DILIP S. ARUMUGAM

dilip@cs.stanford.edu

RESEARCH INTERESTS

Reinforcement Learning, Information Theory, Machine Learning, Artificial Intelligence

EDUCATION

Stanford University
Ph.D. Candidate in Computer Science — *In progress, Fall 2018 -*
Advisor: Benjamin Van Roy

Brown University
M.S. in Computer Science — *May 2018*
B.S. in Computer Science — *May 2017*
Advisor: Michael L. Littman

SELECTED PAPERS & PUBLICATIONS

1. *On Rate-Distortion Theory in Capacity-Limited Cognition & Reinforcement Learning*

   Dilip Arumugam, Mark K. Ho, Noah D. Goodman, Benjamin Van Roy.

   NeurIPS Workshop on Information-Theoretic Principles in Cognitive Systems, 2022.

2. *Deciding What to Model: Value-Equivalent Sampling for Reinforcement Learning*

   Dilip Arumugam, Benjamin Van Roy.

   Advances in Neural Information Processing Systems (NeurIPS), 2022.

   ICML Workshop on Decision Awareness in Reinforcement Learning, 2022.

3. *Planning to the Information Horizon of BAMDPs via Epistemic State Abstraction*

   Dilip Arumugam, Satinder Singh.

   Advances in Neural Information Processing Systems (NeurIPS), 2022.

   NeurIPS Workshop on Ecological Theory of Reinforcement Learning, 2021

4. *Between Rate-Distortion Theory & Value Equivalence in Model-Based Reinforcement Learning*

   Dilip Arumugam, Benjamin Van Roy.

   Multi-disciplinary Conference on Reinforcement Learning & Decision Making (RLDM), 2022.

5. *The Value of Information When Deciding What to Learn*

   Dilip Arumugam, Benjamin Van Roy.

   Advances in Neural Information Processing Systems (NeurIPS), 2021.

6. *Bad-Policy Density: A Measure of Reinforcement Learning Hardness*

   David Abel, Cameron Allen, **Dilip Arumugam**, D. Ellis Hershkowitz, Michael L. Littman, Lawson L.S. Wong.

   ICML Workshop on Reinforcement Learning Theory, 2021.

7. *Deciding What to Learn: A Rate-Distortion Approach*

   Dilip Arumugam, Benjamin Van Roy.

   International Conference on Machine Learning (ICML), 2021.

8. *An Information-Theoretic Perspective on Credit Assignment in Reinforcement Learning*

   Dilip Arumugam, Peter Henderson, Pierre-Luc Bacon.

   NeurIPS Workshop on Biological and Artificial Reinforcement Learning, 2020.

9. *Randomized Value Functions via Posterior State-Abstraction Sampling*

   Dilip Arumugam, Benjamin Van Roy.

   NeurIPS Workshop on Biological and Artificial Reinforcement Learning, 2020.

10. *Value Preserving State-Action Abstractions*

    David Abel, Nate Umbanhowar, Khimya Khetarpal, **Dilip Arumugam**, Doina Precup, Michael L. Littman.

    International Conference on Artificial Intelligence and Statistics (AISTATS), 2020.

    Multi-disciplinary Conference on Reinforcement Learning & Decision Making (RLDM), 2019.

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

11. *Goal-Directed Learning as a Bi-level Optimization Problem*

    Pierre-Luc Bacon, **Dilip Arumugam**, Emma Brunskill

    Multi-disciplinary Conference on Reinforcement Learning & Decision Making (RLDM), 2019.

12. *State Abstraction as Compression in Apprenticeship Learning*

    David Abel, **Dilip Arumugam**, Kavosh Asadi, Yuu Jinnai, Michael L. Littman, Lawson L.S. Wong.

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

13. *Deep Reinforcement Learning from Policy-Dependent Human Feedback*

    Dilip Arumugam, Jun Ki Lee, Sophie Saskin, Michael L. Littman.

    Preprint, 2018.

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

    Dilip Arumugam*, Sidd Karamcheti*, Nakul Gopalan, Eddie Williams, Mina Rhee, Lawson L.S. Wong, Stefanie Tellex.

    Autonomous Robots (AuRo), 2018.

15. *State Abstractions for Lifelong Reinforcement Learning*

    David Abel, **Dilip Arumugam**, Lucas Lehnert, Michael L. Littman.

    International Conference on Machine Learning (ICML), 2018.

16. *Sequence-to-Sequence Language Grounding of Non-Markovian Task Specifications*

    Nakul Gopalan*, **Dilip Arumugam***, Lawson L.S. Wong, Stefanie Tellex.

    Robotics: Science and Systems, 2018.

17. *Mitigating Planner Overfitting in Model-Based Reinforcement Learning*

    Dilip Arumugam, David Abel, Kavosh Asadi, Nakul Gopalan, Chris Grimm, Jun Ki Lee, Lucas Lehnert, Michael L. Littman.

    Preprint, 2017.

18. *Accurately and Efficiently Interpreting Human-Robot Instructions of Varying Granularities*

    Dilip Arumugam*, Siddharth Karamcheti*, Nakul Gopalan, Lawson L.S. Wong, Stefanie Tellex.

    Robotics: Science and Systems, 2017.

19. *Grounding English Commands to Reward Functions*

    James MacGlashan, Monica Babes-Vroman, M. desJardins, Michael L. Littman, Smaranda Muresan, Shawn Squire, Stefanie Tellex, **Dilip Arumugam**, and Lei Yang.

    Robotics: Science and Systems, 2015.

PROFESSIONAL EXPERIENCE

Research Scientist Intern — Summer 2021
DeepMind, London, UK
Mentors: Brendan O'Donoghue & Satinder Singh

Visiting Researcher — Summer 2020
Montreal Institute for Learning Algorithms (MILA), Montreal, QC
Mentor: Pierre-Luc Bacon

Reinforcement Learning Group Research Intern — Summer 2019
Microsoft Research, Redmond, WA
Mentor: Debadeepta Dey

Reinforcement Learning Group Research Intern — Summer 2016
Microsoft Research, Cambridge, UK
Mentors: Matthew Johnson, Katja Hofmann, & Dave Bignell

TEACHING EXPERIENCE

Stanford University ($\star$ = TA, † = Head TA, ‡ = Guest Lecturer)

- † Reinforcement Learning (CS234), Emma Brunskill, Winter 2022-2023.
- ‡ Reinforcement Learning: Frontiers (MS&E338), Benjamin Van Roy, Spring 2021-2022.
- † Reinforcement Learning (CS234), Emma Brunskill, Winter 2021-2022. $\big($Reviews$\big)$
- $\star$ Deep Multi-Task & Meta Learning (CS330), Chelsea Finn & Karol Hausman, Autumn 2021-2022.
- $\star$ Reinforcement Learning (CS234), Emma Brunskill, Winter 2020-2021. $\big($Reviews$\big)$
- $\star$ Deep Multi-Task & Meta Learning (CS330), Chelsea Finn, Autumn 2020-2021.
- $\star$ Artificial Intelligence: Principles & Techniques (CS221), Chelsea Finn & Nima Anari, Spring 2019-2020.

Brown University

- ⋆ Deep Learning (CSCI1470/CSCI2470), Eugene Charniak, Fall 2017.
- ⋆ Deep Learning Seminar (CSCI2950K), Eugene Charniak, Fall 2016.
- ⋆ Applied Artifical Intelligence (CSCI1410), Stefanie Tellex, Fall 2015.

AWARDS & SERVICE

Conference Reviewer - CoRL (2018), ICML (2020 [Top Reviewer Award], 2021, 2022), NeurIPS (2021 [Outstanding Reviewer Award]), ICLR (2022), IEEE TPAMI (2021)
Program Committee - NeurIPS Deep RL Workshop (2020, 2021, 2022), ICML Decision Awareness in RL Workshop (2022)
Student Co-organizer, Stanford Reinforcement Learning Group, 2020 - present.
Brown University CS Department Incoming Graduate Student Orientation Czar, 2017.
Brown University CS Department Senior Prize, 2017.
Sigma Xi Scientific Research Honor Society, 2017.