

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

Note Acceptance rates in green; trainees in blue; star (★) for equal contributions.

Refereed Journal and Conference Articles (35 total)

- ICML-2024a **Elsayed, M., Farrahi, H., Dangel, F., Mahmood, A. R.** (2024). Revisiting
(27%) scalable Hessian diagonal approximations for applications in reinforcement learning. In *The International Conference on Machine Learning*.
- ICML-2024b **Che, F., Xiao, C., Mei, J., Dai, B., Gummadi, R., Ramirez, O. A., Har-**
(27%) **ris, C. K., Mahmood, A. R., Schuurmans, D.** (2024). Target networks and over-parameterization stabilize off-policy bootstrapping with function approximation. In *The International Conference on Machine Learning*.

- Under Revision-2024 [Dohare, S.](#), Hernandez-Garcia, J.F., [Lan, Q.](#), Rahman, P., **Mahmood, A. R.**, Sutton, R. S. (2024) Loss of plasticity in deep continual learning. *Conditionally Accepted by a Leading Journal*, [link](#).
- ICLR-2024a (31%) [Elsayed, M.](#), **Mahmood, A. R.** (2024). Addressing loss of plasticity and catastrophic forgetting in continual learning. In *Proceedings of International Conference on Learning Representations*.
- ICLR-2024b (31%) Ishfaq, H.*, [Lan, Q.*](#), Xu, P., **Mahmood, A. R.**, Precup, D., Anandkumar, A., Azizzadenesheli, K. (2024). Provable and practical: efficient exploration in reinforcement learning via Langevin Monte Carlo. In *Proceedings of International Conference on Learning Representations*.
- AAMAS-2024 (20%) Grooten, B., Tomilin, T., [Vasan, G.](#), Taylor, M. E., **Mahmood, A. R.**, Fang, M., Pechenizkiy, M. Mocanu, D.C. (2024). MaDi: Learning to mask distractions for generalization in visual deep reinforcement learning. In *Proceedings of the 23rd Autonomous Agents and Multiagent Systems*.
- RLC-2024a [Elsayed, M.](#), [Lan, Q.](#), Lyle, C., **Mahmood, A. R.** (2024). Weight clipping for deep continual and reinforcement learning. In *the First Reinforcement Learning Conference*.
- RLC-2024b [Vasan, G.](#), [Wang, Y.](#), [Shahriar, F.](#), Bergstra, J., Jagersand, M., **Mahmood, A. R.**(2024). Revisiting constant negative rewards for goal-reaching tasks in robot learning. In *the First Reinforcement Learning Conference*.
- RLC-2024c Ishfaq, H., Tan, Y., Yang, Y., [Lan, Q.](#), Lu, J., **Mahmood, A. R.**, Precup, D., Xu, P. (2024). More efficient randomized exploration for reinforcement learning via approximate sampling. In *the First Reinforcement Learning Conference*.
- RLC-2024d [Lan, Q.](#), **Mahmood, A. R.**, Yan, S., Xu, Z. (2024). Learning to optimize for reinforcement learning. In *the First Reinforcement Learning Conference*.
- ICML-2023 (28%) [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*.
- 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*.
- UAI-2023 (31%) [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 (43%) [Karimi, A.](#), Jin, J., Luo, J., **Mahmood, A. R.**, Jagersand, M., [Tosatto, S.](#) (2023). Loosely consistent emphatic temporal-difference learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.

- ICRA-2023 (43%) 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 International Conference on Robotics and Automation*.
- 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*.
- JMLR-2022 (<20%) 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*.
- ICML-2022 (22%) 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 (43%) 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-2022a (29%) 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*.
- AISTATS-2022b (29%) 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*.
- ICRA-2021 (49%) 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*.
- RA-L-2020 IROS-2020 (47%) 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.
- IJCAI-2019 (18%) 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*.
- JMLR-2018 (<20%) 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.
- 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). (46%) 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.
- JMLR-2016a Sutton, R. S., **Mahmood, A. R.**, White, M. (2016). (<20%) An emphatic approach to the problem of off-policy temporal-difference learning. *Journal of Machine Learning Research* 17(73):1–29.
- JMLR-2016b van Seijen, H., **Mahmood, A. R.**, Pilarski, P. M., Machado, M. C., Sutton, R. S. (2016). (<20%) True online temporal-difference learning. *Journal of Machine Learning Research* 17(1):5057–5096.
- UAI-2015 **Mahmood, A. R.**, Sutton, R. S. (2015). (34%) 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). (25%) 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). (32%) 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). (37%) A new Q(λ) 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). (49%) Tuning-free step-size adaptation. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Kyoto, Japan.

Teaching Experience

- Win. 2022–24 **CMPUT 340:** Numerical Methods, University of Alberta.
Fall 2023
- Win. 2020–22 **CMPUT 365/397:** Reinforcement Learning, University of Alberta.
Fall 2022
- Sum. **Reinforcement Learning**, Next AI.
2020–22
- Fall 2021,23 **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** (top 10%), ICLR.
- 2022–23 **Top Reviewer** (top 11%), NeurIPS.
- 2022 **Highlighted Reviewer** (top 9%), ICLR.
- 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

- Area Chair ICLR 2023–2024, NeurIPS 2024
- Workshop ICML Reinforcement Learning for Real Life Workshop 2021
- Co-Chair
- Session-Chair ICRA 2021, Conference on Robots and Vision 2021
- SPC Senior Program Committee Member: IJCAI 2019–2021
- Associate Editor IROS 2020

Services: Reviewing

- Conferences NeurIPS, ICLR, ICML, AISTATS, AAAI, CoRL, IJCAI
- Journals Editorial Board at the Artificial Intelligence Journal (AIJ) 2024–Present
- Transactions on Machine Learning Research (TMLR) 2023
- IEEE Transactions on Robotics (T-RO) 2023
- Machine Learning Journal (MLJ), 2020, 2021
- Autonomous Robots (Auro), 2019
- IEEE Transaction on Automatic Control (AC), 2011, 2013
- IEEE Transactions on Systems, Man and Cybernetics: Systems (SMC), 2015
- Journal of Computer and System Sciences (JCSS), 2014
- Journal of Machine Learning Research (JMLR), 2010
- Proceedings of the IEEE, 2013

Grants NSERC Alliance 2023
Swiss National Science Foundation 2021

Services: University of Alberta

Amii Resource Allocation Panel 2023–present
Computing Science Faculty Recruiting Committee 2023, 2024
Computing Science Curriculum Committee 2023–present
Computing Science Graduate Admissions Strategy Committee 2021–22

Invited Talks

- Jul 2024 *Reinforcement Learning 1*, DL-RL Summer School, Toronto, Canada.
- Jun 2024 *Continual Learning under Resource Constraint*, Bayes-Duality Workshop 2024, Tokyo, Japan.
- May 2024 *Continual Robot Learning*, UpperBound (MasterClass in Continual Learning), Edmonton, Canada.
- Mar 2021 *End-to-end Continual RL with Robots*, 18th Conference on Robots and Vision (symposia speaker), Online.
- Dec 2020 *Frontiers of Real-time Robot RL*, Huawei Noah's Ark Workshop on Reinforcement Learning, Edmonton, Canada.
- Aug 2020 *Reinforcement Learning with Robots*, DL-RL Summer School, Montreal.
- Nov 2019 *Reinforcement Learning with Physical Robots*, Taiwan AI Centers Collaboration Workshop, Edmonton, Canada.
- Oct 2019 *General Purpose Minds for Robots*, Huawei R&D Strategic Partnership Workshop, Vancouver, Canada.
- Jul 2019 *Reinforcement Learning with Robots*, DL-RL Summer School, Edmonton.
- Feb 2019 *Overcoming the Challenges of Learning with Physical Robots*, Queen's University Smith School, Toronto, Canada.
- Dec 2018 *The Challenges of Real-World Robot Learning*, Vector Institute Seminar, Toronto, Canada.
- May 2018 *Setting up a Reinforcement Learning Task with a Real-World Robot*, Toronto DL Seminars, Toronto, Canada.
- Apr 2018 *Making Minds for Robots with Reinforcement Learning*, Toronto Synthetic Intelligence Forum, Toronto, Canada.

Research Grants

2024–2026 National Research Council Canada Grant: \$247,500 for two years.

- 2022–present Digital Research Alliance of Canada Resource Allocation Grant:
~\$70,000 per year in computational resources in clusters.
- 2021–2026 NSERC Discovery Grant: \$120,000 for five years.
- 2020–2024 Canada CIFAR AI Chair: \$645,100 for five years.
- 2020–2022 Huawei UofA’s joint UAHJIC fund (Co-PI): \$180,000 for two years.
- 2020–2021 Huawei Noah’s Ark fund: \$110,000 for one year.

Supervision

- Post-Doc Samuele Tosatto, 2021–22
- PhD Students Qingfeng Lan, 2020–Present
Gautham Vasan, 2020–Present
Fengdi Che, 2021–Present
Blanca Miller, 2021–Present
Homayoon Farrahi, 2021–Present
Shibhansh Dohare, 2021–Present (co-supervised with Richard S. Sutton)
Mohamed Elsayed, 2022–Present
- MSc Students Fahim Shahriar, 2022–Present
Alireza Azimi, 2024–Present
Haruto Tanaka, 2024–Present
Jiamin He, 2022-2023
Yan Wang, 2021-2023
Amirmohammad Karimi, 2020-2023 (co-supervised with Martin Jagersand)
Mohamed Elsayed, 2020-2022
Shivam Garg, 2020-2021 (co-supervised with Martha White)
Yufeng Yuan, 2020-2021
Homayoon Farrahi, 2020-2021
Shibhansh Dohare, 2020