# SELECTING A PICTURE FRAME WITH DIGITAL IMAGING

Alan Cole, James Christensen, and Frank P. Giordano
IBM T.J. Watson Research
P.O. Box 218
Yorktown Heights, NY 10598
Tel: 914/945-1033 Fax: 914/945-4003
E-mail: colea@watson.ibm.com, ibmjim@us.ibm.com, giordanf@us.ibm.com

and

David Sheridan
Julius Lowy Frame & Restoring Company, Inc.
223 East 80th Street
New York, NY 10021
Tel: 212/861-8585 Fax: 212/988-0443
E-mail: DavidSheridan@hotmail.com

## ABSTRACT

We describe an integrated hardware and software solution to provide assistance in selecting an appropriate picture frame for a valuable work of art. A custom hardware setup is used to capture separate high-quality, accurate color, images of frames and works of art. A user then applies the software to search a database of frame images for those of an appropriate size and style for a particular work of art, and to digitally compose separate images of frame and art to achieve a *virtual picture framing*: a realistic image of what the painting would look like in the given frame.

## INTRODUCTION

This paper describes a joint project between the Julius Lowy Frame & Restoring Company, located in New York City, and the Image Applications group at the IBM T.J. Watson Research Center. The Julius Lowy Company is an internationally known supplier of antique picture frames for works of fine art. Its clients include art galleries, auction houses, and museums as well as private collectors.

Selecting an appropriate frame for a work of art is a complicated process. Deciding such aesthetic issues as whether a French frame is satisfactory for a Flemish painting is not something for which the imaging technology described here provides answers. There are, however, other more practical problems where the technology does assist.

First, it may not be possible or desirable to bring the work of art to the showroom. In this case, the work may be represented only by a transparency or photograph. Holding a photograph up against the frame gives at best a poor impression of what the final result will be. Second, the picture frames themselves can be large and unwieldy, and are stored in two separate buildings, so that even if the art work is available, a great deal of lifting and carrying is involved in trying out different frames. A further complication is that the picture and the antique frame are likely to be of slightly different sizes. Lowy craftsmen may need to alter a frame to make it fit a particular piece of art, a step they are reluctant to take before the sale! These circumstances conspire to make it difficult to gain an accurate impression of what the final combination of art and frame will look like.

Our project addressed this problem by creating a solution of integrated hardware and application software. In brief, the most significant portion of the inventory of antique frames has been scanned over time to establish a digital library of images. In a typical scenario, a client arrives with a painting (or a photograph of the painting) for which framing advice is sought. While the sales associate discusses options, the scanner technician scans the painting and processes it. The sales associate may now display an image of the painting and search for candidate frames based on

country of origin, size, style, or other attributes. Finally, the software can synthesize a realistic image of the painting in any of the candidate frames, and produce a high-quality print of the image.

In following sections, we describe our approach more fully. We first describe the hardware configuration that was developed for the project, and the behind-the-scenes work that goes into capturing and preparing images for the application software.

We then describe the software that is used on the showroom floor to search for and present images to a client. We discuss in some detail the image processing used to compose a "picture-in-frame" image while preserving high-quality and accurate color, both for on-screen and printed images.

## BACKROOM OPERATIONS

We describe the hardware designed for the project and the tasks of image capture and preparation. Both are "behind the scenes" and not readily apparent to the customer who is interested in acquiring a picture frame.

### Hardware Configuration

Although the IBM side of the project has a long history in high-quality digital image capture, the extreme range in size of the objects to be scanned (from a small transparency up to a 60x75 inch object) precluded use of off-the-shelf hardware. A custom hardware configuration was required.

We did use the existing IBM Pro/3000 Digital Imaging System [Yao91] to capture images. This device (which is a research effort and not a full-fledged IBM product) uses a 3072 by 64-stage TDI CCD to capture images up to 3072 x 4000 pixels in size with a signal to noise ratio of better than 3000:1, and with accurate (colorimetric and calibrated) color. Depending on the size of the object being scanned, either a Rodenstock 105mm or a 150mm lens is used. An image is captured by making three separate scans through colorimetric green, red, and blue filters, designed to mimic human color perception. The camera captures 12 bits per pixel for each color plane, resulting in a 36-bit per pixel TIFF file.

We mounted the camera on a motorized column above a copy stand so that it can swivel through a 90-degree angle. It may be positioned to point down at the top of the stand to scan small objects. When scanning objects larger than 8x10 inches, the camera is swiveled to point at the wall, and the entire stand is rolled to a preset stop (depending on the size of the object) on tracks we embedded in the floor, specially reinforced to reduce vibration. Cables are fed to an overhead track to prevent their trailing across the floor.



Figure 1: Hardware installation

To help reduce stray light, we constructed a partition between the camera stand and the wall where the frame (or art) is hung. This partition has a centered opening, into which one of several inserts with successively smaller apertures may be placed, to yield an opening of an appropriate size for the object being scanned.

To achieve uniform illumination of the wall for objects ranging in size from less than 16x20 up to 60x75 inches, we installed a set of eight quartz halogen lights and diffusers between the partition and the wall. The lights, two each at left, right, top, and bottom, are on movable arms that are adjusted to accommodate the size of object being scanned. To image certain types of varnished oil paintings (which cause significant glare), we also found it necessary to include optional polarizing screens for the lights, used in combination with a polarizing filter on the camera lens.

Because the space between the lights and the wall is limited, and because the larger frames can be quite heavy, an X-Y stage was mounted to the wall with a large moveable panel which can be
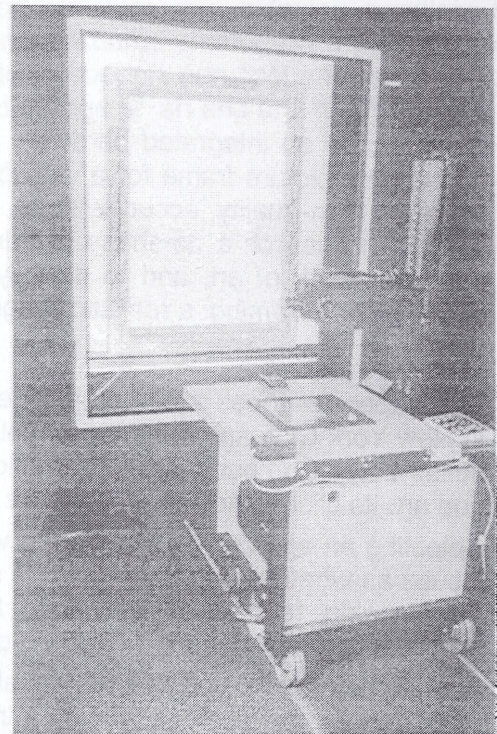
driven left or right as well as up or down by an electric motor. To hang a large frame, the panel is moved out from behind the lights. The frame is hung on the panel, which is then driven back into place, positioned so that the frame is centered with respect to the aperture in the partition (and hence, with respect to the camera). The panel, which provides the background for scanned objects, is, as nearly as possible, a uniform white, while the room itself as well as the partition are painted a flat black. Figure 1 shows the installation. Further details on the background and development of this equipment are found in [Gio99].

Several different formats have been defined (from 16x20 to 60x75 inches for objects hung on the wall, and 4x5 or 8x10 inches for objects photographed on the copy stand). Each has a standard setup in terms of distance of the camera from the wall, position of the camera on the motorized column, and so forth. In practice, throughput is improved by scanning a number of items of approximately the same size with a single setup, eliminating unnecessary switching among configurations.

Each format is color calibrated. A Macbeth ColorChecker color calibration chart is scanned with the given setup. Since the colors on the chart are known precisely, the colors actually captured by the scanner can be compared with the true values, and used to compute a color characterization of the scanning system. This, in combination with the color characterization of each output device, provides the basis for end-to-end color management [Gio97] and accurate color, one of the key features of the overall system.

**Image Capture and Preparation**

Images of frames and works of art are captured with the same physical setup, but fit into the overall workflow in different ways. Frames are digitized in batches ahead of time, while paintings are typically scanned one at a time as each is brought in by a client for framing.

To scan an object, the camera setup is established. A quick initial black and white prescan is performed to check for focus, lighting, and to mark the area to be scanned in detail, and then the three-pass color scan is performed. Each image is inspected for quality and annotated with an inventory reference (for frames) or with an identifying title (if a work of art).

In the case of frame images, the application software must be able to distinguish between the frame and its background, or, in other words, know where the "hole" in the frame is, so that the image of a painting may be inserted. The goal is to produce a mask that differentiates the frame from the frame background.

This is done with the aid of a semiautomatic software tool. There are several modes of operation, but one example may help to illustrate it. Frames are scanned on a background of (more or less) uniform white. In one mode of operation, the operator clicks a mouse in the "hole" of the frame, and all contiguous areas of the same color, or within a specified threshold of the same color, are assumed to be part of the background and temporarily changed to a distinctive color (bright blue, say). This makes it visually apparent which part of the image is considered to be background, and which not. If an area of background is left out because its color is too different, a second mouse click in the omitted area will also mark it as part of the background. In other cases, a portion of the frame (perhaps a bright highlight) may be erroneously included because it is too close to the background in color. The tool is then switched to a subtractive mode, and that portion of the frame marked as not belonging to the background. After a few iterations by the operator, the result is a mask (a 1-bit per pixel version of the image, saved with the image) that defines the frame background. A set of four lines defining the quadrilateral which best defines the interior of the frame is also optionally stored with the image (optional, because a quadrilateral is not always a good description of a frame's interior).

For a work of art, no masking is performed. Instead, the image is cropped to size. If the image is scanned from a photographic transparency or print, the actual size cannot be deduced, and so the operator is prompted to enter the dimensions of the original work.

At this stage, the images, whether frame or art, are submitted to a batch process which reads the image files, converts from the original 36-bit color to a 24-bit color format, and produces three derivative images for each original. The first derivative is a relatively high-resolution image, which,

in the case of a frame, includes the background mask, and is up to 2000 x 2500 pixels in size. This will be used for printing, and also for later production of detail views and picture-in-frame images. The second and third derivative images are designed for viewing on screen; they are about 1100 x 940 (for a high resolution screen) and 830 x 680 pixels (for a laptop with lower resolution screen), and, in the case of frames, have their backgrounds set to black. All derivative images are saved in a device-independent format, and carry with them the color characterization of the scanner, so that they may later be displayed with accurate color.

Producing the derivative images can take a minute or two each, so typically a batch will be processed without operator intervention. After the job is complete, each image may be viewed as a quality assurance check and, if satisfactory, placed on a file server for application use.

## THE SHOWROOM APPLICATION SOFTWARE

### Overall Goals

The application software is intended to allow a sales associate to search the database of frames; to view high-quality, accurate color, images of frames and art; to create and view "picture-in-frame" images; to save selected collections of images in named folders; and to provide high-quality prints of any desired images. It is further designed to allow remote use (on a disconnected laptop, for example) as well as in the LAN environment in the store.

### Architecture of the Application

The application exploits a client/server architecture. The three derivative versions of each image are stored, in device-independent format, on the server's file system; the frame database and application data files are also stored on the server. The application software runs on client computers, which are all connected to the server via a LAN (in the Lowy environment, the server also acts as a client). A laptop computer may be configured as both a server and client, permitting use at remote locations, disconnected from the LAN. We use a variety of Intel-based personal computer models, but timing information mentioned here is measured on a 200 MHz Intel Pentium Pro computer with 96MB of memory.

The application program itself is written in a combination of Java (for the program logic and graphical user interface) and native C++ code (for the image processing algorithms). It allows the user to navigate through a series of screens offering different functions either by touching "buttons" on a touch sensitive screen, or by clicking on the buttons with a mouse.

### Search Capabilities

The program makes use of a preexisting database listing both the frames currently in inventory and those that have been sold. It characterizes each frame by a number of attributes, including such information as an inventory number, period, country of origin, style, and sight size (inner dimensions of the frame).

To make the program accessible to a wide spectrum of users, it was especially desirable to have a simple and intuitive search mechanism that did not require learning a specialized query language.

To achieve this, the user interface represents each possible value of an attribute by a checkbox on the screen. One or more boxes may be checked (by touch or mouse); this represents the logical OR in a search. If multiple attributes are selected, this represents a logical AND in the search. For example, if the "English" and "Spanish" boxes are checked on the country of origin screen, and the "17th century" and "18th century" boxes are checked on the period screen, this represents a query to find all frames that satisfy:

```
(English or Spanish) and (17th or 18th century)
```

As the checkboxes are selected or cleared, a field on the computer screen is immediately updated to show how many frames satisfy the current query. The entire database is held in memory, so that result sets from these simple queries can be computed by bit operations on in-memory bit vectors, giving very fast response times.

Frame size is handled in a different manner. Entry fields allow one to specify the width and height desired for the frame sight size, as well as the allowable variation. These fields are pre-filled with values derived from the size of the most recently examined work of art, so that to find a frame of an appropriate size for a particular painting may require no typing at all. Again, the number of frames within the given size range is shown on the screen, allowing one to widen or narrow the search as appropriate.

This approach to searching, while not providing for arbitrary queries, is straightforward, easy to understand, and fast.

## Viewing Frame and Art Images

The software shows images of the frames resulting from a query. This requires that the device-independent image formats be converted to a device-dependent format, involving two aspects of image processing: color management and image enhancement.

### Color Management

We considered accurate color to be an essential component of the application. Accordingly, each client computer is separately color calibrated. By displaying a series of colors of known values and measuring what is actually displayed, it is possible to compute the transformation produced by each display card/monitor combination [Min92].

Given this information, it is now possible to convert the device-independent image files stored on the server so that the image color is correctly compensated for the monitor on which it is displayed.

### Image Enhancement

The standard operations of brightening, contrast enhancement, and unsharp masking are applied to each image. Several sets of enhancement parameters are used, depending on such factors as whether the scanning was done through polarizing filters, or whether the image is being displayed on a CRT or an LCD screen. In addition, individual images may be "stamped" with unique parameters if the defaults do not seem satisfactory to the scanner operator.

### Image Display

In browsing through the inventory of frame or art images, or the results of a query, the application displays thumbnail images, 12 per screen. It is desirable to be able to do this at speed, so the device-dependent thumbnail version of each image is computed once on each client and saved permanently on a local disk.

To display a full-sized image, the correctly sized TIFF file is read from the server and converted into a device-dependent image for display on a particular monitor, a process that typically requires about six seconds. Display-ready images are cached on a local disk, reducing display time to about a second when an image is found in the cache.

## Creating Portfolios

The software allows the definition of named "portfolios," or folders of images. This is a convenience, which allows a sales associate to collect selected images in a convenient place. For example, the image of a painting and images of several candidate frames, perhaps including some corner detail images, could all be placed into such a group for later viewing by a client.

One of the benefits of using a portfolio is that images that take a relatively long time to produce, such as picture-in-frame images, are saved. Future display of these images can be done immediately, without having to compute them first, and so is considerably faster.

## The "Picture-in-frame" Capability

To find out what a painting would look like in a particular frame, one of the key motivations for our work, the user initiates the process by selecting which frame and which work of art are to be combined. This results in a dialog providing three options.

One option is whether the frame should be rotated or not. By default, the frame will be rotated so that its orientation (landscape or portrait) matches that of the art, but sometimes there are reasons to override this (when the frame has a distinctive top edge, for example).

A second choice is whether to digitally produce a simulated shadow along the edge of the frame. This is the default, as it makes the resulting image look somewhat more realistic.

A third option deals with any possible mismatch between the sizes of frame and picture. In practice, Lowy artisans will modify a frame to fit a desired painting. To simulate this, the software is capable of digitally cutting down or extending the frame. It modifies each side of the frame image in one, two, or four places. By default, this depends on the style of the frame, as indicated in the database, so that a frame with decorative scrollwork in the center of each side would be modified in two places per side, at one-third and two-thirds of the distance along the frame, leaving the scrollwork intact. A frame with demi-centers would be modified in four places, while a frame with plain edges would be modified in the center of each edge.

Another possibility is not to simulate the frame modification at all, but instead to digitally shrink or expand (scale) the whole frame to fit the desired artwork. This is not normally desirable, since the frame will now be displayed at a different scale than the art, providing a physically unrealistic view.

### Picture-in-Frame Processing

Once the options have been selected, image processing begins. First, the high-resolution version of the frame image is read into memory together with its mask. Brightening and contrast enhancement operations are applied, and the frame mask is used to set the background color to uniform black (for on-screen viewing). An exponentially weighted smoothing function is applied to the edge of the frame to eliminate aliasing effects at the boundary between frame and background. The frame image is now rotated if necessary.

In the next step, the art image is examined to compute a target scaling factor that will generate the same resolution as the frame image; this determines the resulting size of the art image, and hence the required size of the frame's interior, in pixels.

At this point, the frame image is modified to make its interior the required size. If the frame is too large, the requisite number of pixels is cut out, evenly divided over the one, two, or four cut locations. If the frame is too small, it is made larger by replicating small bands of pixels one or more times, at the one, two, or four cut locations.

Next, the art image is read into memory and brightening and contrast operations applied. The image is scaled and (if necessary) rotated. The output from this sequence of operations is directed to the interior of the frame, using the frame mask to determine whether each pixel is taken from the frame image or art image.

Now the simulated shadow is generated if requested, and if the optional quadrilateral characterizing the frame's interior is available. An exponentially weighted function is used to darken the four outer edges of the art, using the four mask lines as guides.

Finally, the composite image is scaled to an appropriate size for on-screen display and unsharp masking is applied. The result is written to a device-independent TIFF file (this means that the results of the process may also be used on other client computers without loss of color accuracy).
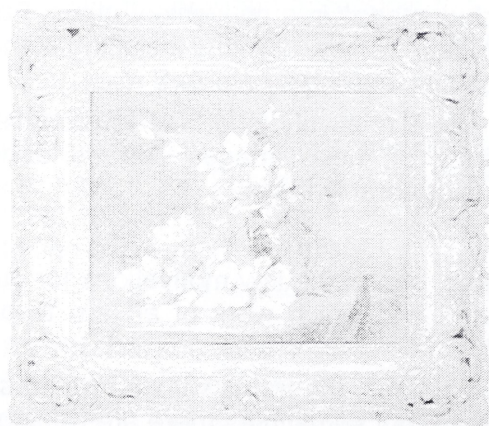


Figure 2: "Picture-in-Frame" image

To actually display the image, the TIFF file is converted to a device-dependent format as described previously; color management is performed, though further image enhancement is bypassed.

This entire process is the most memory- and CPU-intensive part of the application and, even though the image processing algorithms have been carefully optimized in C++ code, computing a picture-in-frame image can take a minute or two. If the image is placed in a portfolio, as described above, then a permanent copy of the image is made and future display of the image on this or any other client will, as with frame and art images, require at worst six seconds or, if the display-ready form is cached, about one second.

**Printing**

The software also provides the capability of printing any of the images. As it is currently configured at Julius Lowy, both a high quality dye-sublimation printer and a lower quality ink jet printer are connected to the server, and either one may be selected.

For printing, the largest (highest resolution) version of the image is used. The image is scaled as necessary, frame backgrounds are set to white (rather than black), color correction is performed, and the usual image enhancement operations (brightness, contrast, and unsharp masking) are performed with settings optimized for printing. The result is a TIFF format image file; the software then converts this to Postscript for actual printing.

For this application, we considered color accuracy of the printed output to be extremely important. Although the printer was calibrated (by printing a set of patches of desired colors, then measuring the actual printed colors with a spectrophotometer and computing the inverse transformation), our initial results were occasionally disappointing. Investigation showed that the main problem was with colors that were out of the printer gamut; that is, colors that were captured by the camera but that could not be produced by the printer. When such a color is to be printed, it must be mapped by the color management software to a color that is realizable. The method in which this is done can make a significant difference. We formulated a new method of handling out-of-gamut colors that now gives consistently improved results.

## CONCLUSION

We have described a project to build a fully integrated hardware and software solution for an antique frame dealer. We have highlighted two significant aspects of the project: first, the custom hardware configuration we designed to capture high-quality digital images of the frames and, second, the techniques we developed to synthesize a highly realistic picture-in-frame image from separate images of the frame and art.

## ACKNOWLEDGMENTS

## REFERENCES

[Gio97] E.J. Giorgianni and T.E. Madden. Digital Color Management: Encoding Solutions, Addison Wesley Longman, Reading, MA (1997).

[Gio99] F.P. Giordano, G.W. Braudaway, J.E. Christensen, J.C. Lee, and F.C. Mintzer, *Evolution of a High-Quality Digital Imaging System*, Proc. IS&T/SPIE International Symposium on Electronic Imaging, vol. 3650 (1999).

[Min92] F.C. Mintzer, G. Goertzel, G.R. Thompson, *Display of images with calibrated color on a system featuring monitors with limited color palettes*, SID International Symposium, Boston, Mass., vol. 23, pp. 377-380, May 17-22, 1992.

[Yao91] Y.L. Yao, F.P. Giordano, H.S. Wong, W. Kang, and J.C. Lee, *Design Considerations for a High Quality Camera Type Scanner*, Proc. IS&T 7th Intl. Congress on Adv. Non-Impact Printing Technology, Portland, Oregon, 441-450, (October 1991).