Course Change Request

New Course Proposal

Date Submitted: 10/09/18 2:36 pm

Viewing: **CDS 403 : Applications of Machine Learning**

Last edit: 10/09/18 2:36 pm

Changes proposed by: blaisten

Are you completing this form on someone else's behalf?

Yes

Requestor:

Name	Extension	Email
Estela Blaisten	31988	blaisten@gmu.edu

Effective Term: Fall 2019

Subject Code: CDS - Computational and Data Sciences **Course Number:** 403

Bundled Courses:

Equivalent Courses:

Catalog Title: Applications of Machine Learning

Banner Title: Applications Machine Learning

Will section titles vary by semester? No

Credits: 3

Schedule Type: Lecture

Hours of Lecture or Seminar per week: 3

Repeatable: May only be taken once for credit (NR)

Default Grade Mode: Undergraduate Regular

Recommended Prerequisite(s):
CDS 230 - Modeling and Simulation or sufficient computing skills
MATH 203 - Linear Algebra
CDS 303 - Scientific Data Mining
Or permission of instructor

Recommended Corequisite(s):

Required Prerequisite(s) / Corequisite(s) (Updates only):

Registrar's Office Use Only - Required Prerequisite(s)/Corequisite(s):

And/Or	(	Course/Test Code	Min Grade/Score	Academic Level	)	Concurrency?

Registration Restrictions (Updates only):

Registrar's Office Use Only - Registration Restrictions:

Field(s) of Study:

Class(es):

Level(s):

Degree(s):

School(s):

In Workflow

1. **CDS Chair**
2. **SC Curriculum Committee**
3. SC Associate Dean
4. Assoc Provost-Undergraduate
5. Registrar-Courses
6. Banner

Approval Path

1. 08/27/18 1:19 pm Jason Kinser (jkinser): Approved for CDS Chair
2. 09/18/18 4:23 pm Gregory Craft (gcraft): Rollback to Initiator
3. 10/09/18 3:13 pm Jason Kinser (jkinser): Approved for CDS Chair

Catalog Description:	Covers practical applications in STEM areas of decision trees, rule-based classification, support vector machines, Bayesian networks, ensemble methods, and Neural Networks. Emphasis resides on the process of applying machine learning effectively to a variety of problems.
Justification:	There is currently no undergraduate machine learning course offered by the Department of Computational and Data Sciences (CDS) at Mason. Yet machine learning is a critical skill not only demanded by the marketplace but also expected of data science and computational science graduates. This course will focus on applying machine learning algorithms to applications relevant in the STEM fields. Machine Learning is concerned with the design and development of programs that automatically improve their performance through experience. The class is meant to teach the practical side of machine learning for applications, such as mining science and engineering data or building adaptive user interfaces. The emphasis will be on learning the process of applying machine learning effectively to a variety of problems rather than emphasizing on the theory behind what makes machine learning work.
Does this course cover material which crosses into another department?	No
Learning Outcomes:	1) The class is meant to teach the practical side of machine learning for applications, such as mining science and engineering data or building adaptive user interfaces. 2) Class emphasis will be on learning the process of applying machine learning effectively to a variety of problems rather than emphasizing on the theoretical aspects.
Attach Syllabus	CDS403_AppliedMachineLearning.pdf
Additional Attachments	
Staffing:	Michael Eagle James Glasbrenner Joseph Marr Jason Kinser
Relationship to Existing Programs:	This course complements the existing CDS BS by providing students with a market-demanded skill. This course fills a critical gap in the existing CDS BS curriculum at the upper-level undergraduate.
Relationship to Existing Courses:	This course does not duplicate any existing undergraduate course. The course introduces new paradigms and new technologies that are applicable in academia as well as the marketplace. The course may share certain components with other courses in the spirit of knowledge transference.

Additional Comments:	
Reviewer Comments	**Gregory Craft (gcraft) (09/18/18 4:23 pm):** Rollback: Per email from Karen that they needed to make adjustments.

Key: 15953

CDS 403 Applications of Machine Learning

SYLLABUS

Monday/Wednesday, TBD, Exploratory Hall,

Fall 2019

INSTRUCTOR: Dr. Michael J. Eagle

OFFICE: Research Hall Room

220

OFFICE HOUR: Tuesdays 10:00 am-12:00 noon

TEXTBOOK:

- *Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems* (Aurélien Géron, O'Reilly Media, 2017)

RELEVANT ADDITIONAL TEXTBOOKS:

- Machine Learning, Tom Mitchell, McGraw Hill, 1997.
- Machine Learning: A Probabilistic Perspective, by Kevin P. Murphy, 2012.
- *Deep Learning* (Ian Goodfellow, Yoshua Bengio, Aaron Courville and Francis Back, 2016.)

Course Description:
Machine Learning (ML) asks "how can we design programs that automatically improve their performance through experience?" Machine Learning is concerned with computer programs that enable the behavior of a computer to be learned from examples or experience rather than dictated through rules written by hand. We will cover a wide range of ML algorithms that can be applied to a variety of problems. In particular, we will cover topics such as decision trees, rule-based classification, support vector machines, Bayesian networks, ensemble methods, and Neural Networks. This class is meant to teach the practical side of machine learning for applications, such as mining scientific databases or building adaptive user interfaces. The emphasis will be on learning the process of applying machine learning effectively to a variety of problems rather than emphasizing an understanding of the theory behind what makes machine learning work.

Grading:

Grades will be based on assignments and quizzes, 2 take-home midterms, and a course project.

The term project will involve applying machine learning to a substantial problem of the student's choice. Several options are found in the Projects subfolder in Blackboard, which include a good collection of datasets. Students may select one of these datasets or may propose one of their own design. Students who wish to design their own project should check in about their plans with the instructor early in the semester.

Grading Criteria

Quizzes (10%)
Assignments (20% total)
Mid-terms (10% each)
Course project (50%)

Week	*Topic*
1	Course goals, philosophy, teaching style, policies and mechanics
2	Introduction: The Machine Learning Process
3	Three Basic Algorithms: regression, k-NN, and k-means
4	Input, Output, and data munging
5	Optimization
6	Classification
7	Advanced Tree and Rule Based Learning
8	Working with Text / NLP
9	Ensemble Learning and Random Forests
10	Feature Selection and Optimization
11	Neural Networks
12	Deep Learning
13	Graphical Models
14	Reinforcement Learning
15	Review
	Wrap-up and final project presentations in mini-symposium format.