

Shimon Whiteson

Research Interests

My research concerns artificial intelligence, with a particular focus on decision-theoretic planning and reinforcement learning. I am applying these methods to problems in information retrieval and robotics.

Education

- 2000–2007 **Ph.D.**, Department of Computer Science, *University of Texas at Austin*.
Dissertation title: *Adaptive Representations for Reinforcement Learning*
Advisor: Peter Stone
- 1996–2000 **B.A.**, *Rice University*, Houston, Texas.
Majors: English and Computer Science
Presidential Honour Roll all eight semesters
Graduated *Magna Cum Laude*

Employment

Department of Computer Science, *University of Oxford*

2015-present **Associate Professor & Tutorial Fellow of St. Catherine's College**

Department of Computer Science, *University of California at Irvine*

2014–2015 **Visiting Faculty Member**

Informatics Institute, *University of Amsterdam*

2014–2015 **Associate Professor, Level II** (*Universitair Hoofddocent II*)

2012 **Interim Group Leader, Intelligent Autonomous Systems Group**

2011–2014 **Tenured Assistant Professor, Level I** (*Universitair Docent I*)

2010–2011 **Tenured Assistant Professor, Level II** (*Universitair Docent II*)

2007–2010 **Temporary Assistant Professor, Level II** (*Universitair Docent II*)

Department of Computer Science, *University of Texas at Austin*

2007 **Postdoctoral Fellow**

2006–2007 **IBM Ph.D. Fellow**

2005–2006 **Assistant Instructor**

2003–2005 **Research Assistant**

2000–2002 **Teaching Assistant**

Professional Training

- 2015 **Undergraduate Admissions Course**, University of Oxford.
- 2010 **Performance Evaluation (Jaargesprek) Course**, University of Amsterdam.
- 2008–2009 **Basic Education Qualification (BKO)**, University of Amsterdam.
- 2002 **Supervised Teaching in Computer Science**, University of Texas at Austin.

Grants Awarded

As Principal Investigator

- 2017 **End-to-End Deep Model-Based Reinforcement Learning**, *Google Faculty Research Award*, only 16% of applicants funded, total budget \$57K.
- 2015 **Coevolutionary Policy Search**, *ERC Starting Grant*, a highly selective 5-year grant awarded annually to the top young researchers in Europe, total budget €1.5M.
- 2014 **VIDI Grant**, *Highly selective 5-year personal grant from the NWO (Dutch national science foundation)*, awarded annually to the best mid-career researchers in the Netherlands, total budget €800K.
- 2013 **Telepresence Reinforcement Learning Social Agent**, *3-year EU-FP7 STREP grant in which I serve as project leader and scientific coordinator*, total budget: €3.8M (€3M EU contribution), UvA budget: €935K (€780K EU contribution).
- 2012 **Self-Optimising Tracking Systems**, *4-year STW (Dutch national technology foundation) Open Technology Program grant for €189K*.
- 2011 **Decision-Theoretic Control for Network Capacity Allocation Problems**, *4-year NWO (Dutch national science foundation) Free Competition grant for €180K*.

As Co-Investigator

- 2014 **Leveraging Data Reuse for Efficient Ranker Evaluation in Information Retrieval**, *3-year Microsoft Research PhD Fellowship*, total budget €103K (with Maarten de Rijke).
- 2012 **Modeling and Learning from Implicit Feedback in Information Retrieval**, *4-year NWO (Dutch national science foundation) Free Competition grant for €180K (with Maarten de Rijke)*.
- 2012 **Multi-Robot Cognitive Systems Operating in Hospitals (MOnarCH)**, *3-year EU-FP7 STREP grant*, total budget: €4.4M (€3.4M EU contribution), UvA budget: €579K (€446K EU contribution) (with Gwenn Englebienne).
- 2012 **Linguistically Motivated Semantic aggregation engines (LiMoSINE)**, *3-year EU-FP7 STREP grant*, total budget: €3.4M (€2.5M EU contribution), UvA budget: €854K (€678K EU contribution) (with Maarten de Rijke and Edgar Meij).

Talks

Invited Talks

- 2017 Microsoft Research AI Summer School, Cambridge
- 2017 Invited Tutorial at British Machine Vision Conference, London
- 2016 NIPS Workshop on Learning, Inference & Control of Multi-Agent Systems, Barcelona
- 2016 AAI Symposium on Challenges for Real-World Multiagent Learning, Stanford, USA
- 2015 Lab Tutorial, Microsoft Research, Cambridge, UK
- 2012 Amsterdam Scientific & Educational Symposium on Humans & Information Technology
- 2011 SIKS Advanced Course on Computational Intelligence, Utrecht, Netherlands
- 2010 Next Generation Infrastructure Academy, Tegelen, Netherlands
- 2009 Foundation for Neural Networks Seminar, Radboud University of Nijmegen, Netherlands
- 2008 Symposium on AI & Traffic, Radboud University of Nijmegen, Netherlands

Seminars

- 2016 Man AHL, London, UK
- 2016 University of Liverpool, UK
- 2015 Google DeepMind, London, UK
- 2015 Robotics Seminar, University of Oxford, UK
- 2015 Forum for Artificial Intelligence, University of Texas at Austin, USA
- 2015 University of California at Irvine, USA
- 2015 University of Southern California, USA
- 2013 Gatsby Institute, University College London, UK
- 2013 University of California at Irvine, USA
- 2012 Erasmus Research Institute of Management, Rotterdam, Netherlands
- 2012 Department of Software Technology, Technical University of Delft, Netherlands
- 2009 Department of Knowledge Engineering, Maastricht University, Netherlands
- 2009 Forum for Artificial Intelligence, University of Texas at Austin, USA
- 2007 Yahoo! Research, Sunnyvale, California, USA
- 2007 IBM Research, Hawthorne, New York, USA
- 2007 Fordham University, New York, New York, USA
- 2007 University of Texas at Arlington, USA
- 2007 Dalle Molle Institute for Artificial Intelligence, Lugano, Switzerland
- 2007 University of Amsterdam, Netherlands
- 2005 Toyota Technological Institute, Chicago, Illinois, USA

Press

- 2016 Article in The Verge: "Can deep learning help solve lip reading?"
- 2016 Article in Live Science: "The Real Reason AI Won't Take Over Anytime Soon"

- 2015 Article in Business Insider: "We are all going to be cyborgs' if humanity wants to solve its biggest problems"
- 2014 Interview in CVZ Magazine about EU telepresence robotics project: "We willen een socialvaardige robot maken" ("We want to make a socially intelligent robot")
- 2013 Interview in Het Parool (main Amsterdam newspaper) about EU telepresence robotics project: "Uiteindelijk worden we allemaal cyborgs" ("Eventually we'll all be cyborgs")

Research Supervision

Current Postdoctoral Researchers

- 2017–2019 Tim Rocktäschel, *Deep model-based reinforcement learning*
- 2015–2017 Kamil Ciosek, *Policy gradient reinforcement learning*

Current PhD Students

- 2011–2017 Matthijs Snel, *Multi-task reinforcement learning*
- 2013–2017 Yash Satsangi, *Decision-theoretic active perception*
- 2014–2018 Kyriacos Shiarlis, *Reinforcement learning for telepresence robots*
- 2015–2018 Yannis Assael, *Deep reinforcement learning*
- 2015–2018 Jakob Foerster, *Deep multi-agent reinforcement learning*
- 2015–2018 Supratik Paul, *Policy search reinforcement learning*
- 2016–2019 Max Igl, *Deep model-based reinforcement learning*
- 2016–2019 Greg Farquhar, *Optimal advantage learning*
- 2016–2019 Tabish Rashid, *Hierarchical reinforcement learning*
- 2016–2019 Matt Fellows, *Deep multi-agent reinforcement learning*

Former Postdoctoral Researchers

- 2016 Diederik Roijers, *Multi-objective decision-making for social robotics*
- 2015–2017 João Messias, *Body-pose control for social robotics*
- 2012–2014 Sander Bakkes, *Machine learning for serious games*
- 2012–2015 Masrour Zoghi, *Relative bandits for information retrieval*

Former PhD Students

- 2012–2016 Diederik Roijers, *Multi-objective reinforcement learning for network capacity allocation*
- 2011–2015 Anne Schuth, *Online evaluation and learning to rank for information retrieval*
- 2011–2015 Guangliang Li, *Socially intelligent reinforcement learning*
- 2009–2013 Katja Hofmann, *Fast and Reliable Online Learning to Rank for Information Retrieval*
- 2008–2011 Harm van Seijen, *Reinforcement Learning under Space and Time Constraints*

Professional Activities

Editing

- 2010-present **Associate Editor**, Journal of Autonomous Agents and Multi-Agent Systems
- 2011-present **Editorial Board Member**, Machine Learning Journal

Department of Computer Science – University of Oxford
 ☎ +44 (0) 7443 287 697 • ✉ shimon.whiteson@cs.ox.ac.uk
 • www.cs.ox.ac.uk/people/shimon.whiteson

- 2010-present **Editorial Board Member**, Journal of Artificial Intelligence Research
- 2009-2011 **Guest Editor**, Machine Learning Journal, Special Issue on Empirical Evaluations in Reinforcement Learning
- Organizing**
- 2016 **Organizer and Presenter**, EASSS Tutorial on Multi-Objective Decision-Making
- 2016 **Organizer and Presenter**, ICAPS Tutorial on Multi-Objective Decision-Making
- 2015 **Organizer and Presenter**, IJCAI Tutorial on Multi-Objective Decision-Making
- 2015 **Primary Organizer**, ICRA Workshop on Machine Learning for Social Robotics
- 2008 **Organizing Chair**, Reinforcement Learning Competition
- 2005 **Organizing Committee Member**, Reinforcement Learning Benchmarks and Bake-Offs

Area Chairs

- 2017 International Conference on Machine Learning (ICML)

Senior Program Committees

- 2015,2016 Int'l. Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)
- 2014-17 AAAI Conference on Artificial Intelligence (AAAI)
- 2009,-13,-15,-16 International Joint Conference on Artificial Intelligence (IJCAI)

Program Committees

- 2012,2013,2014 Neural Information Processing Systems (NIPS)
- '06,09,10,12-16 International Conference on Machine Learning (ICML)
- 2006-2011 AAAI Conference on Artificial Intelligence (AAAI)
- 2004,-06,-07,-11 Int'l. Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)
- 2005,2007,2008 European Conference on Machine Learning (ECML)
- 2004 Genetic and Evolutionary Computation Conference (GECCO)

Proposal Reviewing

- 2014 Dutch National Technology Foundation (STW)
- 2013 Austrian Science Fund
- 2011 United States Air Force Office of Scientific Research
- 2010-2011 Israel Science Foundation
- 2009 US-Israel Binational Science Foundation

Teaching

Tutoring at St. Catherine's College, Oxford

- 2015–2017 **Discrete Mathematics, Continuous Mathematics, Intelligent Systems, Imperative Programming II.**

Courses Given at the University of Amsterdam

- 2013–2014 **Advanced Topics in Autonomous Agents, Master of AI.**
Average student course rating in 2013: 4.7/5.0 (lecturer), 4.1/5.0 (overall course)

Department of Computer Science – University of Oxford

☎ +44 (0) 7443 287 697 • ✉ shimon.whiteson@cs.ox.ac.uk

• www.cs.ox.ac.uk/people/shimon.whiteson

- 2012–2013 **Autonomous Agents**, *Master of AI*.
Average student course rating in 2013: 4.6/5.0 (lecturer), 4.1/5.0 (overall course)
Average student course rating in 2012: 4.8/5.0 (lecturer), 4.5/5.0 (overall course)
- 2010–2012 **Autonomous Agents & Multi-Agent Systems**, *Master of AI*.
(with Albert Salah in 2010 and 2011)
Designed and implemented this new advanced course
Average student course rating in 2012: 4.0/5.0 (lecturer), 3.5/5.0 (overall course)
Average student course rating in 2011: 4.1/5.0 (lecturer), 3.9/5.0 (overall course)
Average student course rating in 2010: 4.3/5.0 (lecturer), 3.5/5.0 (overall course)
- 2009–2011 **Multi-Agent Systems**, *Master of AI*.
Average student course rating in 2011: 4.7/5.0 (lecturer), 4.2/5.0 (overall course)
Average student course rating in 2010: 4.9/5.0 (lecturer), 4.7/5.0 (overall course)
Average student course rating in 2009: 4.8/5.0 (lecturer), 4.5/5.0 (overall course)
- 2009–2010 **Decision Making in Intelligent Systems**, *Bachelor of AI*.
Average student course rating in 2009: 8.4/10.0 (lecturer), 8.4/10.0 (overall course)
Average student course rating in 2008: 8.5/10.0 (lecturer), 7.9/10.0 (overall course)
- 2008–2009 **Multi-Agent Systems and Distributed AI**, *Master of AI*.
Average student course rating in 2009: 4.92/5.0 (lecturer), 4.62/5.0 (overall course)
Average student course rating in 2008: 4.8/5.0 (lecturer), 4.4/5.0 (overall course)

Courses Given at the Amsterdam University College

- 2012 **Machine Learning**, (with Maarten van Someren).
2011 **Text Mining and Collective Intelligence**, (with Maarten de Rijke and Cees Snoek).

Courses Given at the University of Texas at Austin

- 2005–2006 **Computer Programming: C++**, *Bachelor of CS*.
Average student course rating: 3.6/5.0 (2005) and 4.5/5.0 (2006)

Teaching Administration

- 2009–2015 **Gaming Track Coordinator**, *Master of AI*, University of Amsterdam.
2009–2015 **Management Committee Member**, *Master of AI*, University of Amsterdam.

Awards

- 2017 Google Faculty Research Award
2016 Best Paper Nomination, AAMAS Conference
2009,2010 Nominated for Teacher of the Year, University of Amsterdam
2008 First Place (with Rogier Koppejan) in Helicopter Hovering at the RL Competition
2006–2007 IBM PhD Fellowship
2006 Best Paper Award, GECCO Conference, Genetic Algorithms Track
2006 Best Paper Award, GECCO Conference, Graduate Student Workshop
2005 Young Investigator, Cognitive Systems Conference
2000–present Member of Phi Beta Kappa
2000 *Magna Cum Laude* Graduation, Rice University
1996–2000 Presidential Honour Roll, Rice University
1996–2000 Max Roy Full-Tuition Scholarship, Rice University
1998 W. L. Moody, Jr. Scholarship for Engineering, Rice University

Department of Computer Science – University of Oxford
☎ +44 (0) 7443 287 697 • ✉ shimon.whiteson@cs.ox.ac.uk
• www.cs.ox.ac.uk/people/shimon.whiteson

- 1997 Louis Walsh Scholarship for Engineering, Rice University
- 1995 Second Place, National Championships, Lincoln-Douglas Debate
- 1994 First Place, New Mexico State Championships, Extemporaneous Speaking

Publications

Note: *technical reports, workshop papers, other short papers, and some conference papers are omitted here for the sake of brevity. A complete publication list is available [here](#).*

Books

- [1] Diederik Roijers and Shimon Whiteson. *Multi-Objective Decision Making*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan and Claypool, California, USA, 2017. doi:10.2200/S00765ED1V01Y201704AIM034.
- [2] Shimon Whiteson. *Adaptive Representations for Reinforcement Learning*, volume 291 of *Studies in Computational Intelligence*. Springer, Berlin, Germany, 2010.

Journal Articles

- [3] Yash Satsangi, Shimon Whiteson, Frans Oliehoek, and Matthijs Spaan. *Autonomous Robots*, 2017. To appear.
- [4] Guangliang Li, Shimon Whiteson, W. Bradley Knox, and Hayley Hung. Social interaction for efficient agent learning from human reward. *Autonomous Agents and Multi-Agent Systems*, 2017. To appear.
- [5] Guangliang Li, Shimon Whiteson, W. Bradley Knox, and Hayley Hung. Using informative behavior to increase engagement while learning from human reward. *Autonomous Agents and Multi-Agent Systems*, 30(5):826–848, 2016.
- [6] Diederik Roijers, Shimon Whiteson, and Frans Oliehoek. Computing convex coverage sets for faster multi-objective coordination. *Journal of Artificial Intelligence Research*, 52:399–443, 2015.
- [7] Harm van Seijen, Shimon Whiteson, and Leon Kester. Efficient abstraction selection in reinforcement learning. *Computational Intelligence*, 30(4):657–699, 2014.
- [8] Matthijs Snel and Shimon Whiteson. Learning potential functions and their representations for multi-task reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 28(4):637–681, 2014.
- [9] Diederik Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.
- [10] Katja Hofmann, Shimon Whiteson, and Maarten de Rijke. Fidelity, soundness, and efficiency of interleaved comparison methods. *Transactions on Information Systems*, 31(4):17:1–43, 2013.
- [11] Frans Oliehoek, Matthijs Spaan, Christopher Amato, and Shimon Whiteson. Incremental clustering and expansion for faster optimal planning in decentralized POMDPs. *Journal of Artificial Intelligence Research*, 46:449–509, 2013.
- [12] Katja Hofmann, Shimon Whiteson, and Maarten de Rijke. Balancing exploration and exploitation in listwise and pairwise online learning to rank for information retrieval. *Information Retrieval*, 16(1):63–90, 2013.
- [13] Rogier Koppejan and Shimon Whiteson. Neuroevolutionary reinforcement learning for generalized control of simulated helicopters. *Evolutionary Intelligence*, 4:219–241, 2011.
- [14] Shimon Whiteson and Michael L. Littman. Introduction to the special issue on empirical evaluations in reinforcement learning. *Machine Learning*, 84(1):1–6, 2011.

Department of Computer Science – University of Oxford

☎ +44 (0) 7443 287 697 • ✉ shimon.whiteson@cs.ox.ac.uk

• www.cs.ox.ac.uk/people/shimon.whiteson

- [15] Harm van Seijen, Shimon Whiteson, Hado van Hasselt, and Marco Wiering. Exploiting best-match equations for efficient reinforcement learning. *Journal of Machine Learning Research*, 12:2045–2094, 2011.
- [16] Shimon Whiteson, Matthew E. Taylor, and Peter Stone. Critical factors in the empirical performance of temporal difference and evolutionary methods for reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 21(1):1–27, 2010.
- [17] Shimon Whiteson and Daniel Whiteson. Machine learning for event selection in high energy physics. *Engineering Applications of Artificial Intelligence*, 22:1203–1217, 2009.
- [18] Aaltonen et al. (including Shimon Whiteson). Measurement of the top quark mass with dilepton events selected using neuroevolution at CDF. *Physical Review Letters*, 102(15):2001, 2009.
- [19] Shimon Whiteson, Matthew E. Taylor, and Peter Stone. Empirical studies in action selection with reinforcement learning. *Adaptive Behavior*, 15(1):33–50, 2007.
- [20] Shimon Whiteson and Peter Stone. Evolutionary function approximation for reinforcement learning. *Journal of Machine Learning Research*, 7:877–917, 2006.
- [21] Shimon Whiteson, Nate Kohl, Risto Miikkulainen, and Peter Stone. Evolving keepaway soccer players through task decomposition. *Machine Learning*, 59(1):5–30, 2005.
- [22] Shimon Whiteson and Peter Stone. Adaptive job routing and scheduling. *Engineering Applications of Artificial Intelligence*, 17(7):855–869, 2004.

Selected Conference Papers: General Artificial Intelligence

- [23] Jo ao Messias and Shimon Whiteson. Utile context tree weighting. In *NIPS 2017: Proceedings of the Thirty-First Annual Conference on Neural Information Processing Systems*, December 2017. To appear.
- [24] Yash Satsangi, Shimon Whiteson, Frans Oliehoek, and Henri Bouma. Real-time resource allocation for tracking systems. In *UAI 2017: Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 3220–3227, July 2017.
- [25] Jakob Foerster, Nantas Nardelli, Greg Farquhar, Phil Torr, Pushmeet Kohli, and Shimon Whiteson. Stabilising experience replay for deep multi-agent reinforcement learning. In *ICML 2017: Proceedings of the Thirty-Fourth International Conference on Machine Learning*, June 2017.
- [26] Kamil Ciosek and Shimon Whiteson. Offer: Off-environment reinforcement learning. In *AAAI 2017: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 1819–1825, February 2017.
- [27] Jakob Foerster, Yannis Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In *NIPS 2016: Proceedings of the Thirtieth Annual Conference on Neural Information Processing Systems*, December 2016.
- [28] Yash Satsangi, Shimon Whiteson, and Frans Oliehoek. PAC greedy maximization with efficient bounds on information gain for sensor selection. In *IJCAI 2016: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 3220–3227, July 2016.
- [29] Masrour Zoghi, Zohar Karnin, Shimon Whiteson, and Maarten de Rijke. Copeland dueling bandits. In *NIPS 2015: Proceedings of the Twenty-Ninth Annual Conference on Neural Information Processing Systems*, December 2015.
- [30] Diederik Roijers, Shimon Whiteson, and Frans Oliehoek. Point-based planning for multi-objective POMDPs. In *IJCAI 2015: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*, pages 1666–1672, July 2015.
- [31] Yash Satsangi, Shimon Whiteson, and Frans Oliehoek. Exploiting submodular value functions for faster dynamic sensor selection. In *AAAI 2015: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 3356–3363, January 2015.

- [32] Masrour Zoghi, Shimon Whiteson, Remi Munos, and Maarten de Rijke. Relative upper confidence bound for the K-armed dueling bandit problem. In *ICML 2014: Proceedings of the Thirty-First International Conference on Machine Learning*, pages 10–18, June 2014.
- [33] Frans Oliehoek, Shimon Whiteson, and Matthijs Spaan. Exploiting structure in cooperative Bayesian games. In *UAI 2012: Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, pages 654–664, August 2012.
- [34] Shimon Whiteson and Daniel Whiteson. Stochastic optimization for collision selection in high energy physics. In *IAAI 2007: Proceedings of the Nineteenth Annual Innovative Applications of Artificial Intelligence Conference*, pages 1819–1825, July 2007.
- [35] Matthew E. Taylor, Shimon Whiteson, and Peter Stone. Temporal difference and policy search methods for reinforcement learning: An empirical comparison. In *AAAI 2007: Proceedings of the Twenty-Second National Conference on Artificial Intelligence*, pages 1675–1678, July 2007. (Nectar Track).
- [36] Shimon Whiteson and Peter Stone. Sample-efficient evolutionary function approximation for reinforcement learning. In *AAAI 2006: Proceedings of the Twenty-First National Conference on Artificial Intelligence*, pages 518–523, July 2006.
- [37] Shimon Whiteson and Peter Stone. Towards autonomic computing: Adaptive job routing and scheduling. In *IAAI 2004: Proceedings of the Sixteenth Annual Innovative Applications of Artificial Intelligence Conference*, pages 916–922, July 2004.

Selected Conference Papers: Autonomous Agents

- [38] Kyriacos Shiarlis, Jo ao Messias, and Shimon Whiteson. Acquiring social interaction behaviours for telepresence robots via deep learning from demonstration. In *IROS 2017: Proceedings of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems*, September 2017. To appear.
- [39] Kyriacos Shiarlis, Jo ao Messias, and Shimon Whiteson. Rapidly exploring learning trees. In *ICRA 2017: Proceedings of the 2017 IEEE International Conference on Robotics and Automation*, May 2017.
- [40] Kyriacos Shiarlis, Jo ao Messias, and Shimon Whiteson. Inverse reinforcement learning from failure. In *AAMAS 2016: Proceedings of the Fifteenth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 1060–1068, May 2016. **Nominated for Best Student Paper.**
- [41] Diederik Roijers, Joris Scharpff, Matthijs Spaan, Frans Oliehoek, Mathijs De Weerd, and Shimon Whiteson. Bounded approximations for linear multi-objective planning under uncertainty. In *ICAPS 2014: Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling*, pages 262–270, June 2014.
- [42] Diederik Roijers, Shimon Whiteson, and Frans Oliehoek. Linear support for multi-objective coordination graphs. In *AAMAS 2014: Proceedings of the Thirteenth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 1297–1304, May 2014.
- [43] Guangliang Li, Hayley Hung, Shimon Whiteson, and W. Bradley Knox. Using informative behavior to increase engagement in the TAMER framework. In *AAMAS 2013: Proceedings of the Twelfth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 909–916, May 2013.
- [44] Frans Oliehoek, Shimon Whiteson, and Matthijs Spaan. Approximate solutions for factored Dec-POMDPs with many agents. In *AAMAS 2013: Proceedings of the Twelfth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 563–570, May 2013.
- [45] Karun Rao and Shimon Whiteson. V-MAX: Tempered optimism for better PAC reinforcement learning. In *AAMAS 2012: Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 375–382, June 2012.

Department of Computer Science – University of Oxford

☎ +44 (0) 7443 287 697 • ✉ shimon.whiteson@cs.ox.ac.uk

• www.cs.ox.ac.uk/people/shimon.whiteson

- [46] Frans Oliehoek, Shimon Whiteson, and Matthijs Spaan. Lossless clustering of histories in decentralized POMDPs. In *AAMAS 2009: Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 577–584, May 2009.
- [47] Frans Oliehoek, Matthijs Spaan, Shimon Whiteson, and Nikos Vlassis. Exploiting locality of interaction in factored Dec-POMDPs. In *AAMAS 2008: Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 517–524, May 2008.
- [48] Matthew E. Taylor, Shimon Whiteson, and Peter Stone. Transfer via inter-task mappings in policy search reinforcement learning. In *AAMAS 2007: Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 156–163, May 2007.
- [49] Shimon Whiteson and Peter Stone. Concurrent layered learning. In *AAMAS 2003: Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 193–200, July 2003.

Selected Conference Papers: Information Retrieval

- [50] Anne Schuth, Harrie Oosterhuis, Shimon Whiteson, and Maarten de Rijke. Multileave gradient descent for fast online learning to rank. In *WSDM 2016: Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, pages 457–466, February 2016.
- [51] Artem Grotov, Shimon Whiteson, and Maarten de Rijke. Bayesian ranker comparison based on historical user interactions. In *SIGIR 2015: Proceedings of the Thirty-Eighth Annual ACM SIGIR Conference*, pages 273–282, August 2015.
- [52] Masrour Zoghi, Shimon Whiteson, and Maarten de Rijke. MergeRUCB: A method for large-scale online ranker evaluation. In *WSDM 2015: Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, pages 17–26, February 2015.
- [53] Anne Schuth, Floor Sietsma, Shimon Whiteson, Damien Lefortier, and Maarten de Rijke. Multi-leaved comparisons for fast online evaluation. In *CIKM 2014: Proceedings of the Twenty-Third Conference on Information and Knowledge Management*, pages 71–80, November 2014.
- [54] Masrour Zoghi, Shimon Whiteson, Maarten de Rijke, and Remi Munos. Using confidence bounds for efficient on-line ranker evaluation. In *WSDM 2014: Proceedings of the Seventh ACM International Conference on Web Search and Data Mining*, pages 73–82, February 2014.
- [55] Katja Hofmann, Anne Schuth, Shimon Whiteson, and Maarten de Rijke. Reusing historical interaction data for faster online learning to rank for IR. In *WSDM 2013: Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, pages 183–192, February 2013.
- [56] Katja Hofmann, Shimon Whiteson, and Maarten de Rijke. A probabilistic method for inferring preferences from clicks. In *CIKM 2011: Proceedings of the Twentieth Conference on Information and Knowledge Management*, pages 249–258, October 2011.

Selected Conference Papers: Stochastic Optimization

- [57] Thomas van den Berg and Shimon Whiteson. Critical factors in the performance of HyperNEAT. In *GECCO 2013: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 759–766, July 2013.
- [58] Steijn Kistemaker and Shimon Whiteson. Critical factors in the performance of novelty search. In *GECCO 2011: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 965–972, July 2011.
- [59] Matthijs Snel and Shimon Whiteson. Multi-task evolutionary shaping without pre-specified representations. In *GECCO 2010: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1031–1038, July 2010.

- [60] Rogier Koppejan and Shimon Whiteson. Neuroevolutionary reinforcement learning for generalized helicopter control. In *GECCO 2009: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 145–152, July 2009.
- [61] Shimon Whiteson and Peter Stone. On-line evolutionary computation for reinforcement learning in stochastic domains. In *GECCO 2006: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1577–1584, July 2006.
- [62] Matthew E. Taylor, Shimon Whiteson, and Peter Stone. Comparing evolutionary and temporal difference methods in a reinforcement learning domain. In *GECCO 2006: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1321–1328, July 2006. **Best Paper Award, Genetic Algorithms Track.**
- [63] Shimon Whiteson, Peter Stone, Kenneth O. Stanley, Risto Miikkulainen, and Nate Kohl. Automatic feature selection in neuroevolution. In *GECCO 2005: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1225–1232, June 2005.
- [64] Shimon Whiteson, Nate Kohl, Risto Miikkulainen, and Peter Stone. Evolving robocup keep-away players through task decomposition. In *GECCO 2003: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 356–368, July 2003.

Book Chapters

- [65] Shimon Whiteson. *Evolutionary Computation for Reinforcement Learning*, pages 325–355. Springer, Berlin, Germany, 2012.
- [66] Harm van Seijen, Shimon Whiteson, and Leon Kester. *Switching between Representations in Reinforcement Learning*, pages 65–84. Studies in Computational Intelligence. Springer, Berlin, Germany, 2010.
- [67] Bram Bakker, Shimon Whiteson, Leon Kester, and Frans Groen. *Traffic Light Control by Multi-agent Reinforcement Learning Systems*, pages 475–510. Studies in Computational Intelligence. Springer, Berlin, Germany, 2010.

Magazine & Newsletter Articles

- [68] Shimon Whiteson, Brian Tanner, and Adam White. The reinforcement learning competitions. *AI Magazine*, 31(2):81–94, 2010.
- [69] Katja Hofmann, Shimon Whiteson, Anne Schuth, and Maarten de Rijke. Learning to rank for information retrieval from user interactions. *SIGWEB Newsletter*, (Spring):1–7, April 2014.