

Mudit Verma

+1-480-494-7411 • mudit.verma2014@gmail.com • linkedin.com/in/iammudit • famishedover.github.io

EDUCATION

Arizona State University, Tempe, AZ 2019 – Expected Fall 2024

Ph.D in Computer Science, Advisor : Dr. Subbarao Kambhampati, GPA : 4.0/4.0

Delhi Technological University (Delhi College of Engineering), New Delhi, India 2015 – 2019

B.Tech in Information Technology, Advisor : Dr. Seba Susan, GPA : 9.6/10 (**Gold Medalist**)

WORK EXPERIENCE

Apple Inc May 2023 – Sept 2023

Machine Learning Research Intern, Cupertino, CA

Hosts : Machine Learning Research (MLR) Group, Rin Metcalf and Barry Theobald.

- **Published (ICLR 2024)** : Hindsight PRIORs for Reward Learning from Human Preferences.

Apple Inc May 2022 – Aug 2022

Machine Learning Research Intern, Cupertino, CA

Hosts : Machine Learning Research (MLR) Group, Rin Metcalf and Barry Theobald.

- **Published (IROS RLCONFORM (Oral), NeurIPS HILL 2022)** : Symbol Guided Hindsight Priors for Reward Learning from Human Preferences.

Intel Corporation May 2021 – Aug 2021

Deep Learning Software Engineering Intern, Santa Clara (Remote), CA

Deep Learning Acceleration and implications to explainability. Host : Wei Wang

- First analysis of float32 ResNet50 on Intel IceLake (ICX, market impact : \$140M) with proposed BFloat16 optimizations, **enabling CPU deployment of quantized models for four Intel MLPT teams.**
- Revealed differences in Saliency Based explanations in Quantized ResNet50 for CPU ICX deployment.

Samsung Semiconductor India Research May 2018 – Aug 2018

Machine Learning Intern, Bangalore, India

Worked on implementing a DRAM bank simulator followed by Redundancy Analysis . Host : Atishay Kumar

- Implemented C++ based DRAM Simulator, **(400x faster than baseline)** with enhanced Fault Classes.
- Proposed : Monte Carlo Tree Search (with state space reduction) using Residual Network based heuristic.
- This is the **best known solution (as of 2022)**. **Awarded Best Intern Project at SSIR.**

Samsung Semiconductor India Research May 2017 – Aug 2017

Machine Learning Intern, Bangalore, India

Worked on reducing write-wearing and improving garbage collection in SSDs. Host : Sandeep Sammatshetti

- Identified fault classes in SSDs. C++ bindings on Python UI to simulate SSD reads/write/garbage collection.
- Proposed : LSTM based Stream Selection for Smart Data Categorization **(28% improvement over baseline)**
- **Runners up Best Intern Project at SSIR.**

HONORS & AWARDS

ASU SCAI Doctoral Fellowship (\$9300) 2024, 2023, 2019

ASU Engineering Graduate Fellowship (\$3000) 2022

DTU/DCE Gold Medalist 2019

DTU/DCE Department Merit Rank Scholarship (\$1500) 2019, 2018, 2017

Pramod Jain Scholarship, Best Student at DTU (\$1200) 2017

First, Smart India Hackathon. Varanasi, India. (37000+ submissions) (\$1700) 2019

Fourth, Hack in the North (IIIT Allahabad), Allahabad, India (200+ participants)	2018
Education Innovation Mentorship Programme, ReadAlliance USAID (highly selective) (\$17,000)	2018
First, READING Hackathon. USAID, NASSCOM	2017
15th, World Food India Hackathon. 150+ teams pan India	2017

SKILLS

Research Areas :

(2019-Present): Reinforcement Learning from Human Feedback (RLHF), Embodied AI, Interactive Machine Learning, Preference Learning, Safe AI, Human Robot Interaction (HRI), Multimodal Learning, Large Language Models (LLM) and AI Agents

(2015-2019): AutoML, Neural Architecture Search, Applied RL, Search & MCTS

Technicals :

Programming : C++ / Python / PDDL / HDDL

Frameworks :

<i>Deep Learning</i> :	PyTorch, Jax, Tensorflow
<i>Reinforcement Learning</i> :	Stable-Baselines, ACME, Dopamine, Mujoco-Gym, openai-gym, D4RL
<i>Large Models</i> :	HuggingFace Transformers, OpenAI, DeepSeed,
<i>Machine Learning</i> :	Sklearn, Numpy, Pandas, Matplotlib, OpenCV, Einops
<i>Visualization</i> :	Matplotlib, Graphviz

Setups : Docker, Headless Computing, Multi-GPU and Single GPU nodes (Large Scale Training), Multi-Threaded Computing (Search, Rollouts)

Other Interests : Web Development (Flask, D3, HTML/CSS/JS, SQL), Cyber Security (TryHackMe), Android App Development

CONFERENCE PROCEEDINGS AND PREPRINTS: ICLR / NeurIPS / HRI / AAI / AAMAS

Overview: Published >20 papers combined in conferences and workshops at venues such as NeurIPS, ICLR, HRI, AAI, ICML, AAMAS, IROS, ICAPS.

Reinforcement Learning from Human Feedback :

- **Hindsight PRIORS for Reward Learning from Human Preferences.**
Mudit Verma, Katherine Metcalf
International Conference on Learning Representations (ICLR) 2024
- **Theory of Mind abilities of Large Language Models in Human-Robot Interaction : An Illusion**
Mudit Verma, Siddhant Bhambri, Subbarao Kambhampati
Conference on Human Robot Interaction (HRI) 2024 (Oral) ; Invited Talk : SuperAGI Leap Summit 2024
Previously as : Preference Proxies: Evaluating Large Language Models in capturing Human Preferences in Human-AI Tasks.
ICML Theory of Mind Workshop 2023, ICML Many Facets of Preference Learning 2023 (Oral)
- **LLMs Can't Plan, But Can Help Planning in LLM-Modulo Frameworks.**
Subbarao Kambhampati, Karthik Valmeekam, Lin Guan, Kaya Stechly, **Mudit Verma**, Siddhant Bhambri, Lucas Saldyt, Anil Murthy
Position Paper : *arXiv Preprint 2024*
- **Widening the Pipeline in Human-Guided Reinforcement Learning with Explanation and Context-Aware Data Augmentation.**

Lin Guan, **Mudit Verma**, Subbarao Kambhampati

Conference on Neural Information Processing Systems (*NeurIPS*) 2021 (**Spotlight**)

Previously as : Explanation Augmented Feedback in Human-in-the-Loop Reinforcement Learning.

Mudit Verma*, Lin Guan*, Subbarao Kambhampati

ICML Human in the Loop Learning 2020, NeurIPS 2020 Human And Machine in-the-Loop Evaluation and Learning Strategies, NeurIPS 2020 Deep Reinforcement Learning

- **Exploiting Action Distances for Reward Learning from Human Preferences.**

Mudit Verma, Siddhant Bhambri, Subbarao Kambhampati

Many Facets of Preference Learning Workshop at **MFPL ICML 2023**.

Previously as : Exploiting Unlabeled Data for Feedback Efficient Human Preference based Reinforcement Learning.

AAAI Representation Learning for Responsible Human-Centric AI 2023.

Responsible & Explainable AI / Symbolic Lingua Franca

- **Bridging the Gap: Providing Post-Hoc Symbolic Explanations for Sequential Decision-Making Problems with Inscrutable Representations.**

Sarath Sreedharan, Utkarsh Soni, **Mudit Verma**, Siddharth Srivastava, Subbarao Kambhampati

International Conference on Learning Representations (**ICLR**) 2022

Previously in : ICML HILL 2020

- **Symbols as a Lingua Franca for Bridging Human-AI Chasm for Explainable and Advisable AI Systems.**

Subbarao Kambhampati, Sarath Sreedharan, **Mudit Verma**, Yantian Zha, Lin Guan

Association for the Advancement of Artificial Intelligence AAAI 2021 (Blue Sky Track)

Trust Aware AI

- **Trust-Aware Planning: Modeling Trust Evolution in Iterated Human-Robot Interaction.**

Zahra Zahedi, **Mudit Verma**, Sarath Sreedharan, Subbarao Kambhampati.

Conference on Human Robot Interaction (HRI) 2023 (Oral)

- **Modeling the Interplay between Human Trust and Monitoring.**

Zahra Zadehi, Sarath Sreedharan, **Mudit Verma**, Subbarao Kambhampati

Conference on Human Robot Interaction (HRI) 2022 (Oral)

AI in Open Worlds

- **Methods and Mechanisms for Interactive Novelty Handling in Adversarial Environments.**

Mudit Verma, Tung Thai, Utkarsh Soni, Sriram Gopalakrishnan, Ming Shen, Mayank Garg, Ayush Kalani,

Nakul Vaidya, Subbarao Kambhampati, Neeraj Varshney, Chitta Baral, Jivko Sinapov, Matthias Scheutz

International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2023

Others :

- **Fine-grained Language Identification with Multilingual CapsNet Model.**

Mudit Verma, Arun Balaji Buduru

International Conference on Multimedia Big Data (IEEE BigMM) 2020

- **A novel framework for neural architecture search in the hill climbing domain.**

Mudit Verma, Pradyumn Sinha, Karan Goyal, Apoorva Verma, Seba Susan

International Conference on Artificial Intelligence and Knowledge Engineering (IEEE AIKE) 2019

WORKSHOP PROCEEDINGS : ICML / AAI / IROS / NeurIPS / ICAPS

Reinforcement Learning from Human Feedback :

- **Benchmarking Multi-Agent Preference based Reinforcement Learning for Human-AI Teaming.**
Mudit Verma, Siddhant Bhambri, Anil Murthy, Subbarao Kambhampati
(Invited AAAI 2024 Ad Hoc Teamwork Talk)
- **Exploiting Action Distances for Reward Learning from Human Preferences.**
Mudit Verma, Siddhant Bhambri, Subbarao Kambhampati
ICML Many Facets of Preference Learning Workshop (ICML MFPL) 2023.
Previously as : Exploiting Unlabeled Data for Feedback Efficient Human Preference based Reinforcement Learning. AAAI R2HCAI 2023.
- **Data Driven Reward Initialization for Preference based Reinforcement Learning.**
Mudit Verma, Subbarao Kambhampati
AAAI R2HCAI 2023
- **A State Augmentation based approach to Reinforcement Learning from Human Preferences.**
Mudit Verma, Subbarao Kambhampati
AAAI R2HCAI 2023
- **Advice Conformance Verification by Reinforcement Learning agents for Human-in-the-Loop.**
Mudit Verma, Ayush Kharkwal, Subbarao Kambhampati (2022).
IROS RLCONFORM 2022 (Oral)
- **Symbol Guided Hindsight Priors for Reward Learning from Human Preferences.**
Mudit Verma, Katherine Metcalf Susa
IROS RLCONFORM, NeurIPS HILL 2022 (Oral)
- **Towards Customizable Reinforcement Learning Agents: Enabling Preference Specification through Online Vocabulary Expansion.**
Utkarsh Soni, Sarath Sreedharan, **Mudit Verma**, Lin Guan, Matthew Marquez, Subbarao Kambhampati
NeurIPS HILL 2022

Trust aware AI :

- **Trust-Aware Planning: Modeling Trust Evolution in Longitudinal Human-Robot Interaction.**
Zahra Zahedi, **Mudit Verma**, Sarath Sreedharan, Subbarao Kambhampati
ICAPS Workshop on Explainable AI Planning (ICAPS XAIP) 2021 (Oral)

Human Aware AI :

- **Synthesizing Policies That Account For Human Execution Errors Caused By State Aliasing In Markov Decision Processes.**
Sriram Gopalakrishnan, **Mudit Verma**, Subbarao Kambhampati
ICAPS Workshop on Explainable AI Planning (ICAPS XAIP) 2021
- **Making Smart Homes Smarter: Optimizing Energy Consumption with Human in the Loop.**
Mudit Verma, Siddhant Bhambri, Arun Balaji Buduru
arXiv preprint

OTHER 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.
 - Linear Temporal Logic, Reinforcement Learning, Partially Observable Markov Decision Process

- *LTL, PyTorch, OpenAI-Gym, BabyAI, Matplotlib*
- **Randomly Wired Networks are on the rise, have we been creating wrong Networks all along?**
 - With Dr. Joshua Daymude, Proposed an Markov Chain Monte Carlo (MCMC) approach for studying randomly wired neural networks (AutoML) showing similar performance “carefully” orchestrated network architectures on Iris Dataset, ImageNet-Small
 - Erdos Reyni Graphs, MCMC methods
 - *NetworkX, PyTorch, Matplotlib*
- **Diverging Emerging Field of Multi-Task Reinforcement Learning**
 - With Dr. Dimitri Bertsekas, Proposed an umbrella view of several paradigms within multi-task reinforcement learning.
 - Appreciation from Dr. Dimitri Bertsekas on “exhaustive literature coverage” and timeliness of work.
- **Colors of Desert**
 - With Dr. Sharon Hsiao, Proposed D3 based web Visualization of Flora and Fauna of Arizona Desert. We view the desert through the lens of different colors it contains and provide various views to the user to navigate. This is similar to the present day Youtube recommender system recommending videos based on colors.
 - *D3, HTML, JS, CSS, BeautifulSoup, Google-Firebase*
- **User Study Interfaces**
 - Implemented a Flask-Jinja2 based extendable User Study Template to conduct user studies on online services like Prolific, Amazon MTurk
 - Has been used for over six user studies totalling over 700 participants.
 - *Flask, Firebase-Console, Jinja2, HTML, JS, CSS, JQuery*
- **Teach All Children (TAC)**
 - Android App and web interface that adapts and teaches specially abled children (dyslexic) to read/write/recognize using ML methods. Implemented Machine Learning stack.
 - We discussed key challenges to learning abilities with over 15+ specially abled schools.
 - Led the group to win (\$17,000) investment from ReadAlliance USAID
 - *Android App Development, Siamese Networks, Optical Character Recognition, Web Development*

TEACHING & SERVICE

Teaching Assistant, CSE 471 with Dr. Subbarao Kambhampati. Fall 2019

Reviewer/PC Member, at several Computer Science research conferences.

ICML (2024, 2023, 2022); NeurIPS, (2023, 2022); ICLR(2024, 2023, 2022), IJCAI(2024), ICAPS, (2023, 2022, 2021); AAI (2023, 2022), HRI (2022), BayLearn (2023) and several IEEE venues.