

# MUDIT VERMA

@ mudit.verma2014@gmail.com

+1-480-494-7411

Arizona, USA

github.com/famishedrover

famishedrover.github.io

## EDUCATION

---

Ph.D

**Arizona State University**

Aug 2019 – Present

AZ, USA

- Major: Computer Science, GPA: 4.0/4.0
- Ph.D. student advised by Dr. Subbarao Kambhampati. Working on areas of Explainable AI systems specifically Advising RL in the paradigm of Human in the loop Reinforcement Learning.

---

Bachelor of Technology

**Delhi Technological University**

Aug 2015 – June 2019

Delhi, India

- Major: Information Technology, GPA: 9.6/10

## HONORS & AWARDS

---

2022 :

- Engineering Graduate Fellowship, ASU

2019 :

- CIDSE Doctoral Fellowship
- 1st, Prestigious Smart India Hackathon (over 37000+ submissions)

2018 :

- 4th, Hack In The North (IIIT Allahabad)
- Selected for Education Innovation Mentorship Programme, ReadAlliance, USAID
- Department Rank 1, Merit Scholarship, for three consecutive years, DTU

2017 :

- 1<sup>st</sup>, READing Hackathon (USAID)
- 15, World Food India Hackathon
- Pramod Jain Scholarship, Best student at DTU
- Exemplary Contribution, Computer Society of India-DTU
- Interest Development Group Head, CSI-DTU Chapter
- DTU Merit Department Rank Scholarship 2017.

## PROJECTS

---

- **Perfect Observability is a Myth** Proposed a method to deal with partial observability of humans for Reinforcement Learning domains, when providing advice to agents.
- **Term Paper** Randomly Wired Networks are on the rise, have we been creating wrong Networks all along?
- **Term Paper** Diverging Emerging Field of Multi-Task Reinforcement Learning
- **Colors of Desert** Used D3 to highlight deserts are indeed colorful. Term Paper and Project.
- **TAC** App that adapts and teaches children/adults (specially dyslexic) to read/write/recognize using ML Techniques.
- **CookHub** Open Source Community for Recipes where one can chat, push, pull, fork, collaborate & view trending recipes and contributors.
- **Shut The Fake Up** App/Website Wisdom of Majority & AI for Fake News detection.
- **Text Summarization** Human like summarization using Pointer Generator Networks
- **StressOut** App to check one's stress levels and suggest better work timings to bring relief through Machine Learning.

## EXPERIENCE

---

Machine Learning Research Intern

**Apple Inc.**

May 2022 – Present

Cupertino, CA, USA

- Ongoing research work at the Machine Learning Research (MLR) Group with Rin Metcalf & Barry Theobald.

---

Deep Learning Software Engineering Intern

**Intel Corporation**

May – Aug 2021

Santa Clara, CA, USA

- First analysis of float32 ResNet50 architecture on Intel IceLake (ICX) machines.
- Proposed several optimizations BFloat16 performance (as benchmarked on CPX machines) on an ICX cluster. Additionally, provided the then Best Known method for working with ResNet50 architectures on Intel Endeavour Cluster.
- Parallely, worked on Quantized ResNet Models to show discrepancy in Saliency Based explanations between original RN50 and Quantized RN50.
- IntelMPI, Horovod, Tensorflow, PyTorch

---

Machine Learning Intern

**Samsung Semiconductor India Research**

May – July 2018

Bangalore, India

- Created DRAM Bank Simulator, (400 times faster) with enhanced Fault Classes.
- Novel Approach to Redundancy Analysis Algorithms through State Space Reduction schemes & Beating RA through Monte Carlo Tree Search and Residual Networks.
- Awarded Best Intern Project at SSIR.
- C++, Python, PyTorch, UCT Search, MCTS

---

Machine Learning Intern

**Samsung Semiconductor India Research**

May – July 2017

Bangalore, India

- Diagnosed issues with SSDs & Implemented SSD Simulator for Read/Write/Garbage Collection.
- Created an LSTM Algorithm - Stream Selection for Smart Data Categorization (STRASDAC) to reduce write-wearing in SSDs and improve Garbage Collection.
- Reached Best Intern Project Finals at SSIR.
- Tech : C++, PyTorch, Python, LSTM, ML

## SKILLS

---

Confident with C, C++, Python

Others : Java, JS/HTML/CSS, SQL, PDDL.

**Frameworks:** PyTorch, NLTK, Tensorflow, Keras, OpenAI-Gym, OpenCV, Scikit-Learn, D3, Hadoop, Android App Development, Bootstrap, Flask

# PUBLICATIONS

---

- Sreedharan, S., Soni, U., **Verma, M.**, Srivastava, S., & Kambhampati, S. (2022). Bridging the Gap: Providing Post-Hoc Symbolic Explanations for Sequential Decision-Making Problems with Inscrutable Representations. In International conference on learning representations (ICLR).
- Zadehi, Z., Sreedharan, S., **Verma, M.** & Kambhampati, S. (2022) Modeling the Interplay between Human Trust and Monitoring. In Companion of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI).
- Guan, L., **Verma, M.**, & Kambhampati, S. (2021). Widening the Pipeline in Human-Guided Reinforcement Learning with Explanation and Context-Aware Data Augmentation. Advances in neural information processing systems (NeurIPS). (Spotlight)
- Kambhampati, S., Sreedharan, S., **Verma, M.**, Zha, Y., & Guan, L. (2021). Symbols as a Lingua Franca for Bridging Human-AI Chasm for Explainable and Advisable AI Systems. In Proceedings of the AAAI Conference on Artificial Intelligence (Blue Sky Track).
- Zahedi, Z., **Verma, M.**, Sreedharan, S., & Kambhampati, S. (2021). Trust-Aware Planning: Modeling Trust Evolution in Longitudinal Human-Robot Interaction. In ICAPS 2021 Workshop on Explainable AI Planning.
- Gopalakrishnan, S., **Verma, M.**, & Kambhampati, S. (2021, June). Synthesizing Policies That Account For Human Execution Errors Caused By State Aliasing In Markov Decision Processes. In ICAPS 2021 Workshop on Explainable AI Planning.
- **Verma, M.**, & Buduru, A. B. (2020). Fine-grained Language Identification with Multilingual CapsNet Model. In 2020 IEEE Sixth International Conference on Multimedia Big Data (BigMM) IEEE.
- Guan, L. \*, **Verma, M.\***, & Kambhampati, S. (2020). Explanation Augmented Feedback in Human-in-the-Loop Reinforcement Learning. In 2020 ICML Workshop on Human in the Loop Learning (HILL 2020), NeurIPS 2020 HAMLETS, NeurIPS 2020 DRL
- Sreedharan, S., Soni, U., **Verma, M.**, Srivastava, S., & Kambhampati, S. (2020). Bridging the Gap: Providing Post-Hoc Symbolic Explanations for Sequential Decision-Making Problems with Black Box Simulators. In 2020 ICML Workshop on Human in the Loop Learning (HILL 2020)
- **Verma, M.**, Sinha, P., Goyal, K., Verma, A., & Susan, S. (2019, June). A Novel Framework for Neural Architecture Search in the Hill Climbing Domain. In 2019 IEEE Second International Conference on Artificial Intelligence and Knowledge Engineering (AIKE) (pp. 1-8). IEEE.
- **Verma, M.**, Bhambri, S., & Buduru, A. B. (2019). Making Smart Homes Smarter: Optimizing Energy Consumption with Human in the Loop. arXiv preprint arXiv:1912.03298.