

Curriculum Vitae of Timothy A. Mann

Google DeepMind
London, UK

E-mail: timothymann@google.com
Web: <http://www.timothyamann.com>

Education

Ph.D., Computer Science, Texas A&M University, 2012.
Dissertation: Scaling Up Reinforcement Learning without Sacrificing Optimality by Constraining Exploration
Advisor: Dr. Yoonsuck Choe

B.A., Computer Science & Mathematics (Philosophy minor), SUNY Potsdam, 2007.
Magna Cum Laude

Research Interests

adversarial/robust machine learning (decision making under uncertainty and big data)
personalization and recommender systems
reinforcement learning and multiarmed bandits
nonstationary/transfer/multitask/lifelong learning

Experience

Staff Research Scientist November 2018 – Present
DeepMind London, UK
Leading a team of research scientists and engineers to develop machine learning solutions to applied research challenges. Connecting research solutions to products where they are most impactful.

Senior Research Scientist December 2016 – October 2018
Google DeepMind London, UK
Research and development of algorithms for end-to-end optimization of sequential decision making problems that scale to complex Google-sized problems. Incorporating research ideas into products and publishing results. Transitioning reinforcement learning from an interesting research area to a full profession.

Research Scientist February 2015 – December 2016
Google DeepMind London, UK
Research and engineering towards scaling up reinforcement learning to complex Google-scale problems. Moving reinforcement learning from an interesting research area to a full profession.

Senior Researcher June 2014 – January 2015
Technion, Electrical Engineering / Shie Mannor Haifa, Israel
Scaling up reinforcement learning and approximate dynamic programming to large, complex, real-world control problems. Investigating how machine learning can be used to support the needs of on-line social communities.

Postdoctoral Fellow January 2013 – May 2014
Technion, Electrical Engineering / Shie Mannor Haifa, Israel
Theoretical and experimental investigation applying hierarchies of abstract actions to planning and reinforcement learning. Applying reinforcement learning and approximate dynamic programming to solve large scale, complex control problems.

Data Mining Consultant October 2012 – December 2012
Supervisor: Valerie Taylor College Station, TX
Performed exploratory analysis on data collected from a multi-year subprogram of the Super Computing conference (<http://supercomputing.org>). Generated a thorough report examining multiple aspects of how participation in the subprogram impacted participation in the overall conference.

Texas A&M Dissertation Fellowship
Texas A&M University

Spring 2012 – Fall 2012
College Station, TX

Investigated how learning systems can acquire and transfer knowledge from previously encountered tasks to efficiently solve novel tasks. Invented theoretically motivated algorithms for knowledge transfer and compared them empirically with established single task and multitask algorithms to find success and failure conditions. Compiled results into dissertation and articles for publication.

East Asia and Pacific Summer Institutes
National Science Foundation

Summer 2011
Daegu, South Korea

Investigated the role of action in autonomous improvement of binocular depth perception. Discovered actions that preserve perceptual invariance can expose inconsistent depth estimates and are therefore useful for autonomously improving depth estimation.

Computer Vision Consultant
VoRo Technologies LLC.

Fall 2009
Austin, TX

Investigated existing algorithms and developed software tools for markerless pose estimation and action classification from multiple cameras.

Research Experience for Undergraduates
National Science Foundation

Summer 2006
Texas A&M University

Compared biological neural network structure, as a whole, using graph theoretical techniques, with and without delay. Delay was simulated by matrix augmentation. Sum of Shortest paths was used to compare the effect of adding delay to three classes of network: random, scale-free, and small world. Found that small world networks are most affected by introducing delay.

Research Contractor
Naval Research Laboratory

Summer 2005
Washington, DC

Investigated reliable multicast, wrote a survey of available protocols, and made a recommendation as to which protocol should be selected for a project by the Distributed Computing Group at the Naval Research Laboratory.

Publications : Book Chapters

Ji Ryang Chung, Jaerock Kwon, **Timothy A. Mann**, and Yoonsuck Choe. Evolution of time in neural networks: From the present to the past, and forward to the future. In A. Ravishankar Rao and Guillermo A. Cecchi, editors, *The Relevance of the Time Domain to Neural Network Models*, Springer, Berlin, 2011.

Publications : Journal Publications

Timothy A. Mann, Doina Precup, and Shie Mannor. Approximate Value Iteration with Temporally Extended Actions. *Journal of Artificial Intelligence Research*, 2015.

Timothy A. Mann, Yunjung Park, Sungmoon Jeong, Minho Lee, and Yoonsuck Choe. Autonomous and Interactive Improvement of Binocular Visual Depth Estimation through Sensorimotor Interaction. *IEEE Transactions on Autonomous Mental Development*, 2012.

Publications : Conference Proceedings & Other

1. Carlos Riquelme, Hugo Penedones, Damien Vincent, Hartmut Maennel, Sylvain Gelly, **Timothy A. Mann**, Andre Barreto, and Gergely Neu. Adaptive Temporal-Difference Learning for Policy Evaluation with Per-State Uncertainty Estimates. *Advances in Neural Information Processing Systems 32 (NeurIPS)*, 2019. (**Acceptance Rate: 21%**)
2. Sven Gowal, Dj Dvijotham, Robert Stanforth, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, **Timothy A. Mann**, and Pushmeet Kohli. Scalable Verified Training for Provably Robust Image Classification. *Proceedings of the International Conference on Computer Vision*, 2019. (**Acceptance Rate: 25%**)

3. **Timothy A. Mann**, Sven Gowal, András György, Ray Jiang, Huiyi Hu, Balaji Lakshminarayanan, and Prav Srinivasan. Learning from Delayed Outcomes via Proxies with Applications to Recommender Systems. *In Proceedings of the International Conference on Machine Learning (ICML)*, 2019. **(Acceptance Rate: 23%)**
4. Ray Jiang, Sven Gowal, Yuqiu Qian, **Timothy A. Mann**, Danilo Rezende. Beyond Greedy Ranking: Slate Optimization via List-CVAE. *In Proceedings of the International Conference on Representation and Learning (ICLR)*, 2019. **(Acceptance Rate: 31%)**
5. Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy R. Bunel, **Timothy A. Mann**, and Pushmeet Kohli. On the Effectiveness of Interval Bound Propagation for Verifiably Robust Models. *NIPS 2018 Workshop on Security in Machine Learning*, 2018. **(Acceptance Rate: 28%, Best Paper Award)**
6. Krishnamurthy Dvijotham, Robert Stanforth, Sven Gowal, **Timothy A. Mann**, and Pushmeet Kohli. A Dual Approach to Scalable Verification of Deep Networks. *Conference on Uncertainty in Artificial Intelligence*, 2018. **(Acceptance Rate: 31%, UAI Best Paper Award)**
7. Esther Derman, Daniel J. Mankowitz, **Timothy A. Mann**, and Shie Mannor. Soft-Robust Actor-Critic Policy-Gradient. *Conference on Uncertainty in Artificial Intelligence*, 2018. **(Acceptance Rate: 30.86%)**
8. Daniel J. Mankowitz, **Timothy A. Mann**, Pierre-Luc Bacon, Doina Precup, and Shie Mannor. Learning Robust Options. *Proceedings of the Association for the Advancement of Artificial Intelligence*, 2018.
9. Ray Jiang, Sven Gowal, **Timothy A. Mann**, Danilo Rezende. Optimizing Slate Recommendations via Slate-CVAE. arXiv:1803.01682, 2018.
10. Hugo Penedones, Damien Vincent, **Timothy A. Mann**, Sylvain Gelly. Temporal Difference Learning with Neural Networks-Study of the Leakage Propagation Problem. arXiv:1807.03064, 2018.
11. **Timothy A. Mann**, Hugo Penedones, Shie Mannor, and Todd Hester. Adaptive Lambda Least-Squares Temporal Difference Learning. arXiv:1612.09465, 2016.
12. Daniel J. Mankowitz, **Timothy A. Mann**, and Shie Mannor. Adaptive Skills, Adaptive Partitions (ASAP). *Advances in Neural Information Processing Systems 29*, 2016. **(Acceptance Rate: 23%)**
13. Assaf Hallak, Francois Schnitzler, **Timothy A. Mann**, and Shie Mannor. Off-policy evaluation for MDPs with unknown structure. *Proceedings of the 32nd International Conference on Machine Learning*, 2015. **(Acceptance Rate: 26.04%)**
14. Daniel J. Mankowitz, **Timothy A. Mann**, and Shie Mannor. Bootstrapping Skills. *Second Multidisciplinary Conference on Reinforcement Learning and Decision Making*, 2015. **(Selected for Oral Presentation: < 15%)**
15. Nir Levin, **Timothy A. Mann**, and Shie Mannor. Learning to Attract Followers on Twitter. *Second Multidisciplinary Conference on Reinforcement Learning and Decision Making*, 2015.
16. **Timothy A. Mann**, Daniel J. Mankowitz, and Shie Mannor. Learning when to Switch between Skills in a High Dimensional Domain. *AAAI-2015 Workshop on Learning for General Competency in Video Games*, 2015.
17. **Timothy A. Mann**, Doina Precup, and Shie Mannor. Fast MDP planning with landmarks. *NIPS 2014 Workshop on Large-scale reinforcement learning and Markov decision problems*, 2014.
18. Odalric-Ambrym Maillard, **Timothy A. Mann**, and Shie Mannor. “How Hard is My MDP?” The distribution norm to the rescue. *Advances in Neural Information Processing Systems 27 (NIPS)*, 2014. **(Acceptance Rate: 25%, Selected for Full Oral Presentation: 1%)**
19. Daniel J. Mankowitz, **Timothy A. Mann**, and Shie Mannor. Time-Regularized Interrupting Options (TRIO). *Proceedings of the 31st International Conference on Machine Learning (ICML)*, 2014. **(Acceptance Rate: 25%)**
20. **Timothy A. Mann** and Shie Mannor. Scaling Up Approximate Value Iteration with Options: Better Policies with Fewer Iterations. *Proceedings of the 31st International Conference on Machine Learning (ICML)*, 2014. **(Acceptance Rate: 25%)**
21. **Timothy A. Mann** and Shie Mannor. The Advantage of Planning with Options. *Reinforcement Learning and Decision Making (RLDM)*, 2013. **(Best Contributed Paper)**
22. **Timothy A. Mann** and Shie Mannor. Theoretical Analysis of Planning with Options. *European Workshop on Reinforcement Learning*, Dagstuhl, Germany, 2013.

23. **Timothy A. Mann** and Yoonsuck Choe. Directed Exploration in Reinforcement Learning with Transferred Knowledge. *JMLR: Workshop and Conference Proceedings*, 24:59–76, 2013.
24. **Timothy A. Mann** and Yoonsuck Choe. Directed Exploration in Reinforcement Learning with Transferred Knowledge. *European Workshop on Reinforcement Learning*, Edinburgh, Scotland, 2012. (**Acceptance Rate: 68%**)
25. Yoonsuck Choe and **Timothy A. Mann**. From problem solving to problem posing. *Brain Mind Magazine*, 1:7–8, 2012.
26. Sungmoon Jeong, Yunjung Park, Minhoo Lee, **Timothy A. Mann**, and Yoonsuck Choe. Proactive learning mechanism of sensory perception and behavior generation for an autonomous robot. In *Proceedings of the International Conference on Brain-Mind*, 2012.
27. **Timothy A. Mann**, Yunjung Park, Sungmoon Jeong, Minhoo Lee, and Yoonsuck Choe. Autonomously improving binocular depth estimation. In *Proceedings of the Japanese Neural Networks Society*, 2011.
28. **Timothy A. Mann** and Yoonsuck Choe. Scaling up reinforcement learning through targeted exploration. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011. (**Acceptance Rate: 25%**)
29. **Timothy A. Mann**, and Yoonsuck Choe. Grounding the meaning of nonprototypical smiles on motor behavior. *Behavioral and Brain Sciences*, 2010. Commentary on Niedenthal et al. (same volume).
30. **Timothy A. Mann** and Yoonsuck Choe. Prenatal to postnatal transfer of motor skills through motor-compatible sensory representations. In *Proceedings of the Ninth International Conference on Development & Learning*, 2010. (**Acceptance Rate: 67%**)
31. **Timothy A. Mann** and Yoonsuck Choe. Neural Conduction Delay Forces the Emergence of Predictive Function in Simulated Evolution. *BMC Neuroscience*, 11(Suppl 1):P62, 2010.
32. Maarten van Zomeren, Joshua Peschel, **Timothy A. Mann**, Gabe Knezek, James Doebbler, Jeremy Davis, Tracy A. Hammond, Augustinus H. J. Oomes, and Robin R. Murphy. Human-Robot Interaction Observations from a Proto-Study Using UAVs for Structural Inspection. In *HRI '09: Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 235-236, La Jolla, CA, March 2009.
33. Bum Soon Jang, **Timothy A. Mann**, and Yoonsuck Choe. Effects of varying the delay distribution in random, scale-free, and small-world networks. In *Proceedings of the 2008 IEEE International Conference on Granular Computing*, 2008.

Posters

- Timothy A. Mann** and Yoonsuck Choe. Minimizing On-Line Learning by Exploiting Stationary Properties in Nonstationary Environments. *ICML*, Edinburgh, Scotland, July 2012.
- Timothy A. Mann** and Yoonsuck Choe. Learning Visual Distance by Exploring Rewarding Tactile Stimuli. *TBSI*, Texas A&M University, September 2010.
- Timothy A. Mann** and Yoonsuck Choe. Extending Prenatal to Postnatal Motor Skills. *ARMADILLO 2009*, Rice University, October 2009.

Invited Talks

- Time-Regularized Interrupting Options @ Google DeepMind 2014
- An Overview of Reinforcement Learning @ SUNY Potsdam 2013
- The Advantage of Planning with Options
- University of Massachusetts at Amherst 2013
 - Stony Brook University 2013
 - Vassar College 2013
 - Multidisciplinary Conference on Reinforcement Learning and Decision Making 2013
- Convergence Rates of Planning with Macro-Actions in Markov Decision Processes @ The Technion - Israel Institute of Technology 2013

Finite-Sample Reinforcement Learning in Nonstationary Environments @ University of Texas at Austin 2011

Scaling Up Reinforcement Learning through Targeted Exploration

- Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea 2011
- Postech, Pohang, South Korea 2011
- Kyungpook National University, Dague, South Korea 2011

Human Undecidability @ European Conference on Computing and Philosophy 2007

Altruism in Spinoza's Ethics @ SUNY Oneonta Undergraduate Philosophy Conference 2007

Workshops

Abstraction in Reinforcement Learning @ ICML 2016. Co-organized with Daniel J. Mankowitz and Shie Mannor.

From Bad Models to Good Policies @ NIPS 2014. Co-organized with Odalric-Ambrym Maillard, Shie Mannor, and Laurent Orseau.

Honors and Awards

Attended the Heidelberg Laureate Forum, 2014

Best Contributed Paper Award @ RLDM 2013 (Title: The Advantage of Planning with Options)

ICML Travel Scholarship, 2012

Texas A&M University Dissertation Fellowship, 2012

Member of Upsilon Pi Epsilon, 2011

AAAI Doctoral Consortium, 2011 (Invited but could not attend)

NSF East Asia and Pacific Summer Institutes Fellowship, 2011

ICDL Travel Scholarship, 2010

CNS Travel Scholarship, 2010

Honorable Mention, NSF Graduate Research Fellowship, 2009

Benjamin F. Bradley Award, SUNY Potsdam, 2007

Member of Pi Mu Epsilon, 2007

Haden Land Scholarship, SUNY Potsdam, 2006.

MANNY (Multicultural Association of Northern New York) Leadership Award, 2006

NSF John von Neumann Scholarship (Advanced), SUNY Potsdam, 2006

NSF John von Neumann Scholarship (Beginning), SUNY Potsdam, 2005

Service

Vice President, Computer Science Graduate Student Association, Texas A&M University, Fall 2009 – Spring 2010

Graduate Mentor, NSF Research Experience for Undergraduates, Texas A&M University, Summer 2008 & Summer 2009

Vice President, Philosophy Forum, SUNY Potsdam, 2006 – 2007

CSSTEP (Collegiate Science and Technology Entry Program) Friend, SUNY Potsdam, 2006 – 2007

President, ACM Student Chapter, SUNY Potsdam, 2004 – 2006

Peer Review

Conference on Learning Theory 2014
International Conference on Machine Learning 2014
Neural Networks 2013
International Conference on Development and Learning 2012
IEEE Symposium Series on Computational Intelligence 2011
International Joint Conference on Neural Networks 2010
International Conference on Neural Information Processing 2009
Joint Conference on Digital Libraries 2009
Neural Networks 2008
IEEE World Congress on Computational Intelligence 2008

Professional Societies

Association for Computing Machinery (ACM) Member, 2009 - Current

- SIGART 2009 - Current
- SIGEVO 2009 - 2013

IEEE Member, 2009 - Current

Teaching Experience

Teaching Assistant, Texas A&M University,

- CSCE 315: Programming Studio, Fall 2011
- CSCE 315: Programming Studio, Fall 2010 – Spring 2011
- CSCE 606: Software Engineering, Summer 2010
- CSCE 315: Programming Studio, Fall 2009 – Spring 2010
- CPSC 606: Software Engineering, Summer 2009
- CPSC 315: Programming Studio, Fall 2008 – Spring 2009
- CPSC 601: Programming with C and Java, Summer 2008
- CPSC 111: Intro. to Computer Science Concepts and Programming, Fall 2007 – Spring 2008

Teaching Assistant, CIS 206: Data Structures, SUNY Potsdam, Fall 2004

Technical Skills

C/C++, Java, Python, TensorFlow, Go, Lua, Matlab, GNU Octave, Linux, L^AT_EX, Bash, Perl, PHP, Javascript, SQL, Scheme

References

Dr. Shie Mannor, Professor, Electrical Engineering at the Technion, Haifa, Israel.
Dr. Yoonsuck Choe, Associate Professor, Computer Science & Engineering at Texas A&M University.
Dr. Robin Murphy, Raytheon Professor, Computer Science & Engineering at Texas A&M University.
Dr. Peter Stone, Associate Professor, Computer Science at University of Texas at Austin.
Dr. Doina Precup, Associate Professor, School of Computer Science at McGill University, Montreal, Canada.