

CYCLE-CONSISTENT GENERATIVE ADVERSARIAL NETWORKS AS A BAYESIAN APPROXIMATION

PREFACE

This chapter is derived from work previously published as:

Louis C Tiao, Edwin V Bonilla, and Fabio T Ramos. “Cycle-Consistent Adversarial Learning as Approximate Bayesian Inference”. In: *ICML 2018 Theoretical Foundations and Applications of Deep Generative Models*. Stockholm, Sweden, July 2018. (Accepted as *Oral presentation*).

4.1 INTRODUCTION

Learning correspondences between entities from different domains is an important and challenging problem in ML, especially in the *absence of paired data*. Consider for example the task of image-to-image translation where we want to learn a mapping from an image in a source domain, such as a photograph of a natural scene, to a corresponding image in a target domain, such as the realisation of such a scene in an 1860s celebrated artist’s signature impressionistic style. The shortage of ground-truth pairings from the source domain to the target domain renders standard supervised approaches infeasible, thus motivating the need for unsupervised learning.

Within unsupervised approaches, a number of recently proposed CYCLEGAN methods have achieved remarkable success in addressing this problem [125, 302]. As their name suggests, these approaches are based upon two heuristics: (i) adversarial learning and (ii) cycle consistency. The former, adversarial learning [86], allows images in the source domain to be translated to output images that, to an auxiliary discriminator, are indistinguishable from images in the target domain, thereby matching their distributions. However, while distribution matching is necessary, it is insufficient to guarantee one-to-one mappings between the images, as the problem is heavily under-constrained. Briefly stated, the cycle-consistency is the constraint that an image mapped to a target domain should be *representable* in the original domain. It is this constraint that significantly shrinks the space of possible solutions.

Beyond the empirical risk minimisation framework motivated intuitively by the two principles mentioned above, the original CYCLEGAN formulation lacks any further theoretical justification. Besides providing sound quantification of uncertainty, a LVM allows us to disentangle our modelling assumptions from the inference machinery used to

reason about the model variables. Interpreting standard methods from a Bayesian perspective has contributed significantly to the understanding of these methods and to the development of new approaches [74, 261].


In this chapter, we introduce *implicit* LVMS, where the prior over hidden representations can be specified flexibly as an *implicit* distribution. We develop a VI algorithm for this model based on minimisation of the *symmetric* KL divergence between a *variational joint* and the exact joint distribution, in contrast to traditional reverse KL minimisation, which notoriously underestimate posterior's / exact distribution's support. Lastly, we demonstrate that the state-of-the-art CYCLEGANS as proposed contemporaneously by Kim et al. [125] and Zhu et al. [302] can be derived as a special case within our proposed VI framework, thus establishing its connection to approximate Bayesian inference methods.

4.2 IMPLICIT LATENT VARIABLE MODELS

LVMS are an indispensable tool for uncovering the hidden representations of observed data. In a LVM, an observation \mathbf{x} is assumed governed by its underlying hidden variable \mathbf{z} , which is drawn from a prior $p(\mathbf{z})$ and related to \mathbf{x} through the likelihood $p_\theta(\mathbf{x} | \mathbf{z})$. Accordingly, the joint density of \mathbf{x} and \mathbf{z} is given by

$$p_\theta(\mathbf{x}, \mathbf{z}) = p_\theta(\mathbf{x} | \mathbf{z})p(\mathbf{z}). \quad (4.1)$$

Given data distribution $q^*(\mathbf{x})$ and a finite collection $\mathbf{X} = \{\mathbf{x}_n\}_{n=1}^N$ of observations $\mathbf{x}_n \sim q^*(\mathbf{x})$, and the set of corresponding latent variables $\mathbf{Z} = \{\mathbf{z}_n\}_{n=1}^N$, the joint over all variables factorises as, $p_\theta(\mathbf{X}, \mathbf{Z}) = \prod_{n=1}^N p_\theta(\mathbf{x}_n, \mathbf{z}_n)$.



(a) Without amortised inference, each local latent variable is governed by its own local variational parameters. (b) With amortized inference, we condition on observed variables and employ a single set of global variational parameters.

Figure 4.1: Graphical representation of the *generative model* (solid) and the *recognition model* (dashed).

The graphical representation of implicit LVMS is depicted in Figure 4.1. Instead of approximating the exact posterior $p_\theta(\mathbf{z} | \mathbf{x}_n)$ for each \mathbf{x}_n , using a separate variational distribution $q(\mathbf{z}; \lambda_n)$ with local

variational parameters λ_n , we condition on \mathbf{x} and optimise a single set of variational parameters ϕ across all $\mathbf{x} \sim q^*(\mathbf{x})$. Accordingly, the variational distribution is denoted $q_\phi(\mathbf{z} | \mathbf{x}) \triangleq q(\mathbf{z} | \mathbf{x}; \phi)$.

4.2.1 Prescribed Likelihood

We specify the likelihood through a mapping \mathcal{F}_θ that takes as input random noise ξ and latent variable \mathbf{z} ,

$$\mathbf{x} \sim p_\theta(\mathbf{x} | \mathbf{z}) \Leftrightarrow \mathbf{x} = \mathcal{F}_\theta(\xi; \mathbf{z}), \quad \xi \sim p(\xi). \quad (4.2)$$

We shall restrict our attention to *prescribed* likelihoods, where evaluation of their density is tractable. This requires that $\mathcal{F}_\theta(\cdot; \mathbf{z})$ be a diffeomorphism wrt ξ and density $p(\xi)$ be tractable. For example, when $\mathcal{F}_\theta(\cdot; \mathbf{z})$ is a location-scale transform of noise ξ and $p(\xi)$ is Gaussian, we recover Gaussian observation models.

Our model specification is sufficiently general for encapsulating a broad range of familiar latent variable models, even when we make simplifying assumptions on the mapping $\mathcal{F}_\theta(\cdot; \mathbf{z})$. In particular, consider the special case where the mapping is an affine transformation of the noise vector ξ ,

$$\mathcal{F}_\theta(\xi; \mathbf{z}) \triangleq \boldsymbol{\mu}_\theta(\mathbf{z}) + \boldsymbol{\Sigma}_\theta(\mathbf{z})^{\frac{1}{2}}\xi, \quad \xi \sim \mathcal{N}(\mathbf{0}, \mathbf{I}),$$

for functions $\boldsymbol{\mu}_\theta$ and $\boldsymbol{\Sigma}_\theta$ parameterised by θ that take \mathbf{z} as input. To simplify matters further, assume $\boldsymbol{\Sigma}_\theta$ is constant wrt to its input, i. e. $\boldsymbol{\Sigma}_\theta(\mathbf{z}) = \boldsymbol{\Psi}$ for all \mathbf{z} . The likelihood can then be written explicitly as

$$p_\theta(\mathbf{x} | \mathbf{z}) = \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_\theta(\mathbf{z}), \boldsymbol{\Psi}).$$

FACTOR ANALYSIS & PROBABILISTIC PCA. In the case where the mean function $\boldsymbol{\mu}_\theta$ is an affine transformation of \mathbf{z} ,

$$\boldsymbol{\mu}_\theta(\mathbf{z}) \triangleq \mathbf{W}\mathbf{z} + \mathbf{b},$$

and the covariance matrix is diagonal $\boldsymbol{\Psi} = \text{diag}(\psi_1^2, \dots, \psi_D^2)$, we recover FA [8]. Furthermore, when the covariance matrix is isotropic $\boldsymbol{\Psi} = \psi^2\mathbf{I}$, we recover *probabilistic principal component analysis* (PPCA) [261]. In this example, the parameters θ consist of the factor loading matrix \mathbf{W} , the bias vector \mathbf{b} and the covariance matrix $\boldsymbol{\Psi}$.

DEEP AND NONLINEAR LATENT VARIABLE MODELS. By introducing nonlinearities to the mean function, we are able to recover nonlinear factor analysis [139], nonlinear Gaussian sigmoid belief networks [72], and other more sophisticated variants of deep latent variable models. When the mapping is defined by a MLP, we can recover simple instances of a variational autoencoder (VAE) with a Gaussian probabilistic decoder [127, 212].

4.2.2 *Implicit Prior*

In LVMS, the prior typically specified as a prescribed distribution, e. g. a factorised Gaussian centered at zero. Oftentimes, however, the practitioner possesses prior knowledge that simply cannot be embodied within a prescribed distribution. To address this limitation, we introduce *implicit* LVMS, wherein the prior over latent variables is specified as an implicit distribution $p^*(\mathbf{z})$, given only by a finite collection $\mathbf{Z}^* = \{\mathbf{z}_m^*\}_{m=1}^M$ of its samples,

$$\mathbf{z}_m^* \sim p^*(\mathbf{z}). \quad (4.3)$$

This formulation offers the utmost degree of flexibility in the treatment of prior information, the difficulties of which have hindered the application of Bayesian statistics since the time of Laplace [115].

EXAMPLE: UNPAIRED IMAGE-TO-IMAGE TRANSLATION . Suppose we have collections of images \mathbf{X} and \mathbf{Z}^* , which are assumed to be draws from the data distribution $q^*(\mathbf{x})$ and implicit prior distribution $p^*(\mathbf{z})$, respectively. For example, these might be photographs of natural landscapes and the paintings of Van Gogh. The goal of unpaired image-to-image translation is to learn the correspondence between variables \mathbf{x} and \mathbf{z} by capturing the underlying generative process specified by mapping \mathcal{F}_θ . This defines the likelihood $p_\theta(\mathbf{x} | \mathbf{z})$ – a conditional density of \mathbf{x} given \mathbf{z} . Continuing with the above example, the problem amounts to learning parameters θ of the mapping such that this conditional yields photorealistic renderings of scenes portrayed in Van Gogh’s paintings. Furthermore, the resulting posterior on the latent representation $p_\theta(\mathbf{z} | \mathbf{x})$ – a conditional density of \mathbf{z} given \mathbf{x} – should produce renderings of landscape scenery in Van Gogh’s iconic style.

4.3 VARIATIONAL INFERENCE

In this section, we describe the first component of our bipartite VI framework. In traditional VI, one specifies a family \mathcal{Q} of densities over the latent variables and seeks the member $q \in \mathcal{Q}$ closest in KL divergence to the exact posterior $p_\theta(\mathbf{z} | \mathbf{x})$ [17, 119, 276].

4.3.1 *Prescribed Variational Posterior*

We begin by describing the variational family $q \in \mathcal{Q}$. We adopt the common practice of *amortising* inference using an inference network [83]. Namely, instead of approximating the exact posterior $p_\theta(\mathbf{z} | \mathbf{x}_n)$ for each \mathbf{x}_n , using a separate variational distribution $q(\mathbf{z}; \lambda_n)$ with local variational parameters λ_n , we condition on \mathbf{x} and optimise a single set of variational parameters ϕ across all $\mathbf{x} \sim q^*(\mathbf{x})$.

The variational distribution $q_\phi(\mathbf{z} | \mathbf{x})$ is specified through an inverse mapping \mathcal{G}_ϕ that takes as input random noise ϵ and observed variable \mathbf{x} ,

$$\mathbf{z} \sim q_\phi(\mathbf{z} | \mathbf{x}) \Leftrightarrow \mathbf{z} = \mathcal{G}_\phi(\epsilon; \mathbf{x}), \quad \epsilon \sim p(\epsilon). \quad (4.4)$$

Just as mapping \mathcal{F}_θ underpins the generative model, mapping \mathcal{G}_ϕ underpins the *recognition model* [52]. As with the likelihood, *we restrict our attention to prescribed variational distributions*.

As depicted in Figure 4.1b, the dependency relationship between the variational parameters and the latent variables mirrors that of the model parameters and observed variables. This symmetry is crucial to the derivation of CYCLEGAN later in Section 4.5.3.2.

4.3.2 Reverse KL Variational Objective

Minimising the reverse KL between the exact and variational posterior is equivalent to maximising the ELBO, or minimising its *negative*, defined as

$$\begin{aligned} \mathcal{L}_{\text{NELBO}}(\boldsymbol{\theta}, \boldsymbol{\phi}) \triangleq & \mathbb{E}_{q^*(\mathbf{x})q_\phi(\mathbf{z} | \mathbf{x})}[-\log p_\theta(\mathbf{x} | \mathbf{z})] \\ & + \mathbb{E}_{q^*(\mathbf{x})} \text{KL}[q_\phi(\mathbf{z} | \mathbf{x}) \parallel p^*(\mathbf{z})]. \end{aligned} \quad (4.5)$$

The first term of the ELBO is the (negative) ELL, defined as

$$\mathcal{L}_{\text{NELL}}(\boldsymbol{\theta}, \boldsymbol{\phi}) \triangleq \mathbb{E}_{q^*(\mathbf{x})q_\phi(\mathbf{z} | \mathbf{x})}[-\log p_\theta(\mathbf{x} | \mathbf{z})]. \quad (4.6)$$

It is easy to perform stochastic gradient-based optimisation of this term by applying the reparameterisation trick [127, 212],

$$\mathcal{L}_{\text{NELL}}(\boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbb{E}_{q^*(\mathbf{x})p(\epsilon)}[-\log p_\theta(\mathbf{x} | \mathcal{G}_\phi(\epsilon; \mathbf{x}))]. \quad (4.7)$$

However, the second term – the KL divergence between $q_\phi(\mathbf{z} | \mathbf{x})$ and implicit prior $p^*(\mathbf{z})$ – is not so straightforward. In particular, the KL divergence can be expressed as

$$\text{KL}[q_\phi(\mathbf{z} | \mathbf{x}) \parallel p^*(\mathbf{z})] \triangleq \mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x})}[\log r^*(\mathbf{z}; \mathbf{x})], \quad (4.8)$$

where $r^*(\mathbf{z}; \mathbf{x})$ is defined as the ratio of densities,

$$r^*(\mathbf{z}; \mathbf{x}) \triangleq \frac{q_\phi(\mathbf{z} | \mathbf{x})}{p^*(\mathbf{z})}. \quad (4.9)$$

The dependence on this density ratio is problematic since the prior $p^*(\mathbf{z})$ is implicit and cannot be evaluated directly. To overcome this, we resort to methods for approximating f -divergences between implicit distributions, which are inextricably tied to DRE [177, 251].

4.3.3 *Approximate Divergence Minimisation*

Although we are primarily interested in estimating the KL divergence of Equation (4.8), we give a generalised treatment that is applicable to all f -divergences [2, 47]. We denote a generic member of the family of f -divergences between distributions p and q as $\mathcal{D}_f [p \parallel q] \triangleq \mathbb{E}_p[f(q/p)]$, for some convex lower-semicontinuous function $f : \mathbb{R}_+ \rightarrow \mathbb{R}$.

Leveraging results from convex analysis, Nguyen, Wainwright, and Jordan [188] devise a variational lower bound that estimates an f -divergence through samples when either or both of the densities are unavailable. Nowozin, Cseke, and Tomioka [191] extend this framework to derive GAN objectives that minimise arbitrary f -divergences. These results underpin our methodology, and we restate a variant of it here for completeness.

Theorem 4.3.1 (Nguyen, Wainwright, and Jordan [188]). *Let f^* be the convex dual of f and \mathcal{R} a class of functions with codomains equivalent to the domain of f' . We have the following lower bound on the f -divergence between distributions $p(\mathbf{u})$ and $q(\mathbf{u})$,*

$$\mathcal{D}_f [p(\mathbf{u}) \parallel q(\mathbf{u})] \geq \max_{\hat{r} \in \mathcal{R}} \{ \mathbb{E}_{q(\mathbf{u})} [f'(\hat{r}(\mathbf{u}))] - \mathbb{E}_{p(\mathbf{u})} [f^*(f'(\hat{r}(\mathbf{u})))] \},$$

where equality is attained when $\hat{r}(\mathbf{u})$ is exactly the true density ratio $\hat{r}(\mathbf{u}) = q(\mathbf{u})/p(\mathbf{u})$.

Applying Theorem 4.3.1 to $p^*(\mathbf{z})$ and $q_\phi(\mathbf{z} | \mathbf{x}_n)$ for a given \mathbf{x}_n , and optimising over a class of functions indexed by parameters ω_n , we obtain the following lower bound on their divergence,

$$\mathcal{D}_f [p^*(\mathbf{z}) \parallel q_\phi(\mathbf{z} | \mathbf{x}_n)] \geq \max_{\omega_n} \left\{ \mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x}_n)} [f'(r_{\omega_n}(\mathbf{z}))] - \mathbb{E}_{p^*(\mathbf{z})} [f^*(f'(r_{\omega_n}(\mathbf{z})))] \right\}.$$

While this provides a way to estimate any f -divergence between implicit prior $p^*(\mathbf{z})$ and variational distribution $q_\phi(\mathbf{z} | \mathbf{x}_n)$ with only samples, it also requires us to optimise a separate density ratio estimator with parameters ω_n for each observed \mathbf{x}_n . Instead, as with the posterior approximation, we also amortise the density ratio estimator by conditioning on \mathbf{x} and optimising a single set of parameters α across all $\mathbf{x} \sim q^*(\mathbf{x})$. Accordingly, the estimator becomes $r_\alpha(\mathbf{z}; \mathbf{x})$, taking also \mathbf{x} as input. We now maximise an instance of the following generalised objective,

$$\mathcal{L}_f^{\text{latent}}(\alpha | \phi) \triangleq \mathbb{E}_{q^*(\mathbf{x})q_\phi(\mathbf{z} | \mathbf{x})} [f'(r_\alpha(\mathbf{z}; \mathbf{x}))] - \mathbb{E}_{q^*(\mathbf{x})p^*(\mathbf{z})} [f^*(f'(r_\alpha(\mathbf{z}; \mathbf{x})))] \quad (4.10)$$

Corollary 4.3.2. *We have the lower bound,*

$$\mathbb{E}_{q^*(\mathbf{x})} \mathcal{D}_f [p^*(\mathbf{z}) \parallel q_\phi(\mathbf{z} | \mathbf{x})] \geq \max_{\alpha} \mathcal{L}_f^{\text{latent}}(\alpha | \phi), \quad (4.11)$$

with equality at $r_\alpha(\mathbf{z}; \mathbf{x}) = r^*(\mathbf{z}; \mathbf{x})$.

DENSITY RATIO ESTIMATION OBJECTIVE. We write $\mathcal{L}_f^{\text{latent}}(\boldsymbol{\alpha} \mid \boldsymbol{\phi})$ to denote the DRE objective, wherein $\boldsymbol{\phi}$ is fixed, while $\boldsymbol{\alpha}$ is a free parameter that varies as this objective is *maximised*, thus tightening the bound of Equation (4.11) and the estimate of the density ratio $r_\alpha(\mathbf{z}; \mathbf{x})$.

DIVERGENCE MINIMISATION LOSS. Conversely, the *divergence minimisation* (DM) loss, denoted as $\mathcal{L}_f^{\text{latent}}(\boldsymbol{\phi} \mid \boldsymbol{\alpha})$, is *minimised* wrt $\boldsymbol{\phi}$ while $\boldsymbol{\alpha}$ remains fixed, thus approximately minimising the f -divergence. In theory, this should be symmetric to the DRE objective, $\mathcal{L}_f^{\text{latent}}(\boldsymbol{\phi} \mid \boldsymbol{\alpha}) \triangleq \mathcal{L}_f^{\text{latent}}(\boldsymbol{\alpha} \mid \boldsymbol{\phi})$. However, alternative settings are often used in practice to alleviate the problem of vanishing gradients, as we shall see in Section 4.5.

By applying Theorem 4.3.2 for the setting $f_{\text{KL}}(u) \triangleq u \log u$, we instantiate a lower bound on the KL divergence of Equation (4.8) in the following objective,

$$\mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\alpha} \mid \boldsymbol{\phi}) \triangleq \mathbb{E}_{q^*(\mathbf{x})q_\phi(\mathbf{z} \mid \mathbf{x})}[\log r_\alpha(\mathbf{z}; \mathbf{x})] - \mathbb{E}_{q^*(\mathbf{x})p^*(\mathbf{z})}[r_\alpha(\mathbf{z}; \mathbf{x}) - 1]. \quad (4.12)$$

As we discuss in Section 4.A, maximisation of the objective in Equation (4.12) is closely related to the KLIEP [250].

Now, we define the DM loss symmetrically to the DRE objective in Equation (4.12) – terms not involving $\boldsymbol{\phi}$ are omitted,

$$\begin{aligned} \mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\phi} \mid \boldsymbol{\alpha}) &\triangleq \mathbb{E}_{q^*(\mathbf{x})q_\phi(\mathbf{z} \mid \mathbf{x})}[\log r_\alpha(\mathbf{z}; \mathbf{x})] & (4.13) \\ &\approx \mathbb{E}_{q^*(\mathbf{x})\text{KL}}[q_\phi(\mathbf{z} \mid \mathbf{x}) \parallel p^*(\mathbf{z})]. \end{aligned}$$

Combined with the ELL, this estimate of the KL divergence yields an approximation to the ELBO where all terms are tractable. These objectives are summarised in the bi-level optimisation problem below,

$$\max_{\boldsymbol{\alpha}} \quad \mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\alpha} \mid \boldsymbol{\phi}), \quad (4.14a)$$

$$\min_{\boldsymbol{\phi}, \boldsymbol{\theta}} \quad \mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\phi} \mid \boldsymbol{\alpha}) + \mathcal{L}_{\text{NELL}}(\boldsymbol{\theta}, \boldsymbol{\phi}), \quad (4.14b)$$

thus concluding the reverse KL minimisation component of our VI framework.

4.4 SYMMETRIC JOINT-MATCHING VARIATIONAL INFERENCE

We now complete the remaining component of our VI framework. In the previous section, we gave an extension to classical VI, which is fundamentally concerned with approximating the exact posterior. Now, let us instead consider directly approximating the *exact joint* $p_\theta(\mathbf{x}, \mathbf{z})$ through a *variational joint* $q_\phi(\mathbf{x}, \mathbf{z})$.

4.4.1 Variational Joint

Recall that $q^*(\mathbf{x})$ denotes the empirical data distribution. We define a variational approximation to the exact joint distribution of Equation (4.1) as

$$q_\phi(\mathbf{x}, \mathbf{z}) \triangleq q_\phi(\mathbf{z} | \mathbf{x})q^*(\mathbf{x}). \quad (4.15)$$

We approximate the exact joint by seeking a variational joint closest in *symmetric* KL divergence, $\text{KL}_{\text{SYM}} [p_\theta(\mathbf{x}, \mathbf{z}) \| q_\phi(\mathbf{x}, \mathbf{z})]$, where

$$\text{KL}_{\text{SYM}} [p \| q] \triangleq \text{KL} [p \| q] + \text{KL} [q \| p]. \quad (4.16)$$

We first look at the reverse KL divergence ($\text{KL} [q \| p]$) term. When expanded, we see that it is equivalent to the negative ELBO up to additive constants,

$$\text{KL} [q_\phi(\mathbf{x}, \mathbf{z}) \| p_\theta(\mathbf{x}, \mathbf{z})] \triangleq \mathbb{E}_{q_\phi(\mathbf{x}, \mathbf{z})} [\log q_\phi(\mathbf{x}, \mathbf{z}) - \log p_\theta(\mathbf{x}, \mathbf{z})] \quad (4.17)$$

$$= \mathcal{L}_{\text{NELBO}}(\boldsymbol{\theta}, \boldsymbol{\phi}) - \mathbb{H}[q^*(\mathbf{x})], \quad (4.18)$$

where $\mathbb{H}[q^*(\mathbf{x})] \triangleq \mathbb{E}_{q^*(\mathbf{x})} [-\log q^*(\mathbf{x})]$ is the entropy of $q^*(\mathbf{x})$, a constant wrt parameters $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. Hence, minimising the KL divergence of Equation (4.17) can be reduced to minimising $\mathcal{L}_{\text{NELBO}}(\boldsymbol{\theta}, \boldsymbol{\phi})$ of Equation (4.5), without modification.

4.4.2 Forward KL Variational Objective

As for the forward KL divergence ($\text{KL} [p \| q]$) term, we have a similar expansion,

$$\text{KL} [p_\theta(\mathbf{x}, \mathbf{z}) \| q_\phi(\mathbf{x}, \mathbf{z})] \quad (4.19)$$

$$\triangleq \mathbb{E}_{p_\theta(\mathbf{x}, \mathbf{z})} [\log p_\theta(\mathbf{x}, \mathbf{z}) - \log q_\phi(\mathbf{x}, \mathbf{z})] \quad (4.20)$$

$$= \mathbb{E}_{p^*(\mathbf{z})p_\theta(\mathbf{x} | \mathbf{z})} [\log p_\theta(\mathbf{x} | \mathbf{z}) - \log q_\phi(\mathbf{x}, \mathbf{z})] - \mathbb{H}[p^*(\mathbf{z})]. \quad (4.21)$$

In analogy with the ELBO, we introduce a new variational objective that is minimised when the forward KL divergence of Equation (4.19) is minimised. First we define the recognition model analog to the marginal likelihood – the *marginal posterior*, or *aggregated posterior*, given by $q_\phi(\mathbf{z}) \triangleq \int q_\phi(\mathbf{z} | \mathbf{x})q^*(\mathbf{x})d\mathbf{x}$. It can be approximated by the *aggregate posterior lower bound* (APLBO). For consistency, we give its *negative*, written as

$$\begin{aligned} \mathcal{L}_{\text{NAPLBO}}(\boldsymbol{\theta}, \boldsymbol{\phi}) \triangleq & \mathbb{E}_{p^*(\mathbf{z})p_\theta(\mathbf{x} | \mathbf{z})} [-\log q_\phi(\mathbf{z} | \mathbf{x})] \\ & + \mathbb{E}_{p^*(\mathbf{z})} \text{KL} [p_\theta(\mathbf{x} | \mathbf{z}) \| q^*(\mathbf{x})]. \end{aligned} \quad (4.22)$$

Furthermore, minimising the KL divergence of Equation (4.19) can be reduced to minimising $\mathcal{L}_{\text{NAPLBO}}(\boldsymbol{\theta}, \boldsymbol{\phi})$,

$$\text{KL} [p_\theta(\mathbf{x}, \mathbf{z}) \| q_\phi(\mathbf{x}, \mathbf{z})] = \mathcal{L}_{\text{NAPLBO}}(\boldsymbol{\theta}, \boldsymbol{\phi}) - \mathbb{H}[p^*(\mathbf{z})].$$

The first term of the negative APLBO is the (negative) *expected log posterior* (ELP), defined as

$$\mathcal{L}_{\text{NELP}}(\boldsymbol{\theta}, \boldsymbol{\phi}) \triangleq \mathbb{E}_{p^*(\mathbf{z})p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})}[-\log q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})]. \quad (4.23)$$

We emphasise a key advantage of having considered the KL between the joint distributions instead of between the *posteriors*. Computing the *forward* KL divergence between the exact and approximate *posterior* distribution is problematic, since it requires evaluating expectations over the exact posterior $p_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})$, the intractability of which is the reason we appealed to approximate inference in the first place.

In contrast, the forward KL divergence between the exact and approximate *joint* poses no such difficulties – we are able to sidestep the dependency on the exact posterior by expanding it into the form of Equation (4.21). Furthermore, as with the ELBO, we can perform stochastic gradient-based optimisation of the ELP term by applying the same reparameterisation trick as in Equation (4.7).

Now, the KL divergence term of the APLBO in Equation (4.22) can also be expressed as the expected logarithm of a density ratio $r^*(\mathbf{x}; \mathbf{z}) \triangleq p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})/q^*(\mathbf{x})$ that involves an intractable density $q^*(\mathbf{x})$ – the empirical data distribution. To overcome this, we adopt the same approach as outlined in Section 4.3.3. Namely, we apply Theorem 4.3.1 to $q^*(\mathbf{x})$ and $p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}^*)$, and fit an amortised density ratio estimator $r_{\boldsymbol{\beta}}(\mathbf{x}; \mathbf{z})$ to $r^*(\mathbf{x}; \mathbf{z})$ by maximising an instance of the generalised objective,

$$\begin{aligned} \mathcal{L}_f^{\text{observed}}(\boldsymbol{\beta} | \boldsymbol{\theta}) \triangleq & \mathbb{E}_{p^*(\mathbf{z})p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})}[f'(r_{\boldsymbol{\beta}}(\mathbf{x}; \mathbf{z}))] \\ & - \mathbb{E}_{p^*(\mathbf{z})q^*(\mathbf{x})}[f^*(f'(r_{\boldsymbol{\beta}}(\mathbf{x}; \mathbf{z})))]. \end{aligned} \quad (4.24)$$

Corollary 4.4.1. *We have the lower bound,*

$$\mathbb{E}_{p^*(\mathbf{z})}\mathcal{D}_f[q^*(\mathbf{x}) \| p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})] \geq \max_{\boldsymbol{\beta}} \mathcal{L}_f^{\text{observed}}(\boldsymbol{\beta} | \boldsymbol{\theta}), \quad (4.25)$$

with equality at $r_{\boldsymbol{\beta}}(\mathbf{x}; \mathbf{z}) = r^*(\mathbf{x}; \mathbf{z})$.

By applying Theorem 4.4.1 with the previously defined $f_{\text{KL}}(u)$, we obtain lower bound objective $\mathcal{L}_{\text{KL}}^{\text{observed}}(\boldsymbol{\beta} | \boldsymbol{\theta})$ on the KL divergence term in Equation (4.22), and a corresponding DM loss $\mathcal{L}_{\text{KL}}^{\text{observed}}(\boldsymbol{\theta} | \boldsymbol{\beta})$, analogous to the definitions of $\mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\alpha} | \boldsymbol{\phi})$ and $\mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\phi} | \boldsymbol{\alpha})$ in Equations (4.12) and (4.13), respectively. See Table 4.B.3 for a summary of explicit definitions.

Hence, in addition to the bi-level optimisation problems of Equation (4.14) we have,

$$\max_{\boldsymbol{\beta}} \quad \mathcal{L}_{\text{KL}}^{\text{observed}}(\boldsymbol{\beta} | \boldsymbol{\theta}), \quad (4.26a)$$

$$\min_{\boldsymbol{\phi}, \boldsymbol{\theta}} \quad \mathcal{L}_{\text{KL}}^{\text{observed}}(\boldsymbol{\theta} | \boldsymbol{\beta}) + \mathcal{L}_{\text{NELP}}(\boldsymbol{\theta}, \boldsymbol{\phi}). \quad (4.26b)$$

As shown, the minimisations in Equations (4.14b) and (4.26b) corresponds to minimisation of the symmetric KL over the joints $\text{KL}_{\text{SYM}}[p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}) \| q_{\boldsymbol{\phi}}(\mathbf{x}, \mathbf{z})]$,

while the maximisations in Equations (4.14a) and (4.26a) approximates the divergences, or more precisely, the density ratios involving implicit distributions.

4.5 CYCLEGAN AS A SPECIAL CASE

In this section, we demonstrate that CYCLEGAN methods [125, 302] can be instantiated under our proposed VI framework.

4.5.1 Basic CycleGAN Framework

To address the problem of unpaired image-to-image translation as described in Section 4.2.2, the CYCLEGAN model learns two mappings $\mu_\theta : \mathbf{z} \mapsto \mathbf{x}$ and $\mathbf{m}_\phi : \mathbf{x} \mapsto \mathbf{z}$ by optimising two complementary classes of objectives.

DISTRIBUTION MATCHING. The first are the adversarial objectives, which help match the output of mapping μ_θ to the empirical distribution $q^*(\mathbf{x})$, and the output of \mathbf{m}_ϕ to $p^*(\mathbf{z})$. In particular, for mapping \mathbf{m}_ϕ , this involves introducing a discriminator \mathbf{D}_α and the saddle-point adversarial objective,

$$\ell_{\text{GAN}}^{\text{reverse}}(\alpha | \phi) \triangleq \mathbb{E}_{p^*(\mathbf{z})}[\log \mathbf{D}_\alpha(\mathbf{z})] + \mathbb{E}_{q^*(\mathbf{x})}[\log(1 - \mathbf{D}_\alpha(\mathbf{m}_\phi(\mathbf{x})))] \quad (4.27)$$

while minimising it wrt parameters ϕ . This encourages \mathbf{m}_ϕ to produce realistic outputs $\mathbf{z} = \mathbf{m}_\phi(\mathbf{x})$, $\mathbf{x} \sim q^*(\mathbf{x})$ which, to the discriminator \mathbf{D}_α , are “indistinguishable” from $\mathbf{z}^* \sim p^*(\mathbf{z})$. A similar adversarial objective is defined for mapping μ_θ ,

$$\ell_{\text{GAN}}^{\text{forward}}(\beta | \theta) \triangleq \mathbb{E}_{p^*(\mathbf{x})}[\log \mathbf{D}_\beta(\mathbf{x})] + \mathbb{E}_{p^*(\mathbf{z})}[\log(1 - \mathbf{D}_\beta(\mu_\theta(\mathbf{z})))] \quad (4.28)$$

CYCLE-CONSISTENCY. Next are the cycle-consistency losses, which enforce tight correspondence between domains by ensuring that reconstruction $\mathbf{x}' = \mu_\theta(\mathbf{m}_\phi(\mathbf{x}))$ is close to the input \mathbf{x} , and likewise for $\mathbf{m}_\phi(\mu_\theta(\mathbf{z}))$. This is achieved by minimising a reconstruction loss,

$$\ell_{\text{CONST}}^{\text{reverse}}(\theta, \phi) \triangleq \mathbb{E}_{q^*(\mathbf{x})}[\|\mathbf{x} - \mu_\theta(\mathbf{m}_\phi(\mathbf{x}))\|_\rho^\rho] \quad (4.29)$$

where $\|\cdot\|_\rho$ denotes the ℓ_ρ -norm. A similar loss $\ell_{\text{CONST}}^{\text{forward}}(\theta, \phi)$ is defined for the reconstruction of \mathbf{z} ,

$$\ell_{\text{CONST}}^{\text{forward}}(\theta, \phi) \triangleq \mathbb{E}_{p^*(\mathbf{z})}[\|\mathbf{z} - \mathbf{m}_\phi(\mu_\theta(\mathbf{z}))\|_\rho^\rho] \quad (4.30)$$

These objectives are summarised in the following set of optimisation problems,

$$\max_{\alpha} \ell_{\text{GAN}}^{\text{reverse}}(\alpha | \phi), \quad \max_{\beta} \ell_{\text{GAN}}^{\text{forward}}(\beta | \theta), \quad (4.31a)$$

$$\min_{\phi, \theta} \ell_{\text{GAN}}^{\text{reverse}}(\phi | \alpha) + \ell_{\text{CONST}}^{\text{reverse}}(\theta, \phi), \quad (4.31b)$$

$$\min_{\phi, \theta} \ell_{\text{GAN}}^{\text{forward}}(\theta | \beta) + \ell_{\text{CONST}}^{\text{forward}}(\theta, \phi). \quad (4.31c)$$

We now highlight the correspondences between these objectives and those of our proposed VI framework, as summarised in the optimisation problems of Equations (4.14) and (4.26).

4.5.2 Cycle-consistency as Conditional Entropy Maximisation

We now demonstrate that minimising the cycle-consistency losses corresponds to maximising the expected log likelihood and variational posterior of Equations (4.6) and (4.23). This can be shown by instantiating specific classes of $p_{\theta}(\mathbf{x} | \mathbf{z})$ and $q_{\phi}(\mathbf{z} | \mathbf{x})$ that recover $\ell_{\text{CONST}}^{\text{reverse}}(\theta, \phi)$ and $\ell_{\text{CONST}}^{\text{forward}}(\theta, \phi)$ from $\mathcal{L}_{\text{NELL}}(\theta, \phi)$ and $\mathcal{L}_{\text{NELP}}(\theta, \phi)$, respectively.

Proposition 4.5.1. *Consider a typical case where the likelihood and the posterior approximation are both Gaussians,*

$$p_{\theta}(\mathbf{x} | \mathbf{z}) \triangleq \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_{\theta}(\mathbf{z}), \tau^{-1}\mathbf{I}), \quad q_{\phi}(\mathbf{z} | \mathbf{x}) \triangleq \mathcal{N}(\mathbf{z} | \mathbf{m}_{\phi}(\mathbf{x}), t^{-1}\mathbf{I}).$$

In the limit as the posterior precision t tends to ∞ , $\mathcal{L}_{\text{NELL}}(\theta, \phi)$ approaches $\ell_{\text{CONST}}^{\text{reverse}}(\theta, \phi)$ for $\rho = 2$, up to constants¹. More precisely,

$$\mathcal{L}_{\text{NELL}}(\theta, \phi) \rightarrow \frac{\tau}{2} \mathcal{L}_{\text{CONST}}^{\text{reverse}}(\theta, \phi) + \text{const}, \quad \text{as } t \rightarrow \infty$$

Similarly, we have,

$$\mathcal{L}_{\text{NELP}}(\theta, \phi) \rightarrow \frac{t}{2} \mathcal{L}_{\text{CONST}}^{\text{forward}}(\theta, \phi) + \text{const}, \quad \text{as } \tau \rightarrow \infty$$

Proof. First, note the generative mappings underlying the given Gaussian likelihood and approximate posterior are

$$\begin{aligned} \mathbf{z} &\sim p_{\theta}(\mathbf{x} | \mathbf{z}) \triangleq \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_{\theta}(\mathbf{z}), \tau^{-1}\mathbf{I}), \\ \Leftrightarrow \mathbf{z} &= \mathcal{F}_{\theta}(\boldsymbol{\xi}; \mathbf{z}) \triangleq \boldsymbol{\mu}_{\theta}(\mathbf{z}) + \tau^{-\frac{1}{2}}\boldsymbol{\xi}, \quad \boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \end{aligned}$$

and,

$$\begin{aligned} \mathbf{x} &\sim q_{\phi}(\mathbf{z} | \mathbf{x}) \triangleq \mathcal{N}(\mathbf{z} | \mathbf{m}_{\phi}(\mathbf{x}), t^{-1}\mathbf{I}), \\ \Leftrightarrow \mathbf{x} &= \mathcal{G}_{\phi}(\boldsymbol{\epsilon}; \mathbf{x}) \triangleq \mathbf{m}_{\phi}(\mathbf{x}) + t^{-\frac{1}{2}}\boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \end{aligned}$$

¹ we obtain the same result for the case $\rho = 1$ by instead setting both the likelihood and approximate posterior to be Laplace distributions.

respectively. Thus, expanding out $\mathcal{L}_{\text{NELL}}(\boldsymbol{\theta}, \boldsymbol{\phi})$, we have

$$\begin{aligned}
\mathcal{L}_{\text{NELL}}(\boldsymbol{\theta}, \boldsymbol{\phi}) &= \mathbb{E}_{q^*(\mathbf{x})q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})}[-\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})] \\
&= \mathbb{E}_{q^*(\mathbf{x})p(\boldsymbol{\epsilon})}[-\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathcal{G}_{\boldsymbol{\phi}}(\boldsymbol{\epsilon};\mathbf{x}))] \\
&= \frac{\tau}{2}\mathbb{E}_{q^*(\mathbf{x})p(\boldsymbol{\epsilon})}[\|\mathbf{x} - \boldsymbol{\mu}_{\boldsymbol{\theta}}(\mathcal{G}_{\boldsymbol{\phi}}(\boldsymbol{\epsilon};\mathbf{x}))\|_2^2] + \frac{D}{2}\log\frac{2\pi}{\tau} \\
&= \frac{\tau}{2}\mathbb{E}_{q^*(\mathbf{x})p(\boldsymbol{\epsilon})}[\|\mathbf{x} - \boldsymbol{\mu}_{\boldsymbol{\theta}}(\mathbf{m}_{\boldsymbol{\phi}}(\mathbf{x}) + t^{-\frac{1}{2}}\boldsymbol{\epsilon})\|_2^2] + \text{const} \\
&\rightarrow \frac{\tau}{2}\mathbb{E}_{q^*(\mathbf{x})}[\|\mathbf{x} - \boldsymbol{\mu}_{\boldsymbol{\theta}}(\mathbf{m}_{\boldsymbol{\phi}}(\mathbf{x}))\|_2^2] + \text{const}, \quad \text{as } t \rightarrow \infty \\
&= \frac{\tau}{2}\mathcal{L}_{\text{CONST}}^{\text{reverse}}(\boldsymbol{\theta}, \boldsymbol{\phi}) + \text{const}.
\end{aligned}$$

A similar analysis can be carried out for $\mathcal{L}_{\text{NELP}}(\boldsymbol{\theta}, \boldsymbol{\phi})$ and its deterministic counterpart $\ell_{\text{CONST}}^{\text{forward}}(\boldsymbol{\theta}, \boldsymbol{\phi})$. \square

Hence, the cycle-consistency losses can be seen as special cases of the ELL and ELP with *degenerate* conditional distributions. Furthermore, this sheds new light on the roles of the cycle-consistency losses. For example, similar to the ELL, the reverse consistency loss encourages the conditional $q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})$ to place its mass on configurations of latent variables that can explain, or in this case, *represent* the data well.

4.5.3 Distribution Matching as Approximate Divergence Minimisation

We now discuss how the adversarial objectives $\ell_{\text{GAN}}^{\text{reverse}}(\boldsymbol{\alpha}|\boldsymbol{\phi})$ and $\ell_{\text{GAN}}^{\text{forward}}(\boldsymbol{\beta}|\boldsymbol{\theta})$ relate to the KL variational lower bounds of our framework, $\mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\alpha}|\boldsymbol{\phi})$ and $\mathcal{L}_{\text{KL}}^{\text{observed}}(\boldsymbol{\beta}|\boldsymbol{\theta})$, respectively. To reduce clutter, we restrict our discussion to the reverse objective $\ell_{\text{GAN}}^{\text{reverse}}(\boldsymbol{\alpha}|\boldsymbol{\phi})$, as the same reasoning readily applies to the forward $\ell_{\text{GAN}}^{\text{forward}}(\boldsymbol{\beta}|\boldsymbol{\theta})$.

4.5.3.1 As Density Ratio Estimation by Probabilistic Classification

Firstly, the connections between GANS, divergence minimisation and DRE are well-established [177, 191, 251]. Although $\ell_{\text{GAN}}^{\text{reverse}}(\boldsymbol{\alpha}|\boldsymbol{\phi})$ is a scoring rule for probabilistic classification [85], one can readily show that it can also be subsumed as an instance of the generalised variational lower bound $\mathcal{L}_f^{\text{latent}}(\boldsymbol{\alpha}|\boldsymbol{\phi})$. Furthermore, similar to $\mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\alpha}|\boldsymbol{\phi})$, maximising $\ell_{\text{GAN}}^{\text{reverse}}(\boldsymbol{\alpha}|\boldsymbol{\phi})$ corresponds estimating the intractable density ratio $r^*(\mathbf{z};\mathbf{x})$ of Equation (4.9).

Lemma 4.5.2. *By setting $f_{\text{GAN}}(u) = u \log u - (u + 1) \log(u + 1)$ in the generalised objective $\mathcal{L}_f^{\text{latent}}(\boldsymbol{\alpha}|\boldsymbol{\phi})$ of Equation (4.10), we instantiate the objective*

$$\begin{aligned}
\mathcal{L}_{\text{GAN}}^{\text{reverse}}(\boldsymbol{\alpha}|\boldsymbol{\phi}) &\triangleq \mathbb{E}_{q^*(\mathbf{x})p^*(\mathbf{z})}[\log \mathcal{D}_{\boldsymbol{\alpha}}(\mathbf{z};\mathbf{x})] \\
&\quad + \mathbb{E}_{q^*(\mathbf{x})q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})}[\log(1 - \mathcal{D}_{\boldsymbol{\alpha}}(\mathbf{z};\mathbf{x}))],
\end{aligned} \tag{4.32}$$

where $\mathcal{D}_{\boldsymbol{\alpha}}(\mathbf{z};\mathbf{x}) \triangleq 1 - \sigma(\log r_{\boldsymbol{\alpha}}(\mathbf{z};\mathbf{x}))$, and σ is the logistic sigmoid function.

Proof. To instantiate $\mathcal{L}_{\text{GAN}}^{\text{reverse}}(\alpha | \phi)$ of Equation (4.32), it suffices to show that $-f_{\text{GAN}}^*(f'_{\text{GAN}}(r_\alpha(\mathbf{z}; \mathbf{x}))) = \log \mathcal{D}_\alpha(\mathbf{z}; \mathbf{x})$ and $f'_{\text{GAN}}(r_\alpha(\mathbf{z}; \mathbf{x})) = \log(1 - \mathcal{D}_\alpha(\mathbf{z}; \mathbf{x}))$, where $\mathcal{D}_\alpha(\mathbf{z}; \mathbf{x}) \triangleq 1 - \sigma(\log r_\alpha(\mathbf{z}; \mathbf{x}))$. First we compute the first derivative f'_{GAN} and the convex dual f_{GAN}^* of f_{GAN} , which involve straightforward calculations,

$$f'_{\text{GAN}}(u) = \log \sigma(\log u), \quad f_{\text{GAN}}^*(t) = -\log(1 - \exp t).$$

Thus, the composition $(f_{\text{GAN}}^* \circ f'_{\text{GAN}}) : u \mapsto f_{\text{GAN}}^*(f'_{\text{GAN}}(u))$ can be simplified as

$$f_{\text{GAN}}^*(f'_{\text{GAN}}(u)) = -\log(1 - \exp f'_{\text{GAN}}(u)) = -\log(1 - \sigma(\log u)).$$

Applying f'_{GAN} and $f_{\text{GAN}}^* \circ f'_{\text{GAN}}$ to $r_\alpha(\mathbf{z}; \mathbf{x})$, we have

$$f'_{\text{GAN}}(r_\alpha(\mathbf{z}; \mathbf{x})) = \log \sigma(\log r_\alpha(\mathbf{z}; \mathbf{x})) = \log(1 - \mathcal{D}_\alpha(\mathbf{z}; \mathbf{x})),$$

and

$$f_{\text{GAN}}^*(f'_{\text{GAN}}(r_\alpha(\mathbf{z}; \mathbf{x}))) = -\log(1 - \sigma(\log r_\alpha(\mathbf{z}; \mathbf{x}))) = -\log \mathcal{D}_\alpha(\mathbf{z}; \mathbf{x}),$$

respectively, as required. \square

Lemma 4.5.3. *By specifying a discriminator $\mathcal{D}_\alpha(\mathbf{z}; \mathbf{x}) = \mathbf{D}_\alpha(\mathbf{z})$ that ignores auxiliary input \mathbf{x} , and mapping $\mathcal{G}_\phi(\epsilon; \mathbf{x}) = \mathbf{m}_\phi(\mathbf{x})$ that ignores noise input ϵ , $\mathcal{L}_{\text{GAN}}^{\text{reverse}}(\alpha | \phi)$ reduces to $\ell_{\text{GAN}}^{\text{reverse}}(\alpha | \phi)$.*

Proof. Through reparameterisation of $q_\phi(\mathbf{z} | \mathbf{x})$, we have

$$\begin{aligned} \mathcal{L}_{\text{GAN}}^{\text{reverse}}(\alpha | \phi) &= \mathbb{E}_{q^*(\mathbf{x})p^*(\mathbf{z})}[\log \mathcal{D}_\alpha(\mathbf{z}; \mathbf{x})] \\ &\quad + \mathbb{E}_{q^*(\mathbf{x})p(\epsilon)}[\log(1 - \mathcal{D}_\alpha(\mathcal{G}_\phi(\epsilon; \mathbf{x}); \mathbf{x}))]. \end{aligned}$$

By specifying a discriminator $\mathcal{D}_\alpha(\mathbf{z}; \mathbf{x}) = \mathbf{D}_\alpha(\mathbf{z})$ that ignores auxiliary input \mathbf{x} , and mapping $\mathcal{G}_\phi(\epsilon; \mathbf{x}) = \mathbf{m}_\phi(\mathbf{x})$ that ignores noise input ϵ , this reduces to

$$\begin{aligned} \mathcal{L}_{\text{GAN}}^{\text{reverse}}(\alpha | \phi) &= \mathbb{E}_{p^*(\mathbf{z})}[\log \mathbf{D}_\alpha(\mathbf{z})] + \mathbb{E}_{q^*(\mathbf{x})}[\log(1 - \mathbf{D}_\alpha(\mathbf{m}_\phi(\mathbf{x})))] \\ &= \ell_{\text{GAN}}^{\text{reverse}}(\alpha | \phi), \end{aligned}$$

as required. \square

Proposition 4.5.4. *The reverse adversarial objective $\ell_{\text{GAN}}^{\text{reverse}}(\alpha | \phi)$ can be subsumed as an instance of the generalised variational lower bound $\mathcal{L}_f^{\text{latent}}(\alpha | \phi)$.*

Theorem 4.5.4 follows directly from Theorems 4.5.2 and 4.5.3.

Now, by Theorem 4.3.2, objective $\mathcal{L}_{\text{GAN}}^{\text{reverse}}(\alpha | \phi)$ is maximised exactly when $r_\alpha(\mathbf{z}; \mathbf{x}) = r^*(\mathbf{z}; \mathbf{x})$. Hence, we can interpret $\mathcal{L}_{\text{GAN}}^{\text{reverse}}(\alpha | \phi)$ as an objective for density-ratio estimation based on probabilistic classification, while $\mathcal{L}_{\text{KL}}^{\text{latent}}(\alpha | \phi)$ is an objective based on KLIEP.

Now, the default choice of DM loss is $\mathcal{L}_{\text{GAN}_A}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha}) \triangleq \mathcal{L}_{\text{GAN}}^{\text{reverse}}(\boldsymbol{\alpha} | \boldsymbol{\phi})$. Omitting terms not involving $\boldsymbol{\phi}$, this is given by

$$\mathcal{L}_{\text{GAN}_A}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha}) \triangleq \mathbb{E}_{q^*(\mathbf{x})q_{\boldsymbol{\phi}}(\mathbf{z} | \mathbf{x})}[\log(1 - \mathcal{D}_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x}))]. \quad (4.33)$$

Unlike $\mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\phi} | \boldsymbol{\alpha})$, minimising $\mathcal{L}_{\text{GAN}_A}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha})$ does not minimise the KL divergence of Equation (4.8). Hence, the minimisation problem of Equation (4.31b) does not correspond to that of Equation (4.14b), and so does not maximise the ELBO, or any known VI objective.

4.5.3.2 Recovering KL Through Alternative Divergence Minimisation Losses

Although the default choice of DM loss does not yield a tight correspondence to VI, the existing CYCLEGAN frameworks – and indeed most GAN-based approaches – arbitrarily select an alternative DM loss that avoids vanishing gradients, and work well in practice. Hence, one need only choose an alternative that *does* correspond to minimising the KL divergence of Equation (4.8).

Firstly, of the CYCLEGAN methods, Kim et al. [125] adopt the widely-used DM loss originally suggested by Goodfellow et al. [86],

$$\mathcal{L}_{\text{GAN}_B}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha}) \triangleq \mathbb{E}_{q^*(\mathbf{x})q_{\boldsymbol{\phi}}(\mathbf{z} | \mathbf{x})}[-\log \mathcal{D}_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x})], \quad (4.34)$$

while [302] optimise the Least-Squares GAN (LSGAN) objectives of [162].

Consider the *combination* of losses $\mathcal{L}_{\text{GAN}_A}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha})$ and $\mathcal{L}_{\text{GAN}_B}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha})$,

$$\begin{aligned} \mathcal{L}_{\text{GAN}_C}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha}) &\triangleq \mathcal{L}_{\text{GAN}_A}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha}) + \mathcal{L}_{\text{GAN}_B}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha}) \\ &= \mathbb{E}_{q^*(\mathbf{x})q_{\boldsymbol{\phi}}(\mathbf{z} | \mathbf{x})} \left[-\log \frac{\mathcal{D}_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x})}{1 - \mathcal{D}_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x})} \right]. \end{aligned} \quad (4.35)$$

Proposition 4.5.5. *We have $\mathcal{L}_{\text{GAN}_C}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha}) = \mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\phi} | \boldsymbol{\alpha})$.*

Theorem 4.5.5 was originally noted by Sønderby et al. [239] and is shown below.

Proof. Expanding out $\mathcal{L}_{\text{GAN}_C}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha})$, we have

$$\begin{aligned} \mathcal{L}_{\text{GAN}_C}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha}) &= \mathbb{E}_{q^*(\mathbf{x})q_{\boldsymbol{\phi}}(\mathbf{z} | \mathbf{x})} \left[-\log \frac{\mathcal{D}_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x})}{1 - \mathcal{D}_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x})} \right] \\ &= \mathbb{E}_{q^*(\mathbf{x})q_{\boldsymbol{\phi}}(\mathbf{z} | \mathbf{x})} \left[\log \frac{\sigma(\log r_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x}))}{1 - \sigma(\log r_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x}))} \right] \\ &= \mathbb{E}_{q^*(\mathbf{x})q_{\boldsymbol{\phi}}(\mathbf{z} | \mathbf{x})} [\log r_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x})] \triangleq \mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\phi} | \boldsymbol{\alpha}). \end{aligned}$$

Hence, $\mathcal{L}_{\text{GAN}_C}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha}) = \mathcal{L}_{\text{KL}}^{\text{latent}}(\boldsymbol{\phi} | \boldsymbol{\alpha})$ as required. \square

Thus, for the setting of the DM loss $\mathcal{L}_{\text{GAN}_C}^{\text{reverse}}(\boldsymbol{\phi} | \boldsymbol{\alpha})$, the minimisation problem of Equation (4.31b) corresponds to that of Equation (4.14b), and thus maximises the ELBO. This is equivalent to fitting the density ratio estimator $r_{\boldsymbol{\alpha}}(\mathbf{z}; \mathbf{x})$ by maximising the objective $\mathcal{L}_{\text{GAN}}^{\text{reverse}}(\boldsymbol{\alpha} | \boldsymbol{\phi})$

instead of $\mathcal{L}_{\text{KL}}^{\text{latent}}(\alpha | \phi)$, and plugging it back into $\mathcal{L}_{\text{KL}}^{\text{latent}}(\phi | \alpha)$ to approximately minimise the KL divergence of Equation (4.8). Such an approach is prevalent among existing implicit VI methods [110, 172, 202, 267]

SUMMARY OF THEORETICAL CONNECTIONS We have that the cycle-consistency losses are a specific instance of the ELL and ELP, while the adversarial objectives are a specific instance of the variational lower bound for divergence estimation, the maximisation of which can be seen as density ratio estimation by probabilistic classification. By explicitly setting the corresponding divergence minimisation loss such that it leads to minimisation of the required KL divergence terms in the ELBO and APLBO, we subsume the CYCLEGAN model under our proposed VI framework. See Section 4.B for a succinct summary of the relationships.

4.6 RELATED WORK

The work presented in this chapter is closely related to prior efforts to extend the scope of VI to implicit distributions. A recurring theme throughout this line of work is approximation of the ELBO by exploiting the formal connection between density-ratio estimation and GANS [177, 271]. The major axis of variation lies in the choice of the target density ratio being estimated, as dictated by the problem setting. Makhzani et al. [160] and Mescheder, Nowozin, and Geiger [172, AVB] estimate the density ratio $q_{\phi}(\mathbf{z} | \mathbf{x})/p(\mathbf{z})$ in order to allow for expressive sample-based posterior approximations $q_{\phi}(\mathbf{z} | \mathbf{x})$. This corresponds to the reverse KL minimisation component of our approach, in which we accommodate implicit priors $p(\mathbf{z})$.

Similar to BIGAN [62] and adversarially learned inference (ALI) [59], Tran, Ranganath, and Blei [267, LFVI] match a variational joint to an exact joint distribution by estimating the density ratio $p_{\theta}(\mathbf{x}, \mathbf{z})/q_{\phi}(\mathbf{x}, \mathbf{z})$ and use it to approximately minimise the KL divergence. Although this formulation sidesteps the requirement of having *any* tractable densities, their focus is on inference for models with intractable likelihoods $p_{\theta}(\mathbf{x} | \mathbf{z})$ and on incorporating the implicit posteriors *à la* AVB. In contrast, in our framework, the joint’s intractability instead stems from the implicit prior $p^*(\mathbf{z})$. While we also approximate the exact joint, we do so by minimising a *symmetric* KL divergence. Furthermore, since both $p_{\theta}(\mathbf{x} | \mathbf{z})$ and $q_{\phi}(\mathbf{z} | \mathbf{x})$ are prescribed, we evaluate them explicitly as part of our loss functions and estimate a different set of density ratios. This closely resembles the approach of Pu et al. [202], which also minimises the symmetric KL divergence between the joints. However, the focus of their method is not on implicit distributions and thus specifies a different set of losses than ours – one that requires solving more complicated density ratio estimation problems. More

importantly, their method does not yield a tight correspondence to CYCLEGAN models.

A consequence of solely minimising the forward KL (as in frameworks like ALI) rather than minimising the symmetric KL (as in our framework) is *non-identifiability*. This issue has been addressed by Li et al. [144] who proposed a conditional-entropy regulariser to ALI’s objective. Although Li et al. [144] examine the link between their method and CYCLEGAN, unlike our work, the relationship is not made explicit in a mathematically precise manner. Additionally, we derive our full objective from the perspective of approximate Bayesian inference. More recently, Chen et al. [31] also highlight issues associated with ALI concerning the quality of data generated from the inferred latent variables. They propose a symmetric VAE that simultaneously inherits the realistic image generation capabilities of adversarial approaches while overcoming the asymmetry limitations of the forward KL divergence inherent in standard VAEs. Additionally, unlike our approach, they do not provide an explicit relationship with CYCLEGANs.

Finally, similar to InfoGAN [35] and VEEGAN [245], the forward KL minimisation component of our method also optimises a model of the latent variables, which is reminiscent of the wake-sleep algorithm for training Helmholtz machines [52]. This is discussed further by Hu et al. [109], who provide a comprehensive treatment of the links between the work mentioned in this section and importantly, the symmetric perspective of *generation* and *recognition* that underpins our approach.

4.7 EXPERIMENTS

SYNTHETIC DATA. First we consider the problem of reducing the dimensionality of the MNIST dataset to a 2D latent space, wherein the prior distribution on the latent representations is specified by its samples (shown in Figure 4.2a). This “banana-shaped distribution” is a commonly used testbed for adaptive MCMC methods [93, 264]. Its samples can be generated by drawing from a bivariate Gaussian with unit variances and correlation $\rho = 0.95$, and transforming them through mapping $H(z_1, z_2) \triangleq [z_1, z_2 - z_1^2 - 1]^T$. While the density of this distribution can be computed, it is withheld from our algorithm and used only in the VAE baseline, which does not permit implicit distributions.


Qualitative and quantitative results are given Figure 4.2 and Table 4.1, which demonstrate superior performance to competing approaches. Observe that instances of the various digit classes are disentangled in this latent space, while still closely matching the shape of the prior distribution, despite having only access to its samples. The resulting manifold of reconstructions is depicted in Figure 4.2c.

In Table 4.1, we report the mean-squared error (MSE) on the reconstructions of observations from the held-out test set and benchmark

Table 4.1: Mean-squared errors of reconstructions.

METHOD	MSE \mathbf{z}	MSE \mathbf{x}
SJMVI (OURS)	0.17	0.04
VAE [127]	0.88	0.04
AVB [172]	0.29	0.04

against VAE / AVB. Also, for the joint approximation to properly match the support of the exact joint, the latent codes should also be representable by its corresponding observation. Hence, we also report the MSE between samples from the prior and their reconstructions. While we find no improvements on reconstruction quality of observations, our method significantly outperforms others in reconstructing latent codes, suggesting our method has greater capacity to faithfully approximate the exact joint.




(a) 10k samples from implicit prior $p^*(\mathbf{z})$. (b) Mean of $q_\phi(\mathbf{z} | \mathbf{x})$ for every \mathbf{x} from held-out test set of size 10k, colored by digit class. (c) Mean of $p_\theta(\mathbf{x} | \mathbf{z})$ for 20×20 values of \mathbf{z} along a uniform grid.


Figure 4.2: Visualisation of 2D latent space and the corresponding observed space manifold. Instances of the various digit classes are disentangled in this latent space, while still closely matching the shape of the prior distribution, despite having only access to its samples.

IMAGE-TO-IMAGE TRANSLATION. We apply our method to the task of transferring features between images of faces on the CelebA dataset [152]. We consider the case where one feature differs between domains. In particular, distributions $q^*(\mathbf{x})$ and $q^*(\mathbf{z})$ are specified by images of women with blond and black hair, respectively. We specify both $p_\theta(\mathbf{x} | \mathbf{z})$ and $q_\phi(\mathbf{z} | \mathbf{x})$ as a Laplace distribution, with fixed variance, and mean functions $\mu_\theta(\mathbf{z})$ and $\mathbf{m}_\phi(\mathbf{x})$ defined by neural networks. Their architectures, as well as those of discriminators $\mathcal{D}_\alpha(\mathbf{z}; \mathbf{x})$ and $\mathcal{D}_\beta(\mathbf{x}; \mathbf{z})$ are defined in the same way as in [302]. In Figure 4.3, we show the outputs of the mean functions on samples from a hold-out test set, after training for 10 epochs. From Figure 4.3b, we see that

given datapoints x (top) we're able to learn a posterior over latent representations z in the other domain (mean is shown in middle). Furthermore, these latent representations are configured so as to maximise the likelihood of observing the original data, as evident from the reconstructions (bottom). Refer to appendix Figure 4.4 for qualitative results produced by the CYCLEGAN baseline approach [125, 302].



(a) Samples from the prior (top), mean of the likelihood (middle), and the mean reconstruction (bottom).




(b) Samples from the data (top), mean of the posterior (middle), and the mean reconstruction (bottom).

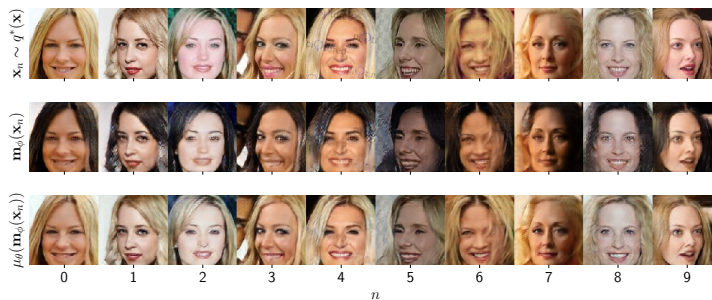
Figure 4.3: Image-to-image translation (blond to black hair) on CelebA dataset [152] performed by our proposed approach.

4.8 SUMMARY

In this chapter, we've provided a theoretical treatment of the link between CYCLEGANS and approximate Bayesian inference. In short, samples from the two domains correspond respectively to those drawn from the data and implicit prior distribution in a implicit LVM (ILVM). Parameter learning in CYCLEGANS corresponds to approximate inference in this ILVM under our proposed VI framework. The forward and reverse mappings in CYCLEGANS arise naturally in the generative and recognition models, while the cycle-consistency constraints correspond to their log probabilities, and the adversarial losses are approximations to an f -divergence. By lifting the limitations of prescribed prior distributions in favour of arbitrarily flexible implicit distributions, we



(a) Samples from the prior (top), output of mapping (middle), and the reconstruction (bottom).



(b) Samples from the data (top), output of mapping (middle), and the reconstruction (bottom).

Figure 4.4: Image-to-image translation (blond to black hair) on CelebA dataset [152] performed by the baseline CYCLEGAN method.

can discover different perspectives on existing learning methods and provide more flexible approaches to probabilistic modelling.

ADDENDUM

4.A RELATION TO KL IMPORTANCE ESTIMATION PROCEDURE (KLIEP)

We now discuss the connections to KLIEP [250]. Consider the same problem setting as in Section 4.3.3 where we wish to use a parameterised function r_α to estimate the exact density ratio,

$$r_\alpha(\mathbf{z}; \mathbf{x}) \approx r^*(\mathbf{z}; \mathbf{x}) \triangleq \frac{q_\phi(\mathbf{z} | \mathbf{x})}{p^*(\mathbf{z})}.$$

We can view $r_\alpha(\mathbf{z}; \mathbf{x})$ as the correction factor required for $p^*(\mathbf{z})$ to match $q_\phi(\mathbf{z} | \mathbf{x})$. This gives rise to an estimator of $q_\phi(\mathbf{z} | \mathbf{x})$,

$$q_\alpha(\mathbf{z} | \mathbf{x}) \triangleq r_\alpha(\mathbf{z}; \mathbf{x})p^*(\mathbf{z}) \approx q_\phi(\mathbf{z} | \mathbf{x}).$$

Although in our specific problem setting, the density $q_\phi(\mathbf{z} | \mathbf{x})$ is tractable, we nonetheless fit an auxiliary model $q_\alpha(\mathbf{z} | \mathbf{x})$ to it as a means of fitting the underlying density ratio estimator $r_\alpha(\mathbf{z}; \mathbf{x})$.

In particular, consider *minimising* the KL divergence between $q_\phi(\mathbf{z} | \mathbf{x})$ and $q_\alpha(\mathbf{z} | \mathbf{x})$ with respect to α ,

$$\begin{aligned} & \mathbb{E}_{q^*(\mathbf{x})} \text{KL} [q_\phi(\mathbf{z} | \mathbf{x}) \parallel q_\alpha(\mathbf{z} | \mathbf{x})] \\ & \triangleq \mathbb{E}_{q^*(\mathbf{x})} \mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x})} \left[\log \frac{q_\phi(\mathbf{z} | \mathbf{x})}{q_\alpha(\mathbf{z} | \mathbf{x})} \right], \\ & = \mathbb{E}_{q^*(\mathbf{x})} \mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x})} \left[\log \frac{q_\phi(\mathbf{z} | \mathbf{x})}{p^*(\mathbf{z})r_\alpha(\mathbf{z}; \mathbf{x})} \right], \\ & = -\mathbb{E}_{q^*(\mathbf{x})} \mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x})} [\log r_\alpha(\mathbf{z}; \mathbf{x})] + \text{const.} \end{aligned}$$

Hence, this is equivalent to *maximising*

$$\mathbb{E}_{q^*(\mathbf{x})} \mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x})} [\log r_\alpha(\mathbf{z}; \mathbf{x})].$$

Now, for the conditional $q_\alpha(\mathbf{z} | \mathbf{x})$ to be a probability density function, its integral must sum to one,

$$\int q_\alpha(\mathbf{z} | \mathbf{x}) q^*(\mathbf{x}) d\mathbf{x} d\mathbf{z} = 1.$$

Rewriting this integral, we have the constraint

$$\begin{aligned} \int q_\alpha(\mathbf{z} | \mathbf{x}) q^*(\mathbf{x}) d\mathbf{x} d\mathbf{z} &= \int r_\alpha(\mathbf{z}; \mathbf{x}) p^*(\mathbf{z}) q^*(\mathbf{x}) d\mathbf{x} d\mathbf{z} \\ &= \mathbb{E}_{q^*(\mathbf{x})} p^*(\mathbf{z}) [r_\alpha(\mathbf{z}; \mathbf{x})] = 1. \end{aligned}$$

Table 4.B.1: Relevant latent and observed space f -divergences instantiated for particular settings of f .

		REVERSE KL	GAN
$f(u)$		$u \log u$	$u \log u - (u + 1) \log(u + 1)$
LATENT	$\mathbb{E}_{q^*(\mathbf{x})} \mathcal{D}_f [p^*(\mathbf{z}) \parallel q_\phi(\mathbf{z} \mathbf{x})]$	$\mathbb{E}_{q^*(\mathbf{x})} \text{KL} [q_\phi(\mathbf{z} \mathbf{x}) \parallel p^*(\mathbf{z})]$	$2 \cdot \mathbb{E}_{q^*(\mathbf{x})} \mathcal{D}_{\text{JS}} [p^*(\mathbf{z}) \parallel q_\phi(\mathbf{z} \mathbf{x})] - \log 4$
OBSERVED	$\mathbb{E}_{p^*(\mathbf{z})} \mathcal{D}_f [q^*(\mathbf{x}) \parallel p_\theta(\mathbf{x} \mathbf{z})]$	$\mathbb{E}_{p^*(\mathbf{z})} \text{KL} [p_\theta(\mathbf{x} \mathbf{z}) \parallel q^*(\mathbf{x})]$	$2 \cdot \mathbb{E}_{p^*(\mathbf{z})} \mathcal{D}_{\text{JS}} [q^*(\mathbf{x}) \parallel p_\theta(\mathbf{x} \mathbf{z})] - \log 4$

Combined, we have the following constrained optimisation problem,

$$\begin{aligned} \max_{\alpha} \quad & \mathbb{E}_{q^*(\mathbf{x})} \mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x})} [\log r_\alpha(\mathbf{z}; \mathbf{x})] \\ \text{subject to} \quad & \mathbb{E}_{q^*(\mathbf{x})} p^*(\mathbf{z}) [r_\alpha(\mathbf{z}; \mathbf{x}) - 1] = 0. \end{aligned}$$

Through the method of Lagrange multipliers, this can be cast as an unconstrained optimisation problem with objective,

$$\begin{aligned} \mathcal{L}_{\text{KLIEP}}^{\text{latent}}(\alpha | \phi) \triangleq & \mathbb{E}_{q^*(\mathbf{x})} \mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x})} [\log r_\alpha(\mathbf{z}; \mathbf{x})] \\ & - \lambda \mathbb{E}_{q^*(\mathbf{x})} p^*(\mathbf{z}) [r_\alpha(\mathbf{z}; \mathbf{x}) - 1], \end{aligned}$$

where λ is the Lagrange multiplier. For $\lambda = 1$, $\mathcal{L}_{\text{KLIEP}}^{\text{latent}}(\alpha | \phi)$ trivially reduces to $\mathcal{L}_{\text{KL}}^{\text{latent}}(\alpha | \phi)$.

4.B SUMMARY OF DEFINITIONS

In this section, we summarise the definitions of the losses defined in the proposed VI framework of Sections 4.3 and 4.4, and underscore the relationships to their respective counterparts in the CYCLEGAN framework of Section 4.5.

Table 4.B.1 summarises the settings of convex function $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ that recover the reverse KL divergence terms within the ELBO and APLBO, and the JS divergence (up to constants) that GANS are known to minimise.

Table 4.B.2 gives the calculations of the terms necessary to explicitly write down instances of the generalised variational lower bound for particular convex functions f – namely the convex dual f^* , the first derivative f' and the composition $f^* \circ f'$.

Table 4.B.3 gives instances of the variational lower bound that approximate the latent and observed space KL divergences within the ELBO and APLBO, respectively. Additionally, it gives generalised *stochastic* formulations of the GAN objectives in the CYCLEGAN framework, while Table 4.B.4 lists their *deterministic* counterpart.

Lastly, Table 4.B.5 gives forward and reverse cycle-consistency constraints in the CYCLEGAN framework, and the specific class of Gaussian

Table 4.B.2: Calculations for convex functions.

	REVERSE KL	GAN
$f(u)$	$u \log u$	$u \log u - (u + 1) \log(u + 1)$
$f^*(t)$	$\exp(t - 1)$	$-\log(1 - \exp t)$
$f'(u)$	$1 + \log u$	$\log \sigma(\log u)$
$f^*(f'(u))$	u	$-\log(1 - \sigma(\log u))$

Table 4.B.3: Instances of variational lower bounds on the relevant latent and observed space f -divergences.

		REVERSE KL	GAN
	$f(u)$	$u \log u$	$u \log u - (u + 1) \log(u + 1)$
LATENT	$\mathcal{L}_f^{\text{latent}}(\alpha \phi) \triangleq \mathbb{E}_{q^*(x)q_\phi(z x)} [f'(r_\alpha(z; \mathbf{x}))] - \mathbb{E}_{q^*(x)p^*(z)} [f^*(f'(r_\alpha(z; \mathbf{x})))]$	$\mathcal{L}_{\text{KL}}^{\text{latent}}(\alpha \phi) \triangleq \mathbb{E}_{q^*(x)q_\phi(z x)} [\log r_\alpha(z; \mathbf{x})] - \mathbb{E}_{q^*(x)p^*(z)} [r_\alpha(z; \mathbf{x}) - 1]$	$\mathcal{L}_{\text{GAN}}^{\text{reverse}}(\alpha \phi) \triangleq \mathbb{E}_{q^*(x)q_\phi(z x)} [\log \sigma(\log r_\alpha(z; \mathbf{x}))] + \mathbb{E}_{q^*(x)p^*(z)} [\log(1 - \sigma(\log r_\alpha(z; \mathbf{x})))]$
OBSERVED	$\mathcal{L}_f^{\text{observed}}(\beta \theta) \triangleq \mathbb{E}_{p^*(z)p_\theta(x z)} [f'(r_\beta(x; \mathbf{z}))] - \mathbb{E}_{p^*(z)q^*(x)} [f^*(f'(r_\beta(x; \mathbf{z})))]$	$\mathcal{L}_{\text{KL}}^{\text{observed}}(\beta \theta) \triangleq \mathbb{E}_{p^*(z)p_\theta(x z)} [\log r_\beta(x; \mathbf{z})] - \mathbb{E}_{p^*(z)q^*(x)} [r_\beta(x; \mathbf{z}) - 1]$	$\mathcal{L}_{\text{GAN}}^{\text{forward}}(\beta \theta) \triangleq \mathbb{E}_{p^*(z)p_\theta(x z)} [\log \sigma(\log r_\beta(x; \mathbf{z}))] + \mathbb{E}_{p^*(z)q^*(x)} [\log(1 - \sigma(\log r_\beta(x; \mathbf{z})))]$

likelihoods and posteriors that instantiates these constraints (in the limit).

Table 4.B.4: General stochastic GAN objectives and their deterministic counterparts.

	STOCHASTIC	DETERMINISTIC
REVERSE	$\mathcal{L}_{\text{GAN}}^{\text{reverse}}(\alpha \phi) \triangleq \mathbb{E}_{q^*(\mathbf{x})p^*(\mathbf{z})}[\log \mathcal{D}_\alpha(\mathbf{z}; \mathbf{x})]$ $+ \mathbb{E}_{q^*(\mathbf{x})p(\epsilon)}[\log(1 - \mathcal{D}_\alpha(\mathcal{G}_\phi(\epsilon; \mathbf{x}); \mathbf{x}))]$	$\ell_{\text{GAN}}^{\text{reverse}}(\alpha \phi) \triangleq \mathbb{E}_{p^*(\mathbf{z})}[\log \mathbf{D}_\alpha(\mathbf{z})]$ $+ \mathbb{E}_{q^*(\mathbf{x})}[\log(1 - \mathbf{D}_\alpha(\mathbf{m}_\phi(\mathbf{x})))]$
FORWARD	$\mathcal{L}_{\text{GAN}}^{\text{forward}}(\beta \theta) \triangleq \mathbb{E}_{p^*(\mathbf{z})q^*(\mathbf{x})}[\log \mathcal{D}_\beta(\mathbf{x}; \mathbf{z})]$ $+ \mathbb{E}_{p^*(\mathbf{z})p(\xi)}[\log(1 - \mathcal{D}_\beta(\mathcal{F}_\theta(\xi; \mathbf{z}); \mathbf{z}))]$	$\ell_{\text{GAN}}^{\text{forward}}(\beta \theta) \triangleq \mathbb{E}_{q^*(\mathbf{x})}[\log \mathbf{D}_\beta(\mathbf{x})]$ $+ \mathbb{E}_{p^*(\mathbf{z})}[\log(1 - \mathbf{D}_\beta(\mu_\theta(\mathbf{z})))]$

Table 4.B.5: Negative expected log conditionals and the cycle-consistency constraints.

	GAUSSIAN	DEGENERATE	
$p_\theta(\mathbf{x} \mathbf{z})$	$q_\phi(\mathbf{z} \mathbf{x})$	$p_\theta(\mathbf{x} \mathbf{z})$	$q_\phi(\mathbf{z} \mathbf{x})$
$\mathcal{N}(\mathbf{x} \mu_\theta(\mathbf{z}), \tau^{-1}\mathbf{I})$	$\mathcal{N}(\mathbf{z} \mathbf{m}_\phi(\mathbf{x}), t^{-1}\mathbf{I})$	$\delta(\mathbf{x} - \mu_\theta(\mathbf{z}))$	$\delta(\mathbf{z} - \mathbf{m}_\phi(\mathbf{x}))$
$\mathcal{L}_{\text{NELP}}(\theta, \phi) \triangleq \frac{\tau}{2} \mathbb{E}_{q^*(\mathbf{x})p(\epsilon)}[\ \mathbf{x} - \mu_\theta(\mathbf{m}_\phi(\mathbf{x}) + t^{-\frac{1}{2}}\epsilon)\ _2^2] + \frac{D}{2} \log \frac{2\pi}{\tau}$		$\ell_{\text{CONST}}^{\text{reverse}}(\theta, \phi) \triangleq \mathbb{E}_{q^*(\mathbf{x})}[\ \mathbf{x} - \mu_\theta(\mathbf{m}_\phi(\mathbf{x}))\ _2^2]$	
$\mathcal{L}_{\text{NELP}}(\theta, \phi) \triangleq \frac{t}{2} \mathbb{E}_{p^*(\mathbf{z})p(\xi)}[\ \mathbf{z} - \mathbf{m}_\phi(\mu_\theta(\mathbf{z}) + \tau^{-\frac{1}{2}}\xi)\ _2^2] + \frac{K}{2} \log \frac{2\pi}{t}$		$\ell_{\text{CONST}}^{\text{forward}}(\theta, \phi) \triangleq \mathbb{E}_{p^*(\mathbf{z})}[\ \mathbf{z} - \mathbf{m}_\phi(\mu_\theta(\mathbf{z}))\ _2^2]$	