Solving High-Dimensional Inverse Problems with Auxiliary Uncertainty via Operator Learning with Limited Data

Joseph Hart, **Mamikon Gulian**, Indu Manickam, and Laura Swiler

Exceptional service in the national interest

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND number: SAND2023-02512C.

Source Inversion / Identification in Climate Systems

- The identification of sources in climate systems is vital for **attribution** and prediction, which inform policy decisions.
- Inability to isolate sources while observing the climate, and the cost of simulating computational models is high.
- **Surrogate models** enable the many-query algorithms required for source identification, but challenges arise due to:
  - High dimensionality of the state and source;
  - Limited ensembles of costly model simulations to train a surrogate model, and
  - Few and potentially noisy state observations for inversion due to measurement limitations.
- The influence of **auxiliary processes** adds an additional layer of uncertainty that further confounds source identification.
- Inversion for a source is an ill-posed problem. It is natural to state the problem in a probabilistic way.
- We develop a **Bayesian method** that predicts the most probable source and quantifies uncertainty in the predicted source.

Summary of Method, Exemplar, and Results

- We introduce a framework based on
  1. Calibrating **deep neural network** surrogates to the flow maps provided by an ensemble of simulations obtained by varying sources, then
  2. Using these surrogates in a Bayesian framework to identify sources from observations via **optimization**.
- We focus on an **atmospheric dispersion exemplar** in which the source represents injection of SO_2 and we observe concentratation.
- The expressive and computationally efficient nature of the deep neural network operator surrogates in **reduced dimension** allows for source identification with uncertainty quantification using limited data.
- To stress the algorithm, we then introduce a variable wind field as an auxiliary process that **confounds** source inversion.
- A Bayesian approximation error (BAE) approach becomes essential for reliable inversion.
- Derivative-based optimization and sampling algorithms leverage algorithmic differentiation tools to efficiently generate approximate samples from the posterior.

Problem Formulation and Solution Outline

- We consider time-dependent models of the form

$$\frac{\partial u}{\partial t} + \mathcal{A}(u, z, w) = f(z, w) \qquad \text{on } \Omega \times (0, \infty) \tag{1}$$

$$\mathcal{B}(u, z, w) = 0 \qquad \text{on } \partial\Omega \times (0, \infty) \tag{2}$$

$$u = u_0 \qquad \text{on } \Omega \times \{0\} \tag{3}$$

  where $u : \Omega \times [0, \infty) \to \mathbb{R}$ represents a state defined for time $t \geq 0$ on a domain Ω.

- We seek to **estimate** z by solving an **inverse problem** which combines the computational model (1) with sparse and noisy observations of the state variable u.
- The other parameter w, sometimes referred to as a **nuisance parameter**, is also assumed to be uncertain but is not of primary interest; it represents an auxiliary process.
- We assume that the computational model (1) may be solved for any z and w, and that we have access to an **ensemble** of such simulations.
- We utilize a DNN approximation of the **flow map** operator of the model (1), mapping the state and parameters at time t to the state at time $t + \Delta t$.

Flow Map & Ensemble of Simulation Data for Training and Validation

- We work with a **discretization** of the PDE model (1). Let

$$t_n, \quad n = 0, 1, 2, \ldots, N \tag{4}$$

  denote increasing points in time, with $t_0 = 0$ and $t_N = T$.

- We seek to build a surrogate for the flow map, defined as the function

$$\mathcal{F} : \mathbb{R}^m \times \mathbb{R}^s \times \mathbb{R}^q \to \mathbb{R}^m$$

  which **evolves the state** from time t_n to t_{n+1}, i.e.

$$\mathbf{u}_{n+1} = \mathcal{F}(\mathbf{u}_n, \mathbf{z}_n, \mathbf{w}_n). \tag{5}$$

- Assume that M samples of $\mathbf{z}$ and $\mathbf{w}$ are given. Denote them as

$$\mathbf{z}_n^i \in \mathbb{R}^s, \qquad \mathbf{w}_n^i \in \mathbb{R}^q, \qquad i = 1, 2, \ldots, M, \qquad n = 0, 1, \ldots, N.$$

- Integrating the model (1) with these parameters yields a set of M trajectories

$$\left\{\mathbf{u}_n^i\right\}_{n=0}^N \subset \mathbb{R}^m, \quad i = 1, 2, \ldots, M.$$

Spatial Dimension Reduction of State and Parameters

- The **high dimensionality** of $\mathbf{u}_n^i \in \mathbb{R}^m$ poses significant challenges to learning $\mathcal{F}$.
- We obtain a lower-dimensional linear subspace using the PCA/POD/EOF method. We compute the singular value decomposition $\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ of the matrix

$$\mathbf{Y} = \begin{pmatrix} \mathbf{u}_0^1 & \mathbf{u}_1^1 & \cdots & \mathbf{u}_N^1 & \mathbf{u}_0^2 & \cdots & \mathbf{u}_N^2 & \cdots & \mathbf{u}_0^M & \cdots & \mathbf{u}_N^M \end{pmatrix} \in \mathbb{R}^{m\times(N+1)M}.$$

- The idea of our approach is that if the state trajectory $\{\mathbf{u}_n\}_{n=0}^N$ is approximately contained in the range of $\mathbf{U}_r$ then the following diagram approximately commutes:

$$\begin{array}{ccc} \mathbf{u}_n & \xrightarrow{\mathcal{F}(\cdot,\mathbf{z}_n,\mathbf{w}_n)} & \mathbf{u}_{n+1} \\ \text{PCA projection} = \mathbf{U}_r^\top \downarrow & & \uparrow \text{PCA reconstruction} = \mathbf{U}_r \\ \mathbf{c}_n & \xrightarrow[\mathcal{F}_r(\cdot,\mathbf{z}_n,\mathbf{w}_n)]{} & \mathbf{c}_{n+1} \end{array}$$

- We seek to learn a DNN approximation of $\mathcal{F}_r$ using the MN training data pairs

$$\left\{(\mathbf{c}_n^i, \mathbf{z}_n^i, \mathbf{w}_n^i), \mathbf{c}_{n+1}^i\right\} \qquad n = 0, 1, \ldots, N-1, \qquad i = 1, 2, \ldots, M,$$

Flow Map Approximation Architecture and Training

- Consider DNN approximations $\mathcal{N} \approx \mathcal{F}_r$ of the form

$$\mathcal{N} : (\mathbf{c}, \mathbf{z}, \mathbf{w}; \xi) \mapsto \mathbf{c} + \Delta t \mathcal{E}(\mathbf{c}, \mathbf{z}, \mathbf{w}; \xi)$$

  where $\mathcal{E}(\mathbf{c}, \mathbf{z}, \mathbf{w}; \xi)$ is a **dense feedforward DNN** with weights and biases $\xi \in \mathbb{R}^\ell$.

- The surrogate must provide stable and accurate evolution of the state starting from $\mathbf{u}_0$, through many timesteps.

- We train with the **loss function** that includes **repeated compositions** of $\mathcal{N}$:

$$\mathcal{L}(\xi) = \sum_{i=1}^{M} \sum_{n=0}^{N-1} \sum_{p=1}^{P(n)} \| \mathbf{c}_{n+p}^i - \mathcal{N}^{[p]}(\mathbf{c}_n^i, \{\mathbf{z}_j^i, \mathbf{w}_j^i\}_{j=n}^{n+p-1}, \xi) \|_{\ell^2}^2 \tag{6}$$

  where $\mathcal{N}^{[p]}(\mathbf{c}_n^i, \{\mathbf{z}_j^i, \mathbf{w}_j^i\}_{j=n}^{n+p-1}, \xi)$ denotes the composition of $\mathcal{N}$ with itself p times.

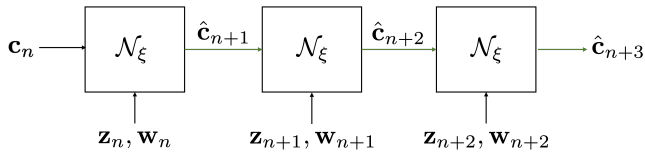

The Inverse Problem

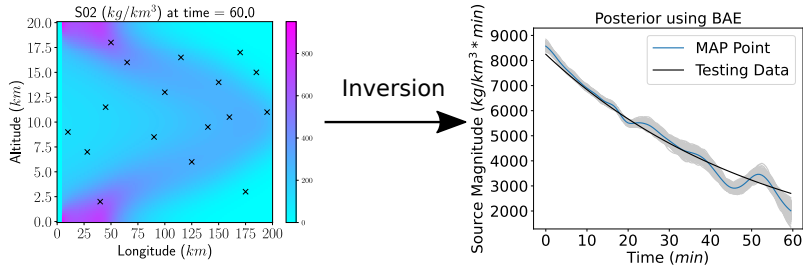

- Our ultimate goal is solving an inverse problem to estimate $\mathbf{z}$ given sparse and noisy measurements of u.
- Let $\mathbf{d}_n \in \mathbb{R}^L$, $n = 1, 2, \ldots, N$, denote observations of the state u at L spatial locations at each n^{th} time step.
- Let $\mathcal{O} : \mathbb{R}^m \to \mathbb{R}^L$ denote the observation operator where $\mathbf{d}_n = \mathcal{O}(\mathbf{u}_n) + \boldsymbol{\epsilon}_n$, for noise vectors $\boldsymbol{\epsilon}_n$, $n = 1, 2, \ldots, N$.
- We model $\{\boldsymbol{\epsilon}_n\}_{n=1}^N$ as independent identically distributed random vectors which follow a mean zero Gaussian distribution.

Bayesian Posterior Distribution

- We assume some **prior knowledge** of possible $\mathbf{z}$'s in the form of a Gaussian prior distribution $\pi_{\text{prior}}(\mathbf{z})$ with mean $\bar{\mathbf{z}}$ and covariance $\mathbf{\Gamma}_{\text{prior}}$.
- With the assumption of additive Gaussian noise ϵ **contaminating** the observed data

$$\mathbf{d} = (\mathbf{d}_1, \mathbf{d}_2, \ldots, \mathbf{d}_N) \in \mathbb{R}^{LN}, \tag{7}$$

  we have a Gaussian **likelihood** $\pi_{\text{like}}(\mathbf{d}|\mathbf{z})$.
- Bayes' Theorem gives the **posterior distribution** of $\mathbf{z}$ as

$$\pi_{\text{post}}(\mathbf{z}) \propto \pi_{\text{prior}}(\mathbf{z})\pi_{\text{like}}(\mathbf{d}|\mathbf{z}). \tag{8}$$

- The point of greatest posterior probability is called the maximum a posteriori probably (MAP) point and gives a **best estimate** of $\mathbf{z}$.
- Drawing samples from the posterior distribution **quantifies uncertainty** about the MAP due to limitations on data availability and quality.

Parameter-to-Observable Map and Bayesian Approximation Error

- Our surrogate model for the mapping from $(\mathbf{z}, \mathbf{w})$ to the observations of the state at time step t_n **given** $\mathbf{u}_0$ is

$$\mathbf{F}_n(\mathbf{z}, \mathbf{w}) = \mathcal{O}\left(\mathbf{U}_r \mathcal{N}^{[n]}(\mathbf{U}_r^\top \mathbf{u}_0, \{\mathbf{z}_j, \mathbf{w}_j\}_{j=0}^n, \xi)\right). \tag{9}$$

- The **parameter-to-observable map** is defined by concatenating $\mathbf{F}_n(\mathbf{z}, \mathbf{w})$ for all time steps,

$$\mathbf{F}(\mathbf{z}, \mathbf{w}) = [\,\mathbf{F}_1(\mathbf{z}, \mathbf{w}),\, \mathbf{F}_2(\mathbf{z}, \mathbf{w}),\, \ldots,\, \mathbf{F}_N(\mathbf{z}, \mathbf{w})\,] \in \mathbb{R}^{LN}.$$

- The Bayesian posterior implicitly depends on $\mathbf{w}$ which is uncertain. We assume that a probabilistic model (for which samples may be computed) for $\mathbf{w}$ is given.
- Merely fixing $\mathbf{w} = \overline{\mathbf{w}}$ to its best estimate and solving the inverse problem for $\mathbf{z}$ using $\mathbf{F}(\mathbf{z}, \overline{\mathbf{w}})$ fails to incorporate uncertainty in $\mathbf{w}$ and **may provide poor UQ**.
- In the Bayesian approximation error (BAE) approach, we model error about $\mathbf{F}(\mathbf{z}, \text{true } \mathbf{w})$ from both noise and misspecification of $\mathbf{w}$ using an empirical Gaussian.
- Rather than a noise model with mean zero and covariance $\mathbf{\Gamma}_{\text{noise}}$, we use the noise model ν with mean $\overline{\mathbf{e}}$ and covariance $\mathbf{\Gamma}_{\text{BAE}} = \mathbf{\Gamma}_{\text{noise}} + \mathbf{\Gamma}_{\mathbf{e}}$.

Maximum *a Posteriori* (MAP) Point Estimation and Posterior Sampling

- To determine the MAP point, we formulate an optimization problem to maximize the posterior PDF given observations:

$$\min_{\mathbf{z}} \left\{ J(\mathbf{z}) = \frac{1}{2} \sum_{n=1}^{N} \|\mathbf{F}_n(\mathbf{z}) + \overline{\mathbf{e}} - \mathbf{d}_n\|^2_{\mathbf{\Gamma}^{-1}_{\text{BAE}}} + \frac{1}{2} \|\mathbf{z} - \overline{\mathbf{z}}\|^2_{\mathbf{\Gamma}^{-1}_{\text{prior}}} \right\}. \tag{10}$$

- We utilized **derivative-based algorithms** in the Rapid Optimization Library (**ROL**), part of the Trilinos package provided by Sandia National Laboratories.
- To obtain required derivatives, we leveraged algorithmic, or automatic, differentiation at each time step when evaluating the flow map approximation.
- To rapidly quantify uncertainty, we compute samples from a Laplace approximation of the posterior.
- This assumes that $\mathbf{F}_n(\mathbf{z})$ is linear in $\mathbf{z}$, which is valid in a neighborhood of the MAP point.

Numerical Example: Atmospheric Aerosol (SO_2) Dispersion Model

- Data is generated from the advection-diffusion-reaction PDE

$$\begin{aligned} &\frac{\partial u}{\partial t} - \kappa \nabla^2 u + \mathbf{v}(w) \cdot \nabla u - S\mathbf{e}_y \cdot \nabla u = \mathcal{R}(u) + f(z) && \text{on } \Omega \times (0, \infty) \\ &\nabla u \cdot \mathbf{n} = 0 && \text{on } \partial\Omega \times (0, \infty) \\ &u = 0 && \text{on } \Omega \times \{0\} \end{aligned} \tag{11}$$

- The source is defined as $f(z) = z(t)F(x, y)$ where $z : [0, T] \to \mathbb{R}$ is the time varying source magnitude being inferred and

$$F(x, y) = \exp\left(-100(x - 5)^2\right) \exp\left(-0.1(y - 9)^2\right).$$

- $\mathbf{z}$ is a vector with length being the number of time steps N in the discretized data.

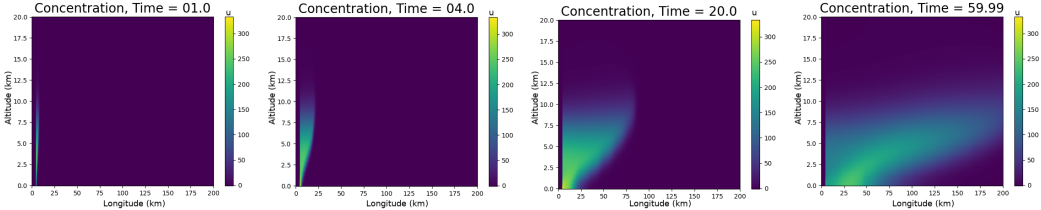

Training & Test Data Generation: 4 Source Magnitudes × 4 Wind Fields

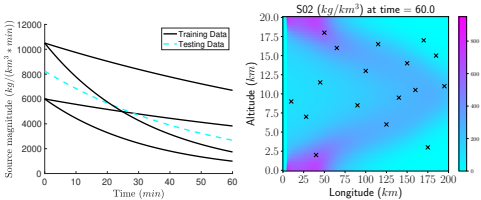

Figure 1: We generate a dataset by solving the model for $M = 16$ combinations. The discretized data has $N = 120$ time nodes and $m = 101101$ spatial nodes.

- $\mathbf{z}$ is represented by a vector in $\mathbb{R}^{120}$ and $\mathbf{u}, \mathbf{w}$ are represented by a matrix in $\mathbb{R}^{101101 \times 120}$.
- To test inversion, we generated a test dataset "in the middle" of the training data.
- The relative ℓ^2 distance between the test and mean wind field is 3.6% and the relative ℓ^2 distances between the test and training wind fields ranges from 4.7% to 6.5%.
- We contaminate the state data with multiplicative Gaussian noise whose mean is 1 and standard deviation is 0.02, i.e. 2% of the data magnitude.
- To mimic realistic uncertainties, we invert with the mean wind field $\overline{\mathbf{w}}$ as our best estimate of the wind despite test data being generated with the testing wind field $\mathbf{w}^\star$.

Building the Surrogate Model: Hyperparameters, Training, and Tuning

- We **compress the state** with a rank $r = 70$ PCA representation, with maximum relative ℓ_2 reconstruction error on the validation set of 0.6%.
- We also compress the wind fields using $r = 10$ modes with 0.7% reconstruction error.
- Since the source magnitude is scalar, the **reduced flow map** $\mathcal{F}_r$ maps $\mathbb{R}^{81}$ into $\mathbb{R}^{70}$.
- We divided our ensemble into a **training dataset** consisting of data from 12 PDE solves and a **validation dataset** consisting of 4 PDE solves.
- The validation set is used to **tune the flow map architecture**. For all experiments, we used the ADAM optimizer and ELU nonlinear activation layers.
- We examined the impact of increasing P, the number of DNN compositions used in the loss function (6) which increases the network accuracy as well as **stability**, and selected $P = 25$
- A **hyperparameter study** suggested optimal width 200, depth 2, and learning rate 0.0008.

Inversion Results: Improvement over Prior, and Importance of BAE

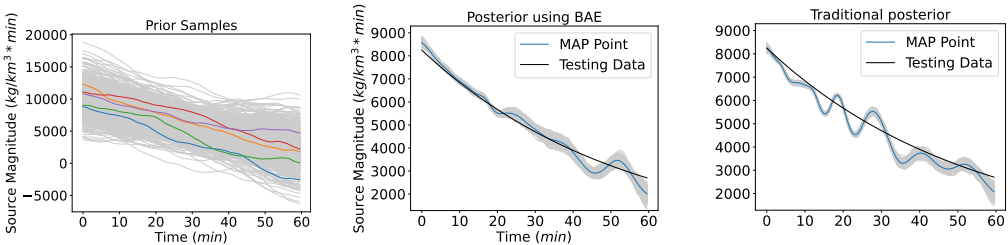

- We observe that the posterior in the traditional Bayesian approach has oscillations due to modeling error which is not characteristic of the testing data.
- It also has a small variance about this solution indicating **confidence in an erroneous MAP point**.
- This highlights the importance of **BAE formulation**, which largely avoided the oscillations and overconfidence seen in the traditional posterior.
- This also highlights the utility of **incorporating the wind**, the auxiliary source of uncertainty, into our flow map training so it could be properly handed in inversion.

Testing with Greater Wind Variability

- In a second study, the **wind field variability** (or range of uncertainty) is **increased** to understand limitations of the proposed approach.
- The relative ℓ_2 differences between the training samples and mean wind range from 9.7% to 10.8% (recall a range of 3.3% to 3.7% in the lesser wind variability case).
- With increased wind variability, we observe **greater error** in the flow map approximation and inversion, as is expected.
- The quality of the result as the posterior shown in Figure 16 is based on using **only** $M = 16$ **PDE solves** to generate training data for this more difficult scenario.

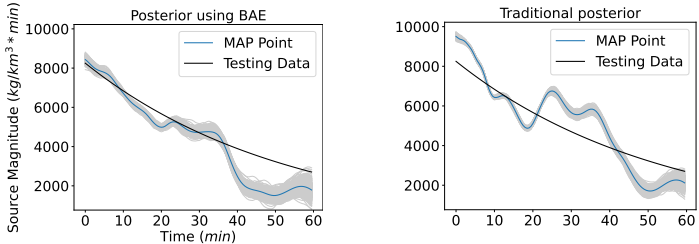

Comparing BAE vs Traditional Posterior using Relative ℓ^2 and Mahalanobis Metrics

- The Mahalanobis distance between the posterior distribution and the test source is a commonly used metric to measure distance between a **distribution** and a **point**.
- The much larger Mahalanobis distance for the traditional posterior demonstrates the **significant advantage of BAE** to provide uncertainty quantification.
- This is crucial in **limited data settings**.

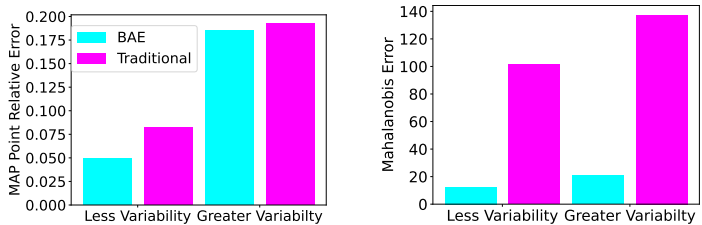

Figure 2: Bar plots comparing the error estimating the test data source. *Left panel:* the relative ℓ^2 error between the MAP point and the test data source. *Right panel:* the Mahalanobis distance between the posterior and the test data source.

Conclusions

- Our approach combines learning **data-driven surrogate operators** parametrized by DNNs with **Bayesian approximation error** for inversion.
- We leveraged **algorithmic differentiation** to enable optimization and approximate posterior sampling without the need for extensive tuning to achieve convergence.
- Rapid source inversion with **UQ** when **confounded** by auxiliary unobserved processes, high state dimension, and small ensembles of training data.
- We **do not require intrusion** to obtain derivative information.
- Method was tested and analyzed in a **setting that mimics the challenges** faced when working with earth system model (ESM) and satellite observation data.

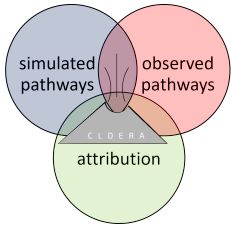

Figure 3: This work was supported by **CLDERA** (**CL**imate impact: **D**etermining **E**tiology th**R**ough p**A**thways) GC LDRD at Sandia National Laboratories. It will be published in a forthcoming article in the Journal of Machine Learning for Modeling and Computing.