

Bayesian Imaging for Intensity Interferometry with Deep Generative Priors



Biwei Dai

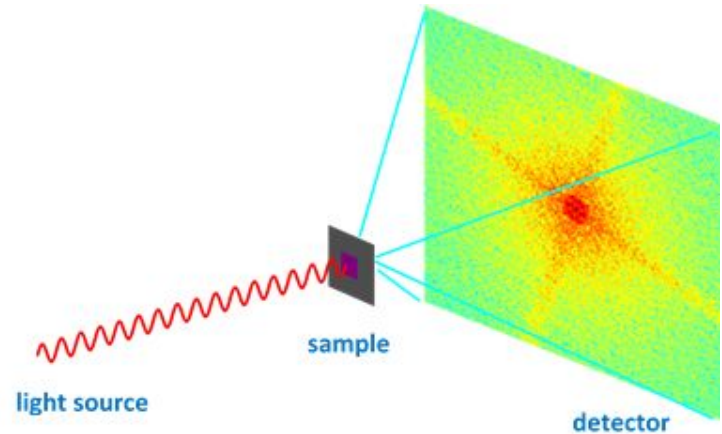
Institute for Advanced Study

November 1 @ Future Prospects of Intensity Interferometry workshop

collaborated with Neal Dalal

Phase Retrieval


- ▷ Intensity interferometer measures the Fourier amplitudes and loses the phase information.
- ▷ The problem of reconstructing a signal from its Fourier magnitude is known as phase retrieval. This reconstruction problem has a rich history and arises in many areas of engineering and applied physics.
- ▷ Coherent diffraction imaging:
 - Idea introduced in 1950s and 1980
 - Phase retrieval algorithms developed in 1970s and 1980s
 - First experimental demonstration in 1999



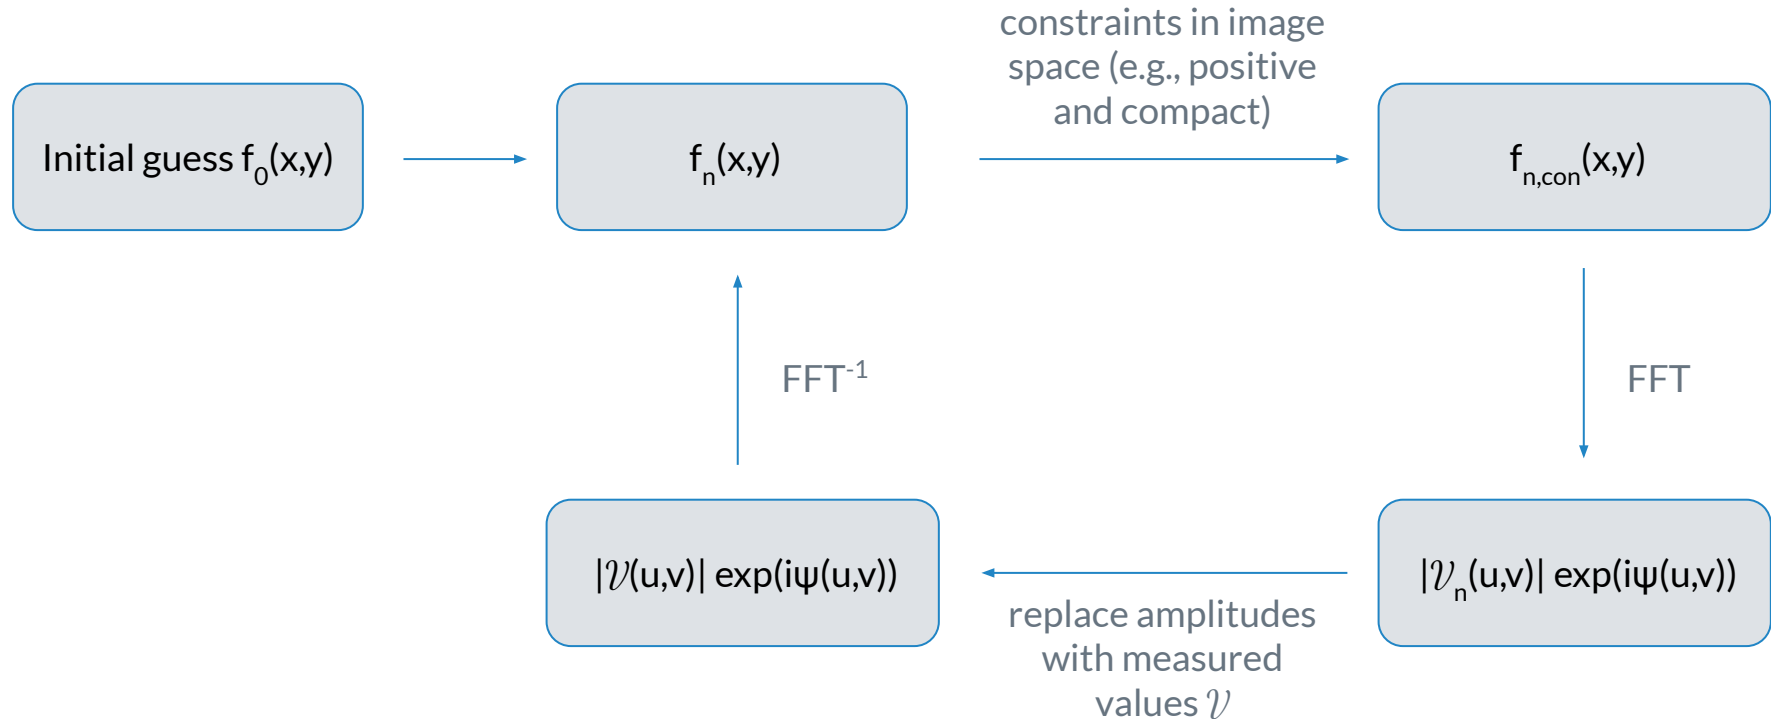
Why is Phase Retrieval possible?

- ▷ Additional constraints / regularizations / prior information are necessary to retrieve the phases
 - Positivity, smoothness, compactness, etc.


- ▷ Theorem: the signal can be uniquely reconstructed (up to some trivial transformations such as translation and inversion), if
 - the Fourier amplitudes are oversampled by at least a factor of 2 (i.e., pixel value non-zero only at the central $(L/2)^2$ pixels)
 - the signal cannot be represented by the convolution of two noncentrosymmetric functions



Iterative phase retrieval algorithm



Hybrid input-out (HIO) algorithm



However...

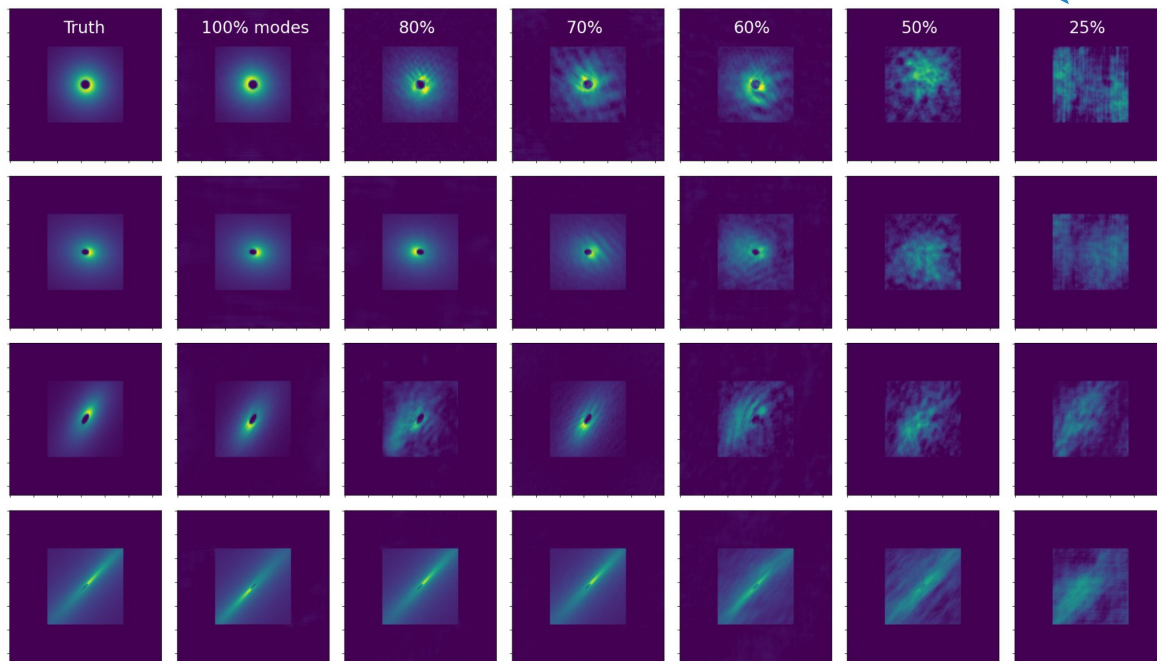
- ▷ The HIO algorithm assumes
 - Oversampling the Fourier amplitudes by at least a factor of 2 (i.e., the image is only non-zero at the central $(L/2)^2$ pixels) (also guarantees unique solution)
 - Noiseless data

- ▷ In reality, our measurements are
 - Sparsely-sampled
 - Noisy

However...

the number of measured amplitudes = the number of pixels
(because the Fourier modes are oversampled)


- ▷ The HIO algorithm assumes
 - Oversampling the Fourier amplitudes by at least a factor of 2 (i.e., the image is only non-zero at the central $(L/2)^2$ pixels) (also guarantees unique solution)
 - Noiseless data



- ▷ In reality, our measurements are
 - Sparsely-sampled
 - Noisy

However...

- ▷ The HIO algorithm assumes
 - Oversampling the Fourier amplitudes by at least a factor of 2 (i.e., the image is only non-zero at the central $(L/2)^2$ pixels) (also guarantees unique solution)
 - Noiseless data



- ▷ In reality, our measurements are
 - Sparsely-sampled
 - Noisy

Phase retrieval in a Bayesian framework

$$\log p(\mathbf{Image} \mid |\mathcal{V}|^2) = \log p(|\mathcal{V}|^2 \mid \mathbf{Image}) + \log p(\mathbf{Image}) - \log p(|\mathcal{V}|^2)$$

Likelihood function given by the measurement noise model. Usually approximated with a Gaussian


Prior function. It could either be some constraints or regularizations, or it could be learned from simulations.

- Easily incorporate measurement noise
- Uncertainty quantification with posterior sampling

Choices of prior

- ▷ HIO algorithm
 - Positive pixel values
 - Compact object (nonzero at the central $(L/2)^2$ pixels / oversampling the Fourier plane)

- ▷ The regularizations in EHT image reconstruction (RML algorithm):
 - Total flux
 - Favors images similar to a “prior image”(a circular Gaussian)
 - Sparsity (L1 norm)
 - Smoothness (total variation, total squared variation)



EHT collaboration 2019. [1906.11241]

- ▷ How to incorporate our physical knowledge (e.g., the AGN disk is thin and roughly follows a Shakura-Sunyaev profile, or predictions from GRMHD simulations) into the prior?


Deep generative priors

- ▷ A generative model aims to model the underlying probability distribution (our prior distribution) of a dataset, given independently and identically distributed (i.i.d.) samples (e.g., simulations, analytical models).
- ▷ Score-based diffusion models estimate the gradient of the log density function (score function, $\nabla \log p$).

$$\nabla \log p(\text{Image} \mid |\mathcal{V}|^2) = \nabla \log p(|\mathcal{V}|^2 \mid \text{Image}) + \nabla \log p(\text{Image})$$

Gradient of Gaussian likelihood.

Learned by diffusion model.




Experiments


We consider two experiment setups:

- Shakura-Sunyaev disk

$$I(R) = I_0 \left[e^{f(R)} - 1 \right]^{-1}$$
$$f(R) = \frac{\nu}{\nu_0(R)} = \left[\left(\frac{R_0}{R} \right)^n \left(1 - \sqrt{\frac{R_{\text{in}}}{R}} \right) \right]^{-1/4}$$




- GRMHD simulations for RIAF (Wong et al. 2022)




Low noise, densely sampled case

- ▷ The reconstructed images converge to the truth when the noise is low $\sigma_{|M|^2} = 10^{-4}$ and the Fourier space is densely sampled.




- ▷ Note that unlike HIO algorithm, here we don't need to oversample the Fourier space. The HIO algorithm fails at this noise level, even though it uses 4 times more measurements.


Changing noise levels




- ▷ The reconstruction is poorer and the algorithm is less certain with more noise. In the high noise limit, the algorithm essentially generates random samples from the prior.
- ▷ Sampling the posterior allows us to quantify the uncertainty of the image reconstruction.



Changing UV coverage




- ▷ Apply a random mask in the Fourier space. The noise level is fixed at $\sigma_{|Y|^2} = 10^{-2}$
- ▷ Preliminary results: the model works well within the UV coverage range we tested.



On GRMHD simulations

▷ GRMHD simulations of RIAF (Wong et al. 2022)

▷ The algorithm is also able to reconstruct small-scale features like rings and spirals, given sufficient SNR (fixed at $\sigma_{|v|^2} = 10^{-4}$ here) and good UV coverage.




Misspecified prior

- Reconstruct Shakura-Sunyaev disk with photon ring GRMHD prior.


- Misspecified prior introduces bias (e.g., spirals, rings) into the images. The reconstruction is a competition between the likelihood (measurement) and the misspecified prior, depending on the SNR.

- As the noise increases, the reconstructions show more features from the prior. Some features like spirals are random and averaged to 0 in posterior mean.




Misspecified prior

- ▷ Feng et al. 2024 [2406.02785] explores reconstructing EHT images with different priors:
 - CIFAR-10: natural images like cars, airplanes, animals, etc.
 - GRMHD photon-ring simulations
 - Radiatively Inefficient Accretion Flow (RIAF) models
 - CelebA: Celebrity faces
- ▷ Different (misspecified) priors introduce different biases, but there are robust features that do not depend on the choice of prior, e.g., ring structure, orientation, asymmetry, etc.



Misspecified prior

- ▷ Feng et al. 2024 [2406.02785] explores reconstructing EHT images with different priors:
 - CIFAR-10: natural images like cars, airplanes, animals, etc.
 - GRMHD photon-ring simulations
 - Radiatively Inefficient Accretion Flow (RIAF) models
 - CelebA: Celebrity faces
- ▷ Different (misspecified) priors introduce different biases, but there are robust features that do not depend on the choice of prior, e.g., ring structure, orientation, asymmetry, etc.
- ▷ The CIFAR-10 prior introduces the least amount of noise, because it covers a wide range of natural images (even though none of them look like a ring or a disk).



Conclusions

- ▷ Phase retrieval (reconstructing the image from Fourier amplitude measurement) is possible.
- ▷ Solving the phase retrieval in a Bayesian framework allows
 - Incorporation of measurement uncertainty
 - Posterior sampling for uncertainty quantification
- ▷ Deep generative diffusion models learn the prior distribution from simulations or models.
- ▷ Misspecified prior may introduce biases into the phase retrieval, but it can be reduced with broad prior choices and high SNR measurements.