Course Syllabus: COMPSCI 791BDL
Bayesian Deep Learning – Spring 2022

Instructor: Benjamin M. Marlin
Email: marlin@cs.umass.edu

First Class: Jan 28, 2021
Time: Fri 2:30-5:15
Location: LGRC 143

Course Description: This seminar will introduce students to research in the area of Bayesian methods applied to deep neural network models. The course will begin with foundational readings on Markov chain Monte Carlo and variational Bayesian methods and proceed to cover recent advances that are enabling the application of Bayesian inference to increasingly large deep learning models. The course will also cover methods for accelerating prediction using Bayesian deep learning models and for evaluating Bayesian deep learning models. Students will need background in deep learning (such as provided by COMPSCI 682 or COMPSCI 689) and probabilistic graphical models (such as provided by COMPSCI 688). The seminar will focus on reading, presenting, and discussing classical and recent papers (1 credit) and a final project focusing on a Bayesian deep learning topic (3 credits). 1-3 credits.

Course Goals: The primary goal for this course is for students to learn about the state-of-the-art methods for approximate Bayesian inference that are driving current research in Bayesian deep learning (1 credit track), and to gain practical experience working with these methods to solve real-world problems (3-credit track).

Learning Outcomes: The primary learning outcomes in this course are as follows:

1. To gain familiarity with classical Markov chain Monte Carlo and Variational Bayesian inference methods and the challenges with applying these methods to large-scale models.
2. To gain familiarity with recent advances in scaling approximate Bayesian inference methods to deep learning models.
3. To gain familiarity with criteria for evaluating Bayesian deep learning models.
4. To gain practical experience implementing and evaluating Bayesian deep learning systems (3-credit project-based option only).

Course Modality: This is a presentation and discussion-based seminar. The course will meet in person. In the event of disruptions due to the on-going COVID-19 pandemic, the course will meet synchronously via Zoom.

Course Materials and Technologies:

- **Readings:** Readings will be from the primary literature. We will supplement with readings from *Machine Learning: A Probabilistic Perspective* by Kevin Murphy (free for UMass students via the UMass Library) for general background.
- **Course Website:** The course website is hosted on UMass' Moodle cloud course management portal https://umass.moonami.com/. Log in to the portal using your campus (Spire) ID and password. The course website will host the reading list and presentation schedule.
- **Course Communications:** Official announcements for the course will go out through Piazza. Piazza will also host discussion forums for the class. Students will be added to Piazza at the start of the semester. Students are expected to check their UMass email for course announcements at least once a day. The instructor will aim to respond to piazza posts within one day.

Coursework and Grading Plan: The course includes two tracks. In the 1-credit track, coursework will consist of reading papers, writing paper responses, and presenting papers. The 3-credit track includes the completion of a substantial course project on Bayesian deep learning. The details are provided below including the weight of each course component and the grade thresholds for the course.

- **1-Credit Track:** The course will cover one or two papers per week. Students will complete all readings and write a response for one research paper each week. Responses will be 250-500 words. Each student will participate in presenting one or two papers depending on final enrollment. Students will prepare an original presentations for the papers they present. All students are expected to participate in paper discussions.
- **3-Credit Track:** Students will complete the work in the 1-credit track in addition to a substantial course project. The deliverables for the course project include a 2-page project proposal due by mid-semester, a final project presentation with demo (10-15 minutes), and a final project report in standard 8-page NeurIPS format (plus additional pages for references). Course projects will be completed individually.

Coursework Item	1-Credit Track	3-Credit Track
Class Participation	20%	10%
Reading Responses	50%	20%
Paper Presentation	30%	10%
Project Proposal	-	10%
Project Presentation	-	20%
Project Final Report	-	30%

Highest	Lowest	Letter
100.00 %	93.00 %	A
92.99 %	90.00 %	A-
89.99 %	87.00 %	B+
86.99 %	83.00 %	B
82.99 %	80.00 %	B-
79.99 %	75.00 %	C+
74.99 %	70.00 %	C
69.99 %	0.00 %	F

Course Policies: Students should make sure they are familiar with all course policies and the relevant University policies linked to below. By staying enrolled in this course, students agree to be bound by all applicable policies.

- **Course Community Code of Conduct:** The instructor is committed to providing a friendly, safe and welcoming environment for all, regardless of gender identity and expression, sexual orientation, disability, personal appearance, body size, race, ethnicity, age, religion, nationality, or other similar characteristic. Please be courteous, respectful, and professional in all of your interactions with other students in all mediums of communication including but not limited to in-person, email, video meetings, chat, and discussion forums.
- **Accommodation Statement:** The University of Massachusetts Amherst is committed to providing an equal educational opportunity for all students. If you have a documented physical, psychological, or learning disability on file with Disability Services (DS), you may be eligible for reasonable academic accommodations to help you succeed in this course. If you have a documented disability that requires an accommodation, please notify the instructor within the first two weeks of the semester so that we may make appropriate arrangements.
- **Class Attendance Policy:** Students are required to attend class in order to participate in discussions. Participation is a graded component of the course. Students should let the instructor know if they are not able to attend a class as early as possible. Approved absences will be granted in accordance with the University attendance policy (e.g., in the case of illness, religious observances, official University travel,

and other extenuating circumstances). In particular, students who are ill should let the instructor know in advance of the class meeting time if possible and should not attend class.

- **Late Work and Rescheduling Policy:** All written work must be submitted by the indicated due date and students must present in their scheduled slots in order for their work to count for credit unless they have obtained an approved extension. Extensions and res-scheduling of presentations will be granted in in accordance with the University attendance policy (e.g., in the case of illness, religious observances, official University travel, and other extenuating circumstances).

- **General academic honesty statement:** Since the integrity of the academic enterprise of any institution of higher education requires honesty in scholarship and research, academic honesty is required of all students at the University of Massachusetts Amherst. Academic dishonesty is prohibited in all programs of the University. Academic dishonesty includes but is not limited to: cheating, fabrication, plagiarism, and facilitating dishonesty. Appropriate sanctions may be imposed on any student who has committed an act of academic dishonesty. Instructors should take reasonable steps to address academic misconduct. Any person who has reason to believe that a student has committed academic dishonesty should bring such information to the attention of the appropriate course instructor as soon as possible. Instances of academic dishonesty not related to a specific course should be brought to the attention of the appropriate department Head or Chair. Since students are expected to be familiar with this policy and the commonly accepted standards of academic integrity, ignorance of such standards is not normally sufficient evidence of lack of intent.

  http://www.umass.edu/dean_students/codeofconduct/acadhonesty/

- **Course Academic Honesty Policy:** Students must write original and individual responses to readings and must produce original presentations. Similarly, course projects must be original work produced by the student for this course. The course project may not overlap with projects for other courses that the student is currently enrolled in. Course projects may build on students' existing research, but must represent new directions developed for this course. Plagiarism and copying of reading responses is considered cheating. Plagiarism of presentation materials is considered cheating. Plagiarism of course project materials is considered cheating.

- **Course Material Intellectual Property Policy:** The presenter retains intellectual property rights for the materials that they present. This holds both for the instructor and for students in the course. Students are not permitted to make recordings of presentations or to share presentation materials without the consent of the presenter. Students are not permitted to make recordings of in-class discussions.

Schedule of Topics: (Subject to change over the semester)

Meeting	Topics/Papers
Week 1	Course Introduction. Review of classical MCMC methods.
Week 2	Neal. *Bayesian learning via stochastic dynamics*. NeurIPS 1993.
Week 3	Welling and The. *Bayesian Learning via Stochastic Gradient Langevin Dynamics.* ICML 2011.
Week 4	Chen, Fox, and Guestrin. *Stochastic Gradient Hamiltonian Monte Carlo.* ICML 2014.
Week 5	Zhang et al. *Cyclical Stochastic Gradient MCMC for Bayesian Deep Learning*. ICLR 2020.
Week 6	Izmailov et al. *Subspace Inference for Bayesian Deep Learning*. UAI 2019.
Week 7	Maddox et al. *A Simple Baseline for Bayesian Uncertainty in Deep Learning*. NeurIPS 2019.
Week 8	Lakshminarayanan. *Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles.* NeurIPS 2017.
Week 9	Review of classical variational inference methods and variational Bayes.

Week 10	Blundell et al. *Weight Uncertainty in Neural Networks.* ICML 2015.
Week 11	Gal and Ghahramani. *Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning.* ICML 2016.
Week 12	Korattikara et al. *Bayesian Dark Knowledge.* NeurIPS 2015.
Week 13	Final project presentations.