
JointGAN: Multi-Domain Joint Distribution Learning with Generative Adversarial Nets

Yunchen Pu¹ Shuyang Dai² Zhe Gan³ Weiyao Wang² Guoyin Wang² Yizhe Zhang³ Ricardo Henao²
Lawrence Carin²

Abstract

A new generative adversarial network is developed for joint distribution matching. Distinct from most existing approaches, that only learn conditional distributions, the proposed model aims to learn a joint distribution of multiple random variables (domains). This is achieved by learning to sample from conditional distributions between the domains, while simultaneously learning to sample from the marginals of each individual domain. The proposed framework consists of multiple generators and a single softmax-based critic, all jointly trained via adversarial learning. From a simple noise source, the proposed framework allows synthesis of draws from the marginals, conditional draws given observations from a subset of random variables, or complete draws from the full joint distribution. Most examples considered are for joint analysis of two domains, with examples for three domains also presented.

1. Introduction

Generative adversarial networks (GANs) (Goodfellow et al., 2014) have emerged as a powerful framework for modeling the draw of samples from complex data distributions. When trained on datasets of natural images, significant progress has been made on synthesizing realistic and sharp-looking images (Radford et al., 2016). Recent work has also extended the GAN framework for the challenging task of text generation (Yu et al., 2017; Zhang et al., 2017b). However, in its standard form, GAN models distributions in one domain, *i.e.*, for a single random variable.

This work was done when Yunchen Pu, Zhe Gan and Yizhe Zhang were Ph.D. students at Duke University. ¹Facebook, Menlo Park, CA, USA ²Duke University, Durham, NC, USA ³Microsoft Research, Redmond, WA, USA. Correspondence to: Yunchen Pu <pyc40@fb.com>, Shuyang Dai <shuyang.dai@duke.edu>, Zhe Gan <zhe.gan@microsoft.com>, Lawrence Carin <lcarin@duke.edu>.

Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, PMLR 80, 2018. Copyright 2018 by the author(s).

There has been recent interest in employing GAN ideas to learn *conditional* distributions for two random variables. This setting is of interest when one desires to synthesize (or infer) one random variable given an instance of another random variable. Example applications include generative models with (stochastic) latent variables (Mescheder et al., 2017; Tao et al., 2018), and conditional data synthesis (Isola et al., 2017; Reed et al., 2016), when both domains consist of observed pairs of random variables.

In this paper we focus on learning the *joint* distribution of multiple random variables using adversarial training. For the case of two random variables, conditional GAN (Mirza & Osindero, 2014) and Triangle GAN (Gan et al., 2017a) have been utilized for this task in the case that paired data are available. Further, adversarially learned inference (ALI) (Dumoulin et al., 2017; Donahue et al., 2017) and CycleGAN (Zhu et al., 2017; Kim et al., 2017; Yi et al., 2017) were developed for unsupervised learning, where the two-way mappings between two domains are learned without any paired data. These models are unified as the joint distribution matching problem by Li et al. (2017a). However, in all previous approaches the joint distributions are *not fully* learned, *i.e.*, the model only learns to sample from the *conditional* distributions, assuming access to the *marginal* distributions, which are typically instantiated as empirical samples from each individual domain (see Figure 1(b) for illustration). Therefore, only conditional data synthesis can be achieved due to the lack of a learned sample mechanism for the marginals.

It is desirable to build a generative-model learning framework from which one may sample from a fully-learned joint distribution. We design a new GAN framework that learns the joint distribution by decomposing it into the product of a marginal and a conditional distribution(s), each learned via adversarial training (see Figure 1(c) for illustration). The resulting model may then be employed in several distinct applications: (i) synthesis of draws from any of the marginals; (ii) synthesis of draws from the conditionals when other random variables are observed, *i.e.*, imputation; (iii) or we may simultaneously draw all random variables from the joint distribution.

For the special case of two random variables, the proposed model consists of four generators and a softmax critic function. The design includes two generators for the marginals, two for the conditionals, and a single 5-way critic (discriminator) trained to distinguish pairs of real data from four different kinds of synthetic data. These five modules are implemented as neural networks, which share parameters for efficiency and are optimized jointly via adversarial learning. We also consider an example with three random variables.

The contributions of this work are summarized as follows. (i) We present the first GAN-enabled framework that allows for full joint-distribution learning of multiple random variables. Unlike existing models, the proposed framework learns marginals and conditionals simultaneously. (ii) We share parameters of the generator models, and thus the resulting model does not have a significant increase in the number of parameters relative to prior work that only considered conditionals (ALI, Triangle GAN, CycleGAN, *etc.*) or marginals (traditional GAN). (iii) Unlike existing approaches, we consolidate all real *vs.* artificial sample comparisons into a single softmax-based critic function. (iv) While the main focus is on the case of two random variables, we extend the proposed model to learning the joint distribution of three or more random variables. (v) We apply the proposed model for both unsupervised and supervised learning paradigms.

2. Background

To simplify the presentation, we first consider joint modeling of two random variables, with the setup generalized in Sec. 3.2 to the case of more than two domains. For the two-random-variable case, consider marginal distributions $q(x)$ and $q(y)$ defined over two random variables $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, respectively. Typically, we have samples but not an explicit density form for $q(x)$ and $q(y)$, *i.e.*, ensembles $\{x_i\}_{i=1}^N$ and $\{y_j\}_{j=1}^M$ are available for learning. In general, their joint distribution can be written as the product of a marginal and a conditional in two ways: $q(x, y) = q(x)q(y|x) = q(y)q(x|y)$. One random variable can be synthesized (or inferred) given the other using conditional distributions, $q(x|y)$ and $q(y|x)$.

2.1. Generative Adversarial Networks

Nonparametric sampling from marginal distributions $q(x)$ and $q(y)$ can be accomplished via adversarial learning (Goodfellow et al., 2014), which provides a sampling mechanism that only requires gradient backpropagation, avoiding the need to explicitly adopt a form for the marginals. Specifically, instead of sampling directly from an (assumed) parametric distribution for the desired marginal, the target random variable is generated as a deterministic transformation of an easy-to-sample, independent

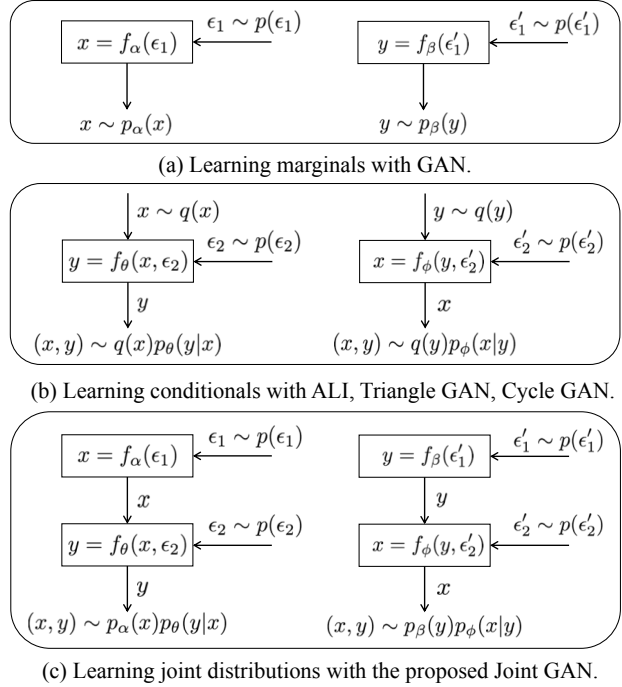


Figure 1. Comparing the generative process of GAN (Goodfellow et al., 2014), ALI (Dumoulin et al., 2017), Triangle GAN (Gan et al., 2017a), CycleGAN (Zhu et al., 2017) and the proposed JointGAN model. GAN in (a) only allows synthesizing samples from marginals, the models in (b) only allow for conditional data synthesis, while the proposed model in (c), allows data synthesis from both marginals and conditionals.

noise source, *e.g.*, Gaussian distribution. The sampling procedure for the marginals, *implicitly* defined as $x \sim p_\alpha(x)$ and $y \sim p_\beta(y)$, is carried out through the following two generative processes:

$$\tilde{x} = f_\alpha(\epsilon_1), \quad \epsilon_1 \sim p(\epsilon_1), \quad (1)$$

$$\tilde{y} = f_\beta(\epsilon'_1), \quad \epsilon'_1 \sim p(\epsilon'_1), \quad (2)$$

where $f_\alpha(\cdot)$ and $f_\beta(\cdot)$ are two *marginal* generators, specified as neural networks with parameters α and β , respectively. $p(\epsilon_1)$ and $p(\epsilon'_1)$ are assumed to be simple distributions, *e.g.*, isotropic Gaussian. The generative processes manifested by (1) and (2) are illustrated in Figure 1(a). Within the models, stochasticity in x and y is manifested via draws from $p(\epsilon_1)$ and $p(\epsilon'_1)$, and respective neural networks $f_\alpha(\cdot)$ and $f_\beta(\cdot)$ transform draws from these simple distributions such that they are approximately consistent with draws from $q(x)$ and $q(y)$.

For this purpose, GAN trains a ω -parameterized critic function $g_\omega(x)$, to distinguish samples generated from $p_\alpha(x)$ and $q(x)$. Formally, the minimax objective of GAN is

$$\min_{\alpha} \max_{\omega} \mathcal{L}_{\text{GAN}}(\alpha, \omega) = \mathbb{E}_{x \sim q(x)} [\log \sigma(g_\omega(x))] + \mathbb{E}_{\epsilon_1 \sim p(\epsilon_1)} [\log(1 - \sigma(g_\omega(\tilde{x})))] \quad (3)$$

with expectations $\mathbb{E}_{\mathbf{x} \sim q(\mathbf{x})}[\cdot]$ and $\mathbb{E}_{\epsilon_1 \sim p(\epsilon_1)}[\cdot]$ approximated via sampling, and $\sigma(\cdot)$ is the sigmoid function. As shown in Goodfellow et al. (2014), the equilibrium of the objective in (3) is achieved if and only if $p_\alpha(\mathbf{x}) = q(\mathbf{x})$.

Similarly, we can design a corresponding minimax objective that is similar to (3) to match the marginal $p_\beta(\mathbf{y})$ to $q(\mathbf{y})$.

2.2. Adversarially Learned Inference

In the same spirit, sampling from conditional distributions $q(\mathbf{x}|\mathbf{y})$ and $q(\mathbf{y}|\mathbf{x})$ can be also achieved as a deterministic transformation of two inputs, the variable in the source domain as a *covariate*, plus an independent source of noise. Specifically, the sampling procedure for the conditionals $\mathbf{y} \sim p_\theta(\mathbf{y}|\mathbf{x})$ and $\mathbf{x} \sim p_\phi(\mathbf{x}|\mathbf{y})$ is modeled as

$$\tilde{\mathbf{y}} = f_\theta(\mathbf{x}, \epsilon_2), \quad \mathbf{x} \sim q(\mathbf{x}), \quad \epsilon_2 \sim p(\epsilon_2), \quad (4)$$

$$\tilde{\mathbf{x}} = f_\phi(\mathbf{y}, \epsilon'_2), \quad \mathbf{y} \sim q(\mathbf{y}), \quad \epsilon'_2 \sim p(\epsilon'_2), \quad (5)$$

where $f_\theta(\cdot)$ and $f_\phi(\cdot)$ are two *conditional* generators, specified as neural networks with parameters θ and ϕ , respectively. In practice, the inputs of $f_\theta(\cdot)$ and $f_\phi(\cdot)$ are concatenated. As in GAN, $p(\epsilon_2)$ and $p(\epsilon'_2)$ are two simple distributions that provide the stochasticity when generating \mathbf{y} given \mathbf{x} , and *vice versa*. The conditional generative processes manifested in (4) and (5) are illustrated in Figure 1(b),

ALI (Dumoulin et al., 2017) learns the two desired conditionals by matching joint distributions $p_\theta(\mathbf{x}, \mathbf{y}) = q(\mathbf{x})p_\theta(\mathbf{y}|\mathbf{x})$ and $p_\phi(\mathbf{x}, \mathbf{y}) = q(\mathbf{y})p_\phi(\mathbf{x}|\mathbf{y})$, using a critic function $g_\omega(\mathbf{x}, \mathbf{y})$ similar to (3). The minimax objective of ALI can be written as

$$\min_{\theta, \phi} \max_{\omega} \mathbb{E}_{(\mathbf{x}, \tilde{\mathbf{y}}) \sim p_\theta(\mathbf{x}, \mathbf{y})} [\log \sigma(g_\omega(\mathbf{x}, \tilde{\mathbf{y}}))] \quad (6)$$

$$+ \mathbb{E}_{(\tilde{\mathbf{x}}, \mathbf{y}) \sim p_\phi(\mathbf{x}, \mathbf{y})} [\log(1 - \sigma(g_\omega(\tilde{\mathbf{x}}, \mathbf{y})))] .$$

The equilibrium of the objective in (6) is achieved if and only if $p_\theta(\mathbf{x}, \mathbf{y}) = p_\phi(\mathbf{x}, \mathbf{y})$.

While ALI is able to match joint distributions using (6), only conditional distributions $p_\theta(\mathbf{y}|\mathbf{x})$ and $p_\phi(\mathbf{x}|\mathbf{y})$ are learned, thus assuming access to (samples from) the true marginal distributions $q(\mathbf{x})$ and $q(\mathbf{y})$, respectively.

3. Adversarial Joint Distribution Learning

Below we discuss *fully* learning the joint distribution of two random variables in both supervised and unsupervised settings. By “supervised” it is meant that we have access to joint empirical data (\mathbf{x}, \mathbf{y}) , and by “unsupervised” it is meant that we have access to empirical draws of \mathbf{x} and \mathbf{y} , but not paired observations (\mathbf{x}, \mathbf{y}) from the joint distribution.

3.1. JointGAN

Since $q(\mathbf{x}, \mathbf{y}) = q(\mathbf{x})q(\mathbf{y}|\mathbf{x}) = q(\mathbf{y})q(\mathbf{x}|\mathbf{y})$, a simple way to achieve joint-distribution learning is to first learn models

for the two marginals separately, using a pair of traditional GANs, followed by training an independent ALI model for the two conditionals. However, such a two-step training procedure is suboptimal, as there is no information flow between marginals and conditionals during training. This suboptimality is demonstrated in the experiments. Additionally, a two-step learning process becomes cumbersome when considering more than two random variables.

Alternatively, we consider learning to sample from conditionals via $p_\theta(\mathbf{y}|\mathbf{x})$ and $p_\phi(\mathbf{x}|\mathbf{y})$, while also learning to sample from marginals via $p_\alpha(\mathbf{x})$ and $p_\beta(\mathbf{y})$. All model training is performed simultaneously. We term our model *JointGAN* for full GAN analysis of joint random variables.

Access to Paired Empirical Draws In this setting, we assume access to samples from the true (empirical) joint distribution $q(\mathbf{x}, \mathbf{y})$. The models we seek to learn constitute two means of approximating draws from the true distribution $q(\mathbf{x}, \mathbf{y})$, *i.e.*, $p_\alpha(\mathbf{x})p_\theta(\mathbf{y}|\mathbf{x})$ and $p_\beta(\mathbf{y})p_\phi(\mathbf{x}|\mathbf{y})$, as shown in Figure 1(c):

$$\tilde{\mathbf{x}} = f_\alpha(\epsilon_1), \quad \tilde{\mathbf{y}} = f_\theta(\tilde{\mathbf{x}}, \epsilon_2), \quad (7)$$

$$\hat{\mathbf{y}} = f_\beta(\epsilon'_1), \quad \hat{\mathbf{x}} = f_\phi(\hat{\mathbf{y}}, \epsilon'_2), \quad (8)$$

where $f_\alpha(\cdot)$, $f_\beta(\cdot)$, $f_\theta(\cdot)$ and $f_\phi(\cdot)$ are neural networks as defined previously. ϵ_1 , ϵ_2 , ϵ'_1 and ϵ'_2 are independent noise. Note that the only difference between $f_\alpha(\cdot)$ and $f_\phi(\cdot)$ is that the function $f_\phi(\cdot)$ has another conditional input \mathbf{y} when compared with $f_\alpha(\cdot)$. Therefore, in implementation, we couple the parameters α and ϕ together. Similarly, β and θ are also coupled together. Specifically, $f_\alpha(\cdot)$ and $f_\beta(\cdot)$ are implemented as

$$f_\alpha(\cdot) = f_\phi(\mathbf{0}, \cdot), \quad f_\beta(\cdot) = f_\theta(\mathbf{0}, \cdot), \quad (9)$$

where $\mathbf{0}$ is an all-zero tensor which has the same size as \mathbf{x} or \mathbf{y} . As a result, (7) and (8) have almost the same number of parameters as ALI-like approaches.

The following notation is introduced for simplicity of illustration:

$$\begin{aligned} p_1(\mathbf{x}, \mathbf{y}) &= q(\mathbf{x})p_\theta(\mathbf{y}|\mathbf{x}), & p_2(\mathbf{x}, \mathbf{y}) &= q(\mathbf{y})p_\phi(\mathbf{x}|\mathbf{y}) \\ p_3(\mathbf{x}, \mathbf{y}) &= p_\alpha(\mathbf{x})p_\theta(\mathbf{y}|\mathbf{x}), & p_4(\mathbf{x}, \mathbf{y}) &= p_\beta(\mathbf{y})p_\phi(\mathbf{x}|\mathbf{y}) \\ p_5(\mathbf{x}, \mathbf{y}) &= q(\mathbf{x}, \mathbf{y}), \end{aligned} \quad (10)$$

where $p_1(\mathbf{x}, \mathbf{y})$, $p_2(\mathbf{x}, \mathbf{y})$, $p_3(\mathbf{x}, \mathbf{y})$ and $p_4(\mathbf{x}, \mathbf{y})$ are implicitly defined in (4), (5), (7) and (8). Note that $p_5(\mathbf{x}, \mathbf{y})$ is simply the empirical joint distribution.

When learning, we wish to impose that the five distributions in (10) should be identical. Toward this end, an adversarial objective is specified. Joint pairs (\mathbf{x}, \mathbf{y}) are drawn from the five distributions in (10), and a critic function is learned to discriminate among them, while the four generators are

trained to mislead the critic. Naively, for JointGAN, one can use 4 binary critics to mimic a 5-class classifier. Departing from previous work such as Gan et al. (2017a), here the discriminator is implemented directly as a 5-way softmax classifier. Compared with using multiple binary classifiers, this design is more principled in that we avoid multiple critics resulting in possibly conflicting (real vs. synthetic) assessments.

Let the critic $g_\omega(\mathbf{x}, \mathbf{y}) \in \Delta^4$ (in the 4-simplex) be a ω -parameterized neural network with softmax on the top layer, i.e., $\sum_{k=1}^5 g_\omega(\mathbf{x}, \mathbf{y})[k] = 1$ and $g_\omega(\mathbf{x}, \mathbf{y})[k] \in (0, 1)$, where $g_\omega(\mathbf{x}, \mathbf{y})[k]$ is an entry of $g_\omega(\mathbf{x}, \mathbf{y})$. The minimax objective for JointGAN, $\mathcal{L}_{\text{JointGAN}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\omega})$, is given by

$$\begin{aligned} \min_{\boldsymbol{\theta}, \boldsymbol{\phi}} \max_{\boldsymbol{\omega}} \mathcal{L}_{\text{JointGAN}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\omega}) & \quad (11) \\ &= \sum_{k=1}^5 \mathbb{E}_{p_k(\mathbf{x}, \mathbf{y})} [\log g_\omega(\mathbf{x}, \mathbf{y})[k]]. \end{aligned}$$

The above objective (11) has taken into consideration the model design that $\boldsymbol{\alpha}$ and $\boldsymbol{\phi}$ are coupled together, with the same for $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$; thus $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are not present in (11). Note that expectation $\mathbb{E}_{p_5(\mathbf{x}, \mathbf{y})}[\cdot]$ is approximated using *empirical* joint samples, expectations $\mathbb{E}_{p_3(\mathbf{x}, \mathbf{y})}[\cdot]$ and $\mathbb{E}_{p_4(\mathbf{x}, \mathbf{y})}[\cdot]$ are both approximated with *purely synthesized* joint samples, while $\mathbb{E}_{p_1(\mathbf{x}, \mathbf{y})}[\cdot]$ and $\mathbb{E}_{p_2(\mathbf{x}, \mathbf{y})}[\cdot]$ are approximated using *conditionally synthesized* samples, given samples from the empirical marginals. The following proposition characterizes the solution of (11) in terms of the joint distributions.

Proposition 1 *The equilibrium for the minimax objective in (11) is achieved if and only if $p_1(\mathbf{x}, \mathbf{y}) = p_2(\mathbf{x}, \mathbf{y}) = p_3(\mathbf{x}, \mathbf{y}) = p_4(\mathbf{x}, \mathbf{y}) = p_5(\mathbf{x}, \mathbf{y})$.*

The proof is provided in Appendix A.

No Access to Paired Empirical Draws When paired data samples are not available, we do not have access to draws from $p_5(\mathbf{x}, \mathbf{y}) = q(\mathbf{x}, \mathbf{y})$, so this term is not considered in (11). Instead, we wish to impose ‘‘cycle consistency’’ (Zhu et al., 2017), i.e., $q(\mathbf{x}) \rightarrow \mathbf{x} \rightarrow p_\theta(\mathbf{y}|\mathbf{x}) \rightarrow \mathbf{y} \rightarrow p_\phi(\mathbf{x}|\mathbf{y}) \rightarrow \hat{\mathbf{x}}$ yields small $\|\mathbf{x} - \hat{\mathbf{x}}\|$, for an appropriate norm. Similarly, we impose $q(\mathbf{y}) \rightarrow \mathbf{y} \rightarrow p_\phi(\mathbf{x}|\mathbf{y}) \rightarrow \mathbf{x} \rightarrow p_\theta(\mathbf{y}|\mathbf{x}) \rightarrow \hat{\mathbf{y}}$ resulting in small $\|\mathbf{y} - \hat{\mathbf{y}}\|$.

In this case, the discriminator becomes a 4-class classifier. Let $g'_\omega(\mathbf{x}, \mathbf{y}) \in \Delta^3$ denote a new critic, with softmax on the top layer, i.e., $\sum_{k=1}^4 g'_\omega(\mathbf{x}, \mathbf{y})[k] = 1$ and $g'_\omega(\mathbf{x}, \mathbf{y})[k] \in (0, 1)$. To encourage cycle consistency, we modify the objective in (11) as

$$\begin{aligned} \min_{\boldsymbol{\theta}, \boldsymbol{\phi}} \max_{\boldsymbol{\omega}} \mathcal{L}_{\text{JointGAN}}(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\omega}) & \quad (12) \\ &= \sum_{k=1}^4 \mathbb{E}_{p_k(\mathbf{x}, \mathbf{y})} [\log g'_\omega(\mathbf{x}, \mathbf{y})[k]] + R_{\boldsymbol{\theta}, \boldsymbol{\phi}}(\mathbf{x}, \mathbf{y}), \end{aligned}$$

where $R_{\boldsymbol{\theta}, \boldsymbol{\phi}}(\mathbf{x}, \mathbf{y})$ in (12) is a cycle-consistency regulariza-




Figure 2. Generative model for the tuple $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ manifested via $p_\alpha(\mathbf{x})p_\nu(\mathbf{y}|\mathbf{x})p_\gamma(\mathbf{z}|\mathbf{x}, \mathbf{y})$. Note the skip connection that naturally arises due to the form of the generative model.

tion term specified as

$$\begin{aligned} R_{\boldsymbol{\theta}, \boldsymbol{\phi}}(\mathbf{x}, \mathbf{y}) &= \mathbb{E}_{\mathbf{x} \sim q(\mathbf{x}), \mathbf{y} \sim p_\theta(\mathbf{y}|\mathbf{x}), \hat{\mathbf{x}} \sim p_\phi(\mathbf{x}|\mathbf{y})} \|\mathbf{x} - \hat{\mathbf{x}}\| \\ &+ \mathbb{E}_{\mathbf{y} \sim q(\mathbf{y}), \mathbf{x} \sim p_\phi(\mathbf{x}|\mathbf{y}), \hat{\mathbf{y}} \sim p_\theta(\mathbf{y}|\mathbf{x})} \|\mathbf{y} - \hat{\mathbf{y}}\|. \end{aligned}$$

3.2. Extension to multiple domains

The above formulation may be extended to the case of three or more joint random variables. However, for m random variables, there are $m!$ different ways in which the joint distribution can be factorized. For example, for joint random variables $(\mathbf{x}, \mathbf{y}, \mathbf{z})$, there are possibly six different forms of the model. One must have access to all the six instantiations of these models, if the goal is to be able to generate (impute) samples from all conditionals. However, not all modeled forms of $p(\mathbf{x}, \mathbf{y}, \mathbf{z})$ need to be considered, if there is not interest in the corresponding form of the conditional. Below, we consider two specific forms of the model:

$$p(\mathbf{x}, \mathbf{y}, \mathbf{z}) = p_\alpha(\mathbf{x})p_\nu(\mathbf{y}|\mathbf{x})p_\gamma(\mathbf{z}|\mathbf{x}, \mathbf{y}) \quad (13)$$

$$= p_\beta(\mathbf{z})p_\psi(\mathbf{z}|\mathbf{y})p_\eta(\mathbf{x}|\mathbf{y}, \mathbf{z}). \quad (14)$$

Typically, the joint draws from $q(\mathbf{x}, \mathbf{y}, \mathbf{z})$ may not be easy to access; therefore, we assume that only empirical draws from $q(\mathbf{x}, \mathbf{y})$ and $q(\mathbf{y}, \mathbf{z})$ are available. For the purpose of adversarial learning, we let the critic be a 6-class softmax classifier that aims to distinguish samples from the following 6 distributions:

$$\begin{aligned} p_\alpha(\mathbf{x})p_\nu(\mathbf{y}|\mathbf{x})p_\gamma(\mathbf{z}|\mathbf{x}, \mathbf{y}), & \quad p_\beta(\mathbf{z})p_\psi(\mathbf{z}|\mathbf{y})p_\eta(\mathbf{x}|\mathbf{y}, \mathbf{z}) \\ q(\mathbf{x})p_\nu(\mathbf{y}|\mathbf{x})p_\gamma(\mathbf{z}|\mathbf{x}, \mathbf{y}), & \quad q(\mathbf{x})p_\psi(\mathbf{z}|\mathbf{y})p_\eta(\mathbf{x}|\mathbf{y}, \mathbf{z}) \\ q(\mathbf{x}, \mathbf{y})p_\gamma(\mathbf{z}|\mathbf{x}, \mathbf{y}), & \quad q(\mathbf{y}, \mathbf{z})p_\eta(\mathbf{x}|\mathbf{y}, \mathbf{z}). \end{aligned}$$

After training, one may synthesize $(\mathbf{x}, \mathbf{y}, \mathbf{z})$, impute (\mathbf{y}, \mathbf{z}) from observed \mathbf{x} , or impute \mathbf{z} from (\mathbf{x}, \mathbf{y}) , etc. Examples of this learning paradigm is demonstrated in the experiments. Interestingly, when implementing a sampling-based method for the above models, skip connections are manifested naturally as a result of the partitioning of the joint distribution,

e.g., via $p_\gamma(z|\mathbf{x}, \mathbf{y})$ and $p_\eta(\mathbf{x}|\mathbf{y}, \mathbf{z})$. This is illustrated in Figure 2, for $p_\alpha(\mathbf{x})p_\nu(\mathbf{y}|\mathbf{x})p_\gamma(\mathbf{z}|\mathbf{x}, \mathbf{y})$.

4. Related work

Adversarial methods for joint distribution learning can be roughly divided into two categories, depending on the application: (i) *generation and inference* if one of the domains consists of (stochastic) latent variables, and (ii) *conditional data synthesis* if both domains consists of observed pairs of random variables. Below, we review related work from these two perspectives.

Generation and inference The joint distribution of data and latent variables or *codes* can be considered in two (symmetric) forms: (i) from observed data samples fed through the encoder to yield codes, i.e., *inference*, and (ii) from codes drawn from a simple prior and propagated through the decoder to manifest data samples, i.e., *generation*. ALI (Dumoulin et al., 2017) and BiGAN (Donahue et al., 2017) proposed fully adversarial methods for this purpose. There are also many recent works concerned with integrating variational autoencoder (VAE) (Kingma & Welling, 2013; Pu et al., 2016) and GAN concepts for improved data generation and latent code inference (Hu et al., 2017). Representative work includes the AAE (Makhzani et al., 2015), VAE-GAN (Larsen et al., 2015), AVB (Mescheder et al., 2017), AS-VAE (Pu et al., 2017), SVAE (Chen et al., 2018), etc.

Conditional data synthesis Conditional GAN can be readily used for conditional-data synthesis if paired data are available. Multiple conditional GANs have been proposed to generate the images based on class labels (Mirza & Osindero, 2014), attributes (Perarnau et al., 2016), text (Reed et al., 2016; Xu et al., 2017) and other images (Isola et al., 2017). Often, only the mapping from one direction (a single conditional) is learned. Triangle GAN (Gan et al., 2017a) and Triple GAN (Li et al., 2017b) can be used to learn bi-directional mappings (both conditionals) in a semi-supervised learning setup. Unsupervised learning methods were also developed for this task. CycleGAN (Zhu et al., 2017) proposed to use two generators to model the conditionals and two critics to decide whether a generated sample is synthesized, in each individual domain. Further, additional reconstruction losses were introduced to impose cycle consistency. Similar work includes DiscoGAN (Kim et al., 2017), DualGAN (Yi et al., 2017) and UNIT (Liu et al., 2017).

CoGAN (Liu & Tuzel, 2016) can be used to achieve joint distribution learning. However, the joint distribution is only roughly approximated by the marginals, via sharing low-layer weights of the generators, hence not learning the true (empirical) joint distributions in a principled way.

All the other previously proposed models focus on learning

to sample from the conditionals given samples from one of the true (empirical) marginals, while the proposed model, to the best of the authors’ knowledge, is the first attempt to learn a full joint distribution of two or more observed random variables. Moreover, this paper presents the first consolidation of multiple binary critics into a single unified softmax-based critic.

We observe that the proposed model, JointGAN, may follow naturally in concept from GAN (Goodfellow et al., 2014) and ALI (Donahue et al., 2017; Dumoulin et al., 2017). However, there are several keys to obtaining good performance. Specifically, (i) the condition distribution setup naturally yields skip connections in the architecture. (ii) Compared with using multiple binary critics, the softmax-based critic can be considered as sharing the parameters among all the binary critics except the top layer. This also imposes the critic embedding the generated samples from different ways into a common latent space and reduces the number of parameters. (iii) The weight-sharing constraint among generators enforces that synthesized images from the marginal and conditional generator share a common latent space, and also further reduces the number of parameters in the network.

5. Experiments

Adam (Kingma & Ba, 2014) with learning rate 0.0002 is utilized for optimization of the JointGAN objectives. All noise vectors ϵ_1 , ϵ_2 , ϵ'_1 and ϵ'_2 are drawn from a $\mathcal{N}(0, \mathbf{I})$ distribution, with the dimension of each set to 100. Besides the results presented in this section, more results can be found in Appendix C.2. The code can be found at <https://github.com/sdai654416/Joint-GAN>.

5.1. Joint modeling multi-domain images

Datasets We present results on five datasets: edges \leftrightarrow shoes (Yu & Grauman, 2014), edges \leftrightarrow handbags (Zhu et al., 2016), Google maps \leftrightarrow aerial photos (Isola et al., 2017), labels \leftrightarrow facades (Tyleček & Šára, 2013) and labels \leftrightarrow cityscapes (Cordts et al., 2016). All of these datasets are two-domain image pairs.

For three-domain images, we create a new dataset by combining labels \leftrightarrow facades pairs and labels \leftrightarrow cityscapes pairs into facades \leftrightarrow labels \leftrightarrow cityscapes tuples. In this dataset, only empirical draws from $q(\mathbf{x}, \mathbf{y})$ and $q(\mathbf{y}, \mathbf{z})$ are available. Another new dataset is created based on MNIST, where the three image domains are the MNIST images, clockwise transposed ones, and anticlockwise transposed ones.

Baseline As described in Sec. 3.1, a two-step model is implemented as the baseline. Specifically, WGAN-GP (Gulrajani et al., 2017) is employed to model the two marginals; Pix2pix (Isola et al., 2017) and CycleGAN (Zhu et al., 2017) are utilized to model the conditionals for the case with and

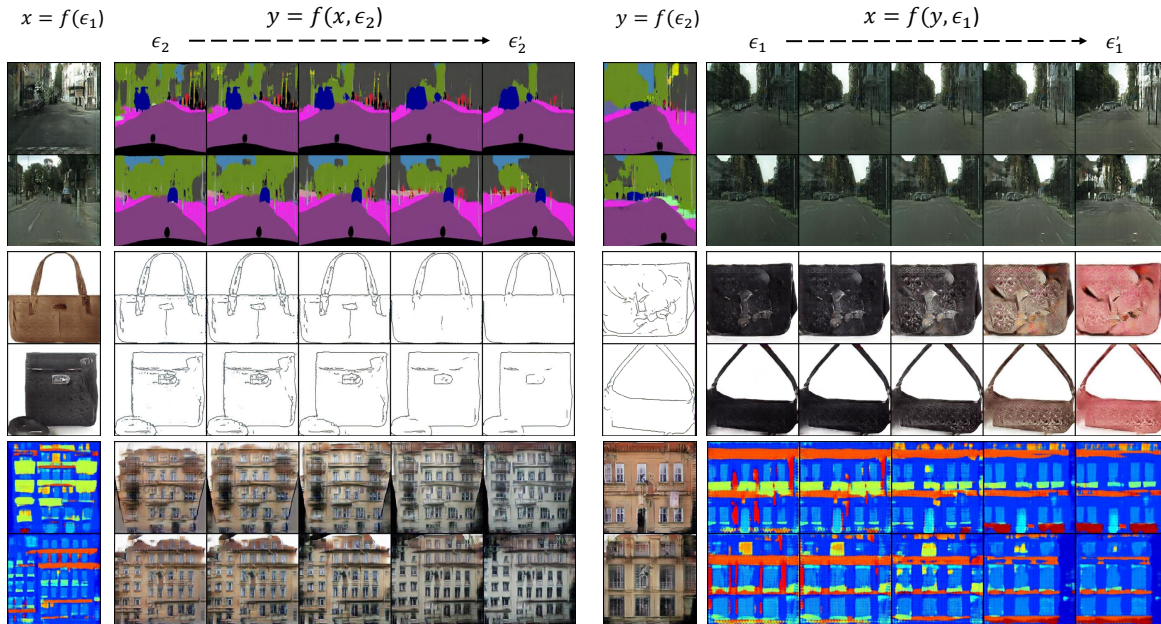


Figure 3. Generated paired samples from models trained on paired data.

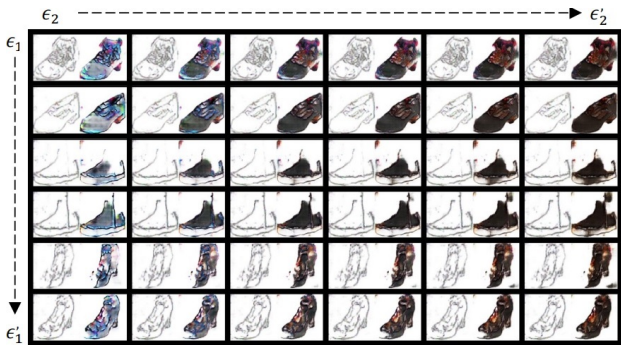


Figure 4. Generated paired samples from the JointGAN model trained on the paired edges \leftrightarrow shoes dataset.

without access to paired empirical draws, respectively.

Network Architectures For generators, we employed the U-net (Ronneberger et al., 2015) which has been demonstrated to achieve impressive results for image-to-image translation. Following Isola et al. (2017), PatchGAN is employed for the discriminator, which provides real vs. synthesized prediction on 70×70 overlapping image patches.

5.1.1. QUALITATIVE RESULTS

Figures 3 and 4 show the results trained on paired data. All the image pairs are generated from random noise. For Figure 4, we first draw (ϵ_1, ϵ_2) and $(\epsilon'_1, \epsilon'_2)$ to generate the top-left image pairs and bottom-right image pairs according to (7). All remaining image pairs are generated from the noise pair made by linear interpolation between ϵ_1 and ϵ'_1 , and between ϵ_2 and ϵ'_2 , respectively, also via (7). For Figure 3, in each row of the left block, the column is first

generated from $p_\alpha(x)$, and then the images of the right part are generated based on the leftmost image and an additional noise vector linear-interpolated between two random points ϵ_2 and ϵ'_2 . The images in the right block are produced in a similar way.

These results demonstrate that our model is able to generate both realistic and highly coherent image pairs. In addition, the interpolation experiments illustrate that our model maintains smooth transitions in the latent space, with each point in the latent space corresponding to a plausible image. For example, in the edges \leftrightarrow handbags dataset, it can be seen that the edges smoothly transforming from complicated structures into simple ones, and the color of the handbags transforming from black to red. The quality of images generated from the baseline is much worse than ours, and are provided in Appendix C.1.

Figure 5 shows the generated samples trained on unpaired data. Our model is able to produce image pairs whose quality are close to the samples trained on paired data.

Figures 6 and 7 show the generated samples from models trained on three-domain images. The generated images in each tuple are highly correlated. Interestingly, in Figure 7, the synthesized labels strive to be consistent with both the generated street scene and facade photos.

5.1.2. QUANTITATIVE RESULTS

We perform a detailed quantitative analysis on the two-domain image-pair task.

Human Evaluation We perform human evaluation using Amazon Mechanical Turk (AMT), and present human eval-



Figure 5. Generated paired samples from models trained on unpaired data.



Figure 6. Generated three-domain samples from models trained on MNIST. The images in each tuple of the left three columns are sequentially generated from left to right, while those in the right three columns are generated from right to left.

uation results on the relevance and realism of generated pairs in both the cases with or without access to paired empirical samples. In each survey, we compare JointGAN and the two-step baseline by taking a random sample of 100 generated image pairs (5 datasets, 20 samples on each dataset), and ask the human evaluator to select which sample is more realistic and the content of which pairs are more relevant. We obtained roughly 44 responses per data sample (4378 samples in total) and the results are shown in Table 1. Clearly, human analysis suggest that our JointGAN produces higher-quality samples when compared with the

Table 1. Human evaluation results on the quality of generated pairs.

Method	Realism	Relevance
<i>Trained with paired data</i>		
WGAN-GP + Pix2pix wins	2.32%	3.1%
JointGAN wins	17.93%	36.32%
Not distinguishable	79.75%	60.58%
<i>Trained with unpaired data</i>		
WGAN-GP + CycleGAN wins	0.13%	1.31%
JointGAN wins	81.55%	40.87%
Not distinguishable	18.32%	57.82%

two-step baseline, verifying the effectiveness of learning the marginal and conditional simultaneously.

Relevance Score We use relevance score to evaluate the quality and relevance of two generated images. The relevance score is calculated as the cosine similarity between two images that are embedded into a shared latent space, which are learned via training a ranking model (Huang et al., 2013). Details are provided in Appendix B. The final relevance score is the average over all the individual relevance scores on each pair. Results are summarized in Table 2. Our JointGAN provides significantly better results than the two-step baselines, especially when we do not have access to the paired empirical samples.

Besides the results of our model and baselines, we also present results on three types of real images: (i) *True pairs*: this is the real image pairs from the same dataset but not used for training the ranking model; (ii) *Random pairs*: the images are from the same dataset but the content of two images are not correlated; (iii) *Other pairs*: the images are correlated but sampled from a dataset different from the training set. We can see in Table 2 that the first one obtains a high relevance score while the latter two have a very low score, which shows that the relevance score metric assigns a low value when either the content of generated image pairs is not correlated or the images are not plausibly like the training set. It demonstrates that this metric correlates well with the quality of generated image pairs.

5.2. Joint modeling caption features and images

Setup Our model is next evaluated on the Caltech-UCSD Birds dataset (Welinder et al., 2010), in which each image of bird is paired with 10 different captions. Since generating realistic text using GAN itself is a challenging task, in this work, we train our model on pairs of caption *features* and images. The caption features are obtained from a pretrained word-level CNN-LSTM autoencoder (Gan et al., 2017b), which aims to achieve a one-to-one mapping between the captions and the features. We then train JointGAN based on the caption features and their corresponding images (the paired data for training JointGAN use CNN-generated text

Table 2. Relevance scores of the generated pairs on the five two-domain image datasets.

	edges↔shoes	edges↔handbags	labels↔cityscapes	labels↔facades	maps↔satellites
True pairs	0.684	0.672	0.591	0.529	0.514
Random pairs	0.008	0.005	0.012	0.011	0.054
Other pairs	0.113	0.139	0.092	0.076	0.081
WGAN-GP + Pix2pix	0.352	0.343	0.301	0.288	0.125
JointGAN (paired)	0.488	0.489	0.377	0.364	0.328
WGAN-GP + CycleGAN	0.203	0.195	0.201	0.139	0.091
JointGAN (unpaired)	0.452	0.461	0.339	0.341	0.299




Figure 7. Generated paired samples from models trained on facades↔labels↔cityscapes. The images in the first row are sequentially generated from left to right while those in the second row are generated from right to left. Both generation starts from a noise vector.

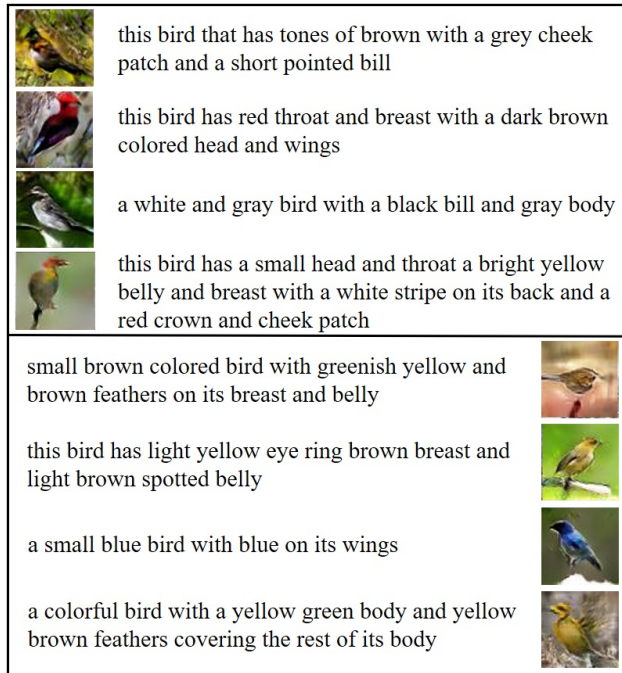


Figure 8. Generated paired samples of caption features and images (all data synthesized). Top block: from generated images to caption features. Bottom block: from generated caption features to images.

features, which avoids issues of training GAN for text generation). Finally to visualize the results, we use the pretrained LSTM decoder to decode the generated features back to captions. We employ StackGAN-stage-I (Zhang et al., 2017a) for generating images from caption features while a CNN is

utilized to generate caption features from images. Details are provided in Appendix D.

Qualitative Results Figure 8 shows the qualitative results of JointGAN: (i) generate images from noise and then conditionally generate caption features, and (ii) generate caption features from noise and then conditionally generate images. The results show high-quality and diverse image generation, and strong coherent relationship between each pair of the caption feature and image. It demonstrates the robustness of our model, in that it not only generates realistic multi-domain images but also handles well different datasets such as caption feature and image pairs.

6. Conclusion

We propose JointGAN, a new framework for multi-domain joint distribution learning. The joint distribution is learned via decomposing it into the product of a marginal and a conditional distribution(s), each learned via adversarial training. JointGAN allows interesting applications since it provides freedom to draw samples from various marginalized or conditional distributions. We consider joint analysis of two and three domains, and demonstrate that JointGAN achieves significantly better results than a two-step baseline model, both qualitatively and quantitatively.

Acknowledgements

This research was supported in part by DARPA, DOE, NIH, ONR and NSF.

References

- Chen, L., Dai, S., Pu, Y., Zhou, E., Li, C., Su, Q., Chen, C., and Carin, L. Symmetric variational autoencoder and connections to adversarial learning. In *AISTATS*, 2018.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., and Schiele, B. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016.
- Donahue, J., Krähenbühl, P., and Darrell, T. Adversarial feature learning. In *ICLR*, 2017.
- Dumoulin, V., Belghazi, I., Poole, B., Mastropietro, O., Lamb, A., Arjovsky, M., and Courville, A. Adversarially learned inference. In *ICLR*, 2017.
- Gan, Z., Chen, L., Wang, W., Pu, Y., Zhang, Y., Liu, H., Li, C., and Carin, L. Triangle generative adversarial networks. In *NIPS*, 2017a.
- Gan, Z., Pu, Y., Henao, R., Li, C., He, X., and Carin, L. Learning generic sentence representations using convolutional neural networks. In *EMNLP*, 2017b.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *NIPS*, 2014.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of Wasserstein GANs. In *NIPS*, 2017.
- Hu, Z., Yang, Z., Salakhutdinov, R., and Xing, E. P. On unifying deep generative models. *arXiv preprint arXiv:1706.00550*, 2017.
- Huang, P.-S., He, X., Gao, J., Deng, L., Acero, A., and Heck, L. Learning deep structured semantic models for web search using clickthrough data. In *CIKM*, 2013.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- Kim, T., Cha, M., Kim, H., Lee, J., and Kim, J. Learning to discover cross-domain relations with generative adversarial networks. In *ICML*, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Larsen, A. B. L., Sønderby, S. K., Larochelle, H., and Winther, O. Autoencoding beyond pixels using a learned similarity metric. *arXiv preprint arXiv:1512.09300*, 2015.
- Li, C., Liu, H., Chen, C., Pu, Y., Chen, L., Henao, R., and Carin, L. Alice: Towards understanding adversarial learning for joint distribution matching. In *NIPS*, 2017a.
- Li, C., Xu, K., Zhu, J., and Zhang, B. Triple generative adversarial nets. In *NIPS*, 2017b.
- Liu, M.-Y. and Tuzel, O. Coupled generative adversarial networks. In *NIPS*, 2016.
- Liu, M.-Y., Breuel, T., and Kautz, J. Unsupervised image-to-image translation networks. In *NIPS*, 2017.
- Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I., and Frey, B. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- Mescheder, L., Nowozin, S., and Geiger, A. Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks. In *ICML*, 2017.
- Mirza, M. and Osindero, S. Conditional generative adversarial nets. In *arXiv preprint arXiv:1411.1784*, 2014.
- Perarnau, G., van de Weijer, J., Raducanu, B., and Álvarez, J. M. Invertible conditional gans for image editing. *arXiv preprint arXiv:1611.06355*, 2016.
- Pu, Y., Gan, Z., Henao, R., Yuan, X., Li, C., Stevens, A., and Carin, L. Variational autoencoder for deep learning of images, labels and captions. In *NIPS*, 2016.
- Pu, Y., Wang, W., Henao, R., Chen, L., Gan, Z., Li, C., and Carin, L. Adversarial symmetric variational autoencoder. In *NIPS*, 2017.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *ICLR*, 2016.
- Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., and Lee, H. Generative adversarial text to image synthesis. In *ICML*, 2016.
- Ronneberger, O., Fischer, P., and Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- Tao, C., Chen, L., Henao, R., Feng, J., and Carin, L. Chi-square generative adversarial network. In *ICML*, 2018.
- Tyleček, R. and Šára, R. Spatial pattern templates for recognition of objects with regular structure. In *GCPR*, 2013.
- Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., and Perona, P. Caltech-UCSD Birds 200. Technical report, California Institute of Technology, 2010.

- Xu, T., Zhang, P., Huang, Q., Zhang, H., Gan, Z., Huang, X., and He, X. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. *arXiv preprint arXiv:1711.10485*, 2017.
- Yi, Z., Zhang, H., Tan, P., and Gong, M. Dualgan: Unsupervised dual learning for image-to-image translation. In *CVPR*, 2017.
- Yu, A. and Grauman, K. Fine-grained visual comparisons with local learning. In *CVPR*, 2014.
- Yu, L., Zhang, W., Wang, J., and Yu, Y. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*, 2017.
- Zhang, H., Xu, T., Li, H., Zhang, S., and Metaxas, D. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, 2017a.
- Zhang, Y., Gan, Z., Fan, K., Chen, Z., Hénao, R., Shen, D., and Carin, L. Adversarial feature matching for text generation. In *ICML*, 2017b.
- Zhu, J.-Y., Krähenbühl, P., Shechtman, E., and Efros, A. A. Generative visual manipulation on the natural image manifold. In *ECCV*, 2016.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *CVPR*, 2017.