NORTHERN ARIZONA UNIVERSITY
College of Engineering, Forestry & Natural Sciences
School of Informatics, Computing, and Cyber Systems

Course Syllabus

CS 570 – Artificial Intelligence		**Spring 2017**
Credits: 3 credits	**Pre-reqs:** Instructor consent	**Co-Reqs:** N/A
Section#: 1	**Co-convened with:** CS470	**Cross-listed with:** N/A

Academic Catalog Description:

Introduces fundamental principles of artificial intelligence, as well as exploration of current research themes and challenges in areas including knowledge representation, planning, game playing, learning, and adaptive algorithms.

Course Purpose:

This course is a CS graduate elective, and provides an introduction to the specialized area of Artificial Intelligence, augmented by deeper exploration of current research directions and challenges across the field. It is therefore suitable for graduate students interested in pursuing MS or Ph.D. projects that touch on or integrate AI techniques, or who are interested in research and development work in intelligent or self-adapting/learning systems.

Expected Student Learning Outcomes:

Upon successful completion of this course, students will be able to:

- Outline the main areas of past and present activity within the broad domain of AI, including expert systems, game playing, machine learning, theorem-proving, and natural language processing.
- Plan and organize exploration of a specialty research area by effective survey reading, development of focused in-depth reading lists, and small prototype exploration.
- Effectively read and analyze research literature to articulate the state-of-the-art in key areas, identify emerging trends, and identify key challenges that define the AI research agenda within the next decade.
- Trace the historical trajectory of AI development across key areas, associate major milestones with related research groups and projects, and relate this trajectory to current efforts, i.e., active research efforts, projects, and researchers on the cutting edge.
- Describe main algorithmic approaches used in computational implementation of AI systems, and demonstrate their mastery in programming projects motivated by current research themes.

- Effectively communicate research and foster learning in others by preparing, presenting, and leading class discussions on a chosen specialty research topic, including demonstration of software developed.
- Communicate research analysis and results in effective written formats matching those used in current research forums (proceedings, journals) in computer science.
- Effectively integrate and apply the above-mentioned research skillsets to analyze current research in one active AI research specialty area in-depth, synthesize a novel project around this topic, and to produce a presentation, formal survey paper, and demonstration project to report results.

Detailed Information for this offering

Time and Location: **11:30-12:20 MWF, Rm 321 Engineering (Bldg 69)**

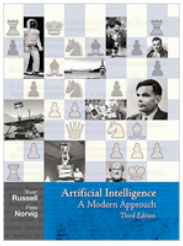

Readings and Materials:

Course Textbook: *Artificial Intelligence: A Modern Approach*, Author: Stuart Russell and Peter Norvig. Publisher: Pearson

Instructor's Name: **Dr. Eck Doerry**
Office Building/Room Number: **SICCS Building (bldg. 90), Rm 217**
Email: **Eck.Doerry@nau.edu**
Instructor Availability:

Office hours:	**Mon, Fri 12:30 – 1:30 (in my lab, Room 104a, Engineering Building)** **Tues 2:00-3:00 (in my office, SICCS Rm 217)** We often change office hours by class vote during the term to accommodate your schedules better. Check the course website for current information!
Other:	• Although you should try hard to make it to scheduled office hours, I am also available at other times by appointment. To schedule, see me before or after class or send email.

Course Structure and Evaluation Mechanisms:

This course offering will consist of a number of elements: Lectures that meet three hours per week to present and discuss core topics relating to Artificial Intelligence; readings from the textbook, augmented with both classic and recent seminal research papers; and an extended research project spanning the term. In addition, various smaller homeworks will allow you to demonstrate mastery of AI algorithms and techniques we discuss, as will a number of smaller programming challenges along the way.

As graduate students, developing strong proficiency in reading, analyzing, and communicating about research literature in AI (or another area!) is absolutely critical. To develop these skills, readings from a set of seminal papers (see Appendix) will be assigned. A lead presenter will be assigned for each, with an assignment to produce a short "summary and analysis" paper and lead a seminar discussion for the class on the topic. In addition, graduate students will need to develop a reading list of 10-15 papers in their chosen specialty area, approved by the instructor, to serve as the basis for their independent research project.

Evaluation Mechanisms: There are several mechanisms by which your course grade will be determined. Your semester research project, which includes a number of deliverables spanning the term,

constitutes the main element in your course grade. In addition, small smaller homework assignments and programming projects allow you to practice and develop AI concepts and techniques that we discuss in lectures. Exams allow you to explore the extent to which you understand the core material presented in lectures, and complement the final project presentation, which will prove your mastery of research analysis, synthesis of your own idea in AI systems, implementation of realistic AI systems, and effective communication of research results.

Individual research project: you will research a chosen specialty area/problem within AI in substantial depth, and develop a novel AI project that embodies your findings. There will be various deliverables throughout the term to mark your progress, culminating in in a submission-ready technical report and class presentation.

Homeworks: Homework will be assigned as needed throughout the semester to develop your understanding of key AI techniques or concepts explored in lecture. All homework must be done individually.

Programming: Three to four small programs will be assigned to more deeply explore some of the concepts discussed in class. Programming assignments must be done individually.

Exams: There will be two midterm exams during the semester.

Class participation: Interaction -- both with the instructor and with other students -- is crucial for understanding and integrating the ideas presented in lecture. To emphasize this, a significant number of points are set aside to reward students who take an active role in the class. Participation points are also used to credit quizzes (in-class or on BBlearn) and other formative exercises.

Important note on graduate student evaluation: There may be some semesters in which this course is co-convened with an undergraduate section to efficiently take advantage of the fact that much of the lecture material on historical developments, core AI techniques, and areas of current development can be shared. In such offerings, homeworks and programming assignments may often contain additional features or complexity targeted for graduate students. In any case, more demanding rubrics appropriate for graduate level work will be applied to graduate students where graduate and undergraduate assignments overlap.

Grading System:

Weighting of Deliverables: The following percentages are used in weighting total points earned on programming, exams, and participation: • Programming Assignments and homeworks=25% • Two midterms = 25% • Individual Research project: 40% total ✓ Literature review and analysis: 5% ✓ Project proposal and intro presentation: 5% ✓ Draft paper: 5% ✓ Final research presentation: 10% ✓ Final paper: 15% • Class participation/Quizzes = 10%	**Grading Scale**: 90-100% = A 80-89% = B 67-79% = C 55-66% = D under 55% = F
Notes: • You're a grad student now! "Mostly correct" solutions will generally get you a "C". To get an "A" or a "B" score, you'll generally have to show "above average" or "excellent" mastery, which in this course means well-developed research reading lists, coherent and insightful analysis, clear written and oral communication of research findings, and feature-rich, elegant implementations.	

Class Outline or Tentative Schedule:

Weeks	Topics	Readings	Notes
1-2	Introduction, Intelligent Agents, History	Ch.1 & 2	
2-3	Problems solving agents: state space search	Ch. 3 & 4	Choose project topic
4-5	Adversarial Search, Game Playing, Constraint probs	Ch.5, 6	Reading list due
	Exam 1 (exact date TBD)		
6-8	Logical Agents & reasoning (K-rep, 1st - order logics)	Ch. 7, 8, 9	Lit. Review, schedule project presentations
9-11	Planning Agents: Basic and Hierarchical planning	Ch. 10, 11 (through 11.2)	
	Exam 2 (exact data TBD)		
12-13	Reasoning under uncertainty	Ch. 13, 14	
14	Learning	Ch. 18	Demo final projects
15-16	Final Projects and Presentations		Final paper due!

Class Policies:

Attendance: You are responsible for all material covered during the lectures whether you attend or not. If you must miss a class, be sure to get the notes from another student. Late arrivals are very disruptive and are not acceptable --- plan to arrive five minutes before the start of class.

Electronic Device usage: All cell phones, PDAs, music players and other electronic devices must be turned off (or in silent mode) during lecture, and may not be used at any time. Laptops are allowed for note-taking only during lectures; no surfing or other use is allowed. I devote 100% of my attention to providing a high quality lecture; please respect this by devoting 100% of your attention to listening and participating.

Late work and Make-ups: Unless otherwise noted, all assigned work is due at the beginning of class on the date they are due. The following specific policies apply:

- **Quizzes:** No make-ups, no late work accepted. These are mostly on BBlearn
- **Exams:** Make-ups only when scheduled/approved in advance or with proper documentation (note from physician, etc.), as required by NAU policy.
- **Programming assignments and written Homeworks:** Late work is accepted, with appropriate deductions as specified in the Late Submissions Policy on the course website.

Grade Challenges: Although I try hard to grade as accurately and fairly as I can, mistakes do occur. If you feel that I owe you some points, or would like to discuss an evaluation, I encourage you to stop by office hours. To avoid loss of context, any grade disputes must be brought to my attention no later than 5 business days after the assignment was returned.

Homework Submission and Format: All homework where a packet of deliverables is required must:

- Have a cover page containing only the Your name and NAU ID number, the class name, the homework name, and the date of submission.
- Have problems and solutions presented in the order deliverables are listed in the assignment spec. All work must be typewritten, with the exception of figures and diagrams which could be hand-drawn, scanned, and included. Clear formatting and presentation counts; illegible work can not be graded.
- Hardcopy submissions must be stapled. No other fastening mechanisms (paper clips, dog-ears, etc.) are acceptable.

Academic Dishonesty: Cheating will not be tolerated and may result in serious sanctions, including immediate failure in the course. Serious incidents of academic dishonesty will also for brought to the attention of the

university and may result in expulsion. All work in this class is meant to be an individual effort by the person receiving the grade. Any variation from this is considered cheating and all parties involved (giving or receiving) will be sanctioned.

Other Important Course Information:

Vital Skills and other Hints for success in CS570, Artificial Intelligence

Be active, not passive.

One of the skills you should develop is an ability to read difficult material on your own - you will exercise this skill in the present course. Artificial Intelligence is a very broad ocean, with many fascinating sub-areas to explore. We will necessarily only have time to sample some subset of the chapters in the book…and will be moving along briskly to do that. Do NOT expect that I will lecture on every detail contained in the chapters; you **must** read the chapters before the lecture, and then come back later to deepen and refine the lecture discussion with material from the chapters. If you find this material as fascinating as I do, reading the chapters should be a knowledge-rich pleasure. In any case, you are responsible for material in any chapters (or sections of chapters) assigned in the course schedule.

Strive for elegance.

Those who have taken my courses before know that I'm a big fan of elegant solutions. This is especially true for programming challenges…and triply true for programming solutions in Artificial Intelligence. It is not a stretch to say that the way you factor up a programmed solution is a directed reflection of how well you really understand the concept you are trying to embody in your program. Take some extra time up front to really think it through so that you truly get what "tasks" your program will need to do…which you can then represent in your code (functions, objects, etc.). Careful thinking through of the problem up-front is the difference between three hours of programming where you just type in what you've planned, test it, and fix a few small things; and a marathon hacking session where you start off more or less blind, hit a roadblock, kludge around it, then repeat that ad infinitum until your code is an ungodly mess that you yourself can barely understand!

Don't procrastinate.

Don't delay working on homeworks and programming until the last minute. We are way way past the days of CS126, where you could blow off the assignment until the night before it's due, then sit down and crank it out in a couple of hours. The programming problems in this course are designed to take student learners between 6-20 hours of work…depending on how carefully you plan and analyze the problem up front, versus hacking away at it. Leave time for exploring, failing, thinking, and learning. Unlike lower division courses, where the goal is always to simply hack up a solution to a narrow complete spec, the programming in this course is truly part of the thinking and learning process. And that just takes time…

Take your research seriously…make it a continuous ramp of learning.

Unlike that freshman English paper, you simply can't grasp and digest an entire research area in one or two days of intensive "cramming". You'll need time to sift through find the truly *relevant* sources to your problem, to read and digest those entirely, and then to synthesize your own broad understanding of the state of the art in that area. This synthesis piece is critical for demonstrating both broad and deep understanding…and it only develops over time as you read, think about, and discuss your progress with your instructor.

University Policies:

This course is conducted in accordance with all applicable university policies which can be found at: http://nau.edu/OCLDAA/_Forms/UCC/SyllabusPolicyStmts2-2014/ including Safe Environment, Students with Disabilities, Academic Contact Hours, Academic Integrity, Research Integrity, Sensitive Course Materials, and Classroom Disruption Policy. The student handbooks are also valuable resources for other policies. The undergraduate student handbook is at http://nau.edu/student-life/student-handbook/.

Appendix: Course Reading List

Seminal research papers spanning the development of the AI area of computer science. Additional papers may be added/substituted by the instructor as the term progresses; students are encouraged to propose (and defend!) ideas for additions as well:

- M. Turing (1950) Computing Machinery and Intelligence.
- Dietterich, T. G. (2003). Machine Learning.
- Kenneth D. Forbus. 1989. Qualitative physics: past present and future.
- McDermott, Drew (1992). Robot Planning
- Bruce G. Buchanan and Reid G. Smith. 1988. Fundamentals of expert systems.(In Annual review of computer science: vol. 3, 1988)
- Judea Pearl. 1996. Decision making under uncertainty.
- Douglas B. Lenat, et al. 1990. Cyc: toward programs with common sense.
- Brin, S. and Page, L. (1998) The Anatomy of a Large-Scale Hypertextual Web Search Engine.
- G. Adomavicius; A. Tuzhilin (2005). Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions
- Itamar Arel et al. (2010). Deep Machine Learning—A New Frontier in Artificial Intelligence Research
- Mark Harman. 2012. The role of artificial intelligence in software engineering.
- **Others, as updated at each offering**