

Arjun Balaji

✉ ab6136@columbia.edu ☎ +1 9294938642 🔗 www.arjunbalaji.in 🌐 ArjunBalaji79

EDUCATION

Columbia University, Masters in Public Administration Aug 2025 – Present | New York City
Concentration in Technology Policy and Innovation

R.V. College of Engineering, BE in Electronics and Telecomm. Dec 2021 – Jun 2025 | Bangalore
Coursework: Multivariable Calculus, Differential Equations and Numerical Methods, Machine Learning, Statistics and Probability for Data Science, Linear Algebra, Integral Transforms and Fourier Series, Deep Learning, Signal Processing with Machine Learning

PUBLICATIONS

- "Sparse Spectral Signatures of Reasoning: Model-Agnostic Verification via Sentence-Level Graph Signals." **ICLR 2026 Workshop on Logical Reasoning of Large Language Models**
- "Encoding Values: Injecting Morality into Machines via Prompt-Conditioned Moral Frames." **NeurIPS 2025 Creative AI Track** (<https://neurips.cc/virtual/2025/loc/san-diego/129264> 🔗)
- "BiasCheck: An Analytical Framework for Contextual Bias Detection in Text, Models and Databases." **AAAI 2025 Workshop on AI for Public Missions**
- "ProteoDockNet: A deep learning graph neural network-based architecture for the prediction of ligand binding affinities via SMILES to key human proteins of liver, kidney and brain using QSAR data." **Elsevier Computational and Structural Biotechnology Reports, 2024 (Q1 Journal)**
<https://doi.org/10.1016/j.csbr.2024.100011> 🔗
- Graph Neural Network Framework for Web-Based Prediction of Protein-Ligand Docking Scores across multiple organs. **Protocols.io, 2023.** (<https://dx.doi.org/10.17504/protocols.io.j8nlkoy9xv5r/v1> 🔗)
- Presented "Developing a Comprehensive AI-Driven Decision Support System for Colorectal Cancer: Integration of Mutation Data and Web Application Implementation" at **ICCPSTMLA-24, New York**

RESEARCH EXPERIENCE

Carnegie Mellon University, Research Intern Sep 2024 – Mar 2025 | Pittsburgh

- Developed Unsupervised Models for Brain MRI Surface Registration in a collaborative project with **Harvard Medical School** at the Xu Lab.
- Designed and Implemented efficient geometric deep learning techniques for deforming a mesh template on to the brain surface.

Nanyang Technological University, Research Intern Oct 2024 – Jun 2025 | Singapore

- Researched advanced segmentation for MRI and echocardiograms at **MVAIT** (Medical Vision and Artificial Intelligence Technologies) Lab under the guidance of Prof. Si Yong.
- Built a dynamic 4D MRI memory bank with slice-level adaptive quantization (8/16/32-bit) for compact storage and novel rapid frame retrieval techniques.
- Designed and benchmarked multi-objective loss functions for custom architecture (Dice, Wasserstein, Cross-Entropy, Normal Consistency, Tversky, etc.) for UNet/CNN cardiac segmentation.

Indian Institute of Science, Research Intern Oct 2024 – Dec 2024 | Bangalore

- Developed graph sampling techniques and anomaly detection algorithms for detecting money laundering in financial transactions in collaboration with the National Payments Corporation of India (**NPCI**) at the VISTA Lab.

Samsung Research Institute, Research Intern Aug 2023 – Aug 2024 | Bangalore

- Researched and applied memory-efficient techniques, including Knowledge Distillation and Quantization, to optimize Vision Transformers (ViTs) for resource-constrained environments.
- Developed lightweight architectures achieving up to 8x memory compression with SOTA accuracy.

RVCE Centre for Computational Genomics, AI/ML Intern Nov 2022 – May 2024 | Bangalore

- Led a Consultancy project for RapidCanvas.ai, USA, implementing and deploying ML-based QSAR models on the RapidCanvas Platform.

- Developed and deployed a comprehensive Graph Neural Network Framework (60 models) for Docking Score Prediction across 3 organs, 10 proteins each on Streamlit and published this work in a Q1 Journal (Link: <https://bit.ly/41gOOVR>)

PROFESSIONAL EXPERIENCE

Incluya, *Founding AI Engineer*

Apr 2025 – Jul 2025

- Designed a personalized financial wellness AI using LLMs that adapt to user traits and emotional triggers.
- Integrated memory layers to retain evolving user profiles and improve long-term personalization.
- Built a dual-retriever system combining vector search and rule-based filters for targeted intervention recommendations.

People+ai, *AI/ML Developer*


Jun 2024 – Mar 2025 | Bangalore

- Developed an application using Google Maps API and GRIHA guidelines to dynamically evaluate proximity to essential services to improve green building planning which has been publicly released and shared to both the Indian Government's GRIHA.
- Developed a multimodal LLM framework to automate and optimize the GRIHA Green Building Certification Process from sitemaps and documents.
- Utilized Stable Diffusion with ControlNet to generate synthetic site map data, leveraging edge maps and layout conditions to create precise and diverse datasets for training AI models.

PROJECTS


Encoding Values

Jul 2025

- Developed Encoding Values, a Python research toolkit to inject explicit ethical value sets (e.g. feminist, ecological, decolonial) into large language models via system-prompt conditioning.
- Enables evaluation of how moral framings shift model behavior measuring effects on tone, framing, inclusivity, semantic similarity, toxicity, and alignment through reproducible metrics across different value-oriented prompts. Link: <http://bit.ly/46IfQb3> 

BiasCheck

May 2025

- Developed BiasCheck, a modular Python library designed to analyze and detect bias in text, models, RAG frameworks and datasets.
- It provides tools for researchers, data scientists, and developers to measure various forms of contextual bias (region-specific terms, etc.) and assess the quality of language model outputs or textual data. Link: <https://bit.ly/41V566E> 

TECHNICAL SKILLS

Languages: Python, C/C++, SQL, MATLAB

Tools: Power BI, VS Code, Git/Github, Streamlit, Apache Spark

Machine Learning: TensorFlow, PyTorch, Torch Geometric, Torch Vision, XGBoost, LangChain, Scikit-Learn

Databases: Neo4j, ChromaDB, MySQL, PostgreSQL

RECOGNITIONS

President of RV Debating Society

Leading one of the oldest and most prestigious debating societies in India, mentoring and coordinating with over 200+ members, and organized one of South India's largest debating tournaments.

LossFunk AI Resident (Batch 2)

(Formerly Turing's Dream) Chosen as one of 10 AI researchers and founders across India for an AI residency program run by Paras Chopra.