Specifics of Imaging Practice

Michael Ester
Luna Imaging, Inc.

As the individual titles make evident, this conference reflects the use of images drawn from a composite of interests and backgrounds. And there is certainly much going on in imaging, whether it is yet another site on the Internet showing pictures, more institutions exploring how to move their visual collections into electronic form, or the crafting of new analytic and interpretative applications. As both the moderator and a speaker, what I believed we could add to these fast moving developments is some reflection on our practical experience to date. The rationale is that every issue or problem we can surface is one that doesn't have to be faced totally afresh by our audience. Where we have developed good solutions, great. But of equal importance is identifying unanticipated situations we have encountered and sharing concepts or frameworks to better understand them. Sometimes the working answers are untidy; the practical moral may be how to compensate for or live with present limitations.

In previous papers on electronic imaging and the arts,[1] I have covered such topics as viewer perception, definitions and distinctions for archival and derivative quality images, descriptive documentation, and the many issues of electronic publications and work environments for research and education. Some of these are theoretical topics, but generally I have favored considerations that are well tethered to practical decision-making in museums and archives. In this paper, I would like to bear down even more on specific practices that in our experience have an important bearing on the resulting image resource and on the future value of visual collections converted into digital form. Taken individually, some of these practices — such as photographic problems, color matching and transformations — seem like minor points, and often are. They are rarely the beacons of technological excitement. In the aggregate, however, they can make an enormous difference.

In covering this terrain, I believe the company I represent, Luna Imaging, brings an interesting perspective. On one side of the house, we provide imaging services for the cultural heritage community. We convert visual collections into digital form and supply systems for using images and integrating

1 Ester, Michael, 'Image Quality and Viewer Perception', *Leonardo: Journal of the International Society for the Arts, Sciences and Technology*, Supplemental Issue 1990, pp. 51-63.; Ester, Michael, 'Digital Images in the Context of Visual Collections and Scholarship', *Visual Resources*, Vol. X, 1990, pp. 11-24.; Ester, Michael, 'Issues in the Use of Electronic Images for Scholarship in the Arts and the Humanities', *Networking in the Humanities*, ed. Stephanie Kenna and Seamus Ross, London: Bowker Saur, 1995, pp. 111-125.

them with text information. From the other side of the house, it is accurate to say that we are our own most demanding client: we use our services to produce digital archives and derivative content for our own line of electronic publications. In short, we have had an on-the-ground look at more, different kinds of production situations than any single organization is apt to encounter, and yet we also use and depend upon the images we create.

The four areas where I would like to share some of what we have learned are photography, color control for collections of visual material, color transformation and image output, and management documentation.

Photography

There is an abiding expectation that sometime in the future, direct digital capture from cultural heritage materials will be a commonplace reproduction method. Even today, a persuasive case can be made for digital capture rivaling and in some cases surpassing 35mm photography. Yet for general reproduction purposes, as opposed to special analytic purposes, direct digital capture plays a very modest role in what we do compared to the vast majority of scanning from film and print sources. Some of the reasons why photographic media will probably predominate over the medium term future are:

- Film remains the transaction medium of professional use for institutions, printers and publishers. Publishers rarely want digital images as a visual source, even though the first thing they may do is scan incoming film. At the moment it is more effective to send off large format film, than it is to wrestle with correspondingly large files and import them correctly.

- Film is a well understood and familiar medium. At the extreme, several granting institutions *require* a film intermediary for digital imaging projects. While this may seem an overly cautious position, funding guidelines are intended to serve the conservative objectives of preserving intellectual access to cultural materials.

- Perhaps most persuasively, there exists a huge base of photographic collections that constitutes a very real and central concern in the humanities. Comprising an assortment of slides, transparencies, and prints in museums and universities, these holdings number into the tens and hundreds of thousands of images. Beyond serving a variety of research, educational, and management needs, visual collections represent major capital investments and operating commitments.

It follows, then, that for much of what we do, the digital image will only be as good as its photographic source; if visual detail or subtlety is not in the photographic medium, neither will it appear in the digital image. Where institutions are preparing original photography, what are the considerations we have

found important in anticipation of digital conversion? Several points simply appeal to good studio practice; others are more closely tied to the digital environment.

Lighting. Very even, medium contrast lighting works best. High contrast lighting, frequently used for dramatic effect, can produce characteristics that become more evident in digital form: for example, bright colors can flare or reflect from light backgrounds, UV and strobe color casts are more likely and more pronounced. High contrast photography also leaves little latitude for error: there may be a trade off between achieving accurate color control and losing already marginal highlight or shadow detail.

Alignment. It is remarkable how many objects are photographed slightly askew. Even in obvious examples, it seems that our eyes are quite forgiving and we mentally ignore or adjust for mis-alignments. Scanning can be performed exactly square to the photographic medium, and a crooked photograph is more noticeable, particularly when multiple images are displayed at once. Etched grid lenses, levels, placement markings, and simply greater attention to this factor are some of the ways to easily manage alignment.

Color control bars. Though less practical for outdoor photography, color control bars should be considered essential for studio work. As discussed more fully below, systematic inclusion of color bars with an object can reap major rewards moving into the digital environment. Improvements in fidelity, efficiency, and cost are possible if these standard devices are used. These advantages can be projected across entire collections if properly applied. Optimally, color bars should be placed in the identical position for every frame, allowing automatic processing of images in electronic form. Standard color bars are readily available; reasonable care should be taken that color bars have not faded over time and are otherwise in good repair.

Gray-scale control bars. Often we find that institutions have included color bars but left out gray-scale bars. When asked why, the answer is usually that the gray-scale bars were considered superfluous with the use of color film. This is not so. Correcting a digital image involves control in three continuous areas: shadows, mid-tones, and highlights. The gray-scale gradient is crucial to this process.

Maximize useful image real estate. In museum photography it is common to see considerable extraneous content toward the outside edges of a photograph. This may include stands, backdrops, or the ambient surroundings. But of course, these parts of the picture can be cropped out later before use. While the situation is similar in the digital medium, the allocation of available pixels to represent a picture draws attention to the premium on area. Ideally, every pixel of an image would record something important. When an image contains considerable extraneous content, pixels are assigned (or cropped out) that could have contributed to the photographic subject. Giving up a few pixels is hardly worth attention. However, the border effect can be deceptive in how much area is actually wasted. For example it is only necessary to give up a quarter inch in either dimension of a 4 x 5 transparency to loose 5% of the area; similarly reducing the outside area a quarter inch on all four sides, reduces the useful area by

25%. Ironically, we have seen the unconsidered placement of color bars — an effort to maintain image quality — take away a third or more of the useful image area. In framing a subject, selecting an orientation, and picking the size and placement of color bars, it is worth considering how such photographic decisions will affect useful image area.

Film processing. Film processing may produce image artifacts that are less noticeable in the photographic original. Extreme chemical streaking and mechanical roller marks have appeared in client collections.

Avoiding darkroom magic. We encountered this issue the hard way. In correcting a group of images based on color bars, we found that we were making certain parts of the image look wrong. Repeating and checking the previous work yielded the same results. After considerable investigation, the cause was this: in the darkroom the photographer had selectively altered sections of the image, and thereby rendered useless the relationship between the control bars taken in the photograph and the film that we scanned. It should be recognized that darkroom modifications of this sort, while desirable for print objectives, undermine the ability for global corrections to images after scanning.

Avoiding duplicate film. We have come to presume very little or no loss in converting film media into digital form. Beginning a project based on second generation film as a source was a vivid reminder of the significant quality degradation resulting from photographic duplication. This difference will, of course, be reflected in digital capture. Especially for long-term archival use, second generation film should be used with considerable reserve. At the least, it is important to know the status of the source; in some cases institutions are unaware that there are duplicate reproductions in their film collections.

Film stock. What would be the characteristics of a film if its major purpose was as a source for digital capture? Although there is some ongoing research in the film industry, simply pursuing this question with commercially available film has proven valuable. Slower, fine-grained film is greatly preferable for studio work, although this is true whether the intended purpose is digital or not. Perhaps the greatest improvement we have witnessed is the use of negative film. Normally, negative film, especially for color photography, would be an unlikely choice because it is not immediately human-readable. This limitation falls away in conversion to digital form. The notable advantage is a smoother curve in the dynamic range of negative film, yielding additional highlight and shadow detail. Our experience is still too limited to know whether equally successful results will hold for different mixes of image content, film formats, and scanning.

To be sure, some of the above problems can be partially corrected through image processing. Slight adjustments in alignment, for example, can be done quite expediently. Other problems, such as proper color controls, have no easy remedy. Even where certain types of corrections are feasible, in general, there can be no assurance of satisfactory improvement for a particular image without introducing artificial information.

Color Matching for Image Collections

Scanning reproductions and working with images in electronic form has been generally available in the pre-press world since the early 1980's. Highly trained technicians typically hand-tune each scan individually, color-correcting the image by eye to meet specific publication objectives. Many of these systems and practices, though now often scaled down to desktop technology, have been adopted for use in creating image databases. What I want to do in this section is discuss how the construction of image archives as *collections* requires very different thinking from a pre-press mentality. Certain conventional approaches can undermine profoundly the prospect of attaining long-term visual resources. At the same time there are essential steps that must be taken if the digital conversion efforts of institutions are expected to endure for the museum and archive community.

The crucial difference between pre-press and creating digital archives turns on two points. First, is the issue of longevity. In the commercial print world the purpose of the image is immediate: it is a production step on the way to a publication. Images may be retained and there is a good deal of talk about "re-purposing" digital material, but the image is a byproduct of a subsequent objective. Moreover, "color-correction" has more to do with the intended visual impact, paper stock, and printing process of a publication than it has anything to do with a concept of image fidelity to a source. By contrast, the implied value of an image archive is directly related to its ability to last. In the context of a technological environment, this must also mean that images can move from system to system and look the same across many display and output devices.

The second point is the issue of contending with many images at once; image collections must behave as a coherent resource. It should also be assumed that image collections will accumulate in segments or phases within institutions, and that there will be a premium on accessing and combining multiple collections.

On both these counts — having image collections behave in unison over the long haul of technological change — the practice of arbitrarily adjusting individual images is highly undesirable. Visually hand-correcting images to the content of a photograph makes each one unique; even a trained operator cannot resolve color, brightness, and contrast combinations for the contents of a photograph in exactly the same way over series of images. Organizations that are storing away their pre-press images from different publication projects, are not accruing the long-term asset they imagine.

How can photographic collections be brought into the digital environment consistently? I will mention two basic rationales for matching a digital image to a reproduction source: "matching to film" and "matching to the scene". There are additional factors which determine how successful and complete the match will be.

Matching to film. Under this logic the goal is to make the digital image look the same as the photographic medium. In favorable circumstances, with color control bars in the photograph, one can take densitometer readings from the film, and knowing their digital counterparts, use these values as references for the scanned image. This process can be applied very consistently to achieve a record of the film. Where color bars are not present there are less precise ways to pursue the same objective by working with the bright and dark points of the film. Without color controls some operator judgment is involved, but the problem is narrowed to matching film values rather than visually correcting for the original objects in the photograph.

Matching to the scene. The logic of this process says that if we have a color bar in a photograph and the digital values of the target are known, then as the color bars in a scanned image are adjusted to their correct values the rest of the image is matched to these controls as well. We call this matching the "scene", because it represents the photographic scene at the moment the photograph was taken. In this case, the correspondence is to the actual objects in the photograph, where one constant object — the color bar — provides the reference.This can be taken a step further if a gray scale bar is included, so that corresponding highlight, mid-tone, and shadow variation also can be corrected.

To contrast the two methods, matching to the scene corresponds to what the camera lens saw when the photograph was taken, while matching to the film corresponds to what the processed film contains. There can be substantial differences between the two results.

Bringing this summary down to familiar ground, the practice and relationship between these two methods emerged from our work in scanning more than 5,000, 4 x 5 transparencies of drawings from the Frank Lloyd Wright Archive. We began the project from the perspective of matching to the film. Bringing in this mass of photographic material as a consistent and accurate digital version of the Archive's film seemed a remarkable goal at the time. This strategy was applied across all of Wright's presentation and conceptual drawings.

However, we then began to think that it might be possible to go beyond the film, and if we could match the Archive's color bars in the image, we could potentially represent what the camera lens saw rather than what the processed film recorded. We created film output from the resulting digital images and sent them to the Frank Lloyd Wright Archive. The response we received back from them was that the digital images we produced were closer to the actual drawings than the Archive's photography. Only through comparison of the images did some of the problems become evident. We see the same characteristics in nearly every project: the Archive's film ranged over twelve years, there were film processing differences, and there were marked differences in film batches.

I should mention that not every image could be fully processed to match the scene. Several fine line drawings had been photographed with slightly overexposed areas. To have corrected to the color bars would have washed out detail, and so we had to back off ideal values in these instances.

It is also interesting to consider the status of images from the two methods. In the course of the Wright project, we concluded that the images matched to the film should always be produced and kept as the archival referent. The thinking was that since the film was the physical source from which digital images were derived, there ought to be a digital version of the drawings that represented this literal association. While the reasoning is defensible and seemed a conservative touchstone for an initial project, we no longer subscribe to this position. As we looked at the practical implications in project after project, images matched to film only enshrined the intervening photographic problems between object and image.

Color Management, Transformation and Image Output

If such production steps as color matching present a certain level of complexity in themselves, the prospect of obtaining a consistent collection of images is greatly compounded in difficulty by the many potential junctures for losing control of how images look. Anytime an image is captured, changed, or transferred to another system, medium, or device, the outcome may vary from the intended result. The industry is presently at the stage where no transition can be taken for granted. Even the seemingly benign act of moving an image file in the same format from one computer platform to another may trigger selective color shifts. In this section, I outline some of the key areas to watch and propose some of the measures for coping with the still unsteady world of digital images.

In our work at Luna Imaging there are several sources of image input. We use Kodak scanners for 35mm and 4 x 5 film; we use several other units for reflective media (such as prints) and for large and odd sized reproduction materials. We also receive image files from clients as the basis for additional processing. One decided source of variation then is the capture device that produces the digital image.

I have discussed elsewhere (see also Ester 1994), how image quality is determined by a mix of resolution, dynamic range, and the process of capture. This latter consideration is actually a short-hand for an array of factors including the sensing device, conversion from analog to digital, and structuring of digital information into an electronic image. For the moment, I only want to consider one dimension of image quality, the impact of scanning systems on consistency. Although the battles waged in promotional literature are typically confined to the numbers game of resolution and density, scanning systems can also be judged on a suite of other characteristics: precision, accuracy, linearity, flare, and optical properties. What is important here is that images will vary between scanning systems, and moreover, how they vary also will be conditioned by specific image characteristics.

Scanning systems are not the only source of variation. Translating into and out of a color space is the other big factor. Color spaces that are familiar to most of you include RGB, CYMK, and Photo CD. Although it would be tempting to control this variation by committing to one of these color schemes, even this tactic will not suffice. Move an image from one system to another or from one software program to another, and you are very likely to have gone through an inadvertent conversion as one vendor's representation and processing of a color space varies from the next. And the thorny part of this problem is that color transformations are by no means simple mechanics. Although the basic equations are well documented, they are non-linear systems, and the ability to successfully generate real-world transformations is the stuff of proprietary experience and color science. As an added twist, there are both computer platforms and software products which optimize performance by constraining color values.

Occasionally, the difficulty of color transformation erupts into spectacular failure, where particular areas of the color gamut go to some anomalous value. These glaring events also reveal another point: the effects of variation are cumulative. Multiple passes into and out of image environments can dramatically amplify small deviations. Blatant examples are hard to miss; the real interest is in how to cope with color control in the more subtle and commonplace situations we do not readily see.

Someday, moving images around reliably in production will be transparent. But in the meantime, the linchpins to managing images are the abilities to detect and measure variation, and to map images into a uniform color space that serves as the clearinghouse for processing image collections. The regimen we follow that works well in our situation can be described briefly as follows:

> We move images into and out of the CIE color space, the international standard for defining color and the widest color gamut available. Common image representations are Uvl or CIELAB formats.
>
> Input sources are characterized in terms of CIE by inputting standard targets, Kodak's Q60 color target in our case. Software measures the deviation of the target from the known digital values and develops a corresponding transformation. Images from the input source can be translated into CIE on this basis.
>
> During production the input source must remain in calibration, which can be determined by independent testing of the associated input device or by reconfirming the characterization process above.
>
> As a safeguard, Luna associates a Q60 image with every production batch; a batch may represent the number of images stored on a particular medium such as CD-ROM, or it may be an arbitrarily fixed number of images that are processed as a unit. Not only is this a standing check on input calibration, but the Q60 can be used to reference any

subsequent change to the group of images, whether this is a processing step or transfer to a new computer platform.

In shaking out a workflow process it is also useful to perform a round-trip check: starting with photographic media; capturing to digital media and processing the image; and outputting back to film or print stock for review. Again, a standard photographic target, such as Kodak's Q60 is useful for this purpose. Similarly, in assessing transfer of digital images across multiple systems, measurement of digital values at each juncture can identify and isolate problem areas. It is even worth running test images through hardware and software transfers multiple times: subtle errors may become additive or self-amplifying over several iterations.

In principle, mapping image output — for monitor displays, hardcopy devices, and other storage formats — involves the same process of characterization and calibration. Exporting to other color spaces and systems is very much the same; monitors and printers are harder to control for metric fidelity. Given our business, it is worth the investment in labor and equipment to maintain exacting standards of image quality and consistency for the institutions we work with and for the contents of our publications. Yet even so, we consider it no small task to ensure that images remain the same as they go out to the variety of computers around the office. But it can be said of imaging technology today, that if there is sufficient incentive to control color and retain fidelity, it can be done.

Controlling Images in User and Distribution Environments

Within our production facilities and the archival systems that support institutional applications, it is realistic to maintain electronic images that are true to the digital values of original capture. It is an extremely crucial point to make, that, unless this level of control is achieved *from the beginning* for entire collections of digital images, there is little prospect of achieving image resources that will survive technological change. But very much along the lines of the distinction between archival and derivative images I have made elsewhere (see also Ester 1990, 1995), the image standards that underwrite the long-term value of an institution's collections are unlikely, for the time being, to find their way into everyday distribution and user applications. As CD's are inserted into home and office systems or as images are spun out to WEB sites, strict control of digital content and viewing conditions are rarely a sufficient priority to prompt the extra trouble of proper transformation and calibration. Most users are content if the reds look moderately red and greens look rather green, etc. In our electronic publication, *Frank Lloyd Wright: Presentation and Conceptual Drawings*, we included a Q60 image on the disc; I suspect few users have used this target for even visual inspection, let alone employed it for calibration. The only controls that are apt to see widespread use are those that are built into applications and underlying software. Such features are poorly developed at the moment.

It is unusual for me to advocate specific products, especially when the parent company is a founding partner of ours. But it is worth saying a few words about Kodak's Photo CD technology because of its considerable relevance in this context. In rough terms, Photo CD consists of three parts: an image color model, a storage method for CD media, and a more loosely connected family of color management systems (CMS). It is the first, but primarily the latter of these parts that are the most relevant here. And Kodak's CMS products are especially interesting because of Kodak's pervasive impact. Kodak has done the leg-work with hardware and software vendors to specify the transforms that map devices and computer systems to the Photo YCC model. The companies that support Photo CD include a legion of equipment manufacturers and software developers. Perhaps most significantly CMS software has been incorporated into the operating systems provided by Sun, Apple, and Microsoft. For practical and theoretical reasons it may be desirable to go into and out of Photo CD at both the upper production end and at the downstream user end. But there is no other conduit across the disparities of the computer industry that preserves as successfully the integrity of image content.

I occasionally hear the plaint that Photo CD is a proprietary, "closed" format, and that "open" international standards such as RGB.TIFF or CYMK.TIFF are therefore preferable. What is missed in these comments is that while the format, or the package of the image is standard, each transformation into RGB and CYMK is vendor determined. The ability to interpret an image does not account for the content. It is precisely because Kodak provides both the format and the CMS controls that there is a healthy prospect of images looking the same in different environments. Although it will certainly be a while before we can aspire to effortless control at the user level, Kodak's Photo CD is one path with fewer bumps in the road.

Production and Management Documentation

Digital images are beginning to stack up like cord wood as museums and archives plunge into more and more electronic applications. Yet from a management standpoint, it is not at all clear that thousands of image files represent an advance over collections of film and prints. At least with photographic materials, institutions have had decades to develop filing and recording methods. I want to emphasize that this is not a problem of how to describe the contents of an image, which is an entirely separate discussion. Rather, production and management data answer the questions of: what is the source of this image; how was it created; what are its characteristics, where can it be found; and what is it called. Very early on we developed a set of text records that represented this information, and our production and resource management system accumulates and tracks this data during every project. For a typical project we will generate and export files containing the following record types for each image:

- the photographic reproduction source of the image;
- the master digital image created from the reproduction;

- the derivative images generated from the master image;
- the media on which images are stored.

Pointers within each record maintain the interconnections between versions of images and where they reside. There is no one time or place where all of this information is acquired: reproduction data is entered when film is received at Luna; image capture variables are recorded during scanning; image characteristics and linking references are added in a subsequent production step. Not incidentally, working with visual materials raises a very problematic association: How do we know that we have connected the right image with the right text record? To address this quandary, we create a small reference image that tags the text record and can be used for Q&A review.

It is hard to communicate in words the magnitude of information that is generated. A project that involves ten thousand images may well generate forty to fifty thousand data records, depending on the number of image versions. The information represents the basic management documentation on an image collection, and is also used to build tables needed for control and access within computer applications.

We deliver management data as a part of every imaging project we do. It is extraordinary to me that I do not hear of such documentation as an essential part of production. I do see people adding extra data fields in databases to record image file names, but these pointers seem a modest step toward what is needed. To state the obvious, unless certain production information is recorded at the time work is done, the opportunity to record the process is lost.

If I see shortcomings in what we are doing in documenting images, they are traceable to the lack of standards in this area. We have responded to a practical need in our work, and have settled on the information we believe is important to record about production and the resulting image resource. These recording procedures have become stable over time, but the data would become even more valuable if there was broad community consensus on a preferred framework. Compatibility of image data from multiple sources and the potential to develop software around access to a common framework would be some of the advantages.

Conclusion

Despite a good deal of arm-waving implications to the contrary in the media industry, I would like to take as a given among this audience that there is little confusion between a work of art or other cultural artifact and a reproduction. Standing architecture and other three-dimensional objects only make obvious what is true of flat or two-dimensional works: no reproduction can substitute for the original regardless of its quality. This is true in so many interesting ways that it deserves an article of its own.

Moreover, what is important about visual archives in a museum may not correspond to popular notions of clone-like fidelity. I am reminded of this when I encounter the horror of outsiders wishing to do a multimedia project, only to discover that black and white prints have been the traditional medium of photo documentation among cultural institutions. The way I prefer to think about it, is that reproductions allow us to *postpone* the need for the original object when the quantity, location, or fragility of material make this advantageous.

The topics discussed in this presentation are about how to perpetuate the value — both cultural and economic — of visual collections in the electronic medium. There is no assumption about the reproduction source: a photograph of an object that has changed or no longer exists may be remarkably important despite its poor quality or condition. The goal is to move the photographic collection to a digital archive with as little loss as possible. And the Emphasis is on *collections*. No institution can afford to ignore issues of consistency and uniformity of method. The archival images serve as the long-term repository of quality from which derivative images are created for current applications. Consistent with an old lesson in a new medium, even lower quality derivative images will be greatly superior if they are generated from a high quality digital source.