



Dynamical Imaging with Interferometry

Michael D. Johnson¹, Katherine L. Bouman^{1,2}, Lindy Blackburn¹, Andrew A. Chael¹, Julian Rosen³, Hotaka Shiokawa¹, Freek Roelofs⁴, Kazunori Akiyama⁵, Vincent L. Fish⁵, and Sheperd S. Doeleman¹

¹Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA; mjohnson@cfa.harvard.edu

²Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory, 32 Vassar Street, Cambridge, MA 02139, USA

³University of Michigan, Department of Mathematics, 530 Church Street, Ann Arbor, MI 48109-1043, USA

⁴Department of Astrophysics, Institute for Mathematics, Astrophysics and Particle Physics, Radboud University, P.O. Box 9010, 6500 GL Nijmegen, The Netherlands

⁵Massachusetts Institute of Technology, Haystack Observatory, Route 40, Westford, MA 01886, USA

Received 2017 July 30; revised 2017 October 27; accepted 2017 October 31; published 2017 November 30

Abstract

By linking widely separated radio dishes, the technique of very long baseline interferometry (VLBI) can greatly enhance angular resolution in radio astronomy. However, at any given moment, a VLBI array only sparsely samples the information necessary to form an image. Conventional imaging techniques partially overcome this limitation by making the assumption that the observed cosmic source structure does not evolve over the duration of an observation, which enables VLBI networks to accumulate information as Earth rotates and changes the projected array geometry. Although this assumption is appropriate for nearly all VLBI, it is almost certainly violated for submillimeter observations of the Galactic center supermassive black hole, Sagittarius A* (Sgr A*), which has a gravitational timescale of only ~ 20 s and exhibits intrahour variability. To address this challenge, we develop several techniques to reconstruct dynamical images (“movies”) from interferometric data. Our techniques are applicable to both single-epoch and multiepoch variability studies, and they are suitable for exploring many different physical processes including flaring regions, stable images with small time-dependent perturbations, steady accretion dynamics, or kinematics of relativistic jets. Moreover, dynamical imaging can be used to estimate time-averaged images from time-variable data, eliminating many spurious image artifacts that arise when using standard imaging methods. We demonstrate the effectiveness of our techniques using synthetic observations of simulated black hole systems and 7 mm Very Long Baseline Array observations of M87, and we show that dynamical imaging is feasible for Event Horizon Telescope observations of Sgr A*.

Key words: accretion, accretion disks – black hole physics – Galaxy: center – techniques: high angular resolution – techniques: interferometric

1. Introduction

Very long baseline interferometry (VLBI) provides exceptional angular resolution but only sparsely samples the Fourier components of an image. A powerful technique to enhance interferometric imaging utilizes Earth’s rotation—as Earth rotates, each baseline connecting an antenna pair tracks through, and samples, a range of image Fourier components (see, e.g., Ryle 1962; Kellermann & Moran 2001; Thompson et al. 2017). In its conventional implementation, Earth-rotation synthesis imaging assumes that the source being imaged is static over the observing duration (typically \sim hours). This assumption is reasonable for nearly all astrophysical sources of interest, although a few sources have shown detectable structural changes within a single observation (e.g., Reid et al. 2014), most commonly through rapid swings of polarization angle (e.g., Gabuzda et al. 2000).

One notable case for which the static-source assumption is likely to fail is the Galactic center supermassive black hole, Sagittarius A* (Sgr A*). Because Sgr A* has a mass of approximately $M \approx 4 \times 10^6 M_{\odot}$ (Ghez et al. 2008; Gillessen et al. 2009), its gravitational timescale is only $GM/c^3 \approx 20$ s and its innermost stable circular prograde orbits have periods of only 4–30 minutes, depending on the black hole spin (Bardeen et al. 1972). In terms of observed variability, Sgr A* regularly flares with \sim hour timescales (e.g., Marrone et al. 2008; Yusef-Zadeh et al. 2009; Brinkerink et al. 2015), and its polarization shows intense variations on similar timescales

(e.g., Marrone et al. 2006; Eckart et al. 2006; Zamaninasab et al. 2010; Johnson et al. 2015a).

Until recently, limitations from optical depth and interstellar scattering have prevented studies of rapid structural variability of Sgr A* using VLBI (e.g., Bower et al. 2006; Lu et al. 2011). However, the advent of 1.3 mm VLBI with the Event Horizon Telescope (EHT) will soon permit imaging Sgr A* on spatial scales for which intrinsic variability may be significant (Doeleman et al. 2009a). Pronounced variability with accompanying structural change has already been seen in the polarization of Sgr A* with the EHT (Johnson et al. 2015a), although the total-intensity structure of Sgr A* has remained comparatively stable (Doeleman et al. 2008; Fish et al. 2011, 2016b). In addition to these observations and the short characteristic timescales of Sgr A*, numerical simulations suggest that conventional VLBI imaging techniques will be inapplicable for EHT observations of Sgr A* (see Figure 1 and, e.g., Broderick & Loeb 2006; Doeleman et al. 2009a; Dexter et al. 2010; Johnson et al. 2015b; Kim et al. 2016; Lu et al. 2016; Medeiros et al. 2016; Gold et al. 2017; Medeiros et al. 2017; Roelofs et al. 2017).

Nevertheless, interferometry provides capabilities to study rapidly varying structures. For example, using simulated observations of a “hot spot” orbiting Sgr A* (Broderick & Loeb 2005, 2006), Doeleman et al. (2009b) and Fish et al. (2009) demonstrated that robust VLBI observables can sensitively detect periodicities associated with these hot spots. More generally, Johnson et al. (2014) showed that polarimetric VLBI

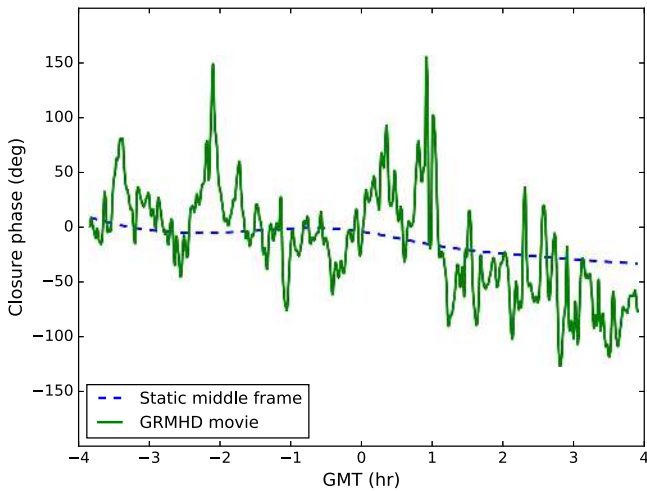


Figure 1. VLBI phase measurement, the closure phase, over time for the SPT-LMT-ALMA triplet of EHT antennas using mock observations of a time-variable general relativistic magnetohydrodynamic (GRMHD) simulation of Sgr A* (Shiokawa 2013). The phases in blue show the array response to a single static frame in the GRMHD movie. The mild variations for this case reflect Earth’s rotation and show the modest additional information available to Earth-rotation synthesis. The phases in green trace the array response to the full simulation, showing that the phase variations are dominated by intrinsic variability of the source. See Roelofs et al. (2017) for additional examples and discussion.

enables microarcsecond astrometry of compact flaring structures, even for faint, nonperiodic flares. Even conventional Earth-rotation synthesis utilizes time-variable Fourier sampling to enhance imaging, and Johnson et al. (2015b) argued that intrinsic variability of a source can be exploited in the same way if the source variability can be modeled (see, e.g., Sault et al. 1997). As Figure 1 shows, while intrinsic variability of Sgr A* may readily break the static-source assumption of conventional imaging, it also provides a rich source of information about the intrinsic variability.

In this paper, we develop techniques to reconstruct dynamical images (i.e., movies) from interferometric data. By accommodating intrinsic variability in the imaging procedure, we can study the dynamical activity of a source while avoiding spurious image features from the static-source assumption of conventional imaging algorithms. In a related approach, Lu et al. (2016) have recently developed a prescription for scaling, averaging, and smoothing interferometric visibilities; the processed visibilities can then be imaged using standard VLBI imaging techniques.⁶ They show that the resulting images are good approximations of the time-averaged image, especially when data from multiple observing epochs can be combined. Our focus is instead on reconstructing dynamical images of the time-variable source, while obtaining reliable approximations of the time-averaged image as a by-product.

Our work is a generalization of the standard regularized minimization approach to VLBI imaging, which includes approaches such as the maximum entropy method (MEM; see, e.g., Narayan & Nityananda 1986) and many other regularization functions. This approach, while not a strictly probabilistic model, can be motivated through a Bayesian framework wherein the minimization corresponds to maximizing the log

⁶ This prescription is motivated by linearity of the Fourier transform: complex visibilities of the time-averaged image are equal to time-averaged visibilities of a variable image. See Shiokawa et al. (2017) for generalized time-domain filtering of images.

posterior probability of a reconstructed image. In a separate paper, we explore an alternative approach to dynamical imaging via a modified Hidden Markov Model with a multivariate Gaussian image prior, and we derive closed-form expressions for both the maximum a posteriori image and its uncertainties (Bouman et al. 2017).

We begin, in Section 2, by reviewing the standard framework and procedure for VLBI imaging through regularized minimization, and we then generalize this framework to accommodate dynamical imaging. Next, in Section 3, we develop three regularizers that can be used for dynamical imaging for a variety of scenarios. In Section 4, we discuss using dynamical imaging for temporal interpolation. In Section 5, we show example results using simulated data and Very Long Baseline Array (VLBA) observations of M87, and in Section 6 we summarize our main results and conclusions.

2. Fundamentals of Interferometric Imaging

2.1. Interferometric Visibilities

Each baseline joining two sites in an interferometer samples complex visibilities. By the van Cittert–Zernike theorem, these visibilities, $V(\mathbf{u})$, are related to the brightness distribution on the sky $I(\mathbf{x})$ via a Fourier transform (Thompson et al. 2017):

$$V(\mathbf{u}) = \int d^2\mathbf{x} I(\mathbf{x}) e^{-2\pi i \mathbf{u} \cdot \mathbf{x}}, \quad (1)$$

where \mathbf{x} is an angular coordinate on the sky, in radians, and $\mathbf{u} \equiv \{u, v\}$ is the dimensionless baseline vector, in wavelengths, projected orthogonal to the line of sight.

Interferometry uses a set of measured visibilities $\{V_i\}$ to estimate the unknown sky image $I(\mathbf{x})$, as we will discuss in Section 2.2. However, when the image is also a function of time, the sampled visibilities at a particular time only represent the corresponding, instantaneous image. In this case, a series of images can be reconstructed if each utilizes only its simultaneous “snapshot” visibility coverage. With N_s participating sites with mutual visibility of the source, there are at most $N_s(N_s - 1)/2$ visibilities in the snapshot coverage, severely limiting the imaging capabilities when N_s is small (see, e.g., Figure 2).

2.2. Interferometric Imaging via Regularized Minimization

We will now review the standard prescription for VLBI imaging via regularized minimization. This prescription encompasses many common approaches to VLBI imaging, such as MEM (see, e.g., Frieden 1972; Cornwell & Evans 1985; Narayan & Nityananda 1986) and many variants (see, e.g., Thiébaud et al. 2013; Honma et al. 2014; Lu et al. 2014; Bouman et al. 2016; Chael et al. 2016; Fish et al. 2016a; Akiyama et al. 2017a, 2017b), but does not describe iterative deconvolution approaches such as CLEAN (Högbom 1974). The flexibility of the regularized minimization framework makes it ideal for sparse and heterogeneous arrays, such as the EHT, and also allows extensions to include, e.g., mitigating the image distortions caused by interstellar scattering (Johnson 2016).

To simplify our presentation, we will represent reconstructed images I as square $N \times N$ arrays, giving flux density per pixel. We denote a sequence of images by $\{I_j\}$, where j indexes the time for N_t different frames. In the following sections, we will generally treat each image as a vector of length N^2 rather than




Figure 2. Top: baseline coverage for the EHT observing Sgr A*. Baselines are colored by Greenwich Sidereal Time (GST) to indicate the snapshot u - v coverage at each time. Bottom: elevation of Sgr A* as a function of GST for each site, with a cutoff of 10° . Current EHT sites are the Atacama Large Millimeter/submillimeter Array (ALMA), the Large Millimeter Telescope (LMT), the Submillimeter Array (SMA), the Submillimeter Telescope (SMT), the Institut de Radioastronomie Millimétrique (IRAM) telescope on Pico Veleta (PV), the IRAM Plateau de Bure Interferometer (PdB), and the South Pole Telescope (SPT). Note that PdB did not participate in 2017 EHT observations.

an $N \times N$ matrix. Linear operators such as the Fourier transform relating images and interferometric visibilities, blurring via convolution, and discrete gradients are linear and can therefore be represented as $N^2 \times N^2$ matrix operators that act on these 1D image vectors (of course, elements of these operators depend on the 2D nature of the images).

Approaches such as MEM estimate the unknown source image \mathbf{I} by numerically minimizing an objective function, $J(\mathbf{I})$. J contains terms that express whether or not an image is consistent with the input VLBI data (a chi-squared term) and also contains terms that favor certain image attributes (such as smoothness or positivity through an entropy or other regularization term). The objective function then takes the form

$$J = \chi^2(\mathbf{I}, \mathbf{d}) - \alpha_S S(\mathbf{I}). \quad (2)$$

In this expression, $S(\mathbf{I})$ denotes the regularization function for the imaging (e.g., $S(\mathbf{I}) \equiv -\sum_{\ell,m} I_{\ell,m} \ln(I_{\ell,m})$ is commonly used for MEM), and χ^2 represents a chi-squared for whatever data products \mathbf{d} are used as part of the imaging. α_S is a “hyperparameter” that controls the relative weighting of the

entropy and data terms. The hyperparameter can be adjusted manually or automatically to yield the expected χ^2 for a satisfactory image (e.g., Cornwell & Evans 1985), or it can be estimated via cross-validation, wherein the data are divided into training and testing sets and the hyperparameters are chosen so that images reconstructed using the training set are compatible with the measurements and errors of the testing set (see Akiyama et al. 2017b). From a probabilistic perspective, the χ^2 term in Equation (2) corresponds to a log-likelihood, while the regularization term corresponds to a log prior distribution of the reconstructed image.

2.3. General Prescription for Dynamical Imaging

We now extend this framework and notation to dynamical imaging. In this case, the imaging problem is to simultaneously reconstruct N_t different frames $\{\mathbf{I}_j\}$. Each frame has an associated entropy, and we will average the frame entropies to give a single representative value. Also, the data chi-squared term must be updated so that each data point is compared with its simultaneous reconstructed image. Finally, we will add a new term $\mathcal{R}_x(\{\mathbf{I}_j\})$ with an associated hyperparameter α_x to regularize the dynamical images (we use x to label different choices for this term). This additional term can enforce expected properties such as continuity from frame to frame, a stable average image, or stable motion. The objective function for dynamical imaging then takes the form

$$J = \chi^2(\{\mathbf{I}_j\}, \mathbf{d}) - \alpha_S \left[\frac{1}{N_t} \sum_{j=1}^{N_t} S(\mathbf{I}_j) \right] + \alpha_x \mathcal{R}_x(\{\mathbf{I}_j\}). \quad (3)$$

Note that multiple dynamical regularizers can easily be combined in this framework, and additional regularization terms could be added (e.g., to mitigate interstellar scattering; Johnson 2016). The main purpose of this paper is to develop effective and efficient choices for the dynamical regularization terms \mathcal{R}_x and to test their performance on a variety of simulated data for the EHT.

2.4. General Considerations for Dynamical Imaging

Before developing specific strategies for dynamical imaging, it is instructive to consider how intrinsic variability can affect image reconstructions that assume a static source. Each baseline changes slowly with Earth’s rotation, so variability of an image on much shorter timescales introduces variations in measured visibilities over small baseline displacements $\Delta \mathbf{u}$. From Equation (1), we see that variations in the visibility over $\Delta \mathbf{u}$ require that the image flux extends over an angular scale $\Delta \mathbf{x} \sim 1/(2\pi|\Delta \mathbf{u}|)$. This mathematical uncertainty relationship arises because the variables of spatial position (\mathbf{x}) and spatial wavenumber (\mathbf{u}) are Fourier conjugates. One consequence of this property is that, for a static image, variations in the complex visibility seen over a baseline displacement of $\Delta \mathbf{u}$ can be used to infer the image field of view (FOV) without requiring detailed imaging. Likewise, the image field of view determines a maximum averaging time for visibilities sampled from a static image (see Section 6.4 of Thompson et al. 2017). However, for a variable source interpreted in the context of a static image, rapid intrinsic variability implies the existence of spurious image structure on large scales.

We now consider some specific examples. First, suppose that variations in the visibility amplitude are seen on a timescale of 5 minutes for an EHT baseline of length $5 \text{ G}\lambda$. In this case, $|\Delta\mathbf{u}| \sim (5 \text{ minutes}) / (24 \text{ hr}) \times (2\pi) \times (5 \text{ G}\lambda) \approx 110 \text{ M}\lambda$. These variations would then imply an image extent of roughly $1 / (2\pi|\Delta\mathbf{u}|) \approx 300 \mu\text{as}$. Note that this inferred extent is an order of magnitude larger than the measured size of Sgr A* at $\lambda = 1.3 \text{ mm}$ ($\approx 40 \mu\text{as}$; Doeleman et al. 2008). In addition, the snapshot visibilities can be compared across the exceptionally wide bandwidths of the EHT (4 GHz in 2017, and 18 GHz of spanned bandwidth in 2018 via dual-sideband recording). These also provide $|\Delta\mathbf{u}| \sim (5 \text{ G}\lambda) \times (4 \text{ GHz}) / (230 \text{ GHz}) \sim 100 \text{ M}\lambda$. Thus, visibilities that vary on timescales of minutes but that are stable across the full EHT bandwidth would provide firm evidence of rapid intrinsic variability.

As another trivial example, no static image can describe data in which the total flux density (i.e., the zero-baseline visibility $V(0)$) is changing with time. The problems of imaging a variable source are further exacerbated with multiple sites because different baseline tracks can cross so that the same spatial Fourier component is sampled at multiple times (see Figure 2).⁷

As these examples illustrate, in some cases intrinsic variability can be robustly decoupled from extrinsic sampling variability (from a changing baseline with Earth’s rotation) by constraining the image FOV (effectively imposing an image prior). In Lu et al. (2016), the authors use temporal filtering and normalization of measured visibilities to mitigate intrinsic variability; their chosen filter parameters effectively impose a maximal FOV. However, the strategy of post-processing visibilities has some limitations relative to an image-based approach; for instance, visibility-domain smoothing with a baseline-based algorithm does not account for mismatched visibilities on crossing baseline tracks. More generally, visibility-domain averaging of robust observables such as closure phases and closure amplitudes can introduce bias in the measurements. Dynamical imaging addresses both these limitations, providing a framework in which the intrinsic variability is incorporated into the imaging model, so that measurements can be directly compared with reconstructed images without additional averaging.

3. Regularizers for Dynamical Imaging

We now derive three regularizers appropriate for dynamical imaging. Our motivation is to identify regularizers that reflect a range of expected properties for astrophysical cases of interest and that also are efficient to implement in a numerical minimization scheme. Our first regularizer only enforces continuity from frame to frame (Section 3.1), the second favors frames that are small perturbations from the time-averaged image (Section 3.2), and the third describes an image that evolves approximately as a fluid with a steady motion field (Section 3.3). We summarize the properties of these regularizers in Section 3.4.

3.1. Smoothly Varying Images over Time

We first develop a generic regularizer that only seeks to enforce continuity from frame to frame in reconstructed

images. Because the motion between frames is unknown and may not be constant in time, this regularizer compares the reconstructed flux density of a pixel at one time with the flux density of nearby pixels at a subsequent time. The appropriate definition of “nearby” depends on the product of the expected velocity of moving features and the frame interval (which could potentially be irregular). Because this strategy is based on enforcing continuity over short time intervals, we denote the regularizer by $\mathcal{R}_{\Delta t}$.

Explicitly, we compute the summed difference among all adjacent images after blurring the frames, $\mathbf{I}_j \rightarrow B(\mathbf{I}_j)$, using a circular Gaussian kernel with standard deviation $\sigma_{\Delta t}$. We will focus on two particular choices to define the distance between a pair of images. First, there is the total pixel-by-pixel squared difference:

$$\mathcal{D}_2(\mathbf{I}, \mathbf{I}') \equiv \|\mathbf{I} - \mathbf{I}'\|^2 = \sum_{m,\ell} (I_{m,\ell} - I'_{m,\ell})^2. \quad (4)$$

A simple generalization of this regularizer is to replace the squared norm $\|\dots\|^2$ with $\|\dots\|_p^p$ for some fixed $p > 0$, $\mathcal{D}_2 \rightarrow \mathcal{D}_p$.

A second option to define an image distance is the relative entropy (i.e., the Kullback–Leibler divergence):

$$\mathcal{D}_{\text{KL}}(\mathbf{I}, \mathbf{I}') = D(\mathbf{I}' \|\mathbf{I}) \equiv \sum_{m,\ell} I'_{m,\ell} \ln \left(\frac{I'_{m,\ell}}{I_{m,\ell}} \right). \quad (5)$$

The relative entropy is frequently used to regularize traditional VLBI imaging against a specified image prior for the reconstruction (see, e.g., Cornwell & Evans 1985) and is also often used for multimodel image registration (Wells et al. 1996; Viola & Wells 1997). Note that the relative entropy is not symmetric, $\mathcal{D}_{\text{KL}}(\mathbf{I}, \mathbf{I}') \neq \mathcal{D}_{\text{KL}}(\mathbf{I}', \mathbf{I})$, and it need not be positive unless the total flux densities of the two images are equal: $\sum_{\ell,m} I_{\ell,m} = \sum_{\ell,m} I'_{\ell,m}$. Thus, useful alternatives include computing the relative entropy with respect to the normalized images (to preserve positivity of the divergence) and symmetrized versions such as $\frac{1}{2}[\mathcal{D}_{\text{KL}}(\mathbf{I}, \mathbf{I}') + \mathcal{D}_{\text{KL}}(\mathbf{I}', \mathbf{I})]$ or $\frac{1}{2}[\mathcal{D}_{\text{KL}}(\mathbf{I}, \bar{\mathbf{I}}) + \mathcal{D}_{\text{KL}}(\mathbf{I}', \bar{\mathbf{I}})]$ with $\bar{\mathbf{I}} \equiv \frac{1}{2}(\mathbf{I} + \mathbf{I}')$ (i.e., the Jensen–Shannon divergence).

The dynamical regularizer then takes the form

$$\mathcal{R}_{\Delta t}(\{\mathbf{I}_k\}) \equiv \sum_{j=1}^{N_t-1} \mathcal{D}(B(\mathbf{I}_j), B(\mathbf{I}_{j+1})). \quad (6)$$

This regularizer thereby penalizes changes between frames, with a steeply decreasing penalty for changes on scales smaller than $\sim \sigma_{\Delta t}$. One limitation of the $\mathcal{R}_{\Delta t}$ regularizer is that it does not favor stable “momentum” of features between frames. In Section 3.3, we will discuss an alternative regularizer that favors reconstructions with smooth and stable motion between frames. In its simplest implementation, this regularization then depends on only two hyperparameters: $\sigma_{\Delta t}$ and $\alpha_{\Delta t}$ (see Section 2.3). However, note that $\mathcal{R}_{\Delta t}$ is meaningful even in the limit $\sigma_{\Delta t} \rightarrow 0$ (i.e., comparing the total difference between adjacent frames with no blurring applied). This limit is appropriate when the expected motion between consecutive frames is smaller than the finest resolution of reconstructed features (comparable to the nominal array resolution).

⁷ For EHT observations of Sgr A*, SPT-PV and SPT-SMT baselines are very close in u - v space but have a 7.1 hr offset in sampling. Also, the ALMA-LMT and ALMA-SMT tracks intersect with a time offset of 1.4 hr. See Figure 2.

This regularization is effective in an imaging framework because the gradient (with respect to changes in each pixel of the $\{\mathbf{I}_k\}$) can be evaluated efficiently. For example, for the \mathcal{D}_{KL} distance function,

$$\frac{\partial \mathcal{R}_{\Delta t}}{\partial \mathbf{I}_k} = B \left(\left[1 + \ln \frac{B(\mathbf{I}_k)}{B(\mathbf{I}_{k-1})} \right] \delta_{k>1} - \frac{B(\mathbf{I}_{k+1})}{B(\mathbf{I}_k)} \delta_{k < N_t} \right), \quad (7)$$

where $\mathbf{1}$ denotes a vector of length N^2 with every element equal to unity and the indicator function δ_x is defined to be unity when the subscripted condition x is satisfied and is zero otherwise. Observe that calculating the gradient via Equation (7) requires roughly $\mathcal{O}(N_t \times N^2)$ computations, while calculating the gradient via finite differences of Equation (6) requires roughly $\mathcal{O}(N_t \times N^4)$ computations. Thus, for typically VLBI image reconstructions, which have $N \sim 10^2\text{--}10^3$, these analytic gradients speed up the imaging by several orders of magnitude.

Note that in Equation (7) and throughout this paper, operations such as quotients, powers, norms ($\|\dots\|$), and products of image vectors are to be computed elementwise. See the [Appendix](#) for corresponding expressions for other distance metrics.

3.2. A Stable Average Image with Small Perturbations

Our next dynamical regularizer is suitable for the case when each snapshot of the time-variable image can be described as a small perturbation from the time-averaged image. This case is applicable for a broad range of stationary processes, such as steady-state accretion or jet systems. Because this regularizer enforces snapshot images to be only small perturbations from the time-averaged image, we denote it $\mathcal{R}_{\Delta t}$.

To proceed, we approximate the time-averaged image by the average of all the reconstructed frames: $\mathbf{I}_{\text{avg}} \equiv \frac{1}{N_t} \sum_{j=1}^{N_t} \mathbf{I}_j$. We then define $\mathcal{R}_{\Delta t}$ to be the summed distance between the estimated time-averaged image and each reconstructed frame:

$$\mathcal{R}_{\Delta t}(\{\mathbf{I}_k\}) = \sum_{j=1}^{N_t} \mathcal{D}(\mathbf{I}_{\text{avg}}, \mathbf{I}_j). \quad (8)$$

As for $\mathcal{R}_{\Delta t}$, a convenient property of this regularization is that the gradient is efficient to compute (see the [Appendix](#)).

Note that this regularization requires only one tunable hyperparameter, $\alpha_{\Delta t}$, determining the overall strength of the regularization. An additional blurring step could be added if individual frames occasionally have flux density in regions that are otherwise empty (e.g., to accommodate flaring behavior), but the average image will tend to act like a blurring operator, so we do not expect that this step will normally be needed. Another difference between $\mathcal{R}_{\Delta t}$ and $\mathcal{R}_{\Delta t}$ is that $\mathcal{R}_{\Delta t}$ is insensitive to abrupt changes between frames or even reordering of frames. In this respect, $\mathcal{R}_{\Delta t}$ is analogous to entropy, which is unaffected by the placement of pixels in an image (see Section 2.2).

3.3. Time-variable Images with Regular Motion

Our third regularizer is motivated by the case when an image evolves according to a regular prescription for motion—i.e., a steady flow of flux density over time. In this case, the appearance at one time largely determines the appearance at

nearby times. A natural example of this case is an accretion flow, and we will denote this regularization by $\mathcal{R}_{\text{flow}}$.

To proceed, we consider the image $I(x, y, t)$ to be an evolving “fluid” with a stable flow vector field $\mathbf{v}(x, y)$. We further assume that the flux density is approximately conserved between nearby frames, so the time-variable images must approximately obey a continuity equation:

$$\begin{aligned} \frac{\partial I(x, y, t)}{\partial t} &= -\nabla \cdot [I(x, y, t)\mathbf{v}(x, y)] \\ &= -[\mathbf{v} \cdot \nabla I + I \nabla \cdot \mathbf{v}], \end{aligned} \quad (9)$$

where $\nabla = \{\partial/\partial x, \partial/\partial y\}$ denotes a 2D spatial gradient operator. Hence, at a given time, the image and flow can be combined to estimate the image at a slightly later time:

$$\begin{aligned} I(x, y, t + \delta t) &\approx I(x, y, t) + \delta t \frac{\partial I(x, y, t)}{\partial t} \\ &= I(x, y, t) - \delta t \times (\mathbf{v} \cdot \nabla I + I \nabla \cdot \mathbf{v}). \end{aligned} \quad (10)$$

We can now use this approximate forward evolution to regularize multiframe imaging. We consider a regularizer $\mathcal{R}_{\text{flow}}$ that is given by the summed difference between each frame and its predicted values based on linearized forward evolution of the previous frame (via Equation (10)) with a discrete spatial gradient operator replacing the continuous gradient. By only comparing adjacent frames, we relax the assumption that the flow field completely determines all forward evolution of a system from an initial state—we only seek to favor series of images that *approximately* respect a stable flow field over short intervals. For specificity, we will work with the \mathcal{D}_2 regularization, in which case

$$\begin{aligned} \mathcal{R}_{\text{flow}}(\{\mathbf{I}_k\}, \mathbf{m}) &= \sum_{j=1}^{N_t-1} \|\mathbf{I}_{j+1} - (\mathbf{I}_j - \nabla \cdot [\mathbf{I}_j \mathbf{m}])\|^2 \\ &\approx \sum_{j=1}^{N_t-1} \|\mathbf{I}_{j+1} - \mathbf{I}_j + \mathbf{m} \cdot \nabla \mathbf{I}_j + (\nabla \cdot \mathbf{m}) \mathbf{I}_j\|^2 \\ &\equiv \sum_{j=1}^{N_t-1} \|\mathbf{I}_{j+1} - \overleftrightarrow{\mathbf{F}}_{\text{flow}} \cdot \mathbf{I}_j\|^2. \end{aligned} \quad (11)$$

Here, we have replaced the velocity field by a dimensionless motion field $\mathbf{m} = \mathbf{v} \delta t / \delta x$, where δx is the discrete grid spacing of reconstructed images, δt is the image spacing in time, and ∇ now denotes a finite-difference operator that approximates the continuous 2D gradient. We have also defined the linear operator $\overleftrightarrow{\mathbf{F}}_{\text{flow}} \equiv 1 - \mathbf{m} \cdot \nabla - \nabla \cdot \mathbf{m}$, and the second line is an approximation that only becomes exact in the continuous limit because identities such as the product rule do not hold exactly for the discrete gradient operator.⁸ In this expression and elsewhere, images are treated as 1D vectors, the 2D vector flow is unwrapped to be a 1D vector of 2D motions $m_i = \{m_{i,x}, m_{i,y}\}$, and products of vectors (e.g., $(\nabla \cdot \mathbf{m}) \mathbf{I}_j$) are to be computed by multiplying the vectors

⁸ For example, the 1D finite forward difference operator $[\nabla_f x]_i \equiv x_{i+1} - x_i$ satisfies $\nabla_f(x y) = y \nabla_f x + x \nabla_f y + (\nabla_f x)(\nabla_f y)$. Likewise, the analogous finite backward difference operator $[\nabla_b x]_i \equiv x_i - x_{i-1}$ satisfies $\nabla_b(x y) = y \nabla_b x + x \nabla_b y - (\nabla_b x)(\nabla_b y)$. In the present work, we keep the gradient operator general and assume smooth images with small fractional gradients so that $\nabla(x y) \approx y \nabla x + x \nabla y$.

point by point (the Hadamard product). For the construction of this regularizer to be valid (i.e., for the linear approximation of Equation (10) to hold), the reconstructed frames must have smooth spatial and temporal gradients; the former is enforced by the image regularization terms $S(I_j)$, while the latter is enforced by the dynamical regularization. More concretely, the time resolution should be fine enough that the vectors of the motion field do not exceed the nominal VLBI beam that describes the angular resolution of the reconstructed images (analogous to the Courant–Friedrichs–Lewy condition for numerical integration of partial differential equations; Courant et al. 1967). Thus, observations with finer angular resolution require correspondingly finer temporal resolution. However, this condition does not require observations that are spaced this closely in time; additional frames can be included that do not have corresponding data constraints (see Section 4).

The major difference between the flow regularizer and our previous dynamical regularizers is that, in addition to estimating all the image frames, this reconstruction strategy must simultaneously estimate the flow vector field \mathbf{m} . In the Appendix, we provide analytic expressions for the gradients of $\mathcal{R}_{\text{flow}}$ with respect to the images and flow, enabling efficient estimation of both in a nonlinear minimization framework.

The motion field can also be regularized (just as the individual frames are regularized) to ensure that it varies smoothly over the image. Because \mathbf{m} is a vector field, one could potentially use the same regularizations as have been proposed for polarimetric synthesis imaging (see, e.g., Chael et al. 2016; Akiyama et al. 2017a). However, most of these choices are insensitive to the polarization direction, with the exception of total variation. We will use a closely related choice, the total squared gradient of the velocity field: $\mathcal{R}_m(\mathbf{m}) = \|\nabla \mathbf{m}_x\|^2 + \|\nabla \mathbf{m}_y\|^2$. This regularizer is also commonly used in studies of optical flow, which reconstruct a flow field from a series of images rather than from sparse Fourier sampling (Horn & Schunck 1981).⁹ The gradient of \mathcal{R}_m with respect to the flow field is simply $\partial \mathcal{R}_m / \partial \mathbf{m}_{x,y} = -2\nabla \cdot (\nabla \mathbf{m}_{x,y})$. Also, \mathcal{R}_m has an associated hyperparameter α_m to govern its overall weight.

Alternatively, in some cases the flow may be known or may be adequately modeled with a small number of parameters (see, e.g., Bouman et al. 2017). In these cases, dynamical imaging is plausible for much sparser arrays. At the other extreme, with sufficient data, the assumption of a stationary flow can be relaxed and the dynamical imaging could allow a smoothly evolving flow field over time.

3.4. Summary and Asymptotic Properties of Dynamical Regularizers

We have developed three regularizers that are suitable for dynamical imaging: $\mathcal{R}_{\Delta t}$, $\mathcal{R}_{\Delta l}$, and $\mathcal{R}_{\text{flow}}$. $\mathcal{R}_{\Delta t}$ favors continuity from frame to frame within a spatial displacement tolerance determined by $\sigma_{\Delta t}$, $\mathcal{R}_{\Delta l}$ favors frames that are small perturbations from the time-averaged image, and $\mathcal{R}_{\text{flow}}$ favors frames that evolve approximately according to a time-independent flow vector field, \mathbf{m} (see Figure 3 for a schematic comparison of these strategies). Each regularizer requires one

associated hyperparameter, α_x , that assigns overall weight to the dynamical regularization. $\mathcal{R}_{\Delta t}$ also requires one parameter describing the expected angular motion of features from frame to frame, and $\mathcal{R}_{\text{flow}}$ requires a hyperparameter α_m to regularize the estimated flow field. $\mathcal{R}_{\Delta l}$ requires no additional hyperparameters. These hyperparameters can be fixed according to a priori expectations, they can be treated as Lagrange multipliers and varied to give properties such as a final reduced chi-squared of unity, or they can be estimated using cross-validation (Akiyama et al. 2017b). As $\alpha_x \rightarrow 0$, each regularization strategy is equivalent to independently imaging a series of frames. Taking $\alpha_{\Delta t} \rightarrow \infty$ or $\alpha_{\Delta l} \rightarrow \infty$ would enforce a static reconstructed image, equivalent to conventional imaging, although this is not necessarily true for $\alpha_{\text{flow}} \rightarrow \infty$.

It is also possible to normalize each regularizer such that its value is unaffected by the choice of temporal (Δt ; the frame spacing) and spatial resolution (Δx ; the pixel linear dimension) of the reconstruction. As these become arbitrarily small, the dynamical reconstruction approaches a continuous representation in time and space. In particular, the limit $\Delta t \rightarrow 0$ is relevant when including interpolating frames (see Section 4), which enable arbitrary temporal resolution. The required normalization factor depends on the chosen distance metric. For the \mathcal{D}_p distance, each regularizer \mathcal{R}_x must be multiplied by $\Delta x^{-2(p-1)} \Delta t^{-(p-1)}$. For the Kullback–Leibler divergence \mathcal{D}_{KL} and its symmetrized variants, the normalization factor is Δt^{-1} . After applying this factor, the associated hyperparameters α_x will be unaffected by the choice of temporal or angular resolution, assuming that the motion in each is well resolved.

4. Dynamical Imaging and Interpolation

Dynamical imaging also serves as a framework for temporal interpolation of images. Namely, image frames can be added, even at times when there are no corresponding data. Without dynamical regularization, these additional frames would default to the image that maximizes the entropy (typically an image with constant brightness, possibly uniformly zero). However, dynamical imaging will favor frames that respect the chosen regularizer. For example, when using $\mathcal{R}_{\Delta t}$, the additional frames will converge toward images that enforce continuity of features with the nearest data-constrained frames. For $\mathcal{R}_{\Delta l}$, frames without data will default to the estimated time-averaged image. For $\mathcal{R}_{\text{flow}}$, unconstrained frames will interpolate according to the derived flow map. In each case, frames with missing data can inherit partial information from other times. Moreover, for the case of $\mathcal{R}_{\text{flow}}$, frames can be intentionally spaced at finer resolution than the sampling time to ensure that the linear approximation of Equation (10) is accurate. All these strategies will produce different results than a straightforward linear interpolation between images, as is commonly used to visualize multiepoch VLBI studies (e.g., Lister et al. 2016).

However, we have found that the interpolated frames can sometimes have a different appearance than data-constrained frames. For instance, when using the $\mathcal{R}_{\Delta t}$ regularization with $\sigma_{\Delta t} > 0$, the interpolated frames are “blurred out” relative to the data-constrained frames. This blurring is unsurprising, as it helps to minimize the mean squared difference among adjacent frames as elements of flux move in time. Consequently, the interpolated frames may achieve continuity of features but may have temporal discontinuities in the total flux density or image entropy.

⁹ Another difference between traditional studies of optical flow and our approach is that the former assume an incompressible flow: $\nabla \cdot \mathbf{m} = 0$.

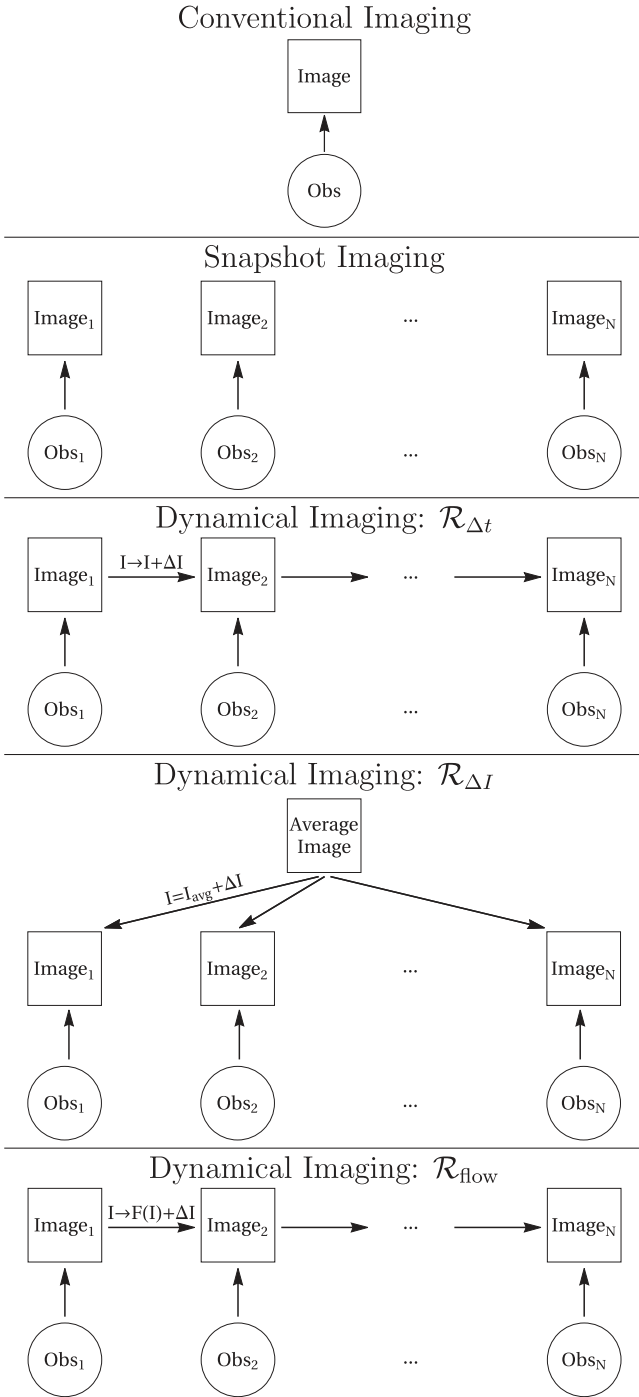


Figure 3. Schematic comparison of our proposed imaging methods. In conventional imaging, a single image is reconstructed from an observation. In snapshot imaging, a set of images is reconstructed from a corresponding set of observations, and each reconstruction is performed independently. For dynamical imaging with $\mathcal{R}_{\Delta t}$ regularization, the images are assumed to be temporally connected, each being a small perturbation of the previous frame. With $\mathcal{R}_{\Delta I}$, each frame is assumed to be a small perturbation of the time-averaged reconstructed frames, and image order is irrelevant. For $\mathcal{R}_{\text{flow}}$, each image is a small perturbation of the previous image after forward evolution with the stationary flow, which must be reconstructed along with the images. For the dynamical reconstructions, images can be meaningfully reconstructed even at times with no corresponding observation (see Section 4). See Bouman et al. (2017) for a discussion relating these schematic diagrams to a probabilistic graphical model for the dynamical imaging problem.

To mitigate the artifacts in interpolated frames, we can directly enforce continuity of quantities such as flux density and entropy. To do so, we add corresponding terms to the

objective function (Equation (3)). For example, to enforce continuity of image entropy, one can add

$$\mathcal{R}_{\Delta S} \equiv \sum_{j=1}^{N_i-1} [S(\mathbf{I}_j) - S(\mathbf{I}_{j+1})]^2, \quad (12)$$

weighted by an associated hyperparameter $\alpha_{\Delta S}$. The hyperparameter can be adjusted so that this term allows continuous variations of the entropy among frames without forcing the entropy of each frame to be equal. The gradient of this term is straightforward to compute in terms of the single-image gradients:

$$\begin{aligned} \frac{\partial \mathcal{R}_{\Delta S}}{\partial \mathbf{I}_k} = & 2 \{ [S(\mathbf{I}_j) - S(\mathbf{I}_{j-1})] \delta_{k>1} \\ & + [S(\mathbf{I}_j) - S(\mathbf{I}_{j+1})] \delta_{k<N_i} \} \frac{\partial S(\mathbf{I}_k)}{\partial \mathbf{I}_k}. \end{aligned} \quad (13)$$

Likewise, to make the total flux continuous from frame to frame, we can add

$$\mathcal{R}_{\Delta F} \equiv \sum_{j=1}^{N_i-1} [F(\mathbf{I}_j) - F(\mathbf{I}_{j+1})]^2, \quad (14)$$

weighted by an associated hyperparameter $\alpha_{\Delta F}$, where $F(\mathbf{I}) \equiv \sum_{\ell,m} I_{\ell,m}$ denotes the total flux density of an image. The gradient is again straightforward to compute:

$$\begin{aligned} \frac{\partial \mathcal{R}_{\Delta F}}{\partial \mathbf{I}_k} = & 2 \{ [F(\mathbf{I}_j) - F(\mathbf{I}_{j-1})] \delta_{k>1} \\ & + [F(\mathbf{I}_j) - F(\mathbf{I}_{j+1})] \delta_{k<N_i} \} \frac{\partial F(\mathbf{I}_k)}{\partial \mathbf{I}_k}, \end{aligned} \quad (15)$$

where $\partial F(\mathbf{I}_k)/\partial \mathbf{I}_k$ is a vector with each element equal to 1 and of length equal to the number of pixels in image \mathbf{I}_k .

5. Examples of Dynamical Imaging

We will now show a few representative examples of dynamical imaging using simulated VLBI observations, and we will discuss general trends that we have identified. We conclude this section with an example showing frames from dynamical imaging of M87.

5.1. Implementation and Procedure

We implemented the dynamical regularizers developed in Section 3 as an extension to the `eht-imaging`¹⁰ Python library, which was originally developed for polarimetric VLBI imaging (Chael et al. 2016). This library provides a modular and flexible imaging framework that can utilize a variety of imaging regularizers (e.g., entropy, total variation, and ℓ_p) and arbitrary combinations of data constraints (e.g., complex visibilities, the bispectrum, or closure quantities). We also used this library for generating synthetic data. Except when noted otherwise, we chose observing parameters that correspond to the 2017 EHT: an observing bandwidth of 4 GHz and site system equivalent flux densities given in Table 1. For simplicity, our simulated observations of Sgr A* account for sensitivity losses from the blurring effects of interstellar scattering (Fish et al. 2014) but not irregular, refractive effects (Johnson & Gwinn 2015).

¹⁰ <https://github.com/achael/eht-imaging>

Table 1
Assumed Site System Equivalent Flux Densities (SEFD)

Site	SEFD (Jy)
SMA/JCMT	4900
SMT	11900
LMT	560
ALMA/APEX	220
SPT	1600
PdB	1600
PV	2900
CA	10000
KP	10000

Note. Most SEFDs match what was specified in the 2016 EHT call for proposals. CA and KP are as-yet-hypothetical EHT sites at the location of the CARMA array and at Kitt Peak, respectively.

The fundamental interferometric data product is the sampled complex interferometric visibility (Equation (1)). However, because of a large stochastic contribution from the atmosphere to each site’s phase, high-frequency VLBI arrays can typically only measure quantities such as closure phase robustly (Thompson et al. 2017). Imaging algorithms can then work with these robust data products directly (see, e.g., Buscher 1994; Baron et al. 2010; Bouman et al. 2016; Chael et al. 2016; Akiyama et al. 2017b). Nevertheless, in the near future, improved techniques such as simultaneous subarrayed observations of calibrators (see Broderick et al. 2011) may provide absolute phase information. Also, we expect dynamical imaging to be applicable at the lower frequencies where phase referencing is routine; e.g., observations with the VLBA at wavelengths of 3 mm and longer. Thus, we will show results both when using complex visibilities and when using only visibility amplitudes and closure phase.

To minimize the objective function given by Equation (3) (i.e., to perform dynamical imaging), we used the nonlinear minimization package `optimize.minimize` of `SciPy` (Jones et al. 2001). We used the Limited-Memory BFGS algorithm (Byrd et al. 1995) except when memory requirements to store the partial Hessian became prohibitive (generally when imaging $\gg 100$ frames simultaneously), in which case we instead used the conjugate gradient algorithm implemented in `SciPy` (which does not compute the Hessian).

Similar to conventional VLBI imaging, convergence to the minimum of the objective function for dynamical imaging can be challenging because of the extremely high-dimensional ($N^2 \times N_t \gtrsim 10^6$) parameter space surveyed. Convergence is especially challenging when using only robust VLBI observables, such as closure phases, rather than complex visibilities because the relative image centroid among frames is only constrained by the dynamical regularization. Consequently, we used a number of strategies to assist convergence, most involving multiple iterations of minimization with modified initial values. One particularly effective strategy for avoiding local minima, following Chael et al. (2016), is to repeatedly image the data, re-initializing the minimization each time to be equal to the previous reconstructed images convolved with the nominal VLBI array resolution. In cases with many high-quality data points for each snapshot, we iterated between imaging all frames and allowing convergence to proceed on individual frames independently. For $\mathcal{R}_{\Delta I}$, we repeatedly

re-initialized all frames to the current average image. For $\mathcal{R}_{\text{flow}}$, we repeatedly re-initialized the flow to be uniformly zero. In all cases, we determined the dynamical imaging hyperparameters α_x by making them as large as possible while still achieving a final reduced χ^2 near unity.

One modification to the prescription outlined above that we did find to be effective for larger arrays, such as the VLBA, was to apply the dynamical regularizers to the logarithm of the reconstructed images rather than to the images when using the \mathcal{D}_2 or \mathcal{D}_p distance function (here, we assume image positivity). This change helps the dynamical regularization to improve time-variable imaging of faint image features and significantly improved reconstructed images with dynamic range $\gtrsim 100$, although it is unnecessary when using the relative entropy distance function \mathcal{D}_{KL} .

To assess the fidelity of reconstructed images when the true (model) image is known, we utilize the (normalized) mean squared error (MSE):

$$\text{MSE} \equiv \left[\sum_{\ell,m} (I_{\ell,m} - I'_{\ell,m})^2 \right] / \left[\sum_{\ell,m} I_{\ell,m}^2 \right]. \quad (16)$$

Here, $I_{\ell,m}$ is the model image and $I'_{\ell,m}$ is the reconstructed image. In cases where closure phases are used for image reconstructions, the image centroid is unconstrained and we report the MSE that is minimum over all shifts $\{\Delta\ell, \Delta m\}$. The precise value of the MSE should not be taken too literally because it is sensitive to sharp features in the original image that the array cannot resolve (see also Gomes et al. 2017). Nevertheless, the MSE does tend to provide a crude characterization of reconstructed image quality.

5.2. Dynamical Imaging of a Steady Accretion Flow

To examine the potential capabilities of dynamical imaging with EHT data, we generated synthetic data from a face-on view of a 3D GRMHD simulation of an accretion flow onto Sgr A* (b0-high from Shiokawa 2013). Figure 4 shows image reconstructions using both snapshot imaging (see Figure 3) and dynamical imaging. These reconstructions use the $\mathcal{R}_{\Delta I}$ regularizer with the \mathcal{D}_2 distance metric and complex visibilities for the data product; we reconstructed 99 frames, each $24t_G \approx 530$ s, for a total duration of 14.6 hr. At times with poor $u-v$ coverage, the snapshot images are uninformative while frames from dynamical imaging are close to the estimated time-averaged image. Dynamical imaging successfully identifies the time-variable regions of enhanced flux density and can also identify the flow direction, which is not apparent in the snapshot reconstructions.

To quantify the improvement of dynamical imaging relative to snapshot imaging, Figure 5 shows the MSE as a function of time for these two reconstructions. Notably, the MSE for snapshot imaging changes significantly over the observation, increasing steeply when the number of sites with mutual visibility of Sgr A* drops. In contrast, the MSE for dynamical imaging is lower overall and is steady, showing the increased resilience to limited data.

5.3. Estimates of Time-averaged Images

Another important utility of dynamical imaging is to estimate the time-averaged image over an observation. For




Figure 4. Example reconstruction of a face-on accretion disk with and without dynamical regularization. The full reconstruction comprised 100 frames of 8.8 minutes each beginning at 16:00 GST and spanning a total of 14.6 hr. The simulated images are from a 3D GRMHD simulation (b0-high from Shiokawa 2013). The above panels show the simulated images, snapshot reconstructions (using conventional maximum entropy imaging with only the instantaneous $u-v$ coverage), dynamical reconstructions (using $\mathcal{R}_{\Delta I}$), and the baseline coverage at three times. The color scale is linear and is consistent among different times but is scaled separately for each case based on the maximum brightness over all frames. Because the early and late frames have few data constraints, the snapshot image reconstructions of those frames are almost entirely uninformative and poorly approximate the true images. In contrast, the dynamical imaging reconstructions at those times appear almost identical, with the data only supporting small perturbations from the estimated time-averaged image.

studies of Sgr A* with the EHT, the time-averaged image is of intense interest because it may reveal distinctive features of the spacetime near a black hole such as the black hole “shadow” (Bardeen et al. 1972; Lunin 1979; Falcke et al. 2000; Takahashi 2004; Johannsen & Psaltis 2010). Yet, as discussed in Section 2.4, imaging techniques that assume a static source can be severely affected by intrinsic source variability; conventional imaging will not simply provide an estimate of the time-averaged image. Instead, the variability must be integrated into the imaging procedure by either preprocessing the data to render it compatible with a static-source assumption (Lu et al. 2016) or modifying the imaging procedure to accommodate image variability, as we propose here.

To test this application of dynamical imaging, we used the simulated EHT data from Lu et al. (2016) for a GRMHD simulation of an accretion flow onto Sgr A*. Note that, in contrast with our other examples, this data set included the CARMA array (the CARMA observatory was shut down in 2015), it sampled the images with 16 GHz of bandwidth rather than 4 GHz, and it used slightly different SEFDs than are given

in Table 1. Because the frame spacing is rather large in this example (3.7 minutes), we again used the regularizer $\mathcal{R}_{\Delta I}$ with the \mathcal{D}_2 distance metric.

Figure 6 compares the time-averaged estimates from dynamical imaging with the time-averaged simulated image. The estimated average image is comparable in quality to the image obtained with the scaling, averaging, and smoothing approach of Lu et al. (2016). We also found that averaging the snapshot images gives an estimated average image with comparable quality to these more sophisticated approaches, especially if periods with poor $u-v$ coverage were down-weighted or omitted. Thus, a weighted average of snapshot images, favoring times with superior $u-v$ coverage, may also produce reliable estimates of the time-averaged image and will provide a useful comparison for these other approaches.

5.4. Dynamical Imaging of Flares

Another important application for dynamical imaging is to study flares of Sgr A* via direct imaging. Figures 7 and 8 show example reconstructions for an orbiting “hot spot” near




Figure 5. MSE as a function of time for all frames of the reconstruction shown in Figure 4 when compared with the simultaneous simulated frames. As expected, the snapshot reconstructions vary erratically, with the most abrupt changes occurring when Sgr A* rises or sets at a participating site (here, we use an elevation limit of 15° ; see Figure 2). In contrast, the MSEs of reconstructions with dynamical imaging are relatively steady, demonstrating the added resilience of snapshot reconstructions when using all data concurrently. Nevertheless, with this sparse array, most of the improvement in MSE comes from the superior estimate of the time-averaged image rather than from precisely tracking the changing features of the image.




Figure 6. Comparing the time-averaged image of a simulated accretion flow (left) to three reconstruction strategies: standard VLBI imaging that assumes a static source (left-center), time-averaged snapshot image reconstructions (right-center), and time-averaged dynamical imaging with $\mathcal{R}_{\Delta t}$ regularization (right). For the standard and snapshot imaging, we used maximum entropy imaging. In this example, we matched the simulation and observing parameters given in Lu et al. (2016) (here, a single 12 hr observation). As expected, conventional VLBI imaging works poorly, but both averaged snapshot imaging and dynamical imaging are comparable in quality to scaling, averaging, and smoothing the interferometric visibilities before static imaging (MSE = 0.14; Lu et al. 2016). Here and throughout this paper, the color scale is linear.

Sgr A* (Broderick & Loeb 2006), using simulated observations that span only 27 minutes. For these examples, we used $\mathcal{R}_{\Delta t}$ regularization with the symmetrized KL divergence as the distance metric. While this simulated observation is too short to build up significant baseline coverage via Earth rotation to estimate an accurate time-averaged image, the reconstructions successfully identify the motion of the hot spot, especially if additional sites at Kitt Peak and CARMA are included. Thus, in the coming years, the EHT may be able to trace rapidly evolving structures and estimate orbital rotation curves from dynamical imaging, especially if the array continues to expand.

5.5. Comparison of Dynamical Imaging Methods

We next compare all three dynamical imaging strategies on simulated observations of an accretion flow viewed at an inclination of 30° with respect to the black hole’s rotation axis. Apart from the viewing inclination, the simulation is identical to the one shown in Figure 4. We used 1400 simulated movie frames, corresponding to 4.3 hr of observations, starting at a GST of 23:00. To simplify the comparison between the three

imaging methods and avoid discrepancies from poor convergence or image misalignment, we used full complex visibilities for dynamical imaging and the \mathcal{D}_2 distance metric for each. Each reconstructed movie has 234 frames (1/6 the time resolution of the input movie).

Figure 9 compares the average image of the simulation with the averaged images from the three dynamical imaging reconstructions using the 2017 EHT configuration (for the reconstruction using $\mathcal{R}_{\text{flow}}$ regularization, the reconstructed motion field is also shown). Figure 10 performs the same comparison for reconstructions that also included sites at Kitt Peak and at the location of the CARMA array. Figure 11 shows the MSE over time for each reconstruction from both array configurations. Despite their differing assumptions about the underlying image variability, the three methods give results that are broadly consistent. With the additional EHT sites of the second example, the flow clearly identifies the correct direction of motion, although the estimated magnitude of the motion underestimates by a factor of several the estimated time-averaged optical flow of the simulated images (Liu 2009). Thus, to recover precise details about the motion will likely require either more stringent dynamical imaging constraints (see, e.g., Bouman et al. 2017) or additional sites added to the EHT.

5.6. Dynamical Imaging of M87

As a final example, we used dynamical imaging on a series of 14 separate 43 GHz VLBA observations of M87 taken over a span of 70 days in 2008 as part of the M87 Movie Project (Walker et al. 2016; R. C. Walker et al. 2017, in preparation). We reconstructed a series of 24 images spaced by 3 days; each observation was associated with the nearest image in time, and the remaining images had no data constraints (see Section 4). We used $\mathcal{R}_{\Delta t}$ regularization with symmetrized Kullback–Leibler divergence as the distance metric and $\sigma_{\Delta t} = 0$ (see Section 3.1). We have found that this overall strategy works well for images with high dynamic range and irregular motion, without requiring fine-tuning of the imaging parameters. We also utilized iterative self-calibration, so that the dynamical imaging serves as both an imaging and calibration framework.

While detailed analysis of these results will be presented separately, Figure 12 shows four of the reconstructed frames with their corresponding static reconstructions over a short time interval (17 days). Even a moderately relativistic component, with an apparent transverse velocity of $2c$, would move by only 0.36 mas over this entire interval. Dynamical imaging successfully finds a series of similar images, each of which is consistent with its respective, self-calibrated data. By eliminating faint spurious features, outward motion along the jet is more readily evident in the dynamical reconstructions. Moreover, this interval includes one epoch (MJD = 54554) for which the original data were adversely affected by poor weather and were not considered to be of adequate quality for inclusion in the final CLEAN data set. Nevertheless, dynamical imaging is able to successfully link this period of inferior data to the higher-quality nearby epochs so that the reconstructed image is not perceptibly degraded.

6. Summary

In summary, we have developed three regularizers that are suitable for dynamical imaging: $\mathcal{R}_{\Delta t}$, $\mathcal{R}_{\Delta l}$, and $\mathcal{R}_{\text{flow}}$. $\mathcal{R}_{\Delta t}$




Figure 7. Example dynamical reconstruction of a simulated flare. This simulation is a “hot spot” orbiting Sgr A* with a period of 27 minutes (model B of Doeleman et al. 2009b). The total observation covers one full orbit of the hot spot. Top panels show six selected frames of the simulated images, middle panels show corresponding reconstructions with the 2017 EHT array, and bottom panels show reconstructions with the 2017 EHT array plus sites at the location of the CARMA array and at Kitt Peak, each with an assumed SEFD of 10,000 Jy. These reconstructions used $\mathcal{R}_{\Delta I}$ regularization and complex visibilities with only thermal noise added. The spurious structure to the southwest in each reconstruction reflects the significantly anisotropic beam, which contains almost no power at this location during the simulated GST range because of a void in the $u-v$ coverage at the corresponding (orthogonal) position angles (see Figure 2).




Figure 8. Same as Figure 7, but using only visibility amplitudes and closure phases for the reconstructed images.

favors continuity from frame to frame, $\mathcal{R}_{\Delta I}$ favors frames that are small perturbations from the time-averaged image, and $\mathcal{R}_{\text{flow}}$ favors frames that approximately evolve according to a stationary flow. For each of these regularizers, we have derived analytic gradients with respect to the unknown image

parameters, so each converges quickly with modest computational resources (e.g., all the reconstructions in this paper were performed on a personal computer). Each can be used with any choice of VLBI data product (e.g., complex visibilities, the bispectrum, or visibility amplitudes and closure phases) and




Figure 9. Comparison of dynamical imaging methods for an accretion flow viewed at 30° off the rotation axis. The left two panels show the average simulated image and the average simulated image blurred with half the nominal observing array beam (the full beam is $26 \mu\text{as} \times 16 \mu\text{as}$, with the major axis at a position angle of 73° east of north). The remaining panels show the average reconstructed image after dynamical imaging with $\mathcal{R}_{\Delta t}$, $\mathcal{R}_{\Delta I}$, and $\mathcal{R}_{\text{flow}}$ regularization, respectively. The flow field that was derived when using $\mathcal{R}_{\text{flow}}$ regularization is overplotted on the final image; flow vectors (scaled in length by a factor of 5) show the derived motion from frame to frame (i.e., over an interval of 67 s). As this example illustrates, EHT coverage in 2017–2018 is unlikely to be sufficient to derive a reliable flow field without additional constraints on the flow structure. Nevertheless, the reconstructed images with $\mathcal{R}_{\text{flow}}$ are broadly consistent with the other methods.




Figure 10. Same as Figure 9, but using the 2017 EHT array plus sites at the location of the CARMA array and at Kitt Peak. In this figure, the derived motion field vectors are scaled by a factor of 10. While the two additional sites hardly change the observing beam (the blurred image is nearly identical to Figure 9), they significantly improve the snapshot baseline coverage and the dynamical imaging.




Figure 11. MSE vs. time for the dynamical imaging reconstructions shown in Figures 9 and 10. Upper, thick lines show results for the 2017 EHT configuration; lower, thin lines include sites at the location of the CARMA array and at Kitt Peak. At early times, when the baseline coverage is minimal, $\mathcal{R}_{\Delta t}$ regularization provides the best results, showing the improvement that can be obtained under the assumption of a stable average image. At later times, $\mathcal{R}_{\Delta t}$ is as good, or slightly advantageous, showing the benefit when enforcing temporal continuity on reconstructed images. With the expanded EHT array configuration, all methods produce accurate and comparable results.

any choice or combination of image regularization for individual frames (e.g., entropy, ℓ_n -norm, or total variation).

For dynamical imaging, the most significant challenge we have encountered is suitable convergence to the optimal reconstruction, especially when using a small array and only VLBI closure quantities. We have discussed a number of strategies to assist convergence, most involving iterative re-imaging with blurring, averaging, individual frame imaging, or

other modifications at each stage. In Bouman et al. (2017), we develop an alternative approach to dynamical imaging that provides an analytic expression for the reconstructed images, lessening the problem of convergence, at the expense of restrictive constraints on the dynamical imaging framework. These methods can potentially be used in sequence to allow a flexible dynamical imaging strategy with reliable convergence.

A major motivation for this work is the possibility of imaging Sgr A* with the EHT. Even in its simplest implementation, dynamical imaging provides a framework to estimate the time-averaged image of Sgr A* and will be significantly more sensitive than static imaging approaches if there is significant intrinsic variability (see Figure 6). Dynamical imaging can also confirm key image features such as the black hole shadow based on their temporal signatures. For example, the region surrounding the shadow is expected to exhibit enhanced high-frequency variability (Shiokawa et al. 2017). Dynamical studies of Sgr A* will be crucial for estimating accretion disk inclinations, breaking a degeneracy in time-averaged images from the near symmetry orthogonal to the rotation axis (see, e.g., Broderick et al. 2011; Johnson et al. 2015a), and they will also be helpful for estimating the black hole spin. Namely, while the shape of the black hole shadow is almost independent of spin (Bardeen et al. 1972; Takahashi 2004; Johannsen & Psaltis 2010), orbital periods at the innermost stable circular orbit vary by nearly a factor of 10 depending on spin (Bardeen 1973).

For EHT imaging, $\mathcal{R}_{\Delta t}$ regularization appears especially promising because inhomogeneous data can be combined regardless of their spacing, specific observing cadence, or participating telescopes. Thus, any robust data products from multiple epochs can be merged to produce an average image

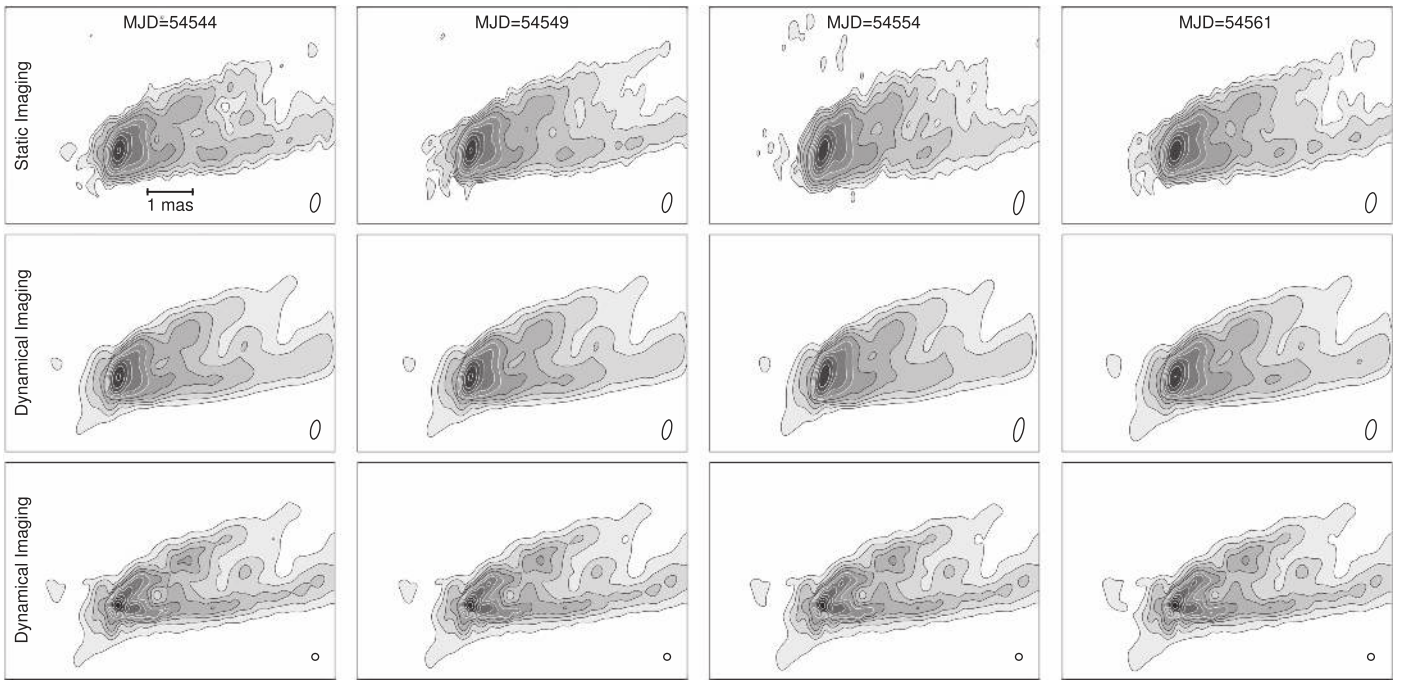


Figure 12. Comparison of static imaging (top) and dynamical imaging (middle, bottom) of M87 for four closely spaced epochs over a span of 17 days. The static images were reconstructed using CLEAN (for details, see Walker et al. 2016; R. C. Walker et al. 2017, in preparation); the dynamical images were reconstructed using $\mathcal{R}_{\Delta t}$ regularization with symmetrized Kullback–Leibler divergence as the distance metric. The restoring beam of the CLEAN images varies from epoch to epoch but is typically $\sim 430 \times 200 \mu\text{as}$ at a position angle of -13° . To simplify comparisons, the dynamical images in the middle row have been convolved with the corresponding CLEAN beam. To highlight evolution of compact structure, the dynamical images in the bottom row have been convolved with a circular Gaussian beam with FWHM of $150 \mu\text{as}$. At the resolution of the CLEAN beam, static imaging and dynamical imaging are broadly consistent, but faint features are more similar from frame to frame in the dynamical reconstruction, more readily identifying physical evolution. Contours in all panels are at equal levels, starting at $10.5 \text{ mJy}/\text{mas}^2$ ($=1 \text{ mJy}/\text{beam}$ in the first CLEAN image) and increasing by factors of 2.

that is not adversely affected by the variability, while also estimating the time-dependent perturbations. Moreover, our approach can be combined with other regularization frameworks, for instance, to simultaneously mitigate the effects of interstellar scattering (Johnson 2016). Looking forward, dynamical imaging may be a rich source of continued study of Sgr A* as millimeter VLBI continues to expand, potentially even to Earth-space baselines (see, e.g., Wild et al. 2009; Smirnov et al. 2012; Kardashev et al. 2014).

Our framework is also suitable for multiepoch VLBI imaging studies, including kinematical studies of relativistic jets (e.g., Kellermann et al. 2004; Jorstad et al. 2005; Lister et al. 2009, 2016; Mertens & Lobanov 2016; Walker et al. 2016), supernovae (e.g., Bartel et al. 2000; Bietenholz et al. 2003; Bartel 2009), and microquasars (e.g., Fomalont et al. 2001; Mioduszewski et al. 2004; Jeffrey et al. 2016). It can also be applied to multiepoch wide-field imaging, where highly sensitive instruments such as the Square Kilometer Array (SKA) are expected to detect a combination of static and variable sources (e.g., Fender et al. 2015; Metzger et al. 2015). For cases with regular motion, dynamical imaging with $\mathcal{R}_{\text{flow}}$ can self-consistently estimate the velocity field, while cases with irregular motion should be imaged with softer regularization, such as $\mathcal{R}_{\Delta t}$, $\mathcal{R}_{\Delta t}$ (see Figure 12), or a nonstationary flow. A benefit of our approach is that epochs with poor u - v coverage or sensitivity can partially inherit the higher resolution of neighboring epochs—our approach does not assume a constant VLBI beam among the epochs or even a constant spacing between epochs (see, e.g., Figure 5). Our approach can also incorporate iterative self-calibration to derive a calibration solution that is compatible with smooth structural evolution

among epochs. And while our focus has been on dynamical imaging of the total flux density, our methods are straightforward to adapt to full-Stokes polarization, which often shows more pronounced variability than the total flux density (e.g., Gabuzda et al. 2000).

We are particularly indebted to Craig Walker for sharing his calibrated M87 data and images and for many illuminating discussions. We gratefully acknowledge helpful conversations with John Wardle, James Guillochon, and Maciek Wielgus. We also thank Alan Marscher, Svetlana Jorstad, Dan Homan, and Matt Lister for making their multiepoch VLBI studies available to help refine and test our code. We thank Avery Broderick for providing the hot spot simulations used in this work. We thank the referee for suggesting the application of this technique to wide-field imaging studies with next-generation arrays. We thank the National Science Foundation (AST-1440254, AST-1614868) and the Gordon and Betty Moore Foundation (GBMF-3561, GBMF-5278) for financial support of this work. This work was supported in part by the black hole Initiative at Harvard University, which is supported by a grant from the John Templeton Foundation. F.R. is supported by the ERC Synergy Grant “BlackHoleCam” (Grant 610058). K.A. is financially supported by the program of Postdoctoral Fellowships for Research Abroad at the Japan Society for the Promotion of Science (JSPS). The VLBA is operated by the Long Baseline Observatory. The Long Baseline Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

Appendix Gradients of Dynamical Regularization Terms

We will now derive analytic expressions for the gradients of the dynamical regularization terms. These gradients depend on the chosen distance function, and so we will provide representative examples.

A.1. Gradients of $\mathcal{R}_{\Delta t}$

The gradient of $\mathcal{R}_{\Delta t}$ when using the \mathcal{D}_2 distance function is given by

$$\frac{\partial \mathcal{R}_{\Delta t}}{\partial \mathbf{I}_k} = 2B^2([\mathbf{I}_k - \mathbf{I}_{k-1}] \delta_{k>1} + [\mathbf{I}_k - \mathbf{I}_{k+1}] \delta_{k<N_t}), \quad (17)$$

where B^2 denotes a blurring operator that is applied twice, and the indicator function δ_x is defined to be unity when the subscripted condition x is satisfied and is zero otherwise. Note that B could be replaced by any $N^2 \times N^2$ matrix operator $\overleftrightarrow{\mathbf{B}}$ (acting on an image vector), in which case B^2 in Equation (17) must be replaced by $\overleftrightarrow{\mathbf{B}} \overleftrightarrow{\mathbf{B}}$.

Likewise, for the \mathcal{D}_p distance function,

$$\begin{aligned} \frac{\partial \mathcal{R}_{\Delta t}}{\partial \mathbf{I}_k} &= pB(|B(\mathbf{I}_k - \mathbf{I}_{k-1})|^{p-1} \\ &\times \text{sgn}(B(\mathbf{I}_k - \mathbf{I}_{k-1})) \delta_{k>1} + |B(\mathbf{I}_k - \mathbf{I}_{k+1})|^{p-1} \\ &\times \text{sgn}(B(\mathbf{I}_k - \mathbf{I}_{k+1})) \delta_{k<N_t}). \end{aligned} \quad (18)$$

Lastly, for the \mathcal{D}_{KL} distance function,

$$\frac{\partial \mathcal{R}_{\Delta t}}{\partial \mathbf{I}_k} = B \left(\left[1 + \ln \frac{B(\mathbf{I}_k)}{B(\mathbf{I}_{k-1})} \right] \delta_{k>1} - \frac{B(\mathbf{I}_{k+1})}{B(\mathbf{I}_k)} \delta_{k<N_t} \right), \quad (19)$$

where $\mathbf{1}$ denotes a vector of length N^2 with every element equal to unity. Gradients for variants of the \mathcal{D}_{KL} function can be computed similarly. We again emphasize that operations such as norms (|...|), quotients, powers, and products of image vectors are to be computed elementwise.

A.2. Gradients of $\mathcal{R}_{\Delta I}$

For the $\mathcal{R}_{\Delta I}$ regularization function, using the \mathcal{D}_2 distance function gives the following gradient:

$$\frac{\partial \mathcal{R}_{\Delta I}}{\partial \mathbf{I}_k} = 2(\mathbf{I}_k - \mathbf{I}_{\text{avg}}). \quad (20)$$

For the \mathcal{D}_p distance function, we obtain

$$\begin{aligned} \frac{\partial \mathcal{R}_{\Delta I}}{\partial \mathbf{I}_k} &= p |\mathbf{I}_k - \mathbf{I}_{\text{avg}}|^{p-1} \text{sgn}(\mathbf{I}_k - \mathbf{I}_{\text{avg}}) \\ &- \frac{p}{N_t} \sum_{j=1}^{N_t} |\mathbf{I}_j - \mathbf{I}_{\text{avg}}|^{p-1} \text{sgn}(\mathbf{I}_j - \mathbf{I}_{\text{avg}}). \end{aligned} \quad (21)$$

Note that the second term is independent of k and is zero for $p = 2$.

Lastly, for the \mathcal{D}_{KL} distance function, we find

$$\frac{\partial \mathcal{R}_{\Delta I}}{\partial \mathbf{I}_k} = 1 - \frac{\mathbf{I}_{\text{avg}}}{\mathbf{I}_k} + \frac{1}{N_t} \sum_{j=1}^{N_t} \ln \left(\frac{\mathbf{I}_{\text{avg}}}{\mathbf{I}_j} \right). \quad (22)$$

A.3. Gradients of $\mathcal{R}_{\text{flow}}$

For $\mathcal{R}_{\text{flow}}$ regularization, we must evaluate the gradients of $\mathcal{R}_{\text{flow}}$ with respect to both the images and the flow. The gradient with respect to the images can be written in a general form that only depends on the linear operator $\overleftrightarrow{\mathbf{F}}_{\text{flow}}$ and its transpose. For instance, using the \mathcal{D}_2 distance metric gives

$$\begin{aligned} \frac{\partial \mathcal{R}_{\text{flow}}}{\partial \mathbf{I}_k} &= 2(\mathbf{I}_k - \overleftrightarrow{\mathbf{F}}_{\text{flow}} \cdot \mathbf{I}_{k-1}) \delta_{k>1} \\ &- 2\overleftrightarrow{\mathbf{F}}_{\text{flow}}^\top \cdot (\mathbf{I}_{k+1} - \overleftrightarrow{\mathbf{F}}_{\text{flow}} \cdot \mathbf{I}_k) \delta_{k<N_t}. \end{aligned} \quad (23)$$

Using the identity that $\nabla^\top = -\nabla$ (appropriate for central finite-difference operators), we obtain $(\mathbf{m} \cdot \nabla)^\top = -(\nabla \cdot \mathbf{m} + \mathbf{m} \cdot \nabla)$. The other elements of $\overleftrightarrow{\mathbf{F}}_{\text{flow}}$ are diagonal matrices, so

$$\begin{aligned} \overleftrightarrow{\mathbf{F}}_{\text{flow}} &= (1 - \mathbf{m} \cdot \nabla - \nabla \cdot \mathbf{m})^\top \\ &= 1 + (\nabla \cdot \mathbf{m} + \mathbf{m} \cdot \nabla) - \nabla \cdot \mathbf{m} \\ &= 1 + \mathbf{m} \cdot \nabla. \end{aligned} \quad (24)$$









Substituting this result into Equation (23), we obtain

$$\begin{aligned} \frac{\partial \mathcal{R}_{\text{flow}}}{\partial \mathbf{I}_k} &= 2(\mathbf{I}_k - \overleftrightarrow{\mathbf{F}}_{\text{flow}} \cdot \mathbf{I}_{k-1}) \delta_{k>1} \\ &- 2(1 + \mathbf{m} \cdot \nabla) \cdot (\mathbf{I}_{k+1} - \overleftrightarrow{\mathbf{F}}_{\text{flow}} \cdot \mathbf{I}_k) \delta_{k<N_t}. \end{aligned} \quad (25)$$

Likewise, the gradient with respect to the flow vector field \mathbf{m} is given by

$$\begin{aligned} \frac{\partial \mathcal{R}_{\text{flow}}}{\partial \mathbf{m}} &= \sum_{j=1}^{N_t-1} \{2[\mathbf{I}_{j+1} - \overleftrightarrow{\mathbf{F}}_{\text{flow}} \mathbf{I}_j] \nabla \mathbf{I}_j \\ &- \nabla [2(\mathbf{I}_{j+1} - \overleftrightarrow{\mathbf{F}}_{\text{flow}} \cdot \mathbf{I}_j) \mathbf{I}_j]\} \\ &\approx -2 \sum_{j=1}^{N_t-1} \mathbf{I}_j \nabla (\mathbf{I}_{j+1} - \overleftrightarrow{\mathbf{F}}_{\text{flow}} \cdot \mathbf{I}_j). \end{aligned} \quad (26)$$

ORCID iDs

Michael D. Johnson  <https://orcid.org/0000-0002-4120-3029>
Lindy Blackburn  <https://orcid.org/0000-0002-9030-642X>
Andrew A. Chael  <https://orcid.org/0000-0003-2966-6220>
Julian Rosen  <https://orcid.org/0000-0003-1086-3224>
Hotaka Shiokawa  <https://orcid.org/0000-0002-8847-5275>
Freek Roelofs  <https://orcid.org/0000-0001-5461-3687>
Kazunori Akiyama  <https://orcid.org/0000-0002-9475-4254>
Vincent L. Fish  <https://orcid.org/0000-0002-7128-9345>

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
Bardeen, J. M. 1973, in *Black Holes (Les Astres Occlus)*, ed. C. DeWitt & B. S. DeWitt (New York: Gordon and Breach), 215
Bardeen, J. M., Press, W. H., & Teukolsky, S. A. 1972, *ApJ*, **178**, 347
Baron, F., Monnier, J. D., & Kloppenborg, B. 2010, *Proc. SPIE*, **7734**, 77342I

- Bartel, N. 2009, in ASP Conf. Ser. 402, Approaching Micro-Arcsecond Resolution with VSOP-2: Astrophysics and Technologies, ed. Y. Hagiwara et al. (San Francisco, CA: ASP), 243
- Bartel, N., Bietenholz, M. F., Rupen, M. P., et al. 2000, *Sci*, **287**, 112
- Bietenholz, M. F., Bartel, N., & Rupen, M. P. 2003, *ApJ*, **597**, 374
- Bouman, K. L., Johnson, M. D., Zoran, D., et al. 2016, in The IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), arXiv:1512.01413
- Bouman, K. L., Michael D. Johnson, M. D., Dalca, A. V., et al. 2017, IEEE Transactions, submitted
- Bower, G. C., Goss, W. M., Falcke, H., Backer, D. C., & Lithwick, Y. 2006, *ApJL*, **648**, L127
- Brinkerink, C. D., Falcke, H., Law, C. J., et al. 2015, *A&A*, **576**, A41
- Broderick, A. E., & Loeb, A. 2005, *MNRAS*, **363**, 353
- Broderick, A. E., & Loeb, A. 2006, *MNRAS*, **367**, 905
- Broderick, A. E., Loeb, A., & Reid, M. J. 2011, *ApJ*, **735**, 57
- Buscher, D. F. 1994, in IAU Symp. 158, Very High Angular Resolution Imaging, ed. J. G. Robertson & W. J. Tango (Dordrecht: Kluwer), 91
- Byrd, R. H., Lu, P., Nocedal, J., & Zhu, C. 1995, *SIAM Journal on Scientific Computing*, **16**, 1190
- Chael, A. A., Johnson, M. D., Narayan, R., et al. 2016, *ApJ*, **829**, 11
- Cornwell, T. J., & Evans, K. F. 1985, *A&A*, **143**, 77
- Courant, R., Friedrichs, K., & Lewy, H. 1967, *IBMJ*, **11**, 215
- Dexter, J., Agol, E., Fragile, P. C., & McKinney, J. C. 2010, *ApJ*, **717**, 1092
- Doeleman, S., Agol, E., Backer, D., et al. 2009a, in Science White Papers, Astro2010: The Astronomy and Astrophysics Decadal Survey, 68
- Doeleman, S. S., Fish, V. L., Broderick, A. E., Loeb, A., & Rogers, A. E. E. 2009b, *ApJ*, **695**, 59
- Doeleman, S. S., Weintraub, J., Rogers, A. E. E., et al. 2008, *Natur*, **455**, 78
- Eckart, A., Schödel, R., Meyer, L., et al. 2006, *A&A*, **455**, 1
- Falcke, H., Melia, F., & Agol, E. 2000, *ApJL*, **528**, L13
- Fender, R., Stewart, A., Macquart, J.-P., et al. 2015, arXiv:1507.00729
- Fish, V., Akiyama, K., Bouman, K., et al. 2016a, *Galax*, **4**, 54
- Fish, V. L., Doeleman, S. S., Beaudoin, C., et al. 2011, *ApJL*, **727**, L36
- Fish, V. L., Doeleman, S. S., Broderick, A. E., Loeb, A., & Rogers, A. E. E. 2009, *ApJ*, **706**, 1353
- Fish, V. L., Johnson, M. D., Doeleman, S. S., et al. 2016b, *ApJ*, **820**, 90
- Fish, V. L., Johnson, M. D., Lu, R.-S., et al. 2014, *ApJ*, **795**, 134
- Fomalont, E. B., Geldzahler, B. J., & Bradshaw, C. F. 2001, *ApJ*, **558**, 283
- Frieden, B. R. 1972, *JOSA*, **62**, 511
- Gabuzda, D. C., Kochenov, P. Y., Kollgaard, R. I., & Cawthorne, T. V. 2000, *MNRAS*, **315**, 229
- Ghez, A. M., Salim, S., Weinberg, N. N., et al. 2008, *ApJ*, **689**, 1044
- Gillessen, S., Eisenhauer, F., Trippe, S., et al. 2009, *ApJ*, **692**, 1075
- Gold, R., McKinney, J. C., Johnson, M. D., & Doeleman, S. S. 2017, *ApJ*, **837**, 180
- Gomes, N., Garcia, P. J. V., & Thiébaud, É 2017, *MNRAS*, **465**, 3823
- Högbom, J. A. 1974, *A&AS*, **15**, 417
- Honma, M., Akiyama, K., Uemura, M., & Ikeda, S. 2014, *PASJ*, **66**, 95
- Horn, B. K., & Schunck, B. G. 1981, *Artificial Intelligence*, **17**, 185
- Jeffrey, R. M., Blundell, K. M., Trushkin, S. A., & Mioduszewski, A. J. 2016, *MNRAS*, **461**, 312
- Johannsen, T., & Psaltis, D. 2010, *ApJ*, **718**, 446
- Johnson, M. D. 2016, *ApJ*, **833**, 74
- Johnson, M. D., Fish, V. L., Doeleman, S. S., et al. 2014, *ApJ*, **794**, 150
- Johnson, M. D., Fish, V. L., Doeleman, S. S., et al. 2015a, *Sci*, **350**, 1242
- Johnson, M. D., & Gwinn, C. R. 2015, *ApJ*, **805**, 180
- Johnson, M. D., Loeb, A., Shiokawa, H., Chael, A. A., & Doeleman, S. S. 2015b, *ApJ*, **813**, 132
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python
- Jorstad, S. G., Marscher, A. P., Lister, M. L., et al. 2005, *AJ*, **130**, 1418
- Kardashev, N. S., Novikov, I. D., Lukash, V. N., et al. 2014, *PhyU*, **57**, 1199
- Kellermann, K. I., Lister, M. L., Homan, D. C., et al. 2004, *ApJ*, **609**, 539
- Kellermann, K. I., & Moran, J. M. 2001, *ARA&A*, **39**, 457
- Kim, J., Marrone, D. P., Chan, C.-K., et al. 2016, *ApJ*, **832**, 156
- Lister, M. L., Aller, M. F., Aller, H. D., et al. 2016, *AJ*, **152**, 12
- Lister, M. L., Cohen, M. H., Homan, D. C., et al. 2009, *AJ*, **138**, 1874
- Liu, C. 2009, PhD thesis, Massachusetts Institute of Technology
- Lu, R.-S., Broderick, A. E., Baron, F., et al. 2014, *ApJ*, **788**, 120
- Lu, R.-S., Krichbaum, T. P., Eckart, A., et al. 2011, *A&A*, **525**, A76
- Lu, R.-S., Roelofs, F., Fish, V. L., et al. 2016, *ApJ*, **817**, 173
- Luminet, J.-P. 1979, *A&A*, **75**, 228
- Marrone, D. P., Baganoff, F. K., Morris, M. R., et al. 2008, *ApJ*, **682**, 373
- Marrone, D. P., Moran, J. M., Zhao, J.-H., & Rao, R. 2006, *Journal of Physics Conference Series*, **54**, 354
- Medeiros, L., Chan, C.-k., Ozel, F., et al. 2016, arXiv:1601.06799
- Medeiros, L., Chan, C.-k., Özel, F., et al. 2017, *ApJ*, **844**, 35
- Mertens, F., & Lobanov, A. P. 2016, *A&A*, **587**, A52
- Metzger, B. D., Williams, P. K. G., & Berger, E. 2015, *ApJ*, **806**, 224
- Mioduszewski, A. J., Rupen, M. P., Walker, R. C., Schillemat, K. M., & Taylor, G. B. 2004, *BAAS*, **36**, 967
- Narayan, R., & Nityananda, R. 1986, *ARA&A*, **24**, 127
- Reid, M. J., McClintock, J. E., Steiner, J. F., et al. 2014, *ApJ*, **796**, 2
- Roelofs, F., Johnson, M. D., Shiokawa, H., Doeleman, S. S., & Falcke, H. 2017, *ApJ*, **847**, 55
- Ryle, M. 1962, *Natur*, **194**, 517
- Sault, R. J., Oosterloo, T., Dulk, G. A., & Leblanc, Y. 1997, *A&A*, **324**, 1190
- Shiokawa, H. 2013, PhD thesis, Univ. Illinois at Urbana-Champaign
- Shiokawa, H., Gammie, C. F., & Doeleman, S. S. 2017, *ApJ*, **846**, 29
- Smirnov, A. V., Baryshev, A. M., Pilipenko, S. V., et al. 2012, *Proc. SPIE*, **8442**, 84424C
- Takahashi, R. 2004, *ApJ*, **611**, 996
- Thiébaud, É 2013, in EAS Publications Ser. 59, New Concepts in Imaging: Optical and Statistical Models, ed. D. Mary, C. Theys, & C. Aime (Les Ulis: EDP Sciences), 157
- Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr. 2017, Interferometry and Synthesis in Radio Astronomy (3rd ed.; Hoboken, NJ: John Wiley & Sons)
- Viola, P., & Wells, W. M., III 1997, *International Journal of Computer Vision*, **24**, 137
- Walker, R. C., Hardee, P. E., Davies, F., et al. 2016, *Galax*, **4**, 46
- Wells, W. M., Viola, P., Atsumi, H., Nakajima, S., & Kikinis, R. 1996, in Medical Image Analysis (Vol. 1; Amsterdam: Elsevier), 35
- Wild, W., Kardashev, N. S., Likhachev, S. F., et al. 2009, *ExA*, **23**, 221
- Yusef-Zadeh, F., Bushouse, H., Wardle, M., et al. 2009, *ApJ*, **706**, 348
- Zamaninasab, M., Eckart, A., Witzel, G., et al. 2010, *A&A*, **510**, A3