A. JAMES CLARK
SCHOOL OF ENGINEERING

Course: ENPM 808F Robot Learning
Instructor: Donald A Sofge
To be offered Fall 2016

Synopsis:

Machine learning may be used to greatly expand the capabilities of robotic systems, and has been applied to a variety of robotic system functions including planning, control, and perception. Machine Learning for Robotic Systems covers the application of a machine learning techniques for which data is used to generate (through induction) a model that is then used by the robot to perform tasks. A wide variety of representations and techniques are available to generate models including multilayer perceptrons (e.g., trained through backpropagation or evolutionary algorithms), Radial Basis Functions, Sparse Representations, Support Vector Machines, Random Decision Forests, Bayesian Networks, and Deep Networks (Convolutional Neural Networks). Ultimately we would like for the robots to expand their knowledge and improve their own performance through learning while operating in the environment (on-line and/or lifelong learning). This graduate course will explore the application of machine learning techniques to robotic systems, focusing primarily on key useful representations and model building techniques for application in non-stationary robotic systems. More attention will be paid to machine learning for robot control than for perception.

Topics:

- Motor Learning and the Evolution of Intelligence
- Behavior Based Robotics
- Robot Shaping and Evolving Behaviors
- Crossing the Reality Gap
- Reinforcement Learning
- Value versus Policy Iteration
- Q-Learning and Actor-Critic Models
- Memory-Based Learning
- Imitation Learning and Learning from Demonstration
- Deep Reinforcement Learning with CNNs
- On-line and Lifelong Learning

References:

- Tom Mitchell. Machine Learning. Cambridge McGraw Hill. 1997.
- White, David A., and Donald A. Sofge, eds. Handbook of Intelligent Control: Neural, Fuzzy, and Adaptative Approaches. Van Nostrand Reinhold Company, 1992.

- Hal Daume III, A Course in Machine Learning, September 2015 (on-line)
- Nikolaus Correll, Introduction to Autonomous Robots, 1st Edition, September 2015 (on-line)
- David Barber, Bayesian Reasoning and Machine Learning, Cambridge University Press, 2012 (on-line)
- Other weekly reading materials to be assigned in class

Homework, Projects, and Examinations:

Course grade will be based on students' performance in the following items:

- Homework (20% of the overall grade)
- Quizzes (10% of the overall grade)
- Course Project Presentation (20% of the overall grade)
- Course Project Final Report (20% of the overall grade)
- Take Home Final Exam (30% of the overall grade)

Code of Academic Integrity:

The University of Maryland, College Park has a nationally recognized Code of Academic Integrity, administered by the Student Honor Council. This Code sets standards for academic integrity at Maryland for all undergraduate and graduate students. As a student you are responsible for upholding these standards for this course. It is very important for you to be aware of the consequences of cheating, fabrication, facilitation, and plagiarism. For more information on the Code of Academic Integrity or the Student Honor Council, please visit http://www.studenthonorcouncil.umd.edu/whatis.html.