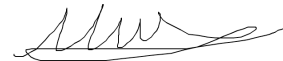


Automating Mode Detection in Quantum Imaging

A thesis submitted in partial fulfillment of the requirement
for the degree of Bachelor of Science in
Physics from the College of William and Mary in Virginia,

by

Caitlyn P. Marat



Advisor: Prof. Eugeniya Mikhailov



Prof. David Armstrong

Williamsburg, Virginia

May 9, 2022

Contents

Acknowledgments	iii
List of Figures	iv
Abstract	v
1 Introduction	1
2 Theory	2
2.1 Quantum Noise	2
2.2 Coherent vs. Squeezed States	3
2.3 Modes	4
2.4 Variance	5
2.5 Masks	7
3 Experimental Procedure	8
3.1 Generating Variance Maps	8
3.2 First Approach: Convolution	10
3.3 Second Approach: Toggling	11
4 Global Toggling Optimization	14
4.1 Testing Larger Masks	14
4.2 Decreasing Rate of Toggle	15
4.3 Clustering	18
4.4 Limiting Pixels in Mask	18
4.5 Applying an Exclusion Mask	19
4.6 Probabilistic Toggling	20

4.7	Additional Approaches	21
5	Results	22
5.1	Squeezed State	22
5.2	Coherent State	25
5.3	Imaging Objects	25
6	Conclusion	29
	References	32
A	Appendix: Simulated Annealing	33

Acknowledgments

I would like to thank the William & Mary Physics Department for the invaluable skills and experience it has afforded me. In particular, I thank Professor Mikhailov for his guidance and for teaching me how a physicist approaches research, as well as his willingness to humor my aspirations as an amateur puppeteer. I especially thank Savannah Cuzzo for her extensive research and expertise and for providing me with the background and data that allowed me to take on this project. Finally, I would like to thank Professor Novikova and the rest of the Quantum Optics Lab for supporting me through this project and exposing me to new fields of research.

List of Figures

1	“Ball” diagram depicting coherent and squeezed states taken from [8]	4
2	Examples of cylindrical mode patterns from [10]	5
3a	Variance map for coherent state without using a mask	8
3b	Variance map for squeezed state without using a mask	8
4	Sample timeline of measurement period with separate intensity maps for each beam	9
5a	Variance map produced by convolution technique on a coherent state	11
5b	Variance map produced by convolution technique on a squeezed state	11
6	Optimization of a 10 by 10 mask over 200 cycles	12
7	Optimization of a 20 by 20 mask over 200 cycles	13
8	Optimization of a 40 by 40 mask over 200 cycles	13
9a	10 by 10 mask resulting from the run shown in Figure 6	14
9b	20 by 20 mask resulting from the run shown in Figure 7	14
9c	40 by 40 mask resulting from the run shown in Figure 8	14
10	128 by 128 mask achieved after 750 cycles on the squeezed state . . .	15
11	Optimization process for 128 by 128 mask	16
12	Optimization process for a 128 by 128 mask with toggle rate inversely proportional to cycle number	17
13	Number of toggles over 1000 cycles.	17
14	Best mask after 1000 cycles using 5 by 5 pixel clusters and a 50 percent cap	19
15a	Mask after 1000 cycles with 10 percent cap and exclusion mask . . .	21
15b	Mask after 1000 cycles with 20 percent cap and exclusion mask . . .	21
16	Classical intensity image for the squeezed state	23

17	Variance map obtained from the convolution approach for the squeezed state	24
18	Sum of 10 masks generated by the optimization on the squeezed state data	24
19	Classical intensity image for the coherent state	25
20	Variance map obtained from the convolution approach for the coherent state	26
21	Sum of 10 masks generated by the optimization on the coherent state data. The color bar indicates the number of masks that share a pixel.	26
22a	Classical intensity image of partial rectangular knife-edge	28
22b	Variance map obtained from the convolution approach for the partial knife-edge image	28
22c	Sum of 20 masks generated by the optimization on the partial knife-edge image	28
23a	Classical intensity image of rectangular knife-edge blocking bottom right corner of the image	29
23b	Variance map obtained from the convolution approach for the knife-edge image	29
23c	Sum of 80 masks generated by the optimization on the knife-edge image	29
24a	Classical intensity image of a cross in the center of the image	30
24b	Variance map obtained from the convolution approach for the cross image	30
24c	Sum of 80 masks generated by the optimization on the cross image .	30

Abstract

Quantum imaging uses non-classical properties of light to image objects in scenarios where classical imaging is insufficient, such as in low-light imaging. Rather than classical intensity measurements, which are highly sensitive to dark noise and background illumination, quantum imaging uses the quantum noise distribution as the metric for depicting the interaction of light with an object. The approach discussed here produces normalized variance maps that can be manipulated to reveal correlated features of the images with corresponding normalized variance. We apply masks of various shapes and sizes to variance maps as a processing technique to bring out these features, known as modes. We then develop an algorithm that optimizes mask selection to automatically apply the best mask to reveal modes. After several rounds of tuning the optimization, we find a ring-shaped mode in our primary squeezed-state test data; we also find that this mode mutates when a blocking object is introduced. Using two approaches, we successfully image an object in a low-light setting using only quantum noise.

1 Introduction

Quantum imaging utilizes the quantum properties of light to overcome many of the limitations of classical imaging for improved imaging performance. This quantum approach to imaging examines the correlations between photons rather than directly measuring intensity using the number of photons [1]. This allows improvements in imaging scenarios with high levels of noise or background illumination and enables low-light imaging, where the number of incident photons is too low for classical detection.

Classical imaging can be limited in resolution, signal-to-noise ratio, contrast, and spectral range [2]; the classical technique is sensitive to background light and sensor noise. Background illumination contaminates an image and blurs the boundary between the blocked and unblocked regions; this makes imaging particularly difficult in low-light conditions. Additionally, a camera or detector possesses an inherent dark noise that can obscure fluctuations in incoming light, an especially pertinent factor when imaging with few photons.

Quantum imaging enables imaging at extremely low-light conditions where background illumination and sensor noise may obscure the few photons used for imaging. Low-light imaging is instrumental for biological imaging, where samples may be highly sensitive to incoming photons, as well as in environments that induce noise, such as in atmospheric conditions like fog or in biological tissue [2]. A related method of imaging is quantum ghost imaging, where an image is gathered by detecting spatially-correlated sets of photons of which only some have actually interacted with the sample [3].

Rather than using intensity measurements as in prior imaging techniques, the method of quantum imaging discussed here examines the quantum noise distribution

[4]. We measure the noise of the field, rather than its mean value, to obtain information about an object being imaged [5]. Using a squeezed vacuum field with few photons, we can study the spatial distribution of the noise and the quantum fluctuations caused by an object. We use a coherent state laser beam as a calibration because the coherent state should maintain fluctuations at the shot noise limit. Meanwhile, a squeezed quantum state reveals fluctuations below the shot noise level, allowing us to image using just this noise data [5].

We search for correlated regions of amplified variance within maps of quantum noise. These correlations reveal modes, which are regions that fluctuate independently of other parts of the system and that encode spatial properties of the field [6]. We use spatial masks to identify modes by selecting pixels over which to calculate the variance. If we select a group of pixels that amplify each other to produce a large variance, this group may indicate a mode. In this paper, we calculate the variance on a quantum noise map with different masks and optimize the search for ideal masks to locate these quantum modes. We develop two approaches to generate and apply masks and compare their effectiveness for low-light imaging applications. For results from both approaches, see Section 5.

2 Theory

2.1 Quantum Noise

Quantum noise is an inherent property of quantum mechanics. Due to the Heisenberg uncertainty principle, which assigns a minimum uncertainty to the measurement of pairs of complementary observables, the accuracy of quantum measurement is limited and subject to quantum noise [7]. In classical imaging, this source of noise is negligible as classical noise sources— laser intensity fluctuations, temperature instability, vibrational noise, etc.— greatly outweigh quantum noise.

However, in low-light imaging, where we rely on quantum fluctuations to extract information, understanding and manipulating this quantum noise becomes extremely important. In quantum imaging, we consider the uncertainty principle as it applies to the product of the amplitude and phase of light. We refer to the standard environment-induced noise from vacuum fluctuations as the “shot-noise limit,” which emerges from the quantum-mechanical nature of light [7]. The shot-noise limit is the minimum amount of noise possible in a classical measurement. This limit can be circumvented using quantum states to reduce noise in an observable and enhance the sensitivity of light-based measurements.

2.2 Coherent vs. Squeezed States

The technique used to manipulate the noise distribution in the phase and amplitude of light is called squeezing. Shot-noise (or *coherent state*) is characterized by equal uncertainties in phase and amplitude; squeezing allows us to decrease one of these uncertainties at the expense of the other. This allows the uncertainty in a select characteristic to decrease below the shot-noise limit in *squeezed states* [8].

Figure 1 presents a standard method of visualizing coherent and squeezed states. In these graphs, the distance from the origin represents the amplitude, while the angle is the phase. δX_1 is the uncertainty in the phase, δX_2 is the uncertainty in the amplitude, and δX_{sq} is the square of δX_1 . In a vacuum, the amplitude is 0 and the ball is centered on the origin. The coherent diagram is circular, showing the equivalent uncertainties in amplitude and phase. The squeezed diagram shows a phase-squeezed state, with compression in the phase direction and corresponding expanding of uncertainty in amplitude.

For this study, we utilize the squeezed state to reveal quantum-mechanical properties that are otherwise obscured by the shot-noise limit. We extract the quantum

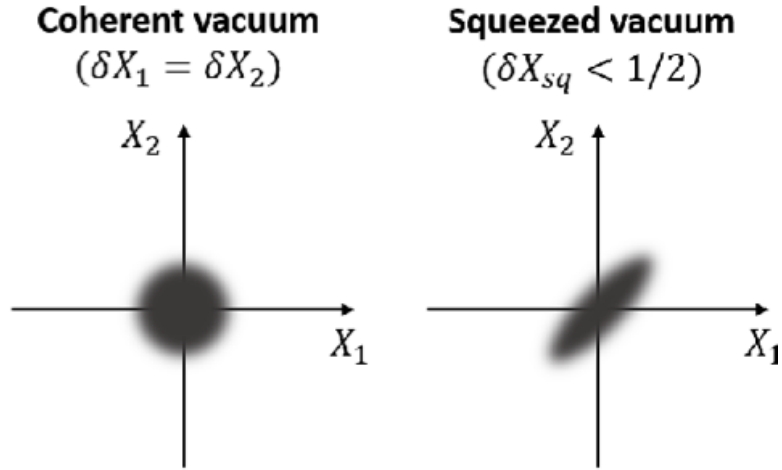


Figure 1: “Ball” diagram depicting coherent and squeezed states taken from [8]

fluctuations and use them to discern modes and recreate classical images with very little light. For information on how squeezed light is generated to collect the data used here, see [9].

2.3 Modes

In this study, we use the quantum-mechanical properties of the squeezed light to locate spatial modes. A mode is a particular pattern in the electromagnetic field. These spatial modes describe the transverse spatial distribution of the laser output, determining the intensity and variance distributions on the cross-sections of the laser beam [10].

The lasers we study here have cylindrical symmetry, so the shapes we look for are combinations of Laguerre polynomials and a Gaussian beam [10]. These will show up as circularly-symmetric patterns on the sensor. Figure 2 shows examples of several cylindrical modes. We look for modes of these shapes with corresponding variance in quantum noise data to find symmetries and characterize the incoming light. This will help to further understand the quantum imaging scheme and identify changes in




Figure 2: Examples of cylindrical mode patterns from [10]

the modes when a blocking object is introduced that may assist in interpreting the images.

2.4 Variance

In this approach, we compare the coherent state shot noise to the squeezed state noise obtained by directing a quantum probe at a camera over a length of time. The key parameter for analysis is the normalized variance. For each pixel, we find the variance between frames; a frame is an individual measurement or “snapshot” at a moment in time. By measuring the intensity at many frames, we track the change in each pixel over the measurement period.

The variance for each pixel is calculated over all frames by

$$V = \sum_{i=1}^{N_{frames}} (N_i - \bar{N})^2 * \frac{1}{N_{frames}} \quad , \quad (1)$$

where N_i is the number of incident photons on a pixel and \bar{N} is the mean number of

photons. Normalized variance is then

$$V_{norm} = \frac{V}{\overline{N}} \quad . \quad (2)$$

A key benefit of using variance rather than the traditional intensity measurement is in the signal-to-noise ratio, which compares the signal being measured to the background noise to describe the uncertainty of the measurement. Using photon count, the signal-to-noise ratio is

$$SNR = \frac{\overline{N}}{\sqrt{\overline{N} + 2(\Delta N_d)^2}}, \quad (3)$$

where N_d is the dark noise [9]. When imaging with very few photons, the dark noise overtakes the average number of photons, creating a high level of uncertainty. On the other hand, the signal-to-noise ratio using variance is

$$SNR = \frac{V - 1}{\sqrt{2 + V^2 + V^4}}; \quad (4)$$

in this calculation, dark noise is eliminated from the signal-to-noise ratio [9]. Therefore, by measuring variance rather than photon number, a high signal-to-noise ratio can be achieved.

For an incoming coherent state, the expectation value of the variance $V = \overline{N}$, the expectation value of the average photon count, so the normalized variance V_{norm} should be 1. On the other hand, for a non-classical squeezed state, $V \neq \overline{N}$, so the normalized variance is not equal to 1 [2]. If an object is placed between the squeezed probe and the sensor, the blocked area is replaced with a coherent vacuum where the normalized variance is expected to be 1, creating a region of variance 1 on the variance map.

2.5 Masks

We calculate variance over subsets of pixels determined by masks. A mask is a two-dimensional matrix that selects pixels over which to calculate the average variance and blocks the intensity from other pixels. Instead of calculating variance for each pixel, we calculate variance for each mask. When we apply a mask to a certain pixel, we center the mask over that pixel and calculate the variance of the region surrounding the pixel using that mask. If pixels included within the mask are correlated, the total variance over the mask is amplified. We hope to find that if the pixels in a certain shape are correlated, the mask that shows the highest variance should mimic that shape to select those pixels, indicating that a mode of that shape has been found.

The masks used here are binary matrices; pixels with a value of 1 are included in the variance calculation, while pixels with a value of 0 are omitted. While we group pixels using masks of different shapes and sizes, we additionally apply a circular exclusion mask to cut off the edges of the map. Because the laser beam focuses on the center of the images, effectively zero photons are received in the corners; this causes division by zero in Equation 2, so the normalized variance in these corner regions blows up. We apply the circular exclusion mask to restrict the range of study to just the circular beam.

Figures 3a and 3b show variance maps where all pixels are included in the calculation for coherent state and squeezed state, before a mask is applied. We apply a circular exclusion mask to assign a variance of 0 to the corners of the variance map. In the region of interest, there is little differentiation between the two variance maps; while the coherent state is expected to have a variance of 1, we expect the squeezed state to have different quantum statistics. These can only be seen by calculating the variance over select groups of pixels.




Figure 3a: Variance map for coherent state without using a mask




Figure 3b: Variance map for squeezed state without using a mask

We test masks of different sizes and shapes to search for the mask that results in the highest average variance for the squeezed state noise map. By deploying a procedure to generate and tests these masks that optimizes for higher variance, we hope to find that the best mask contains areas of pixels indicating a mode.

3 Experimental Procedure

3.1 Generating Variance Maps

The data we examine here were collected by directing quantum squeezed light at a camera. The data are in the form of 128 by 128 pixel arrays; the third dimension holds about 800 frames. The camera collects data in sets of 4 frames that make up kinetic clusters. The squeezed light passes through a beam splitter before being directed at the camera. This creates two simultaneous images of the same beam (referred to in Figure 4 as Beam 1 and Beam 2). We subtract these beams to remove classical noise, which is shared between the beams. This technique differentiates classical noise from quantum noise, as the quantum noise is specific to each beam. Classical noise can come from fluctuations of the laser or the sensor; by splitting the beam and subtracting pairwise, we observe quantum noise properties separate from

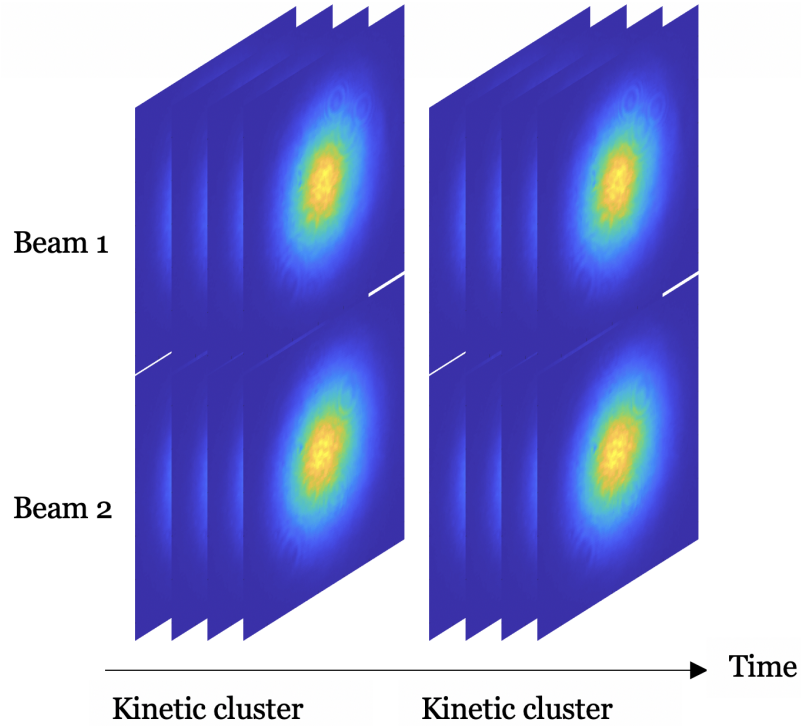


Figure 4: Sample timeline of measurement period with separate intensity maps for each beam

classical noise.

The camera detects the number of photons incident on each pixel in each kinetic cluster as shown in Figure 4. For each pixel in each frame, we take the difference between the intensity measurements from the two beams to get a noise map with reduced classical noise. We track the fluctuations between frames by calculating the variance for each pixel within each cluster and averaging the variance from all the clusters over the measurement period. Each pixel has a resulting variance that can be plotted in the variance map as seen in Figure 1. To normalize, we divide the variance by the mean of the sum of the beams over all included pixels. This process can be summarized as follows:

$$V_{norm} = \frac{\text{variance}(beam1 - beam2)}{\text{mean}(beam1 + beam2)} \quad (5)$$

We calculate variance across the subset of pixels selected by the mask by multiplying the region of pixels by the mask before calculation.

3.2 First Approach: Convolution

The standard approach to group pixels to amplify variance is called convolution. We choose a sampling region over which to group pixels and calculate the variance over just the surrounding pixels; we decided to focus our approach on groups of 20x20 pixels. The assumption here is that we expect pixels within a small region to behave in similar ways, so selecting this group should produce an amplified variance.

To visualise these sampling regions on a variance map, we assign the resulting variance value to the center pixel of the 20x20 box. To apply it to the whole image, we shift the sampling box, centering it on each pixel. This produces a convolved variance map.

The variance maps shown in Figures 5a and 5b are produced using the convolution technique with a 20x20 sampling region. A 10-pixel strip is cut off along each edge because we are assigning the average for a 20x20 pixel block to a single pixel in the center of the block. Figure 5b shows the result of convolution on a squeezed state. In this case, we achieve patterns of high variance as intended. In Figure 5a, which depicts the coherent state, we see the expected variance value of 1, a reasonable check to calibrate the approach.

While this approach does succeed in extracting quantum noise patterns from the squeezed data, it has several drawbacks. First, when shifting the sampling region over the whole image, the variance must be calculated 108x108 (11,664) times, a computationally expensive and slow task. Further, because an average value is calculated




Figure 5a: Variance map produced by convolution technique on a coherent state




Figure 5b: Variance map produced by convolution technique on a squeezed state

for every 20x20 region, any pattern smaller than the sampling region gets averaged out, so this approach does not work for capturing finer details.

3.3 Second Approach: Toggling

Our second approach is to toggle the pixels in the mask to optimize for the best shape of mask. Toggling refers to randomly selecting pixels within the mask region to be turned “on” (given a value of 1) or turned “off” (given a value of zero). When multiplied over a region in the variance map, the on pixels will be included in the pixel group over which we calculate variance, while the off pixels will be excluded. We test different configurations to determine which combination of pixels within the mask produces the highest variance. For an n by n mask, there are 2^{n^2} possible configurations, so to decrease the number of masks we must test, we develop a procedure to optimize masks for maximum average variance.

In short, this procedure consists of taking a random starting mask, toggling to generate a new mask, and comparing the variance from the two masks. The mask with the higher variance is selected and toggled for the next comparison. This method generates new test masks from previous test masks to adjust towards the mask with




Figure 6: Optimization of a 10 by 10 mask over 200 cycles

the highest possible variance. The optimization procedure is described in detail in Appendix A.

The results of applying this approach to different sized masks on the squeezed state variance map are shown in Figures 6, 7, and 8. The horizontal axis tracks the iteration, while the vertical axis is negative variance; the algorithm minimizes the vertical parameter to find the maximum variance. The red line traces the variance for every mask tried, while the yellow line indicates the variance for the mask that is accepted. By testing different sized masks on the squeezed state variance map, we determined that larger masks yield a larger average variance.

For all three mask sizes, the variance levels off by 200 cycles and any further optimization would be insignificant. The 10 by 10 mask finishes with a variance of 1.231 and the 20 by 20 mask achieves an average variance of 1.582. The 40 by 40 mask surpasses these values within 10 cycles, eventually reaching an average variance of 2.008, showing that a larger mask will yield a higher variance.




Figure 7: Optimization of a 20 by 20 mask over 200 cycles




Figure 8: Optimization of a 40 by 40 mask over 200 cycles

Figure 9 shows the corresponding masks that achieved maximum average variance after 200 cycles of annealing. Yellow pixels have a value of 1 and so are included in




Figure 9a: 10 by 10 mask resulting from the run shown in Figure 6




Figure 9b: 20 by 20 mask resulting from the run shown in Figure 7




Figure 9c: 40 by 40 mask resulting from the run shown in Figure 8

the mask, and blue pixels have a value of 0 and are discarded. Each of the optimized masks converged to more 1s than 0s, indicating that including more pixels leads to higher variance. Although we don't see a clear pattern or mode in these initial masks, this indicates that the maximum variance should come from the largest possible mask.

This approach does not reveal a mode or pattern in the optimal mask. Furthermore, the overlap between masks generated through successive iterations of the algorithm is minimal. Each generated optimal mask is highly randomized, so this approach does not succeed in targeting specific pixels that would increase variance. However, it indicates that we should test larger masks and establishes the toggling strategy that we employ in our final approach.

4 Global Toggling Optimization

We build upon our initial toggling approach to develop the global toggling optimization algorithm that we use to identify modes and image objects.

4.1 Testing Larger Masks

Upon completion of the initial run of the algorithm, we determined that a 128 by 128 mask would be the best option to maximize variance and reveal a mode over the whole variance map. The rest of our tuning and testing focuses on masks of this




Figure 10: 128 by 128 mask achieved after 750 cycles on the squeezed state

size. These masks are the same size as the map itself, eliminating the need for a sampling region over which to shift the mask. This should decrease the computation time; however, the average variance must be calculated now for the entire 128 by 128 map for every new mask tested. Because of this, the procedure takes about 10 minutes for every optimal mask generated.

Although still fairly randomized, the larger mask reaches an average variance of 4.083, confirming that the 128 by 128 mask gives the highest variance. Because of this result, we use this size of mask as a starting point for the rest of our adjustments to the optimization procedure.

4.2 Decreasing Rate of Toggle

With the 128 by 128 masks found in Section 4.1, we found that average variance converges after 1000 cycles. However, every optimal mask that was produced was different and shared no defining features. This may be from the constant randomization




Figure 11: Optimization process for 128 by 128 mask

in each toggle; every time a new mask is generated, it is changed significantly by the random toggles, so any feature may be quickly obscured. By flipping fewer pixels as the masks get closer to the optimal mask, we allow the mask to keep many of the pixels that amplifies its variance while still toggling a few pixels to keep improving the variance.

We introduce a method of toggling with a decreasing rate of number of pixels flipped to allow fine-tuning as the optimization proceeds. The new number of pixels toggled is inversely proportional to the iteration number as shown in Figure 13. The number of toggles is 200 divided by cycle number, so after 200 cycles, only one toggle takes place for each new mask tested. With this decreasing toggle rate, later masks being tested make only small changes to the current mask so the progress of the mask is preserved and adjustments are now highly focused.

Figure 12 demonstrates that with this approach, the first few masks (with a large number of pixels changed) result in a large jump in average variance. The variance

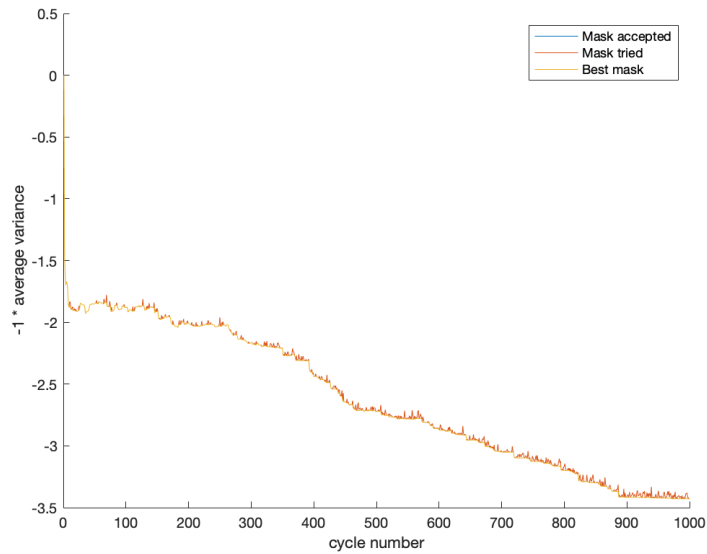


Figure 12: Optimization process for a 128 by 128 mask with toggle rate inversely proportional to cycle number




Figure 13: Number of toggles over 1000 cycles.

then stabilizes as fewer pixels are changed and the masks are adjusted in smaller steps for fine-tuned control over optimization.

4.3 Clustering

When examining the squeezed state variance map, we can reasonably assume that groups of pixels within a small radius have a similar variance. Therefore, when a pixel is selected in the optimization that increases the total average variance, it is likely that the surrounding group of pixels will further amplify the variance. Rather than toggling individual pixels, we now toggle 5 by 5 clusters of pixels. This approach amplifies the effect of a toggle on the total variance; it increases the possibility of converging to the optimal mask by generating test masks that are more likely to maximize variance. Figure 14 shows a mask with 5 by 5 clusters. While a similar mask may be found by toggling individual pixels for many more iterations, clustering greatly speeds up the process to find the optimal mask much faster. While the single-pixel toggle may take up to 30 minutes and requires more than 1000 cycles, the cluster toggle consistently converges within 1000 cycles and takes approximately 10 minutes to find the optimal mask.

4.4 Limiting Pixels in Mask

As discussing in the previous section, a higher average variance correlates with a more pixels in the mask. As such, with sufficient iterations, the optimization algorithm would likely converge to a 128 by 128 mask with every pixel included. To select only the pixels with the highest contribution to average variance, we place a cap on the number of pixels included in the mask. This increases the possibility of a mode emerging in the mask by selecting a percentage of pixels to be “on”. This percentage cap shifts the algorithm from optimizing for the mask with the most pixels to optimizing for the mask with the best pixels. Figure 14 has a 50 percent cap, so only 50 percent of the pixels in the 128 by 128 mask region are allowed to be included. The resulting mask is still spread randomly; a lower percentage cap is more likely to




Figure 14: Best mask after 1000 cycles using 5 by 5 pixel clusters and a 50 percent cap

reveal a mode. This percentage can be tuned to reveal more details in the mask or produce a higher variance.

4.5 Applying an Exclusion Mask

Because we consider the entire variance map in the calculation of average variance for the 128 by 128 masks, the corners of the map have a strong effect on average variance. As shown in Figure 1, the edges of the map have a much higher variance than the circular region of interest. This obscures the smaller fluctuations in the center. With the edges included, changes to the mask that increase the variance in the center, where we hope to see a mode, are discarded by the optimization. Therefore, we must exclude the corners from consideration when generating optimal masks to ensure that the quantum fluctuations within the map are influencing mask selection.

We apply a circular exclusion mask that forces the values outside a central circle of radius 64 pixels to be zero. This exclusion mask is applied to both the starting mask

and the imaging data being tested. This focuses the optimization on the circular region where we hope to find a mode; it cuts off the edges that we consider to be out of range. Figure 15 shows optimal masks with the exclusion mask applied.

4.6 Probabilistic Toggling

We devised a scheme to probabilistically turn pixels on and off to tune the mask towards equilibrium. As previously discussed, a percentage cap on the number of pixels turned on in the mask is needed to reveal a pattern representative of a mode. We treat this percentage as the equilibrium point and probabilistically flip pixels on and off to bring the number of pixels "on" towards that point. In this approach, we use a probability based on the current number of pixels on and the desired percentage to determine the number of flips on and off. The algorithm is as follows:

Probabilistic Toggling

Set the number of flips using the approach from Section 5.2

Set a maximum ratio of "on" pixels to total pixels in mask region (128 by 128)

If more pixels are on than ratio requires, decrease the number of flips by a factor of 10

For each flip:

Toggle off with probability $\frac{Num_On}{Num_Total} - Ratio$

Toggle on with probability $\frac{Num_Off}{Num_Total} - (1 - Ratio)$

This algorithm provides some control over the ratio of pixels being turned on or off. The previous approach selected random pixels to flip, whether they were on or off. This algorithm specifically selects pixels to guide the mask towards equilibrium. It turns pixels on and off based on how far the mask is from equilibrium, but it allows flips in both directions to give the optimization some flexibility. Figure 8 shows the results of the probabilistic toggling method with all of the prior techniques. We test the method with a 10 percent and 20 percent cap and determine that 10 percent is

the ideal percentage to allow details to emerge.




Figure 15a: Mask after 1000 cycles with 10 percent cap and exclusion mask

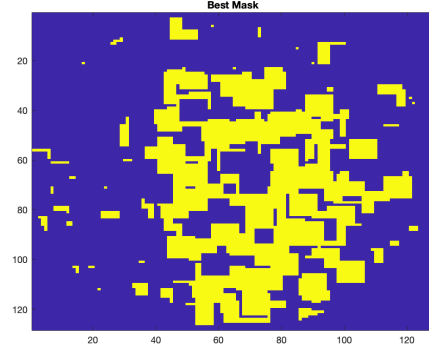


Figure 15b: Mask after 1000 cycles with 20 percent cap and exclusion mask

4.7 Additional Approaches

Sections 4.1 through 4.6 discuss the techniques implemented in the final iteration of the optimization algorithm. We attempted several other approaches that did not lead to modes under our current tests, but which may yield useful results upon further exploration. One of these approaches was changing the optimization from maximizing to minimizing variance. In this case, the optimal case would be a completely “off” mask, so we needed the starting mask to be completely “on” so test masks would be accepted in optimization. The results were not a direct opposite of the maximization as would be expected, but rather yielded no discernible pattern.

Another approach that we tried was including a factor for the number of “on” pixels in the merit function to penalize the mask’s merit if the number of pixels was far off from the target number. Although we kept this contribution small in the merit function by scaling it by a decimal value, it clouded the variance calculation and tended to overtake small fluctuations in the variance map in determining the best mask.

Other techniques we explored were manually assigning a starting mask to highlight the corners (which we eventually cut off with the exclusion mask). Because the corners are known to have relatively high variance, we hoped that groupings of pixels would gather to select the corners. However, the optimization did not appear to accept corner pixels any more than the rest of the variance map.

These techniques, while they did not yield any significant results at this time, may serve as a starting point for further investigation.

5 Results

We tested the optimization procedure on three types of data. The main focus of our investigation was using an incoming squeezed state, which we would expect to have patterns of variance and modes. We also used a coherent state to verify the algorithm. Finally, we ran the procedure on a set of data imaging objects with the squeezed state; these consisted of a rectangular knife edge blocking portions the bottom right quadrant of the camera and a cross-shape in the center of the beam.

We ran the procedure many times on each data set and combined all the masks to depict the overlap in the generated optimal masks. Because all of the generated masks have an element of randomization and probability, the summed masks are a good indication of which groups of pixels, if any, consistently produce a high variance and similar masks.

5.1 Squeezed State

Figures 16 and 17 (previously shown) and 18 show three methods of imaging the same data set. Figure 16 depicts the classical image corresponding to the variance maps obtained in Figures 17 and 18. The incoming light is a round squeezed state beam, so any other patterns are solely quantum results. Figure 17 is the result




Figure 16: Classical intensity image for the squeezed state

of the convolution approach, while Figure 18 arises from the new global toggling optimization method. The color bar in Figure 18 indicates the number of masks that share a pixel.

The global toggling approach on the squeezed state brings out a clear mode with all of the masks sharing pixels in a ring shape around the center of the variance map. This is both a surprising and encouraging result. We saw in the convolution approach that calculating variance over groups of pixels has the ability to amplify the variance. In the quantum toggling approach, we are able to take this a step further and use this amplification to clearly depict previously unknown quantum modes. The optimal mask approach has generated a ring mode given a round squeezed state beam. This demonstrates that global toggling optimization has the ability to discover previously obscured quantum properties of the data and, with more tuning and testing, may reveal even more higher-order modes.




Figure 17: Variance map obtained from the convolution approach for the squeezed state

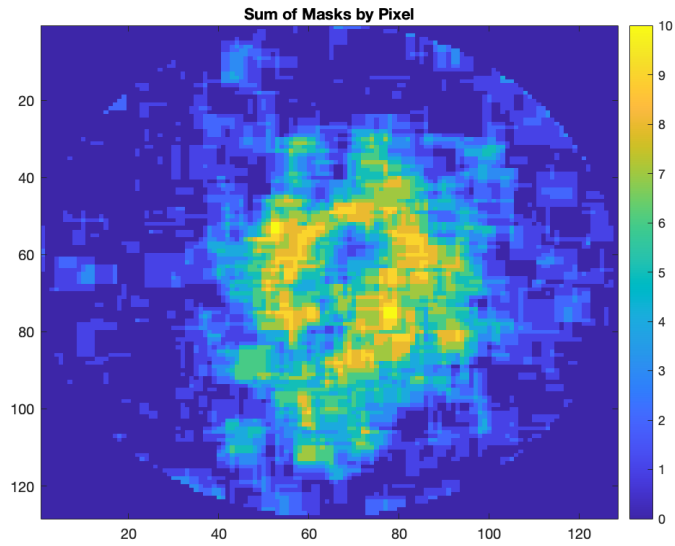


Figure 18: Sum of 10 masks generated by the optimization on the squeezed state data




Figure 19: Classical intensity image for the coherent state

5.2 Coherent State

Figures 19, 20, and 21 refer to the coherent state data that we use to verify our approach. Figure 20 correctly depicts a variance of 1 across the beam shown in Figure 19. Because the coherent state has a constant variance of 1, so we expect the 10 percent of included pixels to be randomly distributed across the circular mask region. After 10 iterations of the global toggling optimization on the coherent state data, the summed mask in Figure 21 shows no significant overlap between individual masks, so no pixels produce a higher variance than others. Therefore, we can assert with confidence that this algorithm works as intended.

5.3 Imaging Objects

To see how the global toggling optimization approach works in practice, we began testing the algorithm on data imaging an obstructing object. All of the following blocked images use an incoming squeezed state. Our hope was that we could recon-




Figure 20: Variance map obtained from the convolution approach for the coherent state

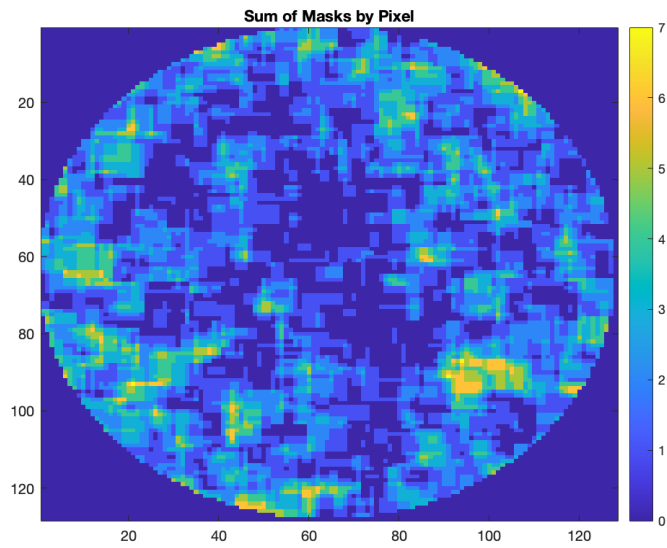


Figure 21: Sum of 10 masks generated by the optimization on the coherent state data. The color bar indicates the number of masks that share a pixel.

struct the block seen in the classical image and observe any changes in modes. In a low-light scenario, the classical image would not be available and quantum imaging would take the place of classical imaging to depict an object. If the global toggling optimization approach can recreate an object as seen in a classical image, this indicates that it has potential to be used to carry out low-light imaging.

Figures 22a, 22b, and 22c show the results of using classical and quantum imaging on a small rectangular knife-edge block on the bottom right of the image. In Figure 22a, the knife-edge is in a poorly-lit area of the image and does not block much of the incoming beam; it is present but barely visible.

We get a large improvement on the visibility of the image using the convolution approach. Figure 22b shows the same image using convolution on equivalent quantum noise data to Figure 22a. The figure shows a clearly-defined shadow in the corner with variance 1 where the incoming squeezed state is blocked. This matches with our expectations of a blocking object with an incoming squeezed state creating a coherent vacuum in the shape of the object. This proves to be a very effective of both recreating and enhancing the classical image.

Figure 22c shows the same image using the global toggling optimization approach. The corner block is less visible, but the masks do avoid grouping in the corner holding the rectangular block. This corner is particularly difficult to see in the classical and global toggling optimization approaches because the block just encroaches the poorly lit outer region of the map. This approach also shows a mode in the center of the image where there is a high level of overlap between masks. While it is not as clear of an image as Figure 22b, it brings out quantum properties that are not visible using convolution.

Similarly, Figures 23a, 23b, and 23c show the same rectangular knife-edge, now blocking about a quarter of the image. Figure 23a shows the classical image that




Figure 22a: Classical intensity image of partial rectangular knife-edge




Figure 22b: Variance map obtained from the convolution approach for the partial knife-edge image




Figure 22c: Sum of 20 masks generated by the optimization on the partial knife-edge image

we try to recreate in Figures 23b and 23c. The variance map from convolution correctly has a shadow of variance 1 in the shape of the rectangular block. For the toggling approach, we expect to see some overlapping outside the obstructed region and missing or scattered pixels inside the obstructed region. We do in fact see some grouping outside of the obstruction, while the obstructed area itself is generally avoided. Notably, the ring-shaped mode from the unblocked squeezed state image has either mutated due to the presence of the blocking object, or a different mode dominates the image due to the block. The toggling approach contains a shadow in the expected spot and begins to show pixel groupings outside the shadow that may be another mode.

Finally, we tested the imaging techniques on a somewhat more complicated image with finer details. This consists of a small cross centered in the beam, as seen in Figure 24a. Small portions of other crosses are present on the top and right edges of the image because the object being imaged was actually a grid of crosses. While we don't expect to see any features of these partial crosses, they show up as shadows on the edges of the variance maps. This particular image has revealed clear differences in the convolution and global toggling optimization approaches.

Shown in Figure 24b, the convolution approach displays a shadow in the correct




Figure 23a: Classical intensity image of rectangular knife-edge blocking bottom right corner of the image




Figure 23b: Variance map obtained from the convolution approach for the knife-edge image




Figure 23c: Sum of 80 masks generated by the optimization on the knife-edge image

central region, as well as an upper shadow from the neighboring cross. However, without prior knowledge of the image itself, it is difficult to discern the object being imaged. The sharp edges and finer details have been blurred out by the averaging of the convolution approach.

The global toggling optimization approach, on the other hand, samples pixels over the entire image and so can maintain a level of detail that is not possible in convolution. Figure 24c clearly achieves a cross-shape in the center of the masks. Surrounding the cross is a grouping similar to the mode from the squeezed state optimization. This image successfully recreates the classical intensity image, which uses hundreds of photons, using only quantum noise measurements, which uses a single photon. The apparent success of imaging the cross using only a few photons with the global toggling optimization methods indicates that this may be a viable approach moving forward to perform quantum imaging.

6 Conclusion

We have presented two approaches, convolution and global toggling optimization, that apply masks to variance maps for quantum imaging in search of high-




Figure 24a: Classical intensity image of a cross in the center of the image




Figure 24b: Variance map obtained from the convolution approach for the cross image




Figure 24c: Sum of 80 masks generated by the optimization on the cross image

variance regions indicative of modes. Both approaches have succeeded in recreating classical images using only quantum noise data gathered in low-light settings.

While the convolution approach produces a clear corner shadow when a rectangular corner block is introduced, there are several drawbacks that cannot be overlooked. First, due to the fact that variance is calculated by average over a block of pixels, this approach can't resolve details smaller than the size of that block. This produces a high level of blur and introduces a factor that must be set by the user: to know what the maximum size block is that can be used to reproduce the image, one must have an idea of the size and level of detail what is being imaged. While this is not a challenge in a testing setting, where we are imaging a known object and have a classical counterpart, in a real scenario where low-light imaging must be used, this prior knowledge is not available. It is also based only on localized correlation, so we can't see modes spanning the whole image from this approach. Finally, because variance must be calculated by centering the block over every pixel, creating a single variance map takes a long time.

The global toggling optimization approach also produces a shadow in blocked regions to recreate the shadow in the classical image. In addition, we are able to observe modes across the span of the image, bringing out a ring-shaped mode that

has not been previously obtained from these data. This approach can also produce sharper images of finer details as each pixel is individually assigned to the mask. The toggling method also does not have the size resolution limitation of the convolution approach, so prior knowledge of the image is not required. This makes it a feasible strategy for low-light quantum imaging where obtaining a classical equivalent is not possible.

One area of further investigation in the global toggling optimization method is the ratio of “on” to “off” pixels. By introducing this limit, we have created another tunable parameter that has the ability to bring out details in the image. Deciding this ratio is a focus for further study; while with our data sets we determined that 10% “on” pixels is appropriate, this number may change based on how much or little of the image is blocked. It may be possible to decide this ratio by looking at the overall variance of the image before masks are applied. Another possibility for this approach is weighting the pixels based on their noise to determine the most “important” pixels that contribute to high variance. These approaches may lead to the discovery of more modes and finer detail in the resulting images.

We developed the global toggling optimization procedure that generates optimal masks with high variance and, when these masks are layered, reveals new modes and displays shadows in blocked regions. We have shown that it is possible to observe modes and recreate classical images using processing techniques on quantum noise data. This technique is a promising method of approaching quantum imaging and may, upon further study and tuning, become a reliable way to produce low-light images.

References

- [1] T. Gregory et al. “Imaging through noise with quantum illumination”. In: *Science Advances* 6.6 (2020), eaay2652. DOI: [10.1126/sciadv.aay2652](https://doi.org/10.1126/sciadv.aay2652).
- [2] Marta Gilaberte Basset et al. “Perspectives for Applications of Quantum Imaging”. In: *Laser & Photonics Reviews* 13.10 (2019), p. 1900097. DOI: <https://doi.org/10.1002/lpor.201900097>.
- [3] Miles J. Padgett and Robert W. Boyd. “An introduction to ghost imaging: quantum and classical”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2099 (2017), p. 20160233. DOI: [10.1098/rsta.2016.0233](https://doi.org/10.1098/rsta.2016.0233).
- [4] Savannah L. Cuozzo et al. “Low-Light Shadow Imaging using Quantum-Noise Detection with a Camera”. 2021. arXiv: [2106.00785](https://arxiv.org/abs/2106.00785) [quant-ph].
- [5] Jeremy B. Clark et al. “Imaging using quantum noise properties of light”. In: *Opt. Express* 20.15 (July 2012), pp. 17050–17058. DOI: [10.1364/OE.20.017050](https://doi.org/10.1364/OE.20.017050).
- [6] Carlos A. Pérez-Delgado, Mark E. Pearce, and Pieter Kok. “Fundamental Limits of Classical and Quantum Imaging”. In: *Phys. Rev. Lett.* 109 (12 Sept. 2012), p. 123601. DOI: [10.1103/PhysRevLett.109.123601](https://doi.org/10.1103/PhysRevLett.109.123601).
- [7] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. “Quantum-Enhanced Measurements: Beating the Standard Quantum Limit”. In: *Science* 306.5700 (2004), pp. 1330–1336. DOI: [10.1126/science.1104149](https://doi.org/10.1126/science.1104149).
- [8] Melissa A. Guidry. “Exploring the Multi-Mode Structure of Atom-Generated Squeezed Light”. William & Mary, 2017. URL: <https://scholarworks.wm.edu/honorstheses/1029>.
- [9] Savannah L. Cuozzo. “Quantum Sensing for Low-Light Imaging”. William & Mary, 2022.
- [10] Wikipedia contributors. *Transverse mode* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 9-May-2022]. 2022. URL: https://en.wikipedia.org/w/index.php?title=Transverse_mode&oldid=1082244321.
- [11] J. E. Gubernatis. “Marshall Rosenbluth and the Metropolis algorithm”. In: *Physics of Plasmas* 12.5 (2005), p. 057303. DOI: [10.1063/1.1887186](https://doi.org/10.1063/1.1887186).
- [12] Eugeny E. Mikhailov. *Programming with MATLAB for Scientists: A Beginner’s Introduction*. CRC Press, 2018.

A Appendix: Simulated Annealing

To search for an optimal mask with maximum average variance, we employ a simulated annealing optimization algorithm, also known as the Metropolis algorithm [11]. The algorithm uses the Boltzmann probability distribution to determine whether to select a mask:

$$p(E) = e^{-(E_{new} - E_{curr})/kT} \quad . \quad (6)$$

The Boltzmann distribution gives the probability of having energy E ; annealing uses this probability to minimize energy $E(X)$. In our case where we optimize for maximum variance, energy corresponds to negative variance and $E(X) = -V(X)$. The algorithm works as follows [12]:

Simulated Annealing

Choose a starting value X_{curr} and calculate $E(X_{curr})$

Set a transition function to generate X_{new} from X_{curr}

For each iteration:

 Generate a candidate X_{new} and calculate $E(X_{new})$

 If $E(X_{new}) < E(X_{curr})$, set $X_{curr} = X_{new}$

 Otherwise, set $X_{curr} = X_{new}$ based on probability $p(E)$

Our starting value X_{curr} is an n by n mask, initialized to zeros and randomly toggled once. The transition function toggles X_{curr} by randomly flipping pixels to generate a candidate X_{new} ; the algorithm determines whether to accept or reject this candidate mask.

In Equation 6 for the Boltzmann distribution, T is the “temperature,” a tunable parameter that here has no physical significance but is manually set by the user. The temperature sets the allowance for how much the variance is allowed to decrease from X_{curr} to X_{new} . If the new mask *increases* the variance, the algorithm immediately accepts the new mask. If the transition *decreases* the variance, the new mask is

accepted or rejected based on the probability [11]. This algorithm optimizes for the maximum average variance. With correctly tuned parameters and sufficient iterations, the mask should converge to the best mask which gives the highest average variance.