Desh Raj

Website: desh2608.github.io
Email: draj@cs.jhu.edu
LinkedIn: rdesh26
GitHub: github.com/desh2608

Education

The Johns Hopkins University — Baltimore, US
Ph.D. in Computer Science — 2018–2023 (anticipated)

- Advisors: Sanjeev Khudanpur, Dan Povey
- Research interests: Multi-talker speech recognition and speaker diarization

Indian Institute of Technology Guwahati — Guwahati, India
B.Tech. in Computer Science and Engineering, GPA: 9.35/10 — 2013–2017

Work Experience

Meta AI — Menlo Park, US
Research Intern, AI Speech (Manager: Ozlem Kalinli) — May 2022 –August 2022

- Designed and implemented target-speaker ASR models to improve transducer performance in background speech and noise
- Obtained 60% and 25% WER reduction on non-overlapping and overlapping test sets based on LibriSpeech utterances, and significant improvement on production data for Portal devices

Microsoft Corporation — Redmond, US
Research Intern, AI Cognitive Services (Manager: Jinyu Li) — May 2021 –August 2021

- Extended Streaming Unmixing and Recognition Transducer (SURT) model for long-form meeting transcription
- Obtained $>$ 20% WER reduction using dual-path LSTM and Transformer models

Samsung Research — Bengaluru, India
Research Engineer in Advanced Technology Lab (ATL) — June 2017 –June 2018

- Implemented features for context engine built for a conversational assistant

Teaching

- **Teaching Assistant** at Johns Hopkins University — Fall 2021
  Information Theory (520.447/647)
- **Teaching Assistant** at Johns Hopkins University — Fall 2020
  Introduction to Human Language Technology (601.467/667)

Mentorship & Professional Services

- **Reviewer**: NeurIPS (2022), ICLR (2022), ICASSP (2022), SLT (2021, 2022), Elsevier CSL, IEEE TASLP
- CLSP Graduate Admissions Committee: 2021, 2022
- CLSP Student Recruitment Committee: 2019, 2020

Skills

- **ML/DL Toolkits:** PyTorch, Scikit-learn
- **ASR Frameworks:** Kaldi, ESPNet, Lhotse, k2
- **Other:** Audacity, Git

Languages

- **Programming:** Python, C++, Bash
- **Natural:** English, Hindi, French (beginner)
  - **TOEFL:** 119/120

Achievements

- Recipient of the JHU+Amazon **AI2AI fellowship** for 2022-23 2022
- JHU nominee for Microsoft Research Fellowship and Apple Scholars in AI/ML 2021
- **ISCA Travel Grant** (registration + membership + travel funds) for attending Interspeech 2021
- Member of Hitachi-JHU team which placed **top 2** in the DIHARD-3 challenge 2020
- Member of JHU team which placed **top 2** in the CHiME-6 challenge track 2 (diarization + ASR) 2020
- **INAE Travel Grant** by Govt. of India (worth INR 50,000) 2017

Publications

[1] **D. Raj**, L. Lu, Z. Chen, Y. Gaur, and J. Li, "Continuous Streaming Multi-talker ASR with Dual-path Transducers", *IEEE ICASSP*, 2022.

[2] M. Wiesner, **D. Raj**, and S. Khudanpur, "Injecting Text and Cross-lingual supervision in few-shot learning from self-supervised models", *IEEE ICASSP*, 2022.

[3] **D. Raj**, P. Denisov, Z. Chen, H. Erdogan, Z. Huang, M. He, S. Watanabe, J. Du, T. Yoshioka, Y. Luo, N. Kanda, J. Li, S. Wisdom, and J. R. Hershey, "Integration of speech separation, diarization, and recognition for multi-speaker meetings: system description, comparison, and analysis", *IEEE SLT*, 2021.

[4] **D. Raj**, P. Garcia, Z. Huang, S. Watanabe, D. Povey, A. Stolcke, and S. Khudanpur, "DOVER-Lap: A method for combining overlap-aware diarization outputs", in *IEEE SLT*, 2021.

[5] **D. Raj**, Z. Huang, and S. Khudanpur, "Multi-class spectral clustering with overlaps for speaker diarization", *IEEE SLT*, 2021.

[6] **D. Raj** and S. Khudanpur, "Reformulating DOVER-Lap Label Mapping as a Graph Partitioning Problem", *INTERSPEECH*, 2021.

[7] Z.-Q. Wang, H. Erdogan, S. Wisdom, K. Wilson, **D. Raj**, S. Watanabe, Z. Chen, and J. R. Hershey, "Sequential multi-frame neural beamforming for speech separation and enhancement", *IEEE SLT*, 2021.

[8] A. Arora, **D. Raj**, A. S. Subramanian, K. Li, B. Ben-Yair, M. Maciejewski, P. Zelasko, P. Garcia, S. Watanabe, and S. Khudanpur, "The JHU Multi-Microphone Multi-Speaker ASR System for the CHiME-6 Challenge", *CHiME-6 Workshop at IEEE ICASSP*, 2020.

[9] **D. Raj**, J. Villalba, D. Povey, and S. Khudanpur, "Frustratingly Easy Noise-aware Training of Acoustic Models", *ArXiv*, 2020.

[10] **D. Raj**, D. Snyder, D. Povey, and S. Khudanpur, "Probing the Information Encoded in X-Vectors", *IEEE ASRU*, 2019.

[11] **D. Raj**, S. K. Sahu, and A. Anand, "Learning local and global contexts using a convolutional recurrent network model for relation classification in biomedical text", in *CoNLL*, 2017.

See Google Scholar for a complete list of publications (450+ citations, h-index=10).