Random mesh projectors for inverse problems

Konik Kothari*, Sidharth Gupta*, Maarten V. de Hoop†, Ivan Dokmanić*

*Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, USA

†Department of Computational and Applied Mathematics, Rice University, USA

S. Gupta and K. Kothari contributed equally

ILLINOIS Coordinated Science Lab COLLEGE OF ENGINEERING

RICE

Learning without ground truth data

- Imaging in applied sciences is exploratory in nature: seismic imaging, remote sensing, molecular imaging etc.
- Ground truth models are unavailable and need to be discovered which limits the use of data-intensive modern machine learning techniques.
- **Formulation:**

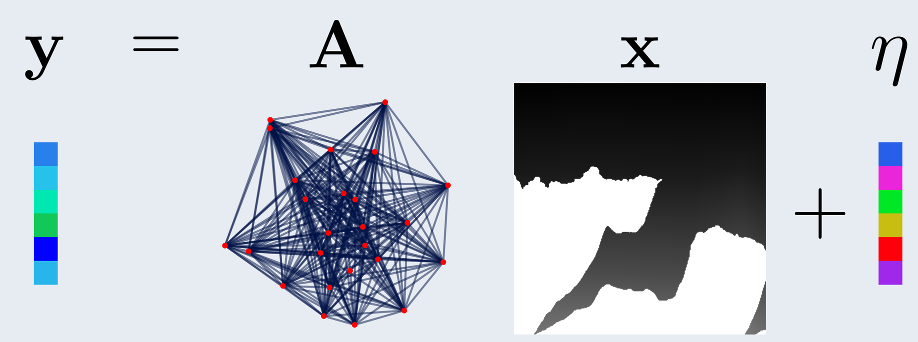

Figure 1: Formulation for linearized seismic traveltime tomography.

- $\boldsymbol{y} \in \mathbb{R}^M$ are noisy measurements.
- $\boldsymbol{A} \in \mathbb{R}^{M\times N}$, $M \ll N$ is the measurement matrix.
- $\boldsymbol{x} \in \mathcal{X} \subset \mathbb{R}^N$ is the required image.

Drawback of existing approaches

- **Classical approaches:** ℓ_1, ℓ_2 regularization and sparsity-based methods fail because we have only a few measurements ($\boldsymbol{A}$ is severely underdetermined).
- **"Deep" approaches:** Modern GAN and U-Net-based methods require lots of training data which we do not have.

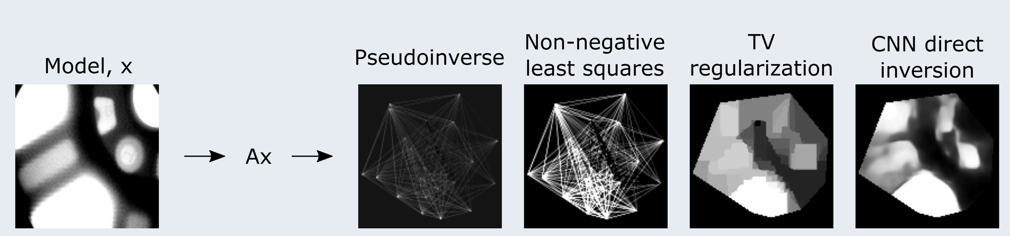

Figure 2: The existing approaches do not perform well with our severely ill-posed seismic traveltime tomography measurement matrix $\boldsymbol{A}$.

Even getting a coarse reconstruction is hard!

Can we *reliably recover* geometries using a CNN trained on a completely different dataset?

Training examples:

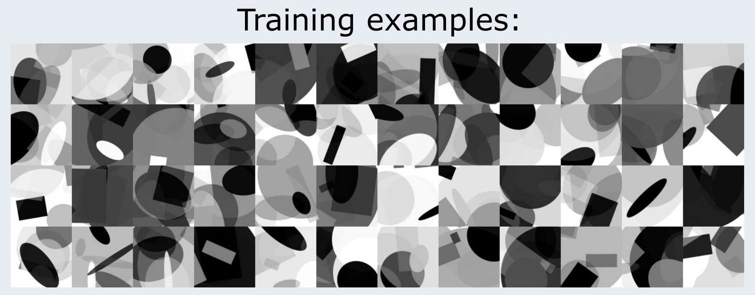

Figure 3: Reconstructions with a competitive U-Net baseline trained on a different dataset. U-Net fails to get the correct structural information.

Solution: Regularization by random projections

Assumption: $\mathcal{X}$ is a low-dimensional manifold and $\boldsymbol{A}$ is injective on $\mathcal{X}$. However, we do not know $\mathcal{X}$ or have samples from it.

Two stage method: Learn to obtain orthogonal projections of $\mathcal{X}$ from measurements, $\boldsymbol{y}$

(i) Decompose a "hard" task of learning the unstable map, $\boldsymbol{y} \to \boldsymbol{x}$ into an ensemble of "easy" tasks of learning more stable maps from $\boldsymbol{y}$ to projections of the unknown model, $\boldsymbol{x}$, into random low-dimensional subspaces.

(ii) Combine the random subspace projection estimates. Here we choose subspaces to be piecewise-constant Delaunay triangulations.

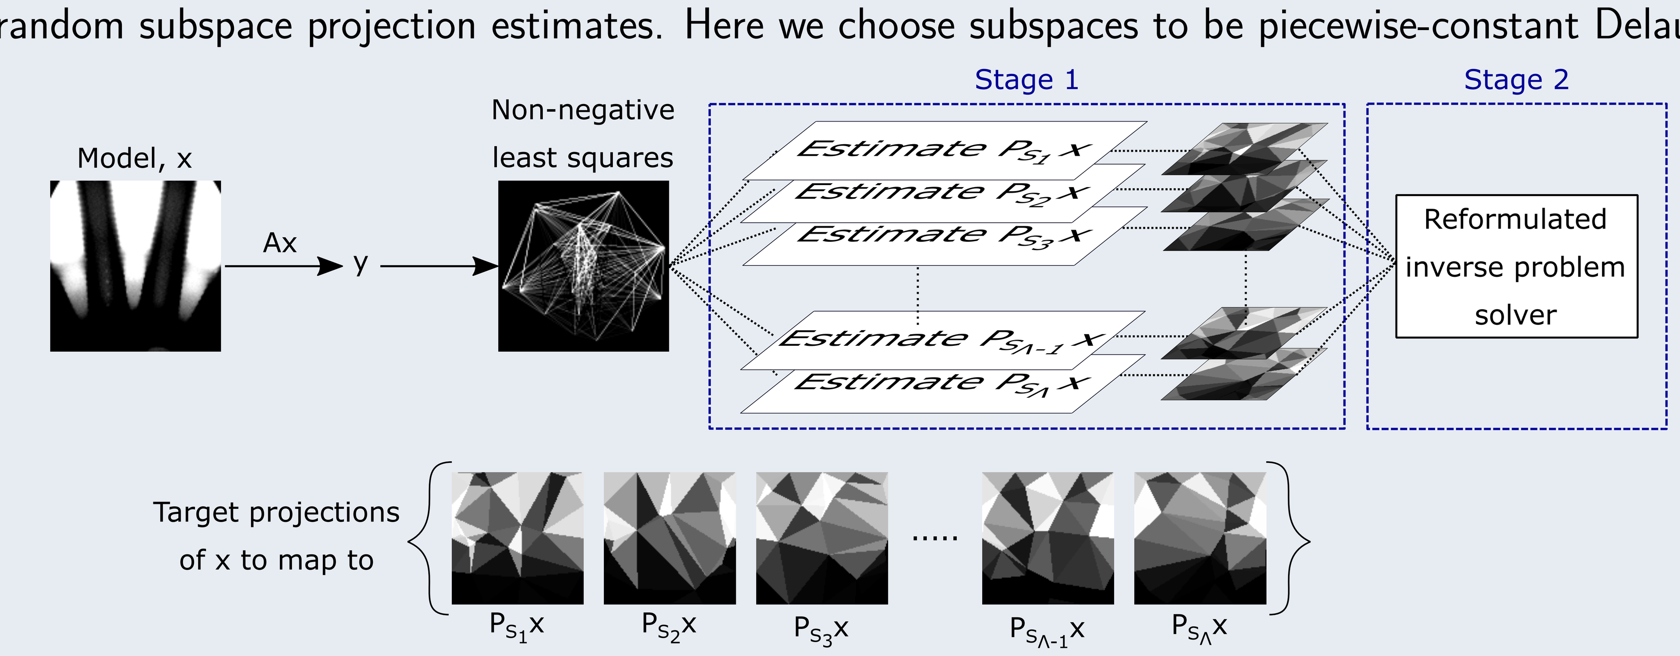

Figure 4: We estimate the projections, $\{\boldsymbol{P}_{S_\lambda}\boldsymbol{x}\}_{\lambda=1}^{\Lambda}$ onto the random subspaces $\{S_\lambda\}_{\lambda=1}^{\Lambda}$ and then combine them using a reformulated inverse problem.

Using projections allows us to generalize to scenarios not seen during training (see Results).

Reformulated inverse problem

- Let $\boldsymbol{B}_\lambda \in \mathbb{R}^{N\times K}$ be an orthogonal basis for the subspace S_λ, and $\boldsymbol{q}_\lambda \overset{\text{def}}{=} \boldsymbol{q}_\lambda(\boldsymbol{y})$ be the estimate of the expansion coefficients of $\boldsymbol{x}$ in basis $\boldsymbol{B}_\lambda$ from the measurements $\boldsymbol{y}$.
- Combine subspace estimates as $\boldsymbol{q} \overset{\text{def}}{=} \boldsymbol{q}(y) \overset{\text{def}}{=} \left[\boldsymbol{q}_1^\top, \boldsymbol{q}_2^\top, \ldots, \boldsymbol{q}_\Lambda^\top\right]^\top \in \mathbb{R}^{K\Lambda}$ and define $\boldsymbol{B} \overset{\text{def}}{=} \left[\boldsymbol{B}_1\ \boldsymbol{B}_2\ \ldots\ \boldsymbol{B}_\Lambda\right] \in \mathbb{R}^{N\times K\Lambda}$ to obtain the following reformulated inverse problem

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{\eta} \Rightarrow \boldsymbol{q} \approx \boldsymbol{B}^\top\boldsymbol{x}.$$

- Solve the reformulated problem using your favorite method, for example $\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}\in[0,1]^N}\left\{\|\boldsymbol{q} - \boldsymbol{B}^\top\boldsymbol{x}\|_2^2 + \lambda\varphi(\boldsymbol{x})\right\}$.

Need for non-linear operators to estimate projections

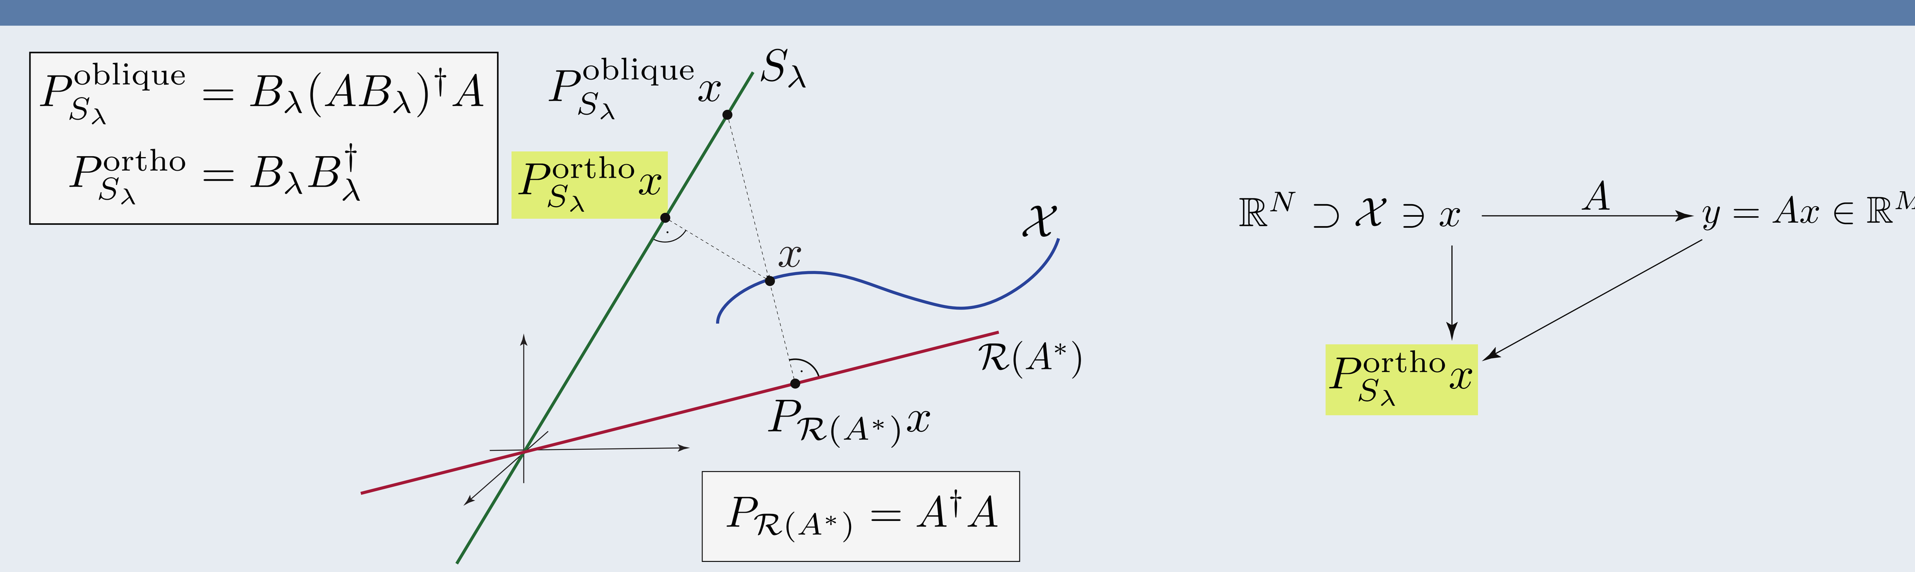

Figure 5: **Left**: Non-linear maps enable us to get null-space information from measurements. **Right**: We estimate orthogonal projections from measurements.

- The best linear operator that estimates projections into S_λ from $\boldsymbol{y}$ is an oblique projection that always lies in null-space of $\boldsymbol{A}$.
- Using non-linear methods to get orthogonal projections of $\boldsymbol{x}$ gives us missing null-space information.

References

- K. Kothari*, S. Gupta*, M. V. de Hoop, I. Dokmanić. "Random mesh projectors for inverse problems", ICLR 2019
- K. H. Jin, M. McCann, E. Froustey, M. Unser. "Deep convolutional neural network for inverse problems in imaging", IEEE Trans. on Image Processing 2017

Results

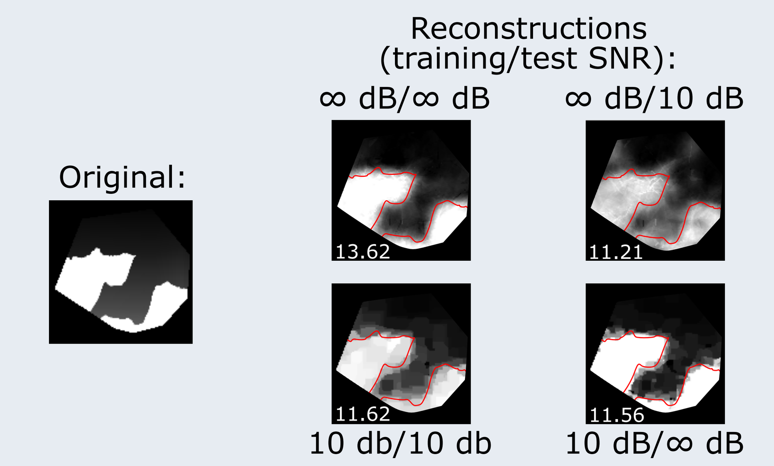

Figure 6: Reconstructions from our method are quantitatively and qualitatively better than the U-net baseline (Figure 3).

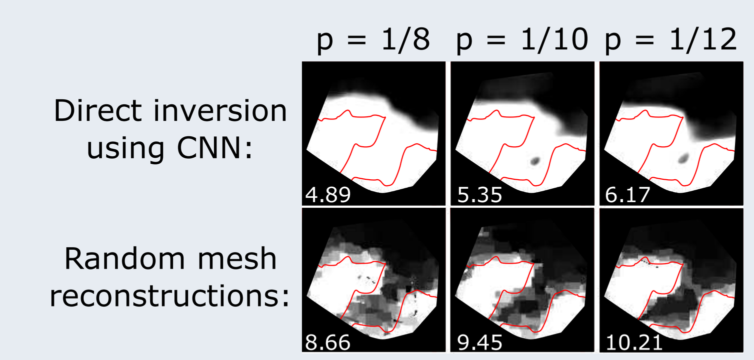

Figure 7: Our method is substantially better for noise models not seen during training. Here measurements are set to zero with probability $\boldsymbol{p}$.

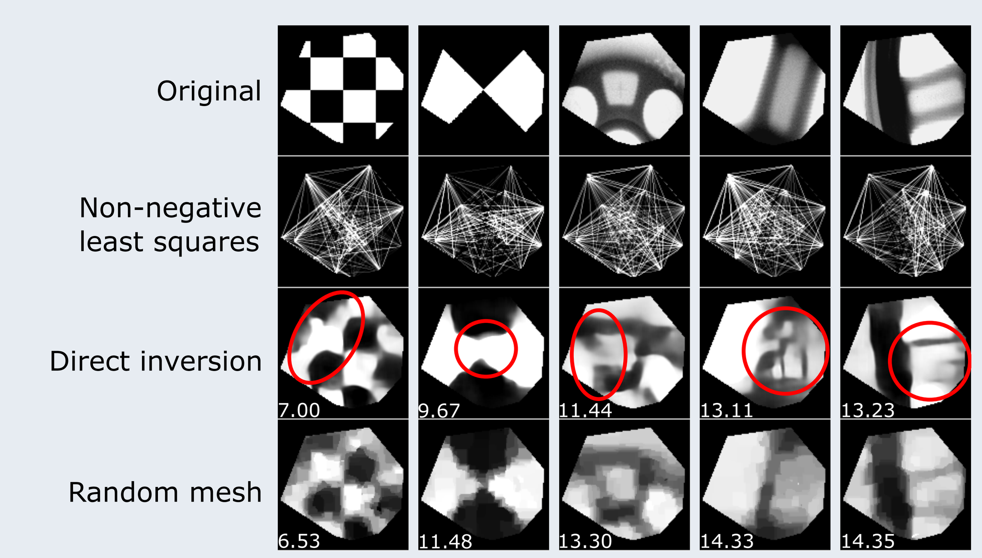

Figure 8: Further reconstructions demonstrate our method's ability to capture correct structural features.

Scenario	Direct inversion	Random mesh
10 db train and ∞ db test	13.78	15.38
∞ db train and 10 db test	10.34	12.88
10 dB train and erasures with $p = 1/8$	9.03	11.09

Table 1: Over a dataset of 102 metal casting x-ray images, our method reports better SNRs in a variety of scenarios.

Our method stabilizes the learning problem via the use of random projections and outperforms baselines even when tested on scenarios not seen in training. Such scenarios are typical in the applied sciences.

Future work

- We want to extend our approach to adversarial imaging scenarios in the applied sciences.
- We are working to improve the second stage of our method using modern regularizers like deep-image prior.