Deep Bayesian Inference of Nearshore Bathymetry

Yizhou Qian,[1] Mojtaba Forghani,[1] Jonghyun Lee[2] Matthew Farthing,[3] Tyler Hesser,[3] Peter K. Kitanidis,[1] Eric F. Darve[1]
[1]Stanford University, [2]University of Hawaii at Manoa, [3]US Army Engineer Research and Development Center

Stanford University

Introduction

Nearshore bathymetry, the topography of the ocean floor in coastal zones, is vital for predicting the surf zone hydrodynamics and for route planning to avoid subsurface features. Hence, it is increasingly important for a wide variety of applications, including shipping operations, coastal management, and risk assessment. However, direct high resolution surveys of nearshore bathymetry are rarely performed due to budget constraints and logistical restrictions. Another option when only sparse observations are available is to use Gaussian Process regression (GPR), also called Kriging. But GPR has difficulties recognizing patterns with sharp gradients, like those found around sand bars and submerged objects, especially when observations are sparse. In this work, we present several deep learning-based techniques to estimate nearshore bathymetry with sparse, multi-scale measurements. We propose a Deep Neural Network (DNN) to compute posterior estimates of the nearshore bathymetry, as well as a conditional Generative Adversarial Network (cGAN) that samples from the posterior distribution. We train our neural networks based on synthetic data generated from nearshore surveys provided by the U.S. Army Corps of Engineer Field Research Facility (FRF) in Duck, North Carolina. We compare our methods with Kriging on real surveys as well as surveys with artificially added sharp gradients. Results show that direct estimation by DNN gives better predictions than Kriging in this application. We use bootstrapping with DNN for uncertainty quantification. We also propose a method, named DNN-Kriging, that combines deep learning with Kriging and shows further improvement of the posterior estimates.

Contribution

In this work, we propose the use of deep learning techniques within a Bayesian framework that provides uncertainty quantification. We apply it to the interpolation problem of predicting nearshore bathymetry, given sparse point-wise measurements and grid cell average measurements. Our main contributions are as follows:

- We propose two deep learning-based approaches that can learn highly non-linear and complex distributions automatically through its training data.
  - We trained a conditional Generative Adversarial Network (cGAN) to learn the posterior distribution of nearshore bathymetry. This allows us to sample directly from the posterior distribution and compute different posterior estimates including mean and standard deviation.
  - We trained a fully connected DNN to estimate the posterior mean. Uncertainty quantification is provided by combining our DNN model with bootstrapping in Kriging. This approach is more computationally feasible than cGAN in terms of both training time and optimization of hyper-parameters.
- We propose a method (DNN-Kriging) that uses Kriging to reduce the error in the DNN's prediction of the posterior mean. The motivation for this approach is that the DNN is capable of learning fine-scale features, while Kriging can accurately capture smooth components in the error. Here by "fine-scale" features we mean features that usually involve rapid oscillations or large values in the derivatives.

Application: Nearshore Bathymetry

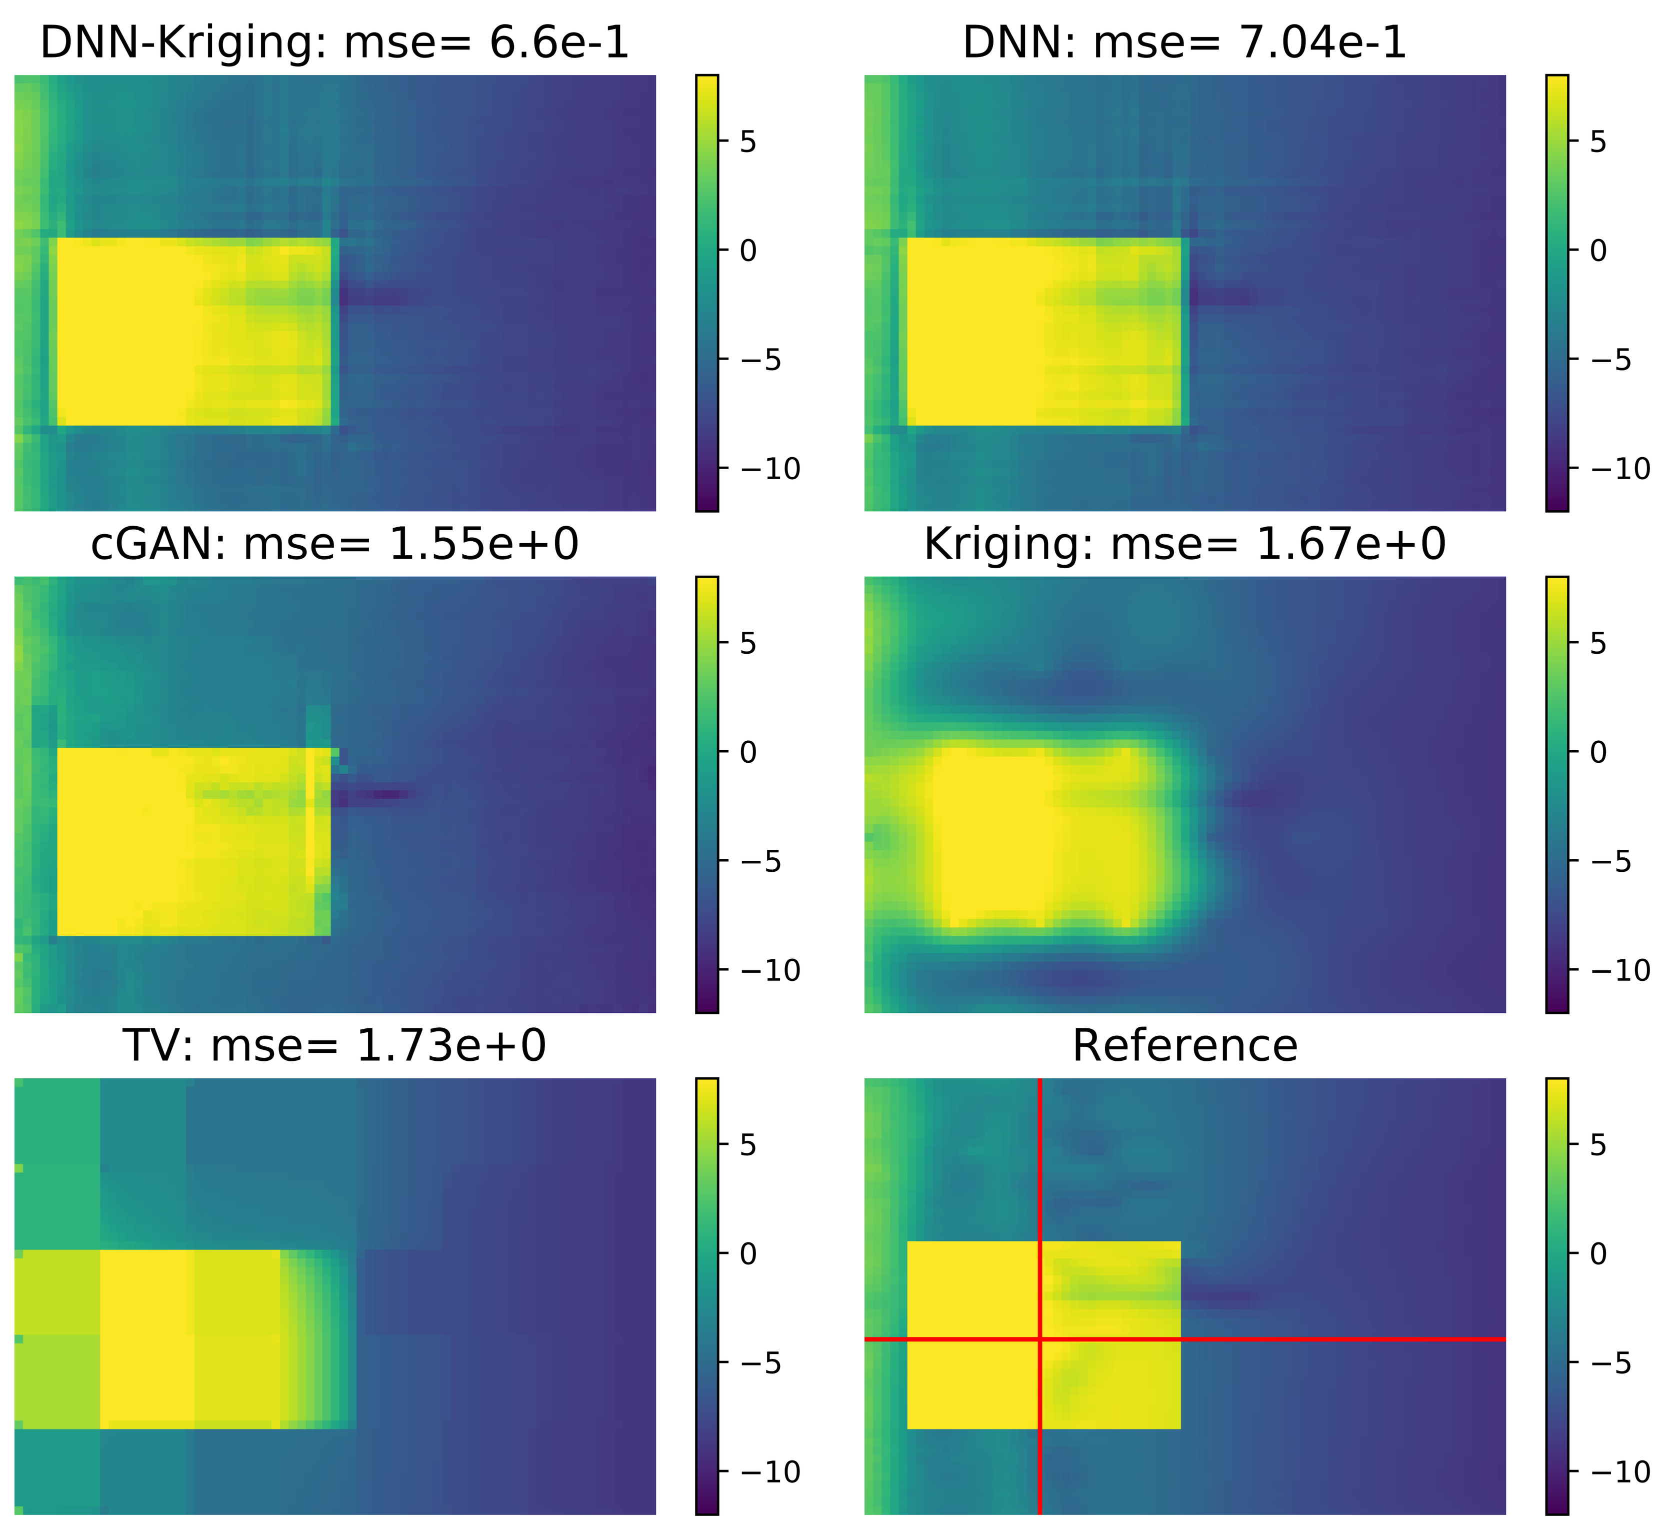

Fig. 1: Nearshore Bathymetry that has features with sharp gradients

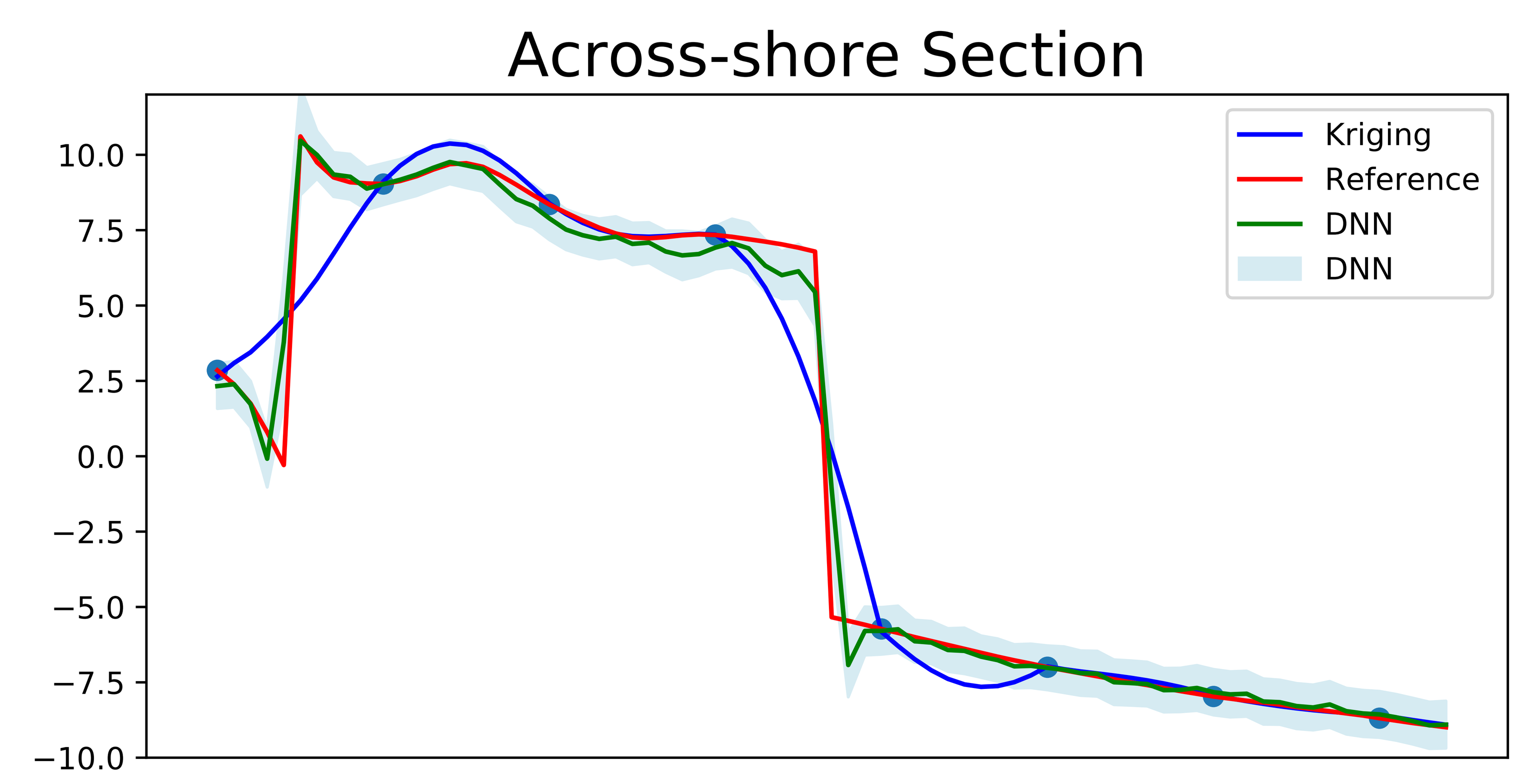

Fig. 2: Across-shore Section Comparisons

Uncertainty Quantification

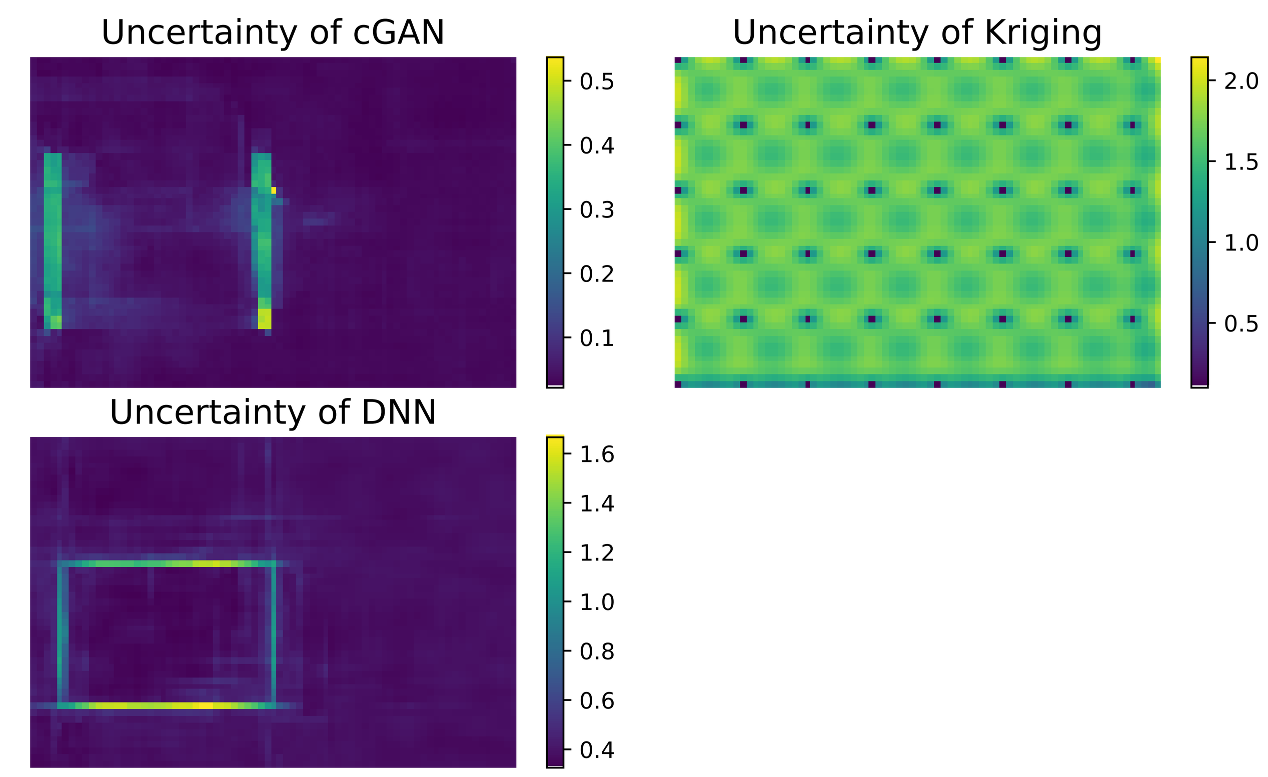

Fig. 3: Uncertainty quantification provided by cGAN, Krigin and DNN

Conclusion

In this work, we have explored the use of deep learning techniques from a Bayesian perspective to estimate nearshore bathymetry with sparse point-wise measurements and grid cell average measurements. We proposed using conditional Generative Adversarial Networks (cGAN) as a purely data-driven approach to directly sample the posterior distribution. This usually has a challenging training process and requires extra effort for tuning of hyper-parameters. We also proposed using a fully connected DNN to directly estimate the posterior mean of the bathymetry, which can be combined with bootstrapping from Kriging to provide uncertainty quantification. Both approaches provide more accurate predictions than Kriging when sharp changes are present in nearshore surveys. Finally, we proposed a method named DNN-Kriging that combines Kriging's ability to model smooth variations in the residuals with DNN's ability to capture fine-scale features. Results show that DNN-Kriging provides the best estimate among all the methods.

Acknowledgements

This research was supported by the U.S. Department of Energy, Office of Advanced Scientific Computing Research under the Collaboratory on Mathematics and Physics-Informed Learning Machines for Multiscale and Multiphysics Problems (PhILMs) project, PhILMS grant DE-SC0019453.