Visualizing and Quantifying Uncertainty of Physics-aware Neural Networks

FASTMath+RAPIDS Exploratory 1yr Project: FY24, $250k

Benjamin Erichson (LBL), Khachik Sargsyan (SNL)

RAPIDS

FASTMATH

FASTMath All Hands, Denver, CO, Nov 2-3, 2023

Main Theses

Trustworthy SciML requires
Uncertainty Quantification (UQ)

Accurate UQ for Neural Networks (NNs)
hinges on the loss surface's behavior

Physics-driven regularization
will improve loss surface and
enable more accurate and efficient UQ

Probabilistic NN == Bayesian NN

Ghahramani, "Probabilistic Machine Learning and Artificial Intelligence". Nature, 2015

"Nearly all approaches to probabilistic programming are ***Bayesian*** *since it is hard to create other coherent frameworks for automated reasoning about uncertainty"*

- Bayesian NN methods have been around since 90s *[MacKay, 1992; Neal, 1996]*
- Full Bayesian treatment was infeasible back then….
  - … and still is, generally, not industry-standard by any means.

UQ-for-NN: Bayesian perspective

Training for NN weights reformulated as a Bayesian inference problem

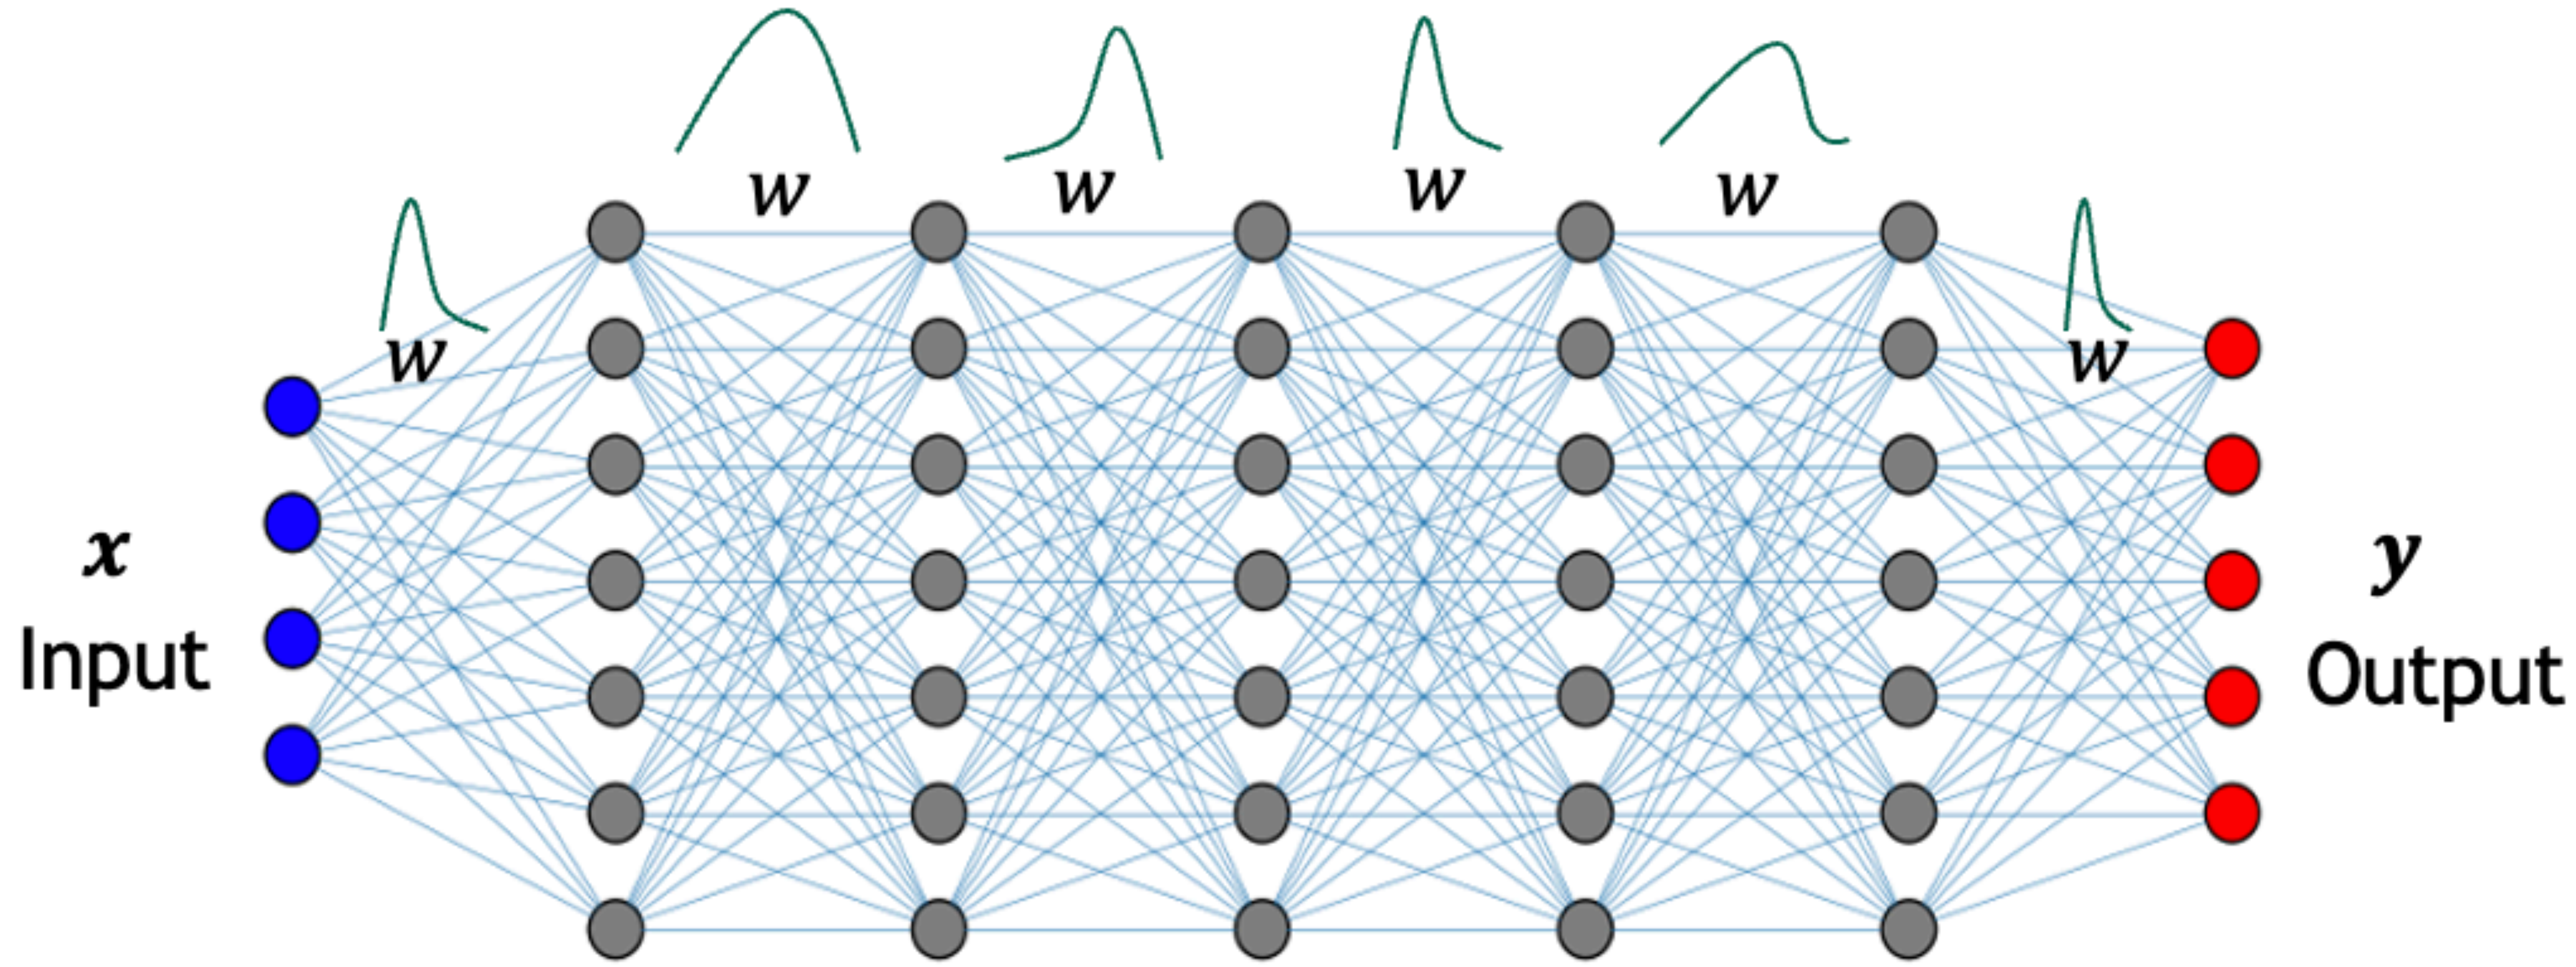

Posterior Prior

$$p(w\,|\,y) \propto p(y\,|\,w)\, p(w)$$

Likelihood

$$\propto \exp\left(-\frac{||y - f_w(x)||^2}{2\sigma^2}\right) \exp\left(-\frac{||w||^2}{2\lambda^2}\right)$$

Negative Log-Posterior $\simeq a\,||y - f_w(x)||^2 + b\,||w||^2 \simeq$ Training Loss Function

- ✓ Markov chain Monte Carlo (MCMC) sampling; Hamiltonian MC *[Levy, 2018]*
- ◉ Tuning is an art: essentially infeasible outside academic examples

UQ-for-NN: approximate methods

- **Variational inference**:
  - ✓ Bayes by Backprop *[Blundell, 2015]*
  - ✓ Probabilistic backprop *[Hernandez-Lobato 2015]*
  - ✓ SVI, BBVI, ADVI, ….
  - ◉ Typically underestimates predictive uncertainty
  - ◉ Restricted to variational class

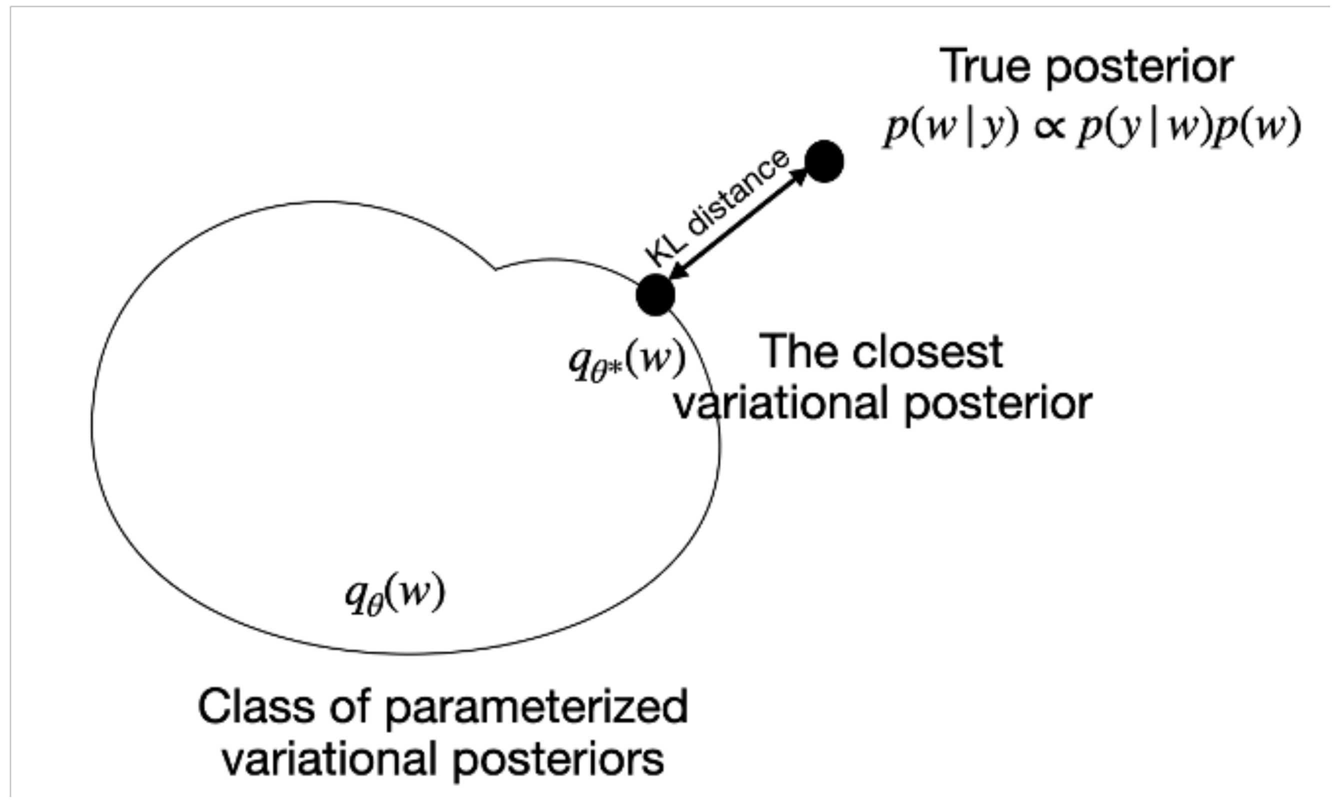

- **Laplace methods**: *[Daxberger, 2021]*
  - ✓ Relies on Gaussian apprx near maximum;
  - ✓ Can be generalized to GMM
  - ◉ Good only locally
  - ◉ Fails to explore the full posterior

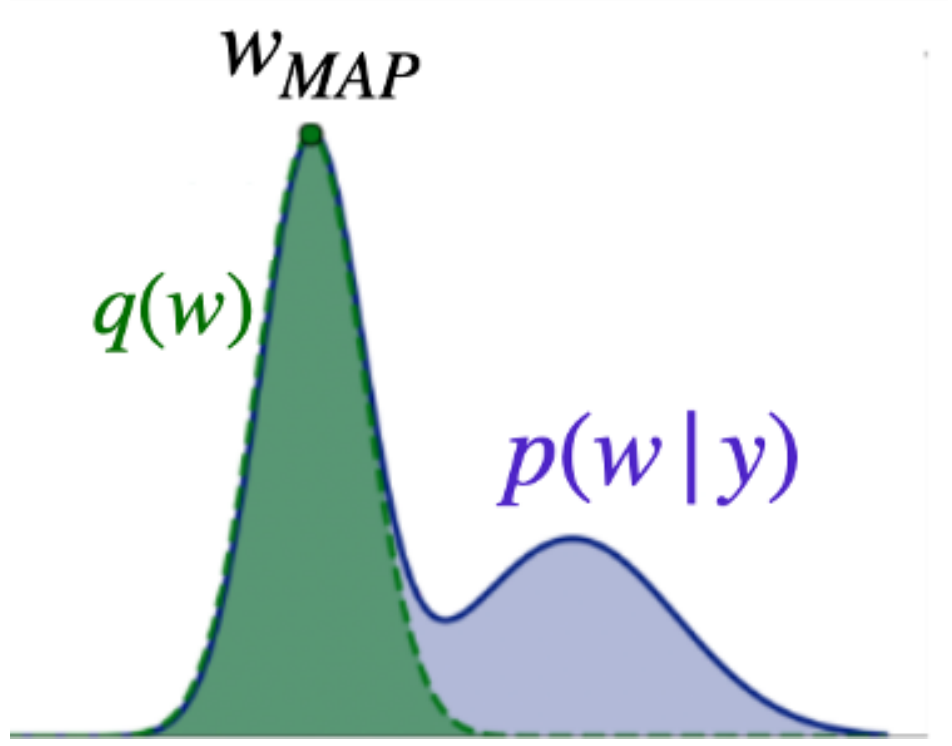

Accurate UQ for Neural Networks (NNs) hinges on the loss surface's behavior

UQ-for-NN: other (more empirical) methods

- **Ensembling methods:** work surprisingly well!
  - ✓ Deep Ensembles *[Lakshminarayanan, 2017]*
  - ✓ Randomized MAP Sampling *[Pearce, 2020]*
  - ✓ MC-Dropout *[Gal, 2015]*
  - ✓ Stochastic Weight Averaging – Gaussian (SWAG) *[Maddox, 2019]*
  - ◉ Little theoretical backing
  - ◉ Too expensive, albeit parallelizable
  - ◉ Some recent work interpreting these from Bayesian perspective

- **Direct learning of predictive RV**
  - ✓ Delta-UQ *[Anirudh, 2021]*,
  - ✓ Conformal UQ *[Hu, 2022]*,
  - ✓ Information-bottleneck UQ *[Guo, 2023]*,
  - ✓ Distance-based methods *[Postels, 2022]*.

Accurate UQ for Neural Networks (NNs) hinges on the loss surface's behavior

Challenges of UQ-for-NN

✓ Complicated posterior distribution (loss surface):
- invariances and symmetries: permuting some weights leads to the same loss,
- multimodality: multiple local minima in the weight space,
- “ridges”: low-d manifolds with same or similar loss.

✓ Prior on weights hard to elicit/interpret/defend
- what does a uniform/gaussian prior on weight matrix elements mean?
- perhaps a prior is needed in the ‘matrix’-space, or…
- driven by outputs, or physics-constraints.

✓ Large number of weights:
- scales linearly with depth and quadratically with width,
- hard to visualize the high-d surface.

Physics-driven regularization should help

- We hypothesize that incorporating prior knowledge of physics will regularize the loss/log-posterior landscapes, making them more amenable to sampling and analysis.

- This means both:
  - *soft* regularization (like PINN) and
  - *hard* architectural changes
    - physics-driven rewiring (invariance, symmetries, positivity, feature extraction),
    - numerical convenience (ResNet/NODE, weight reparameterization, layer/batch normalization).

- This regularization process should enable the derivation of well-calibrated, generalizable, and scalable predictive uncertainties.

ResNet example

ResNets regularize loss landscape compared to MLPs

Conventional MLP: $x_{n+1} = \sigma(W_n x_n + b_n)$

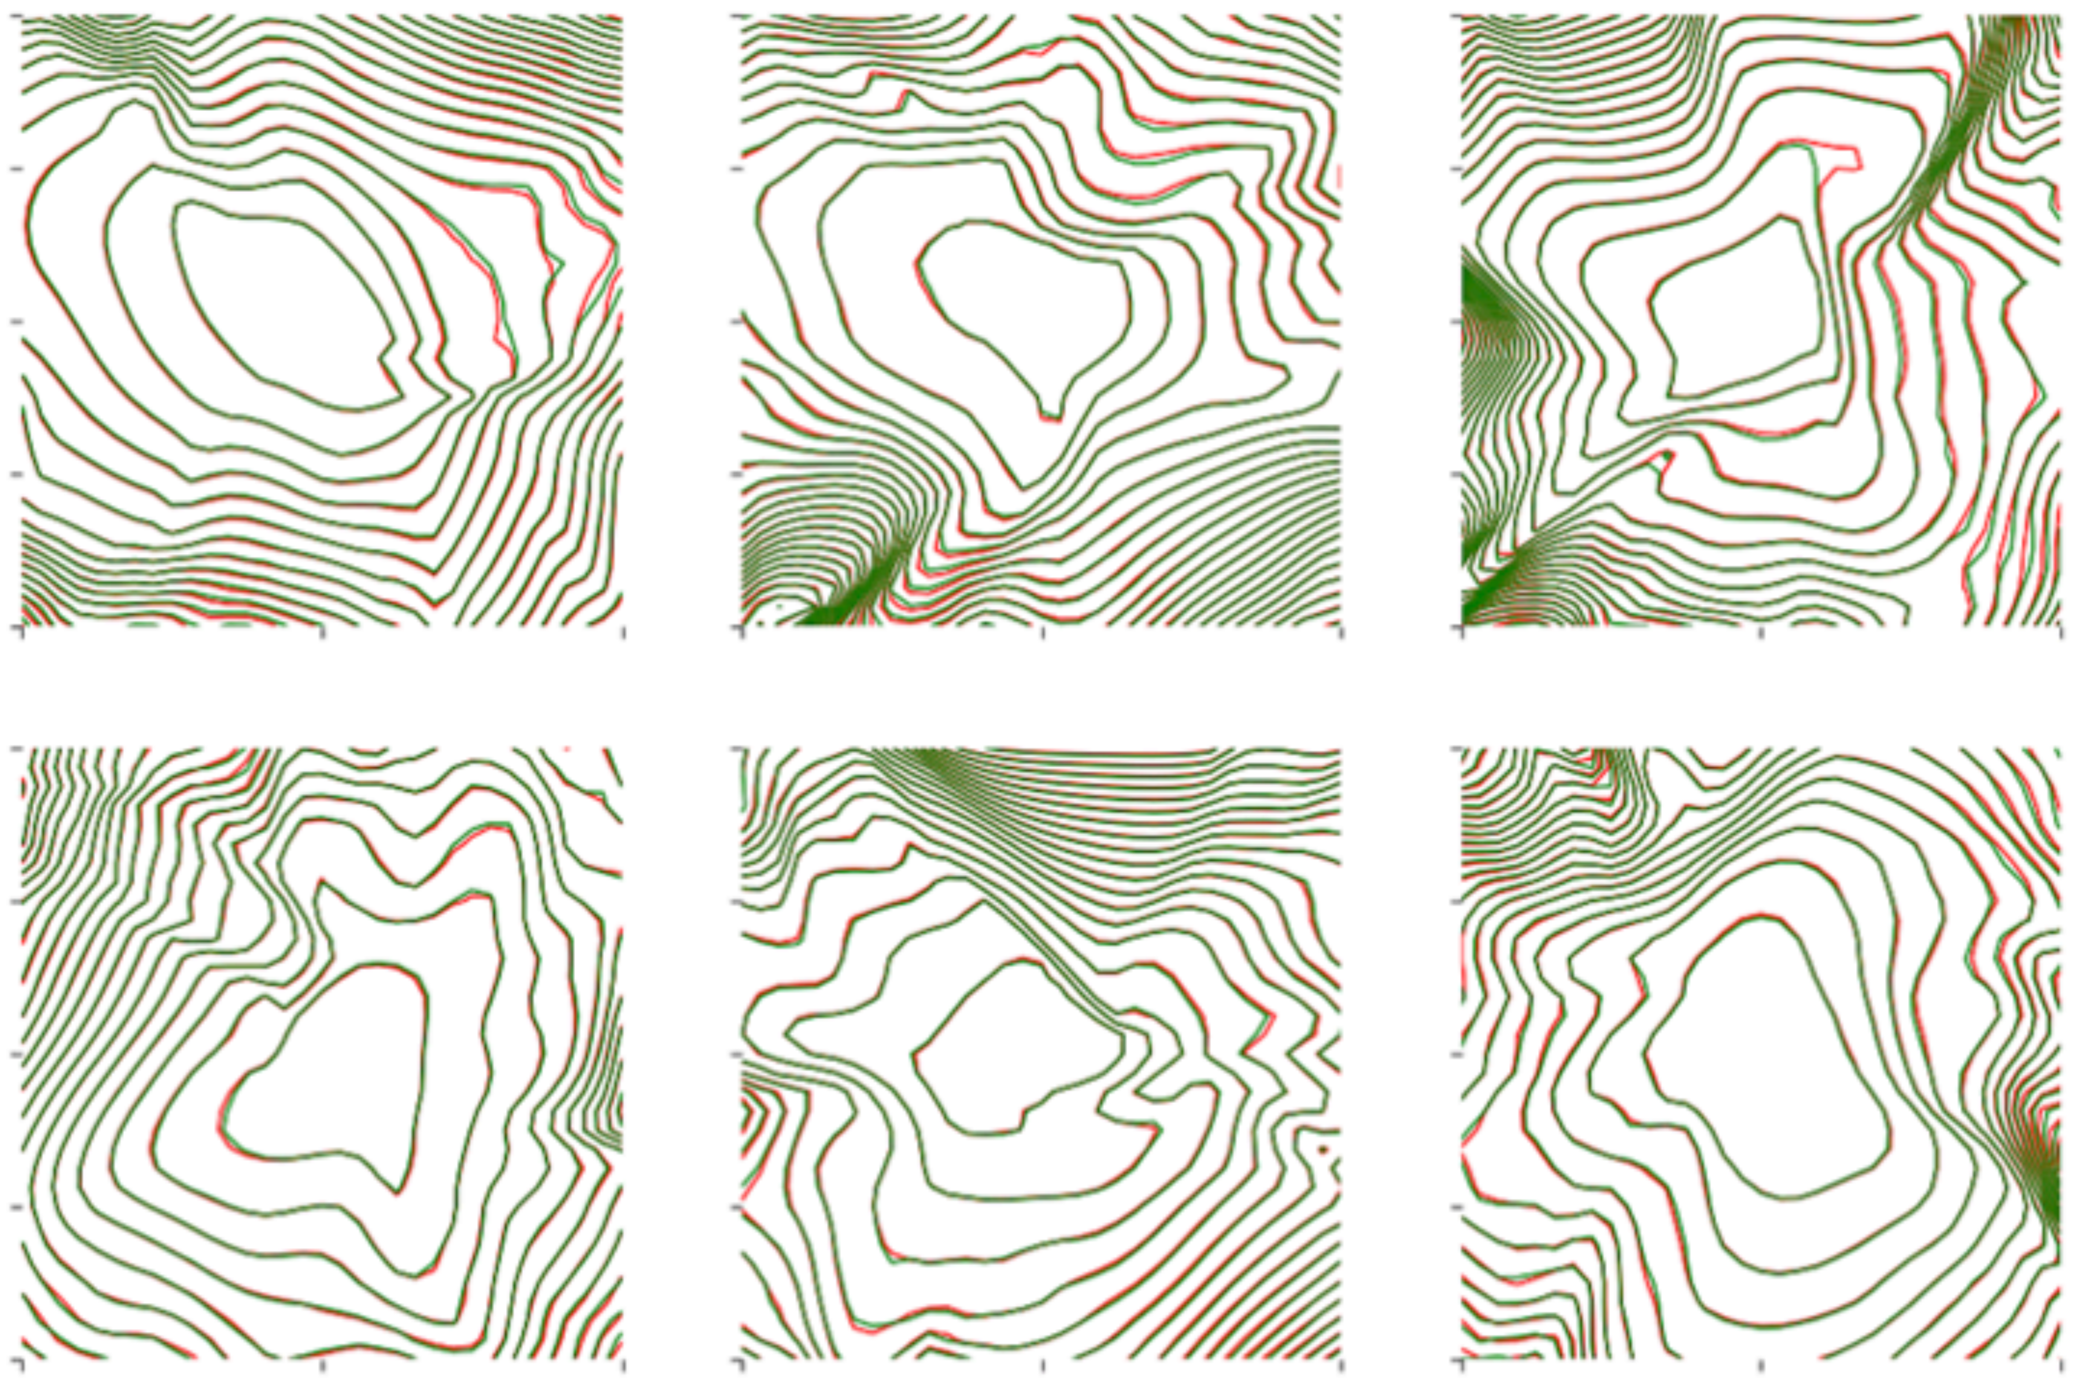

ResNet: $x_{n+1} = x_n + \sigma(W_n x_n + b_n)$

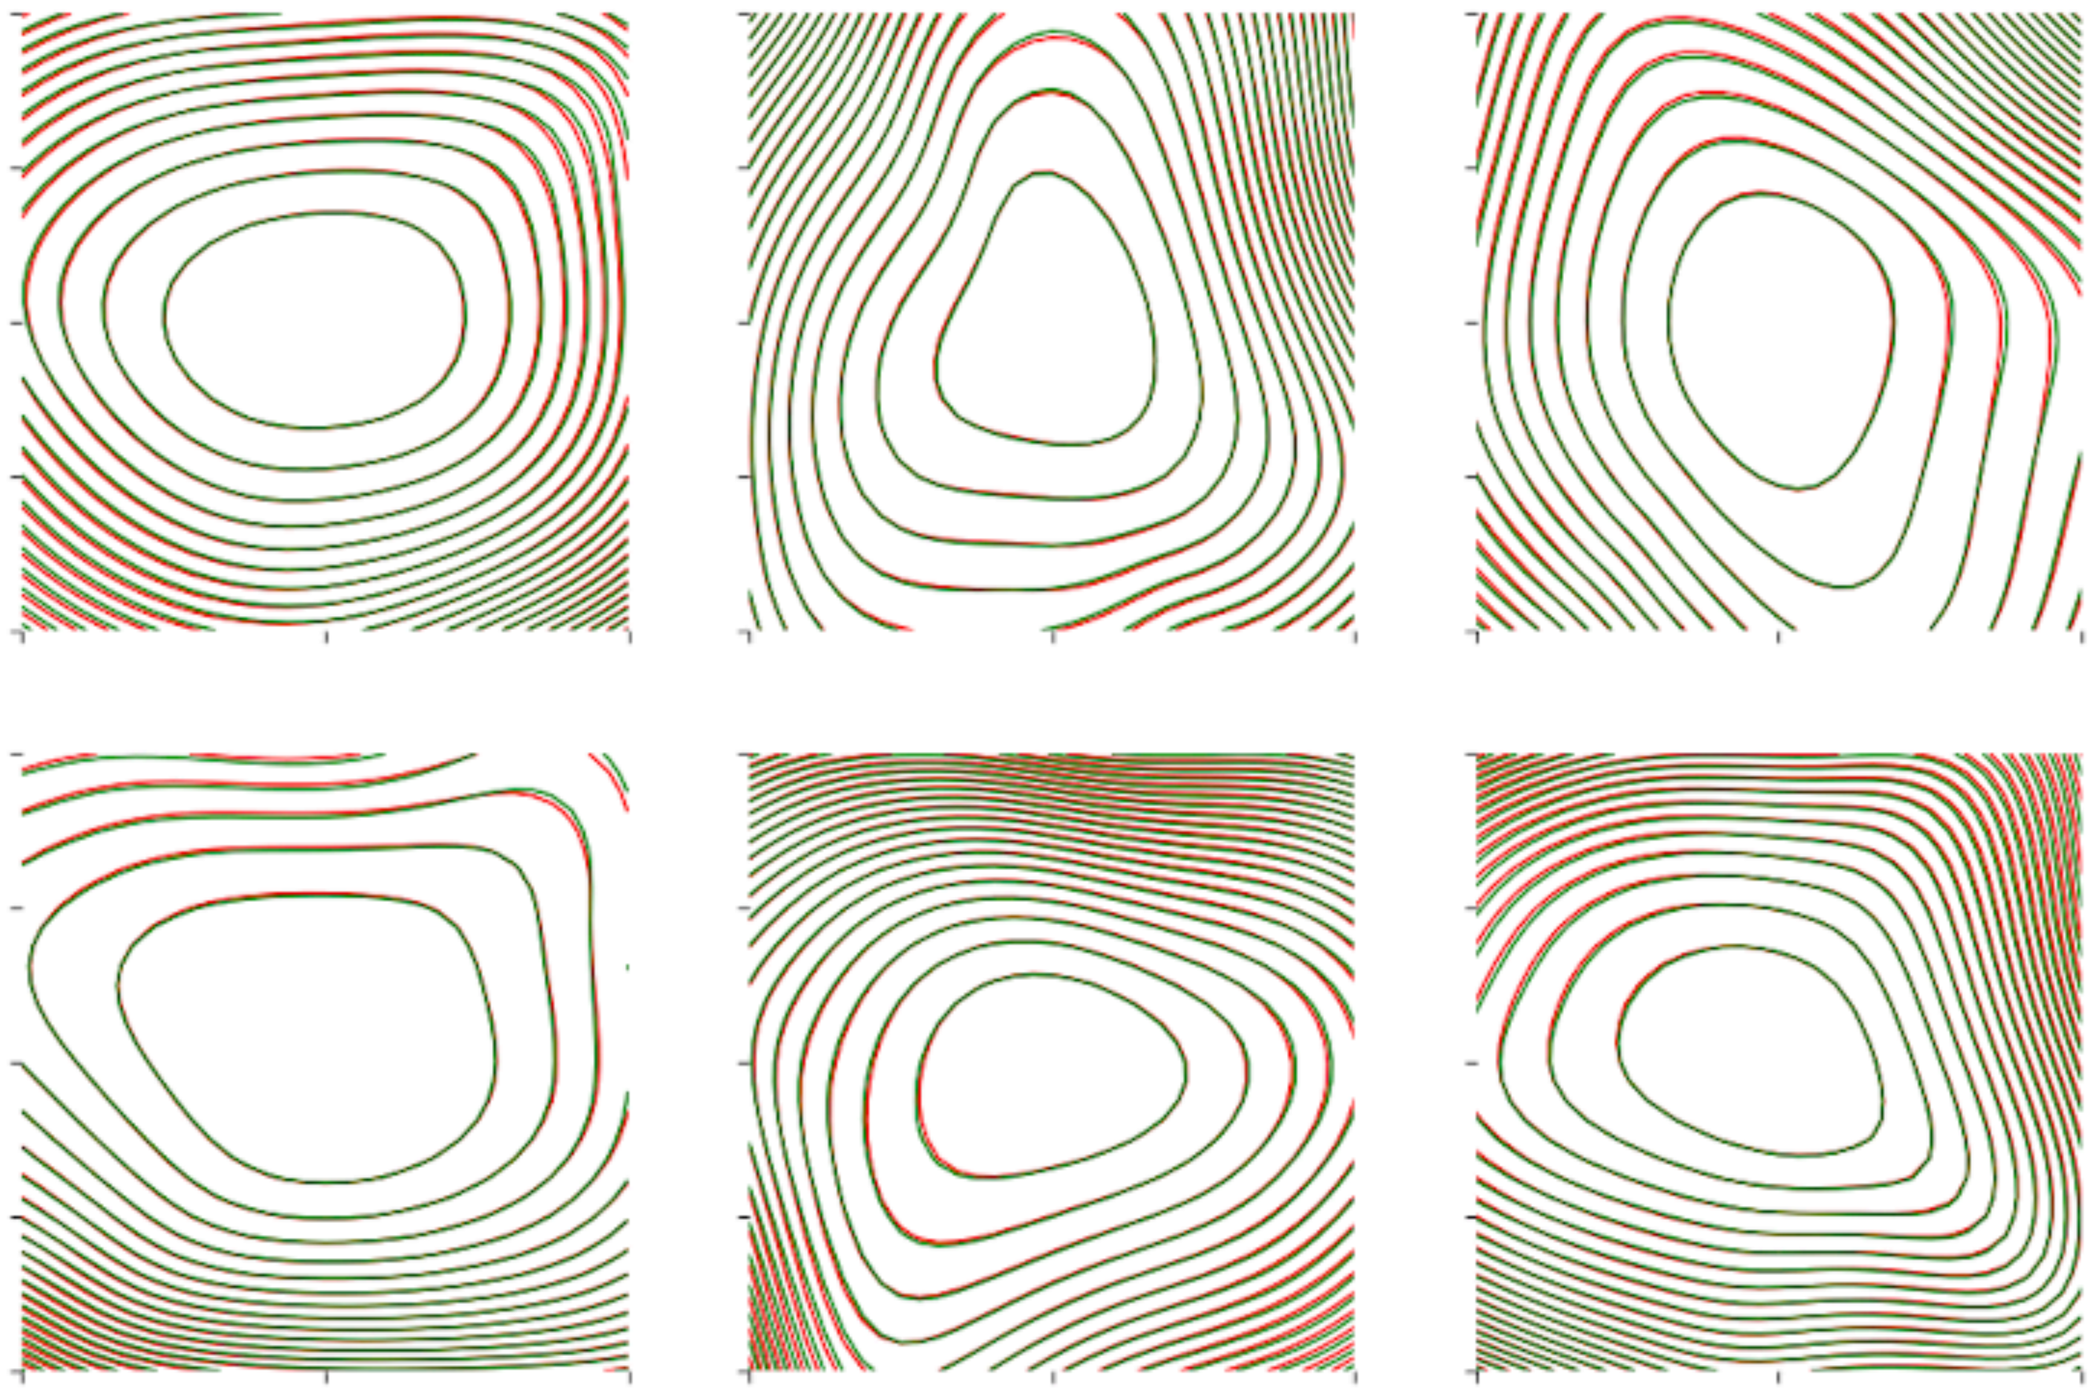

See *[Lee, 2017]* for a more comprehensive study.

Our Plan

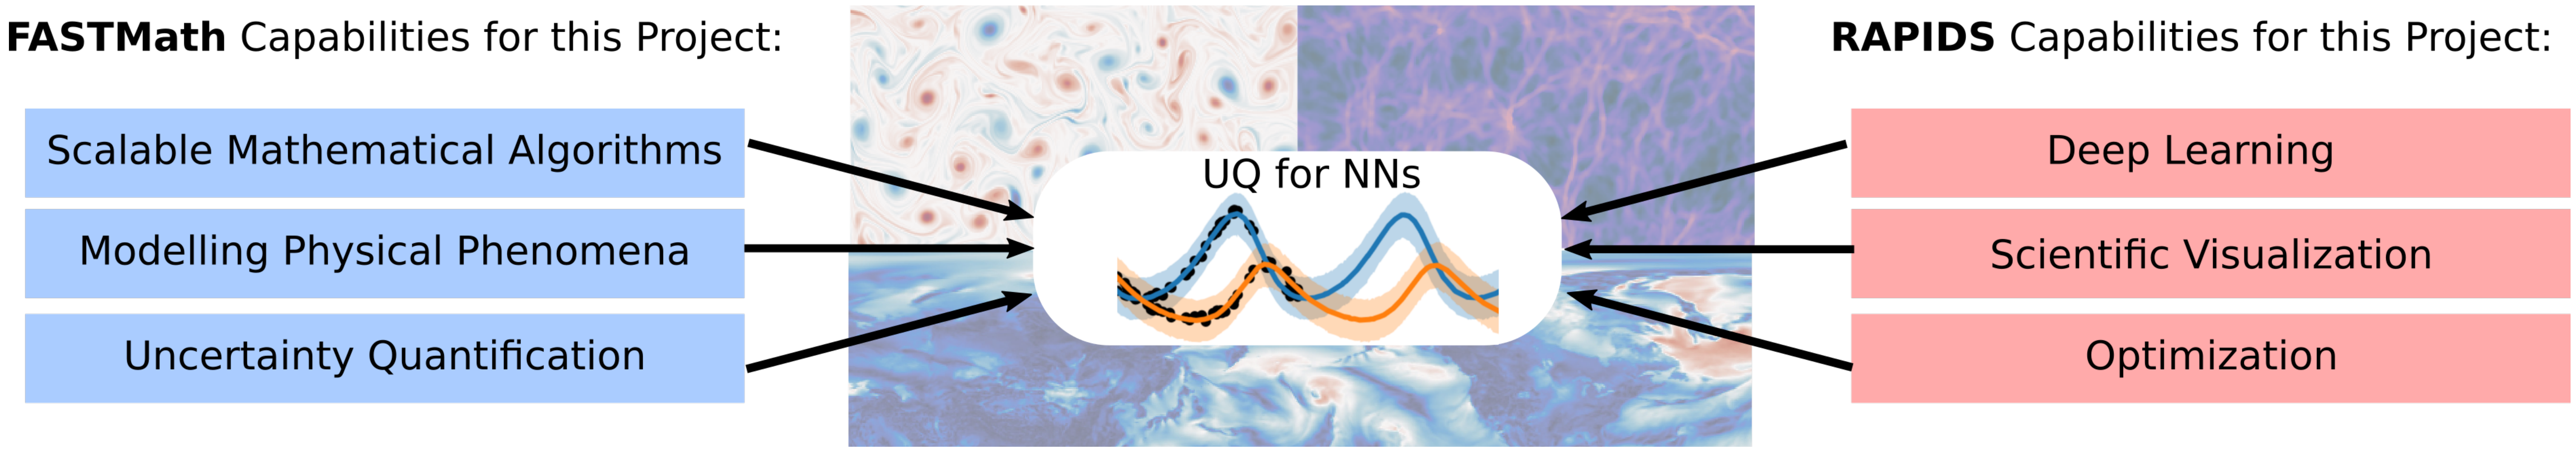

- Task 1: Scientific Visualization of Posterior Distributions of NNs. We aim to develop scientific visualization techniques to understand uncertainties in neural network predictions, and gain insights into the impact of physics-constraints on the shape of posterior distributions.

- Task 2: Mixture of Laplace Approximations for Quantifying Uncertainty in NNs. We aim to develop mixtures of Laplace approximations to model posterior distributions of varying shapes, for computing approximate uncertainty estimates for physics-aware NNs.

Our Plan: Visualization + (Physics) + Laplace

- Visualization of loss surface is key to help understand and characterize NN performance *[Li, 2018; Garipov, 2018; Fort, 2019; Yang, 2021]*,
- We will develop special slicing schemes, anchored at points of interest, such as local minima and saddle points found with conventional SGD methods,
- We will try to develop metrics of regularity, generalizability and "sample-ability" of the loss surface (a.k.a. log-posterior), incl. both local and global features.

- We will establish a systematic approach to categorize and interrogate the loss surface and measure the impact of physics-driven regularization on them,
- We will leverage the idea of Laplace approximation to obtain uncertainty estimates for NNs *[Daxberger, 2021; Graf, 2021; Ott, 2023]*,
- Motivated and informed by the loss surface analysis, we will develop scalable mixture-of-Laplace approximations to model posterior distributions of varying shapes.

Summary

- There is urgent need for principled UQ for NNs in DOE applications
- For moderately-sized NNs there is a lot to be done by loss surface analysis
- Utilize unique strengths of FASTMath and RAPIDS

- Quantifiable metrics for loss surface / log-posterior by utilizing anchored slices
- Measure the effect of hard and soft physics constraints
- Develop scalable UQ algorithms: mixture of Laplace

Impact: <u>methodological development</u> of novel UQ-for-NN approximate Bayesian algorithms, but also <u>practical impact</u> on a range of large-scale DOE applications. This should position us well for further SciML funding.

Logistics: given the short time, we plan to hire internally at each lab