

Building Visual AI

in 10 steps

Team

Form a cross-functional team of computer vision engineers, back-end engineers, mobile engineers, product managers, UX designers, and data labellers. Consider you will need a key stakeholder across all business functions, including operations, customer experience, and product.

Policy Alignment

Define your internal rule set to inform data labeling and business logic, adapting as needed for customer SLAs, UX preferences, driver training levels, package types, and geographical variations (e.g. city center vs. suburban).

Tooling

Develop a complete machine learning tech stack for model development and deployment, factoring in costs for tools including cloud platforms (GCP, AWS), data annotation (Encord, LabelStudio, Scale), and monitoring (Datadog, Weights & Biases).

Data Prep & Machine Learning

Collect and preprocess data for labeling. Set up labeling projects, create guides, and monitor consensus. Consider how your pipeline will integrate multi-modal techniques, including object detection, OCR, face detection, and GPS.

Model Performance

Build a ground truth dataset for evaluating the AI models, taking into account how you will continuously add new edge cases. Consider if you want to prioritise precision or recall and what level of autonomy required.

Monitoring

Setup a dashboard UI for your team to easily search for and audit images. Consider how you are going to build continuous monitoring and alerting for AI accuracy, latency, and system uptime once solution is deployed.

System Performance

Choose if you are going to host your AI models in the cloud or on-device and which devices you will support (Android, iOS, PDAs). Balance critical trade-offs between speed, accuracy, and user experience.

Legal & Data Security

Identify risks related to handling personal data in images and establish risk mitigation processes. Consider how you will anonymise PPI, handle information requests from customers, and maintain data compliance.

Fraud Prevention

Develop strategies to mitigate both end-user and customer fraud. Implement methods such as detecting duplicate images, analyzing images over time to identify changes, and integrating user-level and address-level analytics to generate a risk score.

Business Case

Assess the time, resources, and risks associated with the above steps. On average, it takes an enterprise with an existing machine learning team 12-18 months to deploy an AI model into production.