Artificial Intelligence
CS 300 - Fall 2015

Instructor: George Thomas **Office:** Halsey 218
Email: thomasg@uwosh.edu **Phone:** 424-2069

Office Hours:

M	9:00-10:00	
T	1:00-2:00	
W	10:00-1:00	or **by appointment**.

Lectures: TR 3:00-4:30 Halsey 367

Prerequisites: A grade of C or better in CS271

Textbook: No required textbook

Suggested Reference:
A good encyclopedic reference if you do more with AI after this course is *Artificial Intelligence, A Modern Approach* by Stuart Russell & Peter Norvig, 3rd Edition, Prentice Hall.

Course Website: UWO D2L

Note: If you have special needs, please come and talk to me at the end of the first class.

Course Description:
This course is an introduction to the field of artificial intelligence. A survey of classical search in artificial intelligence and machine learning, and an in-depth examination of a specific application area such as robotics, theorem proving, computer vision, natural language processing, etc. are covered. Students are expected to demonstrate mastery via computer programs using the techniques of artificial intelligence.

Course Outcomes:
1. With respect to the philosophical foundations of Artificial Intelligence (AI):
   a) describe the different characterizations of AI
   b) identify how those characterizations apply to various stages in the historic development of AI as discipline
   c) explain how the PAGE model in stimulus-response agents (Percepts/Actions/Goals/Environments) is used to characterize different types of stimulus-response agents
   d) determine rational actions for an stimulus-response agent based on maximizing goals given a set of percepts

2. With respect to classical search:
   a) formulate certain AI problems as state-space search

b) identify, manually trace and contrast uninformed search methods such as depth-first, breadth-first, uniform-cost and iterative deepening
c) explain the concept of heuristics and manually trace informed search methods such as A-star
d) identify, manually trace and contrast local search methods such as gradient descent, simulated annealing and genetic algorithms
e) formulate constraint satisfaction problems and apply consistency algorithms such as AC-3 to solve them
f) identify and manually trace adversarial search algorithms such as minimax and alpha-beta pruning
g) design heuristic evaluators for a given adversarial search space

3. With respect to learning:
   a) identify and explain general concepts of learning such as function fitting, attributes and classes, bias and over fitting
   b) describe and manually trace the perceptron learning algorithm for perceptrons with a limited number of inputs and outputs
   c) explain the role of an activation function in hidden layer networks
   d) explain how the back propagation training algorithm works in hidden layer networks
   e) represent inductive concept learning via decision trees
   f) apply the information gain method in choosing the "best" attribute of a node and manually trace the construction of a decision tree via the ID-3 algorithm
   g) construct a Bayesian network to model a problem with probabilistic information
   h) represent and answer queries in uncertain circumstances in terms of their conditional probabilities
   i) explain and manually trace the Naïve-Bayes algorithm

4. With respect to ONE selected "big application" area of AI such as Reasoning and Theorem Proving, Natural Language Processing, Computer Vision, Robotics, Multiagent Systems, etc:
   a) identify and describe the motivation, terminology, foundations and key concepts in the selected area
   b) describe and manually trace the essential algorithms in the selected area
   c) use available libraries and software to test the key algorithms and extend them to specific application projects in the selected area

Attendance and Participation: You are expected to not only attend **every** class meeting but also to come **prepared** for and **participate** actively in it. I **strongly encourage you to ask any question** or raise any issue you have with the course either during or at the end of class, or during my office hours. I will also gladly meet with you by appointment.

Course Grading Policy: Your final grade for this course will be based on three components, namely exams, assignments, and review problems. Your overall numerical

grade for the course will be computed as the weighted sum of the component grades using the following weights:

Component	Weight
Exams (3)	50%
Assignments	32%
Review Problems	18%

Tentative Exam Dates are as follows:

- **Exam 1 – Thursday, 10/01**
- **Exam 2 – Thursday, 11/05**
- **Exam 3 – Thursday, 12/17**

Your letter grade for the course will be completed as follows:

Numerical Score	Grade	Numerical Score	Grade
>=92	A	72-78	C
90-92	A-	70-72	C-
88-90	B+	68-70	D+
82-88	B	62-68	D
80-82	B-	60-62	D-
78-80	C+	<60	F

While this overall grading scheme is fixed, I will be happy to discuss any issue you may have with individual grades. If you notice a mistake or have a question regarding a specific grade, please come and talk to me *as soon as possible*. Do not wait until the end of the semester to bring up grading issues.

Assignment/Review Problem Deadlines: Each assignment and review problem will come with a deadline (day and time) by which it must be submitted. Late submissions will NOT be accepted. Extensions on deadlines may be granted at the discretion of the instructor if you provide a valid justification (in the form of a written excuse from a medical doctor or the Dean of Students Office) **before** the due date.

If you miss a scheduled exam (tentative dates are provided), you **may** be able to take a make-up exam provided you give the instructor a valid justification (see above) ahead of time if possible. Only one make-up exam will be given. It will be a comprehensive exam scheduled at the end of the semester. Similarly, there will be no make-up quizzes unless the instructor is provided with a valid justification (see above) for your absence on the day of the quiz.

Collaborating versus Cheating: Unless otherwise stated in the assignment or project, all submissions must be entirely your own work. While it is acceptable to discuss the assignments at a high level (for example, at the design level) with others, you must submit your own work. **You may not "borrow" any piece of code or design of any length from**

someone else, the internet, or any other source, unless you can live with a zero and the other potential academic sanctions of cheating (see UWO Student Discipline Code 2007, Chapter UWS 14).