

AAAI-14 Games Tutorial List of Publications - MCTS

Martin Müller and Akihiro Kishimoto

July 14, 2014

1 MCTS Guide to the Literature

- <http://mcts.ai/bibliography/index.html> maintains a massive bibliography of MCTS-related publications (360 as of July 6, 2014).
- The page <https://chessprogramming.wikispaces.com/Monte-Carlo+Tree+Search#Publications> also has a large selection.
- The list below is a selection of articles which are most relevant for this tutorial.

2 Classics

[34, 66, 51, 48, 35, 83, 67, 7, 98, 21, 1]

3 Survey Papers

[73, 20, 33]

4 Theory

[61, 66, 116, 83, 67, 6, 7, 85, 63, 4]

5 MCTS Enhancements, Algorithms, Analysis

[69, 68, 11, 110, 23, 16, 48, 108, 97, 28, 32, 30, 93, 54, 100, 101, 29, 80, 76, 77, 102, 78, 50, 107, 106, 60, 109]

5.1 Parallel MCTS

[112, 47, 40, 114, 96, 95, 90, 52, 27, 26, 62, 31, 36, 84]

5.2 Adaptive Simulations

[110, 39, 9, 81, 57, 86]

5.3 Machine Learning for MCTS

[53, 111, 92, 91]

6 Specific Games

6.1 Go

[21, 17, 16, 25, 113, 34, 51, 15, 48, 46, 35, 5, 55, 82, 60, 56, 59, 43, 115]

6.1.1 Specific Go Programs

[49, 70, 41, 60, 56, 12, 50]

6.2 Hex

[3, 58]

6.3 General Game-Playing (GGP)

[44, 13, 72, 45]

6.4 Puzzles, Single Player Games

[88, 89, 22, 86, 2]

6.5 Other Specific Games

[18, 69, 68, 75, 110, 87, 71, 42, 99, 109]

7 Non-Game Applications

[19, 53, 10, 74, 38, 94, 64, 65, 8, 24, 85]

8 PhD and MSc Theses

[42, 106, 100, 91, 46, 79, 56]

References

- [1] B. Abramson. Expected-outcome: A general model of static evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:182–193, February 1990.
- [2] H. Akiyama, K. Komiya, and Y. Kotani. Nested Monte-Carlo search with AMAF heuristic. In *TAAI*, 2010.
- [3] B. Arneson, R. B. Hayward, and P. Henderson. Monte Carlo tree search in Hex. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):251–258, 2010.
- [4] J.-Y. Audibert and S. Bubeck. Best arm identification in multi-armed bandits. *Proceedings COLT 2010*, 2010.
- [5] P. Audouard, G. Chaslot, J.-B. Hoock, A. Rimmel, J. Perez, and O. Teytaud. Grid coevolution for adaptive simulations; application to the building of opening books in the game of Go. In *Applications of Evolutionary Computing*, volume 5484 of *Lecture Notes in Computer Science*, pages 323–332. Springer, 2009.
- [6] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2003.
- [7] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, May 2002.
- [8] S. Baba, A. Iwasaki, M. Yokoo, M. C. Silaghi, K. Hirayama, and T. Matsui. Cooperative problem solving against adversary: Quantified distributed constraint satisfaction problem. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '10*, pages 781–788, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems.
- [9] H. Baier and P. Drake. The power of forgetting: Improving the last-good-reply policy in Monte Carlo Go. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):303–309, 2010.
- [10] H. Baier and M. Winands. Nested Monte-Carlo Tree Search for online planning in large MDPs. In L. De Raedt, C. Bessière, D. Dubois, P. Doherty, P. Frasconi, F. Heintz, and P. Lucas, editors, *ECAI*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, pages 109–114. IOS Press, 2012.
- [11] H. Baier and M. Winands. Monte-Carlo Tree Search and minimax hybrids. In *CIG*, pages 1–8. IEEE, 2013.
- [12] P. Baudis. Pachi simple Go/baduk/weiqi bot, 2010. Date retrieved: March 10, 2010.

- [13] Y. Björnsson and H. Finnsson. Cadiaplayer: A simulation-based General Game Player. *Computational Intelligence and AI in Games, IEEE Transactions on*, 1(1):4–15, March 2009.
- [14] A. Bourki, G. Chaslot, M. Coulm, V. Danjean, H. Doghmen, T. Hérault, J.-B. Hoock, A. Rimmel, F. Teytaud, O. Teytaud, P. Vayssière, and Z. Yu. Scalability and Parallelization of Monte-Carlo Tree Search. In *The International Conference on Computers and Games 2010*, Kanazawa, Japan, 2010.
- [15] B. Bouzy. Old-fashioned computer Go vs Monte-Carlo Go. CIG07 Tutorial, Date retrieved: March 10, 2010.
- [16] B. Bouzy. Associating domain-dependent knowledge and Monte Carlo approaches within a Go program. *Information Sciences*, 175:247–257, 2005.
- [17] B. Bouzy and B. Helmstetter. Monte Carlo Go developments. In E. A. Heinz H. J. van den Herik, H. Iida, editor, *Advances in Computer Games conference (ACG-10)*, pages 159–174, 2003.
- [18] S. R. K. Branavan, D. Silver, and R. Barzilay. Non-linear Monte-Carlo search in Civilization II. In T. Walsh, editor, *IJCAI*, pages 2404–2410. IJCAI/AAAI, 2011.
- [19] S. R. K. Branavan, D. Silver, and R. Barzilay. Learning to win by reading manuals in a Monte-Carlo framework. *Journal of Artificial Intelligence Research (JAIR)*, 43:661–704, 2012.
- [20] C. Browne, E. Powley, D. Whitehouse, S. Lucas, P. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. A survey of Monte Carlo tree search methods. *IEEE Trans. Comput. Intellig. and AI in Games*, 4(1):1–43, 2012.
- [21] B. Brüggmann. Monte Carlo Go, March 1993. Unpublished manuscript, <http://www.cgl.ucsf.edu/go/Programs/Gobble.html>.
- [22] T. Cazenave. Nested Monte-Carlo search. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009, Pasadena, California, USA, July 11-17, 2009*, pages 456–461, 2009.
- [23] T. Cazenave. Sequential halving applied to trees, 2014. To appear in IEEE TCIAIG.
- [24] T. Cazenave, F. Balbo, and S. Pinson. Monte-Carlo bus regulation. In *ITSC*, 2009.
- [25] T. Cazenave and B. Helmstetter. Combining tactical search and Monte-Carlo in the game of Go. In *In: CIG 2005*, pages 171–175, 2005.

- [26] T. Cazenave and N. Jouandeau. On the parallelization of UCT. In *Computer Games Workshop*, pages 93–101, Amsterdam, 2007.
- [27] T. Cazenave and N. Jouandeau. A parallel Monte-Carlo tree search algorithm. In van den Herik et al. [105], pages 72–80.
- [28] G. Chaslot, L. Chatriot, C. Fiter, S. Gelly, J. Hoock, J. Perez, A. Rimmel, and O. Teytaud. Combining expert, offline, transient and online knowledge in Monte-Carlo exploration, 2008.
- [29] G. Chaslot, C. Fiter, J.-B. Hoock, A. Rimmel, and O. Teytaud. Adding expert knowledge and exploration in Monte-Carlo Tree Search. In van den Herik and Spronck [104], pages 1–13.
- [30] G. Chaslot, M. Winands, J. Uiterwijk, J. van den Herik, and B. Bouzy. Progressive strategies for Monte-Carlo tree search. *New Mathematics and Natural Computation*, 4(3):343–357, 2008.
- [31] G. Chaslot, M. Winands, and J. van den Herik. Parallel Monte-Carlo tree search. In *Proceedings of the 6th International Conference on Computer and Games*, volume 5131 of *Lecture Notes in Computer Science*, pages 60–71. Springer, 2008.
- [32] B. Childs, J. Brodeur, and L. Kocsis. Transpositions and move groups in Monte Carlo Tree Search. In P. Hingston and L. Barone, editors, *IEEE Symposium on Computational Intelligence and Games*, pages 389–395. IEEE, 2008.
- [33] A. Couëtoux, M. Müller, and O. Teytaud. Monte Carlo Tree Search in Go, 2015? Book chapter, to appear in *Monte Carlo Tree Search for Games and Beyond*.
- [34] R. Coulom. Efficient selectivity and backup operators in Monte-Carlo tree search. In J. van den Herik, P. Ciancarini, and H. Donkers, editors, *Proceedings of the 5th International Conference on Computer and Games*, volume 4630/2007 of *Lecture Notes in Computer Science*, pages 72–83, Turin, Italy, June 2006. Springer.
- [35] R. Coulom. Computing Elo ratings of move patterns in the game of Go. In *Proceedings of the Computer Games Workshop*, 2007.
- [36] R. Coulom. Lockless hash table and other parallel search ideas, 2008. <http://computer-go.org/pipermail/computer-go/2008-March/014537.html>. Date retrieved: April 28, 2009.
- [37] A. Pohorecky, Danyluk, L. Bottou, and M. Littman, editors. *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009, Montreal, Quebec, Canada, June 14-18, 2009*, volume 382 of *ACM International Conference Proceeding Series*. ACM, 2009.

- [38] F. de Mesmay, A. Rimmel, Y. Voronenko, and M. Püschel. Bandit-based optimization on graphs with application to library performance tuning. In Danyluk et al. [37], page 92.
- [39] P. Drake. The Last-Good-Reply policy for Monte-Carlo Go. *ICGA Journal*, 32(4):221–227, 2009.
- [40] M. Enzenberger and M. Müller. A lock-free multithreaded Monte-Carlo tree search algorithm. In van den Herik and Spronck [104], pages 14–20.
- [41] M. Enzenberger, M. Müller, B. Arneson, and R. Segal. Fuego - an open-source framework for board games and Go engine based on Monte Carlo tree search. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):259–270, 2010.
- [42] T. Ewalds. Playing and solving Havannah. Master’s thesis, University of Alberta, 2012.
- [43] S. Fernando and M. Müller. Analyzing simulations in Monte Carlo tree search for the game of Go, 2013. To appear in *Computers and Games 2013*. 12 pp.
- [44] H. Finnsson and Y. Björnsson. Simulation-based approach to General Game Playing. In D. Fox and C. Gomes, editors, *AAAI*, pages 259–264. AAAI Press, 2008.
- [45] H. Finnsson and Y. Björnsson. Simulation control in General Game Playing agents. In *Proceedings of the IJCAI-09 Workshop on General Game Playing (GIGA’09)*, 2009.
- [46] S. Gelly. *A Contribution to Reinforcement Learning; Application to Computer-Go*. PhD thesis, Université Paris-Sud, 2007.
- [47] S. Gelly, J.-B. Hoock, A. Rimmel, O. Teytaud, and Y. Kalemkarian. The parallelization of Monte-Carlo planning - parallelization of MC-planning. In J. Filipe, J. Andrade-Cetto, and J.-L. Ferrier, editors, *ICINCO-ICSO*, pages 244–249. INSTICC Press, 2008.
- [48] S. Gelly and D. Silver. Combining online and offline knowledge in UCT. In *ICML ’07: Proceedings of the 24th international conference on Machine learning*, pages 273–280. ACM, 2007.
- [49] S. Gelly and D. Silver. Achieving master level play in 9 x 9 computer Go. In *Proceedings of the 23rd Conference on Artificial Intelligence, Nectar Track*, 2008.
- [50] S. Gelly and D. Silver. Monte-Carlo Tree Search and Rapid Action Value Estimation in computer Go. *Artificial Intelligence*, 175(11):1856–1875, 2011.

- [51] S. Gelly, Y. Wang, R. Munos, and O. Teytaud. Modification of UCT with patterns in Monte-Carlo Go, 2006. Technical Report RR-6062, INRIA.
- [52] T. Graf, U. Lorenz, M. Platzner, and L. Schaefers. Parallel Monte-Carlo Tree Search for HPC Systems. In *Proc. 17th Int. Euro. Conf. Parallel Distrib. Comput., LNCS 6853*, pages 365–376, Bordeaux, France, 2011.
- [53] A. Guez, D. Silver, and P. Dayan. Scalable and efficient Bayes-adaptive reinforcement learning based on Monte-Carlo Tree Search. *Journal of Artificial Intelligence Research (JAIR)*, 48:841–883, 2013.
- [54] D. Helmbold and A. Parker-Wood. All-moves-as-first heuristics in Monte-Carlo Go, 2009. Date retrieved: March 10, 2010.
- [55] J.-B. Hoock, C.-S. Lee, A. Rimmel, F. Teytaud, O. Teytaud, and M.-H. Wang. Intelligent Agents for the Game of Go. *IEEE Computational Intelligence Magazine*, 11 2010.
- [56] S. Huang. *New Heuristics for Monte Carlo Tree Search Applied to the Game of Go*. PhD thesis, National Taiwan Normal University, 2011.
- [57] S. Huang, R. Coulom, and S. Lin. Monte-Carlo simulation balancing in practice. In van den Herik et al. [103], pages 81–92.
- [58] S.-C. Huang, B. Arneson, R. Hayward, M. Müller, and J. Pawlewicz. MoHex 2.0: a pattern-based MCTS Hex player, 2013. To appear in Computers and Games 2013. 12 pp.
- [59] S.-C. Huang and M. Müller. Investigating the limits of Monte Carlo tree search methods in computer Go, 2013. To appear in Computers and Games 2013. 12 pp.
- [60] K. Ikeda and S. Viennot. Efficiency of static knowledge bias in Monte-Carlo tree search, 2013. To appear in Computers and Games 2013.
- [61] Z. Karnin, T. Koren, and O. Somekh. Almost optimal exploration in multi-armed bandits. In *ICML*, volume 28 of *JMLR Proceedings*, pages 1238–1246. JMLR.org, 2013.
- [62] H. Kato and I. Takeuchi. Parallel Monte-Carlo tree search with simulation servers. In *13th Game Programming Workshop (GPW-08)*, 2008.
- [63] E. Kaufmann, N. Korda, and R. Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In N. Bshouty, G. Stoltz, N. Vayatis, and T. Zeugmann, editors, *ALT*, volume 7568 of *Lecture Notes in Computer Science*, pages 199–213. Springer, 2012.
- [64] T. Keller and P. Eyerich. PROST: Probabilistic planning based on UCT. In L. McCluskey, B. Williams, J. Reinaldo Silva, and B. Bonet, editors, *ICAPS*. AAAI, 2012.

- [65] T. Keller and M. Helmert. Trial-based heuristic tree search for finite horizon MDPs. In D. Borrajo, S. Kambhampati, A. Oddi, and S. Fratini, editors, *ICAPS*. AAAI, 2013.
- [66] L. Kocsis and C. Szepesvári. Bandit based Monte-Carlo planning. In J. Fürnkranz, T. Scheffer, and M. Spiliopoulou, editors, *Machine Learning: ECML 2006*, volume 4212 of *Lecture Notes in Computer Science*, pages 282–293. Springer Berlin / Heidelberg, 2006.
- [67] T. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6:4–22, 1985.
- [68] M. Lanctot, V. Lisý, and M. Winands. Monte Carlo Tree Search in simultaneous move games with applications to Goofspiel. In T. Cazenave, M. Winands, and H. Iida, editors, *CGW@IJCAI*, volume 408 of *Communications in Computer and Information Science*, pages 28–43. Springer, 2013.
- [69] M. Lanctot, A. Saffidine, J. Veness, C. Archibald, and M. Winands. Monte Carlo *-Minimax search. In F. Rossi, editor, *IJCAI*, pages 580–586. IJCAI/AAAI, 2013.
- [70] C.-S. Lee, M.-H. Wang, G. Chaslot, J.-B. Hoock, A. Rimmel, O. Teytaud, S.-R. Tsai, S.-C. Hsu, and T.-P. Hong. The Computational Intelligence of MoGo Revealed in Taiwan’s Computer Go Tournaments. *IEEE Transactions on Computational Intelligence and AI in Games*, 1(1):73–89, 2009.
- [71] R. Lorentz. Amazons discover Monte-Carlo. In van den Herik et al. [105], pages 13–24.
- [72] J. Méhat and T. Cazenave. Combining UCT and NMCS for single-player GGP. *TCIAIG*, 2010.
- [73] R. Munos. From bandits to Monte-Carlo Tree Search: The optimistic principle applied to optimization and planning. *Foundations and Trends in Machine Learning*, 7(1):1–129, 2014.
- [74] H. Nakhost and M. Müller. Monte-Carlo exploration for deterministic planning. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009, Pasadena, California, USA, July 11-17, 2009*, pages 1766–1771, 2009.
- [75] J. Nijssen and M. Winands. Monte Carlo Tree Search for the hide-and-seek game Scotland Yard. *IEEE Trans. Comput. Intellig. and AI in Games*, 4(4):282–294, 2012.
- [76] R. Ramanujan, A. Sabharwal, and B. Selman. On adversarial search spaces and sampling-based planning. In R. Brafman, H. Geffner, J. Hoffmann, and H. Kautz, editors, *ICAPS*, pages 242–245. AAAI, 2010.

- [77] R. Ramanujan, A. Sabharwal, and B. Selman. Understanding sampling style adversarial search methods. In P. Grünwald and P. Spirtes, editors, *UAI*, pages 474–483. AUAI Press, 2010.
- [78] R. Ramanujan and B. Selman. Trade-offs in sampling-based adversarial planning. In F. Bacchus, C. Domshlak, S. Edelkamp, and M. Helmert, editors, *ICAPS*. AAAI, 2011.
- [79] A. Rimmel. *Bandit-based optimization on graphs with application to library performance tuning*. PhD thesis, Université Paris-Sud, Ecole doctorale d’informatique, December 2009.
- [80] A. Rimmel and F. Teytaud. Multiple overlapping tiles for contextual monte carlo tree search. In *2nd European event on Bio-inspired algorithms for GAMES (EvoGAMES’10)*, Istanbul, Turkey, 2010. Springer-Verlag.
- [81] A. Rimmel, F. Teytaud, and O. Teytaud. Biasing Monte-Carlo simulations through RAVE values. In van den Herik et al. [103], pages 59–68.
- [82] A. Rimmel, O. Teytaud, C.-S. Lee, S.-J. Yen, M.-H. Wang, and S.-R. Tsai. Current frontiers in computer Go. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):229–238, 2010.
- [83] H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58:527–535, 1952.
- [84] J. W. Romein, A. Plaat, H. E. Bal, and J. Schaeffer. Transposition Table Driven Work Scheduling in Distributed Search. In *AAAI-99*, pages 725–731, Orlando, 1999.
- [85] C. Rosin. Multi-armed bandits with episode context. *Ann. Math. Artif. Intell.*, 61(3):203–230, 2011.
- [86] C. Rosin. Nested rollout policy adaptation for Monte Carlo Tree Search. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One*, IJCAI’11, pages 649–654. AAAI Press, 2011.
- [87] Y. Sato and D. Takahashi. A shogi program based on Monte-Carlo tree search. In *13th Game Programming Workshop (GPW-08)*, pages 1–8, 2008.
- [88] M. Schadd, M. Winands, M. Tak, and J. Uiterwijk. Single-player Monte-Carlo tree search for SameGame. *Knowledge-Based Systems*, 34:3–11, 2012.
- [89] M. Schadd, M. Winands, H. J. van den Herik, G. Chaslot, and J. Uiterwijk. Single-player Monte-Carlo tree search. In van den Herik et al. [105], pages 1–12.

- [90] R. Segal. On the scalability of parallel UCT. In H. van den Herik, H. Iida, and A. Plaat, editors, *Computers and Games*, volume 6515 of *Lecture Notes in Computer Science*, pages 36–47. Springer Berlin / Heidelberg, 2011.
- [91] D. Silver. *Reinforcement Learning and Simulation-Based Search in Computer Go*. PhD thesis, University of Alberta, 2009.
- [92] D. Silver, R. Sutton, and M. Müller. Temporal-difference search in computer Go. *Machine Learning*, 87(2):183–219, 2012.
- [93] D. Silver and G. Tesauro. Monte-Carlo simulation balancing. In Danyluk et al. [37], pages 945–952.
- [94] D. Silver and J. Veness. Monte-Carlo planning in large POMDPs. In J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, editors, *NIPS*, pages 2164–2172. Curran Associates, Inc., 2010.
- [95] Y. Soejima, A. Kishimoto, and O. Watanabe. Evaluating root parallelization in Go. *Computational Intelligence and AI in Games, IEEE Transactions on*, 2(4):278–287, 2010.
- [96] Y. Soejima, A. Kishimoto, and O. Watanabe. Root parallelization of Monte Carlo tree search and its effectiveness in computer Go. In *14th Game Programming Workshop in Japan*, pages 27–33, 2010.
- [97] N. Sturtevant. An analysis of UCT in multi-player games. In van den Herik et al. [105], pages 37–49.
- [98] G. Tesauro and G. Galperin. On-line policy improvement using Monte-Carlo search. In *NIPS*, 1996.
- [99] F. Teytaud and O. Teytaud. Creating an Upper-Confidence-Tree program for Havannah. In van den Herik and Spronck [104], pages 65–74.
- [100] D. Tom. Investigating UCT and RAVE: steps towards a more robust method. Master’s thesis, University of Alberta, 2010.
- [101] D. Tom and M. Müller. A study of UCT and its enhancements in an artificial game. In van den Herik and Spronck [104], pages 55–64.
- [102] D. Tom and M. Müller. Computational experiments with the RAVE heuristic. In J. van den Herik, H. Iida, and A. Plaat, editors, *Computers and Games*, volume 6515 of *Lecture Notes in Computer Science*, pages 69–80. Springer Berlin / Heidelberg, 2011.
- [103] J. van den Herik, H. Iida, and A. Plaat, editors. *Computers and Games - 7th International Conference, CG 2010*, volume 6515 of *Lecture Notes in Computer Science*. Springer, 2011.

- [104] J. van den Herik and P. Spronck, editors. *Advances in Computer Games. 12th International Conference, ACG 2009, Pamplona, Spain, May 11-13, 2009*, volume 6048 of *Lecture Notes in Computer Science*. Springer, 2010.
- [105] J. van den Herik, X. Xu, Z. Ma, and M. Winands, editors. *Computers and Games, 6th International Conference, CG 2008*, volume 5131 of *Lecture Notes in Computer Science*. Springer, 2008.
- [106] G. Van Eyck. Move groups as a general enhancement for Monte Carlo Tree Search. Master’s thesis, University of Alberta, 2014.
- [107] G. Van Eyck and M. Müller. Revisiting move groups in Monte Carlo Tree Search. In H. van den Herik and A. Plaat, editors, *Advances in Computer Games*, volume 7168 of *Lecture Notes in Computer Science*, pages 13–23. Springer Berlin / Heidelberg, 2012.
- [108] M. Winands, Y. Björnsson, and J.-T. Saito. Monte-Carlo Tree Search solver. In van den Herik et al. [105], pages 25–36.
- [109] M. Winands, Y. Björnsson, and J.-T. Saito. Monte Carlo Tree Search in Lines of Action. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):239–250, 2010.
- [110] M. H. M. Winands and Y. Björnsson. $\alpha\beta$ -based play-outs in Monte-Carlo Tree Search. In S.-B. Cho, S. Lucas, and P. Hingston, editors, *CIG*, pages 110–117. IEEE, 2011.
- [111] M. Wistuba and L. Schmidt-Thieme. Move prediction in Go - modelling feature interactions using latent factors. In Ingo J. Timm and Matthias Thimm, editors, *KI2013*, volume 8077 of *Lecture Notes in Computer Science*, pages 260–271. Springer, 2013.
- [112] H. Yoshimoto, A. Kishimoto, T. Kaneko, and K. Yoshizoe. Depth-first UCT and its application to Go. In *Proc. 12th Game Programming Workshop (GPW 2007)*, pages 30–35. 2007. In Japanese.
- [113] H. Yoshimoto, K. Yoshizoe, T. Kaneko, A. Kishimoto, and K. Taura. Monte Carlo Go has a way to go. In *AAAI*, pages 1070–1075. AAAI Press, 2006.
- [114] K. Yoshizoe, A. Kishimoto, T. Kaneko, H. Yoshimoto, and Y. Ishikawa. Scalable distributed Monte-Carlo tree search. In D. Borrajo, M. Likhachev, and C. Linares López, editors, *Symposium on Combinatorial Search*, pages 180–187. AAAI Press, 2011.
- [115] P. Zhang and K.-H. Chen. Monte Carlo Go capturing tactic search. *New Mathematics and Natural Computation*, 4(3):359–367, 2008.

- [116] M. Zinkevich, M. Johanson, M. Bowling, and C. Piccione. Regret minimization in games with incomplete information. In J.C. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Neural Information Processing Systems (NIPS)*, pages 1729–1736, 2007.