Supplementary information: A high-resolution multispectral macro-imager for geology and paleontology

Ryan A. Manzuk[a], Devdigvijay Singh[a], Akshay Mehra[a,b], Emily C. Geyman[c], Stacey Edmonson[d], Adam C. Maloof[a]

[a]*Department of Geosciences, Guyot Hall, Princeton University, Princeton, NJ 08544, USA*
[b]*Department of Earth Sciences, Dartmouth College, Hanover, NH 03755, USA*
[c]*Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena, CA 91125, USA*
[d]*School of Earth and Ocean Sciences, University of Victoria, Victoria, BC V8W 2Y2, Canada*

This supplement is intended to show our multispectral setup in more detail and explain how we mitigate chromatic aberration. We include a figure with annotated computer-aided design renderings of our transmitted and reflected light setups, and details for our light emission spectra. The text begins with background on the problem of chromatic aberration, details our experimental setup and blur modeling calculations, and discusses our final results.

1. The challenge of chromatic aberration in macro imaging

A major limiting factor when increasing spatial resolution and moving to smaller scales for multispectral imaging setups is chromatic aberration, also referred to as chromatic distortion (Jacobson et al., 2013). Because the refractive index of lens elements is wavelength-dependent, it is impossible to design a lens that focuses all wavelengths of light at a single focal distance. This problem becomes especially apparent in macrophotography, where the refractive differences between wavelengths are relatively large compared to the size of the pixels. The resulting depth of field for any given wavelength is narrower than the focal offset between wavelengths (Figure 2B). In traditional red-green-blue (RGB) photography, lens manufactures mitigate this problem by designing apochromatic lenses with two-node aberration curves that focus red, green, and blue light at the same approximate focal distance (Figure 2A). Designing a lens with a seven-node aberration curve is not practical, and would reduce flexibility of the imaging setup.

To execute multispectral macro imaging in the VNIR range with a standard apochromatic lens, workers are left with two options: (1) move the subject in order to image each wavelength at its respective optimal focal distance, or (2) image all wavelengths at a single focal distance and use image processing post-correction to restore the blurry images. For our setup, as part of the Grinding Imaging and Reconstruction Instrument (GIRI), we have the possibility of moving the subject in increments less than 1 μm to capture sharp images
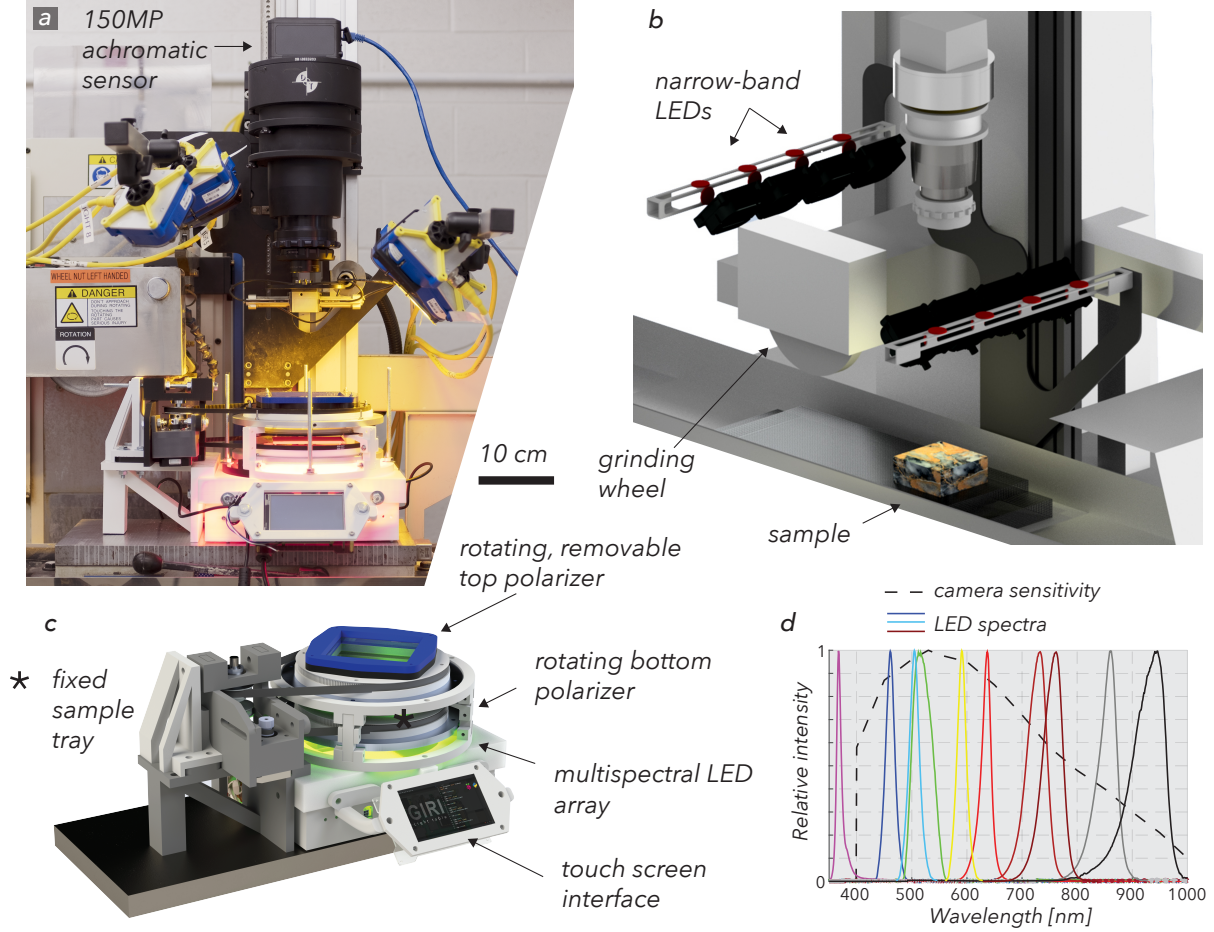
Figure 1: Setups for multispectral petrographic imaging with both reflected and transmitted light. (a) A 150MP achromatic sensor is the basis for capturing multispectral imagery with color provided by either reflected illumination from narrow-band LEDs or thin section transmitted light from a custom light table. (b) The setup for reflected light is integrated with the serial grinding and imaging protocols of GIRI (Mehra and Maloof, 2018), introducing three-dimensional multispectral imagery for sample analysis. (c) Automated multispectral light table with rotating cross-polarizers for sample illumination in transmitted light. (d) Emission spectra for LED lights used in reflected light setup and sensitivity spectrum of the achromatic camera back, adjusted for the 400 nm cut-on filter placed in front of the lens. Of the ten LEDs shown, any 8 can be arranged as in (b) for automated imaging. Preliminary tests to maximize contrast can help select the lights used for any given sample. Note the UV (365 nm) spectrum is outside the camera sensitivity because it is the fluorescence signal of minerals in the visible range of the spectrum that is imaged.

for each wavelength; however, this solution introduces geometric distortions that leave all channels misaligned and slows the image capture protocol. Additionally, we wanted to design a solution that could be reproduced without the need for expensive computer numerical control (CNC) equipment. For this application, we choose to capture all images at a single focal distance where only one wavelength is optimally sharp, and use blur-modeling and image post-correction to restore the other wavelengths. The fine-scale vertical movements possible with our setup are useful in optimization and tuning of this deblurring pipeline.
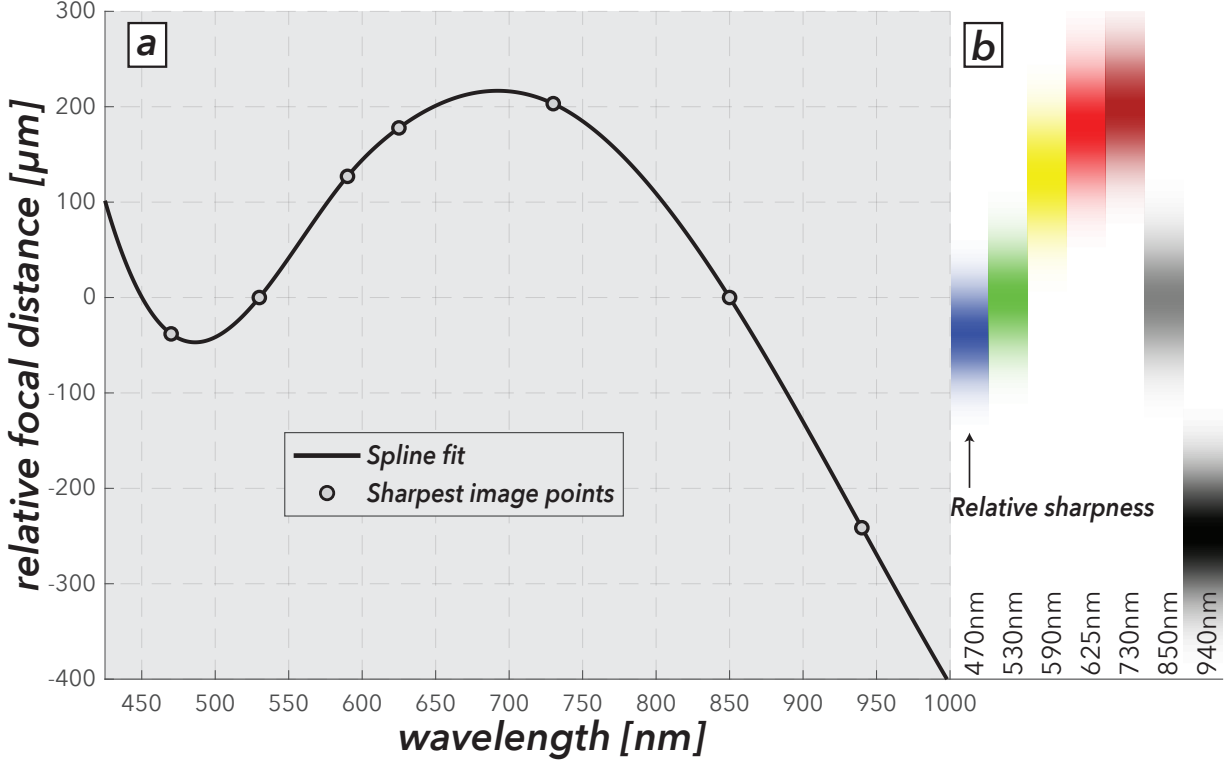
Figure 2: The problem of chromatic aberration in macrophotography. (a) Due to the differences in diffraction angles for each wavelength of light after passing through lens elements, the optimal focal distance is offset between wavelengths by several hundred microns. To produce this chromatic aberration curve, we acquire images for each wavelength at ~12 μm intervals over a vertical range of 2540 μm and estimate sharpness using the total gradient magnitude for each image. The focal distance that produces the largest gradient is the sharpest. The third order polynomial shape of the curve stems from the apochromatic lens used for imaging, which attempts to bring sharp images for red, green, and blue as close as possible to the same focal distance. (b) The range of focal distances at which each wavelength is sharpest do not overlap.

2. Blur modeling and channel sharpening

2.1. Experimental setup

To calibrate our setup, the subject is an Edmund Optics 50 mm x 50 mm white ivory distortion target, consisting of a grid of 125 μm dots with even 250 μm spacing between dot centers. The reported uncertainty is 1 μm for the dot spacing and 2 μm for the dot diameters.

We calibrate this setup for multispectral imaging at a single focal length by modeling the blur kernels that convolve the optimally-sharp channel to produce the blurring effects in all other channels. To adequately constrain these blur kernels, we first image the target illuminated by each light at a range of focal distances. We start the experiment with the green (470 nm) channel and adjust the imaging stage until the focus metric built into the Capture One software reaches its maximum and set that table height as the relative zero-

position for the experiment. After setting the zero-position, we move the table 1270 μm above that position, which is the minimum focal distance tested. We then step the stage down in 12.7 μm increments (increasing focal distance), acquiring an image at every step until the stage is 1270 μm below the zero-position. We repeat this image capture protocol for all wavelengths without resetting the zero-position. We selected the range of table heights based upon previous experimentation to estimate the range of optimal focal distances for all channels, and the increments as the smallest steps at which there were discernible changes in sharpness. Throughout this study, we assess sharpness by taking the sum of the image gradient magnitudes. Because the target used in calibration does not have any activation from UV illumination, we do not include the 365 nm light in the calibration process.

The IQ4 back captures images as IIQ raw files—a proprietary file format unique to Phase One products. For all blur modeling, as well as deconvolution, we export 16-bit TIFF files from the Capture One software without any sharpening or denoising applied. Working with TIFF files that are a strict translation of the sensor values is necessary to model the blur kernels between channels without spurious effects stemming from image processing. In the set of exported images for each channel at each table height, we assess sharpness, and find the table position corresponding to the sharpest image for each channel (sharpest image points in Figure 2a). The final set of images we work with for blur modeling and calibration comprises 49 images with an image from the full set of seven VNIR wavelength channels at all seven optimized table heights.

Blur modeling

Computational methods to deblur images require us to quantify the point spread function of our particular imaging setup—the blur kernel of how a point source of light spreads out when the captured image is not in focus (Zheng et al., 2018). Mansouri et al. (2005) present one way to obtain sharp images from a multispectral image in which only one wavelength is captured in focus. Assuming the blur kernel is small relative to the pixel size, we approximate the blur kernel as a circle of uniform intensity (Mansouri et al., 2005). We adapt the iterative method of Mansouri et al. (2005) to optimize the radius of this circle for each wavelength.

For each multispectral image, there is one sharp image corresponding to the wavelength in focus, six blurry images for the other wavelengths, and one image measuring UV activation. Because the calibration target for this experiment is non-fluorescent, it cannot be used to model the blur for UV activation, but we can apply the same technique for fluorescent samples as needed (the wavelength of fluorescence, and therefore blur kernel will be sample-dependent). For the other six blurry wavelengths, we optimize the blur kernel radius by comparing a computationally-blurred image of the focused wavelength to each image of the

blurred wavelengths by the following process. First, to minimize computational time, we crop our focused-wavelength image to a center 20% window. We contrast stretch all of the cropped 16-bit images by scaling the pixel values to be between 0 and $2^{16} - 1$ (the maxmum 16-bit value) to minimize variations in exposure levels between images. To increase precision in our optimized kernel radius, we upscale this cropped image by a factor of 4 using nearest neighbor interpolation. Given a certain radius, we compute a circular blur kernel with each pixel value in the kernel being initialized to the fraction of the square pixel enclosed within the circle. Next, we normalize the blur kernel such that the sum of all pixel values is one.

Because the artificial blurring process with the simulated kernels smooths out noise that may exist in the actual blurry image, we must be able to estimate and add back any noise present in our images to best represent the blurry wavelengths. We estimate the imaging noise using MATLAB's built-in pretrained image denoising convolutional neural network, DnCNN, by measuring the variance in pixel values between the upscaled sharp image and its denoised counterpart. Next, we convolve the upscaled sharp image using our calculated blur kernel, and then add back simulated Gaussian noise into our artificially-blurred image. The Gaussian noise is parametrized by a mean of 0 and the variance of denoised and noisy pixel values measured earlier. We quantify the error between the artificially-blurred image and the captured blurred image as the mean squared error (MSE) of all pixel values. We iterate this process through a range of radii between 0.1 px and 7 px. Figure 3 illustrates the optimized results for the table height in which green (530 nm) is the focused wavelength. We repeat this process through every table height in which any wavelength is in focus, effectively iterating through all possible combinations of focused and blurry images.

Having optimized the blur kernel for each blurry wavelength at every table height, we restore sharpness in the blurry images using MATLAB's iterative Lucy-Richardson deconvolution method. We optimize the number of iterations to deconvolve each image by minimizing the MSE between the edge spread function of sharp target dots and the deconvolved target dots. This metric rewards restoring contrast within edges of an image and simultaneously penalizes artifacts arising from oversharpening.

Deblurring results

Deconvolving blurry images of our calibration target results in images that are qualitatively and quantitatively closer to the focused wavelength (Fig. 4a-c). The deconvolution pipeline significantly sharpens the target dots when measured by the increase in the sum of gradients over the entire image—a metric that quantifies the increase in edge definition—while ensuring no visible oversharpening artifacts. As seen in Figure 4e, a closer look at the edge spread functions (ESF) of the calibration target dots reveals how the ESF of the
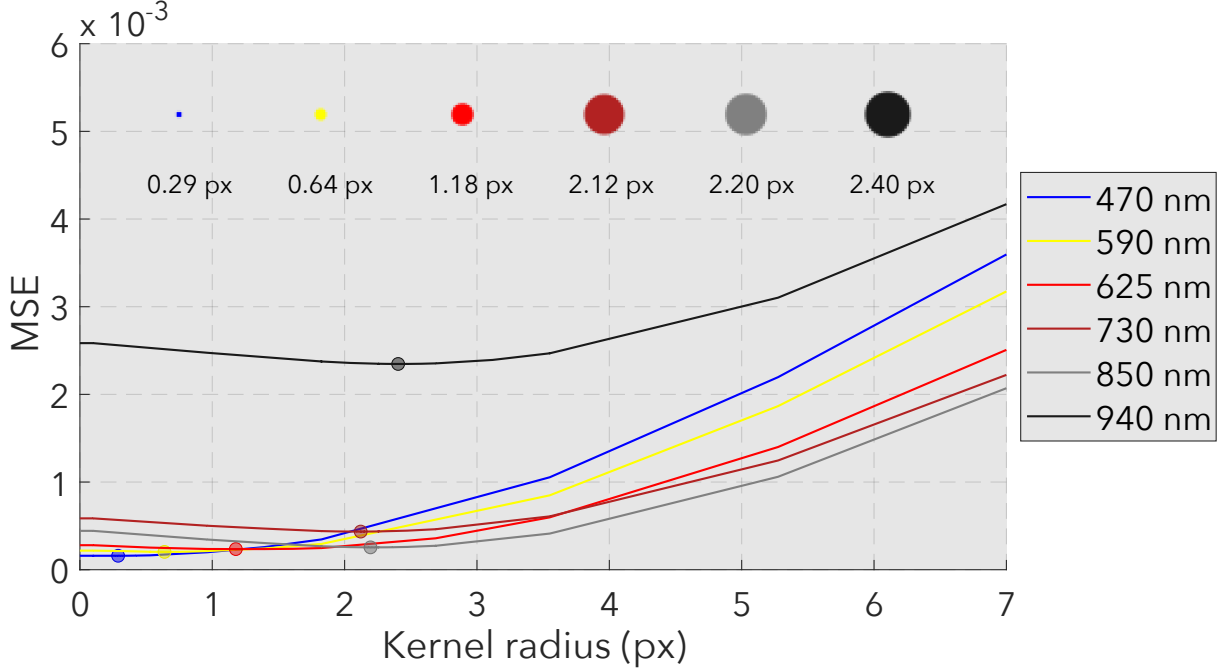
5

Figure 3: The process of optimizing blur kernel radii for each wavelength by comparing an artificially-blurred image (taken at 530 nm at its respective focus height) to the actual blurry image for a given wavelength. The radii that minimizes this difference are the optimal blur kernel radii to use for deconvolving blurry images. The radii are reported relative to the 4x upscaled images, meaning most kernel radii are sub-pixel (3.76 μm) on the original images. Wavelengths further away from the focal distance of 530 nm in Figure 2a require larger blur kernels. The mean squared error (MSE) for 940 nm is relatively higher than the other wavelengths because it has the largest disparity in focal distance (and therefore largest blur kernel), as well as having the lowest signal-to-noise ratio on our sensor.

deconvolved image at 625 nm shifts closer towards the ESF of the focus image at 530 nm. Furthermore, the deconvolved image recovers the contrast within the edge (Fig. 4f) by reducing the width of the blurry ESF (Fig. 4e). It is important to note that the blur observed in Figures 4b and 4h is the blur in a regular RGB image. Thus, all enhancements in sharpness and contrast by deconvolving these multi-spectral images corrects the error due to chromatic aberration in individual pixel values and is an improvement even over traditional RGB images.

With our calibration data set of 49 images (an image of 7 wavelengths at all 7 focused table heights), we choose our imaging table height to be the one that results in the largest overall improvement in sharpness across the entire image set. Since the focal distance for 530 nm lies near the middle of the range of focal distances (Fig. 2A), imaging at this table height results in the smallest blur kernels and the sharpest images across all wavelengths (Fig. 4D).

Having determined an optimal table height for imaging, we find that the sharpness and
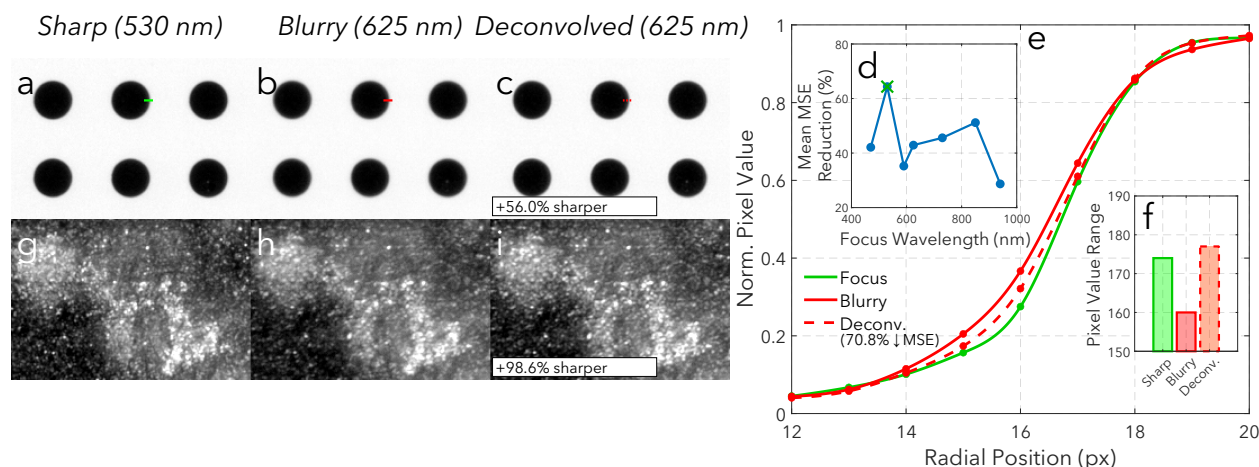
Figure 4: Visual and quantitative comparison of image deblurring between focussed images taken at 530 nm and blurry images at 625 nm. (a)-(c) The deconvolved target image in (c) is visually sharper than the blurry target image in (b), where (a) is a reference for what a focussed image of the target looks like. The sharpness metric in (c), (f), and (i) measures the increase in the sum of gradients over the deconvolved image relative to the blurry image. (d) Focusing at green (530 nm) is optimal for maximizing the average sharpness—as measured by the mean squared error (MSE) between the edge spread functions of the target image—across all wavelengths. (e) The deconvolved edge spread function of a blurry image at 625 nm moves closer to the sharp edge spread function at 530 nm, reducing the MSE by 70.8% over a traditional RGB image. (f) The deconvolved image recovers all of the contrast of the sharp image, as measured by the range of pixel values over the image.

contrast enhancements after deblurring that are observed when imaging the target also carry over to geological samples. For example, deblurring a carbonate sample results in visually better resolving ability and a quantitative increase in edge contrast (Fig. 4g-i). Furthermore, prior to deblurring, information about certain features encoded in the blurry imaging channels, such as the secondary crystals, is spread out into the surrounding pixels. As a result, information in the edges of these features across several wavelengths is lost. Our deblurring pipeline recovers this information by enabling features observed in one wavelength to overlap with other wavelengths. Thus, these deblurred multi-spectral images can be combined to form a vector of quantitative spectral information that can then be used to improve measurements of object size, shape, and density among many other geological features.

References

Jacobson, R., Ray, S., Attridge, G.G., Axford, N., 2013. Manual of Photography. Routledge.

Mansouri, A.A., Marzani, F.S., Hardeberg, J.Y., Gouton, P., 2005. Optical calibration of a multispectral imaging system based on interference filters. Optical Engineering 44, 027004.

Mehra, A., Maloof, A., 2018. Multiscale approach reveals that Cloudina aggregates are detritus and not *in situ* reef constructions. Proceedings of the National Academy of Sciences , 201719911.

144 Zheng, Y., Huang, W., Xu, M., 2018. Optimal PSF Estimation for Simple Optical System Using a Wide-Band
145   Sensor Based on PSF Measurement. Sensors 18, 3552.