UBC ScanCam –
An Inexpensive 122 Million Pixel Scan Camera

Shuzhen Wang and Wolfgang Heidrich
The University of British Columbia

ABSTRACT

We describe the design of a very high resolution, low-cost scan camera for use in image-based modeling and rendering, cultural heritage projects, and professional digital photography. Our camera can acquire black&white, color, and near-infrared images with a resolution of over 122 million pixels and can be readily built from off-the-shelf components for less than $1200. We discuss the construction of the system as well as color calibration and noise removal. Finally, we obtain quantitative measurements of the light sensitivity and the optical resolution of our camera and compare the image quality to a commercial digital SLR camera.

1. MOTIVATION

Although current high-end consumer and semi-professional digital cameras start to approach the resolution of 35mm analog film, there is a strong demand for even higher resolutions in digital photography. Large photographic prints require resolutions closer to analog large format photography. Many electronic cultural heritage applications and research projects in image-based modeling and rendering would benefit from higher resolutions as well.

For example, many museums are currently starting to digitize portions of their collections for research purposes. They are interested in recording as many of the exhibits as possible, and capture the geometry as well as very high resolution images at multiple different wavelengths. The hope is that these databases will be a powerful resource for future generations of researchers in history, art history, and anthropology.

Unfortunately, the resolutions required for these applications are currently only provided by professional digital camera backends for large format view cameras. These can easily cost several 10s of thousands of dollars for the backend alone. The highest image resolutions are achieved with *scan backends*, where a single line of high resolution sensor elements is swept across the image area.

In this paper, we present a scan camera system in which we turn a $100 consumer-grade flatbed scanner into a very high resolution digital camera backend that we then use together with a conventional $8'' \times 10''$ large format view camera.

This allows us to acquire images at more than 122 million pixels (corresponding to an image plane of size $8.5'' \times 10''$ digitized at 1200 DPI). The total system, including large format camera, lens, scanner, and color filters can readily be assembled for less than $1200. Furthermore, the approach will scale up as scanner technology improves - the scanner is only about 10% of the total system cost.

As an additional benefit, we can use the same camera with an infrared filter for photography in the near-infrared (NIR) spectrum, which has proven useful for applications such as object recognition and food processing, as well as digital matting in cinematography.[1, 2]

The remainder of this paper is organized as follows: in Section 2 we briefly review technologies for very high resolution photography. We then introduce our camera in Section 3, in particular discussing

- hardware setup (Section 3.1),
- black&white imaging including calibration and noise removal (Section 3.2), and
- color imaging including alignment of color planes and white balance (Section 3.4).

Afterwards we provide results, including quantitative resolution measurements and a comparison of our camera with a semi-professional digital SLR camera, the Canon EOS D60 (Section 4).

2. RELATED WORK

The technologies for professional and semi-professional digital photography can roughly be grouped into two categories: digital single lens reflex (SLR) cameras, which are similar to conventional analog SLRs in design, and digital backends for medium- and large format cameras. The latest generation of digital SLR cameras achieve resolutions of about 6 million pixels, which produces images roughly comparable in quality to analog 35 mm cameras.

For higher resolutions like the ones required, for example, for large-format prints, the user has to turn to digital backends for medium format (i.e. 6×4.5cm) or large format (typically $4'' \times 5''$) cameras. These backends exist in two basic flavors: *single shot* backends have a 2-dimensional image sensor, similar to the one used in consumer cameras and digital SLRs. This technology can now achieve image resolutions of up to 22 million pixels (for example, with the Sinarback 54,[3] which uses a Kodak imaging sensor).

Even higher resolutions can be achieved with a so-called *scan backend*. In this technology, a single row of sensor elements (*line sensor*) is swept across the image plane similar to a familiar flatbed scanner. For example, the Betterlight Super 8K[4] scan camera is a backend for $4'' \times 5''$ large format cameras, that can achieve an optical resolution of 8000×10660 pixels (12000×15990 interpolated). The down side of this approach is that the technology cannot be used for animated objects and is only useful in studio settings.

Market Segment	Resolution	Price
Consumer	<4M pixels	$200–1,500
Semi-Professional (SLR)	<6M pixels	<$3,000
Professional (single shot)	<22M pixels	<$30,000
Professional (scan)	<100M pixels	<$25,000
Our scan camera	122M pixels	$1,200

Table 1. Overview of resolution and price for different digital camera technologies (the prices for professional backends exclude the actual view camera and lenses).

Unfortunately, both single-shot and scan backends are quite expensive. Table 1 shows a comparison of the resolutions and the price points for the different technologies. Note that the information for the backends does not include the cost of lenses or the camera, but only of the backend itself. In this paper, we describe how to build a scan backend for a $8'' \times 10''$ large format camera with a resolution of up to 10200×12000 (122 million pixels) from a consumer-grade scanner. The whole system including camera and lens can be readily built from off-the shelf components for about $1200.

A research project related to our work is the light field camera built by Yang et al.[5, 6] Their system is based on a flatbed scanner with a lens array attached to the glass surface. Since their optical setup is quite different from ours, the calibration issues differ significantly.

2.1. Color Imaging

Color imaging can be achieved in one of several ways in digital photography. One can split the light into three different color components using a beam splitter and suitable filters, and measure the components with independent sensor chips. Due to alignment problems this procedure is only suitable for relatively low resolution cameras, such as video cameras. In single shot photo cameras there is usually only one chip with different color filters in front of every individual pixel. The different color filters are interleaved between the pixels in a regular pattern (Bayer pattern[7]). This, however, means that only one color channel can actually be measured at every pixel; the remaining channels are interpolated from neighboring pixels. Scan cameras on the other hand often have three rows of pixels that obtain true measurements of all color channels in every pixel as the sensor sweeps across the image plane.

Alternatively, one can use a sensor that measures all wavelengths and take three images with different color filters (usually arranged in a filter wheel) that are then combined into a color image. This is the approach we are using in our camera, since the scanner technology that we base our work on has only a single sensor (see Section 3.1 for details). The method is suitable for both single-shot and scan cameras, although the former can in this way no longer be used for moving targets. While this method adds to the complexity of the acquisition process, it allows for additional flexibility by using non-standard filters that select specific ranges in the light spectrum.

3. CAMERA SYSTEM

At the core of our camera system, we use the Canon LIDE 30 scanner (\$100), which is based on Canon's technology "LED in Direct Exposure" (LIDE[8]). As we will see in Section 3.1 this particular scanner design proves essential for our project. We modify the scanner to work as a backend for a conventional view camera. To keep the cost of the overall system down, we use the $8'' \times 10''$ view camera kit from Bender Photographic[9] (\$400), but any other view camera would work as well. Since view cameras need to work in both landscape and portrait mode, the image plane of an $8'' \times 10''$ camera is actually $10'' \times 10''$ in size. Due to the dimensions of the scanner, we can thus obtain images of up to $8.5'' \times 10''$.

The large-format lens is easily the most expensive part of the system. In our setup we use a Nikkor-M 300mm f/9 (≈\$600), which is a relatively inexpensive portrait-style lens for large-format photography, and has enough angular coverage to illuminate the whole image plane. In addition, we need three color filters (at \$12 each) to create a manual color wheel and a UV/IR cutoff filter (\$40) for color imaging. Figure 1 shows a photograph of the overall camera system.

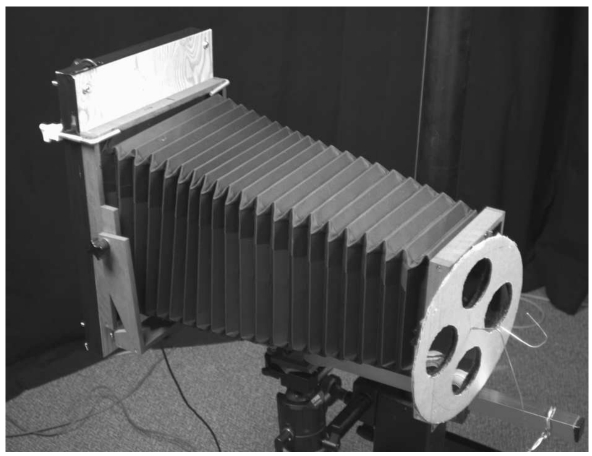

Figure 1. A photo of our camera setup.

3.1. Hardware Setup

As mentioned above, the central part of our camera is a LIDE 30 scanner from Canon. In contrast to other flatbed scanners, this one has a design where the sensor covers the complete width of the paper (about $8.5''$), and focusing is achieved through an array of lenses directly on top of the sensor (rather than a combination of mirrors and a single lens). The conceptual design is depicted on the left of Figure 2. The illumination in these scanners is achieved by illuminating the paper with different color light emitted by colored LEDs.[8] That light is spread out evenly across the width of the paper by a light guide. The sensor itself only has a single row of elements that are sensitive to all visible wavelengths, and in fact also to near-infrared light.

For our purposes, we remove the filter array in front of the sensor, because we want to use the lens of the large format camera for creating a focused image instead. We also disable the LEDs and remove the light guide, so that only light from the lens arrives at the sensor.

Note that a similar approach would be difficult with a more traditional scanner technology, where the imaging sensors is significantly shorter than the full width of the paper. These scanners use an optical system made of a lens and several mirrors to scale the width of the scanline down to the sensor width and at the same time focus onto the paper surface (see right of Figure 2). Because the mirrors are very narrow, they restrict the angular coverage too much. One would therefore have to redesign the complete optical system, which in turn would require changes to the mechanics as well.

For color imaging, we built a filter wheel from cardboard and photographic filters made from polyester. The filter wheel has four slots, three of which we use for Lee Tricolor filters, models 25, 58, and 47B (for red, green, and blue with peak transmission at 630, 540, and 440nm, respectively).[10] The fourth slot either remains empty (for black and white photography), or gets loaded with an infrared filter (number 87 of the same company) to block out frequencies below 730nm. In contrast to some commercial filter wheels, our current wheel is manually operated, that is, the photographer has to rotate it to the next filter slot for every scan.

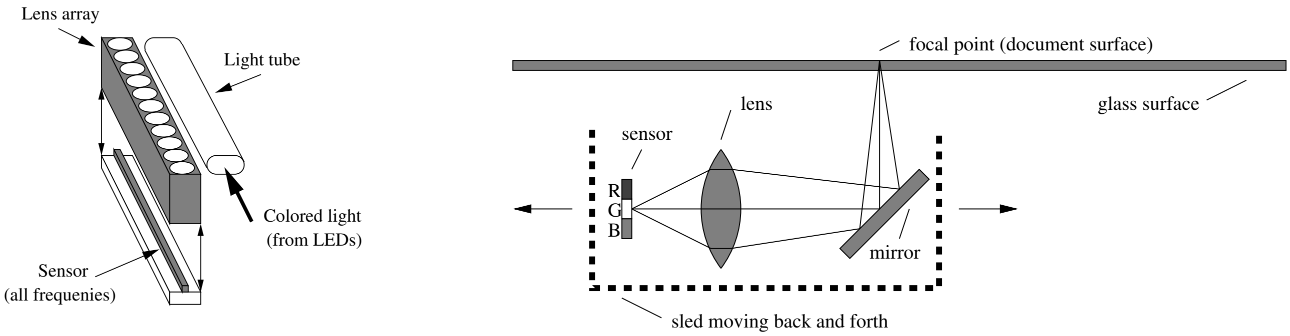

Figure 2. Left: Design of a scanner with the LIDE (LED in Direct Exposure) technology. Right: Design of a conventional flatbed scanner.

3.2. Black&White Imaging and Calibration

For black&white photography, we acquire a single image without a filter. Our scanner takes about 5 minutes to scan the full $8.5'' \times 10''$ visible area at 1200 DPI in grayscale mode and to transfer the resulting data over USB 1.1 to the host. In order to get uncalibrated 16 bit raw sensor data we use the shareware program VueScan[11] instead of vendor-provided scan software. This is necessary since the usual scanner calibration steps will fail in our setup because the LED light source has been removed.

Like all sensor arrays, the individual imaging sensors in the scanner can have slightly different characteristics. The most important sources of variations for both CCD and CMOS sensors are *dark current* and *flat field response*.[12] Dark current results in a non-zero measurement even for pixels that have not been hit by any light at all. It effectively adds a positive constant to all measurements for one particular sensor element. Apart from dark current, the sensor response is very close to linearly proportional to the number of incoming photons. However, the constant of this linear mapping can again vary from element to element in the sensor array. To correct for both effects, we take one dark (i.e. zero aperture) image and one "white" image (i.e. an image of a white piece of paper, out of focus, and exposed to just below saturation). In order to filter out any random noise, we average the pixels in every column of both the dark and the white image (since all pixels in one column were measured with the same sensor). These two average values then define a linear mapping that removes the artifacts from both dark current and flat field response in a single step.

Figure 3. Processing pipeline for B&W images: first, dark-current noise and flat-field response are calibrated with a single linear mapping. Then, broken columns in the image are interpolated. Finally, gamma correction yields the end result.

In addition to these two effects, there are typically also individual sensors that are broken, or little scratches in the transparent plastic housing that are reflecting light away from the underlying sensor element. We deal with these defects by masking the corresponding columns (about 30 defects affecting some 90 sensor elements in our scanner) and interpolating the color from the neighbors on the same scan line. This simple approach works quite well in our case since the defects are very narrow (3-4 pixels maximum) and small compared to the total resolution. Alternatively, one could use a more sophisticated method such as image inpainting[13] to generate the missing information. The final step of the black&white image processing is a gamma correction depending on the characteristics of the desired output device and a subsequent reduction of the image depth to 8 bit. The individual steps of the black&white processing pipeline are depicted in Figure 3, while the results of the process are shown in Figure 4.

Figure 4. Effect of the black&white calibration step. Left: raw sensor data. Center: after dark current subtraction and flat-fielding. Right: calibrated image after interpolating faulty columns.

3.3. Infrared Imaging

Without additional filters, CCD elements are not only sensitive to visible light, but also to light in the near-infrared (NIR) spectrum, up to about 1000nm. While many cameras include infrared blocking filters to avoid wrong colors, scanner sensors do not usually have such a filter because scanner light sources do not emit significant quantities of infrared light. This, however, means that we can use our camera for infrared imaging if we use an IR filter (such as the Lee filter number 87[10]) to block out all visible light. Figure 5 shows a comparison of an infrared and a black&white image of a painted ceramic figurine that we took with our camera. Note that the different color paint used for the figurine is equally reflective in the infrared spectrum, so that pigmentation effects are almost completely removed. The infrared image would be a much better input for a shape-from-shading algorithm than the black&white image.

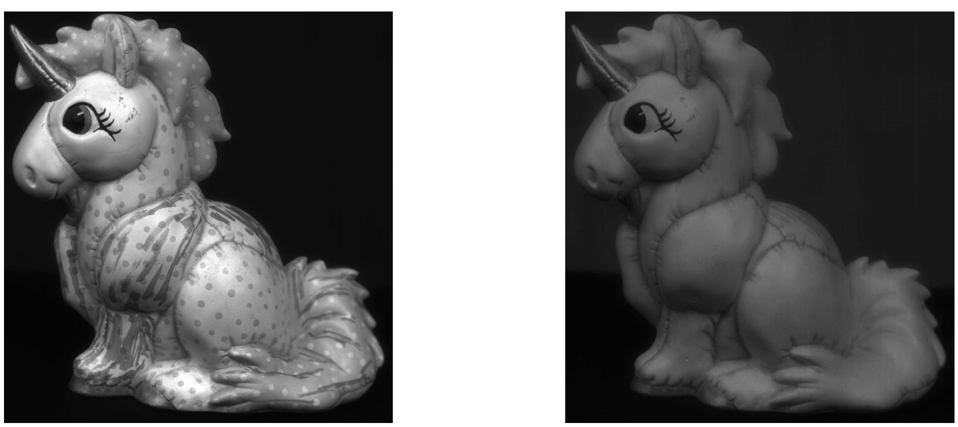

Figure 5. Comparison of a black&white photograph (left) with an infrared photograph (right). Note that except for the black color of the eye, the effect of different paints is almost completely removed in the infrared image.

3.4. Color

For color imaging we have to acquire three different scans with a red, green, and blue filter applied. These individual images go through the black&white calibration process described in Section 3.2, except for the final gamma correction stage. We then have to merge the individual color planes into a single image. We find the scans to be aligned very well in the columns (indicating that the sensor elements repeatably move along the same line in every scan), but the scanlines can be off by about 20-30 pixels, since the initial starting position can vary from scan to scan. To compensate for this, we select a region of the image with high-frequency detail, and then determine the one-dimensional offset by performing an edge detection on all channels of that region, and then running a one-dimensional search for the best offset.

As mentioned in the previous section, the sensor of our scanner is sensitive not only to visible light, but also to infrared. For color imaging, we have to make sure that our color scans do not include any infrared illumination, or the colors will end up being wrong. To compensate for this, we use an infrared cutoff filter (Schneider Optics B+W Filter no. 486) to eliminate the infrared component. This filter might also be useful for B&W photography, but it is essential for color imaging.

After we have merged the three color channels, we need to perform color calibration. A simple white point calibration (relative scaling of the individual color channels) works surprisingly well once the IR light has been filtered out from the color channels. The images in this paper were created with this simple color calibration method. Again, gamma correction according to the desired output device has to be performed at the end. In the future we plan to investigate more sophisticated color calibration methods such as ICC profiles. Since we use filters with fairly standard spectral characteristics, we do not expect any major surprises when applying standard ICC profile generators.

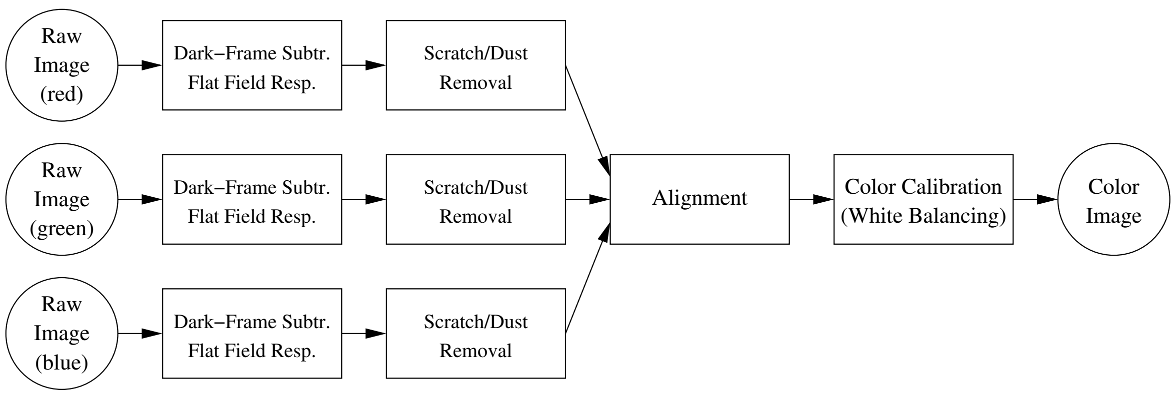

Figure 6. Processing pipeline for color images: first the individual color planes are treated with the calibration stages known from the B&W process. Then, the clor planes are aligned, and a color calibration using ICC profiles is performed. Note that the gamma correction is now implicitly performed in the color calibration stage.

4. RESULTS

In order to obtain quantitative results with our camera, we compared it against the Canon EOS D60 together with a Canon EF 28-105mm zoom lens. The camera is a 6 million pixel digital SLR camera that the second author is very fond of. Note that this camera alone is about twice as expensive as the complete camera/lens system described here.

As a first step, we tested the light sensitivity of our scan camera. To this end, we set both our system and the D60 to an aperture of f/9 (the largest aperture of our large-format lens). We then set the film sensitivity on the D60 to 100 ASA, and adjusted the exposure time until we got an image of a brightness comparable to the one obtained with our scan camera. This way we determined that our camera has roughly the sensitivity of a 100 ASA film operated at an exposure time of 1/60 s, or a 1600 ASA film operated at 1/960 s, and so forth. Like many digital cameras we can simulate a higher sensitivity by linearly scaling up the raw data (i.e. turning up the gain value). This multiplier can be rolled into the linear mapping we use for black&white calibration (Section 3.2). Of course this will result in an amplification of the noise, just like in all digital cameras that offer this feature.

We then analyzed the actual optical resolution we can obtain with our setup, since this effective resolution does not automatically correspond to the nominal resolution of the scanner. The usual way of evaluating the resolution of digital imaging devices is the modulation transfer function (MTF), which analyzes how different frequencies are attenuated by the combination of optics and sensor. A common way of generating this MTF is to analyze a focused image of a slanted edge.[14–16] In the case of scanning devices, the behavior in the rows and in the columns is analyzed separately.[16] Figure 7 shows the MTF for our setup (frequencies are given in cycles/pixel) using the method by Williams and Burns. We computed

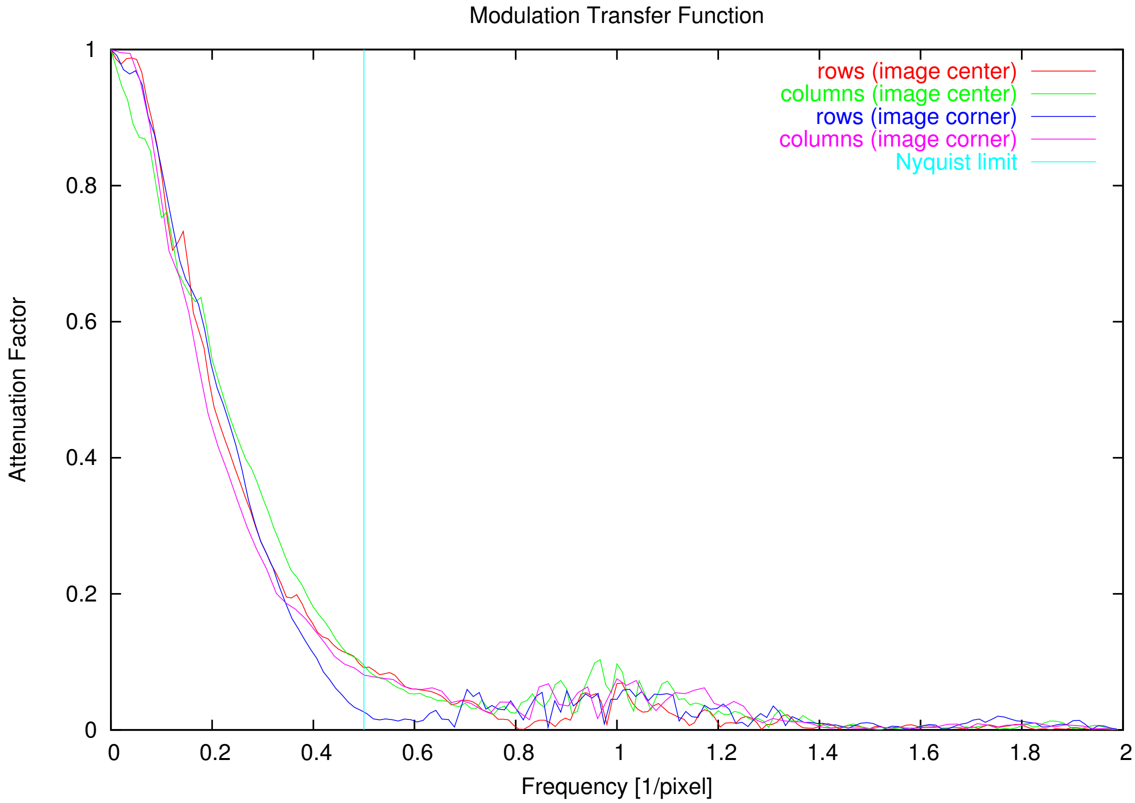

Figure 7. The modulation transfer function for the rows and the columns of our optical system

the MTF for different locations in the image, but could not detect any significant position-dependencies. The MTF shown here is taken roughly from the center of the image plane.

Due to the Nyquist limit, the highest useful frequency any imaging system can have is 0.5 cycles/pixel. In order to evaluate any real-world system, a common criterion is to look at the lowest frequency where the MTF drops below 10%. If this point is at or close to the Nyquist limit, the optical system is said to have a resolution corresponding to its number of pixels.[14, 16] On the other hand, an MTF with higher values for frequencies above the Nyquist limit indicates a potential for aliasing. In our camera, the 10% cutoff lies at about 0.36 cycles/pixel for the rows (i.e. within a scan line). For the columns, the cutoff is at about 0.54 cycles/pixel. Together, these values indicate that the optical system is very good, and operating at close to the nominal resolution we can expect from our scanner.

Figure 8 compares our camera to the Canon EOS D60. We set up both cameras to cover roughly the same field-of-view of the object at their respective full resolution. Unfortunately, this required us to change the setup to accommodate for differences in the lens and depth-of-field, so that the illumination between both images is not directly comparable. The top row of Figure 8 shows almost the full images acquired this way (they were slightly cropped to an aspect ratio that fits the page). The scan camera image illustrates a limitation of the separate acquisition process for every color plane: due to the heat from the light sources the leafs of the bamboo dried somewhat during the photo shoot so that they moved slightly in between the scans. Therefore, the color planes do not line up at the leafs.

For the images in the bottom row we reproduced a region from the scan camera image at 300 DPI, and chose a corresponding region in the EOS D60 image. The differences in detail for both images are quite dramatic. At this scale, the full image from the scan camera would fill a $34'' \times 40''$ poster. Note that the pillow visible in the photos is only about one inch on each side in reality. While the colors do not match exactly, the differences are well within the range of color differences one gets with different commercial cameras. In comparing the images to the real object, we find that the green of the stem is slightly better reproduced by our scan camera, while the red colors are a bit better in the EOS D60 image.

5. DISCUSSION AND FUTURE WORK

In this paper we describe a 122 million pixel scan camera that can be built from off-the-shelf components for about $1200. We discuss calibration issues, and evaluate light sensitivity and optical resolution of the camera system. The results indicate that the resolution of our system is dramatically superior to a digital SLR camera that is twice as expensive as our system. Although it takes scan cameras longer than a single shot camera to acquire an image, this added time is tolerable for applications where setup and lighting take up a significant amount of time (for example, product photography, or digitizing of cultural objects). This even more so if it results in a drastic gain in resolution.

We therefore expect that this camera design will be useful not only for professional photography, but also for research in image-based methods and for cultural heritage applications. In particular, we are currently in contact with several museums that are interested in digitizing as many of the exhibits as possible, and capture the geometry as well as very high resolution images at multiple different wavelengths, including infrared. Our camera should be a perfect match for these requirements.

In addition to seeking more applications for the camera, we also plan further work on the color imaging problem. First, we will investigate the use of ICC profiles with our existing color filters, and the extension to more narrow filters. We then want to explore to what extent the color wheel can be replaced with liquid crystal tunable filters (LCTFs[17]). These filters let only a narrow window of frequencies pass, but that frequency window can be moved to any place in the visible and near-infrared spectrum within milliseconds. This should allow us to take images with arbitrary spectral basis functions.

6. ACKNOWLEDGMENTS

We would like to thank Michael Goesele for help with determining the MTF of our camera. Thanks also to Dave Martindale for providing several discussions, and to Ted Gierschner from Lens&Shutter for help with selecting camera equipment. This work was supported by the BC Advanced Systems Institute and the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

1. R. Fielding, *The Technique of Special Effects Cinematography*, Hastings House, 1965.
2. P. Debevec, A. Wenger, C. Tchou, A. Gardner, J. Waese, and T. Hawkins, “A lighting reproduction approach to live-action compositing,” in *Proc. of ACM SIGGRAPH 2002*, pp. 547–557, July 2002.
3. Sinar AG. http://www.sinar.ch, 2002.
4. Betterlight. http://www.betterlight.com, 2002.
5. J. Yang, “A light field camera for image based rendering,” Master’s thesis, MIT, 2000.
6. J. Yang, C. Lee, A. Isaksen, and L. McMillan, “A low-cost, portable light field capture device,” in *Siggraph Conference Abstracts and Applications*, p. 224, 2000.
7. B. E. Bayer, “Color imaging array.” U.S. Patent 3,971,065.
8. Canon Inc., “LED in Direct Exposure (LIDE).” http://www.canon.com/technology/optics/lide/content.html, 2002.
9. Bender Photographic. http://www.benderphoto.com, 2002.
10. Lee Filter. http://www.leefilter.com, 2002.
11. Hamrick Software. http://www.hamrick.com/vsm.html, 2002.
12. J. Janesick, *Scientific Charge Coupled Devices*, vol. PM83, SPIE PRESS, 2001.
13. M. Bertalmio, G. Sapiro, and V. C. anf C. Ballester, “Image inpainting,” in *Proc. of ACM SIGGRAPH 2000*, pp. 417–424, July 2000.
14. ISO, “ISO 12233:2000 Photography – Electronic still picture cameras – Resolution measurements.”
15. P. Burns, “Slanted-edge MTF for digital camera and scanner analysis,” in *Image Processing, Image Quality, Image Capture Systems Conference (PICS)*, pp. 135–138, Society for ImagingScience and Technology, Apr. 2000.
16. D. Williams and P. Burns, “Diagnostics for digital capture using MTF,” in *Image Processing, Image Quality, Image Capture Systems Conference (PICS)*, pp. 227–232, Society for ImagingScience and Technology, Apr. 2001.
17. H. Morris, C. Hoyt, and P. Treado, “Imaging spectrometers for flourescence and raman microscopy: Acousto-optic and liquid crystal tunable filters,” *Applied Spectroscopy* **48**(7), 1994.

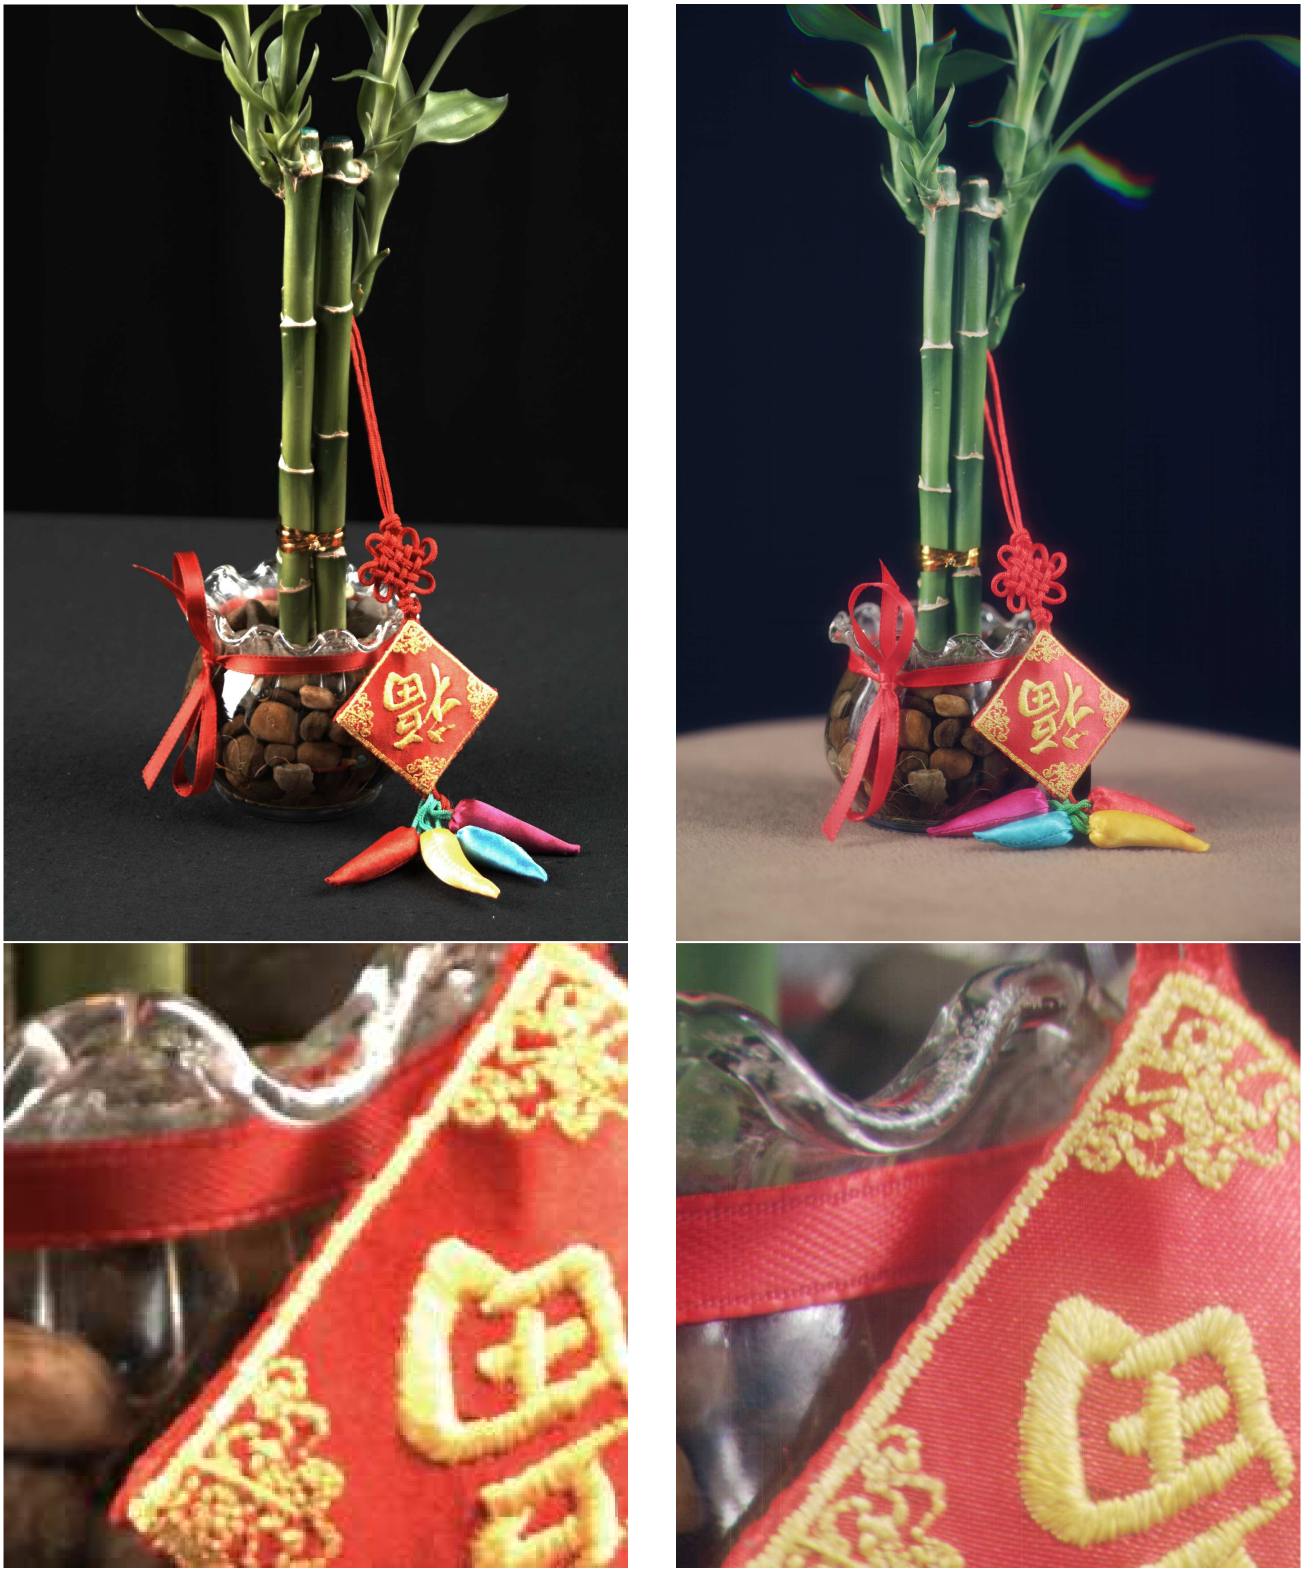

Figure 8. Comparison between the Canon EOS D60 (left) and our scan camera (right). We tried to match the field of view of both cameras as well as possible. Due to the different focal lengths of the lenses used we had to modify the setup and lighting so that the illumination is not directly comparable (see text). The top row shows a slightly cropped image. The bottom row shows a magnification of one region to better compare the resolutions.