John Winder

johnwinderphd@gmail.com
johnwinder.me
john-winder-phd
Baltimore, MD, USA

Artificial Intelligence | Machine Learning | Reinforcement Learning

Mission

I am currently a Senior Staff AI/ML Research Scientist researching reinforcement learning (RL) for real-world systems at the Johns Hopkins University Applied Physics Laboratory (JHU/APL).

The aim of developing complex decision-making agents – capable of long-term reasoning, collaborating with other agents or humans, generalizing to new circumstances while operating under uncertainty in dynamic and open environments – drives my research. I'm working at the confluence of attention (transformers), RL with vision and language (grounded language acquisition), and multimodal deep generative (e.g., text-to-image) models.

Education

Ph.D. in Computer Science, *University of Maryland, Baltimore County (UMBC)* *2019*
Advised by Dr. Marie desJardins, Dr. Cynthia Matuszek
Thesis: *Abstract Decision Making and Concept Formation for Adaptability and Generalization*
Research in hierarchical reinforcement learning, state abstraction, probabilistic planning

M.S. in Computer Science, *UMBC* *2015*
Advised by Dr. Marie desJardins, Dr. Tim Oates
Research in machine learning, computer vision, feature extraction

B.S. in Computer Science, magna cum laude, *UMBC* *2013*

Publications

Proposals (Awarded)

- **John Winder**, Thomas Urban. *Beyond Human Reasoning - Bridging the Information Gap.* Johns Hopkins University Applied Physics Laboratory, Propulsion Grant (Internal R&D). Year 3 (competitively awarded each year), 2022-2023. Award total: $500,000. Topics: *Transformers and attention, graph neural networks (GNNs), variational auto-encoders (VAEs), multi-agent reinforcement learning (MARL), human-machine collaboration in virtual reality (VR) using Unity.*

- **John Winder**. *Towards Zero RL Safety Violations.* Johns Hopkins University Applied Physics Laboratory, AD FACT (Internal R&D). 2022-2023. Award total: approx. $100,000. Topics: *Safe reinforcement learning, simulation-to-reality (sim2real).*

- **John Winder**, Thomas Urban. *Beyond Human Reasoning - Bridging the Information Gap.* Johns Hopkins University Applied Physics Laboratory, Propulsion Grant (Internal R&D). Year 2 (competitively awarded each year), 2021-2022. Award total: $350,000. Topics: *Self-attention, GNNs, VAEs, MARL, spatiotemporal prediction, cognitive modeling.*

- Thomas Urban, Edward White, Matthew Sharp, **John Winder**. *Providentia's Potential.* Johns Hopkins University Applied Physics Laboratory, Propulsion Grant (Internal R&D). Year 1 (competitively awarded each year), 2021-2022. Award total: $250,000. Topics: *Causal inference and machine learning, counterfactual reasoning, Bayesian networks.*

- **John Winder**. *Ender's Dilemma: Defeating the Hive Mind.* Johns Hopkins University Applied Physics Laboratory, AD FACT (Internal R&D). 2021-2022. Award total: $100,000. Topics: *Multi-agent, hierarchical, and model-based reinforcement learning (MARL, HRL, MBRL).*

- **John Winder**, Thomas Urban. *Beyond Human Reasoning - Bridging the Information Gap.* Johns Hopkins University Applied Physics Laboratory, Propulsion Grant (Internal R&D). Year 1 (competitively awarded each year), 2020-2021. Award total: $250,000. Topics: *Novel neural network architectures (memory and self-attention), probabilistic graphical models, human-machine interaction, MARL.*

- **John Winder**. *Adversarial Attacks on RL & Explainable RL Agents.* Johns Hopkins University Applied Physics Laboratory, "I Have Hammer" Projects (Internal R&D). 2021-2022. Award total: $100,000. Topics: *Adversarial attacks on reinforcement learning, explainable machine learning (XAI/XRL).*

- *Contributions to proposals for research with the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL).* Johns Hopkins University Applied Physics Laboratory. 2020-2023. Award total: on the order of $1MM. Topics: *Transformers, world models, MARL, novel metrics for collaboration.*

- Cynthia Matuszek, Francis Ferraro, **John Winder**. *NRI: FND: Semi-Supervised Deep Learning for Domain Adaptation in Robotic Language Acquisition.* National Science Foundation (NSF), Information & Intelligent Systems (IIS). 2020-2023. Award total: $748,724. Topics: *Grounded language learning for robots, manifold alignment, imitation learning, inverse reinforcement learning.*

- Dan Lee, **John Winder**. *STTR Phase I: A Machine Learning Framework for Comprehensive Dental Caries Detection.* National Science Foundation (NSF), Industrial Innovation & Partnerships (IIP). 2020-2021. Award total: $224,999. Topics: *Medical imaging, computer vision for radiology, semantic segmentation, semi- and self-supervised learning.*

- Marie desJardins [and **John Winder** (student co-author)]. *Concept Formation in Partially Observable Domains.* National Science Foundation (NSF), Information & Intelligent Systems (IIS). 2018-2021. Award total: $399,993. Topics: *Concept-based knowledge transfer, state abstraction, online function approximation for contextual bandits.*

Journal Articles

- Karan K Budhraja, **John Winder**, Tim Oates. *Feature Construction for Controlling Swarms by Visual Demonstrations*. ACM Transactions on Autonomous and Adaptive Systems (TAAS), 12(2), 10. 2017.

Conference Papers

- Gaoussou Youssouf Kebe, Padraig Higgins, Patrick Jenkins, Kasra Darvish, Rishabh Sachdeva, Ryan Barron, **John Winder**, Don Engel, Edward Raff, Francis Ferraro, Cynthia Matuszek. *A Spoken Language Dataset of Descriptions for Speech-Based Grounded Language Learning*. Proceedings of The Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS 2021). 2021.

- **John Winder**, Stephanie Milani, Matthew Landen, Erebus Oh, Shane Parr, Shawn Squire, Marie desJardins, Cynthia Matuszek. *Planning with Abstract Learned Models While Learning Transferable Subtasks*. Proceedings of The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20). 2020.

- David Abel*, **John Winder***, Marie desJardins, Michael L Littman. *The Expected-Length Model of Options*. Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19) [*equal contribution]. 2019.

- Nakul Gopalan, Marie desJardins, Michael L Littman, James MacGlashan, Shawn Squire, Stefanie Tellex, **John Winder**, Lawson LS Wong. *Planning with Abstract Markov Decision Processes*. Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling (ICAPS-17). 2017.

- Nicholay Topin, Nicholas Haltmeyer, Shawn Squire, **John Winder**, Marie desJardins, James MacGlashan. Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-15). 2015.

Workshop Papers & Extended Abstracts

- Patrick Jenkins, Rishabh Sachdeva, Gaoussou Youssouf Kebe, Padraig Higgins, Kasra Darvish, Edward Raff, Don Engel, **John Winder**, Francis Ferraro, Cynthia Matuszek. *Presentation and Analysis of a Multimodal Dataset for Grounded Language Learning*. arXiv preprint arXiv:2007.14987. 2020.

- Patrick Jenkins, Padraig Higgins, Rishabh Sachdeva, **John Winder**, Cynthia Matuszek. *GLD: A Grounded Language Dataset of Object Images and Descriptions in Natural Language Text and Speech*. The 8th Mid-Atlantic Student Colloquium on Speech, Language and Learning (MASC-SLL 2020) [Extended Abstract]. 2020.

- Monali Saraf, Padraig Higgins, **John Winder**, Cynthia Matuszek. *A Human-Robot Interaction Data Set: Towards Active Learning*. The 8th Mid-Atlantic Student Colloquium on Speech, Language and Learning (MASC-SLL 2020) [Extended Abstract]. 2020.

- **John Winder**, Stephanie Milani, Matthew Landen, Erebus Oh, Shane Parr, Shawn Squire, Marie desJardins, Cynthia Matuszek. *Planning with Abstract, Learned Models.* Do Good Robotics Symposium (DGRS-19) [Extended Abstract]. 2019.

- **John Winder**, Marie desJardins. *Concept-Aware Feature Extraction for Knowledge Transfer in Reinforcement Learning.* Knowledge Extraction from Games (KEG-18) Workshop at the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18). 2018.

- **John Winder**, Shawn Squire, Matthew Landen, Stephanie Milani, Marie desJardins. *Towards Planning With Hierarchies of Learned Markov Decision Processes.* Integrated Execution of Planning and Acting Workshop (IntEx-17) at the Twenty-Seventh International Conference on Automated Planning and Scheduling (ICAPS-17). 2017.

- **John Winder**. *Anomaly Reasoning through Concept Formation for Planning and Reinforcement Learning.* Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling (ICAPS-17) [Doctoral Consortium]. 2017.

- Shawn Squire, **John Winder**, Matthew Landen, Stephanie Milani, Marie desJardins. *R-AMDP: Model-Based Learning for Abstract Markov Decision Process Hierarchies.* The Third Conference on Reinforcement Learning and Decision Making (RLDM-17) [Extended Abstract]. 2017.

- Nakul Gopalan, Marie desJardins, Michael L Littman, James MacGlashan, Shawn Squire, Stefanie Tellex, **John Winder**, Lawson LS Wong. *Planning with Abstract Markov Decision Processes.* The Third Conference on Reinforcement Learning and Decision Making (RLDM-17) [Extended Abstract]. 2017.

- **John Winder**. *A Framework for Anomaly Reasoning: Interpretation through Concept Formation for Knowledge Transfer and Lifelong Learning.* Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16) [Doctoral Consortium]. 2016.

- Nakul Gopalan, Marie desJardins, Michael L Littman, James MacGlashan, Shawn Squire, Stefanie Tellex, **John Winder**, Lawson LS Wong. *Planning with Abstract Markov Decision Processes.* Abstraction in Reinforcement Learning Workshop at the Thirty-Third International Conference on Machine Learning (ICML-16). 2016.

Work Experience

Johns Hopkins University Applied Physics Laboratory
Advanced AI Algorithms Section, Intelligent Platforms Group
Laurel, MD
Section Supervisor
(Senior Professional Staff Scientist)
Feb. 2021 - Present

Johns Hopkins University Applied Physics Laboratory
Laurel, MD
Senior Professional Staff Scientist
July 2020 - Feb. 2021

Department of Computer Science and Electrical Engineering
UMBC
Adjunct Assistant Professor
October 2020 - Present

Department of Computer Science and Electrical Engineering
UMBC
Faculty Research Assistant
Fall 2019 - July 2020

Interactive Robotics and Language (IRAL) Lab
UMBC
Graduate Research Assistant
Fall 2018 - Summer 2019

Multi-Agent Planning and Learning (MAPLE) Lab
UMBC
Graduate Research Assistant
Fall 2013 - Summer 2018

International Computer Science Institute (ICSI)
Berkeley, CA
Consultant
May - August 2016

CS Matters in Maryland (CSforALL)
Baltimore, MD
Graduate Assistant
May - August 2014, 2015

SAIC (Leidos)
Columbia, MD
Computer Science Intern
May - August 2012

Teaching Experience

Reinforcement Learning and Probabilistic Planning
MAPLE Lab Instructor
Summer, Winter 2018
Summer, Winter 2017

Principles of Operating Systems
Teaching Assistant
Spring 2014

Object Oriented Programming
Teaching Assistant
Fall 2013

Service

AAAI Conference on Artificial Intelligence (AAAI-23)	Senior Program Committee *Fall 2022*
Robotics: Science and Systems (RSS 2020)	Reviewer *Spring 2020*
Conference on Human-Robot Interaction (HRI 2020)	Program Committee (Reviewer) *Fall 2019*
AAAI Conference on Artificial Intelligence (AAAI-20)	Program Committee (Reviewer) *Fall 2019*
Conference on Robot Learning (CoRL-19)	Reviewer *Summer 2019*
Knowledge Extraction from Games (KEG-19) Workshop at AAAI-18	Program Committee (Reviewer) *Fall 2018*
Integrated Execution of Planning and Acting (IntEx-18) Workshop at ICAPS-18	Program Committee (Reviewer) *Spring 2018*
Maryland Computing Education Summit (CE21-Maryland)	Student Organizer, Volunteer *April 2016*

Academic Awards

IJCAI-16 Travel Award	*June 2016*
T. Rowe Price Associates Scholarship	*May 2013*
UMBC Class of 2013 Featured Student	*May 2013*
Marshall Scholar Nominee at UMBC	*May 2013*
Phi Beta Kappa	*Fall 2012*
Undergraduate Research Award	*Spring 2011*
UMBC Honors College	*2009 - 2013*

Knowledge & Skills

Areas: Artificial intelligence / machine learning (AI/ML) and reinforcement learning (RL) algorithms.

- Classic and Deep RL (DRL) algorithms such as Value Iteration, Q-learning, SARSA, Proximal Policy Optimization (PPO) Soft Actor-Critic (SAC), AlphaZero, MuZero;
- NN architectures such as transformers, graph neural networks (GNNs), convolutional neural networks (CNNs), multi-layer perceptrons (MLPs);
- Deep generative models such as Generative Adversarial Networks (GANs, e.g., StyleGAN and BigGAN), variational auto-encoders (VAEs), diffusion models;
- Computer Vision (CV) & Natural Language Processing (NLP) Tasks: object detection, object recognition, semantic segmentation, fine-tuning large CV models, neural machine translation, fine-tuning large language models (LLMs), using foundation models (FMs);
- Techniques: data augmentation (e.g., CutMix), self-supervised training (e.g., triplet loss), model surgery, curriculum learning, population-based training;
- Traditional supervised ML such as decision trees, random forests, logistic regression, naive Bayes, Bayesian networks, support vector machines;
- Unsupervised learning (dimensionality reduction and clustering) such as t-SNE, UMAP, PCA.

AI/ML, Deep learning Frameworks: Python: PyTorch, Tensorflow/Keras, JAX, ray/RLlib, stable-baselines3, OpenAI gym (also Mujoco, DM Control Suite, PettingZoo, OpenSpiel, many more), HuggingFace Transformers, scikit-learn, fast.ai, MONAI); C/C++: Nvidia CUDA; Java: BURLAP, Weka.
Data analysis, visualization: Python (NumPy/SciPy, IPython / Google Colab, Matplotlib, Plotly, Tensorboard, Weights & Biases), D3.js, Excel, R, MATLAB, Mathematica; data de-identification for HIPAA compliance.
General software development: Python (PyQt, Cython), Java/Kotlin, C, C++.
Web development: PHP (Laravel), SQL, JavaScript (Node.js, React), HTML, CSS.
Misc. programming: Java/Kotlin (Android development), libGDX (game development), C# (Unity), Git, LaTeX, Bash, Linux.
Software Architecture (API/Library Design): At JHU/APL I have designed two large-scale libraries for training reinforcement learning agents using a reproducible, configuration-driven paradigm. The libraries feature curriculum learning, population-based training, self-play and mixed heuristic-vs-RL play. They have been used to train hundreds of highly performant agents including those used on projects for DARPA and AFRL. As a PhD student I redesigned and refactored our lab's BURLAP library to include more advanced serialization, parallelization, and deep RL algorithms.
Leadership: Principal investigator or technical lead on 5+ projects; proposal/grant writing, raising funds from customers and investors (Lt Cols, VCs / Angels, CEOs); project scoping, planning, and execution for technical teams with 10+ SMEs.
Operations: Project management, budgeting (DFARs compliance, SBIR/STTR for NSF).