

Received July 23, 2021, accepted August 15, 2021, date of publication September 3, 2021, date of current version September 14, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3110335

Handgun Detection Using Combined Human Pose and Weapon Appearance

JESUS RUIZ-SANTAQUITERIA¹, ALBERTO VELASCO-MATA¹, NOELIA VALLEZ¹, GLORIA BUENO¹, JUAN A. ÁLVAREZ-GARCÍA², AND OSCAR DENIZ¹, (Senior Member, IEEE)

¹ETSI Industriales, University of Castilla-La Mancha, VISILAB, 13071 Ciudad Real, Spain

²Dep. Lenguajes y Sistemas Informáticos, ETSI Informática, University of Seville, 41012 Sevilla, Spain

Corresponding author: Jesus Ruiz-Santaquiteria (Jesus.RAlegre@uclm.es)

This work was supported in part by the Spanish Ministry of Economy and Business under Grant TIN2017-82113-C2-2-R, in part by the Autonomous Government of Castilla-La Mancha under Grant SBPLY/17/180501/000543, and in part by European Regional Development Fund (ERDF). The work of Jesus Ruiz-Santaquiteria was supported by Postgraduate Grant from the Spanish Ministry of Science, Innovation, and Universities under Grant PRE2018-083772.

ABSTRACT Closed-circuit television (CCTV) systems are essential nowadays to prevent security threats or dangerous situations, in which early detection is crucial. Novel deep learning-based methods have allowed to develop automatic weapon detectors with promising results. However, these approaches are mainly based on visual weapon appearance only. For handguns, body pose may be a useful cue, especially in cases where the gun is barely visible. In this work, a novel method is proposed to combine, in a single architecture, both weapon appearance and human pose information. First, pose keypoints are estimated to extract hand regions and generate binary pose images, which are the model inputs. Then, each input is processed in different subnetworks and combined to produce the handgun bounding box. Results obtained show that the combined model improves the handgun detection state of the art, achieving from 4.23 to 18.9 AP points more than the best previous approach.

INDEX TERMS CCTV surveillance, deep learning, handgun detection, human pose estimation.

I. INTRODUCTION

Video surveillance has come a long way in the past decades. Nowadays, public or private spaces such as train stations, airports, museums, banks or government institutional buildings have their own video surveillance systems. These systems are very useful for post-event investigations and also assisting the security personnel to manage crowds, being able to monitor different locations simultaneously. However, the main drawback of these solutions is the need for continuous monitoring by a human operator. The increasing number of areas controlled by video surveillance cameras, as well as factors inherent to human condition such as fatigue or loss of attention over time, make these systems rather inefficient [1], [2].

Related studies in this area show that early detection of security threats or risks is fundamental to mitigate the damage caused as much as possible [3]. Situations involving firearms such as handgun attacks, mass shootings, gunfire incidents on school grounds [4] or terrorist attacks [5] are representative

The associate editor coordinating the review of this manuscript and approving it for publication was Wei Liu.

examples of this kind of threats, which unfortunately have become rather common nowadays.

The development of intelligent systems capable of automatically detecting threats or risk situations involving firearms as soon as possible can provide important advantages in terms of security. Recently, thanks to the momentum achieved by the introduction of deep learning methodologies, remarkable results have been obtained in visual tasks such as image classification or object detection and segmentation. In the particular case of firearm detection, while the results obtained with these novel methods are promising, there are still substantial limitations when they are applied in new scenarios different to those used for training, especially an unacceptable number of false positives [6]. Furthermore, most of the proposed methods for automatic detection of firearms are based only on the appearance of the weapon in the image, without taking into account additional information that may help provide a more robust and accurate detection.

This work proposes the use of the human pose as complementary information to improve the performance of current handgun detectors based on deep learning. The human pose, defined as the relative position of the different joints and

limbs of the human body, is quite characteristic in shootings. On the other hand, the images obtained by CCTV cameras are not generally of high quality, due to their low resolution, the presence of noise or poor lighting conditions. Also, other factors such as distance to the camera, the small visual size of the weapon¹ or a total or partial occlusion can prevent the object from being detected [7]. Our hypothesis in this work is that contextual body information can help improve the robustness of the detection.

The contributions of this paper are as follows: (1) the development of a novel method for detecting hand-held firearms; (2) the performance evaluation of the proposed method in comparison with well-known appearance-based detection methods such as YOLOv3 and also other recent alternatives that consider human pose information and (3) assessment of the robustness of the method in environments under poor illumination conditions, with large camera distances and with different camera perspectives.

The rest of the article is organized as follows. Section II describes previous work related to the task of handgun detection. The datasets used in this study are detailed in Section III. In Section IV the proposed method is explained. The experiments carried out and results obtained are summarized in Section V. Finally, conclusions and future work are presented in Section VI.

II. RELATED WORK

X-ray scanning machines are widespread in public spaces such as airports, train stations or museums with the objective of finding concealed weapons in luggage. The X-ray generated images are manually analyzed by a security operator. In this context, several approaches based on classical vision methods were proposed to automate the detection process. The work Necessian *et al.* [8] introduced a detection system based on image segmentation and edge-based feature vectors. Xiao *et al.* [9] proposed a method based on Haar-like features and AdaBoost classifiers to automatically detect handguns in this kind of images. Also, 3D interest point descriptors have been studied for object classification in 3D baggage security computed tomography imagery [10], [11].

While X-ray imaging-based systems are useful to find weapons in travel bags or luggage, the scope of these solutions is very limited. Additionally, this kind of scanning machines are quite expensive. Using the RGB images captured by the CCTV video surveillance cameras to detect dangerous objects can be a more versatile and economical option. In this respect, several works related to the detection of weapons in RGB images through traditional machine learning methods have been proposed. Tiwari and Verma [12] proposed a method to detect weapons in RGB images which used color segmentation and the k-means algorithm to remove unrelated objects. Then, the Harris interest point detector and Fast Retina Keypoint (FREAK) is used to locate the handgun

¹Powerful weapons like rifles and shotguns, which have a larger visual size, have not been considered in this work.

in the segmented images. Later, Halima and Hosam [13] proposed another method to detect the presence of a handgun in an image. In this case, SIFT features are extracted from the collection of images and clustered by the k-means algorithm. Then, a word vocabulary based histogram is implemented and finally, a Support Vector Machine is used to decide whether the new image contains a weapon.

More recent deep learning based methods have been also applied to this task using different strategies. An important family of them is based on sliding windows. In this case, a large number of regions or windows of different sizes and aspect ratios are generated within the image (on the order of 10^4) and each one is classified individually by a neural network. Several studies have applied this technique to detect handguns in images similar to those captured by CCTV video surveillance cameras [14], [15]. The major drawback of this type of system is the high processing time required to classify these windows, making it difficult to use them in real time. Other solutions are based on region proposals, which instead of using all possible windows in an image select only actual candidates. The first techniques that have used CNNs in this context are the Region-based CNN family of methods [16], [17]. Verma and Dhillon [18] proposed a method based on the Faster-RCNN framework with a VGG-16 backbone as feature extractor trained with the IMFDB dataset [19] to detect hand-held arms. Both sliding window and Faster-RCNN methods were tested and compared by Olmos *et al.* [20] for handgun detection. Faster-RCNN pre-trained with VGG-16 architecture obtained the best results on a custom dataset of 3000 YouTube gun images. Finally, another common approach for detecting objects is based on the YOLO family of methods [21]–[23]. In these architectures a single deep neural network is applied once to the full image instead of multiple region proposals. The image is divided into fixed regions and probabilities and bounding boxes are predicted for each one. Several works have also recently applied YOLOv3 for detecting firearms with promising results [24], [25].

Human pose information has been recently used for handgun detection and threat assessment. Abruzzo *et al.* [26] proposed a method for identifying people and handguns in images and then evaluate the threat level of the person poses based on the their body posture. However, the main limitation of this work is that the handgun detection performance is limited by the handgun detector used (in this case YOLO). In the handgun detection step, no human pose information is considered. Basit *et al.* [27] proposed a method for classifying human-handgun pairs. As in the previous work, human and handgun are separately detected. Then, each detected human is paired with each detected handgun and, finally, a neural network is trained to classify these paired human-handgun bounding boxes into two classes: “carrying handgun” and “not carrying handgun”. This method can be used to remove false handgun detections, but again the detection performance is limited by the handgun detector used and the human pose cannot help to reduce the number of false negatives. More




FIGURE 1. Sample images from Gun Movies Database.

recently, Salido *et al.* [28] analyzed how including body pose information (skeleton keypoints and limbs retrieved by a pose detector) in the input images, as a preprocessing step, can improve the handgun detection performance.

In the closest work to ours, an approach to improve a handgun detector through the integration with the human pose was introduced recently in Velasco-Mata *et al.* [29]. This method used a visual heatmap representation of both the pose and the weapon location, using convolutional layers to obtain a final grayscale image that indicates potential handgun regions on the image.

III. MATERIALS

This section describes the datasets used for assessing the performance of the proposed method. In order to consider different contexts and image features, the images have been collected from different sources, such as public handgun datasets, YouTube clips and even synthetic images obtained from video games.

A. PUBLIC HANDGUN DATASETS

The proposed method is intended to be applied in CCTV surveillance systems on a wide variety of scenarios. Unfortunately, most public handgun datasets contain weapon profile images occupying the whole image and with homogeneous background, which are quite different from the type of images captured with surveillance cameras. Surveillance scenarios are typically characterized by a large distance between the subjects recorded and the camera, low image resolution or quality, and poor lighting conditions, among others. Salazar Gonzalez *et al.* [30] recently introduced a new dataset composed of CCTV images from a real video surveillance system and synthetic images generated with the Unity game engine. However, the CCTV images in this dataset are not realistic enough, showing unnatural poses for handgun attacks or mass shootings. On the other hand, it is possible to find some public datasets or parts thereof which are realistic enough. The first dataset used in our study is composed of 665 images of size 640×480 extracted from videos of the Guns Movies Database [14]. In these clips a man is holding a handgun in a few shooting poses in an indoor room. Camera distance, image resolution and illumination conditions are a good representation of CCTV scenarios. In Figure 1 two example images from this dataset are shown.

Also, 300 images of size 512×512 were obtained from the publicly available Monash Guns Dataset [31] for test




FIGURE 2. Sample images from Monash Guns Dataset.

purposes. These images show different CCTV scenarios with people holding handguns in several body poses. In Figure 2 two example images from this realistic dataset are shown.

B. YouTube VIDEOS

YouTube is another useful source to find videos of people carrying or holding weapons and/or shooting. As in the previous case, it is difficult to find clips of real CCTV footage showing handguns. Nevertheless, there are videos of shooting practice sessions which are suitable for our purposes. This dataset is composed of 952 images of size 1920×1080 extracted from 12 YouTube clips. In these videos there are different camera locations, background scenarios, shooting poses and lighting conditions. Figure 3 shows two example images from this dataset.



FIGURE 3. Sample images from YouTube dataset.

C. SYNTHETIC VIDEO GAME IMAGES

Video games can be also used to create new data for this task. Through specific video games it is possible to recreate representative situations or scenarios and then extract videos or images. In this case, a synthetic dataset was created with the popular shooter video game Watch Dogs 2 on a PC platform. Using the novel NVIDIA Ansel feature,² ingame videos can be recorded from different camera locations, distances or angles. In this way, 4 video sequences were recorded, performing a full camera rotation around the main character with two different heights in various shooting animations. Finally, 650 images of size 3840×2160 were obtained from these video sequences. In Figure 4 two example images of this dataset are shown.

D. DATA AUGMENTATION AND DATASET SPLIT

The use of a large and representative dataset is essential to achieve good performance in novel object detectors based on deep learning and CNNs. YouTube clips or synthetic video

²<https://developer.nvidia.com/ansel>



FIGURE 4. Sample images from Watch Dogs 2 dataset (©2016 Ubisoft Entertainment).

game images can be a solution, but the manual labelling process is a time consuming task, limiting the number of images that can be effectively generated.

Novel deep learning methods need a huge amount of data to be correctly trained, because of the large number of parameters and model complexities. To deal with this problem, data augmentation is a common practice that helps to increase the size and variability of the dataset by applying a set of transformations to the original data. In our case, for each image included in the dataset, another one was generated performing a horizontal flip. In this way, the number of images was doubled, generating new shooting poses.

Moreover, transfer learning is another technique commonly used in classification or object detection tasks to circumvent the need for large training datasets. It is based on adapting the useful features learned in a previous training process with a large public dataset, such as COCO [32] or ImageNet [33], to a new domain. In this way, starting the new training with the previously learned parameters instead of randomly initialized weights, allows us to obtain good results for the specific task even with small specific datasets.

Finally, for this work, the dataset used for training the proposed method after the data augmentation procedure is composed of a total of 3000 randomly selected images, obtained from Guns Movies Database, Watch Dogs 2 and YouTube datasets, including a total of 3160 handguns. The validation set is formed by a total of 300 images obtained from the same datasets (Guns Movies Database, Watch Dogs 2 and YouTube), including 306 handguns. For test, two different sets have been created. The first one (test set A) is formed by a total of 300 images of the YouTube database, including a total of 297 handguns. The second one (test set B) is composed of 300 images from the Monash Guns Dataset, containing another 300 handguns. The test images were selected in this way to evaluate the performance of the methods against the most realistic and heterogeneous datasets available. None of the test images was either in the training or validation set. A summary of the dataset composition is presented in Table 1.

IV. METHODOLOGY

In this section, the different steps involved in the proposed method are detailed, starting from the input image down to the final handgun detections.

A. HUMAN POSE ESTIMATION

The first step consists of collecting the human pose information found in the input image. This was done with the

OpenPose framework [34]. OpenPose is an open-source multi-person pose estimator which is able to predict the 2D keypoints as well as keypoint associations, keeping a high accuracy and a low inference time. In this step, a set of 25 2D keypoints are predicted for each person in the image, along with predicted confidence for each one of them. These keypoints include the necessary human body position information (neck, shoulders, elbows, wrists, etc.) to define the pose of each person.

B. HAND REGION EXTRACTION

In the second step, using the collected pose information, the hand regions for each detected person are inferred and extracted. The elbow and wrist positions, as well as the distances and directions between them are used to generate a set of bounding boxes around all of the hand regions in the input image (see Figure 5).

The confidence score given by OpenPose for each keypoint is applied to filter wrong or inaccurate detections, and an intersection over union (IoU) threshold between the predicted bounding boxes is checked to prevent overlapping areas (e.g., a handgun held with both hands is considered as a single region, since both bounding boxes are overlapping).

C. HAND REGION CLASSIFICATION

For this stage, a convolutional neural network was trained to classify the previously generated hand regions into handgun or no-handgun areas, based on whether there is a handgun inside the region (see Figure 6). The selected network was Darknet-53, the backbone feature extractor used in the YOLOv3 object detector. Henceforth, this hand region classifier will be denoted as HRC (Hand Region Classifier).

The dataset used for training the hand region classifier was composed of 6177 images, generated from the hand areas extracted from the 3000 training images described in Section III. These regions were automatically labelled by comparing the hand areas and the ground truth handgun locations. For this, instead of the IoU score, we followed the overlap measure proposed in Velasco-Mata *et al.* [29]: intersection over minimum area (IoMin), see Equation 1. Usually, ground truth handgun locations are smaller than hand bounding boxes. This metric allows a better overlap measurement in this particular scenario, since bounding boxes of different sizes are not penalized. More details are given in [29].

$$IoMin(A, B) = \frac{area(A \cap B)}{\min(area(A), area(B))} \quad (1)$$

If the regions overlap with a 0.5 IoMin threshold, the hand area is labelled as a handgun area. On the other hand, if there is no overlap or the threshold is below 0.5 IoMin, the hand area is labelled as no-handgun.

Each hand region was also resized to a fixed size of 256×256 . The model was trained with a batch size of 4 in 60 epochs using the Adam optimization algorithm and the categorical cross entropy as loss function.

TABLE 1. Number of images for each dataset after the data agumentation process.

	Training	Validation	Test set A	Test set B	Total
Gun Movies Database	1064	113	0	0	1177
Monash Guns Dataset	0	0	0	300	300
Watch Dogs 2	1071	108	0	0	1179
YouTube	865	79	300	0	1244
Total	3000	300	300	300	3900




FIGURE 5. Pose estimation and hand region extraction.




FIGURE 6. Hand region classification step.

D. POSE COMBINATION METHOD

A further modification of the HRC method described in the previous section was considered. In this case, the network is modified to combine the hand region image with the human pose information obtained with OpenPose. This was done to help the region classifier by exploiting correlation with the individual’s pose information. Pose data is used to create binary images of fixed size 512×512 for each detected person in the input image, drawing the keypoints and the connections between them. A normalization procedure is also applied to focus only on the relative position between the keypoints, removing variable factors such as camera distance and absolute position in the image. For this, the original neck keypoint j_0 is taken as reference and the distance between this point and the lumbar spine keypoint j_1 is used as the scale factor for the normalization. In this way, the new keypoints k_n are calculated following Equation 2:

$$k_n = \frac{j_n - j_0}{|j_0 j_1|} \tag{2}$$

where j_n is the original 2D point and $|j_0 j_1|$ is the distance between j_0 and j_1 .

In Figure 7b an example of a generated pose image is shown, along with the original image (Figure 7a). This binary pose image along with the original hand region image are the inputs to the new classifier. However, note that each pose image is related to two different hand regions (the two hands of an individual), and the pose image generated would be the same in the two cases. The problem then is that one of these

regions could be labelled as handgun and the other as no-handgun, and the net effect of the additional pose information would be ignored (as the two regions have the same pose image). To prevent this, the pose image is divided into two parts of size 256×512 as shown in Figure 8, selecting as input the half in which the hand region is included. Thus, for each detected hand region there are two network inputs, the hand region itself and the pose image half corresponding to this hand region. For those cases in which the handgun is held with the two hands and the bounding boxes are overlapping, as the case shown in Figure 7, we ensure that only a single hand region is taken into account, along with the corresponding pose image half.

The whole network architecture, henceforth named as HRC+P (Hand Region Classifier + Pose data), is divided into two main branches. The first one is the hand region classifier (HRC). On the other hand, the processing of the pose image is carried out by another custom subnetwork. The last feature vectors of the two branches are then joined in a single feature vector connected to the final classification layer. During the training step of the whole framework, the HRC and pose branches are optimized simultaneously. In this way, the model is capable of learning the optimal combination of handgun appearance with human pose information to improve the classification performance. The whole architecture is shown in Figure 9.

E. BOUNDING BOX PREDICTION

The last step of the proposed method consists of generating the handgun predictions in the image. Each hand region



FIGURE 7. Pose classification example.




FIGURE 8. Pose processing steps, including normalization and division in halves.

of each detected person is passed through the classification network to obtain a class label (handgun vs no-handgun). Then, the bounding boxes of the regions classified as handgun are included in the output list of predicted handguns (see Figure 10).

V. RESULTS

This section presents the results obtained in the tests carried out to evaluate the performance of the proposed method. In general, object detection models are evaluated using Precision, Recall, and Average Precision (PASCAL VOC AP50) metrics [35]. In fact, these metrics are based on True Positives (TP), False Positives (FP) and False Negatives (FN). These values are calculated taking into account the overlap between the ground truth bounding boxes and those predicted by the detector. In the same way as in the automatic labeling process for the training of the hand region classifier (Subsection IV-C), the IoMin is the selected criterion to calculate the overlap between the predicted bounding boxes and the ground truth data, due to the size difference between them.

The proposed pose-combined approach (HRC+P) has been compared to three different handgun detectors:

- YOLOv3 [23]: YOLOv3 is one of the fastest and most accurate deep learning-based object detectors. The Darknet-53 CNN backbone is used as feature extractor, which provides an interesting baseline for comparison.

- Basit et al. [27]: This work proposes a method to classify person-handgun pairs detected in an image, between people carrying handguns and those who do not.
- Velasco-Mata et al. [29]: This recent work, the most similar to ours, proposes a method to improve a handgun detector based on a visual heatmap representation of both pose and weapon location.
- Salido et al. [28]: This method overlays body pose information retrieved by the OpenPose framework to the input images. In this way, CNN-based detectors can learn the association of a handgun location with the visual patterns of the pose skeletons included in the images.
- The proposed hand region classifier without pose information (HRC): To check the effect of including the 2D human pose information in the hand region classifier, the hand region processing branch without pose combination is taken for comparison.

All methods were trained and tested using the datasets described in Section III.

A. TEST SET A - ORIGINAL DATA

The results obtained in the test set A are summarized in Table 2. Precision and Recall values are calculated with a 0.5 prediction score threshold. Also, Precision-Recall curves of all methods are shown in Figure 11.

The highest AP score is achieved with the HRC+P method, the pose-combined version of the proposed approach. The AP

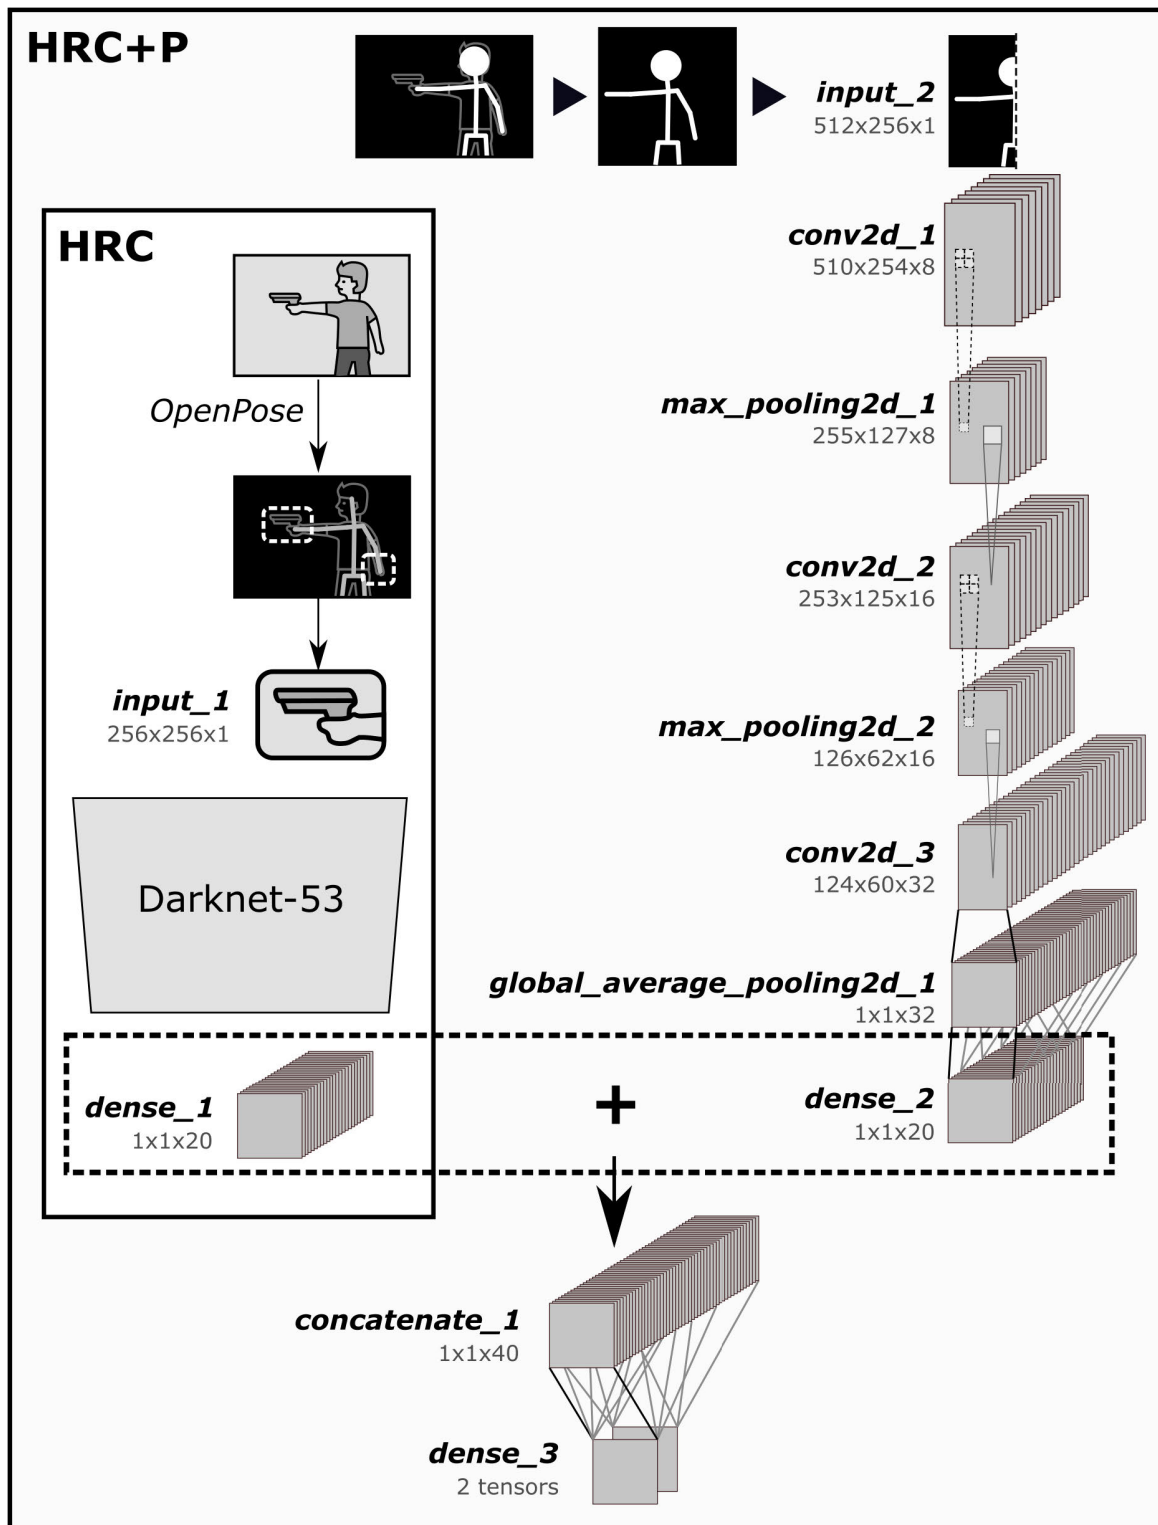


FIGURE 9. Network architecture.

obtained is approximately 5% higher than the no-combined version (HRC). HRC+P is able to detect the largest number of handguns in the dataset. However, the least number of false positives is obtained with the HRC approach.

To better assess the performance of the models under different conditions, two additional versions of the test set A were generated. The results are described in what follows.

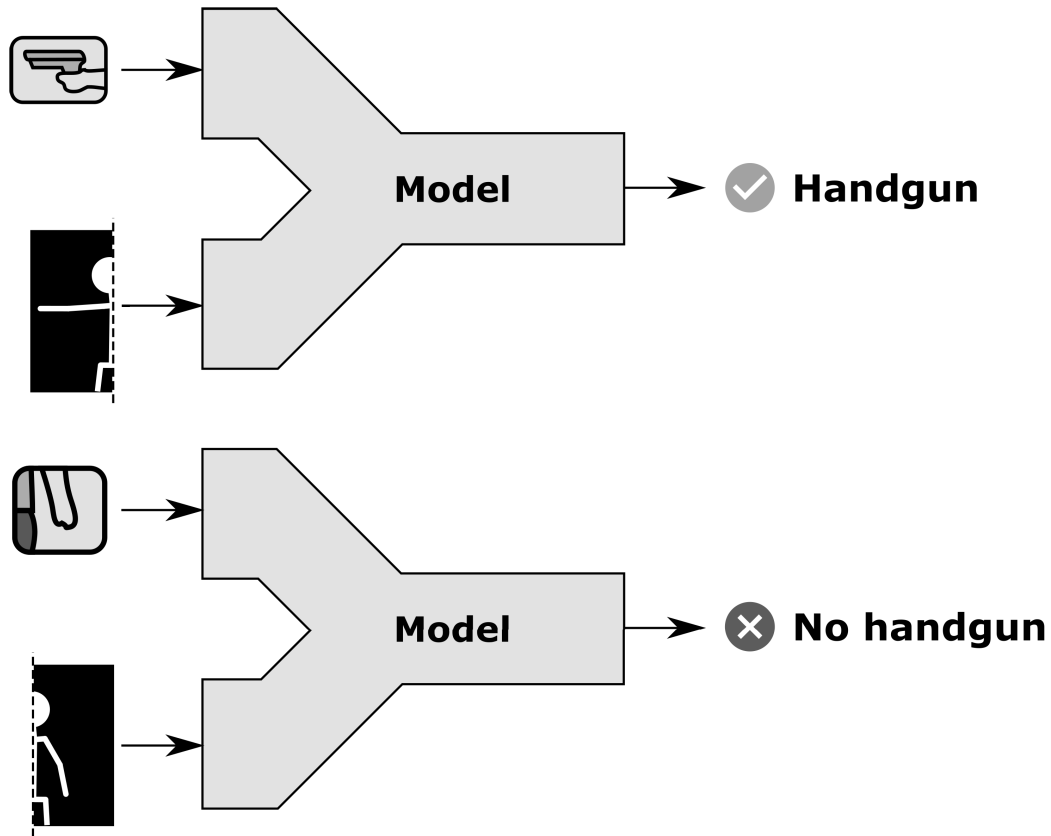


FIGURE 10. Bounding box prediction step.




FIGURE 11. Precision-Recall curves obtained for test set A with the original images.

B. TEST SET A - DARK DATA

A darkened version of the test set A was generated to simulate the performance of the trained models under poor illumination conditions. This scenario was obtained by modifying the Value component in the HSV color space for all images in the test set A. Figure 12 shows an example image.

The results obtained for these darkened images are summarized in Table 3. Precision and Recall values are calculated with a 0.5 prediction score threshold and Precision-Recall curves of all methods are shown in Figure 13.

Again, in these modified images the HRC+P method obtains the highest AP score, with more than 4% of improvement over the second best, showing that the proposed pose-combined method can be useful in unfavorable lighting conditions.

C. TEST SET A - FAR DATA

Camera distance can also be a relevant factor in detecting small objects such as handguns, especially in cases where only the visual appearance of the object is used for detection. To test this scenario, another version of the test set A has been generated, reducing image size by half and filling the rest of the image with black pixels. In Figure 14, an example of this transformation is presented.

The results obtained for these far images are summarized in Table 4. Precision and Recall values are calculated with a 0.5 prediction score threshold and Precision-Recall curves of all methods are shown in Figure 15.

The HRC+P model also presents the highest AP score in this last test scenario, improving in 6% the performance of the HRC method. The results show that this reduction in the size of the objects present in the image significantly affects methods that are based exclusively on the appearance



(a) Original image.



(b) Darkened version.

FIGURE 12. Test set A - Dark data.

TABLE 2. Evaluation metrics for test set A with the original images.

Method	Precision 0.5	Recall 0.5	AP
YOLOv3 [23]	0.7937	0.5959	61.42
Basit et al. [27]	0.8544	0.4545	49.52
Velasco-Mata et al. [29]	0.8692	0.6936	64.06
Salido et al. [28]	0.9351	0.7272	76.34
HRC	0.9733	0.7374	79.11
HRC+P	0.9188	0.8383	83.85

TABLE 3. Evaluation metrics for test set A with the dark images.

Method	Precision 0.5	Recall 0.5	AP
YOLOv3 [23]	0.8043	0.4983	51.86
Basit et al. [27]	0.8561	0.4007	45.58
Velasco-Mata et al. [29]	0.8831	0.6869	62.88
Salido et al. [28]	0.9378	0.6599	72.02
HRC	0.9751	0.6599	69.64
HRC+P	0.8937	0.7643	76.25

of the object, such as YOLOv3. Pose-based methods show even higher metrics. This can be explained due to the fact that the pose estimation is not severely affected by this larger distance to the camera, being able to accurately detect the pose keypoints in the image.

D. TEST SET B - MONASH DATASET

For the last experiment, the realistic Monash Guns Dataset [31] has been used to test the methods compared in this work. This dataset shows people holding handguns in a variety of real-world CCTV surveillance environments. The results obtained for these test images are summarized in Table 5. Precision and Recall values are calculated with a 0.5 prediction score threshold and Precision-Recall curves of all methods are shown in Figure 16.

For this test set, all methods show a significantly lower performance. This can be explained by the differences with respect to the training images in terms of lighting, camera perspective and image distortions. However, the proposed pose-combined method (HRC+P) still shows the highest AP score, improving the no-combined version (HRC) in approximately 10%.




FIGURE 13. Precision-Recall curves obtained for test set A with the dark images.

E. EXAMPLE IMAGES

The two proposed approaches (HRC and HRC+P) present the best average performance in all studied scenarios. The pose-combined method (HRC+P) obtains better results in terms of Recall and AP. In Figure 17 two example images are shown to illustrate these results. In both examples the handgun is not detected with the HRC method. In the first example (Figure 17a and Figure 17b) the handgun is almost completely occluded and in the second example (Figure 17c and Figure 17d) the handgun area is blurry. Conversely,



(a) Original image.



(b) Far version.

FIGURE 14. Test set A - Far data.

TABLE 4. Evaluation metrics for test set A with the distance-simulated images.

Method	Precision 0.5	Recall 0.5	AP
YOLOv3 [23]	0.8766	0.4545	46.93
Basit et al. [27]	0.9350	0.3872	46.28
Velasco-Mata et al. [29]	0.8092	0.7710	62.99
Salido et al. [28]	0.8449	0.6970	68.90
HRC	0.9657	0.7576	81.31
HRC+P	0.9123	0.8754	87.80

TABLE 5. Evaluation metrics for test set B - Monash data.

Method	Precision 0.5	Recall 0.5	AP
YOLOv3 [23]	0.8571	0.1800	20.11
Basit et al. [27]	0.8929	0.0833	9.07
Velasco-Mata et al. [29]	0.7848	0.2067	17.17
Salido et al. [28]	0.8833	0.1767	21.30
HRC	0.9683	0.2033	24.72
HRC+P	0.9018	0.3367	34.68

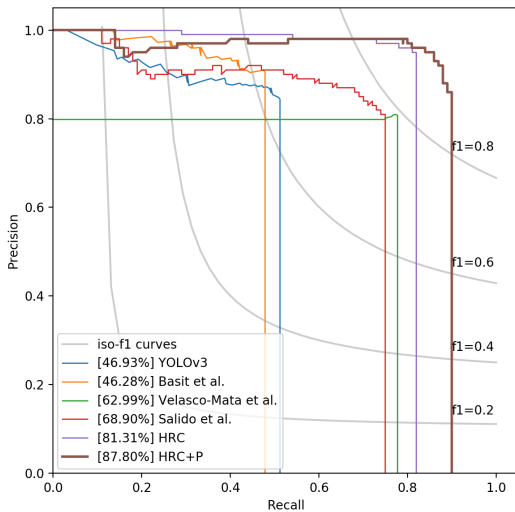


FIGURE 15. Precision-Recall curves obtained for test set A with the distance-simulated images.




FIGURE 16. Precision-Recall curves obtained for test set B with the Monash images.

the HRC+P method is able to locate both of them thanks to the clear shooting poses.

In terms of Precision HRC obtains better results. Figure 18 shows an example of a false positive detection in the HRC+P

method. In this case, both right hand object and body pose of the second subject have caused an incorrect detection. On the

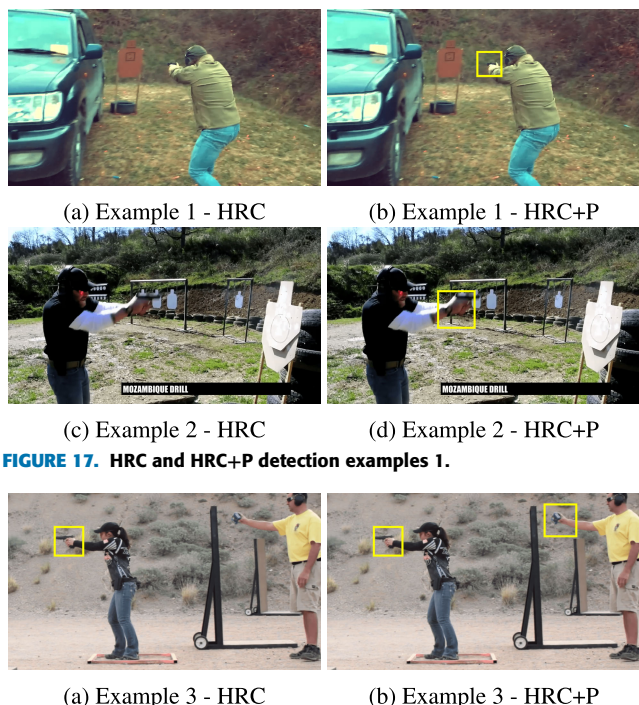


FIGURE 17. HRC and HRC+P detection examples 1.

FIGURE 18. HRC and HRC+P detection examples 2.

other hand, the HRC approach can classify all hand regions in the image correctly.

VI. CONCLUSION

The 2D human pose is widely used in tasks such as action or gesture recognition. However, for the detection of threats or dangerous objects such as firearms, most of the proposed methods are based only on the visual appearance of the objects, without taking into account the human pose or another additional information.

In this work a novel method that combines in the same architecture the visual appearance of the handgun with the 2D human pose information is proposed. There are certain situations in which the object cannot be viewed correctly due to camera distance, poor lighting conditions or partial or total occlusion. In these situations, the human body pose helps in detecting the presence of handguns that would not be detected without this additional information. On the other hand, as the pose information is used to classify only the hand regions of the people detected, it is possible to remove false positives that may appear in other locations of the image.

The tests performed with the different datasets show that the proposed method using the pose combination obtains better results in all cases. Especially interesting is the fact that metrics in the reduced size images are even higher than in the case of the original size images.

Automatic and real time handgun detection in CCTV video surveillance images is still an open problem and there is room for improvement. The authors hope that the proposed work can be used as inspiration for new approaches based on 2D human pose information to improve the overall detector performance in this kind of application.

Finally, note that in real scenarios common hand-held objects such as cell phones, keys or wallets may be an important source of false positives or misclassifications. In future work this aspect will be addressed with more specific methods.

REFERENCES

- [1] M. Mendelsohn, "Buyer beware," *Hair Therapy Transplantation*, vol. 5, no. 2, pp. 18–26, 2015.
- [2] S. A. Velastin, B. A. Boghossian, and M. A. Vicencio-Silva, "A motion-based image processing system for detecting potentially dangerous situations in underground railway stations," *Transp. Res. C, Emerg. Technol.*, vol. 14, no. 2, pp. 96–113, Apr. 2006.
- [3] F. Enríquez et al., "Vision and crowdsensing technology for an optimal response in physical-security," in *Computational Science—ICCS 2019* (Lecture Notes in Computer Science), vol. 11540, J. Rodrigues et al., Eds. Cham, Switzerland: Springer, 2019, doi: 10.1007/978-3-030-22750-0_2.
- [4] (2020). *Gunfire on School Grounds in the United States*. Accessed: Jul. 20, 2021. [Online]. Available: <https://everytownresearch.org/gunfire-in-school/#ns>
- [5] R. A. Tessler, S. J. Mooney, C. E. Witt, K. O'Connell, J. Jenness, M. S. Vavilala, and F. P. Rivara, "Use of firearms in terrorist attacks: Differences between the United States, Canada, Europe, Australia, and New Zealand," *JAMA Internal Med.*, vol. 177, no. 12, pp. 1865–1868, 2017.
- [6] N. Vázquez, G. Bueno, and O. Déniz, "False positive reduction in detector implantation," in *Artificial Intelligence in Medicine* (Lecture Notes in Computer Science), vol. 7885, N. Peek, R. M. Morales, and M. Peleg, Eds. Berlin, Germany: Springer, 2013, doi: 10.1007/978-3-642-38326-7_28.
- [7] A. Castillo, S. Tabik, F. Pérez, R. Olmos, and F. Herrera, "Brightness guided preprocessing for automatic cold steel weapon detection in surveillance videos with deep learning," *Neurocomputing*, vol. 330, pp. 151–161, Feb. 2019.
- [8] S. Nercessian, K. Panetta, and S. Agaian, "Automatic detection of potential threat objects in X-ray luggage scan images," in *Proc. IEEE Conf. Technol. Homeland Secur.*, May 2008, pp. 504–509.
- [9] Z. Xiao, X. Lu, J. Yan, L. Wu, and L. Ren, "Automatic detection of concealed pistols using passive millimeter wave imaging," in *Proc. IEEE Int. Conf. Imag. Syst. Techn. (IST)*, Sep. 2015, pp. 1–4.
- [10] G. Flitton, T. P. Breckon, and N. Megherbi, "A comparison of 3D interest point descriptors with application to airport baggage object detection in complex CT imagery," *Pattern Recognit.*, vol. 46, no. 9, pp. 2420–2436, Sep. 2013.
- [11] G. Flitton, A. Mouton, and T. P. Breckon, "Object classification in 3D baggage security computed tomography imagery using visual codebooks," *Pattern Recognit.*, vol. 48, no. 8, pp. 2489–2499, Aug. 2015.
- [12] R. K. Tiwari and G. K. Verma, "A computer vision based framework for visual gun detection using Harris interest point detector," *Procedia Comput. Sci.*, vol. 54, pp. 703–712, Jan. 2015.
- [13] N. B. Halima and O. Hosam, "Bag of words based surveillance system using support vector machines," *Int. J. Secur. Appl.*, vol. 10, no. 4, pp. 331–346, Apr. 2016.
- [14] M. Grega, A. Matiolański, P. Guzik, and M. Leszczuk, "Automated detection of firearms and knives in a CCTV image," *Sensors*, vol. 16, no. 1, p. 47, 2016.
- [15] F. Gelana and A. Yadav, "Firearm detection from surveillance cameras using image processing and machine learning techniques," in *Smart Innovations in Communication and Computational Sciences* (Advances in Intelligent Systems and Computing), vol. 851, S. Tiwari, M. Trivedi, K. Mishra, A. Misra, and K. Kumar, Eds. Singapore: Springer, 2019, doi: 10.1007/978-981-13-2414-7_3.
- [16] R. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1440–1448.
- [17] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [18] G. K. Verma and A. Dhillon, "A handheld gun detection using faster R-CNN deep learning," in *Proc. 7th Int. Conf. Comput. Commun. Technol. (IC3CT)*, 2017, pp. 84–88.
- [19] (2020). *Database*. Accessed: Jul. 20, 2021. [Online]. Available: http://www.imfdb.org/wiki/Main_Page
- [20] R. Olmos, S. Tabik, and F. Herrera, "Automatic handgun detection alarm in videos using deep learning," *Neurocomputing*, vol. 275, pp. 66–72, Jan. 2018.

- [21] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 779–788.
- [22] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 7263–7271.
- [23] A. Farhadi and J. Redmon, "YOLOv3: An incremental improvement," in *Proc. Comput. Vis. Pattern Recognit.*, 2018, pp. 1–6.
- [24] A. Warsi, M. Abdullah, M. N. Husen, M. Yahya, S. Khan, and N. Jawaid, "Gun detection system using YOLOv3," in *Proc. IEEE Int. Conf. Smart Instrum., Meas. Appl. (ICSIMA)*, Aug. 2019, pp. 1–4.
- [25] R. F. D. A. Kanehisa and A. D. A. Neto, "Firearm detection using convolutional neural networks," in *Proc. ICAART*, 2019, pp. 707–714.
- [26] B. Abruzzo, K. Carey, C. Lowrance, E. Sturzinger, R. Arnold, and C. Korpela, "Cascaded neural networks for identification and posture-based threat assessment of armed people," in *Proc. IEEE Int. Symp. Technol. Homeland Secur. (HST)*, Nov. 2019, pp. 1–7.
- [27] A. Basit, M. A. Munir, M. Ali, N. Werghi, and A. Mahmood, "Localizing firearm carriers by identifying human-object pairs," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2020, pp. 2031–2035.
- [28] J. Salido, V. Lomas, J. Ruiz-Santaquiteria, and O. Deniz, "Automatic handgun detection with deep learning in video surveillance images," *Appl. Sci.*, vol. 11, no. 13, p. 6085, Jun. 2021.
- [29] A. Velasco-Mata, J. Ruiz-Santaquiteria, N. Vallez, and O. Deniz, "Using human pose information for handgun detection," *Neural Comput. Appl.*, vol. 1, pp. 1–4, Jul. 2021.
- [30] J. L. S. González, C. Zaccaro, J. A. Álvarez-García, L. M. S. Morillo, and F. S. Caparrini, "Real-time gun detection in CCTV: An open problem," *Neural Netw.*, vol. 132, pp. 297–308, Dec. 2020.
- [31] J. Lim, M. I. Al Jobayer, V. M. Baskaran, J. M. Lim, J. See, and K. Wong, "Deep multi-level feature pyramids: Application for non-canonical firearm detection in video surveillance," *Eng. Appl. Artif. Intell.*, vol. 97, Jan. 2021, Art. no. 104094.
- [32] T. Y. Lin et al., "Microsoft COCO: Common objects in context," in *Computer Vision—ECCV 2014 (Lecture Notes in Computer Science)*, vol. 8693, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham, Switzerland: Springer, 2014, doi: [10.1007/978-3-319-10602-1_48](https://doi.org/10.1007/978-3-319-10602-1_48).
- [33] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Proc. CVPR*, Jun. 2009, pp. 248–255.
- [34] Z. Cao, G. Hidalgo, T. Simon, S. E. Wei, and Y. Sheikh, "OpenPose: Realtime multi-person 2D pose estimation using part affinity fields," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 1, pp. 172–186, Jan. 2021.
- [35] R. Padilla, S. L. Netto, and E. A. B. da Silva, "A survey on performance metrics for object-detection algorithms," in *Proc. Int. Conf. Syst., Signals Image Process. (IWSSIP)*, Jul. 2020, pp. 237–242.



JESUS RUIZ-SANTAQUITERIA received the bachelor's and master's degrees in computer science from the University of Castilla-La Mancha, Spain, in 2016 and 2018, respectively. He is currently pursuing the Ph.D. degree with UCLM. His research interests include artificial intelligence, computer vision, and machine learning.




ALBERTO VELASCO-MATA received the bachelor's degree in computer science, in 2020. He is currently pursuing the master's degree with the University of Castilla-La Mancha, Spain. He has been working with the VISILAB Group, UCLM, since 2017, as an Assistant Researcher, with a focus on computer vision and neural networks.



NOELIA VALLEZ received the B.S. degree in computer science and the Ph.D. degree from the University of Castilla-La Mancha, in 2009 and 2015, respectively. She is currently working with the VISILAB Group, UCLM, as a Lecturer, and an Assistant Researcher. Her research interests include image processing and machine learning.



GLORIA BUENO received the Ph.D. degree in machine vision from Coventry University, in 1998. She has been a Lecturer and the Principal Research with the School of Engineering, UCLM, Ciudad Real, Spain, since 2002. She has carried on her research activities at different research centres, such as the Centro de Estudios e Investigaciones Técnicas de Guipuzkoa—San Sebastián (E), the Centre National de la Recherche Scientifique, the Hôpitaux Civil & Telecommunication School, University Louis Pasteur, Strasbourg (FR), and Gilbert Gilkes & Gordon Technology, Kendal, U.K. She is leading different national and the European research project in biomedical image processing. Her current research interests include signal and image processing, modeling, and artificial intelligence.



JUAN A. ÁLVAREZ-GARCÍA received the degree in computer engineering from the University of Seville, Spain. He is currently an Associate Professor with the Department of Languages and Computer Systems, University of Seville. His research interests include human behavior analysis and activity recognition using wearable/mobile devices or computer vision from CCTV.



OSCAR DENIZ (Senior Member, IEEE) is currently a Full Professor at UCLM and contributes to VISILAB. He is the author of more than 50 refereed papers in journals and conferences. He has published two books on OpenCV. His research interests include computer vision and pattern recognition. He has served as a Visiting Researcher for Carnegie Mellon University, USA, Imperial College London, U.K., and Leica Biosystems, Ireland. He is affiliated with AAAI, SIAMI, CEA-IFAC, AEPIA, AERFAI-IAPR, and The Computer Vision Foundation. He serves as an Academic Editor for *PLoS ONE* journal. He has been a Coordinator of European Project "Eyes of Things," an Innovation Action funded by the H2020 Program. He is a Reviewer/Technical Expert for EU Programs, such as Eurostars.

...