

The Role of Adaptive Ray Tracing in Analyzing Black Hole Structure

Z. Gelles^{1,3}, B. S. Prather², D. C. M. Palumbo^{1,3}, M. D. Johnson^{1,3}, G. N. Wong², and B. Georgiev^{4,5}

¹Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA

² Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, IL 61801, USA

³ Black Hole Initiative at Harvard University, 20 Garden Street, Cambridge, MA 02138, USA

⁴ Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, ON, N2L 3G1, Canada

Received 2020 December 23; revised 2021 March 9; accepted 2021 March 10; published 2021 May 3

Abstract

The recent advent of the Event Horizon Telescope (EHT) has made direct imaging of supermassive black holes a reality. Simulated images of black holes produced via general relativistic ray tracing and radiative transfer provide a key counterpart to these observational efforts. Black hole images have a wide range of physically interesting image structures, ranging from extremely fine scales in their lensed "photon rings" to the very large scales in their relativistic jets. The multiscale nature of the black hole system is therefore suitable for a multiscale approach to generate simulated images that capture all key elements of the system. Here, we present a prescription for adaptive ray tracing, which enables efficient computation of extremely high-resolution images of black holes. Using the polarized ray-tracing code IPOLE, we image a combination of semianalytic and general relativistic magnetohydrodynamic (GRMHD) models, and we show that images can be reproduced with a mean squared error of less than 0.1% even after tracing $12 \times$ fewer rays. We then use adaptive ray tracing to explore the properties of the photon ring. We illustrate the behavior of individual subrings in GRMHD simulations, and we explore their signatures in interferometric visibilities.

Unified Astronomy Thesaurus concepts: Black holes (162); Radiative transfer simulations (1967); Gravitational lensing (670); Radio interferometry (1346)

1. Introduction

When surrounded by emitting material, black holes imprint distinctive properties of their spacetimes on the image seen by a distant observer. Black hole images can then offer valuable insights into the astrophysical processes that govern the accretion and outflow, the physical processes that produce heating and dissipation in the nearby plasma, and the geometrical lensing of light. Over the past few decades, images of black holes have evolved from being studied primarily for their rich theoretical features (Luminet 1979; Bardeen 1973) to being directly accessible via very long-baseline interferometry (Event Horizon Telescope Collaboration et al. 2019a, 2019b, 2019c, 2019d, 2019e, 2019f). With progressively sharper images of black holes expected as these observations continue to improve, increasingly accurate simulated images of black holes are imperative to guide analysis and interpretation.

One limitation of image accuracy is related to finite image sampling at discrete points on the screen. Namely, the intensity at each point on an image is computed by ray-tracing the path of the corresponding null geodesic and computing the radiative transfer along the trajectory. The computational expense of forming an image then increases with the number of rays at which this intensity function is sampled. A crucial question is how to efficiently distribute a finite sample of rays across an image to reach a prescribed image fidelity.

Black hole ray-tracing programs typically distribute rays on an evenly spaced grid (see, e.g., Gold et al. 2020). In this approach, regions of the image with sharp, bright features are sampled with the same density of rays as the faint regions of only diffuse structure. Black hole images are expected to have regions of both categories. Near the black hole, the accretion flow is turbulent and bright, requiring high resolution to adequately resolve. Far from the black hole, tightly collimated outflows or "jets" produce narrow regions with significant flux. The strong lensing of Kerr black holes is manifest in the "photon ring," a bright ring with a self-similar substructure that emerges in the limit of no absorption and scattering (see, e.g., Luminet 1979; de Vries 2000; Takahashi 2004; Beckwith & Done 2005; Johannsen & Psaltis 2010; Gralla et al. 2019; Johnson et al. 2020). Apart from these distinctive parts of the image, black hole images often have the bulk of their flux density concentrated in a small fraction of the image.

In this paper, we develop a recursive scheme for black hole ray tracing. We begin, in Section 2, by summarizing previous work related to adaptive and high-resolution black hole imaging. Next, in Section 3, we describe expected interpolation errors in black hole images and present a recursive algorithm for efficiently generating high-resolution images. In Section 4, we evaluate the performance of our adaptive ray-tracing algorithm and explore the properties of extremely high-resolution features in the image and interferometric visibility domains. We consistently generate images with $>12\times$ fewer rays and <0.1% mean squared error (MSE) compared to their truth counterparts. In Section 5, we study the specific case of photon subrings in images from general relativistic magnetohydrodynamic (GRMHD) simulations and semianalytic models, and we explore how different averaging prescriptions suppress stochastic image features. In Section 6, we summarize our results.

2. Background and Literature Review

Mathematical and computational techniques in general relativistic radiative transfer have seen tremendous developments during the past century (see, e.g., Connors & Stark 1977; Luminet 1979; Rybicki & Lightman 1979; Broderick & Blandford 2003, 2004; Shcherbakov & Huang 2011; Gammie & Leung 2012; Krawczynski 2012; Beheshtipour et al. 2017).

⁵ Waterloo Centre for Astrophysics, University of Waterloo, Waterloo, ON N2L 3G1 Canada

In order to use these techniques to render the appearance of black holes, the actual conditions of the surrounding plasma/ emitting material are often simulated via GRMHD (see, e.g., Gammie et al. 2003; Narayan et al. 2012). Ray-tracing codes can then combine the output of GRMHD simulations with a radiative transfer scheme to produce black hole images under various astrophysical conditions (see, e.g., Schnittman et al. 2006; Noble et al. 2007; Dolence et al. 2009; Shcherbakov et al. 2012).

Among the GRMHD simulations, adaptive refinement techniques have been used extensively in the past for EHT-related work. Because GRMHD simulations evolve the accretion dynamics in individual "cells" surrounding a black hole, codes can adaptively alter the size of these cells to better capture the relevant physical processes (Porth et al. 2019). However, methods of spatial refinement in the subsequent ray-tracing analysis are not as well documented, despite this method's ability to improve both image generation speed and quality.

Refinement strategies have already found some applications among ray-tracing programs. Many of the ray-tracing packages compared by Gold et al. (2020) use adaptive step sizes to boost efficiency during numerical integration of the geodesic or radiative transfer equation (see, e.g., Chan et al. 2013; Dexter 2016; Pu et al. 2016; Mościbrodzka & Gammie 2017; Bronzwaer et al. 2018).

Chan et al. (2015) implemented a multiscale sampling procedure by separately ray-tracing images with uniform gridding but different pixel sizes and then combining them. In their three-layer scheme, each successive layer had the same number of pixels but increased the field of view by a factor of 4, providing relatively high resolution near the black hole $(\Delta x = M/8)$ and also giving a complete estimate for the X-ray flux from the simulated domain.

Owing to their computational expense, high-resolution images of black holes are sparse in the growing literature related to black hole images. Notable exceptions include Bronzwaer et al. (2018), who compared their RAPTOR code with BHOSS (Younsi et al. 2020) using images with 4096 \times 4096 pixels over a field of view of 60*M*. Similarly, Davelaar et al. (2018) used RAPTOR to generate a high-resolution, virtual reality simulation near a black hole with 2000 \times 1000 pixels per snapshot.

To systematically produce high-resolution images, especially with a resolution to better resolve the substructure of the photon ring, a fully adaptive ray-tracing approach is advisable. Such an algorithm has been studied in the context of cosmological simulations (e.g., Abel & Wandelt 2002; Wise & Abel 2011). Our approach to adaptive ray-tracing more closely parallels the one implemented in the VRT² package within THEMIS (Broderick & Blandford 2003, 2004; Broderick et al. 2020), which chooses whether to ray-trace or interpolate on a pixelby-pixel basis. Examples of high-resolution black hole images with this adaptive ray-tracing scheme can be seen in, e.g., Broderick & Loeb (2006). Parkin (2011) also presented an analogous refinement model for adaptive ray tracing, although it was not developed for black hole images. Wong (2021) presented an adaptive approach that uses only the path length of null geodesics, giving a spacetime-dependent grid that is independent of the image structure. We will next present a new method for adaptive ray tracing, which differs from these previous methods in its sampling methodology, and we will

evaluate its performance using numerical simulations of black holes.

3. Interpolation and Adaptive Image Refinement

The emission from a black hole produces a smooth intensity distribution on the screen of a distant observer, which we denote I(x), where x is position. Ray tracing discretely samples I(x) at a set of specified locations $\{x_i\}$. An output image $\hat{I}(x)$ depends on the set of sampled rays $\{x_i\}$ and the interpolation method used to estimate a smooth distribution from this finite set. We now evaluate the expected interpolation errors for black hole images, and we present a recursive sampling approach that enables efficient estimates of high-resolution images.

3.1. Continuous Error Approximations

The simplest interpolation scheme to generate a full image from a discrete set of rays is a nearest-neighbors approach. Given a ray sampled at $x_1, x_2,...,x_n$, the nearest-neighbors intensity distribution simply finds the closest sampled ray at every location: $\hat{I}(\mathbf{x}) \equiv I(\arg \min_{x_i} |x_i - \mathbf{x}|)$. A second interpolation scheme is a linear/bilinear approach, which is defined so that $\hat{I}(\mathbf{x})$ is an average of its nearest surrounding rays that is weighted by distance. For a smooth intensity distribution, the errors in these two schemes are given by the first and second image spatial derivatives, respectively.

Specifically, consider a point x lying equidistant from four rays spaced evenly around x. Then, to leading order, the interpolation residuals will be (see Appendices A.1 and A.2 for full derivation)

$$|\hat{I}(\mathbf{x}) - I(\mathbf{x})| \sim \begin{cases} |\nabla \hat{I}(\mathbf{x})| \ \Delta x, & \text{Nearest} \\ \frac{1}{4} |\nabla^2 \hat{I}(\mathbf{x})| \Delta x^2, & \text{Linear.} \end{cases}$$
(1)

where Δx denotes the distance between x and each of its four surrounding rays.

In Figure 1, we plot the gradients and Laplacians for ray-traced images of both a semianalytic model and a model from the EHT GRMHD library (Event Horizon Telescope Collaboration et al. 2019e). In particular, the GRMHD model is a Magnetically Arrested Disk (MAD; Yuan & Narayan 2014) of M87^{*}, with dimensionless spin $a_* = +0.94$, inclination angle $i = 17^\circ$, and a field of view of 160 μ as (corresponding to ~44.17*M*/*D*). We also use the electron heating parameter $R_{\rm high} = 20$ (for details, see Mościbrodzka et al. 2016; Event Horizon Telescope Collaboration et al. 2019e).

The semianalytic model mirrors Test 1 of Gold et al. (2020) and models a spherically symmetric fluid distribution around a black hole located at a distance D = 7.78 kpc and with mass $M = 4 \times 10^6 M_{\odot}$. Additional parameters for this semianalytic model include spin $a_* = 0.9$, inclination angle $i = 60^\circ$, and a field of view of ~ 152.289 μ as (30M/D). For both models, we use a camera distance of $r_{\rm cam} = 10^6 M$ to mitigate small errors that may arise from the use of a pinhole camera.

In all cases, the gradient and Laplacian of the intensity increase sharply near the photon ring because of strong gravitational lensing. Indeed, Psaltis et al. (2015) proposed using sharp image gradients as a way to localize the photon ring; the Laplacian would also be an effective choice.



Figure 1. Simulated black hole images (left), their gradients (center), and their Laplacians (right). The top panels show a GRMHD snapshot, while the bottom panels show a semianalytic model (Test 1 of Gold et al. 2020). All panels are rescaled to have units of brightness temperature (by multiplying by an appropriate power of $\Delta x \approx 0.156 \ \mu as$); the center and right panels then give the expected interpolation residual over this interval (Equation (1)).

3.2. Adaptive Ray Tracing by Recursive Subdivision

The highly localized regions of Figure 1 with large derivatives suggest that judicious sampling can be used to significantly improve image generation speed and quality relative to a uniform grid. We will now describe an efficient recursive sampling procedure that is defined by a pair of error tolerances: \mathcal{R}_{abs} and \mathcal{R}_{rel} .

To begin, we ray-trace an image on a grid at a coarse initial resolution, $n_{0x} \times n_{0y}$, where n_{0x} and n_{0y} are any two integers. Next, we selectively populate an image with finer resolution by a factor of two: $(2n_{0x} - 1) \times (2n_{0y} - 1)$. For each point, we either ray-trace or interpolate based on the expected relative and absolute interpolation errors:

$$\epsilon_{\rm rel}(\mathbf{x}) = \left| \frac{I(\mathbf{x}) - \widehat{I}(\mathbf{x})}{I(\mathbf{x})} \right|,\tag{2}$$

$$\epsilon_{\rm abs}(\mathbf{x}) = \left| \frac{I(\mathbf{x}) - \hat{I}(\mathbf{x})}{\overline{I}} \right|.$$
(3)

Both ϵ_{abs} and ϵ_{rel} are dimensionless, but they differ in their normalization: ϵ_{rel} uses the local image intensity $I(\mathbf{x})$, while ϵ_{abs} uses the image-averaged intensity \overline{I} . If $\epsilon_{abs} > \mathcal{R}_{abs}$ and $\epsilon_{rel} > \mathcal{R}_{rel}$, then $I(\mathbf{x})$ is computed by ray tracing. Otherwise, $I(\mathbf{x})$ is computed by interpolation. Hence, rays are computed only where interpolation residuals are expected to be large. This procedure can then be repeated arbitrarily, giving an effective final image resolution that is $n_x \times n_y$, with

$$n_x = 2^N (n_{0x} - 1) + 1$$

$$n_y = 2^N (n_{0y} - 1) + 1.$$
 (4)

Our recursive approach relies on computing estimates $\hat{\epsilon}_{rel}(x)$ and $\hat{\epsilon}_{rel}(x)$ for the interpolation residuals, and these estimates change for different configurations of pixels. As shown in



Figure 2. Recursive gridding approach. Black dots show points sampled at the *n*th refinement level; colored dots show additional points sampled at the (n + 1)th refinement level. Distinct colors represent distinct cases for interpolation error estimation.

Figure 2, each point x lying on the (n + 1)th grid will fall into one of four categories:

- 1. Category 1 (Black): *x* had its intensity computed at the *n*th refinement level; there is no interpolation residual.
- 2. Category 2 (Blue): *x* lies in between two points located above and below, each of whose intensities were computed at the *n*th refinement level.



Figure 3. Corner pixels used to compute the interpolation uncertainties $\epsilon_{abs}(\mathbf{x})$ and $\epsilon_{rel}(\mathbf{x})$. For nearest-neighbor interpolation, only points 1–4 are needed, while linear interpolation requires points 5 – 8 as well.

- 3. Category 3 (Red): *x* lies in between two points located to left and right, each of whose intensities were computed at the *n*th refinement level.
- 4. Category 4 (Green): *x* lies in between four equidistant corner points, each of whose intensities were computed at the *n*th refinement level.

For cases 2–4, we use finite differences to estimate derivatives and then take the appropriate leading term in the Taylor series approximation to evaluate interpolation uncertainties (see, e.g., Appendix A of Pedrola 2015). For example, consider estimating the intensity at the central point of Figure 2, whose location we will call \mathbf{x} . Defining $\hat{I}_i \equiv \hat{I}(\mathbf{x}_i)$ and labeling the points as in Figure 3, we find

$$\hat{\epsilon}_{abs}(\mathbf{x}) = \begin{cases} \left| \frac{1}{\bar{I}_{int}} \sqrt{\left(\frac{\hat{I}_4 - \hat{I}_1}{2}\right)^2 + \left(\frac{\hat{I}_3 - \hat{I}_2}{2}\right)^2} \right|, & \text{Nearest} \\ \left| \frac{-(\hat{I}_1 + \hat{I}_2 + \hat{I}_3 + \hat{I}_4) + (\hat{I}_5 + \hat{I}_6 + \hat{I}_7 + \hat{I}_8)}{16\bar{I}_{int}} \right|, & \text{Linear}, \end{cases}$$
(5)

$$\hat{\epsilon}_{\rm rel}(\mathbf{x}) = \begin{cases} \left| \frac{1}{\hat{I}_1} \sqrt{\left(\frac{\hat{I}_4 - \hat{I}_1}{2}\right)^2 + \left(\frac{\hat{I}_3 - \hat{I}_2}{2}\right)^2} \right|, & \text{Nearest} \\ \left| \frac{-(\hat{I}_1 + \hat{I}_2 + \hat{I}_3 + \hat{I}_4) + (\hat{I}_5 + \hat{I}_6 + \hat{I}_7 + \hat{I}_8)}{4(\hat{I}_1 + \hat{I}_2 + \hat{I}_3 + \hat{I}_4)} \right|, & \text{Linear.} \end{cases}$$
(6)

Here, the nearest neighbor is chosen (arbitrarily) to be \hat{I}_1 . Additionally, \bar{I}_{int} is defined as the interpolated average intensity after the first pass of ray tracing (with resolution $n_{0x} \times n_{0y}$). For a detailed derivation of Equations (5) and (6), see Appendices A.1 and A.2. While we have derived estimates for $\hat{\epsilon}_{abs}(\mathbf{x})$ and $\hat{\epsilon}_{rel}(\mathbf{x})$ using the simplified approximation of a smooth image with interpolation residuals dominated by low-order derivatives, these estimates are also useful for the more general case of images with small-scale structure and sharp gradients. For example, if the small-scale image structure is stochastic with a power-law spectrum, then we can compare the ensembleaverage properties of the exact (Equation (2)) and estimated (Equation (5)) interpolation residuals. For 1D interpolation, we find

$$\frac{\langle \hat{\epsilon}_{abs}^2 \rangle^{\frac{1}{2}}}{\langle \epsilon_{abs}^2 \rangle^{\frac{1}{2}}} = \begin{cases} 2^{\beta/2-2}, & \text{Nearest} \\ 2^{\frac{\beta}{2}-4} \sqrt{\frac{1+2^{\beta-1}-3^{\beta-2}}{1-2^{\beta-4}}}, & \text{Linear}, \end{cases}$$
(7)

where β is the power-law exponent. GRMHD simulations often have $\beta \sim 2.5$, giving ratios $\langle \hat{\epsilon}_{abs}^2 \rangle^{\frac{1}{2}} / \langle \epsilon_{abs}^2 \rangle^{\frac{1}{2}}$ of 0.59 and 0.27 for the 1D nearest and linear interpolation, respectively. Thus, even in this generalized case of stochastic image fluctuations with a power-law spectrum (for which a local Taylor series expansion is poor), our approximate estimates for interpolation error will still provide useful refinement criteria. For additional discussion of the interpolation errors for images with power-law spectra, see Appendix B.

Previous uses of adaptive ray tracing have used more complicated refinement criteria than the ones we have selected. Parkin (2011) also implemented a second-order criterion (the Löhner 1987 criterion), which uses multiple cross derivatives. Similarly, VRT^2 employs a bicubic interpolator. However, we expect the benefit of these higher-order schemes to be marginal in the case of an image with a turbulent power spectrum (see Appendix B). For the remainder of this paper, we use the linear interpolation scheme.

4. Results

4.1. Models and Implementation

We implement the recursive ray-tracing scheme described above into the polarized general relativistic radiative transfer (GRRT) code IPOLE⁶ (Mościbrodzka & Gammie 2017). This implementation allows us to assess the performance of our algorithm and subsequently generate extremely high-resolution images of black holes.

In addition to the MAD model and spherical semianalytic model used in Figure 1, we also generate images using a Standard and Normal Evolution (SANE) model from the EHT GRMHD library (Yuan & Narayan 2014; Event Horizon Telescope Collaboration et al. 2019e), as well as a semianalytic model of a geometrically thin, rotating disk (Test 5 of Gold et al. 2020).

Before generating images, however, we must first determine suitable error tolerances \mathcal{R}_{abs} and \mathcal{R}_{rel} . To evaluate the effect of different tolerances on the resultant image, we use the following three metrics (the second and third of which are also used in Gold et al. 2020):

Interpolation Fraction
$$\equiv \frac{\#\text{pixels interpolated}}{n_x \times n_y}$$
, (8)

⁶ https://github.com/moscibrodzka/ipole. Our adaptive tools were developed at https://github.com/AFD-Illinois/ipole.



Figure 4. Contour plots representing how the error metrics depend on the tolerances \mathcal{R}_{rel} and \mathcal{R}_{abs} for the MAD GRMHD snapshot traced with 1025×1025 pixels on 160 μ as FOV. Colors represent interpolation fraction while dashed lines represent the error (flux error on left and mean squared error on right). The star indicates the error tolerances we choose to generate the subsequent GRMHD images in this paper ($\mathcal{R}_{rel} = 0.001$ and $\mathcal{R}_{abs} = 0.025$).

Flux Error
$$\equiv \left| \frac{F - \hat{F}}{F} \right|,$$
 (9)

Mean Squared Error
$$\equiv \frac{\sum_{\text{pixels}i} |I(\mathbf{x}_i) - \hat{I}(\mathbf{x}_i)|^2}{\sum_i |I(\mathbf{x})|^2}.$$
 (10)

Here, $F \equiv \int d^2 \mathbf{x} I(\mathbf{x})$ is the image flux associated with the model intensity distribution, and $\hat{F} \equiv \int d^2 \mathbf{x} \hat{I}(\mathbf{x})$ is the image flux associated with the estimated intensity distribution. We note that while a high flux error necessarily implies a high MSE, we include the former metric because it has a clear physical interpretation and bears more relevance in potential applications to light curves.

Our goal is to maximize the interpolation fraction (IF) while minimizing the flux error (FE) and MSE, as this achieves both efficiency and accuracy. To this end, we run our adaptive scheme on a single GRMHD snapshot (the same snapshot/ parameters used in Figure 1) with a wide variety of error tolerances, and we evaluate the IF/FE/MSE for each. The contours in Figure 4 show how the error metrics each depend on \mathcal{R}_{rel} and \mathcal{R}_{abs} . For this comparison, we take $n_{0x} = n_{0y} = 65$ pixels and $n_x = n_y = 1025$ pixels, and we compute the model intensity distribution $I(\mathbf{x})$ by ray-tracing each pixel on a 1025×1025 grid.

As expected, as \mathcal{R}_{rel} and \mathcal{R}_{abs} increase, the adaptively sampled image $\hat{I}(\mathbf{x})$ increasingly deviates from the fully sampled $I(\mathbf{x})$. As $\mathcal{R}_{rel} \to \infty$ and $\mathcal{R}_{abs} \to \infty$, all adaptive refinement will cease and the image will be given by the initial sampling. For the initial 65×65 sampling of this snapshot, $MSE \approx 0.22 \approx 10^{-0.65}$ and $FE \approx 0.029 \approx 10^{-1.54}$. These serve as reference values to compare against iterations of the scheme with more stringent tolerances.

The contours in Figure 4 can be used for practical image generation purposes—for the desired accuracy, one can select \mathcal{R}_{abs} and \mathcal{R}_{rel} such that the interpolation fraction and hence the efficiency are maximized. For our subsequent analysis of GRMHD images, we choose $\mathcal{R}_{abs} = 0.025$ and $\mathcal{R}_{rel} = 0.001$ (the red star in Figure 4), as this set of tolerances allows for

>90% interpolation while constraining the MSE and FE to <0.1%.

For subsequent analysis of the semianalytic models, we choose $\mathcal{R}_{abs} = \mathcal{R}_{rel} = 0.001$, as we find that the semianalytic models are slightly more sensitive to the absolute tolerance.

4.2. High-resolution Images and Analysis

Using the tolerances selected described above, we adaptively ray-trace in IPOLE at resolutions of 1025×1025 and 4097×4097 pixels. Images of the four models for the latter resolution are shown in Figure 5, along with their respective "ray-density plots," illustrating where the program is concentrating rays in each image. The ray-density plots are generally consistent with Figure 1—the majority of ray tracing takes place in an annulus about the photon ring, while the background and shadow are predominantly interpolated. As with Figure 1, the annuli for the semianalytic models are significantly thinner and more pronounced.

The corresponding residual images (fully ray-traced images subtracted from adaptively ray-traced images and normalized by the image-averaged intensity: $\left|\frac{\hat{I}(x) - I(x)}{\bar{I}}\right|$) are shown in Figure 6. In all images, the residual amplitudes are relatively constant, which reflects the refinement goal of ϵ_{abs} . In portions of the image that correspond to rays that have been traced, the residuals are zero. In particular, all images contain a base grid of 65×65 evenly spaced rays, giving evenly spaced nulls in the residuals that align with this grid. This feature is most evident in the semianalytic models, where the smooth intensity distribution allows significant interpolation.

The resultant error statistics (Equations (8)–(10)) for these adaptively ray-traced images are presented in Table 1. The 1025×1025 images match the predictions from the contours in Figure 4: both the MSE and FE are lower than 0.1%, while the interpolation fraction exceeds 90%. In Table 1, we also list the total number of rays traced in each image: $(1 - \text{IF}) \times n_x \times n_y$, which is roughly proportional to computational expense.

While these statistics were generated for a specific choice of image parameters, we obtain similar results for other models. We ray-trace with the same tolerances on all combinations of The Astrophysical Journal, 912:39 (14pp), 2021 May 1

Gelles et al.



Figure 5. Adaptively ray-traced images, shown with ray-density plots to the right. Rays are concentrated near the photon ring, with occasional rays scattered across the diffuse emission. To compute the ray density, we use a kernel size equal to the spacing between pixels at the zeroth refinement level: 17×17 pixels for the lower-resolution image and 65×65 pixels for the higher-resolution image.



Figure 6. Normalized residual images for the 4097 × 4097 snapshots of the GRMHD models and semianalytic models. Regions with a high concentration of rays have particularly small residuals, as do regions with a diffuse structure. For reference, the image-averaged intensities \overline{I} for the four images are (after converting to brightness temperature): 5.2×10^8 K, 7.0×10^8 K, 1.6×10^9 K, and 2.5×10^7 K, from left to right.

6

Table 1Statistics for 1025×1025 and 4097×4097 Images

Model—1025 ²	IF	FE	MSE	# Rays
GRMHD—MAD	0.92	$2.5 imes 10^{-4}$	$4.6 imes 10^{-4}$	8.7×10^{4}
GRMHD—SANE	0.92	$2.3 imes 10^{-4}$	$2.5 imes 10^{-4}$	$9.0 imes 10^4$
SEMI-SPHERICAL	0.97	$1.5 imes 10^{-5}$	$5.6 imes 10^{-7}$	2.9×10^4
SEMI—DISK	0.96	$1.5 imes 10^{-4}$	$1.3 imes 10^{-5}$	4.3×10^4
Model-4097 ²	IF	FE	MSE	# Rays
GRMHD—MAD	0.96	$4.2 imes 10^{-4}$	$1.5 imes 10^{-4}$	7.1×10^{5}
GRMHD—SANE	0.95	$2.5 imes 10^{-4}$	$8.3 imes 10^{-5}$	7.9×10^5
SEMI-SPHERICAL	0.99	$2.1 imes 10^{-5}$	$1.8 imes 10^{-7}$	1.0×10^{5}
SEMI—DISK	0.99	$1.3 imes 10^{-4}$	$3.7 imes 10^{-6}$	1.9×10^5

Note. Here, FE and MSE give the values for an adaptively ray-traced image at the specified resolution relative to a fully ray-traced image at that resolution. The last column shows the number of rays traced in the simulation. All GRMHD images used $\mathcal{R}_{rel} = 0.001$ and $\mathcal{R}_{abs} = 0.025$, while semianalytic images used $\mathcal{R}_{rel} = \mathcal{R}_{abs} = 0.001$.

MAD/SANE, $a_* = -0.94$, 0, + 0.94, and $R_{high} = 20$, 40, 80. We find that among these images, IF ranges from 0.877 to 0.929, FE ranges from 3.02×10^{-5} to 8.59×10^{-4} , and MSE ranges from 2.82×10^{-4} to 5.55×10^{-4} .

We may further vary the magnetization (σ) cut (see, e.g., Chael et al. 2019), which is a quantity designed to restrict emission to regions where the fluid density has not been physically invalidated. Regenerating the MAD image from Figure 1 with $\sigma = 1, 5, 10$, and 50, we find that IF ranges from 0.917 to 0.931, FE ranges from 2.50×10^{-4} to 7.50×10^{-4} , and MSE ranges from 2.79×10^{-4} to 4.69×10^{-4} . Thus, for these GRMHD simulations, the accuracy and efficacy of the adaptive refinement criteria are relatively insensitive to the choice of accretion and radiation model parameters.

The semianalytic models have lower errors than the GRMHD models and higher interpolation fractions. This is likely because the smooth underlying structure in semianalytic models is better suited to interpolation.

While IF is a useful metric to quantify sampling efficiency, it does not directly correspond to the reduction in image generation time. Namely, rays near the photon ring are more expensive to trace, as their geodesics have longer paths through the emitting material and thus require more calculations to perform the radiative transfer. Nevertheless, even though our adaptive scheme predominantly samples in this computationally expensive region, we find that the reduction in image generation time is similar to the interpolation fraction (see Figure 7). In general, the relationship between interpolation fraction and computational expense will depend on details of the underlying model and of the GRRT implementation.

We note additionally from Table 1 that as resolution increases, IF increases as well. This behavior is expected because the error estimates decrease at smaller separations, allowing more pixels to satisfy the error tolerances required for interpolation. Indeed, Figure 5 shows that the density of rays in the diffuse parts of the image quickly saturates—upon increasing the spatial resolution from 1025×1025 to 4097×4097 , the sampled ray density only increases in the photon ring region.

Unlike the IF, however, the MSE and FE do not depend strongly on resolution. These quantities instead depend on the tolerances \mathcal{R}_{rel} and \mathcal{R}_{abs} . Thus, even though more pixels are



Figure 7. Adaptive ray-tracing speedup factor as a function of image resolution for the MAD GRMHD model shown in Figure 1. The red line shows the full-to-adaptive ratio of runtime, while the blue line shows the full-to-adaptive ratio of the number of rays. This latter quantity represents the idealized speedup factor if all pixels took equally long to ray-trace. Here, "full" refers to an image that was ray-traced at every pixel (i.e., $\hat{I}(\mathbf{x}) = I(\mathbf{x})$ at all pixels on the image).

interpolated at a resolution of 4097×4097 compared to a resolution of 1025×1025 , the output images have comparable accuracy by these two metrics.

4.3. Visibility-domain Analysis

VLBI directly measures interferometric visibilities, which correspond to complex Fourier components of the sky image (Thompson et al. 2017). Hence, visibility-domain tests are appropriate to assess suitability for direct comparisons with observables. In this section, we analyze the visibility spectra of our adaptively ray-traced images and show that they exhibit the universal properties expected for black hole images.

4.3.1. Expected Visibility Signatures

For all images (both GRMHD and semianalytic), we expect the visibility spectra to reflect signatures of the strong gravitational lensing of light. Namely, in the Kerr spacetime, photons can complete spherical orbits at a fixed set of Boyer– Lindquist radii, and null geodesics near these orbits approach a "critical curve" on an observer's screen upon eventual escape (Bardeen 1973; Teo 2003). Rays that terminate increasingly near the critical curve make increasingly many revolutions around the black hole, producing a bright "photon ring" when the emitting material is optically thin. The visibility spectrum of such a ring should exhibit damped oscillatory behavior with a characteristic period of 1/d, where d is the photon ring's screen diameter projected along the baseline direction (Johnson et al. 2020).

A photon that ends up within the photon ring may be further labeled by a number n representing the number of half-orbits



Figure 8. Visibilities for the adaptively ray-traced MAD model, spherical semianalytic model, and disk semianalytic models. Left: |V(u, 0)| versus u on a log-linear scale. Middle: |V(u, 0)| versus u on a log-log scale. Right: |V(u, v)| as a 2D plot. Visibility residuals are shown in red.

the photon has taken around the black hole between emission and reception at the observer. In the case of a geometrically thin disk of emitting plasma, the photon ring naturally decomposes into a series of overlapping, self-similar "subrings" indexed by n, where each subring comprises the set of all photons labeled n. Because the subrings are exponentially demagnified, the flux from each successive subring will dominate the visibility spectrum for a range of baselines that sample angular scales matched to those of the subring (Johnson et al. 2020).

In addition to the effects of gravitational lensing, we also expect to identify the signature of accretion turbulence present in GRMHD images. These will produce stochastic visibility noise that may be described by their power spectrum (i.e., the squared visibility amplitude).

4.3.2. Visibility Spectra of High-resolution Images

We now explore the expected long-baseline visibility signatures using our high-resolution images computed with adaptive ray tracing. While short interferometric baselines will have a complex visibility structure that depends on the overall image morphology, the visibility signatures from the photon ring and from turbulence will emerge on baselines that heavily resolve the image. Specifically, to accurately estimate the visibility on a baseline with dimensionless length u requires angular resolution $\Delta\theta \sim 1/u$ (Thompson et al. 2017). For our images of both the GRMHD and semianalytic models, the visibility spectra can thus be computed to baselines of $u_{\text{max}} \gtrsim 5000 \text{ G}\lambda$. This value is approximately 600 times larger than the longest current EHT baselines (Event Horizon Telescope Collaboration et al. 2019b). However, in the following analysis, to minimize errors from finite pixel size, we only analyze baselines shorter than 1000 G λ .

Figure 8 shows visibility amplitudes for the adaptively raytraced MAD model, the spherically symmetric semianalytic model, and the disk semianalytic model. Because the semianalytic images have significant emission extending beyond our specified FOV, the sharp artificial cutoffs at the edges of the FOV produce spurious high-frequency visibility power. To suppress this artificial power, we double the FOV and pixel number (keeping the image resolution fixed) and then apply a Gaussian taper with FWHM of FOV/4 before computing visibilities. Figure 8 also shows the visibility errors resulting from the adaptive ray tracing, demonstrating that these errors are a small fraction of the visibility amplitudes on all baselines. Specifically, these errors correspond to the Fourier amplitudes of the residual images defined in Section 4.2.



Figure 9. Visibilities along u and v axes for subring-decomposed disk semianalytic model.

The MAD visibility spectrum (top row) possesses many kinks and does not display a clear-cut pattern. Although a characteristic periodicity may be evident, small-scale image power from turbulence exceeds that of the lensed emission. This turbulent power gives a visibility "noise" that decays on long baselines approximately as $V(u) \sim u^{-p}$ with $p \approx 1.2$.

The spherical semianalytic model (middle row), on the other hand, displays a smooth, turbulence-free spectrum. Just as distinct subrings are not visible in the images of this model due to the spherical symmetry of the fluid distribution (e.g., Narayan et al. 2019), distinct subrings are not visible in the visibilities of this model.

In contrast, the disk semianalytic model (bottom row) does show clear signs of distinct photon subrings in both the images and visibilities. The spectrum falls steeply around ~ 500 G λ before flattening again shortly thereafter, corresponding to the transition between the n = 1 and n = 2 subrings.

5. Applications to High-resolution Science

In this section, we discuss specific applications of adaptive ray tracing to larger problems in black hole simulations. In particular, we show that by generating images with extremely fine resolution, adaptive ray tracing presents a useful tool to examine the substructure of the photon ring and to explore signatures of turbulence in simulations.



Figure 10. Adaptively ray-traced MAD model decomposed into subrings. Images for n = 1, 2, and 3 have resolution $32,769 \times 32,769$ across a field of view of 80 μ as, while the n = 0 image has a resolution of 16385×16385 . The sum of all four subrings is shown in the lower panel. All rings are visible, although the colorbar for the n = 3 image needs a larger range of values so that we can see dim pixels. At this resolution, artifacts related to limitations in the underlying simulation are apparent in the n = 0 image.

5.1. Images of Subrings

Resolving the photon subrings in the image domain is difficult in practice due to the exponential radial demagnification of each successive subimage (Darwin 1959; Luminet 1979; Ohanian 1987). In particular, per Johnson et al. (2020), the width of each subring on the screen scales as $w_n \sim w_0 e^{-\gamma n}$, where the Lyapunov exponent γ is defined as

$$\gamma \equiv \sqrt{\frac{2\mathcal{R}''(r_{\rm c})}{-u_{-}a^2E^2}} K\left(\frac{u_{+}}{u_{-}}\right). \tag{11}$$

Here, *K* is the complete elliptic integral of the first kind (with squared modulus u_+/u_-), $\mathcal{R}(r)$ is the radial effective potential for null geodesics in Kerr, *E* is the conserved energy, r_c is the corresponding radius of the spherical orbit, and u_{\pm} denote roots of the angular effective potential. Furthermore, the radial curve ρ_n of each successive subring exponentially approaches the critical curve ρ_c : $\delta \rho_n \approx \rho_c e^{-\gamma n}$. Hence, the image resolution



Figure 11. Stacked intensities at four locations on the ring for the same snapshot as Figure 10. Different colors show subrings for n = 0 (red), 1 (blue), 2 (green), and 3 (yellow), illustrating a realization of the photon ring substructure. The contributions from individual subrings most clearly align on the bottom and right sides of the image, although the right half of the image is Doppler deboosted due to the orientation of the black hole spin.

required to resolve each successive subring thus grows exponentially for a fixed FOV.

To explore the properties of photon subrings explicitly, we combine our adaptive scheme with a subring decomposition code. The decomposition code can generate images corresponding to the *n*th subring by only including emission from the appropriate segment of the full geodesic. Notice that this definition of subrings differentiates between photons that may follow the same geodesic: the geodesic for a photon that makes n half-orbits will overlap with the geodesic for some other photon that makes n-1 half-orbits. Thus, a pixel that is illuminated in the n=2 image may also be illuminated in the n=1 one, and the total pixel brightness will have contributions from the second and first subrings, respectively (see, e.g., Figure 3 of Johnson et al. 2020).

Figure 9 shows the visibility spectra of the disk semianalytic model decomposed into contributions from individual subrings. These reveal a new feature in the subring visibilities: the different subring widths between the top and bottom of the image lead to intermittent beating along the v axis between the n and n + 1 subrings. When the former is resolved, the beating vanishes, and the visibilities smoothly decay until reaching the level of power from the next subring. This phenomenon does not appear on the u axis, as the ring does not have significant thickness asymmetry on the horizontal axis.

To illustrate the presence of subrings at high resolutions, we again use the adaptive subring decomposition code to generate 32769×32769 images of the n = 1, 2 and 3 subrings of the

MAD model, now with a FOV of 80 μ as. We show these images along with a lower-resolution image of n = 0 in Figure 10. Figure 11 shows stacked cross sections of the brightness profiles and zoom in on four points of the image (bottom, top, left, right). Individual subrings become thinner, with approximately constant peak brightness. Thus, the peak brightness of the sum of the first *n* subimages is approximately proportional to (n + 1).

5.2. Time and Visibility Averaging

We can use adaptive ray tracing to generate high-resolution movies, which can be used to study how various averaging techniques reduce the turbulent noise present in GRMHD visibility spectra. Reducing turbulent noise is necessary to reveal universal signatures of the lensed emission surrounding the black hole, which reflect the purely geometrical properties of the spacetime. Additionally, by quantifying the amount of turbulence present in these simulations, we can obtain a better understanding of the accretion dynamics surrounding black holes and the timescales over which they vary.

In Figure 12, we show the visibility spectra for the MAD model with $t_{avg} = 0M$ (i.e., a single snapshot), as well as $t_{avg} = 100M$ and $t_{avg} = 500M$. The images were generated with a resolution of 1025×1025 and are thus capable of resolving the n = 0 and n = 1 subrings, whose individual visibility spectra are shown in green and blue, respectively.

We see that by an averaging scale of $t_{avg} = 100M$, the visibility spectrum of the full image has mostly converged to



Figure 12. Single snapshot visibilities, along with time-averaged visibilities for scales of t = 100M and t = 500M. The image cadence for this simulation is 5M.

that of the n = 1 subring for $u \leq 100 \text{ G}\lambda$. However, the n = 0 spectrum still neighbors that of n = 1, indicating that turbulent noise is still significant. By 500*M*, the n = 0 spectrum falls below the n = 1 spectrum in the region where we expect n = 1 to dominate the time-averaged spectrum.

6. Summary

In this paper, we have presented a recursive algorithm for adaptive ray tracing, with natural applications to current and future high-resolution black hole imaging efforts. Whereas most conventional ray-tracing programs spread rays evenly across a uniform grid, our method preferentially samples rays in regions of an image with small-scale structure. When applied to both GRMHD and semianalytic models, we find that our algorithm reduces the time required to generate images by an order of magnitude or more.

We then use this code to generate images with resolutions of 1025×1025 , 4097×4097 , and $32,769 \times 32,769$ pixels. These images directly visualize the fine structure present in both the accretion flow and the photon ring, revealing the n = 1, 2, and 3 subrings. Finally, we explored the utility of time-averaging in reducing stochastic noise in high-resolution images.

While our algorithm reduces the computational expense required to produce high-resolution images, significant limitations in the physical modeling remain. Although we verified that the magnetization (σ) cut negligibly alters the MSE of the adaptively ray-traced image, a sharp cut on σ may introduce a spurious high-frequency image power. Moreover, for the images generated in Section 5, we cut at $\sigma = 1$, but this excludes emission near the jet in MAD models and requires additional study (see, e.g., Chael et al. 2019).

We are also fundamentally limited in resolution by the MHD cell size. For the MAD model, we use a GRMHD simulation on a spherical polar grid with a resolution of $384 \times 192 \times 192$ in the radial, polar, and azimuthal directions, respectively (zones are compressed exponentially toward the event horizon and lightly toward the midplane). For the SANE model, we use a resolution of $288 \times 128 \times 128$. The finite MHD cell size may introduce an unphysical, high-power noise from sharp boundaries and will not reproduce subgrid turbulent power.

Future applications of adaptive ray tracing could extend our results to selectively sample the time and frequency domains. For the purposes of this study, we used one frequency (230 GHz), but fine-scale frequency structure is expected from GRMHD simulations (Ricarte et al. 2020). Adaptive sampling

in time would allow efficient generation of high-resolution movies from numerical simulations.

We have integrated this approach into IPOLE, but it should be compatible with any GRRT scheme that ray-traces on a rectangular grid. And while we have used unpolarized transport to generate the images in this paper, the approach generalizes to polarized images as well by simply replacing the total intensity I(x) with Stokes parameters Q(x), U(x), and V(x). Highly lensed structure near the critical curve resolved with adaptive ray tracing may show interesting, spin-dependent symmetries in the polarization (Himwich et al. 2020).

We thank Chi-Kwan Chan for helpful suggestions on our manuscript, and we thank Avery Broderick for useful discussions. We thank the National Science Foundation (AST-1716536, AST-1440254, AST-1935980, OISE-1743747) and the Gordon and Betty Moore Foundation (GBMF-5278) for financial support of this work. G.N.W. was supported by a Donald C. and F. Shirley Jones Fellowship. This work was supported in part by the Black Hole Initiative, which is funded by grants from the John Templeton Foundation and the Gordon and Betty Moore Foundation to Harvard University. Funding for this project was provided in part by the Harvard College Program for Research in Science and Engineering.

Software: eht-imaging library (Chael et al. 2016), IPOLE (Mościbrodzka & Gammie 2017), Numpy (Harris et al. 2020), Matplotlib (Hunter 2007).

Appendix A Estimating Relative and Absolute Errors

Here, we derive the approximations for $\epsilon_{abs}(\mathbf{x})$ and $\epsilon_{rel}(\mathbf{x})$ presented in Equations (5) and (6), respectively. For convenience, we refer to the nearest-neighbor errors as ϵ^{NN} and to the linear errors as ϵ^{lin} .

A.1. Nearest-neighbors Interpolation

Suppose we wish to interpolate the intensity at x directly from its nearest neighbor at x_1 . Because I(x) and $\hat{I}(x)$ are smooth in between x and x_1 , then we may apply Taylor's theorem (or in this case, just the mean value theorem) to see that

$$\widehat{I} - I = \nabla \widehat{I}(p) \cdot (x - x_1), \qquad (A1)$$

for a point p lying on the line segment connecting x to x_1 . To leading order in $|x - x_1|$, we may replace p with x, giving

$$\hat{I} - I \approx \nabla \hat{I}(\mathbf{x}) \cdot (\mathbf{x} - \mathbf{x}_1).$$
 (A2)

Let us now restrict our attention to the rectangular gridding scheme presented in Figure 2. On this grid, Equation (A2) is ambiguous, as \mathbf{x} will be adjacent to multiple equidistant pixels, rendering the location of \mathbf{x}_1 ill defined. Regardless of where we choose to set \mathbf{x}_1 , however, the interpolation residual will be extremized by the quantity $|\nabla \hat{I}(\mathbf{x})||\mathbf{x} - \mathbf{x}_1|$. So, defining $\Delta x \equiv |\mathbf{x} - \mathbf{x}_1|$, we take

$$|I(\mathbf{x}) - \hat{I}(\mathbf{x})| \sim |\nabla \hat{I}(\mathbf{x})| \Delta x, \qquad (A3)$$

which is the first half of Equation (1). This expression is now symmetric—it does not depend on which of the equidistant pixels we define to be x's nearest neighbor.

Because we only have access to I(x) at discretely sampled rays, we approximate the gradient using finite differences, which requires an examination of each of the four categories of pixel locations listed in Section 3.2. Let us suppose that x falls into Category 4, as Categories 1–3 are just simplifications thereof. The arrangement of pixels for Category 4 is shown in Figure 3.

Because $|\nabla I|$ is rotationally invariant, we may evaluate the partial derivatives along rotated axes x' and y' that are aligned with the corner pixels. Then adopting the same notation as Figure 3, the central difference approximations to the derivatives are (e.g., Appendix A of Pedrola 2015)

$$\frac{\partial \hat{I}}{\partial x'} |_{\mathbf{x}} \approx \frac{\hat{I}_4 - \hat{I}_1}{2\Delta x},$$
$$\frac{\partial \hat{I}}{\partial y'} |_{\mathbf{x}} \approx \frac{\hat{I}_3 - \hat{I}_2}{2\Delta x}.$$
(A4)

Because $\hat{I}(\mathbf{x}) \equiv \hat{I}_1$, the nearest-neighbor error estimates become

$$\epsilon_{\rm abs}^{\rm NN}(\mathbf{x}) \approx \left| \frac{1}{\bar{I}_{\rm int}} \sqrt{\left(\frac{\hat{I}_4 - \hat{I}_1}{2}\right)^2 + \left(\frac{\hat{I}_3 - \hat{I}_2}{2}\right)^2} \right|,$$

$$\epsilon_{\rm rel}^{\rm NN}(\mathbf{x}) \approx \left| \frac{1}{\hat{I}_1} \sqrt{\left(\frac{\hat{I}_4 - \hat{I}_1}{2}\right)^2 + \left(\frac{\hat{I}_3 - \hat{I}_2}{2}\right)^2} \right|, \qquad (A5)$$

which are the first half of Equations (5) and (6).

A.2. Linear Interpolation

For linear interpolation on an arbitrary 2D grid, the error estimate does not reduce to an equation as simple as A2. However, by restricting our attention again to the specific gridding scheme in Figure 2, we are able to derive a straightforward error estimate as follows.

Suppose that x falls into Category 4, as once again, Categories 1–3 will just be simplifications thereof. Adopting the same labels as Figure 3, the linear interpolation residual is

Gelles et al.

explicitly given by

$$\hat{I} - I = \frac{1}{4} \left(\sum_{i=1}^{4} \hat{I}(\mathbf{x}_{i}) \right) - I(\mathbf{x})$$

=
$$\frac{\left(\frac{1}{2} [\hat{I}(\mathbf{x}_{1}) + \hat{I}(\mathbf{x}_{4})] - I(\mathbf{x}) \right) + \left(\frac{1}{2} [\hat{I}(\mathbf{x}_{2}) + \hat{I}(\mathbf{x}_{3})] - I(\mathbf{x}) \right)}{2}.$$
(A6)

We recognize this quantity as the average of two 1D linear interpolation residuals along the x' and y' axes. In 1D, linear interpolation residuals scale with the second spatial derivative (see, e.g., Theorem 4.3 of Epperson 2013), with the leading order expansion $|\hat{I}_{1D} - I| \approx \frac{1}{2} |\hat{I}''(\mathbf{x})| \Delta x^2$. Plugging this into the numerator of (A6), the 2D residual becomes

$$\hat{I} - I = \frac{\frac{1}{2}\hat{I}_{x'x'}(\mathbf{x})\Delta x^2 + \frac{1}{2}\hat{I}_{y'y'}(\mathbf{x})\Delta x^2}{2} = \frac{1}{4}\nabla^2 \hat{I}(\mathbf{x})\Delta x^2,$$
(A7)

where in the last step, we used the rotational invariance of the Laplacian. This is the second half of Equation (1).

The final task is now to approximate the Laplacian with finite differences. In doing so, we must be careful not to break the symmetry of the error approximation, and we must rely only on the pixels in Figure 3. To this end, we approximate the second derivatives by averaging the second-order central differences on either side of x. This gives

$$\frac{\partial^2 \hat{I}}{\partial x'^2} \approx \frac{\hat{I}_8 - \hat{I}_4 - \hat{I}_1 + \hat{I}_5}{4\Delta x^2},$$
$$\frac{\partial^2 \hat{I}}{\partial y'^2} \approx \frac{\hat{I}_7 - \hat{I}_3 - \hat{I}_2 + \hat{I}_6}{4\Delta x^2}.$$
(A8)

And because $\hat{I}(\mathbf{x}) \equiv \frac{1}{4}(\hat{I}_1 + \hat{I}_2 + \hat{I}_3 + \hat{I}_4)$, the error estimates become

$$\epsilon_{abs}^{lin}(\mathbf{x}) \approx \left| \frac{-(\hat{l}_{1} + \hat{l}_{2} + \hat{l}_{3} + \hat{l}_{4}) + (\hat{l}_{5} + \hat{l}_{6} + \hat{l}_{7} + \hat{l}_{8})}{16\bar{l}_{int}} \right|,$$

$$\epsilon_{rel}^{lin}(\mathbf{x}) \approx \left| \frac{-(\hat{l}_{1} + \hat{l}_{2} + \hat{l}_{3} + \hat{l}_{4}) + (\hat{l}_{5} + \hat{l}_{6} + \hat{l}_{7} + \hat{l}_{8})}{4(\hat{l}_{1} + \hat{l}_{2} + \hat{l}_{3} + \hat{l}_{4})} \right|,$$
(A9)

which is the second half of Equations (5) and (6).

Appropriate modifications to the scheme are made when the pixels are close to the edge of the image. In this region, we may not have access to $\hat{I}_1 - \hat{I}_8$.

Appendix B Interpolation Errors for Images with a Power-law Fluctuation Spectrum

Our analysis above is appropriate for functions that are smooth and are dominated by linear or quadratic variations in a neighborhood comparable to the final pixel size. More generally, we can describe intensity fluctuations $\Delta I(\mathbf{x})$ by their power spectrum, $Q(\mathbf{u}) \equiv \langle |\Delta V(\mathbf{u})|^2 \rangle$, and we can quantify expected interpolation errors statistically. By the Wiener-Khinchin theorem, the power spectrum of the intensity fluctuations is related to the two-point correlation function $C(\mathbf{x})$ via a Fourier transform:

$$C(\mathbf{x}) \equiv \langle \Delta I(\mathbf{x}_0) \Delta I(\mathbf{x}_0 + \mathbf{x}) \rangle, \qquad (B.1)$$

$$Q(\boldsymbol{u}) \equiv \left\langle \left| \int d^2 \boldsymbol{x} \ \Delta I(\boldsymbol{x}) e^{-2\pi i \boldsymbol{u} \cdot \boldsymbol{x}} \right|^2 \right\rangle$$
$$= \int d^2 \boldsymbol{x} \ C(\boldsymbol{x}) e^{-2\pi i \boldsymbol{u} \cdot \boldsymbol{x}}. \tag{B.2}$$

It is also convenient to define the second-order structure function of the intensity fluctuations,

$$D(\mathbf{x}) \equiv \langle [\Delta I(\mathbf{x}_0) - \Delta I(\mathbf{x}_0 + \mathbf{x})]^2 \rangle$$

= 2[C(0) - C(\mathbf{x})]. (B.3)

For a power spectrum determined by a single, unbroken power law, $Q(\boldsymbol{u}) \propto |\boldsymbol{u}|^{-\beta}$ and $D(\boldsymbol{x}) \propto |\boldsymbol{x}|^{\beta-2}$.

We can express interpolation errors under various schemes in terms of these functions. For instance, the rms absolute error for nearest-neighbor interpolation over a displacement x is

$$\epsilon_{\text{abs}}^{\text{NN}}(\boldsymbol{x}) \propto \langle [\Delta I(\boldsymbol{x}_0) - \Delta I(\boldsymbol{x}_0 + \boldsymbol{x})]^2 \rangle^{1/2}$$

= $\sqrt{D(\boldsymbol{x})}$.
 $\propto |\boldsymbol{x}|^{\beta/2-1}$. (B.4)

For comparison, if a function is smooth and is dominated by linear errors, then $\epsilon_{abs}^{NN}(\mathbf{x}) \propto |\mathbf{x}|$ (see Equation (1)). Hence, when $\beta < 4$, turbulent fluctuations will dominate over errors from interpolating the smooth underlying image in the limit that $|\mathbf{x}| \to 0$. On angular scales that are relevant for interpolation, black hole images are expected to show a shallower spectrum due to turbulent fluctuations in the accretion flow (see, e.g., Balbus & Hawley 1991). For $\beta \sim 2.5$, one has $\epsilon_{\rm abs}^{\rm NN}({\pmb x}) \propto |{\pmb x}|^{1/4}$, in contrast to the factor of $|{\pmb x}|$ predicted by Equation (1).

We can also compare the relative error for linear and nearest interpolation strategies. For 1D linear interpolation, we have

$$\epsilon_{abs}^{lin}(\mathbf{x}) \propto \left\langle \left[\frac{\Delta I(\mathbf{x}_0 - \mathbf{x}) + \Delta I(\mathbf{x}_0 + \mathbf{x})}{2} - \Delta I(\mathbf{x}_0) \right]^2 \right\rangle^{1/2} \\ = \sqrt{\frac{1}{2}D(\mathbf{x}) + \frac{1}{2}C(2\mathbf{x}) - C(\mathbf{x}) + \frac{1}{2}C(\mathbf{0})} \\ = \sqrt{D(\mathbf{x}) - \frac{1}{4}D(2\mathbf{x})} \\ = \sqrt{(\epsilon_{abs}^{NN})^2 - \frac{1}{4}D(2\mathbf{x})}.$$
(B.5)

In the case of a power-law spectrum, for the limit $|\mathbf{x}| \rightarrow 0$, we obtain

$$\frac{\epsilon_{\text{abs}}^{\text{in}}(\mathbf{x})}{\epsilon_{\text{abs}}^{\text{NN}}(\mathbf{x})} = \sqrt{1 - 2^{\beta - 4}}.$$
(B.6)

Thus, for turbulent spectra, the improvement of linear interpolation relative to nearest interpolation is rather modest and (unlike the Taylor series analysis) is independent of the interpolated distance x. For $\beta \sim 2.5$, the linear interpolation error is only smaller than the nearest-neighbor interpolation error by a factor of $(1 - 2^{-3/2})^{1/2} \approx 0.8$.

In short, small-scale turbulence in black hole accretion flows may lead to departures from the error estimates expected for a smooth image, with higher-order interpolation schemes giving less improvement than expected. For instance, while increasing the image resolution by a factor of 10 would decrease residuals by a factor of ~ 10 for nearest-neighbor interpolation and ~ 100 for linear interpolation of a smooth image, it may only decrease residuals by a factor of ~ 2 in turbulent regions of a GRMHD image.

Unlike this statistically isotropic noise model, black hole images are restricted to a finite domain, their stochastic noise is not isotropic, and their power spectra are scale dependent. The primary effect of a finite domain is to introduce correlations among different frequencies (i.e., different baselines will measure correlated fluctuations, with a correlation length given roughly by the inverse spatial extent of the image structure). The effect of position-dependent power spectra will be to blend physically distinct sources of image noise in the visibility domain. A scale-dependent power law will give interpolation errors over an angular interval x that are primarily sensitive to the behavior near $Q(\mathbf{u} = \hat{x}/|\mathbf{x}|)$. Thus, we do not expect any of these effects to seriously modify our conclusions about interpolation errors from image stochasticity.

ORCID iDs

- Z. Gelles https://orcid.org/0000-0001-8053-4392
- D. C. M. Palumbo () https://orcid.org/0000-0002-7179-3816
- M. D. Johnson ^(b) https://orcid.org/0000-0002-4120-3029
- G. N. Wong https://orcid.org/0000-0001-6952-2147

References

- Abel, T., & Wandelt, B. D. 2002, MNRAS, 330, L53
- Balbus, S. A., & Hawley, J. F. 1991, ApJ, 376, 214
- Bardeen, J. M. 1973, Les Astres Occlus (New York: Gordon & Breach), 215
- Beckwith, K., & Done, C. 2005, MNRAS, 359, 1217
- Beheshtipour, B., Krawczynski, H., & Malzac, J. 2017, ApJ, 850, 14
- Broderick, A., & Blandford, R. 2003, MNRAS, 342, 1280 Broderick, A., & Blandford, R. 2004, MNRAS, 349, 994
- Broderick, A. E., Gold, R., Karami, M., et al. 2020, ApJ, 897, 139
- Broderick, A. E., & Loeb, A. 2006, MNRAS, 367, 905 Bronzwaer, T., Davelaar, J., Younsi, Z., et al. 2018, A&A, 613, A2
- Chael, A., Narayan, R., & Johnson, M. D. 2019, MNRAS, 486, 2873
- Chael, A. A., Johnson, M. D., Narayan, R., et al. 2016, ApJ, 829, 11
- Chan, C.-K., Psaltis, D., & Özel, F. 2013, ApJ, 777, 13
- Chan, C.-K., Psaltis, D., Özel, F., Narayan, R., & Sadowski, A. 2015, ApJ, 799, 1
- Connors, P. A., & Stark, R. F. 1977, Natur, 269, 128
- Darwin, C. 1959, RSPSA, 249, 180
- Davelaar, J., Bronzwaer, T., Kok, D., et al. 2018, ComAC, 5, 1
- de Vries, A. 2000, CQGra, 17, 123
- Dexter, J. 2016, MNRAS, 462, 115
- Dolence, J. C., Gammie, C. F., Mościbrodzka, M., & Leung, P. K. 2009, ApJS, 184, 387
- Epperson, J. F. 2013, An Introduction to Numerical Methods and Analysis (2nd ed.; New York: Wiley)
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019a, ApJL, 875, L1
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019b, ApJL, 875, L2
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019c, ApJL, 875, L3
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019d, ApJL, 875, L4
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019e, ApJL, 875, L5
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019f, ApJL, 875, L6
- Gammie, C. F., & Leung, P. K. 2012, ApJ, 752, 123

- Gammie, C. F., McKinney, J. C., & Tóth, G. 2003, ApJ, 589, 444
- Gold, R., Broderick, A. E., Younsi, Z., et al. 2020, ApJ, 897, 148
- Gralla, S. E., Holz, D. E., & Wald, R. M. 2019, PhRvD, 100, 024018
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357
- Himwich, E., Johnson, M. D., Lupsasca, A., & Strominger, A. 2020, PhRvD, 101, 084020
- Hunter, J. D. 2007, CSE, 9, 90
- Johannsen, T., & Psaltis, D. 2010, ApJ, 718, 446
- Johnson, M. D., Lupsasca, A., Strominger, A., et al. 2020, SciA, 6, eaaz1310
- Krawczynski, H. 2012, ApJ, 754, 133
- Löhner, R. 1987, CMAME, 61, 323
- Luminet, J.-P. 1979, A&A, 75, 228
- Mościbrodzka, M., Falcke, H., & Shiokawa, H. 2016, A&A, 586, A38
- Mościbrodzka, M., & Gammie, C. F. 2017, MNRAS, 475, 43
- Narayan, R., Johnson, M. D., & Gammie, C. F. 2019, ApJL, 885, L33
- Narayan, R., Sädowski, A., Penna, R. F., & Kulkarni, A. K. 2012, MNRAS, 426, 3241
- Noble, S. C., Leung, P. K., Gammie, C. F., & Book, L. G. 2007, CQGra, 24, S259
- Ohanian, H. C. 1987, AmJPh, 55, 428
- Parkin, E. R. 2011, MNRAS Lett., 410, L28

- Pedrola, G. L. 2015, Finite Difference Approximations of Derivatives (New York: Wiley), 300
- Porth, O., Chatterjee, K., Narayan, R., et al. 2019, ApJS, 243, 46
- Psaltis, D., Özel, F., Chan, C.-K., & Marrone, D. P. 2015, ApJ, 814, 115
- Pu, H.-Y., Yun, K., Younsi, Z., & Yoon, S.-J. 2016, ApJ, 820, 105
- Ricarte, A., Prather, B. S., Wong, G. N., et al. 2020, MNRAS, 498, 5468
- Rybicki, G. B., & Lightman, A. P. 1979, Radiative Processes in Astrophysics (New York: Wiley)
- Schnittman, J. D., Krolik, J. H., & Hawley, J. F. 2006, ApJ, 651, 1031
- Shcherbakov, R. V., & Huang, L. 2011, MNRAS, 410, 1052
- Shcherbakov, R. V., Penna, R. F., & McKinney, J. C. 2012, ApJ, 755, 133 Takahashi, R. 2004, ApJ, 611, 996
- Teo, E. 2003, GReGr, 35, 1909
- Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr+ 2017, Interferometry and Synthesis in Radio Astronomy (3rd ed.; Hoboken, NJ: Wiley)
- Wise, J. H., & Abel, T. 2011, MNRAS, 414, 3458
- Wong, G. N. 2021, ApJ, 909, 217
- Younsi, Z., Porth, O., Mizuno, Y., Fromm, C. M., & Olivares, H. 2020, in Proc. IAU 342, Perseus in Sicily: From Black Hole to Cluster Outskirts, ed. K. Asada (Cambridge: Cambridge Univ. Press), 9
- Yuan, F., & Narayan, R. 2014, ARA&A, 52, 529