Course Guide

PHIL10134 The Computational Mind

Course Organiser: Dr Mark Sprevak (mark.sprevak@ed.ac.uk)

- Office Location: Dugald Stewart Building room 5.12
- Office hours: Book a slot to see me by clicking here: https://calendly.com/sprevak/office-hours-meeting
- If you cannot find a time that suits, email me to set up another time

---

Department of Philosophy
School of Philosophy, Psychology and Language Sciences
University of Edinburgh

1. Course Aims and Objectives

Computational theories of mind are our best theories of how the mind works. In this course, we will be looking at those computational theories from a philosophical point of view. We will ask foundational questions about the aim, nature, and prospects of these theories. We will ask such questions as:

1. What is a computation?
2. If the mind is a computer, what kind is it?
3. Is computation a real feature of brain, or a projection of our interests?
4. Can consciousness be explained by computation?
5. Are cognitive computations in the brain or do they spill into the environment?

We will gain a lively appreciation of how some of these issues play out using examples of computational models from psychology and cognitive neuroscience.

2. Seminar Content & Readings

For each week, readings are listed below. Readings include *core* and *secondary* readings. The core readings for each week are starred (*).

Core readings are the material that it is your responsibility to read before each class.

Core readings are also the material on which your weekly online discussions will be based (see below). Please do delve into the further reading too; these should be your first port of call when writing your essay. The core readings and as many as possible of the secondary readings are posted as PDFs on Learn.

Some hints: Read the core readings *carefully*. You may find an article challenging or difficult—persist! If you do not understand something, read it again, think about it, try to make sense of it in your own words. If after multiple attempts to make sense of a passage, you still cannot, then there is a good chance that you have identified a real problem in the article—a perfect point to raise in your discussion forum, in the class, or to form the basis of an excellent essay! Jim Pryor has some wonderful tips for reading philosophy (as he says, '*you should expect to read a philosophy article more than once*')

Background reading

The more background you know the better. A good starting point is to read one of the books listed below. Even if you already have a strong background in this topic already, I would encourage you to read one of these books during the semester to consolidate your knowledge.

[1] A. Clark. Mindware: An Introduction to Cognitive Science. Oxford University Press, Oxford, 2 edition, 2014.
[2] T. Crane. The Mechanical Mind. Routledge, London, 2 edition, 2003.

Week 1: Introduction to CTM

* Chapter 1 of A. Clark. Mindware: An Introduction to Cognitive Science. Oxford University Press, Oxford, 2 edition, 2001.
* Pages 134–148 of T. Crane. The Mechanical Mind. Routledge, London, 2 edition, 2003.

[1] J. L. Bermudez. Cognitive Science: An Introduction to the Science of the Mind. Cambridge University Press, Cambridge, 2010.
[2] N. Block. The mind as the software of the brain. In E. E. Smith and D. N. Osherson, editors, An Invitation to Cognitive Science, Vol. 3, Thinking, pages 377–425. MIT Press, Cambridge, MA, 1995.
[3] N. Chomsky. A review of B. F. Skinner's Verbal Behavior. Language, 35:26–58, 1959.
[4] B. J. Copeland. Artificial Intelligence: A Philosophical Introduction. Blackwell, Oxford, 1993.
[5] T. Crane. The Mechanical Mind. Routledge, London, 2 edition, 2003.
[6] M. Davies. Cognitive science. In F. Jackson and M. Smith, editors, The Oxford Handbook of Contem- porary Analytic Philosophy, chapter 14. Oxford University Press, 2005.
[7] J. Haugeland, editor. Mind Design II. MIT Press, Cambridge, MA, 1999.
[8] W. G. Lycan and J. Prinz, editors. Mind and Cognition. Blackwell, Oxford, 2 edition, 2008.
[9] H. Putnam. The nature of mental states. In Mind, Language and Reality, Philosophical Papers, vol. 2, pages 429–440. Cambridge University Press, Cambridge, 1975.
[10] D. M. Rosenthal, editor. The Nature of Mind. Oxford University Press, Oxford, 1991.
[11] J. J. C. Smart. Sensations and brain processes. Philosophical Review, 68:141–156, 1959.
[12] T. van Gelder. What might cognition be, if not computation? The Journal of Philosophy, 91:345–381, 1995.

Week 2: Classical models and mentalese

* Chapters 2 and 3 of J. Haugeland. Artificial Intelligence: The Very Idea. MIT Press, Cambridge, MA, 1985.

[1] M. Aydede. The language of thought. In E. N. Zalta, editor, The Stanford Encyclopedia of Philosophy. Fall 2010 edition, 2010.
[2] L. Barsalou. Perceptual symbol systems. Behavioral and Brain Sciences, 22:577–660, 1999.
[3] E. Camp. Thinking with maps. Philosophical Perspectives, 21:145–182, 2007.
[4] J. A. Fodor. The Language of Thought. The Harvester Press, Sussex, 1975.
[5] Chapter 1 of J. A. Fodor. Psychosemantics. MIT Press, Cambridge, MA, 1987.

[6] J. A. Fodor. LOT2: The Language of Thought Revisited. Oxford University Press, Oxford, 2008.
[7] J. Haugeland. Semantic engines: An introduction to mind design. In J. Haugeland, editor, Mind Design, pages 1–34. MIT Press, Cambridge, MA, 1981.
[8] S. Laurence and E. Margolis. Regress arguments against the language of thought. Analysis, 57:60–66, 1997.
[9] J. Prinz. Furnishing the Mind. MIT Press, Cambridge, MA, 2002.
[10] J. Prinz. Has mentalese earned its keep? on Jerry Fodor's LOT 2. Mind, 120:485–501, 2011.
[11] M. Rescorla. Cognitive maps and the language of thought. The British Journal for the Philosophy of Science, 60:377–407, 2009.
[12] S. Schneider. The language of thought. In P. Calvo and J. Symons, editors, Routledge Companion to Philosophy of Psychology, chapter 17. Routledge, London, 2009.

Week 3: Connectionist models

* 'Connectionism' in A. Clark. Mindware: An Introduction to Cognitive Science. Oxford University Press, Oxford, 2001.

[1] K. Aizawa. Explaining systematicity. Mind and Language, 12:115–136, 1997.
[2] M. Aydede. Language of thought: The connectionist contribution. Minds and Machines, 7:57–101, 1997.
[3] D. J. Chalmers. Syntactic transformations on distributed representations. Connection Science, 2:53–62,1990.
[4] D. J. Chalmers. Connectionism and compositionality: Why Fodor and Pylyshyn were wrong. Philosophical Psychology, 6:305–319, 1993.
[5] J. A. Fodor and B. P. McLaughlin. Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. Cognition, 35:183–204, 1990.
[6] J. A. Fodor and Z. W. Pylyshyn. Connectionism and cognitive architecture. Cognition, 28:3–71, 1988.
[7] R. F. Hadley. Cognition, systematicity and nomic necessity. Mind and Language, 12:137–153, 1997.
[8] G. E. Hinton. How neural networks learn from experience. Scientific American, 267:145–151, 1992.
[9] K. Johnson. On the systematicity of language and thought. The Journal of Philosophy, 101:111–139, 2004.
[10] G. F. Marcus. The Algebraic Mind. MIT Press, Cambridge, MA, 2003.
[11] R. J. Matthews. Can connectionists explain systematicity? Mind and Language, 12:154–177, 1997.
[12] P. Smolensky. On the proper treatment of connectionism. Behavioral and Brain Sciences, 11:1–74, 1988.
[13] M. S. C. Thomas and J. L. McClelland. Connectionist models of cognition. In R. Sun, editor, Cambridge Handbook of Computational Psychology, pages 23–58. Cambridge University Press, Cambridge, 2008.

[14] ‘Are syntactically structured representations needed?’ in W. Bechtel and A. Abrahamsen. Connectionism and the Mind: An Introduction to Parallel Processing in Networks. Oxford University Press, Oxford, 2 edition, 2002.

Week 4: Marr’s levels of explanation

* Chapter 1 of D. Marr. Vision. W. H. Freeman, San Francisco, CA, 1982.

[1] W. Bechtel and O. Shagrir. The non-redundant contributions of Marr’s three levels of analysis for explaining information processing mechanisms. Topics in Cognitive Science, forthcoming.
[2] D. Danks. Rational analyses, instrumentalism, and implementations. In N. Chater and M. Oaksford, editors, The Probabilistic Mind: Prospects for Rational Models of Cognition, chapter 3. Oxford University Press, 2008.
[3] D. C. Dennett. Cognitive science as reverse engineering: Several meanings of ”top-down” and ”bottom- up”. In D. Prawitz, B. Skyrms, and D. Westerstahl, editors, Proceedings of the 9th International Congress of Logic, Methodology and Philosophy of Science. Elsevier, 1994.
[4] M. Jones and B. C. Love. Bayesian fundamentalism or enlightenment? on the explanatory status and theoretical contributions of Bayesian models of cognition. Behavioral and Brain Sciences, 34:169–231, 2011.
[5] A. Newell. The knowledge level. Artificial Intelligence, 18:87–127, 1982.
[6] O. Shagrir. Marr on computational-level theories. Philosophy of Science, 77:477–500, 2010.
[7] O. Shagrir and W. Bechtel. Marr’s computational level and delineating phenomena. In D. Kaplan, editor, Integrating Psychology and Neuroscience: Prospects and Problems. Oxford University Press, Oxford, 2013.

Week 5: Computational implementation

* Chapter 9 of J. R. Searle. The Rediscovery of the Mind. MIT Press, Cambridge, MA, 1992.
* ‘Appendix’ in H. Putnam. Representation and Reality. MIT Press, Cambridge, MA, 1988.

[1] R. L. Chrisley. Why everything doesn’t realize every computation. Minds and Machines, 4:310–333, 1995.
[2] B. J. Copeland. What is computation? Synthese, 108:335–359, 1996.
[3] E. Dresner. Measurement-theoretic representation and computation-theoretic realization. The Journal of Philosophy, 107:275–292, 2010.
[4] G. Piccinini. Computational modelling vs. computational explanation: Is everything a Turing Machine, and does it matter to the philosophy of mind? Australasian Journal of Philosophy, 85:93–115, 2007.
[5] G. Piccinini. Computing mechanisms. Philosophy of Science, 74:501–526, 2007.
[6] G. Piccinini. Computers. Pacific Philosophical Quarterly, 89:32–73, 2008.
[7] G. Piccinini. Computation in physical systems. In E. N. Zalta, editor, The Stanford

Encyclopedia of Philosophy. Fall 2010 edition, 2010.
[8] J. R. Searle. Is the brain a digital computer? Proceedings and Addresses of the American Philosophical Association, 64:21–37, 1990.
[9] O. Shagrir. The rise and fall of computational functionalism. In Y. Ben-Menahem, editor, Contemporary Philosophy in Focus: Hilary Putnam, pages 220–250. Cambridge University Press, 2005.

Week 6: Chalmers on implementation

* D. J. Chalmers. A computational foundation for the study of cognition. Journal of Cognitive Science, 12:323–357, 2012.

[1] C. Brown. Implementation and indeterminacy. In Selected papers from conference on Computers and Philosophy—Volume 37, CRPIT '03, pages 27–31, Darlinghurst, Australia, 2003. Australian Computer Society, Inc.
[2] C. Brown. Combinatorial-state automata and models of computation. Journal of Cognitive Science, 13:51–73, 2012.
[3] D. J. Chalmers. Does a rock implement every finite-state automaton. Synthese, 108:309–333, 1996.
[4] D. J. Chalmers. The varieties of computation: A reply. Journal of Cognitive Science, 13:211–248, 2012.
[5] F. Egan. Metaphysics and computational cognitive science: Let's not let the tail wag the dog. Journal of Cognitive Science, 13:39–49, 2012.
[6] D. Gamez. Are there functional correlates of consciousness? manuscript, 2013.
[7] P. Godfrey-Smith. Triviality arguments against functionalism. Philosophical Studies, 145:273–295, 2009.
[8] S. Harnad. The causal topography of cognition. Journal of Cognitive Science, 13:181–196, 2012.
[9] C. Klein. Two paradigms for individuating implementations. Journal of Cognitive Science, 13:167–179, 2012.
[10] M. Milkowski. Beyond formal structure: A mechanistic perspective on computation and implementation. Journal of Cognitive Science, 12:359–379, 2011.
[11] G. O'Brien. Defending the semantic conception of computation in cognitive science. Journal of Cognitive Science, 12:381–399, 2011.
[12] M. Rescorla. How to integrate representation into computational modelling, and why we should. Journal of Cognitive Science, 13:1–38, 2012.
[13] J. B. Richie. Chalmers on implementation and computational sufficiency. Journal of Cognitive Science, 12:401–417, 2011.
[14] M. Scheutz. What it is not to implement a computation: A critical analysis of Chalmers' notion of implementation. Journal of Cognitive Science, 13:75–106, 2012.
[15] O. Shagrir. Can a brain possess two minds? Journal of Cognitive Science, 13:145–165, 2012.
[16] M. Sprevak. Three challenges to Chalmers on computational implementation. Journal of Cognitive Science, 13:107–143, 2012.
[17] B. N. Towl. Home, pause, or break: A critique of Chalmers on implementation.

Journal of Cognitive Science, 12:419–433, 2011.

Week 7: Essay writing workshop

* Readings for this week are the past essays of students posted on Learn.

Week 8: Computational theories of consciousness

* T. Maudlin. Computation and consciousness. The Journal of Philosophy, 86:407–432, 1989.

[1] M. V. Antony. Against functionalist theories of experience. Mind and Language, 9:105–123, 1994.
[2] E. Barnes. The causal history of computational activity: Maudlin and Olympia. The Journal of Philosophy, 88:304–316, 1991.
[3] G. Bartlett. Computational theories of conscious experience: between a rock and a hard place. Erkenntnis, 76:195–209, 2012.
[4] C. Klein. Dispositional implementation solves the superfluous structure problem. Synthese, 165:141–153, 2008.
[5] C. Klein. Olympia and other O-machines. manuscript, 2012.
[6] M. Muhlestein. Counterfactuals, computation, and consciousness. Cognitive Computation, 5:99–105, 2013.

Week 9: Internalism vs externalism

* T. Burge. Individualism and psychology. Philosophical Review, 95:3–45, 1986.

[1] T. Bontly. Individualism and the nature of syntactic states. The British Journal for the Philosophy of Science, 49:557–574, 1998.
[2] K. Butler. Content, computation, and individuation in vision theory. Analysis, 56:146–154, 1996.
[3] F. Egan. Must psychology be individualistic? Philosophical Review, 100:179–203, 1991.
[4] F. Egan. Individualism and vision theory. Analysis, 54:258–264, 1994.
[5] F. Egan. Computation and content. Philosophical Review, 104:181–204, 1995.
[6] F. Egan. In defence of narrow mindedness. Mind and Language, 14:177–194, 1999.
[7] F. Egan. Naturalistic inquiry: Where does mental representation fit in? In L. M. Antony and N. Hornstein, editors, Chomsky and his Critics, chapter 4. Blackwell, Oxford, 2003.
[8] F. Egan. Computational models: a modest role for content. Studies in History and Philosophy of Science, 41:253–259, 2010.
[9] C. Peacocke. Content, computation and externalism. Mind and Language, 9:303–335, 1994.
[10] C. Peacocke. Computation as involving content: A response to Egan. Mind and Language, 14:195–202, 1999.

[11] G. Segal. Defence of a reasonable individualism. Mind, 100:485–493, 1991.
[12] O. Shagrir. Content, computation and externalism. Mind, 110:369–400, 2001.
[13] L. Shapiro. A clearer vision. Philosophy of Science, 64:131–153, 1997.
[14] R. A. Wilson. Wide computationalism. Mind, 103:351–372, 1994.

Week 10: Mechanistic accounts

* G. Piccinini. Computing Mechanisms. Philosophy of Science 74/4: 501-26, 2007.

[1] M. Chirimuuta. Minimal models and canonical neural computations. Synthese 191: 127-53, 2014.
[2] D. Coelho Mollo. Functional individuation, mechanistic implementation. Synthese 2017.
[3] S. Glennan. Rethinking Mechanistic Explanation. Philosophy of Science 64: 342-53, 2002.
[4] S. Haimovici. A Problem for the Mechanistic Account of Computation. Journal of Cognitive Science 14: 151-181, 2013.
[5] D. Kaplan. Explanation and Description in Computational Neuroscience. Synthese 183: 339-73, 2011.
[6] P. Machamer, L. Darden, & C. Craver. Thinking about Mechanisms. Philosophy of Science 67: 1-25, 2000.
[7] M. Miłkowski. Beyond Formal Structure: A Mechanistic Perspective on Computation and Implementation. Journal of Cognitive Science 12: 359-79, 2011.
[8] M. Miłkowski. Explaining the Computational Mind. Cambridge, MA: MIT Press, 2013.
[9] G. Piccinini. Functionalism, Computationalism, and Mental Contents. Canadian Journal of Philosophy 34: 375-410, 2004.
[10] G. Piccinini. Functionalism, Computationalism, and Mental States. Studies in History and Philosophy of Science 35: 811-33, 2004.
[11] G. Piccinini. Computation without Representation. Philosophical Studies 137/2: 205-41, 2008.
[12] G. Piccinini. Physical Computation. Oxford: OUP, 2015.

Week 11: Hypercomputation

* G. Piccinini. The physical Church–Turing Thesis: Modest or bold? The British Journal for the Philosophy of Science, 62:733–769, 2011.

[1] A. Boucher. Parallel machines. Minds and Machines, 7:543–551, 1997.
[2] T. Button. SAD computers and two versions of the Church-Turing Thesis. The British Journal for the Philosophy of Science, 60:765–792, 2009.
[3] B. J. Copeland. Turing's O-machines, Searle, Penrose and the brain. Analysis, 58:128–138, 1998.
[4] B. J. Copeland. Narrow versus wide mechanism. The Journal of Philosophy,

97:5–32, 2000.
[5] B. J. Copeland. Hypercomputation: philosophical issues. Theoretical Computer Science, 317:251–267, 2004.
[6] B. J. Copeland. The Church-Turing thesis. In E. N. Zalta, editor, The Stanford Encyclopedia of Philosophy. Fall 2008 edition, 2008.
[7] B. J. Copeland and O. Shagrir. Do accelerating Turing machines compute the uncomputable? Minds and Machines, 21:221–239, 2011.
[8] M. L. Hogarth. Non-Turing computers and non-Turing computability. In D. Hull, M. Forbes, and R. M. Burian, editors, PSA 1994, volume 1, pages 126–138. Philosophy of Science Association, East Lansing, MI, 1994.
[9] M. L. Hogarth. Deciding arithmetic using SAD computers. The British Journal for the Philosophy of Science, 55:681–691, 2004.
[10] I. Nemeti and G. David. Relativistic computers and the Turing barrier. Applied Mathematics and Computation, 178:118–142, 2006.
[11] T. Ord. Hypercomputation: Computing more than the Turing machine. Master's thesis, University of Melbourne, Melbourne, 2002.
[12] T. Ord. The many forms of hypercomputation. Applied Mathematics and Computation, 178:143–153, 2006.
[13] O. Shagrir. Two dogmas of computationalism. Minds and Machines, 7:321–344, 1997.

Autonomous Learning Groups (ALGs)

Here are 3 things that you might like to do in an ALG for this course:

1. Watch Christof Koch's YouTube lecture on computation and the brain: https://www.youtube.com/watch?v=dwOGh6qkzG4. Some questions to think about:

    a. Why is it useful to think of the brain as computing?
    b. Which bits of the brain compute?
    c. What assumptions are being made by this approach to brain function that could later be proven wrong? What kinds of empirical evidence could count against it?

2. Read and discuss B. J. Copeland. Narrow versus wide mechanism. The Journal of Philosophy, 97:5–32, 2000. Some questions to think about:

    a. What is the difference between a computation and a mechanism?
    b. What might mechanisms do that computations cannot?
    c. What consequences does this have for CTM?

3. Read and discuss Introduction and Chapter 1 of J. A. Fodor. The Mind Doesn't Work That Way. MIT Press, Cambridge, MA, 2000. Some questions to think about:

    a. What is RTM and CTM? Why is the CTM so good, in Fodor's view?
    b. What is the New Synthesis?

c. What are syntactic properties, according to Fodor?