Omkar Nitsure

omkarnitsure2003@gmail.com | google-scholar | LinkedIn | Webpage

Education

- **ETH Zürich** *September 2025 – Present*
  MSc in Computer Science (Major: Machine Intelligence, Minor: Information Security)
- **Indian Institute of Technology Bombay** *November 2021 – May 2025*
  B.Tech. (Hons.) in Electrical Engineering & Minor in AI and Data Science **GPA: 9.34/10**

Publications

C=Conference, S=In Submission

[S.1] **Omkar Nitsure***, Sampath Kumar Dondapati*, Satish Mulleti. **Learnable Kernels for FRI: Joint Kernel Encoder Optimization and Hardware Validation**

[S.2] **Omkar Nitsure***, Gouranga Bala*, Amit Sethi. **A Benchmarking Framework for Uncertainty Quantification with Application to GMM and Test-Time Augmentation**

[C.1] Sampath Kumar Dondapati*, **Omkar Nitsure***, Satish Mulleti. **Super-Resolution via Learned Predictor** National Conference on Communications (NCC), New Delhi, India, 2025

[C.2] Darshan Prabhu*, Abhishek Gupta*, **Omkar Nitsure**, Preethi Jyothi, Sriram Ganapathy. **Improving Self-supervised Pre-training using Accent-Specific Codebooks**. *Interspeech 2024*, pp. 2310–2314.

* Equal Contribution

Research Experience

- **Certified Unlearning through Differential Privacy** – University of Zurich *September 2025 – Present*
  Guide: Prof. Anastasia Koloskova, Department of Mathematical Modeling and Machine Learning
  - Analyzed per-instance privacy guarantees of **DP-SGD** training and proved a bound on input memorization
  - This bound is **4x** more computationally efficient than the baseline while maintaining strong empirical results
  - Implemented a certified machine-unlearning algorithm that uses per-instance privacy loss to set noise variance during noisy fine-tuning, yielding a **40%** reduction in noise and an **11%** improvement in accuracy across tasks
- **Cross-lingual Knowledge Transfer in Multilingual LLMs** – IIT Bombay *December 2024 – May 2025*
  Guide: Prof. Preethi Jyothi, Department of Computer Science and Engineering
  - Isolated language-specific **knowledge neurons** for diverse LLM families like LLaMA, Mistral, Qwen, and BLOOM, spanning sizes from **1B** to **40B**. Enhanced cross-lingual knowledge transfer through targeted **fine-tuning**
  - Leveraged **causal tracing** to identify English-dominant representations in intermediate LLM layers using the final token of the prompt, resulting in a **22%** average improvement in factual retrieval across evaluated LLM families
  - Analyzed latent activations after patching from translations through few-shot examples and observed an **18%** increase in alignment to the corresponding English representations, indicating effective knowledge transfer
- **Uncertainty Quantification & Benchmarking** – IIT Bombay *December 2024 – May 2025*
  Guide: Prof. Amit Sethi, Department of Electrical Engineering [GitHub]
  - Developed distance-aware out-of-distribution (**OOD**) detection using **RBF** networks through a new loss function
  - Executed **adversarial attacks** for exploiting vulnerable data samples, which resulted in a **2x** inference speedup. Also evaluated the performance of **SOTA** techniques for uncertainty estimation through **novel benchmarks**
  - Employed an adversarial model-based method for **fine-tuning** during **inference** to achieve a **15%** relative improvement in classification accuracy for ImageNet samples with data uncertainty. A report is available here
- **Data-driven High-Resolution Spectral Estimation** – IIT Bombay *May 2024 – May 2025*
  Guide: Prof. Satish Mulleti, Department of Electrical Engineering [GitHub]
  - Designed Machine-learning-based algorithm to resolve frequencies in the **resolution limit** by predicting intermittent samples that optimize frequency estimation, leading to a **66%** reduction in sampling requirements
  - Developed an efficient ML algorithm for live-training the sampling kernel to optimize **reconstruction** of **FRI** signals for specific sampling rates, leading to **13%** improvement over SOTA methods for low sampling rates
- **Adapting ASR to Accented Speech using Codebooks** – IIT Bombay *December 2023 – December 2024*
  Guide: Prof. Preethi Jyothi, Department of Computer Science and Engineering [GitHub]
  - Implemented a **gating** mechanism in codebook cross-attention, trained using the supervision of a soft distribution extracted during beam search decoding to fine-tune **accent codebooks**, reducing active parameter count by **15%**
  - Enhanced HuBERT to have **accent-generic** and **accent-specific** targets and pooled k-means centers across accents to generate a single codebook, encouraging accent information mixing leading to **7%** improvement in WER
  - Reduced the number of codebooks used for cross-attention by **80%** through a custom **budget-based** restriction

Work Experience

- **Machine Learning Engineering Intern** *May 2025 – August 2025*
  Kipling Secure *Remote*
  - Developed an MCP server with gRPC to communicate with Kipling APIs and deployed it in an agentic system using the OpenAI Responses API, increasing the effectiveness of user interactions by **21%** as reported in surveys
  - Designed a UI client for users to interact with their data in the **ClickHouse** database through native MCP server
- **Systems Engineering Intern** *May 2024 – July 2024*
  Texas Instruments *Bangalore, India*
  - Minimized non-linearities with memory, introduced by transceiver Power Amplifier for 5G telecommunication applications using **digital-pre-distortion** (DPD) models, improving the **ACLR** metric by **16%**
  - Developed a Transformer-based model to correct PA non-linearities, achieving **27%** reduction in model size
  - Integrated channel filtering to mitigate **IMD3** interference in **ILC** data extraction during testing

Technical Projects

- **Constrained Sampling from Large Language Models** – IIT Bombay *January 2024 – May 2024*
  Guide: Prof. Sunita Sarawagi, Department of Computer Science and Engineering [GitHub]
  - Improved the **COLD** discrete sampling algorithm by performing beam search decoding in continuous logit space before discretizing into tokens, resulting in improved fluency and coherence of text while satisfying the constraints
  - Enhanced **lexically constrained generation** for words with a **low probability** of **co-occurrence** in the LM output distribution by **6%**, with a **10%** improvement in fluency, as reflected in **perplexity** and **GPT-3.5** assessment scores
  - Achieved **positive sentiment transfer** using a novel constraint function. A detailed report is available here
- **AI vs. Human Text Classifier** – IIT Bombay *August 2023 – December 2023*
  Guide: Prof. Amit Sethi, Department of Electrical Engineering [GitHub]
  - Trained a **SentencePiece**-style tokenizer on OpenWebText-10k dataset chunk with **1024** tokens in the vocabulary
  - Employed **Llama 2-7B-Chat** and **OpenAI GPT API** to rephrase text from **The Pile** & **C4** dataset chunks and trained an encoder-only Transformer with a context length of **512** tokens, correctly classifying **80%** of text samples
  - Deployed the trained model as an easy-to-use classification tool on a live-hosted website using **Flask** and **Docker**
- **Machine Learning for Monitoring and Forecasting of Mumbai Floods** – IIT Bombay *May 2023 – January 2024*
  Guide: Prof. Subimal Ghosh, Interdisciplinary Program in Climate Studies [GitHub]
  - Used rainfall, wind speed, and GFS data available at **15**-minute intervals from **37** AWSs across Mumbai and trained a Transformer-based model to predict rainfall for the next **3** hours using weather data from the past **3** hours
  - Leveraged **2D CNN** backbones like U-Net to extract data-overlap of spatially proximal stations and improved the correlation by **17%** over the entire prediction period while boosting the hit rate of high rainfall durations
- **Memory-efficient Image Classification using Compressed Sensing** – IIT Bombay *January 2024 – May 2024*
  Guide: Prof. Ajit Rajwade, Department of Computer Science and Engineering [GitHub]
  - Deployed a **smashed**-filter-based algorithm for image classification and achieved an accuracy of **91%** on the MNIST dataset, using only **10%** of the original number of camera measurements through compressed sensing
  - Employed deep learning to learn the **sensing matrix**, further boosting classification accuracy from **87%** to **96%** under arbitrary rotations on the MNIST dataset with a minimal increase in coherence. A report is available here

Technical Skills

- **ML Libraries:** PyTorch, JAX, TensorFlow, Hugging Face, NumPy, Pandas, OpenAI Responses API
- **Programming Languages:** Python, Bash, SQL, C++, C, VHDL, Assembly, Embedded C, Go, protocol buffers
- **Software & Frameworks:** MATLAB, Git, Intel Quartus Prime, Keil, Docker, Flask, MCP, React, gRPC, LaTeX

Scholastic Achievements

- Secured an **All India Rank** of **275** in JEE Advanced among **260,000** aspirants *October 2021*
- Secured an **All India Rank** of **120** in JEE Main among **1 million** aspirants *September 2021*
- Placed among **National Top 1%** in Indian Olympiad Qualifier in Chemistry (**IOQC**) Stage 1 *March 2021*
- Awarded with the prestigious **KVPY Fellowship** by **IISc** Bangalore, Government of India *January 2021*

Extracurricular Activities

- Mentored 100+ JEE Advanced 2022/23 aspirants of Vidyalankar Coaching Institute, Thane *October 2021*
- Captained school table tennis team in district-level Inter-school table tennis competition *September 2018*
- Awarded A grade in state-level Elementary drawing examination, Govt. of Maharashtra *March 2016*
- Successfully completed a 2-semester long course of NSO at IIT Bombay *2021 – 2022*