SAND2022-9156C

MLDL

Machine Learning and Deep Learning Conference 2022

Embedded uncertainty estimation for data-driven surrogates to enable trustworthy ML for UQ

- Tim Wildey (1441) and Gianluca Geraci (1463)
- Michael S. Eldred, John Jakeman, Owen Davis and Teresa Portone (1463)
- Tien Yu Yen (1441)
- Bryan Reuter (1446)
- Alex Gorodetsky (University of Michigan)
- Ahmad Rushdi (Stanford University)
- Daniele Schiavazzi, Lauren Partin (University of Notre Dame)

- Funding Sources: LDRD, ASCR and ASC V&V
- DC: Dan Turner (1463)

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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.

Work from a Large Group of Collaborators

1441 - Scientific Machine Learning

Tim Wildey

Tian Yu Yen

1446 – Computational Science

Bryan Reuter

1463 - Optimization and UQ

Owen Davis

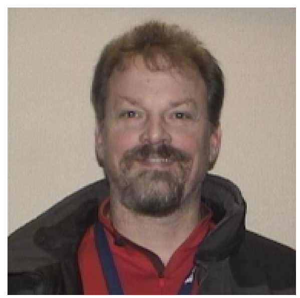

Michael Eldred

Gianluca Geraci

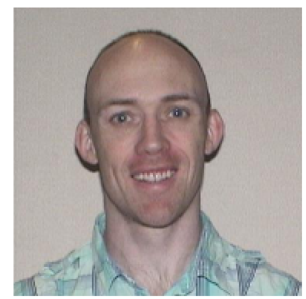

John Jakeman

Teresa Portone

External Collaborators

Ahmad Rushdi (Stanford)

Alex Gorodetsky (Univ. of Michigan)

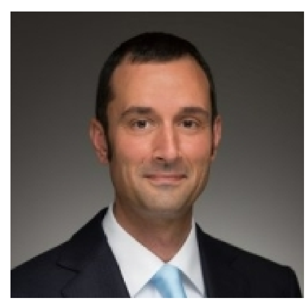

Daniele Schiavazzi (Notre Dame)

Motivation/Goals

We (1441 and 1463) are broadly interested in using ML to accelerate/enhance our ability to perform UQ and solve inverse problems.

Can we estimate/assess the uncertainty in predictions made by data-driven models to improve their trustworthiness?

Data-driven Foundational Activities

- Non-hierarchical approaches (DARPA, ASC, LDRD) [Gorodetsky *et al.*, JCP 2020, Gorodetsky *et al.*, Comp. Mech 2021]
- Hyperparameter optimization (ASC) [Bomarito *et al.*, SciTech 2022]
- Incorporating non-deterministic models (LDRD, ASC, SciDAC) [Geraci *et al.*, SIAM CSE 2021]
- Heterogeneous input parametrization (LDRD, SciDAC) [Geraci, Eldred, SAND Rep. 2018]

Embedded UQ for ML

- Stochastic inverse problems (ASCR) [Butler et al, SISC 2018]
- Ensembles for stochastic inverse problems (ASCR, ASC) [Wildey et al, SIAM UQ 2022]
- MF Convolutional Neural Networks with embedded error estimation (ASC) [Partin et al., SciTech 2022]

Multi-Fidelity UQ threads

2018/2019:

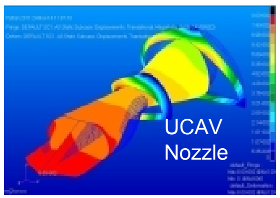

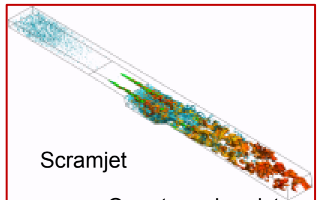

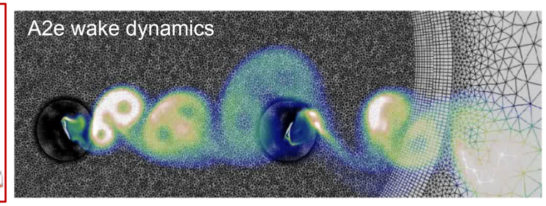

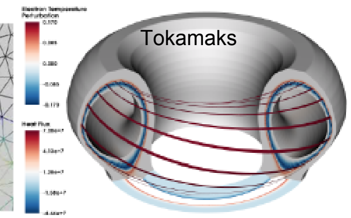

2020/2021:

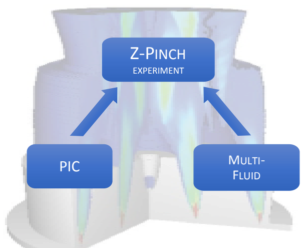

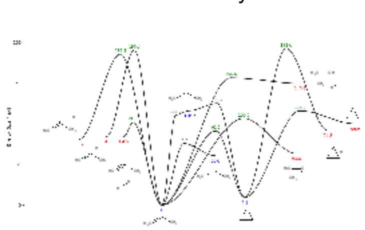

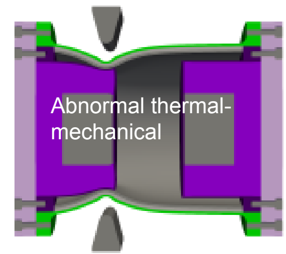

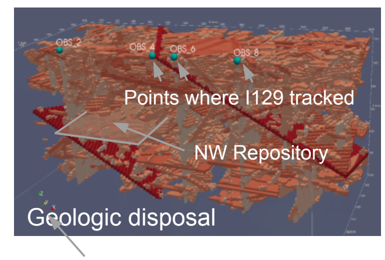

Monte Carlo UQ Methods

- *Production:* optimal resource allocation for multilevel, multifidelity, combined (DARPA EQUiPS, Wind, Cardiovascular)
- *Emerging:* active dimensions (LDRD, SciDAC), generalized fmwk for approx control variates (ASC V&V), goal orientation (rare events), hybrid methods for GSA
- *On the horizon:* control of time avg; model tuning / selection (LDRD)

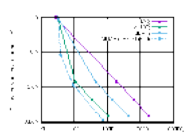

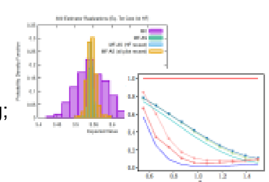

Surrogate UQ Methods (PCE, SC)

- *Production (v6.10+):* ML PCE w/ projection & regression; ML SC w/ nodal/hierarchical interp; greedy ML adaptation (DARPA SEQUOIA), multilevel fn train (ASC V&V)
- *Emerging:* multi-index stochastic collocation; multiphysics/multiscale integration (ASC V&V); new surrogates (GP, ROM, NN) w/ error mgmt. fmwk (LDRD, SciDAC); learning latent variable relationships (MFNets, LDRD)
- *On the horizon:* unification of surrogate + sampling approaches (LDRD)

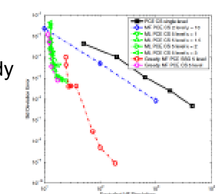

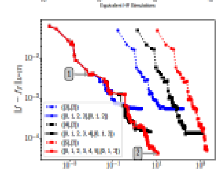

Optimization Under Uncertainty

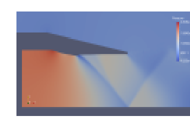

- *Production:* manage simulat and/or stochastic fidelity
- *Emerging:*
  Derivative-based methods (DARPA SEQUOIA)
  - Multigrid optimization (MG/Opt)
  - Recursive trust-region model mgmt.: extend TRMM to deep hierarchies

  Derivative-free methods (DARPA Scramjet)
  - SNOWPAC (w/ MIT, TUM) with goal-oriented MLMC error estimates
- *On the horizon:* Gaussian process-based approaches: multifidelity EGO; Optimal experimental design (OED)

Non-hierarchical data sources

Challenge

Common hierarchical approaches are limited in representing realistic data sources

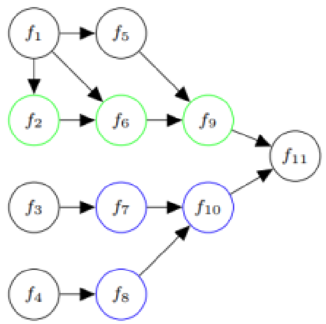

Multifidelity Bayesian Networks (MFNets)

- Linear subspace models represent each model
- Direct Acyclic Graph relationships are encoded via conditional independencies on coefficients
- Similarly to Gaussian Processes, given a prior we can compute the posterior in closed-form
- Efficient inference is possible for sparsely connected Gaussian Bayesian networks
- A shared manifold can be to increase models' correlation

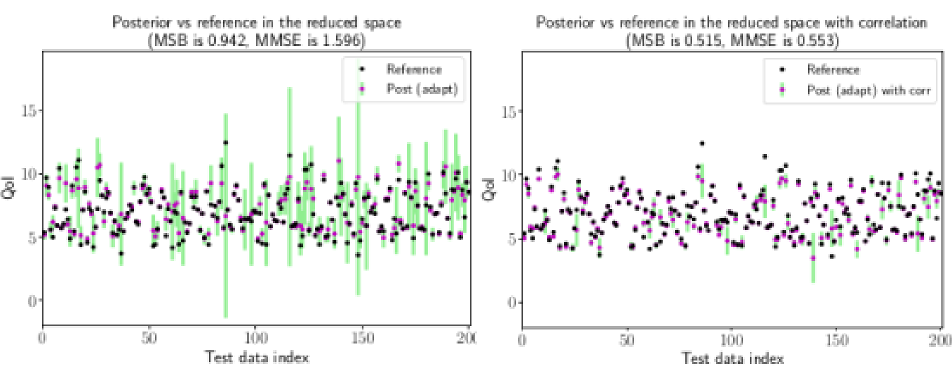

MFNets W/ (left) and W/O (right) shared manifold learning

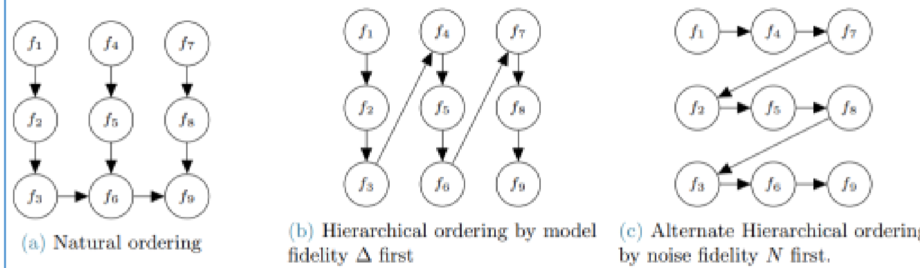

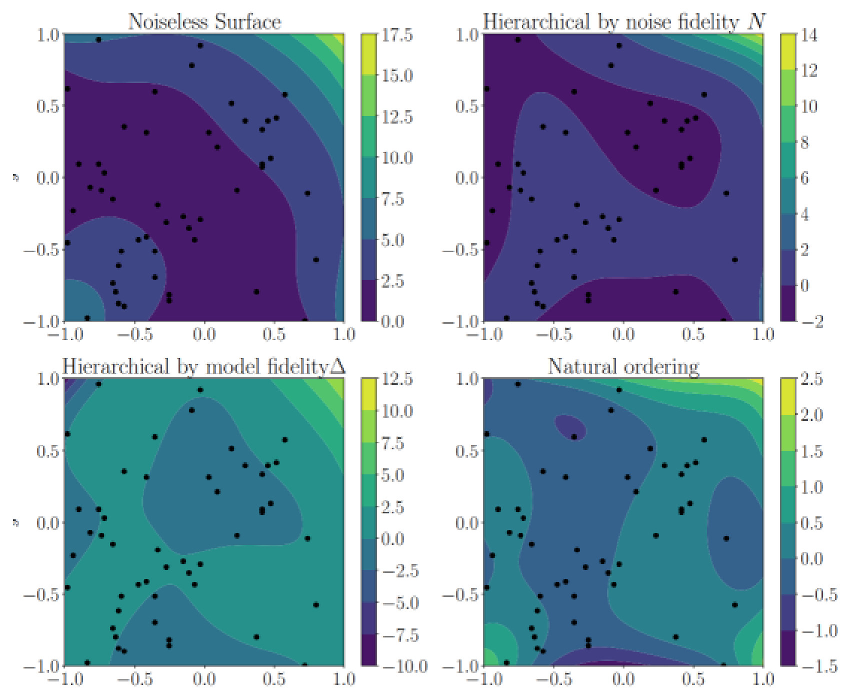

Model problem illustrating surrogate accuracy exploiting different graphs

Hyperparameter Optimization

Can we design LF or data-driven models to be optimal for the MF UQ task?

- Both data-driven and computational models contain hyper-parameters
- Selection of hyper-parameters, within MF UQ, is a function of the models' ensemble and selected method

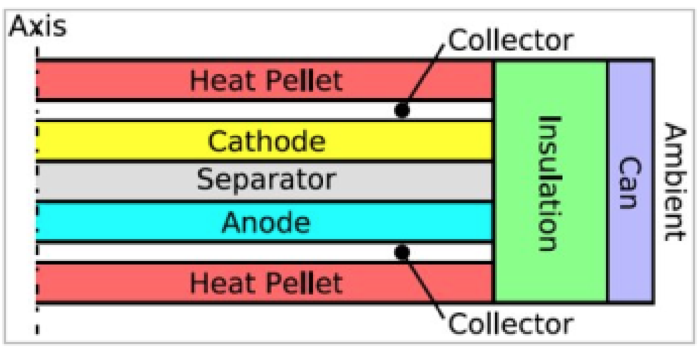

Single-cell thermal battery exemplar.

24.8x

Hand-tuned hyper-parameters:

0.01	initial time step
0.10	predictor-corrector tol
0.10	nonlinear residual tol

Projected ACV Estimator Variance: .053138
Single fidelity accuracy for equiv cost: 1.3178 (1005 HF)
Single fidelity cost for equiv accuracy: 24,925 HF (EstVar .053138)

143x

Automatic tuning hyper-parameters:

0.0067487	initial time step
0.0010880	predictor-corrector tol
0.046707	nonlinear residual tol

Projected ACV Estimator Variance: 0.0092395
Single fidelity accuracy for equiv cost: 1.3192 (1004 HF)
Single fidelity cost for equiv accuracy: 143,340 HF (EstVar 0.0092395)

Incorporating Non-deterministic Models

<u>*What is the effect of non-deterministic models on MF UQ?*</u>

- Several applications at Sandia involve embedded stochasticity, *e.g.* cybersecurity, radiation transport, Particle-in-Cell, *etc.*
- Data-driven methods, e.g., neural networks, reduced order models, *etc.,* are often built using noisy data and/or non-deterministic optimization methods
- Complex models are too expensive and finite averaging need to be handled within MF UQ
- The correlation among models is reduced by the models' stochasticity

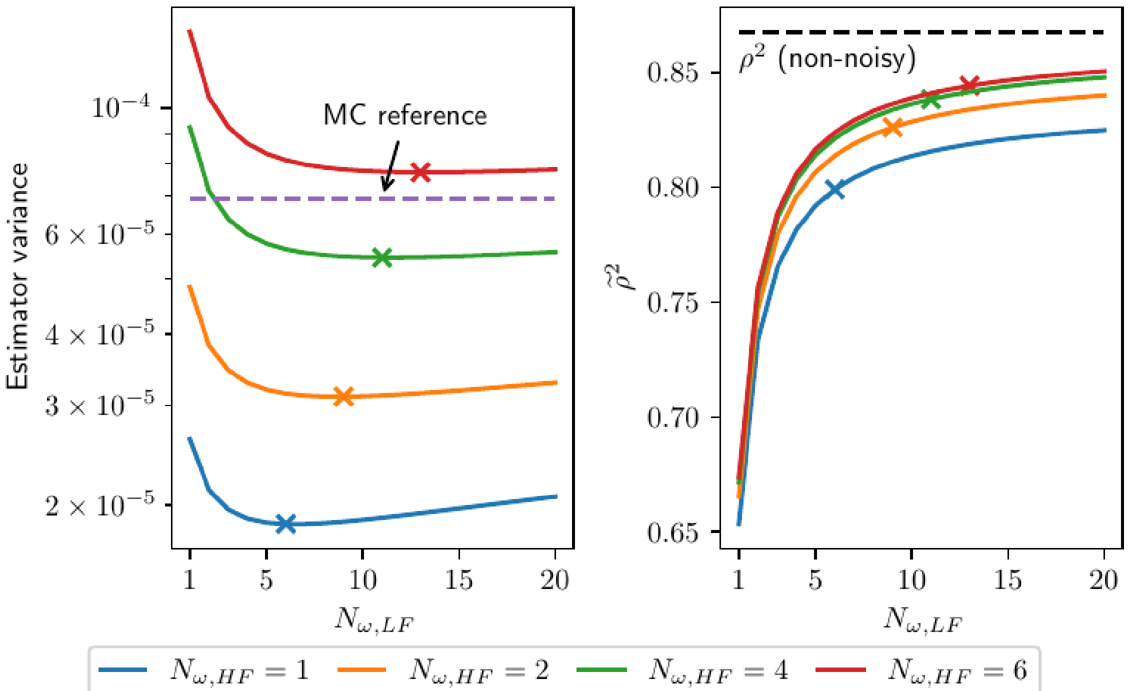

Estimator variance as a function of the averaging (for a fixed computational budget).

MF CNNs with embedded error estimation

How can efficiently train MF CNNs and get an embedded prediction error?

- Convolutional Neural Networks reduce the weights w.r.t. fully connected networks
- CNNs are assembled from *encoders*, *decoders* and *skip connections*
- We propose MF CNNs in which all models are learned simultaneously in an *all-at-once training* paradigm
- We explored the "UQ for ML" aspect by characterizing the uncertainty in predictions that are inherent with the ML surrogate
- We employed DropBlock for the "UQ for ML" task

Application scenarios

- Verification tests:
  - Simple *1D regression* problems (two model forms)
- Computational fluid dynamics (predicting pressure distribution)
  - *Dense regression* – input and output are images of same size
  - Low- to high-dimensional regression – high-dimensional output predicted by low-dimensional input

MF CNNs – Building blocks

DropBlocks
- Dropout layers are widely used for regularization
- Dropout layers operate by dropping neurons at random during training → this corresponds to a simultaneous training of an ensemble of architerctures
- DropBlocks adapt this idea to CNNs by dropping a group of pixels
- *MC DropBlock* consists in activating DropBlocks at evaluation time in conjunction with MC sampling → this approach can be used to quantify the stochasticity in network predicions

Convolutional encoder
Alternating layers of convolutions and pooling (down-sampling)

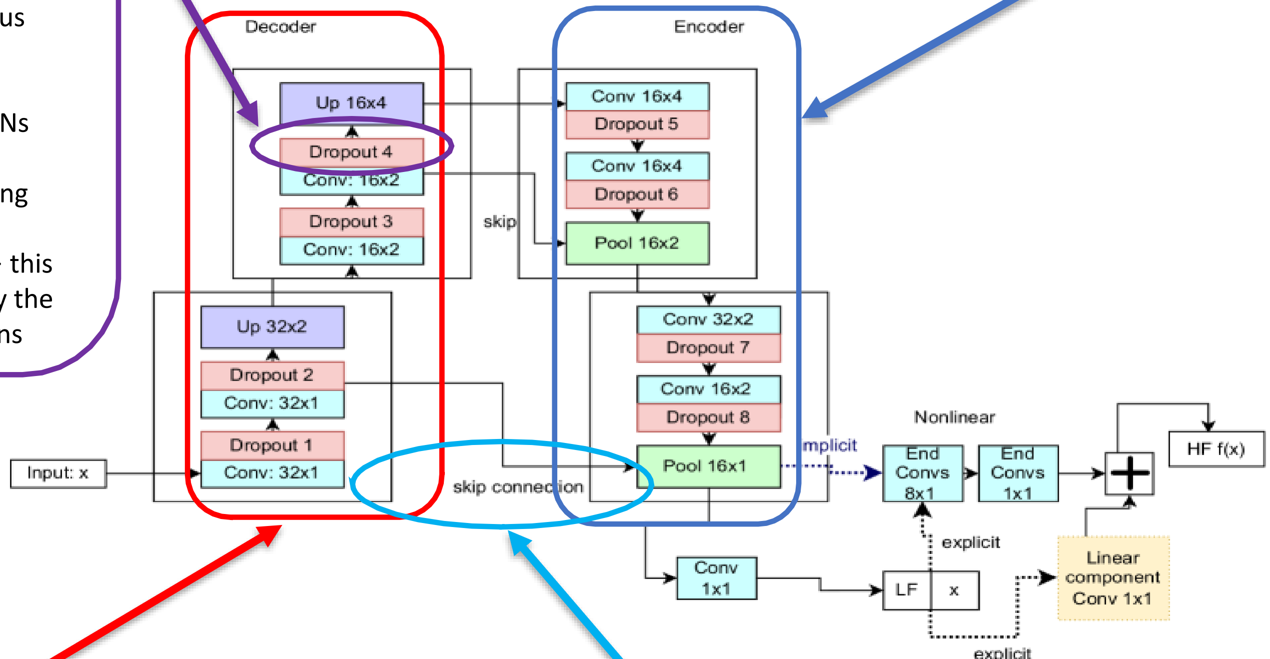

Convolutional decoder
Alternating layers of convolutions and upsampling

Skip connections
Added to mitigate the loss of information due to downsampling

Decoder-Encoder architecture for 1D regression

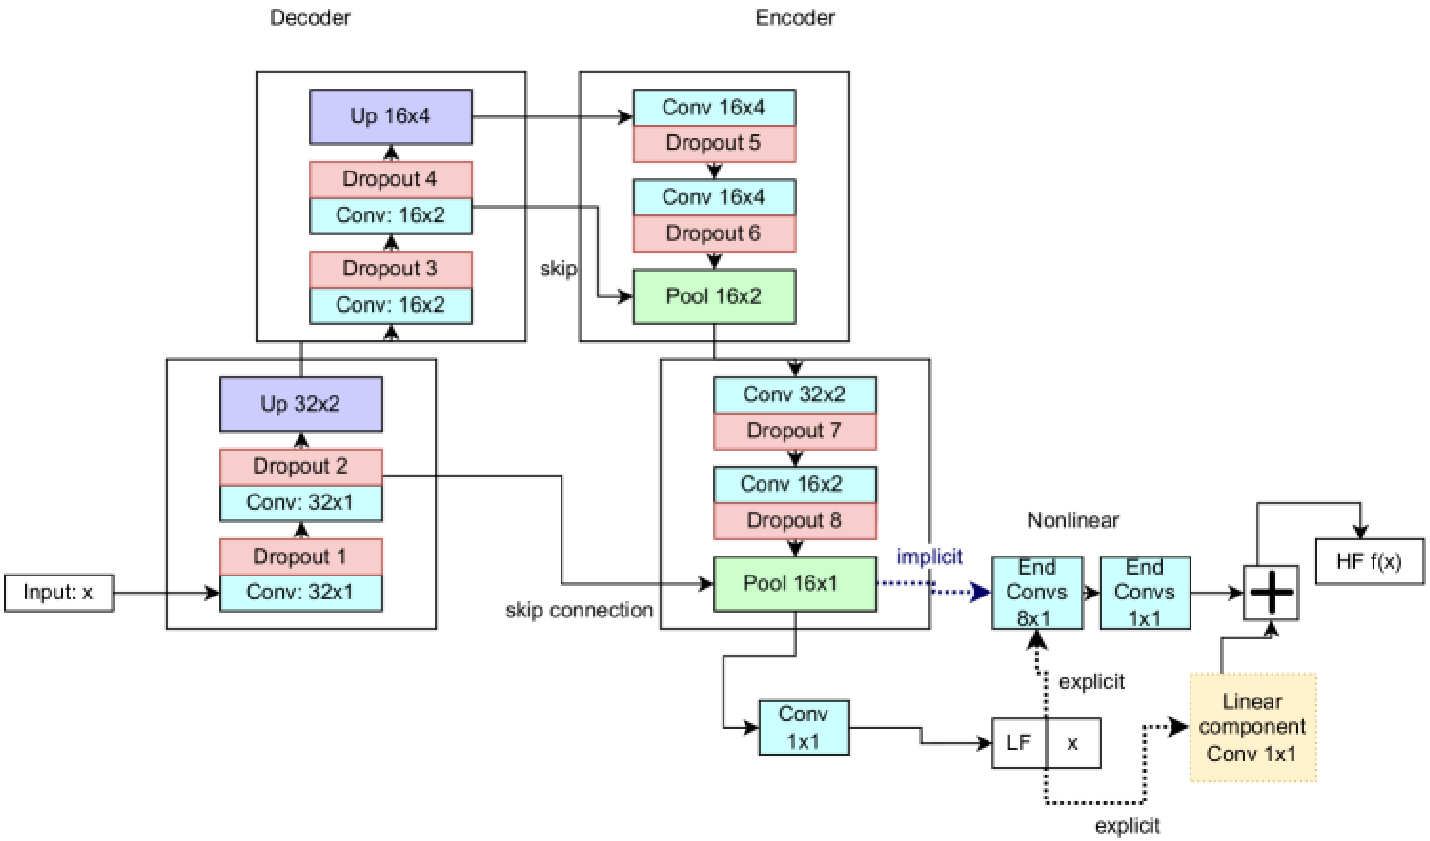

- Concatenation between LF prediction and x coordinate
- HF predictions obtained by summing two contribution
  - Convolution designed to capture HF and LF linear correlation
  - Convolution designed to capture the HF and LF nonlinear correlation

MF CNNs – 1D regression results

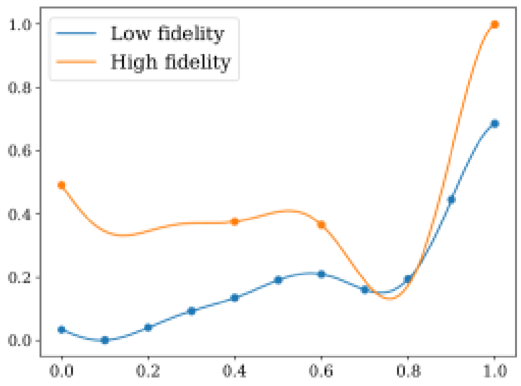

Dataset

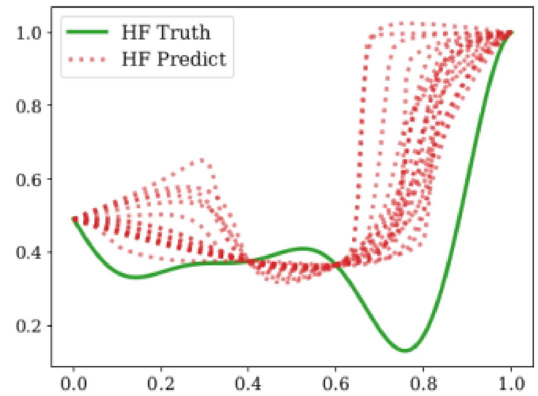

Single-fidelity

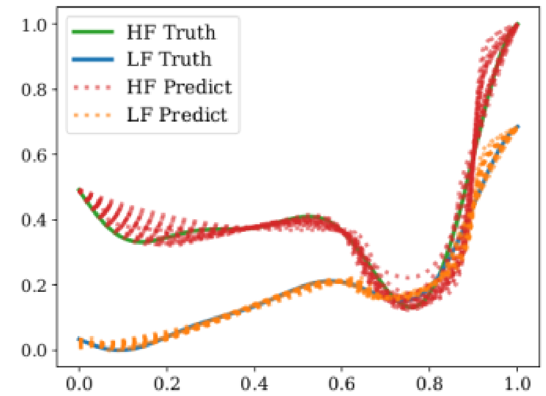

Multi-fidelity

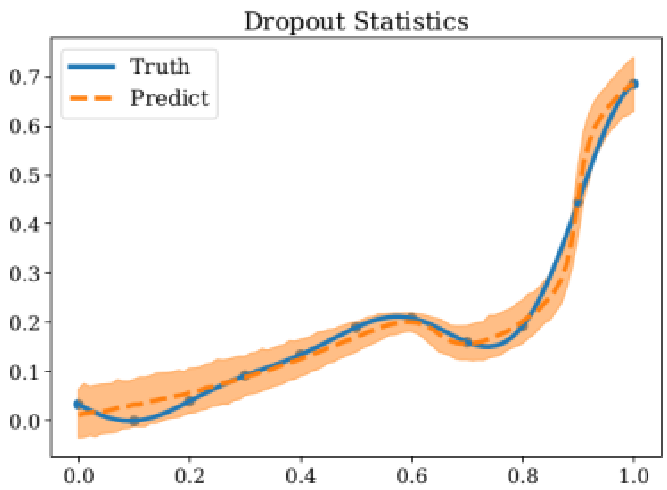

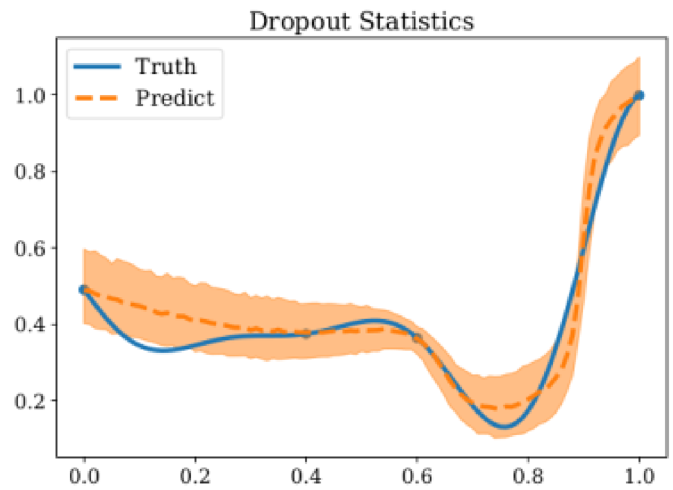

Mean and estimated 5%-95% percentiles for the low- and high-fidelity via DropBlock realizations after every convolution layer

Encoder-Decoder architecture for dense regression

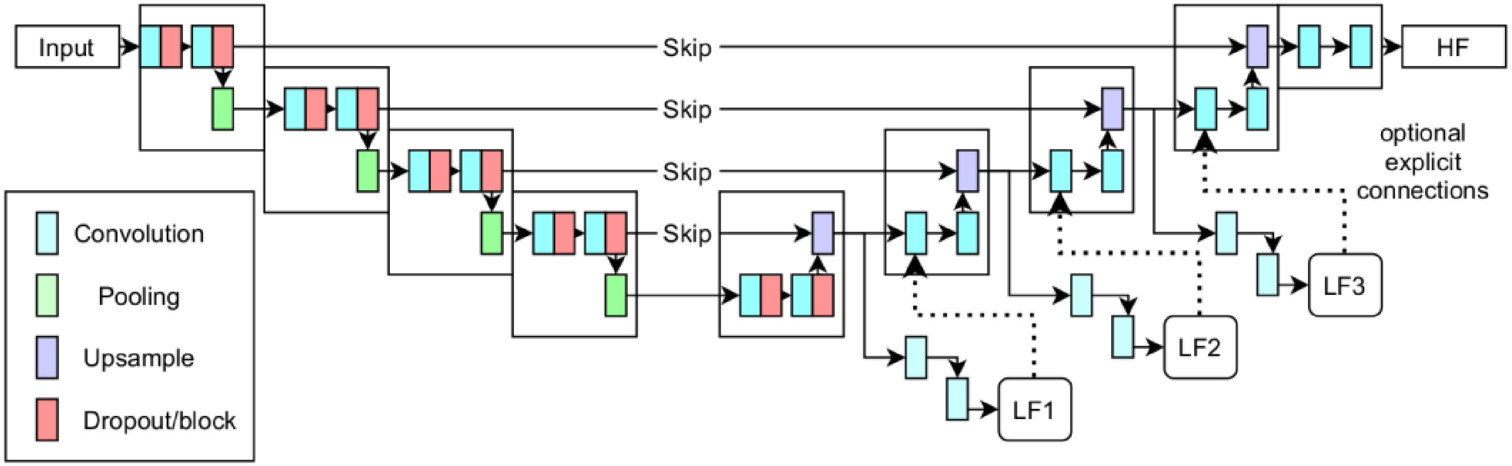

- Input-Output are 64x64 images
- 3 LF models are generated at the decoder stage

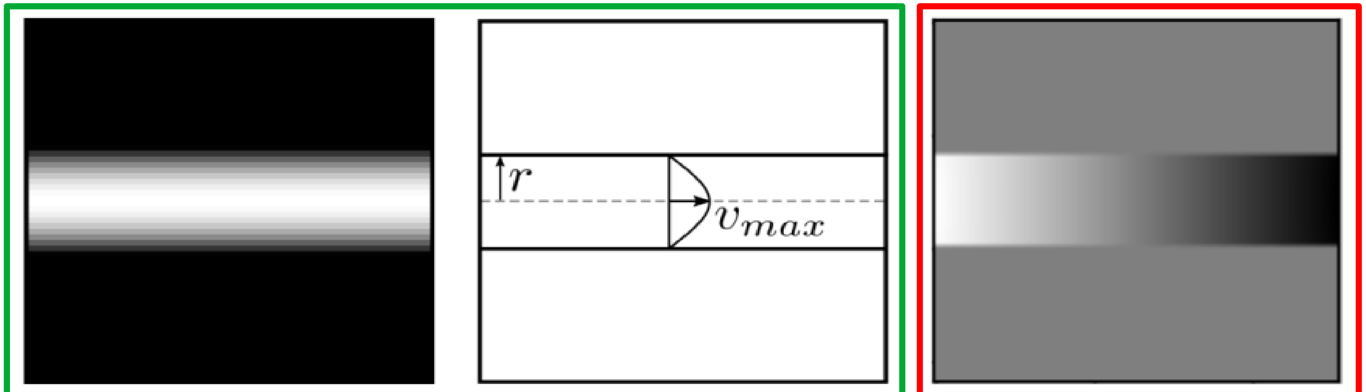

(a) (b) (c)

- **INPUT:** concentration field and velocity
- **OUTPUT:** pressure field
- Dataset: 32 HF + 116 LF (for each resolution)

MF CNNs – Dense regression results

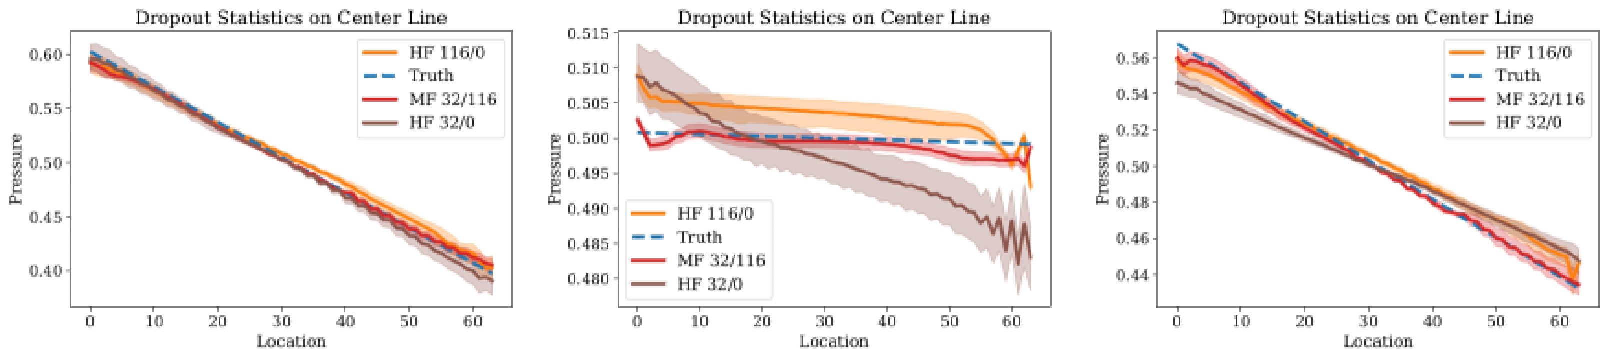

(a) Sample 1 (b) Sample 2 (c) Sample 3

Mean prediction and 5%-95% percentiles from an ensemble of 1000 DropBlock pressure realizations for three test data examples sliced along the cylinder centerline

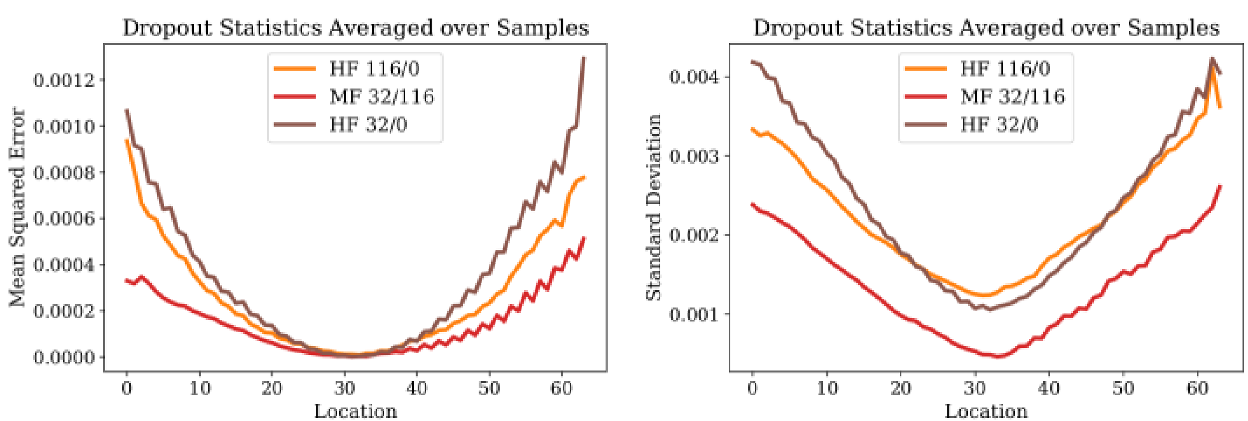

Mean square error and standard deviation resulting from 1000 network evaluations with 10 DropBlock layers in each network.

Low-to-high dimensional architecture for dense regression

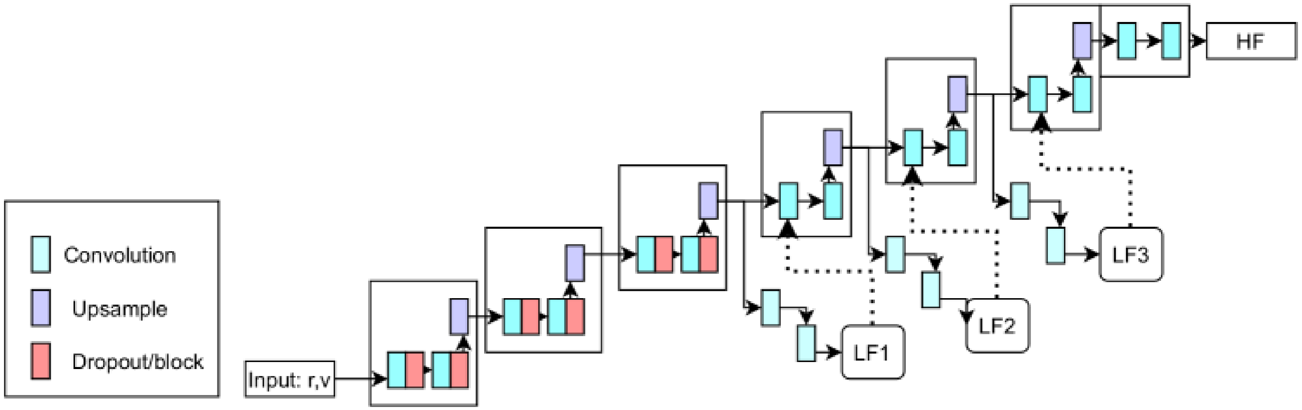

- Input: radius and max velocity
- Output are 64x64 images for the pressure
- 3 LF models are generated at the decoder stage

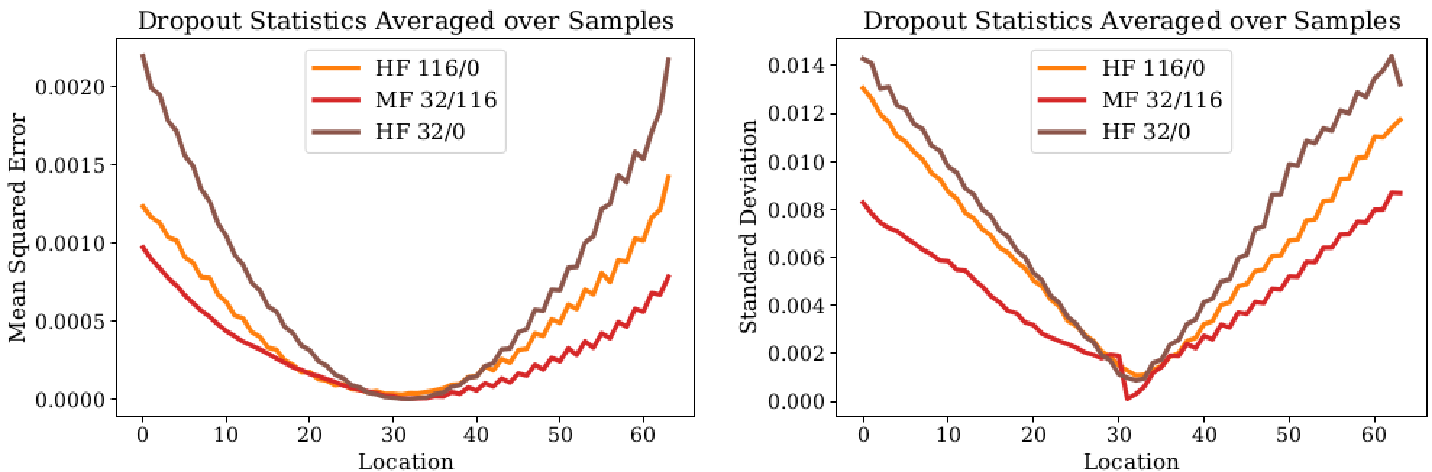

Mean square error and standard deviation resulting from 1000 network evaluations with 6 DropBlock layers in each network.

Low-to-high dense regression results

- HF 32 →
- HF/LF 32/116 →
- HF 32 →
- HF/LF 32/116 →

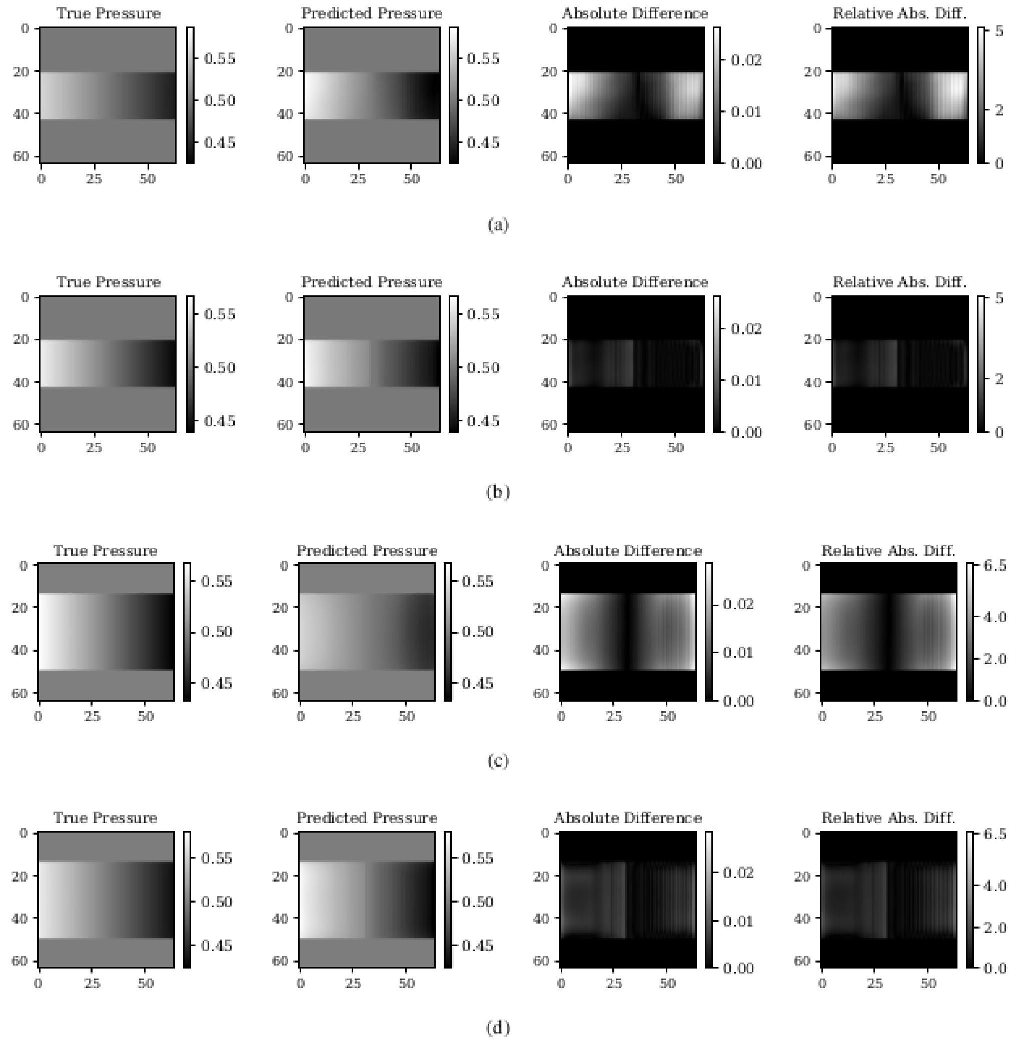

Using Ensembles to Estimate Uncertainty in Solutions to Stochastic Inverse Problems (ASC+ASCR)

- Data-consistent inversion [BJW, 2018a] seeks to solve the following problem:
  Given a target distribution on outputs (QoI), find a distribution on inputs such that the push-forward of this distribution matches the target.
- Solution is given by:

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

- Requires performing a forward uncertainty propagation, which is the dominant cost.
- Using an approximate model [BJW, 2018b] and even a neural network [Zhang, 2021]
- Can we use an ensemble based approach to estimate the error in the inverse solution?

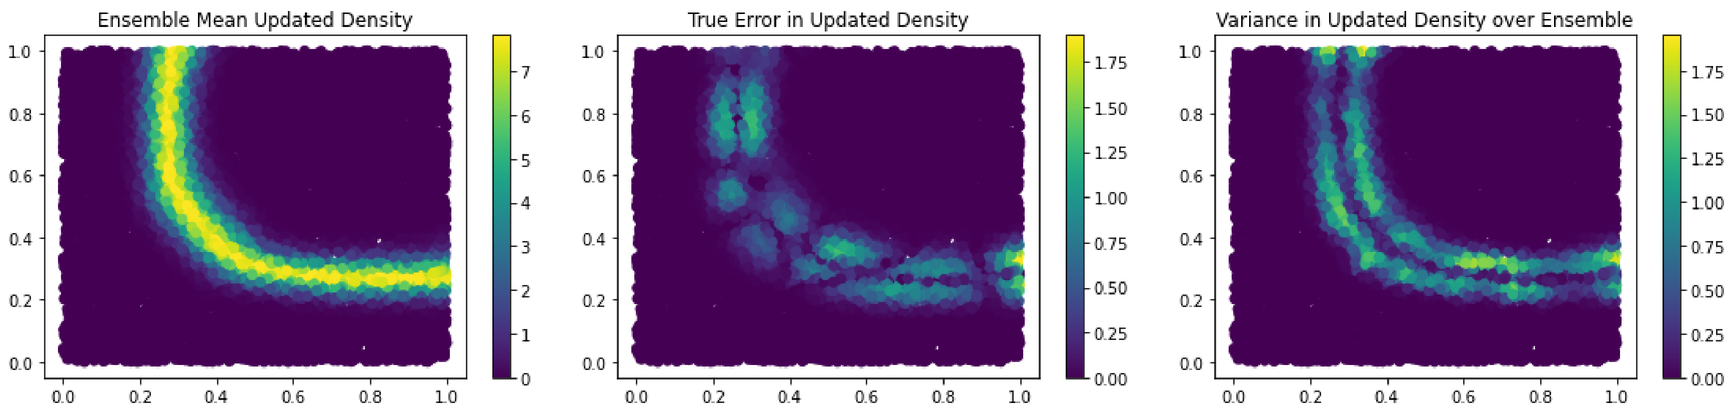

Ensemble averaged solution to a stochastic inverse problem, the true error in the solution, and the variance of the solution over the ensemble.

Using ML to Enhance our Ability to Solve Stochastic Inverse Problems (ASCR)

Motivation

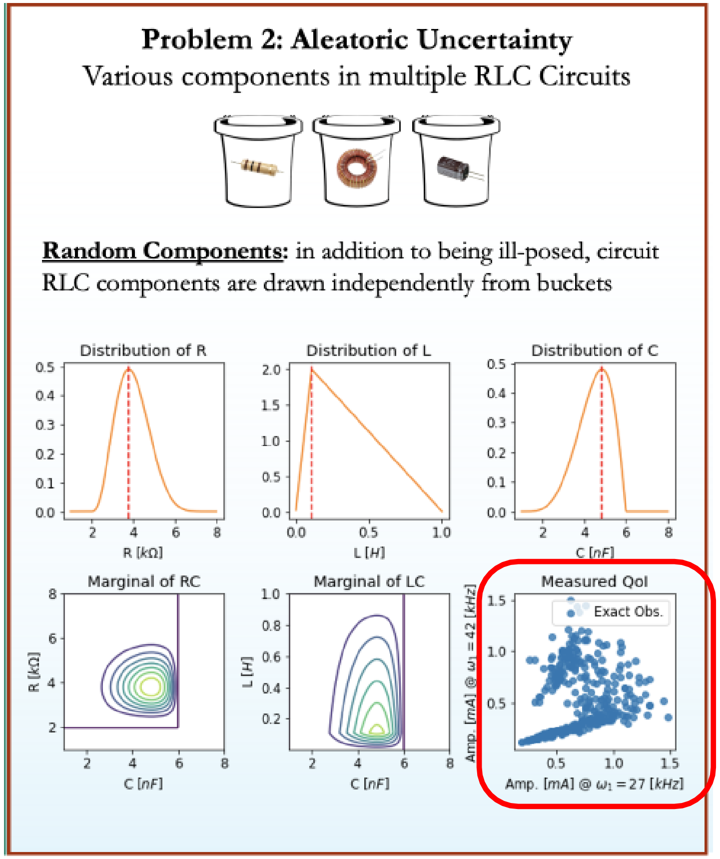

Density estimation is often very difficult!

Approach

- Use ML to improve the density estimation - unsupervised.
- Normalizing flows seek a mapping to a target distribution.
- Works well in high-dimensions, but does not necessarily preserve structure.
- Principal manifold flows seek to find a mapping that also preserves contour structure [Cunningham et al 2022]

Results

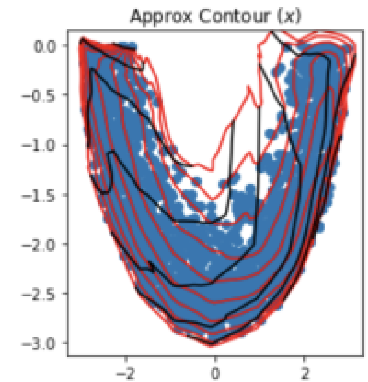

Standard normalizing flow is decent.

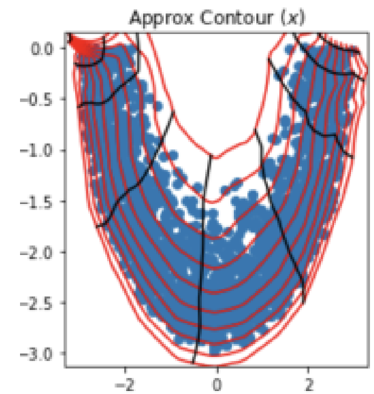

Principal manifold flow preserves structure.

Conclusions

- Data-driven methods have been employed widely to accelerate/enhance UQ.
- Data-driven surrogate models have errors/uncertainties arising from a wide variety of sources.
- Assessing the impact of these errors/uncertainties is necessary to have confidence in ML-based predictions.
- Multi-fidelity data fusion and data-consistent inversion can both utilize ML in various ways.
- The ability to understand and control the sources of uncertainties in data-driven surrogate construction are fundamental technology for both machine learning and uncertainty quantification
- High-fidelity Sandia-relevant applications require strategies with affordable data requirements and embedded measures of trustworthiness.

Comments/Questions?

tmwilde@sandia.gov and ggeraci@sandia.gov

Backup Material

Heterogeneous input parametrization

- *How can we handle heterogeneous models' input parametrization?*

- The presence of different physics (in models used for MF UQ) often leads to a dissimilar input parametrization
- Dimension reduction strategies, *e.g.* Active Subspace and Basis Adaptation, can be used to obtain a shared manifold among models' input
- This shared manifold can be used to link models with dissimilar parametrization