

# Hidden State and Short-Term Memory: A bibliography

Lonnie Chrisman & Michael Littman, CMU

These references were compiled for the ML93 workshop: *Reinforcement Learning: What We Know, What We Need*. “ML93” is the 10th International Conference on Machine Learning, held at the University of Massachusetts, Amherst, June 30th & July 1st, 1993. The “compilers” can be reached at: `lonnie.chrisman@cs.cmu.edu` and `michael.littman@cs.brown.edu`.

## References

- [Baldi and Chauvin, 1993] Pierre Baldi and Yves Chauvin. Smooth on-line learning algorithms for hidden markov models. 1993.
- [Baum *et al.*, 1970] Leonard E. Baum, Ted Petrie, George Soules, and Norman Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *Annals of Mathematical Statistics*, 41(1):164–171, 1970.
- [Chrisman *et al.*, 1991] Lonnie Chrisman, Rich Caruana, and Wayne Carriker. Intelligent agent design issues: Internal agent state and incomplete perception. In *Proc. AAAI Fall Symposium on Sensory Aspects of Robotic Intelligence*, 1991.
- [Chrisman, 1992] Lonnie Chrisman. Reinforcement learning with perceptual aliasing: The perceptual distinctions approach. In *Proc. Tenth National Conference on AI (AAAI)*, 1992.
- [Dean *et al.*, 1992] Thomas Dean, Dana Angluin, Kenneth Bayse, Sean Engelson, Leslie Kaelbling, Evangelos Kokkevis, and Oded Maron. Inferring finite automata with stochastic output functions and an application to map learning. In *AAAI*, 1992.
- [Dempster *et al.*, 1977] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Stat. Soc.*, 39(1):1–38, 1977.
- [Drescher, 1991] Gary L. Drescher. *Made-Up Minds: A Constructivist Approach to Artificial Intelligence*. MIT Press, 1991.
- [Howard, 1966] Ronald A. Howard. Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, SSC-2(1):22–26, August 1966.
- [Lin and Mitchell, 1992] Long-Ji Lin and Tom M. Mitchell. Memory approaches to reinforcement learning in non-markovian domains. Technical Report CMU-CS-92-138, Carnegie Mellon University, School of Computer Science, May 1992.
- [Littman, 1993] Michael L. Littman. An optimization-based categorization of reinforcement learning environments. In I. H. Meyer, H. Roithlat, and S. Wilson, editors, *From Animals to Animats: Proc. Second International Conference on Simulation and Adaptive Behavior*. MIT Press, 1993.
- [Lovejoy, 1991] William S. Lovejoy. A survey of algorithmic methods for partially observable markov decision processes. *Annals of Operations Research*, 28:47–66, 1991.
- [Mahadevan and Connell, 1991] Sridhar Mahadevan and Jonathan Connell. Automatic programming of behavior-based robots using reinforcement learning. In *Proc. Ninth National Conference on Artificial Intelligence*, 1991.
- [McCallum, 1992] R. Andrew McCallum. First results with utile distinction memory for reinforcement learning. Technical Report 446,

- Dept. Comp. Sci., Univ. Rochester, 1992. See also Proceedings of Machine Learning Conference 1993.
- [Monahan, 1982] George E. Monahan. A survey of partially observable markov decision processes: Theory, models, and algorithms. *Management Science*, 28:1–16, January 1982.
- [Moore, 1991] Andrew W. Moore. Variable resolution dynamic programming: Efficiently learning action maps in multivariate real-valued spaces. In *Proc. Eight International Machine Learning Workshop*, 1991.
- [Mozer and Bachrach, 1991] Michael C. Mozer and Jonathan Bachrach. SLUG: A connectionist architecture for inferring the structure of finite-state environments. *Machine Learning*, 7:139–160, 1991.
- [Papadimitriou and Tsitsiklis, 1987] Christos H. Papadimitriou and John N. Tsitsiklis. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, August 1987.
- [Rabiner, 1989] Lawrence R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), February 1989.
- [Ring, 1993] Mark Ring. Two methods for hierarchy learning in reinforcement environments. In I. H. Meyer, H. Roithlat, and S. Wilson, editors, *From Animals to Animats: Proc. Second International Conference on Simulation and Adaptive Behavior*. MIT Press, 1993.
- [Rivest and Schapire, 1987] Ronald L. Rivest and Robert E. Schapire. A new approach to unsupervised learning in deterministic environments. In *Machine Learning Workshop*, Irvine, CA, 1987.
- [Rivest and Schapire, 1989] Ronald L. Rivest and Robert E. Schapire. Inference of finite automata using homing sequences. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, pages 411–420, Seattle, Washington, May 1989.
- [Ross, 1970] S. Ross. *Applied Probability Models with Optimization Applications*. Holden-Day, San Francisco, Calif., 1970.
- [Schmidhuber, 1991] Juergen H. Schmidhuber. Reinforcement learning in markovian and non-markovian environments. In D. S. Lippman, J. E. Moody, and D. S. Touretzky, editors, *Advances in Neural Information Processing Systems 3*, pages 500–506, San Mateo, CA, 1991. Morgan Kaufman.
- [Sondik, 1971] E. Sondik. *The Optimal Control of Partially Observable Markov Processes*. PhD thesis, Stanford University, 1971.
- [Sondik, 1978] Edward J. Sondik. The optimal control of partially observable markov processes over the infinite horizon: Discounted costs. *Operations Research*, 26(2), 1978.
- [Tan, 1991] Ming Tan. Cost-sensitive reinforcement learning for adaptive classification and control. In *Proc. Ninth National Conference on Artificial Intelligence*, 1991.
- [Whitehead and Ballard, 1990] Steven D. Whitehead and Dana H. Ballard. Active perception and reinforcement learning. In *Proc. Seventh International Conference on Machine Learning*, 1990.
- [Whitehead and Ballard, 1991] Steven D. Whitehead and Dana H. Ballard. Learning to perceive and act by trial and error. *Machine Learning*, 7:45–83, 1991.
- [Whitehead, 1991] Steven D. Whitehead. A complexity analysis of cooperative mechanisms in reinforcement learning. In *Proc. Ninth National Conference on Artificial Intelligence*, 1991.
- [Wilson, 1985] Stewart W. Wilson. Knowledge growth in an artificial animal. In *Proceedings of the First International Conference on Genetic Algorithms and Their Applications*, pages 16–23, Hillsdale, NJ, 1985. Lawrence Erlbaum Associates.
- [Wixson, 1991] Lambert E. Wixson. Scaling reinforcement learning techniques via modularity. In *Proc. Eighth International Machine Learning Workshop*, 1991.