© ESO 2012

Improved CLEAN reconstructions for rotation measure synthesis with maximum likelihood estimation

M.R. Bell[⋆][1], N. Oppermann[1], A. Crai[2], and T.A. Enßlin[1]

[1] Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany
[2] Jacobs University, Campus Ring 1, 28759 Bremen, Germany

ABSTRACT

The CLEAN deconvolution algorithm has well-known limitations due to the restriction of locating point source model components on a discretized grid. In this letter, we demonstrate that these limitations are even more pronounced when applying CLEAN in the case of Rotation Measure (RM) synthesis imaging. We suggest a modification that uses Maximum Likelihood estimation to adjust the CLEAN-derived sky model. We demonstrate through the use of mock one-dimensional RM synthesis observations that this technique shows significant improvement over standard CLEAN and gives results that are independent of the chosen image pixelization. We suggest using this simple modification to CLEAN in upcoming polarization sensitive sky surveys.

Key words. Methods: data analysis - Techniques: polarimetric

1. Introduction

CLEAN (Högbom 1974) is a deconvolution algorithm that is widely used throughout astronomy. In CLEAN, a deconvolved image is constructed iteratively by locating the maximum in the image, adding a delta function at that location with some fraction of the peak brightness to a model, and then subtracting a point spread function (PSF) that has been shifted to the location of the model component and scaled by its strength. This is repeated until a user-defined stop condition has been reached.

The algorithm makes an implicit assumption that the image is well described by a collection of statistically independent point sources. In some cases this is an adequate assumption and experience has shown that even in circumstances where this assumption is not ideal, CLEAN is still able to produce qualitatively reasonable results. Because of its flexibility, simplicity, and speed, CLEAN has become the standard tool for image reconstruction in radio astronomical aperture synthesis imaging.

A limitation of CLEAN arises due to the discretization of the image space, and therefore one can only create model point sources on a regularly spaced grid. To properly remove the PSF pattern associated with a strong point source, it is crucial for the delta function representing the source in the CLEAN model to have an accurate position. While the PSF limits the minimum spatial scale of structures that can be resolved, the accuracy with which one can measure the location of a point source depends on the signal-to-noise ratio. In the high signal to noise regime the uncertainty in the location of a source can be much smaller than the PSF, and to measure the source location accurately with a single model point source one would have to create an image with a large number of pixels within the PSF. Such over-resolving is at best computationally wasteful and in some cases may make the imaging problem completely unfeasible. For example, even with a few pixels per resolution element, three-dimensional (3D) Faraday synthesis (Bell & Enßlin 2012)

images will easily exceed the available memory on a single computer.

CLEAN is able to partially overcome this issue by describing a single point source by a cluster of CLEAN components around the true source location. The strongest model point source will lie on the pixel closest to the source location. Weaker model points will be added to neighboring grid points such that when the collection of model sources is convolved by a Gaussian restoring beam that approximates the idealized PSF, the resulting image will be closer to the true source position and flux than that of a single model point located on the grid. Nevertheless, inaccuracies in image reconstructions due to the pixelization of the sky are a well known limiting factor to dynamic range when using CLEAN (Briggs & Cornwell 1992; Perley 1999; Cotton & Uson 2008; Yatawatta 2010).

Rotation Measure (RM) synthesis (Brentjens & de Bruyn 2005) is a promising new technique for studying magnetic fields with the new generation of broadband radio telescopes. RM synthesis allows for the separation of multiple sources of polarized emission along a line of sight when each is Faraday rotated by different amounts. It produces an estimate of the Faraday spectrum, or the polarized emission as a function of the Faraday depth, which is a quantity that measures the amount of Faraday rotation. The Faraday depth, ϕ, is proportional to $\int n_e B_z \mathrm{d}z$ where n_e is the density of thermal electrons and B_z is the component of the magnetic field along the line of sight. RM synthesis is similar to aperture synthesis imaging due to the Fourier relationship between the Faraday spectrum and the polarized intensity as a function of the squared wavelength. CLEAN has been proposed as a deconvolution technique to be applied to RM synthesis imaging (Heald et al. 2009), and is often referred to as RMCLEAN in this context. Several alternative image reconstruction techniques have been proposed (Frick et al. 2010; Li et al. 2011; Andrecut et al. 2012), but RMCLEAN has thus far been used most often as the deconvolution method of choice (Feain et al. 2009; Harvey-Smith et al. 2010; Mao et al. 2010; Bernardi et al. 2010;

⋆ e-mail: `mrbell@mpa-garching.mpg.de`

Brentjens 2011; Van Eck et al. 2011; Heald 2012; Mao et al. 2012a,b; van Weeren et al. 2012; Iacobelli et al. 2012).

In this letter we show that the limitations of CLEAN due to pixelization of the sky are particularly pronounced in the case of RM synthesis, and propose to use a Maximum Likelihood (ML) estimation procedure to improve the model obtained using the standard RMCLEAN algorithm. This method is similar to those proposed previously in the context of high-fidelity aperture synthesis imaging (El-Behery & MacPhie 1980; Yatawatta 2010; Bernardi et al. 2011). It is also similar to the method proposed by O'Sullivan et al. (2012), but is more easily scaled to higher dimensions and includes fewer assumptions about the Faraday spectrum. Our aim is to maintain or improve accuracy of the standard RMCLEAN algorithm while keeping the number of pixels to a minimum.

2. The algorithm

Here we consider the case of 1D RM synthesis where one is attempting to reconstruct the complex valued Faraday spectrum, s, from some polarized sky brightness data, d, that has been taken at many frequencies. We note that the following description is trivially extended to 3D as would be required in the case of Faraday synthesis (Bell & Enßlin 2012), which combines aperture and RM synthesis into a single procedure that provides significantly better results than when they are performed separately.

The Faraday depth coordinate axis (the image coordinate) will be denoted as ϕ, and the data coordinate is λ^2. For RM synthesis, the following measurement equation applies:

$$d(\lambda^2) = S(\lambda^2) \int d\phi\, s(\phi) e^{2i\phi\lambda^2} + n(\lambda^2). \quad (1)$$

Here n is a Gaussian additive noise term and S represents a sampling function defining the discrete λ^2 values at which measurements are obtained.

Using the CLEAN algorithm, one generates a list of model points $M = m_i, \phi_{i,i}$ with flux values and positions given by m_i and ϕ_i, respectively, where i is an index over the list of N_M model points. The representation of the model in data space is

$$\widetilde{d}_j = \sum_i m_i e^{2i\phi_i\lambda_j^2}, \quad (2)$$

where j is an index over the N_d values of λ^2 for which measurements have been made.

We assume that the probability of measuring the full data set, d, given the CLEAN model (i.e. the likelihood), is a product of the probabilities of measuring each individual data point, d_i. In this way, we assume the measurements to be independent of one another, thus making the likelihood

$$P(d|M) = \prod_j P(d_j|M) = \prod_j \sqrt{\frac{1}{2\pi\sigma_j^2}} e^{-\frac{1}{2\sigma_j^2}\left(d_j - \widetilde{d}_j\right)^2}. \quad (3)$$

Here, as mentioned previously, we assume Gaussian distributed noise, and σ_j^2 is the variance of the noise in channel j.

Our approach is to adjust the CLEAN model by maximizing the likelihood with respect to m_i and ϕ_i. For simplicity, we opt to work with the negative log-likelihood

$$H(d|M) = -\ln P(d|M) = \frac{1}{2} \sum_j \left[\frac{\left(d_j - \widetilde{d}_j\right)^2}{\sigma_j^2} + \ln\left(2\pi\sigma_j^2\right) \right]. \quad (4)$$
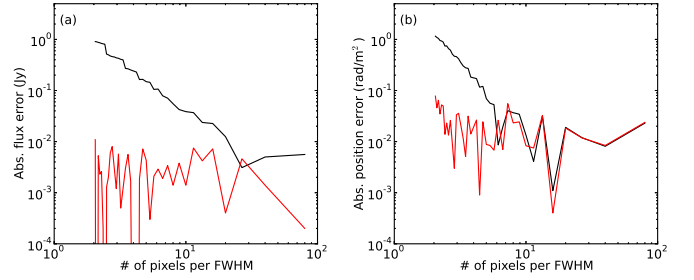
Fig. 1. A comparison of the (a) flux and (b) position errors as a function of image pixel size in reconstructions of a single point source using RMCLEAN (black line) and the ML algorithm described in this paper (red line). The reported errors are the average absolute errors from 100 simulated data sets.

The maximum likelihood solution is obtained by *minimizing* this function. There are many possible approaches to minimizing Eq. 4. The iterative procedure that we use is presented in Appendix A.
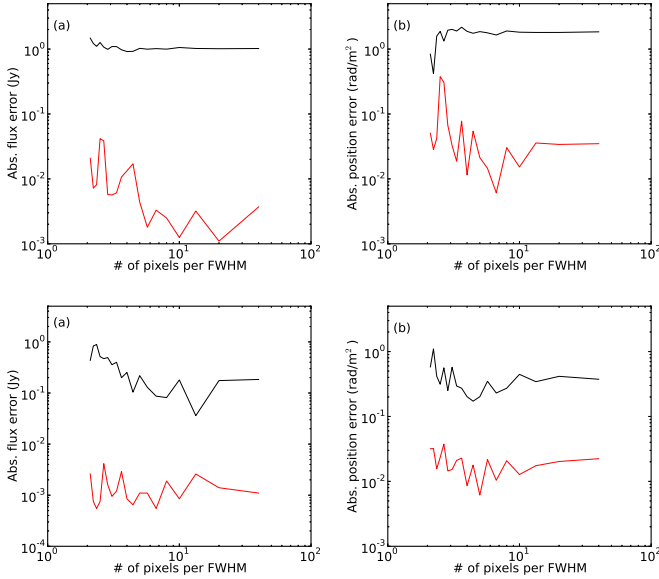
Minimization of Eq. 4 provides an estimate for the optimal model parameters under the assumption that the number of point sources in the initial CLEAN model is appropriate. This, however, is almost certainly not the case. CLEAN will naturally over-estimate the number of point sources that are required to model the data because it needs a cluster of sources in order to model a single source at an arbitrary location.

To find the most appropriate number of free parameters in the model that are supported by the data, we modify H to add a term that penalizes additional degrees of freedom. We use the so-called Bayesian Information Criterion (Schwarz 1978, BIC) to suit this role. The BIC, C, is given by

$$C = 2H(d|m) + 3N_m \ln N_d. \quad (5)$$

The algorithm proceeds as follows:

1. Start with a list of model point sources generated by RMCLEAN. Condense the model such that there is only one entry per pixel location.
2. Throw away any model entries with a flux below a user-defined threshold. This step removes model points that result from cleaning too deeply or that only function to slightly shift a single point source location. We use a threshold of twice the CLEAN image noise level. This step is not necessary, but can speed up computation in case there are many excess point sources in the initial model.
3. Take one or more steps to iteratively minimize Eq. 4 for the current model using e.g. the prescription outlined in Appendix A.
4. Attempt to merge nearby model points. Pairwise combine CLEAN components into a single component having a location at the flux-weighted mean location and a flux that is either the sum of the two model point fluxes or that is solved for using e.g. Eq. A.7. Accept the merged CLEAN component if C is reduced, otherwise revert to the previous model.
5. Iterate steps 3 and 4 until the fractional change in H from one step to the next is below a user-defined threshold.
6. Obtain a residual image by subtracting the new model from the data.
7. Convolve the new model with an idealized PSF and add it to the residual image.
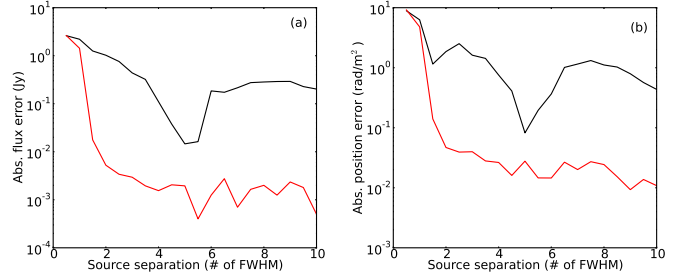
Fig. 3. As Fig. 2 but with fixed resolution and varying source separation, which is given as a multiple of the FWHM of the PSF (40 rad/m^2).


Fig. 2. As Fig. 1 but with the average errors of two sources that are separated by 2×FWHM (top row) and 6×FWHM (bottom row).

3. Demonstration

To demonstrate the effectiveness of the approach outlined above, we present a series of 1D RM synthesis mock observations of simple single and double point-source sky models. We generate simulated data according to Eq. 1 including Gaussian random white noise and sample using a frequency coverage with 800 channels equally spaced between 1 and 4.2 GHz. This frequency range is similar to that of the combination of the upgraded Very Large Array L- and S-band receivers and corresponds to a PSF with a full width at half maximum (FWHM) of the main peak of ∼ 40 rad/m^2. A Clark style CLEAN (Clark 1980) is then performed and the resulting model is used as input for the ML algorithm. The output images from both procedures are stored for comparison.

In Fig. 1 we compare the performance of standard RMCLEAN to that of the ML method at reconstructing the position and flux of a single point source as a function of pixel size. The source is located at roughly the same ϕ location at each resolution, but shifted slightly to ensure that it is always located one-third of the way between two pixels. The source flux is roughly 100 times the noise level. The plotted flux and position errors are the average of the magnitude of those from 100 trials with different noise realizations. Source position and flux are measured by locating the maximum of the modulus of the Faraday spectrum and fitting a Gaussian to the image around this point. We find that for RMCLEAN the position and flux errors increase significantly with pixel size while those for the ML approach remain nearly constant. Six or more pixels per FWHM of the PSF are required before the positional accuracy for RMCLEAN matches that of the ML approach. To obtain accurate fluxes with RMCLEAN, 40 or more pixels per FWHM are required.

For a sky model with a single source, RMCLEAN is found to work well as long as the image is sufficiently over-resolved. With more than one source, however, over-resolution is no longer effective, as shown in Fig. 2. Here we show positional and flux errors in the case where there are two sources along the line of sight. In one case, they are separated by 80 rad/m^2

(roughly 2×FWHM of the PSF) and in the other by 240 rad/m^2 (6×FWHM). The modulus and phase of the complex valued flux is the same for both sources. The errors that are plotted are the average of those from both sources (the errors for the individual sources were qualitatively similar). We find that increasing the resolution does not reduce errors in the RMCLEAN reconstruction and that on average these errors are two or more orders of magnitude larger than in the maximum likelihood reconstructions. This behavior is due to the interactions between the highly structured, complex valued PSFs associated with each point source. The exact dependence on the error as a function of source separation depends on the specific sampling function, and therefore PSF, being used.

We also find that the errors are larger for both approaches when sources are close together relative to the FWHM of the PSF. To test this further, we kept the image resolution fixed (at 5 rad/m^2 per pixel) and varied the separation between the two sources (again, each having the same flux and phase). In Fig. 3 we see that below a source separation of 2×FWHM the errors for both methods increase dramatically. Nevertheless, at separations of 1×FWHM or larger, the maximum likelihood approach fares significantly better than RMCLEAN.

Even though the errors increase dramatically as source separation decreases for both methods, we find that the ML reconstructions are still a significant improvement over RMCLEAN. In Fig. 4 we show reconstructed Faraday spectra from 500 data realizations of the same sky model. Although the ML results show a systematic reduction in flux and a shift in position relative to the sky model, two distinct sources are clearly identifiable. This is not the case in the RMCLEAN image. We performed similar tests for hundreds of combinations of source separation (between 1 and 6×FWHM), relative fluxes (flux ratios between 1 and 10), relative phases, and noise levels. In all cases the ML approach was an improvement over RMCLEAN, and it only showed significant errors in a few cases like the one shown in Fig. 4.

4. Conclusions

We have shown that the well-known limitations of the CLEAN algorithm that arise due to discretization of the sky are particularly pronounced in the case of RM synthesis. We find that accuracy of the results obtained by CLEAN depend strongly on the choice of pixelization, and significant over-resolution is required to obtain accurate measurements of flux and Faraday depth in the simplest case where there is a single source along the line of sight. Furthermore, RMCLEAN is unable to accurately recon-
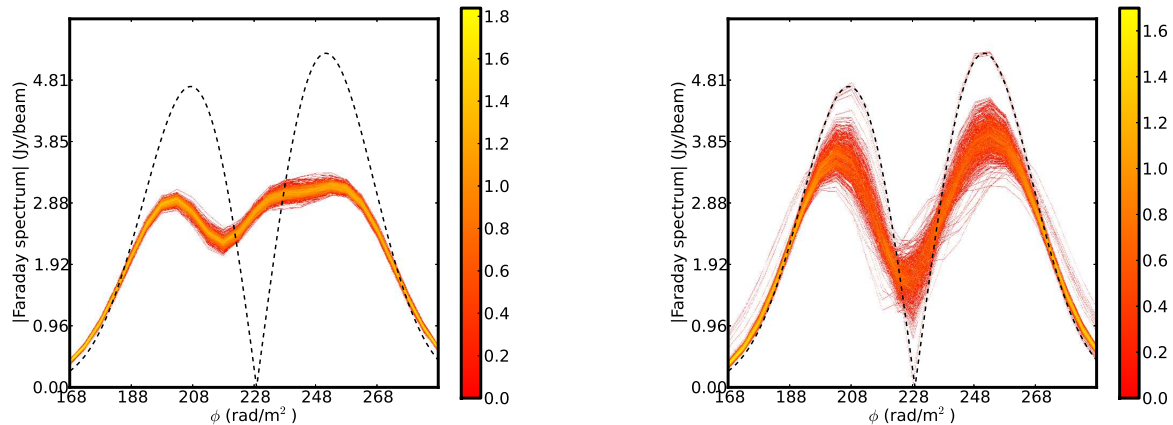
Fig. 4. Modulus of the Faraday spectra of two point sources having similar fluxes and the same phase and that are separated by $1\times$FWHM (40 rad/m^2). The dashed line shows the sky model after convolution with the idealized PSF. The (a) RMCLEAN and (b) ML reconstructions from 500 trials are plotted in red/yellow, with the color scale indicating the log of the number of reconstructions that pass through each location on the figure.

struct fluxes and locations of the individual sources even with over-resolution when there are multiple sources of emission.

We propose a simple algorithm to adjust the CLEAN model parameters using Maximum Likelihood estimation together with a prescription for reducing the degrees of freedom in the model to the minimum number that are supported by the data. We show that this algorithm improves upon the results obtained using RMCLEAN alone and provides highly accurate reconstructions independent of the choice of pixelization of Faraday depth space.

Both algorithms struggle in cases where two point sources are located within $1\times$FWHM of the PSF, but this merely reflects the fundamental limitation of the observations to resolve structures on these scales. However, although errors increase as source separation decreases, the ML algorithm still provides significant improvement over RMCLEAN.

CLEAN is ideally suited for the case where the image to be reconstructed is well-described by a set of independent point sources, but as we have shown, does not perform well even in these ideal circumstances when applied in the case of RM synthesis. The algorithm described herein provides significant improvements over RMCLEAN alone and is easy to append to existing RM synthesis imaging pipelines. Therefore, we recommend its use in upcoming polarization surveys that plan to include RM synthesis imaging.

The ML procedure still assumes that the image is described by a set of independent point sources, and will likely not be ideal in case large-scale, diffuse emission is present. A point-source based image reconstruction method is nevertheless an important tool to have on hand. It can be used as a method for separating point sources from a diffuse background, which may be better reconstructed using other methods that do not handle point sources well. It may also be the optimal deconvolution method for use with instruments that are not sensitive to scales much larger than the PSF, as is the case with e.g. LOFAR (Beck et al. 2012).

References

Andrecut, M., Stil, J. M., & Taylor, A. R. 2012, Astronomical Journal , 143, 33

Beck, R., Frick, P., Stepanov, R., & Sokoloff, D. 2012, Astronomy & Astrophysics , 543, A113

Bell, M. R. & Enßlin, T. A. 2012, Astronomy & Astrophysics , 540, A80

Bernardi, G., de Bruyn, A. G., Harker, G., et al. 2010, Astronomy & Astrophysics , 522, A67

Bernardi, G., Mitchell, D. A., Ord, S. M., et al. 2011, Monthly Notices of the Royal Astronomical Society , 413, 411

Brentjens, M. A. 2011, Astronomy & Astrophysics , 526, A9

Brentjens, M. A. & de Bruyn, A. G. 2005, Astronomy and Astrophysics, 441, 1217

Briggs, D. S. & Cornwell, T. J. 1992, in Astronomical Society of the Pacific Conference Series, Vol. 25, Astronomical Data Analysis Software and Systems I, ed. D. M. Worrall, C. Biemesderfer, & J. Barnes, 170

Clark, B. G. 1980, Astronomy & Astrophysics , 89, 377

Cotton, W. D. & Uson, J. M. 2008, Astronomy & Astrophysics , 490, 455

El-Behery, I. & MacPhie, R. 1980, Antennas and Propagation, IEEE Transactions on, 28, 234

Feain, I. J., Ekers, R. D., Murphy, T., et al. 2009, Astrophysical Journal , 707, 114

Frick, P., Sokoloff, D., Stepanov, R., & Beck, R. 2010, Monthly Notices of the Royal Astronomical Society, 401, L24

Harvey-Smith, L., Gaensler, B. M., Kothes, R., et al. 2010, Astrophysical Journal , 712, 1157

Heald, G., Braun, R., & Edmonds, R. 2009, Astronomy and Astrophysics, 503, 409

Heald, G. H. 2012, Astrophysical Journal Letters , 754, L35

Högbom, J. A. 1974, Astronomy & Astrophysics Supplement, 15, 417

Iacobelli, M., Haverkorn, M., & Katgert, P. 2012, ArXiv e-prints

Li, F., Brown, S., Cornwell, T. J., & de Hoog, F. 2011, Astronomy & Astrophysics , 531, A126

Mao, S. A., Gaensler, B. M., Haverkorn, M., et al. 2010, Astrophysical Journal , 714, 1170

Mao, S. A., McClure-Griffiths, N. M., Gaensler, B. M., et al. 2012a, Astrophysical Journal , 755, 21

Mao, S. A., McClure-Griffiths, N. M., Gaensler, B. M., et al. 2012b, Astrophysical Journal , 759, 25

O'Sullivan, S. P., Brown, S., Robishaw, T., et al. 2012, Monthly Notices of the Royal Astronomical Society , 421, 3300

Perley, R. A. 1999, in Astronomical Society of the Pacific Conference Series, Vol. 180, Synthesis Imaging in Radio Astronomy II, ed. G. B. Taylor, C. L. Carilli, & R. A. Perley, 275

Schwarz, G. 1978, The annals of statistics, 6, 461

Van Eck, C. L., Brown, J. C., Stil, J. M., et al. 2011, Astrophysical Journal , 728, 97

van Weeren, R. J., Röttgering, H. J. A., Intema, H. T., et al. 2012, Astronomy & Astrophysics , 546, A124

Yatawatta, S. 2010, arXiv:1008.1892

Appendix A: Iterative maximization approach

We assume that the optimal model point position is shifted by some small distance, $\delta\phi_i$, relative to the grid point location found using CLEAN, $\overline{\phi_i}$

$$\phi_i = \overline{\phi_i} + \delta\phi_i. \tag{A.1}$$

The maximum likelihood solution is obtained by minimizing H in Eq. 4 with respect to $\delta\phi_i$, which gives

$$\frac{\partial H(d|M)}{\partial\delta\phi_i} = \frac{2}{\sigma_n^2}\sum_j \mathfrak{R}\left[i\lambda_j^2 m_i\left\{m_i^* - \left(d_j - \widetilde{\widetilde{d_j^i}}\right)^* e^{2i\lambda_j^2\phi_i}\right\}\right] = 0 \tag{A.2}$$

where $\widetilde{\widetilde{d_j^i}}$ is the same as $\widetilde{d_j}$, Eq. 2, except that the sum is over $k \neq i$. The operator $\mathfrak{R}$ selects the real part of its argument. This expression cannot be solved analytically for $\delta\phi_i$, but we can use it to search for the solutions iteratively. To do so we Taylor expand the exponential term in Eq. A.2 to first order, thereby making H second order in $\delta\phi_i$. This gives

$$\frac{\partial H}{\partial\delta\phi_i} = \frac{2}{\sigma_n^2}\sum_j \mathfrak{R}\left[\lambda_j^2 m_i i\left\{m_i^* - \left(d_j - \widetilde{\widetilde{d_j^i}}\right)^* e^{2i\lambda_j^2\overline{\phi_i}} - 2i\lambda_j^2\left(\widetilde{\widetilde{d_j^i}} - d_j\right)^* e^{2i\lambda_j^2\overline{\phi_i}}\delta\phi_i\right\}\right] = 0, \tag{A.3}$$

which we can rewrite in the form

$$A = B\delta\phi_i, \tag{A.4}$$

where

$$A = \sum_j \mathfrak{R}\left[\left\{m_i^* - \left(d_j - \widetilde{\widetilde{d_j^i}}\right)^* e^{2i\lambda_j^2\overline{\phi_i}}\right\}\lambda_j^2 m_i i\right],$$

$$B = \sum_j \mathfrak{R}\left[-2\lambda_j^4 m_i \left(d_j - \widetilde{\widetilde{d_j^i}}\right)^* e^{2i\lambda_j^2\overline{\phi_i}}\right]. \tag{A.5}$$

We must also solve for an updated flux. We can again extremize Eq. 4, this time with respect to m_i. We find

$$\frac{\partial H}{\partial m_i} = \frac{1}{2\sigma_n^2}\sum_j[m_i^* - (d_j - \widetilde{\widetilde{d_j}})^* e^{2i\phi_i\lambda_j^2}] \tag{A.6}$$

and thus

$$m_i = \frac{1}{N_d}\sum_{j=1}^{N_d}[(d_j - \widetilde{\widetilde{d_j^i}})e^{-2i\phi_i\lambda_j^2}]. \tag{A.7}$$

In our example implementation, we solve for $\delta\phi$ and m iteratively until convergence is achieved. We also attempt to merge nearby model components to reduce the degrees of freedom in the model according to the prescription described in Sec. 2.

We tried other iterative schemes, e.g. solving for position and flux changes by inverting the Hessian matrix of H, but found that the approach given here is the most stable.