Real-time Embedded Person Detection and Tracking for Shopping Behaviour Analysis

R. Schrijvers^{1,2}, S. Puttemans¹, T. Callemein¹, and T. Goedemé¹

¹ EAVISE, KU Leuven, Jan De Nayerlaan 5, 2860 Sint-Katelijne-Waver, Belgium

² PXL Smart-ICT, Hogeschool PXL, Elfde Liniestraat 24, 35000 Hasselt, Belgium
{`steven.puttemans`, `toon.goedeme`, `timothy.callemein`}@kuleuven.be
`robin.schrijvers@pxl.be`

Abstract. Shopping behaviour analysis through counting and tracking of people in shop-like environments offers valuable information for store operators and provides key insights in the stores layout (e.g. frequently visited spots). Instead of using extra staff for this, automated on-premise solutions are preferred. These automated systems should be cost-effective, preferably on lightweight embedded hardware, work in very challenging situations (e.g. handling occlusions) and preferably work real-time. We solve this challenge by implementing a real-time TensorRT optimized YOLOv3-based pedestrian detector, on a Jetson TX2 hardware platform. By combining the detector with a sparse optical flow tracker we assign a unique ID to each customer and tackle the problem of losing partially occluded customers. Our detector-tracker based solution achieves an average precision of 81.59% at a processing speed of 10 FPS. Besides valuable statistics, heat maps of frequently visited spots are extracted and used as an overlay on the video stream.

Keywords: Person Detection · Person Tracking · Embedded · Real-time

1 Introduction

Mapping the flow and deriving statistics (e.g. the amount of visitors or the time spent in a store) of people visiting shop-like environments, holds high value for store operators. To this day, people counting in retail environments is often being accomplished by using cross-line detection systems [24] or algorithms counting people through virtual gates [32].

To accurately count visitors, one can place a computer in the network of the store with access to already available security cameras, deploying software that detects and tracks people in the store automatically and stores the results on a central storage system (e.g. an in-store server or the cloud). In order to run these software solutions, one needs expensive and bulky dedicated computing hardware, frequently covered by a desktop GPU (e.g. NVIDIA RTX 2080). On the other hand, the recent availability of lightweight and affordable embedded GPU solutions, like the NVIDIA Jetson TX2, can be a valid alternative. This is the main motivator to build an embedded and cost-effective people detection

Fig. 1. Examples of (left) the viewpoint from our camera setup in the store and (right) the generated heat map for the store owner using the combined detector-tracker unit.

and tracking solution. An example of the camera viewpoint from the designed setup can be seen in the left part of Figure 1.

The remainder of this paper is organized as follows. Section 2 discusses related work on person detection and tracking, along with available embedded hardware solutions. Section 3 provides details about the proposed implementation, while experiments and results are discussed in section 4. Finally, section 5 summarizes this work and discusses useful future research directions.

2 Related Work

The related work section is subdivided in three subsections, each focusing on a specific subtopic within this manuscript. Subsection 2.1 starts by discussing literature on person detection, where subsection 2.2 continues on specific optimizations in this technology for embedded hardware. Finally subsection 2.3 focuses on the person tracking part.

2.1 Person Detection

Robust person detection solutions have been heavily studied in literature. Originally person detection made use of handcrafted features, combined with machine learning to generate an abstract representation of the person [6, 9, 47, 15, 12, 13]. While these approaches showed promising results, their top-accuracy and flexibility make that these algorithms are not suited for the very challenging conditions in which they will be deployed in this application. Dynamic backgrounds, illumination changes and different store lay-outs are only some of the reasons of a high rate of false positive detections. While re-training the detectors for every store layout would theoretically solve the issue, this isn't a cost-effective solution for companies.

Convolutional neural networks for detecting persons in images offer potentially more robust solutions [20, 31, 38, 39, 17]. By automatically selecting the most discriminate feature set based on a very large set of training data, they quickly pushed the traditional approaches into the background. In

literature these deep learning based approaches are subdivided in two categories. Single-shot approaches [31, 38–40] solve the detection task by classifying and proposing bounding boxes in a single feed-forward step through the network. Multi-stage approaches [17, 19, 16, 41] include separate networks that first generate region proposals before classifying the objects inside the proposals. A more recent and promising multi-stage approach is Trident-Net [30], where scale-specific feature maps are built into the network. Single-stage approaches tend to have a more compact and faster architecture making them preferred in lightweight embedded hardware solutions.

2.2 Optimizing for embedded hardware

Solving the task of person detection on embedded platforms has been an active research topic in recent years. A common approach is to enable deep learning on embedded platforms by studying more compact models such as Tiny-YOLOv2 [39]. While these compacter models are able to run at decent speeds on embedded hardware, they tend to lose some percentage in accuracy compared to their full counterparts (e.g. YOLOv2). Besides going compact, several approaches [43, 49, 34] optimize the architecture further by looking at the indirect computation complexity and address efficient memory access and platform characteristics. More recent embedded object detection algorithms introduce optimized filter solutions, like depth-wise separable convolutions [22] and inverted residuals [42], to further optimize the performance of these embedded solutions.

With the uprising of FPGA chips for deep learning, several architectures like [23] try to reduce the number of parameters of the models even further, having a 50x reduction in parameters compared to classic models like [28]. These specifically designed hardware chips also allow for fixed-point 16-bit optimizations through OpenCL [48]. While being very promising, these FPGA systems still lack flexibility. They are in most cases designed for a very specific case, and are thus no cost-effective solution for our problem.

While many of these embedded object detectors are explicitly shaped for running on embedded hardware, detection accuracies are still lower compared to their traditional full blown CNN counterparts (eg. YOLOv3). Taking into account that our solution should perform person detection at a minimal processing speed of 10 FPS and a minimal accuracy of 80%, we opt for integrating an optimized embedded implementation of the YOLOv3 object detector [46] in our pipeline. The architecture smartly combines several optimizations based on mobile convolutions in PyTorch and TensorRT compilations.

The introduction of the NVIDIA Jetson embedded GPU enabled balancing of local-processing, low power consumption and throughput in an efficient way. [35] discusses several implementations of deep learning models on the Jetson platform while considering a range of applications, e.g. autonomous driving or traffic surveillance. The work focuses on obtaining low latency, to make detectors useful for providing valuable real-time feedback. The advantages in both FPGA and embedded GPU systems are the fact that they are substantially more space-efficient and less power-consuming than desktop GPUs.

2.3 Person Tracking

Tracking objects in videos has been studied across many fields. Whenever quasi-static environments occur, motion detection through robust background subtraction is used to identify moving objects in images [11, 8, 29, 27]. More challenging cases involve dynamic environments, e.g. tracking objects in autonomous driving applications or drones. [7] solves the tracking task using detection results of a lightweight object detection algorithm combined with euclidean distance equations, GPS-locations and data from an inertial measurement unit. Because research is moving more and more towards highly accurate CNN-based object detectors, a new sort of tracking is introduced, called tracking-by-detection. By calculating an intersection over union (IoU) between the detections of two consecutive frames and applying a threshold, one can decide if we are dealing with the same object.

Where tracking-by-detection works well in many situations, there are also some downsides with these approaches. When people are missed in several subsequent frames by the initiating detector, the person gets lost during tracking. However, in our application of real-time costumer analysis, we want to keep a unique ID for each detected person, and thus need to fill in this gap automatically. [2] solves this issue by integrating a SVM-classifier with an optical-flow-based tracker. With the uprising of CNN-based detectors, CNN-based object trackers have also been proposed. Limb-based detection with tracking-by-detection is proposed in [1], while [4] explains an approach using IoU information from multiple object detectors. In addition, [36] combines location information and similarity measures to perform tracking-by-detection.

Feature-based tracking algorithms are a valid alternative. Sparse optical flow calculates the flow of objects based on a sparse feature set [33, 45] and have been successfully deployed for person tracking [26, 44, 11]. Dense optical flow [14] calculates the flow for every point in the detection and is thus more computationally complex. [37] propose an object tracker based on weakly aligned multiple instance local features and an appearance learning algorithm.

A final range of trackers use online learning, where a tracker learns the representation of the object on the fly through a classifier when initialized with a bounding box [18, 3]. [21] learns more efficiently from these overlapping patches by using kernelized correlation filters, while [25] enables tracking failure detection by tracking objects both forward and backward in time, measuring and qualifying the discrepancies between the two trajectories.

[10] gives a clear general overview of state-of-the-art trackers and their accuracy on a public dataset. For this work, we integrated three object trackers: sparse optical flow, kernelized correlation filters and median flow.

This paper proposes a complete off-the-shelf solution, taking into account the fact that we work on a compact embedded system with limited power consumption, achieve real-time performance (e.g. minimally processing at 10 FPS) and obtain acceptable accuracies (e.g. over 75%) in a single setup. While many of these sub-tasks have been discussed already in literature, to our best knowledge, there isn't a single publication that proposes such end-to-end solution

Fig. 2. Pipeline overview with the proposed batch processing approach, passing frames from the image buffer to the detection or tracking unit, based on the batch iterator.

for in-store customer behaviour analysis. On top of that we are combining all parts of the pipeline in a batch system that consumes the resources (both CPU and GPU) of the host system as efficient as possible. This is the biggest novelty this work introduces and will be discussed in further detail in subsection 3.1.

3 Methodology

The goal of this paper is to map the flow of costumers in shop-like environments, with a special focus on detecting and tracking individual costumers using a unique ID. Preferably the system should run minimally at a processing speed of 10 frames per second on top of a compact, easy-to-use, embedded platform (e.g. Jetson TX2). Therefore, we propose a multi-threaded YOLOv3 TensorRT optimized person detector combined with a fast lightweight object tracker.

Figure 2 gives an overview on how our pipeline is implemented on the embedded Jetson TX2 platform. All sub-parts will be discussed in detail in the following subsections.

3.1 Person detection and batch processing

In applications that require real-time performance, one can either choose to maximize throughput or minimize latency. Throughput focuses on the number of images that are processed in a given time slot, while latency focuses on the time needed to process a single image. Given the goal of analyzing the flow of costumers in a given security camera stream, and in the future even in multiple camera streams, throughput is in our case more important than latency. We maximize throughput by performing a batch processing approach that takes advantage parallel GPU processing. As seen in Figure 2, an image buffer is used to collect all incoming images from the camera stream, at a 1280×720 resolution, which are downsampled to a 512×288 input resolution before moving to the detector and the tracker unit. We could sent these images directly to our optimized YOLOv3 implementation [46], but since it is only capable of processing

data at a maximum speed of 8FPS at this resolution, we skip several frames when collecting a batch that needs to be processed by the YOLOv3 detector.

Take for example a image buffer of 24 images. From those 24 images, only 8 images are selected as a batch that gets passed to the detector unit. The remaining 16 images are sent to the tracker unit, which waits for the processed batch of the detector, so that it can track each detected person throughout the 24 frames series, based on the detections in those 8 reference frames. In order to improve processing speed, both detector and tracker unit are implemented in separate threads, to reduce the processing time delays again as much as possible.

3.2 Person tracking

Given our hardware setup of a Jetson TX2 embedded platform, we need to carefully consider our hardware resources. Since we're using the on-board GPU to efficiently processes batches of 8 images with the object detector unit, the decision was made to implement a CPU-based tracker to divide the workload as good as possible between the different processing units. We augment the detector unit with a lightweight CPU-based Lukas Kanade Sparse Optical Flow tracker (based on the OpenCV4.1 implementation [5]). This allows us to run the tracker on the CPU, while simultaneously running the GPU-focused YOLOv3 object detector. To achieve the minimally required frame rate of 10 FPS, we downscale the original input frames to a resolution of 512×288 pixels.

In between detection frames, two extra frames remain which are only processed by the tracker unit, and thus so far have no knowledge of the location of the detected objects. For each detection inside a detection frame, key feature points are calculated, generating a sparse representation of that detection. Those points are passed into an optical flow mechanism that is used to generate predictions of the locations of these keypoints in the in-between tracker frames. Finally the detections in the next detection frame are used to apply a mapping between the predictions and the detections, possible proposing a slight correction of the bounding box location. Figure 3 illustrates how this detector-tracker combination is working on a sample from the security camera stream. The tracker only produces predictions in the next frame for detections with a detection

Fig. 3. The detector-tracker pipeline, using tracking-by-detection with interleaved tracker proposals: (solid red) detections (interleaved green) tracker proposals (solid green) tracker proposal accepted as detection in case of a missing detection.

probability higher than 10%. This allows us to carefully choose the working conditions of the detector and avoid as much false negative detections (persons getting missed) as possible. Detections with a lower confidence will simply be ignored by the tracker and immediately be removed from the tracker memory.

Applying this approach also introduces issues. When customers move outside the field-of-view of the camera stream, their tracking information is kept in memory by the tracker, clogging up the tracker memory and resulting in locally dangling tracks. In order to avoid these dangling tracks, after 5 consecutive detection misses, we force the tracker to remove the tracking information and forget the track altogether, which will end the tracking of that object.

Besides the sparse optical flow approach, two more CPU-based tracking implementations (Kernelized Correlation Filters and Median Flow) are tested to compare tracking robustness in shop-like environments. Both are initialized by the detections of the neural network, just like the sparse optical flow approach. For both trackers, the same input resolution of 512x288 pixels has been used.

3.3 Tracker memory

Since the selected tracker implementations only takes the information between two consecutive frames into account, we are risking losing a person with a specific ID from tracker memory (e.g. due to exceeding the threshold of missed frames), that can in a later stage, be picked up again by the detector. In order to avoid a new ID being assigned, we introduce a method to keep and match these lost tracks in memory. In case a track is deleted from memory, it first changes its status to lost. A lost track is remembered for maximally 5 seconds, and is recovered whenever a new detection appears in a location close to the last known location of a deleted track. A match between a lost track and a newly generated track is made by meeting one of the following conditions: the new location lies within a radius of 200 pixels around the last known location, or the new location lies in the quadrant of which the gradient of the lost track points to. In case of multiple new detections in the same quadrant, the detection with the shortest (perpendicular) distance to the gradient is preferred.

3.4 Heat map extraction

On top of providing statistics, about the exact location and path followed by customers, users of the system would also like to have some sort of visual confirmation on the activity in their stores. Therefore we propose a heat map based system that gives an overview of the store spots that are most frequently visited within a specific time slot. The heat maps are extracted by mapping the pixels in the detected and tracked bounding boxes on top of a visual layout of the store. The used object detector is known for its jittery bounding boxes, with inconsistent widths and heights. By simply incrementing pixel locations each time a detection box matches a pixel, we end up with a very jittered and visually unpleasing heat map, which we would like to avoid. We solve this by using a weighted increment of the pixels falling within a bounding box. By

taking the ratio between the bounding box width and height into account and by giving the center pixel a maximum weight that degrades towards the borders of the bounding box, we end up with a increment value that is less influenced by the jitter, resulting in a visually pleasing heat map. After all frames within the given time slot have been processed, heat maps are normalized to get a meaningful color overlay, as seen in Figure 1.

4 Experiments and results

In this section, we evaluate the proposed embedded detector-tracker approach used for in-store customer behaviour analysis by performing four different experiments. We start by evaluating the detector unit and decide on the final deep learned architecture for the detector unit. Next, we add different trackers to the problem to cope with missing detections, optimize parameters and have a look at the influence of changing internal parameters. We then integrate our optimal detector-tracker combination into our batch processing pipeline, to gain processing speed. Finally we take a shot at generating visual heat maps, giving confirmation of frequently visited places in the stores.

4.1 Detector and tracker evaluation

A set of 5000 in-store images were collected via the security camera streams inside the store and annotated with ground truth labels. We evaluate both the TensorRT optimized YOLOv2 and YOLOv3 implementations on this set by running them through our NVIDIA Jetson TX2 platform. The left hand side of Figure 4 shows the difference in obtained average precision (AP). We notice that the YOLOv3 architecture (AP=84.36%) is performing much better than the YOLOv2 architecture (AP=64.26%), with a tremendous rise in AP of 20%. This can be explained by the capability of the YOLOv3 architecture to detect persons over a wider variety of scales. Giving the security cameras are mounted to give a single point overview of the store, customers walking further away

Fig. 4. Precision-recall curves of (left) the embedded and optimized YOLOv2 and YOLOv3 implementation (right) the combined detector-tracker implementations based on optical flow, median flow and kernelized correlation filters.

from the camera are not being detected by the YOLOv2 architecture. On top of that, we also notice that the YOLOv3 architecture is better capable of detecting partially occluded persons.

However, even with the YOLOv3 architecture, we do not yet reach the optimal solution (AP=100%). After carefully inspecting the missed detections, we notice the detect still drops persons when they are crossing each other or when they are move behind counters and thus get partially occluded. To solve these issues we run our combined detector-tracker units, for which the AP curves are visible in Figure 4 at the right hand side. We notice that adding the tracker unit slightly helps the detector improve in AP for all the combinations. Based on these experiments we decided to stick with the best performing combination, being YOLOv3 as detector and optical flow as tracker.

One could however argue that the influence of the tracker can also be increased if predicted tracks are kept in memory longer. Until now a track is rejected as soon as the detector unit is unable to find a matching detection in 5 consecutive detection frames. To show the influence of this parameter, we increased this threshold towards 10 missed detection frames and again ran the precision-recall evaluation for the optimal combination again, as seen in Figure 5 on the left hand side. This change again provides a small boost to the AP of the optimal solution towards 85.35%. However, the influence of varying this parameter even further should be investigated in future research.

However, up till now our experiments were targeted at getting the highest AP possible with a detector-tracker combination. Since the original TensorRT YOLOv3 optimized implementation is only able to run the network at an average speed of 3FPS, adding the tracker reduced this further towards 2FPS. Since we want to aim for a minimal processing speed of 10 FPS, this did not suffice. As discussed in subsection 3.1 we applied a batch processing pipeline to overcome this low processing speed. The obtained AP is seen in the right part of Figure 5. While increasing our processing speed towards 9.8 FPS by simultaneously reducing the AP towards 81.59%. In order to see if this change might influence the other detector-tracker combinations, we also evaluated those in the batch

Fig. 5. Influence of (left) increasing the miss threshold for the tracker and (right) applying our batch processing pipeline to gain processing speed.

Optical Flow	Median Filtering	Kernelized Correlation Filters
9.8 FPS	4.7 FPS	2.1 FPS

Table 1. Comparing average detector-tracker processing speeds over the dataset.

processing setup. Results of the obtained average tracking speeds can be seen in Table 1, which confirmed that we picked the correct combination before. Median filtering and kernelized correlation filters perform worse due to the linear scale-up in relation to the amount of detected persons, while this is not true for optical flow. However, we acknowledge that this could depend on the implementation of the tracker used, and that further research on this is necessary.

4.2 Generating heat maps of frequently visited places

Figure 1 illustrates the resulting heat maps that indicate the most frequently visited spots in the store. This gives a clear indication to the shop owners which counters attract customers better than others. Generating heat maps with the detector unit only tends to generate heat maps with visually only very small differences. While the tracker helps in obtaining higher processing speeds in capturing some of the missed detections due to occlusion, it does not really benefit the heatmap generation part.

5 Conclusion

In this work we propose a solution for mapping the flow of customers in shop-like environments in real-time based on person detection and tracking, implemented on an power-efficient and compact embedded platform. In this work we proposed a novel batch processing approach that efficiently uses the power of both GPU and CPU simultaneously, to run a state-of-the-art TensorRT optimized YOLOv3 object detector combined with a sparse optical flow object tracker on the embedded Jetson TX2 platform. We achieve an average precision of 81.59% at a processing speed of 10 FPS on a dataset of challenging real-life imagery acquired in a real shop. While this is not yet the optimal solution, given the challenging conditions (dynamic backgrounds, illumination changes and different store lay-outs), this numbers are quite impressive. On top of that we provide visually pleasing heat maps of the store, giving owners valuable insights to customer behavior. To improve the current results, we are planning to add a person re-identification pipeline to the solution and keep an eye out to new ways of optimizing deep learning pipelines for embedded platforms. This paper creates a proof-of-concept setup that can be further exploited in multiple camera stream setups, exploring how we can use shared memory buffers and shared detection networks between different embedded platforms efficiently.

6 Acknowledgements

This work is supported by PixelVision, KU Leuven and Flanders Innovation & Entrepreneurship (VLAIO) through a Baekelandt scholarship.

References

1. Andriluka, M., Roth, S., Schiele, B.: People-tracking-by-detection and people-detection-by-tracking. In: Proceedings of CVPR. pp. 1–8 (2008)
2. Avidan, S.: Support vector tracking. In: Proceedings of CVPR. pp. 1064–1072 (2001)
3. Babenko, B., Yang, M.H., et al.: Visual tracking with online multiple instance learning. In: Proceedings of CVPR. pp. 983–990 (2009)
4. Bochinski, E., Eiselein, V., Sikora, T.: High-speed tracking-by-detection without using image information. In: Proceedings of AVSS. pp. 1–6 (2017)
5. Bradski, G., Kaehler, A.: Learning OpenCV: computer vision with the OpenCV library. " O'Reilly Media, Inc." (2008)
6. Breiman, L.: Random forests. Machine learning pp. 5–32 (2001)
7. Chen, P., Dang, Y., et al.: Real-time object tracking on a drone with multi-inertial sensing data. Transactions on ITS pp. 131–139 (2017)
8. Choi, J.W., Moon, D., Yoo, J.H.: Robust multi-person tracking for real-time intelligent video surveillance. Proceedings of ETRI pp. 551–561 (2015)
9. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: Proceedings of CVPR. pp. 886–893 (2005)
10. Dendorfer, P., Rezatofighi, H., et al.: CVPR19 tracking and detection challenge: how crowded can it get? arXiv preprint arXiv:1906.04567 (2019)
11. Denman, S., Chandran, V., et al.: Person tracking using motion detection and optical flow. In: Proceedings of DSPCS. pp. 1–6 (2005)
12. Dollár, P., Appel, R., et al.: Fast feature pyramids for object detection. Proceedings of TPAMI pp. 1532–1545 (2014)
13. Dollár, P., Tu, Z., et al.: Integral channel features. Proceedings of BMVC pp. 91.1–91.11 (2009)
14. Farnebäck, G.: Two-frame motion estimation based on polynomial expansion. In: Scandinavian conference on Image analysis. pp. 363–370 (2003)
15. Felzenszwalb, P.F., McAllester, D.A., et al.: A discriminatively trained, multiscale, deformable part model. In: Proceedings of CVPR. pp. 1–8 (2008)
16. Girshick, R.: Fast R-CNN. In: Proceedings of ICCV. pp. 1440–1448 (2015)
17. Girshick, R., Donahue, J., et al.: Rich feature hierarchies for accurate object detection & semantic segmentation. In: Proceedings of CVPR. pp. 580–587 (2014)
18. Grabner, H., Grabner, M., et al.: Real-time tracking via on-line boosting. In: Proceedings of BMVC. pp. 47–56 (2006)
19. He, K., Gkioxari, G., et al.: Mask R-CNN. In: Proceedings of ICCV. pp. 2961–2969 (2017)
20. He, K., Zhang, X., et al.: Deep residual learning for image recognition. In: Proceedings of CVPR. pp. 770–778 (2016)
21. Henriques, J.F., Caseiro, R., et al.: High-speed tracking with kernelized correlation filters. Proceedings of TPAMI pp. 583–596 (2014)
22. Howard, A.G., Zhu, M., et al.: Mobilenets: efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
23. Iandola, F.N., Han, S., et al.: SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and 0.5 MB model size. arXiv preprint arXiv:1602.07360 (2016)
24. Igurnaisi, R., Merad, D., et al.: People counting based on kinect depth data. In: Proceedings of ICPRAM. pp. 364–370 (2018)
25. Kalal, Z., Mikolajczyk, K., et al.: Forward-backward error: automatic detection of tracking failures. In: Proceedings of CVPR. pp. 2756–2759 (2010)

26. Kanagamalliga, S., Vasuki, S.: Contour-based object tracking in video scenes through optical flow and Gabor features. *Optik* pp. 787–797 (2018)
27. Kowcika, A., Sridhar, S.: A literature study on crowd (people) counting with the help of surveillance videos. *International Journal of Innovative Technology and Research* pp. 2353–2361 (2016)
28. Krizhevsky, A., Sutskever, I., et al.: Imagenet classification with deep convolutional neural networks. In: *Advances of NeurIPS*. pp. 1097–1105 (2012)
29. Lefloch, D., Cheikh, F.A., et al.: Real-time people counting system using a single video camera. In: *Real-Time Image Processing* (2008)
30. Li, Y., Chen, Y., et al.: Scale-aware trident networks for object detection. *arXiv preprint arXiv:1901.01892* (2019)
31. Liu, W., Anguelov, D., et al.: SSD: single shot multibox detector. In: *Proceedings of ECCV*. pp. 21–37 (2016)
32. Liu, X., Tu, P.H., et al.: Detecting and counting people in surveillance applications. In: *Advanced Video and Signal Based Surveillance*. pp. 306–311 (2005)
33. Lucas, B.D., Kanade, T., et al.: An iterative image registration technique with an application to stereo vision. *Proceedings of IJCAI* pp. 674–679 (1981)
34. Ma, N., Zhang, X., et al.: ShuffleNet v2: practical guidelines for efficient CNN architecture design. In: *Proceedings of ECCV*. pp. 116–131 (2018)
35. Mittal, S.: A survey on optimized implementation of deep learning models on the NVIDIA Jetson platform. *Journal of Systems Architecture* (2019)
36. Murray, S.: Real-time multiple object tracking - a study on the importance of speed. *arXiv preprint arXiv:1709.03572* (2017)
37. Pernici, F., Del Bimbo, A.: Object tracking by oversampling local features. *Proceedings of TPAMI* pp. 2538–2551 (2013)
38. Redmon, J., Divvala, S., et al.: You only look once: unified, real-time object detection. In: *Proceedings of CVPR*. pp. 779–788 (2016)
39. Redmon, J., Farhadi, A.: YOLO9000: better, faster, stronger. In: *Proceedings of CVPR*. pp. 7263–7271 (2017)
40. Redmon, J., Farhadi, A.: YOLOv3: an incremental improvement. *arXiv preprint arXiv:1804.02767* (2018)
41. Ren, S., He, K., et al.: Faster R-CNN: towards real-time object detection with region proposal networks. In: *Advances of NeurIPS*. pp. 91–99 (2015)
42. Sandler, M., Howard, A., et al.: Mobilenetv2: inverted residuals and linear bottlenecks. In: *Proceedings of CVPR*. pp. 4510–4520 (2018)
43. Shafiee, M.J., Chywl, B., et al.: Fast YOLO: a fast you only look once system for real-time embedded object detection in video. *arXiv preprint arXiv:1709.05943* (2017)
44. Shashev, D., et al.: Methods and algorithms for detecting objects in video files. In: *MATEC Web of Conferences*. p. 01016 (2018)
45. Shi, J., Tomasi, C.: Good features to track. In: *Proceedings of CVPR*. pp. 593–600 (1993)
46. Vandersteegen, M., Vanbeeck, K., et al.: Super accurate low latency object detection on a surveillance UAV. *arXiv preprint:1904.02024* (2019)
47. Viola, P., Jones, M., et al.: Rapid object detection using a boosted cascade of simple features. *Proceedings of CVPR* pp. 511–518 (2001)
48. Wai, Y.J., Mohd Yussof, Z., et al.: Fixed point implementation of Tiny-YOLOv2 using OpenCL on FPGA. *Proceedings of IJACSA* pp. 506–512 (2018)
49. Zhang, X., Zhou, X., et al.: Shufflenet: an extremely efficient convolutional neural network for mobile devices. In: *Proceedings of CVPR*. pp. 506–512 (2018)