SAND2019-1844C

Exploiting Low-dimensional Structure to Efficiently Perform Stochastic Inference for Prediction

Tim Wildey

Sandia National Laboratories
Center for Computing Research
Optimization and UQ Group

SIAM Conference on Computational Science and Engineering
Spokane, WA
February 25 - March 1, 2019

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. SAND2019-**** C

Motivation

Flow in Nuclear Reactor (Turbulent CFD)

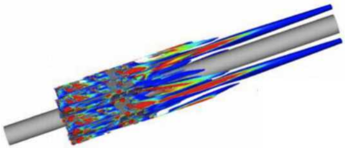

Tokamak Equilibrium (MHD)

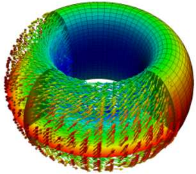

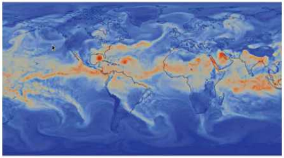

Climate Modeling

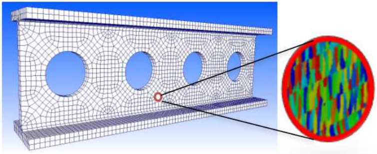

Multi-scale Materials Modeling

Motivation

Flow in Nuclear Reactor (Turbulent CFD)

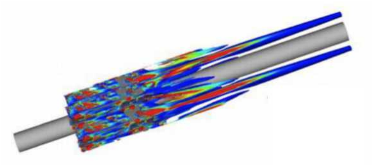

Tokamak Equilibrium (MHD)

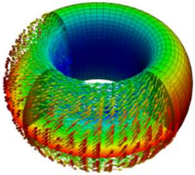

We are working to develop **data-informed** models ...

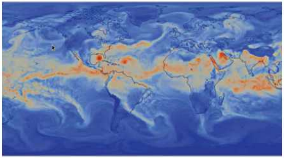

Climate Modeling

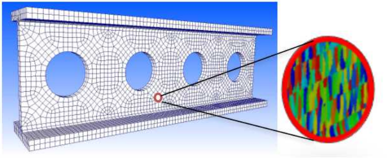

Multi-scale Materials Modeling

Data-informed Physics-Based Predictions

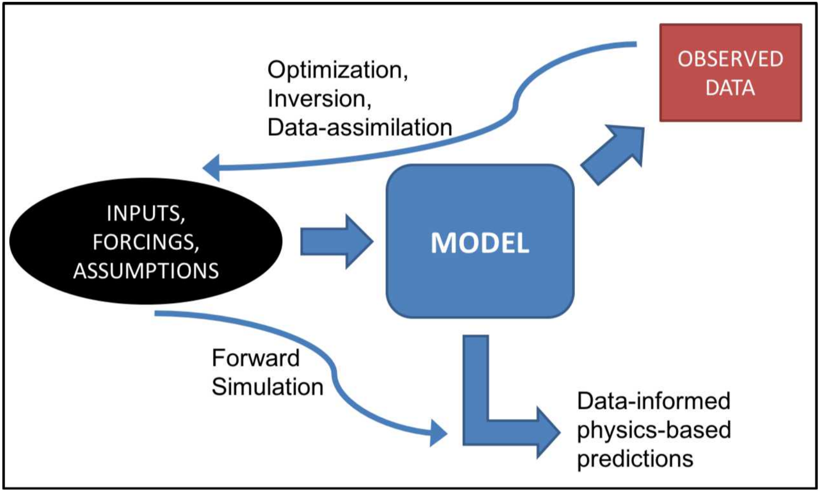

A Deterministic Inverse Problem

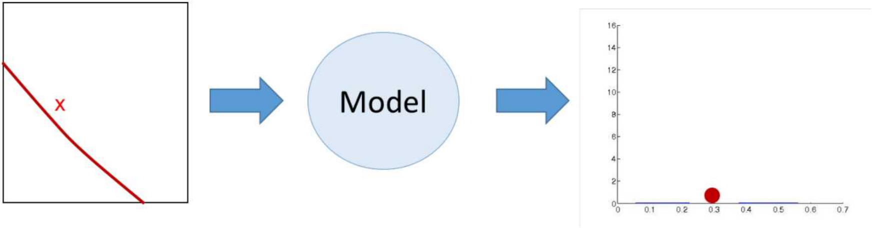

Problem

Given a deterministic observation, $\hat{Q}$, find $\lambda \in \Lambda$ such that $Q(\lambda) = \hat{Q}$.

A Deterministic Inverse Problem

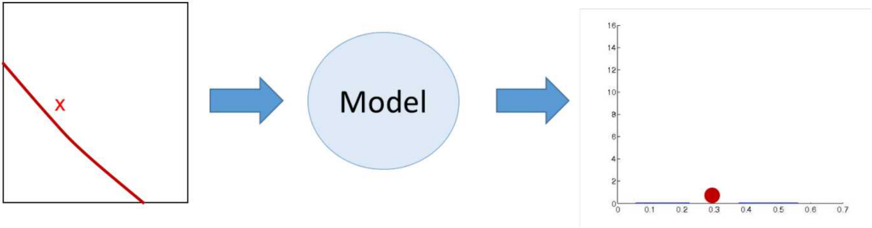

Problem

Given a deterministic observation, $\hat{Q}$, find $\lambda \in \Lambda$ such that $Q(\lambda) = \hat{Q}$.

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

A Stochastic Inverse Problem

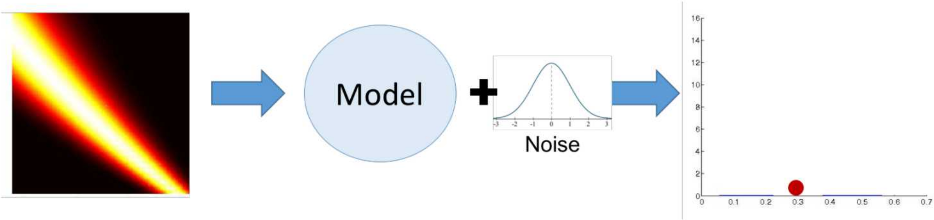

Problem

Given a deterministic observation, $\hat{Q}$, and an assumed noise model, find the parameters that are most likely to have produced the data.

A Stochastic Inverse Problem

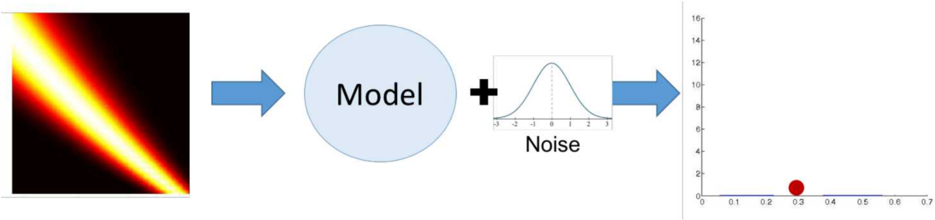

Problem

Given a deterministic observation, $\hat{Q}$, 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

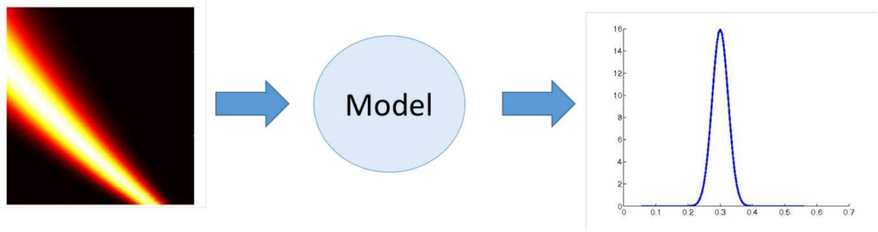

Problem

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

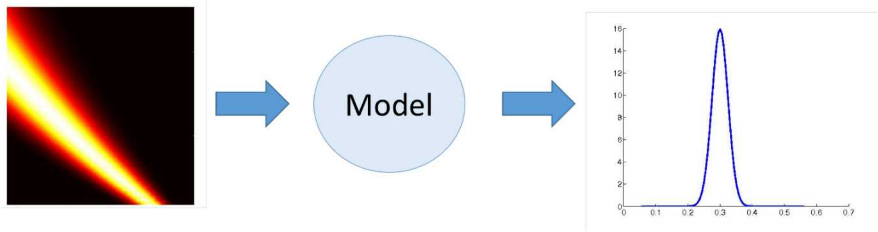

Problem

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. An **observed probability measure** on $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$, denoted $\mathbb{P}_{\mathcal{D}}^{\text{obs}}$, that has a density, $\pi_{\mathcal{D}}^{\text{obs}}$.
4. An **initial probability measure** on $(\Lambda, \mathcal{B}_{\Lambda})$, denoted $\mathbb{P}_{\Lambda}^{\text{init}}$, that has a density, $\pi_{\Lambda}^{\text{init}}$.

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. An **observed probability measure** on $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$, denoted $\mathbb{P}_{\mathcal{D}}^{\text{obs}}$, that has a density, $\pi_{\mathcal{D}}^{\text{obs}}$.
4. An **initial probability measure** on $(\Lambda, \mathcal{B}_{\Lambda})$, denoted $\mathbb{P}_{\Lambda}^{\text{init}}$, that has a density, $\pi_{\Lambda}^{\text{init}}$.

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 $\pi_{\mathcal{D}}^{Q(\text{init})}$ 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, $\mathbb{P}_{\mathcal{D}}^{Q(\text{init})}$.

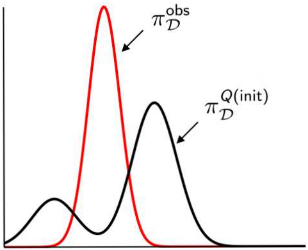

Good Initial

Bad Initial
(Cannot predict all observations)

A Solution to the Stochastic Inverse Problem

Theorem (Butler, Jakeman, Wildey, SISC, 2018a)

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 P_{Λ}^{up} on $(\Lambda, \mathcal{B}_{\Lambda})$ defined by

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

solves the stochastic inverse problem.

A Solution to the Stochastic Inverse Problem

Theorem (Butler, Jakeman, Wildey, SISC, 2018a)

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 P_Λ^{up} on $(\Lambda, \mathcal{B}_\Lambda)$ defined by

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

solves the stochastic inverse problem.

Corollary (BJW., SISC 2018a)

The updated measure of Λ is 1.

A Solution to the Stochastic Inverse Problem

Theorem (Butler, Jakeman, Wildey, SISC, 2018a)

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 P_Λ^{up} on $(\Lambda, \mathcal{B}_\Lambda)$ defined by

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

solves the stochastic inverse problem.

Corollary (BJW., SISC 2018a)

The updated measure of Λ is 1.

Theorem (BJW., SISC 2018a)

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

A Solution to the Stochastic Inverse Problem

Theorem (Butler, Jakeman, Wildey, SISC, 2018a)

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 P_Λ^{up} on $(\Lambda, \mathcal{B}_\Lambda)$ defined by

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

solves the stochastic inverse problem.

Corollary (BJW., SISC 2018a)

The updated measure of Λ is 1.

Theorem (BJW., SISC 2018a)

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

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 (Butler, Jakeman, Wildey, SISC, 2018a)

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 P_Λ^{up} on $(\Lambda, \mathcal{B}_\Lambda)$ defined by

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

solves the stochastic inverse problem.

The updated density is:

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

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

A Parameterized Nonlinear System

Example

Consider a parameterized nonlinear system of equations:

$$\begin{aligned} \lambda_1 x_1^2 + x_2^2 &= 1, \\ x_1^2 - \lambda_2 x_2^2 &= 1 \end{aligned}$$

- The quantity of interest is the second component: $q(\lambda) = x_2$.
- Assume that we observe $q(\lambda) \sim N(0.3, 0.025^2)$.
- We consider a uniform initial density.
- We use 10,000 samples from the initial and a standard KDE to approximate the push-forward.

A Parameterized Nonlinear System

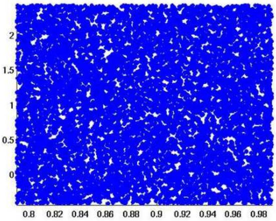

Initial

A Parameterized Nonlinear System

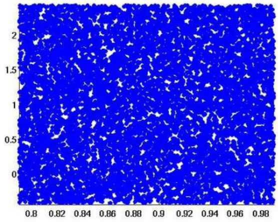

Initial

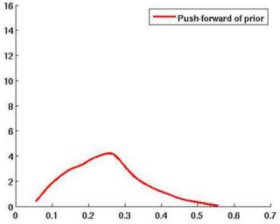

Push-forward of Initial

A Parameterized Nonlinear System

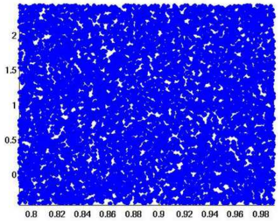

Initial

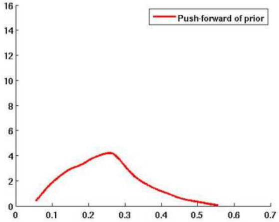

Push-forward of Initial

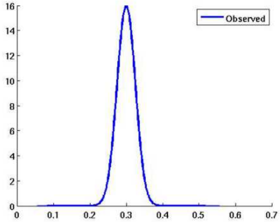

Observed density

A Parameterized Nonlinear System

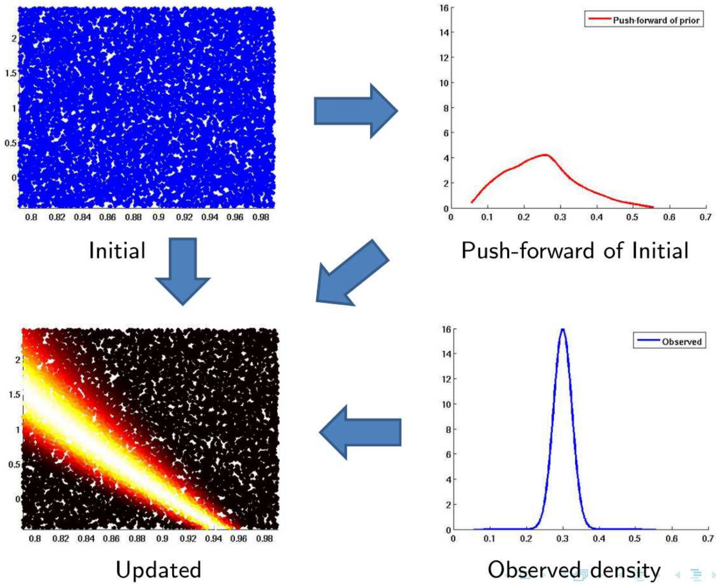

A Parameterized Nonlinear System

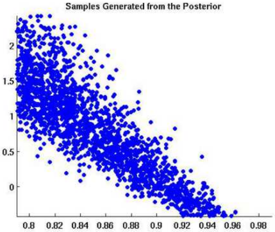

Samples from the updated density

A Parameterized Nonlinear System

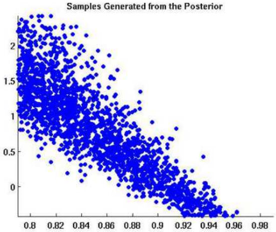

Samples from the updated density

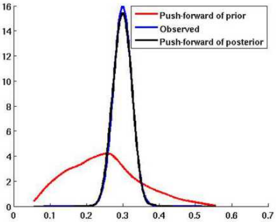

Observed and push-forward densities in $\mathcal{D}$

Nice, but not very practical!

- We cannot expect to be able to generate a large number of samples from a high-fidelity computational model!

Nice, but not very practical!

- We cannot expect to be able to generate a large number of samples from a high-fidelity computational model!
- Can we use approximate models, e.g., discretizations or surrogate models?
  - Yes, see [Butler, Jakeman, W. SISC 2018b].

Nice, but not very practical!

- We cannot expect to be able to generate a large number of samples from a high-fidelity computational model!
- Can we use approximate models, e.g., discretizations or surrogate models?
  - Yes, see [Butler, Jakeman, W. SISC 2018b].
- Can we leverage lower-fidelity models in a multi-fidelity context?
  - Yes, see [Bruder, Gee, W. 2019 (in review)]

Nice, but not very practical!

- We cannot expect to be able to generate a large number of samples from a high-fidelity computational model!
- Can we use approximate models, e.g., discretizations or surrogate models?
  - Yes, see [Butler, Jakeman, W. SISC 2018b].
- Can we leverage lower-fidelity models in a multi-fidelity context?
  - Yes, see [Bruder, Gee, W. 2019 (in review)]
- Can we leverage connections with deterministic optimization with regularization to develop scalable approaches?
  - Yes, see [Marvin, Bui-Thanh, W. CCR Proceedings 2018].

Nice, but not very practical!

- We cannot expect to be able to generate a large number of samples from a high-fidelity computational model!
- Can we use approximate models, e.g., discretizations or surrogate models?
  - Yes, see [Butler, Jakeman, W. SISC 2018b].
- Can we leverage lower-fidelity models in a multi-fidelity context?
  - Yes, see [Bruder, Gee, W. 2019 (in review)]
- Can we leverage connections with deterministic optimization with regularization to develop scalable approaches?
  - Yes, see [Marvin, Bui-Thanh, W. CCR Proceedings 2018].
- **Can we use dimension reduction techniques, e.g., active subspaces?**

A General Framework

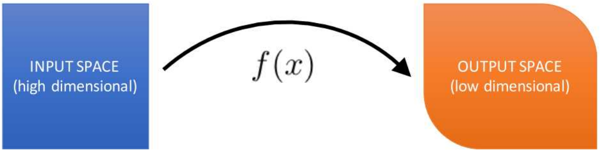

A General Framework

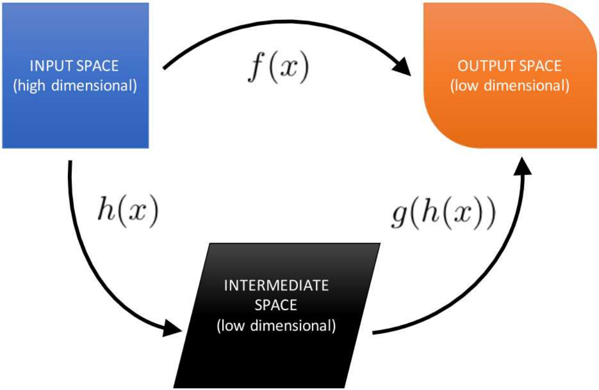

A General Framework

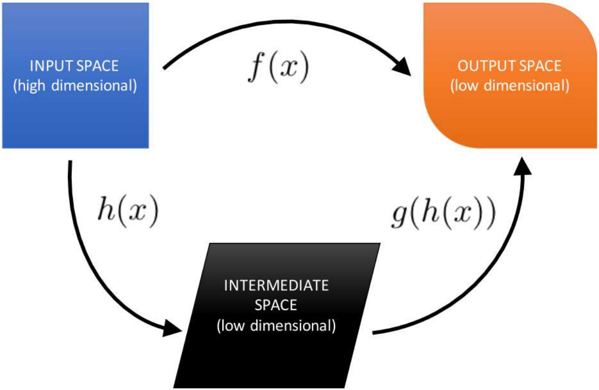

- If $h(x)$ is lower-fidelity model, then we recover a particular multi-fidelity formulation [Koutsourelakis 2009; Biehler, Gee, Wall 2015; Bruder, Gee, W. 2019]

A General Framework

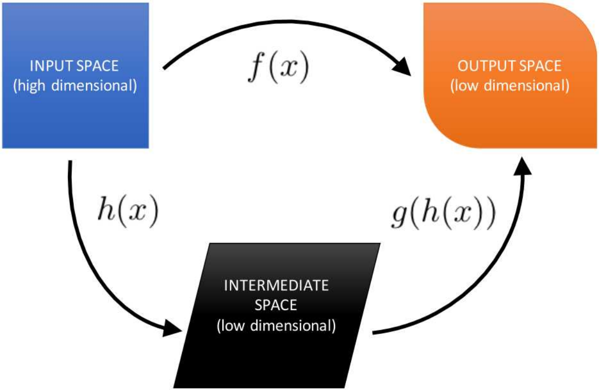

- If $h(x)$ is lower-fidelity model, then we recover a particular multi-fidelity formulation [Koutsourelakis 2009; Biehler, Gee, Wall 2015; Bruder, Gee, W. 2019]
- If $h(x) = W^T x$, then we have a ridge approximation.

Ridge Approximations and Active Subspaces

- What is a *ridge function*?
  - A composite function of the form: $f(x) = g(h(x))$
  - $y = h(x)$ depends *linearly* on x, e.g. $y = \mathbf{W}_A^T x$ where $\mathbf{W_A} \in \mathbb{R}^{n \times m}$.
- What does an active subspace method do?
  - Use evaluations of the function and/or the gradient, $\nabla f(x)$, to find $\mathbf{W}_A$.
  - Define the average outer product of the gradient and its eigendecomposition

$$\mathbf{C} = \int_\Lambda \nabla f(x) \nabla f(x)^T d\mathbb{P}_\Lambda^{\text{init}} = \mathbf{W} \mathbf{\Lambda} \mathbf{W}^T.$$

  - Partition the eigendecomposition,

$$\mathbf{\Lambda} = \begin{bmatrix} \Lambda_A & \\ & \Lambda_I \end{bmatrix}, \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_A & \mathbf{W}_I \end{bmatrix}, \quad \mathbf{W}_A \in \mathbb{R}^{n \times m}$$

  - Define a rotation and partition into *active* and *inactive* directions,

$$x = \mathbf{W}\mathbf{W}^T x = \mathbf{W}_A \mathbf{W}_A^T x + \mathbf{W}_I \mathbf{W}_I^T x = \mathbf{W}_A y + \mathbf{W}_I z$$

  - y denotes the *active* variables and z the *inactive* variables.

Ridge Approximations and Active Subspaces

- Given $\mathbf{W}_A$, we can easily compute $y_i = \mathbf{W}_A^T x_i$ for any sample x_i.
- Plots of $y = h(x) = \mathbf{W}_A^T x$ are helpful:

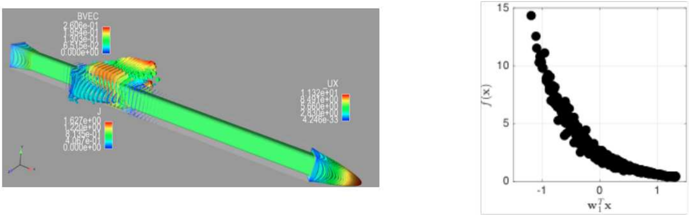

Figure: MHD generator model (left) and the 1-dimensional active subspace studied in [Glaws, Constantine, Shadid, W. 2017]

- Need to define a reasonable approximation of $g(y) = g(h(x))$.
- Let $G(y)$ denote the conditional expectation with respect to the inactive variables:

$$G(y) = \mathbb{E}_z[f|y] = \int f(\mathbf{W}_A y + \mathbf{W}_I z)\pi_{Z|Y}(z)\, dz.$$

- Various approaches can be used to approximate G (e.g., the mean of a GP).

A General Framework

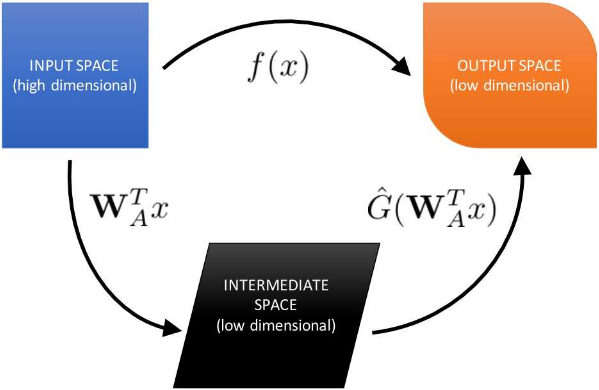

- In general, $\hat{G}(\mathbf{W}_A^T x)$ defines a **different** push-forward probability measure.
- Can be used to solve the stochastic inverse problem if it satisfies a *predictibility assumption.*
- In practice, it is easy to detect if this assumption is satisfied.

Active subspaces provide L^2 error estimate

Theorem (Theorem 3.7, Constantine, Dow, Wang SISC 2014)

The mean-squared error using N Monte Carlo samples to approximate the eigenvalues, perturbed eigenvectors, $\hat{\mathbf{W}}_A$, with error ϵ, and a response surface approximation $G \approx g(y)$ with error δ, is given by

$$\|f(x) - G(\hat{\mathbf{W}}_A^T x)\|^2_{L^2(\Lambda)} \leq$$

$$C_1 \left(1 + \frac{1}{N}\right) \left(\epsilon \left(\lambda_1 + \cdots \lambda_n\right)^{1/2} + \left(\lambda_{n+1} + \cdots \lambda_m\right)^{1/2}\right)^2 + C_2 \delta.$$

Bounds the error in the push-forward

A modification of the arguments in [Butler, Jakeman, W. SISC 2018b] gives the following estimate of the error in the push-forward using an active subspace.

Theorem

The expected error in a kernel density estimate of the push-forward of the initial density, $\hat{\pi}_{\mathcal{D}}^{Q(init)}$*, using a sufficiently smooth kernel of order* s*, and* M *Monte Carlo evaluations of the active subspace model is bounded by,*

$$\mathbb{E}\left[\|\pi_{\mathcal{D}}^{Q(init)} - \hat{\pi}_{\mathcal{D}}^{Q(init)}\|_{L^2(\mathcal{D})}^2\right] \leq C\Big(\underbrace{\frac{\log M}{M^{2s/(2s+m)}}}_{\text{KDE error}} + \underbrace{\|f(x) - G(\hat{\mathbf{W}}_A^T x)\|_{L^2(\Lambda)}^2}_{\text{Active subspace error}}\Big)$$

Bounds the error in the updated density

Similarly, we can bound the error in the updated density using an active subspace.

Theorem

The expected error in the updated density, $\hat{\pi}_{\Lambda}^{up}$, using a kernel density estimate of the push-forward of the initial density[a], $\hat{\pi}_{\mathcal{D}}^{Q(init)}$, *with a sufficiently smooth kernel of order s, and M Monte Carlo evaluations of the active subspace model is bounded by,*

$$\mathbb{E}\left[\|\pi_{\Lambda}^{up} - \hat{\pi}_{\Lambda}^{up}\|_{L^2(\Lambda)}^2\right] \le C\Bigg(\underbrace{\frac{\log M}{M^{2s/(2s+m)}}}_{\text{KDE error}} + \underbrace{\|f(x) - G(\hat{\mathbf{W}}_A^T x)\|_{L^2(\Lambda)}^2}_{\text{Active subspace error}} \Bigg)$$

[a] Not a kernel density estimate of the updated density!

Example

Consider a parameterized nonlinear system of equations:

$$
\begin{aligned}
\lambda_1 x_1^2 + x_2^2 &= 1, \\
x_1^2 - \lambda_2 x_2^2 &= 1
\end{aligned}
$$

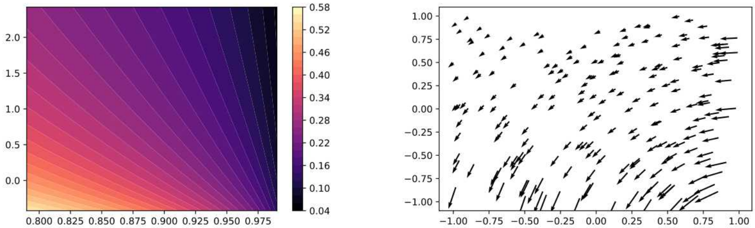

Figure: Response surface for the QoI (left) and samples of the gradient in normalized coordinates (right).

A Parameterized Nonlinear System

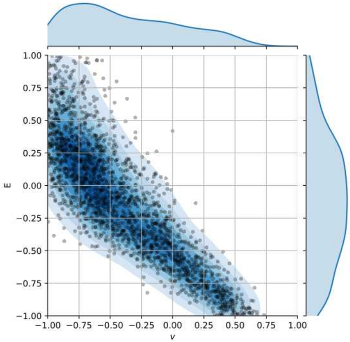

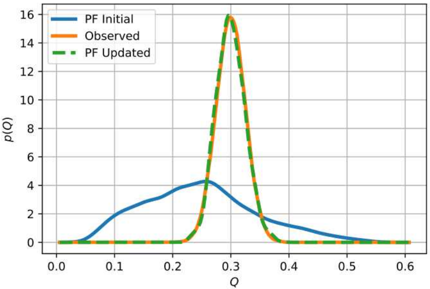

Figure: Samples from the updated density (left) and a comparison of $\pi_{\mathcal{D}}^{\text{obs}}$, $\pi_{\mathcal{D}}^{Q(\text{init})}$ and $\pi_{\mathcal{D}}^{Q(\text{up})}$ (right).

A Parameterized Nonlinear System

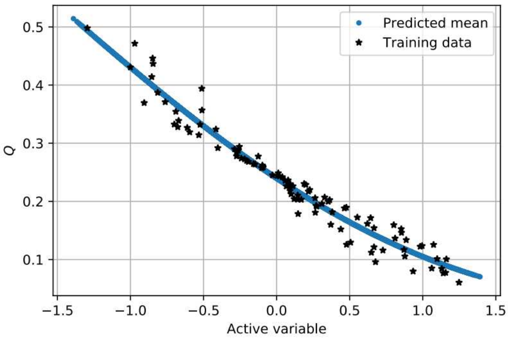

Figure: Samples of the active variable and the true model response as well as the regression model for $N = 100$.

A Parameterized Nonlinear System

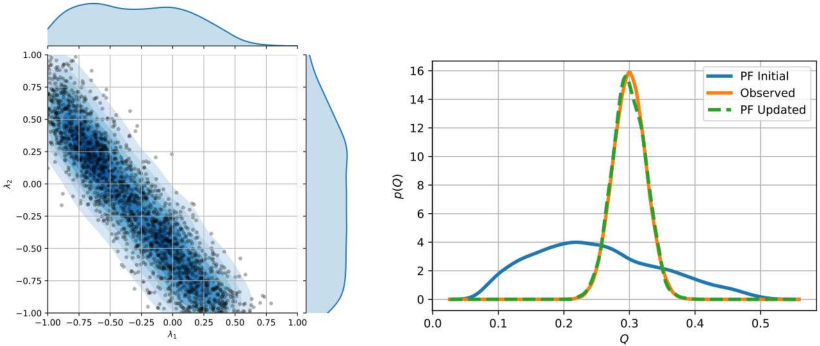

Figure: Samples from the updated density (left) and a comparison of $\pi_{\mathcal{D}}^{\text{obs}}$, $\pi_{\mathcal{D}}^{Q(\text{init})}$ and $\pi_{\mathcal{D}}^{Q(\text{up})}$ (right) using a 1-dimensional active subspace with $N = 100$ and $M = 10,000$.

A Parameterized Nonlinear System

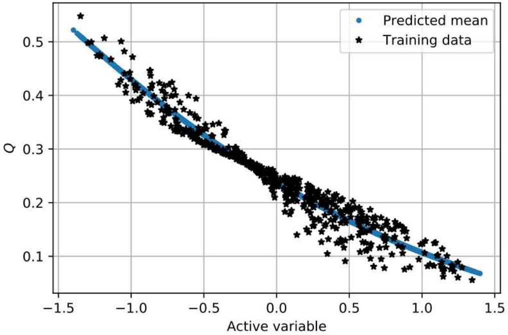

Figure: Samples of the active variable and the true model response as well as the regression model for $N = 500$.

A Parameterized Nonlinear System

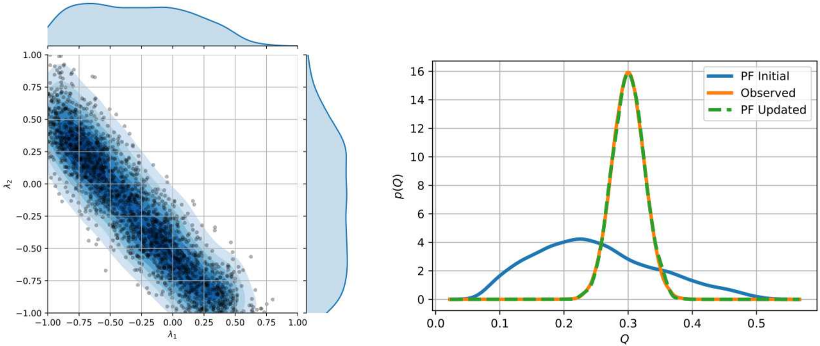

Figure: Samples from the updated density (left) and a comparison of $\pi_{\mathcal{D}}^{\text{obs}}$, $\pi_{\mathcal{D}}^{Q(\text{init})}$ and $\pi_{\mathcal{D}}^{Q(\text{up})}$ (right) using a 1-dimensional active subspace with $N = 500$ and $M = 10,000$.

Comparison with Reference Solution

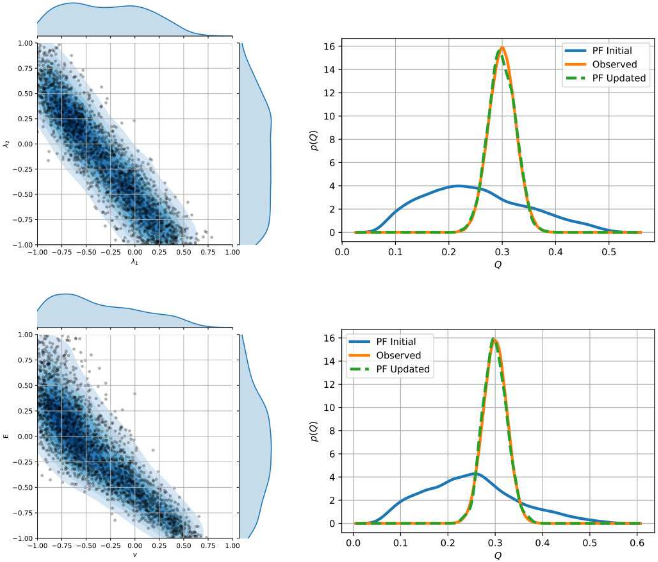

Comparison Using 2-dimensional Active Subspace

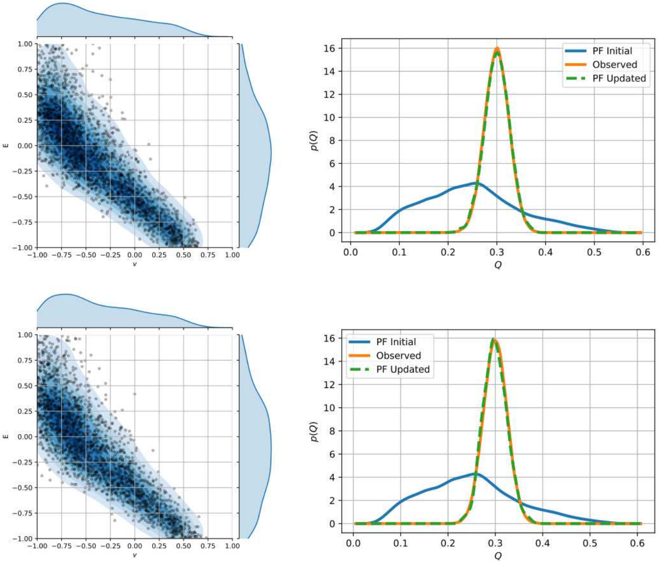

A Predator-Prey System

Example

Consider a Lotka-Volterra system:

$$\frac{\partial u_i}{\partial t} = r_i u_i \left(b_i - \sum_{j=1}^{3} A_{ij} u_j \right), \quad i = 1, 2, 3$$

- The initial conditions, $u_i(0)$, and self-interaction terms, A_{ii}, are known.
- Leaves 9 random parameters, initial distribution assumes independent uniformly distributed over $[0.3, 0.7]$.
- Use 4th-order explicit Runge-Kutta method to solve to $T = 50$ with $\Delta t = 0.01$
- Quantity of interest is $u_3(T)$
- Observed distribution is given by $N(0.3, 0.025^2)$

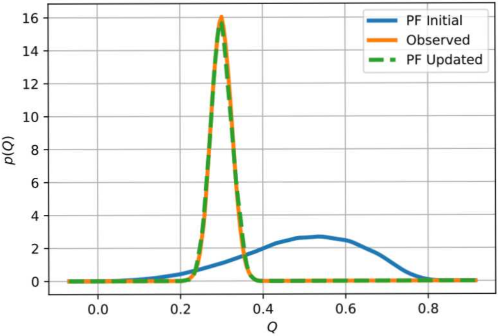

Figure: Samples from the updated density (left) and a comparison of $\pi_{\mathcal{D}}^{\text{obs}}$, $\pi_{\mathcal{D}}^{Q(\text{init})}$ and $\pi_{\mathcal{D}}^{Q(\text{up})}$ (right) using a 1-dimensional active subspace with $N = 500$ and $M = 10,000$.

A Predator-Prey System

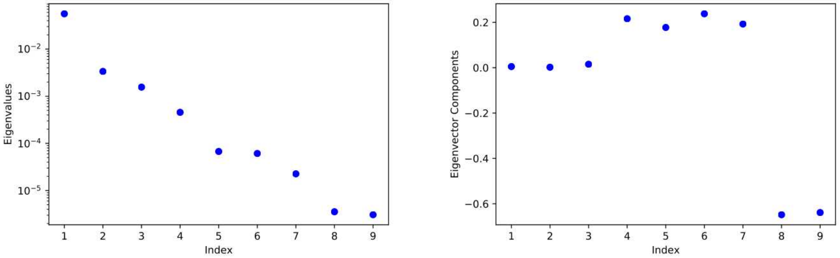

Figure: Samples from the updated density (left) and a comparison of $\pi_{\mathcal{D}}^{\text{obs}}$, $\pi_{\mathcal{D}}^{Q(\text{init})}$ and $\pi_{\mathcal{D}}^{Q(\text{up})}$ (right) using a 1-dimensional active subspace with $N = 500$ and $M = 10,000$.

A Predator-Prey System

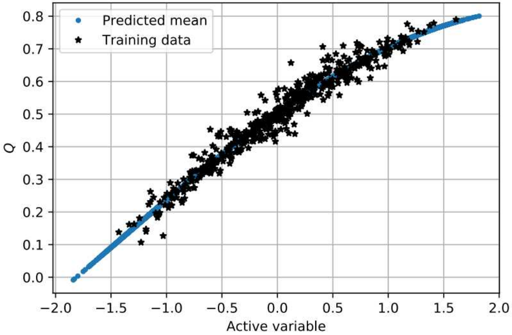

Figure: Samples from the updated density (left) and a comparison of $\pi_{\mathcal{D}}^{\text{obs}}$, $\pi_{\mathcal{D}}^{Q(\text{init})}$ and $\pi_{\mathcal{D}}^{Q(\text{up})}$ (right) using a 1-dimensional active subspace with $N = 500$ and $M = 10,000$.

A Predator-Prey System

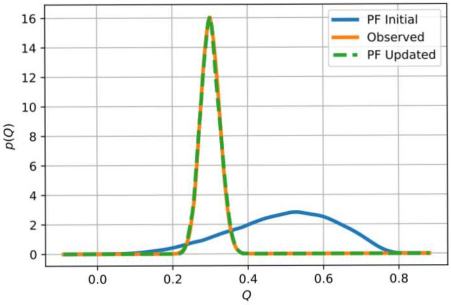

Acceptance rate	0.0675
Mean of $\pi_{\mathcal{D}}^{Q(\text{up})}$	0.2997
St. dev. of $\pi_{\mathcal{D}}^{Q(\text{up})}$	0.0251
Integral of $\pi_{\Lambda}^{\text{up}}$	0.9952
$\text{KL}(\pi_{\Lambda}^{\text{init}} \vert \pi_{\Lambda}^{\text{up}})$	2.1604

A Predator-Prey System

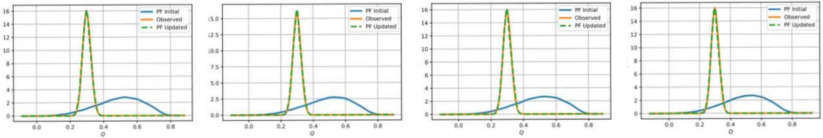

	Dimension of active subspace				Reference
	1	2	3	4	
Acceptance rate	0.0675	0.0687	0.0672	0.0675	0.0663
Mean of $\pi_{\mathcal{D}}^{Q(\text{up})}$	0.2997	0.3000	0.3003	0.2999	0.2999
St. dev. of $\pi_{\mathcal{D}}^{Q(\text{up})}$	0.0251	0.0252	0.0253	0.0254	0.0254
Integral of $\pi_{\Lambda}^{\text{up}}$	0.9952	0.9991	0.9958	0.9979	0.9965
$\text{KL}(\pi_{\Lambda}^{\text{init}} \| \pi_{\Lambda}^{\text{up}})$	2.1604	2.1646	2.1517	2.1804	2.1838

A Computational Mechanics Example

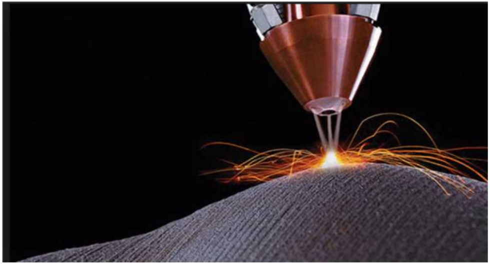

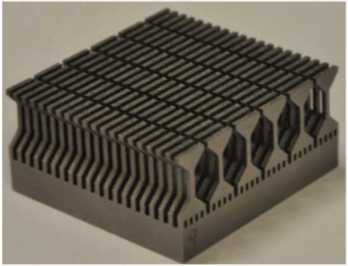

Figure: Additive manufacturing and high-throughput testing provides new data challenges.

A Computational Mechanics Example

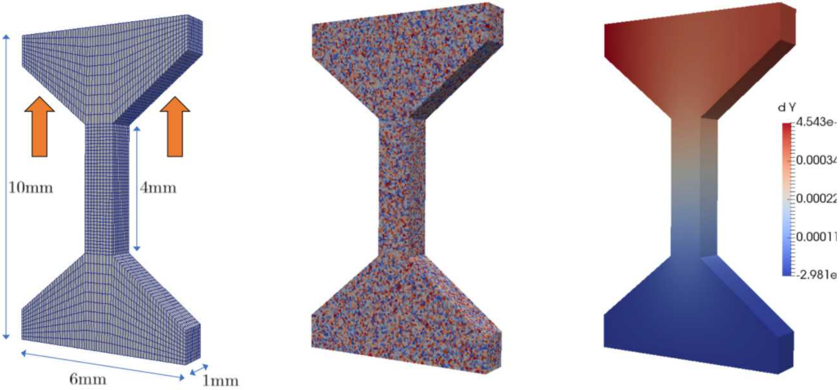

Figure: On the left, an illustration of the computational model on a coarse mesh (16,600 elements). In the middle, the granular microstructure on a finer mesh ($\approx$ 17 million elements). On the right, the vertical displacement using the high-fidelity model and nominal parameter values.

A Computational Mechanics Example

- We assume the Young's modulus, E, and Poisson ratio, ν, are random.
- Each grain has a random orientation defined by 4 independent Gaussian parameters.
- Model has $\approx 225,000$ grains.

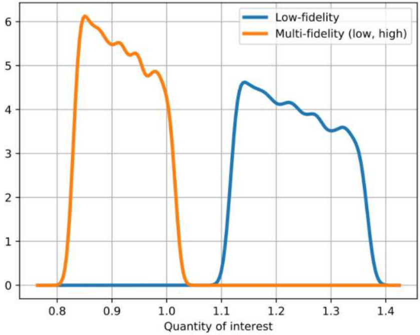

Figure: Comparison of the $\pi_{\mathcal{D}}^{Q(\text{init})}$ using the low-fidelity model and the high-fidelity model.

A Computational Mechanics Example

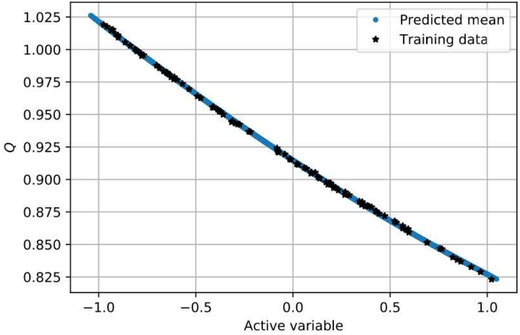

Figure: Samples of the active variable and the QoI as well as the regression model for $N = 100$.

A Computational Mechanics Example

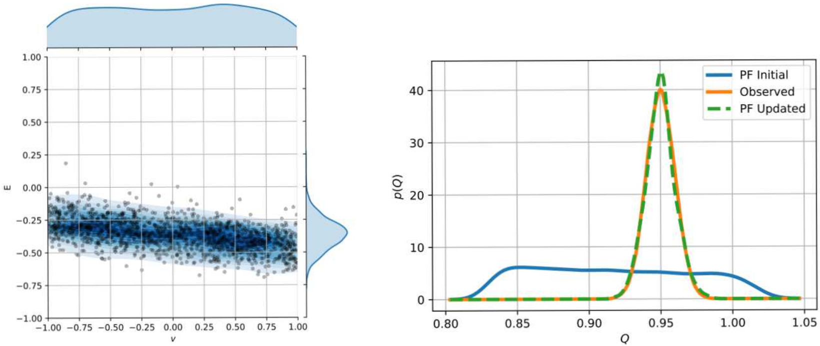

Figure: Samples from the updated density (left) and the comparison of $\pi_{\mathcal{D}}^{\text{obs}}$, $\pi_{\mathcal{D}}^{Q(\text{init})}$ and $\pi_{\mathcal{D}}^{Q(\text{up})}$ (right) using a 1-dimensional active subspace with $N = 100$ and $M = 10,000$.

Conclusions

- Our goal is to develop **data-informed physics-based** models.
- Many approaches exist for incorporating data into a model.
  - Deterministic optimization, Bayesian methods, OUU, data assimilation, etc.
- Our approach provides a solution to a specific stochastic inverse problem.
- Main computational expense is the forward UQ problem to obtain the push-forward of the initial density.
- We demonstrated that **dimension reduction techniques** can be utilized within this framework.
- Theoretical results show that the **errors in the push-forward and in the updated density are bounded by the errors in the active subspace model**.
- Didn't quite get to *inference for prediction* ...

Thanks! Questions?

Acknowledgments

T. Wildey's work was supported by the U.S. Department of Energy, Office of Science, Early Career Research Program.

Thank you for your attention!

Questions?