Principal Component Analysis as a Tool for Characterizing Black Hole Images and Variability

Lia Medeiros[1,2] , Tod R. Lauer[3] , Dimitrios Psaltis[1], and Feryal Özel[1]

[1] Steward Observatory and Department of Astronomy, University of Arizona, 933 N. Cherry Ave., Tucson, AZ 85721, USA
[2] Department of Physics, Broida Hall, University of California Santa Barbara, Santa Barbara, CA 93106, USA
[3] National Optical Astronomy Observatory, Tucson, AZ 85726, USA

Abstract

We explore the use of principal component analysis (PCA) to characterize high-fidelity simulations and interferometric observations of the millimeter emission that originates near the horizons of accreting black holes. We show mathematically that the Fourier transforms of eigenimages derived from PCA applied to an ensemble of images in the spatial domain are identical to the eigenvectors of PCA applied to the ensemble of the Fourier transforms of the images, which suggests that this approach may be applied to modeling the sparse interferometric Fourier-visibilities produced by an array such as the Event Horizon Telescope. We also show that the simulations in the spatial domain can themselves be compactly represented with a PCA-derived basis of eigenimages, which allows for detailed comparisons to be made between variable observations and time-dependent models, as well as for detection of outliers or rare events within a time series of images. Furthermore, we demonstrate that the spectrum of PCA eigenvalues is a diagnostic of the power spectrum of the structure and, hence, of the underlying physical processes in the simulated and observed images.

1. Introduction

The task of imaging and modeling the millimeter emission close to the horizon of an accreting black hole with the Event Horizon Telescope (EHT) encompasses a number of challenges. Interferometric imaging requires accurate synthesis of an image based on a sparse and incomplete set of Fourier-visibilities (see, e.g., Honma et al. 2014; Bouman et al. 2016; Chael et al. 2016; Akiyama et al. 2017). Understanding the morphological diversity of the structure of the emission and its dependence on the physical parameters of the black hole rests on the comparison of such observations to high-fidelity simulations of the accretion flow (see, e.g., Dexter et al. 2009; Mościbrodzka et al. 2009, 2017; Chan et al. 2015b; Kim et al. 2016; Gold et al. 2017). The accretion flow itself is dynamic, potentially causing strong variations in the emission morphology over the very timescales required to synthesize an image with a large baseline interferometer (see, e.g., Lu et al. 2016; Medeiros et al. 2017, 2018). In considering all of these issues, a common thread that emerges is a need to efficiently capture and characterize a complex series of images in diverse contexts. In this and forthcoming papers, we apply Principal Component Analysis (PCA) to General Relativistic magnetohydrodynamic (GRMHD) simulations and simulated EHT observations to explore its utility as an approach to addressing these challenges.

PCA is a mathematical approach to quantifying variability of an ensemble. In our case, the ensemble is a collection of images obtained from time-dependent simulation outputs of black hole accretion flows. PCA is non-parametric and does not incorporate any physical knowledge of the black hole or its accretion physics. Instead, PCA decomposes each image into a sum of orthogonal-basis eigenvectors (i.e., eigenimages), with eigenvalues that correspond to the brightness variance that each eigenimage captures. The eigenimages are then ranked by their eigenvalues, which allows minor variations to be discarded if

desired. In other words, PCA allows for a compact and effective representation of the images in the ensemble. In practice, the implicit compression can be substantial, using perhaps only a dozen eigenvectors to represent over 1000 source images (Boroson & Lauer 2010).

In this initial exploration, we show that PCA is particularly useful to help recognize and characterize the large-scale temporal variability in the morphology of the millimeter emission close to the horizon of black holes such as Sgr A* and M87. Numerous observations and studies in the past few decades have established the fact that black hole accretion flows are highly variable. X-ray observations of galactic black hole binaries reveal a variability spectrum that is characterized by red noise and also distinct high-frequency quasi-periodic components (see, e.g., van der Klis 2000; Remillard & McClintock 2006). Similarly, multi-wavelength observations of nearby AGN, including the Galactic Center black hole Sgr A*, show variability on timescales ranging from hours to months (see, e.g., Genzel et al. 2003; Do et al. 2009; Neilsen et al. 2015). This variability is not surprising. It is understood to be the result of the turbulent accretion flows and a potential manifestation of the unique black hole spacetime near its horizon (see, e.g., Rauch & Blandford 1994; Chan et al. 2015a). Consequently, it is expected that the images for Sgr A* and M87 at mm wavelengths will also be variable at dynamical timescales in the vicinities of the black hole horizons (Medeiros et al. 2017, 2018).

Black hole variability is an important consideration for the EHT. Because interferometry relies on the rotation of the Earth to obtain images, it is critical to understand how the sources may vary over timescales comparable to those observations to design and implement proper image reconstruction algorithms (see, e.g., Lu et al. 2016; Bouman et al. 2017; Johnson et al. 2017 for early attempts). It is also important to study and characterize the variability predicted by different accretion flow models in order to investigate whether the mode and amplitude of predicted

variability agrees with observations (see Kim et al. 2016). Both these issues are especially true for the primary target for the EHT, Sgr A*, for which the characteristic time of variability can be as short as a few minutes.

In this paper, we show that PCA can generate a compact orthogonal set of basis eigenimages that can accurately represent the ensemble of images generated in a suite of high-fidelity GRMHD simulations and facilitate the efficient comparison of models to observations. This basis may also be used to provide a compact rendition of the ensemble in Fourier space and, in turn, a path to efficient representation of the sparse visibility observations. Furthermore, we show that PCA allows us to recognize "outliers" in the typical source morphology and identify, both in simulations and in observations, instances of episodic physical phenomena, such as magnetic reconnection and flaring events.

In parallel to the efforts to characterize and understand its origins, the EHT has a nearly orthogonal interest in the question of black hole variability; i.e., identifying emission signatures that are, in fact, not variable. In particular, the black hole shadow, which offers excellent opportunities for testing the predictions of general relativity (Psaltis et al. 2015), is expected to be both invariant in time and nearly independent of any property of the system other than the mass of the black hole. It is, therefore, valuable to separate time-variable aspects of black hole images, such as turbulence and periodic variabilities in the accretion flow, from the constant signals arising from the black hole spacetime.

The structure of this paper is as follows. We provide a brief development of the PCA formalism in Section 2 and demonstrate that a PCA basis that is derived in the image domain also provides a basis in the Fourier (i.e., visibility) domain. We apply our formalism to a simple ensemble of images in Section 3. In Section 4, we demonstrate the ability of PCA to represent a temporal sequence of high-fidelity simulated images of an accreting black hole. In Section 5, we demonstrate the use of PCA to compactify the space of images using dimensionality reduction and to identify times of rare or unusual activity in the simulated time series. In Section 6, we compare the spectrum of PCA eigenvalues to that of Gaussian and red-noise processes and show how the PCA eigenvalues are related to the underlying power spectrum of structures in the images. We conclude and discuss future applications of our work in Section 7.

2. Principal Component Analysis

Our goal is to use PCA to determine the dominant components in a set of images of black holes. In this section, we give a brief introduction to PCA and show that it may be applied directly to interferometric observables. The majority of this derivation follows Turk & Pentland (1991), with some differences that we will explicitly outline below.

2.1. Introduction to PCA

The principle of PCA is to calculate a set of orthogonal eigenimages (or eigenvectors) from an ensemble of images. We can then utilize this basis to compactly represent all of the images in the original ensemble as a linear combination of those eigenimages.

We denote an ensemble of m images by $I_n(x, y)$, where $n = 1, \ldots, m$ and the pair of coordinates (x, y) are used to

represent the location of each of the $N \times N$ pixels on the image. For simplicity, each image can also be represented as a column vector $\boldsymbol{I}_n$ of length N^2. For our purposes, the ensemble of images will be comprised of a series of snapshots of a black hole accretion flow that are obtained from simulations or observations, although the derivations that we provide below are much more general.

As the basis of our decomposition, we choose to use the m orthogonal eigenimages $\boldsymbol{u}_k$ of the covariance matrix

$$C = \frac{1}{m} \sum_{n=1}^{m} \boldsymbol{I}_n \boldsymbol{I}_n^T$$
$$\equiv AA^T. \qquad (1)$$

In this equation, we defined the $N^2 \times m$ matrix A, such that its columns are the m images of the ensemble; i.e.,

$$A \equiv [\boldsymbol{I}_1 \, \boldsymbol{I}_2 \, \cdots \, \boldsymbol{I}_m]. \qquad (2)$$

Strictly speaking and contrary to the notation of Turk & Pentland (1991), C is not a covariance matrix because we have not subtracted the mean from each image. However, we will refer to C as the covariance matrix throughout this paper to avoid introducing unnecessary terminology.

The covariance matrix C is an $N^2 \times N^2$ matrix that measures how the variation in the brightness of each pixel across the ensemble of images is correlated to the variation in brightness of every other pixel. We can write explicitly each element of the matrix C as

$$C_{ij} = \frac{1}{m} \sum_{n=1}^{m} A_{in} A_{jn}, \qquad (3)$$

where the indices i and j correspond to the N^2 pixels $(i, j = 1, 2, \ldots, N^2)$ and the index $n = 1, 2, \ldots, m$ corresponds to the different images in the ensemble.

In principle, we can then find the eigenimages $\boldsymbol{u}_k$ of the covariance matrix C by diagonalizing it such that

$$C\boldsymbol{u}_k = \mu_k \boldsymbol{u}_k$$
$$AA^T \boldsymbol{u}_k = \mu_k \boldsymbol{u}_k. \qquad (4)$$

However, diagonalizing an $N^2 \times N^2$ matrix is computationally expensive and, in fact, not necessary. Because there are (at most) only m independent images in the ensemble, there are only m non-trivial eigenvalues and eigenvectors for this covariance matrix, which we can compute in an efficient way.[4]

We start by computing the eigenvectors and eigenvalues of the $m \times m$ matrix $L = A^T A$ such that

$$L\boldsymbol{v}_\gamma = \lambda_\gamma \boldsymbol{v}_\gamma$$
$$A^T A \boldsymbol{v}_\gamma = \lambda_\gamma \boldsymbol{v}_\gamma, \qquad (5)$$

where $\boldsymbol{v}_\gamma$ are the m eigenvectors of L, each of dimension m. It is easy to show by multiplying both sides of Equation (5) by A that the matrix L and the covariance matrix C share the same eigenvalues; i.e.,

$$AA^T A \boldsymbol{v}_\gamma = \lambda_\gamma A \boldsymbol{v}_\gamma$$
$$CA \boldsymbol{v}_\gamma = \lambda_\gamma A \boldsymbol{v}_\gamma. \qquad (6)$$

---

[4] See, e.g., Appendix A of Strang (1988) for a discussion of this property.

This equation also demonstrates that the vectors

$$\boldsymbol{u}_\gamma = A\boldsymbol{v}_\gamma, \tag{7}$$

of size N^2, are the eigenimages of the covariance matrix C with corresponding eigenvalues λ_γ.

The normalization of the eigenimages is, in principle, arbitrary. Following standard PCA convention, we normalize each eigenimages such that

$$\boldsymbol{u}_k^2 = \lambda_k \tag{8}$$

and

$$\sum_{k=1}^m \lambda_k = 1. \tag{9}$$

Because the eigenvectors are orthogonal, it also follows that $\boldsymbol{u}_k\boldsymbol{u}_{k'} = \lambda_k \delta_{kk'}$, where $\delta_{kk'}$ is the Kronecker delta. Hereafter, we will use the notation $\boldsymbol{u}_k^2 = \boldsymbol{u}_k\boldsymbol{u}_k^T$ to denote the square of the magnitude of an eigenimage. The overall sign of each eigenimage is arbitrary and, in principle, the mean pixel value of an eigenimage may be negative and this caries no physical meaning. Here, for simplicity, we enforce the mean of each eigenimage to be positive.

Having obtained the eigenimages of our ensemble, we can then express any of its images as the linear combination

$$\boldsymbol{I}_n = \sum_{k=1}^m a_{nk}\boldsymbol{u}_k, \tag{10}$$

where

$$a_{nk} \equiv \frac{\boldsymbol{u}_k^T \boldsymbol{I}_n}{(\boldsymbol{u}_k^2)} \tag{11}$$

are the amplitudes of the projections of the images on the eigenimage basis. The column vectors $\boldsymbol{a}_k$ can be written as

$$\boldsymbol{a}_k = \frac{A^T \boldsymbol{u}_k}{(\boldsymbol{u}_k^2)}, \tag{12}$$

and are equivalent to $\boldsymbol{v}_k$, because

$$\boldsymbol{a}_k = \frac{A^T A \boldsymbol{v}_k}{(\boldsymbol{u}_k^2)} = \frac{L\boldsymbol{v}_k}{(\lambda_k)} = \boldsymbol{v}_k. \tag{13}$$

Therefore, Equation (12) provides a simpler way of calculating the column vectors $\boldsymbol{a}_k$ compared to Equation (11). The square of the magnitude of each image is equal to

$$
\begin{aligned}
\boldsymbol{I}_n^2 &= \boldsymbol{I}_n^T \boldsymbol{I}_n \\
&= \left(\sum_{k=1}^m a_{nk}\boldsymbol{u}_k^T\right)\left(\sum_{k'=1}^m a_{nk'}\boldsymbol{u}_{k'}\right) \\
&= \sum_{k=1}^m a_{nk}^2 \lambda_k.
\end{aligned}
\tag{14}
$$

In our following discussion of outlier detection, we will also use the notion of the fractional contribution of eigenimage k to each snapshot n (cf. Equation (11)), which we define as

$$
\begin{aligned}
a'_{nk} &\equiv \frac{\boldsymbol{u}_k^T \boldsymbol{I}_n}{\sqrt{\boldsymbol{I}_n^2 \boldsymbol{u}_k^2}} \\
&= \frac{a_{nk}\sqrt{\lambda_k}}{\sqrt{\sum_{k=1}^m a_{nk}^2 \lambda_k}}.
\end{aligned}
\tag{15}
$$

In principle, when we use this basis set to reconstruct each original image in the ensemble, we need all m eigenimages. However, depending on the level of fidelity required and on the uniformity of the images in the set, the PCA decomposition makes it possible for us to reduce the dimensionality of the problem by only using the first few eigenimages to reconstruct an approximation of each of the original images. This approach becomes especially useful when only a few eigenimages are significant and the rest are small. Naturally, the number of eigenimages used to construct the model depends on the particular application and does require judgment. For example, in observational data with real noise, the eigenimage expansion can be terminated when the model begins to overfit the noise. Boroson & Lauer (2010) presented a detailed analysis of the optimal method to terminate a PCA expansion given knowledge of the typical S/N of the ensemble images. In characterizing images from simulations that do not include observational noise, the judgment of when to terminate the expansion is one of how much fidelity is required to capture the critical morphology of the image.

Lastly, we emphasize an obvious but important application of PCA. Given a set of eigenimages, the basis can also be used to represent and analyze images that are similar to those in the set used to define the eigenimages but that are not actually in the set itself. In the present context, this means that a basis constructed from a set of simulated images of an accreting black hole should be able to represent observations of the black hole if the simulations are sufficiently realistic.

2.2. PCA in the Fourier Domain

Even though we presented the PCA formalism using a set of images, the data that we ultimately aim to work with are the complex Fourier components of the image; i.e., visibility amplitudes and phases. This is because the EHT is an interferometric array and directly measures the latter quantities. Ideally, we would like to devise a method for characterizing image variability that can be used in both image space and Fourier space and that allows us to move freely between the two.

From a purely mathematical point of view, the image and Fourier domains are highly symmetric and it is straightforward to represent an operation in one domain with a complementary operation in the other domain. In practice, however, the two domains present strongly asymmetric viewpoints. The spatial distribution of radio emission close to the horizon of an accreting black hole is readily formulated and visualized with high-fidelity simulations in the image domain. The observations are obtained in the visibility domain, however, with relatively sparse coverage. Confronting the simulations with the observations requires a sophisticated synthesis of the visibility data into an interpretable form. One path is to use general purpose image reconstruction techniques, although they may suffer from less than optimal use of the expected morphology of the observations. In contrast, our approach will be to develop a basis directly in the visibility domain that

3

encodes the expected behavior of the source as informed by simulations. We thus need to understand how the PCA basis of the simulations will relate to their visibilities. In this section, we show that the visibilities of the principal components of the simulation images are in fact the same as the principal components of the visibilities of these images.

We define the 2D discrete Fourier transform of an image as

$$\tilde{I}_\alpha = \sum_{i=1}^{N^2} F_{\alpha i} I_i, \qquad \alpha = 1, ..., N^2. \tag{16}$$

Here, to account for the folding of the images into one-dimensional vectors, we have written the discrete Fourier operator in the compact form

$$F_{\alpha i} = e^{\frac{-2\pi i}{N}[j\beta + k\delta]}, \tag{17}$$

where the indices j, k, β, and δ in the right-hand side of this relation can be evaluated from the indices α and i via the relations

$$k = [(i - 1) \bmod N] + 1, \tag{18}$$

$$j = \frac{i - k}{N} + 1 \tag{19}$$

and

$$\delta = [(\alpha - 1) \bmod N] + 1, \tag{20}$$

$$\beta = \frac{\alpha - \delta}{N} + 1. \tag{21}$$

Note that in Equation (17) we used the symbol i for the imaginary number to distinguish it from index i.

The Fourier transform of matrix A is simply

$$\tilde{A}_{\alpha n} = \sum_{i=1}^{N^2} F_{\alpha i} A_{in} \tag{22}$$

and we define the $m \times m$ matrix $L' \equiv \tilde{A}^T \tilde{A}$ as

$$L' \equiv \tilde{A}^T \tilde{A} = \begin{bmatrix} \tilde{I}_1 \tilde{I}_1 & \tilde{I}_1 \tilde{I}_2 & \cdots & \tilde{I}_1 \tilde{I}_m \\ \tilde{I}_2 \tilde{I}_1 & \tilde{I}_2 \tilde{I}_2 & \cdots & \tilde{I}_2 \tilde{I}_m \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{I}_m \tilde{I}_1 & \tilde{I}_m \tilde{I}_2 & \cdots & \tilde{I}_m \tilde{I}_m \end{bmatrix}. \tag{23}$$

Our goal here is to show that this matrix is the same as L; i.e., that $L' = L$.

We write the vector product that appears in each element of matrix L' as

$$\tilde{I}_f \tilde{I}_g = \sum_{\alpha=1}^{N^2} \tilde{I}_{\alpha f} \tilde{I}_{\alpha g}^*$$

$$= \sum_{\alpha=1}^{N^2} \left(\sum_{i=1}^{N^2} F_{\alpha i} I_{if} \right) \left(\sum_{i'=1}^{N^2} F_{\alpha i'}^* I_{i'g} \right). \tag{24}$$

Because I_{if} and $I_{i'g}$ do not depend on α, we can rearrange the above equation as follows

$$\tilde{I}_f \tilde{I}_g = \sum_{i=1}^{N^2} I_{if} \sum_{i'=1}^{N^2} I_{i'g} \left(\sum_{\alpha=1}^{N^2} F_{\alpha i} F_{\alpha i'}^* \right). \tag{25}$$

The term in parenthesis above is the 2D Fourier transform of a constant and is equal to $\delta_{ii'}$ such that

$$\tilde{I}_f \tilde{I}_g = \sum_{i=1}^{N^2} I_{if} \sum_{i'=1}^{N^2} I_{i'g} \delta_{ii'}$$

$$= \sum_{i=1}^{N^2} I_{if} I_{ig} = I_f I_g. \tag{26}$$

Therefore, each element of L is equal to that of L' so their eigenvectors and eigenvalues must also be equal ($L' v_\gamma = \lambda_\gamma v_\gamma$).

We now define the covariance matrix for the visibilities in analogy to that of the images as

$$C' = \tilde{A} \tilde{A}^{*T}. \tag{27}$$

As before, we can find the eigenvectors of C' by diagonalizing L' (see Equation (6)), which demonstrates equivalently that $\tilde{A} v_\gamma$ are eigenvectors of C' with eigenvalues λ_γ. To complete our proof, we must show that the eigenvectors $\tilde{A} v_\gamma$ are equal to the Fourier transform of the eigenvectors $A v_\gamma$. In other words, we must show that the principal components of the set of visibility maps is equal to the visibilities of the principal components of the set of images. This can be seen by evaluating each component of the eigenvector $\tilde{A} v_\gamma$ as

$$(\tilde{A} v_\gamma)_\alpha = \sum_{l=1}^{m} \tilde{A}_{\alpha l} (v_\gamma)_l = \sum_{l=1}^{m} \left(\sum_{i=1}^{N^2} F_{\alpha i} A_{il} \right) (v_\gamma)_l$$

$$= \sum_{i=1}^{N^2} F_{\alpha i} \left(\sum_{l=1}^{m} A_{il} (v_\gamma)_l \right) = (A \tilde{v}_\gamma)_\alpha. \tag{28}$$

Therefore, the visibilities of the principal components are indeed equal to the principal components of the visibilities. We have thus shown that the PCA basis can be developed in one domain, such as image space, but is readily applied in the complementary Fourier (or visibility) domain. This will allow us to use PCA to compare and possibly fit EHT data to simulations directly in the visibility domain. The maximally compact basis provided by the PCA approach may be well-suited to address the sparse coverage of the visibilities.

3. An Example of PCA

To elucidate how PCA works in practice, we present, in this section, a simple example that is easy to calculate and understand. We consider a Gaussian spot moving along a circular path and simulate 1080 snapshot images as the spot completes an integer number (3) of orbits. Figure 1 shows a few example snapshots from this model. We calculate the principal components of this image set using the PCA formalism and show in Figure 2 the first few principal components. We also show in Figure 3 the spectrum of eigenvalues that we obtain for this model.

The first principal component, which has the largest flux variance ($\sim$12.7%), amounts to the average image of the various snapshots (modulo a normalization constant); i.e., it represents a ring surrounding the circular path with a width comparable to the width of the Gaussian spot. This is not a general property of a PCA decomposition but is exact in the particular example discussed here and is approximately correct
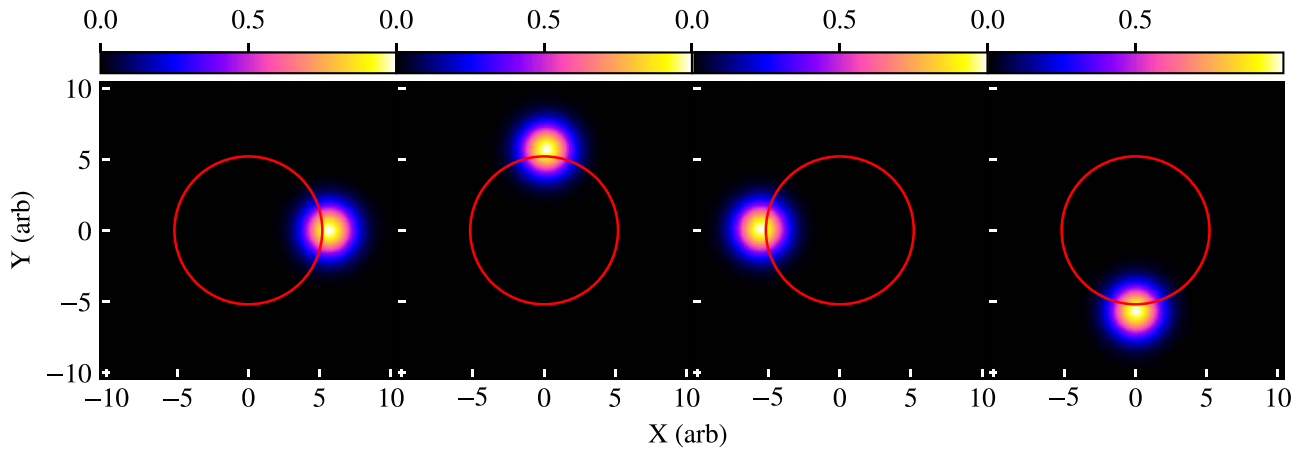
Figure 1. Example snapshots from a simple model of a Gaussian spot moving on a circular path. Here the red circle indicates the approximate trajectory of the center of the Gaussian spot. The linear scale of the image is arbitrary. We present PCA analysis of realistic GRMHD simulations later in the paper, but this simple example is useful for understanding how PCA decomposition of the simulations work.
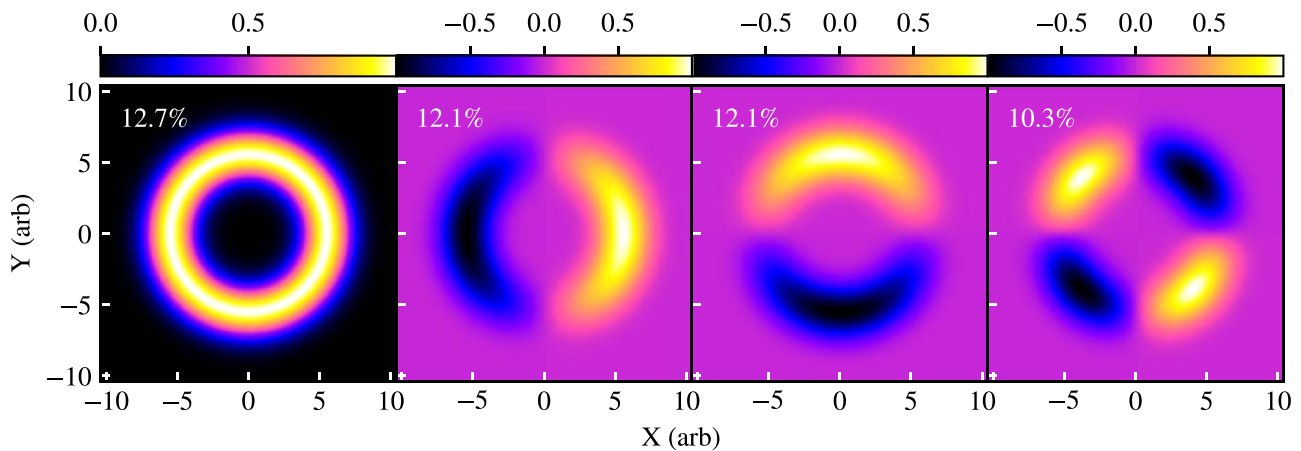

Figure 2. First four components of the PCA decomposition of the Gaussian spot moving on the circular path shown in Figure 1. The eigenvalues which correspond to these four components are shown in the top left-hand of each panel, respectively. Note that in this figure, and in all figures of principal components in the rest of the paper, each component has been normalized independently so that the fluxes cannot be compared between different components.

in the PCA decomposition of the black hole images that we will discuss in the next section. In the example of the orbiting Gaussian spot, all terms in each row of the $m \times m$ matrix L (which is also a covariance matrix) also appear in each other row of the same matrix but are displaced at different columns. This is true because the product of any two images in the ensemble depends only on the relative positions of the Gaussian spots in the two images. In other words, the sum of the elements of each row of matrix L, i.e., $\sum_{n=1}^{m} L_{in}$ is constant. One of the eigenvectors of a matrix with elements that obey this property is a vector that has all elements equal to one (or actually any constant, depending on how the eigenvector is normalized), i.e., $v_1 = [1\ 1\ 1\ ...\ 1]$. The eigenvector of the covariance matrix C that corresponds to this eigenvector of the matrix L is then (see Equation (7))

$$u_1 = Av_1 = \sum_{n=1}^{m} I_m, \tag{29}$$

which (modulo a normalization constant) is nothing but the average image of the ensemble.

Most eigenimages, other than the lowest-order one, have pixels with significantly positive and negative brightness, as one would expect from an eigenvector decomposition.

However, because all images are positive definite and the lowest-order eigenimage is the average of the ensemble of images, it follows that the lowest-order eigenimage is also positive definite. Moreover, because of the symmetry of the ensemble of images, components 2 and 3 of the spectrum of eigenimages (in Figure 2) differ only by a rotation. The eigenvalue connected to each eigenimage is related to its variance (see Equation (8)) and, hence, these two components correspond to the same eigenvalue. This behavior persists with higher components such that components come in pairs with similar eigenvalues. This creates the step pattern that is present in Figure 3.

Here and in all simulations of black hole images discussed in the sections below, the typical values of the amplitudes a_{nk} are very similar between different eigenimages; i.e., the typical values and distribution of the amplitudes a_{nk} depend weakly on k. Because of this and the fact that the eigenimages are normalized according to their eigenvalues (see Equation (8)), the spectrum of eigenvalues, such as the one shown in Figure 3, matches very closely the spectrum of the relative contributions of each eigenimage to the reconstruction of any of the m images in the ensemble. As a result, we can use the spectrum of eigenvalues as a proxy to investigate the relative contribution of each eigenimage to the reconstruction of a typical image in

the ensemble (see below for outlier detection). This is the reason why we normalize all eigenvalues so that they sum to unity (see Equation (9)) and we often quote them as percentages.

In the example of the circulating Gaussian spot that we discuss in this section, the eigenvalues drop dramatically after the first few, indicating that only a few components would be sufficient to reconstruct the original images. Specifically, under the assumption that the parameters a_{nk} are independent of k, we conclude that the first 10 (out of 1080) components account for $\approx 88\%$ of the structures present in the ensemble of images, while ~ 25 components account for nearly all of it. For this particular example, it is straightforward to understand why it takes only a small fraction of the eigenimages to reconstruct most of the structures seen in each of the images in the ensemble by estimating the number of substantially different images that are present in the ensemble. For the parameters used in this model, the FWHM of the Gaussian spot subtends $\sim 13°$, as viewed from the center of the circular path. Therefore, the full circular trajectory can be decomposed into 28 distinct Gaussian spots that are (mostly) not overlapping. In other words, there are only 28 "resolution" elements in the circular trajectory and, therefore, the contribution of all but the first ~ 28 eigenimages can only be negligible.

4. Principal Components Analysis of Simulated Black Hole Images

We now apply PCA to a set of simulated black hole images at the 1.3 mm wavelength of observations for the EHT. We focus on three of the five best-fit models (Chan et al. 2015a; Medeiros et al. 2018, 2017) that were calibrated to reproduce the broadband spectrum of Sgr A* and the size of the 1.3 mm emission region inferred by early EHT observations (Doeleman et al. 2008). These three models span the range of image morphologies and structural variability that we encountered in all of our simulations. Specifically, Model B has a 1.3 mm image that is dominated by the accretion disk region and resembles a crescent shape, Model C has a 1.3 mm image that is dominated by the base of the jet funnel, while Model D has a 1.3 mm image that is a combination of both the base of the funnel and the disk (here we use Chan et al.'s 2015a nomenclature to label the models).

The simulations were generated by performing time-dependent general relativistic magnetohydrodynamic (GRMHD) simulations using the 3D HARM code (Gammie et al. 2003; Narayan et al. 2012; Sądowski et al. 2013) and by carrying out radiative transfer and geodesic ray tracing calculations on the simulation outputs using GRay (Chan et al. 2013). Because GRay is a massively parallel GPU based code, we were able to generate images as a function of time with high spatial and temporal resolution, for a large number of simulations, while varying the black hole spin, the geometry of the magnetic field, and the plasma model (see Chan et al. 2015b for a detailed description).

In Figure 4, we show four example snapshots from Model B highlighting the structural variation in the emission region that is prominent in this model. Hereafter, when displaying images of black holes, we will measure all lengths in units of the gravitational radius GM/c^2, where M is the mass of the black hole and G and c are the gravitational constant and speed of light, respectively. The radius of the black hole shadow is approximately equal to five gravitational radii while the center of the shadow is displaced with respect to the center of gravity,
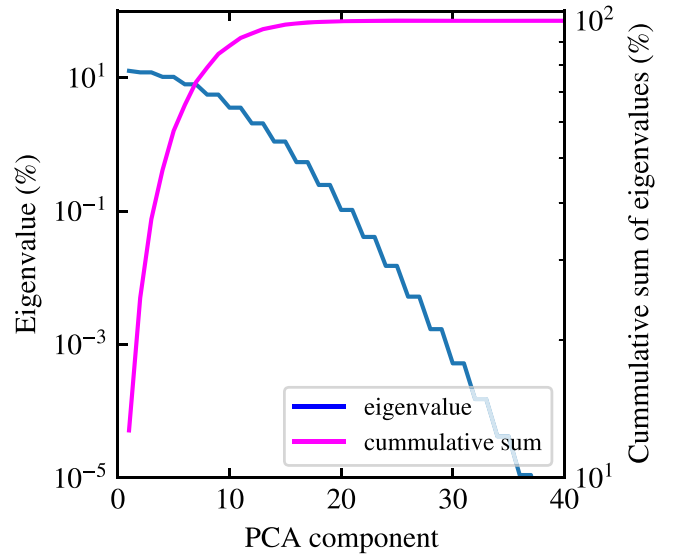

Figure 3. Blue curve: the spectrum of eigenvalues for the PCA decomposition of the Gaussian spot shown in Figure 1; the eigenvalues have been normalized to sum up to 100%. The step-like features in this spectrum are present because the high degree of symmetry in this model causes the principal components to come in pairs with very similar eigenvalues (see the second and third panels of Figure 2). Magenta curve: the cumulative sum of the eigenvalues. Note that only the first ~ 40 of the 1080 components are shown and that the first 10 components contain 88% of the structural information.

depending on the spin of the black hole (see, e.g., Chan et al. 2013).

We perform PCA on the three simulations described above following the procedure outlined in Section 2. Each image set consists of 1024 images corresponding to the number of snapshots obtained from the accretion flow simulations that span ≈ 60 hr. In Figure 5, we show the first four eigenimages and their respective eigenvalues for the PCA decomposition of the three models. For all models, the first eigenimage (left) is similar to the time average of the ensemble of images (see Figure 1 in Medeiros et al. 2017 for a comparison with the time averaged images of these simulations). This is true because all images have a dominant structure (i.e., a crescent or the footprints of the funnel), on top of which the variability of the accretion flow introduces sub-dominant perturbations. Consequently, the correlations between the various snapshots are very similar to each other and the arguments given in Section 3 for the dominant eigenimage also apply here, although only approximately.

Although PCA is a purely mathematical tool and is agnostic about the physics of the system, some of the components do appear to have identified important physical features. For example, the second component in Model B shown in the top row of Figure 5 appears to have identified a region of the Doppler boosted accretion disk (center of the crescent shape) that is very close to the black hole shadow and is highly variable in the simulations (see, e.g., the second and third panels of Figure 4). The third and fourth components of Model B appear to be tracing the Doppler boosted walls of the funnel region. This also matches the behavior that can be seen directly in the simulation, where the relative brightness of the wall of the funnel region is highly variable (see, e.g., the fourth panel of Figure 4). Note that, due to the 60° inclination of the observer relative to the spin axis of the black hole in these simulations, the base of the funnel appears to come from within
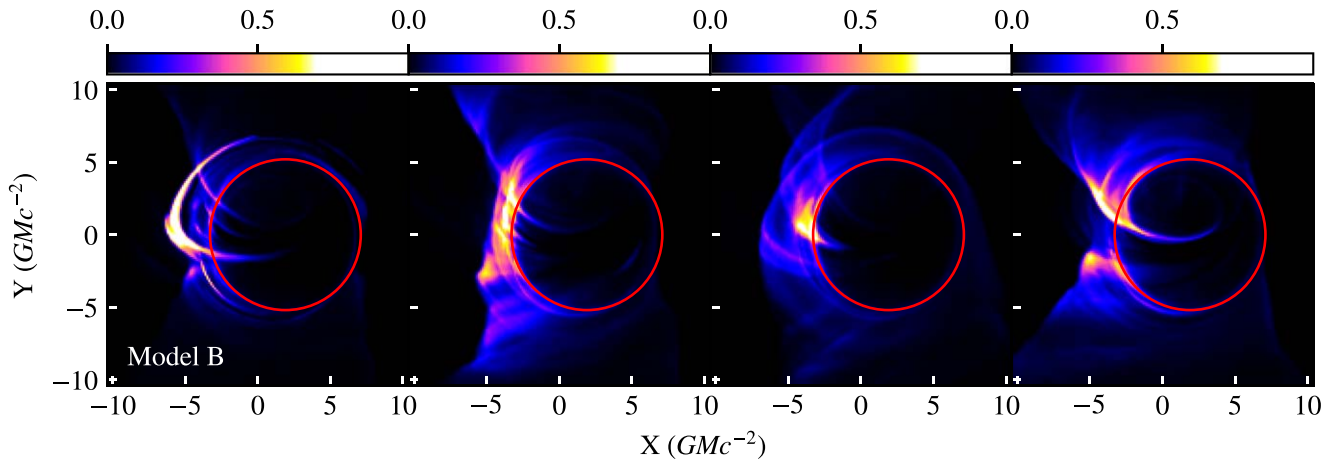
Figure 4. Four example snapshots of the ensemble of black hole images computed using Model B at a wavelength of 1.3 mm (see text for description of the model and of the simulations). None of these snapshots correspond to an instant with a significant flare in the flux of the black hole. For this and the following figures, the peak flux in each panel has been normalized to unity, so changes in overall flux have not been preserved. Although the original images span 32 GM/c^2 on each side and the full size images are used in all of the calculations, we choose to show only the innermost ~20GM/c^2 in the figures throughout the paper so that the black hole shadows are easy to distinguish. The red circles in the figures correspond to the expected size and location of the black hole shadow for each particular model. The location of the circle relative to the center of the image depends on the black hole spin and is not necessarily centered on the location of the black hole itself.

the black hole shadow but is actually positioned between the observer and the black hole.

In Model C (middle row), the second component highlights the edge of the black hole shadow while the other two components show various ways in which the structure of the emission at the base of the funnel varies. The PCA decomposition for Model D (bottom row) shows some features of Model B—i.e., that a crescent shape is present—and also some features of Model C—i.e., that the base of the funnel is an important variable feature in the image.

In Figure 6, we plot the spectra of eigenvalues of the PCA decomposition of the images from the three models. Unlike the example of the Gaussian spot that was discussed in Section 3, the eigenvalue of the first principal component in all three models overwhelms that of the remaining components. For example, the eigenvalue of the first component of Model D corresponds to ~89%, whereas the eigenvalue of the second component drops to only ~2%. In other words, under the assumptions discussed in Section 3, only the first two components (out of 1024) are required to account for 90% of the structures in the images from Model D and only the first three components are required to reach the same level for Model C. This result indicates that PCA can be extremely useful in reducing the dimensionality of the images that arise in these GRMHD simulations and that only the first few components are needed to preserve the majority of the image structure.

Model B differs somewhat from the other two models in this regard. The eigenvalue that corresponds to the first component is equal to ~66%; i.e., it is ~20% less than in the other two models. Correspondingly, as many as 33 components are required to account for 90% of the structure seen in the images for model B, showing that this model contains significantly more structural variability than the other models. This is in agreement with the findings reported in Chan et al. (2015a) and Medeiros et al. (2018, 2017), where the higher level of flux variability and flaring behavior was attributed to structural changes rather than to simple brightness fluctuations. Nevertheless, even for such a simulation that shows more significant structural variability, the required number of components (33)

is significantly smaller than the total number of images, making PCA useful for dimensionality reduction. We explore this result further in the following section.

It is intriguing that despite the differences in the relative importance of the first ~10 eigenvalues, the eigenvalue spectrum declines with the same slope for the higher components in all models. This suggests a common origin for the slope of the eigenvalue spectrum, which we will explore in detail in Section 5.

Finally, we also apply the PCA analysis directly on the complex visibilities of our image set, which are the components of the 2D Fourier transform of each image. As we showed in Section 2.2, we can either calculate the complex visibility maps for each image in our ensemble and then perform PCA or we can directly calculate the complex visibility maps for each PCA component of the ensemble of images; the results will be identical. Given that the images correspond to vectors with real elements whereas the visibilities correspond to vectors with complex elements, we follow the second procedure, which is easier to implement. In Figure 7, we show the first four PCA components of the visibility amplitudes and visibility phases of Model B. As expected, the structures of the visibility amplitudes and phases changes significantly between these four components. In fact, the visibility amplitudes of higher components have more power at longer baselines, which is a direct consequence of the fact that they contain smaller scale structures.

5. Dimensionality Reduction and Outlier Identification

As we discussed in the previous section, only the first few PCA components are required to account for the majority of the structure seen in the images from each GRMHD simulation. Components with smaller eigenvalues contribute less to the brightness of each pixel in the image (see discussion at the end of Section 3) and account primarily for small-scale structures (see Figure 7). As alluded to in the introduction, this conclusion (which is often referred to as dimensionality reduction) also offers the possibility of using a small number of measurements, such as those possible with the sparse coverage of the EHT
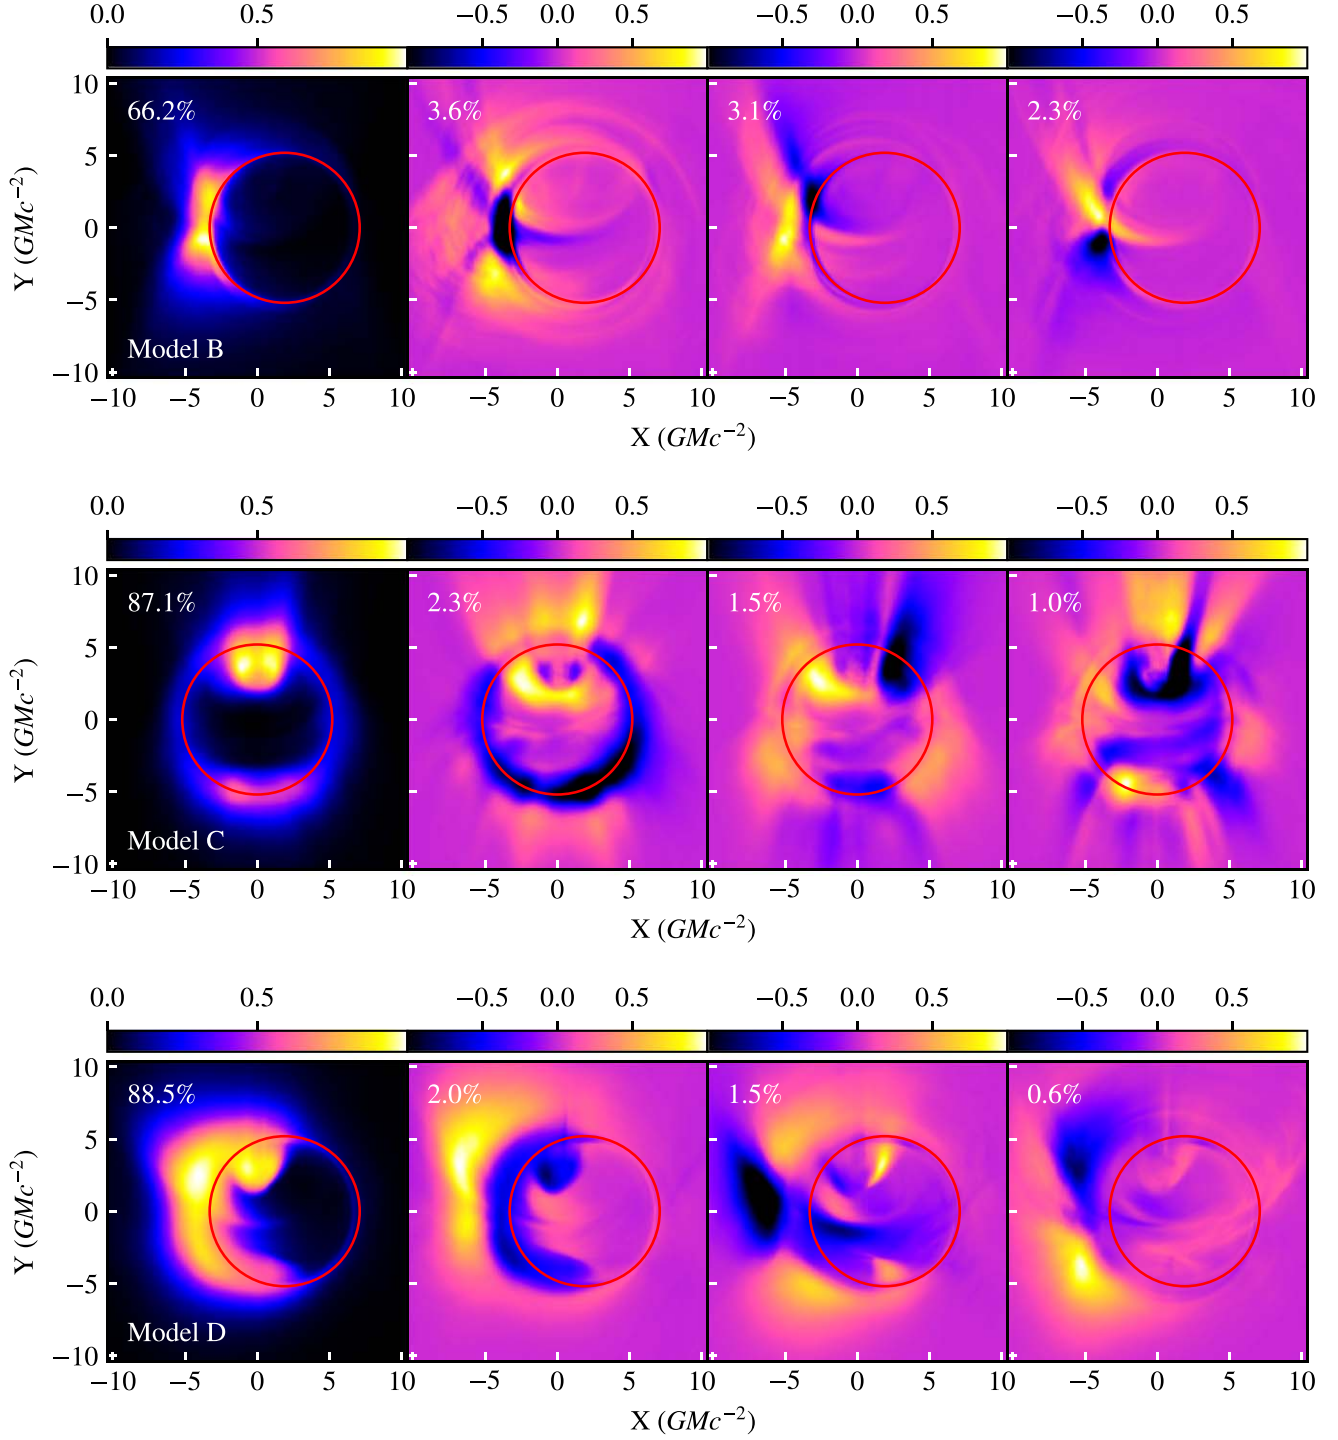
Figure 5. First four principal components and their corresponding eigenvalues for the three GRMHD simulations (models B, C, and D) described in the text. The principal component for each simulation (leftmost panel) is approximately equal to the average image from the simulation (see Figure 1 in Medeiros et al. 2017).

array, to reconstruct the persistent image of a black hole and, therefore, extract the information that is most relevant for detecting its shadow. To explore the idea of using the first few components to describe the persistent structure from the variable flow, we calculate and compare reconstructions of images from the simulations using the first few components of the PCA decomposition to the original images.

Figure 8 shows an example snapshot from the Model B simulation compared to its reconstruction using the first 10, 40,

and 100 out of the 1024 PCA components. Although the reconstructions with only a small number of components do not reproduce the finer details of the images, they do capture their overall structure. The fidelity of reconstruction naturally increases as more components are added. The number of components we may choose to keep in a particular reconstruction and, hence, the degree of dimensionality reduction will naturally depend on the goal of the reconstruction. Nevertheless, even at a qualitative level, this figure suggests that
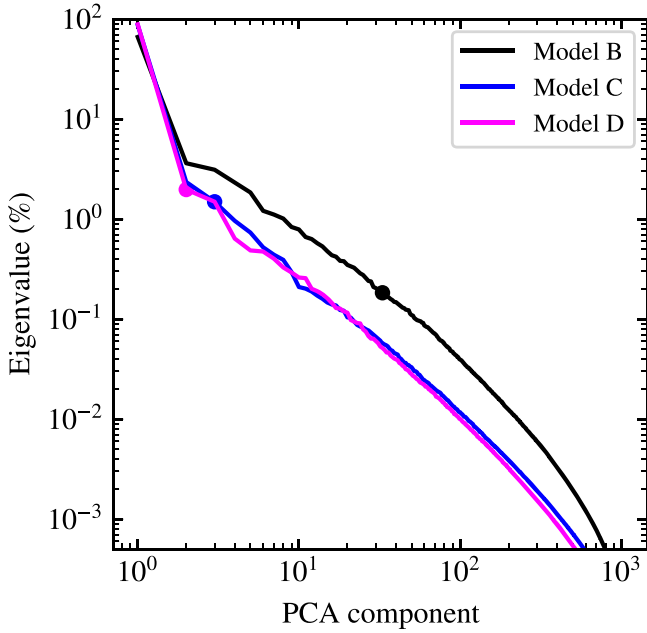
Figure 6. The eigenvalue spectra of the PCA decomposition of the images from the three GRMHD simulations that we consider. The filled circles along each spectrum indicate the number of PCA components that are required to account for 90% of the image structures. The rapid decay of the eigenvalue spectrum indicates that PCA can be used to significantly reduce the dimensionality of the ensemble of images that arise in these simulations.

dimensionality reduction by factors of 10–100 may be achievable in characterizing black hole images with PCA.

The snapshot in Figure 8 is typical and, therefore, can easily be reconstructed using only the first few PCA components. However, there may be snapshots within a given simulation that are much harder to reconstruct because the structure of the image is unusual compared to the rest of the ensemble. For the purposes of this work, we will define an outlier as an image that cannot be easily reconstructed by the first few (or a "typical" number of) eigenimages. As we will show below, PCA allows us to devise an algorithmic approach for outlier detection. When we apply PCA to numerical simulations, this outlier detection will enable us to efficiently identify instances where rare and episodic events occurred in the simulation, such as a flare in the emission properties of the accretion flow. When we apply PCA to observational data, detecting outliers will allow us to identify similar episodic events that may be caused by physical phenomena or data corruption.

There are many ways of using PCA to identify outliers in a set of images. A common method measures the Euclidean distance of each image in the hyperspace spanned by the set of eigenimages (often related to the Mahalanobis distance, Mahalanobis 1936). In implementations of outlier detection based on the Mahalanobis distance, the ensemble of images (or other data) is often standardized so that the distribution of pixel brightness within each image has been mean centered and scaled by its standard deviation. Because we have chosen not to standardize our dataset, applying the Euclidean distance method directly to our PCA implementation would identify as outliers any images that are simply brighter than the average image but without necessarily any substantial structural difference. In the context of using PCA to describe simulations of accreting black holes, we can easily identify such bright events by simply looking at large excursions of the total flux

from the mean value. Instead, our goal is to identify as outliers those snapshots with structures in the images that are substantially different from those of the typical snapshots. Consequently, we define a Euclidean distance using the fractional contribution of each eigenimage to the reconstruction of an image in the ensemble (see Equation (15)).

We will consider a given snapshot (I_n) as typical, if it can be adequately reconstructed by the first l eigenimages (l can be chosen based on the particular distribution and application). To quantify the degree to which a snapshot is atypical, we define the quantity

$$R_{nl} \equiv \frac{1}{\sum_{k=l+1}^{m} w_k} \sum_{k=l+1}^{m} \left(\frac{a'_{nk} - \langle a'_k \rangle}{\sigma^2_{a_k}} \right)^2 w_k \quad (30)$$

that measures the weighted squared Euclidean distance from the mean of the distribution of a'_{nk}, scaled by the standard deviation of the distribution, and projected onto the subspace of eigenimages that is not being used in the reconstruction. Here, the mean of the distribution

$$\langle a'_k \rangle = \frac{1}{m} \sum_{n=1}^{m} a'_{nk} \quad (31)$$

provides a measure of the average coefficient of an eigenimage needed to reconstruct the snapshots in a given set of images,

$$\sigma^2_{a'_k} = \frac{1}{m} \sum_{n=1}^{m} (a'_{nk} - \langle a'_{nk} \rangle)^2, \quad (32)$$

shows the spread in that distribution, and w_k is an appropriate weight function. We set the weight function to $w_k = \lambda_k$ because, in our implementation, the typical contribution of each eigenimage to any reconstruction is proportional to $\sqrt{\lambda_k}$ and we want to give more weight to the most dominant eigenimages when identifying outliers.

In Figure 9, we show R_{nl} for $l = 5$, 10, and 20 for all snapshots in Model B as well as the normalized light curve. A number of time instances can be easily identified as atypical— i.e., with $R_{nl} \gg 1$—but these instances do not necessarily correlate with large brightness excursions. To examine this point further, we show in Figure 10 three original snapshots as well as the reconstructions using the first 5, 10, and 20 principal components. The top row (denoted as "row 1" in Figure 9) shows an example of a time instance that is not identified as an outlier but corresponds to the largest flux excursion in this simulation. Clearly, this snapshot can be easily reconstructed by the first 10 eigenimages and has a low R_{nl} value for all three values of l. This snapshot, despite being substantially brighter than the others, does not correspond to a significant structural change in the image. The second row ("row 2") shows a time instance that is identified as an outlier but does not correspond to a significant flux excursion. The morphology of the image is quite unusual when compared to the rest of the simulation and a reconstruction with 20 eigenimages fails to capture the general structure of the image. The third row ("row 3") shows a time instance that is both identified as an outlier and shows a significant flux excursion. The reconstruction of this snapshot with 20 eigenimages is also inadequate.

These results demonstrate that, in our simulations, flux excursions and unusual image morphologies are not necessarily
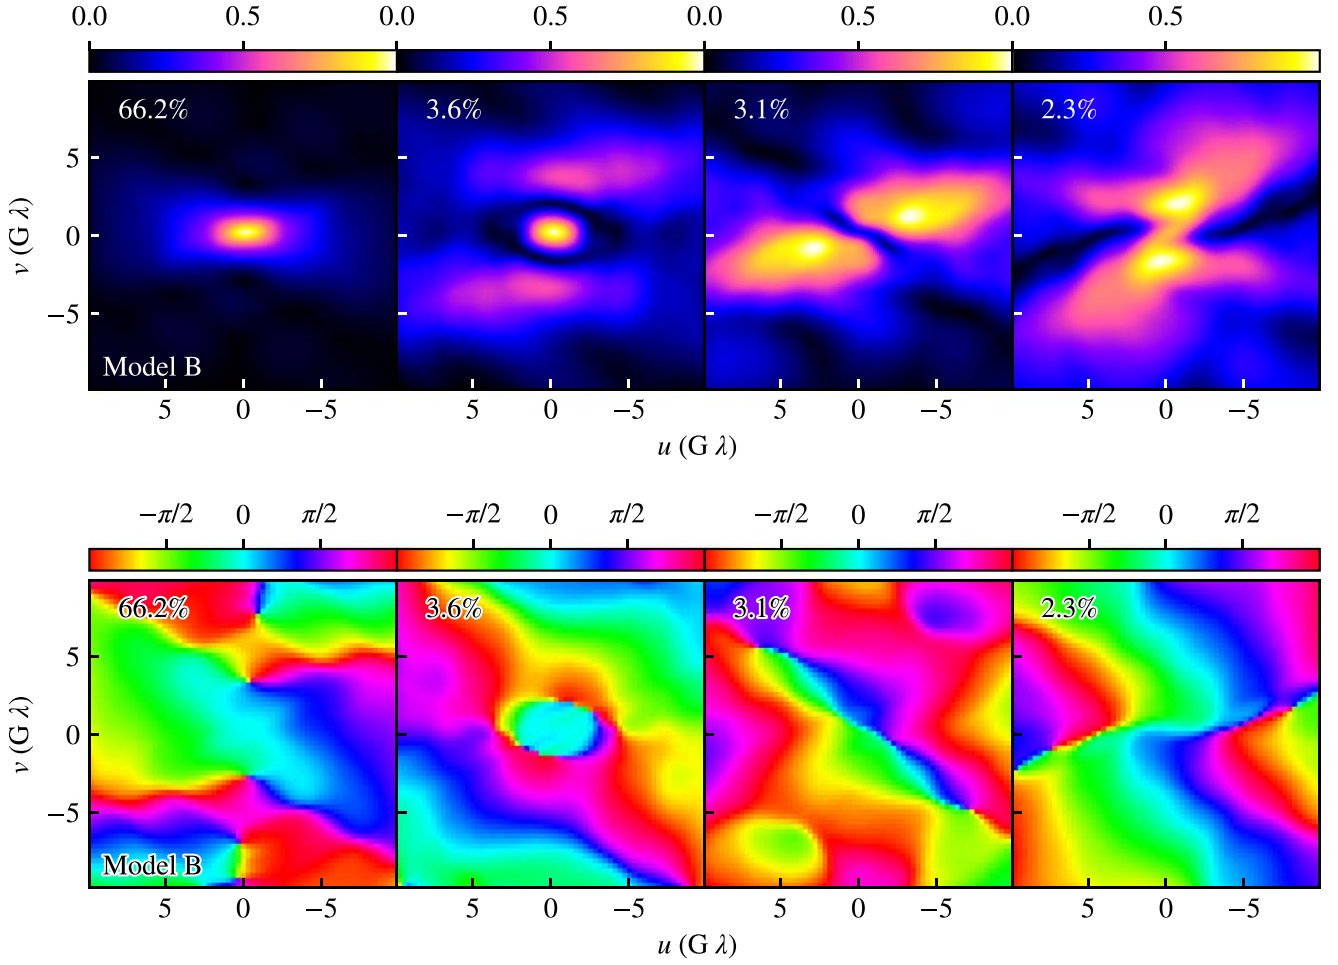
Figure 7. Visibility amplitudes (top row) and visibility phases (bottom row) of the first four components of the PCA decomposition of Model B (cf. the top row of Figure 5). Higher components contribute significantly at increasingly longer baselines.
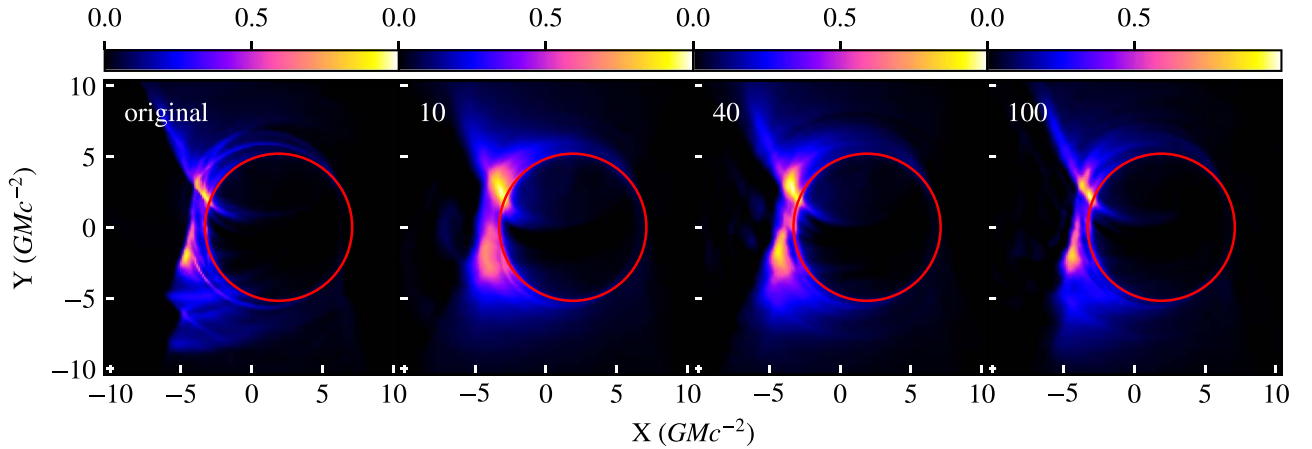

Figure 8. Left panel shows a typical snapshot from Model B. The three right panels show the same snapshot from Model B but reconstructed using only the first 10, 40, and 100 components from the PCA decomposition. The reconstruction using only the first 10 components smooths over the fine scale structure but faithfully reproduces the overall brightness distribution of the full image.

coincident but the two can be disentangled with the use of the quantity R_{nl} that we have introduced here.

6. Understanding the Eigenvalue Spectrum of PCA

In this section, we turn our attention to understanding the behavior of the spectrum of eigenvalues of the PCA decomposition of the GRMHD simulations. Specifically, we focus on the higher-order components, which have small eigenvalues and contribute primarily to the small-scale, variable structures that are seen in the images. We aim to understand the origin of their eigenvalue spectra, which have the intriguing property of being power laws with very similar slopes in all simulations. This allows us to explore whether the spectra of eigenvalues are related to the underlying properties
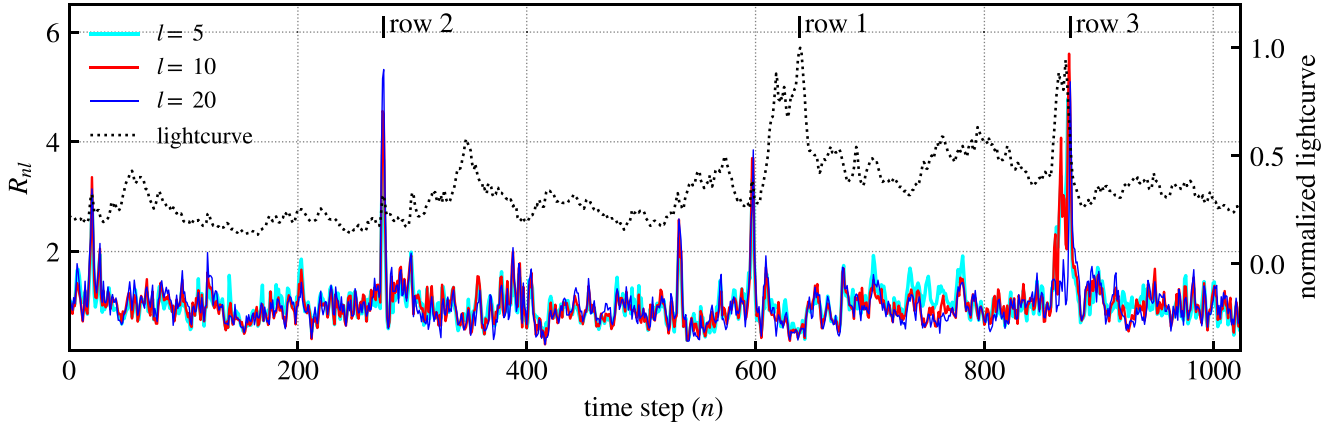
Figure 9. Dotted black line shows the normalized light curve for Model B. The cyan, red, and blue curves correspond to the quantity R_{nl} for $l = 5$, 10, and 20, respectively. The peaks in the cyan, red, and blue curves indicate time instances that cannot be adequately reconstructed by only the first few components and are, hence, identified as outliers. The three time instances identified as "row 1," "row 2," and "row 3" correspond to the images that are shown in Figure 10.
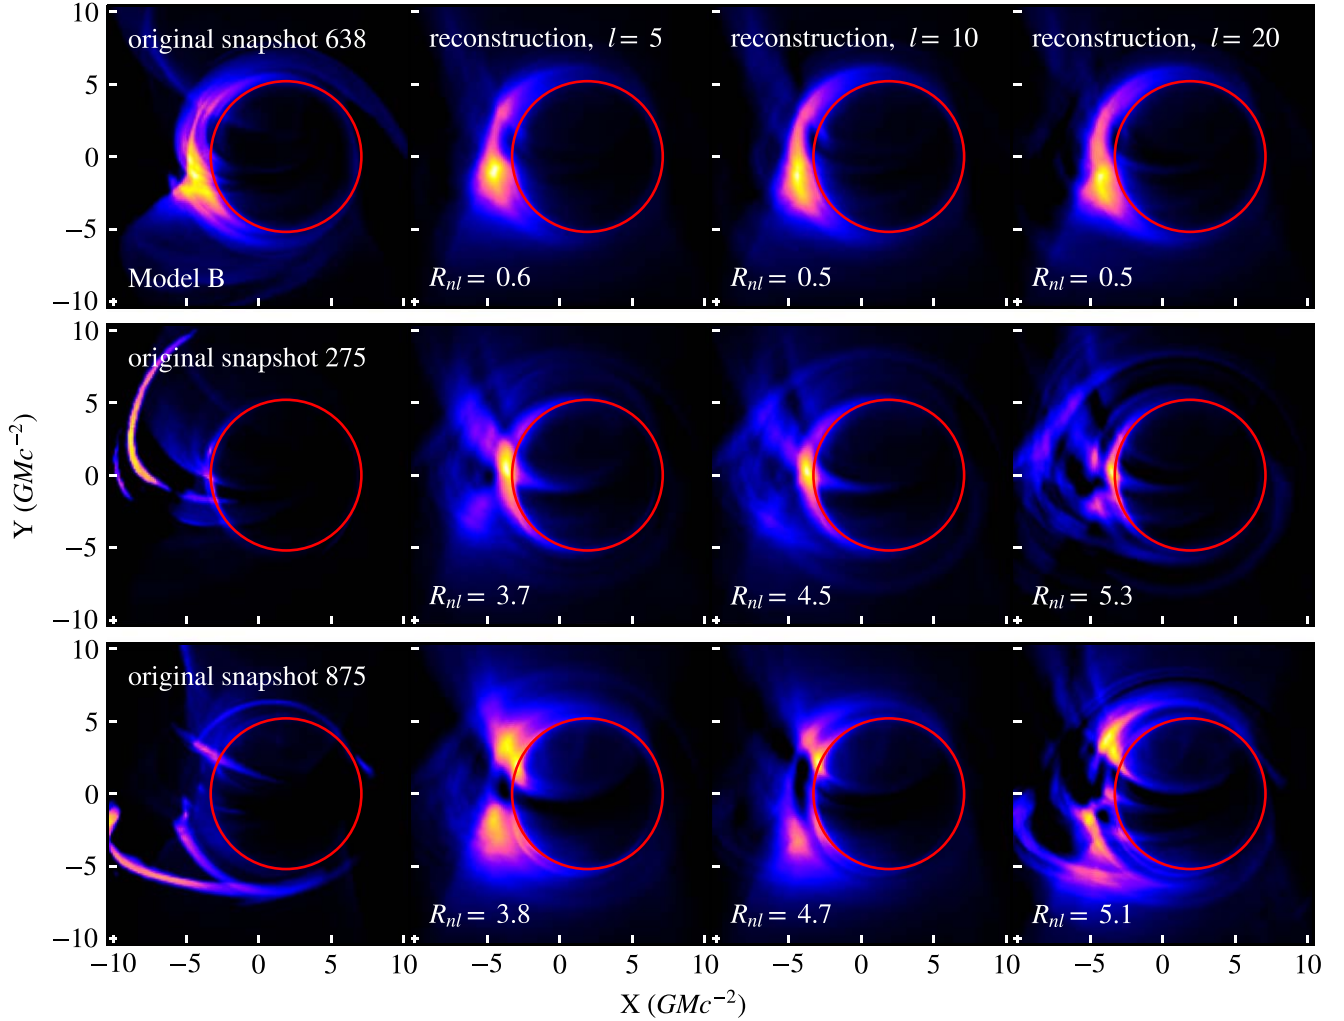

Figure 10. Left-most panels show three snapshots from the ensemble of images calculated for Model B. In each row, the three consecutive panels show reconstructions using the first 5, 10, and 20 principal components. The top row corresponds to a time instance that is not identified as an outlier but which corresponds to a large flux excursion from the accretion flow (this time step is denoted by "row 1" in Figure 9). Note that this image is well-fitted with the first 20 eigenimages and, thus, has a small R_{nl} value. The second row corresponds to a time step which is identified as an outlier with no significant flux excursion. The third row corresponds to a time instance that is both identified as an outlier and shows a large flux excursion. Both of these latter images are poorly fitted by even 20 eigenimages and have been identified as outliers by their R_{nl} values.
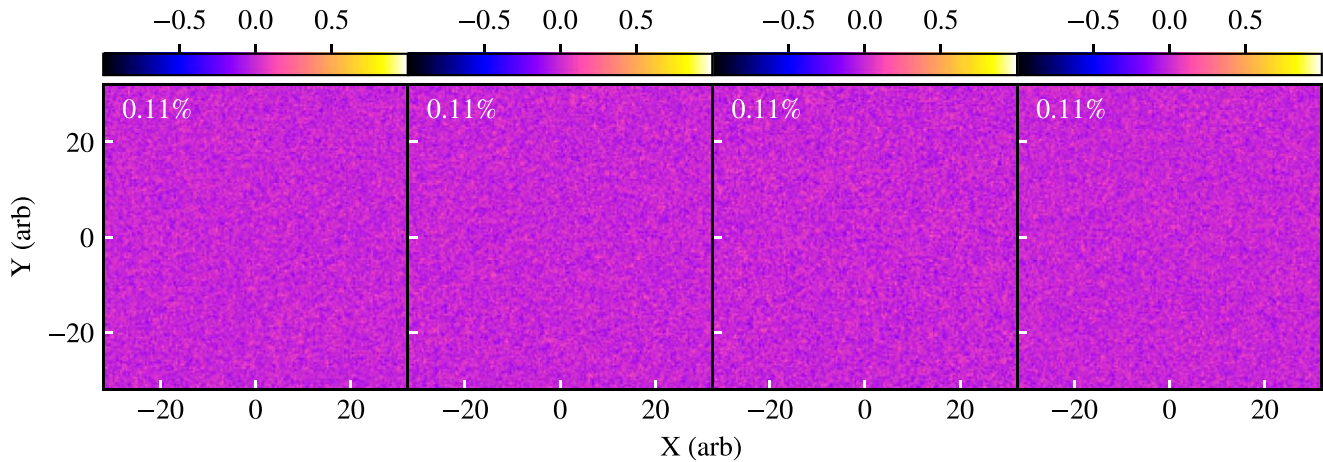
Figure 11. First four components of the PCA decomposition of 1024 images with purely Gaussian noise and their respective eigenvalues. The brightness of each pixel in these images is a random number taken from a Gaussian distribution centered at zero with a width of $\sigma = 0.5$. Because each image is uncorrelated from the rest, all principal components have very similar eigenvalues and dimensionality reduction using PCA is not possible for this system.

of GRMHD turbulence and, hence, whether measuring them in observations can help us better understand turbulence in accretion flows.

The power-law shapes of the eigenvalue spectra are reminiscent of noise processes. Consequently, we begin by exploring the PCA eigenvalue spectrum of Gaussian noise in an image and then continue with a red-noise process.

6.1. Gaussian Noise

We consider a Gaussian noise model where the brightness of each pixel is a random number taken from a Gaussian distribution centered at zero. We perform PCA on 1024 images with independent realizations of Gaussian noise with a standard deviation of $\sigma = 0.5$ over 512×512 pixels. Figure 11 shows the first few principal components and their respective eigenvalues. Any two images in the ensemble are statistically uncorrelated. However, the elements of matrix L are not zero because small, non-zero residual correlations between any two images remain because of the finite number of pixels in each image and the statistical nature of noise. The eigenvalues of all components are similar, indicating that all of the principal components are of similar importance and dimensionality reduction is not possible for this configuration.

The presence of minor correlations between pairs of images leads to a distribution of eigenvalues of finite width. In PCA, we count the eigenvectors in decreasing order of their eigenvalues; therefore, this distribution leads to a spectrum of eigenvalues with a non-zero slope. Figure 12 shows the spectrum of the eigenvalues of our realization of the Gaussian noise model. The eigenvalues are normalized such that they sum to unity (see Equation (9)). Given that our simulation of Gaussian noise involves $m = 1024$ images, there are 1024 non-trivial eigenvalues of similar magnitude with a mean of $1/1024 \simeq 0.098\%$. To estimate the standard deviation of the distribution of eigenvalues, we consider the fact that there are mN^2 individual realizations of the Gaussian noise in the ensemble of $m = 1024$ images with $N^2 = 512^2$ pixels each. Therefore, we expect the standard deviation of eigenvalues to be comparable to

$$\sigma = \frac{1}{\sqrt{mN^2}} \simeq 0.06 \left(\frac{m}{1024} \right)^{-1/2} \left(\frac{N}{512} \right)^{-1} \%. \quad (33)$$

The full range of eigenvalues in our particular realization of images is from 0.0859% to 0.1101%, which corresponds to a width of $\approx 4\sigma$. In Figure 12, we show the range of 4σ around the expected mean magnitude of the eigenvalues to visualize this result.

In contrast to Gaussian noise, the spectra of PCA eigenvalues for our GRMHD simulations, including the power-law tails at large eigenvector numbers, do not depend on either the number of images or the number of pixels. We tested this by decreasing our spatial and temporal resolution by factors of 2 and 4 but preserving the total time span and image size. This behavior indicates that the structures present in our simulations are much larger than the pixel size; consequently, changing the number of pixels does not alter the PCA decomposition. A similar argument is valid for the lack of dependence on the number of images. For this reason, we now turn our attention to noise spectra with maximum power at scales that are larger than the pixel sizes.

6.2. Red Noise

The spatial and time variability of images of accretion flows, such as those from Sgr A*, are expected to be approximated by red-noise power spectra. This is based both on the observationally measured flux variability of Sgr A* (Meyer et al. 2008; Dexter et al. 2014) and on theoretical models (e.g., Dolence et al. 2012 and Chan et al. 2015a). Other physical phenomena that affect black hole images, such as refractive scattering in the intervening medium, are also expected to introduce noise at different characteristic scales (e.g., Johnson & Narayan 2016). Because of such considerations, we consider here an ensemble of images with structure described by an isotropic, red-noise 2D Fourier spectrum given by

$$P(q) = 2^\alpha \pi \alpha e^{-(q/q_{max})^2} (q^2 + q_{min}^2)^{-(1+\alpha/2)} \quad (34)$$

such that the image brightness at a location given by the transverse vector $\boldsymbol{r}$ on the image plane is

$$I(\boldsymbol{r}) = I_0 \int d^2 q P(q) \exp[-i\boldsymbol{q} \cdot \boldsymbol{r}]. \quad (35)$$

Here, q_{min} and q_{max} determine the location of the first and second breaks in the spectrum and consequently the sizes of the
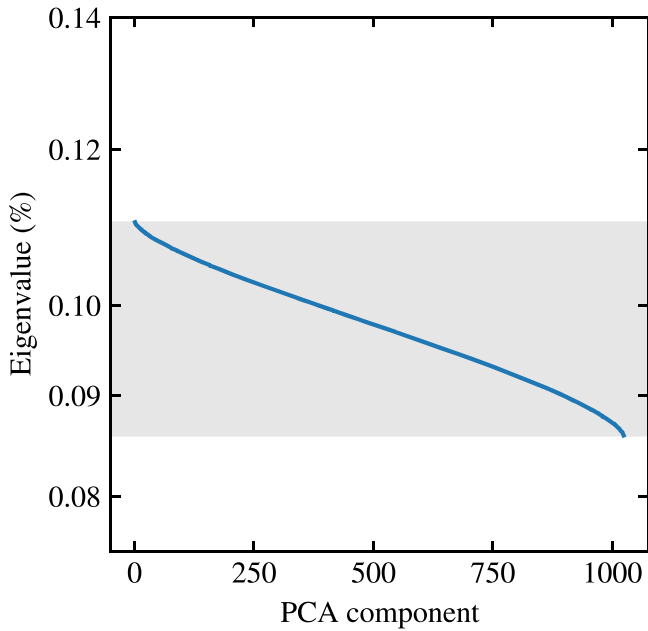
Figure 12. Eigenvalue spectrum for the PCA decomposition of an ensemble of images with Gaussian noise. The gray rectangle represents the expected range of eigenvalues, given the statistical nature of Gaussian noise (see text). Note the very small range of the y-axis.
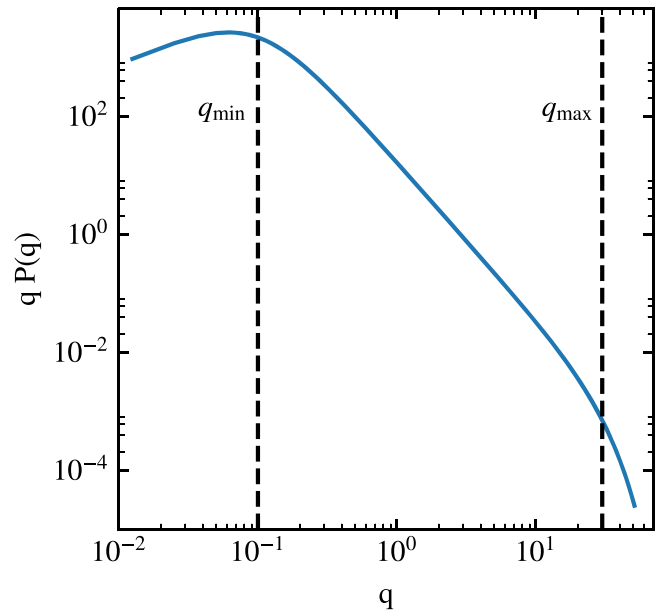

Figure 13. Red-noise spectrum given in Equation (34) with $q_{max} = 30$, $q_{min} = 0.5$, and $\alpha = 5/3$. The parameters q_{min} and q_{max} determine the locations of the first and second break in the spectrum, respectively, and α specifies the slope of the region between q_{min} and q_{max}.

largest and smallest structures in the images, respectively. The parameter α determines the slope of the power spectrum in the region between q_{min} and q_{max}.

Figure 13 shows a plot of this spectrum with $q_{max} = 30$, $q_{min} = 0.5$, and $\alpha = 5/3$. As expected, for $\alpha = 5/3$, there is very little power at the small scales where q_{max} is relevant. The majority of the power is at the larger scales related to q_{min}. Figure 14 shows some examples of images with different red-noise realizations. By construction, the structures in these images are almost entirely resolved in an $N \times N$ image, as long as the size of each pixel is much smaller than $1/q_{max}$.

To investigate the effect of red noise on the PCA of images, we construct numerous sets of 1024 images for different values of the red-noise parameters, such as those in Figure 14, and perform PCA on the set. We now explore the dependence of the PCA decomposition of these images on the parameters of the red-noise spectrum.

The spectrum of PCA eigenvalues of red noise does not depend on the number of pixels N per image, as long as the size of the dominant scale of the noise is fully resolved; i.e., as long as $L/N \ll 1/q_{min}$, where L is the linear size of the image. This is similar to the PCA results for the images of the GRMHD simulations and unlike those of the Gaussian noise simulations that were discussed earlier. The eigenvalue spectrum is also independent of q_{max} as long as $q_{max} \gg q_{min}$ and $\alpha > -1$ because, if these conditions are met, there is negligible power at scales $\sim 1/q_{max}$ to affect the PCA decomposition significantly.

Figure 15 shows the spectrum of PCA eigenvalues and its dependence on q_{min}. For $\alpha > -1$, $1/q_{min}$ determines the size of the dominant scale of the noise structures. The number of dominant noise structures that can fit in an image of size L is

$$n = \left(\frac{L}{1/q_{min}}\right)^2 = L^2 q_{min}^2. \tag{36}$$

Therefore, following the discussion in Section 3, we expect that this number corresponds to the number of dominant PCA components. This is shown in Figure 15, where the filled circle on each spectrum corresponds to the eigenvalue of the nth PCA component given by the previous relation. Clearly, as q_{min} increases, the structures on the images become smaller and more PCA components are necessary to reconstruct with fidelity the original ensemble of images.

For a given value of the parameter q_{min}, the number of images m in the ensemble determines whether or not the spectrum of eigenvalues has converged. Indeed, as we discussed earlier, for small values of q_{min}, which correspond to large dominant noise structures, a small number of eigenimages is required to reconstruct with fidelity the full ensemble of images. In this case, as long as $m \gg n = L^2 q_{min}^2$, the eigenvalue spectrum has converged and its shape depends only very weakly on the number m of images in the ensemble.

Figure 16 shows the spectrum of PCA eigenvalues for images with Fourier spectra characterized by different power-law indices $\alpha > -1$. As in Figure 15, the spectra are relatively flat until the nth PCA component but then turn into power laws with indices that appear to be correlated with α. For $\alpha \lesssim -1$, which we do not show, the dominant noise structures occur at small scales $\simeq 1/q_{max}$ and the resulting eigenvalue spectra are flat with very weak dependence on α.

We further explore the dependence of the eigenvalue spectra on $\alpha > -1$ by fitting the higher components of each eigenvalue spectrum with a power-law function of the form $\lambda_k \sim k^{-\gamma}$ and show in Figure 17 the dependence of the fitted power-law index γ on α. We find this dependence to be

$$\gamma \simeq \frac{5}{4}\alpha + 2. \tag{37}$$

Note that, for the simulations used in generating Figure 17, we set $q_{min} = 0.1$ to force the breaks of the eigenvalue spectra to
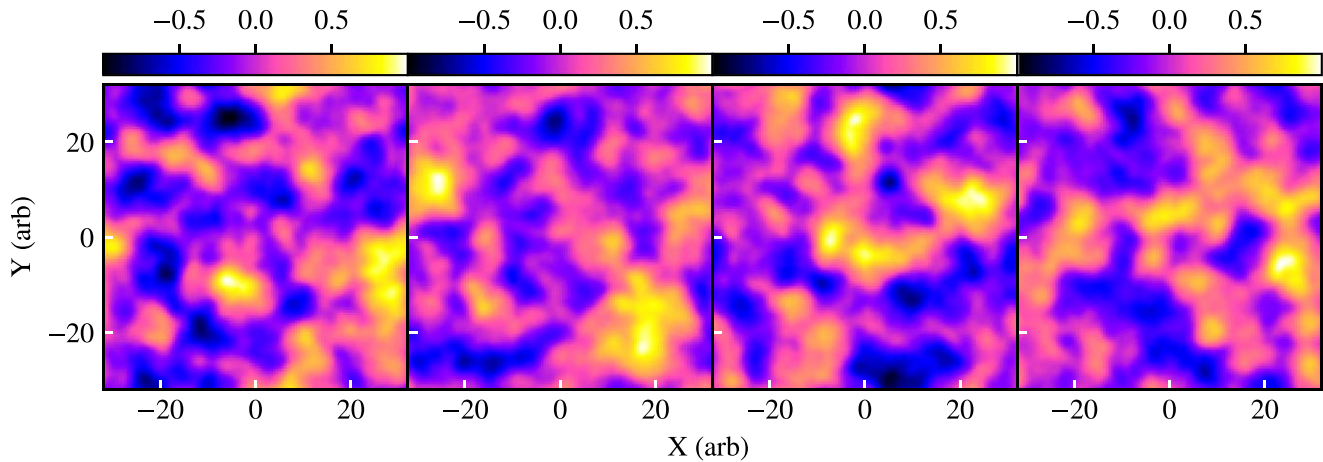
Figure 14. Examples of images with different realizations of red noise with the isotropic power spectrum shown in Figure 13 and random phase fluctuations. As expected for $\alpha = 5/3$, most of the power is at scales $\simeq 1/q_{min}$.
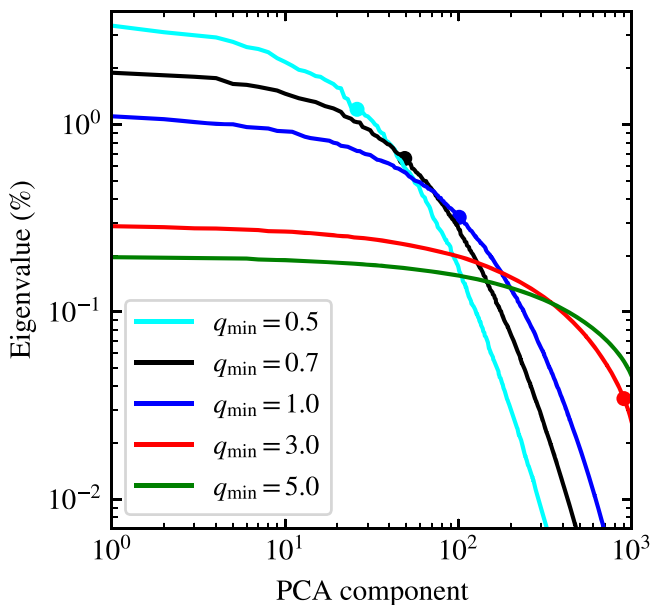

Figure 15. Spectra of PCA eigenvalues for ensembles of images with isotropic red noise and for different values of the parameter q_{min}; the remaining parameters are the same as in Figure 13. For lower values of q_{min} the dominant scale of the structures in the images is larger and fewer PCA components are required to reproduce the majority of structure in the images. The filled circles on each curve indicate the number of PCA components that are equal to the approximate number of different noise structures that can fit in the image; i.e., where the number of components is equal to $L^2 q_{min}^2$, where L is the size of the image.

Figure 16. Spectra of PCA eigenvalues for ensembles of images with isotropic red noise and for different values of the power-law index α; the remaining parameters are the same as in Figure 13. For comparison, the eigenvalue spectra of the GRMHD models B, C, and D are also included in the black, blue, and red dashed lines, respectively. The power-law slope of the eigenvalue spectrum after the break depends strongly on α.

occur at low PCA components and, thus, to allow for a more accurate determination of the power-law index of the spectra. This result demonstrates that the 2D Fourier spectrum of the structures in the image plane determine in a predictable way the high-end spectrum of PCA eigenvalues and, therefore, the latter can be used to infer the former.

6.3. The Small-scale Structures of Black Hole Images from GRMHD Simulations

In Figure 16, we compare the spectra of PCA eigenvalues from the black hole images of GRMHD simulations to those of the images with red-noise Fourier spectra. The large range of eigenvalues in the GRMHD simulations is clearly inconsistent
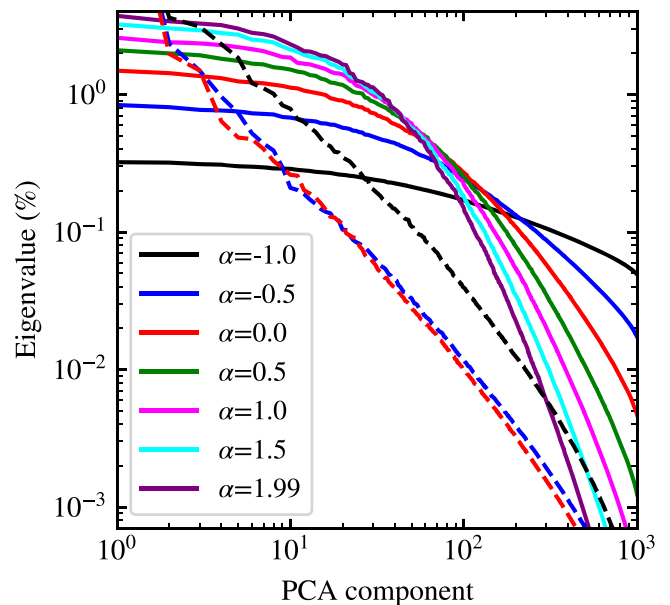
with the small expected range of eigenvalues for purely Gaussian noise (see also Figure 12). This suggests a more complex origin of image structure and variability than what has been assumed in the past (cf. Broderick et al. 2016).

The spectra of PCA eigenvalues for the images of GRMHD simulations become power laws after only the first handful of PCA components. This suggests that Lq_{min} for these simulations is a small number (see Equation (36) and Figure 15) and, therefore, that the typical scale of variable structure in the images is comparable to the size of the images themselves. In other words, it is comparable to the size of the black hole shadow. This is consistent with the discussion in Medeiros et al. (2017, 2018), who attributed the variability of the simulated interferometric amplitudes and closure phases to overall changes in the widths of the crescent-like images, as well as to the appearance and disappearance from the images of
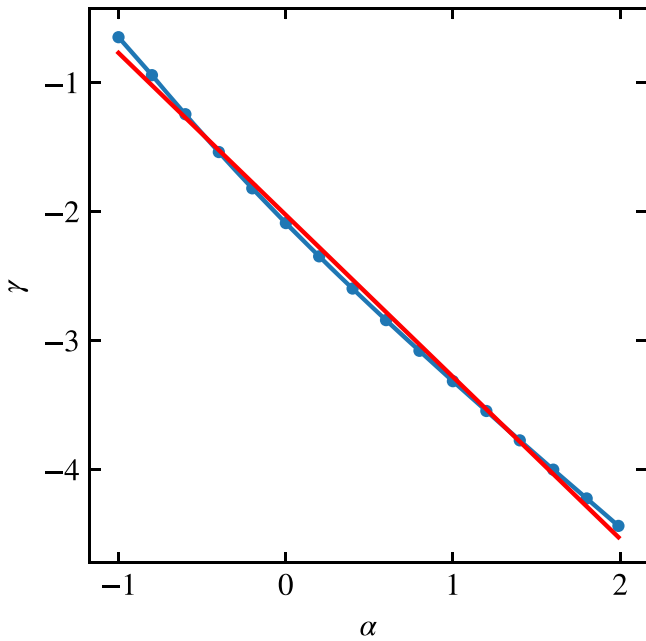
Figure 17. Power-law index γ of the spectra of PCA eigenvalues as a function of the power-law index α of the red-noise Fourier spectra that was used to generate the images. The blue circles are the values obtained numerically from fitting the eigenvalue spectra with power-law functions; the red line is a linear fit to the blue circles.

large, hot, and, therefore, bright magnetic flux tubes that orbit the black hole.

The power-law indices in the eigenvalue spectra of the images from GRMHD simulations are nearly independent of the underlying model and equal to $\gamma \simeq 1.3$. Using Equation (37), we find that this implies a power-law index for the 2D Fourier spectrum of the variable structures of $\alpha \simeq -0.5$. It is important to emphasize here that this power-law index characterizes the 2D Fourier spectrum of the variable structures, which are determined in a complex, nonlinear way by the anisotropies in the density, temperature, emissivity, magnetic field, and lensing in the vicinity of the black hole. It is, therefore, not surprising that the inferred value of α does not reflect the underlying power spectrum of the MHD turbulence in the accretion flow.

7. Conclusion

Understanding the horizon-scale millimeter emission around an accreting black hole requires a two-pronged approach. One component is to utilize our best understanding of the physics to generate high-fidelity GRMHD simulations of the morphology of the emission. The second is to use an interferometer, such as the EHT, to test the understanding of the physics with real observations. There is a common theme in both components: the question of how we characterize and extract the salient information in an ensemble of images. In this paper, we have demonstrated that PCA offers an effective tool for this task over a wealth of different problems.

Focusing purely on the simulations, we showed that PCA offers an extremely compact representation of the theoretical millimeter images. Each simulation comprises over 1000 distinct images, yet we find that we can represent *most* of the images with only a few to a few dozen eigenimages, depending

on the desired fidelity. Moreover, recognizing images poorly represented by the leading eigenimages is critical and represents another useful application of PCA. As detailed in the introduction, the temporal variability of the strength and morphology of the millimeter emission close to the horizon is a phenomenon that can limit or compromise the construction of interferometric images. Knowledge of the amplitudes of the eigenimages needed to represent any given image can be used to define a simple scalar metric that flags outliers in either the simulations or observations. This approach has already provided the realization that outliers may be more subtle than had been presumed. It had been supposed that flares in flux would correspond to events in which the emission morphology would show strong departures from the average form. Yet the outlier metric R_{nl} (see Equation (30)) allowed us to identify both images that had unusual morphology with no significant excursion in flux, as well as flare events that had perfectly ordinary morphology. As useful as this particular metric is in this work, we emphasize that other metrics and classifiers can be constructed from the locations of the simulated images in their eigenspace. Our goal here is not to strongly advocate any particular metric but to provide a useful example of what is possible within the PCA representation.

Apart from the identification of outliers, we have also demonstrated the use of the eigenvalue spectrum to characterize the properties of the noise and turbulent structure in the simulations. This approach shows a path for allowing the rapid quantitative evaluation of GRMHD simulations over a significant time span of accretion. As with the outlier metric R_{nl}, other diagnostic metrics can be built around the locations or trajectories of the simulated images as a function of time in their eigenspace.

Lastly, we showed that PCA may be applied directly to the analysis of interferometric data because the Fourier transform of the principal components of a set of images is equivalent to the principal components of the set of Fourier transformed images. When coupled with the dimensionality reduction that we discussed above, this property opens the possibility of using PCA for efficient image reconstruction from sparse interferometric data. In parallel, the PCA approach can be incorporated into the Bayesian inference method discussed in Kim et al. (2016) to generate efficient comparisons of EHT data to large suites of GRMHD simulations. Additionally, analysis of the temporal variability of the amplitude of each eigenimage may be fruitful for further understanding of the variability in both the simulations and the data. We will explore these avenues in future work.

ORCID iDs

Lia Medeiros https://orcid.org/0000-0003-2342-6728
Tod R. Lauer https://orcid.org/0000-0003-3234-7247

References

Akiyama, K., Ikeda, S., Pleau, M., et al. 2017, AJ, 153, 159
Boroson, T. A., & Lauer, T. R. 2010, AJ, 140, 390
Bouman, K. L., Johnson, M. D., Dalca, A. V., et al. 2017, arXiv:1711.01357
Bouman, K. L., Johnson, M. D., Zoran, D., et al. 2016, arXiv:1512.01413
Broderick, A. E., Fish, V. L., Johnson, M. D., et al. 2016, ApJ, 820, 137
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., et al. 2015a, ApJ, 812, 103
Chan, C.-K., Psaltis, D., Özel, F., Narayan, R., & Sądowski, A. 2015b, ApJ, 799, 1
Dexter, J., Agol, E., & Fragile, P. C. 2009, ApJL, 703, L142
Dexter, J., Kelly, B., Bower, G. C., et al. 2014, MNRAS, 442, 2797
Do, T., Ghez, A. M., Morris, M. R., et al. 2009, ApJ, 691, 1021
Doeleman, S. S., Weintroub, J., Rogers, A. E. E., et al. 2008, Natur, 455, 78
Dolence, J. C., Gammie, C. F., Shiokawa, H., & Noble, S. C. 2012, ApJL, 746, L10
Gammie, C. F., McKinney, J. C., & Tóth, G. 2003, ApJ, 589, 444
Genzel, R., Schödel, R., Ott, T., et al. 2003, Natur, 425, 934
Gold, R., McKinney, J. C., Johnson, M. D., & Doeleman, S. S. 2017, ApJ, 837, 180
Honma, M., Akiyama, K., Uemura, M., & Ikeda, S. 2014, PASJ, 66, 95
Johnson, M. D., Bouman, K. L., Blackburn, L., et al. 2017, ApJ, 850, 172
Johnson, M. D., & Narayan, R. 2016, ApJ, 826, 170
Kim, J., Marrone, D. P., Chan, C.-K., et al. 2016, ApJ, 832, 156
Lu, R.-S., Roelofs, F., Fish, V. L., et al. 2016, ApJ, 817, 173
Mahalanobis, P. C. 1936, Proc. National Institute of Sciences (Calcutta), 2, 49
Medeiros, L., Chan, C.-k., Özel, F., et al. 2017, ApJ, 844, 35
Medeiros, L., Chan, C.-k., Özel, F., et al. 2018, ApJ, 856, 163
Meyer, L., Do, T., Ghez, A., et al. 2008, ApJL, 688, L17
Mościbrodzka, M., Dexter, J., Davelaar, J., & Falcke, H. 2017, MNRAS, 468, 2214
Mościbrodzka, M., Gammie, C. F., Dolence, J. C., Shiokawa, H., & Leung, P. K. 2009, ApJ, 706, 497
Narayan, R., Sądowski, A., Penna, R. F., & Kulkarni, A. K. 2012, MNRAS, 426, 3241
Neilsen, J., Markoff, S., Nowak, M. A., et al. 2015, ApJ, 799, 199
Psaltis, D., Özel, F., Chan, C.-K., & Marrone, D. P. 2015, ApJ, 814, 115
Rauch, K. P., & Blandford, R. D. 1994, ApJ, 421, 46
Remillard, R. A., & McClintock, J. E. 2006, ARA&A, 44, 49
Sądowski, A., Narayan, R., Penna, R., & Zhu, Y. 2013, MNRAS, 436, 3856
Strang, G. 1988, Linear Algebra and its Applications (3rd ed.; Orlando, Fl: Harcourt Brace Jovanovich)
Turk, M., & Pentland, A. 1991, Journal of Cognitive Neuroscience, 3, 71
van der Klis, M. 2000, ARA&A, 38, 717