



DFRWS EU 2026 - Selected Papers from the 13th Annual Digital Forensics Research Conference Europe

## Plug to place: Indoor multimedia geolocation from electrical sockets for digital investigation

Kanwal Aftab<sup>a,\*</sup>, Graham Adams<sup>b</sup>, Mark Scanlon<sup>a</sup>

<sup>a</sup> Forensics and Security Research Group, School of Computer Science, University College Dublin, Belfield, Dublin, 4, Ireland

<sup>b</sup> School of Engineering, Case Western Reserve University, Cleveland, OH, United States



### ARTICLE INFO

#### Keywords:

Multimedia Geolocation  
Computer Vision  
Hotels-50K  
Indoor  
Multimedia forensics  
Human trafficking

### ABSTRACT

Computer vision is a rapidly evolving field, giving rise to powerful new tools and techniques in digital forensic investigation, and shows great promise for novel digital forensic applications. One such application, indoor multimedia geolocation, has the potential to become a crucial aid for law enforcement in the fight against human trafficking, child exploitation, and other serious crimes. While outdoor multimedia geolocation has been widely explored, its indoor counterpart remains underdeveloped due to challenges such as similar room layouts, frequent renovations, visual ambiguity, indoor lighting variability, unreliable GPS signals, and limited datasets in sensitive domains.

This paper introduces a pipeline that uses electrical sockets as consistent indoor markers for geolocation, since plug socket types are standardised by country or region. The three-stage deep learning pipeline detects plug sockets (YOLOv11, mAP@0.5 = 0.843), classifies them into one of 12 plug socket types (Xception, accuracy = 0.912), and maps the detected socket types to countries (accuracy = 0.96 at >90 % threshold confidence). To address data scarcity, two dedicated datasets were created: socket detection dataset of 2328 annotated images expanded to 4074 through augmentation, and a classification dataset of 3187 images across 12 plug socket classes. The pipeline was evaluated on the Hotels-50K dataset, focusing on the TraffickCam subset of crowd-sourced hotel images, which capture real-world conditions such as poor lighting and amateur angles. This dataset provides a more realistic evaluation than using professional, well-lit, often wide-angle images from travel websites. This framework demonstrates a practical step toward real-world digital forensic applications. The code, trained models, and the data for this paper are available open source.

### 1. Introduction

Human trafficking is a severe global crime affecting millions across all ages, genders, and backgrounds, causing deep personal, community, and societal harm (Walby and Francis, 2025). It entails the illegal trade of people through deception, violence, or exploitation, resulting in forced labour, sexual abuse, and organ trafficking (Dimas et al., 2022). Given its severe and long-lasting impact, the fight against human trafficking is explicitly prioritised under 3 of the United Nations Sustainable Development Goals (SDGs) (United Nations Office on Drugs and Crime, 2025). In addition, child sexual exploitation material (CSEM) investigation is one of the most common case types encountered in digital forensics laboratories within law enforcement agencies worldwide (Hargreaves et al., 2024).

With rapid technological advancement, the rise of feature-rich

smartphones, expanded storage capabilities, widespread internet access, and the growing influence of social media, nearly every facet of modern life has become digital (Kumar et al., 2024). However, modern technologies are increasingly exploited by offenders to facilitate serious crimes, underscoring the critical importance of digital forensics (Fähndrich et al., 2023). As the volume of digital data continues to surge, the analysis and interpretation of digital evidence have become indispensable to modern investigations (Khadka et al., 2024). The application of artificial intelligence (AI) to digital forensic investigations is still very much in its infancy (Breitinger et al., 2024), but nonetheless, practitioners have already identified image/media classification as having the most potential for the future use of AI in their investigations (Hargreaves et al., 2024). In parallel, hotel recognition has become a common need for human trafficking investigations, as hotels are frequently used as intermediary stopover locations during the trafficking

\* Corresponding author.

E-mail address: [kanwal.aftab@ucdconnect.ie](mailto:kanwal.aftab@ucdconnect.ie) (K. Aftab).

<https://doi.org/10.1016/j.fsidi.2026.302056>

Available online 24 March 2026

2666-2817/© 2026 The Author(s). Published by Elsevier Ltd on behalf of DFRWS. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

of victims (Kamath et al., 2021). Online human trafficking advertisements or the interception of organised crime gang's electronic communications are often the starting points for many human trafficking investigations. Identifying hotel rooms from these photographs is an extremely arduous task. Indeed, police agencies often resort to crowdsourcing their locations. For example, hotel rooms often feature in the Trace an Object projects run by Europol<sup>1</sup> or the Australian Centre to Counter Child Exploitation,<sup>2</sup> asking for the general public's help in identifying the hotels during investigations of cases involving child sexual exploitation material. In the DFPulse 2024 survey (Hargreaves et al., 2024), digital forensic practitioners identified image/media classification and CSAM investigation as two of the main areas where artificial intelligence has the potential to assist in their future cases.

In terms of automated geolocation, indoor geolocation specifically faces significant challenges. GPS is unreliable indoors. Alternative indoor sensors are relatively costly and have limited range and accuracy (Hromadová et al., 2021), and image metadata is frequently stripped during online or instant messaging sharing – making it difficult to trace an image's origin (Bamigbade et al., 2024). While outdoor environments typically provide clear geolocation cues, such as landmarks and infrastructure (Zhang et al., 2021; Brejcha and Čadík, 2017; Luo et al., 2011), indoor settings are far more complex. Similar layouts, recurring furniture, and inconsistent lighting make it challenging to reliably differentiate between locations (Bhavanasi and Stylianou, 2023). Amid these challenges, it is therefore crucial to identify consistent and distinctive cues for indoor geolocation. One such feature of indoor environments is the presence of electrical sockets, which can serve as distinctive and geographically informative visual markers. Each country adheres to standardised socket designs governed by national or regional electrical regulations (International Electrotechnical Commission, 2015). Plug and socket designs vary considerably in shape, grounding, voltage, and frequency. This study focusses solely on their visual characteristics (McGregor, 2025). Using Computer Vision (CV), detecting and classifying these plug sockets can provide a reliable cue for narrowing down the search space, as illustrated in Fig. 1.

### 1.1. Contribution of this work

This paper makes the following key contributions:

- **Dataset Creation and Release.** Two novel CV datasets have been curated and publicly shared: (i) a socket detection dataset for identifying sockets in indoor room images, and (ii) a socket classification dataset containing 12 socket categories. These datasets provide valuable benchmarks for advancing research in fine-grained indoor object recognition.
- **Algorithmic Pipeline for Detection and Classification.** A comparative study of state-of-the-art detection models and Convolutional Neural Network (CNN) classifiers has been conducted. Based on this analysis, a pipeline was designed to select the most effective combination of detection and classification models for socket recognition.
- **Evaluation of the Proposed Approach on Real-World Data.** The approach is evaluated on the Hotels-50K dataset, specifically the TraffickCam subset, demonstrating its practical utility for law enforcement by narrowing down search spaces in real-world human trafficking investigations. Furthermore, this work aims to lay the foundation for a universal socket detector, enabling broader applications beyond the law enforcement/investigative domains.

All code, trained models, and data created as part of this paper are made available open source and can be accessed from <https://github.com/markscanlonucd/plugtoplace>.

<sup>1</sup> <https://www.europol.europa.eu/stopchildabuse>.

<sup>2</sup> <https://www.acce.gov.au/what-we-do/trace-an-object>.

## 2. Related work

To address the challenge of indoor multimedia geolocation, specifically hotel recognition, researchers have employed a variety of techniques, ranging from hand-crafted feature extraction, image classification to neural networks for automated feature extraction, as well as more advanced approaches such as object-based similarity, image embeddings, and semantic scene understanding. Regardless of the methodology employed, the fundamental building block remains CV. Broadly, researchers approach multimedia geolocation in two ways. When a large and representative reference dataset is available, the task is formulated as a Content based Image Retrieval (CBIR) problem; in the absence of such a dataset, geolocation instead relies on universal visual cues, i.e., landmarks or sun angle in outdoor environments, to infer location directly from the image itself (Hays and Efron, 2008; Kakar and Sudha, 2012).

In this context, high-quality data remains essential. However, obtaining such data can be particularly challenging in sensitive scenarios, such as human trafficking investigations. Recognising these ethical and operational challenges, the European Union formalised AI regulations through the AI Act (Gstrein et al., 2024; European Union, 2024), aiming to maximise the benefits of AI while enforcing rigorous ethical and safety standards.

In the context of data collection and processing, the creation of specialised datasets has been significant. The Hotels-50K dataset (Stylianou et al., 2019) contains over one million labelled images from 50,000 hotels worldwide, sourced from both travel websites and the Traffick Cam mobile application, and is designed specifically for hotel recognition research. The Hotel-ID dataset (Kamath et al., 2021) provides a similarly large-scale resource for the same domain. Both datasets serve as benchmarks for evaluating recognition approaches.

Deep learning models are often black-box systems, and their complex structures make it difficult to interpret how visual inputs lead to an output, obscuring the specific cues that drive similarity judgments. To address this, Black et al. (2022) introduced a paired image similarity visualisation technique tailored for Transformer architectures. Their method enables the inspection of attention maps to reveal which regions contribute most to a given similarity score. When comparing ResNet and Vision Transformer (ViT) on datasets such as Hotels-50K, Clean Google LandmarksV2, and Stanford Online Products, the study found that ViT's attention focused on meaningful cues, such as bed covers, runners, and tiled shower walls. These results emphasise the need for explainable embeddings in visual retrieval.

Similarly, Wazzan et al. (2024) explored how context affects object matching, finding that a moderate amount of context improves annotation efficiency and retrieval accuracy, while excessive context complicates recognition in ambiguous scenes.

Beyond these pixel-level and embedding-based methods, some research focussed on object-centric retrieval. This approach represents images as groups of distinct objects rather than just overall textures, aiming to bridge the gap between simple image features and what people actually see. For instance, Kim et al. (2003) introduced a method to detect and segment the main object in an image, as subjects are often centred in photos. The authors used relevance feedback on these segments to improve retrieval, though their reliance on colour made it less effective in busy scenes. Later, Pradhan et al. (2016) used visual saliency maps to identify object regions and create feature vectors for retrieval. This method worked well for images with a clear main object but was less effective for complex scenes.

Furthermore, Bhavanasi and Stylianou (2023) argued that hotel recognition cannot be treated as a standard image-classification task due to the extremely large number of hotel classes. To address this, the authors proposed an object-centric ensemble approach in which separate models are trained for different object types. The system employed a Hybrid Vision Transformer combining a ResNet-26 backbone for feature extraction with a ViT classifier across 3110 classes. This hybrid design

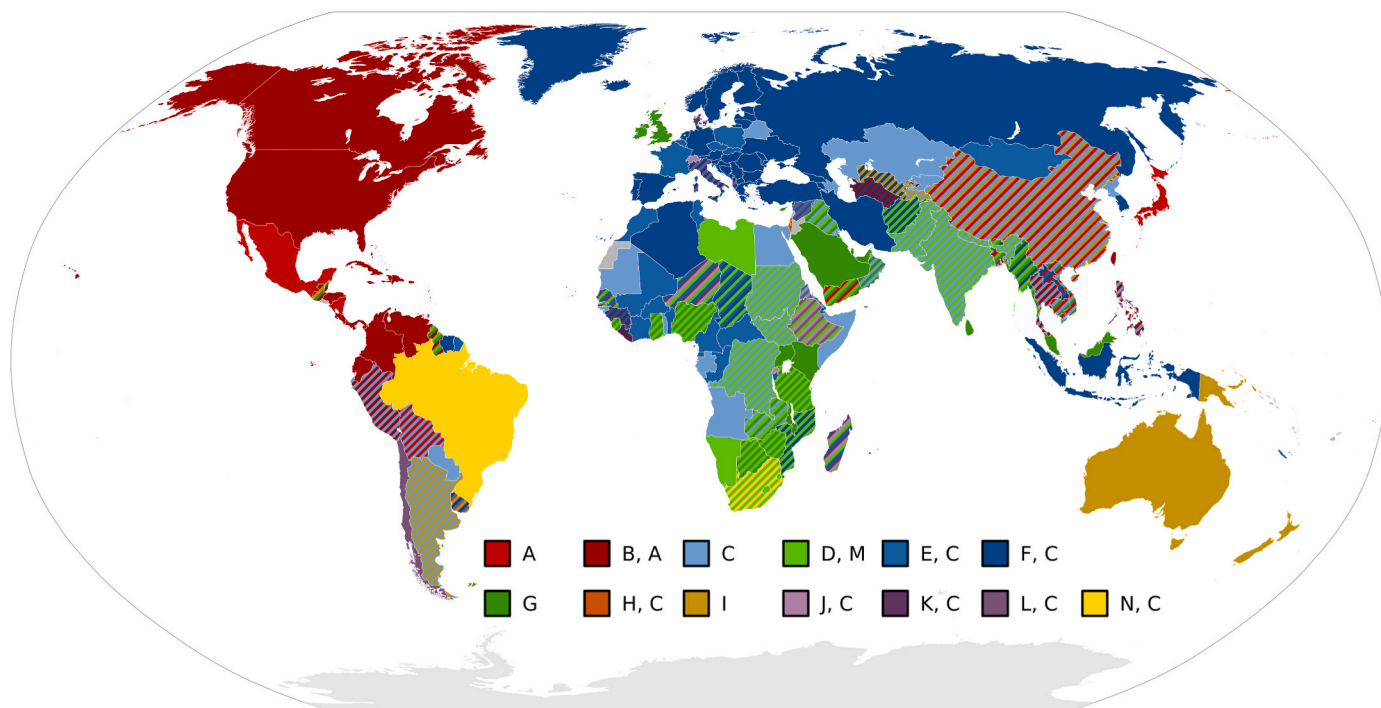


Fig. 1. Worldwide plug type distribution map ©

delivered superior performance, achieving nearly 80 % top-1 accuracy on the Hotel-ID 2022 dataset, compared with only 60 % using a conventional full-image method. This shows the value of pinpointing distinctive features, such as furniture, fixtures, or wall patterns. Additionally, semi-automated labelling pipelines (Wazzan et al., 2024) demonstrate that using less context speeds up labelling without hurting performance, suggesting that focusing on main objects and a moderate amount of context best balances interpretability and efficiency.

Besides these semantic methods, feature-level CBIR focusses on capturing broad visual qualities such as colour, texture, and layout (Cao et al., 2020). Among these attributes, colour is a dominant and widely used image characteristic, as it is both robust, computationally efficient, and independent of orientation and image resolution, making it highly versatile (Shamoi et al., 2022; Bakair et al., 2017). While RGB is common, colour spaces like CIELAB, Munsell, and fuzzy-based models often yield more meaningful results (Shamoi et al., 2022). Using a mix of colour spaces, such as RGB, YCbCr, and Lab\*, can also boost precision (Sangeetha et al., 2022). Specific to hotel room identification, Herrmann et al. (2024) evaluated CBIR systems using colour features on Hotels-50K, achieving over 95 % Top-50 accuracy with just two descriptors, thus supporting faster and more reliable investigative workflows. More recently, Bamigbade et al. (2025) combined major colour palettes and simple histograms with deep metric learning and classification to raise top-20 retrieval accuracy by 17 %. These findings show that blending handcrafted features with deep embeddings can improve both clarity and results.

Even with sophisticated methods, systematic reviews expose key gaps in using CV for social good. Dimas et al. (2022) note that Operations Research and Analytics efforts for anti-human trafficking mostly target sex trafficking and prosecution, with less attention to labour trafficking, prevention, or victim protection. Similarly, Bamigbade et al. (2024) highlighted the benefits of using CV based geolocation in investigations, calling for more data types and clearer deep learning models to extract useful information.

Overall, these studies show a shift in hotel recognition research: from data-driven deep metric learning to more comprehensible, object-focused, and feature-blended CBIR systems. This change not only

improves technical performance but also supports the growing demands for ethical, transparent, and socially responsible AI in sensitive investigations.

### 3. Experimental setup and analysis

The CV methodology proposed as part of this paper is a three-tiered process, as detailed in Fig. 2. In the first stage, YOLO object detection identifies and localises electric sockets in an image, generating cropped regions of interest (ROI). In the second stage, a CNN classifies each cropped socket ROI by type. Finally, in the third stage, the detected socket type is mapped to potential countries, which narrows the list of likely locations and supports law enforcement investigations. Sections 4 to 6 provide a detailed discussion of each stage, including dataset preparation for every stage, the methodologies used, and a comparative analysis of different algorithms along with their results.

### 4. Stage 1: socket detection

Stage 1 of the pipeline focuses on detecting electrical sockets in image to infer indoor geographical location. Electrical sockets are consistent and recognisable markers, as each country or region uses specific socket types defined by distinct pin configurations. These features are visually distinctive in indoor environments such as hotel rooms. A CV-based object detection model is applied to automatically identify socket instances and their positions within an image. The detected socket type is then used to constrain the possible geographical location, supporting subsequent investigative analysis.

#### 4.1. Dataset preparation

High-quality and diverse datasets are critical in CV, as both the quantity and quality of training data directly affect model accuracy and generalisation (Joshi et al., 2024). For this study, a total of 2328 socket images were compiled, of which 1525 were obtained from publicly available Roboflow (Roboflow (2025) socket datasets licenced under CC 4.0, while the remaining 803 images were cropped from hotel room

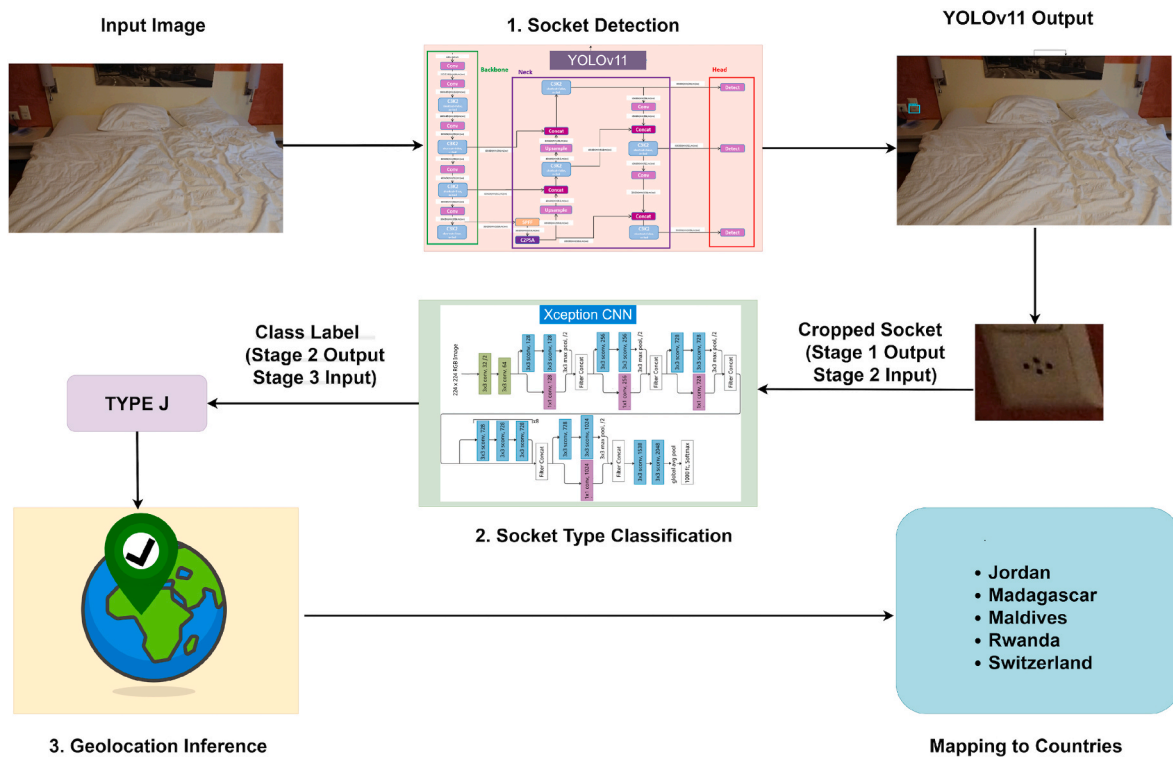


Fig. 2. Architecture of the proposed three-stage pipeline: (1) Socket detection, (2) Socket type classification, and (3) Geolocation.

scenes in the Hotels-50K dataset. This original dataset (Dataset A) was partitioned into training (70 %, 1629 images), validation (20 %, 455 images), and test (10 %, 244 images) subsets. All images were annotated with bounding boxes using the Roboflow platform and manually classified into two categories: class 0 (NA) for switchboards and class 1 for sockets. Although the initial focus was exclusively on sockets, the frequent misclassification of visually similar switchboard buttons necessitated the inclusion of the NA class to better differentiate non-target elements from actual sockets. This confusion was heavily compounded by the small sizes of the objects, their low resolution in cropped images, and the similar colour and high concentration of buttons and sockets on the switchboards.

To enhance dataset diversity and improve model generalisation, data augmentation was applied exclusively to the training set, doubling its size from 1629 to 3258 images. Augmentation techniques included random cropping (0–20 %), rotation ( $-15^\circ$  to  $+15^\circ$ ), grayscale conversion (15 % of images), hue adjustment ( $-24^\circ$  to  $+24^\circ$ ), and brightness adjustment ( $-19\%$  to  $+19\%$ ). After additional images were added, the dataset resulted in training (80 %, 3258 images), validation (10 %, 409 images), and test (10 %, 407 images) subsets, for a total of 4074 images. The validation and test sets were not augmented; however, an additional 117 images cropped from the Hotels-50K dataset were added to ensure an unbiased evaluation of model performance.

#### 4.2. Methodology

Recent advances in CV have positioned object detection as a fundamental task, achieving performance that, in some cases, rivals or even surpasses human capabilities (Borji and Itti, 2014). A wide range of state-of-the-art algorithms are available for object detection, and the choice of an appropriate model for custom datasets depends heavily on the specific requirements of the application. Object detection algorithms are broadly classified as two-stage or single-stage detectors. Two-stage models are more accurate but slower, while single-stage models offer faster inference with only a slight accuracy trade-off (Mohammed, 2025). To handle large volumes of images efficiently, speed was

prioritised, rendering single-stage object detectors a more practical option. The YOLO family was chosen for its combination of high speed and competitive accuracy (Wang et al., 2023), which is suitable for socket detection as a first-stage step before potential refinement. All YOLO models used in this study are initialised with pre-trained weights and subsequently fine-tuned for socket detection, allowing reliable performance with limited task-specific training data.

Although newer YOLO versions introduce architectural improvements, the latest release is not always stable (Jegham et al., 2025). A comparative analysis was then conducted. The experimental outcomes, as outlined in Table 1, demonstrate the performance of each model across diverse metrics, starting with YOLOv8 as a benchmark, followed by YOLOv11 and YOLOv12. Each release provides multiple variants (nano, small, medium, large). This study predominantly focused on small and medium models, and in one case, the large model was also considered to balance computational cost with accuracy. Hyperparameter tuning, particularly the number of training epochs, was evaluated to minimise underfitting and overfitting, and the effect of data augmentation on model performance was also assessed.

#### 4.3. Performance matrix

Mean Average Precision (mAP) is a standard metric in object detection that summarises how accurately a model detects and localises objects. It is based on precision and recall, using Intersection over Union (IoU) to measure the overlap between predicted and ground truth bounding boxes.  $\text{mAP}@0.5$  computes the average precision at an IoU threshold of 0.5, meaning a detection is considered correct if at least 50 % of the predicted box overlaps the ground truth, providing a measure of overall detection success.  $\text{mAP}@0.5:0.95$  averages performance across multiple IoU thresholds from 0.5 to 0.95, evaluating the model's ability to both detect objects and precisely localise them. High  $\text{mAP}@0.5:0.95$  indicates strong localisation accuracy, while  $\text{mAP}@0.5$  primarily reflects general detection performance. These metrics together provide a robust assessment of the model's effectiveness in identifying and locating objects in images.

**Table 1**

Test results for all classes across models. Bold indicates the best value per column.

Version	Type	Epochs	Precision (P)	Recall (R)	mAP@0.5	mAP@0.5:0.95
YOLOv8	s	180	0.843	0.709	0.793	0.529
YOLOv8	s	250	0.834	0.709	0.790	0.524
YOLOv8	m	set = 300 (early stopping at 233)	<b>0.883</b>	0.750	0.809	<b>0.540</b>
YOLOv11	s	180	0.847	0.734	0.804	<b>0.540</b>
YOLOv11	m	250	0.846	<b>0.791</b>	<b>0.832</b>	0.539
YOLOv11	l	set = 250 (early stopping at 234)	0.851	0.764	0.806	0.534
YOLOv12	s	180	0.720	0.650	0.748	0.483

#### 4.4. Comparative analysis and results

This study evaluated the performance of various YOLO models for socket detection using Dataset A, the original dataset, in the first phase, and subsequently assessed the effect of data augmentation in the second phase on the augmented data, Dataset B. All models were trained with a batch size of 16 and an input resolution of 640 pixels. The optimiser was set to AdamW, which automatically tuned hyperparameters and overrode the default learning rate and momentum values, resulting in an effective learning rate of 0.001667 with momentum fixed at 0.9.

The experiment began with YOLOv8, starting with the small variant (YOLOv8s) to assess the effect of increasing the number of training epochs from 180 to 250. The focus then shifted to the medium variant (YOLOv8m) to evaluate the trade-off between model capacity and performance. Although the number of epochs was initially set to 300, training was stopped early at 233 epochs due to the early stopping mechanism designed to prevent overfitting.

The analysis then progressed to YOLOv11, beginning with the small variant (YOLOv11s) trained for 180 epochs, followed by the medium variant (YOLOv11m) trained for 250 epochs. For the large variant (YOLOv11l), training was scheduled for 250 epochs but was stopped early at 234 epochs due to early stopping. Finally, YOLOv12 was evaluated; however, it achieved comparatively lower performance metrics.

Specifically, YOLOv12 underperformed, achieving mAP@0.5 of 0.748 and mAP@0.5:0.95 of 0.483, with a precision of 0.720 and a recall of 0.650. These results indicate weaker socket detection performance, highlighting that cutting-edge models do not always outperform more established versions such as YOLOv8 and YOLOv11. Newly released models are often unstable and improve over time; therefore, it is generally advisable to allow them to mature before deployment in CV applications (Jiang and Zhong, 2025).

Validation accuracy should not be solely relied upon, as it may indicate overfitting on the training data and not necessarily reflect performance on unseen test data (Santos and Papa, 2022). Therefore, test accuracy is considered more decisive than validation accuracy. On the test set, YOLOv8m (early stopping at 233 epochs) achieves the highest precision (0.883), while both YOLOv8m and YOLOv11s (180 epochs) achieve the highest mAP@0.5:0.95 (0.540). This indicates that these models are highly accurate in correctly detecting sockets, though their recall varies (0.750 for YOLOv8m, 0.734 for YOLOv11s). In contrast, YOLOv11m (250 epochs) attains the highest recall (0.791) and the highest mAP@0.5 (0.832), indicating strong overall detection coverage and good localisation.

To further assess the impact of data augmentation, YOLOv11m was trained on Dataset B with augmentation, while YOLOv11s was trained using both augmentation and 5-fold cross-validation. Table 2 summarises the results. Adding only augmentation slightly altered performance, with precision decreasing from 0.846 to 0.783, recall increasing from 0.748 to 0.765, mAP@0.5 decreasing from 0.832 to 0.766, and mAP@0.5:0.95 decreasing from 0.539 to 0.485. These results suggest that augmentation improves recall at a modest cost to precision and overall localisation accuracy.

In contrast, combining K-fold cross-validation with augmentation further enhances model robustness, reducing variance between folds and providing more reliable generalisation to unseen images.

**Table 2**

Comparison of YOLOv11m variant with data augmentation and k fold cross validation on and Dataset B. Bold indicates the best value per dataset.

Dataset	Setting	Precision (P)	Recall (R)	mAP@0.5	mAP@0.5:0.95
Test (Dataset B)	Yolov11m without Aug	0.805	0.748	0.792	0.498
	Yolov11m with Aug	0.783	0.765	0.766	0.485
K fold cross validation	Yolov11s with Aug	<b>0.8675</b>	<b>0.7990</b>	<b>0.8427</b>	<b>0.5771</b>

Specifically, average precision increased from 0.847 to 0.872, and recall increased from 0.734 to 0.756, demonstrating that K-fold training effectively mitigates overfitting on smaller datasets and improves overall detection performance. Based on this analysis, YOLOv11s with augmentation and K-fold cross-validation was identified as the best performer in Stage One socket detection. Figs. 3 and 4 show the visual socket detection results of this best-performing model.

## 5. Stage 2: socket type classification

Stage 2 of the pipeline focuses on classifying detected electric sockets into region-specific socket types based on their visual characteristics. Although 14 domestic socket types are internationally recognised, this stage considers only visual distinctions relevant for image-based classification. As illustrated in Figs. 5 and 1, socket designs exhibit strong regional variation, enabling geographical constraints to be inferred from socket appearance.

In this experiment, CV is used to classify socket types, since the focus is exclusively on visual characteristics, the study considers 12 classes instead of the full set of 14. Types D and M have been merged into a single class (DM) because, despite differences in pin size, their layouts are visually indistinguishable. Type M closely resembles Type D but features larger pins. Similarly, Types J and N are merged into a single class (JN). Both sockets use three round pins with nearly identical configurations, differing only in the precise offset of the earth pin – with Type J, it is offset by 5 mm, and with Type N, it is offset by 3 mm (International Electrotechnical Commission, 2015). This small structural difference makes them electrically incompatible, but the variation was deemed too subtle to be reliably distinguished visually. All other socket types present clear visual differences and are, therefore, treated as separate classes. In addition, a noise class was introduced to exclude non-socket objects, such as light switches, thermostats, or low-quality regions of interest detected by YOLO, thereby further enhancing accuracy.

### 5.1. Dataset preparation

The dataset for socket type classification was constructed from two primary sources. The first source was publicly available datasets from Roboflow, consisting of plug socket images licenced under CC 4.0. These images were downloaded, cleaned, and merged into a consistent




Fig. 3. YOLOv11m detection results on room images (a–b), showing bounding boxes for socket classes.




Fig. 4. YOLOv11m detection results on bathroom images (a–b), showing bounding boxes for socket classes.




Fig. 5. Plug and socket types from type A to type N (McGregor, 2025).

collection. The second source was the Hotels-50K dataset, from which socket regions were automatically detected, cropped, and assigned to their respective classes. After processing, the final dataset comprised 12 socket type classes, with the number of images per class summarised in Table 3. In total, the dataset contains 3187 images, which were

partitioned into training, validation, and test sets using a 70:15:15 split.

### 5.2. Methodology

A transfer learning approach for multi-class image classification

**Table 3**  
Socket types with image counts and corresponding country usage count.

Socket	Images	Countries
A	192	46
B	302	28
C	305	65
D/M	300	12/9
E	304	24
F	303	35
G	304	32
H	138	1
I	268	11
J/N	262	5/4
K	291	6
L	222	9
Total	3187	-

using five state-of-the-art CNN architectures was implemented: VGG16, InceptionV3, Xception, ResNet50, and ResNet101. The dataset, consisting of 3187 images across 12 classes (socket types), was split into training, validation, and test sets, with 2224 images for training, 473 images for validation, and 490 images for testing. Images were pre-processed and augmented with normalisation and horizontal flipping to improve generalisation. Each model was initialised with ImageNet pre-trained weights, and the convolutional base was frozen to leverage pre-learned features while training a new classification head for the target classes. Models were trained independently using categorical cross-entropy loss and the Adam optimiser.

### 5.3. Evaluation and comparative analysis

The performance of the model was evaluated using standard metrics, including accuracy, precision, recall, F1-score, and a confusion matrix to provide detailed insights into class-wise predictions. Among the evaluated models, VGG16, InceptionV3, and Xception achieved the highest accuracies. Xception attained the best overall performance with an accuracy of 91.22 %, consistently demonstrating high precision, recall, and F1-scores across all 12 classes, as summarised in Table 4. In comparison, VGG16 achieved 82.65 % accuracy, while InceptionV3 reached 89.80 %, highlighting the performance differences among the individual models.

To further validate the robustness of Xception, the impact of K-fold cross-validation was evaluated by modifying the dataset. A new Noise class, consisting of 304 images representing non-socket objects potentially missed by Step 1 YOLO, was added to the original 3187 images, resulting in a total of 3491 images across 13 classes (12 socket types + Noise). These images were organised into 13 folders for training and validation in a 5-fold cross-validation setup, while an additional 175 unseen images, spanning all classes, were reserved for testing. This setup allowed us to assess whether cross-validation improves model generalisation.

As summarised in Table 5, the 5-fold cross-validation strategy increased accuracy from 85.4 % to 87.7 %, precision from 87.8 % to 89.4 %, recall from 85.3 % to 88.4 %, and F1-score from 85.5 % to 88.1 %. These results demonstrate that K-fold cross-validation provided a modest but consistent improvement in performance over the single-run Xception model. Note that the test dataset differs from the split of

**Table 4**  
Performance summary of different models.

Model	Accuracy	Precision	Recall	F1-score
VGG16	0.827	0.846	0.816	0.819
InceptionV3	0.898	0.907	0.900	0.901
Xception	<b>0.912</b>	<b>0.914</b>	<b>0.910</b>	<b>0.911</b>
ResNet50	0.492	0.599	0.477	0.466
ResNet101	0.443	0.634	0.429	0.433

**Table 5**  
Performance summary of Xception models.

Model	Accuracy	Precision	Recall	F1
Xception	0.854	0.878	0.853	0.855
Xception (5-Fold CV)	<b>0.877</b>	<b>0.894</b>	<b>0.884</b>	<b>0.881</b>

70:15:15.

Therefore, Xception with 5 cross-validation was identified as the best performer in stage two socket type classification.

## 6. Stage 3: Inferring geolocation through socket Type–Country mapping

The final stage of the proposed pipeline focusses on using socket detection results to infer the geolocation of hotel rooms, highlighting the practical applicability of the proposed approach.

### 6.1. Test dataset preparation

The test dataset for this experiment was derived from the Hotels-50K TraffickCam dataset, available from its official GitHub repository. To prepare the data, the dataset was obtained by modifying the download script to preserve the original image resolution. This was a crucial step because the target object, electric sockets, is often small within the larger image, and resizing could lead to a loss of detail essential for accurate detection.

The Hotels-50K dataset comprises two subcategories: TraffickCam and travel website images. TraffickCam consists of crowd-sourced photos submitted by travellers worldwide, reflecting real-world conditions. After retaining and downloading the images in their original resolution of  $1024 \times 768$ , the dataset provides images in this resolution, with sockets typically appearing in small regions of approximately  $130 \times 87$  pixels.

In contrast, travel website images are professionally captured under ideal lighting and angles, often with colour correction and photo editing. These images are also provided at a lower resolution ( $350 \times 233$ ), making the sockets barely visible. As a result, they were excluded from testing. The TraffickCam subset, by depicting more realistic, non-professionally taken and/or edited photos, is considered to provide a better representation of the conditions encountered in practical investigations and was therefore used exclusively for evaluation.

To establish a ground truth for the experiment, the dataset was processed to associate each image with its corresponding country. This involved a multistep process: Merging Metadata: the various CSV files from the original dataset were consolidated to create a unified file containing image IDs, hotel IDs, and geographic coordinates (latitude and longitude). Geolocation: Using the `geopy.geocoders` library, the geographic coordinates were converted into country names. Standardisation: To ensure consistency, the country names were standardised using the `pycountry` library, as the raw geolocation output sometimes returned names in native languages. Illegal characters were removed, and spaces were replaced with underscores to create valid directory names.

The final dataset was restructured into a clean directory, with images organised into subfolders named after their respective countries. This structure, along with a consolidated CSV file containing all relevant metadata, streamlined the subsequent country-specific analysis. The reason images were arranged in folders instead of being directly taken and compared from the CSV file was to facilitate visual inspection, ensuring that the code functioned correctly and that each country was accurately represented with its respective socket type.

### 6.2. Data analysis, evaluation, and results

The model's performance was evaluated by assessing its ability to

detect a plug socket in an image and then classify it to determine its corresponding country. For each image in which a socket was successfully detected and classified, it was checked whether the predicted plug type and the image's actual country formed a valid pair according to this mapping. A correct match was assigned a score of 1, an incorrect match received a score of -1, and the noise class was assigned a neutral score of 0. These scores were used to calculate key performance metrics, such as the confusion matrix, precision, recall, and F1 score. This approach provided a comprehensive assessment of the model's accuracy. Finally, the results were compiled into a detailed summary report. They were presented with graphical visualisations, such as bar charts, to provide a clear and intuitive interpretation of the model's overall performance.

A total of 44,630 TraffickCam images were processed through the algorithmic pipeline. In the first stage, YOLO detected 3759 potential sockets. To enhance detection accuracy and eliminate false positives, a second-stage classifier was employed to identify and remove noise. Specifically, instances where non-socket objects (e.g., switchboards) were incorrectly detected as sockets in the first stage were classified as noise. This step identified 1393 noisy detections, leaving 2366 valid sockets. These valid detections were subsequently passed to the socket classification stage, where only those with a confidence threshold above 70 % were retained. The results are summarised in Table 6.

When considering different socket classification confidence thresholds, the performance varies. The confidence threshold represents the model's predicted probability that a socket belongs to the predicted class. Without applying any socket confidence threshold, 1967 predictions were correct, and 399 were incorrect, resulting in an accuracy of 83.08 %. At a threshold above 70 %, 1595 predictions were correct, and 146 were incorrect, resulting in an accuracy of 91.61 %. Increasing the threshold above 80 % reduced the number of correct detections to 1,421, while incorrect detections decreased to 95, yielding an improved accuracy of 93.73 %. At the highest threshold of above 90 %, correct detections further decreased to 1,167, with only 45 incorrect predictions, resulting in the highest accuracy of 96.29 %. These results illustrate the trade-off between confidence and accuracy: lower thresholds capture more sockets but result in more false positives, whereas higher thresholds reduce errors at the cost of missing some detections, as shown in Fig. 6.

7. Discussion

The proposed pipeline has been developed with generalisability in mind, aiming to detect electric sockets in images beyond the Hotels-50K dataset. As a universal visual cue, sockets have the potential to support geolocation in any indoor image. However, socket detection is inherently challenging due to their small size and the frequent low resolution of imagery. Despite these challenges, the method performs strongly when sockets are present, with geolocation inference achieving 96.29 % accuracy. The reported precision of 0.85 demonstrates effective detection, but only around 10 % of TraffickCam images contain visible sockets, which explains why applying high confidence thresholds reduces useable images to about 2 % of the dataset. Misidentification of countries is primarily due to the YOLO-based detection stage identifying non-socket objects, such as switchboards or other sources of visual noise, as potential sockets. In stage two, these detections are assigned the NOISE class, resulting in a incorrect socket score. Contributing factors include low image resolution, occlusions, and challenging lighting conditions. These issues indicate that most misclassifications are

Table 6  
Classifier threshold analysis for socket detection.

Class Confidence	Correct	Wrong	Total	Accuracy (%)
≥ 70%	1595	146	1741	91.61
≥ 80%	1421	95	1516	93.73
≥ 90%	1167	45	1212	96.29




Fig. 6. Country prediction accuracy by varying socket type confidence threshold values.

technical limitations of the current detection and classification pipeline, although some inherent limitations remain when atypical or uncommon socket types are present. Improvements in detection accuracy, higher-resolution imagery, and enhanced classification models could reduce such errors.

8. Conclusion

This paper presents a universal pipeline for socket detection and classification, motivated by its potential to enhance geolocation capabilities in investigative contexts where conventional cues, such as metadata, outdoor landmarks, and sensor data, are unavailable. By curating two new datasets, benchmarking multiple detection and classification models, and evaluating performance on the Hotels-50K TraffickCam dataset, this study demonstrates both the feasibility and the challenges associated with socket-based indoor geolocation. Despite challenges such as small object size and low-resolution imagery, the results demonstrate strong detection accuracy and validate the concept's practical potential, enabling investigators to automate geolocation across large volumes of digital evidence and convert it into actionable intelligence.

Beyond technical performance, this work contributes to the emerging field of AI-driven multimedia forensics, where visual scene elements are utilised to support digital investigations. Socket detection offers a unique, region-specific forensic cue that can aid law enforcement agencies in narrowing search regions, corroborating other forms of evidence, and identifying the possible origins of illicit or trafficking-related imagery. The proposed framework, therefore, represents an important step towards scalable, context-aware forensic tools that bridge the gap between CV and real-world investigative practice.

Future work will focus on grouping visually similar scenes based on distinctive feature representations, enabling indirect geolocation inference and the propagation of cues from one image across an entire image set. Such clustering can reveal patterns in large datasets and link unknown or unlabelled images to geographically known locations. Further integration of multimodal cues, including architectural styles, interior designs, or visible textual information, could strengthen geolocation when electrical fixtures are unavailable. Human inspection of representative subsets could also help verify socket presence and detection accuracy, refine confidence thresholds, and assess the pipeline's applicability to diverse, real-world images. Together, these strategies would make the system more scalable and practical for forensic and investigative applications, bridging the gap between AI-based contextual geolocation and digital investigations. to cluster visually similar scenes and strengthen geolocation inference when sockets are not visible. By combining socket-based detection with broader contextual cues, the

framework could evolve into a robust and general-purpose forensic system capable of supporting a wide range of investigative and humanitarian applications.

## References

- Bakair, A.S., Mohamed, T.M., Sadek, R.A., 2017. Enhancing the capability of color based CBIR by voting of color histograms. In: 2017 27th International Conference on Computer Theory and Applications (ICCTA), pp. 4–10. <https://doi.org/10.1109/ICCTA43079.2017.9497165>.
- Bamigbade, O., Scanlon, M., Sheppard, J., 2025. Improving image embeddings with colour features in indoor scene geolocation. *IEEE Access* 13, 79860–79870. <https://doi.org/10.1109/ACCESS.2025.3564496>.
- Bamigbade, O., Sheppard, J., Scanlon, M., 2024. Computer vision for multimedia geolocation in human trafficking investigation: a systematic literature review. URL: <https://arxiv.org/abs/2402.15448>.
- Bhavanasi, S.S., Stylianou, A., 2023. Hotel recognition using object ensembles. In: 2023 IEEE Applied Imagery Pattern Recognition Workshop (AIPR), pp. 1–8. <https://doi.org/10.1109/AIPR60534.2023.10440661>.
- Black, S., Stylianou, A., Pless, R., Souvenir, R., 2022. Visualizing paired image similarity in transformer networks. In: 2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp. 1534–1543. <https://doi.org/10.1109/WACV51458.2022.00160>.
- Borji, A., Itti, L., 2014. Human vs. computer in scene and object recognition. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition, pp. 113–120. <https://doi.org/10.1109/CVPR.2014.22>.
- Breitinger, F., Hilgert, J.N., Hargreaves, C., Sheppard, J., Overdorf, R., Scanlon, M., 2024. DFRWS EU 10-year review and future directions in digital forensic research. *Forensic Sci. Int.: Digit. Invest.* 48, 301685. <https://doi.org/10.1016/j.fsidi.2023.301685>. <https://www.sciencedirect.com/science/article/pii/S2666281723002044>. DFRWS EU 2024 - Selected Papers from the 11th Annual Digital Forensics Research Conference Europe.
- Brejcha, J., Čadík, M., 2017. State-of-the-art in visual geo-localization. *Pattern Anal. Appl.* 20, 613–637. <https://doi.org/10.1007/s10044-017-0611-1>.
- Cao, B., Araujo, A., Sim, J., 2020. Unifying deep local and global features for image search. In: Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX. Springer-Verlag, Berlin, Heidelberg, pp. 726–743. [https://doi.org/10.1007/978-3-030-58565-5\\_43](https://doi.org/10.1007/978-3-030-58565-5_43).
- Dimas, G.L., Konrad, R.A., Lee Maass, K., Trapp, A.C., 2022. Operations research and analytics to combat human trafficking: a systematic review of academic literature. *PLoS One* 17, 1–24. <https://doi.org/10.1371/journal.pone.0273708>.
- European Union, 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). Official Journal of the European Union, L2024/1689. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689>. (Accessed 16 December 2025).
- Fährdrich, J., Honekamp, W., Povalej, R., Rittelmeyer, H., Berner, S., Labudde, D., 2023. Digital forensics and strong AI: a structured literature review. *Forensic Sci. Int.: Digit. Invest.* 46, 301617. <https://doi.org/10.1016/j.fsidi.2023.301617>. URL: <https://www.sciencedirect.com/science/article/pii/S2666281723001294>.
- Gstrein, O., Haleem, N., Zwittner, A., 2024. General-purpose AI regulation and the European Union AI Act. *Internet Policy Review* 13, 1–26. <https://doi.org/10.14763/2024.3.1790>.
- Hargreaves, C., Breitinger, F., Dowthwaite, L., Webb, H., Scanlon, M., 2024. DFPulse: the 2024 digital forensic practitioner survey. *Forensic Sci. Int.: Digit. Invest.* 51, 301844. <https://doi.org/10.1016/j.fsidi.2024.301844>. URL: <https://www.sciencedirect.com/science/article/pii/S2666281724001719>.
- Hays, J., Efron, A.A., 2008. IM2GPS: estimating geographic information from a single image. In: 2008 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8. <https://doi.org/10.1109/CVPR.2008.4587784>.
- Herrmann, J., Bamigbade, O., Sheppard, J., Scanlon, M., 2024. Perceptual colour-based geolocation of human trafficking images for digital forensic investigation. In: 2024 Cyber Research Conference - Ireland (Cyber-RCI). *IEEE*, pp. 1–8. <https://doi.org/10.1109/Cyber-RCI60769.2024.10941203>.
- Hromadová, V., Machaj, J., Brída, P., 2021. Impact of user orientation on indoor localization based on wi-fi. *Transp. Res. Procedia* 55, 882–889. <https://doi.org/10.1016/j.trpro.2021.07.056>. <https://www.sciencedirect.com/science/article/pii/S2352146521004609>, 14th International Scientific Conference on Sustainable, Modern and Safe Transport.
- International Electrotechnical Commission, 2015. Plugs and socket-outlets for domestic and similar general use standardized in member countries of IEC. Technical Report. URL: <https://webstore.iec.ch/publication/23628>.
- Jegham, N., Koh, C.Y., Abdelatti, M., Hendawi, A., 2025. YOLO evolution: a comprehensive benchmark and architectural review of YOLOv12, YOLO11, and their previous versions. URL: <https://arxiv.org/abs/2411.00201>. arXiv:2411.00201.
- Jiang, T., Zhong, Y., 2025. Odverse33: is the new YOLO version always better? A multi domain benchmark from YOLO v5 to v11. *CoRR* abs/2502.14314. <https://doi.org/10.48550/ARXIV.2502.14314>.
- Joshi, S., Jain, A., Payani, A., Mirzasoleiman, B., 2024. Data-efficient contrastive language-image pretraining: prioritizing data quality over quantity. In: Dasgupta, S., Mandt, S., Li, Y. (Eds.), Proceedings of the 27th International Conference on Artificial Intelligence and Statistics. PMLR, pp. 1000–1008. URL: <https://proceedings.mlr.press/v238/joshi24a.html>.
- Kakar, P., Sudha, N., 2012. Authenticating image metadata elements using geolocation information and sun direction estimation. In: 2012 IEEE International Conference on Multimedia and Expo, pp. 236–241. <https://doi.org/10.1109/ICME.2012.82>.
- Kamath, R., Rolwes, G., Black, S., Stylianou, A., 2021. The 2021 Hotel-ID to combat human trafficking competition dataset. URL: <https://arxiv.org/abs/2106.05746>.
- Khadka, S., Singh, R., Begum, M., Muhammad, O., Naem, B., Chaeikar, S.S., 2024. A review on computer crimes and digital forensics. In: 2024 IEEE Consumer Life Tech (ICLT), pp. 1–4. <https://doi.org/10.1109/ICLT63507.2024.11038544>.
- Kim, S., Park, S., Kim, M., 2003. Central object extraction for object-based image retrieval. In: Bakker, E.M., Lew, M.S., Huang, T.S., Sebe, N., Zhou, X.S. (Eds.), Image and Video Retrieval. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 39–49. URL: [https://link.springer.com/chapter/10.1007/3-540-45113-7\\_5#citeas](https://link.springer.com/chapter/10.1007/3-540-45113-7_5#citeas).
- Kumar, S., Verma, A.K., Mirza, A., 2024. Digital Revolution, Artificial Intelligence, and Ethical Challenges. Springer Nature Singapore, Singapore, pp. 161–177. [https://doi.org/10.1007/978-981-97-5656-8\\_11](https://doi.org/10.1007/978-981-97-5656-8_11).
- Luo, J., Joshi, D., Yu, J., Gallagher, A., 2011. Geotagging in multimedia and computer vision—a survey. *Multimed. Tool. Appl.* 51, 187–211. <https://doi.org/10.1007/s11042-010-0623-y>.
- McGregor, C.H., 2025. Plug & socket types around the world. URL: <https://www.worldstandards.eu/electricity/plugs-and-sockets/>. (Accessed 18 September 2025).
- Mohammed, S.Y., 2025. Architecture review: two-stage and one-stage object detection. *Franklin Open* 12, 100322. <https://doi.org/10.1016/j.fraope.2025.100322>. URL: <https://www.sciencedirect.com/science/article/pii/S2773186325001100>.
- Pradhan, J., Pal, A.K., Banka, H., 2016. A prominent object region detection based approach for CBIR application. In: 2016 Fourth International Conference on Parallel, Distributed and Grid Computing (PDGC), pp. 447–452. <https://doi.org/10.1109/PDGC.2016.7913237>.
- Roboflow, 2025. Roboflow: computer vision tools for developers and enterprises. URL: <https://roboflow.com/>. (Accessed 18 September 2025).
- Sangeetha, S.K.B., Mathivanan, S.K., Pandi, T., Arivu selvan, K., Jayagopal, P., Teshite Dalu, G., 2022. An enhanced triadic color scheme for content-based image retrieval. *Math. Probl Eng.* 2022, 5736630. <https://doi.org/10.1155/2022/5736630>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1155/2022/5736630>.
- Santos, C.F.G.D., Papa, J.a.P., 2022. Avoiding overfitting: a survey on regularization methods for convolutional neural networks. *ACM Comput. Surv.* 54. <https://doi.org/10.1145/3510413>.
- Shamoi, P., Sanyszbayev, D., Abiley, N., 2022. Comparative overview of color models for content-based image retrieval. In: 2022 International Conference on Smart Information Systems and Technologies (SIST), pp. 1–6. <https://doi.org/10.1109/SIST54437.2022.9945709>.
- Stylianou, A., Xuan, H., Shende, M., Brandt, J., Souvenir, R., Pless, R., 2019. Hotels-50K: a global hotel recognition dataset. In: Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence. AAAI Press, pp. 726–733. <https://doi.org/10.1609/aaai.v33i01.3301726>.
- United Nations Office on Drugs and Crime, 2025. Human trafficking and the SDGs. <https://www.unodc.org/unodc/human-trafficking/sdgs.html>.
- Walby, S., Francis, B., 2025. Improving the estimate of trafficking in human beings and modern slavery by integrating data from ILO/walk Free/IOM and UNODC. *Soc. Indic. Res.* 176, 669–693. <https://doi.org/10.1007/s11205-024-03474-w>.
- Wang, X., Li, H., Yue, X., Meng, L., 2023. A comprehensive survey on object detection YOLO. In: The 5th International Symposium on Advanced Technologies and Applications in the Internet of Things (ATAIT 2023), pp. 77–89.
- Wazzan, A., Ahmad, I., Macneil, S., Souvenir, R., 2024. Context or clutter? Efficiently matching objects across scenes. In: Proceedings of the 2024 International Conference on Multimedia Retrieval. Association for Computing Machinery, New York, NY, USA, pp. 404–413. <https://doi.org/10.1145/3652583.3658090>.
- Zhang, X., Wang, L., Su, Y., 2021. Visual place recognition: a survey from deep learning perspective. *Pattern Recogn.* 113, 107760. <https://doi.org/10.1016/j.patcog.2020.107760>. URL: <https://www.sciencedirect.com/science/article/pii/S00132032030563X>.