Samanvith Kashyap

Bangalore | samanvithkashyap@gmail.com | linkedin.com/in/samanvith-kashyap | github.com/samanvithkashyap

Education

PES University, B.Tech in Computer Science 2024 – 2028
GPA: 8.14/10

Coursework: Data Structures & Algorithms, Linear Algebra, Probability & Statistics

Experience

Research Intern Dec 2025 – Present
CIE, PES University
Supervised by Prof. Lakshmeesha

- Built a CPU-deployable biometric attendance pipeline combining dlib ResNet-34 embeddings with a k-NN consensus classifier to reduce false positives under variable lighting conditions.
- Implemented blink-based liveness detection using 6-point EAR landmark tracking to counter photo and replay presentation attacks without requiring a GPU.
- Developed a t-SNE embedding visualization tool for pre-training dataset auditing, enabling identification and removal of noisy or mislabeled identity samples.
- Optimized end-to-end inference using HOG face detection and adaptive frame-skipping to sustain real-time throughput on CPU-only hardware.

Projects

ArtExtract GitHub Repo

- Built a multi-task CRNN (EfficientNet-B3 + BiGRU) with three joint classification heads (Style, Artist, Genre) trained on 81,000+ WikiArt images using a freeze-unfreeze strategy; tracked with Weights & Biases.
- Achieved Weighted F1 of 0.787 / 0.833 / 0.771 across Style, Artist, Genre heads; 96.98% top-5 artist accuracy and avg F1 of 0.797.
- Architecture generalizes as a fine-grained texture discriminator; being explored for medical domain transfer (surgical nerve identification, histopathology grading).
- **Tools:** Python, PyTorch, EfficientNet-B3, Weights & Biases, Pandas, PIL

Painting Similarity Retrieval GitHub Repo

- Built a zero-shot visual retrieval pipeline using DINOv2 (ViT-B/14) as a frozen feature extractor to generate 768-d CLS embeddings for 4,000+ paintings from the National Gallery of Art open dataset.
- Indexed embeddings using FAISS cosine similarity search; evaluated retrieval quality using Recall@K with artist attribution as weak supervision (R@1: 0.315, R@5: 0.400, R@10: 0.440).
- Visualized DINOv2 feature saliency maps and UMAP embedding clusters to validate semantic structure in learned representations.
- **Tools:** Python, PyTorch, DINOv2, FAISS, NumPy, UMAP

Technical Skills

Languages: Python, C++, C, MATLAB, JavaScript

ML & DL: PyTorch, Scikit-learn, NumPy, Pandas, OpenCV, FAISS, DINOv2

Tools: Git, Weights & Biases, VS Code, Jupyter, LaTeX, Qiskit

Achievements

1st place — IBM Qiskit Fall Fest Hackathon 2025, PES University (1 / 45+ teams)