

Contrastive Feature Loss for Image Prediction

Alex Andonian^{1,3} Taesung Park^{2,3} Bryan Russell³ Phillip Isola¹ Jun-Yan Zhu^{3,4} Richard Zhang³
¹MIT ²UC Berkeley ³Adobe Research ⁴CMU

Abstract

Training supervised image synthesis models requires a critic to compare two images: the ground truth to the result. Yet, this basic functionality remains an open problem. A popular line of approaches uses the L1 (mean absolute error) loss, either in the pixel or the feature space of pretrained deep networks. However, we observe that these losses tend to produce overly blurry and grey images, and other techniques such as GANs need to be employed to fight these artifacts. In this work, we introduce an information theory based approach to measuring similarity between two images. We argue that a good reconstruction should have high mutual information with the ground truth. This view enables learning a lightweight critic to “calibrate” a feature space in a contrastive manner, such that reconstructions of corresponding spatial patches are brought together, while other patches are repulsed. We show that our formulation immediately boosts the perceptual realism of output images when used as a drop-in replacement for the L1 loss, with or without an additional GAN loss.

1. Introduction

A fundamental requirement for image prediction tasks is an effective loss function, to judge whether the synthesized results are “close” to the ground truth. Straightforward measures – such as the Hamming distance or Euclidean L2 distance – do not work well with high-dimensional and structured signals, such as images.

A predominant approach addressing the high-dimensionality of images is to take advantage of the emergent perceptual similarity found in deep network activations, commonly known as the “perceptual loss” or feature matching loss [9, 40, 11, 19]. As seen in Figure 1 (bottom-left), this loss is obtained by computing the L1 distance between the internal activations of a pretrained network to judge the similarity of two images. Even though the network is pretrained on a seemingly unrelated task (ImageNet [7] classification), the corresponding feature space has been shown to be “unreasonably” effective in matching human perception [48] even compared to a rich

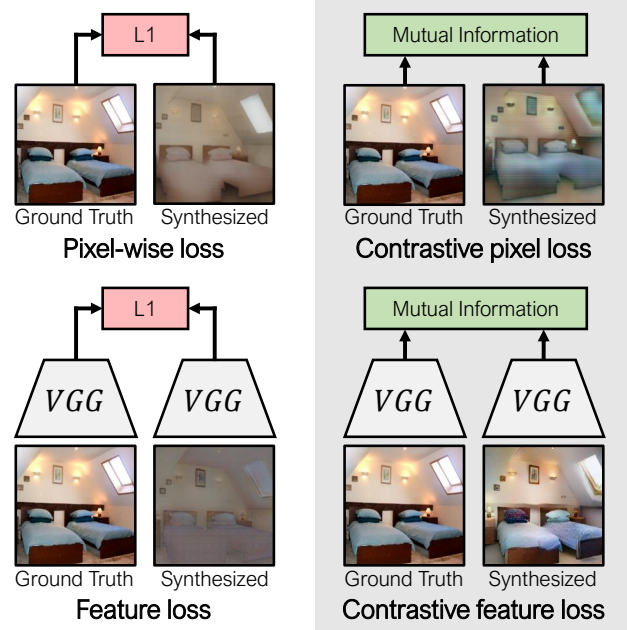


Figure 1. **Patchwise contrastive perceptual loss.** Existing losses for comparing synthesized to ground truth patches use a fixed distance metric such as L1, either in the pixel space [17] or some feature space [3, 43, 28]. Our method aims to replace the L1 or the L2 regression loss with a loss based on mutual information. Namely, our method learns to group the predicted patch and the ground truth together, distinct from other patches. Our method can be used in either the pixel space or the feature space to improve visual realism.

line of hand-engineered similarity metrics [44, 46, 21].

However, despite the success of deep features used for perceptual metrics and losses [43, 28, 1, 33] in supervised image synthesis tasks, the original question of how to measure similarity between images has not been fully resolved. In other words, the representations of the signals have been transformed, but an L1 or L2 loss is ultimately still used to judge similarity.

We observe that using the L1 distance metric for image synthesis tasks often results in undesirable visuals that are too smooth or desaturated, even when used in deep feature space (Figure 1 bottom-left). We hypothesize that this is because the L1 loss estimates the median of the multi-modal

distribution of the ground truth images. For example, the bed in Figure 1 could be in various styles, and L1 finds the median of all, which turns out to be flat and desaturated.

In this work, we view structural similarity from the perspective of mutual information. We posit that a good synthesis is one that is easily associated to the ground truth instance, even if they do not match pixel-by-pixel or feature-by-feature. While directly estimating mutual information is intractable, recent techniques in unsupervised representation learning [15, 4, 26] can be adapted for this purpose. Namely, we use contrastive learning, which learns an embedding to associate corresponding signals, in *contrast* to other samples in the dataset. In our case, corresponding signals are patches at the same spatial location in ground truth and synthesized images, and other samples are patches at different spatial locations in the dataset. Furthermore, we find that defining the contrastive loss bidirectionally – ground truth as the target in one direction, and the synthesized as the target in the other – improves performance.

We investigate whether this contrastive loss formulation serves as a suitable “drop-in” replacement for the standard L1 loss when applied directly on both patches of raw pixels and features extracted by common CNN backbones (Figure 1 right). There are several reasons why a contrastive formulation provides an attractive alternative to L1 loss. In particular, since contrastive learning aims to maximize mutual information, it may potentially capture more desired structure compared to L1. Moreover, the most prominent approaches to contrastive loss lend themselves to being made adaptable to particular data domains, allowing for “calibration” on top of a pretrained feature space via a critic function. That is, a small 1-2 layer network is trained on top to select for features, with the practical purpose of emphasizing those that are important for the task, and the theoretical purpose of serving as a better estimator for mutual information [26].

We note a few key differences between our method and GAN discriminators [12], which also learn a critic. Whereas discriminators are trained adversarially, ours is trained in a cooperative fashion, with both generator and critic aiming to maximize mutual information. Ultimately, the methods can be considered complementary. The goal of our method is assessing the *similarity between two images*, while GANs evaluate the *realism of a single image*. We show that our contrastive feature loss can be used in conjunction with a GAN loss, performing competitively against SPADE [28], which uses a discriminator and the standard perceptual loss.

We test our method across several image translation tasks, purposely choosing tasks with output domains of varying appearances and levels of photorealism: synthesizing photorealistic (Cityscapes [6] and ADE20K [49]) and synthetic scenes (GTA [31]) from labels, predicting depth

maps from indoor images [34]. Our experimental results show that across these image translation tasks, our contrastive loss formulation serves as a strong replacement for the L1 loss as a distance metric, especially when applied to the embedding space of trained deep networks.

2. Related Work

Our work is related to prior work on image similarity, image synthesis, and contrastive feature learning.

Image similarity. Similarity metrics aim to measure how “close” two images are. The simplest method is taking a point-wise difference, using Euclidean ℓ_2 , Manhattan ℓ_1 , or PSNR. Such methods do not have the capacity to consider joint statistics. As such, methods developed in the image processing community, such as SSIM [44], FSIM [46], and HDR-VDP-2 [21] build models on top of patches, taking into consideration the human visual system. These methods are effective in situations where structural ambiguity is not a factor, and factors such as photometric changes dominate. However, in many image translation problems, such as translating labels into photorealistic images, synthesizing structure is the predominant factor. Recently, the computer vision community has found that taking distances in the internal embedding space of a deep network, referred to as a “perceptual loss” or feature matching loss [19, 9, 11, 40], works remarkably well for image synthesis [3]. Improvements include adding transformations to allow for more spatial flexibility [32, 22], calibrating the embeddings using human perceptual judgments [48], or architectural changes [8, 13, 47]. Nonetheless, the backbone remains a network that was trained for a completely different task and data distribution, typically ImageNet classification [7]. In our work, we investigate whether it is possible to move beyond a frozen backbone, and *adapt* the feature matching loss to the data and task distribution at hand.

An aspect of image synthesis algorithms is also how real the image looks, in isolation. A natural method of enforcing realism is with a learned discriminator, or GAN loss [12]. Unlike similarity metrics, which compare two images, a discriminator looks at one image in isolation, in a “no-reference” fashion, and evaluates it. In image synthesis, similarity and realism are both factors that should be accounted for, and are sometimes at odds [2]. Thus, discriminators and similarity metrics are often combined in image translation systems in an additive manner [17, 50, 43, 28]. Recent works have also explored the coupling between the two ideas, such as re-using the discriminator as a perceptual loss [43, 41, 28], as a learned texture descriptor [29], combining their architectures together [38], or a training schedule that anneals from one to the other [37]. We show that our method is complementary to GANs, and can be used effectively in conjunction with a learned discriminator.

Contrastive feature learning. Contrastive feature learning

has primarily looked at learning strong feature representations without supervision for downstream visual recognition tasks [4, 15, 24, 26, 42, 45]. These approaches optimize a loss that aims to map a query feature to a positive target, which is obtained via data augmentation [4, 15, 24] or leveraging a cross-modal signal [23, 25, 39], and contrast against a set of negative targets. Then, a given learned feature representation can be applied to downstream recognition tasks with minimal supervision. Our approach is inspired by these efforts, but our goal is different – we seek to leverage contrastive feature learning for image prediction. Our work is related to CUT [27], as both use a contrastive loss. However, while CUT asks the output to correspond to the input, we ask the network to synthesize details close to the ground truth image. Furthermore, we find that a direct implementation of CUT in the paired data domain suffers from bad local minima due to difficult optimization and show that our proposed *bidirectional* formulation effectively mitigates this problem. Furthermore, since our feature encoder does not share weights with the generator (as it does in CUT), we are able to leverage pretrained feature encoders to produce further downstream performance gains. Lastly, while the CUT system would not function without a discriminator, we show our method can operate in both settings, with and without a discriminator.

3. Methods

We consider problems where a generator G learns a mapping from input $X \in \mathbb{R}^{H \times W \times C_X}$ to $Y \in \mathbb{R}^{H \times W \times C_Y}$. For example, for synthesizing photo from a semantic layout, the input is a label map with resolution $H \times W$ and C_X classes, and the output is an RGB image ($C_Y = 3$). The generator can be any function approximator, for example a deep convolutional neural network. Our goal is to measure similarity between a ground truth signal Y and its reconstruction $\hat{Y} = G(X)$. In this section, we motivate and introduce our bidirectional PatchNCE loss for comparing the similarity of two signals. Moreover, we describe and compare with existing traditional strategies for measuring signal similarity.

3.1. Mutual information maximization

Often, due to the inherent ambiguity of the task, for a given input signal X , there is not just a single ground truth signal but a set of *perceptually plausible* output signals \mathcal{Y} . For example, given the task of generating an image given an input segmentation mask, a “car” mask may correspond to any number of satisfactory distinct instances of the same category. As the desired output is often ill-defined from the input signal alone, we seek to define a loss function comparing a generated result \hat{Y} with a ground truth signal $Y \in \mathcal{Y}$ that allows for such ambiguity in the signal.

An example of a poor loss function is per-pixel regression, such as Euclidean ℓ_2 loss. Using such a loss will aver-

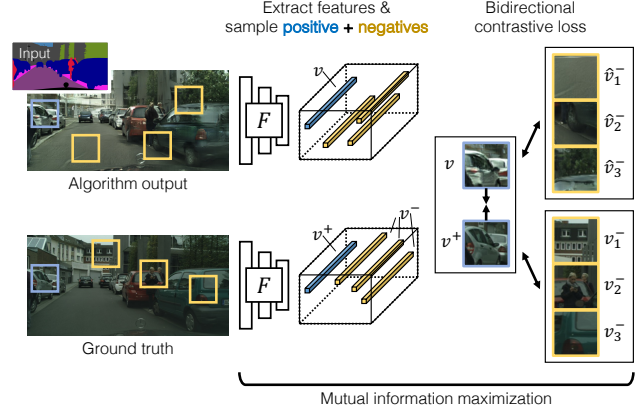


Figure 2. **Bidirectional patchwise contrastive perceptual loss.** Our approach compares spatial patches in a generated signal to patches in a corresponding ground truth. For the depicted label→image generation task, previous work uses L1 or L2 loss on a feature representation (F can be a frozen feature extractor or identity function to represent raw pixels). Based on mutual information, our loss encourages the encoding of corresponding patches in the generated and ground truth signals to be close (blue patches) and non-corresponding patches (yellow) to be far.

age over the outcomes \mathcal{Y} , producing a blurry result. This is undesirable, as there is heavy *loss of information*. A blurry result cannot be associated with the ground truth image, relative to other images. We hypothesize that a property of a good reconstruction is *high mutual information* \mathcal{I} with the ground truth,

$$\mathcal{I}(\hat{Y}; \mathcal{Y}) = \sum_{Y \in \mathcal{Y}} P(\hat{Y}, Y) \log \left(\frac{P(\hat{Y}, Y)}{P(\hat{Y})P(Y)} \right), \quad (1)$$

where P denotes probability distributions over random variables. Note that if the set of plausible signals \mathcal{Y} is a single instance, and reconstruction \hat{Y} perfectly matches this, then the mutual information is infinite.

In many cases, we are given a dataset $\mathcal{D} = \{(X, Y)\}$ of input and output ground truth pairs, which contains a *single* observation of the plausible results. We seek a generator G^* that maximizes the mutual information of reconstructions and ground truth images, within the dataset, using the following objective,

$$G^* = \arg \max_G \mathbb{E}_{(X, Y) \in \mathcal{D}} \mathcal{I}(G(X); Y) \quad (2)$$

Computing the mutual information directly is intractable, as it requires enumerating all possible outcomes and calculating the joint probability. We use recent advances in the unsupervised learning literature, described below [15].

3.2. Bidirectional patchwise contrastive loss

We use a noise contrastive estimation framework [14] as a means to maximizing mutual information between reconstruction and ground truth. The idea is that if a reconstruction looks “similar” to the ground truth, the two should

be able to be embedded closely, relative to other images in the dataset. Furthermore, we note that for an *image* to be reconstructed well, each *patch* within the image should look similar to the corresponding patch in the ground truth. Thus, we can operate on a patchwise level, which enables denser supervision for generator G .

We realize this intuition with a *patchwise, contrastive loss*. Let $\hat{y}_p, y_p \in \mathbb{R}^{H_{\text{patch}} \times W_{\text{patch}} \times C_Y}$ be patches from the two signals, in the same spatial location. Let $v = F(\hat{y}_p), v^+ = F(y_p) \in \mathbb{R}^D$ be the D -dimensional embedding vectors of spatially-corresponding generated and ground truth patches, and $v^- \in \mathbb{R}^{N \times D}$ be the embedding on the N other patches (“negatives”) from the dataset. In this work we investigate two classes of encoder F : (1) a simple linear projection of a patch of pixels, and (2) a pretrained frozen deep network. The loss is defined as an $(N+1)$ -way classification problem, with logits proportional to the similarity between embedded patches,

$$\ell_{\text{contrastive}}(v, v^+, v^-) = -\log \frac{\exp(s(v, v^+)/\tau)}{\exp(s(v, v^+)/\tau) + \sum_{n=1}^N \exp(s(v, v_n^-)/\tau)}, \quad (3)$$

where $s(v_1, v_2) = v_1^\top v_2$ returns the similarity between two encoded patch signals as a dot product, and τ is a parameter to control temperature. This objective realizes our intuition. If a patch is exactly reconstructed, then the embeddings will have perfect similarity and the numerator will be maximized. If the patch is approximately reconstructed, it obtains a high score if it can be easily associated to the ground truth, relative to the non-matching “negative” patches.

Multiscale/Multilayer implementation. In the case where F is a simple linear layer, patches of pixels are cropped at multiple scales and passed through the layer. In contrast, when using a deep network as our encoder, we do not literally crop out patches and encode. Rather, for computational efficiency, we pass the full images \hat{Y} and Y into a network, producing a stack of features. The feature activation at a given spatial location and layer corresponds to a feature representation of a patch in the input image. We illustrate this setup in Figure 2. The size of the patch depends on the receptive field, depending on the architecture, and the layer of the network. Thus, by taking activations at multiple layers of the network, we are able to compute the patchwise contrastive loss at different spatial scales.

More concretely, images \hat{Y}, Y are extracted into sets of L feature tensors $\{\hat{V}_l, V_l\}_{l=1}^L$, where L is a pre-specified number of layers. Each feature tensor is the output of the l^{th} layer of encoder F , and a small projection head (e.g. a linear layer or two-layer MLP) H . There is a slight abuse of notation; the composition of H and F corresponds to F in the previous subsection. Adding a small MLP was shown by Chen et al. [4] to improve performance, so we found

this practice compatible with our framework. The shape of tensor $V_l \in \mathbb{R}^{S_l \times D_l}$ is determined by the architecture of the network, where S_l is the number of spatial locations of the tensor. We index into the tensor with notation $v_l^s \in \mathbb{R}^{D_l}$, which is the D_l -dimensional feature vector at the s^{th} spatial location. We denote $\bar{v}_l^s \in \mathbb{R}^{(S_l-1) \times D_l}$ as the collection of feature vectors at all other spatial locations.

Our loss on a pair of images is written as follows.

$$\mathcal{L}_{\text{contrastive}}(\hat{Y}, Y) = \sum_{l=1}^L \sum_{s=1}^{S_l} \ell_{\text{contrastive}}(\hat{v}_l^s, v_l^s, \bar{v}_l^s) \quad (4)$$

Bidirectional PatchNCE Loss. Drawing inspiration upon classic work in image similarity [35], we investigate a symmetric “bidirectional” variant of the PatchNCE loss. We pursue this with the intuition that not only should generated patches contrast with non-corresponding in training patches, the corresponding training patch should similarly contrast with non-corresponding *generated* patches. We find that this leads to more stable training behavior, faster convergence, and avoids degenerate solutions. We find this to consistently hold true across different settings. Here, a second contrastive term is added with the roles of the generated and ground truth patches reversed, as shown below:

$$\mathcal{L}_{\text{contrastive}}(\hat{Y}, Y) = \frac{1}{2} \sum_{l=1}^L \sum_{s=1}^{S_l} \ell_{\text{contrastive}}(\hat{v}_l^s, v_l^s, \text{sg}(\bar{v}_l^s)) + \ell_{\text{contrastive}}(v_l^s, \hat{v}_l^s, \text{sg}(\hat{v}_l^s)) \quad (5)$$

where sg function indicates a stop gradient operation (i.e., gradients are prevented from flowing through the negatives). In the representation learning context, SiamSiam [5] uses this technique to help prevent degenerate solutions. All experiments take advantage of this bidirectional variant, unless noted otherwise.

Final contrastive objective. Given a dataset $\mathcal{D} = \{(X, Y)\}$ of input and output ground truth pairs, we obtain the generator G^* with the following objective,

$$G^* = \arg \min_G \min_H \mathbb{E}_{(X, Y) \in \mathcal{D}} \mathcal{L}_{\text{contrastive}}(G(X), Y). \quad (6)$$

As encoder H and generator G are neural networks, then the entire system is differentiable and can be optimized jointly using stochastic gradient descent.

3.3. Classic feature matching loss

Our loss has similarities with the classic feature matching loss, or “perceptual loss”. The loss also passes image \hat{Y}, Y through a pretrained network, which we designate as F , to produce a stack of features \hat{V}, V . The perceptual loss uses the ℓ_2 function to compare feature activations,

$$\mathcal{L}_{\text{classic}}(\hat{Y}, Y) = \sum_{l=1}^L \sum_{s=1}^{S_l} \ell_2(\hat{v}_l^s, v_l^s). \quad (7)$$

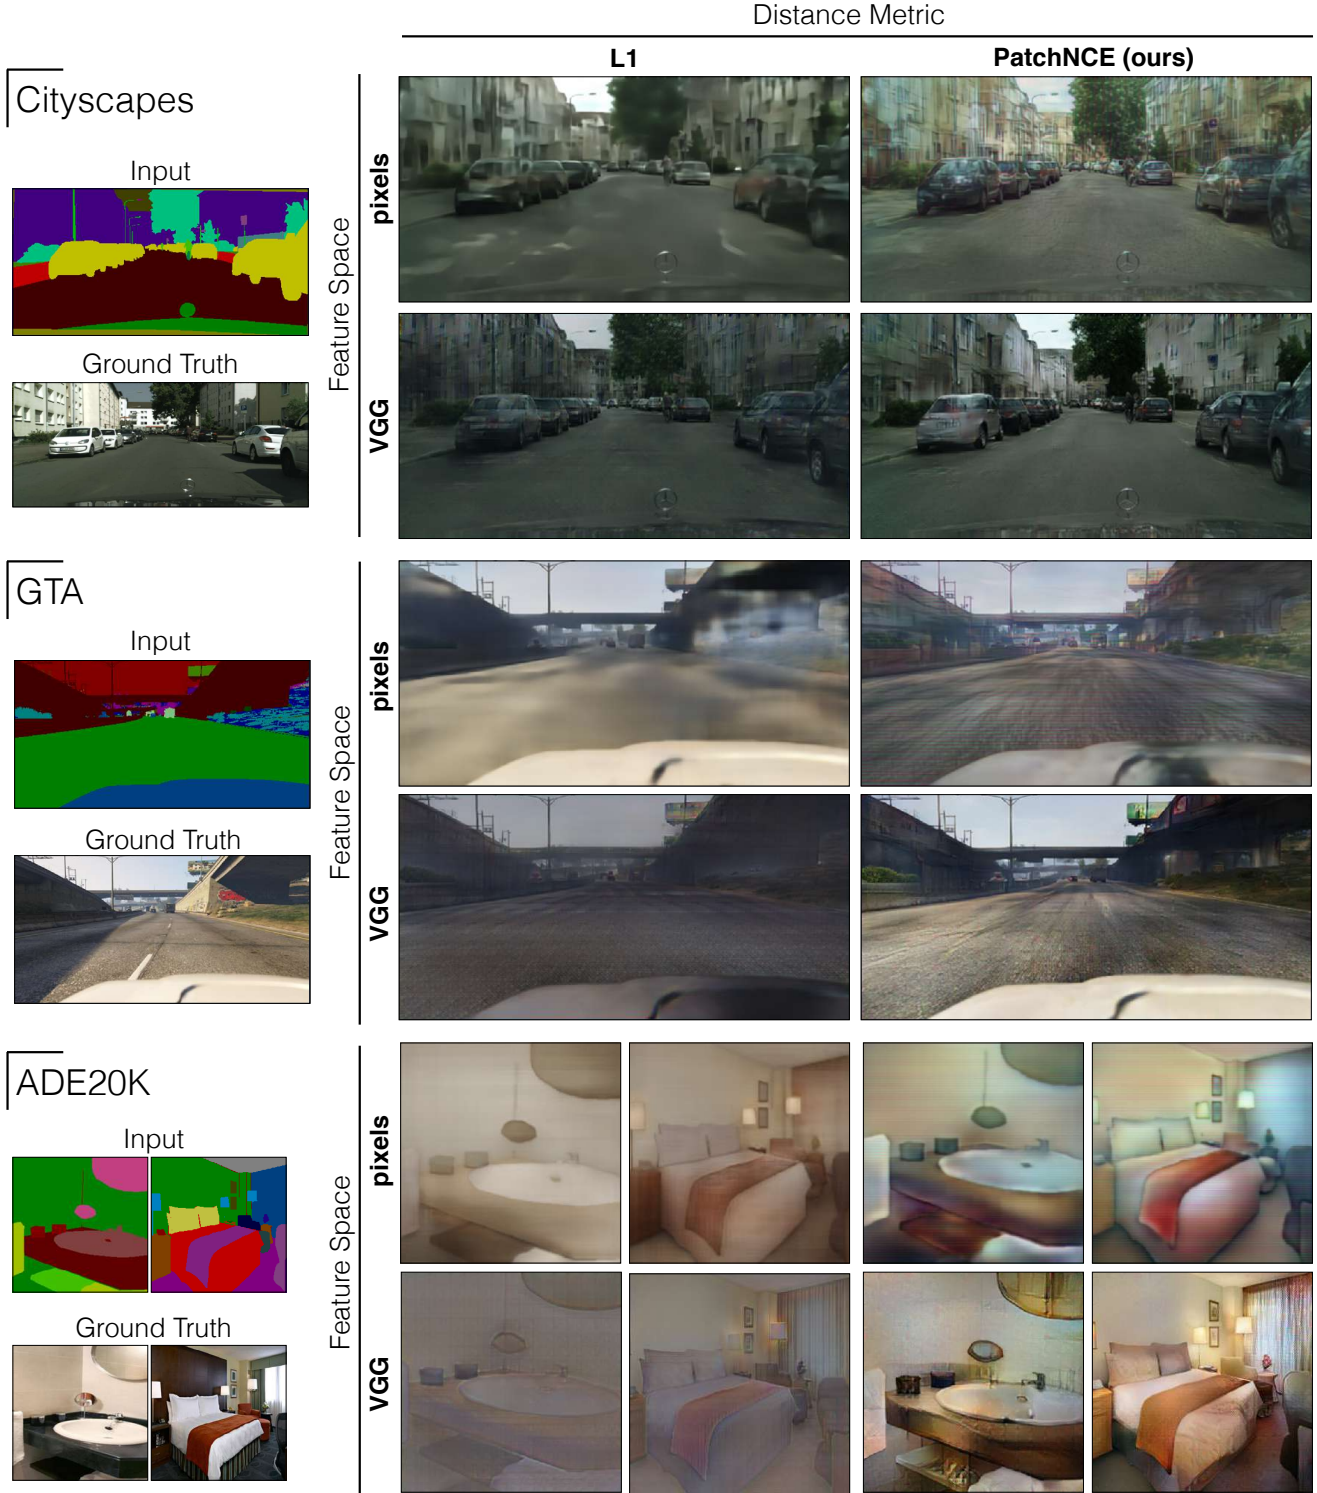


Figure 3. We compare our results against typical regression losses used in supervised images synthesis tasks, such as L1 in pixel space [18] and the feature space of the pretrained VGG network [3, 41, 28], on the Cityscapes[6], GTA[30], and ADE20K[49] dataset. Because the L1 loss encourages regressing toward the median of all possible targets, they end up producing overly blurred or desaturated results. In contrast, our PatchNCE loss, used in either the pixel or feature space, significantly boosts sharpness and colorfulness. Note that other techniques such as progressive growing or GAN are not used in this experiment.

The loss is used to optimize a generator, in the following,

$$G^* = \arg \min_G \mathbb{E}_{(X,Y) \in \mathcal{D}} \mathcal{L}_{\text{classic}}(G(X), Y). \quad (8)$$

Comparing the optimization above with our method in Equation 6, we note two differences. Firstly, rather than only bringing corresponding patches together, ours pushes them apart from non-corresponding patches at the same time. Secondly, we have a shallow, learnable network H on top of the primary feature extractor F , that can select for the more important features for a given task.

3.4. GAN loss

We note that our method provides a lightweight mechanism for calibrating to data, reminiscent of a GAN discriminator. A key difference is that our method judges the similarity between reconstruction and ground truth, whereas a GAN loss judges the “realism” of just the reconstruction conditional on the input. It does so by learning a discriminator D that judges whether an input output pair of $(X, G(X))$ looks like it belongs to a dataset of input-output pairs $(X, Y) \in \mathcal{D}$.

$$\mathcal{L}_{\text{cGAN}}(X, \hat{Y}, Y) = \log D(X, Y) + \log(1 - D(X, \hat{Y})) \quad (9)$$

The generator aims to “fool” the discriminator. The solution to the generator is found via min-max optimization. In our experiments, we show that our method can be used both by itself, or in conjunction with a discriminator,

$$G^* = \arg \min_G \min_{F,H} \max_D \mathbb{E}_{(X,Y) \in \mathcal{D}} [\mathcal{L}_{\text{contrastive}}(G(X), Y) + \mathcal{L}_{\text{cGAN}}(X, G(X), Y)]. \quad (10)$$

4. Experiments

We evaluate our method on several paired image-to-image translation datasets, for synthesizing both photorealistic and non-photorealistic images, and compare against several baselines, followed by model ablations.

Network architecture and training details. As our main interest is investigating the patch-based contrastive loss, we keep the same training procedure and architecture established by Pix2PixHD [43] and SPADE [28], except the loss term. In more detail, for our generator network, we use the ResNet-based SPADE generator, which pairs well with the multiscale PatchGAN discriminator network (when present). We investigate several different backbone architectures for the feature extraction network in our contrastive loss (see supp. for complete list), focusing specifically on the VGG19 [36] network standard “perceptual” loss. Features extracted by these networks are then reduced to low-dimensional embeddings via a simple linear layer projection head and normalized by their L2 norms.

Baselines. Our main goal is to compare the effect of the data-driven contrastive loss function as an improved drop in replacement for the ubiquitous L1/L2 loss. Therefore,

we compare our approach against simple $L1$ pixel distance loss, $L1$ distance of VGG feature responses [3], in combination with the GAN loss [28], while we use the same network architecture for a fair comparison. Therefore, we do not directly compare our method against prior work [3, 43] as they used different architectures. For SPADE, we use publicly available code.

4.1. Paired Image-to-Image Translation

Datasets. We evaluate our method on several datasets with varying amounts of photorealism.

Cityscapes [6] contains street photographs paired with their corresponding semantic label maps for 2,975 training and 500 validation images. Models learn to translate semantic label maps containing up to a total of 19 unique semantic categories to 512×256 resolution output images.

GTA dataset [30] contains synthetic scenes from the widely-popular *Grand Theft Auto V* video game. This dataset is semantically similar to Cityscapes but is larger with 26k training images and 5k testing images. The resolution is same as the Cityscapes dataset.

NYU Depth V2 [34] contains video frames of 464 indoor scenes from 3 cities recorded by the Microsoft Kinect RGB-D cameras. Our task is to synthesize depth maps from RGB images at 256×256 resolution. We create an 80%-20% split of the 1,449 pre-processed pairs to synthesize colored depth map from RGB images.

ADE20K [49] consists of challenging indoor and outdoor scenes labeled for 150 different semantic classes. The dataset is split into 20,210 training images and 2,000 validation images.

Evaluation protocol. We adopt evaluation protocols that are tailored to traditional image translation tasks as well as more domain-specific metrics such as those frequently employed in the pixel labeling and depth estimation literature. To briefly summarize, we first use the standard Fréchet Inception Distance (**FID**) metric [16], which is aimed at assessing the visual quality of generated samples compared to real images by computing the statistics between the distributions of real and generated images estimated by a deep neural network. In the standard setting, this network is an InceptionV3 pretrained on the ImageNet dataset. However, FID does not evaluate the correspondence between inputs and outputs. Therefore, for semantic image synthesis tasks of the Cityscapes, GTA and ADE20K datasets, we adopt prior protocol and evaluate segmentation accuracy [3, 43, 28]. We report the mean average precision (**mAP**) and pixel accuracy (**Acc**). In order to evaluate the performance of our approach on the NYU V2 depth estimation task, we report the standard absolute relative difference (**rel**) and RMSE-linear (**rms**) criteria [10, 20].

Qualitative evaluation As shown in Figure 3, the L1-regression loss in the pixel or the VGG feature space re-

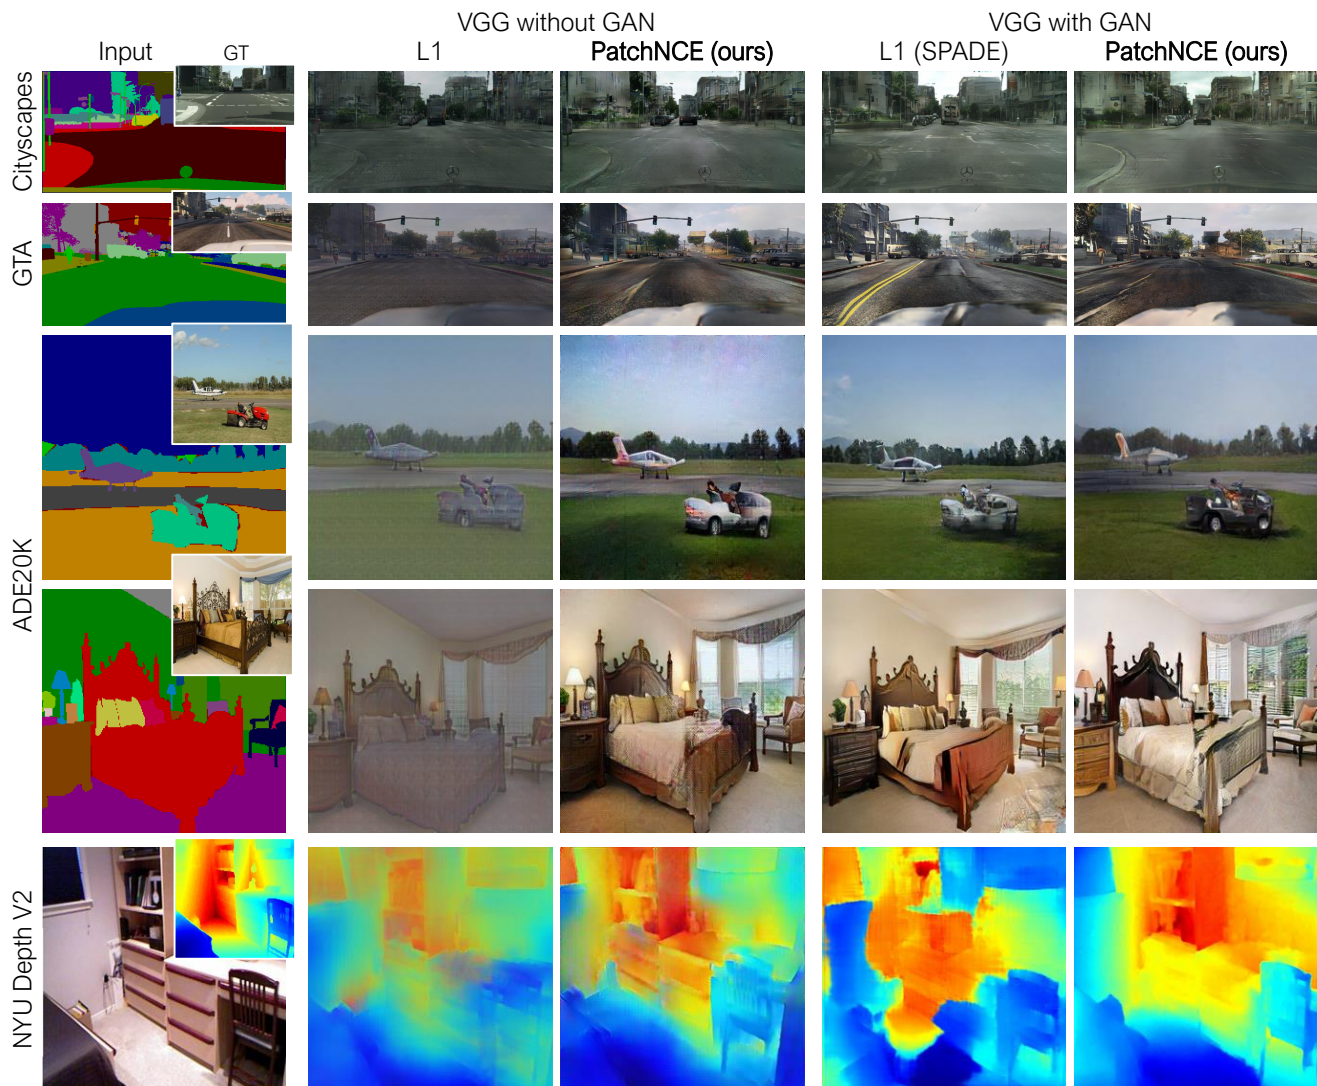


Figure 4. We show the effect of using the PatchNCE loss in comparison to the L1 feature loss on the Cityscapes, GTA, ADE20K, and NYU Depth V2 dataset, without and with the GAN discriminator. Even without GAN, our PatchNCE loss brings out sharpness and saturation, almost matching the quality of SPADE [28], which uses GAN as well as the L1 feature matching loss. Adding GAN to PatchNCE further enhances realism by removing small grid artifacts in the outdoor scene (ADE20K, third row) for example.

sults in desaturated, blurry images, since the information loss in the output image is not penalized. In contrast, our PatchNCE method does not suffer from this same phenomena and produces output images with increased sharpness and color saturation. We observe that the increased sharpness and saturation makes the outputs similar to the results of GAN-based methods like SPADE [28] (Figure 4). Moreover it is seen that our PatchNCE can also further benefit from the addition of GAN.

Quantitative evaluation In Table 1, we show quantitative measures comparing our method variants to alternative approaches for two street scene datasets, one natural (Cityscapes) and one synthetic (GTA), a diverse indoor-outdoor dataset (ADE20K), and RGB - colored depth map dataset (NYU Depth V2). Here, we compare our method

against baselines using $L1$ pixel regression, $L1$ loss on VGG features as in [3], and also in combination with the GAN loss used in [43, 28]. We present three variants of our method: the generator trained with only the patch-based contrastive loss applied directly to raw pixels (Pixel PatchNCE), contrastive loss applied to fixed VGG features (VGG PatchNCE), and our formulation that incorporates a GAN discriminator loss (VGG + GAN PatchNCE). In nearly all cases across these datasets, simply replacing the $L1$ metric with our contrastive loss formulation brings notable improvements to segmentation metrics (mAP, Acc) and depth estimation metrics (rel, rms). Moreover, our method consistently lowers FID scores for the variants that forgo the GAN discriminator loss, suggesting that the PatchNCE loss helps to bring significant improvements in photo-realism.

Feature Space	Loss	Cityscapes [6]			GTA [30]			ADE20K [49]			NYU Depth V2 [34]		
		FID ↓	mAP ↑	Acc ↑	FID ↓	mAP ↑	Acc ↑	FID ↓	mAP ↑	Acc ↑	FID ↓	rel ↓	rms ↓
Pixel	L1	135.1	45.3	76.3	132.1	20.2	33.7	126.5	7.4	43.4	264.6	0.431	1.83
	PatchNCE	107.3	48.2	78.2	107.2	23.1	40.1	125.9	8.7	46.2	170.6	0.434	1.77
VGG	L1 [3]	76.2	60.5	81.9	62.22	31.6	47.7	74.6	20.9	63.6	103.2	0.410	1.80
	PatchNCE	68.3	64.6	82.4	47.3	33.6	50.7	48.9	29.0	69.5	97.2	0.413	1.87
VGG	L1 [28]	58.7	62.3	81.9	36.3	30.8	47.3	33.1	38.5	79.9	86.5	0.454	2.12
+ GAN	PatchNCE	66.7	66.1	82.4	39.4	34.8	51.9	34.9	29.1	70.8	64.0	0.444	2.11

Table 1. Using the PatchNCE loss instead of L1 improves quantitative metrics across different datasets and feature spaces. For each dataset, we compare the results of training with our PatchNCE loss as opposed to the standard L1 regression loss. The experiments are also performed in different feature spaces, including just the pixel space (Pixel), the feature space of a pretrained VGG19 network (VGG), and VGG in combination with a GAN discriminator (VGG + GAN). In absence of GAN (the first four result rows), the advantage of our method over L1 is clear across datasets and various metrics. When GAN is added, we observe that the gap is reduced, likely because the GAN discriminator “fixes” the typical artifacts of the L1 regression. Still, our method attains superior correspondence metrics (mAP, Acc, rel, rms) overall. Note that L1 loss in VGG + GAN corresponds to the original SPADE [28]. Also, the L1 regression loss in VGG space is similar to CRN [3], except that the SPADE architecture was used for fair comparison.

For completeness, we note that published CUT [27] performance on the Cityscapes semantic segmentation metrics (mAP: 24.7, Acc: 68.6) is significantly lower than all other methods in Table 1, underscoring the advantages of using paired data when it is available.

One advantage afforded by our implementation of a contrastive loss over L1 loss is the ability to still take advantage of strong feature extractors pretrained on large datasets while adapting it to new data domains in an efficient way. Specifically, our methods trains a lightweight projection head as it learns to become a better critic (i.e., a better estimator of mutual information). Therefore, we hypothesize that this difference should prove to be advantageous when the target domain diverges from natural images. Indeed, the qualitative and quantitative comparison shown on the GTA and NYU Depth V2 dataset support this hypothesis with our method showing performance improvements on synthesizing CG images or colorized depth maps.

Importance of Bidirectional PatchNCE and pretraining

In Table 2 we show the performance of our method under several ablations with and without a GAN discriminator loss. In particular, we compare standard vs bidirectional PatchNCE loss variants and investigate the effect of unfreezing the underlying feature network, or even training it from scratch. Compared to the baseline method where L1 loss is applied to VGG features, all bidirectional variants achieve higher mAP scores than the baseline (60.5), even if the underlying feature extractor is trained from scratch on the target domain. Keeping all other factors fixed, the bidirectional NCE loss performs consistently better than the standard NCE, but the bidirectional NCE loss becomes increasingly important in stabilizing training when the underlying feature extractor is also being trained. The last row in each section of in Table 2 most closely resembles a direct application of CUT to paired data with/without a discriminator, and its trailing performance further highlights the advantages of our proposed bidirectional formulation.

Method				Performance	
NCE variant	GAN	Pretrained F	Frozen F	FID ↓	mAP ↑
Bidirectional*	✓	✓	✓	66.7	66.1
Standard	✓	✓	✓	69.3	64.9
Bidirectional	✓	✓	✗	63.0	64.6
Standard	✓	✓	✗	68.8	58.8
Bidirectional	✓	✗	✗	76.2	62.0
Standard	✓	✗	✗	87.1	57.7
Bidirectional*	✗	✓	✓	68.3	64.5
Bidirectional	✗	✓	✓	72.1	62.6
Standard	✗	✓	✓	73.8	61.6
Bidirectional	✗	✓	✗	85.9	62.9
Bidirectional	✗	✗	✗	86.0	62.7
Standard	✗	✓	✗	135.7	42.5
Standard [†]	✗	✗	✗	143.0	39.3

Table 2. FID and mAP scores on Cityscapes under different configurations of the VGG-PatchNCE method, including whether the backbone VGG network is frozen and/or pretrained, or if GAN loss is present. Bidirectional* forgoes NCE loss applied to the first layer appearing in the original VGG loss implementation (layer 2) and serves as our default configuration. Standard[†] is analogous to CUT [27] without a discriminator adapted for use with paired data.

5. Conclusion

We have introduced an approach for learning a lightweight image similarity critic via our PatchNCE formulation. We have demonstrated that the standard L1 loss frequently applied in pixel or feature space can be replaced by learning a new distance metric based on maximizing information between patches. Our loss formulation achieves competitive or superior results on synthesizing both natural and non-natural images, and works well in raw pixel space or the feature space of standard deep networks, with or without a GAN discriminator. Our approach opens up new avenues for exploring adaptive critics in other signal modalities, such as video and audio. **Acknowledgments.** We thank Yonglong Tian, Dilip Krishnan and Aude Oliva for helpful discussion. This work was initiated while AA and TP were interns and JYZ was an employee at Adobe, and was supported in part by an Adobe gift.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4432–4441, 2019. 1
- [2] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [3] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2, 5, 6, 7, 8
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning (ICML)*, 2020. 2, 3, 4
- [5] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. *arXiv preprint arXiv:2011.10566*, 2020. 4
- [6] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 5, 6, 8
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Ieee, 2009. 1, 2
- [8] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. *arXiv preprint arXiv:2004.07728*, 2020. 2
- [9] Alexey Dosovitskiy and Thomas Brox. Inverting visual representations with convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4829–4837, 2016. 1, 2
- [10] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in Neural Information Processing Systems*, 2014. 6
- [11] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 2
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014. 2
- [13] Jinjin Gu, Haoming Cai, Haoyu Chen, Xiaoxing Ye, Jimmy Ren, and Chao Dong. Pipal: a large-scale image quality assessment dataset for perceptual image restoration. *arXiv preprint arXiv:2007.12142*, 2020. 2
- [14] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010. 3
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 3
- [16] 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. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 6
- [17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2
- [18] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [19] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, 2016. 1, 2
- [20] Kevin Karsch, Ce Liu, and Sing Bing Kang. Depth extraction from video using non-parametric sampling. In *European Conference on Computer Vision (ECCV)*, 2012. 6
- [21] Rafał Mantiuk, Kil Joong Kim, Allan G Rempel, and Wolfgang Heidrich. Hdr-vdp-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on graphics (TOG)*, 30(4):1–14, 2011. 1, 2
- [22] Roey Mechrez, Itamar Talmi, and Lihi Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [23] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-End Learning of Visual Representations from Uncurated Instructional Videos. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [24] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [25] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. *CoRR*, abs/2004.12943, 2020. 3
- [26] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *CoRR*, abs/1807.03748, 2018. 2, 3
- [27] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for conditional image synthesis. In *European Conference on Computer Vision (ECCV)*, 2020. 3, 8
- [28] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 5, 6, 7, 8

- [29] Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei A Efros, and Richard Zhang. Swapping autoencoder for deep image manipulation. *arXiv preprint arXiv:2007.00653*, 2020. 2
- [30] Stephan R Richter, Zeeshan Hayder, and Vladlen Koltun. Playing for benchmarks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 5, 6, 8
- [31] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [32] Tamar Rott Shaham and Tomer Michaeli. Deformation aware image compression. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [33] Assaf Shocher, Yossi Gandelsman, Inbar Mosseri, Michal Yarom, Michal Irani, William T Freeman, and Tali Dekel. Semantic pyramid for image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7457–7466, 2020. 1
- [34] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision (ECCV)*, 2012. 2, 6, 8
- [35] Denis Simakov, Yaron Caspi, Eli Shechtman, and Michal Irani. Summarizing visual data using bidirectional similarity. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 4
- [36] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015. 6
- [37] Jinxuan Sun, Yang Chen, Junyu Dong, and Guoqiang Zhong. Progressively unfreezing perceptual gan. *arXiv preprint arXiv:2006.10250*, 2020. 2
- [38] Diana Sungatullina, Egor Zakharov, Dmitry Ulyanov, and Victor Lempitsky. Image manipulation with perceptual discriminators. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 579–595, 2018. 2
- [39] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [40] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. In *International Conference on Machine Learning (ICML)*, 2016. 1, 2
- [41] Chaoyue Wang, Chang Xu, Chaohui Wang, and Dacheng Tao. Perceptual adversarial networks for image-to-image transformation. *IEEE Transactions on Image Processing*, 27(8):4066–4079, 2018. 2, 5
- [42] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning (ICML)*, 2020. 3
- [43] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 6, 7
- [44] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 1, 2
- [45] Zhirong Wu, Yuanjun Xiong, X Yu Stella, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [46] Lin Zhang, Lei Zhang, Xuanqin Mou, and David Zhang. Fsim: A feature similarity index for image quality assessment. *IEEE transactions on Image Processing*, 20(8):2378–2386, 2011. 1, 2
- [47] Richard Zhang. Making convolutional networks shift-invariant again. *International Conference on Machine Learning (ICML)*, 2019. 2
- [48] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [49] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba. Scene parsing through ade20k dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5122–5130, 2017. 2, 5, 6, 8
- [50] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2017. 2