Data Reduction and Image Reconstruction Techniques for Non-redundant Masking

S. Sallum and J. Eisner

Astronomy Department, University of Arizona, 933 N. Cherry Ave., Tucson, AZ 85721, USA; ssallum@email.arizona.edu

Received 2017 July 10; revised 2017 September 13; accepted 2017 September 28; published 2017 November 16

Abstract

The technique of non-redundant masking (NRM) transforms a conventional telescope into an interferometric array. In practice, this provides a much better constrained point-spread function than a filled aperture and thus higher resolution than traditional imaging methods. Here, we describe an NRM data reduction pipeline. We discuss strategies for NRM observations regarding dithering patterns and calibrator selection. We describe relevant image calibrations and use example Large Binocular Telescope data sets to show their effects on the scatter in the Fourier measurements. We also describe the various ways to calculate Fourier quantities, and discuss different calibration strategies. We present the results of image reconstructions from simulated observations where we adjust prior images, weighting schemes, and error bar estimation. We compare two imaging algorithms and discuss implications for reconstructing images from real observations. Finally, we explore how the current state of the art compares to next-generation Extremely Large Telescopes.

Key words: methods: data analysis – techniques: high angular resolution – techniques: image processing – techniques: interferometric

1. Introduction

Direct exoplanet studies rely on high-contrast imaging methods used with adaptive optics systems. Techniques such as coronagraphy (e.g., Guyon et al. 2014), coupled with post-processing algorithms such as angular differential imaging (ADI; e.g., Marois et al. 2006), can detect planets around nearby stars (e.g., Macintosh et al. 2015). The theoretical inner working angles of state-of-the-art coronagraphs are $\sim 0.7\text{--}3\lambda/D$ (e.g., Guyon et al. 2006; Mawet et al. 2012). In practice, the achievable inner working angle is a function not only of coronagraph design but also of wavefront control. Residual low-order aberrations such as tip-tilt, which can be caused by seeing variations and vibrations (e.g., Meimon et al. 2010), can leak into the off-axis light (e.g., Mawet et al. 2012). Furthermore, data reduction algorithms that build reference point-spread functions (PSFs) in the image plane (e.g., Lafrenière et al. 2007; Soummer et al. 2012) perform poorly at small separations. Close to the star, the small number of image elements causes small number statistics that degrade the achievable contrast (e.g., Mawet et al. 2014). These factors combine to limit achievable coronagraph performance to $\gtrsim \lambda/D$ (e.g., Mawet et al. 2012; Ruane et al. 2017).

Non-redundant masking (NRM; e.g., Tuthill et al. 2000) is a way to probe smaller separations than more traditional imaging techniques such as coronagraphy. NRM turns a conventional telescope into an interferometer through the use of a pupil-plane mask. No two baselines have the same position angle or separation, meaning that residual wavefront errors do not add incoherently and can be characterized. Thus, despite blocking the majority of incident light, in the presence of noise NRM provides much better PSF characterization than filled-aperture observations. It can detect sources with moderate contrast ($\sim 1\text{:}100$) at separations as small as $0.5\lambda/D$, expanding the companion discovery phase space. NRM has led to the detection of stellar (e.g., Ireland & Kraus 2008; Biller et al. 2012) companions, substellar companions (e.g., Kraus & Ireland 2012; Sallum et al. 2015b), and circumstellar disk features (e.g., Cheetham et al. 2015; Sallum et al. 2015a) at or

even within the diffraction limit. This makes it particularly useful for direct planet formation studies, where most targets are at distances of >100 pc.

Here we describe observational, data reduction, and image reconstruction techniques for NRM observations. We discuss a `python` NRM data reduction pipeline, used primarily on observations from the Large Binocular Telescope Interferometer (LBTI; e.g., Hinz et al. 2008) detector, LMIRCam (Leisenring et al. 2012). However, our pipeline is general and has also been applied to data from Magellan, Keck, and the Very Large Telescope (VLT).

We show the effects of various calibration steps on the phases and squared visibilities. We present reconstructed images from simulated observations of sources with different morphologies. We discuss the effects of different initial images, error bar scalings, and baseline weighting schemes on these simulated reconstructed images. We also compare simulated Giant Magellan Telescope (GMT) reconstructed images to those currently achievable with the 23 m LBTI.

2. Experimental Setup

In NRM observations, the detector records the interference fringes formed by the mask, called “interferograms.” Fourier-transforming the interferograms yields complex visibilities, which have the form $A \exp i\phi$. Since the mask is non-redundant—no two hole pairs have the same separation and orientation—information from each baseline is located at a unique location in Fourier space. Sampling the Fourier transform, we calculate squared visibilities (the total power on each baseline) and closure phases (sums of phases around baselines forming a triangle) (e.g., Jennison 1958; Baldwin et al. 1986). Closure phases are particularly powerful for companion detections since they are sensitive to asymmetries and are unaffected by atmospheric phase offsets. Because closure phases are correlated, we then project them into linearly independent quantities called kernel phases (e.g., Martinache 2010; Ireland 2013; Sallum et al. 2015a). We apply both model fitting and image reconstruction to understand the source morphology.

Figure 1. Example raw science image (left), the same image after dark and background subtraction (middle), and then after flat-fielding (right). The left and middle panels are shown with different color scales since the median value is significantly higher before background subtraction. The right panel is shown on the same color scale as the middle panel. These images are from the 2014 December data set published in Sallum et al. (2015b). In the two rightmost panels, the snowflake patterns show the interferograms formed by the mask. The two sides show the interferograms formed by the left and right LBT primary mirrors. The negative images on top show the results of using the median of one dither to subtract the sky background from the bottom dither.

Non-Photometric Conditions Reduction Steps

Figure 2. From left to right: an example raw science image, the same image after dark subtraction, then after flattening, and lastly after background subtraction. Vignetting can be seen in the rightmost columns, bottom rows, and corners of the CCD. The middle two panels are shown on the same color scales. The outermost two panels are on different color scales since the average pixel value changes significantly after dark and background subtraction. These images were taken after LMIRCam was upgraded to a 2048×2048 HAWAII-2RG detector.

3. Observational Strategy

NRM observations at LBT, Magellan, Keck, and the VLT are carried out in the near-infrared ($\sim 1\text{--}4\ \mu\text{m}$). Since the sky background is high at these wavelengths (especially at wavelengths longer than $2\ \mu\text{m}$), we observe each target with two dither pairs per pointing. We dither the images between the top and bottom halves of the detector. To limit the effects of flat-field errors, we keep the interferograms centered as close to the same pixels as possible for each dither position. In clear conditions where the sky brightness is not variable on short timescales, one dither can be used to form the sky background for the other.

We alternate the science target observations with observations of unresolved calibrator stars. This allows us to account for instrumental contributions to the closure/kernel phases and squared visibilities. We match the target and calibrators' visible fluxes so that they have similar AO correction quality, and their infrared fluxes so that they have similar noise levels on the science detector. Calibrators are also vetted to ensure they do not have significant near-infrared excesses (indicating circumstellar dust) in their spectral energy distributions. We choose multiple calibrators for each science target to minimize the possibility of contamination by binary calibrators.

4. Data Reduction

4.1. Dark Subtraction, Flat-fielding, and Background Subtraction

We create a master dark frame by taking the median of many dedicated dark images taken during the observing night. We then dark-subtract a set of sky flats taken during the night. We median-combine them and divide by the mode to create the master flat.

We apply dark, flat, and background calibrations differently depending on the observing conditions. In photometric conditions, for each dither pair we use one dither to perform dark and background subtraction on the other. We median all the frames in the first dither and subtract this median from each frame in the second dither. We then divide by a flat. Figure 1 shows example images as these steps are applied.

For data sets taken in variable conditions such as intermittent cirrus, one dither may have a very different background level than another. In this case, rather than use one dither to perform the dark and background subtraction for another, we first subtract a median dark from every image in each dither. We then divide by a flat. After flattening, we take the median of all pixels in non-vignetted portions of the CCD to be the median background signal. We subtract this median background value

Channel Bias Subtraction

Figure 3. Top: images before and after channel bias correction. Bottom: complex visibilities (Fourier-transformed images) before and after channel bias correction. The vertical striping present in the uncorrected images creates low spatial frequency noise in the complex visibilities.

from the entire frame. Figure 2 shows example images for each of these steps.

4.2. Channel Bias Correction

LMIRCam is a 2048×2048 HAWAII-2RG (H2RG; Beletic et al. 2008) detector, recently upgraded from a 1024×1024 HgCdTe detector (e.g., Leisenring et al. 2012). Both of these configurations are read out in 64-pixel channels, each of which has its own analog-to-digital converter with a unique bias level. We correct for these different bias levels and any nonlinear bias changes during each exposure by taking the median of each 64-column channel for each image. We subtract the bias from each readout channel, as shown in Figure 3. If uncorrected, these channel biases lead to low spatial frequency noise in the complex visibilities (Figure 3). While this 64-pixel scenario is specific to LMIRCam, H2RGs are used in other instruments with NRM capabilities such as the Gemini Planet Imager (e.g., Ingraham et al. 2014). Regardless of the readout configuration, variable bias levels across the subframe may add systematic errors to the Fourier transform and can be corrected in a similar way.

4.3. Bad Pixel Correction

We create a bad pixel map for each data set using dark frames. We calculate the standard deviation of each pixel across the cube of images. We then label some fraction of pixels with the highest and lowest standard deviations as bad. We allow for asymmetric cuts since the tails on the distribution of standard deviations may be asymmetric (they are for LMIRCam; see Figure 4). We use a variety of upper and lower cuts and choose the bad pixel map that minimizes the scatter in the calibrator observations for the night. We make corrections in the image plane, replacing each bad pixel with the average of the surrounding ones that have not been flagged. For an

Bad Pixel Map

Figure 4. Top: bad pixel map created from a set of 20 dark frames. Bottom: histogram of pixel standard deviations over 20 dark frames. The vertical lines indicate the cuts used to create the bad pixel map shown in the top panel.

Bad Pixel Correction

Figure 5. Example interferogram subframe before (left) and after (right) bad pixel correction.

example data set taken in 2014 December (published in Sallum et al. 2015b; see Figure 5), dropping the top 2% and bottom 1% of pixels resulted in the best bad pixel correction.

4.4. Noise versus Reduction Steps

In Figure 6 we compare the scatter in the calibrated and uncalibrated data at each reduction step for the 2014 December LBT observations published in Sallum et al. (2015b). We use two unresolved stars to calibrate each other, and then we examine the scatter after flattening, channel bias correction, and bad pixel correction. For the closure phases, the channel bias and bad pixel corrections decrease the closure phase scatter by

Figure 6. Comparison of closure phase (left) and squared visibility (right) scatters for observations of unresolved stars from 2014 December. Dashed lines show reductions that include flattening, while solid lines show reductions without flattening. Calibrated data are shown in blue, while uncalibrated data are shown in black.

nearly equal amounts ($\sim 0^\circ.11$ – $0^\circ.14$). For the squared visibilities, reductions with a channel bias correction have orders of magnitude more scatter than those without a channel bias correction. This is because only certain baselines sample the spatial frequencies of the channel bias noise; these will have much more power in them than the baselines that do not.

For the 2014 December L' data, both the uncalibrated and calibrated data have lower scatter when a flat is not applied. We suspect that this is because of the small number of flats used to create the master sky flat, which can be thought of as an image of ones with Gaussian noise added. Flattening will thus add noise to the uncalibrated complex visibilities. If the image on the detector always falls on the exact same pixels, then the flat-field applied to the subframed interferogram is always the same. With pointing inconsistencies, the subframed flat-field is also inconsistent; any noise added by the master flat will then change slightly between pointings and will be more difficult to calibrate out. This increases the scatter in the calibrated data as well. A large number of flats would decrease the noise in the master flat and thus in both the uncalibrated and calibrated observables. Previous simulations of NRM observations suggest that $\sim 10^6$ photoelectrons per pixel are required for flattening to add less than $\sim 0^\circ.06$ (Ireland 2013). While flattening does not significantly change or improve the LMIRCam closure phases and squared visibilities, it may be more useful on other detectors with greater flat-field variations across the subframed interferograms.

5. Closure Phase Calculation

The (u, v) coordinates for a closing triangle of baselines satisfy the following:

$$(u_1, v_1) + (u_2, v_2) + (u_3, v_3) = 0. \quad (1)$$

To calculate a closure phase for a single triangle, we sample the complex visibilities for the three baselines and multiply them to

Figure 7. Example window functions.

form the bispectrum. For each triangle, we average the bispectra measured from all images in a dither and then take the average bispectrum phase as the closure phase. Since the bispectrum has both amplitude and phase, averaging bispectra upweights higher signal-to-noise measurements, which have higher complex visibility amplitudes.

The simplest way of calculating a closure phase in practice is to use only a single pixel for each baseline. However, the finite mask hole size causes information from each baseline to be spread over multiple pixels in the Fourier transform. To use information from more than just the central pixels, we can multiply the subframed interferogram by a window function. Taking the Fourier transform of a windowed image convolves the Fourier transforms of the interferogram and window function, creating interpixel correlations. We can then still sample single pixels but incorporate more information than the unwindowed single-pixel method. One caveat with this method is that narrow enough window functions can cause signals from adjacent baselines to bleed into one another.

An alternative to windowing is to average many bispectra in each individual image for each triangle of baselines (the “Monnier” method; Monnier 1999), as diagrammed in Figure 8. We first take the Fourier transform of the unwindowed interferogram. For each mask hole triplet we find all closing triangles that connect the extended signals from the three baselines. We calculate their bispectra and average them to create a single bispectrum for each triangle of baselines in an image. As in the single-pixel method, we average the bispectra for all images in a dither and then take the bispectrum argument to be the average closure phase.

Figure 7 shows example window functions (Hanning and super-Gaussian), and Figure 8 illustrates the different closure phase calculation methods. A one-dimensional Hanning window of size M has the form

$$0.5 - 0.5 \cos \frac{2\pi n}{M-1}, \quad (2)$$

where n goes from 0 to $M-1$. To create a two-dimensional window, we make two one-dimensional Hanning windows and then take their outer product. We generate super-Gaussian

Figure 8. Closure phase generation methods. The left panel shows the Fourier transform of an interferogram multiplied by a Hanning window. Here we would use a single pixel at the center of each splodge to calculate the bispectrum for a mask hole triplet (white lines). The window function shown here caused correlations between adjacent splodges. The right panel shows the Fourier transform of an unwindowed interferogram. The white lines show example closing triangles for individual pixels within each splodge. Here we average the bispectra of these triangles to calculate the bispectrum for each mask hole triplet.

windows according to the following:

$$\exp\left[-\ln(0.5)\left(\frac{r}{\text{HWHM}}\right)^m\right], \quad (3)$$

where r is the distance from the center of the subframe and HWHM is the window half-width at half-maximum.

We use the scatter in the uncalibrated and calibrated data to compare the various closure phase calculation methods (see Figure 9). For both the calibrated and uncalibrated data, the single-pixel method without windowing results in the highest scatter ($\sigma_{\text{uncal}} = 3^\circ.8$, $\sigma_{\text{cal}} = 3^\circ.4$) and the “Monnier” method without windowing results in the lowest scatter ($\sigma_{\text{uncal}} = 2^\circ.3$, $\sigma_{\text{cal}} = 2^\circ.1$). The single-pixel + window methods have a range of intermediate scatters ($\sigma_{\text{uncal}} = 2^\circ.4$ – $3^\circ.0$, $\sigma_{\text{cal}} = 2^\circ.2$ – $2^\circ.6$). Here, the narrower windows, which have wider Fourier transforms and thus farther-reaching interpixel correlations, lead to lower scatter in both the calibrated and uncalibrated data.

The “Monnier” method is computationally more expensive than windowing, since hundreds or thousands of pixel triangles could exist for each closure phase. However, since windowing blurs the Fourier transform, it could contaminate closure phases by incorporating phase information from spatial frequencies that do not form closing triangles. For most data sets the “Monnier” method also results in the lowest scatter. For these reasons we use it over windowing for closure phase calculations. The single-pixel + window method is only $\sim 0^\circ.1$ – $0^\circ.2$ higher scatter and would be useful with limited computational resources or large data sets.

6. Kernel Phase Projection

We project closure phases into linearly independent quantities called kernel phases (Martinache 2010) so that we can fit uncorrelated observables. To calculate kernel phases, we assume that our observed phase vector, Φ , can be written as a linear combination of instrumental phases, ϕ , plus any intrinsic

Figure 9. Comparison of uncalibrated (left) and calibrated (right) closure phase scatters for different window functions and calculation methods.

source phase Φ_0 (following the notation in Sallum et al. 2015a):

$$\Phi = A \cdot \phi + \Phi_0. \quad (4)$$

For a non-redundant mask, ϕ contains the instrumental phase measured on each of N subapertures. A is an M by N matrix that describes how those are combined to form the phase measured for each of $M = \binom{N}{2}$ hole pairs. We search for the kernel, or null space, of A so that the instrumental phase signal is eliminated, or

$$K \cdot A = 0. \quad (5)$$

We find K using singular value decomposition of A^T :

$$A^T = U \cdot W \cdot V^T. \quad (6)$$

Here U is an $N \times N$ unitary matrix and V is an $M \times M$ unitary matrix. W is an $N \times M$ diagonal matrix with zeros and positive values only. The columns in V corresponding to the zero values in W form the null space of A . These make up the rows of K .

We already have a matrix T that is not linearly independent but that does eliminate instrumental phase signals. T projects the M phases into $\binom{N}{3}$ closure phases:

$$\Phi_{\text{CP}} = T \cdot A \cdot \phi + T \cdot \Phi_0. \quad (7)$$

We can thus project the closure phases into kernel phases using the matrix B such that

$$B \cdot T = K. \quad (8)$$

Since K has full row rank, but not full column rank, it has a right inverse (K_R^{-1}) only:

$$K \cdot K_R^{-1} = I = B \cdot T \cdot K_R^{-1}. \quad (9)$$

B also only has a right inverse:

$$B_R^{-1} = T \cdot K_R^{-1}. \quad (10)$$

We then take the left inverse of Equation (10) to calculate B , and we use it to project the closure phases into kernel phases.

Figure 10. Histograms of uncalibrated (left) and calibrated (right) kernel phases (top) and their fractional errors (bottom) for the “Martinache” (without C_{CP} diagonalization; solid black lines) and “Ireland” (with C_{CP} diagonalization; dashed blue lines) projection methods.

We form the kernel phase variances by taking the diagonal values of the projected closure phase covariance matrix:

$$C_k = B \cdot C_{CP} \cdot B^T. \quad (11)$$

We can form statistically independent kernel phases (Ireland 2013) by applying another projection to diagonalize the kernel phase covariance matrix, C_k . We diagonalize C_k by the spectral theorem:

$$C_k = U \cdot W \cdot U^*, \quad (12)$$

where U is a unitary matrix whose columns contain the eigenvectors of C_k . K_s , a statistically independent kernel phase projection, is then calculated according to

$$K_s = U \cdot K. \quad (13)$$

Figure 10 compares kernel phases calculated using K (the “Martinache” projection) and K_s (the “Ireland” projection). The top two panels show uncalibrated and calibrated observations of two unresolved stars from 2014 December. The Martinache kernel phases have lower calibrated and uncalibrated scatters

Figure 11. Histograms of uncalibrated (left) and calibrated (right) squared visibilities for different window functions.

than the Ireland kernel phases. The two projections may have relative scalings, which would result in kernel phase signals of different magnitudes for a given source morphology. To ensure a fair comparison, we also show their fractional errors (see Figure 10, bottom panels). While the Ireland projection has only slightly higher fractional error for the majority of the kernel phases, it has many more outliers than the Martinache projection. We thus use the Martinache projection, which incorporates information only about the mask and not about the observations, for model fits to kernel phases.

7. Squared Visibility Generation

We calculate squared visibilities by summing the power ($|Ae^{i\phi}|^2$) in all pixels corresponding to each baseline in the complex visibilities. We measure any bias by taking the average power for regions in Fourier space without signal. We subtract this bias and then normalize by the power at zero baseline (equivalent to normalizing by total power in the interferogram).

Figure 11 compares squared visibilities generated after applying the different window functions. The choice of window function has a more dramatic effect on the uncalibrated visibilities than on the calibrated visibilities. Windowed squared visibilities have much higher values before calibration. This is because convolving in Fourier space decreases the power at a spatial frequency of zero relative to the mask spatial frequencies. However, this effect is uniform between the target and calibrator observations and thus calibrates out.

We note that this squared visibility comparison is for a data set with very good sky subtraction; the edges of the subframed images have an average pixel value very close to zero. However, in conditions where sky subtraction is difficult, average pixel values may deviate from zero at the edges of the subframed images. This can cause ringing in the Fourier plane that would contaminate the squared visibilities. Here, windowing helps to remove this noise, which may not be

Figure 12. Polycal example for a single kernel phase. The top panel shows the calibrator (black points) and target (blue points) observations for a single kernel phase as a function of time. The plotted lines show polynomial fits as a function of time with different maximum orders. The bottom panel shows the calibrated kernel phases over time, with different colors indicating fits up to different orders in time. Since the best-fit third-order polynomial was nearly identical to the best-fit second-order polynomial, the purple triangles and purple dotted line are plotted over the green diamonds and green solid line. Here a second-order polynomial was eventually chosen, since it led to the lowest scatter in the calibrated observables.

consistent between target and calibrator observations in variable conditions.

8. Calibration

8.1. Instrumental Signal Fitting

We calibrate the phases and squared visibilities in two different ways. In the first, which we refer to as Polycal, we fit polynomial functions in time to the calibrator observations. A zeroth-order function in time would represent a constant instrumental signal, while a very high order polynomial would indicate a highly variable instrumental signal. We calculate a different set of coefficients for each kernel phase as a function of time. We set the maximum allowed order to be the same for all kernel phases, choosing the one that minimizes the scatter in all of the calibrated observables without overfitting the calibrator measurements. We sample this polynomial function at the time of the target observations to find the instrumental signal. We subtract the instrumental signal from the raw kernel and closure phases and divide it into the raw squared visibilities. Figure 12 shows an example calibration for a single, linearly independent kernel phase (see Section 6). Figure 13 compares the calibrated kernel phases for polynomials with terms up to different orders in time and shows how the observed scatter changes with maximum order.

Polycal is best performed when calibrator observations can be taken before the first and after the last target observation. Figure 14 illustrates the potential issues that can occur when

Figure 13. Comparison of calibrated kernel phases for Polycal. The top panel shows histograms for an entire night of calibrated kernel phases for different maximum-order polynomials. The bottom panel shows the standard deviation of a night of calibrated kernel phases as a function of maximum order.

Figure 14. Polycal example for a single kernel phase. The top panel shows the calibrator (black points) and target (blue points) observations for a single kernel phase as a function of time. The plotted lines show polynomial fits as a function of time with different maximum orders. The bottom panel shows the calibrated kernel phases over time, with different colors indicating fits up to different orders in time. Since the best-fit first-order polynomial was nearly identical to the best-fit zeroth-order polynomial, the red squares and red dot-dashed line are plotted over the gray circles and gray dashed line. Here a second-order polynomial was eventually chosen, since it led to the lowest scatter in the calibrated observables.

calibrator observations do not bookend the target observations. For this data set, the large scatter in the calibrator observations leads to best-fit polynomials that prefer large coefficients for higher-order terms. These high-order polynomials match the first calibrator observation well but have unrealistic values at the earlier time of the first target observation (see purple dotted line and solid green line in Figure 14). Second-order polynomials often provide calibration that is comparable to or better than higher orders. They also tend to avoid these issues when target observations are not bookended by calibrator measurements.

8.2. Optimized Calibrator Weighting

The second calibration method is an optimized calibrator weighting detailed in Kraus & Ireland (2012) and Ireland (2013), similar to the “locally optimized combination of images” (LOCI; Lafrenière et al. 2007) algorithm applied in filled-aperture direct imaging. Here, calibrator weights are first chosen to minimize the χ^2 of the null model (zero phase signal). They are then chosen iteratively to minimize the χ^2 of the global best-fit model until the best fit converges. Following the notation in Ireland (2013), for each uncalibrated set of kernel phases $\mathbf{x}_t$, we search for a set of N weights (a_k , where k goes from 1 to N), to average the N calibrator measurements. The calibrated kernel phases for one pointing ($\mathbf{x}_c$) are

$$\mathbf{x}_c = \mathbf{x}_t - \sum_{k=1}^N a_k \mathbf{x}_k. \quad (14)$$

The χ^2 of the null model is then

$$\chi^2 = \sum \frac{\mathbf{x}_c^2}{\sigma_c^2}, \quad (15)$$

where σ_c are the calibrated kernel phase errors, given by

$$\sigma_c^2 = \sigma_t^2 + \sum_{k=1}^N a_k^2 \sigma_k^2. \quad (16)$$

Several different likelihood functions can be maximized to find the optimal calibrator weights. The simplest likelihood function would be

$$L = \exp(-0.5 \chi^2). \quad (17)$$

However, maximizing this likelihood function could wash out true signals. To try to prevent this, Kraus & Ireland (2012) add a regularizer, π , to the likelihood function. Arbitrary functions can be used for π ; the calibration used in Kraus & Ireland (2012) applies the following regularization:

$$\pi = \exp\left(-0.5 a_k^2 \sum \frac{\sigma_k^2}{\sigma_t^2}\right), \quad (18)$$

which punishes large weights and calibrator measurements with large errors. The likelihood function is then

$$L = \exp(-0.5 \chi^2) \times \pi. \quad (19)$$

This calibration is performed iteratively to constrain an additional “calibration error” term, Δ , intended to account for errors beyond the random component found in σ_t and σ_k . The Δ term is a constant added to all errors (calibrators and target), designed so that the reduced χ^2 of the calibrated kernel phases

Figure 15. LOCI regularization comparison. In the top panel, the black line shows the Polycal kernel phase distribution for data set 2 from Figure 16. The dashed blue line shows the LOCI calibration from Figure 16, the red dotted line a LOCI calibration with no “calibration error” term, and the green solid line a LOCI calibration with no regularization or “calibration error” term. The bottom panel shows the three LOCI calibrations’ absolute value kernel phases plotted against those for the Polycal calibration.

is equal to 1:

$$\chi_r^2 = \frac{1}{N} \sum \frac{\mathbf{x}_c^2}{\sigma_c^2} = 1. \quad (20)$$

Figure 15 demonstrates the ability of the regularizer π and the calibration error term Δ to bias the calibration. The blue line/points show the LOCI calibrated kernel phases for data set 1 in Figure 16. The red line/points show the LOCI calibration without using the calibration error term, and the green line/points show the LOCI calibration without using Δ or π . The discrepancy in the overall distribution of kernel phases (top panel), as well as in the values for the individual kernel phases (bottom panel), shows that the calibration can change significantly depending on the choice of error scaling and regularizer. With no error scaling or regularizer, the LOCI-like calibration could eliminate all astrophysical signal as the number of calibrator measurements becomes large. Thus, care must be taken in choosing an appropriate regularizer.

Figure 16. Comparison of Polycal and LOCI-like calibrated kernel phases for two nights of data. The top two panels show the histograms of calibrated kernel phases for Polycal in black and LOCI in blue. The bottom two panels show the absolute value of the LOCI kernel phases plotted against those for the Polycal calibration. The dashed line shows a 1:1 correlation.

8.3. Calibration Comparison

Figure 16 compares the Polycal kernel phases to the LOCI-like kernel phases using the likelihood function defined by Equation (19) for two different data sets. The histograms in the top two panels show that the LOCI-like calibration does not decrease the scatter much more than the Polycal calibration. The bottom two panels show the absolute value of the LOCI kernel phases plotted against those for the Polycal calibration. Deviations from the dashed line in these panels (indicating a 1:1 relationship) show differences between the two calibration methods for each kernel phase. While LOCI does not change the distribution of all kernel phases significantly, it does change the individual kernel phase values.

To compare the LOCI and Polycal calibrations under different noise conditions, we injected binary signals into L-band observations of two unresolved calibrator stars in two data sets. The mean uncalibrated kernel phase scatter for the two data sets was $1^{\circ}.1$ and $1^{\circ}.5$, leading to calibrated scatters of $0^{\circ}.3$ and $0^{\circ}.6$, respectively, using Polycal. We first checked for both sets of observations that no significant companion signals

existed, by calibrating each star using the other and fitting companion models to the calibrated data. We then chose one star as a mock target in which to inject companion signals, and we used the other to calibrate it using both methods. We fit the Polycal kernel phases a single time using the open-source Markov Chain Monte Carlo (MCMC) package `emcee` (Foreman-Mackey et al. 2013). For LOCI we calibrated the kernel phases iteratively, calibrating toward the global best-fit model until the fit converged.

We injected fake signals at a separation of 80 mas with contrasts ranging from $\Delta_L = 1$ to 7. Polycal could recover slightly higher contrast companions than LOCI, since the optimized calibration washed out faint input signals. LOCI could not recover companions with contrasts fainter than $\Delta_L \sim 5.5$ mag and $\Delta_L \sim 7$ mag for uncalibrated kernel phase scatters of $1^{\circ}.5$ and $1^{\circ}.1$, respectively. Here, the initial LOCI calibration left behind a few noise spikes with likelihoods comparable to the input companion model’s. In this scenario it would be possible to calibrate toward the wrong companion signal by selecting a noise spike rather than the true signal. This highlights the importance of quantifying Type 1 and Type 2

errors in NRM companion modeling (e.g., Sallum et al. 2015a). At these contrasts, Polycal recovers the input signal in both data sets, although the parameter uncertainties are large. For brighter sources, both calibrations recover the true signal, resulting in comparable fit parameter errors for the 1°5 scatter data set. With lower scatter (1°1), the parameter uncertainties from the LOCI calibration were ~ 2 times larger.

The injected companion detections and kernel phase scatter comparison show that Polycal can provide comparable calibration to LOCI without regularization and error scaling. For at least the case of a simple binary, Polycal results in similar or lower parameter uncertainties and recovers high-contrast companion signals more reliably. Furthermore, in this implementation the optimized calibration takes the science observations into account; it thus cannot be a completely unbiased measurement of the instrumental signal. This could be avoided by observing many calibrators and, for each target observation, computing an optimal weighting that minimizes the calibrator signal closest in time to the target observation. However, this is nearly identical to Polycal. Since Polycal depends only on the calibrator observations, requires no regularization or error scaling terms, and performs comparably to LOCI, for inexperienced users we recommend using simple calibrations like this over the optimized calibrator weighting.

We note that using negative calibrator weights is possible only with the LOCI calibration. This is important if the target and calibrator spectra differ, since dispersion can cause the target instrumental kernel phases to be best represented by the difference in calibrator kernel phases (e.g., Ireland 2013). When uncorrected, dispersion can mimic a close-in companion (e.g., Le Bouquin & Absil 2012). We avoid this by choosing calibrators whose fluxes are close to the target at both the wavefront sensing and science wavelengths. In situations where this is not possible, the optimized calibrator weighting may yield better results.

9. Image Reconstruction

Uneven and incomplete Fourier coverage, analogous to imaging in radio interferometry (e.g., Högbom 1974), makes image reconstruction an ill-posed problem. With perfect knowledge of the complex visibilities, synthesizing an image would only require an inverse Fourier transform. However, the number of pixels in a reconstructed image is much larger than the number of Fourier phases and amplitudes that we can constrain. Furthermore, since the Earth's atmosphere corrupts the individual phases, we measure combinations of phases that are robust to atmospheric errors. As a result, we cannot independently measure all of the phases for the array and have less phase information than in the radio case. This lack of Fourier information means that an infinite number of model images can provide comparable fits to the observations. To reconstruct images from NRM observations, we thus use algorithms that maximize the likelihood of the data while also satisfying a regularization constraint (e.g. Thiebaud & Giovannelli 2010). Regularizers, which were first used to compensate for incomplete Fourier coverage in the radio (e.g., Ables 1974; Högbom 1974; Titterton 1985), are chosen by hand to impose prior knowledge such as positivity, smoothness, sharp edges, or sparsity (e.g., Renard et al. 2011).

9.1. Optimization Engines

While many different optimization engines exist for finding the best image, they can be split into two general categories. Deterministic algorithms take steps proportional to the gradient of the likelihood function at the current image state. These include the steepest descent method in the Building Block Method (e.g., Hofmann & Weigelt 1993), the constrained semi-Newton method (used in the algorithm MiRA; e.g., Thiébaud 2008), and the trust region method (used in the algorithm BSMEM; e.g., Buscher 1994; Baron & Young 2008). Stochastic algorithms, on the other hand, involve moving flux elements randomly in the image plane during iterations in MCMC. One stochastic algorithm is MACIM (Ireland et al. 2006), a simulated annealing method that accepts or rejects images based on a temperature parameter. A newer stochastic algorithm is SQUEEZE (Baron et al. 2010), which can perform parallel simulated annealing, where the results of several simulated annealing chains are averaged. SQUEEZE can also perform parallel tempering, where several MCMC chains at different temperatures can exchange information as they run.

9.2. Regularizers

Many regularizers can be used to constrain the image reconstruction in different ways. Some regularizers punish excursions from a prior image, or a default image in the absence of any prior information. An example is the Maximum Entropy, or MEM, regularization, which favors the least informative reconstruction by maximizing an entropy term in the likelihood function (e.g., Haniff et al. 1987; Buscher 1994). The results of regularizers such as MEM depend heavily on the choice of prior image (e.g., Baron 2016).

Compressed sensing regularizations decompose the image into a linear combination of basis functions and then apply constraints. A simple example would be sparsity in the pixel basis. This was first enforced by the Building Block Method (Hofmann & Weigelt 1993) and can result in large areas with low flux. Sparsity can also be imposed in the gradient of the image, which would preserve uniform flux and edges (e.g., Renard et al. 2011). A more recent regularization that can be applied in SQUEEZE is one that imposes sparsity after decomposing the image into a wavelet basis (e.g., Baron et al. 2010). The wide variety of available regularizers, of which these are just a few examples, shows that care must be taken in choosing one and in understanding its effect on the final reconstructed image.

9.3. Degeneracies

Comparing the reconstructed images from LBT NRM observations of MWC 349A (Sallum et al. 2017) and LkCa 15 (Sallum et al. 2015b) illustrates the effects of prior choice and array configuration on the final reconstructed images. The MWC 349A data set has two pointings with $\sim 13^\circ$ change in parallactic angle, while the LkCa 15 data set has 15 pointings with parallactic angles between -65° and 65° . The LkCa 15 observations were taken with the LBT in single-aperture mode, using only baselines within each of the two primary mirrors. The MWC 349A data were taken in dual-aperture mode, but due to the small amount of sky rotation, these data only have

Figure 17. Best-fit skewed ring+delta function model for observations of MWC 349A published in Sallum et al. (2017). This image was used as the input for the simulations shown in the bottom row of Figure 18 and in Figure 19.

Figure 18. From Sallum et al. (2017), reconstructed images from observations of MWC 349A (top row) and simulated observations of the best-fit skewed disk+delta function model (bottom row; see Figure 17). The left and right columns show reconstructions for delta function and Gaussian priors, respectively.

triple the single-aperture resolution for a small range of position angles.

Figure 17 shows the best-fit skewed ring+delta function model for the MWC 349A observations. This model was determined by fitting the complex visibility data directly, and this does not depend on image reconstruction.

The top row of Figure 18 shows images reconstructed using BSMEM for the data, and the bottom row is for simulated observations of the model image in Figure 17. The left and right columns show images reconstructed using delta function

and Gaussian priors. The simulations show that observations of the source shown in Figure 17 lead to different reconstructed images depending on the choice of prior. The delta function prior results in a single, bright pixel, surrounded by a dark hole roughly the same size as the central part of the synthesized beam (contours in Figure 18). The Gaussian prior puts the same fractional flux into the central part of the beam as there is in the single pixel at the center of the delta function reconstruction.

Figure 19 shows the effects of improving the sky rotation (middle panel) and resolution (right panel) of the observations. Reconstructions improve with more complete Fourier coverage and longer baselines, showing the true extent and position angle of the input model despite an inadequate prior. The LkCa 15 reconstructions (see Figure 20) illustrate this as well. Again, the Gaussian prior results in a region of emission roughly the size of the beam that has the same fractional flux as the central pixel in the delta function reconstruction. However here the sources are at large enough separation (compared to the array resolution) that their portion of the image is less affected by the choice of prior. These degeneracies and their different effects on different data sets highlight the need for modeling in interpreting reconstructed images.

9.4. Simulated Data Reconstructions: BSMEM versus SQUEEZE

Since choices made during the image reconstruction process can change the final image, we use simulated data, rather than observations, to compare the reconstructed images to eight different source brightness distributions. The point of these reconstruction tests is to evaluate each algorithm's performance with no prior knowledge of the source morphology. Thus, we run each algorithm in identical configurations regardless of the input model. The detailed implementation and results of these simulations can be found in the Appendix. We compare two image reconstruction algorithms, BSMEM (Buscher 1994; Baron & Young 2008) and SQUEEZE (Baron et al. 2010), that are often used on NRM data sets. BSMEM is a deterministic algorithm with an MEM regularizer, while SQUEEZE is a stochastic algorithm with a variety of available regularizers.

With no knowledge of the source morphology, and with the potential for poor error bar estimation, BSMEM performs better than SQUEEZE. The BSMEM reconstructions vary less when the data error bars are over- or underestimated and also when the relative closure phase and squared visibility errors vary. Weighting the data by baseline length changes the results of both algorithms, with SQUEEZE providing better reconstructions of extended sources than BSMEM when long baselines were downweighted. Both BSMEM and SQUEEZE fail to converge when certain initial images are used for some input models (delta function initial images for BSMEM, and Gaussian initial images for SQUEEZE). Neither algorithm behaves so poorly that the source morphology is completely unrecognizable, but recovering features like close-separation point sources is more difficult with SQUEEZE under certain conditions.

We note that, with prior knowledge of the source morphology, the relative performances of BSMEM and SQUEEZE may change as more optimal regularizations can be applied. For example, using BSMEM with a delta function prior may be optimal when the target is expected to be a collection of point sources. When large areas of extended emission are expected,

Figure 19. Simulated observations of the model shown in Figure 17. The leftmost panel shows the MWC 349A (u, v) coverage and sky rotation. The middle panel shows the same mask configuration but with very dense sky rotation, and the right panel shows the dense sky rotation case with a mask having twice the resolution.

Figure 20. Reconstructed images from the L' LBT NRM data set published in Sallum et al. (2015b). The left panel shows the results of using a delta function prior, and the right panel shows those using a Gaussian prior. The two images are shown on the same color scale.

using SQUEEZE with a weighting scheme that upweights short baselines may yield the best results.

9.5. NRM on GMT: Comparison with LBTI

GMT's expected performance makes NRM appealing for imaging at small angular separations. The expected wavefront error due to imperfect segment phasing is ~ 50 nm rms for bright guide stars (e.g., Quirós-Pacheco et al. 2016). Furthermore, while next-generation adaptive optics systems are designed to reach rms wavefront errors less than ~ 100 nm, their measured wavefront errors are ~ 100 – 150 nm (e.g., Macintosh et al. 2014). These wavefront errors, which would make traditional high-resolution imaging difficult, can be characterized and calibrated using a non-redundant mask.

To compare the current state of the art to future facilities, we simulate observations with the dual-aperture LBT to those for a hypothetical 12-hole mask made for GMT. For both masks, we simulate observations with 20 pointings having parallactic angles between -65° and 65° , comparable to the sky rotation coverage in the 2014 December data set. We also simulate observations with poorer sky rotation—two pointings with

Figure 21. Fourier coverage for the 12-hole LBT mask (left) and a hypothetical 12-hole GMT mask (right). The top row shows parallactic angle coverage between -5° and 5° , and the bottom row shows -65° to 65° .

parallactic angles of -5° and 5° . This sky rotation is comparable to the amount observed for the 2012 dual-aperture LBT observations of MWC 349A (Sallum et al. 2017). Figure 21 shows the Fourier coverage for each facility and sky rotation case.

We reconstruct images for all eight objects in each of these four cases. Figure 22 shows the results using BSMEM. For the good sky rotation case, LBT and GMT produce comparable reconstructed images. With poorer sky rotation, the LBT's uneven Fourier coverage leads to poorer reconstructed images. The lower resolution in one direction smears flux out along the position angle where the beam is wider. However, even with smaller sky rotation coverage, none of the reconstructed images are so poor that the input source morphology is not recovered.

BSMEM LBT vs GMT Fourier Coverage Test

Figure 22. Images reconstructed from simulated observations using BSMEM with different masks and sky rotation coverage. The top row shows the input model image. The next two rows show observations with 20 pointings having sky rotation angles between -65° and 65° for the LBT and hypothetical GMT masks, respectively. The last two rows show observations with two pointings having sky rotation angles of -5° and 5° for the LBT and GMT masks. The point sources in Models 1, 2, 6, 7, and 8 in the top row have scaled brightnesses and sizes for ease of viewing. Figure 21 shows the Fourier coverage of these four data sets.

Since most data sets will have sky rotation coverage between the “Good” and “Bad” cases, Figure 22 shows that the co-phased LBTI can already provide comparable imaging to that possible with GMT.

While LBTI’s long baselines can achieve GMT-like resolution now, GMT’s symmetric aperture will provide more even Fourier coverage. GMT will thus make detections more efficiently for objects that are unresolved with baselines less than ~ 8 m. Here, GMT can provide higher-quality image reconstructions from observations with less sky rotation (see Figure 22). Furthermore, significantly increasing the number of holes in the LBTI mask requires that the holes themselves be much smaller, or else baselines bleed into one another in the Fourier plane. GMT’s larger collecting area would allow for many more holes with the same diameter as the LBTI 12-hole mask. This would increase efficiency by allowing for shorter exposures and by increasing the amount of recoverable phase information; it would also make GMT more sensitive to fainter sources.

9.6. Summary: Image Reconstruction Strategies

The ambiguities described in Sections 9.3 and 9.4 and in the Appendix show that many factors affect the quality of reconstructed images. Different algorithms behave differently depending on error bar over- or underestimation, as well as

the choice of priors and baseline weighting schemes. Furthermore, different sources will reconstruct with different quality depending on their size relative to the array baselines (see Figure 19), as well as the sky rotation coverage (see Figure 22). We thus recommend reconstructing images for simulated data to compare to real observations. This is useful for understanding what degeneracies and ambiguities exist for particular algorithms and observing strategies.

Particular observation strategies can also be useful for interpreting reconstructed images. Comparing reconstructed images from multiwavelength data sets can distinguish between different source morphologies that could lead to similar image reconstructions. For example, the position of a companion candidate should remain fixed with wavelength, while quasi-static speckle locations will change with wavelength. Multiwavelength observations can also be compared to radiative transfer modeling to differentiate scattered light scenarios from thermal emission. Furthermore, passive, thermally emitting disk sizes should increase with increasing wavelength, because dust farther from the star will be cooler and emit at longer wavelengths. Since disks can masquerade as companions for certain array resolutions and sky rotation coverage, this wavelength dependence may be useful for distinguishing between these scenarios. Multiepoch data sets are even more

powerful for breaking the degeneracies between disks and companions, since static disks cannot cause companion signals with smoothly changing position angles (e.g., Sallum et al. 2016).

10. Conclusions

We described observational and data reduction strategies for NRM observations with specific examples for individual data sets. We showed how image calibrations such as channel bias subtraction, flat-fielding, and bad pixel corrections affect the uncalibrated and calibrated observables for LMIRCam. We recommend checking reduction steps in this way for each data set since factors such as observing conditions and detector features can change the relative importance of these calibrations.

We explored different closure phase and squared visibility calculation methods. The closure phases calculated using the unwindowed, “Monnier” method resulted in lower scatter than all single-pixel calculation methods. Kernel phase projections that took into account information from the mask alone (the “Martinache” projection) led to lower scatter and fractional errors than projections that diagonalized the average closure phase covariance matrix. While windowing did not change the squared visibilities dramatically for observations taken in photometric conditions, data sets with more variable background levels could benefit from windowing. We described different calibration strategies for the phases and visibilities and made a case for applying simple calibrations rather than optimized ones that have more arbitrary biases.

We presented image reconstruction tests using the BSMEM and SQUEEZE algorithms for eight models representing potential NRM science targets. Overall, without prior knowledge of the source morphology and with the potential for poorly estimated error bars, BSMEM led to more robust reconstructed images than SQUEEZE. However, neither algorithm led to such poor reconstructed images that the source morphology could not be identified. These simulations demonstrated the utility of comparing simulated reconstructions to those for real observations. We also discussed ways in which multiwavelength and multiepoch data sets can be useful for interpreting reconstructed images.

We compared the imaging capabilities of the LBTI 12-hole mask to a hypothetical 12-hole mask designed for GMT. For good sky rotation coverage and typical noise levels, the dual-aperture LBTI and the GMT produced nearly identical reconstructed images. The dual-aperture LBTI can thus already provide GMT-like NRM imaging. GMT’s larger collecting area and more even Fourier coverage will enable detections at similar angular separations with lower sky rotation and a shorter total integration time. Given the expected performance of GMT’s segment phasing and adaptive optics, NRM remains a promising method for high-resolution imaging.

This material is based on work supported by the National Science Foundation under A.A.G. grant No. 1211329, by the National Science Foundation under grant No. 1228509, and by the National Science Foundation Graduate Research

Figure 23. Comparison of simulated squared visibilities (left) and closure phases (right) to those observed for the unresolved calibrator stars observed in 2014 December. The black histograms show the real observations, and the red dashed curves show the simulated data.

Fellowship under grant No. DGE-1143953. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. This research made use of Astropy, a community-developed core Python package for Astronomy (Astropy Collaboration et al. 2013). This research made use of the open-source software packages Scipy (Jones et al. 2001), Numpy (Van Der Walt et al. 2011), and matplotlib (Hunter 2007).

Appendix Image Reconstructions for Simulated Data

We simulate observations of the following sources to test two image reconstruction algorithms, BSMEM and SQUEEZE, on potential NRM science targets.

1. Model 1: A 1 mag contrast binary with a separation of 100 mas.
2. Model 2: A multiple system with contrasts ranging from 1 to 5 mag and separations ranging from 30 to 70 mas. Keck surveys of young stars (e.g., Kraus et al. 2011) detected companions of roughly these angular separations and magnitude ranges (and those in Model 1).
3. Model 3: A uniform circle of diameter 200 mas with spots having contrasts of 2–4 mag. NRM is used to image circumstellar disks, which may have nonuniformities such as hot spots. This case is also interesting for long-baseline optical interferometry, which is often used to image stellar surfaces.
4. Model 4: A uniform ellipse with a major axis of 200 mas and axis ratio of 0.33.
5. Model 5: A skewed ellipse with a major axis of 200 mas, axis ratio of 0.33, and 50% skew amplitude. This case and Model 4 are meant to represent the circumstellar disks that may have skew or outflows that have been

BSMEM Initial Image Test

Figure 24. Images reconstructed from simulated observations using BSMEM with different initial images. The top row shows the input model image, and the next three rows show reconstructions with flat, Gaussian, and delta function initial images, respectively. The point sources in Models 1, 2, 6, 7, and 8 in the top row have scaled brightnesses and sizes for ease of viewing.

SQUEEZE Initial Image Test

Figure 25. Images reconstructed from simulated observations using SQUEEZE with different initial images (rows). The top row shows the input model image, and the next three rows show reconstructions with flat, Gaussian, and delta function initial images, respectively. The point sources in Models 1, 2, 6, 7, and 8 in the top row have scaled brightnesses and sizes for ease of viewing.

imaged by NRM studies (e.g., Danchi et al. 2001; Sallum et al. 2017).

- Model 6: A point source plus a uniform ring with a diameter of 300 mas and Gaussian cross section, meant to emulate a

transition disk with no detectable companions. This diameter is roughly the angular diameter of transition disk clearings for the most nearby potential targets, corresponding to a hole diameter of 84 au at 140 pc (e.g., Andrews et al. 2011).

Figure 26. Images reconstructed from simulated observations using BSMEM with different error bar scalings. The top row shows the input model image, and the following rows have all closure phase and squared visibility errors scaled by factors of 4.0, 2.0, 1.0, 0.5, and 0.25 from top to bottom. The point sources in Models 1, 2, 6, 7, and 8 in the top row have scaled brightnesses and sizes for ease of viewing.

7. Model 7: A point source plus a skewed Gaussian ring with a diameter of 300 mas and 70% skew amplitude. Skew has been observed in NRM observations of transition disks (e.g., Huélamo et al. 2011; Cheetham et al. 2015; Sallum et al. 2015a).
8. Model 8: The multiple system in Model 2 surrounded by the skewed Gaussian ring in Model 7. This is meant to imitate companion searches in skewed transition disks. Companions with separations and contrasts similar to those in Model 2 have been found in NRM observations of transition and circumbinary disks (e.g., Ireland & Kraus 2008; Kraus & Ireland 2012; Sallum et al. 2015b).

All data sets are generated using the 12-hole mask installed in LBTI/LMIRCam in dual-aperture mode. To create data with realistic amounts of correlated noise, we add Gaussian noise to the complex visibility amplitudes (A) and phases (ϕ) before calculating the squared visibilities and closure phases. As shown in Figure 23, amplitude and phase scatters of 0.065 and $1^\circ.15$ match the observed scatter in the 2014 December data set. The propagated uncertainties for the closure phases and

squared visibilities are then

$$\sigma_{\text{CP}} = \sqrt{3} \sigma_\phi \quad (21)$$

and

$$\sigma_{V^2} = \sqrt{2A} \sigma_A. \quad (22)$$

Unless specified otherwise, we use these noise levels and roughly the same sky rotation coverage as in the 2014 December data set—20 pointings with parallactic angles between -65° and 65° .

We reconstruct images using the BSMEM (Buscher 1994) and SQUEEZE (Baron et al. 2010) algorithms. The total image field of view is always set to 780 mas, corresponding to the resolution of the shortest baseline in the mask, and the pixel scale is always set to 2 mas. To calculate residuals, we first align the output image with the input model via cross-correlation and then sum the square of the images' difference. We perform BSMEM reconstructions in “Classic Bayesian”

Figure 27. Images reconstructed from simulated observations using SQUEEZE with different error bar scalings. The top row shows the input model image, and the following rows have all closure phase and squared visibility errors scaled by factors of 4.0, 2.0, 1.0, 0.5, and 0.25 from top to bottom. The point sources in Models 1, 2, 6, 7, and 8 in the top row have scaled brightnesses and sizes for ease of viewing.

mode, running 200 iterations for every input data set. For SQUEEZE, we run four MCMC chains in parallel tempering mode for 1000 iterations. Since the point of these imaging experiments is to test the algorithms with no knowledge of the input source, we do not optimize any regularization or noise scalings when reconstructing images. The results of all the reconstructions for both algorithms are discussed here, with the images shown in Figures 24–31.

A.1. Initial Images

We test the effect of using different starting images (flat, Gaussian, and δ functions) in both BSMEM and SQUEEZE for each of the eight sources. The Gaussian initial images have FWHMs of 200 mas. Figure 24 shows the results for BSMEM. Here, starting with a flat initial image or a centered Gaussian provides the best reconstruction with no knowledge of the true source morphology. The binary and multiple reconstructions have lower residuals with a delta function initial image. However, for extended sources, BSMEM cannot converge to a realistic image in classic Bayesian mode with a delta function prior. The flat initial image results also have vertical or

horizontal striping in some cases. This is because BSMEM does not automatically recenter when reconstructing images non-interactively. Recentering by hand during the reconstruction would eliminate these artifacts.

Figure 25 shows the same tests using SQUEEZE. Here, only the flat and δ function initial images result in reconstructions that match the input images. With 1000 steps, SQUEEZE does not reproduce the input model starting from a Gaussian initial image. Comparing the reconstructions of models that include compact point sources (i.e., Models 1, 2, and 8) between BSMEM and SQUEEZE shows that SQUEEZE is more likely to blend point sources together. As a result, BSMEM more reliably reproduces the point sources within the ring in Model 8. Conversely, comparing Models 3, 4, and 5 between the two algorithms shows that SQUEEZE is less prone to over-resolve extended emission than BSMEM.

A.2. Total χ^2 Scaling

We scale the total χ^2 of each data set to explore the dependence of each algorithm on correct error bar estimation. We multiply all the closure phase and squared visibility error

BSMEM Weighting by Baseline Length Test

Figure 28. Images reconstructed from simulated observations using BSMEM with different baseline weightings. The top row shows the input model image. The next rows show reconstructions for data sets whose assigned errors depend on baseline length raised to powers of 2, 1, 0, -1 , and -2 from top to bottom. The upper rows weight short baselines more heavily, while the lower rows weight longer baselines more heavily. The point sources in Models 1, 2, 6, 7, and 8 in the top row have scaled brightnesses and sizes for ease of viewing.

bars by constant values of 0.25, 0.5, 1.0, 2.0, and 4.0, without changing the actual noise added to the model. For each algorithm we use the initial image that produced the best results in the previous section—a large Gaussian for BSMEM, and a flat image for SQUEEZE.

Figures 26 and 27 show the results for BSMEM and SQUEEZE, respectively. As the error bar scaling decreases, both algorithms produce tighter reconstructions of the point sources in Models 1, 2, and 8. With overestimated error bars, SQUEEZE’s blurring of close-separation point sources becomes more pronounced (see Model 8 in Figure 27). In this regime, SQUEEZE also puts regions of very low flux directly opposite asymmetric emission (see Model 1, top two panels in Figure 27). While SQUEEZE blurs extended emission (e.g., Models 3, 4 and 5) more with overestimated errors, BSMEM creates blurrier reconstructions of extended emission when the errors are underestimated.

Neither algorithm causes huge qualitative mismatches between the inputs and reconstructed images for any of the

error scalings. This suggests that poorly estimated errors will not degrade reconstructed images to the point that the source morphology could not be recovered. However, for SQUEEZE especially, poorly estimated error bars may result in nondetections of close-in companions. Comparing Figures 26 and 27 shows that BSMEM produces more reliable reconstructions of most models with under- or overestimated errors.

A.3. Weighting by Baseline Length

To test the effects of different (u, v) weighting schemes, we reconstruct images using closure phase and squared visibility errors that depend on baseline length. We assign errors in the following way:

$$\sigma = B^p \sigma_{\text{med}}, \quad (23)$$

where B is the baseline length for squared visibilities and the mean baseline length for closure phases, and σ_{med} is the median observed error bar. Here, $p = 0$ corresponds to uniform

SQUEEZE Weighting by Baseline Length Test

Figure 29. Images reconstructed from simulated observations using SQUEEZE with different baseline weightings. The top row shows the input model image. The next rows show reconstructions for data sets whose assigned errors depend on baseline length raised to powers of 2, 1, 0, -1 , and -2 from top to bottom. The upper rows weight short baselines more heavily, while the lower rows weight longer baselines more heavily. The point sources in Models 1, 2, 6, 7, and 8 in the top row have scaled brightnesses and sizes for ease of viewing.

weighting, $p < 0$ upweights long baselines, and $p > 0$ upweights short baselines. To separate the relative weighting effects from a varying total χ^2 , we scale the errors so that the total χ^2 of the initial image (Gaussian for BSMEM and flat for SQUEEZE) is constant. We reconstruct images for p values of -2 , -1 , 0 , 1 , and 2 .

Figure 28 shows the results for BSMEM, and Figure 29 shows the results for SQUEEZE. Both algorithms generally produce better reconstructions of extended emission when the short baselines are upweighted. SQUEEZE’s reconstructions of Models 4 and 5 show a more marked improvement with upweighted short baselines than BSMEM’s. For both algorithms, upweighting the long baselines improves the recovery of compact sources, especially when they are alongside an extended component like in Model 8 (skewed ring+multiple system). However, both SQUEEZE and BSMEM tend to over-resolve extended components when the long baselines are upweighted (see bottom rows of Figures 28 and 29). As in the total χ^2 tests, no weighting

scheme creates large qualitative differences between the inputs and reconstructions.

A.4. Closure Phase and Squared Visibility Weighting

We test the effect of reconstructing images using different relative closure phase and squared visibility weights. We scale all closure phase errors by a constant factor f_{cp} , while leaving the squared visibility errors alone ($f_{v^2} = 1.0$). We then rescale f_{cp} and f_{v^2} by the same constant to keep the χ^2 of the prior image for each algorithm fixed. We use error scaling ratios (f_{cp}/f_{v^2}) of 4.0, 2.0, 1.0, 0.5, and 0.25, where larger values downweight closure phases and smaller values upweight closure phases relative to squared visibilities.

Figures 30 and 31 show the results for BSMEM and SQUEEZE, respectively. BSMEM shows little to no difference in reconstruction quality across all weighting schemes. The exception to this is the $f_{cp}/f_{v^2} = 4$ reconstruction of Model 3,

BSMEM CP and V^2 Weighting Test

Figure 30. Images reconstructed from simulated observations using BSMEM with different relative closure phase and squared visibility weightings. The top row shows the input model image, and the next rows show reconstructions where the closure phase errors have been scaled by a factor of 4.0, 2.0, 1.0, 0.5, and 0.25 top to bottom, relative to the squared visibilities. The point sources in Models 1, 2, 6, 7, and 8 in the top row have scaled brightnesses and sizes for ease of viewing.

which has a poorer reconstruction of the spots on the uniform disk.

SQUEEZE produces more reliable reconstructions of Models 1, 2, 7, and 8 when the closure phases are downweighted relative to the squared visibilities. It does not recover the compact components in these models when the squared visibilities are downweighted. However, the reconstructions of Model 6 (uniform ring+delta function) are worse with downweighted closure phases. Here SQUEEZE allows compact emission to appear in two locations within the ring. This could result from averaging

subsequent images in the MCMC chains that were not centered consistently. Model 4 appears particularly asymmetric for large closure phase errors. Increasing the power-law exponent of the chain temperature distribution, and thus allowing SQUEEZE to explore larger parts of parameter space more quickly, makes the reconstructions more symmetric. This suggests that it takes longer for SQUEEZE to converge to the true source morphology when there is not much information in the closure phases. Overall, the relative closure phase and squared visibility weights change SQUEEZE's performance more dramatically than BSMEM's.

SQUEEZE CP and V^2 Weighting Test

Figure 31. Images reconstructed from simulated observations using BSMEM with different relative closure phase and squared visibility weightings. The top row shows the input model image, and the next rows show reconstructions where the closure phase errors have been scaled by a factor of 4.0, 2.0, 1.0, 0.5, and 0.25 top to bottom, relative to the squared visibilities. The point sources in Models 1, 2, 6, 7, and 8 in the top row have scaled brightnesses and sizes for ease of viewing.

ORCID iDs

S. Sallum <https://orcid.org/0000-0001-6871-6775>

J. Eisner <https://orcid.org/0000-0002-1031-4199>

References

- Ables, J. G. 1974, *A&AS*, **15**, 383
- Andrews, S. M., Wilner, D. J., Espaillat, C., et al. 2011, *ApJ*, **732**, 42
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Baldwin, J. E., Haniff, C. A., Mackay, C. D., & Warner, P. J. 1986, *Natur*, **320**, 595
- Baron, F. 2016, *ASSL*, **439**, 75
- Baron, F., Monnier, J. D., & Kloppenborg, B. 2010, *Proc. SPIE*, **7734**, 773421
- Baron, F., & Young, J. S. 2008, *Proc. SPIE*, **7013**, 3
- Beletic, J. W., Blank, R., Gulbransen, D., et al. 2008, *Proc. SPIE*, **7021**, 70210H
- Biller, B., Lacour, S., Juhász, A., et al. 2012, *ApJL*, **753**, L38
- Buscher, D. F. 1994, in *IAU Symp. 158, Very High Angular Resolution Imaging*, ed. J. G. Robertson & W. J. Tango (Dordrecht: Kluwer), **91**
- Cheetham, A., Huélamo, N., Lacour, S., de Gregorio-Monsalvo, I., & Tuthill, P. 2015, *MNRAS*, **450**, L1
- Danchi, W. C., Tuthill, P. G., & Monnier, J. D. 2001, *ApJ*, **562**, 440
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Guyon, O., Hinz, P. M., Cady, E., Belikov, R., & Martinache, F. 2014, *ApJ*, **780**, 171
- Guyon, O., Pluzhnik, E. A., Kuchner, M. J., Collins, B., & Ridgway, S. T. 2006, *ApJS*, **167**, 81
- Haniff, C. A., Mackay, C. D., Titterton, D. J., Sivia, D., & Baldwin, J. E. 1987, *Natur*, **328**, 694
- Hinz, P. M., Bippert-Plymate, T., Breuninger, A., et al. 2008, *Proc. SPIE*, **7013**, 28
- Hofmann, K.-H., & Weigelt, G. 1993, *Proc. SPIE*, **1983**, 203
- Högbom, J. A. 1974, *A&AS*, **15**, 417
- Huélamo, N., Lacour, S., Tuthill, P., et al. 2011, *A&A*, **528**, L7
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ingraham, P., Perrin, M. D., Sadakuni, N., et al. 2014, *Proc. SPIE*, **9147**, 914770
- Ireland, M. J. 2013, *MNRAS*, **433**, 1718
- Ireland, M. J., & Kraus, A. L. 2008, *ApJL*, **678**, L59
- Ireland, M. J., Monnier, J. D., & Thureau, N. 2006, *Proc. SPIE*, **6268**, 1T1
- Jennison, R. C. 1958, *MNRAS*, **118**, 276
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, *SciPy: Open Source Scientific tools for Python*, (<https://www.scipy.org>)
- Kraus, A. L., & Ireland, M. J. 2012, *ApJ*, **745**, 5
- Kraus, A. L., Ireland, M. J., Martinache, F., & Hillenbrand, L. A. 2011, *ApJ*, **731**, 8
- Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, **660**, 770
- Le Bouquin, J.-B., & Absil, O. 2012, *A&A*, **541**, A89
- Leisenring, J. M., Skrutskie, M. F., Hinz, P. M., et al. 2012, *Proc. SPIE*, **8446**, 4
- Macintosh, B., Graham, J. R., Barman, T., et al. 2015, *Sci*, **350**, 64
- Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, *PNAS*, **111**, 12661
- Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, **641**, 556

- Martinache, F. 2010, [ApJ](#), **724**, 464
- Mawet, D., Milli, J., Wahhaj, Z., et al. 2014, [ApJ](#), **792**, 97
- Mawet, D., Pueyo, L., Lawson, P., et al. 2012, *Proc. SPIE*, 8442, 844204
- Meimon, S., Petit, C., Fusco, T., & Kulcsar, C. 2010, [JOSAA](#), **27**, A122
- Monnier, J. D. 1999, PhD thesis, Univ. California
- Quirós-Pacheco, F., Conan, R., Bouchez, A., & McLeod, B. 2016, [Proc. SPIE](#), **9906**, 99066D
- Renard, S., Thiébaud, E., & Malbet, F. 2011, [A&A](#), **533**, A64
- Ruane, G., Mawet, D., Kastner, J., et al. 2017, [AJ](#), **154**, 73
- Sallum, S., Eisner, J., Close, L. M., et al. 2016, [Proc. SPIE](#), **9907**, 99070D
- Sallum, S., Eisner, J. A., Close, L. M., et al. 2015a, [ApJ](#), **801**, 85
- Sallum, S., Eisner, J. A., Hinz, P. M., et al. 2017, [ApJ](#), **844**, 22
- Sallum, S., Follette, K. B., Eisner, J. A., et al. 2015b, [Natur](#), **527**, 342
- Soummer, R., Pueyo, L., & Larkin, J. 2012, [ApJL](#), **755**, L28
- Thiébaud, E. 2008, [Proc. SPIE](#), **7013**, 70131I
- Thiébaud, E., & Giovannelli, J.-F. 2010, [ISPM](#), **27**, 97
- Titterton, D. M. 1985, *A&A*, **144**, 381
- Tuthill, P. G., Monnier, J. D., Danchi, W. C., Wishnow, E. H., & Haniff, C. A. 2000, [PASP](#), **112**, 555
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, arXiv:1102.1523