Jay Anil Dodia

+1(445)-208-7662 | jayanildodia@gmail.com | linkedin.com/in/jayanildodia/

Summary

Computer Vision Engineer with hands-on experience building production-grade pipelines for person, cart, and action recognition in real-world retail environments. Skilled in deploying models like YOLOv26, Grounding DINO, and MMAction2, and integrating them with DeepStream, Redis, and Milvus for real-time analytics. Known for translating research into scalable, camera-ready systems and collaborating across MLOps, detection, and hardware teams to deliver robust, end-to-end CV solutions.

Technical Skills

Programming & Scripting: Python, Bash, JavaScript, SQL, HTML/CSS
Machine Learning & Computer Vision: YOLOv8/v11, Grounding DINO, SAM2, ST-GCN, PoseC3D, DeepStream, OpenCV, scikit-learn, NumPy, Pandas, Matplotlib
Cloud : AWS (EC2, S3, Lambda, IAM, CodePipeline), Azure (VMs, AD, DevOps, Blob Storage), GCP (GKE, Cloud SQL)
DevOps & Infrastructure as Code: Terraform, Docker, Kubernetes (EKS/GKE), Jenkins, GitHub Actions, Ansible
MLOps & Data Infrastructure: RabbitMQ, Redis, Milvus, Azure Blob Storage, CRON automation, Git, GitHub
Monitoring & IT Operations: Prometheus, CloudWatch, OpsRamp, ServiceNow, Azure Monitor
AI Tools Productivity: ChatGPT, GitHub Copilot, Claude, Notion AI

Work Experience

Computer Vision Engineer
Oct 2025 – Present

Walmart, *Dallas, TX*

- **Leading a stealth Proof of Concept (POC)** for **AI-driven theft prevention** across Walmart **international stores (Chile, Mexico, Canada)**, owning **end-to-end strategy, research direction, system architecture, and implementation**.
- Driving **project vision and execution** through **daily standups**, milestone planning, task assignment, success metric definition, and **stakeholder demos** to secure funding for production rollout.
- Designing **real-time theft detection pipelines** on **live surveillance camera feeds** using constrained, low-quality infrastructure; optimizing for **suboptimal angles, compression artifacts, and variable lighting conditions**.
- Architecting **action recognition systems** leveraging **pose estimation, human detection, object tracking**, and cart **"safe zone" modeling** to identify **concealment behaviors** (e.g., hiding merchandise inside clothing or jackets).
- Implementing and experimenting with state-of-the-art models including **YOLOv8, YOLOv11, YOLOv26** for detection, **SAM2/SAM3** for segmentation-driven localization, and **MMAction2** for temporal action modeling; evaluating grounding-based approaches using **GroundingDINO**.
- Developing **temporal reasoning pipelines** correlating **pose keypoints, tracked objects, and object disappearance events** to infer potential theft scenarios over time.
- Addressing **data scarcity challenges** by curating limited in-store theft footage, adapting **public datasets** to retail concealment scenarios, **manually annotating** custom training and evaluation datasets, and generating **AI-synthesized theft videos** to improve model robustness.
- Building **employee-facing UI integrations** within existing internal applications to **flag suspicious activity**, generate **short video clips**, and **redirect store associates to specific aisles** for real-time intervention.
- Designing **human-in-the-loop validation workflows** where store associates confirm flagged events, enabling continuous **model refinement and false-positive reduction**.
- Defining a **scalability roadmap** for upgraded camera infrastructure and broader international deployment, positioning the POC as a **multi-million-dollar annual shrink mitigation initiative**.

Computer Vision Engineer
Sept 2024 – Sept 2025

Sams Club (Walmart), *Dallas, TX*

- Developed and fine-tuned **FC-CLIP** for human and cart detection; also experimented with **FrozenSeg** and **X-CLIP**, and integrated **SAM2** for segmentation-based detection — improving detection accuracy by **12%**.

- Implemented **Grounding DINO** and **YOLOv8/v11** for person/cart detection and human pose estimation, achieving **20% accuracy improvement** over FC-CLIP in complex retail scenarios.
- Spearheaded **skeleton-based action recognition** using models like **ST-GCN**, **PoseC3D**, **SlowFast**, and **YOWO** to recognize behaviors (e.g., item pick/place actions), transitioning from fragile rule-based logic to robust deep learning solutions.
- Designed and optimized the in-house **DeepStream-based real-time video pipeline** for live store camera feeds, integrating detection, ReID, and metadata generation through **RabbitMQ** to support downstream behavior analytics in the MTMC system.
- Engineered the **Time Series Service**, a core MTMC module, to ingest real-time detections and synchronize person/cart metadata across **MilvusDB** and **Redis**, enabling consistent global ID assignment across camera views.
- Enhanced **Cart ReID** by implementing **1D/2D ArUco code detection** on price tags and carts for reliable tracking — improving cart identification accuracy and reducing ID ambiguity during tracking tasks.
- Led **failure case analysis** of detection models by comparing **Grounding DINO**, **YOLOv8**, and internal models against curated ground truth, identifying weak points and training on false positives/negatives — boosting accuracy by **5%**.
- Supported critical **manual labeling** using **CVAT** and internal tools to recover lost annotations, and contributed to real-time **item-bin mapping** during live executive demos by accurately masking bins using UPC metadata.
- Collaborated with the **MLOps team** to automate the data pipeline: developed **CRON jobs** to detect new data in **Azure Blob Storage**, update manifest files, and auto-convert datasets — reducing data prep time by **40%**.
- Actively participated in **in-club visits** for camera setup, calibration, and scenario-specific data collection by simulating realistic shopping behavior — supporting teams outside core scope to ensure high-quality data readiness.
- Consistently partnered with cross-functional teams (**Detection, MTMC, MLOps, Action Recognition, Hardware**) to maintain scalable, accurate, and production-grade computer vision systems across Sam's Club retail stores.

Cloud DevOps Engineer
Oct 2023 – October 2024

Virufy, *Los Altos, CA*

- Reviewed and refined the **AWS cloud architecture** and internal documentation to serve as a single source of truth for the MLOps team — improving visibility and reducing handoff errors
- Partnered with senior engineers to design a **CRON-based automation pipeline** that spun up reproducible test environments on demand, accelerating iterative model validation
- Implemented CI/CD workflows with **GitHub Actions** and **AWS CodePipeline**, integrating code quality checks, tests, and zero-downtime deployments
- Adopted **Terraform**-based Infrastructure-as-Code to manage infrastructure changes with version control and rollback support across multiple environments
- Managed cross-functional delivery cycles using **Jira** and **Confluence**, ensuring transparency and alignment across backend, ML, and DevOps teams
- Assisted in establishing IAM policies, secure VPC networking, and encrypted S3 storage policies to align with best practices for production readiness and internal compliance reviews

IT Infrastructure Specialist
Jun 2022 – Dec 2022

(Co-Op), SEPTA, *Philadelphia, PA*

- Designed and deployed a custom **Alexa Skill** for real-time regional rail tracking, streamlining transit information access for Philadelphia commuters via voice interaction
- Built an intuitive voice interface using the **Alexa Developer Console**, integrating with SEPTA's live data feed to provide accurate train schedules and alerts
- Secured the application with role-based access via **Azure AD**, and enforced strict network control using security groups and **WAF** for external traffic filtering
- Simulated real-world commuter flows and conducted performance testing to ensure scalability and consistent response times under varying loads
- Collaborated with engineers, security architects, and operations staff to deliver the end-to-end project under tight timelines, resulting in a stable and reliable production deployment

IT Infrastructure and Operations Analyst *Oct 2020 – Aug 2021*
Wolters Kluwer, *Chennai, India*

- Developed and maintained CI/CD pipelines using **AWS CodePipeline** and **Azure DevOps**, replacing manual deployments with automated workflows across development and production environments
- Designed and deployed scalable VM infrastructures and virtual networks across **Azure**, supporting critical applications used by global clients in the Legal and Regulatory space
- Provided day-to-day operational support for Azure IaaS and PaaS workloads through **OpsRamp** and **ServiceNow**, resolving service requests and infrastructure issues across time zones
- Implemented automated patching schedules and health checks across hybrid environments, ensuring compliance with internal SLAs and minimizing downtime during maintenance windows
- Configured **Azure AD-based RBAC**, security groups, and custom policies to enforce access control across infrastructure layers, improving audit-readiness and user governance
- Collaborated closely with engineering and infrastructure teams to modernize legacy deployment workflows and improve DevOps maturity across departments

Certifications

AWS Certified Cloud Practitioner | Amazon Web Services *Aug 2022*
AZ 900: Azure Fundamentals | Microsoft Azure *Mar 2021*

Education

Drexel University | Philadelphia, USA *Sept 2021 – Jun 2023*
Master of Science in Computer Engineering | *GPA: 3.5*
Relevant Coursework: Applied Machine Learning, Machine Learning and AI, Cloud Technology

Projects

Multi-Cloud Migration for Luxury Hospitality Brand | *Docker (GCR), Kubernetes (GKE), Terraform, GCP, Amazon S3*

- Migrated on-premises applications and databases to a multi-cloud, optimizing resource allocation and operational efficiency.
- Realized significant cost savings through efficient resource allocation and utilization on Google Cloud Platform and AWS.
- Strengthened scalability and resilience with Google Kubernetes Engine, fortified database security within Google Cloud SQL, and ensured high availability of assets through Amazon S3, ensuring uninterrupted service delivery and data integrity.

Secure Django Login and Authentication System | *Python, Django, HTML, CSS*

- Engineered a Secure User Authentication System by implementing a full-stack sign-up and login system using Django, enhancing application security with advanced password hashing, validation, and reset functionalities.
- Configured GitHub Actions for automated testing and deployment, ensuring rapid, reliable, and consistent software delivery.
- Created a user-friendly, responsive interface using HTML and CSS, improving user experience across various devices.

Real-time Twitter Sentiment Analysis Platform on AWS | *Python, AWS Lambda, AWS Kinesis, Amazon QuickSight*

- Developed a function tailored for real-time collection and analysis of up to 1000 tweets/min, employing targeted keyword filtering.
- Utilized AWS services for real-time data streaming, sentiment analysis, and trend visualization, enabling proactive response to emerging trends and public sentiment for brand optimization.

Secure and Highly Available Website Hosting with AWS | *Amazon EC2, Amazon S3, Route53, CloudFront*

- Designed and implemented a highly resilient website hosting architecture on AWS, emphasizing scalability and security to accommodate varying traffic loads, ensure uninterrupted service availability.
- Integrated a bunch of AWS Services to achieve durable, scalable storage, efficient content delivery, and precise DNS management to guarantee a seamless and secure user experience.