The University Of Toledo

Existing Graduate Course Modification Form

* denotes required fields

Contact Person*: Richard Molyet Phone: 530-8143 (xxx - xxxx) Email: richard.molyet@utoledo.edu

Present	**Proposed**
Supply all information asked for in this column.(Supply core, research intensive and transfer module info if applicable)	Fill in appropriate blanks only where entry differs from first column.
College*: College of Engineering	**College**: --Select a College--
Dept/Academic Unit*: Electrical Engineering and Computer Science	**Dept/Academic Unit**: --Select a Department--
Course Alpha/Numeric*: EECS - 5750	**Course Alpha/Numeric:** -
Course Title: Machine Learning	**Course Title:**
Credit hours: Fixed: 3 or Variable: to	**Credit Hours:** Fixed: or Variable: to
CrossListings: EECS 4750	**CrossListings:**
Insert	Insert
To add a course, type in course ID and click the Insert button. To remove a course, select the course on left and click the Remove button.	To add a course, type in course ID and click the Insert button. To remove a course, select the course on left and click the Remove button.
Remove	Remove
Prerequisite(s)(if longer than 50 characters, please place it in Catalog Description): EECS 2100, MATH 2890 and MIME 4000	**Prerequisite(s)**(if longer than 50 characters, please place it in Catalog Description): EECS 2110, MATH 2890 and MIME 4000
Corequisite(s)(if longer than 50 characters, please place it in Catalog Description):	**Corequisite(s)**(if longer than 50 characters, please place it in Catalog Description):

Catalog Description (*only if changed*) 75 words max:

This course emphasizes learning algorithms and theory including concept, decision tree, neural network, computational, Bayesian, evolutionary, and reinforcement learning.

Catalog Description (*only if changed*) 75 words max:

Has course content changed? ○ Yes ◉ No

If course content is changed, give a brief topical outline of the revised course below(less than 200 words)

Proposed effective term*: 201740 (e.g. 201140 for 2011 Fall)

File Type	View File
Syllabus	View

List any course or courses to be deleted. Effective Date:

Effective Date:

Comments/Notes:

Rationale:

Approval:

Department Curriculum Authority:	Richard G. Molyet	Date	2017/03/23
Department Chairperson:	Mansoor Alam	Date	2017/03/23
College Curriculum Authority or Chair:	Efstratios Nikolaidis	Date	2017/03/31
College Dean:	Mohamed Samir Hefzy	Date	2017/04/17
Graduate Council:	Constance Schall, GC mtg 5/2/17	Date	2017/05/03
Dean of Graduate Studies:	Amanda C. Bryant-Friedrich	Date	2017/05/04
Office of the Provost :	marcia king-blandford	Date	2017/05/10

print

Administrative Use Only

Effective Date: 2016/08/22 (YYYY/MM/DD)

CIP Code:

Subsidy Taxonomy:

Program Code:

Instructional Level:

Registrar's Office Use Only

Processed in Banner on: 2017/05/17

Processed in Banner by:	Tasha Woodson
Banner Subject Code:	EECS
Banner Course Number:	5750
Banner Term Code:	201810
Banner Course Title:	Machine Learning

The University of Toledo • 2801 W. Bancroft • Toledo, OH 43606-3390 • 1.800.586.5336

Machine Learning

The University of Toledo
Department of Electrical Engineering and Computer Science
EECS 4750/5750 Section 002 (CRN: 49717/49718)

Instructor:	Dr. Kevin S. Xu	**Term**:	Fall 2016
Email:	kevin.xu@utoledo.edu	**Class Location**:	Palmer Hall 1030
Office Hours:	Mon, Tues: 2:00pm-4:30pm	**Class Day/Time**:	Mon, Wed: 12:30pm-1:45pm
Office Location:	Nitschke Hall 2055	**Credit Hours**:	3
Office Phone:	419-530-8144		

COURSE/CATALOG DESCRIPTION

This course emphasizes learning algorithms and theory including concept, decision tree, neural network, computational, Bayesian, evolutionary, and reinforcement learning.

COURSE OVERVIEW

This course provides students with an introduction to machine learning, a subfield of computer science that focuses on giving computers the ability to make decisions by learning from data rather than by following explicit rules or instructions. Machine learning algorithms are used in many application settings where programming explicit rule-based algorithms is unfeasible or insufficient, including

- Optical character recognition (OCR), including handwriting recognition.
- Filtering spam e-mails and posts on social media.
- Business analytics, such as customer revenue prediction.
- Prediction of disease from biomarkers or other patient data.

These algorithms work by building models of relationships between variables from example inputs (training data) in order to make predictions or decisions. This course will cover both theoretical and practical aspects of machine learning.

STUDENT LEARNING OUTCOMES

Upon successful completion of the course, students should be able to

- Identify a practical machine learning task as a supervised, unsupervised, or reinforcement learning task.
- Explain the underlying assumptions behind a machine learning algorithm.
- Implement and test a variety of machine learning algorithms in Python.
- Evaluate the performance of a machine learning algorithm on a data set.
- Select and apply suitable machine learning algorithms to make data-driven predictions in different application settings.
- Effectively communicate results, predictions, and implications of using a machine learning algorithm in an application setting.

Additionally, the following outcomes apply to graduate students enrolled in the MS-level course (EECS 5750):

- Solve intermediate-level problems in machine learning.
- Solve engineering problems using intermediate-level mathematics.
- Successfully communicate the results of engineering research in oral and written forms.
- Present and publish the results of high-quality engineering research projects.

TEACHING STRATEGIES

Lectures will consist of a mix of teaching on the board, presentation slides, and demonstrations of machine learning algorithms involving some computer programming. Course materials, including lecture notes and homework assignments, will be distributed using Blackboard.

PREREQUISITES AND COREQUISITES

Formal prerequisites:

- Undergraduate level MIME 4000 (Engineering Statistics I): Minimum Grade of D-
- Undergraduate level MATH 2890 (Numerical Methods and Linear Algebra): Minimum Grade of D-
- Undergraduate level EECS 2110 (Computer Architecture and Organization): Minimum Grade of D- (incorrectly listed in myUT as EECS 2100)

Students missing any formal prerequisites may request an override from me; however, students are expected to understand the topics covered in the prerequisite courses. Specifically, students are expected to understand the following topics at an undergraduate EECS level:

- *Linear algebra:* Students should be comfortable with basic matrix manipulations such as matrix multiplications and inversions and should have some familiarity with eigenvalue decompositions.
- *Probability:* Students should be comfortable with probability mass and density functions, Bayes' theorem, and commonly used probability distributions including the binomial, Poisson, and Gaussian distributions.
- *Computer programming:* Students should be comfortable with basic programming in a high-level language, including writing if statements, loops, storing and accessing data in arrays, and file input and output. We will be using Python in this course, and it should be relatively easy to learn for students with prior programming experience. However, the focus of the course is on machine learning algorithms and not on programming concepts, so students are expected to already be competent programmers.

REQUIRED TEXTS AND ANCILLARY MATERIALS

Required textbook: S. Marsland, Machine Learning: An Algorithmic Perspective (Second Edition), CRC Press, 2015. ISBN-13: 978-1466583283

This textbook provides an introduction to machine learning from an algorithmic perspective that is not too heavily obscured by mathematical details and should be suitable for both undergraduates and graduate students. Additional reading material, particularly on topics not covered in Marsland's book, will be assigned and posted on the course website on Blackboard. Other related textbooks for optional reading include

- C. Bishop, Pattern Recognition and Machine Learning, Springer, 2007. ISBN-13: 978-0387310732. This textbook presents machine learning from a Bayesian perspective and is one of the most commonly used textbooks for graduate-level introductory courses in machine learning.
- T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction (Second Edition), 2009. ISBN-13: 978-0387848570. This is a classical textbook on machine learning from a computational statistics perspective. It is available freely in PDF format at http://statweb.stanford.edu/~tibs/ElemStatLearn/.

TECHNOLOGY REQUIREMENTS

Assignments will be issued on the course website on Blackboard. Students must submit their completed assignments in Blackboard by 11:59pm on the due date for full credit. If the Blackboard website is inaccessible near the assignment deadline such that a student is unable to submit his or her completed assignment, the student may email me with the completed assignment.

I will also issue anonymous surveys throughout the course using Blackboard. Unlike with assignments, Blackboard surveys allow me to see only

- The list of students who have completed the survey.
- A set of aggregated responses from all students who completed the survey.

These surveys will be used for assessment purposes and to improve the quality of the course and future offerings.

UNIVERSITY POLICIES

Policy Statement on Non-Discrimination on the basis of Disability (ADA)

The University is an equal opportunity educational institution. Please read The University's Policy Statement on Nondiscrimination on the Basis of Disability Americans with Disability Act Compliance.)

ACADEMIC ACCOMMODATIONS

The University of Toledo is committed to providing equal access to education for all students. If you have a documented disability or you believe you have a disability and would like information regarding academic accommodations/adjustments in this course please contact the Student Disability Services Office.

ACADEMIC POLICIES

Students are expected to comply with all of the university's Undergraduate Academic Policies (EECS 4750 students) or Graduate Academic Policies (EECS 5750 students) in addition to the following course policies.

Missed Class Policy

A student who is unable to submit an assignment by the due date or take an exam on the scheduled date should notify me in advance by sending an e-mail or providing a written statement in person. In the event of an emergency, the student must present me with approved written documentation upon the student's return to class in compliance with the university's Missed Class Policy. Any other absences will be treated as *unexcused absences*, which will result in a grade of zero issued for the particular assignment or exam and may result in an Incomplete grade. Students do not need to notify me for missed lectures; however, they are responsible for the material covered in any missed lecture.

If a student is granted an excused absence for a homework assignment, the student will be given a default grade for the assignment equal to the average of the other homework assignments. In the case of an excused absence for a project, the student will be given an extended due date to complete the project commensurate with the length of the excused absence. If a student is granted an excused absence for an exam, the student will be given the opportunity to make up the exam at a later date. In this case, please contact me as soon as possible to make alternate exam arrangements.

Academic Dishonesty Policy

All work turned in must be your own individual work. Students may share general ideas but not specific approaches to solving a particular problem. In particular, you must write your own solutions and code. You must be prepared to prove that your solutions and code are your own, by explaining it to me or with written evidence. ***Unless otherwise specified*** (e.g. in a particular assignment problem), students are allowed to use code from the textbook as well as publicly available Python packages for machine learning, such as scikit-learn, to assist them in completing assignments ***provided that such usage is documented in the code turned in by the student***. However, students are ***not*** allowed to use code from discussion forums and Q & A websites such as StackExchange. In particular, gaining advantage from posting a homework problem to such a venue constitutes academic dishonesty.

A student found to be academically dishonest will be sanctioned according to the university's Academic Dishonesty Policy. ***A first violation will result in a grade of zero issued for the homework assignment, project, or exam in question. A second violation will result in a grade of F issued for the course.***

GRADING

Undergraduate (EECS 4750) students' numeric grades will be calculated according to the following weighting:

- Homework assignments: 25% (5 total, 5% each)
- Projects: 25% total (2 total, 12.5% each)
- Mid-term exams: 50% (2 total, 25% each)
- Class participation: Up to 5% extra credit

Graduate (EECS 5750) students' numeric grades will be calculated according to the following weighting:

- Homework assignments: 20% (5 total, 4% each)
- Projects: 20% (2 total, 10% each)
- Mid-term exams: 40% (2 total, 20% each)
- Presentation of advanced topic: 20%
- Class participation: Up to 5% extra credit

Unless otherwise noted in the course schedule, homework assignments will be issued on Wednesdays and due the following Tuesday at 11:59pm Eastern. Late assignments, either homework or projects, will be accepted ***up to 24 hours after the assignment deadline with a 25% penalty***.

Homework Assignments

Homework assignments will consist of a mix of conceptual, mathematical, and programming problems related to the material covered in lectures. All homework assignments must be submitted on Blackboard; ***it is the responsibility of the student to ensure that any materials completed on paper are scanned and legible!***

For each homework assignment, ***one randomly selected problem will be carefully graded for correctness***; this will account for 50% of each homework grade. The other problems will be briefly graded for completion and relevance and will account for the other 50% of each homework grade.

Projects

Unlike homework assignments, projects will take on the form of open-ended case studies where you are given a particular task and one or more data sets, and your goal is to build a machine learning system to accomplish the stated task and write a report summarizing your results, predictions, and implications. You are free to use any machine learning algorithms you would like. Projects will be evaluated based on the quality of results including accuracy and complexity of the machine learning system, as well as effective communication of the results.

Midterm Exams

There will be two 75-minute written midterm exams taken during the usual lecture time. Since these will be written exams, the focus will be on conceptual understanding of machine learning algorithms. A dedicated scientific calculator is allowed, but no smartphone or other multi-function device with calculator is allowed without my approval.

Presentation of Advanced Topic (EECS 5750 only)

Graduate students should form groups of 2 (3 only if necessary) to investigate and give a 1-hour tutorial-style presentation on an advanced topic in machine learning. I will provide a list of recommended topics to choose from. Some of the topics are covered briefly in the textbook, while others may not be. Students are expected to read related papers to get a high-level understanding of the topic, then present the topic in a clear and concise manner to the rest of the class. Presentations will be graded on correctness, completeness, and clarity.

Class Participation

A student may receive up to 5% extra credit for participation in lectures and timely completion of surveys.

Letter Grades

Letter grades will be assigned in the following manner. I identify clusters in the numeric grades and use gaps between clusters as thresholds between letter grades. Each student with a numeric grade in a particular cluster will be assigned the same letter grade. This procedure ensures that no student is only a few points away from getting a higher letter grade. Students in EECS 4750 and EECS 5750 will be clustered separately due to the different weights for numeric grades.

Midterm Grading

After the first midterm exam, students will receive a midterm numeric and letter grade calculated using the procedure described above. This letter grade is purely for the student to gauge his or her progress and does not affect the final letter grade.

Final Grading

Students' final letter grades will be computed using the above procedure and weightings. Students who have not completed both of the midterm exams and projects may be assigned a grade of Incomplete, unless prior arrangements have been made.

COMMUNICATION GUIDELINES

I can be reached in my office (NI 2055) during scheduled office hours or via email. I will make my best effort to respond to email inquiries within 48 hours. ***Please place "4750" or "5750" in the subject line of the email to expedite this process.*** If my scheduled office hours do not work for you, please email me to make an appointment for an alternate time.

Students are expected to check their UT email address regularly. If there are any issues with a submitted assignment, for example, I may send you a time-sensitive email allowing you to correct your assignment prior to solutions being uploaded. Announcements to the class will be posted on the course website on Blackboard.

COURSE SCHEDULE

The course schedule is tentative and subject to change. Changes will be announced in lecture and posted to the course website on Blackboard. Assignment due dates and exam dates are shown in **bold**. Additional reading material, including survey papers and papers on advanced topics in machine learning, will be made available on Blackboard. There will be no final exam.

Week	Important Dates	Content	Chapters
1 (8/22, 8/24)	Homework 1 assigned 8/24	Introduction to machine learning Introduction to Python and NumPy	1, Appendix A
2 (8/29, 8/31)		Review of probability and statistics Naïve Bayes' classifier Evaluation metrics	2
3 (9/7)	No lecture or office hours 9/5 **Homework 1 due 9/6** Homework 2 assigned 9/7	Introduction to neural networks The Perceptron	3
4 (9/12, 9/14)		Linear regression Logistic regression	3
5 (9/19, 9/21)	**Homework 2 due 9/20** Homework 3 assigned 9/21	Multi-layer perceptron	4
6 (9/26, 9/28)		Nearest neighbor methods Support vector machines (SVMs)	7, 8
7 (10/5)	No lecture or office hours 10/3 and 10/4 **Homework 3 due 10/5** Project 1 issued 10/6	Decision trees	12
8 (10/10, 10/12)	**Midterm 1 10/10**	Linear discriminant analysis (LDA) Principal components analysis (PCA)	6
9 (10/17, 10/19)		Optimization and search algorithms	9
10 (10/24, 10/26)	**Group presentations 10/24 and 10/26** **Project 1 due 10/25** Homework 4 issued 10/26	Topics in dimensionality reduction, classification, and regression	5, 6, 13
11 (10/31, 11/2)		K-means and hierarchical clustering Gaussian mixture models	7, 14

12 (11/7, 11/9)	**Homework 4 due 11/8** Homework 5 issued 11/9	Kernel density estimation Bayesian networks	16
13 (11/14, 11/16)		Hidden Markov models Deep learning	16, 17
14 (11/21)	**Homework 5 due 11/22** Project 2 issued 11/22 No lecture 11/23	Reinforcement learning	11
15 (11/28, 11/30)	**Midterm 2 11/28** **Group presentations 11/30**	Topics in unsupervised learning, reinforcement learning, evolutionary learning, and graphical models	10, 11, 14, 16
16 (12/5, 12/7)	**Group presentations 12/5 and 12/7** **Project 2 due 12/9**	Topics in unsupervised learning, reinforcement learning, evolutionary learning, and graphical models	10, 11, 14, 16