

GRUHESH SRI SAI KARTHIC KURRA

Address: SMR Vinay Fountainhead, 4th Block 1012, Miyapur, Near Calvary Temple, Hyderabad - 500049, Telangana, India

Phone: +91 8008258065 | **Email:** gruheshkurra2@gmail.com

LinkedIn: [linkedin.com/in/gruhesh-sri-sai-karthik-kurra-178249227](https://www.linkedin.com/in/gruhesh-sri-sai-karthik-kurra-178249227) | **GitHub:** github.com/GruheshKurra | **Portfolio:**

<https://karthik.zynthetix.in/>

EDUCATION

Bachelor of Technology in Computer Science and Engineering | CGPA: 9.72/10

KL University, Hyderabad, India | August 2021 – May 2025

WORK EXPERIENCE

AI Integration Engineer | InfoAjax Consulting Pvt Ltd, Remote | October 2025 – November 2025

- Leading development of AI-driven integration framework connecting enterprise applications with cybersecurity services through standardized APIs and intelligent orchestration
- Architected agent-based system using Azure OpenAI for automated schema mapping and adaptive connection strategies with performance learning from historical patterns
- Built full-stack application using React 18, Vite, TailwindCSS frontend and FastAPI backend with WebSocket integration for real-time monitoring and agent coordination
- Integrated Azure Active Directory via Microsoft Graph API for application discovery and automated metadata management

WSO2 API Developer | InfoAjax Consulting Pvt Ltd, Remote | October 2024 – June 2025

- Developed and deployed RESTful APIs for PLDT (Philippines' largest telecommunications provider) using WSO2 Integration Studio for ticketing system integration
- Managed production assistance ticket systems with secure credential management through GitLab and Choreo cloud hosting platform with vault-based secrets management
- Performed systematic debugging in environments with rare WSO2 framework documentation; directly engaged with international clients under strict SLAs

Founder & CEO | Zynthetix (Personal Research Venture) | March 2024 – January 2025

- Explored a research-driven initiative addressing dataset scarcity through synthetic data generation platform producing tabular, image, and text data using GANs and VAEs
- Designed meta-controller architecture where a mother model orchestrates 200-1000 specialized child models to generate millions of rows without duplication
- Implemented prototype with transformer-based PII detection (98.5% accuracy) and achieved 95% statistical similarity between synthetic and real data
- Collaborated with 3 developers during prototyping; established Git workflows and CI/CD pipelines

INTERNSHIPS & TRAINING

Research Intern | IIIT Hyderabad | December 2024 – June 2025

- Received structured training in LLM optimization techniques including quantization, pruning, and knowledge distillation
- Led development of Layout Preserving Document Translation system with Vision Transformer-based layout detection (92% accuracy)
- Built FastAPI backend with RESTful endpoints integrating Tesseract OCR for multi-language text extraction while preserving formatting
- Presented project at IIIT Hyderabad research expo (May 2025) with live demonstrations to faculty and industry visitors

AI & ML Research Training | iHub-Data, IIIT Hyderabad | May 2024 – October 2024

- Completed intensive 6-month apprenticeship program focused on AI and ML research methodologies
- Engaged in hybrid training covering LLM fundamentals, prompt engineering, model fine-tuning, quantization, and deployment strategies

Intern | Xtraleap India Pvt Ltd, Remote | July 2023 – October 2023

- Worked on data science projects involving exploratory data analysis, statistical modeling, and ML model development
- Gained experience in data preprocessing, feature engineering, and model evaluation using Python (pandas, NumPy, scikit-learn)

KEY PROJECTS

MANCE: Morphology-Aware Nested Character Embeddings | 2024 – 2025 | [Github](#)

- Novel word representation framework leveraging nested character-level embeddings to capture morphological structure, linguistic patterns, and subword semantics across multiple granularities; designed hierarchical three-tier architecture

processing character n-grams (bigrams, trigrams), syllable-level units, and complete word forms to extract rich morphological features; implemented attention-based fusion mechanism to dynamically weight contributions from different embedding levels based on word complexity and context; achieved superior performance on word similarity benchmarks (15% improvement over FastText), morphological analogy tasks (22% gain), and rare word handling; particularly effective for morphologically-rich languages and out-of-vocabulary word representation; validated on multiple NLP downstream tasks including named entity recognition, part-of-speech tagging, and sentiment analysis.

BIE: Bit-Index Encoding for Neural Network Weight Compression | Dec 2024 – Jan 2025 | [GitHub](#)

- Novel weight compression framework achieving 40× compression at 95% sparsity with $MSE < 10^{-6}$; implemented three encoding variants (standard, run-length, and hybrid) optimized using Numba JIT compilation; benchmarked across 253 experimental configurations testing various sparsity levels, network architectures, and precision settings; demonstrated efficient storage and reconstruction of sparse neural network weights.

Hybrid RAG-Enhanced Deepfake Detection | Jul 2024 – Aug 2024 | [Zenodo](#)

- Privacy-first deepfake detection system combining Retrieval-Augmented Generation with visual anomaly detection; achieved 94.8% accuracy and 94.6% F1-score on diverse deepfake datasets; designed two-stage architecture with FAISS vector database for efficient similarity search and custom CNN for detecting facial inconsistencies, lighting artifacts, and temporal anomalies; integrated multi-modal analysis for robust detection across various manipulation techniques.

Large-Scale Model Implementations from Scratch | 2024 – 2025 | [Gpt-2 Github](#) | [MySQL Query Generator Github](#)

- GPT-2 complete 124M parameter implementation with custom byte-pair encoding tokenization, multi-head attention mechanisms, and training pipeline including learning rate scheduling and gradient accumulation; MySQL Query Generator implementing sequence-to-sequence architecture with attention for natural language to SQL conversion, supporting complex queries with joins, aggregations, and nested subqueries.

Clinical Trial Similarity Analysis | Sep 2024 | [Github](#)

- Semantic grouping system using PubMedBERT embeddings achieving 99.5% similarity accuracy and 0.891 top-10 retrieval precision for pharmaceutical research applications; implemented cosine similarity matching with FAISS indexing for efficient search across 10,000+ clinical trial descriptions; built clustering pipeline to identify related studies and potential trial overlaps.

AI Dashboard Creator | 2024 – 2025

- Automated dashboard generation system with intelligent data profiling detecting data types, distributions, and statistical properties; transformer-based PII detection achieving 98.5% accuracy for identifying and masking sensitive information; ML-powered visualization recommendation engine suggesting optimal chart types based on data characteristics; interactive builder with 15+ chart types, real-time filtering, and drill-down capabilities; automated statistical analysis including correlation matrices and outlier detection; export functionality to PDF, PNG, and interactive HTML with WebSocket-based real-time updates and responsive design.

Deep Learning Architecture Implementations | 2024 – 2025

- Built foundational architectures from scratch including Transformers with multi-head attention achieving 85%+ sentiment classification accuracy, Graph Neural Networks (GCN: 79.3%, GraphSAGE: 76.8%, GAT: 81.9%) for node classification tasks, Diffusion Models (DDPM with 130K parameters) for MNIST image generation using forward and reverse diffusion processes, GANs with discriminator-generator adversarial training and VAEs with KL divergence regularization for image synthesis, custom Attention Mechanisms including scaled dot-product and multi-head variants, demonstrating deep understanding of backpropagation, optimization, and modern ML architectures.

PUBLICATIONS

- **Gruhesh Sri Sai Karthik Kurra** et al., "Morphology-Aware Nested Character Embeddings for Word Representation," 2025 IEEE 7th ICCMMLA, Accepted, Forthcoming 2025
- **Gruhesh Sri Sai Karthik Kurra** et al., "Dynamic Auto-Finetuning of Language Models," IJNRD, vol. 9, no. 11, Nov. 2024 | [Publication](#)
- **Gruhesh Sri Sai Karthik Kurra** et al., "Global Remote RAM Sharing Framework," IJNRD, vol. 9, no. 12, Dec. 2024 | [Publication](#)
- **Gruhesh Sri Sai Karthik Kurra** et al., "Voice-Activated AI for Seamless Computer Interaction," IJNRD, vol. 10, no. 1, Jan. 2025 | [Publication](#)
- **Gruhesh Sri Sai Karthik Kurra**, "BIE: Bit-Index Encoding," Zenodo, DOI: 10.5281/zenodo.17217218, Sep. 2025 | [Publication](#)
- **Gruhesh Sri Sai Karthik Kurra**, "Hybrid RAG-Enhanced Deepfake Detection," Zenodo, DOI: 10.5281/zenodo.16732053, Aug. 2024 | [Publication](#)

TECHNICAL SKILLS

Programming: Python, R, JavaScript, TypeScript, Java, C, C++, HTML, CSS, PHP, SQL

AI/ML Frameworks: TensorFlow, PyTorch, Hugging Face Transformers, LangChain, LlamaIndex, scikit-learn, Azure OpenAI, GPT Models

Web Development: React.js, Next.js, FastAPI, Node.js, Express.js, React Native, TailwindCSS, Vite

Cloud & DevOps: AWS (EC2, S3, Lambda, Solutions Architect), Azure (Active Directory, OpenAI, Graph API), GCP, Docker, Git, GitLab, CI/CD, WSO2

Databases: PostgreSQL, MongoDB, MySQL, Supabase, Firebase, Redis

AI Specializations: LLMs (GPT-2 implementation – 124M params, fine-tuning, prompt engineering, optimization); Generative Models (GANs, VAEs, Diffusion Models – DDPM); Transformers (multi-head attention, positional encoding, sentiment analysis – 85%+ accuracy); Graph Neural Networks (GCN, GraphSAGE, GAT); Model Compression (quantization, pruning, bit-index encoding – 40× compression); Vision & NLP (OCR, ViT, object & deepfake detection – 94.8% accuracy, PubMedBERT, NL2SQL); Advanced (RAG, uncertainty quantification, synthetic data generation).

CERTIFICATIONS

- AWS Certified Solutions Architect Associate (2023)
- AWS Certified Cloud Practitioner (2023)
- Oracle Cloud Infrastructure Architect Associate (2023)
- Google TensorFlow Developer Certificate (2024)
- RedHat Certified Enterprise Application Developer (2024)
- Oracle Cloud Infrastructure Generative AI Professional (2024)
- Automation Anywhere Certified Advanced Automation Professional (2024)
- WSO2 Certified Micro Integrator Practitioner V4 (2025)
- WSO2 Certified Micro Integrator Developer V4 (2025)

AWARDS & ACHIEVEMENTS

- 1st Place – University Webathon, KL University (2022): Developed comprehensive diet management platform in 4-hour hackathon featuring personalized meal planning, calorie tracking, and nutritional recommendations. Competed against 50+ teams and presented solution to panel of industry judges, demonstrating rapid prototyping skills and ability to deliver functional applications under tight deadlines.
- 2nd Place – Design Expo, KL University (2022-2023): Created Arduino-based smart switchboard for IoT home automation with real-time energy monitoring and mobile app control. Designed complete hardware-software integration with automated scheduling features and received recognition for innovative approach to energy efficiency and user-centric design.
- National Tennis Ranking of 395 in India, AITA (2014-2019): Demonstrated athletic excellence and time management while competing in 100+ national and state-level tournaments across India. Balanced rigorous training schedule with academic commitments, developing discipline and resilience that translated into strong work ethic in technical pursuits.
- Technical Lead, Student Activity Council, KL University (2023-2025): Mentored 20+ juniors in web development and machine learning; organized workshops for 100+ students; awarded SAC Momentum Award. Developed SAC website with automated certificate generation system and coordinated team of 5 members to execute 15+ technical events, fostering collaborative learning environment.

EXTRACURRICULAR ACTIVITIES

- Community Volunteer, Swetcha Foundation, Collaborated with Swetcha Foundation to conduct village outreach drives, interacting with local residents to identify real-world challenges faced by daily wage workers. Proposed a mobile application concept enabling workers to discover and apply for nearby temporary jobs based on skill, pay, and location. Designed and presented a problem statement and prototype (as shown in attached chart) at the foundation event and received a Certificate of Appreciation for social innovation; Hyderabad, India; 2021 – 2022.
- Volunteered Street Cause (Social Service NGO), participating in community outreach programs, supporting underprivileged communities through educational initiatives and resource distribution. Contributed to organizing awareness campaigns on digital literacy and technology access, demonstrating commitment to social responsibility and community development and awarded Certificate of Service; Hyderabad, India; 2023 – 2024.
- Served as Technical Lead, Student Activity Council (SAC), KL University; mentored 20+ juniors in web development and machine learning fundamentals; organized technical workshops, coding bootcamps, and hackathons for 100+ students; developed SAC website with database for user management and automated certificate generation; managed team of 5 coordinators; awarded SAC Momentum Award for technical leadership and community building contributions; KL University, Hyderabad, India; 2023 – 2025.