Data Science Seminar Hosted by Department of Mathematical Sciences

Interdisciplinary Dean's Speaker Series in Data Science

Reception at CW-112 at 4:30 pm with refreshment available, allowing the audience and the speaker to mingle after the talk.

RSVP at http://bit.ly/DS-TAE-RSVP.

Date: Friday, November 8, 2019

■ Time: 3:30pm - 4:30pm

■ Room: LH-10

 Speaker: Andrew Gordon Wilson (New York University; Courant Institute of Mathematical Sciences and Center for Data Science)

• Title: How do we build models that learn and generalize?

Abstract

To answer scientific questions, and reason about data, we must build models and perform inference within those models. But how should we approach model construction and inference to make the most successful predictions? How do we represent uncertainty and prior knowledge? How flexible should our models be? Should we use a single model, or multiple different models? Should we follow a different procedure depending on how much data are available?

In this talk I will present a philosophy for model construction, grounded in probability theory. I will exemplify this approach for scalable kernel learning and Gaussian processes, Bayesian deep learning, and understanding human learning.

Bio: Andrew Gordon Wilson is faculty in the Courant Institute and Center for Data Science at NYU. Before joining NYU, he was an assistant professor at Cornell University from 2016-2019. He was a research fellow in the Machine Learning Department at Carnegie Mellon University from 2014-2016, and completed his PhD at the University of Cambridge in 2014. Andrew's interests include probabilistic modelling, scientific computing, Gaussian processes, Bayesian statistics, and loss surfaces and generalization in deep learning. His webpage is https://cims.nyu.edu/~andrewgw.

The Interdisciplinary Dean's Speaker Series in Data Sciences is supported by the:

- Dean's Office of Harpur College of Arts and Sciences
- Department of Biological Sciences
- Department of Mathematical Sciences
- Department of Political Science
- Department of Systems Science and Industrial Engineering
- Data Science Transdisciplinary Area of Excellence

For questions, contact Ken Kurtz (kkurtz@binghamton.edu) or Xingye Qiao (qiao@math.binghamton.edu). Contact Ken Kurtz to request meeting time with the speaker.

References:

- Garipov, T., Izmailov, P., Podoprikhin, D., Vetrov, D.P. and Wilson, A.G., 2018. Loss surfaces, mode connectivity, and fast ensembling of dnns. In Advances in Neural Information Processing Systems (pp. 8789-8798).
- Izmailov, P., Podoprikhin, D., Garipov, T., Vetrov, D. and Wilson, A.G., 2018. Averaging weights leads to wider optima and better generalization. arXiv preprint arXiv:1803.05407. (UAI 2018)
- Izmailov, P., Maddox, W.J., Kirichenko, P., Garipov, T., Vetrov, D. and Wilson, A.G., 2019. Subspace Inference for Bayesian Deep Learning. arXiv preprint arXiv:1907.07504. (UAI 2019)
- Gardner, J., Pleiss, G., Weinberger, K.Q., Bindel, D. and Wilson, A.G., 2018. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In Advances in Neural Information Processing Systems (pp. 7576-7586).

From:

http://www2.math.binghamton.edu/ - **Department of Mathematics and Statistics, Binghamton University**

Permanent link:

http://www2.math.binghamton.edu/p/seminars/datasci/191108

Last update: 2019/10/18 15:59