Visionplatform.ai user manual

User management

Creating a Workspace

- When you create a new account, you will be asked to set up a workspace.
- A workspace is a shared space where team members can work together on the same project.

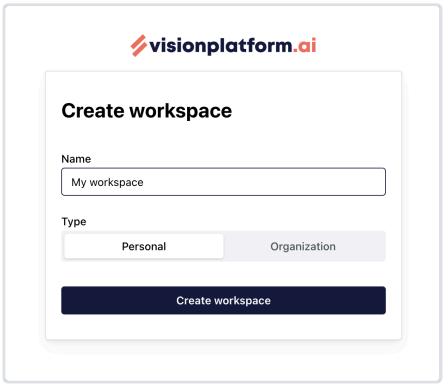

Inviting Members to a Workspace

- After your workspace is set up, you can invite others to join.
- To invite members:
  - Click **Workspaces** in the **User Menu** at the top-right corner.
  - On the Workspaces screen, click the **Invite Members** button.
  - Enter the email addresses of the people you want to invite and assign them a role.
  - Click **Invite Members**. They will get an email with a link to join the workspace.

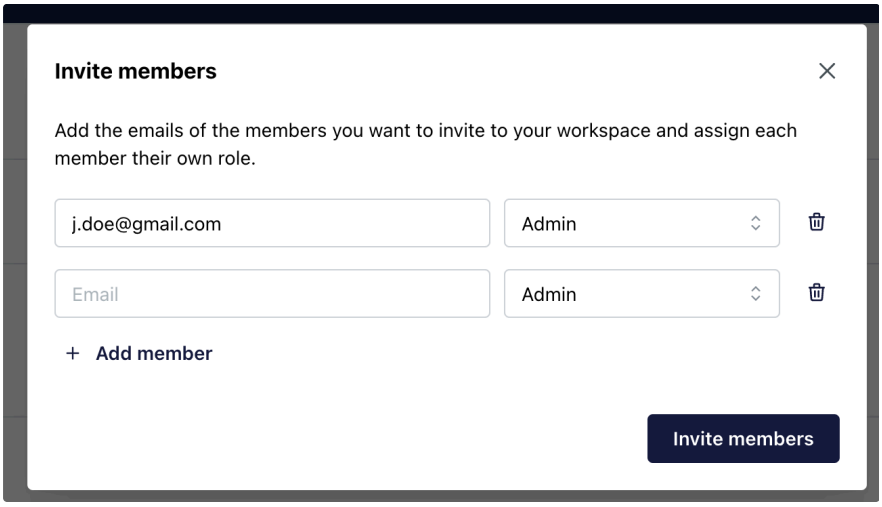

Managing Members

- **Roles**:
  - User
    - Can create applications, can train and deploy ML models.
  - Admin
    - Has all User permissions + can manage workspace users (invite, assign roles, remove members).
- **Deactivate Members**:
  - To remove someone's access, click the **Deactivate (X)** button in the Members Table.
- **Reactivate Members**:
  - To give someone access again, click **Re-send Invite**. They will get a new invitation email.

Note for Invited Members

- Invited members will need to select a **Free License** after accepting the invitation to start using the workspace.
- A **Shared Workspace License** option will be available in future updates for teams.

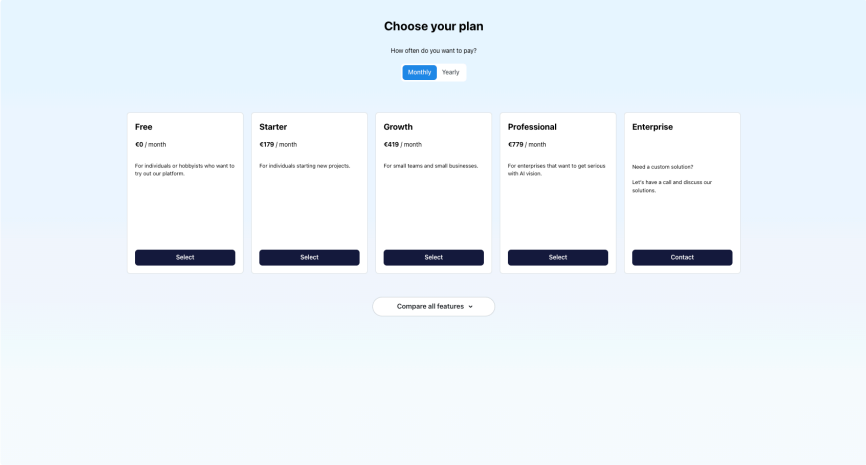

AI Assistant

This whole manual and many more documents are synced in the memory of our AI Assistant. Ask it anything you need. It's best in English but can also speak your local language.
If it answers ''Sorry, I cannot find the relevant information, could you please clarify your question'' then rest assure we will update the memory of our AI assistant.
Try clarifying the question once more with more context in case it doesn't know the answer also helps a lot!

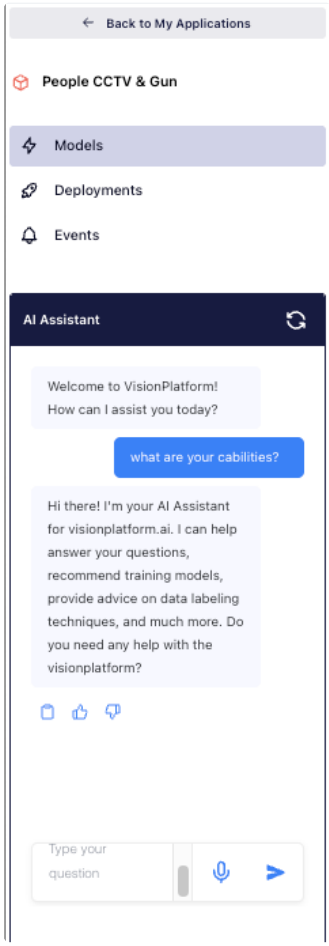

Create an application

Press ''create application'' on the home screen.

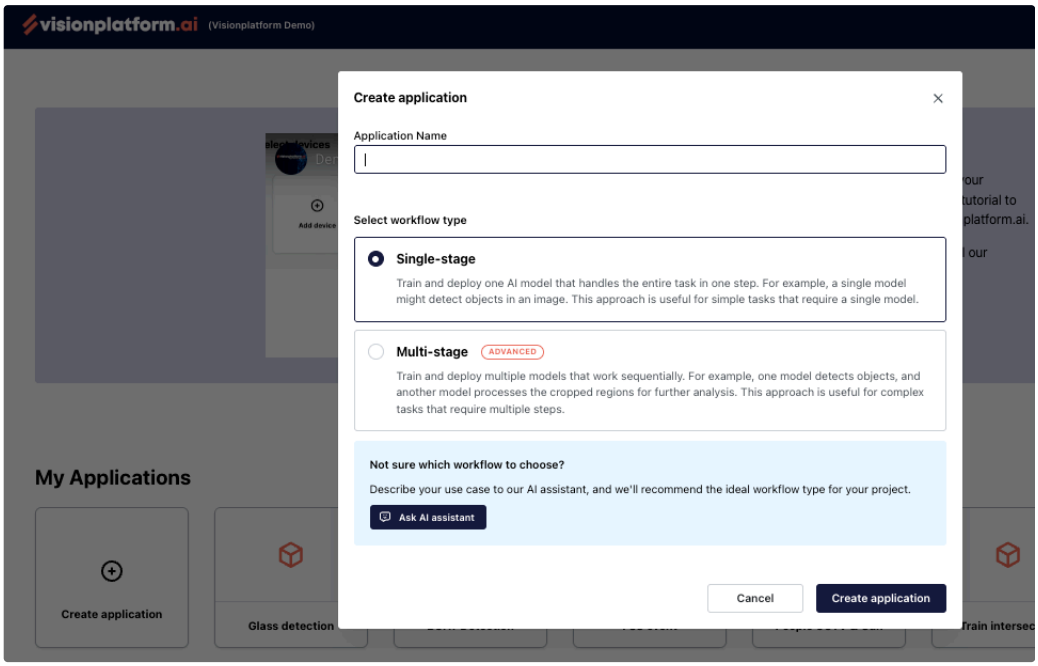

Here, you get the choice to select Single-stage or Multi-stage.

Single-stage

Used if you only want to detect objects without doing extra analyses on the object itself

---

Multi-stage

In case you want to see objects within objects. We crop the 1st object and send this crop to a 2nd algorithm for the next detection.

For example.

- I want to detect a human in a 1st stage and in the 2nd stage we are looking for a helmet or safety vest.
- I want to detect a car in a 1st stage and in the 2nd stage we are looking for the license plate that we then OCR into text.

Models

An AI computer vision **model** is like a smart brain that learns to recognize patterns in your pictures and videos. On visionplatform.ai , the main type of model currently available is **Object Detection**, with plans to add more types like **Segmentation** and **Anomaly Detection** soon.

Types of Models

1. **Object Detection** *(Available Now)*
   - This model learns to find and recognize objects in your pictures or videos, like people, cars, or boxes.
   - **When to Use It:**
     - When you need to track specific objects or people.
     - Example: Detecting cars in a parking lot or spotting people in a restricted area.
   - **When Not to Use It:**
     - If you need to understand very detailed shapes or specific parts of objects (this is where segmentation might be better).
2. **Segmentation** *(Coming Soon)*
   - This model will divide images into detailed sections, like identifying the exact shape of objects (e.g., separating a car from its background).
   - **When to Use It:**
     - When you need precise shapes of objects, like identifying the size of damaged areas on a product.
   - **When Not to Use It:**
     - If you only need to find and count objects, object detection is simpler and faster.
3. **Anomaly Detection** *(Coming Soon)*
   - This model will look for unusual patterns or events, like things that don't normally happen.
   - **When to Use It:**
     - For spotting problems, like damages or missing items. It's typically only used in quality inspection.
     - If you have little data of quality inspection problems. Anomaly detection learns from the good examples to identify what is an anomaly.
   - **When Not to Use It:**
     - If you're only interested in identifying specific objects or tracking known patterns.

How to Create a Model

1. Start by pressing the **'+'** button at **'Train or Choose Model'** to create your first model.
2. If you see the system missing objects or making mistakes (false positives), you can improve it:
   - Add more pictures or videos of those cases to your dataset.
   - Train a new model with the updated data.

By doing this, your model becomes smarter and works better over time!

We suggest giving it a good name with a V1, V2 etc behind it so you keep it neat and easy to manage.
If you give it the same name it will say it's already in use.

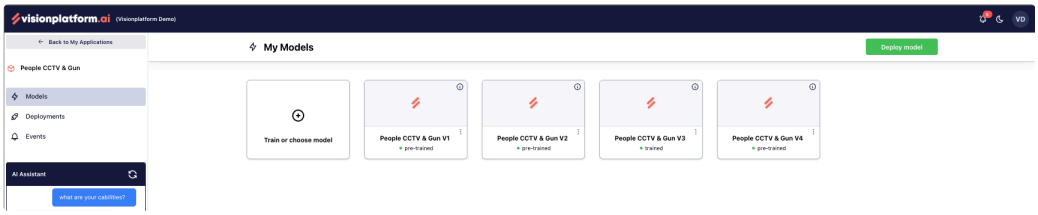

Train or choose your model

After creating your first version on the previous screen, you have three options:

1. **Start from scratch** – Create a completely new model yourself.
2. **Use a ready-made model** – Pick and deploy a model we already trained for you by us
3. **Improve an existing model** – Take one of our pre-trained models and customize it for your needs.

If you're unsure, we recommend asking the chatbot. It's like starting with a recipe instead of inventing a new dish—you'll save time and effort, but you have to cook with the right ingredients.

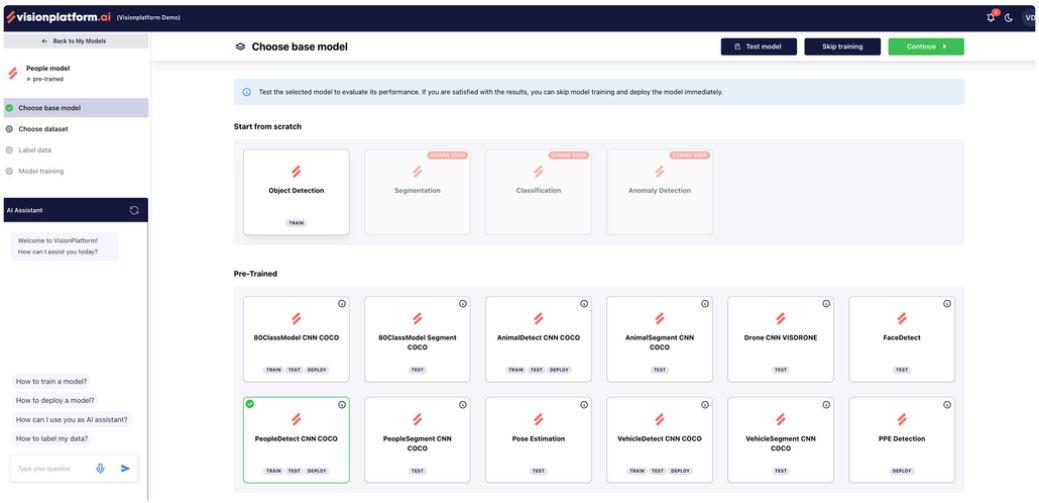

From Scratch

This option means starting fresh to create your own model from the ground up. It's like drawing on a blank page—you'll need to teach the model everything step by step with new labels. For example, you could train the model to detect damages on your product or identify specific components unique to your use case. In case you have a label class that's already in one of visionplatform.ai pre-trained models then always go for the ''build on top'' option explained in the below chapter.

Pick and Deploy

With this option, you can use a model that's already trained and ready to go. It's like picking a ready-made tool and putting it straight to work. No extra training is required. Just select the model and press the 'Skip Training' button to deploy it.

Build on Top of a Pre-Trained Model

This option allows you to start with a model we've already trained and customize it to suit your needs. It's like taking a basic recipe and adding your own ingredients to make it perfect.

For **example**:

- **People and Firefighter Detection**: While we don't have a specific firefighter class in our pre-trained models, we do have people detection. You can select the people detection model and add firefighter-specific data to it. You then improve the people class but can add a separate class for firefigthers.

Choose dataset

Press the button ‘‘Add dataset’’ to create your first own dataset.
Give your dataset a clear name, select object detection (we don’t support other labeling types yet) and pick your labeling classes.

If you selected a pre-trained model from visionplatform.ai you will see all the labeling classes from that model. Pick the ones you need and then type your own labeling classes.

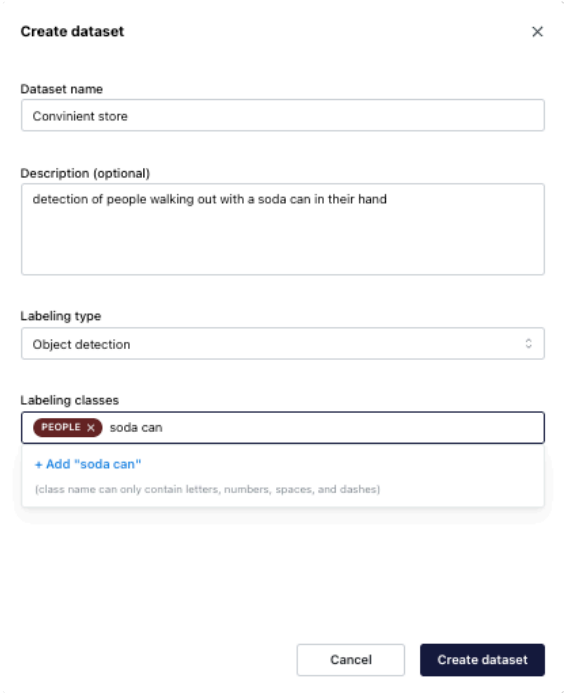

Label data

To get the best results, you need to label your pictures and videos. This helps the system learn better. Start by adding your pictures or videos to a new dataset.

The system learns more if you use a variety of data. You can also find pictures or videos online, as long as they look similar to what your camera sees (same angles, object sizes, etc.).

Tips for Videos
Use short videos, about 10 seconds long. It’s better to have many short videos than a few long ones.

How Much Data Do You Need?

- If your pictures or videos are simple (same background and objects), start with 30–50 items.
- If your pictures or videos are more complex (different objects and backgrounds), you’ll need at least 100 items.

The amount of data you’ll need depends on several factors, like:

- How complicated the detections are.
- The size of the objects in the images.
- The quality and sharpness of the video.
- The accuracy percentage you want to achieve.

Some models can perform well with **100 small videos** and achieve 98% accuracy, while others may require **1,500 small videos** to reach the same level.

Why Use Videos Over Pictures?
Videos are made up of many frames, which means the system can analyze much more data faster. This often makes detections quicker and more efficient compared to single pictures.

Training the System
You need at least the following amount to hit the training button:

- 20 labeled pictures, or
- 5 labeled videos.

Tips for Labeling Data

1. **Label Important Objects**
   - If your model needs to find specific things (like people), label all of them in each picture.
   - If you only want certain areas (like close objects), label just those. Be consistent so the system doesn't get confused.
     Example: If you only want to detect cars in the nearest lane, label only those cars, not distant ones.
2. **Use Relevant Pictures**
   - Add pictures that look like what your camera will see during real use. This helps the system learn better.
3. **Show Different Situations**
   - Train with a mix of examples, like:

- Day and night views.
- Different weather (rain, sun, clouds).
- Few objects vs. crowded scenes.

Include Empty Pictures or videos

Adding empty pictures or videos to your dataset can improve your model. These are images or videos with no objects of interest in them. Save them without adding any labels.

Why Add Empty Pictures or Videos?

They teach the system what not to detect. For example:

- Raindrops on the camera lens.
- Spider webs near the camera.
- Reflections or shadows that might confuse the model.
- Background objects you don't want detected, like trucks in the distance if you're only focusing on cars in the foreground.

Fixing False Detections

If the system makes mistakes (false detections), adding empty pictures or videos with those tricky situations can help the model learn to ignore them. For example:

- If the model mistakenly detects a **shadow** as an object, include an empty picture or video with the shadow to teach it not to detect it.
- If the model detects a **spider web** as a person, include an empty picture or video showing only the spider web.
- If you want to detect **cars in the foreground**, add empty pictures or videos with only trucks driving in the background.

How to Add Empty Pictures or Videos

1. Upload the data (pictures or videos) to your dataset.
2. Open the file to label it.
3. Press the Save button without adding any labels.
4. The picture or video will then get a "labeled" status.

Using these examples and steps, the model can “unlearn” its mistakes and become more accurate over time.

How to Label

Labeling Pictures or Videos

Select a picture or video file.

Press **‘Label File’** in the top-right corner.

- If it’s a **picture**, the labeling tool will open automatically.
- If it’s a **video**, you’ll be prompted to choose the **Frames Per Second (FPS)** before the labeling tool opens.

Why to consider to reduce the Frames per Second (FPS)?

- If your video was recorded with a high FPS and the object moves slowly, many frames will look very similar. Reducing the FPS saves time and effort when labeling without losing important data.
- When you press Label on a video, the system will prompt you to reduce the FPS if needed.

Recommended FPS for Labeling

Our platform will ask you to what FPS you want to convert your model when you want to label a video.

- 2-5 frames per second works well for most scenarios.
- Use higher FPS only if the object moves very quickly.

How to label correctly

1. **Draw Tight Boxes**
   - Make the box fit the object closely. Don’t leave extra empty space around it.
   - Make sure the whole object is labeled.

Wrong:

The person with the gun is not labeled correctly

The whole person with the gun is labeled and that’s correct.

2. **Use Rectangles Only**
   - All labels should be rectangular.

Videos vs images

Using videos for labeling is an excellent way to make your model more robust. The movement of objects provides a lot of learning opportunities for the model. Variations in visibility, reflections, shapes, and sizes help the model understand and detect objects faster and more accurately.

Best Practices for Video Uploads

Short Videos:

- Keep your videos short, ideally no longer than 10 seconds.
- Uploading multiple short videos with different variations is better than a few long ones.

Focus on Variations:

- Aim to include diverse scenarios in your videos, such as:
- Different lighting conditions
- Changing angles or reflections
- Objects appearing in various sizes or shapes

Train model

Now that you've labeled enough data, it's time to train your first model!

How to Train

Click the green "Train Model" button in the top-right corner.

All your labels are automatically selected, and we've pre-configured the settings for you.

You need to make decisions on the following:

- Full Training or Fine Tuning
- Resolution

Choosing Between Full Training and Fine Tuning

When training your model, you can choose between Full Training and Fine Tuning. Here's what they mean and when to use each:

Full Training

This option trains a model from scratch, starting with no prior knowledge.

- **Use Full Training if:**
  - You're starting without a pre-trained model.
  - You're using a pre-trained model from visionplatform.ai but have added **many new object types (classes)** that weren't in the original model.

Fine Tuning

This option uses a pre-trained model as a starting point and updates it with your data.

- **Benefits of Fine Tuning:**
  - Faster training because the model already knows basic object patterns.
  - Better performance with smaller datasets, as it builds on existing knowledge.
- **Use Fine Tuning if:**
  - You're working with a pre-trained model from visionplatform.ai and are using the **default classes** (e.g., people, cars, etc.) without adding new ones.
    - For best results, set **Freeze Layers** to 10 to keep the pre-trained model's knowledge while training on your data.
  - You've added **a few new object classes** and want to see how well the model can adapt to them.

Which Option Should I Choose?

- If you're using the visionplatform.ai **pre-trained model with default classes** (no custom classes), Fine Tuning is the best choice.
- If you've added **new classes**, you may need to experiment with both Fine Tuning and Full Training to see which works better.

Pro Tip: Stay Consistent

When naming your object labels, use the same naming style throughout your dataset for better results.

Choosing the Right Resolution

Default Resolution (640x640):

- Works well for most cases and is the standard option. It's also best of both worlds for performance versus hardware stress.

Higher Resolution (1024x1024):

- Choose this if your use case involves small objects or if you're running only a few cameras on a Jetson device.
- Note: Higher resolution improves accuracy but will make the model slower and increase hardware stress.

Adjust for Higher Resolution:

- If you choose 1024x1024, change the Batch Size to 32 for optimal performance.

Now that your model is training, you can sit back and wait until you receive an email notification that the process is complete.

What to Do While You Wait?

- Gather More Data:
- Collect a few real-life examples of your use case that weren't included in your labeled dataset. These should be good-quality videos or images that represent the scenarios your model will handle.

Testing Your Model

Once your model has finished training you will receive an email confirmation. Go to the platform and:

Press “Test Model”.

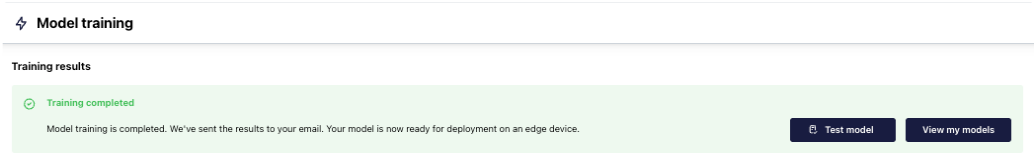

Upload 1-3 videos or pictures from the data you just gathered.

- Keep videos short (15-20 seconds max) to ensure faster testing.
- Pictures have to be JPEG. If you have PNG or another format change the name of the file to .jpg at the end

Wait for the test to complete and review the results.

Evaluating Performance

When testing your model, pay attention to how well it detects objects. Here’s what to look for and how to improve performance:

No Detections or Blinking Detections

What It Means:

- The model is missing objects completely or detecting them inconsistently.

How to Improve:

- Add more examples of this type of data to your dataset.
  - Focus on similar scenarios (e.g., objects in different lighting conditions, angles, or positions).
  - Create a new version of your dataset that includes this additional data.
  - Train a new version of your model (**V2 model**) using the updated dataset.

Low Confidence Detections

What It Means:

- Confidence is how sure the model is about a detection. It ranges from 0.1 (low confidence) to 0.99 (high confidence).
- A confidence score below **0.7** is considered low for production-ready models. Scores above **0.8** usually don’t need improvement.

How to Improve:

- Add more examples of the low-confidence cases to your dataset.
  - Focus on scenarios where the model struggles (e.g., blurry objects, similar-looking objects).
  - Update your dataset version with these new examples.
  - Train a new version of your model to improve accuracy.

Steps to Improve Model Performance

Analyze Test Results:

- Identify where the model struggles (e.g., no detections, low confidence).

Add More Data:

- Add examples of the missing or low-confidence scenarios to your dataset.

Create a New Dataset Version:

- Update your dataset with the new data.

Train a New Model Version:

- Train your updated dataset as a new model (e.g., V2, V3).

By analyzing test results and making these improvements, you can create a smarter, more accurate model with each version! Don't forget to name them properly and consistent 🙂

Deployment

Congratulations, you trained your first AI Computer vision model that is ready to deploy!

Select model

Go to the deployment screen and select the model you trained

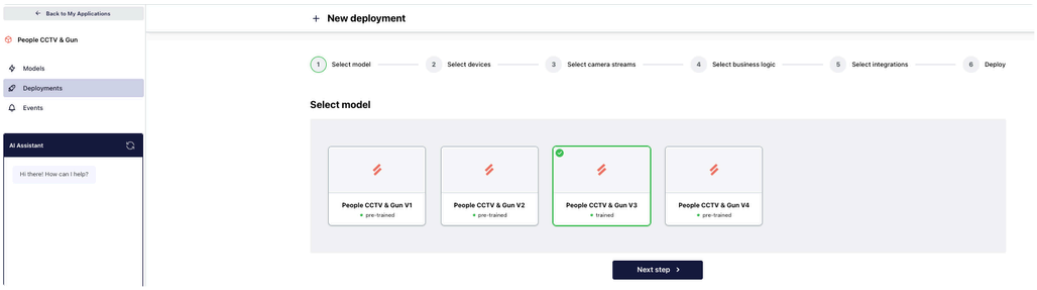

Select devices

Select the device you want to deploy on. You can select 1 or go for mass deployment and select multiple at the same time.

If you haven't set up your device yet press the ''Add device'' button and follow the steps.

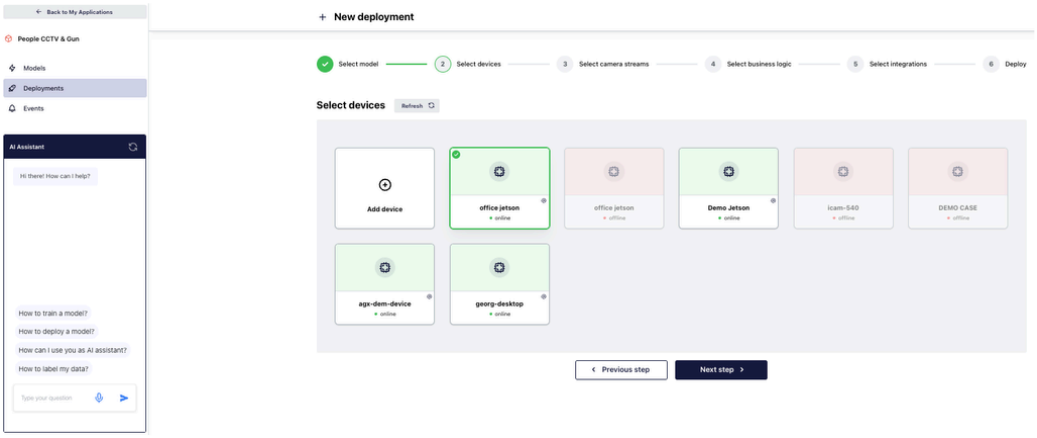

Troubleshootig or updating edge client

On each device there is a small gear icon.

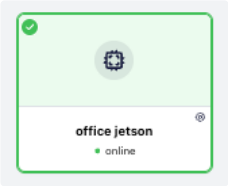

Press the gear icon to see the device client version. You can remotely update your client and or restart the client remotely by pressing the buttons.

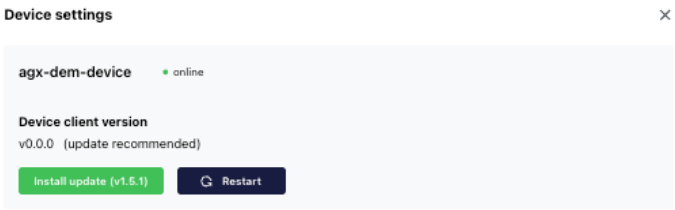

Select camera streams

Manually Add Camera Streams

1. Press the **'Add Camera Stream'** button.
2. A pop-up will appear asking for the stream name. Choose a clear and consistent name.
3. Select the device linked to your camera. Make sure this device is on the same network as the camera or can access it using the IP address.
4. Enter the RTSP URL or USB device path.
5. Press **'Add Camera Stream'** to finish.

Import Camera Streams with Milestone Systems XProtect

1. Enter your Milestone AI Bridge URL and press **Next**.
2. Visionplatform.ai will connect to your system and automatically import all cameras.

For help setting up the Milestone AI Bridge, please contact us for a separate manual.

Troubleshooting

- To check if your RTSP URL works, test it on your laptop:
  a. Download and open **VLC Player**.
  b. Go to **‘File’** > **‘Open Network’**.
  c. Enter the RTSP URL and press **Open**.
- If the stream works in VLC, the RTSP link is correct.

Select business logic

Detect Objects

Your AI model is ready to work! It will now detect objects and follow the rules you set to send you data or alerts.

Add Line Crossing

1. Go to the **Line Crossing** tool and click **‘Add Line Crossing’**.
2. Draw a line where you want to monitor activity. For example, across a door or road. Important: Make sure the line is bigger then you need! Objects need to move through a complete line
3. Save the line, and the system will track when objects cross it.
4. The direction you add have to be in the middle through the line. The direction will determine if it will trigger an event or not.
5. Set alerts or actions for when the line is crossed (e.g., send a notification).

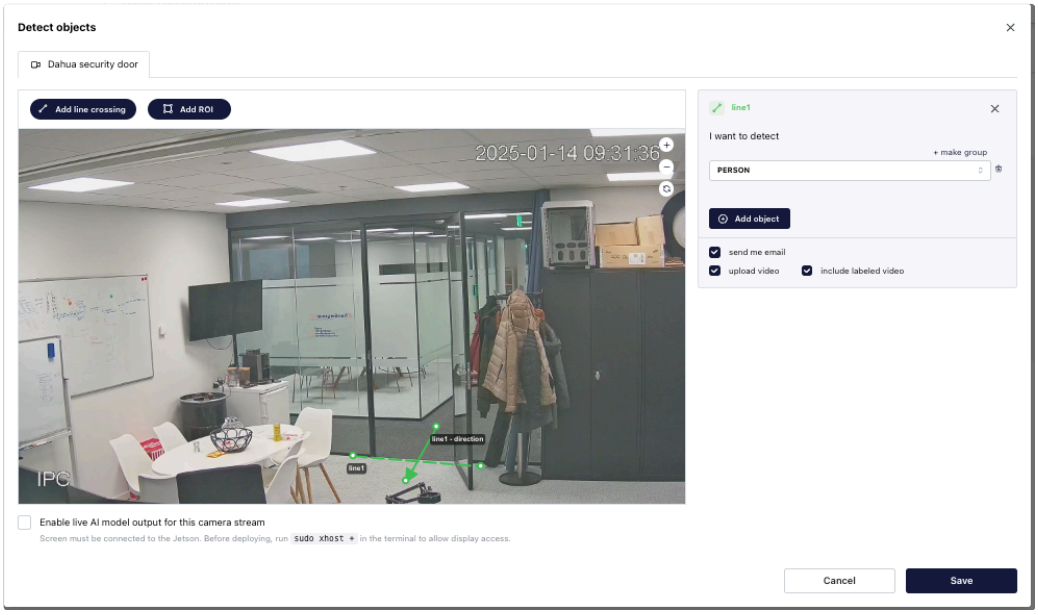

Add ROI (Region of Interest)

1. Open the **ROI** tool and press **‘Add ROI’**.
2. Draw a box or shape around the area you want to monitor. For example, a specific room or parking spot.
3. Save the region, and the system will track what happens in that area.
4. You can use this to focus on important zones and ignore unimportant ones.

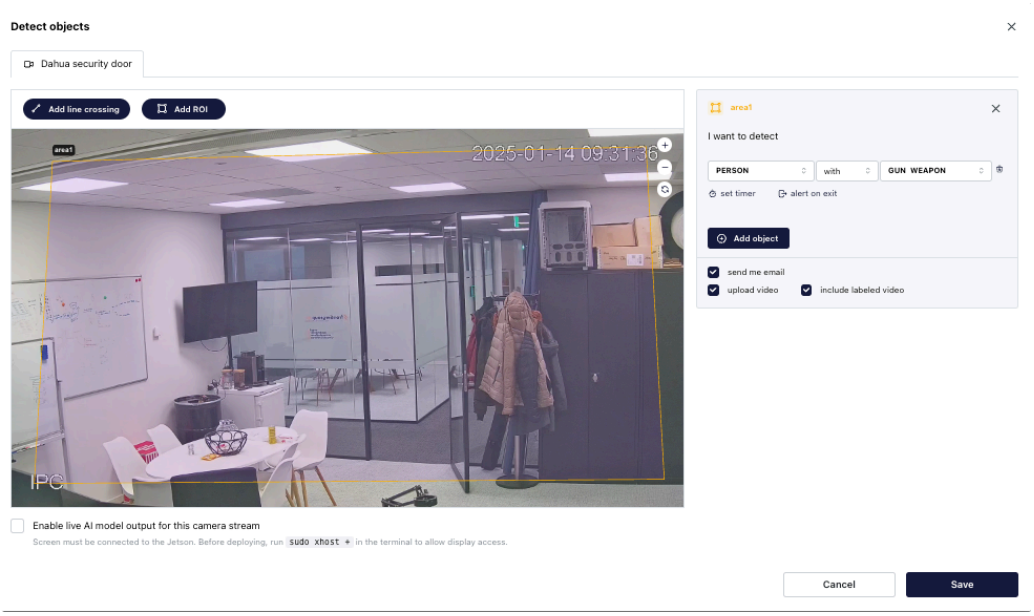

Event notifications

After adding a line crossing or ROI, you can set up event notifications on the right side. These settings let you choose what triggers an event, when it happens, and how it works.

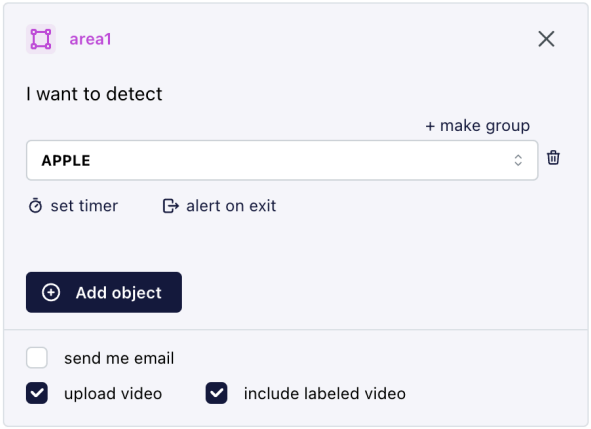

The following notification options can be configured for each ROI (Region of Interest) in the platform:

Option	Description
+ make group	If you want to combine classes. It opens up a drop down with ‘‘with'' and without. In the picture of ROI you can see an example. **Be aware that this option works best for multi stage applications!** It does work for single stage but the logic looks at a class within a class. For example is a person ‘‘with'' a helmet. If the helmet is on the ground not within the person class box it will not make an alert. A car with a person in single stage can be a good example that will work.
+ Add object	If you use a single stage application but you still want to detect 2 objects at the same time within your ROI to create an alert then use this button. Select OR or AND.
Set timer	Set the minimum time (in seconds) an object must stay in the area to trigger an alert.
Alert on exit	Enable this option to trigger an alert when an object exits the defined area (ROI)
Send me email	When an object is detected within the ROI, an email notification will be sent to the configured address.
Upload video	Automatically attach a video recording to the event created on the platform.
Include labeled video	Additionally, include a video recording with bounding boxes highlighting detected objects.

Important Note:
If you request both standard and labeled (bounding box) videos, the app will create two recordings: one with bounding boxes and one without. Be careful when using this on devices with limited resources, like the NVIDIA Jetson Nano, as it can increase system load.

To keep the system running smoothly:

- The app skips event detection during the first 20 frames after starting.
- There is a 2-second pause between events.

Count Objects

If you only want to get alerts when a certain number of objects are in your ROI, use this business logic. For example, you can set it to send an alert only when there are 4 or more people in the ROI.

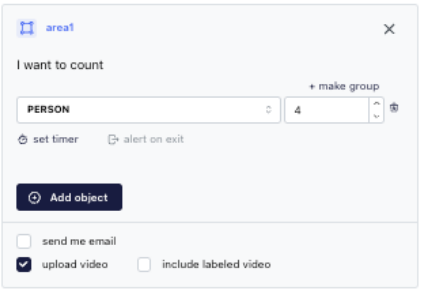

Select integrations and app settings

In the next screen you can select different output integrations and application settings.

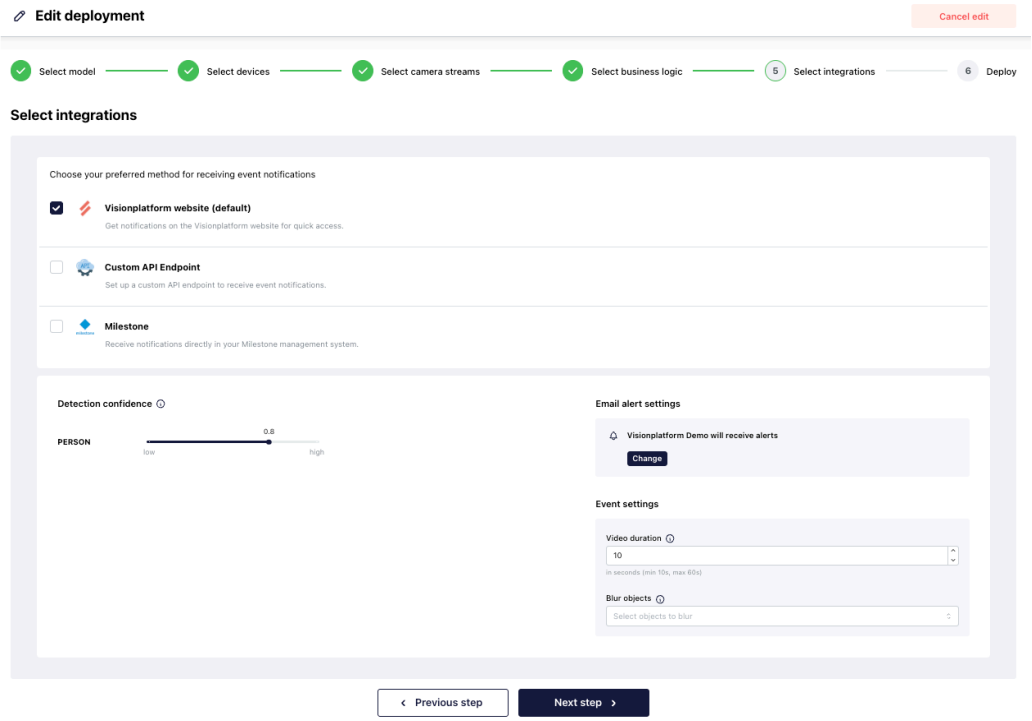

Select Output Integrations

Choose how you want to receive notifications:

- **Visionplatform.ai Website** (default): Get alerts directly on the platform
- **Custom API Endpoint**: Set up your own system to receive notifications.
  - add your REST API endpoint
- **Milestone**: Send notifications to your Milestone management system.
  - Enter the url from the Milestone AI Bridge

Adjust Application Settings

- **Detection Confidence**: Set how sure the system needs to be before detecting an object (low to high).
  - When you make your first deployment it's advised to use 0.4-0.6. It will generate more events but it allows you to collect more data to add to the model
  - Once you are more confident we suggest using 0.7 as minimum
- **Event Settings**:
  - **Video Duration**: Choose how long each event video will be (10–60 seconds).
  - **Blur Objects**: Users can choose to blur faces or license plates in event video recordings to meet privacy standards like GDPR. The blurring is permanent and cannot be undone, as it changes the image pixels directly.

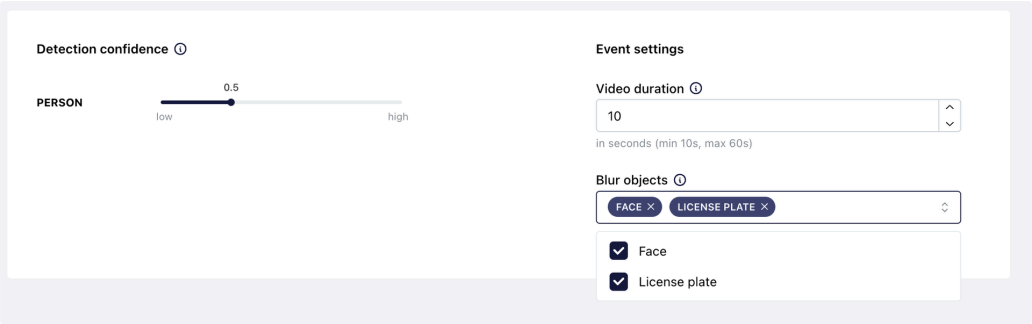

When you're done, click **Next Step** to continue.

Email Alert Settings

By default, email alerts are sent to the user who created the deployment. You can customize these settings to better suit your needs:

- **Assign Users**
  - Select specific users in your workspace to receive email alerts.
- **Set a Schedule**
  - Choose when alerts should be sent, such as specific days and times (e.g., Monday to Friday, 9 AM to 5 PM).
  - For example - Monday-Friday from 17:30 till 08:00 in the morning, split up the shhedule till 23:59

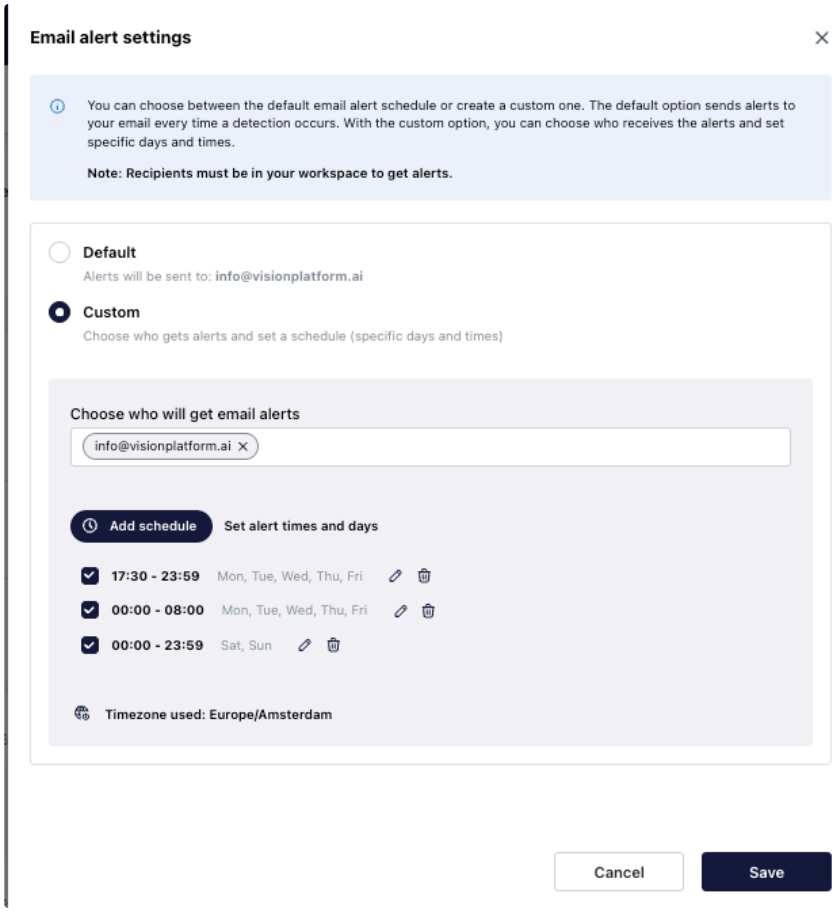

Deploy

Now press **Deploy**, and your entire application will be sent to the hardware.

Events

Your application, if you selected to receive notifications on visionplatform.ai, will now generate events

Events Overview

This screen helps you manage and view all detected events with the following features:

Details Shown for Each Event:

- **Captured On**: Shows the date and time the event was recorded (sortable).
- **Device**: Displays the device that captured the event (sortable).
- **Camera**: Indicates the name of the camera used.
- **Detections**: Lists what was detected (e.g., person, object).
- **Confidence**: Shows the system's confidence in the detection.
- **Actions**: Allows you to:
  - **Add a note** to the event for reference.
  - **Delete the event** if needed.

Top Menu Options:

- **Download Events**: Export event data in PDF or CSV format.
- **Filter Button**: Filter events based on date/time (Captured On) or Device.
- **Customization Options:**
  - **Show or Hide Columns**: Adjust which columns you want to see, like hiding Confidence or Camera Name.
  - **Toggle Event Density**: Switch between compact or more spaced-out rows to suit your preferences.
  - **Toggle Full Screen**: Expand the event view to full screen for better visibility.