CSC250 Computers and Human Intelligence

Professor Clark Elliott

Spring 2018 /2019

Logistics:

Class meets: Wed 5:45 Downtown CDM 222
Professor: Dr. Clark Elliott
Class website: http://condor.depaul.edu/~elliott/250
email: elliott@depaul edu.
[Include "250: " prefix in subject line and MEANINGFUL mail header!]
Grader email: elliottgrading AATT gmail com. (Administration of submissions only)
Course Management: D2L.depaul.edu

This syllabus may be updated during the quarter with respect to learning goals, but with prior notice and in ways that will not affect the grading structure.

Class overview:

This is a lecture and discussion class, with background readings. We will use rudimentary hands-on AI programming assignments to become familiar with a few of the tools that AI researchers use. In addition to covering the basics of traditional AI, including topics like "intelligence as search through the problem space" and "symbolic models of reasoning by heuristic classification (explanation-based reasoning)" we will use this class as a laboratory for a discussion of the nature of science itself. Many very hard, and unsolved problems arise when we attempt to create artificial human-like intelligence.

But our approach is not primarily philosophical. To the degree it is possible, we will illustrate and discuss the fundamental nature of the problems that arise when scientists attempt to represent human intelligence and sentience by using purely computational tools. Thus, this class is necessarily interdisciplinary in nature in that it draws on scientific results in psychology, cognitive science, neuroscience and information processing as well as logical results from philosophy and theoretical principles from physics.

Textbook:

We will not use a textbook this quarter, but will have readings from many sources provided either as web links, or provided at D2L.

Recommended background: Ertel, Wolfgang. *Introduction to Artificial Intelligence*, Weingarten, Germany: Springer, 2011. ISBN-13: 978-0857292988 http://www.amazon.com/Introduction-Artificial-Intelligence-Undergraduate-Computer/dp/0857292986

Grading:

Exams and quizzes 60%
Assignments 40%

```
Grading Scale:
                95%             A
                90%             A-
                86,83,80        B+,B,B-
                78,74,70        C+,C,C-
                65,60           D+,D
```

No late assignments graded for credit.

I reserve the right to raise the grade of a student who has done exceptional work in some aspect of the coursework.

I reserve the right to move credit between exams and quizzes with prior notice. I reserve the right to give "free points" if the class corporate body is demonstrating exception collegiality on the discussion forums.

All grades are subject to Academic Integrity Sanctions. See the class website and the student handbook for further discussion.

"Minor points" notation:

From time to time I use the point box as a communication vehicle in two specific ways, and I reserve the right to add minor points for this purpose:

- One point extra: I am tipping my hat to you for particularly fine work. That is, if you get 101 points on a 100 point programming assignment, I may be saying, "Hey, I noticed the five extra modules your wrote. Good job!"
- Two points extra: If you receive two extra points, I am acknowledging an *exceptional* contribution beyond expectations, so 102 points on a 100 point assignment is something to feel really good about, and is a rare compliment.
- Grade of "1": used as a placeholder to let a student know that I have reviewed an assignment, and am waiting for further information or work as per correspondence. A "1" will *always* be resolved to a different grade.
- Grade of "2": a serious warning that you need to communicate with me about possible plagiarism or some other irregularity that is being investigated.

Practical Survey Topics:

1) Artificial Neural Networks basics
2) Ambiguity in the English language; how world knowledge is required to understand it.
3) Pattern matching and the generation of categories using exemplars and prototypes.
4) Categorization as one basis of human intelligence
5) Reduction of world representation into recursive symbolic languages
6) Search as an organizing principle in AI.
7) The debate about the usefulness of studying (a) symbolic logic and (b) connectionist models.
8) Deductive, inductive and abductive logic
9) Human AI versus Alien AI
10) Procedures versus knowledge versus designing systems that learn.
11) The principles of the plastic brain that constrain fixed design for human AI
12) Forward and backward chaining problem solving design (bottom-up vs. top-down).
13) Depth First Search, Breadth First Search, basic A* through problem spaces
14) Basic pattern matching with variables and Unification
15) Constraint satisfaction for problem solving.
16) Propositional and Predicate logic for problem solving
17) Emotion modeling and building a compassionate computer
18) The translation of visual and auditory input into symbolic meaning
19) AI Models of human memory
20) AI models for parsing language

Overview topics (integrated with the above units):

1) The physics of undecidability and the design of free choice in computational systems
2) Implications for society of machines that can make intelligent and *autonomous* decisions.
3) The role of creativity when approaching the currently intractable problems of AI.

4) Evaluating the promises of the past, and correctly evaluating what we may achieve now
5) What we can learn about ourselves from studying artificial intelligence
6) What evidence, hypotheses and scientific claims look like in the creative study of AI
7) How society, the military and commercial gain may drive the study of AI
8) Grounding the logic of philosophy in absolute computational models
9) Implications of the *singularity* when machines may know more than we do

Sample light programming:

1) Symbolic programming in LISP [Using ABCL over Java http://common-lisp.net/project/armedbear/]
2) Symbolic programming in PROLOG [Using JIPROLOG http://www.jiprolog.com/] http://www.learnprolognow.org/lpnpage.php?pagetype=html&pageid=lpn-htmlse1
3) Simple expert system creation [Using in-house package, or, e.g., http://www.lpa.co.uk]
4) Creation of an artificial neural network using a simulator.
5) Using the online natural language processing tools provided by Stanford and Princeton.

Class structure:

This survey course is primarily a lecture-discussion and reading class, with some hands-on assignments, and exams. Discussions will take place in class, an online using the D2L threaded forums, especially to include our online students.

There are programming assignments, but these will be quite basic, and no prior programming experience is necessary. Students must be familiar with their computing platform such that (a) they can install and run a simple Java virtual machine, and (b) know how to use a simple text editor for writing programs.

There will be one large formal writing assignment, and several smaller written assignments.

Learning Goals:

Borrowing from the Scientific Inquiry structure for both *Science as a Way of Knowing* and *Elective Inquiry* within content domains, we can define the following learning goals for CSC250. At the end of class you will be able to…

Estimate the time table for—and limits on—our ability to understand human intelligence and our ability to mimic that intelligence with computers using a purely scientific approach.

Connect the evidence produced by *running models* of AI systems with the theories and hypotheses that drove the design of those systems and evaluate claims and refutations of the theories that are drawn as conclusions of such systems.

Critically evaluate the many (often erroneous) assumptions made about the staggering complexity of human intelligence, alongside the stunning growth in the computational power of both hardware and AI software that might be used to model it.

Evaluate the evidence of computational successes and failures for both *human AI* (mimicking the design of human systems) and *alien AI* (seeking human computational results without constraints on design).

Evaluate the role of creative approaches to solving these currently intractable problems, and role of truly scientific (and explicitly open-minded) skepticism in the context of what modern neuroscience tells us about human information processing.

Recognize the scientific difficulties in studying human cognition wherein we have a computationally closed system of great complexity that requires us to continually evaluate results with the goal of further minimizing the substantial uncertainty in our theories.

Evaluate the study of AI as a whole, and successes and failures of various broad approaches that have been taken to model human intelligence.

Minimally apply the tools we use in studying AI to a representative set of problems

Apply what we know about human intelligence in a computational framework to other disciplines that also study the human psyche.

Have a solid grounding in the implications for society of building artificially intelligent systems that make autonomous decisions previously exclusively reserved for real humans.

Have a broad understanding of the central problems in artificial intelligence.

Have a basic understanding of algorithmic approaches to human problem solving, and how computational mechanisms function as a mechanism for precisely defined inquiry.

Know how to write simple AI programs.

Write a sophisticated argument paper about applying artificial computational intelligence to systems in our modern world.

Demonstrate a broad survey-level knowledge of artificial intelligence on exams.

These learning goals will be met through lecture and both in-class and online discussion of the above topics that integrate (a) the details of how AI works in practice, and also (b) where the study of AI fits within the greater practice of science in the modern world. The study of AI presents particular problems for science because of the self-reflective nature of intelligent thought itself, and because of the human condition that is embodied in human intelligence and sentience. Neither of

these is well understood at an algorithmic level. There is no unifying algorithmic theory of thought that is even close to being universally accepted.

Office hours for the course are available from my faculty link at cdm.depaul.edu.

All assignments, the assignment schedule, and the course materials, are available online at either D2L or the class website.

Sample D2L Discussion Forums:

See http://d2l.depaul.edu | Discussions

Say Hello—Provide a (possibly brief) biography with which you are comfortable. Meet your classmates. Say hello!

The nature of sentience—Discussion of some of the difficult problems and the seeming paradox of representing sentience through software and hardware implementations of Turing machines.

Neural network connectionist models vs. symbolic models—On the one hand it is clear that the human brain and its neural structure is at least one component of human intelligence (and possibly the only one). But it is also clear that higher-level symbolic processes provide the organizing principles (such as categorization) for the lower-level (and thus possibly irrelevant) implementation of intelligence. Which model should be the locus of AI study?

Human AI vs. Alien AI—Since the human brain is by far the most powerful computational device in the known universe, and because it also implements human intelligence, should we exclusively study the brain's way of implementing true intelligence? Or should we care only about meeting the requirements of intelligence without regard to how it is implemented?

The nature of scientific inquiry—In what ways do the seemingly intractable problems we encounter in modeling human intelligence push the boundaries of science? What is the balance we should be seeking between creative approaches that lead to places not yet known (build it first and see where we end up later), and traditional scientific approaches based on the gathering of evidence (such as through measurements coming into vogue in the new neuroscience) and the assessment of hypotheses? What benefits to society might we find along the way? What is the role of *basic research* into the weighty questions of modeling human cognition?

Submission File Formats:

All submissions to D2L MUST BE IN THE SPECIFIED FORMAT or they will not be accepted for credit. No other formats will be graded. Often this means in STANDARD ZIP FORMAT, including submissions of a single file. No 7zip files, no .rar files. No exceptions. Contained within the ZIP archive, all text submissions must be made in Microsoft Word format, or in plain ascii HTML, or plain ascii text. NO PDF FILES. (Free programs are available to produce each of these formats. In some cases zip files are prohibited.)

Follow the instructions for each assignment.

Students are responsible for downloading their assignments after uploading to D2L, to make sure that files have not been corrupted. Corrupted files will not be graded.

NO LATE ASSIGNMENTS will be accepted for credit, unless otherwise noted. For students that miss assignments, there will typically be limited extra credit assignments available.

Academic Integrity:

Cheating, plagiarism, and unethical conduct are not allowed, and will be sanctioned, including referral to the dean's office, and failure in the class. Please refer to the academic handbook by which rules you are expected to abide.

Violations include, but are not limited to: making claims on any checklist for work that has not been done; including ANY un-cited work of others in any documents you turn in; turning in work, including any program, that has been authored by someone other than yourself; and in some cases including *any* work of others, whether cited or not–see the rules for each assignment.

Instructor:

Dr. Clark Elliott received his degree in Artificial Intelligence from Northwestern University's Institute for the Learning Sciences in 1992. He has served as Associate Director of the Institute for Applied Artificial Intelligence at DePaul.

As part of his thesis work he developed one of the first comprehensive running AI models of human emotion embodied in a program called *The Affective Reasoner*. As an outgrowth of his work with models of emotion in multi-agent social systems he developed one of the world's first master's programs in Distributed Systems at DePaul, based on Web and Internet technologies, which courses have remained in the core of the College of Computing and Digital Media for two decades. In his thirty years at DePaul and as one of the early faculty in CDM he has taught forty different courses including many AI and Cognitive Science courses.

In addition to Computer Science, Dr. Elliott helped to design the neuroscience program curriculum at DePaul and has continuously served on the university committee that oversees the program since its inception. Since the publication of his book *The Ghost in My Brain* in 2015, Dr. Elliott has spoken to more than ten million people via radio and T.V. interviews about the cognitive neuroscience of brain injury. He has been in demand as a speaker at neuroscience, and neuro-optometric rehabilitation conferences and conventions. He has published case-studies in neuroscience relative to movement disorders and the brain.

With DePaul graduate students he is currently developing AI systems to generate automated story content for computer games via *story-morphing* techniques based on automated personality construction, and is working on building a compassionate computer using AI techniques relative to the representation of emotion, text-to-speech, speech recognition, modulated delivery of music, natural language processing, and real-time morphing emotionally expressive faces.

Dr. Elliott also holds a Bachelor of Music degree in Music Education with three concentrations, a Master of Music degree in Music Performance and a Master of Science degree in Artificial Intelligence, with study at the Eastman School of Music, the Juilliard School, San Francisco State University, the University of California at Berkeley and DePaul University.

More Policies

Changes to Syllabus

This syllabus is subject to change as necessary during the quarter. If a change occurs, it will be addressed during class, posted in D2L, and sent via email.

Online Course Evaluations

Evaluations are a way for students to provide valuable feedback regarding their instructor and the course. Detailed feedback will enable the instructor to continuously tailor teaching methods and course content to meet the learning goals of the course and the academic needs of the students. They are a requirement of the course and are key to continue to provide you with the highest quality of teaching. The evaluations are anonymous; the instructor and administration do not track who entered what responses. A program is used to check if the student completed the evaluations, but the evaluation is completely separate from the student's identity. Since 100% participation is our goal, students are sent periodic reminders over three weeks. Students do not receive reminders once they complete the evaluation. Students complete the evaluation online in CampusConnect.

Academic Integrity and Plagiarism

This course will be subject to the university's academic integrity policy. More information can be found at http://academicintegrity.depaul.edu/. If you have any questions be sure to consult with your professor. There is NO CHEATING OF ANY KIND in this class!

Academic Policies

All students are required to manage their class schedules each term in accordance with the deadlines for enrolling and withdrawing as indicated in the University Academic Calendar. Information on enrollment, withdrawal, grading and incompletes can be found at: cdm.depaul.edu/enrollment.

Students with Disabilities

Students who feel they may need an accommodation based on the impact of a disability should contact the instructor privately to discuss their specific needs. All discussions will remain confidential.
To ensure that you receive the most appropriate accommodation based on your needs, contact the instructor as early as possible in the quarter (preferably within the first week of class), and make sure that you have contacted the Center for Students with Disabilities (CSD) at: csd@depaul.edu.

Lewis Center 1420, 25 East Jackson Blvd.
Phone number: (312)362-8002
Fax: (312)362-6544
TTY: (773)325.7296

Other Course Policies

Attendance: Students are expected to attend each class, or view the class online, typically during the week the lecture is presented. Attendance will not be formally taken beyond the start of the quarter, but unless otherwise noted ALL the course material presented in the lectures is suitable for exams. I will typically ask questions of named students NOT present in the classroom, to be viewed online, and answered at the forums. All students matter to me. Many CDM courses are also recorded via Course Online software for later review.

Class Discussion: Student participation in class discussions is expected, and this will take place in class for local students, and online for all students.

Attitude: A professional and academic attitude is expected throughout this course. Measurable examples of non-academic or unprofessional attitude include but are not limited to: talking to others when the instructor is speaking, mocking another's opinion, cell phones ringing, emailing, texting or using the internet whether on a phone or computer. If any issues arise a student may be asked to leave the classroom.

Cell phones / laptops in class: If you need to use your cell phone for any reason, or your laptop for any reason other than following the class slides, and taking notes, *leave the room.* You may quietly leave and re-enter as often as necessary unless I note otherwise. Your peers devote many hours out of their busy lives, and hundreds of dollars, to come to class. They deserve a vibrant, focused, environment. If you have a special case, discuss it with the instructor ahead of time. NO TEXTING, EMAIL, FACEBOOK, etc. in the classroom.

Civil Discourse: DePaul University is a community that thrives on open discourse that challenges students, both intellectually and personally, to be Socially Responsible Leaders. It is the expectation that all dialogue in this course is civil and respectful of the dignity of each student. Any instances of disrespect or hostility can jeopardize a student's ability to be successful in the course.