# Phase retrieval: from computational imaging to machine learning

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

October 2025

PHYS-715, EPFL

# In a previous episode

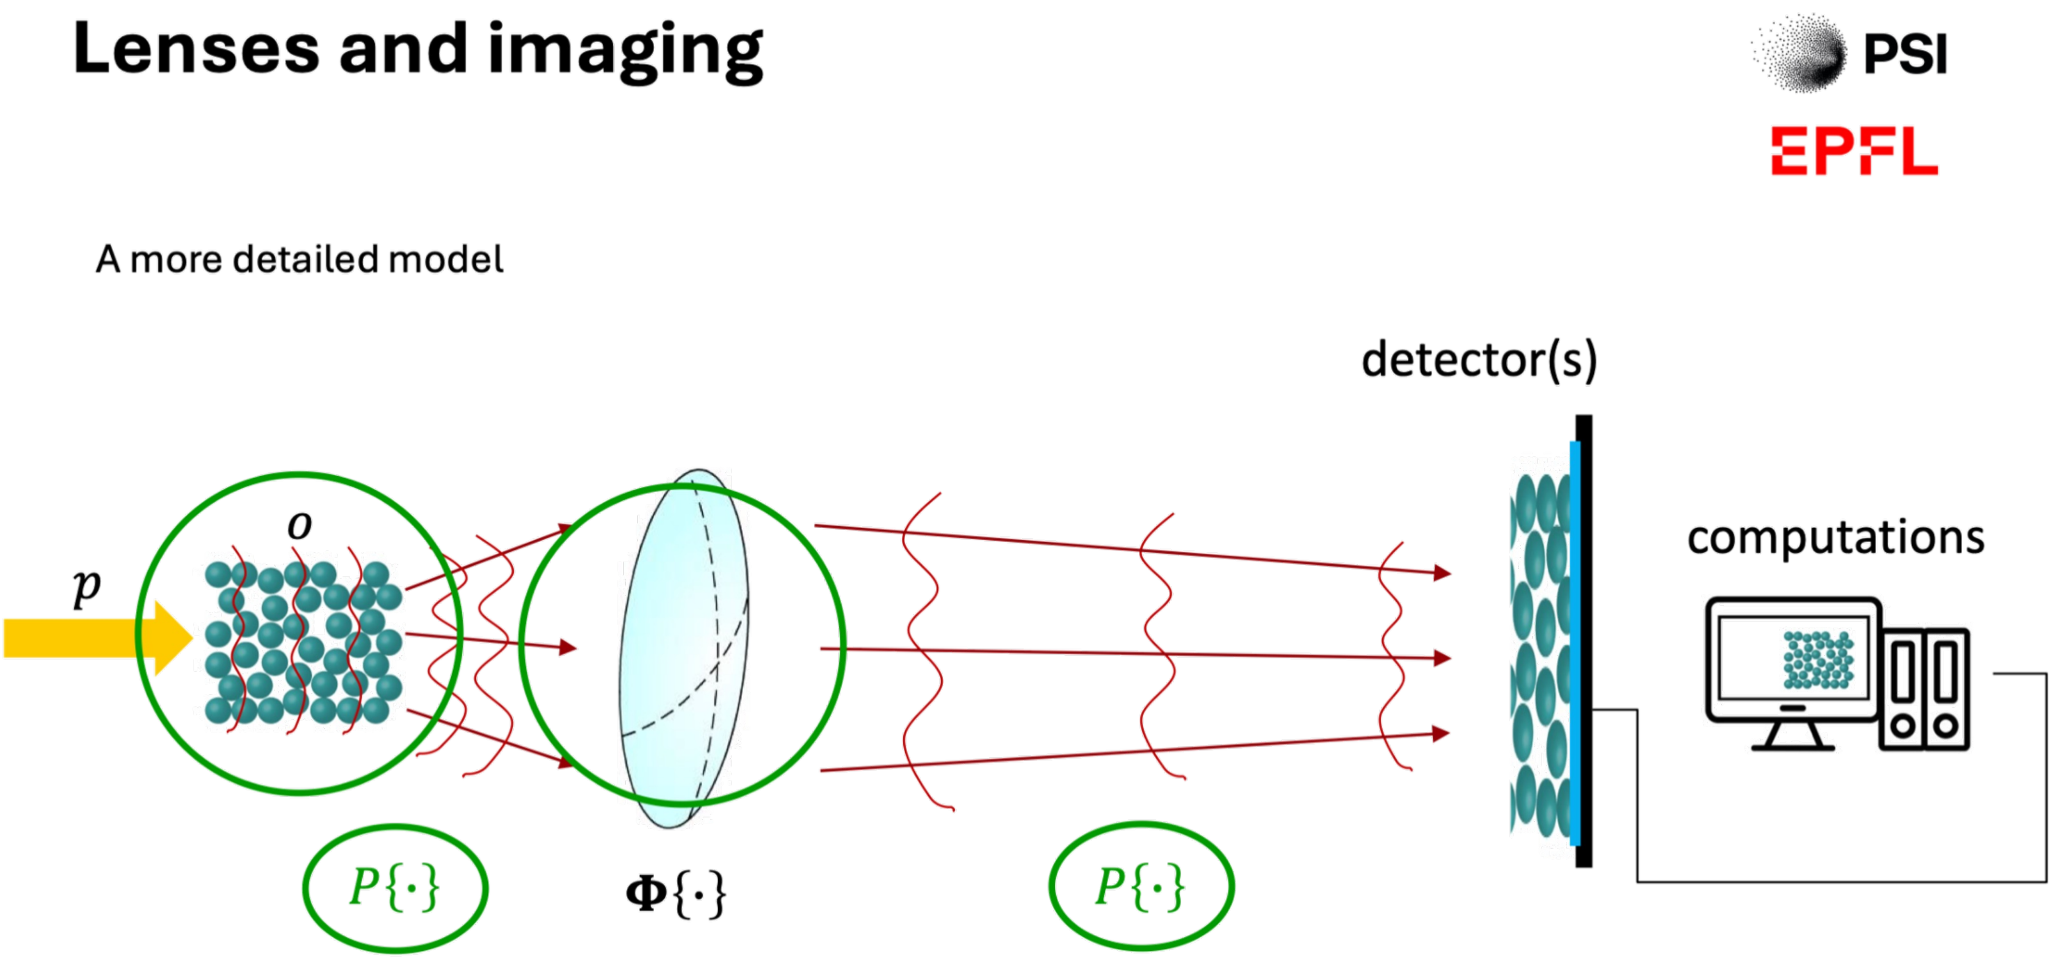

# In a previous episode

How to measure a complex-valued electric field?

What reconstruction algorithm to use?

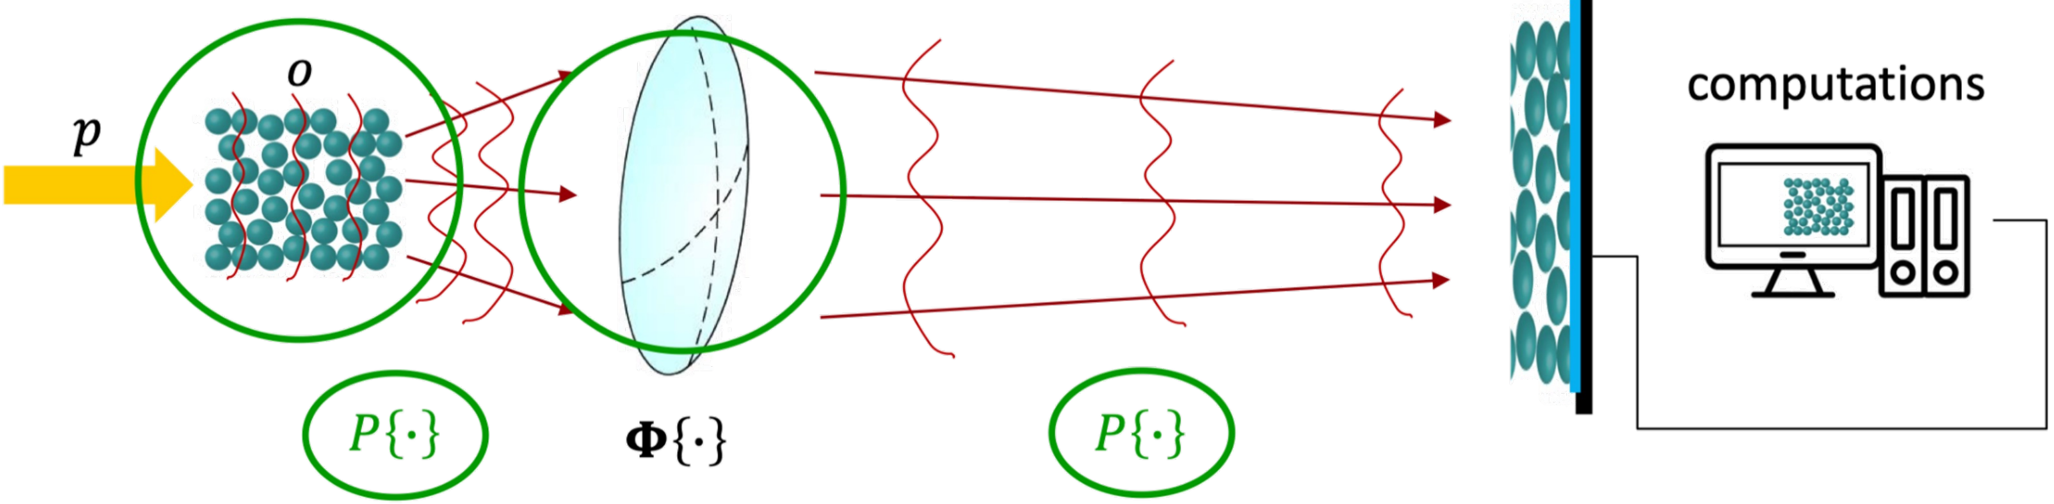

# In a previous episode

How to measure a complex-valued electric field?

What reconstruction algorithm to use?

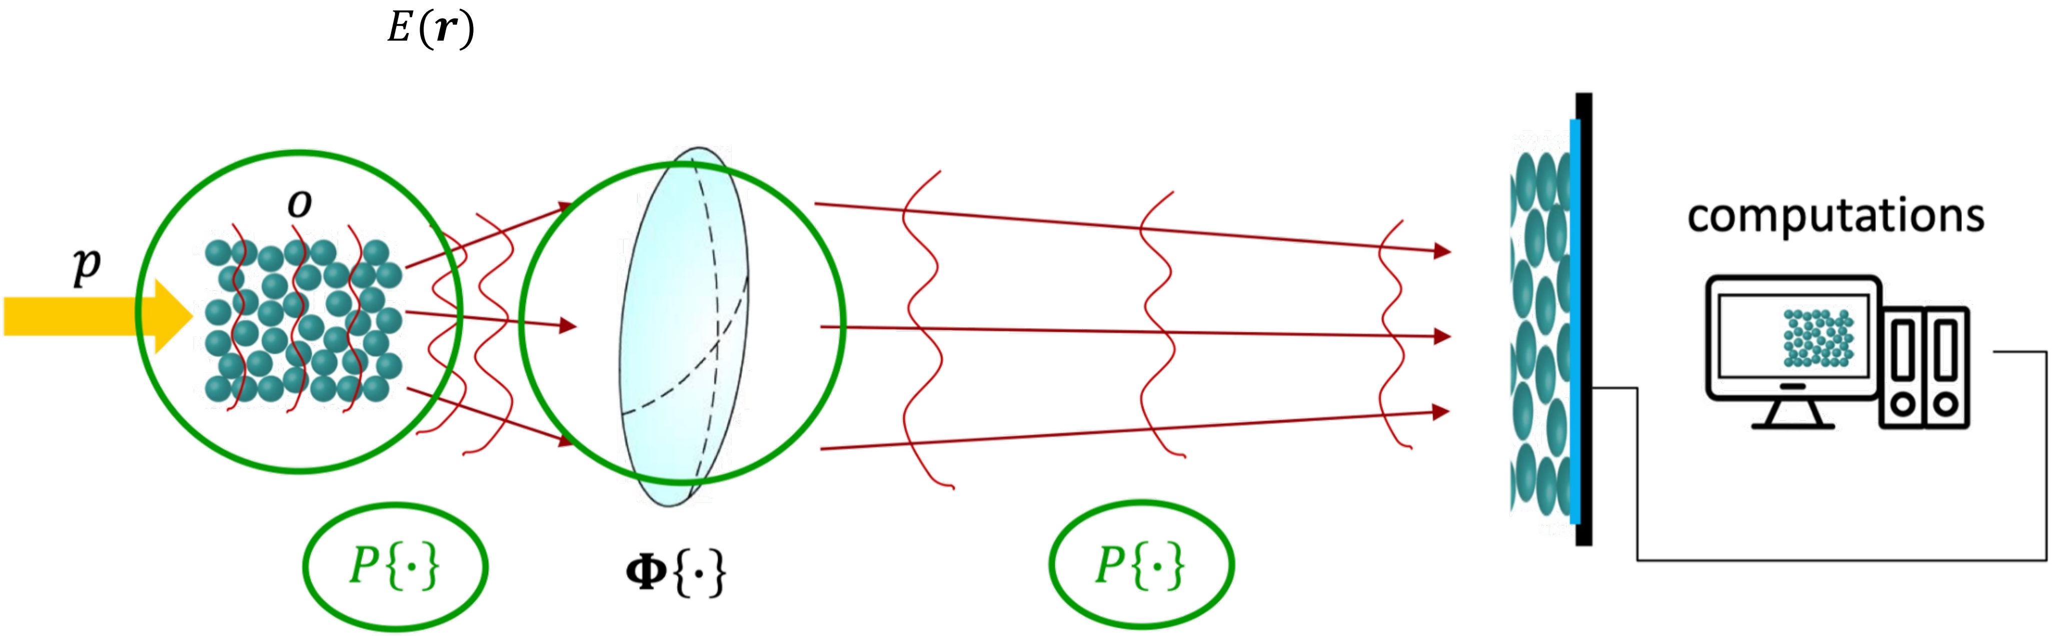

# In a previous episode

Problem: Cameras measure intensities $|E(\boldsymbol{r})|^2$

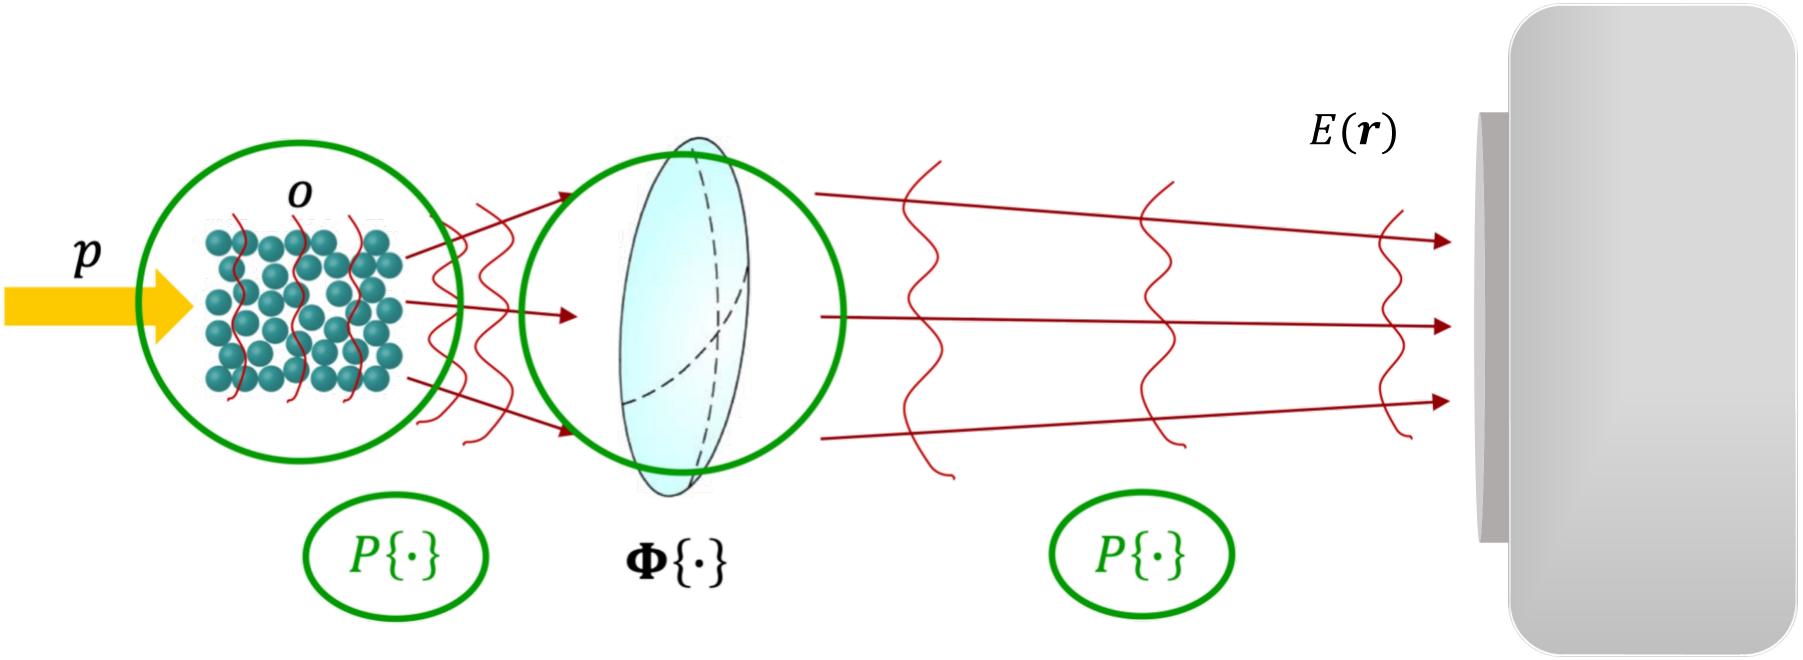

# Phase microscopy

Amplitude

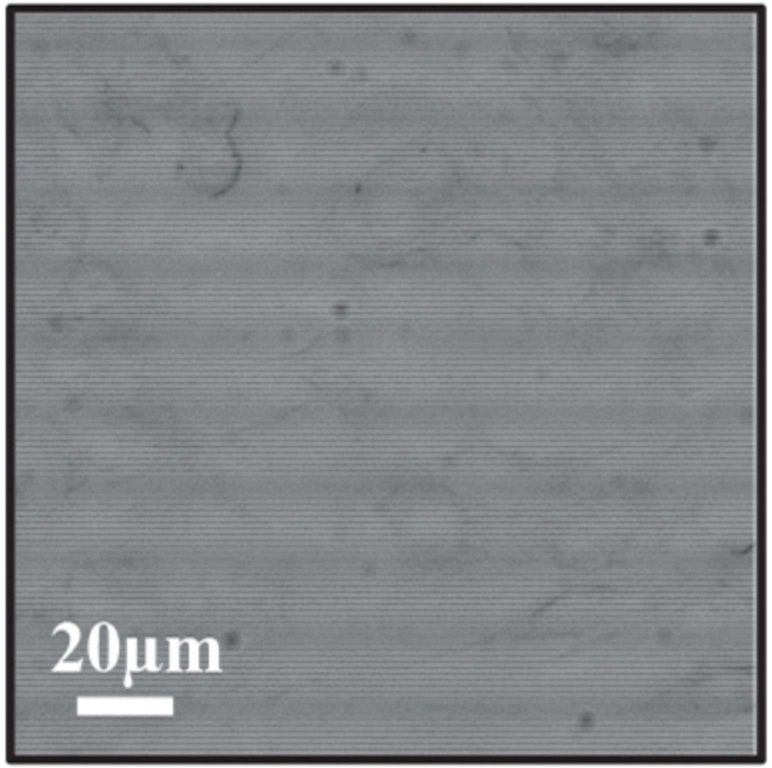


Phase

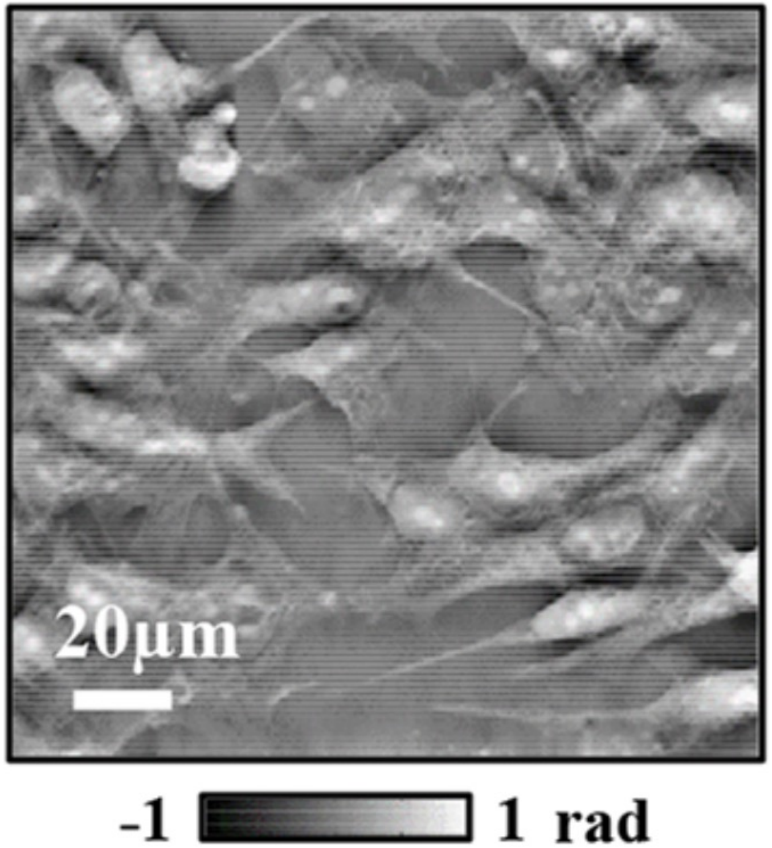


Tian, Waller (2015). *Optics Express*

# 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

# Content

Find $\boldsymbol{x}^*$ in

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

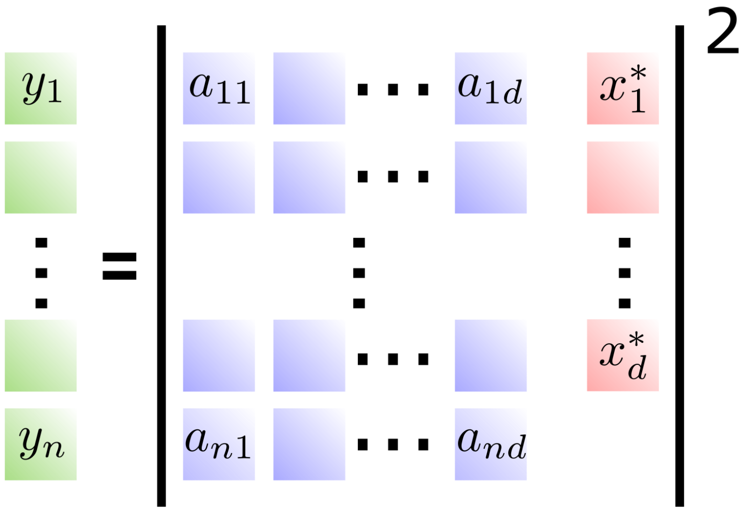


- Physics: Phase imaging devices
- Math:
  - Inverse problems
  - Phase retrieval algorithms
  - Phase retrieval theory
- Machine learning: Regularization

# A basic microscope

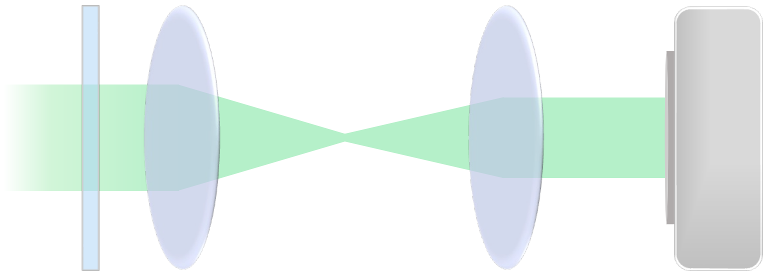

- 4f system:
  - Sample
  - Objective lens
  - Tube lens
  - Camera
- 2 Fourier transforms
- Magnification:
  ratio of the focal lengths

# Digital holographic microscopy

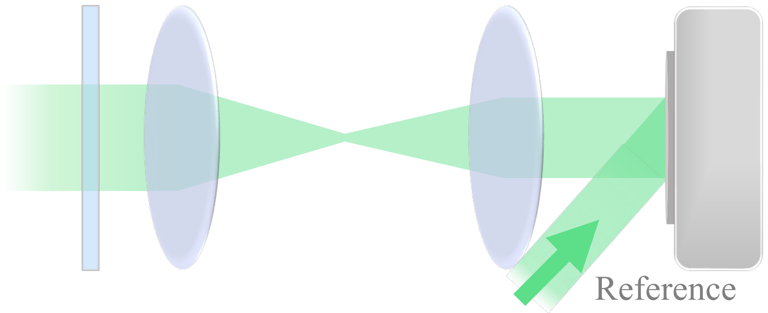


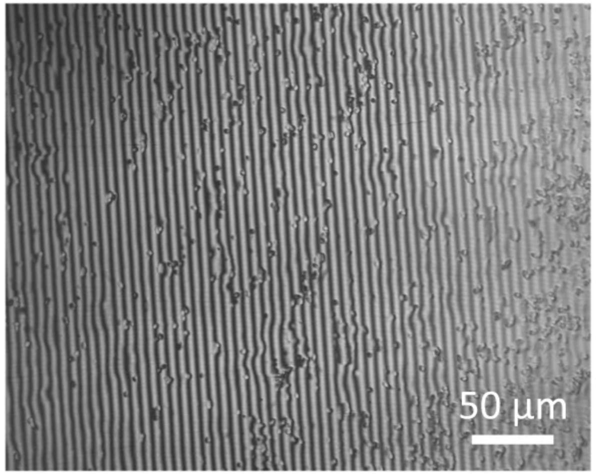


Source: Abbasian, et al, Scientific Data 2024

- Introduction of a reference

$$I = \left|E + E_{ref}\right|^2$$
$$\approx |\mathrm{E}_{\mathrm{ref}}|^2 + 2\left|E_{ref}\right||E| \cos \phi$$

- Off-axis reference:
  - Typical fringe patterns
  - Reference phase varies between 0 and $2\pi$
- Reconstruction: Fourier filtering

# Digital holographic microscopy

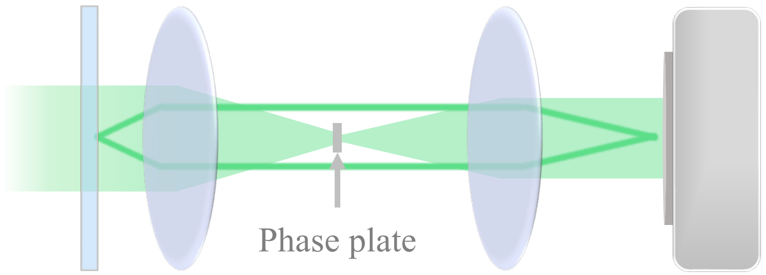


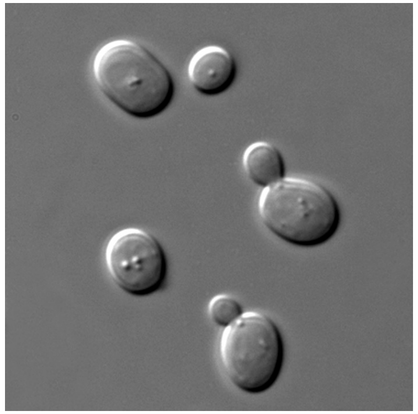

Source: Wikipedia

- On-axis variation
  (Zernike phase microscope)
  but only for weak-phase objects

- Limitations:
  - Expensive
  - Complex to align
  - Coherent illumination only (laser)

# Wavefront sensors

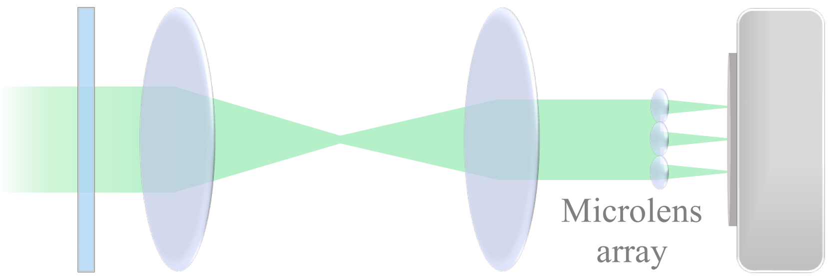


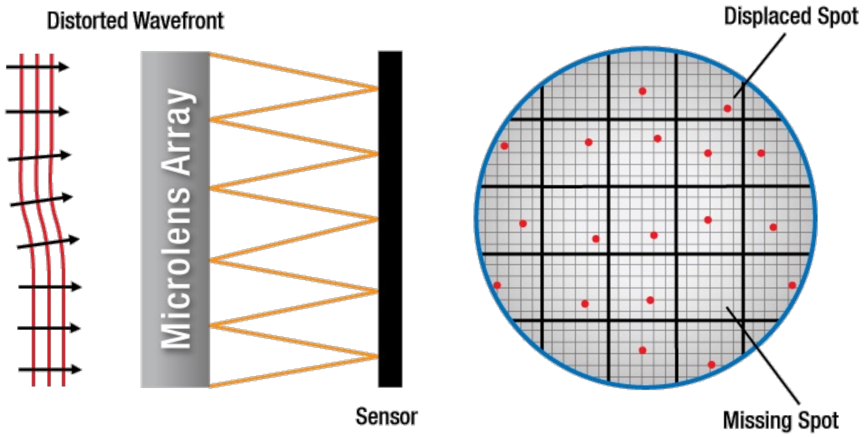


Source: Thorlabs

- Shack-Hartmann wavefront sensor
- Based on a microlens array
- Gradients of the phase => small displacements of focal spots

- Limitations:
  - Small phase shifts only
  - Limited resolution

# Wavefront sensors

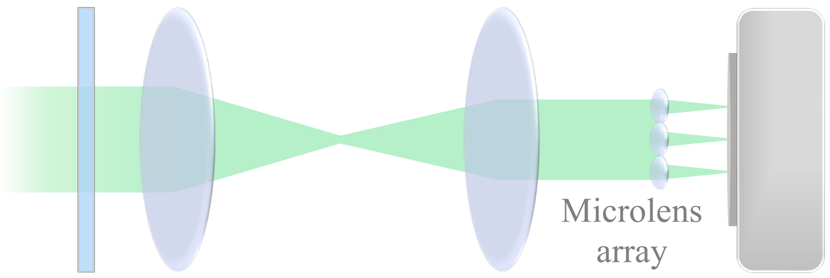


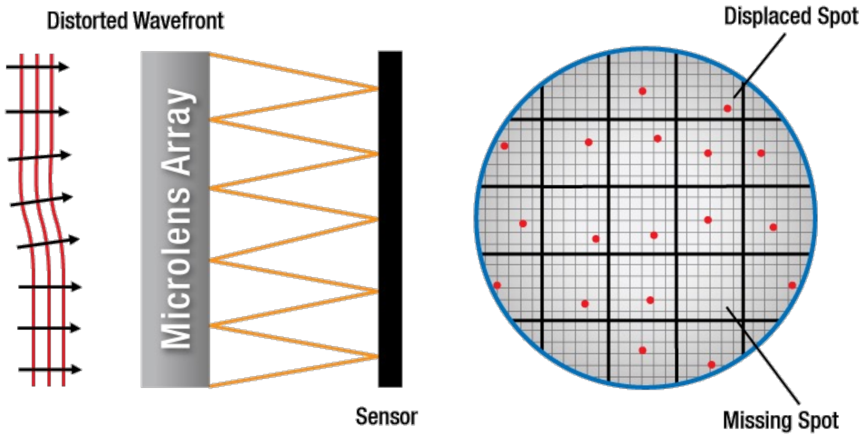


Source: Thorlabs

- Shack-Hartmann wavefront sensor
- Based on a microlens array
- Gradients of the phase => small displacements of focal spots

- Limitations:
  - Small phase shifts only
  - Limited resolution

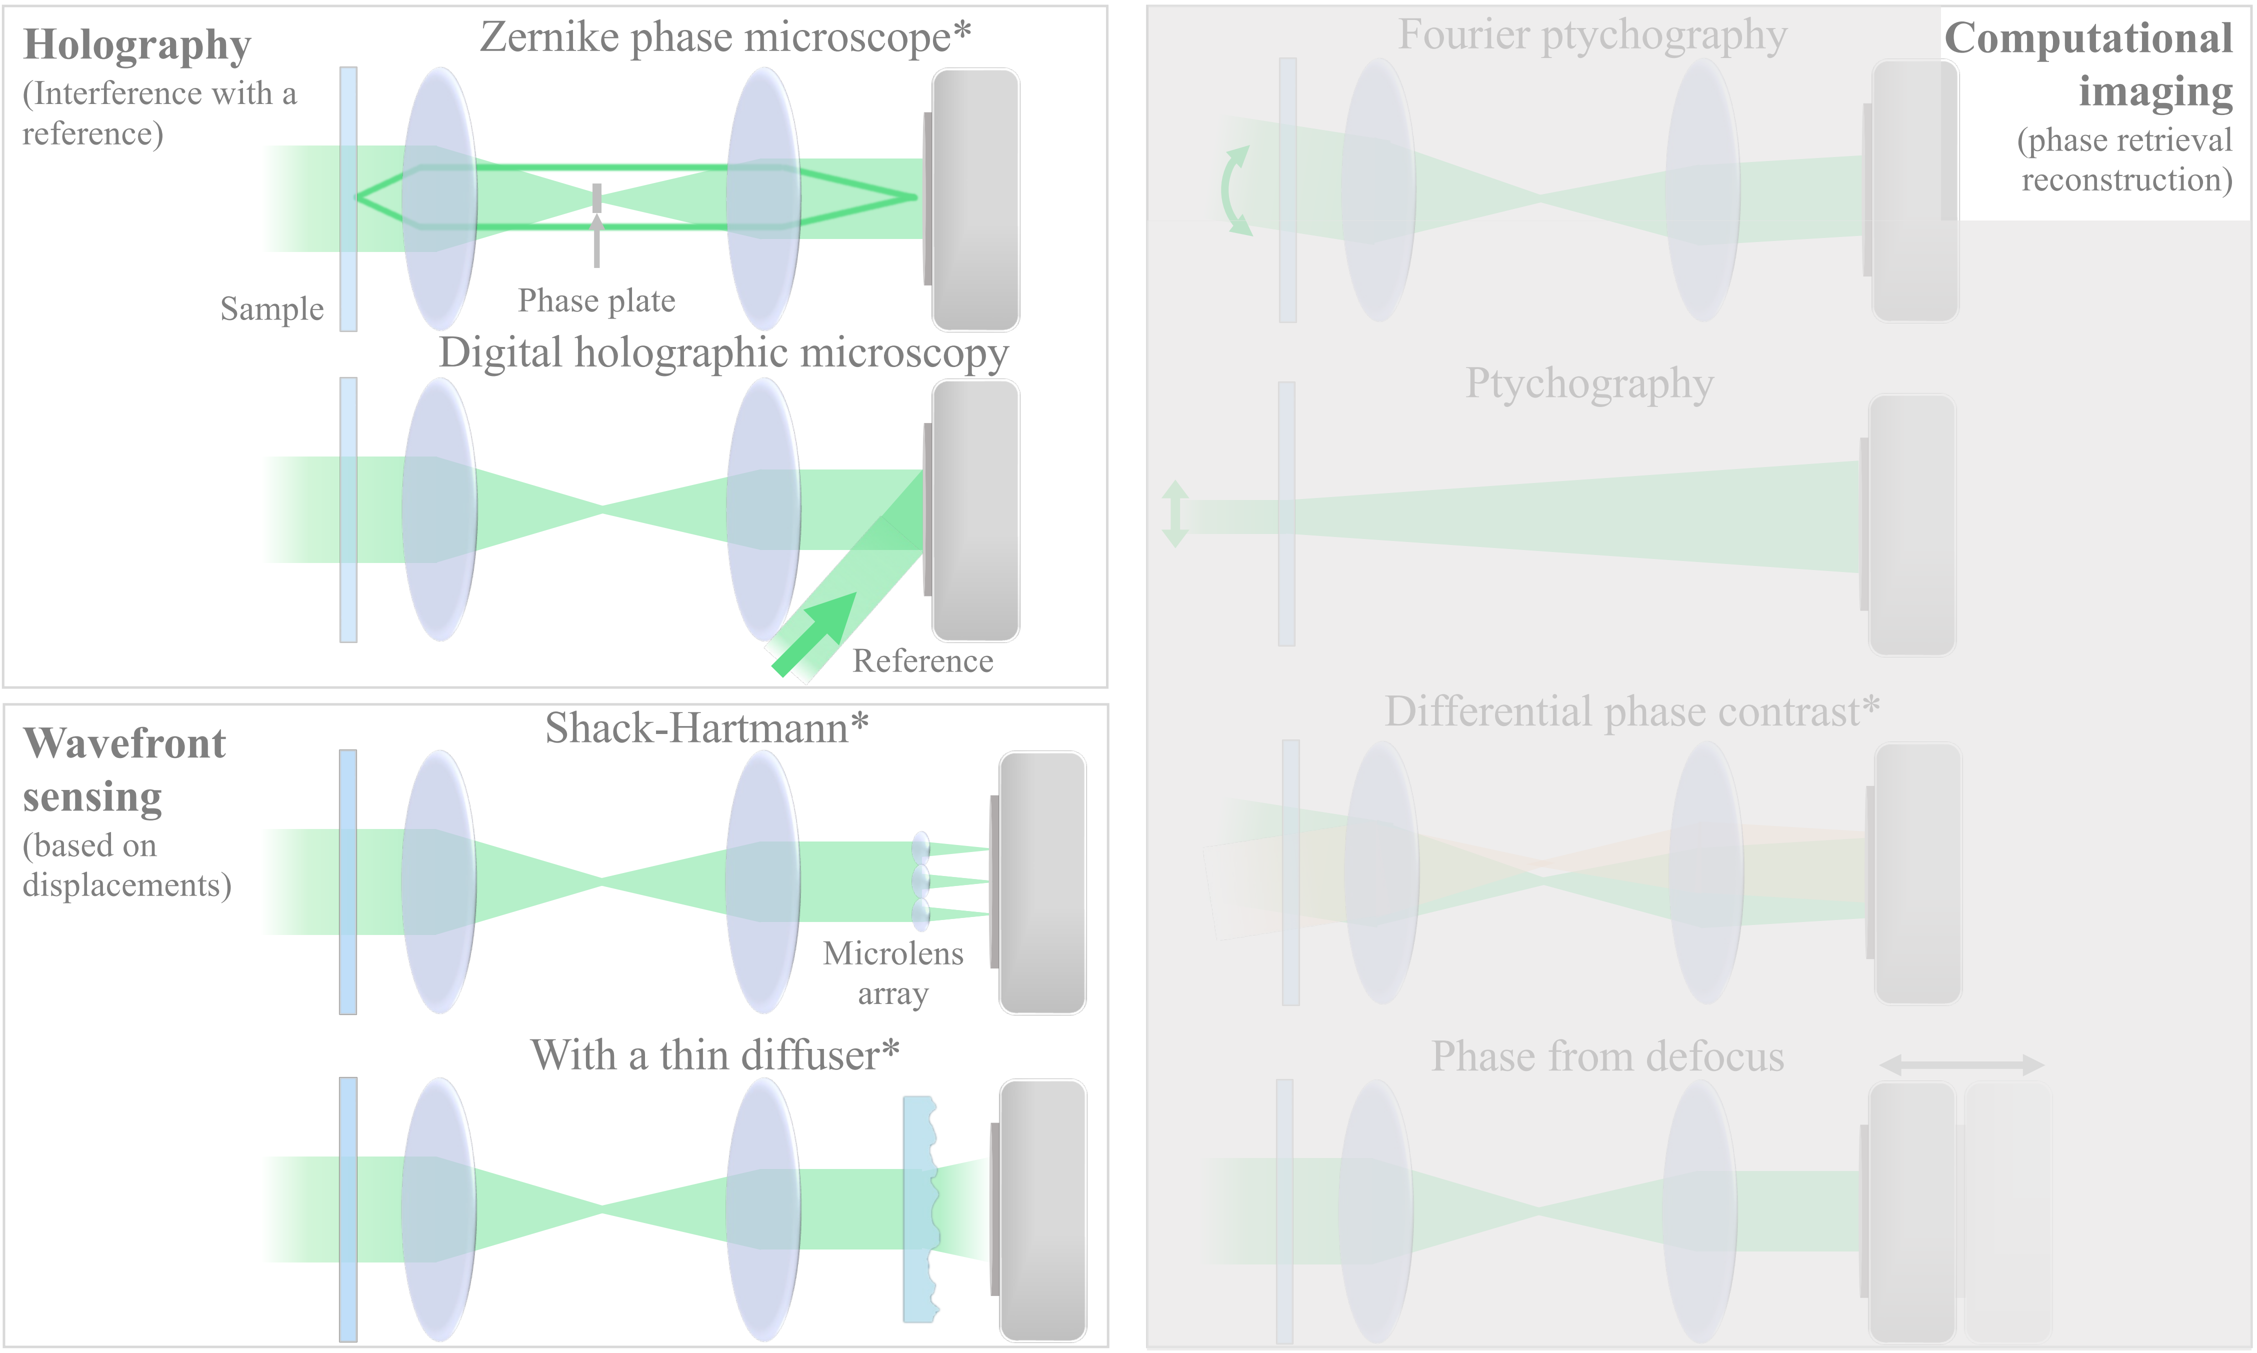

Holography
(Interference with a reference)
Zernike phase microscope*
Sample
Phase plate
Digital holographic microscopy
Reference
Wavefront sensing
(based on displacements)
Shack-Hartmann*
Microlens array
With a thin diffuser*
Computational imaging
(phase retrieval reconstruction)
Fourier ptychography
Ptychography
Differential phase contrast*
Phase from defocus

# Phase retrieval

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

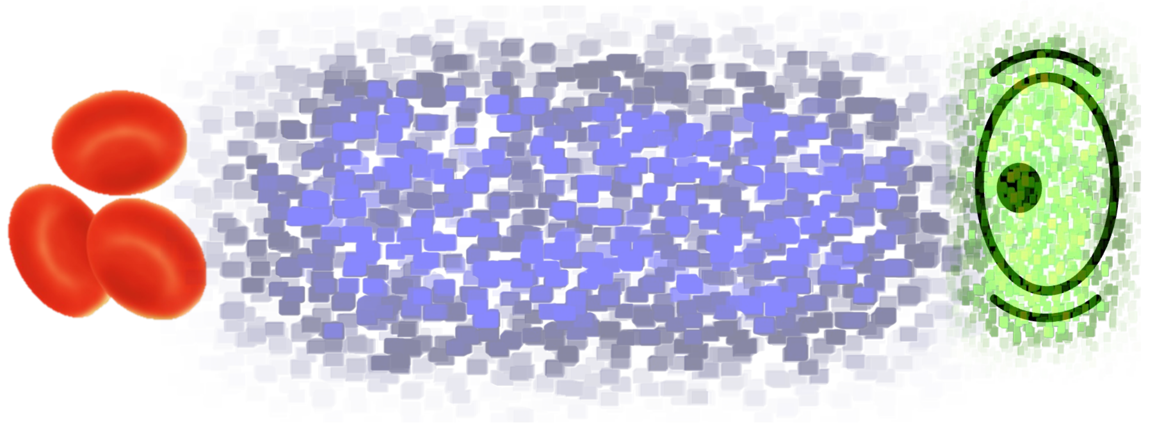

unknown object

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

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

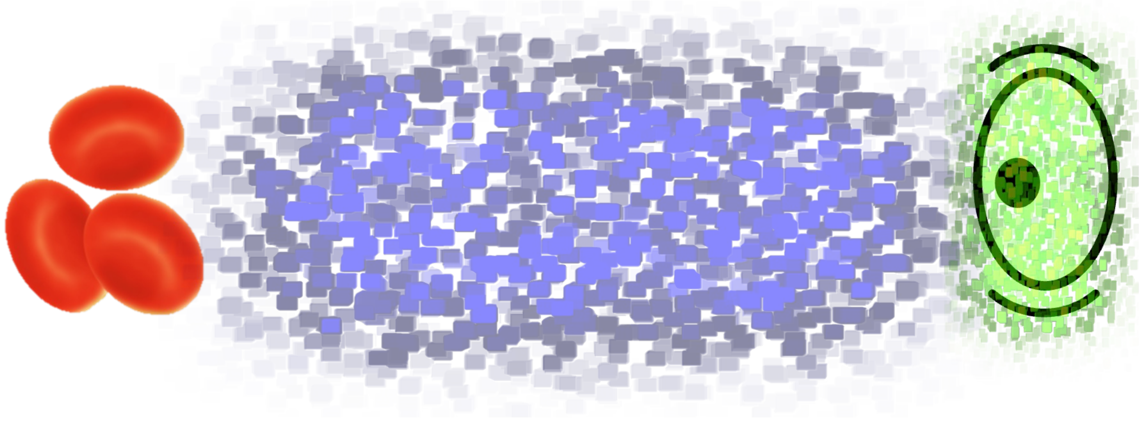

unknown object

$\boldsymbol{x}^*$

$\in \mathbb{C}^d$

imaging system

$\boldsymbol{A}$

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

measurements

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

$\in \mathbb{R}^n$

Amplitude

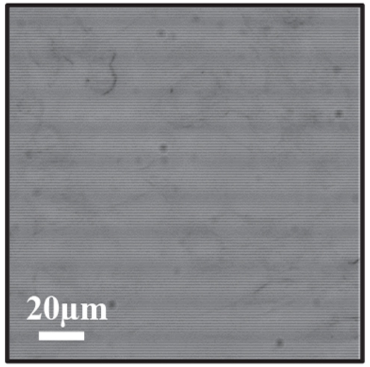


Quantitative Phase Imaging

Phase

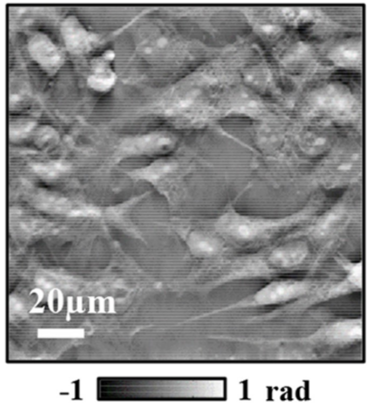


Label-free

Less invasive

No bleaching

Observations over several days

Fast

Videos of samples

Tian, Waller (2015). *Optics Express*

# Phase retrieval applications

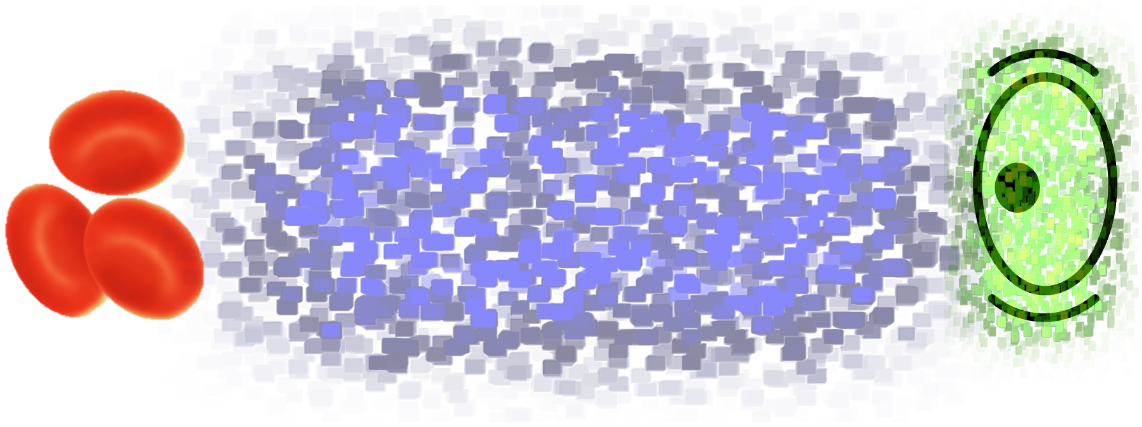

| unknown object | imaging system | measurements |
|---|---|---|
| $\boldsymbol{x}^*$ | $\boldsymbol{A}$ | $\boldsymbol{y} = \lvert \boldsymbol{A}\boldsymbol{x}^* \rvert^2$ |
| $\in \mathbb{C}^d$ | $\in \mathbb{C}^{n \times d}$ | $\in \mathbb{R}^n$ |

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

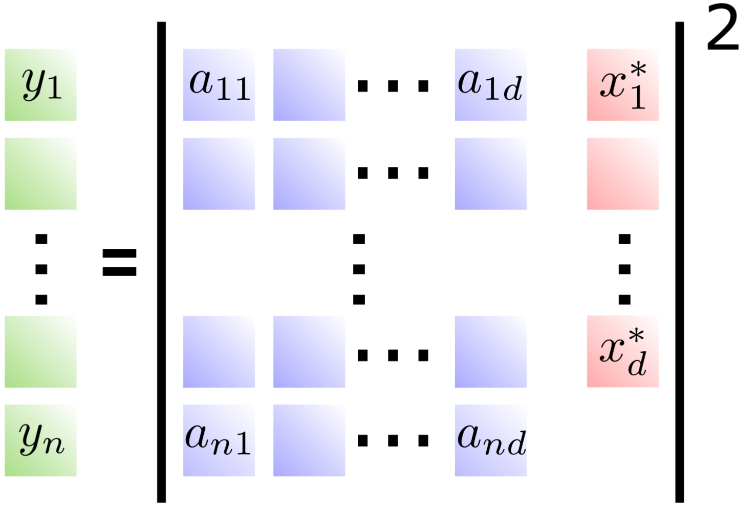


Non-linear equation

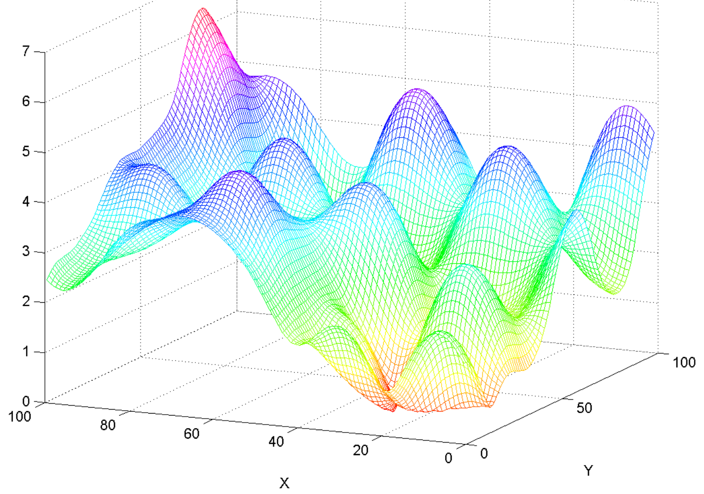

# Phase retrieval and machine learning

Find $\boldsymbol{x}^*$ in

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

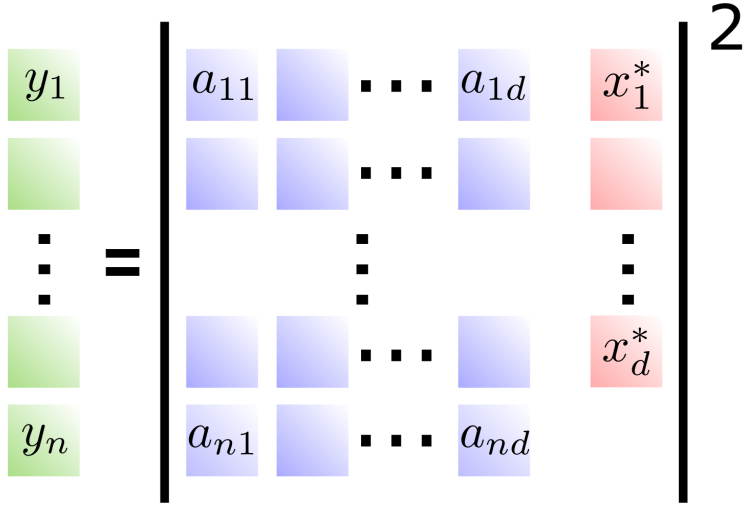


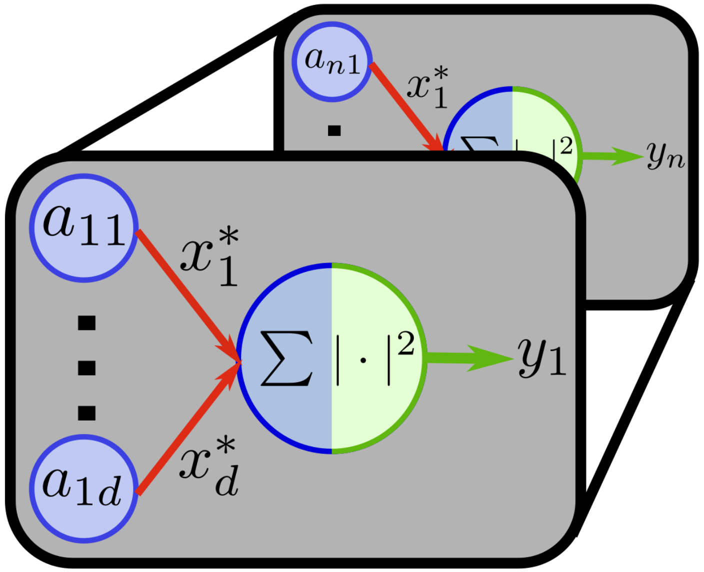

# Machine learning

Single-layer neural network

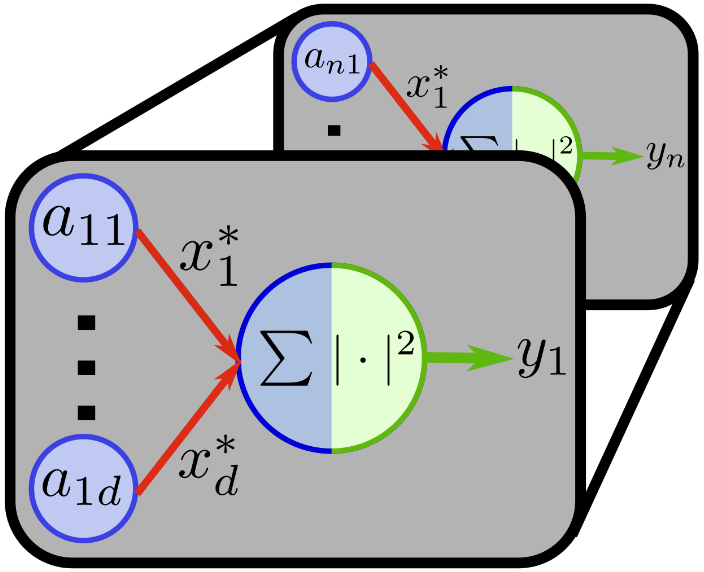


Deep neural network

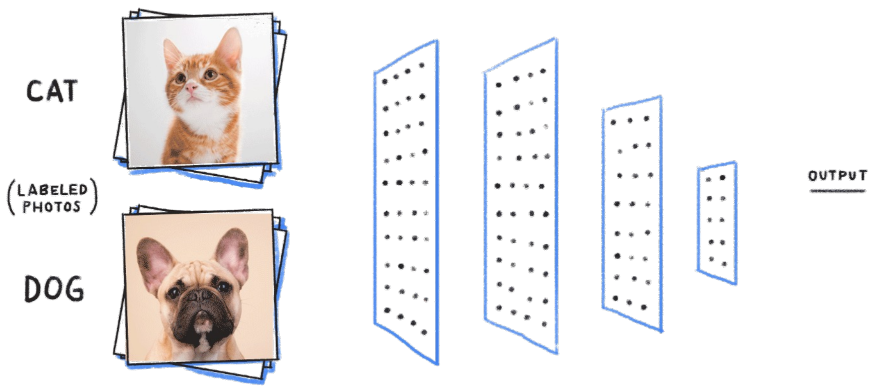

# Machine learning

Single-layer neural network

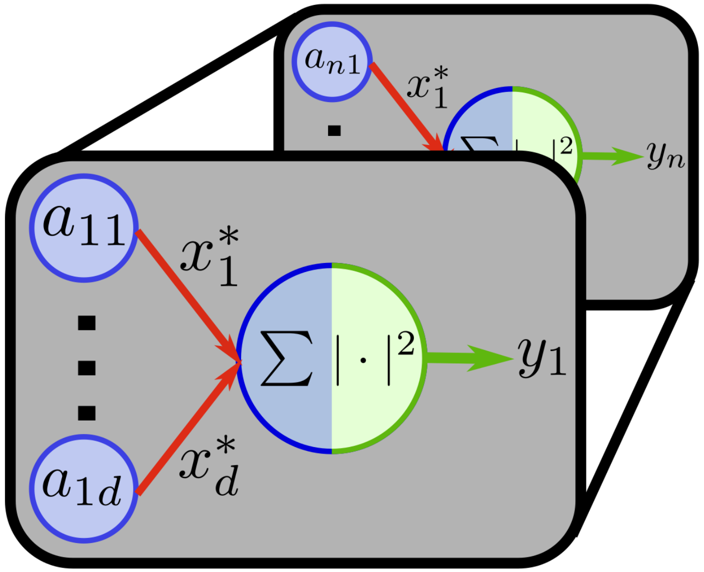


Phase retrieval =
Training a 1-layer neural network

When is it solvable?

What algorithm to use?

Is the solution unique?

# Fourier in optics

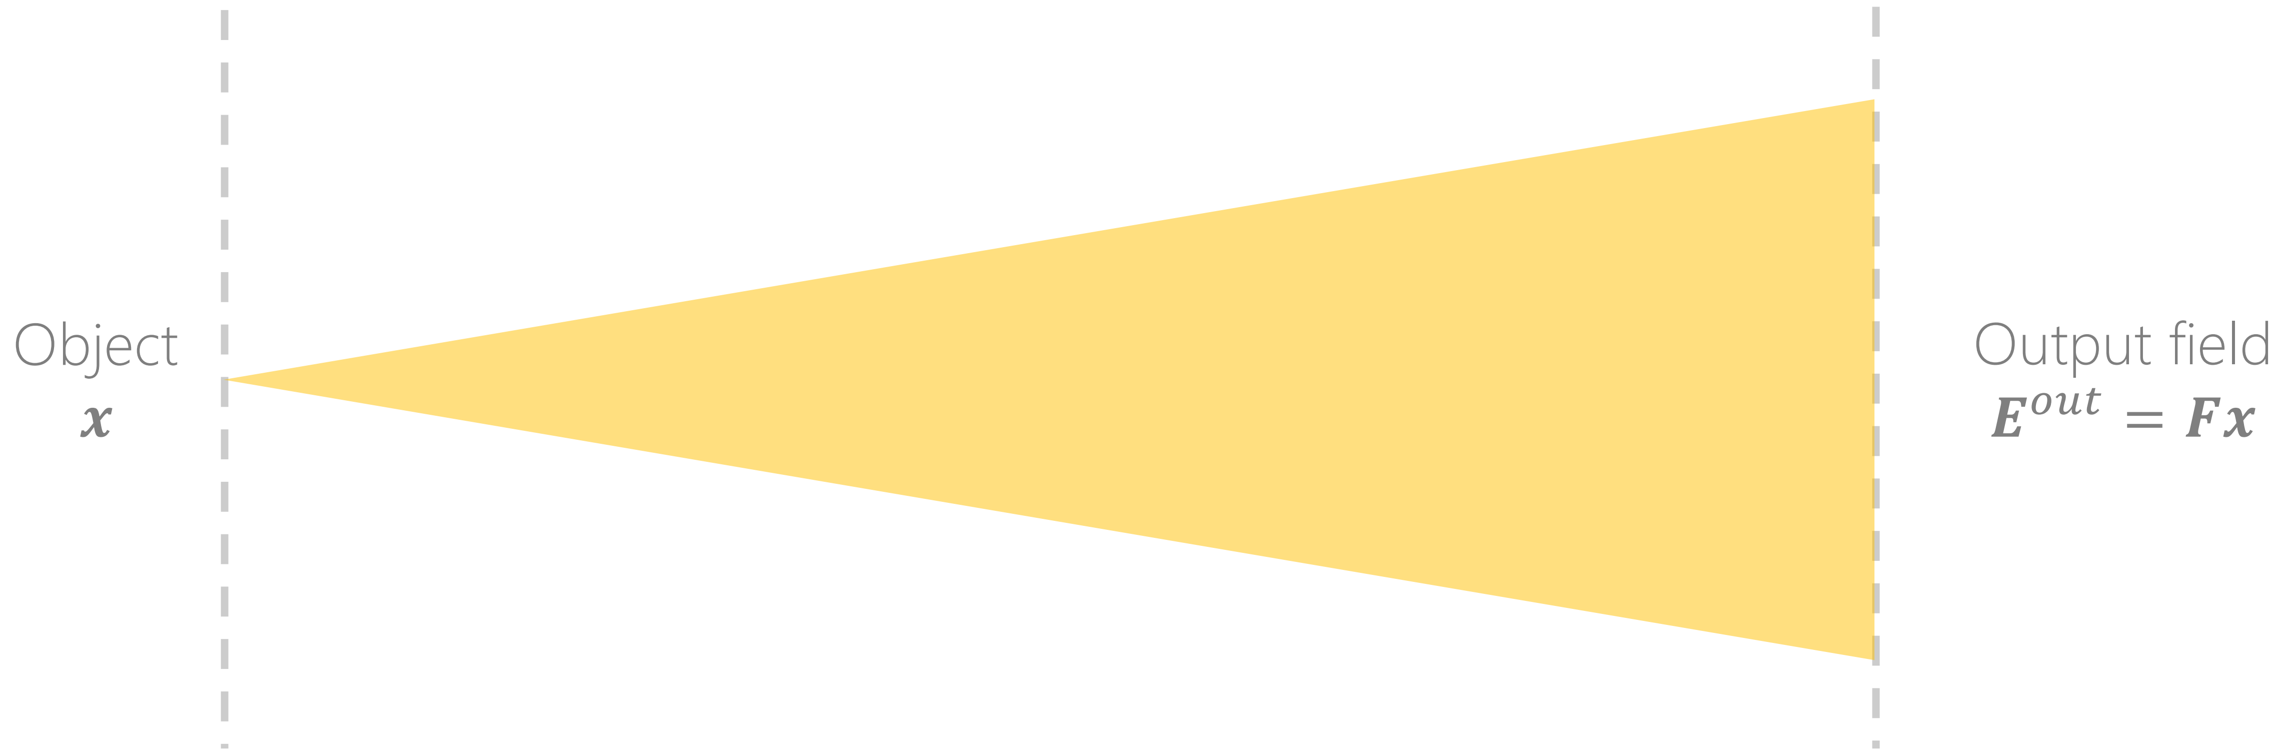


Far-field propagation
Fourier transform $\boldsymbol{F}$

# Fourier in optics

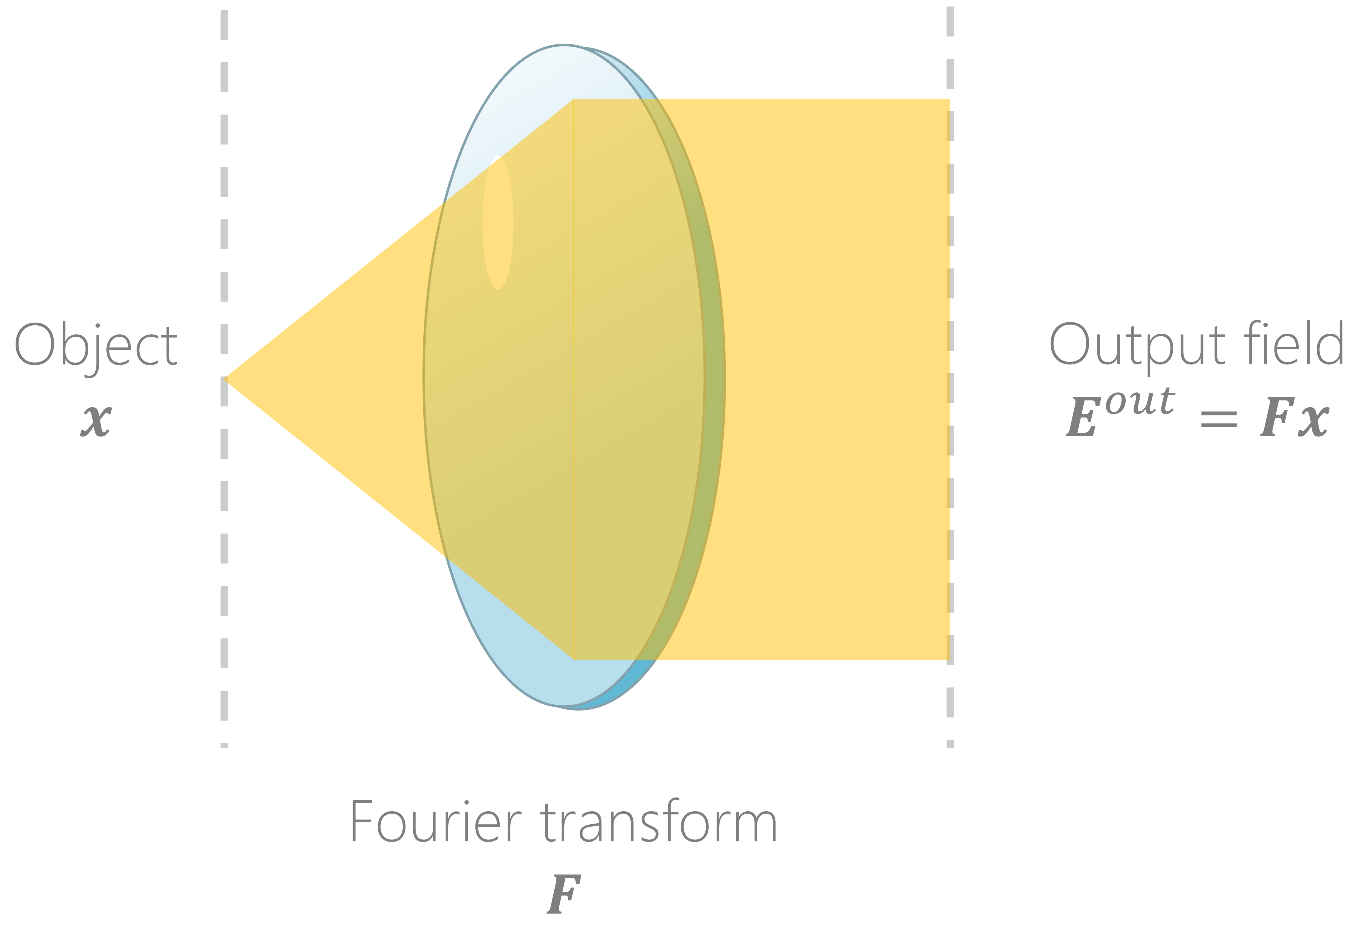

# Fourier in optics

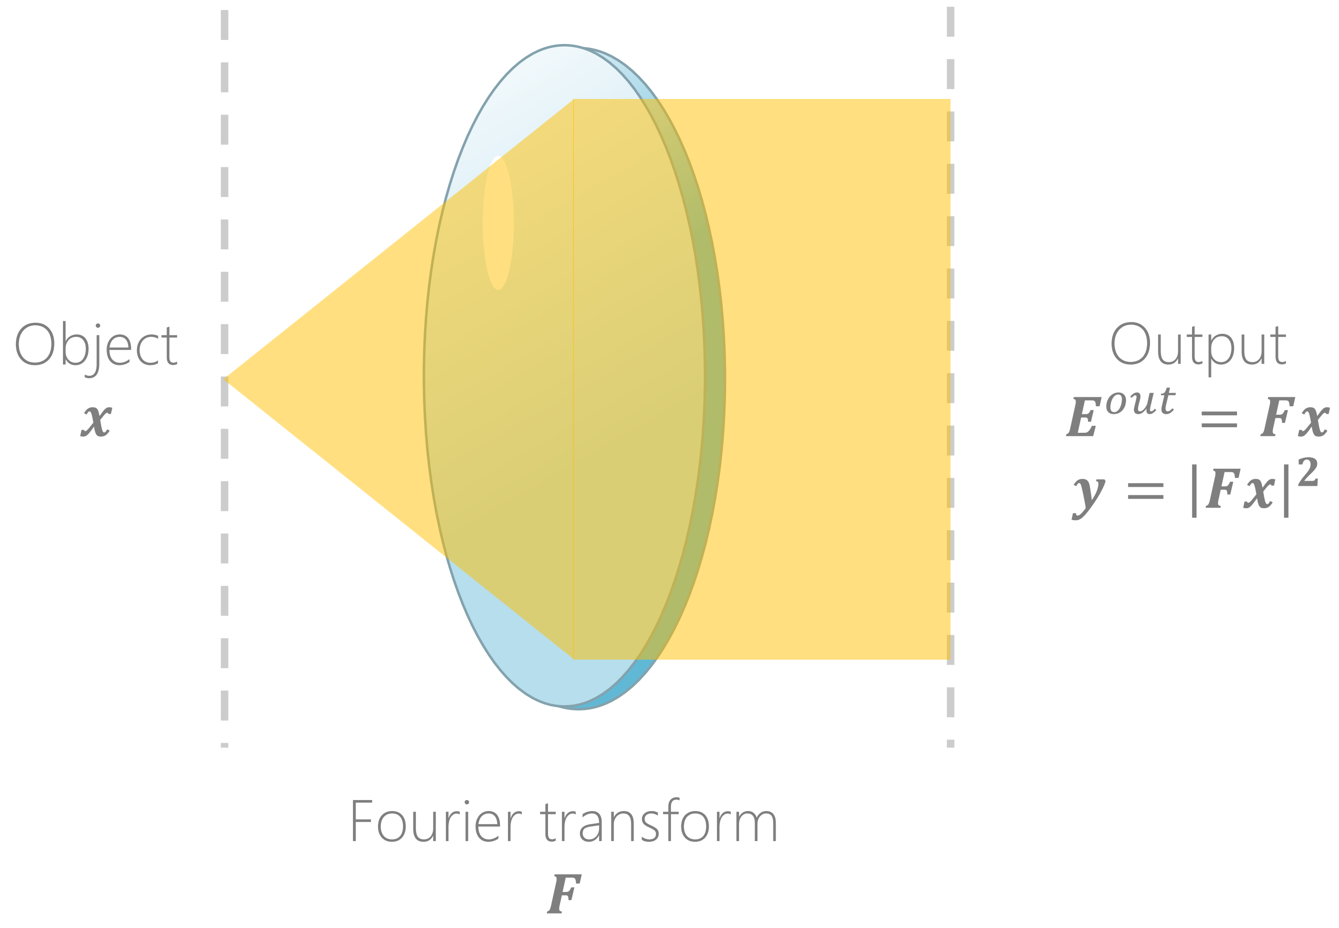


- Applications
  - Crystallography
  - Adaptive optics
  - PSF engineering
  - Complex media imaging
  - Non-line of sight imaging

Fienup (1982). *Applied Optics*

# Coded-illumination: ptychography

Object
$\boldsymbol{x}$

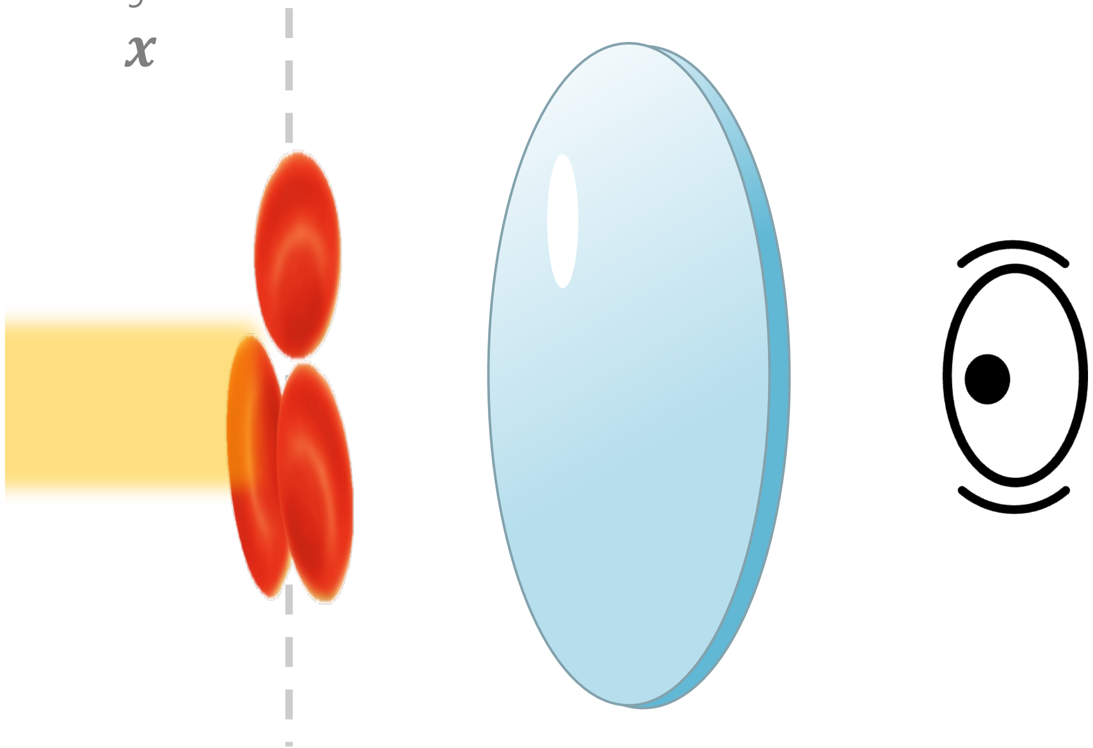

Output field
$\boldsymbol{E}^{out} = \boldsymbol{F}\boldsymbol{D}\boldsymbol{x}$

Illumination probe
$\boldsymbol{D}$

# Coded-illumination: ptychography

Object
$\boldsymbol{x}$

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

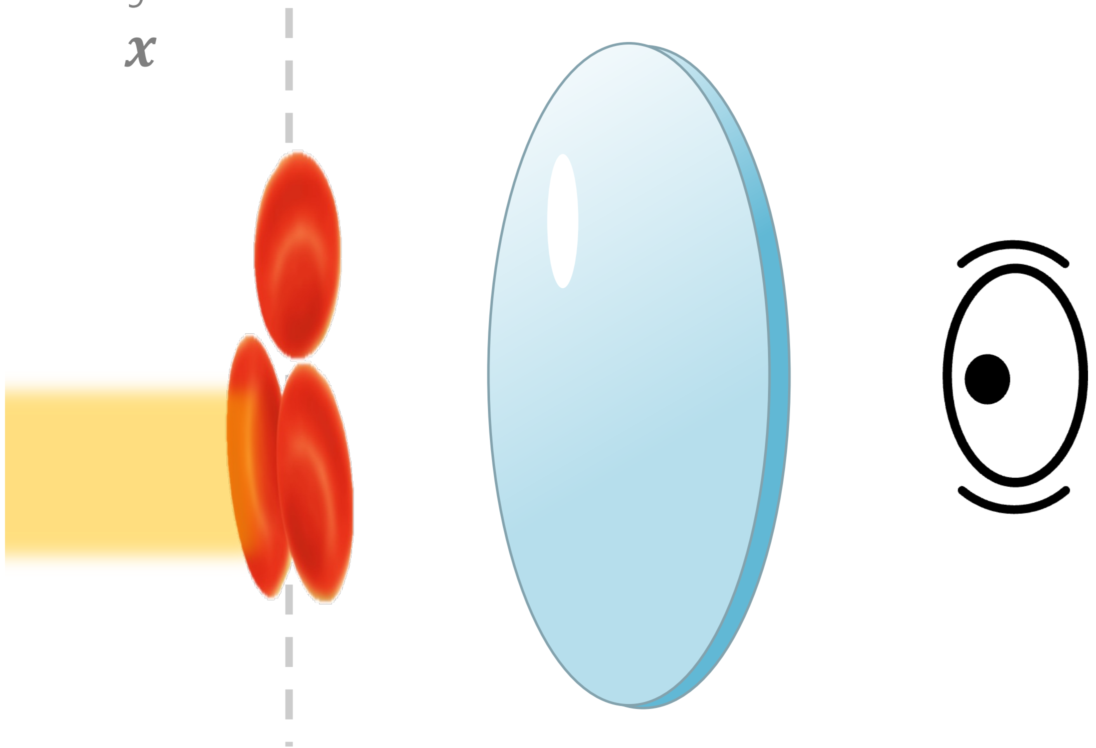

Illumination probe
$\boldsymbol{D}_l$

# Coded-illumination: ptychography

Object
$\boldsymbol{x}$

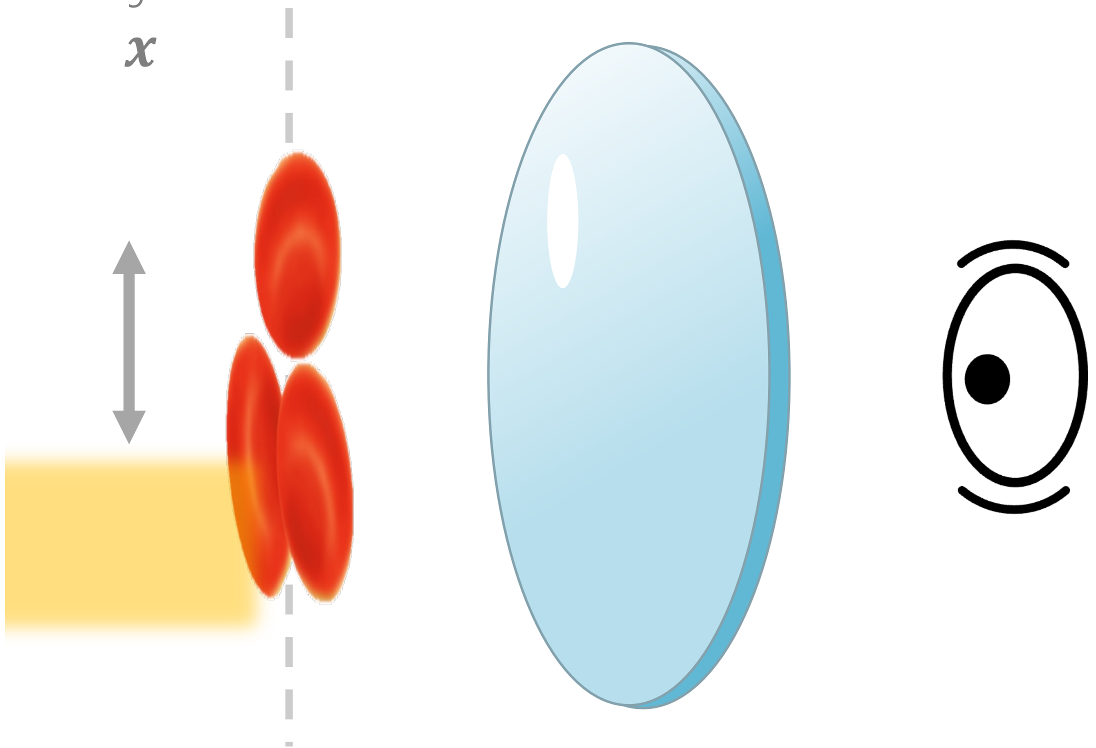

Illumination probe
$\boldsymbol{D}_l$

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

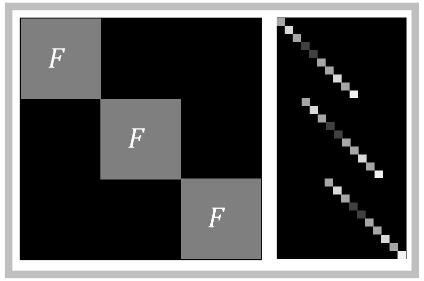


Miao, Charalambous, Kirz, Sayre (1999). *Nature*

# Coded-illumination experiments

Ptychography

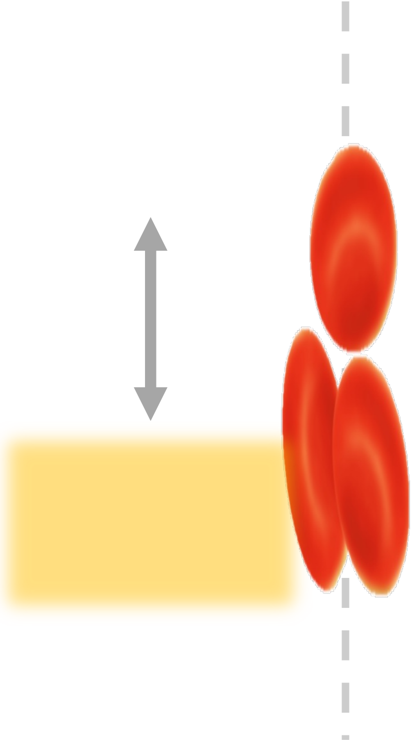

Shifted illumination

Fourier Ptychography

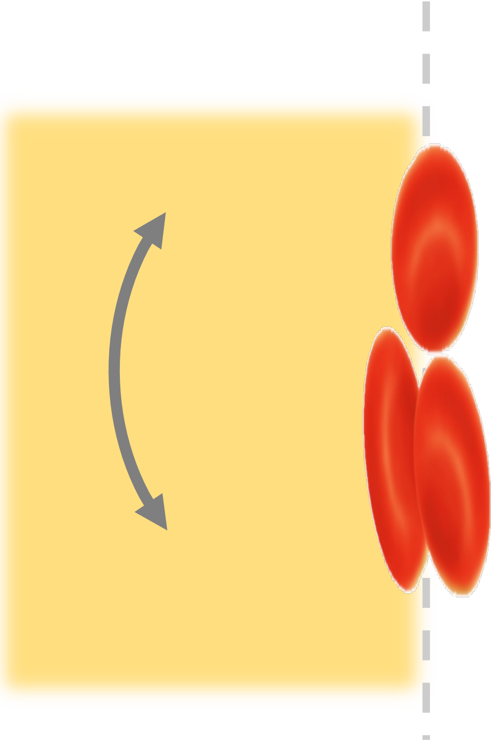

Tilted illumination

Coded Diffraction Imaging

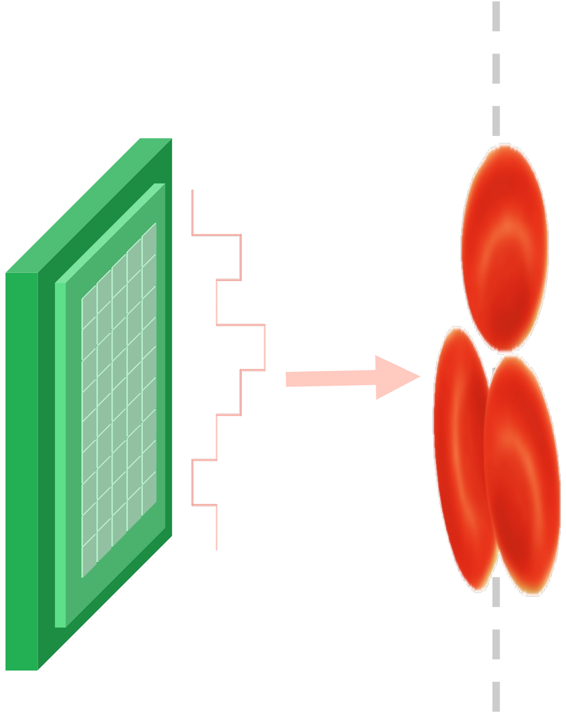

Arbitrary pattern
(Spatial Light Modulator,
mask, grating)

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

# Coded detection

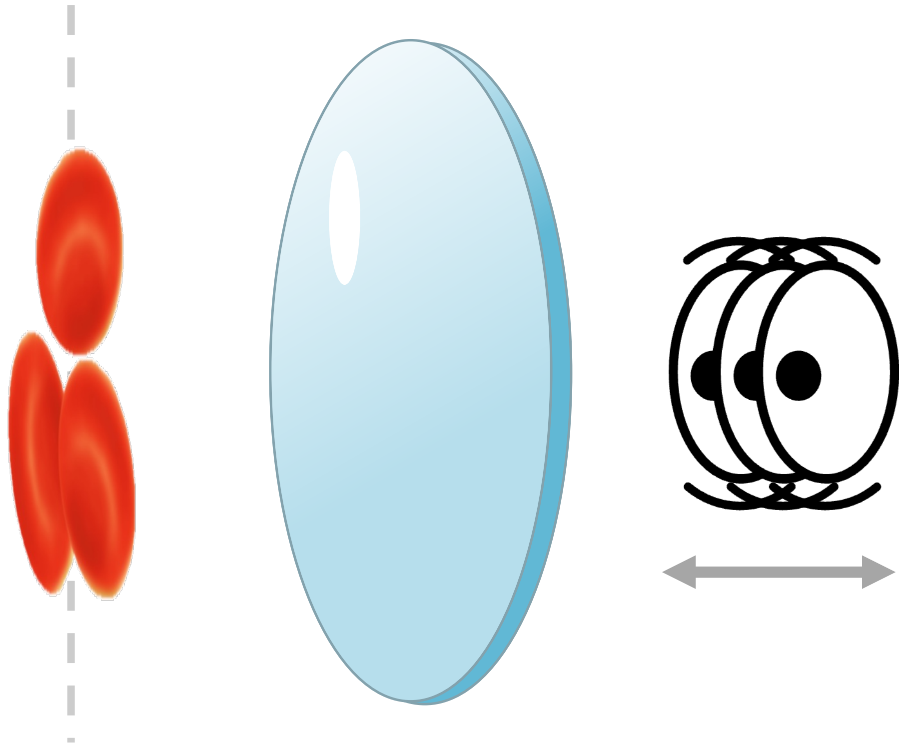

- Applications:
  - Astronomy
  - Non-invasive bioimaging

$\Rightarrow$ 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

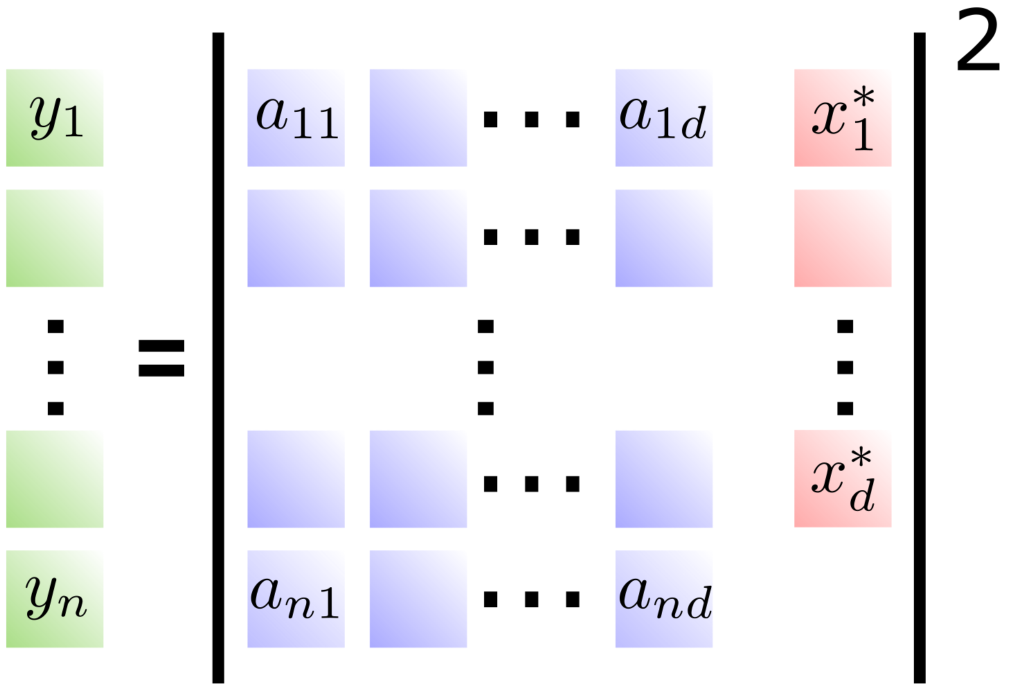


Mondelli, Montanari (2019). *Foundations of Computational Mathematics*

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

## Fourier phase retrieval

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

by a lens

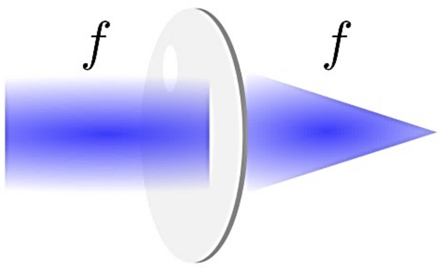


by free-space propagation

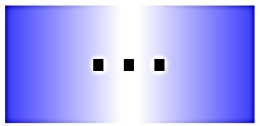

## Coded illumination

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

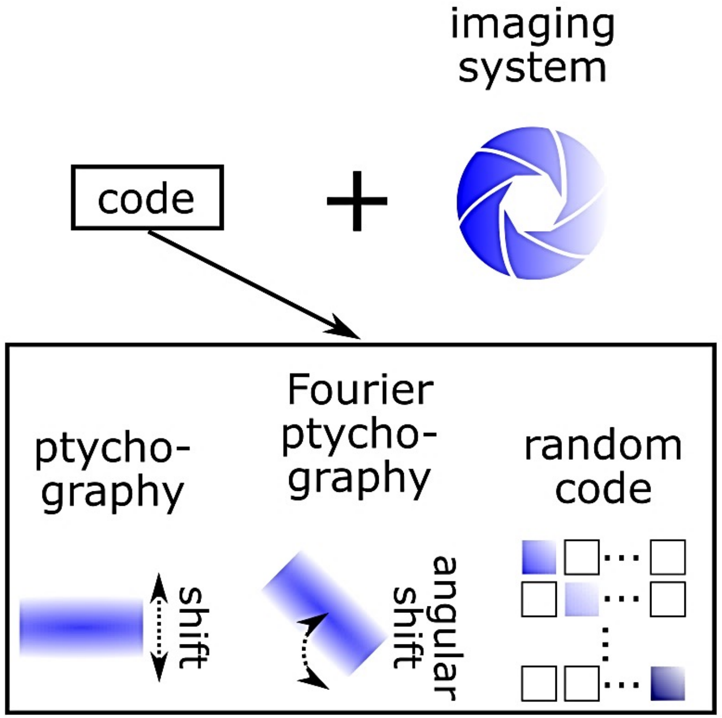


## Coded detection

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

phase diversity

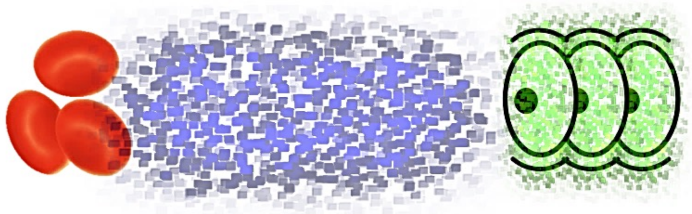

computer-generated holography

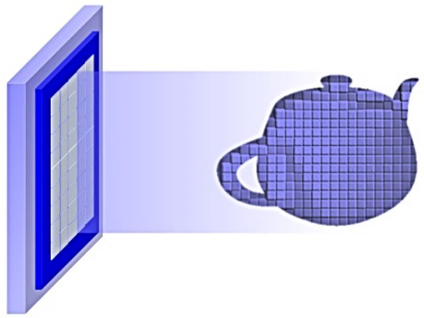

## Random

by random projections

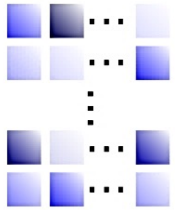

by propagation through complex media

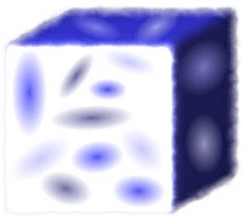

Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). *IEEE Signal Processing Magazine*

# Content

Find $\boldsymbol{x}^*$ in

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

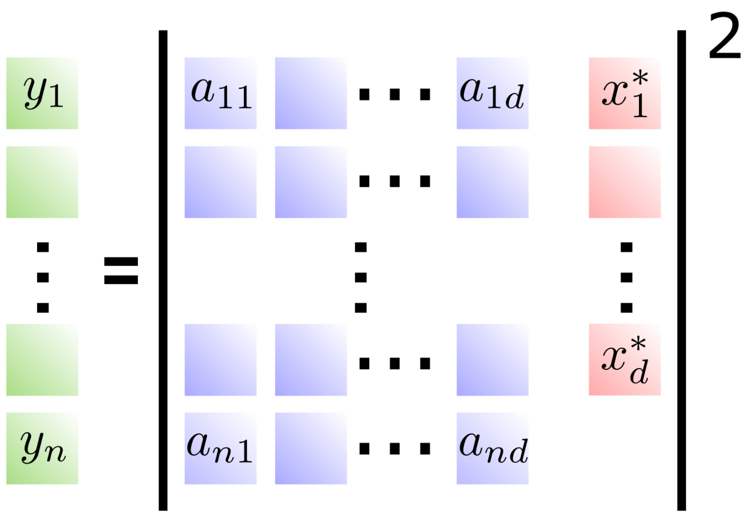


- Physics: Phase imaging devices
- Math:
  - Inverse problems
  - Phase retrieval algorithms
  - Phase retrieval theory
- Machine learning: Regularization

# 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

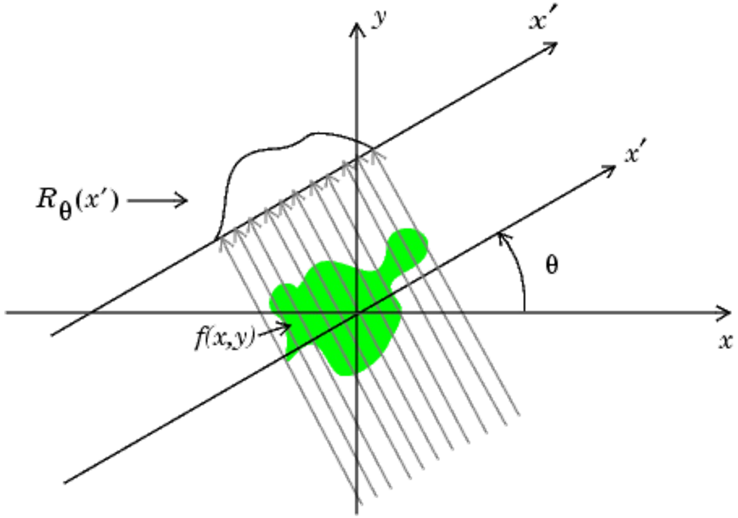


# Linear systems

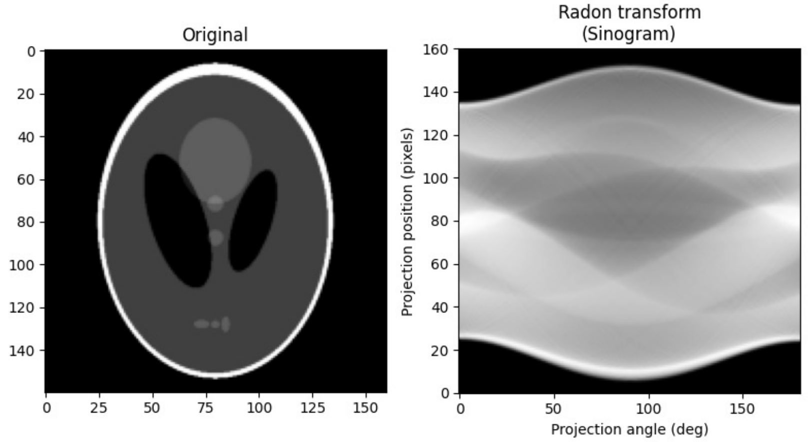


Source: scikit-image

Find $\boldsymbol{x}^*$ in

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

with a linear operator $\boldsymbol{A}$

- Examples:
  - Radon projection: computed tomography

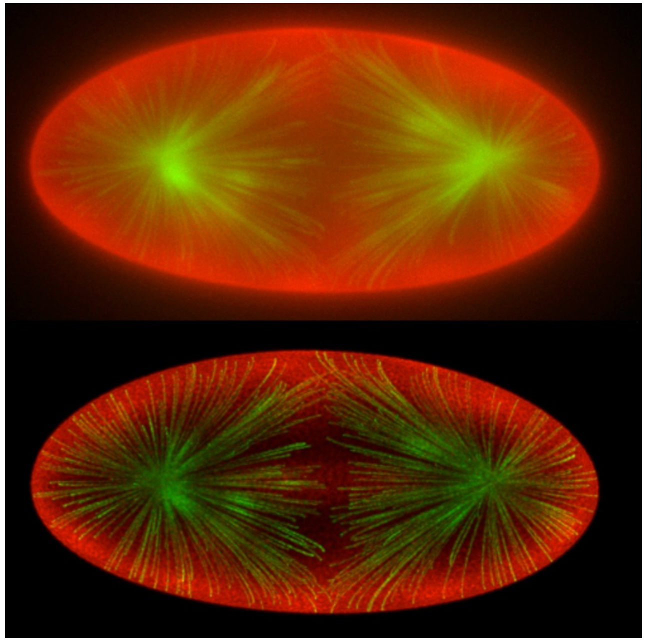

Source: BIG EPFL

# Linear systems

Find $\boldsymbol{x}^*$ in

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

with a linear operator $\boldsymbol{A}$

- Examples:
  - Radon projection: computed tomography
  - Convolution by PSF: superresolution microscopy

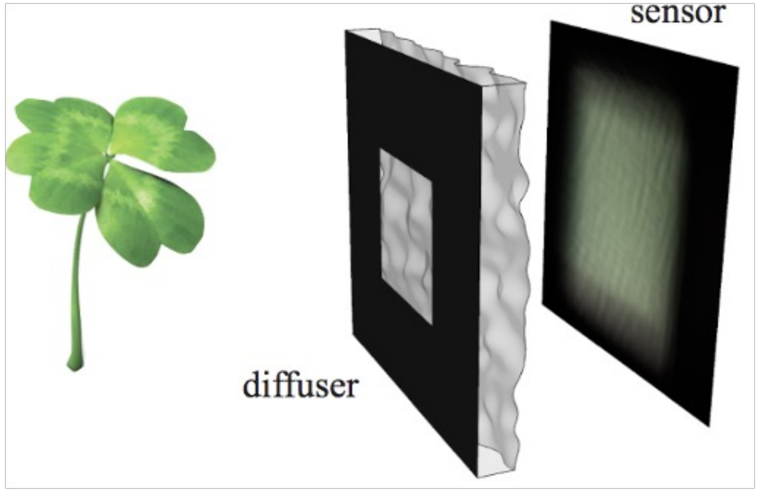

Source: Waller lab, UC Berkeley

# Linear systems

Find $\boldsymbol{x}^*$ in

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

with a linear operator $\boldsymbol{A}$

- Examples:
  - Radon projection: computed tomography
  - Convolution by PSF: superresolution microscopy
  - Random convolution: lensless imaging

# Linear systems

Find $\boldsymbol{x}^*$ in

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

with a linear operator $\boldsymbol{A}$

- Examples:
  - Radon projection: computed tomography
  - Convolution by PSF: superresolution microscopy
  - Random convolution: lensless imaging
  - Random $\boldsymbol{A}$: compressed sensing

# Non-linear systems

Find $\boldsymbol{x}^*$ in

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

with a linear operator $\boldsymbol{A}$

- Other examples:
  - 3D reconstruction (multislice, etc.)
  - Aberration correction
  - Nonlinear optics (fluorescence, spontaneous/coherent Raman, second-/third-harmonic generation)

# Inverse problem framework

Find $\boldsymbol{x}^*$ in

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

- Optimization approach: L2 loss

$$\hat{x} = \mathrm{argmin}_x \|y - A\{x\}\|_2^2$$

- Convex optimization for linear problems
  - But sometimes non-invertible
    (undetermined system with many solutions)
- Non-convex optimization for non-linear problems

# Regularization

Find $\boldsymbol{x}^*$ in

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

- Optimization approach:

$$\hat{x} = \mathrm{argmin}_x f(y, x) + \lambda g(x)$$

- Data fidelity + regularization
  - Tikhonov L2: $g(x) = \|x\|_2^2$
  - Sparsity L1: $g(x) = \|x\|_1$
  - Total variation: $g(x) = \|\nabla x\|_1$
  - Deep learning regularization
- $\lambda$: Regularization constant to tune

# A statistical approach

- $\boldsymbol{y} = \boldsymbol{A}\{\boldsymbol{x}^*\} + \boldsymbol{n}$ with additional noise term $\boldsymbol{n}$

Likelihood $p(\boldsymbol{y}|\boldsymbol{x})$: models noise and forward map $\boldsymbol{A}$

- Prior $p(\boldsymbol{x})$: encodes regularity / structural assumptions
- Posterior distribution (Bayes law):

$$p(\boldsymbol{x}|\boldsymbol{y}) \propto p(\boldsymbol{y}|\boldsymbol{x})p(\boldsymbol{x})$$

- Maximum a posteriori (MAP) estimator:

$$\begin{aligned}\hat{x} &= \text{argmax}_x \ln p(\boldsymbol{x}|\boldsymbol{y}) \\ &= \text{argmax}_x \ln p(\boldsymbol{y}|\boldsymbol{x}) + \ln p(\boldsymbol{x})\end{aligned}$$

# A statistical approach

- $\boldsymbol{y} = \boldsymbol{A}\{\boldsymbol{x}^*\} + \boldsymbol{n}$ with additional noise term $\boldsymbol{n}$
- Maximum a posteriori (MAP) estimator:

$$\hat{x} = \mathrm{argmin}_x - \ln p(\boldsymbol{y}|\boldsymbol{x}) - \ln p(\boldsymbol{x})$$

- Gaussian noise: $p(\boldsymbol{y}|\boldsymbol{x}) = \exp\left(-\frac{\|y-A\{x\}\|^2}{2\sigma^2}\right)$

$$-\ln p(\boldsymbol{y}|\boldsymbol{x}) = \frac{1}{2\sigma^2}\|y - A\{x\}\|^2$$

  We get back the L2 loss
- Other possible models:
  - Poisson distribution
  - Poisson-Gaussian distribution

Seifert, Shao, Bouchet, et al 2023

# A statistical approach

- $\boldsymbol{y} = \boldsymbol{A}\{\boldsymbol{x}^*\} + \boldsymbol{n}$ with additional noise term $\boldsymbol{n}$
- Maximum a posteriori (MAP) estimator:

$$\hat{x} = \operatorname{argmin}_x - \ln p(\boldsymbol{y}|\boldsymbol{x}) - \ln p(\boldsymbol{x})$$

- Gaussian prior: $p(\boldsymbol{x}) = \exp\left(-\frac{\|x\|^2}{2\tau^2}\right)$

$$-\ln p(\boldsymbol{x}) = \frac{1}{2\tau^2}\|x\|_2^2$$

  We get back Tikhonov regularization.
- Laplace prior: $p(\boldsymbol{x}) = \exp\left(-\frac{\|x\|_1}{\tau}\right)$ => L1 regularization (sparsity)

# Content

Find $\boldsymbol{x}^*$ in

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

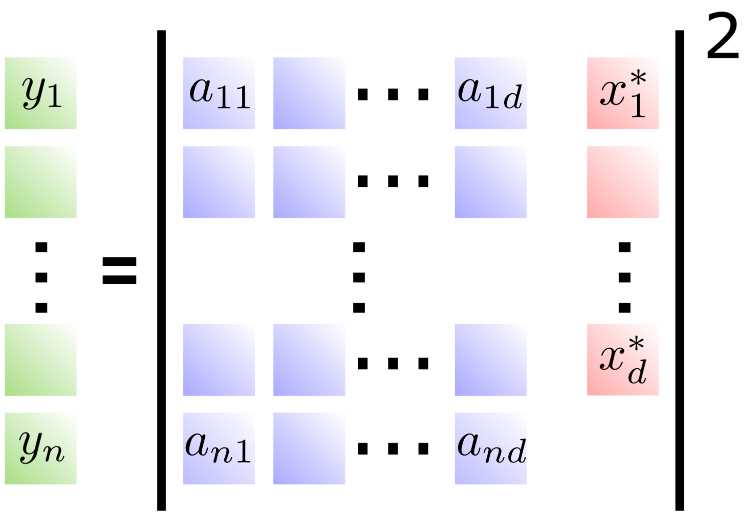


- Physics: Phase imaging devices
- Math:
  - Inverse problems
  - Phase retrieval algorithms
  - Phase retrieval theory
- Machine learning: Regularization

# Gradient descent

Find $\boldsymbol{x}^*$ in

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

- Optimization approach:

$$\hat{x} = \mathrm{argmin}_x\ f(y, x) + g(x)$$

- 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:

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

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:

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

  - 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 $y_i = |a_i^H x|^2$

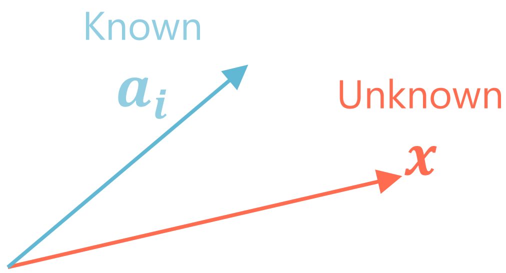


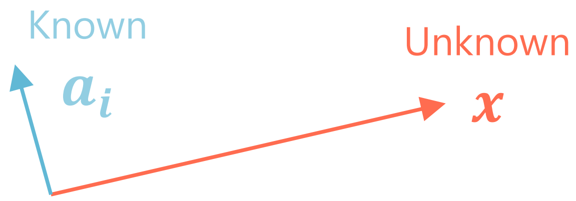


$\boldsymbol{a}_i$ and $\boldsymbol{x}$
correlated
$\Rightarrow$ high intensity $y_i$

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

# Spectral methods

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

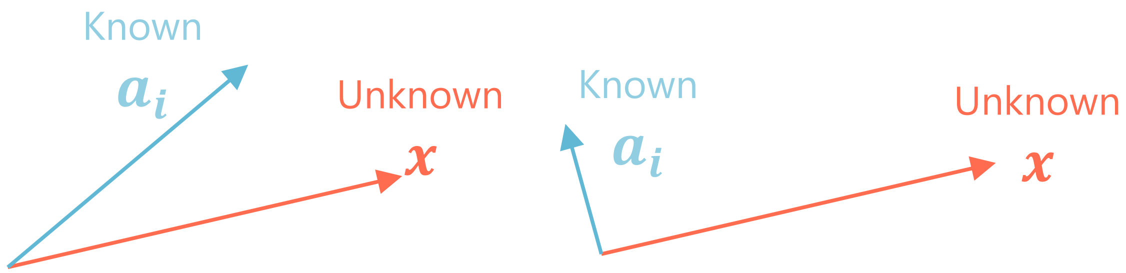


$\boldsymbol{a}_i$ and $\boldsymbol{x}$
correlated
$\Rightarrow$ high intensity $y_i$

$\boldsymbol{a}_i$ 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 $y_i = |a_i^H x|^2$

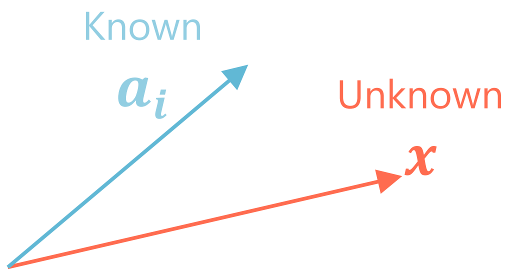


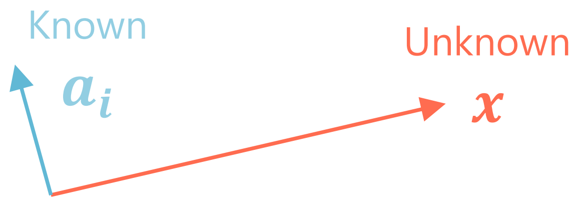


$\boldsymbol{a}_i$ and $\boldsymbol{x}$
correlated
$\Rightarrow$ high intensity $y_i$

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

**Spectral 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 $y_i = \left|a_i^H x\right|^2$

| | |
|---|---|
| $\mathcal{T}_0(y) = y$ <br> $\mathcal{T}_1(y) = (y > T)$ | *E. Candes, et al, IEEE Trans. on Information Theory (2015)* <br> *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**

Fienup '82
Maiden '09

**Gradient-based optimization**

Fienup '93
Yeh, Dong '15
Chen '18

**Convex relaxation**

Candès '11
Waldspurger '12
Goldstein '16

**Bayesian AMP**

Rangan '10
Metzler '17
Barbier '17
Maillard '20

**Spectral methods**

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

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

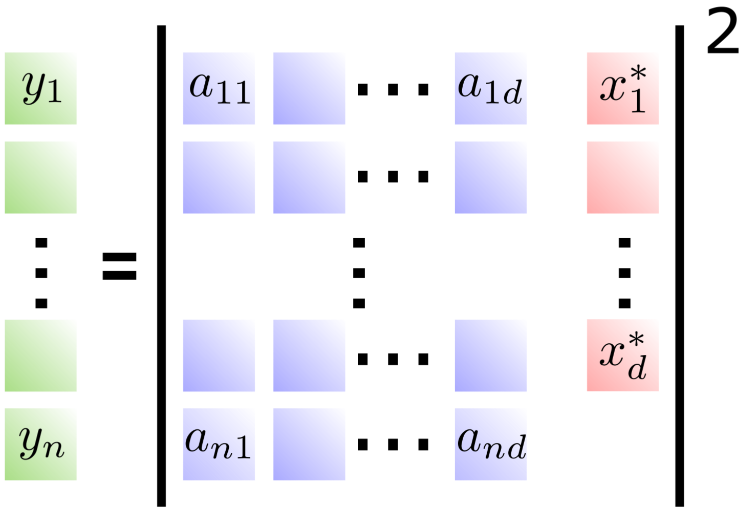


- Physics: Phase imaging devices
- Math:
  - Inverse problems
  - Phase retrieval algorithms
  - Phase retrieval theory
- Machine learning: Regularization

# Unifying framework $y = |Ax|^2$

**Fourier phase retrieval**

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

by a lens

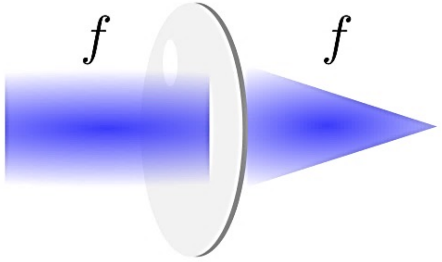


by free-space propagation

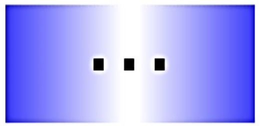

**Coded illumination**

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

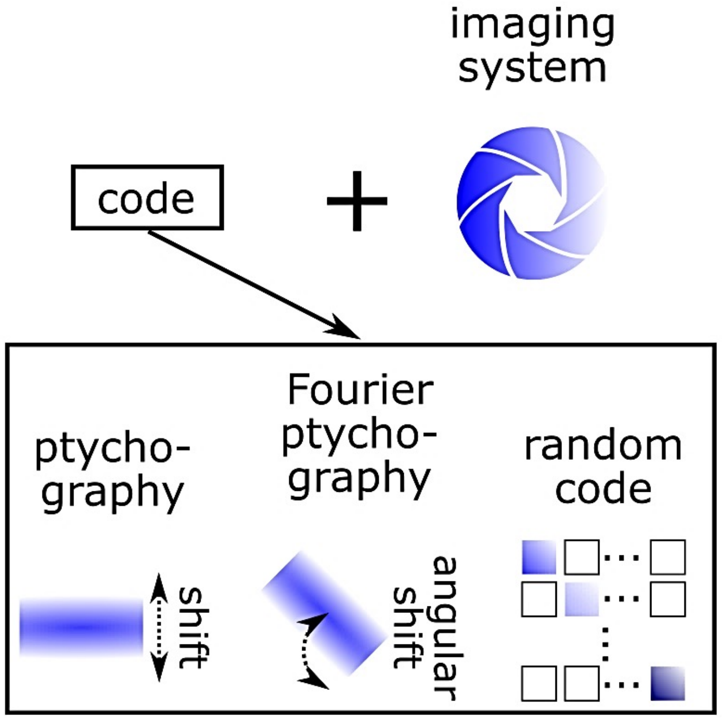


**Coded detection**

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

phase diversity

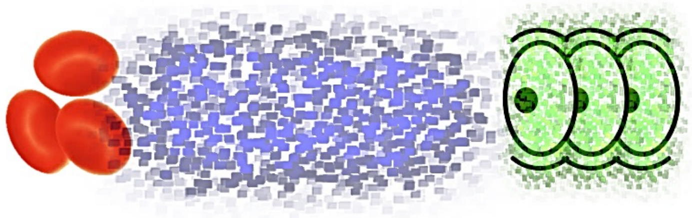

computer-generated holography

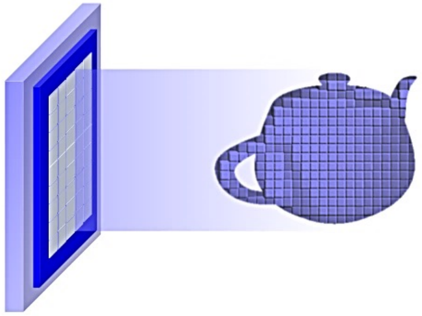

**Random**

by random projections

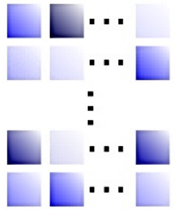

by propagation through complex media

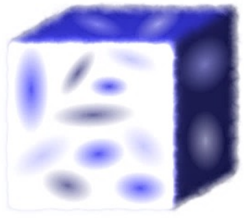

Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). *IEEE Signal Processing Magazine*

# Deeper dive in the random setting

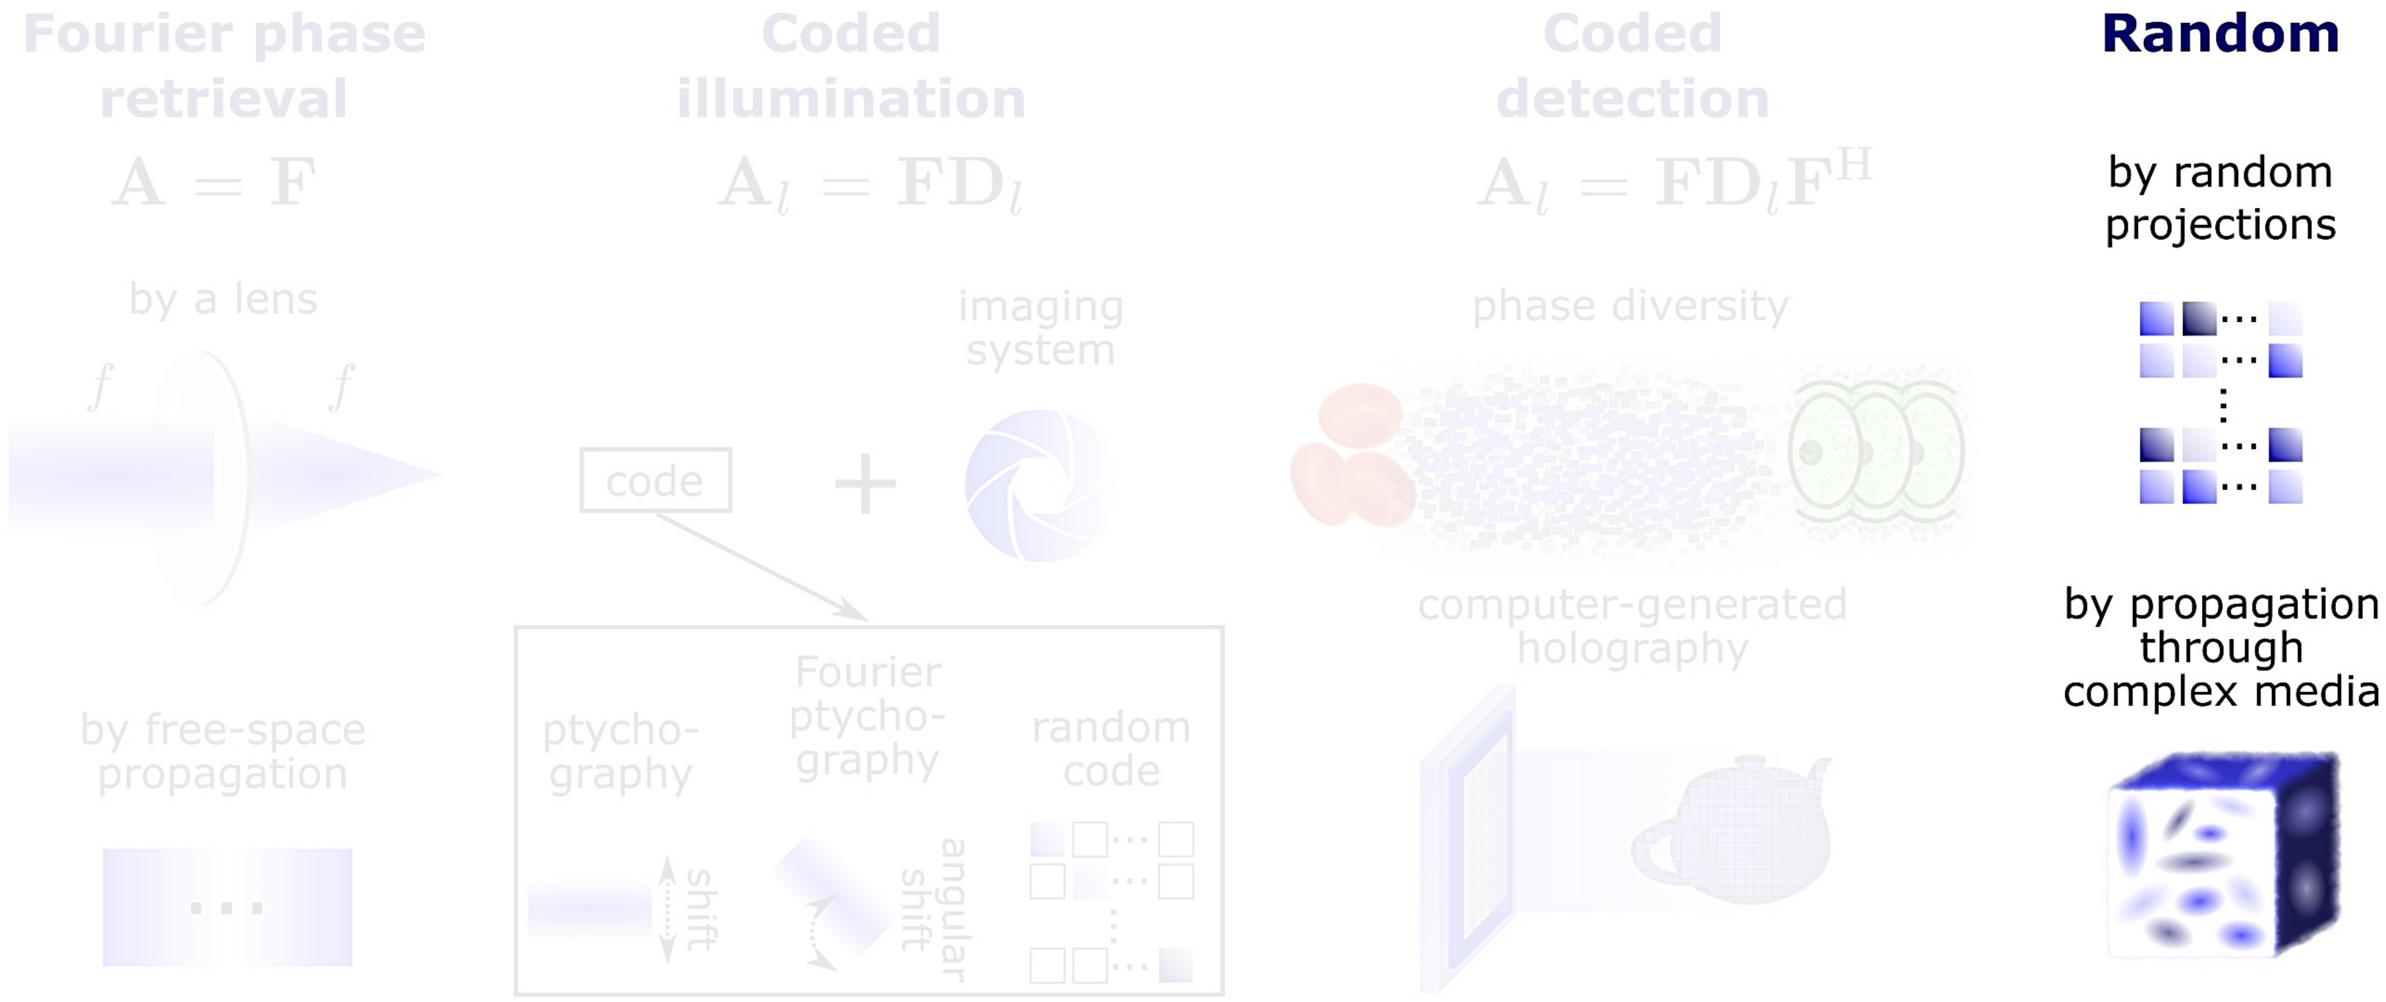


Dong, Valzania, Maillard, Pham, Gigan, Unser (2023). *IEEE Signal Processing Magazine*

# Theory break

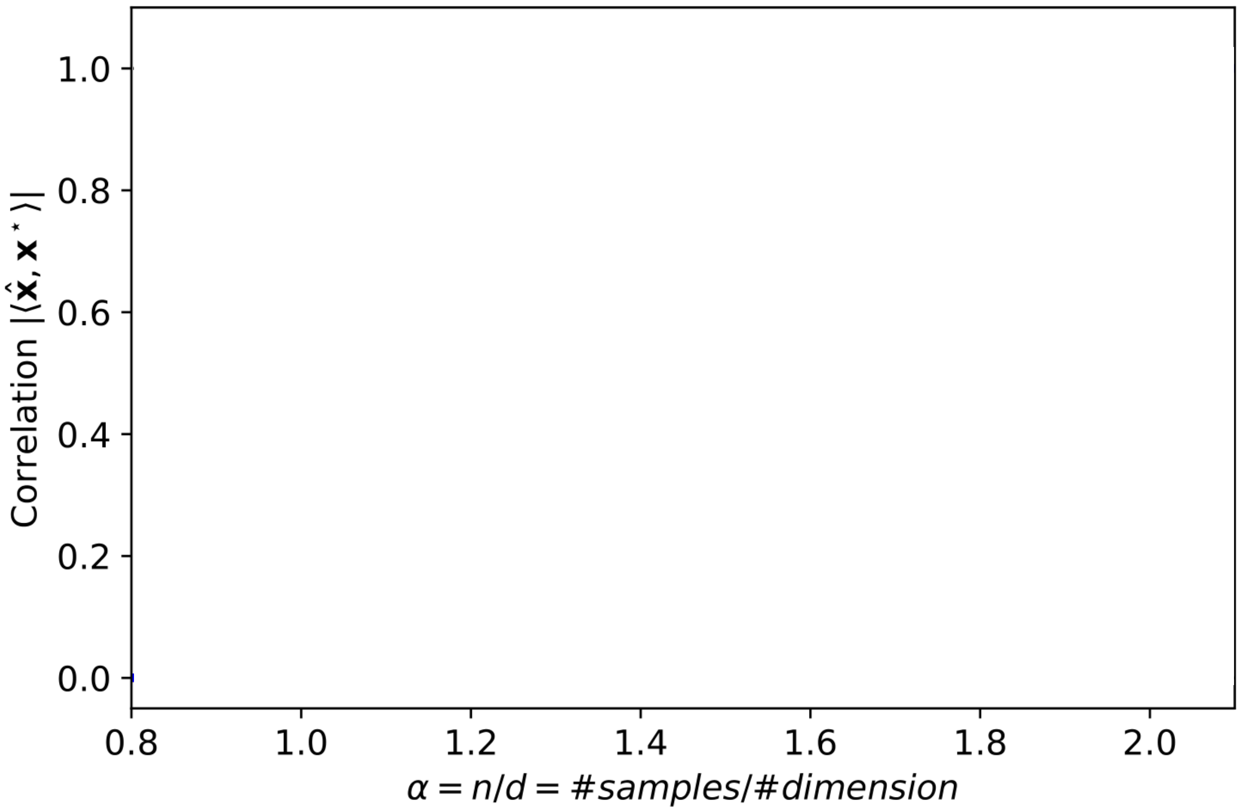


- Can we characterize performance as a function of oversampling $\alpha = n/d$?
- Correlation = higher is better

# Weak recovery

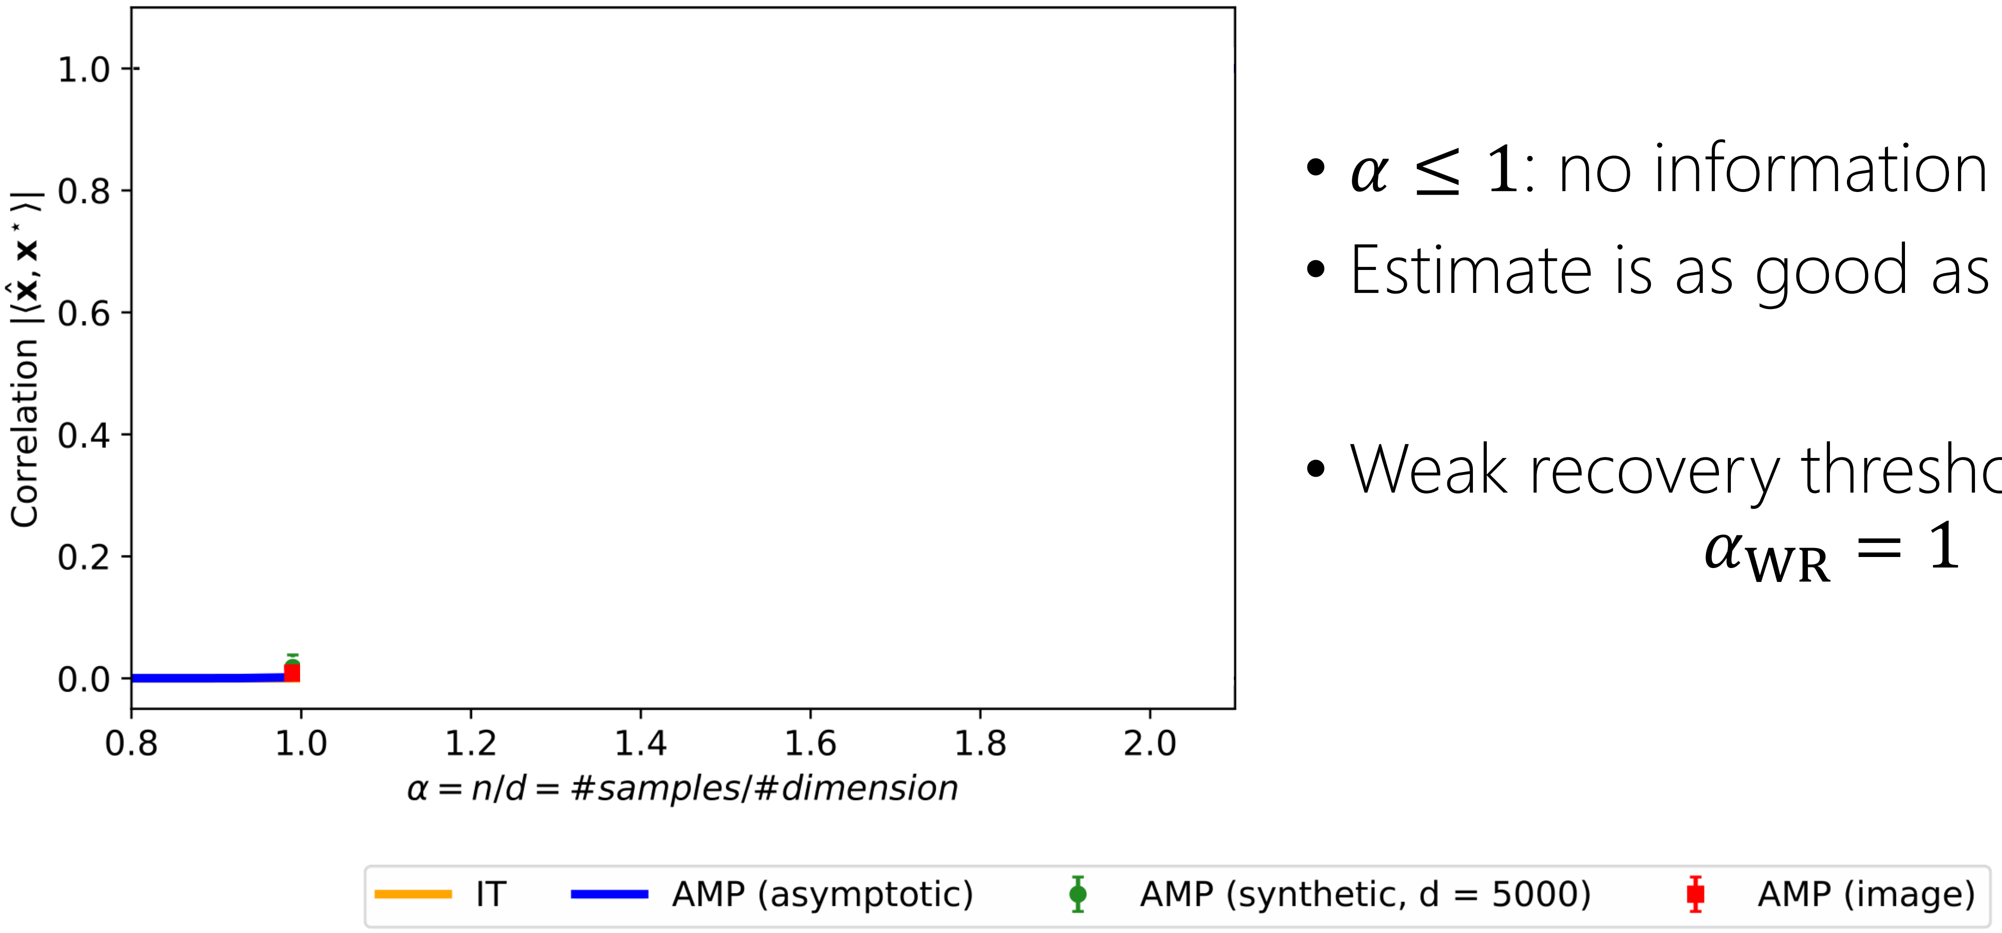


- $\alpha \leq 1$: no information on solution
- Estimate is as good as random

- Weak recovery threshold:

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

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

# Perfect recovery

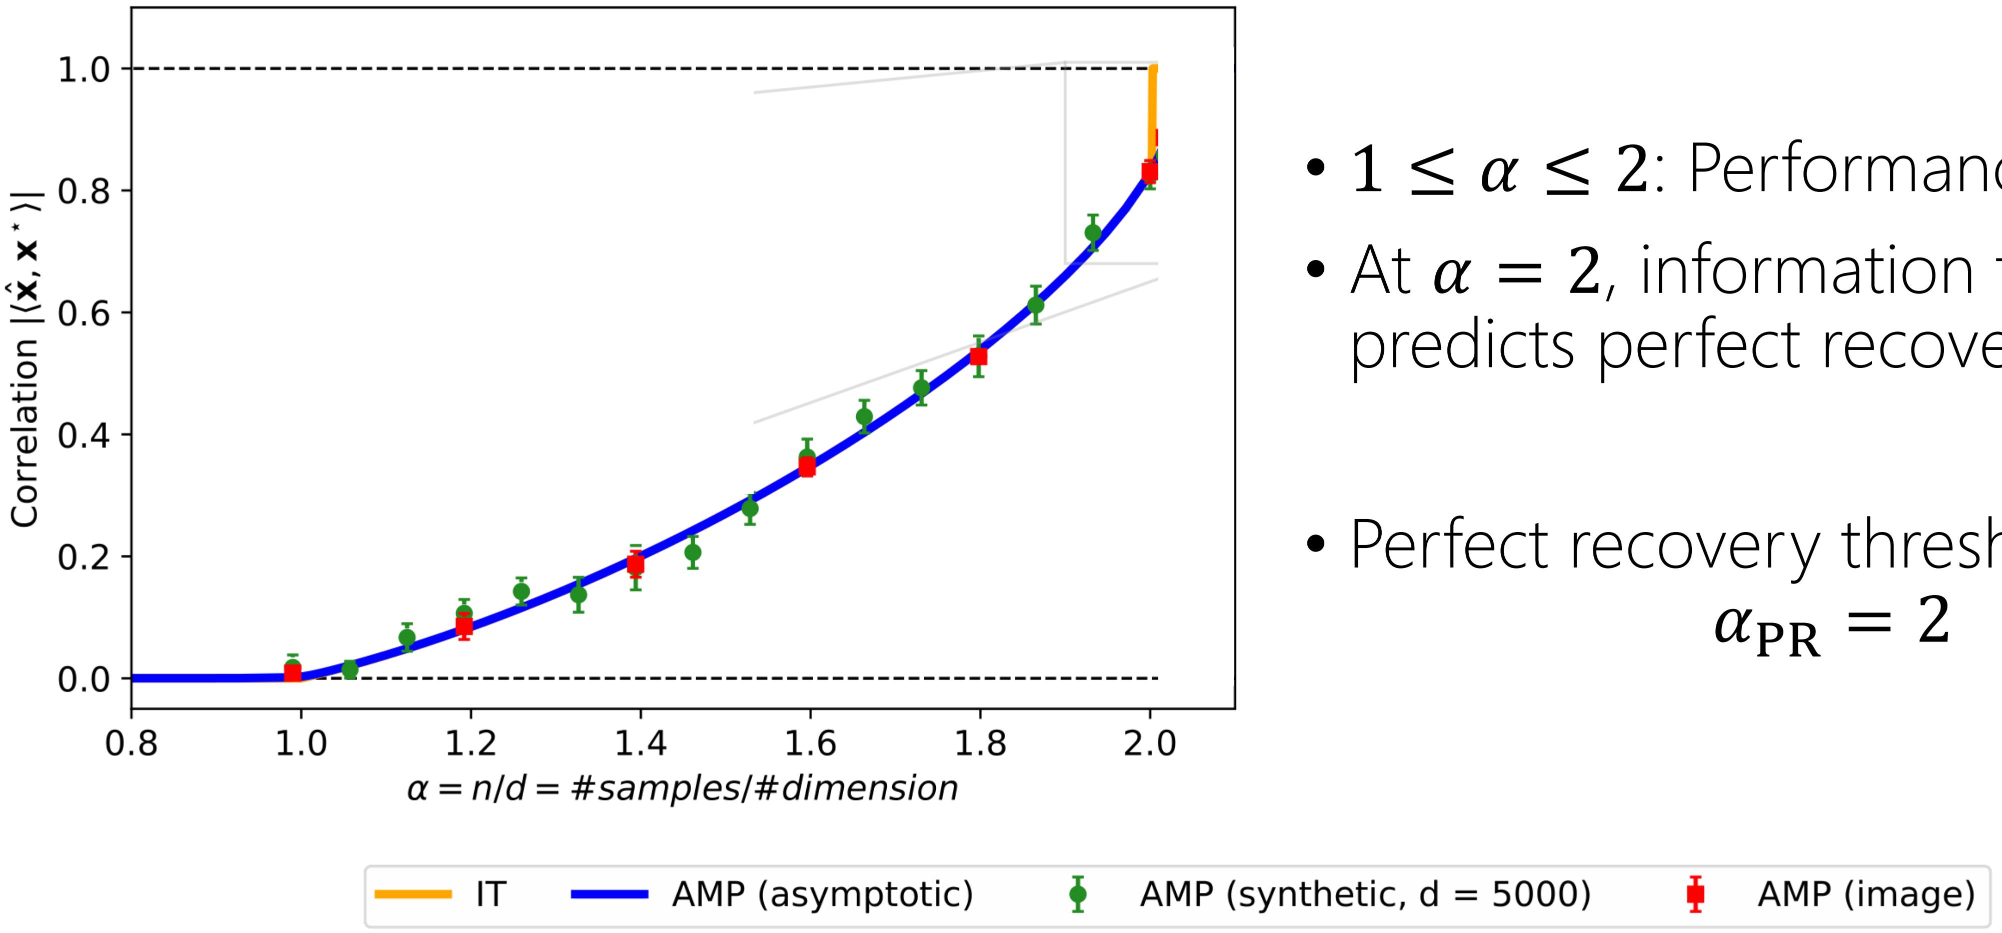


- $1 \leq \alpha \leq 2$: Performance improves
- At $\alpha = 2$, information theory predicts perfect recovery
- Perfect recovery threshold:
  $$\alpha_{\mathrm{PR}} = 2$$

# In practice

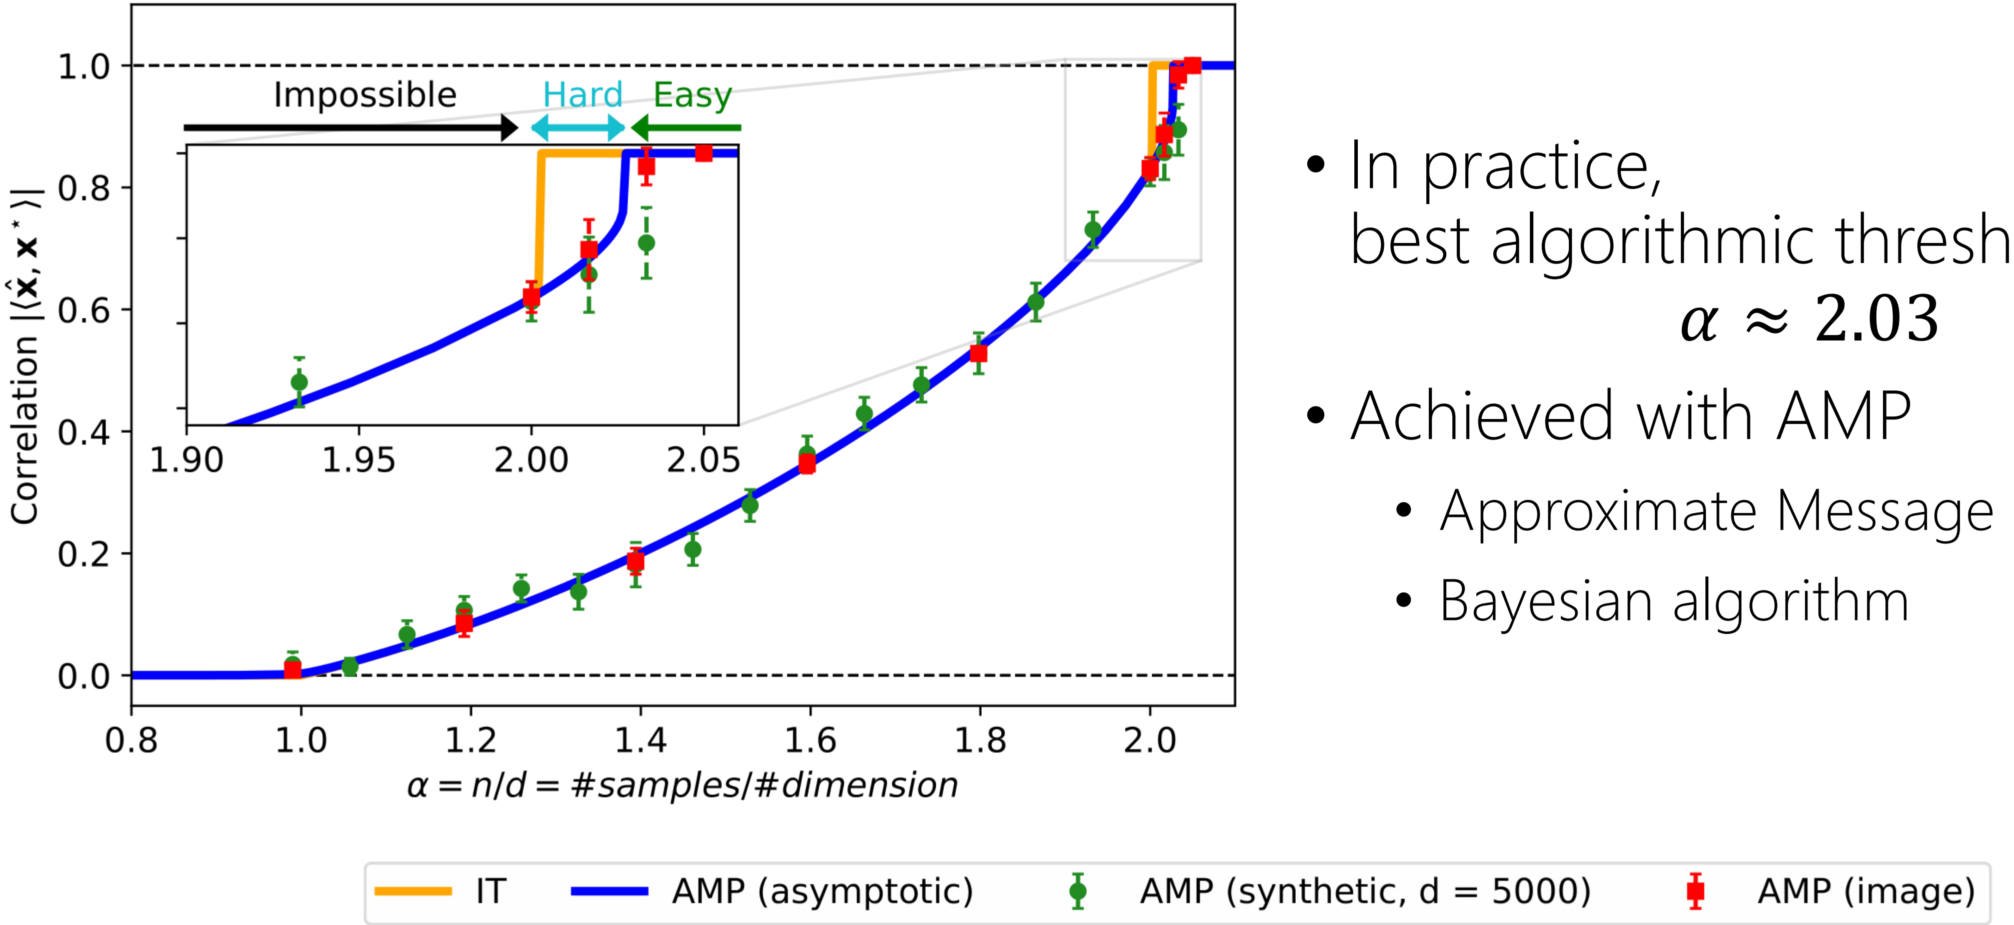


- In practice,
  best algorithmic threshold:
  $$\alpha \approx 2.03$$
- Achieved with AMP
  - Approximate Message Passing
  - Bayesian algorithm

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

# Example

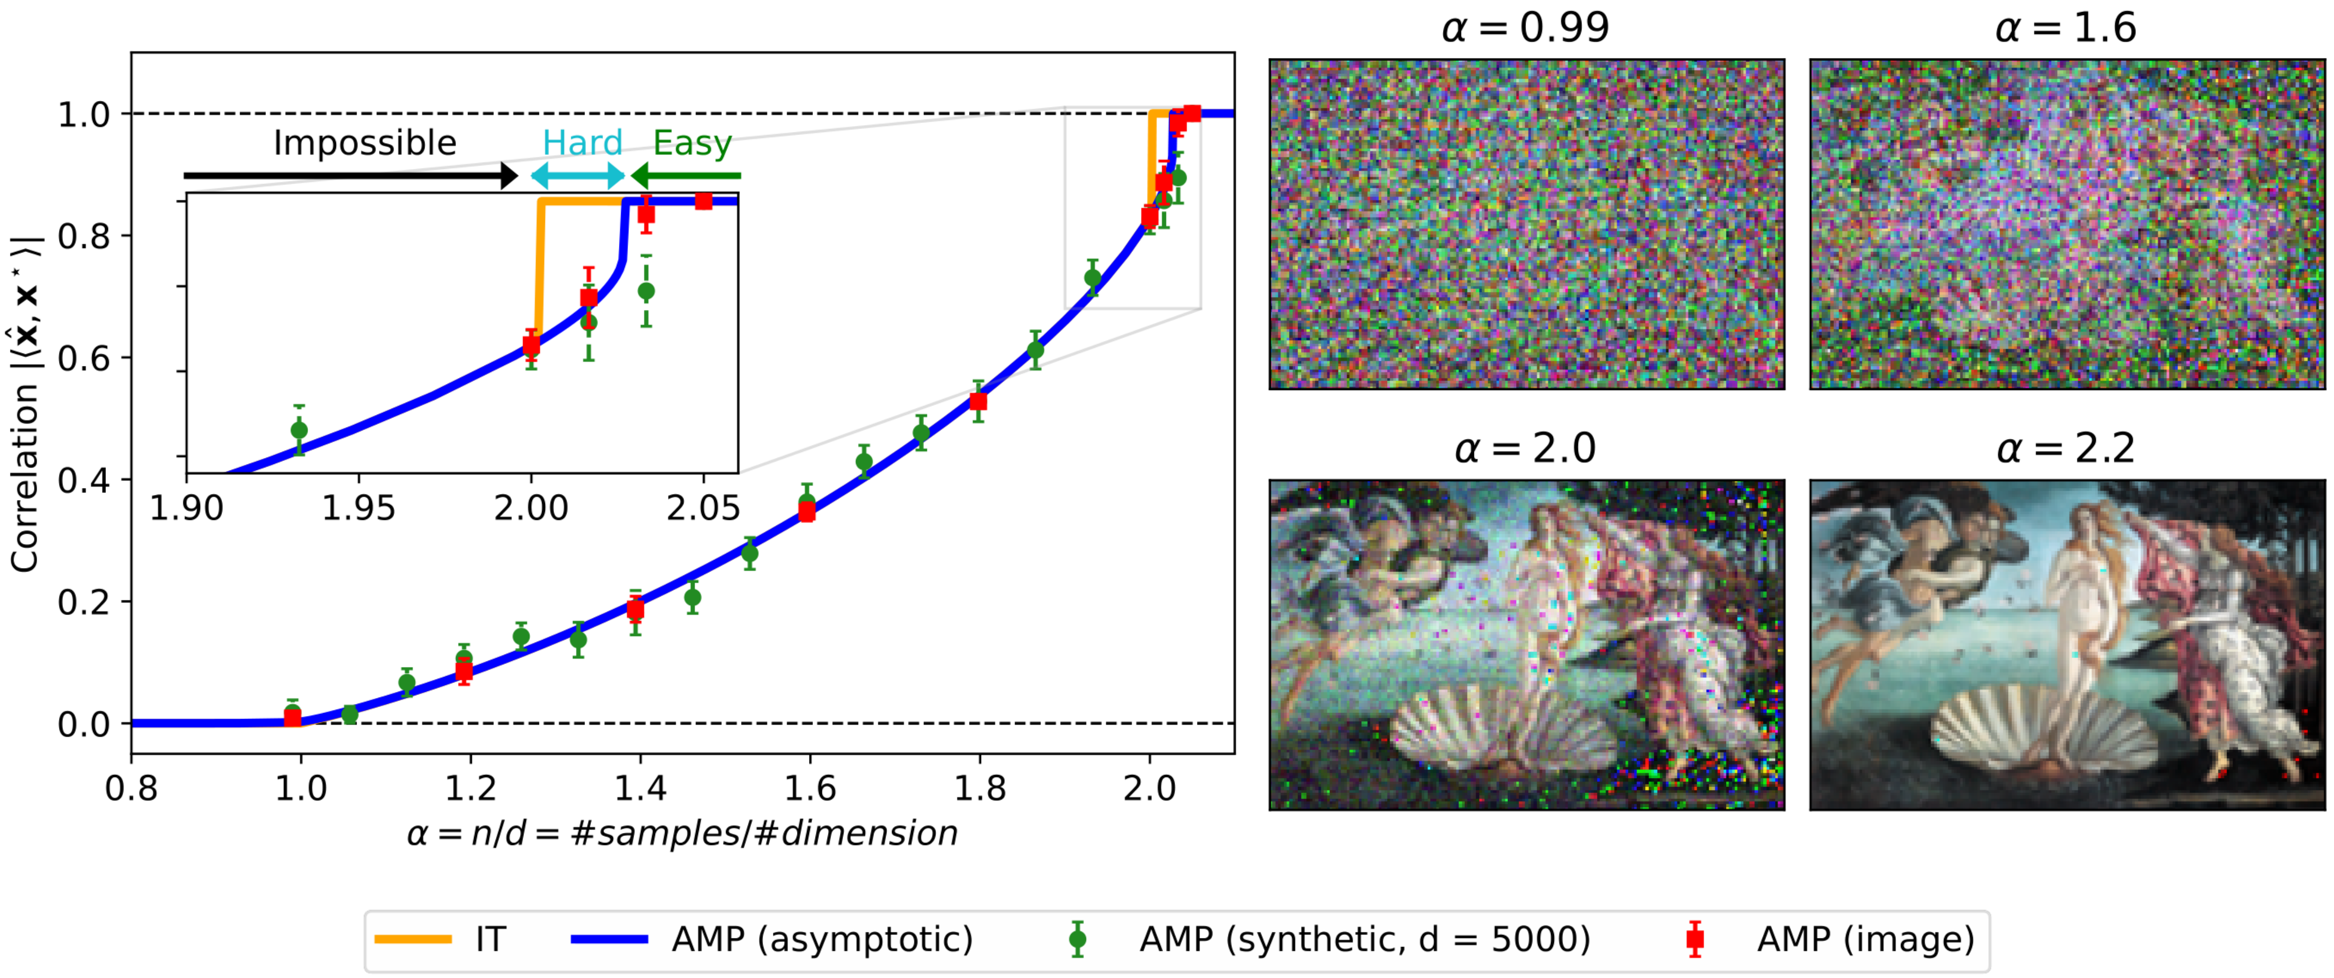

Impossible
Hard
Easy
Correlation $|\langle \hat{\mathbf{x}}, \mathbf{x}^\star \rangle|$
$\alpha = n/d = \#samples/\#dimension$
$\alpha = 0.99$
$\alpha = 1.6$
$\alpha = 2.0$
$\alpha = 2.2$
IT
AMP (asymptotic)
AMP (synthetic, d = 5000)
AMP (image)

# And in practice?

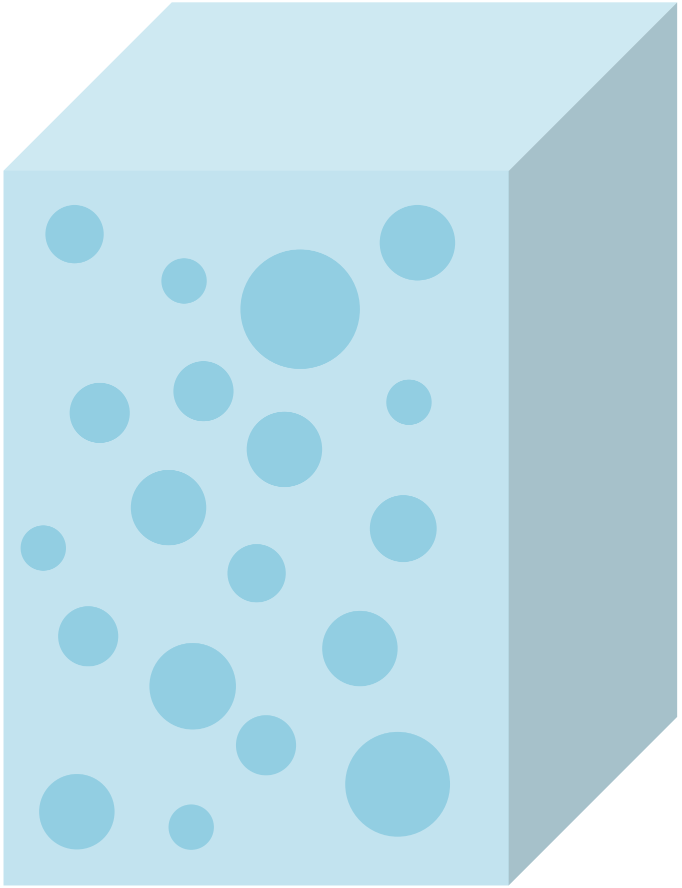

How to get a random matrix in optics?

Thanks to multiple scattering

# Examples

Fog

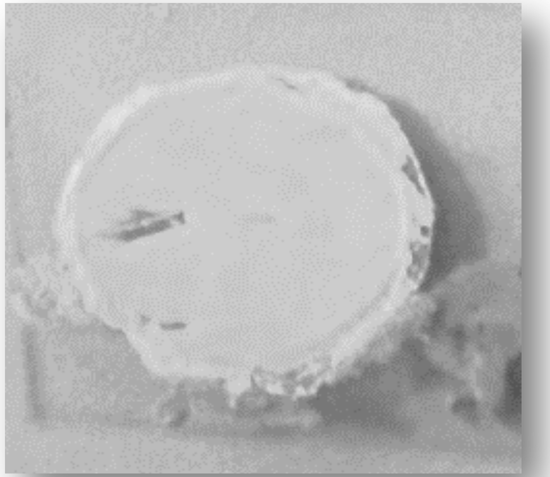

White paint

Biological tissue

# Light scattering

$E^{in}$

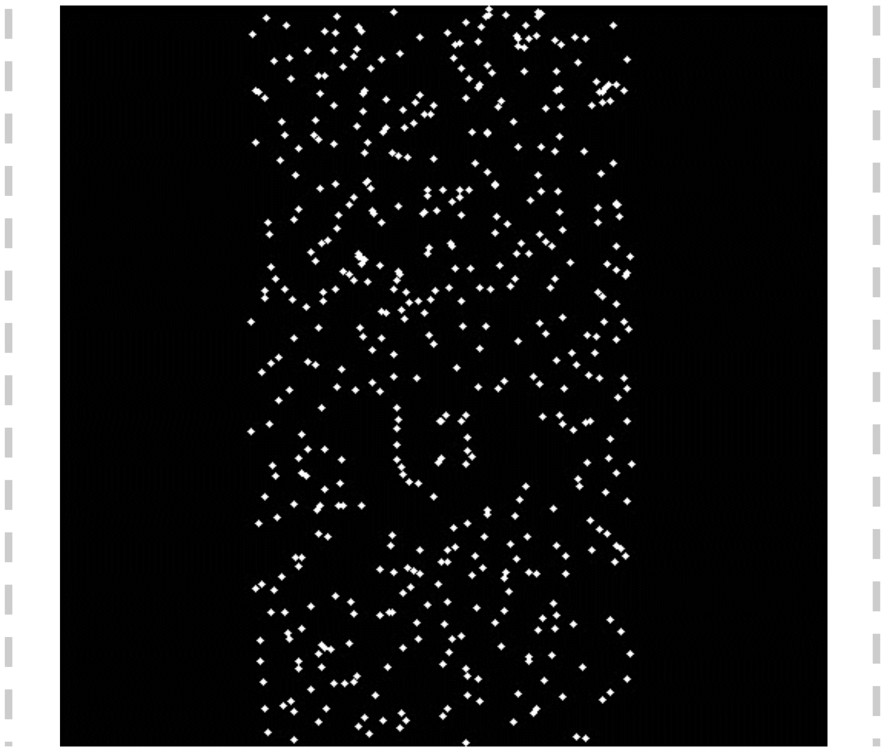

$E^{out}$

Credits: E. Bossy
(LiPHY Grenoble)

Random interference
→ speckle pattern

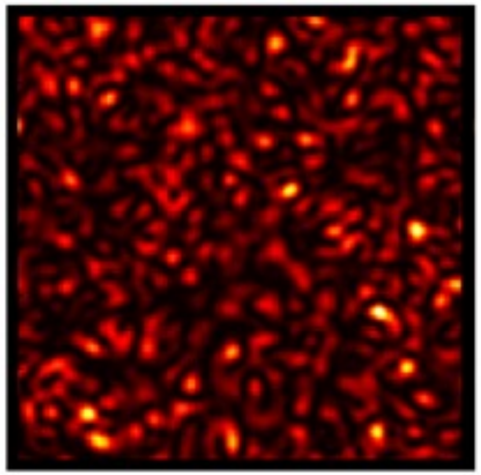

Still linear!

$$E^{out} = AE^{in}$$

(assuming monochromatic coherent light)

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

# Structured-random example

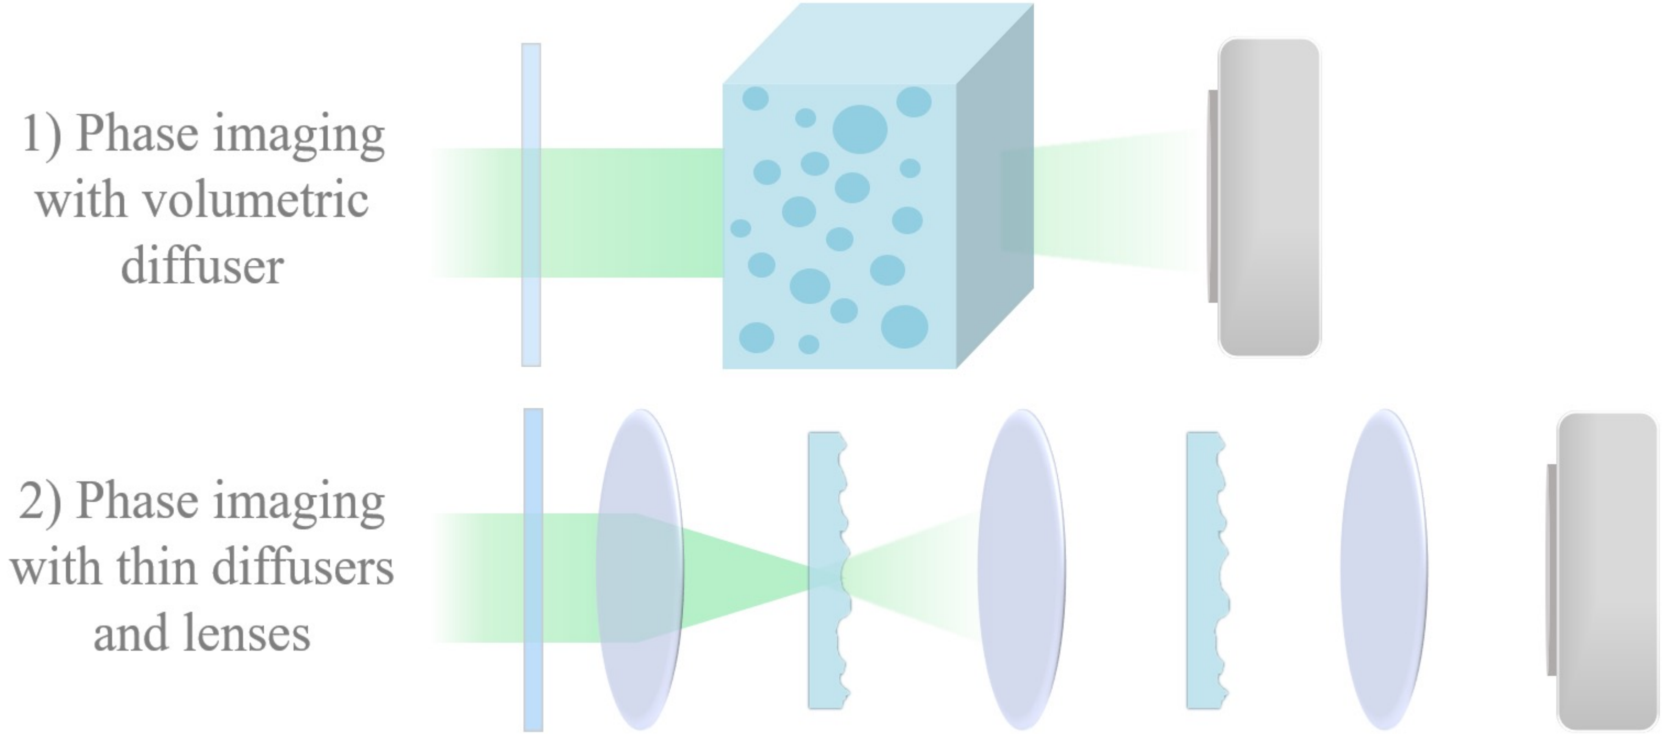


Lens = Fourier transform
Diffuser = Multiplication by diagonal matrix

Final operator: $A = FD_1FD_2F$

# Reconstruction results for $FDFD$

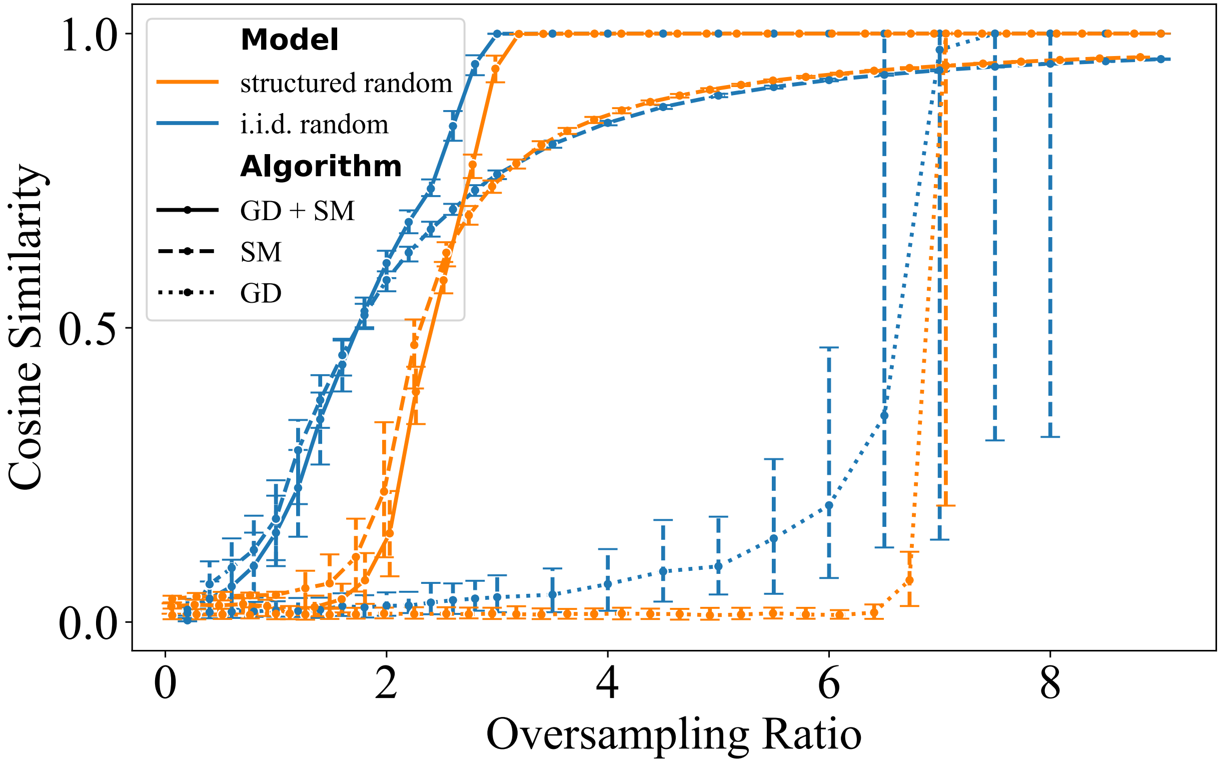


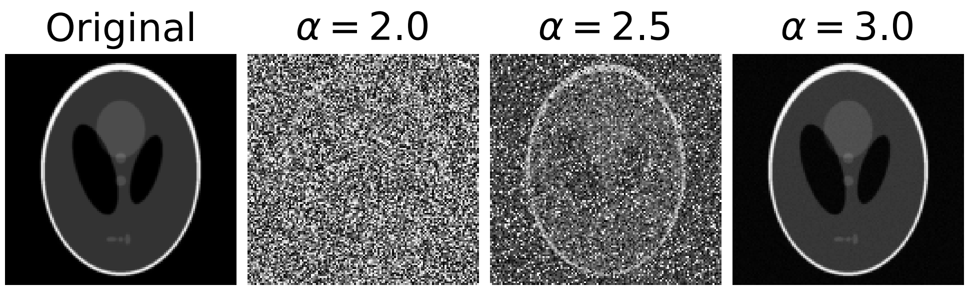


Hu, Tachella, Unser, Dong (2024). *ICASSP*

# Additional results

## How many layers?

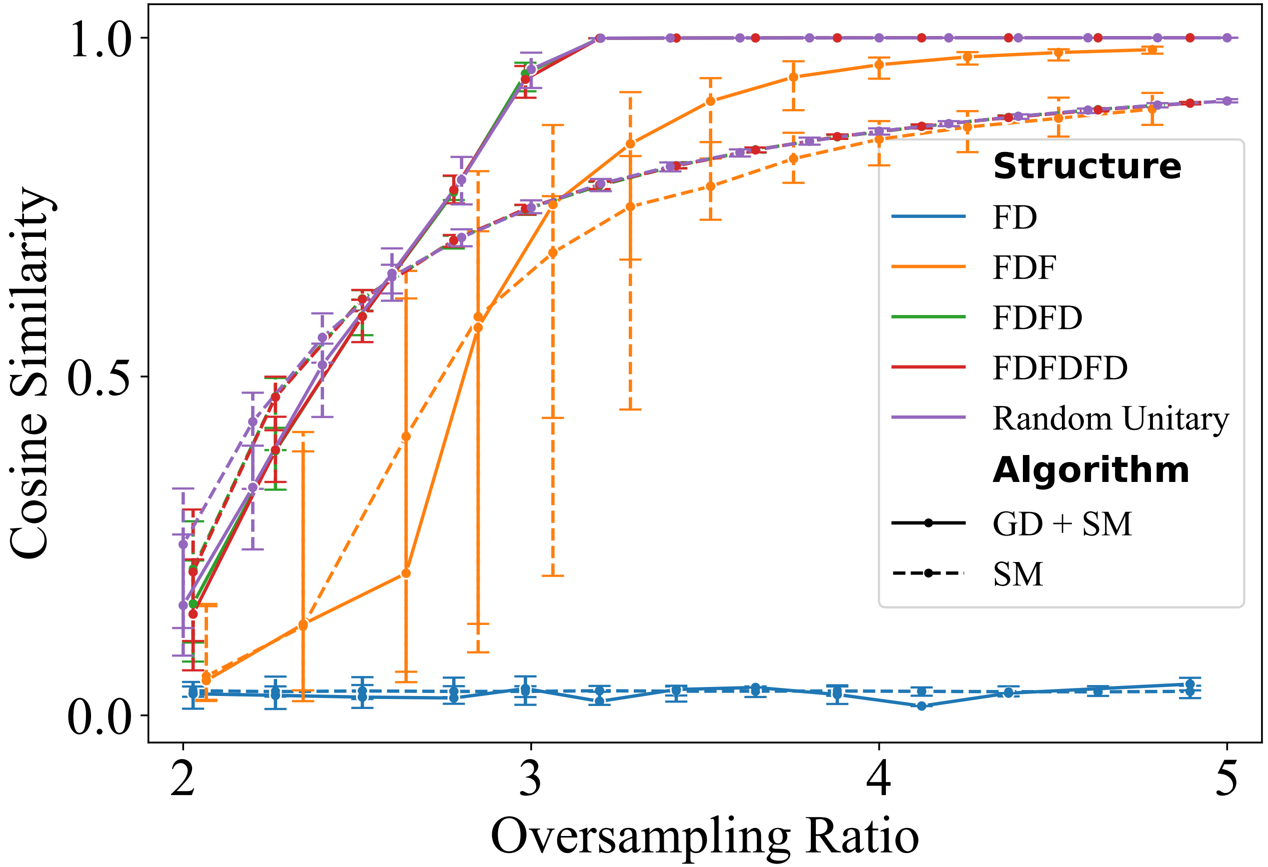


## Speed benchmark

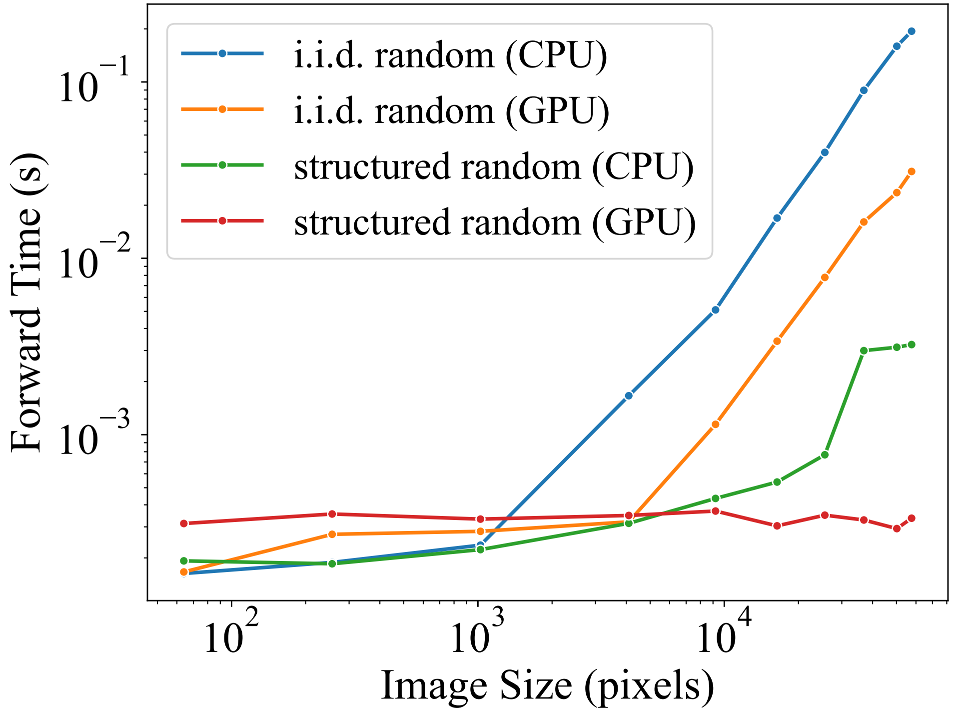

# Content

Find $\boldsymbol{x}^*$ in

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

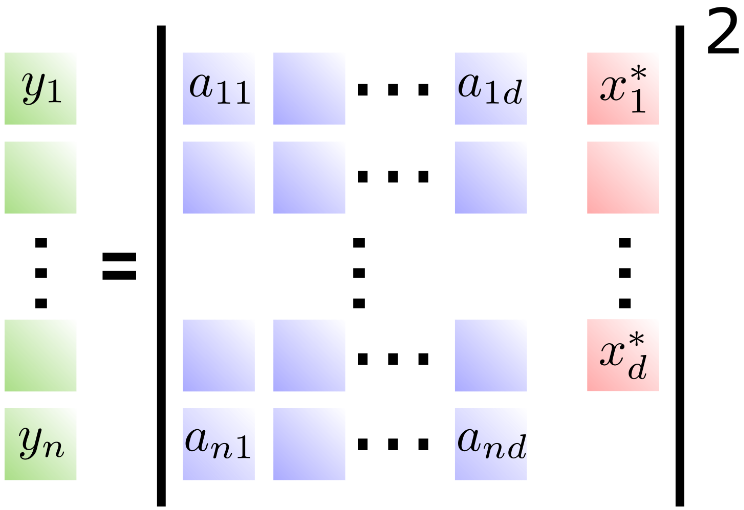


- Physics: Phase imaging devices
- Math:
  - Inverse problems
  - Phase retrieval algorithms
  - Phase retrieval theory
- Machine learning: Regularization

# How to use machine learning for inverse problems?

# Direct inversion

Find $\boldsymbol{x}^*$ in

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

- Train a network $f$ to return $\hat{x} = f(y)$

  Kappeler et al, *IEEE ICIP 2017*

  Rivenson et al, *Light Sci. Appl.* 2018
- Application-specific
  Needs to be retrained for any new microscope
- No robustness guarantee

# Regularization

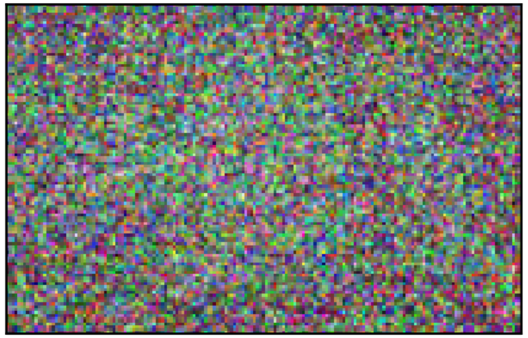

Add information about typical solutions
to help reconstruction

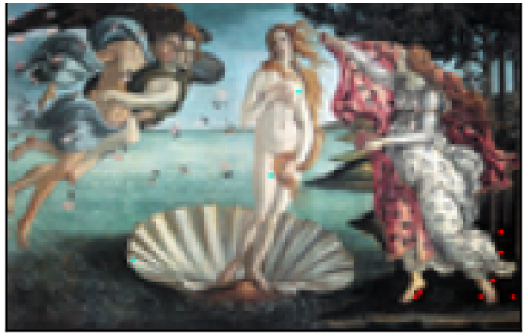

# Deep learning regularization

- Step 1: Train a neural network $f$ for denoising

  *Learn what is a realistic image*

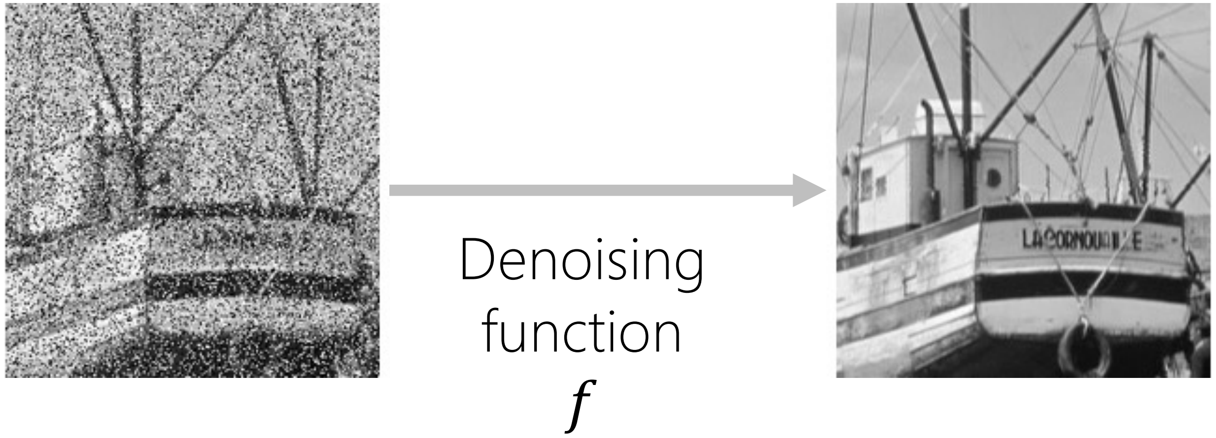


Metzler, Schniter, Veeraraghavan, Baraniuk (2018). *ICML*

# Deep learning regularization

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

*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

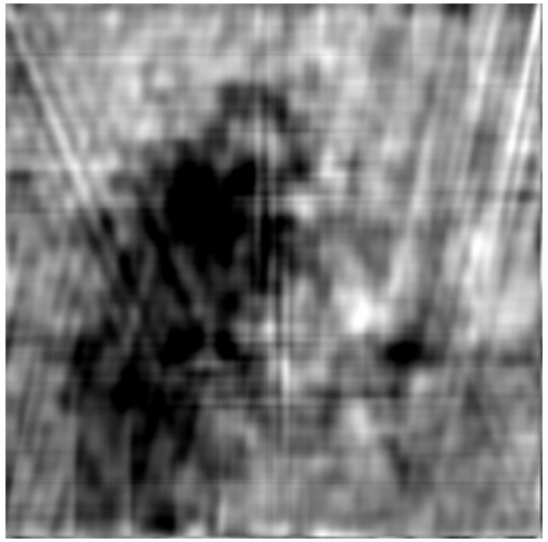

(b) WF (63 sec)

Without regularization

(d) SPAR (294 sec)

Classical regularization

(h) prDeep (345 sec)

Deep learning regularization

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

# Deep learning regularization on the random setting

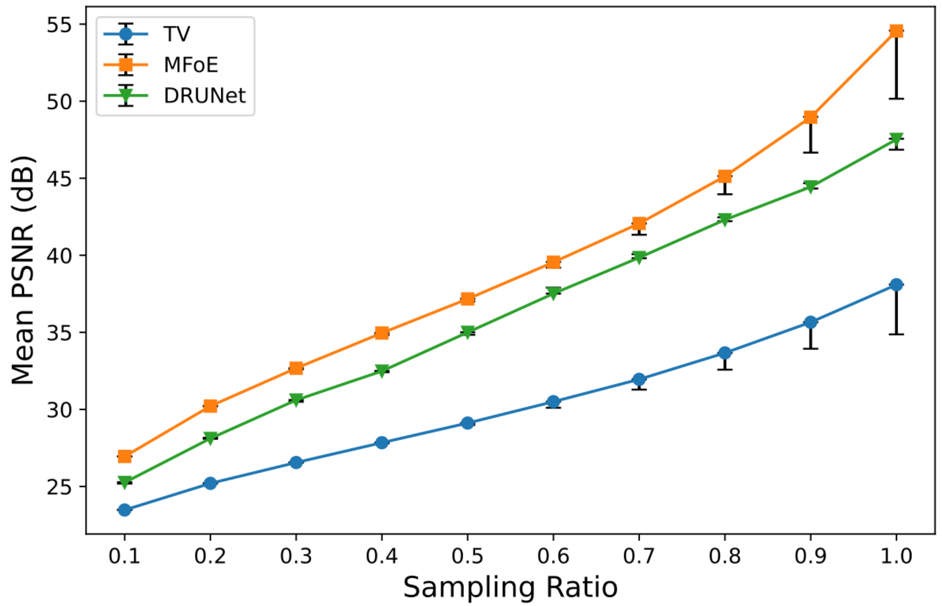


**Fig. 1**. PSNR (dB) for each model and sampling ratio without noise. Error bars show best, median, and worst performance over three random realizations on BSD68.

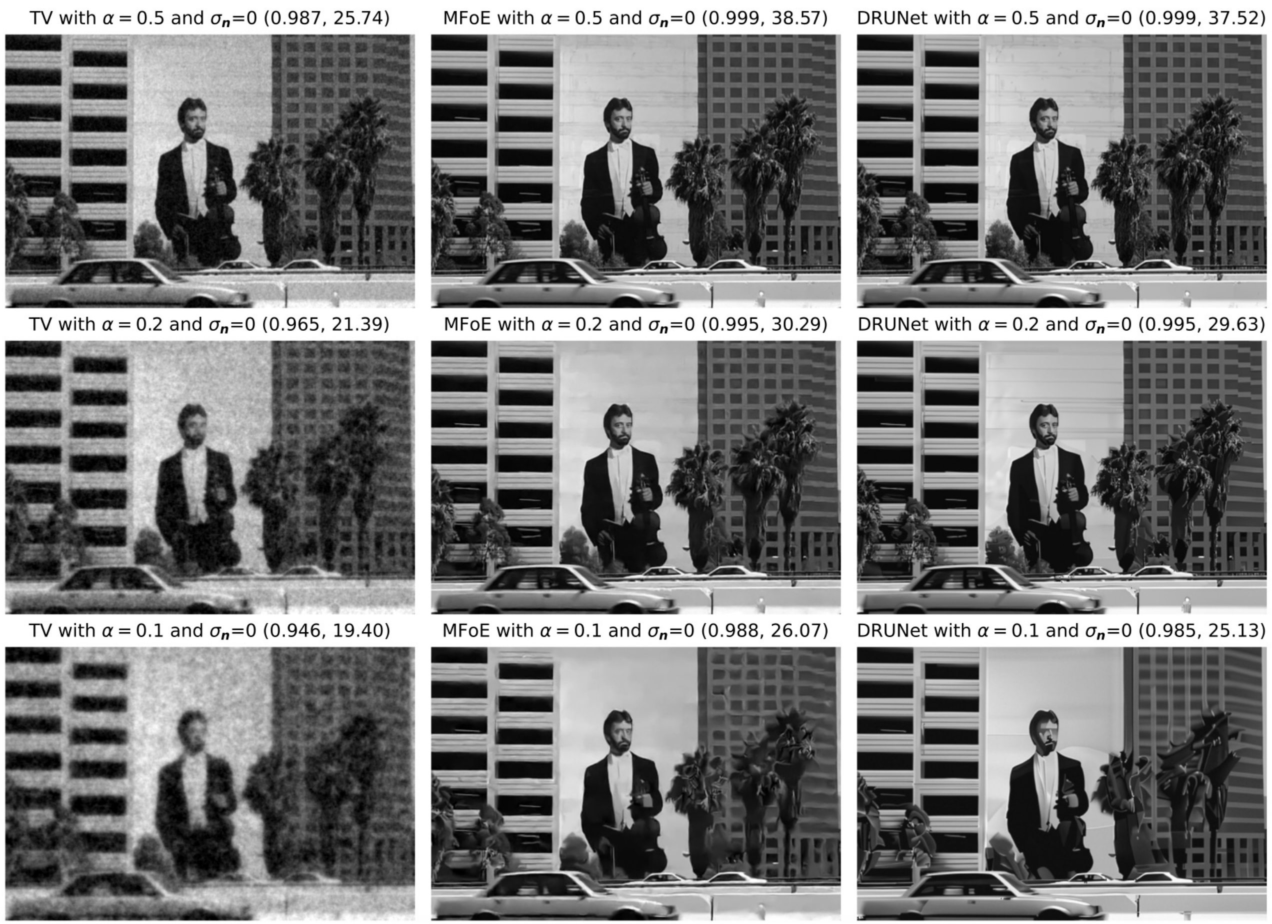


Ducotterd, Hu, Unser, Dong (2025)

# Deep learning regularization

| Favor realistic images | Restrict the search space |
| --- | --- |
| Regularization by denoising<br>Metzler, Schniter, Veeraraghavan, Baraniuk (2018). *ICML* | Generative models<br>Hand, Leong, Voroninski (2018). *NeurIPS* |
| Plug-and-play priors<br>Chang, Bian, Zhang (2021). *eLight* | Deep Image Prior<br>Wang *et al* (2021). *Light: Science & Applications* |

# Generative adversarial networks

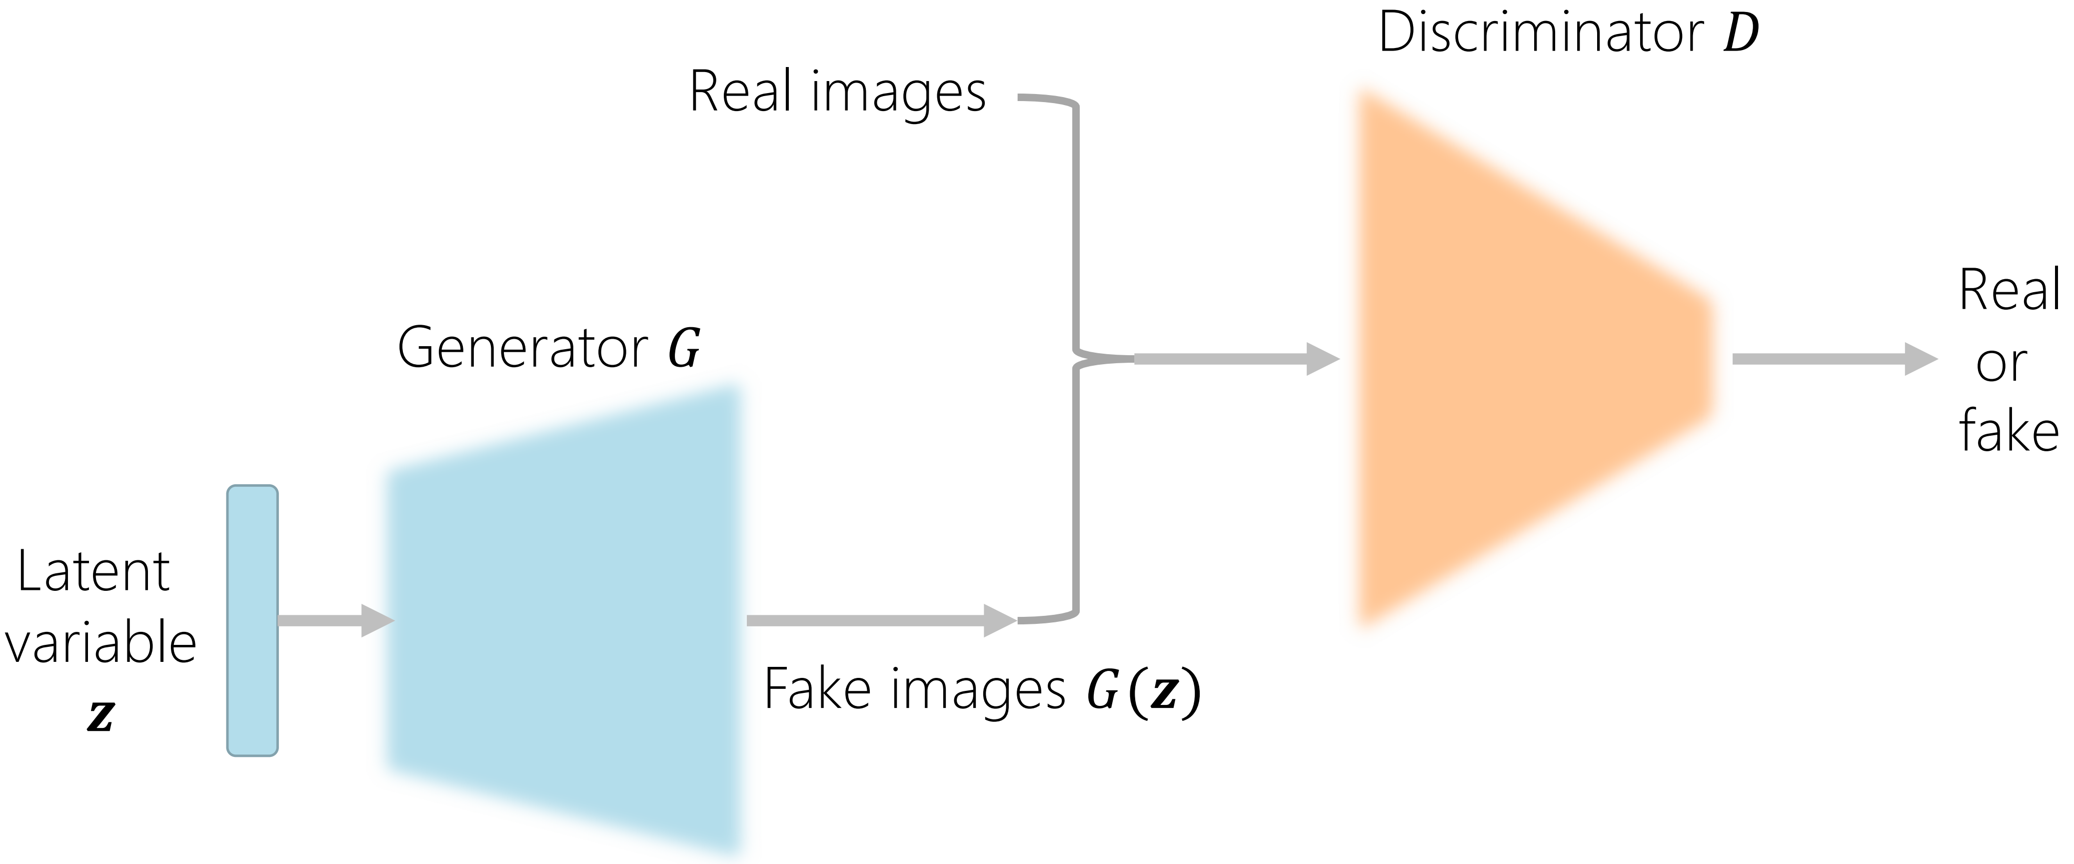


Goodfellow *et al* (2014). *NeurIPS*

# Generative adversarial networks

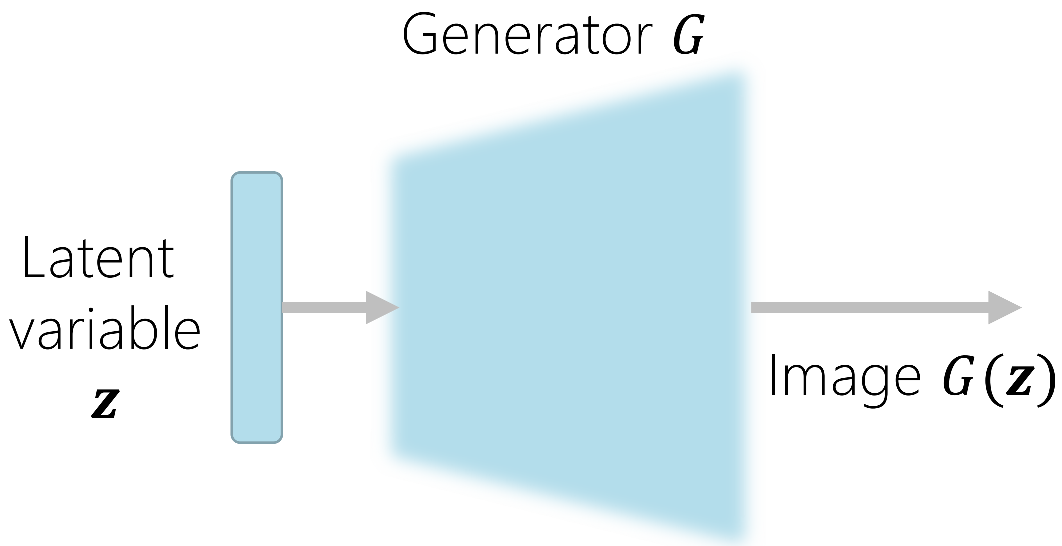


- The generator has learned the *distribution of images*
- Restrict the search space to the generator output

⇒ Gradient descent on $z$ directly

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

# Generative adversarial networks

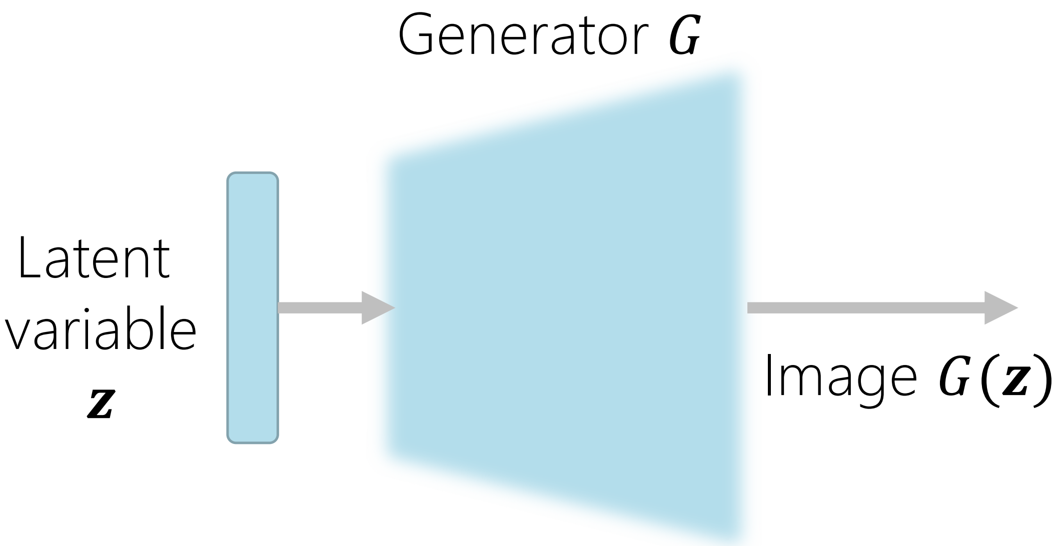


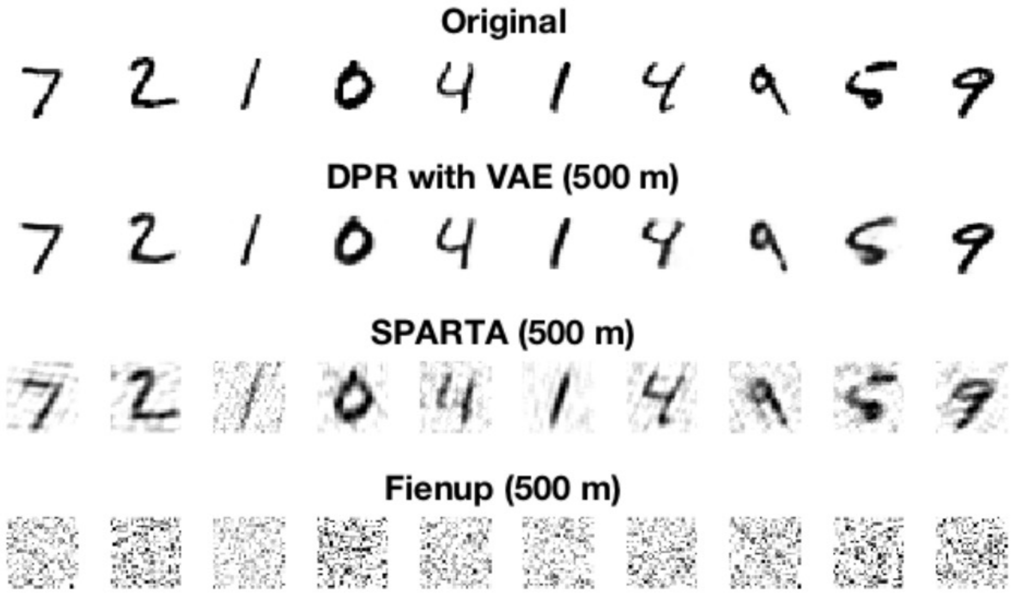


Hand, Leong, Voroninski (2018). *NeurIPS*

# Limits of deep learning

Original phantom

Deep learning reconstruction

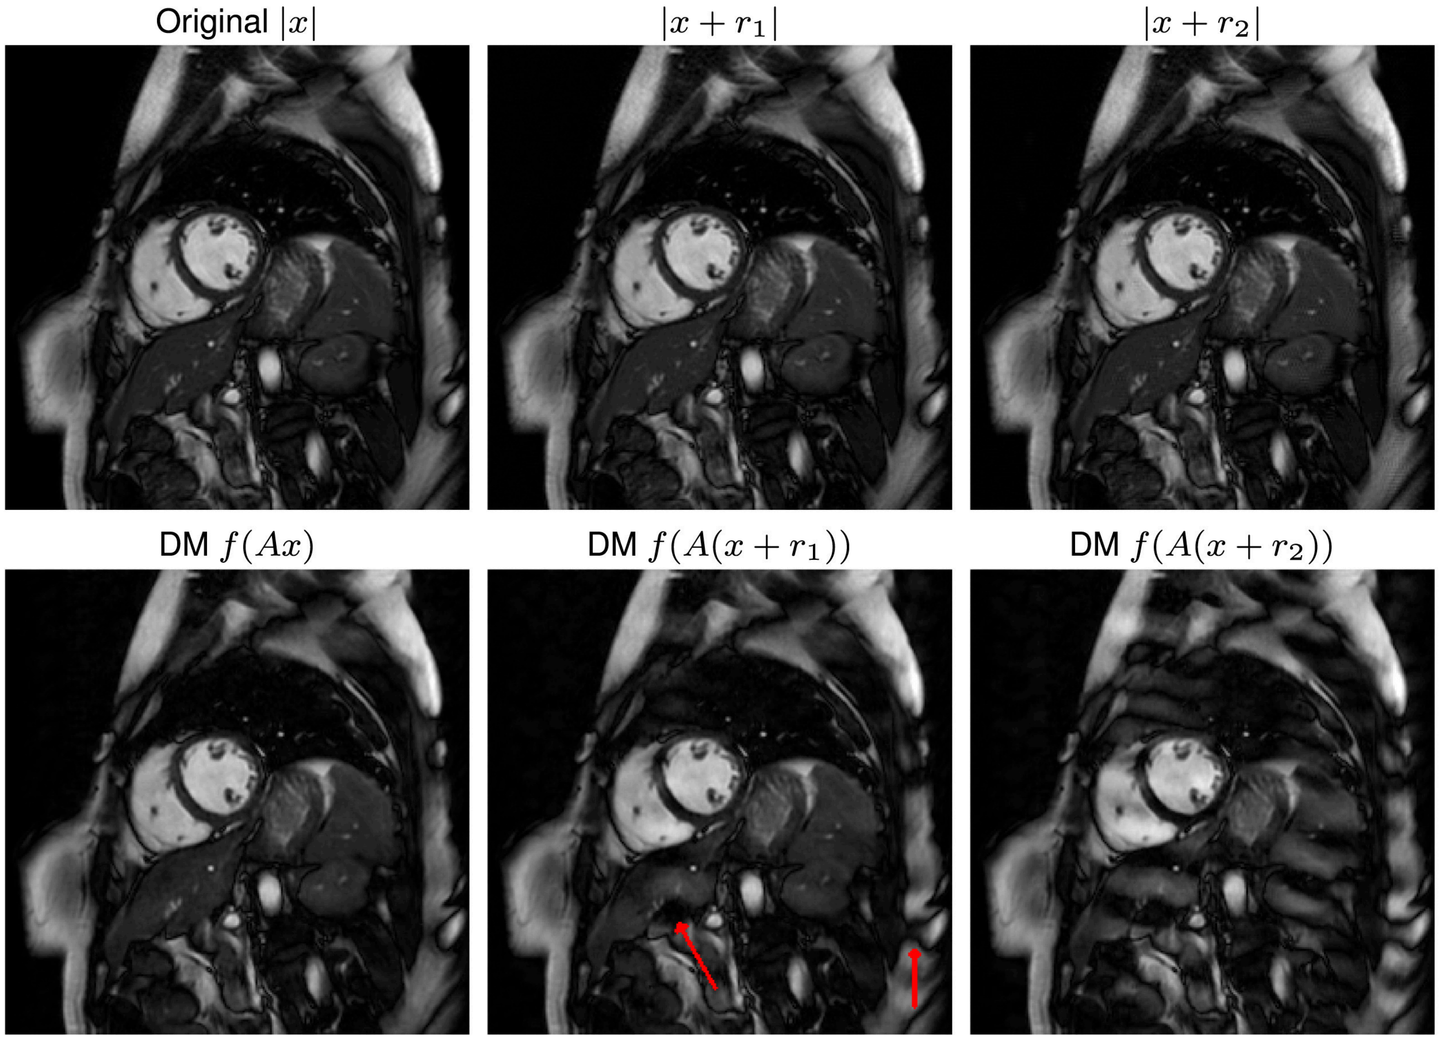


Antun *et al* (2020). PNAS

# 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

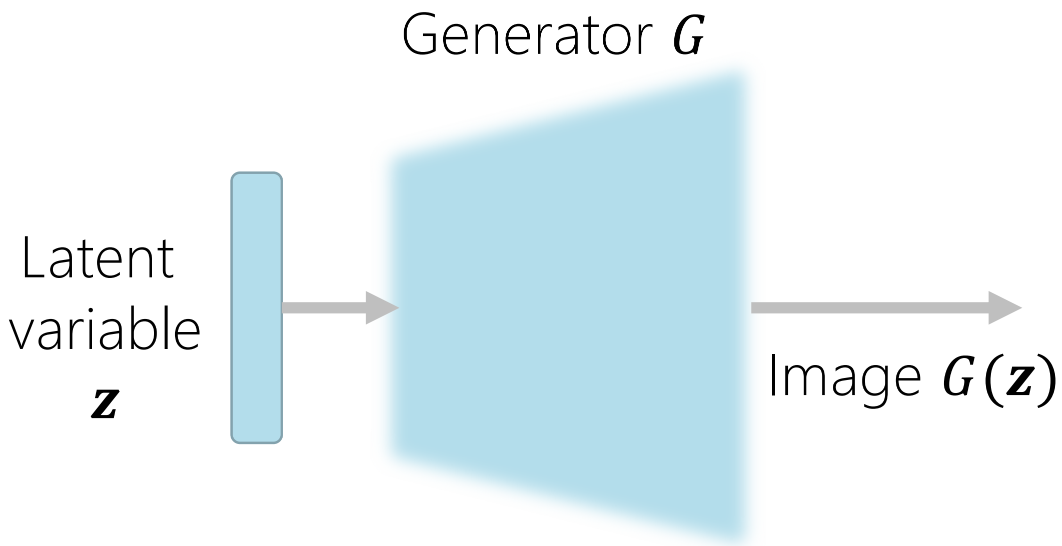


- The generator has learned the *distribution of images*
- Restrict the search space to the generator output

⇒ 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

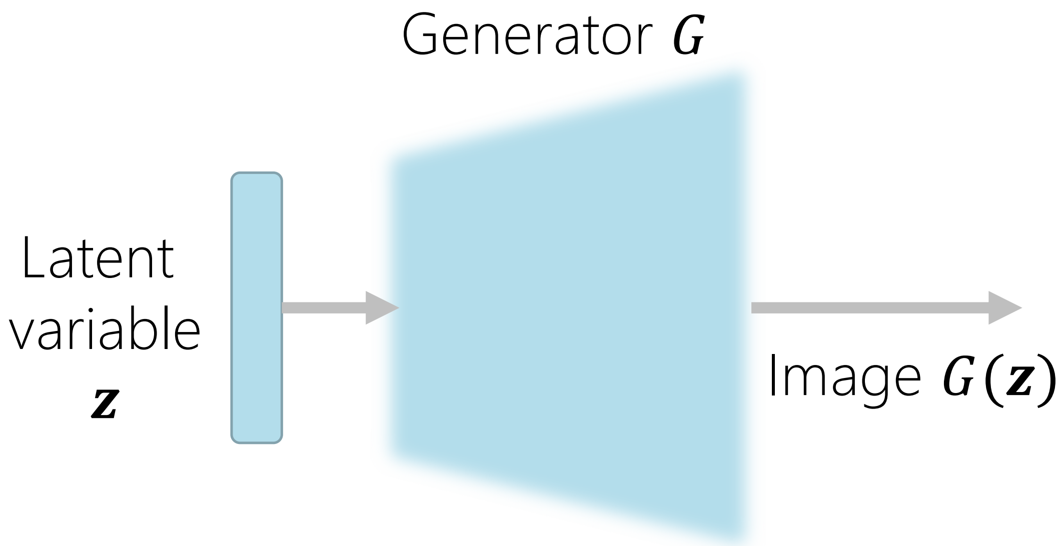


⇒ ~~Gradient descent on $\boldsymbol{z}$ directly (point estimate)~~

Sampling from the posterior distribution

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

Bohra, Pham, Dong, Unser (2022). *IEEE TCI*

# Bayesian GAN

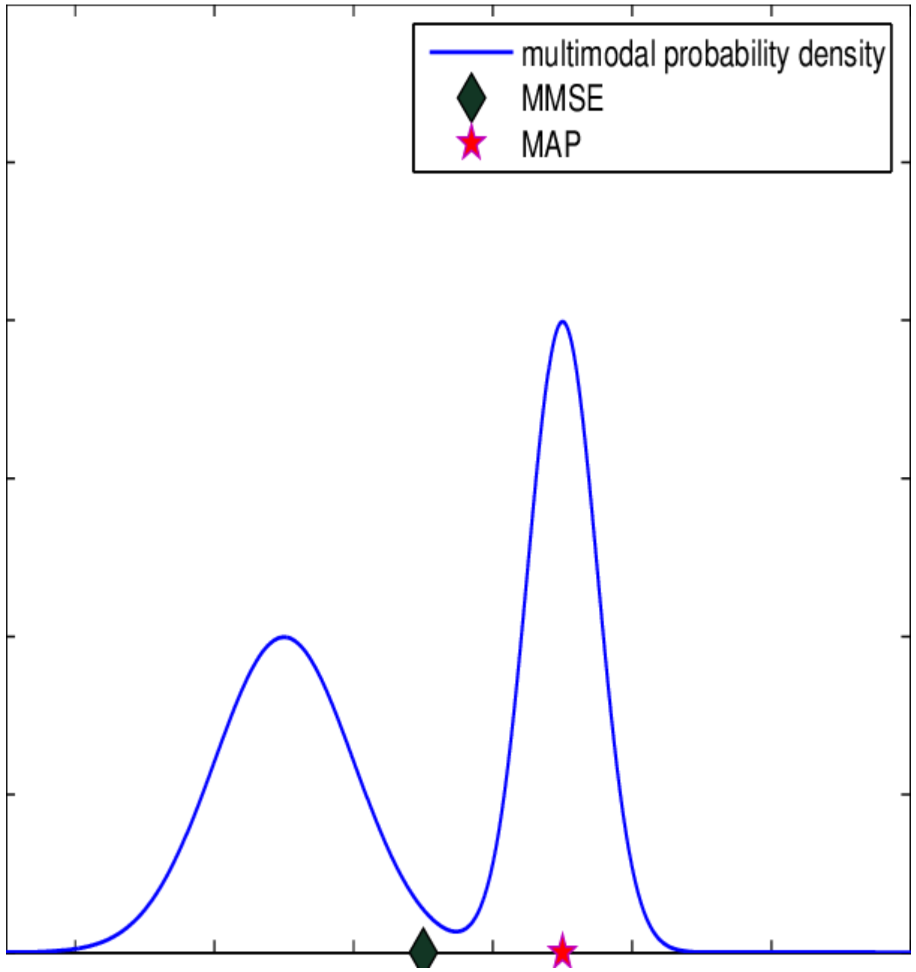


- The generator has learned the *distribution of images*
- Restrict the search space to the generator output

⇒ ~~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). *IEEE TCI*

# Bayesian GAN results

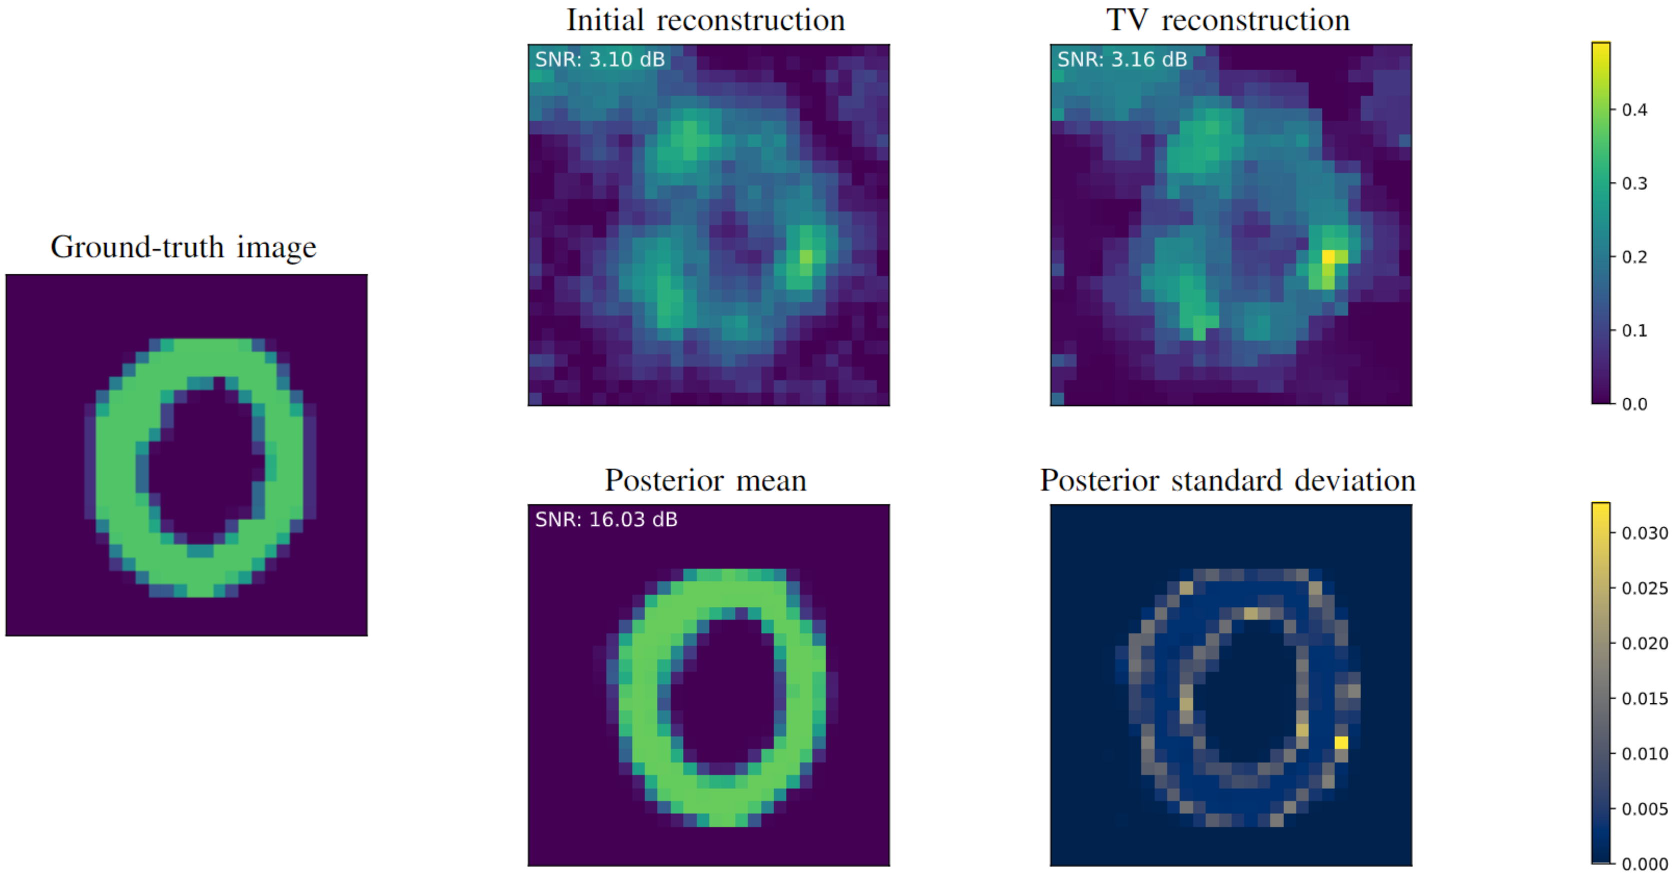


Bohra, Pham, Dong, Unser (2022). *IEEE TCI*

# Bayesian GAN results

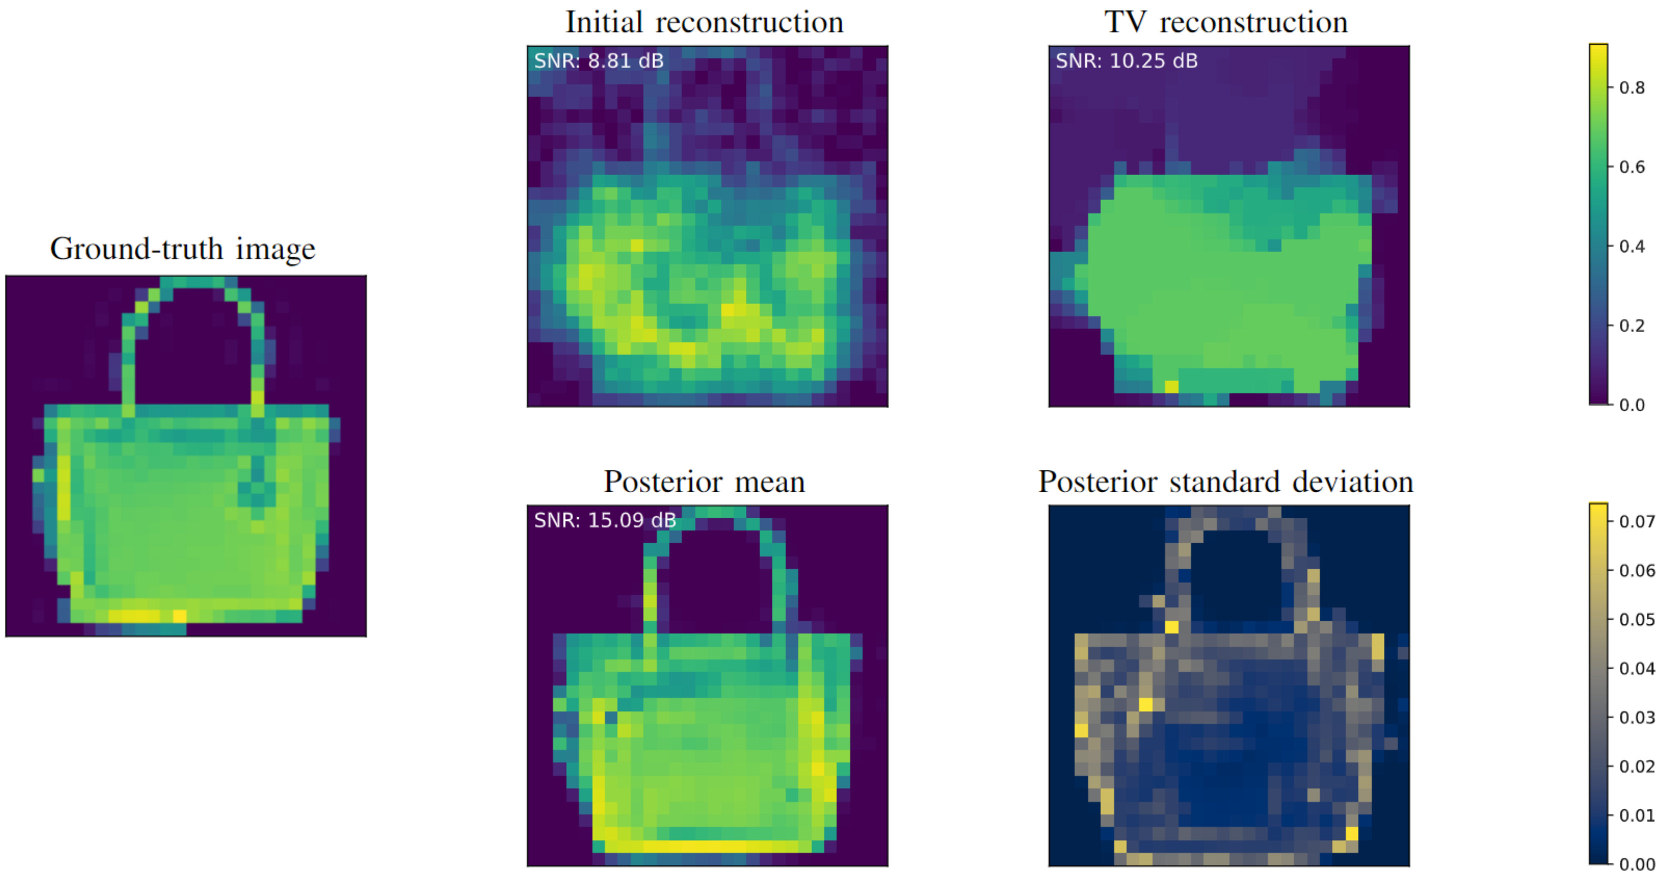


Bohra, Pham, Dong, Unser (2022). *IEEE TCI*

# Open-source libraries

- Deep inverse:
  - Forward models
  - Reconstruction algorithms
  - Pre-trained networks
  - Based on Pytorch
  - https://deepinv.github.io/deepinv/auto_examples/physics/demo_ptychography.html
  - https://deepinv.github.io/deepinv/auto_examples/physics/demo_phase_retrieval.html
- Also:
  - GlobalBioIm: legacy MATLAB library (EPFL)
  - Pyxu: High-performance JAX library (EPFL)
  - Chromatix: JAX optical simulation (Janelia, USA)

Tachella et al (2025)

# 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