

# Learning to Drive by Watching YouTube Videos: Action-Conditioned Contrastive Policy Pretraining

Qihang Zhang<sup>1</sup>, Zhenghao Peng<sup>1</sup>, and Bolei Zhou<sup>2</sup>

<sup>1</sup> The Chinese University of Hong Kong, Hong Kong SAR, China

<sup>2</sup> University of California, Los Angeles, USA

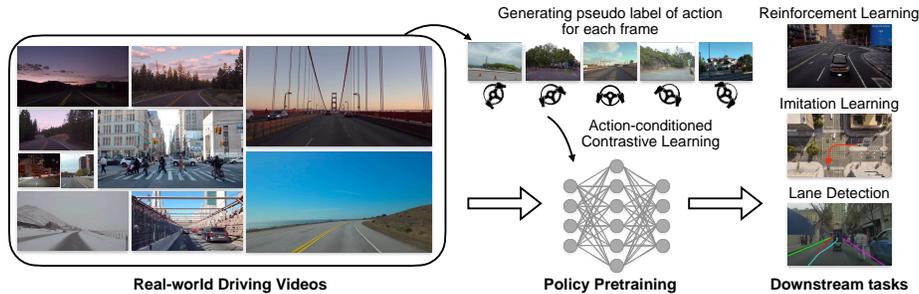
**Abstract.** Deep visuomotor policy learning, which aims to map raw visual observation to action, achieves promising results in control tasks such as robotic manipulation and autonomous driving. However, it requires a huge number of online interactions with the training environment, which limits its real-world application. Compared to the popular unsupervised feature learning for visual recognition, feature pretraining for visuomotor control tasks is much less explored. In this work, we aim to pretrain policy representations for driving tasks by watching hours-long uncurated YouTube videos. Specifically, we train an inverse dynamic model with a small amount of labeled data and use it to predict action labels for all the YouTube video frames. A new contrastive policy pretraining method is then developed to learn action-conditioned features from the video frames with pseudo action labels. Experiments show that the resulting action-conditioned features obtain substantial improvements for the downstream reinforcement learning and imitation learning tasks, outperforming the weights pretrained from previous unsupervised learning methods and ImageNet pretrained weight. Code, model weights, and data are available at: <https://metadriverse.github.io/ACO>.

**Keywords:** Learning From Unlabeled Data, Feature Pretraining

## 1 Introduction

Deep policy learning makes promising progress to many visuomotor control tasks ranging from robotic manipulation [25,20,22,38] to autonomous driving [5,47]. By learning to map visual observation directly to control action through a deep neural network, it mitigates the manual design of controller, lowers the system complexity, and improves generalizability. However, the sample efficiency of the underlying algorithms such as reinforcement learning or imitation learning remains low. It requires a significant amount of online interactions or expert demonstrations in the training environment thus limits its real-world applications.

Many recent works use unsupervised learning and data augmentation to improve the sample efficiency by pretraining the neural representations before policy learning. For example, random background videos are incorporated in the



**Fig. 1.** Hours-long YouTube videos contain a huge variety of driving scenes. By developing an action-conditioned contrastive learning algorithm, our work aims to pretrain the visual representations with pseudo action labels on the diverse real-world data and improve sample efficiency, feature generalizability, and performance of the downstream tasks

policy feature pretraining [16,43,17]. However, the augmented data with random background videos shifts drastically from the original data distribution, which degrades the overall performance of the model. Also, it remains challenging to generalize the learned weights to real world as it is hard to design augmentations that reflect the real-world diversity. In this work, we explore pretraining the neural representation on massive amount of real-world data directly. Figure 1 shows some uncurated YouTube videos, which contain driving scenes all over the world with diverse conditions such as different weathers, urban and rural environments, and various traffic densities. We show that exploiting such real-world data in deep policy learning can substantially improve the generalizability of the learned weight and benefit various downstream tasks.

Learning deep representations from unlabeled data is a popular topic in visual recognition. The learned representations are shown to be generalizable across visual tasks ranging from image classification, semantic segmentation, to object detection [18,42,13,7,39]. However, these methods are mainly designed for learning visual features for recognition tasks, which are different from control tasks where an agent takes actions in uncertain environment. In control tasks, visual information such as the abstraction of appearance and texture of the scene may not be useful for decision making. For example, visual elements like lighting and weather are usually irrelevant to the driving task and might even become confounders in policy learning, negatively affecting the driving performance [46]. On the contrary, it is crucial to learn the features that matter to the output action. For example, at the driving junction, the traffic light and the driving lane occupy only a few pixels in the visual observation but has a significant impact on the driver’s actions.

In this paper, we propose a novel action-conditioned policy pretraining paradigm, which learns to capture important features in the neural representation relevant to decision making and benefits downstream tasks. As shown in Figure 1, we first collect a large corpus of driving videos from YouTube which are

recorded in 68 cities all over the world with a wide range of visual appearances. We then train an inverse dynamics model with a small amount of labeled data and use it to generate pseudo action labels for each video frame. We finally develop a novel policy pretraining method called **A**ction-conditioned **C**ONtrastive Learning (**ACO**) that incorporates the action information in the representation learning. The motivation is to learn the discriminative features of video frames most related to driving actions. Specifically, instead of contrasting images based on different augmented views [18], we define a new contrastive pair conditioned on action similarity. By learning with those action-conditioned contrastive pairs, the representation captures visual elements that are highly correlated to the output actions.

We evaluate the effectiveness of the action-conditioned pretraining for a variety of tasks, such as policy learning through Imitation Learning (IL) and Reinforcement Learning (RL) in end-to-end autonomous driving, and Lane Detection (LD). The experimental results show that ACO successfully learns generalizable features for the downstream tasks. Our contributions are summarized below:

1. We propose a new paradigm of policy pretraining on a massive amount of real-world driving videos.
2. We develop a novel action-conditioned contrastive learning approach ACO to learn action-related features.
3. Experiments of various pretraining methods in downstream policy learning tasks show that the feature resulting from the proposed method achieves sufficient performance gain in the driving tasks.

## 2 Related work

### 2.1 Image-based Policy Learning

**Learning with auxiliary task.** Reconstruction tasks [44,15,24], self-supervised objectives [23,33], and future prediction objectives [15,41,12,29,1] are proposed to mitigate the gap between state-based and image-based Reinforcement Learning (RL). However, unsupervised learning based on visual information may not lead to optimal policy learning performance since the irrelevant visual features might become confounders in the representation and distract policy learning [46].

**Augmenting the input distribution.** Data augmentations like random shift [43,17] and inserting random videos at background [17,16] can improve policy robustness and generalization ability. Visual and physical attributes of the environment and the agent can be randomized through domain randomization [35,2,28]. However, it is hard to design augmentations that cover real world’s variations. Significant computational resources are also required to train the policy with augmented input [35].

**Decoupling visual-level and policy-level learning.** Chen *et al.* [5] first train a policy with full state observation and then use this cheating policy as a teacher to train an image-based policy. Similar teacher-student architecture is adopted in [6] to learn visuomotor grasping policy. These methods decouple

visual knowledge and policy knowledge and learn in a two-stage fashion. However, they require accessing the underlying proprioceptive state as observation to train the teacher agent.

**Pretraining.** Several works pretrain policies on offline demonstrations [45,37]. However, the training data is from the policy learning environment and requires expert policy to collect trajectories, limiting the scalability and efficiency. On the contrary, RRL [32] pretrains ResNet on ImageNet dataset as image encoder and freezes it in following policy learning process. Xiao *et al.* [40] pretrain on in-the-wild frames via masked image modeling. Our proposed ACO conducts pretraining based on YouTube videos which are captured in real world. We train an inverse dynamics model to create pseudo action labels and perform contrastive learning conditioned on these pseudo action labels. As a concurrent work, VPT [3] similarly pretrains on online videos (game playing of Minecraft) and gets pseudo action labels from a trained inverse dynamics model. VPT pretrains by behavior cloning, different from contrastive learning used in our method.

## 2.2 Contrastive Learning

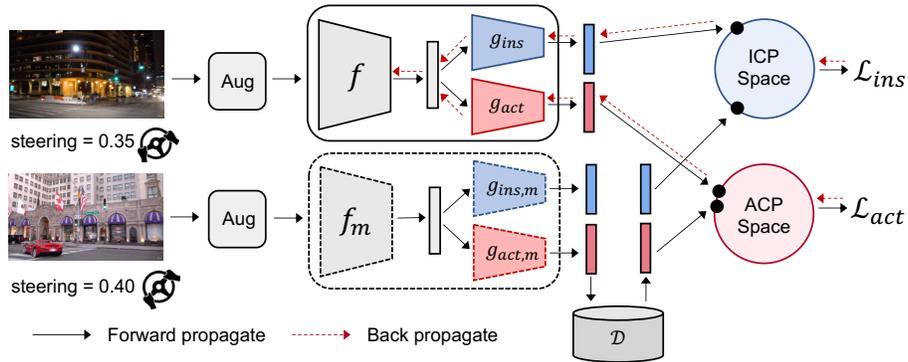
Contrastive learning [18,39] is a popular pretext task for self-supervised learning. It creates supervisory labels via considering each image (instance) in the dataset forms a unique category and applies the learning objective of instance discrimination. Previous works [18,39,7,13] in computer vision use contrastive learning to learn general vision knowledge that transfers well to various downstream tasks. In reinforcement learning, CURL [23] and ATC [33] leverage contrastive learning as an additional signal to learn representation with environmental reward.

Common practice [18] considers two views, *i.e.* different augmentations, of a single image as a positive pair and views of different images as negative pairs. We call this type of positive/negative pair as Instance Contrastive Pair (ICP). However, ICP only preserves knowledge for visual discriminability without task-relevant information. We introduce Action Contrastive Pair (ACP) that contrasts images based on the underlying actions. Learning together with ACP and ICP will force the representation to focus on the visual elements most relevant to the decision-making process.

## 3 Method

We aim to learn generalizable visual feature for visuomotor policy learning, by pretraining on large amounts of uncurated real-world driving videos. We propose a novel contrastive learning algorithm called **A**ction-**C**onditioned **C**ontrastive Learning (**ACO**).

The essential of ACO is that we define two types of contrastive pair: Instance Contrastive Pair (ICP) and Action Contrastive Pair (ACP). Two views of a single image form a positive ICP, while two views of different images form a negative ICP. Only pretraining with ICP may make the representation include



**Fig. 2.** In the training pipeline of ACO, we form both ICP and ACP for the input images and update the model with contrastive losses in ICP space and ACP space. The upper branch shows the calculation flow of query images and the lower shows the flow to generate the features of key images. History key features  $\mathcal{D}$  are stored in a key dictionary and re-used later

unnecessary information for downstream policy learning tasks. For example, visual cues like weather and lighting conditions are essential for forming ICP but contribute little to decision-making in autonomous driving. To focus the feature on policy-relevant properties, we introduce another type of contrastive pair called Action Contrastive Pair (ACP). Positive ACP is composed of two different images showing scenarios where the drivers’ actions are similar. As an example, the two real-world snapshots in Figure 2 form a positive ACP, both showing the first-view images when the drivers are turning left. ACO learns representation based on both ICP and ACP, where ICP focuses on learning discriminative general visual feature and ACP focuses on policy-relevant feature.

Figure 2 illustrates the training pipeline of ACO. In upper branch, we first augment the given images twice to create query and key views and form ICPs. As discussed in Section 3.1, the ICPs are used to compute the ICP loss for visual discriminative feature learning. We create another learning flow based on ACP. To create ACP, each frame in the dataset is tagged with a pseudo action label as introduced in Section 3.2. In Section 3.3, the action labels are used to form ACPs between different images and compute the ACP loss for policy feature learning.

### 3.1 Visual Feature Learning with ICP

As shown in Figure 2, our contrastive learning is performed in ICP space and ACP space. To project given images into ICP space and form ICP, our pipeline contains the following parts:

**Data Augmentation** module  $Aug(\cdot)$ . For each sample image  $x$ , we generate two random augmentations,  $x^q = Aug(x)$  as query view and  $x^k = Aug(x)$  as key view.

**Encoder** module  $f(\cdot)$ . Each query view  $x^q$  is further mapped to a feature vector  $v^q = f(x^q) \in \mathcal{R}^{d_E}$ , where  $d_E$  is the feature dimension. For the key view  $x^k$ , a momentum-updated Encoder module  $f_m(\cdot)$  is used to extract feature vector  $v^k = f_m(x^k) \in \mathcal{R}^{d_E}$ . The update rule of the encoder’s parameters is:

$$\theta_{f_m} \leftarrow \alpha \theta_{f_m} + (1 - \alpha) \theta_f, \quad (1)$$

where  $\alpha$  is the momentum coefficient and  $\theta_f, \theta_{f_m}$  are the parameters of  $f$  and  $f_m$ , respectively. The output of encoder module is the representation passed to policy head in downstream task.

**Projector** module  $g_{ins}(\cdot)$ . We map key and query feature  $v^q$  and  $v^k$  to instance-contrastive space:

$$z_{ins}^q = g_{ins}(v^q), \quad z_{ins}^k = g_{ins,m}(v^k), \quad (2)$$

where  $z_{ins}^q, z_{ins}^k \in \mathcal{R}^{d_P}$  and  $d_P$  is the dimension of the projected vector.  $z_{ins}^q$  and  $z_{ins}^k$  are further normalized to the unit hypersphere in  $R^{d_P}$ . The projector  $g_{ins,m}$  for key representation is also updated in a momentum way:

$$\theta_{g_{ins,m}} \leftarrow \alpha \theta_{g_{ins,m}} + (1 - \alpha) \theta_{g_{ins}}. \quad (3)$$

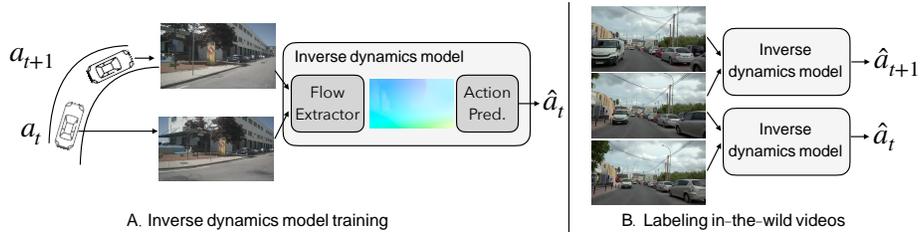
**Key Dictionary** module  $\mathcal{D}$ . Following [18], we use a key dictionary to store historical encoded key  $z_{ins}^k$  to enable larger contrastive batchsize. Concretely, we generate key features  $z^k$  of images following the lower branch in Figure 2. These key features not only form positive pairs with query features at current training epoch, but are also stored into  $\mathcal{D}$  and used to form negative pairs at future training, if sampled. The samples in the dictionary are progressively replaced if  $\mathcal{D}$  exceeds its maximum size.

At each iteration of training, a batch of images are sampled from dataset. The images are sequentially processed by Data Augmentation, Encoder and Projector modules to get query and key features. The query-key pair forms positive ICP since they come from the same image. To create negative ICP, we sample a batch (num= $N$ ) of key features from  $\mathcal{D}$  and form negative pairs with current queries. The negative keys are concatenated with current key features to form the key set  $K$ . Current key features are inserted into key dictionary for later use. Then we perform contrastive learning on ICP between query and key features.

The ICP loss for visual feature learning is computed as:

$$\begin{aligned} \mathcal{L}_{ins} &= -\log \frac{\sum_{z^+ \in P_{ins}(z_{ins}^q)} \exp(z_{ins}^q \cdot z^+ / \tau)}{\sum_{z^- \in N_{ins}(z_{ins}^q)} \exp(z_{ins}^q \cdot z^- / \tau)}, \\ P_{ins}(z^q) &= \{z \mid \text{id}(z) = \text{id}(z^q), z \in K\}, \\ N_{ins}(z^q) &\equiv K \setminus P_{ins}(z^q). \end{aligned} \quad (4)$$

The symbol  $\cdot$  denotes inner (dot) product,  $\text{id}(z^q) = x$  returns the image from which the projected vector is generated,  $P_{ins}(z^q)$  and  $N_{ins}(z^q)$  are the sets of keys of all positive and negative ICP, respectively. Notice that  $P_{ins}(z^q)$  only contains one sample and  $N_{ins}(z^q)$  contains  $N$  samples. This is because only one element in key set comes from the anchor image  $x$  while other key features are sampled from  $\mathcal{D}$ .



**Fig. 3.** **A.** We use real-world driving dataset to train an inverse dynamics model. The inverse dynamics model is composed of a frozen flow extractor and action prediction head. **B.** We use the trained inverse dynamics model to label actions on large corpus of in-the-wild first view driving videos

### 3.2 Generating Pseudo Labels for Video Frames

We propose to conduct contrastive learning based on drivers’ actions. However, uncurated videos on the web do not contain any ground-truth action. We thus train an inverse dynamics model  $\phi$  to predict the pseudo action label [21,36]  $\hat{a}_t$  that has happened between two consecutive frames  $I_t$  and  $I_{t+1}$ . Nevertheless, many visual cues, such as weather and lighting, are irrelevant to action prediction. Using two RGB frames as input may introduce unnecessary distractions to the inverse dynamics model. Instead, we find that employing optical flow, the concise description of interframe movement, improves action prediction accuracy.

As illustrated in Figure 3A, we first extract the optical flow between frames by using an off-the-shelf algorithm RAFT [34]. We then use the optical flow as the input to the inverse dynamics model  $\phi$  to predict the action. We choose NuScenes [4] dataset which contains consecutive first-view driving scenes with labeled steering action to perform supervised learning of  $\phi$ . L1 loss between ground truth action  $a_t$  from the dataset and predicted  $\hat{a}_t$  is optimized. As shown in Figure 3B, we then use the extracted optical flow together with the learned  $\phi$  to label action for each frame in YouTube videos. The pseudo action labels are further used to formulate ACPs to be introduced in the next section.

### 3.3 Policy Feature Learning with ACP

Based on the pseudo action labels, we conduct contrastive learning on ACP space to discover policy-relevant representation. We share the Data Augmentation and Encoder modules with ICP learning in Section 3.1. During learning with ACP, the projector  $g_{act}(\cdot)$  shares the same architecture as ICP’s, but not the weights. A momentum-updated ACP projector  $g_{act,m}(\cdot)$  is updated in a similar way as ICP projector in Equation 3. Thus the key and query vectors in action-contrastive space are computed as:

$$z_{act}^q = g_{act}(v^q) = g_{act}(f(x^q)), \quad z_{act}^k = g_{act,m}(v^k) = g_{act,m}(f_m(x^k)). \quad (5)$$

We also sample a batch of key features from  $\mathcal{D}$  and concatenate them with current key features to form key set  $K$ . The ACP loss is computed as:

$$\mathcal{L}_{act} = -\log \frac{\sum_{z^+ \in P_{act}(z^q_{act})} \exp(z^q_{act} \cdot z^+ / \tau)}{\sum_{z^- \in N_{act}(z^q_{act})} \exp(z^q_{act} \cdot z^- / \tau)}. \quad (6)$$

Different from ICP, the positive and negative keys in ACP are determined based on the action similarity. The projected vectors of frames with similar underlying actions form positive pairs. Therefore, when collecting the key set, we additionally store the action labels  $\hat{a}$  with the key features  $z^k$ . We use L1 distance to measure the similarity of two actions and the positive and negative key sets are determined as:

$$\begin{aligned} P_{act}(z^q) &= \{z \mid \|\hat{a} - \hat{a}^q\| < \epsilon, (z, \hat{a}) \in K\}, \\ N_{act}(z^q) &\equiv K \setminus P_{act}(z^q), \end{aligned} \quad (7)$$

where  $\hat{a}^q$  is the predicted pseudo action label of the query image and  $\epsilon$  is the action distance threshold.

### 3.4 Loss

The overall contrastive loss of ACO is:

$$\mathcal{L} = \lambda_{ins} \mathcal{L}_{ins} + \lambda_{act} \mathcal{L}_{act}, \quad (8)$$

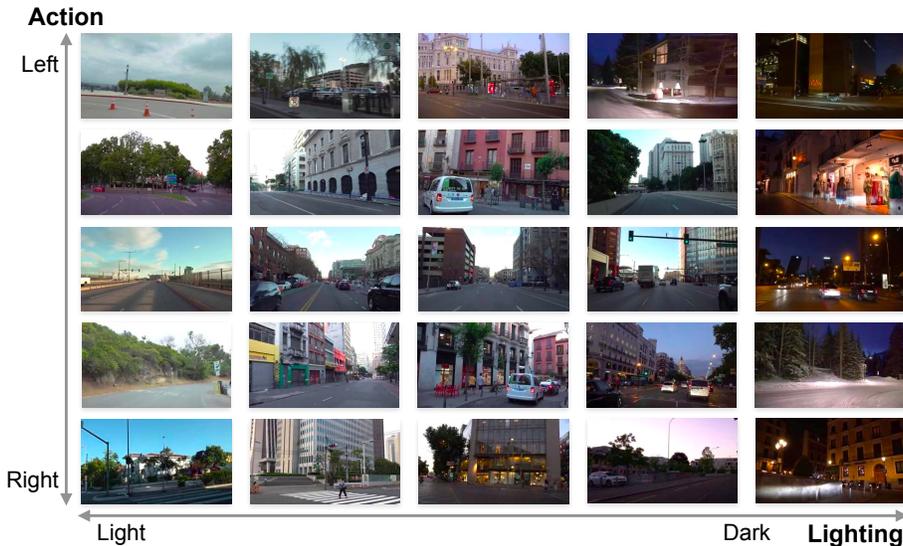
where  $\lambda_{ins}$  and  $\lambda_{act}$  are weighted factors to balance ICP and ACP.

## 4 Experiments

We evaluate the pretrained feature from our method on two main policy learning tasks: Imitation Learning (IL) and Reinforcement Learning (RL) in end-to-end driving. We also evaluate on one related visual recognition task: Lane Detection (LD). Section 4.1 lists the compared baselines. The main experimental results with a comparison of different pretraining methods are presented in Section 4.2 for IL, in Section 4.3 for RL, and in Section 4.4 for LD. Section 4.5 gives a qualitative analysis of the learned action-conditioned representation. Ablation study and further discussion are included in Section 4.6.

**YouTube driving dataset.** We crawl first-view driving videos from YouTube. 134 videos with a total length of over 120 hours are collected. As shown in Figure 1, these videos cover different driving scenes with various weather conditions (sunny, rainy, snowy, etc.) and regions (rural and urban areas). We sample one frame every one second, resulting to a dataset of 1.30 million frames. We split the YouTube driving dataset into training set with 70% data and test set with 30% data and conduct the training of ACO on the training set.

**Inverse dynamics model.** An inverse dynamics model with ResNet-34 architecture is used to predict action labels for frames in YouTube driving dataset.



**Fig. 4. Exemplar video frames organized by their pseudo labels.** We use the inverse dynamics model to label each frame’s action. The predicted action is mostly correct despite diverse visual variations of the image content

We use RAFT [34] as the optical flow extractor and apply Adam optimizer with learning rate of 0.001 and weight decay of 0.001 to learn the model on the training set of NuScenes [4]. We normalize steering value between 0 and 1. In our preliminary experiments, the L1 error of steering prediction averaged over the test set of NuScenes [4] is 0.10 when using consecutive images as input, and 0.04 when using optical flow. Using optical flow as the input brings better inverse dynamics prediction performance as it is more invariant to the image content. Figure 4 shows some sampled frames. The steering action is predicted accurately under different visual conditions.

**Pretraining.** In the representation learning, we largely follow the hyperparameters from the official implementation of MoCo-v2 [8]. We optimize the model with synchronized SGD over 8 GPUs with a weight decay of 0.0001, a momentum of 0.9, and a batch size of 32 on each GPU. We perform 100 epochs optimization with an initial learning rate of 0.03 and a cosine learning rate schedule. Two projectors for ICP and ACP are both instantiated as two-layer MLPs. We maintain a memory queue of 40,960 samples as the Key Dictionary  $\mathcal{D}$ . The momentum coefficient in Equation 1 and Equation 3 is set to 0.999 for updating the key encoder and projector.

**Data augmentations.** During pretraining, we follow MoCo-v2 [8] and use image augmentations to generate key and query views of an image. Specifically, we apply random gray scaling, color jittering and Gaussian blurring. Random resized cropping and horizontal flipping are excluded because we find that they

**Table 1. Success Rate of Imitation Learning.** We evaluate the performance of different pretrained models under IL tasks with different scales of dataset

Pretrain Method	IL Demonstration Size ( $\times 40K$ )			
	10%	20%	40%	100%
Random	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	37.3 $\pm$ 6.6	81.3 $\pm$ 5.2
AutoEncoder	0.0 $\pm$ 0.0	4.0 $\pm$ 3.3	6.7 $\pm$ 2.5	46.0 $\pm$ 5.9
ImageNet	21.3 $\pm$ 7.5	52.7 $\pm$ 13.1	72.0 $\pm$ 4.3	90.7 $\pm$ 7.4
MoCo	19.3 $\pm$ 3.4	60.3 $\pm$ 8.2	76.0 $\pm$ 9.1	80.7 $\pm$ 2.5
<b>ACO</b>	<b>30.7<math>\pm</math>3.4</b>	<b>66.0<math>\pm</math>5.7</b>	<b>82.0<math>\pm</math>5.0</b>	<b>96.0<math>\pm</math>3.3</b>

will largely change the scene’s semantic, thus degrading the overall performance, as shown in the ablation study.

#### 4.1 Baselines

We choose ResNet-34 as the backbone (encoder) architecture and compare the downstream tasks’ performance when using different pretrained models as the initial parameters of the backbone network. Different pretraining methods are as follows:

**Random.** The backbone’s parameters are initialized randomly. Specifically, we use Kaiming normalization [19] to initialize convolution layers. Constant initialization is applied to batchnorm layer (1 for weight and 0 for bias).

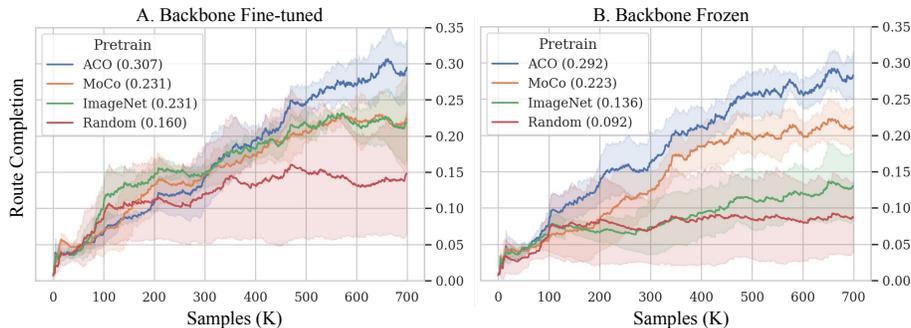
**ImageNet.** We use official ImageNet’s pretrain weights to initialize the backbone. This is considered as the most common approach for current image-based policy learning [32].

**AutoEncoder.** We train an autoencoder on the collected YouTube videos and use the encoder to initialize the backbone, which has been used in previous policy learning work [14].

**MoCo.** As a straightforward comparison to our method, contrastive learning is conducted following MoCo-v2 [8] on the collected YouTube videos. As a baseline, we only use ICP to conduct contrastive learning and set  $\lambda_{act} = 0$  in Equation 8.

#### 4.2 Imitation Learning

**Setup.** We evaluate the performance of different pretrained weights for Imitation Learning in the open-source CARLA [11] simulator. We use the original CARLA [11] benchmark (also known as CORL2017) and collect 50 trajectories of expert demonstration, which amount to 40K transitions, in Town01 with train weather. Town02 with test weather is used for evaluation. The expert is a rule-based PID controller with injected noise [9]. We adopt CILRS [10], a conditional behavior cloning algorithm, to train policies based on different pretrained models. The backbone network is fine-tuned during policy learning. We also conduct IL with frozen backbone and conclude the result in Appendix. All experiments



**Fig. 5. Route completion curves for PPO [31] with different pretrain weights.** In the left panel, the backbone is actively updated during policy learning. In the right panel, the backbone is frozen and not updated. The proposed ACO outperforms all baselines

are repeated three times with different random seeds. Information about other hyper-parameters is given in Appendix.

**Results.** To see how pretrained models improve IL tasks, we experiment on the expert datasets with 10%, 20%, 40% and 100% of all 40K samples. As shown in Table 1, ACO outperforms all other baselines in datasets of all scales. Notably, AutoEncoder is less competitive even compared to random initialization. This is because pixel reconstruction target is misaligned with policy learning target and makes the encoder pay more attention to areas with larger size like buildings and the sky, which however are useless for imitation learning. MoCo and ImageNet perform well at full dataset but have degraded performance when dataset size is reduced. ACO-pretrained models consistently outperform baselines at all different dataset sizes.

### 4.3 Reinforcement Learning

**Setup.** Similar to Imitation Learning, we also use CARLA [11] as the simulator for evaluating pretrained models in Reinforcement Learning. We use Town01 in NoCrash benchmark [10] as the environment and evaluate the performance of Proximal Policy Optimization (PPO) [31] with different pretrained weights.

**Results.** As shown in Figure 5A, PPO with ACO initialization introduces a clear improvement of 7% route completion compared to other baselines when the backbone is fine-tuned during policy learning. This suggests that incorporating policy-relevant objects into pretraining stage benefits downstream RL’s performance. ImageNet and MoCo, which only learn visual knowledge, are less competitive compared to ACO. For random initialization, we observe a poor 15% route completion, in which case the agent can only drive in straight road but fails to learn turning.

To further understand the importance of considering policy-relevant information in pretraining, we freeze the backbone and only finetune the policy head

**Table 2. F1 metric of lane detection**

Pretrain Method	LD Method		
	UFLD [30]	SCNN [27]	RESA [48]
<i>ImageNet</i>	<i>70.9</i>	<i>74.1</i>	<i>76.1</i>
Random	68.4	70.4	73.9
AutoEncoder	68.6	70.1	74.0
MoCo	69.5	<b>72.0</b>	74.1
MoCo-bn	69.8	71.9	74.1
ACO	69.5	71.6	74.5
ACO-bn	<b>70.4</b>	71.8	<b>75.0</b>

during RL. As shown in Figure 5B, compared to backbone fine-tuned case, the performance of PPO declines dramatically when using ImageNet pretraining. This indicates that ImageNet pretrained weight lacks the essential knowledge for policy learning thus leads to a poor RL performance. ACO’s and MoCo’s performance are not affected when freezing the backbone.

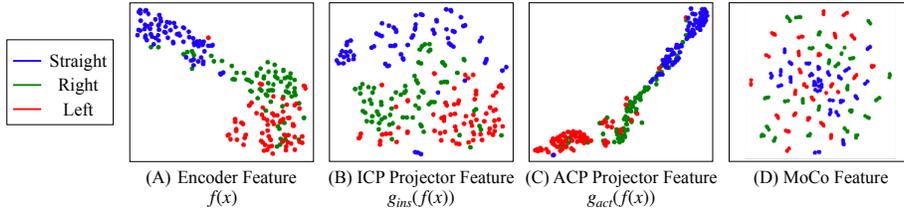
#### 4.4 Lane Detection

**Setup.** To further prove that the representation learned by ACO is generalizable, we conduct lane detection experiments. We train three popular lane detection methods on CULane dataset [27]: UFLD [30], SCNN [27], and RESA [48]. All of these methods employ backbone networks with ResNet architecture. We use ACO and other baselines to initialize the backbone’s parameters before training. We adopt the training settings (*e.g.* learning rate) of three LD methods fine-tuned on ImageNet pretrained backbone. Better performance is expected after tuning the hyper-parameters for different pretrained methods respectively. MoCo [18] suggests that a batch normalization layer after backbone can alleviate the gap between contrastive learning and supervised learning caused by different parameters’ distribution. We thus follow this procedure on ACO and MoCo baseline and present the results with “-bn” suffix.

**Results.** As shown in Table 2, ImageNet pretraining achieves the best performance among all candidates. Although trained in an object-centric dataset, ImageNet still provides high quality feature for lane detection learning. Within the self-supervised baselines, ACO outperforms others in UFLD [30] and RESA [48]. In SCNN [27], ACO is on par with MoCo. Introducing batch normalization improves the performance of ACO.

#### 4.5 Visualization

We perform t-SNE [26] analysis on features extracted by different layers of ACO and compare them with that of MoCo. We randomly select frames from the test set and divide them into three categories: *left turn*, *right turn*, and *go straight*.



**Fig. 6. t-SNE [26] visualization of extracted features.** Red, green, and blue dots indicate frames under *left turn*, *right turn*, and *go straight* scenarios respectively

**Table 3. Ablation study of pretraining datasets.** We use ACO to pretrain models on our YouTube dataset and the NuScenes dataset respectively, and conduct IL training on CARLA

Pretraining Dataset	Frame num (M)	IL Demonstration Size ( $\times 40K$ )			
		10%	20%	40%	100%
NuScenes [4]	0.19@10Hz	13.3 $\pm$ 3.4	27.3 $\pm$ 10.4	67.3 $\pm$ 9.0	82.7 $\pm$ 8.4
<b>YouTube</b>	<b>1.30@1Hz</b>	<b>30.7<math>\pm</math>3.4</b>	<b>66.0<math>\pm</math>5.7</b>	<b>82.0<math>\pm</math>5.0</b>	<b>96.0<math>\pm</math>3.3</b>

Each category has 25 examples. Images within same category form positive ACPs with one another. To build positive ICPs, we sample the previous and next frame of chosen samples and consider them as positive pairs.

As shown in Figure 6A, features extracted by ACO’s encoder, which will be passed to policy head in downstream task, form separable clusters and have clear semantics related to the possible actions of drivers. On the contrary, the action-conditioned clustering does not emerge in MoCo in Figure 6D. Instead, we find that MoCo feature demonstrates the *instance clustering phenomenon*, where the features of three temporal neighboring frames lay closely and form a clique in the t-SNE visualization.

Diving into the contrastive learning modules, we find the ICP projector demonstrates instance clustering as in Figure 6B. The phenomenon does not exist in ACP projected features in Figure 6C, where features cluster more tightly according to action information instead of temporal relationship. t-SNE visualization shows that ICP and ACP in ACO both achieve their respective goals: ICP endows features with the capacity to discriminate across instances, whereas ACP creates features that tightly couple with action information.

#### 4.6 Ablation Study

We conduct ablation studies that examine the effectiveness of collecting and labeling YouTube driving videos rather than using NuScenes dataset [4] and the impact of geometric-aware augmentations (cropping and flipping).

**YouTube Dataset.** To highlight the importance of collecting large dataset from YouTube, we train IL agents based on models pretrained on different datasets

**Table 4. Ablation study of augmentations.** CF stands for Cropping and Flipping

Methods	IL Demonstration Size ( $\times 40K$ )			
	10%	20%	40%	100%
MoCo w/ CF	19.3 $\pm$ 3.4	60.3 $\pm$ 8.2	76.0 $\pm$ 9.1	80.7 $\pm$ 2.5
MoCo w/o CF	13.3 $\pm$ 3.4	47.3 $\pm$ 13.7	72.7 $\pm$ 9.0	84.0 $\pm$ 4.9
ACO w/ CF	27.3 $\pm$ 0.9	36.0 $\pm$ 5.9	74.0 $\pm$ 7.5	89.3 $\pm$ 5.0
ACO w/o CF	<b>30.7<math>\pm</math>3.4</b>	<b>66.0<math>\pm</math>5.7</b>	<b>82.0<math>\pm</math>5.0</b>	<b>96.0<math>\pm</math>3.3</b>

and compare the success rate on IL benchmark. As shown in Table 3, the size of our collected YouTube dataset is six times larger than NuScenes dataset [4] even though we sample frames at 1Hz from the video, compared to 10Hz of NuScenes. The imitation learning performance of ACO trained on YouTube dataset outperforms that on NuScenes dataset [4] at all IL training settings. The result shows that YouTube dataset improves the generalizability of the pretrained models and leads to better performance in the downstream task.

**Augmentations.** We conduct imitation learning experiments to discover the impact of two augmentations, cropping and flipping frames (CF in the following for lucidity). As shown in Table 4, the performance of MoCo without CF is significantly lower than MoCo with these two augmentations, showing that CF is important for instance discriminative learning. However, the result suggests ACO is incompatible with CF. Excluding CF in ACO will affect the feature’s instance discriminative ability but improve the overall performance in downstream driving tasks with action-conditioned contrastive learning.

## 5 Conclusion and Discussion

In this paper, we propose a novel contrastive policy pretraining method ACO using hours-long uncurated YouTube videos. By learning action-conditioned features from unlabeled video frames with pseudo action labels, our methods greatly improve the generalizability of the learned representation and brings substantial improvements to downstream tasks.

**Limitations.** Despite the generalizable representation provided by pretraining with action information, we incorporate a strong assumption: only one action corresponds to each driving scene in a single video frame. This does not hold true as drivers may have various driving intentions in complex scenarios. Future study will concentrate on how to label frames more accurately, *e.g.* labeling each frame with an action distribution, to produce more precise action contrastive pairs. Besides, action-conditioned contrastive learning has conflict with widely used geometric-aware augmentations like cropping and flipping, suppressing feature’s instance discrimination ability. We leave this problem for future study.

## References

1. Agrawal, P., Nair, A.V., Abbeel, P., Malik, J., Levine, S.: Learning to poke by poking: Experiential learning of intuitive physics. *Advances in neural information processing systems* **29** (2016)
2. Andrychowicz, O.M., Baker, B., Chociej, M., Jozefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., et al.: Learning dexterous in-hand manipulation. *The International Journal of Robotics Research* **39**(1), 3–20 (2020)
3. Baker, B., Akkaya, I., Zhokhov, P., Huizinga, J., Tang, J., Ecoffet, A., Houghton, B., Sampedro, R., Clune, J.: Video pretraining (vpt): Learning to act by watching unlabeled online videos. *arXiv preprint arXiv:2206.11795* (2022)
4. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11621–11631 (2020)
5. Chen, D., Zhou, B., Koltun, V., Krähenbühl, P.: Learning by cheating. In: *Conference on Robot Learning*. pp. 66–75. PMLR (2020)
6. Chen, T., Xu, J., Agrawal, P.: A system for general in-hand object re-orientation. In: *Conference on Robot Learning*. pp. 297–307. PMLR (2022)
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International conference on machine learning*. pp. 1597–1607. PMLR (2020)
8. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297* (2020)
9. Codevilla, F., Müller, M., López, A., Koltun, V., Dosovitskiy, A.: End-to-end driving via conditional imitation learning. In: *2018 IEEE international conference on robotics and automation (ICRA)*. pp. 4693–4700. IEEE (2018)
10. Codevilla, F., Santana, E., López, A.M., Gaidon, A.: Exploring the limitations of behavior cloning for autonomous driving. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9329–9338 (2019)
11. Dosovitskiy, A., Ros, G., Codevilla, F., López, A.M., Koltun, V.: CARLA: an open urban driving simulator. *CoRR* **abs/1711.03938** (2017), <http://arxiv.org/abs/1711.03938>
12. Finn, C., Tan, X.Y., Duan, Y., Darrell, T., Levine, S., Abbeel, P.: Learning visual feature spaces for robotic manipulation with deep spatial autoencoders. *arXiv preprint arXiv:1509.06113* **25** (2015)
13. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems* **33**, 21271–21284 (2020)
14. Ha, D., Schmidhuber, J.: World models. *arXiv preprint arXiv:1803.10122* (2018)
15. Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., Davidson, J.: Learning latent dynamics for planning from pixels. In: *International conference on machine learning*. pp. 2555–2565. PMLR (2019)
16. Hansen, N., Jangir, R., Sun, Y., Alenyà, G., Abbeel, P., Efros, A.A., Pinto, L., Wang, X.: Self-supervised policy adaptation during deployment. *arXiv preprint arXiv:2007.04309* (2020)
17. Hansen, N., Wang, X.: Generalization in reinforcement learning by soft data augmentation. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 13611–13617. IEEE (2021)

18. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
19. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: Proceedings of the IEEE international conference on computer vision. pp. 1026–1034 (2015)
20. Kalashnikov, D., Irpan, A., Pastor, P., Ibarz, J., Herzog, A., Jang, E., Quillen, D., Holly, E., Kalakrishnan, M., Vanhoucke, V., et al.: Scalable deep reinforcement learning for vision-based robotic manipulation. In: Conference on Robot Learning. pp. 651–673. PMLR (2018)
21. Kumar, A., Gupta, S., Malik, J.: Learning navigation subroutines by watching videos. corr abs/1905.12612 (2019) (1905)
22. Lange, S., Riedmiller, M., Voigtländer, A.: Autonomous reinforcement learning on raw visual input data in a real world application. In: The 2012 international joint conference on neural networks (IJCNN). pp. 1–8. IEEE (2012)
23. Laskin, M., Srinivas, A., Abbeel, P.: Curl: Contrastive unsupervised representations for reinforcement learning. In: International Conference on Machine Learning. pp. 5639–5650. PMLR (2020)
24. Lee, A.X., Nagabandi, A., Abbeel, P., Levine, S.: Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. Advances in Neural Information Processing Systems **33**, 741–752 (2020)
25. Levine, S., Finn, C., Darrell, T., Abbeel, P.: End-to-end training of deep visuomotor policies. The Journal of Machine Learning Research **17**(1), 1334–1373 (2016)
26. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research **9**(11) (2008)
27. Pan, X., Shi, J., Luo, P., Wang, X., Tang, X.: Spatial as deep: Spatial cnn for traffic scene understanding. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)
28. Peng, X.B., Andrychowicz, M., Zaremba, W., Abbeel, P.: Sim-to-real transfer of robotic control with dynamics randomization. In: 2018 IEEE international conference on robotics and automation (ICRA). pp. 3803–3810. IEEE (2018)
29. Pinto, L., Gandhi, D., Han, Y., Park, Y.L., Gupta, A.: The curious robot: Learning visual representations via physical interactions. In: European Conference on Computer Vision. pp. 3–18. Springer (2016)
30. Qin, Z., Wang, H., Li, X.: Ultra fast structure-aware deep lane detection. In: European Conference on Computer Vision. pp. 276–291. Springer (2020)
31. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017)
32. Shah, R., Kumar, V.: Rrl: Resnet as representation for reinforcement learning. arXiv preprint arXiv:2107.03380 (2021)
33. Stooke, A., Lee, K., Abbeel, P., Laskin, M.: Decoupling representation learning from reinforcement learning. In: International Conference on Machine Learning. pp. 9870–9879. PMLR (2021)
34. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: European conference on computer vision. pp. 402–419. Springer (2020)
35. Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P.: Domain randomization for transferring deep neural networks from simulation to the real world. In: 2017 IEEE/RSJ international conference on intelligent robots and systems (IROS). pp. 23–30. IEEE (2017)
36. Torabi, F., Warnell, G., Stone, P.: Behavioral cloning from observation. In: IJCAI (2018)

37. Wang, C., Luo, X., Ross, K., Li, D.: Vrl3: A data-driven framework for visual deep reinforcement learning. arXiv preprint arXiv:2202.10324 (2022)
38. Wu, B., Nair, S., Fei-Fei, L., Finn, C.: Example-driven model-based reinforcement learning for solving long-horizon visuomotor tasks. arXiv preprint arXiv:2109.10312 (2021)
39. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3733–3742 (2018)
40. Xiao, T., Radosavovic, I., Darrell, T., Malik, J.: Masked visual pre-training for motor control. arXiv preprint arXiv:2203.06173 (2022)
41. Yan, W., Vangipuram, A., Abbeel, P., Pinto, L.: Learning predictive representations for deformable objects using contrastive estimation. arXiv preprint arXiv:2003.05436 (2020)
42. Yang, C., Wu, Z., Zhou, B., Lin, S.: Instance localization for self-supervised detection pretraining. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3987–3996 (2021)
43. Yarats, D., Kostrikov, I., Fergus, R.: Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In: International Conference on Learning Representations (2020)
44. Yarats, D., Zhang, A., Kostrikov, I., Amos, B., Pineau, J., Fergus, R.: Improving sample efficiency in model-free reinforcement learning from images. arXiv preprint arXiv:1910.01741 (2019)
45. Zhan, A., Zhao, P., Pinto, L., Abbeel, P., Laskin, M.: A framework for efficient robotic manipulation. arXiv preprint arXiv:2012.07975 (2020)
46. Zhang, A., McAllister, R., Calandra, R., Gal, Y., Levine, S.: Learning invariant representations for reinforcement learning without reconstruction. arXiv preprint arXiv:2006.10742 (2020)
47. Zhang, Z., Liniger, A., Dai, D., Yu, F., Van Gool, L.: End-to-end urban driving by imitating a reinforcement learning coach. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
48. Zheng, T., Fang, H., Zhang, Y., Tang, W., Yang, Z., Liu, H., Cai, D.: Resa: Recurrent feature-shift aggregator for lane detection (2020)