

Kundan Krishna

PhD student, Language Technologies Institute, CMU

✉ kundank@andrew.cmu.edu

Basic Information

Mailing Address GHC-5515 (LTI), 5000 Forbes Avenue, Pittsburgh, PA-15213

Email kundank@andrew.cmu.edu

Research Interests

- **General:** Machine learning, Natural language processing, Interactive ML systems, Deep learning
- **Focus areas:** Ensuring factual correctness of LLM-generated text, Text summarization, Data-efficient pretraining, Human-in-the-loop based learning
- **Thesis:** Improving the reliability of summarization models (Thesis committee: [Zachary Lipton](#), [Jeffrey Bigham](#), [Sherry Wu](#), [Byron Wallace](#), [Alexander Rush](#))

Education

PhD in Language and Information Technologies

CARNEGIE MELLON UNIVERSITY

August 2018 - present

Pittsburgh, Pennsylvania

- Advisors: Professor Zachary C. Lipton and Professor Jeffrey P. Bigham
- CGPA: 3.96/4.0

B.Tech. in Computer Science and Engineering

INDIAN INSTITUTE OF TECHNOLOGY KANPUR

July 2012 - June 2016

Kanpur, India

- CPI: 9.9/10

Professional Experience

Research Intern

ABRIDGE AI

May 2023 - Sep 2023

- Designed an interactive system to help doctors to easily augment LLM-generated summaries of patient visits.

Research Intern

GOOGLE BRAIN

June 2022 - Aug 2022

- Worked on improving robustness of summarization models to the presence of noise in the input.
- Designed an algorithm to detect and remove arbitrary noise types without prior knowledge of its characteristics.

Applied Scientist Intern

AMAZON ALEXA AI

May 2020 - Aug 2020

- Worked on improving entity linking for video entities requested via Alexa.
- Implemented and improved a vector search based method delivering over 20% improvement in accuracy.

Research Engineer

ADOBE RESEARCH, INDIA

June 2016 - July 2018

- Identified potential applications of machine learning in Adobe's products.
- Designed models and algorithms and implemented research prototypes.

Selected Works

- **GenAudit: Fixing Factual Errors in Language Model Outputs with Evidence.** Kundan Krishna, Sanjana Ramprasad, Praxhar Gupta, Byron Wallace, Zachary C Lipton, Jeffrey P Bigham. [arXiv:2402.12566 \(preprint\)](#)
- **Filling in the Blanks: Helping Clinicians Add Missing Content to Documentation.** Kundan Krishna, Anna von Reden, Davis Liang, Elisa Ferracane, Nathan Price, Zachary C Lipton. *(ongoing work)*

- **USB: A Unified Summarization Benchmark Across Tasks and Domains.** Kundan Krishna, Prakhar Gupta, Sanjana Ramprasad, Byron C Wallace, Jeffrey P Bigham, Zachary C Lipton. *Conference on Empirical Methods in Natural Language Processing (EMNLP): Findings, 2023*
- **Downstream Datasets Make Surprisingly Good Pretraining Corpora.** Kundan Krishna, Saurabh Garg, Jeffrey Bigham, Zachary Lipton. *Annual Meeting of the Association for Computational Linguistics (ACL), 2023*
- **Improving the Robustness of Summarization Models by Detecting and Removing Input Noise.** Kundan Krishna, Yao Zhao, Jie Ren, Balaji Lakshminarayanan, Jiaming Luo, Mohammad Saleh, Peter J Liu. *Conference on Empirical Methods in Natural Language Processing (EMNLP): Findings, 2023*
- **Out-of-Distribution Detection and Selective Generation for Conditional Language Models.** Jie Ren, Jiaming Luo, Yao Zhao, Kundan Krishna, Mohammad Saleh, Balaji Lakshminarayanan, Peter J Liu. *International Conference on Learning Representations (ICLR), 2023*
- **Does Pretraining for Summarization Require Knowledge Transfer?.** Kundan Krishna, Jeffrey Bigham, Zachary Lipton. *Conference on Empirical Methods in Natural Language Processing (EMNLP): Findings, 2021*
- **Generating SOAP Notes from Doctor-Patient Conversations Using Modular Summarization Techniques.** Kundan Krishna, Sopan Khosla, Jeffrey Bigham, Zachary Lipton. *Annual Meeting of the Association for Computational Linguistics (ACL), 2021*
- **Extracting Structured Data from Physician-Patient Conversations By Predicting Noteworthy Utterances.** Kundan Krishna, Amy Pavel, Benjamin Schloss, Jeffrey Bigham, Zachary Lipton. *International Workshop on Health Intelligence at AAAI 2020*
- **Generating topic-oriented summaries using neural attention.** Kundan Krishna, Balaji V. Srinivasan. *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2018*
- **Vocabulary Tailored Summary Generation.** Kundan Krishna, Aniket Murhekar, Saumitra Sharma, Balaji V. Srinivasan. *International Conference on Computational Linguistics (COLING), 2018*
- **An LSTM based method for prediction of human activities with durations.** Kundan Krishna*, Deepali Jain*, Sanket Mehta, Sunav Choudhary. *UbiComp 2018*

Patents

- **Generating a topic-based summary of textual content.** *US Patent 10685050*
- **Bundling online content fragments for presentation based on content-specific metrics and inter-content constraints.** *US Patent 10891667*
- **Constructing content based on multi-sentence compression of source content.** *US Patent 10949452*
- **Visualizing natural language through 3D scenes in augmented reality.** *US Patent 10665030*
- **Generating a Targeted Summary of Textual Content Tuned to a Target Audience Vocabulary.** *US Patent 10534854*
- **Methods of automatically generating formatted annotations of doctor-patient conversations.** *US Patent App. 17736624*

Honors & Awards

2014-2016	Academic Excellence Award for 3 consecutive years for outstanding academic performance	<i>IIT Kanpur</i>
2014	Dr. Elizabeth and Varkey Cherian Award for best research project in the Summer Undergraduate Research Grant for Excellence(SURGE) program (jointly awarded to 2 out of 43 total research projects)	<i>IIT Kanpur</i>

Relevant Courses

- Algorithms for NLP
- Intermediate Statistics
- Grammars and Lexicons
- Convex Optimization
- Computational Ethics for NLP
- Computational Semantics
- Machine Learning Techniques
- Deep Learning
- Speech Recognition