

# Computational [Principles of] Psychology\*

Shimon Edelman  
Department of Psychology  
Cornell University  
Ithaca, NY 14853

<http://kybele.psych.cornell.edu/~edelman>

Spring 2018

## Contents

<b>1 Overview</b>	<b>2</b>
<b>2 Notes for participants</b>	<b>2</b>
2.1 Diversity and inclusion . . . . .	2
2.2 Learning goals and opportunities . . . . .	3
2.3 Best learning practices . . . . .	4
2.4 Credit and grading . . . . .	5
<b>3 Lecture topics by date</b>	<b>6</b>
3.1 Week 1 (1/23; 1/25): introduction and overview . . . . .	6
3.2 Week 2 (1/30; 2/1): methodology . . . . .	7
3.3 Week 3 (2/6; 2/8): universal tools, I: measurement and representation . . . . .	8
3.4 Week 4 (2/13; 2/15): universal tools, II: probability and the Ace of Bayes . . . . .	9
3.5 Week 5 (2/20; 2/22): February break; 1st prelim . . . . .	10
3.6 Week 6 (2/27; 3/1): universal tools, III: similarity; learning and generalization . . . . .	11
3.7 Week 7 (3/6; 3/8): universal tools, IV: dimensionality; memory . . . . .	12
3.8 Week 8 (3/13; 3/15): actions and consequences . . . . .	13
3.9 Week 9 (3/20; 3/22): higher cognition . . . . .	14
3.10 Week 10 (3/27; 3/29): higher cognition (cont.); 2nd prelim . . . . .	15
3.11 Week 11 (4/10; 4/12): intro to neural computation; neurons, I . . . . .	16
3.12 Week 12 (4/17; 4/19): neurons, II . . . . .	17
3.13 Week 13 (4/24; 4/26): neurons, III; brains . . . . .	18
3.14 Week 14 (5/1; 5/3): advanced topics . . . . .	19
3.15 Week 15 (5/8): wrap-up . . . . .	20
<b>4 Computational concepts, principles, and methods: an abecedary</b>	<b>20</b>

---

\*Psych 3140/6140. Spring semester, Tuesdays and Thursdays.

# 1 Overview

This course states, motivates, and offers detailed support for the observation that cognition is fundamentally a computational process [30]. Students are introduced to a number of conceptual tools for thinking about natural behavior and the cognitive information processing that underlies it, including statistical learning from experience and the use of patterns distilled from past experience in guiding future actions. The application of these tools to the understanding of natural minds and to advancing the goals of artificial intelligence is illustrated on selected examples drawn from the domains of perception, memory, motor control, action planning, problem solving, decision making, reasoning, and creativity.

The material is conceptually advanced and moderately to highly technical. It is aimed at advanced undergraduate students, as well as graduate students from psychology, neurobiology, computer science, and other cognitive sciences. Prior exposure to statistical concepts and the scientific method is essential.

## Readings

The recommended textbook is *Computing the Mind: How the Mind Really Works* (Oxford University Press, 2008). Additional readings (a zipped collection of PDFs) are available on the course Blackboard site.

## How to use this syllabus

- For each week, there's a list of readings with references. Some of the readings are required, others are optional. The references are also listed at the end of the syllabus, alphabetically by first author.
- For an alphabetical roster of select key ideas and topics, see section 4.

## Important dates

- **First prelim:** Thursday, 2/22, in class.
- **Second prelim:** Thursday, 3/29, in class.
- **Final exam:** TBA (set by the college).

# 2 Notes for participants


This section contains essential information for participants: the inclusion statement,<sup>1</sup> learning goals and practices, and credit requirements.

## 2.1 Diversity and inclusion

*Computational Psychology* is more diverse than many other courses at Cornell, in at least two respects. First, it purposely ignores the traditional disciplinary boundaries, as cognitive science has done since its inception (see the illustration on the next page). Accordingly, my plan is for us to freely mix concepts and topics from psychology, mathematics, computer science, and neurobiology. Second, this course does not respect college and program boundaries: historically, it has been attended and successfully completed by students from different

---

<sup>1</sup>The remarks in section 2.1, which are specific to this course, are intended to supplement the official Cornell statement on diversity and inclusion, which covers dimensions such as gender, race, socio-economic background, etc., and which can be found here: <http://diversity.cornell.edu/>.



On the left, a group photograph of the (disciplinarily, but not otherwise, diverse) members of the Hixon Symposium, held in 1948, which helped shape the modern integrative approach to brain/mind science. *Left to right:* (seated) Halstead, Lashley, Klüver, Köhler, Lorente de Nó; (standing) Brosin, Jeffres, Weiss, Lindsley, von Neumann, Nielsen, Gerard, Liddell. On the right, the cover of *Embodiments of Mind* (1965; [75]), a still readable collection of papers by Warren McCulloch, the one member of the group who is absent from the group photo.

colleges and a variety of majors, among them psychology, neurobiology, engineering, information science, computer science, and English, as well as by graduate and professional masters students.

This diversity makes inclusion particularly important; interestingly, it also makes it easier to attain, because we can always “triangulate,” and thus better understand, every issue and concern from multiple disciplinary and personal points of view. In such circumstances, the more each of you contributes to the discussion, the more we all learn. Both myself and the TA are fully committed to having everyone’s questions and concerns heard — both in class and outside — and to helping everyone succeed who is willing to invest the effort to do so. In addition to our office hours, we are always available for meetings by appointment; we will also answer promptly any questions posted on Blackboard or emailed to us. Let us know if at any point during the semester you have suggestions for making things work better; these may include additions to the present statement, which will be included in periodic revisions, to be shared on Blackboard.

## 2.2 Learning goals and opportunities

Active participation in this course should improve significantly your understanding of how the brain/mind works, regardless of your disciplinary and personal background. In addition, upon its completion, you should be able to:

- reflect on how different disciplines provide complementary insights into how the brain/mind works, and how scholarship that cuts across disciplines can help integrate those insights;

- reflect on how people with different personal and educational backgrounds can join forces in solving complex problems.

Some of the instructional strategies and learning opportunities that we plan to use and make available to promote these outcomes are as follows:

- The key conceptual framework that is stated, discussed, and applied in this course — the “levels of understanding” of complex information-processing systems (introduced in week 2) — is inherently geared toward integrating multiple sources of data, methods of analysis, and formal models.
- The examples that will be used in class to illustrate various concepts will be drawn from a variety of naturally diverse real-life situations.
- You will have opportunities to come up with your own examples, thus adding to the diversity of viewpoints, enriching the abstract mathematical and computational concepts that are necessarily impersonal.

### 2.3 Best learning practices

In creating this course and keeping it up to date, I have in a sense been reprising the trajectory of my own multidisciplinary educational background and research experience. For my undergraduate degree, I studied electrical engineering. Following a stint in the military, I went back to school for an MSc and then a PhD in computer science, so as to be able to work at the intersection of mathematics, philosophy, art, and science that Doug Hofstadter described so alluringly in his Pulitzer-winning *Gödel, Escher, Bach* [53] (in cognitive science and AI, this became known eventually as “the book that launched a thousand careers”). The main text for this course, *Computing the Mind* [30], has been written with that great book as an inspiration.

The readings that are drawn from the textbook are specified by chapter, section, and subsection, but ideally the book should be read in its entirety. The key concepts are marked in the text by SMALL CAPITALS and are usually included in the index. The many margin notes may be ignored on the first reading to avoid distraction, or you may want to take them in for an extra helping of (sometimes only marginally relevant) information or an entertaining flight of fancy.

Additional readings come from a selection of more than 100 academic papers, some of which are decades-old classics and some published less than a year ago. Don’t panic! Only a few of the papers, listed in the references at the end of this document, are required reading. These are clearly marked in the week-by-week syllabus. Furthermore, for some of the required papers, it will suffice to read only a few select sections. If in doubt, please feel free to ask (in class or on Blackboard).

When reading, always focus on the “big picture” and the key concepts and principles. Concepts and principles are also what I shall be trying to bring out during the lectures, which is why it is important that you attend. Neither technical details (such as proofs of mathematical statements) nor trivia (such as the names of researchers, with the exception of a few luminaries) will appear in the exams. The two prelims and the final will each consist of several multiple-choice questions and several micro-essays on assigned topics. A mock exam will be posted ahead of the first prelim. Please bring any concerns about learning practices to my office hours (or to the TA’s).

## 2.4 Credit and grading

To get credit for this course:

1. Meet with me (during office hours or by appointment) at some point during the first three weeks of classes.
2. Attend the lectures; ask questions and participate in the discussion. If you must miss a lecture, please send along an email with an explanation.
3. Prepare brief summaries of the weekly readings. Each summary should be one (short) paragraph long and should be accompanied by a question based on the readings. Altogether 10 of these should be submitted by the end of the semester (via email, to the TA). Each one will be graded on the great / good / not good basis (2/1/0 points).
4. Take the two prelims and the final exam.
  - Optionally, substitute a 1000-word essay for one of the prelims. The topic must be coordinated with me *ahead of time* and submitted on the day of the prelim via email to the TA.
5. (Graduate students only, enrolled in Psych 6140.) Submit a written summary of your impressions and lessons from the material, in an essay form (about 1000 words), by noon, Tuesday, May 15, via email to the instructor.

Final grade components:

---

First prelim:	20%
Second prelim:	20%
Weekly write-ups:	20%
Final exam:	40% (grads: final 30% + essay 10%)

---

Up to 3 extra credits can be accrued by participating in the Psychology Department subject pool (SONA). These will be added to the final score before it is converted into a letter grade.

The list of lecture topics and readings by week number and date begins on the next page. There is one page per week, with the “Primary” (required) and “Other” (optional) readings listed separately. I have also included a few photographs of the authors of the papers, to offer you a glimpse of the collective face of cognitive science; bibliographical entries for papers thus distinguished are marked in **boldface**. Advanced readings (entirely optional) are marked in *gray*.

### 3 Lecture topics by date

#### 3.1 Week 1 (1/23; 1/25): introduction and overview

1.1 [No meeting — the semester hasn't started yet.]

1.2 The subject matter of psychology. The fundamentality of computation [30].

Examples:

- perception — *lightness*, or estimating surface reflectance from images;
- thinking — *planning*, or estimating the actions needed to attain a goal);
- action — *motor control*, or estimating the signals that need to be sent to the muscles to execute the desired motion.

The nature of behavior:


- Dewey on the “reflex arc” concept [26];
- Thurstone on the “stimulus-response” fallacy [107].

A quick overview of computation:

- dynamical systems [57, sections 1,2];
- Turing Machines [5].

#### Primary readings

- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapters 1,2.
- [26] **J. Dewey. The reflex arc concept in psychology. *Psychological Review*, 3:357–370, 1896.**
- [107] L. L. Thurstone. The stimulus-response fallacy in psychology. *Psychological Review*, 30:354–369, 1923.



John Dewey

#### Other readings

- [111] E. P. Wigner. The unreasonable effectiveness of mathematics in the natural sciences. *Comm. Pure Appl. Math.*, XIII:1–14, 1960.
- [5] D. Barker-Plummer. Turing machines. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. 2007. Available online at <http://plato.stanford.edu/archives/win2007/entries/turing-machine/>.
- [57] S. Hotton and J. Yoshimi. Extending dynamical systems theory to model embodied cognition. *Cognitive Science*, 35:444–479, 2010, sections 1,2.

## 3.2 Week 2 (1/30; 2/1): methodology

### 2.1 The blind men and the elephant.

Four case studies: a computational theory [2]; a single-cell electrophysiology study [94]; an imaging study [51]; an engineering hack [76] (to be revisited later in the semester [35]).

### 2.2 Open your eyes!

- The general methodology: the Marr-Poggio program for neurosciences [30, ch.4].
- A worked-out example: sound localization in the barn owl [64, 74, 52].

## Primary readings

- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapters 3,4.
- [72] D. Marr and T. Poggio. From understanding computation to understanding neural circuitry. *Neurosciences Res. Prog. Bull.*, 15:470–488, 1977.
- [64] P. X. Joris, P. H. Smith, and T. C. T. Yin. Coincidence detection in the auditory system: 50 years after Jeffress. *Neuron*, 21:1235–1238, 1998.

## Other readings

- [2] J. R. Anderson. ACT: A simple theory of complex cognition. *American Psychologist*, 51:355–365, 1996.
- [76] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- [35] S. Edelman. The minority report: some common assumptions to reconsider in the modeling of the brain and behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, 28:751–776, 2016.
- [32] S. Edelman. Vision, reanimated and reimagined. *Perception*, 41:1116–1127, 2012. Special issue on Marr’s *Vision*.
- [51] J. V. Haxby, M. I. Gobbini, M. L. Furey, A. Ishai, J. L. Schouten, and P. Pietrini. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293:2425–2430, 2001.
- [94] C. D. Salzman, K. H. Britten, and W. T. Newsome. Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 346:174–177, 1990.
- [52] Y. Hazan, Y. Kra, I. Yarin, H. Wagner, and Y. Gutfreund. Visual-auditory integration for visual search a behavioral study in barn owls. *Frontiers in Integrative Neuroscience*, 9(11):1–12, 2015
- [41] K. Embar, S. Mukherjee, and B. P. Kotler. What do predators really want? The role of gerbil energetic state in determining prey choice by Barn Owls. *Ecology*, 95: 280–285, 2014
- [74] **C. Massa, F. M. Gabelli, and G. R. Cueto. Using GPS tracking to determine movement patterns and foraging habitat selection of the common barn-owl (*Tyto alba*). *Hornero*, 30:7–12, 2015**



Carolina Massa

### 3.3 Week 3 (2/6; 2/8): universal tools, I: measurement and representation

- 3.1 Perceptual measurement [30, ch.5]. Channel coding and hyperacuity [100]. The fundamental uncertainty [96] and a way around it [104].
- 3.2 Perceptual representation spaces [40]. The face space [63, 43, 78].

#### Primary readings

- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapter 5.
- [40] S. Edelman, K. Grill-Spector, T. Kushnir, and R. Malach. Towards direct visualization of the internal shape representation space by fMRI. *Psychobiology*, 26: 309–321, 1998.
- [63] **F. Jiang, V. Blanz, and A. J. O’Toole. Probing the visual representation of faces with adaptation: A view from the other side of the mean. *Psychological Science*, 17:493–500, 2006.**



Alice O’Toole

#### Other readings

- [1] R. A. Altes. Ubiquity of hyperacuity. *J. Acoust. Soc. Am.*, 85:943–952, 1988.
- [100] H. P. Snippe and J. J. Koenderink. Discrimination thresholds for channel-coded systems. *Biological Cybernetics*, 66:543–551, 1992.
- [96] R. Shahbazi, R. Raizada, and S. Edelman. Similarity, kernels, and the fundamental constraints on cognition. *Journal of Mathematical Psychology*, 70:21–34, 2016.
- [104] R. Talmon and R. R. Coifman. Intrinsic modeling of stochastic dynamical systems using empirical geometry. *Applied and Computational Harmonic Analysis*, 39:138–160, 2015.
- [43] W. A. Freiwald, D. Y. Tsao, and M. S. Livingstone. A face feature space in the macaque temporal lobe. *Nature Neuroscience*, 12:1187–1199, 2009.
- [78] A. Nestor, D. C. Plaut, and M. Behrmann. Feature-based face representations and image reconstruction from behavioral and neural data. *Proceedings of the National Academy of Science*, 113:416–421, 2016.



### 3.4 Week 4 (2/13; 2/15): universal tools, II: probability and the Ace of Bayes

- 4.1 A probabilistic formulation of cognition [17]. The Bayesian framework [49].
- 4.2 The Bayesian approach, applied to lightness perception [10].

#### Primary readings

- [17] N. Chater, J. B. Tenenbaum, and A. Yuille. Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, 10:287–291, 2006.
- [49] T. L. Griffiths and A. Yuille. Technical introduction: A primer on probabilistic inference. *Trends in Cognitive Sciences*, 10, 2006. Supplementary material. DOI: 10.1016/j.tics.2006.05.007.
- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, Appendix A.
- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapter 5.
- [58] **A. Hurlbert. Colour constancy. *Current Biology*, 17:R906–R907, 2007.**
- [10] D. H. Brainard and W. T. Freeman. Bayesian color constancy. *J. Opt. Soc. Am. A*, 14:1393–1411, 1997.



Anya Hurlbert

#### Other readings

- [110] M. Welling. Are ML and statistics complementary?, 2015. Roundtable discussion at the 6th IMS-ISBA meeting on Data Science in the next 50 years.
- [46] Z. Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521:452–459, 2015.
- [66] D. Knill and W. Richards, editors. *Perception as Bayesian Inference*. Cambridge University Press, Cambridge, 1996.
- [42] T. Evgeniou, M. Pontil, and T. Poggio. Regularization networks and support vector machines. *Advances in Computational Mathematics*, 13:1–50, 2000.
- [28] S. Edelman. Representation is representation of similarity. *Behavioral and Brain Sciences*, 21:449–498, 1998.
- [105] J. B. Tenenbaum and T. L. Griffiths. Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, 24:629–641, 2001.

### 3.5 Week 5 (2/20; 2/22): February break; 1st prelim

5.1 [no meeting — February break]

Read: Artificial Intelligence — **David Marr**'s personal view [73].

View: *Artificial Intelligence and the “Barrier of Meaning”*, a talk by **Melanie Mitchell** (2017, Cornell AI seminar).

5.2 **First prelim.**

#### Primary readings

[73] D. C. Marr. Artificial Intelligence — a personal view. *Artificial Intelligence*, 9:37–48, 1977.

#### Other readings

[33] S. Edelman. *The Happiness of Pursuit*. Basic Books, New York, NY, 2012



David Marr



Melanie Mitchell

### 3.6 Week 6 (2/27; 3/1): universal tools, III: similarity; learning and generalization


- 6.1 Similarity and generalization [30, ch.5]. Shepard’s Law [97, up to section “Derivation. . .” only]. Representation of similarities [28].
- 6.2 Learning from examples as function approximation [89]. The challenges of function approximation. Ill-posedness and regularization [6].

#### Primary readings

- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapter 5.
- [97] R. N. Shepard. Toward a universal law of generalization for psychological science. *Science*, 237:1317–1323, 1987.
- [28] S. Edelman. Representation is representation of similarity. *Behavioral and Brain Sciences*, 21:449–498, 1998.
- [89] T. Poggio and S. Edelman. A network that learns to recognize three-dimensional objects. *Nature*, 343: 263–266, 1990.

#### Other readings

- [98] **R. N. Shepard and S. Chipman. Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology*, 1:1–17, 1970.**
- [6] M. Bertero, T. Poggio, and V. Torre. Ill-posed problems in early vision. *Proceedings of the IEEE*, 76: 869–889, 1988.
- [65] D. Kersten, P. Mamassian, and A. Yuille. Object perception as Bayesian inference. *Annual Review of Psychology*, 55:271–304, 2004.
- [22] F. Cutzu and S. Edelman. Faithful representation of similarities among three-dimensional shapes in human vision. *Proceedings of the National Academy of Science*, 93:12046–12050, 1996.
- [38] S. Edelman and R. Shahbazi. Renewing the respect for similarity. *Frontiers in Computational Neuroscience*, 6:45, 2012.
- [34] S. Edelman. Varieties of perceptual truth and their possible evolutionary roots. *Psychonomic Bulletin and Review*, 22:1519–1522, 2015. doi: 10.3758/s13423-014-0741-z



Susan Chipman

### 3.7 Week 7 (3/6; 3/8): universal tools, IV: dimensionality; memory

- 7.1 Managing dimensionality. Dimensionality reduction by random projections [29, 38].
- 7.2 Managing the storage and retrieval of examples. Associative memory [30, ch.6]. Locality-sensitive hashing [4, 38].

#### Primary readings

- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapter 6.
- [38] S. Edelman and R. Shahbazi. Renewing the respect for similarity. *Frontiers in Computational Neuroscience*, 6:45, 2012.
- [4] A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Communications of the ACM*, 51:117–122, 2008.

#### Other readings

- [29] S. Edelman. *Representation and recognition in vision*. MIT Press, Cambridge, MA, 1999.
- [21] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20: 273–297, 1995.



Corinna Cortes

### 3.8 Week 8 (3/13; 3/15): actions and consequences


- 8.1 Managing action. The basics of motor control [112]. Bayesian motor decision making [67].
- 8.2 Action and reward. Reinforcement learning (RL) [114] and its relationship to Bayes. Hierarchical RL [8, 9].

#### Primary readings

- [112] S. P. Wise and R. Shadmehr. Motor control. In V. S. Ramachandran, editor, *Encyclopedia of the Human Brain*, volume 3, pages 137–157. Academic Press, San Diego, CA, 2002.
- [67] K. P. Körding and D. M. Wolpert. Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, 10:319–326, 2006.
- [114] D. M. Wolpert, J. Diedrichsen, and J. R. Flanagan. Principles of sensorimotor learning. *Nature Reviews Neuroscience*, 12:739–751, 2011.
- [8] M. M. Botvinick. Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences*, 12:201–208, 2008.
- [9] M. M. Botvinick, Y. Niv, and A. C. Barto. Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, 113:262–280, 2009.

#### Other readings

- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapter 6.
- [113] W. Woergoetter and B. Porr. Reinforcement learning. *Scholarpedia*, 3(3):1448, 2007.
- [23] P. Dayan. How to set the switches on this thing. *Current Opinion in Neurobiology*, 22:1068–1074, 2012.
- [15] N. Chater. Rational and mechanistic perspectives on reinforcement learning. *Cognition*, 113:350–364, 2009.
- [46] Z. Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521:452–459, 2015.
- [103] **A. Solway, C. Diuk, N. Córdova, D. Yee, A. G. Barto, Y. Niv, and M. M. Botvinick. Optimal behavioral hierarchy. *PLOS Computational Biology*, 10(8):e1003779, 2014.**



Debbie Yee

### 3.9 Week 9 (3/20; 3/22): higher cognition

9.1 Graphical models (Bayesian Networks) and reasoning [30, ch.8].

9.2 Induction [106]. General intelligence and IQ [79].

#### Primary readings

- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapter 8.
- [106] J. B. Tenenbaum, C. Kemp, T. L. Griffiths, and N. D. Goodman. How to grow a mind: statistics, structure, and abstraction. *Science*, 331:1279–1285, 2011.
- [79] R. E. Nisbett, J. Aronson, C. Blair, W. Dickens, J. Flynn, D. F. Halpern, and E. Turkheimer. Intelligence: new findings and theoretical developments. *American Psychologist*, 2012.

#### Other readings

- [84] J. Pearl. Structural counterfactuals: A brief introduction. *Cognitive Science*, 37: 977–985, 2013.
- [48] **L. S. Gottfredson. Life, death, and intelligence. *Journal of Cognitive Education and Psychology*, 1:23–46, 2004.**



Linda Gottfredson

### 3.10 Week 10 (3/27; 3/29): higher cognition (cont.); 2nd prelim

10.1 Problem solving [30, ch.8]. Analogy [55] and creativity [54, 7].


10.2 **Second prelim.**

#### Primary readings

- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapter 8.
- [54] D. R. Hofstadter. Variations on a theme as the crux of creativity. In *Metamagical Themas*, chapter 12, pages 232–259. Viking, Harmondsworth, England, 1985.
- [55] D. R. Hofstadter. Analogy as the core of cognition. In D. Gentner, K. J. Holyoak, and B. N. Kokinov, editors, *The Analogical Mind: Perspectives from Cognitive Science*, pages 499–538. MIT Press, Cambridge MA, 2001.

#### Other readings

- [7] **M. A. Boden. Precis of “The creative mind: Myths and mechanisms”. *Behavioral and Brain Sciences*, 17:519–570, 1994.**



Margaret Boden

### 3.11 Week 11 (4/10; 4/12): intro to neural computation; neurons, I

11.1 Introduction to neural computation. Brains and neurons. Cortical receptive fields (RFs), maps, and hierarchies [30, ch.2,3].

11.2 What do neurons do?

Projection:

- neurons as readout devices [13];
- neurons and projection pursuit [19].
- neurons, random projections, and similarity [38, 70].

Kernels:

- landmarks in representation spaces, similarity, and kernels [96].

#### Primary readings

- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008, chapters 2,3.
- [13] G. Buzsáki. Neural syntax: cell assemblies, synapsembles, and readers. *Neuron*, 68:362–385, 2010.
- [19] L. N. Cooper and M. F. Bear. The BCM theory of synapse modification at 30: interaction of theory with experiment. *Nature Reviews Neuroscience*, 13:798–810, 2012.

#### Other readings

- [38] **S. Edelman and R. Shahbazi. Renewing the respect for similarity. *Frontiers in Computational Neuroscience*, 6:45, 2012.**
- [70] T. P. Lillicrap, D. Cownden, C. J. Akerman, and D. B. Tweed. Multilayer controllers can learn from random feedback weights. In *Proc. Symp. on Translational and Computational Motor Control (TCMC)*, pages 83–84, 2013. Satellite to the annual Society for Neuroscience meeting.
- [96] R. Shahbazi, R. Raizada, and S. Edelman. Similarity, kernels, and the fundamental constraints on cognition. *Journal of Mathematical Psychology*, 70:21–34, 2016.



Reza Shahbazi



### 3.12 Week 12 (4/17; 4/19): neurons, II

#### 12.1 What do neurons do?

Time-dependent dynamic learning:

- spike timing dependent plasticity (STDP), and Hebbian learning (see the Scholarpedia article [99]);
- history-dependent learning with the BCM rule [19];
- learning dynamic embodied control modes [25].

#### 12.2 What do neurons do?

Population dynamics:

- neural trajectories and classification [12].

#### Primary readings

- [19] L. N. Cooper and M. F. Bear. The BCM theory of synapse modification at 30: interaction of theory with experiment. *Nature Reviews Neuroscience*, 13:798–810, 2012.
- [25] R. Der and G. Martius. Novel plasticity rule can explain the development of sensorimotor intelligence. *Proceedings of the National Academy of Science*, pages E6224–E6232, 2015.
- [99] J. Sjöström and W. Gerstner. Spike-timing dependent plasticity. *Scholarpedia*, 5(2):1362, 2010.
- [12] D. V. Buonomano and W. Maass. State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10:113–125, 2009.

#### Other readings

- [60] E. M. Izhikevich and N. S. Desai. Relating STDP to BCM. *Neural Computation*, 15:1511–1523, 2003
- [24] **S. Deneve. Bayesian spiking neurons I: Inference. *Neural Computation*, 20: 91–117, 2008.**
- [59] E. M. Izhikevich. Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cerebral Cortex*, 17:2443–2452, 2007.
- [81] B. J. Fischer J. L. Peña. Owl’s behavior and neural representation predicted by Bayesian inference. *Nature Neuroscience*, 14:1061–1067, 2011.



Sophie Deneve

### 3.13 Week 13 (4/24; 4/26): neurons, III; brains

#### 13.1 What do neurons do?

Population dynamics (cont.):

- ongoing dynamics and chaotic itinerancy [92].

#### 13.2 What does the brain do? Prediction [80, 18].

#### Primary readings

- [92] M. I. Rabinovich, A. N. Simmons, and P. Varona. Dynamical bridge between brain and mind. *Trends in Cognitive Sciences*, 19:453–461, 2015.
- [44] **K. J. Friston. The free-energy principle: a unified brain theory? *Nature Neuroscience*, 11:127–138, 2010.**
- [80] H.-J. Park and K. J. Friston. Structural and functional brain networks: from connections to cognition. *Science*, 342:1238411, 2013. doi: 10.1126/science.1238411.



Karl Friston

#### Other readings

- [18] A. Clark. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36:181–204, 2013.
- [11] J. Bruineberg, J. Kiverstein, and E. Rietveld. The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. *Synthese*, 2016. doi: 10.1007/s11229-016-1239-1.

### 3.14 Week 14 (5/1; 5/3): advanced topics

- 14.1 Bayes and the real world [95]
- 14.2 Real behavior and real brains [76, 35]

#### Primary readings

- [95] A. N. Sanborn and N. Chater. Bayesian brains without probabilities. *Trends in Cognitive Sciences*, 20: 883–893, 2016. doi: 10.1016/j.tics.2016.10.003.
- [76] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- [35] S. Edelman. The minority report: some common assumptions to reconsider in the modeling of the brain and behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, 28:751–776, 2016.

#### Other readings

- [69] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521:436–444, 2015.
- [68] N. Kriegeskorte. Deep neural networks: a new framework for modelling biological vision and brain information processing. *Annual Reviews of Vision Science*, 2017. In press.
- [45] **S. J. Gershman, K. A. Norman, and Y. Niv. Discovering latent causes in reinforcement learning. *Current Opinion in Behavioral Sciences*, 5:43–50, 2015.**



Yael Niv

### 3.15 Week 15 (5/8): wrap-up

15.1 Summary [30, ch.11]. Being human.

15.2 [the semester's over!]

---

## 4 Computational concepts, principles, and methods: an abecedary

Note: the number of important principles in cognitive science is somewhat larger than the number of letters in the alphabet. To make up for this inconvenience, I have highlighted in SMALL CAPITALS every concept of interest that is mentioned, whether or not it has its own entry.

**ANALOGY** — a comparison that involves a structure mapping between complex entities, situations, or domains [55, 56]. Analogy is central to general cognitive function (general fluid intelligence, often referred to as *gF* or IQ [101]) and has been hypothesized to underlie structural **LEARNING** in vision and in language.

**BAYES THEOREM** — a direct consequence of the definition of conditional probability; the basis for the so-called rational theories of cognition. The Bayes Theorem prescribes a way of integrating prior beliefs with new data, in a way that proves useful in all domains, from perception, through thinking and language, to motor control [66, 115, 37, 105, 86, 67].

**CHANNEL CODING** — measuring a stimulus with a set of graded, overlapping filters (receptive fields or “channels”) supports a high degree of resolution, or hyperacuity, that cannot be achieved through dense sampling by pointlike filters [100]. This principle is at work throughout cognition [30, p.90].

**DIMENSIONALITY REDUCTION** — A high-dimensional perceptual measurement space is advantageous because it may capture more of the useful structure of the problems that a cognitive system needs to deal with, such as categorization. **LEARNING** by “tiling” the representation space with examples is, however, infeasible in a high-dimensional space, because the number of required examples grows exponentially with dimensionality [36]. This necessitates dimensionality reduction prior to learning, which, moreover, needs to be done so as to lose as little as possible of the useful information.

**EMBODIMENT AND SITUATEDNESS** — **EVOLUTION** fine-tunes the computations carried out by natural cognitive systems to the mechanics of the bodies that they control and the ecological niche in which they are situated [3].

**FUNCTION APPROXIMATION** — **LEARNING** from examples and generalization to new queries is equivalent to function approximation, a problem in which the values of an unknown function are given at a number of points in its domain and are used to form an estimate that can then support generalization [88].

**GRAPHICAL MODELS** — The relationships among a set of variables of interest to a cognitive system can be conveniently represented in the form of a directed graph, in which the vertices stand for variables (of which some may be observable and others hidden, corresponding to the properties of the world that need to be inferred from sensory data) and the edges — for probabilistic dependencies between pairs of variables [82]. One type of such model is the **BAYES** Network [85]. Graphical models map naturally onto the architecture of the brain [71].

**HOLISM** — The **PATTERN** of causal dependencies in a system of knowledge about the natural world is such that any two items may be potentially interdependent; in this sense, rich cognitive representations are

- holistic [91]. This property of world knowledge gives rise to serious algorithmic challenges in truth maintenance systems, where facts newly acquired through LEARNING can potentially interact with, and cause the revision of, the entire knowledge base.
- ILL-POSEDNESS** — Problems arising in perception, thinking, and action planning are typically ill-posed in the sense that they do not possess a unique solution (e.g., 6). FUNCTION APPROXIMATION, which is central to LEARNING, is ill-posed because an infinite number of mappings may be consistent with a given set of input-output pairs, and so is probability density estimation. Such problems can be made well-posed by REGULARIZATION.
- JOINT PROBABILITY** — The most that can be learned about the world by observing or tracking a set of variables of interest is an approximation to their joint probability density (note that the problem of probability estimation is ILL-POSED). To learn the causal dependencies among the variables, one must go beyond mere observation and intervene on variables of interest [83].
- KERNELS** — A family of mathematical methods that arise from measurements of SIMILARITY of two vectors and that are widely applicable in modeling cognition [61, 62, 96]. Formally, a positive definite kernel is a function of two arguments that represents an inner product (dot product) in some feature space.
- LEARNING AND LEARNABILITY** — Most of the detailed knowledge about how the world works that animals with advanced cognitive systems need to master cannot be “squeezed” through the genomic bottleneck and must therefore be learned from experience. The field of machine learning has amassed a wealth of insights into the computational nature of this process, including constraints and limitations on learning and learnability (e.g., 108, 109).
- MINIMUM DESCRIPTION LENGTH (MDL) PRINCIPLE** — The fundamental principle of the computational theory of learning, due to Solomonoff [102], is that LEARNING is learning of regularities. It is derived from the observation that learning is only useful insofar as it supports generalization and that generalization is only possible if regularities are discovered in the observed data. A modern operationalization of this idea is the Minimum Description Length Principle of Rissanen [93], according to which regularities in the data are best captured by a representation that minimizes the sum of the description lengths of the code and of the training data under that code [50]. A related principle is that of SIMPLICITY [16].
- NAVIGATION** — Finding a route through a representation space, subject to certain CONSTRAINTS, is a paradigm for all sequential behaviors. Thus, in foraging, for instance, the SEARCH space represents the terrain in which the animal is situated; in planning, it may be a graph representing the space of possible solutions to the problem at hand; in language production, the graph would be a representation of the speaker’s knowledge of language [31].
- OPTIMIZATION** — A wide range of tasks in cognition, including perception, thinking (e.g., problem solving and decision making), and motor control reduce to SEARCHING a space of possible solutions for an optimal one [30]. Optimality in this context is imposed by various CONSTRAINTS, which may stem from the nature of the problem, from implementational considerations, from EVOLUTIONARY pressure, or from general requirements of tractability and uniqueness (as in REGULARIZATION).
- PREDICTION** — A true understanding of the world (e.g., one that takes the form of a CAUSAL PROBABILISTIC model) should allow the cognitive system to exercise FORESIGHT: to predict impending events and the consequences of its own actions [30]. Such capacity for prediction turns out to be a very general explanatory principle in cognition [18], which can be linked to other general principles, such as BAYESIAN probability theory.
- QUANTUM PROBABILITY** — Animal behaviors that involve probabilistic assessment of cues and outcomes are often strongly context- and order-dependent. Understanding such behaviors may require positing individual states that are superpositions (i.e., are impossible to associate with specific values), as well as composite systems that are entangled (i.e., that cannot be decomposed into their subsystems). The relevant

- theories are best expressed in terms of quantum probability postulates [90].
- REGULARIZATION** — A problem that is formally ILL-POSED in that it has no unique solution can be turned into a well-posed one by imposing external CONSTRAINTS on the solution space. One class of such constraints, which has a profound grounding in LEARNING theory, is regularization through smoothing, which is related to BAYESIAN probability, to statistical learning theory, and to the MAXIMUM LIKELIHOOD idea [42].
- SIMILARITY** — the most important ultimate use to which sensory data could be put involves estimating the similarity between two stimuli, which constitutes the only principled basis for GENERALIZATION of response from one stimulus to another, and therefore for any non-trivial LEARNING from experience [97, 28, 38].
- TUNING** — Neurons in animal brains, and neuron-like units in artificial distributed cognitive systems, are typically tuned to various features of the perceptual world, of motor behavior, or of the animal’s internal representational states [30, ch.3]. Graded, shallow tuning, with a high degree of overlap between the profiles of adjacent units, is behind perceptual filters or CHANNELS that support hyperacuity. Because a tuned unit effectively represents the SIMILARITY between the actual and optimal stimuli, graded tuning also underlies VERIDICALITY [29].
- UNCERTAINTY** — “The percept is always a wager. Thus uncertainty enters at two levels, not merely one: the configuration may or may not indicate an object, and the cue may or may not be utilized at its true indicative value” [47]. The fundamental uncertainty in dealing with the world is the central motivation for the use of PROBABILISTIC representations and processes by cognitive systems.
- VERIDICALITY** — Perceptual representations based on CHANNEL CODING are provably veridical in that they faithfully reflect the SIMILARITY relationships of the represented items, such as visual objects [29].
- WEIGHT LEARNING** — In computational systems composed of simple, neuron-like elements, LEARNING typically proceeds by adjusting the weights of the synaptic connections, although the threshold of the non-linear transfer function that is part of the standard “formal neuron” can also be adjusted [20]. The rule for experience-based weight adjustment proposed by Hebb in 1949, according to which “neurons that fire together, wire together,” has now been recast and widely accepted as spike timing dependent plasticity, or STDP [14].
- MAXIMUM LIKELIHOOD ESTIMATION (MLE)** — According to the MLE principle, the parameters of a PROBABILISTIC model that is intended to reproduce a set of observations should be tuned so as to make the actual observed data set most likely [77].
- SYNTAX** — The system of PATTERNS and CONSTRAINTS that governs the composition of utterances in a natural language [87, 39].
- DOBZHANSKY** — “Nothing in biology makes sense except in the light of EVOLUTION” [27].

## References

- [1] R. A. Altes. Ubiquity of hyperacuity. *J. Acoust. Soc. Am.*, 85:943–952, 1988.
- [2] J. R. Anderson. ACT: A simple theory of complex cognition. *American Psychologist*, 51:355–365, 1996.
- [3] M. L. Anderson. Embodied cognition: A field guide. *Artificial Intelligence*, 149:91–130, 2003.
- [4] A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Communications of the ACM*, 51:117–122, 2008.

- [5] D. Barker-Plummer. Turing machines. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. 2007. Available online at <http://plato.stanford.edu/archives/win2007/entries/turing-machine/>.
- [6] M. Bertero, T. Poggio, and V. Torre. Ill-posed problems in early vision. *Proceedings of the IEEE*, 76: 869–889, 1988.
- [7] M. A. Boden. Precis of “The creative mind: Myths and mechanisms”. *Behavioral and Brain Sciences*, 17:519–570, 1994.
- [8] M. M. Botvinick. Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences*, 12:201–208, 2008.
- [9] M. M. Botvinick, Y. Niv, and A. C. Barto. Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, 113:262–280, 2009.
- [10] D. H. Brainard and W. T. Freeman. Bayesian color constancy. *J. Opt. Soc. Am. A*, 14:1393–1411, 1997.
- [11] J. Bruineberg, J. Kiverstein, and E. Rietveld. The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. *Synthese*, 2016. doi: 10.1007/s11229-016-1239-1.
- [12] D. V. Buonomano and W. Maass. State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10:113–125, 2009.
- [13] G. Buzsáki. Neural syntax: cell assemblies, synapsembles, and readers. *Neuron*, 68:362–385, 2010.
- [14] N. Caporale and Y. Dan. Spike timing-dependent plasticity: A Hebbian learning rule. *Annual Review of Neuroscience*, 31:25–46, 2008.
- [15] N. Chater. Rational and mechanistic perspectives on reinforcement learning. *Cognition*, 113:350–364, 2009.
- [16] N. Chater and P. Vitányi. Simplicity: a unifying principle in cognitive science? *Trends in Cognitive Sciences*, 7:19–22, 2003.
- [17] N. Chater, J. B. Tenenbaum, and A. Yuille. Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, 10:287–291, 2006.
- [18] A. Clark. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36:181–204, 2013.
- [19] L. N. Cooper and M. F. Bear. The BCM theory of synapse modification at 30: interaction of theory with experiment. *Nature Reviews Neuroscience*, 13:798–810, 2012.
- [20] L. N. Cooper, N. Intrator, B. S. Blais, and H. Z. Shouval. *Theory of Cortical Plasticity*. World Scientific, Singapore, 2004.
- [21] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20:273–297, 1995.
- [22] F. Cutzu and S. Edelman. Faithful representation of similarities among three-dimensional shapes in human vision. *Proceedings of the National Academy of Science*, 93:12046–12050, 1996.
- [23] P. Dayan. How to set the switches on this thing. *Current Opinion in Neurobiology*, 22:1068–1074, 2012.

- [24] S. Deneve. Bayesian spiking neurons I: Inference. *Neural Computation*, 20:91–117, 2008.
- [25] R. Der and G. Martius. Novel plasticity rule can explain the development of sensorimotor intelligence. *Proceedings of the National Academy of Science*, pages E6224–E6232, 2015.
- [26] J. Dewey. The reflex arc concept in psychology. *Psychological Review*, 3:357–370, 1896.
- [27] T. Dobzhansky. Nothing in biology makes sense except in the light of evolution. *The American Biology Teacher*, 35:125–129, 1973.
- [28] S. Edelman. Representation is representation of similarity. *Behavioral and Brain Sciences*, 21:449–498, 1998.
- [29] S. Edelman. *Representation and recognition in vision*. MIT Press, Cambridge, MA, 1999.
- [30] S. Edelman. *Computing the mind: how the mind really works*. Oxford University Press, New York, NY, 2008.
- [31] S. Edelman. On look-ahead in language: navigating a multitude of familiar paths. In M. Bar, editor, *Prediction in the Brain*, chapter 14, pages 170–189. Oxford University Press, New York, 2011.
- [32] S. Edelman. Vision, reanimated and reimagined. *Perception*, 41:1116–1127, 2012. Special issue on Marr’s *Vision*.
- [33] S. Edelman. *The Happiness of Pursuit*. Basic Books, New York, NY, 2012.
- [34] S. Edelman. Varieties of perceptual truth and their possible evolutionary roots. *Psychonomic Bulletin and Review*, 22:1519–1522, 2015. doi: 10.3758/s13423-014-0741-z.
- [35] S. Edelman. The minority report: some common assumptions to reconsider in the modeling of the brain and behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, 28:751–776, 2016.
- [36] S. Edelman and N. Intrator. Learning as extraction of low-dimensional representations. In D. Medin, R. Goldstone, and P. Schyns, editors, *Mechanisms of Perceptual Learning*, pages 353–380. Academic Press, 1997.
- [37] S. Edelman and R. Shahbazi. Survival in a world of probable objects. *Behavioral and Brain Sciences*, 34:197–198, 2011. A commentary on *Bayesian Fundamentalism or Enlightenment? On the Explanatory Status and Theoretical Contributions of Bayesian Models of Cognition* by Jones & Love.
- [38] S. Edelman and R. Shahbazi. Renewing the respect for similarity. *Frontiers in Computational Neuroscience*, 6:45, 2012.
- [39] S. Edelman and H. R. Waterfall. Behavioral and computational aspects of language and its acquisition. *Physics of Life Reviews*, 4:253–277, 2007.
- [40] S. Edelman, K. Grill-Spector, T. Kushnir, and R. Malach. Towards direct visualization of the internal shape representation space by fMRI. *Psychobiology*, 26:309–321, 1998.
- [41] K. Embar, S. Mukherjee, and B. P. Kotler. What do predators really want? The role of gerbil energetic state in determining prey choice by Barn Owls. *Ecology*, 95:280–285, 2014.



- [42] T. Evgeniou, M. Pontil, and T. Poggio. Regularization networks and support vector machines. *Advances in Computational Mathematics*, 13:1–50, 2000.
- [43] W. A. Freiwald, D. Y. Tsao, and M. S. Livingstone. A face feature space in the macaque temporal lobe. *Nature Neuroscience*, 12:1187–1199, 2009.
- [44] K. J. Friston. The free-energy principle: a unified brain theory? *Nature Neuroscience*, 11:127–138, 2010.
- [45] S. J. Gershman, K. A. Norman, and Y. Niv. Discovering latent causes in reinforcement learning. *Current Opinion in Behavioral Sciences*, 5:43–50, 2015.
- [46] Z. Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521:452–459, 2015.
- [47] J. J. Gibson. Survival in a world of probable objects. *Contemporary Psychology*, 2:33–35, 1957.
- [48] L. S. Gottfredson. Life, death, and intelligence. *Journal of Cognitive Education and Psychology*, 1: 23–46, 2004.
- [49] T. L. Griffiths and A. Yuille. Technical introduction: A primer on probabilistic inference. *Trends in Cognitive Sciences*, 10, 2006. Supplementary material. DOI: 10.1016/j.tics.2006.05.007.
- [50] P. Grunwald. Introducing the minimum description length principle. In *Advances in minimum description length: theory and applications*, pages 3–22. MIT, Cambridge, MA, 2005.
- [51] J. V. Haxby, M. I. Gobbini, M. L. Furey, A. Ishai, J. L. Schouten, and P. Pietrini. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293:2425–2430, 2001.
- [52] Y. Hazan, Y. Kra, I. Yarin, H. Wagner, and Y. Gutfreund. Visual-auditory integration for visual search a behavioral study in barn owls. *Frontiers in Integrative Neuroscience*, 9(11):1–12, 2015.
- [53] D. R. Hofstadter. *Gödel, Escher, Bach: an Eternal Golden Braid*. Basic Books, New York, 1979.
- [54] D. R. Hofstadter. Variations on a theme as the crux of creativity. In *Metamagical Themas*, chapter 12, pages 232–259. Viking, Harmondsworth, England, 1985.
- [55] D. R. Hofstadter. Analogy as the core of cognition. In D. Gentner, K. J. Holyoak, and B. N. Kokinov, editors, *The Analogical Mind: Perspectives from Cognitive Science*, pages 499–538. MIT Press, Cambridge MA, 2001.
- [56] D. R. Hofstadter and M. Mitchell. The Copycat project: a model of mental fluidity and analogy-making. In D. R. Hofstadter, editor, *Fluid Concepts and Creative Analogies*, chapter 5, pages 205–265. Basic Books, NY, 1995.
- [57] S. Hotton and J. Yoshimi. Extending dynamical systems theory to model embodied cognition. *Cognitive Science*, 35:444–479, 2010.
- [58] A. Hurlbert. Colour constancy. *Current Biology*, 17:R906–R907, 2007.
- [59] E. M. Izhikevich. Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cerebral Cortex*, 17:2443–2452, 2007.

- [60] E. M. Izhikevich and N. S. Desai. Relating STDP to BCM. *Neural Computation*, 15:1511–1523, 2003.
- [61] F. Jäkel, B. Schölkopf, and F. A. Wichmann. Generalization and similarity in exemplar models of categorization: Insights from machine learning. *Psychonomic Bulletin & Review*, 15:256–271, 2008.
- [62] F. Jäkel, B. Schölkopf, and F. A. Wichmann. Does cognitive science need kernels? *Trends in Cognitive Sciences*, 13:381–388, 2009.
- [63] F. Jiang, V. Blanz, and A. J. O’Toole. Probing the visual representation of faces with adaptation: A view from the other side of the mean. *Psychological Science*, 17:493–500, 2006.
- [64] P. X. Joris, P. H. Smith, and T. C. T. Yin. Coincidence detection in the auditory system: 50 years after Jeffress. *Neuron*, 21:1235–1238, 1998.
- [65] D. Kersten, P. Mamassian, and A. Yuille. Object perception as Bayesian inference. *Annual Review of Psychology*, 55:271–304, 2004.
- [66] D. Knill and W. Richards, editors. *Perception as Bayesian Inference*. Cambridge University Press, Cambridge, 1996.
- [67] K. P. Körding and D. M. Wolpert. Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, 10:319–326, 2006.
- [68] N. Kriegeskorte. Deep neural networks: a new framework for modelling biological vision and brain information processing. *Annual Reviews of Vision Science*, 2017. In press.
- [69] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521:436–444, 2015.
- [70] T. P. Lillicrap, D. Cownden, C. J. Akerman, and D. B. Tweed. Multilayer controllers can learn from random feedback weights. In *Proc. Symp. on Translational and Computational Motor Control (TCMC)*, pages 83–84, 2013. Satellite to the annual Society for Neuroscience meeting.
- [71] S. Litvak and S. Ullman. Cortical circuitry implementing graphical models. *Neural Computation*, 21:3010–3056, 2009.
- [72] D. Marr and T. Poggio. From understanding computation to understanding neural circuitry. *Neurosciences Res. Prog. Bull.*, 15:470–488, 1977.
- [73] D. C. Marr. Artificial Intelligence — a personal view. *Artificial Intelligence*, 9:37–48, 1977.
- [74] C. Massa, F. M. Gabelli, and G. R. Cueto. Using GPS tracking to determine movement patterns and foraging habitat selection of the common barn-owl (*Tyto alba*). *Hornero*, 30:7–12, 2015.
- [75] W. S. McCulloch. *Embodiments of mind*. MIT Press, Cambridge, MA, 1965.
- [76] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- [77] I. J. Myung. Tutorial on maximum likelihood estimation. *Journal of Mathematical Psychology*, 47:90–100, 2003.

- [78] A. Nestor, D. C. Plaut, and M. Behrmann. Feature-based face representations and image reconstruction from behavioral and neural data. *Proceedings of the National Academy of Science*, 113:416–421, 2016.
- [79] R. E. Nisbett, J. Aronson, C. Blair, W. Dickens, J. Flynn, D. F. Halpern, and E. Turkheimer. Intelligence: new findings and theoretical developments. *American Psychologist*, 2012.
- [80] H.-J. Park and K. J. Friston. Structural and functional brain networks: from connections to cognition. *Science*, 342:1238411, 2013. doi: 10.1126/science.1238411.
- [81] B. J. Fischer J. L. Peña. Owl’s behavior and neural representation predicted by Bayesian inference. *Nature Neuroscience*, 14:1061–1067, 2011.
- [82] J. Pearl. Causal inference in statistics: A gentle introduction. In *Computing Science and Statistics: Proceedings of Interface ’01*, volume 33, 2001. Available online at [http://ftp.cs.ucla.edu/pub/stat\\_ser/R289.pdf](http://ftp.cs.ucla.edu/pub/stat_ser/R289.pdf).
- [83] J. Pearl. Causal inference in statistics: an overview. *Statistics Surveys*, 3:96–146, 2009.
- [84] J. Pearl. Structural counterfactuals: A brief introduction. *Cognitive Science*, 37:977–985, 2013.
- [85] J. Pearl and S. Russel. Bayesian networks. In M. A. Arbib, editor, *Handbook of Brain Theory and Neural Networks*, pages 157–160. MIT Press, Cambridge, MA, 2003. Available online at [http://ftp.cs.ucla.edu/pub/stat\\_ser/R277.pdf](http://ftp.cs.ucla.edu/pub/stat_ser/R277.pdf).
- [86] A. Perfors, J. B. Tenenbaum, and T. Regier. Poverty of the stimulus? A rational approach. In R. Sun and N. Miyake, editors, *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, pages 663–668, Austin, TX, 2006. Cognitive Science Society.
- [87] C. Phillips. Syntax. In L. Nadel, editor, *Encyclopedia of Cognitive Science*, volume 4, pages 319–329. Macmillan, London, 2003.
- [88] T. Poggio. A theory of how the brain might work. *Cold Spring Harbor Symposia on Quantitative Biology*, LV:899–910, 1990.
- [89] T. Poggio and S. Edelman. A network that learns to recognize three-dimensional objects. *Nature*, 343:263–266, 1990.
- [90] E. M. Pothos and J. R. Busemeyer. Can quantum probability provide a new direction for cognitive modeling? *Behavioral and Brain Sciences*, 36:255–327, 2013.
- [91] W. V. O. Quine. Two dogmas of Empiricism. In *From a Logical Point of View*, pages 20–46. Harvard University Press, Cambridge, 1953.
- [92] M. I. Rabinovich, A. N. Simmons, and P. Varona. Dynamical bridge between brain and mind. *Trends in Cognitive Sciences*, 19:453–461, 2015.
- [93] J. Rissanen. Minimum description length principle. In S. Kotz and N. L. Johnson, editors, *Encyclopedia of Statistic Sciences*, volume 5, pages 523–527. J. Wiley and Sons, 1987.
- [94] C. D. Salzman, K. H. Britten, and W. T. Newsome. Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 346:174–177, 1990.

- [95] A. N. Sanborn and N. Chater. Bayesian brains without probabilities. *Trends in Cognitive Sciences*, 20: 883–893, 2016. doi: 10.1016/j.tics.2016.10.003.
- [96] R. Shahbazi, R. Raizada, and S. Edelman. Similarity, kernels, and the fundamental constraints on cognition. *Journal of Mathematical Psychology*, 70:21–34, 2016.
- [97] R. N. Shepard. Toward a universal law of generalization for psychological science. *Science*, 237:1317–1323, 1987.
- [98] R. N. Shepard and S. Chipman. Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology*, 1:1–17, 1970.
- [99] J. Sjöström and W. Gerstner. Spike-timing dependent plasticity. *Scholarpedia*, 5(2):1362, 2010.
- [100] H. P. Snippe and J. J. Koenderink. Discrimination thresholds for channel-coded systems. *Biological Cybernetics*, 66:543–551, 1992.
- [101] R. E. Snow, P. C. Kyllonen, and B. Marshalek. The topography of ability and learning correlates. In R. J. Sternberg, editor, *Advances in the psychology of human intelligence*, volume 2, pages 47–104. Erlbaum, Hillsdale, NJ, 1984.
- [102] R. J. Solomonoff. A formal theory of inductive inference, parts A and B. *Information and Control*, 7: 1–22, 224–254, 1964.
- [103] A. Solway, C. Diuk, N. Córdova, D. Yee, A. G. Barto, Y. Niv, and M. M. Botvinick. Optimal behavioral hierarchy. *PLOS Computational Biology*, 10(8):e1003779, 2014.
- [104] R. Talmon and R. R. Coifman. Intrinsic modeling of stochastic dynamical systems using empirical geometry. *Applied and Computational Harmonic Analysis*, 39:138–160, 2015.
- [105] J. B. Tenenbaum and T. L. Griffiths. Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, 24:629–641, 2001.
- [106] J. B. Tenenbaum, C. Kemp, T. L. Griffiths, and N. D. Goodman. How to grow a mind: statistics, structure, and abstraction. *Science*, 331:1279–1285, 2011.
- [107] L. L. Thurstone. The stimulus-response fallacy in psychology. *Psychological Review*, 30:354–369, 1923.
- [108] L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27:1134–1142, 1984.
- [109] V. N. Vapnik. An overview of statistical learning theory. *IEEE Transactions on Neural Networks*, 10: 988–999, 1999.
- [110] M. Welling. Are ML and statistics complementary?, 2015. Roundtable discussion at the 6th IMS-ISBA meeting on Data Science in the next 50 years.
- [111] E. P. Wigner. The unreasonable effectiveness of mathematics in the natural sciences. *Comm. Pure Appl. Math.*, XIII:1–14, 1960.
- [112] S. P. Wise and R. Shadmehr. Motor control. In V. S. Ramachandran, editor, *Encyclopedia of the Human Brain*, volume 3, pages 137–157. Academic Press, San Diego, CA, 2002.
- [113] W. Woergoetter and B. Porr. Reinforcement learning. *Scholarpedia*, 3(3):1448, 2007.

- [114] D. M. Wolpert, J. Diedrichsen, and J. R. Flanagan. Principles of sensorimotor learning. *Nature Reviews Neuroscience*, 12:739–751, 2011.
- [115] A. Yuille and D. Kersten. Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*, 10:301–308, 2006.