

AI/ML Project Portfolio - Extracted Work

Table of Contents

1. [AI/LLM & RAG Systems](#)
 2. [Computer Vision Projects](#)
 3. [Generative AI Projects](#)
 4. [MLOps & Infrastructure](#)
 5. [NLP & Language Processing](#)
 6. [Finance & Trading Applications](#)
 7. [Healthcare & Medical Imaging](#)
 8. [Edge AI & Optimization](#)
-

AI/LLM & RAG Systems

Project 1: Cloud-Agnostic AI Platform

Description: Architected a cloud-agnostic platform deploying 1,000+ models with 99.9% uptime using domain-driven design, KServe, vLLM, and Dynamo KV caching for transformer and mixed architecture models with Kubernetes orchestration.

Technologies: Kubernetes, KServe, vLLM, Dynamo KV caching, Domain-driven design

Impact:

- Deployed 1,000+ models
 - Achieved 99.9% uptime
 - Cloud-agnostic architecture
-

Project 2: Investment Advisor Copilot

Description: Built an investment advisor copilot using LangGraph orchestration and MCP (Model Context Protocol) for agent tool calling with financial assistant capabilities, real-time market analysis, and sentiment-driven portfolio insights.

Technologies: LangGraph, MCP, AI agents, Real-time market analysis, Sentiment analysis

Impact: Real-time financial insights and portfolio optimization

Project 3: Legal M&A Agentic RAG System

Description: Engineered a legal M&A agentic RAG system processing complex merger documents with 95% precision using fine-tuned Llama 3, LoRA, and Graph RAG with vLLM for GPU-optimized document analysis automation.

Technologies: Llama 3, LoRA, Graph RAG, vLLM, GPU optimization

Impact:

- 95% precision in document analysis
 - Automated complex merger document processing
-

Project 4: Compliance Audit Automation

Description: Created compliance audit automation processing 500,000+ SEC documents with 98% accuracy using agentic workflows and multi-level RAG with MongoDB optimization for Fortune 500 financial regulations and risk assessment.

Technologies: Agentic workflows, Multi-level RAG, MongoDB, Document processing

Impact:

- Processed 500,000+ documents
 - 98% accuracy in compliance detection
-

Project 5: AI Supply Chain Optimizer

Description: Built an AI supply chain optimizer using AutoGen agentic frameworks, MCP protocol, and RAG for inventory automation and disruption prediction, reducing operational costs via distributed data pipelines.

Technologies: AutoGen, MCP, RAG, GPT-4, Distributed data pipelines

Impact:

- 20% reduction in operational costs
- 15% improvement in delivery efficiency

Project 6: Investment Research Copilot

Description: Designed an investment research copilot using agentic RAG with MongoDB Atlas, processing 100,000+ financial documents and achieving a 0.95 BLEU score for analysts with real-time search and ranking capabilities.

Technologies: Agentic RAG, MongoDB Atlas, Real-time search, NLP

Impact:

- Processed 100,000+ financial documents
- 0.95 BLEU score
- Real-time search capabilities

Project 7: Portfolio Risk Assessment System

Description: Created a portfolio risk assessment system using fine-tuned LLaMA-3 with LoRA for real-time market volatility analysis, implementing MCP for seamless integration with risk management tools.

Technologies: LLaMA-3, LoRA, MCP, Real-time analysis, Risk management

Impact: Real-time market volatility analysis and risk assessment

Project 8: ESG Compliance Analyzer

Description: Built an ESG compliance analyzer using vLLM-optimized models, processing sustainability reports with 94% accuracy and automating ESG scoring and regulatory compliance via Kubernetes deployment.

Technologies: vLLM, Kubernetes, NLP, Compliance automation

Impact:

- 94% accuracy in ESG analysis
 - Automated sustainability reporting
-

Project 9: Quantitative Research Platform

Description: Engineered a quantitative research platform using the MCP protocol for model orchestration, implementing fine-tuned LLMs for factor analysis and alpha generation with a GPU-accelerated backtesting infrastructure.

Technologies: MCP, Fine-tuned LLMs, GPU acceleration, Backtesting

Impact: Enhanced quantitative research capabilities with automated factor analysis

Project 10: Multi-Tenant AI Backend

Description: Architected a multi-tenant AI backend supporting 10,000+ concurrent users with real-time experiments. Established MLOps pipelines, reducing model deployment time from two hours to one hour using Kubernetes orchestration.

Technologies: Kubernetes, MLOps, Multi-tenant architecture, CI/CD

Impact:

- Supports 10,000+ concurrent users
 - 50% reduction in deployment time
-

Project 11: Sales Intelligence RAG System

Description: Built an RAG system using LangGraph. Optimized LLaMA-3 8B models reduced size by 75% and latency by 50%, maintaining 0.95 recall and 30 BLEU score.

Technologies: LangGraph, LLaMA-3, Model optimization, RAG

Impact:

- 75% size reduction
 - 50% latency reduction
 - 0.95 recall, 30 BLEU score
-

Project 12: Personalized Sales & Marketing Bot

Description: Developed a personalized sales and marketing bot using LangChain with few-shot and chain-of-thought prompting. Improved lead qualification and content personalization.

Technologies: LangChain, Few-shot learning, Chain-of-thought prompting

Impact:

- 40% improvement in lead qualification
 - 25% improvement in content personalization
-

Project 13: Multi-Agent Contract Management Tool

Description: Developed an end-to-end multi-agent contract management tool for extracting, drafting, and managing legal contracts. Utilized a fine-tuned Llama 3 model, Graph RAG with Neo4j via LangChain, and quantized with QLora. Deployed using vLLM.

Technologies: Llama 3, Graph RAG, Neo4j, LangChain, QLora, vLLM

Impact: End-to-end automated contract management

Project 14: Legal Copilot with Multi-Level RAG

Description: Developed a legal copilot using a multi-level RAG approach with MongoDB Atlas and OpenAI Assistant API. The system queries over 100,000 documents for efficient and accurate contract insights, achieving a BLEU score of 0.95.

Technologies: Multi-level RAG, MongoDB Atlas, OpenAI Assistant API

Impact:

- 100,000+ documents processed
- 0.95 BLEU score

Project 15: Automated Contract Drafting System

Description: Designed a system to automate legal contract drafting and review by fine-tuning Llama 3 with LoRA. It generates templates, suggests clauses, and detects inconsistencies.

Technologies: Llama 3, LoRA, NLP, Legal automation

Impact:

- 50% reduction in drafting time
- 35% reduction in errors

Project 16: Insurance Copilot

Description: Created an insurance copilot using a fine-tuned OpenAI GPT-4 API. Enhanced customer interactions by providing accurate, real-time responses and automating routine tasks.

Technologies: OpenAI GPT-4, Fine-tuning, Customer interaction automation

Impact: Enhanced customer experience with real-time AI assistance

Project 17: Financial Advisory AI Assistant

Description: Created an AI assistant using GPT-4. Developed prompt templates with constrained generation for tailored investment advice and market insights.

Technologies: GPT-4, Prompt engineering, Constrained generation

Impact: Personalized financial advisory services

Project 18: Dynamic Lead Scoring System

Description: Engineered a dynamic lead scoring system utilizing RHFL-tuned, fine-tuned Llama 2 with few-shot learning to analyze prospect interactions and prioritize high-value leads.

Technologies: Llama 2, RLHF, Few-shot learning, Lead scoring

Impact: Improved lead prioritization and conversion rates

Project 19: Travel Assistant with RAG

Description: Developed a travel assistant using retrieval-augmented generation (RAG) and agentic workflows. Implemented chain-of-thought prompting and context-aware templates for trip planning, flight bookings, and Q&A.

Technologies: RAG, Agentic workflows, Chain-of-thought prompting, LangChain

Impact: Automated travel planning and booking assistance

Project 20: Virtual House Agent

Description: Developed a virtual house agent via CrewAI and fine-tuned Llama 3 via LoRA to reduce home maintenance costs by 30% and improve service

efficiency. AI processes homeowner requests, then generates quick estimates and schedules professionals.

Technologies: CrewAI, Llama 3, LoRA, Agentic workflows

Impact:

- 30% reduction in maintenance costs
- Improved service efficiency

Project 21: Personal Finance Assistant

Description: Implemented a personal finance assistant using LangChain and PEFT fine-tuned Mistral AI to increase savings. AI categorizes expenses, creates budgets, and optimizes investment portfolios based on risk profiles and market conditions.

Technologies: LangChain, PEFT, Mistral AI, Portfolio optimization

Impact: 15-20% increase in savings

Project 22: Clinical Trials Audit Bot

Description: Created an audit bot for auditing clinical trials, utilizing chain of thought prompting, data-driven prompting, and ElasticDB-based retrieval-augmented generation (RAG).

Technologies: Chain-of-thought prompting, ElasticDB, RAG, Clinical data analysis

Impact: Automated clinical trial auditing

Project 23: Legal Law Chatbot

Description: Developed a legal chatbot leveraging OpenAI's GPT-4, LangChain, and the retrieval-augmented generation (RAG) model, integrated with a Pinecone

database and Streamlit UI, all built on a scalable Azure architecture. Fine-tuned the model on Azure AI Studio and enhanced capability by connecting LLMs with Azure services like Azure AI Search.

Technologies: GPT-4, LangChain, RAG, Pinecone, Streamlit, Azure AI Studio, Azure AI Search

Impact: Precise and context-sensitive legal advice delivery

Project 24: Personality Assessment and Reporting System (PARS)

Description: Developed the automated Personality Assessment and Reporting System using LLMs and retrieval-augmented generation (RAG); Mistral, fine-tuned via PEFT, RAG with LangChain, and Pinecone for data extraction and cleaning. Used OpenAI GPT-4 with chain of thought and LangChain with Pinecone to compare traits and create visual reports.

Technologies: Mistral, PEFT, RAG, LangChain, Pinecone, GPT-4, Chain-of-thought

Impact: 95% accuracy in personality assessment

Project 25: AI-Powered Kubernetes Monitoring System

Description: Developed an intelligent system leveraging LangGraph for agent orchestration and integrated multiple MCP servers from GitHub, including specialized MCP servers for Grafana, Prometheus, and Kubernetes. Automatically remediates high resource spikes through strategic restarts and memory clearing operations. Used fine-tuned Llama 3.1 8B model for agent orchestration and decision-making.

Technologies: LangGraph, MCP, Grafana, Prometheus, Kubernetes, Llama 3.1, Cursor MCP, GitHub integration

Impact: Fully autonomous infrastructure management and code optimization pipeline

Computer Vision Projects

Project 26: Airport Security Surveillance System

Description: Implemented a model that detects suspicious activity at airports using VideoMAE. Prioritized high accuracy, low latency, and efficient deployment on Linux servers.

Technologies: VideoMAE, Computer Vision, Deep Learning, Linux

Impact: High accuracy security monitoring with low latency

Project 27: Kanji Writing Assistant

Description: Created a system to assist human Kanji writing through imitation learning and OpenCV using Kanji videos to generate Kanji images, predicting poses for robotic arms.

Technologies: OpenCV, Imitation learning, Pose prediction, Robotics

Impact: Automated Kanji character generation for robotic writing

Project 28: Table Tennis Shot Detection System

Description: Used YOLOv5 for shot detection, OpenCV for object detection, and VideoMAE for shot recognition in TT Games for sports organization. Developed for the World Table Tennis Organization using OpenCV for ball detection and depth estimation via OpenMiDaS, achieving smooth 60 fps. Implemented deep learning approach using VideoMAE for shot detection with 95% accuracy in predicting forehand and backhand shots.

Technologies: YOLOv5, OpenCV, VideoMAE, OpenMiDaS, Real-time processing

Impact:

- 95% accuracy in shot detection
- 60 fps real-time performance
- Used in television broadcasts

Project 29: Infrared Object Detection System

Description: Integrated an infrared object detection system using the YOLO architecture with DeepSORT, achieving high accuracy with an mAP of 0.88 for detecting and tracking vehicles.

Technologies: YOLO, DeepSORT, Infrared imaging, Object tracking

Impact:

- 0.88 mAP accuracy
- Real-time vehicle detection and tracking

Project 30: Satellite Image Segmentation System

Description: Created a satellite image segmentation system using U-Net and Mask R-CNN models, achieving a Dice score of 0.94, greatly enhancing agricultural field detection and analysis.

Technologies: U-Net, Mask R-CNN, Satellite imagery, Semantic segmentation

Impact:

- 0.94 Dice score
- Enhanced agricultural field detection

Project 31: Motion Transfer System

Description: Implemented a motion transfer system using a first-order model, achieving high-quality motion transfer between faces while preserving the identity and facial expressions of the target face.

Technologies: First-order motion model, Deep Learning, Face animation

Impact: High-quality facial motion transfer with identity preservation

Project 32: 3D Object Detection for Autonomous Vehicles

Description: Rolled out a 3D object detection and tracking system for autonomous vehicles using lidar data and the VoxelNet algorithm, enhancing the vehicles' perception and tracking capabilities in a 3D environment.

Technologies: VoxelNet, Lidar, 3D object detection, Autonomous vehicles

Impact: Enhanced 3D perception for autonomous navigation

Project 33: Real-time Yoga Pose Estimation System

Description: Built a real-time Yoga pose estimation system for Android and iOS using OpenPifPaf. Achieved 95% accuracy and offers detailed feedback to refine users' yoga practice. Optimized for real-time performance on mobile devices.

Technologies: OpenPifPaf, TensorFlow Lite, Mobile optimization, Pose estimation

Impact:

- 95% accuracy
 - Real-time mobile performance
 - Personalized yoga feedback
-

Project 34: Kitchen Item Detection System

Description: Built a computer vision system for accurately detecting items in the kitchen, with high accuracy and low latency. The system was optimized for real-time performance on mobile devices.

Technologies: Computer Vision, Deep Learning, Mobile optimization, Object detection

Impact: Real-time kitchen item detection on mobile devices

Project 35: Interactive NFT (2D-3D Conversion)

Description: Developed a system that converts 2D images of NFTs into immersive 3D models using a combination of selective 3D inpainting via Stable Diffusion and depth estimation techniques. The system ensures stable and consistent generation, delivering high-quality 3D representations.

Technologies: Stable Diffusion, 3D inpainting, Depth estimation, 3D reconstruction

Impact: High-quality 2D to 3D NFT conversion for immersive experiences

Project 36: Toonify Pets

Description: Developed a system to convert animal images into animated cartoons by training a GAN on unpaired images using the StyleGAN architecture. Enhanced with CLIP for better alignment with textual descriptions and a feature extractor for refined details.

Technologies: StyleGAN, GAN, CLIP, Image-to-cartoon conversion

Impact: FID score of 12, demonstrating high visual fidelity

Project 37: Brain Image Segmentation for Alzheimer's Detection

Description: Worked on brain image segmentation using deep learning, training neural networks to accurately identify and classify structures in brain images linked to Alzheimer's disease using segmentation and computer vision techniques.

Technologies: Deep Learning, Medical imaging, Semantic segmentation, CNNs

Impact: Accurate identification of Alzheimer's disease markers

Project 38: ASL-to-Speech Healthcare System

Description: Developed an advanced system that translates American Sign Language (ASL) gestures into natural-sounding speech using computer vision and deep learning. Employs MediaPipe for hand tracking, custom CNNs for spatial relationships, bidirectional LSTM for temporal sequences, Tacotron 2 for mel-spectrogram generation, and WaveNet as vocoder.

Technologies: MediaPipe, CNNs, Bidirectional LSTM, Tacotron 2, WaveNet, Hand tracking

Impact:

- 92% accuracy in ASL translation
 - Natural-sounding speech output
 - Enhanced communication in healthcare
-

Generative AI Projects

Project 39: News to Infographics (Text-to-Video)

Description: Successfully delivered a generative AI model utilizing Stable Diffusion and LLMs that animates stories sourced from news articles. Helped secure Y Combinator funding. News articles are summarized using GPT-3.5 Turbo and Davinci via LangChain, then videos are generated using fine-tuned Stable Diffusion 2.1.

Technologies: Stable Diffusion 2.1, GPT-3.5 Turbo, Davinci, LangChain,
Text-to-video

Impact:

- Secured Y Combinator funding
- Automated news story animation

Project 40: Text-to-Video Generation for Mathematical Equations

Description: Developed a robust diffusion model capable of interpreting English text descriptions of mathematical equations and generating accurate, coherent video representations. Implemented advanced optimization techniques to improve latency and memory footprint for real-time applications.

Technologies: Diffusion models, Text-to-video, Optimization, Real-time processing

Impact: Visual aid for educational settings to understand mathematical concepts

Project 41: Text-to-Video with Character Consistency

Description: Implemented a text-to-video system using MORA and fine-tuning the iv2GenXL model with LoRA, where the system converts text descriptions into short video reels while preserving character consistency across scenes.

Technologies: MORA, iv2GenXL, LoRA, Text-to-video, Character consistency

Impact: 0.9 CLIP score indicating strong text-video alignment

Project 42: Personalized Art Generation Bot

Description: Developed a bot to assist users in generating custom art based on their interactions and provided images. Employed GPT-3.5 as the LLM with soft prompt pipeline considering users' prior interactions. Stable Diffusion 2.1 was

fine-tuned using LoRA, incorporating themes and prompts recommended by the LLM. System handles user-specific data including NSFW content while maintaining privacy.

Technologies: GPT-3.5, Stable Diffusion 2.1, LoRA, Soft prompting, Custom art generation

Impact: Personalized art generation with privacy preservation

MLOps & Infrastructure

Project 43: Quantization SDK for Edge Devices

Description: Built a quantization software development kit (SDK) for 4-bit and 8-bit quantization, enabling efficient implementation and optimization of deep learning models on Edge (CPU-based) hardware, which enhanced performance and capabilities.

Technologies: Quantization, Edge AI, SDK development, Model optimization

Impact: Enhanced performance for edge devices

Project 44: Computer Vision Model Benchmarking

Description: Benchmarked different computer vision and generative models with custom quantization and optimization SDK for IOT and custom Edge devices.

Technologies: Benchmarking, Quantization, IOT, Edge devices

Impact: Optimized model performance for IOT devices

Project 45: Integer-Only Quantization System (Patent)

Description: Developed and patented a system for integer-only quantization-aware training. The system enhances speed and performance of deep learning networks

on low-precision devices. Developed the pseudo-cross entropy loss function and designed quantization scheme. SDK developed to enable utilization on low-power edge computing devices, successfully used on Jetson and custom hardware.

Technologies: Quantization, Edge computing, Patent, SDK development, Low-precision inference

Impact:

- Patented technology
- Enhanced performance on edge devices
- Successfully deployed on Jetson hardware

Project 46: On-Device Edge Learning System (Patent)

Description: Developed and patented a system for device edge learning. The method includes training an AI model for extracting visual embeddings with pre-trained visual deployment networks, performing inference, initiating edge learning, extracting visual embeddings, adding text image embeddings, using text embedders, converting text to image embeddings to generate augmented image embeddings.

Technologies: Edge learning, Visual embeddings, Multi-modal learning, Patent

Impact: Patented on-device learning technology

Project 47: MLOps Pipeline Implementation

Description: Implemented and managed MLOps practices to ensure efficient lifecycle management of ML models. Worked with GitHub Actions to enhance code integration and deployment process automation for CI/CD. Reduced model deployment time from two hours to one hour.

Technologies: MLOps, GitHub Actions, CI/CD, Kubernetes, Automation

Impact:

- 50% reduction in deployment time
- Automated ML lifecycle management

Project 48: Cloud Computing Infrastructure Optimization

Description: Utilized AWS for cloud computing, cutting model deployment time by 50% and improving resource efficiency by 40%. Enhanced model streaming and observability, boosting real-time performance tracking by 30% and reducing operational overhead by 25%.

Technologies: AWS, Cloud computing, Model deployment, Observability

Impact:

- 50% faster deployment
- 40% improved resource efficiency
- 30% better performance tracking
- 25% reduced operational overhead

NLP & Language Processing

Project 49: AI-Powered Language Exchange Platform

Description: Created an AI-powered language exchange platform using React, FastAPI, and NLP models for real-time translation and analysis. Increased employee language proficiency by 45% and improved cross-cultural communication efficiency by 30%.

Technologies: React, FastAPI, NLP, Real-time translation

Impact:

- 45% increase in language proficiency
- 30% improvement in communication efficiency

Project 50: AI-Powered Language Learning Platform

Description: Developed an AI-powered language learning platform using transformer models and SpaCy. Implemented adaptive difficulty and personalized feedback, boosting engagement by 55% and retention by 40%. Integrated 20+ language corpora for diverse content.

Technologies: Transformers, SpaCy, Adaptive learning, NLP

Impact:

- 55% boost in engagement
- 40% increase in retention
- 20+ languages supported

Project 51: Real-Time Social Media Sentiment Analyzer

Description: Developed a real-time social media sentiment analyzer using RoBERTa and Kafka. Achieved 92% accuracy in multilingual emotion detection, reducing customer response time by 33% for major brands.

Technologies: RoBERTa, Apache Kafka, Sentiment analysis, Real-time processing

Impact:

- 92% accuracy
- 33% reduction in response time
- Multilingual emotion detection

Project 52: Topic Modeling System

Description: Created a topic modeling model utilizing LDA and NMF algorithms for latent topic extraction from text corpora, and applied clustering algorithms to group similar topics, providing better understanding and organization of text documents.

Technologies: LDA, NMF, Topic modeling, Clustering, NLP

Impact: Improved text organization and understanding

Project 53: Fake News Classification System

Description: Developed a system to detect and classify fake news articles using machine learning and natural language processing techniques. Employed the SetFit model and LSTM, developed an ensemble of SetFit and LSTM. Additionally, k-means clustering was used to cluster types of fake news.

Technologies: SetFit, LSTM, K-means clustering, Ensemble learning, NLP

Impact: Reliable tool to combat misinformation

Project 54: Text and Voice Assistant for Dyslexia

Description: Developed a text and voice-based assistant for dyslexic students as an MVP. Fine-tuned the Mixtral-8x7B-Instruct model using RLHF and leveraged Azure DevOps for CI/CD and deployment. Implemented RAG to enhance data extraction and query responses. Integrated text-to-speech functionality to aid students with reading difficulties.

Technologies: Mixtral-8x7B, RLHF, RAG, Azure DevOps, Text-to-speech

Impact: Improved accessibility for dyslexic students

Project 55: Patent Search System

Description: Created a tool to search for similar patents using LangChain's OpenAI ADA model embeddings and Pinecone for improved indexing.

Technologies: LangChain, OpenAI ADA embeddings, Pinecone, Vector search

Impact: 90% precision rate for relevant patent identification

Project 56: Recipe Recommendation System

Description: Fine-tuned Llama 3 8B using instruction tuning, preference alignment (ORPO), and hyperparameter optimization for improved recipe recommendations.

Technologies: Llama 3, Instruction tuning, ORPO, Hyperparameter optimization

Impact: Improved recipe recommendations with preference alignment

Finance & Trading Applications

Project 57: Weather Forecasting System

Description: Built a weather forecasting system using Mistral LLM. The LLM enhanced real-time weather insights, improving forecast accuracy to 85% and reducing alert processing time by 35%, ensuring timely and precise weather updates.

Technologies: Mistral, LLM, Time series forecasting, Real-time processing

Impact:

- 85% forecast accuracy
 - 35% reduction in alert processing time
-

Project 58: Fashion Sales Forecasting System

Description: Created a fashion sales forecasting system for an eCommerce platform using CLIP embeddings and GPT-3. Analyzed product images and descriptions to predict trends, improving inventory management and boosting quarterly sales by 18%.

Technologies: CLIP, GPT-3, Sales forecasting, Computer Vision + NLP

Impact:

- 18% boost in quarterly sales

- Improved inventory management

Project 59: Keystroke Dynamics Authentication

Description: Implemented keystroke dynamics for eCommerce authentication with Python for feature extraction and TensorFlow with LSTM networks for deep learning. Achieved 95% accuracy, cut authentication time by 30%, and reduced security breaches by 45%.

Technologies: LSTM, TensorFlow, Biometric authentication, Behavioral biometrics

Impact:

- 95% accuracy
- 30% faster authentication
- 45% reduction in security breaches

Project 60: Autonomous Treasury Management System

Description: Designed and implemented an AI-driven treasury management platform that optimizes corporate cash positioning across multiple currencies and investment vehicles. Engineered hybrid architecture combining LLMs for financial document analysis with traditional statistical methods for cash flow forecasting. Includes natural language interface and interactive dashboard.

Technologies: LLMs, Statistical forecasting, NLP, Dashboard visualization, Multi-currency

Impact:

- 22% improvement in cash utilization efficiency
- Maintained all liquidity constraints

Project 61: Income Tax Fraud Detection

Description: Detected income tax fraud using an ensemble of supervised anomaly detection, unsupervised clustering, and rule-based backtracking.

Technologies: Anomaly detection, Clustering, Rule-based systems, Ensemble methods

Impact: Effective fraud detection system

Healthcare & Medical Imaging

Project 62: Medical Imaging Analysis (Various Projects)

Description: Multiple projects involving brain image segmentation, medical imaging analysis, and Alzheimer's disease detection using deep learning and computer vision techniques.

Technologies: Deep Learning, Medical imaging, Semantic segmentation, CNNs

Impact: Accurate medical diagnosis support

Edge AI & Optimization

Project 63: 3D Reconstruction System

Description: Implemented 3D reconstruction capabilities for various applications including autonomous vehicles and computer vision projects.

Technologies: 3D Reconstruction, Computer Vision, Lidar, Depth estimation

Impact: Enhanced 3D perception capabilities

Project 64: Real-time Processing Optimization

Description: Multiple projects involving real-time processing optimization for mobile devices, edge devices, and embedded systems with focus on latency reduction and efficient resource usage.

Technologies: TensorFlow Lite, Model quantization, Edge AI, Mobile optimization

Impact: Real-time performance on resource-constrained devices

Summary Statistics

Total Projects Extracted: 64 projects

Key Technology Areas:

- AI/LLM & RAG Systems: 25 projects
- Computer Vision: 13 projects
- Generative AI: 4 projects
- MLOps & Infrastructure: 6 projects
- NLP & Language Processing: 8 projects
- Finance & Trading: 5 projects
- Healthcare: 1 major area
- Edge AI & Optimization: 2 projects

Key Technologies Used:

- LLMs: Llama 3, GPT-4, Mistral, Claude
- Frameworks: LangChain, LangGraph, PyTorch, TensorFlow
- RAG: Pinecone, MongoDB Atlas, Neo4j, ChromaDB
- Computer Vision: YOLO, OpenCV, Mask R-CNN, U-Net
- Generative AI: Stable Diffusion, DALL-E, GANs
- MLOps: Kubernetes, Docker, GitHub Actions, MLflow
- Cloud: AWS, Azure, Google Cloud

Notable Achievements:

- 99.9% uptime on 1,000+ model deployments
- 95% precision in legal document analysis
- 98% accuracy in compliance audit automation
- 92% accuracy in ASL translation

- 85% weather forecast accuracy
- Secured Y Combinator funding with generative AI model
- Multiple patented technologies