Learning a Probabilistic Strategy for Computational Imaging Sensor Selection

He Sun, Member, IEEE, and Adrian V. Dalca, Member, IEEE, and Katherine L. Bouman, Member, IEEE

Abstract—Optimized sensing is important for computational imaging in low-resource environments, when images must be recovered from severely limited measurements. In this paper, we propose a physics-constrained, fully differentiable, autoencoder that learns a probabilistic sensor-sampling strategy for optimized sensor design. The proposed method learns a system's preferred sampling distribution that characterizes the correlations between different sensor selections as a binary, fully-connected Ising model. The learned probabilistic model is achieved by using a Gibbs sampling inspired network architecture, and is trained end-to-end with a reconstruction network for efficient co-design. The proposed framework is applicable to sensor selection problems in a variety of computational imaging applications. In this paper, we demonstrate the approach in the context of a very-long-baseline-interferometry (VLBI) array design task, where sensor correlations and atmospheric noise present unique challenges. We demonstrate results broadly consistent with expectation, and draw attention to particular structures preferred in the telescope array geometry that can be leveraged to plan future observations and design array expansions.

Index Terms—Computational Imaging, Optimized Sensing, Ising Model, Deep Learning, VLBI, Interferometry

1 Introduction

▼ OMPUTATIONAL imaging systems tightly integrate algorithm and sensor design, making it possible to observe phenomena previously difficult or impossible to see with traditional sensors. A common constraint in these imaging systems is that they must operate in low-resource environments. For instance, the measurements collected are severely limited due to radiation dosage in computed tomography (CT) [1], speed in Magnetic Resonance Imaging (MRI) [2] and microscopy [3], and cost in very-longbaseline-interferometry (VLBI) [4]. Imaging methods are then designed to intelligently fill in the gaps of missing information in order to recover the targeted image. We propose a framework to jointly *learn* the ideal sampling structure of non-linear, correlated measurements simultaneously with the image recovery procedure, for the purpose of sensor design. We apply this framework to study optimal VLBI telescope array geometries in different observing conditions.

A computational imaging system typically consists of two primary components: the sensing system and the image reconstruction method. The sensing system is most often described by a physics-based forward model y=f(z), which defines how the selected measurements, y, are related to the underlying image, z. These measurements often experience noise and do not contain enough information on their own to fully specify z. Therefore, imaging methods are designed to recover an image, \hat{z} , from the observed measurements by imposing pre-defined assumptions, such as image priors or hard physical constraints.

In addition to developing image reconstruction methods, optimization of the sensor is critical for producing high-quality reconstructions. In many applications the placement (i.e., sampling) of sensors controls the nature of information gathered in measurements. It is especially important that this information be optimized for a specific task when the measurement space can't be fully observed. Previous work in 3D reconstruction [5], microscopy [6], [7], and astronomical imaging [8], [9] have proposed strategies to optimize sensor placement to achieve maximally informative measurements. Seminal work in compressed sensing showed that, under particular sampling distributions, it is possible to recover a sparse signal with near perfect accuracy [10]; this idea revolutionized a number of real-world applications, including MRI [11]. However, these theoretical results only hold under certain assumptions, such as a linear forward model or limited noise. Consequently, these sampling strategies are surely underoptimized and could be improved in complex imaging systems.

The optimal sensing problem is typically idealized and considered independently from the image reconstruction problem. However, these two problems are intimately related: existing imaging algorithms inform what data we are capable of interpreting and therefore should collect, and the particular data we measure informs what imaging algorithms we should use. Inspired by this insight, we address co-designing the sensor-sampling strategy and the image reconstruction algorithm. The complexity of real-world measurements makes it difficult to tackle this co-design problem analytically. Therefore, we propose to solve a joint optimization problem using deep learning, where the optimal sampling policy is represented as a trainable joint probability distribution over candidate sensors. The proposed framework enables 1) tackling complicated sampling problems where measurements are correlated and experience substantial noise, 2) exploring multiple plausible sampling strategies simultaneously rather than focus on a single deterministic design, and 3) qualitatively understanding the importance of each sensor and correlation among sensors.

The content of the paper is organized as follows. In Sec. 2, we review related work on reconstruction algorithms and optimal sensing. In Sec. 3, we present our method that jointly optimizes the sensor-sampling pattern and image reconstruction. In Sec. 4, we demonstrate our method on the astronomical imaging VLBI array

H. Sun and K.L. Bouman are with the Department of Computing and Mathematical Sciences, California Institute of Technology, CA. F-mail: hesun@caltech.edu

A.V. Dalca is with Harvard Medical School and the Massachusetts Institute of Technology, MA.

design problem to identify which telescopes are most important to observe with in different conditions. In Sec 5, we summarize our work and discuss additional sensor-selection applications.

2 RELATED WORK

2.1 Image Reconstruction Techniques

2.1.1 Regularized Inverse Model

A common approach taken in computational imaging is to reconstruct the target image from limited sensor measurements by solving a regularized inverse problem:

$$\hat{z} = \arg\min_{z} \left[\mathcal{L}(y, f(z)) + \alpha \mathcal{R}(z) \right], \tag{1}$$

where f(z) is the forward model of the underlying imaging system, which produces measurements y from the true image z, $\mathcal{L}(\cdot)$ is inspired by the data negative log-likelihood, $\mathcal{R}(\cdot)$ is a regularization function, α is the coefficient balancing the two terms, and \hat{z} is the reconstructed image. The $\mathcal{L}(\cdot)$ and $\mathcal{R}(\cdot)$ terms quantify the image's fit with the measurements and the prior knowledge, respectively. The regularization term shrinks the solution domain, consequently making the solution unique even with an incomplete set of measurements. Typically, this inverse problem can be solved using iterative optimization algorithms when explicit image regularizer functions are imposed, such as total variation (TV) [12], maximum entropy (MEM) [13], sparsity in Fourier or wavelet transformed domain [14], or locally-learned distributions [15].

2.1.2 Learned Inverse Model

Deep learning techniques have also been used for computational imaging reconstruction by either directly training a deep reconstruction neural network [16] or inserting deep learning "priors" into model-based optimization frameworks [17]. Unlike many classic computer vision tasks, such as face or object recognition, computational imaging problems are fundamentally related to the physics of a known imaging system. Therefore, recent deep learning based reconstruction methods focus on incorporating known physics to improve computational efficiency, reconstruction accuracy, and robustness. For example, this can be done by adding physics-informed regularization to the neural network loss function [18], only learning the residuals between ground truth and physics-based solutions [19], or unrolling the physics-based optimization procedure as a neural network [20], [21].

2.2 Sensor-Sampling Strategies

2.2.1 Analytic Strategies

Conventionally, sparse sensor sampling is achieved by adopting the Nyquist-Shannon sampling theorem. For instance, in imaging applications with Fourier component measurements, such as VLBI and MRI, since the images are approximately band-limited, the signal sampling rate in the Fourier domain (e.g., κ -space) can be reduced to twice the maximum frequency [22]. This observation can be used to define metrics that characterize sampling quality independent of image reconstruction techniques [23]. Modern advances in the theory of compressed sensing demonstrate that, when the underlying image is sparse, certain sampling schemes can achieve the same quality reconstruction with an even smaller number of samples [10], [24]. However, both Nyquist sampling and compressed sensing typically consider only linear forward models and assume no noise or only *i.i.d.* Gaussian noise, which is unrealistic for many real-world problems.

2.2.2 Learned Strategies

Recently, approaches have been proposed to explore the joint optimization of a sensing strategy and image reconstruction method using deep neural networks. This idea has been used to design the color multiplexing pattern of camera sensors [25], the LED illumination pattern in Fourier ptychography [7], the optical parameters of a camera lens [26], and the exposure pattern in video compressed sensing [27]. However, all these approaches yield only a single deterministic sensing pattern that ignores the possibility of multiple, equally good sensing strategies. In a recent study of "compressed sensing" MRI, the sensing sampling strategy in the κ -space is modeled as Bernoulli distributions, so that multiple sensing patterns are explored simultaneously [28]. This formulation helps characterize which sampling frequencies are important for different anatomy. However, none of these prior approaches address complicated non-Gaussian noise, or explicitly model the correlations among measurements.

3 PROPOSED METHOD

We propose a learning method that jointly optimizes the sensorsampling distribution and the image reconstruction method for forward models that contain correlated measurements and complicated noise. In Sec. 3.1, we discuss the problem setup. In Sec. 3.2 we introduce the sensor-sampling distribution, formulated as an Ising model. In Sec. 3.3 we discuss the image reconstruction subnetwork architecture.

3.1 Joint Sensing and Imaging Optimization

We formulate the joint sensing and imaging optimization as a deep autoencoder. This autoencoder is trained to identify an efficient image representation that is then used to make limited observations, from which an image can be recovered. The encoder represents the sensing strategy and the decoder represents the image reconstruction method. Both the encoder and decoder are formulated as deep neural networks in this framework. However, to ensure the sensing system design obeys physical constraints, the encoder samples from all possible measurements generated by the physics simulator.

This optimization problem can be written as

$$\underset{\theta,\omega}{\operatorname{arg \, min}} \quad \mathbb{E}_{M \sim p_{\theta}(M)} \left[\sum_{j=1}^{n} s(\hat{z}_{j}, z_{j}) + \lambda \mathcal{R}(M) \right],$$

$$s.t. \quad \hat{z}_{j} = A_{\omega} \{ g(M, f(z_{j})) \},$$

$$(2)$$

where M is a sensor-sampling pattern, whose distribution is $p_{\theta}(\cdot)$ parameterized by θ , $s(\hat{z}_j, z_j)$ is a function defining the similarity between the reconstruction result \hat{z}_j and the true image z_j in the n image training set, $\mathcal{R}(\cdot)$ is the regularization on the sensing samples, $A_{\omega}\{\cdot\}$ is the image reconstruction neural network parameterized by ω , and $g(\cdot,\cdot)$ is the sensing function that produces the observed measurements. In this loss function, the first term quantifies the image reconstruction performance, while the second term quantifies desired properties of the sensor-sampling pattern (e.g., sparsity; see Sec. 3.2.3). These two terms are balanced by the hyper-parameter λ . The sampling pattern $M \in [0,1]^K$ is a binary vector that indicates which of the K sensor sources is selected or excluded from observations. This is equivalent to selecting elements from the forward model's vector of measurements, i.e.

$$g(M, f(z_j)) = S(M) \odot f(z_j), \tag{3}$$

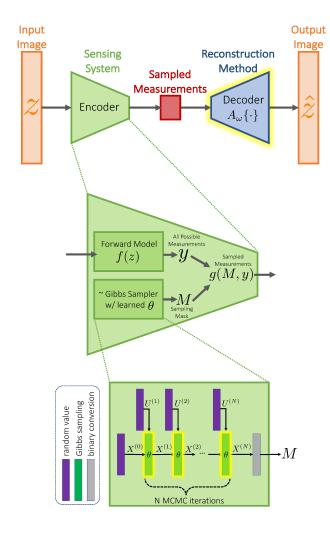


Fig. 1. The architecture of the proposed joint sensing and imaging optimization problem. We draw a connection between optimizing a computational imaging system and learning the parameters of an autoencoder: the encoder consists of a known physics-based forward model that generates all possible measurements and a sensor-sampling block that chooses among those measurements; the decoder drives image reconstruction from measurements. We combine the proposed decoder and physically-motivated encoder into a single, differentiable neural network. By maximizing the similarity between the input and output images, the sensor-sampling strategy and image reconstruction are jointly optimized. The trainable blocks, sensor-sampling distribution and reconstruction method, are highlighted in yellow, where the corresponding trainable parameters are θ and ω .

The architecture of the sampling sub-network is inspired by the Gibbs sampling procedure. Each Gibbs sampling layer acts as an MCMC iteration: it takes samples $X^{(i-1)}$ from the previous layer and generates more likely samples $X^{(i)}$ based on the states' joint distributions (parameterized by θ). The input to the sub-network is a set of uniform random numbers $U^{(i)}$ and an initial random set of states, $X^{(0)}$, $\in [-1,1]$. After N Gibbs layers, the states are converted to approximately binary by using a sigmoid function with large slope. Mathematical representations of the successive Gibbs sampling layers and the binary conversion layer are shown in Eq. (13) and Eq. (14), respectively.

where $S(\cdot)$ is the binary selection function based on the mask M, whose dimensionality is the same as the forward model $f(\cdot)$. This formulation is a generalization of [28], which used a similar optimization designed specifically for MRI acquisition.

Based on Eq. (2) and Eq. (3), we build a joint sensing and imaging autoencoder, as shown in Fig. 1. The trainable blocks in our neural network are the sensor-sampling distribution $p_{\theta}(\cdot)$

(from which the mask M is sampled) and the reconstruction network $A_{\omega}(\cdot)$, where θ and ω are the trainable network parameters. In the following sections we describe how to build Eq. (2) as a single, differentiable neural network and solve for θ and ω simultaneously.

3.2 Physics-Constrained Encoder

The goal of the encoder in Fig. 1 is to select measurements from the physics-based forward model to pass on to the decoder. The forward model and procedure for masking measurements are both well defined; the challenge is to design a network architecture that can sample a mask from a parameterized distribution. In order to capture the correlations among observations, we model the proposed sensor-sampling distribution $p_{\theta}(\cdot)$ in Eq. (2) as a fully-connected, Ising model (binary Markov random field). In this section we explain the Ising model sensor-sampling distribution and describe how to sample from this model using a differentiable network architecture.

3.2.1 Ising model

In an Ising model, the state of an n-element system is represented by a set of binary random variables, $X=(x_1,\cdots,x_n), x_j\in\{+1,-1\}, \forall j\in[1,n]$. The total energy of a fully-connected Ising model is given by the Hamiltonian function,

$$H_{\theta}(X) = -\sum_{j} \theta_{jj} x_j - \sum_{j < k} \theta_{jk} x_j x_k, \tag{4}$$

where θ_{jk} are the Ising model parameters. The first term models the energy of single elements and the second term models the potential from the interactions between pairs of elements. The state's probability can be derived according to the Boltzmann distribution

$$p_{\theta}(X) = \frac{\exp[-H_{\theta}(X)]}{Z_{\theta}},\tag{5}$$

where Z_{θ} is the partition function,

$$Z_{\theta} = \sum_{X} \exp[-H_{\theta}(X)]. \tag{6}$$

The lower the total energy of a state, the higher the probability it will occur.

We use a fully-connected Ising model to model the presence (+1) or absence (-1) of a particular sensor, and capture correlations among sensors. Parameter θ_{jj} characterizes how important it is to include the j-th sensor, and θ_{jk} characterizes how the j-th and the k-th sensor are positively or negatively correlated: if $\theta_{jk}>0$, observing the j-th and the k-th sensor are favored or disfavored together; while if $\theta_{jk}<0$, the j-th sensor is disfavored when the k-th is selected, or visa versa. Although an Ising model only explicitly defines pairwise correlations, higher order correlations among sensors are implicitly captured by the model. These higher order correlations can be partially extracted by interpreting the Ising model as an affinity graph and performing clique analysis on the pruned edges (see an example in Sec. 4.3.3).

Solving the optimization problem in Eq. 2 requires sampling the Ising model according to parameters $\theta = \{\theta_{jk} | \forall j, k \in [1,n]\}$. However, unlike simple distributions such as Bernoulli, Uniform, or Gaussian, it is difficult to sample an Ising model directly, because the binary variables are correlated and the number of possible states grows exponentially with the dimension of elements. To build a computationally efficient and fully differentiable sampling network, we propose a Markov chain Monte Carlo (MCMC) inspired method.

3.2.2 Differentiable Gibbs Sampling Sub-Network

We design a sub-network to perform successive MCMC sampling iterations of an Ising model (parameterized by θ), which enables sampling masks from $p_{\theta}(M)$ and facilitates computing Eq. (2).

MCMC methods have been widely used for approximating complicated statistical distributions. In a commonly used MCMC method, Gibbs sampling, one samples from the joint probability distribution by conditionally sampling each variable in turn until convergence. The detailed procedure of Gibbs sampling for a binary state space can be summarized as:

Randomly draw an initial sample

$$X^{(0)} = (x_1^{(0)}, \cdots, x_n^{(0)}),$$
 (7)

from an independent Bernoulli distribution.

Compute the next sample $X^{(i)}$ based on $X^{(i-1)}$ with each component $x_j^{(i)}$ sequentially updated using its conditional distribution given the other states

$$p_j^{(i)} = p_{\theta}(x_j^{(i)} = 1 | x_1^{(i)}, \cdots, x_{j-1}^{(i)}, x_{j+1}^{(i-1)}, \cdots, x_n^{(i-1)}).$$
(8)

For each $p_i^{(i)}$, draw a random number from a uniform distribution

$$u_j^{(i)} \sim U(0,1).$$
 (9

If $u_j^{(i)} < p_j^{(i)}$, assign $x_j^{(i)} = +1$; otherwise, assign $x_j^{(i)} = -1$. Repeat step 2 for N iterations.

Gibbs sampling is often used to sample from Ising models, as the conditional probability of an Ising model can also be described as an Ising model:

$$H_{\theta}(\lbrace x_{j} \rbrace \in \Lambda_{u} | \lbrace x_{l} \rbrace \in \Lambda_{k})$$

$$= -\sum_{\substack{j \ x_{j} \in \Lambda_{u}}} (\theta_{jj} + \sum_{\substack{l \ x_{l} \in \Lambda_{k}}} \theta_{jl} x_{l}) x_{j} - \sum_{\substack{j < k \ x_{j}, x_{k} \in \Lambda_{u}}} \theta_{jk} x_{j} x_{k}$$

$$= -\sum_{\substack{j \ x_{j} \in \Lambda_{u}}} \theta'_{jj} (\lbrace x_{l} \rbrace \in \Lambda_{k}) x_{j} - \sum_{\substack{j < k \ x_{j}, x_{k} \in \Lambda_{u}}} \theta_{jk} x_{j} x_{k}$$

$$(10)$$

where Λ_k and Λ_u are the known and unknown/to-be-updated subsystems respectively, and θ'_{ij} defines the new Ising model parameter conditioned on the known subsystem. Eq. (10) indicates the conditional distribution of a single element is a Bernoulli distribution, making it possible to implement each Gibbs sampling iteration as a neural network layer.

In order to sample a mask M from the distribution $p_{\theta}(M)$, we design a neural network where each dense layer corresponds to an iteration of Gibbs sampling for an Ising model. Each Gibbs sampling layer can be mathematically described as a function of

$$x_j^{(i)} = \operatorname{sgn}\left(\sigma(2(\sum_{k < j} \theta_{jk} x_k^{(i)} + \theta_{jj} + \sum_{k > j} \theta_{jk} x_k^{(i-1)})) - u_j^{(i)}\right),\tag{11}$$

where $sgn(\cdot)$ is a sign step function and $\sigma(\cdot)$ is a sigmoid function,

$$sgn(y) = \begin{cases} -1 & y \le 0 \\ +1 & y > 0 \end{cases}, \quad \sigma(y) = \frac{1}{1 + \exp(-y)}. \quad (12)$$

The initial sample $X^{(0)}$ is initialized using uniform random values, $x_j^{(0)}=\mathrm{sgn}(0.5-u_j^{(0)}).$ Eq. (11) defines a layered Ising

sampling function, which converges to a sample from the Ising model after a sufficient number N of iterations. However, the sign function in Eq. (11) has no gradient (zero or undefined), rendering back-propagation strategies impossible. To address this, we replace the sign function with a hyperbolic tangent function

$$x_{j}^{(i)} = \tanh_{s_{1}} \left(\sigma(2(\sum_{k < j} \theta_{jk} x_{k}^{(i)} + \theta_{jj} + \sum_{k > j} \theta_{jk} x_{k}^{(i-1)})) - u_{j}^{(i)} \right), \tag{13}$$

where $tanh_{s_1}(a) = tanh(s_1a)$. Using a moderate s_1 slope enables us to approximately binarize the states to [-1, 1] between successive Gibbs layers, while still allowing gradients to propagate through the network.

To obtain the sampling mask, M, we apply a strong sigmoid function to the final sampling layer of the Ising model to convert the states from approximately [-1, 1] to nearly binary states [0, 1]:

$$M = \sigma_{s_2}(X^{(N)}) = \sigma(s_2 X^{(N)}), \tag{14}$$

where s_2 is a slope of the sigmoid function that should be large enough to approximate the Heaviside function well.

For notational simplicity, we represent this sampling procedure as a nested function,

$$M \sim p_{\theta}(M)$$

$$= \sigma_{s_{2}}(X^{(N)})$$

$$= \sigma_{s_{2}}(G_{\theta}(U^{(0)}, \cdots, U^{(N)})),$$
(15)

with Ising model coefficient θ as its parameters, uniformly distributed random numbers, $\{U^{(0)}, \cdots, U^{(N)}\}$ as inputs, and the [0,1] binary samples in M as outputs. The architecture of the Gibbs sampling neural network is shown at the bottom of Fig. 1.

3.2.3 Ising Model Regularization

To encourage sparse samples and promote sample diversity (in order to explore multiple observational strategies), we define the regularization on Ising samples $(M \sim p_{\theta}(M))$ as consisting of two terms:

$$\mathcal{R}(M) = \lambda_1 ||M||_1 - \lambda_2 H_{\theta}(X^{(N)}), \tag{16}$$

where the first term is a ℓ_1 regularization on the number of participating sensors (sparsity loss), and the second term is the negative Hamiltonian of the Ising samples (diversity loss). The Hamiltonian of an Ising model is a good approximation of the distribution entropy as it only differs in its log partition function,

$$S_{\theta} = -\sum_{X} p_{\theta}(X) \log p_{\theta}(X)$$

$$= -\sum_{X} p_{\theta}(X) [-H_{\theta}(X) - \log Z_{\theta}]$$

$$= \mathbb{E}_{X \sim p_{\theta}(X)} [H_{\theta}(X)] + \log Z_{\theta}.$$
(17)

Therefore, maximizing the Hamiltonian typically increases the system entropy, thus diversifying the samples from our Ising model. Substituting Eq. (16) into Eq. (2) and approximating the expectation using the empirical average via MCMC samples, we derive a stochastic formulation of the joint optimization problem under Ising sampling,

$$\arg \min_{\theta,\omega} \quad \sum_{j=1}^{n} s(\hat{z}_{j}, z_{j}) + \lambda_{1} || M_{j} ||_{1} - \lambda_{2} H_{\theta}(X_{j}^{(N)}),$$

$$s.t. \quad \hat{z}_{j} = A_{\omega} \{ S(M) \odot f(z_{j}) \},$$

$$M_{j} = \sigma_{s}(X_{j}^{(N)}),$$

$$X_{j}^{(N)} = G_{\theta}(U_{j}^{(0)}, \cdots, U_{j}^{(N)}),$$

$$U_{j}^{(0)}, \cdots, U_{j}^{(N)} \sim U(0, 1).$$
(18)

3.3 Image Reconstruction Decoder

The image reconstruction decoder sub-network is not restricted to any particular neural architecture, but should have enough model capacity to be able to reconstruct with multiple sensing scenarios simultaneously. For imaging problems with a complicated forward model and noise, physics-based neural networks (PbNN) [7] and unrolled neural networks [20], [29] have proven successful. However, with sufficient training data, general deep learning frameworks such as the U-Net and its variants [30], [31] often perform sufficiently well. The U-Net combines local and global information during image reconstruction by incorporating hierarchical down-sampling and up-sampling blocks and multiple skip connections in its architecture. In our VLBI case study, two variants of the U-Net are used for reconstruction. We discuss the details of their architectures in Sec. 4.2.3.

4 CASE STUDY: VLBI ARRAY DESIGN

Very-Long-Baseline-Interferometry (VLBI) is a radio interferometry technique for high-resolution astronomical imaging. In VLBI, telescopes are linked to form an imperfect, virtual telescope with a maximum resolution defined by the longest baseline between participating telescopes [32]. Recently, this technique was used by Event Horizon Telescope (EHT) to capture the first image of a black hole, the image of M87* [33], by joining telescopes (observing at 230 GHz) from across the globe. Building on this recent success, the EHT is continuing to improve its telescope array by adding new sites, in order to improve M87*'s image quality as well as image additional astronomical sources, such as the Milky Way's black hole, Sagittarius A*.

It is important to carefully study which new sites should be added to the next-generation EHT to best improve future results. Adding a new telescope to the existing EHT VLBI array is extremely expensive: equipping a pre-built telescope to work as part of the array costs over two million dollars, and the construction of a new telescope can easily cost over 10 million dollars [34]. Therefore, limited budgets require selecting only a handful of new sites that can be added to the future array. Since only a select number of ground sites are suitable for constructing telescopes that can observe at 230 GHz (e.g., due to altitude and weather), the problem of choosing telescope locations for the next-generation EHT reduces to a site selection problem. Studying the impact of a telescope on imaging quality is not only important for planning an EHT expansion but also becomes a powerful tool during an observing campaign; when coordinating future observations it is essential to evaluate the impact of losing a particular telescope – should costly observations proceed if weather is poor or there are instrumental problems at a particular site?

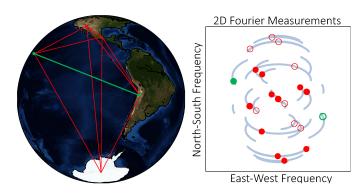


Fig. 2. VLBI baselines and measurements. The left figure indicates the baselines between telescopes in the EHT array at a particular time, where the Earth is being viewed from the direction of Sgr A*. The right figure shows the corresponding 2D spatial-frequency coverage of measurements collected by the telescopes. Each dot represents a visibility measurement, where the 2D spatial-frequency of the visibility is related to the projected baselines between two telescopes at a given time. An example baseline and its Fourier measurement (two symmetric dots due to the constraint the the image is real-valued) are marked in green. As Earth rotates, the projected baseline vectors change, resulting in a more complete frequency coverage (marked with light blue) collected over a night.

We optimize the design of a VLBI array by solving a site selection problem. The goal is to decide which sites from a given list we should build into the array, and which sites can be excluded without substantial loss to the reconstructed image. Although many factors are ultimately important for site selection (e.g., construction cost, weather stability), in this paper we focus solely on the location of sites, as the array geometry usually has the largest influence over resulting image quality. As explained in Sec. 4.1, VLBI measurements are produced from data collected at multiple telescopes, so it is important to consider how telescopes interact with one another. We capture the telescope correlations through the fully-connected binary Ising model, and optimize the telescope sampling distribution jointly with the image reconstruction method using the proposed approach.

4.1 Problem Setup

In radio interferometry (e.g., VLBI), data simultaneously collected from a pair of telescopes is used to produce a single measurement. Each measurement corresponds to a two-dimensional Fourier component of the image, with the frequency related to the baseline between the two participating telescopes, projected in the direction of the observed astronomical source. This measurement is typically referred to as a *visibility*. For a K telescope array, there are $\binom{K}{2}$ possible visibility measurements at a single time. Due to Earth's rotation, the projected baselines change over time, causing sparse elliptical measurements in the spatial-frequency domain to be collected over the course of a night (see Fig. 2).

The measurement equation for each visibility, V, is given by

$$V_{p,q}^{t} = g_{p}^{t} g_{q}^{t} \exp\left[-i(\phi_{p}^{t} - \phi_{q}^{t})\right] F_{p,q}^{t} z + n_{p,q}^{t}, \tag{19}$$

for telescope pair (p,q). $F_{p,q}^tz$ extracts a Fourier component from the image z corresponding to the baseline between telescopes p and q at time t. Each measurement experiences time-dependent telescope-based gain, g, and phase error, ϕ , and baseline-based thermal noise, n, where $n_{p,q}^t \sim \mathcal{N}(0, \nu_{p,q}^2)$. The standard devi-

TABLE 1						
Telescope sites in "EHT+" array.						

Sites	Location	SEFD	
PV*	Pico Veleta, Spain	1400	
PDB	Plateau de Bure, France	1500	
ALMA*	Atacama Desert, Chile	90	
APEX*	Atacama Desert, Chile	3500	
LMT*	Sierra Negra, Mexico	600	
SMT*	Mt. Graham, Arizona	5000	
SPT*	South Pole	5000	
OVRO	Owens Valley, California	10000	
JCMT*	Maunakea, Hawai'i	6000	
SMA*	Maunakea, Hawai'i	4900	
KP	Tucson, Arizona	10000	
GLT	Thule Air Base, Greenland	10000	

ation of the thermal noise depends on each telescope's "system equivalent flux density" (SEFD):

$$\nu_{p,q} \propto \sqrt{SEFD_p \times SEFD_q},$$
 (20)

where a higher SEFD indicates the measurements using that telescope contain more thermal noise (i.e., lower SNR) [32]. When atmospheric turbulence is not considered, g=1 and $\phi=0$. Although the absolute gains g can often be sufficiently calibrated, phase errors due to atmospheric turbulence cannot be calibrated a priori since ϕ varies rapidly and is uniformly sampled from $[0,2\pi)$ [32]. Thus, in the case of atmospheric error, recovering an image amounts to solving a phase retrieval problem.

In the case of atmospheric noise, data products called closure quantities can be used to constrain the image reconstruction [35]. Closure phase is obtained by multiplying three visibilities from telescopes in a closed loop:

$$C_{p,q,b}^t = \angle \left(V_{p,q}^t V_{q,b}^t V_{b,p}^t \right) \tag{21}$$

Since the atmospheric errors are site dependent, $C_{p,q,b}$ is robust to any phase errors ϕ . Data products called closure amplitudes are robust to gain errors [35]. Since amplitudes can often be much better calibrated than phases, in this paper we do not consider closure amplitudes but make use of visibility amplitudes, $|V_{p,q}^t|$.

4.2 Implementation Details

4.2.1 Telescope Arrays

In this paper, we consider two arrays of radio telescopes: "EHT+" and "FUTURE" (Fig. 3). "EHT+" includes twelve telescopes sites, as listed in Table 1. Eight of these sites (marked with a star) were used by the EHT in 2017 [33], while the other four sites host existing radio telescopes that plan to eventually join the EHT. Each telescope's SEFD (i.e., noise coefficient) is reported in the table. The "FUTURE" array consists of nine additional sites that do not currently host a suitable telescope, but have sufficient atmospheric conditions for a constructed telescope to successfully observe at 230 GHz with the EHT.¹

1. Our analysis conservatively assumes each new telescope in the "FUTURE" array has an SEFD of 10000.

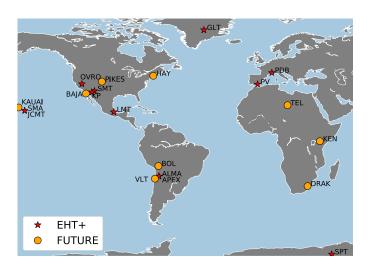


Fig. 3. Site map of potential future EHT telescope locations. Twelve sites ("EHT+") marked with blue stars are existing telescopes currently participating in or planning to join the EHT. The other nine sites ("FUTURE"), marked with orange dots, are potential locations where new telescopes could be added. "FUTURE" sites are selected as locations that can observe at the necessary 230 GHz (1.3 mm wavelength) observed by the EHT.

4.2.2 Ising Model Sampler

The VLBI forward model $f(\cdot)$ produces as a vector of complex visibilities or visibility amplitudes and closure phases. The mask, M, is a binary vector representing whether each telescope is selected for observation. We define the masking function for each VLBI measurement in $f(\cdot)$ as $S_{p,q}(M) = M_p M_q$ for a visibility $V_{p,q}$ or visibility amplitude $|V_{p,q}|$, and $S_{p,q,b}(M) = M_p M_q M_b$ for a closure phase $C_{p,q,b}$. The mask is generated from the Ising model sampler in Section 3.2, which is built using five MCMC layers. The sampling order of telescopes for each MCMC layer is randomly shuffled for each training trial. In Section 4.3, we run the joint optimization five times for each experiment to approximate the mean and standard deviation of the learned Ising model parameters. The slopes for binary conversion layers (Eq. 13 and Eq. 14) are empirically set to $s_1=3$ and $s_2=10$.

4.2.3 Image Reconstruction Architectures

Two different neural network architectures (Fig. 4) are used for the image reconstructions (Section 3.3), with input corresponding to either complex visibilities or visibility amplitudes and closure phases, as shown in Fig. 4. Both networks output a 32×32 pixel image.

In the case of no atmospheric turbulence, Reconstruction network A is used (Fig. 4). The input to this network are complex visibilities. The network architecture consists of one fully connected layer and a U-Net (four down-size layers + four upsize layers). Ideally, the fully connected layer transforms the Fourier measurements from the frequency domain back to the image domain. Then, the U-Net removes aliasing effects caused by missing data and noise. Since the image's phase information is preserved when there is no atmospheric turbulence, we can define the image similarity loss, $s(\cdot,\cdot)$ as an ℓ_1 norm of the difference between the reconstructed and blurred original image:

$$s(\hat{z}_{i}, z_{i}) = \|\hat{z}_{i} - \mathcal{K}(res) * z_{i}\|_{1},$$
 (22)

where K is a Gaussian kernel that defines our desired reconstruction resolution. The nominal resolution of our telescope

Reconstruction Network A U-Net ℓ Los

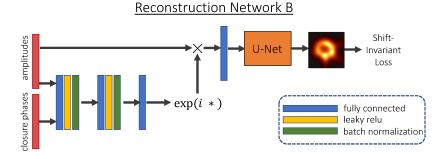


Fig. 4. VLBI reconstruction sub-network architectures. Reconstruction Network A is designed to recover an image from complex visibility measurements that experience, at most, baseline-varying Gaussian thermal noise. Reconstruction Network B is designed to recover an image when measurements experience corrupting atmospheric phase error in addition to thermal noise; since the absolute phase information in the complex visibilities are lost in this case, the network takes visibility amplitudes and closure phases as input (rather than complex visibilities) and uses a shift-invariant loss (rather than a simpler ℓ_1 loss). For each network variant we show a single representative test reconstruction obtained from measurements, which corresponds to the ring truth image shown in Fig. 7.

array ($\approx 25~\mu$ as full-width-half-max (FWHM)) is found using the longest baseline. Image reconstructions that recover structure smaller than the nominal resolution are super-resolving the target.

In the case of atmospheric turbulence, reconstruction network B is used (Fig. 4). The input is constructed from visibility amplitudes and closure phases (rather than complex visibilities). The reconstruction network architecture is designed through physical intuition in order to handle the more challenging phase-retrieval problem. First, the visibility amplitudes and closure phases are sent to three general dense layers whose purpose is to disentangle the corrupted visibility phases. Second, the disentangled phases are combined with the visibility amplitudes to produce modified complex visibilities. These modified complex visibilities are then passed to the same architecture used above (Dense + U-Net). Since the absolute phase is lost in the presence of atmospheric noise, we use a shift-invariant loss that represents the negative inner product between the true and reconstructed images:

$$s(\hat{z}_j, z_j) = 1 - \frac{\max\{\hat{z}_j * (\mathcal{K}(res) * z_j)\}}{\|\hat{z}_j\|_2 \|\mathcal{K}(res) * z_j\|_2}.$$
 (23)

4.2.4 Training

We train our joint optimization problem using images from the "Fashion-MNIST" dataset [36], and validate on the rest of "Fashion-MNIST", classic "MNIST", and simple geometric model images meant to resemble a black hole. Each image is 32×32 pixels and corresponds to a $100~\mu$ -arcsecond field of view (FOV) with a total flux of 1 Jansky (Jy), unless otherwise specified. The training set is augmented by introducing local deformations and random rotations to the original "Fashion-MNIST" images. Data augmentation is found to be especially important for training reconstruction network B to avoid overfitting. Each network is trained until convergence with Adam in TensorFlow [37], requiring between 50-300 epochs with a learning rate of 10^{-3} . VLBI visibity measurements are produced using the eht-imaging Python library [35].

4.3 Experiments

We demonstrate the proposed optimization framework in a variety of different VLBI scenarios. The purpose of the proposed method is to learn sensor-sampling designs in complex scenarios where intuition is limited. We choose to show some examples in simple

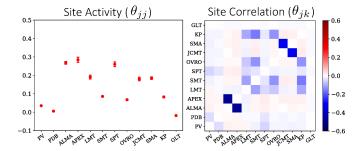


Fig. 5. The learned sensor-sampling Ising model parameters. In the VLBI array design problem, θ_{jj} represents the activity of each site (roughly proportional to how frequently a telescope is sampled), and θ_{jk} represents the correlation between each pair of telescopes. Representative hyper-parameters (explained in Sec. 4.3) are used in this optimization: the "EHT+" telescope array, science target Sgr A*, $\lambda_1=\lambda_2=0.005,\,0.75\times$ the nominal resolution target resolution, and reconstruction network A with complex visibilities. No noise is introduced on the measurements so as to purely investigate the influence of baseline geometry on the recovered distributions.

scenarios where intuition can be used to confirm that our method provides reasonable results (e.g., Section 4.3.2). However, what is most interesting are the more realistic, complicated cases that cannot be as easily judged from intuition (e.g., Section 4.3.3).

Fig. 5 shows a representative Ising model learned from our proposed method. It has been trained using the complex visibility reconstruction network, with representative coefficients $\lambda_1=\lambda_2=0.005$. In this example, we considered the "EHT+" array with the black hole Sagittarius A*(Sgr A*) as the science target. We define $\mathcal{K}(res)$ in $s(\cdot,\cdot)$ as corresponding to a Gaussian kernel with a FWHM of $0.5\times$ the nominal beam. To investigate the influence of baseline geometry on the recovered distributions, we did not include any noise on the measurements.

Note that in the site correlation plot (θ_{jk}) , the two pairs of telescopes at same location, ALMA-APEX in Chile and JCMT-SMA in Hawaii, are both highly negatively correlated. This agrees with intuition, as the baselines from any telescope to either colocated telescope will produce identical measurements (when ignoring noise). The parameters, θ_{jj} , represent the activity of each telescope site, and is roughly proportional to how frequently a telescope is sampled. As can be seen, the most important sites in the EHT network for Sgr A* observations are ALMA, APEX,

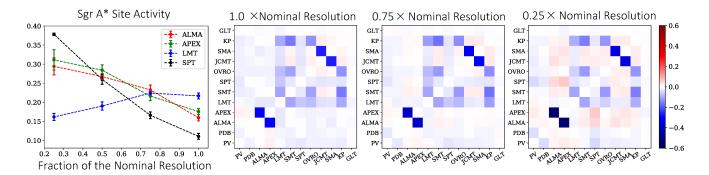


Fig. 6. Effect of target reconstruction resolution on array design. Higher resolution requirements (higher resolution corresponds to a smaller fraction of the nominal resolution) increase the learned activity of sites that participate in long baselines, such as SPT, but decrease the activities of sites that participate in short baselines, such as LMT. The long baseline telescope pairs (SPT-ALMA, SPT-APEX, SPT-JCMT and SPT-SMA) become more positively correlated with increased resolution, while zero baseline pairs (APEX-ALMA and JCMT-SMA) become more negatively correlated.

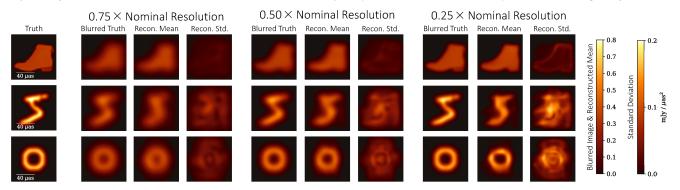


Fig. 7. Image reconstructions for different resolution requirements. Reconstruction Network A (refer to Fig. 4) was trained as part of the full proposed network shown in Fig. 1 at different target resolutions. We show the results of the reconstruction sub-network in three test cases at different target resolutions. The reconstruction mean and standard deviation are computed based on 1000 samples from the learned sensor-sampling distribution. Each reconstruction sample uses a different mask, $M \sim p_{\theta}(M)$, corresponding to a different set of observing telescopes; the sub-network must be able to handle multiple telescope sampling patterns simultaneously. Resolution requirements are specified during training by changing $\mathcal{K}(res)$ in the image similarity loss. The resolution of the array, as defined by the longest baseline (i.e., highest spatial-frequency), is defined as the nominal resolution. Recovering images at a fraction of this resolution (e.g., 0.5 \times and 0.50 \times the nominal array) require super resolving the target.

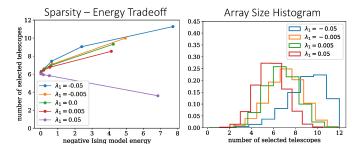


Fig. 8. Effect of regularization weights. In the left sparsity-diversity tradeoff figure, we show that (1) sparsity regularization (λ_1) controls the average number of telescopes selected, and (2) as diversity regularization (λ_2) is increased, the probability a telescope is selected approaches 50%, resulting in the average number of telescopes converging to 6 out of 12. The right figure shows the array size (number of telescopes) histogram for a fixed diversity regularization weight $(\lambda_2=0.005)$ and changing sparsity regularization weights $(\lambda_1=-0.05,-0.005,0.005,0.005)$. Larger sparsity regularization drives the sensor-sampling distribution to learn to sample fewer telescopes, as expected.

SPT and LMT, which produce measurements covering low to high spatial-frequency. To the contrary, since GLT provides no measurements (it is occluded by Earth since Sgr A* is in the southern hemisphere), it is the least important site and has a

negative activity parameter.

4.3.1 Effect of Sparsity and Diversity Regularization

We conduct a grid search on the sparsity and diversity regularization coefficients and explore their influence on the learned optimal sensor-sampling distributions. In our experiments, the diversity regularization coefficient (λ_2) ranges from 0.001 to 0.05, and the sparsity regularization coefficient (λ_1) ranges from -0.05 to 0.05. Negative sparsity regularization favors observing with more telescopes. Apart from the regularization weights, we use the representative parameters described in the beginning of Section 4.3.

Fig. 8 shows the diversity-sparsity relationship recovered from different combinations of regularization weights. The diversity regularizer characterizes the entropy of the Ising model: when the diversity (i.e., the Ising model energy) is large (large λ_2), all telescope sites have around a 50% chance of being selected or excluded. Therefore, the mean number of selected telescopes gradually converges to six out of twelve with increasing λ_2 . As shown by the array size histogram, increasing the sparsity regularizer (λ_1) results in a fewer number of sampled telescopes.

4.3.2 Effect of Reconstruction Resolution

We investigate the optimal sensor-sampling for different reconstruction resolution requirements by optimizing the joint sensing

and imaging network at four different fractions (1, 0.75, 0.5 and 0.25) of the nominal resolution in $\mathcal{K}(res)$. Apart from $\mathcal{K}(res)$, we use the representative parameters described in the beginning of Section 4.3. Fig 6 reports the activity of four telescopes (ALMA, APEX, LMT and SPT) that change most with the target resolution, and the learned correlation matrix at 1x, 0.75x and 0.25x the nominal resolution. The correlation matrix at 0.5x the nominal resolution can be found in Fig. 5.

As the resolution becomes higher (right to left in the activity plot), LMT becomes increasingly less important, but ALMA, APEX and especially SPT become more important. This agrees with intuition, as SPT baselines measure the highest spatial-frequency information for Sgr A* and LMT is contained in many of the shorter baselines probing low spatial-frequencies. The correlation coefficients between short-baseline pairs (such as APEX-ALMA, JCMT-SMA) become lower at higher resolution, while the correlation coefficients between long-baseline pairs (such as SPT-APEX, SPT-ALMA, SPT-JCMT and SPT-SMA) become higher. This satisfies physical intuition as long baselines are most important for high-resolution reconstructions.

Fig. 7 shows the blurred images and reconstruction results at different resolutions. We select one example from each type of validation dataset (Fashion-MNIST, MNIST and geometric models). More precisely, we show the reconstruction mean and standard deviation obtained by taking 1000 telescope samples from each learned model. The higher resolution reconstructions are more challenging as they attempt to super-resolve the image, and as expected they have larger standard deviations. Although the goal of this work is not to develop a new image reconstruction method, we find that the proposed neural network appears to produce promising results for the simplified VLBI reconstruction problem, and is powerful enough to be able to tackle a variety of possible array configurations simultaneously.

4.3.3 Effect of Noise

The previous experiments assume no noise in the VLBI measurements. In this section, we compare the optimal telescope sensor-sampling under six different noise assumptions, (1) no noise, (2) equal thermal noise (using the average thermal noise derived from Table 1), (3) site-varying thermal noise (as derived from Table 1), (4) atmospheric phase noise, (5) atmospheric phase noise & equal thermal noise, and (6) atmospheric phase noise & site-varying thermal noise. For the latter three cases, the visibility amplitudes and closure phases are used for the image reconstruction. All cases assume a target source of Sgr A* with a constant total flux of 1 Jy using the "EHT+" array.

Fig. 9 shows the learned Ising model parameters for these cases. As can be seen in the site activity plots, the sensor-sampling distributions are not significantly influenced by thermal noise only: (1), (2) and (3) are almost identical. However, the atmospheric noise significantly impacts telescope placement. With the presence of atmospheric phase error, LMT and OVRO become much more important, while SPT becomes less important. This effect is further amplified when both thermal noises and atmospheric errors appear. Short baselines, such as OVRO-LMT, are sampled more frequently, suggesting that it is especially important to include shorter baselines in high-resolution VLBI imaging when dealing with atmospheric phase errors.

Based on the undirected graph defined by the Ising model correlation matrix, we can also analyze all "cliques" of size three to understand the importance of each closure triangle (i.e., closure

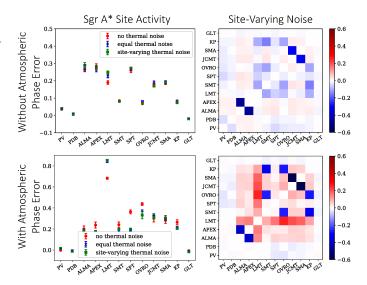


Fig. 9. Effect of measurement noise on array design. Two rows respectively show the Ising model parameters without and with atmospheric phase error. The top row uses Reconstruction Network A from Fig. 4, while the bottom row uses Reconstruction Network B to handle atmospheric noise. Thermal noise, no matter equal or site-varying, does not significantly change the learned telescope sensor-sampling. However, atmospheric phase noise is very impactful. According to a clique analysis in Sec. 4.3.3, nearly colinear telescopes are less favored after introducing atmospheric phase error. In both cases, only the correlation matrices for site-varying noise (cases (3) and (6) in Sec. 4.3.3) are shown.

phase measurement). In the affinity graph, telescopes are vertices and the pairwise correlations define edges. We define a subset of telescopes (x_j, x_k, x_l) as a "three-clique" if the correlations between each pair are significantly positive (larger than a threshold τ), i.e.

$$\theta_{ik} > \tau, \theta_{kl} > \tau, \theta_{il} > \tau.$$
 (24)

For instance, in case (6), eight three-cliques can be found in the corresponding graph given $\tau=0.04$. After ranking them with a simple metric,

$$m_c = \theta_{jk} + \theta_{kl} + \theta_{jl}, \tag{25}$$

we find the five most important three-cliques are "LMT-OVRO-JCMT", "LMT-OVRO-SMA", "APEX-LMT-JCMT", "ALMA-LMT-JCMT" and "LMT-JCMT-KP", all of which form triangles with similar baseline lengths (see Fig. 3). This suggests that closure phases from sites that are nearly colinear may not be as important for image reconstruction with atmospheric phase error.

4.3.4 EHT Design for Different Science Targets

In this section, we compare the results of two different targets, Sgr A* and M87*. They are in the southern and northern celestial hemisphere, respectively (Sgr A* declination: -29.24°, M87* declination: 12.39°). Therefore Sgr A* and M87* should have different preferences for telescope site selections. Fig. 10 reports the learned EHT sensor-sampling distributions considering no noise or only atmospheric phase noise. For both noise assumptions, SPT (occluded for M87* observations) becomes less important for M87* compared with Sgr A*, while PV, PDB and GLT become more important.

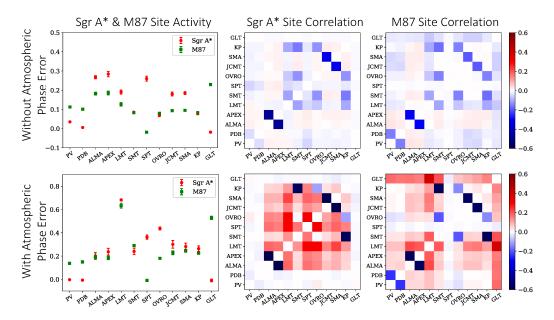


Fig. 10. Comparison of the learned telescope importance for science targets Sgr A* and M87. Because the targets are at different declinations (Sgr A*: -29.24°, M87: 12.39°), different visibility measurements are sampled by the same telescope array (due to a change in the projected baselines). Therefore, different sets of telescopes are selected as important for the targets. For instance, GLT is very important for M87* but is the least important for Sgr A*, because it is occluded by Earth when observing targets in southern celestial hemisphere. The opposite is true for SPT.

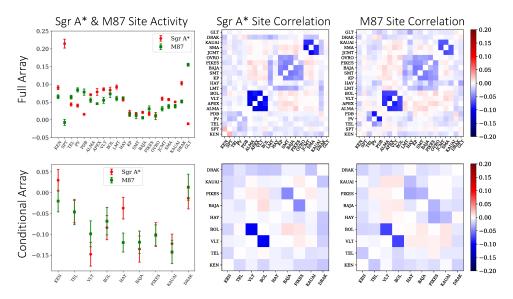


Fig. 11. Ising model parameters for the "FUTURE" array. The top figure shows the joint distribution of all sites in the "FUTURE" array. Negative correlations among close telescopes are strongly detected. The bottom row reports the Ising model distributions of nine additional sites conditioned on all the "EHT+" telescopes being included in the array. We find that adding telescopes in Africa, particularly DRAK, would be particularly helpful in constraining both Sgr A* and M87* reconstructions. BAJA, which is close to KP, SMT and OVRO, has the least conditional activity.

4.3.5 Site Selection for "FUTURE" array

We study the site selection for a "FUTURE" array, which includes nine additional telescope sites in addition to the twelve in "EHT+". The same representative parameters described at the beginning of Sec. 4.3 are applied and no noise is introduced. The learned Ising model parameters for both Sgr A* and M87 are shown in Fig. 11. As can be seen, the results successfully capture the negative correlations among co-located telescopes, such as ALMA-APEX-VLT-BOL and JCMT-SMA-KAUAI.

Since the conditional distribution of a subset of Ising model elements is also an Ising model (shown in Eq. 10), we also report

the conditional activities and correlations of nine additional sites assuming all the "EHT+" telescopes are selected. Although nearly all nine sites have negative activities (due to the sparsity constraint of the original Ising model), the rank of these sites' importance still holds. We find that telescopes in Africa, particularly DRAK, would be helpful in constraining both Sgr A* and M87* reconstructions. However, we caution that more factors should be considered in the model before drawing conclusions about the importance of any one site, such as realistic atmospheric noise, weather assumptions, the science target's expected evolution, and both social and economic costs.

4.3.6 Swapping Sampling Strategies

In this section, we study the interaction of learned sampling strategies (encoders) and reconstruction methods (decoders) in order to demonstrate the advantage of co-design. In particular, we swap the sampling strategy learned for different types of noise (no noise & atmospheric phase noise) and science targets (Sgr A* & M87*), while keeping the reconstruction methods fixed. For instance, for two co-designed networks, we evaluate the quality of images reconstructed with the decoder of network 1, when the telescope measurements were sampled from the policy learned by network 2, and visa versa. To meaningfully compare networks, λ_1 and λ_2 were tuned so that the expected number of sampled telescopes in the EHT+ array for the different networks were nearly equal (\approx 8).

Table 2 reports shift-invariant reconstruction losses (Eq. 23) of various "recombined" autoencoders for a res FWHM of 0.5×10^{-5} the nominal beam. Each row represents a learned reconstruction method and each column represents a learned sensor sampling distribution. The reconstruction metric is computed based on 1000 test images in the "MNIST" dataset. As expected, the optimal sampling strategy and reconstruction method are those that are learned jointly: a particular VLBI array sampling policy works best with its corresponding reconstruction method.

Fig. 12 shows the mean reconstruction and standard deviation resulting from 1000 sampling trials of the same truth image (see Fig. 7) for each "recombined" autoencoder. The standard deviation is significantly lower for autoencoders that were co-designed. This simple tests help to confirm the co-designed sensor-sampling and reconstruction strategies outperform independently learned strategies.

5 DISCUSSION

Optimal sensing is important for resource-limited image reconstruction. We presented an approach to learn an optimal sensor-sampling distribution for a computational imaging system. This method optimizes an Ising sensor-sampling distribution jointly with an image reconstruction method, implemented in a physics-constrained, fully differentiable, autoencoder.

We demonstrated the proposed framework on a VLBI telescope array design task, where sensor correlations and atmospheric noise present unique challenges. Experiments show that the proposed technique can be helpful in planing future telescope array designs and developing observation strategies. We also believe the proposed framework is applicable to a variety of other sensor selection/sampling problems in computational imaging, such as the optimization of LED array illumination patterns in Fourier ptychography [7] and the sampling pattern in κ -space for fast-MRI [28]. Changing the application simply requires changing the physics-based forward model, f(z), the decoder, $A_w\{\cdot\}$, in Eq. 2 and the training data. By co-designing the Ising model sampling distribution simultaneously with the reconstruction decoder, we can better optimize the design of future computational imaging systems.

ACKNOWLEDGMENTS

The authors would like to thank Lindy Blackburn, Alexander Raymond, Michael Johnson, and Sheperd Doeleman for helpful discussions on the constraints of a next-generation EHT array, and

TABLE 2
Reconstruction losses after swapping sampling strategies.

	samp.	Sgr A*, no atmospheric phase error	Sgr A*, atmospheric phase error	M87, no atmospheric phase error	M87, atmospheric phase error
-	Sgr A*, no atmospheric phase error	0.0072	0.0088	0.0168	0.0165
	Sgr A*, atmospheric phase error	0.0199	0.0163	0.0252	0.0230
	M87, no atmospheric phase error	0.0155	0.0164	0.0061	0.0079
	M87, atmospheric phase error	0.0171	0.0152	0.0127	0.0116

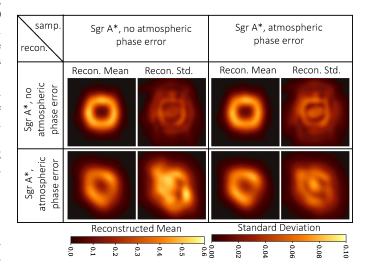


Fig. 12. Reconstruction mean and standard deviation from 1000 sampling trials of the "ring" image in Fig. 7 using "recombined" autoencoders. Reconstruction performance is best for the co-designed networks; the reconstructed standard deviation is significantly higher for the "recombined" autoencoders with independently learned encoders and decoders.

Michael Kellman for helpful discussions on Fourier ptychography. This work was supported by NSF award 1935980: "Next Generation Event Horizon Telescope Design," and Beyond Limits.

REFERENCES

- [1] E. Kobler, M. Muckley, B. Chen, F. Knoll, K. Hammernik, T. Pock, D. Sodickson, and R. Otazo, "Variational deep learning for low-dose computed tomography," in 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2018, pp. 6687–6691.
- [2] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing mri," *IEEE signal processing magazine*, vol. 25, no. 2, p. 72, 2008
- [3] Y. Rivenson, Z. Göröcs, H. Günaydin, Y. Zhang, H. Wang, and A. Ozcan, "Deep learning microscopy," *Optica*, vol. 4, no. 11, pp. 1437–1443, 2017.
- [4] K. L. Bouman, M. D. Johnson, D. Zoran, V. L. Fish, S. S. Doeleman, and W. T. Freeman, "Computational imaging for vlbi image reconstruction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 913–922.
- [5] G. Olague and R. Mohr, "Optimal camera placement for accurate reconstruction," *Pattern recognition*, vol. 35, no. 4, pp. 927–944, 2002.

- [6] L. Tian, Z. Liu, L.-H. Yeh, M. Chen, J. Zhong, and L. Waller, "Computational illumination for high-speed in vitro fourier ptychographic microscopy," *Optica*, vol. 2, no. 10, pp. 904–911, 2015.
- [7] M. Kellman, E. Bostan, M. Chen, and L. Waller, "Data-driven design for fourier ptychographic microscopy," in 2019 IEEE International Conference on Computational Photography (ICCP). IEEE, 2019, pp. 1–8
- [8] D. Woody, "Radio interferometer array point spread functions i. theory and statistics," ALMA Memo Series, vol. 389, 2001.
- [9] F. Boone, "Interferometric array design: Optimizing the locations of the antenna pads," *Astronomy & Astrophysics*, vol. 377, no. 1, pp. 368–376, 2001.
- [10] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, Feb 2006.
- [11] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.21391
- [12] C. Bouman and K. Sauer, "A generalized gaussian image model for edgepreserving map estimation," *IEEE Transactions on image processing*, vol. 2, no. 3, pp. 296–310, 1993.
- [13] J. Skilling and R. Bryan, "Maximum entropy image reconstructiongeneral algorithm," *Monthly notices of the royal astronomical society*, vol. 211, p. 111, 1984.
- [14] E. Candes and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse problems*, vol. 23, no. 3, p. 969, 2007.
- [15] K. L. Bouman, M. D. Johnson, A. V. Dalca, A. A. Chael, F. Roelofs, S. S. Doeleman, and W. T. Freeman, "Reconstructing video of time-varying sources from radio interferometric measurements," *IEEE Transactions on Computational Imaging*, vol. 4, no. 4, pp. 512–527, 2018.
- [16] A. Sinha, J. Lee, S. Li, and G. Barbastathis, "Lensless computational imaging through deep learning," *Optica*, vol. 4, no. 9, pp. 1117–1125, 2017.
- [17] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in 2013 IEEE Global Conference on Signal and Information Processing. IEEE, 2013, pp. 945–948.
- [18] H. Chen, J. Gu, O. Gallo, M.-Y. Liu, A. Veeraraghavan, and J. Kautz, "Reblur2deblur: Deblurring videos via self-supervised learning," in 2018 IEEE International Conference on Computational Photography (ICCP). IEEE, 2018, pp. 1–9.
- [19] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4509–4522, 2017.
- [20] S. Diamond, V. Sitzmann, F. Heide, and G. Wetzstein, "Unrolled optimization with deep priors," arXiv preprint arXiv:1705.08041, 2017.
- [21] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll, "Learning a variational network for reconstruction of accelerated mri data," *Magnetic resonance in medicine*, vol. 79, no. 6, pp. 3055–3071, 2018.
- [22] C. E. Shannon, "Communication in the presence of noise," *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, 1949.
- [23] D. Palumbo, S. S. Doeleman, M. D. Johnson, K. L. Bouman, and A. A. Chael, "Metrics and motivations for earth-space vlbi: Time-resolving sgr a* with the event horizon telescope," arXiv preprint arXiv:1906.08828, 2019.
- [24] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling [a sensing/sampling paradigm that goes against the common knowledge in data acquisition]," *IEEE signal processing magazine*, vol. 25, no. 2, pp. 21–30, 2008.
- [25] A. Chakrabarti, "Learning sensor multiplexing design through backpropagation," in Advances in Neural Information Processing Systems, 2016, pp. 3081–3089.
- [26] Y. P. X. D. S. B. W. H. F. H. G. W. V. Sitzmann, S. Diamond, "End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging," ACM Trans. Graph. (SIGGRAPH), 2018.
- [27] M. Yoshida, A. Torii, M. Okutomi, K. Endo, Y. Sugiyama, R.-i. Taniguchi, and H. Nagahara, "Joint optimization for compressive video sensing and reconstruction under hardware constraints," in *Proceedings* of the European Conference on Computer Vision (ECCV), 2018, pp. 634– 649.
- [28] C. D. Bahadir, A. V. Dalca, and M. R. Sabuncu, "Adaptive compressed sensing mri with unsupervised learning," arXiv preprint arXiv:1907.11374, 2019.

- [29] J. Sun, H. Li, Z. Xu et al., "Deep admm-net for compressive sensing mri," in Advances in neural information processing systems, 2016, pp. 10–18
- [30] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [31] D. Lee, J. Yoo, and J. C. Ye, "Deep residual learning for compressed sensing mri," in 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017). IEEE, 2017, pp. 15–18.
- [32] A. R. Thompson, J. M. Moran, G. W. Swenson et al., Interferometry and synthesis in radio astronomy. Wiley New York et al., 1986.
- [33] K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, D. Bintley et al., "First m87 event horizon telescope results. iv. imaging the central supermassive black hole," The Astrophysical Journal Letters, vol. 875, no. 1, p. L4, 2019.
- [34] L. Blackburn, S. Doeleman, J. Dexter, J. L. Gómez, M. D. Johnson, D. C. Palumbo, J. Weintroub, J. R. Farah, V. Fish, L. Loinard *et al.*, "Studying black holes on horizon scales with vlbi ground arrays," *arXiv preprint arXiv:1909.01411*, 2019.
- [35] A. A. Chael, M. D. Johnson, K. L. Bouman, L. L. Blackburn, K. Akiyama, and R. Narayan, "Interferometric imaging directly with closure phases and closure amplitudes," *The Astrophysical Journal*, vol. 857, no. 1, p. 23, 2018.
- [36] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," arXiv preprint arXiv:1708.07747, 2017.
- [37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.



He Sun is a postdoctoral researcher in Computing and Mathematical Sciences (CMS) at the California Institute of Technology. He received his PhD from Princeton University in 2019 and his BS degree from Peking University in 2014. His research focuses on adaptive optics and computational imaging, especially their applications in astrophysical and biomedical sciences.



Adrian V. Dalca is an assistant professor at Harvard Medical School, Massachusetts General Hospital, and a Research Scientist in CSAIL, Massachusetts Institute of Technology. His research focuses on machine learning techniques and probabilistic models for medical image analysis. Driven by clinical questions, he develops core learning algorithms, as well as registration, segmentation and imputation methods aimed at clinical-sourced datasets and broadly applicable in image analysis.



Katherine L. Bouman is an assistant professor in Computing and Mathematical Sciences (CMS) and Electrical Engineering (EE) and a Rosenberg Scholar at the California Institute of Technology (Caltech) in Pasadena, California. Her group combines ideas from signal processing, computer vision, machine learning, and physics to design computational imaging systems. These systems tightly integrate algorithm and sensor design, making it possible to observe phenomena previously difficult or impossible to

measure with traditional approaches. She was recognized as a corecipient of the Breakthrough Prize in fundamental physics, and was named the 2020 imaging Scientist of the Year by the Society for Imaging Science and Technology for her work capturing the first picture of a black hole with the international Event Horizon Telescope Collaboration.