PS452 Intelligent Behaviour
Academic Year 2017/2018 (spring term)

Overview

The aim of this course is to investigate the various theories intended to account for intelligent behaviour in humans, animals and machines. The approach will be multidisciplinary, looking for common themes and comparing different perspectives.

The first two lectures will briefly survey human intelligent behaviour in the form of the cognitive processes proposed for reasoning and solving problems, along with theories of individual differences in human intelligence from the perspective of cognitive psychology and neuropsychology: why are some people better than others at reasoning and problem solving?

The next four lectures will investigate attempts to make computers behave in intelligent ways. Knowledge of this field is relevant to cognitive psychology because the research highlights just how much we take for granted about human cognition. No knowledge of computer programming will be necessary. We will look at the underlying theories and assumptions of this field and examples of the most well-known AI programs. Reasons for their failure will be discussed along with a survey of new directions and prospects.

The final four lectures will look at attempts to understand non-human animal intelligence and see whether parallels can be drawn with humans in the domains of reasoning, tool use, communication and theory-of-mind. By seeing how animals and humans differ, we may be in a better position to understand our own intelligence and program computers more effectively.

Teaching

The course will be delivered via 10 weekly two hour lectures (within a three-hour slot). These are scheduled for Tuesday mornings in Term 2, between 09:00 and 12:00 in NTC 1.02. An additional class will be scheduled for the final multiple choice test, and some revision suggestions and exam strategies will also be given. This is scheduled for Tuesday 1st May from 09:00 to 11:00 in room TC1.12. If you wish to contact me away from the lectures, use email in the first instance: mjr@essex.ac.uk

Assessment

Assessment is via a two-hour end-of-year essay-based examination (75%), and weekly multiple choice tests (25%), five questions per week, administered before each lecture. The tests are effectively untimed. I will discuss details of them in the first class.

Attendance

Because of the weekly test, it is essential that you only enroll on this course if you are able to attend the scheduled lectures. Some absences can be accommodated, but not complete non-attendance.

Syllabus

The organisation of course topics is as follows:

Humans

Lecture 1: *Tools for intelligent behaviour.* Problem solving strategies for well-defined problems, mental models for reasoning and inference, and schemas as structured knowledge for problem solving. Limitations in human performance, and recurring themes and issues.

Lecture 2: *Theories of human intelligence.* Domain-general versus domain-specific theories of differences in intelligence, working memory capacity accounts, intelligence as inhibition of incorrect answers. Low-level cognition and brain structure and function explanations of differences in intelligence. Synthesis and more recurring themes.

Machines

Lecture 3: *What is Artificial Intelligence?* Artificial Intelligence: process- behaviour- and task-based definitions. The Symbol System Hypothesis and brain-computer equivalence. Intelligent behaviour as heuristic-driven navigation of problem state spaces. The Symbolic Search Space Paradigm as the dominant approach to AI.

Lecture 4: *Artificial Intelligence Landmarks*. A survey of the famous programs; solving games and puzzles, natural language processing and expert systems. Knowledge-light versus knowledge-heavy approaches. The harsh evaluation of the first 25 years of attempts to create Artificial Intelligence.

Lecture 5: *Artificial Intelligence: Observations, objections*. The Turing Test as a behavioural benchmark. The five barriers to true Artificial Intelligence; the Logic Problem, the Knowledge Problem, the Frame Problem, the Intentionality Problem and the Symbol Grounding Problem. The human brain versus the computer.

Lecture 6: *Artificial Intelligence: Where Next*? CYC as the last standing for traditional Artificial Intelligence, connectionism, situated cognition & Artificial Life.

Animals

Lecture 7: *Animal intelligence tests.* Pitfalls in studying animal intelligence. Understanding and measuring individual differences in intelligence between species, and individual differences within species. Working memory capacity and attentional focus as the key to understanding differences in ability between species (including humans).

Lecture 8: *Tools, puzzles, beliefs, and intentions*. Natural tool use, problem solving experiments in the laboratory, the nature of causal beliefs and intentions.

Lecture 9: *Animal communication.* Communication in the wild, attempts to teach animals sign or artificial languages. Implications of their failure for theories of intelligence: words as symbols versus words as tools, and the cognitive fringe-benefits of language.

Lecture 10: *Theory of mind and deception.* Cognitive capacity versus modularity as explanations for social behaviour. Animal (lack of) awareness of other minds. Deception in the wild and in the laboratory. The evolution of general intelligence as a response to deadly technology, not social complexity. Looking back across the ten lectures.

Recommended Reading

There are no textbooks currently available which cover all of the topics on this course, or even two of the three major ones. This means that the course has had to be assembled from several different sources. The major sources for each lecture are identified in the references section. There is one book that considers animal versus machine intelligence, but it dwells on phenomenology more than I intend to. Nonetheless, it is worth a read.

McFarland (2008). *Guilty Robots, Happy Dogs: The Question of Alien Minds.* Oxford: Oxford University Press.

Lectures 1 and 2

Eysenck, M.W. & Keane, M.T. (2010). *Cognitive Psychology: A Student's Handbook (6th Ed).* Hove: Psychology Press. (Part IV).

Duncan, J. (2010). *How Intelligence Happens*. Cambridge: Cambridge University Press.

Lectures 3 to 6

Copeland, B.J. (1993). *Artificial Intelligence: A Philosophical Introduction.* Oxford: Blackwell.

Ekbia, H.R. (2008). *Artificial Dreams: The Quest for Non-Biological Intelligence.* Cambridge: Cambridge University Press.

Lectures 7 to 10

Reznikova, Z. (2007). *Animal Intelligence.* Cambridge: Cambridge University Press.

Pearce, J.M. (2008). *Animal Cognition.* Hove: Psychology Press.

MJ Roberts, 28th January 2018

Past Exam Papers

Summer 2013

1. Give an overview of the various working memory accounts of individual differences in human intelligence. Evaluate the extent to which this knowledge has been of use for EITHER the creation of intelligent machines OR the interpretation of animal behaviour.
2. Discuss the various attempts to define Artificial Intelligence, and the difficulties encountered in providing an adequate definition. Explain the consequences of this for researchers attempting both to create or evaluate Artificial Intelligence.
3. Compare and contrast the traditional approach to Artificial Intelligence to ONE attempt to overcome its problems. Choose from (a) Connectionism; (b) Lenat's CYC project; (c) Artificial Life/Situated Cognition.
4. What has been learnt from the various studies of non-human animal tool use, both in the laboratory and the natural environment. Evaluate the extent to which the ability to use tools by non-human animals depends upon their attentional focus as well as their cognitive capacity.
5. Discuss the extent to which apes who have been taught 'language' by human trainers are engaging in intelligent problem solving rather than communication.
6. Comparing human intelligence with non-human animal and machine intelligence, to what extent are humans 'special', and how? Comparing human intelligence with machine intelligence, what are the prospects for this situation continuing in the future?

Summer 2014

1. Considering people with high intelligence test scores versus low intelligence test scores, give an overview of the claims that such differences can be accounted for by physical brain structure (including research from imaging studies). Evaluate the extent to which such research provides useful additional information compared with traditional cognitive psychological approaches.
2. For many years, researchers attempting to create intelligent machines have based their work on the symbol system hypothesis (a symbol system can be intelligent) and the state space paradigm (intelligence = efficient navigation of a state space). Explain these concepts, and evaluate whether this approach will ultimately prove to be fruitful or futile.
3 Give overviews of (i) Turing's test for computer thinking, and (ii) Searle's Chinese Room thought experiment. Explain the conflict underlying these two conceptualisations and whether these can be reconciled in any way.
4 Many researchers have attempted to evaluate non-human animal intelligence using various simple learning/logic tasks to different species. What have these experiments shown, and what are the pitfalls that can make interpretation of this research difficult?
5 Explain why the concept of Theory of Mind is important for the debate concerning whether cognitive processes are domain-general or domain-specific. Discuss whether any evidence demonstrates that non-human animals might have an awareness of the minds of others.
6 Discuss the extent to which human expectations and prejudices can interfere with attempts to understand and interpret potentially intelligent behaviour displayed by machines and non-human animals.

Winter 2015

1 Give an overview of the debate concerning whether human cognitive processes are domain-free versus domain-specific. Evaluate the extent to which this debate is important in EITHER the creation of intelligent machines OR the interpretation of non-human animal behaviour.

2 Discuss the various attempts to define artificial intelligence, and the difficulties encountered in providing an adequate definition. Explain the consequences of this for researchers attempting both to create or evaluate artificial intelligence.

3 It has been claimed that the difficulties faced by artificial intelligence researchers can be accounted for by a set of fundamental problems associated with attempting to make computers knowledgeable and able to make intelligent inferences. Give an overview of these problems and assess whether there is any likelihood that they will be resolved.

4 The working memory hypothesis can offer a useful account of individual differences in human intelligence. Discuss whether this could also successfully account for differences in problem solving performance between members of the same non-human animal species, and between different animal species.

5 There are many non-human animals that communicate in natural settings, and it is even possible to identify instances of communication that appear to be deceptive. Discuss the difficulties encountered in determining whether these might be signs of intelligent behaviour.

6 Discuss the extent to which human expectations and prejudices can interfere with attempts to understand and interpret potentially intelligent behaviour displayed by machines and non-human animals.

Winter 2016

1 Imagine that aliens from outer space have landed on planet earth. Evaluate the extent to which the cognitive processes that underlie their intelligent behaviour (and the reasons for individual differences) *should* be similar, and *might* be different from humans.

2 Give an overview of the landmark Artificial Intelligence programs, highlighting the recurring problems that have prevented them from emulating true human intelligence.

3 Briefly describe (i) Turing's test for computer thinking, and (ii) Searle's Chinese Room thought experiment. Explain the conflict underlying these two conceptualisations and whether these can be reconciled in any way.

4 Discuss the proposal that the attentional focus of animals is an important factor in determining their problem solving, communicative and social skills.

5 It is possible to identify instances of communication between non-human animals that could be labelled as deceptive. Discuss the difficulties encountered in determining whether these acts might be signs of intelligent behaviour.

6 Discuss the ways in which the word intelligence has different definitions, usage, and connotations in the fields of human, machine, and animal behaviour. To what extent is this variability problematic?