

Prabhat Nagarajan

CONTACT INFORMATION	Engineer Preferred Networks Tokyo, Japan	<i>E-mail:</i> prabhat@prabhatnagarajan.com <i>Webpage:</i> prabhatnagarajan.com
RESEARCH INTERESTS	Reinforcement Learning, Imitation Learning, Robot Learning	
EDUCATION	The University of Texas at Austin , Austin, TX Aug. 2014 - Aug. 2018 <i>Master of Science in Computer Science</i> , Aug. 2018 <ul style="list-style-type: none">• Thesis Title: <i>Nondeterminism as a Reproducibility Challenge for Deep Reinforcement Learning</i>• Committee: Peter Stone (advisor), Scott Niekum• GPA: 3.89 <i>Bachelor of Science in Computer Science</i> , Aug. 2018 <ul style="list-style-type: none">• GPA: 3.92 <i>Bachelor of Science in Mathematics</i> , Dec. 2017 <ul style="list-style-type: none">• GPA: 3.92	
PROFESSIONAL EXPERIENCE	Preferred Networks <i>Engineer</i> Tokyo, Japan ChainerRL & Robotics. Facebook <i>Software Engineer Intern</i> New York, NY Messenger Infrastructure. Microsoft <i>Software Engineer Intern</i> Redmond, WA Visual Studio Team Services. Yahoo! <i>Technical Intern</i> Sunnyvale, CA Ads Targeting.	Sept. 2018 - <i>Present</i> May 2017 - Aug. 2017 May 2016 - Aug. 2016 May 2015 - Aug. 2015
TEACHING EXPERIENCE	The University of Texas at Austin , Austin, TX <i>Tutor</i> Discrete Mathematics for Computer Science (CS 311) <i>Tutor</i> Data Structures (CS 314)	Fall 2015 Fall 2015
SERVICE	The University of Texas at Austin <i>MS Admissions Committee</i>	2018
REVIEWING	• Reviewer - RSS 2020 Workshop on <i>Closing the Academia to Real-World Gap in Service Robotics</i>	
ADVISING	Preferred Networks , Tokyo, Japan <i>Intern Advising</i> <ul style="list-style-type: none">• Aaron Havens - Model-based reinforcement learning from latent reward models• Zhang-Wei Hong - Improving ensemble reinforcement learning with knowledge distillation	

PAPERS UNDER REVIEW	[1] Yasuhiro Fujita, Toshiki Kataoka, Prabhat Nagarajan , and Takahiro Ishikawa. <i>ChainerRL: A Deep Reinforcement Learning Library</i> . Accept after revisions at the Journal of Machine Learning Research.
CONFERENCE PUBLICATIONS	[1] Yasuhiro Fujita, Kota Uenishi, Avinash Ummadisingu, Prabhat Nagarajan , Shimpei Masuda, and Mario Castro. <i>Distributed Reinforcement Learning of Targeted Grasping with Active Vision for Mobile Manipulators</i> . Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), October 2020. [2] Daniel Brown, Wonjoon Goo, Prabhat Nagarajan , and Scott Niekum. <i>Extrapolating Beyond Suboptimal Demonstrations via Inverse Reinforcement Learning from Observations</i> . Proceedings of the 36th International Conference on Machine Learning (ICML), June 2019.
REFEREED WORKSHOP PUBLICATIONS	[1] Zhang-Wei Hong, Prabhat Nagarajan , and Guilherme Maeda. <i>Swarm-inspired Reinforcement Learning via Collaborative Inter-agent Knowledge Distillation</i> . In Workshop on Deep Reinforcement Learning at the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), December 2019. [2] Yasuhiro Fujita, Toshiki Kataoka, Prabhat Nagarajan , and Takahiro Ishikawa. <i>ChainerRL: A Deep Reinforcement Learning Library</i> . In Workshop on Deep Reinforcement Learning at the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), December 2019. [3] Aaron Havens, Yi Ouyang, Prabhat Nagarajan , and Yasuhiro Fujita. <i>Learning Latent State Spaces for Planning through Reward Prediction</i> . In Workshop on Deep Reinforcement Learning at the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), December 2019. [4] Prabhat Nagarajan , Garrett Warnell, and Peter Stone. <i>Deterministic Implementations for Reproducibility in Deep Reinforcement Learning</i> . In AAAI 2019 Workshop on Reproducible AI, January 2019. [5] Prabhat Nagarajan , Garrett Warnell, and Peter Stone. <i>The Impact of Nondeterminism on Deep Reinforcement Learning</i> . In ICML Reproducibility in Machine Learning Workshop, July 2018.
THESES	[1] Prabhat Nagarajan . Nondeterminism as a Reproducibility Challenge for Deep Reinforcement Learning. <i>Master's Thesis</i> , The University of Texas at Austin, August 2018
TALKS	<ul style="list-style-type: none"> • “Deterministic Implementations for Reproducibility in Deep Reinforcement Learning” <i>AAAI 2019 Workshop on Reproducibility in AI</i> Honolulu, HI, January, 2018.
PROFESSIONAL MEMBERSHIPS	<ul style="list-style-type: none"> • AAAI - Member (2018–present)
PERSONAL DETAILS	Citizenship: USA Languages: English (Native), Japanese (JLPT N5)