

Integrating Spectral and Reflectance Transformation Imaging Technologies for the Digitization of Manuscripts and Other Cultural Artifacts

Final Report on Experiments Conducted and Lessons Learned (White Paper)

Todd R. Hanneken, Ph.D., Project Director

This project sought to determine the viability, procedures, and benefits of combining spectral (sometimes called multi-spectral or hyper-spectral) imaging technology with Reflectance Transformation Imaging (RTI) technology. The basic premise is that the two approaches are fundamentally compatible and complementary because spectral imaging primarily addresses spectral reflectance properties of the material (chrominance) and RTI primarily addresses luminance reflectance properties of the surface topography (texture). That is, the color information from spectral imaging can be combined with the texture information (conveyed through highlights and shadows) of RTI. Two experimental approaches were tested and compared to previously established procedures employed for each separate type of imaging. The first experimental approach proved to be a time-efficient and quality-effective means to integrate the interactivity and texture information of RTI with the processed derivatives, spatial resolution, and color resolution of spectral imaging. The second experimental approach had a much higher cost in time and equipment and produced mixed benefits. In particular, the pseudocolor processing based on principal component analysis (PCA) produced inconsistent coloration that washed out when combined into an RTI image. Although various improvements may be possible, the simpler approach suffices for the particular goal of producing RTI images based on PCA pseudocolor. For special situations when the more thorough approach is justified, the large amount of data might be best rendered in ways other than RTI.

In addition to the core objective of integrating spectral and RTI, some new processing techniques were developed that could benefit spectral imaging even without the addition of RTI. First, "Extended Spectrum" uses PCA to reduce a set of images with detailed color resolution to the color resolution of the human eye. Second, pseudocolor images can be created by mapping monochrome images to assigned color spaces that distinguish luminance from chrominance (such as YCbCr or L*A*B*) and then converting to RGB. This approach correlates with Principal Component Analysis in that the first principal component often resembles the luminance range and subsequent principal components represent variation in chrominance.

Selection of Objects

The integration of spectral and RTI is valuable for imaging projects which would benefit from each technology individually, namely the imaging of objects that benefit from analysis of both color and texture. Spectral imaging technology is appropriate for objects with meaningful contrasts in "color," both in the usual sense of spectral properties in the visible wavelength range, and extended to contrasts beyond the range of human vision. RTI is appropriate for objects with meaningful texture, such as inscriptions and coins; but it is also applicable to fine texture variations, including the effects of corrosion of acidic ink on parchment. For our experiments we selected objects that we expected to benefit from enhanced visualization of both

color and texture contrasts. In contrast, coins and inscriptions on homogenous materials might not benefit significantly from spectral imaging, and faded ink on paper (such as the Livingstone diaries) might not benefit significantly from RTI. Among objects that met the criterion of benefiting from both technologies we selected objects reflecting a variety of depths of texture (from a relatively flat palimpsest to a figurine) and complexities of color (from a very faded, non-rewritten palimpsest to a colorful painted portrait). They are as follows:

UCLA Palimpsest (Pal1)

This palimpsest is relatively flat and shows little apparent chromatic variation. The texture conveys fine details such as the corrosion of parchment where ink had been. The primary color contrast is between traces of erased ink and blank parchment. This palimpsest is unusual in that it was erased but not overwritten (and thus might not technically be considered a palimpsest, but is referred to here as such because it served to test technologies for reading palimpsests). On the one hand, the absence of interference from another text eliminated an impediment to reading the erased text. On the other hand, it also eliminated the opportunity to demonstrate a clear advantage of PCA pseudocolor, which is the ability to make two different but visually similar inks appear in highly contrasting false colors. Another advantage of the palimpsest is that quality evaluation could be relatively objective according to the criterion of readability of text. An additional page of the palimpsest was captured as a backup but not fully processed or analyzed. This data can be found in the complete data archive under the code "Pal3."

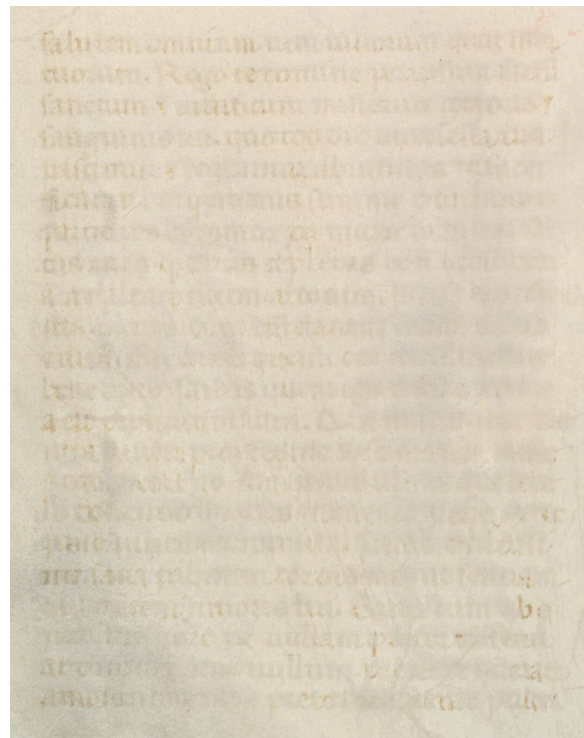


Figure 1: UCLA Palimpsest in Accurate Color
(click image to open full-size image in browser)

The catalog information for the UCLA palimpsest is as follows.

Museum Accession Number: UCLA Rouse Ms. 32

Physical Object Description: Sixty leaves, 56 of which have been scrubbed, sparing the initials and the rubrics; some text present on various pages; parchment (Italian preparation), 240 x 174 (122 x 94) mm. Quires of 8 leaves, primarily from quires VIII-XVII; leaf signatures; catchwords. 22 long lines ruled in hard point. Written by one person in a round Italian littera textualis.

Text: Medieval Latin Book of Hours commissioned in northern Italy, with illuminations painted by the Master of the Brussels Initials. Palimpsest, no overtext.

Provenance: ca. 1400 CE, Italy

InscriptiFact Number: 05001

Bibliography: Kidd, Peter. "UCLA Rouse MS 32: The Provenance of a Dismembered Italian Illuminated Book of Hours Illuminated by the Master of the Brussels Initials," in *Medieval Manuscripts, Their Makers and Users: A Special Issue of Viator in Honor of Richard and Mary Rouse* (Turnhout: Brepols, 2011), pp. 279–91.

British Library Record: <http://www.bl.uk/catalogues/illuminatedmanuscripts/record.asp?MSID=8033&CollID=27&NStart=52539>

USC Mummy Mask (Mask)

This partial mummy mask depicting the head and shoulders of a young man has slightly greater complexity in texture, including strokes of paint, chips of missing paint, and fissures in the underlying wood. The color information is exposed wood and a simple palette of paint. Subtle features that stand out to varying degrees in different renderings include: a chip of paint missing from the left cheek of the figure, a drip of paint on the nose, the contrast along the side of the chin between the paint used for the background and the paint used for the cheek, an indented line in the pupil from the ten-o'clock position to the four-o'clock position, and various other chips, striations, and brush strokes throughout the object.

The catalog information for the USC mummy mask is as follows.

Museum Accession Number: USCARC 7498 (University of Southern California Archaeological Research Collection)

Physical Object Description: Mummy portrait, wood, no text, 33.3 x 7.8 cm, less than 0.5 cm thick

Provenance: Roman Period Egypt, 2nd Century CE, thought to originate in the Fayyum region

InscriptiFact Number: 00556



Figure 2: USC Mummy Mask in Accurate Color

USC Terracotta Figurine of Conqueror and Conquered Soldier (Sold)

The terracotta figurine has significant depth, sufficient to test (and show) the limits of the approach. Limits of concern include the imperfect diffuseness of the Eureka Lights and the slight difference between the points of origin of different colors of light in the Magic Flashlight (see below for descriptions of these light sources). Principal Component Analysis might have been problematic because it does not discriminate between variation due to reflectance properties and variation due to shadows. Deep crevices are found between the figures depicted, namely between the conqueror and the falcon, and between the conqueror and the conquered. There are also finer textures of various types: some of which were intended by the artist, others that result from the manufacture of the object, and still others due to deterioration attributable to the age of the object.

Chrominance comes from terracotta and a few traces of the original paint. The results of an archaeological study of the object by Grant Dixon are available online at

<https://dornsife.usc.edu/what-is-a-king-to-do/>. In particular, evidence from comparison with similar objects confirms the



Figure 3: USC Terracotta Figurine in Accurate Color

identification of some details, such as the anatomically-incorrect position of the conqueror's hand and sword.

The catalog information for the USC terracotta figurine is as follows.

Museum Accession Number: USCARC 10663 (University of Southern California Archaeological Research Collection)

Physical Object Description: Egyptian terracotta figurine; orange clay; molded; decorated on front and sides, smooth on back; red/pink and green polychrome painting; no text; standing man with hand on top of head of much smaller man with shield indicating military defeat or domination

Provenance: Roman Period Egypt, perhaps second century C.E.

InscriptiFact Number: 03625

USC Illuminated Antiphonary Manuscript, Folio 2 Recto (Ant2)

A supplementary object, an illuminated manuscript, has texture from wrinkling of parchment and the thickness of paint. The chrominance is markedly rich. One subtle feature is smudging in the text of the antiphon consistent with contact with another page. Another subtle feature appears to be the signature of the artist beneath the antiphonal text, most prominently two large cursive letters "M." A detail of an additional folio was captured using the MegaVision camera, Eureka Lights, and flash in order to demonstrate the spatial resolution capabilities of the fifty mega-pixel camera (2190 dpi). This data was not fully processed and evaluated but can be found in the data archive under the code "Ant3."

The catalog information for the USC illuminated antiphonary manuscript is as follows.

USC Flewelling Antiphonary

Physical Object Description: Bound parchment manuscript with illuminations. Codex 17 x 23 cm, 3.5 cm thick at first illustration; illustrations have been "improved" (by painting over).

Text: Latin Antiphonary

Provenance: medieval, ca. 15th-16th century, most likely German

InscriptiFact Number: 05002



Figure 4: USC Illuminated Antiphonary Manuscript, Accurate Color

Overview of Methods

The following table provides an overview of the two established approaches and the two experimental approaches. The established conventions for spectral imaging, as practiced by the Early Manuscripts Electronic Library (EMEL), and RTI, as practiced by the West Semitic Research Project at the University of Southern California (WSRP), were conducted at the same time as the two experimental approaches for purposes of comparison.

The control spectral imaging procedure captured eleven images per object using a fifty-megapixel monochrome MegaVision camera and two banks of Eureka Lights (narrow band diffuse illuminators capable of projecting eleven discrete wavelengths from ultraviolet through infrared). The control RTI procedure captured thirty-five images per object using a DSLR

camera and conventional electronic flash. The first experimental method captured forty-six images (eleven plus thirty-five) per object using the MegaVision camera and the illuminators from both of the established technologies (eleven narrow bands of diffuse light from the Eureka Lights and thirty-five angles of the flash). The second experimental method captured 315 images (thirty-five times nine) per object using the MegaVision camera and a prototype Magic Flashlight, which combined the two types of illuminators in hardware. Like the conventional flash, the Magic Flashlight could be moved to thirty-five positions on a virtual hemisphere around the object. Like the Eureka Lights, the Magic Flashlight could illuminate in narrow bands within and beyond the visual spectrum of light (nine bands from ultraviolet through infrared).

Each of the control set procedures is capable of limited output. Spectral imaging produces two-dimensional static images with a variety of enhanced renderings. Reflectance Transformation Imaging produces interactive images in either Polynomial Texture Map (PTM) format or Hemispherical Harmonics (HSH) format. These images include detailed texture information that can be visualized in real time by using a “virtual flashlight” and a variety of enhancements. Both of the experimental procedures are capable of combining the output of the two established procedures. Like conventional RTI, the experimental procedures create interactive PTM files, in addition to the series of images at thirty-five angles. Like spectral imaging, various renderings are possible, three of which are chosen here. The first rendering is accurate color. Because of the precision of illumination, the MegaVision camera and PhotoShoot software can be calibrated to produce exceedingly accurate color. The second rendering, extended spectrum, is new to this project. It seeks to map the contrasts detectable in the dataset to the contrasts detectable to the human eye, while maintaining a semblance of reality (if humans had nine or more color receptors instead of three). The effect is that pixels that reflect ultraviolet appear blue and pixels that reflect infrared appear red. The third rendering is pseudocolor, which uses Principal Component Analysis to find the greatest contrasts in a complex data set and render them in contrasts perceptible to the human eye. Pseudocolor makes no claim to convey color accurately, but solely to illustrate contrasts as starkly as possible.

	Spectral Imaging Control Set	RTI Control Set	Eureka + Flash Experiment	Magic Flashlight Experiment
Camera	MegaVision monochrome	DSLR RGB	MegaVision monochrome	MegaVision monochrome
Lighting	11 narrow bands from UV through IR from two positions with diffusers (Eureka Lights)	Electronic flash (xenon tube) at 35 angles	Eureka Lights for diffuse narrow-band capture and flash for white capture at 35 angles	9 narrow bands from UV through IR at 35 angles
Number of captures	11 monochrome	35 RGB	46 monochrome	315 monochrome
Integration approach	n/a	n/a	software	hardware
Processing Options	accurate color, extended spectrum, principal component analysis pseudocolor	color	accurate color, extended spectrum, principal component analysis pseudocolor	accurate color, extended spectrum, principal component analysis pseudocolor
Product	digital images	PTM and digital images	PTM and digital images	PTM with problems of consistency, a greater number of digital images

Spectral Imaging Control Procedure

Spectral imaging was conducted according to the practices established by team members on other projects including the Archimedes Palimpsest Project, the Sinai Palimpsests Project, and the David Livingstone Diaries Project. The term “spectral imaging” is sometimes used loosely to include capture using a single band of illumination (or filtered broadband illumination) outside the visible spectrum, such as ultraviolet or infrared. But in the case here nine to eleven bands were captured and processed. Sometimes a single monochrome image from a narrow band is useful, but more often the series of bands is processed using Principal Component Analysis (PCA) and rendered into a more useful pseudocolor image. PCA can be thought of as a tool for finding the greatest contrasts in a complex data set.

Capture Method

Like all the procedures, the spectral imaging procedure required a mostly dark room and a rigid stand to hold the camera over the object with no movement throughout the entire procedure (Foba stand, 7.5'). The camera utilized was a MegaVision E7 operated by PhotoShoot Software on a Windows 7 PC. The camera produced 16-bit monochrome images with pixel dimensions of 8176 x 6132 (50 megapixels, 4:3 aspect ratio). Metadata was also managed by the PhotoShoot software. The camera produces images in the Digital Negative format (DNG), and monochrome 16-bit TIFF files were also created at the time of capture.

Illumination depended on two banks of Eureka Lights controlled by the PhotoShoot software. Eureka Lights were developed by William Christens-Barry of Equipose Imaging originally for the Archimedes Palimpsest Project. They use LEDs to emit narrow bands of illumination at eleven wavelengths ranging from 393 nanometers (ultraviolet) to 830 nanometers (infrared).

After setting up the camera stand, camera, object stand, and lights, but prior to capturing images of the objects, a series of captures was taken for calibration and flattening. Flattening refers to the process of correcting for unevenness in the light sources by capturing images of a known uniform white reference object (spectralon) and assuming that any variation in pixel intensity is due to the lighting, rather than the properties of the object. Those reference images (flats) can then be used to correct the images of the captured objects to eliminate variability due to unevenness in the light sources. Flattening requires that the known reference object be imaged in the same plane (not too high or low) as the object to be corrected, that the reference object have the same shape as the flattened object (that is, basically flat), and that the lights do not move throughout the sequence. The shape requirement prevented flattening of the terracotta figurine. Flattening is not always required or even beneficial as it has the potential to introduce more error than it eliminates if the above assumptions are not reliable. It is most valuable when attempting to recover very fine contrasts that may be discernible only in one or two wavelengths of light (as in a palimpsest). In those cases a “dim spot” from one of the lights might otherwise be misinterpreted by PCA as a meaningful variation.

Reference items were included in each frame: a ruler, spectralon, and X-Rite Colorchecker target (a shiny ball was included for other methods, as will be discussed below).

Processing Methods

Three basic processes were applied to render the captured data for display on screen: accurate color, extended spectrum, and pseudocolor.

Accurate color processing is built into the PhotoShoot software. Accurate color processed from seven bands of visible light is more accurate than RGB sensors with three band-pass filters for the reasons described on the Mega-Vision website: http://www.mega-vision.com/cultural_heritage.html. The X-Rite Colorchecker target was used for white balance.

The extended spectrum process is new to this project but can be applied to spectral capture sequences even without the addition of RTI. The premise is to divide the captured wavelengths into three categories. In the case of the captures with the Eureka Lights the eleven wavelengths were divided into the three longest wavelengths, the four middle wavelengths, and the four shortest wavelengths. Within each category PCA was applied to the three or four captures. The first principal component, representing the greatest contrast detectable in that range of color, is mapped to the appropriate range in an RGB image. Thus, the first principal component of the infrared and red captures becomes the red channel, the first principal component of the ultraviolet and blue captures becomes the blue channel, and the first principal component of the remaining middle captures becomes the green channel. This has the effect of using the entire range of captured images and reducing the data in an intuitive manner that maximizes contrast perception within the three channels of human color resolution. This processing was done in ImageJ (Fiji distribution) with the PCA plugin and the IntegratingMacros.ijm script available from the project website and data archive.

The spectral imaging control set includes pseudocolor images created using procedures established prior to the Integrating project. The control-set PCA pseudocolor images were prepared by Christens-Barry and are so indicated in folder and filenames by the initials WCB. The common principle of this approach is to create principal component images and manually select the two or three most useful images to place in the red, green, and blue channels of an RGB image. This makes it possible to visualize two or three principal component images at once. This approach is particularly helpful when one principal component shows the erased text of a palimpsest and another principal component shows the secondary text. The pseudocolor combination makes it easy to see at a glance one layer of writing in blue and the other in red (showing only one layer is less effective because the upper layer tends to obfuscate the lower layer). This method of creating pseudocolor differs from the experimental methods below in two major ways. First, it maps principal components directly to the RGB color space without using the color spaces that distinguish luminance from chrominance (YCbCr or LAB). Second, it involves the creation and evaluation of many images and the manual selection of the best two or three to use; the experimental method below worked with a preference for finding a predictable pattern that could be automated.



Figure 5: Spectral Imaging Extended Spectrum, Terracotta Figurine

Christens-Barry processed the objects Mask, Pal1, and Sold and created a variety of pseudocolor images using ImageJ software with the PCA plugin. PCA produces principal component images equal in number to the number of input images, which was at least eleven but sometimes more when additional images were captured using a filter. Christens-Barry also experimented with modifying the PCA routine to allow negative eigenvalues (indicated “ne” in filenames) and normalization (indicated “N” in filenames). For each object Christens-Barry selected a “first choice” representative of the best that the established technology could do. For the object Pal1 (palimpsest) the method used normalization, PCA, manual evaluation of principal components, and mapping of the first principal component to the R channel and the fifteenth principal component to both the G and B channels. For the object Mask (mummy mask) the method used normalization, PCA allowing negative eigenvalues, manual evaluation of the principal components, and mapping of the first, second, and fifth principal component to the RGB channels respectively. For the object Sold (terracotta figurine) the method used normalization before and after principal component analysis, manual evaluation of the first five principal components, and mapping of principal components two, three, and five to the RGB channels respectively.




Figure 6: Spectral Imaging Control Pseudocolor, Palimpsest

Quality Evaluation

The spectral imaging control set succeeded in doing what it was intended to do. It captured subtle differences in chrominance and rendered them with great precision both for accuracy and for enhancement. For the object Pal1 it should be kept in mind that the palimpsest is thoroughly erased and thus extremely difficult to read in conventional illumination. The accurate color image is in fact an accurate and even ideal representation of the vibrancy of color. The pseudocolor is effective at making apparent the traces of ink where any is left to be discerned. The main limitation in quality compared to the other sets pertains to the lack of texture imaging. The image is not interactive, and it does not discern texture, which is effective in showing where ink had once been even where no ink remains. For the object Mask the color vibrancy is accurate and ideal, and the processed images accentuate differences in materials. For the object Sold it is necessary to qualify the diffuseness of the diffuse light. Two overhead panels with diffusers still show shadows in the deep texture of the terracotta figurine, allowing some of the benefits of shadow-based texture analysis. PCA highlights variation whether it results from chrominance on the object or shadowing from the light source. This phenomenon did not make the pseudocolor images less interesting or useful. The control set images were also highly effective at discerning and visualizing the traces of paint on the terracotta.

Cost in Time and Equipment

The cost in time of conventional spectral imaging consists of three stages: setup, capture, and processing. Compared to most forms of photography, the spectral imaging equipment is portable only in a loose sense. Setup and calibration can easily fill a day or more depending on the challenges of circumstances. Major steps involve room darkening, installation of camera

stand, object stand, and light stands (which must be perfectly stable and immobile), computer, cable connections, and calibration captures. Capturing a series of images of spectralon allows for flattening to compensate for irregularities in the light source.

The PhotoShoot software automates the capture sequence with full control of the lights, sensor, shutter, and filter wheel. Once initiated a capture sequence takes a minute or so. The time required for capture consists mostly in mounting the object. For a steady sequence of objects with moderate conservation requirements one can achieve a pace of four minutes per page.

Processing time depends mostly on the choices and efficiency of the person evaluating the processed images. Recent reasonably powerful computers have adequate processing power, leaving the main consumption of machine time in the transferring of files over Ethernet, USB, or SATA. A fifty-megapixel monochrome image with thirty-two bits per pixel approaches 200 megabytes. The human role in processing the accurate color images involves the identification of a white standard (Colorchecker target) for white balance and confirming the calibration file and output settings. Creation of extended-spectrum processed images benefits from selection of a region of interest so that the brightest whites and darkest blacks in the reference materials do not dominate the normalization. The script also requires standardized file names, which requires some amount of human intervention. PCA and pseudocolor rendering also requires identification of a region of interest. This can go very quickly if one proceeds with assumptions of what methods are likely to be effective, and very slowly if one experiments with many possible variations. The evaluation time also depends significantly on whether the evaluator has a clear sense of what is being evaluated. For example, one can relatively quickly compare two renderings of an erased letter and identify which is more legible. For other objects identifying a “better” image can require more deliberation of what contrasts are most meaningful for scholarly analysis.

The cost of equipment puts spectral imaging out of the price range of casual enthusiasts who rely on consumer electronics. Even as technology constantly advances and prices go down, such that pixel count no longer distinguishes museum quality imaging equipment from a smart phone, a crucial separator of the markets will continue to be the Bayer filter. Consumer imaging equipment captures color in a single exposure by filtering the light to three different sensors for red, green, and blue. While this has obvious advantages in non-laboratory environments, it does introduce error and require more light to enter the aperture. Although hacks of consumer electronics, such as removing the infrared filter, have made it possible to work around some limitations, the Bayer filter cannot be easily removed. Many variables go into providing an exact cost in dollars. We can say that as of 2014 a complete spectral imaging system costs in the range of \$60,000 to \$110,000 depending on options such as filters, lenses, spectral range, and power. The camera alone constitutes \$30,000 to \$50,000 of that amount. Equipment rental is also an option, and would require \$5000 to \$10,000 for equipment for a three to five day week, which is the minimum amount of time an imaging session would require.

Conclusion: Benefits and Challenges Going Forward

The cost of spectral imaging equipment is justified in two major ways. First, the quality and precision in color and spatial resolution are higher than consumer imaging equipment. The accurate color images are calibrated without human subjectivity and capture a remarkable level of detail. Second, the quantity of data captured is not limited to the ability of the human eye, which is particularly valuable when the objective is not only accuracy but recovery of information that is invisible to the human eye, as is the case when reading palimpsests. While the

first justification is an improvement in quality, the second justification introduces an entirely new range of possibilities. One might add additional advantages, such as the fact that spectral imaging requires minimal exposure to light and wastes little or none of the reflected light on filtration. The benefit of spectral imaging is limited to objects that contain meaningful color information, visible or invisible. For example, spectral imaging might not be justified for imaging objects of homogenous material, such as coins or inscriptions (unless variation in oxidation and traces of deposits are important for conservation and study). The current state of the technology also assumes a relatively small target object, not far from a maximum of twelve inches on any dimension. Larger projects would pose challenges to spatial resolution and illumination. Going forward we look for improvements in two directions. First, we hope to see improved accessibility through lower cost and greater simplicity of application. Second, we imagine variations that would use multiple cameras and an advanced mounting apparatus to capture dimensionality of objects that cannot be approximated by a two-dimensional plane.

RTI Control Procedure

Reflectance Transformation Imaging was conducted according to the practices established by the West Semitic Research Project at the University of Southern California. The only variation was to mount the camera on a “dual-head” camera mount alongside the MegaVision camera on the Foba stand. This allowed the captures to be taken at the exact same time without moving the object. The same firing of the flash was captured by both cameras (this was done by setting a 2-second exposure on the MegaVision camera, waiting for the sound of the shutter opening, triggering the DSLR camera from the control computer, which in turn triggered the flash). However, for the initial days of capture this caused some reduction in spatial resolution because the two parallel cameras, though very close together, did not entirely overlap in their fields of view. Consequently both cameras had to zoom out to capture a wider field of view than otherwise might have been required. This was solved by the end of the capture week by renting a tilt-shift lens. This allowed the sensor of the DSLR camera to remain parallel to the object plane while shifting the lens sideways such that the image of object could be moved to a position on the sensor that closely matched the image on the MegaVision camera. The objects imaged with the tilt shift lens were Pal1 and Pal3 (the UCLA palimpsest). The combination also made it important that we used a flash that was triggered by a radio signal rather than an infrared signal from the camera-mounted transmitter (such an infrared signal would be picked up by the Megavision camera and contaminate its image).

Capture Method

The RTI control set used a Canon EOS 5D Mark II DSLR camera controlled by a computer. The lenses used were Canon EF100mm f/2.8 Macro USM (objects Mask and Sold), EF24-105mm f/4L IS USM (object Ant2), and TS-E90mm f/2.8 tilt shift lens (objects Pal1 and Pal3). The flash was a Canon model 600EX-RT flash with ST-E3-RT Speedlite transmitter. Several “shiny balls” (reflective hemispheres) were placed in the frame. Although only one shiny ball is used at this point, redundancy allows for backup and for the future possibility of triangulating the light position at the center of the object from the slightly different light positions detected from shiny balls at the periphery of the frame. Another redundancy executed for the sake of comparison was to calculate the light positions on two criteria: the shiny ball method (which is conventional), and measurements in Cartesian space in x - y - z coordinates. First,

Matlab was used to calculate an even distribution of thirty-five points on a hemisphere excluding the horizon and the camera at the apex. A tape measure, masking tape, and markers were used to identify the relevant points on the floor (the x - y plane). For the z -axis we used a plumb with knots and masking tape labels at the appropriate measurements. It quickly became apparent that the measurements based on a tape measure would not surpass the “shiny ball” for precision or ease in creating the light position (lp) file used by the PTM fitter. It was apparent from the pattern on the shiny ball that we successfully created a more even distribution of light points than would have been achieved with less careful planning. Conventional RTI involves using a string to maintain a roughly consistent distance of the flash from the object. With the addition of a reusable floor pattern showing the ideal positions on the x - y plane, the same string could be used to establish the z coordinate with little or no additional labor. The following table indicates the coordinates we used in inches. The z -axis shows height of the light from the object plane, so the measurement had to add the height from the floor to the object plane.

Position Number	x	y	z
1	-11.1	-8.1	37.6
2	-11.1	8.1	37.6
3	4.2	13.1	37.6
4	13.7	0	37.6
5	4.2	-13.1	37.6
6	19.6	-14.3	31.8
7	19.6	14.3	31.8
8	-7.5	23.1	31.8
9	-24.3	0	31.8
10	-7.5	-23.1	31.8
11	-22.2	-16.1	29.1
12	-22.2	16.1	29.1
13	8.5	26.1	29.1
14	27.5	0	29.1
15	8.5	-26.1	29.1
16	22.2	-26.1	20.6
17	31.7	-13.1	20.6
18	31.7	13.1	20.6
19	22.2	26.1	20.6
20	-2.6	34.2	20.6
21	-18	29.2	20.6
22	-33.3	8.1	20.6
23	-33.3	-8.1	20.6
24	-18	-29.2	20.6
25	-2.6	-34.2	20.6
26	-28.9	-21	17.9

27	-28.9	21	17.9
28	11.1	34	17.9
29	35.8	0	17.9
30	11.1	-34	17.9
31	31.8	-23.1	7.5
32	31.8	23.1	7.5
33	-12.1	37.4	7.5
34	-39.3	0	7.5
35	-12.1	-37.4	7.5

Processing Methods

The first stage of processing creates archival sixteen bit per channel TIFF files for each capture from the proprietary Canon raw format (CR2) files. At the same time, eight bit per channel Jpeg compressed images are created, as required by the software for creating RTI images. This stage of processing and similar conversions in image size and resolution can be done in a Photoshop script (as was used for the RTI control set), or similarly with ImageJ (as was used with the experimental methods). The second stage is to create a light position file recording the position of the light for each of the captures. This is done by identifying the shiny ball in LPtracker or the RTI-Builder software, which then detects highlights on the ball for each of the captures and calculates the position of the light accordingly. The third stage is to process the thirty-five compressed jpeg files and light position file into a single file that contains, for each pixel, information not only about its color in RGB, but also the surface norm of the object at that pixel. The two major options for the final processing are the PTM fitter provided by Hewlett Packard and the HSH (hemispherical harmonics) fitter provided by Cultural Heritage Imaging. We also experimented using both LRGB and RGB formats in the PTM. Overall we preferred the LRGB PTM format for quality and efficiency.

Quality Evaluation

The RTI control set succeeded in doing what it was intended to do. It captured and rendered a range of textures, from the large and deep structures of the terracotta figurine to the surface of the palimpsest where ink had been. The limitations pertain to spatial and color resolution. The fine texture of the grain of the animal skin (parchment) of the antiphonary is more apparent in the fifty-megapixel images (with no Bayer filter) than the nineteen-megapixel images (with Bayer filter). The vibrancy of color is somewhat less than the accurate color renderings possible with spectral imaging. The main limitation is the inability to create enhanced images based on data beyond the range and resolution of the color spectrum natural to the human eye. Although it is technically possible to perform PCA on RGB images with some interesting results, there is no expectation of seeing additional information because the captured data does not surpass the ability of the human eye.



Figure 7: RTI, Terracotta Figurine (click image to download PTM file)

Cost in Time and Equipment

The cost in time of conventional RTI consists of three stages: setup, capture, and processing. The equipment is highly portable and can be set up within an hour under the right conditions. During capture the room must be reasonably well darkened. The camera must be stationary and immobile on a tripod or camera stand and works best when triggered remotely by computer, wireless remote, or electronic shutter release. The object must also be immobile and the image frame must contain shiny balls, scale, and typically a color checker. Exposure settings are fixed at the beginning and retained for the entire image sequence. The time of capture is occupied primarily by the positioning of the flash, which is usually held by hand and must be positioned at a roughly consistent distance from the object with an approximately even distribution of angles. A capture sequence of thirty-five or so images can be performed within fifteen minutes (which includes set-up time). Processing time is reasonable because procedures such as white balance, crop, and export to archival formats can be done once for the entire batch of captures. Care is required in identifying the hemisphere for use in detecting highlights for the light position file. It can take longer to experiment with different processing techniques, but a routine procedure of image processing can be done in approximately fifteen minutes. If a second person is working from a separate networked computer, the processing task can generally keep up with the capture task.

The cost in equipment has the potential to vary significantly. Theoretically any camera capable of manual exposure settings, remote triggering, and firing a remote flash can be used. The string and shiny balls add negligible cost. The requisite processing can be done with freely-available software. Gains in quality and convenience add cost from there. Not counting the software and the computer itself, the equipment used in the RTI control procedure (using the highlight method) can be acquired in the range of \$2,000 to \$20,000 (including tripod, DSLR camera body, lens, flash, etc.).

Conclusion: Benefits and Challenges Going Forward

The benefits of RTI for imaging a variety of objects have been demonstrated over the past ten years and more. Target objects can be as small as a coin or as large as can be surrounded by a virtual hemisphere of light with a radius from the object roughly four times the diameter of the object. Successful RTI images have been made of monumental statues as tall as sixteen feet. The main benefit comes when imaging objects with meaningful texture, such as incised inscriptions or cuneiform impressions, but also including texture as fine as the thickness of ink or impression of now-missing ink on parchment. The main limitation addressed in this project is the limitation in the spatial and color resolution of the camera and lighting. There are other limitations and challenges on the horizon. One pertains to the fixed relationship of the camera and object, which generally assumes a roughly planar shape to the overall object. The ability to virtually move not only the light but the object to see a different side entirely has already been accomplished on one axis and holds promise for additional future development. Another challenge on the horizon pertains to the maintenance and development of the PTM fitter. The Hewlett-Packard PTM fitter is limited to thirty-two bit processing and runs out of memory on very large files. The largest PTMs created for this project were 188 megabytes (Ant2 3952 x 5564 pixels). For now this was not a problem because the reference materials were cropped out before fitting the PTM. Going forward the pixel count of a PTM will be limited by the memory addressing capabilities of the PTM fitter, not the capture equipment, unless a solution, update, or workaround is achieved.

Experimental Software Approach (Eureka+Flash)

The first experimental approach adds together the captures required by the spectral imaging and RTI methods and combines the data to produce files with the advantages of each technology. Namely, the digital image has all the spatial and color resolution of spectral imaging, including the pseudocolor renderings, and all the texture information and interactivity of RTI, including specular and other enhancements.

Capture Method

The first experimental approach adds together the eleven captures with the Eureka Lights at different wavelengths of illumination and the thirty-five captures with the flash at different angles. We experimented with using a white LED flashlight or UV flashlight as alternatives to the flash and found the flash superior. The illumination combines both methods, but the camera follows the specifications of the spectral imaging method (monochrome, no Bayer filter, apochromatic lens). The object and camera remain immobile throughout, such that each pixel in the data set represents the same location on the object. Note that it would not be practical to combine the data from different cameras or different capture sessions. The combination of illumination begins at the time of capture with a single camera.

The field of view includes the shiny balls from RTI along with the reference materials used in spectral imaging. The eleven captures with the Eureka Lights are automated as described above for the spectral imaging control method. The subsequent thirty-five captures at different angles of illumination are a little different in that the camera must wait for the operator to position the flash (assuming the flash is hand-held rather than fixed to a dome or arc). At the experimental phase this was done manually by communication between a person operating the flash and a person operating the camera. The flash operator would position and aim the flash and signal “flash ready” to the camera operator. The camera operator would trigger a one-second exposure. The flash operator would listen for the click of the opening shutter and manually trigger the flash. In the future this could be streamlined by having the PhotoShoot software trigger the flash, and more so by fixing thirty-five light sources to a dome, or five light sources to an arc moved seven times, etc.

Processing Methods

The core principle of the software-based integration is the distinction between luminance and chrominance. The special advantage of RTI is in its ability to determine texture from highlights and shadows, properties of luminance. The chrominance capability of RTI is no greater than the RGB DSLR camera and RGB Jpeg files used to create the RTI files. Conversely, the special advantage of spectral imaging is in its ability to distinguish wavelengths of light (colors) more precisely than and beyond the range of the human eye. Assuming the data is fully registered (by using a single immobile camera and apochromatic lens), the luminance and chrominance properties can be combined in software. This can be easily done by using a color space that distinguishes luminance from chrominance, such as YCbCr or LAB, which can be easily converted to RGB for the PTM fitter.

The first stage of processing, which can be done at time of capture, is the creation of archival 16-bit monochrome TIFFs for all forty-six captures and one RGB color image based on the color calibration file for the Eureka Lights. This is done in the PhotoShoot software. The second stage is the creation of a light position (lp) file from the highlights on the shiny ball in the thirty-five captures with the flash. This is done using the RTI-Builder software, as above in the

RTI control method. The third stage of processing is done in the freely-available ImageJ software, with the option of using more advanced proprietary software (such as ENVI) for PCA. The third stage is scripted using the ImageJ Macro Language, and differs for the rendering desired (accurate color, extended color, pseudocolor, and variants). The processes for these renderings are described immediately below. The fourth stage of processing occurs when the PTM fitter uses the light position file to combine thirty-five RGB Jpeg files into an LRGB PTM file. Note that a custom PTM fitter could conceivably provide a shortcut by generating the RGB channels directly from the Eureka Light data and the L channel directly from the flash data. However, this was not found necessary or viable since the PTM fitter source code is not open. Note also that the final stage of processing reduces the data to compressed 8-bit per channel images. This is acceptable because of the limitations of display hardware and the human eye, but it remains important that calculations “upstream” in the processing sequence be done with uncompressed 16-bit or 32-bit floating point data as long as possible.

The first rendering option provided by the ImageJ macro is accurate color. The first step is for the macro to open the accurate color file created by PhotoShoot from the Eureka Lights. The image is converted to the YCbCr color space, and the Y channel (luminance) is discarded. Then each of the thirty-five flash captures are opened and normalized. There are optional variants for normalization that depend on whether the user desires a natural effect of dimming with lower angles of illumination, or an enhanced effect of maximum contrast at all angles. Using a fixed value for brightness adjustment at all angles is more natural, while using a region of interest on the object for normalization at each angle provides greatest contrast. An intermediate option uses the X-Rite Colorchecker target (Macbeth chart) for the range of normalization at each angle. Once the luminance from the flash is adjusted as desired it is assigned to the Y channel and combined with the Cb and Cr channels acquired above. The YCbCr image is then converted to RGB and exported as a Jpeg file for the PTM fitter. This is repeated for all thirty-five flash captures.

The extended spectrum process is similar to the accurate color process just described, but instead of deriving the Cb and Cr channels directly from the accurate color image created by PhotoShoot, this rendering derives those channels from the extended color process described above under Spectral Imaging Control Set. Briefly, that means creating an RGB image in which R is the first principal component of the longest wavelength captures, G is the first principal component of the middle wavelength captures, and B is the first principal component of the shortest wavelength captures. Again, the RGB image is converted to the YCbCr color space and Y is discarded. For each of the thirty-five angles the flash capture is normalized, combined with the chrominance channels from the Eureka Lights, and exported as an RGB Jpeg for the PTM fitter.



Figure 8: Eureka+Flash, Extended Spectrum, Terracotta Figurine

The pseudocolor rendering also takes luminance from the thirty-five flash captures, but derives chrominance data from PCA. Several variants showed potential. The core method maps the second and third principal component from the Eureka Lights to the Cb and Cr chrominance channels directly, which are then combined with the Y luminance channel from the flash. The core method discards the first principal component because it typically approximates the range of luminance while the deviations in the subsequent principal components are analogous to our perception of color. However, for the object Pal1, which is seriously lacking in color if one focuses on the text as the region of interest, the third principal component was of little value so using the first two principal components for the two chrominance channels proved beneficial (the variant code `yccf01` means Y from flash, Cb from first principal component, Cr from second principal component). Conversely, for color-rich objects like the antiphonary it may be desirable to create additional pseudocolor images utilizing higher principal components to bring out contrasts beyond the most prominent (the variant code `yccf34` means Y from flash, Cb from fourth principal component, Cr from fifth principal component). Although customized variants may be desirable based either on human review of the principal components or prediction based on properties of the object, it is worth emphasizing that the default core method consistently produces a generally useful result.




Figure 9: Eureka+Flash,
Pseudocolor, Palimpsest

The principal component analysis (PCA) can be done in a number of ways. For the core objects (Pal1, Mask, and Sold) PCA was done by Roger Easton, Jr. using ENVI software. Also included in the product summary is a PTM of the antiphonary (Ant2) using PCA performed with the ImageJ PCA plugin. ENVI has clear advantages in ease of use, the ability to customize stats (which is more important for the second experimental method below), and the ability to perform Independent Component Analysis (ICA). Images processed using ICA are included in the data archive and showed some potential, but were not generally pursued because of the scope of the project and a general inconsistency in the most useful independent components. That is, one could generally assume that the second and third principal components would be useful without human review, but utilizing independent components would reduce automation and require more human involvement. One of the key decisions involved in PCA is selection of the region of interest. Because PCA finds the greatest contrasts within the region of interest one would generally not want to highlight the contrasts in the color checker or other background materials. Even within the object, selecting a wide region of interest increases the probability of “distractions.” As discussed above under object selection, in the case of the palimpsest the region of interest was clear (the erased text). For objects like the illuminated manuscript it would make a difference which region one wished to study. Consequently, it would be advantageous to integrate PCA directly into the ImageJ macro for purposes of interactive selection and preview. We believe this can be done with proprietary software such as ENVI and Matlab, but for a wider audience there are clear advantages to using and developing free software such as the PCA plugin available for ImageJ or the R statistical computing environment.

Quality Evaluation

The first experimental method succeeded in the core objective of integrating the interactivity and texture information of RTI with the processed derivatives, spatial resolution, and color resolution of spectral imaging. The core premise that luminance and chrominance can be separated and combined is confirmed. The reduction to compressed 8-bit per channel RGB poses some compromise compared to spectral imaging, but it should be noted that display on screen makes such a reduction anyway. The object Sold was selected as a test of the limits of the method at very deep structures such as the edge of the figurine and the deep indentations within the figurine. As with conventional RTI, there are some counter-intuitive shadows, and the PCA pseudocolor highlights indiscriminately both variations in shadow and variations in material reflectance. Thus one should not be fooled into thinking there is a line of paint in the indentation between the conqueror and the falcon or between the conqueror and the conquered. Overall these artifacts are deemed tolerable or even helpful for showing what they show. If one were determined to eliminate this effect the Eureka Lights could be developed to increase diffuseness (they were designed originally for palimpsests). The other, flatter, objects demonstrate the benefits of the integration without reservation. The palimpsest remains difficult to read in parts, but the integration makes it possible to read some letters by virtue of texture and other letters by virtue of traces of ink remaining, even in the same word. Thus the integration is greater than the sum of the two parts side-by-side. The pseudocolor of the mummy mask fully shows the texture of the wood and paint, and also the otherwise subtle space at the edge of the jaw between the paint of the cheek and the paint of the background. In the reasonably flat portion of the figurine both the thickness and the color of traces of paint are easily visible. The antiphonary clearly shows the grain of the “hair” side of the parchment. Overall the Eureka + flash method produced the most consistent results, and can be highly recommended for imaging objects that conform to the criteria discussed above under Selection of Objects (meaningful color and texture information, range of size and depth).

Cost in Time and Equipment

The cost in time of the software-based integration is essentially the sum of the cost in time of the two control procedures. At the time of capture the positioning of the object and reference materials and focus of the camera occurs once. The time to capture the eleven images with the Eureka Lights is added to the time to capture the thirty-five images with the flash, which consists mostly of the time positioning the flash (absent a fixed structure). We were able to run through the capture sequence within twenty minutes (when we did not simultaneously test the second experimental procedure). Improvements in efficiency are easily imaginable. The processing time is similarly the sum of the processing required by the two established methods. As with the spectral imaging control set, one can try as many or as few variants as one desires. The default method reliably produces useful results. The way the ImageJ macro is currently written the processing is heavily automated but assumes uniformity in input filenames, such that some curation of filenames is required. This could be easily adapted to involve more interaction in the processing script and less preparation.

The cost in equipment is essentially the same as the cost of the spectral imaging method with some minor additions. The electronic flash and radio trigger connected to the control computer could add a few hundred dollars. The shiny balls and positioning materials add no significant cost. The increase in the quantity of data captured will add to storage costs. Each additional capture at fifty megapixels and sixteen bits per pixel adds about 100 megabytes, which

can be doubled if one stores both the original DNG file and the archival TIFF. Quadrupling the storage for captured data compared to a technique developed ten years ago is easily offset by the decreases in storage costs over that time.

Conclusion: Benefits and Challenges Going Forward

The software-based approach is a time-efficient and quality-effective means to integrate the interactivity and texture information of RTI with the processed derivatives, spatial resolution, and color resolution of spectral imaging. The limitations are the same as the integrated components. The approach is not optimal for every object, particularly not objects that are very large, very dimensional, or homogenous in texture or color. One is still constrained by the cost of the spectral imaging equipment and the need for a special (but free) viewer for the PTM files. Progress in the separate technologies will also benefit the integration. Growth and improvements can be expected as the macro script is adjusted to apply to more objects and the number of possible variants is reduced to a smaller number of recommended variants. Scholars and conservationists currently using spectral imaging equipment should be able to justify the extra time to add RTI functionality. Scholars and conservationists currently using RTI equipment will face a greater concern in the higher equipment costs. Color accuracy may be an incremental improvement over consumer electronics, but the ability to create extended spectrum and PCA pseudocolor images will be valuable for the study of objects beyond what is readily apparent to the human eye (or to cameras designed to mimic the capability of the human eye).

Experimental Hardware Approach (Magic Flashlight)

The second experimental approach multiplies the captures required by the spectral imaging and RTI methods ($9 \times 35 = 315$). This involves far more captures than the first experimental approach which added the captures ($11 + 35 = 46$). (For practical reasons only nine bands were used for the second experimental approach, but the principle was the same as the eleven bands of the Eureka Lights.) This approach does in hardware the combination that the first experimental approach did in software. That is, for each angle of illumination nine narrow bands of illumination are captured such that accurate color, extended spectrum, and pseudocolor can be generated at each angle without reference to other angles. In some ways this was the lower-risk approach because it did not rely on the premise that luminance and chrominance could be separated and “grafted” back together. It is also by far the most time- and equipment-intensive approach. In some situations, such as accurate color, the hardware approach produces marginal benefits, but not sufficient to justify the increased effort and expense. In other situations, particularly pseudocolor of chromatically complex objects, the vast data creates problems for PTM rendering that have not been adequately resolved. We can imagine refinements in capture and processing, and more importantly ways other than RTI of rendering the vast quantity of registered data. However, for the task at hand of integrating spectral and RTI, it is precisely the success of the simpler software approach that makes it unlikely that the more complex hardware approach will achieve such improvements in quality as to justify the additional cost in time and equipment.

Capture Method

The main change in capture method for the second experimental approach was the creation of a “Magic Flashlight” that shines nine narrow bands of light from ultraviolet through

infrared. This prototype was built by Christens-Barry. The Magic Flashlight was designed to be controllable by the PhotoShoot software, like the Eureka Lights, such that the sequence of nine captures at each position can be automated. (For purposes of executing both experiments at once, a tenth and eleventh capture used a white LED flashlight and xenon flash for the first experimental procedure.) Whereas the flash fires so quickly that it can be held by hand, the Magic Flashlight needs to be stationary for nearly a minute for the sequence of nine captures. This required a rolling, adjustable boom to hold the Magic Flashlight immobile at each position. The Magic Flashlight also had to be manually aimed at each position so that the object and essential reference materials fell within the beam of the Magic Flashlight, which was considerably more narrow and weaker than the flash. The x - y plane on the floor and the plumb-bob on the z -axis guided the general position of the Magic Flashlight, but the precise orientation of the Magic Flashlight was not so consistent as to permit flattening. A dome or other fixed light mount might have addressed this problem although flattening still would not have been effective on objects of deep texture like the terracotta figurine.

In other ways the capture sequence resembled the procedures above. Once the light was positioned and aimed, the PhotoShoot software controlled the camera and turned the narrow-band illuminators on and off. The field of view included shiny balls, color checker charts, rulers, and spectralon for reference materials.

Processing Methods

The four stages of processing were the same as in the first experimental approach described above, but with some key differences in the third stage. The four stages again are: creation of archival TIFFs and accurate color images in the PhotoShoot software, creation of a light position file from the highlights in the shiny ball at each of the thirty-five positions, the ImageJ macro, and the PTM fitter. The only difference at the first stage is that, instead of creating one accurate color image, thirty-five color images were created separately using the same procedure, including setting the white balance at each position. The ability to create white balance at each position, not just adjust the normalization of the luminance, likely contributed to the advantage of the Magic Flashlight in the accurate color renderings. One lesson may be to pay more attention to white balance at the various angles in the software-based approach. The only difference at the second stage is that the light position was determined based on one point on the Magic Flashlight (the 507 nanometer light). The UV and IR modules were slightly offset from the visible light portion of the Magic Flashlight, but this is not believed to have introduced any perceptible error. The only difference in the fourth stage was that HSH and RGB PTM fittings were tried more exhaustively, but unsuccessfully, in search of a solution to the problems encountered with the pseudocolor PTMs, discussed below. The third stage mimicked the renderings described above but with the following differences.

For creation of the accurate color PTM the third stage consisted merely of creating eight bit per channel RGB Jpeg images from the higher quality color images produced by the PhotoShoot software.

The extended spectrum rendering was executed in the same way as described above without the luminance-chrominance grafting in the YCbCr color space. At each angle the process created one RGB file from nine narrow-band captures. The first principal component of the three captures with the longest wavelength was mapped to the R channel in an RGB image, and similarly with the middle three captures to G and the shortest three captures to B. This processing did require some time because 315 captures at 100 MB each had to be read from the file system (a Network Attached Storage device in this case) and PCA was run in ImageJ 105 times. The process took something like ten minutes, but results could vary with data transfer and processor speed.

The PCA processing was done two ways in anticipation of a problem that inhibited RTI rendering of PCA pseudocolor images. The problem was that the pseudocolor images at different angles did not resemble each other in coloration. This in turn was a result of principal components not resembling each other at different angles. We anticipated that performing PCA separately at each of the thirty-five positions separately could create problems when combining the thirty-five pseudocolor images in the PTM fitter, but we tried it anyway and called this the “local stats” method because the stats used to create the principal components were calculated at each local angle individually. Using ENVI software, Easton was also able to perform PCA on all the angles with a single standard set of stats, called “global stats.” Global stats were calculated from the five positions nearest the camera because those images contained fewest shadows and least gradient as the illumination decreased across the field of view with increased distance from the light source. Processing all the positions with a single set of stats helped, but did not eliminate the problem that a pseudocolor of a pixel could look green from one position and red from a proximate position. When the divergent colorations were combined into a PTM the color washed out. Several variations were attempted. One called “lightText” examined two sample pixels in each image, one representing text and one representing background, and inverted any images for which the text pixel was darker than the background pixel. Another variant mapped the first three principal components to the R, G, and B channels rather than the Y, Cb, and Cr channels (the code rgb210 means RGB mapped to the third, second, and first principal component). The results are discussed below under Quality Evaluation.

PCA was also inhibited by the lack of flattening and the generally predictable pattern that at steep angles the light produced a gradient with the brightest points near the light and the darker points further from the light (the one-over- r -squared problem). PCA easily found the contrast in the shape of the illumination and missed the very subtle contrast of erased ink. One theoretical solution that was not implemented would be to generate artificial flats for each light point based on the




Figure 10: Magic Flashlight, Extended Spectrum, Terracotta Figurine




Figure 11: Magic Flashlight, Pseudocolor, Palimpsest, Component Image

geometry of the illumination. This would only work if the entire region of interest did not intersect the edge of the beam of illumination. A simpler solution was to perform PCA on smaller regions of interest, which could be pieced together in a mosaic or strung together in a panning video. Increases in effort such as these produced modest gains in functionality without approaching the simple benefits of the first experimental method.

Quality Evaluation

The Magic Flashlight demonstrated a slight advantage in accurate color renderings. Since white balance was performed at each angle the steeply raking images appear brighter and therefore more pleasing to the eye. Greater attention, even manual adjustment, to the images in the software-based approach could close this gap and still require less effort than the hardware-based approach.

The extended spectrum renderings were close between the Eureka Lights and Magic Flashlight, with the Eureka Lights the better of the two. This is most likely explained by the fact that the Eureka Lights have a shorter UV and a longer IR wavelength capability compared to the Magic Flashlight. This gap could presumably be closed to bring the two to quality parity, but not to bring the Magic Flashlight to such an advantage as to justify the extra expense in time and equipment, below.

The pseudocolor images produced from the Magic Flashlight were individually stunning but washed out when combined in the PTM (or HSH) fitter. The ghostly washed-out effect could be aesthetically pleasing compared to the deliberately stark contrasts in the other pseudocolor PTMs. In the case of Pall for which the only contrast of interest was the contrast between where ink had been and parchment where it had never been, the wash-out effect did not hinder the core objective. In fact, the interactivity of the RTI viewer brought unintended benefits in that moving the light position had radical effects on the overall contrasts. In this case a slight quality advantage could be granted to the Magic Flashlight. For the other images of more complex chrominance the washed out PTMs offered at best an aesthetic distortion (the visual equivalent of euphonic distortion), but did not aid the study of detail in the imaged object.

However, the general failure of the Magic Flashlight to produce superior (or sufficiently superior as to justify the effort) PTMs is not the last word on the basic concept. The problem is the combination in the PTM fitter, not the quality of the individual images. For example, 20140405-Mask-324-08-Pmpl-03-50m.jpg (Figure 12) may be the most stunning and interesting image from the project. This single image compares favorably to some of the interactive images for highlighting features of the object. The present project has a simple objective of integrating spectral and RTI, but also points to the possibility of alternatives to RTI. If the Magic Flashlight were used without RTI it would no longer be necessary to capture a




Figure 12: Magic Flashlight, Pseudocolor, Mask, Component Image

complete set of thirty-five sequences. One could use the software-based approach as an initial stage to identify useful angles of illumination and then go back with the Magic Flashlight to capture a small number of angular spectral sequences.

Cost in Time and Equipment

The Magic Flashlight approach costs considerably more in time than the Eureka + flash approach. Using a boom to hold the Magic Flashlight it took us nearly four hours to complete a sequence (although a small portion of that was for the white LED and flash used for the other experimental approach). It simply takes longer to position a boom and aim a narrow beam of light than it does to hand-hold a flash. Capture time could be significantly improved by using a dome or a movable arc especially if multiple Magic Flashlights were permanently fixed in absolute position on the dome or relative position on a movable arc. It is conceivable that the Magic Flashlight could be developed past the prototype stage to cost less than a full bank of Eureka Lights, but such a number multiplied by thirty-five and added to the cost of a dome would seriously limit the accessibility of the technology out of the range of most museums and institutions. Even if the entire sequence of 315 captures were fully automated under a single complete dome, five seconds per capture would fill more than twenty-six minutes. The data to be stored and transferred is almost seven times greater than the software-based approach at the capture stage and thirty-five times greater at the PCA stage. It should be recalled that amounts of data that seem astronomical now may seem trivial in a few years. Approaches similar to the Magic Flashlight may well deliver significant success in the conceivable future, particularly for very special projects and media other than PTM. For the present time and simple objective of integrating spectral and RTI, the cost of the hardware-based approach in time and equipment greatly eclipses its benefit, particularly in comparison to the software-based approach.

Conclusion: Benefits and Challenges Going Forward

The benefits of the Magic Flashlight for creating PTMs is limited to a small margin of improvement for creating accurate color interactive images and little or no improvement for creating extended spectrum renderings. For all but the most special of circumstances these benefits do not outweigh the costs. The large data set does offer potential for additional processing approaches. Challenges ahead for anyone seeking to create PCA pseudocolor PTMs include flattening with captured or artificial flats and adjusting the analysis to produce more consistent principal components.

The Magic Flashlight may prove more useful for creating still images with directional spectral illumination, but not a complete set of images processed with a PTM fitter. The large data set could also be mined in other ways. One experiment ran PCA on the complete set of 315 captures for a small area and mapped the first three principal components to a YCbCr pseudocolor. The results were good but not a clear improvement over simpler methods. However, more thinking along these lines may produce better results. Our data is freely available to anyone who wishes to try other processing approaches.

Summary of General Conclusions

The following points summarize the major findings of the project.

- Spectral imaging and Reflectance Transformation Imaging can be efficiently and effectively integrated by grafting the chrominance data from spectral imaging onto the luminance data from RTI.
- Extended spectrum processing effectively visualizes the range of data captured by spectral imaging while maintaining resemblance to the human perception of color.
- The YCbCr color space is an effective way of creating pseudocolor images not only for integrating spectral and RTI, but for other applications, especially when the first principal component approximates the range of luminance and the subsequent principal components approximate other variations in composition.
- The prototype Magic Flashlight, even with conceivable upgrades in spectral range and a dome, is not an efficient way of creating RTI images. However, individual images and other means of rendering the massive data set warrant further exploration.

Appendix 1: Access to Additional Information

- The Integrating Spectral and Reflectance Transformation Imaging project is a phase of the Jubilees Palimpsest Project. Information about the project and images produced by the Integrating phase can be found on the project website: <http://palimpsest.stmarytx.edu>.
- The major PTM images of the project are available through the InscriptiFact Digital Image Library, <http://www.inscriptifact.com> (free registration required).
- The complete archive of all the data captured and processed for the project is stored on magnetic disk (external drive with USB interface) and is expected to be publicly available from special collections at the libraries of St. Mary's University in San Antonio and USC in Los Angeles.
- The project director, Todd Hanneken, can be contacted at thanneken@stmarytx.edu or thanneke@uchicago.edu.

Appendix 2: List of Participants

Todd Hanneken, St. Mary's University – project director, data management, ImageJ processing
 Mike Phelps, Early Manuscripts Electronic Library – sponsoring institution and spectral imaging team lead

Bill Christens-Barry, Equipose Imaging – Illumination, design and construction of Eureka Lights and Magic Flashlight prototype

Ken Boydson, MegaVision – spectral imaging camera and PhotoShoot software operator

Roger Easton, Jr., Rochester Institute of Technology – image processing in ENVI (PCA and ICA)

Bruce Zuckerman, University of Southern California (USC), West Semitic Research (WSR) and InscriptiFact projects – imaging site host and RTI team lead

Ken Zuckerman, WSR – RTI image capture

Marilyn Lundberg, WSR – RTI processing, cataloguing

Leta Hunt, USC, InscriptiFact – InscriptiFact images dissemination, project administration

Matt Klassen, St. Mary's University – student researcher, project scribe, Magic Flashlight operator

Funded by the United States National Endowment for the Humanities, Office of Digital Humanities

Appendix 3: List of Figures

Figure 1: UCLA Palimpsest

<http://palimpsest.stmarytx.edu/integrating/pal1/images/20140218-Pal1-00-Accc-03-50m.jpg>

Figure 2: USC Mummy Mask

<http://palimpsest.stmarytx.edu/integrating/mask/images/20140217-Mask-00-Accc-03-50m.jpg>

Figure 3: USC Terracotta Figurine of Conqueror and Conquered Soldier

<http://palimpsest.stmarytx.edu/integrating/sold/images/20140220-Sold-00-Accc-03-50m.jpg>

Figure 4: USC Illuminated Antiphonary Manuscript, Folio 2 Recto

<http://palimpsest.stmarytx.edu/integrating/ant2/images/20140224-Ant2-00-Accc-03-50m.jpg>

Figure 5: Spectral Imaging Extended Spectrum, Terracotta Figurine

<http://palimpsest.stmarytx.edu/integrating/sold/images/20140427-Sold-045-00-Xepn-03-50m-n.jpg>

Figure 6: Spectral Imaging Control Pseudocolor, Palimpsest

http://palimpsest.stmarytx.edu/integrating/pal1/images/20130725_PAL1-00-Capt-01-50m_WCB_N_PSE-01-15-15.jpg

Figure 7: RTI, Terracotta Figurine

http://palimpsest.stmarytx.edu/integrating/sold/ptms/USC10663_Smiting_Warrior_cropped_2690.ptm

Figure 8: Eureka + Flash, Extended Spectrum, Terracotta Figurine

<http://palimpsest.stmarytx.edu/integrating/sold/images/20140408-Sold-045-29-Xepn-03-50m-n.jpg>

Figure 9: Eureka + Flash, Pseudocolor, Palimpsest

<http://palimpsest.stmarytx.edu/integrating/pal1/images/20140415-Pal1-045-32-Pepl-03-50m-yccf01.jpg>

Figure 10: Magic Flashlight, Extended Spectrum, Terracotta Figurine

<http://palimpsest.stmarytx.edu/integrating/sold/images/20140408-Sold-045-29-Xepn-03-50m-n.jpg>

Figure 11: Magic Flashlight, Pseudocolor, Palimpsest, Component Image

<http://palimpsest.stmarytx.edu/integrating/pal1/images/20140414-Pal1-324-20-Pmpg-03-50m-rgb210.jpg>

Figure 12: Magic Flashlight, Pseudocolor, Mask, Component Image

<http://palimpsest.stmarytx.edu/integrating/mask/images/20140405-Mask-324-08-Pmpl-03-50m.jpg>