Equivariant Imaging: learning to solve inverse problems without ground truth

TéSA Seminar March 2022

Joint work with Dongdong Chen and Mike Davies

Julián Tachella
CNRS
Physics Laboratory
École Normale Supérieure de Lyon

Inverse problems

Definition $$y = Ax + \epsilon$$

where $x \in \mathbb{R}^n$ signal
$y \in \mathbb{R}^m$ observed measurements
$A \in \mathbb{R}^{m \times n}$ ill-posed forward sensing operator ($m \leq n$)
$\epsilon \in \mathbb{R}^m$ is noise

- **Goal:** recover signal x from y
- Inverse problems are **ubiquitous** in science and engineering

Examples

Magnetic resonance imaging

- A = subset of Fourier modes ($k - \text{space}$) of 2D/3D images

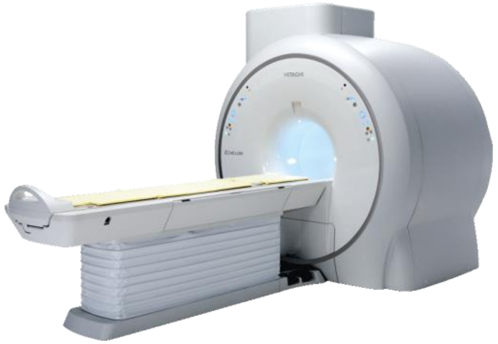

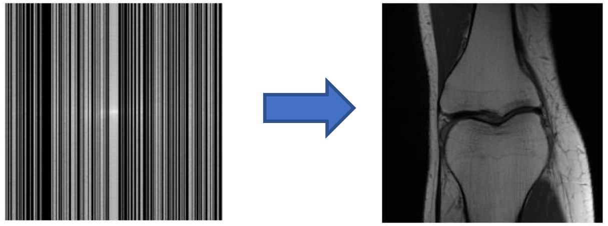

y x

Computed tomography

- A = 2D projections (sinograms) of 2D/3D images

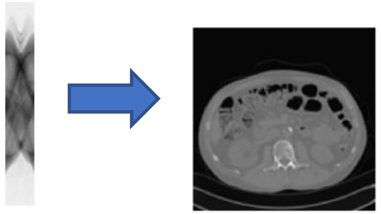

y x

Image inpainting

- A = diagonal matrix with 1's and 0s.

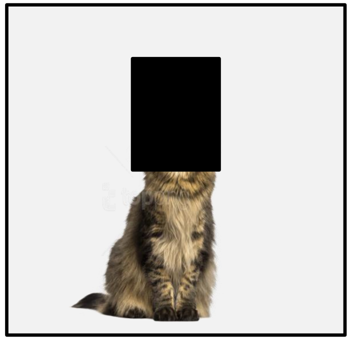

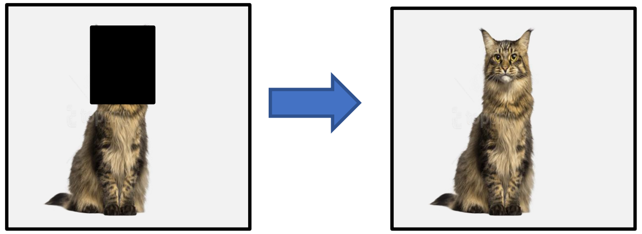

y x

Why it is hard to invert?

Even in the absence of noise, infinitely many $\hat{x}$ consistent with y:

$$\hat{x} = A^{\dagger}y + v$$

where $A^{\dagger}$ is the pseudo-inverse of A and v is any vector in nullspace of A

- *Unique solution only possible if set of plausible x is low-dimensional*

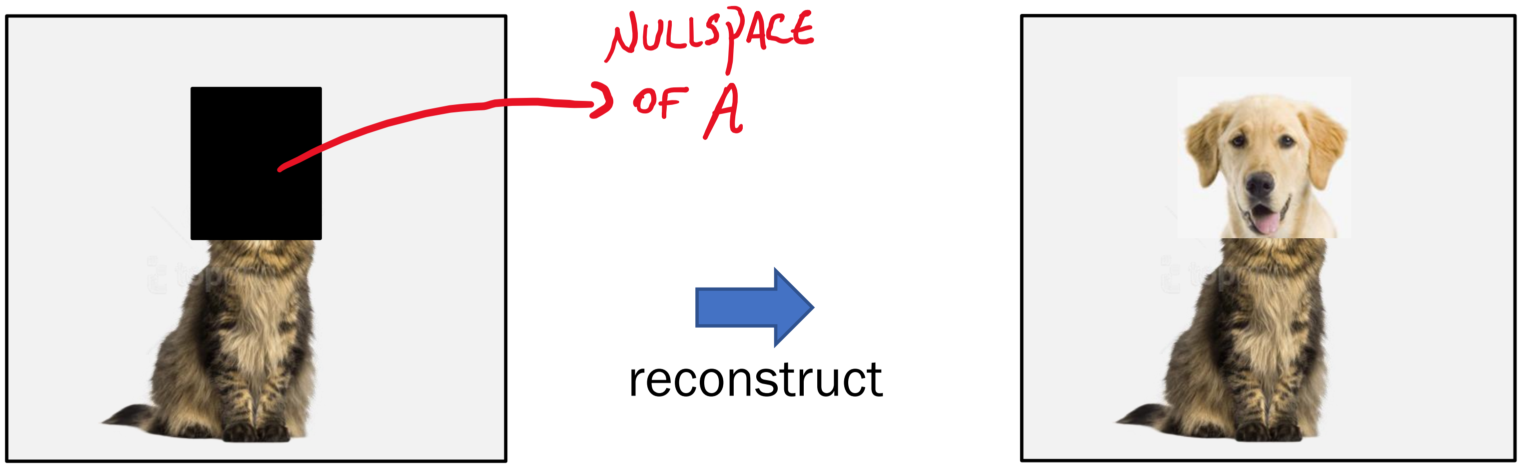

Regularised reconstruction

Idea: define a loss $\rho(x)$ that allows *valid* signals from low-dim set $\mathcal{X}$

$$\hat{x} = \underset{x}{\operatorname{argmin}} \ \left|\left|y - Ax\right|\right|^2 + \rho(x)$$

Examples: total-variation, sparsity, etc.

Disadvantages: hard to define a good $\rho(x)$ in real world problems, loose with respect to the true signal distribution

Learning approach

Idea: use training pairs of signals and measurements (x_i, y_i) to directly learn the inversion function

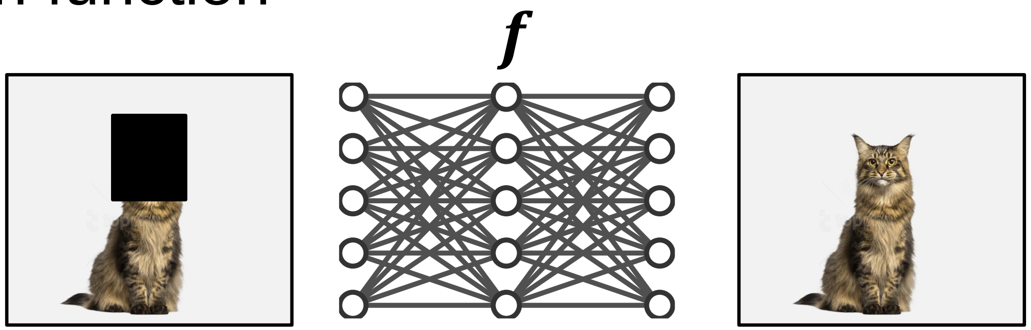

$$\underset{f}{\operatorname{argmin}} \sum_i \left|\left|x_i - f(y_i)\right|\right|^2$$

where $f: \mathbb{R}^m \mapsto \mathbb{R}^n$ is parameterized as a deep neural network.

Learning approach

Advantages:

- State-of-the-art reconstructions
- Once trained, f is easy to evaluate

fastMRI

Accelerating MR Imaging with AI

Ground-truth | Total variation (28.2 dB) | Deep network (34.5 dB)

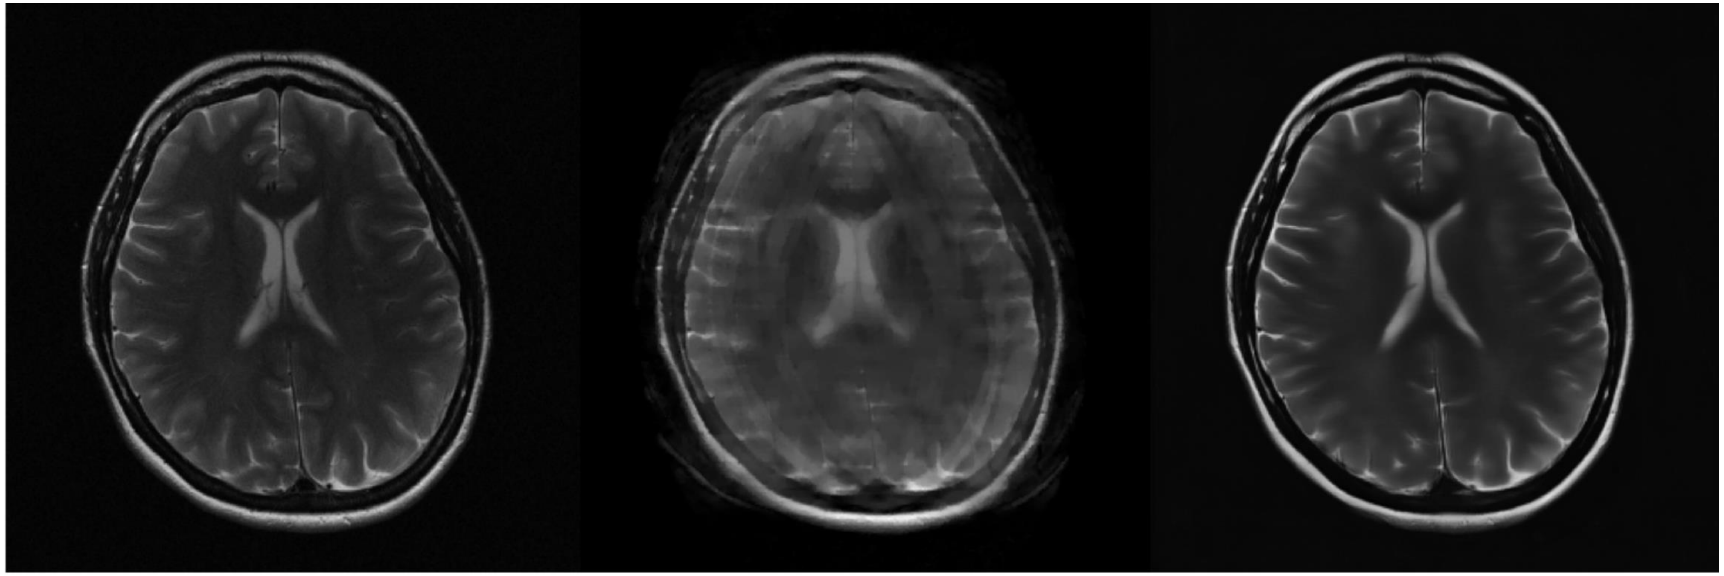

x8 accelerated MRI [Zbontar et al., 2019]

Learning approach

Main disadvantage: Obtaining training signals x_i can be expensive or impossible.

- Medical and scientific imaging
- Only solves inverse problems which we already know what to expect
- Risk of training with signals from a different distribution

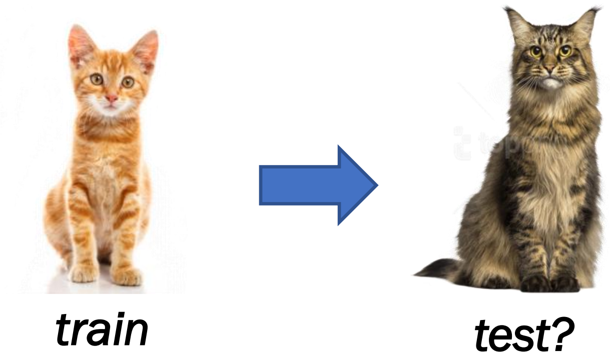

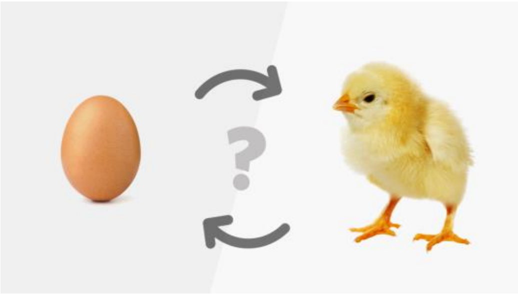

Learning approach

Learning from only measurements y?

$$\underset{f}{\operatorname{argmin}} \sum_i ||y_i - Af(y_i)||^2$$

Proposition: Any reconstruction function $f(y) = A^{\dagger}y + g(y)$ where $g: \mathbb{R}^m \mapsto \mathcal{N}_A$ is any function whose image belongs to the nullspace of A.

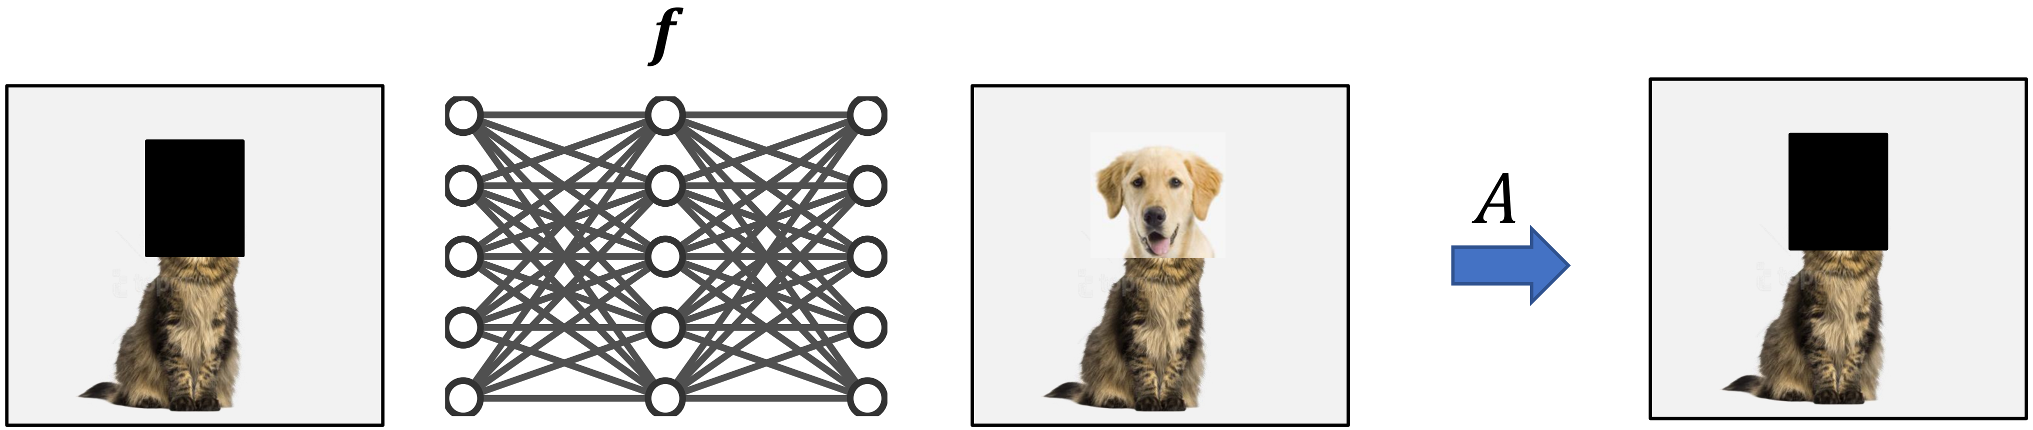

Geometric intuition

Toy example ($n = 3, m = 2$): Signal set is $\mathcal{X} = \mathrm{span}[1,1,1]^{\mathrm{T}}$
Forward operator A keeps first 2 coordinates.

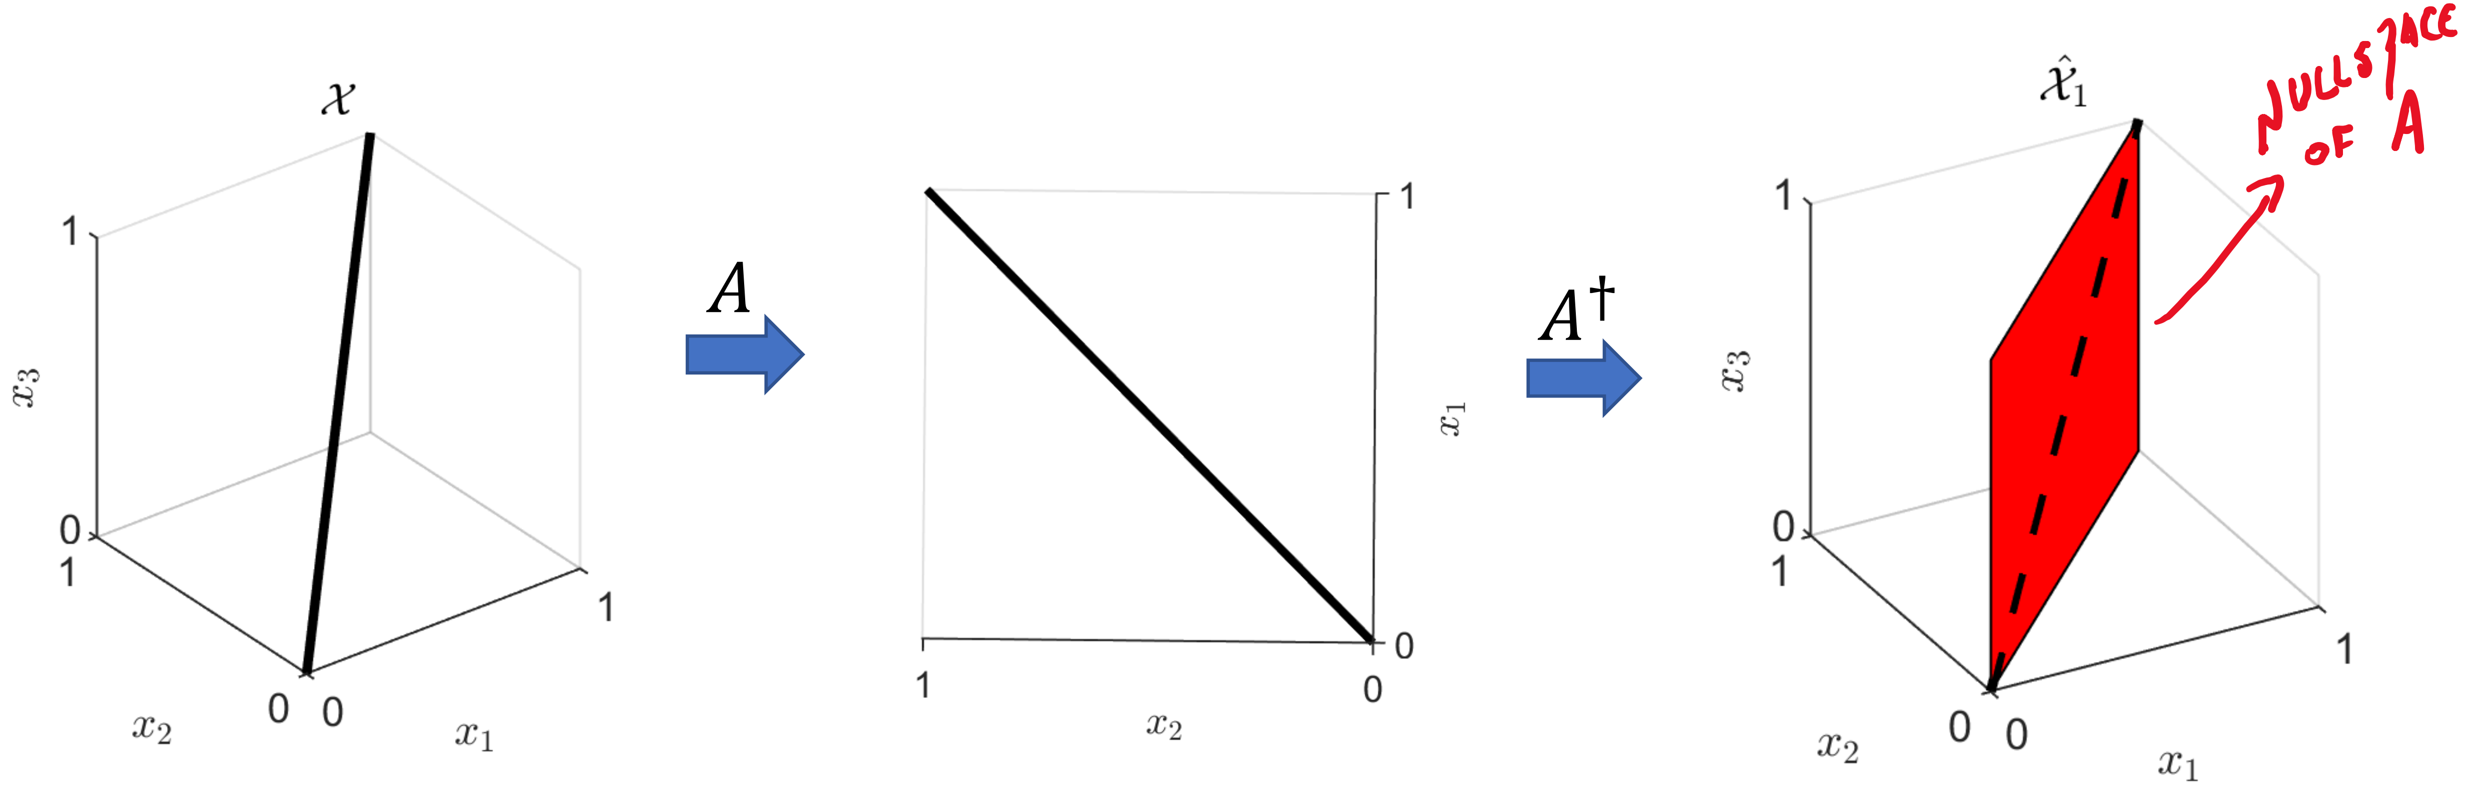

Purpose of this talk

How can we learn reconstruction $f: y \mapsto x$ from noisy measurement data only y_i?

1. Learning with no noise $y = Ax$
2. Learning with noise $y = Ax + \epsilon$

Symmetry prior

How to learn from only y? *We need some prior information*

Idea: Most natural signals distributions are invariant to certain groups of transformations:

$$\forall x \in \mathcal{X},\ \forall g \in G,\ x' = T_g x \in \mathcal{X}$$

$T_g x$ for different $g \in G$

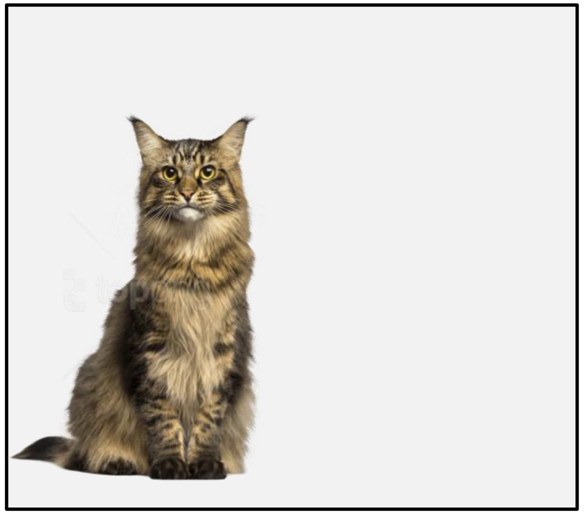

Example: natural images are shift invariant

G = group of 2D shifts

Exploiting invariance

How to learn from only $\boldsymbol{y}$? *We need some prior information*

For all $g \in G$ we have

$$y = Ax = AT_gT_g^{-1}x = A_gx'$$

- Implicit access to multiple operators A_g
- Each operator with different nullspace

AT_gx for different g

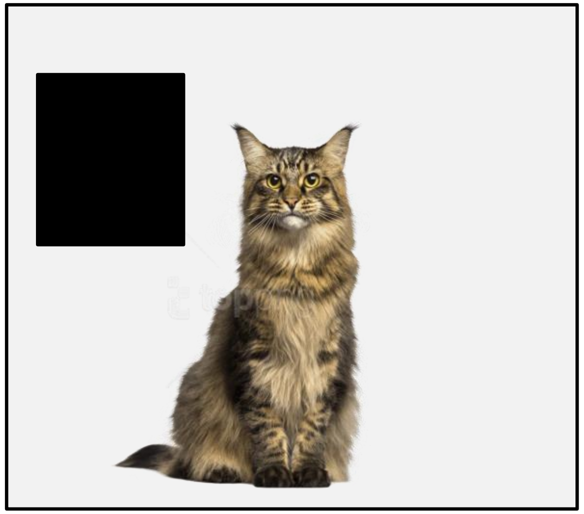

Geometric intuition

Toy example ($n = 3, m = 2$): Signal set is $X = \mathrm{span}[1,1,1]^{\mathrm{T}}$
Forward operator A keeps first 2 coordinates.

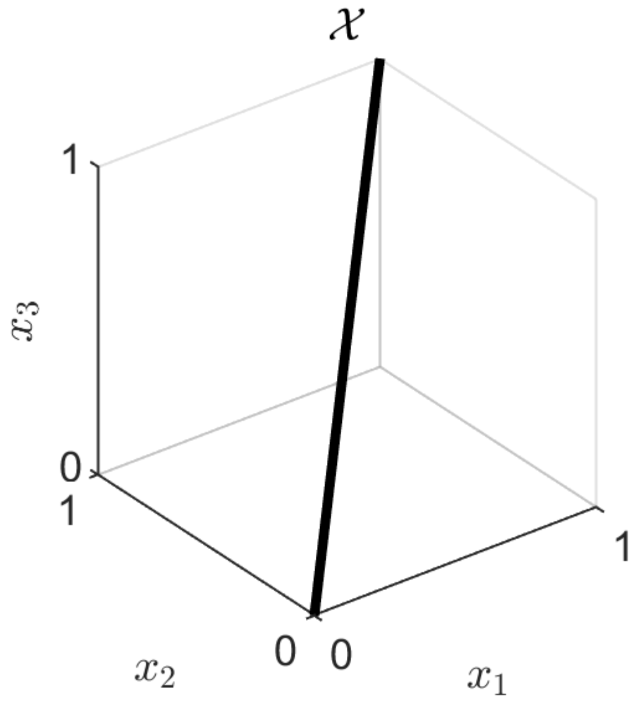

Basic definitions

- Let $G = \{g_1, \dots, g_{|G|}\}$ be a finite group with $|G|$ elements. A **linear representation** of G acting on $\mathbb{R}^n$ is a mapping $T: G \mapsto GL(\mathbb{R}^n)$ which verifies

$$T_{g_1} T_{g_2} = T_{g_1 g_2}$$
$$T_{g_1^{-1}} = T_{g_1}$$

- A mapping $H: \mathbb{R}^n \mapsto \mathbb{R}^m$ is **equivariant** if

$$HT_g = T'_g H \text{ for all } g \in G$$

where $T': G \mapsto GL(\mathbb{R}^m)$ is a linear representation of G acting on $\mathbb{R}^m$.

Necessary conditions

Proposition *[T., Chen and Davies ’22]:*

Learning only possible if $\mathrm{rank}\left(\begin{bmatrix} AT_1 \\ \vdots \\ AT_{|G|} \end{bmatrix}\right) = n$, thus $m \geq n/|G|$

Theorem *[T., Chen and Davies ’22]:* Learning only possible if A is not equivariant to the group action

- If A is equivariant, all AT_g have the same nullspace!

Consequences

Magnetic resonance imaging

- A = subset of Fourier modes ($k-\mathrm{space}$) of 2D/3D images
- Equivariant to shifts
- Not equivariant to rotations, which have $\max c_j \approx \sqrt{n}$

$$m > 2k + \sqrt{n} + 1$$

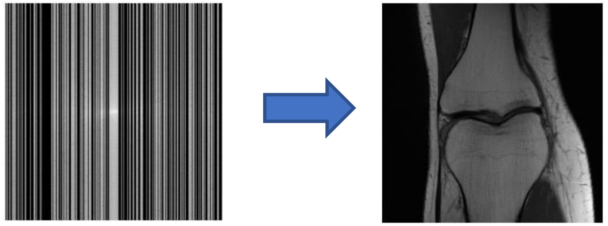

y x

Computed tomography

- A = 2D projections (sinograms) of 2D/3D images
- Equivariant to shifts
- Not equivariant to rotations, which have $\max c_j \approx \sqrt{n}$

$$m > 2k + \sqrt{n} + 1$$

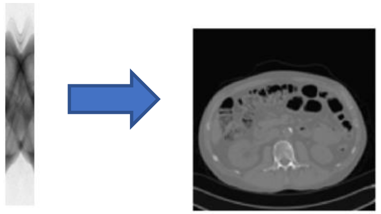

y x

Image inpainting

- A = diagonal matrix with 1's and 0s.
- Not equivariant to shifts, which have $\max c_j \approx 1$

$$m > 2k + 2$$

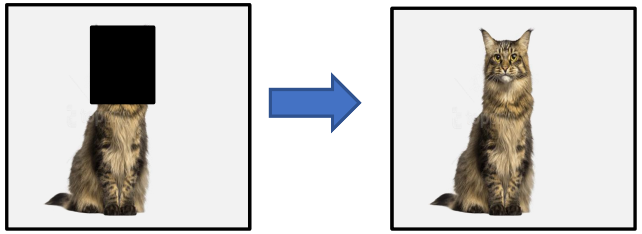

y x

Equivariant imaging loss

How can we enforce invariance in practice?

Idea: we have $f\left(AT_g x\right) = T_g f(Ax)$, i.e. $f \circ A$ should be G-equivariant

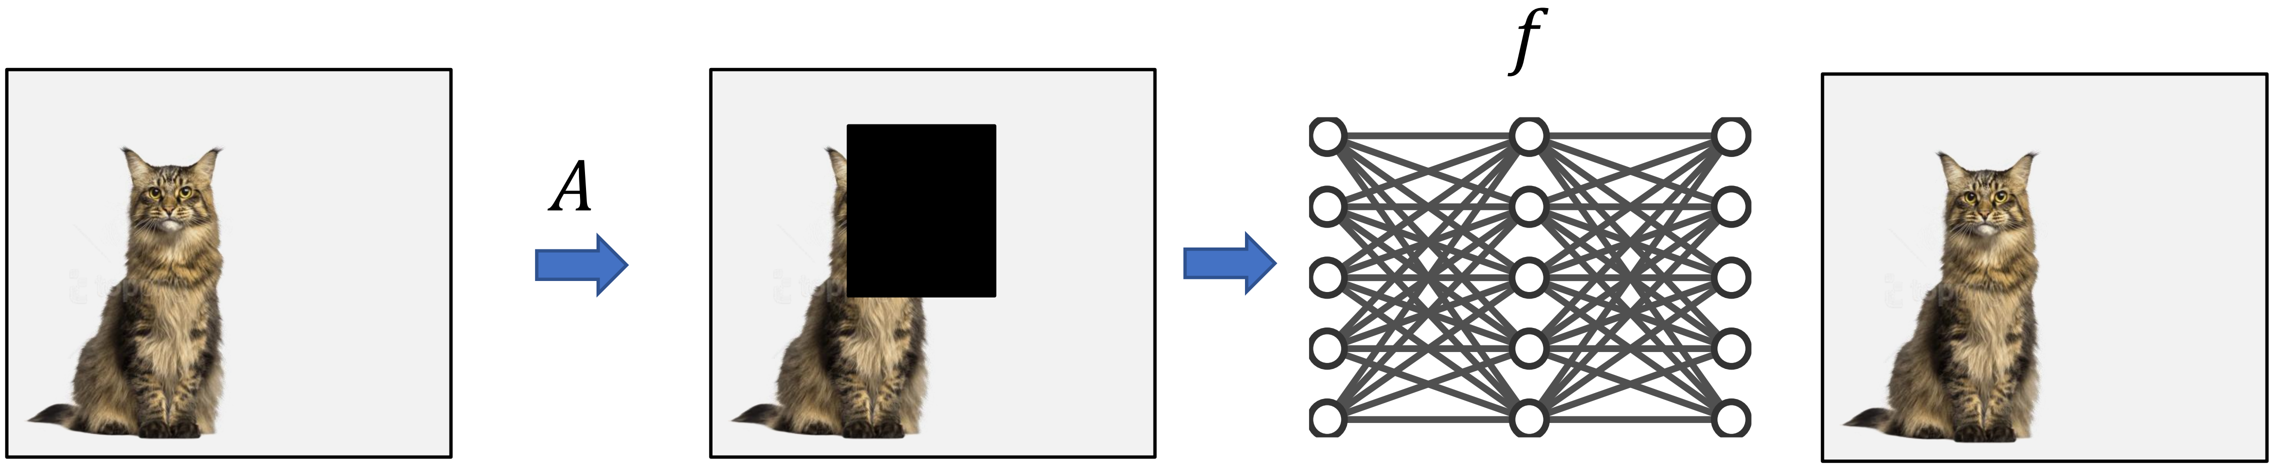

Equivariant imaging

Unsupervised training loss

$$\underset{f}{\operatorname{argmin}}\ \mathcal{L}_{MC}(f) + \mathcal{L}_{EI}(f)$$

- $\mathcal{L}_{MC}(f) = \sum_i \left|\left|y_i - Af(y_i)\right|\right|^2$ measurement consistency
- $\mathcal{L}_{EI}(f) = \sum_{i,g} \left|\left|f\left(AT_g f(y_i)\right) - T_g f(y_i)\right|\right|^2$ enforces equivariance of $A \circ f$

Network-agnostic: applicable to any existing deep model!

Experiments

Tasks:

- Magnetic resonance imaging

Network

- $f = g_\theta \circ A^\dagger$ where g_θ is a U-net CNN

Comparison

- Pseudo-inverse $A^\dagger y_i$ (no training)
- Meas. consistency $Af(y_i) = y_i$
- Fully supervised loss: $f(y_i) = x_i$
- Equivariant imaging (unsupervised) $Af(y_i) = y_i$ and equivariant $A \circ f$

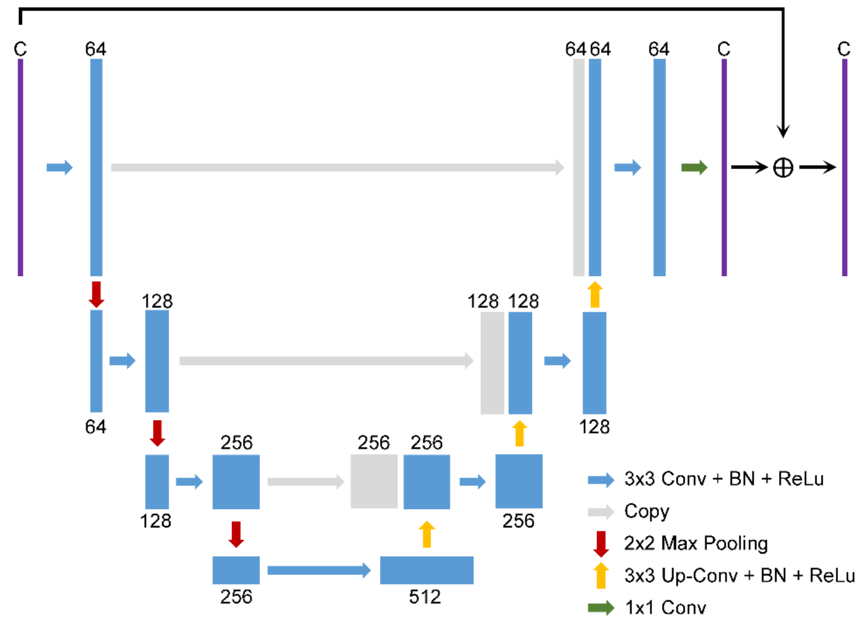

Magnetic resonance imaging

- Operator A is a subset of Fourier measurements (x2 downsampling)
- Dataset is approximately rotation invariant

Signal x

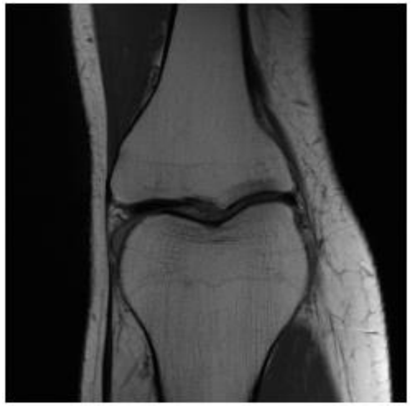

Measurements y

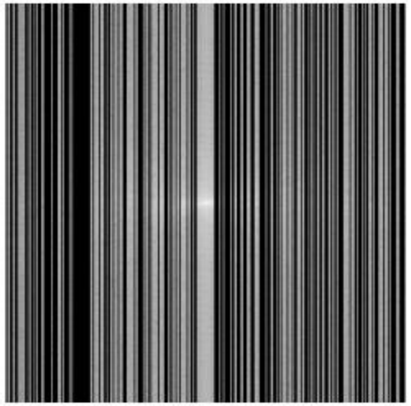

2. Learning from noisy measurements

What about noise?

Noisy measurements

$$y|u \sim q_u(y)$$
$$u = Ax$$

Examples: Gaussian noise, Poisson noise. Poisson-Gaussian noise

MRI with different noise levels:
- EI degrades with noise!

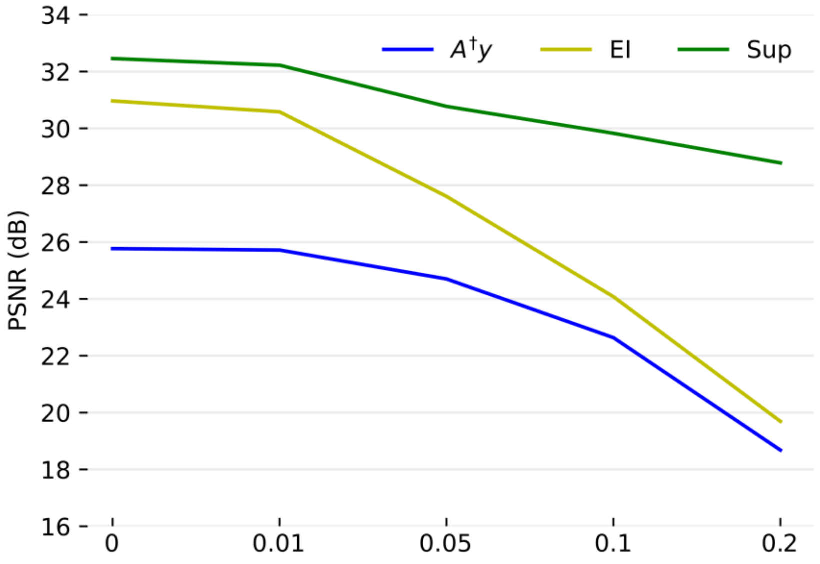

Gaussian noise standard deviation

Handling noise via SURE

Oracle consistency loss with clean/noisy measurements pairs (u_i, y_i)

$$\mathcal{L}_{MC}(f) = \sum_i \left|\left|u_i - Af(y_i)\right|\right|^2$$

However, we don't have clean u_i!

Idea: Proxy unsupervised loss $\mathcal{L}_{SURE}(f)$ which is an unbiased estimator, i.e.

$$\mathbb{E}_{y,u}\{\mathcal{L}_{MC}(f)\} = \mathbb{E}_y\{\mathcal{L}_{SURE}(f)\}$$

Handling noise via SURE

Gaussian noise $y \sim \mathcal{N}(u, I\sigma^2)$

$$\mathcal{L}_{SURE}(f) = \sum_i \left|\left|y_i - Af(y_i)\right|\right|^2 - \sigma^2 m + 2\sigma^2 \mathrm{div}(A \circ f)(y_i)$$

where $\mathrm{div}(h(x)) = \sum_j \frac{\delta h_j}{\delta x_j}$ is approximated with a Monte Carlo estimate which only requires evaluations of h [Ramani, 2008]

Theorem [Stein, 1981] Under mild differentiability conditions on the function $A \circ f$, the following holds

$$\mathbb{E}_{y,u}\{\mathcal{L}_{MC}(f)\} = \mathbb{E}_y\{\mathcal{L}_{SURE}(f)\}$$

Robust EI: SURE+EI

Robust Equivariant Imaging

$$\underset{f}{\operatorname{argmin}}\ \mathcal{L}_{SURE}(f) + \mathcal{L}_{EI}(f)$$

- $\mathcal{L}_{SURE}(f)$: unbiased estimator of oracle measurement consistency
  - noise dependent
  - Gaussian, Poisson, Poisson-Gaussian

- $\mathcal{L}_{EI}(f)$: enforces equivariance of $A \circ f$

Experiments

Tasks:

- Magnetic resonance imaging (Gaussian noise)
- Image inpainting (Poisson noise)
- Computed tomography (Poisson-Gaussian noise)

Network

- $f = g_\theta \circ A^\dagger$ where g_θ is a U-net CNN

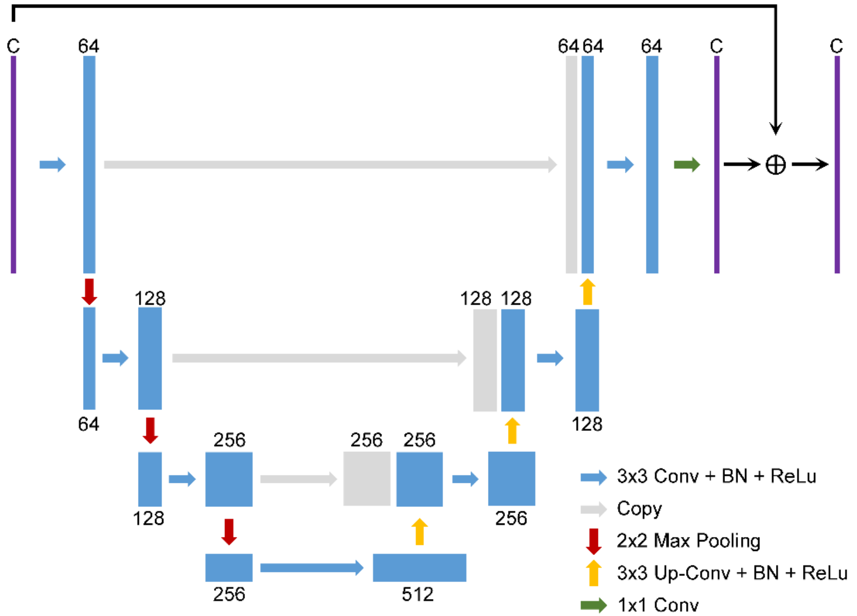

Comparison

- Meas. consistency $Af(y_i) = y_i$
- Fully supervised loss: $f(y_i) = x_i$
- Equivariant imaging (unsupervised)
  $Af(y_i) = y_i$ and equivariant $A \circ f$

Magnetic resonance imaging

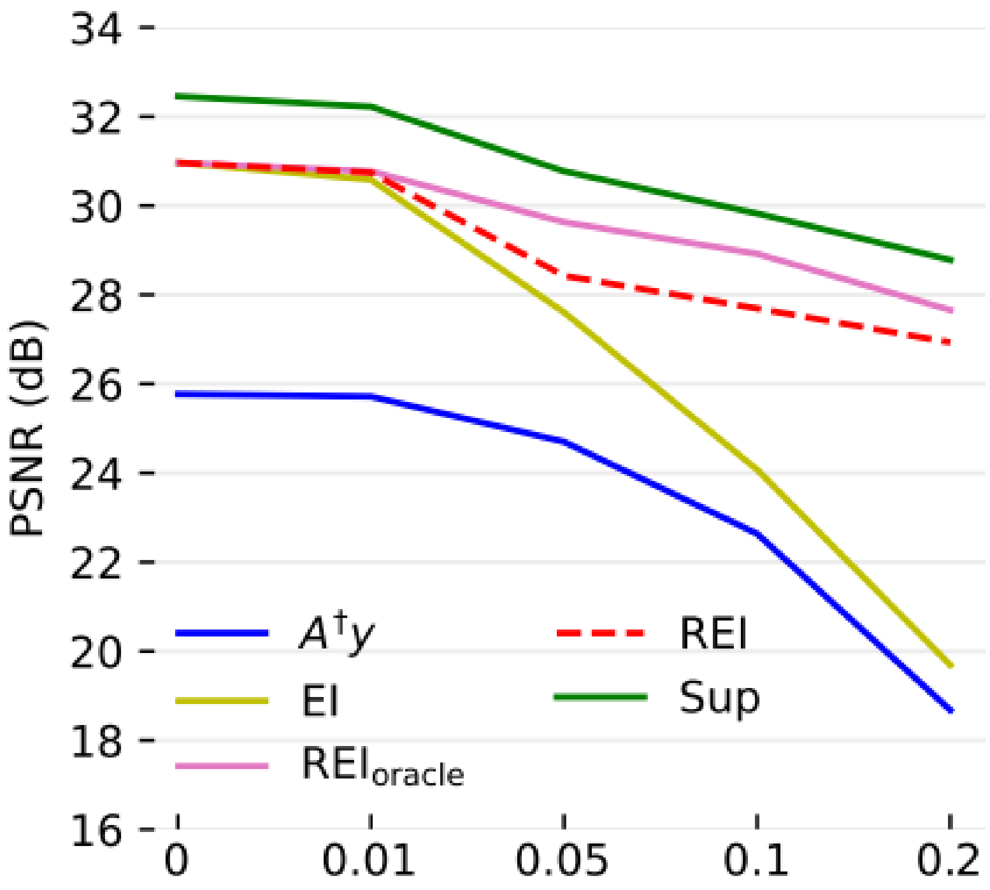

Gaussian noise standard deviation

Magnetic resonance imaging

- Operator A is a subset of Fourier measurements (x4 downsampling)
- Gaussian noise ($\sigma = 0.2$)
- Dataset is approximately rotation invariant

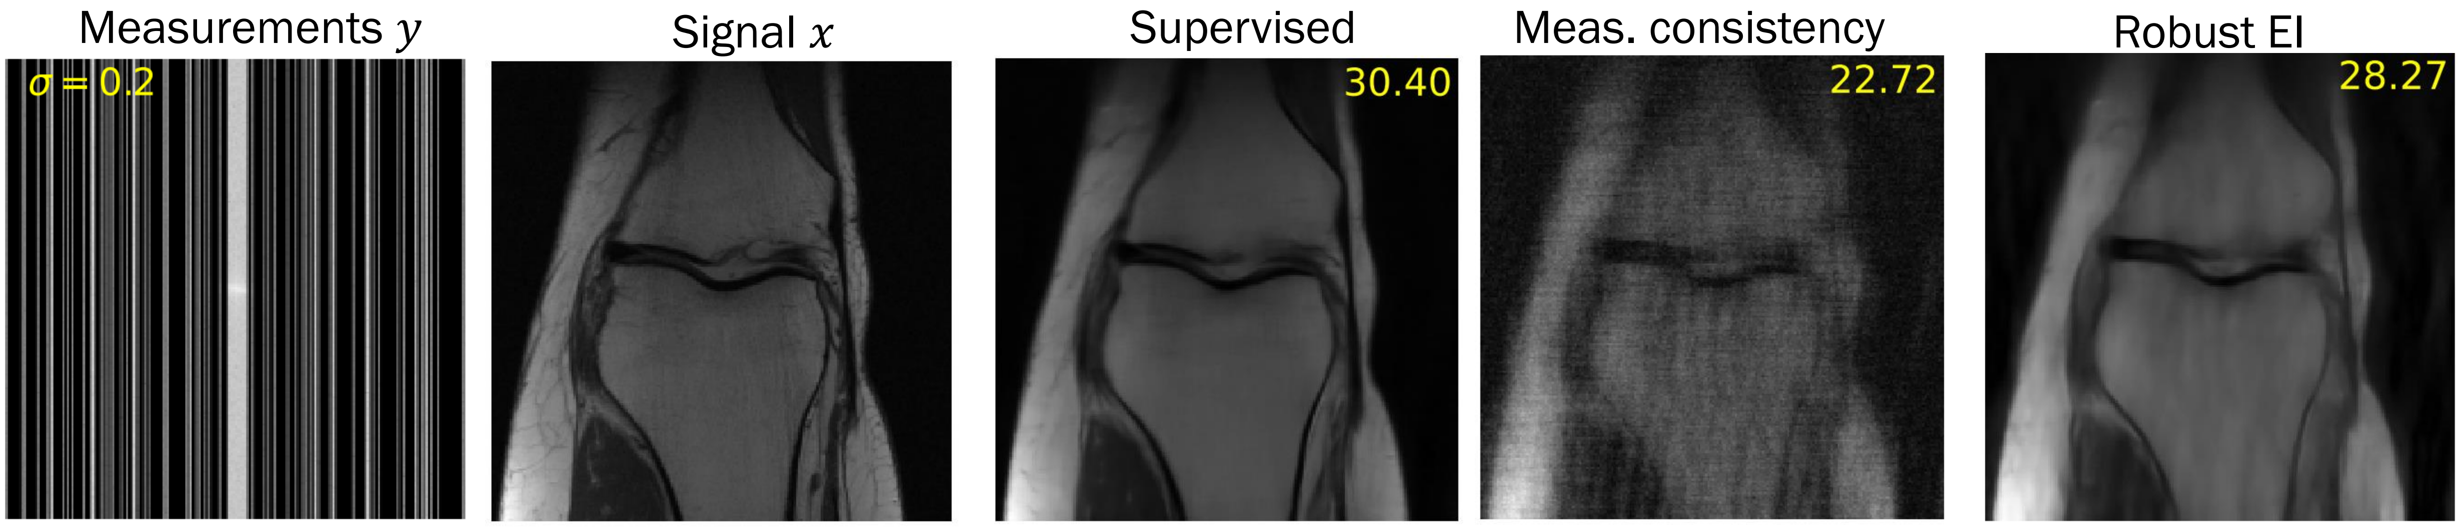

Inpainting

- Operator A is an inpainting mask (30% pixels dropped)
- Poisson noise (rate=10)
- Dataset is approximately shift invariant

Measurements y | Signal x | Supervised | Meas. consistency | Robust EI

Computed tomography

- Operator A is (non-linear variant) sparse radon transform (50 views)
- Mixed Poisson-Gaussian noise
- Dataset is approximately rotation invariant

Noisy measurements y

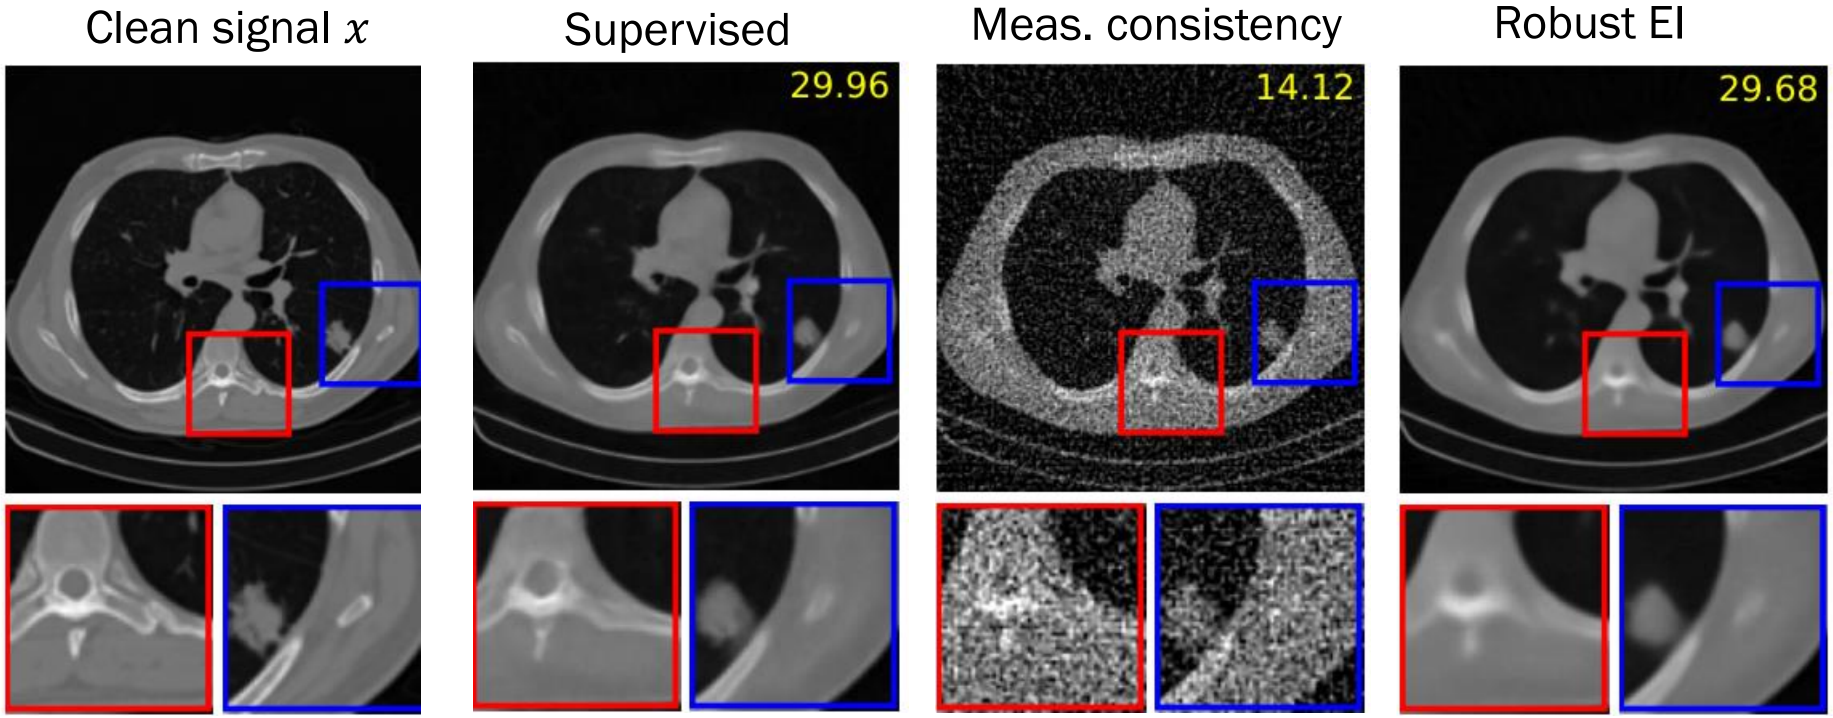

Conclusions

Novel unsupervised learning framework

- **Theory:** Necessary & sufficient conditions for learning
  - Number of measurements
  - Interplay between forward operator/ data invariance
- **Practice:** deep learning approach
  - Unsupervised loss which can be applied to any model

Conclusions

Novel unsupervised learning framework

- **Ongoing/future work**
  - More inverse problems
  - Other signal domains

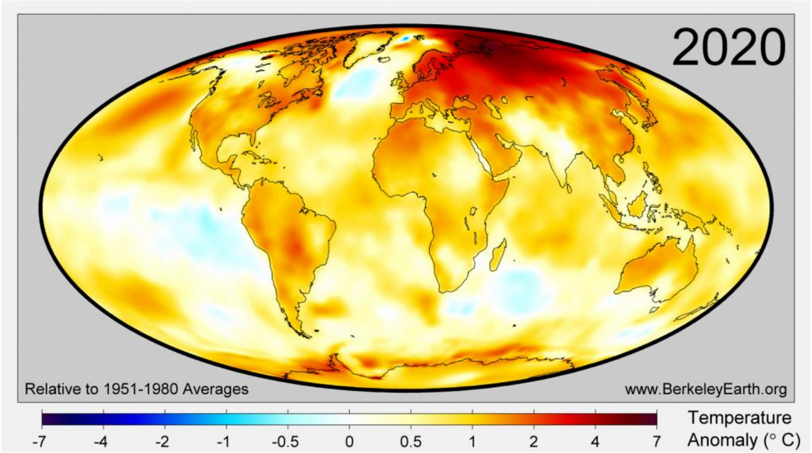

Papers

[1] “Equivariant Imaging: Learning Beyond the Range Space”, Chen, Tachella and Davies, ICCV 2021

[2] “Robust Equivariant Imaging: a fully unsupervised framework for learning to image from noisy and partial measurements”, Chen, Tachella and Davies, CVPR 2022

[3] “Sampling Theorems for Learning from Incomplete Measurements”, Tachella, Chen, Davies, Arxiv, 2022

[4] “Sampling Theorems for Unsupervised Learning in Inverse Problems” , Tachella, Chen and Davies, To Appear.

Thanks for your attention!

Tachella.github.io

- ✓ Codes
- ✓ Presentations
- ✓ … and more