

Transitive Invariance for Self-supervised Visual Representation Learning

Xiaolong Wang¹ Kaiming He² Abhinav Gupta¹
¹Carnegie Mellon University ²Facebook AI Research

Abstract

Learning visual representations with self-supervised learning has become popular in computer vision. The idea is to design auxiliary tasks where labels are free to obtain. Most of these tasks end up providing data to learn specific kinds of invariance useful for recognition. In this paper, we propose to exploit different self-supervised approaches to learn representations invariant to (i) inter-instance variations (two objects in the same class should have similar features) and (ii) intra-instance variations (viewpoint, pose, deformations, illumination, etc.). Instead of combining two approaches with multi-task learning, we argue to organize and reason the data with multiple variations. Specifically, we propose to generate a graph with millions of objects mined from hundreds of thousands of videos. The objects are connected by two types of edges which correspond to two types of invariance: “different instances but a similar viewpoint and category” and “different viewpoints of the same instance”. By applying simple transitivity on the graph with these edges, we can obtain pairs of images exhibiting richer visual invariance. We use this data to train a Triplet-Siamese network with VGG16 as the base architecture and apply the learned representations to different recognition tasks. For object detection, we achieve 63.2% mAP on PASCAL VOC 2007 using Fast R-CNN (compare to 67.3% with ImageNet pre-training). For the challenging COCO dataset, our method is surprisingly close (23.5%) to the ImageNet-supervised counterpart (24.4%) using the Faster R-CNN framework. We also show that our network can perform significantly better than the ImageNet network in the surface normal estimation task.

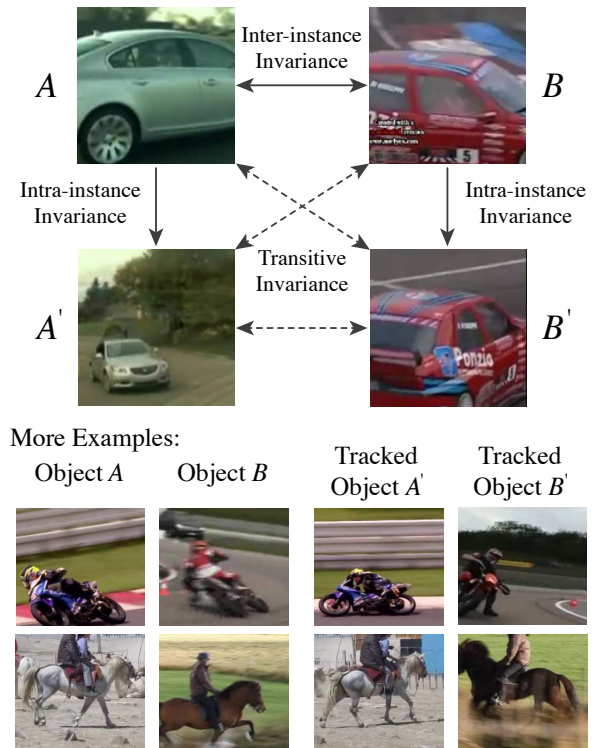


Figure 1: We propose to obtain rich invariance by applying simple transitive relations. In this example, two different cars A and B are linked by the features that are good for inter-instance invariance (e.g., using [9]); and each car is linked to another view (A' and B') by visual tracking [56]. Then we can obtain new invariance from object pairs $\langle A, B' \rangle$, $\langle A', B \rangle$, and $\langle A', B' \rangle$ via transitivity. We show more examples in the bottom.

1. Introduction

Visual invariance is a core issue in learning visual representations. Traditional features like SIFT [36] and HOG [6] are histograms of edges that are to an extent invariant to illumination, orientations, scales, and translations. Modern deep representations are capable of learning high-level invariance from large-scale data [42], e.g., viewpoint, pose, deformation, and semantics. These can also be transferred

to complicated visual recognition tasks [17, 35].

In the scheme of supervised learning, human annotations that map a variety of examples into a single label provide supervision for learning invariant representations. For example, two horses with different illumination, poses, and breeds are invariantly annotated as a category of “horse”. Such human knowledge on invariance is expected to be learned by capable deep neural networks [30, 27] through

carefully annotated data. However, large-scale, high-quality annotations come at a cost of expensive human effort.

Unsupervised or “self-supervised” learning (e.g., [56, 9, 40, 57, 58, 32, 39]) recently has attracted increasing interests because the “labels” are free to obtain. Unlike supervised learning that learns invariance from the semantic labels, the self-supervised learning scheme mines it from the nature of the data. We observe that most self-supervised approaches learn representations that are invariant to: (i) *inter-instance* variations, which reflects the commonality among different instances. For example, relative positions of patches [9] (see also Figure 3) or channels of colors [57, 58] can be predicted through the commonality shared by many object instances; (ii) *intra-instance* variations. Intra-instance invariance is learned from the pose, view-point, and illumination changes by tracking a single moving instance in videos [56, 39]. However, either source of invariance can be as rich as that provided by human annotations on large-scale datasets like ImageNet.

Even after significant advances in the field of self-supervised learning, there is still a long way to go compared to supervised learning. What should be the next steps? It seems that an obvious way is to obtain multiple sources of invariance by combining multiple self-supervised tasks, e.g., via multiple losses. Unfortunately, this naïve solution turns out to give little improvement (as we will show by experiments).

We argue that the trick lies not in the tasks but in the way of exploiting data. To leverage both intra-instance and inter-instance invariance, in this paper we construct a huge affinity graph consisting of two types of edges (see Figure 1): the first type of edges relates “different instances of similar viewpoints/poses and potentially the same category”, and the second type of edges relates “different viewpoints/poses of an identical instance”. We instantiate the first type of edges by learning commonalities across instances via the approach of [9], and the second type by unsupervised tracking of objects in videos [56]. We set up simple transitive relations on this graph to infer more complex invariance from the data, which are then used to train a Triplet-Siamese network for learning visual representations.

Experiments show that our representations learned without any annotations can be well transferred to the object detection task. Specifically, we achieve 63.2% mAP with VGG16 [45] when fine-tuning Fast R-CNN on VOC2007, against the ImageNet pre-training baseline of 67.3%. More importantly, we also report the first-ever result of un-/self-supervised pre-training models fine-tuned on the challenging COCO object detection dataset [34], achieving 23.5% AP comparing against 24.4% AP that is fine-tuned from an ImageNet pre-trained counterpart (both using VGG16). To our knowledge, this is the closest accuracy to the ImageNet pre-training counterpart obtained on object detection tasks.

2. Related Work

Unsupervised learning of visual representations is a research area of particular interest. Approaches to unsupervised learning can be roughly categorized into two main streams: (i) generative models, and (ii) self-supervised learning. Earlier methods for generative models include Auto-Encoders [38, 51, 31, 29] and Restricted Boltzmann Machines (RBMs) [23, 4, 49, 12]. For example, Le *et al.* [29] trained a multi-layer auto-encoder on a large-scale dataset of YouTube videos: although no label is provided, some neurons in high-level layers can recognize cats and human faces. Recent generative models such as Generative Adversarial Networks [20] and Variational Auto-Encoders [26] are capable of generating more realistic images. The generated examples or the neural networks that learn to generate examples can be exploited to learn representations of data [11, 10].

Self-supervised learning is another popular stream for learning invariant features. Visual invariance can be captured by the same instance/scene taken in a sequence of video frames [56, 48, 25, 1, 37, 52, 32, 39]. For example, Wang and Gupta [56] leverage tracking of objects in videos to learn visual invariance within individual objects; Jayaraman and Grauman [25] train a Siamese network to model the ego-motion between two frames in a scene; Mathieu *et al.* [37] propose to learn representations by predicting future frames; Pathak *et al.* [39] train a network to segment the foreground objects where are acquired via motion cues. On the other hand, common characteristics of different object instances can also be mined from data [9, 57, 58]. For example, relative positions of image patches [9] may reflect feasible spatial layouts of objects; possible colors can be inferred [57, 58] if the networks can relate colors to object appearances. Rather than rely on temporal changes in video, these methods are able to exploit still images.

Our work is also closely related to mid-level patch clustering [46, 7, 8] and unsupervised discovery of semantic classes [43, 47] as we attempt to find reliable clusters in our affinity graph. In addition, the ranking function used in this paper is related to deep metric learning with Siamese architectures [5, 21, 19, 54, 24].

Analysis of the two types of invariance. Our generic framework can be instantiated by any two self-supervised methods that can respectively learn inter-/intra-instance invariance. In this paper we adopt Doersch *et al.*’s [9] context prediction method to build inter-instance invariance, and Wang and Gupta’s [56] tracking method to build intra-instance invariance. We analyze their behaviors as follows.

The context prediction task in [9] randomly samples a patch (blue in Figure 3) and one of its eight neighbors (red), and trains the network to predict their relative position, defined as an 8-way classification problem. In the first two

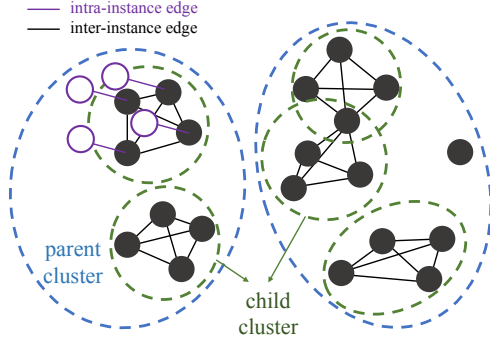


Figure 2: Illustrations for our graph construction. We first cluster the object nodes into coarser clusters (namely “parent” clusters) and then inside each cluster we perform nearest-neighbor search to obtain “child” clusters consisting of 4 samples. Samples in each child cluster are linked to each other with the “inter-instance” edges. We add new samples via visual tracking and link them to the original objects by “intra-instance” edges.

examples in Figure 3, the context prediction model is able to predict that the “leg” patch is below the “face” patch of the cat, indicating that the model has learned some commonality of spatial layout from the training data. However, the model would fail if the pose, viewpoint, or deformation of the object is changed drastically, *e.g.*, in the third example of Figure 3 — unless the dataset is diversified and large enough to include gradually changing poses, it is hard for the models to learn that the changed pose can be of the same object type.

On the other hand, these changes can be more successfully captured by the visual tracking method presented in [56], *e.g.*, see $\langle A, A' \rangle$ and $\langle B, B' \rangle$ in Figure 1. But by tracking an identical instance we cannot associate different instances of the same semantics. Thus we expect the representations learned in [56] are weak in handling the variations between different objects in the same category.

3. Overview

Our goal is to learn visual representations which capture: (i) inter-instance invariance (*e.g.*, two instances of cats should have similar features), and (ii) intra-instance invariance (pose, viewpoint, deformation, illumination, and other variance of the same object instance). We have tried to formulate this as a multi-task (multi-loss) learning problem in our initial experiments (detailed in Table 2 and 3) and observed unsatisfactory performance. Instead of doing so, we propose to obtain a richer set of invariance by performing transitive reasoning on the data.

Our first step is to construct a graph that describes the affinity among image patches. A node in the graph denotes an image patch. We define two types of edges in the graph




Figure 3: The *context prediction* task defined in [9]. Given two patches in an image, it learns to predict the relative position between them.

that relate image patches to each other. The first type of edges, called *inter-instance* edges, link two nodes which correspond to different object instances of similar visual appearance; the second type of edges, called *intra-instance* edges, link two nodes which correspond to an identical object captured at different time steps of a track. The solid arrows in Figure 1 illustrate these two types of edges.

Given the built graph, we want to *transit* the relations via the known edges and associate unconnected nodes that may provide under-explored invariance (Figure 1, dash arrows). Specifically, as shown in Figure 1, if patches $\langle A, B \rangle$ are linked via an inter-instance edge and $\langle A, A' \rangle$ and $\langle B, B' \rangle$ respectively are linked via “intra-instance” edges, we hope to enrich the invariance by simple transitivity and relate three new pairs of: $\langle A', B' \rangle$, $\langle A, B' \rangle$, and $\langle A', B \rangle$ (Figure 1, dash arrows).

We train a Triplet-Siamese network that encourages similar visual representations between the invariant samples (*e.g.*, any pair consisting of A, A', B, B') and at the same time discourages similar visual representations to a third distractor sample (*e.g.*, a random sample C unconnected to A, A', B, B'). In all of our experiments, we apply VGG16 [45] as the backbone architecture for each branch of this Triplet-Siamese network. The visual representations learned by this backbone architecture are evaluated on other recognition tasks.

4. Graph Construction

We construct a graph with inter-instance and intra-instance edges. Firstly, we apply the method of [56] on a large set of 100K unlabeled videos (introduced in [56]) and mine millions of moving objects using motion cues (Sec. 4.1). We use the image patches of them to construct the nodes of the graph.

We instantiate inter-instance edges by the self-supervised method of [9] that learns context predictions on a large set of still images, which provide features to cluster the nodes and set up inter-instance edges (Sec. 4.2). On the other hand, we connect the image patches in the same visual track by intra-instance edges (Sec. 4.3).




Figure 4: Some example clustering results. Each row shows the 4 examples in a child cluster (Sec. 4.2).

4.1. Mining Moving Objects

We follow the approach in [56] to find the moving objects in videos. As a brief introduction, this method first applies Improved Dense Trajectories (IDT) [53] on videos to extract SURF [2] feature points and their motion. The video frames are then pruned if there is too much motion (indicating camera motion) or too little motion (*e.g.*, noisy signals). For the remaining frames, it crop a 227×227 bounding box (from $\sim 600 \times 400$ images) which includes the most number of moving points as the foreground object. However, for computational efficiency, in this paper we rescale the image patches to 96×96 after cropping and we use them as inputs for clustering and training.

4.2. Inter-instance Edges via Clustering

Given the extracted image patches which act as nodes, we want to link them with extra inter-instance edges. We rely on the visual representations learned from [9] to do this. We connect the nodes representing image patches which are close in the feature space. In addition, motivated by the mid-level clustering approaches [46, 7], we want to obtain millions of object clusters with a small number of objects in each to maintain high “purity” of the clusters. We describe the implementation details of this step as follows.

We extract the `pool5` features of the VGG16 network trained as in [9]. Following [9], we use ImageNet *without* labels to train this network. Note that because we use a patch size of 96×96 , the dimension of our `pool5` feature is $3 \times 3 \times 512 = 4608$. The distance between samples is calculated by the cosine distance of these features. We want the object patches in each cluster to be close to each other in the feature space, and we care less about the differences between clusters. However, directly clustering millions of image patches into millions of small clusters (*e.g.*, by K-

means) is time consuming. So we apply a hierarchical clustering approach (2-stage in this paper) where we first group the images into a relatively small number of clusters, and then find groups of small number of examples inside each cluster via nearest-neighbor search.

Specifically, in the first stage of clustering, we apply K-means clustering with $K = 5000$ on the image patches. We then remove the clusters with number of examples less than 100 (this reduces K to 546 in our experiments on the image patches mined from the video dataset). We view these clusters as the “parent” clusters (blue circles in Figure 2). Then in the second stage of clustering, inside each parent cluster, we perform nearest-neighbor search for each sample and obtain its top 10 nearest neighbors in the feature space. We then find any group of samples with a group size of 4, inside which all the samples are each other’s top-10 nearest neighbors. We call these small clusters with 4 samples “child” clusters (green circles in Figure 2). We then link these image patches with each other inside a child cluster via “inter-instance” edges. Note that different child clusters may overlap, *i.e.*, we allow the same sample to appear in different groups. However, in our experiments we find that most samples appear only in one group. We show some results of clustering in Figure 4.

4.3. Intra-instance Edges via Tracking

To obtain rich variations of viewpoint and deformation changes of the same object instance, we apply visual tracking on the mined moving objects in the videos as in [56]. More specifically, given a moving object in the video, it applies KCF [22] to track the object for $N = 30$ frames and obtain another sample of the object in the end of the track. Note that the KCF tracker does not require any human supervision. We add these new objects as nodes to the graph and link the two samples in the same track with an intra-instance edge (purple in Figure 2).

5. Learning with Transitions in the Graph

With the graph constructed, we want to link more image patches (see dotted links in Figure 1) which may be related via the transitivity of invariance. Objects subject to different levels of invariance can thus be related to each other. Specifically, if we have a set of nodes $\{A, B, A', B'\}$ where $\langle A, B \rangle$ are connected by an inter-instance edge and $\langle A, A' \rangle$ and $\langle B, B' \rangle$ are connected by an intra-instance edge, by assuming transitivity of invariance we expect the new pairs of $\langle A, B' \rangle$, $\langle A', B \rangle$, and $\langle A', B' \rangle$ to share similar high-level visual representations. Some examples are illustrated in Figure 1 and 5.

We train a deep neural network (VGG16) to generates similar visual representations if the image patches are linked by inter-instance/intra-instance edges or their transitivity (which we call a positive pair of samples). To avoid a

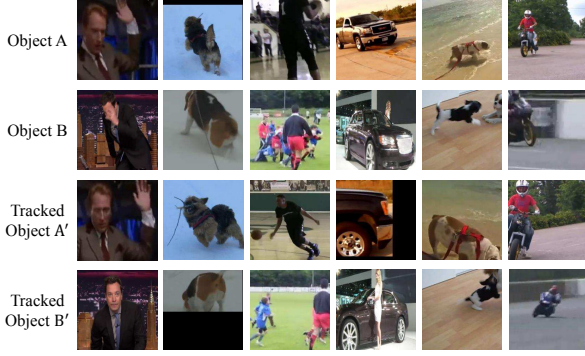


Figure 5: Examples used for training the network. Each column shows a set of image patches $\{A, B, A', B'\}$. Here, A and B is linked by an inter-instance edge, and A'/B' is linked to A/B via intra-instance edges.

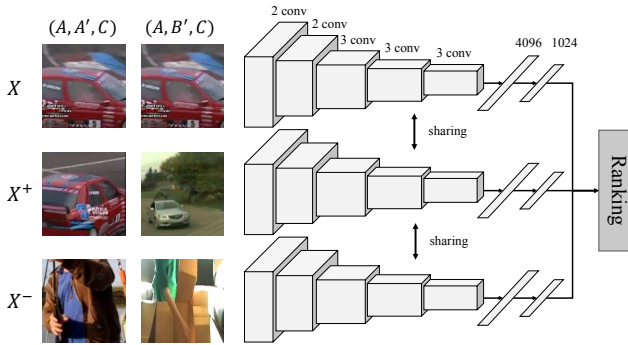


Figure 6: Our Triplet-Siamese network. We can feed in the network with different combinations of examples.

trivial solution of identical representations, we also encourage the network to generate dissimilar representations if a node is expected to be unrelated. Specifically, we constrain the image patches from different “parent” clusters (which are more likely to have different categories) to have different representations (which we call a negative pair of samples). We design a Triplet-Siamese network with a ranking loss function [54, 56] such that the distance between related samples should be smaller than the distance of unrelated samples.

Our Triplet-Siamese network includes three towers of a ConvNet with shared weights (Figure 6). For each tower, we adopt the standard VGG16 architecture [45] to the convolutional layers, after which we add two fully-connected layers with 4096-d and 1024-d outputs. The Triplet-Siamese network accepts a triplet sample as its input: the first two image patches in the triplet are a positive pair, and the last two are a negative pair. We extract their 1024-d features and calculate the ranking loss as follows.

Given an arbitrary pair of image patches A and B , we de-

fine their distance as: $D(A, B) = 1 - \frac{F(A) \cdot F(B)}{\|F(A)\| \|F(B)\|}$ where $F(\cdot)$ is the representation mapping of the network. With a triplet of (X, X^+, X^-) where (X, X^+) is a positive pair and (X, X^-) is a negative pair as defined above, we minimize the ranking loss:

$$\mathcal{L}(X, X^+, X^-) = \max\{0, \mathcal{D}(X, X^+) - \mathcal{D}(X, X^-) + m\},$$

where m is a margin set as 0.5 in our experiments. Although we have only one objective function, we have different types of training examples. As illustrated in Figure 6, given the set of related samples $\{A, B, A', B'\}$ (see Figure 5) and a random distractor sample C from another parent cluster, we can train the network to handle, e.g., view-point invariance for the same instance via $\mathcal{L}(A, A', C)$ and invariance to different objects sharing the same semantics via $\mathcal{L}(A, B', C)$.

Besides exploring these relations, we have also tried to enforce the distance between different objects to be larger than the distance between two different viewpoints of the same object, e.g., $\mathcal{D}(A, A') < \mathcal{D}(A, B')$. But we have not found this extra relation brings any improvement. Interestingly, we found that the representations learned by our method can in general satisfy $\mathcal{D}(A, A') < \mathcal{D}(A, B')$ after training.

6. Experiments

We perform extensive analysis on our self-supervised representations. We first evaluate our ConvNet as a feature extractor on different tasks *without* fine-tuning. We then show the results of transferring the representations to vision tasks including object detection and surface normal estimation with fine-tuning.

Implementation Details. To prepare the data for training, we download the 100K videos from YouTube using the URLs provided by [33, 56]. By mining the moving objects and tracking in the videos, we obtain ~ 10 million image patches of objects. By applying the transitivity on the graph constructed, we obtain 7 million positive pairs of objects where each pair of objects are two different instances with different viewpoints. We also randomly sample 2 million object pairs connected by the intra-instance edges.

We train our network with these 9 million pairs of images using a learning rate of 0.001 and a mini-batch size of 100. For each pair we sample the third distractor patch from a different “parent cluster” in the same mini-batch. We use the network pre-trained in [9] to initialize our convolutional layers and randomly initialized the fully connected layers. We train the network for 200K iterations with our method.

6.1. Qualitative Results without Fine-tuning

We first perform nearest-neighbor search to show qualitative results. We adopt the pool5 feature of the VGG16

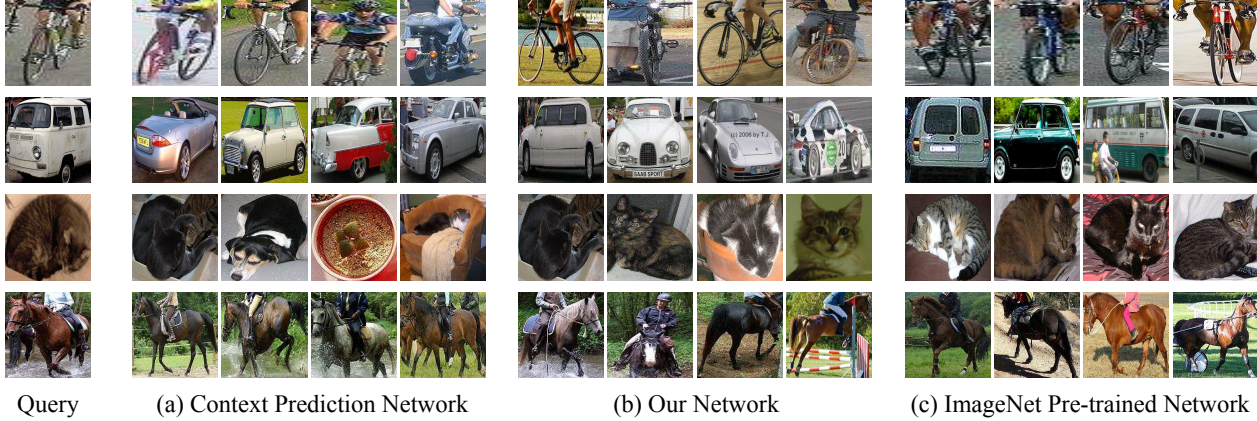


Figure 7: Nearest-neighbor search on the PASCAL VOC dataset. We extract three types of features: (a) context prediction network from [9], (b) network trained with our self-supervised method, and (c) the network pre-trained in the annotated ImageNet dataset. We show that our network can represent a greater variety (e.g., viewpoints) of objects of the same category.




Figure 8: Top 6 responses for neurons in 4 different convolutional units of our network, visualized using [59].

network for all methods without any fine-tuning (Figure 7). We do this experiment on the object instances cropped from the PASCAL VOC 2007 dataset [13] (`trainval`). As Figure 7 shows, given an query image on the left, the network pre-trained with the context prediction task [9] can retrieve objects of very similar viewpoints. On the other hand, our network shows more variations of objects and can often retrieve objects with the same class as the query. We also show the nearest-neighbor results using fully-supervised ImageNet pre-trained features as a comparison.

We also visualize the features using the visualization technique of [59]. For each convolutional unit in `conv5_3`, we retrieve the objects which give highest activation responses and highlight the receptive fields on the images. We visualize the top 6 images for 4 different convolutional units in Figure 8. We can see these convolutional units are corresponding to different semantic object parts (e.g., fronts of cars or buses wheels, animal legs, eyes or faces).

6.2. Analysis on Object Detection

We evaluate how well our representations can be transferred to object detection by fine-tuning Fast R-CNN [16] on PASCAL VOC 2007 [13]. We use the standard `trainval` set for training and `test` set for testing with VGG16 as the base architecture. For the detection network, we initialize the weights of convolutional layers from our self-supervised network and randomly initialize the fully-connected layers using Gaussian noise with zero mean and 0.001 standard deviation.

During fine-tuning Fast R-CNN, we use 0.00025 as the starting learning rate. We reduce the learning rate by 1/10 in every 50K iterations. We fine-tune the network for 150K iterations. Unlike standard Fast R-CNN where the first few convolutional layers of the ImageNet pre-trained network are fixed, we fine-tuned all layers on the PASCAL data as our model is pre-trained in a very different domain (e.g., video patches).

We report the results in Table 1. If we train Fast R-CNN *from scratch* without any pre-training, we can only obtain 39.7% mAP. With our self-supervised trained network as initialization, the detection mAP is increased to 63.2% (with a 23.5 points improvement). Our result compares competitively (4.1 points lower) to the counterpart using ImageNet pre-training (67.3% with VGG16).

As we incorporate the invariance captured from [56] and [9], we also evaluate the results using these two approaches individually (Table 1). By fine-tuning the context prediction network of [9], we can obtain 61.5% mAP. To train the network of [56], we use exactly the same loss function and initialization as our approach except that there are only training examples of the same instance in the same visual track (i.e., only the samples linked by intra-instance edges in our graph). Its results is 60.2% mAP. Our result (63.2%)

method	mAP	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	persn	plant	sheep	sofa	train	tv
from scratch	39.7	51.7	55.8	21.7	24.0	10.5	58.7	59.2	41.1	18.2	32.9	35.6	33.4	60.4	57.3	45.5	19.7	29.2	30.8	61.0	47.3
Vid-Edge [32]	44.2	54.4	58.2	39.6	30.8	12.5	58.7	61.9	51.0	22.0	41.4	47.4	41.5	63.2	58.4	47.5	17.2	27.6	45.4	59.8	45.4
Context [9]	61.5	70.8	72.1	54.7	49.7	31.0	72.3	76.9	70.8	44.6	61.1	59.8	67.0	74.6	72.5	68.3	29.4	58.5	66.9	75.1	54.3
Tracking [56]	60.2	65.7	73.2	55.4	46.4	30.9	74.0	76.9	67.8	40.9	58.0	60.9	65.0	74.1	71.6	67.1	31.5	55.0	61.8	73.9	53.8
Ours	63.2	68.4	74.6	57.1	49.6	34.1	73.5	76.9	73.2	45.8	63.3	66.3	68.6	74.9	74.2	69.5	31.9	57.4	70.3	75.9	59.3
ImageNet	67.3	74.4	78.0	65.9	54.4	39.7	76.4	78.6	82.5	48.6	73.3	67.2	78.4	77.3	75.7	72.2	32.2	65.8	66.8	75.2	62.4

Table 1: Object detection Average Precision (%) on the VOC 2007 test set using **Fast R-CNN** [16] (with selective search proposals [50]): comparisons among different self-supervised learning approaches.

method	mAP	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	persn	plant	sheep	sofa	train	tv
Ours	63.2	68.4	74.6	57.1	49.6	34.1	73.5	76.9	73.2	45.8	63.3	66.3	68.6	74.9	74.2	69.5	31.9	57.4	70.3	75.9	59.3
Multi-Task	62.1	70.0	74.2	57.2	48.4	33.0	73.6	77.6	70.7	45.0	61.5	64.8	67.2	74.0	72.9	68.3	32.4	56.6	64.1	74.1	57.5
Ours (15-frame)	61.5	70.3	74.1	53.3	47.1	33.5	74.6	77.1	67.7	43.3	58.1	65.5	65.8	75.2	72.2	67.6	31.6	55.5	65.6	74.6	57.2
Ours (HOG)	60.4	65.8	73.4	54.7	47.7	30.2	75.6	77.1	67.6	42.0	58.8	63.2	65.3	74.1	72.0	67.2	29.9	54.4	62.1	72.9	53.9

Table 2: More ablative studies on object detection on the VOC 2007 test set using **Fast R-CNN** [16] (with selective search proposals [50]).

is better than both methods. This comparison indicates the effectiveness of exploiting a greater variety of invariance in representation learning.

Is multi-task learning sufficient? An alternative way of obtaining both intra- and inter-instance invariance is to apply multi-task learning with the two losses of [9] and [56]. Next we compare with this method.

For the task in [56], we use the same network architecture as our approach; for the task in [9], we follow their design of a Siamese network. We apply different fully connected layers for different tasks, but share the convolutional layers between these two tasks. Given a mini-batch of training samples, we perform ranking among these images as well as context prediction in each image simultaneously via two losses. The representations learned in this way, when fine-tuned with Fast R-CNN, obtain 62.1% mAP (“Multi-task” in Table 2). Comparing to only using context prediction [9] (61.5%), the multi-task learning only gives a marginal improvement (0.6%). This result suggests that multi-task learning in this way is not sufficient; organizing and exploiting the relationships of data, as done by our method, is more effective for representation learning.

How important is tracking? To further understand how much visual tracking helps, we perform ablative analysis by making the visual tracks shorter: we track the moving objects for 15 frames instead of by default 30 frames. This is expected to reduce the viewpoint/pose/deformation variance contributed by tracking. Our model pre-trained in this way shows 61.5% mAP (“15-frame” in Table 2) when fine-tuned for detection. This number is similar to that of using context prediction only (Table 1). This result is not surprising, because it does not add much new information for training. It suggests adding stronger viewpoint/pose/deformation invariance is important for learning

	All	>c1	>c2	>c3	>c4	>c5
Context [9]	62.6	61.1	60.9	57.0	49.7	38.1
Tracking [56]	62.2	61.5	62.2	61.4	58.9	39.5
Multi-Task [9, 56]	62.4	63.2	63.5	62.9	58.7	27.6
Ours	65.0	64.5	63.6	60.4	55.7	43.1
ImageNet	70.9	71.1	71.1	70.2	70.3	64.3

Table 3: Object detection Average Precision (%) on the VOC 2007 test set using joint training **Faster R-CNN** [41].

better features for object detection.

How important is clustering? Furthermore, we want to understand how important it is to cluster images with features learned from still images [9]. We perform another ablative analysis by replacing the features of [9] with HOG [6] during clustering. The rest of the pipeline remains exactly the same. The final result is 60.4% mAP (“HOG” in Table 2). This shows that if the features for clustering are not invariant enough to handle different object instances, the transitivity in the graph becomes less reliable.

6.3. Object Detection with Faster R-CNN

Although Fast R-CNN [16] has been a popular testbed of un-/self-supervised features, it relies on Selective Search proposals [50] and thus is not fully end-to-end. We further evaluate the representations on object detection with the *end-to-end* Faster R-CNN [41] where the Region Proposal Network (RPN) may suffer from the features if they are low-quality.

PASCAL VOC 2007 Results. We fine-tune Faster R-CNN in 8 GPUs for 35K iterations with an initial learning rate of 0.00025 which is reduced by 1/10 after every 15K iterations. Table 3 shows the results of fine-tuning all

	AP	AP ⁵⁰	AP ⁷⁵	AP ^S	AP ^M	AP ^L
from scratch	20.5	40.1	19.0	5.6	22.5	32.7
Context [9]	22.7	43.5	21.2	6.6	24.9	36.5
Tracking [56]	22.6	42.8	21.6	6.3	25.0	36.2
Multi-Task [9, 56]	22.0	42.3	21.1	6.6	24.5	35.0
Ours	23.5	44.4	22.6	7.1	25.9	37.3
ImageNet (shorter)	23.7	44.5	23.5	7.2	26.9	37.4
ImageNet	24.4	46.4	23.1	7.9	27.4	38.1

Table 4: Object detection Average Precision (%; COCO definitions) on COCO minival using joint training **Faster R-CNN** [41]. “(shorter)” indicates a shorter training time (fewer iterations, 61.25K) used by the codebase of [41].

layers (“All”) and also ablative results on freezing different levels of convolutional layers (*e.g.*, the column $>c3$ represents freezing all the layers below and including conv3_x in VGG16 during fine-tuning). Our method gets even better results of 65.0% by using Faster R-CNN, showing a larger gap compared to the counterparts of [9] (62.6%) and [56] (62.2%). Noteworthy, when freezing all the convolutional layers and only fine-tuning the fully-connected layers, our method (43.1%) is much better than other competitors. And we again find that the multi-task alternative does not work well for Faster R-CNN.

COCO Results. We further report results on the challenging COCO detection dataset [34]. To the best of our knowledge this is the first work of this kind presented on COCO detection. We fine-tune Faster R-CNN in 8 GPUs for 120K iterations with an initial learning rate of 0.001 which is reduced by 1/10 after 80k iterations. This is trained on the COCO trainval35k split and evaluated on the minival5k split, introduced by [3].

We report the COCO results on Table 4. Faster R-CNN fine-tuned with our self-supervised network obtains 23.5% AP using the COCO metric, which is very close ($<1\%$) to fine-tuning Faster R-CNN with the ImageNet pre-trained counterpart (24.4%). Actually, if the fine-tuning of the ImageNet counterpart follows the “shorter” schedule in the public code (61.25K iterations in 8 GPUs, converted from 490K in 1 GPU)¹, the ImageNet supervised pre-training version has 23.7% AP and is comparable with ours. This comparison also strengthens the significance of our result.

To the best of our knowledge, our model achieves the best performance reported to date on VOC 2007 and COCO using un-/self-supervised pre-training.

6.4. Adapting to Surface Normal Estimation

To show the generalization ability of our self-supervised representations, we adopt the learned network to the surface normal estimation task. In this task, given a single

¹<https://github.com/rbgirshick/py-faster-rcnn>

	Mean	Median	11.25°	22.5°	30°
	(lower is better)		(higher is better)		
from scratch	31.3	25.3	24.2	45.6	56.8
Context [9]	29.0	21.6	28.8	51.5	61.9
Tracking [56]	27.8	21.8	27.4	51.1	62.5
Ours	26.0	18.0	33.9	57.6	67.5
ImageNet	27.8	21.2	29.0	52.3	63.4

Table 5: Results on NYU v2 for per-pixel surface normal estimation, evaluated over valid pixels.

RGB image as input, we train the network to predict the normal/orientation of the pixels. We evaluate our method on the NYUv2 RGBD dataset [44] dataset. We use the official split of 795 images for training and 654 images for testing. We follow the same protocols for generating surface normal ground truth and evaluations as [14, 28, 15].

To train the network for surface normal estimation, we apply the Fully Convolutional Network (FCN 32-s) proposed in [35] with the VGG16 network as base architecture. For the loss function, we follow the design in [55]. Specifically, instead of direct regression to obtain the normal, we use a codebook of 40 codewords to encode the 3-dimension normals. Each codeword represents one class thus we turn the problem into a 40-class classification for each pixel. We use the same hyperparameters as in [35] for training and the network is fine-tuned for same number of iterations (100K) for different initializations.

To initialize the FCN model with self-supervised nets, we copy the weights of the convolutional layers to the corresponding layers in FCN. For ImageNet pre-trained network, we follow [35] by converting the fully connected layers to convolutional layers and copy all the weights. For the model trained from scratch, we randomly initialize all the layers with “Xavier” initialization [18].

Table 5 shows the results. We report mean and median error for all visible pixels (in degrees) and also the percentage of pixels with error less than 11.25, 22.5 and 30 degrees. Surprisingly, we obtain much better results with our self-supervised trained network than ImageNet pre-training in this task (3 to 4% better in most metrics). As a comparison, the network trained in [9, 56] are slightly worse than the ImageNet pre-trained network. These results suggest that our learned representations are competitive to ImageNet pre-training for high-level semantic tasks, but outperforms it on tasks such as surface normal estimation. This experiment suggests that different visual tasks may prefer different levels of visual invariance.

Acknowledgement: This work was supported by ONR MURI N000141612007 and Sloan Fellowship to AG.

References

- [1] P. Agrawal, J. Carreira, and J. Malik. Learning to see by moving. 2015. [2](#)
- [2] H. Bay, T. Tuytelaars, and L. V. Gool. Surf: Speeded up robust features. In *ECCV*, 2006. [4](#)
- [3] S. Bell, C. L. Zitnick, K. Bala, and R. Girshick. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. *arXiv*, 2015. [8](#)
- [4] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy layer-wise training of deep networks. In *NIPS*, 2007. [2](#)
- [5] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005. [2](#)
- [6] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005. [1](#), [7](#)
- [7] C. Doersch, A. Gupta, and A. A. Efros. Mid-level visual element discovery as discriminative mode seeking. In *NIPS*, 2013. [2](#), [4](#)
- [8] C. Doersch, A. Gupta, and A. A. Efros. Context as supervisory signal: Discovering objects with predictable context. In *ECCV*, 2014. [2](#)
- [9] C. Doersch, A. Gupta, and A. A. Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [10] J. Donahue, P. Krähenbühl, and T. Darrell. Adversarial feature learning. *arXiv*, 2016. [2](#)
- [11] V. Dumoulin, I. Belghazi, B. Poole, A. Lamb, M. Arjovsky, O. Mastropietro, and A. Courville. Adversarially learned inference. *arXiv*, 2016. [2](#)
- [12] S. M. A. Eslami, N. Heess, and J. Winn. The shape boltzmann machine: a strong model of object shape. In *CVPR*, 2012. [2](#)
- [13] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The Pascal Visual Object Classes (VOC) Challenge. *IJCV*, 2010. [6](#)
- [14] D. F. Fouhey, A. Gupta, and M. Hebert. Data-driven 3D primitives for single image understanding. In *ICCV*, 2013. [8](#)
- [15] D. F. Fouhey, A. Gupta, and M. Hebert. Unfolding an indoor origami world. In *ECCV*, 2014. [8](#)
- [16] R. Girshick. Fast R-CNN. In *ICCV*, 2015. [6](#), [7](#)
- [17] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. [1](#)
- [18] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010. [8](#)
- [19] Y. Gong, Y. Jia, T. K. Leung, A. Toshev, and S. Ioffe. Deep convolutional ranking for multilabel image annotation. *arXiv*, 2013. [2](#)
- [20] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. [2](#)
- [21] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006. [2](#)
- [22] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista. High-speed tracking with kernelized correlation filters. *TPAMI*, 2015. [4](#)
- [23] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 2006. [2](#)
- [24] E. Hoffer and N. Ailon. Deep metric learning using triplet network. *arXiv*, 2015. [2](#)
- [25] D. Jayaraman and K. Grauman. Learning image representations tied to egomotion. In *ICCV*, 2015. [2](#)
- [26] D. Kingma and M. Welling. Auto-encoding variational bayes. In *ICLR*, 2014. [2](#)
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. [1](#)
- [28] L. Ladický, B. Zeisl, and M. Pollefeys. Discriminatively trained dense surface normal estimation. In *ECCV*, 2014. [8](#)
- [29] Q. V. Le, M. A. Ranzato, R. Monga, M. Devin, K. Chen, G. S. Corrado, J. Dean, and A. Y. Ng. Building high-level features using large scale unsupervised learning. In *ICML*, 2012. [2](#)
- [30] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1989. [1](#)
- [31] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *ICML*, 2009. [2](#)
- [32] Y. Li, M. Paluri, J. M. Rehg, and P. Dollár. Unsupervised learning of edges. In *CVPR*, 2016. [2](#), [7](#)
- [33] X. Liang, S. Liu, Y. Wei, L. Liu, L. Lin, and S. Yan. Computational baby learning. *arXiv*, 2014. [5](#)
- [34] T. Lin, M. Maire, S. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *arXiv*, 2014. [2](#), [8](#)
- [35] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. [1](#), [8](#)
- [36] D. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *IJCV*, 2004. [1](#)
- [37] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv*, 2015. [2](#)
- [38] B. A. Olshausen and D. J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997. [2](#)
- [39] D. Pathak, R. Girshick, P. Dollár, T. Darrell, and B. Hariharan. Learning features by watching objects move. In *CVPR*, 2017. [2](#)
- [40] L. Pinto and A. Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *ICRA*, 2016. [2](#)
- [41] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. [7](#), [8](#)
- [42] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge. *IJCV*, 2015. [1](#)
- [43] B. C. Russell, A. A. Efros, J. Sivic, W. T. Freeman, and A. Zisserman. Using multiple segmentations to discover objects and their extent in image collections. In *CVPR*, 2006. [2](#)
- [44] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor

- segmentation and support inference from RGBD images. In *ECCV*, 2012. 8
- [45] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv*, 2014. 2, 3, 5
- [46] S. Singh, A. Gupta, and A. A. Efros. Unsupervised discovery of mid-level discriminative patches. In *ECCV*, 2012. 2, 4
- [47] J. Sivic, B. C. Russell, A. A. Efros, A. Zisserman, and W. T. Freeman. Discovering objects and their location in images. In *ICCV*, 2005. 2
- [48] N. Srivastava, E. Mansimov, and R. Salakhutdinov. Unsupervised learning of video representations using LSTMs. *arXiv*, 2015. 2
- [49] Y. Tang, R. Salakhutdinov, and G. Hinton. Robust boltzmann machines for recognition and denoising. In *CVPR*, 2012. 2
- [50] J. Uijlings, K. van de Sande, T. Gevers, and A. Smeulders. Selective search for object recognition. *IJCV*, 2013. 7
- [51] P. Vincent, H. Larochelle, Y. Bengio, and P. Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, 2008. 2
- [52] J. Walker, C. Doersch, A. Gupta, and M. Hebert. An uncertain future: Forecasting from variational autoencoders. In *ECCV*, 2016. 2
- [53] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, 2013. 4
- [54] J. Wang, Y. Song, T. Leung, C. Rosenberg, J. Wang, J. Philbin, B. Chen, and Y. Wu. Learning fine-grained image similarity with deep ranking. In *CVPR*, 2014. 2, 5
- [55] X. Wang, D. F. Fouhey, and A. Gupta. Designing deep networks for surface normal estimation. In *CVPR*, 2015. 8
- [56] X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. In *ICCV*, 2015. 1, 2, 3, 4, 5, 6, 7, 8
- [57] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In *ECCV*, 2016. 2
- [58] R. Zhang, P. Isola, and A. A. Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, 2017. 2
- [59] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Object detectors emerge in deep scene cnns. In *ICLR*, 2015. 6