

MultiSiam: Self-supervised Multi-instance Siamese Representation Learning for Autonomous Driving

Kai Chen¹ Lanqing Hong² Hang Xu² Zhenguo Li² Dit-Yan Yeung¹
¹Hong Kong University of Science and Technology ²Huawei Noah's Ark Lab

kai.chen@connect.ust.hk {honglanqing, xu.hang, li.zhenguo}@huawei.com dyyeung@cse.ust.hk

Abstract

Autonomous driving has attracted much attention over the years but turns out to be harder than expected, probably due to the difficulty of labeled data collection for model training. Self-supervised learning (SSL), which leverages unlabeled data only for representation learning, might be a promising way to improve model performance. Existing SSL methods, however, usually rely on the single-centric-object guarantee, which may not be applicable for multi-instance datasets such as street scenes. To alleviate this limitation, we raise two issues to solve: (1) how to define positive samples for cross-view consistency and (2) how to measure similarity in multi-instance circumstances. We first adopt an IoU threshold during random cropping to transfer global-inconsistency to local-consistency. Then, we propose two feature alignment methods to enable 2D feature maps for multi-instance similarity measurement. Additionally, we adopt intra-image clustering with self-attention for further mining intra-image similarity and translation-invariance. Experiments show that, when pre-trained on Waymo dataset, our method called **Multi-instance Siamese Network (MultiSiam)** remarkably improves generalization ability and achieves state-of-the-art transfer performance on autonomous driving benchmarks, including Cityscapes and BDD100K, while existing SSL counterparts like MoCo, MoCo-v2, and BYOL show significant performance drop. By pre-training on SODA10M, a large-scale autonomous driving dataset, MultiSiam exceeds the ImageNet pre-trained MoCo-v2, demonstrating the potential of domain-specific pre-training. Code will be available at <https://github.com/KaiChen1998/MultiSiam>.

1. Introduction

Autonomous driving has attracted much attention over the years [19, 15, 32]. However, it becomes harder than people have expected in such an age of AI advances. Fully autonomous cars are still out of reach except in special trial programs, mainly due to the limitation of model per-




Figure 1. **Visualizations of different random crops on Waymo [20] (left) and ImageNet [9] (right).** On ImageNet, images are small and pre-processed to guarantee only one object in the center part of it (*i.e.*, single-centric-object) mostly. However, images from Waymo are of high resolution and contain multiple instances. Different random crops might represent different semantic meanings (*i.e.*, global-inconsistency), which will restrict the effectiveness of current self-supervised learning methods.

formance. One of the main restrictions is that the annotation cost of self-driving datasets is much more expensive than other datasets. Considering that autonomous cars keep collecting unlabeled data when operating, self-supervised learning (SSL) might be a promising way to ease the desire for labeled data and improve model performance, which has achieved remarkable transfer results on different downstream tasks using unlabeled data only.

Existing SSL methods are mainly based on the pretext called instance discrimination and cross-view consistency framework, whose basic assumption is that different views (*e.g.* data augmentation) of a single image should be consistent in the feature space under different metrics, such as cosine distance [10, 7], clustering assignments [4] and discriminability from negative samples [27, 21, 5, 12, 6, 18]. This assumption is satisfied well with *single-centric-object* datasets such as ImageNet. In self-driving, nevertheless, the

data are usually high-resolution images containing multiple instances on a single image (see Figure 1 for an illustration). Here we define *instance* as any individual object regardless of its semantic class following Wu *et al.* [27]. In this case, different crops may correspond to different instances and represent different semantic meanings, which results in the *global-inconsistency* of multi-instance images. The effectiveness of instance discrimination and cross-view consistency can no longer be guaranteed.

To adapt current instance discrimination and cross-view consistency framework to multi-instance circumstances, we need to solve two problems: (1) how to define positive samples for cross-view consistency and (2) how to calculate the similarity of two randomly generated views within multi-instance images. Considering the locality of images, we first add an IoU threshold during random cropping as a proxy to control the two views not too far from each other and transfer global-inconsistency to local-consistency. In order to distinguish different instances, we maintain the final 2D feature maps of the backbone networks and propose the *RoI alignment* and *offset alignment* to solve the feature misalignment introduced by it (see Figure 4(a)), which is usually neglected when global pooling layers are adopted. Moreover, we observe a hierarchy of clusters existing in multi-instance circumstances naturally, so to model the relationships between instances, we perform clustering within a single image not for inter-image similarity [2, 4], but for further mining intra-image similarity. To ease the ambiguity of cluster assignments, we deploy a self-attention mechanism with the predictor for more precise cluster prediction. Translation-invariance is also enhanced in the learned representation, which is beneficial for downstream pixel-level visual tasks like semantic segmentation.

The main contributions of this work contain three parts:

1. We propose the Multi-instance Siamese Network (*MultiSiam*) for extending the cross-view consistency framework to multi-instance circumstances by dealing with positive sample definition and similarity measurement of 2D feature maps.
2. Experiments on Cityscapes [8] and BDD100K [31] show that *MultiSiam* pre-trained on Waymo [20] has stronger generalization ability to multi-instance datasets and achieves state-of-the-art transfer performance on downstream autonomous driving benchmarks compared with MoCo, MoCo-v2 and BYOL. Moreover, *MultiSiam* pre-trained on SODA10M [11], a large-scale autonomous driving dataset, exceeds the ImageNet pre-trained MoCo-v2, revealing the potential of domain-specific pre-training.
3. To the best of our knowledge, our work is the first to perform self-supervised learning on large-scale high-resolution multi-instance street scene datasets (*e.g.*

Waymo), which will be beneficial for empowering SSL in further autonomous driving research.

2. Related Work

Contrastive learning. Contrastive learning is widely used in self-supervised representation learning, which has shown promising performance for various tasks [27, 21, 5, 12, 6, 18, 22]. The main idea is to consider every single image as a separate class and train the model to pull positive sample pairs closer while pushing negative sample pairs away using the InfoNCE loss [23]. MoCo [12, 6] proposes to maintain a queue of negative samples with a Siamese network and shift one branch of the network into a momentum encoder to improve the consistency of the queue. SimCLR [5], however, directly uses negative samples co-existing in the current batch. Contrastive learning requires comparing each positive with many other negative samples to work well, which is usually achieved by crops from a large amount of single-centric-object images. In multi-instance data such as street scenes, however, two random crops from one single image may correspond to different instances, leading to the possible ineffectiveness of existing contrastive methods.

Clustering. Clustering is another popular paradigm for unsupervised representation learning [2, 3, 1, 4]. The main idea is to consider each cluster as a separate class instead of each single image as contrastive learning does to capture the similarity between images by clustering the representations and learning to predict the cluster assignments alternatively. Deep Cluster [2] iteratively clusters all the images based on current representations and treats cluster indexes as pseudo labels to train a classifier from scratch. Although there is no need for negative samples, a costly offline clustering process is introduced. SwAV [4] conducts online clustering with a Siamese network by computing the assignment from one branch and predicting it from the other under a balanced partition constraint for each batch. Different from previous methods, in the proposed *MultiSiam*, we perform clustering within a single image instead of the whole dataset to further discover intra-image similarity in multi-instance circumstances.

Cosine similarity. Both negative samples and clustering are considered useful to prevent model collapse in self-supervised learning. BYOL [10] instead proposes a cross-view consistency framework with cosine similarity only that can produce meaningful representations with the help of a momentum target network. SimSiam [7] further points out that the stop-gradient operation is the critical component to prevent model collapse based on an EM-like hypothesis. Due to its simplicity and effectiveness, we adopt BYOL as our baseline model and show that intra-image clustering is a better similarity metric in multi-instance circumstances.

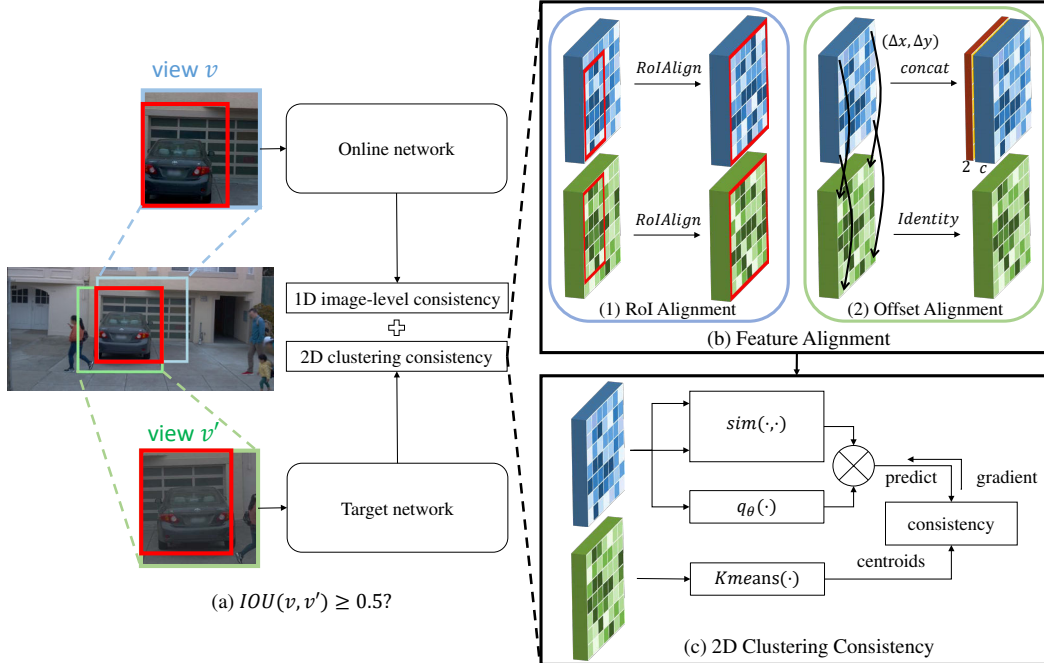


Figure 2. **Model architecture of MultiSiam.** (a) Two views v and v' are fed for further calculation only when their IoU value is larger than a given threshold. (b) We maintain the 2D feature maps and propose two methods for feature alignment after projection and flipping back (more details in Figure 4). *RoI alignment* extracts features of the intersection region only, while *offset alignment* provides the coordinate offset map to the predictor for implicit feature alignment. (c) Intra-image clustering is performed on the aligned target feature map and the online network needs to predict the cluster centroids. We deploy a self-attention mechanism to deal with the ambiguity of clustering.

3. Method

In this section, we will first introduce the BYOL [10] in Section 3.1, which can achieve state-of-the-art transfer performance without the need for either negative samples or clustering. Due to its effectiveness and simplicity, we choose BYOL as our baseline model. Then we extend the cross-view consistency framework to multi-instance circumstances by discussing two problems: how to define positive samples and how to measure similarity in multi-instance images and propose our final *MultiSiam* model in Sections 3.2 and 3.3, respectively.

3.1. Preliminary: BYOL

Given a single image x , BYOL learns the representation by maximizing the similarity of two random views v and v' in the feature space. BYOL first randomly generates two views $v \sim \mathcal{T}(x)$ and $v' \sim \mathcal{T}'(x)$, which are then fed into the online and target networks separately. The online network consists of a backbone with a global pooling layer, an MLP projector, and an MLP predictor, while the target network shares the same architecture except for the final predictor. To prevent model collapse, BYOL adopts a stop-gradient operation on top of the target network. Finally, BYOL minimizes the cosine distance between the online

network's prediction q and the target network's projected feature z' as the 1D image-level consistency by:

$$\mathcal{L}_{1D_img} \triangleq -\cos(q, z') = -\frac{\langle q, z' \rangle}{\|q\|_2 \cdot \|z'\|_2}, \quad (1)$$

where $\cos(\cdot, \cdot)$ is the cosine similarity and q and z' are both 1D feature vectors.

To further improve performance, BYOL updates the parameters of the target network ξ as the exponential moving average of the online network's parameters θ as:

$$\xi \leftarrow \tau \xi + (1 - \tau) \theta, \quad (2)$$

where $\tau \in [0, 1]$ is the momentum and will increase to 1.0 until the end of training.

3.2. Positive Samples in Multi-instance Data

The basic assumption of instance discrimination is that different views of the same image should be consistent in the feature space. However, this strong assumption is only satisfied well on ImageNet [9] with single-centric-object guarantee but not scalable to more realistic datasets with multiple instances on a single image like Waymo [20], as shown in Figure 1. A new definition of positive samples is definitely needed to extend cross-view consistency framework to multi-instance circumstances.




Figure 3. **Random crops with different Intersection-over-Union values on Waymo (each column).** (a) Original image; (b) $IoU = 0.3$; (c) $IoU = 0.5$; (d) $IoU = 0.7$. When $IoU = 0.3$, the two random crops suffer from global-inconsistency. As the IoU value raises, the two random views are restricted within a local region, and the local-consistency increases. When $IoU = 0.7$, the two views seem nearly the same. As shown in Section 4.3, IoU controls a trade-off between noise and data complexity during data augmentation.

Intersection-over-Union as a proxy. Images are continuous natural signals and have strong locality. One of the main reasons why inconsistency happens in multi-instance cases is that the two random crops may be far away from each other. Considering the locality of images, if the two views are “close” enough, it is reasonable to assume that they represent the same semantic meanings in a local region and transfer global-inconsistency to local-consistency.

Here we propose to use the Intersection-over-Union $IoU(v, v')$ as a proxy about how “close” the two random crops are and set an additional IoU threshold during data augmentation. Two views are used for further calculation only when their $IoU(v, v')$ is larger than the pre-defined threshold. We show how different IoU values will affect the consistency of the two random crops in Figure 3. When $IoU(v, v') = 0.3$, we can still see the car in the bottom crop but the top crop has been mainly occupied by the garage. As the $IoU(v, v')$ rises gradually, the semantic meanings of the two crops align better and when $IoU(v, v') = 0.7$, the two views seem nearly the same.

In Section 4.3, we will show that using different IoU thresholds is actually a trade-off between noise and data complexity during data augmentation. If not specified, we set the IoU threshold to be 0.5 in all experiments by default.

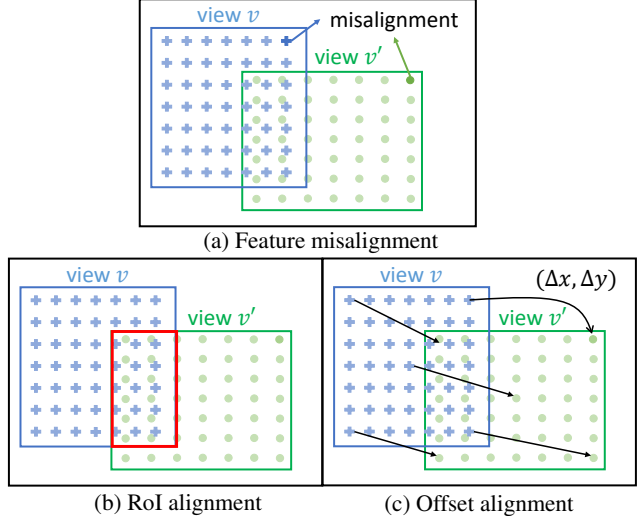


Figure 4. **Feature misalignment** of two random views (*i.e.*, blue and green boxes) in a local region (*i.e.*, the black box). (a) Feature misalignment: top right corners of v and v' represent different pixels, so one-to-one correspondence no longer exists; (b) RoI alignment: use RoI Align [13] to extract features of overlapping regions (*i.e.*, the red box) only; (c) Offset alignment: calculate the coordinate offset of each pixel pair at the same relative position to get the offset map ΔC , which is then concatenated with the projected online 2D feature map $G = g_\theta(F)$ and fed into the predictor.

Consistency learning via 2D feature map. Even with the IoU threshold, there might be multiple instances in a single crop, as shown in Figure 3(d). State-of-the-art methods all deploy a global pooling layer at the end of the backbones to produce a 1D feature vector, which will lose the spatial and structural information of 2D feature space, and the model cannot distinguish different instances any more.

As a result, we discard the global pooling layers and maintain the final 2D feature maps $F, F' \in \mathcal{R}^{H \times W \times C}$ of both online and target backbone networks. Meanwhile, we replace the MLPs in the projectors and predictor with 1×1 convolution layers in order to match the 2D structure while keeping the amount of parameters unchanged.

Feature misalignment. However, there is no free lunch. Using 2D feature maps introduces another feature misalignment problem that has not been taken into account when global pooling layers are available, as shown in Figure 4(a). The top right corners of v and v' lie in different positions. The one-to-one correspondence (*e.g.*, the top right corner of v corresponds to the top right corner of v') no longer exists because we apply different spatial augmentations to the two views. To retrieve the correspondence, we propose two alignment methods: *RoI alignment* and *offset alignment*.

(1) RoI alignment. Since the IoU threshold guarantees a non-trivial overlapping between v and v' , *RoI alignment* treats the overlapping area as the Region of Interest for both views and uses RoI Align [13] to extract the feature of inter-

section region only. We denote the *relative* box coordinates of the overlapping in two views as $B, B' \in \mathcal{R}^4$, then the aligned feature $R, R' \in \mathcal{R}^{H \times W \times C}$ can be represented as:

$$R = \text{RoIAlign}(g_\theta(F), B), \quad (3)$$

$$R' = \text{RoIAlign}(g_\xi(F'), B'), \quad (4)$$

where $g_\theta(\cdot)$ and $g_\xi(\cdot)$ are the projectors of online and target networks. If not specified, we keep the spatial resolution of R and R' the same as F and F' by default.

(2) Offset alignment. Although guaranteeing precise feature alignment, *RoI alignment* does not make full use of the information of non-overlapping areas. Motivated by Liu *et al.* [16], we provide the coordinate offset map $\Delta C \in \mathcal{R}^{H \times W \times 2}$ from the projected online 2D feature map $G = g_\theta(F)$ to the projected target feature map $G' = g_\xi(F')$ as additional information to the predictor for implicit feature alignment, as shown in Figure 4(c). Specifically, for $\forall i \in [1, H], j \in [1, W]$, we define the offset map ΔC as:

$$\Delta C_{i,j} \triangleq \frac{\text{coord}(G'_{i,j}) - \text{coord}(G_{i,j})}{\text{coord}(G_{H,W}) - \text{coord}(G_{1,1})}, \quad (5)$$

where $\text{coord}(\cdot)$ returns the corresponding coordinates of the given pixel in the original image. Note that ΔC should be normalized by the size of v to decrease variance because the predictor with *offset alignment* is actually estimating a *conditional expectation* (see more details in Section 4.3 and Appendix A). Then R and R' can be denoted as:

$$R = \text{concat}(g_\theta(F), \Delta C), \quad (6)$$

$$R' = g_\xi(F'). \quad (7)$$

After feature alignment, the model can recover the one-to-one correspondence between $R_{i,j}$ and $R'_{i,j}$ for $\forall i \in [1, H], j \in [1, W]$.

3.3. Similarity Measurement of 2D Feature Maps

With 2D feature maps, we can model the relationships between different instances. We find there exists a hierarchy of clusters in multi-instance circumstances naturally from *class cluster* down to *instance cluster* and *pixel cluster*, as shown in Figure 5(a). *Instance cluster* means that pixels belonging to the same instance should lie in the same cluster, which is consistent with the basic assumption of instance discrimination. On the other hand, *class cluster* suggests pixels belonging to different instances of the same semantic class should also form a separate cluster. Besides, *pixel cluster* refers to groups of pixels sharing common characteristics like pixel intensity (*e.g.*, super-pixels). This hierarchy is more common in autonomous driving datasets (*e.g.*, Waymo [20]) whose semantic class set is relatively small.

Intra-image clustering. Clustering is previously used in SSL mainly for mining *inter-image similarity* [2, 4] to reduce false negative during contrastive learning. However,

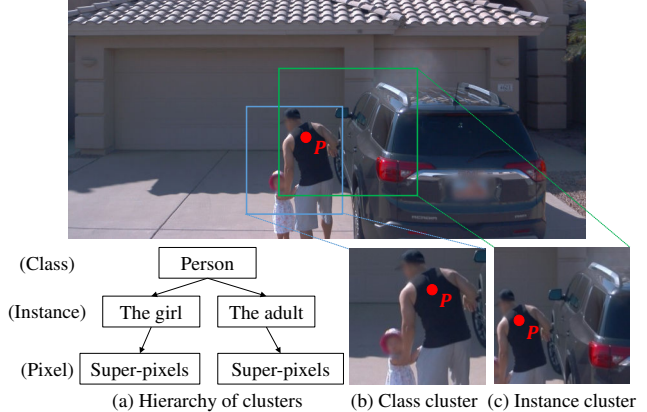


Figure 5. **The hierarchy of clusters** on Waymo (a) from class cluster down to instance and pixel clusters; (b) *class cluster*: pixels belonging to different instances of the same semantic class; (c) *instance cluster*: pixels belonging to the same instance. Cluster is actually a relative concept. The same pixel P might belong to different clusters in different contexts of (b) and (c).

based on the analysis above, we think clustering can also be performed within a single image to capture *intra-image similarity* further.

Specifically, we deploy the K-means algorithm on the aligned target feature map R' to get the assigned cluster **centroid** for each pixel, denoted as $Kmeans(R')$. The aligned online feature map R is then fed into the predictor to predict the cluster assignment $q_\theta(R)$, which should be consistent with $Kmeans(R')$ for every pixel due to the one-to-one correspondence after feature alignment. The network will extract smoother features for pixels in the same cluster so that translation-invariance is also encouraged.

Self-attention. As shown in Figure 5 (b) and (c), cluster is actually a relative concept. The same pixel P might belong to different clusters under different contexts. Here we consider the *Person* class cluster and the *Adult* instance cluster as different clusters since they have different centroids even if ideally *Adult* might be a subset of *Person*. However, the predictor $q_\theta(\cdot)$ only consists of 1×1 convolutions which operate locally. A “global” view is needed to deal with the ambiguity of cluster assignments. Here we deploy the same non-local-network style [24] self-attention module as Xie *et al.* [29]. Specifically, for $\forall i \in [1, H], j \in [1, W]$, the final cluster prediction $Q_{i,j}$ of $R_{i,j}$ is defined as:

$$Q_{i,j} = \sum_{i'=1}^H \sum_{j'=1}^W \text{sim}(R_{i,j}, R_{i',j'}) \cdot q_\theta(R_{i',j'}), \quad (8)$$

where $q_\theta(\cdot)$ is the original local predictor and $\text{sim}(\cdot, \cdot)$ is the similarity function defined as:

$$\text{sim}(R_{i,j}, R_{i',j'}) = (\max(\cos(R_{i,j}, R_{i',j'}), 0))^2. \quad (9)$$

Method	Pre-train Dataset	Epochs	Cityscapes			BDD100K		
			mAP	AP ₅₀	mIoU	mAP	AP ₅₀	mIoU
Random Init	-	-	25.4	51.1	65.3	16.4	30.4	50.7
Supervised	ImageNet	90	32.9	59.6	74.6	21.9	40.0	58.8
InstDist [†] [27]	ImageNet	200	33.0	60.1	73.3	21.4	38.9	57.2
SwAV [†] [4]	ImageNet	200	33.9	62.4	73.0	22.5	40.8	57.1
BYOL [†] [10]	ImageNet	200	33.8	62.9	75.1	21.8	39.3	59.1
MoCo [†] [12]	ImageNet	200	32.3	59.3	75.3	22.4	40.4	59.7
MoCo-v2 [†] [6]	ImageNet	200	33.9	60.8	75.7	23.1	41.3	60.0
BYOL [10]	Waymo	325	28.8	55.7	69.4	18.1	33.8	53.7
MoCo [12]	Waymo	325	30.5	57.1	73.9	21.0	39.1	57.0
MoCo-v2 [6]	Waymo	325	31.4	59.4	73.6	20.9	38.9	56.6
MultiSiam	Waymo	325	31.8 ^{+3.0}	59.6 ^{+3.9}	74.1 ^{+4.7}	21.1 ^{+3.0}	39.3 ^{+5.5}	57.6 ^{+3.9}
MultiSiam ^{††}	Waymo	325	32.2 ^{+3.4}	59.9 ^{+4.2}	75.5 ^{+6.1}	21.8 ^{+3.7}	40.2 ^{+6.4}	56.9 ^{+3.2}
MultiSiam	SODA5M	55	34.1 ^{+5.3}	61.7 ^{+6.0}	75.8 ^{+6.4}	23.5 ^{+5.4}	42.7 ^{+8.9}	60.3 ^{+6.6}

Table 1. **Comparisons on Cityscapes and BDD100K instance and semantic segmentation.** The metrics include mask mAP and AP50 for instance segmentation and mIoU for semantic segmentation. (1) BYOL, MoCo and MoCo-v2 all suffer from global-inconsistency and performance degradation from ImageNet to Waymo, (2) but *MultiSiam* directly based on BYOL recovers its decrease and shows better generalization ability, achieving state-of-the-art performance among Waymo pre-trainings. (3) By pre-training on SODA5M, *MultiSiam* exceeds the ImageNet pre-trained MoCo-v2, revealing the potential of domain-specific pre-training. [†]: we take the officially released pre-trained weights and report fine-tuning results. ^{††}: a simple implementation of MoCo-based *MultiSiam*.

The final 2D clustering consistency loss is defined as:

$$\mathcal{L}_{2D_cluster} \triangleq \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W -\cos(Q_{i,j}, Kmeans(R'_{i,j})). \quad (10)$$

Incorporated with 1D image-level consistency. We also keep the 1D image-level consistency branch as defined in Section 3.1 due to its effectiveness to improve classification performance. Both classification and localization are important for visual perception tasks like semantic segmentation. We formulate *MultiSiam* in the multi-task learning manner and the final loss function is a weighted sum of 1D image-level consistency and 2D clustering consistency as:

$$\mathcal{L}_{MultiSiam} = \lambda \mathcal{L}_{1D_img} + (1 - \lambda) \mathcal{L}_{2D_cluster}, \quad (11)$$

where the balanced weight λ is set to be 0.5 by default.

4. Experiments

4.1. Implementation Details

Dataset. We pre-train our *MultiSiam* on the widely used Waymo Open [20] autonomous driving dataset mainly, which consists of around 790 thousand training images. The image sizes range from (1920, 968) to (1920, 1280).

Data augmentation. We follow the standard data augmentation pipeline in BYOL [10] with the proposed IoU threshold. Two random crops whose IoU is larger than the pre-defined threshold (0.5 by default) are generated and then resized to 224×224 , followed by random horizontal flip, color jitter, Gaussian blur and solarization. We flip

back the projected 2D feature map before feature alignment if the horizontal flip is applied previously. All the augmentation parameters are kept the same with BYOL.

Training details. We adopt standard ResNet-50 [14] as the backbone network. The momentum starts at 0.996 and increases gradually to 1.0 at the end of the training process. We use the LARS [30] optimizer with a cosine learning rate scheduler for large batch training. The base learning rate is set to 1.0, which will scale linearly with the batch size ($lr = lr_{base} \times bs/256$). Weight decay is set to $1e-5$. We use a batch size of 1024 running on 8 Tesla V100 GPUs per experiment. We pre-train 325 epochs on Waymo to maintain similar training iterations with training 200 epochs on ImageNet for a fair comparison. For ablation studies, we adopt 150 epochs pre-training and then evaluate on Cityscapes [8] val set for semantic segmentation.

4.2. Transfer Settings and Results

Baseline. In this paper, we choose the strong BYOL [10], MoCo [12] and MoCo-v2 [6] as our baseline methods. We use OpenSelfSup¹ as our codebase to pre-train all baseline methods on Waymo and tune the hyperparameters according to the transfer results on Cityscapes val set of 150 epochs pre-training for better performance.

Transfer settings. We choose widely used Cityscapes [8] and BDD100K [31] instance and semantic segmentation for autonomous driving as downstream tasks, adapting the settings of Cityscapes in MoCo [12] for both datasets using Detectron2 [26]. For instance segmentation, we fine-tune a

¹<https://github.com/open-mmlab/OpenSelfSup>

IoU Thre	Feature Alignment			mIoU
	RoI	Offset	Offset nonorm	
(a) IoU threshold				
0.3	✓			69.2
0.4	✓			68.6
0.5	✓			70.0
0.6	✓			69.7
0.7	✓			69.2
(b) Feature alignment methods				
0.5				68.8
0.5	✓			70.0
0.5		✓		69.5
0.5			✓	69.4

Table 2. **Ablations on positive sample definition.** (a) IoU threshold; (b) feature alignment. All results are evaluated on Cityscapes val set over three independent trials.

Mask R-CNN detector (FPN-backbone) for 24k iterations, while a FCN-16s [17] is trained for 90k iterations for semantic segmentation. We train the models on both train sets and evaluate on the corresponding val sets separately.

Discussion. We report the final transfer results in Table 1. All baseline methods suffer from significant performance degradation adapted from ImageNet to Waymo, revealing the vulnerability to global-inconsistency of current models. Our *MultiSiam*, however, recovers the decrease and generalizes better to multi-instance circumstances. Compared with our direct BYOL baseline, *MultiSiam* improves **3.0%/3.0%** mAP and **4.7%/3.9%** mIoU for instance and semantic segmentation separately, surpassing the strong MoCo baselines and achieving state-of-the-art performance among Waymo pre-trainings. To verify the guidance effect of clustering, we also train *MultiSiam without K-means* (more details in Appendix B.1) and achieve 71.0% and 54.8% mIoU on Cityscapes and BDD100K, which is significantly worse than *MultiSiam*, suggesting that K-means might help produce more robust targets during self-supervised learning.

Domain-specific pre-training. Although being a self-driving dataset, Waymo still has disadvantages in both quantity (0.79M vs 1.28M) and quality (e.g., imbalance of foreground and background) compared with ImageNet, which may hurt the performance of Waymo representations. We further pre-train *MultiSiam* on SODA10M [11], a large-scale autonomous driving dataset, with similar GPU hours of ImageNet pre-trainings and surpass the ImageNet pre-trained MoCo-v2 as shown in Table 1, revealing the potential of domain-specific pre-training. Due to the hardware resources, here we only use a 5-million subset of SODA10M (split 0, 2, 4, 6, 8), denoted as SODA5M. Note that due to single-centric-object guarantee, it is costly to collect a large amount of ImageNet-like samples, making multi-instance self-supervised learning more valuable in practice.

#Cluster K	Dist metrics	Dense	mIoU
(a) Number of Clusters K			
3	Cosine		70.0
4	Cosine		68.4
5	Cosine		69.7
(b) Dense Clustering & Distance Metrics			
3	Cosine		70.0
3	Cosine	✓	69.0
3	Euclidean		69.3

Table 3. **Ablations on intra-image clustering.** (a) Cluster number K; (b) dense clustering and clustering distance metrics.

Method	Components	mIoU	Δ
BYOL		67.2	
MultiSiam	+ IoU thre 0.5	69.2	+ 2.0
	+ $\mathcal{L}_{cluster}$	70.0	+ 0.8
	+ self-attention	71.2	+ 1.2
	+ offset align	71.9	+ 0.7
	/+ offset align nonorm	71.2	+ 0.0

Table 4. **Ablations on proposed components.** The model’s performance continues to increase after each component is deployed.

Flexibility. As shown in Table 1, BYOL has a significant performance gap with MoCo because pure cross-view consistency framework might be more vulnerable to global-inconsistency than contrastive methods, which is complementary to the improvement of *MultiSiam*. As a flexible plug-and-play module, *MultiSiam* can be naturally extended to contrastive learning based SSL methods. Here we deploy a simple implementation of MoCo-based *MultiSiam* (more details in Appendix B.2) and obtain further improvements (e.g., 0.4% mAP and 1.4% mIoU on Cityscapes in Table 1). More delicate design is definitely needed to achieve the best performance, which we will dig into in future work.

Generalization. Although originally designed for multi-instance circumstances, our *MultiSiam* also demonstrates remarkable generalization ability to single-centric-object datasets. As shown in Appendix C, ImageNet-pretrained *MultiSiam* still achieves state-of-the-art performance even when compared with more recent SSL methods [25, 28].

4.3. Ablation Study and Analysis

Setup. We conduct 150 epochs pre-training with a base learning of 0.3 for all ablation studies. The batch size is set to 1024, and parameters are updated every four iterations to mimic a batch size of 4096 following the BYOL [10]. We report the average result of three independent trials to reduce variance. See more ablations in Appendix D.

IoU threshold. As shown in Table 4, BYOL simply equipped with the proposed IoU threshold during random cropping can achieve a significant improvement of 2.0% mIoU, showing the effectiveness of local-consistency. We further find that there is a balanced point when the IoU

threshold is set to 0.5 in Table 2(a). Interestingly, the model achieves sup-optimal results when the threshold is set to 0.7, which can generate well-aligned random crops in Figure 3(d). We think it is because by controlling how “close” the two crops are, the IoU threshold also controls a trade-off between **noise** and **data complexity**. Remember that the motivation of data augmentation is to generate two **different** but **consistent** views. We do not want to differ the two views too much since it will violate the basic assumption of instance discrimination. Meanwhile, the two views should not be exactly the same. Otherwise, the network would tend to collapse due to the trivial supervision signal.

Feature alignment. We ablate the proposed two feature alignment methods in Table 2(b). The model suffers from a performance degradation of 1.2% mIoU without feature alignment due to the absence of one-to-one correspondence. Both *RoI* and *offset alignment* achieve significant improvements and as we can see in Table 4, *offset alignment* performs much better when self-attention is available.

We also verify the necessity of coordinate normalization during the offset alignment in Equation (5). As shown in Table 2(b) and 4, offset alignment with coordinate normalization performs better whether or not self-attention is available. We analyze effectiveness of offset alignment based on the EM-like hypothesis proposed by Chen *et al.* [7]. Specifically, predictor with offset alignment is actually estimating an expectation of target cluster centroids conditioned on the online feature and offset map (see proofs in Appendix A):

$$q_{\theta}^{optimal}(R, \Delta C) = \mathbb{E}[Kmeans(R')|\Delta C, R]. \quad (12)$$

So to reduce variance during training, using normalized offset maps should be a better choice than the absolute values.

Number of clusters K. We ablate the number of clusters K during the K-means algorithm in Table 3(a). The optimal performance is obtained when $K = 3$. One of the reasons is that self-driving datasets like Waymo usually have a relatively small semantic class set (*e.g.* size 4 for Waymo). Besides, as we discuss in Section 3.3, the definition of clusters in multi-instance circumstances is actually a relative concept dependent on the specific context of a random crop, which is not limited to *instance* or *class* only. Keeping the number of clusters K at a relatively small value will make the model more robust to different image contexts.

Dense clustering. Instead of predicting cluster centroids, dense clustering calculates cosine similarity with every pixel in the corresponding cluster. As shown in Table 3(b), predicting centroids performs better because as discussed in Equation (12), the predictor estimates a conditional expectation while dense clustering will increase variance.

Self-attention. As we can see in Table 4, using the local predictor can only achieve a marginal improvement of 0.8% mIoU upon BYOL with 0.5 IoU threshold because of the



Figure 6. **Clustering ($K = 3$) results of BYOL and MultiSiam.** Different colors represent different clusters. *MultiSiam* can better capture intra-image similarity in different circumstances.

ambiguity of cluster assignments in Section 3.3. After introducing a “global” view by deploying self-attention with the predictor, our model can get more precise clustering predictions and achieve another significant improvement. Note that offset alignment performs much better together with self-attention, because introducing relative offset can make better use of the image locality for clustering prediction.

4.4. Visualization

Figure 6 shows the clustering results on the final backbone 2D feature maps of BYOL and *MultiSiam* in single instance, multi-instance of the same class and multi-instance of different classes circumstances. BYOL shows random clustering pattern while *MultiSiam* can better capture intra-similarity. It recognizes the car in the first row, separates the background crowds from the man in the second row and distinguishes the lady and the car successfully in the third row. The clustering result of *MultiSiam* is much smoother, suggesting that translation-invariance is also enhanced.

5. Conclusion

This paper explores the usage of SSL in general multi-instance autonomous driving circumstances. We show that current methods rely on the single-centric-object ImageNet suffering from performance degradation on multi-instance datasets like Waymo for pre-training. Targeting this issue, our proposed *MultiSiam* demonstrates better generalization ability and achieves state-of-the-art transfer performance by dealing with positive sample definition and adopting intra-image clustering with self-attention. We believe the information in multi-instance data is still under-explored (*e.g.*, long distance similarity). We hope our simple yet effective method can bring researchers’ attention to multi-instance self-supervised learning.

References

- [1] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. *arXiv:1911.05371*, 2019. [2](#)
- [2] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. [2](#), [5](#)
- [3] Mathilde Caron, Piotr Bojanowski, Julien Mairal, and Armand Joulin. Unsupervised pre-training of image features on non-curated data. In *ICCV*, 2019. [2](#)
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv:2006.09882*, 2020. [1](#), [2](#), [5](#), [6](#)
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv:2002.05709*, 2020. [1](#), [2](#)
- [6] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv:2003.04297*, 2020. [1](#), [2](#), [6](#)
- [7] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. *arXiv:2011.10566*, 2020. [1](#), [2](#), [8](#)
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. [2](#), [6](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. [1](#), [3](#)
- [10] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, and others. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv:2006.07733*, 2020. [1](#), [2](#), [3](#), [6](#), [7](#)
- [11] Jianhua Han, Xiwen Liang, Hang Xu, Kai Chen, Lanqing Hong, Chaoqiang Ye, Wei Zhang, Zhenguo Li, Xiaodan Liang, and Chunjing Xu. Soda10m: Towards large-scale object detection benchmark for autonomous driving. *arXiv:2106.11118*, 2021. [2](#), [7](#)
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. [1](#), [2](#), [6](#)
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. [4](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [6](#)
- [15] Joel Janai, Fatma Güney, Aseem Behl, Andreas Geiger, et al. Computer vision for autonomous vehicles: Problems, datasets and state of the art. In *Foundations and Trends® in Computer Graphics and Vision*, 2020. [1](#)
- [16] Rosanne Liu, Joel Lehman, Piero Molino, Felipe Petroski Such, Eric Frank, Alex Sergeev, and Jason Yosinski. An intriguing failing of convolutional neural networks and the coordconv solution. *arXiv:1807.03247*, 2018. [5](#)
- [17] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. [7](#)
- [18] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *CVPR*, 2020. [1](#), [2](#)
- [19] Scott Drew Pendleton, Hans Andersen, Xinxin Du, Xiaotong Shen, Malika Meghiani, You Hong Eng, Daniela Rus, and Marcelo H Ang. Perception, planning, control, and coordination for autonomous vehicles. In *Machines*, 2017. [1](#)
- [20] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, and others. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2020. [1](#), [2](#), [3](#), [5](#), [6](#)
- [21] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv:1906.05849*, 2019. [1](#), [2](#)
- [22] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning. *arXiv:2005.10243*, 2020. [2](#)
- [23] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018. [2](#)
- [24] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. [5](#)
- [25] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. *arXiv:2011.09157*, 2020. [7](#)
- [26] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. 2019. [6](#)
- [27] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. [1](#), [2](#), [6](#)
- [28] Enze Xie, Jian Ding, Wenhai Wang, Xiaohang Zhan, Hang Xu, Zhenguo Li, and Ping Luo. Detco: Unsupervised contrastive learning for object detection. *arXiv:2102.04803*, 2021. [7](#)
- [29] Zhenda Xie, Yutong Lin, Zheng Zhang, Yue Cao, Stephen Lin, and Han Hu. Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. *arXiv:2011.10043*, 2020. [5](#)
- [30] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv:1708.03888*, 2017. [6](#)
- [31] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv:1805.04687*, 2018. [2](#), [6](#)
- [32] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A survey of autonomous driving: Common practices and emerging technologies. In *IEEE Access*, 2020. [1](#)