TO: The Faculty of the College of Engineering

FROM: Elmore Family School of Electrical and Computer Engineering

RE: New Graduate Course, ECE 60131 Inference and Learning in Generative Models

The faculty of the School of Electrical and Computer Engineering has approved the following new course. This action is now submitted to the Engineering Faculty with a recommendation for approval.

ECE 60131 Inference and Learning in Generative Models
Sem. 2, Lecture 3, Cr. 3.
Prerequisite: Linear algebra, multivariable calculus, and basic probability theory

Description: An introduction to modern generative models like diffusion models (e.g., "Stable Diffusion"), variational autoencoders, normalizing flows, and energy-based models, with a focus on derivations from the perspective of statistical learning theory. We build up from the basics, starting with probabilistic graphical models, which provide the framework for many of the ideas in the class. What exactly are generative models? They are a powerful alternative to discriminative models that, when properly specified, estimate their parameters more efficiently and can generate samples from the distribution of their input data, but also can be used (like discriminative models) to infer features or labels from their inputs. However, the generative and inferential faculties typically come at each other's expense. This course will cover five different attempts at finessing this trade-off, and the resulting learning algorithms: exact inference in directed graphical models (EM algorithm); sampling-based methods in undirected (energy-based) models; deterministic approximate inference ("variational" methods, e.g. VAEs); invertible, deterministic models (e.g. ICA, normalizing flows); and adversarial training (GANs).

Reason: Generative models have had an enormous impact on public life in the last year—image, audio, and text synthesis—and this is likely to continue. This is an extremely active research area and the course will prepare students for high level research.

Course Enrollment History: Spring 2021 – 13, Spring 2022 – 10, Spring 2024 - 38

T.S. Mithuna

Mithuna Thottethodi,
Associate Head for Teaching and Learning
Elmore Family School of Electrical and Computer Engineering

ECE 60131: Inference & Learning in Generative Models

Spring 2024

Instructor:	J.G. Makin	**Email:**	jgmakin@purdue.edu
Class time:	MWF 2:30 – 3:20	**Class location:**	Forney Hall B124
Office hours:	M 3:30 – 5:30	**Office:**	BHEE 330

Course Information

- **ECE69500**: Inference and Learning in Generative Models
- **CRN**: 20748
- **Section**: 090
- **Meeting days; times; location**: MWF; 2:30 – 3:20; Forney Hall B124
- **Instructional modality**: In person. Recorded lectures may additionally be provided for some content.
- **Course credit hours**: 3
- **Prerequisites**: Linear algebra, calculus, and basic probability theory

Instructor Information

- **Instructor:** Assistant Professor J.G. Makin, Purdue ECE
- **GTA**: Bilal Ahmed, Purdue ECE

Course Description

An introduction to modern generative models like diffusion models, variational autoencoders, normalizing flows, and energy-based models, with a focus on derivations from the perspective of statistical learning theory. We build up from the basics, starting with probabilistic graphical models, which provide the framework for many of the ideas in the class.

What exactly are generative models? They are a powerful alternative to discriminative models that, when properly specified, estimate their parameters more efficiently and can generate samples from the distribution of their input data, but also can be used (like discriminative models) to infer features or labels from their inputs. However, the generative and inferential faculties typically come at each other's expense. This course will cover different attempts at finessing this trade-off, and the resulting learning algorithms: exact inference in directed graphical models (EM algorithm); sampling-based methods in undirected (energy-based) models; deterministic approximate inference ("variational" methods, e.g. VAEs and diffusion models); invertible, deterministic models (e.g. ICA, normalizing flows); and adversarial training (GANs).

Learning Resources, Technology & Texts

- **Required texts:** The course will be based primarily on unpublished textbooks, chapters of which will be distributed via Brightspace.
- **Recommended texts:** Bishop, Christopher M. *Pattern Recognition and Machine Learning*. New York: Springer, 2006. Print. Information Science and Statistics. ISBN: 9780387310732
- **Additional reading:** Papers, notes, etc., will be posted on the course Brightspace page.
- **Software/web resources:**
  - Programming assignments should be completed in Python or Matlab. If you strongly prefer another programming language (e.g., Julia or C), please contact the instructor.

- **Brightspace learning management system (LMS):** Announcements and lecture notes will be available on the class's Brightspace page.

Learning Outcomes

By the end of the course, you will be able to:

1. *identify the trade-offs in inference and generation in generative models*
   Methods of evaluation: homework assignments, exams
2. *translate a description of data into a probabilistic graphical model*
   Methods of evaluation: homework assignments, exams
3. *identify the appropriate inference algorithm for a dataset*
   Methods of evaluation: homework assignments, exams, project
4. *implement the inference and learning algorithms for common generative models, both classical (Kalman filter, forward-backward, EM, particle filter, etc.) and modern (VAE, diffusion model, etc.)*
   Methods of evaluation: homework assignments, exams
5. *present to their peers novel results that use the ideas in this course*
   Methods of evaluation: project

Assignments

Mastery of the material will be assessed through a combination of homework, projects, and exams. **Homework will generally be due biweekly, at 11:59 p.m. on Sunday evenings.**

Assignments	**Due Date**	**Points**
Homework 1	January 22	7
Homework 2	February 5	7
Homework 3	February 19	7
Exam 1	February 26	16
Homework 4	March 4	7
Homework 5	March 18	7
Homework 6	April 1	7
Homework 7	April 15	7
Exam 2	April 19	16
Project	April 22	21
		Total: 100

Grading Scale

Classroom instructors necessarily serve two masters. They must teach the students, and at the same time evaluate and (implicitly) rank them for future employers/schools/etc. For graduate classes, when these aims are at odds with each other, I try to err on the side of the first.

Using the point breakdowns given above, grades will be assigned as follows:

A+	100
A	90–99
B	80–90
C	70–80
D	60–70
F	0–60

This is an experimental class, and the instructor reserves the right to be more generous with grades later in the semester.

Academic Integrity

Any student found cheating in any way on exams, homework assignments, or the class project will fail the class and be reported to the dean.

Nondiscrimination Statement

A hyperlink to Purdue's full Nondiscrimination Policy Statement is included in our course Brightspace under University Policies.

Accessibility

If you anticipate or experience physical or academic barriers based on disability, please talk to the instructor at the beginning of the term to discuss options. You are also encouraged to contact the Disability Resource Center at: drc@purdue.edu or by phone: 765-494-1247.

Mental Health/Wellness Statement

If you find yourself beginning to feel some stress, anxiety and/or feeling slightly overwhelmed, try WellTrack. Sign in and find information and tools at your fingertips, available to you at any time.

If you need support and information about options and resources, please contact or see the Office of the Dean of Students. Call 765-494-1747. Hours of operation are M–F, 8:00–17:00.

If you find yourself struggling to find a healthy balance between academics, social life, stress, etc. sign up for free one-on-one virtual or in-person sessions with a Purdue Wellness Coach at RecWell. Student coaches can help you navigate through barriers and challenges toward your goals throughout the semester. Sign up is completely free and can be done on BoilerConnect. If you have any questions, please contact Purdue Wellness at evans240@purdue.edu.

If you're struggling and need mental health services: Purdue University is committed to advancing the mental health and well-being of its students. If you or someone you know is feeling overwhelmed, depressed, and/or in need of mental health support, services are available. For help, such individuals should contact Counseling and Psychological Services (CAPS) at 765-494-6995 during and after hours, on weekends and holidays, or by going to the CAPS office on the second floor of the Purdue University Student Health Center (PUSH) during business hours.

Emergency Preparation

In the event of a major campus emergency, course requirements, deadlines and grading percentages are subject to changes that may be necessitated by a revised semester calendar or other circumstances beyond the instructor's control. Relevant changes to this course will be posted on the course's Brightspace page.

Course Schedule

Date	Lecture	Topics	Reading
M 01/08	1	*review of basic probability* (notation; pdfs, pmfs, RVs; marginalization, conditioning, statistical independence)	IPGM [2]
W 01/10	2	*directed graphical models* (parameterization, chain rule of probability, d separation, Bayes ball algorithm)	IPGM [2]
F 01/12	3	*undirected graphical models* (parameterization, conditional independence, energy formulation)	IPGM [2]
M 01/15	–	**NO CLASS — MLK DAY**	–
W 01/17	4	*probabilistic inference* (Bayes rule, the elimination algorithm)	IPGM [3]
F 01/19	5	*probabilistic inference* (graph elimination, sum-product algorithm)	IPGM [3, 4]
M 01/22	6	*probabilistic inference* (factor graphs, sum-product on factor trees)	IPGM [4]
W 01/24	7	*the junction-tree algorithm* (part 1)	IPGM [17]
F 01/26	8	*the junction-tree algorithm* (part 2)	IPGM [17]
M 01/29	9	*exponential families* (form, examples, parameters)	IPGM [8]
W 01/31	10	*exponential families* (cumulants, moment param., sufficient statistics)	IPGM [8]
F 02/02	11	*introduction to learning* (entropy [information/cross/relative], maximum likelihood)	IMSL [4]
M 02/05	12	*discriminative models* (supervision, linear and logistic regression, classification)	IPGM [5, 6], IMSL [5]
W 02/07	13	*generalized linear models (GLiMs)* (form; fitting with IRLS)	IPGM [8]
F 02/09	14	*generative models* (LDA, discriminative vs. generative models)	IPGM [5, 6]
M 02/12	15	*Bayesian [vs. frequentist] statistics*	IPGM [5]
W 02/14	16	*conjugate priors*	IPGMv2 [9]
F 02/16	17	*exact inference in directed models: HMMs*	IMSL [2.2.1], IPGM [12]
M 02/19	18	*exact inference in directed models: linear-Gaussian models*	IMSL [2.1.2], IPGM [13]
W 02/21	19	*Kalman filtering and RTS smoothing*	IMSL [2.2.2], IPGM [15]
F 02/23	–	**EXAM REVIEW**	–
M 02/26	–	**EXAM 1**	–
W 02/28	20	*Unsupervised learning with exact inference: the EM algorithm*	IMSL [6], IPGM [11]
F 03/01	21	*EM algorithm for GMMs and HMMs*	IMSL [7], IPGM [10, 12]
M 03/04	22	*EM algorithm for factor analysis and state-space models*	IMSL [7], IPGM [14]
W 03/06	23	*sparse coding and ICA*	IMSL [2.1.3], IMSL [8.1]
F 03/08	24	*variational autoencoders*	IMSL [8.2]
M 03/11	–	**NO CLASS — SPRING BREAK**	–
W 03/13	–	**NO CLASS — SPRING BREAK**	–
F 03/15	–	**NO CLASS — SPRING BREAK**	–
M 03/18	25	*variational autoencoders* (cont.)	IMSL [8.2]
W 03/20	26	*diffusion models*	IMSL [8.3]
F 03/22	27	*normalizing flows*	Rezende 2015, Dinh 2015/2016
M 03/25	28	*variational inference*	PRML [10]
W 03/27	29	*variational inference* (cont.)	PRML [10]
F 03/29	30	*sampling* (standard distributions, rejection sampling, importance sampling)	PRML [11]
M 04/01	31	*sampling* (Markov chain Monte Carlo & Gibbs sampling)	PRML [11]
W 04/03	32	*sampling* (Hamiltonian Monte Carlo and Langevin dynamics)	Neal 2012, Betancourt 2018
F 04/05	33	*energy-based models* (RBM/EFH; contrastive divergence)	IMSL [10.1]
M 04/08	34	*energy-based models* (deep belief networks)	IMSL [10.1]
W 04/10	35	*noise-contrastive estimation (NCE)*	Gutmann 2012
F 04/12	36	*InfoNCE*	van den Oord 2018
M 04/15	37	*generative adversarial networks*	Goodfellow 2014
W 04/17	–	**EXAM REVIEW**	–
F 04/19	–	**EXAM 2**	–
M 04/22	–	**PROJECT PRESENTATIONS**	–
W 04/24	–	**PROJECT PRESENTATIONS**	–
F 04/26	–	**PROJECT PRESENTATIONS**	–