
Fork or Fail: Cycle-Consistent Training with Many-to-One Mappings

Qipeng Guo^{1,2}
Weinan Zhang⁵

Zhijing Jin^{1,3}
Jun Zhu⁴

Ziyu Wang⁴
Zheng Zhang¹

Xipeng Qiu²
David Wipf¹

¹Amazon Shanghai AI Lab ²Fudan University ³Max Planck Institute for Intelligent Systems, Tübingen
⁴Tsinghua University ⁵Shanghai Jiao Tong University

Abstract

Cycle-consistent training is widely used for jointly learning a forward and inverse mapping between two domains of interest without the cumbersome requirement of collecting matched pairs within each domain. In this regard, the implicit assumption is that there exists (at least approximately) a ground-truth bijection such that a given input from either domain can be accurately reconstructed from successive application of the respective mappings. But in many applications no such bijection can be expected to exist and large reconstruction errors can compromise the success of cycle-consistent training. As one important instance of this limitation, we consider practically-relevant situations where there exists a many-to-one or surjective mapping between domains. To address this regime, we develop a conditional variational autoencoder (CVAE) approach that can be viewed as converting surjective mappings to implicit bijections whereby reconstruction errors in both directions can be minimized, and as a natural byproduct, realistic output diversity can be obtained in the one-to-many direction. As theoretical motivation, we analyze a simplified scenario whereby minima of the proposed CVAE-based energy function align with the recovery of ground-truth surjective mappings. On the empirical side, we consider a synthetic image dataset with known ground-truth, as well as a real-world application involving natural language generation from knowledge graphs and vice versa, a pro-

totypical surjective case. For the latter, our CVAE pipeline can capture such many-to-one mappings during cycle training while promoting textural diversity for graph-to-text tasks.¹

1 Introduction

Given data $\mathbf{x} \in \mathcal{X}$ from domain \mathcal{X} and $\mathbf{y} \in \mathcal{Y}$ from domain \mathcal{Y} , it is often desirable to learn bidirectional mappings $f : \mathcal{Y} \rightarrow \mathcal{X}$ and $g : \mathcal{X} \rightarrow \mathcal{Y}$ such that for matched pairs $\{\mathbf{x}, \mathbf{y}\}$, we have that $\mathbf{x} \approx \hat{\mathbf{x}} \triangleq f(\mathbf{y})$ and $\mathbf{y} \approx \hat{\mathbf{y}} \triangleq g(\mathbf{x})$. When provided with a corpus of suitably aligned data, this amounts to a straightforward supervised learning problem. However, in many applications spanning computer vision (Zhu et al., 2017a), natural language processing (Lample et al., 2018; Artetxe et al., 2018) and speech recognition (Hori et al., 2019), we may only have access to individual samples from \mathcal{X} and \mathcal{Y} but limited or no labeled ground-truth matches between domains, since, for example, the labeling process may be prohibitively expensive. To address this commonly-encountered situation, cycle-consistent training represents an unsupervised means of jointly learning f and g by penalizing the cycle-consistency reconstruction losses $\|\mathbf{x} - f[g(\mathbf{x})]\|$ and $\|\mathbf{y} - g[f(\mathbf{y})]\|$ using non-parallel samples from \mathcal{X} and \mathcal{Y} and some norm or distance metric $\|\cdot\|$ (Zhu et al., 2017a).

However, this process implicitly assumes that there exists a suitable bijection between domains (implying $f = g^{-1}$ and $g = f^{-1}$), an assumption that frequently does not hold for practical applications of cycle-consistent training. As a representative example related to natural language understanding, each \mathbf{x} may represent a text segment while \mathbf{y} corresponds with the underlying knowledge graph describing the text content. The relationship between these domains is

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

¹Our code is available at <https://github.com/QipengGuo/CycleGT>. And for an extended version of this work, please see Guo et al. (2021).

surjective, but not bijective, in the sense that multiple sentences with equivalent meaning but different syntactic structure can be mapped to the same knowledge graph. Hence if we follow any possible learned mapping $\mathbf{x} \rightarrow \hat{\mathbf{y}} \rightarrow \hat{\mathbf{x}}$, there will often be significant error between \mathbf{x} and the reconstructed $\hat{\mathbf{x}}$. In other words, no invertible transformation exists between domains and there will necessarily be information about \mathbf{x} that is lost when we map through \mathcal{Y} space. Additionally, deterministic mappings do not reflect the ground-truth conditional distribution $p_{gt}(\mathbf{x}|\mathbf{y})$, which is necessary for the generation of diverse text consistent with a given knowledge graph.

Despite these limitations, there has been relatively little effort or rigorous analysis devoted to explicitly addressing the lack of a bijection in applications of cycle-consistent training; Section 2 on related work will discuss this point in greater detail. As a step towards filling this void, in Section 3 we will consider replacing the typical deterministic cycle training pipeline with a stochastic model reflecting $p_{gt}(\mathbf{x}|\mathbf{y})$ and $p_{gt}(\mathbf{y}|\mathbf{x})$ for the $\mathbf{x} \rightarrow \hat{\mathbf{y}} \rightarrow \hat{\mathbf{x}}$ and $\mathbf{y} \rightarrow \hat{\mathbf{x}} \rightarrow \hat{\mathbf{y}}$ cycles respectively. In doing so, we apply a conditional variational autoencoder (CVAE) formulation (Doersch, 2016; Sohn et al., 2015) to deal with the intractable integrals that arise. Note that although the proposed CVAE methodology can be generalized, we will herein restrict ourselves to situations where there exists a many-to-one mapping from \mathbf{x} to \mathbf{y} (i.e., a surjection) as originally motivated by our interest in conversions between knowledge graphs and diverse, natural language text.

Proceeding further, Section 4 provides theoretical support by analyzing a simplified scenario whereby minima of the proposed CVAE-based energy function align with the recovery of ground-truth surjective mappings. To the best of our knowledge, this is the only demonstration of a cycle-consistent model with any type of performance guarantee within a non-trivial, non-bijective context. We then turn to empirical validation in Section 5 that corroborates our theory via a synthetic image example and demonstrates real-world practicality on an application involving the conversion between diverse natural language and knowledge graphs taken from the WebNLG dataset. Overall, experimental results indicate that our proposed CVAE pipeline can approximate surjective mappings during cycle training, with performance on par with supervised alternatives, while promoting diversity for the graph-to-text direction.

2 Related Work

General Cycle-Consistent Training The concept of leveraging the transitivity of two functions that serve as inverses to one another has been applied to a variety of tasks. For example, in computer vision, forward-

backward consistency has been used extensively in computer vision (Kalal et al., 2010; Sundaram et al., 2010), and cycle-consistent training pipelines underlie image style transfer (Zhu et al., 2017a; Liu et al., 2017), depth estimation (Godard et al., 2017), and unsupervised domain adaptation (Hoffman et al., 2018) pipelines. Turning to natural language processing (NLP), back translation (Sennrich et al., 2016; Edunov et al., 2018; Jin et al., 2020) and dual learning (Cheng et al., 2016; He et al., 2016) have been widely deployed for unsupervised machine translation. Similar techniques have also contributed to applications such as language style transfer (Shen et al., 2017; Jin et al., 2019). However, the above models primarily rely on deterministic pipelines that implicitly assume a bijection even if one does not actually exist.

And finally, a VAE-inspired model for converting between knowledge graphs and text is considered in Tseng et al. (2020). But again there is no explicit accounting for non-bijective data as a shared latent space is assumed to contain all information from *both* \mathbf{x} and \mathbf{y} domains, and the proposed model is designed and tested for semi-supervised learning (not fully unsupervised cycle-consistency). Moreover, for tractable inference, some terms from the variational bound on the log-likelihood (central to all VAE models) are heuristically removed; hence the relationship with the original, motivating probabilistic model remains unclear.

Non-Bijective Mappings Non-bijective mappings are investigated in applications such as multi-domain image-to-image translation (Choi et al., 2018), voice conversion (Kameoka et al., 2018), multi-attribute text style transfer (Lample et al., 2019), music transfer (Bitton et al., 2018), and multi-modal generation (Shi et al., 2019). Most of this work uses adversarial neural networks, or separate decoders (Lee et al., 2019; Mor et al., 2018), and one case even applies a CVAE model (Jha et al., 2018). However, all the above assume multiple *pre-defined* style domains and require data be clearly separated a priori according to these domains to train non-bijective mappings. In contrast, our proposed model assumes a completely arbitrary surjective mapping that can be learned from the data without such additional domain-specific side information pertaining to styles or related (so ours can fit unknown styles mixed within an arbitrary dataset). One exception is Zhu et al. (2017b), which handles general non-bijective image mappings using a hybrid VAE-GAN model but unlike our approach, it requires matched $\{\mathbf{x}, \mathbf{y}\}$ pairs for training.

3 Model Development

We will first present a stochastic alternative to deterministic cycle-consistency that, while useful in princi-

ple for handling surjective (but explicitly non-bijective) mappings, affords us with no practically-realizable inference procedure. To mitigate this shortcoming, we then derive a tractable CVAE approximation and discuss some of its advantages. Later in Section 4 we will analyze the local and global minima of this cycle-consistent CVAE model in the special case where the decoder functions are restricted to being affine.

3.1 Stochastic Cycle-Consistent Formulation

Although the proposed methodology can be generalized, we will herein restrict ourselves to situations where there exists a many-to-one mapping from \mathbf{x} to \mathbf{y} (i.e., a surjection) and the resulting asymmetry necessitates that the $\mathbf{x} \rightarrow \hat{\mathbf{y}} \rightarrow \hat{\mathbf{x}}$ and $\mathbf{y} \rightarrow \hat{\mathbf{x}} \rightarrow \hat{\mathbf{y}}$ cycles be handled differently. In this regard, our starting point is to postulate an additional latent variable $\mathbf{u} \in \mathcal{U}$ that contributes to a surjective matched pair $\{\mathbf{x}, \mathbf{y}\}$ via

$$\mathbf{x} = h_{gt}(\mathbf{y}, \mathbf{u}) \quad \text{and} \quad \mathbf{y} = h_{gt}^+(\mathbf{x}), \quad (1)$$

where $h_{gt} : \mathcal{Y} \times \mathcal{U} \rightarrow \mathcal{X}$ and $h_{gt}^+ : \mathcal{X} \rightarrow \mathcal{Y}$ represent ground-truth mappings we would ultimately like to estimate. For this purpose we adopt the approximations $h_\theta : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{X}$ and $h_\theta^+ : \mathcal{X} \rightarrow \mathcal{Y}$ with trainable parameters θ , noting that the second input argument of h_θ is now $\mathbf{z} \in \mathcal{Z}$ instead of $\mathbf{u} \in \mathcal{U}$. This is because the latter is unobservable and it is sufficient to learn a mapping that preserves the surjection between \mathbf{x} and \mathbf{y} without necessarily reproducing the exact same functional form of h_{gt} . For example, if hypothetically $\mathbf{u} = \pi(\mathbf{z})$ in (1) for some function π , we could redefine h_{gt} as a function of \mathbf{z} without actually changing the relationship between \mathbf{x} and \mathbf{y} .

We are now prepared to define a negative conditional log-likelihood loss for both cycle directions, averaged over the distributions of \mathbf{y} and \mathbf{x} respectively. For the simpler $\mathbf{y} \rightarrow \hat{\mathbf{x}} \rightarrow \hat{\mathbf{y}}$ cycle, we define

$$\ell_y(\theta) \triangleq - \int \left(\log \int p_\theta(\mathbf{y}|\hat{\mathbf{x}}) p(\mathbf{z}) d\mathbf{z} \right) \rho_{gt}^y(d\mathbf{y}), \quad (2)$$

where $\hat{\mathbf{x}} = h_\theta(\mathbf{y}, \mathbf{z})$, $p_\theta(\mathbf{y}|\hat{\mathbf{x}})$ is determined by h_θ^+ and an appropriate domain-specific distribution, and $p(\mathbf{z})$ is assumed to be fixed and known (e.g., a standardized Gaussian). Additionally, ρ_{gt}^y denotes the ground-truth probability measure associated with \mathcal{Y} . Consequently, $\rho_{gt}^y(d\mathbf{y})$ is the measure assigned to the infinitesimal $d\mathbf{y}$, from which it obviously follows that $\int \rho_{gt}^y(d\mathbf{y}) = 1$. Note that the resulting derivations can apply even if no ground-truth density $p_{gt}(\mathbf{y})$ exists, e.g., a counting measure over training samples or discrete domains can be assumed within this representation.

Given that there should be no uncertainty in \mathbf{y} when conditioned on \mathbf{x} , we would ideally like to learn parameters whereby $p_\theta(\mathbf{y}|\hat{\mathbf{x}})$, and therefore $\ell_y(\theta)$, degenerates

while reflecting a many-to-one mapping. By this we mean that

$$\mathbf{y} \approx \hat{\mathbf{y}} = h_\theta^+(\hat{\mathbf{x}}) = h_\theta^+(h_\theta(\mathbf{y}, \mathbf{z})), \quad \forall \mathbf{z} \sim p(\mathbf{z}). \quad (3)$$

Hence the $\mathbf{y} \rightarrow \hat{\mathbf{x}} \rightarrow \hat{\mathbf{y}}$ cycle only serves to favor (near) perfect reconstructions of \mathbf{y} while ignoring any \mathbf{z} that can be drawn from $p(\mathbf{z})$. The latter stipulation is unique relative to typical deterministic cycle-consistent training, which need not learn to ignore randomness from an additional confounding latent factor.

In contrast, the $\mathbf{x} \rightarrow \hat{\mathbf{y}} \rightarrow \hat{\mathbf{x}}$ cycle operates somewhat differently, with the latent \mathbf{z} serving a more important, non-degenerate role. Similar to before, we would ideally like to minimize the negative conditional log-likelihood given by

$$\ell_x(\theta) \triangleq - \int \left(\log \int p_\theta(\mathbf{x}|\hat{\mathbf{y}}, \mathbf{z}) p(\mathbf{z}) d\mathbf{z} \right) \rho_{gt}^x(d\mathbf{x}), \quad (4)$$

where now $\hat{\mathbf{y}} = h_\theta^+(\mathbf{x})$, $p_\theta(\mathbf{x}|\hat{\mathbf{y}}, \mathbf{z})$ depends on h_θ , and ρ_{gt}^x represents the ground-truth probability measure on \mathcal{X} . Here $\hat{\mathbf{y}}$ can be viewed as an estimate of all the information pertaining to the unknown paired \mathbf{y} as preserved through the mapping h_θ^+ from \mathbf{x} to \mathcal{Y} space. Moreover, if cycle-consistent training is ultimately successful, both $\hat{\mathbf{y}}$ and \mathbf{y} should be independent of \mathbf{z} , and it should be the case that $\int p_\theta(\mathbf{x}|\hat{\mathbf{y}}, \mathbf{z}) p(\mathbf{z}) d\mathbf{z} = p_\theta(\mathbf{x}|\hat{\mathbf{y}}) \approx p_{gt}(\mathbf{x}|\mathbf{y})$.

Per these definitions, it is also immediately apparent that $p_\theta(\mathbf{x}|\hat{\mathbf{y}})$ is describing the distribution of $h_\theta(\mathbf{y}, \mathbf{z})$ conditioned on \mathbf{y} being fixed to $h_\theta^+(\mathbf{x})$. This can be viewed as a stochastic version of the typical $\mathbf{x} \rightarrow \hat{\mathbf{y}} \rightarrow \hat{\mathbf{x}}$ cycle, whereas now the latent \mathbf{z} allows us to spread probability mass across all \mathbf{x} that are consistent with a given $\hat{\mathbf{y}}$. In fact, if we set $\mathbf{z} = \mathbf{z}'$ to a fixed null value (i.e., change $p(\mathbf{z})$ to a Dirac delta function centered at some arbitrary \mathbf{z}'), we recover this traditional pipeline exactly, with $-\log p_\theta(\mathbf{x}|\mathbf{y}, \mathbf{z}')$ simply defining the implicit loss function, with limited ability to assign high probability to multiple different values of \mathbf{x} for any given \mathbf{y} .

3.2 CVAE Approximation

While $\ell_y(\theta)$ from Section 3.1 can be efficiently minimized using stochastic sampling from $p(\mathbf{z})$ to estimate the required integral, the $\ell_x(\theta)$ term is generally intractable. Note that unlike $\ell_y(\theta)$, directly sampling \mathbf{z} is not a viable solution for $\ell_x(\theta)$ since $p_\theta(\mathbf{x}|\hat{\mathbf{y}}, \mathbf{z})$ can be close to zero for nearly all values of \mathbf{z} , and therefore, a prohibitively large number of samples would be needed to obtain reasonable estimates of the integral. Fortunately though, we can form a trainable upper bound using a CVAE architecture that dramatically improves

sample efficiency (Doersch, 2016; Sohn et al., 2015). Specifically, we define

$$\ell_x(\theta, \phi) \triangleq \int \left\{ -\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\hat{\mathbf{y}}, \mathbf{z})] p(\mathbf{z}) d\mathbf{z} + \mathbb{KL}[q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})] \right\} \rho_{gt}^x(d\mathbf{x}), \quad (5)$$

where in this context, $p_\theta(\mathbf{x}|\hat{\mathbf{y}}, \mathbf{z})$ is referred to as a *decoder* distribution while $q_\phi(\mathbf{z}|\mathbf{x})$ represents a trainable *encoder* distribution parameterized by ϕ . And by design of general VAE-based models, we have that $\ell_x(\theta, \phi) \geq \ell_x(\theta)$ for all ϕ (Kingma and Welling, 2014; Rezende et al., 2014). Note that we could also choose to condition $q_\phi(\mathbf{z}|\mathbf{x})$ and $p(\mathbf{z})$ on $\hat{\mathbf{y}}$, although this is not required to form a valid or maximally-tight bound. In the case of $q_\phi(\mathbf{z}|\mathbf{x})$, $\hat{\mathbf{y}}$ is merely a function of \mathbf{x} and therefore contains no additional information beyond direct conditioning on \mathbf{x} (and the stated bound holds regardless of what we choose for q_ϕ). In contrast, $p(\mathbf{z})$ defines the assumed generative model which we are also free to choose; however, conditioning on $\hat{\mathbf{y}}$ can be absorbed into $p_\theta(\mathbf{x}|\hat{\mathbf{y}}, \mathbf{z})$ such that there is no change in the representational capacity of $p_\theta(\mathbf{x}|\hat{\mathbf{y}}) = \int p_\theta(\mathbf{x}|\hat{\mathbf{y}}, \mathbf{z}) p(\mathbf{z}) d\mathbf{z}$.

Given (5) as a surrogate for $\ell_x(\theta)$ and suitable distributional assumptions, the combined cycle-consistent loss

$$\ell_{\text{cycle}}(\theta, \phi) = \ell_x(\theta, \phi) + \ell_y(\theta) \quad (6)$$

can be minimized over $\{\theta, \phi\}$ using stochastic gradient descent and the reparameterization trick from Kingma and Welling (2014) and Rezende et al. (2014). We henceforth refer to this formulation as *CycleCVAE*. And as will be discussed in more detail later, additional constraints, regularization factors, or inductive biases can also be included to help ensure identifiability of ground-truth mappings. For example, we may consider penalizing or constraining the divergence between the distributions of \mathbf{y} and $\hat{\mathbf{y}} = h_\theta^+(\mathbf{x})$, both of which can be estimated from unpaired samples from ρ_{gt}^y and ρ_{gt}^x respectively. This is useful for disambiguating the contributions of \mathbf{y} and \mathbf{z} to \mathbf{x} and will be equal to zero (or nearly so) if $h_\theta^+ \approx h_{gt}^+$ (more on this in Section 4 below).

3.3 CycleCVAE Inference

Once trained and we have obtained some optimal CycleVAE parameters $\{\theta^*, \phi^*\} \approx \arg \min_{\theta, \phi} \ell_{\text{cycle}}(\theta, \phi)$, we can compute matches for test data in either the $\mathbf{x}_{test} \rightarrow \mathbf{y}_{test}$ or $\mathbf{y}_{test} \rightarrow \mathbf{x}_{test}$ direction. For the former, we need only compute $\hat{\mathbf{y}}_{test} = h_{\theta^*}^+(\mathbf{x}_{test})$ and there is no randomness involved. In contrast, for the other direction (one-to-many) we can effectively draw approximate samples from the posterior distribution $p_{\theta^*}(\mathbf{x}|\mathbf{y}_{test})$ by first drawing samples $\mathbf{z} \sim p(\mathbf{z})$ and then computing $\hat{\mathbf{x}}_{test} = h_{\theta^*}(\mathbf{y}_{test}, \mathbf{z})$.

3.4 CycleCVAE Advantages in Converting Surjections to Implicit Bijections

Before proceeding to a detailed theoretical analysis of CycleCVAE, it is worth examining a critical yet subtle distinction between CycleCVAE and analogous, deterministic baselines. In particular, given that the $\mathbf{x} \rightarrow \hat{\mathbf{y}} \rightarrow \hat{\mathbf{x}}$ cycle will normally introduce reconstruction errors because of the lack of a bijection as discussed previously, we could simply augment traditional deterministic pipelines with a \mathbf{z} such that $\mathbf{x} \rightarrow \{\hat{\mathbf{y}}, \mathbf{z}\} \rightarrow \hat{\mathbf{x}}$ forms a bijection. But there remain (at least) two unresolved problems. First, it is unclear how to choose the dimensionality and distribution of $\mathbf{z} \in \mathcal{Z}$ such that we can actually obtain a bijection. For example, if $\dim[\mathcal{Z}]$ is less than the unknown $\dim[\mathcal{U}]$, then the reconstruction error $\|\mathbf{x} - \hat{\mathbf{x}}\|$ can still be large, while conversely, if $\dim[\mathcal{Z}] > \dim[\mathcal{U}]$ there can now exist a many-to-one mapping from $\{\mathbf{y}, \mathbf{z}\}$ to \mathbf{x} , in which case the superfluous degrees of freedom can interfere with our ability to disambiguate the role \mathbf{y} plays in predicting \mathbf{x} . This issue, combined with the fact that we have no mechanism for choosing a suitable $p(\mathbf{z})$, implies that deterministic cycle-consistent training is difficult to instantiate.

In contrast, CycleCVAE can effectively circumvent these issues via two key mechanisms that underpin VAE-based models. First, assume that $\dim[\mathcal{X}] > \dim[\mathcal{Y}]$, meaning that \mathcal{X} represents a higher-dimensional space or manifold relative to \mathcal{Y} , consistent with our aforementioned surjective assumption. Then provided we choose $\dim[\mathcal{Z}]$ sufficiently large, e.g., $\dim[\mathcal{Z}] \geq \dim[\mathcal{U}]$, we would ideally prefer that CVAE regularization somehow prune away the superfluous dimensions of \mathbf{z} that are not required to produce good reconstructions of \mathbf{x} .² For example, VAE pruning could potentially be instantiated by setting the posterior distribution of unneeded dimensions of the vector \mathbf{z} to the prior. By this we mean that if dimension j is not needed, then $q_\phi(z_j|\mathbf{x}) = p(z_j)$, uninformative noise that plays no role in improving reconstructions of \mathbf{x} . This capability has been noted in traditional VAE models (Dai et al., 2018; Dai and Wipf, 2019), but never rigorously analyzed in the context of CVAE extensions or cycle-consistent training. In this regard, the analysis in Section 4 will elucidate special cases whereby CycleCVAE can provably lead to optimal pruning, the first such analysis of CVAE models, cycle-consistent or otherwise. This serves to motivate the proposed pipeline as a vehicle for learning an implicit bijection even without knowing the dimensionality or distribution of data from \mathcal{U} , a particularly relevant notion given the difficulty in directly estimating $\dim[\mathcal{U}]$ in practice.

²Note that $\dim[\mathcal{X}]$ refers to the intrinsic dimensionality of \mathcal{X} , which could be a low-dimensional manifold embedded in a higher-dimensional ambient space; same for $\dim[\mathcal{Y}]$.

Secondly, because the CycleCVAE model is explicitly predicated upon a *known* prior $p(\mathbf{z})$ (as opposed to the unknown distribution of \mathbf{u}), other model components are calibrated accordingly such that there is no need to provide an empirical estimate of an unknown prior. Consequently, there is no barrier to cycle-consistent training or the generation of new \mathbf{x} conditioned on \mathbf{y} .

4 Formal Analysis of Special Cases

To the best of our knowledge, there is essentially no existing analysis of cycle-consistent training in the challenging yet realistic scenarios where a bijection between \mathbf{x} and \mathbf{y} *cannot* be assumed to hold.³ In this section we present a simplified case whereby the proposed CycleCVAE objective (with added distributional constraints) is guaranteed to have no bad local minimum in a specific sense to be described shortly. The forthcoming analysis relies on the assumption of a Gaussian CVAE with an affine model for the functions $\{h_\theta, h_\theta^+\}$; however, the conclusions we draw are likely to be loosely emblematic of behavior in broader regimes of interest. While admittedly simplistic, the resulting CVAE objective remains non-convex, with a combinatorial number of distinct local minima. Hence it is still non-trivial to provide any sort of guarantees in terms of associating local minima with ‘good’ solutions, e.g., solutions that recover the desired latent factors, etc. In fact, prior work has adopted similar affine VAE decoder assumptions, but only in the much simpler case of vanilla VAE models (Dai et al., 2018; Lucas et al., 2019), i.e., no cycle training or conditioning as is our focus herein.

4.1 Affine CycleCVAE Model

For analysis purposes, we consider a CVAE model of continuous data $\mathbf{x} \in \mathbb{R}^{r_x}$, $\mathbf{y} \in \mathbb{R}^{r_y}$, and $\mathbf{z} \in \mathbb{R}^{r_z}$, where r_x , r_y , and r_z are the respective sizes of \mathbf{x} , \mathbf{y} , and \mathbf{z} . We assume $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \mathbf{I})$ and a typical Gaussian decoder $p_\theta(\mathbf{x}|\hat{\mathbf{y}}, \mathbf{z}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_x, \boldsymbol{\Sigma}_x)$, where the mean network satisfies the affine parameterizations

$$\begin{aligned} \boldsymbol{\mu}_x &= h_\theta(\hat{\mathbf{y}}, \mathbf{z}) = \mathbf{W}_x \hat{\mathbf{y}} + \mathbf{V}_x \mathbf{z} + \mathbf{b}_x, \\ \text{with } \hat{\mathbf{y}} &= h_\theta^+(\mathbf{x}) = \mathbf{W}_y \mathbf{x} + \mathbf{b}_y. \end{aligned} \quad (7)$$

In this expression, $\{\mathbf{W}_x, \mathbf{W}_y, \mathbf{V}_x, \mathbf{b}_x, \mathbf{b}_y\}$ represents the set of all weight matrices and bias vectors which define the decoder mean $\boldsymbol{\mu}_x$. And as is often assumed in practical VAE models, we set $\boldsymbol{\Sigma}_x = \gamma \mathbf{I}$, where $\gamma > 0$ is a scalar parameter within the parameter set θ . Despite these affine assumptions, the CVAE energy function can still have a combinatorial number of distinct local minima as mentioned previously. However, we will

³Note that Grover et al. (2020) address identifiability issues that arise during cycle training, but only in the context of strictly bijective scenarios.

closely examine conditions whereby all these local minima are actually global minima that correspond with the optimal inversion of a non-trivial generative model.

Although we could proceed by allowing the encoder to be arbitrarily complex, when the decoder mean function is forced to be affine and $\boldsymbol{\Sigma}_x = \gamma \mathbf{I}$, a Gaussian encoder with affine moments is sufficient to achieve the optimal CVAE cost. Specifically, without any loss of representational capacity, we may choose $q_\phi(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z)$ with $\boldsymbol{\mu}_z = \mathbf{W}_z \mathbf{x} + \mathbf{b}_z$ and a diagonal $\boldsymbol{\Sigma}_z = \text{diag}[\mathbf{s}]^2$, where \mathbf{s} is an arbitrary parameter vector independent of \mathbf{x} .⁴ Collectively, these specifications lead to the complete parameterization $\theta = \{\mathbf{W}_x, \mathbf{W}_y, \mathbf{V}_x, \mathbf{b}_x, \gamma\}$, $\phi = \{\mathbf{W}_z, \mathbf{b}_z, \mathbf{s}\}$, and the CVAE energy given by $\ell_x(\theta, \phi) \equiv$

$$\int \left\{ \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} \left[\frac{1}{\gamma} \|(\mathbf{I} - \mathbf{W}_x \mathbf{W}_y) \mathbf{x} - \mathbf{V}_x \mathbf{z} - \mathbf{b}_x\|_2^2 \right] \right. \quad (8) \\ \left. + d \log \gamma + \sum_{k=1}^{r_z} (s_k^2 - \log s_k^2) + \|\mathbf{W}_z \mathbf{x} + \mathbf{b}_z\|_2^2 \right\} \rho_{gt}^x(d\mathbf{x}),$$

noting that, without loss of generality, we have absorbed a $\mathbf{W}_x \mathbf{b}_y$ factor into \mathbf{b}_x .

And finally, for the corresponding $\ell_y(\theta)$ model we specify $p_\theta(\mathbf{y}|\hat{\mathbf{x}}) = \mathcal{N}(\mathbf{y}|\boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y)$ using a shared, cycle-consistent parameterization borrowed from (7). For this purpose, we adopt

$$\boldsymbol{\mu}_y = h_\theta^+(\hat{\mathbf{x}}) = \mathbf{W}_y \hat{\mathbf{x}} + \mathbf{b}_y, \quad \text{with } \hat{\mathbf{x}} = \mathbf{W}_x \mathbf{y} + \mathbf{V}_x \mathbf{z} + \mathbf{b}_x, \quad (9)$$

and $\boldsymbol{\Sigma}_y = \gamma \mathbf{I}$. Given these assumptions, we have

$$\ell_y(\theta) \equiv \int \left(\boldsymbol{\epsilon}_y^\top \boldsymbol{\Sigma}_{\epsilon_y}^{-1} \boldsymbol{\epsilon}_y + \log |\boldsymbol{\Sigma}_{\epsilon_y}| \right) \rho_{gt}^y(d\mathbf{y}) \quad (10)$$

excluding irrelevant constants, where $\boldsymbol{\epsilon}_y \triangleq (\mathbf{I} - \mathbf{W}_y \mathbf{W}_x) \mathbf{y} - \mathbf{b}_y$, $\boldsymbol{\Sigma}_{\epsilon_y} \triangleq \gamma \mathbf{I} + \mathbf{W}_y \mathbf{V}_x \mathbf{V}_x^\top \mathbf{W}_y^\top$ and again, analogous to before we have absorbed $\mathbf{W}_y \mathbf{b}_x$ into \mathbf{b}_y without loss of generality.

4.2 Properties of Global/Local Minima

As a preliminary thought experiment, we can consider the minimization of $\ell_{\text{cycle}}(\theta, \phi)$, where $\ell_x(\theta, \phi)$ is defined via (8) and $\ell_y(\theta)$ via (10), but no assumptions are placed on the distributions ρ_{gt}^x and ρ_{gt}^y . In this situation, it is obvious that even CycleCVAE global minima, were they obtainable, will not generally recover the ground-truth mappings between paired $\mathbf{x} \sim \rho_{gt}^x$ and $\mathbf{y} \sim \rho_{gt}^y$; there simply will not generally be sufficient capacity. Furthermore, it can even be shown that under

⁴Note also that because $\hat{\mathbf{y}} = \mathbf{W}_y \mathbf{x} + \mathbf{b}_y$ is an affine function of \mathbf{x} , including this factor in the encoder representation is redundant, i.e., it can be absorbed into $\boldsymbol{\mu}_z = \mathbf{W}_z \mathbf{x} + \mathbf{b}_z$ without loss of generality.

quite broad conditions there will exist a combinatorial number of *non-global* local minima, meaning local minimizers that fail to achieve the lowest possible cost.

Hence we now present a narrower scenario with constraints on the ground-truth data to better align with the affine simplification described in Section 4.1. This will allow us to formulate conditions whereby all local minima are actually global minima capable of accurately modeling the ground-truth surjection. To this end, we define the following affine ground-truth model:

Definition 1 (Affine Surjective Model) *We define an affine surjective model whereby all matched $\{\mathbf{x}, \mathbf{y}\}$ pairs satisfy*

$$\mathbf{x} = \mathbf{A}\mathbf{y} + \mathbf{B}\mathbf{u} + \mathbf{c} \quad \text{and} \quad \mathbf{y} = \mathbf{D}\mathbf{x} + \mathbf{e}, \quad (11)$$

with $\mathbf{B} \in \text{null}[\mathbf{D}]$, $\mathbf{D}\mathbf{A} = \mathbf{I}$, $\mathbf{D}\mathbf{c} = -\mathbf{e}$, $\text{rank}[\mathbf{A}] = r_y < r_x$ and $\text{rank}[\mathbf{B}] \leq r_x - r_y$. Furthermore, we assume that $\mathbf{y} \sim \rho_{gt}^y$ and $\mathbf{u} \sim \rho_{gt}^u$ are uncorrelated, and the measure assigned to the transformed random variable $\mathbf{W}\mathbf{y} + \mathbf{V}\mathbf{u}$ is equivalent to ρ_{gt}^y iff $\mathbf{W} = \mathbf{I}$ and $\mathbf{V} = \mathbf{0}$. We also enforce that \mathbf{y} and \mathbf{u} have zero mean and identity covariance, noting that any nonzero mean components can be absorbed into \mathbf{c} . Among other things, the stated conditions of the affine surjective model collectively ensure that the mappings $\mathbf{y} \rightarrow \mathbf{x}$ and $\mathbf{x} \rightarrow \mathbf{y}$ can be mutually satisfied.

Additionally, for later convenience, we also define $r_c \triangleq \text{rank}(\mathbb{E}_{\rho_{gt}^x}[\mathbf{x}\mathbf{x}^\top]) \leq r_x$. We then have the following:

Proposition 2 *Assume that matched pairs $\{\mathbf{x}, \mathbf{y}\}$ follow the affine surjective model from Definition 1. Then subject to the constraint $\rho_{gt}^y = \rho_\theta^y$, where ρ_θ^y defines the θ -dependent distribution of \hat{y} , all local minima of the CycleVAE objective $\ell_{\text{cycle}}(\theta, \phi)$, with $\ell_x(\theta, \phi)$ taken from (8) and $\ell_y(\theta)$ from (10), will be global minima in the limit $\gamma \rightarrow 0$ assuming $r_z \geq r_c - r_y$. Moreover, the globally optimal parameters $\{\theta^*, \phi^*\}$ will satisfy*

$$\begin{aligned} \mathbf{W}_x^* &= \mathbf{A}, \quad \mathbf{V}_x^* = [\tilde{\mathbf{B}}, \mathbf{0}] \mathbf{P}, \quad \mathbf{b}_x^* = \mathbf{c}, \\ \mathbf{W}_y^* &= \tilde{\mathbf{D}}, \quad \mathbf{b}_y^* = -\tilde{\mathbf{D}}\mathbf{c}, \end{aligned} \quad (12)$$

where $\tilde{\mathbf{B}}$ has $\text{rank}[\mathbf{B}]$ columns, $\text{span}[\tilde{\mathbf{B}}] = \text{span}[\mathbf{B}]$, \mathbf{P} is a permutation matrix, and $\tilde{\mathbf{D}}$ satisfies $\tilde{\mathbf{D}}\mathbf{A} = \mathbf{I}$ and $\mathbf{B} \in \text{null}[\tilde{\mathbf{D}}]$.

Note that in practice, we are free to choose r_z as large as we want, so the requirement that $r_z \geq r_c - r_y$ is not significant. Additionally, the constraint $\rho_{gt}^y = \rho_\theta^y$ can be instantiated (at least approximately) by including a penalty on the divergence between these two distributions. This is feasible using only unpaired samples of \mathbf{y} (for estimating ρ_{gt}^y) and \mathbf{x} (for estimating ρ_θ^y), and

most cycle-consistent training pipelines contain some analogous form of penalty on distributional differences between cycles (Lample et al., 2018).

And finally, if $r_y + \text{rank}[\mathbf{B}] = r_x$, then the dual requirements that $\tilde{\mathbf{D}}\mathbf{A} = \mathbf{I}$ and $\mathbf{B} \in \text{null}[\tilde{\mathbf{D}}]$ will ensure that $\tilde{\mathbf{D}} = \mathbf{D}$ and $\mathbf{b}_y^* = \mathbf{e}$. However, even if $\tilde{\mathbf{D}} \neq \mathbf{D}$ it is inconsequential for effective recovery of the ground-truth model since any \mathbf{x} produced by Definition 1 will nonetheless still map to the correct \mathbf{y} when applying $\tilde{\mathbf{D}}$ instead of \mathbf{D} .

Corollary 3 *Given the same setup as Proposition 2, let $\{\mathbf{W}_z^*, \mathbf{b}_z^*\}$ denote the CVAE encoder parameters of any minimum. Then $\mathbf{W}_z^* = \mathbf{P} \begin{bmatrix} \tilde{\mathbf{W}}_z^* \\ \mathbf{0} \end{bmatrix}$ and $\mathbf{b}_z^* = \mathbf{P} \begin{bmatrix} \tilde{\mathbf{b}}_z^* \\ \mathbf{0} \end{bmatrix}$, where $\tilde{\mathbf{W}}_z^*$ has $\text{rank}[\mathbf{B}]$ rows, $\tilde{\mathbf{b}}_z^* \in \mathbb{R}^{\text{rank}[\mathbf{B}]}$, and there exists a bijection between \mathbf{x} and $\{\mathbf{y}, \tilde{\boldsymbol{\mu}}_z\}$, with $\tilde{\boldsymbol{\mu}}_z \triangleq \tilde{\mathbf{W}}_z^* \mathbf{x} + \tilde{\mathbf{b}}_z^*$ (i.e., the nonzero elements of $\boldsymbol{\mu}_z$).*

We will now discuss various take-home messages related to these results.

4.3 Practical Implications

As alluded to in Section 3.4, we will generally not know in advance $p(\mathbf{u})$ or even $\dim[\mathcal{U}]$, which reduces to $r_c - r_y$ in the simplified affine case. Hence inducing a bijection may seem problematic on the surface. Fortunately though, Proposition 2 and Corollary 3 indicate that as long as we choose $r_z \geq \dim[\mathcal{U}]$ in our CVAE model, we can nonetheless still learn an *implicit* bijection between \mathbf{x} and $\{\mathbf{y}, \tilde{\boldsymbol{\mu}}_z\}$, where $\tilde{\boldsymbol{\mu}}_z$ are the informative (nonzero) dimensions of $\boldsymbol{\mu}_z$ that actually contribute to the reconstruction of \mathbf{x} . In the affine case, these are the dimensions of \mathbf{z} aligned with nonzero columns of $\tilde{\mathbf{B}}$, but in general these dimensions could more loosely refer to the degrees-of-freedom in \mathbf{z} that, when altered, lead to changes in $\hat{\mathbf{x}}$. The remaining superfluous dimensions of \mathbf{z} are set to the uninformative prior $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \mathbf{I})$ and subsequently filtered out by the CVAE decoder module parameterized by $h_\theta(\mathbf{y}, \mathbf{z})$. In this diminutive role, they have no capacity for interfering with any attempts to learn a bijection.

Even so, in non-affine cases it is impossible to guarantee that all local minima correspond with the recovery of h_{gt} and h_{gt}^+ , or that these functions are even identifiable. Indeed it is not difficult to produce counterexamples whereby recovery is formally impossible. However, if h_θ and h_θ^+ are chosen with inductive biases reasonably well-aligned with their ground-truth counterparts, up to unidentifiable latent transformations of the unobservable \mathbf{u} , we may expect that an approximate bijection between the true \mathbf{x} and $\{\mathbf{y}, \tilde{\boldsymbol{\mu}}_z\}$ can nonetheless be

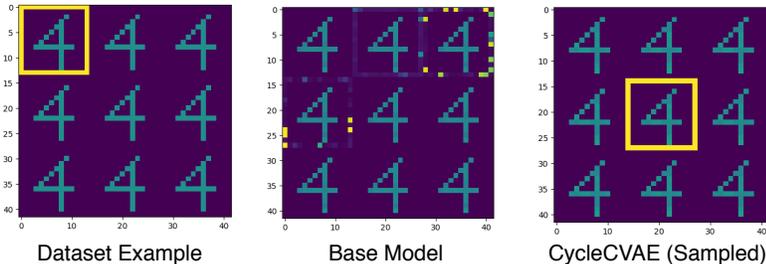


Figure 1: Left: example image from dataset. Middle: image produced by baseline cycle training with $\mathbf{y} = 4$. Right: a sample image generated by CycleCVAE conditioned on $\mathbf{y} = 4$. For the latter, the position of yellow border is random. In contrast, the base model fails to learn the random border distribution.

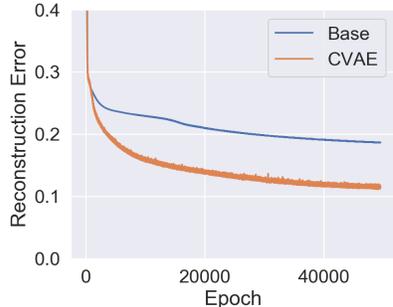


Figure 2: Cycle-consistent reconstruction errors of baseline and CycleCVAE models.

inferred via cycle-consistent training, at least provided we can avoid suboptimal local minimizers that can be introduced by more complex nonlinear decoder models.

5 Experiments

In this section we first consider a synthetic image experiment that supports the theoretical motivation for CycleCVAE. We then turn to practical real-world evaluations involving the conversion between knowledge graphs and text sequences, a core ingredient of natural language understanding/generation and the application that initially motivated our work. We then conclude with an enhanced pipeline involving the recent pre-trained T5 model (Raffel et al., 2020).

5.1 Synthetic Image Experiment

We first conduct an experiment designed such that the surjective conditions of Definition 1 in Section 4.2 are loosely satisfied (see supplementary for reasons). As shown in Figure 1 (left panel), each data sample has three components: a digit, its image, and a decorative yellow border. The digit takes value in $\{0, \dots, 9\}$ and is represented by a 10-dim one-hot vector \mathbf{y} . The corresponding image \mathbf{x} involves 3×3 tiles, and each tile contains the same image of digit \mathbf{y} . One of the 9 tiles is decorated with a 1-pixel-wide yellow border, and which tile will have this border is determined by \mathbf{u} , a 9-dim one-hot vector indicating the 9 possible tiles.

We train two models on this dataset, a base model using standard cycle training, and our CycleCVAE that incorporates the proposed CVAE into a baseline cycle model (see supplementary for network description and details). After training, generated samples of the two approaches when presented with the digit ‘4’ are shown in Figure 1 (middle and right panels). The base model fails to learn the yellow border as it cannot handle the one-to-many mapping from digits to images.

Meanwhile the random CycleCVAE sample correctly places the border around one of the tiles (different CycleCVAE samples simply move the border to different tiles as desired; see supplementary). Finally, consistent with these generation results, the training curves from Figure 2 reveal that the reconstruction error of the base model, which assumes a bijection, plateaus at a significantly higher value than the CycleCVAE model.

5.2 Knowledge Graph to Text Conversion

We now turn to more challenging real-world experiments involving the surjective mapping between knowledge graphs and text sequences. Here the ideal goal is to generate diverse, natural text from a fixed knowledge graph, or extract the knowledge graph from a piece of text. To this end we compare CycleCVAE against SOTA methods on the widely-used WebNLG graph-to-text dataset (Gardent et al., 2017).

WebNLG Dataset and Test Setup WebNLG data is extracted from DBpedia, where each graph consists of 2–7 nodes and the corresponding text is descriptions of these graphs collected by crowd-sourcing. We follow the preprocessing of Moryossef et al. (2019) and obtain 13K training, 1.6K validation, and 5K test text-graph pairs. Please see the supplementary for details of the CycleCVAE architecture explicitly designed for handling text and graph data. Note that we did not include any additional penalty function on the divergence between ρ_{gt}^y and ρ_{θ}^y ; the architecture inductive biases were sufficient for good performance.

Metrics We measure performance using three metrics: (1) text generation quality with the standard BLEU score (Papineni et al., 2002),⁵ (2) graph construction accuracy via the F1 score of the edge predic-

⁵BLEU (%) counts the 1- to 4-gram overlap between the generated sentence and ground truth.

tions among given entity nodes, and (3) text diversity. Text diversity is an increasingly important criterion for NLP because the same meaning can be conveyed in various expressions, and intelligent assistants should master such variations. We evaluate diversity by reporting the number of distinct sentence variations obtained after running the generation model 10 times.

Accuracy Results Since cycle training only requires unsupervised data, we have to break the text-graph pairs to evaluate unsupervised performance. In this regard, there are two ways to process the data. First, we can use 100% of the training data and just shuffle the text and graphs so that the matching/supervision is lost. This is the setting in Table 1, which allows for direct head-to-head comparisons with SOTA supervised methods (assuming no outside training data). The supervised graph-to-text baselines include *Melbourne* (introduced in Gardent et al. (2017)), *StrongNeural*, *BestPlan* (Moryossef et al., 2019), *Seg&Align* (Shen et al., 2020), and *G2T* (Koncel-Kedziorski et al., 2019). Supervised text-to-graph models include *OnePass* (Wang et al., 2019), and *T2G*, a BiLSTM model we implemented. Unsupervised methods include *RuleBased* and *GT-BT* both by Schmitt et al. (2020). Finally, *CycleBase* is our deterministic cycle training model with the architectural components borrowed from CycleCVAE. Notably, from Table 1 we observe that our model outperforms other unsupervised methods, and it is even competitive with SOTA supervised models in both the graph-to-text (BLEU) and text-to-graph (F1) directions.

	Text(BLEU)	Graph(F1)	#Variations
Supervised (100%)			
Melbourne	45.0	-	1
StrongNeural	46.5	-	1
BestPlan	47.4	-	1
Seg&Align	46.1	-	1
G2T	45.8	-	1
OnePass	-	66.2	-
T2G	-	60.6	-
Unsupervised (100%, Shuffled)			
RuleBased	18.3	0	1
GT-BT	37.7	39.1	1
CycleBase (Ours)	46.2	61.2	1
CycleCVAE (Ours)	46.5	62.6	4.67

Table 1: Performance on the full WebNLG dataset.

	Text(BLEU)	Graph(F1)	#Variations
Supervised (50%)			
G2T	44.5	-	1
T2G	-	59.7	-
Unsupervised (first 50% text, last 50% graph)			
CycleBase (Ours)	43.1	59.8	1
CycleCVAE (Ours)	43.3	60.0	4.01

Table 2: Performance on WebNLG with 50% data.

In contrast, a second, stricter unsupervised protocol

involves splitting the dataset into two halves, extracting text from the first half, and graphs from the second half. This is the setting in Table 2, which avoids the possibility of seeing any overlapping entities during training. Although performance is slightly worse given less training data, the basic trends are the same.

Diversity Results From Tables 1 and 2 we also note that CycleCVAE can generate on average more than 4 different sentence types for a given knowledge graph; all other SOTA methods can only generate a single sentence per graph. Additionally, we have calculated that CycleCVAE generates more than two textual paraphrases for 99% of test instances, and the average edit distance between two paraphrases is 12.24 words (see supplementary). Moreover, CycleCVAE text diversity does not harm fluency and semantic relevance as the BLEU score is competitive with SOTA methods as mentioned previously.

Diverse Text Output Generated by CycleCVAE

- The population density of Arlington, Texas is 1472.0.
 - Arlington, Texas has a population density of 1472.0.
 - Alan Bean, who was born in Wheeler, Texas, is now "retired."
 - Alan Bean is a United States citizen who was born in Wheeler, Texas. He is now "retired."
-

Table 3: Every two variations are generated by CycleCVAE from the same knowledge graph.

We list text examples generated by our model in Table 3, with more in the supplementary. The diverse generation is a significant advantage for many real applications. For example, it can make automated conversations less boring and simulate different scenarios. And diversity can push model generated samples closer to the real data distribution because there exist different ways to verbalize the same knowledge graph (although diversity will not in general improve BLEU scores, and can sometimes actually lower them).

5.3 Integrating CycleCVAE with T5

Previous graph-text results are all predicated on no outside training data beyond WebNLG. However, we now consider an alternative testing scenario whereby outside training data can be incorporated by integrating CycleCVAE with a large pretrained T5 sequence-to-sequence model (Raffel et al., 2020). Such models have revolutionized many NLP tasks and can potentially improve the quality of the graph-to-text direction in cycle training on WebNLG. To this end, we trained a CycleCVAE model, with the function $h_{\theta}(\mathbf{y}, \mathbf{z})$ formed from a pretrained T5 architecture (see supplementary for details). Results are shown in Table 4, where un-

supervised CycleCVAE+T5 produces a competitive BLEU score relative to fully supervised T5 baselines. It also maintains diversity of generated text sequences.

Model	BLEU	#Vars.
Supervised T5 (Kale, 2020)	57.1	1
Supervised T5 (Ribeiro et al., 2020)	57.4	1
Supervised T5 (Our Impl.)	56.4	1
Unsupervised CycleCVAE+T5	55.7	3.84

Table 4: Text generation results with T5 on WebNLG.

6 Conclusion

We have proposed CycleCVAE for explicitly handling non-bijective surjections commonly encountered in real-world applications of unsupervised cycle-consistent training. Our framework has both a solid theoretical foundation and strong empirical performance on practical knowledge graph-to-text conversion problems. For future work we can consider extending CycleCVAE to handle many-to-many (non-bijective, non-surjective) mappings, or unsolved applications such as conversions between scene graphs and realistic images (which remains extremely difficult even with supervision).

Acknowledgements

We thank all reviewers for their constructive suggestions.

References

- Artetxe, Mikel, Labaka, Gorka, Agirre, Eneko, and Cho, Kyunghyun (2018). “Unsupervised neural machine translation”. *International Conference on Learning Representations*.
- Bitton, Adrien, Esling, Philippe, and Chemla-Romeu-Santos, Axel (2018). “Modulated variational autoencoders for many-to-many musical timbre transfer”. *arXiv preprint arXiv:1810.00222*.
- Cheng, Yong, Xu, Wei, He, Zhongjun, He, Wei, Wu, Hua, Sun, Maosong, and Liu, Yang (2016). “Semi-supervised learning for neural machine translation”. *Annual Meeting of the Association for Computational Linguistics*.
- Choi, Yunjey, Choi, Minje, Kim, Munyoung, Ha, Jungwoo, Kim, Sunghun, and Choo, Jaegul (2018). “Stargan: unified generative adversarial networks for multi-domain image-to-image translation”. *Computer Vision and Pattern Recognition*.
- Dai, Bin, Wang, Yu, Aston, John, Hua, Gang, and Wipf, David (2018). “Connections with robust PCA and the role of emergent sparsity in variational autoencoder models”. *Journal of Machine Learning Research*.
- Dai, Bin and Wipf, David (2019). “Diagnosing and enhancing VAE models”. *International Conference on Learning Representations*.
- Doersch, Carl (2016). “Tutorial on variational autoencoders”. *arXiv preprint arXiv:1606.05908*.
- Edunov, Sergey, Ott, Myle, Auli, Michael, and Grangier, David (2018). “Understanding back-translation at scale”. *Empirical Methods in Natural Language Processing*.
- Gardent, Claire, Shimorina, Anastasia, Narayan, Shashi, and Perez-Beltrachini, Laura (2017). “The WebNLG challenge: generating text from RDF data”. *International Conference on Natural Language Generation*.
- Godard, Clément, Mac Aodha, Oisín, and Brostow, Gabriel J. (2017). “Unsupervised monocular depth estimation with left-right consistency”. *Computer Vision and Pattern Recognition*.
- Grover, Aditya, Chute, Christopher, Shu, Rui, Cao, Zhangjie, and Ermon, Stefano (2020). “AlignFlow: cycle consistent learning from multiple domains via normalizing flows.” *AAAI Conference on Artificial Intelligence*.
- Guo, Qipeng, Jin, Zhijing, Wang, Ziyu, Qiu, Xipeng, Zhang, Weinan, Zhu, Jun, Zhang, Zheng, and Wipf, David (2021). “Fork or fail: cycle-consistent training with many-to-one mappings”. *arXiv preprint arXiv:2012.07412*.
- He, Di, Xia, Yingce, Qin, Tao, Wang, Liwei, Yu, Nenghai, Liu, Tie-Yan, and Ma, Wei-Ying (2016). “Dual learning for machine translation”. *Neural Information Processing Systems*.
- Hoffman, Judy, Tzeng, Eric, Park, Taesung, Zhu, Jun-Yan, Isola, Phillip, Saenko, Kate, Efros, Alexei A., and Darrell, Trevor (2018). “CyCADA: cycle-consistent adversarial domain adaptation”. *International Conference on Machine Learning*.
- Hori, Takaaki, Astudillo, Ramón Fernández, Hayashi, Tomoki, Zhang, Yu, Watanabe, Shinji, and Roux, Jonathan Le (2019). “Cycle-consistency training for

- end-to-end speech recognition”. *International Conference on Acoustics, Speech and Signal Processing*.
- Jha, Ananya Harsh, Anand, Saket, Singh, Maneesh, and Veeravasarapu, V. S. R. (2018). “Disentangling factors of variation with cycle-consistent variational auto-encoders”. *European Conference on Computer Vision*.
- Jin, Di, Jin, Zhijing, Zhou, Joey Tianyi, and Szolovits, Peter (2020). “A simple baseline to semi-supervised domain adaptation for machine translation”. *arXiv preprint arXiv:2001.08140*.
- Jin, Zhijing, Jin, Di, Mueller, Jonas, Matthews, Nicholas, and Santus, Enrico (2019). “IMaT: unsupervised text attribute transfer via iterative matching and translation”. *Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing*.
- Kalal, Zdenek, Mikolajczyk, Krystian, and Matas, Jiri (2010). “Forward-backward error: automatic detection of tracking failures”. *International Conference on Pattern Recognition*.
- Kale, Mihir (2020). “Text-to-text pre-training for data-to-text tasks”. *arXiv preprint arXiv:2005.10433*.
- Kameoka, Hirokazu, Kaneko, Takuhiro, Tanaka, Kou, and Hojo, Nobukatsu (2018). “StarGAN-VC: non-parallel many-to-many voice conversion with star generative adversarial networks”. *arXiv preprint arXiv:1806.02169*.
- Kingma, Diederik P. and Welling, Max (2014). “Auto-encoding variational bayes”. *International Conference on Learning Representations*.
- Koncel-Kedziorski, Rik, Bekal, Dhanush, Luan, Yi, Lapata, Mirella, and Hajishirzi, Hannaneh (2019). “Text generation from knowledge graphs with graph transformers”. *North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Lample, Guillaume, Conneau, Alexis, Denoyer, Ludovic, and Ranzato, Marc’Aurelio (2018). “Unsupervised machine translation using monolingual corpora only”. *International Conference on Learning Representations*.
- Lample, Guillaume, Subramanian, Sandeep, Smith, Eric Michael, Denoyer, Ludovic, Ranzato, Marc’Aurelio, and Boureau, Y-Lan (2019). “Multiple-attribute text rewriting”. *International Conference on Learning Representations*.
- Lee, Keonnyeong, Yoo, In-Chul, and Yook, Dongsuk (2019). “Many-to-many voice conversion using cycle-consistent variational autoencoder with multiple decoders”. *arXiv preprint arXiv:1909.06805*.
- Liu, Ming-Yu, Breuel, Thomas, and Kautz, Jan (2017). “Unsupervised image-to-image translation networks”. *Neural Information Processing Systems*.
- Lucas, James, Tucker, George, Grosse, Roger B, and Norouzi, Mohammad (2019). “Don’t blame the ELBO! A linear VAE perspective on posterior collapse”. *Neural Information Processing Systems*.
- Mor, Noam, Wolf, Lior, Polyak, Adam, and Taigman, Yaniv (2018). “A universal music translation network”. *arXiv preprint arXiv:1805.07848*.
- Moryossef, Amit, Goldberg, Yoav, and Dagan, Ido (2019). “Step-by-step: separating planning from realization in neural data-to-text generation”. *North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Papineni, Kishore, Roukos, Salim, Ward, Todd, and Zhu, Wei-Jing (2002). “Bleu: a method for automatic evaluation of machine translation”. *Annual Meeting of the Association for Computational Linguistics*.
- Raffel, Colin, Shazeer, Noam, Roberts, Adam, Lee, Katherine, Narang, Sharan, Matena, Michael, Zhou, Yanqi, Li, Wei, and Liu, Peter J (2020). “Exploring the limits of transfer learning with a unified text-to-text transformer”. *Journal of Machine Learning Research*.
- Rezende, Danilo Jimenez, Mohamed, Shakir, and Wierstra, Daan (2014). “Stochastic backpropagation and approximate inference in deep generative models”. *International Conference on Machine Learning*.
- Ribeiro, Leonardo FR, Schmitt, Martin, Schütze, Hinrich, and Gurevych, Iryna (2020). “Investigating pre-trained language models for graph-to-text generation”. *arXiv preprint arXiv:2007.08426*.
- Schmitt, Martin, Sharifzadeh, Sahand, Tresp, Volker, and Schütze, Hinrich (2020). “An unsupervised joint system for text generation from knowledge graphs and semantic parsing”. *Empirical Methods in Natural Language Processing*.

- Senrich, Rico, Haddow, Barry, and Birch, Alexandra (2016). “Improving neural machine translation models with monolingual data”. *Annual Meeting of the Association for Computational Linguistics*.
- Shen, Tianxiao, Lei, Tao, Barzilay, Regina, and Jaakkola, Tommi S. (2017). “Style transfer from non-parallel text by cross-alignment”. *Neural Information Processing Systems*.
- Shen, Xiaoyu, Chang, Ernie, Su, Hui, Niu, Cheng, and Klakow, Dietrich (2020). “Neural data-to-text generation via jointly learning the segmentation and correspondence”. *Annual Meeting of the Association for Computational Linguistics*.
- Shi, Yuge, Narayanaswamy, Siddharth, Paige, Brooks, and Torr, Philip H. S. (2019). “Variational mixture-of-experts autoencoders for multi-modal deep generative models”. *Neural Information Processing Systems*.
- Sohn, Kihyuk, Lee, Honglak, and Yan, Xinchun (2015). “Learning structured output representation using deep conditional generative models”. *Neural Information Processing Systems*.
- Sundaram, Narayanan, Brox, Thomas, and Keutzer, Kurt (2010). “Dense point trajectories by GPU-accelerated large displacement optical flow”. *European Conference on Computer Vision*.
- Tseng, Bo-Hsiang, Cheng, Jianpeng, Fang, Yimai, and Vandyke, David (2020). “A generative model for joint natural language understanding and generation”. *Annual Meeting of the Association for Computational Linguistics*.
- Wang, Haoyu, Tan, Ming, Yu, Mo, Chang, Shiyu, Wang, Dakuo, Xu, Kun, Guo, Xiaoxiao, and Potdar, Saloni (2019). “Extracting multiple-relations in one-pass with pre-trained transformers”. *Annual Meeting of the Association for Computational Linguistics*.
- Zhu, Jun-Yan, Park, Taesung, Isola, Phillip, and Efros, Alexei A. (2017a). “Unpaired image-to-image translation using cycle-consistent adversarial networks”. *International Conference on Computer Vision*.
- Zhu, Jun-Yan, Zhang, Richard, Pathak, Deepak, Darrell, Trevor, Efros, Alexei A., Wang, Oliver, and Shechtman, Eli (2017b). “Toward multimodal image-to-image translation”. *Neural Information Processing Systems*.