KESHAV BHARADWAJ VAIDYANATHAN

(650)-660-0341 | Boston, MA | keshavbharadwaj98@gmail.com| LinkedIn | GitHub | Portfolio

TECHNICAL SKILLS

• **AI & Machine Learning:**, Large Language Models (LLMs), Multi-Agent Systems, RAG, Document Intelligence, Prompt Engineering, Semantic Search, Computer Vision, Generative AI
• **Tools & Frameworks:** LangGraph, LangChain, AWS (SQS, S3, EC2), Docker, Kubernetes, Git, CI/CD, ElasticSearch, Postgres
• **Packages & Libraries:** PyTorch, PyTorch-Lightning, OpenCV, NumPy, Pandas, FastAPI, Flask, SciKit-Learn, Hugging Face, ResNet, YOLO, Mask2Former
• **Languages:** Python, SQL, R, Shell Scripting, HTML

WORK EXPERIENCE

AI Application Engineer | UNAR Labs — *May 2024 – Present*

- **Multi-Agent Remediation:** Architected a distributed Multi-Agent LLM system using LangGraph to parse, audit, and remediate PDF, PPTX, and DOCX files for WCAG/Section 508 compliance, optimizing performance via localized agent execution loops.
- **Document Intelligence:** Engineered a multimodal RAG context pipeline parsing files into granular layout primitives (text, tables, assets) to generate automated document-level and page-level accessibility compliance audit reports.
- **Scalable Infrastructure:** Designed a concurrent microservice orchestration layer using AWS SQS and LangGraph to parallelize agentic processing workflows, ensuring fault-tolerant multi-threading and scalable content refining.
- **Hybrid Image Retrieval System:** Built an ElasticSearch semantic search framework with CLIP embeddings and S3 storage, delivering low-latency multimodal retrieval.
- **AI-driven SVG Vectorization Engine:** Developed an image-to-SVG workflow mapping raster graphics to primitives via ResNet routing, Hough transforms, and LLM feedback.

Computer Vision Research Assistant | Silicon Synapse Lab — *Feb 2024 – May 2024*

- **Perception:** Integrated real-time scene segmentation and object detection into robot perception domains; preprocessed RGB data to achieve a 0.78 IoU for ground classification.
- **Model compression:** Accelerated segmentation models (YOLO, YOSO, Mask2Former) via network quantization for NVIDIA Jetson, yielding a 2x reduction in inference time and memory footprint.

Data Science Co-op | Abiomed Inc (Johnson & Johnson) — *Feb 2023 – Sept 2023*

- **Predictive Modeling:** Created time-series cardiac output prediction models to process high-frequency device data, reducing error by 5.2%.
- **Deep Learning:** Implemented a Domain-Adversarial Neural Network (DANN) framework to forecast aortic pressure across cohorts with a 0.82 RMSE loss.
- **Threshold Tuning:** Optimized deep neural network classification boundaries, improving the Brier skill score by shifting thresholds from 0.23 to 0.5.

Machine Learning Trainee | Mad Street Den — *Dec 2020 – Jul 2021*

- **Recommendation Systems:** Engineered ML models for tag generation in an ElasticSearch recommendation system, fine-tuning scoring functions to achieve a 67% match for Top-1 and 82% match for Top-5 recommendations.
- **Infrastructure:** Deployed Seq2Seq LSTM OCR and CNN text classification models on Kubernetes, reducing monthly OCR operations costs by $10,000.

EDUCATION & CERTIFICATIONS

M.S., Electrical & Computer Engineering | Northeastern University (GPA: 3.88/4.00) — *Sep 2021 – Dec 2023*
B.E., Electronics & Communication | Visvesvaraya Technological University (GPA: 8.43/10) — *Aug 2016 – Aug 2020*
Applied Generative AI Graduate Specialization | Purdue University — *Jul 2025 – Nov 2025*

PATENTS & KEY PROJECTS

Project: Multimodal Visual Question Answering (Transformers) — *Jan 2022 – Apr 2022*

• Developed a transformer-based LxMERT framework, accelerating accuracy from a 57% LSTM baseline to 70.68% on the VQA dataset.

Project: SegMask for 3D Object Detection (Sensor Fusion) — *Sept 2022 – Dec 2022*

• Substituted coordinate masks with segmentation masks in Frustum-PointPillars (RGB/LiDAR), boosting KITTI dataset detection by 3%.

Patent: Q-CerGen (Quick Certificate Generator App) — *Mar 2021*

• Built a Flask, OpenCV, and Tkinter application seamlessly generating 3000+ custom interactive e-certificates/e-trophies.

Patent: ADAM (Automatic Disassemble & Assemble Robotic Machine) — *Jun 2021*

• Directed an autonomous robotic pick-and-place manipulator using SIFT descriptors for structural reassembly via 3D OpenGL UI.