
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** 2019-present
Department of Computing Science, Faculty of Science, University of Alberta

Assistant Professor 2017-2019
Department of Computing Science, Faculty of Science, University of Alberta

Assistant Professor 2015-2017
Department of Computer Science, 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] 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.
- [2] 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.
- [3] Selective Dyna-style Planning Under Limited Model Capacity. M. Zaheer, S. Sokota, E. Talvitie and M. White. *International Conference on Machine Learning (ICML)*, 2020.
- [4] 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.
- [5] 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.
- [6] 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.

- [7] 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.
- [8] Meta-Learning Representations for Continual Learning. K. Javed and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [9] 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.
- [10] 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.
- [11] Planning with Expectation Models. Y. Wan, M. Zaheer, A. White, M. White and R. S. Sutton. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.
- [12] 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.
- [13] Two-Timescale Networks for Nonlinear Value Function Approximation. W. Chung, S. Nath, A. Joseph and M. White. *International Conference on Learning Representations (ICLR)*, 2019.
- [14] 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.
- [15] 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.
- [16] An Off-policy Policy Gradient Theorem Using Emphatic Weightings. E. Imani, E. Graves and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [17] Supervised autoencoders: Improving generalization performance with unsupervised regularizers. L. Le, A. Patterson and M. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [18] 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.
- [19] Improving Regression Performance with Distributional Losses. E. Imani and M. White. *International Conference on Machine Learning (ICML)*, 2018.
- [20] 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.
- [21] Organizing experience: a deeper look at replay mechanisms for sample-based planning in continuous state domains. Y. Pan, M. Zaheer, A. White, A. Patterson, M. White. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.
- [22] High-confidence error estimates for learned value functions. T. Sajed, W. Chung and M. White. *Uncertainty in Artificial Intelligence (UAI)*, 2018.
- [23] 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.

- [24] 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.
- [25] Unifying task specification in reinforcement learning. M. White. *International Conference on Machine Learning (ICML)*, 2017.
- [26] Adapting kernel representations online using submodular maximization. M. Schlegel, Y. Pan and M. White. *International Conference on Machine Learning (ICML)*, 2017.
- [27] Effective sketching methods for value function approximation. Y. Pan, E. Sadeqi Azer and Martha White. *International Conference on Uncertainty in AI (UAI)*, 2017.
- [28] Learning sparse representations in reinforcement learning with sparse coding. L. Le, R. Kumaraswamy, and M. White. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.
- [29] Accelerated Gradient Temporal Difference Learning. Y. Pan, A. White and M. White. *AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- [30] Recovering true classifier performance in positive-unlabeled learning. S. Jain, M. White, P. Radivojac. *AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- [31] 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.
- [32] Investigating practical, linear temporal difference learning. A. White and M. White. *International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2016.
- [33] 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.
- [34] Incremental Truncated LSTD. C. Gehring, Y. Pan and M. White. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- [35] 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.
- [36] Scalable Metric Learning for Co-embedding. F. Mirzazadeh, M. White, A. Gyorgy and D. Schuurmans. *In ECML PKDD*, 2015.
- [37] Optimal Estimation of Multivariate ARMA Models. M. White, J. Wen, M. Bowling and D. Schuurmans. *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.
- [38] Partition Tree Weighting. J. Veness, M. White, M. Bowling, and A. Gyorgy. *Data Compression Conference*, 2013.
- [39] Convex Multiview Subspace Learning. M. White, Y. Yu, X. Zhang, D. Schuurmans. *Advances in Neural Information Processing Systems (NeurIPS)*, 2012.
- [40] Off-Policy Actor-Critic. T. Degris, M. White and R. S. Sutton. *International Conference on Machine Learning (ICML)*, 2012.
- [41] Generalized Optimal Reverse Prediction. M. White and D. Schuurmans. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012.
- [42] 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.

- [43] Interval Estimation for Reinforcement-Learning Algorithms in Continuous-State Domains. M. White and A. White. *Advances in Neural Information Processing Systems (NeurIPS)*, 2010.
- [44] 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.
- [45] 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**
- [46] 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 **Canada CIFAR AI (CCAI) Chair.** 2018-2023
\$650,000 in research funding for five years.

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 total for three years. Sole PI. “Optimizing water treatment operation using reinforcement learning.”

NSF CISE CRII grant. 2016-2018
\$174,616 total for 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
- University of Alberta Killam Accelerator Research Award, 2020
 - Champion of Inclusion, School of Informatics and Computing, Indiana University, 2016
 - Reviewing Award, International Conference on Machine Learning, 2015
 - Faculty of Science Doctoral Dissertation Award, 2015
 - Honourable Mention for Best Paper at the International Conference on Machine Learning, 2009
 - Several national and provincial scholarships for graduate studies, including
 - NSERC CGS D (2012) — 70,000 over two years in PhD, 2012, 233 offered in Canada
 - Honorary Izaak Walton Killam Memorial Scholarship (2012)
 - Alberta Innovates Scholarship for PhD (2010) — 72,000 over two years in PhD, 2010, in top 5% of accepted applicants
 - NSERC CGS M (2008) — 35,000 over two years in MSc, 2008, 713 offered in Canada
 - Alberta Innovates Scholarship (2008) — 25,000 over two years in MSc, 2008, in top 10% of accepted applicants

INVITED
TALKS

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, February 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	<p>Board Member for the International Conference on Learning Representations (ICLR) 2020 - present</p> <p>Program Co-Chair for ICLR with Shakir Mohamed, Kyunghyun Cho and Dawn Song 2019-2020</p> <p>Area Chair 2018 - Present</p> <p>AAAI Conference on Artificial Intelligence (AAAI), 2018, 2019</p> <p>International Conference on Learning Representations (ICLR), 2018, 2019</p> <p>International Conference on Machine Learning (ICML), 2018-2020</p> <p>Advances in Neural Information Processing Systems (NeurIPS), 2018, 2020</p> <p>Program Committee Member 2014 - Present</p> <p>International Conference on Machine Learning (ICML), 2016-2018</p> <p>Advances in Neural Information Processing Systems (NeurIPS), 2015-2018</p> <p>AAAI Conference on Artificial Intelligence (AAAI), 2014-2018</p> <p>International Joint Conference of Artificial Intelligence (IJCAI), 2015-2018</p> <p>International Conference on AI and Statistics (AISTATS), 2017</p> <p>Int. Conf. on Autonomous Agents and Multi-agent Systems (AAMAS), 2017</p> <p>Reviewer 2009 - Present</p> <p>Conferences: AAAI, AISTATS, AAMAS, ICLR, ICML, IJCAI, NeurIPS, UAI, CORL</p> <p>Journals: PNAS 2019-2020</p> <p>Nature 2019</p> <p>Journal of Machine Learning Research, 2014-2017</p> <p>Machine Learning Journal, 2014, 2018-2020</p> <p>Journal of Artificial Intelligence Research, 2014, 2016, 2017</p> <p>Artificial Intelligence Journal, 2014</p> <p>Journal of Autonomous Agents and Multi-agent Systems, 2016</p> <p>Transactions on Image Processing, 2014</p> <p>IEEE Transactions on Neural Networks and Learning Systems, 2014</p> <p>Stochastic Environmental Research and Risk Assessment, 2017</p> <p>Workshop Organization 2019-Present</p> <p>NeurIPS Optimization in RL Workshop, 2019</p> <p>ICML Exploration in RL Workshop, 2019, 2020</p> <p>RLDM Curiosity and Intrinsic Motivation Workshop, 2019</p> <p>NSERC Reviewer 2018- Present</p> <p>Discovery Grants, MITACS grants</p> <p>NSF panel member 2015, 2016</p> <p>Reviewed IIS: Robust Intelligence (RI) proposals on Machine Learning</p>
UNIVERSITY SERVICE	<p>University Service at the University of Alberta</p> <p>Faculty Evaluation Committee, for the Faculty of Science 2020-2022</p> <p>AI4Society Signature Area 2019-present</p> <p>Help manage the cross-department Statistical Machine Learning (SML) program 2018-present</p> <p>Service roles within Amii: Fellows Committee to nominate members and CCAI Chairs (2019-2021), Budgeting Committee (2018), Advisory Committee to CEO (2018-2020), EDI Committee (2020-2021)</p>

Departmental Service at the University of Alberta

AI Curriculum Committee	2020-present
Faculty Recruiting Committee	2017-2020
Data Science Curriculum design	2019-2020

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

Yash Satsangi (now Assistant Professor at Tilburg University)	2019-Present
Ajin George Joseph (now Assistant Professor at IIT Tirupati)	2018-2019
Omid Namaki (now Research Associate at Amii)	2017-2019

PhD students

Ehsan Imani	2019-present
Andrew Jacobsen	2019-present
Raksha Kumaraswamy	2016-present
Vincent Liu	2019-present
Yangchen Pan	2015-present
Andrew Patterson	2018-present
Matthew Schlegel	2017-present
Han Wang	2020-present
Lei Le (Thesis: Identifying Tractable Dictionary Learning Models for Representation Learning)	2015-2019

MSc students

Kirby Banman	2020-present
Shivam Garg	2020-present
Dhawal Gupta	2020-present
Maryam Hashemzadeh	2019-present
Khurram Javed	2019-present
Farzane Aminmansour (<i>Thesis: Inferring Macroscopic Brain Connectomes via Group-Sparse Factorization</i>)	2018-2020
James Bell (<i>Thesis: Electrodiagnostic Nerve Tests: Understanding Healthy Peripheral Nerves</i> , co-supervised Kelvin James)	2017-2019)
Alan Chan (<i>Thesis: Greedification Operators for Policy Optimization: Investigating Forward and Reverse KL Divergences</i>)	2019-2020
Wesley Chung (<i>Thesis: Two-Timescale Networks for Nonlinear Value Function Approximation</i>)	2017-2019
Ehsan Imani (<i>Thesis: Distributional Losses for Regression</i>)	2017-2019
Andrew Jacobsen (<i>Thesis: Meta-descent for online, continual prediction</i>)	2018-2019
Taher Jaferfee (<i>Thesis: Chasing Hallucinated Value: A Pitfall of Dyna Style Algorithms with Imperfect Environment Models</i>)	2017-2019
Sungsu Lim (<i>Thesis: Actor-Expert: A Framework for using Q-learning in Continuous Action Spaces</i>)	2017-2019
Vincent Liu (<i>Thesis: Sparse Representations for Reinforcement Learning</i>)	2017-2019

Andrew Patterson	2017-2018
Somjit Nath (<i>Thesis: Fixed-Point Propagation for Recurrent Neural Networks</i>)	2017-2019
Matthew Schlegel	2016-2017
Sam Sokota (<i>Thesis: Solving Common-Payoff Games with Approximate Policy Iteration</i>)	2019-2020
Han Wang (<i>Thesis: Emergent Representations in Reinforcement Learning and Their Properties</i>)	2018-present
Niko Yasui (<i>Thesis: An Empirical Study of Exploration Strategies for Model-Free Reinforcement Learning</i>)	2017-2019
Muhammad Zaheer (<i>Thesis: Selective Dyna-style Planning Using Neural Network Models with Limited Capacity</i>)	2017-2019

Undergraduate researchers Xinman Liu (2020), Matthew Regehr (2020), Minghan Li (2019, now PhD at University of 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), Jeffrey Kane Johnson (2017, IU), Pegah Fakhari (2018, IU), Shantanu Jain (2018, IU), Mahdi Karami (2020), Katherine Metcalf (2019, IU), Can Liu, Madhavun Candadai Vasu (2020, IU), Nadia Ady, Sina Ghiassian, Chenjun Xiao, Juan Hernandez Garcia

TEACHING EXPERI- ENCE	CMPUT 296: Basics of Machine Learning	Winter 2020
	University of Alberta. Developed this first course, for a planned stream in Machine Learning.	
	Developed a Reinforcement Learning MOOC, on Coursera	
		2019
	30,000+ students registered online.	
	CMPUT 397: Reinforcement Learning	Fall 2019, 2020
	University of Alberta. Developed and introduced this new undergraduate course.	
	CMPUT 466/551: Machine Learning	Fall 2017, 2018, 2019
	University of Alberta.	
	CMPUT 659: Optimization Principles for Reinforcement Learning	Winter 2018, 2019
University of Alberta.		
CSCI B455: Principles of Machine Learning	Spring 2017	
Indiana University.		
CSCI B555: Machine Learning	Fall 2015, 2016	
Indiana University.		
CSCI B659: Stochastic optimization for machine learning	Spring 2016	
Indiana University.		
CSCI B554: Probabilistic Approaches to AI	Spring 2015	
Indiana University.		
CMPUT 379: Operating Systems Concepts	Winter 2013	
University of Alberta. Instructor rating: 4.9/5.0		
OUTREACH	Mentoring	
	Mentor at the Women in Machine Learning (WIML) luncheon	2017, 2019
	Lectures at the AI4Good Summer Lab	2019, 2020

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 **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>