Mann Ahalpara

AI/ML Research | Conversational AI | LLMs | Information Retrieval

Ahmedabad, India | mannahalpara@gmail.com | +91 99044 45765
LinkedIn | GitHub

Research Profile

Final-year B.Tech student in Information & Communication Technology at Adani University, ranked 1st in the ICT Department (CGPA: 9.13/10). Currently a Research Intern with Prof. Erik Cambria's Sentic Team at Nanyang Technological University (NTU), Singapore, working on privacy-preserving long-term Emotional Support Conversation systems, with a focus on conversational memory, memory retrieval, personalization, and memory-guided reasoning in LLM agents. Co-author of an ICIDS 2025 conference paper, with hands-on experience building and deploying retrieval-augmented generation, semantic evaluation, and machine-learning systems.

Research Interests

Large Language Models; Conversational AI; Long-Term Memory in AI Agents; Retrieval-Augmented Generation and Information Retrieval; Agentic Systems and LLM Reasoning; Personalized and Affective AI; Trustworthy, Privacy-Preserving, and On-Device AI.

Education

B.Tech in Information & Communication Technology **2023 – 2027 (Expected)**
Adani University, Ahmedabad, India
CGPA: 9.13/10 | **Ranked 1st in the ICT Department**

Class XII (HSC) – Science Stream **2023**
Gujarat Secondary & Higher Secondary Education Board
Percentile: 97.03

Research Experience

Research Intern (Remote) – Sentic Team, NTU Singapore **May 2026 – Present**
Long-Term Conversational AI and Memory Systems Mentor: Tiesunlong Shen | PI: Prof. Erik Cambria

- Conduct research on privacy-preserving long-term Emotional Support Conversation (ESC) systems, studying how conversational agents can retain relevant information and provide consistent, personalized support across extended, multi-session interactions.
- Investigate hierarchical and agentic memory architectures for storing, retrieving, updating, and selectively using relevant information across conversational sessions.
- Experiment with memory-conditioned AFlow pipelines and Monte Carlo Tree Search (MCTS) using local language models to study contextual continuity and memory-guided reasoning.
- Explore on-device and privacy-preserving approaches for persistent agents, and review recent work on personalized AI agents, affective computing, and retrieval-augmented systems to identify open research gaps.

Publication

Analyzing Indian Stock Markets through Correlations: Comparative Insights **ICIDS 2025**
Proceedings of the International Conference on Infrastructure Development and Sustainability (ICIDS 2025).

- Co-authored a comparative study of relationships across ten major Indian market sectors using Pearson, Spearman, Kendall's Tau, and time-lagged cross-correlation techniques.
- Applied correlation analysis, network representations, and comparative visualization to study inter-sector dependencies, diversification opportunities, and evolving market behavior.
- Extended the research into CorrelationX, an interactive analytics platform for multi-method correlation analysis and sector-relationship visualization.

Projects

AnswerIQ – AI-Powered Answer Evaluation System GitHub
OCR, Sentence Transformers, Gemini, FastAPI, Next.js/TypeScript, Supabase

- Built an answer-sheet evaluation platform that automates the assessment of student responses using OCR, semantic analysis, and LLM-assisted scoring.
- Designed a hybrid, explainable scoring framework combining concept coverage (50%), semantic similarity (30%), and keyword matching (20%) to produce structured, interpretable scores.
- Engineered question-normalization and answer-mapping pipelines that align answers with questions despite inconsistent numbering formats (e.g., Q1, Q-1, Question 1) and varying submission order.
- Developed role-based faculty and student workflows with a FastAPI backend, Supabase data management, and a

Next.js/TypeScript frontend.

Mini-RAG – Citation-Grounded Retrieval-Augmented Generation GitHub | Live Demo

Gemini Embeddings, Qdrant Cloud, FastAPI, Render, Vercel

- Designed and deployed an end-to-end RAG system for querying uploaded PDF and text documents, generating answers grounded in retrieved evidence with inline citations.
- Implemented document ingestion, text extraction, character-based chunking, metadata storage, Gemini embeddings, and top-k cosine-similarity retrieval using Qdrant Cloud.
- Added a lightweight reranking stage to re-score retrieved candidates and filter low-relevance chunks before grounded generation.
- Built FastAPI services deployed on Render with a Vercel frontend, keeping API credentials and vector-database operations server-side.

AirVision – AQI Prediction & Analytics Platform GitHub | Live App

Python, Scikit-learn, XGBoost, LightGBM, CatBoost, Streamlit

- Built a machine-learning pipeline for predicting the Air Quality Index (AQI) from pollutant measurements collected across multiple Indian cities.
- Benchmarked 11 regression algorithms, including Random Forest, XGBoost, LightGBM, CatBoost, SVR, KNN, and Gradient Boosting; CatBoost achieved the strongest validation performance ($R^2 = 0.91$, MAE = 12, RMSE = 19).
- Implemented missing-value imputation, date-based feature extraction, categorical encoding, feature scaling, and an interactive Streamlit dashboard for prediction and model comparison.

AXON – AI-Powered Desktop Assistant GitHub

Python, Gemini, LLaMA, Speech Recognition and Synthesis

- Developed a voice- and text-controlled desktop assistant supporting system commands, application launching, web search, and conversational question answering.
- Designed a dual-LLM architecture routing command interpretation and screen analysis to Gemini and open-ended conversation to LLaMA, with fuzzy command matching for robustness to speech-recognition errors.

Additional Experience

AI & ML Intern **Jan 2026 – Feb 2026**

Edunet Foundation (AICTE × IBM SkillsBuild)

- Developed and deployed Study Buddy, an AI learning assistant using the Gemini API and Streamlit for explanations, summarization, quiz and flashcard generation, and PDF-based learning, built on document-processing and prompt-engineering workflows.

Technical Skills

Programming	Python, SQL, Java
Machine Learning	JAX, Scikit-learn, PyTorch, Regression, Classification, Feature Engineering, Model Evaluation, Statistical Analysis
LLMs & NLP	Transformers, Gemini, LLaMA, Ollama, Prompt Engineering, AI Agents, Conversational AI, Conversational Memory
Retrieval Systems	Retrieval-Augmented Generation, Sentence Transformers, Embeddings, Semantic Search, Vector Search, Reranking, Qdrant
Frameworks	FastAPI, Flask, Streamlit, Next.js, Pandas, NumPy, NetworkX, OSMnx
Databases & Tools	PostgreSQL, MySQL, Supabase, Qdrant, Git, Docker, Render, Vercel
Visualization	Matplotlib, Seaborn, Power BI, Tableau

Achievements & Certifications

- Represented Adani University at National Technology Day 2024, hosted by SAC-ISRO, Ahmedabad, selected based on academic performance.
- Solved 600+ Data Structures and Algorithms problems on LeetCode.
- Completed NVIDIA Deep Learning Institute training in Retrieval-Augmented Generation (RAG).
- Completed Google Quantum AI's hands-on course on Quantum Error Correction.