

Beyond Appearance: a Semantic Controllable Self-Supervised Learning Framework for Human-Centric Visual Tasks

Weihua Chen, Xianzhe Xu, Jian Jia, Hao Luo, Yaohua Wang, Fan Wang, Rong Jin, Xiuyu Sun[†]
Alibaba Group

{kugang.cwh, xianzhe.xxz, jj359864, michuan.lh, xiachen.wyh, fan.w, jinrong.jr, xiuyu.sxy}@alibaba-inc.com

Abstract

Human-centric visual tasks have attracted increasing research attention due to their widespread applications. In this paper, we aim to learn a general human representation from massive unlabeled human images which can benefit downstream human-centric tasks to the maximum extent. We call this method SOLIDER, a Semantic cOntrollable seLf-superviseD lEaRning framework. Unlike the existing self-supervised learning methods, prior knowledge from human images is utilized in SOLIDER to build pseudo semantic labels and import more semantic information into the learned representation. Meanwhile, we note that different downstream tasks always require different ratios of semantic information and appearance information. For example, human parsing requires more semantic information, while person re-identification needs more appearance information for identification purpose. So a single learned representation cannot fit for all requirements. To solve this problem, SOLIDER introduces a conditional network with a semantic controller. After the model is trained, users can send values to the controller to produce representations with different ratios of semantic information, which can fit different needs of downstream tasks. Finally, SOLIDER is verified on six downstream human-centric visual tasks. It outperforms state of the arts and builds new baselines for these tasks. The code is released in <https://github.com/tinyvision/SOLIDER>.

1. Introduction

Human-centric visual analysis plays an important role in widespread applications, such as surveillance, sports, augmented reality, and video production. Person re-identification [13, 14, 41, 82], attribute recognition [72, 78], person search [76, 92], pedestrian detection [3, 24, 32], human parsing [27, 53], and pose estimation [58, 90]




Figure 1. A representation space learned by DINO [6]. Seven human images are represented in seven different colors. Each image is split into four parts according to their semantic regions, *i.e.*, upper body (as ▲), lower body (as +), shoes (as ★) and background (as ×, not visualized to avoid distraction). It can be seen that different parts of a same person are closer to each other even they share different semantic meanings.

have achieved considerable progress in recent years. In another aspect, there are massive human images available in the current computer vision community. For example, even an unlabeled person re-identification dataset, LUPerson [25, 26] (#Img \approx 4.18M) is 4 time larger than the ImageNet dataset (#Img \approx 1M). How to use unlabeled data to build a human representation is challenging, especially when it needs to benefit various downstream tasks.

Self-supervised learning has achieved great developments by using unlabeled data to learn representations. Many pretext tasks have been designed, such as contrastive learning [6, 10, 35] and masking image modeling [2, 34, 77, 95]. Although these methods have achieved great success in learning general image representations, there is a lack of specific design targeting human-centric tasks.

Some researchers [55, 81, 97] focus to extend self-supervised learning methods on human-centric visual tasks. They use DINO [6] with LUPerson [25, 26] dataset to build pre-trained models for person re-identification task. When applying the pre-trained models to other human-centric

[†]Corresponding Author

tasks, such as human parsing and pedestrian detection, we usually get sub-optimal results. It is due to the lack of semantic information in their learned representations. As shown in Fig. 1, in the representation space learned by DINO [6]¹, different parts of a same person are gathered together due to their appearance continuity, no matter what semantic meanings they have.

As we’ve known, semantic information is as important as appearance information for human-centric visual tasks [40, 50, 96]. Therefore, we tend to train the representation with more semantic information to extend the representation to different downstream human-centric visual tasks. In this paper, a Semantic cOntrollable seLf-supervIseD IEaRning framework (SOLIDER) is proposed. In SOLIDER, we take advantage of prior knowledge from human images to discover semantic information, which can produce pseudo semantic labels for every token. And a token-level semantic classification pretext task is imported and supervised by these pseudo labels. With the new pretext task, we can train representation with stronger semantic information.

During the usage of our trained representation on downstream tasks, we find that even though semantic information and appearance information are both important, different downstream tasks require different ratios of them. Adjusting their ratio in the representation would lead to a better performance in downstream tasks. However, as long as the pretext task is trained, the representation can not be changed in current self-supervised learning methods. Different from previous methods, we design SOLIDER as a conditional network involving a semantic controller. The controller takes a value as input and produces a latent representation. In the usage of the pre-trained model from SOLIDER, we send a value (indicating the ratio of semantic information in the representation) to the controller which can adjust the model and output a representation with the required ratio.

In summary, our paper makes four contributions:

- 1) A general human representation is learned in this paper, which is used as a better pre-trained model benefiting to downstream human-centric visual tasks.

- 2) A semantic controllable self-supervised learning framework (SOLIDER) is proposed. It takes advantages of prior knowledge in human images to produce pseudo semantic labels, and utilize it to train the human representation with more semantic information.

- 3) A semantic controller is designed in SOLIDER. With the controller, the pre-trained model can generate representations with various degrees of semantic information that can meet different needs of downstream tasks.

- 4) The effectiveness of the SOLIDER representation is verified on six downstream human-centric tasks. We believe this paper can promote the development of these human-centric tasks in computer vision community.

¹MAE [34] shares a similar phenomenon as DINO [6].

2. Related Work

2.1. Self-supervised learning approaches

Self-supervised learning approaches have attracted significant interests in computer vision, especially when used for learning image representation as the pretext task.

In current self-supervised learning, contrastive methods [5, 6, 10, 15, 35, 95] have achieved great success, and provided state-of-the-art performance for image representation. The goal of contrastive learning is to minimize the distances between two augmented views of a same image [5, 10], and distinguish each image from all the others [6, 15, 35]. MoCo [15, 35] improves the training of contrastive methods by storing representations from a momentum encoder instead of the trained network. SimCLR [10] shows that the memory bank can be entirely replaced with the elements from the same batch if the batch is large enough. DINO [6] combines most of these techniques, including momentum encoder [35], multi-crop training [5], and the use of small patches with ViTs [22], which builds a much better baseline.

Besides contrastive methods, masked image modeling methods [2, 34, 77] arouse extensive attention from researchers. BEiT [2] proposes to predict discrete tokens. SimMIM [77] and MAE [34] find that a moderately large masked size and a light prediction head would be beneficial for the masked image modeling. Although masked language modeling manages to import semantic information into image representations, it can not explicitly figure out the semantic information from the image to supervise the training. In the proposed SOLIDER, we cluster tokens and use human prior to assign semantic labels to these tokens, which can train a stronger human semantic representation.

It is worth noting that DeepCluster [4] also use cluster to learn representations. But it is on image-level and can not produce semantic labels for tokens. Some unsupervised semantic segmentation methods [18] extend DeepCluster on pixel-level. But it still faces problems and can not learn a satisfied representation, which is explained in Section. 3.1

2.2. Human-centric visual tasks

By looking through human-centric visual tasks in computer vision, we find that there are many tasks directly or indirectly related to human², *e.g.*, person re-identification [23, 28, 29, 37, 82], attribute recognition [51, 72, 78], person search [76, 92], pedestrian detection [3, 24, 32], multi-object tracking [11, 12, 48, 49, 67], human parsing [27, 53] and pose estimation [58, 90].

Among them, we sort out six representative tasks, *i.e.*, person re-identification, attribute recognition, person search, pedestrian detection, human parsing and pose

²In this paper, we focus on feature representations for images, therefore, only downstream tasks on 2D image-level are our main concerns, and tasks about 3D data or sequential data are out of our scope.

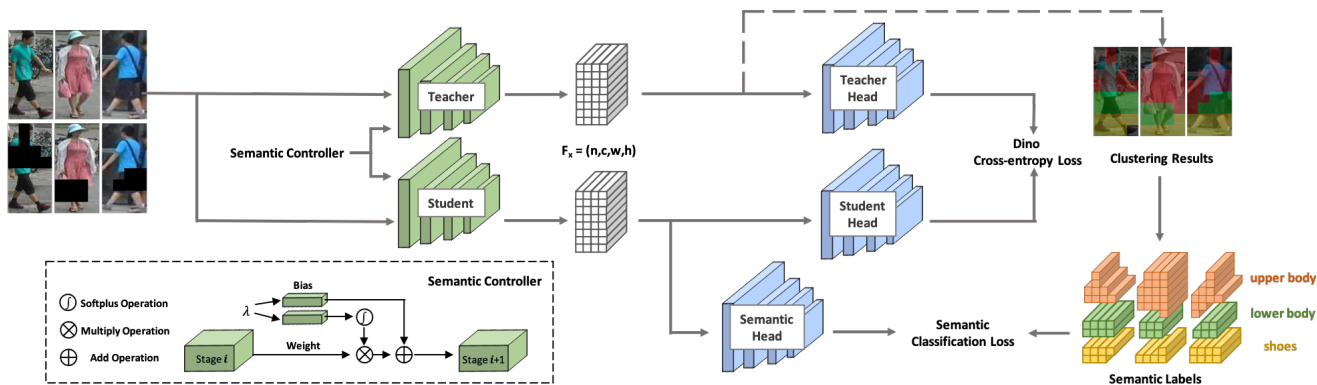


Figure 2. The pipeline of the proposed SOLIDER.

estimation. Person re-identification [8, 39, 59, 69, 89] aims at retrieving a person of interest across multiple non-overlapping cameras. Attribute recognition [20, 42, 51, 72] tries to mine the attributes of target people when given person images, of which the attributes are understandable semantic descriptions. Person search [76, 92] aims to find a probe person from the whole scene which shows great significance in video surveillance community to track lost people. Pedestrian detection [21, 24, 32, 86] focuses on detecting people in general images, which is a key ability for a variety of important applications. Human parsing [27, 43, 53], a sub-task of semantic segmentation, aims to understand human-body parts on pixel level. Human pose estimation [1, 58, 83] dedicates to locate the human body skeletons from images.

Recently, Wang et al. [84] also focus on human-centric visual tasks and design a Token Clustering Transformer (TCFormer), but TCFormer is a supervised method which can not take advantages of massive public unlabeled data.

3. The SOLIDER

The whole pipeline of the proposed SOLIDER is shown in Fig. 2. In this section, we first explain how to generate pseudo semantic labels from human prior knowledge and use it to supervise a token-level semantic classification pretext task. Then, we introduce how the learned representation can be controlled in SOLIDER.

3.1. Semantic Supervision from Human Prior

DINO [6] is a state-of-the-art self-supervised learning method, which is widely used for image representation. As a contrastive learning based method, the visual appearance information is well learned in DINO’s representation. We use DINO as our baseline and plan to involve more semantic information into its representation.

We cluster the token vectors from the learned DINO representation and show in Fig. 3(b). It can be seen that the representation is split into several parts based on its visual appearance. In other words, it can find things in images,

although be not able to tell their meanings.

Some methods [4, 18] go further and try to assign semantic labels for these clustered things. Instead of cluster on single image, they do the clustering across images and aim to build the semantic relationship among images. However, as the features used for clustering are trained from visual appearance clue, the clustered results are dominated by appearance. As shown in Fig. 1, in original DINO space, the features of “blue short” (+) is closer to “blue shirt” (▲) due to similar appearance, but farther to “black short” (+), even though they share the same semantic meaning of “short pants”. After clustering, the blue things (“blue short” and “blue shirt”) will be clustered together, and the things with the same semantic meaning (“blue short” and “black short”) are split apart. In other words, these methods can not produce the semantic labels we wanted.

After scanning through the unlabeled human images, we find that most of images have a fixed pattern [64, 96]: the person body erectly occupies the entire image, and the head is always on the top of the image while the feet is at the bottom of the image. With this observation, we give semantic labels to clustered parts of all images based on the order of their y-axis coordinates, *i.e.*, the top part of all images are labeled as upper body, and the bottom part is marked as shoes. These pseudo semantic labels provide semantic information for every token vector. A token-level semantic classification pretext task is imported and supervised by these semantic labels. To better use these labels, we involve extra modifications, listed as following.

Involving Background&Foreground Clustering. In these labels, we find that there are some background fragments which cause noise disturbance on the clustering results and mislead the alignment. To handle this problem, we introduce another clustering before the semantic clustering. Specifically, we observe that the background tokens always have smaller responses compared to foreground tokens during training [57, 96], as shown in Fig. 3(c). Thus, we cluster the token vectors into two categories based on their vector magnitudes, *i.e.*, foreground and background. The results are shown in Fig. 3(d). Then the semantic clustering




Figure 3. (a) The original images. (b) Initial semantic clustering results. (c) Attention maps. (d) Foreground clustering results. (e) Final semantic clustering results.

is only performed on the foreground tokens, and the new results are listed in Fig. 3(e). It can be clearly seen that the foreground tokens are well clustered into three semantic parts, *i.e.*, upper body, lower body, shoes³.

Involving Masked Image Modeling. Inspired by the masked image modeling methods [34, 77], we would like to further introduce more semantic information into the representation. As we have known, people can easily locate every semantic part in a human image even when some parts are occluded. Therefore, we assume that if a semantic part is missing or occluded in the human image, the model would still be able to predict its semantic meaning based on other surrounding parts. For this purpose, we upgrade our semantic supervision into masked semantic supervision. Specifically, we randomly mask out a semantic part from the image x and re-feed this masked image \tilde{x} through the framework. Then the output tokens are supervised by the original semantic labels, because it is expected that the model would be able to provide the true semantic labels for the masked tokens with the help of other tokens.

The whole semantic self-supervision is presented in Fig. 2. Specifically, during training, for every iteration, we obtain the output feature maps F from backbone with the size of (n, c, h, w) . For each image x , we consider its feature maps⁴ F_t as $w \times h$ token vectors with the size of c . Then we use K-means [31] to cluster them into two categories according to the magnitude (*i.e.*, l2 normalization) of c , and consider the category with larger magnitude as foreground mask M . After that, another K-means is applied on the tokens in foreground mask M to conduct the semantic clustering, to obtain N predefined semantic categories and assign semantic label y for each token. Meanwhile, a semantic head⁵ is involved to classify the vectors F_s from the student branch based on these semantic labels y . The corresponding semantic classification loss is as below:

³In this paper, we treat heads/chests/arms as upper body, waist/thighs as lower body and calves/feet as shoes.

⁴The feature maps F_t from the teacher network is chosen for semantic clustering. The feature maps from the student network is marked as F_s .

⁵The semantic head contains several blocks, and each block includes a fully connected layer, a batch norm layer and a ReLU.

$$L_{sm} = \frac{1}{w \times h} \sum_{\substack{u \in w \\ v \in h}} \sum_{i=1}^{N+1} -y^{(u,v)} \log \frac{f_s(u,v)^{(i)}}{\sum_{k=1}^{N+1} f_s(u,v)^{(k)}} \quad (1)$$

where $f_s = h_{sm}(flatten(F_s))$. $flatten()$ is used to reshape F_s from (n, c, h, w) to $(n * h * w, c)$, and h_{sm} is the semantic head. N means the number of clustered semantic parts. $f_s(u,v)^{(i)}$ indicates the predicted probability of token (u,v) on part i . After image x , we randomly mask out a part of x to obtain image \tilde{x} , and send \tilde{x} to Eq. 1 too.

The total loss for the whole SOLIDER framework is:

$$L = \alpha L_{dino} + (1 - \alpha) L_{sm} \quad (2)$$

where α is a balance weight and set to 0.5 by experience.

3.2. Semantic Controller

When the learned representation is used as the pre-trained model for downstream tasks, we expect the appearance and semantic information learned in the representation can help downstream tasks. However, each downstream task has its own focus. For example, in person re-identification [41, 82], although the semantic information can help to align parts across person images [62, 88], the appearance information is the key clue to distinguish different people, which is the most important information [80, 94]. So a pre-trained model with more appearance information and less semantic information would provide a better startup for person re-identification task. On the contrary, in pedestrian detection [3, 24] and human parsing [27, 43, 53], the semantic information plays a key role for these tasks, while the appearance difference is useless and should be partly ignored. Its pre-trained model is expected to contain enough semantic information.

To fit the needs of different tasks, it requires the pertained model can be adjusted according to the downstream tasks. However, as the pertained model is trained, it is hard to change its parameters for different downstream tasks. Task token [54] is a potential way to solve this problem, which pre-sets an extra one-hot token for each task. But it has several problems. First, the number of task tokens should be pre-defined before learning representation. In real applications, we can not determine what tasks the learned representation would be utilized on in advance. Second, although we can pre-defined task tokens from semantic perspective (two task tokens, *i.e.*, with and without semantic information), the task token is a discrete variable, which leaves the downstream tasks with limited choice of representations. Third, the task token is essentially a re-weighted sum of original learned tokens. It is hard to cooperate with Swin [52] backbone, which is a more friendly transformer backbone for visual tasks, because of the shifted windows in Swin [52].

In this paper, SOLIDER uses a conditional network with a semantic controller to handle this problem. During the pretext task training, due to the unavailability of downstream task, we can not adjust the representation for specific task, so we import the semantic controller to make the pre-trained model conditional. The inputs of the semantic controller are the image feature maps and a continuous value $\lambda \in [0, 1]$. The λ represents the required ratio of semantic information in representation. The output is the new feature maps with the ratio we required. The details of the proposed semantic controller can be seen in Fig. 2. In semantic controller the value λ is encoded into a weight vector and a bias vector. After a Softplus activation function, the weight vector is multiplied onto the original feature maps, and the bias vector is added for final outputs.

We apply the semantic controller after each block of Swin Backbone, and the feature maps $F(\lambda)$ with new ratio λ is sent to the next block. We use the following equation to produce our semantic controllable model:

$$L = \alpha L_{dino}(F(\lambda)) + \lambda(1 - \alpha)L_{sm}(F(\lambda)) \quad (3)$$

In the pretext task training stage, λ is randomly sampled for every iteration. We tried different distributions of λ , *i.e.*, binomial distribution $B(p = 0.5)$, continuous uniform distribution $U[0, 1]$ and beta distribution $\beta(0.2, 0.2)$. $\beta(0.2, 0.2)$ is better than $U[0, 1]$, implying that emphasizing the sampling on two borders is more important for training the controller, which is consistent with the conclusions of Mixup [85]. Finally, $B(p = 0.5)$ is found to perform the best, *a.k.a.*, a binary distribution from $\{0, 1\}$. When the pre-trained model is applied to downstream tasks, λ is set manually to adapt the pre-trained model to each downstream task. The pre-trained model with λ provides a better startup for downstream fine-tuning, and all parameters of the pre-trained model are fine-tuned.

In sum, after a conditional pre-trained model is trained from SOLIDER, we can send different λ to the model to produce representations with different ratios of semantic information for downstream tasks. A training pseudo code is provided in Supplementary Material to further clarify the whole training process of SOLIDER.

4. Experiments

The pre-trained model from SOLIDER is verified on six downstream human-centric visual tasks, including person re-identification, attribute recognition, person search, pedestrian detection, human parsing and pose estimation.

4.1. Training Settings

4.1.1 Datasets

For pretext tasks, LUPerson [25, 26] is used for training, the same as [55, 97]. It contains 4.18M human images

without any label. From each downstream task, we conduct experiments on their commonly-used datasets. Specifically, in person re-identification, the experiments are conducted on Market1501 [91] and MSMT17 [74]. In attribute recognition, PETA_{zs} [38], RAP_{zs} [38] and PA100k [51] are considered. In person search, we adopt CUHK-SYSU [76] and PRW [92] in our experiments. CityPerson [86] is utilized for pedestrian detection. In human parsing and pose estimation, LIP [27] and COCO [47] pose estimation are used respectively.

4.1.2 Evaluation Metric

In person re-identification and person search, mAP/Rank1 [91] are adopted as evaluation metrics. For attribute recognition, the evaluation metrics are mean accuracy (mA) [38]. To evaluate the performance on pedestrian detection, we employ the log-average Miss Rate over false positive on Reasonable and Heavy Occluded (MR⁻² on R/HO) [32]. Human parsing uses mIoU [43] for evaluation, and pose estimation utilizes Average Precision/Recall (AP/AR) [83] as the evaluation metric.

4.1.3 Optimization

We use Swin-Transformer [52] as the backbone throughout all the experiments. In the pretext task training, SGD is used as the optimizer, and the learning rate is 0.0005. The model is trained for 100 epochs, and the learning rate is declined by Cosine Annealing scheduler. The batch size is 48/32/24 when training on 8 Tesla V100 32G GPUs for Swin-Tiny/Small/Base.

It is worth noting that, as the semantic clustering process is time-consuming, we do not directly train SOLIDER with 100 epochs from scratch. Instead, we first train DINO with 100 epochs, and then finetune the SOLIDER on the trained DINO model with another 10 epochs using a smaller learning rate 0.00005.

For downstream tasks, we reproduce a state-of-the-art method in each task as our baseline. The methods we used as baseline are listed in Table. 1. What we do is just replacing their backbones to the Swin-Transformer [52] backbone which is pre-trained by the proposed SOLIDER.

4.2. Qualitative and Quantitative Analysis

Analysis on semantic clustering. To observe the semantic representation ability of our pre-trained model in qualitative analysis, we randomly select some images from the training data, and visualize their representation features learned before and after SOLIDER. The visualization is provided in Fig. 4. It can be found that, before introducing the semantic supervision, the representation features are distributed mainly based on the identities of the images. And the images sharing similar appearance stay closer, even

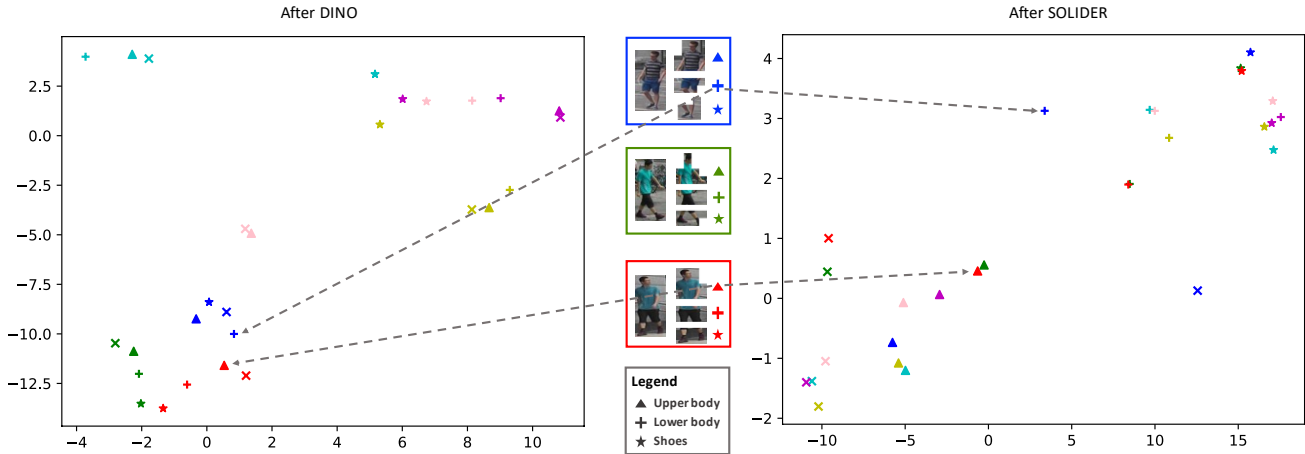


Figure 4. The representation space learned before and after involving SOLIDER. For each image, we split its features into four parts according to their semantic regions, *i.e.*, upper body (as ▲), lower body (as +), shoes (as *) and background (as ×, not shown for a clearer presentation). It can be seen that before SOLIDER is introduced, the features tend to gather by appearance. While after SOLIDER, the features with same semantic meanings are closer to each other.

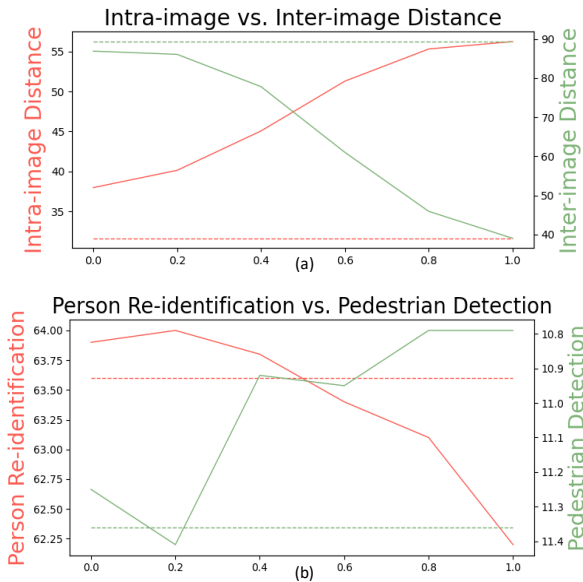


Figure 5. (a) The intra-image and inter-image distances under different λ . (b) The performance of person re-identification and pedestrian detection under different λ . The dash line indicates the distances and two task results from the original DINO model.

with different semantic meanings. For example, the features of “blue short” (+) is closer to “blue shirt” (▲) due to similar appearance “color blue”, but further to “black short” (+), even they share the same semantic meanings “short pants”.

However, after involving the semantic supervision by SOLIDER, the feature distribution pays more attention on the semantic meaning. We can find that the images hold similar semantic meanings are closer to each other, even though they share different appearance. It implies that the human prior discovered by our SOLIDER can provide a better semantic supervision for representation, which helps

the pre-trained model involves more semantic meanings.

Analysis on semantic controller. To further verify the semantic controller, we provide another two experiments. We first define two distance, *i.e.*, intra-image distance and inter-image distance. Intra-image distance is the average distance between any two parts from the same image. Inter-image distance indicates the average distance between parts of the same semantic meaning but coming from different images. Small intra-image distance and large inter-image distance indicate the appearance information is dominant in the representation, otherwise the semantic information dominates. After the SOLIDER is trained, we set the semantic weight λ from 0 to 1, to observe two distances of these different representations in the whole LUP dataset. The results are shown in Fig. 5(a). It can be seen that with the increase of λ , the intra-image distance gets larger, while the inter-image distance becomes smaller, which suggests that more semantic information is involved into the representation and become dominant.

Meanwhile, as we’ve known, person re-identification requires more appearance information, while pedestrian detection is prone to the semantic information. We list the performance of two tasks using the pre-trained representations with different λ As shown in Fig. 5(b), with the increase of λ , the person re-identification performance becomes worse and the pedestrian detection performance turns to better. In other words, the representation with a large λ can provide a better startup for pedestrian detection, while a small λ fits more to person re-identification. This phenomenon further implies the effectiveness of our semantic controller.

4.3. Ablation Study

The ablation study is provide to verify the effectiveness of each module in SOLIDER, and all the experiments are

Table 1. The performance of different pre-trained models on downstream tasks. “Sup” implies the supervised training. “+Clustering” indicate the training with semantic supervision. “+Control” means the semantic controller is involved. For each downstream task, we list its name, evaluation metric and the state of the art that we used as our baseline. \uparrow/\downarrow means the larger/smaller value the better performance.

Pretrain Methods		Sup	DINO [6]	+ Clustering	+ Clustering&Controller
Pretrain Data		ImageNet	LUP1M	LUP1M	LUP1M
Person Re-identification mAP/Rank1 \uparrow TransReID [36]	Market1501	78.1/90.2	89.6/95.9	89.5/95.5	89.9/96.1
	MSMT17	49.7/73.6	63.3/83.2	61.6/82.2	63.9/83.8
Attribute Recognition mA \uparrow RethinkPAR [38]	PETA _{zs}	72.86	73.64	73.90	74.20
	RAP _{zs}	72.10	73.04	73.16	73.21
	PA100k	80.67	82.98	82.98	84.15
Person Search mAP/Rank1 \uparrow SeqNet [45]	CUHK-SYSU	93.0/94.1	93.6/94.3	93.6/94.1	94.0/94.7
	PRW	50.0/84.4	52.9/84.7	53.0/84.0	54.1/85.0
Pedestrian Detection MR ⁻² (R/HO) \downarrow CSP [73]	CityPerson	11.6/43.8	11.4/43.1	11.1/41.7	10.8/40.7
Human Parsing mIOU \uparrow SCHP [43]	LIP	51.10	54.45	55.25	55.45
Pose Estimation AP/AR \uparrow HRFormer [83]	COCO	72.4/78.2	73.1/78.5	73.4/78.7	74.4/79.7

conducted on Swin-Tiny [52] backbone. We use different methods to train the pretext task, and verify pre-trained models on all downstream human-centric visual tasks.

ImageNet and LUP1M are datasets used for the pretext task training. ImageNet indicates the model is trained on ImageNet. This training is supervised with labels in ImageNet. LUP1M is a subset randomly sampled from LUP, which contains 1 million person images and has a similar image number with ImageNet for fair comparison. Due to no labels in LUP1M, the training process is self-supervised. We take DINO [6] as our baseline. “+Clustering” indicates that we involve the human prior into DINO by semantic clustering. “+Clustering&Controller” means that the semantic controller is imported to control the ratio of semantic information in representation. In downstream tasks, we set λ to a small value (0.0-0.2) for person-re-identification, and a large value (0.8-1.0) for pedestrian detection, human attributes, human parsing and pose estimation. For person search, where appearance information is as important as semantic information, we set λ to a moderate value. Details of the λ selection for downstream tasks are provided in Supplementary Material.

From the comparison in Table. 1, it can be found that after involving the semantic information (“+Clustering”), the pre-trained model achieves a better performance on most of downstream tasks. It implies the effectiveness of the

representation with more semantic information. It is worth noticing that the performance on person re-identification and person search are somewhat declined after involving semantic supervision. The reason is that after the semantic information is imported, the model is more inclined to represent semantic information, and its power on distinguishing different identities is weakened to some extent, which needs a more careful balance. The semantic controller is exactly designed for this problem. After involving the semantic control (“+Controller”), the performance further improved, especially for person re-identification and person search. Comparing our SOLIDER results to DINO’s on person re-identification, we can observe that our improvement is limited. Because the key clue in person re-identification is appearance information, which is already well learned in DINO. The semantic information imported from SOLIDER plays an auxiliary role leading to only a slight increment.

It is noteworthy that the larger improvement of DINO LUP1M compared to Sup ImageNet on person ReID is due to the street-view scenarios of LUPerson images which are similar to ReID datasets. Except for the ReID task, the DINO LUP1M brings an average of 1.4 improvement (72.6 vs. 74.0) on five tasks compared to Sup ImageNet, which implies the advantage of the appearance information learned from DINO. With SOLIDER, the average performance is raised to 74.9. The further 0.9 increase

Table 2. The comparison of the proposed SOLIDER with other state of the arts.

Person Re-identification mAP/Rank1 \uparrow		SCSN [16]	ABDNet [9]	TransReID [36]	UP-ReID [81]	PASS [97]	SOLIDER		
	Market1501	88.5/95.7	88.3/95.6	89.5/95.2	91.1/97.1	93.3/ 96.9	91.6/96.1	93.3/96.6	93.9/96.9
	MSMT17	58.5/83.8	60.8/82.3	69.4/86.2	63.3/84.3	74.3/89.7	67.4/85.9	76.9/90.8	77.1/90.7
Attribute Recognition mA \uparrow		MsVAA [61]	VAC [30]	ALM [66]	JLAC [65]	RethinkPAR [38]	SOLIDER		
	PETA _{zs}	71.53	71.91	73.01	73.60	71.62	74.37	76.21	76.43
	RAP _{zs}	72.04	73.70	74.28	76.38	72.32	74.23	76.84	77.06
	PA100k	80.41	79.16	80.68	82.31	81.61	84.14	86.25	86.37
Person Search mAP/Rank1 \uparrow		NAE+ [7]	AlignPS+ [79]	TCTS [68]	SeqNet [45]	GLCNet [93]	SOLIDER		
	CUHK-SYSU	92.1/92.9	94.0/94.5	93.9/95.1	94.8/95.7	95.8/96.2	94.9/95.7	95.5/95.8	94.9/95.5
	PRW	44.0/81.1	46.1/82.1	46.8/87.5	47.6/87.6	47.8/ 87.8	56.8/86.8	59.8/86.7	59.7/86.8
Pedestrian Detection MR ⁻² (R/HO) \downarrow		RepLoss [71]	CSP [73]	NMS-Loss [56]	ACSP [70]	PedesFormer [33]	SOLIDER		
	CityPerson	13.2/56.9	11.0/49.3	10.8/-	9.3/46.3	9.2/36.9	10.3/40.8	10.0/39.2	9.7/39.4
Human Parsing mIOU \uparrow		JPPNet [46]	BraidNet [50]	CE2P [60]	PCNet [87]	SCHP [43]	SOLIDER		
	LIP	51.37	54.40	53.10	57.03	59.36	57.52	60.21	60.50
Pose Estimation AP/AR \uparrow		CPN [17]	SimpleBase [75]	TokenPose [44]	HRNet [63]	HRFormer [83]	SOLIDER		
	COCO	68.6/-	74.3/79.7	75.8/80.9	76.3/81.2	77.2/82.0	74.4/79.6	76.3/81.3	76.6/81.5

shows the success of involving semantic information by our SOLIDER. And this improvement is consistent on all tasks.

Besides, we conduct experiments on the influence of the clustered part number and the semantic head size in SOLIDER, which are present in Supplementary Material.

4.4. Comparison to State of the Arts

We compare our results with state of the arts on six human-centric tasks. The state-of-the-art methods for comparison are listed in Table 2. In person re-identification, we use TransReID without side information as our baseline. Even without side information, we can achieve a better performance than other self-supervised works, *e.g.*, TransReID [36] and PASS [97], which also trained on LUPerson dataset. For the PRW results of person search, the mAP of our results outperform other state of the arts with more than 10%. The mAP criterion reflects the detection ability of models. It implies that our pre-trained model can lead to a better detection result, which may thanks to the imported semantic information. In pedestrian detection, PedesFormer [33] achieves a better performance than ours on pedestrian detection task is because it also involves extra data from autonomous driving datasets which are specific for pedestrian detection task and not for other human-centric tasks. For pose estimation, we report the best performance of HRFormer [83] which is trained on HRFormer backbone instead of Swin. After HRFormer [83] backbone is switched to Swin, the result is 75.9/81.1 re-

ported in MMPose [19], which is lower than ours 76.6/81.5. From the comparison in Table 2, we can see that the pre-trained from SOLIDER can provide a better initialization for these human-centric tasks, and can be used as a new baseline for further works on these tasks.

We also conduct SOLIDER on different backbones, also summarize in Table 2. It can be seen that with the model size increasing in Swin-Transformer backbones, the performance is further improved. In some tasks, Swin-Small achieves a better performance than Swin-Base, it is because of a larger batch size in Swin-Small compared to that in Swin-Base under a limited GPU memory.

5. Conclusion

This paper proposes a semantic controllable self-supervised learning framework called SOLIDER. It can utilize prior knowledge from human images to train representations with more semantic information. Moreover, the pre-trained model from SOLIDER can be adjusted by an input value through the semantic controller, which can produce representations with different ratios of semantic information and to satisfy the requirements on downstream tasks. The human representations from SOLIDER is verified on six human-centric visual tasks, which can promote the development of these human-centric tasks in computer vision community.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014. 3
- [2] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 1, 2
- [3] Antonio Brunetti, Domenico Buongiorno, Gianpaolo Francesco Trotta, and Vitoantonio Bevilacqua. Computer vision and deep learning techniques for pedestrian detection and tracking: A survey. *Neurocomputing*, 2018. 1, 2, 4
- [4] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)*, 2018. 2, 3
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *NIPS*, 2020. 2
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 1, 2, 3, 7
- [7] Di Chen, Shanshan Zhang, Jian Yang, and Bernt Schiele. Norm-aware embedding for efficient person search. In *CVPR*, 2020. 8
- [8] Kai Chen, Weihua Chen, Tao He, Rong Du, Fan Wang, Xiuyu Sun, Yuchen Guo, and Guiguang Ding. Tagperson: A target-aware generation pipeline for person re-identification. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2022. 3
- [9] Tianlong Chen, Shaojin Ding, Jingyi Xie, Ye Yuan, Wuyang Chen, Yang Yang, Zhou Ren, and Zhangyang Wang. Abdnnet: Attentive but diverse person re-identification. In *CVPR*, 2019. 8
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 1, 2
- [11] Weihua Chen, Lijun Cao, Xiaotang Chen, and Kaiqi Huang. A novel solution for multi-camera object tracking. In *ICIP*, 2014. 2
- [12] Weihua Chen, Lijun Cao, Xiaotang Chen, and Kaiqi Huang. An equalised global graphical model-based approach for multi-camera object tracking. *IEEE Transactions on Circuits and Systems for Video Technology*, 2016. 2
- [13] Weihua Chen, Xiaotang Chen, Jianguo Zhang, and Kaiqi Huang. Beyond triplet loss: a deep quadruplet network for person re-identification. In *CVPR*, 2017. 1
- [14] Weihua Chen, Xiaotang Chen, Jianguo Zhang, and Kaiqi Huang. A multi-task deep network for person re-identification. In *AAAI*, 2017. 1
- [15] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 2
- [16] Xuesong Chen, Canmiao Fu, Yong Zhao, Feng Zheng, Jingkuan Song, Rongrong Ji, and Yi Yang. Saliency-guided cascaded suppression network for person re-identification. In *CVPR*, 2020. 8
- [17] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7103–7112, 2018. 8
- [18] Jang Hyun Cho, Utkarsh Mall, Kavita Bala, and Bharath Hariharan. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *CVPR*, 2021. 2, 3
- [19] MMPose Contributors. Openmmlab pose estimation toolbox and benchmark. <https://github.com/open-mmlab/mmpose>, 2020. 8
- [20] Yubin Deng, Ping Luo, Chen Change Loy, and Xiaoou Tang. Pedestrian attribute recognition at far distance. In *Proceedings of the 22nd ACM international conference on Multimedia*, 2014. 3
- [21] Piotr Dollár, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian detection: An evaluation of the state of the art. *PAMI*, 2011. 3
- [22] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [23] Zhaopeng Dou, Zhongdao Wang, Weihua Chen, Yali Li, and Shengjin Wang. Reliability-aware prediction via uncertainty learning for person image retrieval. In *ECCV*, 2022. 2
- [24] Markus Enzweiler and Dariu M Gavrila. Monocular pedestrian detection: Survey and experiments. *PAMI*, 2008. 1, 2, 3, 4
- [25] Dengpan Fu, Dongdong Chen, Jianmin Bao, Hao Yang, Lu Yuan, Lei Zhang, Houqiang Li, and Dong Chen. Unsupervised pre-training for person re-identification. In *CVPR*, 2021. 1, 5
- [26] Dengpan Fu, Dongdong Chen, Hao Yang, Jianmin Bao, Lu Yuan, Lei Zhang, Houqiang Li, Fang Wen, and Dong Chen. Large-scale pre-training for person re-identification with noisy labels. In *CVPR*, 2022. 1, 5
- [27] 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 *CVPR*, 2017. 1, 2, 3, 4, 5
- [28] Jianyang Gu, Weihua Chen, Hao Luo, Fan Wang, Hao Li, Wei Jiang, and Weijie Mao. Multi-view evolutionary training for unsupervised domain adaptive person re-identification. *IEEE Transactions on Information Forensics and Security*, 2022. 2
- [29] Jianyang Gu, Hao Luo, Weihua Chen, Yiqi Jiang, Yuqi Zhang, Shuting He, Fan Wang, Hao Li, and Wei Jiang. 1st place solution to visda-2020: Bias elimination for domain adaptive pedestrian re-identification. *arXiv preprint arXiv:2012.13498*, 2020. 2
- [30] Hao Guo, Kang Zheng, Xiaochuan Fan, Hongkai Yu, and Song Wang. Visual attention consistency under image

- transforms for multi-label image classification. In *CVPR*, 2019. 8
- [31] John A Hartigan and Manchek A Wong. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 1979. 4
- [32] Irtiza Hasan, Shengcai Liao, Jinpeng Li, Saad Ullah Akram, and Ling Shao. Generalizable pedestrian detection: The elephant in the room. In *CVPR*, 2021. 1, 2, 3, 5
- [33] Irtiza Hasan, Shengcai Liao, Jinpeng Li, Saad Ullah Akram, and Ling Shao. Pedestrian detection: Domain generalization, cnns, transformers and beyond. *arXiv preprint arXiv:2201.03176*, 2022. 8
- [34] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 1, 2, 4
- [35] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2
- [36] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15013–15022, 2021. 7, 8
- [37] Takashi Isobe, Dong Li, Lu Tian, Weihua Chen, Yi Shan, and Shengjin Wang. Towards discriminative representation learning for unsupervised person re-identification. In *ICCV*, 2021. 2
- [38] Jian Jia, Houjing Huang, Xiaotang Chen, and Kaiqi Huang. Rethinking of pedestrian attribute recognition: A reliable evaluation under zero-shot pedestrian identity setting. *arXiv preprint arXiv:2107.03576*, 2021. 5, 7, 8
- [39] Yiqi Jiang, Weihua Chen, Xiuyu Sun, Xiaoyu Shi, Fan Wang, and Hao Li. Exploring the quality of gan generated images for person re-identification. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4146–4155, 2021. 3
- [40] Xin Jin, Cuiling Lan, Wenjun Zeng, Guoqiang Wei, and Zhibo Chen. Semantics-aligned representation learning for person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020. 2
- [41] Qingming Leng, Mang Ye, and Qi Tian. A survey of open-world person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology*, 2019. 1, 4
- [42] Dangwei Li, Zhang Zhang, Xiaotang Chen, Haibin Ling, and Kaiqi Huang. A richly annotated dataset for pedestrian attribute recognition. *arXiv preprint arXiv:1603.07054*, 2016. 3
- [43] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Self-correction for human parsing. *PAMI*, 2020. 3, 4, 5, 7, 8
- [44] Yanjie Li, Shoukui Zhang, Zhicheng Wang, Sen Yang, Wankou Yang, Shu-Tao Xia, and Erjin Zhou. Tokenpose: Learning keypoint tokens for human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11313–11322, 2021. 8
- [45] Zhengjia Li and Duoqian Miao. Sequential end-to-end network for efficient person search. In *AAAI*, 2021. 7, 8
- [46] Xiaodan Liang, Ke Gong, Xiaohui Shen, and Liang Lin. Look into person: Joint body parsing & pose estimation network and a new benchmark. *PAMI*, 2018. 8
- [47] 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. 5
- [48] Chong Liu, Yuqi Zhang, Weihua Chen, Fan Wang, Hao Li, and Yi-Dong Shen. Adaptive matching strategy for multi-target multi-camera tracking. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022. 2
- [49] Chong Liu, Yuqi Zhang, Hao Luo, Jiasheng Tang, Weihua Chen, Xianzhe Xu, Fan Wang, Hao Li, and Yi-Dong Shen. City-scale multi-camera vehicle tracking guided by crossroad zones. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 2
- [50] Xinchun Liu, Meng Zhang, Wu Liu, Jingkuan Song, and Tao Mei. Braidnet: Braiding semantics and details for accurate human parsing. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 338–346, 2019. 2, 8
- [51] Xihui Liu, Haiyu Zhao, Maoqing Tian, Lu Sheng, Jing Shao, Shuai Yi, Junjie Yan, and Xiaogang Wang. Hydraplus-net: Attentive deep features for pedestrian analysis. In *ICCV*, 2017. 2, 3, 5
- [52] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 4, 5, 7
- [53] Angélique Loesch and Romaric Audigier. Describe me if you can! characterized instance-level human parsing. In *ICIP*, 2021. 1, 2, 3, 4
- [54] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10437–10446, 2020. 4
- [55] Hao Luo, Pichao Wang, Yi Xu, Feng Ding, Yanxin Zhou, Fan Wang, Hao Li, and Rong Jin. Self-supervised pre-training for transformer-based person re-identification. *arXiv preprint arXiv:2111.12084*, 2021. 1, 5
- [56] Zekun Luo, Zheng Fang, Sixiao Zheng, Yabiao Wang, and Yanwei Fu. Nms-loss: Learning with non-maximum suppression for crowded pedestrian detection. In *Proceedings of the 2021 International Conference on Multimedia Retrieval*, 2021. 8
- [57] Thomas Mauthner, Horst Possegger, Georg Waltner, and Horst Bischof. Encoding based saliency detection for videos and images. In *CVPR*, 2015. 3
- [58] Erik Murphy-Chutorian and Mohan Manubhai Trivedi. Head pose estimation in computer vision: A survey. *PAMI*, 2008. 1, 2, 3, 8
- [59] Xiao Pan, Hao Luo, Weihua Chen, Fan Wang, Hao Li, Wei Jiang, Jianming Zhang, Jianyang Gu, and Peike Li. Dynamic

- gradient reactivation for backward compatible person re-identification. *arXiv preprint arXiv:2207.05658*, 2022. 3
- [60] Tao Ruan, Ting Liu, Zilong Huang, Yunchao Wei, Shikui Wei, and Yao Zhao. Devil in the details: Towards accurate single and multiple human parsing. In *Proceedings of the AAAI conference on artificial intelligence*, 2019. 8
- [61] Nikolaos Sarafianos, Xiang Xu, and Ioannis A Kakadiaris. Deep imbalanced attribute classification using visual attention aggregation. In *ECCV*, 2018. 8
- [62] Yumin Suh, Jingdong Wang, Siyu Tang, Tao Mei, and Kyoung Mu Lee. Part-aligned bilinear representations for person re-identification. In *ECCV*, 2018. 4
- [63] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 2019. 8
- [64] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjing Wang. Beyond part models: Person retrieval with refined part pooling. In *ECCV*, 2018. 3
- [65] Zichang Tan, Yang Yang, Jun Wan, Guodong Guo, and Stan Z Li. Relation-aware pedestrian attribute recognition with graph convolutional networks. In *AAAI*, 2020. 8
- [66] Chufeng Tang, Lu Sheng, Zhaoxiang Zhang, and Xiaolin Hu. Improving pedestrian attribute recognition with weakly-supervised multi-scale attribute-specific localization. In *ICCV*, 2019. 8
- [67] Jiasheng Tang, Fei Du, Weihua Chen, Hao Luo, Fan Wang, and Hao Li SBT. A simple baseline with cascade association for robust multi-objects tracking. In *CVPR RVSU Workshop*, volume 4, 2021. 2
- [68] Cheng Wang, Bingpeng Ma, Hong Chang, Shiguang Shan, and Xilin Chen. Tcts: A task-consistent two-stage framework for person search. In *CVPR*, 2020. 8
- [69] Shanshan Wang, Lei Zhang, Weihua Chen, Fan Wang, and Hao Li. Refining pseudo labels for unsupervised domain adaptive re-identification. *Knowledge-Based Systems*, 2022. 3
- [70] Wenhao Wang. Adapted center and scale prediction: more stable and more accurate. *arXiv preprint arXiv:2002.09053*, 2020. 8
- [71] Xinlong Wang, Tete Xiao, Yuning Jiang, Shuai Shao, Jian Sun, and Chunhua Shen. Repulsion loss: Detecting pedestrians in a crowd. In *CVPR*, 2018. 8
- [72] Xiao Wang, Shaofei Zheng, Rui Yang, Aihua Zheng, Zhe Chen, Jin Tang, and Bin Luo. Pedestrian attribute recognition: A survey. *Pattern Recognition*, 2022. 1, 2, 3
- [73] Liu Wei, Liao Shengcai, Ren Weiqiang, Hu Weidong, and Yu Yanan. High-level semantic feature detection: A new perspective for pedestrian detection. In *CVPR*, 2019. 7, 8
- [74] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*, 2018. 5
- [75] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 466–481, 2018. 8
- [76] Tong Xiao, Shuang Li, Bochao Wang, Liang Lin, and Xiaogang Wang. Joint detection and identification feature learning for person search. In *CVPR*, 2017. 1, 2, 3, 5
- [77] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simsim: A simple framework for masked image modeling. In *CVPR*, 2022. 1, 2, 4
- [78] Ehsan Yaghoubi, Farhad Khezeli, Diana Borza, SV Kumar, João Neves, and Hugo Proença. Human attribute recognition—a comprehensive survey. *Applied Sciences*, 2020. 1, 2
- [79] Yichao Yan, Jinpeng Li, Jie Qin, Song Bai, Shengcai Liao, Li Liu, Fan Zhu, and Ling Shao. Anchor-free person search. In *CVPR*, 2021. 8
- [80] Yang Yang, Zhen Lei, Shifeng Zhang, Hailin Shi, and Stan Z Li. Metric embedded discriminative vocabulary learning for high-level person representation. In *AAAI*, 2016. 4
- [81] Zizheng Yang, Xin Jin, Kecheng Zheng, and Feng Zhao. Unleashing potential of unsupervised pre-training with intra-identity regularization for person re-identification. In *CVPR*, 2022. 1, 8
- [82] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. Deep learning for person re-identification: A survey and outlook. *PAMI*, 2021. 1, 2, 4
- [83] Yuhui Yuan, Rao Fu, Lang Huang, Weihong Lin, Chao Zhang, Xilin Chen, and Jingdong Wang. Hrformer: High-resolution transformer for dense prediction. In *NIPS*, 2021. 3, 5, 7, 8
- [84] Wang Zeng, Sheng Jin, Wentao Liu, Chen Qian, Ping Luo, Wanli Ouyang, and Xiaogang Wang. Not all tokens are equal: Human-centric visual analysis via token clustering transformer. In *CVPR*, 2022. 3
- [85] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv:1710.09412*, 2017. 5
- [86] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Citypersons: A diverse dataset for pedestrian detection. In *CVPR*, 2017. 3, 5
- [87] Xiaomei Zhang, Yingying Chen, Bingke Zhu, Jinqiao Wang, and Ming Tang. Part-aware context network for human parsing. In *CVPR*, 2020. 8
- [88] 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. 4
- [89] Yuqi Zhang, Qi Qian, Chong Liu, Weihua Chen, Fan Wang, Hao Li, and Rong Jin. Graph convolution for re-ranking in person re-identification. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022. 3
- [90] Ce Zheng, Wenhan Wu, Taojiannan Yang, Sijie Zhu, Chen Chen, Ruixu Liu, Ju Shen, Nasser Kehtarnavaz, and Mubarak Shah. Deep learning-based human pose estimation: A survey. *arXiv preprint arXiv:2012.13392*, 2020. 1, 2
- [91] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015. 5
- [92] Liang Zheng, Hengheng Zhang, Shaoyan Sun, Manmohan Chandraker, Yi Yang, and Qi Tian. Person re-identification in the wild. In *CVPR*, 2017. 1, 2, 3, 5

- [93] Peng Zheng, Jie Qin, Yichao Yan, Shengcai Liao, Bingbing Ni, Xiaogang Cheng, and Ling Shao. Global-local context network for person search. *arXiv preprint arXiv:2112.02500*, 2021. 8
- [94] Zhedong Zheng, Xiaodong Yang, Zhiding Yu, Liang Zheng, Yi Yang, and Jan Kautz. Joint discriminative and generative learning for person re-identification. In *CVPR*, 2019. 4
- [95] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021. 1, 2
- [96] Kuan Zhu, Haiyun Guo, Zhiwei Liu, Ming Tang, and Jinqiao Wang. Identity-guided human semantic parsing for person re-identification. In *ECCV*, 2020. 2, 3
- [97] Kuan Zhu, Haiyun Guo, Tianyi Yan, Yousong Zhu, Jinqiao Wang, and Ming Tang. Pass: Part-aware self-supervised pre-training for person re-identification. In *ECCV*. Springer, 2022. 1, 5, 8