Ashwanth Kuppusamy

JR. MACHINE LEARNING ENGINEER

Corvallis OR | 5412862753 | kuppusaa@oregonstate.edu

Profile

Result oriented Machine Learning Engineer with hands-on experience in the full model development lifecycle, from creating **automated labeling frameworks** to optimizing training pipelines. Leverages expertise in **continual learning** of **foundational models**, **pretraining**, and **domain adaptation** to enhance model performance. Proficient in PyTorch, OpenCV, and MLOps tools (Docker, Azure DevOps), with a focus on building efficient computer vision solutions.

Skills & abilities

- Python, C++, SQL, Git, Docker, Azure DevOps, Jenkins, Valohai MLOps, JFrog, JavaScript
- pyTorch, PySide6, OpenCV, CUDA, ONNX, tensorRT, FFCV, pandas, NumPy, matplotlib

Experience

Microtec | Jr. Machine Learning Engineer **Jun 2024 – Present**

Automated Labeling Framework **Part-time / 6 Months Full time**

- Designed and deployed an automated labeling system (object detectors + SAM), cutting labeling time by **50%** and saving **100000**+ dollars.
- Scaled training on dataset with 300,000+ images across 20 wood defect classes under varying lighting, resolutions, and textures to ensure generalization in production.
- Implemented replay-based **continual learning** to prevent catastrophic forgetting during incremental updates.
- **Transfer learned** features from large scale pretrained weights to accelerate model training, improve performance, and mitigate overfitting
- Built end-to-end **MLOps** pipelines using Docker, Jenkins and JFrog for training and deployment.

Training enhancements

- Used supervised and unsupervised pretraining techniques to reduce training data by at least **30%**
- Developed semantic segmentation nets for visually similar wood defects classes with low saliency
- Developed Negative instance mining loss to reduce false positive instances by more than **90%**
- Improved model performance on long-tail distributions by analyzing **t-SNE** visualizations of model embeddings to identify and address underperforming sub-class clusters.
- Created custom metrics for better model selection by combining instance and pixel metrics.
- Conducted **A/B testing** of segmentation architectures, balancing IoU, inference speed, memory efficiency and other custom metrics

Assisted Labelling Tool

- Architected and delivered an end-to-end assisted labeling tool (PySide6 + SAM backend) for prompt-based automated segmentation, **doubling** labeling throughput across the organization

- Implemented novel **caching** algorithm to precompute the SAM embeddings and ensured real-time labeling performance on CPU only machines for images as large as 21000 x 7000 pixels
- Engineered annotation tools including magnetic lasso tool, mask addition and subtraction.
- Led **full product lifecycle** from development to production deployment, managing beta release cycle, incorporating user feedback, and replacing old annotation tool used across the organization
- Built comprehensive performance monitoring system with granular time tracking across operations to identify optimization opportunities to improve tool efficiency by measurable metrics

Education

Oregon State University, MS Computer Science **Sep 2023 – Dec 2025**

Major: Computer Science | Minor: Artificial Intelligence

Master's thesis

Super-resolution attack on Thumbnail preserving encryption **Jun 2024 – Present**

- Implemented a two-stage non-blind image restoration pipeline based on **DiffBIR** for breaking TPE image encryption
- Applied **SwinIR** as the restoration module, achieving +73% SSIM and +33% PSNR improvement over baseline GAN method.
- Trained a **ControlNet-guided Stable Diffusion model** to reconstruct high-resolution details from SwinIR outputs.
- Improved decryptable block size from **8 to 16** over pix2pix GAN baseline for 256 x 256 images

Projects

ConvMAE SSP **Jun 2025**

- Architected novel MAE inspired MIM pre-training for conv transformer hybrid QT-ViT architecture
- Implemented custom masked linear attention and masked convolution blocks for pretraining
- Improved classification accuracy of 12M parameter model from **78.6 to 80.18** on ImageNet-100
- Built multi-scale feature fusion decoder processing 4 resolution levels (H/4 to H/32) with transformer blocks and positional encoding for image reconstruction
- Developed distributed training pipeline using PyTorch DDP with cosine LR schedule, gradient synchronization, and mixed precision for efficient ImageNet pre-training.

Brain tumor segmentation **Jan 2024**

- Investigated 2D/3D UNet, ResUNet and Attention-ResUNet for brain tumor segmentation.
- Developed a Ensemble of 2D Attention ResUNet architecture and improved IOU score by **6%**.
- Applied data augmentation such elastic deformations, intensity shifts to simulate brain tissue variability and scanner inconsistencies to mitigate overfitting on limited medical datasets.
- Obtained IOU score of **0.8** with 3D ResUNet and **0.79** with Ensemble of 2D Attention-ResUNet.

Publications

"Inference at the Edge for Complex Deep Learning Applications with Multiple Models and Accelerators" 2023 14th (ICCCNT), Delhi, India, 2023, pp. 1-7, doi: 10.1109/ICCCNT56998.2023.10306363