JANVI THAKKAR

+447308908696 • janvi.thakkar@alumni.iitgn.ac.in • linkedin.com/in/janvi-thakkar • jvt3112.github.io • London, UK

Summary

Software Engineer and Researcher with hands-on experience in privacy, security, reinforcement learning, robotics, machine learning, and scalable infrastructure. Proven track record in building high-performance software systems, ranging from optimizing simulation algorithms to training end-to-end machine learning policies. Passionate about developing intelligent, reliable systems at the intersection of research and engineering, with a commitment to fast iteration, high ownership, and measurable impact.

Education

Imperial College London
London, UK

MSc Computing (AI & ML), Grade: Distinction

09/2023

- Thesis Title: Towards Unified Defense: Adversarial Training, Watermarking, and Privacy

Indian Institute of Technology (IIT), Gandhinagar
Gandhinagar, India

BTech, Computer Science and Engineering, Grade: 8.99/10

07/2022

- Gold Medal for Outstanding Performance in Arts and Culture (Batch of 2022)
- Ranked in top 0.1% (1.2 million applicants to IIT) JEE Mains and Advanced, 2018

Publications

Differentially Private and Adversarially Robust Machine Learning: An Empirical Evaluation
Janvi Thakkar, Giulio Zizzo, Sergio Maffeis
5th AAAI Workshop on Privacy-Preserving AI (PPAI-AAAI'24) [pdf]

Elevating Defenses: Bridging Adversarial Training and Watermarking for Model Resilience
Janvi Thakkar, Giulio Zizzo, Sergio Maffeis
2nd workshop on Deployable AI in Conjunction with AAAI (DAI-AAAI'24) [pdf]

FedSpectral+: Spectral Clustering using Federated Learning
Janvi Thakkar, Devvrat Joshi
3rd AAAI Workshop on Graphs and more Complex structures for Learning and Reasoning (GCLR-AAAI'23) [pdf]

k-Means SubClustering: A Differentially Private Algorithm with Improved Clustering Quality
Devvrat Joshi, Janvi Thakkar
Workshop on Privacy Algorithms in Systems (PAS), International Conference on Information and Knowledge Management (CIKM 2022) [pdf]

Skills

• Python(Advanced) • JAVA (Intermediate) • SQL • C++ • Bash • HTML • CSS • Git • GitLab CI • PyTorch • TensorFlow • • Google BigQuery • GCP Services • Docker • ROS • RViz • Data Structures • Algorithms • Reinforcement Learning • Robot Learning • Privacy and Security in Machine Learning • CI/CD • Terraform • High Work Ownership • Team Player

Experience

Ocado Technology
London, UK

Software Engineer

09/2023 - Present

- Developed and trained end-to-end reinforcement learning policies for robotic arm pick-and-place tasks. Led experimentation, reward design, and policy optimization, improving overall grasp success rate.
- Implemented real-time action chunking for flow based Vision-Language-Action (VLA) models, enhancing policy efficiency and temporal abstraction in manipulation tasks.

- Enhanced order profile manipulation algorithms used by analysts to simulate large-scale warehouse scenarios, resulting in **4x** improvement across **system performance** via profiling and logic optimization.
- Automated GCP infrastructure provisioning using Terraform and created a self-service portal for cross-team resource requests, enabling **>10 teams** to manage assets independently.

Decimal Point Analytics
Mumbai, India

Data Scientist Intern
05/2021 - 07/2021

- Fine-tuned and evaluated BERTSum and BART models on proprietary news data, achieving ROUGE-1 scores of **54.53** and **49.83**, outperforming existing baselines.
- Conceptualised and deployed a summarization API using the Flask, enabling real-time content processing for **150+ internal users** in the editorial workflow.

Lancaster University
Lancaster, UK

Summer Research Associate
05/2020 - 07/2020

- Developed an online learning and planning framework for multi-agent decision-making under **real-time constraints**, scaling to **50+ agents** in simulation.
- Implemented a Monte Carlo Tree Search (MCTS), a simulation-based approach for the continuous action space with discretization for making sequential decisions.
- Evaluated the technique in an infiltration game scenario, where an agent's goal is to reach the guarded targets with **no pretraining,** demonstrating adaptability against unknown swarm behaviours.

Projects

Unified Defense: Adversarial Training, Watermarking & Privacy
03/2023 - 09/2023

MSc Thesis

- Proposed a unified defense strategy combining adversarial training and adversarial watermarking to enhance model robustness and privacy.
- Achieved state-of-the-art adversarial accuracy of **92.01%** on MNIST and **65.84%** on Fashion-MNIST under PGD attacks, outperforming baseline methods.
- Demonstrated resilience against model stealing attacks, maintaining watermark verification accuracy of **68%** (MNIST) and **60%** (Fashion-MNIST) in grey-box scenarios.

FedSpectral+: Spectral Clustering using Federated Learning
01/2022 - 04/2022

Publication

- Uses the power iteration method to learn the global spectral embedding by incorporating the entire graph data without access to the raw information distributed among the clients.
- Designed a similarity metric to check the clustering quality of the distributed approach to that of the original/non-FL clustering.
- FedSpectral+ obtained a similarity of **98.85%** and **99.8%,** comparable to that of global clustering on the ego-Facebook and email-Eu-core dataset.

Geometrical Homogeneous Clustering for Image Data
01/2021 - 04/2021

Publication

- Presented a novel approach to reduce an image dataset using a Geometrical Homogeneous Clustering.
- Acquired an accuracy of **99.35 %** and **81.10%**, and a training data reduction of **87.27%** and **32.34%**, on MNIST, and CIFAR10 respectively.

Snappy - Command line tool for Snapshot Management
09/2020 - 12/2020

Repository

- Created a command-line tool for the Linux file system that supports snapshot management.
- Used a traditional split-mirror-based approach with some improvisation over the data storage while creating the copy.