A Little Bit Attention Is All You Need for Person Re-Identification

Markus Eisenbach, Jannik Lübberstedt, Dustin Aganian, and Horst-Michael Gross

Abstract-Person re-identification plays a key role in applications where a mobile robot needs to track its users over a long period of time, even if they are partially unobserved for some time, in order to follow them or be available on demand. In this context, deep-learning-based real-time feature extraction on a mobile robot is often performed on specialpurpose devices whose computational resources are shared for multiple tasks. Therefore, the inference speed has to be taken into account. In contrast, person re-identification is often improved by architectural changes that come at the cost of significantly slowing down inference. Attention blocks are one such example. We will show that some well-performing attention blocks used in the state of the art are subject to inference costs that are far too high to justify their use for mobile robotic applications. As a consequence, we propose an attention block that only slightly affects the inference speed while keeping up with much deeper networks or more complex attention blocks in terms of re-identification accuracy. We perform extensive neural architecture search to derive rules at which locations this attention block should be integrated into the architecture in order to achieve the best trade-off between speed and accuracy. Finally, we confirm that the best performing configuration on a re-identification benchmark also performs well on an indoor robotic dataset.

I. INTRODUCTION

In recent years, mobile robots that follow a user have become increasingly important, especially in the field of clinical rehabilitation. One example is a robot coach for walking and orientation training of stroke patients [1], where the robot had to accompany the patients during their walking exercises that they had to perform on their own to improve their mobility as well as their orientation skills. This way, the robot addressed the patients' insecurity and anxiety of not being able to perform the exercises or not being able to find the way back to their apartment, which are possible reasons for neglecting self-training. Another example is a robotic assistant that coached patients during their walking exercises, which were taught to them by physiotherapists after hip endoprosthetics surgery [2], [3]. The robot was intended as a kind of physiotherapist replacement to provide immediate feedback to the patient regarding any deviations from the expected physiological gait pattern.

In both applications, it is mandatory for the robot to permanently be available to the users and to closely follow or guide them when desired. To achieve this, the mobile robot has to track its users for a long time, even if they are occasionally out of sight for some time. Therefore, person re-identification (re-id) plays a key role in such scenarios [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16] either to provide robust features that can be used internally



All authors are with Neuroinformatics and Cognitive Robotics Lab, TU Ilmenau, 98693 Ilmenau, Germany markus.eisenbach@tu-ilmenau.de



Fig. 1. Re-identification performance on the Market-1501 benchmark dataset [4] versus inference speed for different attention blocks and insertion positions in the ResNet-50 (see Fig. 3). You only need to add a single block of the proposed C-NL attention to ResNet-50 to outperform ResNet-101, and this only slightly affects the inference speed.

in a tracker for track continuation in case of temporal total occlusions [17] or to continue a track by re-identifying a person after a track is lost in a separate module [18].

In recent years, ConvNet-based features have proven to be very reliable for non-biometric re-identification after the user has been unobserved for some time. Therefore, these features can be used as visual cues in a probabilistic multi-person tracking framework [17]. A person detector extracts cropped RGB images for all persons in the observed scene. Next, single-shot re-identification features are extracted for these individuals. The tracker then handles the temporal integration and fusion with other tracking inputs.

Fortunately, real-time deep-learning-based feature extraction on a mobile robot is no longer a problem due to the availability of specialized GPU devices such as the NVIDIA Jetson series. However, inference speed still needs to be considered as these devices are typically shared for several neural-network-based classification or regression tasks. This is in contrast to the current trend in person re-identification, where performance is improved at all costs, neglecting the significant slowdown during inference. An example of this is the heavy use of attention blocks, especially non-local attention, which significantly slows down inference speed, as we will show in our experiments (see Sec. IV-B). We will demonstrate that we can improve the re-id performance with attention blocks without any significant slowdown in inference speed, if we find the right kind of attention blocks and a minimum set of appropriate positions in a ResNet-50 architecture. Therefore, our contributions are as follows:

- 1) We investigate the tradeoff between inference speed and re-id performance for different attention blocks, which is neglected in the current state of the art.
- 2) We propose the new attention block C-NL that is faster than current attention blocks and performs better in a regime with few attention

blocks.

- 3) We perform extensive neural architecture search (NAS) to derive a set of rules for where and what type of attention should be integrated into a ResNet-50 to significantly improve re-id performance while only slightly affecting inference speed.
- 4) We confirm that by integrating C-NL attention blocks into the ResNet-50 as specified in the derived set of rules, the re-id performance is also improved on a robotic dataset.

II. RELATED WORK

Nowadays, attention is widely used in computer vision. To improve ConvNet-based single-shot person re-identification (re-id), mainly, three types of attention blocks are available: channel-wise attention [19], spatial attention [20], and nonlocal attention [21], with self-attention of the transformer architecture [22] being a special form of the latter.

Channel-wise attention is based on the principles of the squeeze-and-excitation block [19] (Fig. 2(a)). First, the spatial resolution is reduced to 1×1 by global average pooling. Then, the inter-channel correlation is modeled using a bottleneck in which the number of channels is reduced by r. Finally, the output of the attention block, which is typically used as a per-channel weight, is normalized by a sigmoid activation. Channel-wise attention based on these principles is widely applied for re-id [23], [24], [25], [26], [27], [28]. Fully attention [29] is a variant of channel-wise attention for re-id that deliberately omits the global average pooling. In [30] fully attention was extended to higher-order statistics in the context of re-id. We will show that channel-wise attention is cheap to compute, and should therefore be considered for a mobile robotic application where inference speed plays a key role.

Spatial attention is based on the network proposed in [20]. First, the input volume is reduced to one channel while keeping the spatial resolution. Then, the spatial resolution is reduced and subsequently restored to model the spatial correlation. Finally, a sigmoid activation normalizes the output of the attention block, which is typically used as spatial weighting. Spatial attention is also widely applied for re-id [23], [24], [25], [26], [27], [31], [32]. A variant that deviates from the design principles described above is the use of a foreground mask as spatial attention for re-id [33], [34]. This requires an additional network for extracting a foreground mask, which comes at an additional inference cost. Spatial attention often does not perform that well on re-id benchmarks unless it is used in combination with other attention blocks. Therefore, many re-id approaches combine spatial attention with channel-wise attention. Channel-wise and spatial attention are computed in two branches that are either multiplied element-wise [23], [24], [25] or concatenated at the channel dimension [26], [27].

Non-local attention aims at modelling global inter-pixel correlation. Most often, it follows the design principles of the attention-module in the transformer architecture [22] (Fig. 2(c)). In a query and a key branch similarities of features are computed at different spatial positions which are then used to weight the features calculated in the value branch. For re-id, non-local attention has been applied in this pure form [35], [36], with second-order statistics [37], and

as self-attention [38]. Non-local co-attention is also used to match person images of the probe and gallery [39]. The latter approach is therefore not relevant for our application of re-id features as tracking input, since we aim for extracting features for single person images. Non-local attention achieves the best results in re-id benchmarks, but comes with high inference costs, as we will show in our experiments.

Besides these three main forms of attention blocks that are applicable for single-shot re-id on RGB images, there are also other forms of attention applied for re-id. In [40], [26], [41] reinforcement learning is used to adapt attention weights. In contrast, our focus is on learning the attention weights by error backpropagation. Selecting proper semantic and soft-biometric attributes for re-id by attention has been addressed in [41], [42], [43]. To be able to utilize attributes would require an additional network for attribute extraction. This is in contrast to our goal of affecting the inference speed only little. Qian et al. [44] use a multi-branch multiscale architecture for person re-identification and train an attention block to select the best scale for an input. Multiscale approaches require significantly more operations, and therefore slow down inference, which is contrary to our goal. Attention has also been applied to video-based reid on streams of images [45], [46], [47], [48], [49], [50], [51], [52], [53], [54]. These approaches primarily focus on temporal attention. In addition, attention has been applied to re-id with other input modalities, such as infrared [55] or depth images [56]. Both temporal attention and customized attention for other modalities are not the focus of the work presented in this paper.

Effect of integrating attention blocks on inference speed: The drawback of adding attention blocks, namely slowing down inference, is completely neglected in the current state of the art. However, fast inference is a prerequisite for realtime mobile robot applications. Therefore, in this paper we address this issue. We show that it is possible to integrate attention blocks at only few positions in a ResNet-50 to achieve similar performance to much deeper networks or computationally demanding attention blocks, while only marginally affecting the inference speed, if the correct positions can be identified. Therefore, we perform extensive NAS to derive a set of rules to guide the integration.

Backbone: We decided in favor of ResNet-50 as backbone for our experiments, since it is the de-facto standard for evaluating on person re-identification benchmarks. Therefore, our results are easily comparable to the state of the art. However, since we explore design spaces during NAS, the gained knowledge about integration of attention blocks is also transferable to other similar architectures [57].

III. A FAST AND POWERFUL ATTENTION BLOCK After providing an overview of the state of the art in person re-identification (re-id) with architectures that incorporate attention blocks, we will now describe the attention blocks that we considered in our NAS. Then, we derive a novel attention block that attempts to combine the advantages of these attention blocks in terms of inference speed and reidentification capabilities, as described in the following.



Fig. 2. Layout of attention blocks: Channel-wise attention is realized by the (a) squeeze-and-excitation block [19] (SE) and the (b) harmonious attention channel-wise block [23] (HAC). The prototype design of non-local attention [21] (NL) is shown in (c). Our proposed block (C-NL) that is derived from these channel-wise and non-local attention blocks is shown in (d).

A. Attention Blocks for Re-Identification

Fig. 2 shows the overall layout of channel-wise and nonlocal attention blocks that achieve the best re-id performance in benchmarks. In [23] (harmonious attention channel-wise, HAC), it has been shown that the performance (with channelwise attention) improves most when attention blocks are inserted between residual blocks, as it has been shown first in [21] for another type of attention applied to computer vision applications. We visualize the modified position relative to the residuum in Fig. 2(b) (see Fig. 2(a) for the original position as proposed in the squeeze-and-excitation block [19]). We also examined local attention multi-granularity network (LAG) [28], but its attention block clearly underachieved in comparison to HAC. Therefore, these results are not shown in the following experiments for reasons of clarity.

Also non-local attention blocks are placed between residual blocks (see Fig. 2(c)), which has been shown to be superior in [21]. We consider three non-local attention blocks, namely attention generalized mean pooling with weighted triplet loss (AGW) [21] (applied for re-id in [35]), relationaware global attention (RAGA) [36], and attentive but diverse (ABD) [38]. All non-local attention blocks follow the same layout as shown in Fig. 2(c) with three branches, called query, key, and value in the transformer architecture [22]. In all these three attention blocks, f computes the matrix product of queries q and keys k, with only the normalization of the result being different. RAGA also comes with a further small difference, which translates into additional computational costs.

B. Deriving a Novel Attention Block

Both channel-wise and non-local attention blocks provide advantages for re-identification, which we attempt to combine in our novel attention block as described below. In our experiments, we show that non-local attention blocks slow down inference speed too much to be worth considering for a mobile robotic application. Therefore, our novel block should be designed like the HAC attention block [23]. Its low inference costs are a result of the early reduction of the spatial resolution by global average pooling. The benefit of non-local attention blocks are, however, the three-branch layout to appropriately weight the features in the value branch to estimate global correlations. Therefore, our block uses an early global average pooling to reduce computational costs and then follows the three-branch design of the nonlocal attention blocks to model the inter-channel correlation in a non-local way. Fig. 2(d) shows the derived design of our channel-wise non-local attention block (C-NL). The three-branch layout comes with minor computational costs in this setting as a result of the matrices of reduced size that need to be multiplied (see matrix dimensions shown in gray in Fig. 2(c, d)). Therefore, this design leads to low computational costs and achieves a high re-id performance, as we show in the following experiments.

IV. HOW MUCH ATTENTION DO WE NEED FOR PERSON RE-IDENTIFICATION?

In what follows, we will first address the knowledge gap regarding the extent to which attention blocks slow down inference speed. Then, we will analyze how many attention blocks we really need to significantly improve person reidentification (re-id), by performing extensive NAS over the design space of attention blocks, positions of integration, and hyperparameters. We take inspiration by the Designing Network Design Spaces NAS approach [57] that showed how the design space can gradually be reduced and, as a consequence, a set of design rules can be derived. We primarily reduce the design space regarding positions of integration to derive a set of rules of where attention blocks should be added to get a large improvement with low computational cost.



Fig. 3. ResNet-50 with positions where attention blocks can be inserted marked as diamonds. Applying attention blocks directly on the input is somewhat questionable and has shown to not improve the performance in our experiments. All other 17 positions between the initial stem and the final GAP+FC should be considered. Residual blocks with an 1×1 convolution in the skip connection (gray) mark the beginning of each of the four stages of the ResNet-50.

A. Experimental Setup

As basis for our experiments, we use the strong baseline and training setup of [58] consisting of a modified ResNet-50 [59], which we also deploy in real-time on our mobile robots, in combination with common augmentation and training strategies. For benchmarking, we decided in favor of one of the most popular datasets, namely Market-1501 [4], and used the evaluation protocol of [35]. We could also benchmark on DukeMTMC-reID [60] or CUHK03-NP [61], but instead we decided in favor of a robotic dataset, for which we report results in Sec. V, to show generalization ability.

Note, that better results can be achieved on Market-1501 with more elaborate baselines. But these baselines come with architectural modifications, which result in higher computational costs. This is contrary to our goal of fast inference on a robot. Also by integrating several time frames, like in [62], or other costly techniques, better benchmark results could be achieved, but this is again contrary to our goal of deploying these networks in real-time on a mobile robot.

During NAS, all hyper-parameters are tuned, such that the reported results are trustworthy and not a fragment of the choice of hyper-parameters.

B. Analysis of Inference Speed of Attention Blocks for Person *Re-Identification*

First, we measure the inference speed of attention blocks that have been proposed or applied for re-id as described in Sec. III¹. We add the different attention blocks at single positions in the ResNet-50. The diamonds in Fig. 3 show the locations where attention blocks can be integrated. We measure the inference speed on a Jetson AGX Xavier (Jetpack 4.6, TensorRT 8.0.1.6, Cuda 10.2, 16 bit floating point precision), which is a special device typically used for deeplearning computations on a mobile robot. Since for each image we need to extract features for all persons in the camera, we chose a batch size of 16 for the measurement of inference speed. With this batch size, we would be able to extract features for 16 persons in the field of view simultaneously, which we consider as a good worst-case estimate for crowded scenes, like in hospital hallways. We report the number of batches that can be processed per second during inference. The reported results are the average over 10000 processed batches of the Market-1501 benchmark



Fig. 4. Inference speed on a Jetson AGX Xavier for integrating single attention blocks at different positions measured by number of batches processed per second. The input tensor size for each attention block at the respective positions is listed at the top of the figure.

dataset [4] with cropped and scaled RGB person images of size 256×128 . However, the inference speed does not depend on the dataset, but only on the input image size, which is identical in all our experiments in this paper.

Fig. 4 shows the inference speed for the attention blocks as a function of the single position where the individual attention block is integrated. Note that the inference speed depends only on the size of the input tensor for the attention block (listed at the top of the figure) and is thus identical for several positions. As reference, we show the inference speed of the ResNet-50 without any attention block and of the ResNet-101 (without attention) that is twice as deep. We can see that channel-wise attention blocks shown in blue and green (HAC [23], C-NL) are faster to compute than nonlocal attention blocks shown in orange to red (AGW [35], RAGA [36], ABD [38]). A ResNet-50 with a single nonlocal attention block at an early position (1-4) is significantly slower than a ResNet-101. Thus, these configurations should not be considered for robotics applications. Furthermore, at late positions (9-17) only AGW [35] can cope with the inference speed of channel-wise attention blocks to some degree. Also at middle positions (5-8) channel-wise attention is faster to compute than non-local attention. Comparing the channel-wise attention blocks, for most positions (1-14) the proposed C-NL block achieves a higher inference speed than HAC [23]. To conclude, considering only inference speed so far, channel-wise attention is preferable over non-local attention in robotic applications.

C. How Many Attention Blocks Do We Really Need?

Next, we analyze, how many attention blocks we need to significantly improve the re-id performance. Therefore, inspired by [57], we performed NAS by progressively reducing the design space over attention blocks, positions of integra-

 $^{^1} For$ implementation of the attention blocks, where available, we used publicly available code provided by the authors: AGW, evaluation code – <code>https://github.com/mangyel6/ReID-Survey, RAGA – https://github.com/microsoft/</code> Relation-Aware-Global-Attention-Networks, ABD – <code>https://github.com/VITA-Group/</code> ABD-Net, LAG – <code>https://github.com/SWJTU-3DVision/LAG-Net</code>



Fig. 5. Re-id performance measured by the mean average precision (mAP) on the Market-1501 benchmark dataset [4] for different attention blocks integrated at a single position in the ResNet-50 (as displayed in Fig. 3). Downwards pointing triangles symbolize that the bar is below the displayed range of values. Results for positions where inference with attention blocks is slower than ResNet-101 but does not match its performance, or where its performance is even worse than ResNet-50, are shown in light colors to symbolize that these configurations should not be considered.

tion, and hyperparameters regarding positions of integration while considering computational costs as measured above. Fig. 5 shows the re-id performance when attention blocks are inserted at a single position measured by the mean average precision (mAP) on Market-1501 [4]. Furthermore, the plot shows mean and standard deviation over three training runs.

As we can see, even a single attention block integrated at a suitable position can top the performance of the ResNet-101 that is twice as deep as the baseline. Furthermore, we can see that channel-wise and non-local attention perform very similar in this regime of few attention blocks. By considering the lower inference costs, we clearly should prefer channel-wise attention over non-local attention.

D. Where Should We Add Attention Blocks?

Now that we have seen that adding a single attention block is a powerful approach, we like to investigate whether combinations of computationally effective channel-wise attention blocks at different positions can further improve the performance. From the extensive NAS with plenty of combinations of positions and types of attention, as part of the progressive reduction of the design space, we were able to derive the following set of rules to improve the reidentification performance most by using as few as possible attention blocks:

- 1) Select positions where the single-position performance is good (see Fig. 5). Most often, positions at the end of each stage seem to perform slightly better than earlier positions in each stage.
- 2) The distance between selected positions should be large. Positions in different ResNet stages should be preferred.
- 3) A combination of different types of attention (channel-wise, non-local) performs slightly better than single-type combinations (but is not worth it for robotic applications due to the computational cost of non-local attention blocks).

The best results of the NAS for each type of attention block was achieved for a pair of attention blocks at positions 8 and 14, and for the triplet of attention blocks at positions 6, 8, and 14. Using the same type of attention block at more than three positions did not improve the results further. Tab. I lists the best results for attention blocks at multiple positions. Furthermore, Fig. 1 shows the re-id performance versus inference speed in a graph. We can see that for the channelwise attention blocks HAC and C-NL, the inference speed is affected only slightly. Inference with the proposed C-NL at positions 6, 8, and 14 is only 1.706 ms slower (ca. 5%) than the ResNet-50 baseline without any attention blocks (38.623 ms vs. 36.917 ms). Therefore, while just 1.196 fewer batches can be processed per second, the re-identification performance far exceeds that of ResNet-101 (mAP 0.8806 \pm 0.0007 vs. 0.8707 \pm 0.0006), which is twice as deep and takes 53.454 ms to process a batch.

E. Can the Performance Be Improved Further?

In [64], an alternative way of improving the reidentification performance was examined. By using a largemargin loss function, the recognition rate is increased without affecting the inference speed. We also apply the best performing loss function, namely circle loss [63]. In [64], circle loss achieved a mAP of 0.882 on the Market-1501 dataset using the strong baseline [58] setup as in our experiments. We observed, that the combination of a large-margin loss function and few attention blocks at suitable positions complement each other very well and boosts the performance further (see Tab. I). The proposed C-NL attention block matches the performance of the much more costly AGW non-local attention block, while being faster during inference (highlighted in bold).

TABLE I: Re-id performance on the Market-1501 benchmark dataset [4] and inference speed for attention blocks at single and multiple positions using cross entropy and circle loss [63] respectively.

		Re-id performance	Inference speed
Attention	Loss function	[mAP]	[batches/second]
HAC [23] @ 8	cross entropy	0.8733 ± 0.0006	25.17
AGW [35] @ 7	cross entropy	0.8733 ± 0.0012	22.59
C-NL @ 8	cross entropy	0.8755 ± 0.0007	26.70
HAC [23] @ 8, 14	cross entropy	0.8773 ± 0.0006	25.06
AGW [35] @ 8, 14	cross entropy	0.8805 ± 0.0014	21.09
C-NL @ 8, 14	cross entropy	0.8788 ± 0.0007	25.91
HAC [23] @ 6, 8, 14	cross entropy	0.8780 ± 0.0010	24.93
AGW [35] @ 6, 8, 14	cross entropy	0.8815 ± 0.0008	19.88
C-NL @ 6, 8, 14	cross entropy	0.8806 ± 0.0007	25.89
HAC [23] @ 6, 8, 14	circle loss [63]	0.8897 ± 0.0012	24.93
AGW [35] @ 6, 8, 14	circle loss [63]	0.8916 ± 0.0006	19.88
C-NL @ 6, 8, 14	circle loss [63]	0.8916 ± 0.0005	25.89

V. EXPERIMENTS ON A ROBOTIC DATASET

To examine, whether our improvements by using attention blocks at few positions can also be achieved on robotic data, we conducted an experiment on an extended version of the ROREAS dataset, which was recorded with a mobile robot following a user in a clinical environment [8]. We will refer to this extended version as ROREAS+ in the following.

The ROREAS+ dataset differs significantly from typical re-id benchmark datasets, like Market-1501 [4], DukeMT-MC-reID [60], or CUHK03-NP [61]. These re-id benchmark datasets have been obtained from footage taken from outdoor surveillance cameras mounted at several meters height. In contrast, our target robotic platforms, as in [8], operates inside buildings and the camera is mounted at eye level. This results in completely different images, which differ in perspective, occlusions, lighting conditions, and contrast from the typical benchmark datasets. Through our evaluation on the ROREAS+ robotics dataset, which follows the protocol of [35] as in all experiments in this paper, we should get a clear picture regarding how much attention helps to improve person re-identification in a real-world robotics application scenario. Due to data protection laws, we are not allowed to show pictures, but only report benchmarking results.

The ROREAS+ dataset contains 421 different persons. In order to get a realistic setup, out of 421 recorded persons, we have chosen only 50 persons with the largest amount of recorded images as training data, the rest is taken as test data to get a representative test set with many different (but potentially similarly looking) persons included, where attention can make the difference. Recording many views for 50 people with a mobile robot can already be a challenging and potentially time-consuming task in some scenarios, but we found this number to be a good trade-off between the amount of work needed to record the data and the amount of data available for a deep-learning approach. Therefore, 34012 images of 50 identities are available for training, while the test set is split into 945 query images and 2520 gallery images of 371 identities.

Due to the limited amount of training data, we must use transfer learning to profit from pre-training on larger re-id benchmark datasets. Especially the small amount of different identities can be an issue, otherwise. For transfer learning, we first pre-train on Market-1501 [4] and then use a two-step fine-tuning approach [65], where first the randomly initialized classification layer is fine-tuned on our target dataset ROREAS+, followed by the fine-tuning of all network weights. For training, we follow the strong baseline setup [58]. As loss function, we employ circle loss, as it showed to be beneficial in our previous experiment. Following our best results, we added the best performing

TABLE II: Re-id performance on the robotic dataset ROREAS+ using a ResNet-50 with and without attention.

-	Dataset pre-taining	Dataset fine-tuning	Attention	mAP
(1)	Market-1501 [4]		_	0.6630
(2)	Market-1501 [4]	—	C-NL	0.6684
(3)	Market-1501 [4]	ROREAS+	_	0.7551
(4)	Market-1501 [4]	ROREAS+	C-NL	0.7658

attention block C-NL at positions 6, 8, and 14.

In Tab. II, we report the best results over three runs, since this would be the configuration that we would select for later application. First, we could use a ResNet-50 for feature extraction that was trained on re-id benchmark data only, as in our previous experiments, without any fine-tuning on ROREAS+. We can see, that using a ResNet-50 for feature extraction that was never trained on robotic data (rows 1, 2) performs worse, both with and without attention, compared to the corresponding results with fine-tuning on ROREAS+ (rows 3, 4).

In both cases – with or without finetuning – C-NL noticeably improves the ability to re-identify persons. The standard deviation over three runs is $\sigma = 0.006$. Therefore the results with attention are ca. 1.5σ better than without attention.

VI. CONCLUSION

In the state of the art on person re-identification, inference cost is widely neglected, and adding attention blocks at only few positions was not considered so far. We have shown that adding a channel-wise attention block at a single position in a ResNet-50 is sufficient to perform better in person re-identification than a ResNet-101. Therefore, the common practice of adding attention to each residual block should be questioned.

In addition, due to the progressive reduction of the design space during neural architecture search, we were able to derive a set of rules for where attention blocks should be integrated in a ResNet architecture. By integrating the proposed attention block C-NL at three positions in a ResNet-50 derived by the NAS, we improved the mAP for the Market-1501 dataset, surpassing the performance of the ResNet-101 by a large margin, while slowing down the inference by only 5%, which is marginal compared to the 45% slowdown caused by doubling the depth. Combining this architectural modification with a large-margin loss function further improved the re-identification ability. Compared to HAC block [23], which slows down inference by 9% (nearly twice as much as C-NL) and performs 1.6σ worse in this setup, the modifications of the C-NL attention block building on the channel-wise attention block setup of HAC proved effective.

Finally, we confirmed the performance gain for robotic data when transfer learning with pre-training on a re-id benchmark dataset is applied. The improved mAP shows that few attention blocks in the architecture enable the extracted re-id features to better discriminate individuals in difficult cases where individuals look very similar. This will help the robot to better keep track of its users and will consequently improve clinical applications where a mobile robot needs to follow or guide its user.

In future work it would be interesting to apply the proposed approach to neural architectures designed for mobile devices that have shown fast inference speed on Jetson hardware, such as ShuffleNet v2 [66] or MobileNet v3 [67]. We are curious if the same design rules apply and if a single channel-wise attention block is sufficient for these architectures as well.

REFERENCES

- [1] H.-M. Gross, A. Scheidig, K. Debes, E. Einhorn, M. Eisenbach, S. Mueller, T. Schmiedel, T. Q. Trinh, C. Weinrich, T. Wengefeld, *et al.*, "ROREAS: Robot coach for walking and orientation training in clinical post-stroke rehabilitation: Prototype implementation and evaluation in fields trials," *Autonomous Robots*, vol. 41, no. 3, pp. 679–698, 2017.
- [2] A. Scheidig, B. Schütz, T. Q. Trinh, A. Vorndran, A. Mayfarth, C. Sternitzke, E. Röhner, and H.-M. Gross, "Robot-assisted gait selftraining: Assessing the level achieved," *Sensors*, vol. 21, no. 18, p. 6213, 2021.
- [3] A. Vorndran, T. Q. Trinh, S. Mueller, A. Scheidig, and H.-M. Gross, "How to always keep an eye on the user with a mobile robot?" in *International Symposium on Robotics (ISR)*. VDE, 2018, pp. 1–7.
- [4] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian, "Scalable person re-identification: A benchmark," in *International Conference* on Computer Vision (ICCV), 2015, pp. 1116–1124.
- [5] N. Bellotto and H. Hu, "Multisensor data fusion for joint people tracking and identification with a service robot," in *International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2007, pp. 1494–1499.
- [6] V. Alvarez-Santos, X. M. Pardo, R. Iglesias, A. Canedo-Rodriguez, and C. V. Regueiro, "Feature analysis for human recognition and discrimination: Application to a person-following behaviour in a mobile robot," *Robotics and Autonomous Systems*, vol. 60, no. 8, pp. 1021–1036, 2012.
- [7] J. Satake, M. Chiba, and J. Miura, "Visual person identification using a distance-dependent appearance model for a person following robot," *International Journal of Automation and Computing*, vol. 10, pp. 438– 446, 2013.
- [8] M. Eisenbach, A. Vorndran, S. Sorge, and H.-M. Gross, "User recognition for guiding and following people with a mobile robot in a clinical environment," in *International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 3600–3607.
- [9] K. Koide and J. Miura, "Identification of a specific person using color, height, and gait features for a person following robot," *Robotics and Autonomous Systems*, vol. 84, pp. 76–87, 2016.
- [10] N. An, S.-Y. Sun, X.-G. Zhao, and Z.-G. Hou, "Online context-based person re-identification and biometric-based action recognition for service robots," in *Chinese Control And Decision Conference (CCDC)*. IEEE, 2017, pp. 3369–3374.
- [11] S. Cosar, C. Coppola, N. Bellotto, et al., "Volume-based human reidentification with rgb-d cameras." in *International Joint Conference* on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP), 2017, pp. 389–397.
- [12] M. J. Islam, J. Hong, and J. Sattar, "Person-following by autonomous robots: A categorical overview," *International Journal of Robotics Research*, vol. 38, no. 14, pp. 1581–1618, 2019.
- [13] K. Koide and J. Miura, "Convolutional channel features-based person identification for person following robots," in *International Conference* on *Intelligent Autonomous Systems (IAS)*. Springer, 2019, pp. 186– 198.
- [14] R. Algabri and M.-T. Choi, "Deep-learning-based indoor human following of mobile robot using color feature," *Sensors*, vol. 20, no. 9, p. 2699, 2020.
- [15] —, "Online boosting-based target identification among similar appearance for person-following robots," *Sensors*, vol. 22, no. 21, p. 8422, 2022.
- [16] A. Eirale, M. Martini, and M. Chiaberge, "Human-centered navigation and person-following with omnidirectional robot for indoor assistance and monitoring," *Robotics*, vol. 11, no. 5, p. 108, 2022.
- [17] S. Müller, T. Wengefeld, T. Q. Trinh, D. Aganian, M. Eisenbach, and H.-M. Gross, "A multi-modal person perception framework for socially interactive mobile service robots," *Sensors*, vol. 20, no. 3, p. 722, 2020.
- [18] T. Wengefeld, M. Eisenbach, T. Q. Trinh, and H.-M. Gross, "May i be your personal coach? bringing together person tracking and visual re-identification on a mobile robot," in *International Symposium on Robotics (ISR)*. VDE, 2016, pp. 1–8.
- [19] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), 2018, pp. 7132–7141.
- [20] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang, "Residual attention network for image classification," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), 2017, pp. 3156–3164.

- [21] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 7794–7803.
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.
- [23] W. Li, X. Zhu, and S. Gong, "Harmonious attention network for person re-identification," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 2285–2294.
- [24] F. Yang, K. Yan, S. Lu, H. Jia, X. Xie, and W. Gao, "Attention driven person re-identification," *Pattern Recognition (PR)*, vol. 86, pp. 143– 155, 2019.
- [25] D. Wu, C. Wang, Y. Wu, Q.-C. Wang, and D.-S. Huang, "Attention deep model with multi-scale deep supervision for person reidentification," *Transactions on Emerging Topics in Computational Intelligence (TETCI)*, vol. 5, no. 1, pp. 70–78, 2021.
- [26] G. Chen, C. Lin, L. Ren, J. Lu, and J. Zhou, "Self-critical attention learning for person re-identification," in *International Conference on Computer Vision (ICCV)*, 2019, pp. 9637–9646.
- [27] Z. Ji, X. Zou, X. Lin, X. Liu, T. Huang, and S. Wu, "An attentiondriven two-stage clustering method for unsupervised person reidentification," in *European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 20–36.
- [28] X. Gong, Z. Yao, X. Li, Y. Fan, B. Luo, J. Fan, and B. Lao, "LAG-Net: Multi-granularity network for person re-identification via local attention system," *Transactions on Multimedia (TMM)*, 2021.
- [29] C. Wang, Q. Zhang, C. Huang, W. Liu, and X. Wang, "Mancs: A multi-task attentional network with curriculum sampling for person reidentification," in *European Conference on Computer Vision (ECCV)*, 2018, pp. 365–381.
- [30] B. Chen, W. Deng, and J. Hu, "Mixed high-order attention network for person re-identification," in *International Conference on Computer Vision (ICCV)*, 2019, pp. 371–381.
- [31] S. Zhou, F. Wang, Z. Huang, and J. Wang, "Discriminative feature learning with consistent attention regularization for person re-identification," in *International Conference on Computer Vision* (*ICCV*), 2019, pp. 8040–8049.
- [32] C. Li, S. Jin, E. Chang, S. Xuan, F. Liu, and D. Xu, "SCANet: Person re-identification with semantically consistent attention," in *Chinese Automation Congress (CAC)*. IEEE, 2020, pp. 3424–3428.
- [33] C. Song, Y. Huang, W. Ouyang, and L. Wang, "Mask-guided contrastive attention model for person re-identification," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 1179–1188.
- [34] S. Zhou, J. Wang, D. Meng, Y. Liang, Y. Gong, and N. Zheng, "Discriminative feature learning with foreground attention for person re-identification," *Transactions on Image Processing (TIP)*, vol. 28, no. 9, pp. 4671–4684, 2019.
- [35] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. Hoi, "Deep learning for person re-identification: A survey and outlook," *Transactions* on Pattern Analysis and Machine Intelligence (TPAMI), 2021.
- [36] Z. Zhang, C. Lan, W. Zeng, X. Jin, and Z. Chen, "Relation-aware global attention for person re-identification," in *IEEE/CVF Conference* on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 3186– 3195.
- [37] B. N. Xia, Y. Gong, Y. Zhang, and C. Poellabauer, "Second-order nonlocal attention networks for person re-identification," in *International Conference on Computer Vision (ICCV)*, 2019, pp. 3760–3769.
- [38] T. Chen, S. Ding, J. Xie, Y. Yuan, W. Chen, Y. Yang, Z. Ren, and Z. Wang, "ABD-Net: Attentive but diverse person re-identification," in *International Conference on Computer Vision (ICCV)*, 2019, pp. 8351–8361.
- [39] Y. Zhang, X. Li, and Z. Zhang, "Learning a key-value memory co-attention matching network for person re-identification," in AAAI conference on artificial intelligence (AI), 2019, pp. 9235–9242.
- [40] X. Lan, H. Wang, S. Gong, and X. Zhu, "Deep reinforcement learning attention selection for person re-identification," in *British Machine Vision Conference (BMVC)*, 2017.
- [41] J. Zhang, L. Niu, and L. Zhang, "Person re-identification with reinforced attribute attention selection," *Transactions on Image Processing* (*TIP*), vol. 30, pp. 603–616, 2020.
- [42] J. Xu, R. Zhao, F. Zhu, H. Wang, and W. Ouyang, "Attention-aware compositional network for person re-identification," in *IEEE/CVF*

Conference on Computer Vision and Pattern Recognition (CVPR), 2018, pp. 2119–2128.

- [43] C.-P. Tay, S. Roy, and K.-H. Yap, "AANet: Attribute attention network for person re-identifications," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 7134–7143.
- [44] X. Qian, Y. Fu, T. Xiang, Y.-G. Jiang, and X. Xue, "Leader-based multi-scale attention deep architecture for person re-identification," *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 42, no. 2, pp. 371–385, 2019.
- [45] S. Li, S. Bak, P. Carr, and X. Wang, "Diversity regularized spatiotemporal attention for video-based person re-identification," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 369–378.
- [46] J. Si, H. Zhang, C.-G. Li, J. Kuen, X. Kong, A. C. Kot, and G. Wang, "Dual attention matching network for context-aware feature sequence based person re-identification," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 5363–5372.
- [47] N. Karianakis, Z. Liu, Y. Chen, and S. Soatto, "Reinforced temporal attention and split-rate transfer for depth-based person re-identification," in *European Conference on Computer Vision (ECCV)*, 2018, pp. 715– 733.
- [48] A. Subramaniam, A. Nambiar, and A. Mittal, "Co-segmentation inspired attention networks for video-based person re-identification," in *International Conference on Computer Vision (ICCV)*, 2019, pp. 562– 572.
- [49] Y. Fu, X. Wang, Y. Wei, and T. Huang, "STA: Spatial-temporal attention for large-scale video-based person re-identification," in AAAI conference on artificial intelligence (AI), 2019, pp. 8287–8294.
- [50] G. Chen, J. Lu, M. Yang, and J. Zhou, "Spatial-temporal attentionaware learning for video-based person re-identification," *Transactions* on *Image Processing (TIP)*, vol. 28, no. 9, pp. 4192–4205, 2019.
- [51] Z. Zhang, C. Lan, W. Zeng, and Z. Chen, "Multi-granularity reference-aided attentive feature aggregation for video-based person re-identification," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 10407–10416.
- [52] M. Shim, H.-I. Ho, J. Kim, and D. Wee, "ReAD: Reciprocal attention discriminator for image-to-video re-identification," in *European Conference on Computer Vision (ECCV)*, 2020, pp. 335–350.
- [53] X. Li, W. Zhou, Y. Zhou, and H. Li, "Relation-guided spatial attention and temporal refinement for video-based person re-identification," in *AAAI conference on artificial intelligence (AI)*, 2020, pp. 11434– 11441.
- [54] M. Li, H. Xu, J. Wang, W. Li, and Y. Sun, "Temporal aggregation with clip-level attention for video-based person re-identification," in *IEEE/CVF Winter Conference on Applications of Computer Vision* (WCACV), 2020, pp. 3376–3384.
- [55] M. Ye, J. Shen, D. J Crandall, L. Shao, and J. Luo, "Dynamic dual-attentive aggregation learning for visible-infrared person reidentification," in *European Conference on Computer Vision (ECCV)*, 2020, pp. 229–247.
- [56] A. Haque, A. Alahi, and L. Fei-Fei, "Recurrent attention models for depth-based person identification," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 1229– 1238.
- [57] I. Radosavovic, R. P. Kosaraju, R. Girshick, K. He, and P. Dollár, "Designing network design spaces," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [58] H. Luo, Y. Gu, X. Liao, S. Lai, and W. Jiang, "Bag of tricks and a strong baseline for deep person re-identification," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2019.
- [59] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [60] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi, "Performance measures and a data set for multi-target, multi-camera tracking," in *European Conference on Computer Vision (ECCV)*, 2016, pp. 17– 35.
- [61] W. Li, R. Zhao, T. Xiao, and X. Wang, "Deepreid: Deep filter pairing neural network for person re-identification," in *IEEE/CVF Conference* on Computer Vision and Pattern Recognition (CVPR), 2014.
- [62] G. Wang, J. Lai, P. Huang, and X. Xie, "Spatial-temporal person reidentification," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, 2019, pp. 8933–8940.

- [63] Y. Sun, C. Cheng, Y. Zhang, C. Zhang, L. Zheng, Z. Wang, and Y. Wei, "Circle loss: A unified perspective of pair similarity optimization," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), 2020, pp. 6398–6407.
- [64] D. Aganian, M. Eisenbach, J. Wagner, D. Seichter, and H.-M. Gross, "Revisiting loss functions for person re-identification," in *International Conference on Artificial Neural Networks (ICANN)*, 2021, pp. 30–42.
- [65] H. Chen, Y. Wang, Y. Shi, K. Yan, M. Geng, Y. Tian, and T. Xiang, "Deep transfer learning for person re-identification," in *International Conference on Multimedia Big Data (BigMM)*. IEEE, 2018, pp. 1–5.
- [66] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, "Shufflenet v2: Practical guidelines for efficient cnn architecture design," in *European Conference on Computer Vision (ECCV)*, 2018, pp. 116–131.
- [67] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, *et al.*, "Searching for mobilenetv3," in *International Conference on Computer Vision (ICCV)*, 2019, pp. 1314–1324.