# Viswanath (Vish) Sivakumar

viswanathgs@gmail.com • viswanathgs.github.io
viswanathgs • viswanathgs • hLPvq9AAAAAJhl

I'm an AI research scientist and an engineer with over a decade of experience in foundational and applied research problems spanning vision, reinforcement learning, speech, biosignals, and ML systems. From 2020-2025, I led an interdisciplinary effort to build a neural interface for human-AI input via wrist electromyography (EMG), culminating in the launch of the Meta Neural Band, the first consumer product of its kind. I enjoy conducting principled science, span research and engineering fronts with equal ease, and have traversed the research to production arc a few times over across domains. I'm currently on sabbatical exploring independent directions.

## Experience

**Research Lead, Neural Interfaces**
*Meta Reality Labs* *Jan 2020–June 2025*

- Founding member of an ambitious research effort to build a neural interface text input solution using wrist-based electromyography (EMG) for AI glasses. The Meta Neural Band was unveiled at Connect 2025 by Mark Zuckerberg with a live demo, with the ability to decode handwriting showcased as its signature achievement.
- As the overall research lead for a 30+ person team, took the EMG-handwriting project from extreme ambiguity through various stages of research all the way to a launch-ready consumer product. Spanned wide disciplines (ML, neuroscience, sensing hardware, HCI, edge compute) while remaining deeply hands-on and solving the highest risk problems through principled science; mentored and grew the team; often conducted live demos to senior company leadership to communicate progress.
- Made fundamental research contributions to time-series modeling of neural signals, building on methods from adjacent domains such as speech recognition. This includes architecting and training ultra-low latency transformers optimized for real-time streaming and edge inference, designing evaluations to mitigate offline-online mismatch, and innovating on data-centric modeling strategies to address long-tail failure modes.
- Key contributor towards establishing scaling laws for biosignals; we demonstrated, against prevalent perception in neuroscience literature, that neural decoders can generalize across population given sufficient scale. Our group's foundational paper on this is accepted for publication in *Nature*.
- Advocated for and led an effort to open-source EMG datasets along with baselines to build community interest around practical non-invasive neural interfaces. Published in *NeurIPS*.
- Pioneered exploratory research around multimodal EMG-native LLMs towards the grand goal of a personalized and co-adaptive human-AI neural input interface.

**Research Scientist/Engineer, FAIR (Reinforcement Learning)**
*Meta AI* *Jan 2019–Dec 2019*

- Initiated the AI for Systems pillar in FAIR; conducted research on applying RL methodologies to replace decade-long hand-engineered algorithms underlying internet protocols.
- Researched multimodal reasoning using embedding lookup of external knowledge bases. This was subsequently refined into the now widely popular Retrieval Augmented Generation (RAG) methodology by my collaborators.
- Contributed to torchbeast, an open-source distributed RL platform in PyTorch.

**Research Scientist/Engineer, Applied AI (Computer Vision)**
*Meta AI* *Aug 2016–Dec 2018*

- Founded and built the company's large-scale multilingual OCR system at that time. Deployed across numerous products and served 30+ internal teams. Efforts spanned object detection research, authoring CUDA kernels, datacenter inference optimizations, novel research on quantization methods before the existence of torch.ao, etc.
- Published in KDD; made novel contributions to text detection state-of-the-art.
- In collaboration with FAIR, built a 50 billion scale ultra-fast image similarity search system.
- Contributed to several of Meta AI's widely used open-source projects during their early iterations: PyTorch, Caffe2, Detectron (object detection) and FAISS (GPU similarity search).

**Software Engineer, Infrastructure**
*Meta* *July 2012–July 2016*

- Worked on large-scale storage and networking systems, helped optimize routing and caching layers for ML training patterns, contributed to open-source networking libraries.

**ML Research Intern**
*Technische Universität Dortmund, Germany* *May 2011–July 2011*

- Researched and implemented feature selection algorithms for high-dimensional DNA micro-array data under the mentorship of Dr. Katharina Morik. Co-authored technical report.
- Contributed to RapidMiner, an early ML training platform.

**ML Research Intern**
*Indian Institute of Technology Madras, India* *May 2010–July 2010*

- Researched ML algorithms for classification of protein sequences under the mentorship of Dr. Ashish V. Tendulkar.

## Selected Publications

- CTRL-labs at Meta Reality Labs, **A generic noninvasive neuromotor interface for human-computer interaction**, Nature, 2025.
- **Viswanath Sivakumar**, Jeffrey Seely, Alan Du, Sean Bittner, Adam Berenzweig, Anuoluwapo Bolarinwa, Alex Gramfort, Michael Mandel, **emg2qwerty: A large dataset with baselines for touch typing using surface electromyography**, NeurIPS 2024.
- **Viswanath Sivakumar**, Tim Rocktaschel, Alexander H. Miller, Heinrich Kuttler, Nantas Nardelli, Mike Rabbat, Joelle Pineau, Sebastian Riedel, **MVFST-RL: An Asynchronous RL Framework for Congestion Control with Delayed Actions**, NeurIPS 2019 Machine Learning for Systems Workshop.
- Heinrich Kuttler, Nantas Nardelli, Thibaut Lavril, Marco Selvatici, **Viswanath Sivakumar**, Tim Rocktaschel, Edward Grefenstette, **TorchBeast: A PyTorch Platform for Distributed RL**, 2019.
- Fedor Borisyuk, Albert Gordo, **Viswanath Sivakumar**, **Rosetta: Large scale system for text detection and recognition in images**, KDD 2018.

## Education

**National Institute of Technology, Trichy, India**
*Bachelor of Technology, Computer Science and Engineering* *June 2008–May 2012*

- CGPA: 9.77/10. Graduated top of the class.