Equivariant Bootstrapping for Imaging Inverse Problems

AISTATS 2024, Oral Presentation

Valencia, Spain

Julián Tachella (CNRS & ENS Lyon) & Marcelo Pereyra (Heriot-Watt University)

Inverse Problems

Goal: recover signal x from y

$$\boldsymbol{y} = A\boldsymbol{x} + \boldsymbol{\epsilon}$$

MEASUREMENT OPERATOR $\in \mathbb{R}^{m \times n}$

SIGNAL $\in \mathbb{R}^{n}$

MEASUREMENTS $\in \mathbb{R}^{m}$

NOISE $\in \mathbb{R}^{m}$

Examples

Magnetic resonance imaging

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

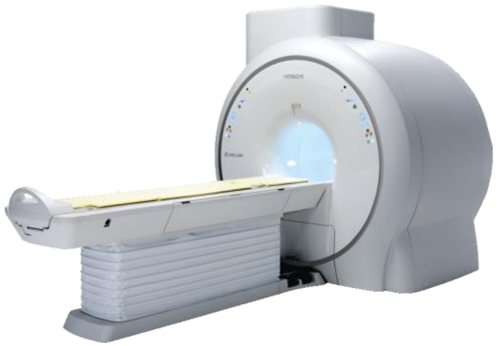

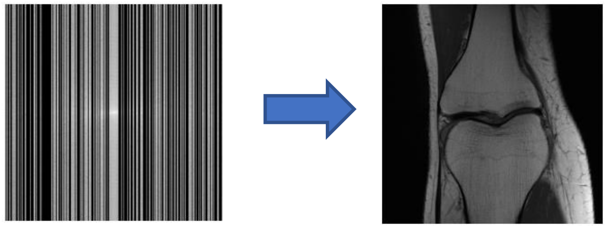

y x

Computed tomography

- A = 1D projections (sinograms) of 2D image

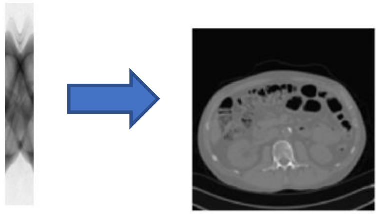

y x

Image inpainting

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

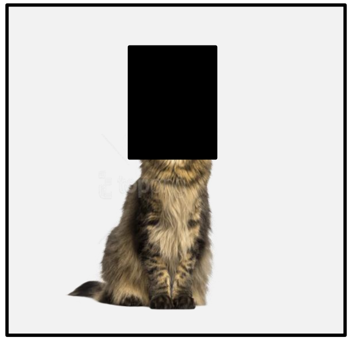

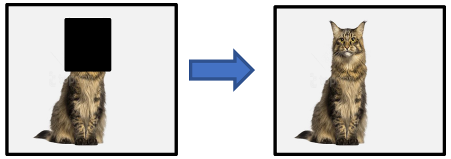

y x

Sampling Algorithms

Recent methods attempt to sample from the posterior distribution

$$-\log p(\boldsymbol{x}|\boldsymbol{y}) \propto \frac{1}{2}||\boldsymbol{y} - A\boldsymbol{x}\,||^2 - \log p(\boldsymbol{x})$$

Algorithms:

- DDRM, DiffPIR, DPS (diffusion-based sampling)
- PnP-ULA (MCMC)
- Normalizing flows

Quantifying UQ

Calibration set $\{(\boldsymbol{x}_i, \boldsymbol{y}_i)\}_{i=1}^{N}$

Empirical α coverage $= \frac{1}{N}\sum_i \mathbf{1}_{\boldsymbol{x}_i \in C_\alpha(\boldsymbol{y}_i)}$

where $C_\alpha(\boldsymbol{y}) = \{\boldsymbol{x} : ||\boldsymbol{x} - \hat{\boldsymbol{x}}(\boldsymbol{y})|| < c_\alpha\}$

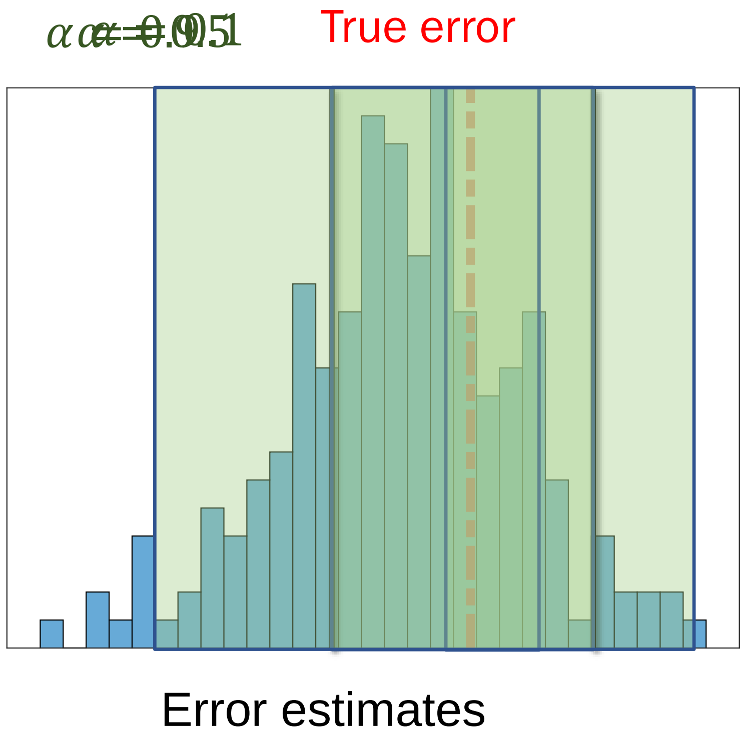

Pitfalls of UQ algorithms

Empirical observations:

- ❌ Require thousands network evals
- ❌ Fail to provide calibrated intervals
- ❌ Conformal calibration helps, but doesn't fix the problem

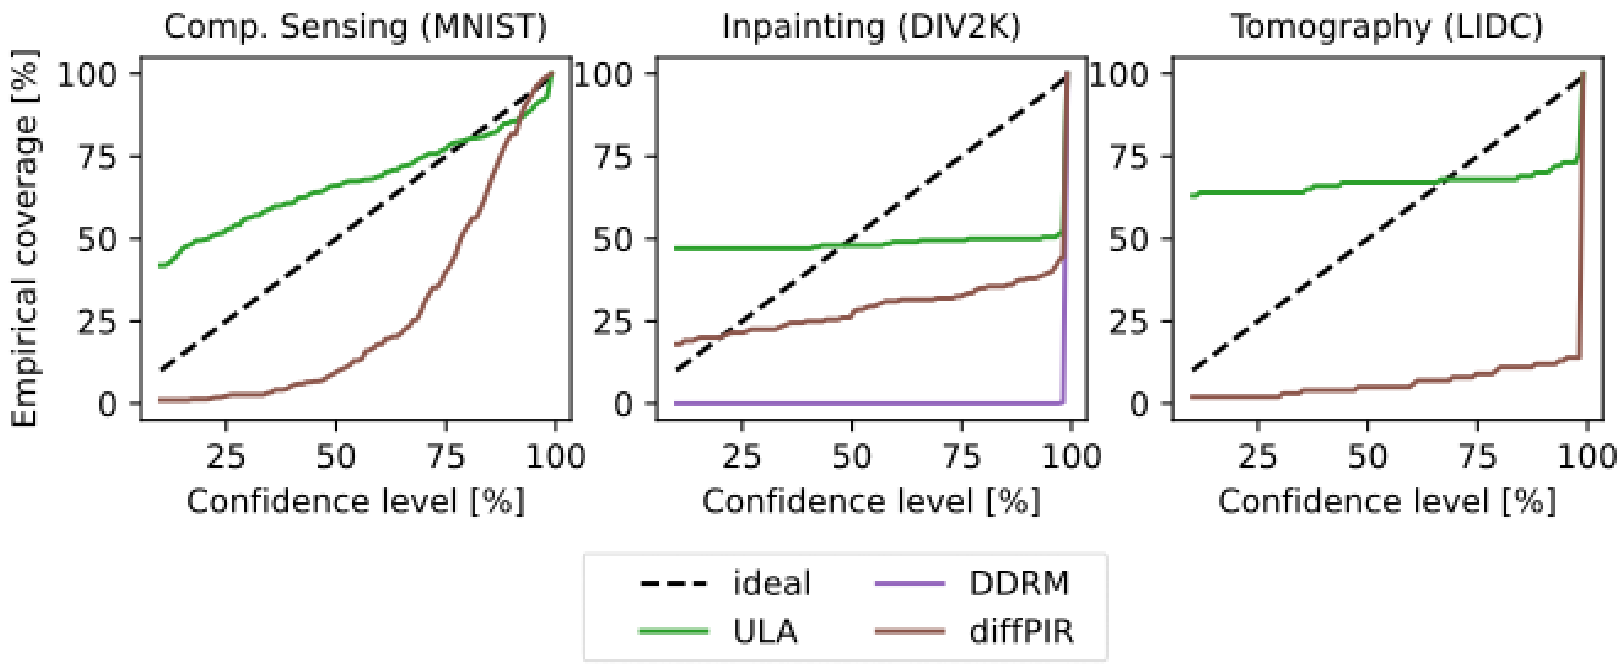

Revisiting the Bootstrap

Parametric bootstrap [Efron, 1986]**:** using $\hat{\boldsymbol{x}} = f(\boldsymbol{y})$ as ‘ground-truth’,

For $j = 1, \dots, N$

- Sample noise $\boldsymbol{\epsilon}_j \sim N(0, I\sigma^2)$
- Bootstrap $\boldsymbol{x}^j = f\left(A\hat{\boldsymbol{x}} + \boldsymbol{\epsilon}_j\right)$
- Error estimates: $\boldsymbol{e}^j = ||\,\hat{\boldsymbol{x}} - \boldsymbol{x}^j||^2$

❌ Bad UQ in the nullspace of A

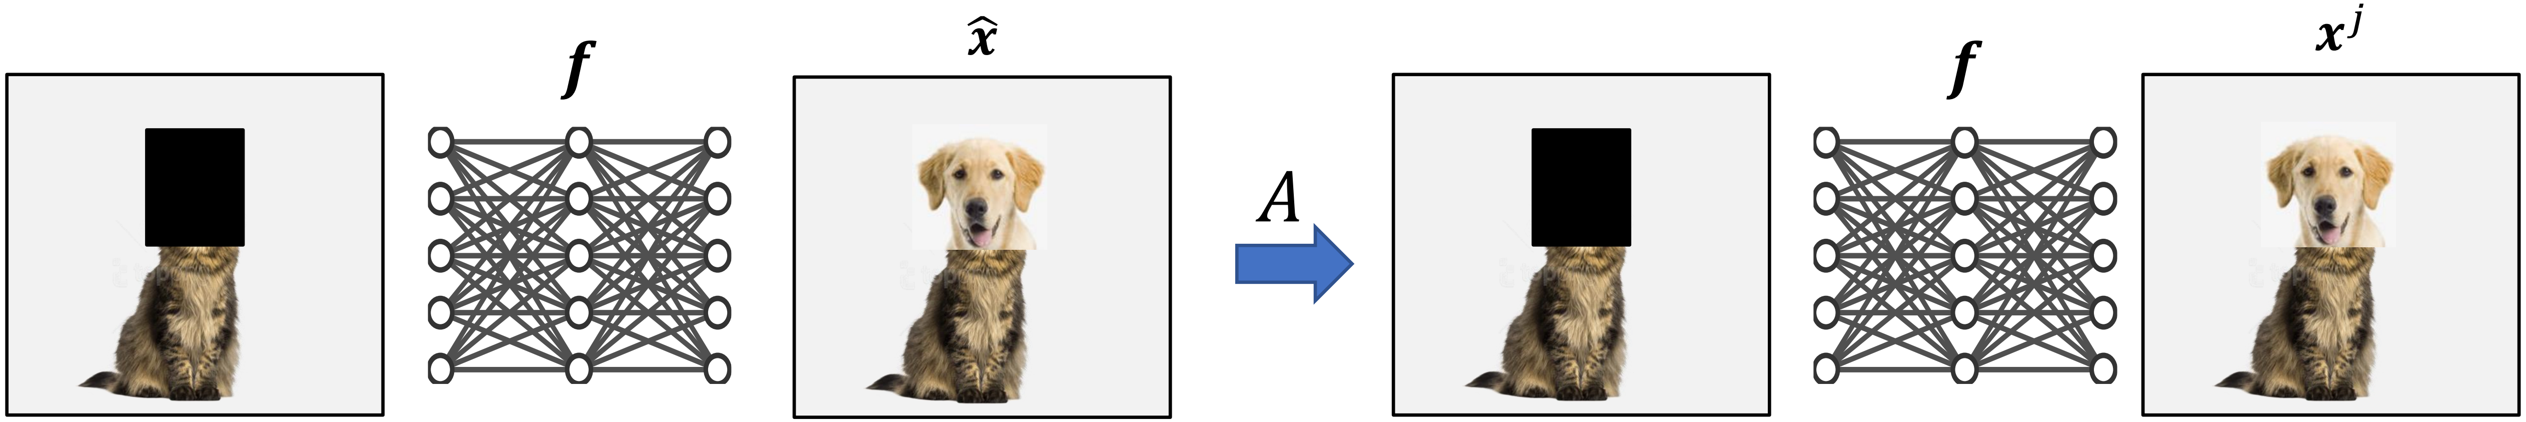

Symmetry Prior

Idea: Most natural signals sets $\mathcal{X}$ are invariant to groups of transformations.

Example: translations, rotations and flips of 2D images

For all $g \in G$ we have

$$\boldsymbol{y} = A\boldsymbol{x} = \underbrace{AT_g}_{A_g}\overbrace{T_g^{-1}\boldsymbol{x}}^{x'} = A_g\boldsymbol{x}'$$

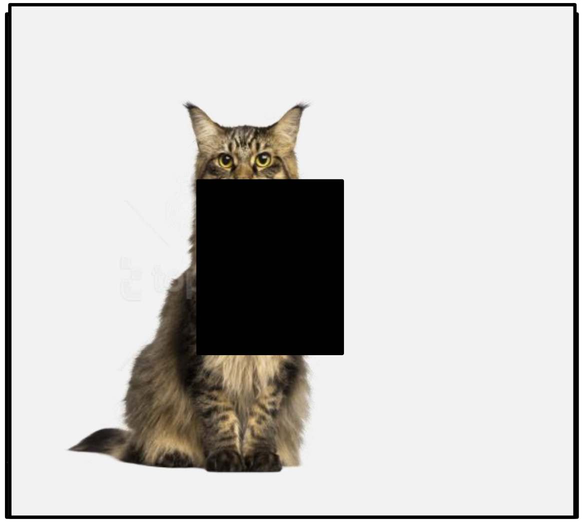

Key observation: nullspace(A) $\neq$ nullspace(AT_g)

Equivariant Bootstrap

Observation model $\begin{cases} g \sim G \\ \boldsymbol{y} = AT_g\boldsymbol{x} + \boldsymbol{\epsilon} \end{cases}$

Using $\hat{\boldsymbol{x}} = f(\boldsymbol{y})$ as ‘ground-truth’:

For $j = 1, \ldots, r$

- Sample transformation $g \sim G$ and noise $\boldsymbol{\epsilon}_j \sim N(0, I\sigma^2)$
- Bootstrap $\boldsymbol{x}^j = T_g^{-1} f(AT_g\hat{\boldsymbol{x}} + \boldsymbol{\epsilon}_j)$
- Error estimates: $\boldsymbol{e}^j = ||\hat{\boldsymbol{x}} - \boldsymbol{x}^j||^2$

Theory insights

- In the **noiseless** case, standard bootstrap gives $||\hat{\boldsymbol{x}} - \boldsymbol{x}|| = 0$ for any measurement consistent estimator verifying $A\hat{\boldsymbol{x}} = y$.

Proposition (informal). *For a linear & measurement consistent operator estimator with no noise, we have*

$$\mathbb{E}_g \, ||\hat{\boldsymbol{x}}\left(T_g A \hat{\boldsymbol{x}}\right) - T_g \boldsymbol{x}|| = ||\hat{\boldsymbol{x}} - \boldsymbol{x}|| + \text{bias}$$

where the bias term is small if $\hat{\boldsymbol{x}} \circ \boldsymbol{A}$ is not equivariant.

- Equivariant bootstrapping is useful when A **is not equivariant** to the transformations.

Equivariance of Forward Operators

	Translation	Rotation	Permutation	Amplitude
Gaussian Blur	★	★		★
Image Inpainting				★
Sparse-view CT	★			★
Accelerated MRI	★			★
Downsampling (no antialias)				★

UQ Experiments

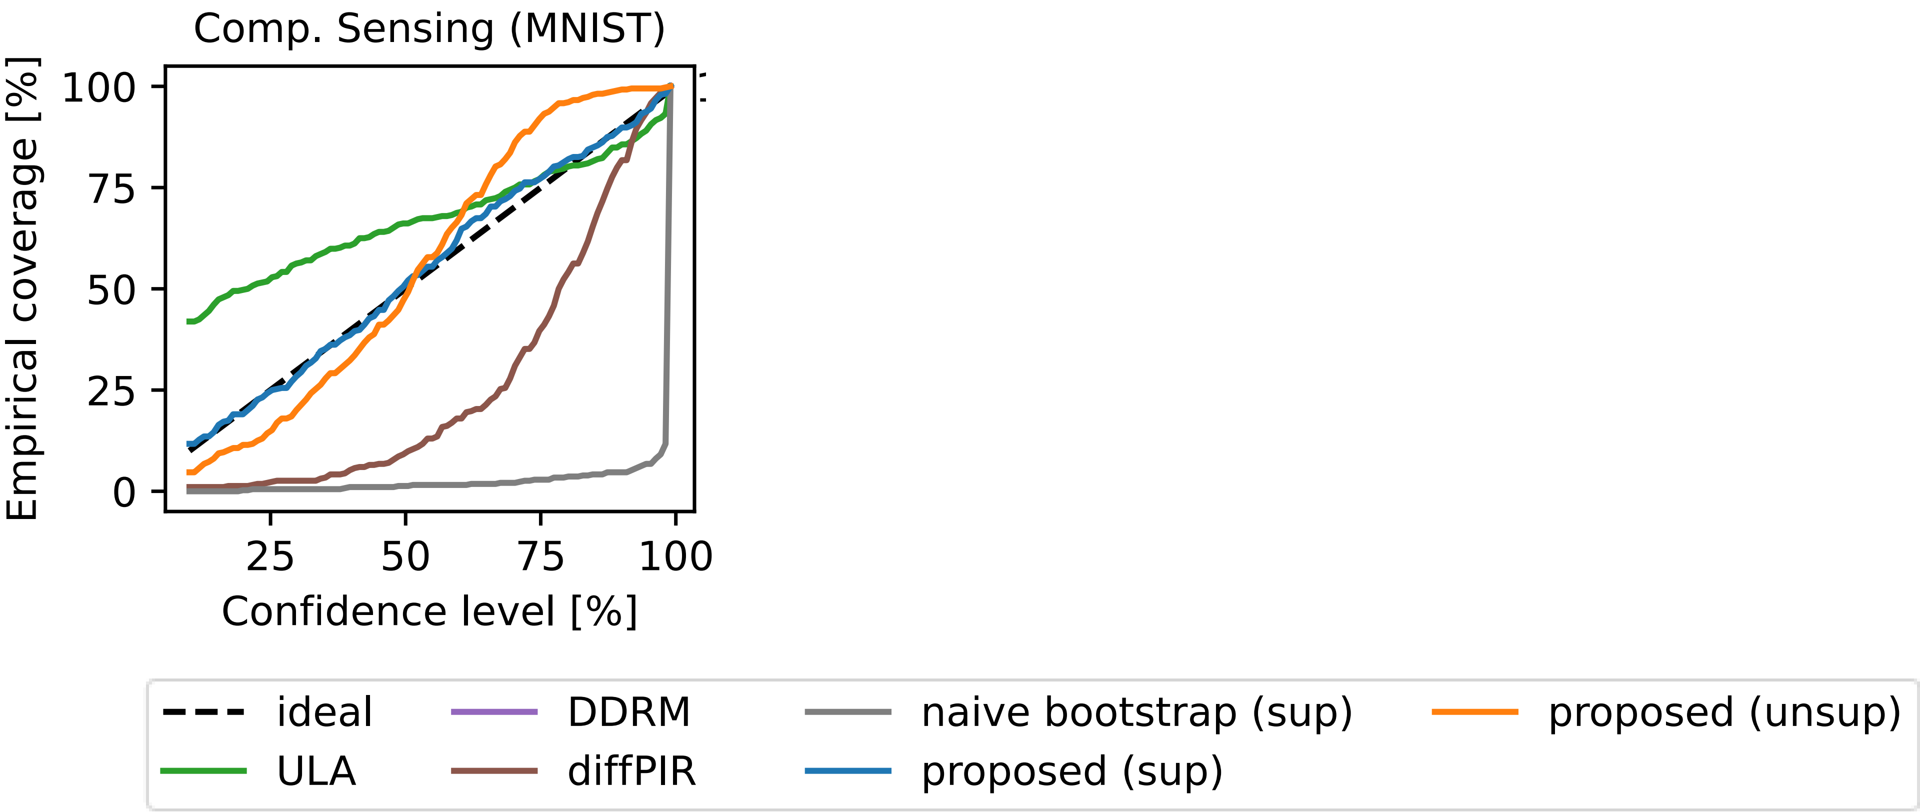

UQ Experiments

Table 1: Average test PSNR in dB for the evaluated methods.

	Diffusion (DDRM)	Diffusion (diffPIR)	ULA	Proposed bstrap (unsup. model)	Proposed bstrap (sup. model)
C. Sensing (MNIST)	-	-	28.54 ± 2.25	$\mathbf{34.11 \pm 2.09}$	33.9 ± 2.32
Inpainting (DIV2K)	32.27 ± 3.95	30.51 ± 3.74	30.52 ± 3.35	31.56 ± 4.12	$\mathbf{32.47 \pm 3.87}$
Tomography (LIDC)	-	37.02 ± 0.79	35.85 ± 0.54	37.38 ± 0.65	$\mathbf{41.03 \pm 0.91}$

Table 2: Neural function evaluations (NFEs) per Monte Carlo (MC) sample.

Method	**Diffusion**	**ULA**	**Bootstrap**
NFEs/MC sample	100	30	1

UQ Experiments

- Blur operators are shift-equivariant, thus shifts do not modify the nullspace of A

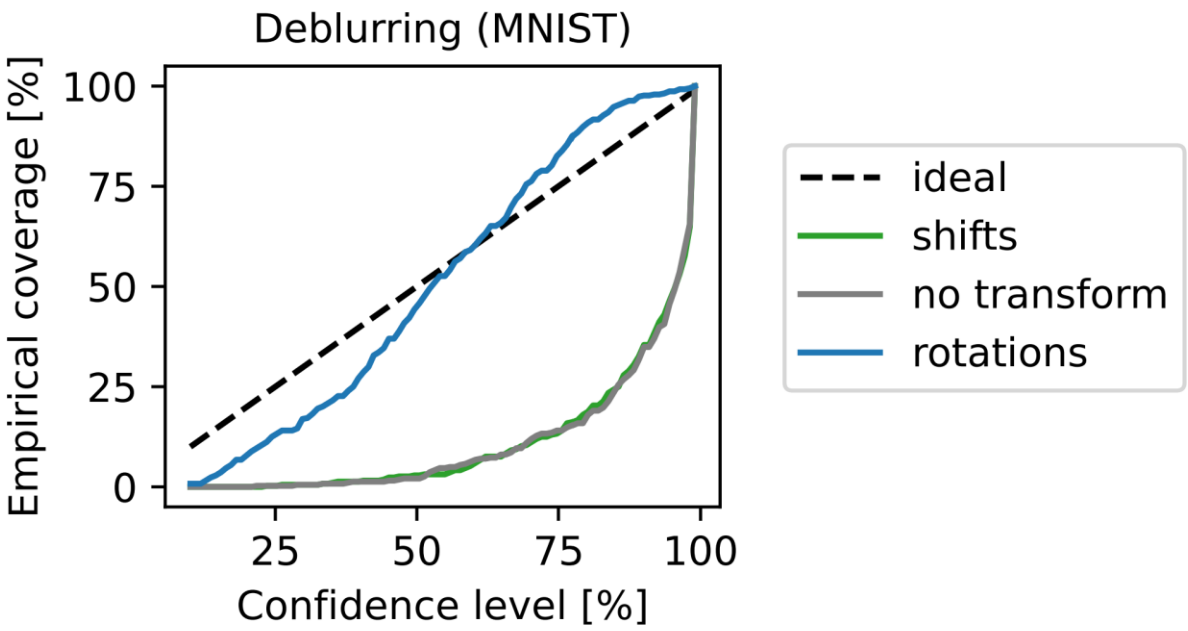

Conclusions

Take home messages

Symmetries in the data can play a key role in measuring uncertainty in the nullspace of the forward operator

Future work

- More efficient Reynolds averaging [Sannai, 2021]
- UQ challenge
- Other groups of transformations/data domains

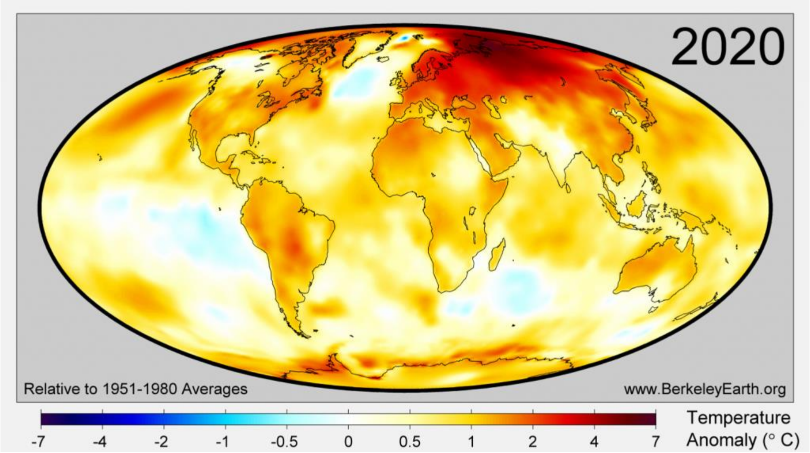

New PyTorch Library

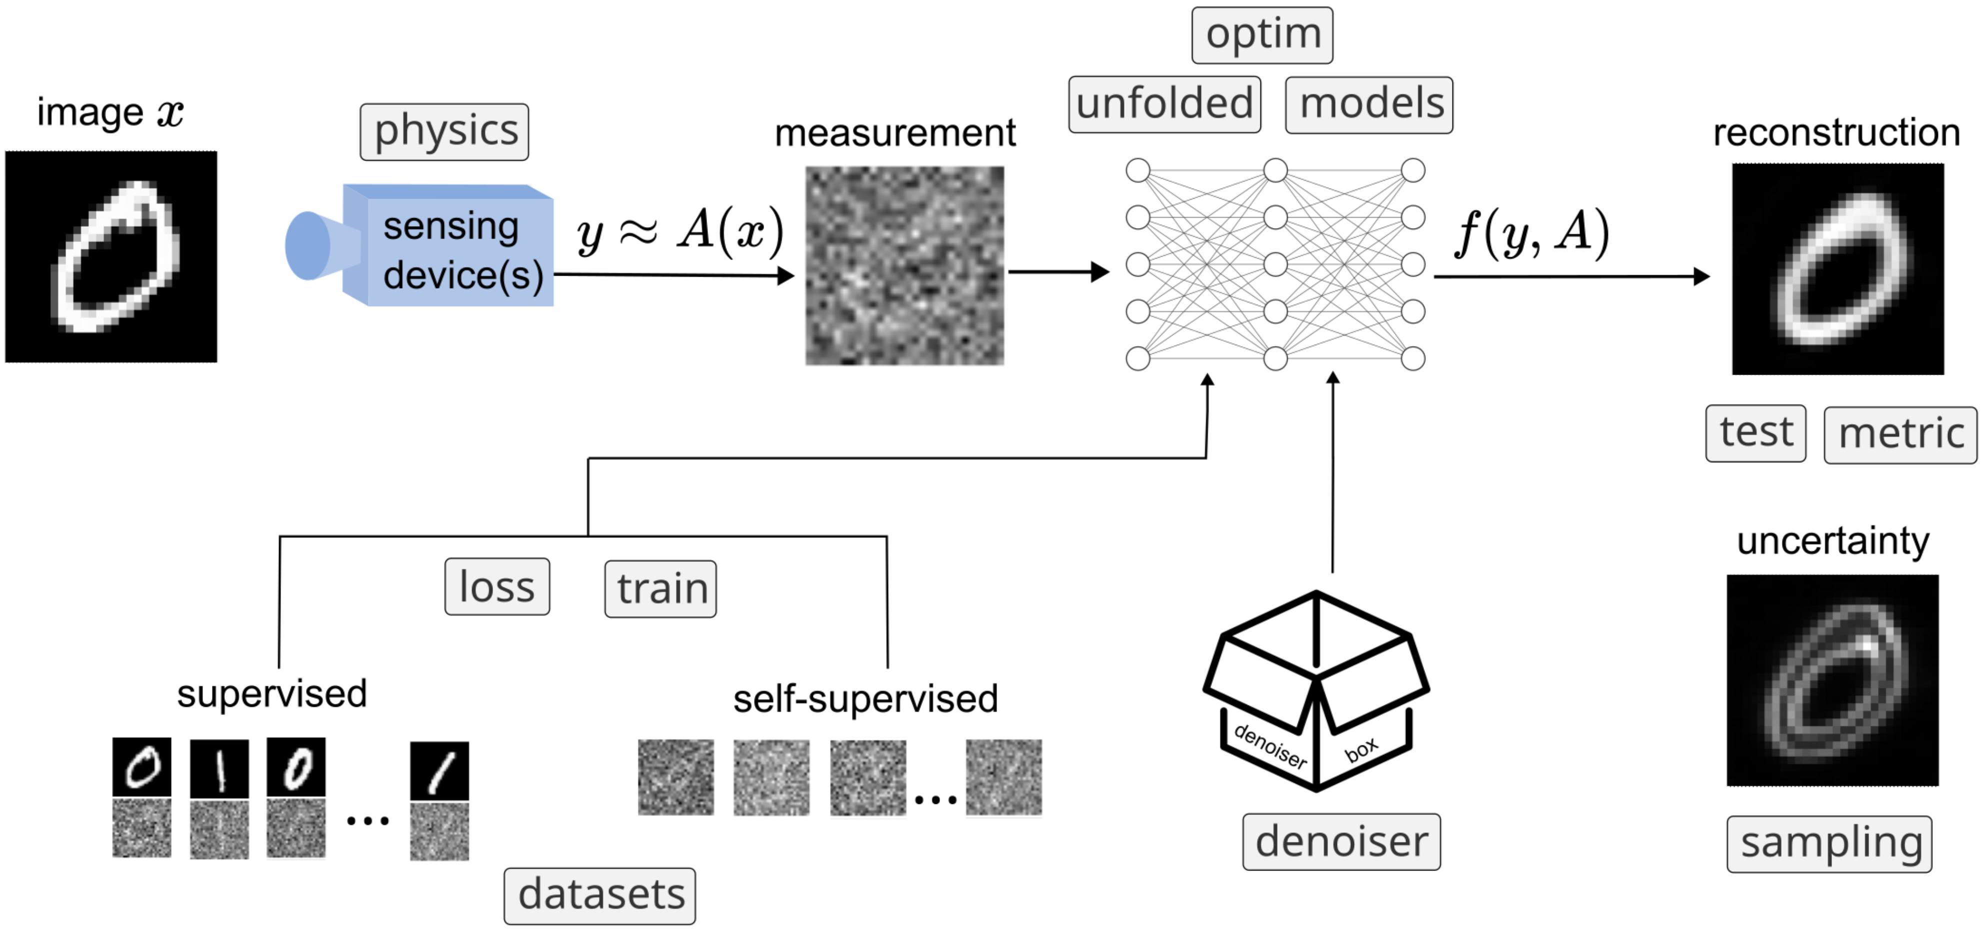

Thanks for your attention!

Tachella.github.io

- ✓ Codes
- ✓ Presentations
- ✓ ... and more

https://github.com/tachella/equivariant_bootstrap