool tailored for use on hyperspectral images of historical documents. Our tool was developed to support the tasks of our collaborators in NAN and we have discussed how various visualization and analysis features inside can assist our collaborators on the analysis and management of historical documents. For future work, we will continue collaborating with the NAN to develop suitable ways of further revealing subtle changes due to environmental aging and to quantify this damage. This will include exploring better spectral similarity measure (e.g., [3, 6]) and higher level data analysis such as principal component analysis (PCA). The ultimate goal is to be able to establish models of document aging to predict how a document will appear after an exhibition or within a certain time span. We envision that this type of models would also lead to developing algorithms for restoring damaged documents.

ACKNOWLEDGMENTS

We gratefully acknowledge the support and efforts from our collaborators Roberto Padoan from the Nationaal Archief of the Netherlands (NAN) and Marvin Klein from Art Innovation. This work was supported in part by the NUS Young Investigator Award, R-252-000-379-101, and the A*Star SERC Grant (SERC Grant No. 092 101 0063).

REFERENCES

- [1] ENVI. <http://www.itervis.com>.
- [2] PCI geomatica. <http://www.pcigeomatics.com>.
- [3] J. C. Anderson, L. J. Gosink, M. A. Duchaineau, and K. I. Joy. Interactive visualization of function fields by range-space segmentation. *Computer Graphics Forum (Proc. of EuroVis)*, 28(3):727–734, 2009.
- [4] E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose. Toolglass and magic lenses: The see-through interface. In *Proc. of SIGGRAPH*, pages 445–446, 1993.
- [5] J. Bogaard and P. M. Whitmore. Explorations of the role of humidity fluctuations in the deterioration of paper. In *Workshop of Art on Paper Books, Documents and Photographs*, pages 11–15, 2002.
- [6] C.-I. Chang. An information-theoretic approach to spectral variability, similarity, and discrimination for hyperspectral image analysis. *IEEE Trans. on Information Theory*, 46(5):1927–1932, 2000.
- [7] M. Choi, R. Y. Kim, M.-R. Nam, and H. O. Kim. Fusion of multispectral and panchromatic satellite images using the curvelet transform. *IEEE Geoscience and Remote Sensing Letters*, 2(2):136–140, Apr. 2005.
- [8] P. Cotte and D. Dupraz. Spectral imaging of Leonardo Da Vinci’s Mona Lisa: An authentic smile at 1523 dpi with additional infrared data. In *IS&T Archiving Conf.*, pages 228–235, 2006.
- [9] M. Cui, A. Razdan, J. Hu, and P. Wonka. Interactive hyperspectral image visualization using convex optimization. *IEEE Trans. on Geoscience and Remote Sensing*, 47(6):1673–1684, 2009.
- [10] Q. Du, N. Raksuntorn, C. Shanshu, and R. J. Moorhead. Color display for hyperspectral imagery. *IEEE Trans. on Geoscience and Remote Sensing*, 46(6):1858–1866, 2008.
- [11] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. *ACM Trans. on Graphics*, 21(3):249–256, 2002.
- [12] R. Gooch. Astronomers and their shady algorithms. In *Proc. of IEEE Visualization Conf. 1995*, pages 374–377. IEEE, Oct/No 1995.
- [13] N. Jacobson and M. Gupta. Design goals and solutions for display of hyperspectral images. *IEEE Trans. on Geoscience and Remote Sensing*, 43(11):2684–2692, Nov. 2005.
- [14] G. Johnson, V. Calo, and K. Gaither. Interactive visualization and analysis of transitional flow. *IEEE Trans. on Visualization and Computer Graphics*, 14(6):1420–1427, Nov.-Dec. 2008.
- [15] J. Kehler, F. Ladstadter, P. Muigg, H. Doleisch, A. Steiner, and H. Hauser. Hypothesis generation in climate research with interactive visual data exploration. *IEEE Trans. on Visualization and Computer Graphics*, 14(6):1579–1586, Nov.-Dec. 2008.
- [16] M. Klein, B. J. Aalderink, R. Padoan, G. de Bruin, and T. A. G. Steemers. Quantitative hyperspectral reflectance imaging. *Sensors*, 8(9):5576–5618, 2008.
- [17] D. Krishnan and R. Fergus. Dark flash photography. *ACM Transactions on Graphics (Proc. of SIGGRAPH)*, 28(3):1–11, 2009.
- [18] D. Krishnan and R. Fergus. Fast image deconvolution using hyper-laplacian priors. In *Proc. of Neural Information Processing Systems*, pages 1033–1041, 2009.
- [19] A. Kruger, C. Kubisch, G. Strauss, and B. Preim. Sinus endoscopy - application of advanced GPU volume rendering for virtual endoscopy. *IEEE Trans. on Visualization and Computer Graphics*, 14(6):1491–1498, Nov.-Dec. 2008.
- [20] A. Levin, A. Zomet, S. Peleg, and Y. Weiss. Seamless image stitching in the gradient domain. In *Proc. of European Conf. on Computer Vision*, pages 377–389, 2004.
- [21] H. Li, C.-W. Fu, and A. J. Hanson. Visualizing multiwavelength astrophysical data. *IEEE Trans. on Visualization and Computer Graphics*, 14(6):1555–1562, Nov./Dec. 2008.
- [22] Z. Lu, Z. Wu, and M. S. Brown. Directed assistance for ink-bleed reduction in old documents. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 88–95, 2009.
- [23] G. Maino. Digitization and multispectral analysis of historical books and archival documents: Two exemplary cases. In *Proc. of the International Conf. of Image Analysis and Processing*, pages 119–124, 2007.
- [24] J. Miller, C. W. Quammen, and M. C. Fleenor. Interactive visualization of intercluster galaxy structures in the horologium-reticulum supercluster. *IEEE Trans. on Visualization and Computer Graphics*, 12(5):1149–1156, 2006.
- [25] R. Padoan, T. Steemers, M. Klein, B. Aalderink, and G. de Bruin. Quantitative hyperspectral imaging of historical documents : Technique and applications. In *Proc. International Conf. on NDT of Art*, pages 445–446, 2008.
- [26] P. Pérez, M. Gangnet, and A. Blake. Poisson image editing. *ACM Trans. on Graphics (Proc. of SIGGRAPH)*, 22(3):313–318, 2003.
- [27] Z. Shi and V. Govindaraju. Historical document image enhancement using background light intensity normalization. *Proc. IEEE International Conf. on Pattern Recognition*, pages 473–476, 2004.
- [28] D. A. Socolinsky and L. B. Wolff. Multispectral image visualization through first-order fusion. *IEEE Trans. on Image Processing*, 11(8):923–931, Aug. 2002.
- [29] C. L. Tan, R. Cao, and P. Shen. Restoration of archival documents using a wavelet technique. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(10):1399–1404, 2002.
- [30] J. Zhao, Q. Pang, J. Ma, X. Zheng, and Q. Meng. Multispectral imaging system applied to element testing of biology. In *Proc. of the 2008 international Conf. on Biomedical Engineering and informatics*, volume 02, pages 648–651, 2008.
- [31] K. Zuzak, M. Schaeberle, I. Levin, N. Lewis, J. Freeman, J. McNeil, and L. Cancio. Visible and infrared hyperspectral visualization of normal and ischemic tissue. In *Proc. of the First Joint BMES/EMBS Conf.*, volume 2, page 1118, 1999.