+91 9611724762
f20221155@goa.bits-pilani.ac.in
a3v1k.com
a3v1k
floofcat
poq7xJsAAAAJ

Advik Raj Basani

Education

2022–2026 **B.E. Computer Science (Honors)**, *BITS Pilani*, Goa, **(Transcript)**
- Achieved a **CGPA of 9.19/10.0**; **Rank 1** in the Department, Fall 2024.
- Recipient of BITS Merit Scholarship, awarded to the top 4% of 1K students, for 2 consecutive semesters in recognition of high academic standing.

Publications & Preprints

[Proceedings] *GASP: Efficient Black-Box Generation of Adversarial Suffixes for Jailbreaking LLMs*; **Advik Raj Basani**, Xiao Zhang. **Accepted to NeurIPS '25 and Oral Presentation at Building Trust in LLMs @ ICLR '25.** [Slides] [GitHub]

[arXiv] *Diversity Boosts AI-Generated Text Detection*; **Advik Raj Basani**, Pin-Yu Chen. **Accepted to Data in Generative Models @ ICML '25 and Oral Presentation @ CLEF '25.** [Slides] [GitHub] [HuggingFace]

[Proceedings] *G-GQSA: Exploiting Feature-Based Vulnerabilities and Enhancing Adversarial Resilience in Android Malware Detection*; **Advik Raj Basani**, Hemant Rathore. **Accepted as an Oral Presentation to 22^{nd} CCNC '25.**

[Proceedings] *When Less is More: Achieving Faster Convergence in Distributed Edge Machine Learning*; **Advik Raj Basani**, Siddharth Chaitra Vivek, Advaith Krishna, Arnab K. Paul. **Accepted to 31^{st} HiPC '24, Best Paper Nominee (top-2.5%).** [GitHub]

Research Experience

Aug. 2024 – Present **IBM Research AI**, *Remote, US*, Research Intern
Advisor: Dr. Pin-Yu Chen (Prinicipal Research Scientist)
- Developed **DivEye**, a framework capturing surprisal-based diversity to identify statistical fingerprints of AI-generated text; accepted to DIG-BUGS Workshop at **ICML 2025**.
  - Attained up to 0.99 **AUROC** on various benchmarks and secured 3^{rd} on the **RAID** leaderboard (2024); invited for an oral presentation at **CLEF 2025**.

Jan. 2024 – Present **CISPA Helmholtz Center for Information Security**, *Germany*, Research Assistant
Advisor: Dr. Xiao Zhang
- Investigating LLM reverse engineering, extending the work of Carlini et al. (arxiv.org:2403.06634) to infer architecture & weights from black-box models.
- Designed **GASP (Generative Adversarial Suffix Prompter)**, an efficient black-box framework for generating coherent adversarial suffixes that expose vulnerabilities in LLM safety mechanisms.
  - Open-sourced code & dataset AdvSuffixes, achieving SoTA performance on proprietary models; accepted at **NeurIPS 2025**.

Sept. 2025 – Present **PALM Lab**, *University of South Florida, Remote*, Research Intern
Advisor: Dr. Anshuman Chhabra
- Reverse engineering factual edits injected into LLMs through knowledge editing methods such as ROME (arxiv.org:2202.05262), by formulating the inversion as a **zeroth-order optimization** problem.

Nov. 2023 – Present **Data, Systems & HPC Lab**, Research Assistant
Advisor: Dr. Arnab K. Paul
- Accelerating GPU I/O operations by evaluating data transfer trade-offs between CPU and **GPUDirect**, aimed at improving throughput and reducing latency in training workloads.
- Developed an open-source framework for **Distributed Machine Learning** on resource-constrained clusters by prioritizing critical gradient updates to accelerate convergence and reduce communication overhead.
  - Achieved a $13.22\times$ reduction in training time and 62.1% lower communication overhead; work accepted as a **Best Paper Nominee** at the 31^{st} **HiPC**.

Work Experience

Oct. 2025 – Present **Trexquant**, *Remote, US*, Global Alpha Researcher – Intern
Supervisors: Saurabh Agarwal, Dr. Yunbo Zhang, Dr. Xin Wang
- Designed and implemented $30+$ research-grade alpha factors on **Pysim** for U.S. equity markets, focused on quarterly earnings prediction, statistically validated for robustness.

May. 2025 – Aug. 2025 **Oracle**, *India*, Member of Technical Staff – Intern
Supervisors: Anurag Sinha, Harish Dalmia (GenAI Team)
- Designed an end-to-end feature importance and factor analysis framework for Oracle's internal **AutoML pipeline**, offering configurable modes (fast vs. comprehensive) with model-agnostic support and scalable interpretability analysis.
- Integrated an extensive evaluation system, improving existing frameworks by $11\times$, for sensitive consumer use-cases such as employee attrition and financial market analysis.

May. 2024 – Aug. 2024 **Centre for Development of Advanced Computing**, *Kolkata, India*, Research Intern
Supervisor: Bibekananda Kundu (Research Scientist)
- Conducted R&D on fine-tuning LLMs to be human-centric via reinforcement learning, integrating emotional intelligence, empathy, and task-oriented conversational abilities.

Selected Projects

[Repo] 2025 ***[Re]*-Teaching Differentially Private Prompt Tuning for LLMs**
- Reimplemented & verified all experiments from **Flocks of Stochastic Parrots: Differentially Private Prompt Learning for LLMs** (arXiv:2305.15594), a study on applying differential privacy to prompt tuning for LLMs, and proposed new techniques that improved benchmark performance by ∼2%. [Slides]

[PR] 2024 **FaustNet: Enabling ML in Faust** **GSoC Contribution**
Mentors: Thomas Rushton, Dr. Stéphane Letz, Dr. Yann Orlarey (INRIA & GRAME)
- Developed an automatic differentiation library for the functional, audio domain-specific language **Faust** during Google Summer of Code, enabling audio engineers to integrate neural networks & other ML techniques.

Teaching Assistantships

BCSZC315 **Multicore & GPGPU Programming**, *Summer, 2025* — *Tasks: Creation of Teaching Material & Grading*
Lead TA; Instructor: Dr. Gargi Alavani, Dr. Kunal Korgaonkar — ∼40 students

CS F242 **Microprocessors & Interfacing**, *Spring, 2024* — *Tasks: Tutorials, Lab Creation & Autograder*
Lead TA; Instructor: Dr. Gargi Alavani — ∼300 students

CS F422 **Parallel Computing**, *Fall, 2024* — *Tasks: Tutorials, Assignments & Docker-based CUDA simulator*
Course Mentor (CM); Instructor: Dr. Gargi Alavani — ∼50 students

CS F111 **Computer Programming**, *Spring, 2023* — *Tasks: Plagiarism Detection & Autograder*
Course Mentor (CM); Instructor: Dr. Arnab K. Paul — ∼1000 students

Relevant Coursework

BITS Computer Networks*, Data Structures and Algorithms, Operating Systems*, Time Series Analysis and Forecasting*, Deep Learning*, Machine Learning, Compiler Construction*, Computer Architecture
* - Top 5 Student

Coursera Specialization in Google Data Analytics [Certificate], Advanced Learning Algorithms [Certificate], Supervised Machine Learning: Regression and Classification [Certificate]

Accomplishments

Grants Recipient of the IEEE TCPP Grant and DDF Grant, NTSE & KVPY Scholarships.

2025 Reviewer for TMLR, ICLR & ACL.

2025 Scored 326 (170Q, 156V) / 340 in GRE & 112 (27R, 28L, 30S, 27W) / 120 in TOEFL.

2024-25 $2\times$ winner of Hackenza, a hackathon organized by ASCII, BITS Goa.

2023 Coordinator & Lead, Google Developer Student Club, BITS Goa.
- Coordinated hackathons, seminars, and events across four verticals; oversaw the **AI/ML vertical** and club operations.

2023 Core Member for BITSKrieg, ranked 1^{st} for performance in CTFs across India.

2022 Developer for **Twitch Rivals: Medieval Mayhem**, **BisectHosting's GameMaster** & several other Minecraft tournaments and events. [Playlist, 2.5M+ views]

Technical Proficiency

Languages Proficient [`Python, Java, C++, LaTeX`], Intermediate [`Faust, Julia, Rust`] & more

Libraries / Softwares PyTorch, Transformers, JAX, Flax, HuggingFace, GitHub, GitLab, Anaconda, SpringBoot, Slurm, Docker, vLLM, Kafka, ZeroMQ, SciKit-Learn, GraphQL, Gemini & OpenAI APIs