

University of Wisconsin – Madison  
Department of Industrial and Systems Engineering  
**ISyE 601: Machine Learning in Action**  
Fall 2019

---

<b>Instructors</b>	Professor Justin J. Boutilier, ME 3246, <a href="mailto:jboutilier@wisc.edu">jboutilier@wisc.edu</a> Office hours: Thursday 1-3pm or by appointment	
<b>Lectures</b>	Tuesday 4-515pm	ME 3121
	Thursday 4-515pm	ME 3121
<b>Course Site</b>	<a href="https://canvas.wisc.edu/courses/171130">https://canvas.wisc.edu/courses/171130</a>	
<b>Description</b>	This course introduces principles, algorithms, and applications of machine learning. We cover predictive analytics, with a focus on combining data and models to improve decision-making. Methods include: statistics, linear regression, logistic regression, regularization, over-fitting, clustering, classification and regression trees, boosting, bagging, deep learning, and neural networks. Applications areas include: healthcare, transportation, and the public sector.	
<b>References</b>	<i>The Analytics Edge</i> by Bertimas, O’Hair, and Pulleyblank, Dynamic Ideas. <i>Pattern Recognition and Machine Learning</i> by Bishop, Springer.	
<b>Prerequisite</b>	ISyE 301, ISyE 312, ISyE 323 or permission of the instructor	
<b>Credits</b>	This is a 3-credit course that meets for two, 75-minute class periods each week and carries the expectation that students will work on course learning activities (reading, labs, homework, etc.) for about 3 hours out of the classroom for every class period.	
<b>Policies</b>	Students are reminded of the importance of academic integrity and seriousness of academic dishonesty in any form, including plagiarism. Please see <a href="https://conduct.students.wisc.edu/misconduct/academic-integrity/">https://conduct.students.wisc.edu/misconduct/academic-integrity/</a> for more information.	
<b>Grading</b>	Homework	40%
	Quizzes	20%
	Final exam	40%
<b>Labs</b>	Labs, in the form of Jupyter Notebooks, will focus on developing proficiency in the analytical methods presented. Labs will use Python to implement methods using real-world data.	

University of Wisconsin – Madison  
Department of Industrial and Systems Engineering  
**ISyE 601: Machine Learning in Action**  
Fall 2019

---

- Homework** Students are encouraged to work in groups of 2 (but no more!). Homework is to be submitted online through the course website. Scans of hand-written work are accepted, but must legible; if we can't read it, we can't grade it. This course follows a strict lateness policy: *late homework will not be accepted*. The homework schedule is as follows:
- *Homework 0*: Assigned on Sept 24, due on Oct 10 before 4pm
  - *Homework 1*: Assigned on Oct 15, due on Oct 31 before 4pm
  - *Homework 2*: Assigned on Nov 5, due on Nov 21 before 4pm
  - *Homework 3*: Assigned on Nov 26, due on Dec 11 before 4pm
- Quizzes** Short (15-20min) quizzes will be administered in class throughout the semester. The quiz schedule is as follows:
- *Quiz 0 – Sept 26*: Concepts and EDA
  - *Quiz 1 – Oct 8*: Linear and logistic regression
  - *Quiz 2 – Oct 17*: Clustering and KNN
  - *Quiz 3 – Nov 5*: Cart and random forest
  - *Quiz 4 – Nov 19*: Ensembles and SVMs
- Final Exam** The final exam will be cumulative and will take place in class on December 18, 2019 from 725-925pm.
- Learning Objectives** By the end of this course, you will be able to:
- Identify problems amenable to machine learning and the techniques required to solve those problems (regression vs. classification, regularization, bagging vs. boosting, etc.)
  - Apply appropriate analytical models to solve problems and improve decision-making using Python
  - Effectively communicate findings through both oral and written formats

University of Wisconsin – Madison  
Department of Industrial and Systems Engineering  
**ISyE 601: Machine Learning in Action**  
Fall 2019

Week	Dates	Tuesday	Thursday	Lab
1	Sept 2-6	NO CLASS	The language of machine learning	Python refresher
2	Sept 9-13	Linear regression	Trade-offs and model training	EDA, feature engineering
3	Sept 16-20	Trade-offs and model training	Exploratory data analysis	
4	Sept 23-27	Feature engineering	The Framingham heart study	Logistic regression
5	Sept 30 – Oct 4	Logistic regression	Logistic regression	
6	Oct 7-11	Clustering	Clustering / K-nearest neighbors	Clustering, KNN
7	Oct 14-18	K-nearest neighbors	Supreme court predictions	
8	Oct 21-25	Guest speaker	Classification and regression trees	CART, random forest
9	Oct 28 – Nov 1	Classification and regression trees	Random forest	
10	Nov 4-8	Ensemble methods	Ensemble methods	Ensembles, SVM
11	Nov 11-15	Support vector machines	Support vector machines	
12	Nov 18-22	The history of deep learning	Neural networks	Neural networks
13	Nov 25-29	Neural networks	NO CLASS	
14	Dec 2-6	Communicating analytics – written	Communicating analytics – oral	
15	Dec 9-13	Exam review	NO CLASS	