

Learning to Transfer: Unsupervised Domain Translation via Meta-Learning

Jianxin Lin,¹ Yijun Wang,² Zhibo Chen,^{1*} Tianyu He³

¹CAS Key Laboratory of Technology in Geo-spatial Information Processing and Application System,
University of Science and Technology of China

²University of Science and Technology of China, ³DAMO Academy, Alibaba Group
{linjx, wyjun}@mail.ustc.edu.cn, chenzhibo@ustc.edu.cn, timhe.hty@alibaba-inc.com

Abstract

Unsupervised domain translation has recently achieved impressive performance with Generative Adversarial Network (GAN) and sufficient (unpaired) training data. However, existing domain translation frameworks form in a disposable way where the learning experiences are ignored and the obtained model cannot be adapted to a new coming domain. In this work, we take on unsupervised domain translation problems from a meta-learning perspective. We propose a model called Meta-Translation GAN (MT-GAN) to find good initialization of translation models. In the meta-training procedure, MT-GAN is explicitly trained with a primary translation task and a synthesized dual translation task. A cycle-consistency meta-optimization objective is designed to ensure the generalization ability. We demonstrate effectiveness of our model on ten diverse two-domain translation tasks and multiple face identity translation tasks. We show that our proposed approach significantly outperforms the existing domain translation methods when each domain contains no more than ten training samples.

1 Introduction

Unsupervised domain translation tasks (Zhu et al. 2017; Choi et al. 2018), which aim at learning a mapping that can transfer images from a source domain to a target domain using unpaired training data only, have been widely investigated in recent years. However, current literature focuses on learning a model for a specific translation task without considering the generalization ability to other tasks. In comparison, human intelligence has the ability to quickly learn new concepts with the prior experiences. Taking painting as an example, after being taught how to paint in Monet’s style, people have learned the basic skills of painting with usage of painting brush, palette, etc. When being required to learn painting in Van Gogh’s style, we do not need to learn from scratch for all these skills. Instead, we can quickly adapt to this new task by viewing only a few Van Gogh’s paintings, extracting his painting style, combining the basic painting knowledge we learned before with the new style and eventually know how to draw in Van Gogh’s way.

In this paper, we take a step towards unsupervised domain translation (UDT) problems from a meta-learning perspective, aiming to effectively leverage learning experiences from previous domain translation tasks. Briefly, the problem is that we have several domain translation tasks where each task only consists of limited samples. We have to learn a meta-model from these tasks, that can be quickly adapted to unseen domain translation tasks. To tackle this problem, we propose a method called Meta-Translation Generative Adversarial Network (MT-GAN) that are robust to different task contexts. The proposed model contains two meta-learners: a meta-generator G which keeps the memory of prior translation experiences and a meta-discriminator D which teaches G how to quickly generalize to a new task. Our approach leverages both model-agnostic meta-learning algorithm (MAML) (Finn, Abbeel, and Levine 2017) and Generative Adversarial Network (GAN) (Goodfellow et al. 2014) to iteratively update G and D . To achieve that, within a meta-training iteration, for a specific translation task, we synthesize its dual translation task with current states of MT-GAN and train these two tasks in a dual learning form (Zhu et al. 2017; He et al. 2016a). Then we design a meta-optimization objective to evaluate the performance of fine-tuned MT-GAN, and minimize the expected losses on the meta-testing samples with respect to parameters of MT-GAN, which ensures that the direction taken to fine-tuning leads to a good generalization performance.

We extensively evaluate the effectiveness and generalizing ability of the proposed MT-GAN algorithm on two kinds of translation task distributions. The first one contains 10 diverse two-domain translation tasks covering a wide range of scenarios, including labels \leftrightarrow photos, horses \leftrightarrow zebras, summer \leftrightarrow winter, etc. The second one is established by repeatedly sampling two arbitrary identities, which forms a two-domain translation task, from a multiple identity dataset (Ng and Winkler 2014). For a translation task T , we take the other 9 different tasks as the training dataset and test the meta-learned parameter initialization on task T . Our experiments only use 10 samples at most in an image domain of a translation task and show that the proposed meta-learning approach outperforms ordinary domain translation models, such as CycleGAN (Zhu et al. 2017) and StarGAN (Choi et

*Corresponding author

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

al. 2018).

We summarize our contributions as follow: 1) We approach the unsupervised domain translation problem from a meta-learning perspective, which allows effective usage of learning experiences from previous domain translation tasks when taking up a new domain translation task; 2) We propose a method (MT-GAN) that jointly trains two meta-learners in an adversarial and dual form, which to the best of our knowledge has not been explored; 3) We extensively verify the effectiveness of our meta-learning based approach on a wide range of translation tasks.

2 Related works

Generative Adversarial Network In recent years, the generative adversarial network (GAN) model (Goodfellow et al. 2014) has gained a wide range of interests in generative modeling. In a GAN, a generator is trained to produce fake but plausible images, while a discriminator is trained to distinguish difference between real and fake images. The conditional version of a GAN, called a conditional generative adversarial network (CGAN) (Mirza and Osindero 2014), is a model in which the generator is feeded with noise vector together with additional data (e.g., class labels) that conditions on both the generator and discriminator. Deep convolutional generative adversarial network (DCGAN) (Radford, Metz, and Chintala 2015) is an extensive exploration of convolution neural network architectures in GANs and contributes to improve the quality of image synthesis. GANs have been successfully leveraged in many image generation applications (Pumarola et al. 2018; Wang et al. 2018; Wu et al. 2016; Yu et al. 2018). Our method adopts the adversarial loss to render images from the generators to be close to real in the target domain and make meta-training performance improve meta-learners’ generalization.

Unsupervised Domain Translation Domain translation has also achieved impressive performance thanks to recent development of GANs and availability of sufficient training data. (Isola et al. 2017) proposed a general conditional GAN (Pix2Pix) framework for a wide range of supervised domain translation tasks. Since obtaining an amount of paired training data can be difficult and impractical for many domain translation tasks, DualGAN (Yi et al. 2017), DiscoGAN (Kim et al. 2017) and CycleGAN (Zhu et al. 2017) were proposed to learn two cross-domain translation models that obey the cycle consistent rule from unpaired data. (Choi et al. 2018) further proposed a unified unpaired domain translation model (StarGAN) to perform domain translation for multiple domains. (Liu et al. 2018) also proposed a method called UFDN to learn domain-invariant representation for multiple domain translation and can perform diverse domain translation and manipulation. A related work similar to our work may be (Benaim and Wolf 2018), in which they proposed a one-shot cross-domain translation which transfers one and only one image in a source domain to a target domain with sufficient data in the target domain. However, all these existing unsupervised domain translation models mainly rely on training data of current translation task, and omit to utilize the meta-knowledge from prior learning experiences like humans. In this work, we focus on image translation that incorporates

the prior learning experiences from other translation tasks for new translation tasks’ learning.

Meta-Learning Meta-learning, which aims to learn a particular process to adjust meta-learners that perform well on a new task, can be traced back to early works (Schmidhuber 1987; Bengio, Bengio, and Cloutier 1990; Bengio et al. 1992). Some recent meta-learning studies have focused on learning a shared metric by comparing similarity among data samples. Specifically, (Vinyals et al. 2016) proposed Matching Networks that learns an embedding function and measures similarity using the cosine distance in an attention kernel. (Snell, Swersky, and Zemel 2017) also proposed to compare new examples in a learned metric space but used the Euclidean distance with a linear classifier. Another popular approach to meta-learning is to learn a shared initialization of network parameters. For example, (Finn, Abbeel, and Levine 2017) presented model-agnostic meta-learning (MAML) to optimize the parameters of a meta-learner with the objective of maximizing its performance on a new task after a small number of gradient steps. Several other methods (Andrychowicz et al. 2016; Mishra et al. 2017) utilized an additional memory-based network (e.g., LSTM) as the meta-learner. Observing that meta-learning methods usually require labeled datasets, recent works (Hsu, Levine, and Finn 2018; Metz et al. 2018) also proposed to tackle the unsupervised meta-learning. In this paper, we extend the concept of meta-learning to image translation. Specifically, we jointly train two meta-learners in an adversarial and dual form, which to the best of our knowledge has not been done before.

3 Unsupervised Domain Translation via Meta-Learning

3.1 Problem Formulation

Domain Translation is the problem of finding a meaningful correspondence between two domains. Since paired supervision is not available in a majority of settings, many works focus on Unsupervised Domain Translation (UDT) where data samples from each domain are unpaired (de Bézenac, Ayed, and Gallinari 2019). Existing unsupervised domain translation models mainly rely on training data of a specific translation task, lacking the ability of utilizing the knowledge from prior learning experiences when taking up alternative translation tasks. To tackle this problem, we approach the unsupervised domain translation problem from a meta-learning perspective, which allows effective usage of learning experiences from previous domain translation tasks. The goal of our method is to first leverage unpaired data for efficient training, and then the obtained model can be applied on a wide range of new domain translation tasks.

Formally, in our meta-learning scenario, assuming that there are a series of tasks following distribution $P(T)$ over the task space \mathcal{T} . For a specific task $T \sim P(T)$, intuitively task T aims at finding a meaningful correspondence between two domains. Specifically, both domains can be expressed as probability distributions $P_T(x)$ and $P_T(y)$ supported over domain spaces \mathcal{X}_T and \mathcal{Y}_T . A task consists in finding a mapping: $\mathcal{X}_T \rightarrow \mathcal{Y}_T$ such that the mapping yields semantically meaningful pairings.

Since we aim at utilizing the learning experience from previous domain translation tasks when take up new translation tasks, in our meta-learning scenario, the training sample consists of a finite set of tasks $\{T_n\}_{n=1}^N$ drawn from $P(T)$, where N denotes the number of training tasks. Concretely, each training task T_n is a tuple $T_n = (S_{T_n}, Q_{T_n})$, where S_{T_n} denotes the support set and Q_{T_n} denotes the query set. Specifically, the support set $S_{T_n} = \{\{x_k\}_{k=1}^K \in \mathcal{X}_{T_n}, \{y_k\}_{k=1}^K \in \mathcal{Y}_{T_n}\}$ contains $2K$ unpaired samples from two different domains \mathcal{X}_{T_n} and \mathcal{Y}_{T_n} . $Q_{T_n} = \{\{x_l\}_{l=1}^L \in \mathcal{X}_{T_n}, \{y_l\}_{l=1}^L \in \mathcal{Y}_{T_n}\}$ is the query set that contains $2L$ unpaired samples from the same two domains. S_{T_n} and Q_{T_n} are disjoint. Our algorithm takes $\{T_n\}_{n=1}^N$ as inputs and produces a learning strategy for two meta-learners, i.e., G and D . In general, the meta-learners iteratively adjust the parameters on data from support set and assess their generalization performance by calculating meta-objective with data from query set. The meta-learners are then improved by considering how the test error changes with respect to the parameters. In effect, the test error on sampled tasks T_n serves as the training error of the meta-learning process. At inference time, suppose that we have a new translation task $T_{N+1} = (S_{T_{N+1}}, Q_{T_{N+1}})$, where $S_{T_{N+1}} = \{\{x_k\}_{k=1}^K \in \mathcal{X}_{T_{N+1}}, \{y_k\}_{k=1}^K \in \mathcal{Y}_{T_{N+1}}\}$ and $Q_{T_{N+1}} = \{\{x_l\}_{l=1}^L \in \mathcal{X}_{T_{N+1}}, \{y_l\}_{l=1}^L \in \mathcal{Y}_{T_{N+1}}\}$ ($S_{T_{N+1}}$ and $Q_{T_{N+1}}$ are disjoint.), the learning strategy should learn a fine-tuned G and a fine-tuned D with $S_{T_{N+1}}$ and accomplishes the translation $X \rightarrow Y$ on $Q_{T_{N+1}}$. That is, the meta-performance is measured by the meta-learners' performance after learning from K samples in each domain. We denote the above process as K -shot domain translation.

3.2 Our Approach

We introduce the formulation of MT-GAN as following: for a K -shot domain translation problem and dataset $\{T_n\}_{n=1}^N$, our goal is to find a meta-generator G that keeps the memory of prior translation experiences and a meta-discriminator D that teaches G how to quickly generalize to a new task. Unlike CycleGAN or StarGAN that utilizes multiple generators (and discriminators) or multiple domain codes to tackle different translation tasks, we only develop one meta-generator G and one meta-discriminator D for different tasks. This is because a initialization is responsible to the learning strategy of all translation tasks and is independent with a specific task. The full algorithm of MT-GAN is outlined in Algorithm 1 in a general case.

Formally, the generator G and discriminator D are parameterized by θ_g and θ_d respectively. In each meta-batch, we sample J tasks for training. For a specific translation task $T \sim P(T)$ in one meta-batch, $T = (S_T, Q_T)$, there are two image domains \mathcal{X}_T and \mathcal{Y}_T . Our primary target is to learn a mapping $F: \mathcal{X}_T \rightarrow \mathcal{Y}_T$ that is derived from G . Since there is no paired data for training, observing that there is naturally a dual task which learns another mapping $H: \mathcal{Y}_T \rightarrow \mathcal{X}_T$ in the reverse direction, we utilize these two translation tasks as a two-agent game for fine-tuning in the meta-training period. Specifically, two discriminators D_Y and D_X are used to render images from the generators to be real in corresponding domain.

Algorithm 1 MT-GAN training process

Require: : Distribution $P(T)$ over domain translation tasks
Require: : Hyperparameters $\alpha, \beta, \lambda_{\text{cyc}}, \lambda_{\text{idt}}, K, J$

- 1: Randomly initialize parameters θ_g of G and θ_d of D
- 2: **while** not converged **do**
- 3: Sample batch of tasks $\{T_j\}_{j=1}^J \sim P(T)$, where J is the meta-batch size.
- 4: Split support set and query set: S_{T_j} and $Q_{T_j} \leftarrow T_j$
- 5: **for all** S_{T_j} **do**
- 6: **Meta-training:**
- 7: Compute initialized parameters $\theta'_{d_Y,0}, \theta'_{d_X,0}, \theta'_{f,0}, \theta'_{h,0}$ with gradient descent by Eqn.(1) and Eqn.(2)
- 8: **for** i in iterations I **do**
- 9: Compute fine-tuned parameters $\theta'_{d_Y,i+1}, \theta'_{d_X,i+1}, \theta'_{f,i+1}, \theta'_{h,i+1}$ with gradient descent by Eqn.(7) and Eqn.(8)
- 10: **end for**
- 11: **Meta-testing:**
- 12: Compute meta-objective $\mathcal{L}_{T_j,I+1}^q$ on Q_{T_j} according to Eqn.(12)
- 13: **end for**
- 14: **Meta-Optimization:**
- 15: Update $\theta_d = \theta_d + \beta \nabla_{\theta_d} \sum_{T_j \in \{T_j\}_{j=1}^J} \mathcal{L}_{T_j,I+1}^q$
- 16: Update $\theta_g = \theta_g - \beta \nabla_{\theta_g} \sum_{T_j \in \{T_j\}_{j=1}^J} \mathcal{L}_{T_j,I+1}^q$
- 17: **end while**

In the initialization step, we initialize the parameters of discriminators and generators as:

$$\theta'_{d_Y,0} = \theta_d + \alpha \nabla_{\theta_d} \mathcal{L}_{T,0}; \quad (1)$$

$$\begin{aligned} \theta'_{d_X,0} &= \theta_{d,wc} + \alpha \nabla_{\theta_{d,wc}} \mathcal{L}_{T,0}, \\ \theta'_{f,0} &= \theta_g - \alpha \nabla_{\theta_g} \mathcal{L}_{T,0}; \\ \theta'_{h,0} &= \theta_{g,wc} - \alpha \nabla_{\theta_{g,wc}} \mathcal{L}_{T,0}, \end{aligned} \quad (2)$$

where α is the learning rate during the meta-training period; $\theta'_{d_Y,0}, \theta'_{f,0}, \theta'_{d_X,0}$ and $\theta'_{h,0}$ are the parameters of D_Y, F, D_X and H respectively after initialization; $\theta_{d,wc}$ and $\theta_{g,wc}$ are the parameters of $wc(D)$ and $wc(G)$, where wc is the network weights copy operation that detaches back-propagation gradient from meta-optimization objectives. G and D can be parameter initialization for any translation task in practice. However, in a specific task T , using G for fine-tuning of both F and H is ambiguous since the meta-optimization objective would become to require G and D to be well adapted for both $\mathcal{X}_T \rightarrow \mathcal{Y}_T$ and $\mathcal{Y}_T \rightarrow \mathcal{X}_T$ at the same time. Therefore, we update G and D only for $\mathcal{X}_T \rightarrow \mathcal{Y}_T$ translation with wc operation, and utilize $\mathcal{Y}_T \rightarrow \mathcal{X}_T$ translation as a dual task. The overall objective $\mathcal{L}_{T,0}$ for training the discriminators and generators at initialization step is given as:

$$\begin{aligned} \mathcal{L}_{T,0}(G, D, \mathcal{X}_T, \mathcal{Y}_T) &= \mathcal{L}_{\text{adv}}(G, D, \mathcal{X}_T, \mathcal{Y}_T) \\ &\quad + \mathcal{L}_{\text{adv}}(wc(G), wc(D), \mathcal{Y}_T, \mathcal{X}_T) \\ &\quad + \lambda_{\text{cyc}} \mathcal{L}_{\text{cyc}}(G, wc(G), \mathcal{X}_T, \mathcal{Y}_T) \\ &\quad + \lambda_{\text{idt}} \mathcal{L}_{\text{idt}}(G, wc(G), \mathcal{X}_T, \mathcal{Y}_T), \end{aligned} \quad (3)$$

where the factors λ_{cyc} and λ_{idt} are used to balance different loss terms \mathcal{L}_{adv} , \mathcal{L}_{cyc} and \mathcal{L}_{idt} . Specifically, the loss

terms \mathcal{L}_{adv} , \mathcal{L}_{cyc} and \mathcal{L}_{idt} are the adversarial loss, the cycle-consistency loss and the identity loss respectively, which are defined as follows:

$$\mathcal{L}_{\text{adv}}(G, D, \mathcal{X}_T, \mathcal{Y}_T) = \mathbb{E}_{y \sim P_T(y)} [\log D(y)] + \mathbb{E}_{x \sim P_T(x)} [\log(1 - D(G(x)))], \quad (4)$$

$$\mathcal{L}_{\text{cyc}}(G_1, G_2, \mathcal{X}_T, \mathcal{Y}_T) = \mathbb{E}_{x \sim P_T(x)} [\|G_2(G_1(x)) - x\|_1] + \mathbb{E}_{y \sim P_T(y)} [\|G_1(G_2(y)) - y\|_1], \quad (5)$$

$$\mathcal{L}_{\text{idt}}(G_1, G_2, \mathcal{X}_T, \mathcal{Y}_T) = \mathbb{E}_{x \sim P_T(x)} [\|G_2(x) - x\|_1] + \mathbb{E}_{y \sim P_T(y)} [\|G_1(y) - y\|_1], \quad (6)$$

where $P_T(x)$ and $P_T(y)$ are the distributions of samples in domain spaces \mathcal{X}_T and \mathcal{Y}_T respectively. Then, we estimate the expectation terms in above equations using the samples in support set S_T .

At iteration $i \geq 0$, we follow the popular unsupervised domain translation model, e.g., CycleGAN, to fine-tune the initialized $D_{\mathcal{X}}$, $D_{\mathcal{Y}}$, F and H to quickly adapt to the task T . Formally, we update the parameters as follow:

$$\theta'_{d_{\mathcal{Y}}, i+1} = \theta'_{d_{\mathcal{Y}}, i} + \alpha \nabla_{\theta'_{d_{\mathcal{Y}}, i}} \mathcal{L}_{T, i+1}; \quad (7)$$

$$\theta'_{d_{\mathcal{X}}, i+1} = \theta'_{d_{\mathcal{X}}, i} + \alpha \nabla_{\theta'_{d_{\mathcal{X}}, i}} \mathcal{L}_{T, i+1},$$

$$\theta'_{f, i+1} = \theta'_{f, i} - \alpha \nabla_{\theta'_{f, i}} \mathcal{L}_{T, i+1}; \quad (8)$$

$$\theta'_{h, i+1} = \theta'_{h, i} - \alpha \nabla_{\theta'_{h, i}} \mathcal{L}_{T, i+1}.$$

The training objective $\mathcal{L}_{T, i+1}$ to fine-tune the generators and discriminators is given as:

$$\begin{aligned} \mathcal{L}_{T, i+1}(F_i, H_i, D_{\mathcal{X}, i}, D_{\mathcal{Y}, i}, \mathcal{X}_T, \mathcal{Y}_T) \\ = \mathcal{L}_{\text{adv}}(F_i, D_{\mathcal{Y}, i}, \mathcal{X}_T, \mathcal{Y}_T) + \mathcal{L}_{\text{adv}}(H_i, D_{\mathcal{X}, i}, \mathcal{Y}_T, \mathcal{X}_T) \\ + \lambda_{\text{cyc}} \mathcal{L}_{\text{cyc}}(F_i, H_i, \mathcal{X}_T, \mathcal{Y}_T) + \lambda_{\text{idt}} \mathcal{L}_{\text{idt}}(F_i, H_i, \mathcal{X}_T, \mathcal{Y}_T), \end{aligned} \quad (9)$$

where $D_{\mathcal{Y}, i}$, F_i , $D_{\mathcal{X}, i}$ and H_i are the state of the $D_{\mathcal{Y}}$, F , $D_{\mathcal{X}}$ and H at meta-training iteration i . We estimate the expectation terms in $\mathcal{L}_{T, i+1}$ also using the samples in support set S_T .

For meta-optimization, we minimize the expected loss on query set Q_T with updated discriminators and generators across the task T to train the initial parameters of D and G . Our MT-GAN model can be trained as follows:

$$\theta_d = \theta_d + \beta \nabla_{\theta_d} \mathcal{L}_{T, I+1}^q, \quad (10)$$

$$\theta_g = \theta_g - \beta \nabla_{\theta_g} \mathcal{L}_{T, I+1}^q, \quad (11)$$

where β is the learning rate for the meta-optimization, and I is the overall iteration number of meta-training. The overall meta-optimization objective $\mathcal{L}_{T, I+1}^q$ for the meta-generator and meta-discriminator is given as:

$$\begin{aligned} \mathcal{L}_{T, I+1}^q(F_I, H_I, D_{\mathcal{X}, I}, D_{\mathcal{Y}, I}, \mathcal{X}_T, \mathcal{Y}_T) \\ = \mathcal{L}_{\text{adv}}(F_I, D_{\mathcal{Y}, I}, \mathcal{X}_T, \mathcal{Y}_T) + \mathcal{L}_{\text{adv}}(H_I, D_{\mathcal{X}, I}, \mathcal{Y}_T, \mathcal{X}_T) \\ + \lambda_{\text{cyc}} \mathcal{L}_{\text{cyc}}(F_I, H_I, \mathcal{X}_T, \mathcal{Y}_T) + \lambda_{\text{idt}} \mathcal{L}_{\text{idt}}(F_I, H_I, \mathcal{X}_T, \mathcal{Y}_T), \end{aligned} \quad (12)$$

where we estimate the expectation terms in $\mathcal{L}_{T, I+1}^q$ using the samples in query set Q_T .

At inference time, for an unseen translation task $T_{N+1} = (S_{T_{N+1}}, Q_{T_{N+1}})$, where $S_{T_{N+1}} = \{\{x_k\}_{k=1}^K \in \mathcal{X}_{T_{N+1}}, \{y_k\}_{k=1}^K \in \mathcal{Y}_{T_{N+1}}\}$ and $Q_{T_{N+1}} = \{\{x_l\}_{l=1}^L \in \mathcal{X}_{T_{N+1}}, \{y_l\}_{l=1}^L \in \mathcal{Y}_{T_{N+1}}\}$, we iteratively fine-tune the obtained G and D with meta-training steps from Eqn.(1) to Eqn.(9) to obtain an F that transfers $\{x_k\}_{k=1}^K$ to $\mathcal{Y}_{T_{N+1}}$ domain and an H that transfers $\{y_k\}_{k=1}^K$ to $\mathcal{X}_{T_{N+1}}$ domain.

4 Experiments

4.1 Experimental Setup

We extensively evaluate the effectiveness and generalizing ability of the proposed MT-GAN algorithm for UDT problem on two kinds of translation task distributions. The first one (denoted as $P_1(T)$) contains 10 diverse translation tasks collected by (Zhu et al. 2017): labels \leftrightarrow photos, horses \leftrightarrow zebras, summer \leftrightarrow winter, apple \leftrightarrow orange, monet \leftrightarrow photo, cezanne \leftrightarrow photo, ukiyoe \leftrightarrow photo, vangogh \leftrightarrow photo, photos \leftrightarrow maps and labels \leftrightarrow facades. In addition, the Facescrub dataset (Ng and Winkler 2014), which comprises 531 different celebrities, is utilized as another collection of domain translation tasks (denoted as $P_2(T)$) that are less diverse, in which different identities are viewed as different domains. Then we can sample arbitrary two identities to form a two-domain translation task that aims to transfer the identity of face images while preserving original face orientation and expression.

In our experiments, for both $P_1(T)$ and $P_2(T)$, we simulate the meta domain translation scenarios by randomly select $N = 9$ tasks as a training dataset and select the other 1 task as the testing dataset/task. This procedure could be seen as task-level 10 fold cross-validation. We establish 10 training datasets and 10 corresponding testing datasets for both $P_1(T)$ and $P_2(T)$. For each training dataset, we randomly select overall 2000 meta batches from the 9 tasks for model training. We set the meta-batch size J to 2 to fit the memory limit of the GPU. Following the common settings of few-shot learning, we mainly focus on 5-shot domain translation and 10-shot domain translation in all experiments. Moreover, we set the query set's size L to 10. For the testing task, we randomly select 5 meta batches from the 1 task for model testing.

For each meta-training period, we use stochastic gradient descent (SGD) with learning rate $\alpha = 0.0001$ to fine-tune the generators and discriminators. At meta-optimization time, we use the Adam optimizer (Kingma and Ba 2014) with learning rate $\beta = 0.0002$ to update both meta-generator and meta-discriminator. For model fine-tuning on the testing tasks, we also use the Adam optimizer with learning rate $\beta = 0.0002$ to fine-tune both meta-generator and meta-discriminator. The overall iteration number of meta-training T is set to 100. We set the loss function balance parameters λ_{cyc} and λ_{idt} to be 10 and 5.

For each K -shot domain translation task, we fine-tune the trained MT-GAN on each meta batch of the testing dataset, and report the average score and its standard deviation. The Frechet Inception Distance (FID) (Heusel et al. 2017) that measures similarity between generated image dataset and real image dataset is used to evaluate translation results' quality.



Figure 1: Translation results of various testing tasks using CycleGAN, StarGAN and our MT-GAN. Left part: 5-shot domain translation. Right part: 10-shot domain translation. From left to right, the columns represent inputs from the source domain, CycleGAN’s results, StarGAN’s results, our results and examples in the target domain respectively.

The lower the FID is, the better the translation results are. In addition, we perform face classification experiments on face identity translation tasks. We re-trained VGG-16 network (Simonyan and Zisserman 2014) on Facescrub, and compute the top-1 and top-5 classification accuracy rates of the translation results.

4.2 Model configuration

We follow (Zhu et al. 2017) to configure the models. The meta-generator G network consists of two convolution layers with stride 2 and kernel size 3×3 , six residual blocks (He et al. 2016b) with kernel size 3×3 and two transposed convolution layers with stride 0.5 and kernel size 3×3 . For

Table 1: Average FID scores ($\times 10$) of various 10-shot testing tasks. \leftarrow represents the reverse translation direction, such as labels \leftarrow photos, and \rightarrow represents the forward translation direction, such as labels \rightarrow photos. For each translation direction, the best FID scores are in bold.

	CycleGAN (Zhu et al. 2017)		StarGAN (Choi et al. 2018)		Ours	
	\leftarrow	\rightarrow	\leftarrow	\rightarrow	\leftarrow	\rightarrow
labels \leftrightarrow photos	15.10 ± 1.72	28.69 ± 1.54	22.56 ± 1.55	39.27 ± 1.48	12.16 ± 1.19	26.08 ± 1.25
horses \leftrightarrow zebras	31.96 ± 1.63	31.45 ± 2.29	34.94 ± 1.72	35.83 ± 1.81	31.80 ± 1.28	30.68 ± 2.09
summer \leftrightarrow winter	23.43 ± 2.12	19.24 ± 2.02	32.18 ± 1.98	34.95 ± 1.89	21.50 ± 1.71	17.95 ± 1.93
apple \leftrightarrow orange	33.16 ± 1.83	34.30 ± 1.51	38.93 ± 1.69	43.58 ± 1.87	31.34 ± 1.57	31.93 ± 1.92
monet \leftrightarrow photo	18.97 ± 1.75	19.93 ± 2.21	33.38 ± 2.27	35.95 ± 2.18	18.42 ± 2.11	19.73 ± 2.08
cezanne \leftrightarrow photo	22.85 ± 1.71	23.24 ± 1.89	29.23 ± 1.85	31.37 ± 1.97	22.00 ± 1.66	22.57 ± 1.78
ukiyo \leftrightarrow photo	22.90 ± 2.05	21.16 ± 2.14	24.87 ± 2.25	28.99 ± 2.23	21.88 ± 2.18	20.86 ± 1.96
vangogh \leftrightarrow photo	22.65 ± 2.29	19.26 ± 2.18	34.91 ± 2.26	36.68 ± 2.32	21.57 ± 2.14	18.42 ± 2.05
photos \leftrightarrow maps	25.57 ± 1.96	27.17 ± 1.72	35.61 ± 1.88	31.36 ± 1.95	19.64 ± 1.64	23.44 ± 1.72
labels \leftrightarrow facades	23.06 ± 1.86	26.35 ± 1.73	26.55 ± 2.03	32.80 ± 2.11	21.72 ± 2.09	25.34 ± 1.96

Table 2: Average classification accuracy of 5-shot and 10-shot face identity translation tasks. The best top-1 and top-5 classification accuracy are in bold.

		CycleGAN (Zhu et al. 2017)	StarGAN (Choi et al. 2018)	Ours
5-shot	Top-1	$11.21 \pm 0.91\%$	$4.36 \pm 0.89\%$	$13.05 \pm 1.06\%$
	Top-5	$37.12 \pm 1.63\%$	$15.35 \pm 1.11\%$	$40.02 \pm 1.22\%$
10-shot	Top-1	$19.04 \pm 1.21\%$	$10.38 \pm 0.78\%$	$21.12 \pm 1.03\%$
	Top-5	$45.56 \pm 1.56\%$	$37.34 \pm 1.43\%$	$48.27 \pm 1.54\%$

meta-discriminator D , we use PatchGANs (Isola et al. 2017) that consists of five convolution layers with stride 2 and kernel size 4×4 . For both G and D , we use batch normalization (Ioffe and Szegedy 2015) among network layers.

4.3 Results

Qualitative and Quantitative Evaluation We compare our method with two baseline domain translation models, i.e., CycleGAN (Zhu et al. 2017) and StarGAN (Choi et al. 2018). We retrain both CycleGAN and StarGAN on each meta-batch in testing dataset of a given K -shot domain translation task, and report their average performance with the retrained models. We show qualitative comparison results of the testing tasks in Figure 1. We observe that StarGAN typically produces quite blurry and noisy outputs, and obviously suffers from the limited training samples. CycleGAN maintains the main structure of the source inputs in most cases and transfers some domain-specific features of the target domains in the translation results. However, CycleGAN still fails to locate the accurate regions that domain-specific features should be transferred in, and produces unnatural images. For example, the translated apple by CycleGAN in 10-shot orange \rightarrow apple is surrounded by inaccurate apple features, and the translated map by CycleGAN in 10-shot photos \rightarrow maps mistakenly transfers the land and houses to water label. On the contrary, most of our results well preserve the domain-invariant features (Lin et al. 2018; Huang et al. 2018) and accurately transfer the domain-specific features (Lin et al. 2018; Huang et al. 2018) in the translation results. It should be noticed that, even with limited unpaired training samples, our model is still able to detect semantic regions of source inputs. For instance, our model successfully detects the water area and land in 10-shot domain translation of the Figure 1 (f) and (g).

For the quantitative evaluation, we present the FID score results of various testing tasks in Table 1. For face identity translation, we report the top-1 and top-5 face recognition accuracy of generated images from CycleGAN, StarGAN and our model in Table 2. We observe that the quantitative results are quite related to the qualitative results in Figure 1, in which our model consistently outperforms CycleGAN and StarGAN. We also find that improvement brought by our model on some natural image generation tasks, such as four painting \leftrightarrow photo tasks, is less significant than other tasks, such as photos \leftrightarrow maps and labels \leftrightarrow photos. Such result is not surprising because only limited samples are hard to include all patterns for natural image generation, while the patterns of photos \leftrightarrow maps or labels \leftrightarrow photos are more simple and regularized.

Comparing the performance of MT-GAN on 5-shot and 10-shot domain translation tasks, we can see that the proposed meta-learning approach is quite robust to the drop in the amount of training samples. With only 5 training samples, MT-GAN still successfully transfers source inputs to target domains in most cases. With the increase of training samples, MT-GAN steadily improves performance on 10-shot domain translation tasks compared with CycleGAN and StarGAN.

Convergence Rate Our meta-learning based approach has demonstrated that, with several domain translation tasks, it can incorporate the prior learning experiences on these tasks and generalize to a new task with better performance than ordinary translation models in the above experiments. We show that meta-learning brings another benefit, i.e., faster convergence rate in a training process. We show the training curves of cycle-consistency loss with respect to training steps on the different testing tasks in Figure 2. We choose cycle-consistency loss to reflect the convergence rate because




Figure 2: The training curves of cycle-consistency loss with respect to training step on the different testing tasks.

a smaller cycle-consistency loss indicates that the two translation models well relate two domains. Comparing training curves of CycleGAN and our model, we observe that our model rapidly minimizes the cycle-consistency loss in the first several steps. In addition, our model achieves lower cycle-consistency loss than CycleGAN after numerous iteration steps in most cases. When training abnormality of GAN occurs, we see that our model can recover to original training states more quickly than CycleGAN. These results demonstrate that our model indeed learns adaptation strategies from previous translation tasks, and helps to converge more quickly in current tasks.

Saturation Study In order to investigate when our model will saturate with the increasing shots, we compare our model with CycleGAN which has shown superior performance in many-shot domain translation. We conduct the experiment on the face identity translation task. As shown in Table 3, we find that our model consistently outperforms CycleGAN in all shots, and performance of two models becomes closer with more shots. This is because with more training samples, a model is able to learn from scratch without prior experiences.

5 Conclusions

In this work, we manage the unsupervised domain translation (UDT) problem from a meta-learning perspective which

Table 3: Average top-1 classification accuracy of face identity translation tasks with different shots.

	10-shot	15-shot	20-shot	25-shot
CycleGAN	19.04%	25.95%	31.84%	35.62%
Ours	21.12%	27.63%	32.96%	36.03%
Improvement	2.08%	1.68%	1.12%	0.41%

aims to effectively incorporate prior domain translation experiences. Accordingly, we propose a model called MT-GAN to find the initialization of a meta-generator and a meta-discriminator that can be used for initialization of any translation task. We jointly train two meta-learners in an adversarial and dual form. We demonstrate our model on ten diverse domain translation tasks and face identity translation tasks. Both qualitative and quantitative results show that the meta-learning based approach significantly outperforms ordinary translation models. In addition, we show that our model can achieve faster convergence rate than CycleGAN, which further demonstrates MT-GAN indeed learns adaptation strategies from previous learning experiences.

For future works, it would be interesting to extend the training paradigm of MT-GAN to other image generation or domain transfer learning tasks. In addition, how to learn the adaptation strategies from many-shot domain translation tasks will be worthy to explore.

Acknowledgments

This work was supported in part by NSFC under Grant 61571413, 61632001.

References

- Andrychowicz, M.; Denil, M.; Gomez, S.; Hoffman, M. W.; Pfau, D.; Schaul, T.; Shillingford, B.; and De Freitas, N. 2016. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, 3981–3989.
- Benaïm, S., and Wolf, L. 2018. One-shot unsupervised cross domain translation. In *Advances in Neural Information Processing Systems*, 2104–2114.
- Bengio, Y.; Bengio, S.; and Cloutier, J. 1990. *Learning a synaptic learning rule*. Université de Montréal, Département d’informatique et de recherche.
- Bengio, S.; Bengio, Y.; Cloutier, J.; and Gecsei, J. 1992. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*, 6–8. Univ. of Texas.
- Choi, Y.; Choi, M.; Kim, M.; Ha, J.-W.; Kim, S.; and Choo, J. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- de Bézenac, E.; Ayed, I.; and Gallinari, P. 2019. Optimal unsupervised domain translation. arXiv:1906.01292.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1126–1135. JMLR. org.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2672–2680.
- He, D.; Xia, Y.; Qin, T.; Wang, L.; Yu, N.; Liu, T.; and Ma, W.-Y. 2016a. Dual learning for machine translation. In *Advances in Neural Information Processing Systems*, 820–828.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016b. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*. Curran Associates, Inc. 6626–6637.
- Hsu, K.; Levine, S.; and Finn, C. 2018. Unsupervised learning via meta-learning. arXiv:1810.02334.
- Huang, X.; Liu, M.-Y.; Belongie, S.; and Kautz, J. 2018. Multimodal unsupervised image-to-image translation. In *ECCV*.
- Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv:1502.03167.
- Isola, P.; Zhu, J.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5967–5976.
- Kim, T.; Cha, M.; Kim, H.; Lee, J. K.; and Kim, J. 2017. Learning to discover cross-domain relations with generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, 1857–1865.
- Kingma, D., and Ba, J. 2014. Adam: A method for stochastic optimization. arXiv:1412.6980.
- Lin, J.; Xia, Y.; Qin, T.; Chen, Z.; and Liu, T.-Y. 2018. Conditional image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Liu, A. H.; Liu, Y.-C.; Yeh, Y.-Y.; and Wang, Y.-C. F. 2018. A unified feature disentangler for multi-domain image translation and manipulation. In *Advances in Neural Information Processing Systems*, 2595–2604.
- Metz, L.; Maheswaranathan, N.; Cheung, B.; and Sohl-Dickstein, J. 2018. Meta-learning update rules for unsupervised representation learning. arXiv:1804.00222.
- Mirza, M., and Osindero, S. 2014. Conditional generative adversarial nets. arXiv:1411.1784.
- Mishra, N.; Rohaninejad, M.; Chen, X.; and Abbeel, P. 2017. A simple neural attentive meta-learner. arXiv:1707.03141.
- Ng, H.-W., and Winkler, S. 2014. A data-driven approach to cleaning large face datasets. In *2014 IEEE International Conference on Image Processing (ICIP)*, 343–347.
- Pumarola, A.; Agudo, A.; Martinez, A.; Sanfeliu, A.; and Moreno-Noguer, F. 2018. Ganimation: Anatomically-aware facial animation from a single image. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Radford, A.; Metz, L.; and Chintala, S. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv:1511.06434.
- Schmidhuber, J. 1987. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. Ph.D. Dissertation, Technische Universität München.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. arXiv:1409.1556.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, 4077–4087.
- Vinyals, O.; Blundell, C.; Lillicrap, T.; and Wierstra, D. 2016. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, 3630–3638.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8798–8807.
- Wu, J.; Zhang, C.; Xue, T.; Freeman, B.; and Tenenbaum, J. 2016. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in Neural Information Processing Systems*, 82–90.
- Yi, Z.; Zhang, H.; Tan, P.; and Gong, M. 2017. Dualgan: Unsupervised dual learning for image-to-image translation. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Yu, J.; Lin, Z.; Yang, J.; Shen, X.; Lu, X.; and Huang, T. S. 2018. Generative image inpainting with contextual attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5505–5514.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *The IEEE International Conference on Computer Vision (ICCV)*.