

Machine Learning in Business (MAR6930)

Course Syllabus

Overview

Information.

Instructor: Dr. Mingzhang Yin

Email: mingzhang.yin@warrington.ufl.edu

Office Hour: after class to 4:20 PM

Course time and location: Tuesday 12:50 PM - 3:50 PM at STZ 267G.

Course objective. The course is a research-oriented course on probabilistic machine learning methods and their adoption in business and marketing research. We will focus on building probabilistic models, algorithms for approximate Bayesian inference, and causal inference. Some of the models we will study include classic Bayesian mixture and regression models, hierarchical models, factor models, and topic models. Alongside these models, we will study algorithms for approximate Bayesian inference including Markov Chain Monte Carlo and variational inference algorithms. We will study recent research on uncertainty quantification such as Bayesian optimization and conformal prediction, generative models such as diffusion models and language models, and methods at the intersection of causal inference and machine learning. Finally, we will study business research papers related to these methodology, completing a virtuous cycle of machine learning in business.

At the end of this course, students will be familiar with the probabilistic machine learning paradigm, and will be able to build connections with business research.

Communications.

Please send all course-related communications via **Slack**. You can expect an answer within 24 hours during weekdays. You are welcome to contact the instructor using the above email address in emergency.

Class notes, readings, and assignments will be available on **Github**.

Prerequisite.

Students are expected to have mathematical and statistical backgrounds, including integral calculus courses. Basic coding skills in Python is preferred.

Textbooks and resources. Through out the course, we will refer to the following books. Some books have public version online.

- Bishop. Pattern Recognition and Machine Learning. New York: Springer, 2006. (PRML) [PDF](#)
- Murphy. Probabilistic Machine Learning: Advanced Topics. MIT Press, 2023. [PDF](#) (PML)
- Pearl, Glymour, and Jewell. Causal Inference in Statistics A Primer. WILEY, 2016. [PDF](#) (PRIMER)

Schedule. (The schedule is subject to change.)

Week	Number of Contact Hours Related to AI	AI-Related Topic	Details on AI-Related Readings, Projects, and Assignments
Week 1	3h	Probability: Bayes' Theorem, Probabilistic model, Graphical model, Conjugate Model, Exponential Family	
Week 2	3h	Mixture model, hierarchical model, factor model, topic model, block model and mixed membership model	
Week 3	3h	MCMC, Metropolis Hasting and Gibbs sampling	
Week 4	3h	EM and variational inference	
Week 5	3h	VI, Neural networks	Exercise, Project Proposal
Week 6	3h	<i>Presentation</i>	
Week 7	3h	<i>Presentation</i>	
Week 8	3h	Variational autoencoders	
Week 9	3h	VAE cont., Diffusion model	Project Milestone
Week 10	3h	Causal Inference	
Week 11	3h	Causal Representation Learning	
Week 12	3h	<i>Presentation</i>	
Week 13	3h	<i>Presentation</i>	
Week 14	3h	<i>Presentation & Final Project Presentation</i>	Project Report

RealDeal™ rules: no phones, no email, no text, etc; tablets ok.

Requirements and Grades

The course assessment is based on Attendance and Active Discussion (10%), Exercises (10%), Presentation (40%), and Final Project (40%). Please refer to [UF grading policies](#).

Excercise.

There will be one light loaded exercise at the starting weeks of the course.

Presentation.

You will present several papers in business research that relate to the machine learning methods we studied. You will prepare slides for the presentation that cover the whole paper. You will play two roles: explain the paper like the authors and comment on the paper as the reviewers. The presentation should aim to answer the following questions:

- What are the business problems the paper addresses?
- Why the problem is important, to whom?
- What are the current solutions/literature background to the problem; why they are not enough? What is the novelty of the proposed approach?
- Why the proposed method is a better fit?
- Which component(s) of the machine learning model makes it suitable to the business problem?
- What are the contributions (methodological, substantive, etc)? listing one by one.
- What are the empirical findings? What evidences support each finding?
- What are the managerial implications?
- What are your comments on the paper: e.g., what are the strongest and weakest parts? methods and data are proper? how to improve the shortcomings? what applications can be extended to?

Final Project.

The final project is an integral part of this course. There is no better way to learn about machine learning than by applying it to a business research question of your own. The final project is an individual project. You have a lot of freedom in choosing a topic for your final project. The only criterion is that it deeply involves applying probabilistic machine learning to a real-world

business-related problem. You choose a dataset, an interesting question about it, and address it with machine learning modeling.

The project involves three assignments: project proposal, project milestone, and project report. For each, **please use 12-point font, single-spaced, and 1-inch margins. Page limits are without figures**; include as many pages of figures as needed; please place all figures and tables at the end of the report with proper labels and captions. We offer a LaTeX template; you are encouraged to use it but not required.

- **Project Proposal.** The project proposal is an abstract (i.e. a paragraph) that imagines the completed project. We understand that your project will evolve and change over the semester, but writing an abstract early is a good way to plan and think about what it would mean to successfully complete it. ([Kaggle](#) is a good place to find interesting datasets.)

- **Project Milestone.** The project milestone describes the problem you are addressing and discusses some preliminary results. Include what you have completed and what you plan to finish by the end of the semester. As a guideline, the project milestone should (at least) answer the following questions.

- Clarify the purpose of data analysis:
 - What is the problem of interest?
 - What is the final goal you want to achieve?
- To the best of your knowledge, explain the data as thoroughly as possible:
 - How is the data collected?
 - Is the data a simple random sample from some population?
 - What is the potential defect in the data?
 - If you were able to collect the data, what would you have done differently?
 - Are you going to account for the data collecting process in your analysis?
 - Describe simple facts about the data: size, data structure & source, etc.
 - If you are not going to use real data, explain how you are going to generate “realistic” synthetic data that represents your knowledge of the problem.
- Draft a plan for data analysis:
 - What model are you going to use for the data analysis?
 - Have you done any preliminary analysis of the data?

- What is your timeline for (1) proposing and fitting a preliminary model, (2) examining initial results and adjusting models further if needed, (3) summarizing and writing a final report.
- What are the (potential) difficulties with your data analysis?

The project milestone is 1–2 pages long.

- **Project Report.** The project report is due at the end of the semester. The report should have a similar format to that of an academic paper. (We may not read the appendix when grading your work.)

The project report should (at least) include

- Introduction: Clear description of the problem. Describe the problem you are trying to solve, why it is important or useful, and summarize any important pieces of prior work that you are building upon.
- Dataset: Clear description/visualization of the data. Describe the dataset or datasets you are working with. Show examples from the datasets. If you collected or constructed your own dataset, explain the process you used to collect.
- Method: Clear and thorough description of statistical analysis. Describe the method you are using; this may also contain parts of the implementation of your model, loss function, or other components along with sanity checks to ensure that those components are correctly implemented.
- Experiments: Clear and thorough interpretation of results. Describe the experiments you did, and key results and figures that you obtained. This may interleave explanations of the experiments you run and figures you generate as a result of those experiments.
- Conclusion / Future work: What did you learn in doing this project? What are the shortcomings or failure cases of your work? If you had more time or resources, how would you continue or expand upon the work you have already done?

You should turn in a .pdf file containing your final report, together with the code and the generated results (tables, figures etc) that are included in the report. You will present your final report at the last class either by pdf file or slides.

The project report is 3–5 pages long, excluding figures or references.

The deadline of the final project report is strict. All late final projects receive a score of zero. In case of requesting exceptions due to severe medical reasons, a doctor's note and a signature from your graduate (or undergraduate) advisor is needed.

Additions.

Auditing. Auditing is welcome. All auditors should present at least one paper as part of interactions.

AI Credit Know-AI: Know the basic functions of AI and how to correctly use AI applications. AI course content is over 50%.

AI Student Learning Outcomes (SLOs)

SLO1. Identify, describe, and/or explain the components, requirements, and/or characteristics of AI.

SLO2. Recognize, identify, describe, define, and/or explain applications of AI in multiple domains.

Regrade Policy. You may submit a regrade request if you believe the course staff made an error in grading. Any regrade requests should be emailed to the course instructor within ten days of receiving your grade. Please try to be as specific as possible with your regrade request. For fairness, the regrade will be based on the original submission, and no additional explanations or materials will be considered. The new grade may be higher, lower, or the same as the original grade. The results of assignments and exams are finalized after ten days of delivery if no further concerns are raised.

Attendance. Since the course mainly relies on the instructor's lecturing and does not follow any textbooks, it will be important to attend the class to get the most out of this course. If you should need to be absent from a class, you need to contact the instructor before the class. The attendance policy will be consistent with the university. For more information, see <https://catalog.ufl.edu/UGRD/academic-regulations/attendance-policies/>.

Final letter grade. The final grades will be based on your relative performance in the class. The final grade will be ranked and set so that the distribution of final grades approximately matches that of previous offerings of the course. That said, your final grade is guaranteed to be the same as or better than the following assignment mechanism of letter grades: [1%, 40%]: (A, A-), (40%, 60%): B+, (60%, 90%): (B, B-); (90%, 99%): B- or lower.

Support Resources

Academic integrity. UF students are bound by The Honor Pledge which states, “We, the members of the University of Florida community, pledge to hold ourselves and our peers to the highest standards of honor and integrity by abiding by the Honor Code. On all work submitted for credit by students at the University of Florida, the following pledge is either required or implied: “On my honor, I have neither given nor received unauthorized aid in doing this assignment.” The Honor Code (<http://www.dso.ufl.edu/sccr/process/student-conduct-honorcode/>) specifies a number of behaviors that are in violation of this code and the possible sanctions. Furthermore, you are obligated to report any condition that facilitates academic misconduct to appropriate personnel. If you have any questions or concerns, please consult with the instructor in this class.

Accommodations for students with disabilities. Students with disabilities requesting accommodations should first register with the Disability Resource Center (www.dso.ufl.edu/drc/) by providing appropriate documentation. Once registered, students will receive an accommodation letter which must be presented to the instructor when requesting accommodation. Students with disabilities should follow this procedure as early as possible in the semester. For more information, see <https://disability.ufl.edu/students/get-started/>.

GatorEvals. You are expected to provide professional and respectful feedback on the quality of instruction in this course by completing course evaluations online. Guidance on how to provide constructive feedback is available at <https://gatorevals.aa.ufl.edu/students/>. Students will be notified when the evaluation period opens. Summaries of course evaluation results are available to students at <https://gatorevals.aa.ufl.edu/public-results/>.

Business Career Services : A significant aspect of your overall success as a business Gator is your journey towards securing your career upon graduation. As the only College at UF with its own dedicated career services office, Business Career Services (BCS) exists to help you explore career paths that fit your major and interests, assist you in developing job and internship search strategies, provide you support in developing an impactful resume, prepare you for interviews, and much more. For more information on identifying your dedicated BCS Career Coach and other career resources offered, please visit <https://warrington.ufl.edu/career/students/>. Interested in scheduling an appointment with your Career Coach or applying for an internship or full-time job, visit HIREWarrington (<https://warrington-ufl.12twenty.com/login>). Business Career Services is exclusive to currently enrolled, Warrington degree seeking students only. Non-Warrington students are encouraged to visit University of Florida’s Career Connection Center for career assistance.

Additional Resources.

U Matter, We Care: If you or someone you know is in distress, please contact umatter@ufl.edu, 352-392-1575, or visit [U Matter, We Care](#) website to refer or report a concern and a team member will reach out to the student in distress.

Counseling and Wellness Center: Visit the [Counseling and Wellness Center website](#) or call 352-392-1575 for information on crisis services as well as non-crisis services.

Student Health Care Center: Call 352-392-1161 for 24/7 information to help you find the care you need, or visit the [Student Health Care Center website](#).

University Police Department: Visit [UF Police Department website](#) or call 352-392-1111 (or 9-1-1 for emergencies).

UF Health Shands Emergency Room / Trauma Center: For immediate medical care call 352-733-0111 or go to the emergency room at 1515 SW Archer Road, Gainesville, FL 32608; Visit the UF Health [Emergency Room and Trauma Center website](#).

E-learning technical support: Contact the UF Computing Help Desk at 352-392-4357 or via e-mail at helpdesk@ufl.edu.

Career Connections Center: Reitz Union Suite 1300, 352-392-1601. Career assistance and counseling services.

Library Support: Various ways to receive assistance with respect to using the libraries or finding resources.

Teaching Center: Broward Hall, 352-392-2010 or to make an appointment 352- 392-6420. General study skills and tutoring.

Writing Studio: 2215 Turlington Hall, 352-846-1138. Help brainstorming, formatting, and writing papers.

Student Complaints On-Campus: [Visit the Student Honor Code and Student Conduct Code webpage for more information](#).

Acknowledgments

A part of course materials are adapted from the related courses offered by David Blei, Yixin Wang and Aaron Schein.