



Computer
Science

CSC696H: Probabilistic Methods in Machine Learning

Term Projects Information

Prof. Jason Pacheco

Outline

- Project Proposal
- Project Presentation
- Project Report

Motivation

Why do we do seminar courses? Projects?

Main Goal Kickstart a project that can lead to a publishable piece of work and potential contribution to the field.

Other Goals

- Opportunity to explore topic in-depth and do some research in it
- Incentive to look at other avenues of your research that have not been prioritized—e.g. ideas you have had but haven't looked at
- For some this may be a first taste of research in an area
- Relatively low-risk setting, negative results are acceptable if you did good research, you can be ambitious

Project Proposal

Assignment Details

- Due : **Friday, Mar. 8**
- Upload to Github repo as PDF

Format (2-3 pages + references)

- Project Summary : What are you doing and why?
- Previous Work : What has been done in the literature?
- Approach : How do you plan to do it?
- Evaluation Methodology : How will you evaluate it?
- Include any figures that are helpful in describing the project, approach, or evaluation.

Resources

Collection of proceedings and other electronic resources...

- [Advances in Neural Information Processing Systems](#) (NeurIPS)
- [Proceedings of Machine Learning Research](#) (PMLR)
 - Artificial Intelligence and Statistics (AISTATS)
 - International Conference on Machine Learning (ICML)
 - Journal of Machine Learning Research (JMLR)
 - Uncertainty in Artificial Intelligence (UAI)
 - Conference on Learning Theory (COLT)
 - Knowledge Discovery and Data Mining (KDD)
- [International Conference on Learning Representations](#) (ICLR)
- Also look at recent workshops at these conferences

Model Development

Propose a new model for a problem of your choice and work out the inference / estimation

Typical Approach:

- Explain why this model is needed—should improve on existing work
- Define model and all distributions
- What data you will use? synthetic or real data?
- What inference algorithm will you use? Typically MCMC sampling
- If possible start by simulating data and visualizing it
- Fit model, do inference, show results compared to some baseline
- LDA paper is a good example of the whole process

Algorithm Development

Extend existing inference / estimation algorithm

Typical Approach:

- Explain how this should improve existing methodology
- Look through existing methods in Variational inference, MCMC, implicit likelihood inference (ABC), Bayesian Deep Learning, Bayesian Optimization
- Define algorithm
- Determine how you will evaluate results. What is your baseline?
- Implement, debug, fix...repeat
- Typical comparison : predictive accuracy, held out log-marginal likelihood, other domain-specific metrics

Data Analysis

Use / extend existing models / algorithms to evaluate a new dataset

Typical Approach:

- Some combination of model and algorithm development
- What existing methods are there to analyze this data? Why is this better?
- Are you going to collect new data? If so, how? What if you can't get the data or it takes too long?
- Collect data, fit model, visualize and compare to existing methods

Theory

Prove results that have not yet been established

Typical Approach:

- Identify existing theoretical gap in the literature
- May be proving tighter bounds, better algorithm guarantees, or a new framework altogether
- Identify the sequence of steps, intermediate results (i.e. lemmas), main results (theorems)
- What if you cannot prove (or disprove) your main results? What is your fallback?
- Some empirical validation should show how closely (or not closely) theory holds in practice

Try to Avoid These Projects

I generally discourage these, but if you have good justification then I may allow it...

- Implementing existing algorithm without extension
 - Only makes sense if no code is available
 - Should be some benefit to having an implementation—is anybody asking for an implementation?
- Topic survey
 - Project report will need to be much more extensive
 - Needs to be a recent research area that lacks existing survey
 - Better suited for a 500-level class

Project Presentations

3 days to present 10 projects

The math...

- We can fit all 10 presentations in 3 days with 20min per
- Lecture is 70min (3 lectures @ 20min = 60min with 15min to spare)
- Can do 4 one day with no room

Project Reports

- Due Fri 5/3 @ 11:59pm
- Upload all your code and PDF report to Github
- I recommend (but do not require) using NeurIPS format here:
<https://neurips.cc/Conferences/2023/PaperInformation/StyleFiles>
- If you work in another discipline you may use a format for a top venue in your field (i.e. CVPR, IEEE, COLT, etc.)
- Maximum 6 pages

Outline should follow some variation of the following:

1. **Introduction** Give an overview of the problem, what you will show, contributions
2. **Background** Give all of the technical foundations necessary to understand the paper
3. **Approach** All of the technical details of your approach
4. **Related Work** How does this work relate to existing publications
5. **Experimental Results** All of your validation