

Vidhisha Balachandran

GRADUATE RESEARCH ASSISTANT · CARNEGIE MELLON UNIVERSITY

5000 Forbes Avenue, Language Technologies Institute, Carnegie Mellon University

□ (+1) 412-961-2637 | □ vidhishanair@gmail.com | □ vidhishanair.github.io | □ vidhishanair

Education

Carnegie Mellon University

Pittsburgh, Pennsylvania

PHD IN LANGUAGE TECHNOLOGIES (GPA: 4.08/4.0)

August 2019 - Current

- Advisor: Prof Yulia Tsvetkov
- Primary Interests: Trustworthy and Safe NLP, Factual Text Generation, Model Interpretability, Evaluation, Fairness.

Carnegie Mellon University

Pittsburgh, Pennsylvania

MASTERS IN LANGUAGE TECHNOLOGIES (GPA: 3.89/4.0)

August 2017 - August 2019

- Advisors: Prof Jaime Carbonell and Prof William Cohen
- Recipient of Research Fellowship for the entire course duration (Monthly Stipend + Full Tuition Fee Waiver)
- Courses: Algorithms for NLP 11-711, Introduction to Machine Learning 11-701, Neural Networks for NLP 11-747, Structured Prediction for Language and Other Discrete Data 11-763, Probability and Statistics 36-700, Advanced Multimodal Machine Learning 11-777, Human Languages for AI 11-724, Topics in Deep Learning 10-707

PES Institute of Technology

Bangalore, India

BACHELORS IN COMPUTER SCIENCE AND ENGINEERING (GPA: 9.6/10)

September 2011 - May 2015

- Within the top 5% of the class
- Recipient of MHRD Scholarship 2011-2015 (Full Tuition Fee Waiver)
- Courses : Algorithms, Data Structures, Operating Systems, Data Mining, Natural Language Processing, Database Management Systems, Big Data

Experience

Allen Institute of Artificial Intelligence

Remote (Pittsburgh, Pennsylvania)

RESEARCH INTERN (HOSTS: DR MATTHEW PETERS, DR PRADEEP DASIGI)

May 2021 - August 2021

- Representation Learning for Long Documents in Scientific Domain
- Extended Transformer architectures to process documents of ~16K token length
- Developed pretraining techniques to encourage long range dependencies in representations

Google Brain

Remote (Pittsburgh, Pennsylvania)

RESEARCH INTERN (HOSTS: NIKI PARMAR, DR ASHISH VASWANI)

May 2020 - August 2020

- Developed Scalable Open-Domain QA system with Retrieval Augmented Generation models.
- Implemented efficient passage reranking to scale to ~13M Web documents.
- Leveraged the computational power of TPUs for faster Top-K MIPS Search

Google AI

Pittsburgh, Pennsylvania

RESEARCH INTERN (HOSTS: DR WILLIAM COHEN, DR MICHAEL COLLINS)

June 2019 - August 2019

- Developed KBQA models to reason using Text + Wikidata Background Knowledge
- Trained Fact-Aware text representations for different downstream tasks
- Implemented scalable module to align text with high-precision related facts from Wikidata

Flipkart Pvt Limited

Bangalore, India

SOFTWARE DEVELOPMENT ENGINEER 2

July 2015 - July 2017

- Developed logistical regression models to produce quality scores for E-Commerce products and listings.
- Built a scalable platform to scoring ~133M entities with low latency and self-learning feedback loops.
- Incorporated scores in product search ranking significantly reducing the produce return percentage by 60 basis points and increasing customer satisfaction by 0.5 Net Promoter Score
- Mentored interns on the NLP projects of text analysis; inferring causes for product returns; and contextual keyword extraction using RAKE and Doc2Vec

Publications

2023	Feng S, Balachandran V , Bai Y, Tsvetkov Y. : FactKB: Generalizable Factuality Evaluation using Language Models Enhanced with Factual Knowledge.	ArXiv
2023	Feng S, Shi W, Bai Y, Balachandran V , He T, Tsvetkov Y. : CooK: Empowering General-Purpose Language Models with Modular and Collaborative Knowledge.	ArXiv
2023	Derczynski L, Kirk HR, Balachandran V , Kumar S, Tsvetkov Y, Leiser MR, Mohammad S. : Assessing Language Model Deployment with Risk Cards.	ArXiv
2023	Ahia O, Gonen H, Balachandran V , Tsvetkov Y, Smith N. : LEXplain: Improving Model Explanations via Lexicon Supervision.	*SEM
2023	Balachandran V* , Kumar S*, Njoo L, Anastasopoulos A, Tsvetkov Y. : Language Generation Models Can Cause Harm: So What Can We Do About It? An Actionable Survey. (* equal contribution)	EACL
2023	Balachandran V* , Joshi R*, Saldanha E, Glenski M, Volkova S, Tsvetkov Y. : Unsupervised Keyphrase Extraction via Interpretable Neural Networks. (* equal contribution)	EACL
2022	Balachandran V , Hajishirzi H, Cohen W, Tsvetkov Y. : Correcting Diverse Factual Errors in Abstractive Summarization via Post-Editing and Language Model Infilling.	EMNLP
2021	Balachandran V , Pagnoni A, Lee JY, Rajagopal D, Carbonell J, Tsvetkov Y. : StructSum: Incorporating Latent and Explicit Sentence Dependencies for Single Document Summarization.	EACL
2021	Balachandran V , Vaswani A, Tsvetkov Y, Parmar N. : Simple and Efficient ways to improve REALM.	MRQA, EMNLP
2021	Balachandran V , Dhingra B, Sun H, Collins M, Cohen W. : Investigating the Effect of Background Knowledge on Natural Questions.	DeeLIO, NAACL
2021	Joshi R, Balachandran V , Vashisht S, Black A, Tsvetkov Y. : DialoGraph: Incorporating Interpretable Strategy-Graph Networks into Negotiation Dialogues.	ICLR
2021	Pagnoni A, Balachandran V , Tsvetkov Y. : Understanding Factuality in Abstractive Summarization with FRANK: A Benchmark for Factuality Metrics.	NAACL
2021	Rajagopal D, Balachandran V , Tsvetkov Y, Hovy E. : SelfExplain: A Self-Explaining Architecture for Neural Text Classifiers.	EMNLP
2020	Dhingra B, Zaheer M, Balachandran V , Neubig G, Salakhutdinov R, Cohen W. : Differentiable Reasoning over a Virtual Knowledge Base.	ICLR
2020	Radhakrishnan K, Chakravarthy S, Kanakagiri T, Balachandran V . : “A Little Birdie Told Me ... ” - Social Media Rumor Detection.	WNUT, EMNLP
2018	Balachandran V , Rajagopal D, Catherine R, Cohen W. : Learning to Define Terms in the Software Domain.	WNUT, EMNLP
2015	Sitaram D, Phalachandra HL, Harwalkar S, Murugesan S, Sudheendra P, Ananth R, Balachandran V , Kanji AH, Bhat SC, Kruti B. : Simple Cloud Federation.	AMS IEEE

Technical Skills

Programming Languages: Python, Java, C
Deep Learning Libraries: PyTorch, Tensorflow
Database Technologies: MySQL, Redis
Big Data Technologies: Hadoop, Hive, Spark
Operating Systems: OSX, Linux, Windows

Teaching

Upcoming	EMNLP 2023 , Mitigating Societal Impacts of Language Models	Tutorial
Spring 2022	The Web Conference 2022 , Mitigating Societal Impacts of Language Models	Tutorial
Spring 2022	Introduction to NLP (Undergraduate) (15-681) , Self Attention and Transformers	Lecture
Fall 2019	Artificial Intelligence Course (Graduate) (15-681) , Conducted Recitation, Designed Exams & Assignments	TA
Spring 2020	Neural Networks for NLP (Graduate) (11-747) , Assignment Grading, Project Mentorship	TA

Mentorship

Spring 23-Present	Varich Boonsanong , UW Undergraduate Student
Spring 22-Summer 23	Orevaoghene Ahia , UW PhD Student
Spring 22-Present	Krithika Ramesh, Gauri Gupta , Manipal Institute of Technology (India) Undergraduate Students
Fall 21-Fall 22	Kayo Yin , CMU Masters Student
Fall 20-Spring 22	Rishabh Joshi , CMU Masters Student
Spring 21-Spring 22	Luyu Gao , CMU Masters Student
Spring 21-Fall 21	Maxine Lui , CMU Undergraduate Student
Fall 19-Spring 21	Artidoro Pagnoni , CMU Masters Student
Spring 20-Fall 20	Karthik Radhakrishnan, Sharanya Chakravarthy, Tushar Kanakagiri , CMU Masters Students

Invited Talks

June, 2023	Actionable Directions for Reporting and Mitigating Language Model Harms , Center for Security and Emerging Technology	Georgetown University
May, 2023	Generalizable Factual Error Correction of Model Generated Summaries , SemaFor Working Group	DARPA
Jun, 2021	Simple and Efficient ways to improve REALM , N2Formal Reading Group	Google
Apr, 2021	On the Transparency and Reliability of Automatic Summarization , CRIM Seminar Series	CRIM Montreal
Jul, 2020	Incorporating External Background Knowledge into Natural Questions , Google News, Brain	Google

Service

Fall 2022	Workshop Organizer , COLING 2022 Workshop on Performance and Interpretability Evaluations of Multimodal, Multipurpose, Massive-Scale Models
2018-Present	Reviewer , ACL, EMNLP, NAACL, EACL, NeurIPS, SRW
Fall 2020 - Present	Committee Head , CMU LTI Mentoring Program
Fall 2020, 2021	Organizing Committee Member , CMU SCS Graduate Application Support Program
Fall 2020-Fall 2021	Member , CMU LTI DEI Committee
Spring, Fall 2021	Member , CMU SCS PhD Dean's Advisory Committee

Honours & Awards

2023	EECS Rising Star , RisingStars 2023	Georgia Tech, USA
2022	Scholarship Recipient , Cadence Diversity in Technology Scholarship	San Jose, USA
2014	Scholarship Recipient , Google Anita Borg Memorial Award Asia Pacific	Tokyo, Japan
2014	Scholarship Recipient , Grace Hopper Conference	Bangalore, India
2011-15	First Class Honours with Distinction , All eight semesters during Bachelors	Bangalore, India

References

Dr. Yulia Tsvetkov	University of Washington
ASSOCIATE PROFESSOR	yuliats@cs.washington.edu
Dr. William Cohen	Google Inc
PRINCIPAL SCIENTIST	wcohen@google.com
Dr. Hannaneh Hajishirzi	University of Washington, AI2
ASSOCIATE PROFESSOR, RESEARCH MANAGER	hannaneh@cs.washington.edu