

# Exploring Patch-wise Semantic Relation for Contrastive Learning in Image-to-Image Translation Tasks

Chanyong Jung<sup>\*1</sup>Gihyun Kwon<sup>\*1</sup>Jong Chul Ye<sup>2,1</sup>Department of Bio and Brain Engineering<sup>1</sup>, Kim Jaechul Graduate School of AI<sup>2</sup>, KAIST

{jcy132, cyclomon, jong.ye}@kaist.ac.kr

## Abstract

Recently, contrastive learning-based image translation methods have been proposed, which contrasts different spatial locations to enhance the spatial correspondence. However, the methods often ignore the diverse semantic relation within the images. To address this, here we propose a novel semantic relation consistency (SRC) regularization along with the decoupled contrastive learning, which utilize the diverse semantics by focusing on the heterogeneous semantics between the image patches of a single image. To further improve the performance, we present a hard negative mining by exploiting the semantic relation. We verified our method for three tasks: single-modal and multi-modal image translations, and GAN compression task for image translation. Experimental results confirmed the state-of-art performance of our method in all the three tasks.

## 1. Introduction

One of the main objectives of image-to-image translation tasks is to learn a mapping function from a source domain to a target domain so that it preserves contents while converting the appearance similar to target domain. The cycle consistency loss, which enforces the consistency between the input image and reconstructed image by an inverse mapping of the converted image, is widely used in various frameworks [8, 14, 19, 36]. However, it requires an additional generator and discriminator to learn the inverse mapping. Also, [15, 20, 26] claimed that the cycle-consistency constraint may produce distortion due to its overly restrictive constraint. Thus, one-sided image translation methods are suggested to bypass the cycle-consistency constraint by enhancing the correspondence between the input and output in various ways [1, 4, 9].

Recently, inspired by the success of contrastive learning, CUT [20] firstly utilized the contrastive learning to maxi-

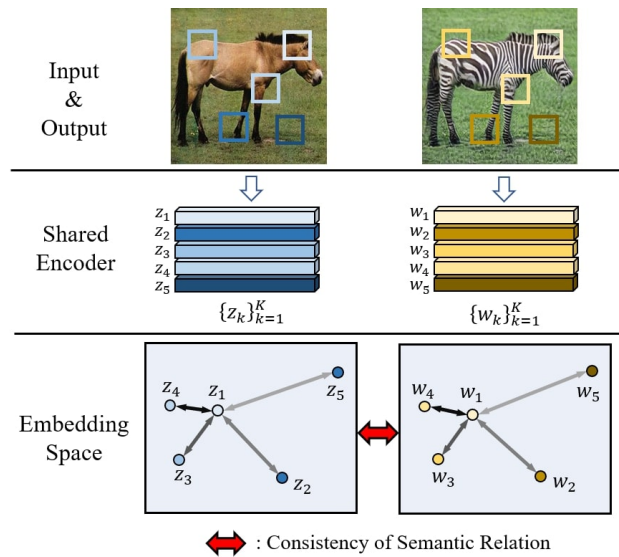


Figure 1. Concept of the proposed method. Decoupled contrastive learning forms the embedding. Consistency regularization of diverse semantic relation is imposed to enhance the correspondence between the input and output.

mize the mutual information between the same location of input and output images. The authors of NEG CUT [26] proposed a contrastive learning method using hard negative samples generated by the negative generator. However, the method requires additional training procedure for the negative generator from random vectors, which may not guarantee to follow the real negative sample distribution and cause the instability of training. F-LSesim [34] utilized the patch-wise similarity map for the contrastive learning, but they ignored the semantic relation between the patches, taking all negative patches as equal negative.

In this paper, we propose a novel contrastive learning method to utilize the heterogeneous semantics within an image. Specifically, as shown in Fig. 1, the key idea is to impose the consistency regularization on the similarity relation that preserves the spatial semantic relation during the

<sup>\*</sup>Co-first authors.

image translation tasks. Specifically, we capture the patch-wise semantic relationship with an image in terms of distributional similarity and enforce it to be preserved during the image translation tasks. This semantic regularization prevents from generating image artifacts which violate semantic relationship. Furthermore, we propose a hard negative mining strategy based on the spatially varying semantic relations between image patches. This strategy further improves the performance by avoiding sampling ‘easy’ negative samples which have unrelated semantic information, but more focusing on ‘hard’ negative samples.

Experimental results using single- and multi-modal translation task and GAN compression confirmed that our method produces the state-of-the-art (SOTA) performance thanks to its capability of utilizing semantic relationship.

## 2. Related works

### 2.1. One-sided image translation

In order to replace the cycle-consistency, many one-sided image translation methods utilize the relational knowledge and correspondency between the input and translated images. For example, GcGAN [9] utilized the consistency for the geometric transformation of images, and DistanceGAN [4] imposed the consistency regularization for the mutual information within a group of images. The approach in TraVeLGAN [1] preserved the arithmetic property of embedding vectors.

Many contrastive learning methods have been recently suggested to maximize the mutual information between the same location of input and output images [20,26,34]. However, each method has major shortcomings as described before. In fact, our method is designed to overcome the limitation of the existing works by exploiting on the relational knowledge transfer as described below.

### 2.2. Relational knowledge transfer

The relational knowledge that captures the structured interdependencies between the data is useful for knowledge distillation (KD) tasks [21,25,29,35]. In particular, the student model additionally leverages the inter-sample relation learned by the teacher model, so the knowledge transfer is proceeded in a more effective way.

Various methods have been proposed to capture the relational knowledge, such as the angle-wise and distance-wise relation [21], instance-wise correlation within the feature space [22], contrastive relation [29], learned relation for features and gradients by a network [35], etc.

Accordingly, we are interested in utilizing the semantic relational knowledge to enhance the input and output correspondence for the image translation tasks. As many previous KD approaches transferred the relational knowledge to enhance the correspondence between student features and

teacher features, we match the semantic relational knowledge to enhance the correspondence between input features and output features for image translation tasks.

Specifically, our method focused on the heterogeneous semantic property of the image patches, which was not considered in the previous works [20,34]. Also, in contrast to NEG CUT [26], the negative samples are obtained from the ‘real’ data distribution and a hardness of the negative mining can be controlled by curriculum learning scheme.

### 2.3. Contrastive learning

Contrastive learning method is a framework which obtains useful representation by using the relation between the positive and negative pairs. The InfoNCE is widely used for the contrastive learning in many previous works [3,6,7,11,24].

Recently, various approaches are suggested to further improve the performance. One approach is to sample informative negatives using hard negative mining by von Mises-Fisher distribution [23], adversarial training [13] and learning to rank the samples [27]. Another approach is to improve the infoNCE loss. For example, FlatNCE [5] is proposed to alleviate the degradation caused by small number of negatives. Decoupled infoNCE loss [31] is suggested as the remedy for a negative-positive coupling (NPC) effect, which results the equivalent loss function of FlatNCE.

As it is undesirable to treat the negatives equally regardless of their semantic relation, some works focused on the heterogeneous semantic relation of negative samples. For example, CO2 [28] proposed the consistency regularization for positive pairs to have same semantic relation with the negatives. PCL [17] utilized EM algorithm to encode the semantic structures of a data.

## 3. Main Contribution

In natural images, image patches from various locations have heterogeneous semantic relation. For example, in Fig. 2, some patches are from horse, while some others represents the background which are unrelated to the horse. Also, even the patches of the horse have diverse semantic information as they represents each part of the horse such as head and legs. The diverse semantic relation should be considered and preserved for accurate image translation tasks.

In the following subsections, we present our method that utilizes the relational knowledge formed by the patch-wise heterogeneous semantics.

### 3.1. Consistency of semantic relation distribution

Let  $z_k$  and  $w_k$  denote the embedded vectors of image patch  $x_k \in \mathcal{X}$  and  $y_k \in \mathcal{Y}$ , where  $\mathcal{X}$  and  $\mathcal{Y}$  are input image and translation output image, respectively. The indices  $k$  refer to the patch location.

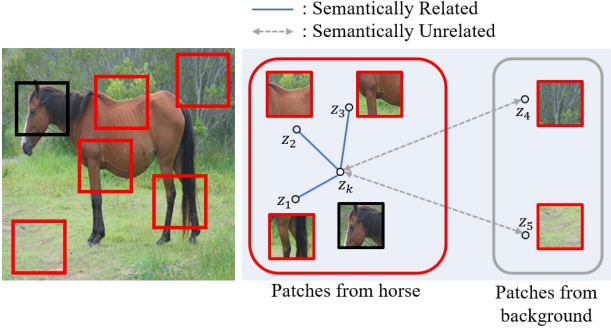


Figure 2. Heterogeneous semantic relation between patches. The black patch is the query and red patches are negatives. Patches from horse are semantically related, but background patches are not.

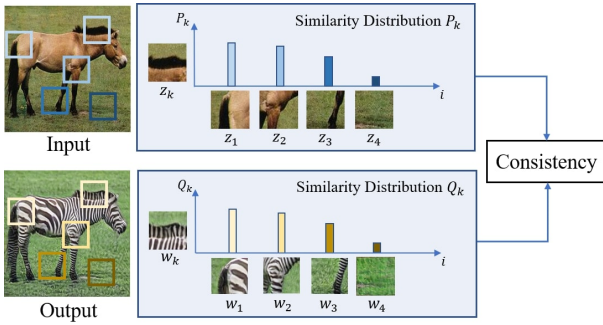


Figure 3. Semantic relation through similarity distribution.

From the perspectives that the contrastive learning is an instance classification [16, 17, 28, 31] (i.e. assigning label 1 for positive sample and 0 for negatives), the distribution of similarity can be viewed as a soft label which reveals the structure for semantic relation between the samples [28, 30]. We capture the patch-wise semantic relation as a distribution of similarities, as shown in Fig. 3.

Accordingly, for a given patch  $x_k$  of the input image, the similarity relation with the negative patch  $x_i$  is defined using the soft-max:

$$P_k(i) = \frac{\exp(z_k^\top z_i)}{\sum_{j=1}^K \exp(z_k^\top z_j)} \quad (1)$$

where the  $z_k$  and  $z_i$  are the corresponding embedding vectors. Then,  $P_k(i)$  is the distribution to capture the semantic closeness between the  $i$ -th location and  $k$ -th location patches within the input image. Similarly, the distribution  $Q_k(i)$  for the output image similarity relationship is defined as:

$$Q_k(i) = \frac{\exp(w_k^\top w_i)}{\sum_{j=1}^K \exp(w_k^\top w_j)} \quad (2)$$

where the  $w_i$  and  $w_k$  are the embedding vector for the corresponding patches.

To preserve the diverse semantic relation between the patches before and after the image translation, we impose the consistency regularization on the similarity relation for all  $K$  sampled vectors by Jensen-Shannon Divergence (JSD) [28]:

$$L_{SRC} = \sum_{k=1}^K JSD(P_k || Q_k) \quad (3)$$

The minimization of (3) therefore enforces the consistency of the diverse semantic relationship.

### 3.2. Hard negative mining by semantic relation

In this section, we explain how the semantic relation can be utilized for a hard negative mining. We design the distribution for the hard negative mining and derive the connection with the similarity relation. We used decoupled infoNCE loss (DCE) [31] for a patch-wise contrastive loss to prevent the negative-positive coupling (NPC) effect which is discussed in detail in the Section 3.3.

Specifically, for a given positive pair  $(z, w) \sim p_{ZW}$ , the hard negative contrastive loss  $L_{hDCE}$  using DCE [31] with  $N$  negatives and temperature parameter  $\tau$  is defined as

$$L_{hDCE}(\gamma, \tau) = \mathbb{E}_{(z, w) \sim p_{ZW}} \left[ -\log \frac{\exp(w^\top z / \tau)}{N \mathbb{E}_{z^- \sim q_{Z^-}} [\exp(w^\top z^- / \tau)]} \right] \quad (4)$$

where the negative sampling is modeled by von Mises-Fisher distribution [23]:

$$z^- \sim q_{Z^-}(z^-; z, \gamma) = \frac{1}{N_q} \exp\{\gamma(z^\top z^-)\} p_Z(z^-) \quad (5)$$

where  $N_q$  is a normalization constant and  $\gamma$  is a hyperparameter to determine the hardness of the negative samples.

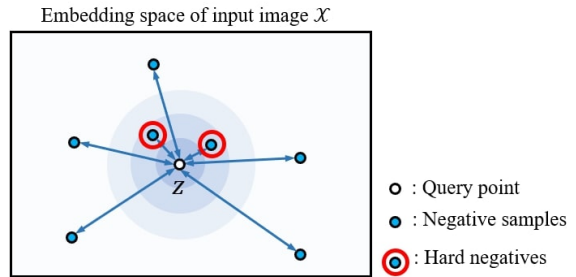


Figure 4. Hard negative mining with semantic relation. The color represents the probability to be sampled.

Accordingly, the designed distribution  $q_{Z^-}$  assigns higher probability to be sampled for the negatives  $z^-$  which

are closer to the  $z$ , as shown in Fig. 4. By employing this sampling strategy, we can avoid ‘easy’ negative samples that do not contribute the contrastive mechanism significantly. Instead, by allowing contrastive mechanism between the positive and ‘hard’ negatives, the network can be trained to have more discriminative power.

In terms of implementation,  $\mathbb{E}_{z^- \sim q_{z^-}} [\exp(w^\top z^- / \tau)]$  is approximated with the importance sampling [23]:

$$\begin{aligned} & \mathbb{E}_{z^- \sim q_{z^-}} [\exp(w^\top z^- / \tau)] \\ &= \mathbb{E}_{z^- \sim p_Z} \left[ \exp(w^\top z^- / \tau) \cdot \frac{q_{Z^-}(z^-)}{p_Z(z^-)} \right] \\ &= N_q^{-1} \mathbb{E}_{z^- \sim p_Z} [\exp(w^\top z^- / \tau) \cdot \exp\{\gamma(z^\top z^-)\}] \end{aligned}$$

Again, the similarity relation  $\exp\{\gamma(z^\top z^-)\}$  becomes the weight of importance sampling for the hard negative mining. Finally, curriculum learning is implemented, controlling the hardness of negatives with  $\gamma$ . Since the semantic relation is unstable at the early of the training,  $\gamma$  is initialized as small value. Then, we gradually increase  $\gamma$  to sample harder negatives as the training proceeds.

The explicit control of the hardness for the hard negative mining improved the performance, compared to the NEG-CUT [26] which used implicit control of the hardness by the negative generator. In the Supplementary Material, we present the details for the experimental settings and the effect of  $\gamma$  to the performance.

In summary, the total loss considering the patch-wise heterogeneous semantic relation is given by:

$$L_{semantic}(\gamma, \tau) = \lambda_{SRC} L_{SRC} + \lambda_{hDCE} L_{hDCE}(\gamma, \tau)$$

where  $\lambda_{SRC}$  and  $\lambda_{hDCE}$  are weighting parameters.

### 3.3. Effect of DCE for heterogeneous semantics

Decoupled contrastive learning using DCE [31] is proposed to remove the negative-positive coupling (NPC) effect. Specifically, the NPC effect refers the diminishing gradient of the InfoNCE by the easy negatives and positive samples, hindering the update by the other informative samples.

Likewise, for the InfoNCE based image translation methods [20, 26, 34], the NPC effect is problematic due to the heterogeneous semantic relation between the image patches. Since we randomly sample the negatives, semantically unrelated image patches can be included as the negatives. For example,  $z_4$  and  $z_5$  in the Fig. 2 are unrelated easy negatives which cause the NPC effect. To prevent the NPC effect, as discussed before, decoupled InfoNCE (DCE) loss should be employed in this paper. This section is dedicated to provide explicit explanation why NPC effect in the classical InfoNCE loss is detrimental for image translation, and why DCE is a better alternative for image translation.

Specifically, for a positive pair  $(z_k, w_k)$ , the DCE loss is defined as:

$$L_{DCE} = -\log \frac{\exp(z_k^\top w_k / \tau)}{\sum_{j \neq k} \exp(z_j^\top w_k / \tau)} \quad (6)$$

which removes the positive pair term in the denominator from the InfoNCE loss:

$$\begin{aligned} & L_{\text{InfoNCE}} \\ &= -\log \frac{\exp(z_k^\top w_k / \tau)}{\exp(z_k^\top w_k / \tau) + \sum_{j \neq k} \exp(z_j^\top w_k / \tau)} \quad (7) \end{aligned}$$

As discussed in [31], the gradient from the loss function  $\nabla_{w_k} L$  for the contrastive loss  $L$  in (6) and (7) is given as:

$$\nabla_{w_k} L = -\frac{\alpha}{\tau} \left\{ z_k - \sum_{j \neq k} \frac{\exp(z_j^\top w_k / \tau)}{\sum_{m \neq k} \exp(z_m^\top w_k / \tau)} \cdot z_j \right\}$$

where

$$\alpha := \begin{cases} q_{NPC} & \text{if } L = L_{\text{InfoNCE}} \\ 1 & \text{if } L = L_{DCE} \end{cases} \quad (8)$$

and

$$q_{NPC} \simeq 1 - \frac{\exp(z_k^\top w_k / \tau)}{\exp(z_k^\top w_k / \tau) + \sum_{j \neq k} \exp(z_j^\top z_k / \tau)} \quad (9)$$

where we use  $\|z_j - w_j\|_2^2 = 2 - 2z_j^\top w_j$  for the normalized feature vectors.

If  $z_j$  for  $j \neq k$  has unrelated semantic information with  $z_k$ ,  $z_j$  will have a small value of  $\exp(z_j^\top z_k / \tau)$ . Hence, the denominator in (9) is dominated by  $\exp(z_k^\top w_k / \tau)$ , decreasing the value of  $q_{NPC}$  and  $\alpha$  through (8), which again results the diminishing gradient for InfoNCE loss. However, in case of DCE, the gradients are not related to the  $q_{NPC}$  as  $\alpha = 1$ , which prevents the gradient vanishing by the easy negatives. Therefore, DCE loss is employed in our method to prevent the NPC effect caused by the semantically unrelated negative samples.

## 4. Experiments

To show the versatility of our model, we apply our proposed method to various frameworks for image-to-image translation (I2I) tasks, which includes single-modal and multi-modal image translation, and I2I GAN model compression. The brief introduction of our experiment is in Fig. 5. Additionally, we provide more results for single-image translation task (i.e. *painting*  $\rightarrow$  *photo*) in the Supplementary Material.






Figure 5. Experiment schematics for (a) single-modal translation, (b) multi-modal translation, and (c) model compression.  $F$  consists of sub-networks  $F^l$  which is 2-layer MLP for the  $l$ -th layer. 256 vectors are randomly sampled (i.e  $K = 256$ ).

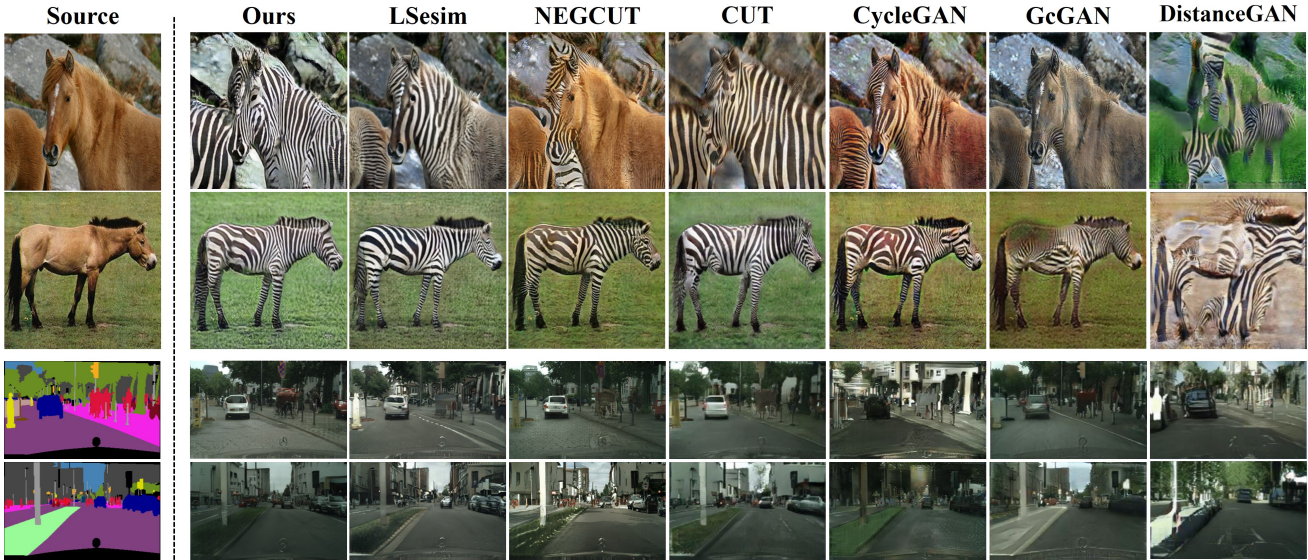


Figure 6. Qualitative comparison on single-modal image translation. Our method shows cleaner and semantically realistic outputs compared with those of baseline methods.

#### 4.1. Single modal image translation

First, we evaluate our method on unpaired image translation task which is designed to translate input images to a single output target domain.

Methods	H→Z		Cityscapes		
	FID↓	mAP↑	pAcc↑	cAcc↑	FID↓
CycleGAN [36]	77.2	20.4	55.9	25.4	76.3
MUNIT [14]	133.8	16.9	56.5	22.5	91.4
DRIT [15]	140.0	17.0	58.7	22.2	155.3
Distance [4]	72.0	8.4	42.2	12.6	81.8
GcGAN [9]	86.7	21.2	63.2	26.6	105.2
CUT [20]	45.5	24.7	68.8	30.7	56.4
NEG CUT [26]	39.6	27.6	71.4	35.0	48.5
LSeSIM* [34]	38.0	-	73.2	-	49.7
OURS	<b>34.4</b>	<b>29.0</b>	<b>73.5</b>	<b>35.6</b>	<b>46.4</b>

\* LSeSim did not report mAP&cAcc

Table 1. Quantitative comparison of single-modal image translation. H→Z refers to Horse→Zebra dataset. Our method outperforms baseline models.

**Experiment details:** Our implementation of unpaired translation model is based on source code of recent image translation model of CUT<sup>1</sup>. We evaluate our method on datasets of horse→zebra and cityscapes label→images. In the framework of CUT [20], we replace the PatchNCE loss with our proposed loss ( $L_{semantic}$ ). Detailed settings are provided in our Supplementary Materials.

**Evaluation metrics:** For evaluation, we calculate the quantitative metrics used in the previous works. For the models trained on horse→zebra dataset, we use Frchet Inception Distance (FID) [12] to evaluate the translated image quality. In the case of cityscapes (label→image), we additionally measure the correspondence between the segmented maps of the outputs and their ground truth maps. Specifically, we use pre-trained model of DRN [32] for segmentation, and calculate mean average precision recall (mAP), pixel-wise accuracy (pixAcc), and average class accuracy (clasAcc) [20].

<sup>1</sup><https://github.com/taesungp/contrastive-unpaired-translation>

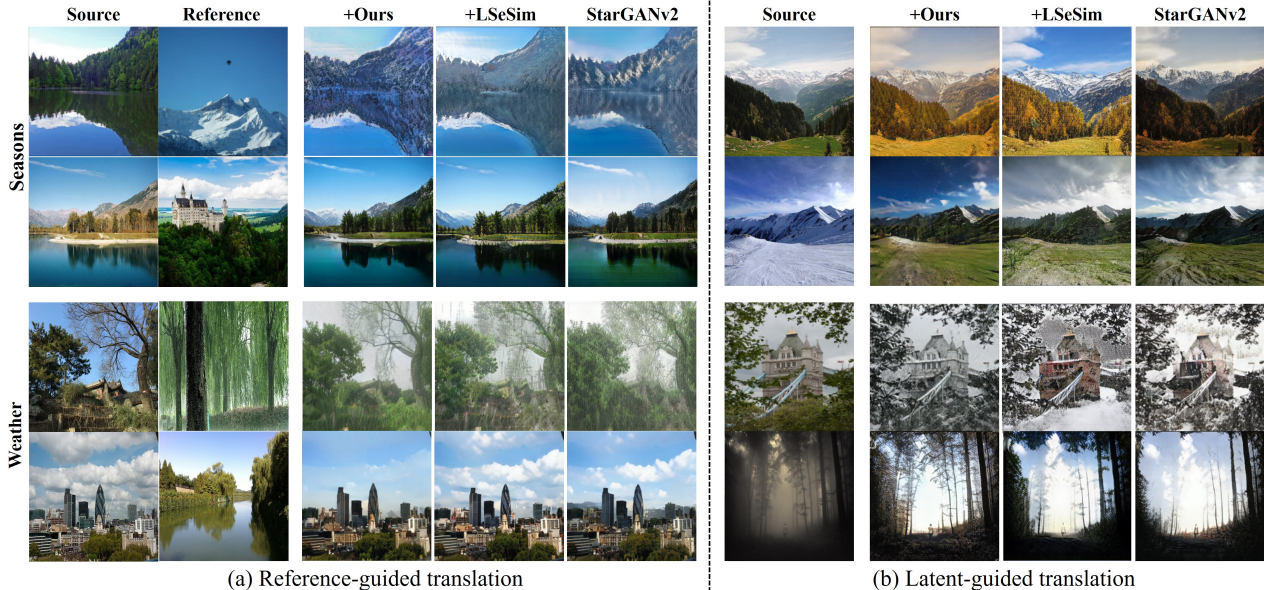


Figure 7. Qualitative comparison on multi-modal image translation. Our method can translate source images to arbitrary domains with diverse styles. (a) Reference-guided translation: source images are translated to target domains with reflecting reference style. (b) Latent-guided translation: each method translate source image to target domain with random styles. The rows from the top show translated outputs in the following order: summer2autumn, winter2spring, cloudy2snowy, and foggy2sunny.

**Results:** Table 1 shows the quantitative comparison between the proposed method and the previous works. In all of the datasets and metrics, our method outperforms the baseline methods. Specifically, our method outperformed the existing models with large margin when compared with one-sided image translation methods such as DistanceGAN [4] and GcGAN [9] and two-sided models such as CycleGAN [36], MUNIT [14], and DRIT [15]. Comparing to recent methods which use contrastive learning (CUT [20], NEGCUT [26], LSeSim [34]), we also achieved better performance in overall metrics.

For qualitative comparison, we show the results in Fig. 6. When we compare the results with those of the baselines, in the case of horse→zebra, the images generated by our method fully reflect the semantic texture of the target domain while the spatial structure of the input source is well preserved. In the experiments on cityscapes, the baselines often produce class mismatch between the input labels and the outputs. On the other hand, our model generates realistic images with right correspondence to the labels.

## 4.2. Multi-modal image translation

For further evaluation, we apply our method to multi-modal image translation model which is a framework for translating input to diverse outputs with multiple domains.

**Experiment details:** Our method is implemented using the official source code<sup>2</sup> of the state-of-the-art diverse transla-

<sup>2</sup><https://github.com/clovaai/stargan-v2>

Dataset	Method	Latent		Reference	
		FID↓	LPIPS↑	FID↓	LPIPS↑
Seasons	StarGANv2 [8]	63.06	0.413	61.19	0.346
	+LSeSim [34]	61.50	0.378	60.40	0.302
	+Ours	<b>54.70</b>	<b>0.496</b>	<b>54.23</b>	<b>0.365</b>
Weather	StarGANv2 [8]	62.45	0.415	64.20	0.342
	+LSeSim [34]	60.07	0.335	62.17	0.286
	+Ours	<b>54.02</b>	<b>0.470</b>	<b>56.91</b>	<b>0.362</b>

Table 2. Quantitative comparison of multi-modal image translation.

tion model StarGANv2 [8]. Specifically, on top of the loss functions of basic StarGANv2, we add our proposed loss ( $L_{semantic}$ ). Similar to single-modal translation, we calculate losses by sampling the embedded vectors from the input and output features (Fig. 5(b)). For the comparison with LSeSim [34], we implemented the StarGANv2 model with LSeSim following the official source code<sup>3</sup>.

**Dataset:** We choose two large-scale multi-domain datasets : *Seasons* [2] and *Weather* [10]. *Seasons* dataset consists of 4 domains (spring, summer, autumn, winter), each domain has 1208,1322,1460,and 1055 images. *Weather* dataset consists of 5 domains (sunny, cloudy, rainy, foggy, snowy) of which each domain has 4,000 images. As validation sets, we choose 100 images per each domain for *Seasons* dataset, and 400 images per each domain for *Weather* dataset.

**Evaluation metrics:** We measure the image quality with FID and generation diversity through LPIPS [33]. We fol-

<sup>3</sup><https://github.com/lyndonzheng/F-LSeSim>



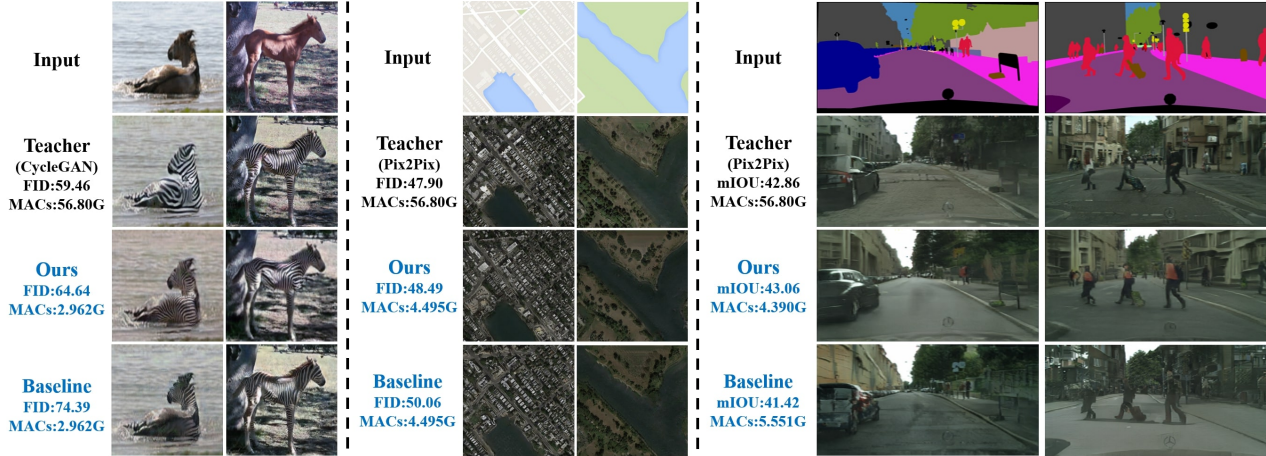


Figure 8. Qualitative comparison on model compression. (left) horse→zebra, (center) map→satellite aerial, and (right) cityscapes label→image. Our method shows high performances with small network size.

low the evaluation scenarios of StarGANv2: 1) latent-based image translation, which converts the style of input image to a random style, and 2) reference-based image translation, in which we convert the styles of inputs to those of the reference images. To calculate metrics, we first generate 10 different outputs per single input. We calculate LPIPS distance for 40 samples, and the average score is obtained by repeating the process for all generated images. We also calculate FID between generated target domain outputs and the training images of the corresponding domain. We report the averaged FID score for all domain cases.

**Results:** We show the quantitative comparison results in Table 2. Our method produced superior performance for all metrics. More specifically, the model trained with SeSim has improved performance in image quality with lower FID scores, but it shows weakness in terms of diversity with decreased LPIPS scores. However, when our method is used, both of image quality and diversity are improved in both datasets with a large margin.

Fig. 7 shows qualitative results on two different datasets. We compared the results of images translated into various domains. In case of the baseline models, the structure correspondence between input and output is not maintained, and some artifacts are included in the output as the information of target domain is not fully reflected. On the other hand, the result of our method better preserved the spatial information of input source, showing the successful translation with diverse styles.

### 4.3. Compression of image translation model

In recent model distillation methods, the delivery of the relational knowledge learned by the teacher model improves the student model, providing the additional information. We apply our method for a GAN compression framework, transferring the patch-wise semantic relational knowledge from the teacher to student. We demonstrate the

Dataset	Model	Metrics			
		#Param↓	MACs↓	FID↓	mIoU↑
H→Z	Teacher	11.38M	56.80	59.46	-
	Baseline	0.412M	2.962	74.39	-
	+Ours	0.412M	2.962	<b>64.64</b>	-
M→S	Teacher	11.38M	56.80	47.90	-
	Baseline	0.667M	4.495	50.06	-
	+Ours	0.667M	4.495	<b>48.49</b>	-
city	Teacher	11.38M	56.80	60.38	42.86
	Baseline	0.730M	5.551	85.24	41.42
	+Ours	<b>0.685M</b>	<b>4.390</b>	<b>72.41</b>	<b>43.06</b>

Table 3. Quantitative comparison of model compression. H→Z refers to Horse→Zebra dataset, M→S refers to Map→Satellite dataset, and city refers to cityscape label→image dataset. Our method outperforms baseline method.

effectiveness of our method for the GAN compression.

**Experiment details:** For the experiment, we adopt Fast GAN Compression [18], a SOTA compression method for image translation models. The method consists of two steps. In step 1, we first train the student network (a.k.a super-net) by distilling the feature of pre-trained teacher generator to student generator. In this step, the student model is a once-for-all network which supports different channel numbers. In step 2, using evolution search, we find the optimal channel number configurations of the student based on the target metric (e.g FID, mIoU). Experimental details are in the Supplementary Material.

We include our proposed method in step 1 of the model distillation part. Specifically, in Fig. 5(c), we introduce header net  $F$  and matched the semantic relation between the embedded features of the teacher and the student. The proposed loss ( $L_{semantic}$ ) is used along with the loss functions of the baseline framework of Fast GAN Compression [18]. The details of experimental setting are in the Supplementary Material. For implementation, we referenced the original

source code of GAN compression-fast version<sup>4</sup>.

**Dataset and models:** To evaluate the performance, we conduct experiments on 3 different settings : compression of CycleGAN trained on Horse→Zebra, Pix2Pix trained on Map→Satellite aerial, Pix2Pix trained on Cityscapes label→images. We retrained the baseline models using the official code for fair comparison.

**Evaluation metrics:** To compare the generation quality of the compressed model, we measure FID values. In the case of Cityscape, mIoU between output and GT is additionally measured. We also compare the compression performance with calculating multiply-accumulate operations (MACs) and the number of model parameters (#Param).

**Results:** Table 3 shows the quantitative performance of model compression. The results show that our model has better FID and mIoU scores while having similar or less MACs and #Param compared to the baseline student model. Particularly, in case of cityscapes dataset, we obtain better mIoU than teacher model. This suggests that our proposed method improves the student model, and additionally delivers the relational knowledge of teacher.

Fig. 8 is provided for qualitative comparison. Although our compressed model size is much smaller than the teacher model, the quality of the generated images is not worsened. Comparing with the baseline outputs, we can see that the output images have clear boundaries and better target domain textures.

#### 4.4. Ablations studies

We compare the quantitative performance of models trained on different settings using Horse→Zebra and cityscapes datasets for single-modal image translation task.

Specifically, we progressively add the components of our method, and observe the corresponding performance. More specifically, we first start from the basic framework using infoNCE, which is identical to CUT [20]. Then we compare the infoNCE with DCE, sequentially including semantic relation consistency (SRC) and hard negative mining (Hneg).

The results in Table 4 shows the meaningful improvement by our method. Specifically, when each component is added to DCE, we observe the improvement in both datasets: Horse→Zebra and cityscapes. When all the components are added, the results showed the best performance.

In addition, in case of InfoNCE, the SRC loss and the hard negative mining contributes to the improvement. However, compared with our best model with DCE, the results with the InfoNCE show worse performance in all metrics.

To further show the effect of proposed methods, we show the similarity relation in Fig. 9. The learned similarity between the query point and the other locations is calculated and mapped for input image (A) and translation output (B). In case of InfoNCE and basic DCE, high similarity is shown




Figure 9. Similarity maps of models trained on horse→zebra datasets. Red dot at the front leg is the query point.

Settings				H→Z		Cityscapes		
InfoNCE	DCE	SRC Loss	Hard Neg Mining	FID↓	mAP↑	pAcc↑	cAcc↑	FID↓
✓	×	×	×	45.5	24.7	68.8	30.7	56.4
✓	×	×	×	42.7	27.2	73.0	33.1	52.7
✓	×	✓	✓	40.5	27.4	72.8	33.5	49.9
×	✓	×	×	41.2	27.8	72.8	33.4	52.0
×	✓	×	✓	37.6	27.3	71.4	33.4	50.1
×	✓	×	×	36.2	27.9	73.3	34.2	49.9
×	✓	✓	✓	<b>34.4</b>	<b>29.0</b>	<b>73.5</b>	<b>35.6</b>	<b>46.4</b>

Table 4. Quantitative results of ablation studies.

at the points not related to the query (especially severe for infoNCE), and (A) and (B) show different similarity patterns. When SRC is used, the similarity aspect of (A) and (B) become similar. When the hard negative mining is further applied, only the points that are closely related to the query point are embedded with high similarity. It shows that we achieve our goals: consistency of patch-wise semantic relation between (A) and (B) and get advantages from hard negatives.

More ablation studies and additional results on our three different tasks are in Supplementary Materials.

## 5. Conclusion

In this paper, we proposed a novel method utilizing the relational knowledge formed by the heterogeneous semantics of image patches for image translation tasks. We presented the decoupled infoNCE with hard negatives (hDCE) along with the regularization of semantic relation consistency (SRC). We verified our methods by various tasks, which are image translation tasks and GAN compression. The results have shown noticeable performance increase, achieving state-of-the-art scores compared to the baseline models. Discussions about limitations and negative social impacts are in our Supplementary Materials.

**Acknowledgements:** This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2019-0-00075, Artificial Intelligence Graduate School Program(KAIST)). This work was also supported by the National Research Foundation of Korea under Grant NRF-2020R1A2B5B03001980.

<sup>4</sup><https://github.com/mit-han-lab/gan-compression>



## References

- [1] Matthew Amodio and Smita Krishnaswamy. Travelgan: Image-to-image translation by transformation vector learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 2
- [2] Asha Anoosheh, Eirikur Agustsson, Radu Timofte, and Luc Van Gool. Combogan: Unrestrained scalability for image domain translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 783–790, 2018. 6
- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 2
- [4] Sagie Benaim and Lior Wolf. One-sided unsupervised domain mapping. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 1, 2, 5, 6
- [5] Junya Chen, Zhe Gan, Xuan Li, Qing Guo, Liqun Chen, Shuyang Gao, Tagyoung Chung, Yi Xu, Belinda Zeng, Wenlian Lu, Fan Li, Lawrence Carin, and Chenyang Tao. Simpler, faster, stronger: Breaking the log-k curse on contrastive learners with flatnce, 2021. 2
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 13–18 Jul 2020. 2
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning, 2020. 2
- [8] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 6
- [9] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, Kun Zhang, and Dacheng Tao. Geometry-consistent generative adversarial networks for one-sided unsupervised domain mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 2, 5, 6
- [10] Lin Gao. Five class weather image dataset, 2019. 6
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. pages 6626–6637, 2017. 5
- [13] Qianjiang Hu, Xiao Wang, Wei Hu, and Guo-Jun Qi. Adco: Adversarial contrast for efficient learning of unsupervised representations from self-trained negative adversaries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1074–1083, June 2021. 2
- [14] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 1, 5, 6
- [15] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 1, 5, 6
- [16] Kibok Lee, Yian Zhu, Kihyuk Sohn, Chun-Liang Li, Jinwoo Shin, and Honglak Lee.  $\$i$ -mix: A domain-agnostic strategy for contrastive representation learning. In *International Conference on Learning Representations*, 2021. 3
- [17] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *International Conference on Learning Representations*, 2021. 2, 3
- [18] Muyang Li, Ji Lin, Yaoyao Ding, Zhijian Liu, Jun-Yan Zhu, and Song Han. Gan compression: Efficient architectures for interactive conditional gans. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 7
- [19] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 1
- [20] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 319–345, Cham, 2020. Springer International Publishing. 1, 2, 4, 5, 6, 8
- [21] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [22] Baoyun Peng, Xiao Jin, Jiaheng Liu, Dongsheng Li, Yichao Wu, Yu Liu, Shunfeng Zhou, and Zhaoning Zhang. Correlation congruence for knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [23] Joshua David Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. In *International Conference on Learning Representations*, 2021. 2, 3, 4
- [24] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019. 2
- [25] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *International Conference on Learning Representations*, 2020. 2

- [26] Weilun Wang, Wengang Zhou, Jianmin Bao, Dong Chen, and Houqiang Li. Instance-wise hard negative example generation for contrastive learning in unpaired image-to-image translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14020–14029, October 2021. [1](#), [2](#), [4](#), [5](#), [6](#)
- [27] Xiaolong Wang and Abhinav Gupta. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015. [2](#)
- [28] Chen Wei, Huiyu Wang, Wei Shen, and Alan Yuille. {CO}2: Consistent contrast for unsupervised visual representation learning. In *International Conference on Learning Representations*, 2021. [2](#), [3](#)
- [29] Guodong Xu, Ziwei Liu, Xiaoxiao Li, and Chen Change Loy. Knowledge distillation meets self-supervision. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 588–604, Cham, 2020. Springer International Publishing. [2](#)
- [30] Chuanguang Yang, Zhulin An, Linhang Cai, and Yongjun Xu. Mutual contrastive learning for visual representation learning, 2021. [3](#)
- [31] Chun-Hsiao Yeh, Cheng-Yao Hong, Yen-Chi Hsu, Tyng-Luh Liu, Yubei Chen, and Yann LeCun. Decoupled contrastive learning, 2021. [2](#), [3](#), [4](#)
- [32] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. [5](#)
- [33] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. pages 586–595, 2018. [6](#)
- [34] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. The spatially-correlative loss for various image translation tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16407–16417, June 2021. [1](#), [2](#), [4](#), [5](#), [6](#)
- [35] Jinguo Zhu, Shixiang Tang, Dapeng Chen, Shijie Yu, Yakun Liu, Mingzhe Rong, Aijun Yang, and Xiaohua Wang. Complementary relation contrastive distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9260–9269, June 2021. [2](#)
- [36] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [1](#), [5](#), [6](#)