

Deep Multi-Camera People Detection

Tatjana Chavdarova and François Fleuret
Idiap Research Institute and
École Polytechnique Fédérale de Lausanne
Email: {firstname.lastname}@idiap.ch

Abstract—This paper addresses the problem of multi-view people occupancy map estimation. Existing solutions either operate per-view, or rely on a background subtraction pre-processing. Both approaches lessen the detection performance as scenes become more crowded. The former does not exploit joint information, whereas the latter deals with ambiguous input due to the foreground blobs becoming more and more interconnected as the number of targets increases.

Although deep learning algorithms have proven to excel on remarkably numerous computer vision tasks, such a method has not been applied yet to this problem. In large part this is due to the lack of large-scale multi-camera data-set.

The core of our method is an architecture which makes use of monocular pedestrian data-set, available at larger scale than the multi-view ones, applies parallel processing to the multiple video streams, and jointly utilises it. Our end-to-end deep learning method outperforms existing methods by large margins on the commonly used PETS 2009 data-set. Furthermore, we make publicly available a new three-camera HD data-set.

I. INTRODUCTION

Due to the high demand in applications, pedestrian detection differentiated itself as a separate class from object detection, enjoying separate attention from the research community. In spite of recent advances, performance of monocular methods remains limited due to the occlusions that often occur among the individuals. Multi-camera approaches offer a promising extension to resolve the one-view detection ambiguities, provided that the multi-stream information is used jointly to yield the detection estimation. In this paper, we confine our discussion to a set-up of either a single calibrated static camera, or several synchronized cameras with overlapping fields of view.

Surprisingly, currently only few multi-view methods jointly utilize the information across views, herein referred to as *joint* methods. Each of the existing joint methods performs background subtraction pre-processing of the input images, where the goal is to segment the moving objects out of the background, while taking into account pixel-wise time consistency. There exist a vast catalog of background-subtraction methods, specific to certain applications or even illumination regimes, which are often shown to be difficult to tune, error-prone and noisy. Regarding multi-camera people detection in particular, these segmenting methods introduce ambiguities when the foreground blobs are highly interconnected, hence limiting the success of the joint methods to less-crowded applications. Also, apart from persons, other moving objects are being segmented as well.

On the other hand, state of the art monocular detectors leverage the full signal either by using a deep Convolutional

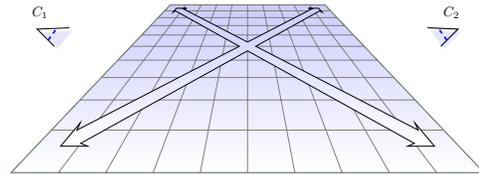


Fig. 1: Illustration of the people occupancy map estimation problem in a static camera setup. As an example we depict two cameras and the area is discretized into a 10×10 grid cells, thus following the notation in the text: $C=2$, $G=100$.

Neural Network (CNN), or by building predictors on top of “deep features” extracted by such a network. Surprisingly, the multi-view occupancy map estimation problem has not been re-visited to incorporate these outperforming methods.

Quite the contrary, the ongoing research on multi-camera people detection incorporates hand-crafted features, whereas to the best of our knowledge real-world industrial applications use monocular CNN pedestrian detector, and yield the final estimation by averaging the separate per-view ones. We presume that this is mostly due to the non-existence of a large-scale multi-camera data-set which would allow for training a multi-input architecture.

As a solution, we propose a method which consists of: (i) fine-tuning a state of the art object detection network on monocular pedestrian detection; (ii) combining several instances of the early layers of that network into a multi-view deep network whose outer layers are trained for multi-view appearance-based joint detection on relatively smaller multi-camera data-set. Our proposed architecture allows for processing the input from the separate views in parallel, and the subsequent interconnection layers allow for automatically learning how features across different views map into each other. In this paper, we focus on per-frame processing, whereas utilising time consistency could be a further extension.

In summary, we: (i) propose the first full deep learning multi-camera people detector; (ii) we empirically demonstrate the superiority of such a method; and we (iii) provide a new three-view data-set with fairly more accurate calibration in terms of consistency of a projection across all of the views.

The monocular methods overview in § II is of those methods which utilize deep learning, whereas a complete overview is given for the more closely related methods to this work - the joint multi-camera people detectors. The problem on

which this paper focuses is formally stated in § III, where we also introduce the multi-view architecture that we propose. Our empirical evaluation presented in § IV demonstrates the encouraging superiority of the deep multi-view detection, and also gives insights about practical considerations which arise.

II. RELATED WORK

A. Deep monocular pedestrian detection

Applying the R-CNN algorithm [1] to monocular pedestrian detection [2] exceeded the state of the art methods at that time on the Caltech pedestrian data-set [3], being the first demonstration that CNNs are well-suited to the task. Later it was shown that the explicit dealing with pedestrian occlusions, by fine-tuning a separate fully convolutional network for sub-parts of the bounding boxes, provides a 50% relative improvement [4]. On top of the part-designated CNNs the authors employ a linear SVM, sparsified to reduce the computational load.

In [5] a Boosting algorithm is proposed, which merges the proposal generation and the detection steps. The cascade learning is formulated as Lagrangian optimisation of a risk accounting for both accuracy and complexity, and integrates both hand-crafted and convolutional features. The downstream classifier of the extension of R-CNN known as Faster R-CNN [6] was recently shown to degrade the performance, and it was proposed to be replaced by boosting forests [7]. The resulting method does not use hand-crafted features and reaches state of the art accuracy.

B. Multi-view occupancy map estimation methods

The first method which uses multi-view streams jointly is the Probabilistic Occupancy Map (POM) [8]. Based on a crude generative model, it estimates the probabilities of occupancy through mean-field inference, naturally handling occlusions. Further, it can be combined with a convex max-cost flow optimization to leverage time consistency [9]. In [10] the problem is re-casted as a linear inverse, regularized by enforcing a sparsity constraint on the occupancy vector. It uses a dictionary whose atoms approximate multi-view silhouettes. To elevate the need of O-Lasso computations, in [11] a regression model was derived which includes solely Boolean arithmetic and sustains the sparsity assumption of [10]. In addition, the iterative method is replaced with a greedy algorithm based on set covering.

In [12] the occlusions are explicitly modeled per view by a separate Bayesian Network, and a multi-view network is then constructed by combining them, based on the ground locations and the geometrical constraints. Although considering crowd analysis, the multi-view image generation of [13] is with a stochastic generative process of random crowd configurations, and then maximum a posteriori (MAP) estimate is used to find the best fit with respect to the image observations.

All of these methods utilise background subtraction pre-processing.

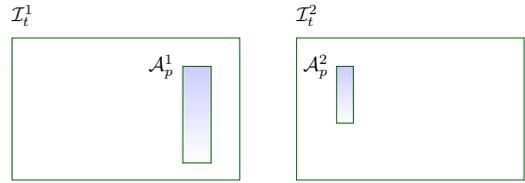


Fig. 2: The input to the model when estimating q_p consists of the cropped regions \mathbf{A}_p of \mathbf{I}_t (see § III-A). The illustrated example is in-line with the one in Fig. 1.

III. DEEP MULTI-VIEW PEOPLE DETECTION

A. Problem definition

We discretize the area common to the fields of view of C cameras in a regular grid of G points or interchangeably - cells (Fig. 1). To estimate the occupancy of a cell p , we consider a cylinder centred at the p th position, whose height corresponds to the one of the humans' average height. We use the cameras' calibration to obtain the cylinder's projections into the views where it is visible. These rectangular projections yield the cropped regions \mathbf{A}_p of \mathbf{I}_t , as illustrated in Fig. 2.

At each time step t we are given a set of images $\mathbf{I}_t = \mathcal{I}_t^1, \dots, \mathcal{I}_t^C$ taken synchronously. Hence, given sub-images cropped for a particular position $p = 1, \dots, G$ in the separate views $\mathbf{A}_p = (\mathcal{A}_p^1, \dots, \mathcal{A}_p^C)$, we aim at estimating the probability of it being occupied, that is

$$q_p = p(X_p = 1 \mid \mathbf{A}_p),$$

where X_p on $\{0, 1\}$ stands for the position p being free and occupied, respectively. In the following, let \mathcal{I} be the set of all possible cropped sub-images.

B. Method

Ideally, given a large scale multi-camera data-set, a multi-stream processing model can be directly trained. However, due to the lack of such data-set, as well as the fact that in practice the annotated data is often scarce, the method we propose takes advantage of the existing larger monocular pedestrian data-set Caltech [3] what allows for better generalization.

1) *Monocular fine-tuning with input-dropout*: Primarily, to learn discriminative features of the pedestrian class, we train a CNN on monocular pedestrian binary classification. To ease the training procedure we recommend starting from a pre-trained object detection CNN. In particular, in our experiments we used either GoogLeNet [14] or AlexNet [15], trained on ImageNet [16] and we replaced the last fully connected layer with a randomly initialized one with two output units. We then fine-tune the resulting network on the Caltech data-set (see § IV-A).

Experimental results on monocular pedestrian detection indicate that state of the art detection accuracy is achieved by explicitly dealing with the occlusions (for example by detecting parts of the body), as explained in § II-A. As we shall see in the following sections, the multi-view architecture finally contains components initialized with the weights of the

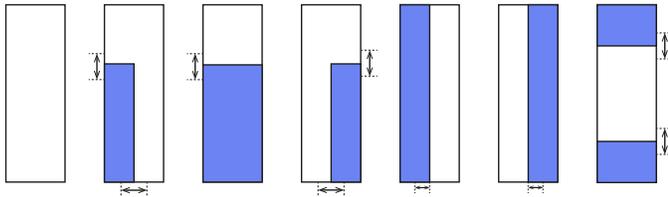


Fig. 3: Input occlusion masks. Among these, the first one represents no occlusion. As the arrows indicate, their width and/or height are restrictively randomized within a margin.

monocular detector. Hence, it is necessary that the robustness to occlusion of the model is implemented efficiently. Thus, we propose a novel technique dubbed *input dropout*, which consists in augmenting the input data by artificially masking part of the signal. While this procedure is adapted to the morphology of the positive class, it has to be carried out on both the positive and the negative samples while training, so as to avoid providing erroneous discriminative cues.

In our experiments in particular we defined 7 rectangular occluding masks—#1 being “no occlusion at all”—(Fig. 3), and for each sample, negative or positive, we pick one of the masks uniformly at random, and replace the pixel it masks with white noise. As for drop-out, this “input drop-out” forces the network toward greater redundancy between representations.

2) *Multi-camera architecture*: Given a CNN monocular detector, we retain d of its layers, resulting in an embedding $\psi : \mathcal{I} \rightarrow \mathbb{R}^Q$, where Q is the number of output units of the d th layer. The concatenation of C such embeddings is a multi-view embedding $\Psi : \mathcal{I}^C \rightarrow \mathbb{R}^{CQ}$, on top of which we train a binary classifier $\Phi : \mathbb{R}^{CQ} \rightarrow \mathbb{R}^2$, as illustrated in Fig. 4.

Type of classifier of Φ . Particular object seen from different perspectives/cameras can be discriminated both with features which are consistent across views: *e.g.* color; and with features that differ across the views: for example a line curved to the left would be curved to the right viewed from the opposite side. The former is more likely when the two cameras are quite close to each other, and the latter in the opposite case. The goal of Φ is to learn how features from different views map to each other. Thus, a natural choice for ϕ is a multi-layer perceptron (MLP), in aim to “wire together what fires together”. In our experiments the output of Ψ is flattened, and Φ is a MLP, and also for demonstrative analysis Φ is sometimes a (forest of) decision tree(s).

Full or partial fine-tuning. To allow the embedding to capture cues discarded for monocular detection but which are informative for multi-view, the complete multi-stream architecture, that is both Ψ and Φ , can be trained, when the latter is differentiable. On the other hand, keeping Ψ fixed makes the capacity of the predictor family low, and hence prevents over-fitting when the training multi-view data-set is small, *e.g.* in the order of few hundreds of examples. Given the size of the data-sets we have used, this approach performed worse than keeping Ψ fixed.

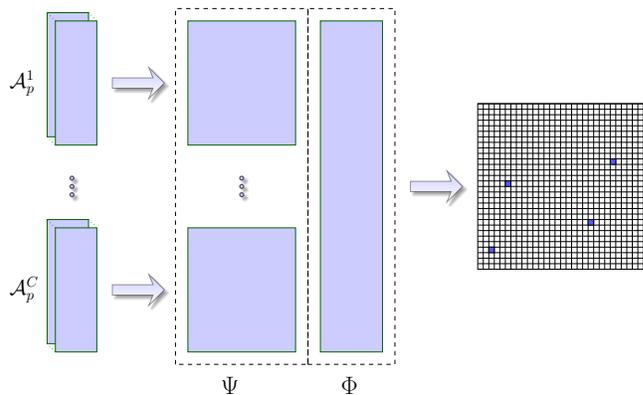


Fig. 4: Multi-view joint occupancy estimation. The input is illustrated on the left; followed by the multi-view architecture consisting of two conceptual parts Ψ and Φ (§ III-B2); and on the rightmost side a top-view grid of the predictions is plotted where a coloured cell represents a detection.

3) *Non Maxima Supression (NMS)*: This step selects the final detections so that out of those candidates overlapping more than a predefined threshold in a particular view, only one remains. Differently than standard implementations, we take into account the detections’ scores, in such a way that the priority of a detection candidate to be selected is proportional to its detection score.

IV. EXPERIMENTS

A. Data-sets and metrics

Caltech10x [3]. For monocular training, we use the Caltech-USA pedestrian data-set and its associated benchmarking toolbox [17]. This set consists of fully annotated 30 Hz videos taken from a moving vehicle in a regular traffic, and as in other recent works, we use 10 fps sampling. This increases the training data to $\simeq 2 \cdot 10^4$ detections. We generate proposals with the *SquaresChnFtrs* detector [18], and use as positive examples the proposals with 0.5 overlap threshold.

PETS 2009 S1 L2 data-set [19]. For comparison with existing methods we use the PETS 2009 data-set, which is a sequence of 795 frames recorded with seven static outdoor cameras. As noted by other authors: [12, p. 10], or [13, p. 10], we also observe inconsistencies both in terms of calibration and synchronization.

EPFL-RLC data-set [20]. This new corpus was captured with three calibrated HD cameras, with a frame rate of 60 frames per second. Currently the annotations represent a balanced set of 4088 multi-view examples. Note that a negative multi-camera example may contain a pedestrian in one of its views. Thus, for each view the negative samples contain a bit of annotated information wheter it contains a pedestrian or not. This allows the data-set to be used for monocular pedestrian training in which case its size is increased three-fold.

Full ground-truth annotations are provided for the last 300 frames of the sequence, intended to be used for testing while ensuring diversity of the appearance with respect to the

training data. As for future work, we also make available the full sequence of 8000 synchronized frames of each view.

Metrics. Apart from standard classification measures, we use the Multiple Object Detection Accuracy (MODA) metric, accounting for the normalized missed detections and false positives, as well as the Multiple Object Detection Precision (MODP) metric, which assesses the localization precision [21]. We also estimate the empirical precision and recall, calculated by $P=TP/(TP+FP)$ and $R=TP/(TP+FN)$ respectively, where: TP , FP , and FN are true positives, false positives and false negatives, respectively.

B. Monocular pedestrian detection

We conducted our experiments using the Torch framework [22], and mainly GoogLeNet [14] and AlexNet [15]. The input dimension for these models are respectively 224×224 and 227×227 . As a performance measure we use the log average miss rate (MR), as it is widely adopted.

When using GoogLeNet, batch size of 64, learning rate 0.005, SGD with momentum 0.9, as well as a proportion of positive samples per mini-batch of $r = 0.33$, we observe mean MR of 19.81 ($\sigma=0.58$) on the test data when the learning starts to saturate. When input dropout is applied the mean MR drops to 17.32 ($\sigma=0.75$), and further continues to decrease, although with slight variance. Applying it allowed for reaching MR performances which were never reached with standard training, and the observed variance suggests that combining it with more sophisticated learning rate policy can provide further performance gain. When using AlexNet, the mean MR is 24.04 ($\sigma=0.39$), and 22.43 ($\sigma=0.51$), without and with input dropout, respectively.

In summary, we consistently observed performance improvement when using: (i) GoogLeNet, (ii) square cropping versus warping the region, (iii) proportion of $r=0.33$ positive samples per mini-batch, as well as (iv) input dropout (§ III-B1). We reduced the MR to 15.61%, which although is higher than the results reported in [4], requires six-fold less computation and is simpler to implement and deploy.

C. Multi-camera people detection

In our experiments, we discretize the ground surface of the EPFL-RLC and PETS data-sets to grids of 45×55 and 140×140 positions, respectively. To train the multi-stream model on the PETS data-set we automatically extract negative examples which outnumber the positive ones provided by the annotated detections. During one epoch of training, these negatives are sampled without replacement. Both on EPFL-RLC and PETS data-sets we observe improved performances when training with forced ratio of increased negative samples per mini-batch.

Contrary to the monocular training, we observed no performance drop of training Φ if rectangular input is being used, which is why we extract features with reduced input in the width dimension by omitting 50 pixels of both sides. In our implementation, we zero-pad the input of the view for which a particular position is out of its field of view.

Unless otherwise stated, we use MLP with 3 fully connected layers, ReLU non-linearities and log-softmax. Natural consideration is the use of regularization as intuitively the weight vector of the fully connected layers would be sparse. We experimented with the p -norm regularization term: $\|w\|_p = (\sum_{i=1}^N |w_i|^p)^{1/p}$, both with $p = 1, 2$, while using SGD optimization, but did not observe a substantial improvement. In fact, by visualizing the weights when training without regularization we observe that major part of the weight vector tends towards zero, and that the training was more stable.

We provide bellow an empirical analysis of the impact on performance of some aspects of our setup.

1) Training with hard-negative auto-generated examples:

Since the EPFL-RLC data-set demonstrates more accurate joint-view calibration than PETS'09 we are able to automatically generate multi-view negative samples which we refer to as *hard negatives* and which force Φ to compare the multi-stream extracted features in a way that generalizes better.

We perform this automatically in two different ways: (i) by shifting the rectangles of a positive detection in one or two of the views, or (ii) by combining positive detections of different persons. Training with these hard negatives allows respectively for: (i) sharper detections in the ground plan, which makes the method less sensitive to the non-maxima suppression post-processing; as well as (ii) forcing the classifier Φ to learn a stronger joint appearance model. The former improve localization and the latter to decrease the false positives of the detector. Note that in the experimental results presented in the following, changing the ratio r implies changing the difficulty of the classification for this data-set and as we shall see hurts the accuracy.

From the results illustrated in Fig. 5 we observe that training with such negatives improves the most important MODA metric and reaches almost perfect precision, at the cost of marginally decreasing the MODP metric, and decreased recall due to the slightly increased number of missed detections. However, as the camera calibration of the PETS data-set demonstrates non-negligible joint-camera inconsistency of the projections, we did not experiment with the multi-view hard negatives for it, to avoid increasing the noise level of the input.

2) Monocular Vs. Multi-view classification:

We compare performance with and without using multiple views. As the performance using a MLP for Φ is potentially highly dependent on the persistence of tuning the hyper-parameters and/or the choice of optimization methods, in this section we use solely random forests for Φ .

For this purpose, we use the accuracy and the ROC while considering all possible (sub)sets of the available views, both on PETS and EPFL-RLC, as well as for different values of the proportion r of positive samples. In all experiments, we consistently observe performance improvement as more views are being added, as illustrated on Fig. 6.

3) Impact of the Selected Layer on the Accuracy:

We use GoogLeNet and compare using the first 11, 15, 20 and 23 layers. After flattening the output of each of the layers, we add fully connected layers, each of 1024, 512 and 2 units.

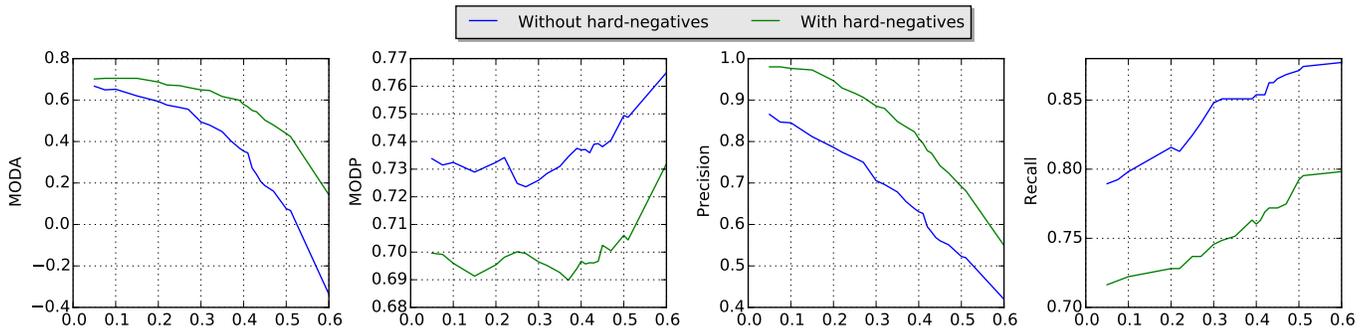


Fig. 5: Impact of training with hard-negative multi-view samples, versus the NMS threshold (x-axis) on the EPFL-RLC data-set.

We experiment with: (i) mini-batch size of: 32, 64, 128; (ii) optimisation technique: SGD w/o momentum, Adam [23], Adadelta [24] and RMSProp [25]; as well as (iii) different initial values for the learning rate when necessary. We run each experiment for 60 epochs on the EPFL-RLC data-set, and we list in Table I the best test accuracy at that point for each of the selected layers. In the majority of the experiments we observe best performances using Adadelta or ADAM, and batch size of 64.

Our experiments indicate that the depth of the CNN used as feature extractor does not impact much the performance. This is encouraging since shallower structures can be used what decreases the computation time when the model is deployed.

4) *Multi-view Information*: Herein we focus on the question if multi-view joint information exploitation allows for easier discrimination among the two classes versus monocular

classification. To this end, we analyze if features from the different views are used uniformly, or if features from one of the views dominate the decision. The most straightforward way to illustrate this is by using a single decision tree for Φ . Legitimately, we define the importance of a feature in terms of its depth in the decision tree: the higher, or the closest to the root, the more important a feature is.

We used the EPFL-RLC data-set and GoogLeNet with $d = 23$, which for the three views provides us with 3×1024 features. We visit the nodes from the root to the leaves, and count the features from each view. We find that among the 50, 100 and 150 top-nodes, the numbers of features from each view are 16/15/19, 36/32/32, and 55/51/44, respectively. This shows that the classifier exploits the views in quite a balanced manner, what motivates multi-view joint detection.

D. Comparison to existing methods on PETS

Finally, in Table II we compare our method with the current state of the art methods. We refer as *ours* when the full PETS sequence is used as testing and the model has never being trained on it; and as *ours-ft* the experiments where we divide the sequence to training for fine-tuning the model, and testing. In the latter we consider all the possible splits to train and test frames, hence we list the mean and the standard deviation of the particular performance measure estimated through these multiple folds.

As the results in Table II show, without data normalization specific to this dataset and without fine-tuning we outperform the existing methods, what indicates generalization property which is of interest. These results justify the multiple steps of our approach.

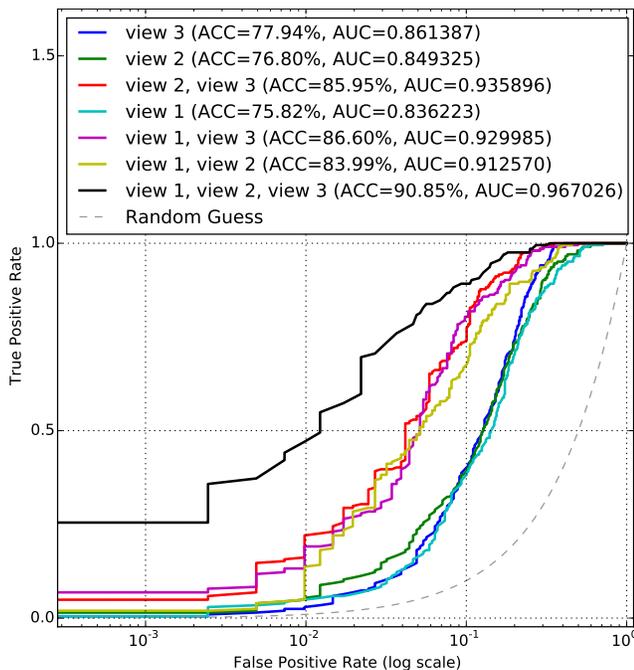


Fig. 6: Accuracy, ROC and AUC on the EPFL-RLC test data, using random forest of 100 trees, with $d = 23$, $r = 0.33$ and GoogLeNet (§ IV-C2).

TABLE I: Test accuracy on EPFL-RLC, for d layers, proportion r of positive samples per batch (see § III-B2) and with different level of difficulty (see § IV-A).

$r \backslash d$	11	15	20	23
0.50	99.75%	100.00%	99.75%	98.77%
0.33	96.24%	97.05%	96.73%	94.60%
0.25	91.78%	83.57%	83.70%	81.74%

TABLE II: Comparison on the PETS data-set (see § IV-D). The views’ enumeration omits the 2nd view.

Method	Views	MODA	MODP	Precision	Recall
[12]	1,4,5,7	0.79	0.74	0.92	0.91
[13]	1,4,5,7	0.75	0.68	0.85	0.89
Ours-ft	1	0.94 (0.03)	0.61 (0.02)	0.97 (0.02)	0.96 (0.02)
Ours	1	0.88	0.75	0.97	0.91
Ours-ft	1,4,5,7	0.90 (0.05)	0.66 (0.02)	0.94 (0.03)	0.96 (0.02)
Ours-ft	1-7	0.94 (0.02)	0.68 (0.02)	0.98 (0.01)	0.97 (0.02)

We recall the calibration joint inaccuracy of PETS which becomes more prominent as the number of used views of it increases. This homograph mapping deterioration originates in part from the presence of a slope in the scene [12]. To the best of our knowledge researchers avoid using the seven views at the same time. Instead, most often reported is the case of using the views 1, 4, 5 and 7 (or 1, 5, 6, 8 in the original views’ enumeration). It is interesting that in our experiments we do not observe performance drop when all of the seven views are used, and that in fact we observe marginal improvement. This could be due to the contextual padding that we recommended earlier. Further, the fact that the MODP metric increases, indicates that Φ is able to learn these calibration inconsistencies.

Our method outperforms the existing state of the art baselines, largely owing to its ability to leverage appearance, and to filter out non-human foreground objects that a background-subtraction procedure may trigger.

V. CONCLUSION AND FUTURE WORK

Our main contribution is proposing end-to-end deep learning method for the multi-camera people detection problem whose core is an architecture adapted to make use of existing monocular data-sets as for improved generalisation, and which jointly leverages the multi-stream deep features. As such it outperforms the existing approaches.

It was shown that: (i) multi-view classification increases the accuracy and the confidence of the classifier over the monocular case, as well as that (ii) appearance features extracted jointly from multiple views provide more easily separable embedding that in turn allows for more accurate classification than the monocular one. We also discussed various implementation insights.

Given larger scale datasets, extensions such as: (i) more explicit multi-view occlusion reasoning; and (ii) training domain-adaptation modules, can further be done. The former would allow for merging the two steps of generating detection candidates and selecting them, whereas the latter would allow for easier adaptation for the task at hand while preserving automatically learnt multi-view features.

ACKNOWLEDGMENT

This work was supported by the Swiss National Science Foundation, under the grant CRSII2-147693 ”WILDTRACK”. We also gratefully acknowledge NVIDIA’s support through their academic GPU grant program.

REFERENCES

- [1] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik, “Rich feature hierarchies for accurate object detection and semantic segmentation,” *CoRR*, vol. abs/1311.2524, 2013.
- [2] J. Hosang, M. O., R. Benenson, and B. Schiele, “Taking a deeper look at pedestrians,” in *Conference on Computer Vision and Pattern Recognition*, 2015.
- [3] P. Dollár, C. Wojek, B. Schiele, and P. Perona, “Pedestrian detection: A benchmark,” in *Conference on Computer Vision and Pattern Recognition*, 2009, pp. 304–311.
- [4] Y. Tian, P. Luo, X. Wang, and X. Tang, “Deep learning strong parts for pedestrian detection,” in *International Conference on Computer Vision*, 2015, pp. 1904–1912.
- [5] Z. Cai, M. J. Saberian, and N. Vasconcelos, “Learning complexity-aware cascades for deep pedestrian detection,” *CoRR*, vol. abs/1507.05348, 2015.
- [6] S. Ren, K. He, R. Girshick, and J. Sun, “Faster R-CNN: Towards real-time object detection with region proposal networks,” in *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [7] L. Zhang, L. Lin, X. Liang, and K. He, “Is faster r-cnn doing well for pedestrian detection?” in *European Conference on Computer Vision*, 2016.
- [8] F. Fleuret, J. Berclaz, R. Lengagne, and P. Fua, “Multi-camera people tracking with a probabilistic occupancy map,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 2, pp. 267–282, 2008.
- [9] J. Berclaz, F. Fleuret, and P. Fua, “Multiple object tracking using flow linear programming,” *Idiap, Idiap-RR Idiap-RR-10-2009*, 6 2009.
- [10] A. Alahi, L. Jacques, Y. Boursier, and P. Vandergheynst, “Sparsity driven people localization with a heterogeneous network of cameras,” *Journal of Mathematical Imaging and Vision*, vol. 41, no. 1-2, pp. 39–58, 2011.
- [11] M. Golbabaee, A. Alahi, and P. Vandergheynst, “Scoop: A real-time sparsity driven people localization algorithm,” *Journal of Mathematical Imaging and Vision*, vol. 48, no. 1, pp. 160–175, 2014.
- [12] P. Peng, Y. Tian, Y. Wang, J. Li, and T. Huang, “Robust multiple cameras pedestrian detection with multi-view bayesian network,” *Pattern Recogn.*, vol. 48, no. 5, pp. 1760–1772, May 2015.
- [13] W. Ge and R. T. Collins, “Crowd detection with a multiview sampler,” in *European Conference on Computer Vision*, 2010.
- [14] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *Conference on Computer Vision and Pattern Recognition*, 2015.
- [15] A. Krizhevsky, I. Sutskever, and G. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Neural Information Processing Systems (NIPS)*, 2012.
- [16] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *CVPR09*, 2009.
- [17] P. Dollár, “Piotr’s Computer Vision Matlab Toolbox (PMT),” <https://github.com/pdollar/toolbox>.
- [18] R. Benenson, M. Omran, J. Hosang, and B. Schiele, “Ten years of pedestrian detection, what have we learned?” in *European Conference on Computer Vision*, 2015.
- [19] J. Ferryman and A. Shahrokni, “Pets2009: Dataset and challenge,” in *2009 Twelfth IEEE International Workshop on Performance Evaluation of Tracking and Surveillance*, Dec 2009, pp. 1–6.
- [20] “EPFL-RLC data-set,” <http://cvlab.epfl.ch/data/rlc>.
- [21] R. Kasturi, D. Goldgof, P. Soundararajan, V. Manohar, J. Garofolo, R. Bowers, M. Boonstra, V. Korzhova, and J. Zhang, “Framework for performance evaluation of face, text, and vehicle detection and tracking in video: Data, metrics, and protocol,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 2, pp. 319–336, Feb 2009.
- [22] R. Collobert, K. Kavukcuoglu, and C. Farabet, “Torch7: A matlab-like environment for machine learning,” in *BigLearn, NIPS Workshop*, 2011.
- [23] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014.
- [24] M. D. Zeiler, “ADADELTA: an adaptive learning rate method,” *CoRR*, vol. abs/1212.5701, 2012.
- [25] T. Tieleman and G. Hinton, “RMSprop Gradient Optimization.” [Online]. Available: http://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf