

Alex Mathew John

mail@l3xj.com | l3xj.com | linkedin.com/in/lexicalyoda/ | github.com/lexical-yoda | +919496551881

Machine Learning Engineer with 4+ years of experience building production computer vision systems from first principles. Built an entire real-time golf shot analysis platform from scratch — no prior models, datasets, or benchmarks existed — as part of a 2-person engineering team. Core expertise in custom neural architecture design, camera geometry, physics simulation, and on-device ML deployment (CoreML/iOS).

EXPERIENCE

Senior Machine Learning Engineer, AsmiiovX

August 2022 – Present

System Architecture & Neural Networks

- Designed and built the full real-time golf shot analysis pipeline from zero: ball detection (YOLOv5/v8/v11), club impact detection (custom ChangeNet), ball flight segmentation, 3D speed/angle computation, and carry distance prediction. Achieved sub 5% error against professional golf monitors (GC2/GC3) with sub-5ms inference on iPhone.
- Built a custom dual-stream change detection network with MobileNet encoder, multi-scale feature concatenation, and deconvolution decoder. Also developed a 3D CNN with temporal convolution for multi-frame impact classification and experimented with a Vision Transformer variant (ChangeFormer) with Siamese encoding.
- Trained MobileNetV2 segmentation models (Dice/Tversky loss), ResNet50 binary classifiers, and YOLO detectors across 5 environments (indoor, net indoor, net outdoor, range, putting). All production models converted to CoreML.
- Built adaptive image preprocessing including Rec.709 luminance-based brightness scoring with percentile trimming, and automatic brightness correction to normalise inputs across varying lighting conditions on the range.

Applied Physics & Camera Geometry

- Built a carry distance physics engine using RK4 numerical integration with aerodynamic drag/lift modelling, Reynolds number-dependent coefficients, spin decay, and terrain bounce simulation across multiple surface types.
- Developed ball speed and launch angle calculation from first principles: 3D reconstruction from 2D pixel detections via camera intrinsics and lens distortion correction. Engineered and benchmarked 5 independent distance estimation methods across thousands of shots to determine production method.
- Built the camera calibration pipeline: checkerboard/ChArUco detection, intrinsic matrix computation, per-device distortion profiles, and YOLO bounding box undistortion. Supported both OpenCV and Apple's additive radial lens distortion models.
- Developed a 3D trajectory visualisation system for the customer portal, rendering full ball flight paths with spin, altitude, and landing data from the physics simulation output.

Data Infrastructure & Leadership

- Designed the data pipeline: S3 storage architecture, DynamoDB shot indexing, concurrent sync systems, and automated dataset generation with train/test splitting and date-based cutoffs across 9,000+ annotated shots.
- Built a 5-path model evaluation framework (Reference, On-Device Log, Recalculated, CoreML, PyTorch) with automated error analysis, per-phone-model breakdowns, and grouped Excel reporting for systematic accuracy tracking.
- Owned the entire ML stack as one half of a 2-person team for 2+ years. Later hired and managed a third engineer and annotation team. Implemented AES-256 encryption for CoreML model IP protection.

AI Engineer, Accubits Technologies

March 2020 – July 2021

- Built face mask detection (ResNet50, 25K samples/class, 96% accuracy) and food image classification pipelines with automated data collection, improving accuracy by 20%.
- Deployed ML services on AWS (Docker, Lambda, EC2) for scalable inference across multiple projects.

SKILLS

Computer Vision: YOLO (v5/v8/v11), OpenCV calibration, lens distortion correction, segmentation, frame differencing

Deep Learning: Custom ChangeNet, 3D CNNs, Vision Transformers, MobileNet, ResNet, U-Net, Siamese networks

Deployment: CoreML (iOS), model quantisation, AES-256 encryption, TensorFlow, PyTorch, Keras

Applied Math: Projectile motion, aerodynamic simulation (RK4), camera geometry, 3D reconstruction, pinhole model

Infrastructure: AWS (S3, DynamoDB, Lambda, EC2), Docker, Git, Python, Flask, FFmpeg, Pandas, NumPy

HomeLab: Self-hosted 4-node infrastructure — Proxmox, Docker, WireGuard, nginx, Prometheus/Grafana, local LLM inference

EDUCATION

University of Exeter — MSc in Data Science (2:1) — Dec 2022

APJ Abdul Kalam Technological University — B.Tech in Computer Science — July 2019