
References that are cited in **red** will likely be required reading. All cited books are available at the reserve desk in the Bailey-Howe Library. This outline is subject to modification.

1 Introduction

- Philosophical Foundations
 - Analog and digital computation [von Neumann, 1958].
 - Artificial intelligence and perceptual problems [McCulloch and Pitts, 1947; Marr, 1982]
 - Adaptive systems, and the curse of dimensionality [Bellman, 1961].
- Biological Neural Networks
 - Anatomy & physiology of neurons and synapses [Dowling, 1992, pp. 1-124] [Shepherd, 1988, pp. 1-150?]
- Artificial neural network models
 - McCulloch-Pitts neurons [Minsky, 1967, Chapter 3][McCulloch and Pitts, 1943]
 - Linear threshold units (LTUs) and perceptrons [Nilsson, 1990, pp. 1-13]; [Rosenblatt, 1962]
 - Sigmoidal threshold functions [Rumelhart et al., 1994]
 - Spiking units [Maass and Bishop, 1999, pp. 3-54] [Rieke et al., 1999]
 - Practical applications [Widrow et al., 1994].

2 Hopfield Networks and Associative Memory

- Discrete time networks [Hopfield, 1982]
- Continuous time networks [Hopfield, 1985]
- Constraint satisfaction problems [Tank and Hopfield, 1987].

3 Perceptrons and Linear Threshold Units

- Feature space vs. weight space [Nilsson, 1990, pp. 65-68]
- Rosenblatt's perceptron [Rosenblatt, 1962; Minsky and Papert, 1969] [Nilsson, 1990, pp. 69-94]
- The LMS rule [Widrow and Hoff, 1988]
- The statistical capacity of an LTU [Nilsson, 1990, pp. 15-42] [Cover, 1965].
- Computational-geometric limitations of a single perceptron. [Minsky and Papert, 1969]

4 Feed-Forward Neural Networks

- Multilayer architectures for pattern classification and function approximation.
- Kolmogorov's theorem [Cybenko, 1989]
- Back-propagation algorithm [Rumelhart et al., 1986a,b]
- NETtalk [Sejnowski and Rosenberg, 1986]
- Radial basis functions [Broomhead and Lowe, 1988].
- Statistical capacity of feed-forward networks [Baum, 1988]
- The problem of generalization:
 - Cross-validation
 - Learning invariants using the tangent distance [Simard et al., 1993]
 - Optimal brain damage [Le Cun et al., 1990]
 - Early stopping rules [Wang and Venkatesh, 1995]
 - Support vector machines [Boser et al., 1992; Cristianini and Shawe-Taylor, 2000]
 - Vapnik-Chervonenkis's theory of generalization [Abu-Mostafa, 1989][Baum and Haussler, 1990]

5 Optimization by Simulated Annealing

- Monte Carlo computations [Metropolis et al., 1953]
- Simulated annealing [Kirkpatrick et al., 1983]
- Boltzmann machines [Ackley et al., 1985] [Hinton and Sejnowski, 1986]

6 Local Learning Rules and Self-Organization

- Principal component analysis (PCA) [Oja, 1982, 1989]
- Self-organizing feature maps [Kohonen, 1989, Chapter 5]

7 Reinforcement Learning

- Temporal difference algorithm [Sutton, 1988][Sutton and Barto, 1998]
- TD-Gammon [Tesauro, 1995]

8 The Complexity of Learning

- Supervised learning is NP-complete [Blum and Rivest, 1989; Judd, 1990]
- The PAC model of learning [Valiant, 1984][Kearns and Vazirani, 1995]

References

- Y. S. Abu-Mostafa. The Vapnik-Chervonenkis dimension: information versus complexity in learning. *Neural Computation*, 1:312-317, 1989.
- D. H. Ackley, G. E. Hinton, and T. J. Sejnowski. A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1):147-169, 1985.
- J. A. Anderson and E. Rosenfeld, editors. *Neurocomputing: Foundations of Research*, volume 1. MIT Press, Cambridge, MA, 1988. ISBN 0-262-01097-6. BH: QUARTO QP363.3 .N46 1988.
- M. Anthony and P. L. Bartlett. *A Theory of Learning in Artificial Neural Networks*. Cambridge University Press, Cambridge, U.K., 1999. BH: QA76.87 .A58 1999.
- E. B. Baum. On the capabilities of multilayer perceptrons. *Journal of Complexity*, 4:193-215, 1988.
- E. B. Baum and D. Haussler. What size net gives valid generalization? *Neural Computation*, 1(1):151-160, 1990. ISSN 0899-7667.
- R. E. Bellman. *Adaptive Control Processes: A Guided Tour*. Princeton University Press, Princeton, NJ, 1961. BH: QA264 .B35 1961.
- C. M. Bishop. *Neural networks for pattern recognition*. Oxford University Press, Oxford, U.K., 1995. QA76.87 .B574 1995.
- A. Blum and R. Rivest. Training a 3-node neural net is np-complete. In D. Touretzky, editor, *Advances in Neural Information Processing Systems I*, pages 494-501. Morgan Kaufmann, 1989.
- B. E. Boser, I. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In *Fifth Annual Workshop on Computational Learning Theory*, pages 144-152. ACM Press, 1992.
- D. S. Broomhead and D. Lowe. Multivariate functional interpolation and adaptive networks. *Complex Systems*, 2:321-355, 1988.
- T. Cover and J. Thomas. *Elements of Information Theory*. John Wiley & Sons, New York, 1991. BH: Q360 .C68 1991.
- T. M. Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Transactions on Electronic Computers*, 14:326-334, 1965.

- N. Cristianini and J. Shawe-Taylor. *An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods*. Cambridge University Press, Cambridge, U.K., 2000.
- G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989. Correction appears in 5 (1995) p. 455.
- P. Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press, Cambridge, MA, 2001.
- J. E. Dowling. *Neurons and Networks: An Introduction to Neuroscience*. Harvard University Press, Cambridge, MA, 1992. BH: QP355.2 .D68 1992.
- R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. John Wiley & Sons, New York, second edition, 2001. BH: Q327 .D83 2000.
- T. Hastie, R. Tibshirani, and J. H. Friedman. *The Elements of Statistical Learning*. Springer Verlag, New York, 2001.
- S. Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice Hall, Upper Saddle River, NJ, second edition, 1999. request.
- J. Hertz, A. Krogh, and R. G. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, Redwood City, CA, 1991. BH: QA76.5 .H475 1991.
- G. E. Hinton and T. J. Sejnowski. *Learning and relearning in Boltzmann machines*, pages 282–317. Volume 1 of , Rumelhart et al. [1986c], 1986. ISBN 0–262–68053–X. BH: BF455 .R853 1986.
- J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the USA*, 79:2554–2558, 1982.
- J. J. Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences of the USA*, 81:3088–3092, 1985.
- J. S. Judd. *Neural Network Design and the Complexity of Learning*. MIT Press, Cambridge, MA, 1990.
- M. J. Kearns and U. Vazirani. *Introduction to Computational Learning Theory*. MIT Press, Cambridge, MA, 1995.
- S. Kirkpatrick, C. D. G. Jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- C. Koch. *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press, Oxford, U.K., 1998.
- T. Kohonen. *Self-Organization and Associative Memory*. Springer Verlag, New York, third edition, 1989. BH: Q325 .K64 1989.
- Y. Le Cun, J. S. Denker, and S. A. Solla. Optimal brain damage. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 2, pages 598–605, San Mateo, CA, 1990. Morgan Kaufmann.
- W. Maass and C. M. Bishop, editors. *Pulsed Neural Networks*. MIT Press, Cambridge, MA, 1999. BH: QA76.87 .P85 1999.
- D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W. H. Freeman, San Francisco, CA, 1982. BH: QP475 .M27 1982.
- W. S. McCulloch and W. Pitts. A logical calculus of ideas imminent in nervous activity. *Bulletin of Mathematical Biophysics*, 5:115–133, 1943.
- W. S. McCulloch and W. Pitts. How we know universals: the perception of auditory and visual forms. *Bulletin of Mathematical Biophysics*, 9:127–147, 1947.
- N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.

- M. L. Minsky. *Computation: Finite and Infinite Machines*. Prentice Hall, Englewood Cliffs, NJ, 1967. BH: QA267 .M55.
- M. L. Minsky and S. A. Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA, 1969. BH: Q327 .M55.
- N. J. Nilsson. *The Mathematical Foundations of Learning Machines*. Morgan Kaufmann, San Mateo, CA, 1990. BH: Q335 .N5 1990.
- E. Oja. A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15: 267-273, 1982.
- E. Oja. Neural networks, principal components, and subspaces. *International Journal of Neural Systems*, 1: 61-68, 1989.
- F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the Neural Code*. MIT Press, Cambridge, MA, 1999.
- B. D. Ripley. *Pattern Recognition and Neural Networks*. Cambridge University Press, Cambridge, U.K., 1996. BH: QA76.87 .R56 1996.
- F. Rosenblatt. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books, Washington, D.C., 1962. BH: QP376 .R7 1962.
- D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by back-propagating errors. *Nature*, 323(99):533-536, 1986a.
- D. E. Rumelhart, G. E. Hinton, and R. J. Williams. *Learning internal representations by error propagation*, chapter 8, pages 318-362. Volume 1 of , Rumelhart et al. [1986c], 1986b. ISBN 0-262-68053-X. BH: BF455 .R853 1986.
- D. E. Rumelhart, J. L. McClelland, and C. P. R. Group, editors. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, volume 1. MIT Press, Cambridge, MA, 1986c. ISBN 0-262-68053-X. BH: BF455 .R853 1986.
- D. E. Rumelhart, B. Widrow, and M. A. Lehr. The basic ideas in neural networks. *Communications of the ACM*, 37(3):87-92, 1994. ISSN 0001-0782. doi: <http://doi.acm.org/10.1145/175247.175256>.
- B. Schölkopf, C. Burgess, and A. Smola, editors. *Advances in Kernel Methods: Support Vector Learning*. MIT Press, Cambridge, MA, 1999.
- T. J. Sejnowski and C. R. Rosenberg. *NETtalk: a parallel network that learns to read aloud*, chapter 40, pages 661-672. Volume 1 of , Anderson and Rosenfeld [1988], 1986. ISBN 0-262-01097-6. The Johns Hopkins University Electrical Engineering and Computer Science Technical Report JHU/EECS-86/01.
- G. M. Shepherd. *Neurobiology*. Oxford University Press, Oxford, U.K., third edition, 1988. Dana: WL 102 S549n 1994.
- P. Simard, Y. L. Cun, and J. Denker. Efficient pattern recognition using a new transformation distance. In S. J. Hanson, J. D. Cowan, and C. L. Giles, editors, *Advances in Neural Information Processing Systems*, volume 5, pages 50-58, San Mateo, CA, 1993. Morgan Kaufmann.
- R. A. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3:9-44, 1988. URL <http://www.cs.ualberta.ca/~sutton/papers/sutton-88.pdf>.
- R. A. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998. URL <http://www-anw.cs.umass.edu/%7Eerich/book/the-book.html>.
- D. W. Tank and J. J. Hopfield. Neural computation by concentrating information in time. *Proceedings of the National Academy of Sciences of the USA*, 84:1896-1900, 1987.
- G. Tesauro. Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3):58-68, 1995. URL <http://www.research.ibm.com/massive/td1.html>.

- L. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134-1142, 1984.
- J. von Neumann. *The Computer and the Brain*. Yale University Press, New Haven, CT, 1958.
- C. Wang and S. S. Venkatesh. Temporal dynamics of generalization in neural networks. In D. S. Touretzky, G. Tesauro, and T. K. Leen, editors, *Advances in Neural Information Processing Systems*, volume 7, Cambridge, MA, 1995. MIT Press.
- B. Widrow and M. E. Hoff. *Adaptive switching circuits*, chapter 10, pages 123-134. Volume 1 of , Anderson and Rosenfeld [1988], 1988. ISBN 0-262-01097-6. BH: QUARTO QP363.3 .N46 1988.
- B. Widrow, D. E. Rumelhart, and M. A. Lehr. Neural networks: applications in industry, business and science. *Communications of the ACM*, 37(3):93-105, 1994. ISSN 0001-0782. doi: <http://doi.acm.org/10.1145/175247.175257>.