Harsh Trivedi

Website: http://harshtrivedi.me **Github:** https://github.com/harshtrivedi

Education

PhD in Computer Science, Stony Brook University, US (**GPA: 4.0/4.0**) Jan'19- Present
Advisor: Professor Niranjan Balasubramanian

M.S in Computer Science, Stony Brook University, US (**GPA: 4.0/4.0**) Jan'17- Dec'18
Advisor: Professor Niranjan Balasubramanian

B. Tech in Information & Communication Technology, Aug'12-May'16
Dhirubhai Ambani Institute of Information & Communication Technology, India (**GPA: 8.23/10**)
Advisor: Professor Prasenjit Majumder

Research Internships

Aristo, Allen Institute of Artificial Intelligence (Ai2) Oct'21-March'22
Mentors: Ashish Sabharwal and Tushar Khot

CILVR lab, New York University Jun'20-Jun'22
Mentor: Sam Bowman

Aristo, Allen Institute of Artificial Intelligence (Ai2) Mar'19-July'19
Mentors: Ashish Sabharwal and Tushar Khot

Conference or Journal Papers

AppWorld: A Controllable World of Apps and People for Benchmarking Interactive Coding Agents
Annual Meeting of the Association for Computational Linguistics **(ACL) 2024** (**Best Resource Paper Award**)
H. Trivedi, T. Khot, M. Hartmann, R. Manku, V. Dong, E. Li, S. Gupta, A. Sabharwal, N. Balasubramanian

Interleaving Retrieval with Chain-of-Thought Reasoning for Knowledge-Intensive Multi-Step QA
Annual Meeting of the Association for Computational Linguistics **(ACL) 2023**
H. Trivedi, N. Balasubramanian, T. Khot, A. Sabharwal

Decomposed Prompting: A Modular Approach for Solving Complex Tasks conference
International Conference on Learning Representations **(ICLR) 2023**
T. Khot, **H. Trivedi**, M. Finlayson, Y. Fu, K. Richardson, P. Clark, A. Sabharwal

Teaching Broad Reasoning Skills for Multi-Step QA by Generating Hard Contexts
Transactions of the Association for Computational Linguistics **(EMNLP) 2022**
H. Trivedi, N. Balasubramanian, T. Khot, A. Sabharwal

MuSiQue: Multihop Questions via Single-hop Question Composition
Transactions of the Association for Computational Linguistics **(TACL) 2022**
H. Trivedi, T. Khot, A. Sabharwal, N. Balasubramanian

Summarize-then-Answer: Generating Concise Explanations for Multi-hop Reading Comprehension
Conference on Empirical Methods in Natural Language Processing **(EMNLP) 2021**
N. Inoue, **H. Trivedi**, S. Sinha, N. Balasubramanian, and K. Inui

What Ingredients Make for an Effective Crowdsourcing Protocol for Difficult NLU Data Collection?
Annual Meeting of the Association for Computational Linguistics **(ACL) 2021**
N. Nangia, S. Sugawara, **H. Trivedi**, A. Warstadt, C. Vania and S. Bowman

IrEne: Interpretable Energy Prediction for Transformers
Annual Meeting of the Association for Computational Linguistics **(ACL) 2021**
Q. Cao, Y. Lal, **H. Trivedi**, A. Balasubramanian, N. Balasubramanian

Is Multihop QA in DiRe Condition? Measuring and Reducing Disconnected Reasoning
Conference on Empirical Methods in Natural Language Processing **(EMNLP) 2020**
H. Trivedi, N. Balasubramanian, T. Khot, A. Sabharwal

DeFormer: Decomposing Pre-trained Transformers for Faster Question Answering
Annual Meeting of the Association for Computational Linguistics **(ACL) 2020**
H. Kwon, **H. Trivedi**, T. Khot, A. Sabharwal, N. Balasubramanian

Repurposing Entailment for Multi-Hop Question Answering Tasks
North American Chapter of the Association for Computational Linguistics **(NAACL) 2019**
H. Trivedi, H. Kwon, T. Khot, A. Sabharwal, N. Balasubramanian

Controlling Information Aggregation for Complex Question Answering
European Conference on Information Retrieval **(ECIR) 2018**
H. Kwon, **H. Trivedi**, P. Jansen, M. Surdeanu, N. Balasubramanian

Workshop or Demo Papers

Two-Turn Debate Does Not Help Humans Answer Hard Reading-Comprehension Questions
Conference on Neural Information Processing Systems **(NeurIPS) 2022** (Best Paper Award)
A. Parrish*, **H. Trivedi***, N. Nangia, V. Padmakumar, J. Phang, A. Saimbhi, S. Bowman

Single-Turn Debate Does Not Help Humans Answer Hard Reading-Comprehension Questions
A. Parrish*, **H. Trivedi***, E. Perez, A. Chen, N. Nangia, J. Phang, S. R. Bowman Annual Meeting of the Association for Computational Linguistics **(EMNLP) 2022**

IrEne-viz: Visualizing Energy Consumption of Transformer Models
Conference on Empirical Methods in Natural Language Processing **(EMNLP) 2021**
K. Lal, R. Singh, **H. Trivedi**, Q. Cao, A. Balasubramanian, N. Balasubramanian

→ * indicates equal contribution

Invited Talks

AppWorld: Reliable Evaluation of Interactive Agents in a World of Apps and People

- University of Waterloo — 27 September 2024
- University of Pennsylvania — 30 September 2024
- University of Southern California — 4 October 2024
- University of California Irvine — 7 October 2024
- Stanford University — 10 October 2024
- University of California Berkeley — 11 October 2024
- University of California Santa Barbara — 15 October 2024
- University of California San Diego — 17 October 2024
- Salesforce — 18 October 2024
- Johns Hopkins University — 23 October 2024
- Columbia University — 24 October 2024
- Princeton University — 25 October 2024
- University of California Santa Cruz — 29 October 2024
- Allen Institute for AI — 1 November 2024
- Apple — 6 November 2024
- Google DeepMind — 7 November 2024
- University of North Carolina Chapel Hill — 7 November 2024
- New York University — 8 November 2024
- Semantic Machines — 15 November 2024
- Camel AI — 10 December 2024

Awards and Recognition

- Invited to present NAIRR proposal work in the White House in 2024
- Best Resource Paper Award for AppWorld in ACL 2024
- Catacosinos Fellowship 2023
- Best Paper Award in ML Safety workshop at Neurips 2022
- Google Fellowship nominee from Stony Brook University in 2021

Professional Service

Program Committee: CoNLL'19, EMNLP'20, AAAI'21, NAACL'21, ACL'21, EMNLP'21, EMNLP'22, ACL'22, ICLR'23, EMNLP'23, ACL'23, ACL'24, COLM'24, ARR

Graduate Teaching Assistantship

Natural Language Processing, Stony Brook — Aug'20-Dec'20

Technical Skills

Languages: Python, Ruby, HTML and Javascript
Frameworks: Pytorch, Tensorflow, Transformers (Huggingface), Pandas and Scikit-learn