Shimon Whiteson

Research Interests

My research concerns artificial intelligence, with a particular focus on machine learning techniques such as reinforcement learning and learning from demonstration. I am applying these methods to problems in robotics, autonomous vehicles, and multi-agent systems.

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*

2018-present **Professor & Tutorial Fellow of St. Catherine's College**
2015-2018 **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**

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

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

2018 NIPS Workshop on Adaptive Learning Agents, Stockholm
2018 Keynote speaker at the British Machine Vision Association Technical Meeting, London
2017,2018 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 AAAI 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

2018 Podcast in Oxford Sparks Big Questions: "How do you teach a robot social cues?"

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

2018-2019 Wendelin Böhmer, *Multi-agent reinforcement learning*

2015-2018 Kamil Ciosek, *Policy gradient reinforcement learning*

Current PhD Students

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*

2017-2020 Anuj Mahajan, *Variational reinforcement learning*

2017-2020 Luisa Zintgraf, *Meta reinforcement learning*

2017-2020 Jelena Luketina, *NLP for reinforcement learning*

Former Postdoctoral Researchers

2017–2018 Tim Rocktäschel, *Deep model-based reinforcement learning*

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

2014–2018 Kyriacos Shiarlis, *Reinforcement learning for telepresence robots*

2013–2017 Yash Satsangi, *Decision-theoretic active perception*

2011–2017 Matthijs Snel, *Multi-task reinforcement learning*

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
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

2017,2018 International Conference on Machine Learning (ICML)

Senior Program Committees

2015,2016,2018 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

2018 Polish Science Foundation
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

2015–2018 **Discete Mathematics, Continuous Mathematics, Imperative Programming II, Intelligent Systems, Machine Learning**.

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)

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

2018	The Outstanding Student Paper of AAAI-2018, AAAI Conference
2017	Google Faculty Research Award
2017	NVIDIA's NVAIL Pioneering Research Award, NIPS Conference
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
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

A complete publication list is available here.