Matthew E. Taylor

Education

2003–2008 **Ph.D.** Department of Computer Science, *University of Texas at Austin*
Dissertation title: *Autonomous Inter-Task Transfer in Reinforcement Learning Domains*

1997–2001 **A.B.** *Amherst College*
Majors: Computer Science and Physics (with distinction); *Magna Cum Laude*

Professional Appointments

2020–current **Associate Professor**, University of Alberta, Department of Computing Science.

2018–2020 **Adjunct Professor**, University of Alberta, Department of Computing Science.

2017–2020 **Principal Researcher**, Borealis AI.

2018–current **Adjunct Professor**, Washington State University, School of Electrical Engineering and Computer Science.

2014–current **Affiliated Assistant Professor**, Center for Precision and Automated Agricultural Systems (courtesy appointment) *Washington State University*, College of Agricultural, Human, and Natural Resource Sciences.

2013–2018 **Assistant Professor**, Allred Distinguished Professorship in Artificial Intelligence, *Washington State University*, School of Electrical Engineering and Computer Science.
Director of the Intelligent Robot Learning Laboratory: `http://irll.eecs.wsu.edu`
Approved for tenured as of summer 2018, but I did not return to WSU

2010–2012 **Assistant Professor**, *Lafayette College*, Computer Science Department

2008–2010 **Postdoctoral Research Assistant**, *University of Southern California*, Department of Computer Science, with Milind Tambe

2003–2008 **Graduate Research Assistant**, *University of Texas at Austin*, Department of Computer Science, with Peter Stone

2001–2003 **Lead Software Developer**, *Epic Systems Corp.*, Madison, WI

Summer Internships

2006 **Research Intern**, *Cycorp*, Austin, TX, with Michael Witbrock

2000 **Software Development Intern**, *Microsoft*, Redmond, WA

1999 **NSF Research Experience for Undergraduates** (physics), *University of Nebraska Lincoln*, Lincoln, NE, with Robert Hilborn

Honors

2018 Early Career Spotlight talk at IJCAI-18

2018 Awarded AAAI Senior Member status

2015 Google Faculty Research Award

2015 WSU EECS Early Career Award

2014 Award for Academic Advisor Excellence, WSU Graduate & Professional Students Association

2014 Elected Board Member of the International Foundation of Autonomous Agents and Multi-Agent Systems (IFAAMAS), 6 year term

✉ *metaylor@gmail.com* ▪ *Last Updated: April, 2020*

2012 NSF CAREER Award
2012 Upsilon Pi Epsilon honor society
2006 Best Paper Award, GECCO-06, Genetic Algorithms Track
2003–2004 MCD Fellowship, UT-Austin's Department of Computer Sciences
2003 Dean's Excellence Award, UT-Austin's College of Natural Sciences
2001 Sigma Xi honor society
1998 NSF/STEMTEC Teaching Fellowship

Publications

Book

[1] Matthew E. Taylor. *Transfer in Reinforcement Learning Domains*, volume 216 of *Studies in Computational Intelligence*. Springer-Verlag, 2009. ISBN 978-3-642-01881-7.

Edited Volume

[2] Matthew E. Taylor and Karl Tuyls, editors. *Adaptive Agents and Multi-Agent Systems IV*, volume 5924 of *Lecture Notes in Computer Science*. Springer-Verlag, 2010. ISBN 978-3-642-11813-5.

Book Chapters: Refereed

[3] Haitham Bou Ammar, Matthew E. Taylor, Karl Tuyls, and Gerhard Weiss. Reinforcement Learning Transfer using a Sparse Coded Inter-Task Mapping. In *LNAI Post-proceedings of the European Workshop on Multi-agent Systems*. Springer-Verlag, 2013.

[4] Anestis Fachantidis, Ioannis Partalas, Matthew E. Taylor, and Ioannis Vlahavas. Transfer Learning via Multiple Inter-Task Mappings. In Scott Sanner and Marcus Hutter, editors, *Recent Advances in Reinforcement Learning*, volume 7188 of *Lecture Notes in Artificial Intelligence*, pages 225–236. Springer-Verlag, Berlin, 2012. ISBN 978-3-642-29945-2.

[5] Peter Stone, Gregory Kuhlmann, Matthew E. Taylor, and Yaxin Liu. Keepaway Soccer: From Machine Learning Testbed to Benchmark. In Itsuki Noda, Adam Jacoff, Ansgar Bredenfeld, and Yasutake Takahashi, editors, *RoboCup-2005: Robot Soccer World Cup IX*, volume 4020, pages 93–105. Springer-Verlag, Berlin, 2006. 28% acceptance rate at RoboCup-2005.

Book Chapters: Invited

[6] Matthew E. Taylor, Christopher Kiekintveld, and Milind Tambe. Evaluating Deployed Decision Support Systems for Security: Challenges, Arguments, and Approaches. In Milind Tambe, editor, *Security Games: Theory, Deployed Applications, Lessons Learned*, pages 254–283. Cambridge University Press, 2011. ISBN 978-1-107-09642-4.

[7] Matthew E. Taylor, Manish Jain, Christopher Kiekintveld, Jun young Kwak, Rong Yang, Zhengyu Yin, and Milind Tambe. Two decades of multiagent teamwork research: Past, present, and future. In C. Guttmann, F. Dignum, and M. Georgeff, editors, *Collaborative Agents - REsearch and Development (CARE) 2009-2010*, volume 6066 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2011.

[8] Marc Ponsen, Matthew E. Taylor, and Karl Tuyls. Abstraction and Generalization in Reinforcement Learning. In Matthew E. Taylor and Karl Tuyls, editors, *Adaptive Agents and Multi-Agent Systems IV*, volume 5924, pages 1–33. Springer-Verlag, 2010.

Journal Articles

[9] Yang Hu, Rachel Min Wong, Olusola Adesope, and Matthew E. Taylor. Effects of a computer-based learning environment that teaches older adults how to install a smart home system. *Computers and Education*, 149, 2020. https://doi.org/10.1016/j.compedu.2020.103816.

[10] Behzad Ghazanfari, Fatemeh Afghah, and Matthew E. Taylor. Sequential association rule mining for autonomously extracting hierarchical task structures in reinforcement learning. *IEEE Access*, January 2020.

[11] Bikramjit Banerjee, Syamala Vittanala, and Matthew E. Taylor. Team learning from human demonstration with coordination confidence. *The Knowledge Engineering Review*, 34, 2019.

[12] Sarah Morton, Julie Kmec, and Matthew E. Taylor. It's what you call it: Gendered framing and women's and men's interest in a robotics instruction task. *International Journal of Gender, Science and Technology*, 11(2), 2019. Accepted.

[13] Pablo Hernandez-Leal, Bilal Kartal, and Matthew E. Taylor. A survey and critique of multiagent deep reinforcement learning. *Journal of Autonomous Agents and Multiagent Systems*, 33:750–797, October 2019.

[14] Gabriel V. de la Cruz Jr., Yunshu Du, and Matthew E. Taylor. Pre-training with non-expert human demonstration for deep reinforcement learning. *The Knowledge Engineering Review*, 34, 2019.

[15] Santosh Bhusal, Kapil Khanal, Shivam Goel, Manoj Karkee, and Matthew E. Taylor. Bird deterrence in a vineyard using an unmanned aerial system (UAS). *Transactions of the ASABE*, 62(2):561–569, 2019.

[16] Garrett Wilson, Christopher Pereyda, Nisha Raghunath, Gabriel de la Cruz Jr., Shivam Goel, Sepehr Nesaei, Bryan Minor, Maureen Schmitter-Edgecombe, Matthew E. Taylor, and Diane J. Cook. Robot-enabled support of daily activities in smart home environments. *Cognitive Systems Research*, 54:258–272, May 2019.

[17] Yunshu Du, Assefaw Gebremedhin, and Matthew E. Taylor. Analysis of university fitness center data uncovers interesting patterns, enables prediction. *IEEE Transactions on Knowledge and Data Engineering*, 31:1478–1490, 2019.

[18] Bei Peng, James MacGlashan, Robert Loftin, Michael L. Littman, David L. Roberts, and Matthew E. Taylor. Curriculum design for machine learners in sequential decision tasks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2:268–277, 2018.

[19] Ariel Rosenfeld, Moshe Cohen, Matthew E. Taylor, and Sarit Kraus. Leveraging human knowledge in tabular reinforcement learning: a study of human subjects. *The Knowledge Engineering Review*, 33, 2018.

[20] Yunxiang Ye, Zhaodong Wang, Dylan Jones, Long He, Matthew E. Taylor, Geoffrey A. Hollinger, and Qin Zhang. Bin-dog: A robotic platform for bin management in orchards. *Robotics*, 6(2), 2017. ISSN 2218-6581. URL `http://www.mdpi.com/2218-6581/6/2/12`.

[21] Yusen Zhan, Haitham Bou Ammar, and Matthew E. Taylor. Non-convex policy search using variational inequalities. *Neural Computation*, 2017.

[22] Tim Brys, Anna Harutyunyan, Peter Vranck, Matthew E. Taylor, and Ann Nowe. Multi-objectivization and ensembles of shapings in reinforcement learning. *Neurocomputing*, 2017.

[23] Yang Hu and Matthew E. Taylor. A computer-aided design intelligent tutoring system teaching strategic flexibility. *Transactions on Techniques for STEM Education*, October–December 2016.

[24] Chris Cain, Anne Anderson, and Matthew E. Taylor. Content-independent classroom gamification. *Computers in Education Journal*, October–December 2016.

[25] Pablo Hernandez-Leal, Yusen Zhan, Matthew E. Taylor, L. Enrique Sucar, and Enrique Munoz de Cote. Efficiently detecting switches against non-stationary opponents. *Autonomous Agents and Multi-Agent Systems*, 2016. ISSN 1387-2532. URL `http://dx.doi.org/10.1007/s10458-016-9352-6`.

[26] Pablo Hernandez-Leal, Yusen Zhan, Matthew E. Taylor, L. Enrique Sucar, and Enrique Munoz de Cote. An exploration strategy for non-stationary opponents. *Autonomous Agents and Multi-Agent Systems*, pages 1–32, 2016. ISSN 1573-7454. URL `http://dx.doi.org/10.1007/s10458-016-9347-3`.

[27] Anestis Fachantidis, Ioannis Partalas, Matthew E. Taylor, and Ioannis Vlahavas. Transfer learning with probabilistic mapping selection. *Adaptive Behavior*, 23(1):3–19, 2015. URL http://adb.sagepub.com/content/23/1/3.abstract.

[28] Robert Loftin, Bei Peng, James MacGlashan, Michael L. Littman, Matthew E. Taylor, Jeff Huang, and David L. Roberts. Learning behaviors via human-delivered discrete feedback: modeling implicit feedback strategies to speed up learning. *Journal of Autonomous Agents and Multi-Agent Systems*, pages 1–30, 2015. URL http://link.springer.com/article/10.1007\%2Fs10458-015-9283-7.

[29] Matthew E. Taylor, Nicholas Carboni, Anestis Fachantidis, Ioannis Vlahavas, and Lisa Torrey. Reinforcement learning agents providing advice in complex video games. *Connection Science*, 26(1):45–63, 2014. URL http://dx.doi.org/10.1080/09540091.2014.885279.

[30] Tim Brys, Tong T. Pham, and Matthew E. Taylor. Distributed learning and multi-objectivity in traffic light control. *Connection Science*, 26(1):65–83, 2014. URL http://dx.doi.org/10.1080/09540091.2014.885282.

[31] Marcos A. M. Vieira, Matthew E. Taylor, Prateek Tandon, Manish Jain, Ramesh Govindan, Gaurav S. Sukhatme, and Milind Tambe. Mitigating Multi-path Fading in a Mobile Mesh Network. *Ad Hoc Networks Journal*, 2011.

[32] Matthew E. Taylor, Manish Jain, Prateek Tandon, Makoto Yokoo, and Milind Tambe. Distributed On-line Multi-Agent Optimization Under Uncertainty: Balancing Exploration and Exploitation. *Advances in Complex Systems*, 2011.

[33] Matthew E. Taylor and Peter Stone. An Introduction to Inter-task Transfer for Reinforcement Learning. *AI Magazine*, 32(1):15–34, 2011.

[34] Matthew E. Taylor, Christopher Kiekintveld, Craig Western, and Milind Tambe. A Framework for Evaluating Deployed Security Systems: Is There a Chink in your ARMOR? *Informatica*, 34(2):129–139, 2010.

[35] Shimon Whiteson, Matthew E. Taylor, and Peter Stone. Critical Factors in the Empirical Performance of Temporal Difference and Evolutionary Methods for Reinforcement Learning. *Journal of Autonomous Agents and Multi-Agent Systems*, 21(1):1–27, 2010.

[36] Matthew E. Taylor and Peter Stone. Transfer Learning for Reinforcement Learning Domains: A Survey. *Journal of Machine Learning Research*, 10(1):1633–1685, 2009.

[37] Matthew E. Taylor, Peter Stone, and Yaxin Liu. Transfer Learning via Inter-Task Mappings for Temporal Difference Learning. *Journal of Machine Learning Research*, 8(1):2125–2167, 2007.

[38] Shimon Whiteson, Matthew E. Taylor, and Peter Stone. Empirical Studies in Action Selection for Reinforcement Learning. *Adaptive Behavior*, 15(1), 2007.

Conference Papers

[39] Pablo Hernandez-Leal, Bilal Kartal, and Matthew E. Taylor. A Very Condensed Survey and Critique of Multiagent Deep Reinforcement Learning. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*. Auckland, New Zealand, May 2020.

[40] Sriram Ganapathi Subramanian, Pascal Poupart, Matthew E. Taylor, and Nidhi Hegde. Multi type mean field reinforcement learning. In *Proceedings of the 19th International Conference on Autonomous Agents and Multi-Agent Systems*, May 2020. 23% acceptance rate.

[41] Felipe Leno Da Silva, Pablo Hernandez-Leal, Bilal Kartal, and Matthew E. Taylor. Uncertainty-aware action advising for deep reinforcement learning agents. In *Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*, January 2020. 21% acceptance rate.

[42] Nathan Douglas, Dianna Yim, Bilal Kartal, Pablo Hernandez-Leal, Matthew E. Taylor, and Frank Maurer. Towers of saliency: A reinforcement learning visualization using immersive environments. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*, ISS '19. ACM, New York, NY, USA, November 2019. ISBN 978-1-4503-6891-9/19/11.

[43] Weixun Wang, Jianye Hao, Yixi Wang, and Matthew E. Taylor. Achieving cooperation through deep multiagent reinforcement learning in sequential prisoner's dilemmas. In *Proceedings of the 1st International Conference on Distributed Artificial Intelligence (DAI-19)*, October 2019. 35% acceptance rate.

[44] Pablo Hernandez-Leal, Bilal Kartal, and Matthew E. Taylor. Agent modeling as auxiliary task for deep reinforcement learning. In *Proceedings of the 15th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE'19)*, October 2019. 25% acceptance rate for oral presentation.

[45] Bilal Kartal, Pablo Hernandez-Leal, and Matthew E. Taylor. Terminal prediction as an auxiliary task for deep reinforcement learning. In *Proceedings of the 15th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE'19)*, October 2019. 25% acceptance rate for oral presentation.

[46] Chao Gao, Bilal Kartal, Pablo Hernandez-Leal, and Matthew E. Taylor. On hard exploration for reinforcement learning: a case study in pommerman. In *Proceedings of the 15th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE'19)*, October 2019. 25% acceptance rate for oral presentation.

[47] Bilal Kartal, Pablo Hernandez-Leal, and Matthew E. Taylor. Action guidance with mcts for deep reinforcement learning. In *Proceedings of the 15th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE'19)*, October 2019. Poster: 25% acceptance rate for oral presentation, additional 25% acceptance rate for posters.

[48] Zhaodong Wang and Matthew E. Taylor. Interactive reinforcement learning with dynamic reuse of prior knowledge from human and agent demonstrations. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019. 18% acceptance rate.

[49] Kenny Young, Baoxiang Wang, and Matthew E. Taylor. Metatrace: Online step-size tuning by meta-gradient descent for reinforcement learning control. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019. 18% acceptance rate.

[50] Nadia V. Panossian, Dustin McLarty, and Matthew E. Taylor. Artificial neural network for unit commitment on networks with significant energy storage. In *Proceedings of the IEEE Green Technologies Conference (GREENTECH)*, April 2019.

[51] Konstantin I. Matveev, John P. Swensen, and Matthew E. Taylor. Modeling of decelerated descent of an elongated body with vectored thrust. In *Proceedings of American Society of Mechanical Engineer's 5th Joint US-European Fluids Engineering Division Summer Meeting*, July 2018.

[52] Felipe Leno Da Silva, Matthew E. Taylor, and Anna Helena Reali Costa. Autonomously reusing knowledge in multiagent reinforcement learning. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 5487–5493. International Joint Conferences on Artificial Intelligence Organization, 7 2018. URL `https://doi.org/10.24963/ijcai.2018/774`.

[53] Matthew E. Taylor. Improving reinforcement learning with human input. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.

[54] Shivam Goel, Santosh Bhusal, Matthew E Taylor, and Manoj Karkee. Detection and localization of birds for Bird Deterrence using UAS. In *2017 ASABE Annual International Meeting*. American Society of Agricultural and Biological Engineers, 2017.

[55] Santosh Bhusal, Shivam Goel, Kapil Khanal, Matthew E. Taylor, and Manoj Karkee. Bird Detection, Tracking and Counting in Wine Grapes. In *2017 ASABE Annual International Meeting*. American Society of Agricultural and Biological Engineers, 2017.

[56] James MacGlashan, Mark Ho, Robert Loftin, Bei Peng, Guan Wang, David L. Roberts, Matthew E.Taylor, and Michael L. Littman. Interactive learning from policy-dependent human feedback. In *Proceedings of the International Conference on Machine Learning (ICML)*, August 2017. 25% acceptance rate.

[57] Ariel Rosenfeld, Matthew E. Taylor, and Sarit Kraus. Leveraging human knowledge in tabular reinforcement learning: A study of human subjects. In *Proceedings of the 26th International Conference on Artificial Intelligence (IJCAI)*, August 2017. 26% acceptance rate.

[58] Zhaodong Wang and Matthew E. Taylor. Improving reinforcement learning with confidence-based demonstrations. In *Proceedings of the 26th International Conference on Artificial Intelligence (IJCAI)*, August 2017. 26% acceptance rate.

[59] Salam El Bsat, Haitham Bou Ammar, and Matthew E. Taylor. Scalable multitask policy gradient reinforcement learning. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, February 2017. 25% acceptance rate.

[60] David Isele, José Marcio Luna, Eric Eaton, Gabriel V. de la Cruz Jr., James Irwin, Brandon Kallaher, and Matthew E. Taylor. Lifelong Learning for Disturbance Rejection on Mobile Robots. In *Proceedings of the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, October 2016. 48% acceptance rate.

[61] Chris Cain, Anne Anderson, and Matthew E. Taylor. Content-Independent Classroom Gamification. In *Proceedings of the ASEE's 123rd Annual Conference & Exposition*. New Orleans, LA, USA, June 2016.

[62] Yang Hu and Matthew E. Taylor. Work In Progress: A Computer-Aided Design Intelligent Tutoring System Teaching Strategic Flexibility. In *Proceedings of the ASEE's 123rd Annual Conference & Exposition*. New Orleans, LA, USA, June 2016.

[63] Yusen Zhan, Haitham Bou Ammar, and Matthew E. Taylor. Theoretically-Grounded Policy Advice from Multiple Teachers in Reinforcement Learning Settings with Applications to Negative Transfer. In *Proceedings of the 25th International Conference on Artificial Intelligence (IJCAI)*, July 2016. 25% acceptance rate.

[64] Halit Bener Suay, Tim Brys, Matthew E. Taylor, and Sonia Chernova. Learning from Demonstration for Shaping through Inverse Reinforcement Learning. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2016. 24.9% acceptance rate.

[65] Bei Peng, James MacGlashan, Robert Loftin, Michael L. Littman, David L. Roberts, and Matthew E. Taylor. A Need for Speed: Adapting Agent Action Speed to Improve Task Learning from Non-Expert Humans. In *Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2016. 24.9% acceptance rate.

[66] Tim Brys, Anna Harutyunyan, Halit Bener Suay, Sonia Chernova, Matthew E. Taylor, and Ann Nowé. Reinforcement Learning from Demonstration through Shaping. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2015. 28.8% acceptance rate.

[67] Tim Brys, Anna Harutyunyan, Matthew E. Taylor, and Ann Nowé. Policy Transfer using Reward Shaping. In *The 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2015. 25% acceptance rate.

[68] Haitham Bou Ammar, Eric Eaton, Paul Ruvolo, and Matthew E. Taylor. Unsupervised Cross-Domain Transfer in Policy Gradient Reinforcement Learning via Manifold Alignment. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, January 2015. 27% acceptance rate.

[69] Matthew E. Taylor and Lisa Torrey. Agents Teaching Agents in Reinforcement Learning (Nectar Abstract). In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECMLPKDD)*, September 2014. Nectar Track, 45% acceptance rate.

[70] Chris HolmesParker, Matthew E. Taylor, Adrian Agogino, and Kagan Tumer. CLEANing the Reward: Counterfactual Actions Remove Exploratory Action Noise in Multiagent Learning. In *Proceedings of the 2014 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT)*, August 2014. 43% acceptance rate.

[71] Robert Loftin, Bei Peng, James MacGlashan, Michael Littman, Matthew E. Taylor, David Roberts, and Jeff Huang. Learning Something from Nothing: Leveraging Implicit Human Feedback Strategies. In *Proceedings of the 23rd IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, August 2014.

[72] Tim Brys, Ann Nowé, Daniel Kudenko, and Matthew E. Taylor. Combining Multiple Correlated Reward and Shaping Signals by Measuring Confidence. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, July 2014. 28% acceptance rate.

[73] Robert Loftin, Bei Peng, James MacGlashan, Machiael L. Littman, Matthew E. Taylor, Jeff Huang, and David L. Roberts. A Strategy-Aware Technique for Learning Behaviors from Discrete Human Feedback. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, July 2014. 28% acceptance rate.

[74] Tim Brys, Anna Harutyunyan, Peter Vrancx, Matthew E. Taylor, Daniel Kudenko, and Ann Nowé. Multi-Objectivization of Reinforcement Learning Problems by Reward Shaping. In *Proceedings of the IEEE 2014 International Joint Conference on Neural Networks (IJCNN)*, July 2014. 59% acceptance rate.

[75] Haitham Bou Ammar, Eric Eaton, Paul Ruvolo, and Matthew E. Taylor. Online Multi-Task Learning for Policy Gradient Methods. In *Proceedings of the 31st International Conferences on Machine Learning (ICML)*, June 2014. 25% acceptance rate.

[76] Anestis Fachantidis, Ioannis Partalas, Matthew E. Taylor, and Ioannis Vlahavas. An Autonomous Transfer Learning Algorithm for TD-Learners. In *Proceedings of the 8th Hellenic Conference on Artificial Intelligence (SETN)*, May 2014. 50% acceptance rate.

[77] Haitham Bou Ammar, Decebal Constantin Mocanu, Matthew E. Taylor, Kurt Driessens, Karl Tuyls, and Gerhard Weiss. Automatically Mapped Transfer Between Reinforcement Learning Tasks via Three-Way Restricted Boltzmann Machines. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, September 2013. 25% acceptance rate.

[78] Tong Pham, Aly Tawfika, and Matthew E. Taylor. A Simple, Naive Agent-based Model for the Optimization of a System of Traffic Lights: Insights from an Exploratory Experiment. In *Proceedings of Conference on Agent-Based Modeling in Transportation Planning and Operations*, September 2013.

[79] Lisa Torrey and Matthew E. Taylor. Teaching on a Budget: Agents Advising Agents in Reinforcement Learning. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2013. 23% acceptance rate.

[80] Haitham Bou Ammar, Karl Tuyls, Matthew E. Taylor, Kurt Driessen, and Gerhard Weiss. Reinforcement Learning Transfer via Sparse Coding. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, June 2012. 20% acceptance rate.

[81] Matthew E. Taylor, Halit Bener Suay, and Sonia Chernova. Integrating Reinforcement Learning with Human Demonstrations of Varying Ability. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2011. 22% acceptance rate.

[82] Matthew E. Taylor, Brian Kulis, and Fei Sha. Metric Learning for Reinforcement Learning Agents. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2011. 22% acceptance rate.

[83] Jason Tsai, Natalie Fridman, Emma Bowring, Matthew Brown, Shira Epstein, Gal Kaminka, Stacy Marsella, Andrew Ogden, Inbal Rika, Ankur Sheel, Matthew E. Taylor, Xuezhi Wang, Avishay Zilka, and Milind Tambe. ESCAPES: Evacuation Simulation with Children, Authorities, Parents, Emotions, and Social Comparison. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2011. 22% acceptance rate.

[84] Matthew E. Taylor, Katherine E. Coons, Behnam Robatmili, Bertrand A. Maher, Doug Burger, and Kathryn S. McKinley. Evolving Compiler Heuristics to Manage Communication and Contention. In *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence (AAAI)*, July 2010. Nectar Track, 25% acceptance rate.

[85] Matthew E. Taylor, Manish Jain, Yanquin Jin, Makoto Yooko, and Milind Tambe. When Should There be a "Me" in "Team"? Distributed Multi-Agent Optimization Under Uncertainty. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2010. 24% acceptance rate.

[86] Pradeep Varakantham, Jun young Kwak, Matthew E. Taylor, Janusz Marecki, Paul Scerri, and Milind Tambe. Exploiting Coordination Locales in Distributed POMDPs via Social Model Shaping. In *Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling (ICAPS)*, September 2009. 34% acceptance rate.

[87] Manish Jain, Matthew E. Taylor, Makoto Yokoo, and Milind Tambe. DCOPs Meet the Real World: Exploring Unknown Reward Matrices with Applications to Mobile Sensor Networks. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, July 2009. 26% acceptance rate.

[88] Matthew E. Taylor, Nicholas K. Jong, and Peter Stone. Transferring Instances for Model-Based Reinforcement Learning. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, pages 488–505, September 2008. 19% acceptance rate.

[89] Katherine K. Coons, Behnam Robatmili, Matthew E. Taylor, Bertrand A. Maher, Kathryn McKinley, and Doug Burger. Feature Selection and Policy Optimization for Distributed Instruction Placement Using Reinforcement Learning. In *Proceedings of the Seventh International Joint Conference on Parallel Architectures and Compilation Techniques (PACT)*, pages 32–42, October 2008. 19% acceptance rate.

[90] Matthew E. Taylor, Gregory Kuhlmann, and Peter Stone. Autonomous Transfer for Reinforcement Learning. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 283–290, May 2008. 22% acceptance rate.

[91] Matthew E. Taylor, Gregory Kuhlmann, and Peter Stone. Transfer Learning and Intelligence: an Argument and Approach. In *Proceedings of the First Conference on Artificial General Intelligence (AGI)*, March 2008. 50% acceptance rate.

[92] Matthew E. Taylor and Peter Stone. Cross-Domain Transfer for Reinforcement Learning. In *Proceedings of the Twenty-Fourth International Conference on Machine Learning (ICML)*, June 2007. 29% acceptance rate.

[93] Matthew E. Taylor, Shimon Whiteson, and Peter Stone. Temporal Difference and Policy Search Methods for Reinforcement Learning: An Empirical Comparison. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI)*, pages 1675–1678, July 2007. Nectar Track, 38% acceptance rate.

[94] Matthew E. Taylor, Shimon Whiteson, and Peter Stone. Transfer via Inter-Task Mappings in Policy Search Reinforcement Learning. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 156–163, May 2007. 22% acceptance rate.

[95] Mazda Ahmadi, Matthew E. Taylor, and Peter Stone. IFSA: Incremental Feature-Set Augmentation for Reinforcement Learning Tasks. In *Proceedings of the the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1120–1127, May 2007. 22% acceptance rate, Finalist for Best Student Paper.

[96] Matthew E. Taylor, Cynthia Matuszek, Pace Reagan Smith, and Michael Witbrock. Guiding Inference with Policy Search Reinforcement Learning. In *Proceedings of the Twentieth International FLAIRS Conference (FLAIRS)*, May 2007. 52% acceptance rate.

[97] Matthew E. Taylor, Cynthia Matuszek, Bryan Klimt, and Michael Witbrock. Autonomous Classification of Knowledge into an Ontology. In *Proceedings of the Twentieth International FLAIRS Conference (FLAIRS)*, May 2007. 52% acceptance rate.

[98] Matthew E. Taylor, Shimon Whiteson, and Peter Stone. Comparing Evolutionary and Temporal Difference Methods for Reinforcement Learning. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pages 1321–28, July 2006. 46% acceptance rate, Best Paper Award in GA track (of 85 submissions).

[99] Matthew E. Taylor, Peter Stone, and Yaxin Liu. Value Functions for RL-Based Behavior Transfer: A Comparative Study. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI)*, July 2005. 18% acceptance rate.

[100] Matthew E. Taylor and Peter Stone. Behavior Transfer for Value-Function-Based Reinforcement Learning. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 53–59, July 2005. 25% acceptance rate.

Short Conference Papers

[101] Nathan Douglas, Dianna Yim, Bilal Kartal, Pablo Hernandez-Leal, Matthew E. Taylor, and Frank Maurer. Towers of saliency: A reinforcement learning visualization using immersive environments. In *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*, ISS '19. ACM, New York, NY, USA, 2019. ISBN 978-1-4503-6891-9/19/11. URL `http://doi.acm.org/10.1145/3343055.3360747`.

[102] Chao Gao, Pablo Hernandez-Leal, Bilal Kartal, and Matthew E. Taylor. Skynet: A top deep rl agent in the inaugural pommerman team competition. In *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2019.

[103] Pablo Hernandez-Leal, Bilal Kartal, and Matthew E. Taylor. Opponent modeling with actor-critic methods in deep reinforcement learning. In *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2019.

[104] Bilal Kartal, Pablo Hernandez-Leal, and Matthew E. Taylor. Predicting when to expect terminal states improves deep rl. In *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2019.

[105] Sriram Ganapathi Subramanian, Pascal Poupart, Matthew E. Taylor, and Nidhi Hegde. Multi type mean field reinforcement learning. In *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2019.

[106] Bei Peng, James MacGlashan, Robert Loftin, Michael L. Littman, David L. Roberts, and Matthew E. Taylor. Curriculum Design for Machine Learners in Sequential Decision Tasks (Extended Abstract). In *Proceedings of the 2017 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2017. Extended abstract: 26% acceptance rate for papers, additional 22% for extended abstracts.

[107] Pablo Hernandez-Leal, Benajamin Rosman, Matthew E. Taylor, L. Enrique Sucar, and Enrique Munoz de Cote. A Bayesian Approach for Learning and Tracking Switching, Non-stationary Opponents (Extended Abstract). In *Proceedings of 15th International Conference on Autonomous Agents and Multiagent Systems*. Singapore, May 2016.

[108] Tim Brys, Anna Harutyunyan, Matthew E. Taylor, and Ann Nowé. Ensembles of Shapings. In *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2015. 15% acceptance rate for oral presentations.

[109] Halit Bener Suay, Tim Brys, Matthew E. Taylor, and Sonia Chernova. Reward Shaping by Demonstration. In *The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)*, 2015.

[110] Pablo Hernandez-Leal, Matthew E. Taylor, Enrique Munoz de Cote, and L. Enrique Sucar. Bidding in Non-Stationary Energy Markets. In *The 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2015. Extended Abstract: 25% acceptance rate for papers, additional 22% for extended abstracts.

[111] Gabriel V. de la Cruz Jr., Bei Peng, Walter S. Lasecki, and Matthew E. Taylor. Towards Integrating Real-Time Crowd Advice with Reinforcement Learning. In *The 20th ACM Conference on Intelligent User Interfaces (IUI)*, March 2015. Poster: 41% acceptance rate for poster submissions.

[112] Tim Brys, Matthew E. Taylor, and Ann Nowé. Using Ensemble Techniques and Multi-Objectivization to Solve Reinforcement Learning Problems. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, August 2014. 41% acceptance rate for short papers.

[113] Tim Brys, Kristof Van Moffaert, Ann Nowe, and Matthew E. Taylor. Adaptive Objective Selection for Correlated Objectives in Multi-Objective Reinforcement Learning (Extended Abstract). In *The 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2014. Extended abstract: 24% acceptance rate for papers, additional 22% for extended abstracts.

[114] Chris HolmesParker, Matthew E. Taylor, Adrian Agogino, and Kagan Tumer. CLEANing the Reward: Counterfactual Actions Remove Exploratory Action Noise in Multiagent Learning (Extended Abstract). In *The Thirteenth International Joint Conference on Autonomous Agents and Multiagent Systems*, May 2014. Extended abstract: 24% acceptance rate for papers, additional 22% for extended abstracts.

[115] Haitham Bou Ammar, Decebal Constantin Mocanu, Matthew E. Taylor, Kurt Driessens, Karl Tuyls, and Gerhard Weiss. Automatically Mapped Transfer Between Reinforcement Learning Tasks via Three-Way Restricted Boltzmann Machines. In *The 25th Benelux Conference on Artificial Intelligence (BNAIC)*, November 2013.

[116] Lisa Torrey and Matthew E. Taylor. Towards Student/Teacher Learning in Sequential Decision Tasks. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, June 2012. Extended Abstract: 20% acceptance rate for papers, additional 23% for extended abstracts.

[117] Haitham Bou Ammar, Matthew E. Taylor, and Karl Tuyls. Common Sub-Space Transfer for Reinforcement Learning Tasks (Poster). In *The 23rd Benelux Conference on Artificial Intelligence (BNAIC)*, November 2011. 44% overall acceptance rate.

[118] Jun young Kwak, Rong Yang, Zhengyu Yin, Matthew E. Taylor, and Milind Tambe. Towards Addressing Model Uncertainty: Robust Execution-time Coordination for Teamwork (Short Paper). In *The IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT)*, August 2011. Short Paper: 21% acceptance rate for papers, additional 28% for short papers.

[119] Jun young Kwak, Rong Yang, Zhengyu Yin, Matthew E. Taylor, and Milind Tambe. Teamwork in Distributed POMDPs: Execution-time Coordination Under Model Uncertainty (Poster). In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2011. Extended Abstract: 22% acceptance rate for papers, additional 25% for extended abstracts.

[120] Matthew E. Taylor and Peter Stone. Towards Reinforcement Learning Representation Transfer (Poster). In *The Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 683–685, May 2007. Poster: 22% acceptance rate for talks, additional 25% for posters.

Refereed Workshop and Symposium Papers

[121] Felipe Leno da Silva, Pablo Hernandez-Leal, Bilal Kartal, and Matthew E. Taylor. Receiving uncertainty-aware advice in deep reinforcement learning. In *Proceedings of the 5th Annual NeurIPS Deep Learning Workshop*, December 2019.

[122] Gabriel de la Cruz Jr., Yunshu Du, and Matthew E. Taylor. Jointly pre-training with supervised, autoencoder, and value losses for deep reinforcement learning. In *Proceedings of the Adaptive and Learning Agents Workshop at AAMAS*, May 2019.

[123] Bilal Kartal, Pablo Hernandez-Leal, Chao Gao, and Matthew E. Taylor. Safer deep rl with shallow mcts: A case study in pommerman. In *Proceedings of the Adaptive and Learning Agents Workshop at AAMAS*, May 2019.

[124] Bilal Kartal, Pablo Hernandez-Leal, and Matthew E. Taylor. Using monte carlo tree search as a demonstrator within asynchronous deep rl. In *Proceedings of the AAAI Workshop on Reinforcement Learning in Games*, January 2019.

[125] Alex Kearney, Anna Koop, Craig Sherstan, Johannes Gunther, Richard S. Sutton, Patrick M. Pilarski, and Matthew E. Taylor. Evaluating predictive knowledge. In *Proceedings of the Fall AAAI Symposium on Reasoning and Learning in Real-World Systems for Long-Term Autonomy*, October 2018.

[126] Weixun Wang, Jianye Hao, Yixi Wang, and Matthew E. Taylor. Achieving cooperation through deep multi-agent reinforcement learning in sequential prisoner's dilemmas. In *Proceedings of the Adaptive and Learning Agents Workshop at AAMAS*, 2018.

[127] Bikramjit Banerjee and Matthew E. Taylor. Coordination confidence based human-multi-agent transfer learning for collaborative teams. In *Proceedings of the Adaptive and Learning Agents Workshop at AAMAS*, 2018. Best paper award nominee.

[128] A. Leah Zulas, Kaitlyn I. Franz, Darrin Griechen, and Matthew E. Taylor. Solar decathlon competition: Towards a solar-powered smart home. In *Proceedings of the AI for Smart Grids and Buildings Workshop (at AAAI)*, February 2017.

[129] Bei Peng, James MacGlashan, Robert Loftin, Michael Littman, David Roberts, and Matthew E. Taylor. Curriculum design for machine learners in sequential decision tasks. In *Proceedings of the Adaptive and Learning Agents Workshop at AAMAS*, 2017.

[130] Ariel Rosenfeld, Matthew E. Taylor, and Sarit Kraus. Speeding up tabular reinforcement learning using state-action similarities. In *Proceedings of the Adaptive and Learning Agents Workshop at AAMAS*, 2017. Best paper award nominee.

[131] Matthew E. Taylor and Sakire Arslan Ay. AI Projects for Computer Science Capstone Classes (Extended Abstract). In *Proceedings of the Seventh Symposium on Educational Advances in Artificial Intelligence*, February 2017.

[132] Timothy Lewis, Amy Hurst, Matthew E. Taylor, and Cynthia Matuszek. Using Language Groundings for Context-Sensitive Text Prediction. In *Proceedings of EMNLP 2016 Workshop on Uphill Battles in Language Processing*. Austin, TX, USA, November 2016.

[133] Robert Loftin, James MacGlashan, Bei Peng, Matthew E. Taylor, Michael L. Littman, and David L. Roberts. Towards Behavior-Aware Model Learning from Human-Generated Trajectories. In *AAAI Fall Symposium on Artificial Intelligence for Human-Robot Interaction*. Arlington, VA, USA, November 2016.

[134] William Curran, Tim Brys, David Aha, Matthew E. Taylor, and William D. Smart. Dimensionality Reduced Reinforcement Learning for Assistive Robots. In *AAAI 2016 Fall Symposium on Artificial Intelligence: Workshop on Artificial Intelligence for Human-Robot Interaction*. Arlington, VA, USA, November 2016.

[135] James MacGlashan, Michael L. Littman, David L. Roberts, Robert Loftin, Bei Peng, and Matthew E. Taylor. Convergent Actor Critic by Humans. In *Workshop on Human-Robot Collaboration: Towards Co-Adaptive Learning Through Semi-Autonomy and Shared Control (at IROS)*, October 2016.

[136] Ruofei Xu, Robin Hartshorn, Ryan Huard, James Irwin, Kaitlyn Johnson, Gregory Nelson, Jon Campbell, Sakire Arslan Ay, and Matthew E. Taylor. Towards a Semi-Autonomous Wheelchair for Users with ALS. In *Proceedings of Workshop on Autonomous Mobile Service Robots (at IJCAI)*. New York City, NY, USA, July 2016.

[137] Yunshu Du, Gabriel V. de la Cruz Jr., James Irwin, and Matthew E. Taylor. Initial Progress in Transfer for Deep Reinforcement Learning Algorithms. In *Proceedings of Deep Reinforcement Learning: Frontiers and Challenges workshop (at IJCAI)*. New York City, NY, USA, July 2016.

[138] Bei Peng, James MacGlashan, Robert Loftin, Michael L. Littman, David L. Roberts, and Matthew E. Taylor. An Empirical Study of Non-Expert Curriculum Design for Machine Learners. In *Proceedings of the Interactive Machine Learning workshop (at IJCAI)*. New York City, NY, USA, July 2016.

[139] Yunshu Du and Matthew E. Taylor. Work In-progress: Mining the Student Data for Fitness . In *Proceedings of the 12th International Workshop on Agents and Data Mining Interaction (ADMI) (at AAMAS)*. Singapore, May 2016.

[140] David Isele, José Marcio Luna, Eric Eaton, Gabriel V. de la Cruz Jr., James Irwin, Brandon Kallaher, and Matthew E. Taylor. Work in Progress: Lifelong Learning for Disturbance Rejection on Mobile Robots. In *Proceedings of the Adaptive Learning Agents (ALA) workshop (at AAMAS)*. Singapore, May 2016.

[141] Zhaodong Wang and Matthew E. Taylor. Effective Transfer via Demonstrations in Reinforcement Learning: A Preliminary Study. In *AAAI 2016 Spring Symposium*, March 2016.

[142] Pablo Hernandez-Leal, Matthew E. Taylor, Benjamin Rosman, L. Enrique Sucar, and Enrique Munoz de Cote. Identifying and Tracking Switching, Non-stationary Opponents: a Bayesian Approach. In *Proceedings of the Multiagent Interaction without Prior Coordination workshop (at AAAI)*. Phoenix, AZ, USA, February 2016.

[143] Yusen Zhan and Matthew E. Taylor. Online Transfer Learning in Reinforcement Learning Domains. In *Proceedings of the AAAI Fall Symposium on Sequential Decision Making for Intelligent Agents (SDMIA)*, November 2015.

[144] Mitchell Scott, Bei Peng, Madeline Chili, Tanay Nigam, Francis Pascual, Cynthia Matuszek, and Matthew E. Taylor. On the Ability to Provide Demonstrations on a UAS: Observing 90 Untrained Participants Abusing a Flying Robot. In *Proceedings of the AAAI Fall Symposium on Artificial Intelligence and Human-Robot Interaction (AI-HRI)*, November 2015.

[145] William Curran, Tim Brys, Matthew E. Taylor, and William D. Smart. Using PCA to Efficiently Represent State Spaces. In *ICML-2015 European Workshop on Reinforcement Learning*. Lille, France, July 2015.

[146] Yawei Zhang, Yunxiang Ye, Zhaodong Wang, Matthew E. Taylor, Geoffrey A. Hollinger, and Qin Zhang. Intelligent In-Orchard Bin-Managing System for Tree Fruit Production. In *Proceedings of the Robotics in Agriculture workshop (at ICRA)*, May 2015.

[147] Bei Peng, Robert Loftin, James MacGlashan, Michael L. Littman, Matthew E. Taylor, and David L. Roberts. Language and Policy Learning from Human-delivered Feedback. In *Proceedings of the Machine Learning for Social Robotics workshop (at ICRA)*, May 2015.

[148] Pablo Hernandez-Leal, Matthew E. Taylor, Enrique Munoz de Cote, and L. Enrique Sucar. Learning Against Non-Stationary Opponents in Double Auctions. In *Proceedings of the Adaptive Learning Agents (ALA) workshop 2015*. Istanbul, Turkey, May 2015. Finalist for Best Student Paper.

[149] Gabriel V. de la Cruz Jr., Bei Peng, Walter S. Lasecki, and Matthew E. Taylor. Generating Real-Time Crowd Advice to Improve Reinforcement Learning Agents. In *Proceedings of the Learning for General Competency in Video Games workshop (AAAI)*, January 2015.

[150] James Macglashan, Michael L. Littman, Robert Loftin, Bei Peng, David Roberts, and Matthew E. Taylor. Training an Agent to Ground Commands with Reward and Punishment. In *Proceedings of the Machine Learning for Interactive Systems workshop (at AAAI)*, July 2014.

[151] Haitham Bou Ammar, Eric Eaton, Matthew E. Taylor, Decibal C. Mocanu, Kurt Driessens, Gerhard Weiss, and Karl Tuyls. An Automated Measure of MDP Similarity for Transfer in Reinforcement Learning. In *Proceedings of the Machine Learning for Interactive Systems workshop (at AAAI)*, July 2014.

[152] Chris HolmesParker, Matthew E. Taylor, Yusen Zhan, and Kagan Tumer. Exploiting Structure and Agent-Centric Rewards to Promote Coordination in Large Multiagent Systems. In *Proceedings of the Adaptive and Learning Agents workshop (at AAMAS)*, May 2014.

[153] Yusen Zhan, Anestis Fachantidis, Ioannis Vlahavas, and Matthew E. Taylor. Agents Teaching Humans in Reinforcement Learning Tasks. In *Proceedings of the Adaptive and Learning Agents workshop (at AAMAS)*, May 2014.

[154] Tong Pham, Tim Brys, and Matthew E. Taylor. Learning Coordinated Traffic Light Control. In *Proceedings of the Adaptive and Learning Agents workshop (AAMAS)*, May 2013.

[155] Nicholas Carboni and Matthew E. Taylor. Preliminary Results for 1 vs. 1 Tactics in Starcraft. In *Proceedings of the Adaptive and Learning Agents workshop (AAMAS)*, May 2013.

[156] Anestis Fachantidis, Ioannis Partalas, Matthew E. Taylor, and Ioannis Vlahavas. Autonomous Selection of Inter-Task Mappings in Transfer Learning (extended abstract). In *The AAAI 2013 Spring Symposium — Lifelong Machine Learning*, March 2013.

[157] Ravi Balasubramanian and Matthew E. Taylor. Learning for Mobile-Robot Error Recovery (Extended Abstract). In *The AAAI 2013 Spring Symposium — Designing Intelligent Robots: Reintegrating AI II*, March 2013.

[158] Sanjeev Sharma and Matthew E. Taylor. Autonomous Waypoint Generation Strategy for On-Line Navigation in Unknown Environments. In *IROS Workshop on Robot Motion Planning: Online, Reactive, and in Real-Time*, October 2012.

[159] Matthew Adams, Robert Loftin, Matthew E. Taylor, Michael Littman, and David Roberts. An Empirical Analysis of RL's Drift From Its Behaviorism Roots. In *Proceedings of the Adaptive and Learning Agents workshop (AAMAS)*, June 2012.

[160] Lisa Torrey and Matthew E. Taylor. Help an Agent Out: Student/Teacher Learning in Sequential Decision Tasks. In *Proceedings of the Adaptive and Learning Agents workshop (AAMAS)*, June 2012.

[161] Haitham Bou Ammar, Matthew E. Taylor, Karl Tuyls, and Gerhard Weiss. Reinforcement Learning Transfer using a Sparse Coded Inter-Task Mapping. In *Proceedings of the European Workshop on Multi-agent Systems*, November 2011.

[162] Anestis Fachantidis, Ioannis Partalas, Matthew E. Taylor, and Ioannis Vlahavas. Transfer Learning via Multiple Inter-Task Mappings. In *Proceedings of European Workshop on Reinforcement Learning (ECML)*, September 2011.

[163] W. Bradley Knox, Matthew E. Taylor, and Peter Stone. Understanding Human Teaching Modalities in Reinforcement Learning Environments: A Preliminary Report. In *Proceedings of the Agents Learning Interactively from Human Teachers workshop (IJCAI)*, July 2011.

[164] Matthew E. Taylor. Model Assignment: Reinforcement Learning in a Generalized Mario Domain. In *Proceedings of the Second Symposium on Educational Advances in Artificial Intelligence*, August 2011.

[165] Matthew E. Taylor. Teaching Reinforcement Learning with Mario: An Argument and Case Study. In *Proceedings of the Second Symposium on Educational Advances in Artificial Intelligence*, August 2011.

[166] Paul Scerri, Balajee Kannan, Pras Velagapudi, Kate Macarthur, Peter Stone, Matthew E. Taylor, John Dolan, Alessandro Farinelli, Archie Chapman, Bernadine Dias, and George Kantor. Flood Disaster Mitigation: A Real-world Challenge Problem for Multi-Agent Unmanned Surface Vehicles. In *Proceedings of the Autonomous Robots and Multirobot Systems workshop (AAMAS)*, May 2011.

[167] Haitham Bou Ammar and Matthew E. Taylor. Common subspace transfer for reinforcement learning tasks. In *Proceedings of the Adaptive and Learning Agents workshop (AAMAS)*, May 2011.

[168] Jun young Kwak, Zhengyu Yin, Rong Yang, Matthew E. Taylor, and Milind Tambe. Robust Execution-time Coordination in DEC-POMDPs Under Model Uncertainty. In *Proceedings of the Workshop on Multiagent Sequential Decision Making in Uncertain Domains (AAMAS)*, May 2011.

[169] Scott Alfeld, Kumera Berkele, Stephen A. Desalvo, Tong Pham, Daniel Russo, Lisa Yan, and Matthew E. Taylor. Reducing the team uncertainty penalty: Empirical and theoretical approaches. In *Proceedings of the Workshop on Multiagent Sequential Decision Making in Uncertain Domains (AAMAS)*, May 2011.

[170] Shimon Whiteson, Brian Tanner, Matthew E. Taylor, and Peter Stone. Protecting Against Evaluation Overfitting in Empirical Reinforcement Learning. In *Proceedings of the IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, April 2011.

[171] Matthew E. Taylor, Halit Bener Suay, and Sonia Chernova. Using Human Demonstrations to Improve Reinforcement Learning. In *The AAAI 2011 Spring Symposium — Help Me Help You: Bridging the Gaps in Human-Agent Collaboration*, March 2011.

[172] Matthew E. Taylor and Sonia Chernova. Integrating Human Demonstration and Reinforcement Learning: Initial Results in Human-Agent Transfer. In *Proceedings of the Agents Learning Interactively from Human Teachers workshop (AAMAS)*, May 2010.

[173] Scott Alfeld, Matthew E. Taylor, Prateek Tandon, and Milind Tambe. Towards a Theoretic Understanding of DCEE. In *Proceedings of the Distributed Constraint Reasoning workshop (AAMAS)*, May 2010.

[174] Samuel Barrett, Matthew E. Taylor, and Peter Stone. Transfer Learning for Reinforcement Learning on a Physical Robot. In *Proceedings of the Adaptive and Learning Agents workshop (AAMAS)*, May 2010.

[175] Jason Tsai, Emma Bowring, Shira Epstein, Natalie Fridman, Prakhar Garg, Gal Kaminka, Andrew Ogden, Milind Tambe, and Matthew E. Taylor. Agent-based Evacuation Modeling: Simulating the Los Angeles International Airport. In *Proceedings of the Workshop on Emergency Management: Incident, Resource, and Supply Chain Management*, November 2009.

[176] Matthew E. Taylor, Manish Jain, Prateek Tandon, and Milind Tambe. Using DCOPs to Balance Exploration and Exploitation in Time-Critical Domains. In *Proceedings of the IJCAI 2009 Workshop on Distributed Constraint Reasoning*, July 2009.

[177] Matthew E. Taylor, Chris Kiekintveld, Craig Western, and Milind Tambe. Is There a Chink in Your ARMOR? Towards Robust Evaluations for Deployed Security Systems. In *Proceedings of the IJCAI 2009 Workshop on Quantitative Risk Analysis for Security Applications*, July 2009.

[178] Matthew E. Taylor and Peter Stone. Categorizing Transfer for Reinforcement Learning. In *Poster at the Multidisciplinary Symposium on Reinforcement Learning*, June 2009.

[179] Shimon Whiteson, Brian Tanner, Matthew E. Taylor, and Peter Stone. Generalized Domains for Empirical Evaluations in Reinforcement Learning. In *Proceedings of the Fourth Workshop on Evaluation Methods for Machine Learning at ICML-09*, June 2009.

[180] Manish Jain, Matthew E. Taylor, Makoto Yokoo, and Milind Tambe. DCOPs Meet the Real World: Exploring Unknown Reward Matrices with Applications to Mobile Sensor Networks. In *Proceedings of the Third International Workshop on Agent Technology for Sensor Networks (AAMAS)*, May 2009.

[181] Jun young Kwak, Pradeep Varakantham, Matthew E. Taylor, Janusz Marecki, Paul Scerri, and Milind Tambe. Exploiting Coordination Locales in Distributed POMDPs via Social Model Shaping. In *Proceedings of the Fourth Workshop on Multi-agent Sequential Decision-Making in Uncertain Domains (AAMAS)*, May 2009.

[182] Matthew E. Taylor, Chris Kiekintveld, Craig Western, and Milind Tambe. Beyond Runtimes and Optimality: Challenges and Opportunities in Evaluating Deployed Security Systems. In *Proceedings of the AAMAS-09 Workshop on Agent Design: Advancing from Practice to Theory*, May 2009.

[183] Matthew E. Taylor. Assisting Transfer-Enabled Machine Learning Algorithms: Leveraging Human Knowledge for Curriculum Design. In *The AAAI 2009 Spring Symposium on Agents that Learn from Human Teachers*, March 2009.

[184] Matthew E. Taylor, Nicholas K. Jong, and Peter Stone. Transferring Instances for Model-Based Reinforcement Learning. In *The Adaptive Learning Agents and Multi-Agent Systems (ALAMAS+ALAG) workshop at AAMAS*, May 2008.

[185] Matthew E. Taylor, Katherine E. Coons, Behnam Robatmili, Doug Burger, and Kathryn S. McKinley. Policy Search Optimization for Spatial Path Planning. In *NIPS-07 workshop on Machine Learning for Systems Problems*, December 2007. (Two page extended abstract.).

[186] Matthew E. Taylor, Gregory Kuhlmann, and Peter Stone. Accelerating Search with Transferred Heuristics. In *ICAPS-07 workshop on AI Planning and Learning*, September 2007.

[187] Matthew E. Taylor and Peter Stone. Representation Transfer for Reinforcement Learning. In *AAAI 2007 Fall Symposium on Computational Approaches to Representation Change during Learning and Development*, November 2007.

[188] Shimon Whiteson, Matthew E. Taylor, and Peter Stone. Adaptive Tile Coding for Reinforcement Learning. In *NIPS workshop on: Towards a New Reinforcement Learning?*, December 2006.

[189] Matthew E. Taylor, Shimon Whiteson, and Peter Stone. Transfer Learning for Policy Search Methods. In *ICML workshop on Structural Knowledge Transfer for Machine Learning*, June 2006.

[190] Matthew E. Taylor and Peter Stone. Speeding up Reinforcement Learning with Behavior Transfer. In *AAAI 2004 Fall Symposium on Real-life Reinforcement Learning*, October 2004.

Technical Reports

[191] Lonny Simonian and Matthew E. Taylor. Applications for UAVs in electric utility construction. Technical Report F3415, The Foundation for Electricial Construction Inc. (ELECTRI), January 2017. 1–69 pp. http://electri.org/research/applications-uavs-electric-utility-construction.

[192] Shimon Whiteson, Matthew E. Taylor, and Peter Stone. Adaptive Tile Coding for Value Function Approximation. Technical Report AI-TR-07-339, University of Texas at Austin, 2007.

Dissertation

[193] Matthew E. Taylor. *Autonomous Inter-Task Transfer in Reinforcement Learning Domains*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, August 2008.

`scholar.google.com` statistics (`http://bit.ly/uWsh6m`):
Citation count: 5175, h-index: 34, i10-index: 83

Patents Pending, with Borealis AI

Opponent Modeling with Asynchronous Methods in Deep RL
System and Method for Deep Reinforcement Learning
System and Method for Multi-Type Mean Field Reinforcement Machine Learning

Major Federal Grants

8/17–8/20 *NAVAIR*. Cognitive Adaptation and Mission Optimization (CAMO) for Autonomous Teams of UAS Platforms

PI: HolmesParker, OSKI. Co-PI: Taylor. $366,690 awarded to Taylor Lab, **$974,990** total

8/17–8/20 *NSF*. NRI: INT: Learning-Enabled Robot Support of Daily Activities for Successful Activity Completion.

PI: Taylor. Co-PIs: Cook & Schmitter-Edgecombe. $400,000 awarded to Taylor Lab, **$999,996** total

8/16–8/17 *NSF IIS1643614*. EAGER: Income Learning: A New Model for Behavior-Analysis-Inspired Learning from Human Feedback.

PI: Taylor. **$70,000** awarded to Taylor Lab

5/16–5/18 *NASA Phase-II SBIR, NNX16CD07C*. Command and Control Software for Single-Operator Multiple UAS Missions.

PI: HolmesParker, OKSI. Co-PI: Taylor. $133,390 awarded to Taylor Lab, **$729,831** total

10/14–8/17 *USDA 2014-67021-22174* (NSF National Robotics Initiative). Intelligent In-Orchard Bin Managing System for Tree Fruit Production.

PI: Zhang. Co-PIs: Hollinger and Taylor. $208,215 awarded to Taylor Lab, **$1,010,169** total

2/14–2/16 *AFRL FA8750-14-1-0069*. Lifelong Transfer Learning for Heterogeneous Teams of Agents in Sequential Decision Processes.

PI: Taylor. Co-PIs: Eaton and Ruvolo. $277,935 awarded to Taylor lab, **$608,182** total

2/14–2/16 *AFRL FA8750-14-1-0070*. Curriculum Development for Transfer Learning in Dynamic Multiagent Settings.

PI: Stone. Co-PI: Taylor. $200,772 awarded to Taylor lab, **$640,439** total

9/13–8/16 *NSF IIS-1319412*. RI: Small: Collaborative Research: Speeding Up Learning through Modeling the Pragmatics of Training.

PI: Roberts. Co-PIs: Littman and Taylor, $135,000 awarded to Taylor Lab, **$439,203** total

9/12–8/17 *NSF IIS-1149917*. CAREER: A Multiagent Teacher/Student Framework for Sequential Decision Making Tasks.

PI: Taylor. **$402,065** awarded to Taylor Lab

Other Research Grants

7/17–6/18 *The Joint Center for Aerospace Technology Innovation*. Precision Aerial Delivery via Low-Weight, Low-Cost UAS.

PI: Taylor. Co-PI: Matveev and Swensen. $72,430 awarded to Taylor Lab, **$79,441** total

7/16 *NVidia GPU Grant Program*. Training and Transfer in Deep Convolutional Neural Networks.

PI: Taylor. NVidia Titan X awarded to Taylor Lab, **$863** in donated hardware

6/16–6/17 *The Joint Center for Aerospace Technology Innovation*. Precision Aerial Delivery via Low-Weight, Low-Cost UAS.

PI: Taylor. Co-PI: Matveev. $61,551 awarded to Taylor Lab, **$75,000** total

4/16–4/17 *WSU Emerging Research Issues in Washington Agriculture*. A Proof of Concept System using Autonomous Unmanned Aerial Systems for Mitigating Bird Damage in Fruit and Berry Crops.
PI: Karkee. Co-PI: Taylor. $40,000 awarded to Taylor Lab, **$80,000** total

9/15–8/16 *Google Faculty Research Award*. Robot Control via RGB Video and Convolutional Neural Networks.
PI: Taylor. Co-PI: Crandall. **$37,763** awarded to Taylor Lab

2015 *Washington State University's Energy Systems Innovation Center*. Greater Renewable Integration and Deployment using Model predictive ENergy Dispatch GRID-MEND.
PI: McLarty. Co-PIs: Mehrizi-Sani and Taylor. $3,000 awarded to Taylor Lab, **$10,000** total

2014 *Washington State University's Energy Systems Innovation Center*. Unmanned Aerial Systems for Autonomous Power Line Monitoring.
PI: Taylor. Co-PIs: Leachman and Mehrizi-Sani. $3,000 awarded to Taylor Lab, **$9,000** total

2014 *Washington State Blueberry Commission*. Unmanned Aerial Systems (UASs) for Mitigating Bird Damage in Blueberry Crops: Proof of Concept.
PI: Karkee. Co-PIs: Leachman, Taylor, and Zhang. $9,771 awarded to Taylor Lab, **$19,543** total

2012 *Lee Pesky Learning Center Foundation*. Examination of cognitive performance of individuals with developmental dyslexia on a visuo-spatial virtual maze task.
Co-PIs: Gabel and Taylor. $10,869 awarded to Taylor Lab, **$49,853** total

2011 *Lafayette College*. Think Tank Grant: Creation of a virtual Hebb-Williams maze to compare cognitive performance of dyslexics with a mouse model of dyslexia.
Co-PIs: Gabel and Taylor. $12,500 awarded to Taylor Lab, **$25,000** total

Non-Research Grants

(All single PI and no funds were allocated for my research.)

2016 *WSU Student Technology Fee Allocation*. Tutorial and Development Support for WSU Robotics Club. **$14,778** awarded to WSU Robotics Club for hardware

2015 *Artificial Intelligence Journal Sponsorship*. AAAI Spring Symposium Series. **€7,000** for student travel

2014 *NSF IIS-1444754*. 19th Annual SIGART/AAAI Doctoral Consortium. **$17,610** for student travel

2013 *Journal of Artificial Intelligence*. 2013 AAAI Fall Symposium Student Funding. **€8,000** for student travel

2012 *NSF IIS-1231124*. EAAI-12: The Third Symposium on Educational Advances in Artificial Intelligence. **$17,000** for professor and student travel

2011 *NSF IIS-1125979*. Doctoral Mentoring Consortium at the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2011). **$20,000** for student travel

2010 *Lafayette College*. Curricular and Pedagogical Development Grant for developing CS 102 – Personal Robotics. **$10,400** for hardware

Total funds raised: $6.3M USD
Total funds raised as PI: $2.5M USD

Teaching

Courses at University of Alberta

Independent Study—Learning from Demonstrations Using Only State Transitions, *F18*.
Independent Study—Learning Representations, *W18*.

Courses at Washington State University

Reinforcement Learning, *Lecture*, S14, S15, S17.
Introduction to Robotics, *Lecture*, F13, F14, S16, F16, F17.
Independent Study—Machine Vision, *S17*.

Courses at Lafayette College

Artificial Intelligence, *Lecture*, F11.
Computer Organization, *Lecture and Lab*, F10, F11, F12.
Computers and Society, *Lecture*, S12.
Independent Study—Machine Learning, *F11,S12*.
Introduction to Computational Science, *Lecture and Lab*, F12.
Introduction to Machine Learning, *Lecture*, F10.
Personal Robotics, *Lecture and Lab*, S12.
Principles of Computer Science 1, *Lecture and Lab*, S11.

Courses at UT-Austin

Software Systems: Unix, *Lecture*, Assistant Instructor (Instructor of record), F07, S08.

Other

Teaching Assistant: Computer Fluency, *Lab and Discussion Section*, UT-Austin, F05.
TA and Tutor, *Multiple computer science and physics courses*, Amherst College, S98–F00.

Tutorials

2018 **Introduction to Reinforcement Learning**.
One day tutorial on RL. Presented to RBC audience

2017 **Interactive Machine Learning: From Classifiers to Robotics**.
Co-taught with Bradley H. Hayes and Ece Kamar. Presented at AAAI-17, AAMAS-17, ICJNN-17, and IJCAI-17

2015 **Speeding up Reinforcement Learning at Adaptive and Learning Agents Workshop**.
Co-taught with Tim Brys. Workshop at AAMAS conference

2010 **Transfer Learning and other RL Speed-up Techniques at AAMAS**.
Co-taught with Dr. Alessandro Lazaric. Part of a 1-day tutorial on Reinforcement Learning

2009 **Transfer Learning for Reinforcement Learning Domains at ECML**.
Co-taught with Dr. Alessandro Lazaric

2009 **Transfer Learning in Reinforcement Learning at AAMAS**.
Co-taught with Dr. Alessandro Lazaric. Part of a 1.5-day tutorial on Reinforcement Learning

Graduated Students (PhD)

2019 Zhaodong Wang, "Knowledge Transfer in Reinforcement Learinng: How Agents Should Benefit from Prior Knowledge". WSU, EECS

2019 Yang Hu, "Teaching Effectiveness of Intelligent Tutoring Systems". WSU, EECS

2018 Bei Peng, "Learning from Human Teachers: Supporting How People Want to Teach in Interactive Machine Learning". WSU, EECS

2018 Chris Cain, "TINGLE — Topic-Independent Gamification Learning Environment". WSU, EECS

2016 Yusen Zhan, "Policy Advice, Non-convex and Distributed Optimization in Reinforcement Learning". WSU, EECS

2016 Tim Brys, "Reinforcement Learning with Heuristic Information". Vrije Universiteit Brussel, Computer Science (Co-Supervisor with Ann Nowé)

2013 Haitham Bou Ammar, "Automated Transfer for Reinforcement Learning". Maastricht University, Information and Knowledge Systems (Co-Supervisor with Karl Tuyls and Gerhard Weiss)

Thesis Committee (PhD)

2019 Selina Akter: WSU Computer Science

2018 External Member: Chris Reinke, Okinawa Institute of Science and Technology

2018 Tao Zeng: WSU Computer Science

2017 Brian Thomas: WSU Computer Science

2016 External Member: Halit Bener Suay, Worcester Polytechnic Institute, Department of Computer Science

2016 External Member: Adam Taylor, University of Dublin, School of Computer Science and Statistics

2015 Yibo Yao, WSU Computer Science

2015 Jeyanthi Narasimhan, WSU Computer Science

2015 Abhik Ray, WSU Computer Science

2015 External Member: Mayank Daswani, Australian National University, Research School of Computer Science

2014 Kyle Feuz: WSU Computer Science

2012 External Member: Mihail Mihaylov, Vrije Universiteit Brussel, Computer Science

Graduated Students (MS)

2020 Paniz Behboudian, "Useful Policy Invariant Shaping from Arbitrary Advice". U of Alberta, CS (Co-supervisor with Michael Bowling and Yash Satsangi)

2019 Gabriel V. de la Cruz Jr., "Accelerate the Learning Speed of Deep Reinforcement Learning by Pre-training with Non-Expert Human Demonstrations". WSU, EECS

2017 Viresh Duvvuri, "BATON: A Low-cost, Low-weight Drone for Precision Delivery". WSU, EECS

2017 Shivam Goel, "Bird Deterrence with Drones: Engineering a Smart Scarecrow". WSU, EECS

2017 Leah Amanda Zulas, "Modifying Smart Home to Smart Phone Notifications Using Reinforcement Learning Algorithms". WSU, EECS

Thesis Committee (MS)

2018 Mitchell Scott: WSU Materials and Mechanical Engineering

2016 Heon Jo: WSU Materials and Mechanical Engineering

2014 Sal Bagaveyev: WSU Computer Science

Thesis Committee (BS)

2011 Reader: Miguel Haruki Yamaguchi, Lafayette College, Computer Science

Major Talks

Invited talk in the AAAI-18 Emerging Topic Track on Human-AI Collaboration

February, 2018 Improving Reinforcement Learning with Human Input. New Orleans, LA.

Keynote talk at the 2nd IEEE International Conference on Agents (ICA)

July, 2017 Improving Reinforcement Learning with Human Input. Tsinghua University, Beijing, China.

Keynote talk at the 14th Bar-Illan Symposium on the Foundations of Artificial Intelligence (BISFAI)

June, 2017 Improving Reinforcement Learning with Human Input. Bar-Illan University, Tel Aviv, Israel.

Invited Talk at NIPS-16 workshop on the Future of Interactive Machine Learning

December, 2016 What's in a Reward? Teaching Agents and Robots with Non-Expert Reinforcement. Barcelona, Spain.

Invited Talk at Global Artificial Intelligence Conference (GAIC)

November 2016 Speeding up sequential decision tasks via transfer learning. Jiading, China.

Invited Talks & Colloquia

Human-in-the-Loop RL

July, 2019 Deep Learning, Reinforcement Learning Summer School: Edmonton, Alberta

Learning Sequential Tasks from Human Feedback

October, 2018 Re-Work Deep Learning Summit: Toronto, Ontario

A Whirlwind Introduction to Artificial Intelligence & Machine Learning

May, 2018 Gerber Foundation Board: Grand Rapids, MI

AI & Wealth Management

February, 2018 Royal Bank of Canada: London, UK. Lunch & Learn

Next Steps in Transfer and Lifelong Learning

May, 2017 AAMAS-17 workshop: 1st Workshop on Transfer in Reinforcement Learning (TiRL): Sao Paulo, Brazil. Invited talk

Defining Tasks with Human-Provided Reward

November 2016 Vrije Universiteit Brussel: Brussels, Belgium. Department Seminar

Learning from others: Speeding up sequential decision tasks via transfer learning

November 2016 Nanjing University: Nanjing, China. Departmental Seminar

Improving Reinforcement Learning with Human Input

November 2018 The Council for Scientific and Industrial Research: Pretoria, South Africa. Invited Talk
October 2018 University of the Witwatersrand: Johannesburgh, South Africa. Invited Talk
May 2018 University of Calgary: Calgary, AB. Department Seminar
March 2018 University of Alberta: Edmonton, AB. AI Seminar
June 2017 New Mexico State University: Las Curces, NM. Invited Talk
June 2017 University of Texas at el Paso: el Paso, TX. Invited Talk
May 2017 SRI: Menlo Park, CA. Invited Talk
February 2017 Google Research: Mountain View, CA. Invited Talk
April 2017 Pacific Northwest National Labs: Richland, WA. AIM Seminar
September 2016 Microsoft Research: Redmond, WA. Invited Research Talk. Talk online at: `https://www.youtube.com/watch?v=7B1fcITbyWo`

Learning from Demonstration, Human Feedback, and Environmental Rewards

May 2016 UMass-Amherst: Amherst, MA. Machine Learning and Friends Lunch

Dude, Where's My Bias? Improving RL with Human Demonstrations

April 2016 10th Barbados Workshop on Reinforcement Learning: Holetown, Barbados

Learning from Others: Speeding up Sequential Decision Tasks with Knowledgable Agents and Humans

March 2016 Naval Research Laboratory: Washington, DC. Seminar for David Aha's group
January 2016 University of Dublin: Dublin, Ireland. Departmental seminar

Agents as Teachers and Learners

April 2015 Google Deepmind: London, UK. Research Seminar

Agents: Teaching and Transfer

January 2015 AAAI Conference: Austin, TX. Invited talk in the AAAI-15 workshop on Knowledge, Skill, and Behavior Transfer in Autonomous Robots

Planning and Learning in Multi-agent Systems

June 2014 University of Washington: Seattle, WA. William E. Boeing Department of Aeronautics and Astronautics Seminar

Agents as Teachers and Learners

October 2013 Microsoft: Redmond, WA. Tech Talk
August 2013 University of Southern California: Los Angeles, CA. Computer Science Colloquium
March 2013 AAAI Spring Symposium: Palo Alto, CA. Invited talk at Lifelong Machine Learning Symposium
March 2013 Oregon State University: Corvalis, OR. Artificial Intelligence Colloquium

Real Life Reinforcement Learning

October 2012 Washington State University: Pullman, WA. Department Seminar
October 2012 Drexel University: Philadelphia, PA. Department Seminar
June 2012 Vrije Universiteit Brussel: Brussels, Belgium. Department Seminar

May 2012 Maastricht University: Maastricht, The Netherlands. Department Colloquium
April 2012 Portland State University: Portland, PA. Department Seminar
February 2012 Stevens Institute of Technology: Hoboken, NJ. Department Seminar
February 2012 Virginia Tech: Blacksburg, VA. Department Seminar

Towards a Student/Teacher framework for Reinforcement Learning Tasks

October 2011 Rutgers University: Piscataway, NJ. Department Seminar

Help an Agent Out: Learning From the Environment and Humans

November 2011 Bryn Mawr College: Bryn Mawr, PA. FLiCS Colloquium, sponsored by Bryn Mawr, Haverford, Swathmore, and Villanova
October 2011 Dickinson College: Carlisle, PA. Departmental Seminar

Integrating Reinforcement Learning and Human Demonstration

January 2011 UT-Austin: Austin, TX. Peter Stone's group meeting

Transfer Learning and Multi-Agent Exploration: Towards Real-Life Learning Agents

May 2010 Caltech: Pasadena, CA. Machine Learning Lunch
February 2010 USC Information Sciences Institute: Marina del Ray, CA. AI Seminar

Real-Life Learning Agents

December 2009 Lafayette College: Easton, PA. Departmental Seminar

Balancing Multi-agent Exploration and Exploitation in Time-Critical Domains

May 2009 UT-Austin: Austin, TX. Forum for Artificial Intelligence

Transfer Learning with Inter-Task Mappings

April 2008 Barbados, Bellairs Research Institute. Third Barbados Workshop on Reinforcement Learning

[Towards] Autonomous Inter-Task Transfer in Reinforcement Learning Domains

December 2008 Lockheed Martin Advanced Technology Laboratories: Cherry Hill, NJ. Advanced Technology Seminar
June 2008 USC: Los Angeles, CA. Departmental Seminar
April 2008 Carnegie Mellon University: Pittsburgh, PA. Manuela Veloso's group meeting
October 2007 University of Wisconsin: Madison, WI. Artificial Intelligence Seminar
October 2007 University of Washington: Seattle, WA. Pedro Domingos's group meeting
October 2007 Microsoft Research: Redmond, WA. Machine Learning Reading Group
October 2007 Harvard University: Cambridge, MA. AI Research Series
October 2007 Massachusetts Institute of Technology: Cambridge, MA. Nicholas Roy's and Leslie Pack Kaelbling's group meeting
October 2007 Brown University: Providence, RI. Seminar
October 2007 University of Alberta: Edmonton, CA. Artificial Intelligence Seminar Series

Faster Inference through Reinforcement Learning

August 2006 Cycorp: Austin, TX. Seminar

Speeding up Reinforcement Learning via Behavior Transfer

September 2005 Amherst College: Amherst, MA. Departmental Colloquium

September 2005 University of Massachusetts at Amherst: Amherst, MA. Machine Learning and Friends Lunch

June 2005 Stanford University: Stanford, CA. Pat Langley's group meeting

Other Talks

Using human demonstrations to improve reinforcement learning

March 2011 Use-inspired Agents and Multiagent Systems Workshop: Los Angeles, CA

Artificial Intelligence Methods for Risk Management and Analysis

April 2009 Risk Analysis Symposium (RISK-09): Sante Fe, NM

Panel Moderator

December 2016 Panel at the NIPS-16 Workshop on the Future of Interactive Machine Learning: Barcelona, Spain

Panel Member

June 2018 Investure$ panel on "Innovation in Alberta: The Artificial Intelligence Ecosystem": Calgary, AB

July 2017 ICA-17 Panel on "Standard rules for multiple AIs": Beijing, China

November 2013 AAAI Symposium Panel on "How Should Intelligence be Abstracted in AI Research: MDPs, Symbolic Representations, Artificial Neural Networks, or — ?": Washington, DC

August 2011 AAAI Doctoral Consortium Panel: San Francisco, CA

November 2010 Workshop on Reasoning in Adversarial and Noncooperative Environments: Durham, NC

Professional Activities

Editing

2020–current Associate editor for the Journal of Artificial Intelligence Research (JAIR)
2019 Guest editor for special issue of Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)
2011–2020 Editorial board member for the Journal of Artificial Intelligence Research (JAIR)
2011 Guest editor for topical issue in Advances in Complex Systems (ACS)
2010–current Associate editor for Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)

Conference Organization

2020 DAI-20 conference — Program co-chair
2019 AAMAS-19 conference — Program co-chair
2019 AAAI-19 conference — Workshops co-chair
2018 AAAI-18 conference — Workshops co-chair
2015 AAAI-15 conference — Doctoral Consortium co-chair
2014 AAAI-14 conference — Doctoral Consortium co-chair
2015 AAMAS-16 conference — Scholarship co-chair
2015 AAMAS-15 conference — Sponsorship co-chair
2014–2020 International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS) Board of Directors
2016–2018 IFAAMAS Treasurer
2014 AAMAS-14 conference — Publicity chair
2013 AAMAS-13 conference — Doctoral Consortium co-chair
2013 SAC-13 conference — Track co-chair: Cooperative Multi-Agent Systems and Applications
2011 AAMAS-11 conference — Scholarships co-chair

Summer School Organization

2019 Deep Learning Reinforcement Learning Summer School (DLRLSS), Edmonton AB, 7/24/19–8/2/19

Topical Committees

2019–current AAAI Publications Committee

Symposium Series Organization

2012–2015 AAAI symposium series chair
2011 AAAI symposium series co-chair

Workshop Organization

2013 ICML workshop — *Theoretically Grounded Transfer Learning*, co-chair
2009, 2010 AAMAS workshop — *Adaptive and Learning Agents*, co-chair
2008 AAAI workshop — *Transfer Learning for Complex Tasks*, co-chair

Workshop/Symposium Organizing Committee

2012–current AAMAS workshop — *Adaptive and Learning Agents*, senior steering committee
2012, 2013 EAAI symposium — Organizing Committee

2011 IJCAI workshop — *Agents Learning Interactively from Human Teachers*
2010 AAMAS workshop — *Agents Learning Interactively from Human Teachers*
2008, 2009 ICML workshop — *The Annual Reinforcement Learning Competition*

Book Reviewer

Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers: 2016

Book Proposal Reviewer

Cambridge University Press: 2017

Journal Reviewer

Adaptive Behavior: 2009, 2014
Advances in Complex Systems: 2009
Applied Intelligence Journal: 2018
Artificial Intelligence Journal: 2007, 2008, 2009, 2010, 2013, 2014
Communications Biology: 2020
Connection Science: 2015
IEEE Access: 2020
IEEE Transactions on Cybernetics: 2014
IEEE Robotics and Automation Letters: 2016
International Journal of Agent Technologies and Systems: 2008
International Journal of Social Robotics: 2012
Journal of Artificial Intelligence Research: 2007, 2010, 2011, 2012, 2013, 2015, 2016, 2020
Journal of Autonomous Agents and Multi-Agent Systems: 2008, 2009, 2011, 2012, 2013, 2015
Journal of Intelligent Transportation Systems: Technology, Planning, and Operations: 2013
Journal of Machine Learning Research: 2009, 2011,2013, 2014
Journal of Science Education and Technology: 2017
Knowledge Engineering Review: 2013
Machine Learning Journal: 2007, 2010, 2012, 2013, 2014, 2020
Neural Computation: 2013
Neural Networks: 2018
Risk.net: 2020 (a quantitative finance publication)
Robotics: 2013, 2017
Robotics and Autonomous Systems: 2011, 2016
Scientific Reports by (Springer Nature): 2017
Sensors: 2019
Transactions on Computational Intelligence and AI in Games: 2013
Transaction on Services Computing: 2015

Conference Area Chair

AAAI Conference on Artificial Intelligence (AAAI): 2018
Intl. Joint Conf. on Autonomous Agents & Multiagent Systems (AAMAS): 2018

Conference Senior Program Committee

AAAI Conference on Artificial Intelligence (AAAI): 2016, 2017

Intl. Joint Conf. on Autonomous Agents & Multiagent Systems (AAMAS): 2011, 2012, 2015, 2017, 2020

Intl. Joint Conf. on Artificial Intelligence (IJCAI): 2013, 2015, 2018, 2020

IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), "associate editor": 2014, 2016

Conference Program Committee

AAAI Conf. on Artificial Intelligence (AAAI): 2010, 2011, 2012, 2013, 2015

Intl. Joint Conf. on Autonomous Agents & Multiagent System (AAMAS): 2009, 2010, 2013, 2014

Intl. Conf. on Digital Health (DH): 2015, 2016

Conf. of the Spanish Association for Artificial Intelligence (CAEPIA): 2007

Conf. on Robot Learning (CoRL): 2019

European Conf. on Machine Learning (ECML): 2007, 2013

IEEE International Conference on Humanized Computing and Communication (HCC): 2019

Intl. Conf. on Automated Planning and Scheduling (ICAPS): 2017

Intl. Conf. on Machine Learning (ICML): 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2018

Intl. Joint Conf. on Artificial Intelligence (IJCAI): 2009, 2011, 2015

Neural Information Processing Systems (NeurIPS): 2008, 2009, 2011, 2018, 2019

International Conference on Pervasive Technologies Related to Assistive Environments (PETRA): 2017

Robotics Science and Systems (RSS): 2014

ACM Symposium on Applied Computing (SAC), Intelligent Robotics and Multi-Agent Systems Track: 2015, 2016, 2017, 2019

Conference Reviewer

IEEE Intl. Conf. on Robotics and Automation (ICRA): 2011, 2012, 2013, 2014

Intl. Joint Conf. on Artificial Intelligence (IJCAI): 2007

International Conference on Intelligent Robots and Systems (IROS): 2012, 2013, 2020

Intl. Semantic Web Conf. (ISWC): 2007

Book Chapter Reviewer

Cambridge University Press: 2015

Workshop/Symposium Program Committee

AAAI *Doctoral Consortium*. 2012, 2013, 2017, 2020

AAAI workshop — *Applied Adversarial Reasoning and Risk Modeling*. 2011

AAAI symposium — *Artificial Intelligence and Human-Robot Interaction*. 2015

AAAI symposium — *Knowledge, Skill, and Behavior Transfer in Autonomous Robots*. 2014

AAAI symposium — *Learning from Observation of Humans*. 2017

AAAI workshop — *Lifelong Learning from Sensorimotor Experience*. 2011

AAMAS workshop — *Adaptive and Learning Agents*. 2008, 2012

AAMAS *Doctoral Consortium*. 2016

AAMAS workshop — *Human-Agent Interaction Design and Models*. 2015

AAMAS workshop — *Issues with Deployment of Emerging Agent-based Systems*. 2015

AAMAS workshop — *International Joint Workshop on Optimization in Multi-Agent Systems and Distributed Constraint Reasoning*. 2014, 2015, 2016

ECML workshop — *European Workshop on Reinforcement Learning*. 2011

Evolutionary and Reinforcement Learning for Autonomous Robot Systems. 2011, 2013

ICML workshop — *Structural Knowledge Transfer for Machine Learning*. 2006
IEEE SSCI Doctoral Consortium. 2014
IJCAI workshop — *Distributed Constraint Reasoning*. 2013
IJCAI workshop — *Human-Agent Interaction Design and Models*. 2016
IJCAI workshop — *Interactions with Mixed Agent Types*. 2016
IJCAI workshop — *Interactive Machine Learning*. 2016
IJCAI workshop — *Quantitative Risk Analysis for Security Applications*. 2009
IROS workshop — *AI-based Robotics*. 2014
International Symposium on Distributed Autonomous Robotic Systems. 2012
Robocup Symposium. 2013
UbiComp workshop — *Smart Health Systems and Applications*. 2014

Proposal Reviewing

Netherlands Organization for Scientific Research: Research project proposal. 2014
Research Foundation Flanders: Research project proposal. 2013
NSF Panel. 2010, 2014, 2015, 2016, 2017, 2018
NSF Proposal Review. 2015

Consulting

2016–2017 Pullman Police Department: Drone Program
2015–2016 2nd Sight BioScience
2016 ELECTRI International: Applications for Unmanned Aerial Vehicles (UAVs) in Electric Utility Construction

Community Service & Outreach

2017 Talk to the Seattle Robotics Society
2013-14 Mentor for Pullman high school's Imagine Tomorrow team
2013 Mentor for local FIRST Robotics team, Sci-Borgs
2013 Lectured to undergraduate students attending REU programs at WSU in CS, EE, ME, MSE, and Physics
2013, 2014 Judge for Imagine Tomorrow, a high school renewable energy contest
2012 Taught a lecture/lab to incoming minority and female students interested in science and engineering enrolled in Lafayette's *Summer Program to Advance Leadership* program
2012 Lectured to low-income high school students in the *Princeton University Preparatory Program*
2011, 2012 Taught lectures/labs to at-risk inner city middle school students in the *Higher Achievement* program
2010 Supervised AJ Piergiovanni: Easton high school senior studying Bayesian networks
2009 Lectured at the Port of Los Angeles High School
2007 Assisted running labs in a senior physics course and updating the school computer lab at San Juan Diego Catholic High School in Austin, TX
2004 Assisted in First Bytes lab, a summer program for high school women at UT-Austin

University/College Service: Washington State University

2017 Presented twice to high school students and parents on Future Cougar Day
2015 Presented twice at the spring Junior Preview
2015 Internal reviewer for pre-proposals to the Murdock Science or Engineering Research Equipment grant
2014, 2015 Presented at VCEA Week of Welcome, representing EECS
2014 Presented to four groups of accepted high school students at Destination WSU

2013–2017 Faculty advisor for the school's Robotics Club
2013–2017 Faculty advisor for the school's RoboSub Club of the Palouse
2013 Presented to high school students for three Fall Preview recruiting events
2013 Assisted with presentation to Highline Community College

Department Service: Washington State University

F15-S16,S17-F17. Senior design group: Gamification in a Classroom Setting
F14-S15, F15-S16, F16-S17. Senior design group: Autonomous wheelchair navigation
S14-F14. Senior design group: Using UAVs for patrolling power lines
2017 Advisor for Restore Robotics. 4th place in WSU Business Plan Competition, $5,500 won.
2015 Judge at the ACM's annual Hackathon event
2014–2016 Faculty Search Committee: Machine Learning
2014 Lectured to WSU's ACM club on artificial intelligence and machine learning
2014 Faculty Search Committee: Software Engineering
2013-2014 Faculty Search Committee: Machine Learning
2013 Preliminary Exam Committee: Kyle Feuz
2013 QE Committee Chair: Daniel Olivares
2013 QE Committee: Jennifer Williams

College Service: Lafayette College

2012 Faculty advisor for the college's Robotics Club
2011-2012 Information Technology and Library Committee: Natural Science Representative
2012 Search Committee: Visitor in Computer Science
2011, 2012 Taught "mini courses" to accepted students during Experience Lafayette Day
2011, 2012 On panels for Lafayette's Center for the Integration of Teaching, Learning, and Scholarship

Departmental Service: Lafayette College

2012 Ran the Annual Robotics Competition
2011–2012 Academic advisor for all 3rd year computer science majors
2011–2012 Faculty advisor for the local chapter of the ACM
2011 Assisted with the Annual Robotics Competition

Departmental Service: UT-Austin

2007–2008 Organized the Transfer Learning Reading Group
2007 Graduate Student Faculty Recruiting Committee (selected by faculty for position)
2006 Computer Sciences Space Committee (volunteer position)
2004–2005 Graduate Representative Association of Computer Sciences (elected position)
2003–2004 Founded and organized the Reinforcement Learning Reading Group