Introduction to Cognitive Science

Gabe Dupre

February 24, 2019

Course Overview

Cognitive Science is the interdisciplinary study of the mind. Research in this area draws on psychology, philosophy, linguistics, neuroscience, computer science, anthropology, biology, and various other areas. The aim of this area of inquiry is to create an understanding of the mental life of humans and other animals which is able to integrate insights from all these disparate but overlapping disciplines. Of course, this is a massive project, and takes at its target what may be the most complex set of systems in the known universe. This makes cognitive science difficult, but also hugely exciting and rewarding.

While the study of the human mind has roots tracing back thousands of years, cognitive science as a distinct approach has emerged only relatively recently, around the middle of the 20th Century. This makes it very young for a science. There is therefore little in the way of consensus in the field with respect to either theory or methodology. This course will give students the opportunity to grapple with foundational scientific debates as they are currently developing.

There are two primary aims of this course. Firstly, students taking this course should come away with a knowledge of some of the major achievements and proposals within this branch of science. We will aim for generality, covering a wide range of theories and approaches, giving students a taste of what work in cognitive science involves. Secondly, we will focus on the puzzles that arise when attempting to integrate this work. One particular difficulty that arises in doing cognitive science is that the different strands of research on which cognitive science draws often suggest quite different ways of looking at the mind. For example, one and the same mental process can appear very different when viewed from a psychological and neuroscientific perspective. It is almost always difficult to see how these different perspectives can be fit together so as to create a unified account of their target, the mind. While we develop an understanding of seminal work in these distinct areas, we must always keep in mind that the goal is ultimately to provide as unified an understanding as possible.

Reading Schedule

Week 1: Introduction

- Pylyshyn, Z. *What is Cognitive Science?*
- Block, N. *The Computer Model of the Mind*
- Gallistel, C.R. *Representation in Animal Cognition: An Introduction*

This week we shall introduce the aims and methods of cognitive science, from a high-level perspective. Block and Pylyshyn describe the ‘information processing’ approach to the mind. According to this approach, we explain agential behavior and capacities with reference to the capacities that the mind/brain has for receiving, storing, manipulating, and responding to relevant information in the environment. Gallistel provides a neat case study involving applying this approach to explain the ability of ants to return to their nests upon finding food in their environment along the simplest path, even though they arrived at the food via a circuitous route. This information processing approach is widespread in cognitive science, and forms the background for much of the work we will look at. However, as we shall see later on in the course, it is not universally accepted.

Week 2: Levels of Explanation

- Marr, D. *Vision* (Chapter 1)
- Carandini, M. *From Circuits to Behavior: A Bridge too Far?*

Given the enormous complexity if our target, human and non-human minds, it is customary in cognitive science to divide up the target into different 'levels of explanation'. The first, and most influential, such division was given by Marr. Marr divided the task of psychology into three subtasks: describing the computational, algorithmic, and implementational levels of mental capacities. At the topmost, computational, level, the aim is to describe the task that the cognitive system is supposed to complete. This is done by describing a function from environmental situations (inputs) to the system's responses (outputs). At the intermediate, algorithmic, level, is described the step-by-step process by which this function is computed by the system. At the lowest, implementational, level, we describe the ways in which this algorithm is realized by a physical, e.g. neurobiological, system. Carandini provides a more up-to-date discussion of how analyses of mental processes at these different levels are able to inform one another.

Week 3: Minds and/as Machines

- Haugeland, J. *What is Mind Design?*
- Newell, A. and Simon, H. *Computer Science as Empirical Inquiry*

This week we aim to precisify our characterization of the mind as an information processing system, in particular by viewing the mind as a computer, a paradigmatic and well-understood device for manipulating information. Haugeland describes the methodology by which we are able to analyze the mind as a computational system, and Newell and Simon likewise show how central notions from computer science (symbolic systems and heuristic search) can be used to illuminate the workings of natural systems like the mind.

Week 4: Symbolic and Connectionist Systems

- Rumelhart, D. *Architecture of Mind*
- Sejnowksi, T. *The Deep Learning Revolution* (Excerpts)
- Hinton, G. *How Neural Networks Learn from Experience*
- Marcus, G. *Deep Learning: A Critical Appraisal*

We look here at one of the central debates in the history of cognitive science: between 'classical theorists' who view the mind as a symbol manipulation system, in line with Newell and Simon's proposal, and 'connectionists', who propose a radically different, and potentially more neurobiologically plausible model. Connectionist models of the mind involve positing large numbers of 'nodes' interacting with one another to form 'neural networks'. These networks display surprising and impressive abilities to learn on the basis of environmental input. Rumelhart, Sejnowski, and Hinton each describe how these networks function and why they can behave in surprisingly rational ways. Marcus offers a critique of this approach.

Week 5: Vision

- Kitcher, P. *Marr's Computational Theory of Vision*
- Stefanucci, J.K. and Geuss, M.N. *Big People, Little World: The Body Influences Size Perception*
- Firestone, C. and Scholl, J. *"Top-Down" Effects Where None Should Be Found: The El Greco Fallacy in Perception Research*

Kitcher describes Marr's application of his 'levels of explanation' approach to the study of human vision. She shows how and why this program was one of the most significant breakthroughs in the study of perception in the history of cognitive science. We then turn to the question of the extent to which what we see is influenced by the rest of our cognitive system. Stefanucci and Geuss present some results that suggest that people view openings as narrower when they are holding long rods. Firestone and Scholl respond to this argument, aiming to defend the view that perception operates independently of our other cognitive systems.

Week 6: Are You as Smart as You Think You Are?

- Tversky, A. and Kahneman, D. *Judgement Under Uncertainty: Heuristics and Biases*
- Tversky, A. and Kahneman, D. *Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment*
- Gigerenzer, G. *How to Make Cognitive Illusions Disappear: Beyond "Heuristics and Biases"*

Tversky and Kahneman have famously shown that people are susceptible to a wide range of biases in their decision-making and inferential procedures. For example, they show that people are prone to violate laws of statistical inference. In particular, people can be led to view conjunctive sentences of the form 'P and Q' as *more* likely true than sentences expressing just one of the conjuncts ('P'), when the other conjunct ('Q') is deemed highly likely by their background beliefs. But it is a law of probability that P is always at least as likely as P and Q. They infer from this that human beings are in general much less rational they we might have thought. Gigerenzer argues that Tversky and Kahneman have misinterpreted the results of their experiments, and that these experiments instead point to different *kinds* of rationality.

Week 7: Competence and Performance Theories of Language

- Chomsky, N. *Aspects of a Theory of Syntax* (Chapter 1)
- Manning, C. *Probabilistic Syntax*

We turn now to one of the most (in-)famous distinctions from the history of linguistics and cognitive science. Noam Chomsky argues that it is essential for linguists to distinguish competence, the structure of an internal psychological system, from performance, the behavior of the organism whose system it is. Computational linguist Chris Manning responds by saying that, since performance provides the central evidential base for cognitive science, this distinction is illegitimate.

Week 8: Language as a Mental Organ

- Pietroski, P. and Crain, S. *The Language Faculty*
- Petitto, L. *How the Brain Begets Language*
- Goldin-Meadow, S. and Feldman, H. *The Development of Language-like Communication Without a Language Model*

Following Chomsky's work, various cognitive scientists have argued that natural language is not learned, in the way that we might learn the piano, but developed, in the way that biological organs like hearts develop. That is, they argue that human language should be viewed as provided by our innate genetic endowment. We look this week at three of the most powerful arguments for this surprising conclusion. Pietroski and Crain present the 'Poverty of the Stimulus' argument, which aims to show that the language we acquire is not made available by our developmental environment (the 'stimulus') and so must be provided by our own mind. Pettito describes the instinctive linguistic behavior of infants, and Goldin-Meadow and Feldman describe the remarkable linguistic capabilities of people who, in virtue of being deaf, have received no environmental linguistic input at all but nonetheless develop complex linguistic systems.

Week 9: Language as Learned

- Evans, N. and Levinson, S. *The Myth of Language Universals*
- Rumelhart, M.E. and McClelland *On Learning The Past Tenses of English Verbs*
- Cowie, F. *What's Within? Nativism Reconsidered* (Excerpts)

We then turn to some of the chief reasons for being skeptical about claims that language is innate. Evans and Levinson catalog the enormous variety observed across the world's languages, and argue that this much variation is incompatible with the claim that the language we speak is largely determined by our species-universal biological development. Rumelhart and McClelland produce a connectionist model which is capable of learning linguistic rules, about the formation of past participles in English, suggesting that language acquisition based on environmental patterns is possible. The philosopher Fiona Cowie provides a general discussion of the merits (or lack thereof) of the innateness hypothesis.

Week 10: Non-Human Minds

- Tomasello, M. and Herrmann, E. *Ape and Human Cognition: What's the Difference?*
- Cheney, R, and Seyfarth, D. *Baboon Metaphysics* (Excerpts)
- Tetzlaff, M. and Rey, G. *Systematicity and Intentional Realism in Honeybee Navigation*

At this point we broaden our target, so as to include a discussion of the differences and similarities between human minds and those of non-human animals. Tomasello and Herrmann argue that most of our cognitive capacities are shared with our nearest biological relatives, the great apes. They argue that we primarily differ with respect to certain kinds of social cognition, such as the desire and ability to share our mental states with others. Cheney and Seyfarth describe the surprisingly complex social cognition of baboons. Tetzlaff and Rey describe the 'waggle dance' of the honeybee, which enables creatures as simple as bees to communicate complex information about the distance and direction of food sources to one another using dances.

Week 11: Modularity

- Fodor, J. *The Modularity of Mind* (Excerpts)
- Karmiloff-Smith, A. *Beyond Modularity: A Developmental Perspective on Cognitive Science* (Excerpts)
- Prinz, J. *Is the Mind Really Modular?*

Jerry Fodor's hugely influential work argues that the mind should be viewed as containing a variety of broadly independent systems, each specialized for a particular task, such as vision or language. He also argues that these systems, or 'modules', are hardwired or determined by our biology. Karmiloff-Smith agrees that the mind is modular, but present an alternative story about how modules develop, according to which capacities can become modularized over time. Prinz rejects the entire picture, and argues for a non-modular picture of the mind.

Week 12: The Adapted Mind

- Tooby, J. and Cosmides, L. *Evolutionary Psychology: A Primer*
- Dupre, J. *Humans and Other Animals* (Excerpt)

One of the most prominent research programs that developed in response to Fodor's claims about modularity is Evolutionary Psychology. This is the attempt to explain contemporary human minds with reference to human evolutionary history. Tooby and Cosmides provide an overview of the evidence, methods, and results of this research program. Dupre provides a critique of the style of reasoning used in this program, arguing that even though the mind did, of course, evolve, this doesn't tell us much about how it currently behaves.

Week 13: Integrating Mind and Brain

- Churchland, P.S. and Sejnowski, T.J. *Perspectives on Cognitive Neuroscience*
- Nathan, M. and Del Pinal, G. *Mapping the Mind: Bridge Laws and the Psycho-Neural Interface*

Cognitive neuroscience is the attempt to integrate the broadly information-theoretic approach to psychology with developments of contemporary understanding of the brain. Churchland and Sejnowski, and Nathan and Del Pinal describe contemporary approaches within this research program, shedding light on the ways in which theories of the mind and of the brain can be mutually illuminating.

Week 14: Challenges to Computationalism

- van Gelder, T. *The Dynamical Hypothesis in Cognitive Science*
- Barrett, L. *Why Brains Are Not Computers, Why Behaviorism Is Not Satanism, and Why Dolphins Are Not Aquatic Apes*
- Clark, A. *Being There: Putting Brain, Body and World Together Again* (Excerpts)
- Brooks, R. *Intelligence Without Representation*

We have mostly been assuming that the mind can be adequately treated as a computer. We look this week at a variety of reasons that people have begun to question this assumption. van Gelder argues that treating the mind as a dynamical system, and thus using differential equations to explain its behavior, is a more profitable strategy than the standard symbolic systems approach. Barrett, Brooks, and Clark argue that abstracting away from the interactions organisms have with their environments, as many computational approaches do, undermines the prospect of a wide range of explanations. They argue for a cognitive science which is embodied, embedded, extended, and enactive: embodied in that cognitive processes rely on the relations between the brain and the rest of the organism, embedded in that agent-world interactions are crucial, extended in that objects in the environment are typically involved in cognitive processes, and enactive in that the mind functions to enable action, not merely passively respond to external stimuli.

Week 15: What is it Like?

- Chalmers, D. *The Puzzle of Conscious Experience*
- Gazzaniga, M. *The Split Brain in Man*
- Johnson, G. *Implicit Bias and the Domain of Consciousness*

We close the course with a discussion of perhaps the most difficult problem in cognitive science: why do our mental states *feel* a certain way? That we experience our mental states is, from a scientific perspective, deeply puzzling: why should electrico-chemical activity in a piece of flesh give rise to *experience* at all? Chalmers provides a general outline for this puzzle. Gazzaniga describes some very surprising experiments in which a single human being seems to have two distinct streams of conscious experience, raising even deeper puzzles about the relation between conscious experience and the physical brain. Johnson describes some recent work, on unconscious biases, suggesting that the boundary between conscious experience and mere physical processing may not be where we previously thought.