
Multi-Task Network Pruning and Embedded Optimization for Real-time Deployment in ADAS

Flora Dellinger

Valeo, France

flora.dellinger@valeo.com

Thomas Boulay

Valeo, France

thomas.boulay@valeo.com

Diego Mendoza Barrenechea

Valeo, France

diego.mendoza@valeo.com

Said El-Hachimi

Valeo, France

said.el-hachimi@valeo.com

Isabelle Leang

Valeo, France

isabelle.leang@valeo.com

Fabian Bürger

Valeo, France

fabian.burger@valeo.com

Abstract

Camera-based Deep Learning algorithms are increasingly needed for perception in Automated Driving systems. However, constraints from the automotive industry challenge the deployment of CNNs by imposing embedded systems with limited computational resources. In this paper, we propose an approach to embed a multi-task CNN network under such conditions on a commercial prototype platform, i.e. a low power System on Chip (SoC) processing four surround-view fisheye cameras at 10 FPS.

The first focus is on designing an efficient and compact multi-task network architecture. Secondly, a pruning method is applied to compress the CNN, helping to reduce the runtime and memory usage by a factor of 2 without lowering the performances significantly. Finally, several embedded optimization techniques such as mixed-quantization format usage and efficient data transfers between different memory areas are proposed to ensure real-time execution and avoid bandwidth bottlenecks. The approach is evaluated on the hardware platform, considering embedded detection performances, runtime and memory bandwidth. Unlike most works from the literature that focus on classification task, we aim here to study the effect of pruning and quantization on a compact multi-task network with object detection, semantic segmentation and soiling detection tasks.

1 Introduction

For **Advanced Driver-Assistance Systems (ADAS)**, an accurate perception of the surrounding environment is of utmost importance. Accurate localization and proper classification of objects such as cars and pedestrians, as well as road, lane markings and sidewalks are essential for path planning and obstacle avoidance. Image-based perception is used for increasingly complex tasks due to the competitive cost of cameras and the advances in computer vision based on neural networks.

Neural networks, particularly **Convolutional Neural Networks (CNN)** offer state-of-the-art performance in a wide range of computer vision tasks, such as image classification, object detection and semantic segmentation. Therefore, they are particularly interesting for perception in ADAS. Yet,

top performing CNNs on common benchmarks such as ImageNet [9] and COCO [21] are usually computationally intensive and require a large amount of memory and power, which makes deployment on embedded systems a challenge, particularly in ADAS. Indeed, for reducing costs and power consumption, automotive manufacturers prefer cheaper low power chips with limited computational resources rather than standard GPUs. Also, ADAS are generally based on a multi-camera system whose information must be processed at a high frame rate with high performance, making the problem even more complex. Although computational cost can be decreased by reducing the CNN network size and depth, it can negatively impact the performance, which makes embedding CNNs on such devices so difficult.

Based on these specific conditions, we have designed a **compact multi-task network** that can be embedded on a low power System on Chip (SoC). Three tasks are integrated: object detection, semantic segmentation and soiling detection (see Figure 2). The soiling detection task is useful in adverse weather conditions by detecting dirt and water on the camera lens. Such network is more efficient for embedded applications compared to multiple single task networks, since the use of a single shared encoder reduces significantly the number of operations. To fit the hardware computational budget, a large effort was made on network compression and embedded optimization. A **pruning algorithm** is applied to the network to reduce the FLOPs (floating point operations) by almost a factor of 2 with limited impact on performances. Several **embedded optimizations** are performed to ensure an efficient execution of the network inside the SoC environment. This is one of the first articles introducing the effect of pruning and quantization on a compact multi-task CNN.

The rest of the paper is structured as follows. Section 2 introduces the overall perception system. Our proposals for a multi-task network and CNN compression are presented respectively in Section 3 and 4. Section 5 details our embedded optimization. Section 6 discusses the experimental setup and results. Finally, section 7 summarizes the paper and provides potential future directions.

2 Perception system overview

Our perception system performs several operations starting with sensor data rendering, CNN execution and finally CNN output interpretation before translating them into actions for our global automotive system. To reach 10 frames per second (FPS) on the entire system, a System on Chip (SoC) capable of optimizing the whole video pipeline is required. The SoCs available on the automotive market include, on the same chip, several types of computing units dedicated to the image processing and computer vision operations, in addition to the more commonly found CPU or GPU cores. Moreover, the manufacturer must compromise and set the number of units on the SoC to the minimum needed to keep the power consumption under 10W.



Figure 1: Video pipeline on the SoC

Figure 1 summarizes the video pipeline stages and corresponding units on the SoC. One or multiple Image Signal Processors (ISPs) transform raw pixels captured by four surround-view 2 MPixel fisheye cameras into high quality images. ISPs are part of the global image processing block used for demosaicing, noise reduction or lens imperfection correction on the raw signal. A cylindrical fisheye correction is performed by a dedicated Image Rendering (IR) core to reduce distortion before feeding the image to the CNN.

The MTL CNN is executed on a dedicated CNN core which performs convolutions of kernel size 5×5 on 4 inputs and 8 output channels in a single clock-cycle, reaching about 1 Tera operations per second (TOPS). Finally, the outputs of the MTL network are post-processed by CPU cores or other computer vision cores (CV cores) available on the SoC. From a global automotive system perspective, the post-processed outputs of the CNN are then fused with several other outputs from

other sensors and are finally translated into active safety actions.

Consequently, this perception system configuration imposes some constraints to the architecture of the multi-task CNN. Based on these specifications, the number of FLOPs must be limited to 3.5 GFlops to ensure a deployment on the embedded hardware at 10 FPS. Starting from these requirements, we will present our approach to design and embed such a multi-task network in the following sections.

3 Multi-task CNN design

3.1 Related work

Object Detection State-of-the-art CNNs for object detection are based on a two-stage approach, with region proposals followed by a detection module [28]. However, faster algorithms have been proposed based on a single step approach [26] [27] [22], which are more adapted to real-time applications.

Semantic Segmentation State-of-the-art approaches to semantic segmentation trained end-to-end are based on Fully Convolutional Networks [24] using transposed convolutions [35] to recover the spatial information needed for pixel-level classification [29].

Soiling detection Soiling detection [31] can be approached as a coarse semantic segmentation task, based on a 4×4 grid for localization as described in [8].

Multi-task learning For problems where multiple tasks are required, such as object detection and semantic segmentation in autonomous driving applications, using a single neural network with a common encoder and multiple heads presents many benefits. Low-level features can be useful for multiple tasks, and multi-task training can provide a regularization that improves the test set performance [30]. Moreover, a common encoder requires less memory and computational resources.

3.2 Our approach: Overview of 3-task CNN

Our 3-task CNN as shown in Figure 2 is composed of a U-Net-like encoder optimized by a NAS (Neural Architecture Search) algorithm based on Evolution Strategies [5]. Particularly the network depth, the number of filters and the position of skip connections and other layers have been optimized to boosted performance while limiting the network complexity. We used a YOLO-like decoder for object detection, a FCN-like decoder for semantic segmentation and for soiling we used the decoder described in [8]. The kernel size was set to 5×5 , which has an efficient implementation on the used embedded system. Given the computational cost of each convolutional layer, it is important to have as few layers as possible. As a consequence, initial experiments with a limited number of layers could hardly detect objects in the near-field which often cover large parts of the image. To overcome this, we determined a minimal number of convolutional layers required so that each task head output has a receptive field larger than the input image size, based on receptive field computation methods described in [2].

While compact, this network still requires too much computation for our embedded setup (6 GFLOPs while the maximal limit of our setup is 3.5). The next section will introduce the concepts of pruning and quantization for network compression, followed by our chosen pruning approach for runtime reduction.

4 Network compression

4.1 Related work

Network compression techniques [7, 12], such as quantization and pruning, have gained a lot of interest in recent years in order to be able to use CNNs on edge devices for real time applications.

Quantization Low-power hardware use fixed-point precision operations to execute the network. Therefore, a step of quantization [16] is mandatory to deploy the network trained with floating-point precision. Quantization is also a way to reduce network size, by representing weights and activations of the model with lower-precision numerical formats (commonly 8/16-bit integers instead of 32-bit floating point). The work of [17, 20, 1] proposed post-training quantization methods to efficiently

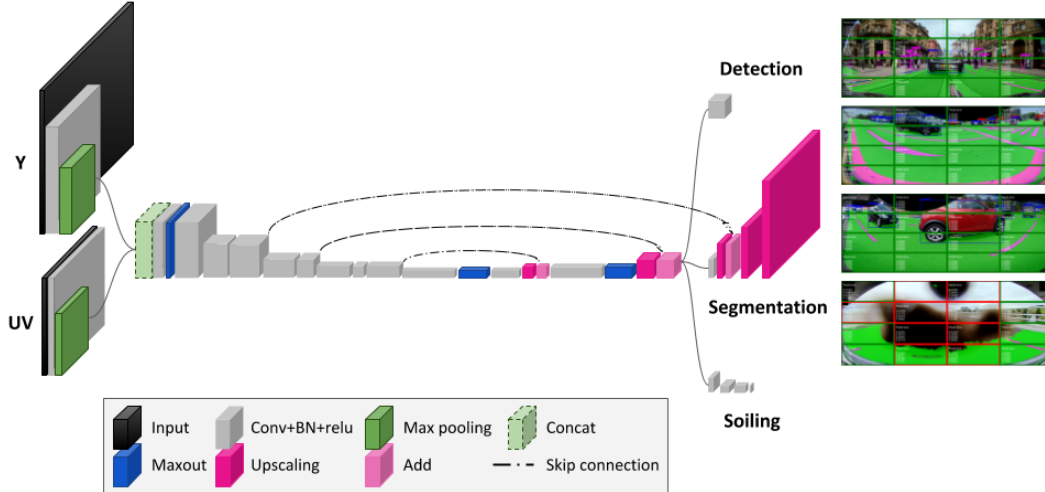


Figure 2: Multi-task visual perception network architecture with Y (gray) and UV (color) input

represent the information of the floating-point numbers during training into fixed-point integers. This technique is widely used to embed single-task CNNs for storage, bandwidth and runtime reduction. However, the impact of quantization on multi-task CNNs has rarely been studied.

Pruning is a compression technique that removes less important connections in a neural network. The idea behind is that neural networks are over-parametrized and thus some neurons don't contribute much to the final result. Starting from a large network with a good performance, the goal is to remove neurons with few contribution in order to obtain a smaller network with similar accuracy [3]. Like quantization, this is useful to reduce network storage, bandwidth and inference runtime.

Pruning has proven its effectiveness through a large number of methods in literature. While it is difficult to compare them due to a lack of benchmark datasets [3], nearly all follow the same steps introduced in [12] and stated in Figure 3. That being said, pruning methods can be differentiated based on their pruning granularity, their pruning criterion and their scheduling.

The **pruning granularity** defines the type of structure pruned in the network.

- **Weight pruning** (*unstructured pruning*) [12] consists in inducing sparsity in weights and activations of a model. Weights are transformed into sparse matrices, by setting some elements to zeros. This is effective to reduce the number of parameters and thus network storage. However, most hardware and libraries do not handle sparse convolutions, thus this won't entail a runtime reduction. For a vast majority of hardware and applications, this solution is not helpful.
- **Filter pruning** (*structured pruning*) [19], on the other hand, allows to prune entire filters or channels (see Figure 4). A direct reduction of the FLOPs is thus obtained, leading to a decrease of inference runtime. However, accuracy can strongly be reduced and fine-tuning is usually needed to recover the performance drop.

The **pruning criterion** allows to compare and choose which elements to prune. A score is assigned to each weight or filter to assess its contribution to the network. Scores are then compared locally (within a layer) or globally (across all layers), and elements with scores below a threshold are pruned. Choosing the right criterion is key in pruning success. Various methods have been investigated:

- **Data-dependent criteria** are related to some statistics on a dataset, like the average or percentage of zeros among activations [14, 25].
- **Data-independent criteria** are based on the state of the network: l_1 -norm on filters [19], Taylor expansion [25] or scaling factors from Batch Normalization [23].

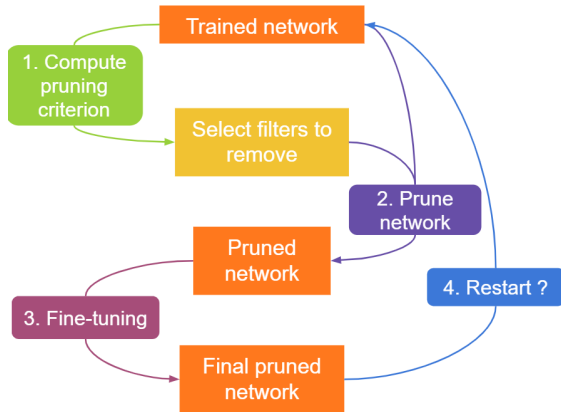


Figure 3: Steps of pruning methods.

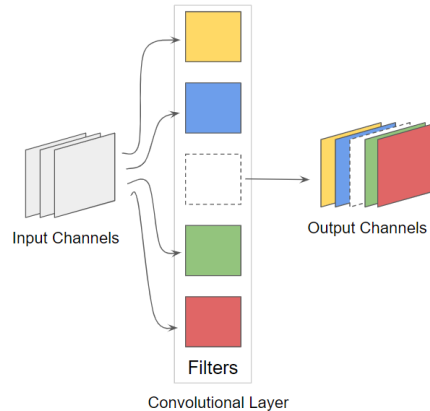


Figure 4: Example of filter pruning.

The **pruning scheduling**: Some methods conduct one-shot pruning by compressing the network in a single step [23], while others prefer to remove a percentage of the network in an iterative manner [12]. Most methods carry out some fine-tuning after each pruning step to recover performances.

As stated in [3], there is no perfect pruning method. Results will be different regarding the network size and amount of pruning. Choosing the right method will depend on the type of network, the compression goal and the hardware implementation. We can also note that most methods have only been applied to image classification, which is not representative of the tasks used for CNN. There are only few publications on the use of pruning on other tasks, like semantic segmentation in [6] or visual attribute prediction in [13]. It is thus necessary to study the impact of pruning on compact networks with complex and multiple visual tasks.

4.2 Our approach: Pruning of 3-task CNN

We have chosen a filter pruning method as our target hardware is optimized for dense convolutions. Two pruning criteria common in the literature are compared: FilterNorm based on the $l1$ -norm of weights [19] and BatchNorm based on the γ scaling factor [23]. Since we have observed that some layers (especially deeper layers) are more pruned than others, filter scores are then compared globally.

The steps of the pruning method are the same as in Figure 3: starting from a pre-trained model, the model is iteratively pruned and fine-tuned until the target runtime is achieved. When the network has reached the desired compression, a longer fine-tuning is carried out to aim for best performances. A trade-off is needed between the amount of network pruned at each step and the fine-tuning length. Experiments have shown that longer fine-tuning with a large compression rate should be preferred over a small compression rate with short fine-tuning. An iterative pruning is necessary here since a high amount of the network needs to be trimmed. A one-shot pruning is too dramatic for accuracy in our case, however, it can be applicable for small amounts of pruning.

After compression, the 3-task CNN can be executed with a limited number of FLOPs. However, some embedded optimizations are still required to fit the target in our case.

5 Embedded optimization

An embedded CNN deployment and optimization for a real product always start with a clear definition of the target to achieve. It must be delineated by taking into account the full system requirements (CNN is usually not the only algorithm running on the SoC). Generally, the most important metrics to track module by module on the hardware are the runtime, the DDR bandwidth and the DDR footprint, DDR (Double Data Rate) being commonly the memory resource shared by all algorithms. For our application, the budget allocated to the CNN is defined based on these 3 metrics. First, the CNN must be executed to reach 10 FPS (a frame is defined as a batch of 4 images coming from our surround-view fisheye cameras). Secondly, taking into consideration the full system bandwidth and other algorithms running in parallel on the CNN, the bandwidth allocated to our CNN module is

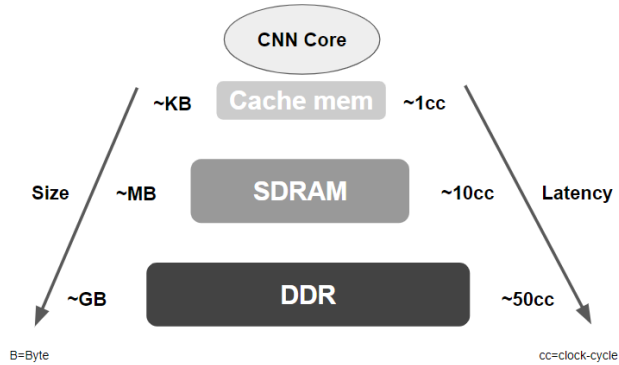


Figure 5: Memory areas for the CNN core.

fixed to 1 Giga Bytes per second (GBps). In concrete terms, it means that the throughput measured on the SoC for the CNN module should always be less than 1GBps. Finally, to work on the SoC, the CNN module needs an amount of allocated memory on the DDR called DDR memory footprint or DDR peak usage. For our application, the DDR memory footprint for the CNN part (4 cameras) is limited to 14MB.

5.1 Network optimization

Network optimization for an embedded deployment often starts as soon as the network architecture is being developed. The paper [4] describes several important network architecture optimizations like input/convolution layer shape optimization or layer sequence optimization. The optimizations described in this paper have all been applied to our network enabling to reduce drastically the runtime and the throughput of the network. For example, a YUV input layer has been used instead of the common RGB input layer to reduce by a factor of two the memory transfers needed to process the network's input data.

Once the network architecture is frozen, the filter pruning technique enables a very considerable reduction of the runtime and the bandwidth. A convolutional layer calculation is subdivided into a series of small calculations on the SoC. The spatial dimension is processed by sliding a 5×5 window on the feature maps and in the depth dimension only a small group of channels is handled in parallel. One CNN core execution convolves 4 input channels by 8 filters to produce 8 output channels. The layer computation is completed when all the channels (depth dimension) and pixels (spatial dimension) were handled. Hence by definition, when the number of the input/output channels is reduced, the number of CNN core runs and the amount of data to transfer to the CNN core are reduced. Consequently, the runtime and bandwidth are cut down. Furthermore, pruning in early layers of the encoder leads to the strongest reductions since these layers handle high resolution feature maps that require a lot of CNN core executions and data transfers. Table 2 shows the achieved optimization results.

5.2 Efficient data transfers

A well-known challenge to run optimally an embedded CNN is to efficiently arrange the data for triggering the computation on the core as fast as possible. When the CNN core ends its current calculation, the goal is to reduce as much as possible the waiting time or latency before launching the next calculation. On the SoC, several memory areas are coexisting. The closer the memory area is to the computation core, the lower the latency is, but the lower the size is, too (see Figure 5). The DDR memory area is usually the biggest area on most of the SoCs. Consequently, it is also a memory area shared by multiple algorithms which can run simultaneously. Hence, our optimization for this application is mainly focused on reducing the DDR memory traffic by cleverly arranging the data along the CNN layer executions. The main idea behind this optimization is to avoid as much as possible to transfer the CNN core output data into DDR. This principle can be applied at two levels: inside a layer when the intermediate outputs have to be stored to be accumulated with the next convolution channel outputs and across the layers when the final result of a layer is saved to be reuse as input of the next layer. We worked on these two levels by building chains of CNN

core execution nodes. These chains have been built in order to keep the data inside the SDRAM memory. So, we alternate two types of execution scheme according to the remaining memory space in SDRAM. A common scheme of tile execution (horizontal/spatial execution) when SDRAM cannot contain the intermediate results and a vertical scheme where a tile of data is processed through all layers in the chain before moving to the next spatial tile and saving in DDR memory. The effects of these optimizations are presented in Table 2.

5.3 Mixed-quantization optimization

The CNN core used for our application is hard-wired for executing 16-bit operations. Consequently, the opportunity to reduce the bit-width is not offered by the hardware. So, we cannot take advantages of approaches like [32, 33] because the operations supported by our hardware accelerator are only 16-bit operations. However, outside the CNN core, a reduction of the bit-width can be considered for reducing the amount of transferred data, especially from and to the DDR memory. Therefore, we propose a 8-bit/16-bit mixed-quantization scheme for our network to reach the embedded target in regards of DDR bandwidth and DDR footprint. The principle is to identify the layers with a high data consumption from and to the DDR and quantify the input and output feature maps of these layers in 8-bit. Following this scheme, 8-bit instead of 16-bit numbers are transported around the memory areas until they are used by the CNN core which can cast the numbers into the hard-wired bit-width of 16-bit. The impact of this optimization is described in Table 2 regarding the embedded metrics and in Table 3 regarding the algorithm performances.

6 Experiments and results

6.1 Training and pruning protocol

We run our experiments on the Woodscape autonomous driving dataset [34] using a split of 80/10/10 for train/val/test and a reduced number of classes: road, lane marking, curb, vehicle, two wheeler and person. The joint training policy consists of minimizing a multi-task loss which is an arithmetic weighted sum of individual task losses $L = \sum_i w_i L_i$, i representing a task. Task balancing is performed by a static weighting whose values have been optimized by automatic tuning described in [18]. Each batch includes images with annotations for the 3 tasks. The loss is backpropagated only through the relevant task heads and to the common encoder. The network is trained end-to-end using the Adam optimizer [15]. Our data augmentation policy includes modifications in brightness, contrast, saturation, hue, noise, scaling, spatial translation and horizontal flip. Our main metrics are mean average precision (mAP) for object detection, mean intersection over union (mIoU) for semantic segmentation and a tile-level mean F1-score for soiling detection. We apply a L2 regularization of 10^{-4} and train the network for 64 epochs. To assess the impact of using a multi-task network, we have trained three independent single-task networks under the same conditions.

After training, the network is pruned in an iterative manner. At each step, an amount of the network is trimmed to reduce the number of GFLOPs and a fine-tuning of 32 epochs is carried out on the intermediate network. The pruning stops when the number of GFLOPs is below the embedded target (here 3.5 GFLOPs). The final model is fine-tuned for 64 epochs. Two pruning experiments are conducted with the two studied pruning criteria.

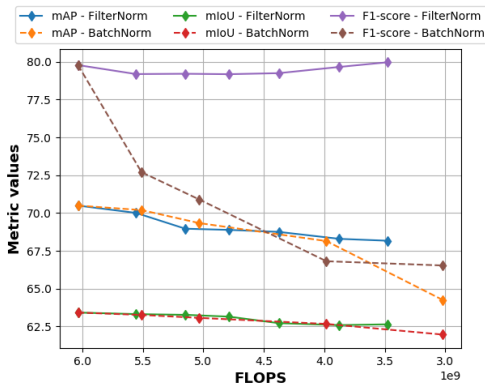
6.2 Multi-task versus single-task networks

Table 1 summarizes metric values for single-task and multi-task networks after training. Metrics obtained by single-task networks are all higher than corresponding ones from the multi-task network. However, the gap is limited for object detection (+0.26) and semantic segmentation (+1.78). Only soiling detection benefits strongly of this single training (+14.28), even though performances are still competitive with joint training (79.78 for F1-score). Moreover, using single-task networks would multiply the number of FLOPS by three in comparison to the multi-task network (17.8 instead of 6.0 GLOPS).

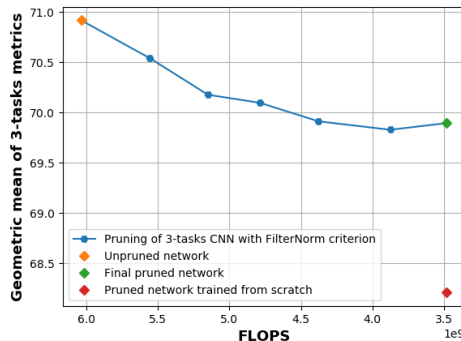
Choosing a multi-task network helps us then to reach competitive performances with limited computation time.

Network version	Detection mAP	Segmentation mIoU	Soiling F1-score	Parameters	GFLOPs
Object detection	70.75	-	-	2,490,700	5.9
Semantic segmentation	-	65.20	-	2,511,320	6.0
Soiling detection	-	-	94.06	2,515,028	5.9
Multi-task unpruned	70.49	63.42	79.78	2,540,208	6.0
Multi-task pruned with FilterNorm	68.17	62.64	79.97	741,956	3.5
Multi-task pruned with FilterNorm, re-trained from scratch	65.15	61.74	78.90	741,956	3.5

Table 1: Performance comparison between single-task and multi-task networks, before and after pruning. The evaluation set contains 5K (detection/segmentation) and 18K images (soiling).



(a) Evolution of metrics over multi-task pruning for two criteria : FilterNorm (l_1 -norm) and BatchNorm (γ scaling factor).



(b) Interest of pruning and comparison with the final pruned network trained from scratch.

Figure 6: Influence of pruning on multi-task networks. For both figures, the curves display values from the unpruned network (far left) to the final pruned network (far right). In between, there are intermediate networks.

6.3 Network compression experiments

Figure 6 displays the multi-task network behavior over pruning. The evolution of metrics for the 3 tasks are depicted in Figure 6a for the two pruning experiments. While object detection and semantic segmentation metrics behave in a similar way for both criteria (at least until 4.0 GFLOPs), pruning with BatchNorm criterion leads to a strong degradation of the soiling metric. Thus, the FilterNorm criterion seems more suited to this network and later experiments and discussions will be conducted with this criterion.

Figure 6b provides an overview of the network accuracy over FilterNorm pruning via the geometric mean of the 3 metrics, while Table 1 gathers detailed metric values. Iterative pruning with FilterNorm allows to strongly reduce the number of FLOPs and parameters, by respectively 42% and 71%, with minimum effect on performances. The object detection task (-2.3% on mAP) is more impacted than semantic segmentation (-0.78% on mIoU), while the soiling task is even slightly improved ($+0.19\%$ on F1-score).

To assess the utility of pruning compared to training a more compact network directly, the final pruned network has been retrained from scratch on 64 epochs. For all three tasks, the values are below the ones obtained with pruning (-3.02% , -0.90% and -1.07% , respectively on the mAP, mIoU, F1-score metrics). Consequently, pruning is useful and efficient to compress the multi-task CNN. While a small degradation of performances may be observed, the impact is less important on soiling and segmentation tasks than on object detection.

	Runtime (FPS)	Bandwidth (GBps)	Footprint (MB)
	Goal		
	10	1	14
Network version	Measured		
Unpruned network	8	1.6	30
Pruned network	14	1.2	15
Data transfer optimized network	10	1	14.4
Mixed-precision network	10.5	0.96	13.2

Table 2: Embedded optimization results.

Network version	Detection mAP	Segmentation mIoU	Soiling F1-score	Parameters	GFLOPs
float32 precision	68.17	62.64	79.97	741,956	3.5
16-bit precision	67.85	62.64	78.68	741,956	N.A
8-bit/16-bit mixed-precision	67.76	62.60	77.33	741,956	N.A

Table 3: Performance comparison of different embedded versions of our pruned multi-task network. The evaluation set contains 5K (detection/segmentation) and 18K images (soiling).

6.4 Embedded optimization results

Table 2 presents the embedded optimization impact on the embedded metrics. The impact of pruning is very significant on the 3 metrics. The pruned network offers a runtime and memory footprint reduction by almost a factor of 2 compared to the unpruned network. The DDR bandwidth and footprint are respectively reduced by 25% and 4% by the data transfer optimized network (see Section 5.2). However, a side effect of using the low-latency memory is that smaller tiles (block of pixels loaded in memory) have to be defined to process the network in a smaller amount of memory. It leads to more CNN-core runs and consequently to lower runtime performances. Finally, with the mixed-precision network DDR bandwidth and footprint are reduced and overcome respectively the 1GBps and 14MB targets. The runtime also takes advantage of this optimization with a gain of 5% which offers some room on the 10 FPS goal. In terms of performances (see Table 3), the mixed-precision network very slightly penalizes the 3 task performances compared to the float32 network (-0.6% , -0.06% and -3% respectively on the mAP, mIoU, F1-score metrics). This small degradation is a good trade-off in regards of the gain observed on the embedded metrics.

7 Conclusion

In this paper, we presented a series of algorithm and embedded optimization techniques necessary for deploying a CNN capable of processing three perception tasks in parallel with reasonable performance on a low power SoC at 10 FPS for four surround-view cameras in a vehicle. First of all, we designed an initial multi-task network with a NAS-based architecture considering a compromise between the hardware constraints and perception performances. Then, we proposed an iterative filter pruning method to greatly compress the network reducing the runtime execution by a factor of 2 without significant performance loss. Finally, we proposed several embedded optimization techniques to achieve real-time execution on the SoC such as efficient data transfers and mixed 8/16-bit quantization for reducing DDR bandwidth and footprint. This is one of the first papers presenting pruning and quantization results on a compact multi-task network for automotive applications.

For future work, we would like to explore other pruning methods to improve the network performance, such as the lottery ticket hypothesis [10] or the automated gradual pruning [36]. We also would like to study the impact of such compression on network robustness [11] and thus safety for automotive applications. To further optimize the network quantization step, the study of other calibration methods would be useful for minimizing the performance degradation even more.

References

- [1] S. Anwar, K. Hwang, and W. Sung. Fixed point optimization of deep convolutional neural networks for object recognition. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1131–1135, 2015.
- [2] André Araujo, Wade Norris, and Jack Sim. Computing receptive fields of convolutional neural networks. *Distill*, 2019. <https://distill.pub/2019/computing-receptive-fields>.
- [3] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Gutttag. What is the state of neural network pruning? *arXiv preprint arXiv:2003.03033*, 2020.
- [4] Thomas Boulay, Said El-Hachimi, Mani Kumar Suriseti, Pullarao Maddu, and Saranya Kandan. Yuvmulti-net: Real-time yuv multi-task cnn for autonomous driving, 2019.
- [5] Fabian Bürger and Josef Pauli. Understanding the interplay of simultaneous model selection and representation optimization for classification tasks. In *ICPRAM*, pages 283–290, 2016.
- [6] Xinghao Chen, Yunhe Wang, Yiman Zhang, Peng Du, Chunjing Xu, and Chang Xu. Multi-task pruning for semantic segmentation networks. *arXiv preprint arXiv:2007.08386*, 2020.
- [7] William Dally. High-performance hardware for machine learning. Tutorial, NIPS 2015.
- [8] Arindam Das, Pavel Krizek, Ganesh Sistu, Fabian Burger, Sankaralingam Madasamy, Michal Uricar, Varun Ravi Kumar, and Senthil Yogamani. Tiledsoilingnet: Tile-level soiling detection on automotive surround-view cameras using coverage metric. *arXiv preprint arXiv:2007.00801*, 2020.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [10] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Training pruned neural networks. *CoRR*, abs/1803.03635, 2018.
- [11] Shupeng Gui, Haotao Wang, Chen Yu, Haichuan Yang, Zhangyang Wang, and Ji Liu. Adversarially trained model compression: When robustness meets efficiency. *CoRR*, abs/1902.03538, 2019.
- [12] Song Han, Huizi Mao, and William Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *ICLR 2016*, 10 2016.
- [13] Xiaoxi He, Dawei Gao, Zimu Zhou, Yongxin Tong, and Lothar Thiele. Disentangling redundancy for multi-task pruning. *CoRR*, abs/1905.09676, 2019.
- [14] Hengyuan Hu, Rui Peng, Yu-Wing Tai, and Chi-Keung Tang. Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. *CoRR*, abs/1607.03250, 2016.
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *CoRR*, abs/1806.08342, 2018.
- [17] Liangzhen Lai, Naveen Suda, and Vikas Chandra. Deep convolutional neural network inference with floating-point weights and fixed-point activations, 2017.
- [18] Isabelle Leang, Ganesh Sistu, Fabian Burger, Andrei Bursuc, and Senthil Yogamani. Dynamic task weighting methods for multi-task networks in autonomous driving systems. *arXiv preprint arXiv:2001.02223*, 2020.
- [19] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *CoRR*, abs/1608.08710, 2016.
- [20] Darryl D. Lin, Sachin S. Talathi, and V. Sreekanth Annapureddy. Fixed point quantization of deep convolutional networks, 2016.
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

- [22] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [23] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. *CoRR*, abs/1708.06519, 2017.
- [24] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [25] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient transfer learning. *CoRR*, abs/1611.06440, 2016.
- [26] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [27] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [30] Marvin Teichmann, Michael Weber, Marius Zoellner, Roberto Cipolla, and Raquel Urtasun. Multinet: Real-time joint semantic reasoning for autonomous driving. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1013–1020. IEEE, 2018.
- [31] Michal Uříčář, Pavel Křížek, Ganesh Sistu, and Senthil Yogamani. Soilingnet: Soiling detection on automotive surround-view cameras. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 67–72. IEEE, 2019.
- [32] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision, 2019.
- [33] Linjie Yang and Qing Jin. Fracbits: Mixed precision quantization via fractional bit-widths, 2020.
- [34] Senthil Yogamani, Ciarán Hughes, Jonathan Horgan, Ganesh Sistu, Pádraig Varley, Derek O’Dea, Michal Uricár, Stefan Milz, Martin Simon, Karl Amende, et al. Woodscape: A multi-task, multi-camera fisheye dataset for autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9308–9318, 2019.
- [35] Matthew D Zeiler, Dilip Krishnan, Graham W Taylor, and Rob Fergus. Deconvolutional networks. In *2010 IEEE Computer Society Conference on computer vision and pattern recognition*, pages 2528–2535. IEEE, 2010.
- [36] Michael Zhu and Suyog Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression, 2017.