

Bandhav Veluri

bandhav@cs.washington.edu

<https://homes.cs.washington.edu/~bandhav/>

RESEARCH INTERESTS

Speech/Audio Language Modeling, Conversational AI and Human-Computer Interaction

EDUCATION

University of Washington, Seattle, WA Jul 2020 to present
Ph.D, Computer Science and Engineering
Advisors: Shyam Gollakota and Michael Taylor

University of Washington, Seattle, WA, Jun 2020
M.S, Electrical and Computer Engineering

Indian Institute of Technology Roorkee, Roorkee, India, May 2016
Bachelor of Technology, Electrical Engineering

WORK EXPERIENCE

Meta Fundamental AI Research, Seattle, WA Sep 2023 to Present
Visiting Researcher

Meta Fundamental AI Research, Seattle, WA July 2024 to Sep 2024
Research Intern

Texas Instruments, Bangalore, India July 2016 to July 2018
Digital Signal Processing Engineer

Janacare, Bangalore, India May 2015 to July 2015
Embedded Software Intern

PUBLICATIONS

Beyond Turn-Based Interfaces: Synchronous LLMs as Full-Duplex Dialogue Agents

Bandhav Veluri, Benjamin Peloquin, Bokai Yu, Hongyu Gong, Shyamnath Gollakota
Conference on Empirical Methods in Natural Language Processing (EMNLP Main 2024)

IRIS: Wireless Ring for Vision-based Smart Home Interaction

Maruchi Kim, Antonio Glenn*, Bandhav Veluri*, Yunseo Lee, Eyoel Gebre, Aditya Bagaria, Shwetak Patel, Shyamnath Gollakota* (*equal contribution)
ACM Symposium on User Interface Software and Technology (UIST 2024)

Look Once to Hear: Target Speech Hearing with Noisy Examples

Bandhav Veluri, Malek Itani*, Tuochao Chen, Takuya Yoshioka, Shyamnath Gollakota* (*equal contribution)
ACM Conference on Human Factors in Computing Systems (CHI 2024)

Scalable, Programmable and Dense: The HammerBlade Open-Source RISC-V Manycore

Dai Cheol Jung, Max Ruttenberg, Paul Gao, Scott Davidson, Daniel Petrisko, Kangli Li, Aditya K Kamath, Lin Cheng, Shaolin Xie, Peitian Pan, Zhongyuan Zhao, Zichao Yue, Bandhav Veluri, Sripathi Muralitharan, Adrian Sampson, Andrew Lumsdaine, Zhiru Zhang, Christopher Batten, Mark Oskin, Dustin Richmond, Michael Bedford Taylor
International Symposium on Computer Architecture (ISCA 2024)

SeamlessExpressiveLM: Speech Language Model for Expressive Speech-to-Speech Translation with Chain-of-Thought

*Hongyu Gong, **Bandhav Veluri***

Arxiv pre-print 2024

Semantic Hearing: Programming Acoustic Scenes with Binaural Hearables

Bandhav Veluri*, Malek Itani*, Justin Chan, Takuya Yoshioka, Shyamnath Gollakota (*equal contribution)

ACM Symposium on User Interface Software and Technology (UIST 2023)

REAL-TIME TARGET SOUND EXTRACTION

***Bandhav Veluri**, Justin Chan, Malek Itani, Tuoqiao Chen, Takuya Yoshioka, Shyamnath Gollakota*

IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2023)

NeuriCam: Key-Frame Video Super-Resolution and Colorization for IoT Cameras

***Bandhav Veluri**, Ali Saffari, Collin Pernu, Joshua Smith, Michael Taylor, Shyamnath Gollakota*

Annual International Conference On Mobile Computing And Networking (MobiCom 2023)

A Tensor Processing Framework for CPU-Manycore Heterogeneous Systems

*Lin Cheng, Peitian Pan, Zhongyuan Zhao, Krithik Ranjan, Jack Weber, **Bandhav Veluri**, Seyed Borna Ehsani, Max Rutenberg, Dai Cheol Jung, Preslav Ivanov, Dustin Richmond, Michael B Taylor, Zhiru Zhang, Christopher Batten*

IEEE TCAD 2021

BlackParrot: An Agile Open-Source RISC-V Multicore for Accelerator SoCs

*Daniel Petrisko, Farzam Gilani, Mark Wyse, Dai Cheol Jung, Scott Davidson, Paul Gao, Chun Zhao, **Bandhav Veluri**, Tavio Guarino, Zahra Azad, Sadullah Canakci, Ajay Joshi, Mark Oskin and Michael Bedford Taylor*

IEEE Micro 2020

The Celerity Open-Source 511-Core RISC-V Tiered Accelerator Fabric

*Scott Davidson, Shaolin Xie, Christopher Torng, Khalid Al-Hawaj, Austin Rovinski, Tutu Ajayi, Luis Vega, Chun Zhao, Ritchie Zhao, Steve Dai, Aporva Amarnath, **Bandhav Veluri**, Paul Gao, Anuj Rao, Gai Liu, Rajesh K. Gupta, Zhiru Zhang, Ronald G. Dreslinski, Christopher Batten, Michael Bedford Taylor*

IEEE Micro 2018

AWARDS

Best paper honorable mention for "Look Once to Hear" at CHI 2024

Recipient of the Meta AI Mentorship (AIM) Scholarship 2023-24.

Recipient of the Paul G. Allen School First-Year Ph.D. Fellowship 2020-21.

Recipient of the Merit-Cum-Means Scholarship at IIT Roorkee 2013-16.

Ranked 1556 among half a million students in Indian Institute of Technology Joint Entrance Examination (IITJEE) 2012.

**MEDIA
COVERAGE**

Noise-canceling headphones use AI to let a single voice through

MIT Technology Review, May 23, 2024

Related: US News, UW News, GeekWire

Noise-canceling headphones could let you pick and choose the sounds you want to hear

MIT Technology Review, Nov 9, 2023

Related: UW News, New Atlas

OPEN SOURCE

LookOnceToHear (★ 544, 🗄️ 56)

A novel human-interaction method for enhancing target speaker in real-time.

Waveformer (★ 287, 🗄️ 36)

A deep neural network architecture for low-latency audio processing.

NeuriCam (★ 90, 🗄️ 9)

Deep learning based video sensing method for low-power IoT cameras.