

MagnifierNet: Towards Semantic Adversary and Fusion for Person Re-identification

Yushi Lan*, Yuan Liu*¹

¹ SenseTime Group Limited

lanyushi15@gmail.com, yliu050@e.ntu.edu.sg

² The Chinese University of HongKong

Xinchi Zhou, Maoqing Tian, Xuesen

Zhang, Shuai Yi¹

zhouxinchi, tianmaoqing, zhangxuesen, yishuai@sensetime.com

Hongsheng Li²

hsl@ee.cuhk.edu.hk

Abstract

Although person re-identification (ReID) has achieved significant improvement recently by enforcing part alignment, it is still a challenging task when it comes to distinguishing visually similar identities or identifying the occluded person. In these scenarios, magnifying details in each part features and selectively fusing them may provide a feasible solution. In this work, we propose MagnifierNet, a triple-branch network which accurately mines details from whole to parts. Firstly, the holistic salient features are encoded by a global branch. Secondly, to enhance detailed representation for each semantic region, the "Semantic Adversarial Branch" is designed to learn from dynamically generated semantic-occluded samples during training. Meanwhile, we introduce "Semantic Fusion Branch" to filter out irrelevant noises by selectively fusing semantic region information sequentially. To further improve feature diversity, we introduce a novel loss function "Semantic Diversity Loss" to remove redundant overlaps across learned semantic representations. State-of-the-art performance has been achieved on three benchmarks by large margins. Specifically, the mAP score is improved by 6% and 5% on the most challenging CUHK03-L and CUHK03-D benchmarks.

1 Introduction

Person Re-Identification (ReID) has attracted significant attention for its critical role in video surveillance and public security. Specifically, given a query image, the ReID system amounts to retrieve all the images of the same identity according to their semantic similarity from a large gallery, typically captured by distinctively different cameras from various viewpoints.

Although advancement has been witnessed, there are still some challenging issues to be better solved, such as pose variation, body occlusion and the part misalignment. Former works attempt to tackle these problems via methods including attention mechanism [26, 46], body feature cropping [32, 45], and human parsing [24]. Most of these works mainly focus on the enhancement of feature alignment, while the exploitation of fine-grained details is often ignored. Detailed semantic information, however, is critical for further improvements, especially in cases where identities are heavily occluded or only share minor differences in certain parts.

*Y.Lan and Y.Liu share equal contribution. This work was done when Y.Lan and Y.Liu were interns at SenseTime Group Limited.

To effectively address the challenges mentioned above, we propose MagnifierNet, a novel triple-branch network that not only extracts aligned representation but also magnifies fine-grained details and selectively fuses each semantic region. The framework is constructed from whole to parts. Apart from holistic representation provided by conventional global branch, the Semantic Adversarial Branch (*SAB*) learns fine-grained representation for each semantic region, and the Semantic Fusion Branch (*SFB*) fuses each semantic feature selectively to focus only on beneficial information. Meanwhile, a light-weight mask module is applied to impose alignment constraint on the feature map and provide segmentation masks during inference. In this way, our model is able to distinguish similar identities, recognize occluded samples with semantic details, as well as alleviate negative influence from noises.

Notably, the semantic features under different body part masks might share overlapping information, as these masks are generated at high-level feature maps with large receptive fields. These overlaps could reduce the network’s capability to capture distinctive details of different semantic regions. Hence, we propose a novel Semantic Diversity (*SD*) Loss to improve diversity among different parts, which further enhances the network performance as shown in our experiments.

Our main contribution can be summarized as follows:

- We propose MagnifierNet, a novel triple-branch ReID framework to improve ReID performance by magnifying nontrivial pedestrian details, selectively fusing features and improving feature diversity.
- Technically, our proposed Semantic Adversarial Branch magnifies fine-grained details in each semantic region and regularizes network learning. Our Semantic Fusion Branch which selectively fuses semantic information sequentially, focusing on beneficial features while filtering out noises. We further improve semantic feature diversity by introducing a novel Semantic Diversity Loss.
- Experimental results show that our proposed MagnifierNet achieves state-of-the-art performance on three ReID benchmark datasets including Market-1501, DukeMTMC-reID, and CUHK03-NP.

2 Related Work

ReID and Part Based Methods Person ReID task aims to retrieve target images belonging to the same person based on their similarity. In terms of *feature representation*, the great success of deep convolution networks has pushed ReID benchmarks to a new level [6, 34, 46, 52]. Recently, researchers have intensively focused on extracting local person features to enforce part alignment and feature representation capability. For instance, Zhao *et al.* [45] proposed a body region proposal network which utilized human landmark information to compute relative aligned part representation. Zhao *et al.* [46] shared a similar idea, but with feature representation generated from K part detectors. Sun *et al.* [34] proposed a Part-based Convolution Baseline (PCB) network which focuses on the consistency of uniform partition part-level feature with refined stripe pooling. Based on PCB, MGN [37] and Pyramidal Model [47] explored multi-branch networks to learn features of different granularity, which attempted to incorporate global and local information. However, their methods result in complicated structures and seriously increased parameters, and also lacks the ability to represent accurate human regions. In comparison, our approach aggregates multi-granularity features in a more efficient way, resulting in smaller network size and better performance. Recently,

researchers also tried to address current ReID limitations by exploring 3D representation [50], federated learning [58] and unsupervised learning [12, 15].

Semantic Alignment in Retrieval Retrieval tasks often suffer from serious occlusions and huge variances in viewpoints, both in person ReID and face recognition. Under such circumstances, many researchers aim to push the margin of retrieval tasks via feature or image alignment. Zhou *et al.* [53] propose a method to generate photo-realistic rotated faces, acting as an effective data augmentation strategy for face recognition. The development of human parsing methods [16, 24] and pose estimation [3, 40, 41] facilitate the use of human semantic information as external cues to promote the performance of ReID task. Xu *et al.* [44] resort to the assistance of predicted keypoints confidence map to extract aligned human parts representation. Similarly, Huang *et al.* [23] proposed person-region guided pooling which shows significant improvement in cross-domain ReID task. Kalayel *et al.* [24] and Tian *et al.* [36] adopts body regions as supervision to drive the model to learn alignment representation and solve background problem. However, these methods treat each semantic region equally and did not further explore the details in each region.

Metric Learning for ReID With the adoption of large datasets and deeper convolution networks, ReID task has gradually evolved to a metric learning problem, which aims to retrieve target images based on their feature similarity. For this purpose, recent works in metric learning have paid intensive attention to loss function design, such as triplet loss [17], center loss [39] and margin-based loss [11, 22, 38], etc. Hard triplet mining strategy has also been successfully utilized in ReID [8, 10]. Nevertheless, the above-mentioned approaches mostly focus on image-level feature refinement. Meanwhile, orthogonality constraints are explored by [53] for ReID to encourage the learning of informative and diverse features. To further exploit the motivation, the Semantic Diversity Loss in our work encourages orthogonality directly on semantic feature level, which accurately improves diversity among body regions that are of great significance for the ReID task.

3 Methodology

This section introduces our proposed framework, including the Semantic Adversarial Branch (*SAB*) to magnify attentive feature of each region, the Semantic Fusion Branch (*SFB*) to fuse and filter extracted features and a commonly adopted Global Branch to capture holistic representations. A light-weight mask module (*M*) is introduced to parse semantic information of human parts for *SAB* and *SFB*. The whole network is trained end-to-end in the first training stage. The Semantic Diversity Loss is then added to fine-tune the network in the second training stage. The overview of the proposed network structure is shown in Figure 1.

3.1 Mask Module and Semantic Aligned Representation

The main idea of the mask module is to extract accurate human landmark features, through which pixel-level representations belonging to the same human semantic landmark can be aggregated. Given the binary segmentation masks for all body parts R_{seg} and the feature map F_{cm} of a given pedestrian image x , we can generate the *Semantic Aligned Representation* via the outer product (\otimes). The semantic information can then be further leveraged by other components of the network, along with a pooling operation. We adopt global max-pooling operation ($GMP(\cdot)$) in our experiments.

Specifically, we rescale R_{seg} into the same size as F_{cm} using bilinear interpolation and calculate the *Semantic Aligned Representation* X^{ID} as shown in Equation 1,

$$X^{ID} = \{x_k^{ID}\}_{k=1}^K = F_{cm} \otimes R_{seg} \quad (1)$$

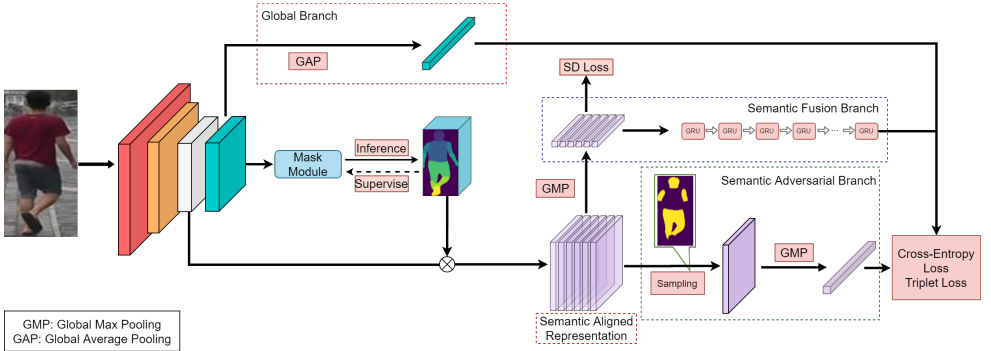


Figure 1: The overall architecture of the proposed MagnifierNet. The Semantic Adversarial Branch explores the hidden details in each semantic region by limiting available information for training. The Semantic Fusion Branch encodes beneficial features from each semantic region sequentially while filtering out noises. The *SD Loss* further improves representation diversity among all semantic features.

where K is the region number and x_k^{ID} is essentially the aligned representation of the semantic region it belongs to, $x_k \in R^c$ and c is the channel number. For the pixels predicted as background, we still aggregate the representations of all pixels as the background can provide complementary information and augment the original representations.

We obtain R_{seg} through a pre-trained Human Parsing model [13] for all datasets as the pseudo mask label during training. To obtain R_{seg} during inference and reinforce alignment across identities [23], we include a light-weight segmentation head after backbone output in the mask module with the structure proposed in [13]. It predicts K human region masks from input feature maps, and is supervised by the generated pseudo labels during training. More complex segmentation head designs have the potential to improve performance but are not the focus of our work.

3.2 Semantic Adversarial Branch

Though *Semantic Aligned Representations* have been acquired, we argue that detailed feature representations haven't been captured in all semantic regions. Further magnification of semantic details is necessary to maintain robust model performance, especially in crowded venues where the identities are usually partially visible. Inspired by the success of adversarial samples, we propose **Semantic Adversarial Branch (SAB)** which dynamically generates semantic-level perturbed samples to drive the network to learn discriminative representation for each semantic region. These carefully tailored adversarial samples are generated online and trained along with original images in a separate branch, which provides complementary information to strengthen the network training.

Adversarial Sample Generation Given an input image, we dynamically occlude some semantic regions with a constant following certain strategy. The samples we generate are adversarial to the network, but with more analogy to real-world scenarios which can drive the network to learn more attentive features of each semantic region. To exploit the advantages of the multi-branch network and utilize the computed *Semantic Aligned Representations*, we directly generate adversarial samples on learned feature maps.

Concretely, given an image in a batch, we sample a semantic region group \hat{O} made up of aligned semantic representations X^{ID} formulated in Equation 1. Semantic aligned feature maps corresponding to the selected regions in \hat{O} are zeroed out. This operation is applied to all feature channels and the generated adversarial samples of an identity can be formulated

as Equation 2,

$$F_{cnn}^{adversarial} = GMP\left(\sum_{k \notin \hat{O}} x_k^{ID}\right) \quad (2)$$

Region Sampling Strategy The intuition behind the occlusion sampling is to reinforce attentive representation learning of all regions via erasing the same semantic aligned features in each batch. Therefore, the most straightforward option is to randomly sample and occlude region(s) from X^{ID} and only keep \hat{K} regions. We call this sampling strategy **Random-baseline**.

Beyond this, we note that the remaining semantic regions should maintain sufficient yet dynamic cues to represent correct identity to drive the network to learn fine-grained representation for all semantic regions. Therefore, we manually split the aligned semantic regions into upper torso regions U_{torso} and lower torso regions L_{torso} . In this way, we can perform sampling in two groups independently, avoiding the situations where the feature of an identity may lose discrimination under **Random-baseline** strategy. Specifically, we sample in a way such that equal number of regions remain in U_{torso} and L_{torso} , namely $\hat{K}/2$ each where \hat{K} is an even number. We name this sampling strategy **Random-torso**, which shows better results in our experiments.

As semantic regions are dynamically occluded in each batch, limited number of semantic regions are left to represent an identity. Hence, the network is forced to capture fine-grained details in each semantic region instead of relying on certain discriminative region(s). Therefore, the model can learn more detailed and attentive representations on viable semantic regions and naturally alleviate occlusion or similar identity problems.

3.3 Semantic Fusion Branch

To yield correct predictions even in cases of similar identities or noisy backgrounds, the model is supposed to focus more on informative semantic regions instead of equally considering noisy parts with misleading information.

To address this problem, we propose the **Semantic Fusion Branch (SFB)** to selectively fuse all semantic regions, focusing on informative features while filtering out noises. Considering the gated units of GRU [9] can estimate the importance of a given semantic feature conditioned by itself and other encoded features, it is natural to apply GRU to encode these semantic features sequentially to focus more on relevant information.

We perform $GMP(\cdot)$ on the *Semantic Aligned Representation* to obtain a feature for each body part, and then feed the semantic aligned features sequentially into a one-layer GRU cell, and adopt the last GRU output as the semantic fusion of the input image. To compute the loss, we adopt the ‘‘BNNeck’’ method proposed by [29] to jointly minimize cross-entropy loss and triplet loss. Besides, we apply ReID supervision for each semantic feature to stabilize the training process, but the separate features are not utilized during inference.

Notably, the two branches *SFB* and *SAB* share the same backbone features and segmentation masks. Hence, *SFB* benefits directly from *SAB*’s semantic detail magnification process. As the feature representation under each semantic mask grows more informative, *SFB* is able to accurately capture this information and selectively fuses them into a better feature embedding for the ReID task.

3.4 Semantic Feature Diversification

As the semantic masks are generated at high-level feature maps, features under different semantic masks might share overlapping information due to large receptive fields, which re-




Figure 2: Visualization of feature map activation with different network components. Both Semantic Adversarial and Fusion Branch are able to help the network capture fine-grained semantic details.

duces the network’s capability to capture details from different semantic regions. To improve diversity among semantic representations, we propose a novel **Semantic Diversity Loss (SD Loss)** to remove redundancy among region features.

For a pair of semantic regions, we aim to diversify them by increasing their distance in the feature space. Specifically, we minimize the pair-wise cosine similarity among all semantic features as shown in Equation 3.

$$L_{SD} = \frac{1}{N * C \binom{K}{2}} \sum_{k=1}^N \sum_{j=1}^{K-1} \sum_{i=j+1}^K \frac{(GMP(x_i^{ID}))^T GMP(x_j^{ID})}{\max(\|GMP(x_i^{ID})\| * \|GMP(x_j^{ID})\|, \epsilon)} \quad (3)$$

where x_i^{ID} and x_j^{ID} are a pair of semantic aligned features and GMP stands for Global Max-pooling. N is batch size, K is the number of semantic regions, T is matrix transpose, $\| \cdot \|$ is vector magnitude and ϵ is a small positive value used to avoid division by zero.

As SFB naturally generates individual semantic features $p(x_k^{ID})$ from *Semantic Aligned Representation*, we calculate SD Loss via SFB on these pooled semantic features. Notably, as SFB and SAB share the same backbone layers and segmentation masks, minimizing SD Loss through SFB forces the backbone to extract diverse features which supports SAB to further magnify semantic details. The overall loss function of the network is shown in Equation 4, which is shared by all the network branches.

$$L_{Total} = L_{cls} + L_{tri} + \gamma L_{SD} + \lambda L_{mask}, \quad (4)$$

where the first two terms stand for cross-entropy and triplet loss that optimize all branches, L_{SD} is SD Loss and L_{mask} distills segmentation information which is used during inference. Weight coefficients γ and λ balance the importance of semantic diversification and segmentation respectively.

4 Experiments

4.1 Implementation Details

We implement our model with reid-strong-baseline [29] framework and train it on a single Tesla V100 GPU. Both SAB and SFB are connected to ResNet-50 stage 3’s output with the final down-sampling layer removed to maintain more spatial information. All person images are re-sized to 384×192 as in [10]. Coefficient λ is set to 2 and γ is set to $2e - 3$, which are verified on validation set and yields robust results across all datasets. The training procedure is two-stage. Firstly, the model is trained near convergence without SD Loss for 460 epochs. We simply adopt the optimizer settings of reid-strong-baseline [29] for fair comparison. Then we fine-tuned the network on Equation 3 for another 40 epochs with learning rate 0.001. The feature embeddings from all branches are concatenated as the final

person representation during inference. The whole training procedure has 500 epochs and can be finished within 4 hours.

Method	CUHK03-L		CUHK03-D		DukeMTMC		Market1501	
	R1	mAP	R1	mAP	R1	mAP	R1	mAP
AOS [22]	-	-	47.1	43.3	79.2	62.1	86.5	70.4
MGCAM [32]	50.14	50.21	46.7	46.9	79.2	62.1	83.8	74.3
PCB [34]	-	-	63.7	57.5	83.3	69.2	93.8	81.6
MGN [57]	68.0	67.4	66.8	66.0	88.7	78.4	95.7	86.9
Pyramid [47]	78.9	76.9	78.9	74.8	89.0	79.0	95.7	88.2
DG-Net [53]	-	-	65.6	61.1	86.6	74.8	94.8	86.0
CAMA [42]	70.1	66.5	66.6	64.2	85.8	72.9	94.7	84.5
CASN [49]	73.7	68.0	71.5	64.4	87.7	73.7	94.4	82.8
BDB [11]	73.6	71.7	72.8	69.3	86.8	72.1	94.2	84.3
DSA [24]	78.9	75.2	78.2	73.1	86.2	74.3	95.7	87.6
AA-Net [55]	-	-	-	-	87.7	74.3	93.9	83.4
IANet [21]	-	-	-	-	87.1	73.4	94.4	83.1
MHN [9]	77.2	72.4	71.7	65.4	89.1	77.2	95.1	85.0
P^2 -Net [17]	78.3	73.6	74.9	68.9	86.5	73.1	95.2	85.6
OS-Net [56]	-	-	72.3	67.8	88.6	73.5	94.8	84.9
ABD-Net [2]	-	-	-	-	89.0	78.6	95.6	88.3
MMGA [2]	-	-	-	-	89.5	78.1	95.0	87.2
FPR [19]	-	-	76.1	72.3	88.6	78.4	95.4	86.6
SCAL [8]	74.8	72.3	71.1	68.6	89.0	79.6	95.8	89.3
CAR [57]	-	-	-	-	86.3	73.1	96.1	84.7
Ours	82.4	79.6	80.2	77.1	90.0	80.7	95.8	89.6
AOS*	-	-	54.6	56.1	84.1	78.2	88.7	83.3
AlignedRe-ID [43]*	-	-	-	-	-	-	94.4	90.7
MGN*	-	-	-	-	-	-	96.6	94.2
TriNet [20]*	-	-	-	-	-	-	86.6	81.1
AA-Net*	-	-	-	-	90.4	86.9	95.1	92.4
SPT [23]*	-	-	-	-	88.3	83.3	93.5	90.6
Ours*	87.3	89.1	86.3	87.5	91.8	90.6	96.5	95.1

Table 1: Comparison with other State-Of-the-Art methods. Re-ranking [60, 52] is applied for method with “*”. Top performances on each dataset are in **bold**.

Semantic mask labels are transformed from the Densepose [11] dataset. For Semantic Adversarial Branch, U_{torso} contains {head, upper arm, lower arm, chest} and L_{torso} includes {upper leg, lower leg, foot} respectively.

4.2 Datasets and Protocols

Market1501 [48] dataset contains 32,688 images of 1,501 identities captured by 6 cameras. The training set contains 751 identities with 12,936 images, the testing set contains 750 identities with 3,368 query images and 15,913 gallery images.

CUHK03-NP [25] dataset contains two sub-sets, namely CUHK03-L (labelled set) and CUHK03-D (detected set) according to data generation method. The CUHK03-L contains 14,096 images and CUHK03-D contains 14,097 images. We adopt the new train / test split protocol proposed by [25], which contains 767 identities for training and 700 identities for testing.

DukeMTMC-reID [31, 51] dataset is collected from 8 cameras. contains 16,522 training images of 702 identities, 222,8 query images and 17,661 gallery images in the testing set of the other 702 identities.

Protocols. We adopt two commonly used evaluation protocols in our experiments. Rank-1 identification rate (R1) and the mean Average Precision (mAP).

4.3 Comparison with other State-of-the-Arts

We compare our method against 21 State-of-the-Art methods and present the results in Table 1. Our method achieves the best results on all the datasets except for Rank-1 on Market1501. Notably, our method outperforms other networks by a large margin on the small and challenging dataset CUHK03 with relatively heavy viewpoint limitation and occlusion. Furthermore, the MagnifierNet outperforms part-alignment related methods including BDB [44] and P^2 -Net [47], which validates the effectiveness of our method to go beyond alignment and further improve Re-ID performance via semantic adversarial and fusion. Meanwhile, similar body-part masks are adopted [4, 52] to improve ReID performance. However, the granularity of their method is coarser compared to ours since the individual body parts are neglected. In addition, the results also showcase the superiority of our approach over other methods with feature regularization such as AOS [22], as our method extracts attentive features accurately from each aligned semantic region and perform end-to-end adversarial training directly on the semantic level.

4.4 Ablation Experiments

We have carried out extensive ablation studies to validate the effectiveness of each module in MagnifierNet, which will be covered in the following sections.

Improvement from each Network Component There are three important components in our framework: M to impose alignment constraint, SAB to encourage attentive semantic feature learning, and SFB to filter and fuse learned representations. We add network components one by one onto the baseline which contains only the Global Branch. The performance improvements from each component are presented in Table 2. It is shown that SAB and SFB are able to boost the performance both individually and simultaneously to a large margin. The proposed SD Loss is able to attain further gain on all datasets, which validates its effectiveness across different domains.

In addition, to further elaborate the benefits of each component qualitatively, we visualize the saliency maps of some randomly selected images using different settings in Figure 2. The baseline alone can only capture a coarse representation of the image. Both SAB and SFB are able to significantly improve model’s activation on each semantic region, while SAB tends to highlight every region and SFB tends to selectively focus on certain semantic parts. When applying dual branches simultaneously, their individual advantages complement each other which yields a joint-representation that selectively highlights hidden details in each informative region. Our proposed model only adds 6% additional time cost compared to the baseline methods on a single 1080 GPU, which is reasonable considering our significant performance gain.

Effect of Semantic Adversarial We visualize the top 1 querying results between the baseline and baseline + SAB for two probe images in Figure 3. As shown in the feature activations, SAB can significantly improve model’s sensitivity on semantic details. For the first query on top, SAB model magnifies the top 1 gallery image’s upper body and foot region for comparison. For the second query at the bottom, SAB model captures its top 1’s collar and back

Components					CUHK03-L		DukeMTMC	
G	M	SAB	SFB	SD	R1	mAP	R1	mAP
✓					69.8	67.4	82.7	70.8
✓	✓				71.3 (+1.5)	69.0 (+1.6)	83.9 (+1.2)	72.9 (+2.1)
✓	✓	✓			76.6 (+6.8)	74.6 (+7.2)	87.1 (+4.4)	76.7 (+5.9)
✓	✓		✓		77.4 (+7.6)	75.9 (+8.5)	86.7 (+4.0)	75.2 (+4.4)
✓	✓	✓	✓		81.2 (+11.4)	78.5 (+11.1)	88.6 (+5.9)	79.6 (+8.8)
✓	✓	✓	✓	✓	82.4 (+12.6)	79.6 (+12.2)	90.0 (+7.3)	80.7 (+9.9)

Table 2: Ablation study on the impact of different network components. We consider the Global Branch G as the baseline performance. The improvement from each component over the baseline is indicated in the bracket.




Figure 3: Top 1 search results with Baseline and Baseline + SAB . The baseline fail to retrieve the correct image and capture distinctive features as shown in red, while SAB manages to generate correct result and locate attentive features as shown in green.

region when retrieving. However, the baseline method produces incorrect results on these cases due to inability to highlight critical details as shown in its activation maps.

In addition, we also analyze different sampling strategies for the partial semantic representation. Specifically, we train MagnifierNet with different \hat{K} in SAB for CUHK03-L dataset. As shown in Figure 5, the mAP of SAB with **Random-torso** is consistently superior to that with **Random-baseline**. We therefore demonstrate the effectiveness of our proposed sampling strategy, while the best performance is both achieved when \hat{K} is 4, which is the setting for SAB in our experiments. The study is done without SD Loss to highlight the impact of semantic adversarial itself.

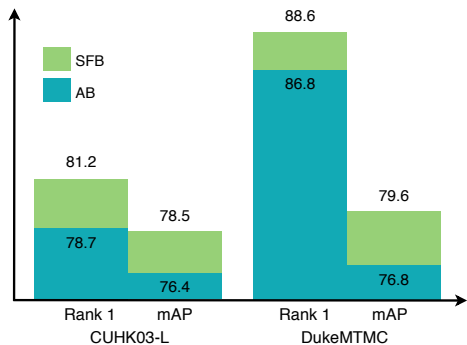
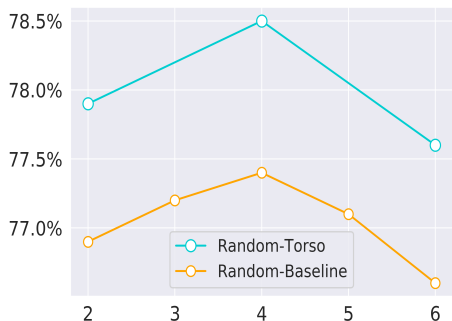


Figure 4: Comparison between SFB and its Ab- Figure 5: Effect of different number of remained lation Branch, where SFB significantly improves semantic regions \hat{K} . Note that \hat{K} is even for Random-torso as discussed in Section 3.2.



Effect of Semantic Fusion To validate the benefit of SFB , we replace it with an “ablation

branch" (AB) and compare their performances. To construct the AB , instead of feeding semantic features into GRU sequentially, we directly perform a 1×1 convolution to reduce their dimensions followed by concatenation to ensure AB has the same output feature dimension as SFB . We train both models until converge without SD Loss to validate the effectiveness of network structure alone. As shown in Figure 4, the proposed semantic fusion method SFB surpasses AB on all metrics, which validates its positive influence on our framework.

5 Conclusion

In this paper, we propose a novel network MagnifierNet that improves ReID performance beyond pure alignment. The Semantic Adversarial Branch mines the fine-grained details in each semantic region by learning with limited semantic representation, the Semantic Fusion Branch selectively encodes semantic features by filtering out noises and focusing only on beneficial information. We further improve the model performance by introducing a novel Semantic Diversity Loss which promotes feature diversity among semantic regions. MagnifierNet achieves State-of-the-Art performance on three major datasets Market1501, DukeMTMC-reID, and CUHK03, which showcases the effectiveness of our method.

References

- [1] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018.
- [2] Honglong Cai, Zhiguan Wang, and Jinxing Cheng. Multi-scale body-part mask guided attention for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [4] Binghui Chen, Weihong Deng, and Jiani Hu. Mixed high-order attention network for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 371–381, 2019.
- [5] Dapeng Chen, Dan Xu, Hongsheng Li, Nicu Sebe, and Xiaogang Wang. Group consistent similarity learning via deep crf for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8649–8658, 2018.
- [6] Guangyi Chen, Chunze Lin, Liangliang Ren, Jiwen Lu, and Jie Zhou. Self-critical attention learning for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9637–9646, 2019.
- [7] Tianlong Chen, Shaojin Ding, Jingyi Xie, Ye Yuan, Wuyang Chen, Yang Yang, Zhou Ren, and Zhangyang Wang. Abd-net: Attentive but diverse person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8351–8361, 2019.

- [8] De Cheng, Yihong Gong, Sanping Zhou, Jinjun Wang, and Nanning Zheng. Person re-identification by multi-channel parts-based cnn with improved triplet loss function. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1335–1344, 2016.
- [9] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.
- [10] Zuozhuo Dai, Mingqiang Chen, Xiaodong Gu, Siyu Zhu, and Ping Tan. Batch drop-block network for person re-identification and beyond. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3691–3701, 2019.
- [11] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [12] Shengyong Ding, Liang Lin, Guangrun Wang, and Hongyang Chao. Deep feature learning with relative distance comparison for person re-identification. *Pattern Recognition*, 48(10):2993 – 3003, 2015. ISSN 0031-3203. doi: <https://doi.org/10.1016/j.patcog.2015.04.005>. URL <http://www.sciencedirect.com/science/article/pii/S0031320315001296>. Discriminative Feature Learning from Big Data for Visual Recognition.
- [13] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3146–3154, 2019.
- [14] Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. *arXiv preprint arXiv:2001.01526*, 2020.
- [15] Yixiao Ge, Feng Zhu, Rui Zhao, and Hongsheng Li. Structured domain adaptation with online relation regularization for unsupervised person re-id. *arXiv: Computer Vision and Pattern Recognition*, 2020.
- [16] Ke Gong, Xiaodan Liang, Dongyu Zhang, Xiaohui Shen, and Liang Lin. Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [17] Jianyuan Guo, Yuhui Yuan, Lang Huang, Chao Zhang, Jin-Ge Yao, and Kai Han. Beyond human parts: Dual part-aligned representations for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3642–3651, 2019.
- [18] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [19] Lingxiao He, Yinggang Wang, Wu Liu, Xingyu Liao, He Zhao, Zhenan Sun, and Jiashi Feng. Foreground-aware pyramid reconstruction for alignment-free occluded person re-identification. *arXiv preprint arXiv:1904.04975*, 2019.

- [20] Alexander Hermans*, Lucas Beyer*, and Bastian Leibe. In Defense of the Triplet Loss for Person Re-Identification. *arXiv preprint arXiv:1703.07737*, 2017.
- [21] Ruibing Hou, Bingpeng Ma, Hong Chang, Xinqian Gu, Shiguang Shan, and Xilin Chen. Interaction-and-aggregation network for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9317–9326, 2019.
- [22] Houjing Huang, Dangwei Li, Zhang Zhang, Xiaotang Chen, and Kaiqi Huang. Adversarially occluded samples for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5098–5107, 2018.
- [23] Houjing Huang, Wenjie Yang, Xiaotang Chen, Xin Zhao, Kaiqi Huang, Jinbin Lin, Guan Huang, and Dalong Du. Eanet: Enhancing alignment for cross-domain person re-identification, 2018.
- [24] Mahdi M. Kalayeh, Emrah Basaran, Muhittin Gökmen, Mustafa E. Kamasak, and Mubarak Shah. Human semantic parsing for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [25] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 152–159, 2014.
- [26] Wei Li, Xiatian Zhu, and Shaogang Gong. Harmonious attention network for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2285–2294, 2018.
- [27] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheroface: Deep hypersphere embedding for face recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [28] Chuanchen Luo, Yuntao Chen, Naiyan Wang, and Zhaoxiang Zhang. Spectral feature transformation for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4976–4985, 2019.
- [29] Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.
- [30] Danfeng Qin, Stephan Gammeter, Lukas Bossard, Till Quack, and Luc Van Gool. Hello neighbor: Accurate object retrieval with k-reciprocal nearest neighbors. In *CVPR 2011*, pages 777–784. IEEE, 2011.
- [31] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *European Conference on Computer Vision workshop on Benchmarking Multi-Target Tracking*, 2016.
- [32] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Mask-guided contrastive attention model for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1179–1188, 2018.

- [33] Yifan Sun, Liang Zheng, Weijian Deng, and Shengjin Wang. Svdnet for pedestrian retrieval. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3800–3808, 2017.
- [34] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [35] Chiat-Pin Tay, Sharmili Roy, and Kim-Hui Yap. Aanet: Attribute attention network for person re-identifications. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7134–7143, 2019.
- [36] Maoqing Tian, Shuai Yi, Hongsheng Li, Shihua Li, Xuesen Zhang, Jianping Shi, Junjie Yan, and Xiaogang Wang. Eliminating background-bias for robust person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5794–5803, 2018.
- [37] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *2018 ACM Multimedia Conference on Multimedia Conference*, pages 274–282. ACM, 2018.
- [38] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [39] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European conference on computer vision*, pages 499–515. Springer, 2016.
- [40] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [41] Jing Xu, Rui Zhao, Feng Zhu, Huaming Wang, and Wanli Ouyang. Attention-aware compositional network for person re-identification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [42] Wenjie Yang, Houjing Huang, Zhang Zhang, Xiaotang Chen, Kaiqi Huang, and Shu Zhang. Towards rich feature discovery with class activation maps augmentation for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1389–1398, 2019.
- [43] Xuan Zhang, Hao Luo, Xing Fan, Weilai Xiang, Yixiao Sun, Qiqi Xiao, Wei Jiang, Chi Zhang, and Jian Sun. Alignedreid: Surpassing human-level performance in person re-identification. *arXiv preprint arXiv:1711.08184*, 2017.
- [44] Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Densely semantically aligned person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 667–676, 2019.
- [45] Haiyu Zhao, Maoqing Tian, Shuyang Sun, Jing Shao, Junjie Yan, Shuai Yi, Xiaogang Wang, and Xiaoou Tang. Spindle net: Person re-identification with human body region guided feature decomposition and fusion. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

- [46] Liming Zhao, Xi Li, Yueting Zhuang, and Jingdong Wang. Deeply-learned part-aligned representations for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3219–3228, 2017.
- [47] Feng Zheng, Cheng Deng, Xing Sun, Xinyang Jiang, Xiaowei Guo, Zongqiao Yu, Feiyue Huang, and Rongrong Ji. Pyramidal person re-identification via multi-loss dynamic training. *arXiv preprint arXiv:1810.12193*, 2018.
- [48] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.
- [49] Meng Zheng, Srikrishna Karanam, Ziyang Wu, and Richard J Radke. Re-identification with consistent attentive siamese networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5735–5744, 2019.
- [50] Zhedong Zheng and Yi Yang. Person re-identification in the 3d space. *arXiv preprint arXiv:2006.04569*, 2020.
- [51] Zhedong Zheng, Liang Zheng, and Yi Yang. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- [52] Zhedong Zheng, Liang Zheng, and Yi Yang. A discriminatively learned cnn embedding for person re-identification. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 14(1):13, 2018.
- [53] Zhedong Zheng, Xiaodong Yang, Zhiding Yu, Liang Zheng, Yi Yang, and Jan Kautz. Joint discriminative and generative learning for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2138–2147, 2019.
- [54] Zhun Zhong, Liang Zheng, Donglin Cao, and Shaozi Li. Re-ranking person re-identification with k-reciprocal encoding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1318–1327, 2017.
- [55] Hang Zhou, Jihao Liu, Ziwei Liu, Yu Liu, and Xiaogang Wang. Rotate-and-render: Unsupervised photorealistic face rotation from single-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5911–5920, 2020.
- [56] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Omni-scale feature learning for person re-identification. *arXiv preprint arXiv:1905.00953*, 2019.
- [57] Sanping Zhou, Fei Wang, Zeyi Huang, and Jinjun Wang. Discriminative feature learning with consistent attention regularization for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8040–8049, 2019.
- [58] Weiming Zhuang, Yonggang Wen, Xuesen Zhang, Xin Gan, Yin Daiying, Dongzhan Zhou, Shuai Zhang, and Shuai Yi. Performance optimization of federated person re-identification via benchmark analysis. In *2020 ACM Multimedia Conference on Multimedia Conference*. ACM, 2020.