AI CCTV Monitoring Platform

White-label ready

Model guide

Choose, install and use the right model

What a model is, how to install one in a few minutes, and every model that could be added, job by job, in plain words. No machine-learning knowledge needed.

Version 1.1.804 Updated 8 November 2026 664 models for 63 jobs

For anyone setting up or running the platform

Start here: what a model is

A model is the platform's eyes

A model is a trained program. It looks at each picture from a camera and says what is in it: a person, a helmet, a vehicle, smoke.

One model does one job

Finding people is one job. Checking their PPE is another. Each job has one model running at a time, and you choose which.

Ready-made models install with one click

You do not train anything to begin. The platform's Model library holds ready-made models for finding people, checking PPE, spotting vehicles and more.

Try it before you trust it

The Test bench checks your own photos and videos with a model, so you see how it does on your site before it judges a live camera.

Every model has a licence

The licence says whether the model may be used in a business. Each model in this guide carries that answer in plain words.

Your site makes it better

When your team reviews events, the platform collects pictures to learn from. A model trained on your own cameras fits your site best.

Install a model in five steps

1. **Open Models.** It is in the sidebar under Improve. Administrators see it.
2. **Pick a model and a size.** In the **Model library**, find the job you need (people, or PPE) and choose a size. The one marked **Recommended** is a good start.
3. **Choose Download.** The model downloads and installs in the background. Its card shows the progress and then says Installed.
4. **Choose Activate.** The first time, the platform asks for a reason, because the model has not been scored on your own pictures yet. Enter one such as "first install".
5. **Check it on the Test bench.** Drop a few photos or a stored video from your site and look at what the model finds. Then let it run on your cameras.

Have a model file of your own? On the Models page, **Import model** takes a `.pt` or `.onnx` file, and **Add from URL** downloads one from a link.

Note. Downloading a ready-made model needs the internet. After that the site can run offline.

Which size to choose

Most ready-made models come in sizes. A bigger model looks more carefully and needs more computing power. Start small and go up only if the Test bench shows you need to.

Nano	**Choose it when** You have many cameras, or no graphics card. The fastest. **Needs** Runs on a normal processor (CPU)
Small	**Choose it when** You are not sure. The balance of speed and care that suits most sites. **Needs** A graphics card (GPU) helps
Medium	**Choose it when** People are far from the camera, or Small misses things. **Needs** A GPU is recommended for several cameras

Large	**Choose it when** You need more care still and have the hardware. **Needs** A GPU is recommended
XLarge	**Choose it when** Few cameras and the most careful result. The slowest. **Needs** A GPU is required

Speed and accuracy figures for any model are its makers' own. Nothing in this guide was measured on your site: the Test bench is where you find out.

How to read a model in this guide

Start here

The recommended first choice for the job.

Licence-safe alternative

The one to take when the first choice's licence has conditions you cannot meet.

Avoid

Listed so that nobody picks it by accident. The reason is beside it.

Allowed

Commercial use: the licence lets a business use it.

With conditions

Commercial use depends on something, and the list says what: for example a paid licence. Ask your supplier before you rely on it.

Not allowed

For research and trials only. Do not use it in a business.

Yes under Loads today

The platform can run it now.

Not yet

The model exists, and your supplier must add support for it before the platform can run it.

Not in the platform yet

The job itself is not part of the platform today. The models are listed for planning.

Unconfirmed

Its details could not be confirmed from the makers' own pages.

Good to know

A short technical note for whoever installs the model. If it means nothing to you, skip it: the labels above are what matter.

Models the platform can load today

These 40 load in the platform now. The ones for finding people and checking PPE install from the Model library; ask your supplier about the others.

Model	Job	Commercial use	Good to know
RF-DETR (N / S / M / L detection, plus Seg N to 2XL)	Person detection	✓ **Allowed**	The person detector in the app's catalogue (N, S, M, L; sha256 checked). Objects365 provenance is unconfirmed.
D-FINE (N / S / M / L / X)	Person detection	✓ **Allowed**	In the app's catalogue (S, M). Take the COCO-only checkpoints; the README says Objects365 ones are not cleared.
RT-DETR / RT-DETRv2 (Baidu / PekingU)	Person detection	✓ **Allowed**	Use the COCO-only checkpoints; first-class Transformers support.
Ultralytics YOLO11 / YOLO26 (and YOLOv8, v5)	Person detection	⚠ **With conditions** Enterprise License or AGPL terms.	Works today; AGPL applies even to models trained from scratch; no public price.
RF-DETR N/S/M/L (Apache sizes) fine-tuned on PPE data	PPE detection	⚠ **With conditions** Pretraining data terms not stated.	Licence-clean base to fine-tune on site crops plus CC BY / CC0 sets; no ready PPE weights exist. Training page: Detector, RF-DETR (learning rate 1e-4); the result is an ONNX model.
killuminati1/construction-ppe-yolov8	PPE detection	⚠ **With conditions** Enterprise License or AGPL terms.	Richest taxonomy among ready weights (19 classes); ear protection and boots are weak.
Ultralytics YOLO11 / YOLO26 trained in-house on PPE data	PPE detection	⚠ **With conditions** Enterprise License.	Fastest way to a working detector today; not for a closed sale without the Enterprise licence.
Hexmon/vyra-yolo-ppe-detection	PPE detection	⚠ **With conditions** Enterprise License or AGPL terms.	14 classes with "No-" forms; no metrics on the card.
thalostech2025/dfine-small-construction-ppe-v1	PPE detection	⚠ **With conditions** Fine-tune data source unknown.	Only permissive-tagged ready PPE detector with five gear classes; weak outside person, hard hat and vest. The author's own scores, out of 1, are vest 0.48 and hard hat 0.39: check it on your cameras.
keremberke/yolov8{n,s,m}-protective-equipment-detection	PPE detection	⚠ **With conditions** Enterprise License or AGPL terms.	Already in the catalogue; low accuracy, near-silent on site frames.
keremberke/yolov8{n,s,m}-hard-hat-detection	PPE detection	⚠ **With conditions** Enterprise License or AGPL terms.	Helmets only, already in the catalogue; reads caps as helmets on wide CCTV views.
Hansung-Cho/yolov8-ppe-detection	PPE detection	⚠ **With conditions** Enterprise License or AGPL terms.	Already in the catalogue; the MIT tag is wrong for an ultralytics model.
mancalazure/construction-ppe-safety	PPE detection	⚠ **With conditions** Enterprise License or AGPL terms.	Ready ONNX in the stock YOLOv8 layout; demo-grade and unmeasured.
MobileSAM and EfficientSAM	Open-vocabulary detection	✓ **Allowed**	Cheap box-to-mask on CPU; upstream LICENSE not fetched; EfficientSAM not checked.
YOLOE (THU-MIG YOLOE and Ultralytics YOLOE-11 / YOLOE-26)	Open-vocabulary detection	⊘ **Not allowed**	Fastest text-prompt detector, for labelling only; bake prompts offline and do not ship it.
RTMPose (MMPose) via rtmlib	Pose	⚠ **With conditions** Training-set terms not read.	Per-person keypoints on crops; the paired RTMDet detector is trained on Objects365.

Model	Job	Commercial use	Good to know
YOLO-pose (Ultralytics YOLO11 / YOLO26 pose)	Pose	⚠ **With conditions** Enterprise License or AGPL terms.	Only pose model that loads today, for development.
HF community YOLO face models (AdamCodd/YOLOv11n-face-detection, Bingsu/adetailer face_yolov8*, arnabdhar/YOLOv8-Face-Detection)	Face detection	⚠ **With conditions** Enterprise License or AGPL terms.	Only face detector that loads today, for internal tests; an Apache tag does not remove the AGPL.
Ultralytics YOLO11 / YOLO26 COCO weights, class 67 "cell phone"	Phone use	⚠ **With conditions** Enterprise License or AGPL terms.	In-house spike: no false alarms on 72 negatives, about half the phones found at about 220 px person height.
Beehzod/smoking-detection-finetuned	Smoking	⚠ **With conditions** Enterprise License or AGPL terms.	In the model library as "Cigarette — YOLO11m (experimental)", pinned to commit d81c4396. The smoking check runs it on person crops at conf 0.5, and a phone or food box overlapping the hit vetoes it. Finds about half of site-scale cigarettes.
Pose rule (hand-to-mouth) on RTMPose keypoints plus a small own-trained cigarette/vape detector	Smoking	⚠ **With conditions** Body7 training-set terms not read.	Licence-clean way to build smoking; needs your own labelled clips.
Beehzod/smoke_cigarette-detection2-yolo11m and best_smoke_cigarette-detection	Smoking	⚠ **With conditions** Enterprise License or AGPL terms.	Second cigarette detector to compare; a 0.994 mAP50 on near-duplicate frames says little.
rabahdev/fire-smoke-yolov8n	Fire and smoke	⚠ **With conditions** Enterprise License or AGPL terms.	Works today; use conf 0.5 or more for fire.
xiazh0219/yolo26m-UAV-fire-smoke-detection	Fire and smoke	⚠ **With conditions** Enterprise License or AGPL terms.	Aerial and outdoor smoke; not for ceiling or wall cameras.
Pyronear YOLO (yolov8s, yolov11s/n, yolo11s_*)	Fire and smoke	⚠ **With conditions** Enterprise License or AGPL terms.	Early smoke at long range from fixed towers; indoor steam and dust will trigger it.
RTMPose (rtmlib / MMPose) plus fall rules	Falls	⚠ **With conditions** Body7 training-set terms not read.	Loads now; the person-down pose rule is built on it.
Ultralytics YOLO11-pose	Falls	⚠ **With conditions** Enterprise License or AGPL terms.	Lets you try the fall rules before the Apache route exists.
Ultralytics YOLO26-pose	Falls	⚠ **With conditions** Enterprise License or AGPL terms.	In the model library as "Pose — YOLO26 (official)" (weights from ultralytics assets v8.4.0). Feeds the Person down check, method pose; s is the default.
melihuzunoglu/human-fall-detection	Falls	⚠ **With conditions** Enterprise License or AGPL terms.	Quick fall-by-posture detector; cannot tell lying down from a fall without a time rule.
mustafahacioglu/yolo11s-forklift-person-detection	Forklift	⚠ **With conditions** Enterprise License or AGPL terms.	Works today; its card adds a box-containment rule for operator presence; the map_up licence is unread.
yihong1120/Construction-Hazard-Detection	Forklift	⚠ **With conditions** Enterprise License or AGPL terms.	Works today as an alternative forklift source: it reports forklifts as "machinery", and the catalogue ignores its "vehicle" class because it false-alarms on site.
keremberke/yolov8{n,s,m}-forklift-detection (and yolov5 versions)	Forklift	⚠ **With conditions** Enterprise License or AGPL terms.	Tiny model for a first demo; a Roboflow-derived forklift model once boxed the whole frame.
RT-DETR and RT-DETRv2	Vehicles	✓ **Allowed**	Easiest DETR to prototype with in Transformers; avoid the AGPL ultralytics port.

Model	Job	Commercial use	Good to know
Ultralytics YOLO11 and YOLO26	Vehicles	⚠ **With conditions** Ultralytics Enterprise licence.	Only family that loads today; AGPL applies even to internal use or training from scratch.
PaddleOCR PP-OCRv5 recognisers (default, latin, arabic, devanagari, cyrillic, ta, te, th, korean, el)	Plate reading	✓ **Allowed**	Only permissive route to non-Latin scripts; not tuned for plates, so fine-tune on cropped plates.
OpenCV classical checks (Laplacian variance, histograms, edge density, ORB or phase-correlation shift, SSIM)	Camera health	✓ **Allowed**	Blur, dark, blocked, moved and frozen frames. Tune thresholds per camera and for day and night.
OpenCV image checks (Laplacian blur, brightness, SSIM, frozen frame)	Camera tamper and picture quality	✓ **Allowed**	Near-zero cost on every camera every minute. Thresholds per camera.
OpenCV ORB or AKAZE matching with a homography	Camera tamper and picture quality	✓ **Allowed**	Detects that the camera moved and by how much, so zones can be marked invalid. Needs background features.
SAM 2.1 Hiera Tiny (zone outline)	Zone outlines and depth hints	✓ **Allowed**	No pinned ONNX file exists, so the install downloads the commit-pinned checkpoint, checks its SHA-256, and exports the encoder and decoder on the site with the trainer's torch. Small, base+ and large are not offered: too slow on a processor.
Depth Anything V2 Small (depth hint)	Zone outlines and depth hints	✓ **Allowed**	One commit-pinned ONNX file from the onnx-community mirror, checked by SHA-256. Relative depth only, never metres.

Every model, by job

664 models for 63 jobs, in eight groups. Choose a group, or search. In the app, the same list is on **Models → All models**, where you can also shortlist one model per job.

People and PPE

Person detection 31 models · 4 load today

What it does: Finds every person in the picture, so each one can be followed and checked.

Model	Commercial use	Loads today	Good to know
RF-DETR (N / S / M / L detection, plus Seg N to 2XL) ☆ Start here	✓ **Allowed**	✓ **Yes**	The person detector in the app's catalogue (N, S, M, L; sha256 checked). Objects365 provenance is unconfirmed.
D-FINE (N / S / M / L / X) Licence-safe alternative	✓ **Allowed**	✓ **Yes**	In the app's catalogue (S, M). Take the COCO-only checkpoints; the README says Objects365 ones are not cleared.
RT-DETR / RT-DETRv2 (Baidu / PekingU)	✓ **Allowed**	✓ **Yes**	Use the COCO-only checkpoints; first-class Transformers support.
YOLOX (Megvii)	✓ **Allowed**	Not yet Needs a loader first.	Cheap CPU fallback; its ONNX head has an objectness score, not the ultralytics layout.
Ultralytics YOLO11 / YOLO26 (and YOLOv8, v5)	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Works today; AGPL applies even to models trained from scratch; no public price.

Model	Commercial use	Loads today	Good to know
RT-DETRv4	⚠ **With conditions** Clear the DINOv3 teacher licence first.	Not yet Needs a loader first.	Whether a student of a DINOv3-licensed teacher is a derivative work is unsettled.
RT-DETRv3 (Baidu, PaddlePaddle) Unconfirmed	⚠ **With conditions** Weights licence not stated.	Not yet Needs a loader first.	Gives nothing the PyTorch RT-DETR family does not; Paddle dependency.
LW-DETR	✓ **Allowed**	Not yet Needs a loader first.	Ancestor of RF-DETR; take RF-DETR instead.
DEIM (original, D-FINE lineage)	✓ **Allowed**	Not yet Needs a loader first.	The same lab's DEIMv2 is non-commercial; check the exact repo and tag.
RTMDet (MMDetection) Unconfirmed	✓ **Allowed**	Not yet Needs a loader first.	Weights licence is inferred from the repo; needs MMDeploy.
PaddleDetection PP-YOLOE+ / PicoDet Unconfirmed	⚠ **With conditions** Weights licence not stated.	Not yet Needs a loader first.	PicoDet suits very small edge boxes; Paddle toolchain.
YOLOv9 MIT re-implementation (MultimediaTechLab/YOLO)	✓ **Allowed**	Not yet Needs a loader first.	The one YOLO-style CNN under MIT; not ultralytics-compatible; check that it shares no GPL code.
LibreYOLO (wrapper library) Unconfirmed	⚠ **With conditions** Check each wrapped weight's licence.	Not yet Needs a loader first.	One API over several detectors; its open-vocabulary DEIM weights are CC-BY-NC-4.0.
DAMO-YOLO (Alibaba)	⚠ **With conditions** Weights licence not stated.	Not yet Needs a loader first.	The OSS weights URL tried returned 404; last release April 2023.
NVIDIA TAO PeopleNet / PeopleNet-Transformer Unconfirmed	⚠ **With conditions** NVIDIA licence text not read.	Not yet Needs a loader first.	Tuned for overhead and CCTV views; NVIDIA GPUs only.
torchvision Faster R-CNN / RetinaNet / SSD Unconfirmed	⚠ **With conditions** Weights licence not stated.	Not yet Needs a loader first.	A stable baseline to compare against; older and slower than DETR and YOLO.
Intel Open Model Zoo person-detection-0200 to 0203 Unconfirmed	⚠ **With conditions** Weights licence inferred, data undisclosed.	Not yet Needs a loader first.	Tiny CPU-only detector for Intel hardware; the zoo is archived.
DETR and Deformable / Conditional DETR (original)	✓ **Allowed**	Not yet Needs a loader first.	Reference baseline only; slow to converge.
Co-DETR, DINO and detrex models (large DETRs) Unconfirmed	⚠ **With conditions** Weights licences not confirmed.	Not yet Needs a loader first.	Offline teacher or auto-labeller only; far too heavy for live streams.
MediaPipe object detector (EfficientDet-Lite, SSD MobileNet) Unconfirmed	⚠ **With conditions** Weights terms not read.	Not yet Needs a loader first.	For phones and tiny CPUs; low accuracy.
dfine-person-s	⚠ **With conditions** 0.633.	Not yet 238.0.	
rf-detr-person-s	⚠ **With conditions** 0.863.	Not yet 300.3.	
rf-detr-person-n	⚠ **With conditions** 0.856.	Not yet 214.2.	
D-FINE-S (yardstick)	⚠ **With conditions** 0.633.	Not yet 201.6.	

Model	Commercial use	Loads today	Good to know
D-FINE-N (`dfine_n_coco`)	⚠ **With conditions** 0.317.	Not yet 84.1.	
RT-DETRv2-R18 (`rtdetr_v2_r18vd`)	⚠ **With conditions** 0.659.	Not yet 321.3.	

Avoid 5

Model	Commercial use	Loads today	Good to know
RF-DETR Plus (Atto / Femto / Pico / XL / 2XL) Avoid	Not rated	Not rated	PML 1.0, With conditions, and only worth it with a Roboflow plan; the PML text was not found.
DEIMv2 and EdgeCrafter Avoid	Not rated	Not rated	Non-commercial licence; was Apache-2.0 until 2026-08-24.
YOLOv10, YOLOv12, YOLOv13 Avoid Unconfirmed	Not rated	Not rated	AGPL-3.0 with no enterprise route.
YOLOv9 (WongKinYiu), YOLOv7, YOLOv6, Gold-YOLO, MMYOLO Avoid	Not rated	Not rated	GPL-3.0 copyleft.
YOLO-NAS (Deci / NVIDIA) Avoid	Not rated	Not rated	The weights licence bars commercial use.

PPE detection 19 models · 9 load today

What it does: Checks each person for helmet, vest, gloves, goggles, mask and boots.

Model	Commercial use	Loads today	Good to know
RF-DETR N/S/M/L (Apache sizes) fine-tuned on PPE data Start here	⚠ **With conditions** Pretraining data terms not stated.	✓ **Yes**	Licence-clean base to fine-tune on site crops plus CC BY / CC0 sets; no ready PPE weights exist. Training page: Detector, RF-DETR (learning rate 1e-4); the result is an ONNX model.
D-FINE / DEIM / RT-DETRv2 (Apache) as fine-tune bases Licence-safe alternative	✓ **Allowed**	Not yet Needs a loader first.	Second route to a clean PPE model; use COCO-only checkpoints.
killuminati1/construction-ppe-yolov8	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Richest taxonomy among ready weights (19 classes); ear protection and boots are weak.
Ultralytics YOLO11 / YOLO26 trained in-house on PPE data	⚠ **With conditions** Enterprise License.	✓ **Yes**	Fastest way to a working detector today; not for a closed sale without the Enterprise licence.
Hexmon/vyra-yolo-ppe-detection	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	14 classes with "No-" forms; no metrics on the card.
thalostech2025/dfine-small-construction-ppe-v1	⚠ **With conditions** Fine-tune data source unknown.	✓ **Yes**	Only permissive-tagged ready PPE detector with five gear classes; weak outside person, hard hat and vest. The author's own scores, out of 1, are vest 0.48 and hard hat 0.39: check it on your cameras.
keremberke/yolov8{n,s,m}-protective-equipment-detection	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Already in the catalogue; low accuracy, near-silent on site frames.
keremberke/yolov8{n,s,m}-hard-hat-detection	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Helmets only, already in the catalogue; reads caps as helmets on wide CCTV views.
Hansung-Cho/yolov8-ppe-detection	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Already in the catalogue; the MIT tag is wrong for an ultralytics model.

Model	Commercial use	Loads today	Good to know
mancalazure/construction-ppe-safety	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Ready ONNX in the stock YOLOv8 layout; demo-grade and unmeasured.
KuBaN1900/rtdetr-v2-r50vd-hardhat	⚠ **With conditions** Training set has no licence tag.	Not yet Needs a loader first.	Helmet versus bare head only; two classes.
Qualcomm PPE-Detection (gear_guard_net)	⊘ **Allowed**	Not yet Needs a loader first.	Tiny detector for Qualcomm-class edge boxes; the source reportedly describes a person/face detector.
PP-YOLOE+ (PaddleDetection) Unconfirmed	⚠ **With conditions** Weights terms not checked.	Not yet Needs a loader first.	Apache code route if the Paddle toolchain is acceptable; needs its own loader.

Avoid 6

Model	Commercial use	Loads today	Good to know
SH17 YOLOv8/ v9 / v10 weights Avoid	Not rated	Not rated	Trained on CC BY-NC-SA 4.0 data and published for evaluation only.
leeyunjai/yolo11-ppe Avoid	Not rated	Not rated	No licence tag, no data or metrics, card copied from another project.
YOLO-NAS PPE fine-tunes (SuperGradients) Avoid	Not rated	Not rated	The Deci weights licence forbids distribution and any fine-tune inherits it.
Community uploads with a PPE name but no card or files (for example VirnectUnityTeam/rfdetr-ppe-detection, gtensorapp/worksite-ppe-*) Avoid	Not rated	Not rated	An apache-2.0 tag on an empty repo tells you nothing.
YOLOv9 (WongKinYiu) and YOLOv7 Avoid Unconfirmed	Not rated	Not rated	GPL-3.0 by the GitHub record, no easier than ultralytics AGPL.
YOLOv10 (THU-MIG) Avoid Unconfirmed	Not rated	Not rated	AGPL-3.0 with no enterprise route.

PPE datasets 13 models

What it does: Labelled pictures for teaching a PPE model your own site. Data, not a model.

Model	Commercial use	Loads today	Good to know
Hard Hat Workers (andrewmvd, Kaggle) and Voxel51 mirror ☆ Start here	⊘ **Allowed**	Data, not a model	Cleanest licence for helmet and bare-head examples; close-up construction images.
Open Images V7 (Helmet, Goggles, Glove, Boot, Hat) Licence-safe alternative Unconfirmed	⊘ **Allowed**	Data, not a model	Hard negatives and helmets from varied scenes; Google does not vouch for each image.
Roboflow "Hard Hats" via keremberke/hard-hat-detection	⊘ **Allowed**	Data, not a model	Largest permissive hard-hat set; the licence is the Roboflow project's own claim.
Roboflow "ppes-kaxsi" via keremberke/protective-equipment-detection	⊘ **Allowed**	Data, not a model	Gear plus explicit negatives; labels are noisy.
Roboflow "PPE Combined Model" (roboflow-universe-projects) Unconfirmed	⚠ **With conditions** Licence only confirmed second-hand.	Data, not a model	Broad class set with negatives; check the Universe page before use.

Model	Commercial use	Loads today	Good to know
Roboflow "construction-site-safety" via keremberke/construction-safety-object-detection	Allowed	Data, not a model	Tiny set that mixes PPE and site vehicles; useful for tests.
Construction-PPE (Ultralytics)	With conditions Ultralytics Enterprise licence.	Data, not a model	Good negative-class design to imitate; Ultralytics publishes no weights for it.
SHWD (Safety Helmet Wearing Dataset) Unconfirmed	With conditions Picture copyright not stated.	Data, not a model	Many bare-head negatives from CCTV-like views.
CPPE-5 (medical PPE)	With conditions Licence unknown.	Data, not a model	Only medical PPE (hospital preset).
Hairnet and kitchen PPE sets (Roboflow Universe, IEEE DataPort) Unconfirmed	With conditions Licences not re-verified.	Data, not a model	Hygiene preset (hairnet, mask, gloves, apron), where nothing else exists.

Avoid 3

Model	Commercial use	Loads today	Good to know
SH17 Avoid	Not rated	Not rated	CC BY-NC-SA 4.0 plus Pexels licence, research only. Copy its class list, not its images.
CHV (ZijianWang-ZW/PPE_detection) and Pictor-PPE (ciber-lab/pictor-ppe) Avoid	Not rated	Not rated	No licence file; the GitHub record shows license null for pictor-ppe.
Hugging Face-uploaded PPE sets with a licence tag (for example VincentGOURBIN/ppe-detection, HB1204/PPE_Detection) Avoid Unconfirmed	Not rated	Not rated	The card holds only a cc-by-4.0 tag and no source.

Open-vocabulary detection 19 models · 2 load today

What it does: Finds things you name in plain words with no training, for labelling and rare items.

Model	Commercial use	Loads today	Good to know
Grounding DINO (open Swin-T and Swin-B) Start here	Allowed	Not yet Needs a loader first.	Standard permissive text-prompt labeller; slow, use offline.
OmDet-Turbo Start here	Allowed	Not yet Needs a loader first.	Fast enough for near frame rate; quality on PPE untested.
MM-Grounding-DINO (tiny / base / large) Licence-safe alternative	Allowed	Not yet Needs a loader first.	Open re-implementation with Transformers support; PPE quality untested.
OWLv2 (and OWL-ViT v1) Licence-safe alternative	Allowed	Not yet Needs a loader first.	Text or image-example queries; heavy ViT, not real-time.
LLMDet	Allowed	Not yet Needs a loader first.	Newer labeller; training data licence not confirmed.
Florence-2 (base / large)	Allowed	Not yet Needs its own integration.	Grounding, captioning and OCR from one small model; boxes less precise than a detector.

Model	Commercial use	Loads today	Good to know
SAM 3 / 3.1 (text and exemplar concept prompts)	⚠ **With conditions** Meta may change the licence; ITAR and weapons-use bar.	Not yet Needs its own integration.	Best-quality labeller for masks and boxes; keep offline and out of weapon-detection workflows.
SAM 2 / 2.1	✓ **Allowed**	Not yet Needs its own integration.	Turns a box into a mask or tracks a mask through video; no text prompts.
Grounded-SAM / Grounded-SAM-2 Unconfirmed	✓ **Allowed**	Not yet Needs its own integration.	Ready text to box to mask to track pipeline; glue-repo licence not re-read.
Roboflow Autodistill Unconfirmed	✓ **Allowed**	Not yet Needs its own integration.	The pattern: label with a foundation model, ship only the student; its YOLO-World and YOLOE plug-ins are copyleft.
MobileSAM and EfficientSAM Unconfirmed	✓ **Allowed**	✓ **Yes**	Cheap box-to-mask on CPU; upstream LICENSE not fetched; EfficientSAM not checked.
YOLOE (THU-MIG YOLOE and Ultralytics YOLOE-11 / YOLOE-26)	⊘ **Not allowed**	✓ **Yes**	Fastest text-prompt detector, for labelling only; bake prompts offline and do not ship it.
Detic (Meta) Unconfirmed	⚠ **With conditions** Weights licence not confirmed.	Not yet Needs a loader first.	Historic baseline, superseded by Grounding DINO and OWLv2.
GLIP (Microsoft) Unconfirmed	⚠ **With conditions** Weights licence not checked.	Not yet Needs a loader first.	Historic baseline, superseded by Grounding DINO.

Avoid 5

Model	Commercial use	Loads today	Good to know
YOLO-World (Tencent; Ultralytics port yolov8-world) Avoid	Not rated	Not rated	GPL-3.0 (Tencent) and AGPL-3.0 (port); a commercial licence on request is from the survey only.
Grounding DINO 1.5 / 1.6, DINO-X (IDEA Research API) Avoid Unconfirmed	Not rated	Not rated	Cloud API only, no weights; it would also send site frames off the site.
Rex-Omni and T-Rex / T-Rex2 (IDEA Research) Avoid	Not rated	Not rated	IDEA License 1.0, non-commercial research.
OV-DEIM / LibreOVDEIM Avoid	Not rated	Not rated	CC-BY-NC-4.0 weights.
FastSAM Avoid	Not rated	Not rated	AGPL-3.0, looks like a cheap segmenter but is copyleft.

Pose 16 models · 2 load today

What it does: Finds the body's joints, the basis for falls, posture and reach.

Model	Commercial use	Loads today	Good to know
RTMPose (MMPose) via rtmlib ☆ Start here	⚠ **With conditions** Training-set terms not read.	✓ **Yes**	Per-person keypoints on crops; the paired RTMDet detector is trained on Objects365.
RF-DETR keypoints (preview) Licence-safe alternative	✓ **Allowed**	Not yet Needs a loader first.	Detector and pose in one pass once the DETR decoder exists; it is a preview.
RTMO (one-stage multi-person)	⚠ **With conditions** Body7 training-set terms not read.	Not yet Needs a loader first.	For crowded scenes; heavier than RTMPose crops for few people.

Model	Commercial use	Loads today	Good to know
ViTPose / ViTPose+	⚠ **With conditions** AIC terms unread.	Not yet Needs a loader first.	Highest accuracy when speed is secondary; per-person crops.
LibreYOLO pose checkpoints (YOLO9-s-pose, RF-DETR n/s/m/l/x-pose)	⚠ **With conditions** YOLO9 variants derive from YOLOv9, believed GPL-3.0.	Not yet Needs a loader first.	The cards call them "EXTREMELY experimental".
DWPose / RTMW whole-body	⚠ **With conditions** RTMW cocktail14 data terms not read.	Not yet Needs a loader first.	Adds hand and face keypoints.
MediaPipe Pose / BlazePose Landmarker Unconfirmed	⚠ **With conditions** Weights licence and data unconfirmed.	Not yet Needs a loader first.	One close subject, 33 landmarks; single-person.
Skeleton action models for falls: PoseC3D, ST-GCN (MMAction2, PYSKL)	⚠ **With conditions** Stock weights carry NTU terms.	Not yet Needs a loader first.	Learn falls from pose sequences if a site has labelled falls; most teams use a rule instead.
HRNet, Intel OMZ human-pose-estimation, PP-TinyPose, Lightweight OpenPose	⚠ **With conditions** No weights licence stated.	Not yet Needs a loader first.	References and edge targets; older and less accurate than RTMPose.
YOLO-pose (Ultralytics YOLO11 / YOLO26 pose)	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Only pose model that loads today, for development.
MoveNet (Google; Qualcomm AI Hub port) Unconfirmed	⚠ **With conditions** Read through a third-party port; data cannot be audited.	Not yet Needs a loader first.	Single-person, very fast keypoints on edge devices.

Avoid 5

Model	Commercial use	Loads today	Good to know
EdgeCrafter ECPose (and the LibreEC-pose mirrors) 👎 Avoid	Not rated	Not rated	EdgeCrafter License, non-commercial; the mirror cards' Apache-2.0 tag contradicts upstream.
OpenPose and AlphaPose 👎 Avoid	Not rated	Not rated	Academic non-profit licence for OpenPose; AlphaPose non-commercial per the survey only.
Sapiens (Meta) 👎 Avoid	Not rated	Not rated	CC BY-NC 4.0.
YOLO-NAS-pose (Deci) and ED-Pose (IDEA) 👎 Avoid	Not rated	Not rated	The YOLO-NAS licence bars commercial use; ED-Pose is under the non-commercial IDEA License.
openpifpaf 👎 Avoid	Not rated	Not rated	AGPL-3.0 with a commercial licence on request from EPFL-TTO, plus a "may not be used to harm any person deliberately" line.

Crowd density and counting 10 models · Not in the platform yet

What it does: Counts and density maps where a detector merges bodies. **Needs:** Elevated view, validation against manual counts. **Main risk:** Large published errors; research-only datasets.

Model	Commercial use	Good to know
RF-DETR person detection with tiling ☆ Start here	✓ **Allowed**	Counts up to a few dozen visible people with exact boxes. Yes applies to N to L only. Would run on: Consumer GPU
CLIP-EBC	⚠ **With conditions** Weights built on research datasets; fine-tune on site frames.	Best open density model. NWPU-Crowd MAE 61.2 (ViT-B/16), 75.8 (ResNet-50). Would run on: Consumer GPU

Model	Commercial use	Good to know
DM-Count	⚠ **With conditions** Weights terms not stated.	Simple baseline that is easy to retrain on site data. Would run on: Consumer GPU
PET (point-query transformer)	⚠ **With conditions** Licence template unfilled; ask the authors.	Gives crowd points, not just a count. Would run on: Consumer GPU
APGCC	⚠ **With conditions** Weights and dataset terms unread.	Point-based localisation, stable under noise. Would run on: Consumer GPU
STEERER	⚠ **With conditions** Licence text still names another owner.	Handles large size change within one view. Would run on: Consumer GPU
CountGD, DAVE, LOCA, GeCo2 and CounTR (class-agnostic counters) Unconfirmed	⚠ **With conditions** Weights terms differ; CountGD builds on Grounding DINO.	Counts any item from examples or text: boxes on a pallet, bags, tyres. Not for dense crowds. Would run on: Consumer GPU to large GPU
Open crowd datasets (ShanghaiTech, UCF-QNRF, NWPU-Crowd, JHU-Crowd++) Unconfirmed	⚠ **With conditions** Terms unread; research-only is usual.	Benchmarking only. Label site frames instead. Would run on: n/a

Avoid 2

Model	Commercial use	Good to know
P2PNet (Tencent) Avoid	Not rated	Academic research only.
CSRNet (leeyeehoo port) Avoid	Not rated	No licence file, so no right to use.

Hands and gestures 5 models · Not in the platform yet

What it does: Warns when a hand enters a hazard zone; spots a help signal. **Needs:** Close cameras, drawn zones, machine-state signal. **Main risk:** Not a safeguard; gloves and small hands.

Model	Commercial use	Good to know
RF-DETR or D-FINE fine-tuned on "hand" and the hazard, plus zone geometry Start here Unconfirmed	✓ **Allowed**	Simplest and most robust with gloves. Needs a few hundred labelled gloved hands at real distance. Advisory only. Would run on: Consumer GPU
MediaPipe Hand Landmarker and Gesture Recognizer Unconfirmed	⚠ **With conditions** Weights licence unconfirmed.	21 landmarks and a built-in open-palm or fist classifier. Fails on small hands in wide views. Would run on: CPU
RTMPose hand and RTMW / DWPose whole-body Unconfirmed	⚠ **With conditions** Hand and whole-body training-set terms not read.	Hand points inside a whole-body pass. Only as good as the pixels available. Would run on: Consumer GPU
HaGRID gesture dataset (and models trained on it)	⚠ **With conditions** ShareAlike may bind derived models; read the PDF closely.	Starting images for a gesture classifier. Indoor, near-camera photos. Would run on: n/a

Avoid 1

Model	Commercial use	Good to know
HaMeR and other MANO-based hand mesh models Avoid Unconfirmed	Not rated	Need the MANO model, which is for non-commercial research.

Ergonomic risk from pose 5 models · Not in the platform yet

What it does: Station or task trend of REBA or RULA scores. **Needs:** Side camera, tall enough person, ergonomist thresholds. **Main risk:** 2D angle error; health inference on named workers.

Model	Commercial use	Good to know
RTMPose / RTMW keypoints plus REBA or RULA angle rules ☆ Start here · Unconfirmed	⚠ With conditions Pose training-set terms not read.	Joint angles from a side view. An ergonomist must review the REBA table. Would run on: Consumer GPU, or CPU for few people
RF-DETR keypoints (preview) plus angle rules · Unconfirmed	⚠ With conditions Keypoint weights licence not read from a primary source.	Permissive person detection and 17 keypoints in one pass. Preview; fewer foot and hand points. Would run on: Consumer GPU
MotionBERT (2D to 3D lifting)	⚠ With conditions Weights terms unstated and data is academic; treat as no until authors answer.	3D joints for better trunk and neck angles. Needs 243-frame clips; mocap data does not look like a plant. Would run on: Consumer GPU
MediaPipe Pose Landmarker (BlazePose) · Unconfirmed	⚠ With conditions Weights licence not found.	One close subject with 33 landmarks and a 3D estimate. Single person per crop. Would run on: CPU

Avoid 1

Model	Commercial use	Good to know
MeTRAbs, VideoPose3D and SMPL-based body models (4D-Humans, WHAM) 👎 Avoid	Not rated	Non-commercial weights and academic training data; the best-known 3D pose models cannot be sold here.

Action and behaviour recognition from video 9 models · Not in the platform yet

What it does: Spots a person down, climbing, a fight or entry into a danger zone. **Needs:** Pose stage, GPU, site-labelled clips, review loop. **Main risk:** False alarms; not a safety function; no ranking of workers.

Model	Commercial use	Good to know
MMAction2 skeleton models (PoseC3D, ST-GCN, CTR-GCN) trained on site clips ☆ Start here	✓ Allowed	Person down, climbing, running, crouching over 1 to 3 seconds. Viewpoint-specific, which suits fixed CCTV. Would run on: CPU for the classifier; consumer GPU for the pose stage
V-JEPA 2 / 2.1 (frozen backbone plus a small probe)	⚠ With conditions Pretraining data mix has research-oriented terms.	Strongest open video features; a small head on 50 to 200 clips per action. Benchmarks are gym-style, not CCTV. Would run on: Consumer GPU (ViT-L); large GPU (ViT-g)
X3D and SlowFast (PyTorchVideo) · Unconfirmed	⚠ With conditions Weights licence silent; Kinetics provenance.	Small classifiers to fine-tune when pose is not enough (carrying, ladder use). Older than V-JEPA 2. Would run on: CPU or small GPU (X3D-XS, S); consumer GPU (SlowFast)
ViViT (google/vivit-b-16x2-kinetics400)	⚠ With conditions Kinetics provenance.	Plainly MIT-tagged video transformer; slow, no reason to prefer it over V-JEPA 2. Would run on: Consumer GPU
X-CLIP (microsoft/xclip-*) zero-shot	⚠ With conditions Kinetics training data.	Name an action in English and score a clip with no training; for prototypes. No Sinhala or Tamil prompts. Would run on: Consumer GPU
Qwen3-VL (2B, 8B) asked about a short clip	✓ Allowed	Second opinion on a clip an alert already flagged, and alert text. Not alarm-grade; assume English. Would run on: Consumer GPU
SmolVLM2 (256M, 500M, 2.2B)	✓ Allowed	Smallest video-aware model; weak on safety detail. Scraped video terms unread; record the risk decision. Would run on: CPU or small GPU

Avoid 2

Model	Commercial use	Good to know
VideoMAE, VideoMAE V2 and TimeSformer (stock weights) Avoid	Not rated	CC BY-NC 4.0 weights. VideoMAE V2 code is MIT but its weights are not.
Community violence, fight and fall detectors (for example HappyGook/videomae-violence-detector) Avoid	Not rated	Card tag MIT, but fine-tuned from a CC BY-NC base on UCF-Crime.

Tracking and identity

Tracking 18 models

What it does: Follows each person from picture to picture, so one person is one track and one event.

Model	Commercial use	Loads today	Good to know
trackers (Roboflow): SORT, ByteTrack, OC-SORT, BoT-SORT, C-BIoU, McByte Start here	Allowed	Not yet Needs its own integration.	Default tracker library for a closed-source product; no ReID branches; McByte needs SAM and Cutie.
norfair (Tryolabs) Licence-safe alternative	Allowed	Not yet Needs its own integration.	Simple distance-based tracking; last release 2025-04-30.
BoT-SORT (NirAharon)	Allowed	Not yet Needs its own integration.	Best for cameras that move or shake; use the trackers port.
supervision ByteTrack (Roboflow)	Allowed	Not yet Needs its own integration.	Ready ByteTrack plus zones, lines and counters; overlaps with trackers.
ByteTrack (ifzhang original)	Allowed	Not yet Needs its own integration.	Reference code; use the maintained trackers package in production.
OC-SORT (noahcao)	Allowed	Not yet Needs its own integration.	Better recovery after occlusion; also inside trackers.
Deep OC-SORT	Allowed	Not yet Needs its own integration.	Adds appearance for crowds; research code.
BoostTrack	Allowed	Not yet Needs its own integration.	For experiments with a newer association scheme; small project.
TrackTrack	Allowed	Not yet Needs its own integration.	Benchmark comparison on your own footage; scores not checked.
deep_sort_realtime	Allowed	Not yet Needs its own integration.	DeepSORT-style tracker with a permissive licence; an old method.
CAMELTrack	With conditions Weights licence and data terms not checked.	Not yet Needs its own integration.	Learned fusion of motion and appearance; needs a ReID model.
MOTRv2 / MOTIP (end-to-end transformer trackers)	With conditions Weights licence unknown.	Not yet Needs its own integration.	Research only; heavy, one model per dataset.
SAM 2 video object tracking (SAMURAI builds on it)	Allowed	Not yet Needs its own integration.	Tracks one or a few chosen objects, not every person; far heavier than box trackers.

Model	Commercial use	Loads today	Good to know
DeepStream trackers (NvDCF, NvSORT, NvDeepSORT, MV3DT) Unconfirmed	⚠ **With conditions** NVIDIA EULA terms not read.	Not yet Needs its own integration.	Only if the product moves to DeepStream on NVIDIA GPUs.

Avoid 4

Model	Commercial use	Loads today	Good to know
SORT (abewley, original) and DeepSORT (nwojke) Avoid	Not rated	Not rated	GPL-3.0; the DeepSORT embedder was trained on MARS. Use SORT from the trackers package.
StrongSORT (dyhBUPT) Avoid	Not rated	Not rated	GPL-3.0, bundled in many YOLO-tracking repos.
BoxMOT Avoid	Not rated	Not rated	AGPL-3.0; fine for prototyping, never shipped.
Ultralytics built-in trackers (ByteTrack, BoT-SORT) Avoid	Not rated	Not rated	AGPL-3.0, covered only by the Ultralytics commercial licence.

Re-identification 11 models

What it does: Recognises the same person again on another camera or later the same day.

Model	Commercial use	Loads today	Good to know
OSNet architecture (torchreid), trained on your own tracklets Start here	✓ **Allowed**	Not yet Needs its own integration.	Start from random or ImageNet initialisation, not the stock Market-1501 checkpoints, or the row becomes With conditions.
DINOv2 embeddings as a re-ID baseline	✓ **Allowed**	Not yet Needs its own integration.	Zero-effort appearance embedding to test against OSNet; not trained for clothing-based identity.
OSNet / OSNet-AIN stock zoo weights	⚠ **With conditions** Trained on research-only datasets.	Not yet Needs its own integration.	For quick experiments only; weights are on Google Drive, not a stable URL.
FastReID (JD AI)	⚠ **With conditions** Stock weights carry research-only dataset terms.	Not yet Needs its own integration.	Complete training pipeline for your own weights; last push about 2024-07.
TransReID (ViT re-ID)	⚠ **With conditions** Stock weights carry dataset terms.	Not yet Needs its own integration.	Higher accuracy than OSNet when GPU time is not tight; heavy for many cameras.
CLIP-ReID and SOLIDER-REID	⚠ **With conditions** Weights licences and LUPerson terms not read.	Not yet Needs its own integration.	Benchmark leaders if you train your own on clean data.
Intel Open Model Zoo person-reidentification-retail-0277 / 0286 / 0287 / 0288	⚠ **With conditions** Training data undisclosed.	Not yet Needs its own integration.	Tiny CPU embedding for a quick demo; accuracy below OSNet.
NVIDIA TAO ReIdentificationNet Unconfirmed	⚠ **With conditions** NGC page could not be read.	Not yet Needs its own integration.	Only on the NVIDIA / DeepStream path.
DINOv3 and SigLIP 2 as embedding baselines	⚠ **With conditions** DINOv3 needs approval and has conditions.	Not yet Needs its own integration.	Neither is trained for person identity; SigLIP 2 is the clean one.

Avoid 2

Model	Commercial use	Loads today	Good to know
BPBreID (body-part re-ID) Avoid	Not rated	Not rated	Hippocratic License 3.0, an ethical-source licence with use restrictions a closed product cannot easily meet.
Community Hugging Face re-ID checkpoints (for example PersonViT MSMT17, OSNet MSMT17 copies) Avoid	Not rated	Not rated	The uploader's tag cannot override the MSMT17 terms; the same weights are tagged MIT by one uploader and GPL-3.0 by another.

Face detection 15 models · 1 load today

What it does: Finds faces in the picture, the first step before any face is matched.

Model	Commercial use	Loads today	Good to know
YuNet (OpenCV Zoo) Start here	⚠ **With conditions** Trained on WIDER FACE.	Not yet Needs a loader first.	Small, fast default with 5 landmarks for alignment; WIDER FACE terms not re-read.
MediaPipe BlazeFace / Face Detector Unconfirmed	⚠ **With conditions** Weights licence rests on an unread model card.	Not yet Needs a loader first.	Permissive, no-WIDER option for close faces (doors, kiosks); misses small CCTV faces.
RetinaFace (biubug6 Pytorch_Retinaface; yakhyo/retinaface-pytorch)	⚠ **With conditions** Weights licence not stated; WIDER FACE.	Not yet Needs a loader first.	Accurate 5-point detection when YuNet misses; the yakhyo fork was not read.
libfacedetection	⚠ **With conditions** WIDER FACE provenance.	Not yet Needs a loader first.	Source project of YuNet, for retraining.
AMD RetinaFace int8 ONNX	⚠ **With conditions** WIDER FACE provenance.	Not yet Needs a loader first.	Ready quantised ONNX, built for AMD NPUs.
UltraFace (Linzaer) and CenterFace (Star-Clouds)	⚠ **With conditions** WIDER FACE provenance.	Not yet Needs a loader first.	Very light CPU detection, less accurate than YuNet; CenterFace licence not read.
MTCNN (ipazc/mtcnn, facenet-pytorch)	⚠ **With conditions** Weights provenance.	Not yet Needs a loader first.	Familiar and easy, but slow and dated; ipazc/mtcnn licence not read.
Intel OMZ face-detection-0200 / 0202 / 0204 / 0205 / 0206, adas-0001, retail-0004	⚠ **With conditions** Training data undisclosed.	Not yet Needs a loader first.	For OpenVINO on Intel CPUs.
dlib mmod CNN face detector	⚠ **With conditions** Source image terms not itemised.	Not yet Needs a loader first.	Explicit public-domain grant on the model files; needs dlib.
HF community YOLO face models (AdamCodd/YOLOv11n-face-detection, Bingsu/adetailer face_yolov8*, arnabdhar/YOLOv8-Face-Detection)	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Only face detector that loads today, for internal tests; an Apache tag does not remove the AGPL.

Avoid 5

Model	Commercial use	Loads today	Good to know
dlib 68-point landmark model (also used by the face_recognition package) Avoid	Not rated	Not rated	The author states it cannot be used commercially. Use the 5-point landmarks from YuNet.
SCRFD (InsightFace) Avoid	Not rated	Not rated	Non-commercial research weights.

Model	Commercial use	Loads today	Good to know
YOLO face models (YOLOv8-face by derronqi and akanametov, YOLOv5-Face by deepcam-cn) Avoid	Not rated	Not rated	GPL-3.0 repos on Ultralytics / YOLOv5 code.
LibreYOLO face-rfdetr-nano Avoid	Not rated	Not rated	A toy checkpoint trained on 100 images.
OpenCV res10 SSD face detector (Caffe) Avoid Unconfirmed	Not rated	Not rated	A forum thread says no licence is given for the weights; replaced by YuNet.

Face recognition 11 models

What it does: Matches a face to the people you have named.

Model	Commercial use	Loads today	Good to know
SFace (OpenCV Zoo, MobileFaceNet) Start here	⚠ **With conditions** Training data may include research-only sets.	Not yet Needs its own integration.	Small baseline embedder through OpenCV FaceRecognizerSF.
AuraFace-v1 (glintr100) Start here	⚠ **With conditions** Training dataset undisclosed.	Not yet Needs its own integration.	Stated commercial intent; the repo also holds InsightFace-looking files that must not be loaded.
FaceNet (davidsandberg, facenet-pytorch)	⚠ **With conditions** VGGFace2 / CASIA-WebFace terms not read.	Not yet Needs its own integration.	A familiar 512-d baseline; dated.
dlib ResNet v1 face recognition (and face_recognition package)	⚠ **With conditions** Source dataset terms (FaceScrub, VGG).	Not yet Needs its own integration.	CPU-friendly fallback; the wrappers also load the non-commercial 68-point model.
Intel OMZ face-reidentification-retail-0095	⚠ **With conditions** Training data undisclosed.	Not yet Needs its own integration.	For Intel CPU deployments; accuracy not read.

Avoid 6

Model	Commercial use	Loads today	Good to know
InsightFace packs (buffalo_l/s/m/sc, antelopev2, ArcFace R50/R100, w600k_r50) Avoid	Not rated	Not rated	Non-commercial research weights; commercial licence sold by InsightFace; the python package auto-downloads them.
EdgeFace (Idiap) Avoid	Not rated	Not rated	CC BY-NC-SA 4.0 weights.
AdaFace Avoid	Not rated	Not rated	Weights trained on research-only face sets and published with no licence.
ONNX Model Zoo ArcFace (LResNet100E-IR) and OMZ face-recognition-resnet100-arcface-onnx Avoid	Not rated	Not rated	Apache-2.0 label on weights trained on MS-Celeb-1M.
MagFace, GhostFaceNets, CurricularFace and similar research repos Avoid	Not rated	Not rated	Weights trained on MS1M data and published without a licence.
Weight bundles (UniFace weights mirror, DeepFace) Avoid	Not rated	Not rated	Mixed per file; they contain InsightFace non-commercial files, GPL-3.0 YOLOv5-Face and unlicensed weights.

Anti-spoofing 6 models

What it does: Tells a real face from a photo or a screen held up to the camera.

Model	Commercial use	Loads today	Good to know
Silent-Face-Anti-Spoofing (MiniFASNetV1SE / V2) Start here	Allowed	Not yet Needs its own integration.	Print and screen-replay checks; RGB only and needs a near-frontal face.
MiniFASNet ONNX and LiteRT ports (UniFace mirror, litert-community)	Allowed	Not yet Needs its own integration.	Runs without PyTorch; the ONNX file sits in the mixed-terms UniFace bundle, so take only that file.

Avoid 4

Model	Commercial use	Loads today	Good to know
CDCN / CDCN++ (ZitongYu) Avoid	Not rated	Not rated	MIT text overridden by "just for research purpose, and commercial use should be contacted with authors first".
light-weight-face-anti-spoofing (kprokofi) and hairymax/Face-AntiSpoofing Avoid	Not rated	Not rated	Weights trained on CelebA-Spoof; hairymax has no licence file.
Hugging Face community liveness models (nguyenkhoa ViT/DINOv2, biometric-ai-lab/Antispoofing) Avoid	Not rated	Not rated	Trained on an unknown dataset, so the numbers and the licence mean little.
Spoof datasets Avoid	Not rated	Not rated	CelebA-Spoof, OULU-NPU, CASIA-FASD (unconfirmed): terms not read; the survey records them as research and non-commercial.

Safety checks

Phone use 6 models · 1 load today

What it does: Notices a phone in a person's hand where phones are not allowed.

Model	Commercial use	Loads today	Good to know
Ultralytics YOLO11 / YOLO26 COCO weights, class 67 "cell phone" Start here	With conditions Enterprise License or AGPL terms.	Yes	In-house spike: no false alarms on 72 negatives, about half the phones found at about 220 px person height.
YOLOX (Megvii) COCO weights Licence-safe alternative	Allowed	Not yet Needs a loader first.	Clean-tier detector with ready ONNX files; the ONNX provider cannot decode its output as is.
RF-DETR N/S/M/L COCO weights (class "cell phone")	With conditions Pretraining data not stated.	Not yet Needs a loader first.	Licence-clean phone detector from the same model that finds people; fine-tune on Open Images phone boxes.
D-FINE / DEIM / RT-DETRv2 COCO weights	Allowed	Not yet Needs a loader first.	Alternative to RF-DETR with a COCO-only checkpoint; same decoder gap.
Open Images V7 "Mobile phone" boxes (fine-tuning data) Unconfirmed	Allowed	Data, not a model	Clean extra phone examples to raise recall; mostly hand-held close-ups.

Avoid 1

Model	Commercial use	Loads today	Good to know
sreenathsree1578/mobile-phone-usage-detector Avoid	Not rated	Not rated	A whole-image classifier with unknown data, no licence and no phone location.

Smoking 5 models · 3 load today

What it does: Notices smoking where it is not allowed.

Model	Commercial use	Loads today	Good to know
Beehzod/smoking-detection-finetuned Start here	With conditions Enterprise License or AGPL terms.	Yes	In the model library as "Cigarette – YOLO11m (experimental)", pinned to commit d81c4396. The smoking check runs it on person crops at conf 0.5, and a phone or food box overlapping the hit vetoes it. Finds about half of site-scale cigarettes.
Pose rule (hand-to-mouth) on RTMPose keypoints plus a small own-trained cigarette/vape detector Licence-safe alternative	With conditions Body7 training-set terms not read.	Yes	Licence-clean way to build smoking; needs your own labelled clips.
Beehzod/smoke_cigarette-detection2-yolo11m and best_smoke_cigarette-detection	With conditions Enterprise License or AGPL terms.	Yes	Second cigarette detector to compare; a 0.994 mAP50 on near-duplicate frames says little.
mbehbooei/vit-base-patch16-224-in21k-finetuned-smoking	With conditions Unknown training data.	Not yet Needs its own integration.	Whole-crop second opinion; no localisation, a lead only.
Roboflow Universe cigarette sets (xu-wei/cigarette-all and others) Unconfirmed	With conditions Licence per set unknown.	Data, not a model	Only source of cigarette boxes the community models use; read each page before training.

Fire and smoke 13 models · 3 load today

What it does: Notices flames or smoke in the picture.

Model	Commercial use	Loads today	Good to know
rabahdev/fire-smoke-yolov8n Start here	With conditions Enterprise License or AGPL terms.	Yes	Works today; use conf 0.5 or more for fire.
RF-DETR or D-FINE fine-tuned on D-Fire + Pyro-SDIS + site negatives (recipe)	With conditions RF-DETR pretraining data unstated; D-Fire authors hold no copyright over the images.	Not yet Needs a loader first.	The licence-clean detector the product can own; the D-FINE COCO-only base avoids the RF-DETR question.
pyronear/rexnet1_0x and other Pyronear classifiers (mobilenet_v3, resnet18/34)	Allowed	Not yet Needs its own integration.	Cheap whole-frame wildfire yes/no gate for an expensive detector; no boxes, outdoor imagery.
D-Fire dataset	Allowed	Data, not a model	Main permissive fire and smoke set.
Pyro-SDIS dataset	Allowed	Data, not a model	Distant smoke from fixed cameras; outdoor wildfire views only.
xiazh0219/yolo26m-UAV-fire-smoke-detection	With conditions Enterprise License or AGPL terms.	Yes	Aerial and outdoor smoke; not for ceiling or wall cameras.
Pyronear YOLO (yolov8s, yolov11s/n, yolo11s_*)	With conditions Enterprise License or AGPL terms.	Yes	Early smoke at long range from fixed towers; indoor steam and dust will trigger it.
fireviewer/rtdetr-v2-r50-fire-smoke	With conditions No redistribution, public credit; get a written agreement first.	Not yet Needs a loader first.	Fine for a hosted service you run yourself; shipping it inside installed software is redistribution.

Model	Commercial use	Loads today	Good to know
EdBianchi/vit-fire-detection and fireviewer/dinov2-small-fire-associated-smoke-v2.5	⚠ **With conditions** Fine-tune data unclear.	Not yet Needs its own integration.	Whole-frame fire classifiers to confirm a detector alarm.
Intel/fire-and-smoke-detection (Qwen2-VL-2B, OpenVINO)	✓ **Allowed**	Not yet Needs its own integration.	Language-prompted verifier for an alarm candidate on Intel hardware; slow and gives no box.
FASDD Unconfirmed	⚠ **With conditions** Licence carried from the survey, not re-verified.	Data, not a model	Large aerial and ground fire/smoke set.

Avoid 2

Model	Commercial use	Loads today	Good to know
Simuletic CCTV smoke/fire and wildfire smoke datasets Avoid Unconfirmed	Not rated	Not rated	CC BY-NC 4.0, even though other Simuletic sets are CC BY 4.0.
Community fire/smoke YOLO fine-tunes (mfranzon, nobita111222, fireviewer yolo11m, Beehzod yolov26…) Avoid	Not rated	Not rated	All AGPL-3.0 by inheritance; a CC BY tag does not lift the framework's AGPL.

Falls 11 models · 4 load today

What it does: Notices a person who has fallen and stays down.

Model	Commercial use	Loads today	Good to know
RTMPose (rtmlib / MMPose) plus fall rules ☆ Start here	⚠ **With conditions** Body7 training-set terms not read.	✓ **Yes**	Loads now; the person-down pose rule is built on it.
ViTPose (usyd-community, Hugging Face)	⚠ **With conditions** Pose training-set terms not read.	Not yet Needs a loader first.	More accurate keypoints than RTMPose when the box is large; top-down, heavier.
Ultralytics YOLO11-pose	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Lets you try the fall rules before the Apache route exists.
Ultralytics YOLO26-pose	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	In the model library as "Pose – YOLO26 (official)" (weights from ultralytics assets v8.4.0). Feeds the Person down check, method pose; s is the default.
Simuletic CCTV Incident Dataset: Fall / Lying Down (synthetic) Unconfirmed	✓ **Allowed**	Data, not a model	Overhead CCTV lying-down poses to test the rules; small and synthetic.
MediaPipe Pose Landmarker (BlazePose) Unconfirmed	✓ **Allowed**	Not yet Needs a loader first.	Tiny single-person pose on a person crop, runs on CPU.
melihuzunoglu/human-fall-detection	⚠ **With conditions** Enterprise License or AGPL terms.	✓ **Yes**	Quick fall-by-posture detector; cannot tell lying down from a fall without a time rule.
pat2echo/fall-detection-posture-classification (dataset) Unconfirmed	✓ **Allowed**	Data, not a model	Clean clips with consent stated; a single subject in one setting.
PoseC3D / ST-GCN skeleton classifiers (MMAction2) Unconfirmed	⚠ **With conditions** Public checkpoints carry research terms.	Not yet Needs a loader first.	A learned classifier once you have your own clips; do not use the public checkpoints.

Avoid 2

Model	Commercial use	Loads today	Good to know
kamalchibrani/yolov8_fall_detection Avoid	Not rated	Not rated	No documentation, and the mit tag is unreliable on an AGPL-3.0 base.
UR Fall Detection and OmniFall datasets Avoid Unconfirmed	Not rated	Not rated	CC BY-NC-SA 4.0 and cc-by-nc-4.0, non-commercial. Retire the internal URFD test.

Forklift 6 models · 3 load today

What it does: Finds forklifts, and people too close to one.

Model	Commercial use	Loads today	Good to know
RF-DETR or D-FINE fine-tuned on forklift + person site frames (recipe) Start here	With conditions RF-DETR pretraining data unstated.	Not yet Needs a loader first.	The licence-clean route; D-FINE COCO-only is the alternative base.
mustafahacioglu/yolo11s-forklift-person-detection Start here	With conditions Enterprise License or AGPL terms.	Yes	Works today; its card adds a box-containment rule for operator presence; the map_up licence is unread.
yihong1120/Construction-Hazard-Detection	With conditions Enterprise License or AGPL terms.	Yes	Works today as an alternative forklift source: it reports forklifts as "machinery", and the catalogue ignores its "vehicle" class because it false-alarms on site.
keremberke/yolov8{n,s,m}-forklift-detection (and yolov5 versions)	With conditions Enterprise License or AGPL terms.	Yes	Tiny model for a first demo; a Roboflow-derived forklift model once boxed the whole frame.
Roboflow forklift datasets (keremberke/forklift-object-detection, forklift-dsitv) Unconfirmed	Allowed	Data, not a model	Seed images; add your own site frames.

Avoid 1

Model	Commercial use	Loads today	Good to know
akdve/ot-forklift-person-yolo11s-smoke Avoid	Not rated	Not rated	A one-epoch, two-image pipeline test.

Blocked exits and clear walkways 5 models · Not in the platform yet

What it does: Alerts when an exit, aisle or lane is blocked, with a dated record. **Needs:** Fixed camera view, zone, empty reference. **Main risk:** People walking through; light and camera shift.

Model	Commercial use	Good to know
Zone comparison against an empty reference (OpenCV differencing or DINOv2 patch features) Start here	Allowed	Is the exit zone different from its empty reference for 3 minutes with nobody passing? Needs no labels. Would run on: CPU
RF-DETR-Seg fine-tuned on "obstruction" (pallet, trolley, boxes) Unconfirmed	Allowed	Names the blocker and gives its mask for the audit record. Needs labelled images of real clutter. Would run on: Consumer GPU
SAM 2 for labelling and mask clean-up	Allowed	Speeds up labelling; not run live on every camera. Would run on: Consumer GPU

Model	Commercial use	Good to know
OneFormer (ADE20K) and Mask2Former (COCO)	⚠ **With conditions** ADE20K terms unread; Mask2Former licence unresolved.	Labels "floor" and "wall" for walkable-area masks without site labels. Would run on: Consumer GPU

Avoid 1

Model	Commercial use	Good to know
SegFormer (NVIDIA) 👎 Avoid	Not rated	Non-commercial NVIDIA licence; a popular default in tutorials.

Left-behind and removed objects 4 models · Not in the platform yet

What it does: Alerts when something appears and stays, or something that should stay is gone. **Needs:** Reference frame, zone, dwell time; CPU. **Main risk:** Light changes and moved stock cause false alarms.

Model	Commercial use	Good to know
OpenCV background subtractors (MOG2, KNN) ☆ Start here	✓ **Allowed**	A slow and a fast background: a blob only in the slow one is left behind, one missing from it is removed. Would run on: CPU
DINOv2 patch-feature comparison with a reference	✓ **Allowed**	Zone-level "same as the reference" check that tolerates lighting change. Coarse (patch size 14). Would run on: Consumer GPU or CPU (small)
RF-DETR (COCO classes such as backpack, suitcase, handbag, bottle)	✓ **Allowed**	Names the thing left behind. Yes for N to L only; a pallet or tool is not a COCO class. Would run on: Consumer GPU
ChangeFormer	⚠ **With conditions** Weights licence not stated.	Change between two pictures; trained on satellite pairs, a poor match to ground-level CCTV. Would run on: Consumer GPU

Objects left behind or removed 6 models · Not in the platform yet

What it does: Alerts when a bag, box or tool stays or disappears. **Needs:** Fixed camera, zones, dwell time, reference image. **Main risk:** Shadows and pausing people cause false alarms.

Model	Commercial use	Good to know
Detector plus tracker plus dwell-time rule (RF-DETR or D-FINE with the trackers library) ☆ Start here Unconfirmed	✓ **Allowed**	Reuses the person detector and tracker. Needs a person-proximity test so someone sitting by a bag is not an alert. Would run on: Consumer GPU
OpenCV background subtraction (MOG2, KNN) with static-foreground logic	✓ **Allowed**	Finds anything new and still in a zone, whatever it is. Cheapest first filter. Would run on: CPU
BGSLibrary (40 or more algorithms)	⚠ **With conditions** Bundles third-party algorithm code with its own terms; some patented.	Try many algorithms to find one that survives your lighting. Would run on: CPU
SAM 2 / SAM 2.1 (object masks for "what changed")	✓ **Allowed**	Cuts a clean mask of a changed region and drafts labels. Not for every frame on every camera. Would run on: Consumer GPU
DINOv2 patch-feature difference against an empty-zone reference	✓ **Allowed**	More robust to light and shadow than pixel differencing. Needs a reference frame per camera. Would run on: Consumer GPU, or CPU at low rates

Avoid 1

Model	Commercial use	Good to know
CoTracker3 (point tracking) 👎 Avoid	Not rated	CC BY-NC 4.0 weights.

Distance to a hazard from one camera (depth) 5 models · Not in the platform yet

What it does: A rough near-or-far hint without a floor map. **Needs:** Depth model; floor reference for metres. **Main risk:** Relative depth only; never a safe-distance decision.

Model	Commercial use	Good to know
Depth Anything V2 Small (relative depth) Start here	Allowed	Which of two things is nearer. Larger models are non-commercial. Training data not itemised; record the risk decision. Would run on: CPU at low frame rates, or small GPU
MiDaS and DPT (Intel), and ZoeDepth	With conditions Mixed training-data terms; KITTI is CC BY-NC-SA.	Mature baselines; older and heavier than Depth Anything V2 Small. Would run on: Consumer GPU
Depth Anything V2 Metric, Small Unconfirmed	With conditions Virtual KITTI 2 and Hypersim terms.	Metres directly, indoors. Outdoor weights likely inherit non-commercial terms. Would run on: Small GPU
Metric3D v2	With conditions Weights terms ambiguous.	Metric depth using known intrinsics; heavy. Would run on: Large GPU

Avoid 1

Model	Commercial use	Good to know
Depth Anything V2 Base, Large, Giant and Apple Depth Pro Avoid	Not rated	CC BY-NC 4.0 and research-only licences.

Vehicles and plates

Vehicles 23 models · 2 load today

What it does: Finds vehicles and says what kind, for gates, yards and loading bays.

Model	Commercial use	Loads today	Good to know
RF-DETR (Nano, Small, Medium, Large) Start here	Allowed	Not yet Needs a loader first.	Default detector once the DETR decoder exists; the weights licence rests on the README table.
D-FINE (N, S, M, L, X) Licence-safe alternative	Allowed	Not yet Needs a loader first.	Second detector to benchmark against RF-DETR; use the plain COCO files.
YOLOX (nano to X) Licence-safe alternative	Allowed	Not yet Needs a loader first.	Runs well on CPU at small sizes; needs a small decoder for its ONNX output.
DINOv2 features plus a small classifier head (train on detector crops)	Allowed	Not yet Needs its own integration.	Vehicle type and colour, including local classes, from a few hundred crops per class; do not swap in DINOv3.
RT-DETR and RT-DETRv2	Allowed	Yes	Easiest DETR to prototype with in Transformers; avoid the AGPL ultralytics port.
YOLOv9 MIT re-implementation (MultimediaTechLab/YOLO, mirrored as LibreYOLO9)	With conditions Weights licence only asserted by a mirror.	Not yet Needs a loader first.	Nearest licence-clean drop-in for ultralytics weights; not the GPL-3.0 WongKinYiu repo.
RF-DETR Atto, Femto, Pico, XL, 2XL (rfdetr_plus) Unconfirmed	With conditions PML 1.0 text not found.	Not yet Needs a loader first.	Only with a Roboflow plan; the Apache sizes cover the same need.
LW-DETR (T, S, M, L, X)	Allowed	Not yet Needs a loader first.	The architecture RF-DETR builds on; research reference only.

Model	Commercial use	Loads today	Good to know
DAMO-YOLO (Nano to L) Unconfirmed	Allowed	Not yet Needs a loader first.	No reason to prefer it over YOLOX or RF-DETR.
PP-YOLOE+ and PP-Vehicle (PaddleDetection) Unconfirmed	Allowed	Not yet Needs a loader first.	Only if the team already runs Paddle; PaddleYOLO ports (YOLOv5/6/7/8) are GPL-3.0, do not use those.
RTMDet (MMDetection) Unconfirmed	With conditions No weights licence stated.	Not yet Needs a loader first.	Mature CPU-friendly detector; last push 2024-08.
Ultralytics YOLO11 and YOLO26	With conditions Ultralytics Enterprise licence.	Yes	Only family that loads today; AGPL applies even to internal use or training from scratch.
NVIDIA TrafficCamNet Unconfirmed	With conditions Licence text unread.	Not yet Needs a loader first.	Cars, persons, road signs and two-wheelers; US training data.
NVIDIA VehicleTypeNet and VehicleMakeNet Unconfirmed	With conditions Licence text not read.	Not yet Needs its own integration.	Six US-style classes; no motorcycle, three-wheeler or minibus taxi.
DINOv3 (facebook/dinov3-*) Unconfirmed	With conditions Licence text not read.	Not yet Needs its own integration.	Stronger features than DINOv2; DINOv2 is the safe choice.
OWLv2 (google/owlv2-*)	Allowed	Not yet Needs a loader first.	Text-prompted pre-labelling of local vehicle types; heavy, not for live streams.
Grounding DINO tiny (IDEA-Research)	With conditions Training sets not checked.	Not yet Needs a loader first.	Text-prompted boxes to bootstrap labels for local classes; not a live detector.

Avoid 6

Model	Commercial use	Loads today	Good to know
YOLOv10, YOLOv12 and YOLOv13 research repos Avoid	Not rated	Not rated	AGPL-3.0 with no commercial route short of the repo owners.
YOLOv7, official YOLOv9 and YOLOv6 Avoid	Not rated	Not rated	GPL-3.0; the same authors' MultimediaTechLab/YOLO is MIT, do not mix them up.
YOLO-NAS Avoid	Not rated	Not rated	The Deci licence bars commercial use; the Apache code does not cover the weights.
DEIMv2 (and EdgeCrafter) Avoid	Not rated	Not rated	DEIMv2 License, non-commercial.
Community YOLO vehicle detectors on Hugging Face (for example Intel/vehicle-detection) Avoid	Not rated	Not rated	The card says MIT but the base is Ultralytics, so the AGPL stays.
YOLOv5 and YOLO-World (AGPL and GPL) Avoid	Not rated	Not rated	Copyleft.

Plate detection 12 models

What it does: Finds number plates in the picture.

Model	Commercial use	Loads today	Good to know
RF-DETR fine-tuned on local plate boxes ☆ Start here	✓ Allowed	Not yet Needs a loader first.	Clean route for local plate shapes and camera angles; needs labelled plates.
autolane/yolox-s-alpr 🛡 Licence-safe alternative	⚠ With conditions Dataset terms unverified.	Not yet Needs a loader first.	Permissive detector with an INT8 TFLite; US plates only in training.
open-image-models YOLOv9 plate detectors Unconfirmed	⚠ With conditions Weights licence and data not written down.	Not yet Needs a loader first.	Small, fast first detector (t-384 is 7.8 MB); ONNX has NMS baked in, not YOLO-style raw output.
autolane/rfdetr-alpr	⚠ With conditions Training data undisclosed.	Not yet Needs a loader first.	Shows an RF-DETR plate detector already exported to ONNX and INT8.
justjuu/rtdetr-v2-license-plate-detection	⚠ With conditions Dataset licence not read.	Not yet Needs a loader first.	Transformers-loadable plate detector to compare against; small, narrow training set.
HyperLPR3 Unconfirmed	⚠ With conditions Weights licence and data not read.	Not yet Needs a loader first.	Only for Chinese plates.
NVIDIA LPDNet Unconfirmed	⚠ With conditions Licence text not retrievable.	Not yet Needs a loader first.	Only in a TAO and TensorRT workflow for US plates.

Avoid 5

Model	Commercial use	Loads today	Good to know
Community ultralytics plate detectors on Hugging Face (YOLOv5, v8, v11, v26) 👎 Avoid	Not rated	Not rated	Cards vary (agpl-3.0, mit, none) but all sit on AGPL-3.0 code.
Intel license-plate-recognition (YOLOv8plus PP-OCRv4) 👎 Avoid	Not rated	Not rated	MIT card label over an ultralytics-trained detector; use the recogniser idea, not the detector.
Nomeroff-Net 👎 Avoid	Not rated	Not rated	GPL-3.0 copyleft covers any product that links it.
OpenALPR 👎 Avoid	Not rated	Not rated	AGPL-3.0, old, and the maintainer sells commercial products.
alpr-unconstrained / WPOD-NET 👎 Avoid	Not rated	Not rated	GPL-3.0 research code from 2018.

Plate reading 16 models · 1 load today

What it does: Reads the characters on a number plate.

Model	Commercial use	Loads today	Good to know
fast-plate-ocr (cct-s-v2-global, cct-xs-v2-global and older) ☆ Start here Unconfirmed	⚠ With conditions Weights licence and data not written down.	Not yet Needs its own integration.	Fast Latin-script reader that can be fine-tuned on local plates; Latin letters and digits only.
RapidOCR 🛡 Licence-safe alternative	✓ Allowed	Not yet Needs its own integration.	Runs PP-OCR models without Paddle installed; same non-plate tuning.
PaddleOCR PP-OCRv5 recognisers (default, latin, arabic, devanagari, cyrillic, ta, te, th, korean, el)	✓ Allowed	✓ Yes	Only permissive route to non-Latin scripts; not tuned for plates, so fine-tune on cropped plates.
Awiros-ANPR-OCR (Indian plates)	⚠ With conditions Upstream data licences not stated.	Not yet Needs its own integration.	Indian plates, one or two rows; self-reported accuracy.

Model	Commercial use	Loads today	Good to know
EasyOCR	Allowed	Not yet Needs its own integration.	Quick baseline with Bengali, Hindi, Marathi, Nepali, Tamil, Telugu, Urdu, Arabic and Persian; no Sinhala.
Tesseract with tessdata_best	Allowed	Not yet Needs its own integration.	Ready packs for Bengali, Hindi, Sinhala, Tamil, Nepali and Urdu; poor on small blurry plate crops.
docTR Unconfirmed	With conditions Weights licences and data never read.	Not yet Needs its own integration.	Document-style OCR; not plate-specific.
PaddleOCR-VL	Allowed	Not yet Needs its own integration.	Second-opinion read of hard plates on a GPU, not every frame.
GOT-OCR2.0	Allowed	Not yet Needs its own integration.	General OCR second opinion; mostly English and Chinese.
DeepSeek-OCR	Allowed	Not yet Needs its own integration.	Heavy GPU fallback for odd plates; too slow for live reading.
NVIDIA LPRNet Unconfirmed	With conditions Licence text not retrievable.	Not yet Needs its own integration.	Only US or Chinese plates inside a TAO workflow.
LPRNet_Pytorch (sirius-ai) Unconfirmed	With conditions Weights licence not stated.	Not yet Needs its own integration.	A small architecture to retrain on local plates; shipped weights read Chinese plates only.
TrOCR (microsoft/trocr-base-printed) and plate fine-tunes Unconfirmed	With conditions No weights licence stated.	Not yet Needs its own integration.	Appears in many plate tutorials; plate fine-tunes carry no stated licence.
hezarai crnn-fa-license-plate-recognition-v2 (Persian plates) Unconfirmed	With conditions No licence field.	Not yet Needs its own integration.	Only non-Latin (Persian) plate reader found; a pattern for an Arabic-script model.

Avoid 2

Model	Commercial use	Loads today	Good to know
Surya OCR Avoid	Not rated	Not rated	Modified AI Pubs Open RAIL-M weights, free for research, personal use and startups under 5 M funding or revenue.
HunyuanOCR Avoid	Not rated	Not rated	Tencent Hunyuan Community License does not apply in the EU, UK and South Korea.

Picture and camera

Camera health 10 models · 1 load today

What it does: Tells the operator when a camera is blurred, blocked, moved, dark or frozen. **Main risk:** False alarms at dusk, rain and glare; reference must not refresh itself.

Model	Commercial use	Loads today	Good to know
OpenCV classical checks (Laplacian variance, histograms, edge density, ORB or phase-correlation shift, SSIM) Start here	Allowed	Yes	Blur, dark, blocked, moved and frozen frames. Tune thresholds per camera and for day and night.
DINOv2 embedding drift against a reference frame Licence-safe alternative	Allowed	Not yet No.	Scene-level check that survives lighting change. Cannot tell a real scene change from tamper.

Model	Commercial use	Loads today	Good to know
HyperIQA	⚠ **With conditions** Weights licence not stated; trained on KonIQ-10k.	Not yet No.	Learned quality score as a second opinion beside the classical checks.
BRISQUE (OpenCV contrib quality module)	⚠ **With conditions** Trained on LIVE-R2; its terms not read.	Not yet No.	One no-reference sharpness number. Built for photos, not CCTV; plain Laplacian variance is usually enough.
MUSIQ (Google Research)	⚠ **With conditions** Checkpoint data terms not read.	Not yet No.	Scores images at any resolution. Porting effort is real.
LIQE	⚠ **With conditions** Weights and dataset terms unclear.	Not yet No.	Quality score plus a distortion label (blur, noise, over-exposure) that explains the flag.
ARNIQA	⚠ **With conditions** No weights licence in the README.	Not yet No.	Learns quality without human scores; a second opinion.

Avoid 3

Model	Commercial use	Loads today	Good to know
Q-Align / OneAlign 👎 Avoid	Not rated	Not rated	S-Lab code licence (non-commercial); a 7B model to find a smudged lens is far too heavy.
IQA-PyTorch (pyiqa) 👎 Avoid	Not rated	Not rated	PolyForm Noncommercial code and CC BY-NC-SA weights; easy to pull in by accident. Use the individual MIT or Apache repositories.
DOVER and CLIP-IQA (NTU S-Lab) 👎 Avoid	Not rated	Not rated	S-Lab License 1.0, non-commercial.

● Camera tamper and picture quality 6 models · 2 load today

What it does: Flags blocked, dirty, moved or frozen cameras before an incident. **Main risk:** Rain, glare and infrared switching cause false alarms.

Model	Commercial use	Loads today	Good to know
OpenCV image checks (Laplacian blur, brightness, SSIM, frozen frame) ☆ Start here	✓ **Allowed**	✓ **Yes**	Near-zero cost on every camera every minute. Thresholds per camera.
OpenCV ORB or AKAZE matching with a homography	✓ **Allowed**	✓ **Yes**	Detects that the camera moved and by how much, so zones can be marked invalid. Needs background features.
XFeat	⚠ **With conditions** MegaDepth and COCO terms unread.	Not yet No.	Learned features that survive light change better than ORB; ORB is usually enough.
DINOv2 linear probe for fog, rain, lens dirt and glare	✓ **Allowed**	Not yet No.	Names the cause of a poor picture; needs labelled examples per condition.
LightGlue (with DISK or ALIKED)	⚠ **With conditions** Never ship the SuperPoint variant; training data unread.	Not yet No.	Accurate view-change estimate when the scene is hard for ORB.

Avoid 1

Model	Commercial use	Loads today	Good to know
pyiqa (MUSIQ, CLIP-IQA, TOPIQ, NIMA) 👎 Avoid	Not rated	Not rated	PolyForm Noncommercial library; classical checks do the job.

Picture improvement 20 models · Not in the platform yet

What it does: Easier review of dark, noisy or hazy footage. **Needs:** CPU for classical steps; GPU for learned models.
Main risk: Learned models invent detail; never evidence.

Model	Commercial use	Good to know
OpenCV classical enhancement (CLAHE, gamma, non-local-means, dehaze) ☆ Start here	✓ Allowed	Contrast lift and denoise with no invented detail. Dehaze struggles with steam and bright sky. Would run on: CPU
Real-ESRGAN (x4plus, general-x4v3)	⚠ With conditions Weights licence unstated; source photo terms unread.	Upscales a small crop for a human viewer. Invents texture; no identification or evidence. Would run on: Consumer GPU
SwinIR and Swin2SR	✓ Allowed	More faithful than GAN upscalers; Swin2SR loads in Transformers. Training set terms unread. Would run on: Consumer GPU
SPAN and community SPAN models	⚠ With conditions Training photo sets unread; no LICENSE file at HEAD.	Near real-time upscaling and denoising of a live view. Read each model card. Would run on: Consumer GPU
SCUNet	⚠ With conditions Weights terms not stated.	Blind denoising of high-ISO night footage. Would run on: Consumer GPU
NAFNet	⚠ With conditions Weights terms unstated; benchmark sets are research-only.	Fast denoise and motion-blur removal; best speed-quality trade-off. Would run on: Consumer GPU
Retinexformer	⚠ With conditions Weights terms not stated.	Low-light enhancement with a tiny network. Photographic sets, not infrared CCTV. Would run on: Consumer GPU
OpenCV videostab (opencv_contrib)	✓ Allowed	Stabilises a shaking pole-mounted camera. Cropping cuts the field of view. Would run on: CPU
HAT (Hybrid Attention Transformer)	⚠ With conditions Weights licence unstated.	Highest still-image super-resolution among open transformers; heavy for video. Would run on: Consumer GPU (slow)
Restormer	⚠ With conditions Weights terms not stated.	Deblur, denoise, deraining for stills; too slow for many streams. Would run on: Consumer GPU (heavy)
HVI-CIDNet	⚠ With conditions Weights terms not stated.	Low-light with fewer colour shifts; newer, less tested. Would run on: Consumer GPU
DehazeFormer	⚠ With conditions Trained on RESIDE; weights terms not stated.	Haze on outdoor yards. Never dehaze the stream a fire detector reads. Would run on: Consumer GPU
BasicVSR++ (via MMagic)	⚠ With conditions Weights terms not stated.	Offline clean-up of a short clip. The licence shown is MMagic's. Would run on: Consumer GPU (offline)
vid.stab	⚠ With conditions LGPL; dynamic linking only, take advice.	Mature two-pass stabiliser used inside ffmpeg. Would run on: CPU

Avoid 6

Model	Commercial use	Good to know
Zero-DCE, Zero-DCE++ and SCI 👎 Avoid	Not rated	No licence file in either repository, so no right to use. Use Retinexformer or CLAHE.
VRT and RVRT 👎 Avoid	Not rated	CC BY-NC 4.0, non-commercial.
DiffBIR 👎 Avoid	Not rated	Diffusion invents plausible detail, the worst fit for evidence; built on Stable Diffusion weights with OpenRAIL++-M terms.

Model	Commercial use	Good to know
GFPGAN Avoid	Not rated	Generates a face that is not the person's; a confident wrong identity is serious harm.
CodeFormer Avoid	Not rated	S-Lab non-commercial licence and the same invented-face problem.
DUT video stabilisation Avoid	Not rated	LICENSE says MIT but the README says research only; the conflict is the risk.

Image restoration for CCTV frames 4 models · Not in the platform yet

What it does: Denoise, deblur and upscale saved frames. **Needs:** A GPU; side-by-side original; enhanced watermark. **Main risk:** Invented detail; face restorers are banned.

Model	Commercial use	Good to know
Real-ESRGAN Start here	With conditions Weights trained on DIV2K, which is academic-only.	Upscales and cleans a still. Mark restored images; it invents texture. Would run on: Consumer GPU (CPU for single stills)
SwinIR, HAT, NAFNet, Restormer, SCUNet, BSRGAN Unconfirmed	With conditions Training sets are academic-only.	Denoise, deblur and super-resolution. Retrain on licensed or own data. Would run on: Consumer GPU
RealBasicVSR, FlashVSR and DiffBIR Unconfirmed	With conditions Weights and data not read.	Cleans short clips. Diffusion restoration hallucinates the most. Would run on: Consumer GPU

Avoid 1

Model	Commercial use	Good to know
CodeFormer, GFPGAN and SUPIR Avoid	Not rated	They rebuild faces and can make a stranger look like a suspect; avoid on ethical grounds whatever the licence.

Depth and geometry from one camera 17 models · Not in the platform yet

What it does: Distances and floor-plan positions from one fixed camera. **Needs:** One floor calibration per camera; floor plan. **Main risk:** Model depth is off by percent to tens of percent; location tracking of workers.

Model	Commercial use	Good to know
Ground-plane homography from clicked floor points (OpenCV) Start here	Allowed	The reliable way to measure floor distance in metres. Flat floor only; redo when the camera moves. Would run on: CPU
Depth Anything V2 Small	Allowed	Relative depth only. Base, Large and Giant are CC-BY-NC-4.0. Would run on: CPU or consumer GPU
Depth Anything 3 (Small, Base, Metric-Large, Mono-Large)	Allowed	Metric depth (DA3METRIC-LARGE) for yards. Take only the Apache checkpoints; check metres against a tape. Would run on: Consumer GPU
GeoCalib	Allowed	Estimates focal length, distortion and tilt from one picture. Weak on featureless scenes. Would run on: Consumer GPU or CPU
MoGe and MoGe-2	Allowed	Point map with metric scale, normals and field of view in one pass. Training-data terms unread. Would run on: Consumer GPU
MiDaS and DPT (Intel)	Allowed	Older relative depth with first-class Transformers support. Yes is at licence level only. Would run on: CPU or consumer GPU
Depth Anything V2 metric (Hypersim, Virtual KITTI 2) Unconfirmed	With conditions Dataset terms unread; outdoor inherits a non-commercial set.	Metres directly. Prefer DA3METRIC-LARGE or MoGe-2. Would run on: Consumer GPU

Model	Commercial use	Good to know
ZoeDepth	⚠ **With conditions** KITTI is non-commercial.	Early metric depth; superseded. Would run on: Consumer GPU
Marigold (diffusion depth)	⚠ **With conditions** Training data includes a non-commercial set.	Very sharp relative depth for one still. Slow. Would run on: Consumer GPU (slow)
Metric3D v2	⚠ **With conditions** Weights terms unconfirmed; authors invite commercial enquiries.	Metric depth that needs true focal length. Would run on: Large GPU
Video Depth Anything	✓ **Allowed**	Steady depth for video; use Small only. Would run on: Consumer GPU
WildCamera	⚠ **With conditions** Weights licence not found.	Focal length from an in-the-wild image; an alternative to GeoCalib. Would run on: Consumer GPU

Avoid 5

Model	Commercial use	Good to know
Apple Depth Pro Avoid	Not rated	Apple research-only licence.
UniDepth (V1, V2) Avoid	Not rated	CC BY-NC 4.0.
VGGT (including VGGT-1B-Commercial) Avoid Unconfirmed	Not rated	VGGT-1B is CC BY-NC 4.0; the commercial licence text is gated and unread; multi-view anyway.
DUSt3R and MASt3R Avoid	Not rated	CC BY-NC-SA 4.0.
Perspective Fields Avoid	Not rated	Adobe noncommercial research licence.

Segmentation for zones and masks 14 models · Not in the platform yet

What it does: One click on a floor or door gives a clean zone outline. **Needs:** A setup screen where a human accepts each outline. **Main risk:** Rough edges at shadows and glass; some models are non-commercial or AGPL.

Model	Commercial use	Good to know
SAM 2 and SAM 2.1 (Tiny, Small, Base+, Large) ☆ Start here	✓ **Allowed**	Click a point to get a zone outline; tracks a mask through a clip. The dataset chain is a theoretical risk. Would run on: Consumer GPU (Tiny and Small on a modest one)
RF-DETR Seg (N to 2XL)	✓ **Allowed**	Person and vehicle masks for foot-in-zone tests. COCO classes only. Would run on: Consumer GPU
Segment Anything (SAM 1: ViT-B, L, H)	✓ **Allowed**	Well-known still-image click-to-mask; SAM 2.1 is the better start. Would run on: Consumer GPU (ViT-H heavy)
MobileSAM	✓ **Allowed**	Click-to-mask without a strong GPU. Rougher edges. Would run on: CPU or any GPU
EfficientViT-SAM and EfficientSAM	✓ **Allowed**	Faster alternatives to SAM. Port effort. Would run on: Consumer GPU
SAM 3	⚠ **With conditions** Custom licence with termination and litigation clauses; read it closely.	Segments every instance from a text phrase. Would run on: Consumer GPU to large GPU
Florence-2 (base, large)	✓ **Allowed**	Referring segmentation in one small model; masks less accurate than SAM. Would run on: CPU or consumer GPU

Model	Commercial use	Good to know
Grounded-SAM-2 recipe (Grounding DINO plus SAM 2)	Allowed	Type "conveyor belt" and get a mask for a zone proposal. A human confirms every zone. Would run on: Consumer GPU
OneFormer and Mask2Former (ADE20K semantic)	With conditions ADE20K and Cityscapes terms are research-oriented.	Labels floor, wall, door and stairs to suggest walkable zones. Would run on: Consumer GPU
EoMT (Encoder-only Mask Transformer)	With conditions ADE20K and COCO terms unconfirmed.	Fast plain-ViT segmentation. Would run on: Consumer GPU
CLIPSeg	With conditions Conflicting statements about the weights.	Cheap text-prompt segmentation. Get a clear answer first. Would run on: CPU or consumer GPU

Avoid 3

Model	Commercial use	Good to know
SegFormer (NVIDIA) Avoid	Not rated	NVIDIA Source Code License, research or evaluation only; easy to adopt by accident.
FastSAM Avoid	Not rated	AGPL-3.0 (YOLOv8-seg based); MobileSAM does the job under Apache.
EdgeSAM and Bria RMBG Avoid	Not rated	S-Lab and Bria non-commercial terms. BiRefNet (MIT) covers background removal.

Zone outlines and depth hints 3 models · 2 load today

What it does: Two helpers the Models page installs. Neither runs on every frame and neither enables a safety check by itself.

Model	Commercial use	Loads today	Good to know
SAM 2.1 Hiera Tiny (zone outline) Start here	Allowed	Yes	No pinned ONNX file exists, so the install downloads the commit-pinned checkpoint, checks its SHA-256, and exports the encoder and decoder on the site with the trainer's torch. Small, base+ and large are not offered: too slow on a processor.
Depth Anything V2 Small (depth hint)	Allowed	Yes	One commit-pinned ONNX file from the onnx-community mirror, checked by SHA-256. Relative depth only, never metres.

Avoid 1

Model	Commercial use	Loads today	Good to know
Depth Anything V2 Base and Large Avoid	Not rated	Not rated	CC BY-NC-4.0, non-commercial. They are not offered.

Optical flow and motion 10 models · Not in the platform yet

What it does: Direction and speed of movement, a stopped conveyor, a moved camera. **Needs:** Frame pairs; floor calibration for real units. **Main risk:** Noise in rain and glare; must not score workers.

Model	Commercial use	Good to know
OpenCV DIS and Farneback dense optical flow Start here	Allowed	Direction and speed per zone, camera shake, conveyor stopped. Noisy on plain surfaces and in rain. Would run on: CPU
SEA-RAFT (S, M, L)	With conditions Pick after a licence answer; KITTI and Sintel terms.	Accurate learned flow at modest cost for vehicle speed and direction. Would run on: Consumer GPU

Model	Commercial use	Good to know
RAFT (original and torchvision port)	⚠ **With conditions** KITTI is non-commercial; synthetic-only weights avoid it.	The reference learned flow with first-class torchvision support. Would run on: Consumer GPU
NeuFlow v2 Unconfirmed	⚠ **With conditions** Weights licence not found.	Flow on edge hardware. The two speed claims disagree; test on the real board. Would run on: Jetson-class board or consumer GPU
WAFT	⚠ **With conditions** Weights licence not found.	Newer warping-based flow from the RAFT group. Would run on: Consumer GPU
UniMatch (flow, stereo, depth)	⚠ **With conditions** Weights licence unstated.	One codebase for flow, stereo and depth. Would run on: Consumer GPU
GMFlow, FlowFormer and MemFlow Unconfirmed	⚠ **With conditions** Weights terms not stated.	Research-grade; heavier than SEA-RAFT for little gain on CCTV. Would run on: Consumer GPU to large GPU
PTLFlow (library of flow models)	⚠ **With conditions** Every bundled model keeps its own licence.	Try many flow models behind one interface. Would run on: Consumer GPU
TAPIR and TAPNext (point tracking) Unconfirmed	⚠ **With conditions** Weights licence not found.	Follows chosen points to measure vibration or drift. Would run on: Consumer GPU

Avoid 1

Model	Commercial use	Good to know
CoTracker Avoid	Not rated	CC BY-NC 4.0, non-commercial.

Reading text in a scene and documents 14 models · Not in the platform yet

What it does: Reads wagon numbers, signs, displays and scanned forms. **Needs:** Text detector and reader on a crop; check-digit rules. **Main risk:** Confident wrong digits; Sinhala reading is weak.

Model	Commercial use	Good to know
PaddleOCR PP-OCRv5 (detection and recognition) ☆ Start here	✓ **Allowed**	Wagon, container and sign text in a crop; English and Tamil (94.2 on Baidu's own Tamil set, vendor figure). No Sinhala recogniser. Would run on: CPU (mobile); consumer GPU (server)
PaddleOCR-VL (0.9B) and PaddleOCR-VL-1.5	✓ **Allowed**	Whole-page forms and tables. Card says 109 languages but names neither Sinhala nor Tamil; test both. Would run on: Consumer GPU
Tesseract 5 with tessdata_best (sin, tam, eng)	✓ **Allowed**	The only engine checked with a ready Sinhala model. Weak on photos and low light; Sinhala accuracy unmeasured. Would run on: CPU
GLM-OCR (0.9B)	✓ **Allowed**	Document parsing with tables; small enough for a site GPU. Sinhala and Tamil not read. Would run on: Consumer GPU
dots.ocr	✓ **Allowed**	Layout plus text for many scripts; worth a Sinhala and Tamil test. Use the original rednote-hilab repo. Would run on: Consumer GPU
Granite-Docling-258M with Docling	✓ **Allowed**	PDF and scanned permits to tables on a CPU-only server. English-centric. Would run on: CPU possible; small GPU faster
DeepSeek-OCR	✓ **Allowed**	Long dense documents. Heavier than PaddleOCR-VL; Sinhala and Tamil not read. Would run on: Consumer GPU
olmOCR 2 (7B)	✓ **Allowed**	Scanned pages and forms to clean text; English-centric and heavy. Would run on: Consumer GPU (16 GB class)

Model	Commercial use	Good to know
EasyOCR Unconfirmed	⚠ **With conditions** No weights licence stated.	Easy API with a Tamil model; no Sinhala. Would run on: CPU or consumer GPU
docTR (Mindee) Unconfirmed	⚠ **With conditions** Pretrained-weight terms not read.	Latin-script documents and receipts only. Would run on: CPU or consumer GPU

Avoid 4

Model	Commercial use	Good to know
Surya and Chandra (Datalab) Avoid	Not rated	Modified OpenRAIL-M weights, free only below a funding or revenue cap; a closed product sold to many customers needs a paid licence.
HunyuanOCR (Tencent) Avoid	Not rated	Licence excludes the EU, UK and South Korea and caps users at 100 million a month.
Nanonets-OCR (s and 2-3B) Avoid	Not rated	No licence stated and the base model is research-only.
Nougat (Meta) Avoid	Not rated	CC BY-NC 4.0 weights; academic PDFs only.

Barcodes and QR codes 4 models · Not in the platform yet

What it does: Gate staff scan permits, passes and labels with any camera. **Needs:** A sharp frame; crop for distant codes. **Main risk:** Codes can be copied; LGPL needs care.

Model	Commercial use	Good to know
ZXing-cpp Start here	✓ **Allowed**	Reads QR, Data Matrix, Code 128, EAN and more, with Python and C++ bindings. Would run on: CPU
WeChat QR code detector (OpenCV contrib) Unconfirmed	⚠ **With conditions** Weights licence not read.	QR only; handles tiny or blurred codes. Built on ZXing. Would run on: CPU
ZBar	⚠ **With conditions** LGPL needs dynamic linking.	Old, fast and widely packaged; use ZXing-cpp first. Would run on: CPU
OpenCV QRCodeDetector	✓ **Allowed**	Used only at kiosk cameras while a visitor's pass is being scanned; no extra install. Would run on: CPU

Visual anomaly detection against a normal reference 6 models · Not in the platform yet

What it does: Flags anything that differs from a normal view. **Needs:** 100 to 500 clean frames per zone; a GPU to train. **Main risk:** Alert fatigue; benchmarks do not carry over to CCTV.

Model	Commercial use	Good to know
PatchCore Start here	⚠ **With conditions** ImageNet backbone terms are research-oriented; use a DINOv2 backbone.	Few-shot "something differs" on a fixed view, with a heat map. Up to 99.6 percent image AUROC on MVTec AD (not re-read). Would run on: Consumer GPU to build; CPU or GPU to infer
Anomalib (framework)	✓ **Allowed**	Runs PatchCore, EfficientAD, Dinomaly, WinCLIP and AnomalyDINO. Each model inside has its own backbone terms. Would run on: CPU or consumer GPU (OpenVINO export)
Dinomaly	✓ **Allowed**	One model for many zones. Yes only if trained on site frames; do not ship the published checkpoints. Would run on: Consumer GPU

Model	Commercial use	Good to know
EfficientAD	⚠ **With conditions** ImageNet images in training; check the teacher weights.	Very fast inference, suits many streams. Would run on: CPU or consumer GPU
WinCLIP and AnomalyDINO (zero or few-shot) Unconfirmed	⚠ **With conditions** Ports not read; CLIP trained on undisclosed web data.	Few or no normal examples needed, to start. Less precise than PatchCore. Would run on: Consumer GPU
RD4AD (reverse distillation)	⚠ **With conditions** ImageNet backbone terms unread.	A classic reconstruction-style baseline. Older. Would run on: Consumer GPU

Anomaly detection in video without labels 5 models · Not in the platform yet

What it does: A ranked review list of moments that look different. **Needs:** Days of normal footage per fixed camera; review screen. **Main risk:** Many false alarms; it says different, not dangerous.

Model	Commercial use	Good to know
Anomalib (PatchCore, EfficientAD, FastFlow, Dinomaly, STFPM and others) ☆ Start here	✓ **Allowed**	"This view looks different" on a still scene. Image-level, not motion-aware. Never use the MVTec-trained checkpoints; record the ImageNet risk decision. Would run on: CPU for tens of images; GPU speeds it; OpenVINO export
AI-VAD (in anomalib, video)	⚠ **With conditions** Component weight terms unread.	Flags unusual movement against a fitted normal; explainable. Heavy; report per zone, never per person. Would run on: Consumer GPU
AnomalyDINO (few-shot, DINOv2 patch features)	✓ **Allowed**	Works from 1 to 10 normal reference images; quick to try on one zone. Sensitive to camera shift and light. Would run on: Consumer GPU, or CPU for small inputs
WinCLIP (zero and few-shot, in anomalib) Unconfirmed	⚠ **With conditions** CLIP weights licence not read.	Scores from a text description plus a few normal images. English prompts only. Would run on: Consumer GPU

Avoid 1

Model	Commercial use	Good to know
Supervised crime and anomaly video models (VadCLIP, Holmes-VAD and others trained on UCF-Crime, XD-Violence) 👎 Avoid Unconfirmed	Not rated	"crime" classes do not match plant events and the data is research-only.

Industrial inspection and condition 9 models · Not in the platform yet

What it does: Watches the plant: stopped conveyor, guard state, lamps, spills. **Needs:** Fixed cameras and site labels; thermal camera for hot spots. **Main risk:** A wrong all-clear is worse than no alert.

Model	Commercial use	Good to know
Anomalib PatchCore or EfficientAD trained on good parts ☆ Start here	✓ **Allowed**	Surface defects and odd states on a fixed-pose station from 20 to 200 good images. Needs identical pose and lighting. Would run on: CPU to consumer GPU
RF-DETR or RF-DETR-Seg fine-tuned on site states (lamp colour, open door, missing guard, spill, jammed belt) Unconfirmed	✓ **Allowed**	Lamp lit or not, door open, guard missing, liquid on the floor. Needs 100 to 500 labelled images per state. Would run on: Consumer GPU
OpenCV optical flow (DIS, Farneback) for conveyor stall and jam	✓ **Allowed**	Is the belt moving at normal speed? Confirm with a PLC signal; a plain belt has little texture. Would run on: CPU

Model	Commercial use	Good to know
DINOv2 or SigLIP 2 linear probe for door, guard or valve state	✓ Allowed	A two-class check from about 50 images per state. Camera shift ruins it; retrain per camera. Would run on: CPU for small models
Florence-2 and PaddleOCR on a cropped display, nameplate or digital gauge	✓ Allowed	Reads digits on a seven-segment display or plate. Yes applies to Florence-2 only; needle dials are not its strength. Would run on: Consumer GPU or CPU
SEA-RAFT and RAFT optical flow	⚠ With conditions Weights provenance (Sintel and others) unread.	Dense motion field for jams and spillage when classical flow is too weak. Heavy for a stall check. Would run on: Consumer GPU
Thermal datasets and models (Teledyne FLIR ADAS, Roboflow Universe thermal sets) Unconfirmed	⚠ With conditions FLIR terms unread; per-project licences.	Pre-training a thermal person or vehicle detector. For equipment hot spots use the camera's radiometric threshold. Would run on: n/a (dataset)

Avoid 2

Model	Commercial use	Good to know
Analog gauge readers (ETH analog_gauge_reader, hcltech-robotics/gauge_reader) Avoid	Not rated	No code licence, an AGPL dependency and untraceable weights. Build your own on a detector plus keypoints.
MVTec AD and VisA anomaly datasets Avoid	Not rated	MVTec AD is CC BY-NC-SA 4.0 and checkpoints trained on it are non-commercial; VisA is CC BY 4.0 but the row mixes the two, so no is the safe label.

Sound and speech

Sound events 11 models · Not in the platform yet

What it does: Hears alarms, breaking glass, bangs and calls for help. **Needs:** A microphone per zone; per-site calibration. **Main risk:** False alarms; audio capture is regulated.

Model	Commercial use	Good to know
EfficientAT (MobileNetV3 and dynamic MN, AudioSet) Start here	⚠ With conditions Weights licence not stated; ask the author or retrain.	Small 527-class tagger; best accuracy per CPU cycle here. mAP 43.2 (mn04) to about 48 (paper). Would run on: CPU
YAMNet Unconfirmed	⚠ With conditions Weights licence not confirmed.	Easiest start: 521 classes, 0.96 s windows, runs on a Raspberry Pi 3. Older and less accurate; needs TFLite or ONNX conversion. Would run on: CPU
AST (MIT/ast-finetuned-audioset-10-10-0.4593)	✓ Allowed	Higher accuracy (mAP 0.459) and a strong base to fine-tune a site alarm. Same AudioSet caveat. Would run on: Consumer GPU
PANNs (CNN14 and relatives)	✓ Allowed	Mature; gives frame-level timestamps. mAP 0.431 to 0.439. Attribution needed. Would run on: CPU
BEATs (Microsoft) Unconfirmed	⚠ With conditions Checkpoint licence unconfirmed; ask Microsoft.	Top accuracy (mAP about 48.6, paper) and good embeddings for machine sound. Would run on: Consumer GPU
PaSST Unconfirmed	⚠ With conditions Weights licence not confirmed.	Fast transformer tagger, mAP 0.476 by checkpoint name. No updates since 2024-01. Would run on: Consumer GPU
CED (mispeech)	✓ Allowed	Distilled for speed; mAP about 0.48 for base. Code licence not confirmed. Would run on: CPU (tiny, mini); GPU (base)

Model	Commercial use	Good to know
LAION CLAP	⚠ With conditions Training collection mixes source licences.	Type a sound and match audio with no training; good for testing new classes. Less precise than a fine-tuned tagger. Would run on: Consumer GPU
Microsoft CLAP	⚠ With conditions Some training sets are non-commercial.	Audio-text retrieval and captioning; repository archived 2024-09. Would run on: Consumer GPU
Qwen2-Audio-7B-Instruct	✓ Allowed	Ask questions about a flagged clip; too heavy for always-on use and answers can be invented. Would run on: Large GPU (about 16 GB in half precision)

Avoid 1

Model	Commercial use	Good to know
Community AST fine-tunes for gunshot Avoid · Unconfirmed	Not rated	0 to 14 downloads each, no false-alarm rate, undocumented data. Train your own on site recordings.

Speech to text 17 models · Not in the platform yet

What it does: Spoken notes become searchable text on the incident log. **Needs:** Push-to-talk button, ASR on the server, site tests. **Main risk:** Push-to-talk only; Sinhala is the weak spot; invented text on noise.

Model	Commercial use	Good to know
Whisper (large-v3, large-v3-turbo, medium, small, base, tiny) ☆ Start here	✓ Allowed	Mature English and radio-quality notes. Tamil is weaker; Sinhala is unreliable out of the box. Run voice-activity gating first. Would run on: CPU (tiny to small); consumer GPU (large)
faster-whisper and whisper.cpp (runtimes) Unconfirmed	✓ Allowed	Runs Whisper several times faster in less memory. Licence files not re-read. Would run on: CPU or consumer GPU
Omnilingual ASR (Meta)	✓ Allowed	The only permissive model found that lists Sinhala and Tamil. 7B CER: Sinhala 7.2, Tamil 5.2 (researcher's figure; CER is not WER). Needs fairseq2. Would run on: Consumer GPU (300M to 1B); large GPU (7B)
Moonshine (Useful Sensors)	✓ Allowed	Very small English streaming ASR for a kiosk. Do not use the non-commercial language models. Would run on: CPU
NVIDIA Parakeet TDT 0.6B and Canary-1B-v2	✓ Allowed	Best permissive English accuracy and speed. No Sinhala or Tamil. Needs NeMo; attribution required. Would run on: Consumer GPU
AI4Bharat IndicConformer 600M	⚠ With conditions Gated download; training-data terms not read.	Strong Tamil candidate. Sinhala is not covered. Would run on: Consumer GPU
sherpa-onnx (k2-fsa) Unconfirmed	✓ Allowed	One CPU runtime for ASR, TTS, VAD and keyword spotting. Check each model you load. Would run on: CPU
IBM Granite Speech (3.3 8B, 4.1 2B)	✓ Allowed	English-first speech with a clear licence trail. Would run on: Consumer GPU (2B); large GPU (8B)
distil-whisper large-v3	✓ Allowed	English only; about six times faster than large-v3. Would run on: Consumer GPU or fast CPU
Kyutai STT (stt-2.6b-en, stt-1b-en_fr)	✓ Allowed	Low-latency streaming English and French. Would run on: Consumer GPU
Qwen3-ASR (0.6B, 1.7B)	✓ Allowed	30 languages; no Tamil or Sinhala on the card. A community Sinhala fine-tune exists, quality and licence unknown. Would run on: Consumer GPU
Mistral Voxtral (Mini 3B, Small 24B)	✓ Allowed	Speech plus questions about audio. No Sinhala or Tamil. Would run on: Consumer GPU (3B); large GPU (24B)

Model	Commercial use	Good to know
Community Sinhala fine-tunes of Whisper and XLS-R Unconfirmed	⚠ **With conditions** CC BY-SA training data; WER unverified.	Shows Sinhala fine-tuning is feasible. Fine-tune your own on data you hold. Would run on: CPU (small); consumer GPU (large)
Vosk Unconfirmed	⚠ **With conditions** Weights licence not read.	Small offline recognisers on weak CPUs. Tamil or Sinhala models not confirmed. Would run on: CPU
SenseVoice Small (Alibaba FunASR)	⚠ **With conditions** Custom weights licence.	Fast Chinese, English, Japanese and Korean. No Sinhala or Tamil; switch off its emotion and event tags. Would run on: CPU or consumer GPU

Avoid 2

Model	Commercial use	Good to know
Meta MMS (mms-1b-all) Avoid	Not rated	CC-BY-NC-4.0 weights. Covers Sinhala and Tamil, so use only as a benchmark to beat.
SeamlessM4T v2 large Avoid	Not rated	CC-BY-NC-4.0, non-commercial.

Text to speech 13 models · Not in the platform yet

What it does: The kiosk or a site speaker says a message aloud. **Needs:** Speaker or paging unit; recorded phrases for fixed alarms. **Main risk:** No good licensed Sinhala voice; no voice cloning.

Model	Commercial use	Good to know
Kokoro-82M Start here	✓ **Allowed**	Small, fast, natural English voice. No Sinhala or Tamil. Data includes audio from closed TTS models. Would run on: CPU
Indic Parler-TTS (AI4Bharat)	⚠ **With conditions** Gated; training-data terms not read.	Tamil (voices Kavitha and Jaya) and English. No Sinhala. Takes seconds, so pre-render fixed messages. Would run on: Consumer GPU (CPU slow)
Piper	⚠ **With conditions** Active code is GPL-3.0; voice licences differ.	Very fast; 38 English voices; community Sinhala voices exist (OpenSLR 30, CC BY-SA). No official Sinhala or Tamil voice. Check the GPL terms. Would run on: CPU
Parler-TTS Mini v1	✓ **Allowed**	English voices with controllable tone. Would run on: Consumer GPU
Chatterbox (Resemble AI)	✓ **Allowed**	Very natural English. Leave voice cloning disabled. Would run on: Consumer GPU
Orpheus TTS 3B	⚠ **With conditions** Built on Llama 3.2; the Llama licence may bind.	Expressive English. Would run on: Consumer GPU
Dia 1.6B and VibeVoice 1.5B	✓ **Allowed**	Long multi-speaker English dialogue; not needed for alarms. Can make convincing fake voices. Would run on: Consumer GPU
espeak-ng Unconfirmed	⚠ **With conditions** GPL-3.0; call as a separate process and read the licence first.	Last-resort robotic Sinhala and Tamil for variable words. Would run on: CPU
Community Sinhala TTS (VITS and Piper on OpenSLR 30) Unconfirmed	⚠ **With conditions** CC BY-SA data; unknown quality and voice consent.	Shows Sinhala neural TTS exists; hobby quality. Better to commission your own recordings. Would run on: CPU

Avoid 4

Model	Commercial use	Good to know
Meta MMS-TTS (Tamil, Sinhala) Avoid	Not rated	CC-BY-NC-4.0 weights; benchmark only.

Model	Commercial use	Good to know
Coqui XTTS-v2 Avoid	Not rated	Coqui Public Model License, non-commercial; the company stopped.
F5-TTS and Fish OpenAudio S1-mini Avoid	Not rated	CC-BY-NC weights, trained on Emilia (non-commercial).
OpenVoice (voice cloning) Avoid	Not rated	MIT and commercial use is allowed, but avoid on purpose. Cloning a supervisor's voice is impersonation and a way to fake an emergency order.

Voice activity and wake-word 7 models · Not in the platform yet

What it does: A hands-free kiosk and a speech gate. **Needs:** A kiosk microphone; a trained wake word. **Main risk:** Always-on microphone; shipped wake-word models are non-commercial.

Model	Commercial use	Good to know
Silero VAD Start here	Allowed	Detects speech, not words, so it works for Sinhala and Tamil. Under 1 ms per chunk on one CPU thread (README). The separate silero-models repo has an unclear licence. Would run on: CPU
openWakeWord	With conditions Shipped models are non-commercial; use only the training pipeline.	Train your own wake word, even in Sinhala or Tamil. 15 to 20 models on one Pi 3 core (README). Check the licences of the data you feed it. Would run on: CPU
WebRTC VAD (py-webrtcvad) Unconfirmed	Allowed	Cheapest speech gate, no weights. Noisier than Silero on a loud floor. Would run on: CPU
microWakeWord Unconfirmed	With conditions Pre-trained model terms unconfirmed.	Tiny wake-word models for microcontrollers. Would run on: CPU (microcontroller)
TEN VAD (Agora)	With conditions May not compete with Agora's offerings.	Very low latency. The competition clause is vague for a white-label product; Silero does the job. Would run on: CPU
Picovoice Porcupine Unconfirmed	With conditions Paid plan needed; key may phone home.	Best wake-word accuracy, but a poor fit for an offline site. Would run on: CPU
pyannote segmentation-3.0	With conditions Gated; training-corpus terms unclear.	Voice activity and overlap detection inside a recording; Silero is simpler. Would run on: CPU

Machine sound anomaly detection 3 models · Not in the platform yet

What it does: Flags a machine that sounds different from its normal. **Needs:** Close microphone and weeks of normal audio per machine. **Main risk:** Per-machine tuning; alarms on every shift change.

Model	Commercial use	Good to know
Dasheng (mispeech) audio encoder Start here	Allowed	General embeddings; measure distance from each machine's normal. No published machine-fault accuracy found. Would run on: Consumer GPU (CPU works)
EfficientAT or BEATs embedding with a per-machine outlier score Unconfirmed	With conditions Weights licences of both unconfirmed.	The sound-event models used as frozen feature extractors. Needs per-machine thresholds. Would run on: CPU
MIMII dataset and DCASE baselines Unconfirmed	With conditions Share-alike; record your own machines.	Public normal and faulty valves, pumps, fans and rails. Test a pipeline with it, do not ship it. Would run on: n/a

Search and language

Video search and embeddings 20 models

What it does: Turns pictures into numbers, so recorded video can be searched by words or by example.

Model	Commercial use	Loads today	Good to know
SigLIP 2 (google/siglip2-*) ☆ Start here	✓ Allowed	Not yet Needs its own integration.	Default text-to-image search over sampled frames; the base model runs on CPU.
SigLIP (original, google/siglip-*) ☆ Start here	✓ Allowed	Not yet Needs its own integration.	Mature, widely exported predecessor of SigLIP 2; English text only.
Perception Encoder PE-Core (Meta)	✓ Allowed	Not yet Needs its own integration.	Best-tested option for clips and actions; needs Meta's perception_models code; do not confuse with Perception-LM.
OpenAI CLIP	✓ Allowed	Not yet Needs its own integration.	Cheapest baseline (B/32 runs on CPU); older and weaker than SigLIP 2, English only.
OpenCLIP LAION models	✓ Allowed	Not yet Needs its own integration.	Open recipe and many sizes if you fine-tune CLIP on site footage.
Qwen3-VL-Embedding (2B, 8B)	✓ Allowed	Not yet Needs its own integration.	One space for text, frames and short clips; heavy, use it on candidate clips.
DINOv2 (search by example image)	✓ Allowed	Not yet Needs its own integration.	Image-to-image lookups (find this vehicle or person again); no text encoder.
bge-m3 and multilingual-e5-large (text embeddings)	✓ Allowed	Not yet Needs its own integration.	Multilingual text embeddings for event notes and plain-words questions.
EVA-CLIP	✓ Allowed	Not yet Needs its own integration.	Accuracy research; the file is 30.1 GB, too big for a site server.
TinyCLIP	✓ Allowed	Not yet Needs its own integration.	For CPU-only sites on a tight budget; accuracy clearly below SigLIP 2.
jina-clip-v1	✓ Allowed	Not yet Needs its own integration.	Text-to-text and text-to-image in one model with an ONNX export; do not confuse with v2.
nomic-embed-vision-v1.5	✓ Allowed	Not yet Needs its own integration.	Shares a vector space with nomic's text embeddings; smaller community.
GME-Qwen2-VL-2B	✓ Allowed	Not yet Needs its own integration.	Superseded by Qwen3-VL-Embedding.
LongCLIP-L	⚠ With conditions Built on OpenAI CLIP; training data not checked.	Not yet Needs its own integration.	Long, detailed text queries (about 248 tokens); the file is 1.7 GB.
X-CLIP (microsoft/xclip-*)	⚠ With conditions Trained on Kinetics (YouTube clips, mixed rights).	Not yet Needs its own integration.	Video clip to text matching on 8-frame inputs; built for action labels.
InternVideo2-CLIP	⚠ With conditions Gated download; web video training data.	Not yet Needs its own integration.	Native video-text retrieval research; heavy for live use.

Avoid 4

Model	Commercial use	Loads today	Good to know
jina-clip-v2 Avoid	Not rated	Not rated	CC BY-NC 4.0; commercial download needs contacting Jina sales. v1 is Apache-2.0.
MobileCLIP, MobileCLIP2 and DFN5B-CLIP (Apple) Avoid	Not rated	Not rated	Apple AMLR, research purposes only; tempting because it is fast on phones.
MetaCLIP and MetaCLIP 2 (facebook) Avoid	Not rated	Not rated	CC BY-NC 4.0.
jina-embeddings-v4 and nomic-embed-multimodal-3b Avoid	Not rated	Not rated	Qwen Research License, non-commercial, inherited from Qwen2.5-VL-3B.

Vision-language models 24 models

What it does: Answers questions about a picture or a clip in plain words.

Model	Commercial use	Loads today	Good to know
Qwen3-VL (2B, 4B, 8B, 32B, 30B-A3B, 235B-A22B) Start here	Allowed	Not yet Needs its own integration.	First model to try for event descriptions; 4B Q4 fits a 6 GB GPU, 8B Q4 an 8 to 12 GB GPU (file sizes only).
Gemma 4 (E2B, E4B, 12B, 26B-A4B, 31B) Licence-safe alternative	Allowed	Not yet Needs its own integration.	Strong small model with an official 4-bit build; download Gemma 4, not 3.
Qwen3.5 (0.8B, 2B, 4B, 9B, 27B, 35B-A3B, 122B-A10B, 397B-A17B)	Allowed	Not yet Needs its own integration.	Newer natively multimodal line; no official GGUF repo exists, check runtime support.
SmolVLM2 (256M, 500M, 2.2B) and SmolVLM	Allowed	Not yet Needs its own integration.	Smallest video-capable model, runs on CPU; machine-generated caption terms are mixed.
MiniCPM-V 4.6 and 4.5	Allowed	Not yet Needs its own integration.	Image and video understanding small enough for a CPU box; only the 4.x weights were checked.
InternVL3.5 (1B to 38B)	Allowed	Not yet Needs its own integration.	Alternative to Qwen3-VL at 2B to 8B; check the exact repo's licence.
Ministral 3 (3B, 8B, 14B)	Allowed	Not yet Needs its own integration.	Image understanding plus chat in one model; image only, no video.
Qwen3.6 (27B, 35B-A3B)	Allowed	Not yet Needs its own integration.	Strongest Apache option if the site has a 24 GB+ GPU; too big for small boxes.
Molmo2 (4B, 8B, O-7B)	Allowed	Not yet Needs its own integration.	Pointing at and tracking objects in video; no official GGUF.
Florence-2 (base, large)	Allowed	Not yet Needs its own integration.	Cheap captions, grounding and OCR on a still frame; not conversational, no video.
Phi-4-multimodal-instruct	Allowed	Not yet Needs its own integration.	MIT image plus audio plus text; image only, weaker than Qwen3-VL at the same size.
Mistral Small 3.2 24B and Pixtral 12B	Allowed	Not yet Needs its own integration.	Larger Apache vision model for a 24 GB GPU; Pixtral is older and was not re-read.
Granite Vision 4.1 4B	Allowed	Not yet Needs its own integration.	For tables, forms and charts, not scenes.

Model	Commercial use	Loads today	Good to know
GLM-4.6V-Flash and GLM-4.1V-9B-Thinking	Allowed	Not yet Needs its own integration.	MIT at 9B for a 16 GB GPU; thinking variants write long reasoning first.
Moondream 2	Allowed	Not yet Needs its own integration.	Tiny captioner and pointer; older and weaker than newer small models.
Moondream 3 Preview	With conditions Must not compete with M87 Labs' paid offerings; check local rules.	Not yet Needs its own integration.	A CCTV analytics product is probably not competing, but embedding it in a paid product is the clause to check.
LFM2-VL and LFM2.5-VL (Liquid AI)	With conditions Revenue threshold of 10 million US dollars.	Not yet Needs its own integration.	Very efficient on CPU; a revenue cap is a poor fit for a reseller.
NVIDIA Cosmos-Reason1 and Cosmos-Reason2	With conditions Read the full agreement; Cosmos-Reason2-8B is gated.	Not yet Needs its own integration.	Reasoning about what happens in a video clip (safety events).
Qwen2.5-VL 7B and 32B (Apache) versus 3B and 72B	Allowed	Not yet Needs its own integration.	Older Apache option; Qwen3-VL replaces it.
DeepSeek-VL2 (tiny, small, base)	With conditions Use-based restrictions (Attachment A) must be passed on.	Not yet Needs its own integration.	Little reason to pick it over Apache models; the record shows license "other".

Avoid 4

Model	Commercial use	Loads today	Good to know
Gemma 3, Gemma 3n and PaliGemma 2 Avoid	Not rated	Not rated	Gemma Terms of Use; Google may restrict usage remotely, awkward for an offline site. Use Gemma 4.
Llama 3.2 Vision (11B, 90B) and Llama 4 Scout and Maverick Avoid	Not rated	Not rated	The 3.2 licence withholds multimodal rights from EU-based persons and companies.
Qwen2.5-VL 3B (and 72B) Avoid	Not rated	Not rated	The 3B is under the Qwen Research License, non-commercial; the 72B licence text was not read.
Meta Perception-LM Avoid	Not rated	Not rated	FAIR Noncommercial Research License; easy to confuse with PE-Core.

Language models 20 models

What it does: Writes summaries and answers questions about the site's records.

Model	Commercial use	Loads today	Good to know
Qwen3 (0.6B to 32B and 30B-A3B) Start here	Allowed	Not yet Needs its own integration.	Plain-words questions and shift summaries; official GGUF builds exist.
gpt-oss-20b and gpt-oss-120b	Allowed	Not yet Needs its own integration.	Strong reasoning on a single 16 GB GPU; replies carry thinking tokens; 120b not re-read.
Qwen3.5 and Qwen3.6 text use	Allowed	Not yet Needs its own integration.	Newest Qwen generation; no official GGUF repos.
Gemma 4 (text use)	Allowed	Not yet Needs its own integration.	Good multilingual chat at 4B to 12B with official 4-bit builds.
Phi-4 and Phi-4-mini	Allowed	Not yet Needs its own integration.	Small MIT model for summaries and structured output; English-centric.

Model	Commercial use	Loads today	Good to know
IBM Granite 4.x (micro, h-small, 3B, 8B, 30B)	Allowed	Not yet Needs its own integration.	Enterprise-minded model for tool use and structured answers; smaller community.
Ministral 3 (3B, 8B, 14B)	Allowed	Not yet Needs its own integration.	European-built, strong European language support; the older 2410 release is research-only.
Qwen3-Embedding (0.6B, 4B, 8B)	Allowed	Not yet Needs its own integration.	Text embeddings so questions can find event descriptions; a retrieval model, not a chat model.
OLMo 3 (7B, 32B)	Allowed	Not yet Needs its own integration.	For buyers who want fully documented training data; behind Qwen3 on benchmarks.
SmolLM3-3B	Allowed	Not yet Needs its own integration.	Fully open 3B model for CPU-only sites; six languages.
GLM-4-9B-0414	Allowed	Not yet Needs its own integration.	MIT 9B with Chinese and English strength.
Qwen2.5 (7B and 14B Apache)	Allowed	Not yet Needs its own integration.	Proven older model, widely quantised; Qwen3 supersedes it.
DeepSeek-R1-Distill-Qwen-7B	Allowed	Not yet Needs its own integration.	Math-style reasoning; long chain-of-thought output makes it slow.
Meta Llama 3.1, 3.2 and 3.3 text models	With conditions "Built with Llama" attribution and Meta's terms.	Not yet Needs its own integration.	Only where the buyer already approved Meta's terms; Apache alternatives are as good.
Falcon3 and Falcon-H1 (TII)	With conditions TII licence text not read.	Not yet Needs its own integration.	Nothing the Apache models do not cover.
NVIDIA Nemotron Nano and Nemotron 3 Nano Unconfirmed	With conditions Nemotron 3 licence text not read.	Not yet Needs its own integration.	Efficient hybrid models if NVIDIA's terms are accepted.
Liquid LFM2 and LFM2.5 text models	With conditions Revenue threshold of 10 million US dollars.	Not yet Needs its own integration.	Very fast on CPU; a revenue cap suits a reseller poorly.

Avoid 3

Model	Commercial use	Loads today	Good to know
Ministral-8B-Instruct-2410 Avoid	Not rated	Not rated	Mistral Research License, research only. Use Ministral 3.
Qwen2.5-3B-Instruct Avoid	Not rated	Not rated	Qwen Research License; the 7B and 14B of the same family are Apache-2.0.
Cohere Command R7B and Aya Expanse Avoid	Not rated	Not rated	CC BY-NC 4.0 (Aya Expanse not re-read).

Embeddings and re-rankers 15 models · Not in the platform yet

What it does: Search notes, reports and manuals by meaning. **Needs:** Vector store, the text, small GPU or fast CPU.

Main risk: Confident wrong results; Sinhala and Tamil retrieval unmeasured.

Model	Commercial use	Good to know
Qwen3-Reranker (0.6B, 4B, 8B) Start here	Allowed	Best-scoring open Apache re-ranker in the card table (MTEB-R 65.8 for 0.6B). Sinhala and Tamil not benchmarked. Would run on: Consumer GPU; 0.6B on CPU for small lists

Model	Commercial use	Good to know
bge-reranker-v2-m3	Allowed	Pairs with bge-m3, already named in Part 1. Behind Qwen3-Reranker (MTEB-R 57.0). Would run on: CPU for short lists; consumer GPU
gte-multilingual-reranker-base	Allowed	Small re-ranker for CPU-only sites; 70+ languages. Would run on: CPU
mxbai-rerank-base-v2	Allowed	Strong English re-ranking in a small footprint; weaker outside English and Chinese. Would run on: Consumer GPU
Snowflake arctic-embed-l-v2.0	Allowed	Multilingual retrieval embeddings; alternative to bge-m3. Sinhala and Tamil coverage not read. Would run on: Consumer GPU or CPU
gte-multilingual-base	Allowed	Small CPU embedding model with long context; 70+ languages claimed. Would run on: CPU
nomic-embed-text-v2-moe	Allowed	Multilingual embeddings in the same family as the nomic vision model in Part 1. Would run on: CPU or consumer GPU
LaBSE	Allowed	Matches a Sinhala note to an English manual line. Weaker as a question-to-passage retriever. Would run on: CPU
paraphrase-multilingual-MiniLM-L12-v2	Allowed	Cheapest multilingual option for duplicate-note detection. Older; may omit Sinhala. Would run on: CPU
granite-embedding-278m-multilingual	Allowed	Clean provenance for English; Sinhala and Tamil are outside its 12 languages. Would run on: CPU
cross-encoder/ms-marco-MiniLM (L6, L12) Unconfirmed	With conditions Trained on MSMARCO, whose data terms are not cleared.	Very fast English re-ranker. English only. Would run on: CPU

Avoid 4

Model	Commercial use	Good to know
ColPali v1.3 and ColQwen2.5 Avoid Unconfirmed	Not rated	MIT adapters on PaliGemma or Qwen2.5-VL-3B backbones with restrictive terms; get exact terms before any use.
bge-reranker-v2-gemma Avoid	Not rated	Gemma base brings the Gemma Terms; the Qwen3 re-ranker does the same job.
Jina re-rankers (v2, v3) and jina-embeddings-v3 Avoid	Not rated	CC BY-NC 4.0.
EmbeddingGemma-300m Avoid	Not rated	Gemma Terms of Use with manual approval; does not suit an offline product.

Translation 12 models · Not in the platform yet

What it does: Alerts, reports and notes in English, Sinhala and Tamil. **Needs:** 8 GB+ GPU or 8-bit CPU; native-speaker review. **Main risk:** Sinhala quality unmeasured; best-known models are non-commercial.

Model	Commercial use	Good to know
MADLAD-400 (3B, 7B, 10B MT) Start here	Allowed	Sinhala and Tamil are among its languages (not re-read for the exact pair). Card says research users and not assessed for production. Would run on: Consumer GPU; 3B on CPU slowly
Qwen3 and Gemma 4 prompted as translators	Allowed	Free-text translation with a glossary in the prompt; nothing extra to ship. Qwen3 names Sinhala among 119 languages; no score read. Would run on: Consumer GPU, already planned for the local language model

Model	Commercial use	Good to know
IndicTrans2 (en-indic 1B and 200M, indic-en 200M)	⚠ **With conditions** Gated download; training-data terms not read.	Strong Tamil to and from English. The card does not mention Sinhala. Would run on: CPU (200M); consumer GPU (1B)
M2M-100 (418M, 1.2B) and SMaLL-100	✓ **Allowed**	Permissive baseline for Sinhala and Tamil on small hardware. Older; test before use. Would run on: CPU (418M) or consumer GPU
CTranslate2 (inference engine)	✓ **Allowed**	Runs translation models in 8-bit on CPU; makes MADLAD 3B practical. The model's own licence governs. Would run on: CPU or GPU
mBART-50 many-to-many Unconfirmed	⚠ **With conditions** Weights licence not stated.	Another candidate with both languages. Would run on: Consumer GPU
Opus-MT (Helsinki-NLP) multilingual models	⚠ **With conditions** Mixed OPUS training-data terms.	Tiny CPU translators; no dedicated Sinhala or Tamil model found, and multilingual ones are weak. Would run on: CPU

Avoid 5

Model	Commercial use	Good to know
NLLB-200 Avoid	Not rated	CC BY-NC 4.0 weights. The tempting choice; use MADLAD-400.
SeamlessM4T v2 Avoid	Not rated	CC BY-NC 4.0.
TranslateGemma (4B, 12B) Avoid	Not rated	Gemma Terms of Use with manual access approval; awkward for an offline product.
Sarvam-Translate, Tower+, Tiny Aya Avoid Unconfirmed	Not rated	Sarvam is GPL-3.0 on a Gemma base; the others are tagged non-commercial (not re-read).
LibreTranslate server Avoid	Not rated	AGPL-3.0 reaches a closed product that serves it to users.

Text classification and extraction 12 models · Not in the platform yet

What it does: Types, severity, places and names pulled from notes. **Needs:** A few hundred labelled notes; spot-checked redaction. **Main risk:** Missed names; never label named workers.

Model	Commercial use	Good to know
GLiNER v2.1 (multi, large) ☆ Start here	✓ **Allowed**	Pulls names, dates, permit numbers and places with labels you choose at run time. Pick v2.1, not v0 or v1. Would run on: CPU
GLiNER PII models (urchade multi_pii-v1, knowledgator gliner-pii-base-v1.0)	✓ **Allowed**	Masks names, phones, IDs and addresses; ONNX builds exist. English-first; spot-check a sample. Would run on: CPU
NuExtract 2.0 (2B and 8B)	✓ **Allowed**	Fills a JSON template from a note, form or photo. Never use the 4B (Qwen Research License). Would run on: Consumer GPU
Microsoft Presidio	✓ **Allowed**	Finds and redacts personal data; add patterns for national ID and phone formats. No Sinhala or Tamil recognisers. Would run on: CPU
MoritzLaurer zero-shot NLI classifiers (deberta-v3-large-zeroshot-v2.0, bge-m3-zeroshot-v2.0) Unconfirmed	⚠ **With conditions** Training-data terms vary.	Sorts a note into "near miss", "injury", "equipment" with no labels. The bge-m3 version is multilingual. Would run on: CPU or small GPU
XLM-RoBERTa (base, large) as a fine-tuning base	✓ **Allowed**	Fine-tune on a few hundred labelled notes for a reliable incident-type classifier. Would run on: CPU or consumer GPU

Model	Commercial use	Good to know
spaCy English pipelines Unconfirmed	⚠ **With conditions** Pipeline weights and data terms not read.	Fast English entities and rules. No Sinhala or Tamil pipeline checked. Would run on: CPU
IndicBERTv2 and SinBERT	✓ **Allowed**	Fine-tuning bases for Tamil and Sinhala. SinBERT is from 2022 and news-domain. Would run on: CPU
mmBERT-base	✓ **Allowed**	Newer multilingual encoder to fine-tune; Sinhala and Tamil coverage not read. Would run on: CPU

Avoid 3

Model	Commercial use	Good to know
piiranha-v1 Avoid	Not rated	CC BY-NC-ND 4.0; non-commercial and no derivatives.
GLiNER v0 (gliner_base, gliner_multi) and NVIDIA gliner-PII Avoid	Not rated	V0 is CC BY-NC 4.0; the NVIDIA licence text was not read. Use the Apache v2.1 or knowledgator models.
dslim/bert-base-NER Avoid Unconfirmed	Not rated	Trained on CoNLL-2003, whose data terms are unclear.

Language identification 3 models · Not in the platform yet

What it does: Routes a note to the right translator or index. **Needs:** A short text sample; script checks cover most cases. **Main risk:** Romanised Sinhala and Tamil; short strings.

Model	Commercial use	Good to know
GlotLID (v3) ☆ Start here	✓ **Allowed**	Covers Sinhala and Tamil among 1,000+ languages. Keep the copyright notice on training data in the product notices. Would run on: CPU
fastText lid.176	⚠ **With conditions** Share-alike on the model file; legal review.	176 languages including si and ta. Would run on: CPU

Avoid 1

Model	Commercial use	Good to know
Facebook fasttext-language-identification (NLLB LID) Avoid	Not rated	CC BY-NC 4.0.

Tool-calling models 14 models · Not in the platform yet

What it does: Plain-words requests become validated calls on product queries. **Needs:** Fixed tool catalogue, constrained decoding, a test set. **Main risk:** Wrong filters give a confident wrong list.

Model	Commercial use	Good to know
Qwen3-Instruct-2507 (4B, 8B, 30B-A3B) ☆ Start here	✓ **Allowed**	Start here: strong tool calling at 4B. BFCL-v3 61.9 (4B, card). Sinhala and Tamil quality not measured. Would run on: 4B at 4-bit on CPU or 6 GB GPU; 30B-A3B on 24 GB GPU
Gemma 4 (E2B, E4B, 12B, 26B-A4B, 31B)	✓ **Allowed**	Native function calling plus image, video and audio input. Sinhala and Tamil not confirmed. Gemma 3 and FunctionGemma are not Apache. Would run on: E4B on 8 GB GPU at 4-bit; 26B-A4B on 16 to 24 GB
gpt-oss-20b	✓ **Allowed**	Native function calling with adjustable reasoning; good for several calls in a row. English-first; reasoning adds latency. Would run on: 16 GB consumer GPU

Model	Commercial use	Good to know
IBM Granite 4.0 (Micro 3B, H-Small)	Allowed	Enterprise-tuned tool calling in low memory. Sinhala and Tamil not among its main languages. Would run on: Micro on CPU or 6 GB GPU
Mistral Small 3.2 (24B) and Ministral 3 (8B)	Allowed	Native function calling and JSON output. Lists Hindi, Bengali, Nepali, not Sinhala or Tamil. Would run on: 8B on 12 GB GPU; 24B at 4-bit on 24 GB
Microsoft Phi-4-mini-instruct	Allowed	Documented tool-calling format. Expect weak Sinhala and Tamil. Would run on: CPU or 6 GB GPU
Constrained decoding (XGrammar, Outlines, llama.cpp grammars)	Allowed	Forces output into a JSON schema so a 4B model never returns a malformed call. Still validate times, camera ids and ranges in code. Would run on: CPU
SmolLM3-3B and Hermes 4 14B	Allowed	Open-recipe 3B and a 14B tuned for tool use. Fine-tune data licences not read. Would run on: 3B on CPU; 14B on 16 GB GPU at 4-bit
Google FunctionGemma (270M)	With conditions Gated; Google may restrict use.	Tiny model to fine-tune for one fixed tool set. Gemma 4 E2B has Apache terms. Would run on: CPU
NVIDIA Nemotron Nano 9B v2	With conditions Guardrail-bypass termination clause.	Reasoning with tool calling on one GPU. No Sinhala or Tamil. Would run on: 12 to 16 GB GPU at 4-bit
Liquid LFM2 (1.2B-Tool, 8B-A1B)	With conditions Revenue threshold of 10 million US dollars.	Very small and fast. Threshold applies to the licensee and affiliates. Would run on: CPU
Llama-based tool tuners (ToolACE-2-Llama-3.1-8B, MeetKai Functionary small v3.2)	With conditions The Llama licence still applies.	Mature parallel function calling. Qwen 3 and Gemma 4 do the same without it. Would run on: 8B on 12 GB GPU

Avoid 2

Model	Commercial use	Good to know
Katanemo Arch-Function 3B Avoid	Not rated	Research-named licence on a research-licensed Qwen base.
Salesforce xLAM-2 and Hammer 2.1 (7B, 3B) Avoid	Not rated	CC-BY-NC-4.0 or Qwen research licence; top of leaderboards but not for sale.

Omni models 9 models · Not in the platform yet

What it does: One model describes video and sound together. **Needs:** 24 GB GPU or larger; control of the audio path.

Main risk: Audio is the legal problem; no Sinhala or Tamil speech.

Model	Commercial use	Good to know
MiniCPM-o 4.5 (9B) Start here	Allowed	Smallest practical omni model: live video and audio in, text out (11 GB int4, 19 GB bf16, card). Disable or gate the audio path. Would run on: 12 GB GPU at int4; 24 GB at bf16
Qwen3-Omni-30B-A3B (Instruct, Thinking, Captioner)	Allowed	Best open sound-plus-video reasoning; Captioner describes audio. Speech: 19 languages, no Sinhala or Tamil. Would run on: Large GPU (80 GB) at bf16; 24 GB at int4 for short clips
Gemma 4 E2B / E4B / 12B (audio and video input)	Allowed	Takes text, image, video frames and audio. The audio encoder is built for speech; use video only unless audio is cleared against local rules. Would run on: E4B on 8 GB GPU at 4-bit
Qwen2.5-Omni-7B	Allowed	Previous generation; superseded by Qwen3-Omni and MiniCPM-o 4.5. Would run on: 24 GB GPU at bf16
Microsoft Phi-4-multimodal-instruct	Allowed	Compact image, audio and text with tool calling. Images not video; no Sinhala or Tamil. Would run on: 12 GB GPU

Model	Commercial use	Good to know
Ming-flash-omni Preview and Ming-Lite-Omni 1.5	Allowed	Permissive omni with video, audio and image input. Smaller community; base-model data terms not read. Would run on: Large GPU (Preview)
Video-SALMONN 2 and Ola-7B Unconfirmed	With conditions Training-set terms unchecked (LLaVA-Video-178K and others).	Research-grade video with audio at 7B. Tag only. Would run on: 24 GB GPU

Avoid 2

Model	Commercial use	Good to know
NVIDIA OmniVinci Avoid	Not rated	The README says NVIDIA One Way Noncommercial License while the record says apache-2.0; the tag contradicts the card.
Baichuan-Omni-1.5 (7B) Avoid	Not rated	Community licence excludes software and cloud service providers; the record says apache-2.0 but the card says otherwise.

Assistant safety 12 models · Not in the platform yet

What it does: Screens what the assistant is asked, answers and reads. **Needs:** CPU or small GPU; a written site policy.

Main risk: False blocks; most injection filters are English only.

Model	Commercial use	Good to know
Qwen3Guard (Gen 0.6B/4B/8B, Stream 4B) Start here	Allowed	Three severity levels, 119 languages claimed. Best start for Sinhala, Tamil and English; accuracy in them unmeasured. Would run on: 0.6B on CPU or small GPU; 8B on a 12 GB+ GPU
ProtectAI deberta-v3-base-prompt-injection-v2 (and small v2)	Allowed	Fast detector of injected instructions in text the assistant reads. English only: a Sinhala or Tamil injection passes. Would run on: CPU
gpt-oss-safeguard (20B, 120B)	Allowed	Reads your own written policy at run time. Slower than a small classifier; competes with the main assistant for GPU memory. Would run on: 16 GB consumer GPU (20B)
IBM Granite Guardian (3.3 8B, 3.2 5B, 3.2 3B-A800M, HAP 38M)	Allowed	Jailbreak, groundedness and tool-call hallucination checks. English only. Would run on: HAP on CPU; 8B on a 12 GB+ GPU
PIGuard (leolee99)	Allowed	Tuned to reduce over-defence, useful when alert text has words like "stop". English only. Would run on: CPU
Guardrail frameworks (NeMo Guardrails, LLM Guard, Guardrails AI)	Allowed	Plumbing for input and output checks and tool allow-lists. A few hundred lines of own code may be enough. Would run on: CPU
Meta Llama Guard 4 (12B) and Llama Guard 3 (8B, 1B)	With conditions Gated; Llama licence (700 million users clause, attribution).	Published hazard taxonomy. English only; no Sinhala or Tamil. Qwen3Guard does the same job. Would run on: 1B on CPU; 8B on 12 GB GPU; 12B on 16 GB+ GPU
Google ShieldGemma (2B, 9B) and ShieldGemma 2 (4B, images)	With conditions Gated; Google may restrict usage.	Text and image safety classification. English only (secondary). Would run on: 2B on CPU or small GPU; 9B on 12 GB+ GPU
NVIDIA Nemotron Safety Guard 8B v3 and Content Safety Reasoning 4B	With conditions Stacks NVIDIA and Llama licences; guardrail-bypass termination.	Multilingual (nine languages) but no Sinhala or Tamil. Would run on: 8B on a 16 GB GPU
AI2 WildGuard (7B)	Allowed	Prompt harm, response harm and refusal detection. English; older. Dataset terms not read. Would run on: 12 GB+ consumer GPU at 4-bit
Meta Llama Prompt Guard 2 (86M, 22M)	With conditions Llama terms (700M MAU, attribution); gated.	Small, fast injection classifier. No Sinhala or Tamil; ProtectAI and PIGuard cover the same ground. Would run on: CPU

Model	Commercial use	Good to know
Small injection classifiers (deepset deberta-v3-base-injection, fmops distilbert, jackhhao jailbreak-classifier)	Allowed	Baselines only; small English sets from 2023 to 2024. Would run on: CPU

Data, privacy and forecasts

Personal-data detection and redaction 10 models · Not in the platform yet

What it does: Masks faces, plates, screens and personal text before data leaves. **Needs:** Face and plate detector, masks, Presidio, review queue. **Main risk:** Missed detections; no Sinhala or Tamil name models.

Model	Commercial use	Good to know
Microsoft Presidio (analyzer, anonymizer, image redactor) Start here	Allowed	Framework for text masking with pattern and model recognisers. English out of the box; no Sinhala or Tamil names. Would run on: CPU
EgoBlur Gen 2 (face and licence plate) Unconfirmed	With conditions Read the weights download agreement first.	Detects faces and plates for blurring (about 21 ms on GPU for 2 MP). Recall on tiny faces in wide CCTV views not published. Would run on: Consumer GPU live; CPU offline
OpenAI Privacy Filter	Allowed	Token classifier for 8 privacy categories. Weak on Sinhala and Tamil; fixed label set. Would run on: CPU or small GPU
GLiNER-PII (Knowledgator base and large) and gliner_multi_pii-v1	Allowed	Zero-shot entity finder. F1 83.25 percent (large) on a synthetic set, not site logs. No Sinhala or Tamil claim. Would run on: CPU
SAM 2.1 (Hiera tiny to large)	With conditions Trained on SA-1B, whose terms nobody read.	Pixel masks of a person or object tracked through a clip. Needs a prompt box. Would run on: Consumer GPU
docTR (text detection and OCR) Unconfirmed	With conditions Weights licence and data unread.	Finds text on monitors and whiteboards to blank. Sinhala and Tamil not covered. Would run on: CPU or consumer GPU
NVIDIA GLiNER-PII	With conditions NVIDIA licence text not read.	55+ PII and health categories. English only; the Apache GLiNER models do the same job. Would run on: CPU or small GPU
Meta SAM 3	With conditions Custom licence with trade-control terms.	Text-prompted masks ("hard hat", "screen"). Would run on: Consumer GPU
deface (CenterFace)	With conditions Trained on WIDER FACE, which is CC BY-NC-ND.	Ready-made command-line face blurring; EgoBlur is cleaner. Would run on: CPU

Avoid 1

Model	Commercial use	Good to know
Piiranha, ai4privacy and Isotonic PII models Avoid	Not rated	CC-BY-NC-ND-4.0 and CC-BY-NC-4.0 weights; popular but not usable in a sold product.

Privacy masking and redaction 5 models · Not in the platform yet

What it does: Covers faces, screens and bystanders in exports. **Needs:** Detectors, solid fill, human review, audit entry. **Main risk:** A missed face leaks identity; blur can be undone.

Model	Commercial use	Good to know
RF-DETR Seg for person and vehicle masks ☆ Start here	✓ Allowed	Covers everyone except the person of interest with a solid shape. Misses partly hidden people. Would run on: Consumer GPU
SAM 2.1 Tiny or Small (mask tracking through a clip)	✓ Allowed	Click once on a screen or person and carry the mask through the clip (about 85 to 91 FPS per README). A reviewer checks the whole clip. Would run on: Consumer GPU
PaddleOCR PP-OCRv5 text detection	✓ Allowed	Finds text on screens, papers and badges to cover. Also finds signs that should stay visible. Would run on: CPU or consumer GPU
OpenCV opaque fill, pixelate and blur	✓ Allowed	Applies the mask. Use solid fill; it only applies what detectors found. Would run on: CPU
BiRefNet (and lite)	✓ Allowed	Clean foreground cut-out to keep one person and hide the rest. Built for objects, not crowds. Would run on: Consumer GPU

Synthetic data for training 10 models · Not in the platform yet

What it does: Rare PPE and hazard pictures to widen training sets. **Needs:** 16 to 24 GB GPU, mask model, human review. **Main risk:** Wrong PPE geometry; generator provenance.

Model	Commercial use	Good to know
FLUX.2 [klein] 4B ☆ Start here	✓ Allowed	Text-to-image and editing in one model. Only the 4B is Apache; FLUX.2 [dev] is non-commercial. Would run on: 16 GB consumer GPU (about 13 GB VRAM)
Copy-paste compositing with SAM 2.1 masks (Albumentations)	⚠ With conditions The mask model inherits the SA-1B question; masks can be drawn by hand.	Cut real PPE from site frames and paste into others; labels stay exact. Do not paste a person's face into other scenes. Would run on: CPU for compositing; GPU for masks
FLUX.1 [schnell]	✓ Allowed	Four-step text-to-image for scene variety. No editing; dev, Kontext and Fill are non-commercial. Would run on: 24 GB GPU, less when quantised
Qwen-Image (20B) and Qwen-Image-Edit-2511	✓ Allowed	Instruction editing ("put a hard hat on the second worker"). Run offline on a workstation. Would run on: Large GPU; 24 GB quantised
Z-Image-Turbo (6B)	✓ Allowed	Fast text-to-image in 8 steps. Editing is a separate variant. Would run on: 16 GB consumer GPU
OmniGen2, BAGEL-7B-MoT, Step1X-Edit	✓ Allowed	Unified generation and editing; alternatives to Qwen-Image-Edit. Base-model data chain not traced. Would run on: 24 GB GPU
HiDream-I1 and Sana 1.6B	⚠ With conditions Sana ships a Gemma text encoder under Gemma terms.	HiDream is high quality; Sana is very fast. Would run on: Sana on 8 GB; HiDream on a large GPU
Stable Diffusion XL and 3.5, with ControlNet	⚠ With conditions Use-restriction clauses pass to customers; SD 3.5 free below 1 million dollars revenue.	Largest ecosystem: depth or pose-controlled scenes and mask inpainting. Older than FLUX.2 klein. Would run on: 12 GB consumer GPU (SDXL)
NVIDIA Cosmos Transfer 2.5 and Predict 2.5 (2B)	⚠ With conditions Read the full agreement.	Turns a rendered scene into photoreal video with exact labels. Needs a 3D scene source. Would run on: Large GPU

Avoid 1

Model	Commercial use	Good to know
FLUX.1 dev / Kontext / Fill, FLUX.2 dev, SDXL-Turbo, BRIA FIBO 👎 Avoid	Not rated	Non-commercial licences; outputs may not be cleared for training a product either.

Image and video generation or editing 6 models · Not in the platform yet

What it does: Clean a floor-plan photo for drawing cameras. **Needs:** Often no model; a mask and inpainting. **Main risk:** Altered pictures passed as evidence; invented pixels.

Model	Commercial use	Good to know
FLUX.2 [klein] 4B (editing mode) Start here	Allowed	Instruction and multi-reference editing to clean or restyle a plan. Show a before and after. Would run on: 16 GB consumer GPU
Qwen-Image-Edit-2511	Allowed	High-quality edits with fewer unwanted changes. Too heavy for a site server. Would run on: Large GPU
LaMa (big-lama)	With conditions Trained on Places2 (non-commercial research terms); check the terms or retrain.	Paint a mask and LaMa fills it; small and fast. Would run on: CPU
Wan 2.2 (TI2V-5B, I2V-A14B) and Wan 2.1 VACE	Allowed	Text-to-video, image-to-video and video editing for offline demos only. Never on real footage in the product. Would run on: 24 GB GPU (5B); large GPU (A14B)
Other video generators (LTX-2, HunyuanVideo, CogVideoX-5B, Stable Video Diffusion, Mochi 1, Open-Sora v2)	With conditions Licences differ; LTX-2 paid above 10 million dollars; Hunyuan excludes EU, UK, South Korea.	Not needed by the product. Would run on: Large GPU

Avoid 1

Model	Commercial use	Good to know
ProPainter and DiffuEraser (video inpainting) Avoid	Not rated	ProPainter is S-Lab non-commercial and DiffuEraser needs its weights; altering video of real people also carries the evidence risk.

Evidence integrity and AI-image labelling 2 models · Not in the platform yet

What it does: Signed exports that show who exported what and when. **Needs:** A signing certificate and key storage. **Main risk:** A signature proves no change, not a real scene.

Model	Commercial use	Good to know
c2pa-rs (Content Credentials library) Start here	Allowed	Signs exports with who, when, camera and redaction state, and verifies later. Outside parties trust only C2PA-listed certificates. Would run on: CPU

Avoid 1

Model	Commercial use	Good to know
AI-generated-image detectors (Ateeqq, haywoodsloan, umm-maybe, Organika sdxl-detector) Avoid	Not rated	Unreliable on camera pictures, and accusing a person of a faked image on a detector's score is unsafe; Organika is CC-BY-NC-3.0.

Forecasting and anomaly detection 11 models · Not in the platform yet

What it does: Forecasts footfall and queues; flags an odd hour. **Needs:** Weeks of counts, a site calendar. **Main risk:** Planned changes cause false alarms.

Model	Commercial use	Good to know
Chronos-2 (and Chronos-Bolt) Start here	Allowed	Footfall and queue forecasts with day of week and shift plan as covariates; quantile ranges. Would run on: CPU (Bolt, small); consumer GPU

Model	Commercial use	Good to know
statsforecast, Prophet, darts and sktime Unconfirmed	Allowed	Seasonal-naive, ETS and ARIMA baselines; often as good as foundation models on a stable site. Licences not re-read. Would run on: CPU
TimesFM 2.5 (200M)	Allowed	Zero-shot forecasts with long context. Training-data terms not read. Would run on: CPU or consumer GPU
Granite TinyTimeMixer (TTM-r2)	Allowed	Right size for a CPU-only site; fine-tunes on a few percent of the data. Would run on: CPU
PyOD, scikit-learn Isolation Forest, STUMPY, river Unconfirmed	Allowed	Flags an unusual hour, odd stretches or drift on counts and camera health. Tune per site. Would run on: CPU
Toto (Datadog)	Allowed	Built for server metrics; useful for the product's own health numbers. Would run on: Consumer GPU
Sundial and Timer (Tsinghua)	Allowed	Alternative zero-shot forecasters. LOTSA has mixed terms. Would run on: CPU or consumer GPU
MOMENT-1 and Time-MoE	Allowed	MOMENT also does anomaly detection. Time-MoE licence not confirmed. Would run on: Consumer GPU
Lag-Llama	Allowed	Earlier probabilistic forecaster; behind Chronos-2 and TimesFM. Would run on: CPU
TiRex (NXAI)	With conditions Commercial licence needed above 100 million EUR revenue.	Strong forecaster; the revenue condition scales with the owner and customers. Would run on: Consumer GPU

Avoid 1

Model	Commercial use	Good to know
Moirai 1.1, 2.0 and MoE (Salesforce) Avoid	Not rated	CC BY-NC 4.0 weights.

Risk scoring on tables 6 models · Not in the platform yet

What it does: Ranks zones and shifts by where problems cluster. **Needs:** Event table, months of history, per-score reason. **Main risk:** Never per person; small groups single out individuals.

Model	Commercial use	Good to know
LightGBM, XGBoost and CatBoost Start here Unconfirmed	Allowed	Strong default for counts per zone and shift. Needs an explanation layer such as SHAP. Places and shifts only. Would run on: CPU
Explainable Boosting Machine (InterpretML) and logistic regression (scikit-learn) Unconfirmed	Allowed	Glass-box models whose effects can be shown to an auditor; easier to defend. Would run on: CPU
SHAP Unconfirmed	Allowed	Explains each score from a boosted-tree model. Would run on: CPU
TabPFN v2 (Prior Labs)	With conditions Enhanced attribution requirement.	Pretrained tabular model for small tables with no training. The 2.5 release is non-commercial. Would run on: Consumer GPU for larger tables
TabICL Unconfirmed	With conditions Official weights licence not confirmed.	Pretrained tabular classifier; newer, less proven. Would run on: Consumer GPU

Avoid 1

Model	Commercial use	Good to know
TabPFN-2.5 Avoid	Not rated	TabPFN-2.5 Non-Commercial License; no production use.

Words used here

Model	A trained program that looks at a picture, a sound or a text and answers a question about it.
Detector	A model that draws a box around each thing it finds: a person, a helmet.
Job	One thing the platform does with a model: finding people, checking PPE, reading a number plate.
Size	How big the model is. Bigger is more careful and slower.
Outline	The exact shape of an object, following its edge, not a box around it. Click to outline draws a zone this way from one click.
Rotated box	A box turned to hug an object that sits at an angle.
Body points	The joints of a person, as points, for checks such as a person who has fallen.
Whole-picture label	One label for the whole picture, with no box around anything.
Fit	Whether a size runs on the processor (CPU), runs faster on a graphics card, or needs one. A size that needs a GPU warns you on a server without one.
Heavy model	A model that takes much of the server's computing power. It starts switched off until you turn it on.
Type what to find	For a model that takes typed words, the names of the things to look for. The platform makes a copy that looks only for those names.
GPU	A graphics card. It runs models many times faster than a normal processor (CPU).
Licence	The rules its makers set for using the model.
Commercial use	Using the model in a business. Some licences allow it, some allow it with conditions, some do not.
Loader	The part of the platform that knows how to run one kind of model. A model whose kind has no loader yet cannot run.
Class	One kind of thing a model can find, such as "helmet" or "no helmet".
Class map	The table that tells the platform which of a model's classes mean which PPE item. The platform suggests one when a model is installed.
Activate	To make an installed model the one that runs for its job.
Test bench	The page where you check photos and stored videos with a model before it runs on live cameras.
Training	Teaching a model with labelled pictures. Fine-tuning is training a ready-made model a little more on your own pictures.

AI CCTV Monitoring Platform

A guide to licences. Licences change: read a model's licence again before you rely on it in a business.

Document	Software	Code owner	Author
Model guide Updated 8 November 2026	Version 1.1.804 Code name Sentinel	Altavision altavision.lk	Deeghayu Adhikari deeghayu.netlify.app

© 2026 Altavision. All rights reserved.

The list was researched on 2026-10-13