Low Spatial Frequency Noise Reduction with Applications to Light Field Moment Imaging

Christopher Madsen
Stanford University
cmadsen@stanford.edu

Abstract

This project involves the implementation of multiple denoising algorithms on images created using a Light Field Moment Imaging (LMI) approach. LMI allows for a shift in perspective using two images captured through a single lens, but suffers from color cast created during the required mathematical approximations. This paper proposes the use of several denoising filters to remove the noise prior to construction of the photo. Quantifying the ability of each denoising filter will be done through the use of PSNR, computation time, and gauge of overall image quality.

1. Introduction

Developed by Orth and Crozier in 2012, LMI provides a method of mimicking depth and perspective through a single lens by simulating the light field mathematically using two captured images taken slightly out of focus from one another, with a focal plane difference (Δz) typically on the order of μm. In its current state, LMI provides intriguing results that have already found use in multiple microscopy applications. However, the method remains somewhat in its infancy. In the introductory paper [1], Orth states that further research should be conducted on the Gaussian distribution assumption used in developing the light field matrix. In addition and most importantly for this paper, high levels of noise are present in the extracted photos that can cause a significant color shift, as shown in Fig. 1. At a low enough value of Δz, this noise can bury the signal, as shown in Fig. 2. This project seeks to implement a noise reduction in the frequency domain that will improve color reproduction while minimizing the effects of denoising on the apparent perspective shift.

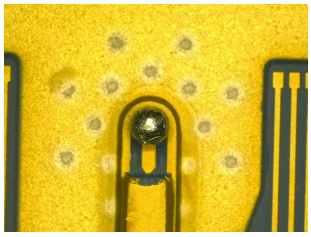

(a) Original image with normal levels of additive noise.

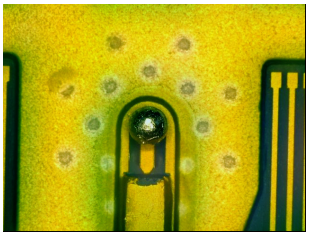

(b) Derived image showing severe color cast due to noise.

Figure 1: Original and LMI derived images showing differences in noise level.

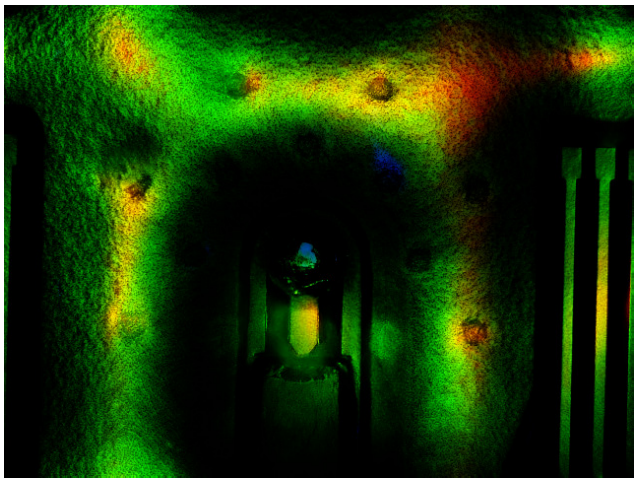

Figure 2: Extremely low PSNR (5.12 dB) resulting in near complete masking of underlying image, Δz = 0.2 μm

2. Background

2.1. LMI

LMI utilizes an approximation of the light field within a scene to provide simulated perspective shifting. By determining the angular moments of each ray, an approximation of the scene at a given angle θ can be made math-

ematically. A full derivation of the mathematical system is presented in the report [1], but a brief overview will be given here in order to provide some insight into the noise reduction algorithm. The foundation of LMI is based on the parameterization of the light field within a given scene. Here, Orth represents the density of rays at a given point as $\bar{L}(x,y,z,\tan\theta_X,\tan\theta_Y)$. A simple assumption about the physics of light at a given point and some mathematical manipulation yields a Poisson equation:

$$\frac{\partial I(x,y;z)}{\partial z} = -\nabla_{\perp}^{2} U(x,y;z) \tag{1}$$

Where $U(x,y;z)$ is defined as the scalar potential of the light field. This equation is solved in Fourier space, implemented in the form of a filter H:

$$H(f_x,f_y) = [-4\pi^2(f_x^2+f_y^2)]^{-1} \quad \text{for } (f_x^2+f_y^2) \neq 0 \tag{2}$$

$$H(f_x,f_y) = 1 \text{ for } (f_x^2+f_y^2) = 0 \tag{3}$$

$$U = \mathscr{F}^{-1}[H \times \mathscr{F}\{(I_1 - I_2)/\Delta z\}] \tag{4}$$

Where f_x and f_y are the spatial frequencies in the x and y-directions, respectively. Here, the derivative is approximated by operating under the assumption that Δz is a very small value (typically on the order of μm). While this approximation works well for perspective shifting, it allows for noise to be amplified in a manner similar to that seen in inverse filtering and deconvolution: a small value in the denominator results in a large amplification of the signal. Ultimately, the amplified noise presents itself as a color shift in the low spatial frequency areas of the final photo. As the value of Δz decreases, the noise can completely overwhelm the signal, resulting in effects such as that shown in Fig. 2.

2.2. Filtering

In order to minimize noise in the final image, this project proposes a noise reduction algorithm applied to the image difference matrix J, derived from Eq. 4:

$$J = \mathscr{F}^{-1}\{H \times \mathscr{F}\{(I_1 - I_2)\}\} \tag{5}$$

Which can then be applied to Eq. 4 in lieu of the I_1/I_2 difference equation:

$$U = \mathscr{F}^{-1}[H \times \mathscr{F}\{J/\Delta z\}] \tag{6}$$

At this point, Δz is applied, but is not as effective in noise amplification due to reduction by the algorithm.

This method of noise reduction poses some inherent problems. First, since the noise doesn't manifest in the spatial domain as the high frequency "speckling" typical of a normal distribution, clean up in the spatial domain proves difficult as a typical measure, compute, and adjust filter will likely prove ineffective. As such, filtering will take place in the frequency domain where it can be assumed that noise manifests itself in both the low and high frequency portions of the spectrum, relatively separate from the majority of the underlying image. Second, overfiltering of the difference matrix can result in a near-zero gradient matrix being extracted. When this occurs, no perspective shifting will occur and the PSNR will approach infinity, i.e. the output photo will perfectly match the input photo I_2. With these complications, filtering in this application has to be a balance between computation time, PSNR, and image quality.

While the filtering algorithm is going to take place wholly in Fourier space, the filters themselves will be created in the spatial domain. By taking a 2D FFT of the image and shifting low frequencies toward the center of the matrix, a spatial bandpass filter (BPF) can be configured and applied to remove high frequency noise (speckling) toward the edges of the transformed image and low frequency noise (color shifting) concentrated at the center. In order to experiment with the maintenance of simplicity and reduction of the required computation time, three of the implemented filters will be an ideal BPF, a Gaussian BPF, and a Butterworth BPF. The ideal BPF manifests itself as two simple circles imposed on a black background. At specific radii, the circles will attenuate both low spatial frequency data and high spatial frequency data located in the corners of the FFT. Gaussian filters are rare in signal processing, but find use in image processing due to their ability to simulate natural blurring effects. This will not be as critical in the frequency domain as it would be in the spatial domain, but the filter's simplicity and minimal computation tax warranted some experimentation. While similar in nature to the Gaussian BPF, the Butterworth filter allows for a flatter frequency response in the more critical low frequency passband which will constitute a large portion of the image, resulting in higher image integrity.

The fourth and final filter will be an implementation of Frequency Domain Wiener Filtering (FDWF) [2]. While Wiener filtering in general is suited for inverse deconvolution and typically finds use in the spatial domain, this particular algorithm is well suited for this task due to its adaptive nature. To avoid amplification of noise, approximations of the noise level are made prior to masking and filtering of the primary image. Additionally, in order to reduce the required computation time, the image is broken up into zones rather than relying on a pixel-by-pixel computation. The algorithm is purported to be highly effective in noise reduction while maintaining the integrity of edge details and fine textures.

3. Methods and Procedures

Due to the relationship between the perspective shifting and the PSNR of the resulting photo, PSNR was considered to be of little use for a definitive assessment of filter quality. However, by attempting to set the PSNR the same in all filters and comparing other characteristics such as computation time and image quality, a solid quantitative analysis of each filter type can be performed. To perform just this, the filter sigmas, cutoff frequencies, window sizes, and other parameters were adjusted and run until the PSNRs of each closely matched each other. This method caused a few issues, which will be discussed in detail in the results section. The PSNR was chosen as a moderate improvement that removed the color shift, but still maintained the perspective shifting algorithm. This also allowed inspection for artifacts such as ringing that is present in the FDWF at high β (division factor) levels. Through experimentation, a PSNR of 26.0 dB was found to be a good compromise when looking at the aforementioned characteristics. When discussing each filter, this value will be the one sought after when developing the filtered images and will be responsible for driving each filter's parameters.

It was noticed during the filter development process that the color cast could be removed, but with the side effect of an excessive amount of noise in the blue channel that revealed itself in the low spatial frequency regions of the image. It is believed that this noise is due to lower spatial frequency pixel pairs in the blue channel that are slightly higher in magnitude than seen in the red and green channels. This would push them just outside of the blocked zero frequency area in the middle. To account for this, each filter was given a higher window value for the blue channel than for the others. The listed values for the low frequency portion of the filter will correspond to the values used in the green and red channels.

3.1. Ideal BPF

The ideal bandpass filter is a simple design that is implemented by placing two circles in the middle of a matrix sized the same as the image to be filtered. Though basic in nature, the parameter (specifically the circle radii) proved to be incredibly difficult to control for and made for poor results, which will be discussed in detail later. Figure 3 shows the spatial representation of the filter. Note that the vast majority of the image is completely unattenuated while only a small section of the corners and center receive any treatment from the filter. In order to provide a PSNR approximating 26.0 dB, filter radii of 395 pixels (outside) and 1.6 pixels (inside) were used.

Figure 3: Ideal filter, 640x480

Figure 4: Gaussian filter, 640x480

3.2. Gaussian Filter

Gaussian filters find use in spatial image processing, especially when a natural blurring effect is sought in an image. They have very little application to spatial denoising applications, but work well as bandpass filters in the frequency domain, which is why the method was chosen as a starting point. Construction of the Gaussian filter is simple, requiring only a filter size and standard deviation (σ). In order to achieve the required PSNR of approximately 26.0 dB, σ values of 2.2 and 200 were found to work well. A spatial representation of the Gaussian filter is shown in Fig. 4.

3.3. Butterworth Filter

The Butterworth filter finds frequent use in signal processing due to its optimal flat response in the passband. This is ideal for a bandpass filter application as it would

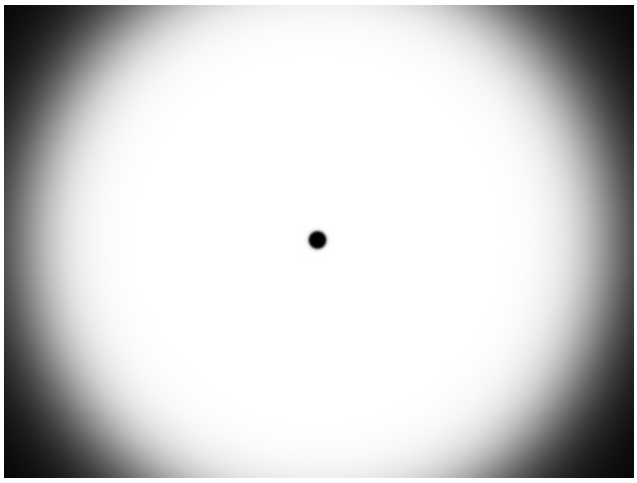

Figure 5: Butterworth filter, 640x480

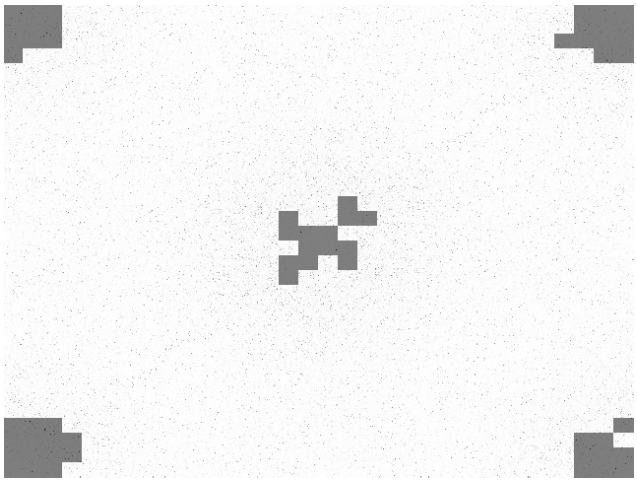

Figure 6: FDWF filter, 640x480

be preferable to maintain as much of the low frequency signal information as possible since this is where the majority of the image lies. A downside of the Butterworth filter is its increased computation time over the simpler Gaussian distribution. Figure 5 shows the spatial representation of the Butterworth filter. It's immediately evident that the Butterworth filter provides an extremely flat response in the passband area where the majority of the image will be concentrated in Fourier space. There is a harder cutoff toward the higher frequency corners and immediate center, where a significant portion of the noise is located. This is far different from the constant gradient nature of the Gaussian filter and will likely provide a better representation of low spatial frequency areas in the final image. Design of the Butterworth filter's characteristics required a balance between PSNR, computation time, and proper shaping of the transition band. Too hard a cutoff is not necessarily a wanted feature, as will be discussed in the ideal LPF results. However, a 6th order filter showed a substantial reduction in noise and preservation of fine details within the image. To reach a PSNR of 26.0, the cutoffs were found to be 1.5 and 300 pixels.

3.4. Frequency Domain Wiener Filter

Application of the FDWF proved to be more complicated than the other three filters. The FDWF filter operates on the assumption that noise is heavily concentrated in the corners of an image in Fourier space, which makes sense for most applications. However, here there is an assumption that there is noise concentrated both in the high and low spatial frequencies. In order to get the filter to operate under this assumption, two FFTs were performed: one shifted with the low spatial frequencies at the center of the image, and another unshifted with the low frequencies concentrated in the corners. These two images were then combined to form an FDWF bandpass filter, whose spatial representation is shown in Fig. 6. Dividing factors of 2 and 5 were found to provide a good balance in the wanted characteristics.

One benefit to this filter is the omission of the blue channel noise compensation that was used in previous designs. The algorithm is able to detect and eliminate this noise in the channel without the need for adjustment of the dividing factor value, β.

4. Results

As the perspective is shifted, the PSNR between the shifted image and the original image understandably decreases. At an angle of -0.0616 radians, the maximum PSNR while maintaining perspective shifting is approximately 32 dB. With no noise filtering in place and the value of Δz fixed, the minimum PSNR is approximately 19 dB. Approximately halfway between these two values is 26 dB, the value chosen to gauge filter effectiveness. This was done to ensure that noise reduction was implemented without completely removing or heavily attenuating the intended perspective shifting. Because PSNR cannot be used as a decisive quantification method, computation time and the admittedly subjective image quality comparison must be used. The original image, available for comparison, is provided in Fig. 7. The noisy image is shown in Fig. 8. To create this noise, a Δz value of $1.5\,\mu\mathrm{m}$ was used. Both images depict the RF input of an ASIC produced by Keysight Technologies.

4.1. Ideal BPF

The ideal BPF, with image shown in Fig. 9, was implemented using the parameters found in the Methods and Procedures section. While the green color cast has been reduced, immediately noticeable is further induced color cast

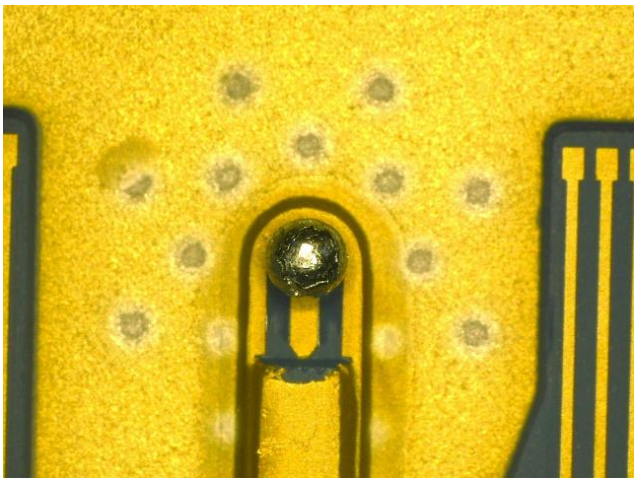

Figure 7: Original image used for PSNR comparison, 640x480

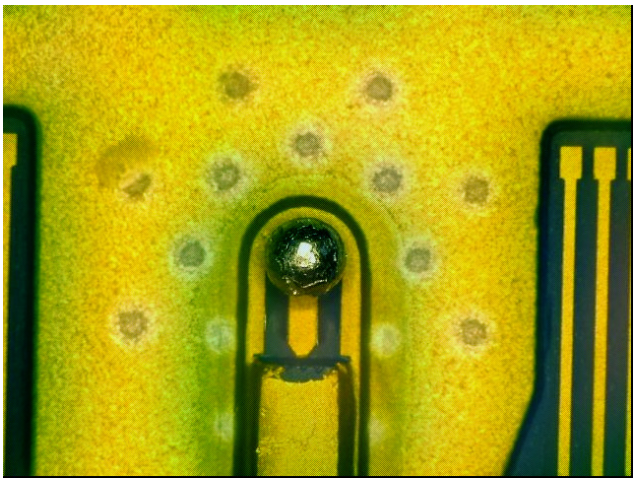

Figure 8: Derived image showing color cast due to noise.

near the center of the image, in addition to noise in the blue and red channels in low spatial frequency areas surrounding the center structure. Noise is also evident on the gold layer toward the top of the image, which gives the appearance of a more textured structure than actually exists there. Some green can still be seen around the structures on the left and right sides of the image. Additionally, adjustment of this filter proved difficult. Fine adjustment to remove noise required the use of non-integer values, something unique to this filter. Minute changes on the scale of 10^{-1} resulted in substantial differences in the filter output. Removal of noise in the red and blue channels resulted in a significant jump in the PSNR, but a reduction or complete elimination of the perspective shift.

PSNR	Computation Time
26.87 dB	0.09 s

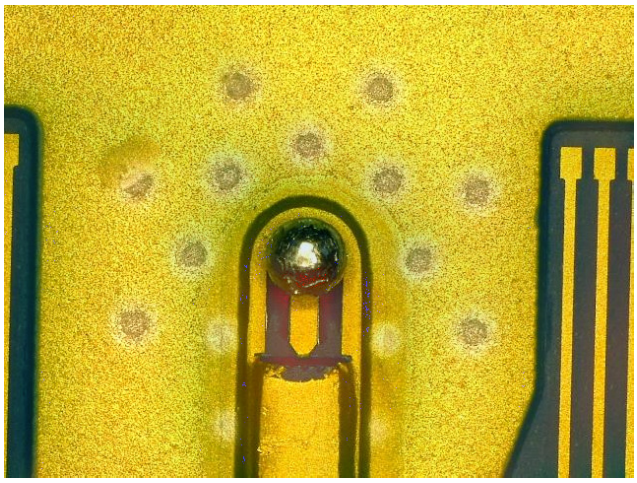

Figure 9: Ideal BPF applied to image, 640x480

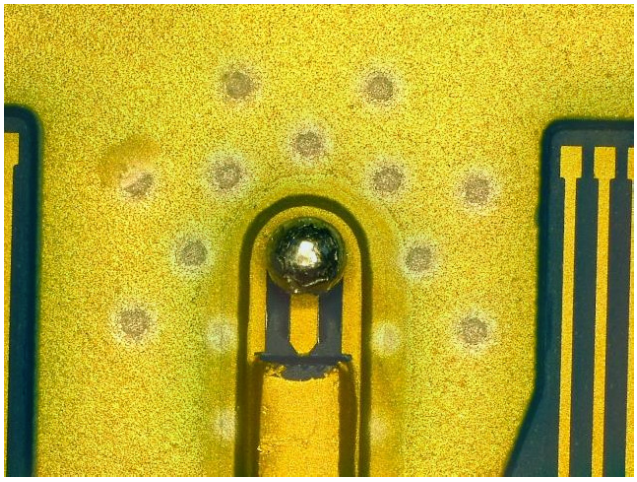

Figure 10: Gaussian BPF applied to image, 640x480

4.2. Gaussian Filter

The Gaussian filter, shown in Fig. 10, fared better than the ideal BPF. The green color cast has been reduced to a minimal level, with some still visible in the low spatial frequency areas surrounding structures. As with the ideal BPF, some level of noise has been introduced which gives an illusion of an enhancement to sharpness or a change in texture. Despite the uneven attenuation in the passband that the Gaussian filter provides, there appears to be only a minimal loss of detail in critical areas. An example of this is the solder ball in the middle of the photo, whose features have been slightly muted by the filter.

PSNR	Computation Time
26.42 dB	0.13 s

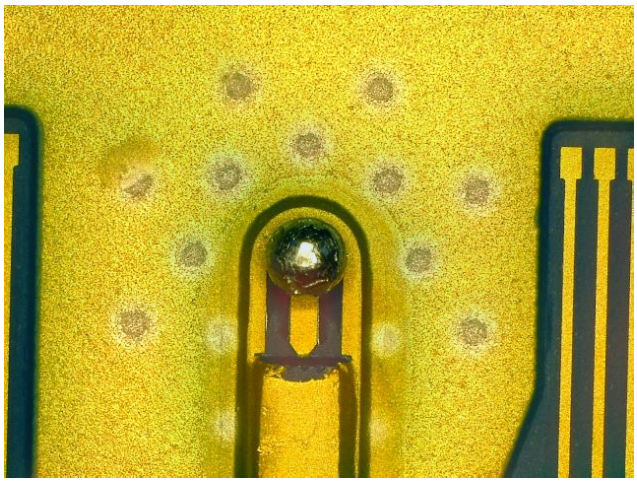

Figure 11: Butterworth BPF applied to image, 640x480

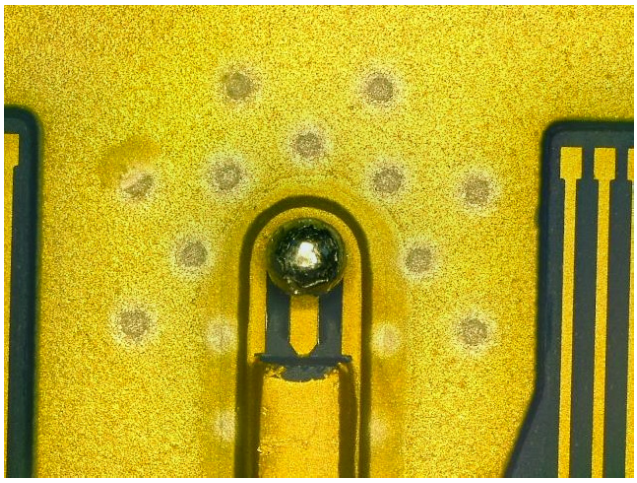

Figure 12: Bandpass FDWF applied to image, 640x480

4.3. Butterworth Filter

The Butterworth filter, shown in Fig. 11, appears similar in nature to the Gaussian filter, but has some critical differences. The green color cast has been reduced, but not to the same level provided by the Gaussian filter. However, the Butterworth filter has done a better job of maintaining details, such as on the solder ball in the center of the photo. Unfortunately, it seems to have enhanced some features, such as the imperfections on the gold trace just below the solder ball. While these show up accurately in the original photo, they look washed out in the filtered image. The substrate features seem neither enhanced nor muted from the Gaussian filter. Given the computation time involved, it actually appears the Gaussian filter provides a more true-to-form result.

PSNR	Computation Time
26.40 dB	0.85 s

4.4. Frequency Domain Wiener Filter

The FDWF, shown in Fig. 12, appears at first glance to be the best representative of the original photo. The green color cast has been reduced to a minimum and there doesn't appear to be a notable level of noise or induced casting. Details are maintained in most of the image, with a small level of blurring expected when high frequency components are removed. When comparing this and previous images to the original, contrast appears to have suffered slightly from the original, as has brightness. Both of these issues are correctable in post processing and should not be considered problematic. Computation time of this filter is half of what was required for the Butterworth filter. This includes the extra FFT derivation and multiple calls to the filter function, making the required time even more impressive.

PSNR	Computation Time
26.75 dB	0.46 s

4.5. Overall Results

In reviewing the results for each filter type, it's readily apparent that there are some critical differences in the filter outputs. Overall, the FDWF appeared to provide the highest attenuation of the color cast while remaining true to the original photo in terms of color and details. The FDWF did not experience any extraneous noise or color casting that was associated with the three previously discussed filters. Coming in close second is the Gaussian filter, which does a surprisingly decent job of maintaining the critical details of the image (with some level of blurring to fine detail), but suffers from an excess level of residual color cast.

5. Conclusion

In this project, a method for denoising LMI images was developed and four different filter types implemented, including the ideal BPF, Gaussian BPF, Butterworth BPF, and a modified bandpass FDWF. All filters showed themselves to be effective, with some providing a higher image quality than others. Specifically, the ideal BPF induced a notable color cast after reducing the cast created by the implementation of LMI while the Gaussian filter displayed a minor color cast but overall excellent quality in terms of image details. The Butterworth filter showed minimization of the color cast and preserved more details than the Gaussian filter, but didn't have the same level of attenuation as the Gaussian. Finally, the FDWF displayed an extremely high quality photo with preserved details and no notable color cast over what the LMI algorithm produces.

This project has shown that through the use of simple filtering algorithms, noise reduction in LMI applications can

be performed with minimal use of computation time or sacrifice of image quality. This report has shown only a small sample of bandpass filters used in a signal processing application. Research into other filtering algorithms is warranted to determine which is best in terms of computation time and image quality.

References

[1] A. Orth and K. B. Crozier. Light field moment imaging. *Opt. Lett.*, 38(15):2666–2668, Aug 2013.

[2] S. Sari and T. Shimamura. Frequency domain wiener filter for image denoising: Derivation of a new power spectrum estimation method. *Journal of Signal Processing*, 16(1):79–85, 2012.