(Random) phase retrieval: theory, algorithms, applications

Jonathan Dong
Biomedical Imaging Group, EPFL, Lausanne
jonathan.dong@epfl.ch

October 2024

PHYS-715, EPFL

Phase imaging

Amplitude

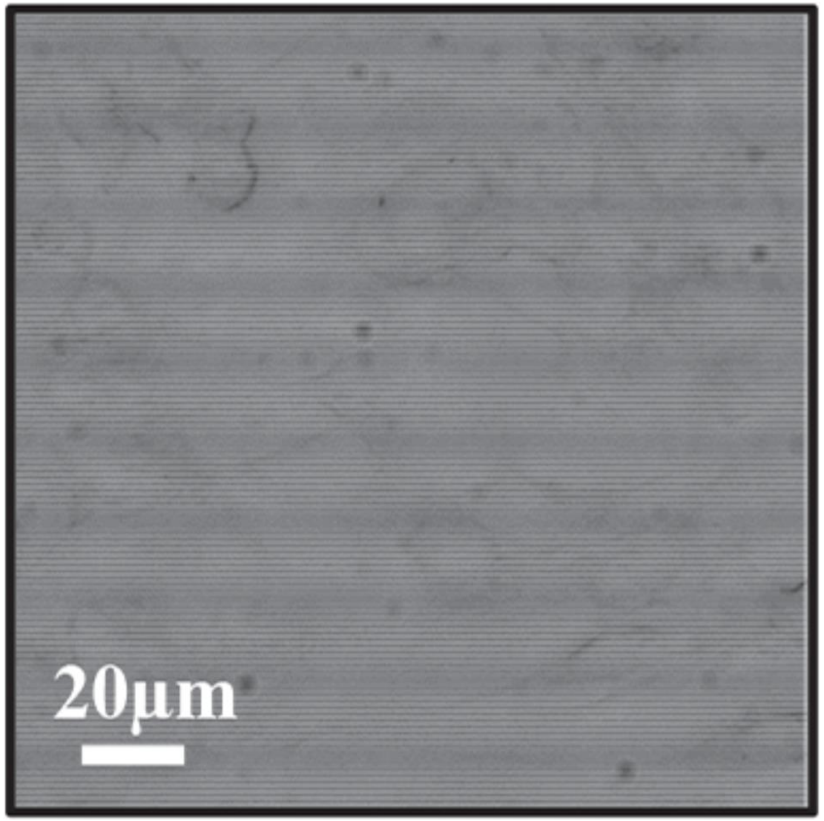

Phase

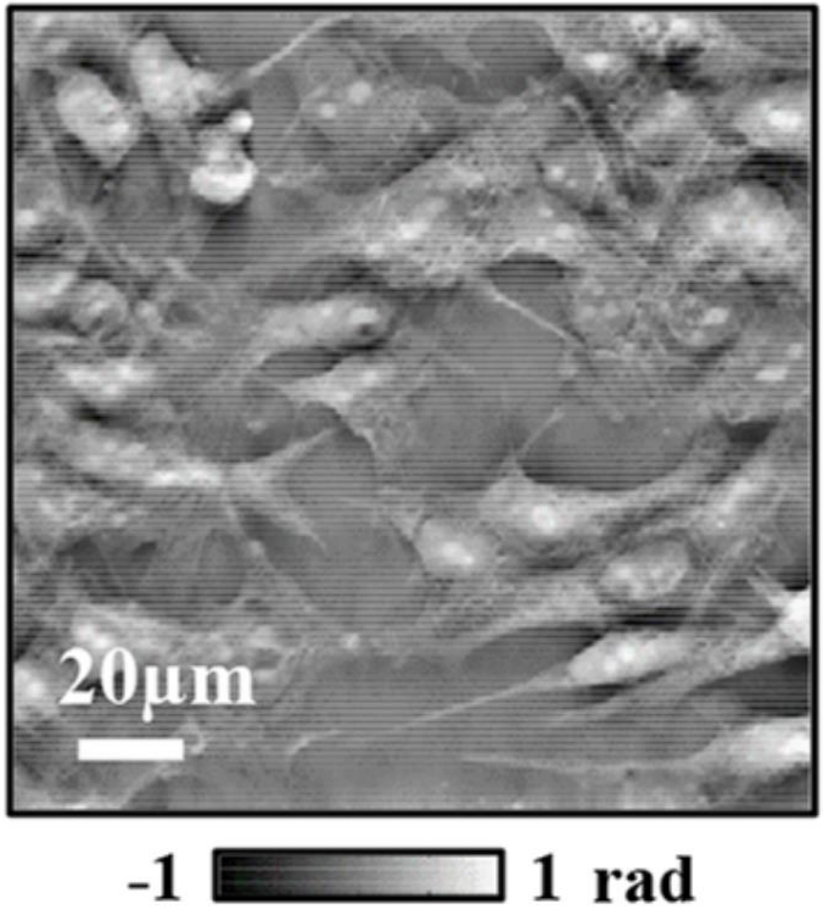

Tian, Waller (2015). *Optics Express*

Optical phase imaging

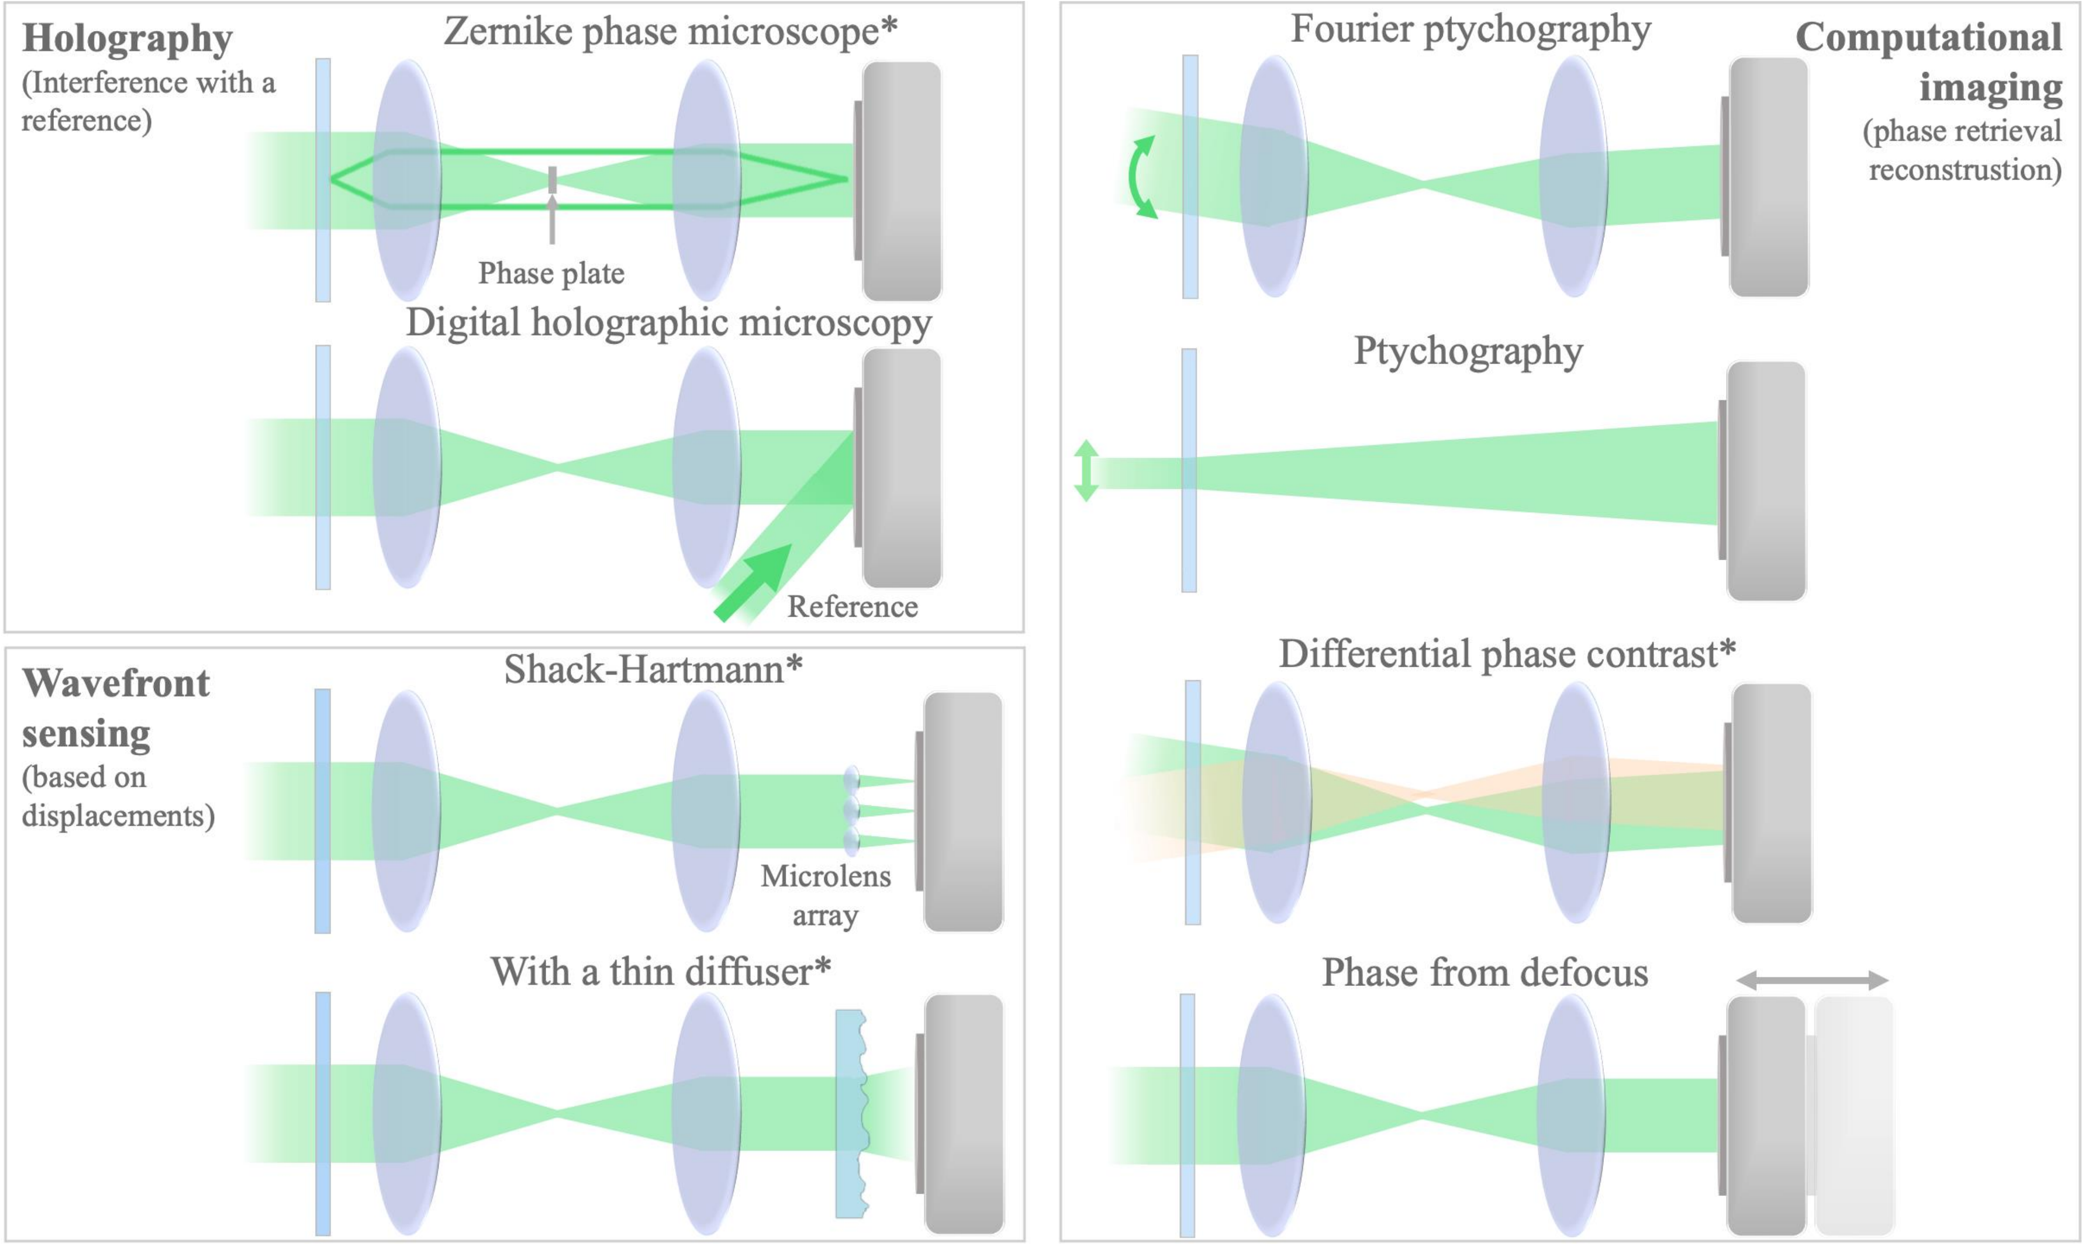

The Biomedical Imaging Group (BIG)

Michael Unser

Image reconstruction
Advanced algorithms
Machine learning for imaging
Collaboration with imaging groups

Image analysis
Digital histopathology
Localization microscopy
Tools for biologists / doctors

Daniel Sage

Phase retrieval

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

$$\begin{matrix} y_1 \\ \ \\ \vdots \\ \ \\ y_n \end{matrix} = \left| \begin{matrix} a_{11} & & \cdots & a_{1d} \\ & & \cdots & \\ & \vdots & & \\ & & \cdots & \\ a_{n1} & & \cdots & a_{nd} \end{matrix} \quad \begin{matrix} x_1^* \\ \ \\ \vdots \\ x_d^* \\ \ \end{matrix} \right|^2$$

Phase retrieval

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

And in the previous episode...

Imaging as an inverse problem

Physical model of interaction between incoming wave and object

Diversity - Can be used to encode more information, e.g. rotation, scanning, energy

r (dimensionality) can be 2D, 3D, time, spectra

O (contrast) matter-wave interaction can be a scalar, vector, tensor
Captures local anisotropy

M (modulation and measurement) images, diffraction, fluorescence, photo-electrons

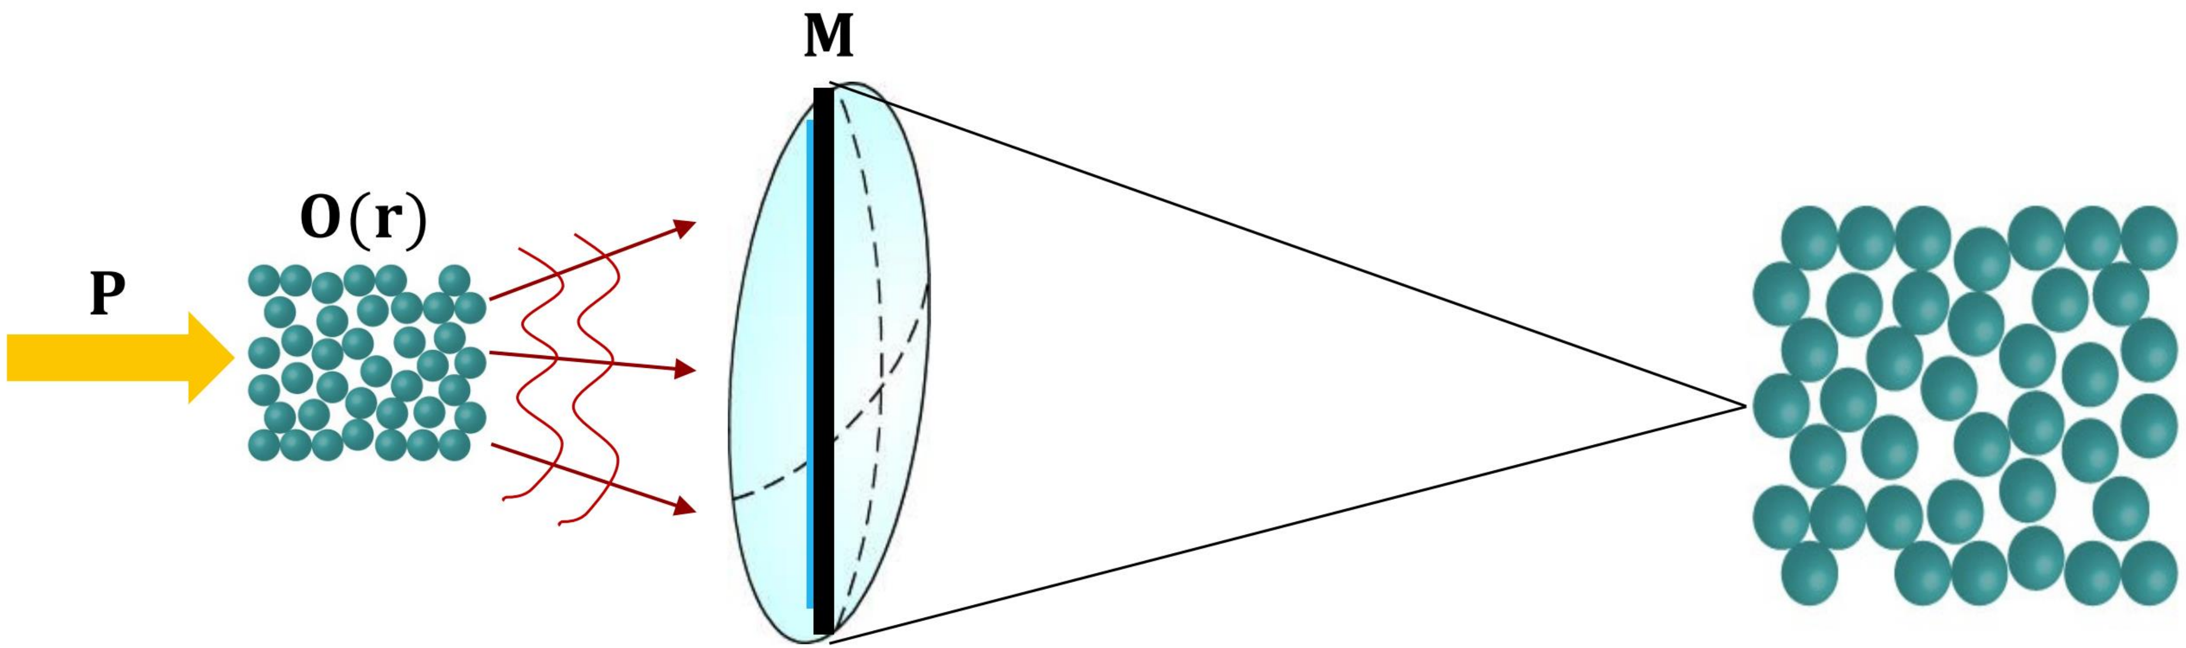

Inverse problem framework

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = \boldsymbol{A}\{\boldsymbol{x}^*\}$$

- $\boldsymbol{x}^*$: image to recover, parameters to estimate
- $\boldsymbol{A}$: physical model
- $\boldsymbol{y}$: measurements

Inverse problem framework

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = \boldsymbol{A}\{\boldsymbol{x}^*\}$$

- Optimization approach:

$$\hat{x} = \mathrm{argmin}_x \, l(y, x)$$

- We often split the loss l into $l(y, x) = f(y, x) + g(x)$:
  - Data fidelity, e.g. L2: $f(y, x) = \|y - A\{x\}\|_2^2$
  - Regularization, e.g. L1: $g(x) = \|x\|_1$

Linear systems

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}^*$$

with a linear operator $\boldsymbol{A}$

- If A is invertible, $\hat{x} = A^{\dagger}y$
- If A is not invertible, add prior on x:
  - Sparsity with L1: $g(x) = \|x\|_1$
  - Total variation: $g(x) = \|\nabla x\|_1$
  - Deep learning regularization

Content

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

$$\begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} = \left| \begin{bmatrix} a_{11} & \cdots & a_{1d} \\ \vdots & & \vdots \\ a_{n1} & \cdots & a_{nd} \end{bmatrix} \begin{bmatrix} x_1^* \\ \vdots \\ x_d^* \end{bmatrix} \right|^2$$

- General inverse problem framework
- Phase retrieval (PR) applications
- PR algorithms
- PR theory: random model
- Machine learning

Phase retrieval

Find $\boldsymbol{x}^*$ in
$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

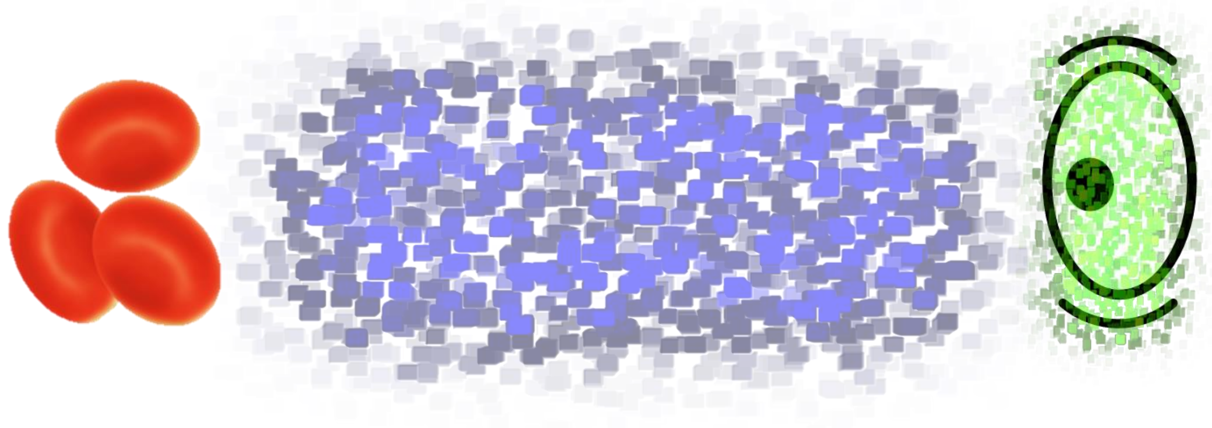

unknown object

$\boldsymbol{x}^*$
$\in \mathbb{C}^d$

imaging system

$\boldsymbol{A}$
$\in \mathbb{C}^{n \times d}$

measurements

$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$
$\in \mathbb{R}^n$

Intensity measurement

Phase retrieval applications

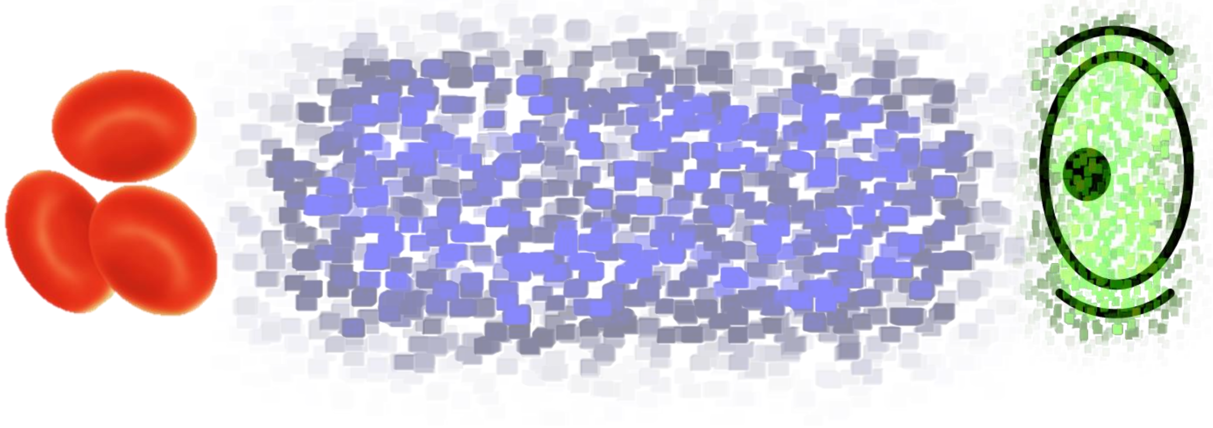

unknown object

$\boldsymbol{x}^*$

$\in \mathbb{C}^d$

imaging system

$\boldsymbol{A}$

$\in \mathbb{C}^{n\times d}$

measurements

$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$

$\in \mathbb{R}^n$

Amplitude

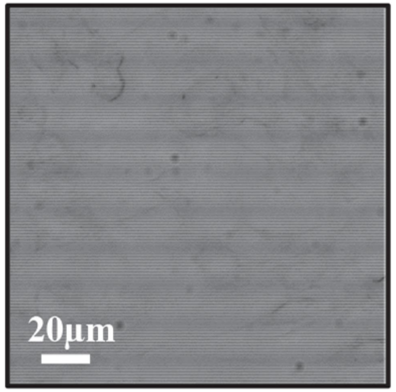

Quantitative Phase Imaging

Phase

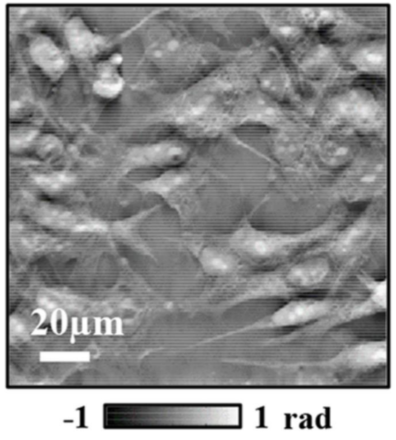

Label-free

Less invasive

No bleaching

Observations over several days

Fast

Videos of samples

Tian, Waller (2015). *Optics Express*

Phase retrieval applications

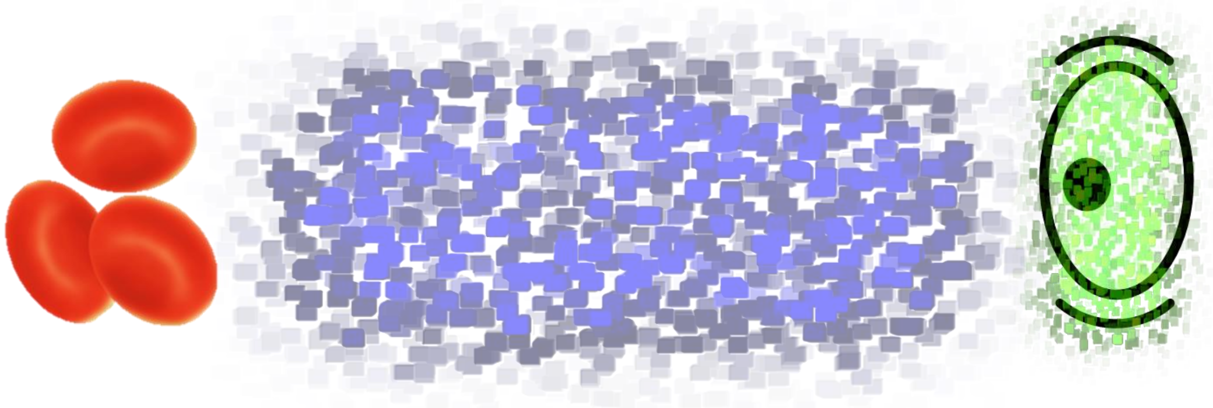

unknown object

$\boldsymbol{x}^*$

$\in \mathbb{C}^d$

imaging system

$\boldsymbol{A}$

$\in \mathbb{C}^{n\times d}$

measurements

$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$

$\in \mathbb{R}^n$

Quantitative Phase Imaging

lyncée tec DHM®

NANOLIVE
Looking inside life

phi

Tomocube

Disclaimer: They mainly use holographic approaches for 3D measurements

Phase retrieval

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

$$\begin{bmatrix} y_1 \\ \ \\ \vdots \\ \ \\ y_n \end{bmatrix} = \left| \begin{bmatrix} a_{11} & & \cdots & a_{1d} \\ & & \cdots & \\ & & \vdots & \\ & & \cdots & \\ a_{n1} & & \cdots & a_{nd} \end{bmatrix} \begin{bmatrix} x_1^* \\ \ \\ \vdots \\ x_d^* \end{bmatrix} \right|^2$$

Non-linear equation

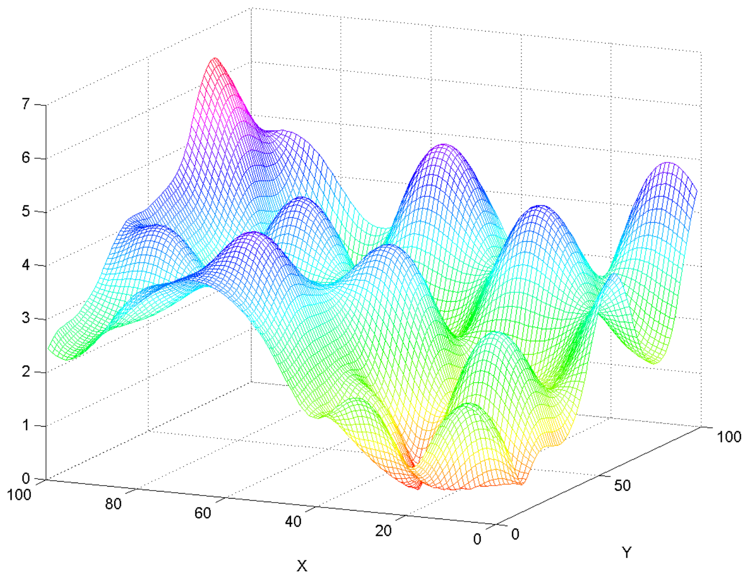

Phase retrieval and machine learning

Find $\boldsymbol{x}^*$ in
$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$

$$\begin{bmatrix} y_1 \\ \\ \vdots \\ \\ y_n \end{bmatrix} = \left| \begin{bmatrix} a_{11} & & \cdots & a_{1d} \\ & & \cdots & \\ & & \vdots & \\ & & \cdots & \\ a_{n1} & & \cdots & a_{nd} \end{bmatrix} \begin{bmatrix} x_1^* \\ \\ \vdots \\ x_d^* \end{bmatrix} \right|^2$$

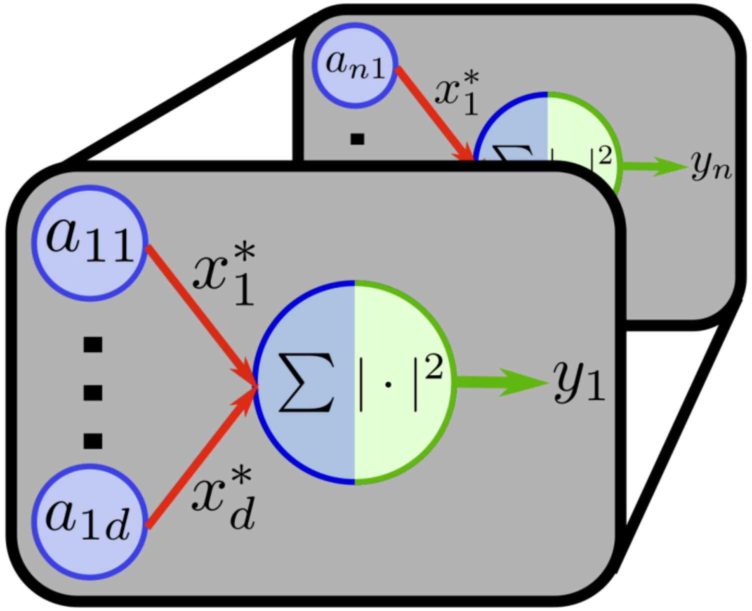

Machine learning

Single-layer neural network

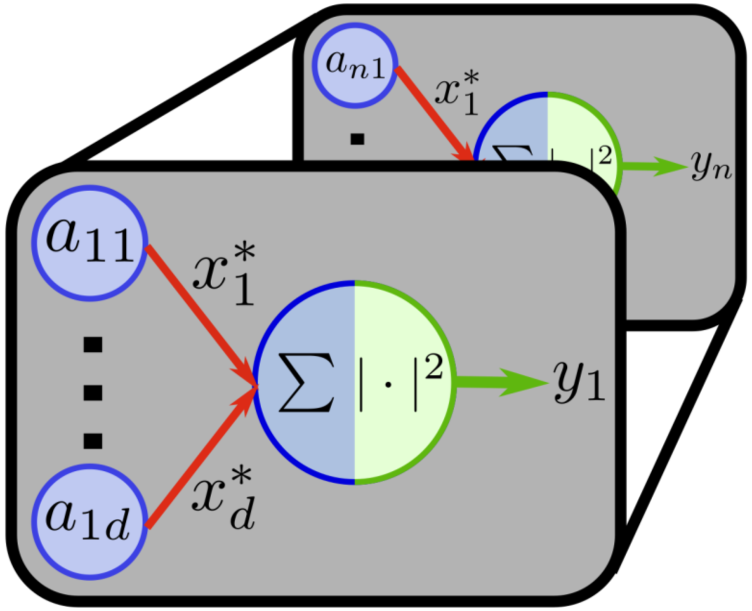

Deep neural network

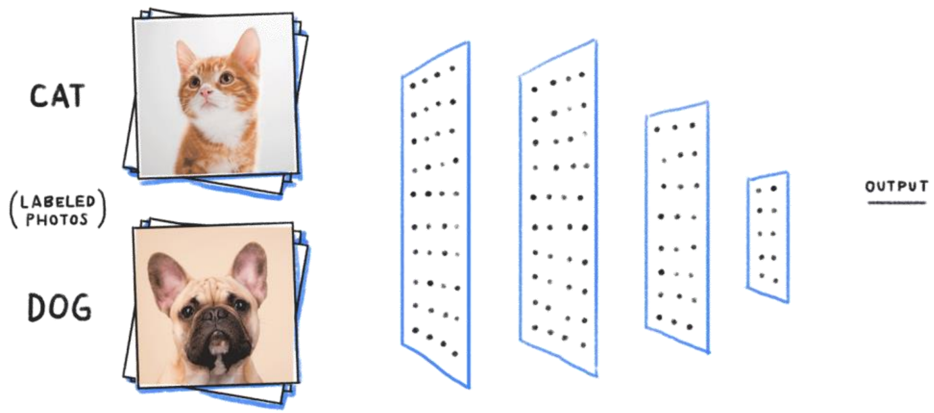

Machine learning

Single-layer neural network

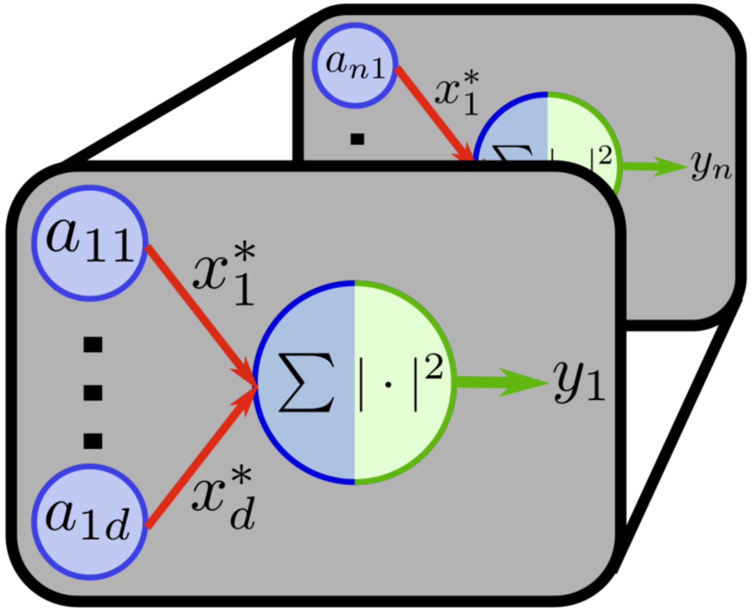

Phase retrieval =
Training a 1-layer neural network

When is it solvable?

What algorithm to use?

Is the solution unique?

Fourier in optics

Object
$\boldsymbol{x}$

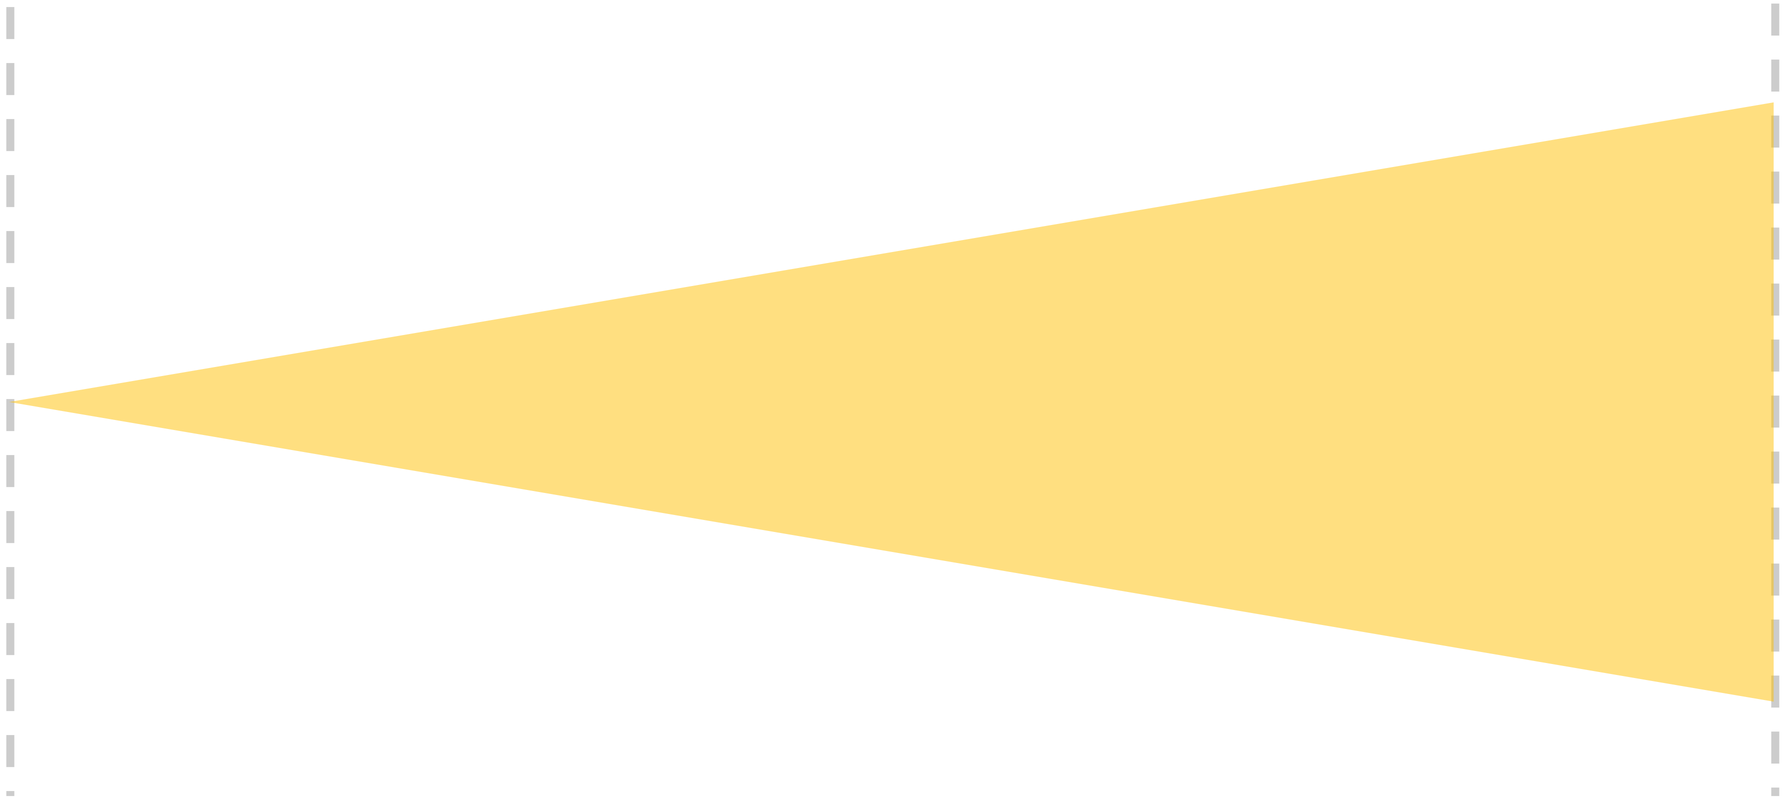

Output field
$\boldsymbol{E}^{out} = \boldsymbol{F}\boldsymbol{x}$

Far-field propagation
Fourier transform $\boldsymbol{F}$

Fourier in optics

Object
$\boldsymbol{x}$

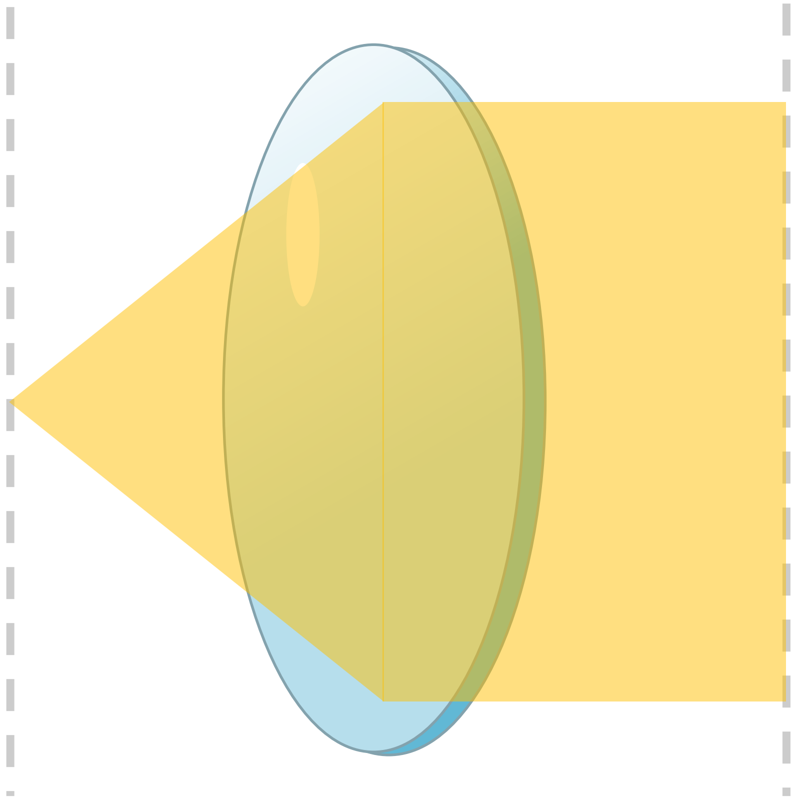

Output field
$\boldsymbol{E}^{\boldsymbol{out}} = \boldsymbol{F}\boldsymbol{x}$

Fourier transform
$\boldsymbol{F}$

Fourier in optics

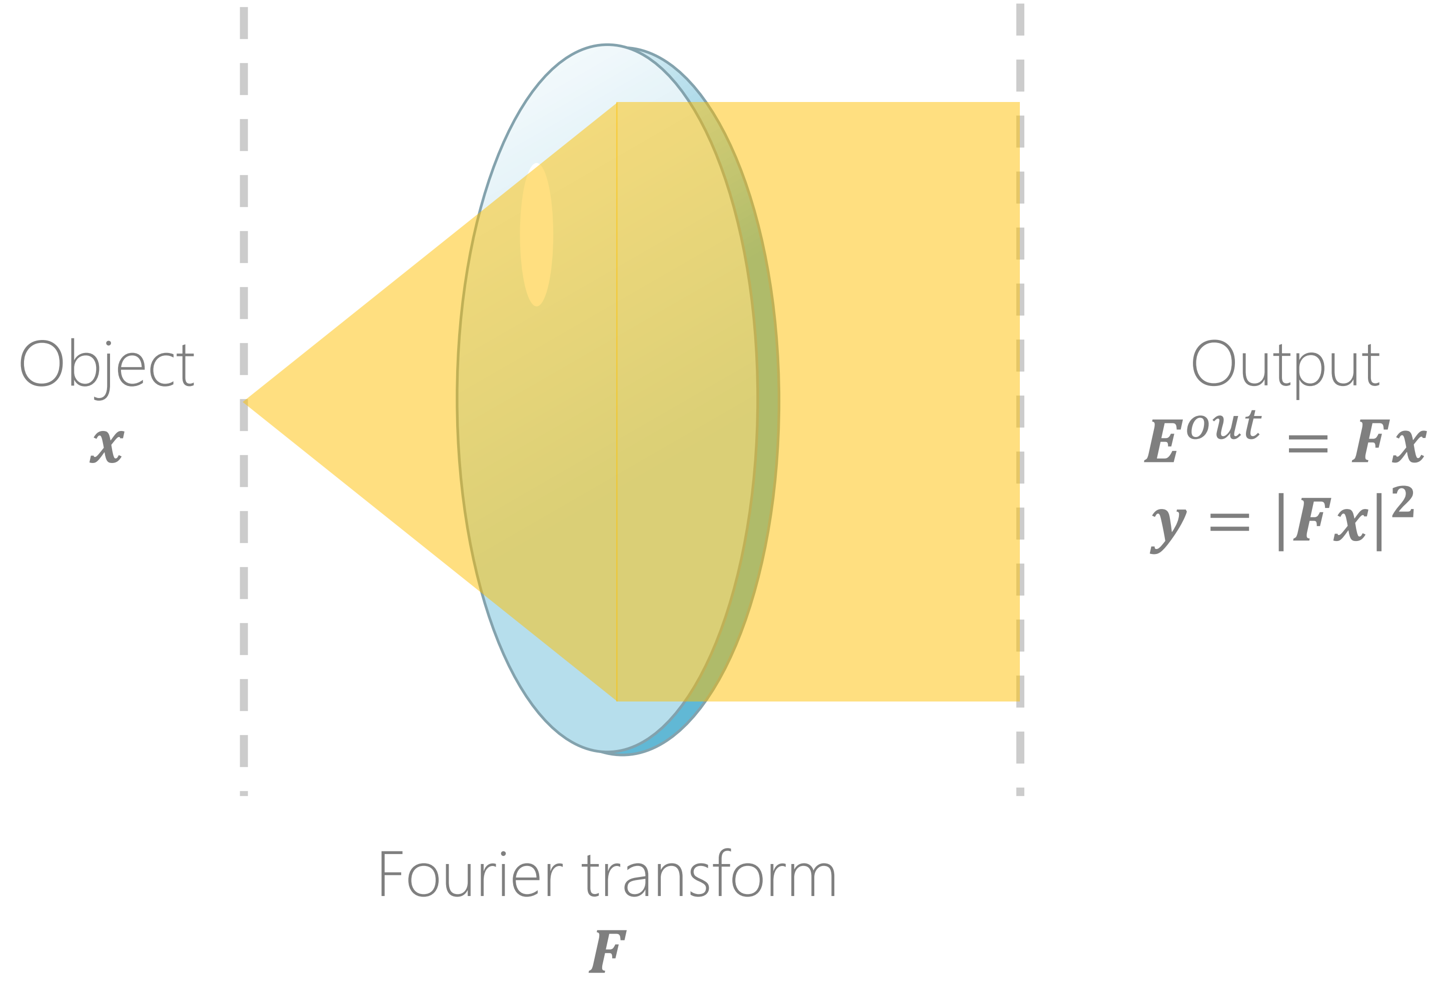

- Applications
  - Crystallography
  - Adaptive optics
  - PSF engineering
  - Complex media imaging
  - Non-line of sight imaging

Fienup (1982). *Applied Optics*

Coded-illumination: ptychography

Object
$\boldsymbol{x}$

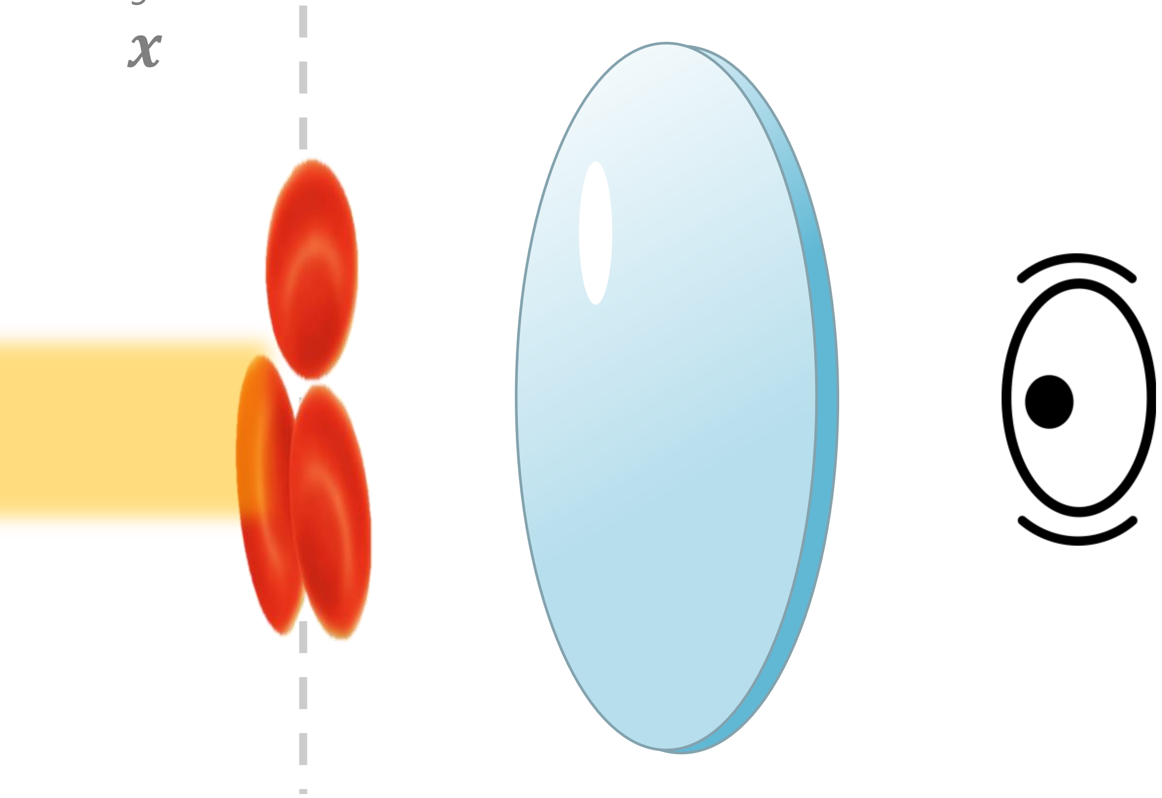

Output field
$\boldsymbol{E}^{out} = \boldsymbol{F}\boldsymbol{D}\boldsymbol{x}$

Illumination probe
$\boldsymbol{D}$

Coded-illumination: ptychography

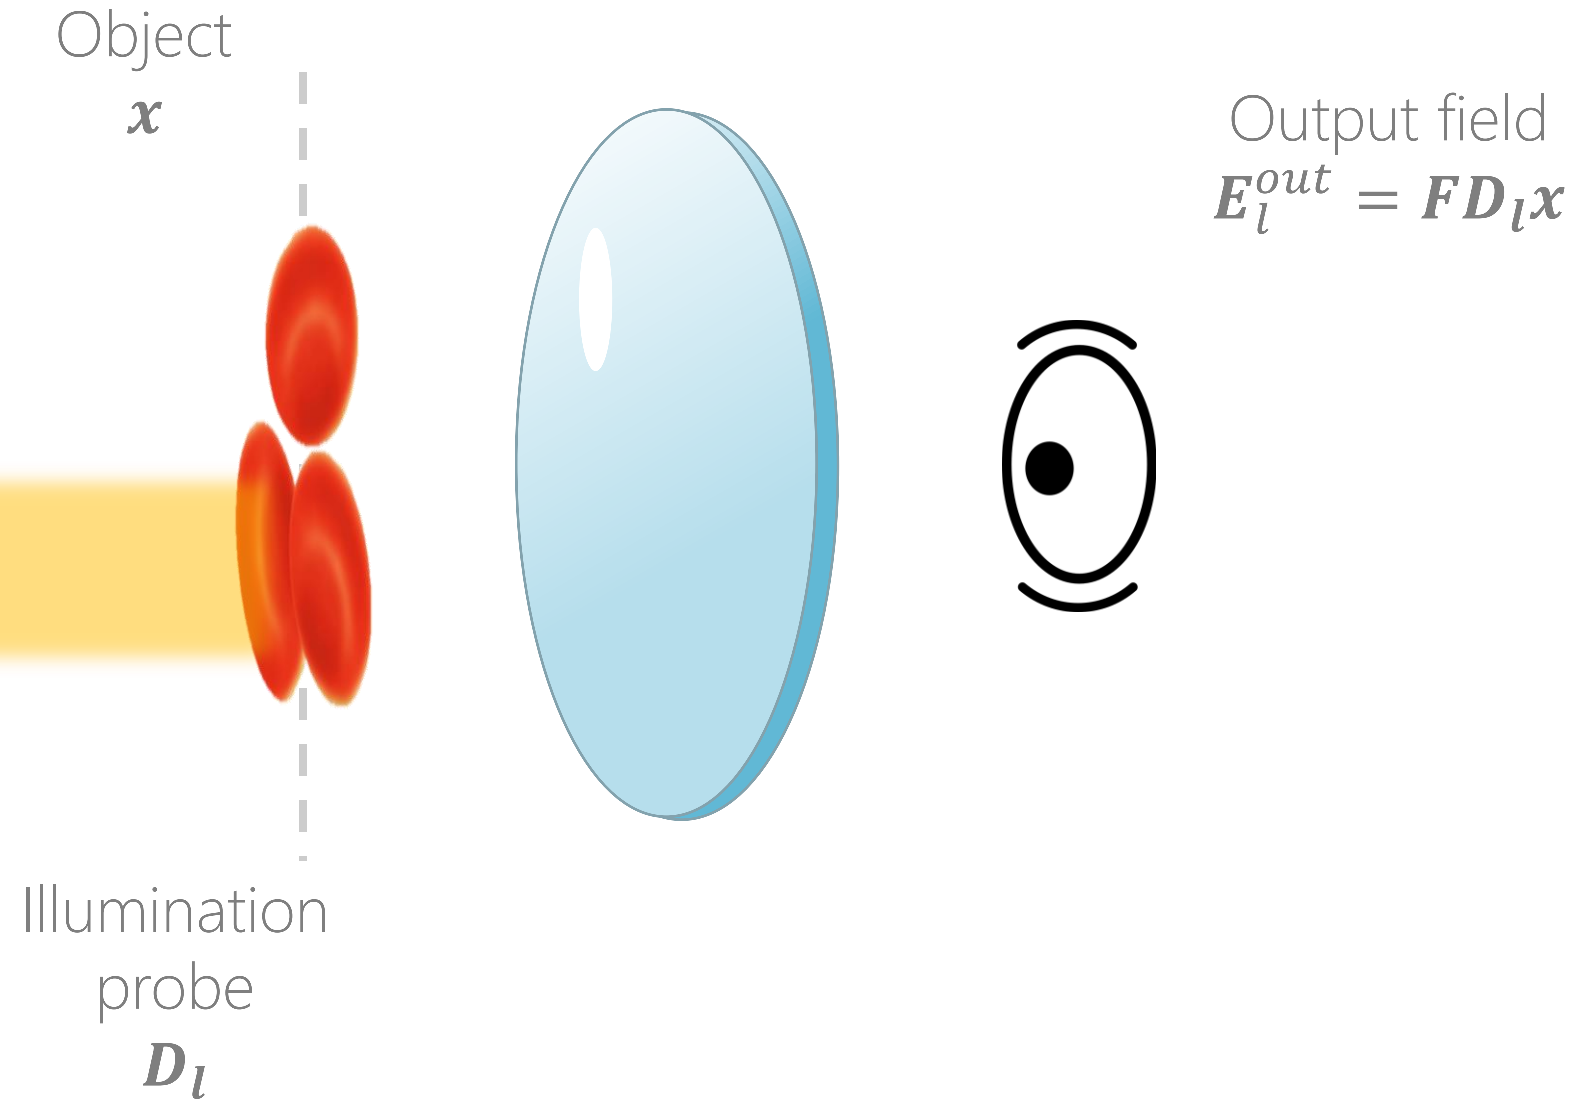

Coded-illumination: ptychography

Object
$\boldsymbol{x}$

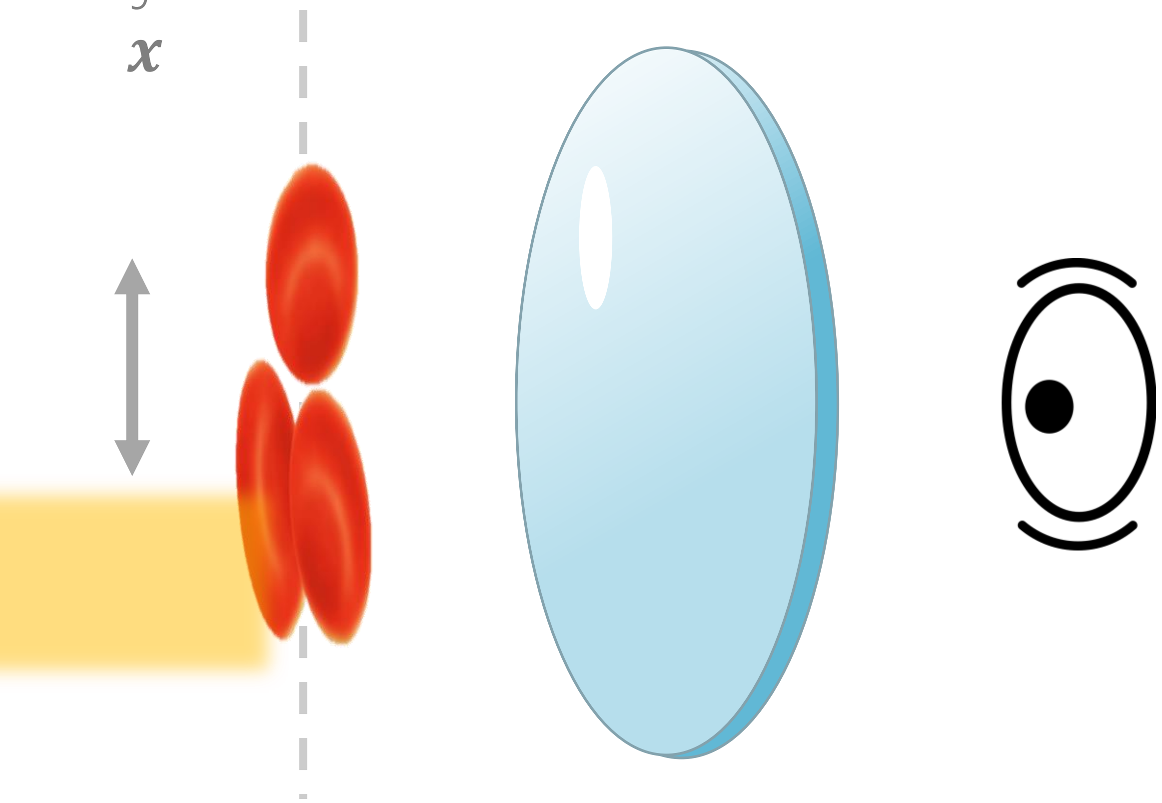

Illumination
probe
$\boldsymbol{D}_l$

Output field
$\boldsymbol{E}_l^{out} = \boldsymbol{F}\boldsymbol{D}_l\boldsymbol{x}$

$$\boldsymbol{y} = \left| \begin{pmatrix} \boldsymbol{A}_1 \\ \vdots \\ \boldsymbol{A}_L \end{pmatrix} \boldsymbol{x} \right|^2$$

with $\boldsymbol{A}_l = \boldsymbol{F}\boldsymbol{D}_l$

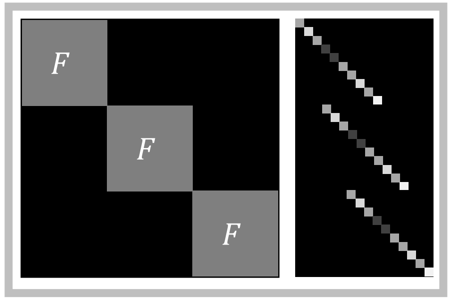

Miao, Charalambous, Kirz, Sayre (1999). *Nature*

Coded-illumination experiments

Ptychography

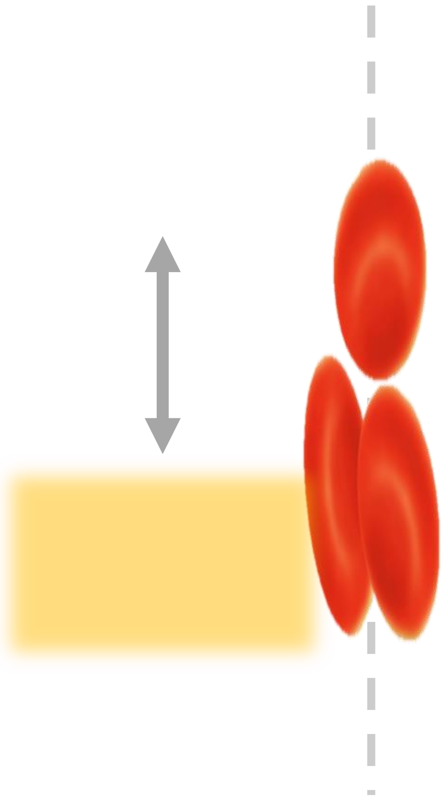

Shifted illumination

Fourier Ptychography

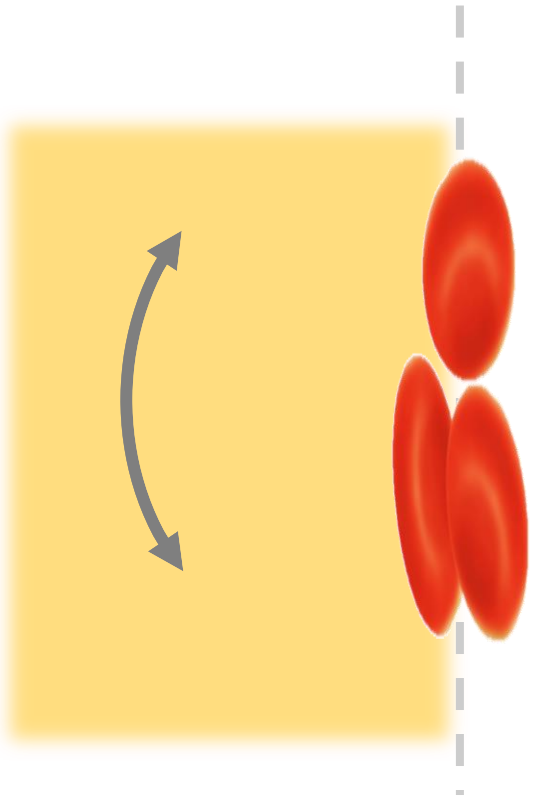

Tilted illumination

Coded Diffraction Imaging

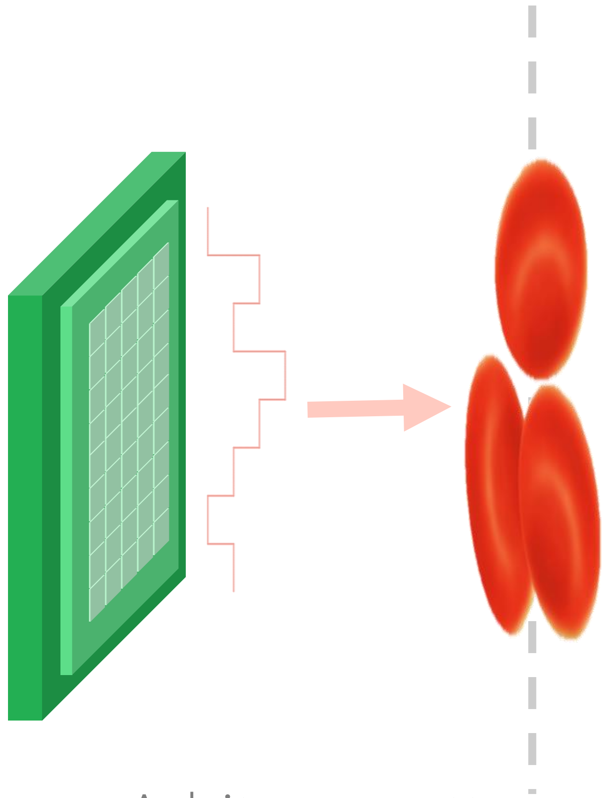

Arbitrary pattern
(Spatial Light Modulator,
mask, grating)

Zheng, Horstmeyer, Yang (2013). *Nature Photonics*

Coded detection

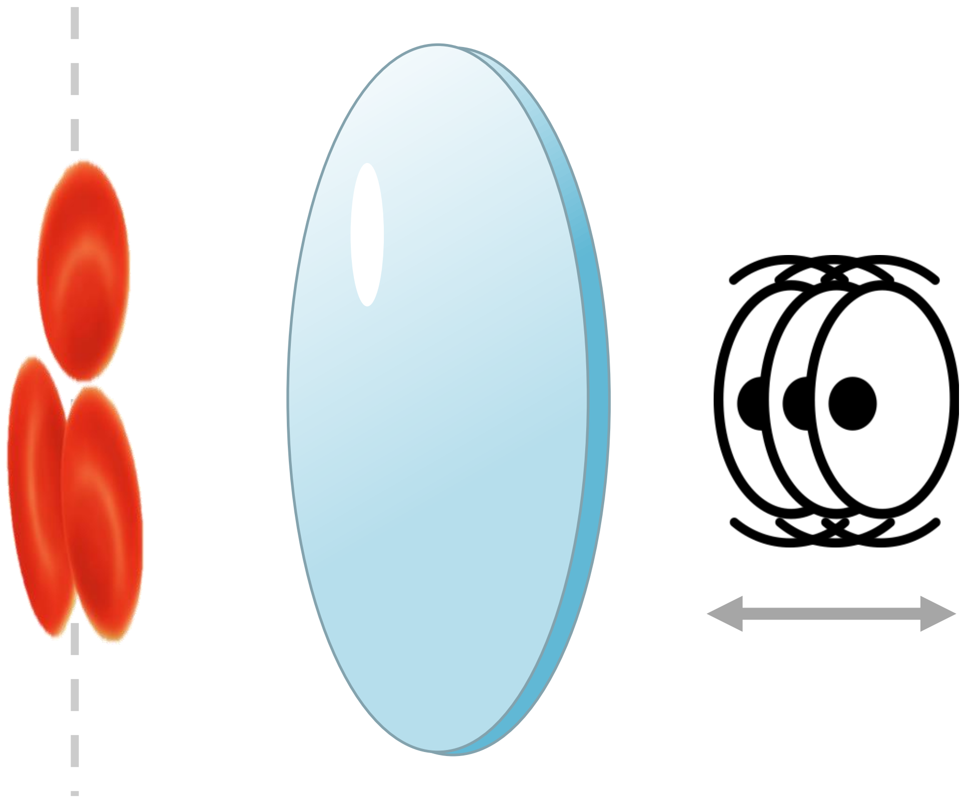

- Applications:
  - Astronomy
  - Non-invasive bioimaging

$\Rightarrow$ Modulate on the detection side

- Model $\boldsymbol{A}_l = \boldsymbol{F}\boldsymbol{D}_l\boldsymbol{F}^H$
  with defocus phase in Fourier space
  $$\boldsymbol{D}_l = \mathrm{Diag}\left(e^{iz_l\sqrt{1-\boldsymbol{u}^2}}\right)$$

Paxman, Schulz, Fienup (1992). *JOSA A*

The random model

- A is an i.i.d. random matrix
- $a_{ij} \sim \mathcal{N}\left(0, \frac{1}{d}\right)$
- Canonical setting for theory

- Applications:
    - Compressed sensing
    - Imaging in complex media

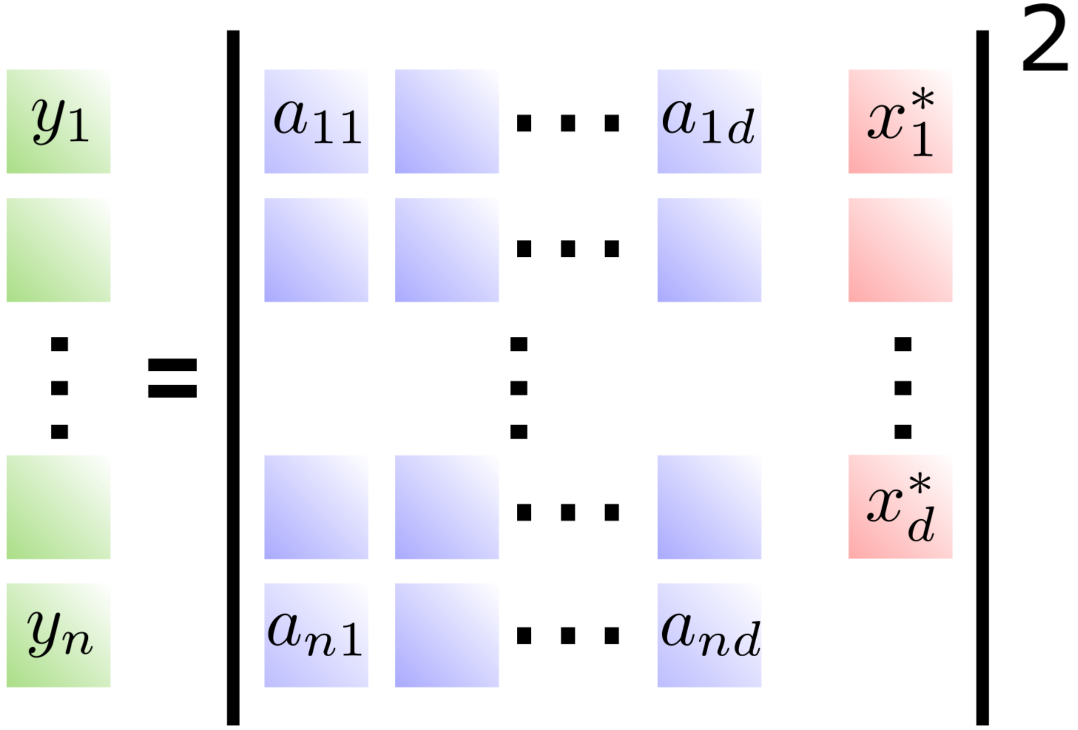

Mondelli, Montanari (2019). *Foundations of Computational Mathematics*

Unifying framework $\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}|^2$

Fourier phase retrieval

$$\mathbf{A} = \mathbf{F}$$

by a lens

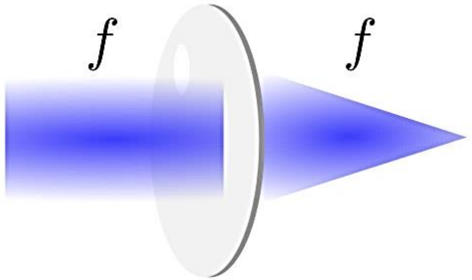

by free-space propagation

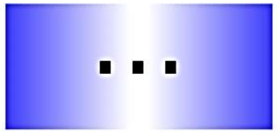

Coded illumination

$$\mathbf{A}_l = \mathbf{F}\mathbf{D}_l$$

imaging system

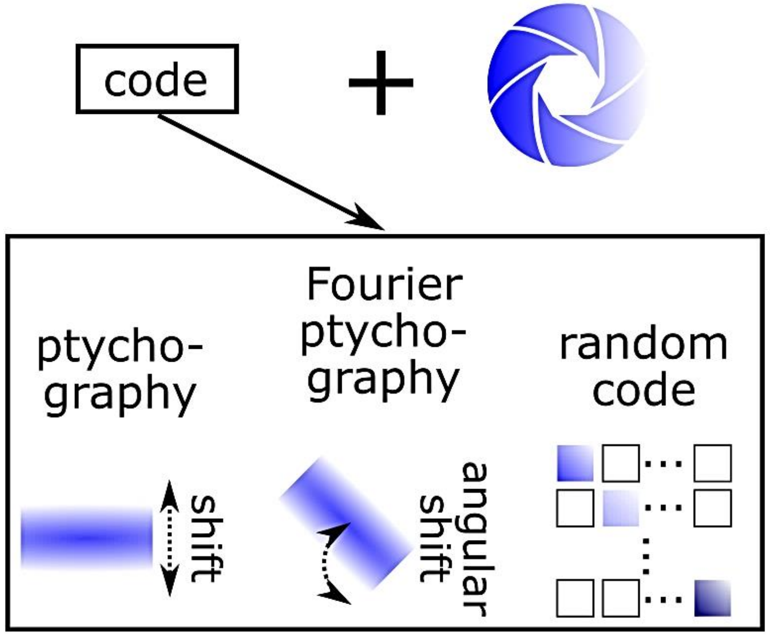

Coded detection

$$\mathbf{A}_l = \mathbf{F}\mathbf{D}_l\mathbf{F}^{\mathrm{H}}$$

phase diversity

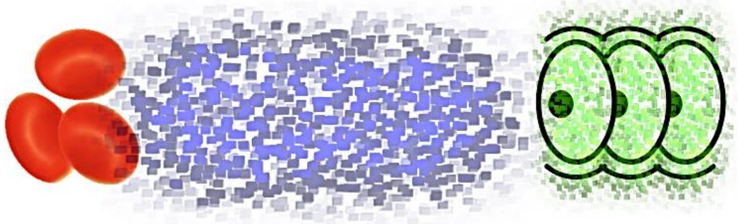

computer-generated holography

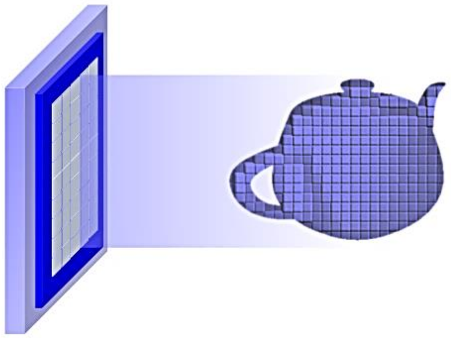

Random

by random projections

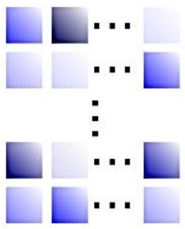

by propagation through complex media

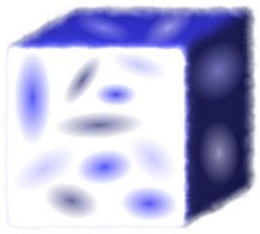

Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). *IEEE Signal Processing Magazine*

Content

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

$$\begin{bmatrix} y_1 \\ \\ \vdots \\ \\ y_n \end{bmatrix} = \left| \begin{bmatrix} a_{11} & & \cdots & a_{1d} \\ & & \cdots & \\ & & \vdots & \\ & & \cdots & \\ a_{n1} & & \cdots & a_{nd} \end{bmatrix} \begin{bmatrix} x_1^* \\ \\ \vdots \\ x_d^* \end{bmatrix} \right|^2$$

- General inverse problem framework
- Phase retrieval (PR) applications
- PR algorithms
- PR theory: random model
- Machine learning

Projection-based algorithms

$$\Psi_{n+1} = \Psi_n + \Pi_{\mathrm{F}}[2\Pi_{\mathrm{O}}(\Psi_n) - \Psi_n] - \Pi_{\mathrm{O}}(\Psi_n)$$

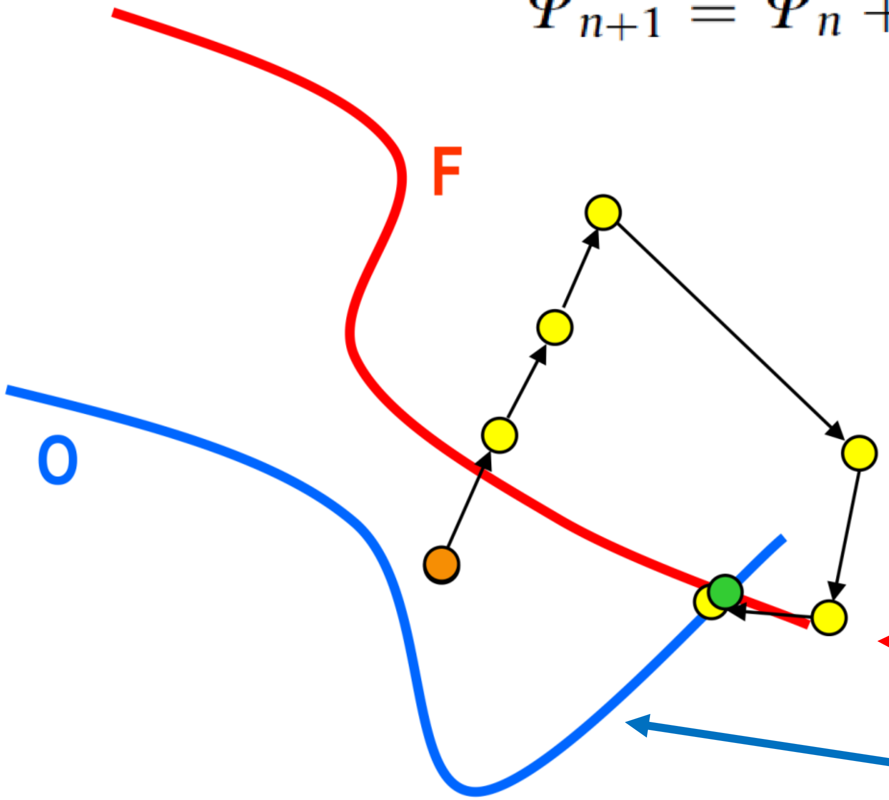

Use projections to find intersection between sets of constraints

Explore solution space and escapes local minima

Data fidelity (to match the measurements)

Object constraints (known support, positivity)

From previous lecture

Fienup '82
Maiden '09

Gradient descent

Find $\boldsymbol{x}^*$ in
$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

- Optimization approach:
$$\hat{x} = \mathbf{argmin}_x\, f(y, x) + g(x)$$
- Gradient descent to solve non-linear optimization problem

- Includes regularization
- Many variants / acceleration strategies

- Lacks theoretical guarantees (in general)

Fienup ’93
Yeh, Dong ’15
Chen ‘18

Convex relaxation

- Example: Phaselift (Candès '11)
- Trick:

$$
\begin{aligned}
y_i &= \left|a_i^H x\right|^2 \\
&= \left(a_i^H x\right)\left(x^H a_i\right) \\
&= a_i^H X a_i \text{ with } X = x x^H
\end{aligned}
$$

Candès '11
Waldspurger '12
Goldstein '16

Convex relaxation

- Example: Phaselift (Candès '11)
- Trick:

$$y_i = a_i^H X a_i \text{ with } X = xx^H$$

- Optimization:

$$\hat{X} = \mathrm{argmin}_{X \text{ rank } 1} \sum_i \left\| y_i - a_i^H X a_i \right\|^2$$

Candès '11
Waldspurger '12
Goldstein '16

Convex relaxation

- Example: Phaselift (Candès '11)
- Trick:

$$y_i = a_i^H X a_i \text{ with } X = xx^H$$

- Optimization with convex relaxation:

$$\hat{X} = \text{argmin}_X \sum_i \left\| y_i - a_i^H X a_i \right\|^2 + \lambda \mathrm{Tr}(X)$$

  - Can avoid local minima
  - Memory intensive (quadratic vs linear)

Candès '11
Waldspurger '12
Goldstein '16

See also PhaseCut, Phasemax, etc.

Spectral methods

- Intuition based on $y_i = |a_i^H x|^2$

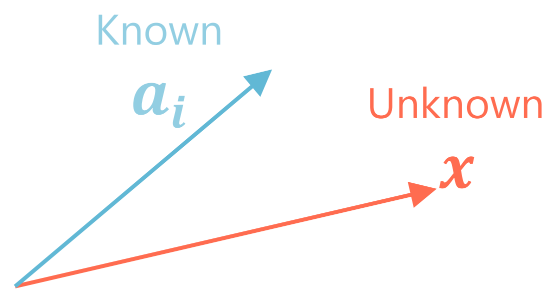

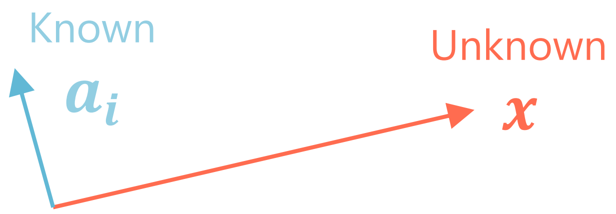

$\boldsymbol{a}_i$ and $\boldsymbol{x}$
correlated
$\Rightarrow$ high intensity y_i

$\boldsymbol{a}_i$ and $\boldsymbol{x}$
uncorrelated
$\Rightarrow$ low intensity y_i

Spectral methods

- Intuition based on $y_i = \left|a_i^H x\right|^2$

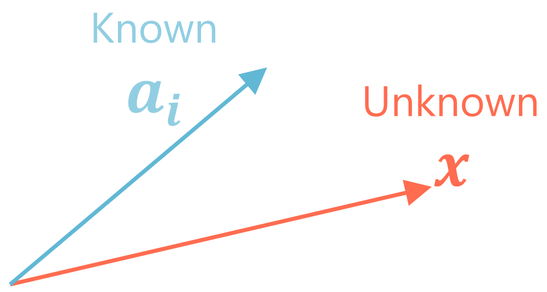

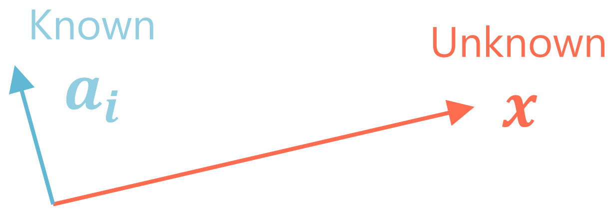

$\boldsymbol{a}_i$ and $\boldsymbol{x}$
correlated
$\Rightarrow$ high intensity y_i

$\boldsymbol{a}_i$ and $\boldsymbol{x}$
uncorrelated
$\Rightarrow$ low intensity y_i

Spectral method

"Construct a matrix giving more weight to a_i correlated with the unknown x"

Returns the leading eigenvector
of the weighted covariance matrix:

$$\boldsymbol{Z} = \sum_i y_i \boldsymbol{a}_i \boldsymbol{a}_i^H$$

Spectral methods

- Intuition based on $y_i = \left|a_i^H x\right|^2$

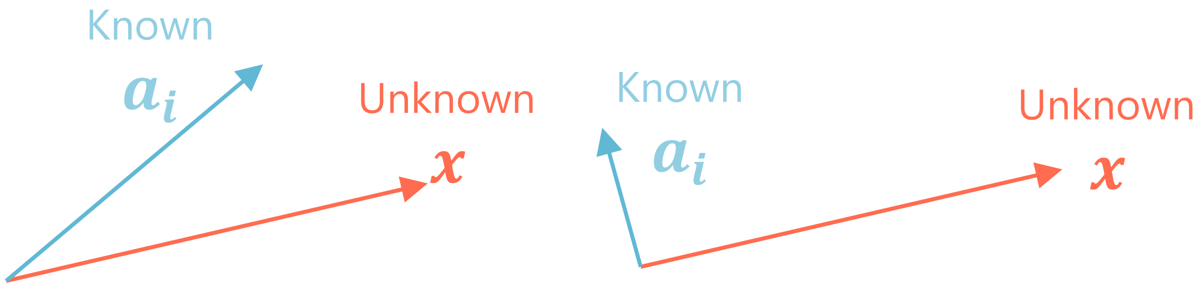

$\boldsymbol{a}_i$ and $\boldsymbol{x}$
correlated
$\Rightarrow$ high intensity y_i

$\boldsymbol{a}_i$ and $\boldsymbol{x}$
uncorrelated
$\Rightarrow$ low intensity y_i

Spectral method

"Construct a matrix giving more weight to a_i correlated with the unknown x"

Returns the leading eigenvector of the weighted covariance matrix:

$$\boldsymbol{Z} = \sum_i \mathcal{T}(y_i)\boldsymbol{s_i}\boldsymbol{s_i}^\dagger$$

for $\mathcal{T}$ an increasing preprocessing function

Spectral methods

- Intuition based on $y_i = \left|a_i^H x\right|^2$

$\mathcal{T}_0(y) = y$	*E. Candes, et al, IEEE Trans. on Information Theory (2015)*
$\mathcal{T}_1(y) = (y > T)$	*S. Marchesini, et al, Applied and Comput. Harmonic Analysis (2016)*
$\mathcal{T}_{\text{optim}}(y) = 1 - \frac{1}{y}$	*W. Luo, et al, IEEE Trans. on Signal Processing (2019)*

Spectral method

"Construct a matrix giving more weight to a_i correlated with the unknown x"

Returns the leading eigenvector of the weighted covariance matrix:

$$\boldsymbol{Z} = \sum_i \mathcal{T}(y_i)\boldsymbol{s_i}\boldsymbol{s_i}^\dagger$$

for $\mathcal{T}$ an increasing preprocessing function

Algorithm taxonomy

Alternating projections	Gradient-based optimization	Convex relaxation	Bayesian AMP	Spectral methods
Fienup ‘82 Maiden ‘09	Fienup ’93 Yeh, Dong ’15 Chen ‘18	Candès ‘11 Waldspurger ‘12 Goldstein ‘16	Rangan ’10 Metzler ‘17 Barbier ‘17 Maillard ‘20	Candès ‘15 Lu ‘17 Mondelli ‘18 Luo ‘18

Algorithms comparison

Name	Computational speed	Performance	Designed for the random setting
Alternating projections	★★★	★☆☆	
Gradient-based optimization	★★★	★☆☆	
Convex relaxation	★★☆	★★☆	
Approximate Message Passing	★☆☆	★★★	Yes
Spectral methods	★★★	★★☆	Yes

Content

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

$$\begin{bmatrix} y_1 \\ \\ \vdots \\ \\ y_n \end{bmatrix} = \left| \begin{bmatrix} a_{11} & & \cdots & a_{1d} \\ & & \cdots & \\ & & \vdots & \\ & & \cdots & \\ a_{n1} & & \cdots & a_{nd} \end{bmatrix} \begin{bmatrix} x_1^* \\ \\ \vdots \\ x_d^* \end{bmatrix} \right|^2$$

- General inverse problem framework
- Phase retrieval (PR) applications
- PR algorithms
- PR theory: random model
- Machine learning

Unifying framework $y = |Ax|^2$

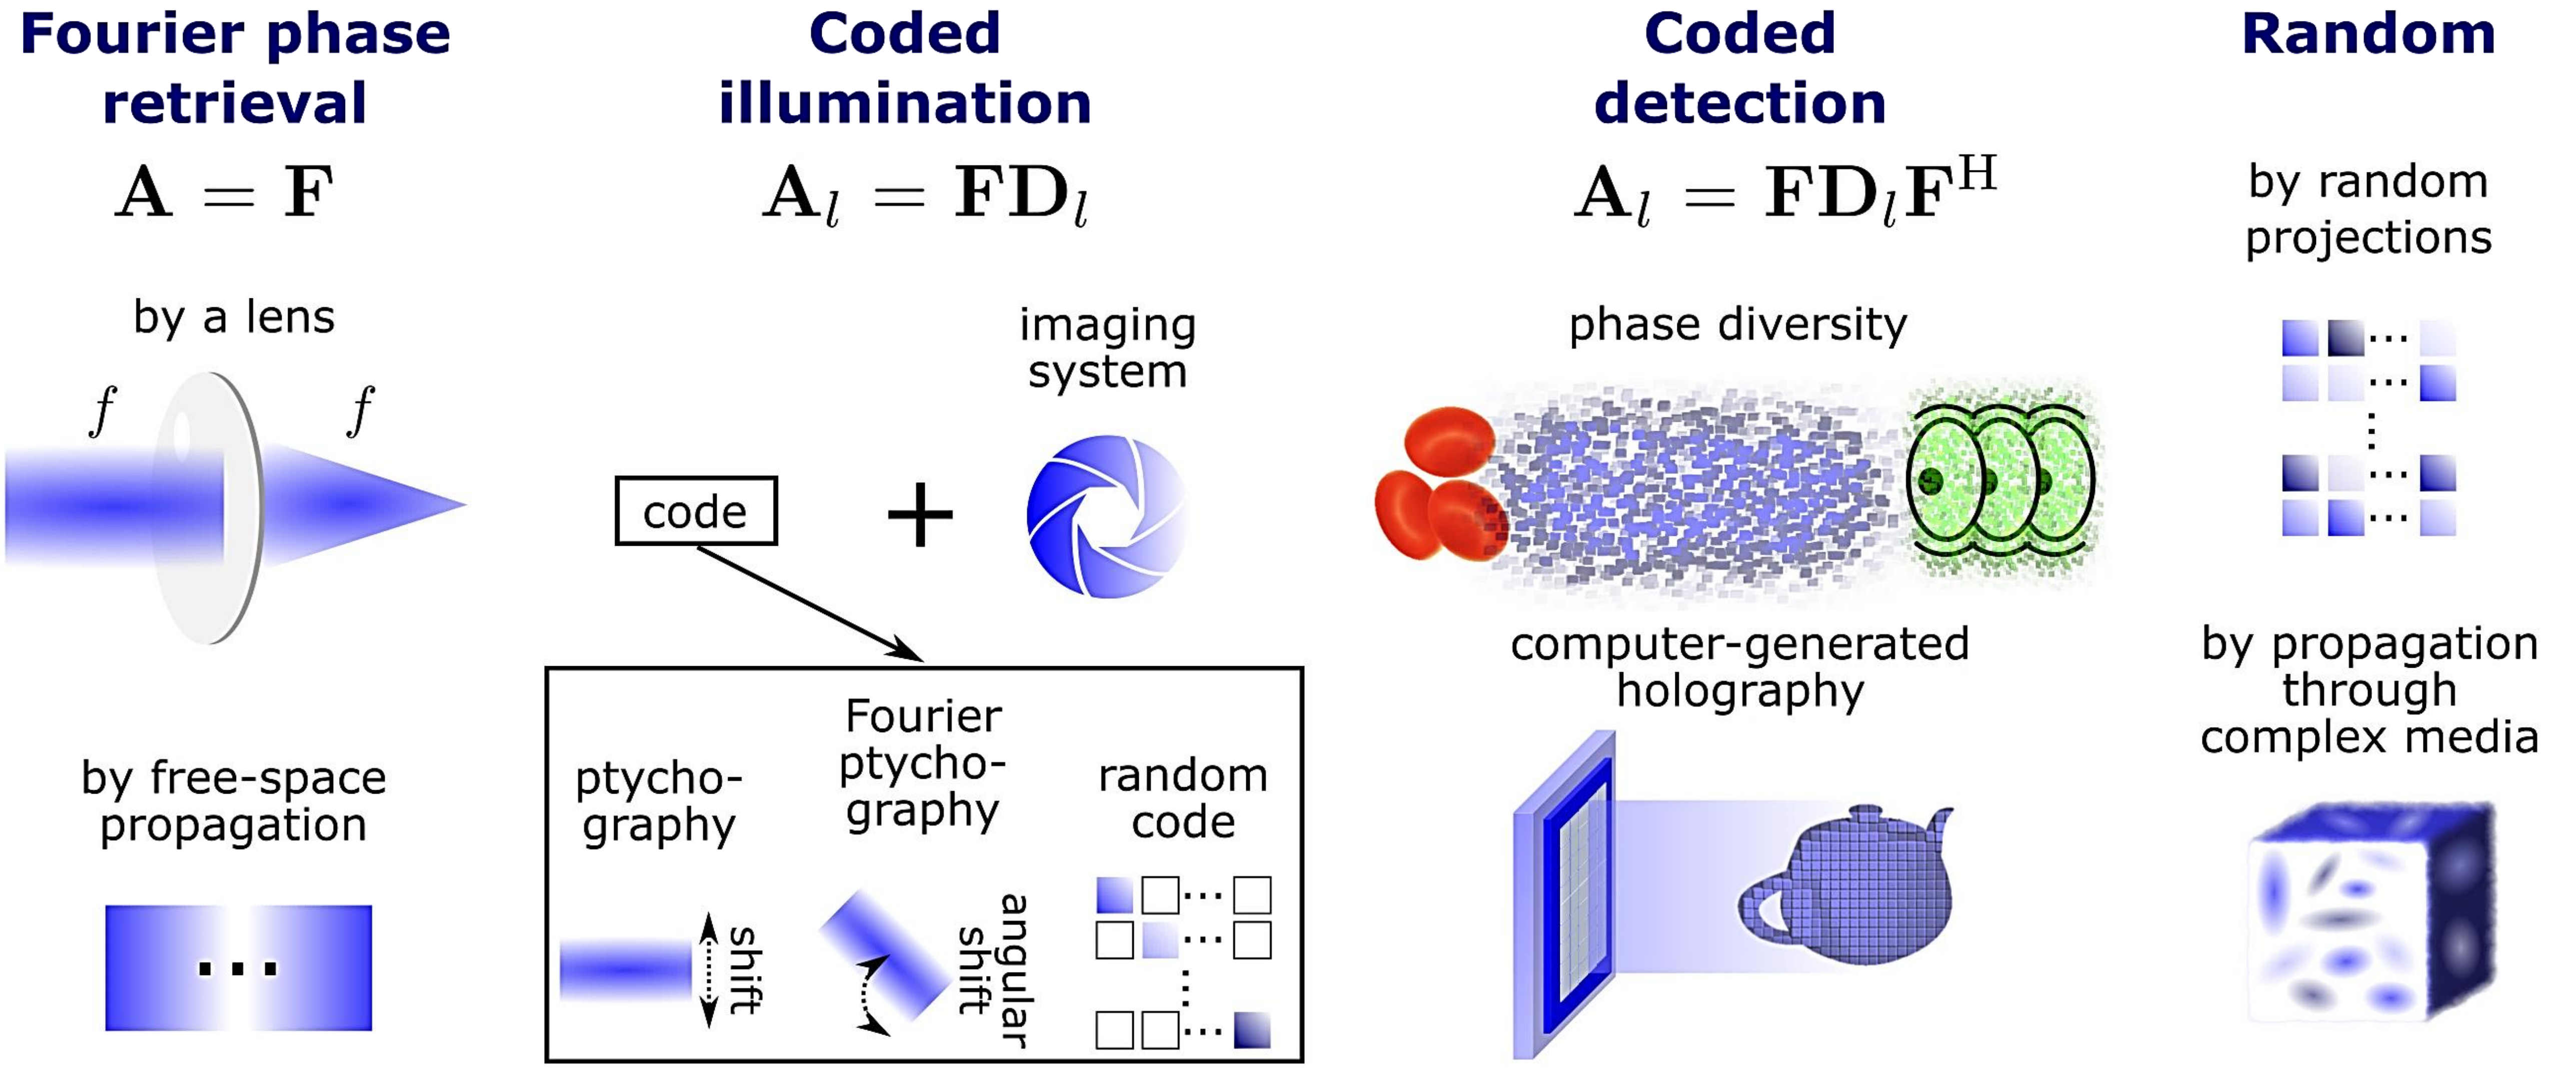

Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). *IEEE Signal Processing Magazine*

Deeper dive in the random setting

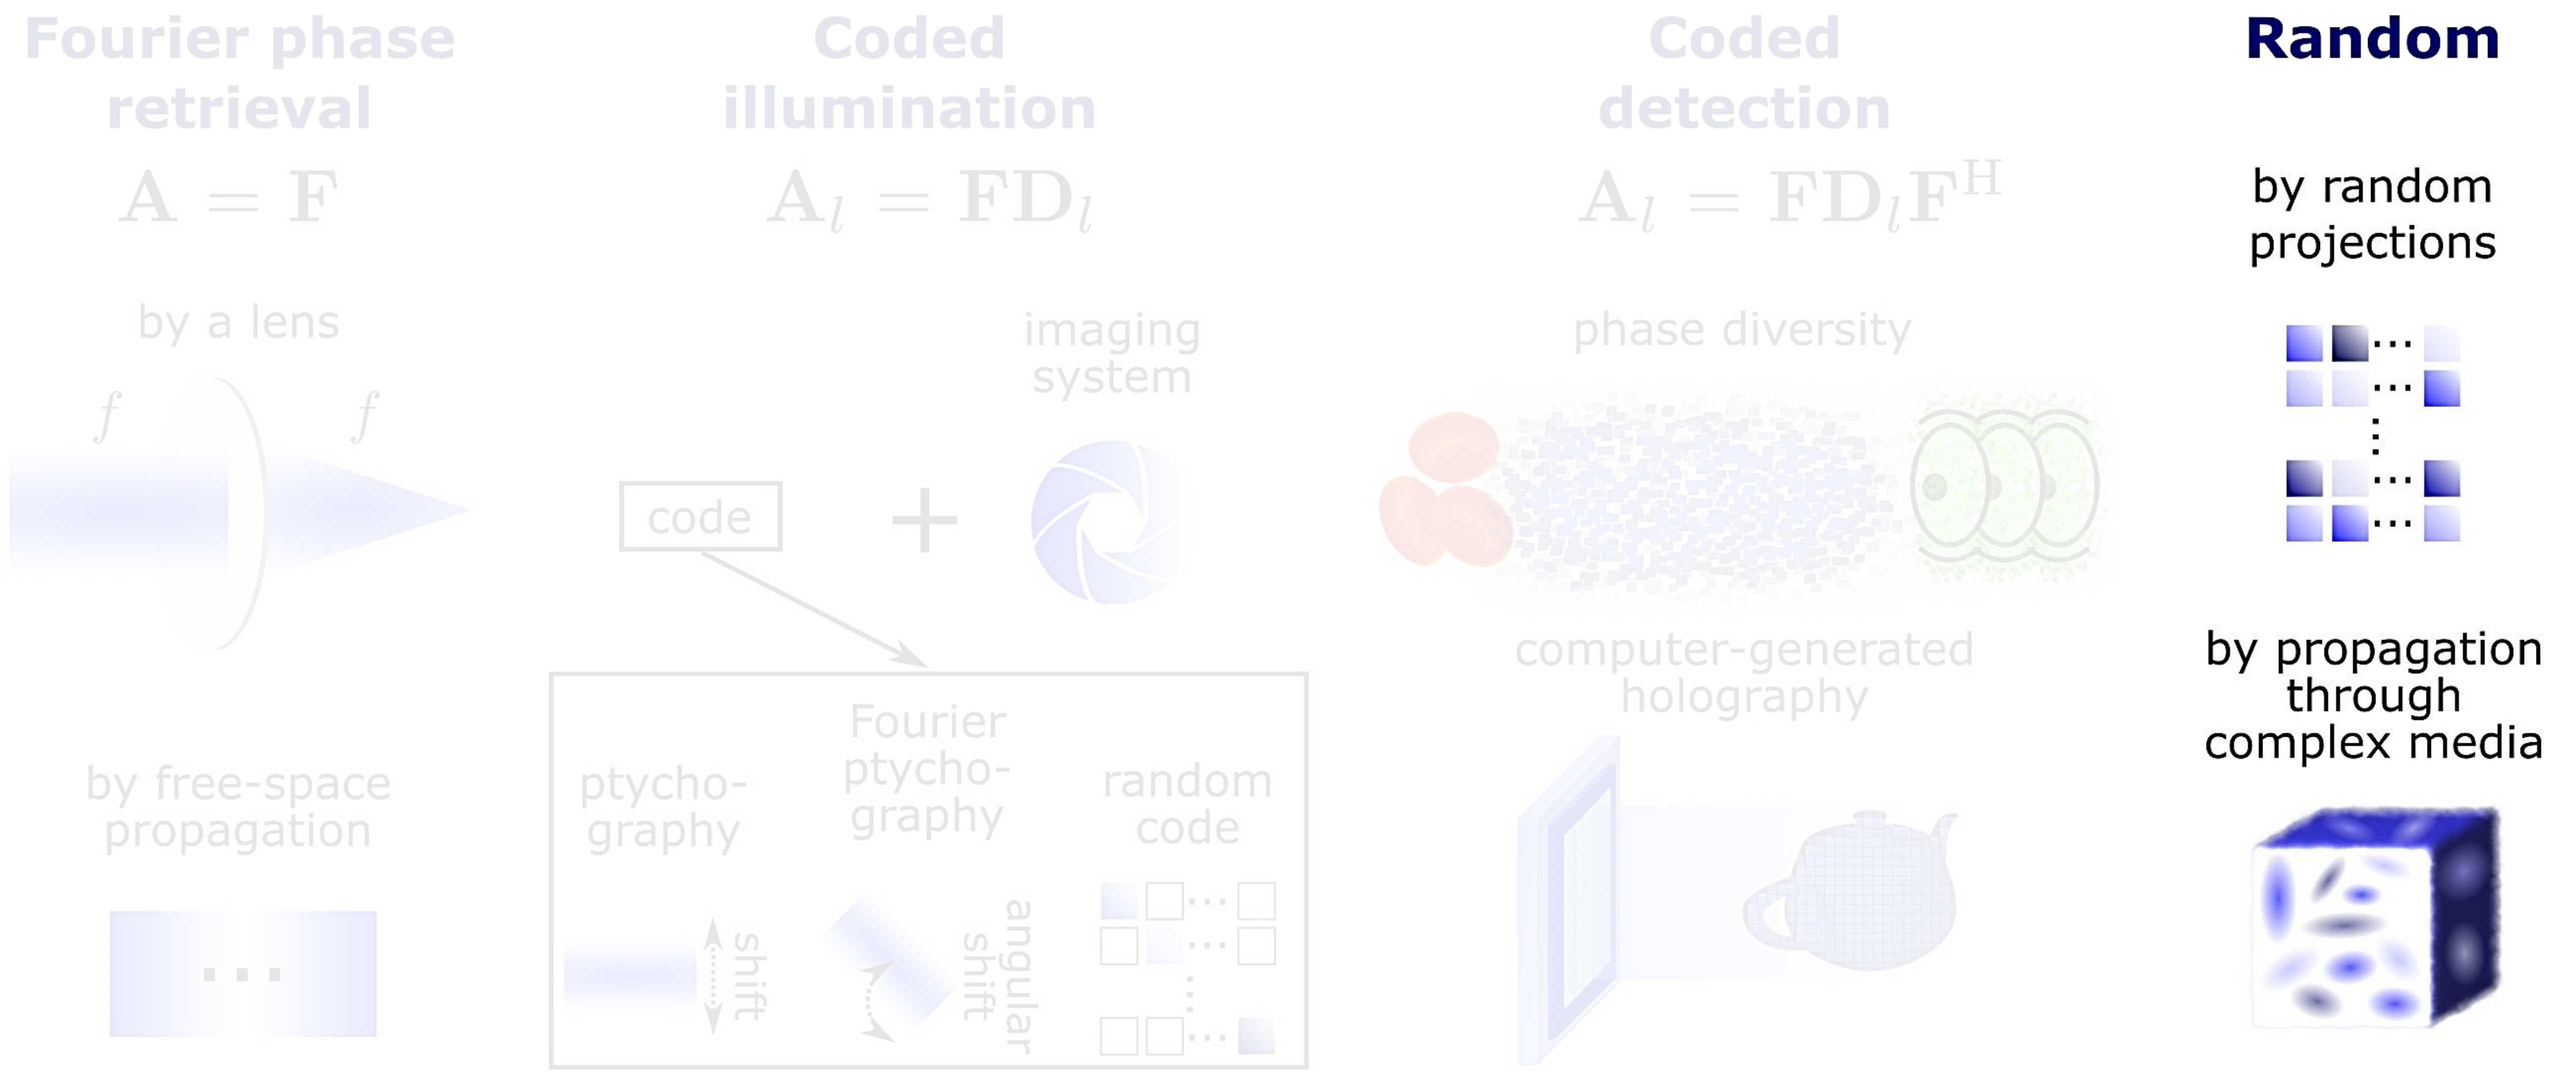

Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). *IEEE Signal Processing Magazine*

Theory break

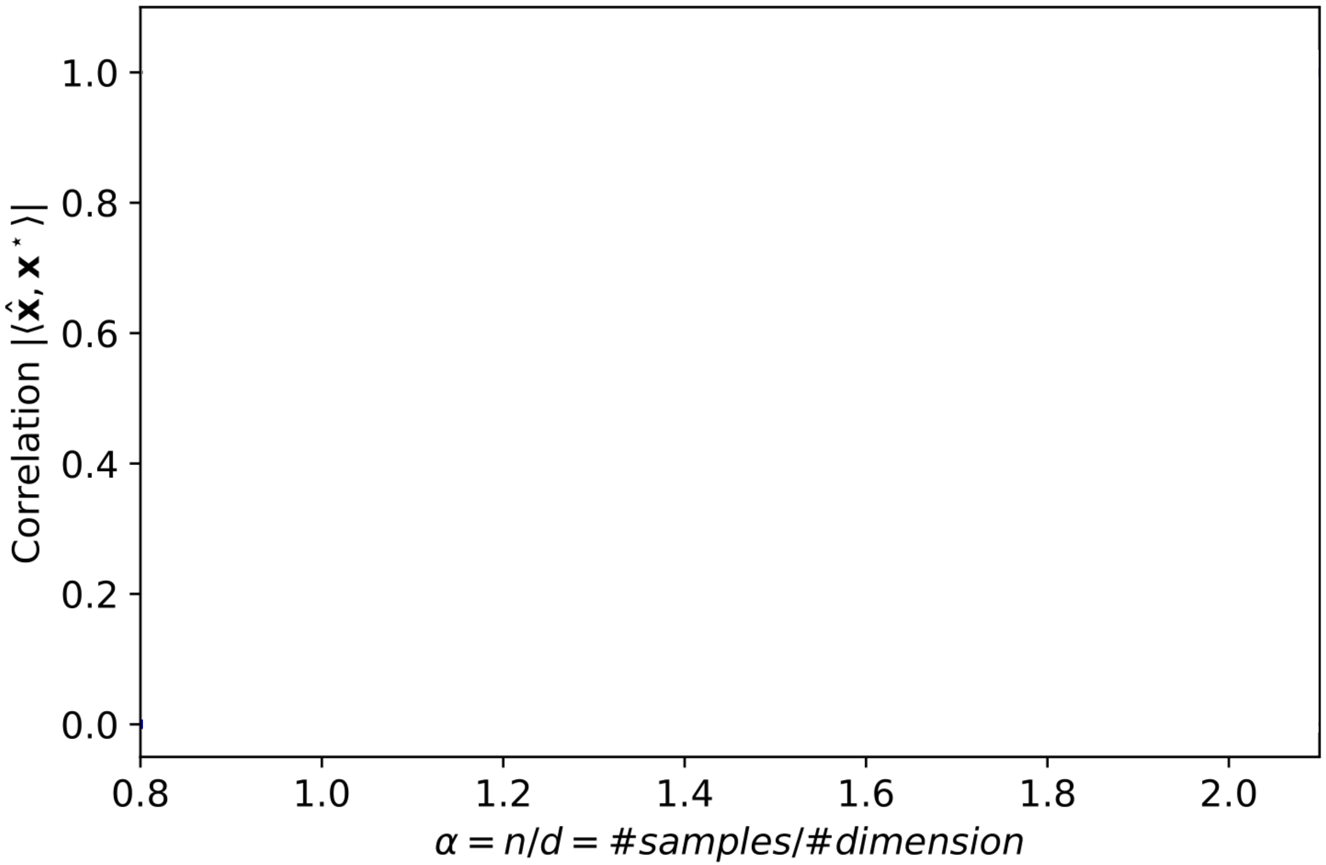

- Can we characterize performance as a function of oversampling $\alpha = n/d$?
- Correlation = higher is better

Weak recovery

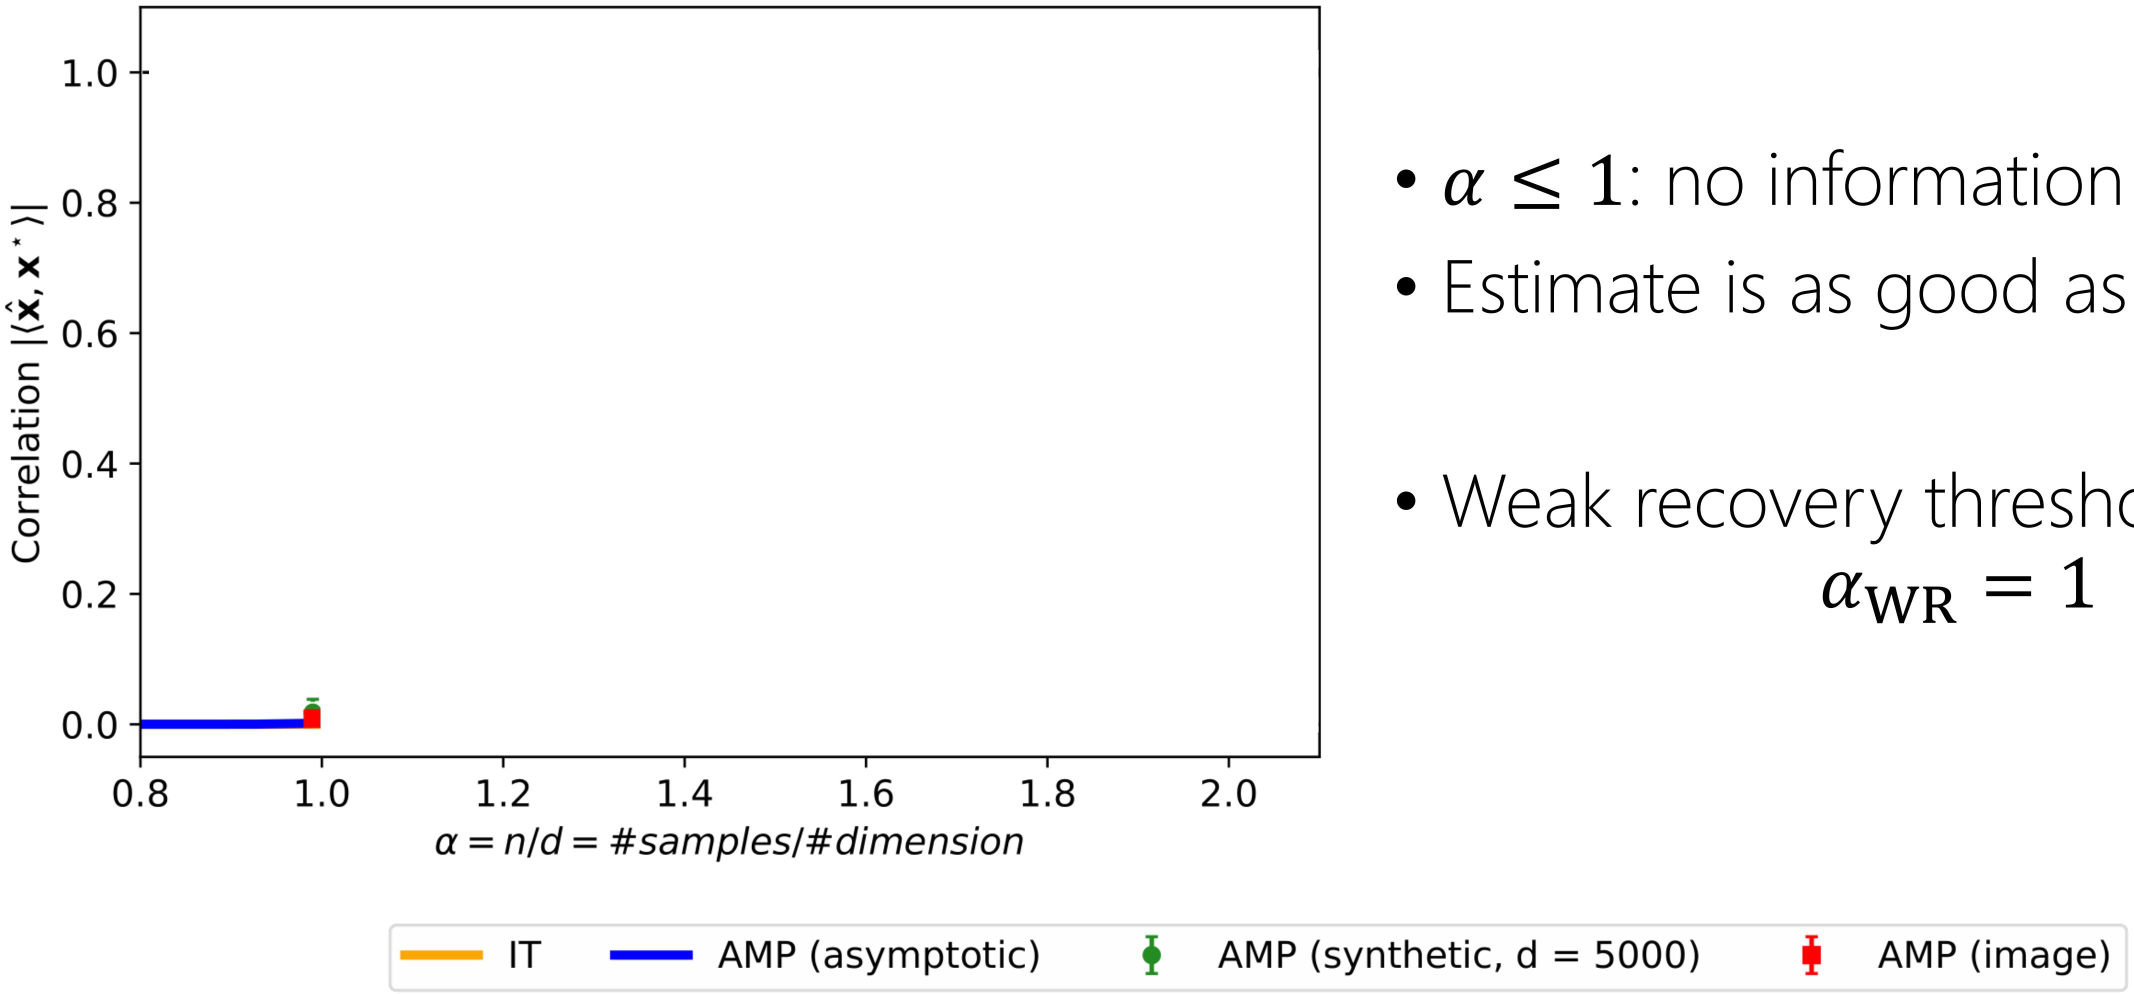

- $\alpha \leq 1$: no information on solution
- Estimate is as good as random

- Weak recovery threshold:

$$\alpha_{\mathrm{WR}} = 1$$

Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). IEEE Signal Processing Magazine

Perfect recovery

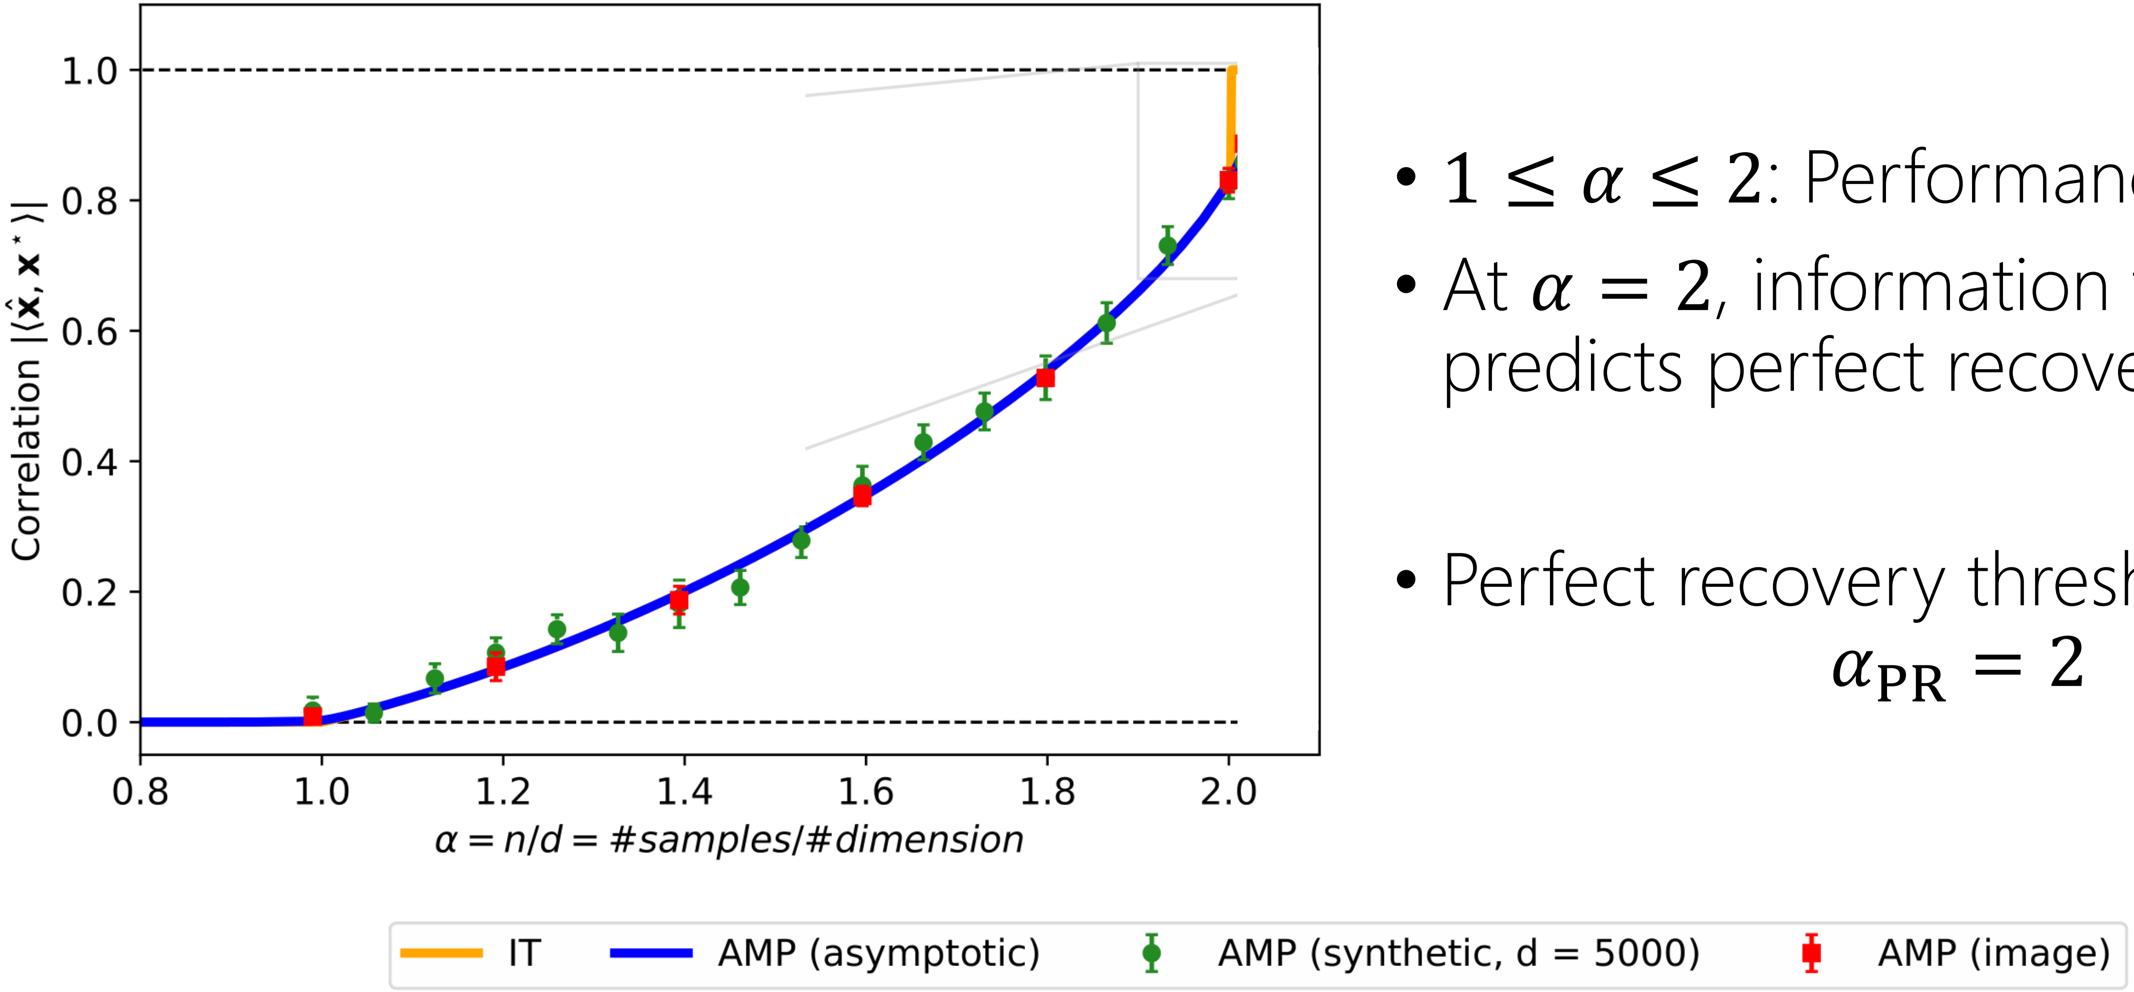

- $1 \leq \alpha \leq 2$: Performance improves
- At $\alpha = 2$, information theory predicts perfect recovery
- Perfect recovery threshold:

$$\alpha_{\mathrm{PR}} = 2$$

Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). IEEE Signal Processing Magazine

In practice

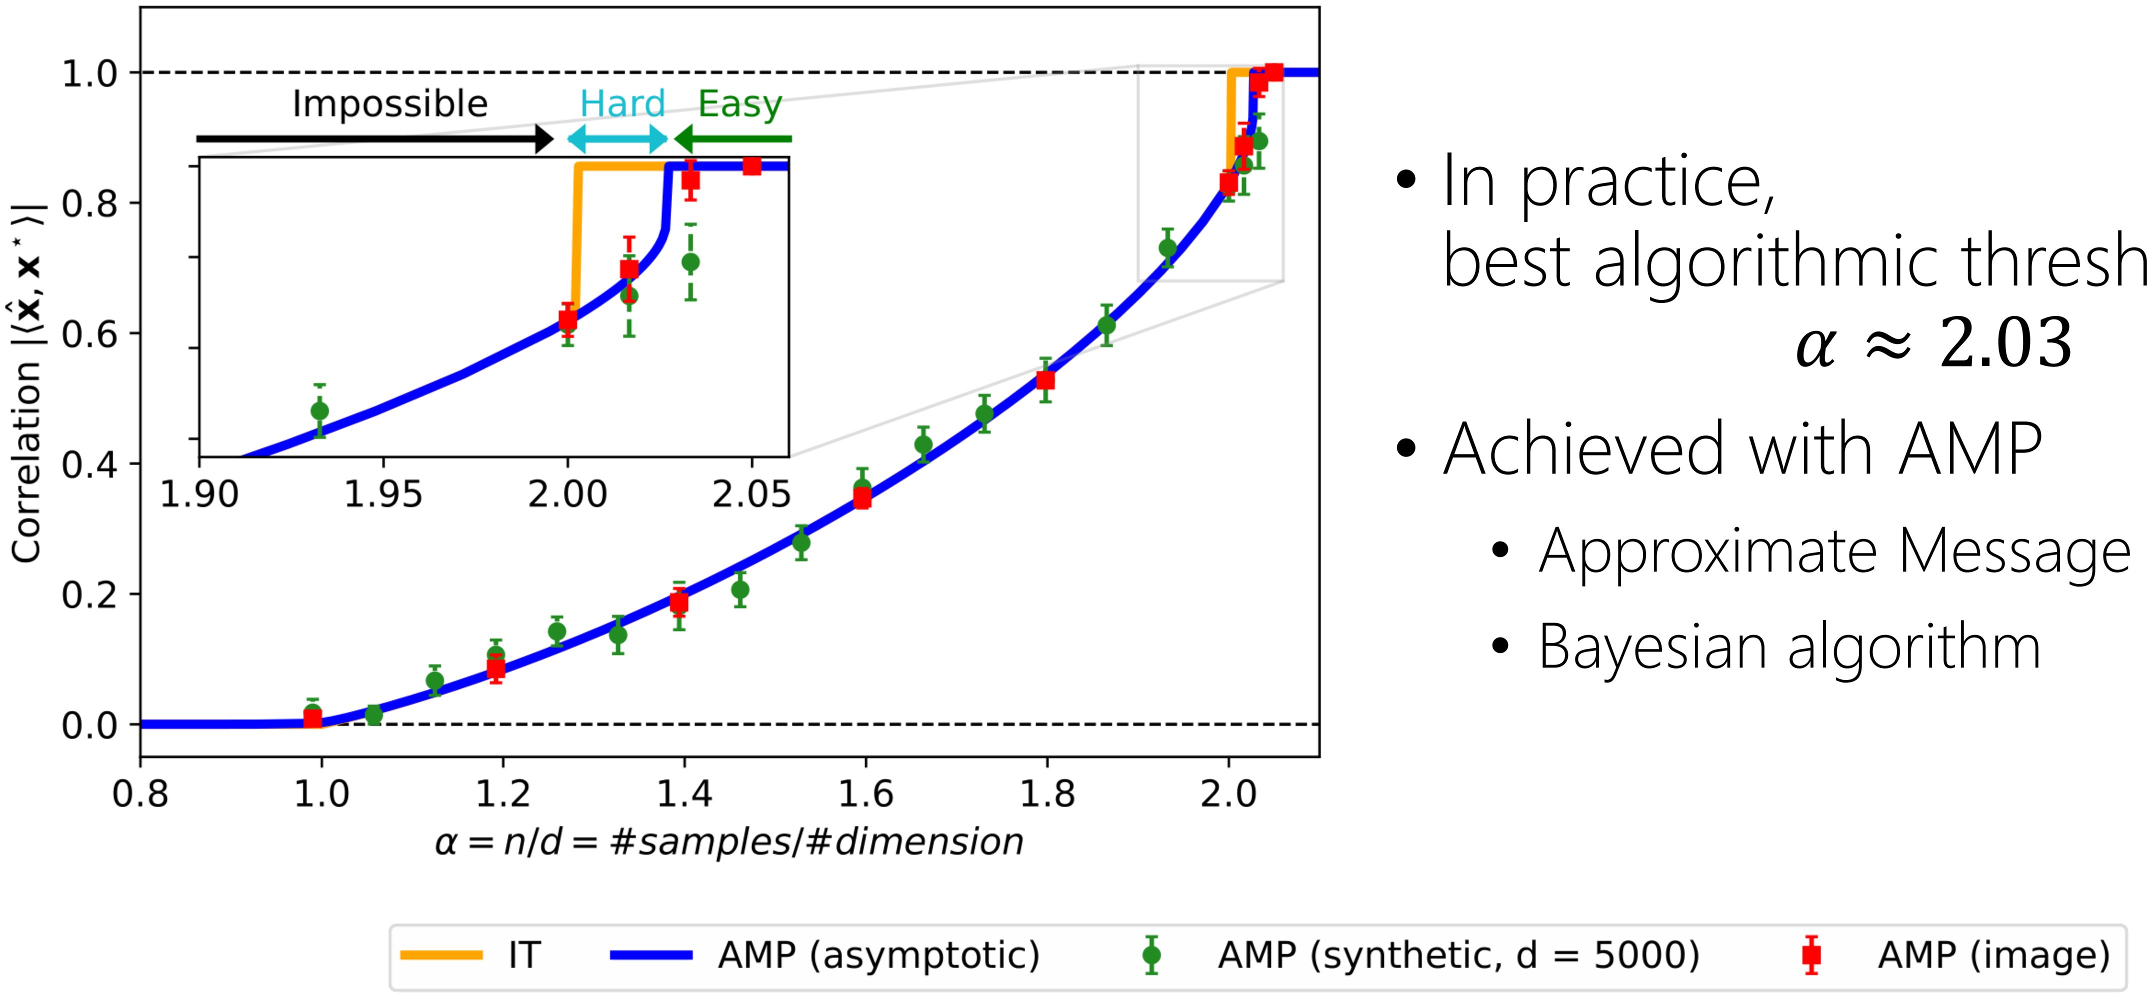

- In practice,
  best algorithmic threshold:
  $$\alpha \approx 2.03$$
- Achieved with AMP
  - Approximate Message Passing
  - Bayesian algorithm

Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). IEEE Signal Processing Magazine

Example

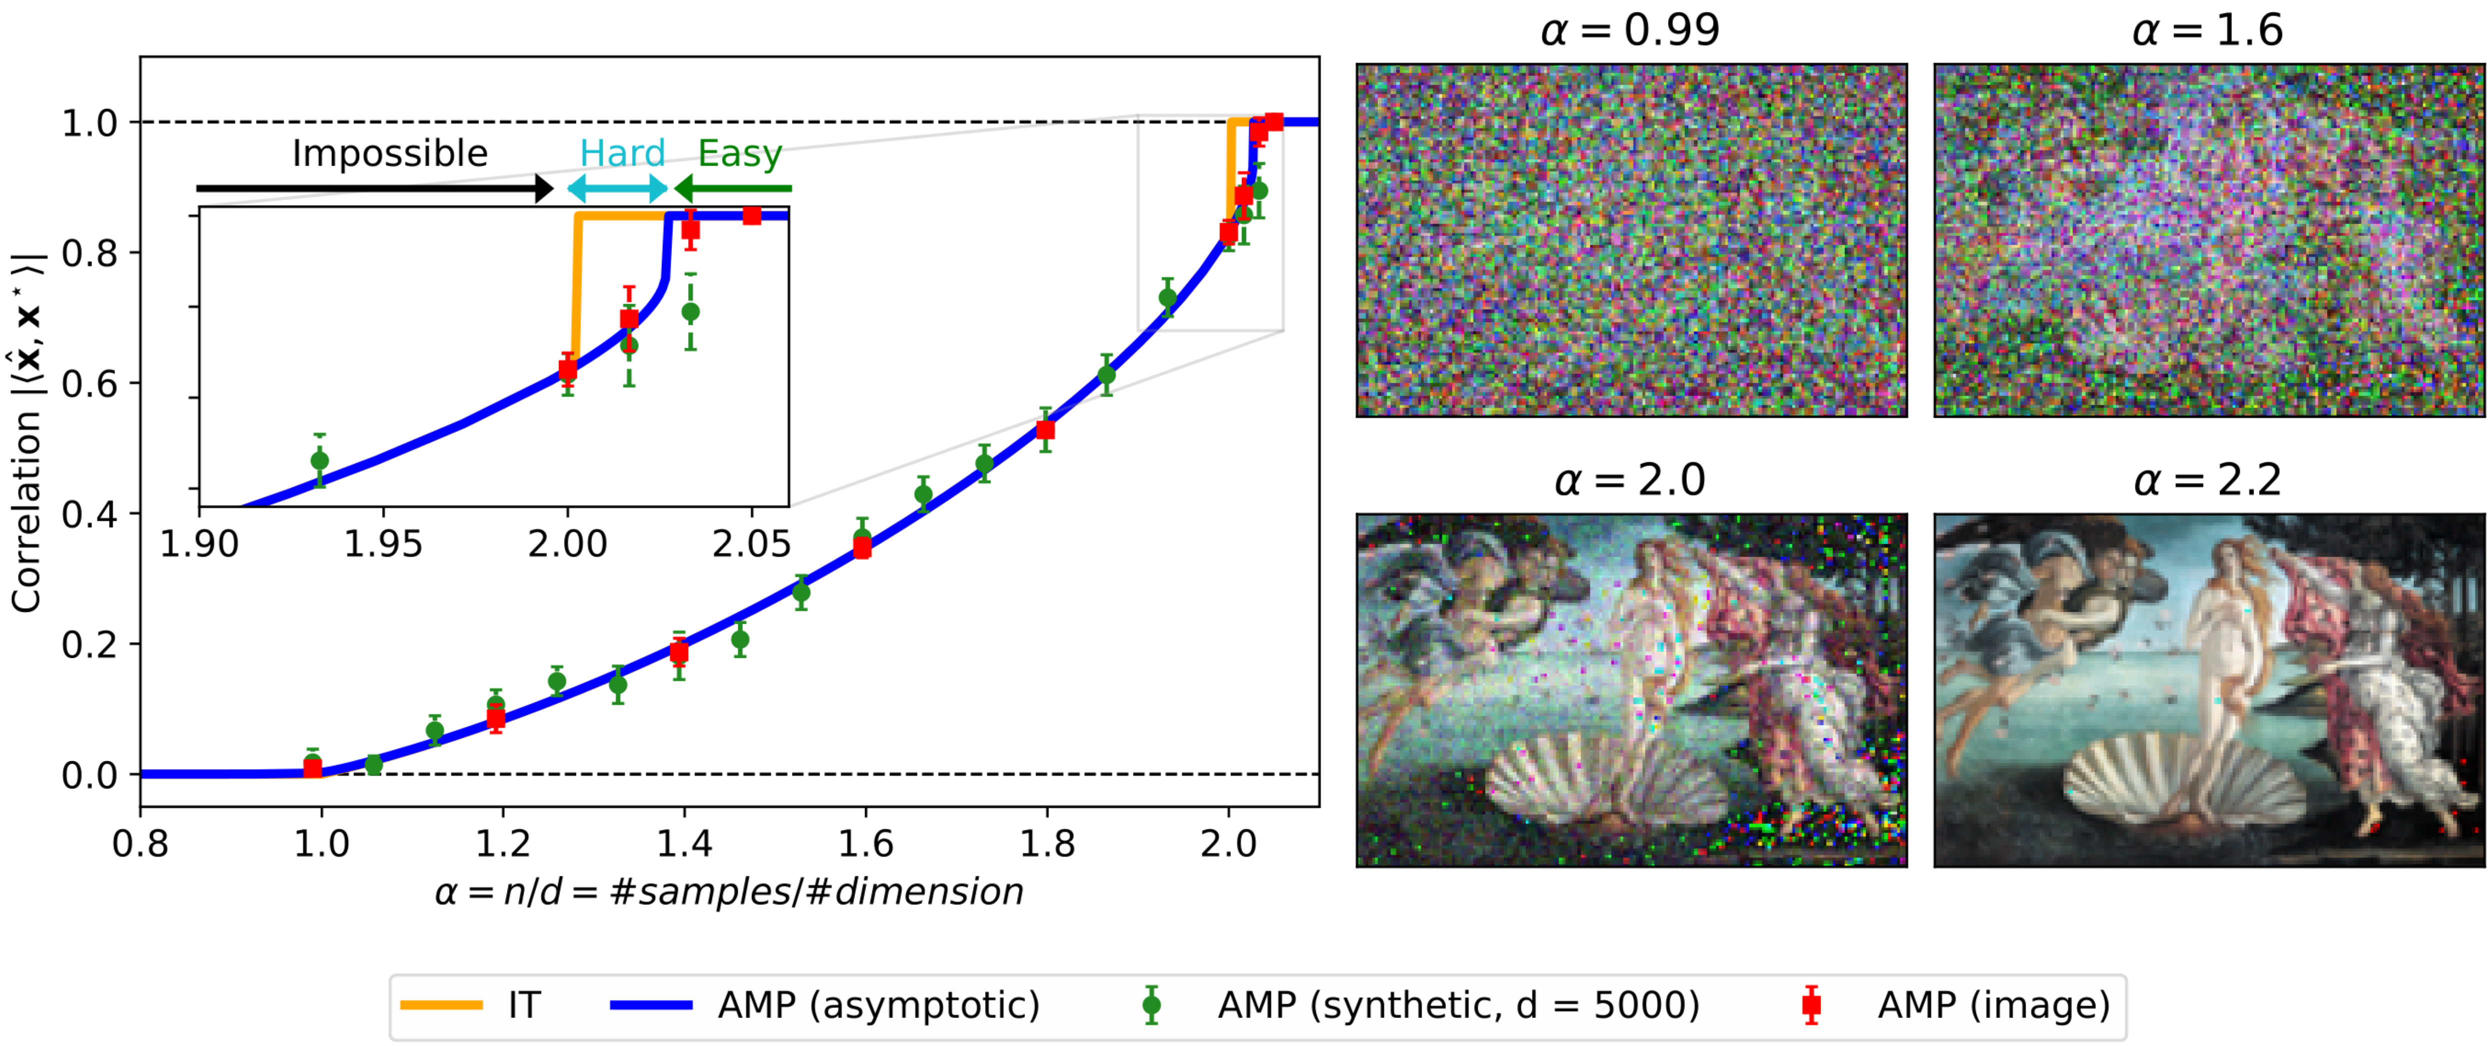

Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). IEEE Signal Processing Magazine

And in practice?

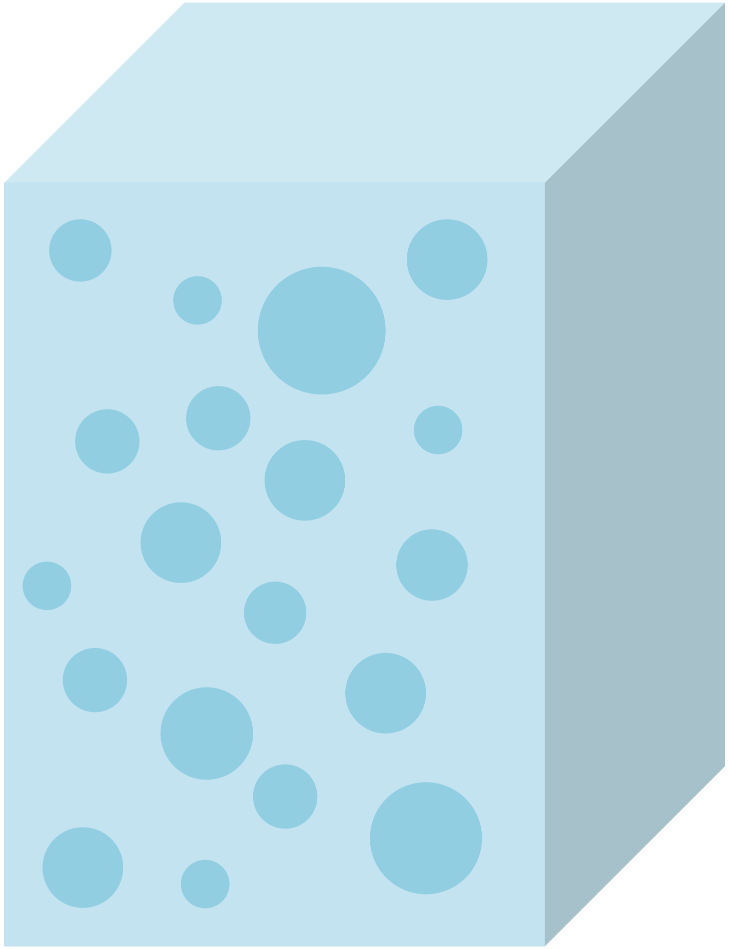

How to get a random matrix in optics?

Thanks to multiple scattering

Examples

Fog

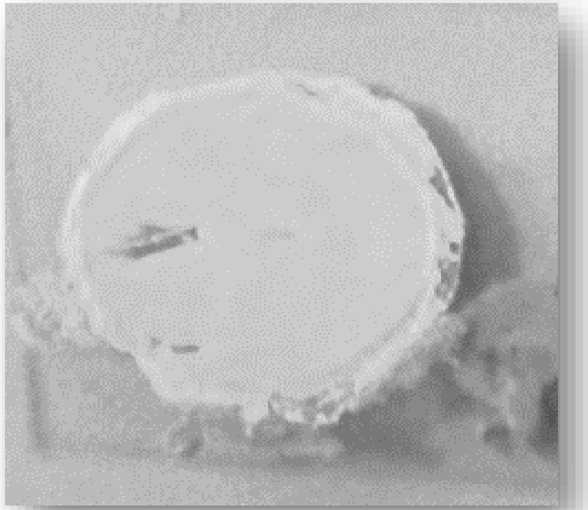

White paint

Biological tissue

Light scattering

E^{in}

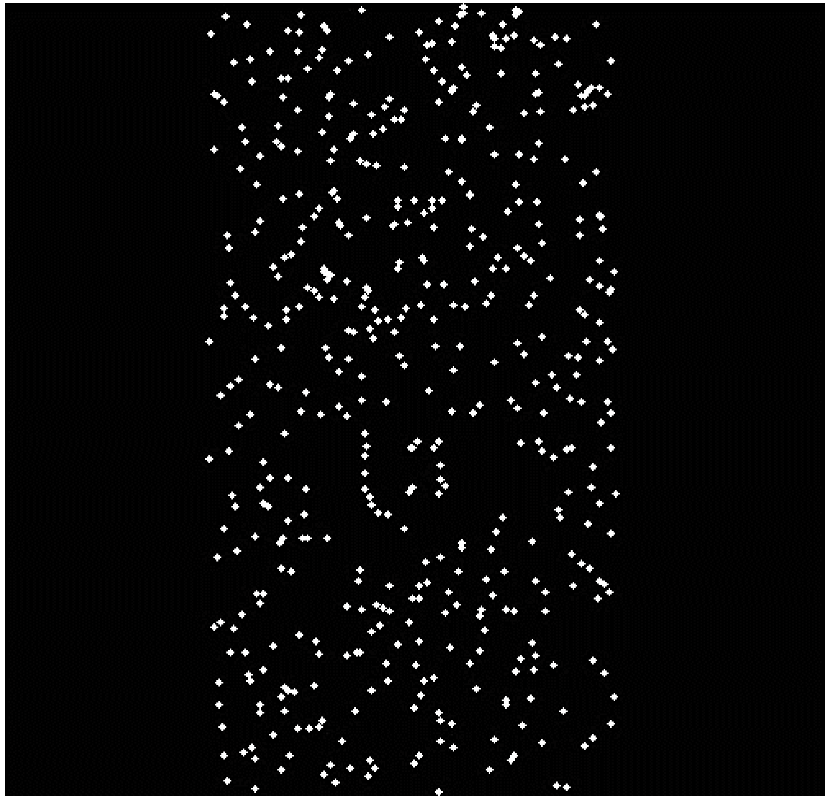

E^{out}

Credits: E. Bossy
(LiPHY Grenoble)

Random interference
→ speckle pattern

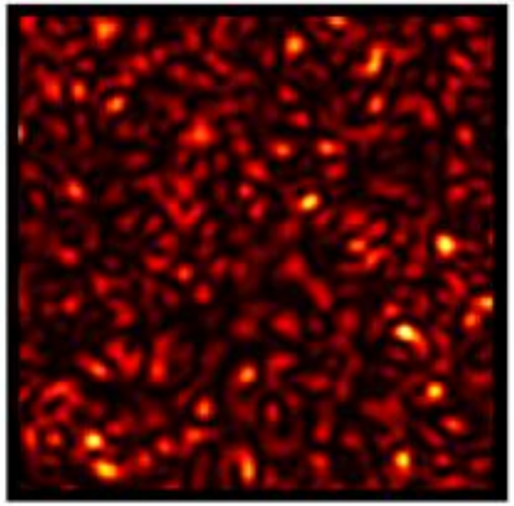

Still linear!
$E^{out} = AE^{in}$

(assuming monochromatic coherent light)

Popoff, Lerosey, Carminati, Fink, Boccara, Gigan (2010). *Physical Review Letters*

Structured-random example

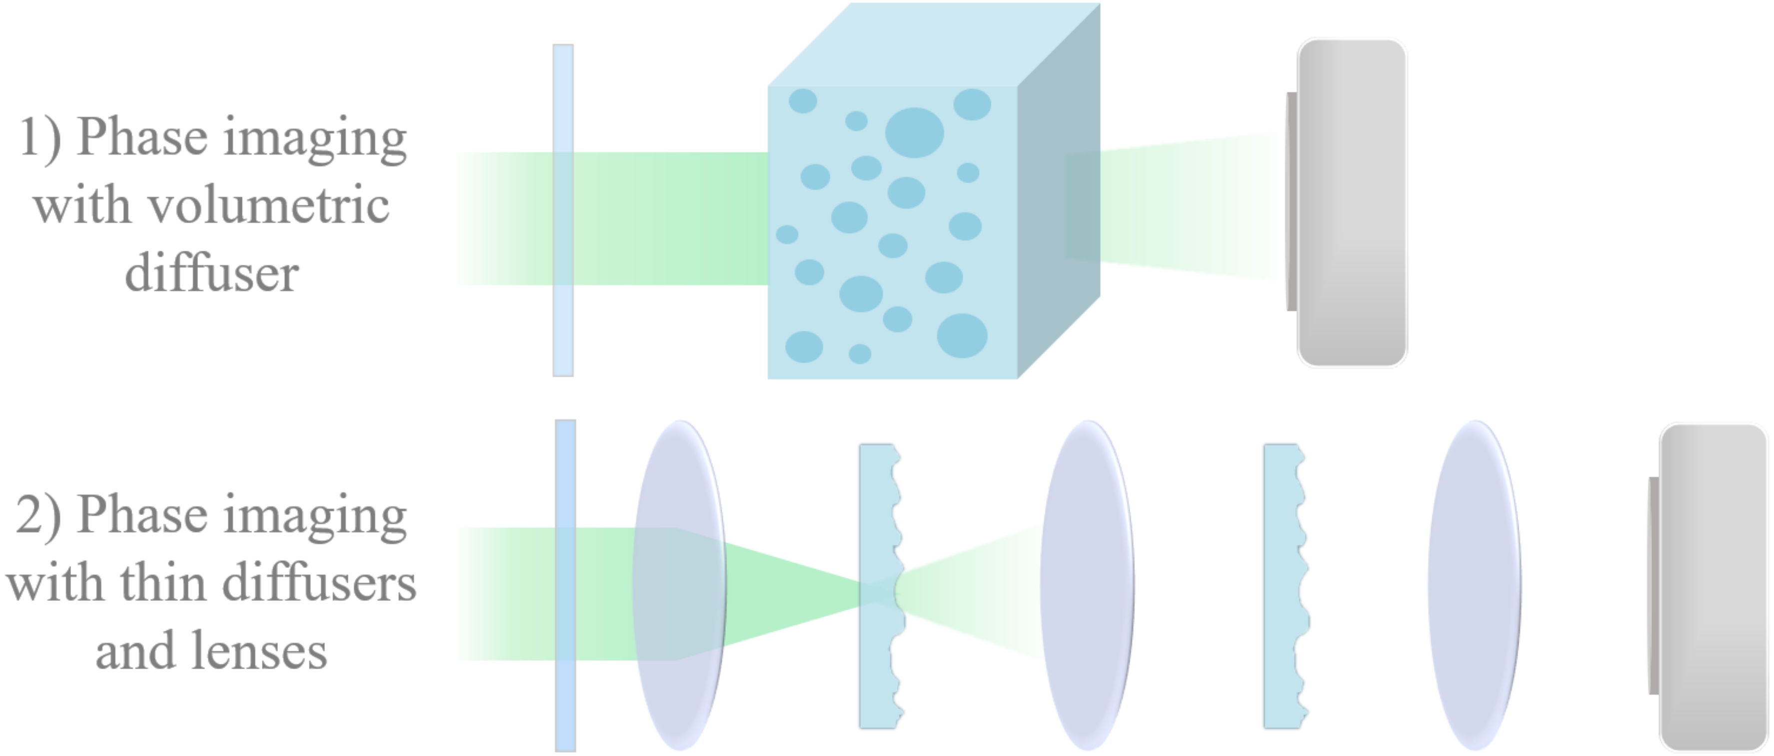

Lens = Fourier transform
Diffuser = Multiplication by diagonal matrix

Final operator: $A = FD_1FD_2F$

Structured-random models

Name	Model	Computational complexity	Storage complexity	Compatible with spectral methods	Compatible with AMP
Random model	$A_{ij} \sim p(a)$ i.i.d.	$O(n^2)$	$O(n^2)$	Yes	Yes
Pseudo-random models (3)	$A = FD_1FD_2FD_3F'$ with D_k random diagonal and F' upsampled Fourier	$O(n \log n)$	$O(n)$	Yes?	Yes?
Pseudo-random models (2)	$A = FD_1FD_2F'$ with D_k random diagonal and F' upsampled Fourier	$O(n \log n)$	$O(n)$	Yes?	Yes?
Pseudo-random models (1)	$A = FD_1F'$ with D_k random diagonal and F' upsampled Fourier	$O(n \log n)$	$O(n)$	?	?
Random Coded Diffractive Imaging	Concatenation of $A_l = FD_l$ with D_l random diagonal	$O(n \log n)$	$O(n)$	?	No?
Random Probe Ptychography	Concatenation of $A_l = FD_l$ with D_l shifted probe vector	$O(n \log n)$	$O(n)$	?	No?

Reconstruction results for $FDFD$

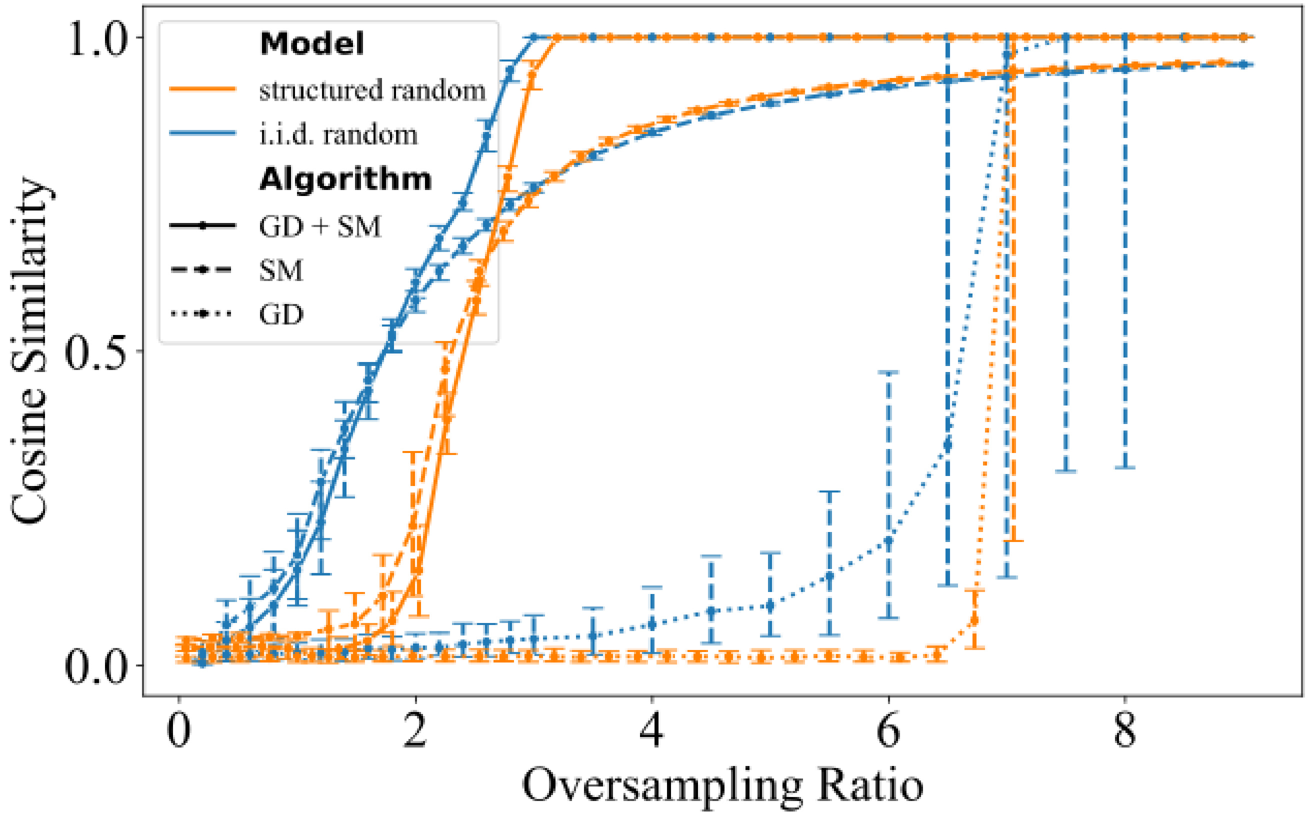

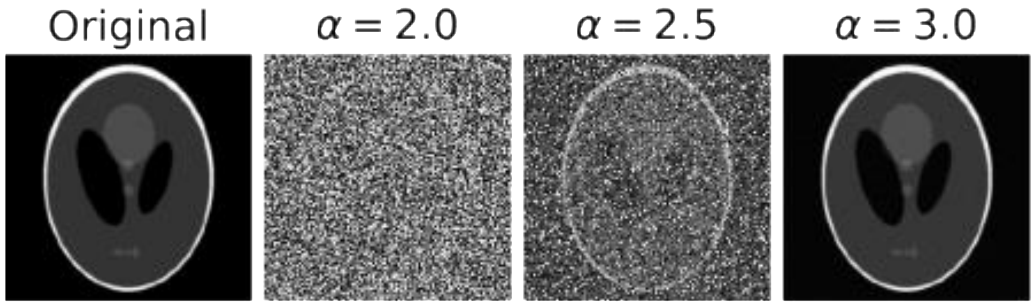

Hu, Tachella, Unser, Dong (2024). *submitted*

Additional results

How many layers?

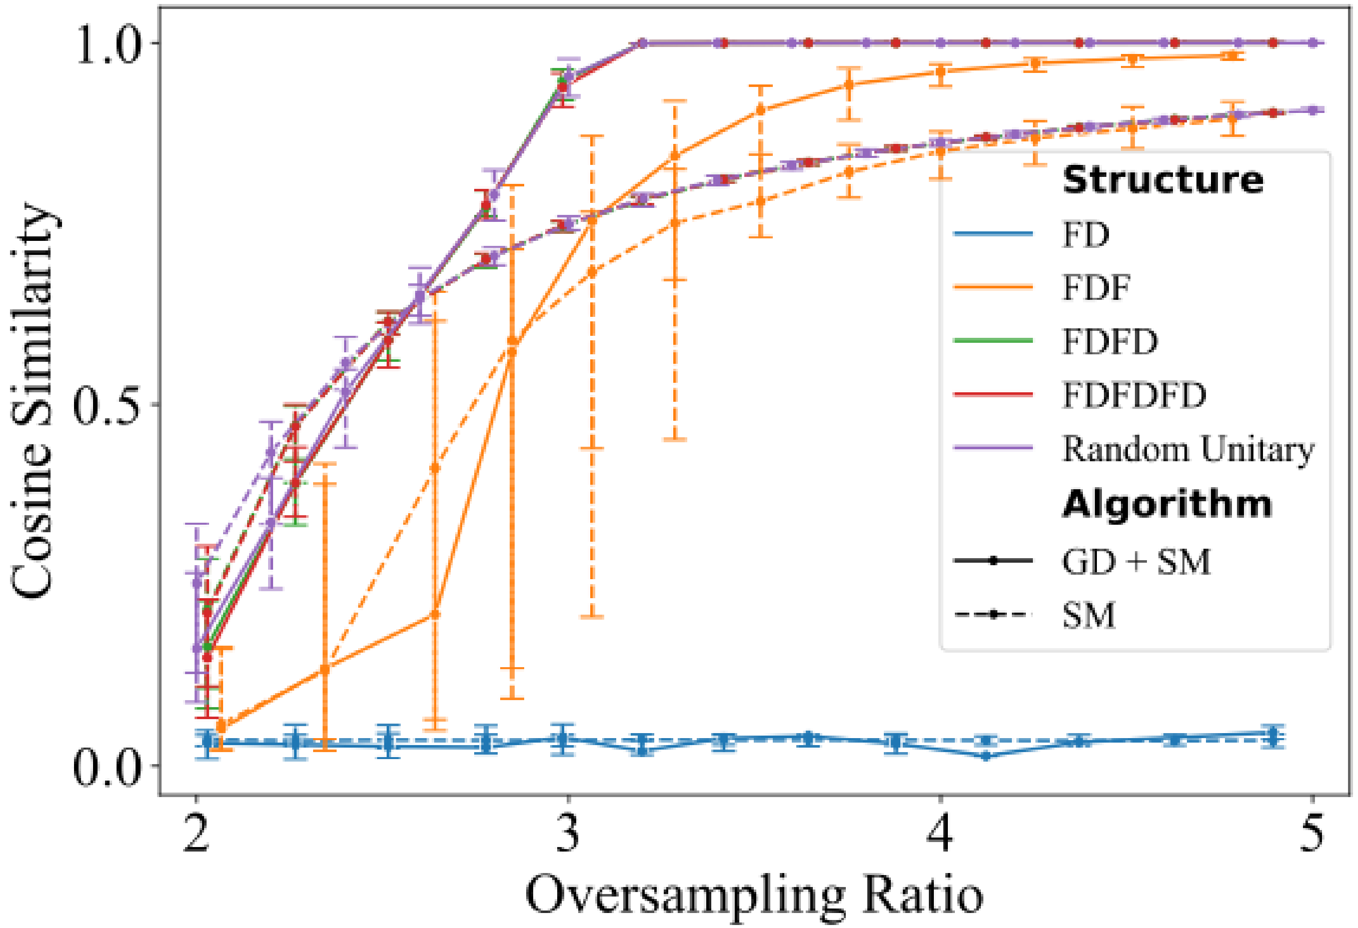

Speed benchmark

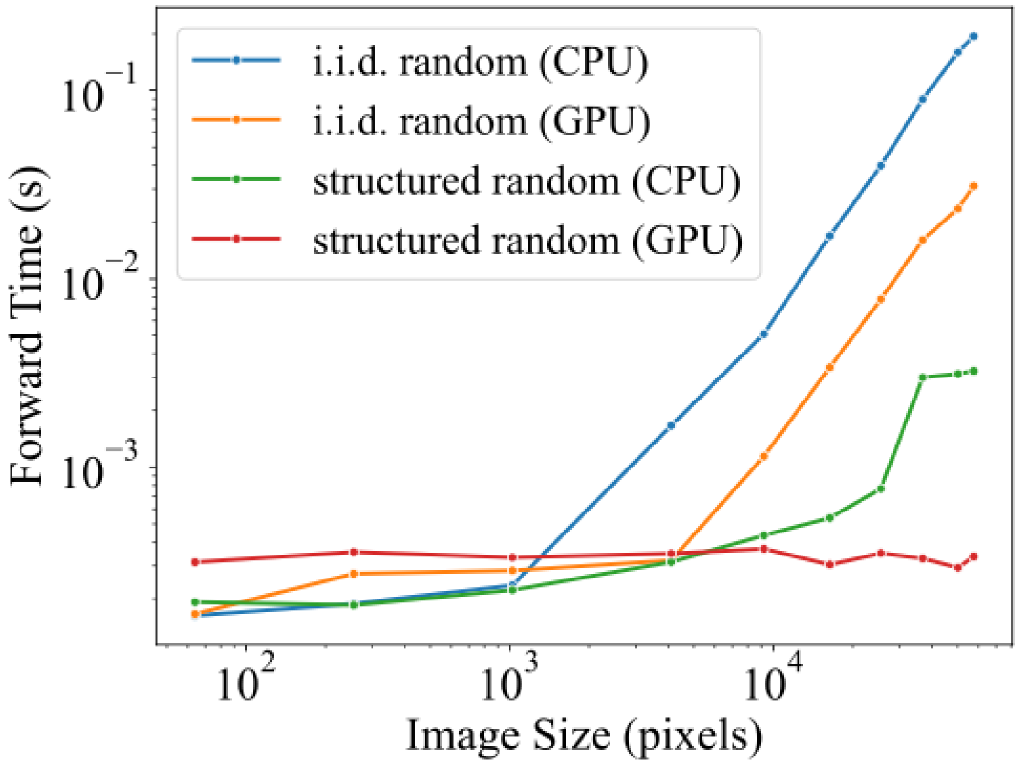

Content

Find $\boldsymbol{x}^*$ in

$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

$$\begin{bmatrix} y_1 \\ \\ \vdots \\ \\ y_n \end{bmatrix} = \left| \begin{bmatrix} a_{11} & & \cdots & a_{1d} \\ & & \cdots & \\ & & \vdots & \\ & & \cdots & \\ a_{n1} & & \cdots & a_{nd} \end{bmatrix} \begin{bmatrix} x_1^* \\ \\ \vdots \\ x_d^* \end{bmatrix} \right|^2$$

- General inverse problem framework
- Phase retrieval (PR) applications
- PR algorithms
- PR theory: random model
- Machine learning

Regularization

Add information about typical solutions to help reconstruction

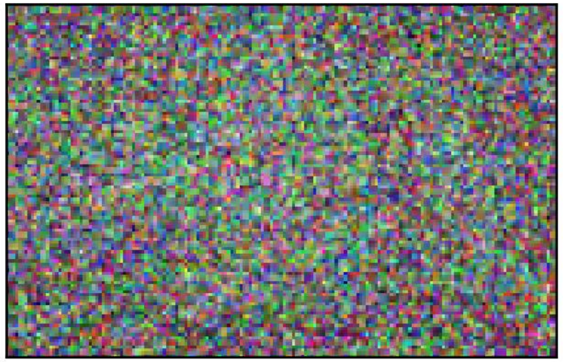

Regularization

- Non-linear optimization formulation:

$$\mathcal{L}(\boldsymbol{x}, \boldsymbol{y})$$

Data-consistency term

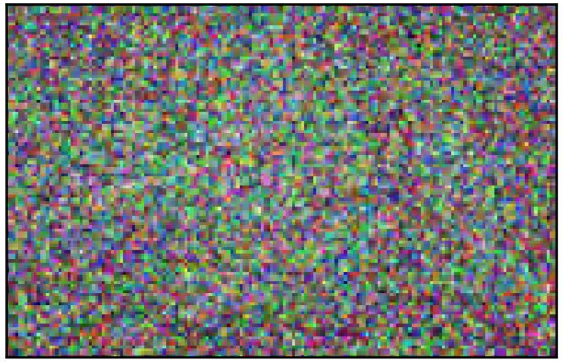

Regularization

- Add a regularization term:

$$\mathcal{L}(\boldsymbol{x}, \boldsymbol{y}) + \mathcal{R}(\boldsymbol{x})$$

Promotes realistic images

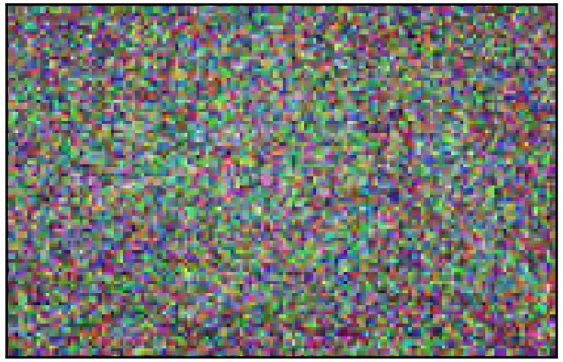

- Classical regularization
  - Sparsity $\mathcal{R}(\boldsymbol{x}) = \|\boldsymbol{x}\|_1$
  - Total variation $\mathcal{R}(\boldsymbol{x}) = \|\boldsymbol{\nabla x}\|_1$

Shechtman, Eldar, Cohen, Chapman, Miao, Segev (2015). *IEEE Signal Processing Magazine*

Deep learning regularization

- Step 1: Train a neural network f for denoising

Learn what is a realistic image

Denoising function f

Metzler, Schniter, Veeraraghavan, Baraniuk (2018). *ICML*

Deep learning regularization

- Step 1: Train a neural network $\boldsymbol{f}$ for denoising
- Step 2: Regularization by denoising (RED) $\mathcal{R}(\boldsymbol{x}) = \boldsymbol{x}\big(\boldsymbol{x} - f(\boldsymbol{x})\big)$

Plug in a deep learning denoiser

Denoising
function
f

Metzler, Schniter, Veeraraghavan, Baraniuk (2018). *ICML*

Deep learning regularization

Phase retrieval reconstruction
from noisy oversampled Fourier measurements

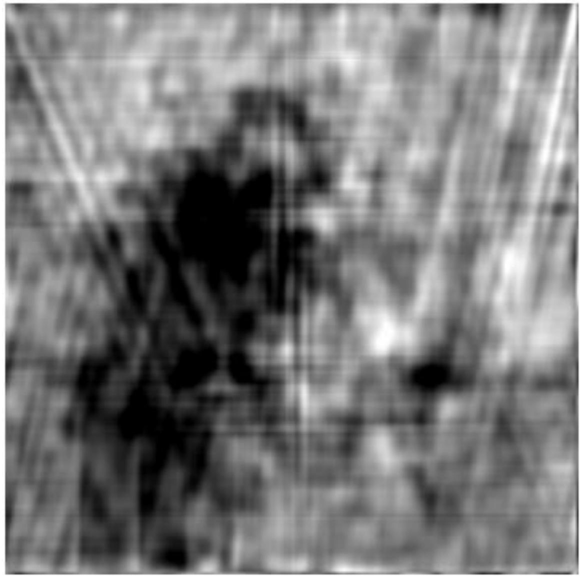

(b) WF (63 sec)

Without regularization

(d) SPAR (294 sec)

Classical regularization

(h) prDeep (345 sec)

Deep learning regularization

Metzler, Schniter, Veeraraghavan, Baraniuk (2018). *ICML*

Deep learning regularization

Favor realistic images

Restrict the search space

Regularization by denoising
Metzler, Schniter, Veeraraghavan, Baraniuk (2018). *ICML*

Plug-and-play priors
Chang, Bian, Zhang (2021). *eLight*

Generative models
Hand, Leong, Voroninski (2018). *NeurIPS*

Deep Image Prior
Wang *et al* (2021). *Light: Science & Applications*

Generative adversarial networks

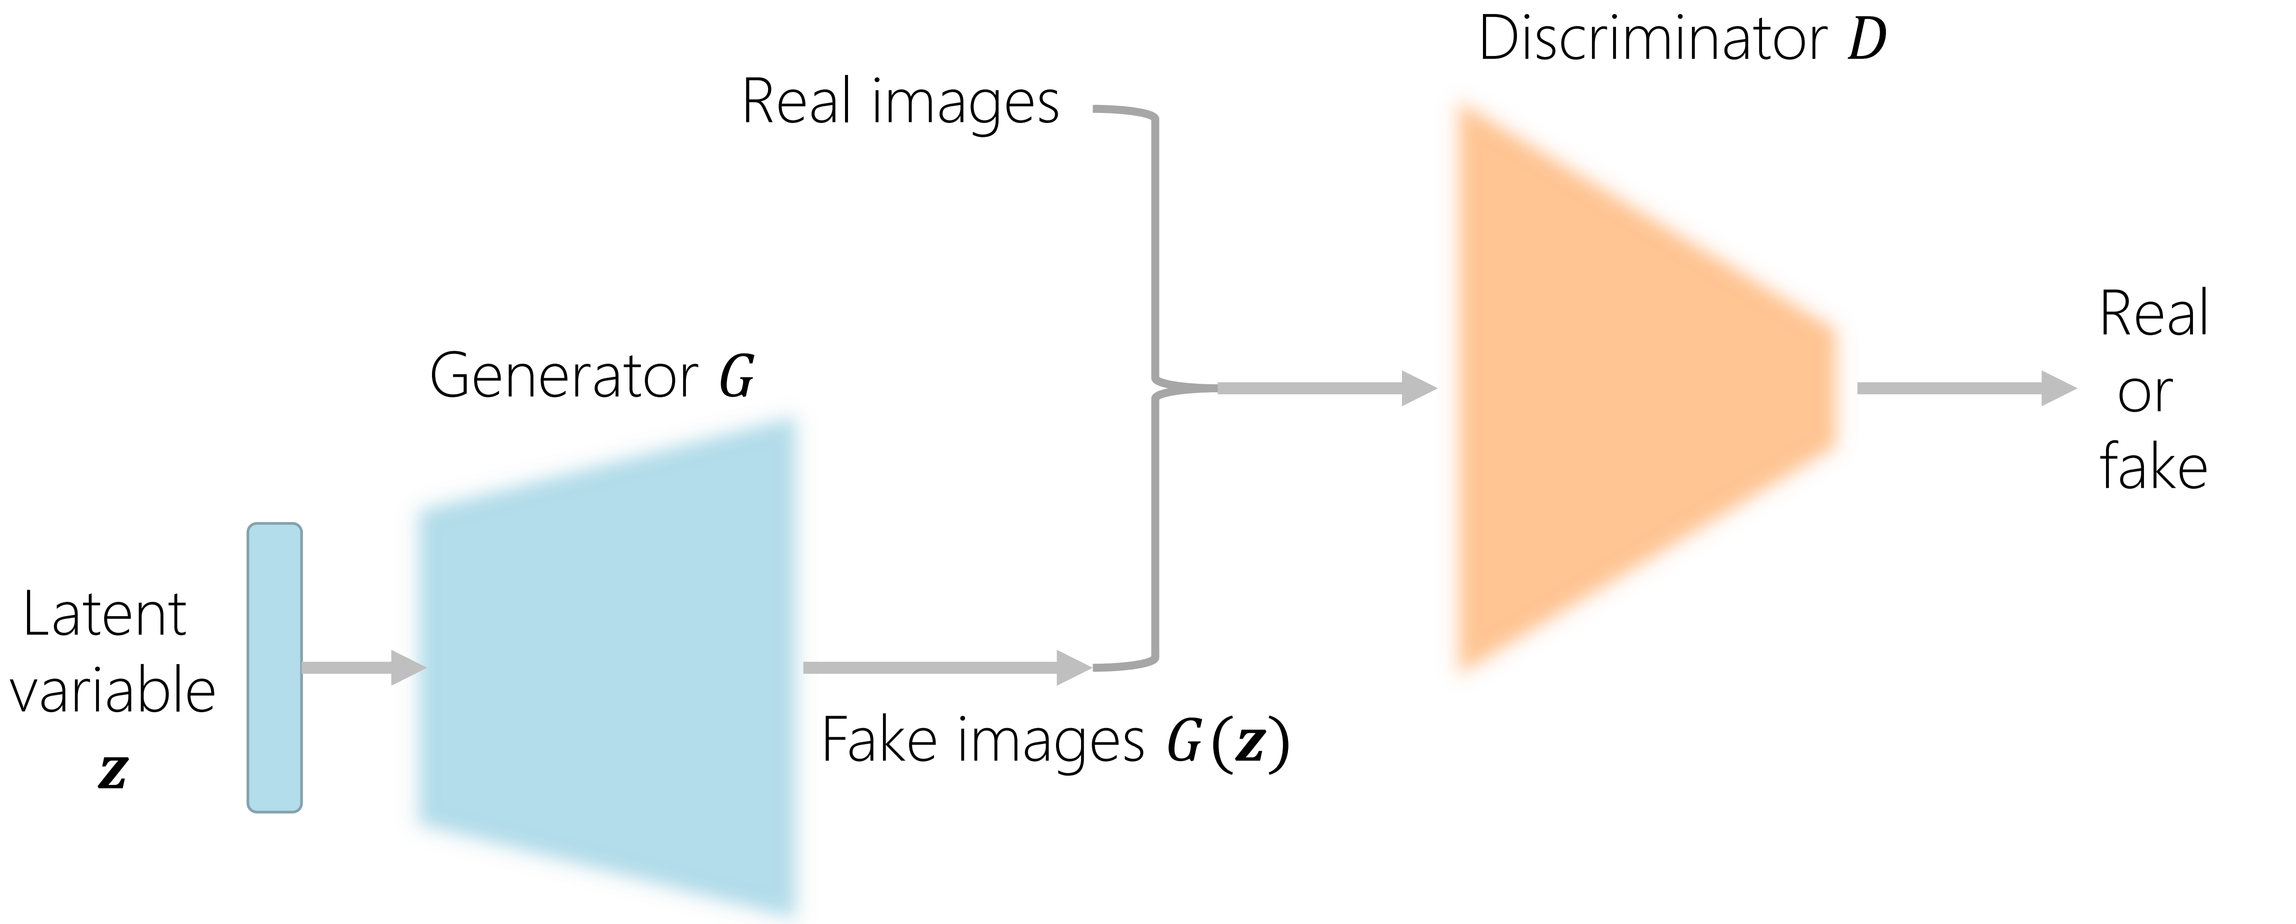

Goodfellow *et al* (2014). *NeurIPS*

Generative adversarial networks

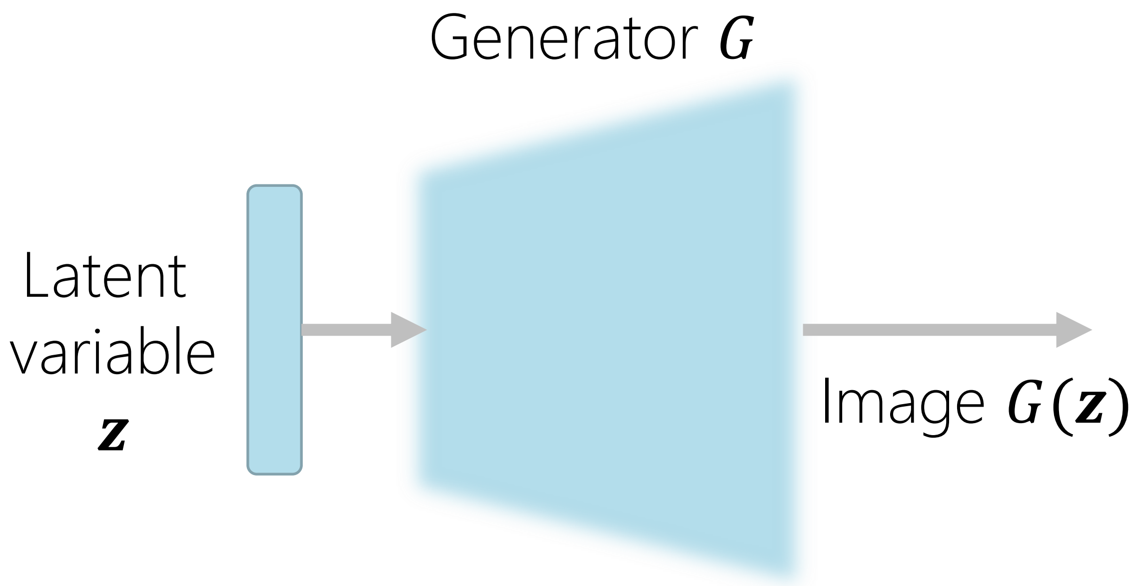

- The generator has learned the *distribution of images*
- Restrict the search space to the generator output

⇒ Gradient descent on $\boldsymbol{z}$ directly

Hand, Leong, Voroninski (2018). *NeurIPS*

Generative adversarial networks

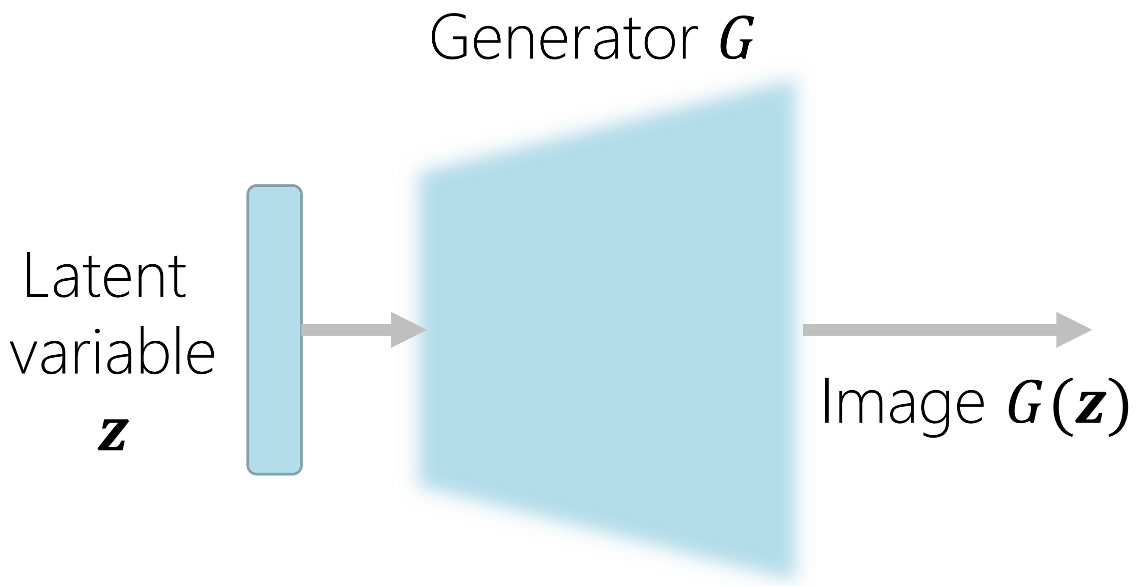

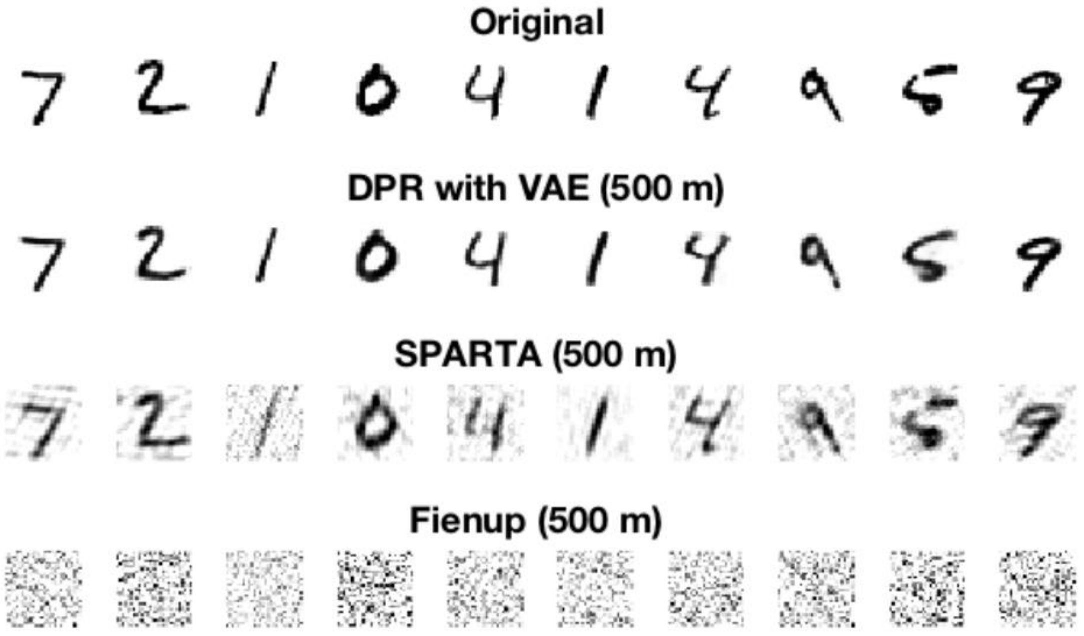

Hand, Leong, Voroninski (2018). *NeurIPS*

Limits of deep learning

Original phantom

Deep learning reconstruction

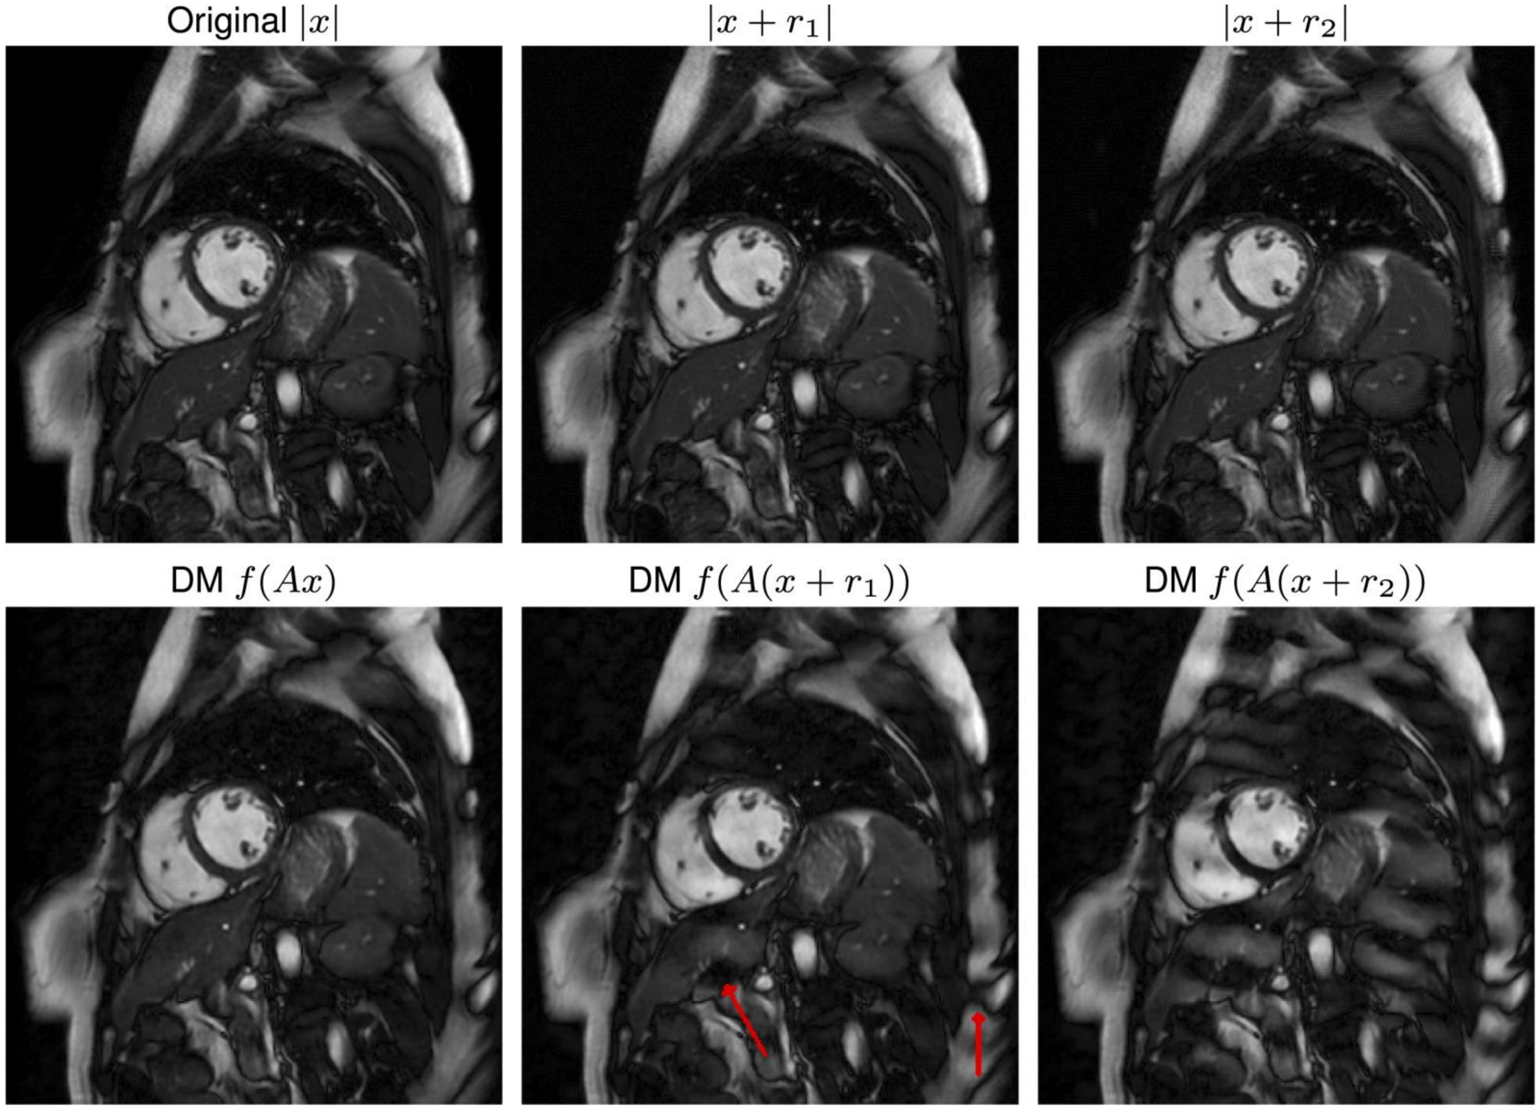

Antun *et al* (2020). PNAS

Limits of deep learning

Unstable, sensitive to perturbations

Erase outliers (e.g., tumors)

Sensitive to acquisition parameters (noise, sampling)

Often returns a realistic image

Bayesian GAN

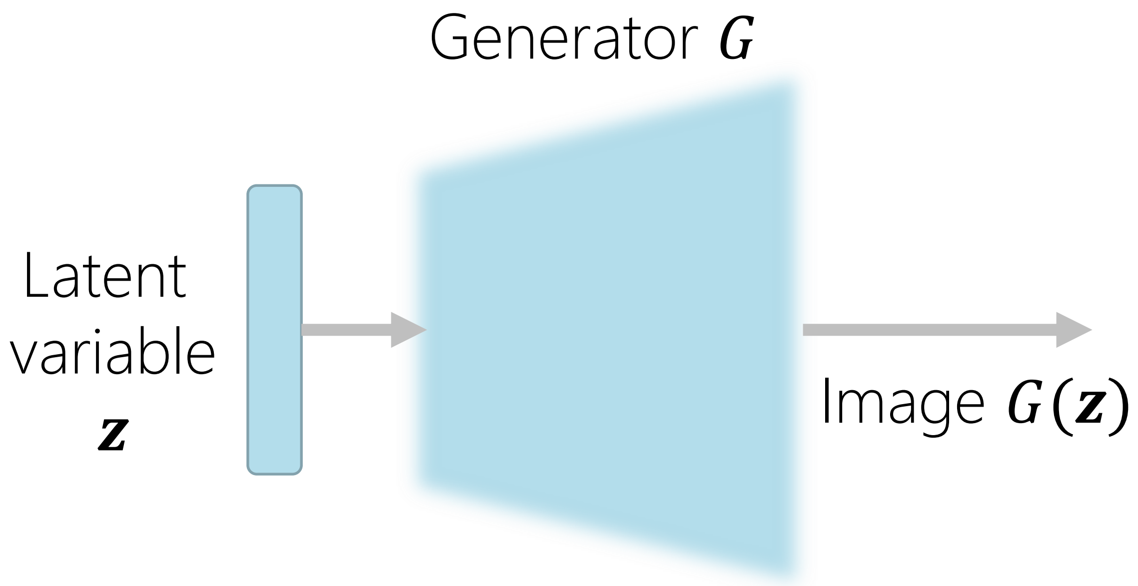

- The generator has learned the *distribution of images*
- Restrict the search space to the generator output

$\Rightarrow$ Gradient descent on $\boldsymbol{z}$ directly (point estimate)

Bayesian GAN

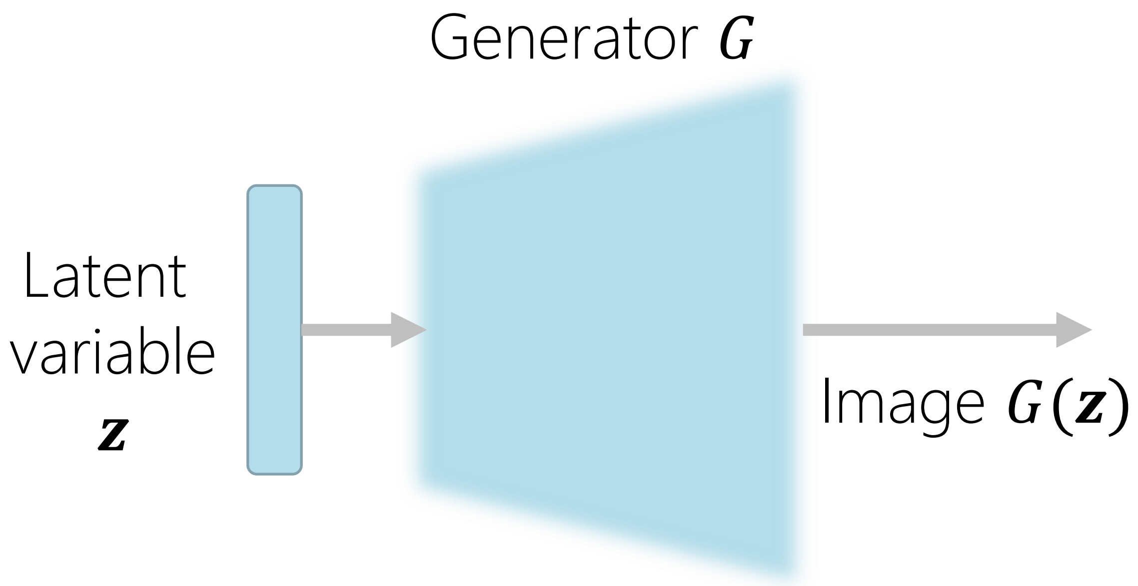

- The generator has learned the *distribution of images*
- Restrict the search space to the generator output

⇒ ~~Gradient descent on $\boldsymbol{z}$ directly (point estimate)~~

Sampling from the posterior distribution

Markov-Chain Monte-Carlo on $\boldsymbol{z}$

Bohra, Pham, Dong, Unser (2022). *Submitted*

Bayesian GAN

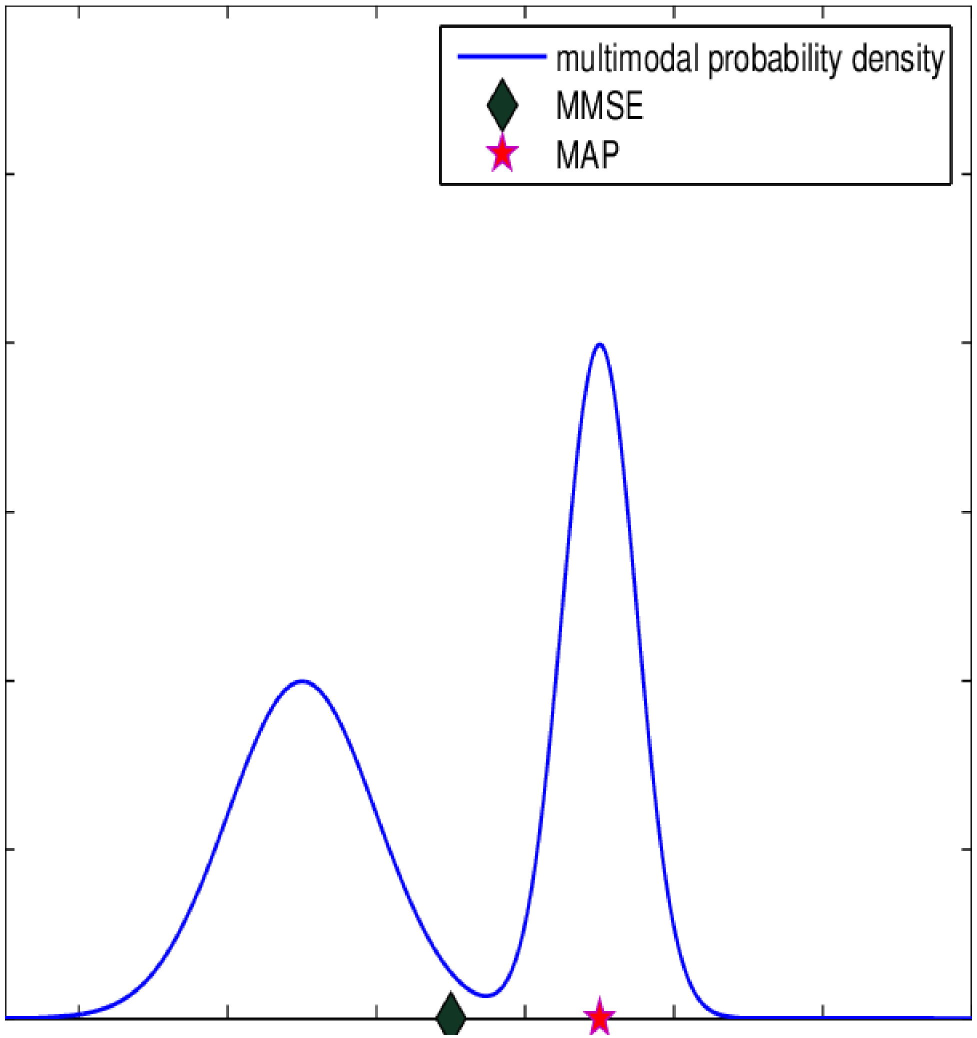

- The generator has learned the *distribution of images*
- Restrict the search space to the generator output

⇒ ~~Gradient descent on $\boldsymbol{z}$ directly (point estimate)~~

Sampling from the posterior distribution

Markov-Chain Monte-Carlo on $\boldsymbol{z}$

Bohra, Pham, Dong, Unser (2022). *Submitted*

Bayesian GAN results

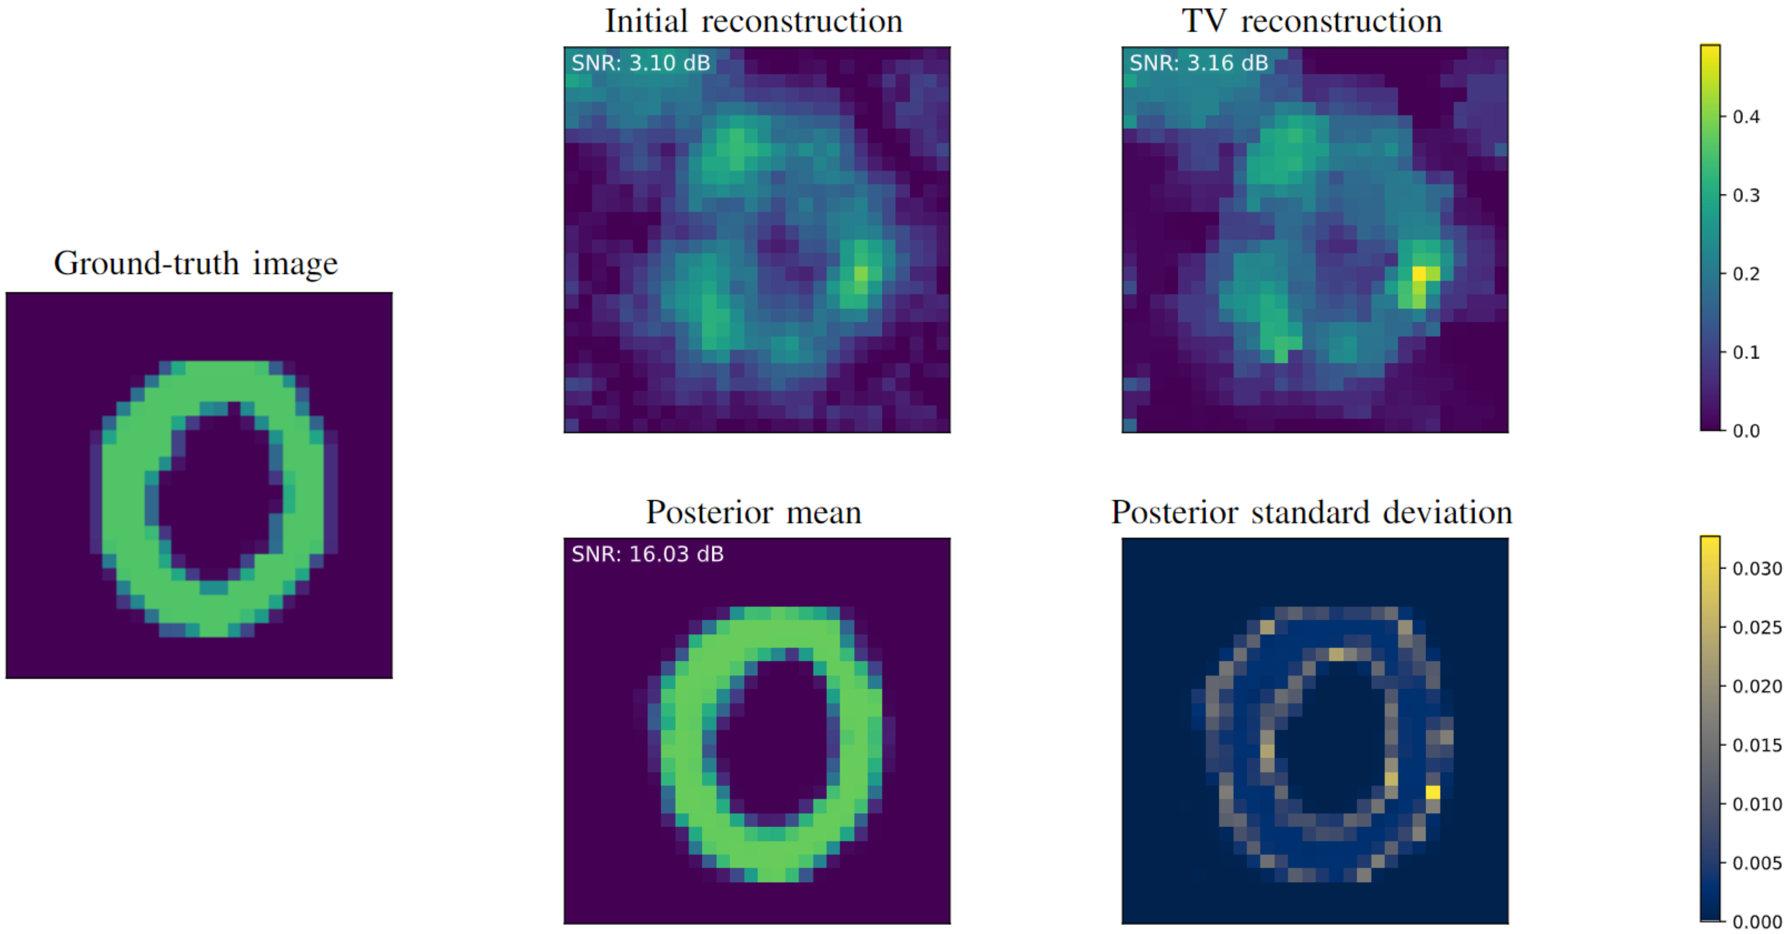

Bohra, Pham, Dong, Unser (2022). *Submitted*

Bayesian GAN results

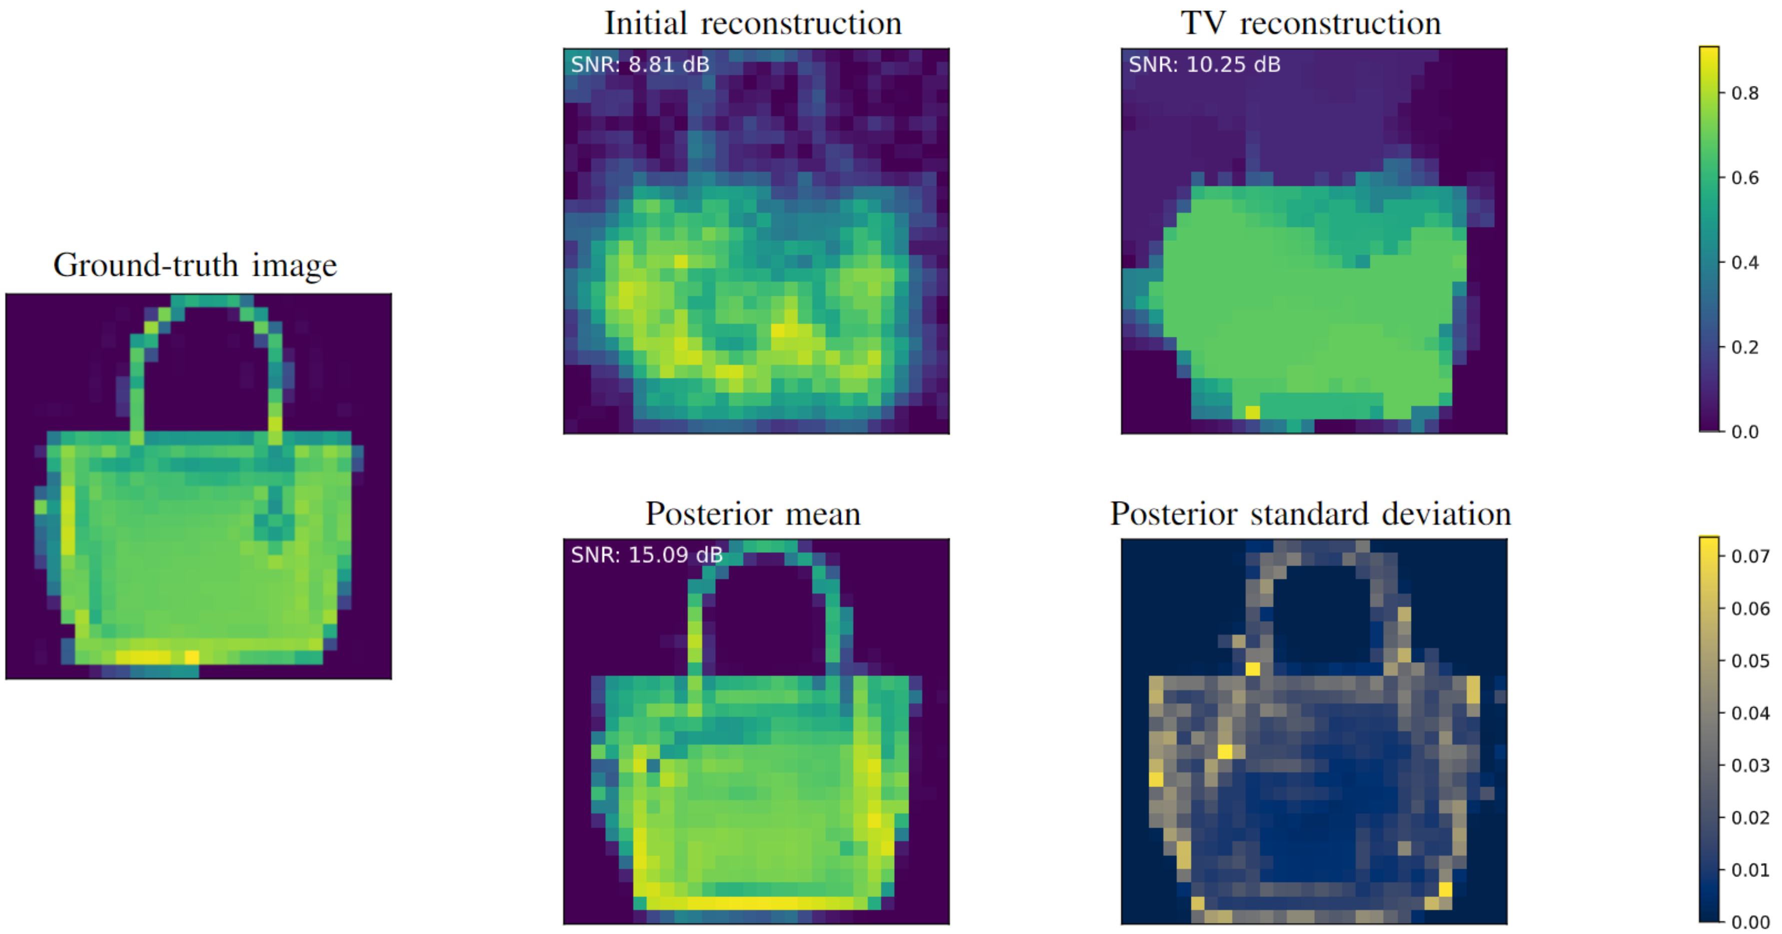

Bohra, Pham, Dong, Unser (2022). *Submitted*

Content

Find $\boldsymbol{x}^*$ in
$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

$$\begin{bmatrix} y_1 \\ \\ \vdots \\ \\ y_n \end{bmatrix} = \left| \begin{bmatrix} a_{11} & & \cdots & a_{1d} \\ & & \cdots & \\ & & \vdots & \\ & & \cdots & \\ a_{n1} & & \cdots & a_{nd} \end{bmatrix} \begin{bmatrix} x_1^* \\ \\ \vdots \\ x_d^* \end{bmatrix} \right|^2$$

- General inverse problem framework
- Phase retrieval (PR) applications
- PR algorithms
- PR theory: random model
- Machine learning

Conclusion

$$\boldsymbol{y} = |\boldsymbol{A}\boldsymbol{x}^*|^2$$

Rich history of phase retrieval

- From Fourier phase retrieval
- To computational imaging

Advancing fast

- Many algorithmic improvements
- Powerful algorithms for random setting

An interdisciplinary topic

- Many applications (astronomy, biomedical, etc.)
- Deep learning regularization

Thanks for your attention