

SIDDHARTH GANDHI

✉ ssgandhi602@gmail.com ≡ ssgandhi.com 📺 ssgandhi1 🌐 Siddharth-Gandhi

Education

Carnegie Mellon University - School of Computer Science

Expected Dec. 2024

Master of Science in Intelligent Information Systems (Deep Learning) — GPA: 4.04/4.33

Pittsburgh, PA

- Coursework: Search Engines (Teaching Assistant), Deep Learning, Machine Learning, Visual Learning, Multimodal Machine Learning, Advanced Natural Language Processing, Deep Learning Systems

Vellore Institute of Technology - School of Computer Science & Engineering

July 2023

Bachelor of Technology in Computer Science & Engineering — GPA: 9.69/10 (Department Rank 3)

Vellore, India

Experience

Carnegie Mellon University

Aug 2024 – Present

Research Assistant (Advisor: Mihir Prabhudesai and Deepak Pathak)

Pittsburgh, PA

- Worked on the first unified multimodal discrete diffusion model and their tradeoffs (preprint submitted to ICLR 2025).
- Previously explored a new regularization method in autoregressive models and its generalization impact.

Cash Flow Portal (YC W22)

May. 2024 – Aug. 2024

Software Engineering Intern

Remote

- Implemented key features for Pallas, a real-estate marketplace, including a homepage redesign, listing pages for LPs/GPs, a review system to address the cold start problem, and advanced searching, filtering, and sorting capabilities.
- Developed an Excel-like cell feature and a papertrail feature for the classification system in the AI-Underwriting product for T12 sheets, allowing users to view/modify classifications, preserving the original structure of sheets.
- Optimized load times, search speed, added tests, and organized both internal and user notifications, resulting in improved performance and user experience.

Carnegie Mellon University

Sept. 2023 – May 2024

Research Assistant (Advisor: Prof. Jamie Callan) | [github](#), [report](#)

Pittsburgh, PA

- Developed a multi-stage reranking system for repository-level code search, leveraging commit histories of popular Github projects, leading to 2x improvement over the widely used BM25 approach.

Technical University of Munich

June 2022 – Aug. 2022

Research Intern (Advisor: Dr. Felix Dietrich) | [talk](#), [report](#)

Munich, Germany

- New approach learn Stochastic Differential Equations (SDEs) for aerosol cloud growth using neural networks, Blender simulations, U-Net preprocessing, and a custom loss function, resulting in an 80% less MSE compared to the baseline.

Technical Projects

Deep Learning from Scratch

March 2024

- Implemented various deep learning fundamentals from scratch in numpy namely - MLP (SGD+Backprop+Batch norm), CNNs (conv/pooling/sampling), GRU with Beam Search, transformers and autograd engine.
- Applied deep learning to several scenarios such as speech recognition (with MLPs, GRUs and encoder-decoder models), image classification/verification (ResNets & ConvNeXts), & image generation (GANs, VAEs, & diffusion models).
- Pretrained a 110M GPT-style Language Model on OpenWebText dataset and supervised finetuned on CNN and Daily Mail, Stanford QA datasets and more.

QryEval - Search Engine from Scratch

Dec 2023

- Implemented BM25 & Indri rankers to search over 500K documents, and experimented with various multi-stage reranking pipelines using SVMRank (LTR), BERT & dense rerankers, to achieve 50% higher MAP/P@10 over BM25 baseline.

Refpred - A Literature Recommender (Bachelor's Thesis) | [github](#), [report](#)

July 2023

- Made a research paper search system leveraging semantics and citation history, featuring an async crawler data collection, SPECTER embeddings and neural reranking KNN candidates (+70% F1).

Skills

Languages: Python, C/C++, Javascript | **Databases:** PostgreSQL, MongoDB | **Frameworks:** PyTorch, React, GraphQL, Flask, Numpy, Pandas, Docker, AsyncIO, ExpressJS, Git, HF (transformers, trl), Slurm, W&B

Awards & Achievements

(2019-23) Merit Scholarship & Award: For being consistently ranked in the top 5 of CS Department.

(2022) DAAD WISE Scholarship: One of 130 selected nationally to pursue a fully-funded internship (by the German Government) at the Technical University of Munich, the top-ranked CS institute in Germany.