



Semantic-driven multi-camera pedestrian detection

Alejandro López-Cifuentes¹  · Marcos Escudero-Viñolo¹ · Jesús Bescós¹ · Pablo Carballeira¹

Received: 12 July 2021 / Revised: 28 February 2022 / Accepted: 5 March 2022 /
Published online: 9 April 2022
© The Author(s) 2022

Abstract

In the current worldwide situation, pedestrian detection has reemerged as a pivotal tool for intelligent video-based systems aiming to solve tasks such as pedestrian tracking, social distancing monitoring or pedestrian mass counting. Pedestrian detection methods, even the top performing ones, are highly sensitive to occlusions among pedestrians, which dramatically degrades their performance in crowded scenarios. The generalization of multi-camera setups permits to better confront occlusions by combining information from different viewpoints. In this paper, we present a multi-camera approach to globally combine pedestrian detections leveraging automatically extracted scene context. Contrarily to the majority of the methods of the state-of-the-art, the proposed approach is scene-agnostic, not requiring a tailored adaptation to the target scenario—e.g., via fine-tuning. This noteworthy attribute does not require *ad hoc* training with labeled data, expediting the deployment of the proposed method in real-world situations. Context information, obtained via semantic segmentation, is used (1) to automatically generate a common area of interest for the scene and all the cameras, avoiding the usual need of manually defining it, and (2) to obtain detections for each camera by solving a global optimization problem that maximizes coherence of detections both in each 2D image and in the 3D scene. This process yields tightly fitted bounding boxes that circumvent occlusions or miss detections. The experimental results on five publicly available datasets show that the proposed approach outperforms state-of-the-art multi-camera pedestrian detectors, even some specifically trained on the target scenario, signifying the versatility and robustness of the proposed method without requiring ad hoc annotations nor human-guided configuration.

✉ Alejandro López-Cifuentes
alejandro.lopezc@uam.es

Marcos Escudero-Viñolo
marcos.escudero@uam.es

Jesús Bescós
j.bescos@uam.es

Pablo Carballeira
pablo.carballeira@uam.es

¹ Video Processing and Understanding Lab (VPULab), Universidad Autónoma de Madrid, Madrid, Spain

Keywords Pedestrian detection · Multi-camera systems · Semantic segmentation · Video surveillance

1 Introduction

In the current worldwide situation, pedestrian detection has reemerged as a pivotal tool for intelligent video-based systems aiming to solve tasks such as pedestrian tracking, social distancing monitoring or pedestrian mass counting. Automatic people detection is generally considered a solid and mature technology able to operate with nearly human accuracy in generic scenarios [10, 16, 30]. However, the handling of severe occlusions is still a major challenge [28]. Occlusions occur due to the projection of the 3D objects onto a 2D image plane. Although recent deep-learning-based methods are able to cope with partial occlusions, the detection process fails when only a small part or no part of the person is visible. To cope with severe occlusions, a potential solution is the use of additional cameras: If they are adequately positioned, the different points of view might allow for disambiguation.

Disambiguation is generally achieved by projecting every camera's detections on a common reference plane. The ground plane is usually the preferred option as it constitutes a common reference in which people's height can be disregarded. Per-camera detections can then be combined on the ground plane to refine and complete pedestrian detection. However, there are several challenges to be addressed during this combination or fusion process. Among the striking ones are: the convenience to define common visibility areas where cameras' views overlap, and how to cope with camera calibration errors and persons' self-occlusions. See Fig. 1 for visual examples of these challenges, which we detail below:

In multi-camera approaches a common strategy is to define an **operational area** or area of interest \mathcal{AOI} on the ground plane. This area represents the overlapping field-of-view of all the involved cameras. It can be used to reduce the impact of calibration errors in the process and to generally ease the fusion of per-camera detections. This area is generally manually defined for each scenario, precluding the automation of the process.

Scene calibration is a well-known task [17] which can be performed either manually or using automatic calibration methods based on image cues. In both cases, small perturbations in the calibration process may cause uncertainty in the fusion of the detections on the ground plane. The impact of calibration errors increases with the distance to the camera: Generally, calibration is more accurate for pixels belonging to objects close to the camera.

Self-occlusions are caused by the intrinsic three-dimensional nature of people, resulting in the occlusion of some human parts by some others. If the visible parts are different for different cameras and these are used to project a person location on the ground plane, the cameras' projections will diverge, hindering their fusion.

To cope with these challenges, in this paper we present a multi-camera pedestrian detection method which is driven by semantic information automatically extracted from the 2D images and transferred to the 3D ground plane and includes the following novel contributions:

1. A novel approach to globally combine pedestrian detections in a multi-camera scenario by creating connected components in a graph representation of detections.
2. An height-adaptive optimization algorithm which uses semantic cues to globally refine the location and size of people detections by aggregating information from all the cameras.

The proposed method is applied over an operational area in the ground plane, which is automatically defined by an adaptation of the method described in [26].

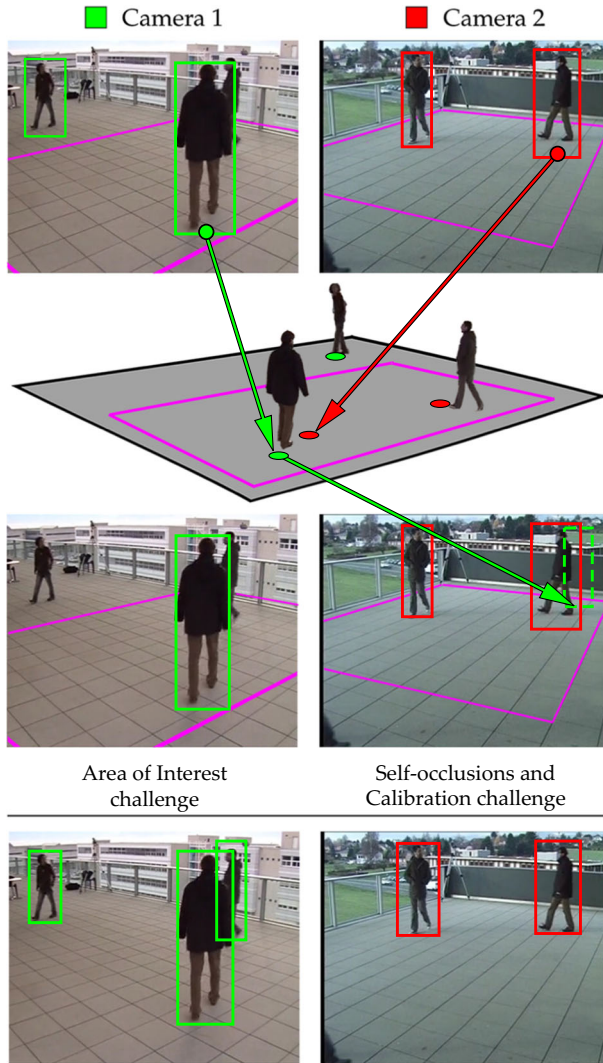


Fig. 1 Common challenges of people detection in multi-camera scenarios. First row: Per-camera people detection by Faster-RCNN [32] (solid bounding boxes) with superimposed—in pink, manually annotated area of interest (*AOI*) from [29]. Second row: Reference ground plane with projected *AOI* and detections (color dots). Area of interest challenge: one projected detection lays outside of the *AOI* and is filtered-out. Self-occlusions and Calibration challenges: projected detections from different camera views diverge in the common plane due to self-occlusions and calibration errors. The back-projection of a detection from Camera 1 onto Camera 2 creates a miss-aligned bounding box (dotted line). Fourth row: qualitative results of the proposed multi-camera pedestrian detection method using the whole imaged floor as *AOI* and aligned back-projected detections. Better viewed in color (Color figure online)

The experimental results on public datasets (PETS 2009 [12], EPFL RLC [3, 6], EPFL Terrace [13, 14] and EPFL Wildtrack [4, 5]) prove that the proposed method: (1) outperforms state-of-the-art monocular pedestrian detectors [31, 32], (2) outperforms state-of-the-art scene-agnostic multi-camera detection approaches and (3) results in a performance com-

parable, and even better, to deep-learning multi-camera detection approaches trained and fine-tuned to the target scenario, while not requiring neither a manually annotated operational area nor a specific training on that scenario.

The rest of the paper is organized as follows: Sect. 2 reviews the state of the art, Sect. 3 describes the proposed method, Sect. 4 presents and discusses experimental results and Sect. 5 concludes the paper.

2 Related Work

Multi-camera people detection faces the combination, fusion and refinement of visual cues from several individual cameras to obtain more people locations. A common pathway in existing approaches starts by defining an operational area, either manually or, as we propose, based on a semantic segmentation. Then, approaches combining detections using a common reference plane, usually follow a three-stage strategy: (1) extract detections on each camera frame, (2) project detections onto the common plane and (3) combine detections and back-project them to the individual views to obtain per-camera people detections. Finally, obtained detections are sometimes post-processed to further refine their localization.

2.1 Definition of the operational area

Some approaches [29, 37] rely on manually annotated operational areas where evaluation is performed. An advantage of these *ad hoc* areas is that the impact of camera calibration errors is limited and controlled. Besides, these areas are defined to maximize the overlapping between the field of view of the involved cameras. However, the manual annotation of these operational areas hinders the generalization of people detection approaches. Our previous work in this domain [26] resulted in an automatic method for the cooperative extraction of operational areas in scenarios recorded with multiple moving cameras: Semantic evidences from different junctures, cameras and points-of-view are spatiotemporally aligned on a common ground plane and are used to automatically define an operational area or *Area of Interest (AOI)*.

2.2 Semantic segmentation

Semantic segmentation is the task of assigning a unique object label to every pixel of an image. During the last years, top performing strategies evolved from the seminal fully convolutional network scheme [24] and the use of dilated convolutions [40]. For instance, Zhao et al. [42] proposed to implicitly use contextual information by including relationships between different labels—e.g., an airplane is likely to be on a runway or flying in the sky but not on the water. These relationships reduce the inner complexity of datasets with large sets of labels, generally improving performance. Lately, the development and use of new backbones for feature extraction has benefited the task. Zhang et al. [41] proposed a new ResNet modification called ResNeSt that uses channel-wise attention to capture cross-feature interactions and learn diverse object representations. Similarly, Tao et al. [36] proposed the use of a hierarchical attention to combine multi-scale predictions, increasing the performance on the small object instances, as those in PASCAL VOC dataset [11].

2.3 Monocular people detection

As stated in Sect. 1, automatic monocular pedestrian detection is considered a mature technology able to obtain accurate results in a broad range of scenarios. Well established object detectors based on CNNs as Faster-RCNN [32] and YOLOv3 [31] have demonstrated their reliability during the last years. Adapting their core schemes, recent approaches have further increased their performance. Specifically, YOLOv3 has been improved by both decreasing the complexity of the model through new architecture designs [25] and by efficient model scaling [38].

Alternatively, novel detectors—also based on CNNs, have been proposed. Tan et al. [35] proposed a weighted bidirectional feature pyramid network allowing easy and fast multi-scale feature fusion and obtaining a new family of detectors called EfficientDet that achieved a new state-of-the-art performance in the COCO dataset [23]. Alternatively, Zhu et al. [44] proposed to use attention in the form of deformable transformers to also obtain the state-of-the-art results.

Nevertheless, even though the most recent works have demonstrated really high performances, in scenarios with severe occlusions the performance of these algorithms decreases.

2.4 Projection of per-camera detections

Multi-camera pedestrian detection is fundamentally based on the projection of monocular detections onto a common reference plane. Projection is typically achieved either by using calibrated camera models that relate any 2D image point with a corresponding referenced 3D world direction [37] or by relying on homographic transformations that project image pixels to a specific 3D plane [29]. In both cases, the ground plane, where people is usually standing on, is chosen as reference for simplicity reasons.

2.5 Fusion and refinement of per-camera detections

Fusion and refinement approaches can be mainly divided into three different groups depending on how global detections are obtained. The first group encompasses **geometrical** methods, which combine detections based on the geometrical intersections between image cues. The second group embraces **probabilistic** methods that combine detections via optimization frameworks and statistical modeling of the image cues. The third group is composed of solutions based on the ability of **deep learning** architectures to model occlusions and achieve accurate pedestrian detection at scene level.

Regarding **geometrical** methods, detections are combined by projecting foreground masks to the ground plane in a multi-view scenario: The intersection of foreground regions leads to pedestrian detection [1]. Accuracy can be increased by projecting the middle vertical axis of pedestrians, leading to a more accurate intersection on the ground plane and, therefore, to a better estimation of the pedestrian's position [21]. Following the same hypothesis, the use of a space occupancy grid to combine silhouette cues has been proposed: Each ground pixel is considered as an occupancy sensor and observations are then used to infer pedestrian detection [15]. All of these approaches outperform single-camera pedestrian detection algorithms by the use of ground plane homography projections. Nevertheless, the evaluation of foreground intersections in crowded spaces may lead to the appearance of *phantoms* or false detections. To handle this problem, the general multi-camera homography framework has been extended by using additional parallel planes to the ground plane [8, 20]. The intersection of the image

cues with these parallel planes is expected to suppress these *phantoms*. Similarly, parallel planes can be also used to create a full 3D reconstruction of pedestrians that can then be back-projected to each of the camera views, improving monocular pedestrian detection [2]. Finally, Lima et al. [22] replicates a preliminar version of the method proposed in this paper, which is available as a preprint [27], with the addition of people re-identification features to guide the fusion of per-camera detections.

Among **probabilistic** methods, an interesting example is the use of a multi-view model shaped by a Bayesian network to model the relationships between occlusions [29]. Detections are here assumed to be images of either pedestrians or *phantoms*, the former differentiated from the latter by inference on the network.

Recent approaches are focused on **deep learning** methods. The combination of CNNs and conditional random fields (CRF) can be used to explicitly model ambiguities in crowded scenes [3]. High-order CRF terms are used to model potential occlusions, providing robust pedestrian detection. Alternatively, multi-view detection can be handled by an end-to-end deep learning method based on an occlusion-aware model for monocular pedestrian detection and a multi-view fusion architecture [7].

2.6 Improving detection's localization

Algorithms in all of these groups require accurate scene calibration: Small calibration errors can produce inaccurate projections and back-projections which may contravene key assumptions of the methods. These errors may lead to misaligned detections, hindering their later use. To cope with this problematic, one can rely on an height-adaptive projection (HAP) procedure in which a gradient descent process is used to find both the optimal pedestrian's height and location on the ground plane by maximizing the alignment of their back-projections with foreground masks on each camera [29].

3 Proposed pedestrian detection method

The proposed method is depicted in Fig. 2. First, state-of-the-art algorithms for monocular pedestrian detection and semantic segmentation are used to extract people detections and the semantic cues for each camera, respectively. These cues drive the automatic definition of the *AOT*, and detections outside this area are discarded. Surviving per-camera detections are combined to obtain global 3D detections by establishing rules and constraints on a disconnected graph. These detections are back-projected to their original camera views in order to further refine their location and height estimates.

3.1 Preliminaries

Monocular Pedestrian Detection: is performed using a state-of-the-art detector. In order to avoid a potential height-bias, we ignore the height and width of the detected bounding boxes, i.e., the j^{th} pedestrian detection at camera k is just represented by the middle point of the base of its bounding box: $\mathbf{p}_{j,k} = (x, y, 1)^T$, in homogeneous coordinates.¹

¹ we use common notation, upper case to denote 3D points/coordinates and lower case to denote 2D camera plane points/coordinates.

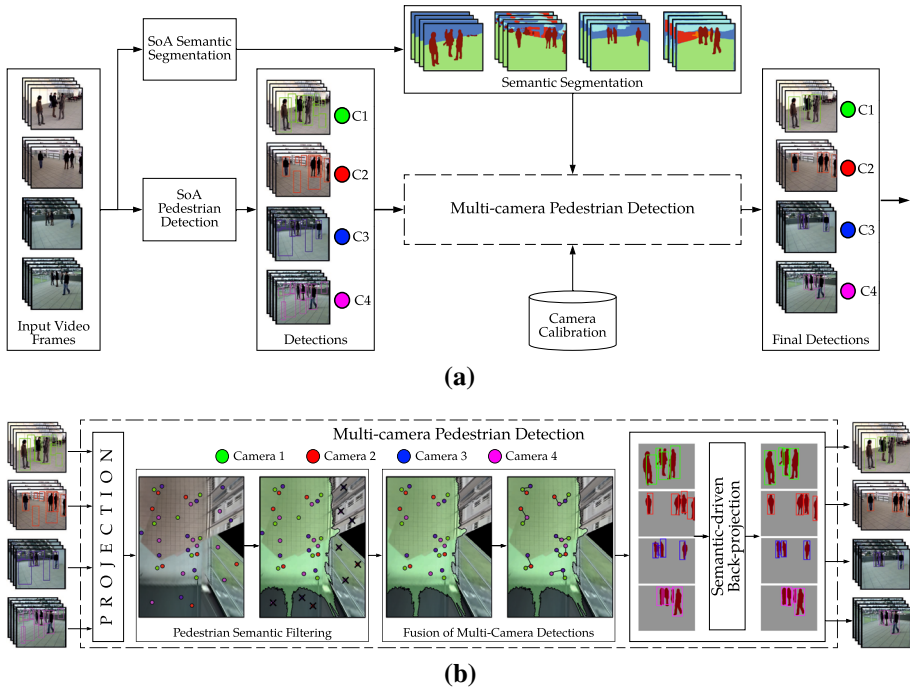


Fig. 2 Overall pedestrian detection method. Top **a**: processing starts performing both a semantic segmentation and a pedestrian detection over a set of cameras (four, in the illustration) with overlapping fields of view. The segmentation, the detections and the camera calibration parameters feed the multi-camera pedestrian detection module which is described in detail in bottom **b**: detections are projected onto a 3D reference plane; a pedestrian semantic filtering module is used to remove detections located out of the automatically generated \mathcal{AOI} ; the remaining detections are combined, based on a disconnected graph, to obtain global detections. The so-obtained global detections are back-projected to the camera views, and the semantic-driven back-projection module globally refines the location of these detections by also using semantic cues. Better viewed in color

Semantic Segmentation: is performed using a state-of-the-art semantic segmentation algorithm. The method is used to label each image pixel \mathbf{p}_k for every camera k and every frame n : $l_n(\mathbf{p}_k) = s_i$, where s_i is one of the L pre-trained semantic classes: $S = \{s_i\}$, where $i \in [1, L]$, i.e., *floor, building, wall...* Fig. 3 depicts examples of semantic labels for selected camera frames of the Terrace Dataset [13, 14].

Projection of People Detections: Let \mathcal{H}_k be the homography matrix that transforms points from the image plane of camera k to the world ground plane. The j^{th} detection of camera k , $\mathbf{p}_{j,k} = (x, y, 1)'$ is projected onto the ground plane by:

$$\mathbf{P}_{j,k} = \mathcal{H}_k \times \mathbf{p}_{j,k} = (X, Y, T)', \tag{1}$$

which corresponds to a $(X = X/T, Y = Y/T, Z = 0)'$ 3D point of the ground plane.

3.2 Pedestrian semantic filtering

Automatic definition of the \mathcal{AOI}

To obtain a semantic partition of the ground plane, an adaptation of [26] for static-camera scenarios is carried out. We first project *every* image pixel \mathbf{p}_k via \mathcal{H}_k . Every projected point




Fig. 3 Top row represents RGB frames from the Terrace Dataset [13, 14]. Bottom row represents the correspondent semantic labels obtained by the PSP-Net algorithm [42]. Columns from left to right represent cameras 1 to 4 of this dataset. The bottom legend indicates the detected semantic classes. Better viewed in color

\mathbf{P}_k inherits the semantic label assigned to \mathbf{p}_k :

$$l_n(\mathbf{P}_k) = l_n(\mathbf{p}_k) = s_i \in S. \quad (2)$$

Thereby, a semantic locus—a ground plane semantic partition, is obtained for each camera. The extent of each locus is defined by the image support, and missing points inside the locus are completed by nearest neighbor interpolation.

In order to globally reduce the impact of moving objects and segmentation errors, we propose to temporally aggregate each locus along several frames. In a set of T loci obtained for T consecutive frames, a given point on the ground plane \mathbf{P}_k is labeled with T semantic labels, which may be different owing to inaccuracies in the semantic segmentation or to the presence of moving objects. A single temporally smoothed label $\bar{l}_n(\mathbf{P}_k)$ is obtained as the mode value of this set. Examples of these per-camera obtained smoothed loci are included in the first four columns of Fig. 4.

We propose to combine these loci to define the \mathcal{AOI} . The definition of the \mathcal{AOI} is scenario-dependent but can be generalized by defining a set \mathcal{G} of ground-related semantic classes: *floor*, *grass*, *pavement*, etc. The operational area \mathcal{AOI} is obtained as the union of the projected pixels from any camera which are labeled with any class in \mathcal{G} :

$$\mathcal{AOI} = \bigcup_k^K \mathbf{P}_k, \quad s.t. \quad \bar{l}_n(\mathbf{P}_k) \in \mathcal{G}. \quad (3)$$

An example of a so-obtained \mathcal{AOI} is included in the right most column of Fig. 4.

Detection filtering

Projected detections $\mathbf{P}_{j,k}$ lying outside the operational area, $\mathbf{P}_{j,k} \notin \mathcal{AOI}$, are filtered out and so, discarded for forthcoming stages.




Fig. 4 Temporally smoothed projected loci for each camera (columns 1 to 4) of the Terrace Dataset [13, 14], both in the RGB domain (top) and the semantic labels domain (bottom). The last column depicts, again in both domains, the resulting \mathcal{AOT} which, in the example, consists of the combined *floor* class of the four smoothed loci. Better viewed in color

3.3 Fusion of multi-camera detections

We propose a **geometrical** approach to combine detections on the ground plane. Every camera single detection is considered a vertex of a disconnected graph located in the reference plane. Vertices are then joined generating connected components C_m , each representing a joint 3D global detection. The whole fusion process is summarized in Fig. 5. The conditions that shall be satisfied to join two vertices or detections, $\mathbf{P}_{j,k}$ and $\mathbf{P}_{j',k'}$, are:

1. That vertices in a connected component are *close enough*. The l_2 -norm between any two vertices in C_m shall be smaller than a predefined distance R_1 : $\|\mathbf{P}_{j,k}, \mathbf{P}_{j',k'}\|_2 \leq R_1$ (Fig. 5a). R_1 may be fixed in the interval between 2.5 and 3.5 with no influence in the results. We experimentally set $R_1 = 3$ meters to: 1) reduce the computational cost of the final stage (see below) assuming that vertices separated R_1 do not belong to the same object and 2) protect against calibration errors, assuming that they are not larger than R_1 .
2. That vertices in a connected component come from different cameras. This condition prevents the joining of two different detections from the same camera which are near in the ground plane. (Fig. 5b)

To avoid ambiguities, the creation of connected components is performed in order, according to the spatial position of the detections: Those with a lower module are combined first.

The outcome of the fusion process for K cameras is a set of M connected components $\{C_m, m = 1, \dots, M\}$, each containing K_m detections: $|C_m| = K_m \leq K$, where $K_m < K$ when a person is occluded or not detected in one or more cameras.

As each connected component is assumed to represent a single person, an initial ground-position of the person $\mathbf{P}_m^G = (X_m, Y_m, Z_m = 0)^T$ is obtained by simply computing the arithmetic mean of all the detections in the C_m connected components (Fig. 5c).

3.4 Semantic-driven back-projection

To obtain correctly positioned detections, i.e., visually precise detections, in each camera, ground plane detections need to be back-projected to each camera and 2D bounding boxes enclosing pedestrians need to be outlined based on these projections.

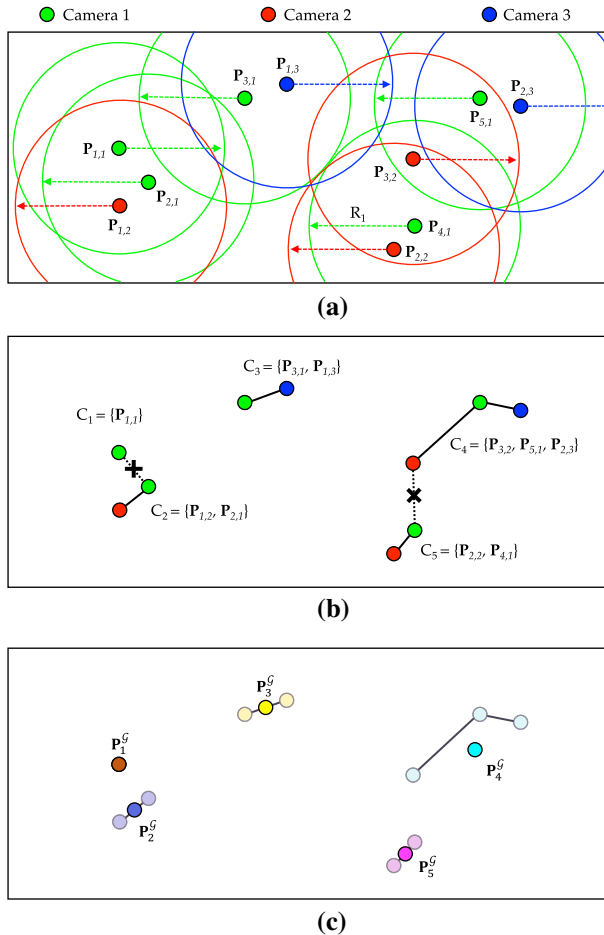


Fig. 5 Fusion of multi-camera detections in the ground plane. **a** The distance R_1 , depicted here as circumferences around detections, defines neighbors for each detection $\mathbf{P}_{j,k}$. **b** Connected components C_m are defined for detections: (i) which l_2 -norm is lower than R_1 and (ii) that are projected from different cameras. Connected components fulfilling (i) but not (ii) are represented by dashed lines crossed out. **c** The ground plane detection \mathbf{P}_m^G is obtained as the arithmetic mean of all the detections in a connected component C_m . Better viewed in color

The problem of back-projecting 3D detections

Let $\bar{\mathbf{P}}_m$ be an orthogonal line segment to the ground plane which represents the detected pedestrian and extends from the detection \mathbf{P}_m^G to a 3D point h_m meters above. Using the camera calibration parameters, the segment $\bar{\mathbf{P}}_m$ can be back-projected onto camera k . This back-projection defines a 2D line segment $\bar{\mathbf{p}}_{m,k}$, which extends between $\mathbf{p}_{m,k}$ and $\mathbf{p}_{m,k} +$ (see Fig. 6a).

We propose to create 2D bounding boxes around these back-projected 2D line segments. To this aim, each segment is used as the vertical middle axis of its associated 2D bounding-box $\mathbf{b}_{m,k}$. For simplicity, the width of $\mathbf{b}_{m,k}$ is made proportional to its height. Due to pedestrian self-occlusion, calibration errors and the uncertainty on the pedestrians' height, this back-

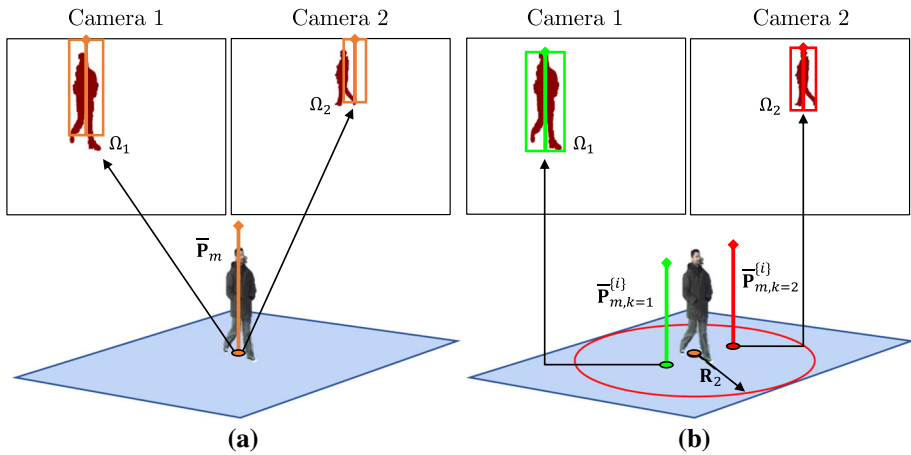


Fig. 6 **a** Back-projecting global segment $\bar{\mathbf{P}}_m$ results in misaligned bounding boxes due to pedestrian self-occlusion, calibration errors and the uncertainty on the pedestrians’ height. **b** The proposed optimization process results in the best-aligned segments $\mathbf{P}_{m,k}^{(i)}$ for each camera. Better viewed in color

projection process results in misaligned bounding-boxes (see Fig. 6a), hindering their later use and degrading camera-wise performance.

To handle this problematic, we define an iterative method which aims to globally optimize the alignment between all 3D detections and their respective views or back-projections in all cameras. This method is based on the idea proposed in [29]. While the referenced method is guided by a foreground-segmentation, we instead propose to use a cost-function driven by the set of pedestrian-labeled pixels Ω_k from the semantic segmentation (e.g., see person label in Fig. 3). Next, we detail the full process for the sake of reproducibility.

Method overview

As a 3D detection $\bar{\mathbf{P}}_m$, with height h_m , inevitably results in misaligned back-projected 2D detections, the proposed method tries to adapt the 3D detection segment to each camera, generating a set of 3D detection segments, $\bar{\mathbf{P}}_{m,k}$, for each 3D detection and iteratively modifying their positions and height to maximize 2D detections’ alignment with the semantic segmentation masks, while constraining all the segments to have the same final height h'_m (as they are all projections of a same pedestrian) and to be located sufficiently close to each other. This process is not performed independently for each 3D detection but jointly and iteratively for all 3D detections. Observe that the joint nature of the optimization problem for all 3D detections is a key step as pedestrian pixels Ω_k may contain segmentations from more than one pedestrian.

For each 3D segment $\bar{\mathbf{P}}_m$, the method starts by initializing (i.e., iteration $i = 0$) the per-camera adapted segments:

$$\bar{\mathbf{P}}_{m,k}^{(i=0)} = \bar{\mathbf{P}}_m, k = 1 \dots K. \tag{4}$$

Iterative steepest ascent algorithm

For each 3D segment, let $\mathcal{P}_k^{(i)} = \{\bar{\mathbf{P}}_{m,k}^{(i)}, m = 1 \dots M\}$ be the set of adapted detections to camera k at iteration i , and let $\mathbb{P}^{(i)} = \{\mathcal{P}_k^{(i)}, k = 1 \dots K\}$ be the set of camera-adapted segments for all cameras at the same iteration.

The optimization process aims to find \mathbb{P}^* , the solution to the constrained optimization problem:

$$\mathbb{P}^* = \arg \max_{\mathbb{P}} \Psi(\mathbb{P}), \quad s.t. \quad \|\bar{\mathbf{P}}_m, \bar{\mathbf{P}}_{m,k}\|_2 \leq R_2 \forall(m, k), \tag{5}$$

where R_2 defines the maximum distance between 3D projections of a single pedestrian, which we set to twice the average width of the human body, i.e., 1 meter, to forestall the effect of nearby pedestrians in the image plane. Performed experiments suggest that variations in R_2 value have no significant influence on the results.

$\Psi(\mathbb{P})$ is defined as the cost function to maximize and is based on the alignment of the back-projected bounding boxes with the set of pedestrian-labeled pixels in each camera: Ω_k . The cost function considers the information from all the cameras.

$$\Psi(\mathbb{P}^{(i)}) = - \sum_{k=1}^K \frac{\sum_{\mathbf{p}} \gamma(\mathbf{p}, \Omega_k) \Phi(\mathbf{p}, \mathbb{P}^{(i)})}{|F_k|}, \tag{6}$$

where $\gamma(\mathbf{p}, \Omega_k)$ is a weight for pixel \mathbf{p} : ω for pedestrian and $\omega/3$ for non pedestrian pixels, $\omega = 1$ in our setup, $|F_k|$ is the number of pixels in the camera image plane and $\Phi(\mathbf{p}, \mathbb{P})$ is the loss function of pixel \mathbf{p} with respect to \mathbb{P} :

$$\Phi(\mathbf{p}, \mathbb{P}^{(i)}) = \begin{cases} \prod_{m|\mathbf{p} \in \mathbf{b}_{m,k}^{(i)}} (1 - 1/d_{m,k}) & , \text{ if } l_k(\mathbf{p}) \in \Omega_k \\ 1 - \prod_{m|\mathbf{p} \in \mathbf{b}_{m,k}^{(i)}} (1 - 1/d_{m,k}) & , \text{ if } l_k(\mathbf{p}) \notin \Omega_k, \end{cases} \tag{7}$$

where $d_{m,k}$ is the distance from \mathbf{p} to the vertical middle axis $\bar{\mathbf{p}}_{m,k}^{(i)}$ of the back-projected bounding box $\mathbf{b}_{m,k}^{(i)}$.

At each iteration i , the set of camera-adapted segments is moved toward the direction of maximum increment:

$$\mathbb{P}^{(i)} = \mathbb{P}^{(i-1)} + \tau_i \vec{\nabla} \Psi(\mathbb{P}^{(i-1)}), \tag{8}$$

where $\tau_i \in \mathbb{R}_+$ is the gradient step that makes $\mathbb{P}^{(i)} \geq \mathbb{P}^{(i-1)}$. This gradient step is initialized, $\tau_0 = 5$, and updated following a decrease schedule of 50% every 3 iterations, to ease convergence. The gradient $\vec{\nabla} \Psi(\mathbb{P}^{(i)})$ in the i -th iteration is approximated by forward difference approximation:

$$\vec{\nabla} \Psi(\mathbb{P}^{(i)}) = \frac{\Psi(\mathbb{P}^{(i)}) - \Psi(\mathbb{P}^{(i)} - \epsilon)}{\epsilon}. \tag{9}$$

The algorithm continues until convergence is reached or the R_2 -constrain is violated.

4 Results

This section addresses the evaluation of the proposed method. To this aim, we first describe the evaluation framework; then, in the ablation studies, we measure the performance improvement of each of the method's stages; finally, we finish by comparing our approach with alternative state-of-the-art approaches in classic and recent multi-camera datasets.

4.1 Evaluation framework

Datasets

The results are obtained by evaluating the proposed method over five scenarios extracted from four publicly available multi-camera datasets in which cameras are calibrated and temporally synchronized:

- EPFL Terrace [13, 14]: Generally used in the state of the art to evaluate multi-camera approaches. It consists of a 5000 frames sequence per camera showing up to eight people walking on a terrace captured by four different cameras. All the cameras record a close-up view of the scene.
- EPFL RLC [3, 6]: consists of an indoor sequence of 2000 frames per camera recorded in the EPFL Rolex Learning Center using three static HD cameras with overlapping field of views. All these cameras represent close-up views of the scene.
- EPFL Wildtrack [4, 5]: A challenging multi-camera dataset which has been explicitly designed to evaluate deep learning approaches. It has been recorded with 7 HD cameras with overlapping fields of view. Pedestrian annotations for 400 frames are provided. All of them are used to define the evaluation set used in this paper.
- PETS 2009 [12]: The most used video sequences from this widely used benchmark dataset have been chosen.
 - PETS 2009 S2 L1, which contains 795 frames recorded by eight different cameras of a medium density crowd—in this evaluation, we have just selected 4 of these cameras: view 1 (far field view) and views 5, 6 and 8 (close-up views).
 - PETS 2009 City Center (CC), recorded only using two far-field view cameras with around 1 minute of annotated recording (400 frames per camera).

Table 1 contains a comparative description of these datasets including the type of data and annotations provided, as well as a subjective indication of their complexity for the pedestrian detection task.

Performance indicators

To obtain quantitative performance statistics according to an experiment-based evaluation criterion the following state-of-the-art performance indicators have been selected: Precision (P), Recall (R), F-Score (F-S), Area Under the Curve (AUC), N-MODA (N-A) and N-MODP (N-P) [9, 33]. To globally assess performance, a single value for each statistic and each configuration is provided by averaging per-camera ones.

Table 1 Dataset description

	Resolution	Cameras	Frames	Pedestrian Density	Occlusions Level	Camera PoV	GT Annotations (Cameras)
PETS 2012 CC [12]	768×576	2	795	+	+	High	Frame (1)
PETS 2012 S2 L1 [12]	768×576	4	795	+	+	High	Frame (1)
EPFL Terrace [13, 14]	366×288	4	5008	++	++	Low	Frame (4)
EPFL RLC [3, 6]	480×270	3	1197	++	++	Low	Frame (1)
EPFL Wildtrack [4, 5]	1920×1080	7	401	+++	+++	Medium	Ground Plane (NA)

Pedestrian Density and Occlusions Level are ranked subjectively in terms of the number of pedestrians with respect to the scenario space and the number and detection difficulty of the occlusions (+ stands for easy, ++ stands for medium and +++ stands for difficult). The camera PoV (Point of View), which affects the level of pedestrians' occlusions, is categorized as low, medium or high according to the height of the camera with respect to the ground plane. Figure 8 depicts frame examples for these analyzed datasets

4.2 System setup

A common setup has been used for all the presented results. Faster-RCNN [32], YOLOv3 [31] and EfficientDet-D7 [35] are used as baseline algorithms to obtain monocular pedestrian detections. The three object detectors are pre-trained on the COCO dataset [23] and we do not fine-tune nor adapt them to any of the faced scenarios. For the semantic segmentation, the Pyramid Scene Parsing Network (PSP-Net) [42], pre-trained on the ADE20K dataset [43] ($L = 150$), has been selected considering a trade-off between performance and efficiency.

In the pedestrian semantic filtering stage, all frames in each sequence are used for temporal and spatial semantic aggregation, i.e. $T = N$. For the semantic-driven back-projection stage, the initial height estimation h_m has been set to an average pedestrian height of 1.7m. Besides, for all the datasets, convergence in the iterative steepest ascent algorithm has been reached before or at the 8th iteration.

4.3 Results overview

The evaluation has been performed carrying out two different studies:

- The ablation studies aim to gauge the impact of the different stages in the performance of the proposed approach. To this end, the following versions of the proposed method are compared:
 1. “Baseline (Faster-RCNN, YOLOv3 and EfficientDet-D7)”, provides reference results of monocular pedestrian detectors.
 2. “Baseline + Filtering (*Filt*)” is a simplified version of our method which aims to independently evaluate the effect of the proposed automatic *AOI* computation obtained by the “Pedestrian Semantic Filtering” stage.
 3. “Baseline + Filtering (*Filt*) + Fusion (*Fus*) + Back-Projection (*BP*)” is the full version of the proposed method, which additionally evaluates the “Fusion of Multi-Camera Detections” and “Semantic-Driven Back-Projection” stages.

Ablation Studies are conducted on four of the described datasets: Terrace, PETS 2009 S2 L1, PETS 2009 CC and RLC.

- State-of-the-art comparison results analyze the proposed method with respect to several non-deep-learning state-of-the-art multi-camera pedestrian detectors on the same four scenarios used in the ablation studies. Additionally, the method is compared with novel deep-learning methods on the Wildtrack dataset.

4.4 Ablation studies

Evaluation criterion

The availability of bounding-box annotations permits to use the classic performance criterion [16]: A detection is considered a TP one if the Intersection Over Union (IoU) with a ground-truth bounding box is higher than 0.5.

Results

Table 2 agglutinates the method’s performance on a per-stage basis. Qualitative examples of automatically generated *AOIs* and algorithm results are depicted in Figs. 7 and 8, respec-

Table 2 Ablation Studies: Stage-wise performance of the proposed method when Faster-RCNN [32], YOLOv3 [31] and EfficientDet [35] are used as baselines

	Dataset																							
	PETS 2009 CC						PETS 2009 S2 L1						EPFL Terrace						EPFL RLC					
	Filt	Fus & BP	AUC	F-S	NA	NP	AUC	F-S	NA	NP	AUC	F-S	NA	NP	AUC	F-S	NA	NP	AUC	F-S	NA	NP		
Faster-RCNN			0.90	0.91	0.85	0.76	0.90	0.91	0.85	0.76	0.82	0.84	0.71	0.74	0.77	0.78	0.58	0.69	0.77	0.78	0.58	0.69		
	✓		0.90	0.91	0.85	0.76	0.90	0.91	0.85	0.76	0.84	0.85	0.73	0.74	0.80	0.82	0.68	0.70	0.80	0.82	0.68	0.70		
	✓	✓	0.94	0.94	0.88	0.79	0.92	0.93	0.89	0.79	0.87	0.90	0.83	0.77	0.81	0.81	0.68	0.70	0.81	0.81	0.68	0.70		
YOLOv3			0.92	0.92	0.87	0.79	0.96	0.96	0.92	0.67	0.83	0.87	0.76	0.73	0.80	0.78	0.59	0.72	0.80	0.78	0.59	0.72		
	✓		0.92	0.92	0.87	0.79	0.96	0.96	0.92	0.67	0.84	0.87	0.77	0.73	0.85	0.83	0.66	0.72	0.85	0.83	0.66	0.72		
	✓	✓	0.94	0.94	0.88	0.79	0.93	0.92	0.89	0.67	0.86	0.89	0.85	0.76	0.85	0.83	0.68	0.72	0.86	0.83	0.68	0.72		
Efficient Det			0.97	0.97	0.94	0.66	0.97	0.97	0.94	0.67	0.82	0.84	0.71	0.78	0.82	0.80	0.61	0.72	0.82	0.80	0.61	0.72		
	✓		0.97	0.97	0.94	0.66	0.97	0.97	0.94	0.66	0.82	0.87	0.76	0.78	0.85	0.83	0.68	0.72	0.85	0.83	0.68	0.72		
	✓	✓	0.95	0.94	0.88	0.75	0.94	0.94	0.88	0.67	0.86	0.89	0.83	0.76	0.83	0.84	0.71	0.72	0.83	0.84	0.71	0.72		

Bold values indicate best result in terms of N-MODA. Indicators are Area Under the Curve (AUC), F-Score (F-S), N-MODA (N-A) and N-MODP (N-P). *Filt* stands for the “Pedestrian Semantic Filtering” stage and *Fus* & *BP* stands for the “Fusion of Multi-Camera Detections (Fus) and Semantic-Driven Back-Projection (BP)” stages. Datasets are sorted from left to right in terms of complexity according to Table 1




Fig. 7 Ablation Studies: Automatically obtained \mathcal{AOI} (superimposed in green) compared to the \mathcal{AOI} manually annotated (red box) by the authors of EPFL Terrace [13, 14] (left), EPFL RLC Dataset [3, 6] (middle) and PETS2009 [12] (right). Better viewed in color (Color figure online)

tively. A visual example of the limitations of the semantic-driven back-projection stage is included in Fig. 9.

Discussion

Table 2 shows that filtering-out detections using automatically generated \mathcal{AOI} s (Baseline + Filtering) improves the performance of all the baselines for datasets where the ground plane area does not cover the whole image representation, i.e., datasets containing close-up views of the scene as EPFL Terrace and RLC. In these datasets, our precise \mathcal{AOI} s reduce *phantom* detections obtained by the baseline detectors. Although \mathcal{AOI} s are automatically computed, they are more precise (tighter to real scene edges) than those defined in the dataset.

Overall, in the EPFL Terrace Dataset, the performance of Faster-RCNN + Filtering improves Faster-RCNN by 2.44% and 2.82% in terms of AUC and N-MODA, respectively. YOLOv3 + Filtering presents relative increments over YOLOv3 baseline of 1.20% and 1.31% for AUC and N-MODA, respectively. Finally, EfficientDet + Filtering also overcomes its baseline results by a 7.04% for N-MODA.

For the EPFL RLC dataset with the proposed \mathcal{AOI} , Faster-RCNN is improved by a 5.13% regarding AUC and by a 17.24% concerning N-MODA. For YOLOv3, relative increments of a 6.25% and a 11.86% in terms of AUC and N-MODA are achieved. EfficientDet gains relative increments of a 3.65% and a 11.47% for AUC and N-MODA.

The proposed filtering stage does not improve baselines' performance for those datasets in which the ground plane dominates the scene, i.e., those recorded with far-field view cameras as both scenarios from PETS 2009. In these cases, although the baseline pedestrian detectors may create *phantom* detections, those lie inside the proposed \mathcal{AOI} and no false-pedestrians are suppressed. However, as depicted in Fig. 7, the automatically obtained \mathcal{AOI} s are larger and more precise than the original operational areas in the datasets, thereby obtaining a more realistic and exhaustive evaluation. Furthermore, observe how the proposed generation method effectively handles multi-class ground partitions as in the PETS 2009 dataset, where the proposed \mathcal{AOI} encompasses *road*, *grass*, *pavement* and *side-walks* classes enabling a high adaptability to unseen scenarios (see Fig. 7 right).

Table 2 also shows that the complete method (Baseline + Filtering + Fusion + Back-Projection) notably improves Faster-RCNN baseline's performance, mainly in scenarios with heavy occlusions, i.e., EPFL Terrace and EPFL RLC (See Table 1 for details). Specifically, for the EPFL Terrace Dataset results are relatively increased a 6.14%, a 7.14% and a 16.90% in terms of AUC, F-Score, and N-MODA, respectively, whereas relative improvements are of




Fig. 8 Ablation Studies: Proposed method qualitative results on selected frames of the EPFL Terrace, PETS S2 L1, PETS CC and EPFL RLC datasets (Faster-RCNN baseline is here used). From left to right: First three columns depict a same time frame captured by three available cameras, showing color bounding boxes (a color per pedestrian) corresponding to the final per-camera detections. The most-right column depicts obtained detections—one per pedestrian in the scene—on the ground plane, conserving the identifying colors. Better viewed in color

a 5.19%—in AUC, a 5.13%—in F-Score terms—and a 20.69% in N-MODA, for the EPFL RLC dataset.

For YOLOv3 and EfficientDet detectors a similar analysis arises. In scenarios where heavy occlusions are present—EPFL Terrace and RLC datasets, performance is increased. For the EPFL Terrace Dataset, relative increments of a 2.38%, a 2.29% and a 11.84% are obtained when using YOLOv3 detections in terms of AUC, F-Score and N-MODA, respectively. In the case of EfficientDet, increments of a 4.87%, a 5.95% and a 16.90% are obtained with respect to the same metrics. For the EPFL RLC dataset, the improvement increases to a 6.25%, a 6.41% and a 15.25% for YOLOv3, whereas a 1.21%, a 3.75% and a 13.11% relative increase is obtained for EfficientDet in terms of AUC, F-Score and N-MODA, respectively.

For both PETS scenarios the performance of the EfficientDet mono-camera detector is saturated (97% F-Score). The specific characteristics of this dataset: low pedestrian density over a wide space, low level of occlusions and a high point of view due to cameras being hanged up in streetlights (see Table 1 and Fig. 8), turns it in the least complex dataset among those analyzed. The generation of new false positive detections and the optimization process

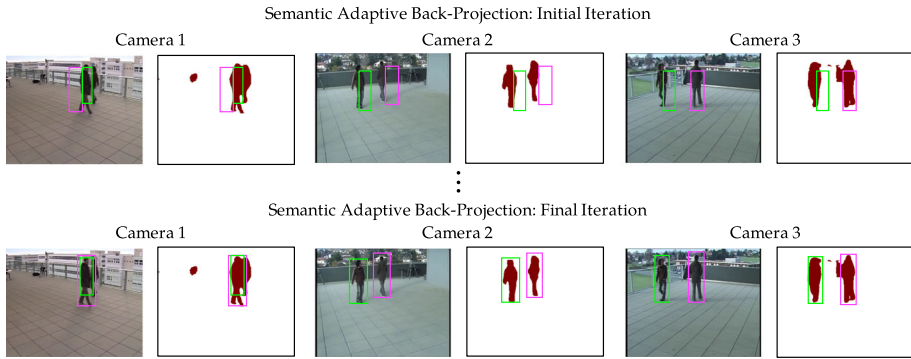


Fig. 9 Ablation Studies: Semantic-Driven Back-Projection. First row: back-projected bounding boxes at the initial iteration of the optimization algorithm. Global detections obtained by the multi-camera detection fusion algorithm are displaced with respect to real pedestrian when back-projected to each camera. Second row: The semantic-driven optimization algorithm correctly refines locations and heights for the bounding boxes in Camera 2 and 3. However, when semantic pedestrian cues are highly overlapped some bounding boxes might be refined to an incorrect location (Camera 1, green bounding-box). Better viewed in color (Color figure online)

related problems ((Fig. 9)) may lead to a slight decrease when saturated baselines are used in low complex datasets. Leaving these specific situations aside, the benefits of the proposed method are evident if one accounts for both performance indicators and qualitative results (see Table 2 and Fig. 8, respectively): the proposed multi-camera detection approach is able to cope with partial, severe and complete occlusions by combining detections from all the cameras through the proposed semantic-guided process leading to an increase of all the reported metrics.

Focusing specifically on the semantic-driven back-projection process, the results in Fig. 8 depict highly tight pedestrian bounding boxes, independently from people’s height, self-occlusions and calibration problems, suggesting that the optimization process is able to automatically adapt bounding boxes by jointly estimating pedestrian heights and world positions. The results in Table 2 corroborate this observation. Semantic-driven back-projection leads to a higher overlap between detections and ground-truth annotations: In terms of the N-MODP metric, the proposed method achieves relative improvements of a 4.05% for EPFL Terrace, a 3.95% for both PETS 2009 S2 L1 and PETS 2009 CC and a 1.45% for the RLC dataset when Faster-RCNN is used as the baseline detector. When YOLOv3 is used as the baseline detector, our method achieves a N-MOPD increment of a 4.10% for EPFL Terrace whereas the N-MODP metric remains stable for PETS 2009 S2 L1, PETS 2009 CC and RLC datasets, suggesting that YOLOv3 individual performance for these datasets is already heaped. A similar result arises when using EfficientDet detector which by default is highly tight to pedestrians. The results are increased only for PETS CC dataset by a 13.63% while maintained for the rest of the datasets. It is important to remind that even tough the N-MODP metrics are sometimes slightly reduced or maintained, without the proposed semantic-driven back-projection process the back-projected bounding boxes and ground-truth would be mis-aligned (see Fig. 6) decreasing the performance in terms of all the accuracy metrics of the proposed method.

Nevertheless, the optimization cost function aims to maximize the 2D detections’ alignment with the semantic segmentation masks, leading to a bias toward wider pedestrians by design, a situation that may result sometimes into wrong relocations of the back-projected

bounding-boxes. Figure 9 shows an example of this case: notice the erring behavior in Camera 1 when there is an extreme overlapping.

4.5 State-of-the-art comparison

Evaluation criterion

The same criterion used in the ablation studies applies for the Terrace, PETS and RLC datasets. However, in the Wildtrack dataset, as the ground truth is provided via detections on the world ground plane (i.e., no bounding boxes are provided), the evaluation criterion is different. Specifically, a detection is considered a TP if it lies at most $r = 0.5\text{m}$ to a ground-truth annotated point [4]. This radius roughly corresponds to the average width of the human body. Due to the absence of bounding-boxes, for this dataset the semantic-driven back-projection stage is not included.

State-of-the-art algorithms

The following multi-camera algorithms have been selected to carry out the comparison:

- POM [14]. This algorithm proposes to estimate the marginal probabilities of pedestrians at every location inside an *AOI*. It is based on a preliminary background subtraction stage.
- POM-CNN [14]. An upgraded version of POM in which the background subtraction stage is performed based on an encoder–decoder CNN architecture.
- MvBN+HAP [29]. Relies on a multi-view Bayesian network model (MvBN) to obtain pedestrian locations on the ground plane. Detections are then refined by a height-adaptive projection method (HAP) based on an optimization framework similar to the one proposed in this paper, but driven by background-subtraction cues.
- RCNN-Projected [39]. The bottom of bounding boxes obtained through per-camera CNN detectors are projected onto ground plane, where 3D proximity is used to cluster detections.
- Deep-Occlusion [3] is an hybrid method which combines a CNN trained on the Wildtrack dataset and a conditional random fields (CRF) method to incorporate information on the geometry and calibration of the scene.
- DeepMCD [7] is an end-to-end deep learning approach based on different architectures and training scenarios:
 - Pre-DeepMCD: a GoogleNet [34] architecture trained on the PETS dataset.
 - Top-DeepMCD: a GoogleNet [34] architecture trained on the Wildtrack dataset.
 - ResNet-DeepMCD: a ResNet-18 [18] architecture trained on the Wildtrack dataset.
 - DenseNet-DeepMCD: a DenseNet-121 [19] architecture trained on the Wildtrack dataset.

Results

Table 3 includes performance indicators for the proposed method compared with multi-camera algorithms POM [14] and MvBN+HAP [29] on the Terrace, PETS and RLC scenarios. (The results for the compared methods are extracted from [29].) Table 4 compares the performance of the proposed approach against deep-learning methods, some of them explicitly

Table 3 State-of-the-art Comparison: Comparison with baseline (Faster-RCNN, YOLOv3 and EfficientDet) and multi-camera state-of-the-art methods non based on deep-learning (POM [14] and MvBN + HAP [29])

Algorithm	Dataset											
	PETS CC			PETS S2 L1			EPFL Terrace			EPFL RLC		
	F-S	N-A	N-P	F-S	N-A	N-P	F-S	N-A	N-P	F-S	N-A	N-P
Faster-RCNN [32]	0.91	0.85	0.76	0.91	0.85	0.76	0.84	0.71	0.74	0.78	0.58	0.69
YOLO v3 [31]	0.92	0.87	0.79	0.96	0.92	0.67	0.87	0.76	0.73	0.78	0.59	0.72
Efficient Det [35]	0.97	0.94	0.66	0.97	0.94	0.67	0.84	0.71	0.78	0.80	0.61	0.72
POM [14]	-	0.70	0.55	-	0.65	0.67	-	0.19	0.56	-	-	-
MvBN + HAP [29]	-	0.87	0.78	-	0.87	0.76	-	0.82	0.73	-	-	-
Ours (Faster-RCNN)	0.93	0.88	0.79	0.93	0.89	0.79	0.90	0.83	0.77	0.81	0.68	0.70
Ours (YOLOv3)	0.94	0.88	0.79	0.92	0.89	0.67	0.89	0.85	0.76	0.83	0.68	0.72
Ours (Efficient Det)	0.94	0.88	0.75	0.94	0.88	0.67	0.89	0.83	0.76	0.84	0.71	0.72

Bold values indicate best results in terms of N-MODA. Indicators are F-Score (F-S), N-MODA (N-A) and N-MODP (N-P). Datasets are sorted from left to right in terms of complexity according to Table 1

Table 4 State-of-the-art Comparison: Wildtrack Dataset Comparison Results

Algorithm	EPFL Wildtrack				
	Authors \mathcal{AOI}	Fine-Tuned	F-Score	N-MODA	N-MODP
Deep-Occlusion [3]	✓	✓	0.86	0.74	0.53
ResNet-DeepMCD [5]	✓	✓	0.83	0.67	0.64
Ours* (EfficientDet)	✓		0.81	0.65	0.63
DenseNet-DeepMCD [5]	✓	✓	0.79	0.63	0.66
Top-DeepMCD [7]	✓	✓	0.79	0.60	0.64
GMC-3D [22]	✓		0.78	0.56	0.67
Ours (EfficientDet)			0.74	0.48	0.63
Ours (YOLOv3)			0.71	0.42	0.60
Ours (Faster-RCNN)			0.69	0.39	0.55
Pre-DeepMCD [7]	✓		0.51	0.33	0.52
POM-CNN [14]	✓		0.63	0.23	0.30
RCNN-Projected [39]	✓		0.52	0.11	0.18

“Authors- \mathcal{AOI} ” stands for algorithm performance evaluated using the \mathcal{AOI} proposed by the authors. “Fine-tuned” denotes that the algorithm has been explicitly trained on Wildtrack dataset. Bold values indicate best results in terms of N-MODA

trained with data from the Wildtrack dataset (which we denote as *Fine-Tuned*) and others trained with data from other datasets or not even trained (which we denote as not *Fine-Tuned*). Performance indicators for these methods are extracted from [5]. In addition, the qualitative results for the Wildtrack dataset are presented in Fig. 10, including obtained detections in camera frames, global detections on the ground plane and the automatically computed \mathcal{AOI} .

Discussion

The results in Table 3 show that the proposed approach (Baseline + Filtering + Fusion + Back-Projection), either with Faster-RCNN, YOLOv3 or EfficientDet baseline, outperforms the MvBN + HAP and the POM-CNN methods in terms of N-MODA metric. The proposed method obtains better results in terms of N-MODA which, precisely, measures detection accuracy along the whole video sequences. Best results are obtained when EfficientDet is used to extract mono-camera detections. Specifically, N-MODA is increased a 1.21% for EPFL Terrace and a 1.14% for both PETS 2009 S2 L1 and CC. Moreover, it obtains the higher performance on the heavily occluded RLC dataset. Besides, N-MODP results, i.e., the overlapping between detections and ground-truth, are better than those obtained by the HAP method [29]. This suggests that our use of semantic segmentation masks instead of foreground masks (HAP method) benefits the optimization process. Relative increments in N-MODP performance of a 5.48% for EPFL Terrace, a 3.95% for PETS 2009 S2 L1 and a 1.28% for PETS 2009 CC support this assumption.

The presented results in Table 3 suggest that the proposed method is able to obtain reliable pedestrian detections in a variety of scenarios with a variety of pedestrian detection algorithms in terms of performance.

Finally, the results on the Wildtrack dataset (Table 4) indicate that the proposed method, operating on detections from a Faster-RCNN, a YOLOv3 or a EfficientDet model, is able to outperform deep-learning approaches that have not been specifically trained using Wild-

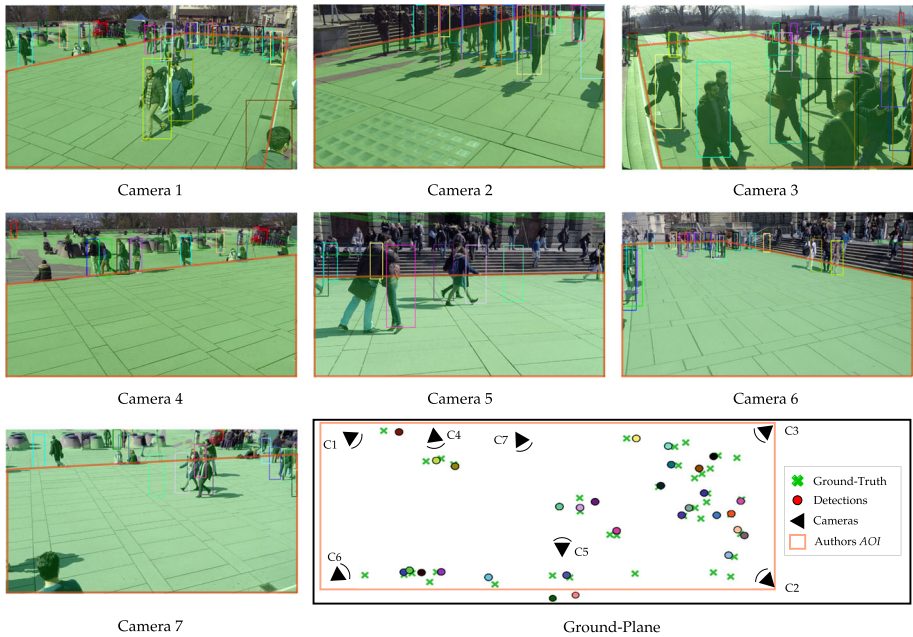


Fig. 10 State-of-the-art Comparison: Qualitative results from a sample frame on Wildtrack dataset. For representation reasons, camera frames depict adapted bounding boxes via the semantic-driven back-projection stage, although this stage is not used for evaluation in the Wildtrack dataset. In addition, the figure depicts the automatically obtained AOI superimposed in green and, finally, the manually annotated AOI proposed by the authors [4, 5] (area delimited by red lines). The last image represents the cameras’ positions, the obtained detections and the authors’ ground truth and AOI over the ground plane. Pedestrians are identified with different colors (one per detection) along views and ground plane. Better viewed in color

track data and use manually annotated detection constrains. Our method, using EfficientDet detections, improves a 45.45% with respect to Pre-DeepMCD—the second ranked, which is an end-to-end deep learning architecture trained on the PETS dataset. However, algorithms explicitly trained on data from the Wildtrack dataset, i.e., DenseNet-DeepMCD, ResNet-DeepMCD, Top-DeepMCD and Deep-Occlusion, outperform the proposed method, in our opinion for two main reasons:

- First, the qualitative results presented in Fig. 10 suggest that the results in Table 4 are highly biased by the authors’ manually annotated area. The proposed method obtains a broader AOI (Fig. 10, green area) than the one provided by the authors (Fig. 10, red area). Although the automatically obtained AOI seems to be better fitted to the ground floor in the scene than the manually annotated one, the performance of our method decreases because ground-truth data is reported only on the manually annotated area. Thereby, our true positive detections out of this area result in false positives in the statistics (see Fig. 10, cameras 1 and 4).
- Second, they learn their occlusion modeling and their inference ground occupancy probabilistic models specifically on the Wildtrack scenario using samples from the dataset. This training, as any fine-tuning procedure in deep neural networks, is highly effective, as indicated by the increase in performance resulting from the use of the same architecture but adapted for the Wildtrack scenario (compare the results of Pre-DeepMCD and Top-DeepMCD). This training requires the use of human-annotated detections in each

scenario, hindering the scalability of these solutions and its application to the real world. The proposed approach, on the other hand, performs equally without the need of being adapted for every target scenario reported in this paper.

Respect to the first issue, i.e., the effect of using the proposed automatically extracted AOI instead of the one provided by the dataset and used by the rest of the methods, in order to obtain a fairer comparison, we have included in Table 4 the result of the proposed method evaluated on the authors' AOI using our top ranked method, i.e., using EfficientDet baseline. As it can be observed, when using the authors' AOI for evaluation, the proposed method outperforms TOP-DeepMCD by a 8.33% and DenseNet-DeepMCD by a 3.17% ranking the third best method on the Wildtrack dataset without requiring a dataset specific fine-tuning stage as the two above it. In addition, performance with respect to GMC-3D [22], which replicates the previous version of the proposed method with the addition of person re-identification features, is increased a 17.85%.

On average, and contrary to state-of-the-art approaches, the proposed method adapts to different target scenarios without needing a separate training stage for each situation, with the consequent reduction of computational resources and time, and neither requiring a manually annotated area of interest.

5 Conclusions

This paper describes a novel approach to perform pedestrian detection in a multi-camera recorded scenario. First, the adapted strategies for the temporal and spatial aggregation of semantic cues, along with homography projections, are used to obtain an estimation of the ground plane. Through this process, a broader, accurate and role-annotated area of interest (AOI) is automatically defined. Per-camera detections, obtained by a state-of-the-art detector, are projected to the reference plane, and those laying outside the obtained AOI are filtered-out. A fusion approach based on creating connected components on a graph representation of the detections is used to combine per-camera detections yielding global pedestrian detection. Then, a semantic-driven back-projection method handles occlusions and uses semantic cues to globally refine the location and size of the back-projected detections by aggregating information from all the cameras. The results on a broad set of scenarios confirm that the method outperforms every other compared multi-camera not deep-learning method and also every deep-learning method not adapted to the target dataset, even with different baseline algorithms. The proposed method performs close to scenario-tailored methods, but without their training stage, which highly hinders their straight use in new scenarios. In overall, the results suggest that the proposed approach is able to obtain accurate, robust, tight-to-object and generic pedestrian detection in varied scenarios, included crowded ones.

Acknowledgements This study has been partially supported by the Spanish Government through its TEC2017-88169-R MobiNetVideo project.

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory

regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alahi A, Jacques L, Boursier Y, Vanderghenst P (2011) Sparsity driven people localization with a heterogeneous network of cameras. *J Math Imaging Vision* 41(1):39–58
2. Aliakbarpour H, Prasath VBS, Palaniappan K, Seetharaman G, Dias J (2016) Heterogeneous multi-view information fusion: review of 3D reconstruction methods and a new registration with uncertainty modeling. *Science* 4:8264–8285
3. Baqué P, Fleuret F, Fua P (2017) Deep occlusion reasoning for multi-camera multi-target detection. In: *IEEE international conference on computer vision (ICCV)*, pp 271–279
4. Chavdarova T, Baqué P, Bouquet S, Maksai A, Jose C, Bagautdinov T, Letry L, Fua P, Van Gool L, Fleuret F (2018) Wildtrack dataset. <https://cvlab.epfl.ch/data/wildtrack>
5. Chavdarova T, Baqué P, Bouquet S, Maksai A, Jose C, Bagautdinov T, Letry L, Fua P, Van Gool L, Fleuret F (2018) WILDTRACK: a multi-camera HD dataset for dense unscripted pedestrian detection
6. Chavdarova T, Fleuret F (2018) Epfl rlc dataset. <https://cvlab.epfl.ch/data/rlc>
7. Chavdarova T, et al (2017) Deep multi-camera people detection. In: *2017 16th IEEE international conference on machine learning and applications (ICMLA)*, pp 848–853. *IEEE*
8. Delannay D, Danhier N, De Vleeschouwer C (2009) Detection and recognition of sports (wo)men from multiple views. In: *ACM/IEEE international conference on distributed smart cameras (ICDSC)*, pp 1–7. *IEEE*
9. Dollár P, Wojek C, Schiele B, Perona P (2009) Pedestrian detection: a benchmark. In: *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pp 304–311. *IEEE*
10. Dollar P, Wojek C, Schiele B, Perona P (2012) Pedestrian detection: an evaluation of the state of the art. *IEEE Trans Pattern Anal Mach Intell* 34(4):743–761
11. Everingham M, Van Gool L, Williams CK, Winn J, Zisserman A (2010) The pascal visual object classes (voc) challenge. *Int J Comput Vision* 88(2):303–338
12. Ferryman J, L. Crowley J, Shahrokni A (2018) Pets 2009 dataset. <http://www.cvg.reading.ac.uk/PETS2009/a.html>
13. Fleuret F, Berclaz J, Lengagne R (2018) Epfl terrace dataset. <https://cvlab.epfl.ch/data/pom>
14. Fleuret F, Berclaz J, Lengagne R, Fua P (2008) Multicamera people tracking with a probabilistic occupancy map. *IEEE Trans Pattern Anal Mach Intell* 30(2):267–282
15. Franco J.S, Boyer E (2005) Fusion of multiview silhouette cues using a space occupancy grid. In: *IEEE international conference on computer vision (ICCV)*, vol 2, pp 1747–1753. *IEEE*
16. García-Martín Á, Martínez JM (2015) People detection in surveillance: classification and evaluation. *IET Comput Vision* 9(5):779–788
17. Hartley R, Zisserman A (2003) *Multiple view geometry in computer vision*. Cambridge University Press, Cambridge
18. He K, Zhang X, Ren S, Sun, J (2016) Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pp 770–778
19. Huang G, Liu Z, Van Der Maaten L, Weinberger KQ (2017) Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 4700–4708
20. Khan SM, Shah M (2009) Tracking multiple occluding people by localizing on multiple scene planes. *IEEE Trans Pattern Anal Mach Intell* 31(3):505–519
21. Kim K, Davis L.S (2006) Multi-camera tracking and segmentation of occluded people on ground plane using search-guided particle filtering. In: *European conference on computer vision (ECCV)*. Springer, pp 98–109
22. Lima J.P, Roberto R, Figueiredo L, Simoes F, Teichrieb V (2021) Generalizable multi-camera 3d pedestrian detection. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 1232–1240
23. Lin T.Y, Maire M, Belongie S, Hays J, Perona P, Ramanan D, Dollár P, Zitnick C.L (2014) Microsoft coco: Common objects in context. In: *European conference on computer vision*. Springer, pp 740–755
24. Long J, Shelhamer E, Darrell T (2015) Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 3431–3440
25. Long X, Deng K, Wang G, Zhang Y, Dang Q, Gao Y, Shen H, Ren J, Han S, Ding E, et al (2020) Pp-yolo: An effective and efficient implementation of object detector. *arXiv preprint arXiv:2007.12099*

26. Lopez-Cifuentes A, Escudero M, Bescos J (2018) Automatic semantic parsing of the ground-plane in scenarios recorded with multiple moving cameras. *IEEE Signal Process Lett* 25(10):1495–1499
27. López-Cifuentes A, Escudero-Viñolo M, Bescós J, Carballeira P (2018) Semantic driven multi-camera pedestrian detection. *arXiv preprint* [arXiv:1812.10779v1](https://arxiv.org/abs/1812.10779v1)
28. Ning C, Menglu L, Hao Y, Xueping S, Yunhong L (2020) Survey of pedestrian detection with occlusion. *Complex Intell Syst* 5:1–11
29. Peng P, Tian Y, Wang Y, Li J, Huang T (2015) Robust multiple cameras pedestrian detection with multi-view bayesian network. *Pattern Recogn* 48(5):1760–1772
30. Priscilla C.V, Sheila S.A (2019) Pedestrian detection-a survey. In: *International conference on information, communication and computing technology*. Springer, pp 349–358
31. Redmon J, Farhadi A (2018) Yolov3: an incremental improvement. *arXiv preprint* [arXiv:1804.02767](https://arxiv.org/abs/1804.02767)
32. Ren S, He K, Girshick R, Sun J (2015) Faster R-CNN: Towards real-time object detection with region proposal networks. In: *Advances in neural information processing systems (NIPS)*
33. Stiefelhagen R, Bernardin K, Bowers R, Garofolo J, Mostefa D, Soundararajan P (2006) The clear 2006 evaluation. In: *International evaluation workshop on classification of events, activities and relationships*. Springer, pp 1–44
34. Szegegy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A (2015) Going deeper with convolutions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 1–9
35. Tan M, Pang R, Le Q.V (2020) Efficientdet: Scalable and efficient object detection. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 10781–10790
36. Tao A, Sapra K, Catanzaro B(2020) Hierarchical multi-scale attention for semantic segmentation. *arXiv preprint* [arXiv:2005.10821](https://arxiv.org/abs/2005.10821)
37. Utasi Á, Benedek C (2013) A bayesian approach on people localization in multicamera systems. *IEEE Trans Circuits Syst Video Technol* 23(1):105–115
38. Wang C.Y, Bochkovskiy A, Liao H.Y.M (2020) Scaled-yolov4: Scaling cross stage partial network. *arXiv preprint* [arXiv:2011.08036](https://arxiv.org/abs/2011.08036)
39. Xu Y, Liu X, Liu Y, Zhu S.C (2016) Multi-view people tracking via hierarchical trajectory composition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pp 4256–4265
40. Yu F, Koltun V (2016) Multi-scale context aggregation by dilated convolutions. In: *ICLR*
41. Zhang H, Wu C, Zhang Z, Zhu Y, Zhang Z, Lin H, Sun Y, He T, Mueller J, Manmatha R, et al (2020) Resnest: Split-attention networks. *arXiv preprint* [arXiv:2004.08955](https://arxiv.org/abs/2004.08955)
42. Zhao H, Shi J, Qi X, Wang X, Jia J (2017) Pyramid scene parsing network. In: *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pp 2881–2890
43. Zhou B, Zhao H, Puig X, Fidler S, Barriuso A, Torralba A (2017) Scene parsing through ade20k dataset. In: *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pp 633–641
44. Zhu X, Su W, Lu L, Li B, Wang X, Dai J (2020) Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint* [arXiv:2010.04159](https://arxiv.org/abs/2010.04159)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Alejandro López-Cifuentes obtained a Degree in Telecommunications Engineering in 2015 at the Universidad Autónoma de Madrid (Spain). In 2017 he received the titles belonging to the International Joint Master Program in Image Processing and Computer Vision (IPCV) at the universities of Budapest (Hungary), Bordeaux (France) and Madrid (Spain). From 2017 to the present, he has been with the Video Processing and Understanding Lab (VPU-Lab) at the Universidad Autónoma de Madrid as a researcher. He has participated in a national project related to video surveillance and analysis of long-term video sequences (HA-Video). His lines of research are focused on context-driven image and video analysis.



Marcos Escudero-Viñolo was born in Madrid, Spain in 1984. He made his under grade studies in the Ramiro de Maeztu high school in Madrid. Then he started his studies in Electric Engineering and Computer Science at the Universidad Autónoma de Madrid in 2002 and in 2006 he spent a year as an Erasmus student in the Mälardälens Högskola at Västerås, Sweden while he received his title as Ingeniero de Telecomunicación (M.S. degree in Electrical Engineering) in 2008 and his MPhil degree in Multimedia Signal Processing at 2009, both from the Universidad Autónoma de Madrid. In 2007 he joined the Video Processing and Understanding Lab (VPU-Lab) to develop his grade Dissertation under the supervision of PhD student Fabrizio Tiburzi and PhD professor Jesús Bescós. From then in advance he has been working as a PhD student in the VPU-Lab under the supervision of PhD professor Jesús Bescós. In the line of his thesis, he has collaborated in European projects (MESH) and National projects (CENIT-VISION). According to his research lines, they are mainly focussed in the area of video anal-

ysis, both in the compress and uncompress domains, targeting to develop automatic video content translators. With this aim, his efforts are focused in the detection, extraction, localization, tracking and characterization of the objects present in the scene.



Jesús Bescós graduated in Telecommunication Engineering in 1993, he obtained the PhD in the same field in 2001, both titles from the Polytechnic University of Madrid (UPM). He was a member of the Image Processing Group of the UPM (1993-2002) and an Assistant Professor at the UPM (1997-2002), and he is Associate Professor at the Autonomous University of Madrid (since 2003), where he jointly coordinates the Video Processing and Understanding Lab. His research interests include video sequence analysis, content-based video indexing, 2D and 3D computer vision, etc. He has been involved in European projects related to museums (eg, RACE-Rama), education (eg, ET-Trends), audiovisual content analysis (eg, ACTS-Hypermedia, ICTS-AceMedia and Mesh), virtual reality (eg, IST) -Slim-VRT), etc., and in national projects, both in the framework of public calls for R&D&I and in the framework of collaborations with the private sector, leading work teams, several doctoral theses, and giving rise to more than fifty publications in conferences and scientific journals. He is a regular evaluator of calls for national and international projects (ICIP, ICASSP, CBMI, SAMT, etc.), and articles sent to conferences and journals (IEEE Tr. On CSVT, Tr. On Multimedia, Tr. On Image Processing, etc.). Within his management activity he has held the post of Deputy Director of Postgraduate and New Degrees of the Polytechnic School (EPS-UAM) (2008-2012), the coordination of the Doctoral Program in Computer Engineering and Telecommunications, the position of Secretary of the Department of Electronic Technology and Communications (2012-2014), and Vice-Rector of Postgraduate Studies of the Autonomous University of Madrid (2014-2017), with competences in Doctoral Schools and Doctorate Programs.



Pablo Carballeira received the Telecommunication Engineering degree (five years engineering program) in 2007 from the Universidad Politécnica de Madrid (UPM). He received the Communications Technologies and Systems Master degree (two year MS program) and the Ph.D. degree in Telecommunication from the UPM in 2010 and 2014 respectively. Since October 2017 he is an Assistant Professor in Universidad Autónoma de Madrid (UAM) and a member of the Video Processing and Understanding Lab (VPULab). His research interests include image processing and video coding, focusing on multiview and free-navigation video, computer vision and compressed sensing. He has been involved since 2008 in the standardization activities from ISO's Moving Picture Experts Group (MPEG), related to multiview and free-navigation video.