

DILIP S. ARUMUGAM

dilip@cs.stanford.edu

RESEARCH INTERESTS

Reinforcement learning, representation learning, natural language processing, machine learning

EDUCATION

Stanford University

Stanford, CA

Ph.D. in Computer Science

Rotation Advisors: Emma Brunskill, Noah D. Goodman, Benjamin Van Roy, Dan Yamins

Fall 2018 - present

Brown University

Providence, RI

M.S. in Computer Science

B.S. in Computer Science

Advisor: Michael L. Littman

May 2018

May 2017

PUBLICATIONS

Reparameterized Variational Divergence Minimization for Stable Imitation

D. Arumugam, D. Dey, A. Agarwal, A. Celikyilmaz, E. Nouri, E. Horvitz, B. Dolan.

Preprint, 2019.

Goal-Directed Learning as a Bi-level Optimization Problem

P.L. Bacon, **D. Arumugam**, E. Brunskill

Multidisciplinary Conference on Reinforcement Learning and Decision Making, 2019.

Value Preserving State-Action Abstractions

D. Abel, N. Umbanhowar, K. Khetarpal, **D. Arumugam**, D. Precup, M. Littman

Multidisciplinary Conference on Reinforcement Learning and Decision Making, 2019.

ICLR Workshop on Structures and Priors in Reinforcement Learning, 2019.

State Abstraction as Compression in Apprenticeship Learning

D. Abel, **D. Arumugam**, K. Asadi, Y. Jinnai, M. Littman, L. Wong

Association for the Advancement of Artificial Intelligence (AAAI) Conference, 2019.

Neural Discrete State Abstraction for Interpreting Deep Reinforcement Learning

D. Arumugam, M. Littman

Preprint, 2018.

Deep Reinforcement Learning from Policy-Dependent Human Feedback

D. Arumugam, J. Lee, S. Saskin, M. Littman

Preprint, 2018.

Grounding Natural Language Instructions to Semantic Goal Representations for Abstraction and Generalization

D. Arumugam*, S. Karamcheti*, N. Gopalan, E. Williams, M. Rhee, L. Wong, S. Tellex
Autonomous Robots (AuRo), 2018.

State Abstractions for Lifelong Reinforcement Learning

D. Abel, **D. Arumugam**, L. Lehnert, M. Littman
International Conference on Machine Learning (ICML), 2018.

Sequence-to-Sequence Language Grounding of Non-Markovian Task Specifications

N. Gopalan*, **D. Arumugam***, L. Wong, S. Tellex
Robotics: Science and Systems, 2018.

Toward Good Abstractions for Lifelong Learning

D. Abel, **D. Arumugam**, L. Lehnert, M. Littman
NIPS Workshop on Hierarchical Reinforcement Learning, 2017.

Modeling Latent Attention Within Neural Networks

C. Grimm, **D. Arumugam**, S. Karamcheti, D. Abel, L. Wong, M. Littman
Preprint, 2017.

A Tale of Two DRAGGNs: A Hybrid Approach for Interpreting Action-Oriented and Goal-Oriented Instructions

S. Karamcheti, E. Williams, **D. Arumugam**, M. Rhee, N. Gopalan, L. Wong, S. Tellex
ACL Workshop on Language Grounding for Robotics, 2017. [*Best Paper Award*]

Mitigating Planner Overfitting in Model-Based Reinforcement Learning

D. Arumugam, D. Abel, K. Asadi, N. Gopalan, C. Grimm, J. Lee, L. Lehnert, M. Littman
Preprint, 2017.

Accurately and Efficiently Interpreting Human-Robot Instructions of Varying Granularities

D. Arumugam*, S. Karamcheti*, N. Gopalan, L. Wong, S. Tellex
Robotics: Science and Systems, 2017.

Grounding English Commands to Reward Functions

J. MacGlashan, M. Babes-Vroman, M. desJardins, M. Littman, S. Muresan, S. Squire, S. Tellex,
D. Arumugam, and L. Yang
Robotics: Science and Systems, 2015.

(* denotes equal contribution)

PROFESSIONAL EXPERIENCE

Reinforcement Learning Group Research Intern
Microsoft Research, Redmond, WA
Mentor: Debadepta Dey

Summer 2019

Reinforcement Learning Group Research Intern
Microsoft Research, Cambridge, UK
Mentors: Matthew Johnson, Katja Hofmann, & Dave Bignell

Summer 2016

R&D Software Engineering Intern
dMetrics Inc., Brooklyn, NY

Summer 2015

TEACHING EXPERIENCE

Teaching Assistant

Deep Learning (CSCI1470/CSCI2470), Professor Eugene Charniak, Fall 2017.

Deep Learning Seminar (CSCI2950K), Professor Eugene Charniak, Fall 2016.

Applied Artificial Intelligence (CSCI1410), Professor Stefanie Tellex, Fall 2015.

SKILLS

<i>Mathematics</i>	Probability, Calculus, Graph Theory, Information Theory
<i>Programming Languages</i>	Python, Java, Matlab, Racket, \LaTeX
<i>Libraries</i>	Tensorflow, PyTorch, Keras, Weka, BURLAP, Mallet, Sci-kit Learn, OpenCV, OpenAI Gym, ROS
<i>Relevant Coursework</i>	Learning & Sequential Decision Making, Deep Learning, Advanced Algorithms, Machine Learning, Applied Artificial Intelligence, Probability and Computing, Randomized Algorithms and Probabilistic Analysis, & Computational Linguistics

AWARDS & SERVICE

Conference Reviewer: Conference on Robot Learning (CoRL) 2018.

Conference Travel Awards: ICML 2018

Brown University CS Department Incoming Graduate Student Orientation Czar, 2017.

Brown University CS Department Senior Prize, 2017.

Sigma Xi Scientific Research Honor Society, 2017.

Brown University CS Department Spring Undergraduate Research Symposium Award, 2017.