Cambridge NeurIPS Meet-Up

Imperial College London

OccaMLab at NeurIPS

GPs, Equivariance, Causality

Mark van der Wilk

Department of Computing
Imperial College London
`https://mvdw.uk`

@markvanderwilk
`m.vdwilk@imperial.ac.uk`

Dec 8, 2022

A Group Effort

All this work done by brilliant PhD candidates, collaborators, together with me.

Anish Dhir

Artem Artemev

Jose Pablo Folch

Ruby Sedgwick

Seth Nabarro

Tycho van der Ouderaa

Baselines and Benchmarking GP Approximations

Recommendations for Baselines and Benchmarking Approximate Gaussian Processes (Workshop on GPs)
Sebastian W. Ober, David R. Burt, Artem Artemev, MvdW

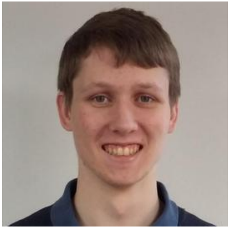  

Baselines and Benchmarking GP Approximations

Literature is very unclear!

- Many approximations, inconsistent parameter selection.
- Run on all UCI datasets, without understanding them.
- Big table, new always outperforms old on "most datasets".

⏩ Methods are so good, that benchmarking is holding progress back!

Take home 1: Benchmarking should include automatic tuning of approximation params.

Take home 2: We should compare on:

- Compute needed to obtain ϵ-optimal performance
  ⏩ Good approximations should have **no** performance difference!
- Performance after various compute limits

SGPR as a Baseline

Take home 3: If you run SGPR in the right way, it's near-exact.

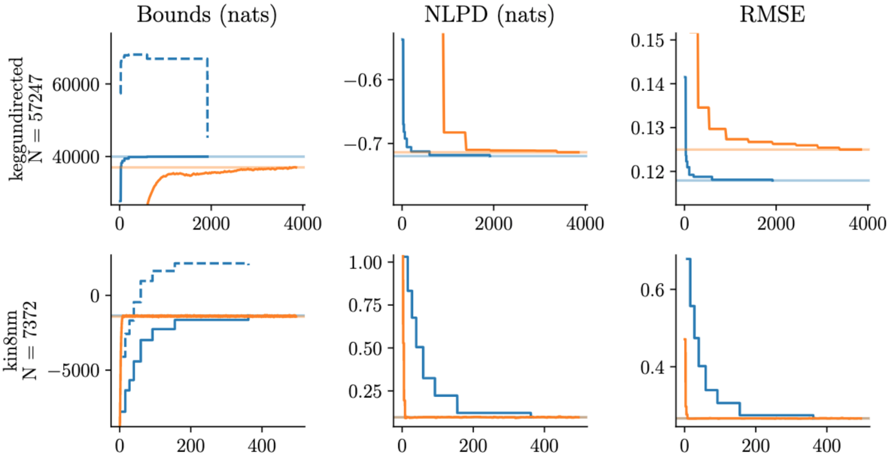

SGPR as a Baseline

Take home 4: Approximations and modelling are related. We need to understand our datasets to make progress.

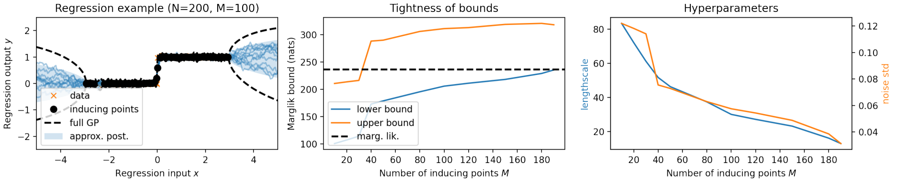

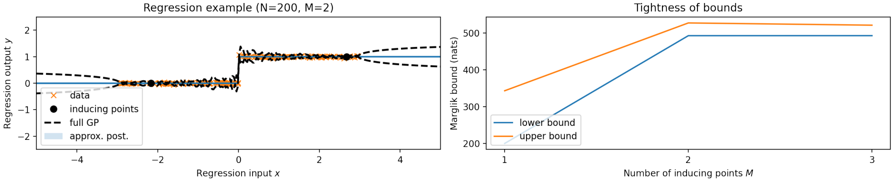

Removing Memory Constraints

Memory Safe Computations with XLA Compiler (NeurIPS conf)
Artem Artemev · Yuze An · Tilman Roeder · MvdW

- Methods are often **memory limited** (GPs, Transformers)
- Memory efficient methods exist, but are hard to implement
- We make an extension for XLA that automatically finds memory bottlenecks in the computational graph, and removes them.
- Remove Out of Memory errors with 0 changes to code.

Learning Equivariance in Neural Networks

Invariance Learning in Deep Neural Networks with Differentiable Laplace Approximations (NeurIPS conf)
Alexander Immer · Tycho van der Ouderaa · Gunnar Rätsch · Vincent Fortuin · MvdW

- Learn invariance parameters (expressed like data augmentation) through backprop
- Laplace approximation to the marginal likelihood
- Bayesian Model Selection for interesting inductive biases seems viable for DNNs!

Learning Equivariance in Neural Networks

Relaxing Equivariance Constraints with Non-stationary Continuous Filters (NeurIPS conf)
Tycho van der Ouderaa · David W. Romero · MvdW

Sparse Convolutions on Lie Groups (NeurReps workshop)
Tycho van der Ouderaa · MvdW

- Parameterise a layer that can interpolate between fully-connected, and various group equivariance
- Towards layer-by-layer inductive bias learning!

Bayesian Model Selection for Causality

Causal Discovery using Marginal Likelihood (CML4Impact workshop)
Anish Dhir · MvdW

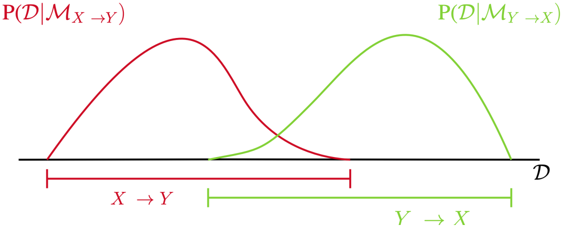

Additive Noise Model

Common assumption:

Join us!

Anish Dhir

Artem Artemev

Jose Pablo Folch

Ruby Sedgwick

Seth Nabarro

Tycho van der Ouderaa

- Good time to be applying through StatML CDT (`https://statml.io/`).
- Check my website (`https://mvdw.uk/`) for tips on applying, and how to get in touch.
- And do find me to chat if you have questions.