




Embedded AI: detecting abandoned objects in public spaces

21 October 2025, 14:48

Nicolás González-Deleito

Giulia Murtas

A bag left on a railway platform, an abandoned suitcase in the middle of an airport concourse: this type of situations are worrying and tie up precious security resources. Surveillance teams must quickly establish whether the item has been deliberately abandoned or left unattended accidentally. With passenger numbers increasing and cameras recording thousands of images every day, it is becoming difficult to analyse all of this manually.

The Reconnaissance project is tackling this problem by developing video and audio analysis systems based on artificial intelligence and deep learning. They can run directly on cameras or local gateways (edge computing, i.e. processing data on-site without sending it to the cloud).

From research to deployable technology for public safety

This is a collaborative project funded by the Walloon Region as part of the MecaTech cluster. It brings together several industrial and academic partners: Phoenix AI (the project coordinator), Thelis and WNM, Chapsvision/ACIC, the University of Liège, UCLouvain and Sirris.

One of the project's use cases, on which ACIC and Sirris have worked, focuses on the detection of abandoned objects in public spaces, a vital issue for transport operators, public space managers and local authorities. The objective is clear: to detect abandoned objects automatically and alert security teams in real time, while avoiding false alarms that result in a site being shut down unnecessarily.

The technical challenges of abandoned object detection

Detecting an abandoned object seems simple, but the challenges are numerous:

- **Lack of available data:** public image databases rarely show abandoned objects in real-life situations (railway stations, shopping centres etc.)
- **Rarity and variety:** objects are rarely abandoned, and real image databases are proprietary, unlabelled and contain personal data, making it harder to train models
- **Different viewing angles:** security cameras often film from above, unlike standard images taken at a person's height
- **Required reliability:** the system must detect real objects without generating too many false alerts
- **Technical constraints:** the model must be lightweight in order to run on platforms such as NVIDIA's Jetson Orin (a very powerful minicomputer designed for embedded AI) without relying on a cloud server

How Sirris improves detection with AI

At Sirris, we contributed to data preparation and the development of learning methods suited to this project. We drew on our expertise in data processing and artificial intelligence.

1. Creating a customised dataset

Public datasets such as COCO 2017 contain numerous images in contexts far removed from settings such as railway stations or shopping centres. The inclusion of these out-of-context images was unnecessarily complicating model training. We therefore selected the most relevant images (including by filtering COCO 2017's "stuff" categories in order to eliminate the least relevant images) and explored other datasets (i-LIDS, YouTube-8M Segments, Open Images, Visual Genome, PETS2006) to create a more representative dataset.

We then used an image classification model (ResNet50) to generate "embeddings": digital representations of the images. These embeddings allowed us to measure the similarity between images and select those that provided the best fit for our use case.


This gave us a much more representative dataset, making it easier to train the model and improving its detection performance.

2. Implementing a semi-supervised and active learning pipeline

Labelling thousands of images manually is time-consuming. Drawing on the latest research (such as that of Chen et al.), we designed a labelling pipeline that combines semi-supervised and active learning. In concrete terms, the model also learns from unlabelled images and only seeks help from a human expert for the most complex cases.

The images are first augmented (variations in brightness, contrast, mirroring etc.). If the model provides stable predictions between the original and augmented versions, the simple images are automatically labelled and added to the dataset. By contrast, complex or ambiguous images are sent to the active learning phase and submitted to a human expert for annotation. This means that far fewer images need to be labelled manually, making it faster to create the training set.

The diagram below illustrates this process and shows how the pipeline automatically selects the easiest images and involves a human expert only for the most complex cases.



3. Field testing and validation

To simplify the validation of the pipeline, we packaged the solution in Docker components. Tests carried out by ACIC on NVIDIA's Jetson Orin platform showed that the system can run on edge computing – directly on local cameras or gateways without sending images to the cloud. This ensures better responsiveness and data confidentiality.

Thanks to these contributions (and those made by ACIC), the system is more reliable, faster to train and ready for implementation in real-world settings such as railway stations or airports.

Continuous improvement through active learning

We are currently developing an innovative active learning approach that capitalises on the difficulty object detection models have in correctly detecting the different objects shown in an image. We hope that this will further improve accuracy while limiting the number of annotations required.

A collaborative industry-driven project

The complementary nature of the industrial stakeholders, integrators, researchers and technology centre (Sirris) means that the results will not remain at the prototype stage. They can be quickly integrated into operational products and services.

Discuss your AI projects with our experts

Would you like to explore how AI and edge computing can keep your facilities secure or optimise your video analysis? Contact our experts.

[Get in touch with Giulia Murtas](#) [Get in touch with Nicolás González-Deleito](#)

Source

- [1] S. Chen, Y. Yang, and Y. Hua, "Semi-Supervised Active Learning for Object Detection", Electronics 2023,12, 375.

Authors



Nicolás González-Deleito



Giulia Murtas