University at Buffalo
The State University of New York

Course Title: Advanced Topics in Scalable Bayesian Methods
Semester: 2018 Fall

COURSE INFORMATION

Course #: CSE 706-23458

The course is for 1-3 credit and will meet on Mondays at 11:00AM – 1:00 PM @ Alumni 90.

INSTRUCTOR

Dr. Changyou Chen, Davis 338L, changyou@buffalo.edu
Office hours: 3:00 PM – 5:00 PM (Mondays)

COURSE DESCRIPTION

Bayesian methods form an important topic in modern machine learning, especially in deep learning. Compared to optimization, Bayesian methods deal with data from a probability point of view, thus are more robust to noisy data, and are able to model some unique qualities such as prediction uncertainty. However, when facing with big data, traditional Bayesian methods such as the Metropolis–Hastings algorithm and variational Bayesian are not scalable enough in real applications. The development of scalable Bayesian methods thus has become an important and urgent task in modern machine learning.

In this seminar, we will review recent developments on this topic, mainly covering the following two lines of research: scalable Bayesian sampling and scalable variational-inference-based methods. Algorithms, theory, and applications in deep learning will be covered.

The following topics will be covered:

- Concepts and algorithms of Bayesian methods
- Bayesian modeling in feedforward deep neural networks (uncertainty modeling)
- Bayesian modeling in deep generated models
- Bayesian modeling in deep reinforcement learning

The students are expected to have basic knowledge of machine learning and deep learning.

SYLLABUS

Date	Topics and Papers
Week 1 (Aug 27)	Introduction
Week 2 (Sep 3)	No class
Week 3 (Sep 10)	Basics on sampling and variational inference
Week 4 (Sep 17)	Basics on sampling and variational inference
Week 5 (Sep 24)	Bayesian for uncertainty modeling: • SG-MCMC for DNN
Week 6 (Oct 1)	Bayesian for uncertainty modeling:

	• SVGD for DNN
Week 7 (Oct 8)	Bayesian for uncertainty modeling: • Variational Bayes based methods
Week 8 (Oct 15)	Bayesian for uncertainty modeling: • Dropout, batch normalization
Week 9 (Oct 22)	Bayesian for deep generative models: • VAE
Week 10 (Oct 29)	Bayesian for deep generative models: • GAN
Week 11 (Nov 5)	Bayesian for deep generative models: • VAE & GAN
Week 12 (Nov 12)	Bayesian for reinforcement learning: • Policy gradient
Week 13 (Nov 19)	Bayesian for reinforcement learning: • Soft-Q learning and related
Week 14 (Nov 26)	Bayesian for reinforcement learning: • Exploration
Week 15 (Dec 3)	• Others & project presentation

COURSE PURPOSE

Exposes students to recently developed methods on Bayesian methods for deep learning.

COURSE GOALS

The course will be in the form of seminar, where instructor and students take turns to present recent papers on these topics, and take questions by other students.

The goals of the course are:

- To introduce students concepts and main methods/algorithms of Bayesian methods for deep learning.
- To enable students to be good readers, critical thinkers and effective communicators on research ideas.
- To enable students to gained a horizon on an up-to-date Bayesian methods for deep learning, and equipped knowledge for their future research.

REFERENCE

- I. Goodfellow, Y. Bengio and A. Courville, Deep Learning, MIT Press, 2016.
- R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, MIT Press, 2017.
- Recent papers at ICML, NIPS, ICLR, and etc.

GRADING POLICY

- Only pass and fail.

ACADEMIC INTEGRITY

Academic dishonesty will not be tolerated. Please see the CSE department policy below on the topic; this policy specifies penalties for violations.

What is academic dishonesty? To hand in any work which is not 100% the student's creation, unless you are explicitly allowed to do so.

Reference to the university academic integrity (Graduate): http://grad.buffalo.edu/Academics/Policies-Procedures/Academic-Integrity.html.

ACCESSIBILITY RESOURCES

If you have any disability which requires reasonable accommodations to enable you to participate in this course, please contact the Office of Accessibility Resources, 25 Capen Hall, 645-2608, and also the instructor of this course. The office will provide you with information and review appropriate arrangements for reasonable accommodations. http://www.student-affairs.buffalo.edu/ods/