

# One-to-one Mapping for Unpaired Image-to-image Translation

Zengming Shen

Yifan Chen

Thomas S. Huang

University of Illinois at Urbana-Champaign

zshen5, yifanc3, t-huang1@illinois.edu

S.Kevin Zhou

Institute of Computing Technology, Chinese Academy of Sciences

zhoushaohua@ict.ac.cn

Bogdan Georgescu

Siemens Healthineers

bogdan.georgescu@siemens-healthineers.com

Xuqi Liu

Rutgers University

x1325@scarletmail.rutgers.edu

## Abstract

Recently image-to-image translation has attracted significant interests in the literature, starting from the successful use of the generative adversarial network (GAN), to the introduction of cyclic constraint, to extensions to multiple domains. However, in existing approaches, there is no guarantee that the mapping between two image domains is unique or one-to-one. Here we propose a self-inverse network learning approach for unpaired image-to-image translation. Building on top of CycleGAN, we learn a self-inverse function by simply augmenting the training samples by swapping inputs and outputs during training and with separated cycle consistency loss for each mapping direction. The outcome of such learning is a proven one-to-one mapping function. Our extensive experiments on a variety of datasets, including cross-modal medical image synthesis, object transfiguration, and semantic labeling, consistently demonstrate clear improvement over the CycleGAN method both qualitatively and quantitatively. Especially our proposed method reaches the state-of-the-art result on the cityscapes benchmark dataset for the label to photo unpaired directional image translation.

## 1. Introduction

Image-to-image translation (or cross-domain image synthesis) learns a mapping function from an input image to an output image or vice versa. It can be grouped into two categories: supervised [12] vs unsupervised (or unpaired) [39].

The task of learning mappings between two domains from unpaired data has attracted a lot of attention, especially in the form of unpaired image-to-image translation




Figure 1. A comparison of our one2one CycleGAN with the original CycleGAN [40] for the mapping between two domains X and Y. (a) Original CycleGAN model. It contains two separated mapping functions  $G : X \rightarrow Y$  and  $F : Y \rightarrow X$ . (b) Our One2One CycleGAN. We propose to realize one-to-one mapping by learning ONLY one self-inverse function  $G$  for the mapping between two domains bidirectionally. It contains only one mapping function  $G : X \leftrightarrow Y$ .

[40, 41, 15, 21]. Thanks to the pioneer work of GAN[10] and cycleGAN [39], recent works [32, 28, 38, 37, 6, 11, 18, 1, 35, 31] have shown promising result for unpaired image-to-image translation. This task is very important because paired data are not available in many cases and the paired information are difficult or time-consuming to get. For example, in the medical field of cross domain medical image segmentation[37, 38]: with the brain CT image semantic label and without the brain MRI semantic label, the goal is to generate semantic label for the brain MRI image. Amazing works [37, 38] like this cross domain image segmentation task in the medical image application could be further improved if unpaired image translation can be unique and more accurate. In many cases, the information source like patient is unique. For example, there is only a

brain MRI image for a patient, but there should be a unique CT brain image for the same patient. This uniqueness requirement can be called one-to-one mapping of the brain CT and the brain MRI image from the same patient. If the unpaired image-to-image translation can achieve this one-to-one mapping, the cross-domain medical image segmentation performance can be further improved. However, existing method can not meet this requirement.

As mentioned above, the major limitation of existing methods for unpaired image-to-image translation like CycleGAN is that they can not realize one-to-one mapping which is necessary in many cases like the information source of the unpaired image is unique. Without the pairing information, CycleGAN using the distribution constraint allows many-to-many mappings. To reduce the space of possible mappings and improve in finding a more unique mapping, their models add an essential cycle-consistency constraint. The cycle-consistency constraint enforces a stronger connection cross domains by requiring the input image and the output image to be close. The output image is generated by first mapping from source domain to target domain then mapping back to the source domain. But this only reduces the many-to-many mapping to many-to-one mapping or one-to-many mapping cross domains, which will be illustrate in detail in section 3.2 and 3.3.

With the success of image generation[10, 29] model Generation Adversarial Networks(GANs) and unsupervised mapping methods like CycleGAN [39], motivated by the recent works [30, 13] of exploring invertibility of convolutional neural networks (CNNs), we propose to learn a one-to-one mapping between domains from unpaired data to compensate the limitation of the existing methods such as CycleGAN.

Specifically, we enforce the generator of the CycleGAN as a self-inverse function to realize a one-to-one mapping. So we call our proposed method One2one CycleGAN. When a function  $G$  is self-inverse, illustrated as

$$G = G^{-1}, \quad (1)$$

it guarantees a one-to-one mapping. We use the CycleGAN [39] as the baseline framework for image-to-image translation. To impose the self-inverse property, we implement a *simple idea* of augmenting the training samples by switching inputs and outputs during training. However, as we will demonstrate empirically, *this seemingly simple idea makes a genuinely big difference!*

The distinct feature of our self-inverse network is that it learns one network to perform both forward ( $X \rightarrow Y$ : from X to Y) and backward ( $Y \rightarrow X$ : from Y to X) translation tasks. It contrasts with the state-of-the-art approaches which typically learn two separate networks, one for forwarding translation and the other for backward translation. As a result, it enjoys several benefits. First, it halves the

necessary parameters, assuming that the self-inverse network and the two separate networks share the same network architecture. Second, it automatically doubles the sample size, a great feature for any data-driven models, thus becoming less likely to over-fit the model.

One key question arises: Is it feasible to learn such a self-inverse network for image-to-image translation? We can not theoretically prove this existence; however, we experimentally demonstrate so. Intuitively, such an existence is related to the redundancy in the expressive power of the deep neural network. Even given a fixed network architecture, the function space for a network that translates an image from  $A$  to  $B$  is large enough, that is, there are many neural networks with different parameters capable of doing the same translation job. The same holds for the inversion network. Therefore, the overlap between these two spaces, in which the self-inverse network resides, does exist.

Our contribution are as follows: (i) We introduce the One2one CycleGAN model for learning one-to-one mappings cross domains in an unsupervised way. (ii) We show that our model can learn mappings that generate a more accurate output for each input. (iii) We evaluate our method in extensive experiments on a variety of datasets, including cross-modal medical image synthesis, object transfiguration, and semantic labeling, consistently demonstrate clear improvement over the CycleGAN method both qualitatively and quantitatively. Especially our proposed method reaches the state-of-the-art result on the cityscapes benchmark dataset for the label to photo unpaired directional image translation.

## 2. Literature Review

Iosla et al. [12] presented the seminar work of image-to-image translation that offered a general-purpose solution, and Goodfellow et al. proposed to use the generative adversarial network (GAN) [10] for the first time in the literature. While paired data are assumed in [12], later Zhu et al. [39] proposed the CycleGAN approach for addressing the unpaired setting using the so-called cyclic constraints. There are many recent advances that use guidance information [32, 28], impose different constraints [9, 22, 36], or deal with multiple domains[41, 6, 11, 18], etc. In this paper, we study unpaired image-to-image translation.

In addition to using the GAN that essentially enforces similarity in image distribution, other guidance information is used such as landmark points [32], contours [8], sketches [23], anatomical information [28] etc. In addition to cyclic constraint [39], other constraints like ternary discriminative function [9], optimal transport function [22], smoothness over the sample graph [36] are used as well.

Also, extensions were proposed to deal with video inputs [33, 3], to synthesize images in high resolution [34], to seek for diversity [25] and to handle more than two im-




Figure 2. (a) The mapping routes of CycleGAN. The limitations of the CycleGAN model is that it allows biased and non-unique unpaired image translation. For the mapping route  $x \rightarrow x'$ , the mapping  $G: x \rightarrow y'$  is a one to many mapping which result that  $x$  can be mapped to infinity possible  $y'$ . Let's denote the unique target is  $y'_t$  and the actually mapped result is  $y'_k$ ; The mapping  $F: y'_k \rightarrow x'$  is a many to one mapping. As a result, there is allowable bias between the target  $y'_t$  and the prediction  $y'_k$ . Similarly, for the mapping route  $y \rightarrow y'$ , the mapping  $F: y \rightarrow x'$  is a one to many mapping which result that  $y$  can be mapped to infinity possible  $x'$ . Let's denote the unique target is  $x'_t$  and the actually mapped result is  $x'_k$ ; The mapping  $G: x'_k \rightarrow y'$  is a many to one mapping. As a result, there is allowable bias between the target  $x'_t$  and the prediction  $x'_k$ . (b) The mapping routes of one2one CycleGAN. The motivation of one2one CycleGAN is to realize unique and accurate unpaired image translation. The mapping function  $G$  is self inverse function with the one-to-one mapping property. For the mapping route  $x \rightarrow x'$ , the mapping  $G: x \rightarrow y'$  is a one to one mapping which result that  $x$  is only mapped to the unique target is  $y'_t$ . The mapping  $F: y'_t \rightarrow x'$  is also a one to one mapping. As a result, there is no bias between the target and the prediction. Similarly, for the mapping route  $y \rightarrow y'$ , the mapping  $F: y \rightarrow x'$  is a one to one mapping which result that  $y$  can only be mapped to the unique target is  $x'_t$ . The mapping  $G: x'_t \rightarrow y'$  is also a one to one mapping. As a result, there is no bias between the target and the prediction.

age domains [41, 6, 11, 18]. Furthermore, there are methods that leverage attention mechanism [24, 5, 26] and mask guidance [20]. Finally, disentangling is a new emerging direction [11, 18].

In terms of works about inverse problem with neural networks, [13] makes the CNN architecture invertible by providing an explicit inverse. Ardizzone et.al [2] prove the invertibility theoretically. More specifically, Kingma [17] shows the benefit of a invertible  $1 \times 1$  convolution.

Different from a one-to-one mapping function are one-to-many, many-to-one, and many-to-many [1]<sup>1</sup> mapping functions. In [12], the well-studied scenarios of labels-to-scenes, edge-to-photo are more likely one-to-many map-

<sup>1</sup>It is worth noting that recently there are quite some works focusing on addressing image-to-image translation among many domains, also the so-called one-to-many .

ping as it is possible that multiple photos (scenes) have the same edge (label) information. The colorization example is also one-to-many. From an information theory perspective, the entropy of the edge map (label) is low while that of the photo is high. When an image translation goes from an information-gaining direction, that is, from low-entropy to high-entropy, its mapping leans towards one-to-many. Similarly, if it goes from an information-losing direction, then its mapping leans toward many-to-one. If the information level of both domains is close (or information-similar), then the mapping is close to one-to-one. In [12], the examples of Monet-to-photo, summer-to-winter are closer to one-to-one mapping as the underlying contents of both images before and after translation are regarded the same but the styles are different, which does not change the image entropy significantly. For image to image translation, many works has been done to diversify the output [1, 21, 18, 11, 41, 19], while not too many work has been done to make the output unique [31]. Our work goes to the latter direction.

Although there are so many research works on image-to-image translation, the perspective of learning a one-to-one mapping network has not been fully investigated, with the exception of [22]. In [22], Lu et al. show that CycleGAN can not theoretically guarantee the one-to-one mapping property and propose to use an optimal transport mechanism to mitigate this issue. However, like GAN, the optimal transport method also measures the similarity in image distribution; hence the one-to-one issue is not fully resolved. By contrast, our self-inverse learning comes with a guarantee that the learned network realizes a one-to-one mapping.

### 3. Unsupervised Learning of One-to-one Mappings between Domains

#### 3.1. Problem setting

For any two domains  $X$  and  $Y$  with only unpaired elements available, we assume there exists a mapping, potentially one-to-one mapping, between the elements of each domain. The goal to make sure there is a unique target element in the target domain to match a element in the source domain. The objective is to recover this mapping. Since there are only unpaired samples available, this goal is realized by matching the distributions  $p_d(x)$  and  $p_d(y)$  of each domain. This can be treated as a conditional generating task. The true conditionals  $p(x|y)$  and  $p(y|x)$  are estimated from the true marginals. To be able to uncover this mapping, the elements in both domain  $X$  and domain  $Y$  are highly dependent.

#### 3.2. CycleGAN model

As shown in Figure 1, the CycleGAN model [39] solves this problem by estimating these two conditionals with two

separated mappings functions  $G : X \rightarrow Y$  and  $F : Y \rightarrow X$ . Both of the mapping functions are parameterized with identical neural networks and constrained by the followings:

- **Distribution matching:** The distribution of the each mapping output should match the distribution of the target domain. This constrain allows many-to-many mappings between the source domain  $X$  and the target domain  $Y$  and vice versa.
- **Cycle-consistency:** Each element is mapped from the source domain to the target domain, then mapped back to source domain. The output should be close to the input element. This constrains one-to-many mapping from source domain and many-to-one mapping from target domain to source domain.

### 3.3. Limitations of CycleGAN for one-to-one mapping

The main weakness of the CycleGAN model is that it can not realize one-to-one mapping for accurate and unique unpaired image translation. Based on the above constrains and the illustration in Figure 2, the CycleGAN model can not satisfy our problem's goal. Next, we show how to modify CycleGAN to meet the goal of our problem.

The distribution matching is implemented by GAN[10]. The two mapping functions  $G$  and  $F$  implemented by neural networks are trained to fool the Discriminator  $D_Y$  and  $D_X$  respectively. The adversarial loss[10] for mapping function  $G$  is

$$\mathcal{L}_{GAN}(G, D_X, X, Y) = \mathbb{E}_{x \sim p_{data}(x)} [\log D_Y(x)] + \mathbb{E}_{y \sim p_{data}(y)} [\log(1 - D_x(G(y)))]. \quad (2)$$

The cycle consistency loss is:

$$\mathcal{L}_{cyc}(G, F) = \mathbb{E}_{x \sim p_{data}(x)} [\|F(G(x)) - x\|_1] + \mathbb{E}_{y \sim p_{data}(y)} [\|G(F(y)) - y\|_1]. \quad (3)$$

The final objective for the mapping function  $G$  and  $F$  is

$$\mathcal{L}(G, F, D_X, D_Y) = \mathcal{L}_{GAN}(G, D_Y, X, Y) + \mathcal{L}_{GAN}(F, D_X, Y, X) + \lambda \mathcal{L}_{cyc}(G, F) \quad (4)$$

and we aim to solves

$$(G^*, F^*) = \arg \min_{G, F} \max_{D_X, D_Y} \mathcal{L}(G, F, D_X, D_Y). \quad (5)$$

## 4. Self-inverse Learning for Unpaired Image-to-image Translation

In the section, we first show the property that the self-inverse function guarantees one-to-one (one2one) mapping. Then we discuss how to train a self-inverse CycleGAN network for image-to-image translation

### 4.1. One-to-one property

In image-to-image translation, we define a forward function as  $Y = f_{X \rightarrow B}(X)$  that maps an image  $X$  on domain  $A$  to another image  $Y$  on domain  $B$  and, similarly, an inverse function as  $X = f_{B \rightarrow A}^{-1}(Y)$ . When there is no confusion, we will skip the subscript (e.g.,  $A \rightarrow B$ ).

**Property:** If a function  $Y = f(X)$  is self-inverse, that is  $f = f^{-1}$ , then the function  $f$  defines a one-to-one mapping, that is,  $Y_1 = Y_2$  if and only if  $X_1 = X_2$ .

**Proof:**

$[\Rightarrow]$  If  $X_1 = X_2$ , then  $Y_1 = f(X_1) = f(X_2) = Y_2$ .

$[\Leftarrow]$  If  $Y_1 = Y_2$ , then  $X_1 = f^{-1}(Y_1) = f^{-1}(Y_2) = X_2$  as long as the inverse function exists, which is the case for a self-inverse function as  $f^{-1} = f$ . #

### 4.2. One-to-one benefits

There are several advantages in learning a self-inverse network to have the one-to-one mapping property.

(1) From the perspective of the application, only one self-inverse function can model both tasks  $A$  and  $B$  and it is a novel way for multi-task learning. As shown in Figure 1, the self-inverse network generates an output given an input, and vice versa, with only one CNN and without knowing the mapping direction. It is capable of doing both tasks within the same network, simultaneously. In comparison to separately assigning two CNNs for tasks  $A$  and  $B$ , the self-inverse network halves the necessary parameters, assuming that the self-inverse network and the two CNNs share the same network architecture as shown in Figure 1.

(2) It automatically doubles the sample size, an important feature for any data-driven models, thus it is less likely to over-fit the model. The self-inverse function  $f$  has the co-domain  $Z = X \cup Y$ . If the sample size of either domain  $X$  or  $Y$  is  $N$ , then the sample size for domain  $Z$  is  $2N$ . As a result, the sample size for both tasks  $A$  and  $B$  are doubled, becoming a novel method for data augmentation to mitigate the over-fitting problem.

(3) As shown in Figure 2, In the unpaired image-to-image translation setting, the goal is to minimize the distribution gap between the two domains. The state-of-art methods can realize this but can not guarantee an ordered mapping or bijection between the two domains. This results in variations for the generated images.

(4) The one-to-one mapping is a strict constraint. Therefore, forcing a CNN model as a self-inverse function can shrink the target function space.

### 4.3. One-to-one CycleGAN

We are inspired by the basic formulation of CycleGAN [39]. In CycleGAN, there are two generators  $Y = F(X)$  and  $X = G(Y)$ , two discriminators  $D_x$  and  $D_y$ , and one joint object function. In our one2one CycleGAN, we



have one shared generator  $G$  and still two discriminators  $D_x$  and  $D_y$ . Instead of having a joint objective for the dual-mappings, our proposed method has two separate objective functions, one for each of two mapping directions.

#### 4.3.1 Separated loss functions

Compared to CycleGAN that uses a joint loss for both image transfer directions, our method have two separate losses, one for each image transfer direction. For the mapping function  $G : X \rightarrow Y$  and its discriminator  $D_Y$ , the adversarial loss is

$$\mathcal{L}_{GAN}(G, D_Y, X, Y) = \mathbb{E}_{y \sim p_{data}(y)} [\log D_Y(y)] + \mathbb{E}_{x \sim p_{data}(x)} [\log(1 - D_Y(G(x)))]. \quad (6)$$

The cycle consistency loss is

$$\mathcal{L}_{cyc}^x(G) = \mathbb{E}_{x \sim p_{data}(x)} [\|G(G(x)) - x\|_1]. \quad (7)$$

For the mapping function  $G : Y \rightarrow X$  and its discriminator  $D_X$ , the adversarial loss is:

$$\mathcal{L}_{GAN}(G, D_X, X, Y) = \mathbb{E}_{x \sim p_{data}(x)} [\log D_X(x)] + \mathbb{E}_{y \sim p_{data}(y)} [\log(1 - D_X(G(y)))]. \quad (8)$$

The cycle consistency loss is:

$$\mathcal{L}_{cyc}^y(G) = \mathbb{E}_{y \sim p_{data}(y)} [\|G(G(y)) - y\|_1]. \quad (9)$$

So, the final objective for the mapping function  $X \rightarrow Y$  is

$$\mathcal{L}(G, D_Y) = \mathcal{L}_{GAN}(G, D_Y, X, Y) + \lambda_x \mathcal{L}_{cyc}^x(G), \quad (10)$$

and the minimax optimization solves

$$(G^*, D_Y^*) = \arg \min_{\mathbf{G}} \max_{\mathbf{D}_Y} \mathcal{L}(G, D_Y). \quad (11)$$

Similarly, the final objective for the mapping function  $Y \rightarrow X$  is

$$\mathcal{L}(G, D_X) = \mathcal{L}_{GAN}(G, D_X, X, Y) + \lambda_y \mathcal{L}_{cyc}^y(G), \quad (12)$$

and the minimax optimization solves

$$(G^*, D_X^*) = \arg \min_{\mathbf{G}} \max_{\mathbf{D}_X} \mathcal{L}(G, D_X). \quad (13)$$

#### 4.4. Self-inverse implementation

We apply the proposed method based on the framework of CycleGAN [39]. To have a fair comparison with CycleGAN, we adopt the architecture of (Johnson et al., 2016) as the generator and the PatchGAN [12] as the discriminator. The log likelihood objective in the original GAN is replaced with a least-squared loss [14] for more stable training. We

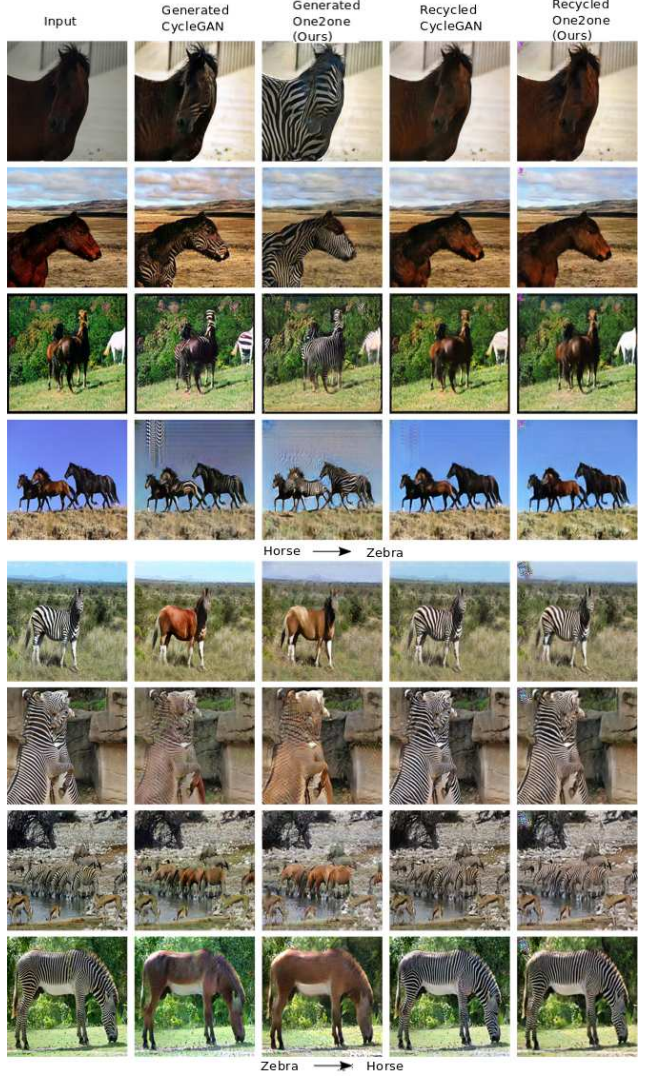


Figure 3. Visual comparison for horse $\leftrightarrow$ zebra.

resize the input images to  $256 \times 256$ . The loss weights are set as  $\lambda_x = \lambda_y = 10$ . Following CycleGAN, we adopt the Adam optimizer [16] with a learning rate of 0.0002. Similarly, we use a pool size of 50. The learning rate is fixed for the first 100 epochs and linearly decayed to zero over the next 100 epochs on Yosemite and apple2orange datasets. The learning rate is fixed for the first 4 epochs and linearly decayed to zero over the next 3 epochs on the BRATS dataset. The learning rate is fixed for the first 90 epochs and linearly decayed to zero over the next 30 epochs on the Cityscapes dataset.

#### 4.5. Training details and optimization

In our experiments, we use a batch size of 1. At each iteration, we randomly sample a batch of pair  $(x_i, y_i)$ , where

samples  $\{x_i\}_{i=1}^N \in X$  and  $\{y_i\}_{i=1}^M \in Y$ . At any iteration  $j$ , we perform the following three steps:

- Firstly, we feed  $x_i$  as the input and  $y_i$  as the target, then forward  $G$  and back-propagate  $G$ ;
- Secondly, we feed  $y_i$  as the input and  $x_i$  as the target, then forward  $G$  and back-propagate  $G$ ;
- Finally, we back-propagate  $D_Y$  and  $D_X$  individually.

## 5. Experiments

In order to test the effect of the proposed method, we evaluate it on an array of applications: cross-modal medical image synthesis, object transfiguration, and style transfer. Also we compare against several unpaired image-to-image translation methods: CycleGAN [39], DiscoGAN [15], DistanceGAN [4], and UNIT [21]. We conduct a user study when the ground truth images are unknown and perform quantitative evaluation when the ground truth images are present.

### 5.1. Datasets and results

**Object transfiguration.** we test our method on the horse  $\leftrightarrow$  zebra task used in CycleGAN paper [39] with 2401 training images (939 horses and 1177 zebras) and 260 test images (120 horses and 140 zebras). This task has no ground truth for generated images and hence no quantitative evaluation is feasible. So we provide the qualitative results obtained in a user study. In the user study, we ask a user to rate his/her preferred image out of three randomly positioned images, one obtains from CycleGAN, one from DistanceGAN, and the other from one2one CycleGAN. Figure 4 shows examples of input and synthesized images and Table 1 summarize the use study results.

Figure 4 tells that one2one CycleGAN likely generates better quality images in an unsupervised fashion, especially in terms of the quality of zebra synthesis from the horse (refer to the first four rows). Our method generated more real and complete zebra content. From Table 1, it is clear that our one2one CycleGAN is the most favorable with a 75% (77%) preference percentage for the horse2zebra (zebra2horse) mapping direction. and DistanceGAN is the least favorable.

we test our method on the apple  $\leftrightarrow$  orange task [39] with 2014 training images (995 apples and 1019 orange) and 514 test images (248 apples and 266 oranges). This task has no ground truth for generated images and hence no quantitative evaluation is feasible. Figure 5 shows examples of input and synthesized images. There are failure cases in rows 1,2,4 from CycleGAN while our model generates normal images.

**Cross-modal medical image synthesis.** This task evaluates cross-modal medical image synthesis. The models are




Figure 4. Visual comparison for summer $\leftrightarrow$ winter on yosemite.

trained on the BRATS dataset [27] which contains paired MRI data to allow quantitative evaluation. It contains ample multi-institutional routine clinically-acquired pre-operative multi modal MRI scans of glioblastoma (GBM/HGG) and lower grade glioma (LGG) images. There are 285 3D volumes for training and 66 3D volume for the test. The  $T_1$  and  $T_2$  images are selected for our bi-directional image synthesis. All the 3D volumes are preprocessed to one channel image of size 256 x 256 x 1. We use the Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) to evaluate the quality of generated images.

As shown in Table 2, on the  $T_1 \rightarrow T_2$  image synthesis direction, our one2one model outperforms the CycleGAN model on PSNR by 6.0%. The qualitative result is shown in columns 3 and 4 in Figure 7. On the  $T_2 \rightarrow T_1$  image synthesis direction, our one2one model outperforms the CycleGAN model on PSNR by 5.0%. The qualitative result is shown in columns 7 and 8 in Figure 7.

**Semantic labeling.** We also test our method on the labels  $\leftrightarrow$  photos task using the Cityscapes dataset [7] under the unpaired setting as in the original CycleGAN paper. For quantitative evaluation, in line with previous work, for labels  $\rightarrow$  photos we adopt the “FCN score” [12], which eval-



	Label $\rightarrow$ Photo			Photo $\rightarrow$ Label		
Method	Pixel Acc. $\uparrow$	Class Acc. $\uparrow$	Class IoU $\uparrow$	Pixel Acc. $\uparrow$	Class Acc. $\uparrow$	Class IoU $\uparrow$
CycleGAN	52.7	15.2	11.0	<b>57.2</b>	<b>21.0</b>	<b>15.7</b>
DiscoGAN	45.0	11.1	7.0	45.2	10.9	6.3
DistanceGAN	48.5	10.9	7.3	20.5	8.2	3.4
UNIT	48.5	12.9	7.9	56.0	20.5	14.3
One2one CycleGAN (ours)	<b>58.2</b>	<b>18.9</b>	<b>14.3</b>	52.7	18.1	13.0

Table 1. Results of Photo  $\leftrightarrow$  Label translation on the Cityscapes dataset.

Direction	Metric	Cycle	Distance	One2one
horse2zebra	Prefer pct. $\uparrow$	25%	0	75%
zebra2horse	Prefer pct. $\uparrow$	23%	0	77%

Table 2. Results of user study on the horse to zebra dataset.



Figure 5. Visual comparison for apple $\leftrightarrow$ orange.

uates how interpretable the generated photos are according to a semantic segmentation algorithm. For photos  $\leftarrow$  labels, we use the standard segmentation metrics, including per-pixel accuracy, per-class accuracy, and mean class Intersection-Over-Union (Class IoU). The quantitative result is shown in Table 3. Our model reaches the state-of-

Direction	Method	PSNR $\uparrow$	SSIM $\uparrow$
T1 $\rightarrow$ T2	CycleGAN	20.79	0.85
T1 $\rightarrow$ T2	One2one CycleGAN	<b>22.03</b>	<b>0.86</b>
T2 $\rightarrow$ T1	CycleGAN	17.47	0.81
T2 $\rightarrow$ T1	One2one CycleGAN	<b>18.31</b>	<b>0.82</b>

Table 3. Evaluation of cross-modal medical image synthesis on the BRATS dataset.

Direction	Metric	Cycle	One2one
summer2winter	Prefer pct. $\uparrow$	34%	66%
winter2summer	Prefer pct. $\uparrow$	41%	59%

Table 4. Results of user study on the summer to winter Yosemite dataset.

the-art on the label  $\rightarrow$  photo direction image synthesis under this unpaired setting. The pixel accuracy outperforms the second best result by 10.4 %; The class accuracy outperforms the second best result by 24.3 %; The class IoU outperforms the second best result by 30.0 %. On the photo  $\rightarrow$  label direction, our model reaches comparable results.

The qualitative result is shown in Figure 6. Compared with CycleGAN which is the second best result in the label  $\rightarrow$  photo direction, our model has clearly better visual results. On the photo  $\rightarrow$  label direction, our model also have a comparable or better result.

**Style Transfer.** We also test our method on the summer  $\leftrightarrow$  winter style transfer task using the Yosemite dataset under the unpaired setting as in the original CycleGAN paper. As shown in Figure 4 for the qualitative result, our method has better visual result in both directions of style transfer. We also do a similar user study by providing the generated image from the test set by our model and the CycleGAN to users. The result is in Table 4. The user study results show that our model has a higher preference than CycleGAN.

## 6. Conclusions

We have presented an approach for enforcing the learning of a one-to-one mapping function for unpaired image-to-image translation. The proposed one-to-one CycleGAN consistently outperforms the baseline CycleGAN model




Figure 6. Visual comparison for photo $\leftrightarrow$ label on the Cityscapes.




Figure 7. Qualitative comparison for T1 $\leftrightarrow$ T2 on BRATS datasets.

and other state-of-the-art unsupervised approaches in terms of various qualitative and quantitative metrics.

## 7. Acknowledgment

The work was supported by grants from Siemens Healthineers (Siemens 082387).



## References

- [1] A. Almahairi, S. Rajeswar, A. Sordoni, P. Bachman, and A. Courville. Augmented cyclegan: Learning many-to-many mappings from unpaired data. *arXiv preprint arXiv:1802.10151*, 2018.
- [2] L. Ardizzone, J. Kruse, S. Wirkert, D. Rahner, E. W. Pellegrini, R. S. Klessen, L. Maier-Hein, C. Rother, and U. Köthe. Analyzing inverse problems with invertible neural networks. *arXiv preprint arXiv:1808.04730*, 2018.
- [3] A. Bansal, S. Ma, D. Ramanan, and Y. Sheikh. Recycle-gan: Unsupervised video retargeting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 119–135, 2018.
- [4] S. Benaïm and L. Wolf. One-sided unsupervised domain mapping. In *Advances in neural information processing systems*, pages 752–762, 2017.
- [5] X. Chen, C. Xu, X. Yang, and D. Tao. Attention-gan for object transfiguration in wild images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 164–180, 2018.
- [6] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8789–8797, 2018.
- [7] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [8] T. Dekel, C. Gan, D. Krishnan, C. Liu, and W. T. Freeman. Smart, sparse contours to represent and edit images. *arXiv preprint arXiv:1712.08232*, 2017.
- [9] Z. Gan, L. Chen, W. Wang, Y. Pu, Y. Zhang, H. Liu, C. Li, and L. Carin. Triangle generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 5247–5256, 2017.
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [11] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018.
- [12] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [13] J.-H. Jacobsen, A. Smeulders, and E. Oyallon. i-revnet: Deep invertible networks. *arXiv preprint arXiv:1802.07088*, 2018.
- [14] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [15] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim. Learning to discover cross-domain relations with generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1857–1865. JMLR.org, 2017.
- [16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [17] D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.
- [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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 35–51, 2018.
- [19] H.-Y. Lee, H.-Y. Tseng, Q. Mao, J.-B. Huang, Y.-D. Lu, M. Singh, and M.-H. Yang. Dri++: Diverse image-to-image translation via disentangled representations. *arXiv preprint arXiv:1905.01270*, 2019.
- [20] X. Liang, H. Zhang, L. Lin, and E. Xing. Generative semantic manipulation with mask-contrasting gan. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 558–573, 2018.
- [21] M.-Y. Liu, T. Breuel, and J. Kautz. Unsupervised image-to-image translation networks. In *Advances in Neural Information Processing Systems*, pages 700–708, 2017.
- [22] G. Lu, Z. Zhou, Y. Song, K. Ren, and Y. Yu. Guiding the one-to-one mapping in cyclegan via optimal transport. *arXiv preprint arXiv:1811.06284*, 2018.
- [23] Y. Lu, S. Wu, Y.-W. Tai, and C.-K. Tang. Image generation from sketch constraint using contextual gan. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 205–220, 2018.
- [24] S. Ma, J. Fu, C. Wen Chen, and T. Mei. Da-gan: Instance-level image translation by deep attention generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5657–5666, 2018.
- [25] Q. Mao, H.-Y. Lee, H.-Y. Tseng, S. Ma, and M.-H. Yang. Mode seeking generative adversarial networks for diverse image synthesis. *arXiv preprint arXiv:1903.05628*, 2019.
- [26] Y. A. Mejjati, C. Richardt, J. Tompkin, D. Cosker, and K. I. Kim. Unsupervised attention-guided image-to-image translation. In *Advances in Neural Information Processing Systems*, pages 3693–3703, 2018.
- [27] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2015.
- [28] A. Pumarola, A. Agudo, A. M. Martinez, A. Sanfeliu, and F. Moreno-Noguer. Ganimation: Anatomically-aware facial animation from a single image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 818–833, 2018.
- [29] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

- [30] Z. Shen, Y. Chen, S. K. Zhou, B. Georgescu, X. Liu, and T. S. Huang. Towards learning a self-inverse network for bidirectional image-to-image translation. *arXiv preprint arXiv:1909.04104*, 2019.
- [31] Z. Shen, M. Huang, J. Shi, X. Xue, and T. Huang. Towards instance-level image-to-image translation. *arXiv preprint arXiv:1905.01744*, 2019.
- [32] L. Song, Z. Lu, R. He, Z. Sun, and T. Tan. Geometry guided adversarial facial expression synthesis. In *2018 ACM Multimedia Conference on Multimedia Conference*, pages 627–635. ACM, 2018.
- [33] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, G. Liu, A. Tao, J. Kautz, and B. Catanzaro. Video-to-video synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [34] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8798–8807, 2018.
- [35] R. Zhang, T. Pfister, and J. Li. Harmonic unpaired image-to-image translation. *arXiv preprint arXiv:1902.09727*, 2019.
- [36] R. Zhang, T. Pfister, and J. Li. Harmonic unpaired image-to-image translation. *CoRR*, abs/1902.09727, 2019.
- [37] Y. Zhang, S. Miao, T. Mansi, and R. Liao. Task driven generative modeling for unsupervised domain adaptation: Application to x-ray image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 599–607. Springer, 2018.
- [38] Z. Zhang, L. Yang, and Y. Zheng. Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9242–9251, 2018.
- [39] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
- [40] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.
- [41] J.-Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, and E. Shechtman. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*, pages 465–476, 2017.