

VENKAT SUPRABATH BITRA

☎ +1 (646) 537 0987 ✉ vsh2127@columbia.edu in [linkedin.com/in/venkat-suprabath](https://www.linkedin.com/in/venkat-suprabath) 🌐 venkatsbitra.github.io

Education

Columbia University, USA

Sep 2026 – Present

PhD in Electrical Engineering

Columbia University, USA

Sep 2024 – Dec 2025

MS in Computer Science

GPA: 4.231

IIIT Bangalore, India

Aug 2019 – Jul 2024

Integrated M.Tech. (B.Tech. + M.Tech.) in Computer Science and Engineering

GPA: 3.85 / 4.00

Experience

Columbia University

Jan 2026 – Aug 2026

Research Assistant

New York, NY

- Designed a **Frequency Convolution Network (FCN)** using dilated 1D convolutions along the log-frequency axis of a **Variable-Q Transform (VQT)** frontend, achieving a large harmonic receptive field with only **17.8K parameters** for real-time monophonic pitch and voicing estimation.
- Developed **spectral-domain noise and frequency-scaling augmentation** strategies to induce noise invariance and pitch equivariance, evaluated on MDB-stem-synth, PTDB-TUG, and MIR-1K with CHiME-Home as a noise source.
- Achieved **98.8% RPA50 on MDB-stem-synth, 93.1% on PTDB-TUG, and 97.5% on MIR-1K** in clean conditions while keeping RPA50 above **80% at 0 dB SNR** across datasets, outperforming a matched-parameter CNN baseline and remaining competitive with substantially larger pitch-estimation models.
- Proposed **Counterpoint Informed Neural Networks (CINN)**, embedding Fux's species counterpoint rules as a differentiable, physics-informed penalty on a **MinGRU encoder-decoder** that generates a soprano accompaniment against a cantus firmus.
- Introduced a **joint pitch-attack coupling loss** (via Jensen-Shannon divergence between neighboring pitch distributions) and an attack-prediction head to keep rhythms and articulations plausible, preventing collapse into rests or rapid spurious attacks.
- Implemented a **joint beam search over pitch and attack** at inference to keep rhythm aligned with the cantus firmus, benchmarking against a MinGRU baseline and a Transformer variant on a held-out composer, **reducing rule-based violations** across harmonic, melodic, and rhythmic categories with minimal cross-entropy trade-off.

Columbia University

Jan 2025 – Dec 2025

Master's Thesis Project

New York, NY

- Developed a **Variational Autoencoder (VAE)**-based Audio Encoder with a **Mamba-enhanced Differentiable Digital Signal Processing (DDSP) Decoder** for waveform analysis, enabling music synthesis and timbre transfer.
- Explored **minimal Recurrent Neural Networks (minLSTMs and minGRUs)**, which significantly improve sequence modeling performance and efficiency by leveraging fewer parameters.
- Developed **Bidirectional Minimal GRU (BiMinGRU)** to effectively incorporate past and future context in the sequence model, removing the reliance on causal convolutions for multi-level architectural designs.
- Implemented a **self-supervised Expectation-Maximisation (EM) based learning approach** to enhance fundamental frequency (F0) and voicing estimation, capitalizing on the co-dependent relationship between existence of fundamental frequency and voicing to filter unvoiced frames and preventing bad gradient updates.
- Engineered a highly parameter-efficient **frame-level F0 and voicing extraction model with only ~9K parameters** using Constant Q-Transform (CQT) features for lightweight deployment.
- Demonstrated a **fully end-to-end trainable DDSP-based music synthesis and timbre transfer system** that accepts any reference monophonic audio, removing dependence on external F0 or loudness models.

Complex Resilient Intelligent Systems (CRIS) Lab, Columbia University

Sep 2024 – Present

Research Assistant

New York, NY

- Investigated interpretability of LLMs (GPT-2, Gemma 2B Instruct) by analyzing hidden representations using **dictionary learning and sparse autoencoders** to uncover monosemantic features.
- Implemented **feature clamping approach** (by Anthropic) to steer model outputs toward specific features, improving controllability and interpretability.
- Applied **Sparse Autoencoders (SAEs) and BERTopic to 500 books across 20 genres**, extracting monosemantic features from dense embeddings at multiple granularities to map semantic structure in decoder-based LLMs.
- Uncovered unexpected semantic relationships across genres through **centroid analysis**, validated by narrative archetype theory.

- Leading team to investigate **emergent scientific understanding in LLMs** by training domain-specific and shared-domain SAEs on physics, chemistry and biology textbooks, analyzing how models represent concepts (e.g., atomic structure) differently across disciplines while identifying shared fundamental features.
- Developing **dead neuron mitigation techniques and hierarchical architectures in SAEs** to control learned feature granularity, reducing reconstruction loss while improving interpretability and sparsity.

Accessible and Accelerated Robotics (A²R) Lab, Columbia University

Sep 2024 – May 2025

Research Assistant

New York, NY

- Optimized and implemented **MobileNet and PyDNet depth estimation models onto resource-constrained microcontrollers** (Crazyflie, ESP32), achieving real-time, on-device depth perception for autonomous navigation.
- Deployed **YOLOv11 for high-speed, real-time object detection in low-power embedded systems**, enhancing autonomous target navigation and obstacle avoidance.
- Developed a **custom lightweight framework leveraging Teacher-Student Distillation** to compress deep learning models into a low-parameter configuration, ensuring reliable operation across diverse microcontroller hardware.

Digital Video and Multimedia (DVMM) Lab, Columbia University

Sep 2024 – Feb 2025

Research Assistant

New York, NY

- Developed and implemented a **novel Vision Transformer approach for Causal Action Understanding**, predicting both the action and its resulting object state changes by analyzing start/end frames and bounding box localizations.
- Developed an **advanced data curation pipeline integrating LLaMA 3.1 via few-shot prompting** to resolve issues like object naming inconsistencies and improve action extraction accuracy from raw Ego4D and YouCook2 narrations.
- Resolved a critical temporal misalignment in the Ego4D dataset by developing and implementing a **binary search algorithm to efficiently map narrations to corresponding video clips**, ensuring data coherence.
- Established performance benchmarks by **training and evaluating multimodal baselines (Clip, BridgeFormer, Stateformer, ACTRON)** on the enhanced dataset, comparing action-prediction and state-change performance.

Laser Materials Processing Division, RRCAT

Jul 2024 – Aug 2024

Research Intern

Indore, IN

- Adapted the **YOLO object detection framework for generalized spectra classification**, achieving **mAP >0.992 for 90 spectra classes** and eliminating the traditional need for extensive data preprocessing.
- Introduced a **novel SVG-based method for automated generation of diverse, annotated spectral image datasets** with varied configurations and backgrounds, overcoming limitations of traditional data creation.
- Evaluated the integration of **attention mechanisms (CBAM and SE) into the YOLO architecture** to further enhance feature extraction and improve classification performance in spectral image analysis.
- Conducted **perturbation analysis on the trained YOLO model** to understand and validate how the network was classifying spectra, providing interpretability for the detection mechanism.
- Successfully deployed trained YOLO models on smartphones for **rapid, real-time spectra classification**, benchmarking YOLOv8m and YOLOv8n while outperforming traditional 1D CNNs in efficiency.

International Institute of Information Technology Bangalore

Jan 2024 – Jun 2024

Master's Thesis Project

Bangalore, IN

- Designed and implemented a **Panoptic Segmentation Guided Moving Object Segmentation (MOS) framework** to correct false positives in the state-of-the-art (SOTA) RigidMask model, improving performance for SLAM.
- Localized the primary cause of SOTA model failure to **anomalies (spikes) in 3D P+P cost maps** for static objects near frame edges or under occlusion on the challenging VKITTI dataset.
- Integrated a **top-down information fusion pipeline combining panoptic segmentation (PanFPN), monocular depth estimation (Metric3D), and dense optical flow (RAFT)** as a robust post-processing correction step.
- Developed an **IoU-based panoptic instance tracking algorithm** and used dense optical flow matching to estimate per-object motion in depth, achieving higher reliability than SIFT for objects 10m-30m away.
- Engineered a **final correction layer using 3D P+P cost histogram analysis** to filter residual false positives caused by occlusions or close-range objects, cases where depth-based motion estimation fails.
- Achieved substantial quantitative improvements across the VKITTI dataset, including **increasing Background IoU to over 99%** in tested scenes (e.g., from 87.67 to 99.41 in Scene 01).
- Improved the system's robustness by introducing a **Reliability Score metric** to penalize false positives in frames with no moving objects, with scores reaching 100.0 in scenes where the baseline had high false positive rates.

Statistics and Machine Learning Laboratory, UBC Okanagan Campus

Aug 2023 – Dec 2023

Research Intern

Kelowna, CA

- Developed and implemented a **Finite Gaussian Mixture Model (FGMM) with Linear Spatial Covariance structure** to perform unsupervised classification of high-dimensional spatial data (MC-simulated radiation dose grids).

- Engineered a **highly efficient covariance parameterization with only four free parameters**, making the model scalable regardless of the number of pixels (p).
- Applied an **Expectation-Maximization (EM) algorithm to classify biomedical MC simulation data** (e.g., 10 mGy vs. 500 mGy dose levels) based on differences in spatial patterns (covariance) rather than global intensity, confirming that higher dose levels exhibit lower spatial variance.
- Evaluated two decay functions (Quadratic and Sigmoid) for the spatial covariance term, finding the **Sigmoid decay better approximates complex spatial patterns** (“many bumps” vs. “single bump”) in radiation data.

Adobe Research

May 2023 – Jul 2023

Research Intern

Bangalore, IN

- Engineered and optimized **diffusion models with multimodal constraints (text and visual inputs)** to gain granular control over image synthesis, accelerating and streamlining the product design creation workflow.
- Produced **visually cohesive, marketing-ready images** by leveraging the constrained generative models, ensuring strict adherence to brand references and style guidelines to drive visual consistency across product assets.
- Executed **advanced multimodal image generation techniques** to enhance control and quality, eliminating manual design iterations and significantly improving the efficiency of creating branded visual assets.

Multimodal Perception Lab, IIITB

May 2022 – May 2023

Research Intern

Bangalore, IN

- Investigated and benchmarked **state-of-the-art GAN-based models for audio-video generation and style transfer**, focusing on applications like realistic lip synchronization (e.g., Wav2Lip) and semantic face editing.
- Implemented and optimized **various published GAN architectures for face manipulation** (e.g., InterFaceGAN, and models similar to Only a Matter of Style for age regression/transformation), and performed comprehensive testing and comparative analysis using custom datasets.
- Combined traditional computer graphics techniques with deep learning by **incorporating the FLAME 3D head model into GAN workflows**, which improved how realistic and structurally consistent the face edits looked, especially when moving beyond flat 2D images.
- Developed **SemUV, a deep learning system that generates and edits 3D human heads by working directly with UV texture maps using StyleGANv2**, bypassing traditional 2D image manipulation limits.
- Demonstrated **improved identity preservation and structural consistency compared to 2D methods through user study validation**, successfully editing facial features (age, gender, facial hair) while maintaining realistic appearance across different 3D poses.
- Built a **flexible and efficient solution that works independently of 3D structure, lighting, and rendering settings**, allowing quick integration into existing graphics pipelines for VR, AR, gaming, and visual effects applications.

Laser Materials Processing Division, RRCAT

Dec 2022 – Jan 2023

Research Intern

Indore, IN

- Optimized and deployed an **8-bit quantized CNN onto resource-constrained TinyML hardware (Arduino Nano 33 BLE)** to achieve accurate classification of 22 complex spectral classes.
- Achieved **>99.5% classification accuracy on the quantized model** by demonstrating the critical role of data normalization, enabling high-precision spectral analysis for field-deployable analytical instrumentation.
- Designed and implemented the **complete spectral analysis pipeline on-device (preprocessing, ML classification, and analyte quantification)** to ensure real-time, autonomous decision-making without external computational or cloud resources.

Laser Materials Processing Division, RRCAT

Dec 2021 – Jan 2022

Research Intern

Indore, IN

- Developed an **innovative, low-cost automated system (Arduino UNO & CNC shield)** that enabled high-throughput spectral data acquisition (up to 250 spectra/hour), generating large labeled datasets for ML.
- Implemented a **novel raster scanning technique and vision-based automatic sample focus distance control**, improving SERS reproducibility via dynamic 3D hotspots and enabling **detection down to a 5 nM limit**.
- Created an **interactive Principal Component Analysis (PCA) GUI** for accessible data exploration, allowing non-experts to dynamically evaluate the impact of normalization methods on spectral data clustering.
- Developed a **novel SERS approach integrating machine learning (PCA & LDA)** to achieve **>97% accuracy in identifying complex pesticide mixtures**, addressing a gap in current analytical methods.
- Conducted ML classification and performance evaluation (SVM, kNN, etc.) across diverse datasets, employing **Recursive Feature Elimination (RFE) for optimal feature selection** for robust classification with limited spectral data.

- Developed and implemented **two novel phenotype-specific filtering methods (network-based and correlation-based)** to rank and select significant genes, reducing the gene space to an optimal size for visualization.
- Constructed a **diseasome (disease-gene network) as an unbalanced bipartite graph** by integrating multi-omics data (DNA methylation and mRNA expression) for seven cancer phenotypes from TCGA (The Cancer Genome Atlas) using the Heterogeneous Correlation Network Model (HCNM).
- Designed a systematic **“filter-funnel-focus” workflow for gene prioritization** to ensure clutter-free visualization and focused analysis of key genes (e.g., 100+ genes).
- Conducted extensive downstream analysis, including **pathway enrichment (WikiPathways, Reactome, and KEGG) and functional enrichment (DAVID)**, to validate biological relevance and compare the performance of the two filtering methods across cancer types.
- Developed **RadTrixVis, an interactive web-based visualization tool using RadTrix, a novel hybrid graph layout**, to display complex disease-gene networks. The layout places a matrix at the center to dramatically reduce visual clutter and edge crossings when showing relationships between diseases and genes.
- Integrated an **intuitive GUI with interactive features** (e.g., phenotype filtering, gene selection, color customization, node reordering) to facilitate disease-specific and gene-specific analysis.

Web Science Laboratory, IIITB

Apr 2020 – Jun 2020

Research Intern

Bangalore, IN

- Engineered and deployed a **full-stack, interactive COVID-19 trend simulator using React, Leaflet, and Recharts** for frontend and a **Python/Flask REST API backend** to efficiently generate and visualize predictions.
- Validated and refined the **core SEIR epidemic model implementation**, ensuring the correct integration of complex network diffusion and community interaction theory into the production code for accurate trend simulation.
- Executed comprehensive testing and troubleshooting of the parameter system and prediction generation logic, guaranteeing **reliable simulation results across 10+ user-controlled variables** (e.g., mobility, school closures).

Projects

Linear Regression Accelerator for FPGA (Embedded Systems Course Project) | C, Verilog, Quartus Jan - May 2025

- Designed a memory-mapped 4-bit linear regression accelerator on Intel FPGA with a pipelined datapath for high-throughput accumulation and closed-form inference
- Achieved near-floating-point performance ($R^2 = 0.949$) using quantized integer arithmetic in hardware
- Wrote a custom C-based Linux device driver and interface to stream training data via AXI-lite to ARM Cortex-A9

ArcCell: RNA-seq Cell Type Annotation Tool (Computational Genomics Course Project) | Python Jan - May 2025

- Developed a robust cell type annotation tool for sparse scRNA-seq data using ArcFace-based supervised embeddings
- Implemented a preprocessing pipeline with GF-ICF weighting and over/under-sampling for rare cell types
- Achieved 96% F1-score on test data and 94% on out-of-sample validation, outperforming CellTypist baseline (62%)
- Integrated SHAP-based interpretability for gene attribution, finding biologically relevant marker genes across datasets

RAG based LLM Chatbot | Python, LangChain, Ollama

Nov - Dec 2024

- Developed a Retrieval-Augmented Generation (RAG) chatbot using LangChain and Ollama for enhanced context-aware responses
- Implemented vector embedding and efficient document retrieval to improve answer relevance and accuracy
- Integrated local LLM deployment with Ollama, reducing latency and ensuring data privacy

Reinforcement Learning based Chatbot | Python, PyTorch, TensorFlow

Jan - May 2023

- Design two chatbots (Agent A and Agent B) using policy gradient based training strategies so that they can handle prolonged meaningful natural conversations between them
- Implement training techniques to give two separate personas for Agent A and Agent B.

Abacus based Mental Calculation using Neuromorphic Architecture | Python, Nengo

Aug - Dec 2022

- Implementation of different abacus based arithmetic operations to mimic mental calculations using Neural Engineering Framework.

Oceanic Data Visualization (Data Visualization Course Project) | Python, Svelte, Flask, D3.js

Aug 2022

- Designed a web application to visualize various features of Indian Ocean and Arabian Sea during the period of 29 Dec 2003 to 29 Dec 2005.

- 2048 Playing Agent** (Personal Project) | [Python, PyTorch, Jupyter](#) **Jun 2022**
- Designed a reinforcement learning agent to play 2048 using different algorithms DQN, PPO, NEAT and Markov Trees.
- Visual Programming Language** (Programming Languages Course Project) | [JavaScript, Blockly](#) **Jan 2022**
- Web based visual programming language to emulate Turtle using Blockly transpiler for equivalent JS code generation.
- Comprehensive Spectroscopic Data Processing and Visualization Tool** | [Python, React, ECharts](#) **Mar - May 2022**
- Designed a full-stack novel software tool for comprehensive data analysis of batch Raman spectra with features like background correction, peaks analysis, heatmaps, SNV, PCA-SVM analysis, spectral data libraries etc.
 - Utilised React and Apache ECharts for designing frontend and data visualizations of the analysed data provided by backend built using Python and Flask.
 - Embedded various machine learning models using Python and scikit-learn for determination of compounds for applications like identifying pesticides.
- Image Captioning Tool** (Visual Recognition Course Project) | [Python, TensorFlow, Jupyter](#) **May 2022**
- A CNN-LSTM network was implemented using InceptionNetv3 as feature extractor to label various images provided in the publically available Flickr8K dataset.
- Auto Rickshaw Detection** (Visual Recognition Course Project) | [Python, PyTorch, Jupyter](#) **Mar 2022**
- Designed object detection model using YOLOv5 on self curated and labelled set of images of auto rickshaw.
- Web Scraping Tool** (Personal Project) | [Python, Selenium, SQLite](#) **Jan 2022**
- Program to scrape publication DOI and other details from different publisher websites based on user given keywords and record them in a database for ease of querying for researchers.
- NLP based Sentiment Analysis** (Machine Learning Course Project) | [Python, Scikit-Learn, Jupyter](#) **Dec 2021**
- Machine Learning model for multiclass classification of a sentence into six different types of sentimentally hurting classes; hosted as Community Kaggle Competition.
- Face Recognition based Security System** (Hackathon Project) | [PyTorch, Scikit-Learn, OpenCV](#) **Jan 2020**
- Real time image pre-processing and face detection using OpenCV and Python integrated with pretrained Deep Learning model to determine strangers and owners.

Publications

- Lightweight Self-Supervised Detection of Fundamental Frequency and Accurate Probability of Voicing in Monophonic Music** **Mar 2026**
- Venkat Suprabath Bitra, Homayoon Beigi
- 15th International Conference on Pattern Recognition Applications and Methods (ICPRAM 2026) [↗](#)
- Spatial Covariance Constraints for Gaussian Mixture Models** **Jan 2026**
- Hanzhang Lu, Keiran Malott, Venkat Suprabath Bitra, et al. (full author list in paper)
- Submitted to Bioinformatics (Arxiv Preprint [↗](#))
- YOLO-spectra: A generalized framework for rapid simultaneous detection and classification of Raman spectra in images with mobile devices for enhancing on-site applications** **Jan 2026**
- Venkat Suprabath Bitra, Shweta Verma and B. Tirumala Rao
- Analytica Chimica Acta Volume 1383, 15 January 2026, 344914, ISSN 0003-2670 [↗](#)
- HiT-IP Raman: An innovative high-throughput spectra acquisition system with interactive PCA combined machine learning for instant classification of various analytes and mixtures** **Nov 2025**
- Venkat Suprabath Bitra, Shweta Verma and B. Tirumala Rao
- Microchemical Volume 219, December 2025, 116054, 0026-265X [↗](#)
- Network-based Disease Construction from Multi-omics Data and RadTrix Visualization** **Nov 2025**
- Venkat Suprabath Bitra, Reddy Rani Vangimalla, Jaya Sreevalsan-Nair
- IEEE Transactions on Computational Biology and Bioinformatics, vol. 22, no. 6, pp. 3550-3556, Nov-Dec 2025 [↗](#)
- Machine learning driven trace detection of pesticide mixtures using citrate optimized Au nanoparticles based in-expensive efficient micro-drop SERS with portable spectrometer** **Nov 2025**
- Shweta Verma, Venkat Suprabath Bitra, B. Tirumala Rao
- Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy, Volume 340, 2025, 126333, ISSN 1386-1425 [↗](#)

DiffCraft: A Modular Differentiable Framework for Music Synthesis with Timbre Transfer May 2025

Venkat Suprabath Bitra, Homayoon Beigi

IEEE Conference on Artificial Intelligence (CAI 2025) [↗](#)

SemUV: Deep Learning based semantic manipulation over UV texture map of virtual human heads Dec 2024

Anirban Mukherjee, Venkat Suprabath Bitra, Vignesh Bondugula, Tarun Reddy Tallapureddy, Dinesh Babu Jayagopi

9th International Conference on Computer Vision & Image Processing (CVIP 2024) [↗](#)

- Won IAPR Best Student Paper Award

TinyML-Raman: A Novel IoT Based Field-Deployable Spectra Analysis for Accurate Identification of Pharmaceuticals and Trace Dye-Pesticide Mixtures from Facile SERS method Aug 2024

Venkat Suprabath Bitra, Shweta Verma, B. Tirumala Rao

Analytica Chimica Acta, 2024, 343063, ISSN 0003-2670 [↗](#)

Optical response of Au films for reproducible Si nano-structuring and its application for efficient micro-drop SERS with portable Raman system Sep 2023

Shweta Verma, Venkat Suprabath Bitra, R. Singh, B. Tirumala Rao

Materials Chemistry and Physics, Volume 306, 2023, 128058, ISSN 0254-0584 [↗](#)

Machine Learning and Convolutional Neural Network based Trace Identification of Pesticides and Dye Mixtures from Raman Spectra Jan 2023

Venkat Suprabath Bitra, Shweta Verma, B. Tirumala Rao

6th Joint International Conference on Data Science & Management of Data (CODS-COMAD 2023) [↗](#)

Software for Comprehensive Batch Analysis of Raman/SERS Spectra with Real-time Interactive Visualization and Machine Learning Predictions Dec 2022

Venkat Suprabath Bitra

IX International Conference on Perspectives in Vibrational Spectroscopy (ICOPVS-2022) [↗](#)

An Interactive Simulator for COVID-19 Trend Analysis Jan 2021

Jayati Deshmukh, Raksha Pavagada Subbanarasimha, Pooja Bassin, Venkat Suprabath Bitra, Srinath Srinivasa, and

Anupama Sharma

3th Joint International Conference on Data Science & Management of Data (CODS-COMAD 2021) [↗](#)

Mentorship

Classical Control Systems (EEME 3601) <i>Teaching Assistant II</i>	Columbia University, US <i>Sep 2025 – Dec 2025</i>
--	---

Classical Control Systems (EEME 3601) <i>Teaching Assistant II</i>	Columbia University, US <i>Jan 2025 – May 2025</i>
--	---

Linear Algebra (SM 102) <i>Teaching Assistant</i>	IIIT Bangalore, India <i>Jan 2024 – May 2024</i>
---	---

Digital Design (ESS 102) <i>Teaching Assistant</i>	IIIT Bangalore, India <i>Aug 2023 – Dec 2023</i>
--	---

Visual Recognition (AI 825) <i>Teaching Assistant</i>	IIIT Bangalore, India <i>Jan 2023 – May 2023</i>
---	---

Introduction to Machine Learning (AI 511) <i>Teaching Assistant</i>	IIIT Bangalore, India <i>Aug 2022 – Dec 2022</i>
---	---

Programming in C++ and Java (ESS 201) <i>Teaching Assistant</i>	IIIT Bangalore, India <i>Aug 2021 – Dec 2021</i>
---	---

Technical Skills

Languages: Python, Java, C, C++, Verilog, Dart, JavaScript, TypeScript, R, SQL, HTML, CSS, MATLAB

Technologies/Frameworks: PyTorch, TensorFlow, Linux, Git, GitHub, Arduino, FPGA, Arduino Nano Sense, Raspberry Pi, Svelte, Nengo, ReactJS, Redux, Android, NextJS, NodeJS, ExpressJS, D3.js, Jenkins, MongoDB, Flutter, Blockly

Relevant Coursework

Artificial Intelligence and Machine Learning: Machine Learning, Math for Machine Learning (MML), MML for Signal Recognition, Computer Vision, Natural Language Processing, Reinforcement Learning, Few-shot Learning, Causal Inference and Learning, Neuromorphic Computing

Computer Science: Topics in Computability and Learning, Data Visualization, Software Testing, Operating Systems, Programming Languages, Software Engineering, Introduction to Automata Theory and Computability, Database Systems, Object Oriented Programming, Design and Analysis of Algorithms, Data Structures and Algorithms, Programming

Electronics and Communication: Signals and Systems, Computer Architecture, Digital Design, Embedded Systems, Control Systems, Computer Networks

Mathematics and Basic Sciences: Graph Theory, Probability, Statistics, Linear Algebra, Computational Genomics, Calculus, Real Analysis, Physics, Chemistry

Social Sciences: Software Product Management, Social Complexity and Systems Thinking, History of Ideas, Technical Communication, Economics, English

Achievements

- IAPR Best Student Paper Award at CVIP 2024 [↗](#)
- Graduate Aptitude Test in Engineering (GATE) 2023; All India Rank: 271 GATE Score: 742 [↗](#)
- Deans Merit List and Scholarship Winner in Academic Years 2019-24, IIITB
- Second Prize by Cartesi in ETHIndia 2022 Hackathon
- Second Prize in Team Hackathon Inclusive Stem Confluence 2020

Certifications

- Machine Learning Verzeo (2020) [↗](#)
- NVIDIA DLI: Fundamentals of Deep Learning (2021) [↗](#)
- NVIDIA DLI: Building Transformer-Based Natural Language Processing Applications (2021) [↗](#)
- Crash Course on Python (by Google, Coursera) [↗](#)
- Mathematics for Machine Learning: Linear Algebra (by Imperial College London, Coursera) [↗](#)
- Problem Solving (Basic) by HackerRank [↗](#)
- Java (Basic) by HackerRank [↗](#)