



# A D-term Modeling Code (DMC) for Simultaneous Calibration and Full-Stokes Imaging of Very Long Baseline Interferometric Data

Dominic W. Pesce<sup>1,2</sup> 

<sup>1</sup> Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA; [dpesce@cfa.harvard.edu](mailto:dpesce@cfa.harvard.edu)

<sup>2</sup> Black Hole Initiative at Harvard University, 20 Garden Street, Cambridge, MA 02138, USA

Received 2020 October 20; revised 2021 January 31; accepted 2021 February 5; published 2021 March 11

## Abstract

In this paper we present DMC, a model and associated tool for polarimetric imaging of very long baseline interferometry data sets that simultaneously reconstructs the full-Stokes emission structure along with the station-based gain and leakage calibration terms. DMC formulates the imaging problem in terms of posterior exploration, which is achieved using Hamiltonian Monte Carlo sampling. The resulting posterior distribution provides a natural quantification of uncertainty in both the image structure and the data calibration. We run DMC on both synthetic and real data sets, the results of which demonstrate its ability to accurately recover both the image structure and calibration quantities, as well as to assess their corresponding uncertainties. The framework underpinning DMC is flexible, and its specific implementation is under continued development.

*Unified Astronomy Thesaurus concepts:* [Radio interferometry \(1346\)](#); [Very long baseline interferometry \(1769\)](#); [Polarimetry \(1278\)](#); [Interferometry \(808\)](#)

## 1. Introduction

Interferometric observations in radio astronomy natively access the so-called “visibility domain,” with each visibility determined by the complex correlation between the electric fields incident at a pair of telescopes (Thompson et al. 2017, hereafter TMS). These visibilities provide information about the Fourier transform of the incident flux distribution via the van Cittert–Zernike theorem, and radio interferometric imaging—i.e., the process by which the visibility measurements are translated into a sky-plane image—presents an example of an ill-posed inverse problem. The combination of sparse Fourier-plane sampling and uncertain calibration, both of which are exacerbated for very long baseline interferometric (VLBI) observations, prevents a direct inversion of the visibilities to produce a unique image. Instead, images must be “reconstructed” with the aid of additional assumptions about the image structure (e.g., flux positivity, source sparsity) to overcome this nonuniqueness.

A variety of algorithms exist for reconstructing images in radio interferometry, and these algorithms can be broadly classified into the two categories established in Event Horizon Telescope Collaboration et al. (2019a). “Inverse modeling” schemes, exemplified by the CLEAN algorithm and its variants (Högbom 1974; Clark 1980; Schwab 1984), operate directly with the inverse Fourier transform of the visibility measurements and seek to iteratively deconvolve the effects of the finite sampling from the reconstructed image. “Forward modeling” schemes, such as the maximum entropy (e.g., Nityananda & Narayan 1982; Cornwell & Evans 1985) and regularized maximum likelihood (e.g., Chael et al. 2016; Akiyama et al. 2017) methods, instead parameterize the image structure (typically using a grid of pixels) and Fourier transform it to predict the values of the visibility measurements. The image parameters are then varied so as to optimize some objective function, typically consisting of a data comparison term (e.g., a  $\chi^2$  metric), along with one or more regularization terms.

From the perspective of computational speed, the CLEAN approach has historically been a clear favorite for VLBI imaging. The forward modeling schemes, though generally

more computationally taxing, benefit from the ability to enforce various nonlinear constraints (such as flux positivity) on the image and to fit directly to nonvisibility data products (such as closure quantities; e.g., Chael et al. 2018). Typical implementations of both classes of algorithm, however, share a mixed relationship with data calibration whereby image reconstruction steps are iterated with interleaving “self-calibration” steps that attempt to solve for station-based calibration terms (e.g., Readhead et al. 1980). Furthermore, both the inverse and forward modeling classes of image reconstruction algorithm classically lack a natural quantification of uncertainty in the image.

Recent developments have yielded a new class of image reconstruction algorithms, based on posterior exploration or parameterization techniques, that aim to overcome the aforementioned shortcomings (e.g., Cai et al. 2018a, 2018b; Arras et al. 2019; Broderick et al. 2020b). From the perspective of statistical integrity, an image reconstruction algorithm should solve simultaneously for the ensemble of both sky-plane emission structures and requisite calibration terms that are permissible, given the uncertainties in the data and any sources of prior knowledge about the parameters. In this paper we present such an algorithm in terms of a model for simultaneous calibration and full-Stokes imaging of VLBI data, along with an implementation of this model within a generic posterior exploration framework. An implementation of an analogous model within THEMIS (Broderick et al. 2020a) is presented in a separate paper, A. E. Broderick et al. (2021, in preparation).

This paper is organized as follows. In Section 2 we provide a detailed description of the model, specify our likelihood construction, and describe its software implementation. In Section 3 we demonstrate the results of fitting this model to both synthetic and real VLBI data sets. We summarize and conclude in Section 4.

## 2. Model Specifications

For the compact sources and small fields of view typically considered in VLBI, the observed visibilities are related to the

Fourier transform of the sky brightness distribution by the van Cittert–Zernike theorem (TMS),

$$\tilde{I}(u, v) = \iint I(x, y) e^{-2\pi i(ux+vy)} dx dy, \quad (1)$$

where we use a tilde ( $\sim$ ) to denote a transformed quantity. In this section and throughout the paper, unless otherwise specified, we consider observations made with only a single frequency channel.

We have developed a new publicly available D-term Modeling Code (DMC)<sup>3</sup> that implements the polarimetric image model detailed in this section. DMC is implemented in Python, and it fits the model using a Bayesian formalism in which the posterior distribution  $\mathcal{P}(\Theta)$  of the parameter vector  $\Theta$  is related to the likelihood  $\mathcal{L}(\Theta)$  and prior  $\pi(\Theta)$  via Bayes's theorem,

$$\mathcal{P}(\Theta) \propto \mathcal{L}(\Theta) \pi(\Theta). \quad (2)$$

Model parameters and their associated priors are aggregated in Table 1. DMC uses the eht-imaging library (Chael et al. 2016, 2018) for internal organization and manipulation of VLBI data, and it uses the PyMC3 library (Salvatier et al. 2016) for sampling. Because it uses a Markov Chain Monte Carlo (MCMC) sampler, the output of running DMC on a VLBI data set is an ensemble of images that are drawn from the posterior distribution of the model. From this ensemble it is possible to compute various useful statistics (e.g., means, variances), which we demonstrate in Section 3.

### 2.1. Image Model

We model the image as a Cartesian grid of  $N_{\text{pix}}$  pixels, with the grid axes aligned with the equatorial coordinate axes. Each pixel has a location  $(x_j, y_j)$  and a Stokes I intensity  $I_j$ . These intensities are constrained to sum to a total flux density  $F$ ,

$$F = \sum_{j=1}^{N_{\text{pix}}} I_j, \quad (3)$$

with  $F$  specifiable but by default sampled from a normal prior truncated at zero to ensure positivity. The constrained sum in Equation (3) is imposed via a Dirichlet prior on the pixel intensity values,

$$\frac{\mathbf{I}}{F} \sim \text{Dir}(N_{\text{pix}}, \mathbf{a}), \quad (4)$$

where  $\mathbf{I} = (I_1, I_2, \dots, I_{N_{\text{pix}}})$  is the vector of pixel intensities. The concentration parameter vector  $\mathbf{a}$  is specifiable but defaults to  $\mathbf{a} = \mathbf{1} \equiv (1, 1, \dots, 1)$ , which corresponds to a flat Dirichlet prior with the flux-normalized pixel intensities sampled uniformly on the  $(N_{\text{pix}} - 1)$ -dimensional simplex. Setting smaller values for the concentration parameters encourages sparsity in the image, while setting larger values encourages diffuse flux.

In each pixel, we can relate the Stokes I intensity  $I_j$  to the other Stokes parameters by the inequality

$$I_j^2 \geq Q_j^2 + U_j^2 + V_j^2 \equiv p_j^2 I_j^2, \quad (5)$$

**Table 1**  
Model Parameters and Priors

Parameter	Description	Default Prior
$F$	Image-integrated Stokes I flux density	$\mathcal{N}_0(\check{F}, [0.1\check{F}]^2)$
$I_j / F$	Fraction of Stokes I flux density contained in pixel $j$	$\text{Dir}(N_{\text{pix}}, \mathbf{1})$
$p_j$	Polarization fraction in pixel $j$	$\mathcal{U}(0, 1)$
$\alpha_j$	Azimuthal angle Poincaré coordinate in pixel $j$	$\mathcal{U}_{\text{per}}(-\pi, \pi)$
$\beta_j$	Polar angle Poincaré coordinate in pixel $j$	$\cos(\beta_j) \sim \mathcal{U}(-1, 1)$
$\Sigma$	FWHM of Gaussian blurring kernel, in $\mu\text{as}$	$\delta(\Sigma - \check{\Sigma})$
$x_0$	Overall image centroid shift along the R.A. axis, in $\mu\text{as}$	$\delta(x_0 - \check{x}_0)$
$y_0$	Overall image centroid shift along the decl. axis, in $\mu\text{as}$	$\delta(y_0 - \check{y}_0)$
$g_{R,a}$	Right-hand gain amplitude for station $a$	$\mathcal{N}_0(1, \check{\sigma}_{R,a}^2)$
$\theta_{R,a}$	Right-hand gain phase for station $a$	$\mathcal{U}_{\text{per}}(-\pi, \pi)$
$g_{L,a}$	Left-hand gain amplitude for station $a$	$\mathcal{N}_0(1, \check{\sigma}_{L,a}^2)$
$\theta_{L,a}$	Left-hand gain phase for station $a$	$\mathcal{U}_{\text{per}}(-\pi, \pi)$
$d_{R,a}$	Right-hand leakage amplitude for station $a$	$\mathcal{U}(0, 1)$
$\delta_{R,a}$	Right-hand leakage phase for station $a$	$\mathcal{U}_{\text{per}}(-\pi, \pi)$
$d_{L,a}$	Left-hand leakage amplitude for station $a$	$\mathcal{U}(0, 1)$
$\delta_{L,a}$	Left-hand leakage phase for station $a$	$\mathcal{U}_{\text{per}}(-\pi, \pi)$
$f$	Fractional systematic uncertainty	$\mathcal{U}(0, 1)$
$N_x$	Number of image pixels along the R.A. axis	...
$N_y$	Number of image pixels along the decl. axis	...
$\text{FOV}_x$	Field of view along the R.A. axis, in $\mu\text{as}$	...
$\text{FOV}_y$	Field of view along the decl. axis, in $\mu\text{as}$	...

**Note.** A list of the model parameters and their corresponding prior distributions. The top portion of the table lists the parameters associated with the image, the middle portion lists parameters associated with the calibration, and the bottom portion lists the hyperparameters. We use a breve ( $\check{\cdot}$ ) to denote user-specified quantities. We use a number of different prior classes:  $\delta(x - a)$  denotes a Dirac delta prior over  $x$  such that it takes on the fixed value  $a$ ,  $\mathcal{U}(a, b)$  denotes a uniform prior on the range  $[a, b]$ ,  $\mathcal{U}_{\text{per}}(a, b)$  denotes a periodic (or “wrapped”) uniform prior on the range  $[a, b]$ ,  $\mathcal{N}(\mu, \sigma^2)$  denotes a normal (Gaussian) distribution with mean  $\mu$  and variance  $\sigma^2$ ,  $\mathcal{N}_0(\mu, \sigma^2)$  denotes a normal distribution (with mean  $\mu$  and variance  $\sigma^2$ ) that has a lower-bound truncation at zero,  $\mathcal{N}_c(\mu, \sigma^2)$  denotes a circularly symmetric complex normal distribution with (complex) mean  $\mu$  and variance  $\sigma^2$  along both the real and imaginary directions, and  $\text{Dir}(N, \mathbf{a})$  denotes a Dirichlet distribution in  $N$  dimensions with concentration parameter vector  $\mathbf{a}$ .

where we have introduced the polarization fraction  $p_j \leq 1$ . This spherical relationship lends itself naturally to a Poincaré parameterization in terms of angular variables,

$$\begin{pmatrix} I_j \\ Q_j \\ U_j \\ V_j \end{pmatrix} = I_j \begin{pmatrix} 1 \\ p_j \cos(\alpha_j) \sin(\beta_j) \\ p_j \sin(\alpha_j) \sin(\beta_j) \\ p_j \cos(\beta_j) \end{pmatrix}, \quad (6)$$

where  $-\pi \leq \alpha_j \leq \pi$  is an azimuthal angle and  $0 \leq \beta_j \leq \pi$  is a polar angle. The angle  $\alpha_j$  determines the orientation of the polarization ellipse, and it is related to the usual electric vector

<sup>3</sup> <https://github.com/dpesce/eht-dmc>

position angle (EVPA)  $\chi_j$  by

$$\chi_j = \frac{1}{2} \tan^{-1} \left( \frac{U_j}{Q_j} \right) = \frac{\alpha_j}{2}. \quad (7)$$

The angle  $\beta_j$  determines the degree of circular polarization, with purely linear polarization having  $\beta_j = \pi/2$  and purely circular polarization having  $\beta_j = 0$  or  $\beta_j = \pi$ . We sample  $p_j$  from a unit uniform distribution, and we sample the angular variables uniformly on the unit sphere. Our polarized image model thus consists of the four quantities ( $I_j, p_j, \alpha_j, \beta_j$ ) for every pixel, which, together with the total flux  $F$ , amount to  $4N_{\text{pix}}$  model parameters.

From the parameters ( $I_j, p_j, \alpha_j, \beta_j$ ), we determine the Stokes parameters in each pixel using Equation (6). We then compute the Fourier transforms of these Stokes parameters via

$$\begin{pmatrix} \tilde{I}_k \\ \tilde{Q}_k \\ \tilde{U}_k \\ \tilde{V}_k \end{pmatrix} = S_k \sum_{j=1}^{N_{\text{pix}}} A_{jk} \begin{pmatrix} I_j \\ Q_j \\ U_j \\ V_j \end{pmatrix}, \quad (8)$$

where

$$A_{jk} = \exp(-2\pi i [u_k(x_j - x_0) + v_k(y_j - y_0)]) \quad (9)$$

are elements of the discrete Fourier transform matrix, ( $u_k, v_k$ ) are the Fourier-plane coordinates for visibility measurement  $k$  in units of the observing wavelength, ( $x_0, y_0$ ) are the image-plane coordinates of the image origin (or “phase center”), and

$$S_k = \exp \left[ -\frac{\pi^2 \Sigma^2 (u_k^2 + v_k^2)}{4 \ln(2)} \right] \quad (10)$$

is a circularly symmetric Gaussian smoothing kernel with FWHM  $\Sigma$  in the image plane that serves to maintain image continuity.

## 2.2. Corruption Model

For an array observing with circularly polarized feeds, the measured quantities are parallel- and cross-hand correlation products; we denote the parallel-hand visibilities as  $RR_{12} \equiv \langle E_{R,1} E_{R,2}^* \rangle$  and  $LL_{12} \equiv \langle E_{L,1} E_{L,2}^* \rangle$ , and we denote the cross-hand visibilities as  $RL_{12} \equiv \langle E_{R,1} E_{L,2}^* \rangle$  and  $LR_{12} \equiv \langle E_{L,1} E_{R,2}^* \rangle$ . The measured visibilities on a baseline  $ab$  are related to the Stokes visibilities on that same baseline by

$$\begin{pmatrix} RR_{ab} \\ LL_{ab} \\ RL_{ab} \\ LR_{ab} \end{pmatrix} = \begin{pmatrix} \tilde{I}_{ab} + \tilde{V}_{ab} \\ \tilde{I}_{ab} - \tilde{V}_{ab} \\ \tilde{Q}_{ab} + i\tilde{U}_{ab} \\ \tilde{Q}_{ab} - i\tilde{U}_{ab} \end{pmatrix}. \quad (11)$$

In real interferometric observations, the measured visibilities are corrupted by a combination of a priori unknown signal propagation effects. Following the radio interferometer measurement equation (RIME) formalism developed by Hamaker et al. (1996)—and in particular the  $2 \times 2$  matrix extension described in Hamaker (2000)—we relate the incident and measured visibilities using Jones matrix transformations of the

“coherency matrix,”<sup>4</sup>

$$\mathbf{V}_{ab} \equiv \begin{pmatrix} RR_{ab} & RL_{ab} \\ LR_{ab} & LL_{ab} \end{pmatrix}. \quad (12)$$

Within the RIME formalism, the (complex) Jones matrix  $\mathbf{J}_a$  captures all linear transformations undergone by the incident astrophysical signal at a station  $a$ , such that

$$\hat{\mathbf{V}}_{ab} = \mathbf{J}_a \mathbf{V}_{ab} \mathbf{J}_b^\dagger, \quad (13)$$

where a dagger ( $\dagger$ ) denotes a conjugate transpose and a hat ( $\hat{\cdot}$ ) denotes an observed quantity.

DMC incorporates a minimal but standard (see, e.g., TMS) threefold decomposition of  $\mathbf{J}_a$ ,

$$\mathbf{J}_a = \mathbf{G}_a \mathbf{D}_a \mathbf{F}_a, \quad (14)$$

where

$$\mathbf{G}_a = \begin{pmatrix} G_{R,a} & 0 \\ 0 & G_{L,a} \end{pmatrix} \quad (15)$$

contains the station gain terms,

$$\mathbf{D}_a = \begin{pmatrix} 1 & D_{R,a} \\ D_{L,a} & 1 \end{pmatrix} \quad (16)$$

contains the polarimetric leakage terms, and

$$\mathbf{F}_a = \begin{pmatrix} e^{-i\phi_a} & 0 \\ 0 & e^{i\phi_a} \end{pmatrix} \quad (17)$$

applies the feed rotation angle,  $\phi_a$ . Note that this decomposition incorporates only station-based corruptions, and it does not account for direction-dependent effects (e.g., Smirnov 2011) or for other baseline-based corruptions. The feed rotation angle  $\phi_a$  depends on the station mount properties and on the source parallactic and elevation angles as a function of time, but for most radio interferometers it is well known a priori; we thus assume the  $\phi_a$  to be given and therefore do not incorporate them as model parameters. The station gain and leakage terms, however, are typically imperfectly calibrated, and so we retain both as model parameters.

We parameterize the complex station gains using amplitude and phase,

$$G_{R,a} = g_{R,a} e^{i\theta_{R,a}}, \quad (18a)$$

$$G_{L,a} = g_{L,a} e^{i\theta_{L,a}}. \quad (18b)$$

Our priors on the gain amplitudes are normal with a lower-bound truncation at zero, and we impose periodic uniform<sup>5</sup> priors on the gain phases with the range  $(-\pi, \pi)$ . For each observation we select a single “reference station” for which both the right and left gain phases are fixed to be zero (i.e.,  $\theta_R = \theta_L$ ).<sup>6</sup> We permit all gains other than those of the reference station to be independent across stations and across time

<sup>4</sup> This particular (sky-intrinsic) coherency matrix is also referred to in the literature as the “brightness matrix,” and its measured counterpart has been referred to as the “visibility matrix” (Smirnov 2011).

<sup>5</sup> A periodic uniform distribution is one that is uniformly distributed on the unit circle.

<sup>6</sup> Note that if the reference station does not actually have a zero-valued phase difference between its right- and left-hand gains, this treatment will result in an overall image EVPA rotation that must be absolutely calibrated (Brown et al. 1989).

stamps. We note that in real-world arrays the station gain amplitudes are not expected to fluctuate wildly from one time stamp to the next. In this sense, DMC aims to provide a conservative treatment of the gain behavior; i.e., DMC permits—though it does not impose—large gain amplitude fluctuations between time stamps, such as may occur when a telescope is repointed.

We use an analogous parameterization for the complex leakage terms,

$$D_{R,a} = d_{R,a} e^{i\delta_{R,a}}, \quad (19a)$$

$$D_{L,a} = d_{L,a} e^{i\delta_{L,a}}. \quad (19b)$$

We impose unit uniform priors on the leakage amplitudes and periodic uniform priors on the leakage phases with the range  $(-\pi, \pi)$ . We assume that the leakage terms are constant in time (see, e.g., Conway & Kronberg 1969; Roberts et al. 1994) and so assign only a single  $\mathbf{D}_a$  for every station.

We note that the default priors described in this section for the gain and leakage terms may be overridden when running DMC to incorporate any a priori knowledge of the station properties.

### 2.3. Likelihood Construction

The thermal noise in any single visibility measurement depends on various factors, including the collecting area of the telescopes constituting the baseline and the averaging time and bandwidth of the observation, but for the vast majority of sources of interest<sup>7</sup> this thermal noise is normally distributed and statistically independent across different baselines (TMS). In the absence of any other corruptions, a measured visibility is drawn from a circularly symmetric complex normal distribution, e.g.,

$$\hat{R}R_k \sim \mathcal{N}_c(\mathcal{R}R_k, \sigma_{\text{th},RR,k}^2), \quad (20)$$

where  $\mathcal{R}R_k$  is the “true” visibility value on baseline  $k$  and  $\sigma_{\text{th},RR,k}^2$  is the thermal variance in the corresponding visibility measurement. Note that the presence of gain corruptions does introduce covariance between visibility measurements, but by explicitly modeling these gains, we account for this covariance and ensure that the remaining differences between modeled and observed visibilities will be independently distributed (Blackburn et al. 2020).

Though we explicitly incorporate a number of known corrupting effects into the model (see Section 2.2), we also permit an additional multiplicative systematic noise component,

$$\begin{pmatrix} \sigma_{RR,k}^2 \\ \sigma_{LL,k}^2 \\ \sigma_{RL,k}^2 \\ \sigma_{LR,k}^2 \end{pmatrix} = \begin{pmatrix} \sigma_{\text{th},RR,k}^2 + f^2 I_k^2 \\ \sigma_{\text{th},LL,k}^2 + f^2 I_k^2 \\ \sigma_{\text{th},RL,k}^2 + f^2 I_k^2 \\ \sigma_{\text{th},LR,k}^2 + f^2 I_k^2 \end{pmatrix}, \quad (21)$$

where  $f$  is sampled from a unit uniform distribution. This systematic component aims to account for any uncalibrated nonclosing errors that cannot be described by thermal noise or

leakage. We then construct the likelihood of a particular set of model visibilities given the visibility measurements using

$$\mathcal{L}_{RR} = \prod_k \frac{1}{2\pi\sigma_{RR,k}^2} \exp \left[ \frac{(\hat{R}R_k - RR_k)(RR_k - \hat{R}R_k)^*}{2\sigma_{RR,k}^2} \right], \quad (22a)$$

$$\mathcal{L}_{LL} = \prod_k \frac{1}{2\pi\sigma_{LL,k}^2} \exp \left[ \frac{(\hat{L}L_k - LL_k)(LL_k - \hat{L}L_k)^*}{2\sigma_{LL,k}^2} \right], \quad (22b)$$

$$\mathcal{L}_{RL} = \prod_k \frac{1}{2\pi\sigma_{RL,k}^2} \exp \left[ \frac{(\hat{R}L_k - RL_k)(RL_k - \hat{R}L_k)^*}{2\sigma_{RL,k}^2} \right], \quad (22c)$$

$$\mathcal{L}_{LR} = \prod_k \frac{1}{2\pi\sigma_{LR,k}^2} \exp \left[ \frac{(\hat{L}R_k - LR_k)(LR_k - \hat{L}R_k)^*}{2\sigma_{LR,k}^2} \right], \quad (22d)$$

where the products are taken over all visibility measurements. The final likelihood expression is then simply the product of the individual visibility likelihoods,

$$\mathcal{L} = \mathcal{L}_{RR} \mathcal{L}_{LL} \mathcal{L}_{RL} \mathcal{L}_{LR}. \quad (23)$$

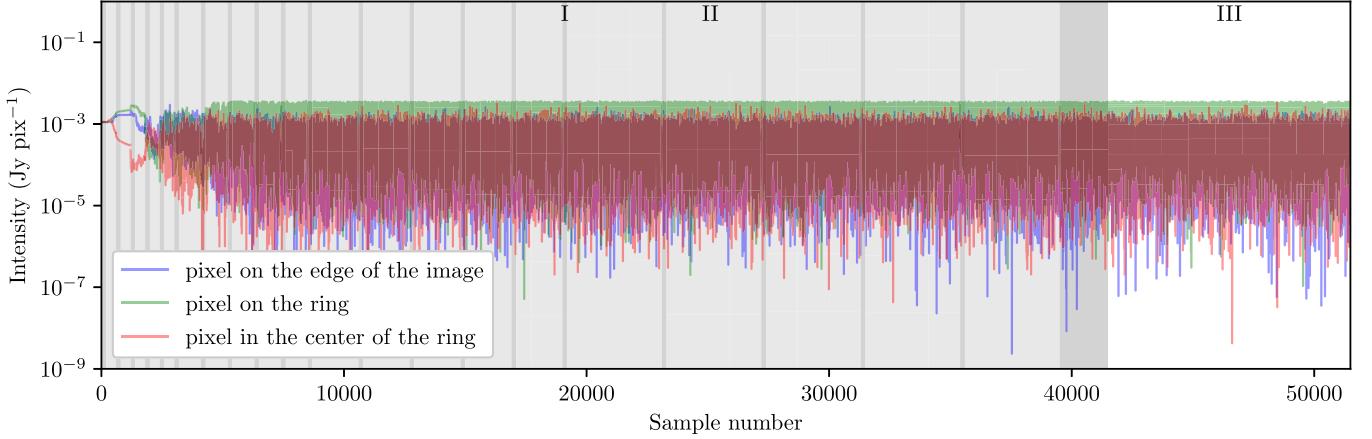
### 2.4. Sampler and Tuning

DMC uses the Hamiltonian Monte Carlo (HMC; Duane et al. 1987) No U-Turn Sampler (NUTS; Hoffman & Gelman 2014) implemented within the PyMC3 Python package (Salvatier et al. 2016) to explore the posterior space. HMC is an MCMC method whose output product is an ensemble of samples from the posterior distribution. Detailed descriptions of the HMC method can be found in, e.g., Neal (2011) and Betancourt (2017). PyMC3 is a probabilistic programming tool that leverages Theano (Bergstra et al. 2010; Bastien et al. 2012) to automatically differentiate the posterior density when computing model gradients.

As an HMC sampler, PyMC3 exploits model gradient information to efficiently explore the high-dimensional posterior space presented by the polarized image model. The number of tunable hyperparameters is minimized through the use of NUTS, but there remain two key hyperparameters that need to be adaptively tuned during sampling itself: a “step size” hyperparameter that sets the discretization interval for trajectory integrations, and a “mass matrix” hyperparameter (actually a collection of hyperparameters, the elements of the matrix) that determines the Gaussian distribution from which the momentum parameters are sampled. PyMC3 natively adapts the step size hyperparameter during sampling (see Hoffman & Gelman 2014), but its default functionality only adapts the diagonal elements of the mass matrix. Strong correlations in the posterior distribution can therefore lead to decreased sampling efficiency.

To mitigate this potential deficiency, we divide the sampling period into multiple tuning windows during which both the step size and (dense) mass matrix are adaptively determined. We have designed these windows to mimic the “warm-up epochs” used in the Stan package (Carpenter et al. 2017). An initial “fast” window is used to tune the step size parameter, after which a series of increasingly heavily sampled “slow” windows are used to estimate the mass matrix using the parameter covariances measured from the set of posterior samples in the previous window. Each slow window is preceded by a brief

<sup>7</sup> The “self-noise” of Kulkarni (1989) introduces statistical dependence in the noise measured across multiple baselines, but this contribution only becomes relevant for extremely bright sources.



**Figure 1.** Example set of parameter traces from a DMC fit to the EHT-like synthetic data set described in Section 3.1, showing the Stokes I intensity of a pixel located toward the edge of the image (in blue), a pixel situated on the ring (in green), and a pixel located in the center of the ring (in red). The shaded regions highlight the different tuning windows described in Section 2.4; the dark-gray shading indicates a “fast” tuning window (one example is labeled as “I”), the light-gray shading indicates a “slow” tuning window (one example is labeled as “II”), and the unshaded region indicates the main sampling phase (labeled as “III”).

fast window to permit the step size to adapt to the new mass matrix. A final fast window follows the last slow window, after which the main sampling phase proceeds using the tuned values for both hyperparameters. This tuning procedure is illustrated in Figure 1.

In practice, we find that tuning of both the step size and the mass matrix is essential for polarized imaging using DMC. If, e.g., the mass matrix tuning is restricted to only the diagonal elements, parameter autocorrelation times are liable to increase by several orders of magnitude and the sampler will effectively stall.

### 3. Demonstrations

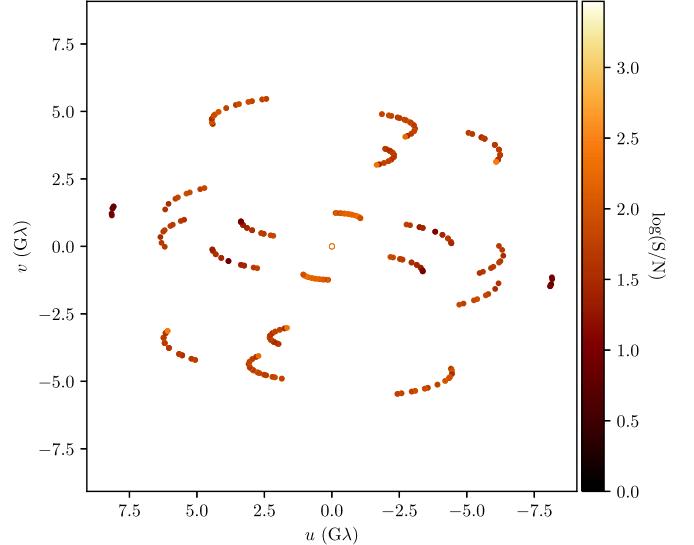
In this section we demonstrate the imaging capabilities of DMC on both synthetic and real data.

#### 3.1. Synthetic Data Construction

We first run DMC on a synthetic data set constructed to have properties similar to the 2017 observations of the M87 black hole with the EHT (Event Horizon Telescope Collaboration et al. 2019a, 2019b, 2019c, 2019d, 2019e, 2019f). The baseline coverage and signal-to-noise ratio distribution for this data set are shown in Figure 2, and the input source model Stokes images are shown in the left panels of Figure 3. The visibility data are generated in a circular polarization basis, corresponding to the state of the real EHT data after fringe fitting has been performed (Event Horizon Telescope Collaboration et al. 2019d).

The input source structure is a circular crescent with a diameter of  $40 \mu\text{as}$  and a Gaussian FWHM of  $10 \mu\text{as}$ , oriented such that the brightest region of the crescent is located toward the north. The image-integrated flux density is  $1.0 \text{ Jy}$ , and it is polarized at the 10% level in linear polarization and at the 2% level in circular polarization. We construct the linear polarization structure to have an approximately threefold azimuthal symmetry (see the left panel of Figure 6), corresponding to nonzero  $\beta_{-1}$  and  $\beta_1$  modes in the decomposition developed by Palumbo et al. (2020, hereafter PWP).

In addition to the thermal noise, we add gain and leakage corruptions to each of the seven stations in the synthetic data set. The gain amplitudes are permitted to vary at the 10% level, while the gain phases are unconstrained (i.e., they are sampled

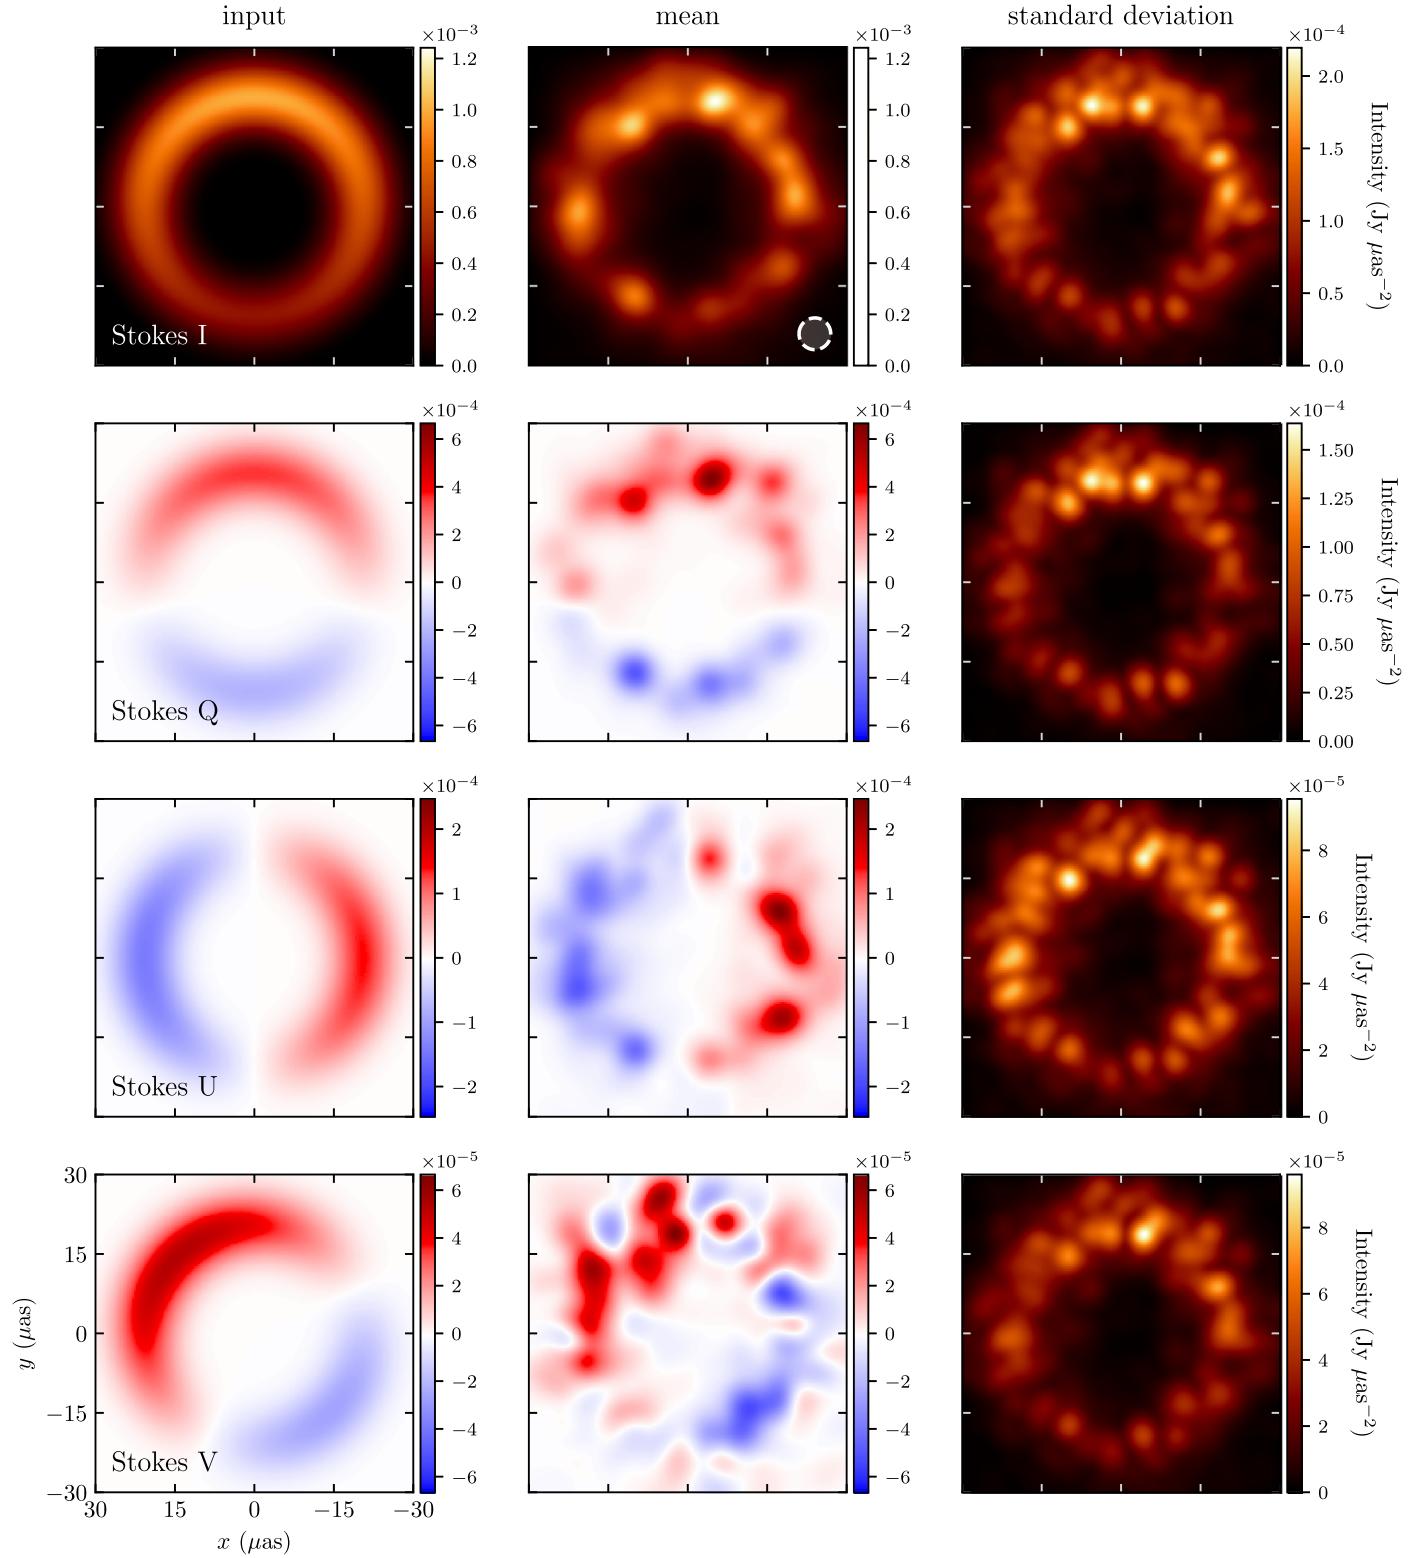


**Figure 2.**  $(u, v)$ -coverage for the EHT-like synthetic data set, with points colored by the base-10 logarithm of their Stokes I signal-to-noise ratio.

uniformly on the unit circle); each station has independent gains, but these gains are drawn from the same distribution and thus share the same magnitude of fluctuations. The same gain corruptions are used for both right- and left-hand polarization. Each station has an associated complex station leakage in both right- and left-hand polarization that is at the level of 0%–10%. The station gain amplitudes and phases are independently generated for each station at each observing time stamp, while the complex right- and left-hand leakage terms are held constant for each station across the synthetic observation.

#### 3.2. Imaging Synthetic Data

We use DMC to image the synthetic data set described in the previous section, setting  $\text{FOV}_x = \text{FOV}_y = 60 \mu\text{as}$ ,  $N_x = N_y = 30$ ,  $\check{x}_0 = \check{y}_0 = 0 \mu\text{as}$ , and  $\Sigma = 6 \mu\text{as}$ . The initial total flux estimate  $F$  is set equal to the visibility amplitude on the shortest baseline, and the gain amplitude prior standard deviations are fixed to  $\check{\sigma}_R = \check{\sigma}_L = 0.1$  for all stations. We select the ALMA station to be our reference antenna, both because it is present at every time stamp and because it provides an absolute EVPA standard (Event



**Figure 3.** DMC image reconstructions of the polarized EHT-like synthetic data set described in Section 3.1. Each row shows the results for a different Stokes parameter. The same field of view is used for all plots and is explicitly labeled in the bottom left panel. The left column shows the ground-truth input images, the middle column shows the mean posterior images for each of the Stokes parameters, and the right column shows the standard deviations of the image posteriors. Because these data sets included arbitrary gain phase corruptions, absolute position information is not uniquely recovered, and so the reconstructed images have been shifted to the location that maximizes the normalized cross-correlation between the ground-truth Stokes I image and the mean of the Stokes I image posterior. The 6  $\mu$ as FWHM of the Gaussian smoothing kernel (i.e.,  $\Sigma$  from Equation (10)) is shown in the lower right corner of the top middle panel. Note that this kernel does not represent a typical “restoring beam” that gets applied after imaging has been completed; rather, it is a convolving function that gets self-consistently applied during the imaging process itself (see Section 2.1).

Horizon Telescope Collaboration et al. 2019d; Goddi et al. 2019).<sup>8</sup>

DMC is able to find good fits to this data set, with reduced- $\chi^2$  values near unity and an additional systematic noise parameter consistent with zero (i.e., the model is able to describe the data to within thermal errors without requiring an additional source of systematic uncertainty). The image-integrated linear polarization fraction is recovered as  $9.96\% \pm 0.03\%$  (compared to an input value of 10%), while the image-integrated circular polarization fraction is found to be  $1.3\% \pm 1.0\%$  (compared to an input value of 2%).

In Figure 3, we show maps of both the mean and the standard deviation of the image ensemble output from DMC—where each image in the ensemble corresponds to a single sample from the posterior distribution—and we compare these against the input images of all four Stokes parameters. Because DMC produces an estimate of the posterior distribution for every image pixel, arbitrarily complicated single- or multipoint statistics may be computed that self-consistently capture the uncertainties and correlations between pixels. For instance, image-domain feature extraction techniques—such as those employed in Event Horizon Telescope Collaboration et al. (2019a, 2019f) to measure the properties (e.g., diameter, width, orientation) of the ring-like structure seen in M87—can be applied to the ensemble of images in the DMC posterior to yield corresponding distributions for any derived values.

We demonstrate an example of the latter capability in Figure 4, which shows the derived posteriors for the PWP  $\beta_1$  and  $\beta_{-1}$  parameters. In this paper, we adopt

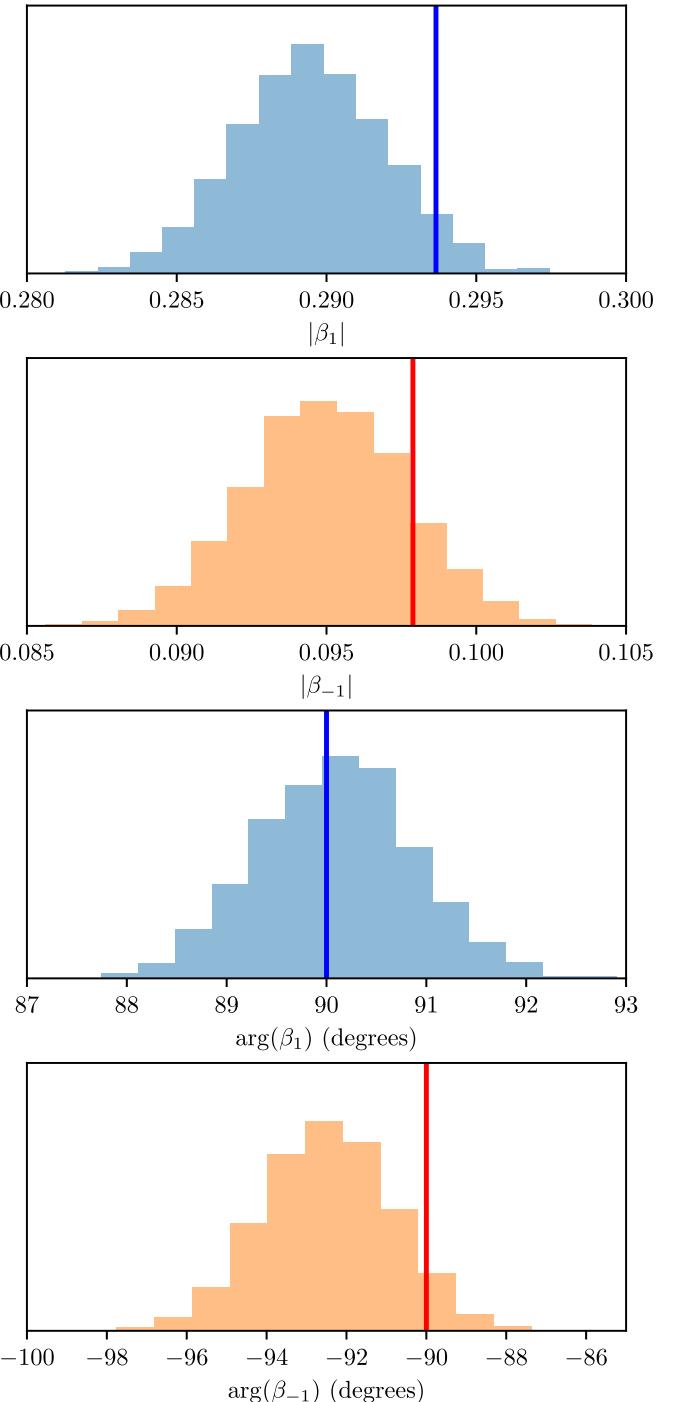
$$\beta_m = \frac{1}{F} \iint (Q + iU) e^{-im\varphi} r dr d\varphi \quad (24)$$

as our working definition for  $\beta_m$ , with  $F$  the total image flux density and the integral taken over the entire range of image-domain spatial coordinates  $(r, \varphi)$ . The PWP  $\beta_m$  decomposition permits a quantitative characterization of the polarization structure on a ring, and for the input synthetic data set only  $\beta_1$  and  $\beta_{-1}$  modes are nonzero. As shown in Figure 4, we find that the DMC fits return posteriors on these two modes that are consistent with the input values.

Furthermore, because DMC simultaneously models both the image structure and the station-based corruptions, we are also able to investigate the posterior behavior of the latter. Of most interest for polarimetric reconstructions are the leakage terms, which we list in Table 2. Figure 5 shows the leakage posterior distributions for each station in the array. We see that DMC is able to recover the correct leakages for all stations and to determine which stations are most well constrained.

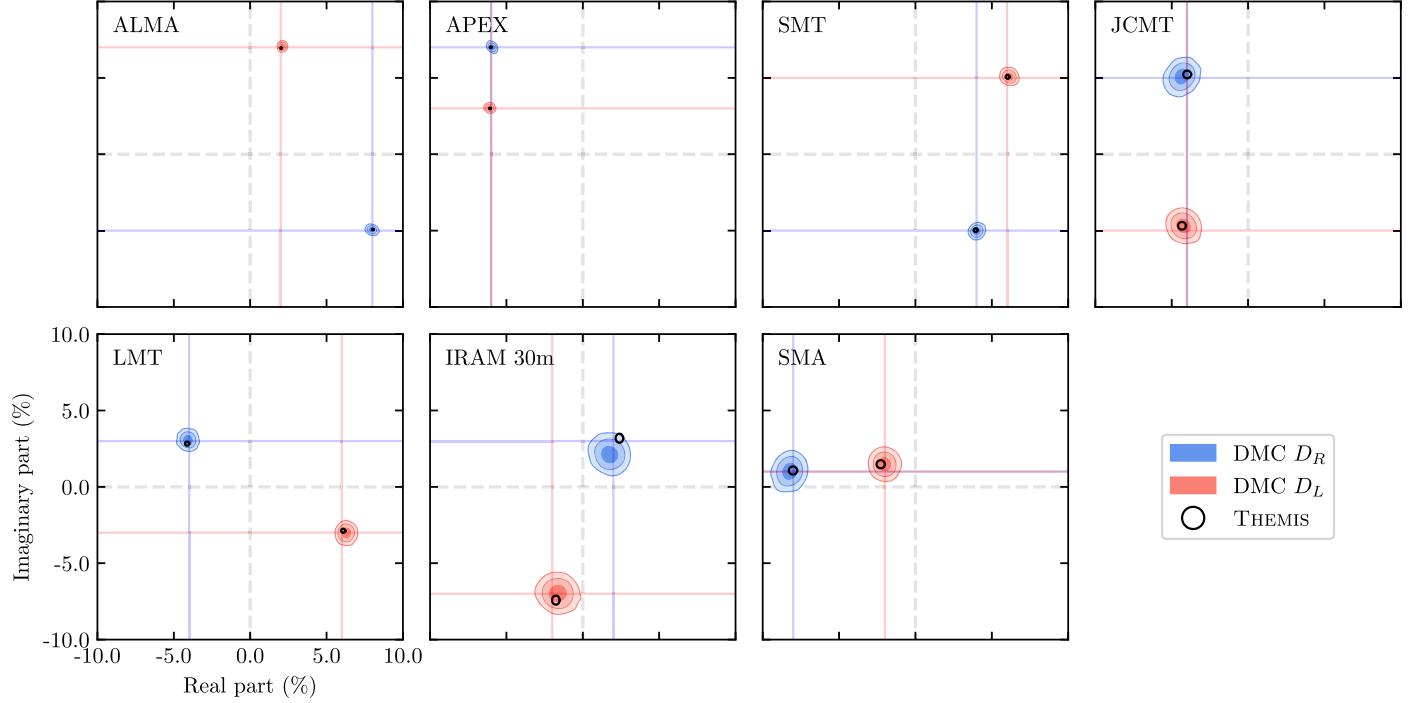
We also compare the DMC results with those of THEMIS (Broderick et al. 2020a), an independent code capable of producing simultaneous image and leakage posteriors for VLBI data sets (A. E. Broderick et al. 2021, in preparation). In Figures 5 and 6, the DMC posteriors are compared against those from THEMIS for the station leakages and image structure, respectively. Although the two softwares employ

<sup>8</sup> We note that for this synthetic data set all of the stations are effectively calibrated in absolute EVPA. For real EHT data only the ALMA station has this property, which it inherits from application of the PolConvert procedure (Martí-Vidal et al. 2016) that gets applied after correlation to convert the mixed-basis polarization products on ALMA baselines to the circular polarization products used by the rest of the array (Event Horizon Telescope Collaboration et al. 2019d; Goddi et al. 2019).



**Figure 4.** Posterior distributions for the  $\beta_1$  (in blue) and  $\beta_{-1}$  (in orange) values corresponding to the DMC reconstruction of the EHT-like synthetic data set. The top two panels show the amplitudes of both quantities, while the bottom two panels show their phases. In all panels, the value derived from the input image is shown as a vertical line.

distinct model specifications and exploration schemes, we find that the posterior reconstructions from both codes nevertheless show good agreement. The THEMIS fit typically shows only small ( $\lesssim 1\sigma$ ) shifts in posterior mean relative to the DMC fit, and the systematically wider posterior widths from the DMC fit stem from the more permissive station gain priors (DMC permits the right- and left-hand station gains to vary independently, while THEMIS enforces equality between both hands).



**Figure 5.** Leakage posteriors for the DMC fit to the EHT-like synthetic data set described in Section 3.1; each panel shows the result for an individual station, and ground-truth input values are marked using solid vertical and horizontal lines. The same axis ranges are used for all panels and are explicitly labeled in the bottom left panel. Blue contours show the right-hand leakage posteriors, and red contours show the left-hand leakage posteriors. In all panels, the plotted contours enclose 50%, 90%, and 99% of the posterior probability. For comparison, we have overplotted the leakage posteriors from THEMIS as black ellipses, with the sizes of the ellipses along each axis corresponding to the posterior standard deviation.

**Table 2**  
Station Leaks for Synthetic Data

Station	Input $D_R$	Posterior $D_R \pm 1\sigma$	Input $D_L$	Posterior $D_L \pm 1\sigma$
ALMA	$8.0 - 5.0i$	$(7.97 - 4.94i) \pm (0.14 + 0.11i)$	$2.0 + 7.0i$	$(2.11 + 7.06i) \pm (0.11 + 0.11i)$
APEX	$-6.0 + 7.0i$	$(-5.93 + 7.00i) \pm (0.11 + 0.12i)$	$-6.0 + 3.0i$	$(-6.09 + 3.04i) \pm (0.12 + 0.10i)$
SMT	$4.0 - 5.0i$	$(4.02 - 5.03i) \pm (0.16 + 0.16i)$	$6.0 + 5.0i$	$(6.14 + 5.08i) \pm (0.19 + 0.18i)$
JCMT	$-4.0 + 5.0i$	$(-4.34 + 5.08i) \pm (0.35 + 0.37i)$	$-4.0 - 5.0i$	$(-4.22 - 4.69i) \pm (0.33 + 0.33i)$
LMT	$-4.0 + 3.0i$	$(-4.07 + 3.08i) \pm (0.22 + 0.22i)$	$6.0 - 3.0i$	$(6.29 - 3.01i) \pm (0.23 + 0.26i)$
IRAM 30 m	$2.0 + 3.0i$	$(1.75 + 2.13i) \pm (0.43 + 0.43i)$	$-2.0 - 7.0i$	$(-1.63 - 6.97i) \pm (0.43 + 0.41i)$
SMA	$-8.0 + 1.0i$	$(-8.23 + 1.01i) \pm (0.34 + 0.40i)$	$-2.0 + 1.0i$	$(-2.02 + 1.48i) \pm (0.29 + 0.32i)$

**Note.** A list of the leakage values derived from the DMC fit to the EHT-like synthetic data set, compared against the input values for all stations. The second and fourth columns list the input values, while the third and fifth columns list the posterior values for the right- and left-hand leakage terms, respectively. The leakage values in the third and fifth columns are quoted as posterior means, with the posterior standard deviation in both the real and imaginary components following the  $\pm$  symbol. All leakage values are quoted as percentages.

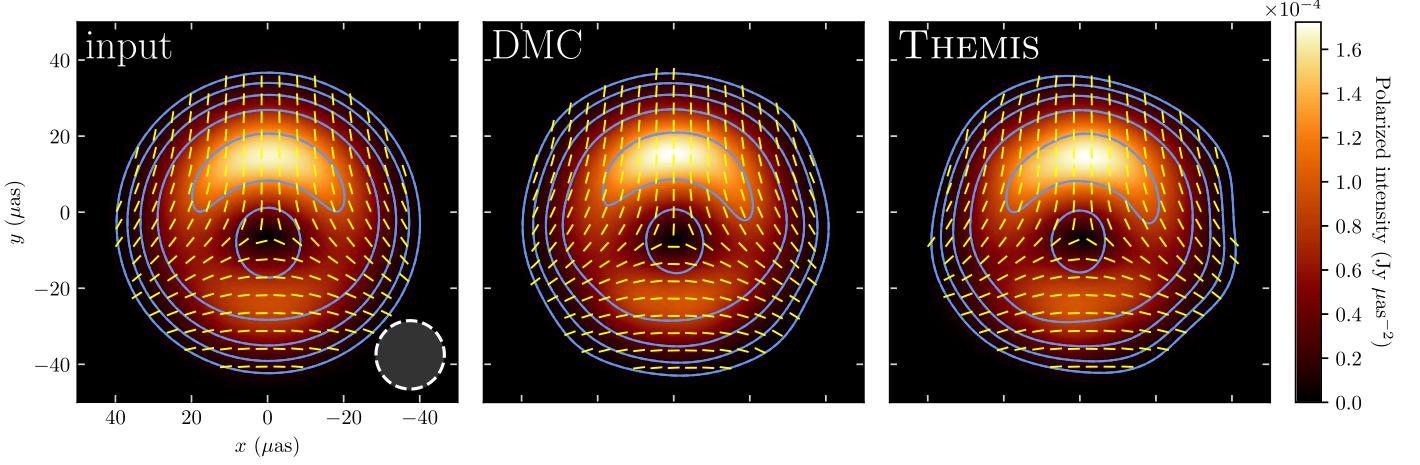
### 3.3. Imaging Real Data

We also apply DMC to a real data set, for which we use a polarimetric VLBI observation of the blazar OJ 287 obtained from the Boston University Blazar Research Group (Jorstad & Marscher 2016; Jorstad et al. 2017). A fully calibrated version of this data set is publicly available,<sup>9</sup> but we wish to demonstrate DMC’s ability to perform calibration itself alongside image reconstruction. Furthermore, applying DMC to a precalibrated data set may violate our assumption that the leakage terms are constant in time (see the Appendix). In this paper we thus use an earlier version of the OJ 287 data set for which only the a priori amplitude and phase calibrations have been applied (S. Jorstad 2020, private communication), but for which the full set of gain and leakage corruptions have not yet

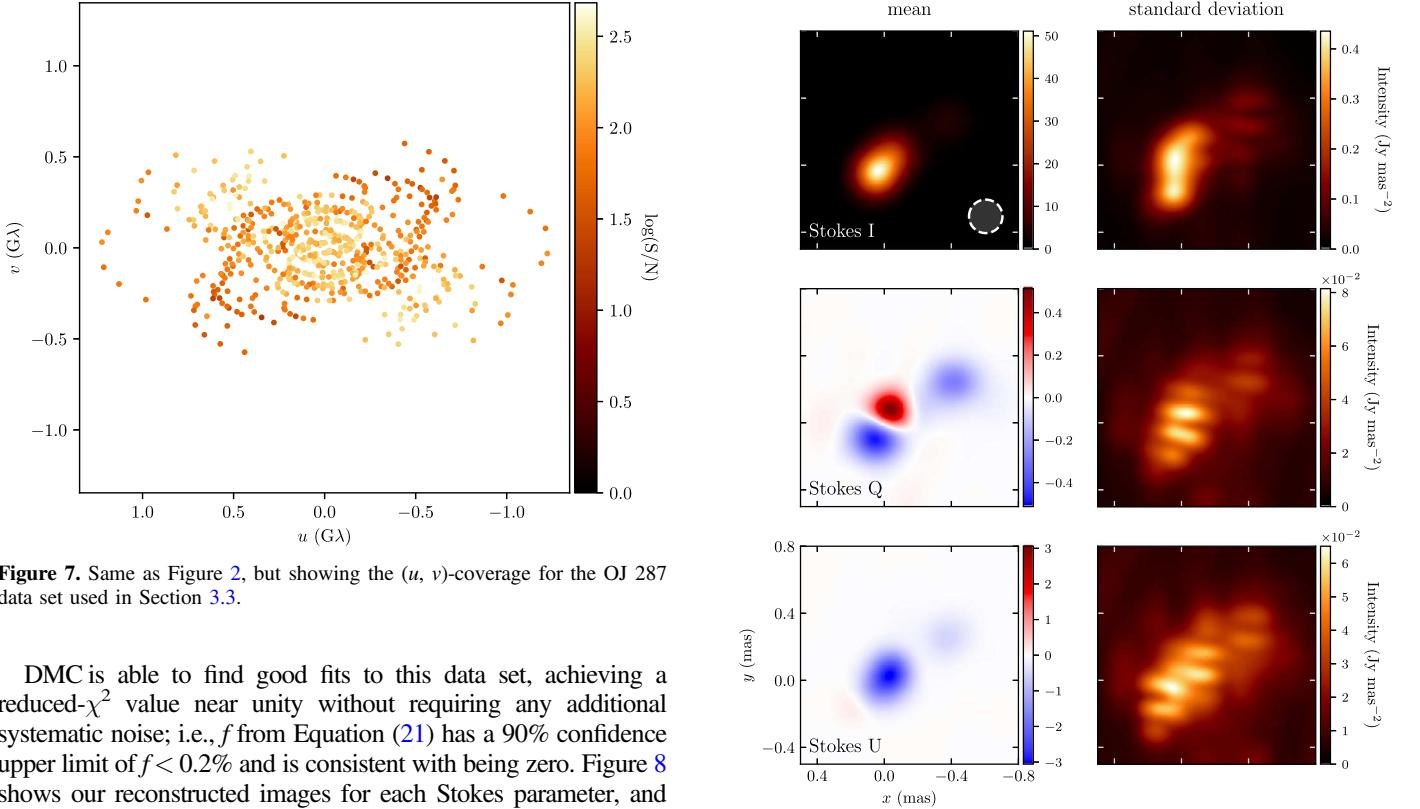
been removed. The observation was carried out with the Very Long Baseline Array (VLBA) on 2020 January 3 at an observing frequency of 43 GHz, and the data reduction procedure is described in Jorstad et al. (2005). The original data set is split into four separate frequency bands, but we combine bands during imaging to produce a single image and set of leakage terms. To reduce data volume, we coherently average the visibilities on a per-scan basis prior to imaging; the  $(u, v)$ -coverage of the resulting data set is shown in Figure 7.

We run DMC on the OJ 287 data set using  $\text{FOV}_x = \text{FOV}_y = 2$  mas,  $N_x = N_y = 20$ ,  $\check{x}_0 = \check{y}_0 = 0$  mas, and  $\Sigma = 0.2$  mas. As with the synthetic data set, we take the initial total flux estimate  $\check{F}$  to be equal to the visibility amplitude on the shortest baseline. The gain amplitude prior standard deviations are fixed to  $\check{\sigma}_R = \check{\sigma}_L = 0.05$  for all stations, and we select the Pie Town (PT) station as our reference antenna.

<sup>9</sup> <http://www.bu.edu/blazars/VLBAproject.html>



**Figure 6.** Comparison of the linear polarization structure in the EHT-like data set as recovered using DMC (middle) and THEMIS (right), with the input image shown in the left panel. The Stokes I structure is marked using blue contours, with the outermost contour levels starting at 5% of the peak intensity and increasing inward by factors of 2. The background color map indicates the linearly polarized intensity, and the overlaid tick marks show the EVPA direction. All three panels share a common field of view (explicitly labeled in the left panel), and all three images have been convolved with the EHT beam (circular Gaussian of width 18  $\mu$ as; Event Horizon Telescope Collaboration et al. 2019a) shown in the lower right corner of the left panel.



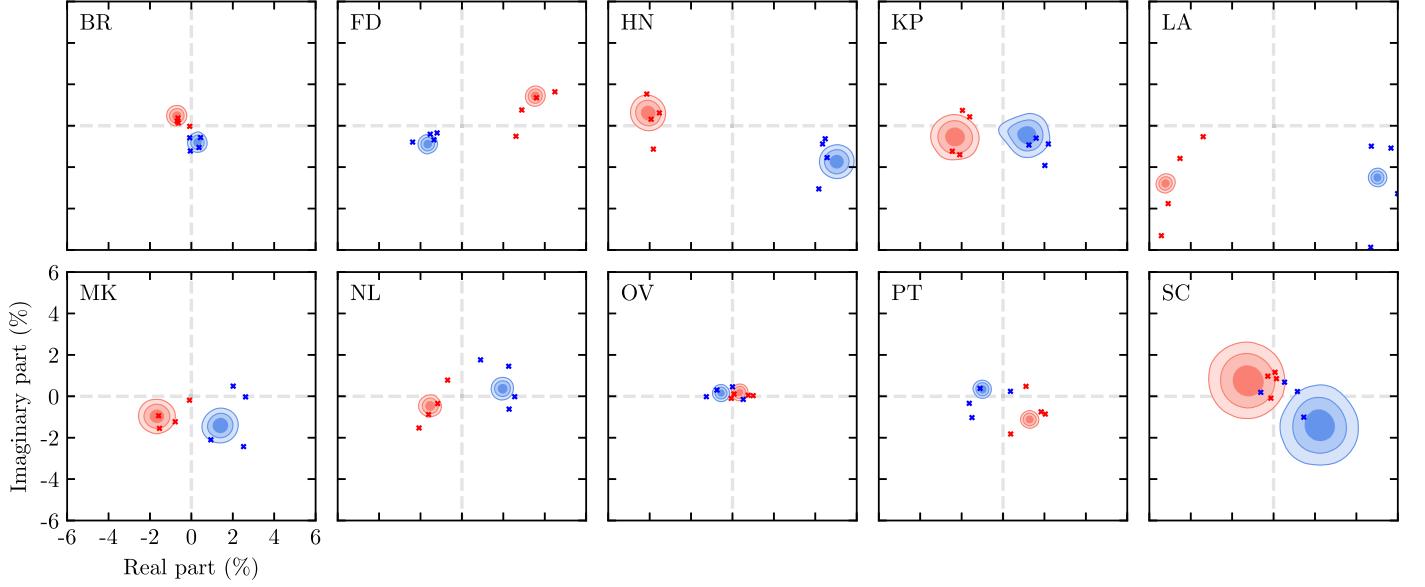
**Figure 7.** Same as Figure 2, but showing the  $(u, v)$ -coverage for the OJ 287 data set used in Section 3.3.

DMC is able to find good fits to this data set, achieving a reduced- $\chi^2$  value near unity without requiring any additional systematic noise; i.e.,  $f$  from Equation (21) has a 90% confidence upper limit of  $f < 0.2\%$  and is consistent with being zero. Figure 8 shows our reconstructed images for each Stokes parameter, and Figure 9 shows the derived station leakages and compares them against the results from LPCAL. The LPCAL leakages were determined separately for each of the four frequency bands in the original data set and are the result of averaging leakage solutions across 15 different observational targets (including OJ 287). We find that the leakage posteriors recovered by DMC—which are determined by solving for only a single leakage term for all four frequency bands at once—behave similarly to the LPCAL values for all stations and are largely consistent with the average of the LPCAL values across bands.

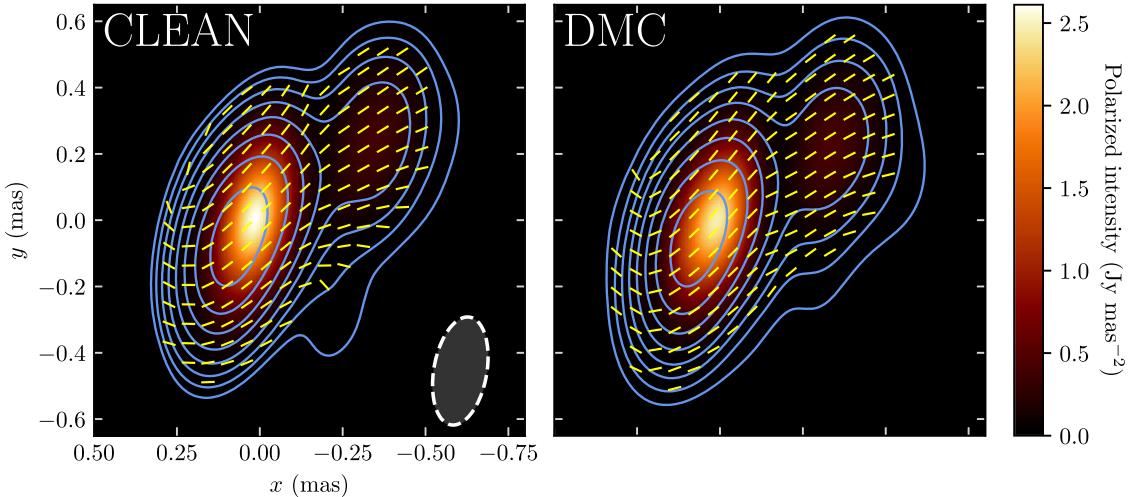
Figure 10 shows a comparison between the DMC and CLEAN image reconstructions for the OJ 287 data set, after restoring both with the CLEAN beam. We see good agreement

**Figure 8.** Similar to Figure 3, but for the DMC fit to the OJ 287 data set (for which we have no ground-truth image). No significant Stokes V signal is detected in this data set, so we show here only the Stokes I, Q, and U maps. The 0.2 mas FWHM of the Gaussian smoothing kernel is shown in the lower right corner of the top left panel, and we note again that this kernel does not represent a “restoring beam” but rather a convolving function that is self-consistently applied during imaging.

between both the Stokes I and linearly polarized image structures across most of the image, with noticeable deviations manifesting only at the  $\lesssim 1\%$  Stokes I contour level (which, as can be seen in Figure 8, is at the level of the uncertainty in the DMC Stokes I map).



**Figure 9.** Similar to Figure 5, but for the DMC fit to the OJ 287 data set; as in Figure 5, the blue contours correspond to right-hand leakages and the red contours to left-hand leakages. The overplotted blue and red crosses indicate the right- and left-hand leakage solutions, respectively, derived from the same data set using CLEAN and LPCAL (S. Jorstad 2020, private communication). There are four crosses associated with each leakage term because the original CLEAN analysis imaged each of the four frequency bands separately, while for the current DMC analysis we image all bands simultaneously and solve for a single set of leakages.



**Figure 10.** Analogous to Figure 6, but showing a comparison of the linear polarization structure in the OJ 287 data set as recovered using CLEAN (left) and DMC (right). The outermost Stokes I contour levels start at 0.5% of the peak intensity, and the contoured value increases inward by factors of 2. Both images have been convolved with the CLEAN beam (dimensions  $0.33 \times 0.16$  mas, position angle  $-10^\circ$  east of north) shown in the lower right corner of the left panel.

#### 4. Summary

In this paper we have presented DMC, a Python-based software package for performing polarimetric imaging of VLBI data. DMC simultaneously reconstructs the full-Stokes image structure and solves for the station-based gain and leakage calibration terms within a probabilistic formalism. The output of DMC is a sample of parameter values drawn from the posterior distribution, such that instead of a single image and associated set of calibration quantities one obtains an ensemble of images and sets of calibration quantities. DMC explores this posterior space using the NUTS HMC sampler implemented within PyMC3. Using this posterior distribution, it is possible to determine not only a best-fit value but also a self-consistent measure of uncertainty in any of the modeled parameters, including the sky-plane emission structure in all four Stokes

parameters, as well as the aforementioned calibration quantities. DMC therefore provides a tool for combining the information across different regions of an image (e.g., for averaging low-flux regions) or across multiple images (e.g., for spectral index or Faraday rotation analyses) in a manner that properly accounts for the associated uncertainties and correlations.

We have demonstrated the effectiveness of DMC on both synthetic and real data. The results from running DMC on an EHT-like synthetic data set show that DMC is capable of accurately recovering both the image structure and the ground-truth calibration quantities. Our results are compatible with those of THEMIS, which is also capable of exploring the full posterior space and returns distributions that are consistent with DMC's. Similarly, we find that running DMC on a real VLBA

data set recovers an image structure and leakage terms that agree well with those derived from an independent analysis using CLEAN and LPCAL.

DMC is built on a highly flexible modeling framework, and it continues to be developed. Future DMC capabilities may include more sophisticated priors on both the image and calibration terms, a more physically motivated calibration model, hybrid imaging (i.e., simultaneous imaging and modeling), compound sampling options, and time-, frequency-, and scale-dependent imaging.

I thank the anonymous referee for a constructive report that improved the quality of this paper. I am grateful to Lindy Blackburn, Avery Broderick, Andrew Chael, Shep Doeleman, Michael Johnson, Iván Martí-Vidal, and Daniel Palumbo for helpful discussions and comments. I'd like to further thank Andrew Chael for his guidance on generating synthetic data using the `eht-imaging` library, Avery Broderick for providing THEMIS results, and Svetlana Jorstad and Alan Marscher for providing the uncalibrated data set, LPCAL leakage solutions, and polarized CLEAN images of OJ 287. Support for this work was provided by the NSF through grants AST-1952099, AST-1935980, AST-1828513, and AST-1440254 and by the Gordon and Betty Moore Foundation through grant GBMF-5278. This work has been supported in part by the Black Hole Initiative at Harvard University, which is funded by grants from the John Templeton Foundation and the Gordon and Betty Moore Foundation to Harvard University. This study makes use of 43 GHz VLBA data from the VLBA-BU Blazar Monitoring Program (VLBA-BU-BLAZAR; <http://www.bu.edu/blazars/VLBAproject.html>), funded by NASA through the Fermi Guest Investigator Program. The VLBA is an instrument of the National Radio Astronomy Observatory. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated by Associated Universities, Inc.

*Software:* `eht-imaging` (Chael et al. 2016, 2018), `ehtplot`,<sup>10</sup> PyMC3 (Salvatier et al. 2016), THEMIS (Broderick et al. 2020a; A. E. Broderick et al. 2021, in preparation).

## Appendix

### Residual Jones Matrices

In this section we explore the notion of a “residual” Jones matrix; i.e., what is the form of the Jones matrix for the remaining calibrations after an imperfect calibration has been applied? Consider a data set that has had such an attempted (but imperfect) calibration applied to it. The true Jones matrix  $\mathbf{J}$  looks like (see Equation (14))

$$\mathbf{J} = \begin{pmatrix} G_R & 0 \\ 0 & G_L \end{pmatrix} \begin{pmatrix} 1 & D_R \\ D_L & 1 \end{pmatrix} \begin{pmatrix} e^{-i\phi} & 0 \\ 0 & e^{i\phi} \end{pmatrix}, \quad (\text{A1})$$

but the applied inverse Jones matrix  $\hat{\mathbf{J}}$  has some imperfections in it. We can assume that these imperfections manifest in the station gain and leakage terms, but not in the field rotation

<sup>10</sup> <https://github.com/liamedeiros/ehtplot>

angles.  $\hat{\mathbf{J}}$  may thus be written as

$$\hat{\mathbf{J}} = \begin{pmatrix} G_R(1 + \epsilon_R) & 0 \\ 0 & G_L(1 + \epsilon_L) \end{pmatrix} \times \begin{pmatrix} 1 & D_R(1 + \delta_R) \\ D_L(1 + \delta_L) & 1 \end{pmatrix} \begin{pmatrix} e^{-i\phi} & 0 \\ 0 & e^{i\phi} \end{pmatrix}. \quad (\text{A2})$$

The act of applying the imperfect calibration corresponds to attempting to “invert”  $\mathbf{J}$  using  $\hat{\mathbf{J}}^{-1}$ , which yields a residual Jones matrix,

$$\mathbf{R} \equiv \hat{\mathbf{J}}^{-1} \mathbf{J} = \begin{pmatrix} \Gamma_R & 0 \\ 0 & \Gamma_L \end{pmatrix} \begin{pmatrix} 1 & \Delta_R \\ \Delta_L & 1 \end{pmatrix} \begin{pmatrix} e^{-i\phi} & 0 \\ 0 & e^{i\phi} \end{pmatrix}. \quad (\text{A3})$$

Here we have defined the residual gains,

$$\Gamma_R \equiv e^{i\phi} \left( \frac{r D_R D_L (1 + \delta_R) - 1}{[D_R (1 + \delta_R) D_L (1 + \delta_L) - 1] (1 + \epsilon_R)} \right), \quad (\text{A4a})$$

$$\Gamma_L \equiv e^{-i\phi} \left( \frac{D_R D_L (1 + \delta_L) - r}{[D_R (1 + \delta_R) D_L (1 + \delta_L) - 1] (1 + \epsilon_R)} \right); \quad (\text{A4b})$$

the residual leakages,

$$\Delta_R \equiv D_R \left( \frac{1 + \delta_R - r}{D_R D_L (1 + \delta_R) - r} \right), \quad (\text{A5a})$$

$$\Delta_L \equiv D_L \left( \frac{r (1 + \delta_L) - 1}{r D_R D_L (1 + \delta_L) - 1} \right); \quad (\text{A5b})$$

and the fractional gain residuals,

$$r \equiv \frac{1 + \epsilon_L}{1 + \epsilon_R}. \quad (\text{A6})$$

We see that while the residual Jones matrix  $\mathbf{R}$  can be cast in the form of the original Jones matrix, the residual gain and leakage terms may no longer adhere to the same set of assumptions that apply to the true gains and leakages. In particular, the residual leakages depend on the fractional gain residuals,  $r$ , which may not be constant in time. It may therefore be preferable to redo calibration entirely rather than attempting to incrementally calibrate a partially calibrated data set; in Section 3.3, we pursue the former strategy.

## ORCID iDs

Dominic W. Pesce  <https://orcid.org/0000-0002-5278-9221>

## References

Akiyama, K., Ikeda, S., Pleau, M., et al. 2017, *AJ*, 153, 159  
 Arras, P., Frank, P., Leike, R., Westermann, R., & Enßlin, T. A. 2019, *A&A*, 627, A134  
 Bastien, F., Lamblin, P., Pascanu, R., et al. 2012, arXiv:1211.5590  
 Bergstra, J., Breuleux, O., Bastien, F., et al. 2010, in Proc. 9th Python for Scientific Computing Conf., ed. S. van der Walt & J. Millman (Austin, TX: SciPy), 18  
 Betancourt, M. 2017, arXiv:1701.02434  
 Blackburn, L., Pesce, D. W., Johnson, M. D., et al. 2020, *ApJ*, 894, 31  
 Broderick, A. E., Gold, R., Karami, M., et al. 2020a, *ApJ*, 897, 139  
 Broderick, A. E., Pesce, D. W., Tiede, P., Pu, H.-Y., & Gold, R. 2020b, *ApJ*, 898, 9  
 Brown, L. F., Roberts, D. H., & Wardle, J. F. C. 1989, *AJ*, 97, 1522  
 Cai, X., Pereyra, M., & McEwen, J. D. 2018a, *MNRAS*, 480, 4154  
 Cai, X., Pereyra, M., & McEwen, J. D. 2018b, *MNRAS*, 480, 4170

Carpenter, B., Gelman, A., Hoffman, M., et al. 2017, *JOSS*, 76, 1

Chael, A. A., Johnson, M. D., Bouman, K. L., et al. 2018, *ApJ*, 857, 23

Chael, A. A., Johnson, M. D., Narayan, R., et al. 2016, *ApJ*, 829, 11

Clark, B. G. 1980, *A&A*, 89, 377

Conway, R. G., & Kronberg, P. P. 1969, *MNRAS*, 142, 11

Cornwell, T. J., & Evans, K. F. 1985, *A&A*, 143, 77

Duane, S., Kennedy, A., Pendleton, B. J., & Roweth, D. 1987, *PhLB*, 195, 216

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019a, *ApJL*, 875, L4

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019b, *ApJL*, 875, L1

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019c, *ApJL*, 875, L2

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019d, *ApJL*, 875, L3

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019e, *ApJL*, 875, L5

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019f, *ApJL*, 875, L6

Goddi, C., Martí-Vidal, I., Messias, H., et al. 2019, *PASP*, 131, 075003

Hamaker, J. P. 2000, *A&AS*, 143, 515

Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, *A&AS*, 117, 137

Hoffman, M., & Gelman, A. 2014, *JMLR*, 15, 1593, <https://jmlr.org/papers/volume15/hoffman14a/hoffman14a.pdf>

Högberg, J. A. 1974, *A&AS*, 15, 417

Jorstad, S., & Marscher, A. 2016, *Galax*, 4, 47

Jorstad, S. G., Marscher, A. P., Lister, M. L., et al. 2005, *AJ*, 130, 1418

Jorstad, S. G., Marscher, A. P., Morozova, D. A., et al. 2017, *ApJ*, 846, 98

Kulkarni, S. R. 1989, *AJ*, 98, 1112

Martí-Vidal, I., Roy, A., Conway, J., & Zensus, A. J. 2016, *A&A*, 587, A143

Neal, R. M. 2011, in *Handbook of Markov Chain Monte Carlo*, ed. S. Brooks, A. Gelman, G. Jones et al. (Boca Raton, FL: CRC Press), 113

Nityananda, R., & Narayan, R. 1982, *JApA*, 3, 419

Palumbo, D. C. M., Wong, G. N., & Prather, B. S. 2020, *ApJ*, 894, 156

Readhead, A. C. S., Walker, R. C., Pearson, T. J., & Cohen, M. H. 1980, *Natur*, 285, 137

Roberts, D. H., Wardle, J. F. C., & Brown, L. F. 1994, *ApJ*, 427, 718

Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. 2016, *PeerJ Computer Science*, 2, e55

Schwarb, F. R. 1984, *AJ*, 89, 1076

Smirnov, O. M. 2011, *A&A*, 527, A106

Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr. 2017, *Interferometry and Synthesis in Radio Astronomy* (3rd ed.; Berlin: Springer)