University of Strathclyde Science

Lecture 1: Introduction to the Class

Dr John Levine

CS310 Foundations of Artificial Intelligence
January 19th 2016

University of Strathclyde Science

Learning Outcomes

At the end of this lecture, you will know:

1. What you need to do pass this class.
2. How to define the term “Artificial Intelligence”.
3. How AI has developed since its birth in 1956.
4. The state of the art: what can AI do today?

The Rules of the Game

- CS310: Foundations of Artificial Intelligence
- 20 credits, Semester 2
- Prerequisites:
  - CS207 Advanced Programming
  - CS208 Logic and Algorithms
- Lecturer: John Levine
- Assessed by:
  - coursework (40%)
  - degree exam (60%)

The Rules of the Game

- Lectures are Tue 1-2 (TG 314) and Thu 2-3 (CV 512).
- Tutorials: Tue 2-3 (TG 310) from Week 2 (next week) onwards.
- Labs: Mon 3-5 (LT12.01) and Tue 11-1 (LT11.05) from Week 3 onwards.
- Web page: http://www.cis.strath.ac.uk/~johnl/CS310
- Email: John.Levine@strath.ac.uk
- Room LT1420, extension 4524

The Rules of the Game

- There will be four practical assignments for the class, each worth 10% of the total mark.
- For the practicals, you'll get your code signed off in the lab and then submit it online for performance testing.
- A ticket system will operate in the labs for getting your code signed off.
- The tutorial work is not assessed, but the degree exam will be based on the tutorial work.
- The degree exam contributes 60% of the final mark for the class.

Time Management

- You should spend about 10-15 hours per week on this class (20 credits = 200 hours, spread over 16 weeks)
- A typical week might look like this:
  - 2 hours of lectures
  - 1 hour at the tutorial
  - 2 hours in the lab
  - 2-3 hours of practical work
  - 2-3 hours of tutorial preparation
  - 1-3 hours of online study

University of Strathclyde Science

Udacity Online Class

- Udactiy "Intro to Artificial Intelligence" by Sebastian Thrun and Peter Norvig.
- We won't cover the whole class, just selected topics
- http://www.udacity.com/course/cs271
- Study for this week: Unit 1, "Welcome to AI"

Course Aims

From the syllabus page:

To give the student a broad appreciation of the scale and nature of the problems within Artificial Intelligence and a detailed understanding of the fundamental techniques used to address those problems.

But before we start…

What exactly is "Artificial Intelligence"?

What is Artificial Intelligence?

- Should a successful AI be an artificial human being?
- Does it need to have emotions in order to be human?
- Does it need to have a human form?
- Is Intelligence to do with thinking clever thoughts or is it to do with doing clever things?
- Does successful AI need to think in the way we think?
- When does a computer program turn into “an AI”?

What is Artificial Intelligence?

Four definitions:

- "AI systems are designed to behave in exactly in the same way that humans behave"
- "AI systems seek to model human thought processes precisely and solve problems in the same way we do"
- "AI systems are those which can make exact logical deductions based on their knowledge of the world"
- "AI systems make rational decisions about how to act in the world to gain the best possible reward"

Acting Humanly: the Turing Test

- If a computer can imitate a human being well enough to fool an observer, then it must be intelligent
- Not the easiest thing to do – we're nowhere near at the moment
- Is the goal of artificial flight to make "machines that fly so exactly like pigeons that they even fool other pigeons"?

Thinking Humanly: Cognitive Science

- Do experiments to find out how the human mind works (psychology)
- Build computational models to test those theories
- Clearly related to the goals of artificial intelligence, but not directly trying to make better computer programs
- Cognitive Science can clearly inform AI, but we may want to ignore the advice it gives us
- We may get better behaviour by not thinking humanly

Thinking Rationally: Laws of Thought

- Rational arguments and logical deductions:
  - Socrates is a man.
  - All men like football.
  - Therefore, Socrates likes football.
- Problem one: it's really not at all easy to translate all the knowledge in the world into logical formulae.
- Problem two: even if you can, the problem of inference turns out to be computationally nigh on impossible.
- Problem three: even if you are the world's greatest logician, does this help you decide what to do next?

Acting Rationally: The Rational Agent

- An *agent* is just something which lives in an environment and performs actions
- A *rational agent* is one that acts in order to achieve the best outcome whenever it needs to choose one action over another
- As an observer of such an agent, we say that it is acting in an intelligent way: it exhibits intelligent *behaviour*
- To act in a rational way requires aspects of the other approaches mentioned, but it is more general, and it is more amenable to scientific experiment

What does “Rational” mean?

- A *rational agent* is one that acts in order to achieve the best outcome whenever it needs to choose one action over another
- Rational means that, given all the things the agents has perceived in its lifetime and its current goals, it chooses its action to maximise its expected reward
- Rationality is not the same as omniscience: the actual reward may turn out to not as good as expected
- Rationality means making an *informed* decision – given what we know, what’s the best thing to do?

What's Needed for Rationality?

- Problem-solving skills: being able to formulate problems and search for solutions to those problems
- Game playing: knowing what action to do next in an adversarial situation in order to win
- Planning: deciding on a sequence of actions that will achieve a goal or a set of goals
- Knowledge representation: holding knowledge about the way the world works
- Learning: finding out how to perform better

The Gestation of AI (1943-1955)

- Early work in the 1940s concentrated on trying to build electronic brains
- Minsky and Edmonds build the first neural network computer in 1950 – the SNARC
  - 3000 vacuum tubes
  - a surplus automatic pilot mechanism from a B-24 bomber
  - simulated a network consisting of 40 neurons
- Alan Turing’s 1950 paper “Computing Machinery and Intelligence” contains a complete vision of AI

The Birth of AI (1956)

- John McCarthy: the Dartmouth workshop:

  “We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.”

- Allen Newell and Herbert Simon: the Logic Theorist

Look Ma, No Hands! (1952-1969)

- General Problem Solver (GPS)
- Arthur Samuel: early checkers-playing program
- New programming language to process symbols – LISP
- The split between those working on logical systems and those working on “getting the program to work”
- Early successes with Natural Language understanding (SHRDLU by Terry Winograd, 1972) and AI planning

A Dose of Reality (1966-1973)

- Early researchers were not shy of predicting success:

  "It is not my aim to surprise or shock you – but the simplest way I can summarise is to say that there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until – in a visible future – the range of problems they can handle will be coextensive with the range to which the human mind has been applied." Herbert Simon, 1957

- However, efforts to automatically translate scientific texts from Russian into English were very unsuccessful

- 1966: U.S. Government pulls plug on all funding for research into machine translation

Expert Systems (1969-present)

- Move to stronger methods to tackle particular tasks
- Diagnosis, configuration, story understanding
- Emphasis on knowledge required to solve such tasks (e.g. rules for prediction and diagnosis)
- Early expert systems:
  - DENDRAL (interpretation of mass spectra)
  - MYCIN (medical diagnosis)
  - R1 (configuration of computers for DEC)
- AI becomes an industry (briefly)

The return of Neural Networks (1986)

- A natural competitor to symbolic systems for doing tasks involving learning and adapting
- Seen as a way of coping with uncertain environments
- Seen by some as an alternative to symbolic systems
- 1986: Parallel Distributed Processing by Rumelhart and McClelland creates much interest
- ANNs are used today to classify patterns in data and to reactively direct the behaviour of an agent

AI Becomes a Science (1987-present)

- Emphasis in AI is now on experimentation, building on previous theories and improving performance
- Success comes with maturity:
  - Speech understanding based on HMMs
  - Data mining based on statistical methods
  - Rational diagnosis based on decision theory
- Very large datasets are now playing a key role in the performance of AI systems:
  - Word sense disambiguation
  - Filling in “holes” in digital images
  - Analysis of DNA data (bioinformatics)

The State of the Art

- Successful AI systems:
  - NASA Remote Agent experiment
  - Driverless cars (DARPA Grand Challenge)
  - Deep Blue (Chess)
  - Spam detection algorithms
  - The Roomba vacuum cleaner
  - Watson wins at “Jeopardy!” in 2011