
Augmented CycleGAN: Learning Many-to-Many Mappings from Unpaired Data

Amjad Almahairi^{1†} Sai Rajeswar¹ Alessandro Sordoni² Philip Bachman² Aaron Courville^{1,3}

Abstract

Learning inter-domain mappings from unpaired data can improve performance in structured prediction tasks, such as image segmentation, by reducing the need for paired data. CycleGAN was recently proposed for this problem, but critically assumes the underlying inter-domain mapping is approximately deterministic and one-to-one. This assumption renders the model ineffective for tasks requiring flexible, many-to-many mappings. We propose a new model, called Augmented CycleGAN, which learns many-to-many mappings between domains. We examine Augmented CycleGAN qualitatively and quantitatively on several image datasets.

1. Introduction

The problem of learning mappings between domains from unpaired data has recently received increasing attention, especially in the context of image-to-image translation (Zhu et al., 2017a; Kim et al., 2017; Liu et al., 2017). This problem is important because, in some cases, paired information may be scarce or otherwise difficult to obtain. For example, consider tasks like face transfiguration (male to female), where obtaining explicit pairs would be difficult as it would require artistic authoring. An effective unsupervised model may help when learning from relatively few paired examples, as compared to training strictly from the paired examples. Intuitively, forcing inter-domain mappings to be (approximately) invertible by a model of limited capacity acts as a strong regularizer.

Motivated by the success of Generative Adversarial Networks (GANs) in image generation (Goodfellow et al., 2014; Radford et al., 2015), existing unsupervised mapping meth-




Figure 1: (a) Original CycleGAN model. (b) We propose to learn many-to-many mappings by cycling over the original domains augmented with auxiliary latent spaces. By marginalizing out auxiliary variables, we can model many-to-many mappings in between the domains.

ods such as CycleGAN (Zhu et al., 2017a) learn a generator which produces images in one domain given images from the other. Without the use of pairing information, there are many possible mappings that could be inferred. To reduce the space of the possible mappings, these models are typically trained with a *cycle-consistency* constraint which enforces a strong connection across domains, by requiring that mapping an image from the source domain to the target domain and then back to source will result in the same starting image. This framework has been shown to learn convincing mappings across image domains and proved successful in a variety of related applications (Tung et al., 2017; Wolf et al., 2017; Hoffman et al., 2017).

One major limitation of CycleGAN is that it only learns one-to-one mappings, i.e. the model associates each input image with a single output image. We believe that most relationships across domains are more complex, and better characterized as *many-to-many*. For example, consider mapping silhouettes of shoes to images of shoes. While the mapping that CycleGAN learns can be superficially convincing (e.g. it produces a single reasonable shoe with a particular style), we would like to learn a mapping that can capture diversity of the output (e.g. produces multiple shoes with different styles). The limits of one-to-one mappings are more dramatic when the source domain and target domain substantially differ. For instance, it would be difficult to learn a CycleGAN model when the two domains are descriptive facial attributes and images of faces.

¹Montreal Institute for Learning Algorithms (MILA), Canada. ²Microsoft Research Montreal, Canada. ³CIFAR Fellow. [†]Work partly done at MSR Montreal. Correspondence to: Amjad Almahairi <amjad.almahairi@umontreal.ca>.

We propose a model for learning many-to-many mappings between domains from unpaired data. Specifically, we “augment” each domain with auxiliary latent variables and extend CycleGAN’s training procedure to the augmented spaces. The mappings in our model take as input a sample from the source domain and a latent variable, and output both a sample in the target domain and a latent variable (Fig. 1b). The learned mappings are one-to-one in the augmented space, but many-to-many in the original domains after marginalizing over the latent variables.

Our contributions are as follows. (i) We introduce the Augmented CycleGAN model for learning many-to-many mappings across domains in an unsupervised way. (ii) We show that our model can learn mappings which produce a diverse set of outputs for each input. (iii) We show that our model can learn mappings across substantially different domains, and we apply it in a semi-supervised setting for mapping between faces and attributes with competitive results.

2. Unsupervised Learning of Mappings Between Domains

2.1. Problem Setting

Given two domains A and B , we assume there exists a mapping, potentially many-to-many, between their elements. The objective is to recover this mapping using *unpaired* samples from distributions $p_d(a)$ and $p_d(b)$ in each domain. This can be formulated as a conditional generative modeling task where we try to estimate the *true conditionals* $p(a|b)$ and $p(b|a)$ using samples from the true marginals. An important assumption here is that elements in domains A and B are highly dependent; otherwise, it is unlikely that the model would uncover a meaningful relationship without any pairing information.

2.2. CycleGAN Model

The CycleGAN model (Zhu et al., 2017a) estimates these conditionals using two mappings $G_{AB} : A \mapsto B$ and $G_{BA} : B \mapsto A$, parameterized by neural networks, which satisfy the following constraints:

1. **Marginal matching:** The output of each mapping should match the empirical distribution of the target domain, when marginalized over the source domain.
2. **Cycle-consistency:** Mapping an element from one domain to the other, and then back, should produce a sample close to the original element.

Marginal matching in CycleGAN is achieved using the generative adversarial networks framework (GAN) (Goodfellow et al., 2014). Mappings G_{AB} and G_{BA} are given by neural networks trained to fool adversarial discriminators D_B and

D_A , respectively. Enforcing marginal matching on target domain B , marginalized over source domain A , involves minimizing an adversarial objective with respect to G_{AB} :

$$\mathcal{L}_{\text{GAN}}^B(G_{AB}, D_B) = \mathbb{E}_{b \sim p_d(b)} [\log D_B(b)] + \mathbb{E}_{a \sim p_d(a)} [\log(1 - D_B(G_{AB}(a)))] \quad (1)$$

while the discriminator D_B is trained to maximize it. A similar adversarial loss $\mathcal{L}_{\text{GAN}}^A(G_{BA}, D_A)$ is defined for marginal matching in the reverse direction.

Cycle-consistency enforces that, when starting from a sample a from A , the reconstruction $a' = G_{BA}(G_{AB}(a))$ remains close to the original a . For image domains, closeness between a and a' is typically measured with L_1 or L_2 norms. When using the L_1 norm, cycle-consistency starting from A can be formulated as:

$$\mathcal{L}_{\text{CYC}}^A(G_{AB}, G_{BA}) = \mathbb{E}_{a \sim p_d(a)} \|G_{BA}(G_{AB}(a)) - a\|_1 \quad (2)$$

And similarly for cycle-consistency starting from B . The full CycleGAN objective is given by:

$$\mathcal{L}_{\text{GAN}}^A(G_{BA}, D_A) + \mathcal{L}_{\text{GAN}}^B(G_{AB}, D_B) + \gamma \mathcal{L}_{\text{CYC}}^A(G_{AB}, G_{BA}) + \gamma \mathcal{L}_{\text{CYC}}^B(G_{AB}, G_{BA}), \quad (3)$$

where γ is a hyper-parameter that balances between marginal matching and cycle-consistency.

The success of CycleGAN can be attributed to the complementary roles of marginal matching and cycle-consistency in its objective. Marginal matching encourages generating realistic samples in each domain. Cycle-consistency encourages a tight relationship between domains. It may also help prevent multiple items from one domain mapping to a single item from the other, analogous to the troublesome mode collapse in adversarial generators (Li et al., 2017).

2.3. Limitations of CycleGAN

A fundamental weakness of the CycleGAN model is that it learns deterministic mappings. In CycleGAN, and in other similar models (Kim et al., 2017; Yi et al., 2017), the conditionals between domains correspond to delta functions: $\hat{p}(a|b) = \delta(G_{BA}(b))$ and $\hat{p}(b|a) = \delta(G_{AB}(a))$, and cycle-consistency forces the learned mappings to be inverses of each other. When faced with complex cross-domain relationships, this results in CycleGAN learning an arbitrary one-to-one mapping instead of capturing the true, structured conditional distribution more faithfully. Deterministic mappings are also an obstacle to optimizing cycle-consistency when the domains differ substantially in complexity, in which case mapping from one domain (e.g. class labels) to the other (e.g. real images) is generally one-to-many. Next, we dis-

cuss how to extend CycleGAN to capture more expressive relationships across domains.

2.4. CycleGAN with Stochastic Mappings

A straightforward approach for extending CycleGAN to model many-to-many relationships is to equip it with stochastic mappings between A and B . Let Z be a latent space with a standard Gaussian prior $p(z)$ over its elements. We define mappings $G_{AB} : A \times Z \mapsto B$ and $G_{BA} : B \times Z \mapsto A$ ¹. Each mapping takes as input a vector of auxiliary noise and a sample from the source domain, and generates a sample in the target domain. Therefore, by sampling different $z \sim p(z)$, we could in principle generate multiple b 's conditioned on the same a and vice-versa. We can write the marginal matching loss on domain B as:

$$\mathcal{L}_{\text{GAN}}^B(G_{AB}, D_B) = \mathbb{E}_{b \sim p_d(b)} \left[\log D_B(b) \right] + \mathbb{E}_{\substack{a \sim p_d(a) \\ z \sim p(z)}} \left[\log(1 - D_B(G_{AB}(a, z))) \right]. \quad (4)$$

Cycle-consistency starting from A is now given by:

$$\mathcal{L}_{\text{CYC}}^A(G_{AB}, G_{BA}) = \mathbb{E}_{\substack{a \sim p_d(a) \\ z_1, z_2 \sim p(z)}} \left\| G_{BA}(G_{AB}(a, z_1), z_2) - a \right\|_1 \quad (5)$$

The full training loss is similar to the objective in Eqn. 3. We refer to this model as *Stochastic CycleGAN*.

In principle, stochastic mappings can model multi-modal conditionals, and hence generate a richer set of outputs than deterministic mappings. However, Stochastic CycleGAN suffers from a fundamental flaw: the cycle-consistency in Eq. 5 encourages the mappings to ignore the latent z . Specifically, the unimodality assumption implicit in the reconstruction error from Eq. 5 forces the mapping G_{BA} to be many-to-one when cycling $A \rightarrow B \rightarrow A'$, since any b generated for a given a must map to $a' = G_{BA}(b, z) \approx a$, for all z . For the cycle $B \rightarrow A \rightarrow B'$, G_{AB} is similarly forced to be many-to-one. The only way for G_{BA} and G_{AB} to be both many-to-one and mutual inverses is if they collapse to being (roughly) one-to-one. We could possibly mitigate this degeneracy by introducing a VAE-like encoder and exchanging the L_1 error in Eq. 5 for a more complex variational bound on conditional log-likelihood. In the next section, we discuss an alternative approach to learning complex, stochastic mappings between domains.

3. Approach

In order to learn many-to-many mappings across domains, we propose to learn to map between pairs of items $(a, z_b) \in$

¹To avoid clutter in notation, we reuse the same symbols of deterministic mappings.

$A \times Z_b$ and $(b, z_a) \in B \times Z_a$, where Z_a and Z_b are latent spaces that capture any missing information when transforming an element from A to B , and vice-versa. For example, when generating a female face ($b \in B$) which resembles a male face ($a \in A$), the latent code $z_b \in Z_b$ can capture female face variations (e.g. hair length or style) independent from a . Similarly, $z_a \in Z_a$ captures variations in a generated male face independent from the given female face. This approach can be described as learning mappings between *augmented spaces* $A \times Z_b$ and $B \times Z_a$ (Figure 1b); hence, we call it *Augmented CycleGAN*. By learning to map a pair $(a, z_b) \in A \times Z_b$ to $(\tilde{b}, z_a) \in B \times Z_a$, we can (i) learn a stochastic mapping from a to multiple items in B by sampling different $z_b \in Z_b$, and (ii) infer latent codes z_a containing information about a not captured in the generated b , which allows for doing proper reconstruction of a . As a result, we are able to optimize both marginal matching and cycle consistency while using stochastic mappings. We present details of our approach in the next sections.²

3.1. Augmented CycleGAN

Our proposed model has four components. First, the two mappings $G_{AB} : A \times Z_b \mapsto B$ and $G_{BA} : B \times Z_a \mapsto A$, which are the conditional generators of items in each domain. These models are similar to those used in Stochastic CycleGAN. We also have two encoders $E_A : A \times B \mapsto Z_a$ and $E_B : A \times B \mapsto Z_b$, which enable optimization of cycle-consistency with stochastic, structured mappings. All components are parameterized with neural networks – see Fig. 2. We define mappings over augmented spaces in our model as follows. Let $p(z_a)$ and $p(z_b)$ be standard Gaussian priors over Z_a and Z_b , which are independent from $p_d(b)$ and $p_d(a)$. Given a pair $(a, z_b) \sim p_d(a)p(z_b)$, we generate a pair (\tilde{b}, \tilde{z}_a) as follows:

$$\tilde{b} = G_{AB}(a, z_b), \quad \tilde{z}_a = E_A(a, \tilde{b}). \quad (6)$$

That is, we first generate a sample in domain B , then we use it along with a to generate latent code \tilde{z}_a . Note here that by sampling different $z_b \sim p(z_b)$, we can generate multiple \tilde{b} 's conditioned on the same a . In addition, given the pair (a, \tilde{b}) , we can recover information about a which is not captured in \tilde{b} , via \tilde{z}_a . Similarly, given a pair $(b, z_a) \sim p_d(b)p(z_a)$, we generate a pair (\tilde{a}, \tilde{z}_b) as follows:

$$\tilde{a} = G_{BA}(b, z_a), \quad \tilde{z}_b = E_B(b, \tilde{a}). \quad (7)$$

Learning in Augmented CycleGAN follows a similar approach to CycleGAN – optimizing both marginal matching and cycle-consistency losses, albeit over augmented spaces.

²Our model captures many-to-many relationships because it captures both one-to-many and many-to-one: one item in A maps to many items in B , and many items in B map to one item in A (cycle). The same is true in the other direction.




Figure 2: Cycles starting from augmented spaces in Augmented CycleGAN. Model components identified with color coding.

Marginal Matching Loss We adopt an adversarial approach for marginal matching over $B \times Z_a$ where we use two independent discriminators D_B and D_{Z_a} to match generated pairs to real samples from the independent priors $p_d(b)$ and $p(z_a)$, respectively. Marginal matching loss over B is defined as in Eqn 4. Marginal matching over Z_a is given by:

$$\begin{aligned} \mathcal{L}_{\text{GAN}}^{Z_a}(E_A, G_{AB}, D_{Z_a}) = & \mathbb{E}_{z_a \sim p(z_a)} [\log D_{Z_a}(z_a)] + \\ & \mathbb{E}_{\substack{a \sim p_d(a) \\ z_b \sim p(z_b)}} [\log(1 - D_{Z_a}(\tilde{z}_a))], \end{aligned} \quad (8)$$

where \tilde{z}_a is defined by Eqn 6. As in CycleGAN, the goal of marginal matching over B is to insure that generated samples \tilde{b} are realistic. For latent codes \tilde{z}_a , marginal matching acts as a regularizer for the encoder, encouraging the marginalized encoding distribution to match a simple prior $p(z_a)$. This is similar to adversarial regularization of latent codes in adversarial autoencoders (Makhzani et al., 2016). We define similar losses $\mathcal{L}_{\text{GAN}}^A(G_{BA}, D_A)$ and $\mathcal{L}_{\text{GAN}}^{Z_b}(E_B, G_{BA}, D_{Z_b})$ for marginal matching over $A \times Z_b$.

Cycle Consistency Loss We define two cycle-consistency constraints in Augmented CycleGAN starting from each of the two augmented spaces, as shown in Fig. 2. In cycle-consistency starting from $A \times Z_b$, we ensure that given a pair $(a, z_b) \sim p_d(a)p(z_b)$, the model is able to produce a faithful reconstruction of it after being mapped to (\tilde{b}, \tilde{z}_a) . This is achieved with two losses; first for reconstructing $a \sim p_d(a)$:

$$\begin{aligned} \mathcal{L}_{\text{CYC}}^A(G_{AB}, G_{BA}, E_A) = & \mathbb{E}_{\substack{a \sim p_d(a) \\ z_b \sim p(z_b)}} \|a' - a\|_1, \\ \tilde{b} = G_{AB}(a, z_b), \tilde{z}_a = E_A(a, \tilde{b}), a' = G_{BA}(\tilde{b}, \tilde{z}_a). \end{aligned} \quad (9)$$

The second is for reconstructing $z_b \sim p(z_b)$:

$$\begin{aligned} \mathcal{L}_{\text{CYC}}^{Z_b}(G_{AB}, E_B) = & \mathbb{E}_{\substack{a \sim p_d(a) \\ z_b \sim p(z_b)}} \|z'_b - z_b\|_1, \\ z'_b = E_B(a, \tilde{b}), \tilde{b} = G_{AB}(a, z_b). \end{aligned} \quad (10)$$




Figure 3: Augmented CycleGAN when pairs $(a, b) \sim p_d(a, b)$ from the true joint distribution are observed. Instead of producing \tilde{b} and \tilde{a} , the model uses samples from the joint distribution.

These reconstruction costs represent an autoregressive decomposition of the basic CycleGAN cycle-consistency cost from Eq. 2, after extending it to the augmented domains. Specifically, we decompose the required reconstruction distribution $p(b, z_a | a, z_b)$ into the conditionals $p(b | a, z_b)$ and $p(z_a | a, z_b, b)$.

Just like in CycleGAN, the cycle loss in Eqn. 9 enforces the dependency of generated samples in B on samples of A . Thanks to the encoder E_A , the model is able to reconstruct a because it can recover information loss in generated \tilde{b} through \tilde{z}_a . On the other hand, the cycle loss in Eqn. 10 enforces the dependency of a generated sample \tilde{b} on the given latent code z_b . In effect, it increases the mutual information between z_b and b conditioned on a , i.e. $I(b, z_b | a)$ (Chen et al., 2016; Li et al., 2017).

Training Augmented CycleGAN in the direction $A \times Z_b$ to $B \times Z_a$ is done by optimizing:

$$\begin{aligned} \mathcal{L}_{\text{GAN}}^B(D_B, G_{AB}) + \mathcal{L}_{\text{GAN}}^{z_a}(D_{Z_a}, E_A, G_{AB}) + \\ \gamma_1 \mathcal{L}_{\text{CYC}}^A(G_{AB}, G_{BA}, E_A) + \gamma_2 \mathcal{L}_{\text{CYC}}^{z_b}(G_{AB}, E_B), \end{aligned} \quad (11)$$

where γ_1 and γ_2 are a hyper-parameters used to balance objectives. We define a similar objective for the direction going from $B \times Z_a$ to $A \times Z_b$, and train the model on both objectives simultaneously.

3.2. Semi-supervised Learning with Augmented CycleGAN

In cases where we have access to paired data, we can leverage it to train our model in a semi-supervised setting (Fig. 3). Given pairs sampled from the true joint,

i.e. $(a, b) \sim p_d(a, b)$, we can define a supervision cost for the mapping G_{BA} as follows:

$$\mathcal{L}_{\text{SUP}}^A(G_{BA}, E_A) = \mathbb{E}_{(a,b) \sim p_d(a,b)} \|G_{BA}(b, \tilde{z}_a) - a\|_1, \quad (12)$$

where $\tilde{z}_a = E_A(a, b)$ infers a latent code which can produce a given b via $G_{BA}(b, \tilde{z}_a)$. We also apply an adversarial regularization cost on the encoder, in the form of Eqn. 8. Similar supervision and regularization costs can be defined for G_{BA} and E_B , respectively.

3.3. Modeling Stochastic Mappings

We note here some design choices that we found important for training our stochastic mappings. We discuss architectural and training details further in Sec. 5. In order to allow the latent codes to capture diversity in generated samples, we found it important to inject latent codes to layers of the network which are closer to the inputs. This allows the injected codes to be processed with a larger number of remaining layers and therefore capture high-level variations of the output, as opposed to small pixel-level variations. We also found that Conditional Normalization (CN) (Dumoulin et al.; Perez et al., 2017) for conditioning layers can be more effective than concatenation, which is more commonly used (Radford et al., 2015; Zhu et al., 2017b). The basic idea of CN is to replace parameters of affine transformations in normalization layers (Ioffe & Szegedy, 2015) of a neural network with a learned function of the conditioning information. We apply CN by learning two linear functions f and g which take a latent code z as input and output scale and shift parameters of normalization layers in intermediate layers, i.e. $\gamma = f(z)$ and $\beta = g(z)$. When activations are normalized over spatial dimensions only, we get Conditional Instance Normalization (CIN), and when they are also normalized over batch dimension, we get Conditional Batch Normalization (CBN).

4. Related Work

There has been a surge of interest recently in unsupervised learning of cross-domain mappings, especially for image translation tasks. Previous attempts for image-to-image translation have unanimously relied on GANs to learn mappings that produce compelling images. In order to constrain learned mappings, some methods have relied on cycle-consistency based constraints similar to CycleGAN (Kim et al., 2017; Yi et al., 2017; Royer et al., 2017), while others relied on weight sharing constraints (Liu & Tuzel, 2016; Liu et al., 2017). However, the focus in all of these methods was on learning conditional image generators that produce single output images given the input image. Notably, Liu et al. (2015) propose to map inputs from both domains into a

shared latent space. This approach may constrain too much the space of learnable mappings, for example in cases where the domains differ substantially (class labels and images).

Unsupervised learning of mappings have also been addressed recently in language translation, especially for machine translation (Lample et al., 2017) and text style transfer (Shen et al., 2017). These methods also rely on some notion of cycle-consistency over domains in order to constrain the learned mappings. They rely heavily on the power of the RNN-based decoders to capture complex relationships across domains while we propose to use auxiliary latent variables. The two approaches may be synergistic, as it was recently suggested in (Gulrajani et al., 2016).

Recently, Zhu et al. (2017b) proposed the BiCycleGAN model for learning multi-modal mappings but in fully supervised setting. This model extends the pix2pix framework in (Isola et al., 2017) by learning a stochastic mapping from the source to the target, and shows interesting diversity in the generated samples. Several modeling choices in BiCycleGAN resemble our proposed model, including the use of stochastic mappings and an encoder to handle multi-modal targets. However, our approach focuses on unsupervised many-to-many mappings, which allows it to handle domains with no or very little paired data.

5. Experiments

5.1. Edges-to-Photos

We first study a one-to-many image translation task between edges (domain A) and photos of shoes (domain B).³ Training data is composed of almost 50K shoe images with corresponding edges (Yu & Grauman, 2014; Zhu et al., 2016; Isola et al., 2017), but as in previous approaches (e.g. (Kim et al., 2017)), we assume no pairing information while training unsupervised models. Stochastic mappings in our Augmented CycleGAN (AugCGAN) model are based on ResNet conditional image generators of (Zhu et al., 2017a), where we inject noise with CIN to all intermediate layers. As baselines, we train: CycleGAN, Stochastic CycleGAN (StochCGAN) and Triangle-GAN (Δ -GAN) of (Gan et al., 2017) which share the same architectures and training procedure for fair comparison.⁴

Quantitative Results First, we evaluate conditionals learned by each model by measuring the ability of the model of generating a specific edge-shoe pair from a test set. We follow the same evaluation methodology adopted in (Metz et al., 2016; Xiang & Li, 2017), which opt for an

³ Public code available at: https://github.com/aalmah/augmented_cyclegan

⁴ Δ -GAN architecture differs only in the two discriminators, which match conditionals/joints instead of marginals.




Figure 4: Shoes reconstruction error given a generated edge as a function of the Gaussian noise ϵ injected in the generated edge.

Model (Paired %)	Avg. L_1
CycleGAN (0%)	0.1837
StochCGAN (0%)	0.0794
Δ -GAN [†] (10%)	0.0748
AugCGAN (0%)	0.0698
AugCGAN (10%)	0.0562

Table 1: Reconstruction error for shoes given edges in the test set. [†]Same architecture as our model.

Model (Paired %)	MSE
Δ -GAN [*] (10%)	0.0102
Δ -GAN [†] (10%)	0.0096
Δ -GAN [*] (20%)	0.0092
AugCGAN (0%)	0.0079
AugCGAN (10%)	0.0052

Table 2: MSE on edges given shoes in the test set. ^{*} From (Gan et al., 2017). [†]Same architecture as our model.

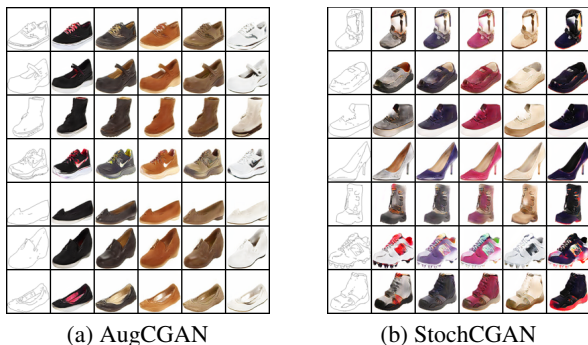


Figure 5: Given an edge from the data distribution (leftmost column), we generate shoes by sampling five $z_b \sim p(z_b)$. Models generate diverse shoes when edges are from the data distribution.

inference-via-optimization approach to estimate the reconstruction error of a specific shoe given an edge. Specifically, given a trained model with mapping G_{AB} and an edge-shoe pair (a, b) in the test set, we solve the optimization task $z_b^* = \arg \min_{z_b} \|G_{AB}(a, z_b) - b\|_1$ and compute reconstruction error $\|G_{AB}(a, z_b^*) - b\|_1$. Optimization is done with RMSProp as in (Xiang & Li, 2017). We show the average errors over a predefined test set of 200 samples in Table 1 for: AugCGAN (unsupervised and semi-supervised with 10% paired data), unsupervised CycleGAN and StochCGAN, and a semi-supervised Δ -GAN, all sharing the same architecture. Our unsupervised AugCGAN model outperforms all baselines including semi-supervised Δ -GAN, which indicates that reconstruction-based cycle-consistency is more effective in learning conditionals than the adversarial approach of Δ -GAN. As expected, adding 10% supervision to AugCGAN improves shoe predictions further. In addition, we evaluate edge predictions given real shoes from test set as well. We report mean squared error (MSE) similar to (Gan et al., 2017), where we normalize over all edge pixels. The Δ -GAN model with our architecture outperforms



Figure 6: Cycles from both models starting from a real edge and a real shoe (left and right respectively in each subfigure). The ability for StochCGAN to reconstruct shoes is surprising and is due to the “steganography” effect (see text).

the one reported in (Gan et al., 2017), but is outperformed by our unsupervised AugCGAN model. Again, adding 10% supervision to AugCGAN reduces MSE even further.

Qualitative Results We qualitatively compare the mappings learned by our model AugCGAN and StochCGAN. Fig. 6 shows generated images of shoes given an edge $a \sim p_d(a)$ (row) and $z_b \sim p(z_b)$ (column) from both model, and Fig. 5 shows cycles starting from edges and shoes. Note that here the edges are sampled from the data distribution and not produced by the learnt stochastic mapping G_{BA} . In this case, both models can (i) generate diverse set of shoes with color variations mostly defined by z_b , and (ii) perform reconstructions of both edges and shoes.

While we expect our model to achieve these results, the fact that StochCGAN can reconstruct shoes perfectly without an inference model may seem at first surprising. However, this can be explained by the “steganography” behavior of CycleGAN (Chu et al., 2017): the model hides in the generated edge \tilde{a} imperceptible information about a given shoe b (e.g. its color), in order to satisfy cycle-consistency without being

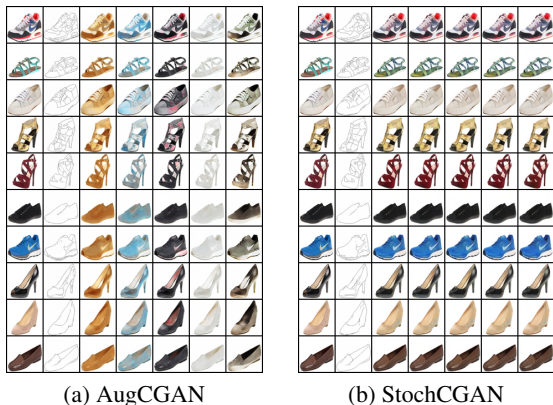


Figure 7: Given a shoe from the data distribution (leftmost column), we generate an edge using the model (second column). Then, we generate shoes by sampling five $z_b \sim p(z_b)$. When edges are generated by the model, StochCGAN collapses to a single mode of the shoes distribution and generate the same shoe.

penalized by the discriminator on A . A good model of the true conditionals $p(b|a)$, $p(a|b)$ should reproduce the hidden joint distribution and consequently the marginals by alternatively sampling from conditionals. Therefore, we examine the behavior of the models when edges are generated from the model itself (instead of the empirical data distribution). In Fig. 7, we plot multiple generated shoes given an edge generated by the model, i.e. \tilde{a} , and 5 different z_b sampled from $p(z_b)$. In StochCGAN, the mapping $G_{BA}(\tilde{a}, z_b)$ collapses to a deterministic function generating a single shoe for every z_b . This distinction between behaviour on real and synthetic data is undesirable, e.g. regularization benefits of using unpaired data may be reduced if the model slips into this regime switching style. In AugCGAN, on the other hand, the mapping seem to closely capture the diversity in the conditional distribution of shoes given edges. Furthermore, in Fig. 8, we run a Markov chain by generating from the learned mappings multiple times, starting from a real shoe. Again AugCGAN produces diverse samples while StochCGAN seems to collapse to a single mode.

We investigate “steganography” behavior in both AugCGAN and StochCGAN using a similar approach to (Chu et al., 2017), where we corrupt generated edges with noise sampled from $\mathcal{N}(0, \epsilon^2)$, and compute reconstruction error of shoes. Fig. 4 shows L_1 reconstruction error as we increase ϵ . AugCGAN seems more robust to corruption of edges than in StochCGAN, which confirms that information is being stored in the latent codes instead of being completely hidden in generated edges.

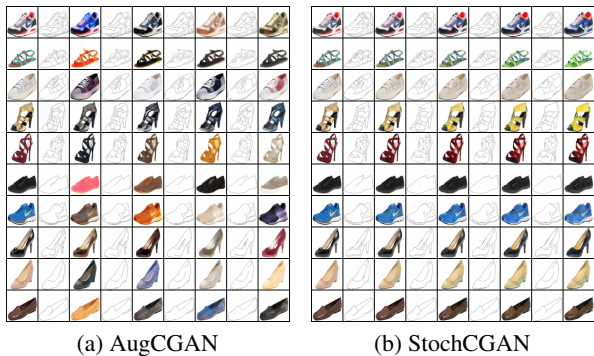


Figure 8: We perform multiple generation cycles from the model by applying the learned mappings in turn. StochCGAN cycles collapse to the same shoe at each step which indicates that it doesn’t capture the data distribution.




Figure 9: Given a male face from the data distribution (leftmost column), we generate 8, 128×128 female faces with AugCGAN by sampling $z_b \sim p(z_b)$.

5.2. Male-to-Female

We study another image translation task of translating between male and female faces. Data is based on CelebA dataset (Liu et al., 2015) where we split it into two separate domains using provided attributes. Several key features distinguish this task from other image-translation tasks: (i) there is no predefined correspondence in real data of each domain, (ii) the relationship is many-to-many between domains, as we can map a male to female face, and vice-versa, in many possible ways, and (iii) capturing realistic variations in generated faces requires transformations that go beyond simple color and texture changes. The architecture of stochastic mappings are based on U-NET conditional image generators of (Isola et al., 2017), and again with noise injected to all intermediate layers. Fig. 9 shows results of applying our model to this task on 128×128 resolution CelebA images. We can see that our model depicts meaningful variations in generated faces without compromising their realistic appearance. In Fig. 10 we show 64×64 generated samples in both domains from our model ((a) and (b)), and compare them to both: (c) our model but with noise injected noise only in last 3 layers of the G_{AB} ’s

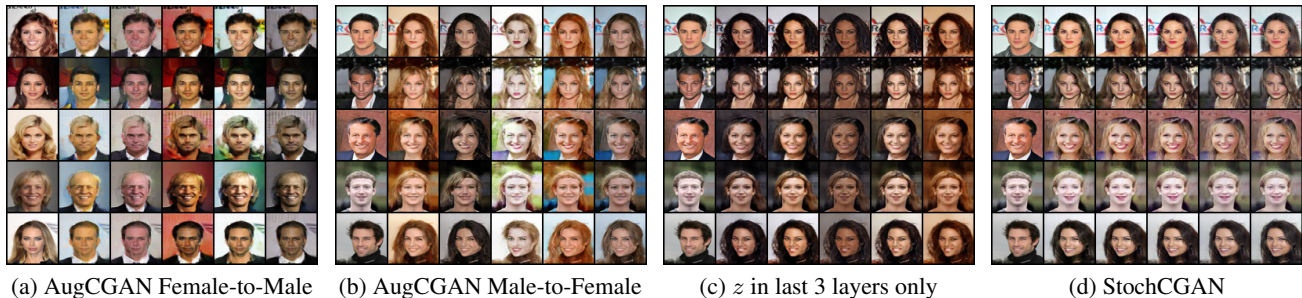


Figure 10: Generated 64×64 faces given a real face image from the other domain and multiple latent codes from prior.

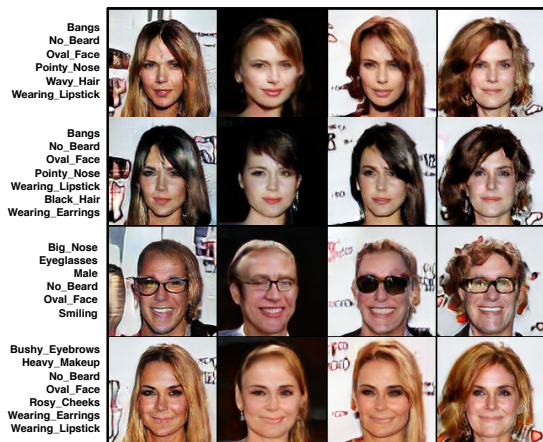


Figure 11: Conditional generation given attributes learned by our model in the Attributes-to-Faces task. We sample a set of attributes from the data distribution and generate 4 faces by sampling latent codes from $z_b \sim p(z_b)$.

network, and (d) StochCGAN with the same architecture. We can see that in Fig. 10-(c) variations are very limited, which highlights the importance of processing latent code with multiple layers. StochCGAN in this task produces almost no variations at all, which highlights the importance of proper optimization of cycle-consistency for capturing meaningful variations. We verify these results quantitatively using LPIPS distance (Zhang et al., 2018), where we average distance between 1000 pairs of generated female faces (10 random pairs from 100 male faces). AugCGAN (Fig. 10-(b)) achieves highest LPIPS diversity score with 0.108 ± 0.003 , while AugCGAN with z in low-level layers (Fig. 10-(c)) gets 0.059 ± 0.001 , and finally StochCGAN (Fig. 10-(d)) gets 0.008 ± 0.000 , i.e. severe mode collapse.

5.3. Attributes-to-Faces

In this task, we make use of the CelebA dataset in order map from descriptive facial attributes A to images of faces B and vice-versa. We report both quantitative and qualitative results. For the quantitative results, we follow (Gan et al.,

Model	P@10 / NDCG@10	
	$s = 1\%$	$s = 10\%$
Triple-GAN [†]	40.97 / 50.74	62.13 / 73.56
Δ -GAN [†]	53.21 / 58.39	63.68 / 75.22
Baseline Classifier	63.36 / 79.25	67.34 / 84.21
AugCGAN	64.38 / 80.59	68.83 / 85.51

Table 3: CelebA semi-supervised attribute prediction with supervision $s = 1\%$ and 10% . [†] From (Gan et al., 2017).

2017) and test our models in a semi-supervised attribute prediction setting. We let the model train on all the available data without the pairing information and only train with a small amount of paired data as described in Sec. 3.2. We report Precision (P) and normalized Discounted Cumulative Gain (nDCG) as the two metrics for multi-label classification problems. As an additional baseline, we also train a supervised classifier (which has the same architecture as G_{BA}) on the paired subset. The results are reported in Table 3. In Fig. 11, we show some generation obtained from the model in the direction attributes to faces. We can see that the model generates reasonable diverse faces for the same set of attributes.

6. Conclusion

In this paper we have introduced the Augmented CycleGAN model for learning many-to-many cross-domain mappings in unsupervised fashion. This model can learn stochastic mappings which leverage auxiliary noise to capture multi-modal conditionals. Our experimental results verify quantitatively and qualitatively the effectiveness of our approach in image translation tasks. Furthermore, we apply our model in a challenging task of learning to map across attributes and faces, and show that it can be used effectively in a semi-supervised learning setting.

Acknowledgements

Authors would like to thank Zihang Dai for valuable discussions and feedback. We are also grateful for ICML anonymous reviewers for their comments.

References

- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2172–2180, 2016.
- Chu, C., Zhmoginov, A., and Sandler, M. Cyclegan: a master of steganography. *arXiv preprint arXiv:1712.02950*, 2017.
- Dumoulin, V., Shlens, J., and Kudlur, M. A learned representation for artistic style.
- Gan, Z., Chen, L., Wang, W., Pu, Y., Zhang, Y., Liu, H., Li, C., and Carin, L. Triangle generative adversarial networks. In *Advances in Neural Information Processing Systems*, pp. 5253–5262, 2017.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Gulrajani, I., Kumar, K., Ahmed, F., Taiga, A. A., Visin, F., Vazquez, D., and Courville, A. Pixelvae: A latent variable model for natural images. *arXiv preprint arXiv:1611.05013*, 2016.
- Hoffman, J., Tzeng, E., Park, T., Zhu, J.-Y., Isola, P., Saenko, K., Efros, A. A., and Darrell, T. Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213*, 2017.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456, 2015.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017.
- Kim, T., Cha, M., Kim, H., Lee, J., and Kim, J. Learning to discover cross-domain relations with generative adversarial networks. *arXiv preprint arXiv:1703.05192*, 2017.
- Lample, G., Denoyer, L., and Ranzato, M. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*, 2017.
- Li, C., Liu, H., Chen, C., Pu, Y., Chen, L., Heno, R., and Carin, L. Alice: Towards understanding adversarial learning for joint distribution matching. In *Advances in Neural Information Processing Systems*, pp. 5501–5509, 2017.
- Liu, M.-Y. and Tuzel, O. Coupled generative adversarial networks. In *Advances in neural information processing systems*, pp. 469–477, 2016.
- Liu, M.-Y., Breuel, T., and Kautz, J. Unsupervised image-to-image translation networks. In *Advances in Neural Information Processing Systems 30*, pp. 700–708, 2017.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- Makhzani, A., Shlens, J., Jaitly, N., and Goodfellow, I. Adversarial autoencoders. In *International Conference on Learning Representations*, 2016.
- Metz, L., Poole, B., Pfau, D., and Sohl-Dickstein, J. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016.
- Perez, E., Strub, F., De Vries, H., Dumoulin, V., and Courville, A. Film: Visual reasoning with a general conditioning layer. *arXiv preprint arXiv:1709.07871*, 2017.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Royer, A., Bousmalis, K., Gouws, S., Bertsch, F., Moressi, I., Cole, F., and Murphy, K. Xgan: Unsupervised image-to-image translation for many-to-many mappings. *arXiv preprint arXiv:1711.05139*, 2017.
- Shen, T., Lei, T., Barzilay, R., and Jaakkola, T. Style transfer from non-parallel text by cross-alignment. *arXiv preprint arXiv:1705.09655*, 2017.
- Tung, H.-Y. F., Harley, A. W., Seto, W., and Fragkiadaki, K. Adversarial inverse graphics networks: Learning 2d-to-3d lifting and image-to-image translation from unpaired supervision. In *The IEEE International Conference on Computer Vision (ICCV)*, volume 2, 2017.
- Wolf, L., Taigman, Y., and Polyak, A. Unsupervised creation of parameterized avatars. *arXiv preprint arXiv:1704.05693*, 2017.
- Xiang, S. and Li, H. On the effects of batch and weight normalization in generative adversarial networks. *stat*, 1050:22, 2017.

- Yi, Z., Zhang, H., Gong, P. T., et al. Dualgan: Unsupervised dual learning for image-to-image translation. *arXiv preprint arXiv:1704.02510*, 2017.
- Yu, A. and Grauman, K. Fine-grained visual comparisons with local learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 192–199, 2014.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep networks as a perceptual metric. In *CVPR*, 2018.
- Zhu, J.-Y., Krähenbühl, P., Shechtman, E., and Efros, A. A. Generative visual manipulation on the natural image manifold. In *European Conference on Computer Vision*, pp. 597–613. Springer, 2016.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint arXiv:1703.10593*, 2017a.
- Zhu, J.-Y., Zhang, R., Pathak, D., Darrell, T., Efros, A. A., Wang, O., and Shechtman, E. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*, pp. 465–476, 2017b.