

VLBInet: Radio Interferometry Data Classification for EHT with Neural Networks

JOSHUA Y.-Y. LIN ^{1,2,3,4} DOMINIC W. PESCE ^{5,6} GEORGE N. WONG ^{1,2,7,8} AJAY UPPILI ARASANIPALAI ¹
BEN S. PRATHER ^{1,2} AND CHARLES F. GAMMIE ^{9,2,3,1}

¹*Department of Physics, University of Illinois, 1110 West Green Street, Urbana, IL 61801, USA*

²*Illinois Center for Advanced Study of the Universe, 1110 West Green Street, Urbana, IL 61801, USA*

³*National Center for Supercomputing Applications, 605 East Springfield Avenue, Champaign, IL 61820, USA*

⁴*Center for Computational Astrophysics, Flatiron Institute, Simons Foundation, 162 5th Ave, New York, NY 10010, USA*

⁵*Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA*

⁶*Black Hole Initiative at Harvard University, 20 Garden Street, Cambridge, MA 02138, USA*

⁷*School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA*

⁸*Princeton Gravity Initiative, Princeton University, Princeton, New Jersey 08544, USA*

⁹*Department of Astronomy, University of Illinois, 1002 West Green Street, Urbana, IL 61801, USA*

Submitted to ApJ

ABSTRACT

The Event Horizon Telescope (EHT) recently released the first horizon-scale images of the black hole in M87. Combined with other astronomical data, these images constrain the mass and spin of the hole as well as the accretion rate and magnetic flux trapped on the hole. An important question for the EHT is how well key parameters, such as trapped magnetic flux and the associated disk models, can be extracted from present and future EHT VLBI data products. The process of modeling visibilities and analyzing them is complicated by the fact that the data are sparsely sampled in the Fourier domain while most of the theory/simulation is constructed in the image domain. Here we propose a data-driven approach to analyze complex visibilities and closure quantities for radio interferometric data with neural networks. Using mock interferometric data, we show that our neural networks are able to infer the accretion state as either high magnetic flux (MAD) or low magnetic flux (SANE), suggesting that it is possible to perform parameter extraction directly in the visibility domain without image reconstruction. We have applied VLBInet to real M87 EHT data taken on four different days in 2017 (April 5, 6, 10, 11), and our neural networks give a score prediction 0.52, 0.4, 0.43, 0.76 for each day, with an average score 0.53, which shows no significant indication for the data to lean toward either the MAD or SANE state.

1. INTRODUCTION

The Event Horizon Telescope (EHT) is a globe-spanning network of millimeter wavelength observatories (EHTC I; EHTC II). Data from individual observatories are combined to measure the Fourier components of a source *intensity* on the sky (EHTC III). The sparse set of Fourier components together with a regularization procedure can then be used to reconstruct an image of the source (EHTC IV). The resulting images of the black hole at the center of M87 (hereafter M87*) have a ringlike structure—attributed to emission from hot plasma surrounding the black hole—with an asymmetry that contains information about the motion of plasma around the hole (EHTC V; EHTC VI). Combined with data from other sources, the EHT images constrain the black hole spin and mass as well as properties of the surrounding magnetic field structure and strength (EHTC V; EHTC VI). More recently, the EHT has produced polarized emission maps of M87* (EHTC VII), and these imply a highly organized magnetic field structure in the source (EHTC VIII).




Figure 1. Example synthetic images based on numerical models of black hole accretion flows with M87-like parameters. Top row: sample images of MAD accretion flows. Bottom row: sample images of SANE flows.

Black hole accretion flows can be divided into two qualitatively different states, called MAD and SANE, depending on the strength of the magnetic field near the event horizon. In the magnetically arrested disk (MAD) state, magnetic fields near the horizon are dynamically important enough to limit the flow of plasma onto the hole (Bisnovatyi-Kogan & Ruzmaikin 1974; Ichimaru 1977; Igumenshchev et al. 2003; Narayan et al. 2003). When MAD, the magnetic flux Φ_B through the black hole horizon is at the maximum value sustainable by the accretion flow; additional flux escapes outward through the inflowing plasma. The flux can be characterized by the dimensionless flux $\phi = \Phi_B / \sqrt{\dot{M} R_g^2 c}$ ($\dot{M} \equiv$ accretion rate; $R_g \equiv GM/c^2$). In the MAD state $\phi \simeq 15$.¹ In the standard and normal evolution (SANE) state, by contrast, $\phi < \phi_c$ and the flow is organized in a conventional, geometrically thick, centrifugally supported disk. For a fixed accretion rate and black hole spin, MAD flows have stronger jets driven by the Blandford–Znajek effect. Is it possible to identify whether a source is in the MAD or SANE state directly from VLBI data?

Various methods have been proposed for using black hole images or VLBI data to constrain system parameters—see (Johnson et al. 2020; Gralla & Lupsasca 2020; Lupsasca et al. 2018) for measuring black hole parameters using long-baseline visibilities and (Wong 2021; Hadar et al. 2021; Chesler et al. 2021) through autocorrelations of position-dependent light echoes—however, the general parameter extraction problem is challenging because simulated images contain an abundance of information about not only the black hole spacetime and gross accretion flow properties, but also about fluctuating, turbulent structure in the accreting plasma. Past efforts to compare simulated GRMHD images with VLBI data have typically either fit to the visibilities constructed from simulated images (e.g., the “average image scoring” procedure from EHTC V, which runs into issues of attempting to match or marginalize over specific instantiations of turbulence between simulation and observation) or else have focused on observationally accessible proxies for physical quantities (e.g., the metric derived in Palumbo et al. 2020 and used in EHTC VIII). Here we consider a neural network (NN)-based approach to parameter extraction from VLBI data. This approach has the benefit of using all the data (as opposed to observationally accessible proxies) and can ideally distill from the data all information that constrains a model parameter.

Earlier work by van der Gucht et al. (2020) found evidence that neural networks could be used to infer values for a limited set of parameters in a library of synthetic black hole images of SANE flows. Work by Lin et al. (2020a) explored both MAD and SANE models and investigated neural network interpretability. Both these earlier efforts worked in the image domain, though the real data correspond to sparse sampling in the Fourier domain. In recent years, work by Sun & Bouman (2020); Sun et al. (2020); Morningstar et al. (2018) have shown that deep neural networks can be used for visibility-to-image reconstruction, as well as tracking their underlying distribution in a probabilistic way using certain evaluation loss metrics. Work by Popov et al. (2021) used convolutional neural networks to estimate an inclination angle from complex visibility with synthetic disk-based model. Here we explore the possibility of using neural networks to perform an end-to-end analysis of synthetic M87* data in the visibility domain.

Our paper is organized as follows. In Section 2, we describe our procedure for generation of synthetic observations. In Section 3 we describe our neural network, and we present results in Section 4 and provide a discussion in Section 5.

¹ We use Lorentz–Heaviside units rather than Gaussian units to be consistent with Porth 2019. In the latter system, $\phi_c \approx 50$.




Figure 2. Visibility amplitude versus baseline length for a randomly-selected single training image with different noise components applied. Panel (a) shows the noise-free Fourier samples, panel (b) shows these samples after the addition of complex Gaussian noise, and panel (c) shows the same as (b) after additionally modulating each visibility by station-based complex gain noise (see subsection 2.2)

2. SYNTHETIC DATA GENERATION

2.1. Simulating black hole images with GRMHD/GRRT

The training data for our supervised machine learning pipeline consist of synthetic images based on general relativistic magnetohydrodynamic (GRMHD) simulations of black hole accretion. The image dataset used in this work is the same as the one used in (Lin et al. 2020a). We obtained 40000 images from the image library, and split them into half for training and testing.

The GRMHD simulations were generated with the *iharm3D* code (Prather et al. JOSS, submitted) for five black hole spins ($a_* = -0.94, -0.5, 0, 0.5, 0.94$) in both MAD and SANE states. Notice that there is a continuum of possible SANE states; our SANE models have $\phi \simeq 1$. For selected spins and magnetic states we consider multiple GRMHD simulations initialized with different seed perturbations to assess the likelihood that our neural networks are over-fitting the simulations.

We then imaged the GRMHD simulation output using *ipole* (Mościbrodzka & Gammie 2018). To do this we must specify additional physical parameters: the approximate size of the black hole and its distance from the observer; the accretion rate (set so that the intensity integrated over the images matches the mm flux density of M87*); and a parameter R_{high} that regulates assignment of the electron temperature from fluid variables (see EHTC V). R_{high} is drawn uniformly from a discrete set of allowed values [1, 10, 40, 160]. In addition we must specify parameters related to the imaging process: the position angle (PA) of the source, the angular resolution of the image, and the coordinates of the black hole center on the image. To prevent the neural network from correlating imaging parameters with source parameters we introduce image-by-image variations: the PA is drawn uniformly from $[0, 2\pi]$; the angular scale is drawn uniformly from a small interval corresponding to a $\pm 10\%$ change in angular scale of the source (equivalently a 10% range in black hole mass or distance); the image center is drawn uniformly from a square region corresponding to a $[-10, 10] \mu\text{as}$ displacement in each coordinate.

2.2. Synthetic Data Generation

For each image in the training and testing sets, we generated synthetic VLBI observations using *eht-imaging* (Chael et al. 2016, 2018) in a manner similar to that described in EHTC IV. The baseline coverage and observing cadence match the 2017 April 11 observations of M87 carried out by the EHT after coherently time-averaging each baseline’s complex visibilities on a per-scan basis (EHTC II; EHTC III). The source image in each case is assumed to be located at the sky position of the M87 black hole, and the baseline thermal noise values have been taken from the original calibrated observations provided by the EHT collaboration (Event Horizon Telescope Collaboration 2019f).

In the absence of noise the complex visibilities are, according to the van Cittert–Zernike theorem (Thompson et al. 2017), samples of the Fourier transform of the sky emission. Real-world telescope arrays suffer from a number of a




Figure 3. VLBI Pipeline.




Figure 4. Two neural network classes used in this paper. The first accepts complex visibilities as input, and the second accepts closure quantities as input.

priori unknown systematics, however, the most severe of which are variations in the phase and amplitude of the signal received at each station caused by atmospheric turbulence and other signal path effects.

We have thus generated three classes of mock observations for each of the training and testing images: (1) a “noise free” observation that contains only the Fourier transform of the input image sampled on the EHT baselines, (2) a “thermal noise” observation in which each complex visibility has been modulated according to its measurement uncertainty, and (3) a “full noise” observation that includes station-based complex gain fluctuations. The synthetic gain amplitudes are modeled as Gaussian-distributed with unit mean and a standard deviation of 10% for all stations, while the synthetic gain phases are drawn uniformly from $[0, 2\pi)$; both gain amplitudes and phases are sampled independently for each station at each timestamp.

Gain variations are particularly severe at the short (~ 1.3 mm) observing wavelength of the EHT, motivating the construction and use of so-called “closure” data products that are immune to station-based corruption. The two standard sets of closure quantities are closure phases and closure amplitudes, which are constructed using triangles and quadrangles of baselines, respectively. Denoting a visibility measurement on the baseline between stations i and j as V_{ij} , the closure phase on triangle ijk is given by the argument of the directed triple product,

$$\psi_{ijk} = \arg(V_{ij}V_{jk}V_{ki}). \quad (1)$$

The closure amplitude on a quadrangle $ijkl$ is given by

$$A_{ijkl} = \left| \frac{V_{ij}V_{kl}}{V_{il}V_{jk}} \right|. \quad (2)$$

Both closure phases and closure amplitudes have the property that station-based noise cancels, making them robust observables. In this paper we construct two classes of neural networks, one to treat complex visibilities and the other to treat closure quantities.

3. NEURAL NETWORKS

We decided to use fully connected neural networks, also known as multilayer perceptrons (MLPs), for the MAD/SANE classification task. Previous work used convolutional NNs (CNNs) and image domain data (van der Gucht et al. 2020; Lin et al. 2020a). We use sparsely sampled visibility domain data—where the input is not translation invariant—and therefore selected the MLP architecture as one of the simplest architectures that avoids possible inductive biases associated with the translational invariance embedded in, for example, CNNs (Battaglia et al. 2018). We will refer to this arrangement as VLBI-net.




Figure 5. Histogram of scores for MAD and SANE models. If the neural network could identify the model perfectly then all MADs would score 0 and all SANEs would score 1. The grey dashed line is the decision boundary: data with score > 0.5 are classified as SANE model, and < 0.5 are classified as MAD model.

When working with complex visibilities, we decompose the complex numbers into amplitudes and phases and then project the phases into sines and cosines to enable phase wrapping. For both complex visibilities and closure quantities the data are organized in a consistent way so that each neuron in the input layer is associated with a particular point in the uv domain (complex visibilities) or triangle or quadrangle (closure quantities).

We trained 6 different versions of VLBI-net, ranging over 2 different visibility inputs (complex/closure) and 3 different noise scenarios. All versions consist of 7 layers. For complex visibilities, the numbers of neurons at each layer are 648, 256, 64, 64, 64, 64, 1 and the full model has 195,137 trainable parameters. For closure quantities the numbers of neurons at each layer are 381, 256, 64, 64, 64, 64, 1 and the model has 126,785 trainable parameters. In MLPs each connection contains a weight and a bias. We use tanh as the activation function. We use binary cross-entropy (BCE) loss for the MAD/SANE since the output is expected to be in one of these two categories: $\text{BCE Loss} = t \log(p) + (1 - t) \log(1 - p)$, where $t = 0$ for MAD and $t = 1$ for SANE. Here p is the score prediction with range from 0 – 1. VLBI-net is built using the Python 3 deep learning library *pytorch* (Paszke et al. 2019).

4. RESULTS

4.1. Results on synthetic data

Figure 5 shows the distribution of MAD/SANE scores for 20,000 samples drawn with uniform probability from the test set. Notice that the vertical axis is logarithmic. Evidently the data contain information about whether the source is MAD or SANE, and this information can be detected by the NN. Figure 6 shows the performance of our VLBI-net with different noise models and the form of VLBI data (complex visibilities/ closure quantities) are given, and we calculate their area under curve (AUC) (Bradley 1997).




Figure 6. We calculate the area under Receiver operating characteristic (ROC) curve as a performance measure for different noise models along with complex visibilities/ closure quantities. Here we treat SANE model as positive.

To quantify how well the NN performs in the classification task, we split the data into 20 subsets and calculate the mean and standard deviation of the prediction across the subsets. The resulting uncertainty estimates are shown in Table 1.

For complex visibilities it is evident that almost no information is lost even when thermal noise is added; the real loss of information is due to gain corruption. By contrast, the closure quantities lose information via thermal noise but lose nothing through gain corruption. This is unsurprising since closure quantities are unaffected by gain errors.

For complex visibility, the MAD/SANE classifier network is able to correctly identify MAD/SANE with $79.67\% \pm 1.13\%$ accuracy without noise, $79.13\% \pm 1.25\%$ accuracy with thermal noise presence and achieve $70.99\% \pm 1.59\%$ accuracy with complex visibility, thermal noise + amplitude gain + phase error. For closure visibility, the MAD/SANE classifier network is able to correctly identify MAD/SANE with $76.57\% \pm 1.31\%$ accuracy without noise, $73.55\% \pm 1.14\%$ accuracy with thermal noise presence and achieve $73.24\% \pm 1.57\%$ accuracy with complex visibility, thermal noise + amplitude gain + phase error.

4.2. Results on real M87* data

We ran the trained VLBI net (trained on models with added noise and using closure quantities) on real M87 VLBI data taken on April 5th, 6th, 10th, and 11th, 2017. VLBI net output scores of 0.52, 0.4, 0.43, 0.76, respectively. Among the 20,000 data in our test set, there are 140 instances that are similar to the prediction scores (within 0.142, one sigma from the score distribution for 4 days). Among the 140 that are similar, 47% of them are “MAD”.

It is interesting that the results are similar across days, suggesting that our data are consistent with no strong indication (e.g., over 2σ) with the MAD/SANE models. If the models faithfully represent the physical situation in M87*, improvements in data quality (e.g., better baseline coverage) are likely to enable a confident classification of the source as MAD or SANE.

5. DISCUSSION

We have presented a neural network pipeline that can perform data classification to identify the accretion state of M87* as either MAD or SANE using VLBI observations, after training on a library of 2×10^4 simulations. Our pipeline operates directly on interferometric data products—either complex visibilities or closure quantities—and thus




Figure 7. Visibility amplitude of M87* taken at 2017 April 5th, 6th, 10th and 11th. The corresponding MAD/SANE scores are 0.52, 0.4, 0.43, 0.76 for each day's observation.

Classification accuracy			
Visibility form	Noise free	Thermal noise	Thermal noise + complex gains
Complex	$79.7\% \pm 1.1\%$	$79.1\% \pm 1.3\%$	$71.0\% \pm 1.6\%$
Closure	$76.6\% \pm 1.3\%$	$73.6\% \pm 1.1\%$	$73.2\% \pm 1.6\%$

Table 1. Summary Statistics

Neural Networks MAD/SANE prediction score with closure quantities		
Observational date	Score	Prediction
2017 April 5	0.52	SANE
2017 April 6	0.40	MAD
2017 April 10	0.43	MAD
2017 April 11	0.76	SANE

Table 2. Neural Networks MAD/SANE prediction score for M87 on all four observational dates

does not rely on an image reconstruction procedure. We find that the networks are able to achieve a $>70\%$ classification accuracy across the board, even in the presence of realistic levels of both thermal and complex gain data corruptions.

Our network performs its most accurate classification ($\sim 80\%$) when the data are provided as noise-free complex visibilities and its least accurate classification ($\sim 71\%$) when the data are provided as complex visibilities containing both thermal and complex gain uncertainties. The network achieves intermediate performance ($\sim 76\%$) when it is provided with noise-free closure quantities and when the closure quantities contain realistic levels of noise the performance is more similar ($\sim 73\%$) to that of the corrupted complex visibilities. These performance differences are qualitatively consistent with the expected relative amounts of source information retained in visibilities versus closure quantities (Blackburn et al. 2020). We also note that the classification accuracy using any variant of interferometric data product is considerably worse than that expected when performing classification using the original images, which can achieve essentially perfect ($\sim 99\%$; Lin et al. 2020a) classification accuracy; again, this difference reflects the loss of information due to sparse Fourier sampling and the reduced angular resolution of the data.

Although we did not provide the neural network with explicit priors, our input training data set covered a limited set of parameter values and thus produced an implicit, data-driven prior through the machine learning process.

5.1. Comparison with traditional analysis pipelines

The initial theoretical analysis performed by the EHT included a comparison between observational data products and simulated images of black hole accretion flows. In a comparison such as this a χ^2_ν (reduced chi squared) distance

comparison of real and synthetic data is not useful because intra-model variations—due to changes in image features caused by turbulent fluctuations in the source, for example—are large compared to measurement errors.

In contrast to snapshot-by-snapshot comparisons, the average image scoring (AIS) procedure described in EHTC V was used to compute the likelihood the data would be drawn from a given model. The AIS procedure ruled out a few retrograde-spin accretion flow models based on variability, but it left much of the parameter space unconstrained. Compared to our data-driven approach, the AIS procedure is less constraining.

5.2. Future work

We have trained VLBI-net only to classify M87 source models as MAD or SANE based on the 2017 EHT configuration. In future work we plan to extend VLBI-net to estimate black hole mass and spin as well as parameters describing the state of the plasma, and we plan to consider other arrays such as the EHT 2021 configuration and prospective ngEHT configurations (Blackburn et al. 2019). We will also introduce more realistic data corruptions using a more sophisticated end-to-end VLBI synthetic data generation pipeline (e.g., Roelofs et al. 2020).

In this work, we applied the neural network to M87 VLBI data. There are few limitations, however: 1) if our GRMHD and GRRT simulations are flawed then the neural network, which has been trained on a simulation-based synthetic dataset, may not be able to generalize to realistic data. 2) Similarly if the noise models used in the generation of synthetic VLBI observables are flawed then the neural network, which has been trained on the flawed noise model in the synthetic dataset, may not be able to generalize to realistic noise (e.g., realistic atmospheric, instrumental, and calibration effects) and hence could predict spurious results. 3) Our training set covers the model parameter space sparsely—for example, we use only five values for black hole spin—and the prediction could therefore be biased.

Although we have demonstrated the utility of our framework only for the EHT, we believe that it should in principle work for general classification analyses using interferometric data. For instance, a similar network trained on an appropriate set of simulated protoplanetary disk observations using ALMA could plausibly constrain properties such as disk inclination, thickness, and mass. In practice, using the network in such a fashion is likely to be limited by the availability of suitable simulations, perhaps motivating the development of simulation capabilities. This could have broad impact on the community, as interferometric techniques have been crucial for many astrophysical observational cases, including studies of dark matter substructure in strong gravitational lensing (Hezaveh et al. 2016; Lin et al. 2020b) and starburst galaxies (Vieira et al. 2013).

Since the EHT captures information about source polarization and time dependence, we also plan to study whether these additional data can further constrain physical parameters of the source. Synchrotron emission is naturally polarized and the local properties of the magnetic field in the plasma alter the electric vector position angle as radiation propagates through the plasma. Thus, the structure of the magnetic field in the plasma influences the orientation of linear polarization observed in black hole images. Palumbo et al. (2020) and EHTC VIII showed that the qualitative and quantitative differences between the near-horizon magnetic field strength and structure suggest that linear polarization data might be particularly useful in differentiating between MAD and SANE models. We will apply interpretability methods (Ghorbani et al. 2019; Lin et al. 2020a) to gain a more quantitative understanding of the features in the visibility domain. It is critical to understand the origin of accurate neural network classifications/regressions so that we can be confident that we are not overfitting.

As pointed out by Sun et al. (2020), the VLBI observing process can be viewed as a physics-constrained autoencoder in which the encoder is the sampling strategy through the radio telescope. Our work could potentially be helpful for evaluating telescope site candidates in a fast and automatic way (Raymond et al. 2021). Visibility in some sense is similar to the hidden representation/bottleneck in the bottleneck of variational autoencoder. Traditionally, one would need to reconstruct the image and do data analysis on top of the reconstructed images. Since the whole VLBI process looks like an auto-encoder, we could in principle not only optimize the decoder (from measurement to parameters) but also optimize the encoder (from source to encoded representation that preserve the most information through sparse sampling) (Sun & Bouman 2020; Bakker et al. 2020).

ACKNOWLEDGMENTS

The authors thank He Sun, Katie Bouman, Michael Janssen, Alexander Raymond, Bart Ripperda, Shep Doeleman, Gil Holder, Max Welling, Abhishek Joshi, Vedant Dhruv, David Ruhe, Michael Eickenberg, Shirley Ho, and David Spergel for useful discussion. The authors thank John Wardle, Shiro Ikeda for useful feedback. JL, GNW, BP, and CFG were supported by the National Science Foundation under grants AST 17-16327, OISE 17-43747, and AST 20-34306. GNW was supported in part by a Donald C. and F. Shirley Jones Fellowship and the Institute for Advanced Study. DWP acknowledges support provided by the NSF through grants AST-1952099, AST-1935980, AST-1828513, and AST-1440254, and by the Gordon and Betty Moore Foundation through grant GBMF-5278. This work has been supported in part by the Black Hole Initiative at Harvard University, which is funded by grants from the John Templeton Foundation and the Gordon and Betty Moore Foundation to Harvard University. This work used the Extreme Science and Engineering Discovery Environment (XSEDE) resource stampede2 at TACC through allocation TG-AST170024. This work utilizes resources supported by the National Science Foundation’s Major Research Instrumentation program, grant 1725729, as well as the University of Illinois at Urbana–Champaign. JL thanks the AWS Cloud Credits for Research program. JL and CFG thank the GCP Research Credits program.

REFERENCES

- Bakker, T., van Hoof, H., & Welling, M. 2020, *Advances in Neural Information Processing Systems*, 33
- Battaglia, P. W., Hamrick, J. B., Bapst, V., et al. 2018, *arXiv preprint arXiv:1806.01261*
- Bisnovatyi-Kogan, G. S., & Ruzmaikin, A. A. 1974, *Ap&SS*, 28, 45, doi: [10.1007/BF00642237](https://doi.org/10.1007/BF00642237)
- Blackburn, L., Pesce, D. W., Johnson, M. D., et al. 2020, *The Astrophysical Journal*, 894, 31
- Blackburn, L., Doeleman, S., Dexter, J., et al. 2019, *arXiv preprint arXiv:1909.01411*
- Bradley, A. P. 1997, *Pattern recognition*, 30, 1145
- Chael, A. A., Johnson, M. D., Bouman, K. L., et al. 2018, *ApJ*, 857, 23, doi: [10.3847/1538-4357/aab6a8](https://doi.org/10.3847/1538-4357/aab6a8)
- Chael, A. A., Johnson, M. D., Narayan, R., et al. 2016, *ApJ*, 829, 11, doi: [10.3847/0004-637X/829/1/11](https://doi.org/10.3847/0004-637X/829/1/11)
- Chesler, P., Blackburn, L., Doeleman, S., et al. 2021, *Classical and Quantum Gravity*
- Collaboration, E. H. T. 2019, *ApJ*, 875, L1, doi: [10.3847/2041-8213/ab0ec7](https://doi.org/10.3847/2041-8213/ab0ec7)
- Event Horizon Telescope Collaboration. 2019a, *ApJ*, 875, L2, doi: [10.3847/2041-8213/ab0c96](https://doi.org/10.3847/2041-8213/ab0c96)
- . 2019b, *ApJ*, 875, L3, doi: [10.3847/2041-8213/ab0c57](https://doi.org/10.3847/2041-8213/ab0c57)
- . 2019c, *ApJ*, 875, L4, doi: [10.3847/2041-8213/ab0e85](https://doi.org/10.3847/2041-8213/ab0e85)
- . 2019d, *ApJ*, 875, L5, doi: [10.3847/2041-8213/ab0f43](https://doi.org/10.3847/2041-8213/ab0f43)
- . 2019e, *ApJ*, 875, L6, doi: [10.3847/2041-8213/ab1141](https://doi.org/10.3847/2041-8213/ab1141)
- . 2019f, *First M87 EHT Results: Calibrated Data*, *CyVerse Data Commons*, doi: [10.25739/g85n-fl34](https://doi.org/10.25739/g85n-fl34)
- . 2021a, *The Astrophysical Journal Letters*, 910, L12
- . 2021b, *ApJ*, 910, L13, doi: [10.3847/2041-8213/abe4de](https://doi.org/10.3847/2041-8213/abe4de)
- Ghorbani, A., Wexler, J., Zou, J. Y., & Kim, B. 2019, in *Advances in Neural Information Processing Systems*, 9273–9282
- Gralla, S. E., & Lupsasca, A. 2020, *Physical Review D*, 102, 124003
- Hadar, S., Johnson, M. D., Lupsasca, A., & Wong, G. N. 2021, *Phys. Rev. D*, 103, 104038, doi: [10.1103/PhysRevD.103.104038](https://doi.org/10.1103/PhysRevD.103.104038)
- Hezaveh, Y. D., Dalal, N., Marrone, D. P., et al. 2016, *The Astrophysical Journal*, 823, 37
- Ichimaru, S. 1977, *ApJ*, 214, 840, doi: [10.1086/155314](https://doi.org/10.1086/155314)
- Igumenshchev, I. V., Narayan, R., & Abramowicz, M. A. 2003, *ApJ*, 592, 1042, doi: [10.1086/375769](https://doi.org/10.1086/375769)
- Johnson, M. D., Lupsasca, A., Strominger, A., et al. 2020, *Science advances*, 6, eaaz1310
- Lin, J. Y.-Y., Wong, G. N., Prather, B. S., & Gammie, C. F. 2020a, *arXiv preprint arXiv:2007.00794*
- Lin, J. Y.-Y., Yu, H., Morningstar, W., Peng, J., & Holder, G. 2020b, *arXiv preprint arXiv:2010.12960*
- Lupsasca, A., Porfyriadis, A. P., & Shi, Y. 2018, *Physical Review D*, 97, 064017
- Morningstar, W. R., Hezaveh, Y. D., Levasseur, L. P., et al. 2018, *arXiv preprint arXiv:1808.00011*
- Mościbrodzka, M., & Gammie, C. F. 2018, *MNRAS*, 475, 43, doi: [10.1093/mnras/stx3162](https://doi.org/10.1093/mnras/stx3162)
- Narayan, R., Igumenshchev, I. V., & Abramowicz, M. A. 2003, *PASJ*, 55, L69, doi: [10.1093/pasj/55.6.L69](https://doi.org/10.1093/pasj/55.6.L69)
- Palumbo, D. C. M., Wong, G. N., & Prather, B. S. 2020, *ApJ*, 894, 156, doi: [10.3847/1538-4357/ab86ac](https://doi.org/10.3847/1538-4357/ab86ac)
- Paszke, A., Gross, S., Massa, F., et al. 2019, in *Advances in Neural Information Processing Systems* 32, ed. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Curran Associates, Inc.), 8024–8035
- Popov, A., Stokov, V., & Surdyayev, A. 2021, *Astronomy and Computing*, 100467

- Porth, O. e. a. 2019, ApJS, 243, 26,
doi: [10.3847/1538-4365/ab29fd](https://doi.org/10.3847/1538-4365/ab29fd)
- Raymond, A. W., Palumbo, D., Paine, S. N., et al. 2021,
The Astrophysical Journal Supplement Series, 253, 5
- Roelofs, F., Janssen, M., Natarajan, I., et al. 2020,
Astronomy & Astrophysics, 636, A5
- Sun, H., & Bouman, K. L. 2020, arXiv preprint
arXiv:2010.14462
- Sun, H., Dalca, A. V., & Bouman, K. L. 2020, in 2020
IEEE International Conference on Computational
Photography (ICCP), IEEE, 1–12
- Thompson, A. R., Moran, J. M., & Swenson, George W., J.
2017, Interferometry and Synthesis in Radio Astronomy,
3rd Edition, doi: [10.1007/978-3-319-44431-4](https://doi.org/10.1007/978-3-319-44431-4)
- van der Gucht, J., Davelaar, J., Hendriks, L., et al. 2020,
Astronomy & Astrophysics, 636, A94
- Vieira, J., Marrone, D. P., Chapman, S., et al. 2013,
Nature, 495, 344
- Wong, G. N. 2021, ApJ, 909, 217,
doi: [10.3847/1538-4357/abdd2d](https://doi.org/10.3847/1538-4357/abdd2d)