The Astrophysical Journal, 940:182 (12pp), 2022 December 1
https://doi.org/10.3847/1538-4357/ac9cd2

Measuring the Ellipticity of M87* Images

Paul Tiede[1,2,3,4,5], Avery E. Broderick[3,4,5], Daniel C. M. Palumbo[1,2], and Andrew Chael[6,7]
[1] Black Hole Initiative at Harvard University, 20 Garden Street, Cambridge, MA 02138, USA; ptiede@perimeterinstitute.ca
[2] Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA
[3] Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, ON, N2L 2Y5, Canada
[4] Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, ON, N2L 3G1, Canada
[5] Waterloo Centre for Astrophysics, University of Waterloo, Waterloo, ON N2L 3G1 Canada
[6] Princeton Center for Theoretical Science, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA

Received 2022 August 9; revised 2022 October 18; accepted 2022 October 21; published 2022 December 5

Abstract

The Event Horizon Telescope (EHT) images of the supermassive black hole at the center of the galaxy M87 provided the first image of the accretion environment on horizon scales. General relativity (GR) predicts that the image of the shadow should be nearly circular given the inclination angle of the black hole M87*. A robust detection of ellipticity in image reconstructions of M87* could signal new gravitational physics on horizon scales. Here we analyze whether the imaging parameters used in EHT analyses are sensitive to ring ellipticity, and measure the constraints on the ellipticity of M87*. We find that the top set is unable to recover ellipticity. Even for simple geometric models, the true ellipticity is biased low, preferring circular rings. Therefore, to place a constraint on the ellipticity of M87*, we measure the ellipticity of 550 synthetic data sets produced from GRMHD simulations. We find that images with intrinsic axis ratios of 2:1 are consistent with the ellipticity seen from EHT image reconstructions.

Unified Astronomy Thesaurus concepts: Black hole physics (159); Submillimeter astronomy (1647); Astronomy data analysis (1858); Astronomy image processing (2306)

1. Introduction

The Event Horizon Telescope (EHT) can resolve the emission around the event horizon of the supermassive black hole M87*, and it directly measured its mass for the first time in 2019 (Event Horizon Telescope Collaboration et al. 2019a, 2019b, 2019c, 2019d, 2019e, 2019f, hereafter Paper I, Paper II, Paper III, Paper IV, Paper V, Paper VI, respectively). Using both imaging and modeling techniques, a ringlike emission structure was observed. The measured ring radius is consistent with a central black hole with mass $6.5 \times 10^9 M_\odot$. The direct mass measurement was the first direct observation of the accretion disk on horizon scales. In addition to the size of the ring, which correlates to the mass (Paper VI), the shape or ellipticity of the ring is theoretically interesting.

There are a number of factors that could cause ellipticity in the measured ring. First are the shadow (Falcke et al. 2000) and $n = 1$ photon ring[8] (Gralla et al. 2020; Johnson et al. 2020), both of which are related to the existence of spherical photon orbits around Kerr black holes. Given the low inclination of M87* (Mertens et al. 2016), general relativity (GR) predicts that the observed shadow and photon ring should be highly symmetric. Due to this symmetry, it has been suggested that image ellipticity be used to constrain deviations from GR near the event horizon. If the no-hair theorem breaks down near the event horizon of supermassive black holes, the shadow may appear more elliptic (see, e.g., Johannsen & Psaltis 2010; Broderick et al. 2014; Johannsen et al. 2016; Medeiros et al. 2020).

A second source of noncircularity could be the imprint of the horizon on the image and has been called the *inner shadow* (Dokuchaev & Nazarova 2019; Chael et al. 2021). The inner shadow occurs from light rays that do not pass through the equatorial plane before hitting the horizon. Generally, the emission from the black hole must be concentrated in the equatorial plane for the inner shadow to be visible. For instance, magnetically arrested disk general relativistic magnetohydrodynamic (GRMHD) simulations, which are preferred for M87* (Event Horizon Telescope Collaboration et al. 2021), tend to display the inner shadow feature (Chael et al. 2021). Interestingly, the location of the inner shadow relative to the photon ring is a function of spin and inclination. During image reconstructions, the displacement of the inner shadow could manifest as a source of noncircularity in the image reconstruction.

A third, more mundane origin of noncircularity is the accretion flow itself. The accretion flow is expected to be highly turbulent and can cause extended features in the image, creating highly elliptical ring reconstructions. Therefore, any measurement of ring ellipticity needs to account for the accretion through calibration or other means.

To constrain the ellipticity of M87* there are two forms of uncertainty that need to be considered: astrophysical/accretion noise and image uncertainty. While the EHT can resolve the horizon-scale structure, its dynamic range and visibility coverage are poor (Paper II, Paper III). Therefore, there are infinitely many images that can reproduce observations. This uncertainty makes measuring the ellipticity of the ring uncertain, and requires measuring an ensemble of images to quantify how well we can constrain ellipticity. In Paper IV, regularized maximum likelihood (RML) methods were applied to the M87* data. RML introduces additional assumptions through regularizers that enforce features such as image smoothness (Bouman et al. 2016; Chael et al. 2016;

---

[7] NASA Hubble Fellowship Program Einstein Fellow.

[8] We follow the Johnson et al. (2020) definition, where $n = 1$ corresponds to photons that have made a half-orbit around the black hole.

Kuramochi et al. 2018), sparseness (Wiaux et al. 2009a, 2009b; Honma et al. 2014; Akiyama et al. 2017), and similarity to some fiducial image (Narayan & Nityananda 1986). However, how best to choose the weights and functional forms of these regularizers is unknown. To combat this problem, Paper IV performed a parameter survey using simulated data sets similar in quality to the 2017 M87* data. The parameter surveys were then used to assess the relative performance of the different regularizer weights and cuts were placed on different choices based on reconstructed image fidelity metrics to give the optimal "top set." The final set of images was considered the set on which feature extraction should be applied and ring ellipticity should be measured. We note that while these images do give a distribution of images, they do not form a posterior. As a result, interpreting quantitative measurements often requires some a priori calibration to assess the reliability.

In Paper VI, a preliminary attempt at measuring the ellipticity of image reconstructions of M87* was presented in Figure 18. To measure the ellipticity the authors measured the fractional spread in the ring radius from the reconstructions and found it was similar to that of the intrinsic images having an axis ratio of 4:3. This axis ratio suggests that the M87* images were highly symmetric. However, there are two potential issues. First, the parameter survey used in Paper IV did not include an elliptic ring in the simulated tests. Therefore, the reliability of the ellipticity measured in the images is unclear. Second, the results presented in Figure 18 of Paper IV were based not on the measured ellipticity of the reconstructed GRMHD images, but rather on that of the actual true image blurred to 20 μas. Whether the imaging parameter survey used in Paper IV can reliably recover ellipticity, and how to interpret the measured ellipticity of M87* are then open questions.

In this paper, we assess whether the M87* top set used in Paper IV can constrain the ellipticity of the M87* shadow across a set of simulated data tests. To evaluate whether the surveys can reliably recover image ellipticity, we consider a new set of geometric tests specifically targeted to measure ellipticity. We then use the image feature extraction tools `REx` (Chael 2019) and `VIDA` (Tiede et al. 2022) to measure the ellipticity of the reconstructions. In the first part of the paper we show that the measured ellipticity of image reconstructions does not reliably recover the true on-sky value. To overcome this issue, we then run the imaging surveys on a set of GRMHD simulations to calibrate for the imaging ellipticity bias. This is similar in spirit to the mass calibration procedure done in Paper VI, and is necessary to interpret the ellipticity results from the imaging pipelines.

The layout of the paper is as follows: In Section 2 we review the M87* imaging top set, feature extraction techniques, and M87* results. In Section 3 we explore whether the current parameter survey used in M87* is able to recover an elliptical geometric ring. Afterward, we use a set of GRMHD simulations to calibrate the image reconstruction ellipticity of the M87* results. In Section 5 we review the results and provide an ellipticity constraint for M87*.

2. Background

This section will review the standard image reconstruction techniques used by the EHT in Paper IV. The imaging techniques used in this paper will be identical to the `eht-imaging` pipeline used in Paper IV. We will then review the two feature extraction techniques used in this paper, `REx` and `VIDA`. These feature extraction techniques are needed since imaging is nonparametric. Therefore, an additional processing step is needed to extract image features of interest, e.g., ring ellipticity. Finally, we will apply `REx` and `VIDA` to the M87* top set. We will reproduce the results from Paper VI and extend the analysis to include the orientation of M87*'s ellipticity.

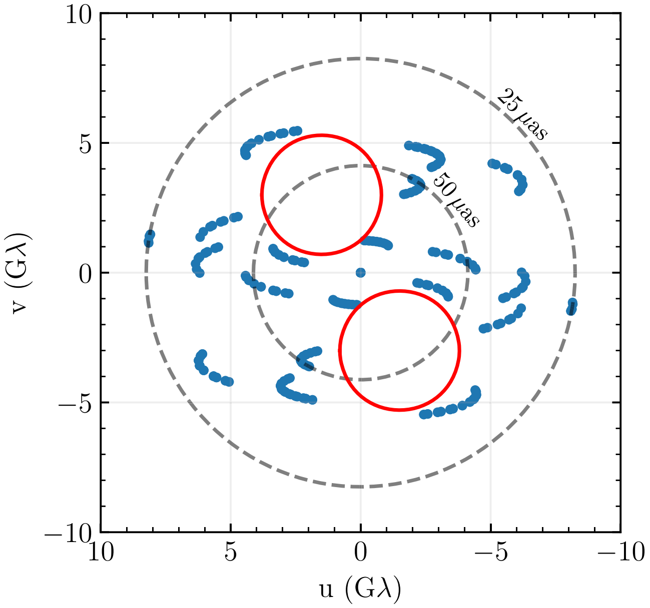

Figure 1. UV coverage of the EHT 2017 array on April 11. The blue dots show where the 2017 array samples in the UV plane in units of Gλ. The gray dashed circles show the characteristic location of the image features of 50 μas and 25 μas in the UV domain. The red circles highlight the coverage gap in M87* in the north–south direction.

2.1. Image Reconstructions and the M87* Top Set

In this paper, we will focus on the RML methods used in Paper IV. These methods attempt to make imaging tractable by forward-modeling the image, I, and minimizing the objective function:

$$J(I) = \sum_{\text{data}} \alpha_d \chi_d^2(I) - \sum_{\text{regularizers}} \beta_r S_r(I). \tag{1}$$

Following Paper IV, each χ^2 is defined solely from the data products from the EHT, e.g., complex visibilities. The second term encapsulates our additional assumptions or regularizers that are placed on the image. The variables α_d and β_r are the "hyperparameters" that control the relative weighting of the data products and regularizers. For a list of the regularizers used, see Paper IV.

The regularizers are important for the EHT, given its sparse coverage (see Figure 1) and poor dynamic range (Paper II, Paper III). Unfortunately, there is no canonical statistical framework[9] for how to choose the relative weights α_d and β_r. Instead, a series of heuristics and data quality metrics across a parameter survey of different regularizers are employed. In Paper IV, a survey of different hyperparameters was performed for each imaging method: the DIFMAP, the `SMILI`, and the `eht-imaging` pipeline. This paper will focus on the

---

[9] See Akiyama et al. (2017) for an approach using cross-validation that requires more data than currently available for the EHT.

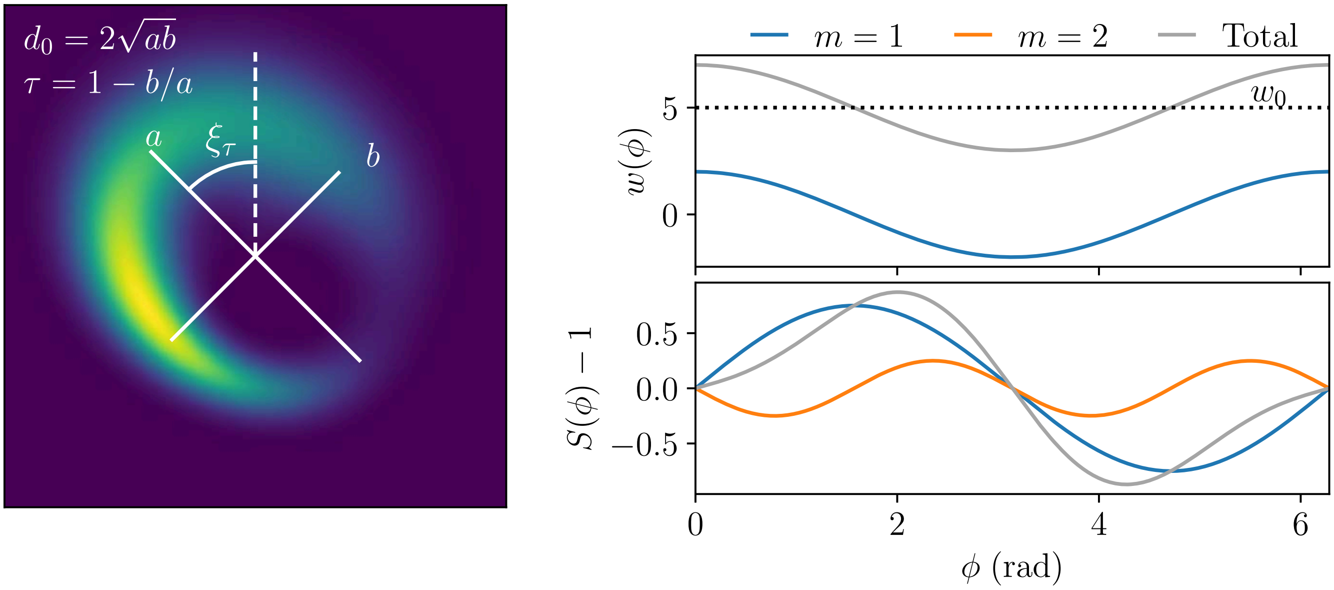

Figure 2. A visualization of the `VIDA` `CosineRing`{1, 2} template with parameters $d_0 = 42$ μas, $w_0 = 11.8$ μas, $w_1 = 4.7$ μas, $\xi_1^{(w)} = 0$, $\tau = 0.2$, $\xi_\tau = \pi/4$, $s_1 = 0.75$, $s_2 = -0.25$, $\xi_1^{(s)} = \pi/2$, and $\xi_2^{(2)} = -\pi/4$. The left panel shows the visual appearance of the template normalized to have unit flux. The top right panel is the width profile as a function of the azimuthal angle of the $N = 1$ expansion in Equation (2), where the black dotted line gives the ring width w_0. The bottom right panel is the brightness profile as a function of the azimuthal angle of the $M = 2$ expansion given in Equation (3).

`eht-imaging` pipeline, but we have found similar results for `SMILI`. The parameter surveys are run on a set of simple geometric models—a ring, a crescent, a disk, and a pair of small Gaussians—to find an appropriate set of hyperparameters for the M87* observations. These synthetic models are constructed with angular scales that approximately mimic the scale inferred from the EHT, thereby providing a test of whether the pipelines can differentiate complex structures at equivalent angular scales.

Every combination of parameters in the survey is mapped to an effective resolution on each geometric model by interpolating normalized cross-correlations to the true image. The median performance across all M87* observation days is used. This process enables comparisons of parameter performance across models; the effective resolution for each parameter combination is then averaged across models and used to rank all sets of parameters. The sets of parameters that produce an average effective resolution better than the EHT nominal resolution are used to form a "top set." This top set is the effective set of parameters that we will use to construct our ensembles of image reconstructions.

2.2. Review of Feature Extraction Techniques

2.2.1. Variational Image Domain Analysis

Variational image domain analysis (VIDA) (Tiede et al. 2022) requires three ingredients: an image I, a template image f_θ, and a probability divergence $\mathcal{D}$ that measures the difference between f and I. `VIDA` relies on the template function f_θ being a reasonable approximation to the true image. Given that we are interested in ring morphologies, we will use the `CosineRing`{N, M} template (see Figure 2 for a visualization).[10] The `CosineRing`{N, M} template is an elliptical Gaussian ring template whose azimuthal brightness and thickness are described by a cosine expansion. More specifically, the model is described by

1. a diameter $d_0 = \sqrt{ab}$, where a and b are the semiminor and semimajor axes, respectively;
2. a thickness function

$$w(\phi) = w_0 + \sum_{n=1}^{N} w_n \cos\left[n(\phi - \xi_m^{(w)})\right]; \tag{2}$$

3. an ellipticity $\tau = 1 - b/a$, with the orientation ξ_τ measured east of north;
4. a slash function

$$S(\phi) = 1 - \sum_{m=1}^{M} s_m \cos\left[m(\phi - \xi_m^{(s)})\right]; \tag{3}$$

5. x_0, y_0, as the center of the ring.

How to pick N and M is left to the user and the imaging problem at hand. Section 3 is interested in recovering the profile for a simple elliptical ring with a slash. Therefore, we will take $N = 0$ and $M = 1$. We also try higher-order mode expansions and find that they are typically much smaller than the first mode and do not change the results. For the GRMHD reconstructions in Section 4 we take $N = 1$, $M = 4$, given their complicated azimuthal structure. This template has 16 parameters in total. Typically, when we refer to the thickness or slash strength of the template, we are referring to the w_0 or s_1 parameter, respectively. In addition to the ring template, we include a constant-intensity background, where the background intensity is also a parameter. This intensity floor models the diffuse intensity that is typically deposited across the image in reconstructions.

The optimal template is found by minimizing the divergence between the image reconstruction, normalized to the unit flux, and the specified template. We will use the Bhattacharyya

---

[10] For information about other templates present in `VIDA`, please see https://github.com/ptiede/VIDA.jl.

divergence,

$$\mathrm{Bh}(f_\theta||\mathcal{I}) = -\log\int\sqrt{f_\theta(x, y)\mathcal{I}(x, y)}\,dxdy, \tag{4}$$

in all analyses below. To minimize the Bh divergence we use the Julia package BlackBoxOptim.jl,[11] which uses an adaptive genetic algorithm for optimization. For more information on the validation of `VIDA` and the optimization strategy see Tiede et al. (2022).

2.2.2. `REx`

The other image feature extraction method used in this paper is the *ring extractor* or `REx` algorithm used in Paper IV and described in detail in Chael (2019). The first step in `REx` (see Paper IV for details) is to identify the dominant ring in the image. First, for each pixel x_i, y_j in the image, an intensity map $I(r, \theta|x_i, y_j)$ with radius and angle is defined relative to x_i, y_j. For each map, the "radius" of the image is given by

$$\begin{aligned} r_{\mathrm{pk}}(\theta|x, y) &= \mathrm{argmax}[I(r, \theta|x, y)]_r \\ \bar{r}_{\mathrm{pk}} &= \langle r_{\mathrm{pk}}(\theta|x, y)\rangle_{\theta\in[0,2\pi]}. \end{aligned} \tag{5}$$

This provides a radius for each pixel center. The optimal center is then selected according to

$$(x_0,\ y_0) = \mathrm{argmin}\left[\frac{\sigma_{\bar{r}(x,y)}}{\bar{r}_{\mathrm{pk}}(x, y)}\right], \tag{6}$$

where $\sigma_{\bar{r}}(x, y) = \langle(r_{\mathrm{pk}}(\theta|x, y) - \bar{r}_{\mathrm{pk}})^2\rangle$ is the radial dispersion. This specifies the image center and all future quantities will be defined relative to this center.

The diameter of the image is

$$d = 2\bar{r}_{\mathrm{pk}}(x_0, y_0). \tag{7}$$

Following Paper VI, `REx` characterizes the ellipticity of the ring structure by the radial fractional dispersion:

$$f_d = \frac{\sigma_{\bar{r}}}{\bar{r}_{\mathrm{pk}}}. \tag{8}$$

The width of the ring is defined by finding the FWHM at a fixed θ ray, and then averaging over θ:

$$w = \langle\mathrm{FWHM}_r[I(r, \theta|x_0, y_0) - I_{\mathrm{floor}}]\rangle_\theta. \tag{9}$$

The intensity floor is given by $I_{\mathrm{floor}} = \langle I(r = 50\mu\mathrm{as}, \theta)\rangle_\theta$ and is included to avoid biasing the measurement due to the low-level intensity present in the image. This is similar to including the constant-intensity template during the `VIDA` extraction.

To characterize the azimuthal profile of the ring (ξ_s and s for `VIDA`) we consider the azimuthal moments of the ring. As in `VIDA` we will only be interested in the first azimuthal moment. The orientation, ξ_s, of the first moment is given by

$$\xi_s = \left\langle \mathrm{Arg}\left[\int_0^{2\pi} I(r, \theta|x_0, y_0)e^{i\theta}d\theta\right]\right\rangle_{r\in[r_{\mathrm{in}},r_{\mathrm{out}}]}, \tag{10}$$

where $r_{\mathrm{in}} = (d - w)/2$ and $r_{\mathrm{out}} = (d + w)/2$. The coefficient of the first moment or slash is given by

$$s = 2\left\langle \frac{\left|\int_0^{2\pi} I(r, \theta|x_0, y_0)e^{i\theta}d\theta\right|}{\int_0^{2\pi} I(r, \theta|x_0, y_0)d\theta}\right\rangle. \tag{11}$$

2.2.3. Relating `VIDA` and `REx` Parameters

`VIDA` and `REx` assume different parameterizations of the intrinsic structure, and use different optimization strategies. Therefore, we expect the resulting distributions to differ slightly. One prominent difference between `VIDA` and `REx` is how they parameterize ellipticity. To compare both `VIDA`'s and `REx`'s ellipticity measurements, we need to relate `REx`'s fractional dispersion f_d to `VIDA`'s ellipticity τ. To accomplish this we will convert `REx`'s fractional diameter spread f_d into `VIDA`'s τ. First, we consider an ellipse with semimajor axis a and semiminor axis b. Then `VIDA` parameterizes this ellipse with $d_0 = 2r_0 = 2\sqrt{ab}$ and $\tau = 1 - b/a$. It can then be shown (see Tiede et al. 2022, Appendix A, for a derivation) that the fractional dispersion is related to ellipticity τ through

$$f_d(\tau) = \frac{\sqrt{1 - \epsilon(\tau)^2 - 4/\pi^2E^2(\epsilon(\tau))}}{\sqrt{1 - \tau}}, \tag{12}$$

where $E(x)$ is the complete elliptic integral of the second kind and $\epsilon(\tau) = \sqrt{1 - (1 - \tau)^2}$ is the orbital eccentricity. Using linear interpolation we invert the function achieving a map from f_d to τ.

This relation assumes that the image feature is a perfect ellipse. Given that `REx` defines the radius of the ring in terms of its peak, we expect that the shape will not be a perfect ellipse. This additional nonellipticity will cause `REx`'s ellipticity to be typically greater than `VIDA`'s, which we empirically find below.

2.3. Review of M87* Ellipticity Measurement

Figure 18 of Paper VI showed the measured M87* fractional deviation from `REx` applied to the `eht-imaging` top-set parameters. To reproduce these results, we apply `VIDA` and `REx` to the `eht-imaging` top set. The results are shown in Figure 3. The top row shows the fiducial image reconstruction from the `eht-imaging` top set across each observation day. The ellipticity of the top-set images is shown in the middle row. We find identical results to those in Paper VI for `REx`. The `VIDA` results are systematically lower than the `REx` results, as expected from the discussion in the previous section. The bottom row presents, for the first time, the orientation of this ellipticity from `VIDA`. Note that `REx` cannot currently measure this orientation. Overall, we find that the ellipticity measurement is stable across all four days, giving $\tau = 0.05$–0.2, and orientation $\xi_\tau = -75°$–$0°$ east of north.

While it is interesting that the ellipticity measurements are consistent across days, it is not clear whether this result is intrinsic to the source. In Tiede et al. (2022), we found statistically identical ellipticities and orientations for the symmetric crescent models and GRMHD models. Given that the crescent models are symmetric, the measured ellipticity may be an imaging artifact. Furthermore, the ellipticity orientation does align with a coverage gap (see the red circles in Figure 1). Finally, since the top set used for M87* did not

---

[11] https://github.com/robertfeldt/BlackBoxOptim.jl

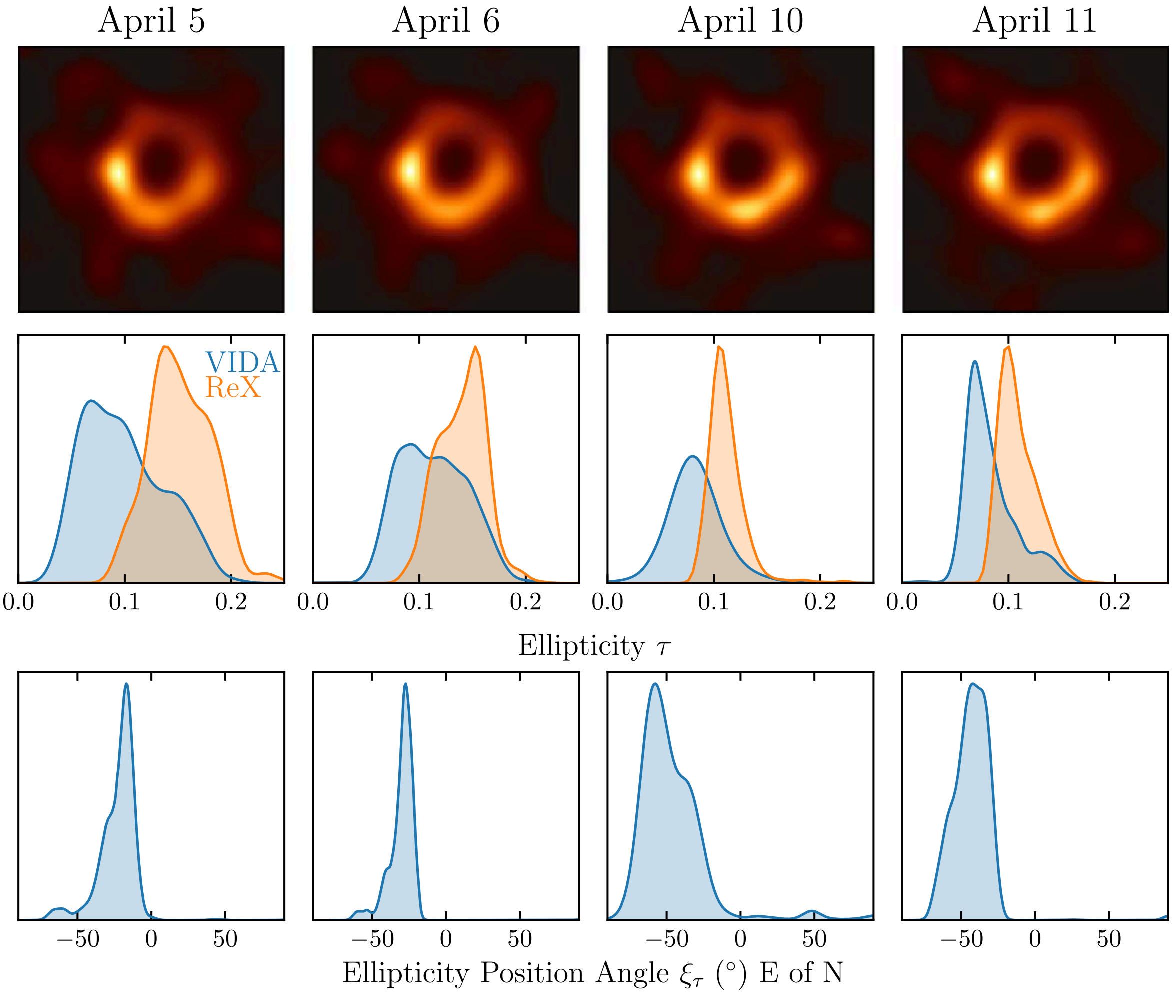

Figure 3. The `eht-imaging` reconstructions and ellipticity features of M87* from the Paper IV top set across the different observations during the 2017 EHT campaign. The top row shows the fiducial `eht-imaging` image from Paper IV. The middle row shows the measured ellipticity of the `eht-imaging` top set from `VIDA` (blue) and `REx` (orange). The bottom row shows the measured ellipticity orientation angle east of north from `VIDA`. Note that `REx` is not currently able to measure the orientation angle. Overall we see that M87* has consistent ellipticity around $\tau = 0.05$–0.2, around −50°–0° east of north.

include an elliptical ring, it is not clear whether ellipticity can reliably be recovered. To investigate the fidelity of ellipticity reconstruction using the reported top set, we will apply it to simple elliptical ring images in the next section.

3. Geometric Test

One of the potential issues with the M87* top set is that no elliptical rings were included in the simulated data tests. Given that the identification of the top set was defined by its performance on simulated data tests, the top set may not accurately recover ring ellipticity even in simple cases. This section will analyze simulated data from an elliptical ring model using the M87* top set.

3.1. Elliptical Image Test

For the elliptical image, we use the `CosineRing{0, 1}` template described in Section 2.2.1, with parameters $d_0 = 37.56\,\mu$as, $w = 7.9\,\mu$as, $\tau = 0.187$, and $s = 0.5$. The ring flux is set to 0.6 Jy, matching the measured compact flux of M87*. We also align the orientation of the slash and ellipse, i.e., we set $\xi_s = \xi_\tau = \xi$. To test the impact of different orientations of the ellipticity we consider $\xi = 0°$–$360°$ in steps of 45°. A subset of the ground-truth images are shown in the top row of Figure 4. For each rotated ring we use the `eht-imaging` top-set pipeline from Paper IV to create 1572 reconstructions.

3.2. Geometric Results

Given the elliptical ring reconstructions, we use `VIDA` and `REx` to extract the relevant image features. Since the `VIDA` template is identical to the on-sky image, we expect that the true parameter value will lie within the distribution of recovered image features. The results for each orientation are summarized in Table 1. The ring diameter and width are consistent across the rotation angles, and are consistent with the truth. The slash and its orientation are similarly recovered.

However, the ellipticity, τ, is significantly biased when $\xi_\tau = 90°$, i.e., when the semimajor axis of the ellipse is aligned in the east–west direction. Furthermore, looking at the bottom row of Figure 4, we see that the orientation of the ellipticity is consistently biased toward $\xi_\tau = 0$, i.e., the north–south direction. This bias can be visually confirmed by looking at the reconstructions, e.g., the middle row of Figure 4. Furthermore, we see that the true ellipticity and orientation are only recovered in the $\xi = 0$ case.

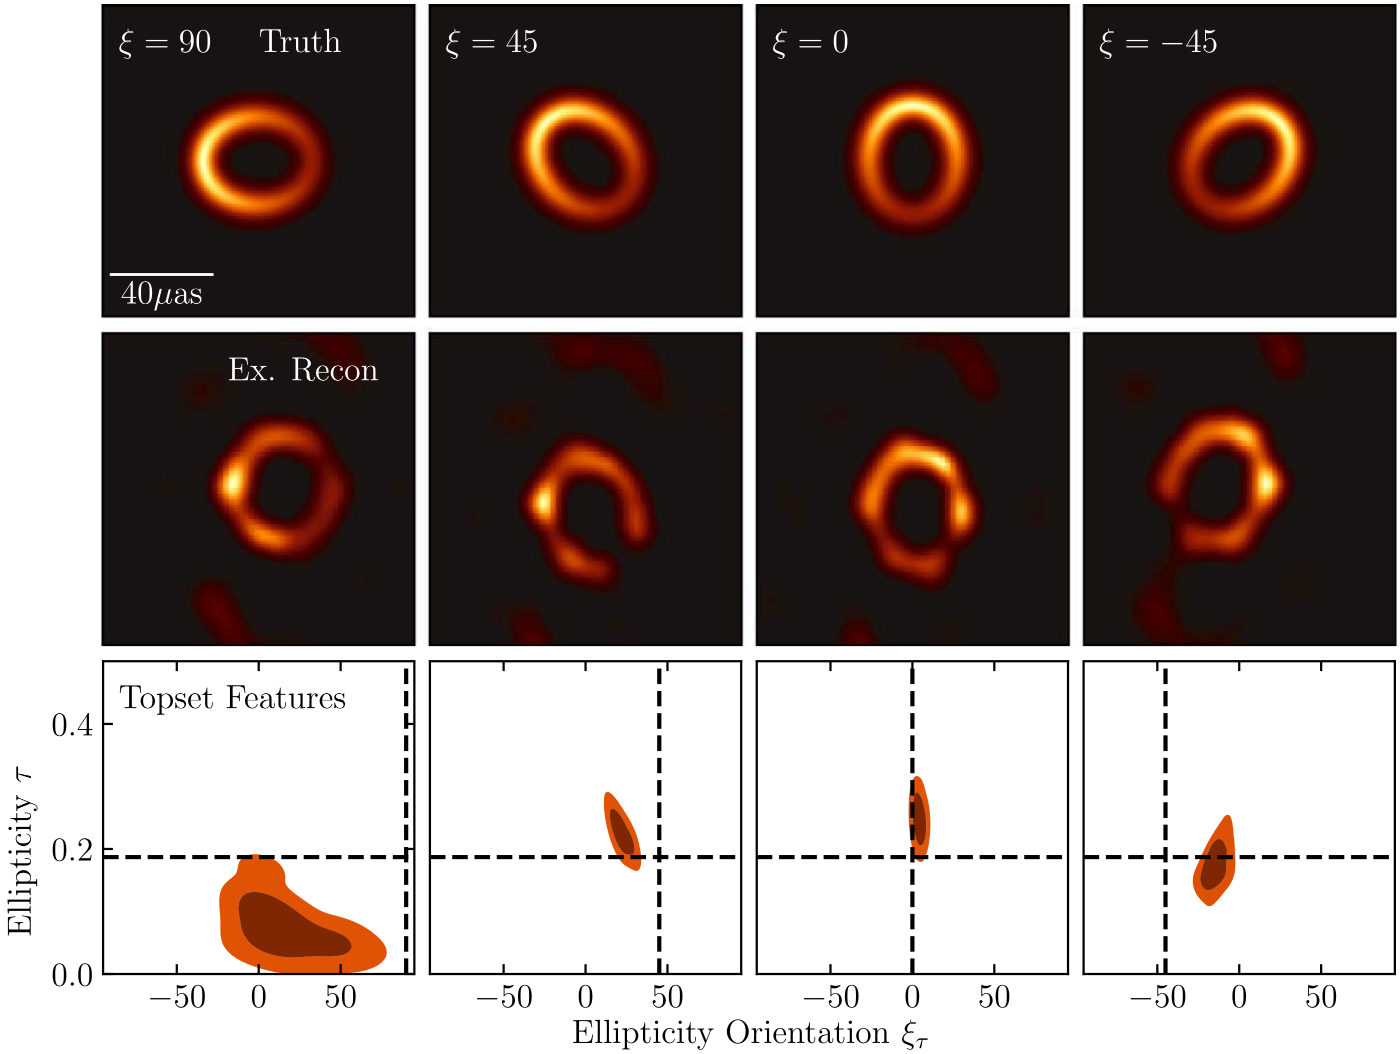

Figure 4. Examples of image reconstructions and `VIDA` fits for different elliptical rings. The top row shows the truth images at position angles $\xi_s = 0°, 45°, 90°, 135°$ north of east. The middle row shows an example reconstruction from the M87* top set for each ring orientation. The bottom row shows the `VIDA` results for the ellipticity and its position angle from the top set. We find significant bias in τ and ξ_τ for rings whose semimajor axis is aligned in the east–west direction.

Therefore, it appears that imaging creates a preferred ellipticity direction $\xi_\tau \approx 0$, τ. That is, as τ increases, the ellipticity orientation tends to point in the north–south direction. This orientation does approximately align with the large gap in EHT coverage for M87* (see the red circles in Figure 1).

Given the inability of the top-set images to faithfully recover the ellipticity in our geometric tests, it is not clear that the measurement found in Figure 3 represents a constraint on the intrinsic ellipticity and not simply an artifact of the top set itself. To address this possibility, we will calibrate the procedure applied above to a large number of GRMHD simulations.

4. Calibrating the M87* Ellipticity Measurement

To calibrate for the uncertainty in the M87* top-set ellipticity, we will use a similar procedure to the mass calibration done in Paper VI. First, we select a number of GRMHD simulations from the Paper V library to provide physically motivated images. Second, we construct simulated data matching the 2017 M87* observations and reconstruct the images using the M87* `eht-imaging` top set. Using this set of reconstructions, we measure the ellipticity of each reconstruction and compare it to the ground truth constructing a "theoretical" uncertainty budget for the ellipticity. This uncertainty is then included in the uncertainty of M87*'s ellipticity.

4.1. Scaled Set

To construct the GRMHD images used in this paper, we first cut the simulations from Paper V based on whether the total jet power is consistent with the observed jet power of M87 (Table 2 of Paper V). From the remaining set of simulations, we randomly select 100 snapshots and randomly assign them to a 2017 M87* observation day. Each selected snapshot is rescaled to its best-fit value (according to the average image-scoring results of Paper VI), and randomly rotated. To include the effects of the mass uncertainty of M87*, we then further scale the intrinsic image by a factor of 0.8, 0.9, 1.1, and 1.2 in both the x- and y-direction. The net result is 500 images uniformly sampled over days and orientations and gridded in mass relative to the M87* best-fit value. We will refer to this list as the *scaled set.*

Table 1
Recovered Parameters for the Slashed Elliptical Ring Test Set

Image			d_0	w	τ	ξ_τ (deg)	s	ξ_s (deg)
$\xi_s = 90°$	$\xi_\tau = 90°$	`REx`	$37.3^{+1.2}_{-1.0}$	$12.4^{+4.4}_{-4.4}$	$0.09^{+0.07}_{-0.04}$	...	$0.54^{+0.07}_{-0.09}$	$90.68^{+3.9}_{-7.8}$
		`VIDA`	$37.5^{+1.3}_{-1.2}$	$11.1^{+3.5}_{-4.0}$	$0.07^{+0.09}_{-0.05}$	$12.6^{+48.8}_{-22.8}$	$0.54^{+0.05}_{-0.07}$	$91.4^{+3.9}_{-10.4}$
$\xi_s = 45°$	$\xi_\tau = 45°$	`REx`	$37.2^{+0.8}_{-1.2}$	$11.2^{+4.7}_{-3.6}$	$0.20^{+0.04}_{-0.04}$	...	$0.52^{+0.11}_{-0.12}$	$48.9^{+5.3}_{-7.8}$
		`VIDA`	$37.7^{+0.9}_{-1.4}$	$10.0^{+3.9}_{-3.4}$	$0.23^{+0.04}_{-0.03}$	$21.9^{+8.1}_{-7.2}$	$0.52^{+0.18}_{-0.12}$	$48.0^{+6.5}_{-9.8}$
$\xi_s = 0°$	$\xi_\tau = 0°$	`REx`	$37.0^{+0.9}_{-1.3}$	$11.2^{+4.3}_{-3.7}$	$0.24^{+0.05}_{-0.03}$	...	$0.45^{+0.15}_{-0.07}$	$-5.7^{+33.3}_{-20.6}$
		`VIDA`	$37.5^{+0.7}_{-1.3}$	$9.8^{+3.8}_{-3.3}$	$0.25^{+0.04}_{-0.04}$	$3.4^{+5.3}_{-3.6}$	$0.45^{+0.13}_{-0.07}$	$-3.4^{+27.1}_{-18.0}$
$\xi_s = -45°$	$\xi_\tau = -45°$	`REx`	$37.5^{+1.1}_{-1.1}$	$11.7^{+4.3}_{-3.9}$	$0.18^{+0.04}_{-0.04}$	...	$0.52^{+0.09}_{-0.09}$	$-49.4^{+22.1}_{-5.6}$
		`VIDA`	$38.3^{+0.9}_{-1.3}$	$10.5^{+3.5}_{-3.5}$	$0.17^{+0.06}_{-0.05}$	$-15.8^{+8.8}_{-9.2}$	$0.53^{+0.13}_{-0.04}$	$-50.7^{+19.8}_{-6.0}$
$\xi_s = -90°$	$\xi_\tau = 90°$	`REx`	$37.5^{+1.1}_{-1.2}$	$12.4^{+4.4}_{-4.4}$	$0.09^{+0.07}_{-0.05}$	...	$0.54^{+0.07}_{-0.12}$	$-90.3^{+6.4}_{-6.6}$
		`VIDA`	$37.8^{+1.2}_{-1.2}$	$11.1^{+3.5}_{-3.9}$	$0.06^{+0.09}_{-0.05}$	$12.6^{+42.4}_{-18.6}$	$0.54^{+0.06}_{-0.08}$	$-90.0^{+6.2}_{-8.6}$
$\xi = -135°$	$\xi_\tau = 45°$	`REx`	$37.1^{+0.7}_{-1.2}$	$11.2^{+4.6}_{-3.6}$	$0.20^{+0.15}_{-0.03}$	...	$0.47^{+0.12}_{-0.09}$	$-131.4^{+5.4}_{-6.0}$
		`VIDA`	$37.6^{+0.8}_{-1.2}$	$9.9^{+3.8}_{-3.5}$	$0.22^{+0.04}_{-0.03}$	$21.3^{+8.8}_{-7.1}$	$0.52^{+0.19}_{-0.09}$	$-134.4^{+4.4}_{-7.7}$
$\xi = -180°$	$\xi_\tau = 0°$	`REx`	$37.0^{+0.9}_{-1.3}$	$11.1^{+4.2}_{-3.7}$	$0.24^{+0.04}_{-0.04}$	...	$0.47^{+0.16}_{-0.08}$	$-184.1.2^{+20.4}_{-27.0}$
		`VIDA`	$37.6^{+0.8}_{-1.2}$	$9.8^{+3.9}_{-3.3}$	$0.25^{+0.03}_{-0.03}$	$4.3^{+3.7}_{-3.5}$	$0.49^{+0.16}_{-0.07}$	$-183.1^{+32.6}_{-12.9}$
$\xi = 135°$	$\xi_\tau = -45°$	`REx`	$37.6^{+1.2}_{-1.1}$	$11.9^{+4.0}_{-4.3}$	$0.18^{+0.04}_{-0.03}$	...	$0.53^{+0.08}_{-0.05}$	$127.9^{+8.3}_{-5.3}$
		`VIDA`	$38.3^{+0.8}_{-1.2}$	$10.5^{+3.9}_{-3.3}$	$0.17^{+0.03}_{-0.03}$	$-13.2^{+3.7}_{-3.5}$	$0.56^{+0.16}_{-0.07}$	$127.9^{+7.5}_{-6.1}$
Truth			37.56	7.9	0.187	...	0.5	...

Note. The parameters are the median value and the 95% interval around the median. `REx` and `VIDA` give very similar results for all parameters, although no results for ξ_τ are given for `REx` since it cannot recover it. All the parameters except the ellipticity τ and its orientation ξ_τ contain the true values.

4.2. Stretched Set

While the scaled set measures the expected ellipticity due to imaging and accretion turbulence, it does not measure how sensitive imaging is to additional intrinsic ellipticity that may occur from, e.g., non-GR spacetimes. To assess the ability to measure an elliptical shadow, we randomly select 10 additional GRMHD snapshots that fit M87's jet power. For each image we scale it to its best-fit mass, and randomly rotate it. Ellipticity is added by picking two random orthogonal directions in the image r_x and r_y and applying the transformation $r_x \to \alpha r_x$ and $r_y \to r_y/\alpha$. To create different amounts of ellipticity we let $\alpha = 0.8$, 0.9, 1.0, 1.1, 1.2, giving a $\tau = 0.36$, 0.19, 0.0, 0.17, 0.31, respectively. We will refer to this as the *stretched set* in the remainder of the paper.

4.3. Removing Failed Top-set Reconstructions

While inspecting the top-set reconstructions of the scaled and stretched GRMHD simulations we notice that a large number of images fail to show a ringlike feature. Instead the image reconstructions have intensity distributed across the image in a pattern similar to that of the EHT dirty beam. This is commonly known as "waffling" and is symptomatic of a poorly chosen set of hyperparameters. For these reconstructions, `VIDA` and `REx` would give nonsensical results since no dominant ring feature exists. To remove this bias these reconstructions need to be removed. Unfortunately, no single set of hyperparameters is identified to have caused the waffling. Therefore, we turn to machine-learning techniques to remove any waffled images.

We create a deep convolutional neural network (CNN) implemented in `Flux.jl` (Innes et al. 2018) to classify and remove the waffled images. More details about the network and image classification are given in the Appendix. The trained neural network outputs a number γ between 0 and 1 that measures the confidence that the image has waffled. We then cut any image reconstructions where $\gamma > \gamma_{\rm thresh} = 0.1$. This threshold cuts 42% of the images in the scaled set and 21% of the images in the stretched set. Some image classification examples are shown in Figure 8. The impact of the value of $\gamma_{\rm thresh}$ is shown in Figure 9. As an additional test of the network, we run the classifier on the elliptical Gaussian reconstructions. We find that only ∼5% of the elliptical images are cut, which is consistent with our visual inspection.

4.4. Scaled Set Results

To analyze the scaled set of images, we use `VIDA`'s `CosineRing{1,4}` template and Bh divergence. Using the CNN image classifier (Section 4.3) we remove the "bad" reconstructions, leaving 454,888 images. To compare the results to the GRMHD simulations' ellipticity we first blur the ground-truth snapshots with a Gaussian kernel with FWHM 15 μas to model the finite resolution of the EHT array. Then we fit the blurred images with the `VIDA CosineRing{1,4}` template. This forms our ground-truth ellipticity to which we will compare all results below.

The results for ξ_τ and ellipticity τ for the entire scaled set are shown in the upper left panel of Figure 5. We find that the ellipticity is quite uncertain in both `VIDA` and `REx` extending to $\tau = 0.3$, which is approximately a 5:3 axis ratio. This ellipticity is larger than the ground-truth ellipticity measured from `VIDA` (right panel of Figure 5), which is quite concentrated at $\tau \approx 0.1$ and extends up to $\tau = 0.25$. As mentioned in Section 1, the spin of the black hole is expected to add a small amount of ellipticity to the on-sky image. We find no evidence for any ellipticity–spin

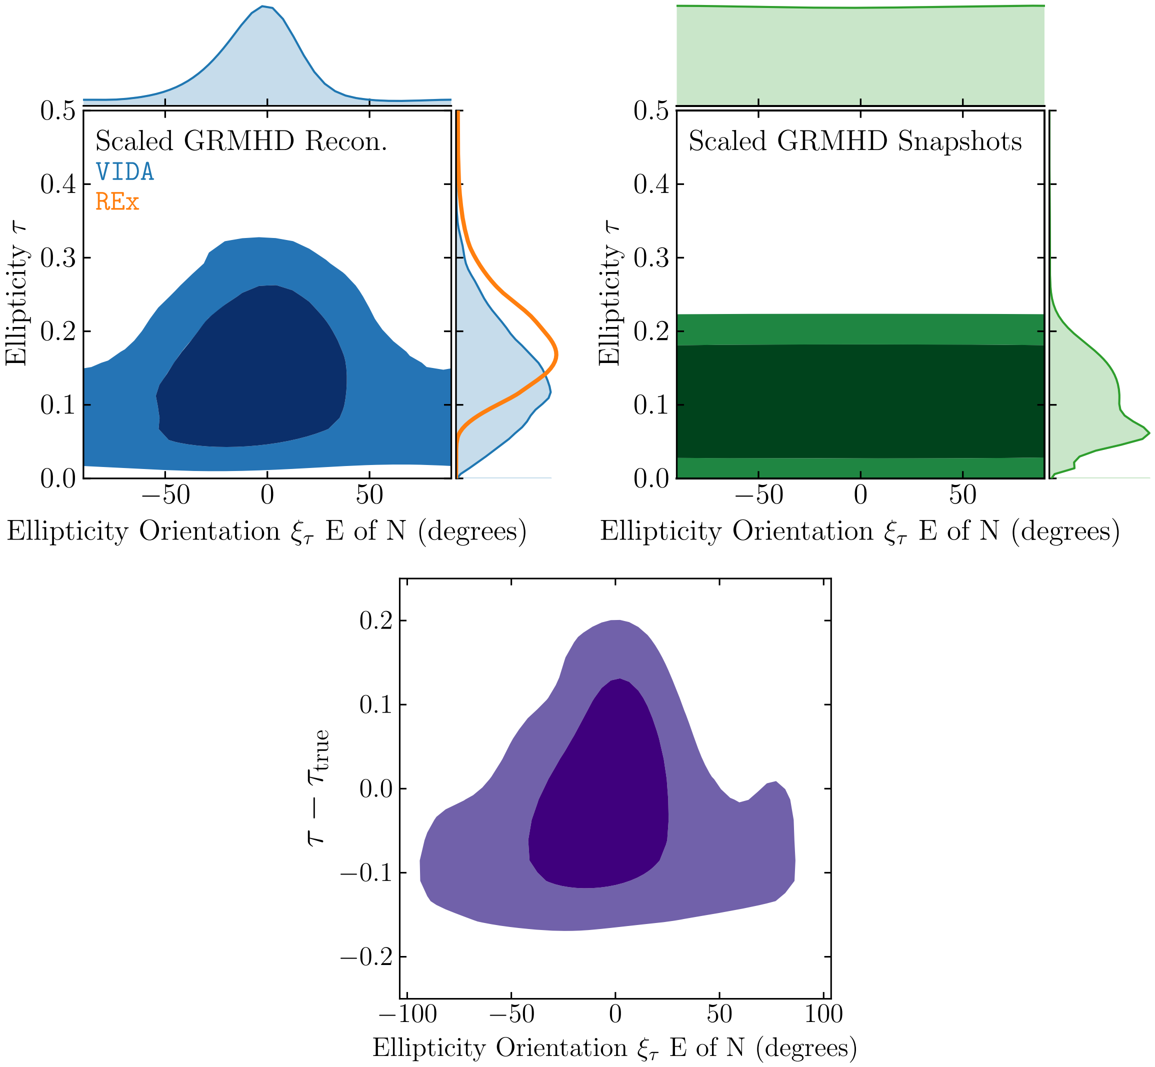

Figure 5. Joint marginal distribution between τ and ξ_τ, where the contours are the 68% and 95% regions. The upper left panel shows the results for the entire reconstructed scaled GRMHD set that satisfies the $a_{\rm waffle} < a_{\rm thresh} = 0.1$ threshold. The upper right panel is when `VIDA` is applied directly to the GRMHD snapshot blurred with a 15 μas FWHM Gaussian kernel. As in the geometric results, we see a preference for $\xi_\tau \approx 0°$, i.e., the north–south direction, regardless of the intrinsic image distribution. The bottom figure shows the measured ellipticity orientation on the x-axis with the measured ellipticity residual distribution.

correlation. The lack of correlation provides evidence that the dominant source of ellipticity in the reconstructions is the accretion flow itself.

Focusing on the ellipticity orientation ξ_τ, we find the ellipticity is strongly biased in the north–south direction. Furthermore, when $\xi_\tau \approx 90°$, τ extends to its highest values. This orientation distribution is inconsistent with the GRMHD distribution (right panel of Figure 5), which is uniform in ξ_τ. A uniform distribution is the expected result since each simulation is randomly rotated before imaging. The north–south bias is similar to the results found for the elliptical ring in Section 3, and for the circular crescents in Figure 3.

Taken together, the measured ellipticity and orientation bias suggests that the imaging algorithms create a preferred ellipticity orientation, and along this direction the ellipticity uncertainty is maximized. This is shown in the bottom panel of Figure 5. Here we see that when the ellipticity orientation is aligned in the north–south direction, the recovered ellipticity is very uncertain, and can be quite different from the truth. Additionally, for ξ_τ approaching ±90°, the top set tends to prefer overly circular images.

To add this ellipticity uncertainty to the results for M87* we take each recovered ellipticity and orientation from the M87* top set and add the theoretical uncertainty. This gives an ellipticity for M87* of $\tau \in [0.0,\ 0.3]$. Note that this result assumes the accretion flow around M87* is well described by a GRMHD simulation, and there are no non-Kerr effects that add ellipticity to the image. In the next section, we will analyze what happens for the stretched set of GRMHD simulations, which include additional ellipticity.

4.5. Stretched Set Results

While the results from the scaled set of GRMHD simulations suggest an upper bound of $\tau \lesssim 0.35$, they do not answer what happens when there is additional ellipticity not due to the accretion disk. For instance, non-GR metrics (e.g., Johannsen & Psaltis 2010) can cause the black hole photon ring and/or event horizon to appear more elliptical than what the Kerr metric would predict. To test this, we will use the GRMHD stretched data set described above. To extract ellipticity and orientation from the stretched set, we again use `VIDA`'s

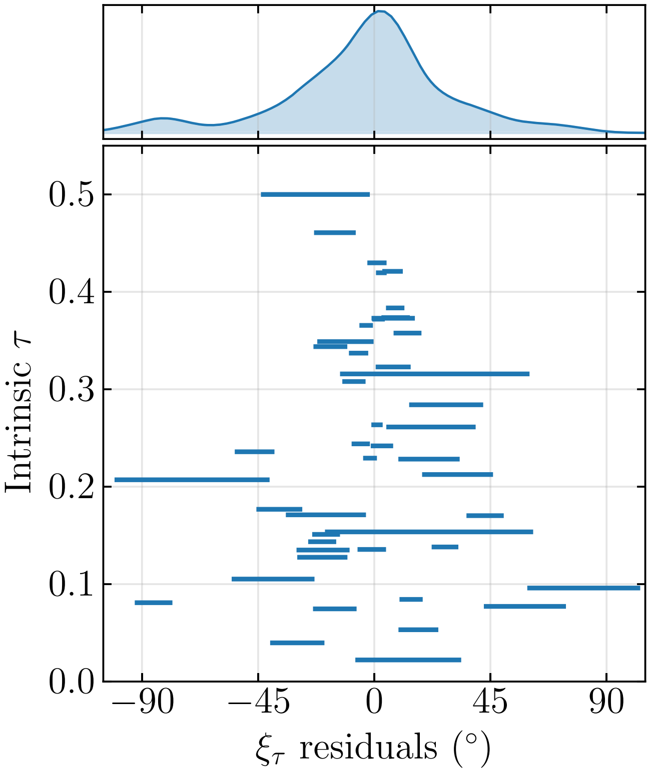

Figure 6. Residual distributions of the stretched GRMHD data set for the ellipticity position angle ξ_τ. The top panel shows the marginal distribution of the ellipticity residuals across the entire stretched GRMHD set. The bottom panel shows the 95% probability interval of the position angle residuals for each simulation, separated by the intrinsic τ on the *y*-axis. The intrinsic τ and ξ_τ are found by applying `VIDA` to the GRMHD snapshots blurred by a Gaussian with an FWHM of 10 μas. The large residuals for small τ are a result of the ξ_τ being heavily biased north–south, similar to the elliptical ring, and to the scaled GRMHD results.

`CosineRing{1,4}` template with a constant background and the Bh divergence. However, we find that a small subset of the image reconstructions have an additional circular blob present. Due to the second component, `VIDA` would sometimes report an artificially high ellipticity since the ring template would try to cover both the central ring and the Gaussian blob. To remedy this issue, we add a Gaussian component to the template to model the nonring intensity. We find that this works as expected and the ring template gives a more reliable ellipticity measurement.

The residuals for the ellipticity position angle are shown in Figure 6. Note that we only show the results for `VIDA` since `REx` is unable to measure the position angle of the ellipticity. Looking at the bottom panel we see that for intrinsic $\tau \lesssim 0.2$, the position angle residuals are very broad, and can be significantly biased from zero. However, as the intrinsic τ increases we find that the residuals improve. This bias for smaller ellipticity is due to ξ_τ being strongly biased toward 0°, and therefore away from a small residual. This bias continues as τ increases but its impact is lessened.

The left panel of Figure 7 shows a map from the recovered τ on the *x*-axis and the intrinsic τ. The horizontal bars are the 95% confidence regions about the median of the top-set images. From this we see that the recovered ellipticity is largely independent of the intrinsic ellipticity when the intrinsic $\tau \lesssim 0.325$. Furthermore, even a GRMHD simulation with intrinsic $\tau = 0.475$ can have a recovered $\tau = 0.1$, which is similar to the value obtained for the M87* top set.

To quantify the maximum allowed ellipticity that is consistent with M87* we first make two cuts on the stretched GRMHD reconstructions—namely, we remove any reconstructions where $\tau > 0.2$ or $\xi_\tau \notin [-75°, 0°]$. The remaining simulations, therefore, match the M87* top-set measurements. The intrinsic ellipticity distribution from the remaining simulations is shown in the right panel of Figure 7. From this distribution we see that stretched simulations with ellipticity as high as 0.5 can have ellipticity similar to the observed M87* results. Note that this is the highest intrinsic ellipticity considered in the stretched set.

5. Summary and Conclusions

The ellipticity of the accretion flow around M87* is a theoretically interesting property related to the nature of the accretion flow and structure of spacetime. While the results in Paper VI measured ellipticity in the image reconstruction of M87*, no attempt was made to interpret or calibrate this result. However, we have shown that the top set used for the M87* images cannot directly measure the on-sky image ellipticity. We demonstrated that even for simple geometric models, the `eht-imaging` top set fails to recover the correct ellipticity in 8/10 test cases.

To account for ellipticity bias, we calibrated the M87* ellipticity using a set of 550 GRMHD images. Assuming that the ellipticity in the reconstructions is due to accretion turbulence, we found that accounting for imaging bias the ellipticity of M87* could be anywhere from $\tau = 0$ to $\tau = 0.3$. However, if there is additional non-accretion-induced ellipticity from, for example, some non-GR effect, we found that ellipticity as high as $\tau \approx 0.5$ could have a recovered ellipticity similar to the M87* results.

The reason for the ellipticity uncertainty is twofold. First, the UV coverage for the EHT 2017 data is very sparse, having significant gaps in the north–south direction. Second, the top set used for M87* is inadequate for both the geometric and GRMHD simulations analyzed in this paper. This can be seen from the strong north–south ellipticity bias. Further evidence of the top-set inadequacy comes from the roughly 20%–40% of reconstructed images that needed to be removed from the results as detailed in Section 4.3.

The results of this paper also stress the importance of defining parameter surveys that include image features that are of interest. If the survey does not include the impact of these image features, it is not clear whether the resulting reconstruction measurements of those features will be reliable predictors of the true on-sky image. This partially requires some preliminary understanding of the true on-sky image features. Unfortunately, this is often unclear for the EHT, making designing parameter surveys difficult. Fundamentally, the reason for this difficulty is that while the top set provides an ensemble of image reconstructions, they are decidedly not a posterior over image structures. Instead, they rely on training sets to decide which set of images meet some heuristic threshold. Using such heuristics can cause the resulting distribution of recovered image features to be biased from the true values. Therefore, any quantity of interest derived from the top set needs careful validation to assess the fidelity of the measurement. In general, this validation is necessary for any imaging technique (e.g., inverse imaging (DIFMAP), RML, and Bayesian imaging), given the ill-posed nature of VLBI imaging.

Future observations will greatly help reduce the ellipticity uncertainty and disentangle the different potential causes of any

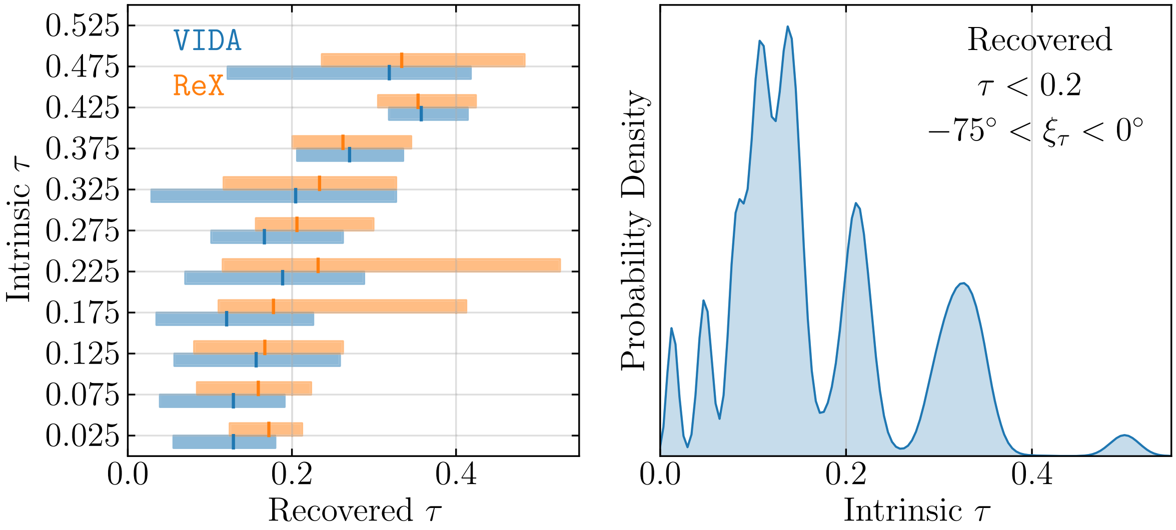

Figure 7. Left: Intrinsic ellipticity τ vs. recovered ellipticity of the stretched GRMHD images. The intrinsic ellipticity is found by applying `VIDA` to the GRMHD snapshot images blurred with a 15 μas Gaussian kernel. The solid line shows the median recovered ellipticity from `VIDA` (blue) and `REx` (orange). The filled bars are the 95% confidence intervals around the median. Right: The resulting intrinsic ellipticity distribution (measured by `VIDA`) for reconstructions with ellipticity and orientation similar to the observed M87* top-set values.

measured ellipticity and increase our knowledge of black hole parameters and accretion (Roelofs et al. 2021). The Next-generation EHT (ngEHT) will potentially add 10 more sites across the globe, significantly increasing coverage and dynamic range. By improving the dynamic range, detection of the inner shadow with the ngEHT (Doeleman et al. 2019; Raymond et al. 2021) may be possible (Chael et al. 2021). The improved coverage will also reduce the need for the ellipticity calibration that is required due to image artifacts. Proposed space-based arrays like the Event Horizon Explorer (Kurczynski et al. 2022) will also have the potential to greatly deepen our understanding of accretion (Gurvits et al. 2022), directly model or image the $n = 1$ photon ring (Gralla et al. 2020; Paugnat et al. 2022), and constrain its shape.

P.T. receives support from the Natural Science and Engineering Research Council through the Alexander Graham Bell CGS-D scholarship. This work was made possible by the facilities of the Shared Hierarchical Academic Research Computing Network (SHARCNET: www.sharcnet.ca) and Compute/Calcul Canada (www.computecanada.ca). Computations were made on the supercomputer Mammouth Parallèle 2 from the University of Sherbrooke, managed by Calcul Québec and Compute Canada. The operation of this supercomputer is funded by the Canada Foundation for Innovation; Ministère de l'Économie, de la Science et de l'Innovation du Québec (MESI); and Fonds de Recherche du Québec–Nature et Technologies (FRQ-NT). This work was supported in part by the Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science, and Economic Development and by the Province of Ontario through the Ministry of Economic Development, Job Creation, and Trade. A.E.B. thanks the Delaney family for their generous financial support via the Delaney Family John A. Wheeler Chair at Perimeter Institute. A.E.B. and P.T. receive additional financial support from the Natural Sciences and Engineering Research Council of Canada through a Discovery Grant. D.C.M.P. and P.T. were supported by the Black Hole Initiative at Harvard University, which is funded by grants from the John Templeton Foundation and the Gordon and Betty Moore Foundation to Harvard University. D.C.M.P. and P.T. were also supported by National Science Foundation grants AST 19-35980 and AST 20-34306. A.C. is supported by Hubble Fellowship grant HST-HF2-51431.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555.

Software: BlackBoxOptim.jl (BlackBoxOptim.jl Developers 2018), `eht-imaging` (Chael et al. 2018), Flux.jl (Innes 2018a), GR (Heinen et al. 1985–2019), Julia (Bezanson et al. 2017), matplotlib 3.3 (Hunter 2007), Pandas (McKinney 2010), Python 3.8.3 (Van Rossum & Drake 2009), SciPy (Virtanen et al. 2020), ThemisPy (ThemisPy Developers 2020), VIDA.jl (Tiede et al. 2022).

Appendix
Image Reconstruction Classification Using CNNs

While analyzing the image reconstructions of the GRMHD simulations, we find that a significant portion of the images fail to show a dominant ringlike feature. This failure can occur when the training set of images used in the top set's construction is different from the true image (e.g., the scaled GRMHD images). When this occurs, the reconstruction will appear to "waffle." The waffling results in flux being spread throughout the imaging plane, in a pattern similar to that of the EHT's dirty beam. Unfortunately, no global section of the top set is able to remove the waffled images. Therefore, we turn to machine-learning techniques to classify the reconstructions as having succeeded or failed.

Machine learning and neural networks have been used numerous times in astrophysical settings to classify images (e.g., Dieleman et al. 2015; Davelaar et al. 2017). To classify our images we have decided to use CNNs (e.g., Goodfellow et al. 2016). CNNs break the images into features of different scales and then group these features to classify the image. For our neural network we use the Julia package Flux (Innes 2018b; Innes et al. 2018).

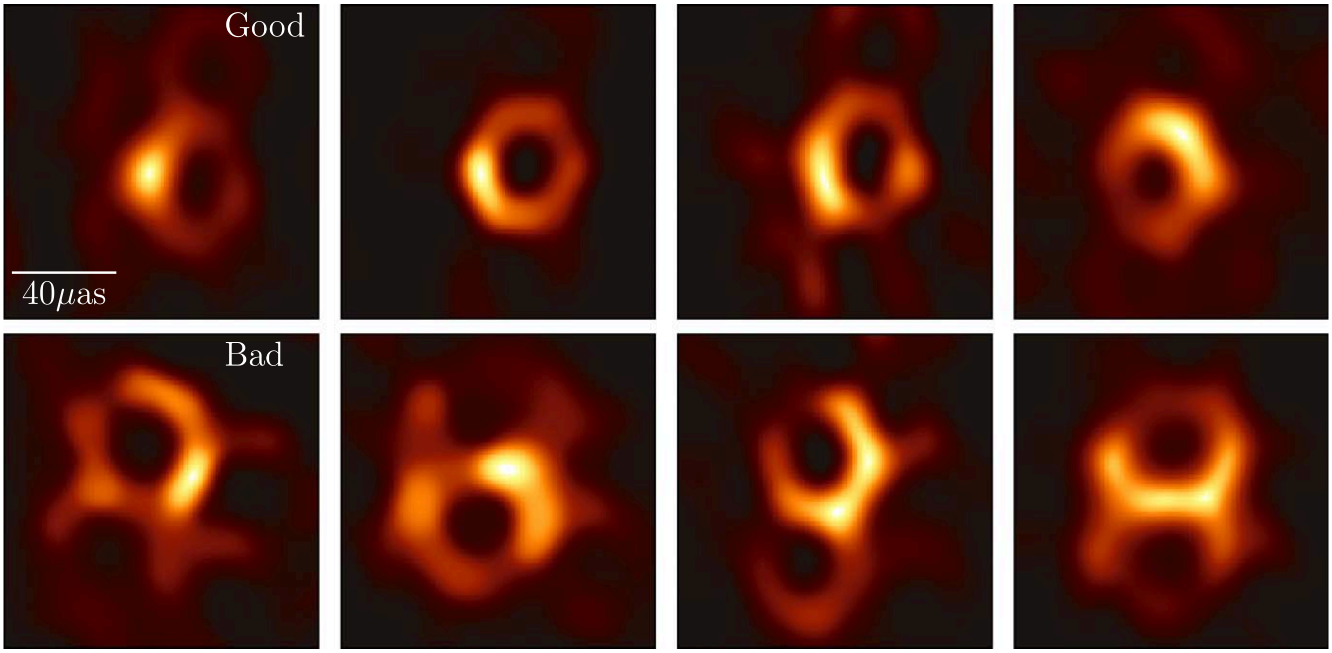

Figure 8. Examples of image classification from the trained CNN. The top row shows image reconstructions that pass the threshold, and the bottom reconstructions are those that fail. Most of the failed reconstructions fail to demonstrate a single dominant ringlike feature in the image.

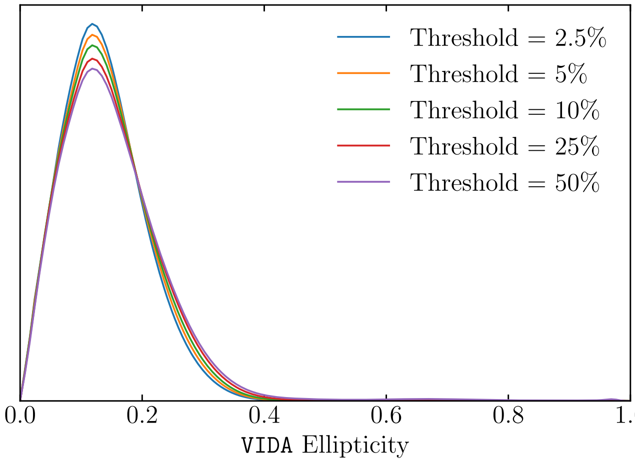

Figure 9. Impact of the threshold on the inferred ellipticity of the scaled reconstructions. Overall there is very little change in the distribution if the threshold is $\leqslant 0.5$. In this paper we use the threshold of 0.10, which removes 42% of the scaled reconstructions.

For our image classifier we use a relatively shallow network using three convolutional layers with a 2×2 max pooling. For the final layer we use a fully connected network to the two-dimensional classification space. Between each convolutional layer we use the ReLu activation function: $f(x) = \max(0, x)$. For the final fully connected layer no activation function is used. Since we are interested in binary classification we use `logitbinarycrossentropy` in Flux, which is given by

$$H(q) = -\sum_{i=1}^{N} y_i \log(\sigma(q)) + (1 - y_i)\log(1 - \sigma(q)), \quad (13)$$

where $\sigma(x) = (1 + e^{-x})^{-1}$ and y_i are the labels (1 for an image that waffles and 0 otherwise). This choice of the loss function is equivalent to using a sigmoid activation function in the last layer of the neural network but has better numerical stability.

CNNs are a form of supervised learning. Therefore, we first have to label a subset of the image reconstructions by hand. To find the labels, y_i, we analyze 5000 random images from the scaled GRMHD set and an additional 500 from the stretched set. We then classify each image by whether it visually has a dominant ringlike feature or not. Some examples of images that pass or fail are shown in Figure 8. Two-thirds of these classified images are used for training, and the rest are used for our test set. To combat overfitting, we also augment the images by adding Gaussian random noise to each image when evaluating the loss function. Finally, ADAM (Kingma & Ba 2014) is used to optimize the network using the options defined in the Flux model zoo package[12] with some minor changes—namely, we break our images into batches, with each batch containing 256 images, and use a learning rate of 3×10^{-3}. However, if predictive performance on the training set does not improve after 10 epochs we drop the learning rate by a factor of 10. The optimizer is run for 100 epochs and achieves an accuracy of 94% and 92% on the training and testing sets, respectively.

ORCID iDs

Paul Tiede https://orcid.org/0000-0003-3826-5648
Avery E. Broderick https://orcid.org/0000-0002-3351-760X
Daniel C. M. Palumbo https://orcid.org/0000-0002-7179-3816
Andrew Chael https://orcid.org/0000-0003-2966-6220

References

Akiyama, K., Ikeda, S., Pleau, M., et al. 2017, AJ, 153, 159
Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. 2017, SIAMR, 59, 65
BlackBoxOptim.jl Developers 2018, robertfeldt/BlackBoxOptim.jl, 0.6.1, https://github.com/robertfeldt/BlackBoxOptim.jl, GitHub
Bouman, K. L., Johnson, M. D., Zoran, D., et al. 2016, in Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (Piscataway, NJ: IEEE), 4

---

[12] https://github.com/FluxML/model-zoo/blob/master/vision/mnist/conv.jl

Broderick, A. E., Johannsen, T., Loeb, A., & Psaltis, D. 2014, ApJ, 784, 7
Broderick, A. E., Pesce, D. W., Tiede, P., Pu, H.-Y., & Gold, R. 2020, ApJ, 898, 9
Chael, A. 2019, PhD thesis, Harvard University
Chael, A., Johnson, M. D., & Lupsasca, A. 2021, ApJ, 918, 6
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
Davelaar, J., Moscibrodzka, M., Bronzwaer, T., & Falcke, H. 2018, A&A, 612, A34
Dieleman, S., Willett, K. W., & Dambre, J. 2015, MNRAS, 450, 1441
Doeleman, S., Blackburn, L., Dexter, J., et al. 2019, BAAS, 51, 256
Dokuchaev, V. I., & Nazarova, N. O. 2019, JETP, 128, 578
Event Horizon Telescope Collaboration, Akiyama, K., & Alberdi, A. 2019a, ApJL, 875, L1
Event Horizon Telescope Collaboration, Akiyama, K., & Alberdi, A. 2019b, ApJL, 875, L2
Event Horizon Telescope Collaboration, Akiyama, K., & Alberdi, A. 2019c, ApJL, 875, L3
Event Horizon Telescope Collaboration, Akiyama, K., & Alberdi, A. 2019d, ApJL, 875, L4
Event Horizon Telescope Collaboration, Akiyama, K., & Alberdi, A. 2019e, ApJL, 875, L5
Event Horizon Telescope Collaboration, Akiyama, K., & Alberdi, A. 2019f, ApJL, 875, L6
Event Horizon Telescope Collaboration, Akiyama, K., Algaba, J. C., et al. 2021, ApJL, 910, L13
Falcke, H., Melia, F., & Agol, E. 2000, ApJL, 528, L13
Goodfellow, I., Bengio, Y., & Courville, A. 2016, Deep Learning (Cambridge, MA: MIT Press)
Gralla, S. E., Lupsasca, A., & Marrone, D. P. 2020, PhRvD, 102, 124004
Gurvits, L. I., Paragi, Z., Amils, R. I., et al. 2022, AcAau, 196, 314
Heinen, J., et al. 1985-2019, GR Framework, https://gr-framework.org/
Honma, M., Akiyama, K., Uemura, M., & Ikeda, S. 2014, PASJ, 66, 95
Hunter, J. D. 2007, CSE, 9, 90
Innes, M. 2018a, JOSS, 3, 602
Innes, M. 2018b, JOSS, 3, 602
Innes, M., Saba, E., Fischer, K., et al. 2018, arXiv:1811.01457
Johannsen, T., & Psaltis, D. 2010, ApJ, 718, 446
Johannsen, T., Wang, C., Broderick, A. E., et al. 2016, PhRvL, 117, 091101
Johnson, M. D., Lupsasca, A., Strominger, A., et al. 2020, SciA, 6, eaaz1310
Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
Kuramochi, K., Akiyama, K., Ikeda, S., et al. 2018, ApJ, 858, 56
Kurczynski, P., Johnson, M. D., Doeleman, S. S., et al. 2022, Proc. SPIE, 12180, 121800M
McKinney, W. 2010, Proc. of the 9th Python in Science Conf., Vol. 445 (Austin, TX: SciPy), 51
Medeiros, L., Psaltis, D., & Özel, F. 2020, ApJ, 896, 7
Mertens, F., Lobanov, A. P., Walker, R. C., & Hardee, P. E. 2016, A&A, 595, A54
Narayan, R., & Nityananda, R. 1986, ARA&A, 24, 127
Paugnat, H., Lupsasca, A., Vincent, F., & Wielgus, M. 2022, arXiv:2206.02781
Raymond, A. W., Palumbo, D., Paine, S. N., et al. 2021, ApJS, 253, 5
Roelofs, F., Fromm, C. M., Mizuno, Y., et al. 2021, A&A, 650, A56
ThemisPy Developers 2020, aeb/ThemisPy, 0.4.0, GitHub, https://github.com/aeb/ThemisPy
Tiede, P., Broderick, A. E., & Palumbo, D. C. M. 2022, ApJ, 925, 122
Van Rossum, G., & Drake, F. L. 2009, Python 3 Reference Manual (Scotts Valley, CA: CreateSpace)
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261
Wiaux, Y., Jacques, L., Puy, G., Scaife, A. M. M., & Vandergheynst, P. 2009a, MNRAS, 395, 1733
Wiaux, Y., Puy, G., Boursier, Y., & Vandergheynst, P. 2009b, MNRAS, 400, 1029