Max Olan Smith

Ph.D. Candidate, Computer Science and Engineering
University of Michigan, Ann Arbor

max.olan.smith@gmail.com
maxosmith.com
2260 Hayward Street
Ann Arbor, MI 48109 USA

Research Interests

Reinforcement Learning, Multiagent Learning, Deep Learning, Empirical Game Theory, and Education.

Education

University of Michigan, Ann Arbor, MI (2017–present)
Ph.D. Candidate in Computer Science (degree expected Spring 2023)
Advisor: Michael P. Wellman
Committee: Satinder Singh, Honglak Lee, Grant Schoenebeck

University of Michigan, Ann Arbor, MI (2014–2016)
B.S.Eng. in Computer Science
Summa Cum Laude

Awards and Honors

2021 Finalist representing the AI Lab, CSE Graduate Honors Competition

2018 Honorable Mention, NSF Graduate Research Fellowship

2016 EECS Outstanding Research Award, University of Michigan

2015 3rd Place, Information and Technology Services: Mobile App Challenge, University of Michigan

2014 1st Place, Microsoft Developer's Challenge

2014 IBM Sponsor Prize, MHacks IV

Professional Experience

2022 May – 2022 Aug **Research Scientist Intern**, DeepMind (Paris, France)
Host: Daniel Hennes
Researched methods for continual multiagent reinforcement learning. Investigated methods for policy evaluation through population-based metrics.

2017 May – 2017 Aug **Research Intern**, Montréal Institute for Learning Algorithms (Montréal, Québec, Canada)
Host: Aaron Courville
Built new Diplomacy multiagent dataset and environment, and performed preliminary studies on it resulting in a NeurIPS publication.

2016 May – 2016 Aug **Software Engineering Intern**, Google (Mountain View, CA, USA)
Host: Edward Lu
Expanded the travel team's conversion model to utilize additional advanced features resulting in higher model performance. Implemented RPC for serving conversion simulation data to partners.

2015 May – 2015 Aug **ORISE DHS HS-STEM Intern**, Sandia National Laboratories (Livermore, CA, USA)
Host: Nerayo Teclemariam
Created census data model with support for geo-fence queries of demographic information. Designed and implemented a learning to rank solution for searching through system models.

Teaching Experience

University of Michigan

2017 Fall **Graduate Student Instructor**, EECS 498/598: Reinforcement Learning†

Primary Instructor: Satinder Singh
Lecture: 63

2016 Fall **Undergraduate Teaching Assistant**, EECS 280: Programming and Data Structures

Primary Instructors: James Juett, Andrew DeOrio, Andrew Lukefahr, and Amir Kamil
Discussion: 27, Lecture: 652

2016 Winter **Undergraduate Teaching Assistant**, EECS 398: Computing for Computer Scientists†

Primary Instructors: Pat Pannuto, Marcus Darden
Lecture: 340

This class attempts to address the experience gap that exists across the spectrum of incoming Computer Science students. While driven by tools (shells, build systems, debuggers, version control), it explores how and why computer scientists interface with computers differently in their day-to-day activities, how to apply principles learned in courses to everyday activities, and ultimately how to be a more efficient user of computing resources.

This course has been adopted as part of the permanent curriculum at the University of Michigan as EECS 201: Computing Pragmatics, an advised co-requisite for first-year EECS majors.

2016 Winter **Undergraduate Teaching Assistant**, EECS 280: Programming and Data Structures

Primary Instructors: James Juett, Amir Kamil, and Edwin Olson
Discussion: 37, Lecture: > 500

2015 Fall **Undergraduate Teaching Assistant**, EECS 280: Programming and Data Structures

Primary Instructors: James Juett, Edwin Olson, Andrew DeOrio, and Rada Mihalcea
Discussion: 39, Lecture: > 500

† *Denotes the first offering of a course.*

Workshop

2018 **Instructor**, Big Data Summer Institute

2018 **Instructor**, Sports Analytics Summer Camp, Exercise & Sports Science Initiative

Professional Service

2023 Reviewer, Intl. Joint Conf. on Artificial Intelligence (IJCAI)

2022 Reviewer, Intl. Conf. on Learning Representations (ICLR)

Highlighted Reviewer, ICLR'22 (8% 543/6217)

2021–2023 Reviewer, Intl. Conf. on Machine Learning (ICML)

2020–2023 Reviewer, Conf. on Neural Information Processing Systems (NeurIPS)

2018–2022 Reviewer, NeurIPS Deep Reinforcement Learning Workshop

2017 Poster Chair, Michigan AI Symposium: AI for Society

Continued Education

2021 Preparing Future Faculty Seminar, University of Michigan

Advising and Mentoring

2021 – 2022 Elijah Soba (Masters → Deloitte, Software Engineer)

2020 – 2021 Yimin Zhu (Undergraduate → UC Berkeley, M.S. Financial Engineering)

Manuscripts

[M0] Co-Learning Empirical Games and World Models
Max Olan Smith and Michael P. Wellman
In Submission (NeurIPS). 2023.

[M1] Population-based Evaluation in Repeated Rock-Paper-Scissors as a Benchmark for Multiagent Reinforcement Learning
Marc Lanctot, John Schultz, Neil Burch, **Max Olan Smith**, Daniel Hennes, Thomas Anthony, and Julien Perolat
arXiv. 2023.

[M2] Learning to Play Against Any Mixture of Opponents
Max Olan Smith, Thomas Anthony, and Michael P. Wellman
In Submission (Frontiers in Artificial Intelligence). 2021.

Journal Publications

[J1] Strategic Knowledge Transfer
Max Olan Smith, Thomas Anthony, and Michael P. Wellman
Journal of Machine Learning Research (2023), to appear.

[J2] Long Term Effects of Pair Programming
Max Olan Smith, Andrew Giugliano, and Andrew DeOrio
IEEE Transactions on Education 61.3 (2017), pp. 187–194.

Conference Publications

[C1] Iterative Empirical Game Solving via Single Policy Best Response
Max Olan Smith, Thomas Anthony, and Michael P. Wellman
9th International Conference on Learning Representations. ICLR '21. 2021.
Acceptance: 687 / 2594 (26%).
Spotlight Presentation (3.9%).

[C2] No Press Diplomacy: Modeling Multi-Agent Gameplay
Philip Paquette, Yuchen Lu, Stephen Bocco, **Max Olan Smith**, Satya Ortiz-Gagne, Jonthan K. Kummerfeld, Satinder Singh, Joelle Pineau, and Aaron Courville
33rd Conference on Neural Information Processing Systems. NeurIPS '19. 2019.
Acceptance: 1428 / 6743 (21%).

[C3] Speaker Naming in Movies
Mahmoud Azab, Mingzhe Wang, **Max Olan Smith**, Noriyuki Kojima, Jia Deng, and Rada Mihalcea
16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. NAACL-HLT '18. 2018.
Acceptance: 207 / 647 (32%).

[C4] A Unified Framework for Automatic Wound Segmentation and Analysis with Deep Convolutional Neural Networks
Changhan Wang, Xinchen Yan, **Max Olan Smith**, Kanika Kochkar, Marci Rubin, Stephen M. Warren, James Wrobel, and Honglak Lee
37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. EMBC '15. 2015.

Other Articles (Blogs, Magazines, Newspapers, Etc.)

[O1] Learning in Multi-Agent Systems: Challenges and Considerations
Max Olan Smith
Dec. 2020. URL: https://ai.engin.umich.edu/2020/12/05/learning-in-multi-agent-systems/ (visited on 02/08/2020).