Seeing Beyond the Blur: Imaging Black Holes with Increasingly Strong Assumptions

Katie Bouman
Caltech Departments of Computing and Mathematical Sciences,
Electrical Engineering, and Astronomy

[redacted]

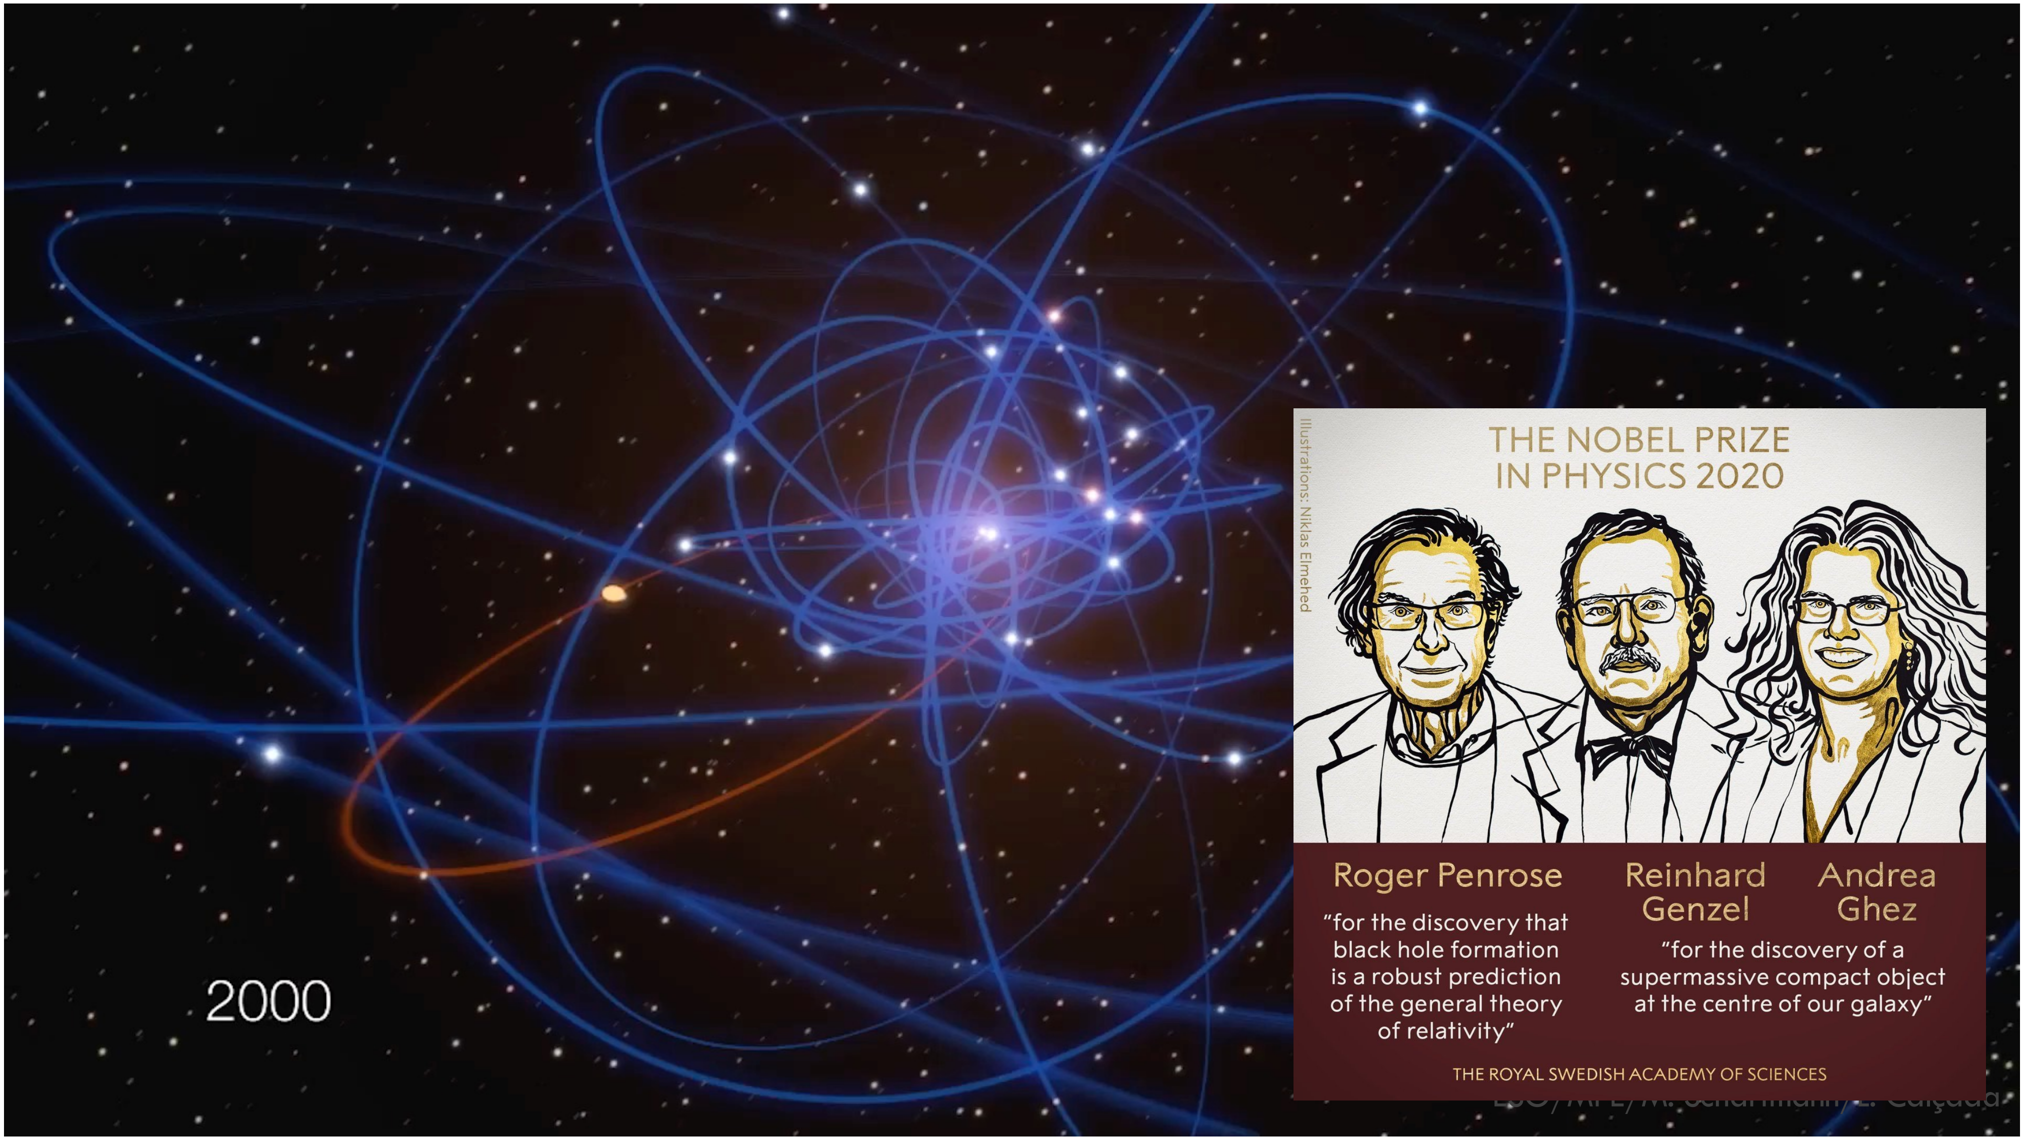
2000
THE NOBEL PRIZE
IN PHYSICS 2020
Illustrations: Niklas Elmehed
Roger Penrose
"for the discovery that black hole formation is a robust prediction of the general theory of relativity"
Reinhard Genzel
Andrea Ghez
"for the discovery of a supermassive compact object at the centre of our galaxy"
THE ROYAL SWEDISH ACADEMY OF SCIENCES

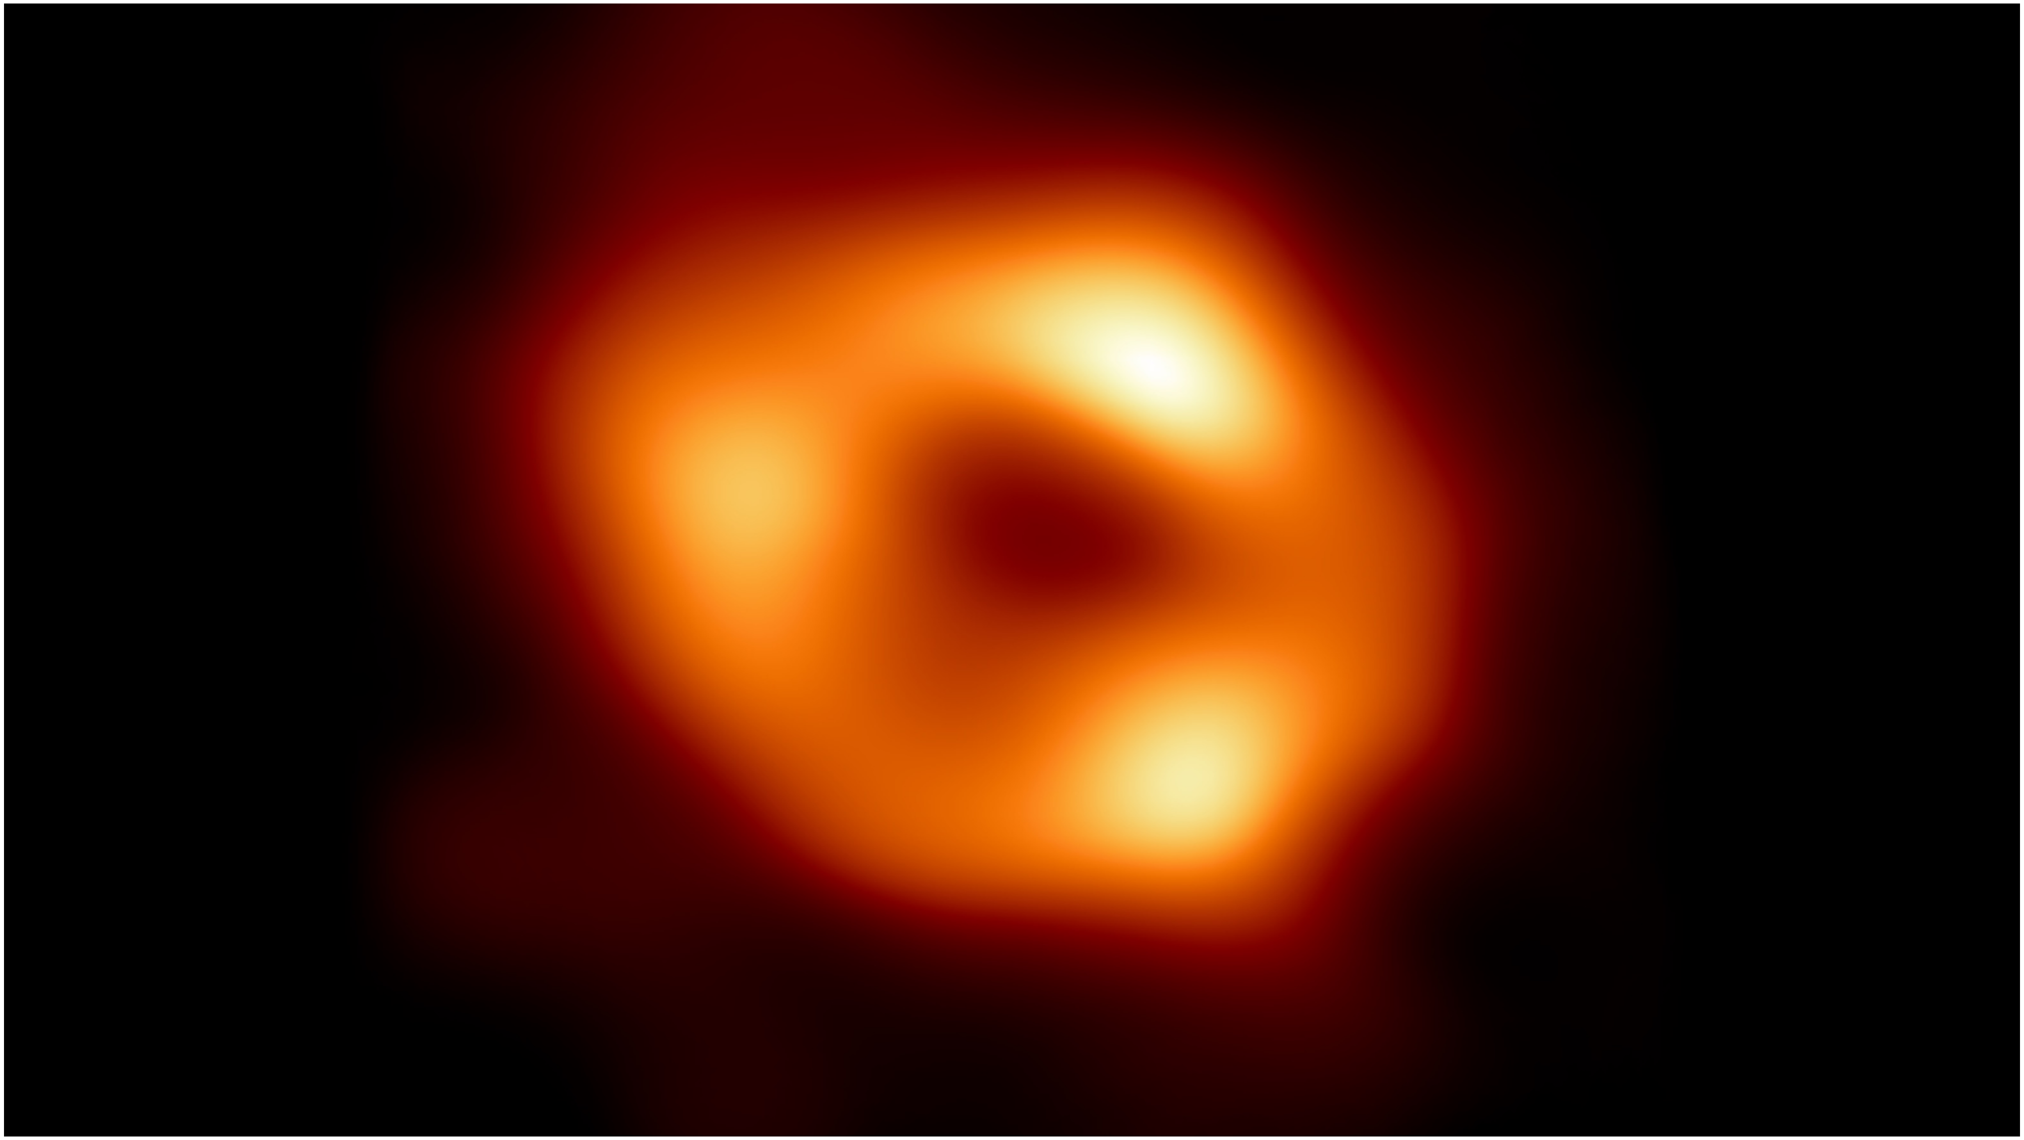

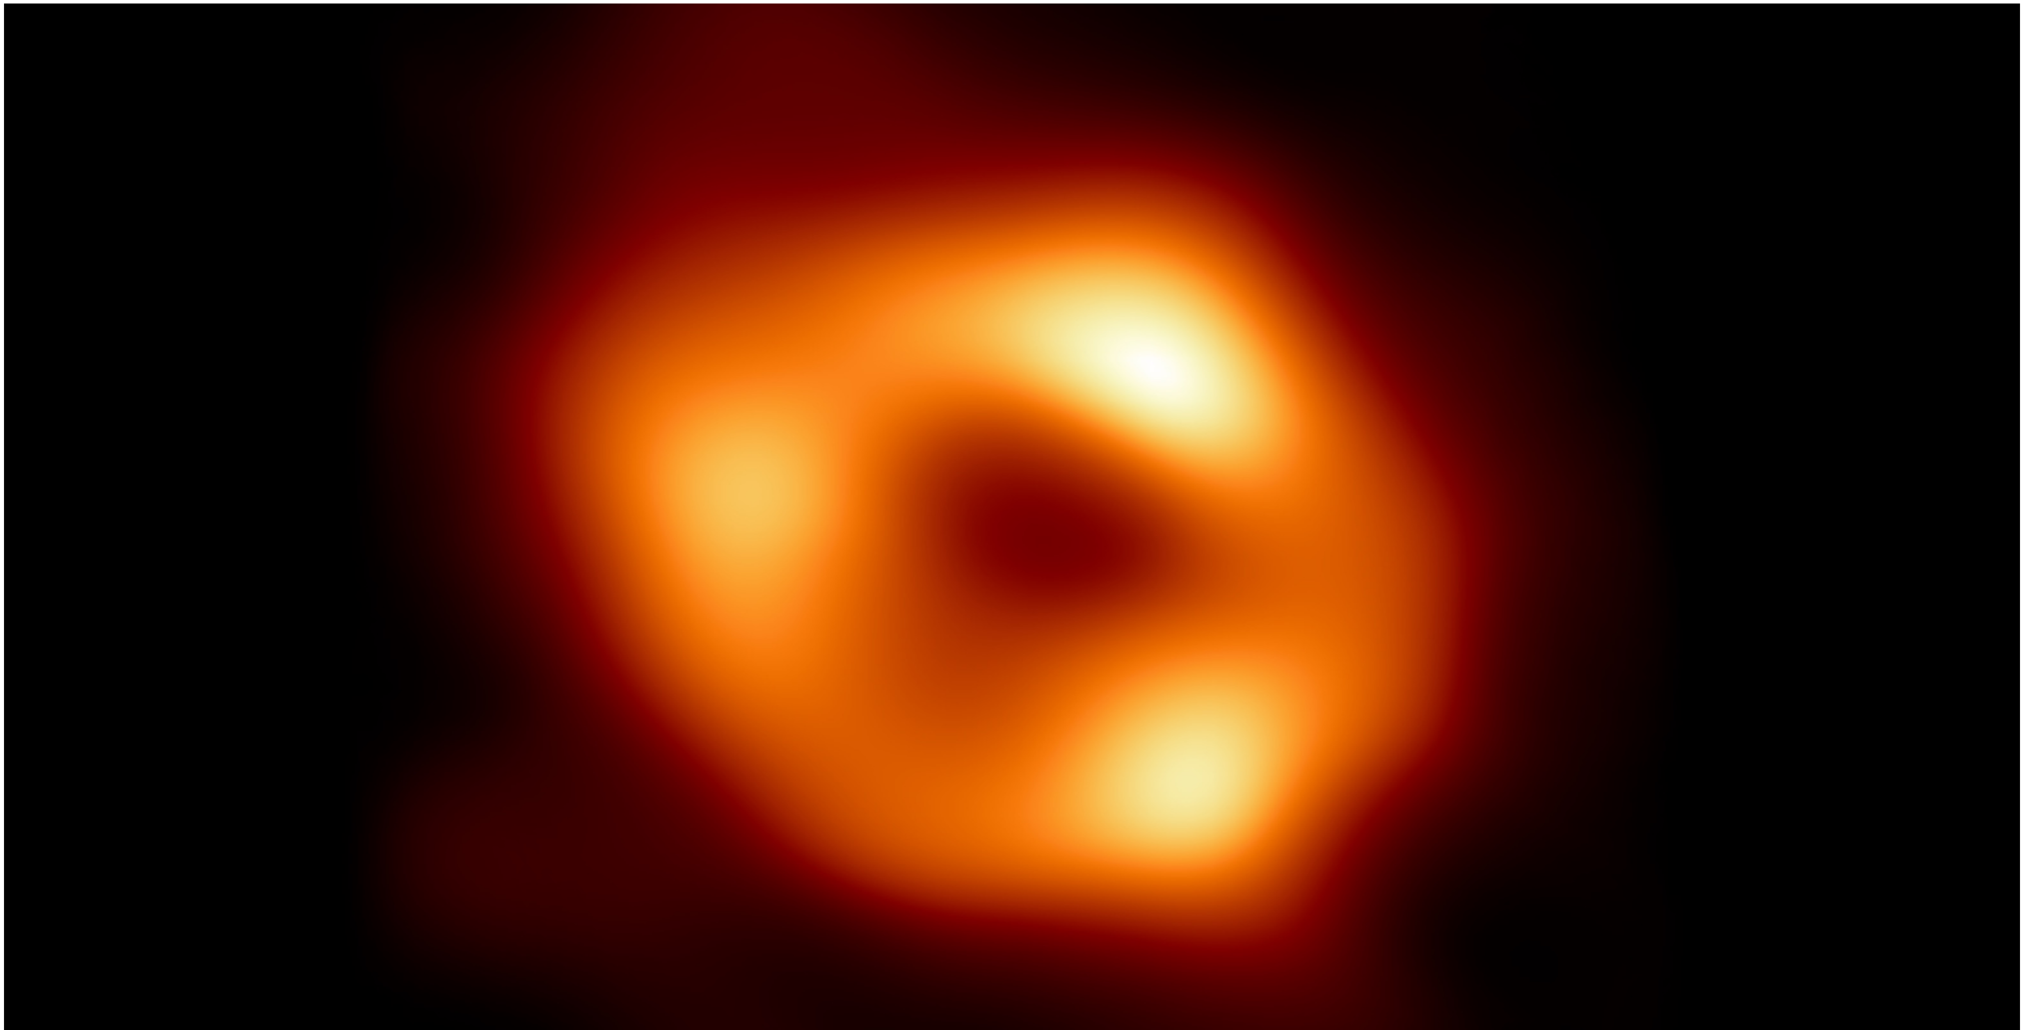

Sagittarius A* (Sgr A*): Black Hole at the Center of the Milky Way

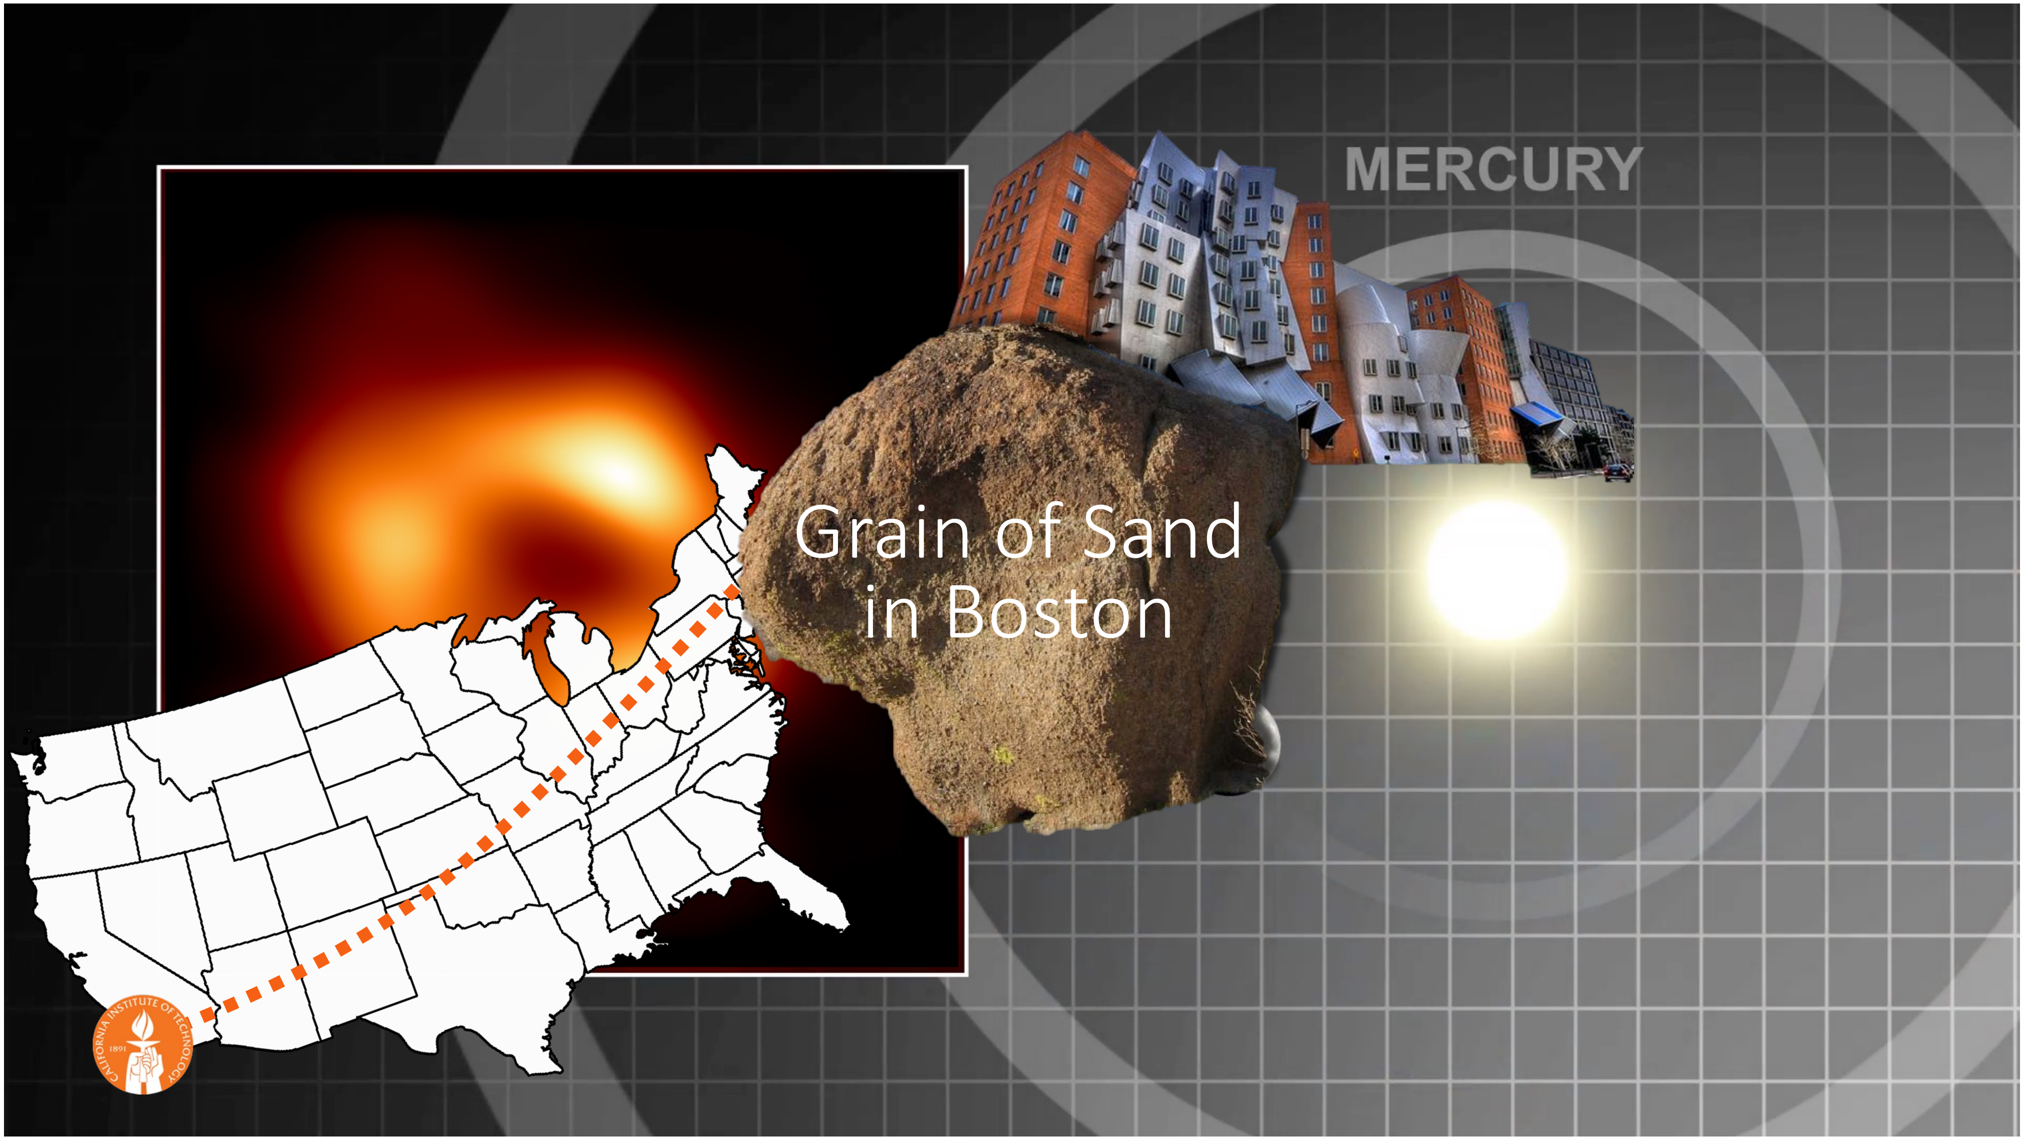
MERCURY
Grain of Sand
in Boston


The Event Horizon Telescope Collaboration
Over 300 Scientists from 80 institutes in countries spanning
Europe, Asia, Africa, North and South America
(along with ~23K Community Contributors from Open-Source Projects)

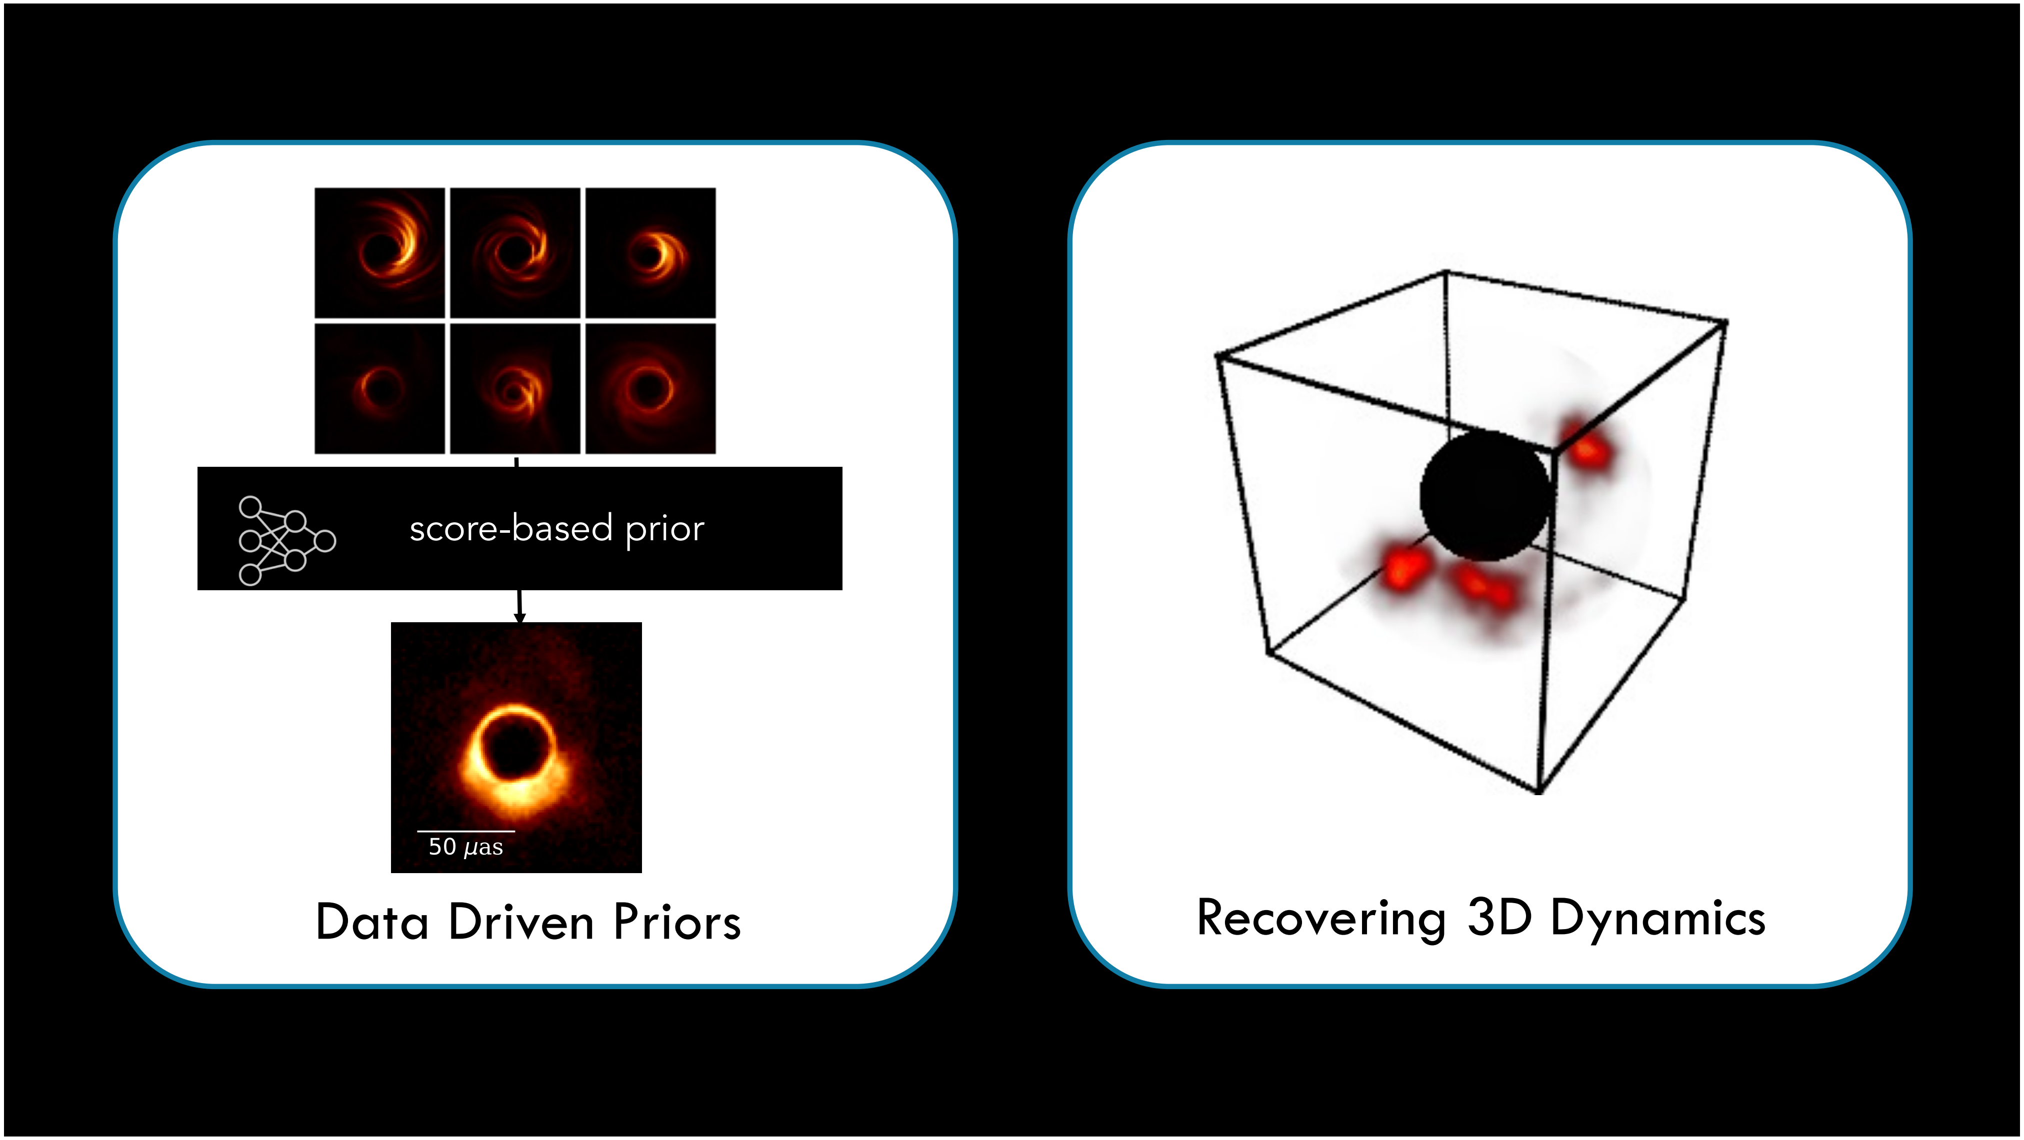
score-based prior
50 μas
Data Driven Priors
Recovering 3D Dynamics

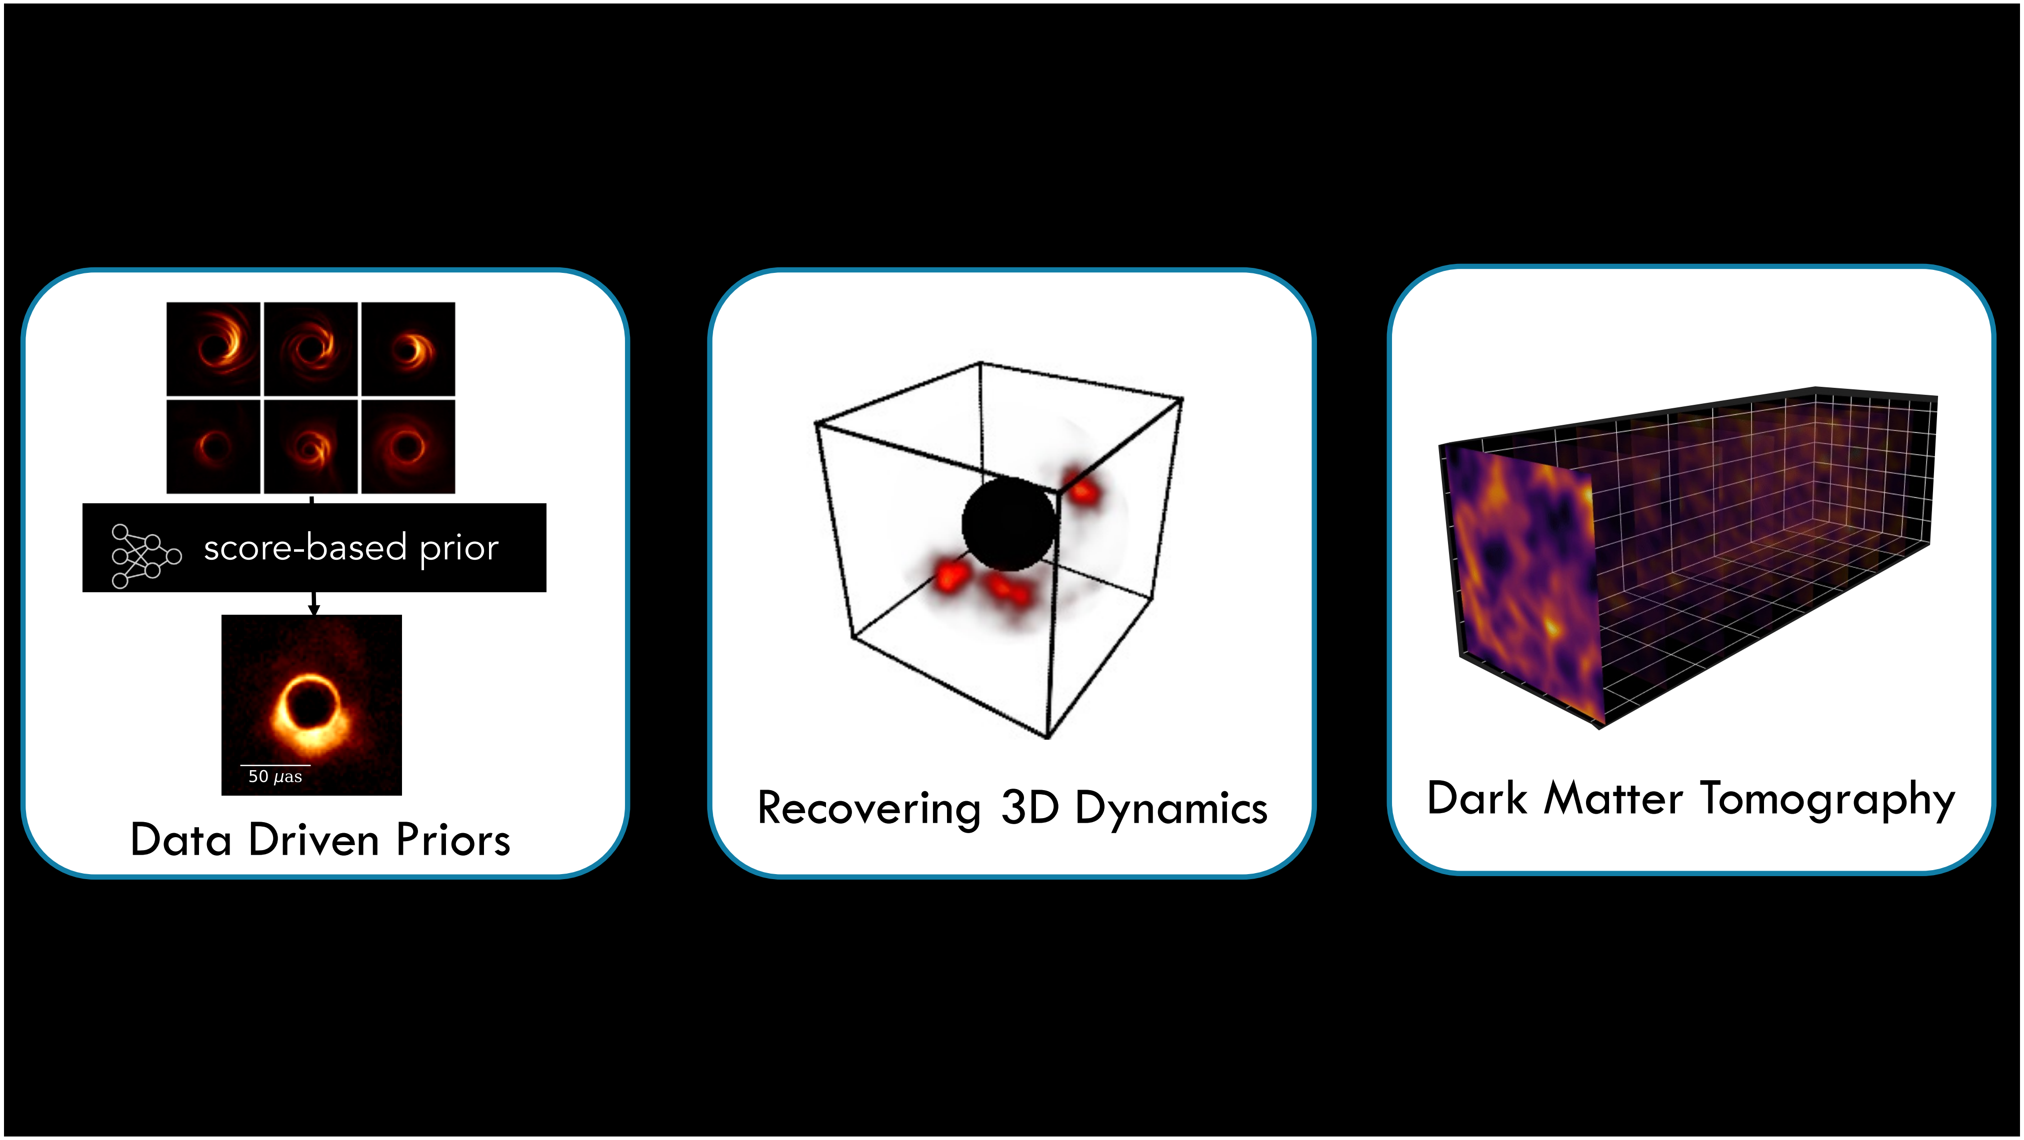
score-based prior
50 μas
Data Driven Priors
Recovering 3D Dynamics
Dark Matter Tomography

How Big Must Our Telescope Be?

$$\text{Telescope Size} \propto \frac{\text{Wavelength}}{\text{Angular Resolution}}$$

$$13 \text{ million meters} \propto \frac{1.3 \text{ mm}}{20\ \mu as}$$

Black Hole Simulation

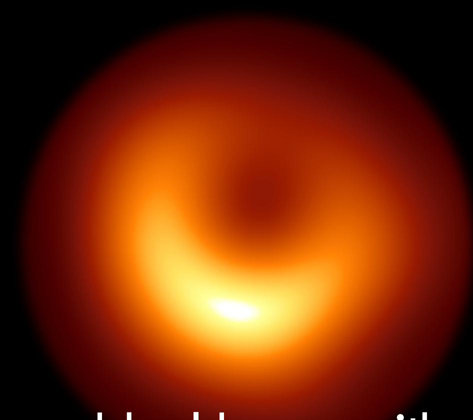

Ideal Image with
Earth-Sized Telescope


The Event Horizon Telescope (EHT)

The Event Horizon Telescope (EHT)

Black Hole Image

Frequency Measurements

North-South Frequency (v)

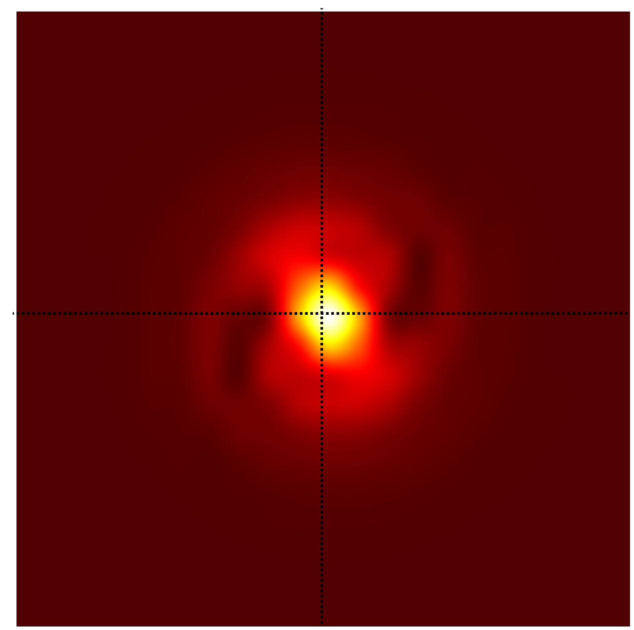

East West Frequency (u)

The Event Horizon Telescope (EHT)

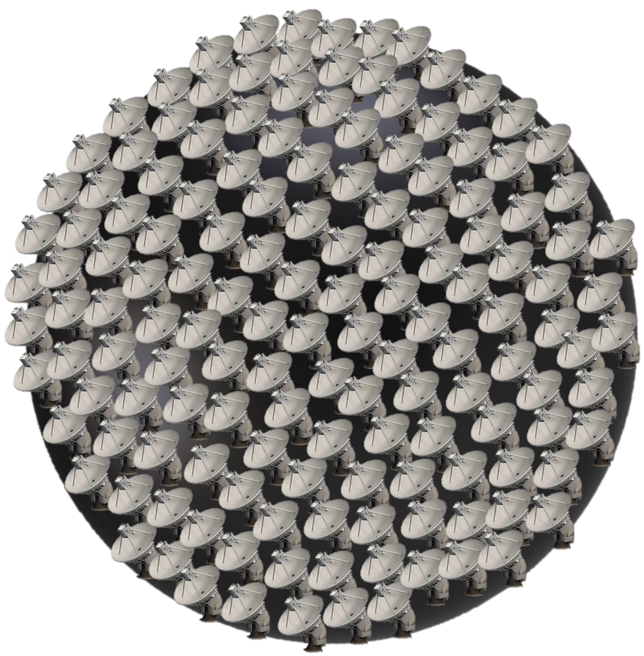

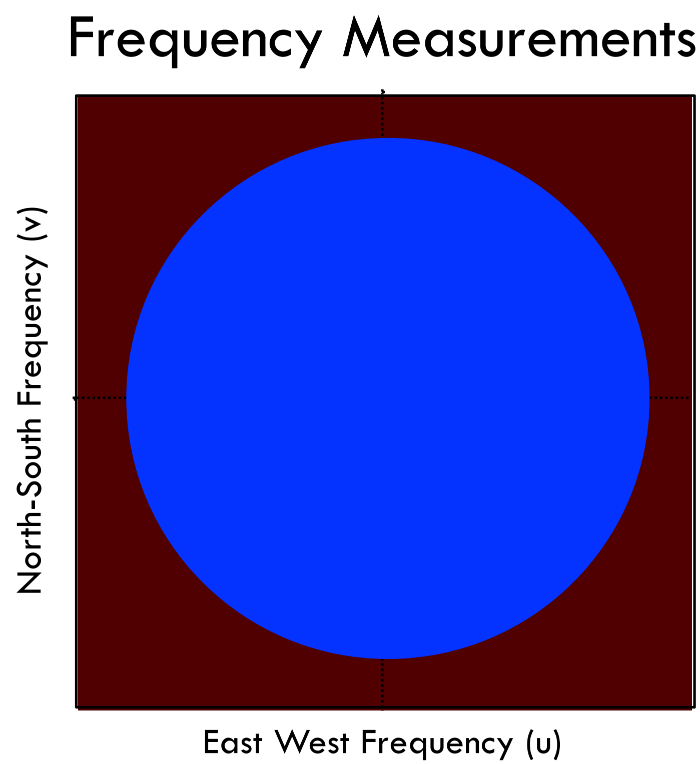

The Event Horizon Telescope (EHT)

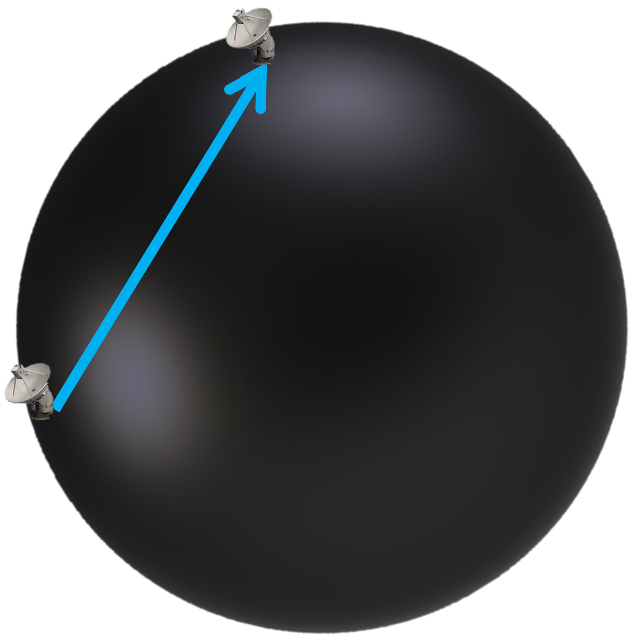

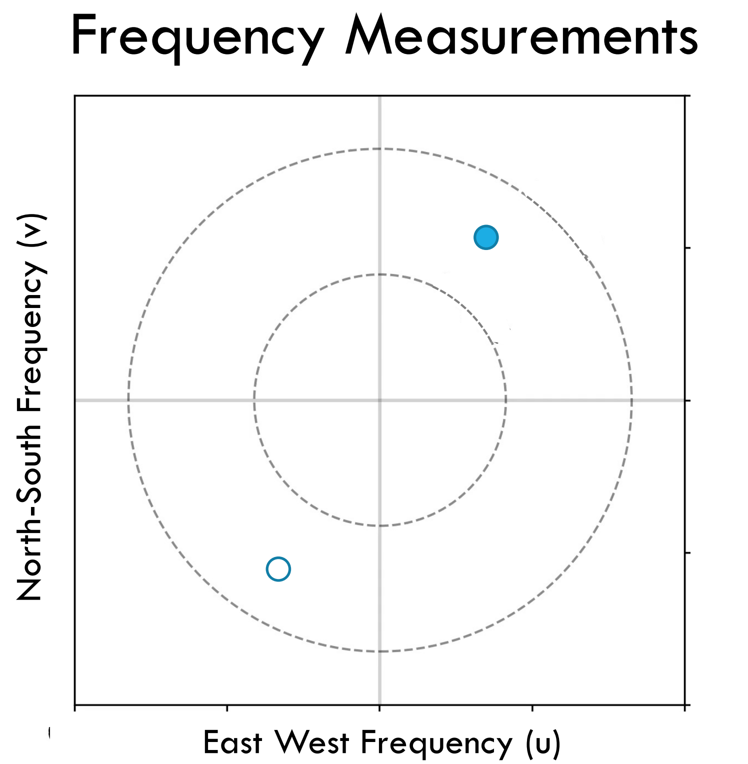

The Event Horizon Telescope (EHT)

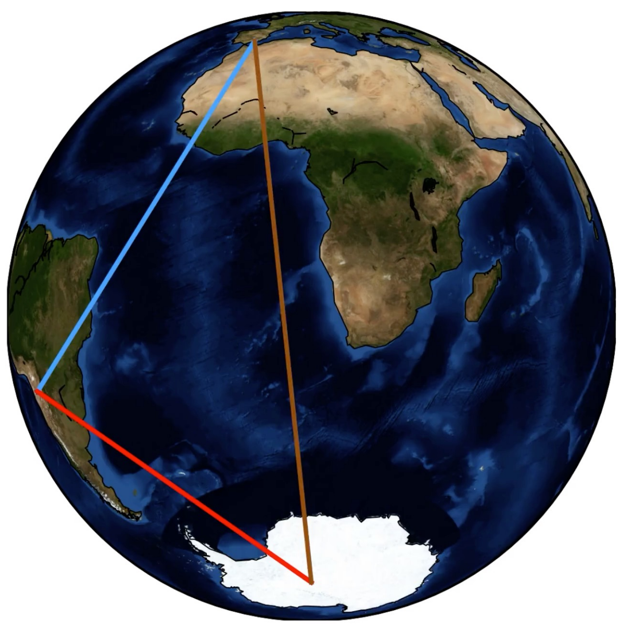

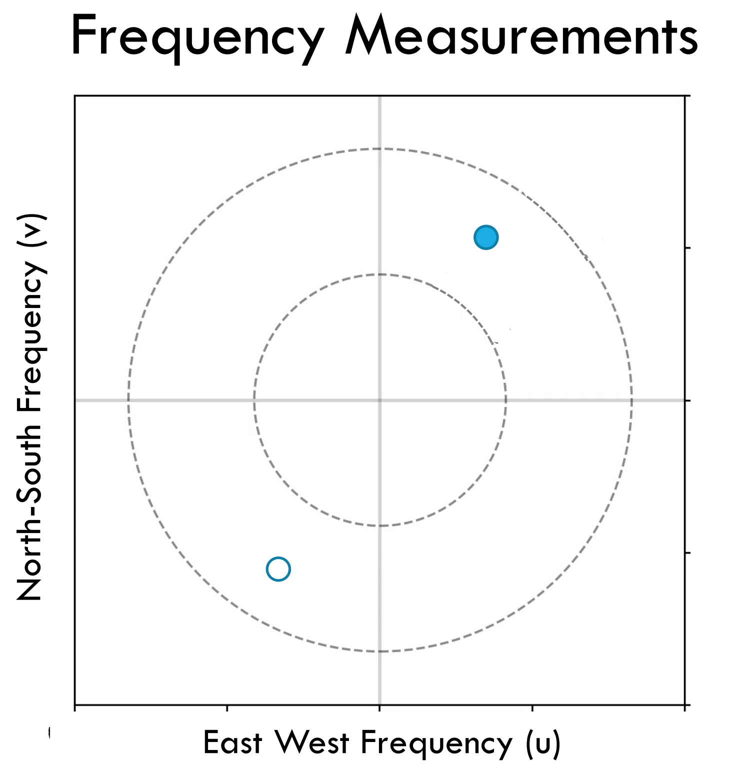

Solving for the Black Hole Image

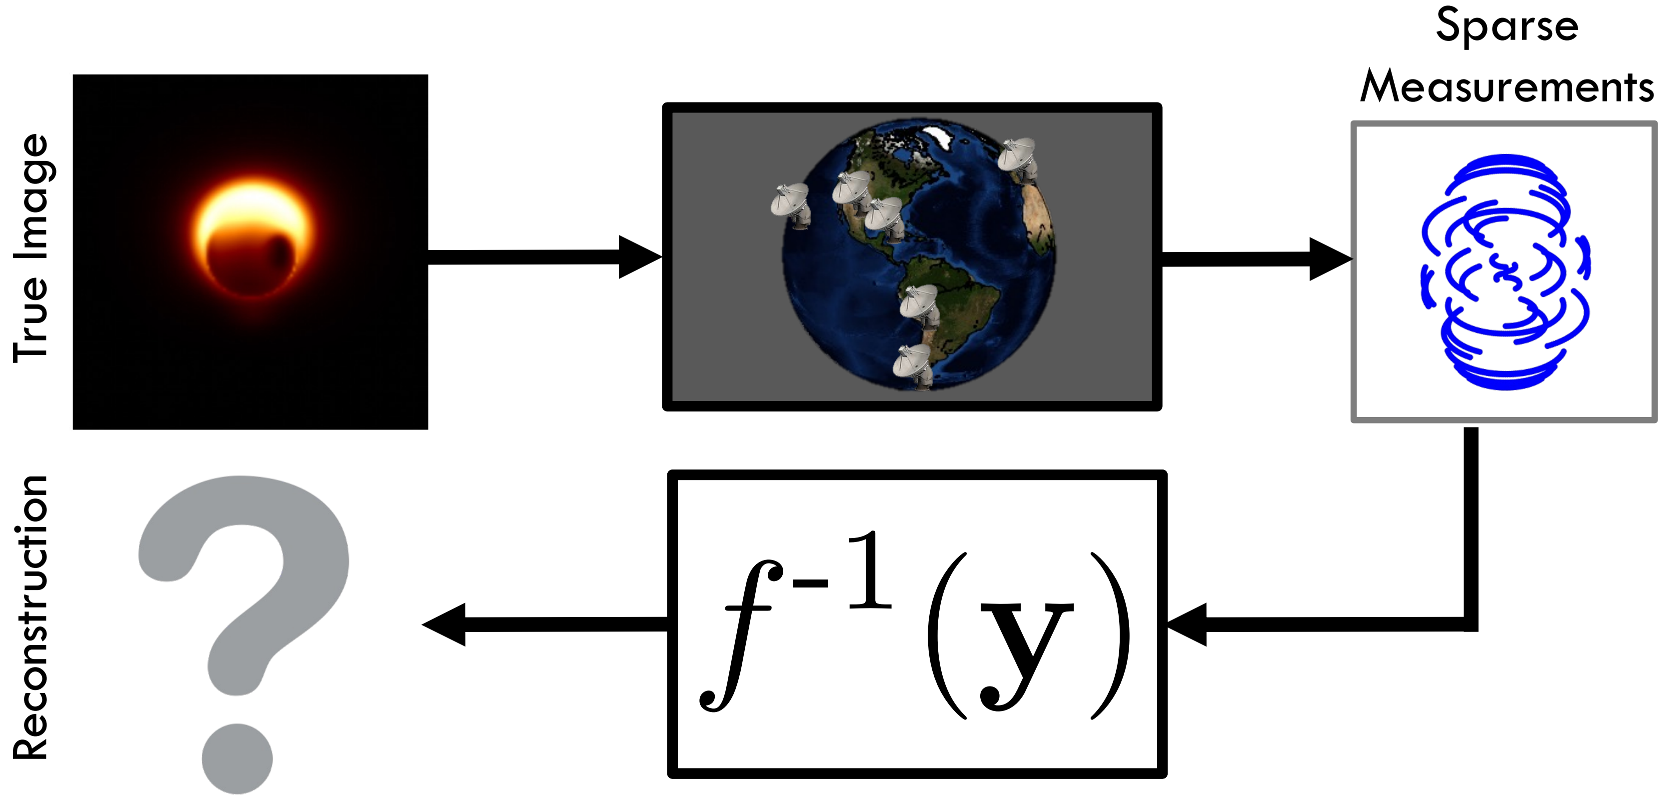

Solving for the Black Hole Image

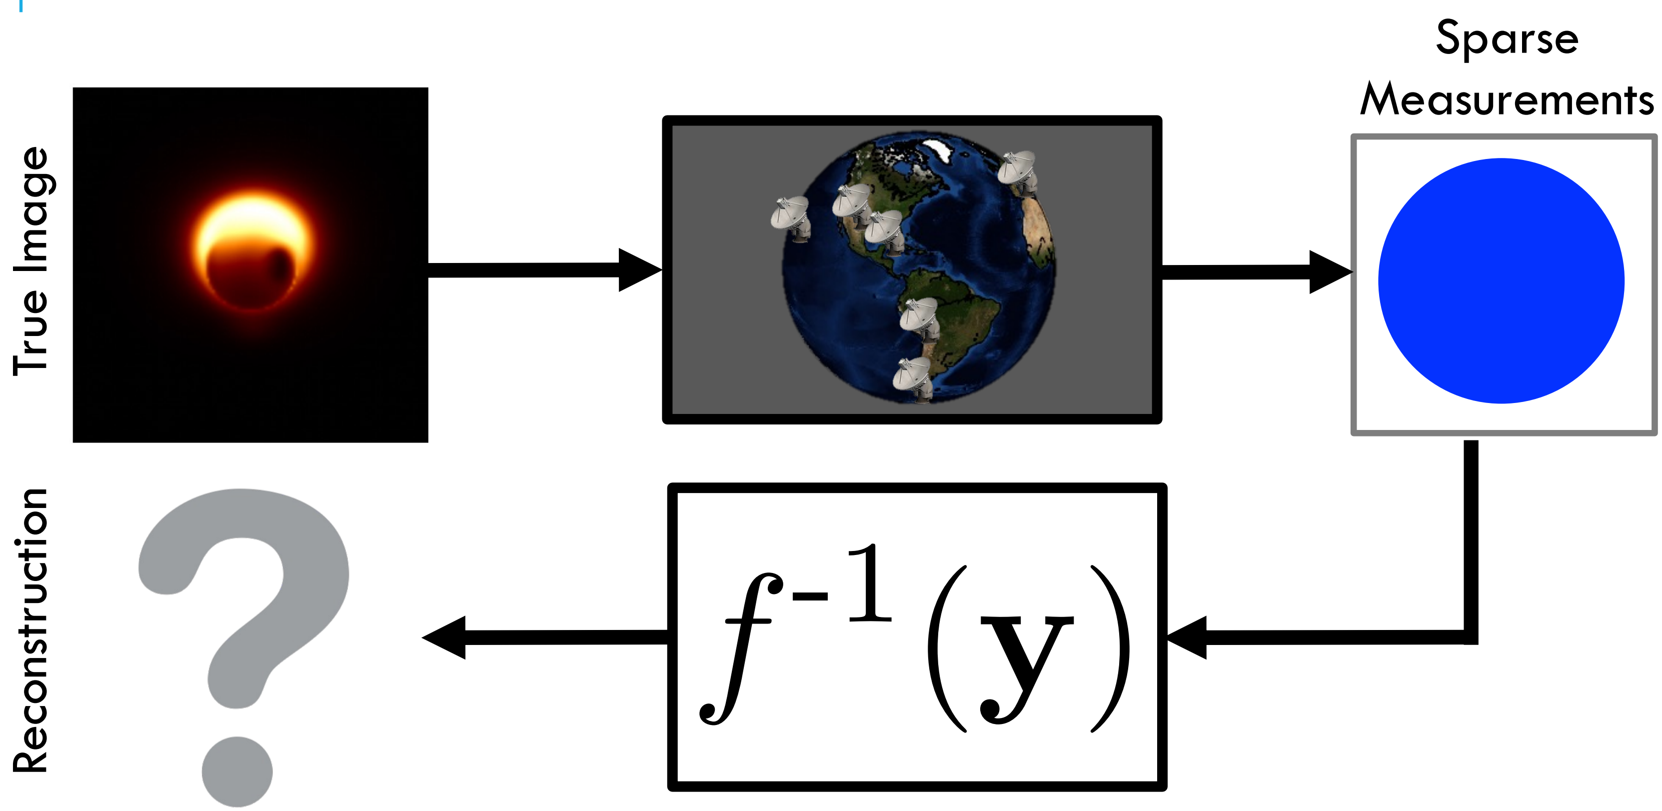

Solving for the Black Hole Image

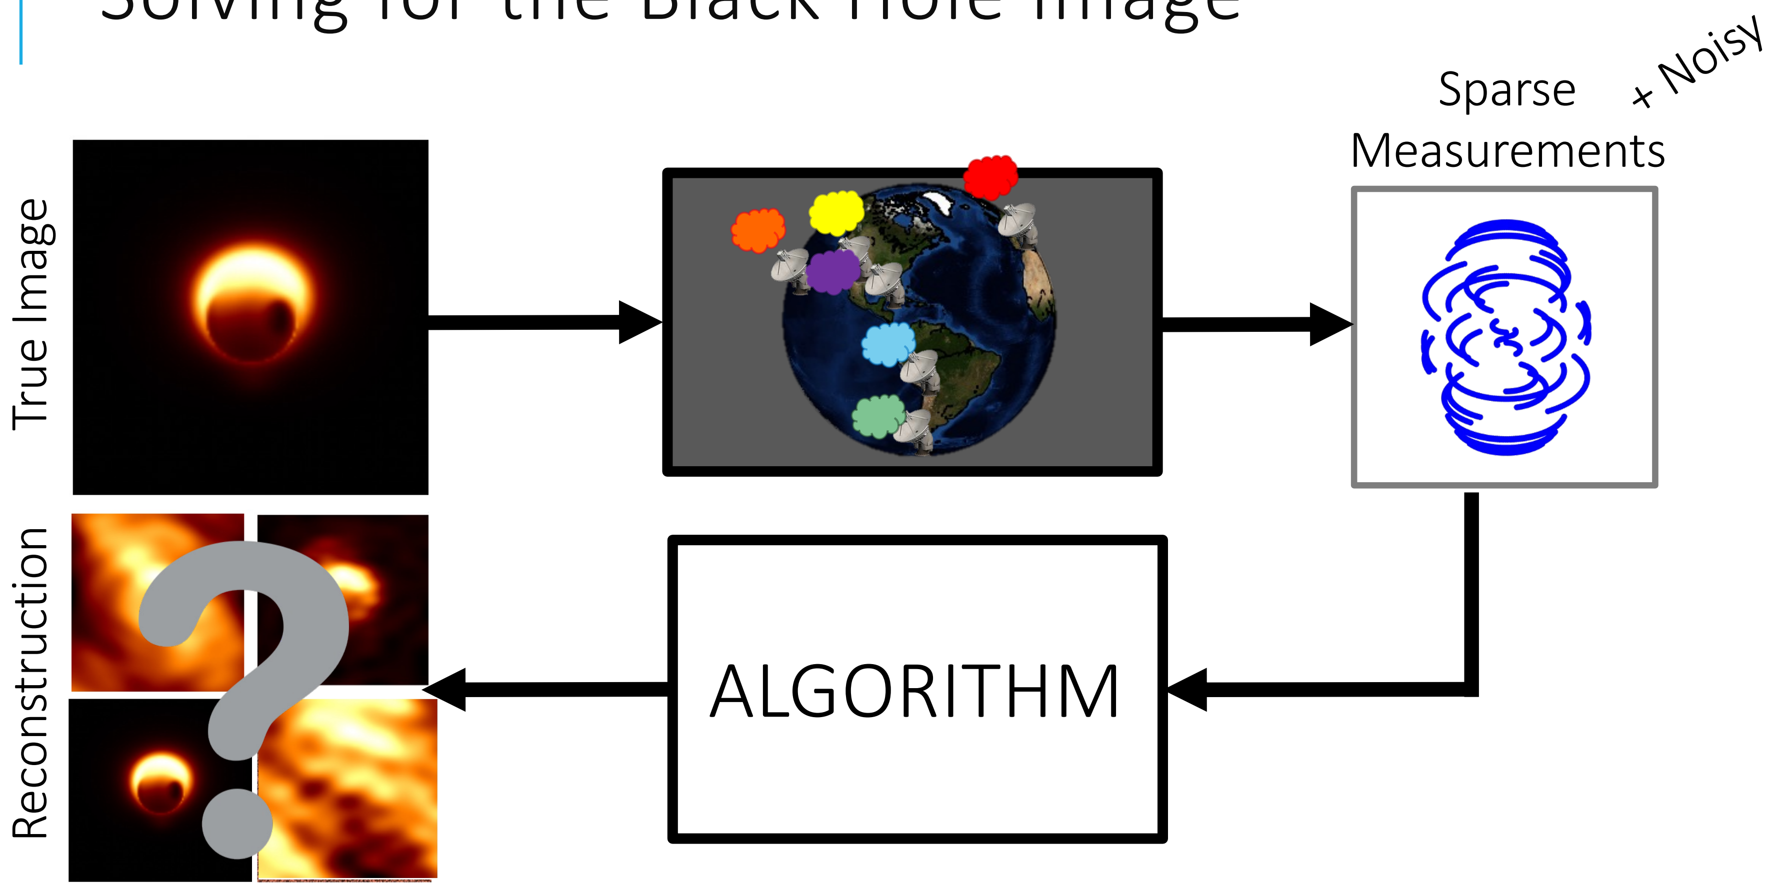

Solving for the Black Hole Image

< <

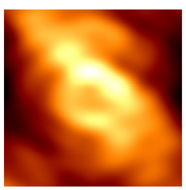

Unlikely More Likely Very Likely

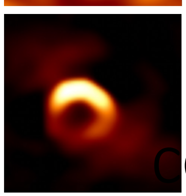

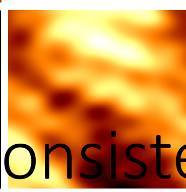

Find an image that is both :

consistent with the data & looks like an image

Solving for the Black Hole Image

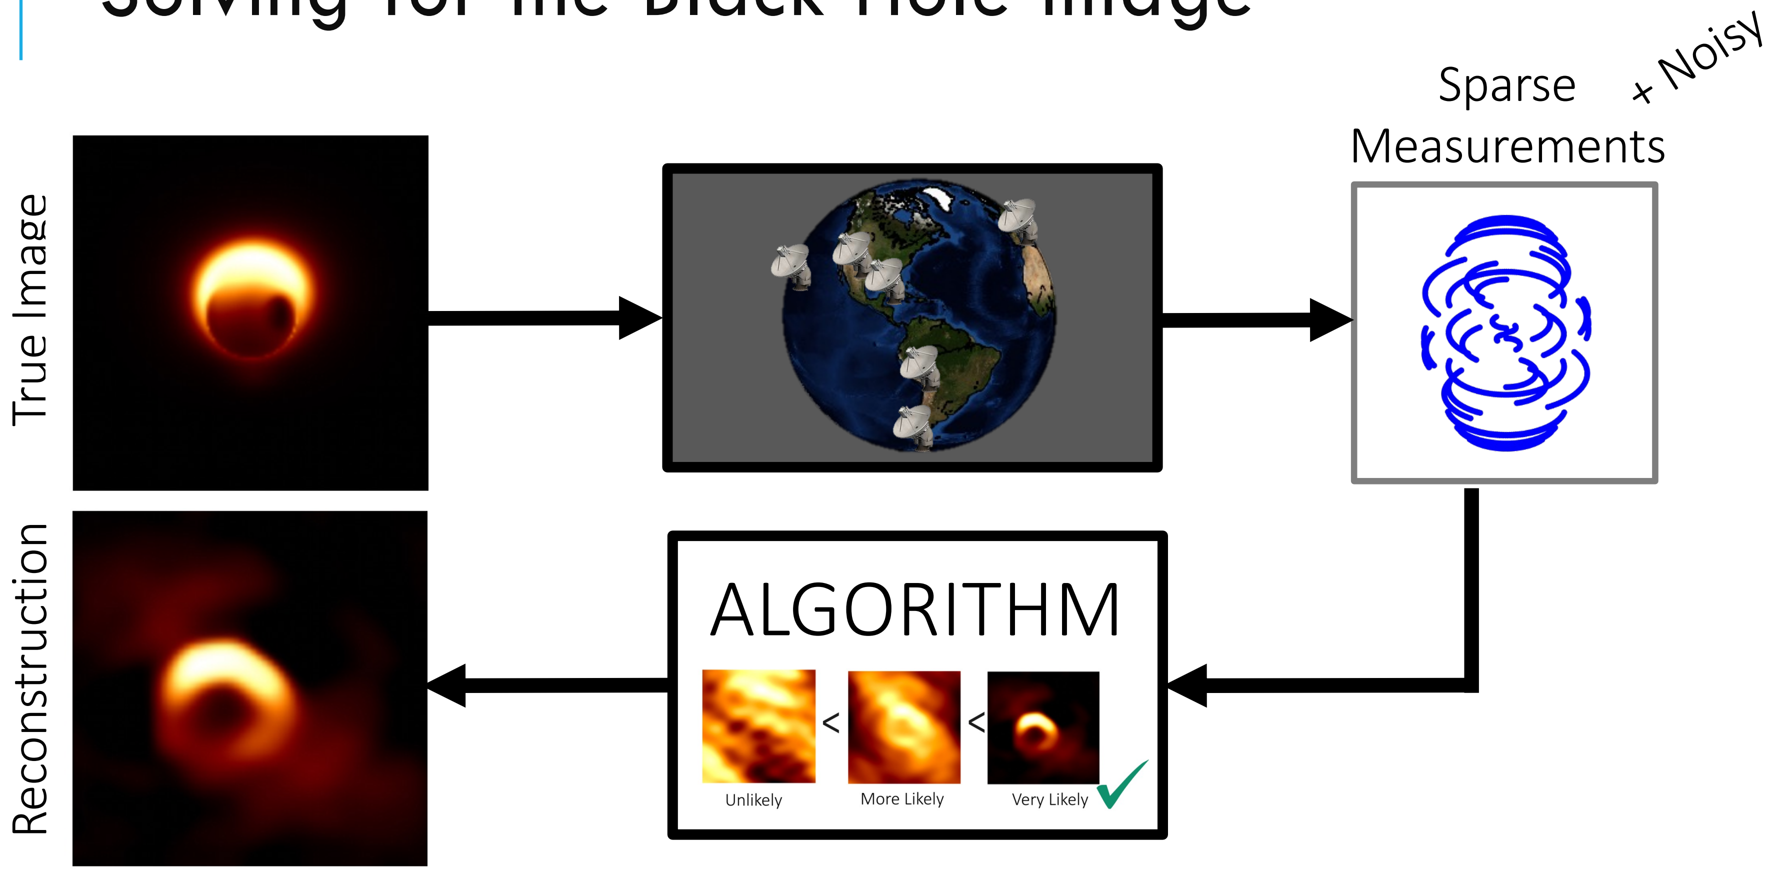

Regularized Maximum Likelihood

Best Image

image

measurements

$$\underbrace{\hat{\mathbf{x}}_{\mathrm{MAP}}}_{\text{Best Image}} = \mathrm{argmax}_{\mathbf{x}} \left[\log p(\mathbf{x}|\mathbf{y})\right]$$

Regularized Maximum Likelihood

Best Image

image

measurements

$$\hat{\mathbf{x}}_{\mathrm{MAP}} = \mathrm{argmax}_{\mathbf{x}} \left[\log p(\mathbf{x}|\mathbf{y}) \right]$$

Bayes Rule

$$= \mathrm{argmax}_{\mathbf{x}} \left[\log p(\mathbf{y}|\mathbf{x}) + \log p(\mathbf{x}) \right]$$

Likelihood

Prior

Regularized Maximum Likelihood

Best Image

image

measurements

$$\hat{\mathbf{x}}_{\mathrm{MAP}} = \mathrm{argmax}_{\mathbf{x}} \left[\log p(\mathbf{x}|\mathbf{y})\right]$$

$$= \mathrm{argmax}_{\mathbf{x}} \left[\log p(\mathbf{y}|\mathbf{x}) + \log p(\mathbf{x})\right]$$

Likelihood

Prior

Imaging Pipelines

DIFMAP
CLEAN + Self Calibration

- Systematic Error
- Scattering Prescription
- Variability Model
- Time Averaging
- ALMA Weight
- Mask Diameter
- Data Weights

eht-imaging
Regularized Max Likelihood

- Systematic Error
- Scattering Prescription
- Variability Model
- Data Weight
- Regularizes
  - MEM
  - TV
  - TSV
  - L1

SMILI
Regularized Max Likelihood

- Systematic Error
- Scattering Prescription
- Variability Model
- Regularizes
  - TV
  - TSV
  - L1

174,720 Imaging Hyper-parameters Surveyed

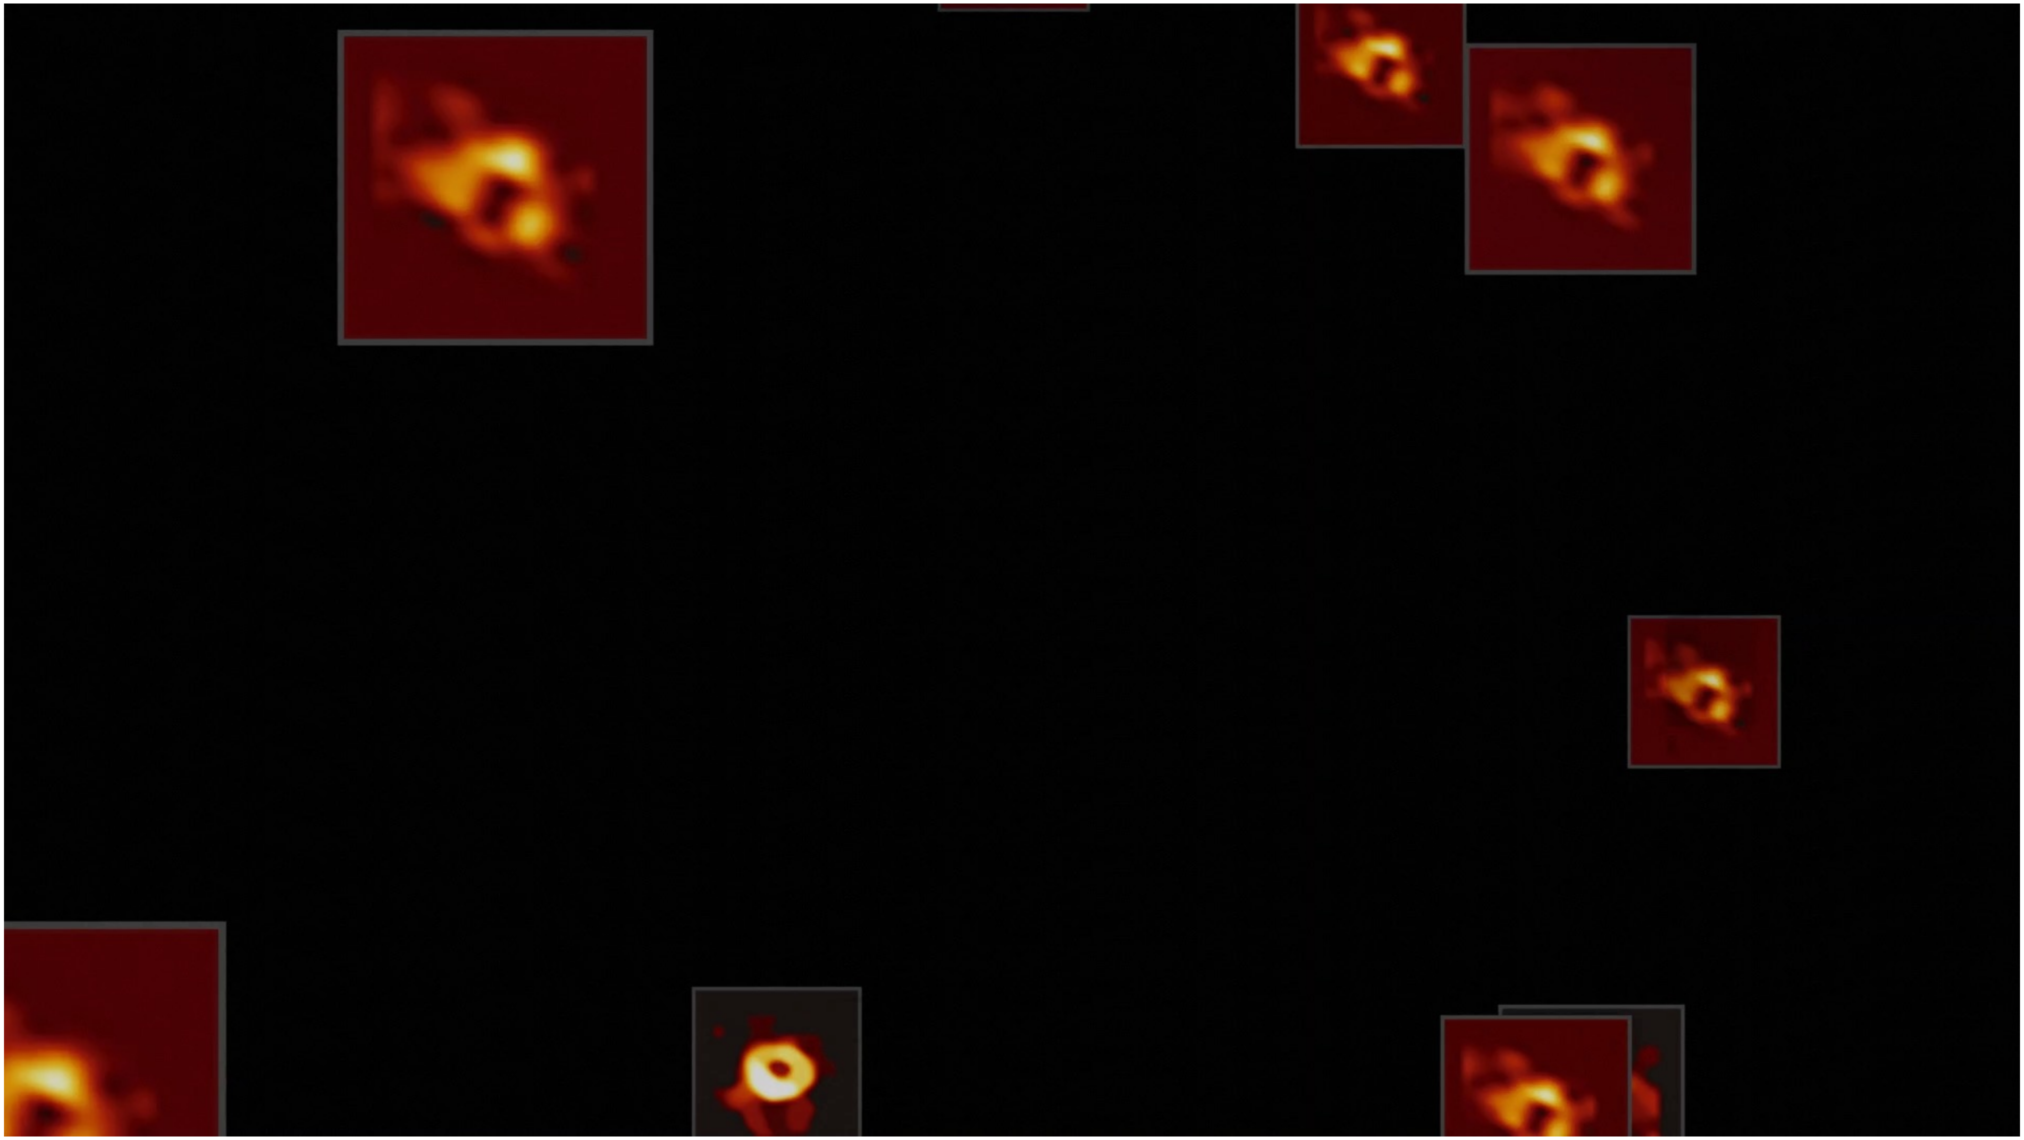

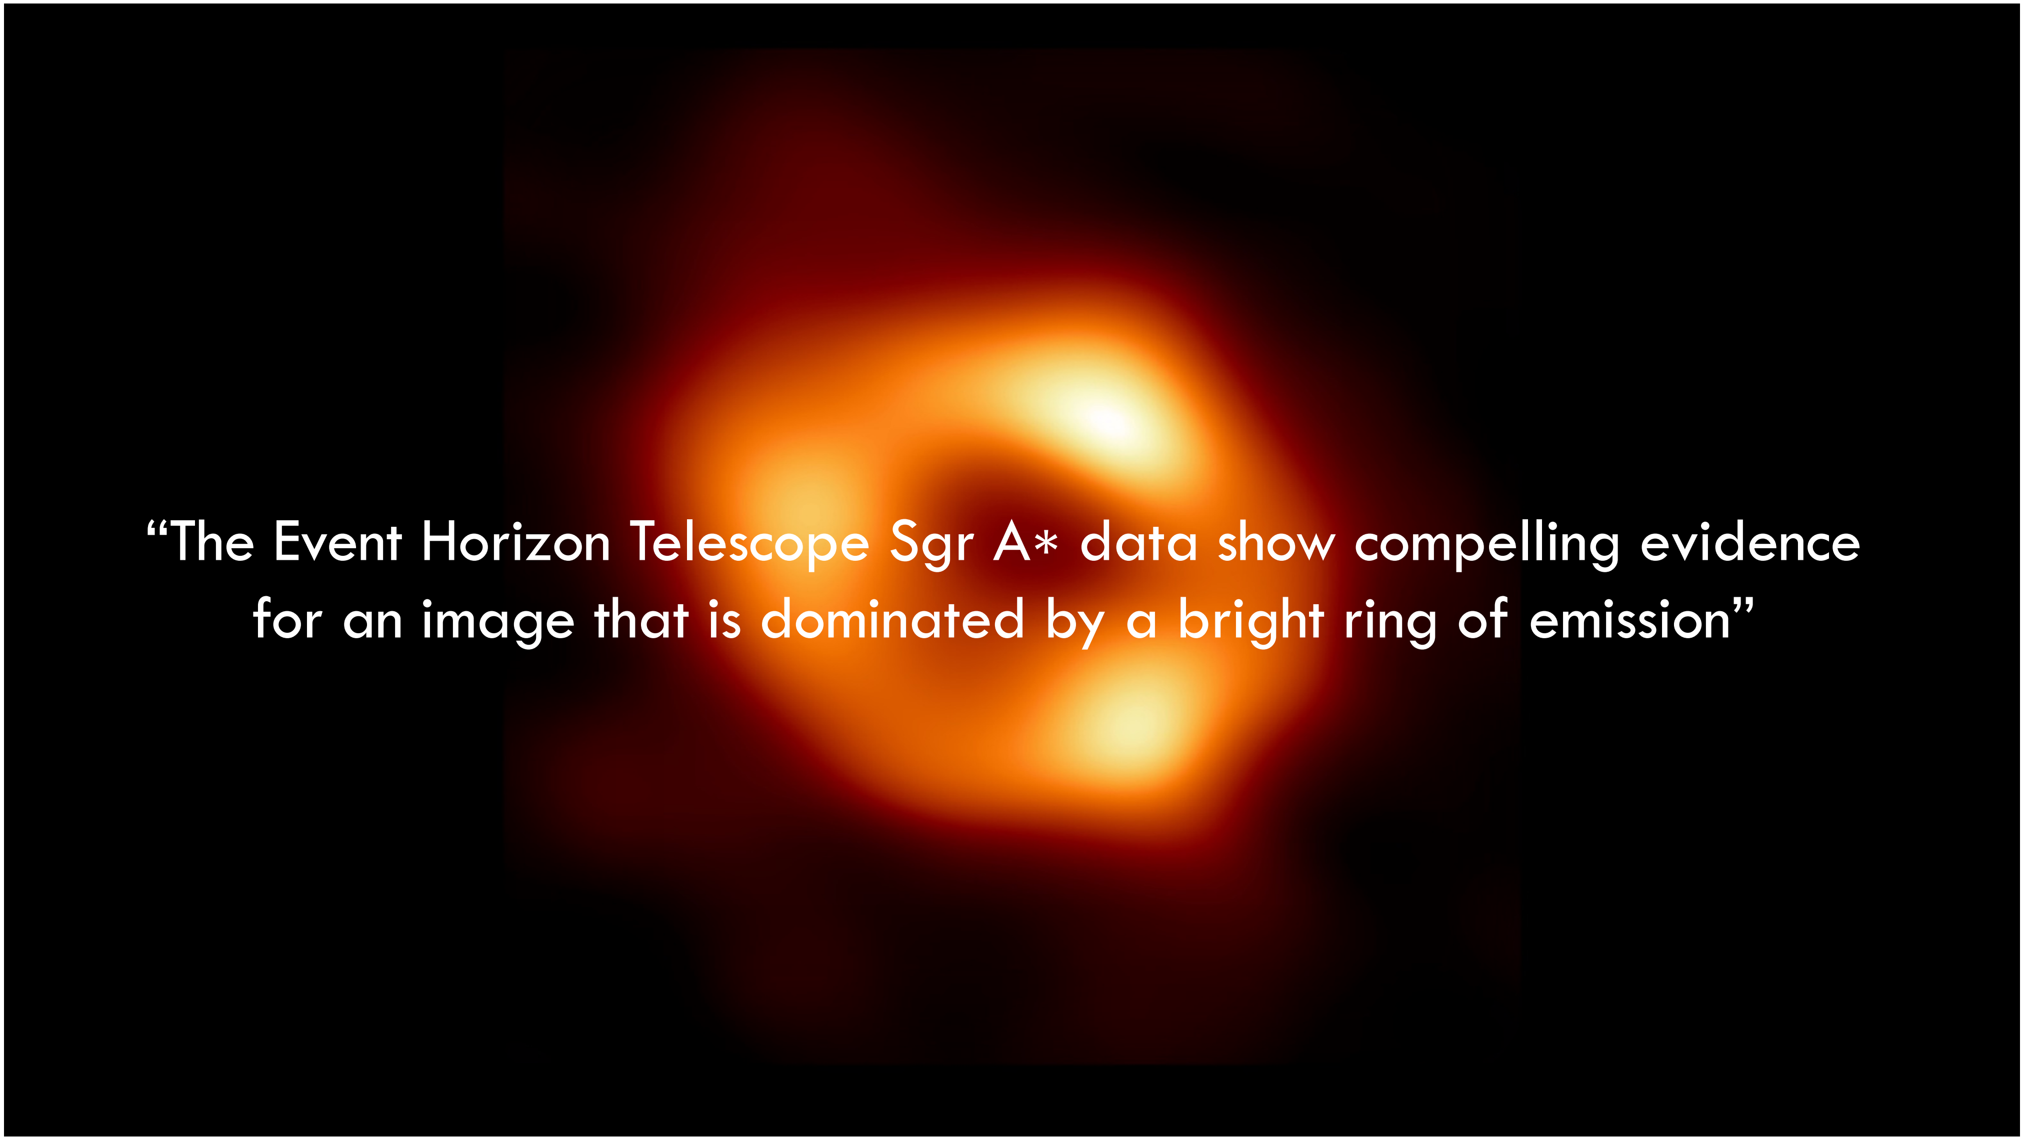
"The Event Horizon Telescope Sgr A* data show compelling evidence
for an image that is dominated by a bright ring of emission"

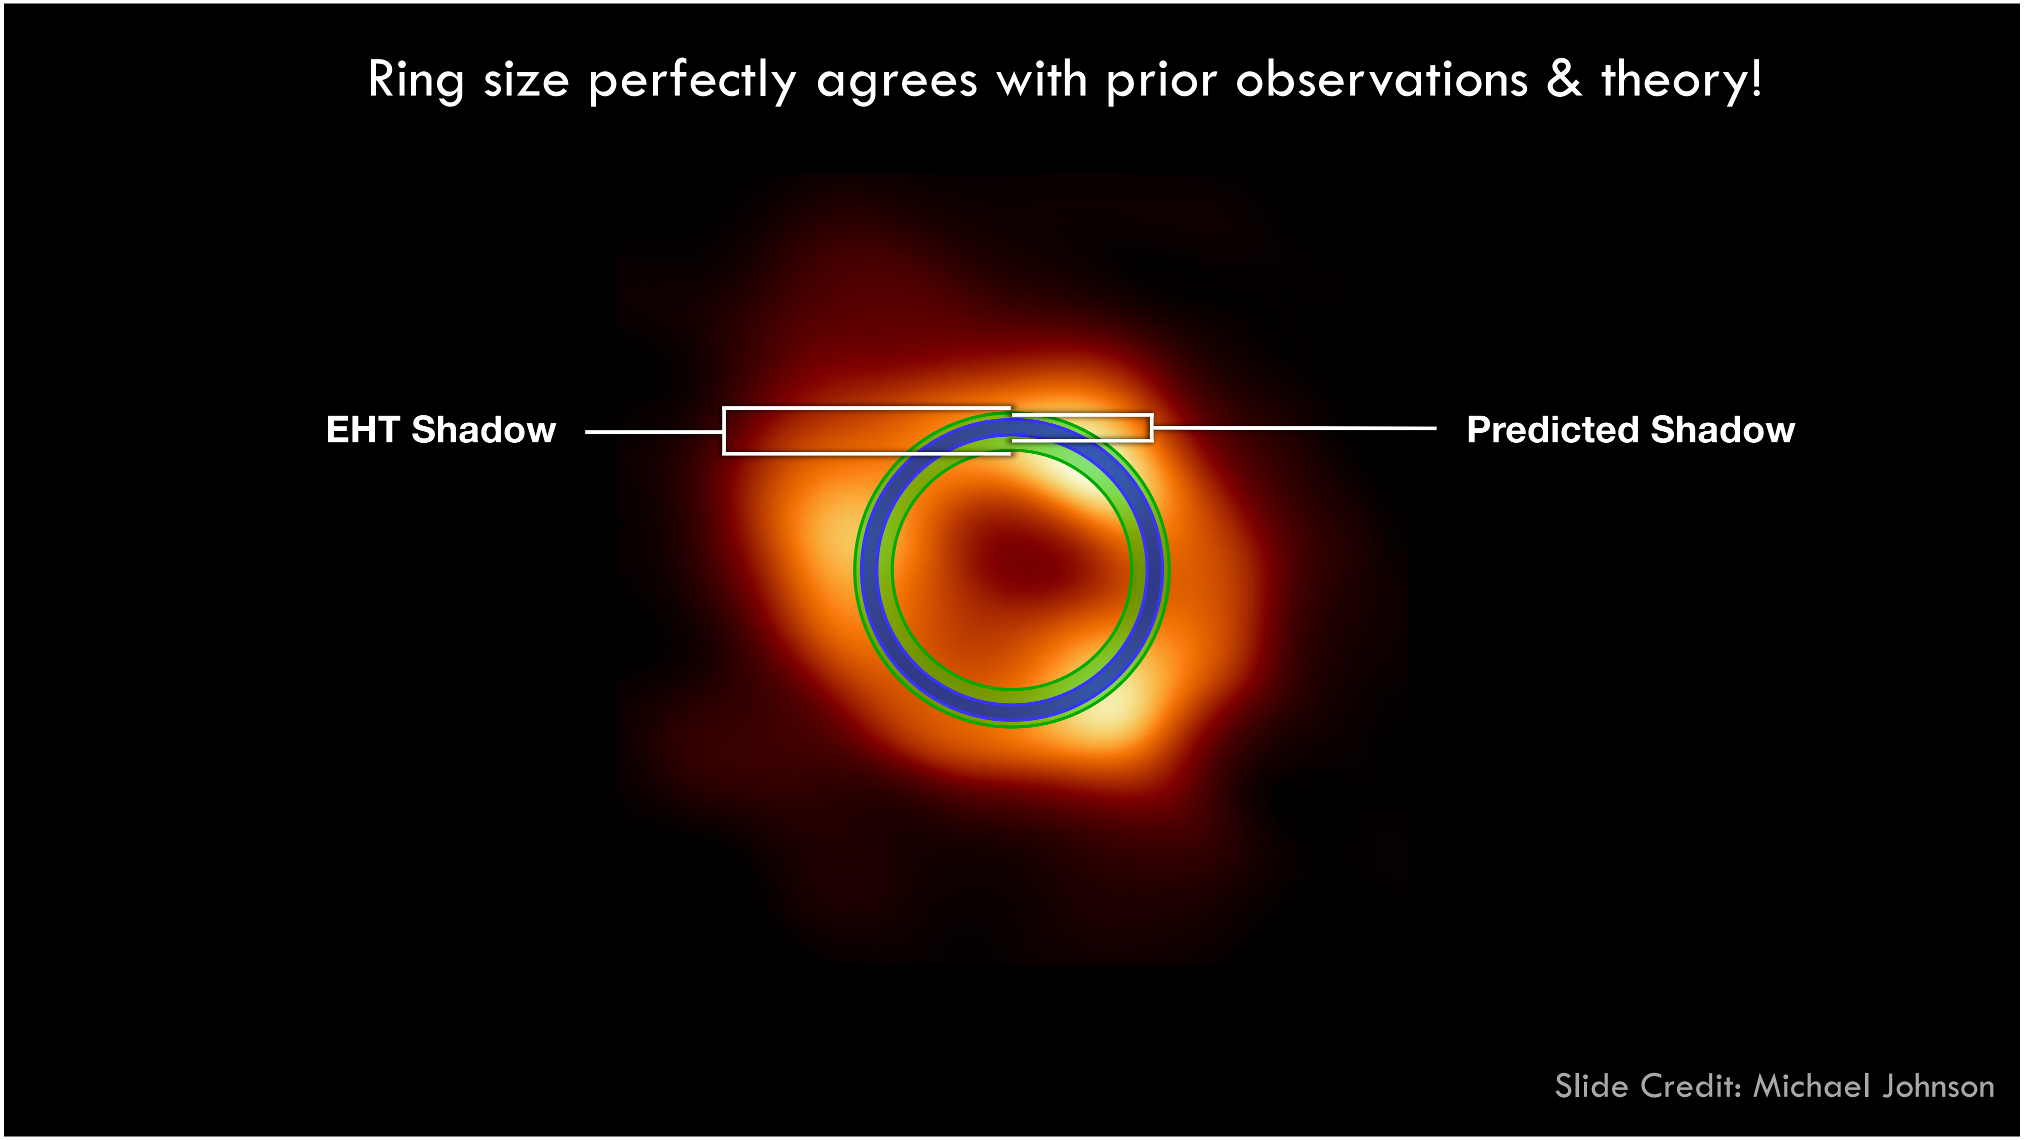
Ring size perfectly agrees with prior observations & theory!
EHT Shadow
Predicted Shadow
Slide Credit: Michael Johnson

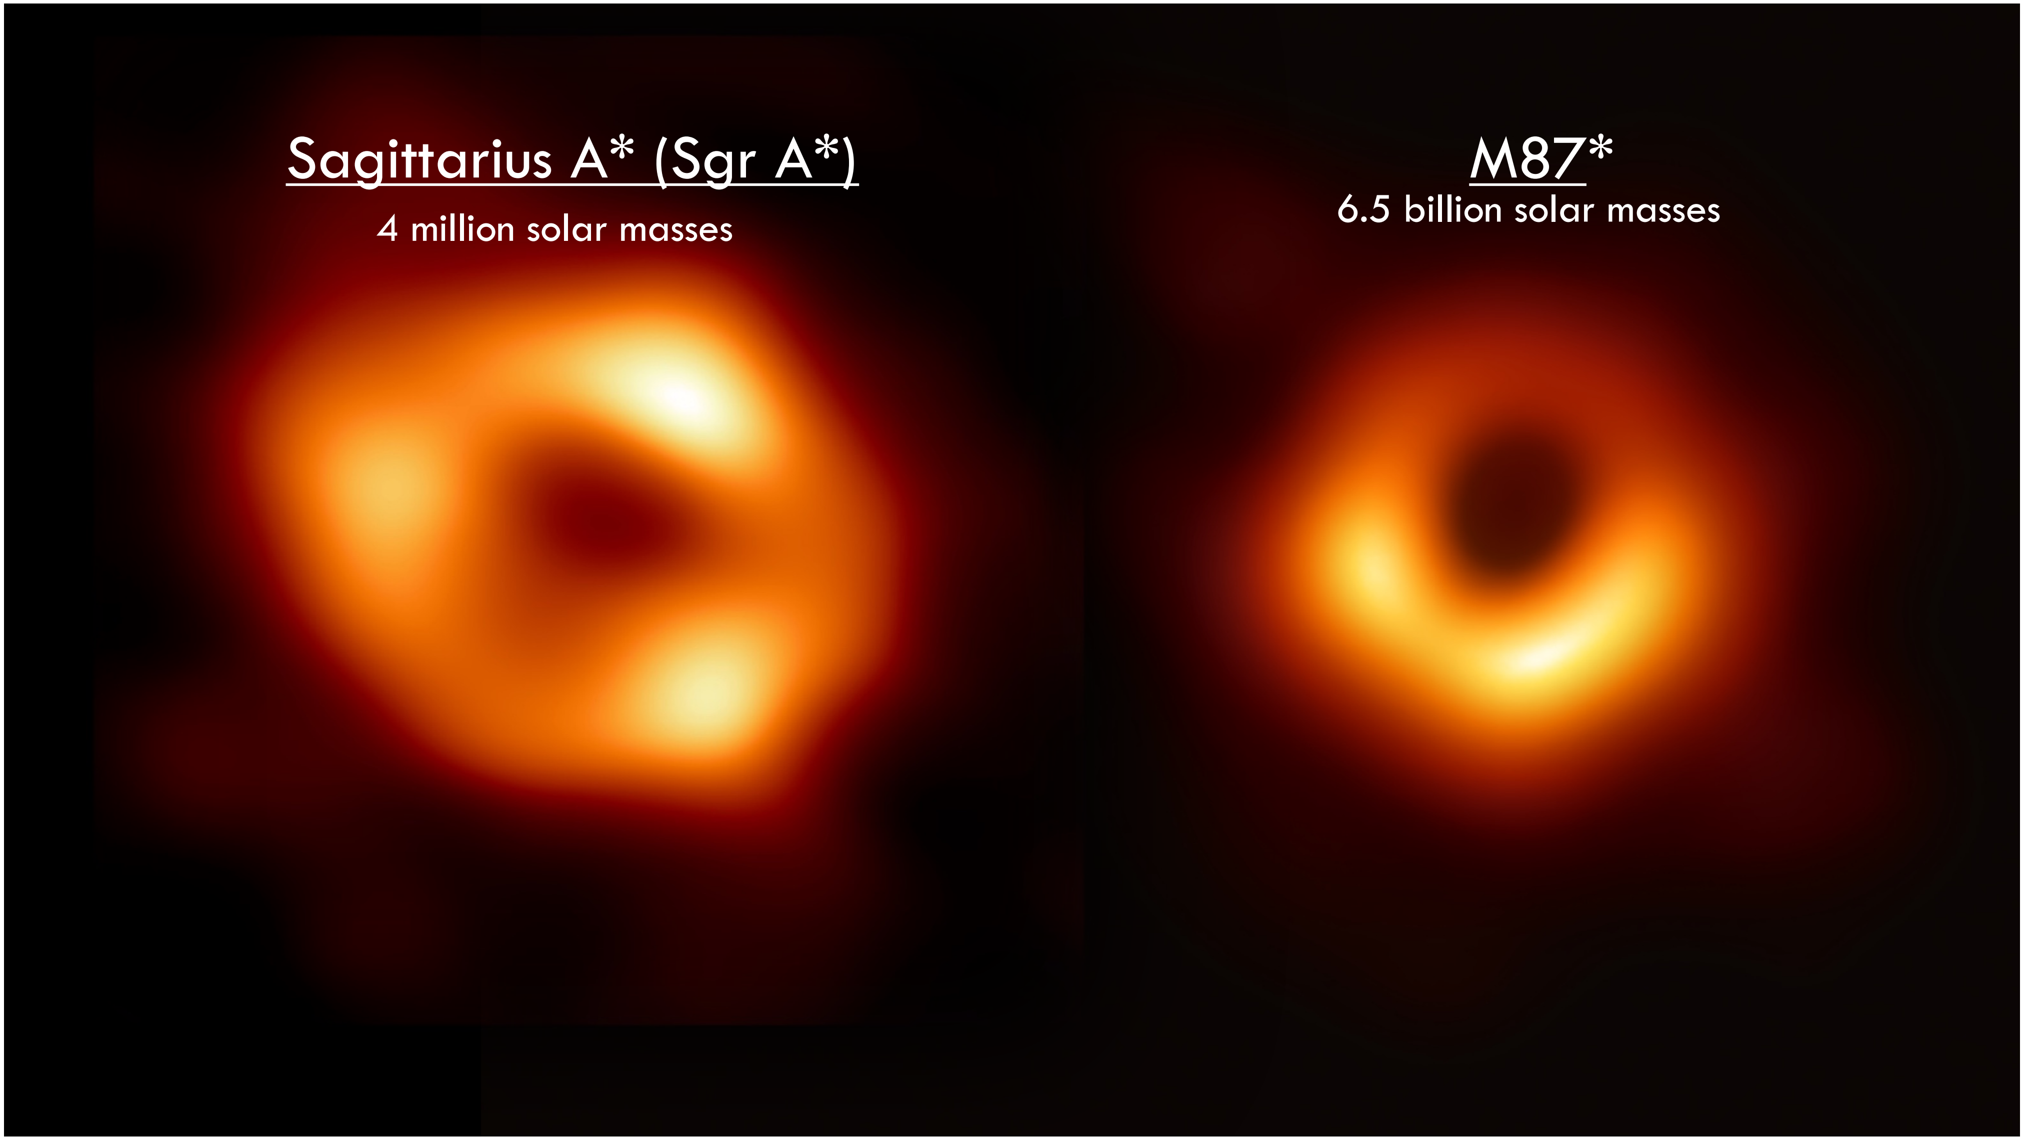
Sagittarius A* (Sgr A*)
4 million solar masses
M87*
6.5 billion solar masses


National Science Foundation
National Science Foundation
National Science Foundation
NSF
EHT


nce Foundation
NSF
MICHAEL JOHNSON
CENTER FOR ASTROPHYSICS | HARVARD & SMITHSONIAN
NSF
EHT

How to increase spatial resolution?

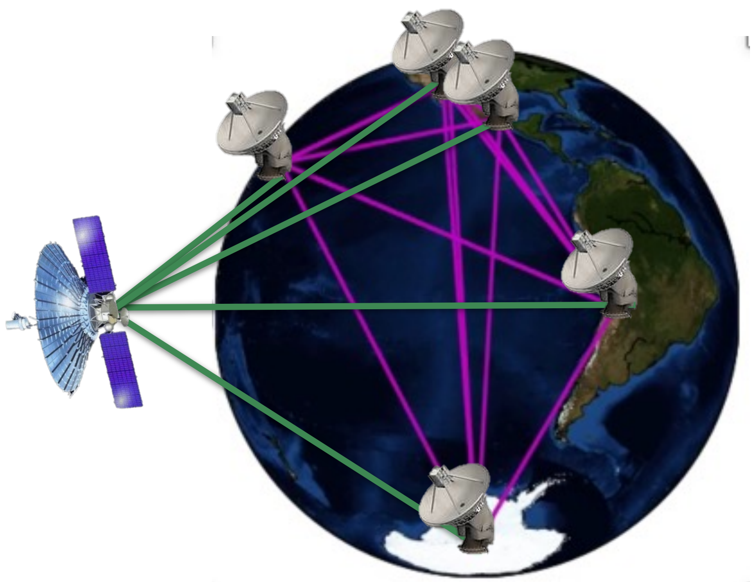

$$\uparrow \text{telescope size} \propto \frac{\text{wavelength}}{\text{angular resolution}\downarrow}$$

To increase spatial resolution (e.g., lower angular resolution) ….

….we would have to go to space

Image Reconstruction by Assuming **Weak Image Structure**

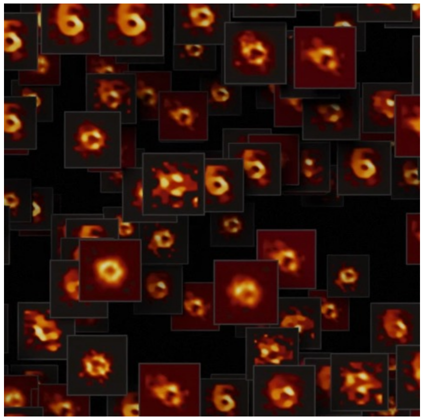

Event Horizon Telescope Collaboration, 2022

Increasingly Strong Assumptions

Image Reconstruction by Assuming **Weak Image Structure**

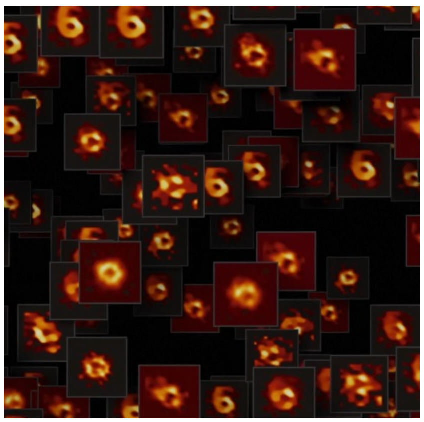

Event Horizon Telescope Collaboration, 2022

Image Reconstruction by Assuming **Data Driven Priors**

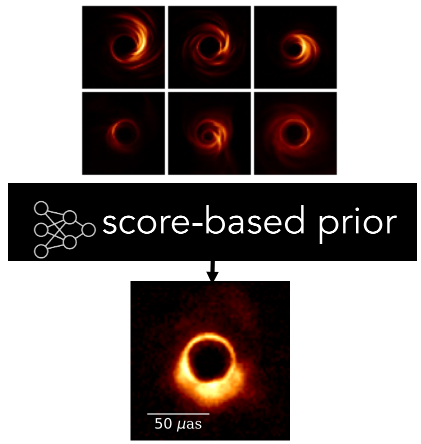

Feng, et al, ICCV, 2023
Feng, et al, ApJ, 2023 (in submission)
Wu, et al, 2024 (in submission)

Evolving Volume Reconstruction by Assuming **General Relativity**

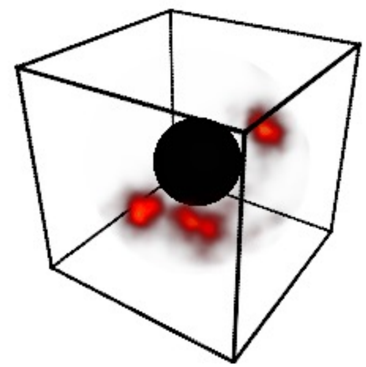

Levis*, Srinivasan*, et al, CVPR, 2022
Levis, et al, Nature Astronomy, 2024

Increasingly Strong Assumptions

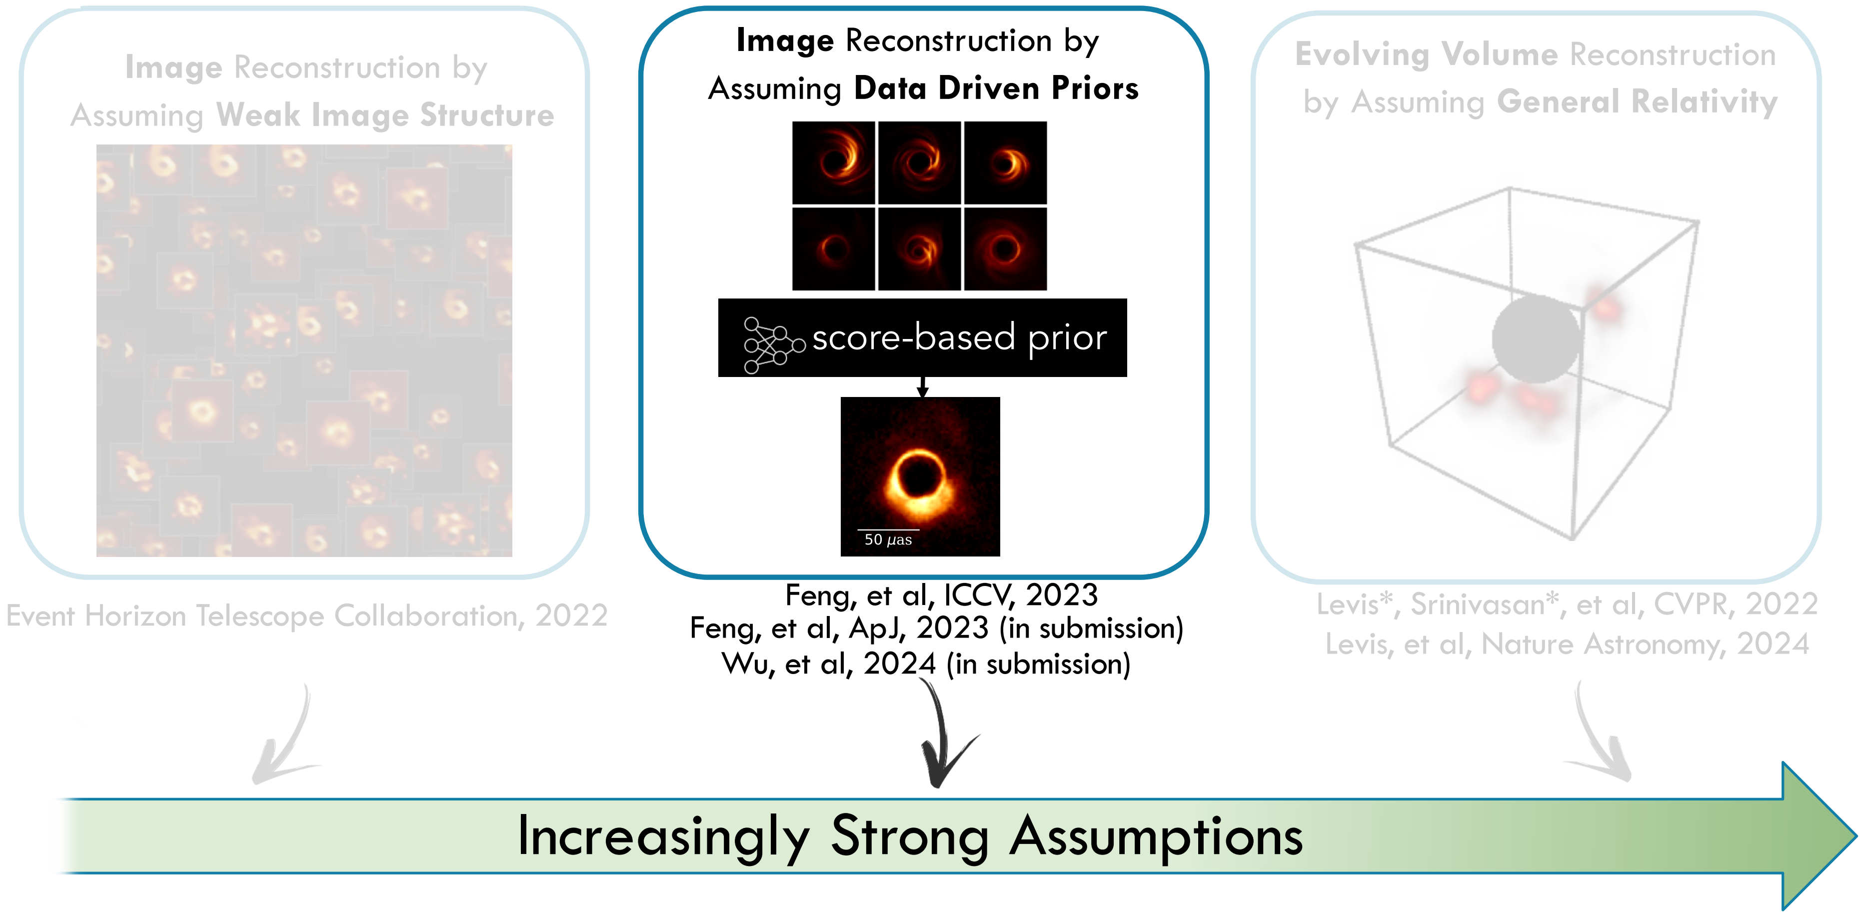
Image Reconstruction by
Assuming Weak Image Structure
Event Horizon Telescope Collaboration, 2022
Image Reconstruction by
Assuming Data Driven Priors
score-based prior
50 μas
Feng, et al, ICCV, 2023
Feng, et al, ApJ, 2023 (in submission)
Wu, et al, 2024 (in submission)
Evolving Volume Reconstruction
by Assuming General Relativity
Levis*, Srinivasan*, et al, CVPR, 2022
Levis, et al, Nature Astronomy, 2024
Increasingly Strong Assumptions

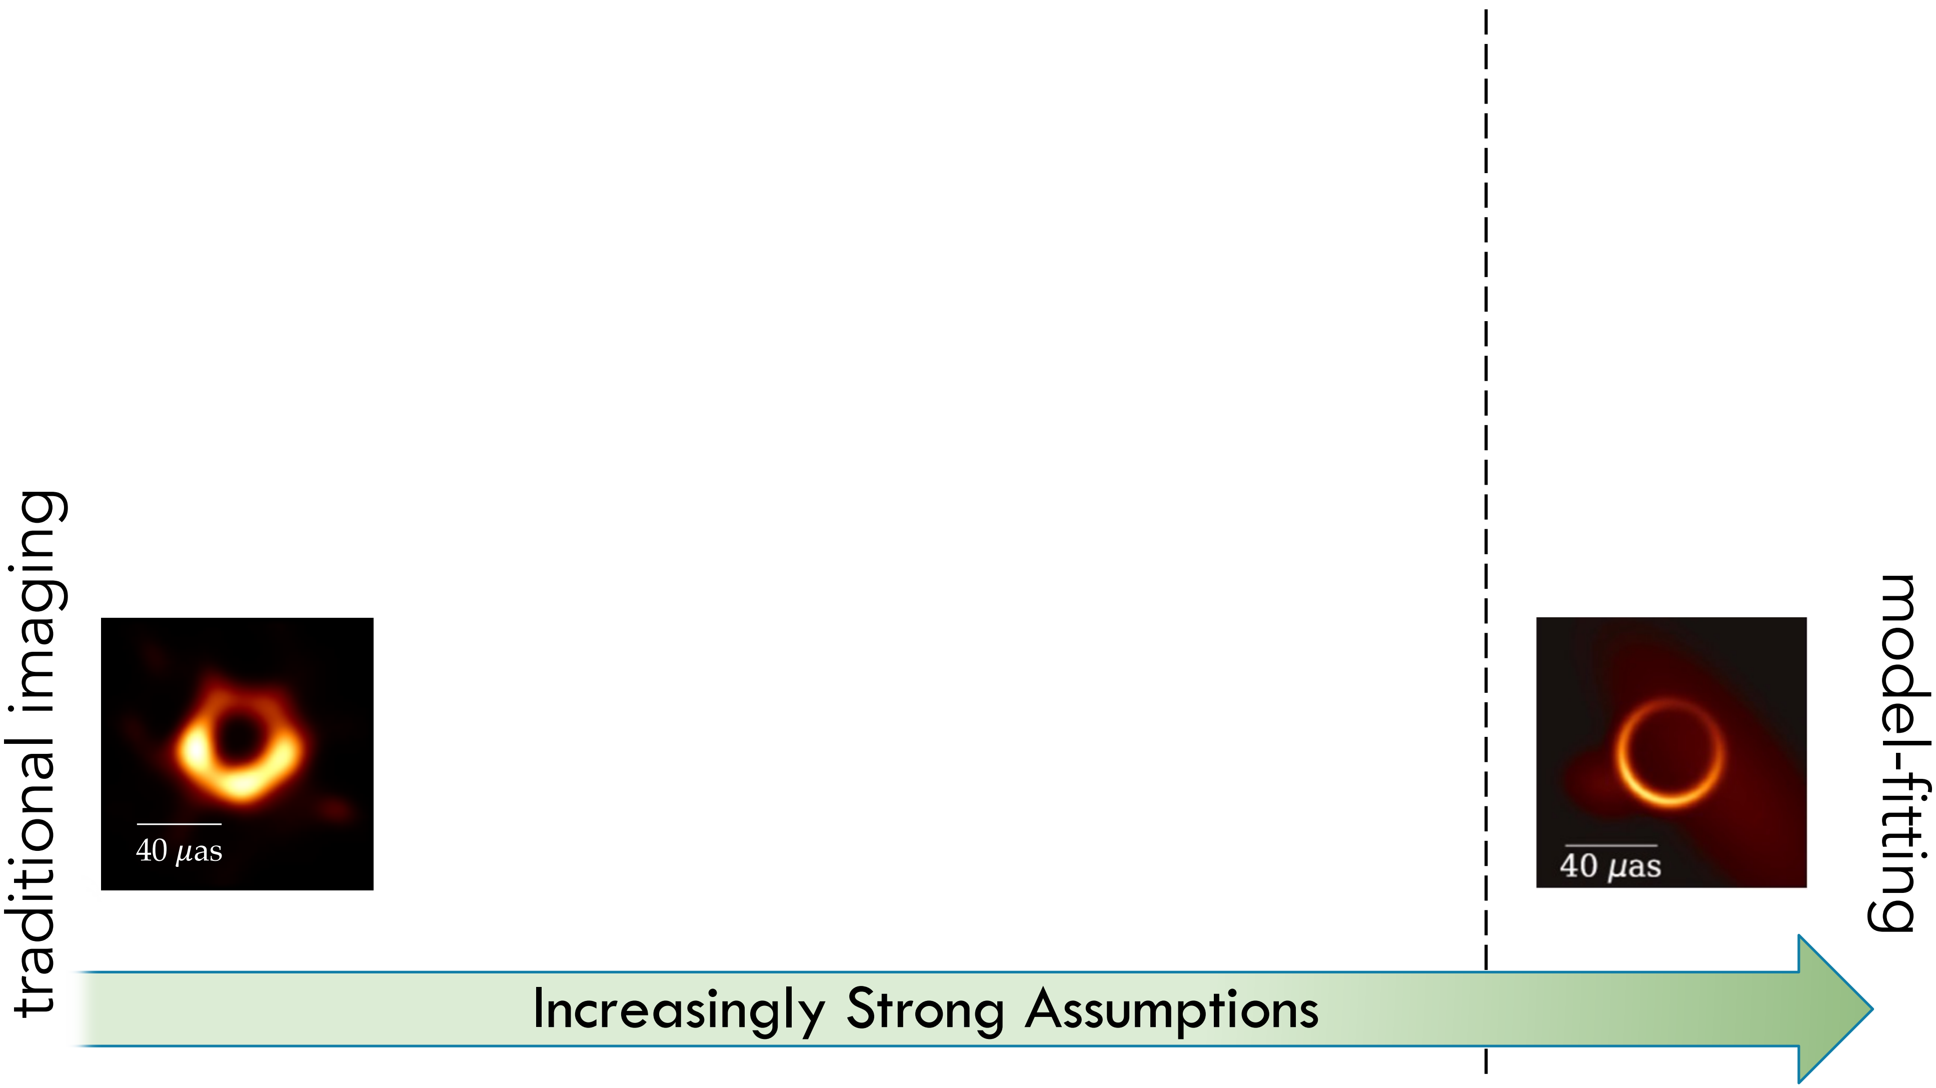
traditional imaging
40 μas
40 μas
model-fitting
Increasingly Strong Assumptions

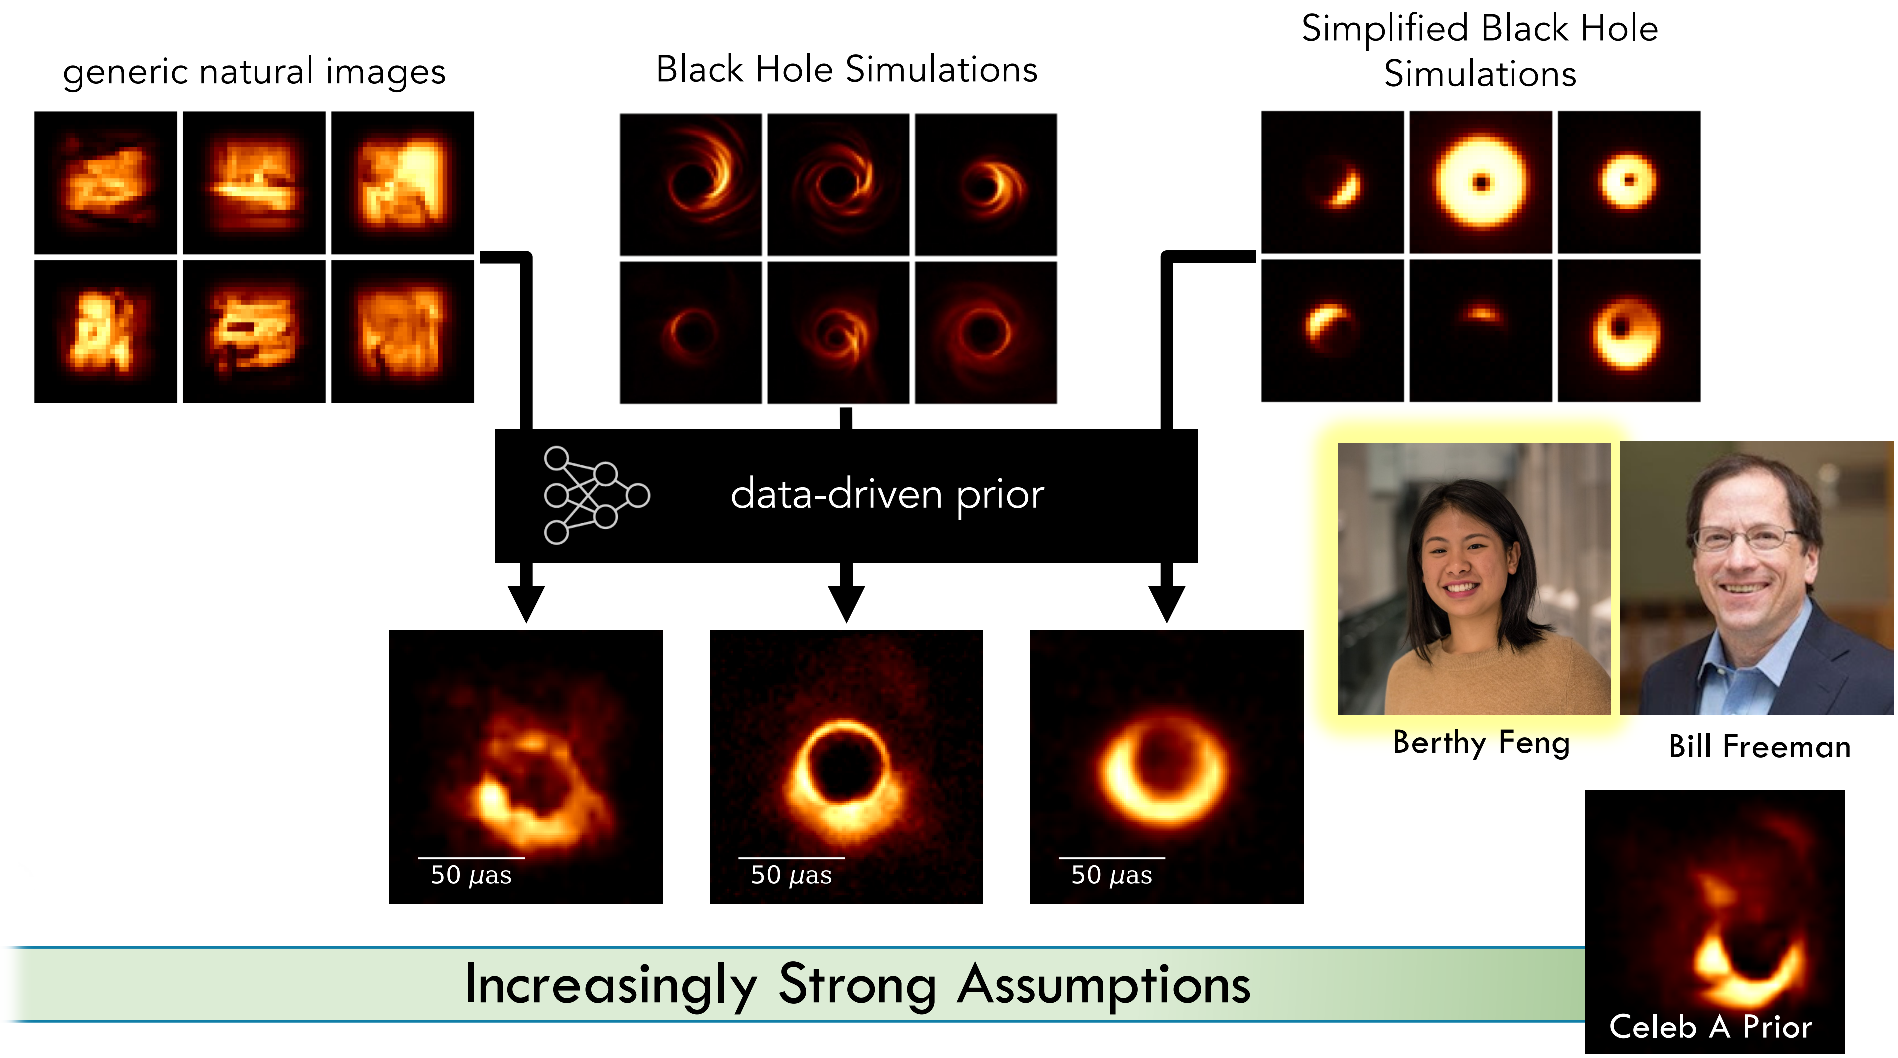
generic natural images
Black Hole Simulations
Simplified Black Hole Simulations
data-driven prior
50 μas
50 μas
50 μas
Berthy Feng
Bill Freeman
Celeb A Prior
Increasingly Strong Assumptions

Diffusion Model

Forward Noising Process: $dx_t = f(t)x_t + g(t)dw$

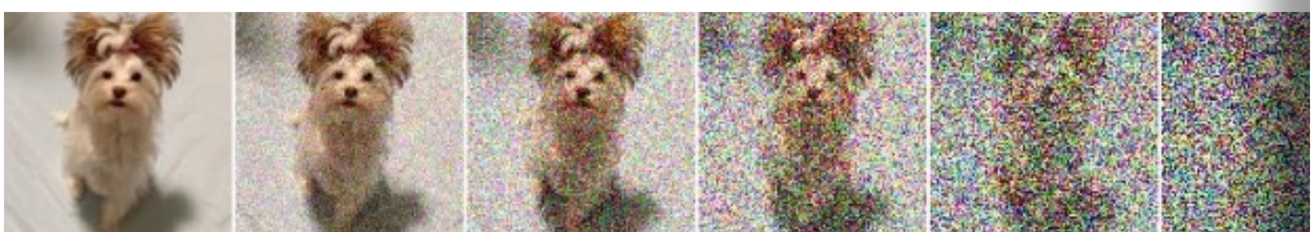

x_0 x_1 x_2 x_t x_{T-1} x_T

Diffusion Posterior Sampling for General Noisy Inverse Problems

Hyungjin Chung[*1,2], **Jeongsol Kim**[*1], **Michael T. Mccann**[2], **Marc L. Klasky**[2] & **Jong Chul Ye**[1]
[1]KAIST, [2] Los Alamos National Laboratory
{hj.chung, jeongsol, jong.ye}@kaist.ac.kr, {mccann, mklasky}@lanl.gov

Abstract

Diffusion models have been recently studied as powerful generative inverse problem solvers, owing to their high quality reconstructions and the ease of combining existing iterative solvers. However, most works focus on solving simple linear inverse problems in noiseless settings, which significantly under-represents the complexity of real-world problems. In this work, we extend diffusion solvers to efficiently handle general noisy (non)linear inverse problems via approximation of the posterior sampling. Interestingly, the resulting posterior sampling scheme is a blended version of diffusion sampling with the manifold constrained gradient without a strict measurement consistency projection step, yielding a more desirable generative path in *noisy* settings compared to the previous studies. Our method demonstrates that diffusion models can incorporate various measurement noise statistics such as Gaussian and Poisson, and also efficiently handle noisy *nonlinear* inverse problems such as Fourier phase retrieval and non-uniform deblurring. Code is available at `https://github.com/DPS2022/diffusion-posterior-sampling`.

1 Introduction

Diffusion models learn the implicit prior of the underlying data distribution by matching the gradient of the log density (i.e. Stein score; $\nabla_{\boldsymbol{x}} \log p(\boldsymbol{x})$) (Song et al., 2021b). The prior can be leveraged when solving inverse problems, which aim to recover $\boldsymbol{x}$ from the measurement $\boldsymbol{y}$, related through the forward measurement operator $\mathcal{A}$ and the detector noise $\boldsymbol{n}$. When we know such forward models, one can incorporate the gradient of the log likelihood (i.e. $\nabla_{\boldsymbol{x}} \log p(\boldsymbol{y}|\boldsymbol{x})$) in order to sample from the posterior distribution $p(\boldsymbol{x}|\boldsymbol{y})$. While this looks straightforward, the likelihood term is in fact analytically intractable in terms of diffusion models, due to their dependence on time t. Due to its intractability, one often resorts to projections onto the measurement subspace (Song et al., 2021b; Chung et al., 2022b; Chung & Ye, 2022; Choi et al., 2021). However, the projection-type approach fails dramatically when 1) there is noise in the measurement, since the noise is typically amplified during the generative process due to the ill-posedness of the inverse problems; and 2) the measurement process is nonlinear.

One line of works that aim to solve noisy inverse problems run the diffusion in the spectral domain (Kawar et al., 2021; 2022) so that they can tie the noise in the measurement domain into the spectral domain via singular value decomposition (SVD). Nonetheless, the computation of SVD is costly and even prohibitive when the forward model gets more complex. For example, Kawar et al. (2022) only considered *seperable* Gaussian kernels for deblurring, since they were restricted to the family of inverse problems where they could effectively perform the SVD. Hence, the applicability of such methods is restricted, and it would be useful to devise a method to solve noisy inverse problems *without* the computation of SVD. Furthermore, while diffusion models were applied to various inverse problems including inpainting (Kadkhodaie & Simoncelli, 2021; Song et al., 2021b; Chung et al., 2022b; Kawar et al., 2022; Chung et al., 2022a), super-resolution (Kadkhodaie & Simoncelli, 2021; Choi et al., 2021; Chung et al., 2022b; Kawar et al., 2022), colorization (Song et al., 2021b; Kawar et al., 2022; Chung et al., 2022a), compressed-sensing MRI (CS-MRI) (Song et al., 2022; Chung & Ye, 2022; Chung et al., 2022b), computed tomography (CT) (Song et al., 2022; Chung et al., 2022a), etc., to our best knowledge, all works so far considered linear inverse problems only, and have not explored *nonlinear* inverse problems.

*Joint first authors

Reverse Denoising Process:

$$dx_t = [f(t)x_t + g(t)^2 \nabla \log p_t(x_t)] + g(t)dw_t$$

Conditional Diffusion Models

Unconditional reverse diffusion

$$dx_t = [f(t)x_t + g(t)^2 \nabla \log p_t(x_t)] + g(t)dw_t$$

Conditional reverse diffusion

$$dx_t = [f(t)x_t + g(t)^2 \nabla \log p_t(x_t|y)] + g(t)dw_t$$

↓ Bayes rule

$$dx_t = [f(t)x_t + g(t)^2 \underbrace{\nabla \log p_t(x_t)}_{\text{Unconditional score}} + g(t)^2 \underbrace{\nabla \log p_t(y|x_t)}_{\text{Likelihood at time t}}] + g(t)dw_t$$

Unconditional score — **Pre-trained diffusion models**

Likelihood at time t — **Intractable in general**

Plug-and-Play Diffusion Models (PnP-DM)

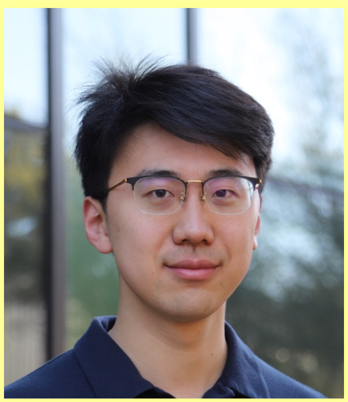

Zihui (Ray) Wu

Yu Sun

Yifan Chen

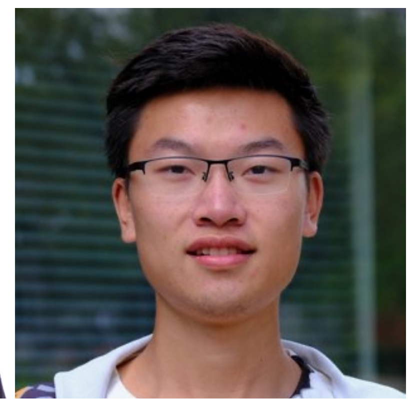

Bingliang Zhang

Yisong Yue

Sample the Bayesian Posterior

$$p(x|y) \propto p(y|x)\ p(x)$$

image measurements

Sample the Bayesian Posterior

$$p(x|y) \propto \quad p(y|x) \quad p(x)$$

$$= \exp(\log p(y|x)) \exp(\log p(x))$$

Combining the exponents

$$= \exp(\log p(y|x) + \log p(x))$$

Split Gibbs Sampler (SGS) [Vono, et al, 2019]

$$p(x|y) \propto \quad p(y|x) \quad p(x)$$

$$= \exp(\log p(y|x)) \exp(\log p(x))$$

$$= \exp(\log p(y|x) + \log p(x))$$

Introduce z

$$= \exp(\log p(y|z) + \log p(x) - \frac{1}{2\rho^2}|x - z|_2^2) \quad \text{as} \quad \rho \to 0$$

Split Gibbs Sampler (SGS) [Vono, et al, 2019]

$$p(x|y) \propto \exp(\log p(y|z) + \log p(x) - \frac{1}{2\,\rho^2}|x - z|_2^2) \quad \text{as} \quad \rho \rightarrow 0$$

Alternate Between 2 Steps:

Likelihood Step: fix x , sample z

Prior Step: fix z , sample x

Split Gibbs Sampler (SGS) [Vono, et al, 2019]

constant

$$p(x|y) \propto \exp\left(\log p(y|z) + \log p(x) - \frac{1}{2\rho^2} |x - z|_2^2 \right) \quad \text{as} \quad \rho \to 0$$

Alternate Between 2 Steps:

Likelihood Step: fix x , sample z

Prior Step: fix z , sample x

Split Gibbs Sampler (SGS) : the Prior Step

constant

$$p(x|y) \propto \exp(\log p(y|z) + \log p(x) - \frac{1}{2\,\rho^2}|x - z|_2^2) \quad \text{as} \quad \rho \to 0$$

Alternate Between 2 Steps:

Likelihood Step: fix x , sample z

Prior Step: fix z , sample x

Split Gibbs Sampler (SGS) : the Prior Step

$$\exp\left(\underbrace{\log p(x)}_{\text{prior}} - \underbrace{\frac{1}{2\,\rho^2} |x - z|_2^2}_{\text{denoising measurement likelihood}} \right)$$

Prior Step: fix z , sample x

Equivalent to sampling the posterior in a denoising problem with measurement z and noise standard deviation of ρ!

EDM Diffusion Model Rigorously Solves Prior Step

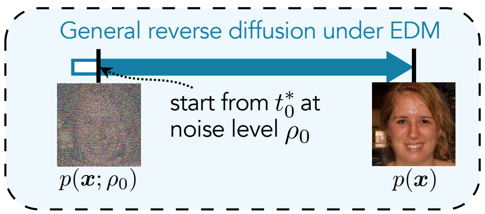

Large ρ → nearly image generation

Observation

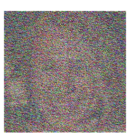

Denoising posterior samples

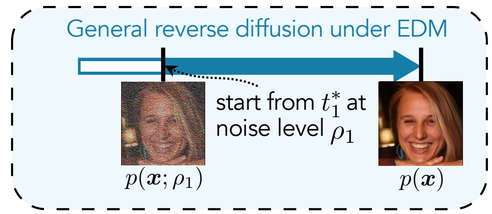

Small ρ → image denoising

Observation

Denoising posterior samples

EDM Diffusion Model Rigorously Solves Prior Step

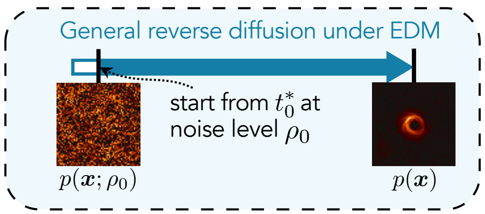

Large ρ → nearly image generation

Observation

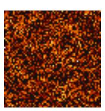

Denoising posterior samples

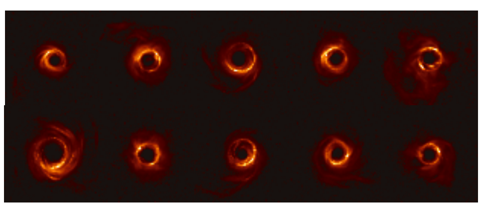

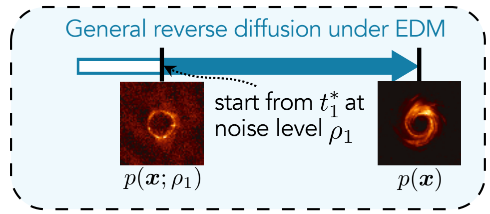

Small ρ → image denoising

Observation

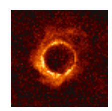

Denoising posterior samples

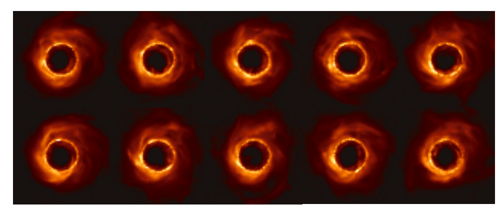

Plug-and-Play Diffusion Model (PnP-DM)

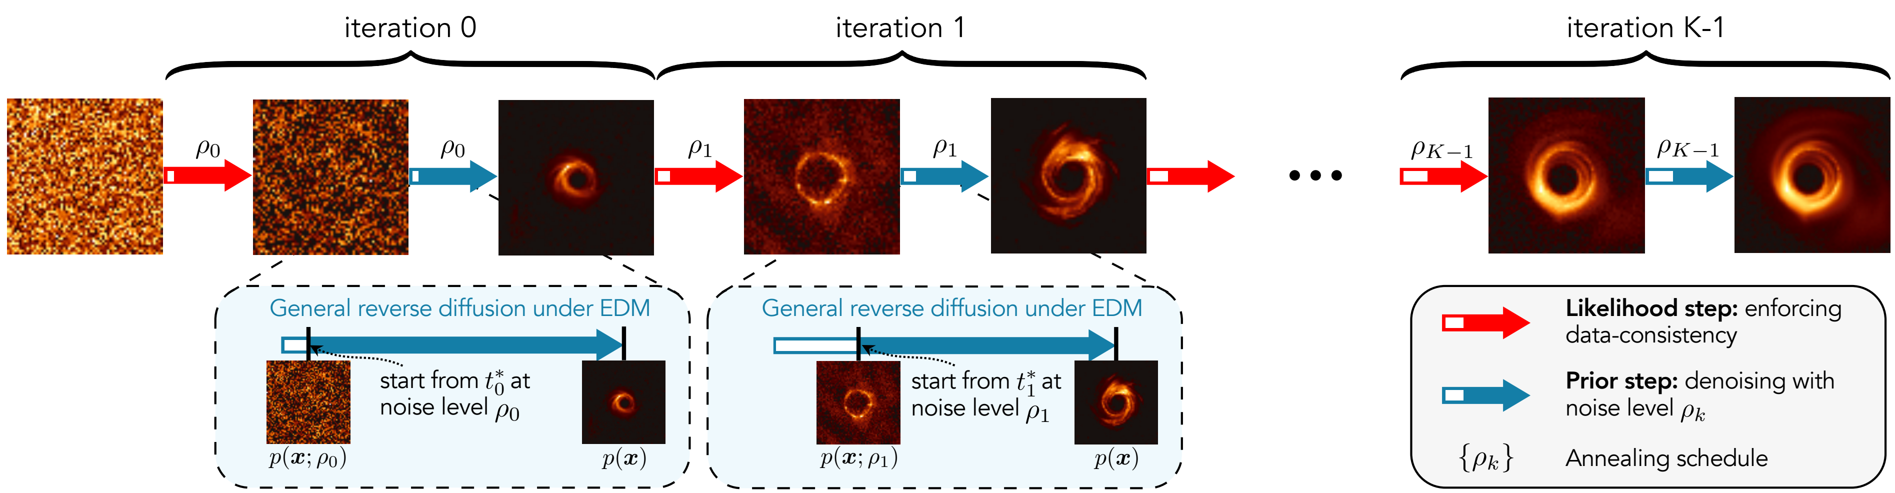

Real Data Reconstruction using Black Hole Prior

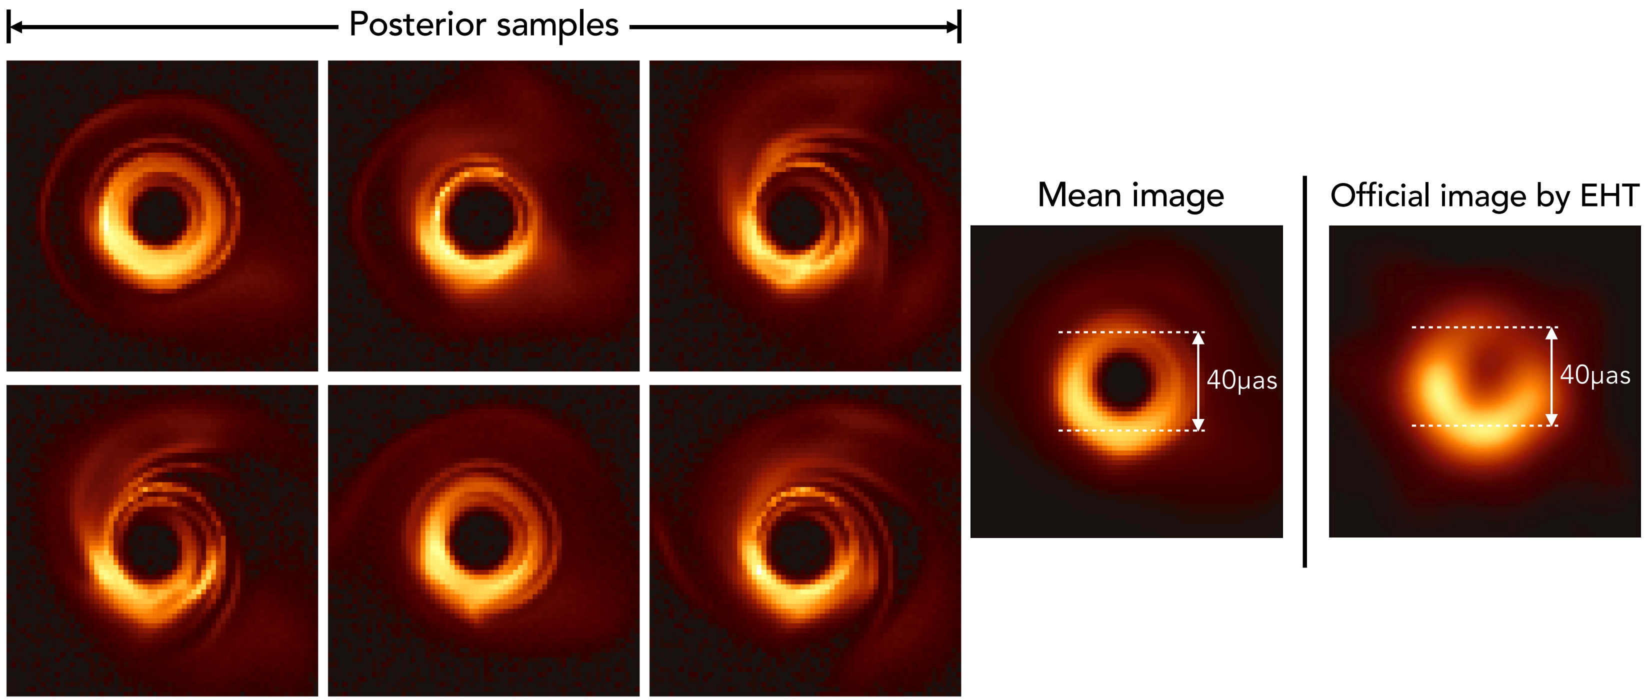

Experiment is performed with real data for the M87 black hole with non-convex constraints

Image Reconstruction by Assuming **Weak Image Structure**

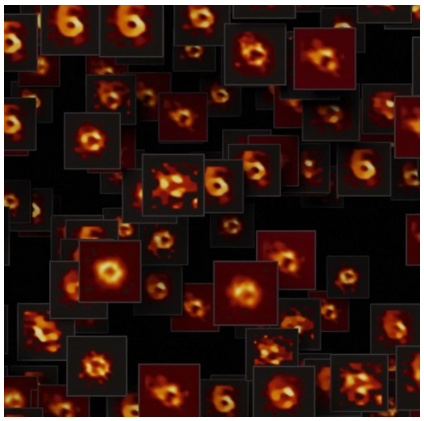

Event Horizon Telescope Collaboration, 2022

Image Reconstruction by Assuming **Data Driven Priors**

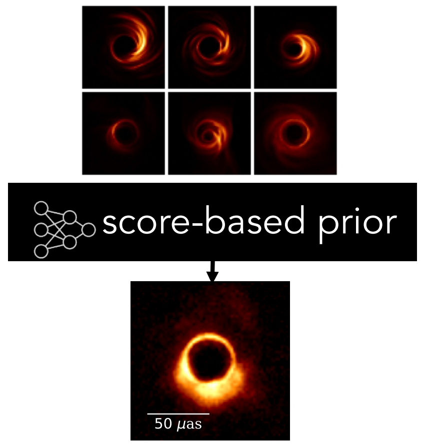

Feng, et al, ICCV, 2023
Feng, et al, ApJ, 2023 (in submission)
Wu, et al, 2024 (in submission)

Evolving Volume Reconstruction by Assuming **General Relativity**

Levis*, Srinivasan*, et al, CVPR, 2022
Levis, et al, Nature Astronomy, 2024

Increasingly Strong Assumptions

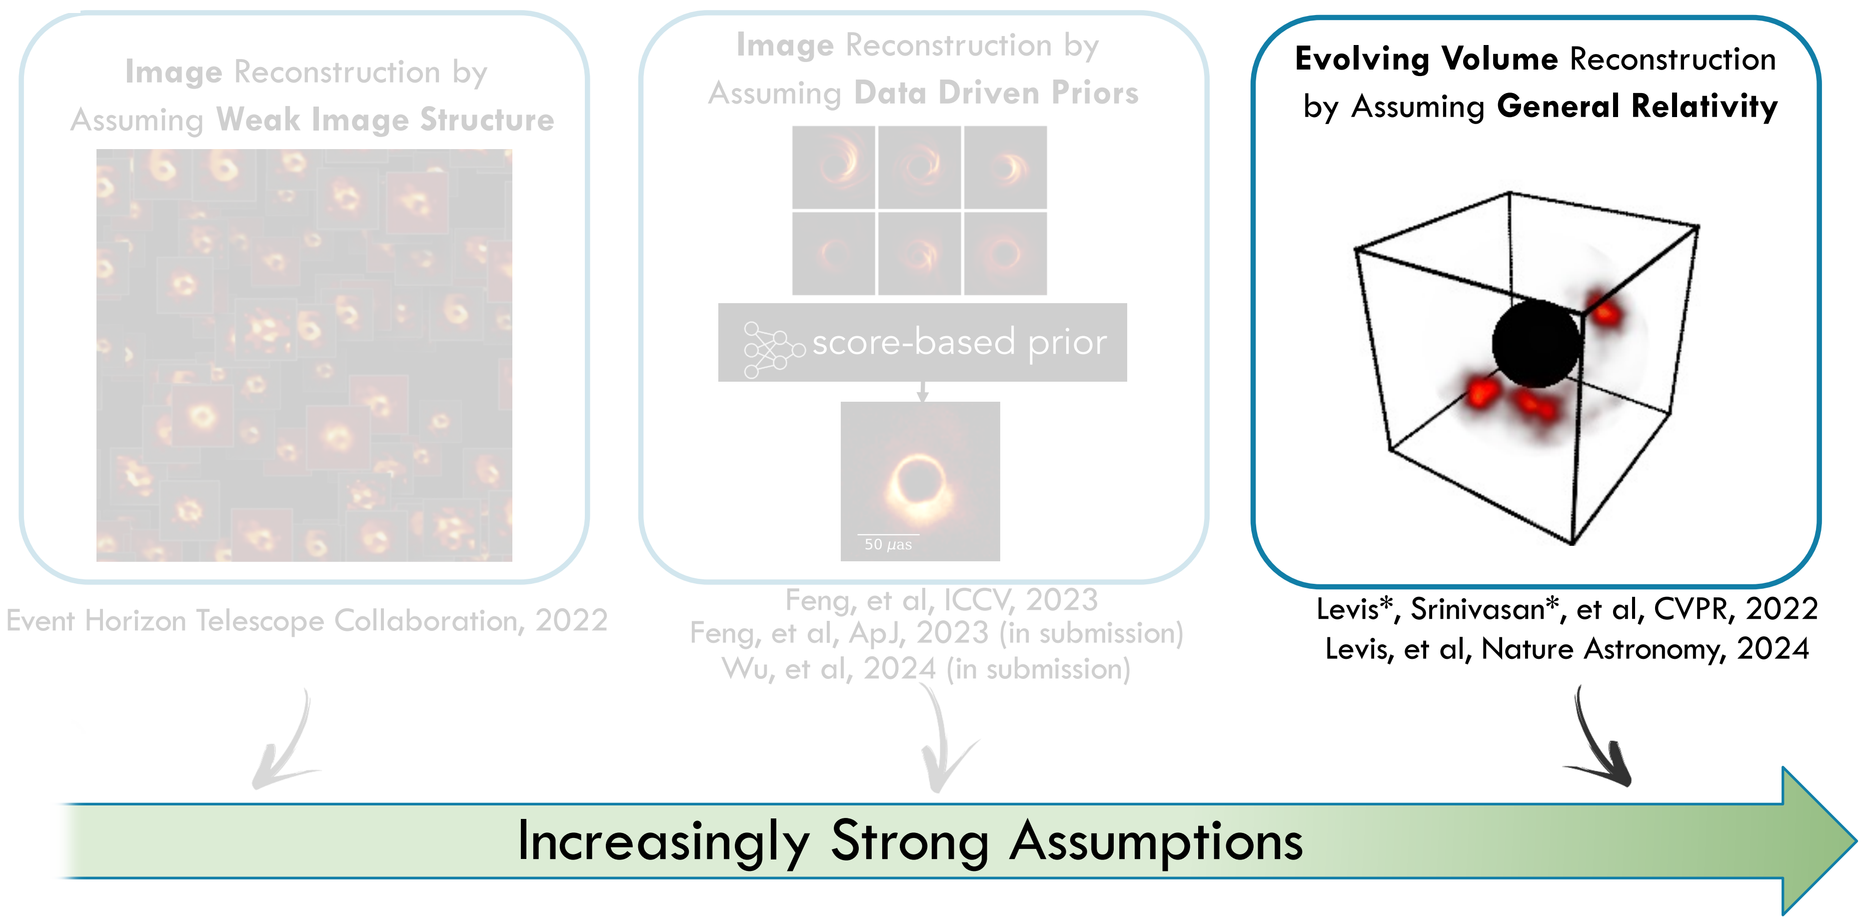
Image Reconstruction by
Assuming Weak Image Structure
Event Horizon Telescope Collaboration, 2022
Image Reconstruction by
Assuming Data Driven Priors
score-based prior
50 μas
Feng, et al, ICCV, 2023
Feng, et al, ApJ, 2023 (in submission)
Wu, et al, 2024 (in submission)
Evolving Volume Reconstruction
by Assuming General Relativity
Levis*, Srinivasan*, et al, CVPR, 2022
Levis, et al, Nature Astronomy, 2024
Increasingly Strong Assumptions

Traditional vs Black Hole Tomography

Computed Tomography (CT)

<u>Challenge 1</u>

Curved Rays

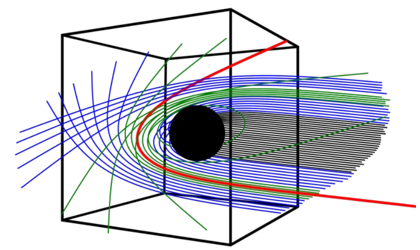

<u>Challenge 2</u>

Single View

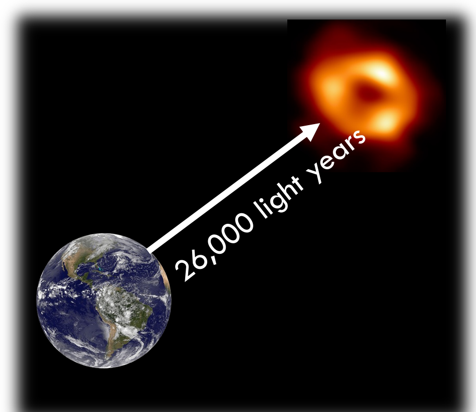

Traditional vs Black Hole Tomography

Computed Tomography (CT)

<u>Challenge 1</u>

Curved Rays

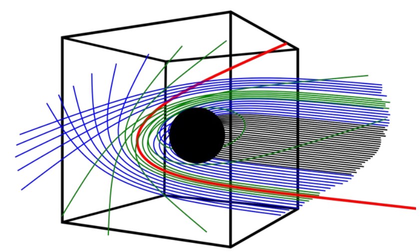

<u>Challenge 2</u>

Single View

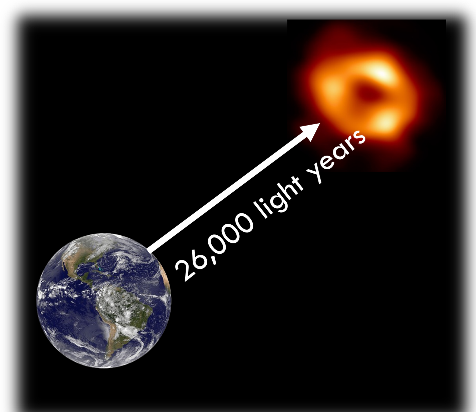

Gravitational Lensing Black Hole Emission Tomography

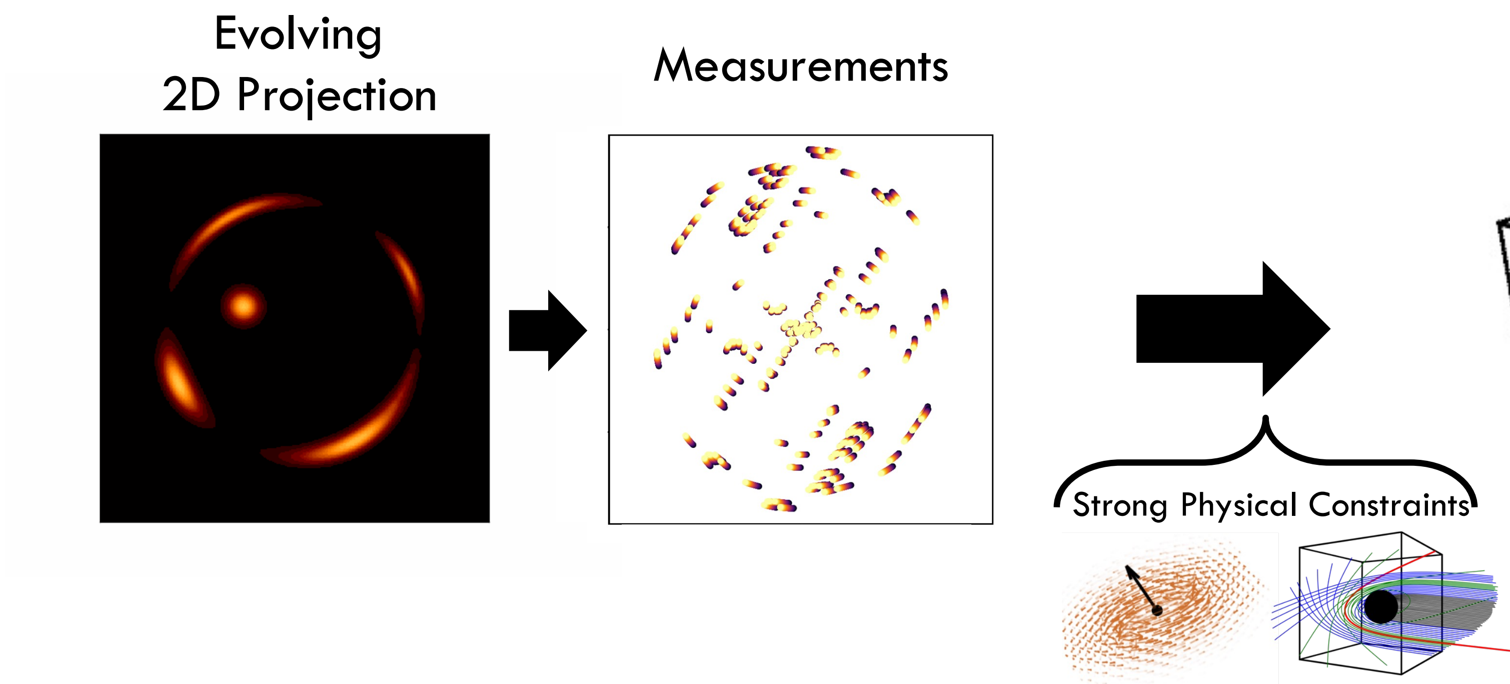

Gravitational Lensing Black Hole Emission Tomography

Aviad Levis

Pratul Srinivasan

Andrew Chael

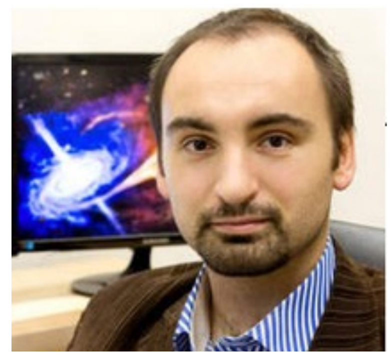

Maciek Weilgus

Ren Ng

Levis*, Srinivasan*, et al, CVPR, 2022

Levis, et al, Nature Astronomy, 2024

Gravitational Lensing Black Hole Emission Tomography

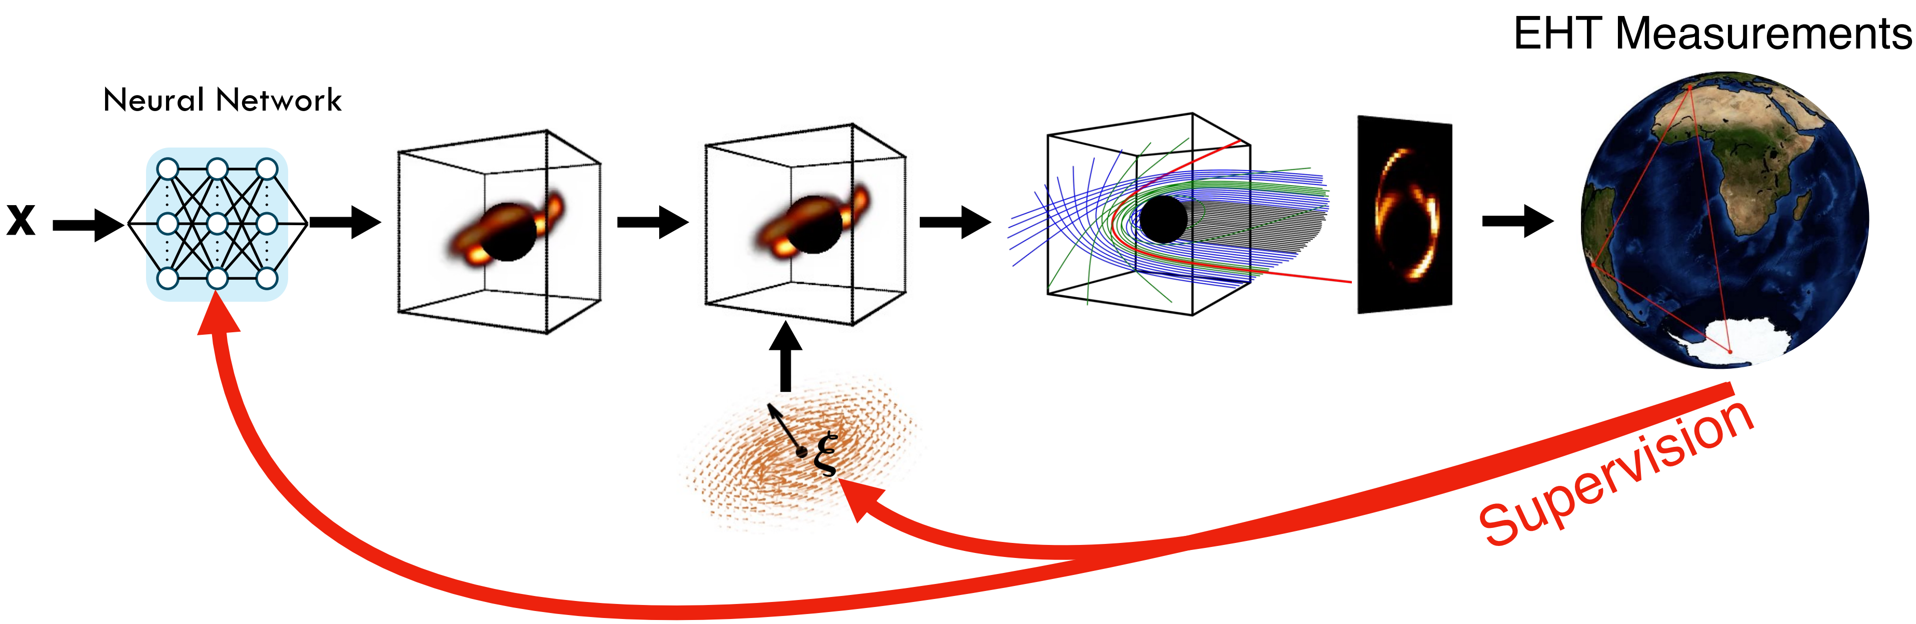

Levis*, Srinivasan*, et al, CVPR, 2022


Galactic Center on April 7: Explosive Day
Galactic Center on April 11: Quiet Days

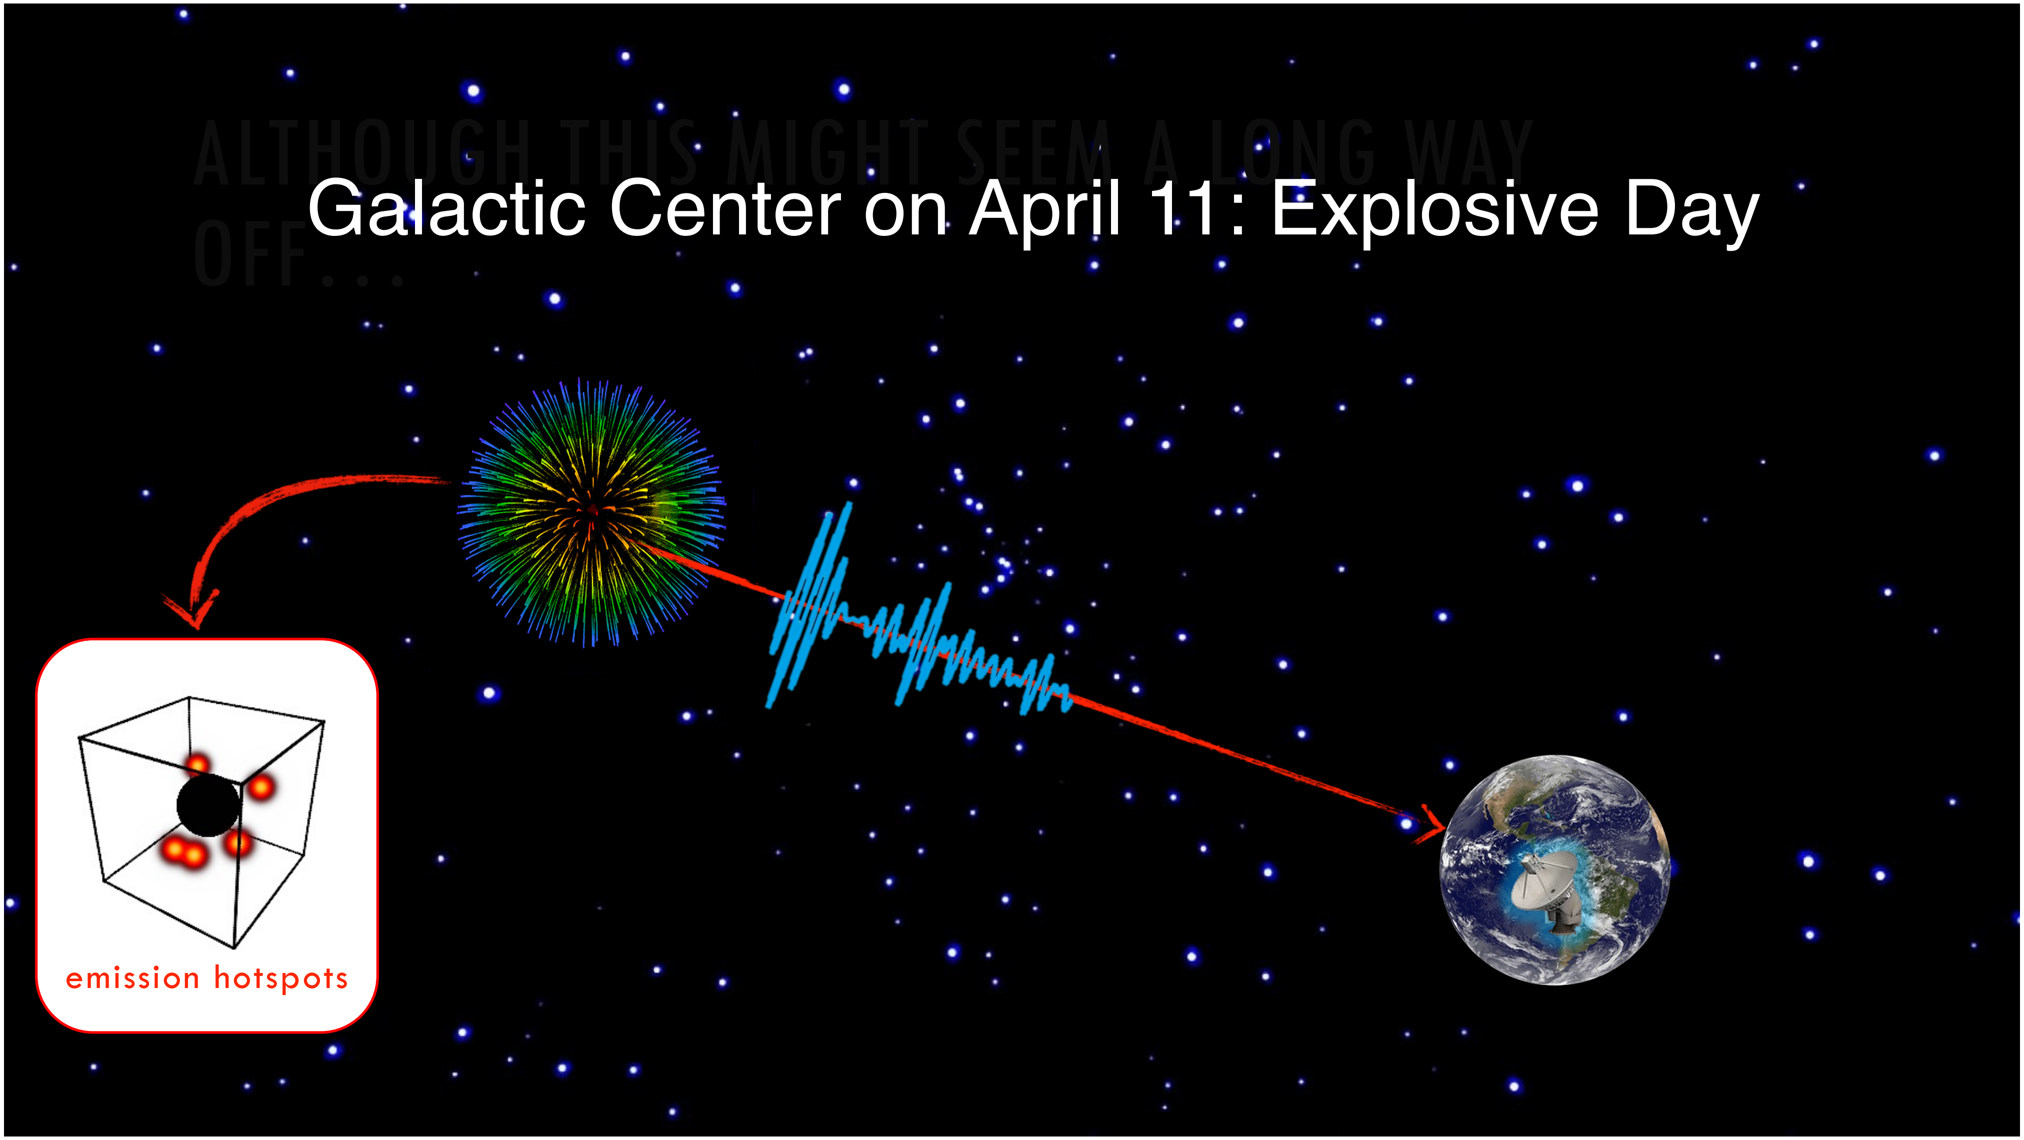
ALTHOUGH THIS MIGHT SEEM A LONG WAY OFF...
Galactic Center on April 11: Explosive Day
emission hotspots

ALMA Observatory

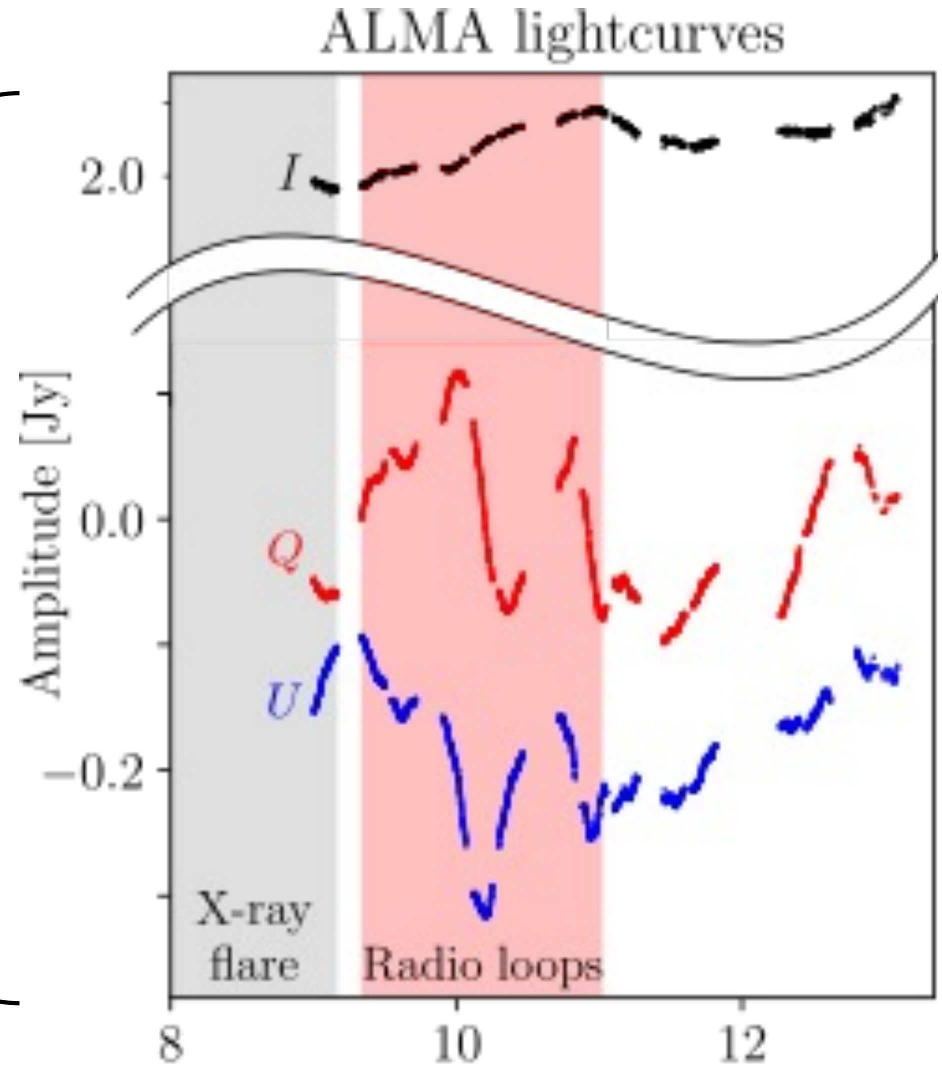

The Black Hole Lightcurve

Evolving
3D Emission

Evolving
2D Projection

Measurements

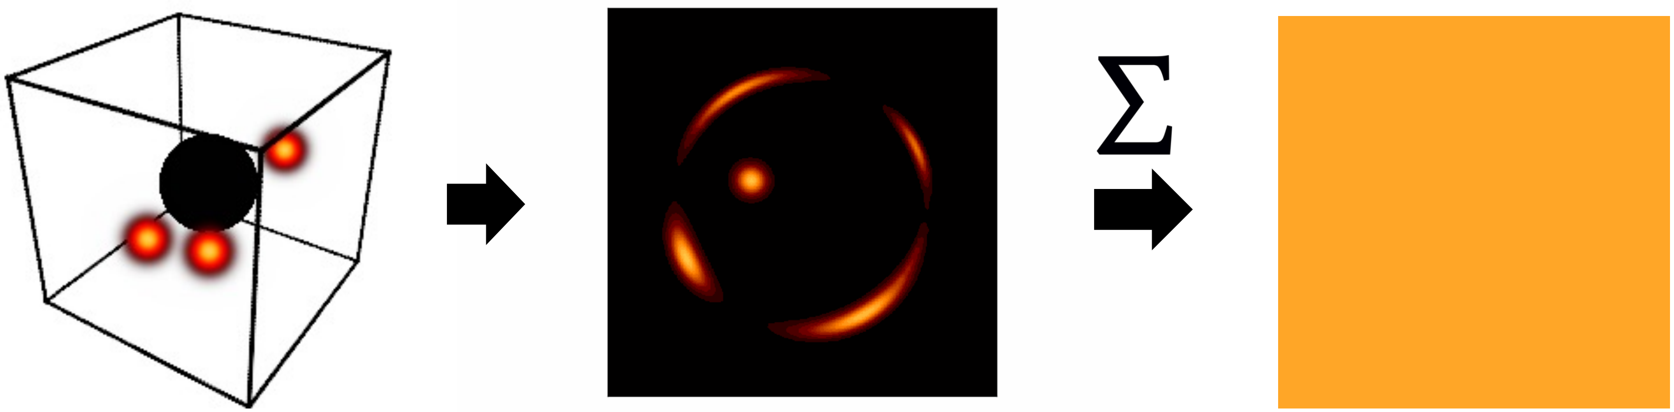

“Lightcurve” :
integrate image to form
a single pixel video

The *Polarized* Black Hole Lightcurve

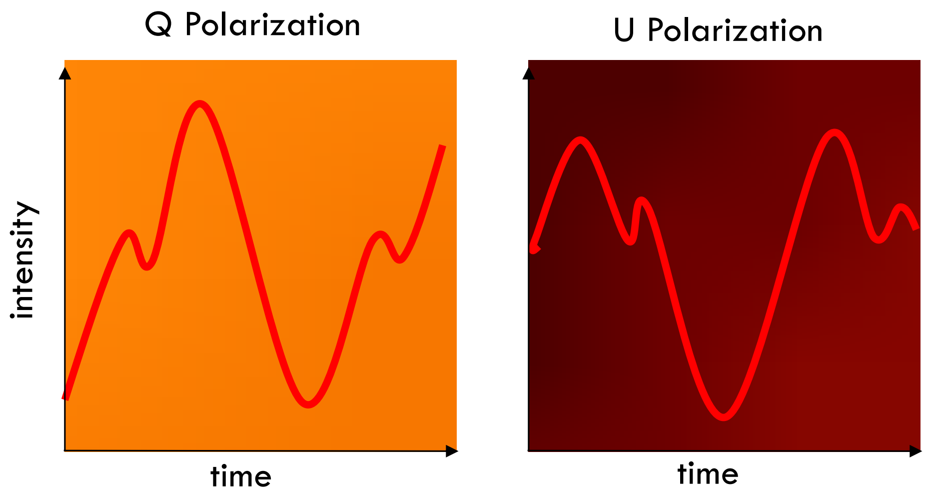

Black Hole Flare Tomography

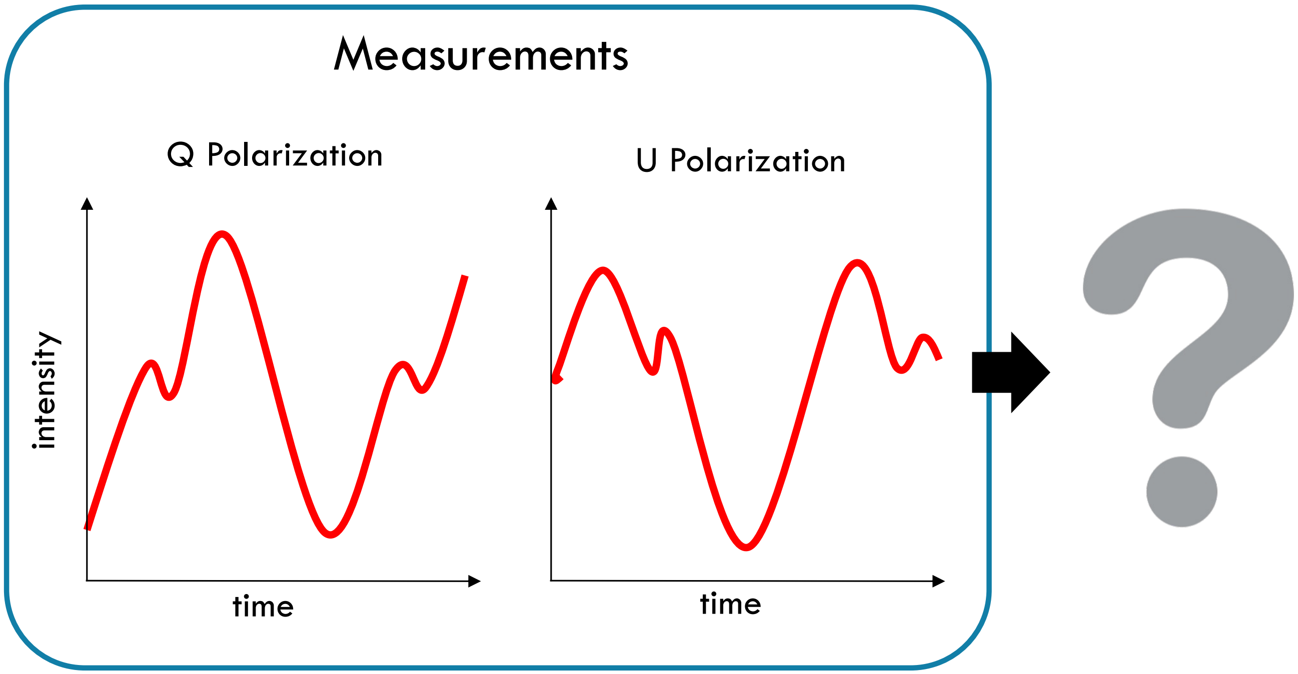

Black Hole Flare Tomography

Measurements

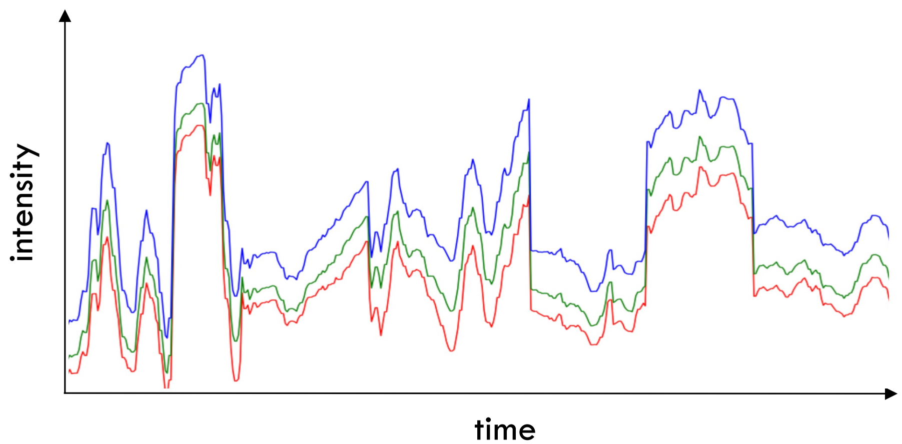

Black Hole Flare Tomography

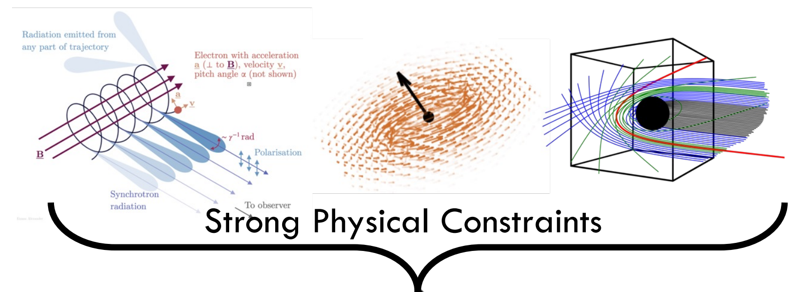

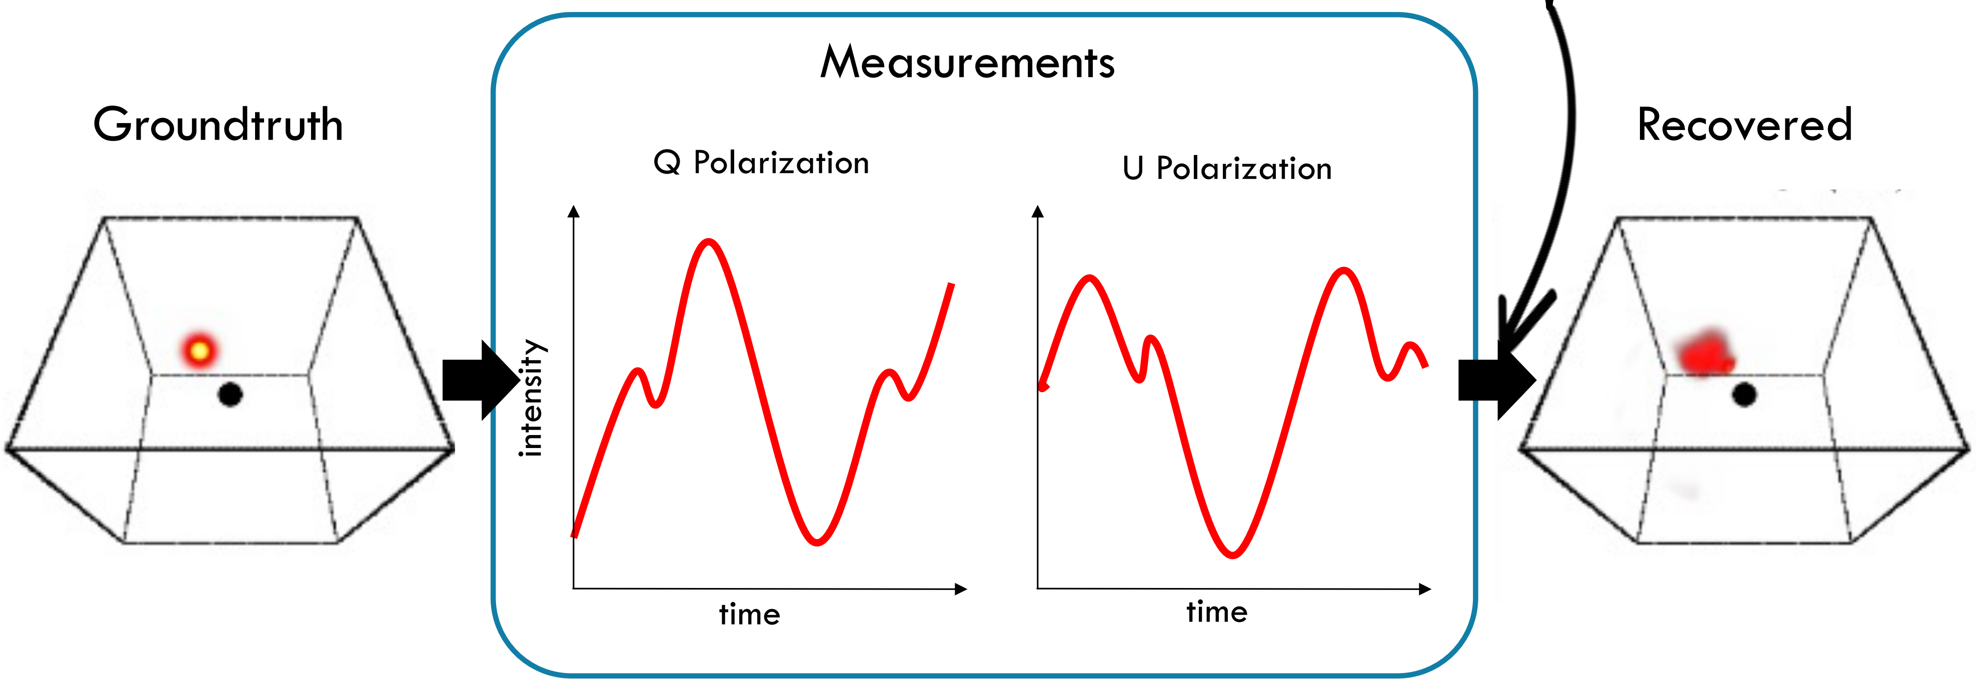

Levis, et al, Nature Astronomy, 2024

Black Hole Flare Tomography

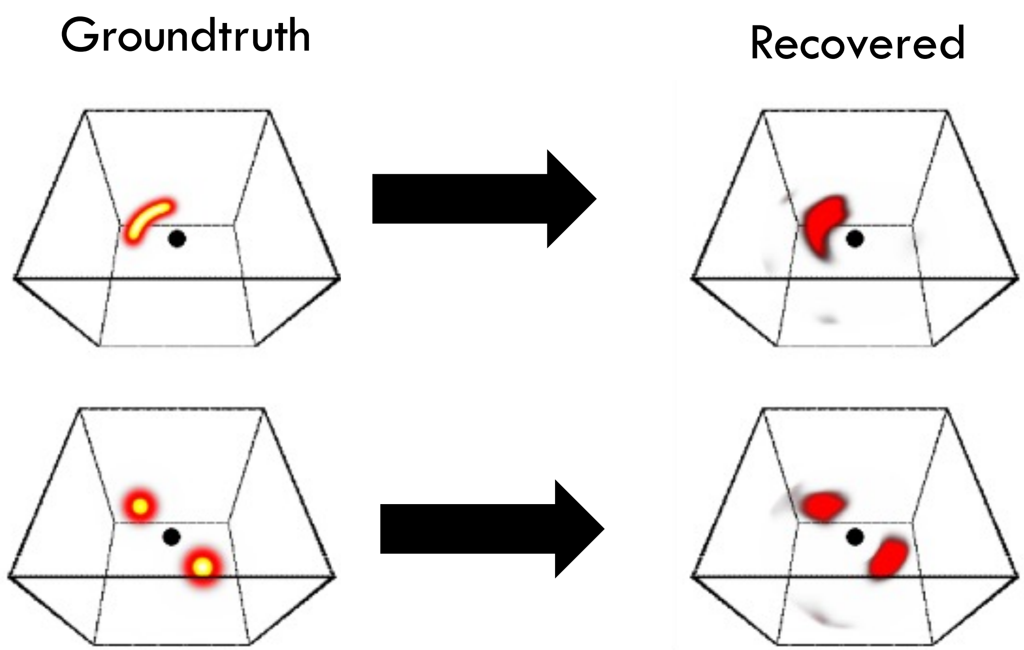

Levis, et al, Nature Astronomy, 2024

Sgr A* Tomography Reconstruction (Real Data!)

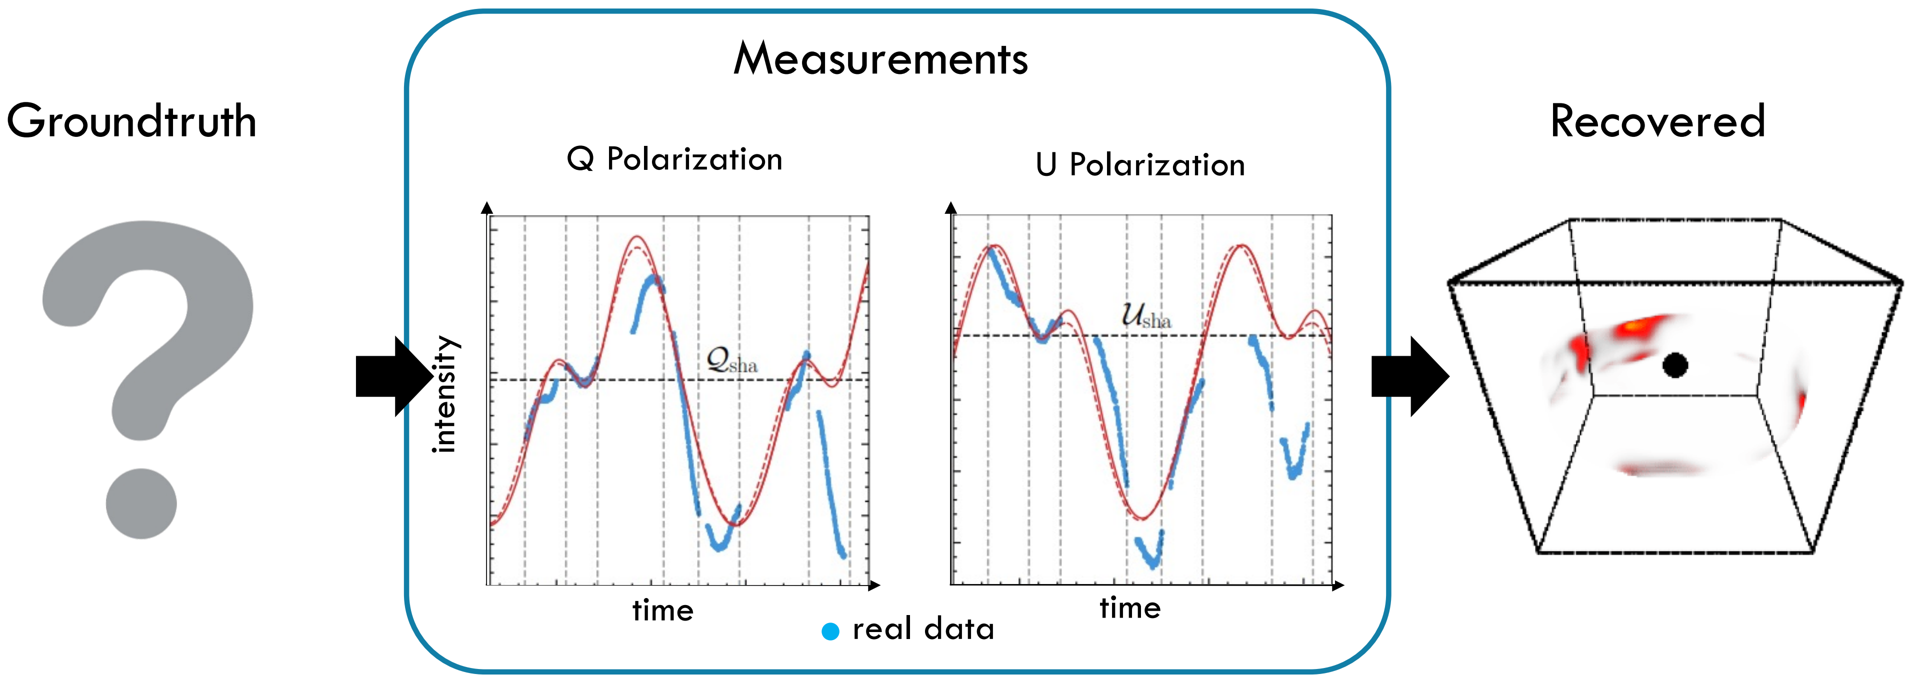

Levis, et al, Nature Astronomy, 2024

Sgr A* Tomography Reconstruction (Real Data!)

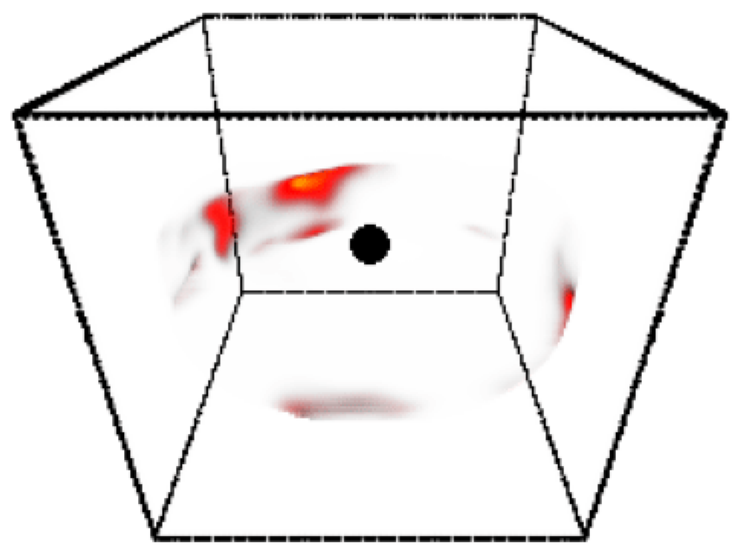

Fixed at Time 9:20 UT

Levis, et al, Nature Astronomy, 2024

Sgr A* Tomography Reconstruction (Real Data!)

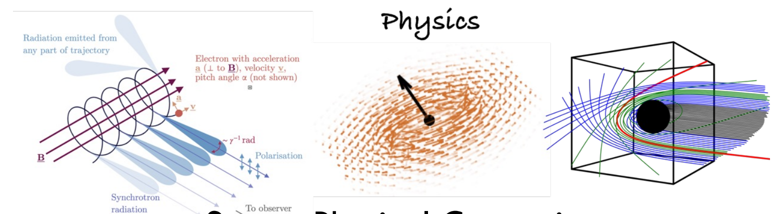

Strong Physical Constraints

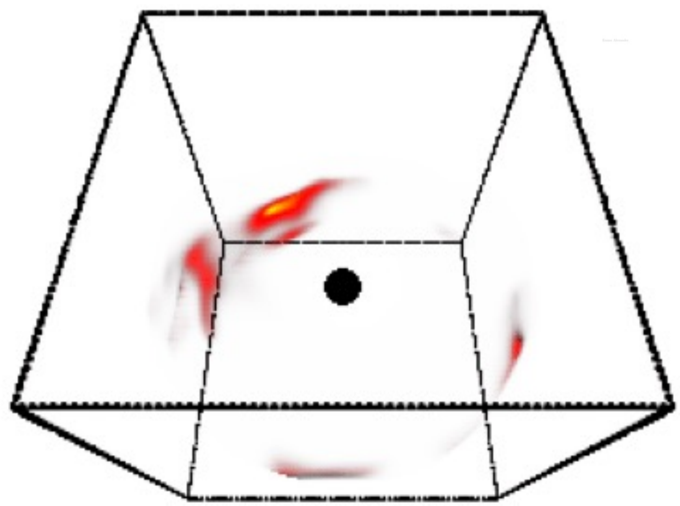

Progression over 100 minutes

Levis, et al, Nature Astronomy, 2024

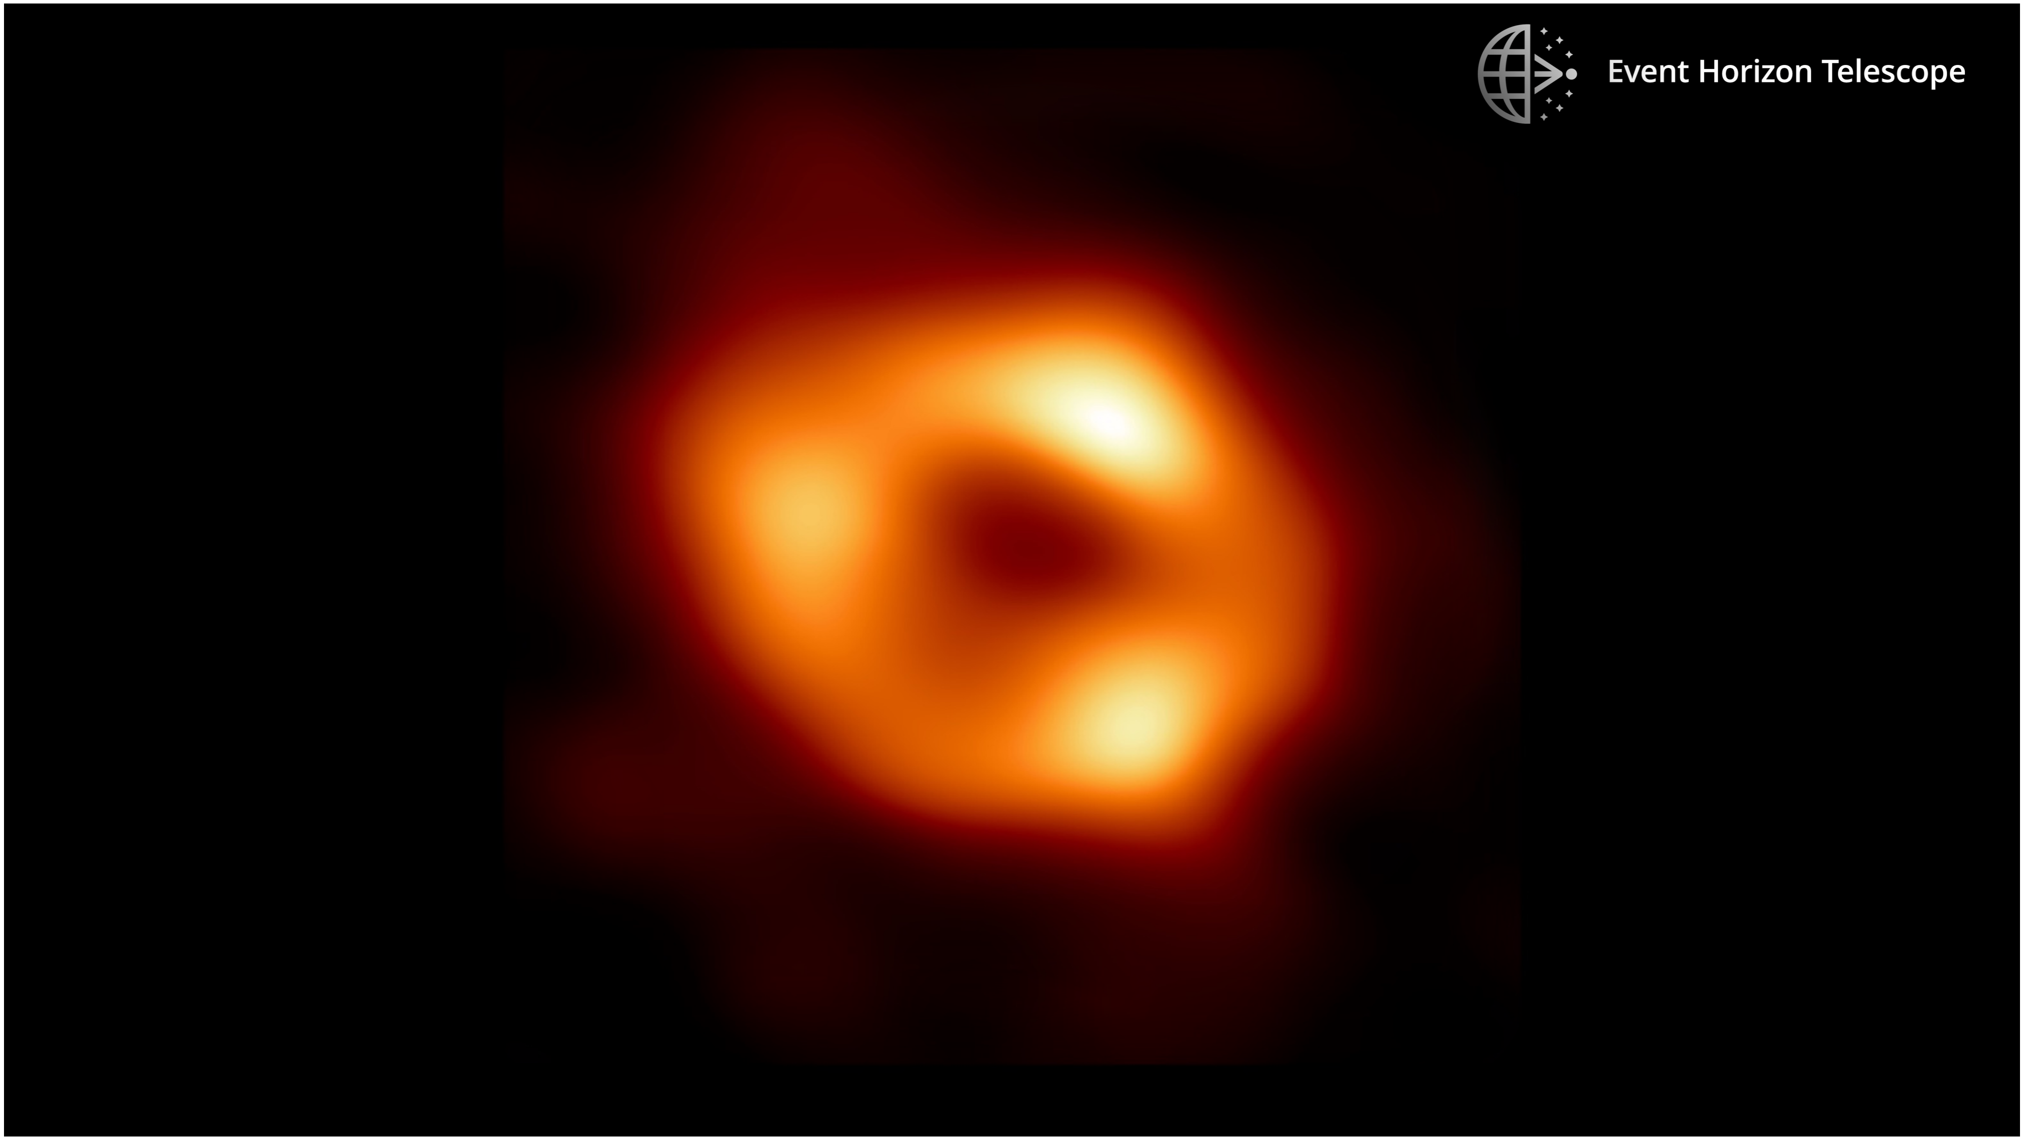
Event Horizon Telescope

The 2-Way Street Between Science and Algorithms

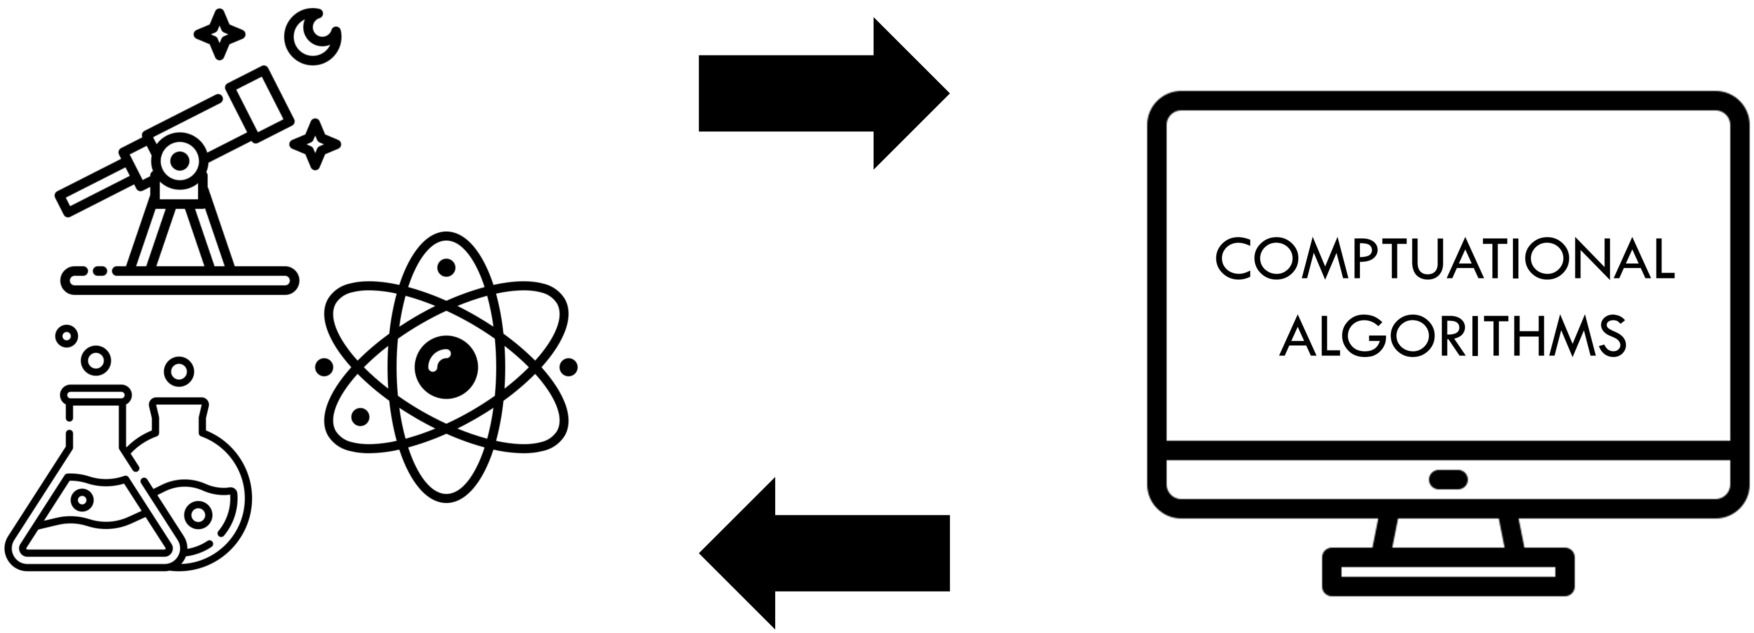

The 2-Way Street Between Science and Algorithms

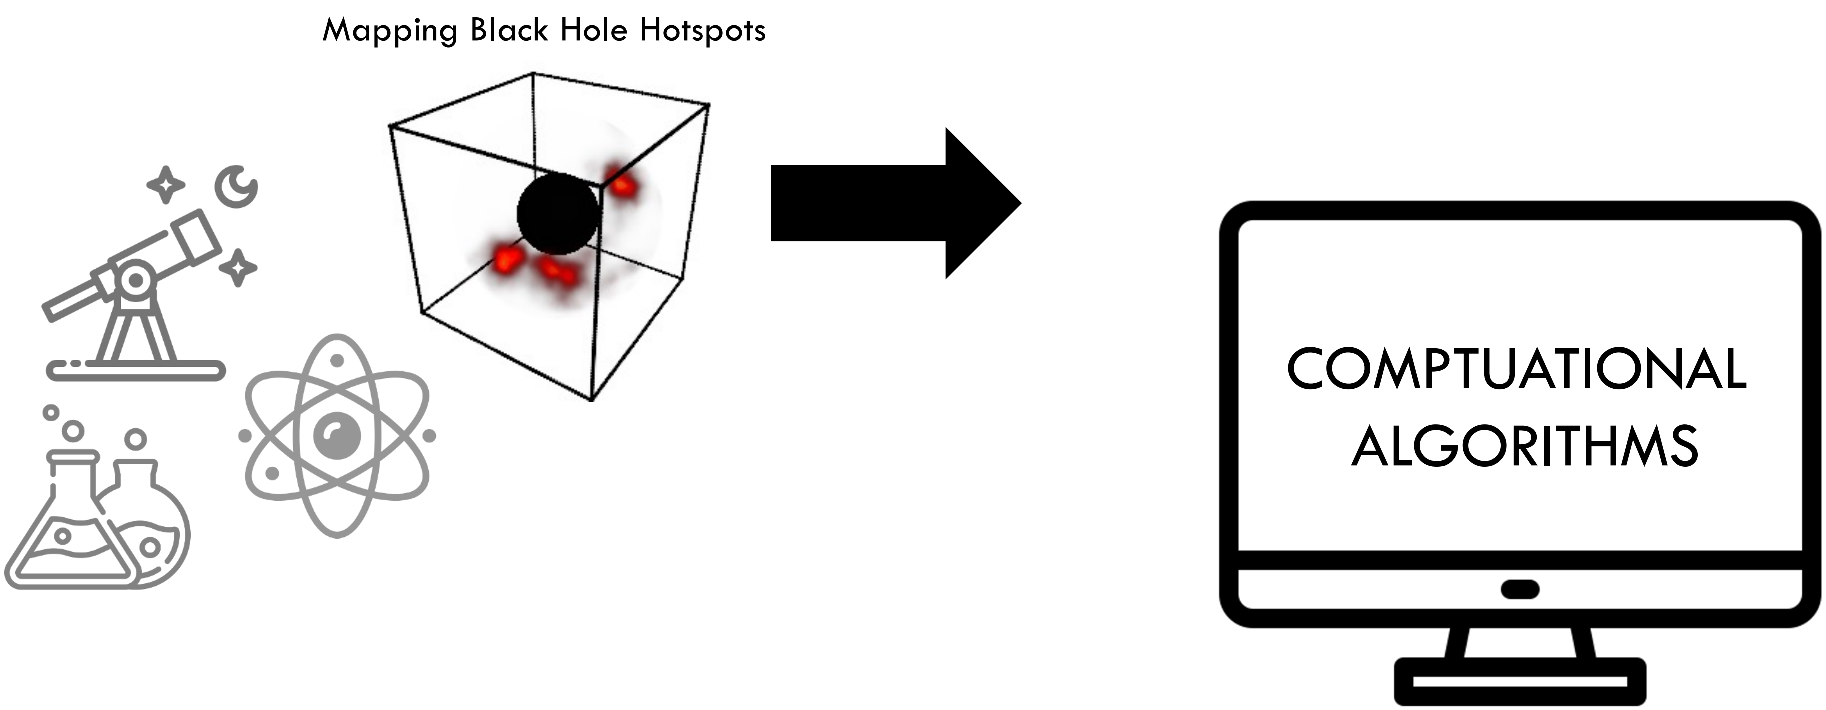

The 2-Way Street Between Science and Algorithms

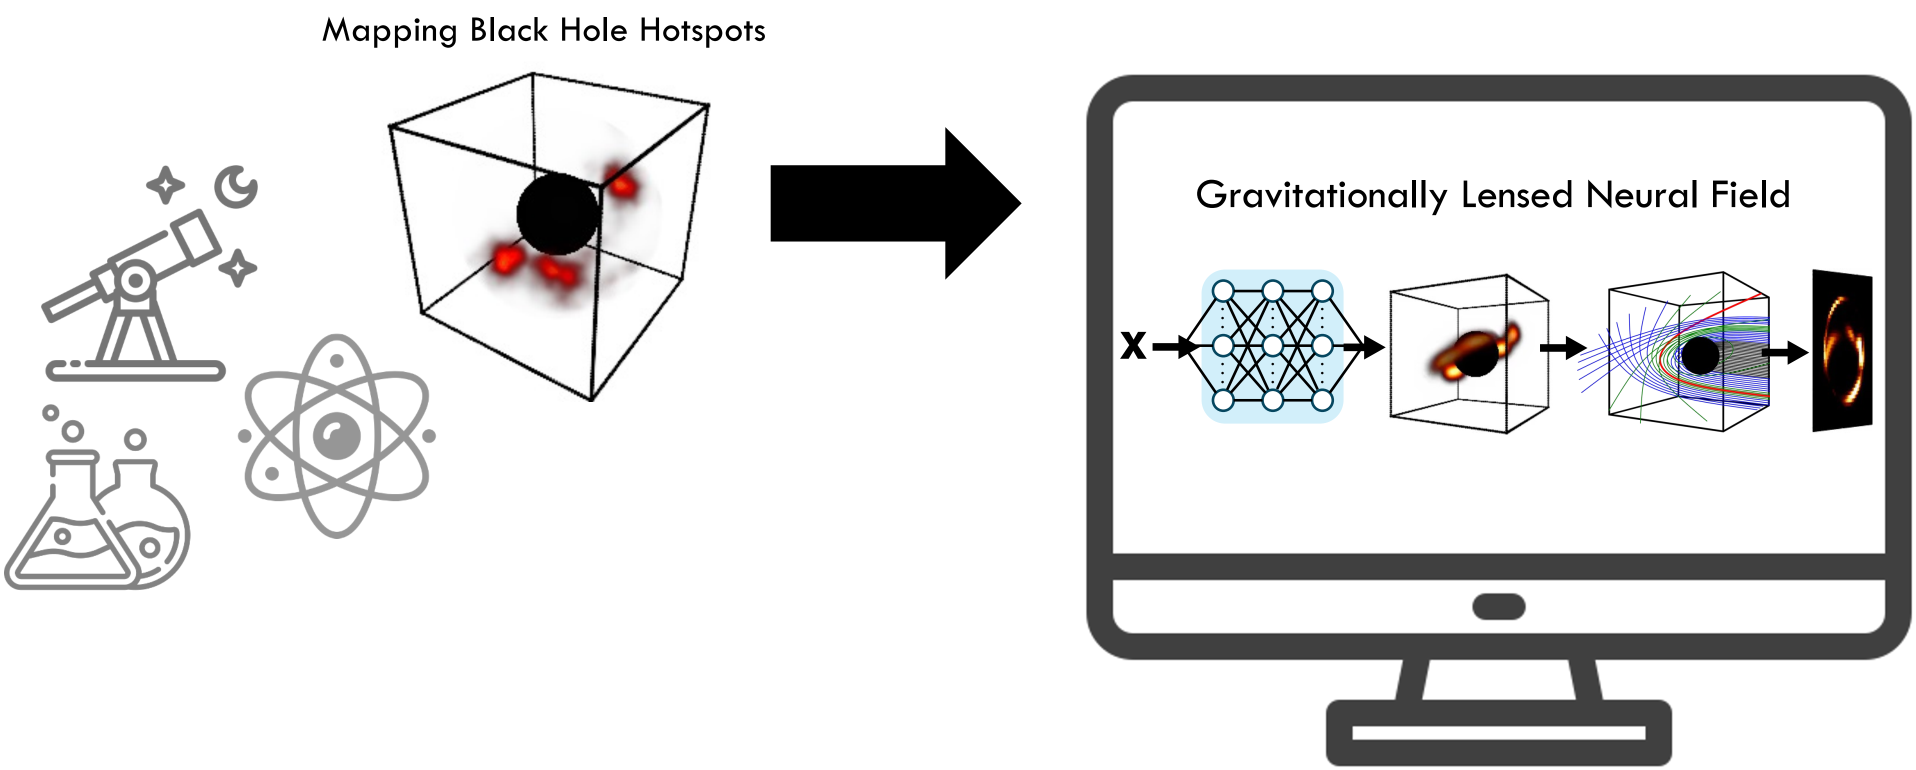

The 2-Way Street Between Science and Algorithms

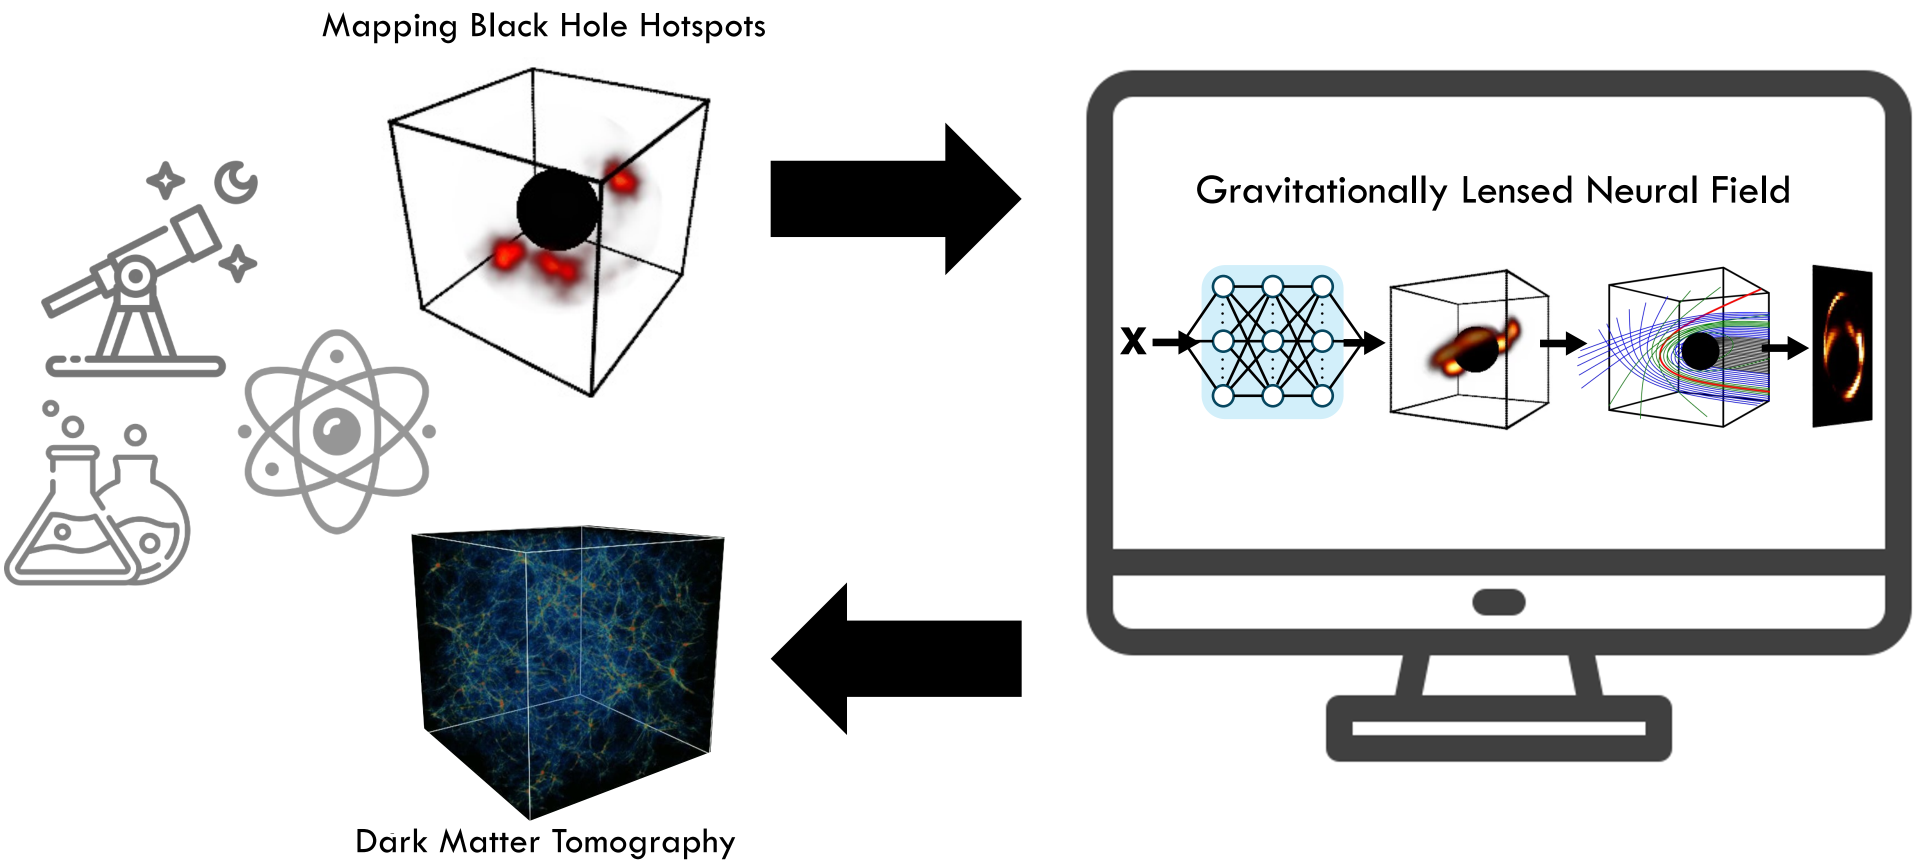

Revealing the 3D Cosmic Web through Gravitationally Constrained Neural Fields

Brandon Zhao

Aviad Levis

Liam Connor

Pratul P. Srinivasan

Zhao, et al, CVPR, 2024

Zhao, et al, in prep

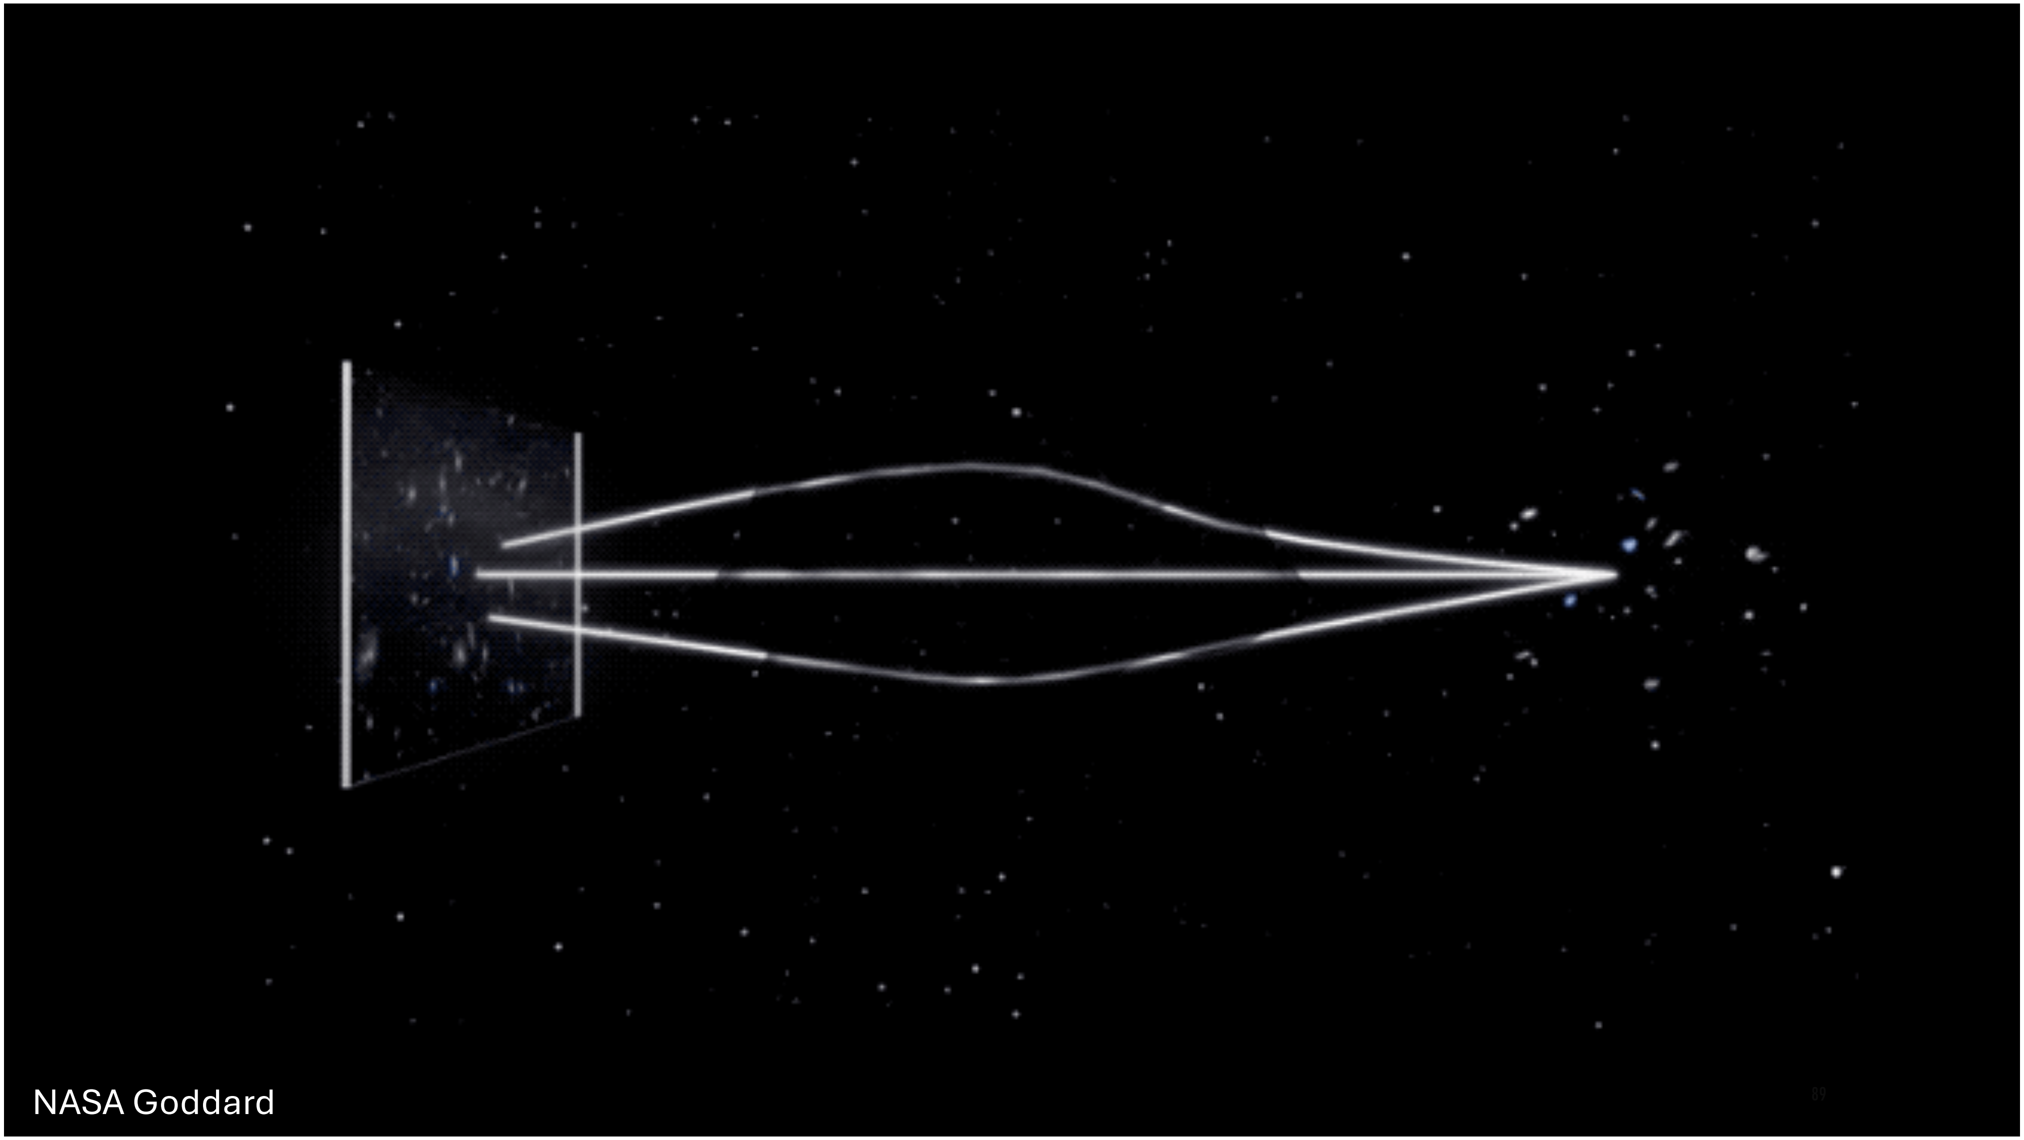
NASA Goddard

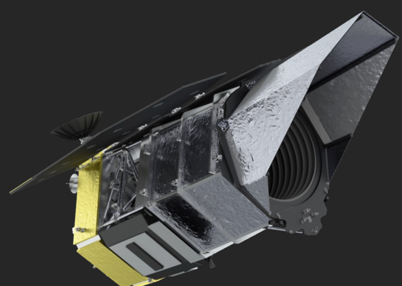

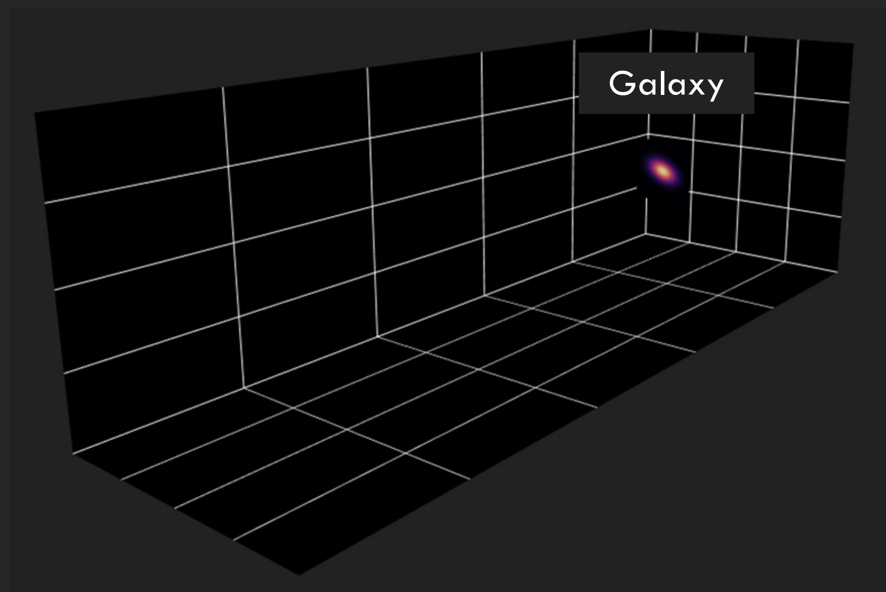

Telescope

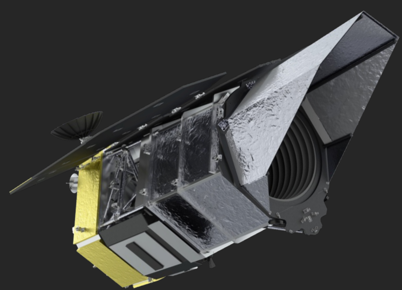

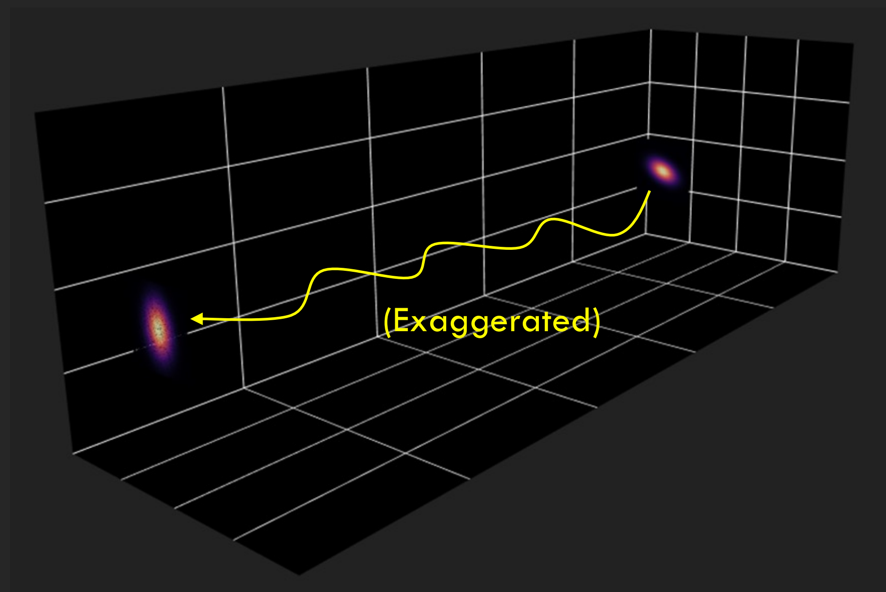
(Exaggerated)

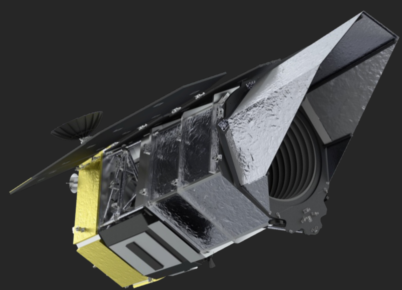

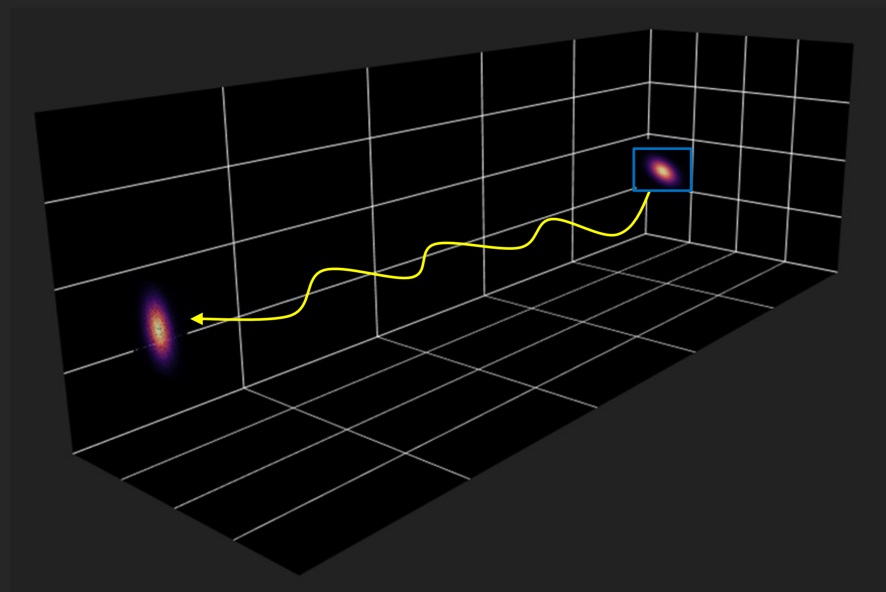

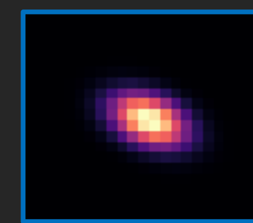

“Intrinsic” Shape

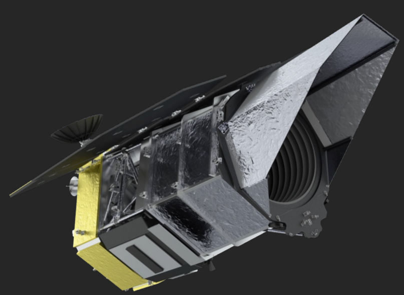

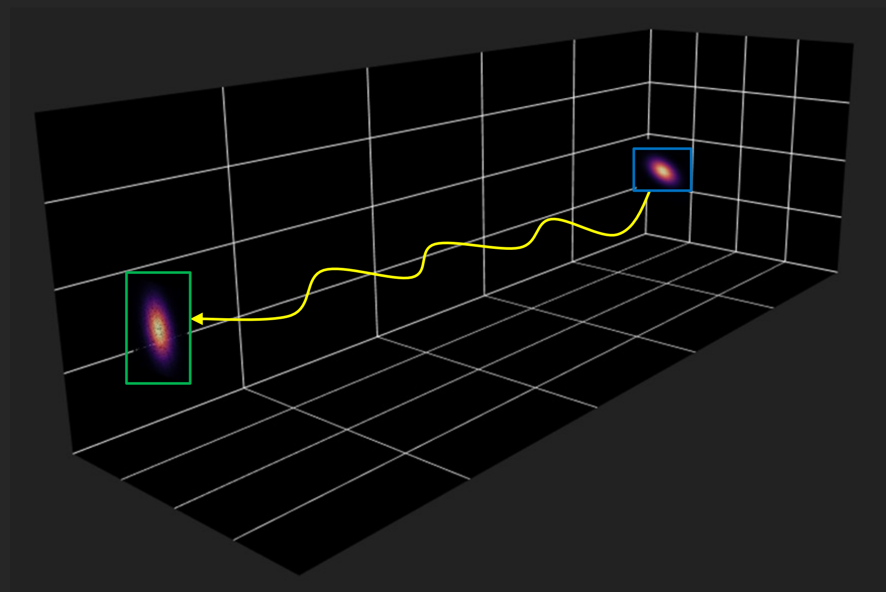

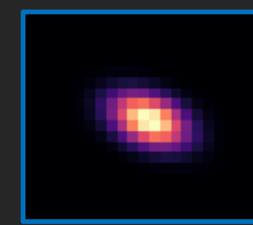

“Intrinsic” Shape

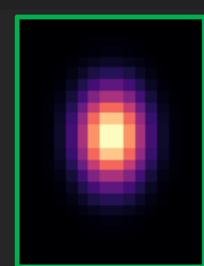

Observed (Lensed) Shape

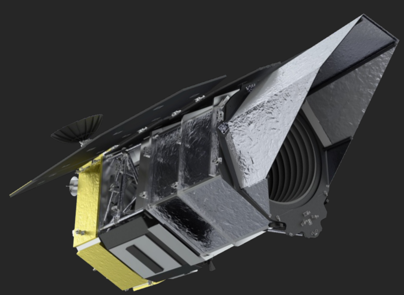

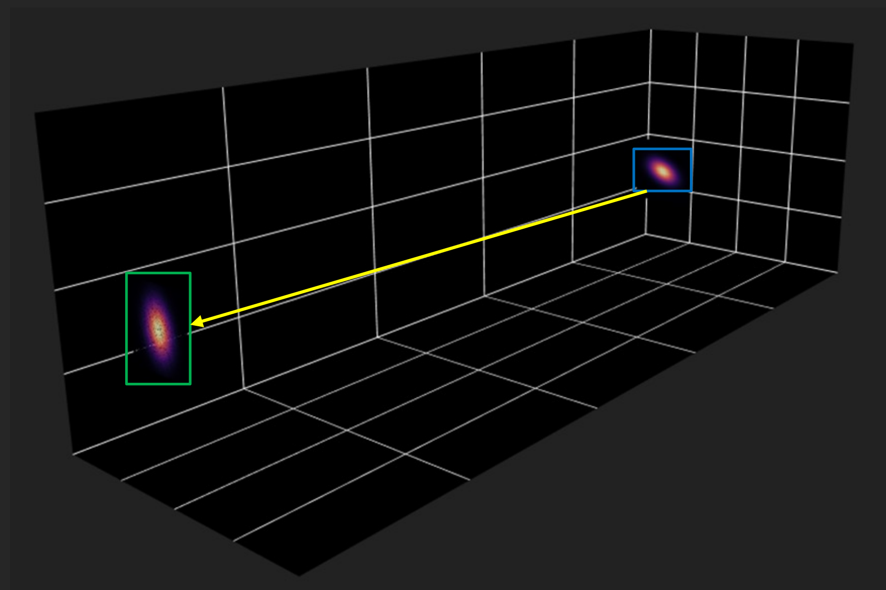

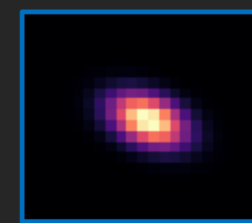

“Intrinsic” Shape

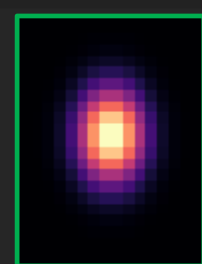

Observed (Lensed)
Shape

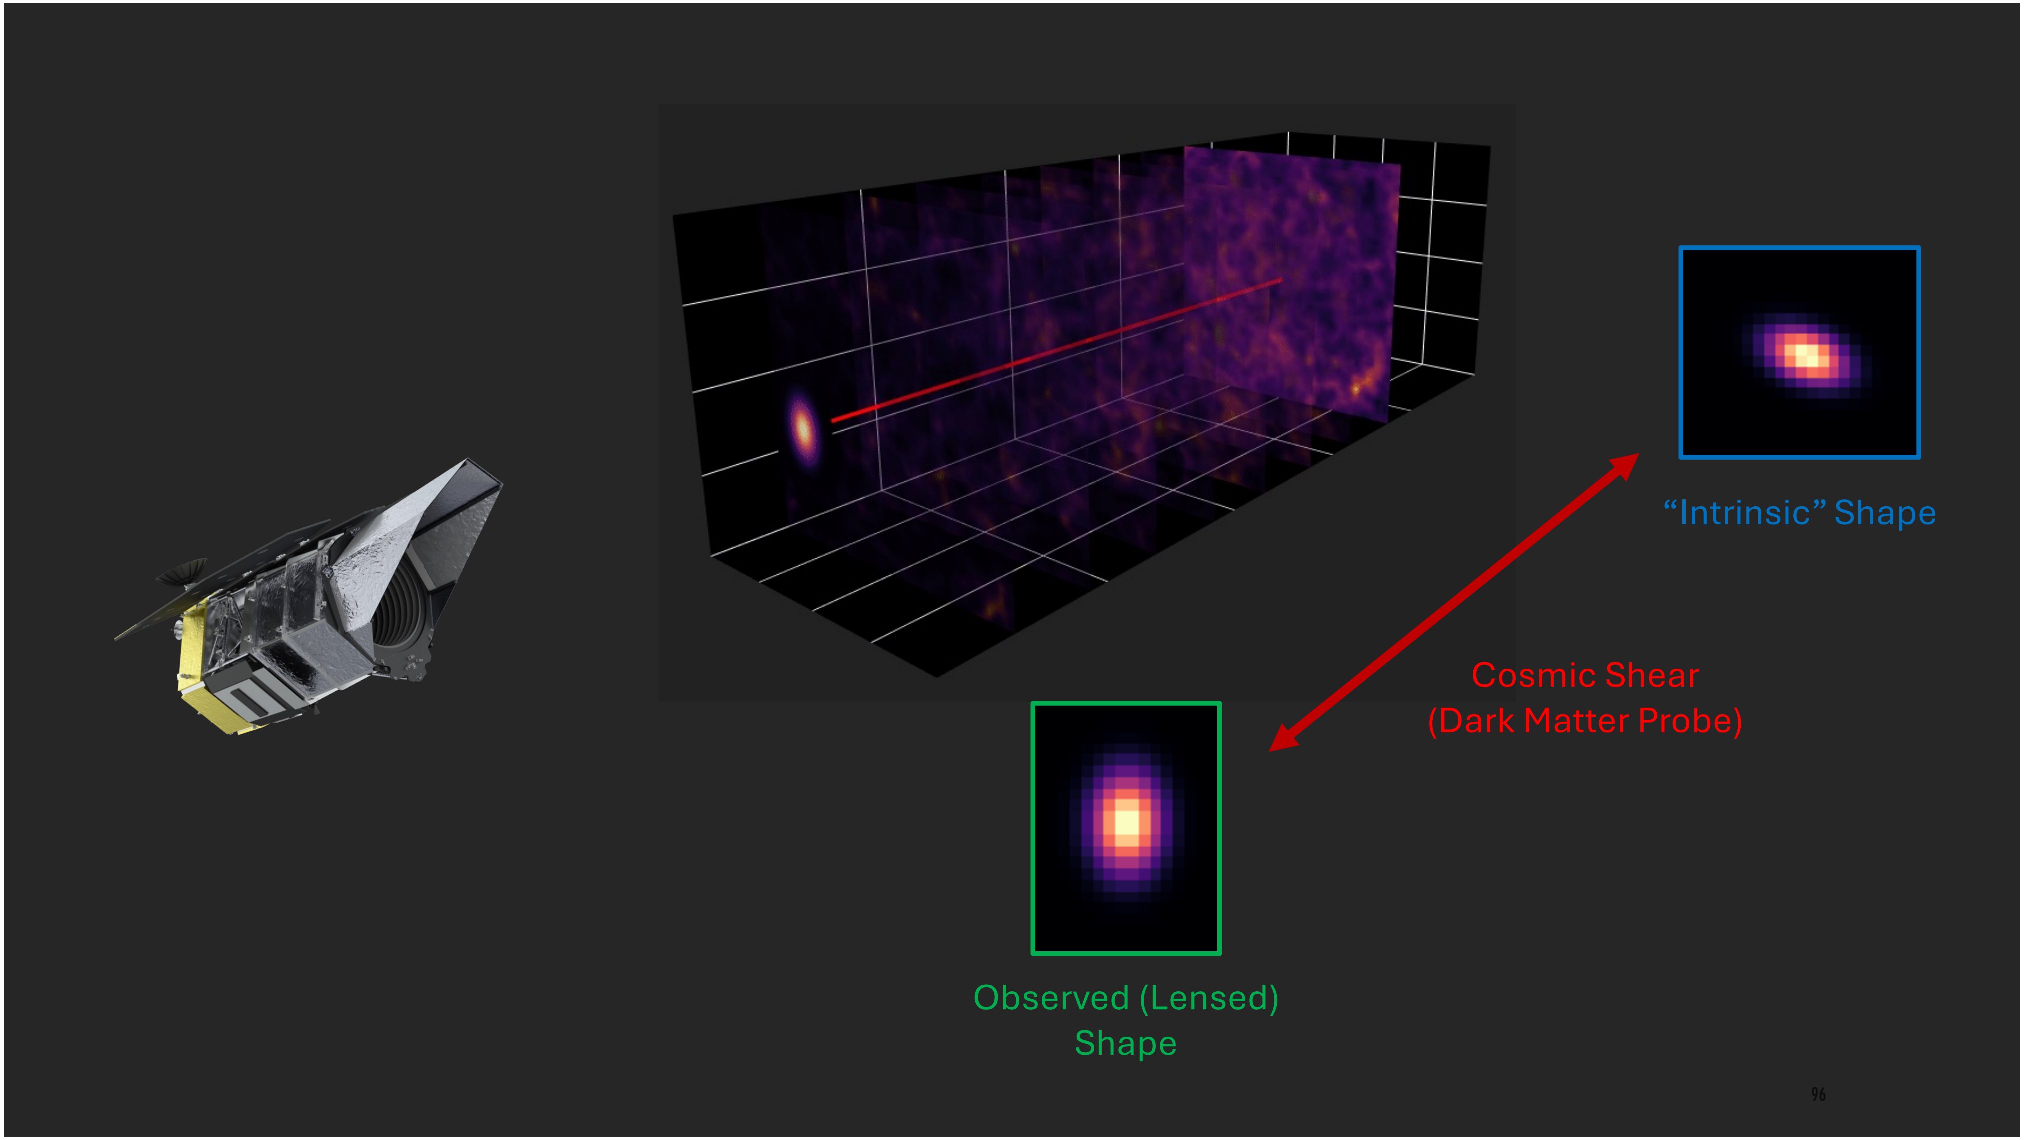
"Intrinsic" Shape
Cosmic Shear
(Dark Matter Probe)
Observed (Lensed)
Shape
96

The Elliptical Parameterization of Galaxies

To describe an ellipse, define its **complex ellipticity:**

$$e = e_1 + ie_2$$

Where the **magnitude** and **phase** determine its **axis ratio** ***r*** and **orientation angle** $\boldsymbol{\phi}$:

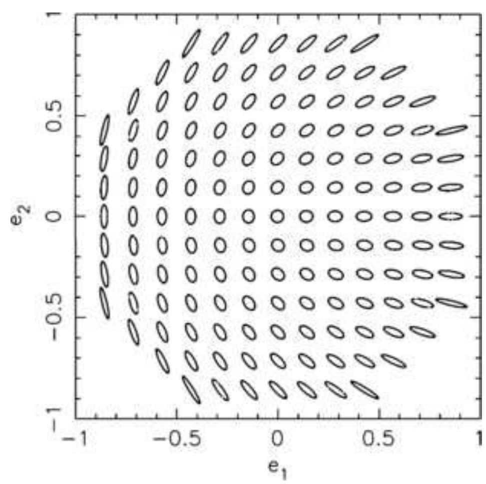

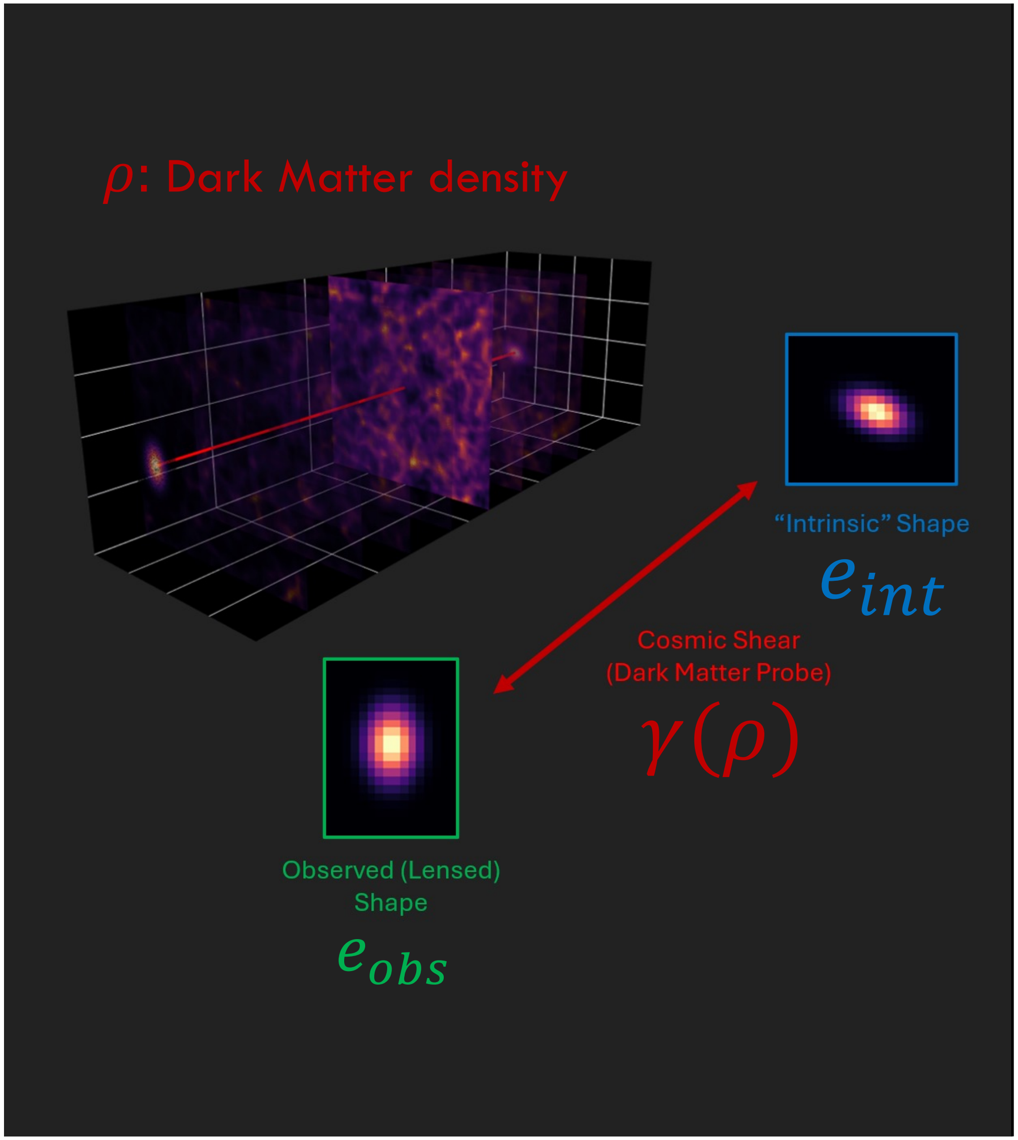

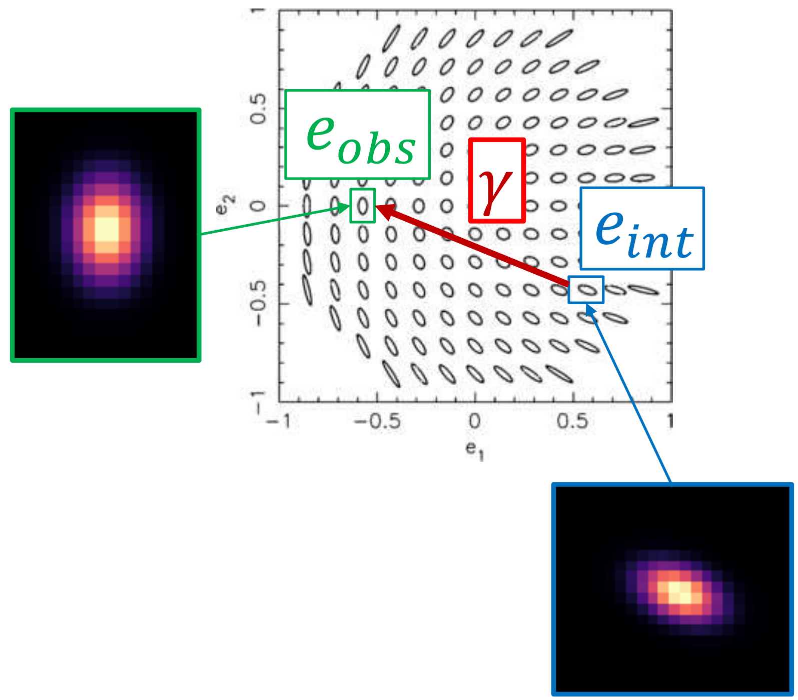

$$e_{obs} - e_{int} = \gamma(\rho)$$

what we want

Estimates are Noisy: "Shape Noise"

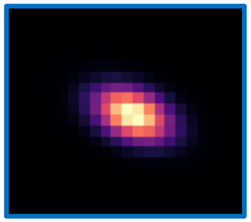

"Intrinsic" Shape

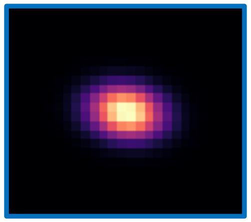

Estimated Shape

Shape Noise

$$e_{int} = \hat{e}_{int} - \epsilon \quad \text{where} \quad \epsilon \sim \mathcal{N}(0, \sigma_{shape})$$

$$\boxed{\gamma_{meas} = e_{obs} - \hat{e}_{int}}$$

Estimates are Noisy: “Shape Noise”

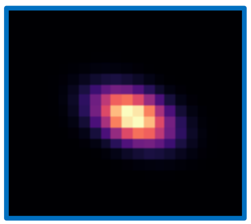

“Intrinsic” Shape

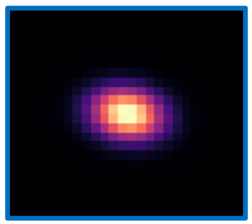

Estimated Shape

Shape Noise

$$e_{int} = \hat{e}_{int} - \epsilon \quad \text{where} \quad \epsilon \sim \mathcal{N}(0, \sigma_{shape})$$

$$\gamma_{meas} = e_{obs} - (e_{int} - \epsilon)$$

Estimates are Noisy: "Shape Noise"

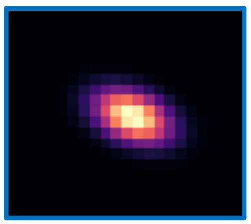

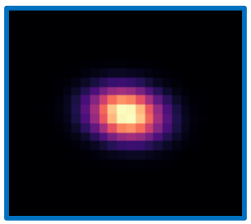

"Intrinsic" Shape — Estimated Shape — Shape Noise

$$e_{int} = \hat{e}_{int} - \epsilon \quad \text{where} \quad \epsilon \sim \mathcal{N}(0, \sigma_{shape})$$

$$\gamma_{meas} = e_{obs} - e_{int} + \epsilon = \gamma(\rho) + \epsilon$$

Galaxy Catalog

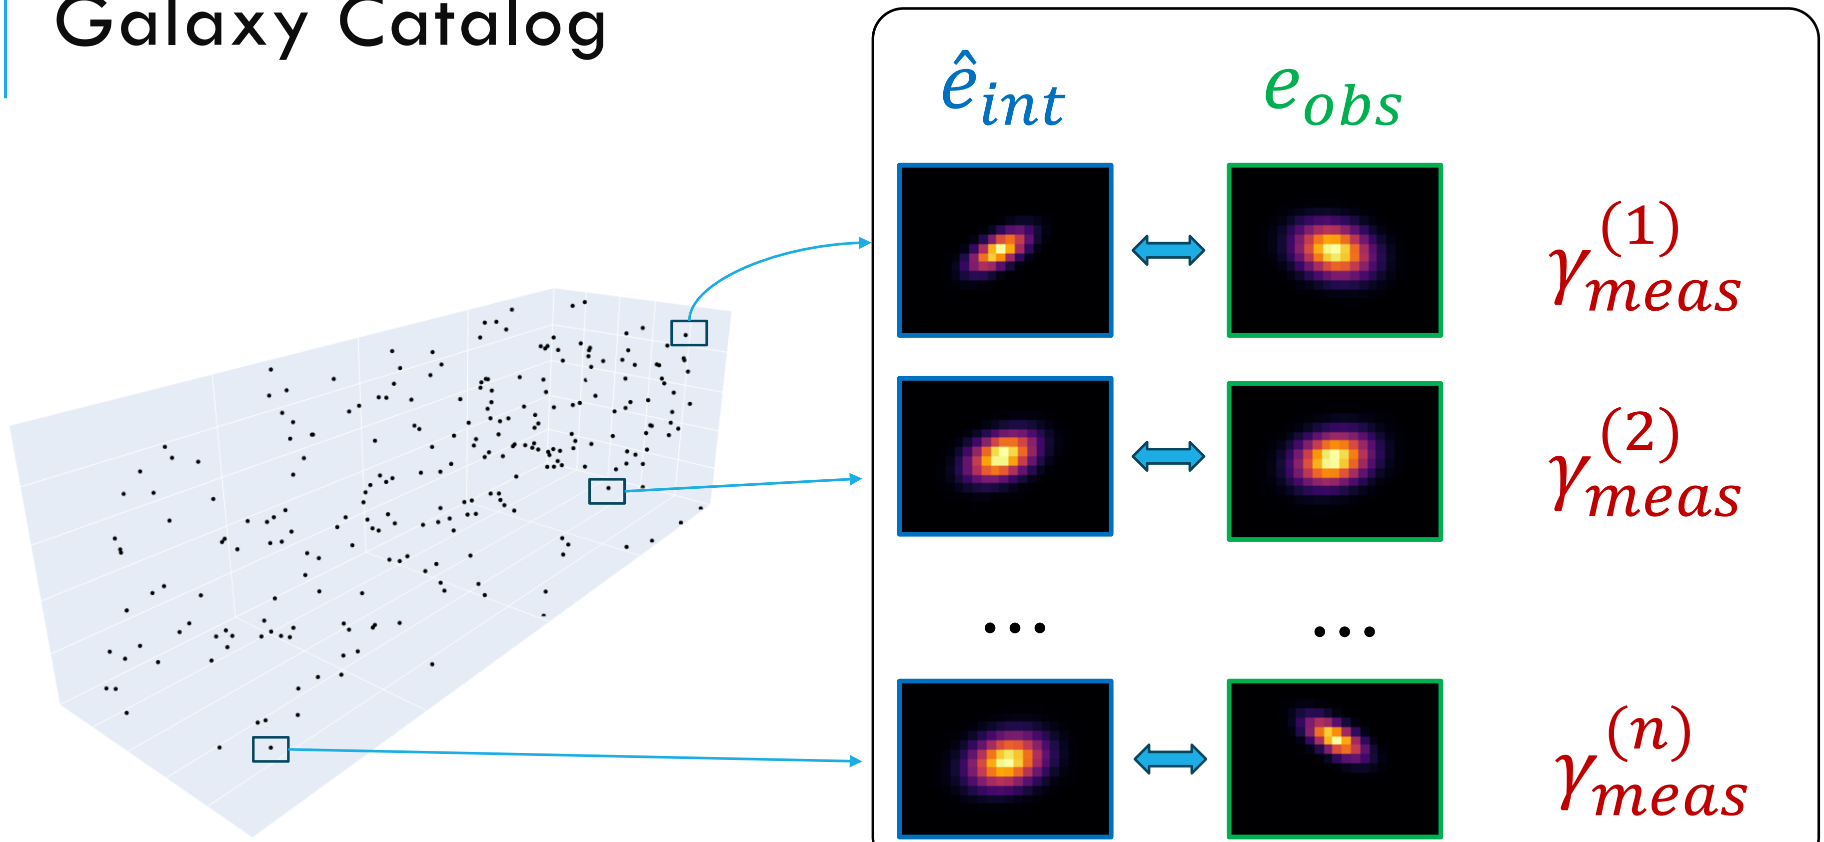

Mass Density Neural Field

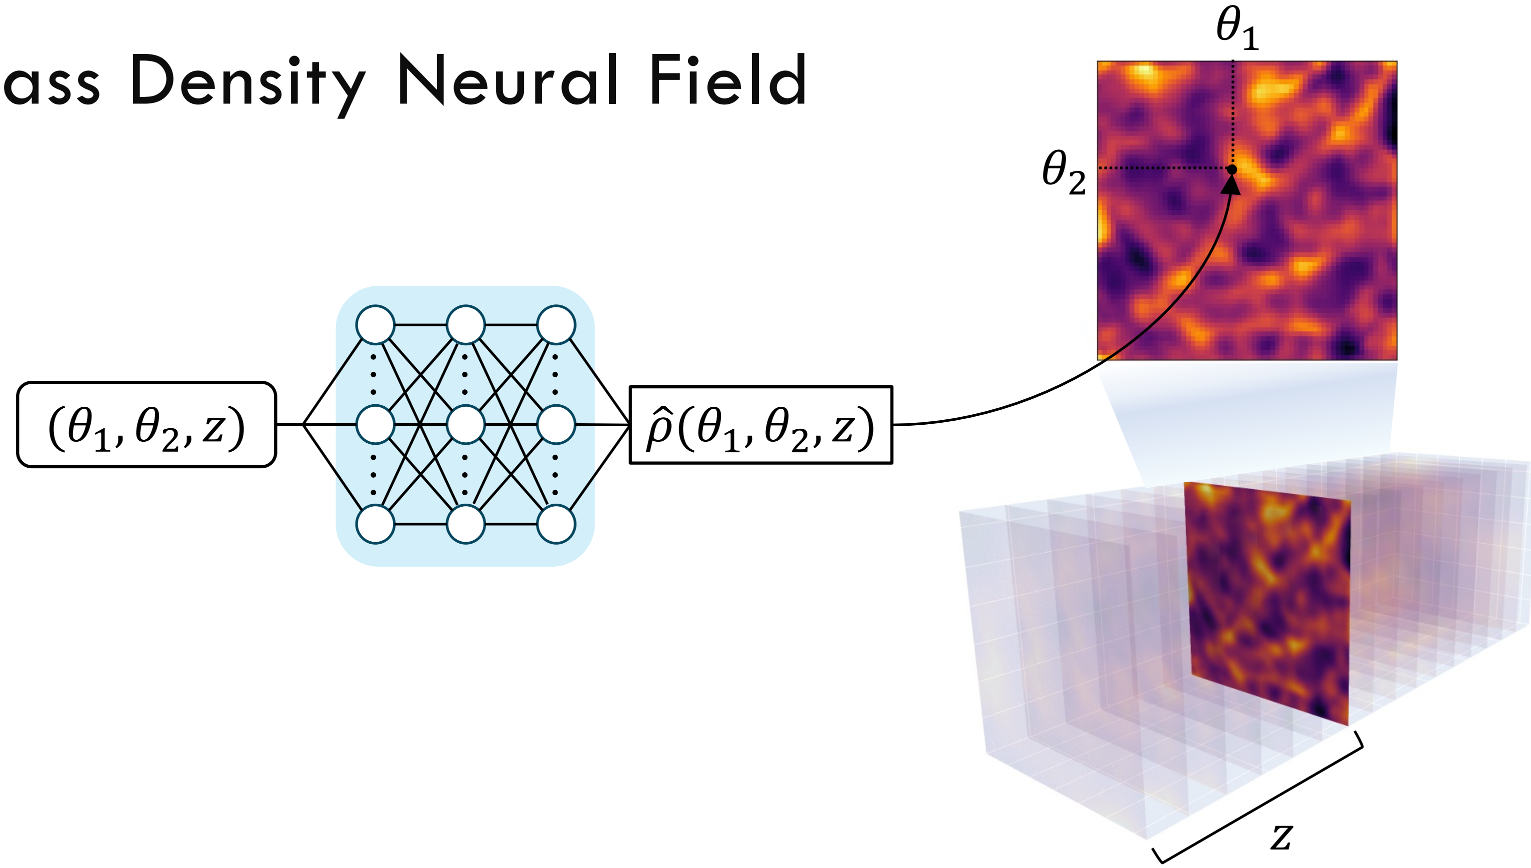

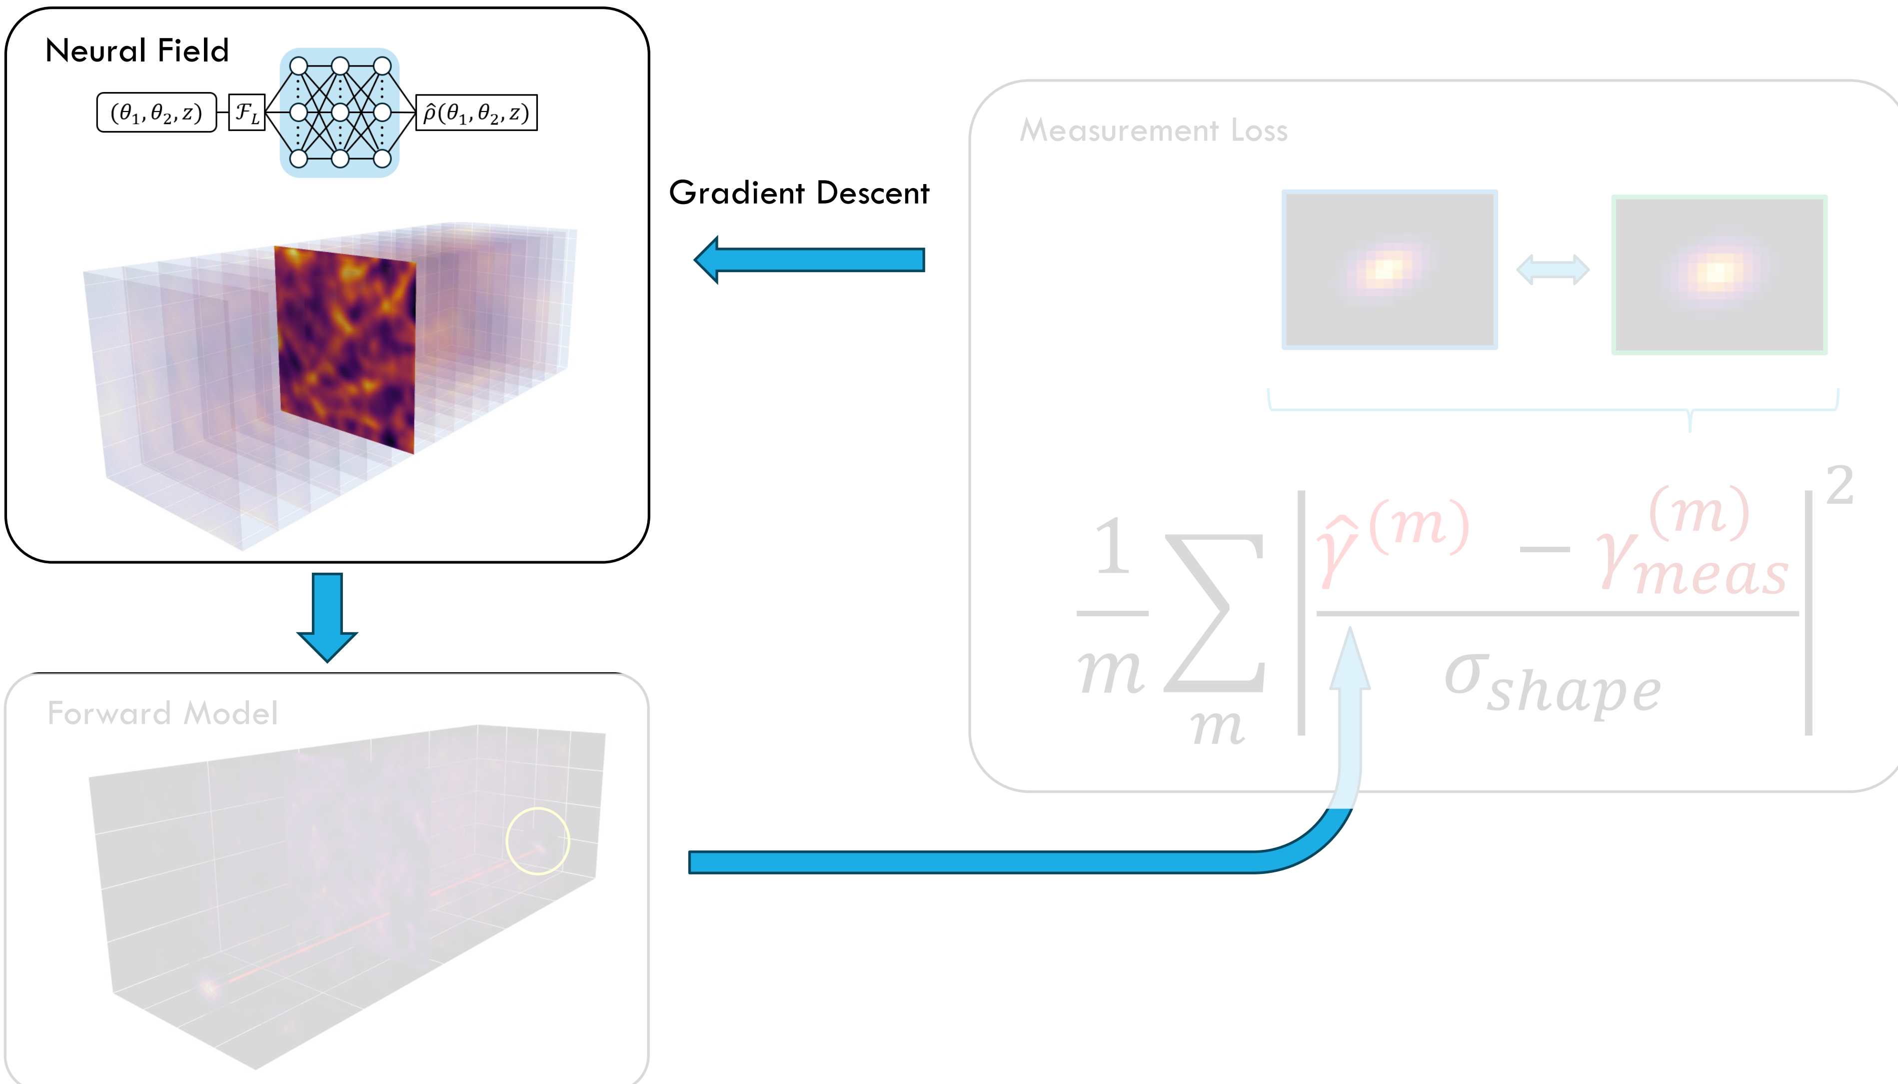

Neural Field
(θ_1, θ_2, z)
$\mathcal{F}_L$
$\hat{\rho}(\theta_1, \theta_2, z)$
Gradient Descent
Measurement Loss
$\frac{1}{m}\sum_m \left| \frac{\hat{\gamma}^{(m)} - \gamma_{meas}^{(m)}}{\sigma_{shape}} \right|^2$
Forward Model

Simulated Dark Matter Field
(Blurred)

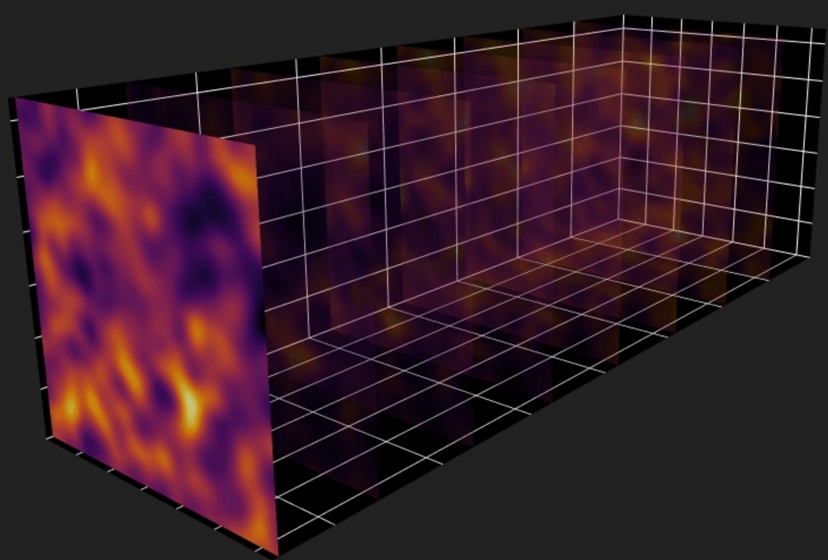

Reconstruction

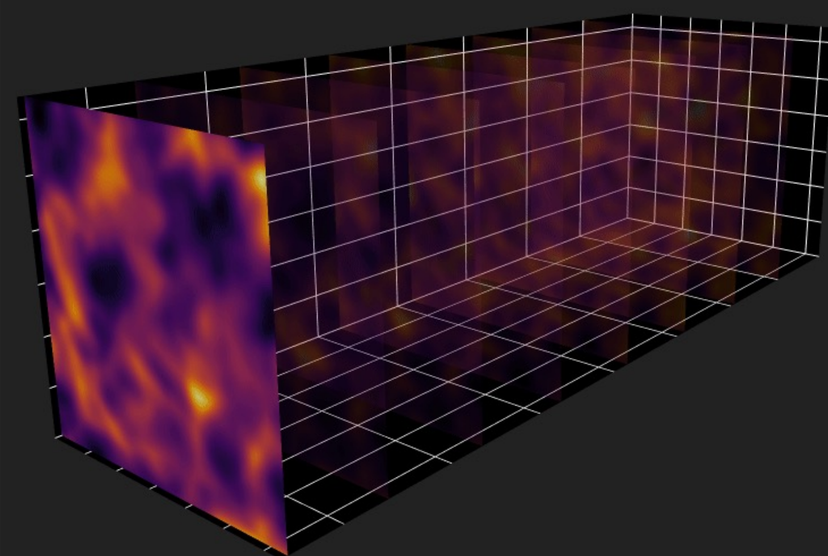

Zhao, et al, in prep

The 2-Way Street Between Science and Algorithms

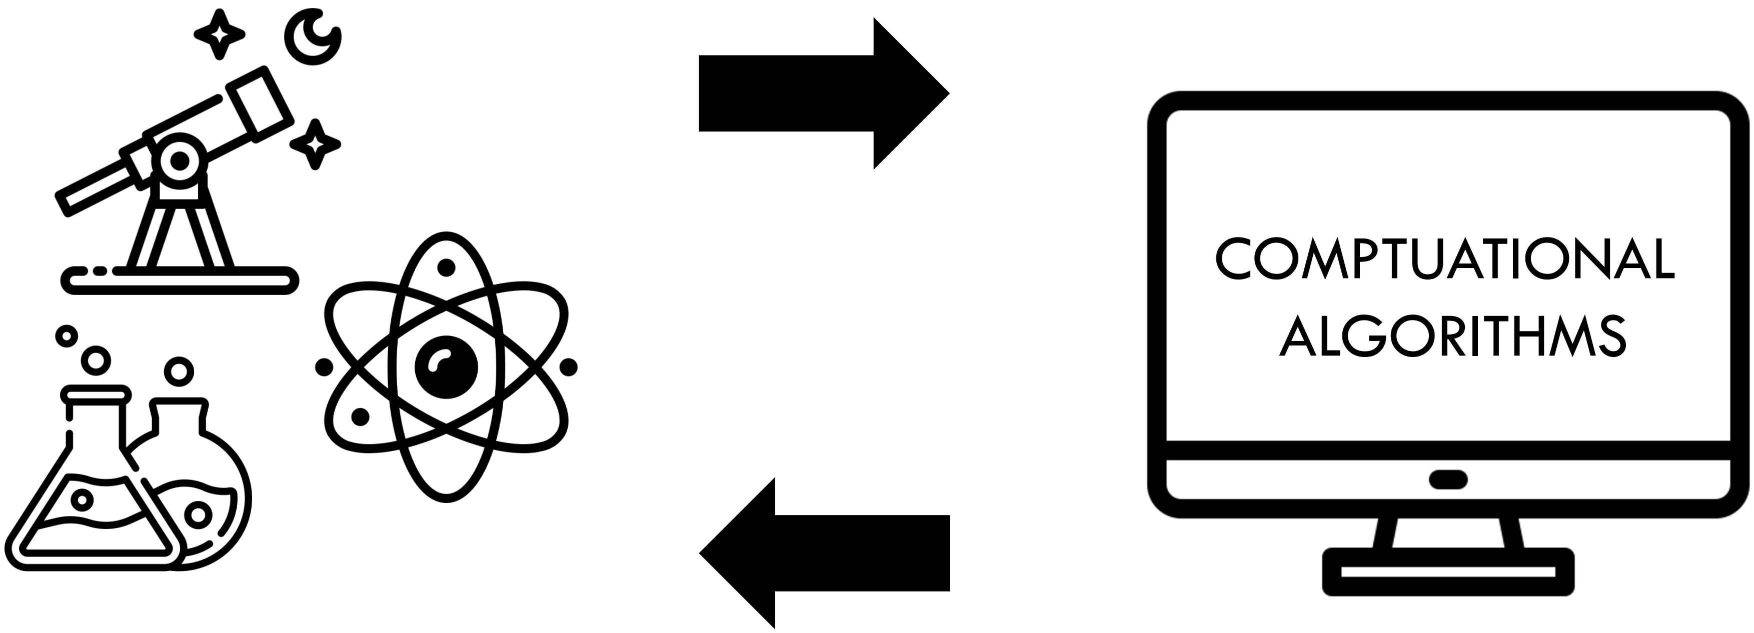


The Event Horizon Telescope Collaboration
Over 300 Scientists from 80 institutes in countries spanning
Europe, Asia, Africa, North and South America
(along with ~23K Community Contributors from Open-Source Projects)


Zihui (Ray) Wu
Berthy Feng
Aviad Levis
Brandon Zhao