



Markov Chains for Horizons MARCH. I. Identifying Biases in Fitting Theoretical Models to Event Horizon Telescope Observations

Dimitrios Psaltis¹ , Feryal Özel¹ , Lia Medeiros^{1,2,5} , Pierre Christian¹ , Junhan Kim^{1,3} , Chi-kwan Chan¹ ,
Landen J. Conway¹, Carolyn A. Raithel¹ , Dan Marrone¹ , and Tod R. Lauer⁴

¹Department of Astronomy, University of Arizona, 933 N. Cherry Ave, Tucson, AZ 85721, USA

²School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA

³California Institute of Technology, 1200 E. California Blvd., MC 367-17, Pasadena, CA 91125, USA

⁴National Optical Astronomy Observatory, 950 North Cherry Ave., Tucson, AZ 85719, USA

Received 2020 May 19; revised 2021 September 3; accepted 2021 October 1; published 2022 March 25

Abstract

We introduce a new Markov Chain Monte Carlo (MCMC) algorithm with parallel tempering for fitting theoretical models of horizon-scale images of black holes to the interferometric data from the Event Horizon Telescope (EHT). The algorithm implements forms of the noise distribution in the data that are accurate for all signal-to-noise ratios. In addition to being trivially parallelizable, the algorithm is optimized for high performance, achieving 1 million MCMC chain steps in under 20 s on a single processor. We use synthetic data for the 2017 EHT coverage of M87 that are generated based on analytic as well as General Relativistic Magnetohydrodynamic (GRMHD) model images to explore several potential sources of biases in fitting models to sparse interferometric data. We demonstrate that a very small number of data points that lie near salient features of the interferometric data exert disproportionate influence on the inferred model parameters. We also show that the preferred orientations of the EHT baselines introduce significant biases in the inference of the orientation of the model images. Finally, we discuss strategies that help identify the presence and severity of such biases in realistic applications.

Unified Astronomy Thesaurus concepts: Supermassive black holes (1663); Astrophysical black holes (98); Astrostatistics (1882); Algorithms (1883); Very long baseline interferometry (1769); Interferometry (808)

1. Introduction

Images of low-luminosity black holes at millimeter wavelengths, where the accretion flow is transparent down to horizon scales (Özel et al. 2000), have long been expected to be compact, dominated by the shadow cast by the black hole on the surrounding plasma emission (Jaroszynski & Kurpiewski 1997; Falcke et al. 2000; Dexter et al. 2009; Broderick et al. 2009; Mościbrodzka et al. 2009, 2014; Chan et al. 2015). Recent Event Horizon Telescope (EHT) observations of the black hole in the center of the M87 galaxy have confirmed this expectation, revealing an image consisting of a narrow ring of emission surrounding the black hole shadow (Event Horizon Telescope Collaboration et al. 2019a, 2019b, 2019c, 2019d, 2019e, 2019f).

As a Very Long Baseline Interferometric (VLBI) array, the EHT measures the complex Fourier components of the brightness distribution in the sky, also called the interferometric visibilities, at a distinct set of spatial frequencies determined by the baselines between the various stations in the array. In principle, the image of the source can then be obtained by constructing the inverse Fourier transform of the observed visibilities. In practice, however, the sparse coverage of the Fourier space by the EHT, which consists of a small number of baselines separated by large distances, makes such a direct reconstruction of the image impossible.

Traditionally, VLBI measurements have been converted to images with CLEAN algorithms (see, e.g., Högbom 1974; Clark 1980), which identify the minimum set of point sources

in the image plane with Fourier transforms that are consistent with the visibility data. These point sources are then convolved with a filter to broaden the points sources to the effective resolution of the interferometric array.

Alternatively, interferometric images have been obtained with regularized maximum likelihood methods, in which a pixel-based model is defined in the image plane with its parameters being the brightness at each pixel location. The Fourier transform of this model is compared to the observed visibilities, and the brightness of each pixel is adjusted to minimize an appropriate χ^2 statistic in the visibility space. Because the number of pixels of the model image is typically much larger than the number of measurements, regularizers are employed to decrease the effective number of degrees of freedom. This approach was introduced in radio interferometry imaging with maximum-entropy regularizers (see Narayan & Nityananda 1986 and references therein), has been developed extensively in optical/IR interferometry (Thiébaut 2009; Berger et al. 2012), and has been employed in imaging black holes with the EHT (see, e.g., Honma et al. 2014; Chael et al. 2016, 2018; Akiyama et al. 2017a, 2017b; Event Horizon Telescope Collaboration et al. 2019d).

The regularized maximum likelihood methods are fundamentally nonparametric model-fitting algorithms. As such, they are agnostic in regards to the underlying image structure. However, this freedom comes with the usual “curse of dimensionality”: the requirement of quantifying the brightness on a very large number of pixels typically leads to presenting a single (or a very small number of) best-fit image with marginal exploration of the parameter space of possible solutions that might fit the data equally well. Moreover, the image properties show a dependence on the regularizers, and the strength of the regularizers is usually chosen such that they suppress structures at scales that are significantly smaller than the effective beam of the array (see,

⁵ NSF Astronomy and Astrophysics Postdoctoral Fellow.



e.g., Event Horizon Telescope Collaboration et al. 2019d). This limits the ability of the algorithms to provide precise measurements of physical quantities of interest.

In order to elucidate the limitations of an agnostic imaging algorithm, we can consider interferometric observations of a simple source consisting of two point sources separated by some angular distance. An imaging algorithm will only be able to measure the separation of the two point sources down to an accuracy comparable to the effective beam size of the array, which depends mainly on the wavelength of observation and the separations of the interferometric stations. However, if we know a priori that we are looking at an image with two point sources, we can use the same interferometric data to measure the angular separation of the point source with a much higher accuracy that will be limited primarily by the signal-to-noise ratio of the measurements. Achieving such a measurement for a realistic image, of course, requires both an accurate parametric model for the underlying image as well as an algorithm to explore the posterior distribution of the statistical comparison of the model to the data.

In the first article of this series, we describe the new Markov-Chain Bayesian algorithm MARCH (MArkov Chains for Horizons) that we have developed to fit models of black hole images to the interferometric data at millimeter wavelengths obtained with the EHT. Similar Bayesian approaches have already been applied to early EHT observations of the galactic center black hole Sgr A* (see, e.g., Broderick et al. 2016; Lu et al. 2018) as well as to the 2017 EHT observations of M87 that have led to the initial measurements of the mass of the latter black hole (Event Horizon Telescope Collaboration et al. 2019f). Our algorithm aims to address a number of complications to this otherwise straightforward statistical problem introduced by theoretical expectations on the images of black holes and the particular characteristics of the EHT array.

First, the structure of the millimeter image of Sgr A* is expected to show substantial variability within the few-hour time span of the observations because of its small mass (see, e.g., Medeiros et al. 2017, 2018; Roelofs et al. 2017). As a result, fitting a static model image to variable interferometric data will introduce biases to the inferred parameters and artificially reduce their inferred uncertainties. This complication was addressed in an earlier work (Kim et al. 2016), and we will not revisit it here.

Second, the images of black holes with horizon-scale resolution are expected to have ring-like structures with very sharp boundaries at the locations of the black hole shadows (see, e.g., Kamruddin & Dexter 2013; Psaltis et al. 2015). The interferometric visibilities of these structures show a number of salient features, such as deep minima of the visibility amplitudes at particular baseline lengths, with rapid swings of the visibility phases across them. The location of the first minimum depends very strongly on the size of the ring of emission, and the location of the second minimum depends on its fractional width (Event Horizon Telescope Collaboration et al. 2019f). While the strong dependence of the salient features on model parameters allows us to measure them accurately, the most likely values of the model parameters and their uncertainties depend disproportionately on how well the model describes the very small number of data points near the salient features. As we will show in Section 3, the combination of the simplicity of the model images with the high signal-to-noise ratio of the measurements causes a straightforward application of Bayesian posteriors to return biased measurements of the model parameters.

Finally, the location of landmass and high mountains on the globe forces the EHT baselines to lie primarily along two

orientations, north–south and east–west (Event Horizon Telescope Collaboration et al. 2019b). As a result of the projection-slice theorem, the measured interferometric visibilities are primarily sensitive only to the projections of the underlying image along these two orientations. We will show in Section 4 that this biases the inferred orientations of the model images to line up with the major orientations of the baselines. The origin of this bias is similar to the development of spurious “knots” along the baseline orientations in the reconstruction of otherwise smooth images from synthetic EHT data (see Figure 10 of Event Horizon Telescope Collaboration et al. 2019d). Additional evidence for this effect comes from the initial application of geometric and GRMHD model fitting to the 2017 EHT data of M87, which revealed that the most likely black hole orientation appears aligned with the orientations of the EHT baselines (see Figure 6 of Event Horizon Telescope Collaboration et al. 2019f) and away from the orientation of the long-wavelength jet that emerges from the central black hole (see Figure 9 of Event Horizon Telescope Collaboration et al. 2019e).

In Section 2 below, we introduce the basic formalism of our Bayesian Markov Chain algorithm and its implementation. In Sections 3–5, we apply this algorithm to various cases of simple synthetic data chosen to explore the issues discussed above and provide strategies for their mitigation. In Section 6, we explore realistic EHT data based on GRMHD simulations of M87 and, in Section 7, we conclude by discussing our results and providing an outlook on the accuracy of measurements with the EHT. In subsequent papers in this series, we will discuss our implementation of modeling the unknown gains of the various interferometric stations and the biases introduced by this approach; we will also explore different information criteria that allow us to quantify the necessity of increasing the complexity of models in order to fit realistic data.

2. The MARCH Algorithm

2.1. Interferometric Observables

The goal of the MARCH algorithm is to estimate the posterior distributions over a number of parameters of a model by comparing to EHT observations the model image brightness in the sky $I(x, y; \vec{\theta})$; here x and y are angular coordinates, and $\vec{\theta}$ is the vector of the model parameters. As an interferometer, the EHT does not directly observe the brightness distribution of the sky but rather the complex visibilities defined by the van Cittert-Zernike theorem as

$$\mathcal{V}(u, v; \vec{\theta}) \equiv \iint e^{-2\pi i(xu + yv)} I(x, y; \vec{\theta}) dx dy. \quad (1)$$

In this equation, u and v are equal to the two components of the baseline vectors \mathbf{b} that connect every pair of sites in the array, projected orthogonally to the line of sight to a particular source, and divided by the wavelength λ of observation, i.e., $(u, v) \equiv \mathbf{b}/\lambda$. The brightness distribution in the image domain and the visibilities in the u – v domain are clearly connected via a two-dimensional Fourier transform.

For a general, nonsymmetric brightness distribution, the interferometric visibilities are complex numbers that can be described either in terms of their real and imaginary components or in terms of their amplitudes and phases. The former description is useful in quantifying the thermal noise in the measurements, which can be well approximated by independent Gaussian distributions in the real and imaginary components of

the visibilities (see Thompson et al. 2017). However, because of atmospheric and instrumental effects, VLBI measurements seldom allow for a precise determination of the complex visibilities. In practice, we write the measured visibility between the i -th and j -th station as

$$V_{ij} \equiv V(u_{ij}, v_{ij}) = g_i g_j^* \mathcal{V}(u_{ij}, v_{ij}), \quad (2)$$

where g_i and g_j are complex “gains” of the two telescopes, and a star superscript denotes complex conjugates. Hereafter, as in this equation, we will denote the measured visibilities between the two stations simply as V_{ij} and suppress for brevity the argument of this quantity.

The magnitudes of the gain factors for the various telescopes can often be constrained with a priori calibration to within $\sim 10\%$ levels (Event Horizon Telescope Collaboration et al. 2019c). On the other hand, the complex phases of the gains can fluctuate wildly and over very short timescales because of atmospheric instabilities and instrumental effects. To overcome this problem, VLBI measurements quote the amplitudes $|V_{ij}|$ of the complex visibilities between pairs of stations and the arguments of the complex bispectra (also known as closure phases) along a closed triangle of stations (Jennison 1958; Thompson et al. 2017). The latter is simply

$$\begin{aligned} \text{Arg}[V_{ij} V_{jk} V_{ki}] &= \text{Arg}[g_i g_j^* g_j g_k^* g_k g_i^* \mathcal{V}_{ij} \mathcal{V}_{jk} \mathcal{V}_{ki}] \\ &= \text{Arg}[|g_i|^2 |g_j|^2 |g_k|^2] + \text{Arg}[\mathcal{V}_{ij} \mathcal{V}_{jk} \mathcal{V}_{ki}] \\ &= \text{Arg}[\mathcal{V}_{ij} \mathcal{V}_{jk} \mathcal{V}_{ki}], \end{aligned} \quad (3)$$

i.e., the closure phases are not affected by the unknown phases of the telescope gains.

We pay three penalties by using visibility amplitudes and closure phases instead of the real and imaginary components of the complex interferometric visibilities. First, by combining three complex visibilities to form one closure quantity, we lose some of the available information. Second, the transformation between the brightness in the image domain and the observables in the visibility domain is no longer linear. As a result, adding a perturbation to an image that would only introduce minor variations in the complex visibilities could still lead to order unity changes in visibility amplitudes and closure phases at various baseline lengths. This is important to understanding the consequences of the strong dependence of the model parameters on the salient features of the visibilities. Third, the error distributions of the visibility amplitudes and closure phases are not Gaussian, even for purely thermal noise.

The likelihood of making a visibility amplitude measurement V_{ij} given a true visibility amplitude of $V_{ij,0}$ is the Rice distribution (Thompson et al. 2017)

$$\mathcal{L}_A(V_{ij}; V_{ij,0}, \sigma_{ij}) = \frac{V_{ij}}{\sigma_{ij}^2} \exp\left[-\frac{V_{ij}^2 + V_{ij,0}^2}{2\sigma_{ij}^2}\right] I_0\left(\frac{V_{ij} V_{ij,0}}{\sigma_{ij}^2}\right), \quad (4)$$

where $I_0(x)$ is the modified Bessel function of zeroth order, and σ_{ij} is the dispersion of thermal noise introduced to the real and imaginary components of the visibility by the array configuration used. At large signal-to-noise ratios $V_{ij,0}/\sigma_{ij}$, the Rice distribution reduces to a Gaussian with the same dispersion. The computational cost of using the Rice distribution is only marginally higher than using the Gaussian approximation.

The likelihood of making a closure phase measurement Φ_{ijk} is given by the triple convolution

$$\begin{aligned} \mathcal{L}_\Phi(\Phi_{ijk}; \Phi_{ijk,0}, \sigma_{ij}, \sigma_{jk}, \sigma_{ki}) &= \int d\theta \int d\phi L_{\text{ph}}(\theta; \text{SNR}_{ij}) \\ &\quad L_{\text{ph}}(\phi - \theta; \text{SNR}_{jk}) L_{\text{ph}}(\Phi_{ijk} - \Phi_{ijk,0} - \phi; \text{SNR}_{ki}), \end{aligned} \quad (5)$$

where θ and ϕ are dummy integration variables; they represent the fact that two of the visibility phases are independent but the third one is constrained by the closure phase sum. Here $\Phi_{ijk,0}$ is the true closure phase, and $L_{\text{ph}}(\theta; \text{SNR}_{ij})$ is the likelihood of measuring for the phase of the ij – baseline an angle θ away from the true value. This last likelihood is given by Thompson et al. (2017)

$$\begin{aligned} L_{\text{ph}}(\theta; \text{SNR}_{ij}) &= \frac{e^{-\frac{\text{SNR}_{ij}^2}{2}}}{4\pi\sigma_{ij}} \left\{ \sqrt{2\pi} \text{SNR}_{ij} \cos \theta e^{\frac{\text{SNR}_{ij}^2 \cos^2 \theta}{2}} \right. \\ &\quad \times \left[\text{Erf}\left(\frac{\text{SNR}_{ij} \cos \theta}{\sqrt{2}}\right) + 1 \right] + 2 \left. \right\}, \end{aligned} \quad (6)$$

where $\text{SNR}_{ij} \equiv V_{ij,0}/\sigma_{ij}$ is the signal-to-noise ratio of the measurement. The triple convolution necessary to calculate the likelihood for the closure phase is too expensive computationally. Christian & Psaltis (2020) developed an efficient numerical algorithm that allows us to obtain an accurate approximation of the complete likelihood (5), which we use here. At the limit of large signal-to-noise ratios, this distribution can be well approximated by a Gaussian with dispersion

$$\sigma_{\Phi_{ijk}} = \left(\frac{\sigma_{ij}^2}{V_{ij,0}^2} + \frac{\sigma_{jk}^2}{V_{jk,0}^2} + \frac{\sigma_{ki}^2}{V_{ki,0}^2} \right)^{1/2}. \quad (7)$$

2.2. Bayesian Posteriors

Given a set of interferometric data points, we calculate the posterior distribution of the model parameters $\vec{\theta}$ using Bayes’ theorem, i.e.,

$$P(\vec{\theta}|\text{data}) = C P_{\text{pri}}(\vec{\theta}) \mathcal{L}(\text{data}|\vec{\theta}). \quad (8)$$

Here, $P_{\text{pri}}(\vec{\theta})$ is the prior distribution over the model parameters, $\mathcal{L}(\text{data}|\vec{\theta})$ is the likelihood that the set of observed data can be obtained from a given set of model parameters, and C is a normalization constant.

The likelihood function is obtained by a proper combination of the likelihoods of the visibility amplitude and closure phase data in Equations (4) and (5), respectively. Writing schematically, the likelihood of the i -th data point as $\mathcal{L}_i(\text{data}|\vec{\theta})$ and assuming that all likelihoods are independent of each other (see Blackburn et al. 2020), we obtain

$$\mathcal{L}(\text{data}|\vec{\theta}) = \prod_i \mathcal{L}_i(\text{data}|\vec{\theta}). \quad (9)$$

The priors over the model parameters depend on the specifics of the model considered. For the analytic crescent model we will be using below, we set all the priors to boxcar functions with limits that are much larger than the widths of the expected posteriors. For reasons that will be discussed in the following sections, the posteriors have very small fractional widths (often less than 1%), and therefore the precise function

form of such noninformative priors does not affect the final results.

2.3. Markov Chains, Parallel Tempering, and Performance

The sparse coverage of the u - v plane with the EHT baselines results in a small number of independent measurements and hence warrants fitting models to the data that only have a relatively small number of parameters. Moreover, the salient features of the interferometric data, such as the baseline lengths that correspond to minima in the visibility amplitude, allow for relatively accurate initial guesses for the values of several model parameters (see the discussion in Sections 3 and 4 below). For these reasons, we employ a simple Metropolis–Hastings stepping algorithm for each chain of the MCMC algorithm. In particular, we choose steps for each model parameters from independent Gaussian distributions with dispersions set equal to a fraction of the expected uncertainties in the corresponding model parameters.

The presence of salient features in the interferometric data often leads to narrow widths for the posteriors and strong correlations between some model parameters (see Section 4 below). In order to both efficiently explore the broad parameter space and take small steps to sample the narrow posteriors, we employ a parallel tempering strategy, as discussed, e.g., in Earl & Deem (2005). In this approach, the posteriors are sampled at various levels, i.e., replicas, each of which has been artificially broadened by some a priori chosen temperature. Each successive replica in the ladder has an increasing temperature, i.e., corresponding to a larger broadening. Only the lowest level replica, which corresponds to no broadening, is used for the final calculation of the posteriors. However, the other replicas help sample the parameter space more efficiently and increase the chance that a Markov chain step will get closer to and sample a region of high posterior. The number and spacing of the parallel tempered ladders (the replicas) we use depend on the complexity of the data and of the model. For the examples with the synthetic data we use in this paper, we typically use a ladder of 40 replicas with temperatures starting from unity and increasing from this lowest replica in multiples of $\sqrt{3}$.

Like all MCMC algorithms, MARCH is trivially parallelizable. However, special care was given in its development to boost its computational efficiency and performance in order to enable fitting complex, numerical models to the interferometric data. In particular, we ensured that computationally expensive operations (such as logs, powers, and trigonometric functions) are only sparsely used and organized the algorithmic flow in such a way that the same quantity is not repeatedly computed. As a result, for analytic models, MARCH achieves very high performance even on a single CPU.

Figure 1 shows the CPU time (on a Macbook Air with an 1.6 GHz Intel Core i5 processor) per temperature replica as a function of the number of MCMC steps, for the synthetic data problem discussed in Section 3 below. In this example, 152 data points were fit with an analytic crescent model described by five model parameters, with the MCMC stepping per parameter adjusted so that the average acceptance ratio of the zero-temperature chain is equal to 20%. The figure demonstrates the linear scaling of the algorithm on the number of MCMC steps and its high performance, which reaches (for this example) $\simeq 17$ seconds per 10^6 MCMC steps. The figure also compares the performance of MARCH to that of a commonly used Python MCMC algorithm (emcee; Foreman-Mackey

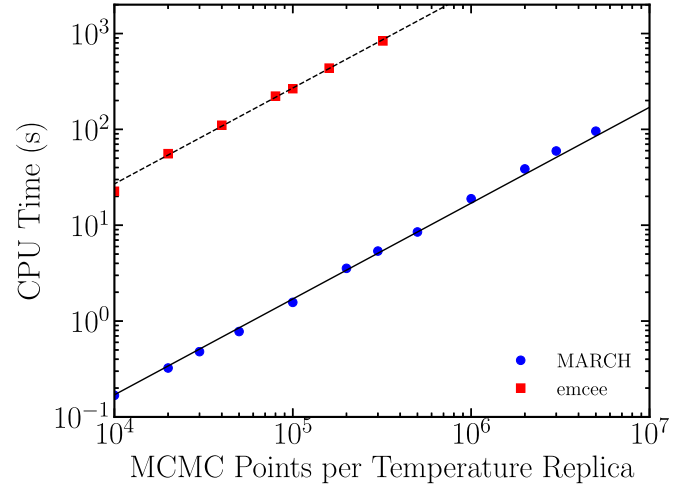


Figure 1. The CPU time for the execution of the MARCH algorithm on a single CPU, as a function of the number of MCMC steps per temperature replica for the synthetic data problem discussed in Section 3. The red points show the corresponding performance of the commonly used Python emcee algorithm. MARCH is optimized to achieve high performance even on a single CPU, which reaches ~ 17 s per 10^6 MCMC steps for this example.

et al. 2013), running the identical problem. The fact that MARCH was optimized for the problem under consideration leads to an increase in performance by a factor of ~ 150 compared to the more general emcee algorithm.

3. An Analytic Model of Crescent Images

In order to explore some of the characteristics of fitting models to the EHT visibility data, we will employ here the geometric model of a crescent image devised by Kamruddin & Dexter (2013). This model creates a crescent by subtracting the brightness of two uniform but displaced disks (see Figure 1 of Kamruddin & Dexter (2013) and Figure 7 of Medeiros et al. (2018) for different visual representations of this geometric model image). Additional complexities to the model were added in Benkevitch et al. (2016), who allowed for a gradient in the brightness across the crescent, and in Event Horizon Telescope Collaboration et al. (2019f), where a number of displaced Gaussian components were included to the image. Albeit simple, this model captures all the basic characteristics of the images encoded in the first EHT data on M87 (see Event Horizon Telescope Collaboration et al. 2019f).

The complex visibility of a simple crescent image is given by Kamruddin & Dexter (2013)

$$V(u, v) = 2\pi I_0 \left[\frac{R_p J_1(2\pi k R_p)}{2\pi k} - e^{-2\pi i(au+bv)} (1-f) \frac{R_n J_1(2\pi k R_n)}{2\pi k} \right], \quad (10)$$

where R_p and R_n are the outer and inner disk angular radii, respectively, $J_1(x)$ is the Bessel function of the first kind, I_0 is the uniform brightness of the crescent, f measures the fractional depression of brightness at its center,

$$k \equiv \sqrt{u^2 + v^2}, \quad (11)$$

is the baseline length in units of the wavelength, and the coordinate pair (a, b) measures the displacement of the center of




Figure 2. Normalized visibility amplitude as a function of baseline length for symmetric rings with different fractional widths ψ and radii such that they all show the first minimum at the same baseline length. The radius of the ring with infinitesimal width is set to $R = 21 \mu\text{as}$, to resemble the 2017 EHT image of M87. The fractional width of the ring determines the relative amplitude of the maximum around $\sim 6 \text{ G}\lambda$ and, for $\psi \gtrsim 0.5$, the location of the second minimum.

the inner disk with respect to that of the outer disk along the two orientations in the sky.

In order to simplify our notation, we express the characteristic radius of the crescent in terms of the average radii of the two uniform disks

$$R \equiv \frac{R_p + R_n}{2} \quad (12)$$

and the fractional width of a symmetric crescent (hereafter a ring) as their normalized difference⁶

$$\psi \equiv \frac{R_p - R_n}{R_p}. \quad (13)$$

For nonsymmetric crescents (i.e., when $a \neq 0$ or $b \neq 0$), we define the orientation of the crescent as

$$\phi \equiv \tan^{-1}\left(\frac{b}{a}\right) \quad (14)$$

and the symmetry of the crescent as

$$\tau \equiv 1 - \frac{\sqrt{a^2 + b^2}}{R_p - R_n}. \quad (15)$$

Finally, we normalize the complex visibilities so that they are equal to unity at zero baseline length, i.e., set

$$I_0 = \frac{1}{\pi R_p^2 - (1 - f)\pi R_n^2}. \quad (16)$$

Unless otherwise stated, we will set $f = 0$ hereafter.

The visibility amplitudes (10) of symmetric rings show the typical “ringing” of the Bessel functions, with minima at a series of characteristic baseline lengths. For a ring of infinitesimal width, the first minimum occurs at a baseline length of $k_0 \approx 0.383/R$. As the fractional width of the ring increases, the baseline length of the




Figure 3. The (u, v) coverage used in the synthetic data of the symmetric crescent model discussed in Section 3. The $(u - v)$ coverage is comparable to that of the 2017 EHT observations of M87. The dashed circle shows the baseline length of the location of the first minimum in visibility amplitude.

first minimum also increases. This introduces a correlation between the radii and fractional widths of rings that show the first visibility minimum at the same baseline length. An approximate fitting formula for this correlation is (see Event Horizon Telescope Collaboration et al. 2019f)

$$R(\psi) = R_0 \frac{1 - \psi/2}{1 - 0.48\psi + 0.11\psi^5}, \quad (17)$$

where R_0 is the radius of the infinitesimal ring with the same location of the first visibility minimum. Setting $0 \leq \psi \leq 1$, we get $0.64R_0 \leq R(\psi) \leq R_0$, for rings that show the first visibility minimum at the same baseline length.

Figure 2 shows the normalized visibility amplitudes of symmetric rings with different fractional widths but with radii that follow the correlation given by Equation (17). By construction, all these visibility amplitudes have the first visibility minimum at the same baseline length ($\approx 3.75 \text{ G}\lambda$). As the fractional width of the ring increases, the local maximum at $\approx 6 \text{ G}\lambda$ decreases and, for $\psi \gtrsim 0.5$, the location of the second minimum moves toward smaller baseline lengths.

In order to explore the performance of our MCMC algorithm and understand the qualitative aspects of fitting crescent models to EHT data, we generated a set of synthetic visibility amplitude and closure phase data using the crescent model (10) for a symmetric ring ($a = b = 0$) with parameters $R = 21 \mu\text{as}$, $\psi = 0.1$, $\phi = 0$, $\tau = 1$, and a flux of 1.3 Jy . We generated data on the (u, v) locations of the baselines of the 2017 April 5 EHT observations of M87 (see Figure 3) and, for the purposes of this exploration, introduced only thermal (Gaussian) errors with standard deviation equal to 10% of the model visibility amplitudes. We then used MARCH to fit the data with the same crescent model. The synthetic visibility amplitudes, the predictions of the best-fit model, and their residuals are shown in Figure 4. As discussed in Section 1, we do not consider here the effects of time-variable telescope gains but leave the exploration of those effects to a subsequent paper.

⁶ Note that, in Event Horizon Telescope Collaboration et al. (2019f), the fractional width is defined in terms of the mean diameter, i.e., $f_w \equiv (R_p - R_n)/(2R)$, so that $f_w = \psi/(2 - \psi)$.




Figure 4. Synthetic data (blue error bars) and best-fit model (red points) for a symmetric crescent image with $R = 21 \mu\text{as}$, $\psi = 0.1$, and the $(u - v)$ coverage of the 2017 EHT observations of M87 shown in Figure 3. For this test, the uncertainties are purely thermal (Gaussian) with a standard deviation equal to 10% of the model visibility amplitudes. The residuals are calculated as (model-data)/error.

Figure 5 shows the 68% and 95% credible intervals for the two primary parameters of the model: the radius R of the ring and its fractional width ψ . The values of these parameters and their correlation can be understood completely in terms of the location of the minimum in the visibility amplitudes that occurs at a baseline length of $\simeq 3.75 \text{ G}\lambda$ (see Figure 4). The red dashed line in Figure 5 shows the expected correlation between the fractional width and radius for all symmetric rings that have the first minimum of the visibility amplitude occurring at that baseline length. (Note that, for the purposes of this figure, which requires a relative accuracy of better than 1 part in 1000, we used the numerical solution for the correlation between fractional width ψ and ring radius R and not the analytic approximation (17).)

Marginalizing the posterior over all model parameters other than the radius results in $R = 20.98^{+0.02}_{-0.03}$, which corresponds to a fractional uncertainty of 0.1%. This is a surprisingly small uncertainty given that there are 152 visibility amplitude and 119 closure phase data points, the uncertainty of each complex visibility was set to 10%, and the model has only a handful of free parameters. For comparison, if this was a power-law model in visibility, i.e., $V(b) = Ab^c$, and there were $N = 152 + 119 = 271$ equidistant data points in $\log(b)$ between a minimum baseline $b_{\min} = 1 \text{ G}\lambda$ and a maximum baseline $b_{\max} = 8 \text{ G}\lambda$, all with fractional uncertainty σ_V/V ; then the uncertainty in the coefficient c of that model would be




Figure 5. The correlated 68% (dark blue) and 95% (light blue) credible intervals in the inferred radius R and fractional width ψ for the synthetic data shown in Figure 4. The dashed line shows the expected correlation between the two model parameters such that the visibility amplitudes show a minimum at a baseline length of $\simeq 3.75 \text{ G}\lambda$. It is the baseline length of this first minimum that almost single-handedly determines the most likely model parameters and their correlations.

$$\sigma_c \simeq \left(\frac{12}{N}\right)^{1/2} \frac{\sigma}{\ln(b_{\max}/b_{\min})} \simeq 1.0 \left(\frac{N}{271}\right)^{-1/2} \left(\frac{\sigma_V/V}{10\%}\right) \%. \quad (18)$$

i.e., an order of magnitude larger.

The fact that the uncertainties in the measurement of one of the key model parameters is substantially smaller than what one might expect suggests that the model predictions for at least some of the data points depend very strongly on that model parameter. This might be viewed naively as good news, since it allows for strong constraints to be placed on the model parameters with only marginal uncertainties in the data. However, it also implies that the inferred uncertainties in the model parameters rely extremely strongly on the ability of the model to predict accurately the salient features of the visibility data. The latter is unlikely to be true when we are using simple geometric models (e.g., crescent shapes) to fit visibility data that describe images of turbulent accretion flows around black holes.

4. The Role of Influential Data Points

The example discussed in the previous section demonstrates that a very small number of data points near the location of salient features in the visibility data obtained with an interferometer such as the EHT have an uneven degree of influence on the measurements and uncertainties of model parameters. This is true even for our example that was based on synthetic data generated by the actual crescent model, with the same fractional uncertainties for each data point. The effect gets exacerbated when the data points have different uncertainties (i.e., they are heteroscedastic), and the actual data are only broadly described by the simple geometric model that we try to fit to them, as is the case in real applications to EHT data.




Figure 6. A toy example of the effect of influential points on linear regression (based on the data set discussed in Millar & Stewart 2007). The black line is the best-fit linear model to all the data points. The red and green lines are the best-fit models after removing the red or the green point, respectively. The model parameters are given in Table 1. The red data point is an outlier, and its presence increases substantially the inferred uncertainties of the model parameters. The green data point, because of its location, introduces a significantly larger leverage in determining the model parameters than the remaining data points.

Quantifying the degree of influence of different data points on model fitting has been extensively explored in both the frequentist (see, e.g., Cook 1977, 1979) and the Bayesian frameworks (see, e.g., Millar & Stewart 2007). Both approaches generate very similar outcomes as they identify two types of influential data points. The first type are the outliers, i.e., the ones that deviate from model predictions by a degree that is statistically inconsistent with the uncertainties. The presence of such data points alters the inferred uncertainties in the model parameters. The second type is the subset of data points that depend substantially more strongly on the model parameters than the remaining data points and therefore exert uneven leverage in determining the best-fit values of the model parameters.

A toy example with both types of data points is shown in Figure 6, where a linear model is fit to the data set used in the Bayesian analysis of Millar & Stewart (2007); the best-fit line is shown in black, and its parameters are shown in Table 1. The red data point is a clear outlier. Removing it introduces only a small change to the parameters of the best-fit line (red curve); however, removing just that single point reduces the inferred uncertainties of the model parameters by as much as a factor of 2. The green data point appears at a large horizontal distance from the other data points and therefore exerts substantial leverage in determining the model parameters. Removing the green point changes substantially the best-fit model parameters, i.e., leads to a best-fit line (green curve) that is significantly flatter than when the full data set is used. This toy example demonstrates how removing even a single influential point changes substantially the parameters of the best-fit model and their uncertainties. For this reason, identifying such influential points for EHT data is critical in assessing the robustness of the results of model fits.

4.1. Identifying Influential Data Points

In order to quantify and formally measure the influence of each data point on the results of model fitting, we follow the

Table 1
Linear Model Parameters (Figure 6)

Data Set	Slope	y-intercept
Full Data Set	-1.09 ± 0.40	110.2 ± 6.5
Red Point Excluded	-1.19 ± 0.24	109.3 ± 4.0
Green Point Excluded	-0.61 ± 0.66	104.4 ± 9.2

Bayesian approach of Millar & Stewart (2007). The qualitative idea is to measure the change in the information content of the resulting posterior when we artificially broaden (or shrink) the likelihood that corresponds to each one of the data points. If this change is similar among all data points, this would imply that all of them have similar influences on the result. If, on the other hand, this change is substantially larger for some of the data points, we will conclude that the latter have high influence on the result.

We define the quasi-likelihood $\mathcal{L}(\text{data}|\vec{\theta}; w_i)$ as the likelihood of obtaining the data set from a particular set of model parameters $\vec{\theta}$ but with the likelihood of the i -th data point having been broadened by a factor w_i , i.e., (see Equation (9))

$$\mathcal{L}(\text{data}|\vec{\theta}; w_i) \equiv \mathcal{L}_i(\text{data}|\vec{\theta})^{w_i} \prod_{j \neq i} \mathcal{L}_j(\text{data}|\vec{\theta}). \quad (19)$$

If the likelihood of the i -th data point is a Gaussian with dispersion σ_i , the contribution of this data point to the total likelihood in Equation (19) is also a Gaussian but with a dispersion equal to $\sigma_i/w_i^{1/2}$. When $w_i \rightarrow 0$, the i -th data point does not contribute to the determination of the best-fit parameters.

Using this quasi-likelihood, we define a new quasi-posterior by effectively elevating the parameter w_i to a model parameter as (see Equation (8))

$$P(\vec{\theta}|\text{data}; w_i) = C P_{\text{pri}}(\vec{\theta}) \mathcal{L}(\text{data}|\vec{\theta}; w_i); \quad (20)$$

here we have incorporated the prior $P_{\text{pri}}(\vec{\theta})$ on the model parameters. When $w_i=1$, the above quasi-posterior is equal to the true posterior (8). In order to quantify the change in the information content of the posterior when we artificially change the likelihood of the i -th data point, i.e., when we set $w_i = 1 \pm \epsilon$, we use the Kullback–Leibler (KL) divergence

$$D_{KL,i}[P(\vec{\theta}|\text{data}; w_i), P(\vec{\theta}|\text{data}; 1)] = \int d\vec{\theta} P(\vec{\theta}|\text{data}; w_i) \log \left[\frac{P(\vec{\theta}|\text{data}; w_i)}{P(\vec{\theta}|\text{data}; 1)} \right]. \quad (21)$$

By definition, when $w_i = 1$, the KL divergence is equal to zero as is its first derivative with respect to w_i , under normal conditions (Millar & Stewart 2007). The second derivative of the KL divergence with respect to w_i offers the lowest-order nontrivial measure of the variation with w_i . After some simple manipulations, the second derivative becomes

$$\ddot{D}_{KL,i}(\vec{\theta}|\text{data}) = \int d\vec{\theta} P(\vec{\theta}|\text{data}) \left[\frac{\partial \log P(\vec{\theta}|\text{data}; w_i)}{\partial w_i} \bigg|_{w_i=1} \right]^2. \quad (22)$$

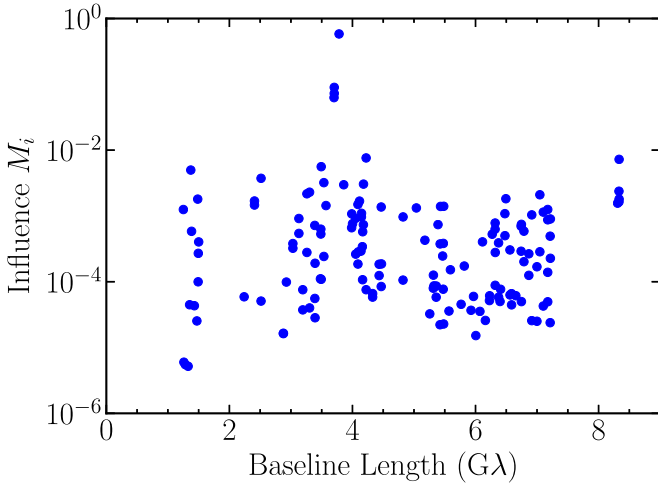


Figure 7. The measure of the influence M_i of each data point on the outcome of fitting a crescent model to the synthetic EHT data shown in Figure 4, plotted in logarithmic scale as a function of the baseline length that corresponds to each point. The three data points near the location of the minimum in visibility amplitude at $\approx 3.5 \text{ G}\lambda$ exert disproportionate influence on the result of the model fitting. Although not as large as the influence of the points near the first minimum, data points at the location of the second minimum at baseline lengths $\sim 8.5 \text{ G}\lambda$ also exert a disproportionately large influence.

This is the well-known result that the curvature of the KL divergence is equal to the Fisher information with respect to the relevant parameter (here w_i).

Using Equations (19), (20), and (22) we obtain

$$\ddot{D}_{KL,i}(\vec{\theta}|\text{data}) = \int d\vec{\theta} P(\vec{\theta}|\text{data}) [\log \mathcal{L}_i(\vec{\theta}|\text{data})]^2. \quad (23)$$

This is the measure of the influence of the i -th data point on the results, and it is equal to the variance of the log-likelihood of this data point across the space of the model parameters. Millar & Stewart (2007) suggest as a normalized measure of the influence of the i -th data point the ratio

$$M_i \equiv \frac{\int d\vec{\theta} P(\vec{\theta}|\text{data}) [\log \mathcal{L}_i(\vec{\theta}|\text{data})]^2}{\int d\vec{\theta} P(\vec{\theta}|\text{data}) [\log \mathcal{L}(\vec{\theta}|\text{data})]^2}, \quad (24)$$

where the denominator involves the combined log-likelihood for all data points. In general, the sum $\sum_i M_i$ will not be equal to unity, because the individual likelihoods may be correlated. Also note that Spiegelhalter et al. (2002) (see also Gelman et al. 2013) identified the denominator of Equation (24) as a measure of the complexity of the model given the data, i.e., of the effective number of degrees of freedom. Among a set of N data points, those with $M_i \gg 1/N$ will be the ones with disproportionate influence on the result of the model fitting.

4.2. Influential Data Points in Crescent Images

In this subsection, we apply the formalism we developed in Section 4.1 to the test problem of fitting crescent models to the interferometric data that we discussed in Section 3 (see Figures 4–5). Figure 7 shows the influence M_i (Equation (24)) of each of the $N = 152$ visibility amplitude data points on the result as a function of its baseline length. For the vast majority of data points, $M_i \lesssim 1/N = 1/152$. This is not true, however, for the three data points that lie at the first minimum of the visibility amplitude curve. The large values of M_i for these three data

points, which are several orders of magnitude larger than the others, demonstrate their disproportionate influence on the result. Indeed, the fact discussed in Section 3 that the best-fit values of and correlations between the inferred radius and fractional width of the crescent are determined almost entirely by the location of the first minimum in the visibility amplitude data is directly related to this disproportionate influence. Note that the influence values M_i of the corresponding closure-phase data do not single out any additional influential points for this synthetic data set.

In order to understand the origin of the disproportionate influence of the data points near the visibility minima on the outcome of the model comparison, we explore the sensitivity of the inference of model parameters on the location of data points of the u - v plane. As we did above, we focus on the visibility amplitude data, although similar considerations apply to the closure phase data as well. The sensitivity in the inference of a particular model parameter θ_i on the measurement at the location (u_j, v_j) of a given data point will depend on the uncertainty σ_j of the measurement (the higher the uncertainty, the weaker the sensitivity) and on the derivative of the model prediction with respect to θ_i , evaluated at the location of the data point, i.e.,

$$\sigma_j^{-1} \frac{\partial V}{\partial \theta_i} \bigg|_{u_j, v_j} = \left(\frac{\sigma_j}{V_j} \right)^{-1} V_j^{-1} \frac{\partial V}{\partial \theta_i} \bigg|_{u_j, v_j}. \quad (25)$$

This term can be recognized as the contribution of the j -th data point on the Cramér–Rao bound of the uncertainty in the inference of the i -model parameter. In the last expression, we have normalized the measurement uncertainty of each data point by the model predictions V_j in order to be consistent with our synthetic data that have constant fractional measurement uncertainties. With this normalization, a sensitivity value of unity implies that making, e.g., a change of 1% to the model parameter will cause a change of 1% to the model prediction.

Figure 8 shows the sensitivity of inferring the radius and fractional width of the symmetric ring model using data points with constant fractional uncertainties, as a function of the baseline length of each data point. As expected, both sensitivity curves show very strong peaks near the locations of the minima in the visibility amplitude. Moreover, the inference of the fractional width of the ring has a higher sensitivity on the measurements than the inference of its radius. The location and magnitude of the peaks in these two sensitivity curves match the locations and relative magnitudes of the influence M_i of the various data points shown in Figure 7.

To illustrate the potential biases introduced by the disproportionate influence of a small number of data points on the model parameters, we generate a new set of synthetic data from the ring model with the same parameters and u - v coverage as those discussed in Section 3 and shown in Figures 3–5. We add to this ring model, however, a single point source at the center of the image that carries a flux of only 0.02 Jy. This is about 1.5% of the total flux of the image and effectively adds a constant ≈ 0.02 to all visibility amplitudes, thus disproportionately affecting the complex visibilities of the data points near the minima in the visibility amplitude curve. The change introduced to the complex visibilities of the vast majority of the remaining data points is within their statistical uncertainties. We then attempt to fit this new synthetic data set with a crescent model that does not include this weak point source. The goal of this test is to simulate the outcome of trying




Figure 8. The sensitivity of predicted visibility amplitudes V for a symmetric ring model on two of the model parameters, the radius R (solid blue line) and the fractional width ψ (dashed red line) of the ring, as a function of the baseline length. For this example, the radius is set to $R = 21 \mu\text{as}$ and the fractional width to $\psi = 0.1$. The peak sensitivity occurs, as expected, at baseline lengths that correspond to the minima in the visibility amplitude.




Figure 9. As in Figure 5 but including, in magenta, the posteriors for the synthetic data in which a point source was added to the center of the ring image carrying a flux of 0.02 Jy; the parameters of the ring remain the same. The blue areas represent the credible intervals when there is no point source in the image and are the same as those in Figure 5. In both cases, darker colors correspond to the 68% intervals and lighter colors to 95%. Even though the point source carries $\sim 1.5\%$ of the total flux of the image and affects predominantly the complex visibilities of the handful of data points near the visibility minima, it modifies the inferred model parameters by a degree that is substantially larger than the statistical uncertainties. This is a direct consequence of the disproportionate influence of the data points near the visibility minima on the fit.

to fit a simple model to an image that has marginally more complex structures than the model allows.

The results for the radius and fractional width for this revised data set are shown in Figure 9. Even though the addition of the weak point source affects only a handful of interferometric data points, the high degree of influence of these data points causes a bias to the inferred model parameters to a degree that is substantially larger than the formal uncertainties from the fitting process. The correlated credible levels for the two parameters do not include the ring parameters that were assumed in the synthetic




Figure 10. The distribution of orientation angles (measured in degrees east of north) of the red- and blue-color coded EHT baselines in the 2017 observations of M87, which span a broad range of baseline lengths, as shown in Figure 3. The green- and magenta-color coded baselines are highly localized and trace out short tracks and, for this reason, are not included here. As can be also seen from the u - v map in the latter figure, most baselines fall primarily along the east-west orientation with position angles $\sim 80^\circ$ – 100° and along the north-south orientation with position angles $\sim 150^\circ$.




Figure 11. The posterior over the orientation of a crescent model after fitting to the synthetic data shown in Figure 4. The vertical blue and red hatched regions show the main orientations of the baselines used in generating the data displayed with the same colors in Figure 3. Even though the underlying model is a fully symmetric ring, the posterior is negligible at positional angles that are nearly perpendicular to the major orientations of baseline coverage in the u - v plane, biasing substantially the measurement.

data. They fall near but not on the expected correlation shown in the red dashed line, given the location of the minimum in the visibility amplitude. As expected from our discussion of the sensitivity of the various model parameters on the measurements, the bias introduced to the fractional width is particularly large (0.4 instead of 0.1 of the underlying image) and is substantially larger than the bias in the radius of the ring.

Note that, in the synthetic data set we used in these examples, we have set the uncertainties of all data points equal to the same fractional value (10%). This is not the case for the real EHT data, which are widely heteroscedastic; the signal-to-noise ratio for some ALMA baselines can be as large as 300, whereas for some other baselines involving smaller telescopes it is equal to only a




Figure 12. (Left) Snapshot image from a MAD GRMHD+GR Radiative Transfer simulation with parameters that are relevant to the M87 black hole; this image, which would have been obtained by a dense array with infinite resolution, is dominated by a crescent structure. (Middle) The same snapshot filtered with an $n = 2$ Butterworth filter with a cutoff at $15\text{ G}\lambda$. (Right) The best-fit crescent image to the synthetic EHT data generated from the snapshot on the left panel. For this snapshot that is dominated by a crescent structure, fitting a crescent model to synthetic data results in a model image that fairly closely matches the one used to generate the synthetic data.

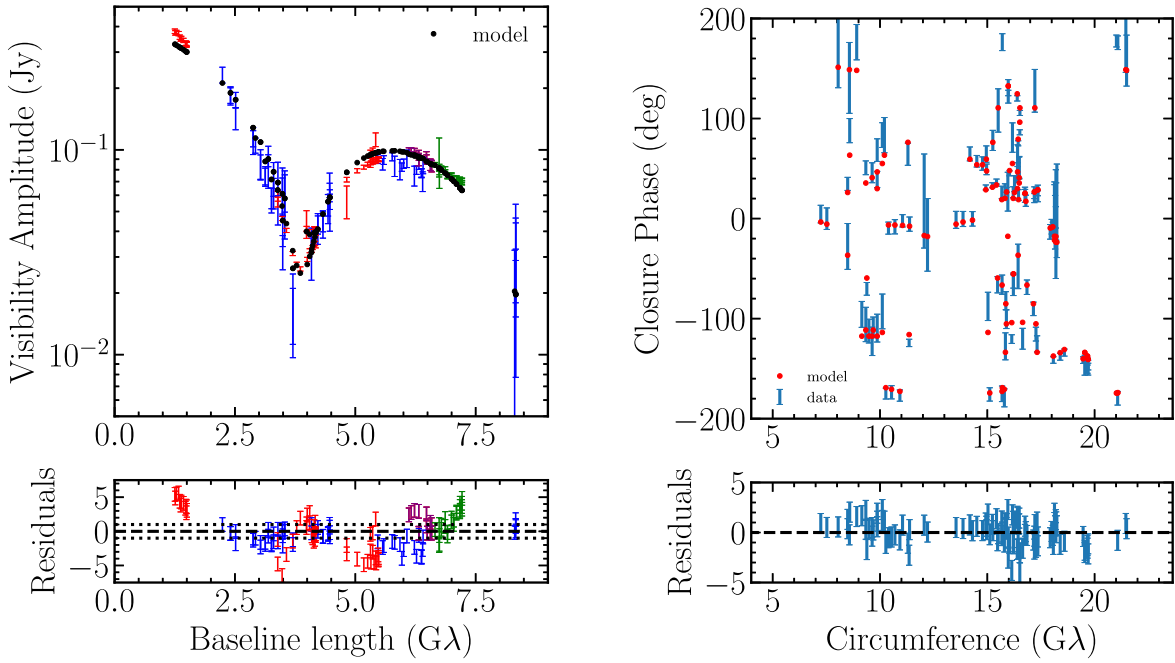


Figure 13. Synthetic data (shown with error bars, color coded in the left panel as in Figure 3 according to the orientations of the various baselines) and best-fit model (shown with black points) for fitting a crescent model to synthetic data generated from the image shown in Figure 12. The residuals are calculated as $(\text{model}-\text{data})/\text{error}$. The left panel shows the visibility amplitude as a function of baseline length, and the right panel shows the closure phase as a function of the circumference of the triangle of baselines along which the closure phase was measured. For this snapshot that is dominated by a crescent structure, fitting a crescent model to synthetic data results in only mild residuals.

few (see Event Horizon Telescope Collaboration et al. 2019c). In this case, the combination of the leverage of each data point (because of its location) and its signal-to-noise ratio will determine its influence: data points with high signal-to-noise ratios in locations of high leverage will have the largest influence. We will explore this more realistic case in Section 6.

5. The Effect of Sparse Coverage

A second important feature of the EHT data that affects the outcome of fitting theoretical models to them is the sparse coverage of the interferometric u - v plane (see, e.g., Figure 3). As Figure 10 shows, in the current form of the array, most of the EHT baselines lie primarily along the north-south and east-west orientations, with substantial gaps between them, at all

baseline lengths. This preferential orientation of the baselines can introduce strong biases in some of the inferred parameters of models that are fit to the EHT data, as we will now show.

Figure 11 shows the marginalized posterior over the position angle of the crescent model fit to the synthetic data introduced in Section 3 (see also Figures 3–5); for consistency with earlier literature, we report the position angle in degrees east of north, i.e., as $90^\circ - \phi$. Even though the underlying image in this example is a fully symmetric ring with no preferred orientation, the posterior over orientation angle is negligible for values that are not similar to those of the primary orientations of the EHT baselines. Note that, as discussed above, we have artificially fixed the uncertainties in all data points to 10%, in order to ensure that the heteroscedastic nature of the data is not responsible for a bias in orientation.




Figure 14. Posteriors of the various crescent model parameters used to fit the synthetic data from the image shown in Figure 12. The red dashed line in the panel showing the correlated posterior between the fractional width ψ and the radius of the crescent corresponds to Equation (17), which connects these two parameters for all crescents with the first minimum in the visibility amplitude occurring at the same baseline length.

The resulting bias is a direct consequence of the projection-slice theorem, which states that the cross section of a two-dimensional Fourier transform along a particular orientation is equal to the Fourier transform of the projection of the image along the same orientation. As discussed in Section 3, the analytic crescent model has a preferred orientation along which the cross section of the Fourier transform is independent of the degree of symmetry (τ) of the crescent. In other words, crescents that are mildly different than the symmetric ring (i.e., with a mild asymmetry) will lead to the same predictions as the symmetric model for the observations along the EHT baselines as long as their symmetry axes are lined up with the

orientations of the baselines. When marginalized over all other model parameters, the large volume in the space of models occupied by asymmetric crescents that are aligned with the EHT baselines and can reproduce the data within their uncertainties gives rise to a posterior over orientation ϕ that is aligned with the EHT baselines.

It is important to emphasize here that not only crescent models but also quasi-analytic (see, e.g., Broderick et al. 2009) and even GRMHD model images (see, e.g., Event Horizon Telescope Collaboration et al. 2019e) of black hole shadows have intrinsic symmetry axes dictated by the spin of the black hole and the brightness asymmetry introduced by Doppler

effects. As a result, similar considerations apply when fitting such models to the EHT data. Moreover, the synthetic data set we used here is characterized by the same fractional uncertainties for all data points, which is not realistic. In a more realistic situation, a complex image of a black hole will be observed with the EHT and will result in a heteroscedastic data set. In this case, it will be a complex combination of the preferred orientations of the baselines with the largest signal-to-noise ratio, of the underlying symmetry and orientation of the image, as well as of the symmetry properties of the model that will determine potential biases in the posterior over orientation angle.

6. Application to Synthetic EHT Data

In order to explore a more realistic situation that includes small-scale structure as well as realistic measurement uncertainties, we used two of the GRMHD+GR radiative transfer simulations calculated by Chan et al. (2015), but for parameters relevant to the M87 black hole. In particular, we used one simulation with the magnetically arrested disk (MAD) field structure (Narayan 2012). We set the mass of the black hole to $M = 6.5 \times 10^9 M_\odot$, the spin of the black hole to $a = 0.9$, the inclination of the observer to $i = 17^\circ$, the position angle of the black hole spin axis to 290° east of north, and the electron temperature to follow a plasma β -dependent prescription with $R_{\text{high}} = 20$ (see Event Horizon Telescope Collaboration et al. 2019e). We created images at 1.3 mm using the GPU-accelerated GRay algorithm (Chan et al. 2013) and calculated angular sizes in the sky for a distance of 16.8 Mpc. Finally, we generated synthetic EHT data from images generated with these simulations, with a u - v coverage and thermal noise levels that are identical to those of the 2017 EHT observations of M87 (Event Horizon Telescope Collaboration et al. 2019c; see Figure 3 for the u - v coverage).

In the following subsection, we discuss a number of cases that correspond to different morphologies of the underlying images.

6.1. Images with Dominant Crescent Structures

The low inferred inclination of the black hole spin in M87 substantially reduces the brightness asymmetry that is caused by the Doppler effect (compare, e.g., with the case of high-inclination simulations for Sgr A* reported in Chan et al. 2015). As a result, if the emission is highly localized near the black hole horizon, the resulting image will be dominated by a crescent structure with mild asymmetry. An example of such a snapshot from our MAD simulation is shown in the left panel of Figure 12. The north-south brightness asymmetry in this snapshot is 2:1, similar to what is inferred in the M87 image. The middle panel of the same figure shows the same snapshot but filtered with a $n = 2$ Butterworth filter and a cutoff baseline of 15 G λ . This filter and choice of parameters has been shown to suppress structures that are not accessible to the EHT array without, at the same time, affecting the interferometric observables within the range of baseline lengths and formal uncertainties of the EHT (see Psaltis et al. 2020).

Figure 13 shows the synthetic data for visibility amplitudes and closure phases generated from this image. For the assumed parameters of the black hole, the radius of the shadow is equal to $18.7 \mu\text{as}$ and, for an image dominated by an infinitesimal bright ring surrounding the shadow, the first minimum in the




Figure 15. The influence M_i for each of the synthetic data from the image shown in Figure 12. The data points are color coded as in Figure 3 to show the orientations of the various baselines. Filled circles denote points that are fit well by the crescent model, whereas open circles denote points that are not (i.e., the difference between the measurement and the model prediction is less than or more than 2.5 times the measurement uncertainty, respectively). The horizontal dashed line corresponds to $1/N$, where N is the total number of data points and separates the most influential (above the line) from the least influential (below the line) points. The influential data points at ~ 4 G λ are fit well by the model and have high leverage because they are near the location of the minimum in the visibility amplitude. The other data points above the dashed line are influential because they are outliers.

visibility amplitude is expected at a baseline length of $\simeq 4.2$ G λ . The synthetic data show instead a visibility amplitude minimum at a baseline length of $\simeq 3.8$ G λ , which corresponds to a radius of the emitting region of $\simeq 20.5 \mu\text{as}$, signifying a broader emission region.

Figure 13 also shows the best-fit crescent model that we obtained after running MARCH for 50 million steps and 40 replica temperatures. For the purposes of this test, we allowed for the crescent image to have softer edges by multiplying the complex visibilities of the model (Equation (10)) with a Gaussian function, which is equivalent to convolving the crescent image by the corresponding Gaussian kernel. The right panel of Figure 12 shows the image that corresponds to the best-fit model parameters. Clearly, when the underlying image is dominated by a crescent structure, then fitting a simple crescent model to synthetic EHT data results in only small residuals and a model image that fairly closely matches the properties of the one used to generate the synthetic data.

Figure 14 shows the posteriors of the model parameters. The very high signal-to-noise ratio of the data, which exceeds 100 for some of the data points, together with the small number of model parameters, leads to highly constrained values for the latter. Large correlations can be seen primarily between two pairs of model parameters. The posteriors in the fractional width ψ and the radius R of the crescent are correlated according to Equation (17), as their values are primarily determined by the baseline length of the first minimum in the visibility amplitude.⁷ At the same time, the posteriors in the FWHM of the broadening Gaussian kernel and the fractional width of the crescent are anticorrelated because the same effective broadening of the crescent can be achieved by either

⁷ The definition of the fractional width as $\psi \sim 1/R$ could introduce an anticorrelation, but this explains neither the shape nor the range of values of the correlated posteriors.




Figure 16. (Upper Left) A second snapshot image from a MAD GRMHD+GR radiative transfer simulation with parameters that are relevant to the M87 black hole. (Upper Right) The same snapshot filtered with an $n = 2$ Butterworth filter with a cutoff at $15 \text{ G}\lambda$. (Lower Left) The best-fit crescent image to the synthetic EHT data generated from the snapshot on the upper left panel; the width and orientation of the crescent image do not match the ones of the image used to generate the synthetic data. (Lower Right) The best-fit crescent image, when an additional elliptical Gaussian component is added to the model. The addition of this component leads to a more reasonable crescent width and an orientation that resembles that of the original image. However, the width and orientation of the Gaussian component do not match any particular structure in the original image, and its minor axis is lined up with one of the prominent orientations of EHT baselines.

increasing the fractional width of the crescent or by broadening it by a wider Gaussian kernel.

Figure 15 shows the relative influence that each of the synthetic data points with realistic uncertainties exerts on the inference of the model parameters, calculated using Equation (24). In order to separate the two types of influential data points (i.e., those with intrinsically large leverage versus the outliers that are not well described by the model), we have used filled and open circles to denote the data points for which the difference between the measurement and the model prediction is less than or more than 2.5 times the measurement uncertainty, respectively. We have also color coded the data points as in Figure 3 to show the orientations of the various baselines. As expected from the discussion in Section 4, the data points near the minimum of the visibility amplitude at $\sim 4 \text{ G}\lambda$ are described well by the model and have high leverage in determining the model parameters. In contrast, a number of data points exist at various baselines that are not described well by the model (i.e., they have large residuals) but exert substantial influence in determining the model parameters because they are outliers. In this case of a realistic image, which is not described adequately by a simple crescent model,

and of heteroscedastic data, the posteriors of the model parameters are determined by a combination of the influence of data points of high leverage and of the outliers.

The fact that an image appears to be dominated by a crescent structure does not mean that the resulting synthetic data will always be well described by a crescent model. This is demonstrated in Figure 16, which shows a snapshot with what appears to be a dominant crescent structure. However, fitting the resulting synthetic data with a crescent model results in substantial residuals (see left panel of Figure 17), and the corresponding best-fit image does not match the width or the orientation of the original image (see lower left panel of Figure 16).

It is tempting to try to improve the quality of the fit by adding elliptical Gaussian components to the crescent model. Adding even a single such elliptical Gaussian increases the number of model parameters from 6 (see Figure 14) to 12, with the additional six parameters being the centroid positions along the two axes of the Gaussian component, its major and minor axes, its position angle, and its flux. This large increase in the number of model parameters indeed leads to substantially smaller residuals, as can be seen in the right panel of Figure 17.

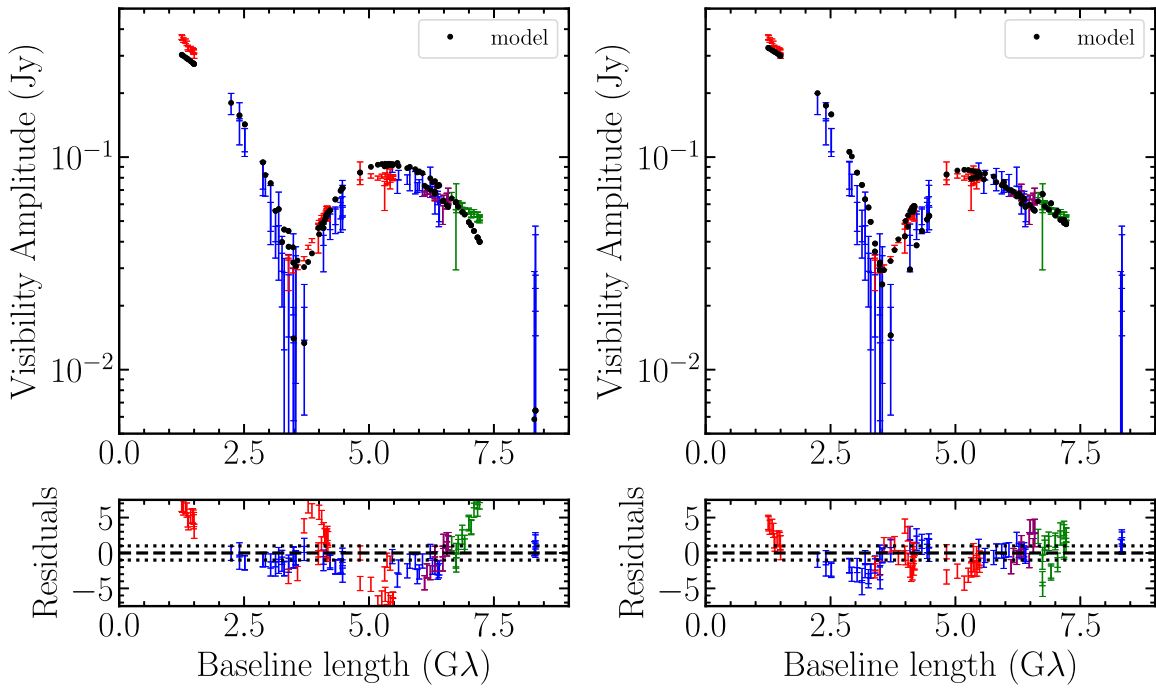


Figure 17. Same as Figure 13 but for the synthetic data from the snapshot shown in the upper left panel of Figure 16. (*Left*) Fit to the synthetic data with a simple broadened crescent model; substantial residuals exist, especially for the red and green baselines, and are the reason behind the biased width and orientation of the best-fit image (see bottom left panel of Figure 16). (*Right*) Fit to the same synthetic data but with a broadened crescent model to which an asymmetric Gaussian component was added; the addition of this component substantially reduces the residuals.

Comparing the best-fit image to the original snapshot from the simulation (see lower right panel of Figure 16) reveals that the additional Gaussian component seems to be adding flux at the lower-left quadrant of the image that has additional structure. However, it is extremely narrow, and its minor axis in the image domain (i.e., the major axis in the visibility domain) is aligned with the orientation of the green baselines. This is not surprising, given that the main utility of the additional elliptical Gaussian component in this particular example has been to reduce the residuals along the red and green color-coded baselines, which have the highest signal-to-noise ratios, without affecting substantially the blue color-coded ones, which have lower signal-to-noise ratios and were well described by the simple crescent model (see Figure 3 for the color code of the baselines). The size, asymmetry, and orientation of the additional Gaussian component are simply artifacts of the sparse coverage of the EHT array and the orientations of its baselines.

Nevertheless, even though the properties of the additional Gaussian component are not reliable, its presence allowed the best-fit crescent model to have a width and an orientation that more closely matches that of the original snapshot compared to the best-fit model without this extra component. In both examples described in this subsection, the dominant crescent shapes in the images caused deep minima in the visibility amplitudes at baselines $\simeq 3.8\text{--}4\text{ G}\lambda$ and, for this reason, the inferred sizes of the crescents are comparable to the true values. It is worth emphasizing here, however, that in fitting real data, there is no guarantee that adding extra components will always bring the gross features of the reconstructed images to better agreement with those of the underlying image. This is especially true when significant salient features, such as deep visibility minima, are not present in the observed interferometric data.

6.2. Images with Subdominant Crescent Structures

Figure 18 shows a third example of a snapshot from a MAD simulation with characteristics that are frequently encountered. In this example, there is substantial structure beyond the simple crescent shape that persists even after we filtered the image with the Butterworth filter in order to suppress visually the structures that are not accessible to the EHT data.

The synthetic data from this snapshot (Figure 19) indeed do not show the characteristic ringing structure with pronounced deep minima and the overlapping secondary maxima of the crescent shapes along all orientations. It is therefore expected that fitting a simple broadened crescent model to the synthetic data (left panel of Figure 19) leaves substantial residuals at all baselines. The resulting best-fit image (lower-left panel of Figure 18) does not match any of the structures in the image from which the synthetic data were generated. Moreover, it is not obvious whether the addition of other simple structures would improve the agreement with the underlying image.

Adding an additional Gaussian component to the model does improve the quality of the fit and reduces the residuals, especially along the red- and green-baselines that have the data with the highest signal-to-noise ratios (right panel of Figure 19). The best-fit image (lower right panel of Figure 18), however, shows structures that again do not match those of the image from which the synthetic data were generated. The width of the crescent shape is reasonable but neither its orientation nor its size are correct. Moreover, as in Figure 16 in the last example discussed in the previous subsection, the minor axis of the additional Gaussian component in the image domain (i.e., the major axis in the visibility domain) is aligned with the orientation of the green baselines.




Figure 18. (*Upper Left*) A third snapshot image from a MAD GRMHD+GR radiative transfer simulation with parameters that are relevant to the M87 black hole. This snapshot has substantial structure beyond the crescent shape. (*Upper Right*) The same snapshot filtered with an $n = 2$ Butterworth filter with a cutoff at $15 \text{ G}\lambda$. (*Lower Left*) The best-fit crescent image to the synthetic EHT data generated from the snapshot on the upper left panel; the width and orientation of the crescent image do not match the one used to generate the synthetic data. (*Lower Right*) The best-fit crescent image, when an additional elliptical Gaussian component is added to the model. The addition of this component leads to a crescent width (but not size or orientation) that resembles that of the original image. However, this additional component does not match in detail any particular structure in the original image and, as in Figure 16, its minor axis is again lined up with one of the prominent orientations of EHT baselines.

7. Conclusions

In this paper, we introduced MARCH, a new Markov Chain Monte Carlo algorithm with parallel tempering for fitting theoretical models to the interferometric data from the EHT. We then used this algorithm to explore the biases in the inferred parameters of crescent models that we fit to synthetic EHT data based on images from GRMHD simulations. We explored, in particular, biases that are introduced by the uneven influence of data points, the heteroscedastic nature of the EHT data, and the sparse coverage of the interferometric u - v plane with the EHT baselines.

When the images are dominated by strong crescent structures, we find that there are a number of salient features in the interferometric data, such as a deep minimum in the visibility amplitude at a baseline length determined primarily by the size of the structure. Data obtained in the vicinity of such salient features exert a disproportionate influence on the values of the inferred model parameters, artificially reducing the widths of their posteriors.

When the underlying images have structures that are not well characterized by crescent shapes, fitting crescent models to

synthetic EHT data reveals substantial residuals, as expected. Attempting to ameliorate these residuals by adding compact (e.g., Gaussian) components to the models may improve the quality of the fit and reduce the residuals. However, the properties of the additional components are strongly biased by the locations and orientations of the EHT baselines and are not trustworthy. Such Gaussian components with similarly large aspect ratios and orientations were deemed necessary to fit the actual M87 data observed in 2017 with the EHT (see Figure 5 of Event Horizon Telescope Collaboration et al. 2019f).

When the underlying image has a high degree of symmetry, the sparse baseline coverage of the EHT array often biases even the inferred overall orientation of the underlying image to align with the orientations of the primary baselines (see, e.g., Figure 11). If we were to match the features of a crescent model with those of a snapshot from a GRMHD simulation, we would identify the “orientation” of the crescent structure (which measures the orientation of displacement between the outermost, positive disk and the innermost, negative one) with the perpendicular orientation to that of the black hole spin. Figure 11 shows that the orientations of a crescent model that

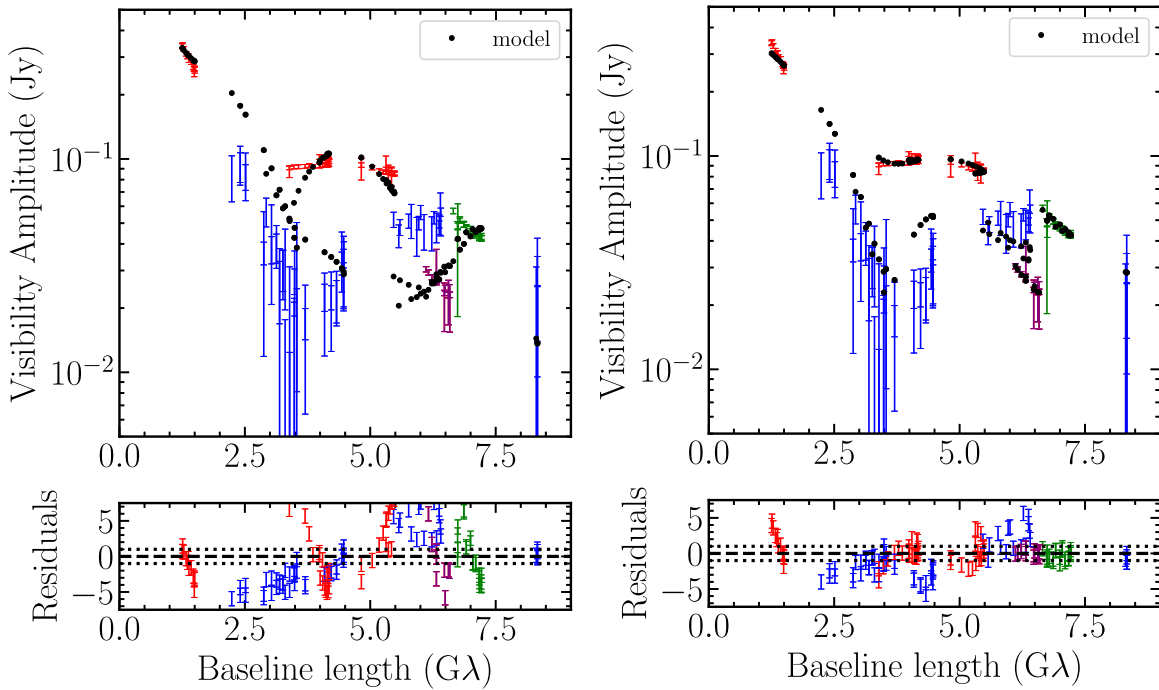


Figure 19. Same as Figure 13 but for the synthetic data from the snapshot shown in the upper left panel of Figure 18. (Left) Fit to the synthetic data with a simple broadened crescent model; substantial residuals exist, especially for the red and green baselines, and are the reason behind the biased width and orientation of the best-fit image (see bottom left panel of Figure 18). (Right) Fit to the same synthetic data but with a broadened crescent model to which an asymmetric Gaussian component was added; the addition of this component substantially reduces the residuals.

was fit to synthetic EHT data of a symmetric structure appear to get biased toward a position angle of $\simeq 150^\circ$ east of north. This would imply a bias in the inferred orientation of the black hole spin toward an orientation $\simeq 150 + 90 = 240^\circ$ east of north, which is very similar to the value with the highest posterior inferred from fitting crescent models to the 2017 EHT data and casts doubt on that inference (see, e.g., Figure 9 of Event Horizon Telescope Collaboration et al. 2019e).

In all cases, the robust detection of salient features leads to the determination of the characteristic sizes of crescent images that are within a small fractional distance from (albeit often statistically inconsistent with) the true sizes of the underlying images. When such salient features are detected, as was the case with the 2017 EHT data on M87, the size of the black hole shadow and, hence, of the black hole mass is not significantly biased. In order to obtain accurate measurements of additional aspects of the image, such as its orientation, asymmetry, and presence of additional structures, requires a more complete coverage of the u - v space with the addition of more stations in the EHT array. A number of such stations have been introduced since the 2017 observations, such as the Greenland telescope, the 12 m telescope on Kitt Peak, and the Northern Extended Millimeter Array in the French Alps, which will help substantially in alleviating the biases introduced by the sparse baseline coverage of the current EHT.

We thank K. Satapathy and T. Trent for useful discussions. This work was supported in part by NSF PIRE grant 1743747 and NSF grant AST 1715061. L.M. acknowledges support from an NSF Astronomy and Astrophysics Postdoctoral Fellowship under award no. AST-1903847. C.R. acknowledges support from NSF Graduate Research Fellowship Program Grant DGE-1746060. All ray-tracing calculations were

performed with the El Gato GPU cluster at the University of Arizona that is funded by NSF award 1228509.

ORCID iDs

Dimitrios Psaltis <https://orcid.org/0000-0003-1035-3240>
 Feryal Özel <https://orcid.org/0000-0003-4413-1523>
 Lia Medeiros <https://orcid.org/0000-0003-2342-6728>
 Pierre Christian <https://orcid.org/0000-0001-6820-9941>
 Junhan Kim <https://orcid.org/0000-0002-4274-9373>
 Chi-kwan Chan <https://orcid.org/0000-0001-6337-6126>
 Carolyn A. Raithel <https://orcid.org/0000-0002-1798-6668>
 Dan Marrone <https://orcid.org/0000-0002-2367-1080>
 Tod R. Lauer <https://orcid.org/0000-0003-3234-7247>

References

- Akiyama, K., Ikeda, S., Pleau, M., et al. 2017a, *AJ*, **153**, 159
- Akiyama, K., Kuramochi, K., Ikeda, S., et al. 2017b, *ApJ*, **838**, 1
- Benkevitch, L., Akiyama, K., Lu, R., Doeleman, S., & Fish, V. 2016, *arXiv:1609.00055*
- Berger, J. P., Malbet, F., Baron, F., et al. 2012, *A&ARv*, **20**, 53
- Blackburn, L., Pesce, D. W., Johnson, M. D., et al. 2020, *ApJ*, **894**, 31
- Broderick, A. E., Fish, V. L., Doeleman, S. S., & Loeb, A. 2009, *ApJ*, **697**, 45
- Broderick, A. E., Fish, V. L., Johnson, M. D., et al. 2016, *ApJ*, **820**, 137
- 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
- Chan, C.-k., Psaltis, D., & Özel, F. 2013, *ApJ*, **777**, 13
- Chan, C.-K., Psaltis, D., Özel, F., Narayan, R., & Sadowski, A. 2015, *ApJ*, **799**, 1
- Christian, P., & Psaltis, D. 2020, *ApJ*, **159**, 226
- Clark, B. G. 1980, *A&A*, **89**, 377
- Cook, R. D. 1977, *Technometrics*, **19**, 15
- Cook, R. D. 1979, *J. Amer. Stat. Assoc.*, **73**, 365
- Dexter, J., Agol, E., & Fragile, P. C. 2009, *ApJL*, **703**, L142
- Earl, D. J., & Deem, M. W. 2005, *PCCP*, **7**, 3910
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019a, *ApJL*, **875**, L1

- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019b, [ApJL](#), **875**, L2
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019c, [ApJL](#), **875**, L3
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019d, [ApJL](#), **875**, L4
- 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
- Falcke, H., Melia, F., & Agol, E. 2000, [ApJL](#), **528**, L13
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, [PASP](#), **125**, 306
- Gelman, A., Hwang, J., & Vehtari, A. 2013, arXiv:1307.5928
- Högbom, J. A. 1974, [A&AS](#), **15**, 417
- Honma, M., Akiyama, K., Uemura, M., & Ikeda, S. 2014, [PASJ](#), **66**, 95
- Jaroszynski, M., & Kurpiewski, A. 1997, [A&A](#), **326**, 419
- Jennison, R. C. 1958, [MNRAS](#), **118**, 276
- Kamruddin, A. B., & Dexter, J. 2013, [MNRAS](#), **434**, 765
- Kim, J., Marrone, D. P., Chan, C.-K., et al. 2016, [ApJ](#), **832**, 156
- Lu, R.-S., Krichbaum, T. P., Roy, A. L., et al. 2018, [ApJ](#), **859**, 60
- Medeiros, L., Chan, C.-k., Özel, F., et al. 2017, [ApJ](#), **844**, 35
- Medeiros, L., Chan, C.-k., Özel, F., et al. 2018, [ApJ](#), **856**, 163
- Millar, R. B., & Stewart, W. S. 2007, [Bayes. Anal.](#), **2**, 365
- Mościbrodzka, M., Falcke, H., Shiokawa, H., & Gammie, C. F. 2014, [A&A](#), **570**, A7
- Mościbrodzka, M., Gammie, C. F., Dolence, J. C., Shiokawa, H., & Leung, P. K. 2009, [ApJ](#), **706**, 497
- Narayan, R., & Nityananda, R. 1986, [ARA&A](#), **24**, 127
- Narayan, R., Śa dowski, A., Penna, R. F., & Kulkarni, A. K. 2012, [MNRAS](#), **426**, 3241
- Özel, F., Psaltis, D., & Narayan, R. 2000, [ApJ](#), **541**, 234
- Psaltis, D., Medeiros, L., & Lauer, T. R. 2020, arXiv:2004.06210
- Psaltis, D., Özel, F., Chan, C.-K., & Marrone, D. P. 2015, [ApJ](#), **814**, 115
- Roelofs, F., Johnson, M. D., Shiokawa, H., Doeleman, S. S., & Falcke, H. 2017, [ApJ](#), **847**, 55
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. 2002, [Journal Royal Stat. Soc. B](#), **64**, 583
- Thiébaud, E. 2009, [NewAR](#), **53**, 312
- Thompson, A. R., Moran, J. M., Swenson, G. W. J., et al. 2017, *Interferometry and Synthesis in Radio Astronomy* (3rd ed.; Berlin: Springer)