

A Novel Conditional Wasserstein GAN for Bayesian Inference

Deep Ray
Department of Mathematics
University of Maryland, College Park

Email: deepray@umd.edu
Website: deepray.github.io

Supported by ARO grant W911NF2010050 and CARG-USC.

SIAM UQ - 24
Trieste, February 27, 2024

Motivation: Forward and inverse problems

Consider a forward problem

$$\mathcal{F} : \mathbf{x} \in \Omega_x \mapsto \mathbf{y} \in \Omega_y, \quad \Omega_x \in \mathbb{R}^{N_x}, \quad \Omega_y \in \mathbb{R}^{N_y}$$

For example the **heat conduction** PDE:

$$\frac{\partial u(\mathbf{s}, t)}{\partial t} - \kappa \Delta u(\mathbf{s}, t) = 0$$


where

$$u(\mathbf{s}, 0) = u_0(\mathbf{s})$$


$u(\mathbf{s}, t)$ → temperature at location \mathbf{s} at time t

$u_0(\mathbf{s})$ → initial temperature at location \mathbf{s}

κ → thermal conductivity of material



Forward problem \mathcal{F} : Given $u_0(\mathbf{s})$ at the sensor nodes determine $u(\mathbf{s}, T)$



Motivation: Forward and inverse problems

Consider a forward problem

$$\mathcal{F} : \mathbf{x} \in \Omega_x \mapsto \mathbf{y} \in \Omega_y, \quad \Omega_x \in \mathbb{R}^{N_x}, \quad \Omega_y \in \mathbb{R}^{N_y}$$

For example the **heat conduction** PDE:

$$\frac{\partial u(\mathbf{s}, t)}{\partial t} - \kappa \Delta u(\mathbf{s}, t) = 0$$


where

$$u(\mathbf{s}, 0) = u_0(\mathbf{s})$$


$u(\mathbf{s}, t)$ → temperature at location \mathbf{s} at time t

$u_0(\mathbf{s})$ → initial temperature at location \mathbf{s}

κ → thermal conductivity of material



Inverse problem \mathcal{F}^{-1} : Given $u(\mathbf{s}, T)$ at the sensor nodes infer $u_0(\mathbf{s})$



Motivation: Forward and inverse problems

Consider a forward problem

$$\mathcal{F} : \mathbf{x} \in \Omega_x \mapsto \mathbf{y} \in \Omega_y, \quad \Omega_x \in \mathbb{R}^{N_x}, \quad \Omega_y \in \mathbb{R}^{N_y}$$

For example the **heat conduction** PDE:

$$\frac{\partial u(\mathbf{s}, t)}{\partial t} - \kappa \Delta u(\mathbf{s}, t) = 0$$


where

$$u(\mathbf{s}, 0) = u_0(\mathbf{s})$$


$u(\mathbf{s}, t)$ → temperature at location \mathbf{s} at time t

$u_0(\mathbf{s})$ → initial temperature at location \mathbf{s}

κ → thermal conductivity of material



Inverse problem \mathcal{F}^{-1} : Given **noisy** $u(\mathbf{s}, T)$ infer $u_0(\mathbf{s})$



Challenges with inverse problems:


- ▶ Inverse map is **not well posed**.
- ▶ **Noisy measurements**.
- ▶ Need to encode **prior knowledge** about x .

Bayesian framework: x and y modelled by random variables X and Y .

AIM: Given a measurement $Y = y$ approximate the conditional (posterior) distribution

$$P_{X|Y}(x|y)$$

and sample from it.



- ▶ Posterior sampling techniques, such as Markov Chain Monte Carlo, are prohibitively expensive when **dimension of X is large**.
- ▶ Characterization of **priors for complex data**

For example, x data might look like:




Representing this data using simple distributions is **hard!**

Resolve both issues using deep learning

Conditional Generative Adversarial Network (cGAN)

- ▶ Proposed by Mirza et al. (2014)
- ▶ Learning distributions conditioned on another field ($P_{X|Y}$)
- ▶ Comprises two neural networks, g and d .



Generator network:

- ▶ $g : \Omega_z \times \Omega_y \rightarrow \Omega_x$.
- ▶ Latent variable $Z \sim P_Z$, e.g. $N(0, I)$. Also $N_z \ll N_x$.
- ▶ (x, y) sampled from true P_{XY}

Critic network:

- ▶ $d : \Omega_x \times \Omega_y \rightarrow \mathbb{R}$.
- ▶ d tries to detect fake samples.
- ▶ $d(x, y)$ large for real x , small otherwise.

- ▶ Objective function

$$\Pi(\mathbf{g}, d) = \mathbb{E}_{\substack{(\mathbf{X}, \mathbf{Y}) \sim P_{\mathbf{X}\mathbf{Y}} \\ \mathbf{Z} \sim P_{\mathbf{Z}}}} [d(\mathbf{X}, \mathbf{Y}) - d(\mathbf{g}(\mathbf{Z}, \mathbf{Y}), \mathbf{Y})]$$

- ▶ Define

$$\text{Lip} = \{f : \Omega_x \times \Omega_y \rightarrow \mathbb{R} \text{ s.t. } f \text{ is 1-Lipschitz in } x \text{ and } y\}$$

* *Solution of physics-based inverse problems using conditional generative adversarial networks with full gradient penalty* (R., Esandi, Dasgupta, Oberai); CMAME, 2023.

- ▶ Objective function

$$\Pi(\mathbf{g}, d) = \mathbb{E}_{\substack{(\mathbf{X}, \mathbf{Y}) \sim P_{\mathbf{X}\mathbf{Y}} \\ \mathbf{Z} \sim P_{\mathbf{Z}}}} [d(\mathbf{X}, \mathbf{Y}) - d(\mathbf{g}(\mathbf{Z}, \mathbf{Y}), \mathbf{Y})]$$

- ▶ Define

$$\text{Lip} = \{f : \Omega_x \times \Omega_y \rightarrow \mathbb{R} \text{ s.t. } f \text{ is 1-Lipschitz in } x \text{ and } y\}$$

- ▶ **Train cGAN:** Find \mathbf{g}^* and d^* by solving the minmax problem

$$d^*(\mathbf{g}) = \arg \max_{d \in \text{Lip}} \Pi(\mathbf{g}, d)$$
$$\mathbf{g}^* = \arg \min_{\mathbf{g}} \Pi(\mathbf{g}, d^*(\mathbf{g})).$$

* *Solution of physics-based inverse problems using conditional generative adversarial networks with full gradient penalty* (R., Esandi, Dasgupta, Oberai); CMAME, 2023.

- ▶ We can show

$$\Pi(\mathbf{g}, d^*(g)) = W_1(P_{\mathbf{XY}}, P_{\mathbf{XY}}^g)$$

where $P_{\mathbf{XY}}^g = P_{\mathbf{X}|Y}^g P_Y$. Thus,

$$\mathbf{g}^* = \arg \min_g W_1(P_{\mathbf{XY}}, P_{\mathbf{XY}}^g)$$

- ▶ We can show

$$\Pi(\mathbf{g}, d^*(\mathbf{g})) = W_1(P_{\mathbf{XY}}, P_{\mathbf{XY}}^{\mathbf{g}})$$

where $P_{\mathbf{XY}}^{\mathbf{g}} = P_{\mathbf{X}|\mathbf{Y}}^{\mathbf{g}} P_{\mathbf{Y}}$. Thus,

$$\mathbf{g}^* = \arg \min_{\mathbf{g}} W_1(P_{\mathbf{XY}}, P_{\mathbf{XY}}^{\mathbf{g}})$$

- ▶ If we can find the sequence $\{(\mathbf{g}_n^*, d_n^*)\}_n$ such that

$$\lim_{n \rightarrow \infty} \Pi(\mathbf{g}_n^*, d_n^*) = \lim_{n \rightarrow \infty} W_1(P_{\mathbf{XY}}, P_{\mathbf{XY}}^{\mathbf{g}_n^*}) = 0,$$

then we weakly converge to the true joint density, $P_{\mathbf{XY}}^{\mathbf{g}_n^*} \xrightarrow{weak} P_{\mathbf{XY}}$

$$\iff \lim_{n \rightarrow \infty} \mathbb{E}_{P_{\mathbf{XY}}^{\mathbf{g}_n^*}} [\ell(\mathbf{X}, \mathbf{Y})] = \mathbb{E}_{P_{\mathbf{XY}}} [\ell(\mathbf{X}, \mathbf{Y})] \quad \forall \ell \in C_b(\Omega_x \times \Omega_y).$$

But our goal is to approximate the conditional density, or rather estimate the conditional expectations

$$\mathbb{E}_{P_{\mathbf{X}|\mathbf{Y}}} [q(\mathbf{X})|\mathbf{y}].$$

But our goal is to approximate the conditional density, or rather estimate the conditional expectations

$$\mathbb{E}_{P_{\mathbf{X}|\mathbf{Y}}} [q(\mathbf{X})|\mathbf{y}].$$

The following result shows us how to do this **robustly**:

Theorem: R., Esandi, Dasgupta, Oberai (2023)

Let $\hat{\mathbf{y}}$ be such that $P_{\mathbf{Y}}(\hat{\mathbf{y}}) \neq 0$ and $q \in C_b(\Omega_x)$. Then, given $\epsilon > 0$ (and under some mild assumptions), there exists $\sigma > 0$ and an integer N such that

$$\left| \mathbb{E}_{P_{\mathbf{X}|\mathbf{Y}}} [q(\mathbf{X})|\hat{\mathbf{y}}] - \mathbb{E}_{P_{\mathbf{X}\mathbf{Y}}^{\sigma,n}} [q(\mathbf{X})] \right| < \epsilon \quad \forall n \geq N,$$

where $P_{\mathbf{X}\mathbf{Y}}^{\sigma,n}(\mathbf{x}, \mathbf{y}) = P_{\mathbf{X}|\mathbf{Y}}^{g_n^*}(\mathbf{x}|\mathbf{y})P_{\mathbf{Y}\sigma}(\mathbf{y})$ and $P_{\mathbf{Y}\sigma}(\mathbf{y}) \equiv N(\hat{\mathbf{y}}, \sigma^2 \mathbf{I})$.


Implication: Instead of feeding the measurement $\hat{\mathbf{y}}$ to g^* , feed $\hat{\mathbf{y}} + \delta_{\mathbf{y}}$ where $\delta_{\mathbf{y}} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$

Approximating posterior expectations

Steps:

- ▶ Acquire samples $\mathcal{S}_x = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, where $\mathbf{x}_i \sim P_X^{\text{prior}}$.
- ▶ Use forward map \mathcal{F} to generate dataset $\mathcal{S} = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N)\}$
- ▶ Train a cWGAN on \mathcal{S} .
- ▶ For a test $\hat{\mathbf{y}}$, generate samples by passing \mathbf{z} and “perturbed” \mathbf{y} samples through g^* .
- ▶ Approximate expectation using Monte Carlo.

$$\mathbb{E}_{\mathbf{X} \sim P_{\mathbf{X}|\mathbf{Y}}}[\ell(\mathbf{X})] \approx \frac{1}{K} \sum_{i=1}^K \ell(g^*(\mathbf{z}^{(i)}, \mathbf{y}^{(i)})), \quad \mathbf{z}^{(i)} \sim P_{\mathbf{Z}}, \quad \mathbf{y}^{(i)} \sim N(\hat{\mathbf{y}}, \sigma^2 \mathbf{I})$$




Consider the pair of 1D random variables defined by:

Tanh + Γ : $x = \tanh(y) + \gamma$ where $\gamma \sim \Gamma(1, 0.3)$ and $y \sim U(-2, 2)$


Bimodal : $x = (y + w)^{1/3}$ where $y \sim \mathcal{N}(0, 1)$ and $w \sim \mathcal{N}(0, 1)$

Swissroll : $x = 0.1t \sin(t) + 0.1w$, $y = 0.1t \cos(t) + 0.1v$, $t = 3\pi/2(1 + 2h)$,
where $h \sim U(0, 1)$, $w \sim \mathcal{N}(0, 1)$ and $v \sim \mathcal{N}(0, 1)$

True joints



Simple 1D problems




Solving the inverse heat conduction equation

Goal: Infer initial temp. field from noisy final temp. field

Assuming x to given by a rectangular inclusion and $N_x = N_y = 28 \times 28 = 784$


Training samples:



cWGANs trained using latent dimension $N_z = 3!$

Solving the inverse heat conduction equation



Test sample, whose reference mean and SD are available



SD map can be used to **quantify uncertainty locally**

Solving the inverse heat conduction equation

L^2 error in mean and SD




- ▶ Substantial increase in wildfire activity around the globe.
- ▶ Complicated physics coupling atmosphere and wildfire dynamics.
- ▶ Correct initial state of wildfire and atmosphere variables required for successful simulations.
- ▶ Mandel et al. (2012) found that
 - Precise wildfire history during initial spread – key for model initialization.
 - History well represented by arrival time map.

Data assimilation problem: Given satellite measurements of active fire during initial spread, determine high resolution fire arrival map for initial period.

* *Generative Algorithms for Fusion of Physics-Based Wildfire Spread Models with Satellite Data for Initializing Wildfire Forecasts* (Shaddy, R., Faruell, Calaza, Mandel, Haley, Hilburn, Mallia, Kochanski, Oberai); preprint on arXiv, 2023.

Predicting arrival times for wildfire spread


- ▶ Synthetic training data generated using **WRF-SFIRE**.
- ▶ $N_x = N_y = 512 \times 512 = 262,144$.
- ▶ Trained cWGAN $N_z = 100$.



Fire arrived first in the darkest regions of the plot

Predicting arrival times for wildfire spread

- ▶ Tested on real wildfire data for fires in California between 2020 - 2022.
- ▶ Data collected from Suomi-NPP satellite, detections 2-4 times a day.
- ▶ High confidence measurements (top row); high+nominal confidence measurements (bottom row)




IR Time: Number of hours since start of fire.

Predicting arrival times for wildfire spread

- ▶ 200 realization for each type of input measurement.
- ▶ Weighted combination of realizations

$$\mathbf{x}_i = 0.2 \times \mathbf{x}_i^{\text{high}} + 0.8 \times \mathbf{x}_i^{\text{high+nom}}$$

used to compute pixel-wise mean and SD



Estimate ignition time based on smallest arrival time compared with California Department of Forestry and Fire Protection (CAL FIRE) reporting and another SVM based method by Farguell et al. (2021).


Wildfire	CAL FIRE	cWGAN	SVM	cWGAN Offset	SVM Offset
Tennant	23:07	23 : 48	21:11	41 minutes	1 hour 56 minutes
Oak	21:10	21 : 30	20:45	20 minutes	25 minutes
Mineral	23:40	23 : 04	27:53	36 minutes	4 hours 13 minutes

See preprint for additional details and comparisons (eg. F-score, false alarm ratio, etc)

- ▶ A cWGAN algorithm for Bayesian inference
- ▶ What do we gain?
 - Ability to represent and encode complex prior data.
 - Dimension reduction since $N_z \ll N_x$.
 - Sampling from cGAN is quick and easy.
 - Uncertainty quantification in terms of SD.
- ▶ Need (x, y) pairs to train – supervised algorithm.
- ▶ Convergence theory that leads to a robust sampling algorithm.
- ▶ Currently testing algorithm on several other physics-based and medical applications.

Questions?

Errors between $P_{X|Y}$ and $P_{X|Y}^g$:



Can expect benefit of Full-GP approach on multi-modal problems!

- ▶ 20 simulations in total
- ▶ 2 day fire spread over flat terrain with a uniform fuel type of brush.
- ▶ Initial wind profile: logarithmic profile up to 2 km, constant wind speed above 2 km, uniform prescription in one direction
- ▶ Wind magnitude 10m from the surface varied randomly from a uniform distribution between $0 - 5 \text{ m s}^{-1}$.

Strategy:

- ▶ Generated 20 fire simulations using WRF-SFIRE
- ▶ Data augmented by rotations and translations to generate 10,000 high-resolution arrival maps $\mathbf{x}_i \in \mathbb{R}^{512 \times 512}$
- ▶ Corresponding measurements $\mathbf{y}_i \in \mathbb{R}^{512 \times 512}$ obtained by coarsening and occluding.
- ▶ 8000 training samples, 2000 validation samples (to tune hyper-parameters)