A Novel BiLevel Paradigm for Image-to-Image Translation

Liqian Ma[1]   Qianru Sun[2,3*]   Bernt Schiele[3]   Luc Van Gool[1,4]

[1]KU-Leuven/PSI, Toyota Motor Europe (TRACE)   [2] National University of Singapore
[3]Max Planck Institute for Informatics, Saarland Informatics Campus   [4]ETH Zurich
{liqian.ma, luc.vangool}@esat.kuleuven.be
{qsun, schiele}@mpi-inf.mpg.de

Abstract

Image-to-image (I2I) translation is a pixel-level mapping that requires a large number of paired training data and often suffers from the problems of high diversity and strong category bias in image scenes. In order to tackle these problems, we propose a novel BiLevel (BiL) learning paradigm that alternates the learning of two models, respectively at an instance-specific (IS) and a general-purpose (GP) level. In each scene, the IS model learns to maintain the specific scene attributes. It is initialized by the GP model that learns from all the scenes to obtain the generalizable translation knowledge. This GP initialization gives the IS model an efficient starting point, thus enabling its fast adaptation to the new scene with scarce training data. We conduct extensive I2I translation experiments on human face and street view datasets. Quantitative results validate that our approach can significantly boost the performance of classical I2I translation models, such as PG^2 [22] and Pix2Pix [14]. Our visualization results show both higher image quality and more appropriate instance-specific details, e.g., the translated image of a person looks more like that person in terms of identity.

1. Introduction

Humans have the impressive ability to imagine new scenes from a few descriptions or reference images. For example, given a single picture of a butterfly and a single overview picture of a garden, we can easily imagine video sequences of this butterfly flying around in the garden. This is, however, a very challenging image-to-image (I2I) translation task for machine learning models, as only a single training image is available. Some works call this one-shot image translation [24, 3].

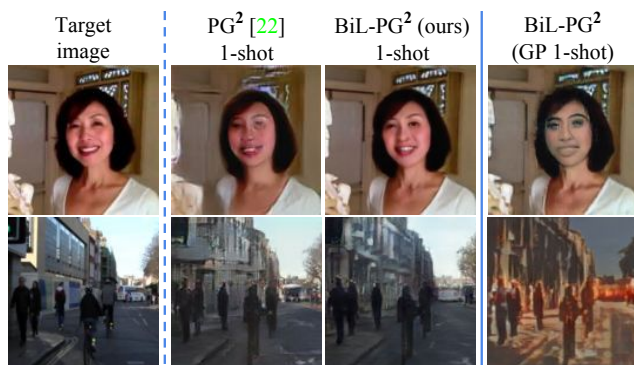By contrast, traditional I2I translation models [14, 43,

Figure 1. Our approach boosts the performance for the classical translation model PG^2 [22] for both face data and street view data. "GP" presents the generalizable translation knowledge we learned from diverse scenes.

52, 22] are usually trained on large-scale datasets of real-world images. These datasets, however, usually contain highly diverse image categories with only few samples in each category. It is difficult for the model to grasp the details of all categories. In other words, it is hard to learn a general I2I translation model. For example in Fig. 1, the pose guided person image generation model (PG^2) [22] suffers from artifacts, blurriness and the loss of individual characteristics in the translated images.

In order to tackle these problems, we propose a novel BiLevel (BiL) image translation paradigm that alternates the learning between an instance-specific (IS) model and a general-purpose (GP) model. The GP model aims to learn the generalizable translation knowledge across all image scenes, while each scene has an IS model that aims to maintain the specific attributes of the scene during translation. In this manner, the parameters GP model can be quickly adapted to be an IS model each time when coming a new scene. We achieve this by alternating the training processes of the IS and GP models. Specifically, "alternating" means that when the training of the IS model finishes, the loss of a validation set (in the same scene) is used to optimize the GP

model. In turn, the optimized GP parameters are used to initialize the IS model for the next scene. The validation loss is actually a generalization loss to optimize the GP model towards having a better generalization ability.

The GP model offers a warm start for the IS model to model a new scene and thus reduces the artifacts or blurriness caused by lacking data in the scene. The IS model is fine-tuned from the GP parameters to the particular scene. Therefore, it is more likely to generate instance-specific results. As we can see from Fig. 1, our approach successfully transfers the girl's identity features, such as skin color and hair style, to the output image and resolves the artifacts around the eyes, compared to PG^2 [22]. Our result on the more challenging street view maintains a sharper and clearer appearance, e.g. for the glass windows.

At testing time, the IS model takes the GP model as point of departure and adapts to the new scene based on few training images. An intuitive way to enhance IS training is to use these training images to query the available database for auxiliary data in similar scenes. It is not obvious how to retrieve the most helpful data and how to utilize them, however, given that there are only a few images for each scene. We propose a scene similarity metric that uses structure information and is computationally efficient. We also propose to use the similar scenes to fine-tune the GP model - thus later providing a better point of departure - rather than directly merge their data for IS training. We call this method Auxiliary fine-tuning (called Aux in this paper).

Our contribution is thus three-fold. (1) A novel and efficient BiLevel (BiL) approach that alternates the learning of GP and IS models. It is a training paradigm orthogonal to the specific model architecture and can thus boost the performance of classic models. (2) A simple and efficient Auxiliary fine-tuning (Aux) method that leverages data of similar scenes to enhance the few-shot learning in test scenes. (3) Extensive experiments on two challenging I2I translation datasets – FaceForensics [34] and BDD100 [47]. Our results are based on two classic models, i.e. PG^2 [22] and Pix2pix [14], and show that our approach can significantly boost performance both qualitatively and quantitatively.

2. Related work

Image-to-image translation. Image-to-image (I2I) translation aims to learn a mapping function to translate an image from a source domain to a target domain, *e.g.* semantic maps to real images [14, 43], real images to cartoon images [39], gray-scale to color images [49] and multi-domain translation [51]. Isola *et al.* [14] proposed the Pix2pix method which solves the I2I problem with a conditional GAN [25]. This currently is the most popular model. Pix2pixHD [43] is the high resolution version. To alleviate the need for paired training data, Zhu *et al.* proposed the self-supervising cycle-consistency loss [52]. Both

Pix2pix and CycleGAN learn deterministic mapping functions, *i.e.* given one input, there is only one possible output. To achieve multi-modal outputs, an implicit latent code is combined with the input [53, 1]. Most recently, disentangled structure and appearance representations are explored to gain explicit control over the translation process. Some works use one image as structure reference (also called content) and another image as appearance reference (also called style or attribute) to generate a novel image [18, 13, 21, 15, 44]. Some other works propose to extract body landmarks from images as the structural reference [22, 23, 20]. Generally, the training of I2I models requires a large dataset that encourages the model to learn the structure of different objects. Yet, this makes it hard to train a unified model as there are usually a handful of images for each object. The intuitive way to address the problem is to fine-tune a large-scale trained model with few-shot data of a specific scene. This is not effective as the model tends to easily overfit to few samples.

Few-shot I2I translation. Deep neural networks are data hungry, so moving to few-shot adaptation is an important improvement. Specific few-shot methods have been developed for different tasks such as recognition [42, 10, 29, 8, 32, 37, 19, 11], segmentation [31] and generative modeling [33, 7, 12], but for I2I translation there is little work. Benaim *et al.* [4] recently proposed a two-steps learning pipeline for one-shot unsupervised image translation. Their assumption is that the source domain is one-shot but the target domain has abundant samples. In contrast, our one-shot setting only contains a single pair of images, which is the most challenging case to study. In terms of the optimization method, our alternating learning paradigm follows the same gradient descent method as [8, 32, 19, 11, 37]. These works focus on image classification and their base-learning task uses a small dataset randomly sampled from a large one. Our instance-specific learning focuses on specific objects, *e.g.* a person identity and a street scene.

Transfer learning Transfer learning transfers knowledge between related source and target domains [28] and it is quite popular for addressing the small data problem [27, 45, 48]. The most intuitive and successful transfer fine-tunes a pre-trained model using the data of a new task [26]. Wang *et al.* [45] evaluated how GANs can be transferred. They showed that using pre-trained networks boosts training and improves the quality of generated images. We also initialize our bilevel training from a pre-trained network.

Active learning. In our test phase, we select auxiliary data to fine-tuning the GP model, akin to active learning [35, 40, 9]. The goal is to optimize the model within a limited time and annotation budget by selecting the most informative instances to annotate. Vijayanarasimhan *et al.* [40] showed how to select a set of examples meeting a given budget of supervision. Gavves *et al.* [9] proposed
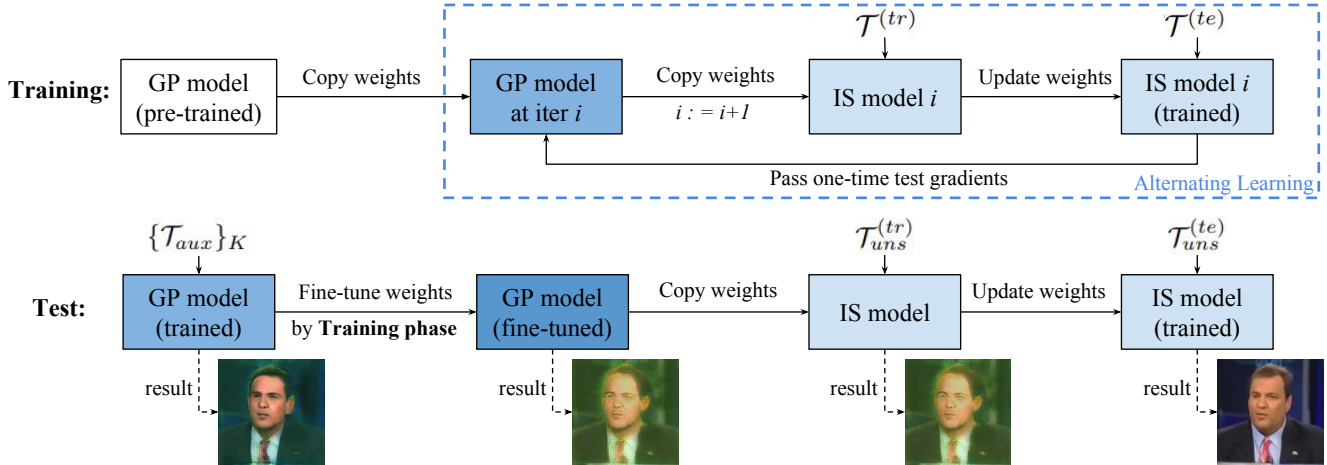
Figure 2. The proposed BiL image translation paradigm pipeline. It contains two phases: training phase and test phase. During the training phase, the GP model is first initialized through pre-training. Afterwards, in each iteration of alternating learning, the IS and GP models are updated in turn. Specifically, the IS model initialized by the GP model is updated with episode training data $\mathcal{T}^{(tr)}$ and then used to calculate the test loss for GP model update. During GP-Testing, for each unseen task, the top K most similar tasks $\{\mathcal{T}_{aux}\}_K$ are selected from the training set for fine-tuning the GP model according to which the new IS model is initialized and further updating.

an active learning method to reuse existing knowledge (*i.e.* available datasets). This is close to our idea of reusing training data during testing.

3. BiLevel I2I translation paradigm

In this section, we first explain the problem setup for instance-specific image translation with few training images. We then introduce our BiLevel (BiL) learning paradigm and the notations used for the two levels: general-purpose (GP) and instance-specific (IS).

3.1. Problem setup

The BiL learning paradigm contains two levels: the IS level for translating images in an individual scene; and the GP level for learning across scenes aiming to capture generalizable image translation knowledge. We thus denote the space of all these scenes as $p(\mathcal{T})$ and each point $\mathcal{T}$ in $p(\mathcal{T})$ represents an I2I image translation task in an individual scene. For any $\mathcal{T}$, the goal is to translate the input image to the target image that shares information with the input, e.g. the structure [14, 43] or the object identity [22, 38].

This task specific definition follows the unified episodic formulation first proposed for few-shot image classification [41]. Yet, there are several differences to the traditional paradigm: (1) the training of the GP model consists of a number of IS-level training-test procedures; (2) the training of an IS model is based on one episode with few-shot training image pairs and test image pairs; (3) the objective of the GP model is to initialize an IS model that should adapt to the new scene quickly; and (4) the final evaluation is the test result of the IS model adapted to unseen scenes.

3.2. BiLevel training and test

The proposed BiLevel paradigm contains two phases: training and test, as shown in Fig. 2. Note that the terms "training" and "test" are used only for the GP model. Because for each task, there are training and test procedures at the IS level. We call them IS training and IS test for clarity. There are multiple episodes $\{\mathcal{T}\}$ for training and one or multiple unseen scene episodes $\{\mathcal{T}_{uns}\}$ for test. Each training episode $\mathcal{T}$ contains an IS training split $\mathcal{T}^{(tr)}$ as well as an IS test split $\mathcal{T}^{(te)}$. The IS training loss on $\mathcal{T}^{(tr)}$ is used to optimize the IS model, while the IS test loss on $\mathcal{T}^{(te)}$ is used to optimize the GP model. In the test phase, an IS model is initialized by the GP model to quickly adapt to the unseen scene through fine-tuning on $\mathcal{T}_{uns}^{(tr)}$, then the IS test result on $\mathcal{T}_{uns}^{(te)}$ is reported as the final evaluation. If there are multiple unseen scene episodes, the average is reported. Specific operations of BiL training and test are given as follows.

Training phase. Deep models trained from scratch usually converge slowly [28]. Besides, it is notoriously difficult to train image generation models, e.g. Generative Adversarial Networks (GAN) based models, from scratch [2]. Therefore, we first pre-train the GP model using the images from all training scenes, with the traditional training paradigm [22, 14]. The pre-trained weights are then simply taken as the initialization of the GP model which is used in the following training phase (see Fig. 2). Afterwards, the GP model learns through a number of episodes in an alternating manner. In each episode, it first initializes an IS model which is then adapted to a specific scene. Finally, the validation loss computed from the learned IS model is used to update the GP model by meta gradient descent which un-

rolls the entire adaptation procedure of the IS model [8].

Given a training episode $\mathcal{T}$, the IS model parameters θ_{IS} are initialized by the parameters θ_{GP} of the current GP model which is learned by former tasks. Then, two stages, *i.e.* episode training and episode test, will be executed. In episode training, for each datapoint x in $\mathcal{T}^{(tr)}$ the translation loss $L_{x \in \mathcal{T}^{(tr)}}(\theta_{IS})$ is used to optimize the IS model and the model parameters θ_{IS} become θ'_{IS} (see Eq. 1). After several epochs of learning on $\mathcal{T}^{(tr)}$, the updated θ_{IS} will be frozen for episode test. During episode test, the losses of all datapoints $L_{\mathcal{T}^{(te)}}(\theta_{IS})$ are computed on $\mathcal{T}^{(te)}$, and the averaged loss is used to optimize θ_{GP} to be θ'_{GP} (see Eq. 2). This is done once in each episode.

$$\theta'_{IS} = \theta_{IS} - \alpha \nabla_{\theta_{IS}} L_{\mathcal{T}^{(tr)}}(\theta_{IS}), \qquad (1)$$

$$\theta'_{GP} = \theta_{GP} - \beta \sum_{\mathcal{T} \sim p(\mathcal{T})} \nabla_{\theta'_{IS}} L_{\mathcal{T}^{(te)}}(\theta'_{IS}), \qquad (2)$$

where α and β are the step sizes for optimizing the IS and GP models, resp.

Test phase. In this phase, a new IS model for the unseen translation task $\mathcal{T}_{uns}$ will be trained with θ_{IS} initialized by the learned θ_{GP}. Then, using the learned θ_{IS}, the image translation results on $\mathcal{T}_{uns}^{(te)}$ provide the final evaluation. Here, we propose an intuitive fine-tuning method to enhance θ_{GP} before IS training. We leverage $\mathcal{T}_{uns}^{(tr)}$ as reference to quickly query auxiliary tasks $\{\mathcal{T}_{aux}\}_K$ (with similar scenes) from training data. Then, we use the data to fine-tune θ_{GP}. We propose an efficient scene similarity metric using the structure information in images. For example, the input of our model consists of image segmentation maps. We calculate the similarity Sim of two tasks by summing up the intersection-over-union (IoU) scores over all segmentation classes, as follows,

$$Sim(x_{uns}, x_{aux}) = \sum_{c=1}^{C} \frac{\{x_{uns} = c\} \cap \{x_{aux} = c\}}{\{x_{uns} = c\} \cup \{x_{aux} = c\}}, \quad (3)$$

where C is the number of object classes. x_{uns} and x_{aux} denote the input segmentation maps of the unseen task and a candidate task, resp.

After retrieval, we use the auxiliary data to fine-tune the GP model that later initializes a better starting point for the IS model (we call this method Aux.), rather than directly merge their data into the unseen task. We will show the comparison in our experiments.

4. BiLevel I2I translation models

To evaluate the effectiveness of our BiLevel image translation paradigm, we choose two popular I2I translation problems: pose guided I2I translation [22] and pix2pix I2I translation [14]. We test our approach with two classical models, *i.e.* PG^2 model [22] and Pix2pix model [14],

on mid-resolution images. Note that it is straightforward to apply our approach to higher resolution models, *e.g.* Pix2pixHD [43].

4.1. Architectures

Our GP model and IS model share the same architecture which contains a generator G and a discriminator D. We apply the proposed BiLevel paradigm based on this architecture, and thus obtain four networks: GP-generator G_{GP} and GP-discriminator D_{GP}; IS-generator G_{IS} and IS-discriminator D_{IS}. Specifically, G_{IS} and D_{IS} work for the instance-specific translation and executes the translation task in a single scene. G_{GP} and D_{GP} provide a general-purpose basis as initialization for fast adaptation with only a few images in a new scene.

For the PG^2 model, the input of the generator is composed of a reference image and structure information. The reference image indicates the instance-specific content, *e.g.* a street scene or a person identity. The generator G contains a residual U-net architecture consisting of an encoder and a decoder, both of which have several modules. Each module consists of one down/up-sampling convolution block and one residual block following [22]. As to the Pix2pix model, the architecture is similar except that the input only contains the structure information and the U-net architecture consists of convolution blocks instead of residual blocks, following [14].

For both PG^2 and Pix2pix models, the discriminator D is a fully convolutional binary classifier whose input is either a fake generated image or a real ground truth image. It contains a series of down-sampling convolution blocks following [14]. Such fully convolutional discriminators can not only adapt to different image resolutions, but also penalize local structures.

4.2. Losses

We apply the L1 loss L_1 and the adversarial loss L_{adv} to optimize the models as follows,

$$L_1(G) = \mathbb{E}_{x_S, x_R, y}\big[\|G(x_S, x_R) - y\|_1\big], \qquad (4)$$

$$\begin{aligned} L_{adv}(G, D) = & \mathbb{E}_{x_S, y}\big[\log D(x_S, y)\big] \\ & + \mathbb{E}_{x_S, x_R}\big[\log\big(1 - D(G(x_S, x_R))\big)\big], \quad (5) \end{aligned}$$

where x_S and x_R denote the input structure information and reference image, resp., and y denotes the target image. In the adversarial loss, G tries to minimize this objective against an adversarial D that tries to maximize it. In addition, we apply a recently proposed perceptual loss L_P, the Learned Perceptual Image Patch Similarity (LPIPS) metric [50], which is obtained by computing the $L2$ distance between the weighted deep features of images, as follows

$$L_P(G) = \mathbb{E}_{x_S, x_R, y}\big[\|\phi(G(x_S, x_R)) - \phi(y)\|_2^2\big], \qquad (6)$$

where ϕ denotes the deep feature extractor which is an ImageNet pre-trained VGG16 [36] in our experiments. The full objective is thus as follows,

$$\min_G \max_D L(G, D) = L_1(G) + \lambda_a L_{adv}(G, D) + \lambda_b L_P(G), \quad (7)$$

where λ_a and λ_b denote weighting hyperparameters.

The learning objective is the same for the GP model and IS model. However, the optimization steps are different for them. Details are given in the following Sec. 4.3.

4.3. Algorithms

Alg. 1 and Alg. 2 summarize the optimization procedures of training phase and test phase, resp.

Training phase. The details of the training phase are given in Alg. 1. To provide a good initialization for the GP model, we first pre-train it on a large-scale dataset (lines 1-5). The pre-trained model is actually the conventional PG^2 or Pix2pix model, without BiLevel learning paradigm. Then, we alternate the update of the GP model and the IS model (lines 6-22).

Test phase. The details of the test phase are given in Alg. 2. For each test task, we first select the auxiliary tasks and then fine-tune the GP model (lines 2-4). We use the fine-tuned GP model parameters to initialize a new IS model (lines 5-7). After the IS model has been adapted to the target task, we apply it to the generated images for the final evaluation (lines 8-13).

The optimization step sizes α and β are adaptively controlled by the Adam [16] method. Basically, Adam [16] maintains a per-parameter learning rate adapted based on the average of the recent first and second moments of the gradients. This design enables fast and stable convergence of the optimization of both GP and IS models.

5. Experiments

We evaluate our BiLevel image translation paradigm in the instance-specific settings [1]. We conduct extensive experiments for tackling two challenging problems: image-to-image translation (*i.e.* Pix2pix model [14]) and structure guided image translation (*i.e.* PG^2 model [22]). We perform 1-shot and 5-shot I2I translation on face images, and 1-shot translation on street view images for which we conduct the model test in a challenging cross-dataset setting [2].

5.1. Datasets

Face translation. We use the FaceForensics [34] dataset which contains 704 training videos and 150 test videos

---

[1] In each task, an IS model initialized by the trained GP model gets a fast adaptation based on a few training images, then its result images are used for the evaluation.

[2] For more results, we refer the reader to the supplementary material.

Algorithm 1: Training phase

Input: Translation training tasks $\{\mathcal{T}\}$ and corresponding dataset $\mathcal{D}$; initial learning rate α, β for IS model and GP model;

Output: Generator G_{GP}, Discriminator D_{GP}

1 %Pre-training
2 Randomly initialize G_{GP} and D_{GP};
3 **for** *samples in $\mathcal{D}$* **do**
4   Update G_{GP}, D_{GP} by Eq. 7 with step size α;
5 **end**
6 %Alternating learning
7 **for** *meta-batches* **do**
8   %IS-update
9     $G_{IS} \leftarrow G_{GP}$;
10     $D_{IS} \leftarrow D_{GP}$;
11   Pick up a task $\mathcal{T}$ from $\{\mathcal{T}\}$;
12   Sample training/test data $\mathcal{T}^{(tr)}/\mathcal{T}^{(te)}$ from $\mathcal{T}$;
13   **for** *samples in $\mathcal{T}^{(tr)}$* **do**
14   Update (G_{IS}, D_{IS}) by $L_{\mathcal{T}}(G_{IS}^{\mathcal{T}}, D_{IS}^{\mathcal{T}}; \mathcal{T}^{(tr)})$ in Eq. 7 with step size α.;
15   **end**
16   %GP-update
17   **for** *samples in $\mathcal{T}^{(te)}$* **do**
18   Compute $L_{\mathcal{T}}(G_{IS}^{\mathcal{T}}, D_{IS}^{\mathcal{T}}; \mathcal{T}^{(te)})$ by Eq. 7;
19   **end**
20     $L_{GP} \leftarrow Average(\{L_{\mathcal{T}}(G_{IS}^{\mathcal{T}}, D_{IS}^{\mathcal{T}}; \mathcal{T}^{(te)})\})$;
21   Update (G_{GP}, D_{GP}) by L_{GP} with step size β;
22 **end**

---

from news flashes by different reporters for face translation. We apply OpenPose [6] to detect facial landmarks for constructing the semantic segmentation maps with 7 classes (regions): eyebrows, eyes, nose, lips, inner-mouth, face and background as shown in Fig. 3. In addition, the exact face region is cropped according to the coordinate boundary of facial landmarks and then resized to 128×128 pixels.

To increase the number of training tasks, we divide each video in the training set into 7 short clips and apply different random image augmentations to 6 of them. For each video in the test set, we just use the original one. After filtering out the facial landmark failures, we have 4,767 video clips for training and 148 for test, resp.

Street view translation. We use the training set of the Berkeley Deep Drive 100K (BDD100K) dataset [47] for training. BDD100K is a real-world dataset with 7,000 segmentation-image training pairs which are captured at diverse driving scenes and under various weather conditions. Since almost each segmentation-image pair comes from different video sequences captured at different scenes, we treat each pair as a single task for the 1-shot experiment. For testing, the performance is evaluated in a cross-dataset set-

Algorithm 2: Test phase

Input: Translation train tasks $\{\mathcal{T}\}$, test tasks $\{\mathcal{T}_{uns}\}$ and corresponding dataset $\mathcal{D}, \mathcal{D}_{uns}$; initial learning rate α, β for IS model and GP model;

Output: Generated results $\hat{y}$

1 **for** *each unseen task $\mathcal{T}_{uns}$* **do**
2   %GP-finetune
3   Select K auxiliary tasks $\{\mathcal{T}_{aux}\}_K$ from $\{\mathcal{T}\}$ by Eq. 3;
4   Alternatingly update (G_{GP}, D_{GP}) with $\{\mathcal{T}_{aux}\}_K$ by Alg. 1 line 6-22;
5   %IS-update
6     $G_{IS} \leftarrow G_{GP}$;
7     $D_{IS} \leftarrow D_{GP}$;
8   Sample training/test data $\mathcal{T}_{uns}^{(tr)}/\mathcal{T}_{uns}^{(te)}$ from $\mathcal{T}_{uns}$;
9   **for** *samples in $\mathcal{T}_{uns}^{(tr)}$* **do**
10   Update (G_{IS}, D_{IS}) by $L_{\mathcal{T}}(G_{IS}^{\mathcal{T}}, D_{IS}^{\mathcal{T}}; \mathcal{T}^{(tr)})$ in Eq. 7 with step size α;
11   **end**
12   %Generate results
13     $\hat{y} \leftarrow G_{IS}(x_S, x_R)$
14 **end**

---

ting, where the model trained on the BDD dataset is tested on another dataset – CamVid dataset [5]. CamVid contains 704 annotated images with 33 semantic classes captured in 5 different street view videos. Specifically, we divide the CamVid videos into small video clips and each clip contains 10 successive annotated images. To compromise between the CamVid and BDD datasets, we map the 33 classes of CamVid to the 20 classes of BDD. Images are resized and then center-cropped to the size of 128×256. For each video clip, we randomly select 1 image for 1-shot learning and 5 images for final evaluation.

5.2. Metrics

We provide both qualitative and quantitative results. As to quantitative evaluation, three well-known objective image quality metrics are used: Structural Similarity (SSIM) [46], Mean Square Error (MSE), Peak Signal-to-Noise Ratio (PSNR), and one recently proposed perceptual distance: the Learned Perceptual Image Patch Similarity distance (LPIPS) [50] which has been demonstrated to correlate well with human perceptual similarity [50]. The LPIPS is given by a weighted L2 distance between deep features of images, where we use AlexNet [17] pre-trained on ImageNet as feature extractor, similar to [13]. For SSIM and PSNR, higher scores are better. For LPIPS and MSE, lower scores are better. We report the mean scores across 2960 and 1380 randomly selected pairs (20 pairs per task) for the test set for face and street view translation, resp.

5.3. Implementation and setting details

For model optimization, we use the Adam [16] optimizer with $\beta_1 = 0.5$, $\beta_2 = 0.999$ and the initial learning rate of 0.0001. The mini-batch size for optimizing the general-purpose (GP) model is set to 5 (*i.e.* task number) and the mini-batch size for optimizing the instance-specific (IS) model is also set to 5 (*i.e.* image pair number). The GP model is optimized with $50k$ and $20k$ iterations during pre-training and meta-training, resp. The IS model is optimized with 20 iterations for each new task. The loss weights are set to $\lambda_a = 10$ and $\lambda_a = 2$.

Data augmentation. To encourage the network to handle large displacements, we perform data augmentation techniques: horizontal flip, crop, rotate. As to the PG^2 model, we only have one segmentation-image pair of a specific instance. Therefore, we manually generate the reference image from the ground truth image via data augmentation, *i.e.* random horizontal flip[3], crop and rotation, similar to the method used in [30].

Ablation settings.

PG^2: the PG^2 model [22] trained with training data.

Pix2pix: the Pix2pix model [14] trained with training data.

PG^2 (n-shot): the PG^2 model trained with training data, and then fine-tuned with n-shot target scene data during test phase.

Pix2pix (n-shot): the Pix2pix model trained with training data, and then fine-tuned with n-shot target scene data during test phase.

BiL-X w/o Aux (n-shot): the model X (*i.e.* PG^2 or Pix2pix) trained by our BiLevel paradigm (without using auxiliary tasks) with training data, and then fine-tuned with n-shot target scene data during test phase.

BiL-X (n-shot): the model X (*i.e.* PG^2 or Pix2pix) trained by our BiLevel paradigm with training data, and then fine-tuned with n-shot target scene data during test phase.

BiL-X (GP n-shot): the GP model of model X (*i.e.* PG^2 or Pix2pix) trained by our BiLevel paradigm with training data.

5.4. Face image translation

For face translation, we evaluate our BiLevel approach in all ablative settings. We use the 5-shot and 1-shot data on the FaceForensics dataset.

Qualitative evaluation. As shown in Fig. 3, the results generated by our BiL-PG^2(5-shot) model are more realistic and look quite closer to the target images. For example, comparing to the results of PG^2(5-shot), our results have the more and better-looking instance-specific details, such as the hair style and gender of ID-1, and the skin color, facial feature and age of ID-2. Compared to the ablative model named

---

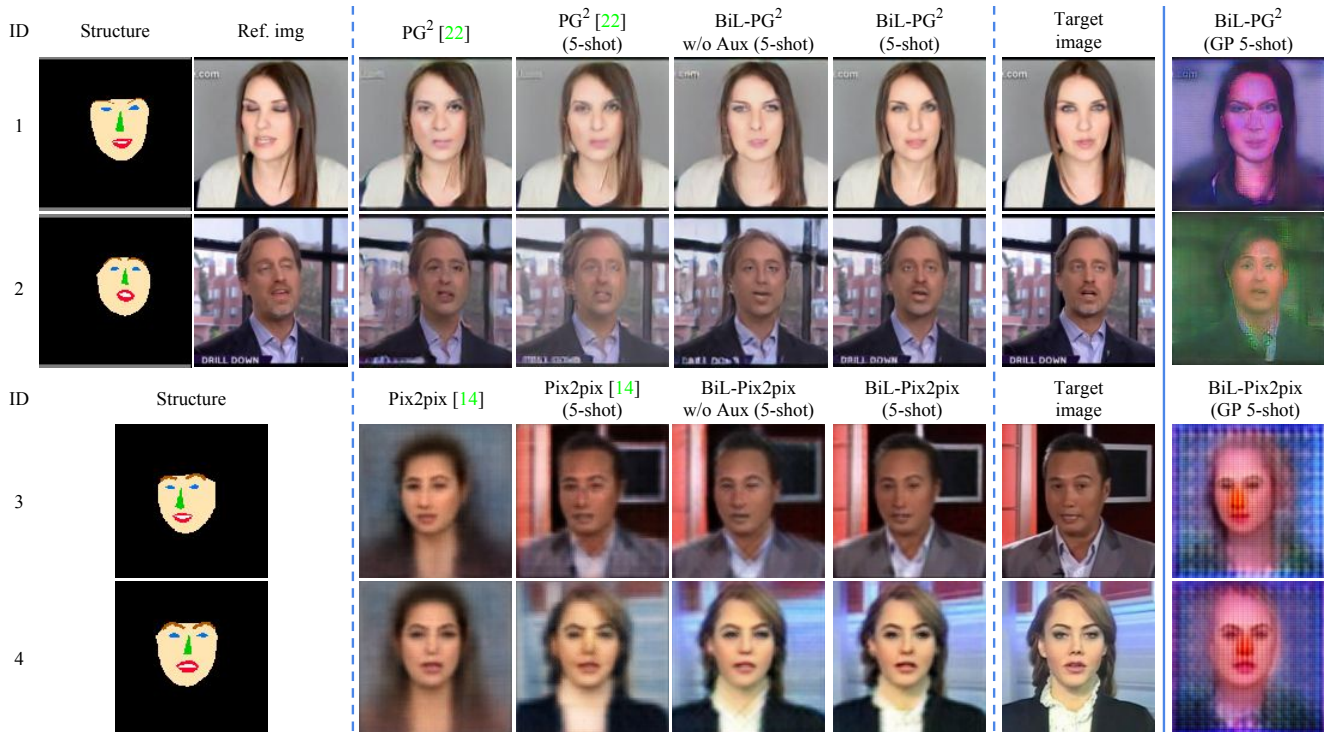[3]The horizontal flip is not applied to the reference image of face translation.

Figure 3. The 5-shot face image translation results on the FaceForensics dataset. From left to right: inputs, results of baselines and our methods, ground truth, and the output of GP model. Among the results, the second column denotes the baseline that is fine-tuned by 5-shot images of the target identity.
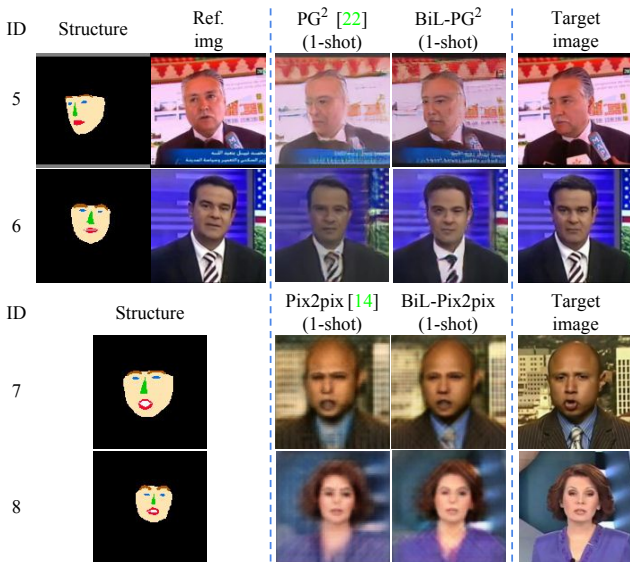

Figure 4. The 1-shot face image translation results on the Face-Forensics dataset. From left to right: inputs, results of fine-tuned baseline and ours, and ground truth.

BiL-PG^2 w/o Aux(5-shot), we have the conclusion that the auxiliary fine-tuning on the GP model greatly reduces the blurriness and artifacts.

As to Pix2pix model, the overall translation quality is inferior to PG^2 model due to the lack of reference image as input. Relatively, our BiL-Pix2pix(5-shot) generates better results than the fine-tuned Pix2pix model (Pix2pix(5-shot)). For example, the facial parts and hair styles of ID-3 and ID-4 have sharper appearances.

Additionally, Fig. 4 shows that our approach also shows reasonably good results in the most extreme 1-shot case. For example, BiL-PG^2(1-shot) successfully translates the poses of ID-5 and ID-6 It preserves better skin colors and more natural appearances of the faces, than the fine-tuned baseline, PG^2(1-shot). Overall, the results above show that our BiL training approach brings impressive improvements for image-to-image translation in the extreme settings with scarce training data. Its results are not only with higher image quality but also maintaining more identity-specific details.

Quantitative evaluation. The quantitative results are consistent with the qualitative visualizations. The 5-shot results are given in Tab. 1. We can see that our BiL-PG^2(5-shot) greatly reduce the LPIPS score from 0.076 to 0.053 (30.2%) and increase SSIM score from 0.681 to 0.756 (11.0%), comparing to the baseline model PG^2(5-shot). Our BiL-Pix2pix(5-shot) also significantly reduces the LPIPS score by 40.4% and increases the SSIM score by 19.4%, comparing to the baseline model Pix2pix(5-shot). In addition, the results of 1-shot setting are given in Tab. 2. As we men-
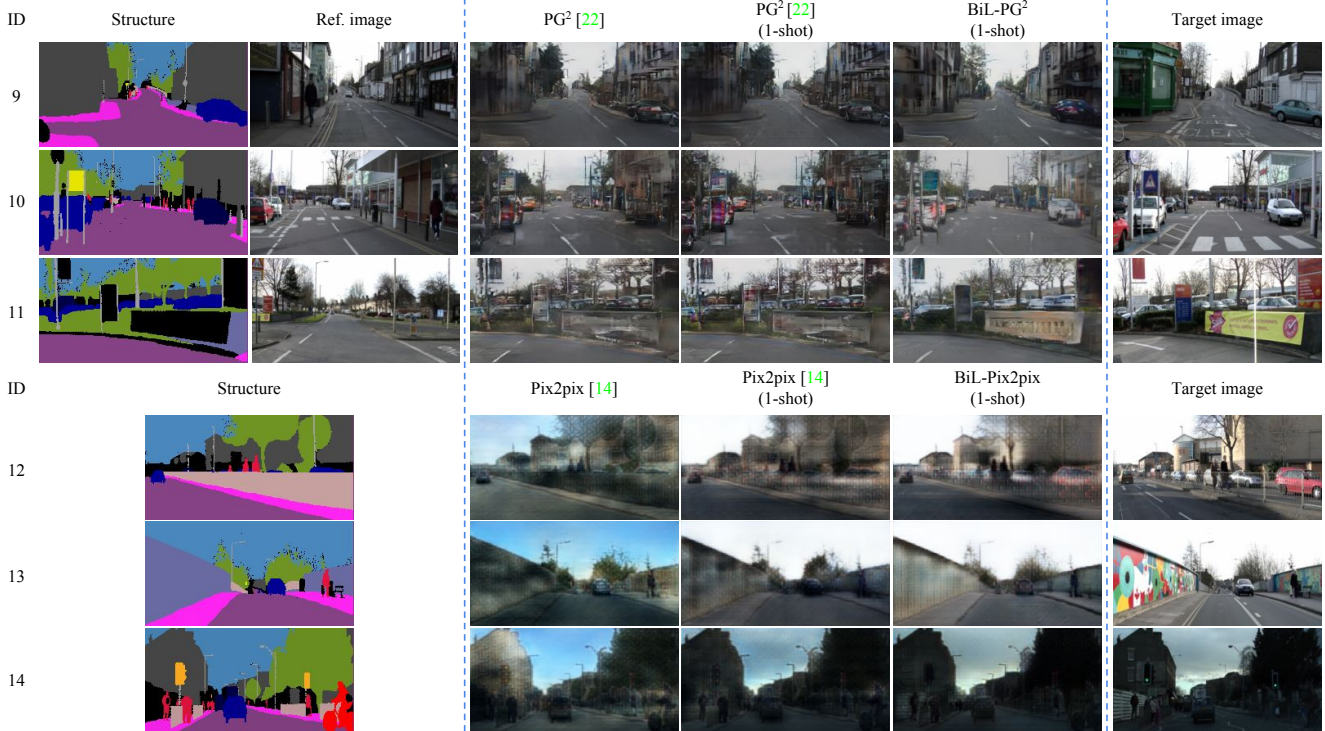
Figure 5. The 1-shot street-view image translation results in the challenging cross-dataset setting. From left to right: inputs, results of baselines and ours, and the ground truth.

Model	LPIPS↓	SSIM↑	MSE↓	PSNR↑
PG^2 [22]	0.106	0.579	0.066	18.549
PG^2 [22] (5-shot)	0.076	0.681	0.040	20.967
BiL-PG^2 (5-shot) w/o Aux	0.054	0.755	0.028	22.987
BiL-PG^2 (5-shot)	**0.053**	**0.756**	**0.027**	**23.176**
Pix2pix [14]	0.357	0.357	0.271	12.106
Pix2pix [14] (5-shot)	0.198	0.638	0.044	20.108
BiL-Pix2pix w/o Aux (5-shot)	0.118	0.757	0.027	22.725
BiL-Pix2pix (5-shot)	**0.116**	**0.762**	**0.026**	**22.845**

Table 1. The evaluation scores for 5-shot face image translation on the FaceForensics dataset.

Model	LPIPS↓	SSIM↑	MSE↓	PSNR↑
PG^2 [22] (1-shot)	0.100	0.584	0.065	18.685
BiL-PG^2 (1-shot) w/o Aux	0.085	0.630	0.050	19.887
BiL-PG^2 (1-shot)	**0.076**	**0.655**	**0.048**	**20.134**
Pix2pix [14] (1-shot)	0.217	0.579	0.058	18.880
BiL-Pix2pix (1-shot) w/o Aux	0.170	0.639	0.048	20.011
BiL-Pix2pix (1-shot)	**0.162**	**0.655**	**0.046**	**20.253**

Table 2. The evaluation scores for 1-shot face image translation on the FaceForensics dataset.

tioned, the LPIPS score is quite close to humans' evaluation. The large gains of LPIPS scores brought by our approach are greatly consistent with our improved visualization results in the figures.

Analysis of the general translation knowledge. In the last column of Fig. 3, we show the GP model results which demonstrate the visualization of the general knowledge encoded in our GP model, noting that the detailed appearance on the image is conditioned on the input image. We can see that this knowledge meets our expectation of an average-looking face with a random blurry background.

5.5. Street view image translation

We evaluate our BiL paradigm on the more challenging street view image translation task, for which the images usually contain a large variety of objects. We use a challenging cross-dataset setting: the GP model is trained on the BDD100K dataset then is tested on the CamVid dataset. Due to the lack of annotated segmentation, we conduct only the 1-shot experiments.

Qualitative results. The 1-shot results are given in Fig. 5. Comparing with the fine-tuned baseline models PG^2(1-shot) and Pix2pix(1-shot), we can observe that our BiL models generate clearly more realistic images which contain a lot of street objects with various appearances. For example, BiL-PG^2(1-shot) generates the trees with very similar appearances as in the target image (ID-9), buildings with sharp outlines (ID-10), cars with clear shapes (ID-11).

Quantitative evaluation. The overall quantitative results of street view translation are provided in Tab. 3. It is easy to see that these are globally inferior to those of face transla-

Model	LPIPS↓	SSIM↑	MSE↓	PSNR↑
PG^2 [22]	0.194	0.373	0.060	15.652
PG^2 [22] (1-shot)	0.174	0.400	0.052	16.429
BiL-PG^2 (1-shot) w/o Aux	0.166	0.434	0.045	16.933
BiL-PG^2 (1-shot)	**0.160**	**0.445**	**0.040**	**17.448**
Pix2pix [14]	0.248	0.347	0.090	13.721
Pix2pix [14] (1-shot)	0.203	0.459	0.040	17.883
BiL-Pix2pix (1-shot) w/o Aux	0.184	0.507	0.028	**19.327**
BiL-Pix2pix (1-shot)	**0.184**	**0.509**	**0.028**	19.322

Table 3. The evaluation scores for 1-shot street-view image translation in the challenging cross-dataset setting.

tion (Tab. 2). This reflects the greater challenge brought by the more complex data and more challenging cross-dataset setting. When comparing the results in Tab. 3, we can conclude that our BiL models still generalize well in this harder setting, e.g. BiL-PG^2(1-shot) achieves a relative improvement rate of 8.0% (from 0.174 to 0.160) for reducing LPIPS, and 11.3% (from 0.400 to 0.445) for improving SSIM.

6. Conclusion

We introduced a BiLevel learning paradigm to learn general-purpose knowledge and fast model adaptability. With the help of our BiLevel paradigm, the IS model of a new task can obtain the general-purpose experience from the GP model, and then quickly learn new instance-specific knowledge with only few-shot data. Numerous quantitative and qualitative results demonstrate the effectiveness of our method. In addition, the images generated by our GP model can visualize the general-purpose knowledge automatically.

Acknowledgments

References

[1] A. Almahairi, S. Rajeswar, A. Sordoni, P. Bachman, and A. Courville. Augmented cyclegan: Learning many-to-many mappings from unpaired data. In *ICML*, 2018. 2

[2] M. Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. In *ICLR*, 2017. 3

[3] S. Benaim and L. Wolf. One-shot unsupervised cross domain translation. 2018. 1

[4] S. Benaim and L. Wolf. One-shot unsupervised cross domain translation. In *NIPS*, 2018. 2

[5] G. J. Brostow, J. Shotton, J. Fauqueur, and R. Cipolla. Segmentation and recognition using structure from motion point clouds. In *ECCV*, 2008. 6

[6] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017. 5

[7] H. Edwards and A. Storkey. Towards a neural statistician. In *ICLR*, 2017. 2

[8] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 2, 4

[9] E. Gavves, T. Mensink, T. Tommasi, C. G. Snoek, and T. Tuytelaars. Active transfer learning with zero-shot priors: Reusing past datasets for future tasks. In *ICCV*, 2015. 2

[10] S. Gidaris and N. Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018. 2

[11] E. Grant, C. Finn, S. Levine, T. Darrell, and T. L. Griffiths. Recasting gradient-based meta-learning as hierarchical bayes. In *ICLR*, 2018. 2

[12] L. B. Hewitt, M. I. Nye, A. Gane, T. Jaakkola, and J. B. Tenenbaum. The variational homoencoder: Learning to learn high capacity generative models from few examples. In *UAI*, 2018. 2

[13] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018. 2, 6

[14] P. Isola, J. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 1, 2, 3, 4, 5, 6, 8, 9

[15] D. Joo, D. Kim, and J. Kim. Generating a fusion image: One's identity and another's shape. In *CVPR*, 2018. 2

[16] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5, 6

[17] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 6

[18] H.-Y. Lee, H.-Y. Tseng, J.-B. Huang, M. Singh, and M.-H. Yang. Diverse image-to-image translation via disentangled representations. In *ECCV*, 2018. 2

[19] Y. Lee and S. Choi. Gradient-based meta-learning with learned layerwise metric and subspace. In *ICML*, 2018. 2

[20] L. Liu, W. Xu, M. Zollhoefer, H. Kim, F. Bernard, M. Habermann, W. Wang, and C. Theobalt. Neural animation and reenactment of human actor videos. *arXiv preprint arXiv:1809.03658*, 2018. 2

[21] L. Ma, X. Jia, S. Georgoulis, T. Tuytelaars, and L. Van Gool. Exemplar guided unsupervised image-to-image translation with semantic consistency. *arXiv preprint arXiv:1805.11145*, 2018. 2

[22] L. Ma, X. Jia, Q. Sun, B. Schiele, T. Tuytelaars, and L. Van Gool. Pose guided person image generation. In *NIPS*, 2017. 1, 2, 3, 4, 5, 6, 8, 9

[23] L. Ma, Q. Sun, S. Georgoulis, L. Van Gool, B. Schiele, and M. Fritz. Disentangled person image generation. In *CVPR*, 2018. 2

[24] A. Mehrotra and A. Dukkipati. Generative adversarial residual pairwise networks for one shot learning. *arXiv preprint arXiv:1703.08033*, 2017. 1

[25] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 2

[26] M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *CVPR*, 2014. 2

[27] S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang. Domain adaptation via transfer component analysis. *IEEE Trans. Neural Networks*, 22(2):199–210, 2011. 2

[28] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2010. 2, 3

[29] S. Qiao, C. Liu, W. Shen, and A. Yuille. Few-shot image recognition by predicting parameters from activations. *arXiv preprint arXiv:1706.03466*, 2, 2017. 2

[30] A. Raj, P. Sangkloy, H. Chang, J. Lu, D. Ceylan, and J. Hays. Swapnet: Garment transfer in single view images. In *ECCV*, 2018. 6

[31] K. Rakelly, E. Shelhamer, T. Darrell, A. A. Efros, and S. Levine. Few-shot segmentation propagation with guided networks. *arXiv preprint arXiv:1806.07373*, 2018. 2

[32] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. 2

[33] D. J. Rezende, S. Mohamed, I. Danihelka, K. Gregor, and D. Wierstra. One-shot generalization in deep generative models. In *ICML*, 2016. 2

[34] A. Rössler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Nießner. Faceforensics: A large-scale video dataset for forgery detection in human faces. *arXiv preprint arXiv:1803.09179*, 2018. 2, 5

[35] B. Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009. 2

[36] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 5

[37] Q. Sun, Y. Liu, T.-S. Chua, and B. Schiele. Meta-transfer learning for few-shot learning. In *CVPR*, 2019. 2

[38] Q. Sun, L. Ma, S. J. Oh, L. V. Gool, B. Schiele, and M. Fritz. Natural and effective obfuscation by head inpainting. In *CVPR*, 2018. 3

[39] Y. Taigman, A. Polyak, and L. Wolf. Unsupervised cross-domain image generation. In *ICLR*, 2017. 2

[40] S. Vijayanarasimhan, P. Jain, and K. Grauman. Far-sighted active learning on a budget for image and video recognition. In *CVPR*. IEEE, 2010. 2

[41] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra. Matching networks for one shot learning. In *NIPS*, 2016. 3

[42] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, et al. Matching networks for one shot learning. In *NIPS*, 2016. 2

[43] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *CVPR*, 2018. 1, 2, 3, 4

[44] W. Wang, X. Alameda-Pineda, D. Xu, P. Fua, E. Ricci, and N. Sebe. Every smile is unique: Landmark-guided diverse smile generation. In *CVPR*, 2018. 2

[45] Y. Wang, C. Wu, L. Herranz, J. van de Weijer, A. Gonzalez-Garcia, and B. Raducanu. Transferring gans: generating images from limited data. In *ECCV*, 2018. 2

[46] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Processing*, 13(4):600–612, 2004. 6

[47] F. Yu, W. Xian, Y. Chen, F. Liu, M. Liao, V. Madhavan, and T. Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2018. 2, 5

[48] A. R. Zamir, A. Sax, W. Shen, L. J. Guibas, J. Malik, and S. Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*, 2018. 2

[49] R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In *ECCV*, 2016. 2

[50] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 4, 6

[51] B. Zhao, B. Chang, Z. Jie, and L. Sigal. Modular generative adversarial networks. In *The European Conference on Computer Vision (ECCV)*, 2018. 2

[52] J. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 1, 2

[53] J.-Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, and E. Shechtman. Toward multimodal image-to-image translation. In *NIPS*, 2017. 2