Dhananjay Ashok

PhD Student at the University of Southern California

dhananjay.ashok99@gmail.com | LinkedIn, Website, Google Scholar | +1 347 723 2785

FOCUS: Factually Grounded Language Model and Multimodal Model Systems

EDUCATION

Ph.D. in Computer Science, University of Southern California (Ongoing)
Researching Factual Grounding in LLMs and Multimodal LMs
Supervisor: Prof. Jonathan May

M.Sc. in Machine Learning, Carnegie Mellon University
Researched Distribution Shift and LLMs for Science
Supervisor: Prof. Zack Lipton

B.Sc. CS and Econ, University of Toronto
Researched Robotic Control and Neurosymbolic-AI
Supervisors: Prof. Animesh Garg, Prof. Vijay Ganesh

AWARDS

- Annenberg Fellowship, USC
- Dean Honour Scholar, UofT
- Valerie Brooks Scholarship
- William Kingston Scholarship

SKILLS

- Algorithms and Data Structures
- Python, C/C++, Bash, Java, Perl
- PyTorch, TensorFlow, Jax
- HuggingFace, Accelerate, DeepSpeed
- Distribution, Parallelization and Quantization of LLMs
- Fine-tuning, Prefix Tuning and Preference Optimization of LLMs
- Independent research

DEBATE

First speaker from a developing country to be judged Best Speaker at the World School Debating Championship

INDUSTRY EXPERIENCE

Machine Learning Research Engineer, Apple Inc. (Summer 2023)
- Developed systems for automated understanding and processing of log files
- Implemented MultiAgent RL Solutions to 6G Cellular Networking Problems

Accelerate AI Research Intern, Borealis AI. (Summer 2022)
- Developed new algorithms for gradient free training of Neural Networks
- Created GDSolver, the first Hybrid Solver+GD Framework for Fine-tuning NNs

Research Engineering Intern, AWS. (Summer 2021)
- Utilized C to verify security of safety critical AWS services and protocols

RESEARCH EXPERIENCE

CUTELABNAME, Prof. Jonathan May (2024-Current)
- Investigating problems related to Factual Grounding of Language Model systems

AutonLab, Prof. Barnabas Poczos (2022-2024)
- Researched Scientific Error Correction, developing a method that outperformed GPT3 despite having only 0.1% as many parameters

ACMI Lab, Prof. Zachary Chase Lipton (2022-2024)
- Created a State-of-the-art Few Shot NER System using LLMs
- Developed a principled Distribution Shift detection and mitigation method

Vector Institute, Prof. Animesh Garg (2019-2022)
- Applied methods from causal discovery for robotic manipulation and control

Vijay Ganesh, Prof. Vijay Ganesh (2019-2022)
- Created algorithms for verifiably compliant ML systems

SELECTED FIRST AUTHOR PUBLICATIONS

SciFix: Outperforming GPT3 on Scientific Factual Error Correction: **EMNLP**
- Conducted a detailed study of Controllable Text Generation methods, showing that instruction tuning consistently outperforms most approaches
- Introduced ConGenBench, a benchmark with hard controllable generation problems to facilitate future research

PromptNER: Prompting For Named Entity Recognition
- Developed a State-of-the-Art FewShot NER system, outperforming all prior methods on 6 different FewShot NER benchmarks using 2% of the available data

Controllable Text Generation in the Instruction Tuning Era
- Conducted a detailed study of Controllable Text Generation methods, showing that instruction tuning consistently outperforms most approaches
- Introduced ConGenBench, a benchmark with hard controllable generation problems to facilitate future research