Martha White

Curriculum vitae

Associate Professor
Department of Computing Science
University of Alberta
Edmonton, Alberta

whitem@ualberta.ca
marthawhite.ca
(587) 590-9940

RESEARCH INTERESTS **Reinforcement learning**, representation learning, time series prediction.

ACADEMIC POSITIONS

Associate Professor 2020-present
Department of Computing Science, Faculty of Science, University of Alberta

Assistant Professor 2017-2020
Department of Computing Science, Faculty of Science, University of Alberta

Assistant Professor 2015-2017
School of Informatics and Computing, Indiana University Bloomington

EDUCATION **The University of Alberta**, Edmonton, Alberta, Canada

Ph.D., Computing Science, January 2010-December 2014
Supervisors: Professor Michael Bowling and Professor Dale Schuurmans

M.Sc. Computing Science, September 2008-December 2009

B.S.c, Honors Mathematics with First Class Honors, 2008

B.S.c, Honors Computing Science with First Class Honors, 2008

PUBLICATIONS **Refereed Journal and Conference Articles**

[1] Exploiting Action Impact Regularity and Exogenous State Variables for Offline Reinforcement Learning. V. Liu, James Wright and M. White. *Journal of Artificial Intelligence Research (JAIR)*, 2023.

[2] Off-Policy Actor-Critic with Emphatic Weightings. E. Graves, E. Imani, R. Kumaraswamy and M. White. *Journal of Machine Learning Research (JMLR)*, 2023.

[3] Greedy Actor-Critic: A New Conditional Cross-Entropy Method for Policy Improvement. S. Neumann, S. Lim, A. G. Joseph, Y. Pan, A. White and M White. *International Conference on Representation Learning (ICLR)*, 2023.

[4] The In-Sample Softmax for Offline Reinforcement Learning. C. Xiao, H. Wang, Y. Pan, A. White and M. White. *International Conference on Representation Learning (ICLR)*, 2023.

[5] Asymptotically Unbiased Off-Policy Policy Evaluation when Reusing Old Data in Nonstationary Environments. V. Liu, Y. Chandak, P. Thomas and M. White. *International Conference on AI and Statistics (AISTATS)*, 2023.

[6] Representation Alignment in Neural Networks. E. Imani, W. Hu and M. White. *Transactions on Machine Learning Research (TMLR)*, 2022.

[7] Robust Losses for Learning Value Functions. A. Patterson, V. Liao and M. White. *Transactions on Pattern Analysis and Machine Learning (TPAMI)*, 2022.

[8] Greedification Operators for Policy Optimization: Investigating Forward and Reverse KL Divergences. A. Chan, H. Silva, S. Lim, T. Kozuno, A. R. Mahmood, M. White.*Journal of Machine Learning Research (JMLR)*, 2022.

[9] No More Pesky Hyperparameters: Offline Hyperparameter Tuning for Reinforcement Learning. H. Wang, A. Sakhadeo, A. White, J. Bell, V. Liu, X. Zhao, P. Liu, T. Kozuno, A. Fyshe, M. White. *Transactions on Machine Learning Research (TMLR)*, 2022.

[10] Understanding and Mitigating the Limitations of Prioritized Replay. J. Mei, Y. Pan, A. Farahmand, H. Yao, M. White. *Uncertainty in Artificial Intelligence (UAI)*, 2022.

[11] A Temporal-Difference Approach to Policy Gradient Estimation. S. Tosatto, A. Patterson, M. White and A. R. Mahmood. *International Conference on Machine Learning (ICML)*, 2022.

[12] A Generalized Projected Bellman Error for Off-policy Value Estimation in Reinforcement Learning. A. Patterson, A. White and M.White. *Journal of Machine Learning Research (JMLR)*, 2022.

[13] An Alternate Policy Gradient Estimator for Softmax Policies. S. Garg, S. Tosatto, Y. Pan, M. White and A. R. Mahmood. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.

[14] Resonance in Weight Space: Covariate Shift Can Drive Divergence of SGD with Momentum. K. Banman, L. Peet-Pare, N. Hegde, A. Fyshe and M. White. *International Conference on Learning Representations (ICLR)*, 2022.

[15] Sim2Real in Robotics and Automation: Applications and Challenges. S. Hofer, K. Bekris, A. Handa, J.C. Gamboa, M. Mozifian, F. Golemo, C. Atkeson, D. Fox, K. Goldberg, J. Leonard, C, Karen Liu, J. Peters, S. Song, P. Welinder, Peter and M. White. *IEEE Transactions on Automation Science and Engineering*, 2021.

[16] Continual Auxiliary Task Learning. M. McLeod, C. Lo, M. Schlegel, A. Jacobsen, R. Kumaraswamy, M. White and A. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

[17] Structural Credit Assignment in Neural Networks using Reinforcement Learning. D. Gupta, G. Mihucz, M. Schlegel, J. Kostas, P. Thomas, M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

[18] Fuzzy Tiling Activations: A Simple Approach to Learning Sparse Representations Online. Y. Pan,K. Banman, and M. White. *International Conference on Learning Representations (ICLR)*, 2021.

[19] General Value Function Networks. M. Schlegel, A. Jacobsen, Z. Abbas, A. Patterson, A. White and M.White. *Journal of AI Research (JAIR)*, 2021.

[20] From Language to Language-ish: How Brain-Like is an LSTM?s Representation of Atypical Language Stimuli? M. Hashemzadeh, G. Kaufeld, M. White, A. E. Martin, A. Fyshe. *Findings of ACL: Empirical Methods in Natural Language Processing (EMNLP)*. 2020.

[21] Towards Safe Policy Improvement for Non-Stationary MDPs. Y. Chandak, S. Jordan, G. Theocharous, M. White, P. S. Thomas. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[22] An implicit function learning approach for parametric modal regression. Y. Pan, E. Imani, A. Farahmand and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[23] Adapting Behaviour via Intrinsic Reward: A Survey and Empirical Study. C. Linke, N. M. Ady, M. White, T. Degris, and A. White. *Journal of AI Research (JAIR)*, 2020.

[24] Gradient Temporal-Difference Learning with Regularized Corrections. S. Ghiassian, A. Patterson, S. Garg, D. Gupta, A. White and M. White. *International Conference on Machine Learning (ICML)*, 2020.

[25] Selective Dyna-style Planning Under Limited Model Capacity. Z. Abbas, S. Sokota, E. Talvitie and M. White. *International Conference on Machine Learning (ICML)*, 2020.

[26] Optimizing for the Future in Non-Stationary MDPs. Y. Chandak, G. Theocharous, S. Shankar, M. White, S. Mahadevan, P. S. Thomas. *International Conference on Machine Learning (ICML)*, 2020.

[27] Training Recurrent Neural Networks Online by Learning Explicit State Variables. S. Nath, V. Liu, A. Chan, A. White and M. White. *International Conference on Learning Representations (ICLR)*, 2020.

[28] Maxmin Q-learning: Controlling the Estimation Bias of Q-learning. Q. Lan, Y. Pan, A. Fyshe and M. White. *International Conference on Learning Representations (ICLR)*, 2020.

[29] Maximizing Information Gain in Partially Observable Environments via Prediction Rewards. Y. Satsangi, S. Lim, S. Whiteson, F. Oliehoek and M. White. *International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2020.

[30] Meta-Learning Representations for Continual Learning. K. Javed and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[31] Importance Resampling for Off-policy Prediction. M. Schlegel, W. Chung, D. Graves, J. Qian and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[32] Learning Macroscopic Brain Connectomes via Group-Sparse Factorization. F. Aminmansour, A. Patterson, L. Le, Y. Peng, D. Mitchell, F. Pestilli, C. Caiafa, R. Greiner and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[33] Planning with Expectation Models. Y. Wan, Z. Abbas, A. White, M. White and R. S. Sutton. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.

[34] Hill Climbing on Value Estimates for Search-control in Dyna. Y. Pan, H. Yao, A. Farahmand and M. White. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.

[35] Two-Timescale Networks for Nonlinear Value Function Approximation. W. Chung, S. Nath, A. Joseph and M. White. *International Conference on Learning Representations (ICLR)*, 2019.

[36] The Utility of Sparse Representations for Control in Reinforcement Learning. V. Liu, R. Kumaraswamy, L. Le, and M. White. *AAAI Conference on Artificial Intelligence (AAAI)*, 2019.

[37] Meta-descent for Online, Continual Prediction. A. Jacobsen, M. Schlegel, C. Linke, T. Degris, A. White and M. White. *AAAI Conference on Artificial Intelligence (AAAI)*, 2019.

[38] An Off-policy Policy Gradient Theorem Using Emphatic Weightings. E. Imani, E. Graves and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

[39] Supervised autoencoders: Improving generalization performance with unsupervised regularizers. L. Le, A. Patterson and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

[40] Context-dependent upper-confidence bounds for directed exploration. R. Kumaraswamy, M. Schlegel, A. White and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

[41] Improving Regression Performance with Distributional Losses. E. Imani and M. White. *International Conference on Machine Learning (ICML)*, 2018.

[42] Reinforcement Learning with Function-Valued Action Spaces for Partial Differential Equation Control. Y. Pan, A. Farahmand, M. White, S. Nabi, P. Grover, D. Nikovski. *International Conference on Machine Learning (ICML)*, 2018.

[43] Organizing experience: a deeper look at replay mechanisms for sample-based planning in continuous state domains. Y. Pan, Z. Abbas, A. White, A. Patterson, M. White. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.

[44] High-confidence error estimates for learned value functions. T. Sajed, W. Chung and M. White. *Uncertainty in Artificial Intelligence (UAI)*, 2018.

[45] Comparing Direct and Indirect Temporal-Difference Methods for Estimating the Variance of the Return. C. Sherstan, D. Ashley, B. Bennet, K. Young, A. White, M. White, R. Sutton. *Uncertainty in Artificial Intelligence (UAI)*, 2018.

[46] Multi-view Matrix Factorization for Linear Dynamical System Estimation. M. Karami, M. White, D. Schuurmans and C. Szepesvari. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[47] Unifying task specification in reinforcement learning. M. White. *International Conference on Machine Learning (ICML)*, 2017.

[48] Adapting kernel representations online using submodular maximization. M. Schlegel, Y. Pan and M. White. *International Conference on Machine Learning (ICML)*, 2017.

[49] Effective sketching methods for value function approximation. Y. Pan, E. Sadeqi Azer and Martha White. *International Conference on Uncertainty in AI (UAI)*, 2017.

[50] Learning sparse representations in reinforcement learning with sparse coding. L. Le, R. Kumaraswamy, and M. White. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.

[51] Accelerated Gradient Temporal Difference Learning. Y. Pan, A. White and M. White. *AAAI Conference on Artificial Intelligence (AAAI)*, 2017.

[52] Recovering true classifier performance in positive-unlabeled learning. S. Jain, M. White, P. Radivojac. *AAAI Conference on Artificial Intelligence (AAAI)*, 2017.

[53] Estimating the class prior and posterior from noisy positives and unlabeled data. S. Jain, M. White, P. Radivojac. *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.

[54] Investigating practical, linear temporal difference learning. A. White and M. White. *International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2016.

[55] A greedy approach to adapting the trace parameter for temporal difference learning. A. White and M. White. *International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2016.

[56] Incremental Truncated LSTD. C. Gehring, Y. Pan and M. White. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

[57] An emphatic approach to the problem of off-policy temporal-difference learning. R. Sutton, A.R. Mahmood, M. White. *Journal of Machine Learning Research*, 2016.

[58] Scalable Metric Learning for Co-embedding. F. Mirzazadeh, M. White, A. Gyorgy and D. Schuurmans. *In ECML PKDD*, 2015.

[59] Optimal Estimation of Multivariate ARMA Models. M. White, J. Wen, M. Bowling and D. Schuurmans. *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

[60] Partition Tree Weighting. J. Veness, M. White, M. Bowling, and A. Gyorgy. *Data Compression Conference*, 2013.

[61] Convex Multiview Subspace Learning. M. White, Y. Yu, X. Zhang, D. Schuurmans. *Advances in Neural Information Processing Systems (NeurIPS)*, 2012.

[62] Off-Policy Actor-Critic. T. Degris, M. White and R. S. Sutton. *International Conference on Machine Learning(ICML)*, 2012.

[63] Generalized Optimal Reverse Prediction. M. White and D. Schuurmans. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012.

[64] Convex Sparse Coding, Subspace Learning, and Semi-Supervised Extensions. X. Zhang, Y. Yu, M. White, R. Huang, and D. Schuurmans. *AAAI Conference on Artificial Intelligence (AAAI)*, 2011.

[65] Interval Estimation for Reinforcement-Learning Algorithms in Continuous-State Domains. M. White and A. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2010.

[66] Relaxed Clipping: A Global Training Method for Robust Regression and Classification. Y. Yu, M. Yang, L. Xu, M. White, D. Schuurmans. *Advances in Neural Information Processing Systems (NeurIPS)*, 2010.

[67] Optimal Reverse Prediction: A Unified Perspective on Supervised, Unsupervised and Semi-supervised Learning. L. Xu, M. White and D. Schuurmans. *International Conference on Machine Learning (ICML)*, 2009. **Honourable Mention for Best Paper**

[68] Learning a Value Analysis Tool For Agent Evaluation. M. White and M.l Bowling. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2009.

Theses

M. White. **Regularized factor models.** PhD thesis, University of Alberta. Received the Faculty of Science Doctoral Dissertation Award.

M. White. **A General Framework for Reducing Variance in Agent Evaluation**. Master's thesis, University of Alberta.

Research Grants

Killam Accelerator Research Award. 2021-2024
$225,000 in research funding for three years.

Canada CIFAR AI (CCAI) Chair. 2018-2023
$650,000 in research funding for five years.

Compute Canada Resource Allocation Grant. 2019-2022
$100,000 per year in additional computing resources on Compute Canada clusters.

NSERC Discovery Grant and Discovery Accelerator. 2018-2021
$120,000 + 150,000 total for three years. Sole PI. "Sparse representations for reinforcement learning."

NSERC CRD 2018-2021
$536,700 over 3 years. Sole PI. "Optimizing water treatment operation using reinforcement learning."

NSF CISE CRII grant. 2016-2018
$174,616 over two years. Sole PI. "Accelerated stochastic approximation for reinforcement learning."

Precision Health Initiative. 2016-2020
$60,000 per year (funding for two students), for four years. Joint with the School of Informatics and Computing and the School of Medicine.

Awards

Faculty of Science Research Award, University of Alberta, 2022

AIs 10 to Watch: The Future of AI, IEEE Intelligent Systems, 2020

University of Alberta Killam Accelerator Research Award, 2020

Outstanding Reviewer Award, International Conference on Learning Representations (ICLR), 2021

Canada CIFAR AI Chair, 2018

Champion of Inclusion, School of Informatics and Computing, Indiana University, 2016

Reviewing Award, International Conference on Machine Learning (ICML), 2015

Faculty of Science Doctoral Dissertation Award, 2015

Honourable Mention for Best Paper at the International Conference on Machine Learning, 2009

Invited Talks

[Keynote] *Using Logged Data to Calibrate Online Reinforcement Learning Agents.* RL Alpes Symposium, March 2023.

Developing Reinforcement Learning Agents that Learn Many Subtasks. Apple Machine Learning Research Seminar AI, March 2023.

[Keynote] *Planning with Models Based on Value Functions.* Reinforcement Learning and Decision-Making Conference (RLDM), Providence, USA, June 2022.

[Keynote] *Reinforcement Learning: the Next Big Thing in AI.* AIWeek, Edmonton, May 2022

Developing Reinforcement Learning Agents that Learn Many Subtasks. Waterloo AI Institute Seminar, February 2022; at the UC Irvine CS Seminar, February 2022; and at TechAide Montreal, April 2022.

Advances in Off-policy Value Estimation in Reinforcement Learning. MIT Computational Sensorimotor Learning Seminar, November 2021; and at London Machine Learning Meetup, April 2022.

Using Logged Data to Calibrate Online Reinforcement Learning Agents. SIAM Symposium on Data-Driven Decision Control for Complex Systems, July 2021.

A Generalized Objective for Off-policy Value Estimation in Reinforcement Learning. CAIMS Annual Meeting, June 2021.

Learning-to-learn with Recurrent Algorithms in Online Learning. ELLIS Workshop on Meta-learning in AI and Cognitive Science, March 2021.

Learning Representations for Reinforcement Learning. Deep Learning 2.0 Virtual Summit, REWORK, January 2021.

Generalizing the Projected Bellman Error Objective for Nonlinear Value Estimation. Deep Reinforcement Learning, Theory of Reinforcement Learning, Simons Institute, September 2020.

A New RNN Algorithm Using the Computational Inductive Bias of Span Independence. Inductive Biases, Invariances and Generalization in RL, ICML Workshop, July 2020.

An Off-policy Policy Gradient Theorem: A Tale About Weightings. Theoretical Foundations of Reinforcement Learning, ICML Workshop, July 2020.

Rethinking the Objective for Policy Optimization in Reinforcement Learning. CAIDA Seminar Series, University of British Columbia, June 2020.

Understanding Inductive Biases for Betrrl Agents. Beyond Tabula Rasa in Reinforcement Learning Workshop (BeTR-RL), ICLR Workshop, April 2020.

Model-Based Reinforcement Learning. Reinforcement Learning Summer School, Edmonton, August 2019.

Learning Representations for Continual Learning. Multi-task and Lifelong Learning Workshop, ICML Workshop, Long Beach, June 2019.

Using Auxiliary Variables to Train Recurrent Neural Networks. AI for Good Summer Lab, Montreal, June 2019.

The Utility of Sparse Representations for Control in RL. Deep Reinforcement Learning Workshop, NeurIPS Workshop, Montreal, December 2018.

Another Laundry List for Continual Learning. Continual Learning Workshop, NeurIPS Workshop, Montreal, December 2018.

Off-policy Learning. Reinforcement Learning Summer School, Toronto, August 2018.

Upper Confidence Bounds on Action-Values. Exploration in Reinforcement Learning, ICML Workshop, Sweden, July 2018.

An RNN Architecture using Value Functions. Credit assignment in Deep Learning and Deep Reinforcement Learning, ICML Workshop, Sweden, July 2018.

Predictive representations. Learning in Machines and Humans, Bloomington, USA, May 2018.

Planning in reinforcement learning with learned models in Dyna. Generative Models for Reinforcement Learning, Data Learning and Inference Workshop, Canary Islands, April 2018.

General Value Function Networks. The Barbados Workshop on Reinforcement Learning, Feb. 2018.

Unifying task specification in reinforcement learning. Oxford (Department colloquium), Imperial College London and Google Deepmind London, July 2017.

Adapting kernel representations online using submodular maximization. Washington University in St. Louis, Machine Learning colloquium, March 2017.

Insights on learning representations with dictionary learning and autoencoders. University of Maryland, Computational Linguistics and Information Processing colloquium, November 2016.

Accelerated Gradient Temporal Difference Learning. Presented at University of Texas at Austin, November 2016 and then again at the Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM), December 2016.

Beyond experts and engineering: exploiting data for automated control. Presented at multiple universities in 2014-2015, including University of Texas at Austin, University of Rochester, Virginia Tech, University of Virginia, Worcester Polytechnic Institute, University of Iowa, University of Connecticut and Dartmouth College.

Generalized Optimal Reverse Prediction. Google New York, August 2012

Linear Off-Policy Actor Critic. The 7th Barbados Workshop on RL, April 2012

Learning a Value Analysis Tool For Agent Evaluation. International Joint Conference of Artificial Intelligence, July 2009 and MITACS, June 2009

Academic Service

Action Editor for JMLR 2022-present

Associate Editor for TPAMI 2020-present

Action Editor for TMLR 2022-present

Board Member for the Int. Conf. on Learning Representations (ICLR) 2020-present

Program Co-Chair for ICLR with Shakir Mohamed, Kyunghyun Cho, Dawn Song 2019-2020
First ML conference to go virtual for 2020; many later conferences built on strategies we introduced

Senior Area Chair 2021-present
Advances in Neural Information Processing Systems (NeurIPS), 2022
International Conference on Learning Representations (ICLR), 2022

Area Chair 2018-present
International Conference on Machine Learning (ICML), 2018-2021
Advances in Neural Information Processing Systems (NeurIPS), 2018, 2020-2021
International Conference on Learning Representations (ICLR), 2018, 2019, 2021
AAAI Conference on Artificial Intelligence (AAAI), 2018, 2019

Program Committee Member (Reviewer) 2009-present
Conferences: AAAI, AISTATS, AAMAS, ICLR, ICML, IJCAI, NeurIPS, UAI, CORL
Journals: PNAS 2019-2020
Nature 2019, Nature Machine Intelligence 2022
Royal Statistical Society 2022
Journal of Machine Learning Research, 2014-2017,2021
Machine Learning Journal, 2014, 2018-2020

Journal of Artificial Intelligence Research, 2014, 2016-2018,2021,2022
Artificial Intelligence Journal, 2014,2021
Journal of Autonomous Agents and Multi-agent Systems, 2016
Transactions on Image Processing, 2014
IEEE Transactions on Neural Networks and Learning Systems, 2014
Stochastic Environmental Research and Risk Assessment, 2017

Workshop and Tutorial Organization 2019-present
Deep Reinforcement Learning Workshop, NeurIPS, 2021, 2022
Self-Supervision for Reinforcement Learning, ICLR 2021
Policy Optimization in Reinforcement Learning Tutorial, NeurIPS, 2020
NeurIPS Optimization in RL Workshop, 2019
ICML Exploration in RL Workshop, 2019, 2020
RLDM Curiosity and Intrinsic Motivation Workshop, 2019

NSERC Reviewer 2018-present
Discovery Grants, MITACS grants

NSF panel member 2015, 2016
Reviewed IIS: Robust Intelligence (RI) proposals on Machine Learning

University Service

University Service at the University of Alberta

Faculty Evaluation Committee, for the Faculty of Science 2020-2022
Data Science Program Committee Member, joint between MSS and CS 2020-2022
AI4Society Signature Area 2019-present
Help manage the cross-department Statistical Machine Learning (SML) program 2018-present
Service roles within Amii: Fellows Membership Committee to nominate members and CCAI Chairs (2019-2023), Budgeting Committee (2018), Advisory Committee to CEO (2018-2020)

Departmental Service at the University of Alberta

AI Curriculum Committee 2020-present
- Introduced AI Certificate, including course redesign
- Developing an MSc in AI

Faculty Recruiting Committee 2017-2020, 2023
Data Science Curriculum design 2019-2020
- Designed and continue to teach the first two courses in a stream of three

University Service at Indiana University

Faculty Hiring Committee for Statistics (consulting role) 2015-2016
Panel for new Phd Students 2016
Panel for Women in Computing 2015

Departmental Service at Indiana University

Faculty Affairs Committee 2016-present
Undergraduate Education Committee 2016-present
- Redesigned Artificial Intelligence specialization for undergrads
- Introduced two courses, Data Mining (B365) and Machine Learning (B455)

Faculty Hiring Committee 2015-2016

Supervision

Post-Doctoral Fellows

Scott Jordan 2022-present
Lingwei Zhu 2022-present
Tadashi Kozuno (now Research Scientist at Omron Sinic X) 2021-2022
Yash Satsangi (now Assistant Professor at Tilburg University) 2019-2020
Ajin George Joseph (now Assistant Professor at IIT Tirupati) 2018-2019
Omid Namaki (now Lead Data Scientist for ATB) 2017-2019

PhD students

Farzane Aminmansour 2020-present
Brett Daley 2022-present
Ehsan Imani 2019-present
Andrew Jacobsen 2019-present
Vincent Liu 2019-present
Prabhat Nagarajan 2021-present
Andrew Patterson 2018-present
Matthew Schlegel 2017-present
Han Wang 2020-present
Raksha Kumaraswamy (*Towards Sample-Efficient Control with Directed Exploration Under Linear Function Approximation*) 2016-2021
Lei Le (*Identifying Tractable Dictionary Learning Models for Representation Learning*) 2015-2019
Yangchen Pan (*Improving Sample Efficiency of Online Temporal Difference Learning*) 2015-2021

MSc students

Alvina Awwal 2022-present
Esraa Elelimy 2022-present
Anna Hakhverdyan 2023-present
Kamran Janjua 2022-present
Hugo Luis Silva 2021-present
Olya Mastikhina 2023-present
Ndidi Obinwanne 2022-present
Kevin Roice 2023-present
Haseeb Shah 2021-present
Abdul Wahab 2022-present
Farzane Aminmansour (*Inferring Brain Connectomes via Group-Sparse Factorization*) 2018-2020
Kirby Banman (*Strange springs in many dimensions: how parametric resonance can explain divergence under covariate shift*) 2020-2021
James Bell (*Electrodiagnostic Nerve Tests: Understanding Healthy Peripheral Nerves*, co-supervised Kelvin James) 2017-2019)
Alan Chan (*Greedification Operators for Policy Optimization: Investigating Forward and Reverse KL Divergences*) 2019-2020
Wesley Chung (*Two-Timescale Networks for Nonlinear Value Function Approximation*) 2017-2019
Shivam Garg (*Analysis of an Alternate Policy Gradient Estimator for Softmax Policies*, co-supervised Rupam Mahmood, CAIAC 2022 Best Masters Thesis Award) 2020-2021
Dhawal Gupta (*Structural Credit Assignment in Neural Networks using Reinforcement Learning*) 2020-2021
Maryam Hashemzadeh (*How Brain-Like is an LSTM's Representation of Nonsensical Language Stimuli?*, co-supervised Alona Fyshe) 2019-2021
Ehsan Imani (*Distributional Losses for Regression*) 2017-2019
Andrew Jacobsen (*Meta-descent for online, continual prediction*, co-supervised Adam White) 2018-2019
Taher Jaferfee (*Chasing Hallucinated Value: A Pitfall of Dyna Style Algorithms with Imperfect Environment Models*, co-supervised Michael Bowling) 2017-2019
Khurram Javed (*Learning Online-Aware Representations using Neural Networks*) 2019-2020
Sungsu Lim (*Actor-Expert: A Framework for using Q-learning in Continuous Action Spaces*) 2017-2019
Vincent Liu (*Sparse Representations for Reinforcement Learning*) 2017-2019
Chunlok Lo (*Goal-Space Planning with Subgoal Models*) 2021-2022
Erfan Miahi (*Measuring the Properties of Deep RL Representations that Do and Do Not Generalize Well*) 2021-2022
Gabor Mihucz (*Dyna with Options: Incorporating Temporal Abstraction into Planning*) 2021-2022
Andrew Patterson 2017-2018
Somjit Nath (*Fixed-Point Propagation for Recurrent Neural Networks*) 2017-2019
Matthew Schlegel 2016-2017
Sam Sokota (*Solving Common-Payoff Games with Approximate Policy Iteration*) 2019-2020

Han Wang (*Emergent Representations in Reinforcement Learning and Their Properties*) 2018-present
Niko Yasui (*An Empirical Study of Exploration Strategies for Model-Free Reinforcement Learning*) 2017-2019
Zaheer Abbas (*Selective Dyna-style Planning Using Neural Network Models with Limited Capacity*) 2017-2019

Undergraduate researchers Kai Luedemann (2023), Sam Scholnick-Hughes (2023), Thang Chu (2022), Robert Joseph (2022), Vlad Tkachuk (2021), Shaurya Seth (2021), Zonglun Li (2021), Victor Liao (2020-2021), Xinman Liu (2020, Now MSc at U Toronto), Matthew Regehr (2020, now MSc at U Toronto), Minghan Li (2019, now PhD at U Toronto), Jian Qian (2018, now PhD at MIT), Andrew Jacobsen (2018), Wenzhang Qian (2017), Andrew Patterson (2015-2017), Abraham Dasilvio (2016),Tyrese Taylor (2016)

Supervisory committee (PhD): Jiecao Chen (2019, IU), Sara Elkerdawy (2022), Pegah Fakhari (2018, IU), Sina Ghiassian (2022), Shantanu Jain (2018, IU), Jeffrey Kane Johnson (2017, IU), Mahdi Karami (2020), Jincheng Mei (2021), Katherine Metcalf (2019, IU), Madhavun Candadai Vasu (2020, IU), Chenjun Xiao (2022), Nadia Ady, Juan Hernandez Garcia, Eric Graves, Khurram Javed, Chen Ma, Samuel Neumann, Roshan Shariff, Kenny Young

Examining committee (PhD): Emmanuel Bengio (McGill, 2022), Adam Earle (Witwatersrand, 2019), Maximilian Igl (Oxford, 2021), Mao Li (UIUC, 2021), Khimya Khetarpal (McGill, 2022), Michael Mitchley (Witwatersrand, 2015), Anjana Puliyanda (2022), Jinnie Shin (2021), Samuele Tosatto (TU Darmstadt, 2020), Nino Vielliard (INRIA, 2022), Amy Zhang (McGill, 2021)

Teaching Experience

CMPUT 267: Basics of Machine Learning — Winter 2020-2022, Fall 2021
University of Alberta. Developed this first course, for a planned stream in Machine Learning.

CMPUT 367: Intermediate Machine Learning — Fall 2021,2022
Developed this second course, for a planned stream in Machine Learning.

Developed a Reinforcement Learning MOOC, on Coursera — 2019
60,000+ students registered online.

CMPUT 365: Reinforcement Learning — Fall 2019, 2020
Developed and introduced this new undergraduate course, based on the MOOC.

CMPUT 655: Reinforcement Learning I — Fall 2020
Taught 65 graduate students about fundamentals in RL and research in RL.

CMPUT 466/551: Machine Learning — Fall 2017, 2018, 2019

CMPUT 659: Optimization Principles for Reinforcement Learning — Winter 2018, 2019

CSCI B455: Principles of Machine Learning — Spring 2017 Indiana University.

CSCI B555: Machine Learning — Fall 2015, 2016

CSCI B659: Stochastic optimization for machine learning — Spring 2016

CSCI B554: Probabilistic Approaches to AI — Spring 2015

Outreach

Mentoring
Mentor at the Women in Machine Learning (WIML) luncheon — 2017, 2019
Lectures at the AI4Good Summer Lab — 2019-2022
Reviewing Mentor for ICLR — 2022

Workshops for youth
Pilot to develop Computing Camps for Aboriginal students in high school — 2018-2019
Presented to high school students about life as an undergraduate and graduate student in Computing Science (WP Wagner panel for Physical Sciences). — 2011
Volunteered for a Women in Scholarship, Engineering, Science and Technology (WISEST) open house

promoting diversity in Computing Science. 2011
Read to grade 3-6 students for a Read-In program promoting literacy. 2010
Held a workshop for junior high girls illustrating interesting aspects of theoretical Computing Science, under Women in Technology (WIT). 2007

Tutor
Tutoring children in an aboriginal high school with Frontier College. 2013
Tutored children from grades 1 to 12 and first year university in mathematics, physics, statistics, chemistry, biology, English and French. 2006
Tutored grade 5 girls in mathematics for the *Studdy Buddy Program*. 2005

Industrial Experience

Technical Board for Awenyx. 2021-present

Software Engineering Internship at Google. Summer 2012

PFM Scheduling company. 2010-2012
Part of the initial technical team for nurse scheduling for Alberta Health Services, that led to the spin-off for this company. `http://pfmscheduling.com`