Efficient Posterior Inference & Generalization in Physics-based Bayesian Inference with Conditional GANs

USC University of Southern California

Deep Ray, Dhruv V. Patel, Harisankar Ramaswamy and Assad A. Oberai

{deepray, dhruvvpa, hramaswa, aoberai}@usc.edu

Motivation

Goal: Given $\boldsymbol{y} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{\eta}$, infer $p(\boldsymbol{x}|\boldsymbol{y})$.

Challenges:
- Measurement may be corrupted by unknown noise.
- Inverse map may not be well-posed.
- Uncertainty in inferred solution critical for applications with high-stake decisions.

Example: medical imaging

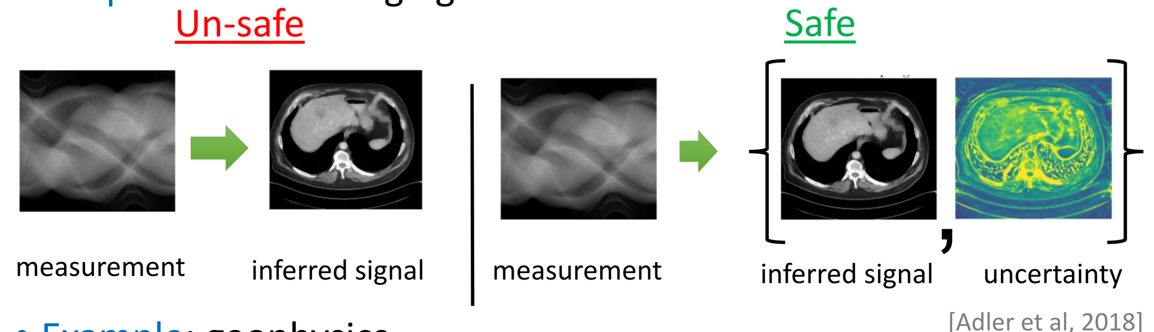

- Example: geophysics

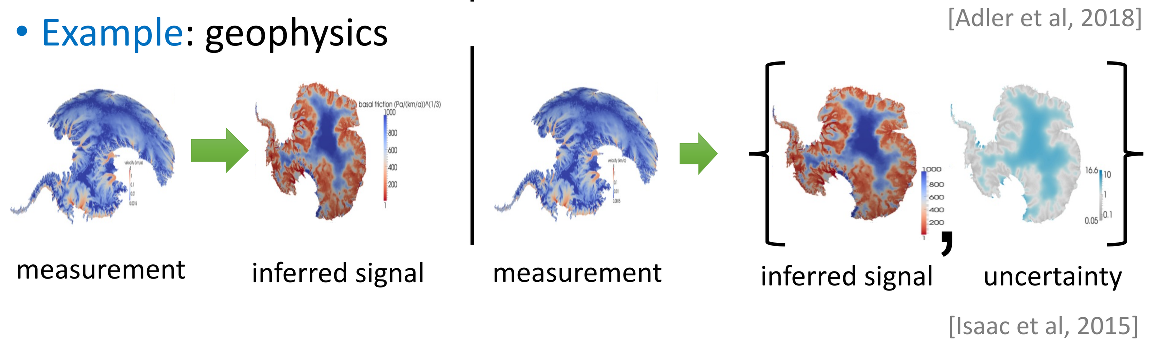

Bayesian inference

- A principled approach to account for uncertainty in an inverse problem.
- Gives probability distribution over inferred field given some measurement.

Posterior distribution

$$p_X^{post}(\boldsymbol{x}|\boldsymbol{y}) = \frac{1}{Q} p^{like}(\boldsymbol{y}|\boldsymbol{x})\, p_X^{prior}(\boldsymbol{x}) \propto p_\eta(\boldsymbol{y} - \boldsymbol{f}(\boldsymbol{x}))\, p_X^{prior}(\boldsymbol{x})$$

Challenge I : Sampling
- Inferred signal is high dimensional (10^3-10^7).
- Difficult to sample from high dimensional posterior space using sampling-based methods like MCMC.

Challenge II : Priors

Difficult to characterize p_X^{prior} analytically when the prior knowledge is given in terms of complex data or constraints.

Challenge III : Forward map
- The forward model may be unknown (or a black-box simulator)
- What is available? Pairwise $(\boldsymbol{x}, \boldsymbol{y})$ data.

Challenge IV : Generalizability

Algorithms trained using dataset A may fail on dataset B containing out-of-distribution (OOD) samples.

cGAN as Posterior

Key idea: Using a conditional GAN for sampling from the posterior in physics-based inverse problems.

Let

$$\boldsymbol{z} \in \Omega_Z \subset \mathbb{R}^n, \quad \boldsymbol{x} \in \Omega_X, \quad \boldsymbol{y} \in \Omega_Y$$
$$\boldsymbol{z} \sim p_Z, \quad (\boldsymbol{x}, \boldsymbol{y}) \sim p_{XY}$$

Consider the generator and critic

$$d: \Omega_X \times \Omega_Y \mapsto \mathbb{R}, \quad \boldsymbol{g}: \Omega_Z \times \Omega_Y \mapsto \Omega_X$$

Define the loss

$$\mathcal{L}(d, \boldsymbol{g}) = \underset{\substack{(\boldsymbol{x},\boldsymbol{y})\sim p_{XY} \\ \boldsymbol{z}\sim p_Z}}{\mathbb{E}} [d(\boldsymbol{x},\boldsymbol{y}) - d(\boldsymbol{g}(\boldsymbol{z},\boldsymbol{y}),\boldsymbol{y})] = \underset{\substack{\boldsymbol{x}\sim p_{X|y} \\ \boldsymbol{y}\sim p_Y}}{\mathbb{E}} [d(\boldsymbol{x},\boldsymbol{y})] - \underset{\substack{\boldsymbol{x}^g\sim p^g_{X|y} \\ \boldsymbol{y}\sim p_Y}}{\mathbb{E}} [d(\boldsymbol{x}^g,\boldsymbol{y})]$$

Under the constraint that d is 1-Lipschitz, solve

$$(d^*, \boldsymbol{g}^*) = \underset{\boldsymbol{g}}{\text{argmin}}\, \underset{d}{\text{argmax}}\, \mathcal{L}(d, \boldsymbol{g})$$
$$\Rightarrow \boldsymbol{g}^*(.,\boldsymbol{y}) = \underset{\boldsymbol{g}}{\text{argmin}}\, W_1(p_{X|y}, \boldsymbol{g}_\#(.,\boldsymbol{y})p_Z)$$

Convergence in Wasserstein-1 metric implies weak convergence

$$\underset{\boldsymbol{x}\sim p_{X|y}}{\mathbb{E}} [m(\boldsymbol{x})] = \underset{\boldsymbol{x}^g\sim p^g_{X|y}}{\mathbb{E}} [m(\boldsymbol{x}^g)] = \underset{\boldsymbol{z}\sim p_Z}{\mathbb{E}} [m(\boldsymbol{g}^*(\boldsymbol{z},\boldsymbol{y}))] \quad \forall m \in C_b(\Omega_X)$$

Procedure

Algorithm
1. Generate $\{\boldsymbol{x}^{(i)}\}_{i=1}^M$ samples from p_X.
2. For each $\boldsymbol{x}^{(i)}$, solve the forward model to obtain (noisy) $\boldsymbol{y}^{(i)}$.
3. Train the cGAN using the set of pairwise samples $\{(\boldsymbol{x}^{(i)}, \boldsymbol{y}^{(i)})\}_{i=1}^M$.
4. For a given $\boldsymbol{y}$, sample $\boldsymbol{z} \sim p_Z$ and push it through $\boldsymbol{g}^*(\cdot,\boldsymbol{y})$ to sample from the posterior.
5. Compute the posterior MC statistics on the set $\{\boldsymbol{g}^*(\boldsymbol{z}^{(j)},\boldsymbol{y})\}_{j=1}^N$.

Network architecture:

Generator:
- U-net architecture with convolution-residual blocks on each level.
- The latent information injected at various levels through conditional instance normalization.

Critic:
- Down-sampling with residual blocks, followed by dense layers.

Results

Inverse heat conduction: *Inferring conductivity*

$$-\nabla.(x(\boldsymbol{s})\nabla y(\boldsymbol{s}))\big|_\Omega = b(\boldsymbol{s}), \qquad y(\boldsymbol{s})|_{\partial\Omega} = 0$$

$x :=$ thermal conductivity
$y :=$ temperature
$b :=$ heat source

Training samples = 8000, dim(z) = 50

Same distribution as training set

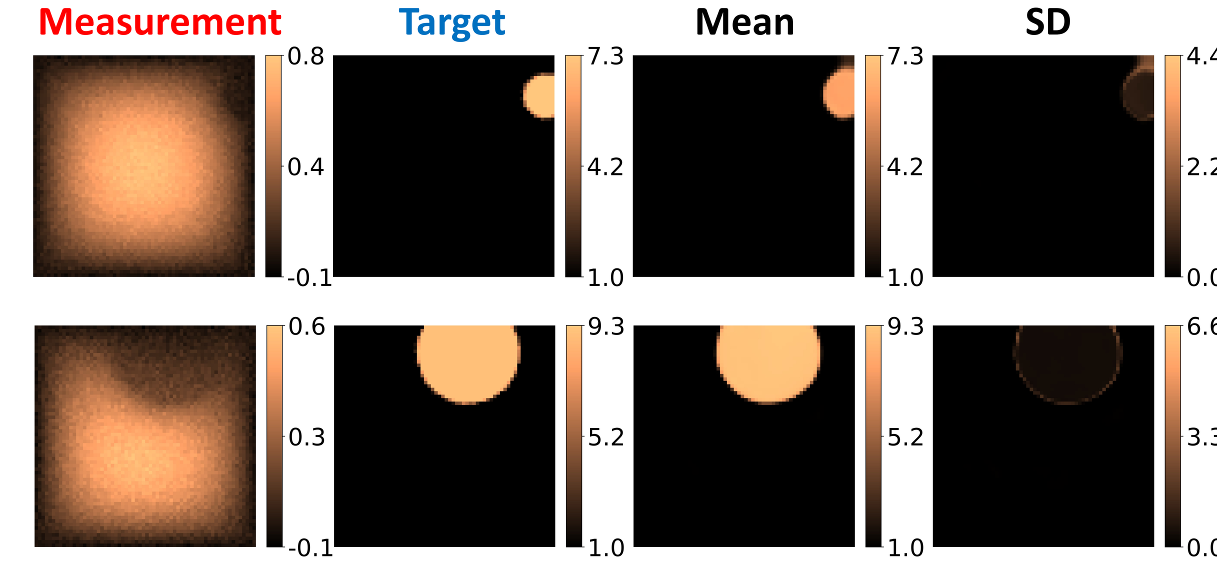

Out of distribution (OOD)

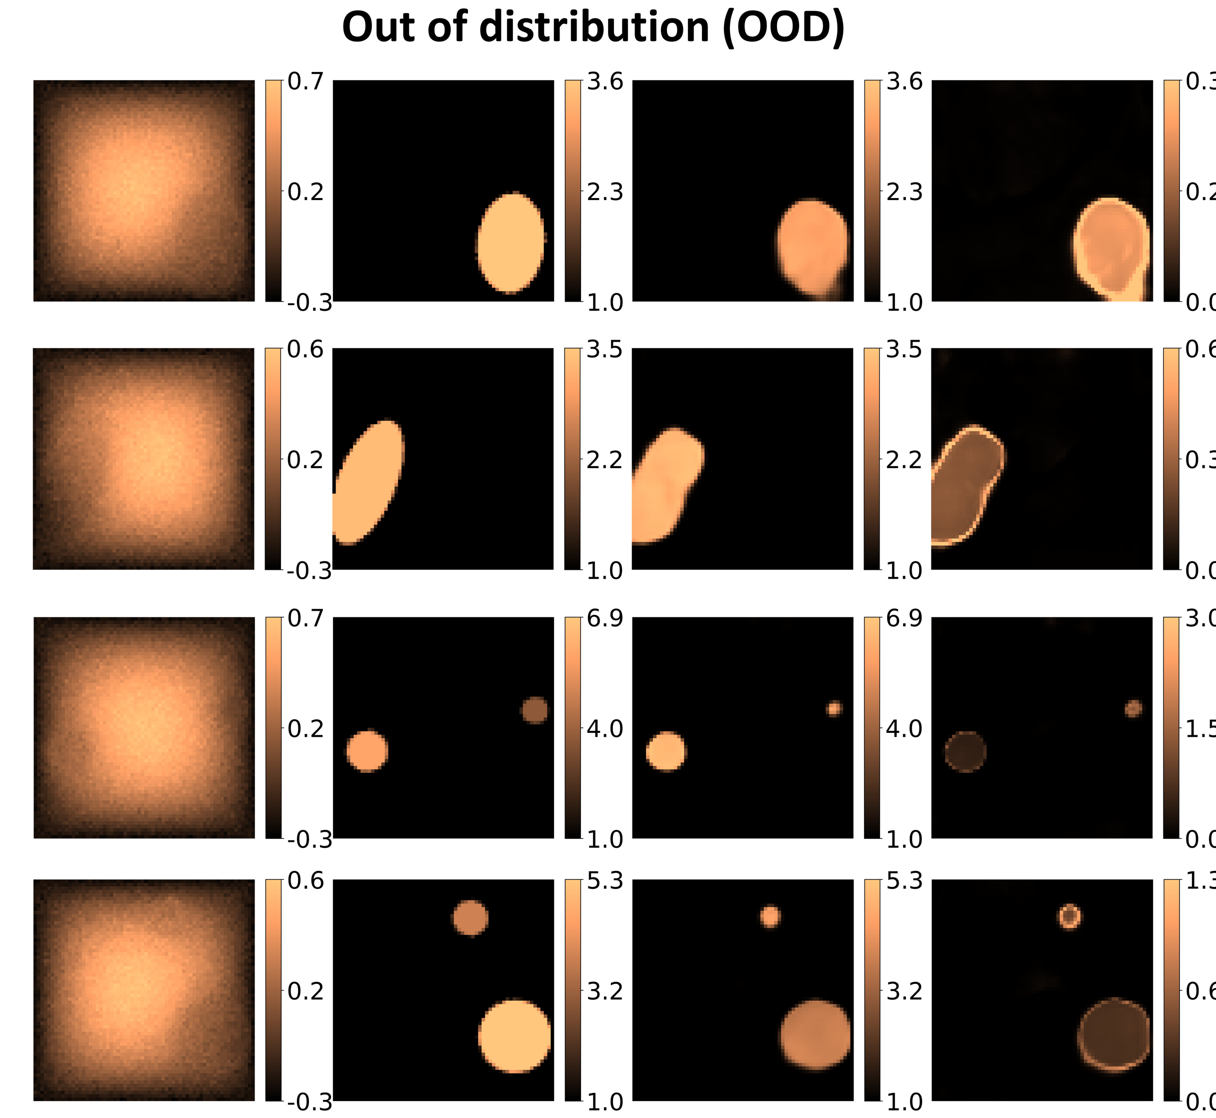

Inverse heat conduction: *Inferring initial condition*

$$\frac{\partial u(\boldsymbol{s},t)}{\partial t} - \nabla.(k(\boldsymbol{s})\nabla u(\boldsymbol{s},t))\big|_\Omega = b(\boldsymbol{s}), \quad u(\boldsymbol{s},0)|_\Omega = x(\boldsymbol{s}), \quad u(\boldsymbol{s},t)|_{\partial\Omega} = 0$$

$k :=$ thermal conductivity
$u :=$ temperature
$x := u(\boldsymbol{s},0)$
$y := u(\boldsymbol{s},1)$
$b :=$ heat source

Training samples = 8000, dim(z) = 100

Same distribution as training set

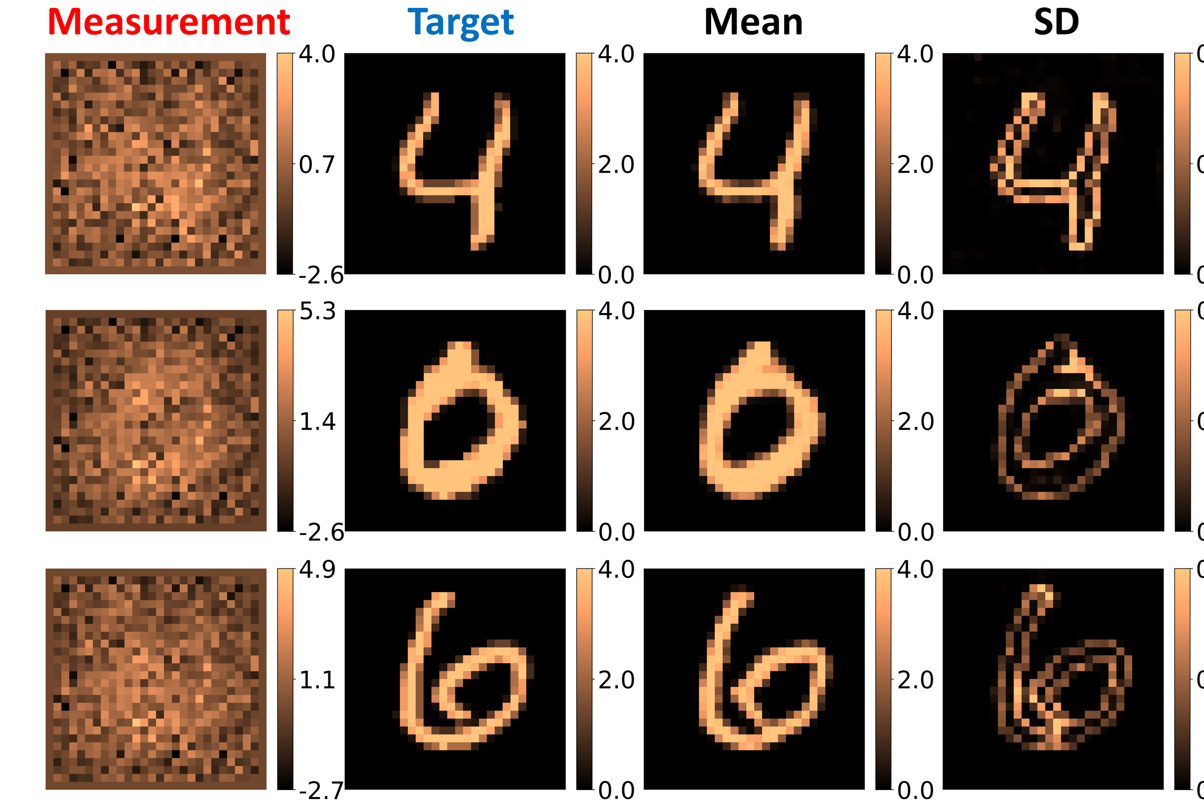

Out of distribution (OOD)

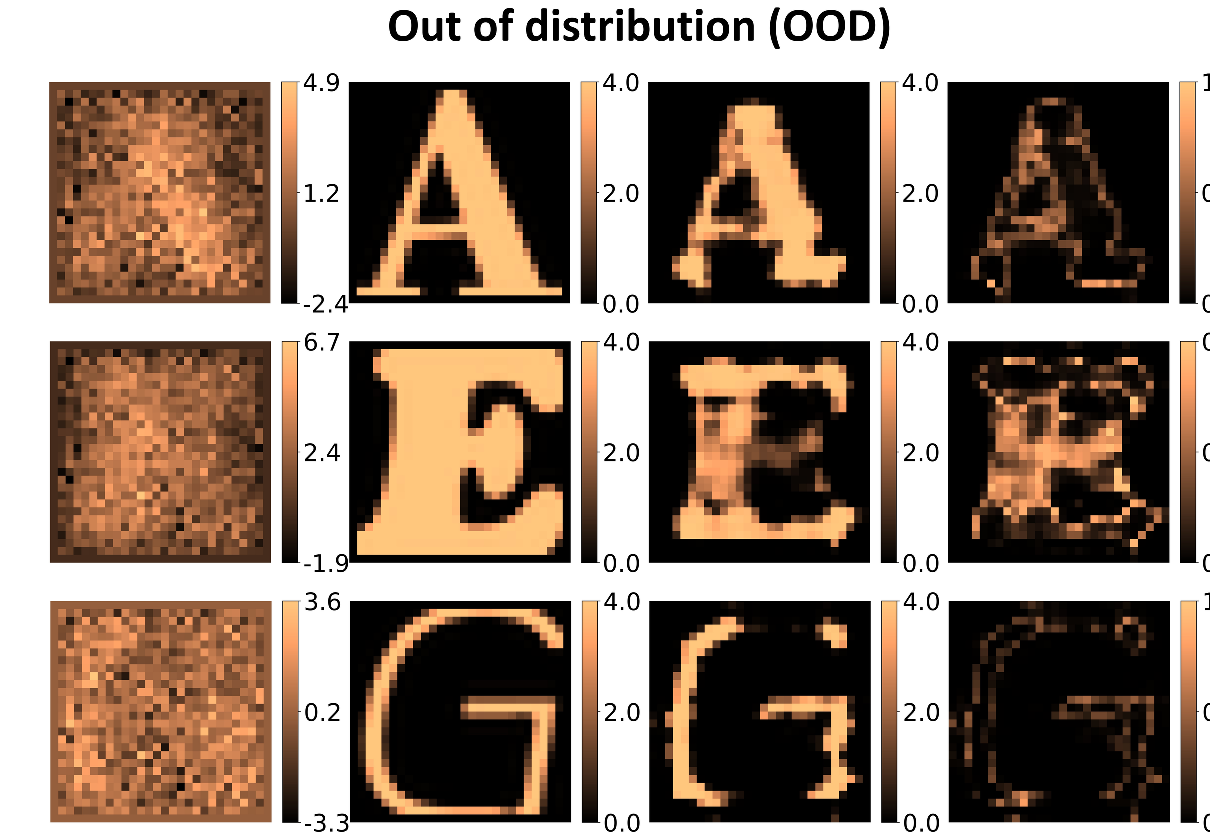

Gradients: *Locality of inverse map*

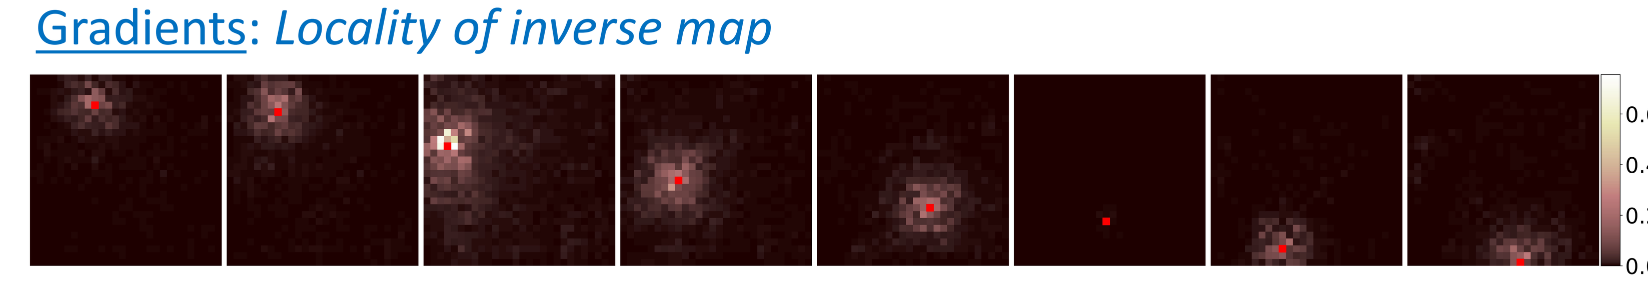

$$\frac{\partial x_k}{\partial \boldsymbol{y}} \approx \text{grad}_k = \frac{1}{100}\sum_{i=1}^{10}\sum_{j=1}^{10}\left|\frac{\partial g_k}{\partial \boldsymbol{y}}(\boldsymbol{z}^{(j)},\boldsymbol{y}^{(i)})\right|$$
$$\boldsymbol{y}^{(i)} \sim p_Y, \quad \boldsymbol{z}^{(j)} \sim p_Z$$

Influence of x_k is the neighborhood of y_k !

Key takeaways:
- ✓ cGANs can learn the posterior distribution in physics-informed inference.
- ✓ cGANs have the capacity to generalize on OOD measurements.
- ✓ Locality of the learned inverse map (promoted by architecture) helps with generalizability.

Paper: D. Ray ,D. Patel, H. Ramaswamy, A. Oberai, "*Efficient Posterior Inference & Generalization in Physics-based Bayesian Inference with Conditional GANs*", Deep inverse workshop (NeurIPS 2021): https://openreview.net/pdf?id=VC7UtQ2j0XW

Acknowledgement: The support from ARO grant W911NF2010050 is acknowledged.