Estimating the Error in Solutions to Stochastic Inverse Problems When Using Machine Learning Surrogates

Tim Wildey

Sandia National Laboratories
Center for Computing Research
Scientific Machine Learning Department

SIAM Conference on the Mathematics of Data Science
September 25-30, 2022

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-NA-0003525. SAND2022-****

boratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LL
Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE

Data-informed Physics-Based Predictions

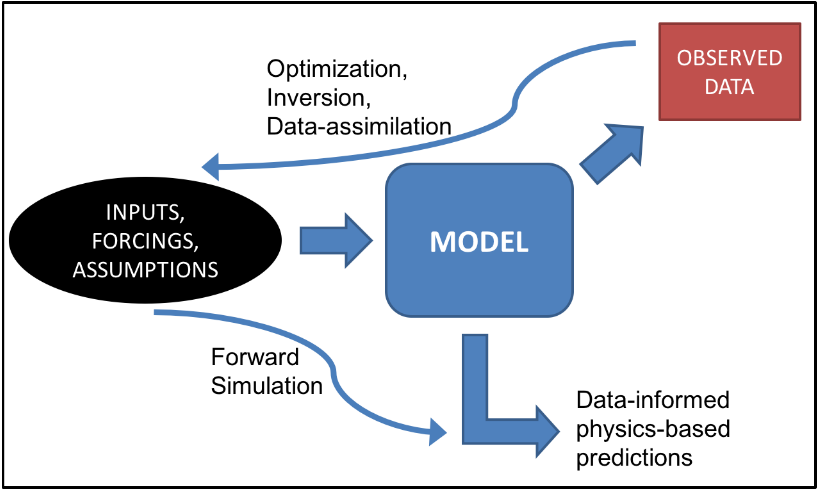

Data-informed Physics-Based Predictions

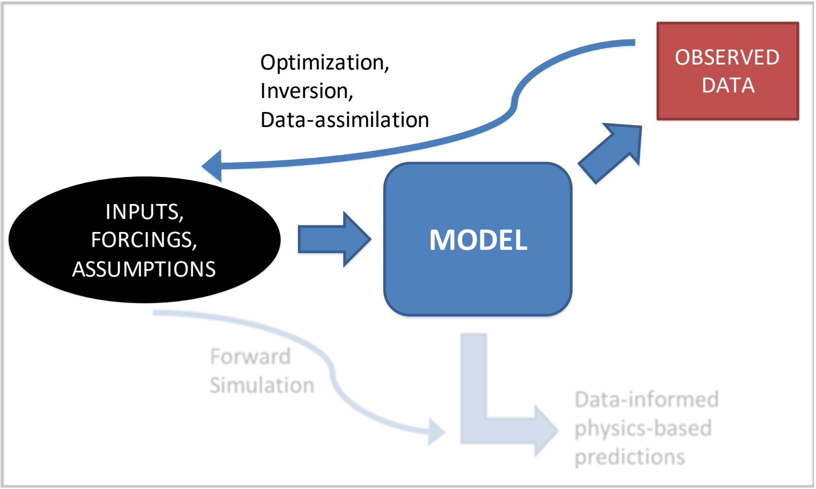

Data-informed Physics-Based Predictions

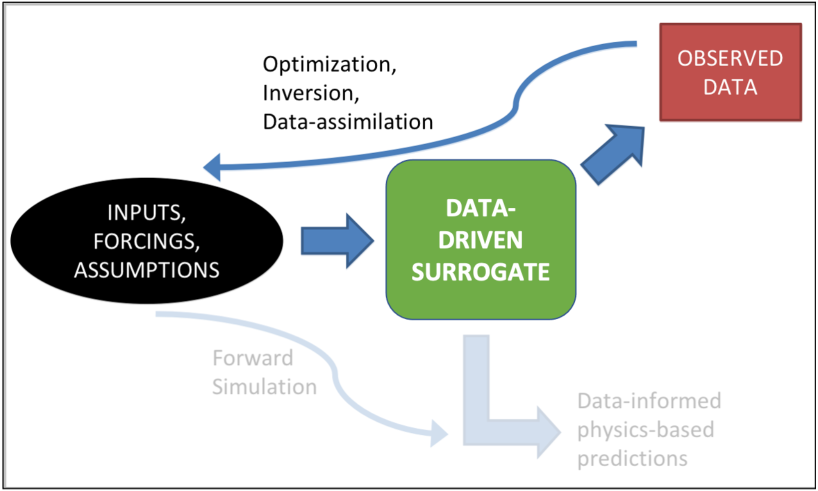

A Deterministic Inverse Problem

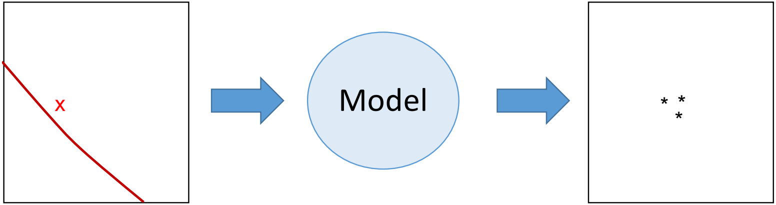

Problem

Given some observed data, find $\lambda \in \Lambda$ that best predicts the data.

A Deterministic Inverse Problem

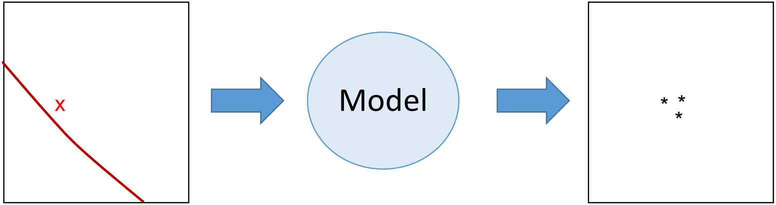

Problem

Given some observed data, find $\lambda \in \Lambda$ that best predicts the data.

- Solutions may not be unique without additional assumptions.
- Requires solving several deterministic forward problems.

A Stochastic Inverse Problem

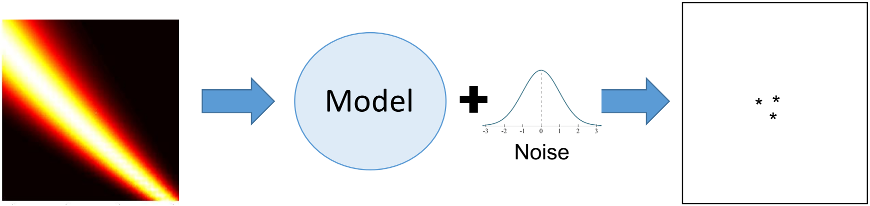

Problem

Given some observed data and an assumed noise model, find the parameters that are most likely to have produced the data.

A Stochastic Inverse Problem

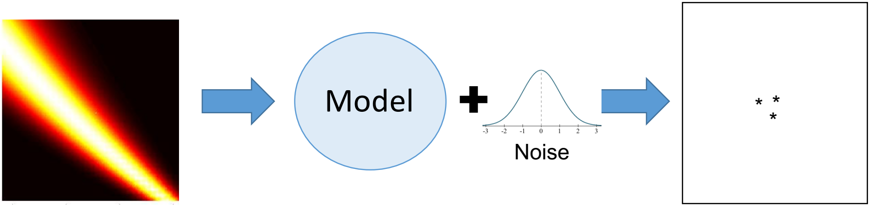

Problem

Given some observed data and an assumed noise model, find the parameters that are most likely to have produced the data.

- Solutions may not be unique without additional assumptions.
- Requires solving several deterministic forward problems.

A Different Stochastic Inverse Problem

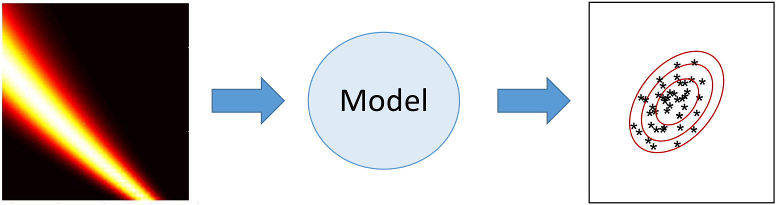

Problem

Given a probability density on observations, find a probability density on Λ such that the push-forward matches the given density on the observed data.

A Different Stochastic Inverse Problem

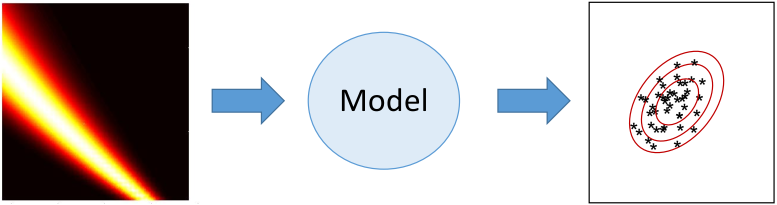

Problem

Given a probability density on observations, find a probability density on Λ such that the push-forward matches the given density on the observed data.

- Solutions may not be unique without additional assumptions.
- **We only need to solve a single stochastic forward problem.**

Notation

We assume we are given:

1. A finite-dimensional **parameter space**, Λ.
2. A **parameter-to-observation/data map**, $Q : \Lambda \to \mathcal{D} = Q(\Lambda)$
3. A **observed/target probability measure** on $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$, denoted $\mathbb{P}_{\mathcal{D}}^{\text{obs}}$, with density $\pi_{\mathcal{D}}^{\text{obs}}$ (typically from experimental data)
4. An **initial probability measure** on $(\Lambda, \mathcal{B}_{\Lambda})$, denoted $\mathbb{P}_{\Lambda}^{\text{init}}$, with density $\pi_{\Lambda}^{\text{init}}$ (typically from prior beliefs or expert knowledge)

Notation

We assume we are given:

1. A finite-dimensional **parameter space**, Λ.
2. A **parameter-to-observation/data map**, $Q : \Lambda \to \mathcal{D} = Q(\Lambda)$
3. A **observed/target probability measure** on $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$, denoted $\mathbb{P}_{\mathcal{D}}^{\text{obs}}$, with density $\pi_{\mathcal{D}}^{\text{obs}}$ (typically from experimental data)
4. An **initial probability measure** on $(\Lambda, \mathcal{B}_{\Lambda})$, denoted $\mathbb{P}_{\Lambda}^{\text{init}}$, with density $\pi_{\Lambda}^{\text{init}}$ (typically from prior beliefs or expert knowledge)

We need to compute:

1. The **push-forward of the initial density** through the model.

- In other words, **we need to solve a forward UQ problem using the initial.**
- We use $\pi_{\mathcal{D}}^{\text{pred}}$ to denote this push-forward density.

A Key Assumption

Predictability Assumption

We assume that the observed probability measure, $\mathbb{P}_{\mathcal{D}}^{\text{obs}}$, is absolutely continuous with respect to the push-forward of the initial, $\mathbb{P}_{\mathcal{D}}^{\text{pred}}$.

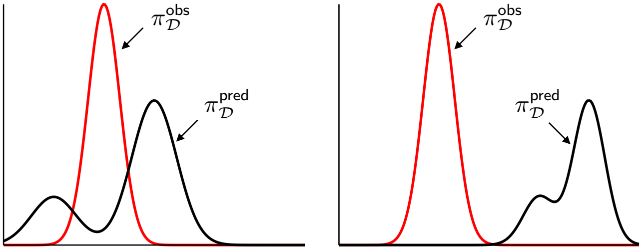

Good Initial

Bad Initial
(Cannot predict all observations)

A Solution to the Stochastic Inverse Problem

Theorem

Given an initial probability measure, $\mathbb{P}_\Lambda^{init}$ *on* $(\Lambda, \mathcal{B}_\Lambda)$ *and an observed probability measure,* $\mathbb{P}_\mathcal{D}^{obs}$*, on* $(\mathcal{D}, \mathcal{B}_\mathcal{D})$*, the probability measure* $\mathbb{P}_\Lambda^{up}$ *on* $(\Lambda, \mathcal{B}_\Lambda)$ *defined by*

$$\mathbb{P}_\Lambda^{up}(A) = \int_\mathcal{D} \left(\int_{A \cap Q^{-1}(q)} \pi_\Lambda^{init}(\lambda) \frac{\pi_\mathcal{D}^{obs}(Q(\lambda))}{\pi_\mathcal{D}^{pred}(Q(\lambda))} d\mu_{\Lambda,q}(\lambda) \right) d\mu_\mathcal{D}(q), \; \forall A \in \mathcal{B}_\Lambda$$

solves the stochastic inverse problem.

A Solution to the Stochastic Inverse Problem

Theorem

Given an initial probability measure, $\mathbb{P}_\Lambda^{init}$ on $(\Lambda, \mathcal{B}_\Lambda)$ and an observed probability measure, $\mathbb{P}_\mathcal{D}^{obs}$, on $(\mathcal{D}, \mathcal{B}_\mathcal{D})$, the probability measure $\mathbb{P}_\Lambda^{up}$ on $(\Lambda, \mathcal{B}_\Lambda)$ defined by

$$\mathbb{P}_\Lambda^{up}(A) = \int_\mathcal{D} \left(\int_{A \cap Q^{-1}(q)} \pi_\Lambda^{init}(\lambda) \frac{\pi_\mathcal{D}^{obs}(Q(\lambda))}{\pi_\mathcal{D}^{pred}(Q(\lambda))} d\mu_{\Lambda,q}(\lambda) \right) d\mu_\mathcal{D}(q), \ \forall A \in \mathcal{B}_\Lambda$$

solves the stochastic inverse problem.

Corollary

The updated measure of Λ is 1.

A Solution to the Stochastic Inverse Problem

Theorem

Given an initial probability measure, $\mathbb{P}_\Lambda^{init}$ on $(\Lambda, \mathcal{B}_\Lambda)$ and an observed probability measure, $\mathbb{P}_\mathcal{D}^{obs}$, on $(\mathcal{D}, \mathcal{B}_\mathcal{D})$, the probability measure $\mathbb{P}_\Lambda^{up}$ on $(\Lambda, \mathcal{B}_\Lambda)$ defined by

$$\mathbb{P}_\Lambda^{up}(A) = \int_\mathcal{D} \left(\int_{A \cap Q^{-1}(q)} \pi_\Lambda^{init}(\lambda) \frac{\pi_\mathcal{D}^{obs}(Q(\lambda))}{\pi_\mathcal{D}^{pred}(Q(\lambda))} d\mu_{\Lambda,q}(\lambda) \right) d\mu_\mathcal{D}(q), \ \forall A \in \mathcal{B}_\Lambda$$

solves the stochastic inverse problem.

Corollary

The updated measure of Λ is 1.

Theorem

$\mathbb{P}_\Lambda^{up}$ is stable with respect to perturbations in $\mathbb{P}_\mathcal{D}^{obs}$ and in $\mathbb{P}_\Lambda^{init}$.

For details: [Combining Push-forward Measures and Bayes' Rule to Construct Consistent Solutions to Stochastic Inverse Problems, BJW. SISC 40 (2), 2018.]

A Solution to the Stochastic Inverse Problem

Theorem

Given an initial probability measure, $\mathbb{P}_\Lambda^{init}$ on $(\Lambda, \mathcal{B}_\Lambda)$ and an observed probability measure, $\mathbb{P}_\mathcal{D}^{obs}$, on $(\mathcal{D}, \mathcal{B}_\mathcal{D})$, the probability measure $\mathbb{P}_\Lambda^{up}$ on $(\Lambda, \mathcal{B}_\Lambda)$ defined by

$$\mathbb{P}_\Lambda^{up}(A) = \int_\mathcal{D} \left(\int_{A \cap Q^{-1}(q)} \pi_\Lambda^{init}(\lambda) \frac{\pi_\mathcal{D}^{obs}(Q(\lambda))}{\pi_\mathcal{D}^{pred}(Q(\lambda))} d\mu_{\Lambda,q}(\lambda) \right) d\mu_\mathcal{D}(q), \ \forall A \in \mathcal{B}_\Lambda$$

solves the stochastic inverse problem.

A Solution to the Stochastic Inverse Problem

Theorem

Given an initial probability measure, $\mathbb{P}_\Lambda^{init}$ on $(\Lambda, \mathcal{B}_\Lambda)$ and an observed probability measure, $\mathbb{P}_\mathcal{D}^{obs}$, on $(\mathcal{D}, \mathcal{B}_\mathcal{D})$, the probability measure $\mathbb{P}_\Lambda^{up}$ on $(\Lambda, \mathcal{B}_\Lambda)$ defined by

$$\mathbb{P}_\Lambda^{up}(A) = \int_\mathcal{D} \left(\int_{A \cap Q^{-1}(q)} \pi_\Lambda^{init}(\lambda) \frac{\pi_\mathcal{D}^{obs}(Q(\lambda))}{\pi_\mathcal{D}^{pred}(Q(\lambda))} d\mu_{\Lambda,q}(\lambda) \right) d\mu_\mathcal{D}(q), \ \forall A \in \mathcal{B}_\Lambda$$

solves the stochastic inverse problem.

The updated density is:

$$\pi_\Lambda^{\text{up}}(\lambda) = \pi_\Lambda^{\text{init}}(\lambda) \frac{\pi_\mathcal{D}^{\text{obs}}(Q(\lambda))}{\pi_\mathcal{D}^{\text{pred}}(Q(\lambda))}.$$

- Both $\pi_\Lambda^{\text{init}}$ and $\pi_\mathcal{D}^{\text{obs}}$ are given.
- Computing $\pi_\mathcal{D}^{\text{pred}}$ requires a forward propagation of the initial density.

A Parameterized Nonlinear System

Example

Consider a parameterized nonlinear system of equations:

$$\begin{aligned} \lambda_1 u_1^2 + u_2^2 &= 1, \\ u_1^2 - \lambda_2 u_2^2 &= 1 \end{aligned}$$

- Quantity of interest is the second component: $Q(\lambda) = u_2$.
- Given $\pi_{\mathcal{D}}^{\text{obs}} \sim N(0.3, 0.025^2)$.
- Given a uniform initial density.
- Use 10,000 samples from the initial and a standard KDE to approximate the push-forward.
- Use standard rejection sampling to generate samples from $\pi_{\Lambda}^{\text{up}}$.

A Parameterized Nonlinear System

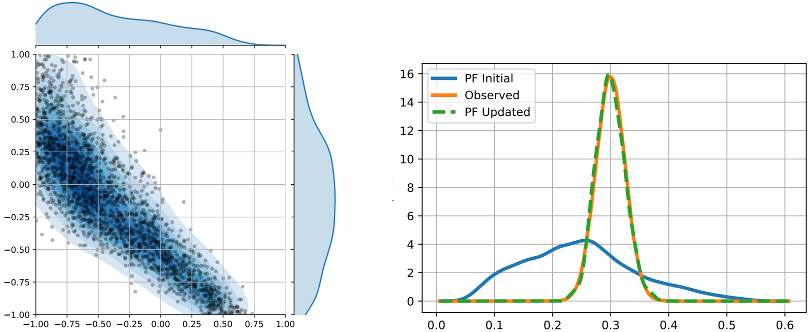

Figure: Samples from the updated density (left) and a comparison of $\pi_{\mathcal{D}}^{\text{obs}}$, $\pi_{\mathcal{D}}^{\text{pred}}$ and push-forward of the updated density (right).

Why do we care about approximate models?

Flow in Nuclear Reactor (Turbulent CFD)

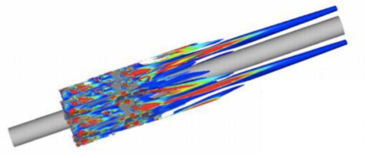

Tokamak Equilibrium (MHD)

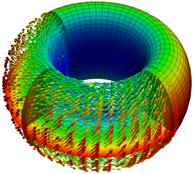

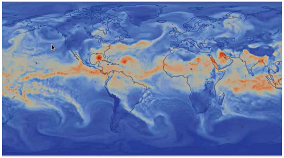

Climate Modeling

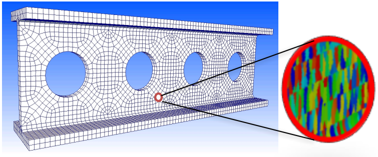

Multi-scale Materials Modeling

Flow in Nuclear Reactor (Turbulent CFD)

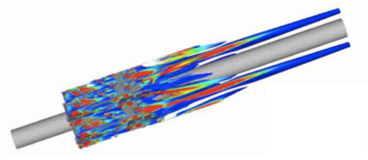

Tokamak Equilibrium (MHD)

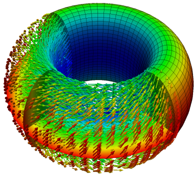

All are computationally expensive and require some form of approximation ...

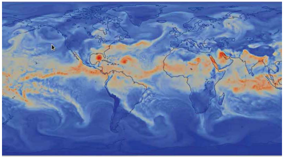

Climate Modeling

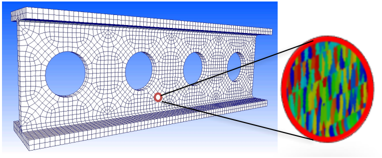

Multi-scale Materials Modeling

Convergence of Inverse Solutions

Recall that the updated density is given by

$$\pi_{\Lambda}^{\text{up}}(\lambda) = \pi_{\Lambda}^{\text{init}}(\lambda) \frac{\pi_{\mathcal{D}}^{\text{obs}}(Q(\lambda))}{\pi_{\mathcal{D}}^{\text{pred}}(Q(\lambda))}$$

The updated density using a surrogate model, $Q_S(\lambda)$, is given by

$$\pi_{\Lambda}^{\text{up},S}(\lambda) = \pi_{\Lambda}^{\text{init}}(\lambda) \frac{\pi_{\mathcal{D}}^{\text{obs}}(Q_S(\lambda))}{\pi_{\mathcal{D}}^{\text{pred},S}(Q_S(\lambda))}$$

Convergence of Inverse Solutions

Recall that the updated density is given by

$$\pi_\Lambda^{\text{up}}(\lambda) = \pi_\Lambda^{\text{init}}(\lambda)\frac{\pi_\mathcal{D}^{\text{obs}}(Q(\lambda))}{\pi_\mathcal{D}^{\text{pred}}(Q(\lambda))}$$

The updated density using a surrogate model, $Q_S(\lambda)$, is given by

$$\pi_\Lambda^{\text{up},S}(\lambda) = \pi_\Lambda^{\text{init}}(\lambda)\frac{\pi_\mathcal{D}^{\text{obs}}(Q_S(\lambda))}{\pi_\mathcal{D}^{\text{pred},S}(Q_S(\lambda))}$$

Theorem (B.J.W. SISC 2018b)

Under the the assumptions in [B.J.W., 2018b], $Q_S(\lambda) \to Q(\lambda)$ in $L^\infty(\Lambda) \implies \pi_\Lambda^{up,S}(\lambda) \to \pi_\Lambda^{up}(\lambda)$ in $L^1(\Lambda)$.

Extensions to convergence in L^p have also been developed recently [Butler, Wildey, Zhang, IJUQ, 2022].

Does this include data-driven models?

Theorem (W. Zhang Thesis 2021)

Suppose $Q \in C(\Lambda)$ *and the assumptions in [B.J.W., 2018b] are satisfied. Then* **there exists** *a sequence of single hidden layer Neural Networks defined on* Λ *such that (amongst other results):*

$$\pi_{\Lambda}^{up,S}(\lambda) \rightarrow \pi_{\Lambda}^{up}(\lambda) \text{ in } L^1(\Lambda).$$

Similar results can be shown for Neural Networks with arbitrary depth and fixed width by combining this result with the UAT from [Zhou et al 2017].

Does this include data-driven models?

Theorem (W. Zhang Thesis 2021)

Suppose $Q \in C(\Lambda)$ and the assumptions in [B.J.W., 2018b] are satisfied. Then **there exists** *a sequence of single hidden layer Neural Networks defined on Λ such that (amongst other results):*

$$\pi_\Lambda^{up,S}(\lambda) \to \pi_\Lambda^{up}(\lambda) \text{ in } L^1(\Lambda).$$

Similar results can be shown for Neural Networks with arbitrary depth and fixed width by combining this result with the UAT from [Zhou et al 2017].

Does this help me add **error bars or confidence intervals** to the stochastic inverse problem using **a given surrogate**?

Does this include data-driven models?

Theorem (W. Zhang Thesis 2021)

Suppose $Q \in C(\Lambda)$ and the assumptions in [B.J.W., 2018b] are satisfied. Then **there exists** *a sequence of single hidden layer Neural Networks defined on Λ such that (amongst other results):*

$$\pi_{\Lambda}^{up,S}(\lambda) \to \pi_{\Lambda}^{up}(\lambda) \text{ in } L^1(\Lambda).$$

Similar results can be shown for Neural Networks with arbitrary depth and fixed width by combining this result with the UAT from [Zhou et al 2017].

Does this help me add **error bars or confidence intervals** to the stochastic inverse problem using **a given surrogate**?

No, but let's see what we can do ...

Estimating Error/Uncertainty in Surrogate Models

Data-driven models tend to have **many** sources of error/uncertainty:

- Discretization/architecture (epistemic)
- Sparse/uninformative data (epistemic)
- Noisy data (aleatoric)
- Optimization/solver variability (aleatoric)
- Extrapolation/OoD (epistemic)

Estimating Error/Uncertainty in Surrogate Models

Data-driven models tend to have **many** sources of error/uncertainty:

- Discretization/architecture (epistemic)
- Sparse/uninformative data (epistemic)
- Noisy data (aleatoric)
- Optimization/solver variability (aleatoric)
- Extrapolation/OoD (epistemic)

From [Hüllermeier and Waegeman 2021]:

> *... a trustworthy representation of uncertainty is desirable and should be considered as a key feature of any machine learning method ...*

From [Abdar et al 2021]:

> *... predictions made without UQ are usually not trustworthy.*

Dropout/Bayesian [Neal 2012; Gal et al 2016; ...] and ensemble-based [Lakshminarayanan et al 2017; Ashukha et al 2021, ...] approaches are the most common.

Using the proper ensemble for DCI

Suppose we compute an ensemble of data-driven surrogate models, $\left\{ Q_S^{(i)}(\lambda) \right\}_{i=1}^{M}$.

Let $\overline{g}$ denote an ensemble-averaged quantity, e.g.,

$$\overline{Q}_S(\lambda) = \frac{1}{M} \sum_{i=1}^{M} Q_S^{(i)}(\lambda)$$

Using the proper ensemble for DCI

Suppose we compute an ensemble of data-driven surrogate models, $\left\{ Q_S^{(i)}(\lambda) \right\}_{i=1}^{M}$.

Let $\overline{g}$ denote an ensemble-averaged quantity, e.g.,

$$\overline{Q}_S(\lambda) = \frac{1}{M} \sum_{i=1}^{M} Q_S^{(i)}(\lambda)$$

How can we construct a data-consistent measure/density?

Using the proper ensemble for DCI

Suppose we compute an ensemble of data-driven surrogate models, $\left\{ Q_S^{(i)}(\lambda) \right\}_{i=1}^{M}$.

Let $\overline{g}$ denote an ensemble-averaged quantity, e.g.,

$$\overline{Q}_S(\lambda) = \frac{1}{M} \sum_{i=1}^{M} Q_S^{(i)}(\lambda)$$

How can we construct a data-consistent measure/density?

Each member of the ensemble can be used to compute a data-consistent solution:

$$\pi_\Lambda^{\text{up},S,i}(\lambda) = \pi_\Lambda^{\text{init}}(\lambda) \frac{\pi_\mathcal{D}^{\text{obs}}(Q_S^{(i)}(\lambda))}{\pi_\mathcal{D}^{\text{pred},S,i}(Q_S^{(i)}(\lambda))}$$

Using the proper ensemble for DCI

Suppose we compute an ensemble of data-driven surrogate models, $\left\{ Q_S^{(i)}(\lambda) \right\}_{i=1}^{M}$.

Let $\overline{g}$ denote an ensemble-averaged quantity, e.g.,

$$\overline{Q}_S(\lambda) = \frac{1}{M} \sum_{i=1}^{M} Q_S^{(i)}(\lambda)$$

How can we construct a data-consistent measure/density?

Each member of the ensemble can be used to compute a data-consistent solution:

$$\pi_\Lambda^{\text{up},S,i}(\lambda) = \pi_\Lambda^{\text{init}}(\lambda) \frac{\pi_\mathcal{D}^{\text{obs}}(Q_S^{(i)}(\lambda))}{\pi_\mathcal{D}^{\text{pred},S,i}(Q_S^{(i)}(\lambda))}$$

Use the ensemble-averaged surrogate model, $\overline{Q}_S(\lambda)$, to compute the update:

$$\pi_\Lambda^{\text{up},S}(\lambda) = \pi_\Lambda^{\text{init}}(\lambda) \frac{\pi_\mathcal{D}^{\text{obs}}(\overline{Q}_S(\lambda))}{\pi_\mathcal{D}^{\text{pred},S}(\overline{Q}_S(\lambda))}$$

A simple example

Consider the following partial differential equation used in [Butler, W., IJUQ 2018]

$$\begin{cases} -\nabla \cdot (\mathsf{K}\nabla u) + \mathsf{b}(\lambda_1, \lambda_2, x) \cdot \nabla u = g(x), & x \in \Omega = (0,1) \times (0,1) \\ u = 0, & x \in \partial\Omega \end{cases}$$

The quantity of interest is a mollified point-evaluation:

$$Q(u(\lambda)) = \int_\Omega \frac{100}{\pi} e^{-100(x_1 - 0.5)^2 - 100(x_2 - 0.5)^2} u(x)\, dx.$$

Discretization details:

- Finite element on 50×50 mesh,
- $\pi_\Lambda^{\text{init}}$ is uniform on $[0,1]^2$
- Build feedforward NN surrogate: $2 \to 20 \to 20 \to 1$ with ReLU activation
- Use 1,000 samples split 80/20 for training/testing
- Use 20,000 samples evaluated using surrogate to approximate push-forward
- Observed distribution is $N(0.033, 0.001^2)$

Approximations and Errors

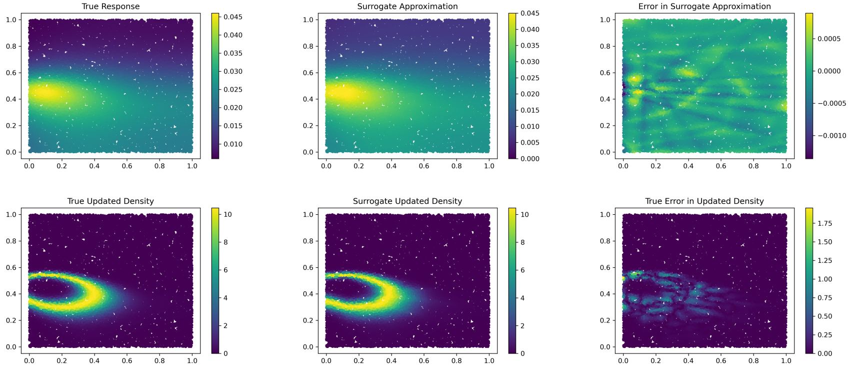

Figure: Top row: true response, approximation and error. Bottom row: true solution to the inverse problem, approximation and error.

Uncertainty Characterization Using Dropout(0.01)

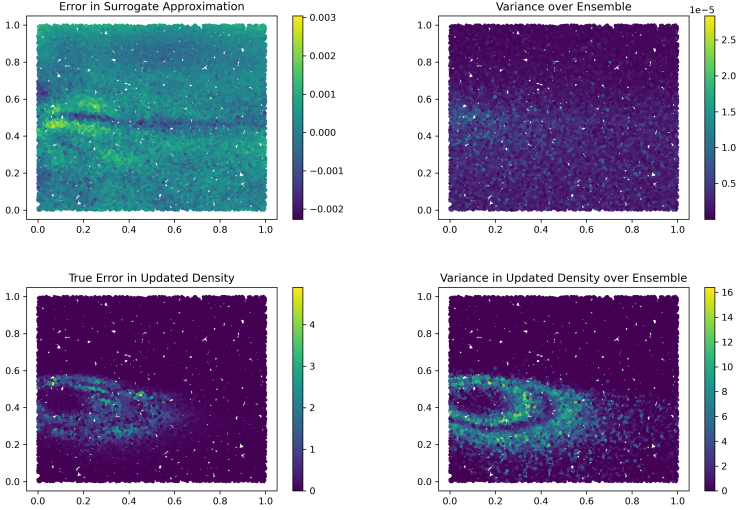

Uncertainty Characterization Using Ensembles

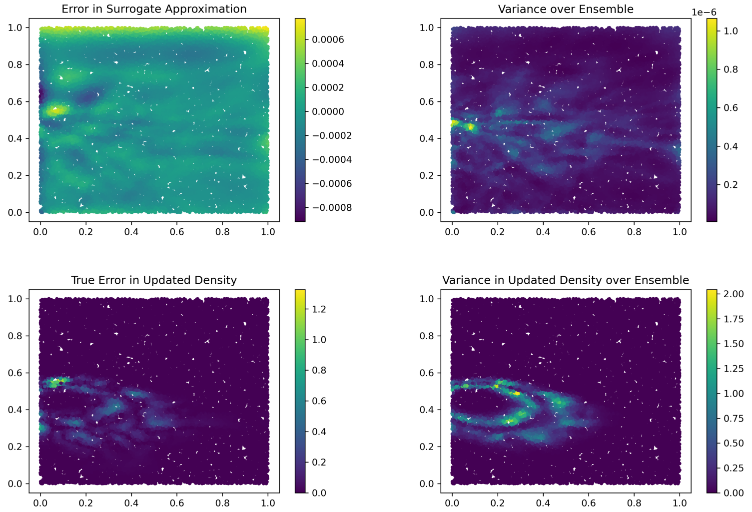

Uncertainty Characterization Using Ensemble of Dropouts

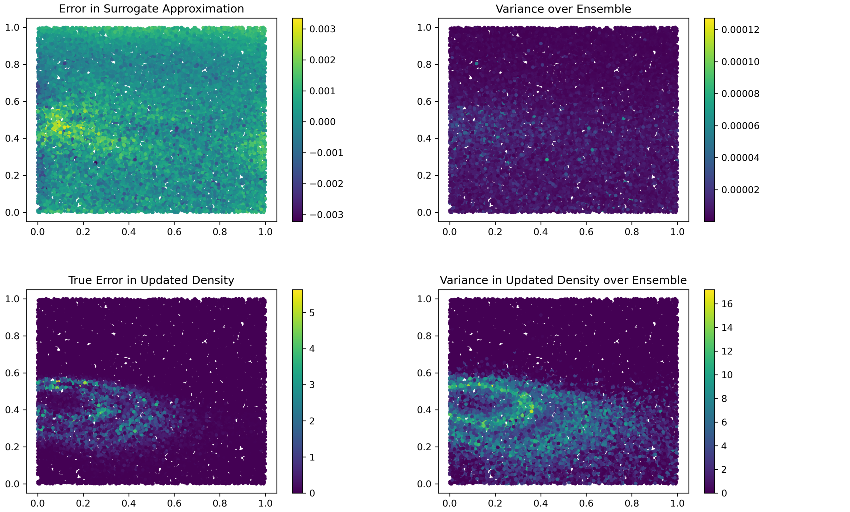

More Formal Verification Techniques

- Dropout and ensembles characterize the *predictive uncertainty*, i.e., the *precision* of the model.
- We are more interested in the *accuracy* of a particular surrogate model.

More Formal Verification Techniques

- Dropout and ensembles characterize the *predictive uncertainty*, i.e., the *precision* of the model.
- We are more interested in the *accuracy* of a particular surrogate model.

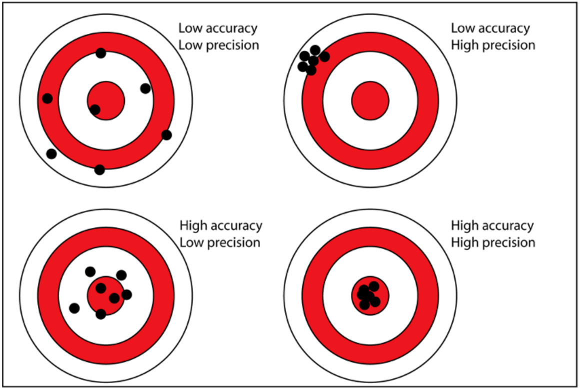

More Formal Verification Techniques

- Dropout and ensembles characterize the *predictive uncertainty*, i.e., the *precision* of the model.
- We are more interested in the *accuracy* of a particular surrogate model.
- Why not use a formal **solution verification** [Eca et al 2010, Xind et al 2010, Rider et al 2016] procedure?
  - Richardson extrapolation, regularized/weighted least squares, etc.

More Formal Verification Techniques

- Dropout and ensembles characterize the *predictive uncertainty*, i.e., the *precision* of the model.
- We are more interested in the *accuracy* of a particular surrogate model.
- Why not use a formal **solution verification** [Eca et al 2010, Xind et al 2010, Rider et al 2016] procedure?
  - Richardson extrapolation, regularized/weighted least squares, etc.
- All require an ansatz and perform best in asymptotic regime.
  - More work leads to smaller error
- Not necessarily true for NN surrogates!

More Formal Verification Techniques

- Dropout and ensembles characterize the *predictive uncertainty*, i.e., the *precision* of the model.
- We are more interested in the *accuracy* of a particular surrogate model.
- Why not use a formal **solution verification** [Eca et al 2010, Xind et al 2010, Rider et al 2016] procedure?
  - Richardson extrapolation, regularized/weighted least squares, etc.
- All require an ansatz and perform best in asymptotic regime.
  - More work leads to smaller error
- Not necessarily true for NN surrogates!

Can we develop an error estimation scheme that does not require monotonic behaviour?

Error Estimates for Surrogates of Quantities of Interest from Physics-based Models

Let's assume:

- we have a QoI from a deterministic physics-based model,
- we have an adjoint for the physics-based model,

Error Estimates for Surrogates of Quantities of Interest from Physics-based Models

Let's assume:

- we have a QoI from a deterministic physics-based model,
- we have an adjoint for the physics-based model,

What is an adjoint?

Error Estimates for Surrogates of Quantities of Interest from Physics-based Models

Let's assume:

- we have a QoI from a deterministic physics-based model,
- we have an adjoint for the physics-based model,

What is an adjoint?

Definition

Let X and Y be Banach spaces and L denote a linear operator $L : X \to Y$. The *adjoint operator* $L^* : Y^* \to X^*$ is defined such that

$$\langle Lx, y \rangle = \langle x, L^* y \rangle, \quad \forall x \in X, y \in Y.$$

Error Estimates for Surrogates of Quantities of Interest from Physics-based Models

Let's assume:

- we have a QoI from a deterministic physics-based model,
- we have an adjoint for the physics-based model,

What is an adjoint?

Definition

Let X and Y be Banach spaces and L denote a linear operator $L : X \to Y$. The *adjoint operator* $L^* : Y^* \to X^*$ is defined such that

$$\langle Lx, y \rangle = \langle x, L^* y \rangle, \quad \forall x \in X, y \in Y.$$

Given a functional of the forward state, $J(u)$, the adjoint problem is given by:

$$L^* \phi = D_u J.$$

Often used in optimization and a posteriori error estimation.

Error Estimates for Surrogates of Quantities of Interest from Physics-based Models

We can use a generalization of adjoint-based techniques to estimate the error in **point-wise evaluations** of the surrogate model [Butler, Dawson, W. 2011].

Let u denote the true solution to the model, $\tilde{U}$ be an approximation and $R(\tilde{U})$ the residual.

The error in a functional of the solution is given by:

$$J(u) - J(\tilde{U}) = \left\langle R(\tilde{U}), \phi \right\rangle + \text{higher order terms},$$

where ϕ is the adjoint solution.

Given an approximate adjoint solution, $\tilde{\phi}$, we have:

$$J(u) - J(\tilde{U}) \approx \left\langle R(\tilde{U}), \tilde{\phi} \right\rangle + \underbrace{\left\langle R(\tilde{U}), \phi - \tilde{\phi} \right\rangle}_{\text{higher order}},$$

Error Estimates for Surrogates

Such error estimates are higher-order and can be used to:

- Define an improved surrogate model [Butler, Dawson, W. 2013]
- Drive adaptivity in the surrogate model [Jakeman, W. 2015]
- Decompose errors into various contributions [Bryant, Prudhomme, W. 2015]
- Derive better MCMC sampling strategies [Butler, Dawson, W. 2015]
- Estimate errors in probabilities of rare events [Butler, W. 2018]

Error Estimates for Surrogates

Such error estimates are higher-order and can be used to:

- Define an improved surrogate model [Butler, Dawson, W. 2013]
- Drive adaptivity in the surrogate model [Jakeman, W. 2015]
- Decompose errors into various contributions [Bryant, Prudhomme, W. 2015]
- Derive better MCMC sampling strategies [Butler, Dawson, W. 2015]
- Estimate errors in probabilities of rare events [Butler, W. 2018]

But what is the drawback?

Requires a surrogate of the forward and adjoint states!

Error Estimates for Surrogates

Such error estimates are higher-order and can be used to:

- Define an improved surrogate model [Butler, Dawson, W. 2013]
- Drive adaptivity in the surrogate model [Jakeman, W. 2015]
- Decompose errors into various contributions [Bryant, Prudhomme, W. 2015]
- Derive better MCMC sampling strategies [Butler, Dawson, W. 2015]
- Estimate errors in probabilities of rare events [Butler, W. 2018]

But what is the drawback?

Requires a surrogate of the forward and adjoint states!

- Not a significant issue for GPCE, pseudo-spectral projection, sparse grids, etc.
- Challenging for NN models ...

Compression/Recovery of States

We seek to build a **compressed representation** of the states and a map from parameters to the latent space.

Set up and train:

- Autoencoders for compression into the latent space
- Feedforward NN for the parameter-to-latent mapping

Repeat for adjoint states

For a new parameter $\lambda \in \Lambda$, we

- Evaluate the parameter-to-latent maps
- Pass latent representations through decoders
- Compute approximate QoI
- Compute error estimate

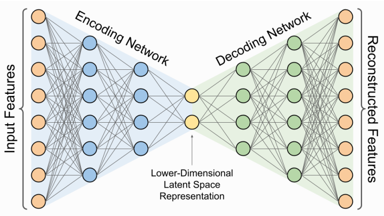

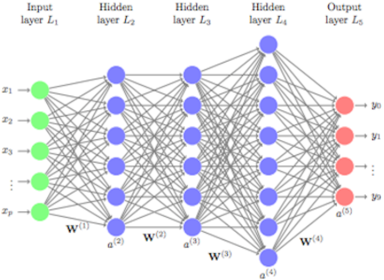

Validation of Forward Autoencoder

Autoencoder architecture: $2601 \rightarrow 128 \rightarrow 16 \rightarrow 128 \rightarrow 256 \rightarrow 2601$ with ReLU activations in hidden layers and tanh output activation.

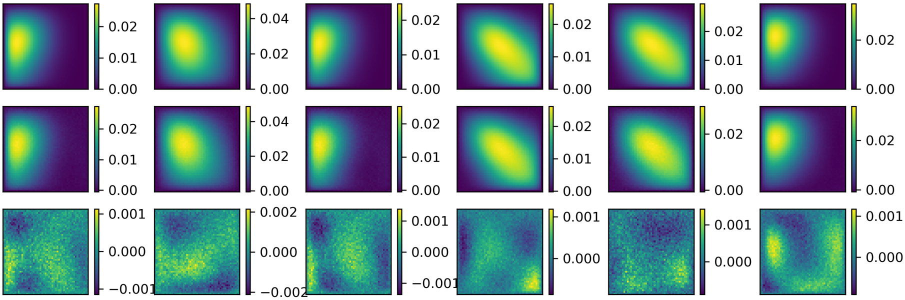

Figure: The true states (top row), the recovered states (middle row) and the error (bottom row) for validation states.

Validation of Adjoint Autoencoder

Autoencoder architecture: $2601 \rightarrow 128 \rightarrow 16 \rightarrow 128 \rightarrow 256 \rightarrow 2601$ with ReLU activations in hidden layers and tanh output activation.

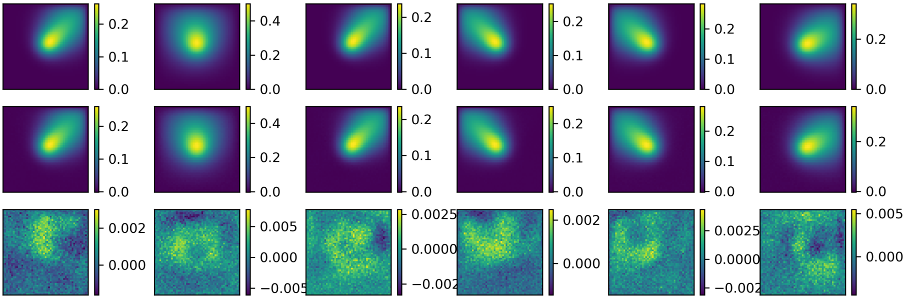

Figure: The true states (top row), the recovered states (middle row) and the error (bottom row) for validation states.

Validation of Forward Parameter-to-latent Map/Decoder

Parameter-to-latent architecture: $2 \rightarrow 32 \rightarrow 32 \rightarrow 16$ with ReLU activations in hidden layers and tanh output activation.

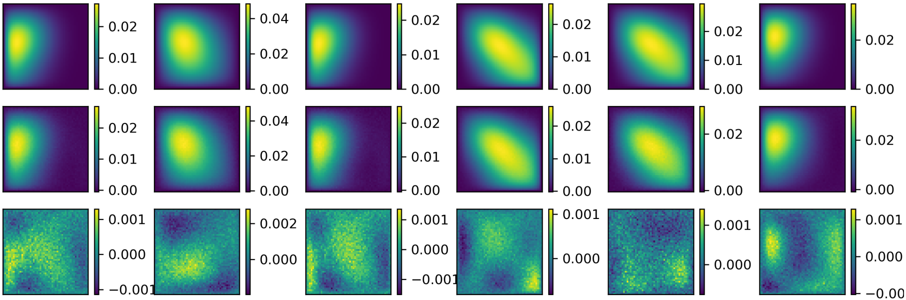

Figure: The true states (top row), the recovered states (middle row) and the error (bottom row) for validation states.

Validation of Adjoint Parameter-to-latent Map/Decoder

Parameter-to-latent architecture: $2 \to 32 \to 32 \to 16$ with ReLU activations in hidden layers and tanh output activation.

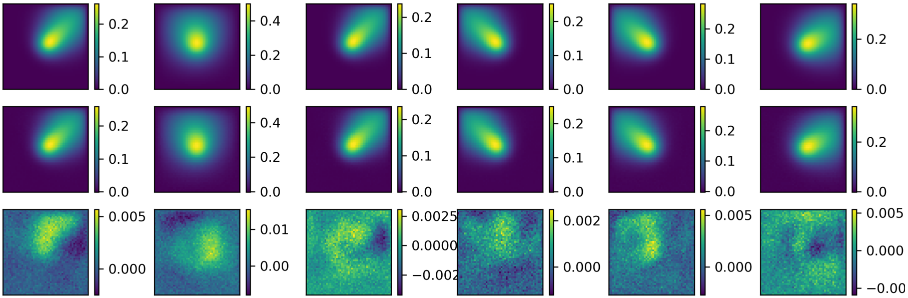

Figure: The true states (top row), the recovered states (middle row) and the error (bottom row) for validation states.

Estimating the Error in the Surrogate

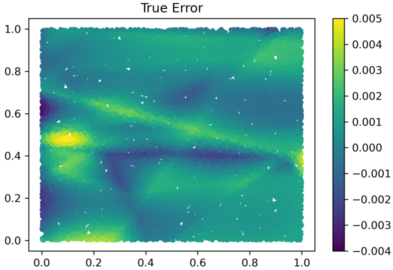

Figure: The true error (left) and the estimated error (right).

Estimating the Error in the Surrogate

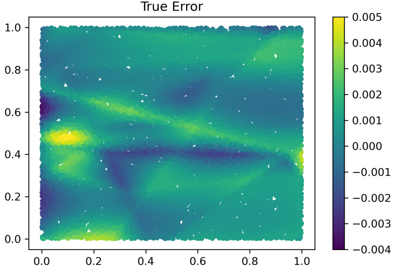

Figure: The true error (left) and the estimated error (right).

Error Estimates for Data-consistent Solutions

Suppose we are given

- A surrogate model, $Q_S(\lambda) \approx Q(\lambda)$.
- A set of samples (not training data), $\{\lambda_i\}_{i=1}^N$, generated from $\pi_\Lambda^{\text{init}}$, where we want to evaluate $Q_S(\lambda)$.
- An estimate of the error $e_i \approx Q(\lambda_i) - Q_S(\lambda_i)$

Error Estimates for Data-consistent Solutions

Suppose we are given

- A surrogate model, $Q_S(\lambda) \approx Q(\lambda)$.
- A set of samples (not training data), $\{\lambda_i\}_{i=1}^N$, generated from $\pi_\Lambda^{\text{init}}$, where we want to evaluate $Q_S(\lambda)$.
- An estimate of the error $e_i \approx Q(\lambda_i) - Q_S(\lambda_i)$

Then, we can defined the **improved surrogate approximation**:

$$Q_{S+}(\lambda_i) = Q_S(\lambda_i) + e_i,$$

and the **improved data-consistent solution**:

$$\pi_\Lambda^{\text{up},S+}(\lambda_i) = \pi_\Lambda^{\text{init}}(\lambda_i) r_{S+}(\lambda_i), \quad r_{S+}(\lambda_i) = \frac{\pi_{\mathcal{D}}^{\text{obs}}(Q_{S+}(\lambda_i))}{\pi_{\mathcal{D}}^{\text{pred},S+}(Q_{S+}(\lambda_i))}$$

Error Estimates for Data-consistent Solutions

The **improved ratio**, $r_{S+}(\lambda_i)$, can be used to estimate the error in the updated density in the total variation metric:

$$\int_\Lambda \left| \pi_\Lambda^{\text{up}}(\lambda) - \pi_\Lambda^{\text{up},S}(\lambda) \right| \, d\mu_\Lambda \approx \int_\Lambda \left| \pi_\Lambda^{\text{up},S+}(\lambda) - \pi_\Lambda^{\text{up},S}(\lambda) \right| \, d\mu_\Lambda$$

$$\approx \frac{1}{N} \sum_{i=1}^{N} |r_{S+}(\lambda_i) - r_S(\lambda_i)|$$

We can also it to evaluate the **reliability** in the updated density on a point-wise basis.

Estimating the Error in the Surrogate

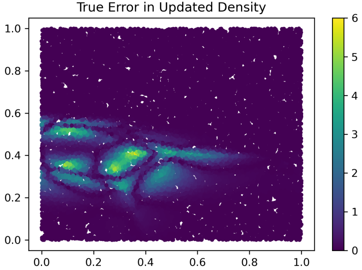

Figure: The true error (left) and the estimated error (right).

Estimating the Error in the Surrogate

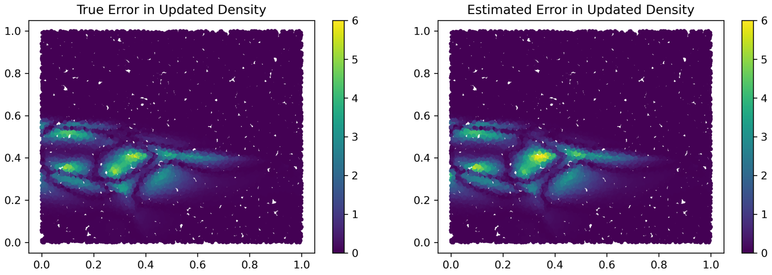

Figure: The true error (left) and the estimated error (right).

True L_1 Error	0.29135
Estimated L_1 Error	0.30101

Can We Assess OOD Errors?

Can We Assess OOD Errors?

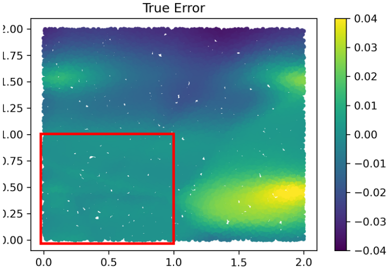

Can We Assess OOD Errors?

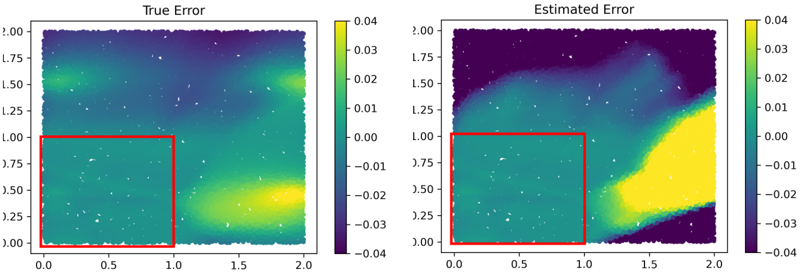

Can We Assess OOD Errors?

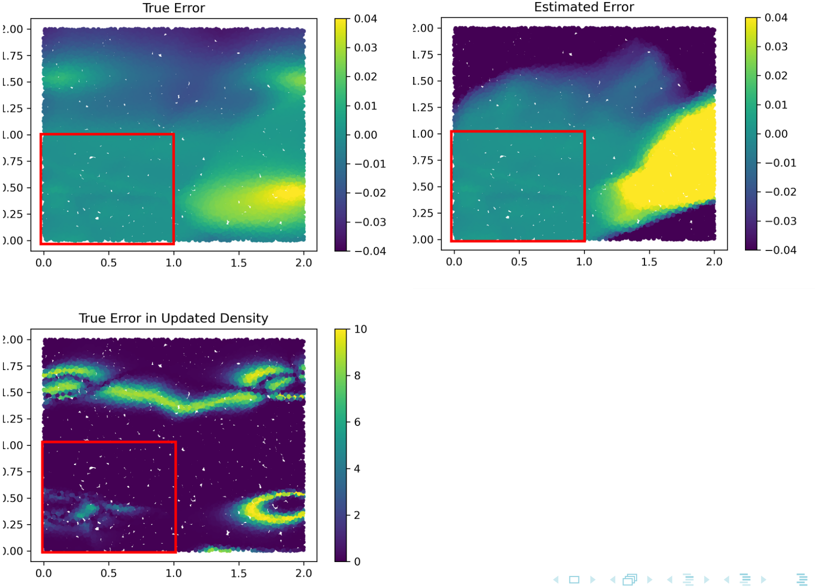

Can We Assess OOD Errors?

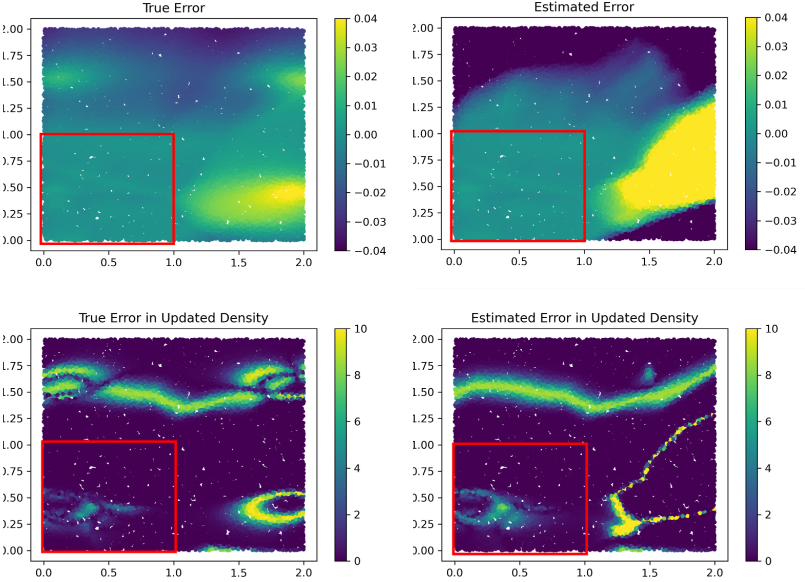

Conclusions and Future Work

- **Errors and uncertainties** can significantly affect the solution to inverse problems.
  - Affects the accept/reject of samples
  - Affects subsequent predictions
- If an adjoint model is available, then the affect of surrogate errors on updated density can be estimated by using dual-weighted residuals.
- Requires forward and adjoint state approximations.
  - We used standard autoencoders with parameter-to-latent NN surrogates.
  - Better compression methods may be required for transient and multiple QoI.
- Future work to limit dependence on dual-weighted residual for each evaluation.
  - Previous papers limited these evaluations by projecting error onto higher-order surrogate.

Thanks! Questions?

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, ASCR, Early Career Research Program.

Thank you for your attention!

Questions?