KASHINATH REMESHKUMAR

Alappuzha, Kerala, India | +91 7034219692 | kashinathremeshkumar2004@gmail.com | github.com/kashinathRemeshkumar | kashinath.qzz.io

SUMMARY

AI/ML engineer (B.Tech CSE – AI & ML, 2026) with a published research paper on zero-shot diffusion-based image enhancement, and a RAG-based tutoring system built with sentence-transformers, ChromaDB and a Local/Cloud LLM. Comfortable with the Python ML stack (PyTorch, Pandas, model evaluation) and with exposing features through REST APIs, with hands-on experience on the backend of a newly launched social media app, including fixing bugs, building features and writing API documentation. Looking to build and deploy machine learning and Generative AI solutions.

TECHNICAL SKILLS

ML / Generative AI: PyTorch, Stable Diffusion, CLIP, LLMs, RAG, Embeddings (sentence-transformers), ChromaDB, NLP, Model Evaluation (PSNR, SSIM, LPIPS, MUSIQ), Pandas
Languages: Python, C/C++, Java, JavaScript, Kotlin
Backend & APIs: Django REST Framework, REST API design, Swagger/OpenAPI (drf-spectacular), PostgreSQL, Git/GitHub, Render
Tools & Platforms: Linux, Google Colab, Cloudflare Tunnel, Gradio
Other: Android (Jetpack Compose, Media3), ESP8266/ESP32 microcontroller programming

AI / ML PROJECTS

Zero-Shot Low-Light Image Enhancement (Published Research) — *PyTorch, Stable Diffusion v1.5, CLIP, Gradio*

- Designed a zero-shot enhancement pipeline that combines CLIP ViT-B/32 semantic guidance with Stable Diffusion v1.5 to brighten low-light images with no paired training data or fine-tuning.
- Added a semantic-consistency filter (CLIP cosine similarity > 0.85) to reject outputs that drift from the original scene, and cut diffusion sampling from 1000 steps to 40–50 with DDIM and FP16 mixed precision.
- Evaluated on the LOL and ExDark datasets using PSNR, SSIM, LPIPS and MUSIQ; reached 15.56 dB PSNR and 0.792 SSIM versus 15.40 dB and 0.515 for a retinex-based baseline on the evaluated samples.
- Built a Gradio demo hosted on Colab and co-authored the paper (see Publication).

Exam-Compass — RAG-Based AI Study Tutor — *Python, ChromaDB, sentence-transformers, LLM*

- Built a retrieval-augmented generation system for NEET/JEE preparation, so the tutor answers from the study material instead of relying only on the LLM's own knowledge.
- Embedded the study material with sentence-transformers and stored it in ChromaDB; at question time the most relevant passages are retrieved and passed to the LLM to ground its answer.
- Served the LLM backend from Google Colab through a Cloudflare Tunnel, giving on-demand inference without dedicated server cost.

IoT Crop Recommendation System — *ESP8266/ESP32, Deep Neural Network*

- Built an ESP32-based soil monitoring device that reads NPK values, moisture, temperature and soil resistance from an RS485 NPK sensor and a resistance probe, and shows the recommended crop and soil readings on an OLED display
- Trained a deep neural network on the public Kaggle crop recommendation dataset (2,000 records augmented to about 22,000), reaching about 99.3% accuracy.
- Ran inference fully on-device, giving near real-time crop recommendations from live sensor data without relying on a server or network connection.

OTHER PROJECTS

Zone Music — Local-First Android Music Player — *Kotlin, Jetpack Compose, Media3 (ExoPlayer)*

- Built an Android music player in Kotlin where you pick your music folders with the system folder picker, so it needs no broad storage permission. Song tags are read once and saved to a JSON cache, so the library shows up instantly and refreshes in the background.
- Added background playback using Media3's *MediaLibraryService*, with notification controls. The app remembers your queue, shuffle and repeat settings when you reopen it.

- Added Android Auto support so songs can be browsed and played from the car display, and tested it in a car. Source: github.com/kashinathRemeshkumar/Zone-Music.

Stego-Suite — Multi-Format Steganography Toolkit — *C++*, *Python*, *NumPy*, *SciPy*, *FFmpeg*

- Built a toolkit that hides files inside images and audio using least-significant-bit (LSB) manipulation. The C++ image module hides any file in a PNG and stores a small header (magic bytes, file size, filename) so the decoder can confirm hidden data is present before extracting it.
- Wrote the audio module in Python, which hides one audio clip in the Side channel of a stereo cover using Mid/Side processing, with an adjustable bits-per-sample setting, and saves the result as a lossless 16-bit WAV.
- Documented the limits, such as JPEG or MP3 re-compression destroying the hidden data. Source: github.com/kashinathRemeshkumar/stego-suite.

PDF Watermarking Tool — Batch Watermark & Merge (Client Project) — *Python*, *GUI*

- Built a Python desktop app for a client that watermarks and merges PDFs in bulk with one click, replacing a manual process.
- Supports clickable text watermarks with hyperlinks, background image watermarks with scale and transparency control, and optional front or appended pages. Source: github.com/kashinathRemeshkumar/PDF_Watermarking_tool.

PUBLICATION

Kashinath Remeshkumar, Abhijith R, Dan Philip Bobby, Kevin Varghese Theveril, Hema H. "Zero Shot Low Light Image Enhancement using Vision Language Models and Semantic Diffusion."
International Journal on Emerging Research Areas (IJERA), vol. 6, no. 1, pp. 77–83, May 2026. ISSN 2230-9993.
ijera.in/index.php/IJERA/article/view/380

EXPERIENCE

Django Developer (Intern) — *Assistrend.ai* May 2026 – Aug 2026

- Built and fixed backend features for a social media app focused on youth mental health, using Django REST Framework and PostgreSQL, deployed on Render.
- Diagnosed and fixed bugs across accounts, social, personalization, messaging and notifications, including input validation, case-insensitive email and username lookups, and message delivery/read status.
- Documented API endpoints in Swagger with drf-spectacular and worked through Git/GitHub branches, pull requests and code review.

Full Stack Web Developer (Intern) — *Diviso Softtech*, *Palakkad* Jun 2025 – Jul 2025

- Built responsive features in HTML, CSS and JavaScript for a game-based learning framework, working with the team on architecture, UI/UX and data handling.

Python Application Developer *(Freelance)* — *KTU Special*, *Alappuzha* Jan 2022 – Jan 2024

- Built and deployed a Python GUI application that automated the website's document branding workflow, replacing manual batch watermarking and merging with a one-click process.
- Developed the PDF watermarking engine: clickable hyperlinked text watermarks, scale- and transparency-controlled image watermarks, and optional front or appended pages.

EDUCATION

B.Tech, Computer Science (AI & ML) — *Sree Buddha College of Engineering*, *Kerala* 2026
CGPA: 7.86 / 10 | Class XII (CBSE): 86.33% | Class X (CBSE): 85.6%

CERTIFICATIONS

Introduction to Programming in C – NPTEL, Elite rank (Jan–Mar 2024) | Python (Basic) – HackerRank | Python & AI Bootcamp (Google, Microsoft) | Hands-on IoT Workshop (Open Leaves)