Appsin.Ai

SERVICE

Object Recognition and Computer Vision Development for Real Cameras

Custom computer vision models that detect, count, read and inspect what your cameras see, trained on your own images and run in the cloud or on edge devices.

https://appsin.ai · +91-9025555584 · sales@appsin.ai
AI development for businesses in India, the GCC and the USA

OVERVIEW

Computer Vision

We build computer vision systems that recognise what matters in your images and video, from defects on a line to plates at a gate or text on a label. Models are trained and tested on footage from your own cameras and run in the cloud, on-premises or on edge devices.

Detection finds where objects are, recognition says what they are and segmentation outlines them exactly. We choose the approach and model, from YOLO-family detectors to segmentation and OCR models, built with PyTorch, TensorFlow and OpenCV, then optimise it for the hardware it has to run on.

Accuracy depends on your data more than the algorithm. Lighting, camera angle, rare defects and look-alike products all matter, so we start by reviewing sample images, agree how results will be measured, and test on footage the model has never seen before anything goes live.

BUILT FOR

- Manufacturers with visual quality checks
- Retailers and warehouses tracking stock
- Parking, logistics and fleet operators
- Facilities and site security teams
- Healthcare and agriculture teams working with images

FEATURES

What you get with Computer Vision

01 **Object detection and counting**
Locate and count products, vehicles, people, animals or parts in images and live video, with bounding boxes and a confidence score for each.

02 **Visual inspection and defect detection**
Spot scratches, dents, missing parts, wrong labels and packaging faults on production lines, and flag items for a person to check.

03 **Number plate recognition**
Read vehicle plates for parking, gated access, tolling and yard management, trained on the plate formats used in your region.

04 **OCR and text reading**
Extract text from documents, labels, meters, containers and screens in several languages and fonts, with fields mapped into your system.

05 **Face recognition with consent**
Face matching for access control and attendance where people have agreed to it, with face blurring, on-device processing and minimal storage options.

06 **Segmentation and measurement**
Pixel-level outlines to measure area, size or coverage, such as the damaged area on a vehicle panel or the spread of disease on a crop.

07 **Safety and site monitoring**
Detect missing safety gear, people in restricted zones, smoke or spills, and send an alert with the image to the right person.

08 **Edge and mobile deployment**
Models compressed and converted with ONNX or TensorRT to run on NVIDIA Jetson devices, existing CCTV servers or inside a mobile app.

09 **Dashboards, alerts and APIs**
Results saved with their image evidence, shown in a dashboard and sent to your own systems through APIs, webhooks or messages.

WHY IT MATTERS

The difference it makes to your business

01 **The same check on every item**
A model checks every item or frame in the same way, instead of a sample checked by a tired person at the end of a shift.

02 **Use the cameras you already have**
Many projects run on existing CCTV or line cameras, with new hardware only where image quality is not good enough.

03 **Evidence behind every result**
Each detection is saved with its image and timestamp, so decisions can be reviewed and disputes settled.

04 **Privacy designed in**
On-device processing, blurring and short retention keep personal images out of places they do not need to be.

HOW WE DELIVER

How we deliver Computer Vision

01 **Review sample images**
We study images or video from your real cameras and conditions to judge what is achievable and what data is missing.

02 **Collect and label data**
We plan which images are needed and label them, or build on your team's labels, keeping a separate set back for testing.

03 **Train and compare models**
Several approaches are trained and compared on precision, recall and speed on the hardware the model will run on.

04 **Pilot on site**
The model runs on one line, store or gate alongside current checks, so its results can be compared with what people find.

05

Deploy and retrain

We roll out, monitor accuracy and drift, and retrain with new images as products, lighting or cameras change.

FAQ

Questions about Computer Vision

What is the difference between object detection and object recognition?

Detection finds where objects are in an image and draws a box around each one. Recognition, or classification, says what each object is. Most projects use both, and some add segmentation, which outlines the exact shape of each object for measurement.

How many images do we need?

It depends on how varied your objects and conditions are. A narrow task with fixed cameras can often start with a few hundred labelled images per class, while rare defects or many look-alike products need more. Pre-trained models and augmentation reduce the amount, and we give you a plan after reviewing samples.

How accurate will the model be?

We cannot promise a figure before seeing your data. We agree measures with you, such as precision and recall for each class, and report them on a test set from your own cameras that the model never saw during training. A pilot then confirms performance in live conditions.

Can it work with our existing cameras?

Often, yes. Resolution, frame rate, angle and lighting decide whether existing CCTV or line cameras are good enough. We check sample footage first and recommend new cameras or lighting only where they are needed.

Do we need a trained model, or can a general AI vision model do it?

For low volumes or varied tasks, a vision-capable model such as GPT or Gemini can read images without training. For high volumes, real-time video or offline sites, a trained detector is faster, cheaper per image and more consistent. We often prototype with the first and move to the second.

How much does a computer vision project cost?

The main drivers are the number of object classes, the data collection and labelling effort, the accuracy you need, real-time requirements, the hardware and the number of sites. You get a fixed quote after we review your requirements and sample images.

Who owns the model and the data?

You do. On completion and final payment, the trained model weights, training code, labelled dataset and documentation we build for you are handed over, as set out in the contract, so you can retrain or move them as you choose. Third-party and open-source models stay under their own licences.

Will it work at night, in poor lighting or in bad weather?

It can, if the training and test images include those conditions. We collect footage across shifts, seasons and weather, report accuracy separately for difficult conditions, and recommend infrared cameras or extra lighting where the image itself is too poor for any model to read.

Can it run without an internet connection?

Yes. Models can run on an edge device or a server on site, so detection carries on when the connection drops. Results and selected images are stored locally and sent on to your dashboard and systems when the connection returns.

NEXT STEP

Show us what your cameras see

Send a few sample images or a short clip and we will tell you what can be detected reliably, what data is missing and how we would test it.

Book a free consultation: https://appsin.ai/book-a-consultation/

WhatsApp or call: +91-9025555584

About AppsInAi

AppsInAi Private Limited is an AI development company in Coimbatore, India. We build generative AI assistants, AI agents, chatbots, machine learning models, computer vision and data platforms on our clients' own data, from a first proof of concept to production, with monitoring and support after launch.

Contact
sales@appsin.ai
+91-9025555584
Housing Unit, Sf No. 393, CMS School Backside, 3rd Street, Nehru Nagar, Ganapathy, Coimbatore 641006, India
https://appsin.ai