VIDHI JAIN

Pittsburgh, PA | Mobile: 412.954.8216 | ✉vidhij2@cs.cmu.edu | Google Scholar
🔗vidhij.com/ | linkedin.com/in/vidhijain23/ | github.com/vidhsss/

EDUCATION

Carnegie Mellon University **Pittsburgh, PA, USA**
Master's in Machine Learning Dec 2025
- **Current Coursework:** Advanced Machine Learning (PhD), Statistics and Probability, Deep RL, Probabilistic Graphical Models
- **Teaching Assistant** for Deep Learning systems under the guidance of **Dr. Tianqi Chen** and **Dr. Tim Dettmers**

Netaji Subhas University of Technology (formerly NSIT/DIT) **New Delhi, India**
B.Tech. in Electrical Engineering with a minor in Machine Intelligence May 2023
Relevant Coursework: Neural Networks & Fuzzy Logic, Soft Computing, Big Data, Data Warehouse, Data Mining, ML in Practice

EXPERIENCE

Carnegie Mellon University - AI SDM
Research Assistant - Professor Aarti Singh [Github] Sep 2024 - Present
- **Cognitive Alignment & Messaging:** Introduced **Integrated Cognitive-Persona Alignment Framework (ICPAF** to align LLM reasoning with human cognition**)**; modeled latent **personas** (CRT, moral foundations, etc.) via **UMAP + HDBSCAN**; studied message framing effects on **beliefs/intentions**; proposed novel metrics: **Cognitive Alignment Score/ Dual-Process Balance/ ASI**.
- **Maternal Healthcare Agent:** Built multilingual, safety-critical **RAG** (**HNSW indexing**, hybrid semantic chunking, **BM25 + dense + cross-encoder**), delivering **+18% retrieval relevance** and **–25% latency** via **vLLM**; added emergency-aware reasoning (intent classification, symptom triage, guardrails with distilled LLMs). Created specific evaluation criteria and trained LLM-as-judge (**κ=0.92**).

Tempus AI
Machine Learning Summer Associate - Applied AI and Research Jun 2025 - Aug 2025
- Architected **agentic clinical LLM** with structured tool use (**planner → router → function-calling**) for diagnosis/medication extraction from unstructured EHR (retrieval, normalization, **citation-matching**), reaching **96% micro-F1**.
- Added explainable decision framework (logic chains, **provenance** tracking, **citation verification**, document-type prioritization), achieving **100% coverage**, **93% faithfulness**, and **–25% latency**.
- Led large-scale **evaluation** (10k+ samples, multi-diagnosis cohorts) combining ground-truth comparison with rubric-driven **LLM-as-a-judge**; built context-aware, reproducible scoring for **clinical safety** and **reasoning quality**.

Wells Fargo
Program Associate- Foreign Exchange Team - **Wells Spot Award** Jul 2023 - Jul 2024
- Developed **Spring Boot microservices** and Angular UIs, optimizing backend performance and reducing API response time by **28%**.
- Automated end-to-end testing using **TestNG and Selenium**, enhancing code reliability and operational efficiency by **32%**.

Inria France
Research Intern - Mnemosyne Lab [Github] Jan 2023 - Mar 2023
- Built a **speech analytics pipeline** to analyze 400+ hours of birdsong audio using **MFCC, STFT, and syllable-level clustering**; modeled effects of sleep deprivation on vocal behavior.
- Applied **unsupervised learning (DTW, HDBSCAN**) to achieve **87% clustering consistency** and proposed **bio-inspired representations** for low-resource speech systems with parallels to human speech and cognitive plasticity

Wells Fargo
Intern Analyst - **Pre-Placement Offer, White Paper** May 2022 - Jul 2022
- Built an AI-driven multi-stage applicant scoring pipeline using **NER, BERT** embeddings**, and TF-IDF** for semantic matching, reducing screening time by **60% (F1: 0.88, MRR: 0.82)**, and deployed via **Dockerized REST APIs on Kubernetes** with A/B testing.

Google ExploreCS Research
Research Intern (IIIT), - Brain, Cognition and Computation Lab [Github] [PDF] Apr 2022 - Jun 2022
- Advanced **human visual attention** modeling using **PredNet,** achieving **23% improvement** over baseline **saliency predictions.**
- Investigated AI-visual human **perception gaps** in **self-supervised learning** systems **(SimCLR, BYOL, etc),** developing **StyleGAN-based** data augmentation techniques that **improved model robustness by 18%,** mitigating perceptual phenomena.

SELECTED PUBLICATIONS

1. ***V Jain**** *et al*, Mind the Gap: Bridging AI-Human Cognitive Misalignment Through Psychological Personas, **AAAI** [in review]
2. *J. Xu*,* ***V. Jain *,*** Contextual Bandits with Online Arm Generation, **NeurIPS** [in review]
3. ***V Jain**** *et al*, Ambient Intelligence-based multimodal human action recognition for autonomous systems, ***ISA Transactions*** [Paper]
4. ***V Jain**** *et al*, Detecting Abnormal Activity in Daily Living: A Deep Learning Approach with RAT-CNN, ***IJHCI*** [Paper]

PROJECTS

Simulating Human-Like Theory of Mind for Social Agents in Visual Environments **Ongoing**
- Building a **visual ToM module** that infers per-agent beliefs/visibility/intent from **gaze, body orientation, and occlusions.**
- Training an **RL policy conditioned on mental states** (not just physical states) to imitate human-like, socially compliant behavior; evaluated via task success, false-belief handling, and plan legibility metrics.

Neurascribe AI: [Link] **Spring 2025**
- Voice-to-voice journaling/wellness platform with **hybrid memory** (Pinecone vectors + knowledge graph) and **episodic recall**.
- **Multi-stage reranking** (semantic + cross-encoder) and **LoRA-tuned LLMs** for emotional inference; deployed as **FastAPI microservices** with LangChain orchestration; **vLLM + async I/O** cut latency ~**40%**.

Cognitive Multimodal Speech Recognition in Noisy Environments [Github] **Spring 2025**
- Built multimodal pipeline combining **vision (Transformer lip-reading, ViT-B/16)** and **self-supervised audio (HuBERT, wav2vec 2.0)**.
- Applied **knowledge distillation** to compress large multimodal models, reducing size by 60% with minimal WER/CER drop.

SKILLS

Languages & Tools: Python, C++, Java, MATLAB, FastAPI, Flask, Spring Boot, Google Cloud, Docker, Kubernetes, Git
Libraries & Frameworks: PyTorch, TensorFlow, JAX, OpenCV, Scikit-learn, Keras, HuggingFace, LangChain, FAISS, Selenium, TestNG