

Elective Sequence: Artificial Intelligence and Data Science

1. Rationale.

Computer programs that make decisions or aid in decision making, that recognize or aid in understanding patterns in data, or that learn from data are becoming more prevalent and important in the modern world. Moreover, the U.S. Navy and U.S. Government recognize the future importance of such systems (see for example the report "Rise of the Machines: Artificial Intelligence and its Growing Impact on U.S. Policy," from the Congressional Subcommittee on Information Technology, September 2018). Therefore, the Computer Science department would like to offer students an opportunity for increased depth of study in key fields associated with these activities: Artificial Intelligence (AI), Machine Learning (ML), and Data Science (DS).

The Computer Science department has considerable faculty expertise in these areas and a number of courses that address them or address key enabling technologies (e.g. High Performance Computing). The sequence will contain one new course ("Machine Learning and Data Science") which derives in part from recent independent study and research courses.

2. Required Courses.

a. SI420 Artificial Intelligence. This is an existing course that has run regularly for many years. Recently, it has been a Fall semester course. As part of establishing this sequence, the department will start offering it as a Spring semester course instead, appropriate timing for a first elective for second-class midshipmen. As a survey course on the whole field, AI has always contained a unit on Bayesian inference, another on machine learning, and another on graph search. These units would be codified as a requirement for the course as building blocks for sequence electives which will build off these concepts.

3. Sequence Electives. None of the sequence electives need be assigned to particular semesters. Instead they should be offered in response to student and faculty interest. At least two should be offered each year.

a. Existing courses. The following courses all count as sequence electives:

1) SI425 Natural Language Processing (NLP). NLP will remain as a standalone elective with no sequence prerequisite.

2) SI475 Intelligent Robotics. Robotics will change to use SI420 as a prerequisite, taking advantage of the Bayesian inference and graph search concepts taught in that course.

3) SI458 High Performance Computing (HPC). HPC will remain as a stand-alone elective, with no sequence prerequisite. While the same concepts will be taught (MPI, OpenMP, CUDA, distributed file systems, etc.), projects should be targeted at Big Data applications.

NLP and HPC are standard courses in Data Science Masters programs. Intelligent Robotics applies AI and ML to solve problems in robotics.

b. New course. The following course will count as a sequence elective:

1) SI424 Machine Learning and Data Science. Machine Learning and Data Science will be a new course with SI420 as a prerequisite. It will build on the Bayesian inference and machine learning introduction in that course to provide deep coverage of these concepts.

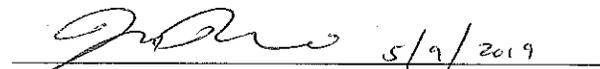
c. Criteria for evaluating future courses.

A course can count as the sequence elective if approved by the department Curriculum Committee. To count, a course should cover material not already covered in the required sequence class (AI) and it should include material with strong connections to AI, ML, or DS. Examples: a research course that applies and adapts existing machine learning algorithms to a novel problem; an experimental elective in graphics with an emphasis on data visualization; or an independent study on planning in uncertain environments. For a course from another department to count, it must not substantially duplicate material from other courses in the CS or IT curricula. For example, a math course in scientific programming that devotes substantial time to teaching programming would not be eligible.

Curriculum Committee Approval:

 5/9/2019

Department Chair Approval:

 5/9/2019