Bayesian Inference with GAN Priors

Dhruv Patel and Assad A Oberai

{dhruvvpa, aoberai}@usc.edu

USC University of Southern California

Why care about Bayesian inversion?

Classical (regularization) approach:

Goal: Recover image x* from measurement y.

$$x^* = \arg\min_x \frac{1}{2}\|f(x) - \hat{y}\|_w^2 + \|x - x_{ref}\|_R^2$$

➢ *Applications where high-stake decisions are made based on the solution of an inverse problem, it is critical to account for uncertainty in the inferred solution.*

Example 1: *medical imaging*

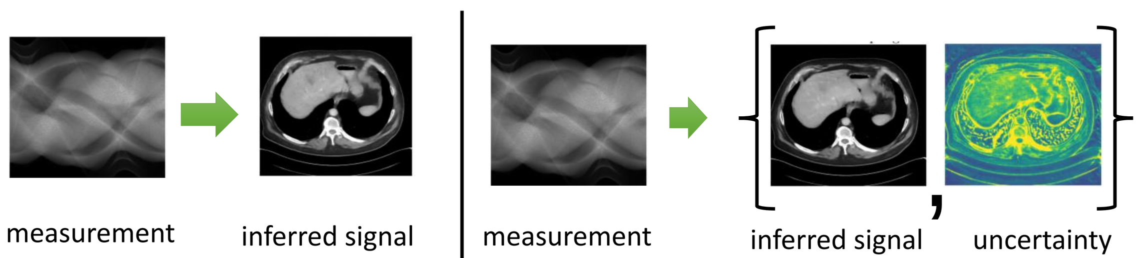

Example 2: *self-driving cars*

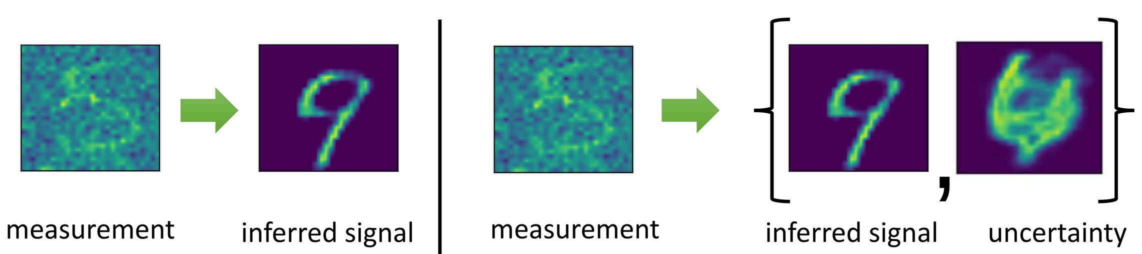

Motivation

Bayesian inference:

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

Posterior distribution

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

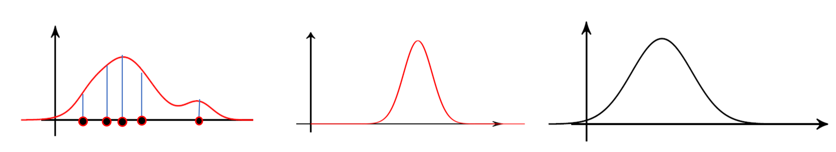

Challenge I : Priors

Finding a quantitative description of informative and feasible priors.

Typical priors..

$$p^{prior}(x) = exp\left(-\frac{1}{\sigma^2}\|x\|^2\right)$$

However, what if..

- the prior knowledge is more complex and difficult to characterize analytically.
- not enough domain knowledge is available to construct informative priors.

Challenge II : Sampling

- Inferred signal is high dimensional (10^3-10^7).
- Difficult to sample from high dimensional posterior space using sampling-based methods like MCMC.
- An efficient sampler is difficult to design in high-dimension.

GAN as Prior in BI

Key idea: Use the distribution learned by GAN as a surrogate for prior distribution and reformulate the inference problem in low-dimensional latent space of GAN.

$$\mathbb{E}_{x\sim p_X^{post}}[m(x)] = \frac{1}{\mathbb{Z}} \mathbb{E}_{x\sim p_X^{prior}}[m(x)p_\eta(\hat{y} - f(x))]$$
$$= \frac{1}{\mathbb{Z}} \mathbb{E}_{x\sim p_X^{data}}[m(x)p_\eta(\hat{y} - f(x))]$$
$$= \frac{1}{\mathbb{Z}} \mathbb{E}_{z\sim p_Z}[m(g(z))p_\eta(\hat{y} - f(g(z)))]$$
$$= \frac{1}{\mathbb{Z}} \mathbb{E}_{z\sim p_Z^{post}}[m(g(z))]$$

where,

$$p_Z^{post} = p_\eta\left(\hat{y} - f(g(z))\right) p_Z(z)$$

Steps:

1. Learn the prior distribution: train a GAN using samples from data distribution $p_X^{data}(x)$.
2. Characterize the posterior distribution: for a given measurement $\hat{y}$, evaluate any statistic of interest $\mathbb{E}_{x\sim p_X^{post}}[m(x)]$.

MCMC Sampling: $p_Z^{MCMC}(z|\hat{y}) \approx p_Z^{post}(z|\hat{y})$.

Evaluate any point estimate $s(x)$ as,

$$s(x) \approx \frac{1}{N}\sum_{n=1}^{N} s(g(z)), \qquad z \sim p_Z^{MCMC}(z|\hat{y})$$

We use Hamiltonian Monte Carlo (HMC) with burn-in period of 0.5

Results

Denoising:

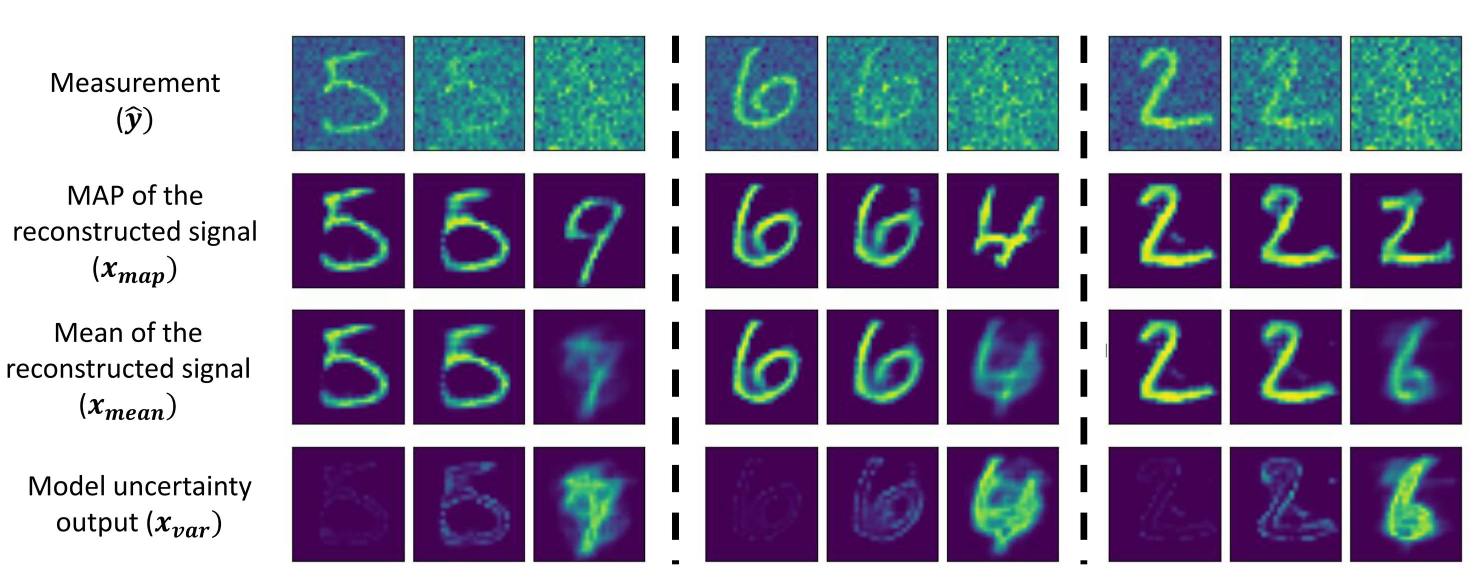

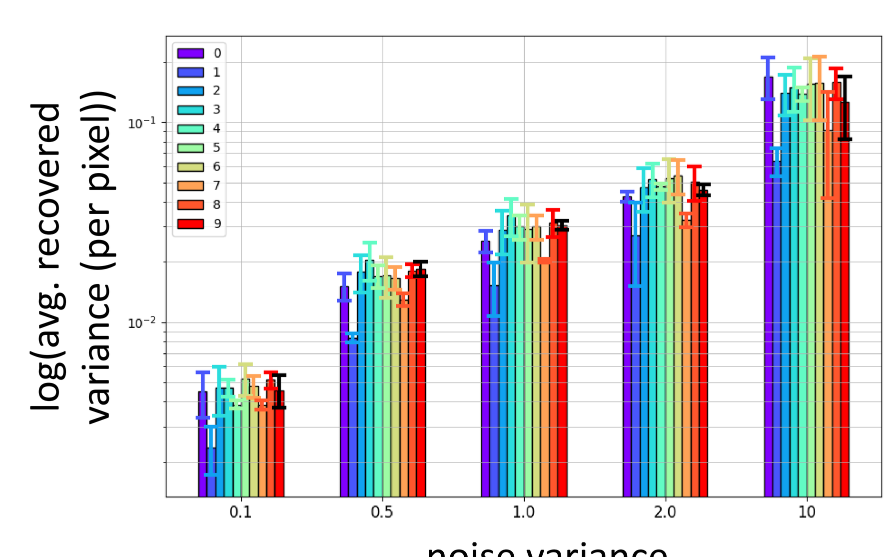

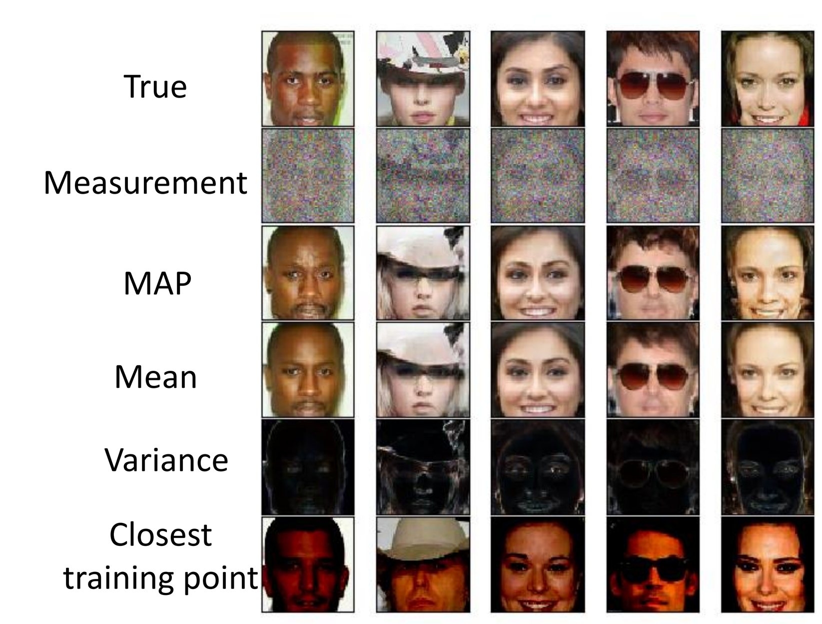

Active learning / Design of Experiments:

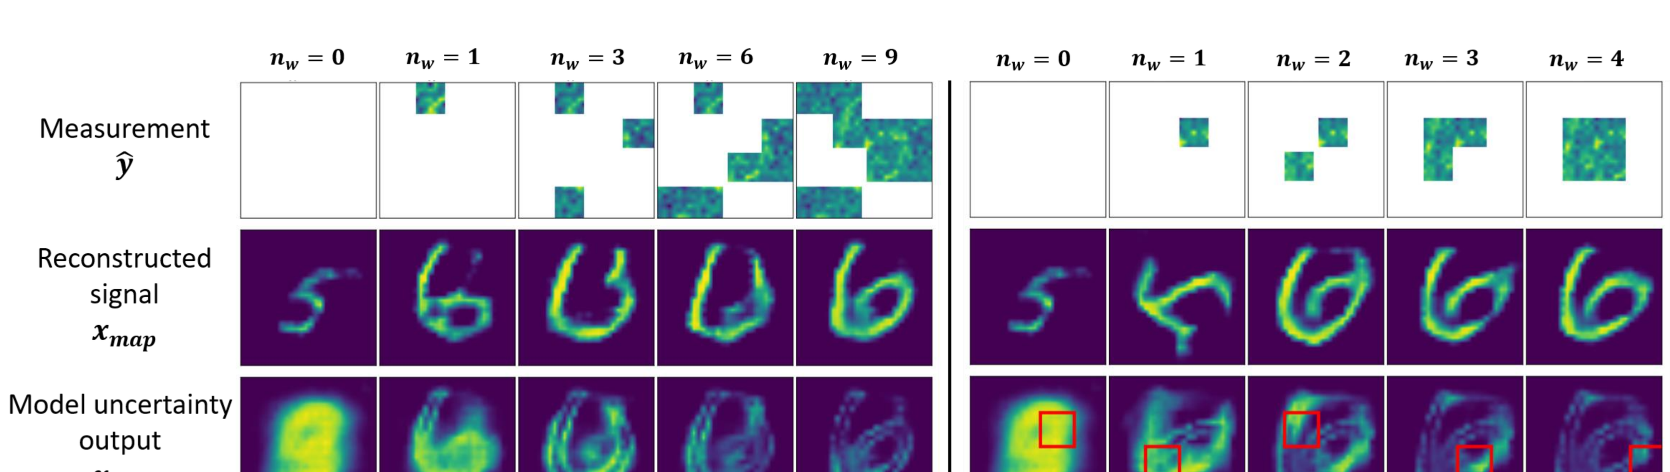

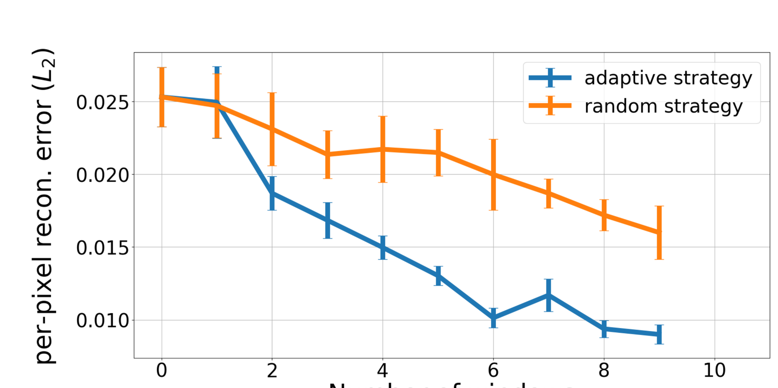

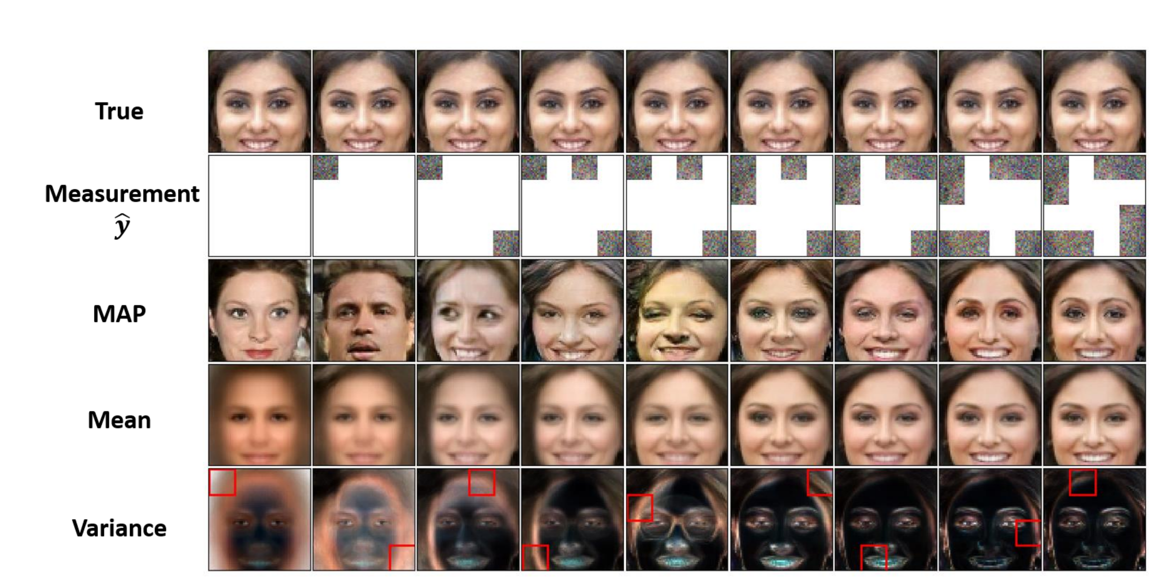

Inpainting:

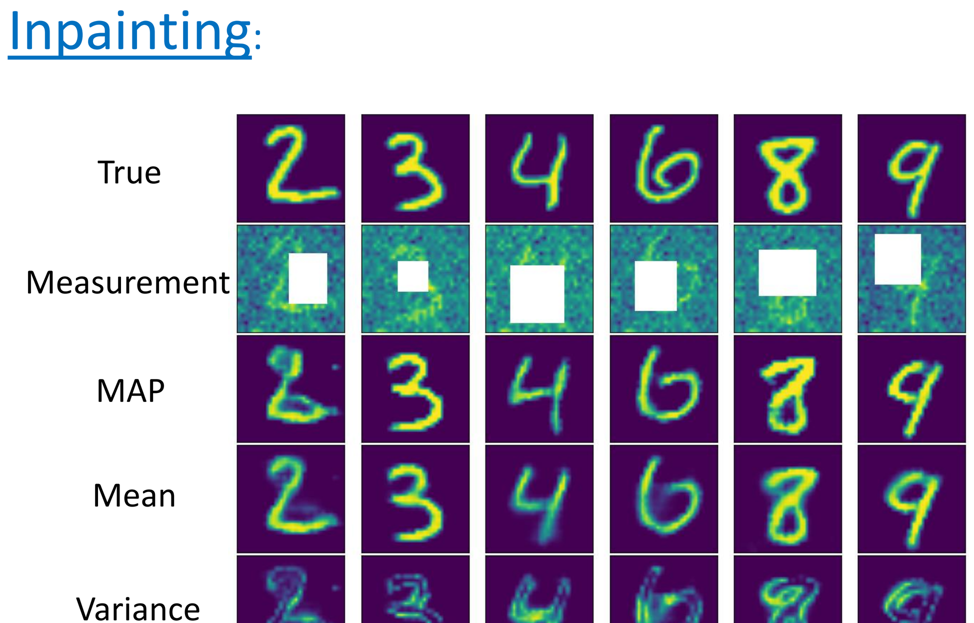

Physics-driven inference:

Key takeaways:

- ✓ A novel method for performing Bayesian inference involving complex priors and high-dimensional posterior.
- ✓ Novel unsupervised probabilistic field inference algorithm.
- ✓ Demonstration of the utility of uncertainty quantification to facilitate active learning.

Paper:

D. Patel, A. Oberai, "*GAN priors for Bayesian inference*", Deep inverse workshop (NeurIPS 2019): https://openreview.net/pdf?id=HJlL2Q2qLS