ECE 57000: Artificial Intelligence Syllabus (Tentative)

Course Information

ECE 57000
CRN: 34846
Fall, 2019
MWF, 12:30pm-1:20pm
ME 1061
3 Credit Hours
Course website including schedule and due dates:
https://www.davidinouye.com/course/ece57000-fall-2019/
Piazza:
https://piazza.com/purdue/fall2019/ece57000

Instructor

David I. Inouye
WANG 3063
765-496-0238
dinouye@purdue.edu
Office Hours: Mon 3-4pm, Thurs 3-4pm, or by appointment in WANG 3063
Instructor's Webpage: https://www.davidinouye.com/

Teaching Assistant

Liming Wu Sr
wu1114@purdue.edu
Office Hours: Tues/Thurs 2-3pm, EE 134

Course Description

This course will provide a graduate-level introduction to artificial intelligence (AI) with a primary focus on unsupervised learning. Topics will include clustering, mixture models, density estimation, representational learning, and deep generative models. The lecture content will focus on key concepts and intuitions. The course project will enable students to dive deeper into a topic of their choice.

Expectations / Prerequisites

This class is oriented towards first-year graduate students. While this class does not have official prerequisites, the course will expect basic knowledge of linear algebra, probability / random variables, and Python programming. I will try to make this course as accessible as possible by briefly reviewing some of these concepts, but if you are not familiar with these subjects, this may not be the right class for you; or, you will have to (re-)learn these concepts on your own.

This syllabus is required reading and you will be required to know the policies outlined in this syllabus. Questions about the syllabus may appear on quizzes.

Learning Outcomes

1. Basic understanding of unsupervised learning and density estimation, subareas of AI. (lecture content)
2. Deeper understanding of one area of interest in AI. (course project)
3. Basic research skills and hands-on experience using current toolkits. (assignments + course project)

Textbooks

No required textbook. We will not follow any particular textbook but related reading will be posted on the course schedule (see course website above) if appropriate. Below are a few supplemental textbooks that may be useful.

(Optional) *Deep Learning* by Ian Goodfellow, Yoshua Bengio, and Aaron Courville, 2016. –Available for free online at http://www.deeplearningbook.org or physical copy available on Amazon. Part I is a good introduction to core concepts that covers the fundamental mathematical, computational, and machine learning concepts useful for the rest of this class. Particularly, chapter 2 on linear algebra and chapter 3 on probability will likely be the references for the relevant content covered in class. Note: We will not be covering the rest of the book because only the last few topics are directly related to deep learning.

(Optional) *Machine Learning: A Probabilistic Perspective* by Kevin P. Murphy, 2012. Available online via Purdue's library system: https://ebookcentral.proquest.com/lib/purdue/detail.action?docID=3339490. – This is a good reference book and some chapters may help supplement the lectures. Chapter 2 on probability may be a useful reference.

(Optional) *Python Data Science Handbook* by Jake VanderPlas, 2016. Available for reading online at https://jakevdp.github.io/PythonDataScienceHandbook/. This could be a useful book for some of the Python tools used in this course including Jupyter notebooks (which are great for interactive and visual python data analysis and are widely used), NumPy (the fundamental library for fast numerical computation), and scikit-learn (the standard machine learning library in Python for traditional non-deep models). We may also use Pandas for data loading and preprocessing. This is not a general reference for learning Python from scratch but can be useful for understanding common tools. If you are trying to learn Python from scratch, it might be useful to work through this book but supplement it with general material online about Python programming.

Computing resources

All programming assignments will be designed to execute on standard Purdue ECE hardware such as https://ecegrid.ecn.purdue.edu or similar. No GPUs or advanced hardware will be required for the programming assignments or course project.

Depending on the course project, high-performance or GPU hardware may be useful for reducing development and evaluation time but the instructor will not be able to provide access to this specialized hardware.

Topics (tentative)

1. Introduction to Artificial Intelligence
2. Clustering
    a. K-means
    b. Applications of clustering
3. Brief review of some linear algebra
    a. Vector and matrices
    b. Inner product
    c. Norms
    d. Orthogonal matrix
    e. SVD Decomposition
4. Brief review of probability basics
    a. Random variable
    b. Expectation of random variables
    c. Probability mass/density functions (PDF, PMF)
    d. Cumulative Distribution Function (CDF)
    e. Inverse CDF
    f. Multivariate distributions
    g. Conditional distributions
    h. Marginal distributions
    i. Bayes Rule
5. Gaussian Mixture Models
    a. Expectation Maximization (EM algorithm)
6. Density estimation
    a. Histograms
    b. Kernel density
    c. Tree density
7. Introduction to probabilistic graphical models
    a. Directed/Bayesian networks
    b. Undirected/Markov networks
    c. Image application
8. Topic modeling
    a. Multinomial distribution
    b. Mixtures of multinomial distributions
    c. Latent Dirichlet Allocation (Admixture)
9. Representational learning
    a. Principal Component Analysis (PCA)
    b. Independent Component Analysis (ICA)
    c. Destructive deep learning
10. Other deep generative models
    a. Variational Auto-encoders (VAE)
    b. Generative Adversarial Networks (GAN)
    c. Invertible flows
11. (maybe) Explainable AI
12. (maybe) Self-supervised learning

Course Assessment and Grading

No +/- for grades; only letter grades A, B, C, D, or F. Letter grades will be assigned at the end of the semester. The instructor will determine the final grade cutoffs but the cutoffs will be equal to

or lower than the standard cutoffs. For reference, the standard cutoffs are A if >= 90%, B if >= 80%, C if >= 70%, D if >= 60%, and F otherwise (e.g. if you have >= 90% you are guaranteed an A).

- 15% - Programming assignments
- 15% - Quizzes
- 70% - Course project
    - 5% - Selection of 3 research papers
    - 15% - Implementation
    - 15% - Term paper
    - 15% - Presentation (5min video)
    - 10% - Reviews
    - 10% - Attendance at project presentations

Missed or Late Work

Quizzes - Quizzes cannot be made up. See "Quizzes" section below.

Programming assignments and course project deliverables will be due **on or before 12:00pm (noon)** on the day they are due (see class schedule).

- One day late: 80% of original score
- More than one day: 0%

Quizzes

There will be 5-10 short quizzes throughout the semester. There are several reasons for quizzes: (1) assess knowledge of content, (2) provide feedback (both for students and instructor), and (3) encourage attendance at course lectures. The quizzes may be related to lecture or programming assignment content. The quizzes will be designed to be relatively easy if you have attended the lectures and have completed the programming assignments. They will likely be multiple choice.

Quizzes cannot be made up under any circumstance. To account for legitimate absences, I will drop the two lowest quiz grades in the final grade calculation for everyone.

Programming Assignments

There will be 3-5 programming assignments throughout the semester. The assignments are meant to provide hands-on experience with a few common tools of AI/ML as well as develop intuitions regarding the content.

Course Project

(Instructions below adapted from Prof. Jeffrey Siskind's ECE57000 syllabus from last year.) Logistics and mechanisms for submitting various parts of the project will be finalized later in the semester.

The course project will be individual. No group projects. However, I encourage you to discuss your papers and project with other students. Explaining your papers or project to someone else can help test your own understanding.

Research Paper Selection (at least 3)

Students will be required to select and read at least three recent conference or journal research papers in the fields of AI, computer vision, natural language processing, or machine learning. In general, the papers should have been published within the **last 3 years** in one of the following venues: AAAI, IJCAI, CVPR, ICCV, ECCV, ACL, NAACL, EMNLP, NIPS/NeurIPS, ICML, ICLR, AISTATS, or JMLR. One paper can be older than 3 years but must have > 500 citations on Google scholar (if older, the citation count must be included in the "note" filed of the BibTeX entry. Finally, given the implementation requirement below, one of the papers must have material that can be implemented.

You are welcome to read more than three papers. Your selection of papers must be approved by me. I may be willing to accept papers even if they do not fit the criteria above, but all paper selections, whether or not they meet the above criteria, must be approved by me. If you intend to select papers that are older than three years old or from a venue other than listed above, I suggest that you discuss this with me prior to the due date.

You will be required to submit a BibTeX entry for each paper including a URL for downloading a pdf of the paper. For a journal paper, this should contain (at least) the paper title, paper abstract, authors, journal, volume, year, pages and pdf URL. For a conference paper, this should contain (at least) the paper title, paper abstract, authors, conference, year, and pdf URL. Please note that you should not send me the paper itself, just the BibTeX entry.

BibTeX journal example:

```
@article{inouye2017review,
 abstract = {The Poisson distribution has been widely studied … [truncated but you should include full abstract]},
 author = {Inouye, D. I. and Yang, E. and Allen, G. I. and Ravikumar, P.},
 journal = {Wiley Interdisciplinary Reviews: Computational Statistics},
 number = {3},
 pages = {e1398},
 title = {A Review of Multivariate Distributions for Count Data Derived from the Poisson Distribution},
 volume = {9},
 year = {2017},
 note = {XXX citations},
 url = {https://arxiv.org/pdf/1609.00066.pdf}
}
```

Implementation

You will be required to do at least one of the following:

1. Reimplement paper idea: Implement and evaluate the ideas from at least one paper. If an implementation of the paper is already available (e.g. from the author's website), you **must state this in your report** and compare your implementation to the existing implementation, both in terms of code and performance.
2. Implement new idea (much harder, research-oriented): Propose, implement, and evaluate an extension or novel idea related to the papers you read. For this option, you can build off of any existing implementation but your implementation must extend or alter the original idea in a significant way.

The implementation must be nontrivial. A good guideline is that the implementation should be at least four pages of code. This is not a strict guideline. Ultimately, I will determine whether or not the implementation meets the non-triviality requirement. If you have questions about your implementation, please contact me to discuss.

The implementation must be in Python. **(UPDATE)** You may use whatever libraries you would like such as PyTorch, TensorFlow, and Keras. However, I am only familiar with PyTorch and we will likely have at least one assignment in PyTorch so if all things are equal, I would suggest using PyTorch. I encourage you to have your implementation conform to the **scikit-learn estimator API** (see the APIs of scikit-learn objects, Rolling your own estimator and other scikit-learn documentation). This would enable interfacing with other parts of scikit-learn including hyperparameter tuning or pipelining. **(UPDATE) Conforming to the scikit-learn estimator API is no longer required but is encouraged.**

You must conduct a substantive evaluation of your implementation to determine how well it solves the intended problem. Ideally, you should replicate the experiments presented in the paper but I will not require this.

Term Paper

You will be required to write a six page paper in LaTeX meeting the typesetting conventions of ICML (see ICML 2019 author instructions for LaTeX template and instructions). Approximately three pages of this paper should be a substantive critique of the three (or more) papers that you have read. And approximately three pages of this paper should be a description of your implementation and

evaluation of the material from one of the papers. (Note: **You should not put your name on the term paper to accommodate the reviews.** See "Reviews" section below.)

5 Minute Spotlight Video

You will also be required to create a 5 minute video presentation of your term paper. The video should be between 4-5 minutes but not longer than 5 minutes. This "spotlight" presentation should cover

1. A brief summary and critique of the three papers that you have read
2. A description of your implementation.
3. A description and discussion of your evaluation.
4. (Optional) Any concluding thoughts or future directions.

You should only have between 5-10 slides to fit within the 5 minutes. One minute per slide is usually reasonable.

You can see real examples of video abstracts from top machine learning conference at https://nips.cc/Conferences/2018/Schedule?type=Poster (look for "3-min video" links). I do not expect your video to be at the quality level of these videos, but it hopefully gives you some ideas. Note your video will be 5 minutes instead of 3 minutes so you will have a little more time than the examples above. The video can be in any format you want including animations, slides with narration, video of you presenting, or combination of the above. However, the focus should be on **clarity** rather than *fanciness*.

Project Presentation

Given the large number of students in the class, we will randomly select students to present during the presentation class periods (see class schedule). The first presentation class period will be Nov 18 so **all students should be ready to present by Nov 18**. We will announce 1-2 hours before class which students will be presenting. The format will be 5 minutes (like the video) + 2 minutes for questions while the next presenter sets up.

The presentation must be given from a laptop or the computer that is built-in to the lectern in the lecture hall. You can use whatever tools you wish to prepare your presentation (i.e. LaTeX/beamer or PowerPoint). If you are scheduled to give your presentation, you should arrive in class early to make sure that your presentation setup works and that you are prepared to give your presentation in the allotted time slot.

Reviews

The term papers will be reviewed by other students in the class. Like all conferences, this process will be double blind: reviewers will not know the identity of authors and vice versa. To support this, like all conferences, **you should NOT put your name on the term paper submission or on your reviews.** Also, like conferences, reviews will be confidential. The only person who will be privy to the reviews will be the reviewer, the instructor, and the author.

I will assign each term paper to five students to read. Each student will be required to read five student term papers (other than their own) and prepare conference-style reviews, primarily indicating clarity and the quality of the implementation effort. The exact format for the review will be determined later in the semester.

Attendance Policy

Students are expected to be present for every class. If a student misses a class, **the student** is still responsible for knowing all missed content (including announcements). The instructor **will not** provide lecture notes to students who have missed a class. Rather, the student should obtain lecture notes from other classmates for missed classes. Also, see quiz section for additional policies.

Attendance for project presentations (Nov 18, 20, 22, 25 and Dec 2, 4, 6) is mandatory. Any student who cannot make class those dates must notify the instructor via email as soon as possible. The penalty for missing even a single project presentation class without obtaining prior approval from the instructor will be at least 2.5%.

Communication Policies

All students are expected to sign up for the Piazza class this week or earlier and regularly check for announcements. Piazza will be used for all announcements and class discussions. Students are expected to check Piazza regularly for announcements. Please post any questions to Piazza instead of email unless the issue is confidential in nature. This should help you receive answers faster. Additionally, you can post anonymously if you feel uncomfortable posting with your name. If you email the instructor or TA with a question, we will likely ask you to re-ask the question on Piazza if appropriate.

I will be available via email during normal business hours (9am-5pm, Mon-Fri), and try to respond as soon as possible (generally within 1-2 business days). When emailing me, please place the course number/section in the subject line of the email (*e.g.*, ECE 57000 …). This will help me locate and respond to your emails quickly. **Please use Piazza for asking questions whenever possible so that others can benefit from your question and my response.**

Academic Integrity

Any cheating or dishonesty will be penalized with a failing grade. All acts of cheating or dishonesty will be reported to the Dean of Students.

Academic integrity is one of the highest values that Purdue University holds. Individuals are encouraged to alert university officials to potential breaches of this value by either emailing **integrity@purdue.edu** or by calling 765-494-8778. While information may be submitted anonymously, the more information that is submitted provides the greatest opportunity for the university to investigate the concern. Please see Purdue's student guide for academic integrity for more information. When in doubt, please discuss with the instructor. It is far better to discuss difficult circumstances with the instructor, than to be dishonest.

Academic Integrity Policy for Programming Assignments

For programming assignments, feel free to talk with other classmates about the assignments. However, do not view, share, or copy code in any way (even just looking at someone's screen or writing on a whiteboard). Only talk, no code. Neither should you talk in such detail that you would both write the same programs. I reserve the right to run code similarity checks on all submitted code to find instances of academic dishonesty.

Also, **sharing your code** with other students is also an academic violation. Do not share code with anyone if they ask. You should refer them to the instructor if there are any extenuating circumstances.

Academic Integrity Policy for Course Project

Because each project should be unique, I will allow more freedom in collaborating or discussing the course project code. You can help each other debug or discuss the actual code. Both students will likely learn from this exercise. However, the student associated with the project should be the one actually writing all the code. Your project implementation should be your work. Additionally, no student should feel any obligation to help another student.

If you let others view your code and help in any substantial way, you must acknowledge their help in your term paper under an "Acknowledgements" section at the end of the term paper (e.g. "John Doe helped me debug my code."). If there is any doubt, please discuss with me before engaging in that activity.

Emergency Statement

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 onto the course website or can be obtained by contacting the instructors or TAs via email or phone. *You are expected to read your @purdue.edu email on a frequent basis.*

Nondiscrimination Statement

Purdue University is committed to maintaining a community which recognizes and values the inherent worth and dignity of every person; fosters tolerance, sensitivity, understanding, and mutual respect among its members; and encourages each individual to strive to reach his or her own potential. In pursuit of its goal of academic excellence, the University seeks to develop and nurture diversity. The University believes that diversity among its many members strengthens the institution, stimulates creativity, promotes the exchange of ideas, and enriches campus life. Purdue's nondiscrimination policy can be found here.

Accessibility and Accommodations

Purdue University strives to make learning experiences as accessible as possible. If you anticipate or experience physical or academic barriers based on disability, you are welcome to let me know so that we can discuss options. You are also encouraged to contact the Disability Resource Center via email or by phone: 765-494-1247.

Mental Health 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 see the Office of the Dean of Students, http://www.purdue.edu/odos, for drop-in hours (M-F, 8 am- 5 pm).
- **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 of the second floor of the Purdue University Student Health Center (PUSH) during business hours.

Disclaimer

This syllabus is tentative and is subject to change. See course website for course schedule including due dates; the course schedule is also subject to change. Any substantive changes will be announced in class and on Piazza.