The Digital Atheneum: New Approaches for Preserving, Restoring and Analyzing Damaged Manuscripts *

Michael S. Brown and W. Brent Seales
Department of Computer Science
University of Kentucky
Hardymon Building, 2nd Floor
301 Rose St.
Lexington, KY 40506, USA
{mbrown,seales}@dcs.uky.edu

ABSTRACT

This paper presents research focused on developing new techniques and algorithms for the digital acquisition, restoration, and study of damaged manuscripts. We present results from an acquisition effort in partnership with the British Library, funded through the NSF DLI-2 program, designed to capture 3-D models of old and damaged manuscripts. We show how these 3-D facsimiles can be analyzed and manipulated in ways that are tedious or even impossible if confined to the physical manuscript. In particular, we present results from a restoration framework we have developed for "flattening" the 3-D representation of badly warped manuscripts. We expect these research directions to give scholars more sophisticated methods to preserve, restore, and better understand the physical objects they study.

Keywords

Digital Preservation, Humanities Computing, Image Restoration, Document Analysis, Digital Libraries

1. INTRODUCTION

There are now major efforts being undertaken throughout the world to digitize and preserve significant materials [13, 8]. Digital acquisition, which is the conversion of physical materials into a digital format, allows the possibility of efficient dissemination, and serves as a means of preservation. In addition, the digital facsimile can be manipulated in ways that are not possible for a fragile, physical artifact. Such manipulation can be used to digitally restore or enhance damaged materials. This is particularly true for digitized handwritten documents, where image processing algorithms can enhance illegible materials and provide improved data for the interested scholarly community [10, 3].

Traditionally, digitization and subsequent digital enhancement has been limited to 2-D images. This limitation is now changing, with several recent digitization efforts [2, 12, 5] focused on capturing highly detailed facsimiles using 3-D acquisition techniques. As the media stored in the digital library evolves into new and more expressive forms, we must develop new approaches and algorithms for manipulating, processing, and enhancing it.

In this paper we present research results from aspects of the *Digital Atheneum* [1], a National Science Foundation Digital Library Initiative Phase Two project. The *Digital Atheneum* encompasses research into new techniques to restore and analyze digitized collections. In particular, we are interested in new methods for acquiring and manipulating realistic facsimiles of damaged manuscripts for the purpose of enabling scholars to use these facsimiles in new ways to gain a better understanding of the physical items. Because many damaged manuscripts are no longer flat, our work involves capturing both the *images* of the manuscript, and the *three-dimensional structure* of the manuscript in the form of a high resolution shape model. Such 3-D models offer an array of uses beyond the 2-D images. For example, an accurate 3-D representation allows metric measurements to be made on the surface of the model. As described in Section 3, such measurements are valuable in a number of contexts. Furthermore, in the case of warped and crinkled documents, our recent research shows how to use the 3-D model for "virtual" flattening.

The remainder of this paper details three aspects of our research. Section 2 presents results from a 3-D acquisition effort in conjunction with collaborators at the British Library. Section 3 gives examples of how the 3-D data can be analyzed via user-specified measurements, and Section 4 presents a technical framework for restoring warped documents by flattening their 3-D facsimile.

2. 3-D ACQUISITION

2.1 Creating Digital Facsimiles

Digital photography is the most common means of creating digital content from non-traditional library materials, such as items found in special collections. While the 2-D image provides a representation that is familiar and widely accepted, it has fundamental limitations. A solitary image cannot unambiguously represent metric scale for all points within the photograph. The usual solution to this problem is to insert meta-data that describes dimensions, or to visually place a ruler next to the object during imaging. This

* We gratefully acknowledge support for this work by the NSF DLI-2 award #9817483

[1] www.digitalatheneum.org

JCDL'01, June 24-28, 2001, Roanoke, Virginia, USA.

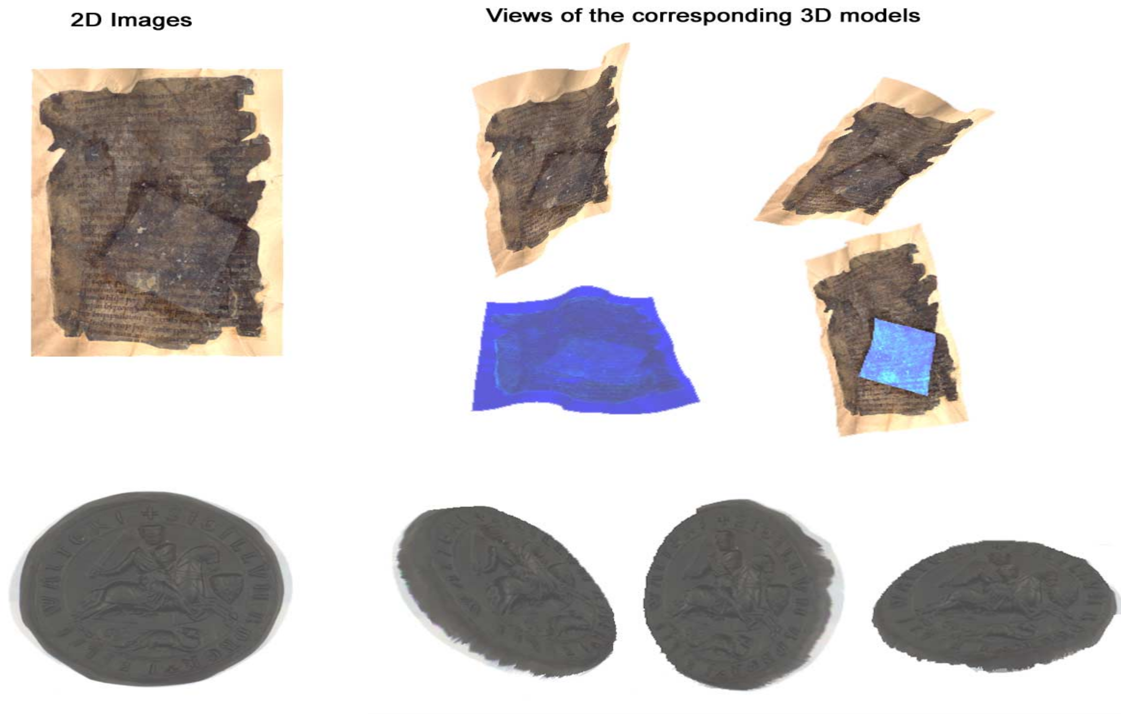

Figure 1: Top Row: A 2-D image juxtaposed with renderings of an acquired 3-D model of a manuscript shows the amount of relief the manuscript contains. This manuscript was imaged under white-light and UV light. The two images can be composited together to form a new texture for the 3-D model. Bottom Row: This acquired 3-D model of a wax seal captures detailed metric shape information.

approach to digitization makes the assumption that the object is flat, which is reasonable when considering many printed materials. There are many older, damaged texts, however, which have become warped and crinkled from age and deterioration. In addition, there are hosts of other items that have inherent 3-D shape, such as wax seals, coins, tablets, leather book bindings, etc. For such items, the image alone is insufficient to capture true 3-D shape.

We are addressing this acquisition problem as part of the Digital Atheneum [5], and have developed a structured-light computer vision technique which uses a light projector and camera to capture 3-D models. In this technique, the projector projects vertical or horizontal stripes of light onto the object. The camera observes these projected stripes and can determine the 3-D shape of the illuminated object by measuring the warp in the stripes. In the following section, we discuss issues in using this technique, and present results from manuscripts scanned at the British Library.

2.2 Acquiring 3-D Materials

The British Library has imaged a number of collections with its ultra high resolution digital color camera from Kontron Elektronik GmbH [11]. This camera is capable of capturing images at a pixel resolution of roughly $4K \times 3K$. Unfortunately, the interface for acquiring an image is proprietary, and a software development kit (SDK) is not available. In addition, capturing an image at the highest resolution takes several minutes and must be performed through an Adobe Photoshop plugin. As a result, it was impractical to use this camera for capturing a large number of images for the purpose of recovering a 3-D representation.

The Kontron camera has a continuous PAL signal that can be used for external monitoring of the camera field of view. This feature allows continuous feedback when positioning and aligning the materials beneath the camera. We captured this PAL signal (768×576 pixel resolution), which is generated from the same optical path used to scan high-resolution data, at 24 frames per second. Using the PAL signal we were able to recover the 3-D shape of manuscripts using structured light [5]. Because the PAL signal and the high-resolution images are created by the same optical pathway (i.e. same lens and same sensor), registration between the imagery is straightforward. In this way we acquired 3-D data using the PAL video signal, and acquired higher resolution imagery for textures.

2.3 Results

Figure 1 shows views of the some of the acquired 3-D models. Many of the manuscripts were photographed using both white light and ultra-violet (UV) light. UV light has been successfully used to enhance certain texts that are badly damaged and difficult to see with the unaided eye [15]. One advantage of our 3-D acquisition technique over commercial laser scanners is the ability to register multiple textures easily and accurately to a single 3-D model, thus allowing for accurate compositing of textures. Figure 2 shows a visualization application for these models. This tool allows the user to select a particular model, and choose from any number of corresponding textures. The user can interactively rotate, translate, and zoom the 3-D model.

In addition to manuscript pages, we tested the acquisition system on other items, such as a wax seal (Figure 1). Overall, we acquired 3-D shape and accompanying texture models for twenty three items.

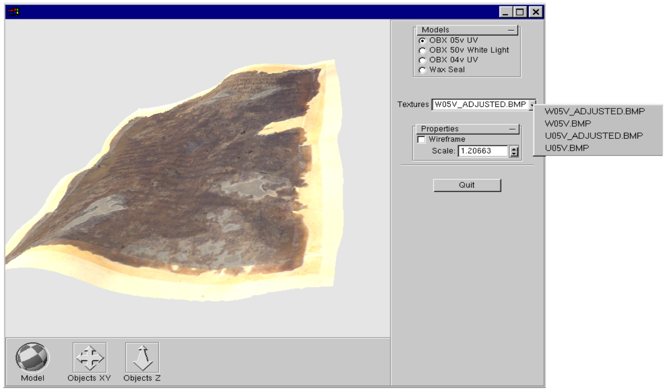

Figure 2: This screen-shot shows one of our applications for viewing an acquired 3-D model. The tool allows the user to manipulate the view of selected models and to visualized their structure in 3-D with any number of corresponding textures.

2.4 Improving the 3-D Scanner

Although we obtained good 3-D results using the PAL signal from the Kontron digital camera, a preferable solution that is likely to be more reliable and accurate is to use a high resolution camera that is supported by an available application programming interface. For example, the Kodak Professional DCS series of cameras use the high-speed IEEE 1394 interface (commonly called *firewire* or *iLink*). These cameras are available at megapixel resolutions, and Kodak provides an SDK. We are currently designing a new scanner using the Kodak DCS 330 camera, which is capable of capturing a $2K \times 1.5K$ color image and transferring the image to a host machine in roughly 10 seconds. The IEEE 1394 interface allows the camera to be driven by a notebook computer, such as a Sony VAIO. With a compact light source such as the 5 lb. Epson Powerlite projector, which can easily be mounted on a tripod, the entire system becomes even more portable. Our goal is to develop a compact, fully portable 3-D acquisition setup, which is affordable and can produce very accurate digital facsimiles.

3. ANALYSIS

The 3-D model that we acquire captures the metric scale of the original[2]. This model can be converted into a depth image. A depth image is an extended image where each pixel is given an associated "depth" value. Thus each image value $I(u,v)$ is represented by a tuple (r,b,g,x,y,z), where (u,v) is the depth image coordinate, r,g,b represent the pixel's Red, Green, and Blue color values, and x,y,z is the 3-D point recovered for the pixel position. Although this representation is larger than a standard intensity image, it directly incorporates a recovered 3-D depth representation and is easy to manipulate.

The user can perform a number of interesting operations using the depth image, which tightly couples 3-D points to pixels in the image. For instance, if the user selects two image points, $I_1(u,v)$ and $I_2(s,t)$, the metric distance, d, between these two points can be calculated directly as

$$d = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2} \qquad (1)$$

[2] We acquire models at correct metric scale, within an error tolerance. We have estimated the mean value of this error to be 0.3mm for the 3-D models acquired at the British Library. See [5] for further details regarding how these error estimates are made.

where x_k, y_k, z_k corresponds to the respective metric 3-D coordinates for pixel I_k stored in the depth image. The ability to make such direct metric measurements provides users with a powerful means to analyze digital facsimiles.

3.1 Examples of Metric Measurements

Figure 3 shows some examples of measurements made using 3-D facsimiles. These measurements can be computed as the direct distance between two points, or calculated as the distance along the surface of the object. In addition, irregular regions, such as the holes in the manuscript in Figure 3(c) and (d), can be selected by the user and measured. This is done by specifying a region with several connected line segments. The overall distance is simply the sum of the individual segments. Making these same measurements on the real object would be tedious if not impossible, when performed with standard tools such as a caliper or ruler.

3.2 Uses of Measurements

We envision that metric measurements may be useful in the following instances:

Monitoring Damage: From the measurements made on the surface of an object, it may be possible to monitor damaged areas over time. For example, the hole measured in Figure 3(c) could be measured periodically to see if has become enlarged. Such measurements could be made before and after an item is loaned to another institution to monitor damage from shipping and handling.

Surface Area and Volume: In addition to measuring surface distances, the 3-D representation makes it possible to determine the surface area and volume of objects. This data, combined with weight measurements, can be used to determine an object's density and thereby possible composition.

Handwriting Analysis: Brush stroke metrics and metric letter form analysis can be performed. These measurements may by useful as another tool for making arguments about authorship. Moreover, accurate measurements may help determine how to re-assemble or re-associate fragments that are physically separate but may be part of the same collection.

We have provided a technical framework that will allow scholars to perform metric measurements on collections. Our framework is independent of the importance and semantics of a particular collection. We believe that by placing this new capability into the hands of scholars who are keenly interested in the content and meaning of various objects, we will enable them to conduct a substantially more sophisticated study.

4. RESTORATION: VIRTUAL FLATTENING

Although we are able to create a 3-D model that encodes the shape of a manuscript, it is quite desirable to produce a flat facsimile even when the physical manuscript is no longer flat. A flat facsimile would make a warped document easier to read. In addition, subsequent image processing operations that derive features from a digital image, such as Optical Character Recognition (OCR) [14] and Hand Writing Recognition [7] algorithms, rely on the assumption that the input images are of flat documents.

We have developed a framework to help restore an image of a warped document by virtually flattening its 3-D model. This is achieved using a physically-based mass-spring system. Physically-based systems are typically used in computer graphics algorithms to simulate the dynamic deformation of 3-D models over time. One

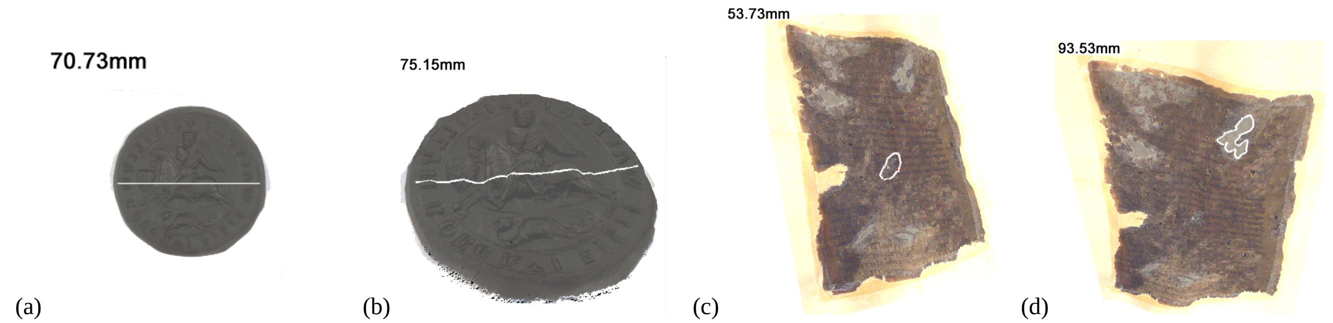

Figure 3: (a) The distance measurement, shown in the upper left corner of the image, can be specified by the user. (b) Using the same user-selected points, the measurement can be made along the 3-D *surface* of the seal, giving a slightly larger distance. This measurement would be extremely difficult to make on the physical object. (c) and (d) show circumference measurements of irregularly shaped holes on the manuscript.

notable application is in *cloth modeling*, where a flat sheet, representing a piece of cloth, is "dropped" and deforms as it hits obstacles in the simulated environment [17, 1]. The shape of the cloth changes according to the geometry of the colliding obstacles and properties of the simulation, such as gravity, the elasticity of the cloth, and so on. The manuscript flattening process can be cast as the inverse problem: given a sheet (a manuscript) in which deformations have already been applied, how can the simulation undo them to obtain the original, flat shape? The starting point for the simulation is the exact, warped 3-D shape, which we can obtain with our acquisition system. We initialize a mass-spring "sheet" with this warped 3-D shape, and force it to collide with a flat plane, which unwarps the manuscript.

The next section gives an overview of the mass-spring system and shows how we apply it to obtain results using this approach. Further details and experiments can be found in [4, 6].

4.1 Mass-Spring System

Recovered 3-D points on the surface of a manuscript form what can be considered as a system of *particles* that are able to move in 3-space. A particle system is governed by the classic second order Newtonian equation, $f = ma$, where f is a force, m is the mass of a particle, and a is an acceleration. A particle modeled by this equation can be described by its *phase state* with six variables $[x_1, x_2, x_3, v_1, v_2, v_3]$, where x_i represents the particle's 3-space position, and v_i represents its velocity. The phase state derivative with respect to time, and the subsequent motion equation, is $[v_1, v_2, v_3, f_1/m, f_2/m, f_3/m]$. This system describes a particle's mass, position and velocity at a given instance in time. During simulation, dynamic external forces such as gravity and collision forces are exerted on these particles over time. New particle positions are calculated based on these forces applied according to the equations as the time variable advances.

In a basic particle system, individual particles respond only to *external* forces, and have no influence on other particles. However, this basic system can be extended to incorporate forces between particles. One common extension, referred to as a mass-spring particle system, is formulated by logically connecting particles together via *springs*. The resulting forces in such a system can be classified into two types: *internal*, or forces between particles; and *external* forces. The slightly modified equation expressing this is

$$F_{int} + F_{ext} = ma \qquad (2)$$

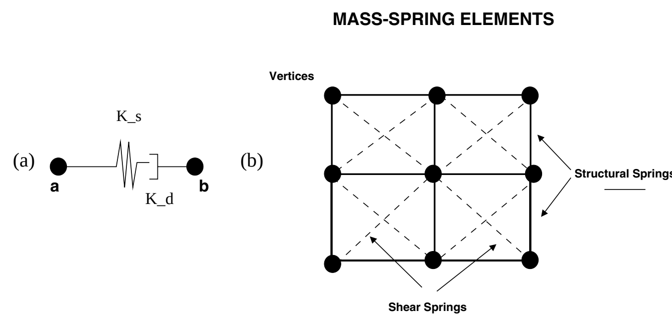

Figure 4: (a) The ideal Hookian spring, with damper, acts on two particles. K_s is the stiffness coefficient of the spring, and K_d is the dampening coefficient. (b) The finite element structure of the particles consists of structural and shear springs.

Figure 4 shows the finite elements of the mass-spring model. Particles form the vertices of quadrangles in which springs are attached. Using Provot's [16] naming convention, each element is composed of *structural* springs, which form the quadrangles' hull, and two *shear* springs, which connect diagonally. This structure is robust for modeling flexible sheet materials, such as cloth. More springs may be used to create additional rigidity if required [16].

The springs exert forces on connected particles when the two particles are moved from their resting length. These forces, governed by the ideal Hookian spring (shown in Figure 4(a)) act to keep the particles together. The Hook spring coefficients can be adjusted to control spring stiffness.

4.2 Flattening

The finite element structure described above is initialized using the acquired 3-D shape model for a manuscript. The manuscript shape is sub-sampled producing a "sheet" at a particular resolution (for example, we used 45×45 particles). This sheet is textured with the acquired 2-D image. As described in Section 3, texturing is straightforward using the depth image. Figure 5(Row II) shows examples of models viewed as non-planer sheets.

A flat collision plane is placed directly below the manuscript. A downward force (gravity) is exerted on the sheet. As the particles move downward, they collide with the plane. While this collision force tends to move particles away from one another, the internal spring forces tend to keep connected particles together. Eventually the surface of the sheet will come to rest on the collision plane when all of the internal and external forces have been minimized. At this point, the manuscript's 3-D structure has been unwarped and is flat. The flattened sheet can be textured with the original image, and the result is an unwarped 2-D image.

4.3 Experiments and Results

4.3.1 Controlled Trials

The first experiment is intended to quantify the ability of the mass-spring system to restore a deformed document to its original planar shape. Figure 5 shows images of two documents: one document is a checkerboard pattern, and the other is a set of printed letters. The documents are imaged while they are flat, serving as the experimental control. The documents are then crumpled by hand and imaged. The 3-D shape models of the documents are acquired as described in Section 2 (shown in Row II). These models provide the starting point for the mass-spring system. These initialized mass-spring meshes are subsequently flattened using the technique previously described.

The resulting *restored* images are compared to their respective control images. For the checkerboard image, we compare how closely the corners of the checkerboard align. We found that the mass-spring system provides a mean alignment error between corners in the restored image and corners in the original (control) image of $0.25mm$.

For the documents with printed letters, we compare the results under a commercial optical character recognition (OCR) package, Readiris Pro [9]. OCR is performed on the control image, the *unrestored* image, and the *restored* image. We compare the number of misses made by the OCR algorithm for these three documents. A miss is defined as any letter that is misclassified and any "noise" letters that are inserted by the character recognition algorithm. There are 176 letters present in the document. The control image was recognized with 100% accuracy, i.e., 0 misses. The *unrestored* image had 39 misses. The *restored* image was recognized with 100% accuracy (0 misses). These experiments were performed a number of times, with repeatable results [4].

4.3.2 Experiments With Manuscripts

The second experiment flattens a manuscript. Since there is no ground-truth for such an experiment, it is not possible to compare the simulation results to what the manuscript looked like before it became warped. However, this experiment shows the flexibility of the mass-spring framework for restoring such data. Two different materials are present in the scanned item. The original velum[3] document is embedded in a paper sleeve to preserve it and allow it to be bound without directly binding the vellum. These two materials, the vellum and the paper sleeve, have very different properties. Their interaction is often a cause for the overall page deformation. In cases such as this, where mixed materials must be modeled, the user can experiment with the flattening process by setting different internal force coefficients at portions of the mesh corresponding to each material. To demonstrate this, we first model the velum material with stiffness values making it *stiffer* than the surrounding paper. We compare this to the inverse setting, where the paper sleeve is made to be stiffer than the velum material. Figure 6 shows the results. Notice that the difference image between these two settings shows large variations between the *restored* images from the two experiments.

[3]parchment made from animal skin

4.4 Restoration Summary

Our restoration framework performs well with objective measures on controlled experiments when documents have undergone rigid deformations, such as paper being crumpled by hand. For a decaying manuscript, however, it may be impossible to model all of the physical phenomena contributing to the deformed state. For such items, we are interested in manipulating the model in a reasonable and flexible way to help restore the *perceptual* quality of the digital representation. Our hope is to extend the current framework to allow user-specified constraints, which can be supplied by scholars who have specific knowledge about the *content* of the imagery. Experts who understand the intricacies of letter forms and page layout, for example, may be able to use this framework to direct the "flattening" simulation for better restoration.

5. CONCLUSION

This paper has presented several aspects of the research being conducted by the DLI-2 *Digital Atheneum* project. We have presented results from a novel 3-D acquisition effort, deployed and tested at the British Library, where several high quality 3-D models of manuscripts and similar artifacts were acquired. In addition, we presented *(1)* how metric measurements, corresponding to the real metric distances on an object's surface, can be calculated using the 3-D facsimile, and *(2)* how the 3-D representation of a warped document can be "virtually" flattened. This research is part of a broader effort to establish sound principles and practices for the creation, restoration, and manipulation of quality archives, thereby aiding those communities that increasingly rely on digital content in their scholarly activities.

6. REFERENCES

[1] D. Baraff and A. Witkin. Large steps in cloth simulation. In *Computer Graphics (Proc. SIGGRAPH)*, pages 43–52, August 1998.

[2] F. Bernardini, J. Mittleman, and H. Rushmeier. Case Study: Scanning Michelangelo's Florentine Pietà. In *SIGGRAPH 99 Course 8*, Los Angeles, August 1999.

[3] G. Braudaway. Restoration of faded photographic transparencies by digital image processing. In *Proceedings of IS and T's 46 Annual Conference*, pages 287–289, 1993.

[4] M. S. Brown and W. B Seales. Document restoration using 3d shape. Technical report no. 312-01, University of Kentucky, Lexington, Kentucky, Jan 2001.

[5] M. S. Brown and W. B. Seales. Beyond 2D images: effective 3D imaging for library materials. In *Proceedings of the Fifth ACM Conference on Digital Libraries*, pages 27–36, June 2000.

[6] M. S. Brown and W. B Seales. Document restoration using 3d shape. In *International Conference on Computer Vision (ICCV)*, June 2001. to appear.

[7] K. W. Cheung, Yeung D. Y., and Chin R. T. A Bayesian framework for deformable pattern recognition with application to handwritten character recognition. In *IEEE Transaction on Pattern Analysis and Machine Intelligence*, pages 12(20):1382–1388, Dec 1998.

[8] H. M. Gladney, F. Mintzer, Schiattarella F., Bescos J., and Treu M. Digital Access to Antiquities. In *Communications of the ACM*, pages 41(4):49–57, April 1998.

[9] Image Recognition Integrated Systems (I.R.I.S). Readiris Pro. *Rue Du Bosquet 10, 1348 Louvain-la-Nueve - Belgium*, http://www.irislink.com.

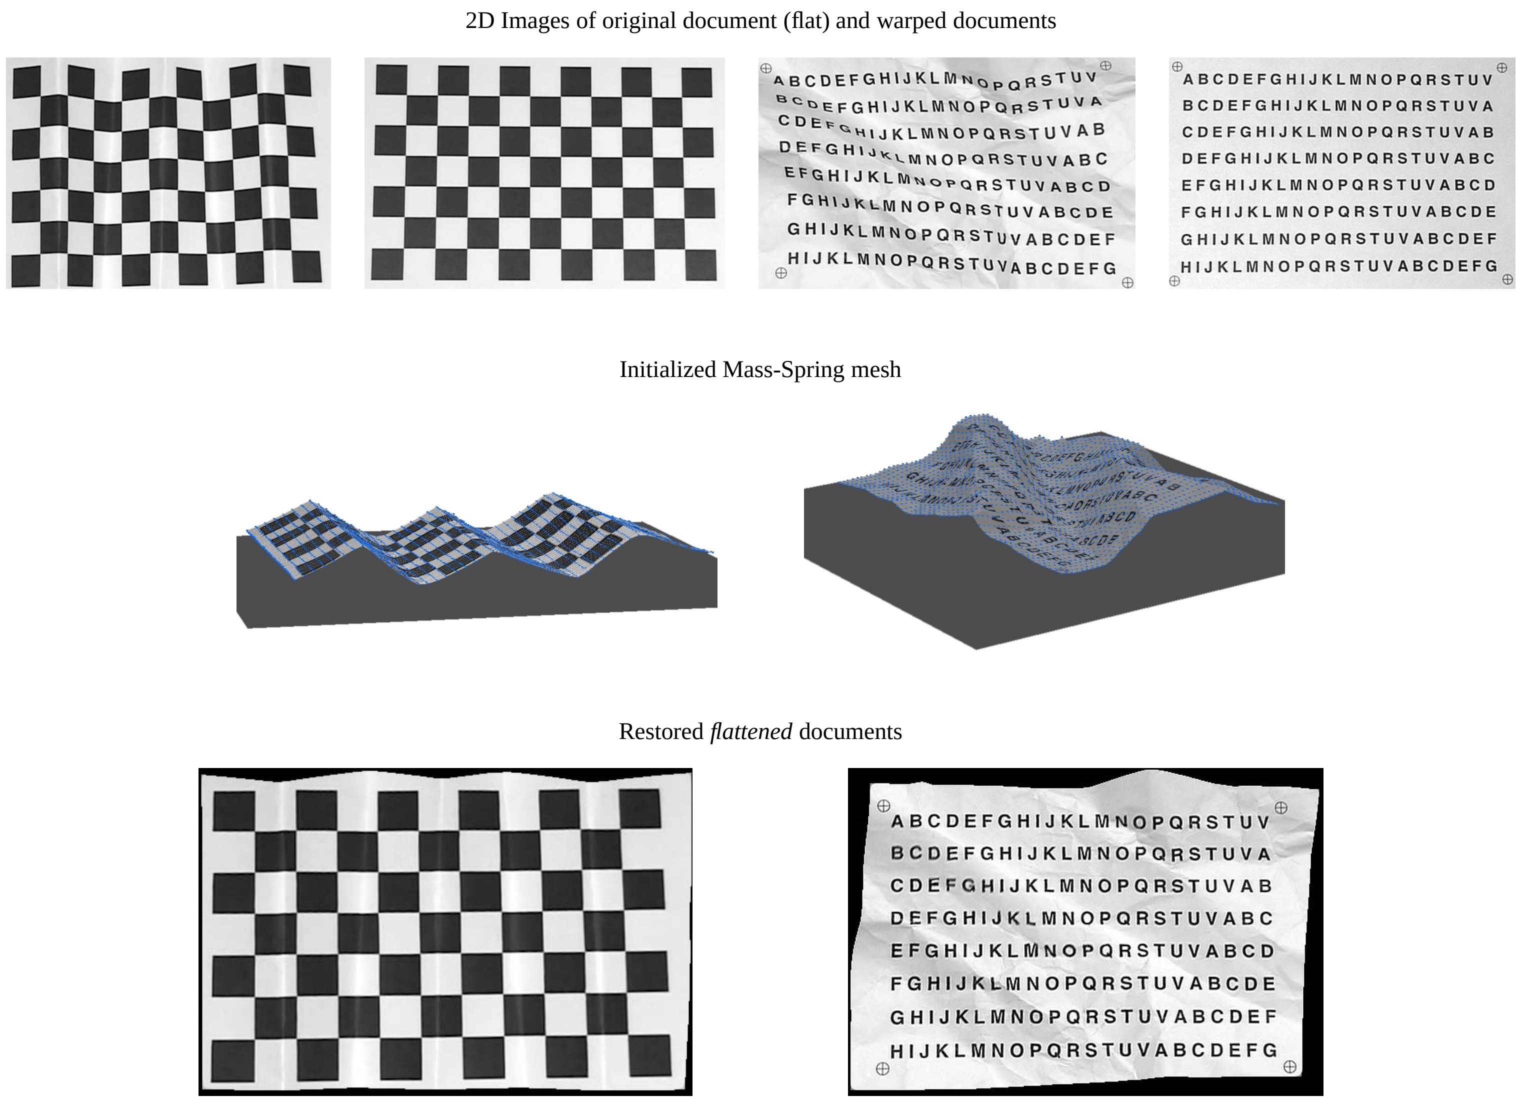

Figure 5: Experiment I Row I: 2-D images of the original flat documents serve as control before crumpling them by hand. Row II: the Mass-Spring finite-element mesh is initially structured from the corresponding 3-D facsimile. Row III: the original 2-D image is correctly texture-mapped onto the restored (*flattened*) shape model.

[10] K.S. Kiernan. Digital Image Processing and the Beowulf Manuscript. *Literary and Linguistic Computing: Special Issue on Computers and Medieval Studies*, 1991.

[11] Kontron Elektronic GmbH. Kontron embedded computers ag. *Oskar-von-Miller-Strabe 1, 85386 Eching - Germany*, http://www.kontron.com.

[12] M. Levoy. The Digital Michelangelo Project. In *Proceedings of the Second International Conference on 3D Digital Imaging and Modeling*, October 5-8, 1999.

[13] B. Mannoni. Bringing Museums Online. In *Communications of the ACM*, pages 39(6):100–105, June 1996.

[14] Pavlidis, T. and Mori, S. (eds.). Optical character recognition. In *Special Issue of Proceedings of the IEEE*, pages 7(80), 1027–1209, July 1992.

[15] A. Prescott. The electronic Beowulf and digital restoration. *Literary and Linguistic Computing*, pages 12(197),185–95, 1997.

[16] Xavier Provot. Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *Graphics Interface*, pages 174–155, 1995.

[17] D. Terzopoulos, J.C. Platt, and A.H. Barr. Elastically deformable models. In *Computer Graphics (Proc. SIGGRAPH)*, pages 21:205–214, 1987.

Mass-Spring meshes with non-uniform spring coefficients

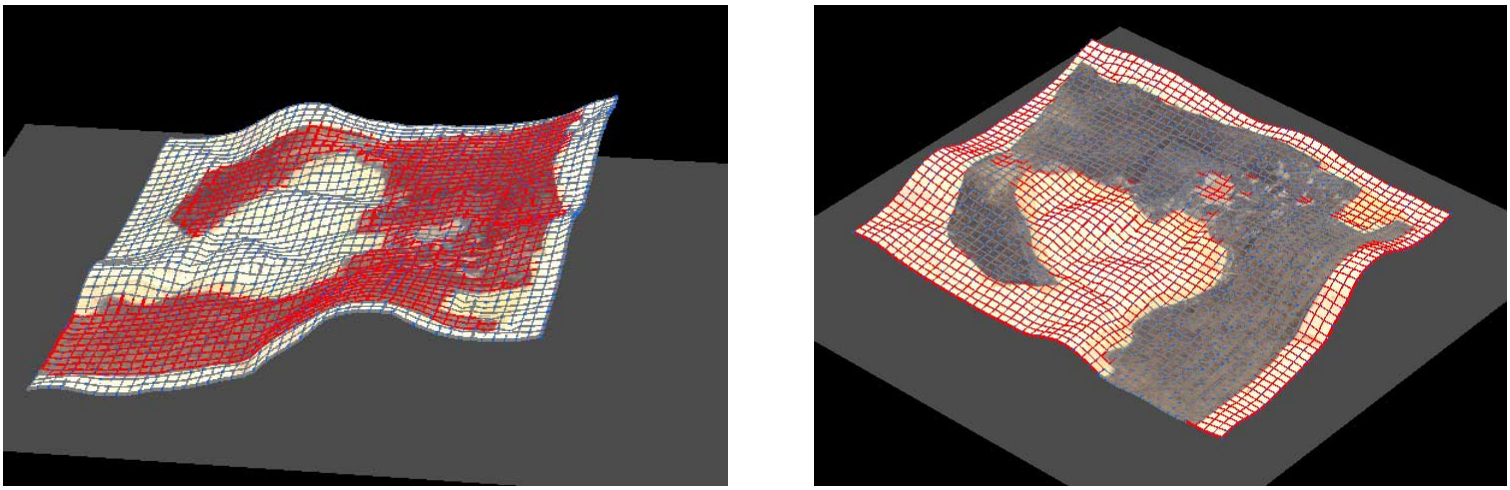

Restored 2-D manuscript and difference image

Figure 6: Experiment II Top Row: The spring stiffness coefficients are non-uniform across the Mass-Spring finite-element meshes for a manuscript. The first mesh has stiffer spring parameters for the velum portion, and the second mesh is stiffer in the paper portion. Bottom Row: restored images and difference images between the two simulations.