

A. Rupam Mahmood

CURRICULUM VITAE

✉ armahmood@ualberta.ca

🌐 www.armahmood.com

Objective Developing a computational and scientific understanding of general-purpose goal-driven systems by building them with physical robots

Employment

- 2019 - Present **Assistant Professor**, *Computing Science, University of Alberta*.
2018 - 2019 **Lead, AI Research**, *Kindred Inc.*, www.kindred.ai, Toronto, Canada.
2017 **Research Scientist**, *Kindred Inc.*, www.kindred.ai, Toronto, Canada.
2009-2016 **Teaching & Research Assistant**, *University of Alberta*, Edmonton, Canada.
2006 - 2008 **Quantitative Software Developer**, *Stochastic Logic Ltd.*, Bangladesh.

Education

- 2017 **Ph.D. in Statistical Machine Learning**, *University of Alberta*, Canada.
Thesis Incremental Off-policy Reinforcement Learning Algorithms
Advisor Professor Richard S. Sutton
- 2010 **M.Sc. in Computing Science**, *University of Alberta*, Canada.
Thesis Automatic Step-size Adaptation in Incremental Supervised Learning
Advisor Professor Richard S. Sutton
- 2006 **B.Sc. in Computer Science & Engineering**, *Bangladesh University of Engineering and Technology*, Bangladesh.
Thesis Designing Neural Networks using Evolutionary Algorithms
Advisor Professor Md. Monirul Islam

Publications

Refereed Journal Articles (6 total)

- TMLR-2023 Lan, Q., Pan, Y., Luo, J., **Mahmood, A. R.** (2023). Memory-efficient reinforcement learning with value-based knowledge consolidation. *Transaction of Machine Learning Research*.
- JMLR-2022 Chan, A., Silva, H., Lim, S., Kozuno, T., **Mahmood, A. R.**, White, M. (2022). On generalized Bellman equations and temporal-difference learning. *Journal of Machine Learning Research*.

- IEEE RA-L-2020 Limoyo, O., Chan, B., Maric, F., Wagstaff, B., **Mahmood, A. R.**, Kelly, J. (2020). Robust generative latent dynamics via novelty detection. *IEEE Robotics and Automation Letters* 5(4): 6654–6661.
- JMLR-2018 Yu, H., **Mahmood, A. R.**, Sutton, R. S. (2018). On generalized Bellman equations and temporal-difference learning. *Journal of Machine Learning Research* 19(48):1–49.
- JMLR-2016 Sutton, R. S., **Mahmood, A. R.**, White, M. (2016). An emphatic approach to the problem of off-policy temporal-difference learning. *Journal of Machine Learning Research* 17(73):1–29.
- JMLR-2016 van Seijen, H., **Mahmood, A. R.**, Pilarski, P. M., Machado, M. C., Sutton, R. S. (2016). True online temporal-difference learning. *Journal of Machine Learning Research* 17(1):5057–5096.
- [Refereed Conference Articles \(20 total\)](#)
- ICML-2023 Che, F., Vasan, G., **Mahmood, A. R.** (2023). Correcting discount-factor mismatch in on-policy policy gradient methods. In *Proceedings of the 40th International Conference on Machine Learning*.
- UAI-2023 He, J., Che, F., Wan, Y., **Mahmood, A. R.** (2023). Loosely Consistent emphatic temporal-difference learning. In *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence*.
- IROS-2023 Karimi, A., Jin, J., Luo, J., **Mahmood, A. R.**, Jagersand, M., Tosatto, S. (2023). Variable-decision frequency option critic. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- IJCNN-2023 Farrahi, H., **Mahmood, A. R.** (2023). Reducing the cost of cycle-time tuning for real-world policy optimization. In *Proceedings of the 2023 International Joint Conference on Neural Networks* (Accepted).
- ICRA-2023 Wang, Y.^{*}, Vasan, G.^{*}, **Mahmood, A. R.** (2023). Real-time reinforcement learning for vision-based robotics utilizing local and remote computers. In *Proceedings of the 2023 International Conference on Robotics and Automation*.
- ICML-2022 Tosatto, S., Patterson, A., White, M., **Mahmood, A. R.** (2022). A temporal-difference approach to policy gradient estimation. In *Proceedings of the 39th International Conference on Machine Learning*.
- ICRA-2022 Yuan, Y., **Mahmood, A. R.** (2022). Asynchronous reinforcement learning for real-time control of physical robots. In *Proceedings of the 2022 International Conference on Robotics and Automation*.
- AISTATS-2022 Lan, Q., Tosatto, S., Farrahi, H., **Mahmood, A. R.** (2022). Model-free policy learning with reward gradients. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*.

- AISTATS-2022 Garg, S., Tosatto, S., Pan, Y., White, M., **Mahmood, A. R.** (2022). An alternate policy gradient estimator for softmax policies. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*.
- ICRA-2021 Przystupa, M., Dehghan, M., Jagersand, M., **Mahmood, A. R.** (2021). Analyzing neural Jacobian methods in applications of visual servoing and kinematic control. In *IEEE International Conference on Robotics and Automation*.
- IROS-2020 Limoyo, O., Chan, B., Maric, F., Wagstaff, B., **Mahmood, A. R.**, Kelly, J. (2020). Robust generative latent dynamics via novelty detection. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- IJCAI-2019 Korenkevych, D., **Mahmood, A. R.**, Vasan, G., Bergstra, J. (2019). Autoregressive policies for continuous control deep reinforcement learning. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*.
- CoRL-2018 **Mahmood, A. R.**, Korenkevych, D., Vasan, G., Ma, W., Bergstra, J. (2018). Benchmarking reinforcement learning algorithms on real-world robots. In *Proceedings of the 2nd Annual Conference on Robot Learning*.
- IROS-2018 **Mahmood, A. R.**, Korenkevych, D., Komer, B. J., Bergstra, J. (2018). Setting up a reinforcement learning task with a real-world robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- CAI-2017 Yu, H., **Mahmood, A. R.**, Sutton, R. S. (2017). On Generalized Bellman Equations and Temporal-Difference Learning. In *Proceedings of the 30th Canadian Conference on Artificial Intelligence*, Edmonton, Canada.
- UAI-2015 **Mahmood, A. R.**, Sutton, R. S. (2015). Off-policy learning based on weighted importance sampling with linear computational complexity. In *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence*, Amsterdam, Netherlands.
- NeurIPS-2014 **Mahmood, A. R.**, van Hasselt, H., Sutton, R. S. (2014). Weighted importance sampling for off-policy learning with linear function approximation. *Advances in Neural Information Processing Systems 27*, Montreal, Canada.
- UAI-2014 van Hasselt, H., **Mahmood, A. R.**, Sutton, R. S. (2014). Off-policy TD(λ) with a true online equivalence. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, Quebec City, Canada.
- ICML-2014 Sutton, R. S., **Mahmood, A. R.**, Precup, D., van Hasselt, H. (2014). A new $Q(\lambda)$ with interim forward view and Monte Carlo equivalence. In *Proceedings of the 31st International Conference on Machine Learning*, Beijing, China.
- ICASSP-2012 **Mahmood, A. R.**, Sutton, R. S., Degris, T., Pilarski, P. M. (2012). Tuning-free step-size adaptation. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Kyoto, Japan.

Teaching Experience

- Win. 2020-22 **CMPUT 365/397:** Reinforcement Learning, University of Alberta.
Fall 2022
- Win. 2022-23 **CMPUT 340:** Numerical Methods, University of Alberta.
- Sum. 2020-22 **Reinforcement Learning**, Next AI.
- Fall 2021 **CMPUT 653:** Real-Time Policy Learning, University of Alberta.
- Fall 2020 **CMPUT 653:** Deep Policy Gradient Methods, University of Alberta.
- Fall 2019 **CMPUT 652:** Reinforcement Learning with Robots, University of Alberta.

Honors, Awards & Distinctions

- 2023 **Notable Area Chair Award**, International Conference on Learning Representations (ICLR).
- 2022 **Top Reviewer Award**, Advances in Neural Information Processing Systems (NeurIPS).
- 2020-Present **Faculty**, NextAI, Toronto, Canada.
- 2019-Present **Canada CIFAR AI Chair**, Canadian Institute for Advanced Research
- 2019-Present **Fellow**, Alberta Machine Intelligence Institute (Amii), Edmonton, Canada
- 2019-Present **Scientific Advisor**, Kindred Inc., Toronto, Canada
- 2014 **Computing Science GPA Award**, University of Alberta
- 2010 **Provost Doctoral Entrance Award**, University of Alberta
- 2009 **M.Sc. Academic Achievement Award**, University of Alberta
- 2004 **University Dean's List Scholarship**, Bangladesh Univ. of Engg. & Tech.
- 2001 **University Merit Scholarship**, Bangladesh Univ. of Engg. & Tech.

Services - Organizing

- ICLR **Area Chair:** International Conference on Learning Representations, 2023
- ICML **Co-Chair:** Reinforcement Learning for Real Life Workshop at the International Conference on Machine Learning, 2021
- ICRA **Session-Chair:** International Conference on Robotics and Automation, 2021
- CRV **Session-Chair:** Conference on Robots and Vision, 2021
- IJCAI **Senior Program Committee Member:** International Joint Conference of Artificial Intelligence, 2019-2021
- IROS **Associate Editor:** IEEE/RSJ International Conference on Intelligent Robots and Systems, 2020

Services - Reviewing

- MLJ Machine Learning Journal, 2020, 2021

ICLR	International Conference on Learning Representations, 2021-2022
ICML	International Conference on Machine Learning, 2013 - 2021
NeurIPS	Neural Information Processing Systems, 2015 - 2022
AISTATS	International Conference on Artificial Intelligence and Statistics, 2017-2022
AAAI	AAAI Conference on Artificial Intelligence, 2015, 2017-2019
Auro	Autonomous Robots, 2019
CoRL	The Conference on Robot Learning, 2017
IJCAI	International Joint Conference of Artificial Intelligence, 2016
AC	IEEE Transaction on Automatic Control, 2011, 2013
SMC	IEEE Transactions on Systems, Man and Cybernetics: Systems, 2015
JCSS	Journal of Computer and System Sciences, 2014
JMLR	Journal of Machine Learning Research, 2010
IEEE	Proceedings of the IEEE, 2013

Invited Talks

Mar 2021	<i>End-to-end continual RL with robots</i> , 18th Conference on Robots and Vision (symposia speaker).
Dec 2020	<i>Frontiers of Real-time Robot RL</i> , Huawei Noah's Ark Workshop on Reinforcement Learning, Edmonton.
Aug 2020	<i>Reinforcement Learning with Robots</i> , DL-RL Summer School, Montreal.
Nov 2019	<i>Reinforcement Learning with Physical Robots</i> , Taiwan AI Centers Collaboration Workshop, Edmonton.
Oct 2019	<i>General Purpose Minds for Robots</i> , Huawei R&D Strategic Partnership Workshop, Vancouver.
Jul 2019	<i>Reinforcement Learning with Robots</i> , DL-RL Summer School, Edmonton.
Feb 2019	<i>Overcoming the Challenges of Learning with Physical Robots</i> , Queen's University Smith School, Toronto.
Dec 2018	<i>The Challenges of Real-World Robot Learning</i> , Vector Institute Seminar, Toronto.
May 2018	<i>Setting up a Reinforcement Learning Task with a Real-World Robot</i> , Toronto DL Seminars, Toronto.
Apr 2018	<i>Making Minds for Robots with Reinforcement Learning</i> , Toronto Synthetic Intelligence Forum, Toronto.

Personal Information

Residency	Citizen of Canada & Bangladesh
Hobbies	Cycling; hiking; conversing passionately about the future of AI and humanity