

# Soft-IntroVAE: Analyzing and Improving the Introspective Variational Autoencoder

Tal Daniel

Department of Electrical Engineering  
Technion, Haifa, Israel

taldanielm@campus.technion.ac.il

Aviv Tamar

Department of Electrical Engineering  
Technion, Haifa, Israel

avivt@technion.ac.il

## Abstract

*The recently introduced introspective variational autoencoder (IntroVAE) exhibits outstanding image generations, and allows for amortized inference using an image encoder. The main idea in IntroVAE is to train a VAE adversarially, using the VAE encoder to discriminate between generated and real data samples. However, the original IntroVAE loss function relied on a particular hinge-loss formulation that is very hard to stabilize in practice, and its theoretical convergence analysis ignored important terms in the loss. In this work, we take a step towards better understanding of the IntroVAE model, its practical implementation, and its applications. We propose the Soft-IntroVAE, a modified IntroVAE that replaces the hinge-loss terms with a smooth exponential loss on generated samples. This change significantly improves training stability, and also enables theoretical analysis of the complete algorithm. Interestingly, we show that the IntroVAE converges to a distribution that minimizes a sum of KL distance from the data distribution and an entropy term. We discuss the implications of this result, and demonstrate that it induces competitive image generation and reconstruction. Finally, we describe an application of Soft-IntroVAE to unsupervised image translation, and demonstrate compelling results. Code and additional information is available on the project website - [taldatech.github.io/soft-intro-vae-web](https://taldatech.github.io/soft-intro-vae-web).*

## 1. Introduction

Two popular approaches for learning deep generative models are generative adversarial training (e.g., GANs [14]), and variational inference (e.g., VAEs [33]). VAEs are known to have a stable training procedure, display resilience to mode collapse, and enable amortized inference. Moreover, the VAE’s inference module makes them prominent in many domains, such as learning disentangled representations [2] and reinforcement learning [17]. GANs, on the other hand, lack an inference module, but are capa-

ble of generating images of higher quality and are popular in computer vision applications, but can suffer from training instability and low sampling diversity [41, 34].

Narrowing the gap between VAEs and GANs has been the aim of many works, in an attempt to combine the best of both worlds: building a stable and easy to train generative model that allows efficient amortized inference and high-quality sampling [36, 11, 40, 58, 44]. While progress has been made, the search for better generative models is an active research field.


Recently, Huang et al. [25] proposed IntroVAE – a VAE that is trained adversarially, and demonstrated outstanding image generation results. A key idea in IntroVAE is introspective discrimination – instead of training a separate discriminator network to discriminate between real and generated samples, as in a GAN, the output of the encoder acts as the discriminatory signal, based on the Kullback-Leibler divergence between the approximate posterior of a sample and the prior latent distribution. Intuitively, this signal can be understood as making the generated samples less likely, as their posterior is more distant from the prior.

Importantly, IntroVAE uses a *hard margin*,  $m$ , as a threshold on the KL divergence for which above it a ‘fake’ sample no longer affects the loss function. This approach leads to difficult training in practice – the optimization is very sensitive to the values of  $m$ , and extensive tuning is required to find a good, stable, margin. Theoretically, [25] proved that IntroVAE converges to the data distribution, but their analysis ignored several terms in the loss function that are important for correct operation of the algorithm.

We aim to provide a better understanding of the introspective training paradigm and improve its stability. To that end, we introduce Soft-IntroVAE (S-IntroVAE) – an introspective VAE that utilizes the evidence lower bound (ELBO) as the discriminatory signal, and replaces the hard margin with a soft threshold function, making it more stable to optimize. This new formulation allows us to analyze the convergence properties of the complete algorithm, and provide new insights into introspective VAEs.



(a) FFHQ dataset – samples from S-IntroVAE (FID: 17.55).



(b) FFHQ – reconstructions.

Figure 1: Generated samples (left) and reconstructions (right) of test data (left: real, right: reconstruction) from a style-based S-IntroVAE trained on FFHQ at 256x256 resolution.

Interestingly, our theoretical analysis shows that, in contrast to the original IntroVAE, the S-IntroVAE encoder converges to the true posterior, thus maintaining the inference capabilities of the VAE framework. Our analysis also reveals that the S-IntroVAE model converges to a generative distribution that minimizes a sum of KL divergence from the data distribution and an entropy term, in contrast to GANs, where the distribution of the generator converges to the data distribution. We further analyze the consequences of this result and its practical implications, and show that the model produces high-quality generations from a distribution with a sharp support. In practice, S-IntroVAE is much more stable than IntroVAE, as it does not involve the sensitive threshold parameter, and we rigorously validate this claim in experiments, ranging from inference on 2D distributions to high-quality image generation.

Finally, we demonstrate a practical application of our model to the task of unsupervised image translation. We exploit the fact that S-IntroVAE has both good generation quality and strong inference capabilities, and combine it with an encoder-decoder architecture that induces disentanglement. Inductive bias of this sort is required for unsupervised learning of disentangled representations [38], and our results show that using this architecture, S-IntroVAE is indeed capable of successfully transferring content between two images without any supervision.

Our contribution is summed as follows: (1) We propose Soft-IntroVAE, a modification of the original IntroVAE that utilizes the evidence lower bound (ELBO) as a discriminatory signal, and does not require a hard-margin threshold; (2) We provide a deeper theoretical understanding of introspective VAEs; (3) We validate that training our model is significantly more robust than training the original IntroVAE; (4) We show that our method is capable of high-quality image synthesis; (5) We demonstrate a practical application of our model to unsupervised image translation.

## 2. Related Work

Studies on enhancing VAEs can be divided to methods that either improve the network’s architecture [54, 49], incorporate stronger priors [52, 45, 27], add regularization [13, 55], or incorporate adversarial objectives [25, 40, 18, 44, 22, 11, 36]. Our work belongs to the latter group. Adversarial Autoencoders [40] add an adversarial loss on the latent code using a discriminator in addition to the typical encoder/decoder training. VAE/GAN chains VAEs and GANs by adding a discriminator on top of the decoder [36]. BiGAN [10] and ALI [11] simulate autoencoding by training with adversarial objectives in both the latent and data spaces. VEEGAN [50] introduces an additional ‘reconstructor’ network that is trained to map the real data to a Gaussian and to approximately invert the generator network. Recently, [44] proposed ALAE, a latent adversarial autoencoding method that adds a style-based encoder to the generator of StyleGAN [30]. AGE [53], IAN [5], PIONEER [21], B-PIONEER [22], and IntroVAE [25] propose an introspective training of VAEs, and the latter is the current state of the art in image generation. While the idea of introspective training is shared between the aforementioned methods, the difference is summed as follows: in AGE, the adversarial game between the encoder and generator is performed in the latent space (minimizing and maximizing divergences). In addition, the encoder minimizes the reconstruction error in *data-space* (e.g. pixel-space), while the generator minimizes the reconstruction error in the *latent-space*. PIONEER’s objective is similar to AGE’s but uses the cosine distance applied to reconstructions in the latent space. In IntroVAE, as in the previous methods, the encoder also maximizes divergence of the generated data latent representation, but uses a hard-margin over it, and both encoder and decoder minimize the reconstruction error in data-space. Our focus is IntroVAE, as it is state-of-the-art in image synthesis. We provide a deeper theoretical analysis

Autoencoding method	(a) Data Distribution	(b) Latent Distribution	(c) Reciprocity Space
VAE [33]	similarity	imposed/divergence	data
AAE [40]	similarity	imposed/adversarial	data
VAE/GAN [36]	similarity	imposed/divergence	data
VampPrior [52]	similarity	learned/divergence	data
BiGAN [10]	adversarial	imposed/adversarial	adversarial
ALI [11]	adversarial	imposed/adversarial	adversarial
VEEGAN [50]	adversarial	imposed/divergence	latent
AGE [53]	adversarial	imposed/adversarial	latent
PIONEER [21]	adversarial	imposed/adversarial	latent
IntroVAE [25]	adversarial	imposed/adversarial	data
ALAE [44]	adversarial	learned/divergence	latent
Soft-IntroVAE (ours)	adversarial	imposed/adversarial	data

Table 1: Comparison between AE methods based on criteria used: (a) for matching the real to the synthetic data distribution; (b) for setting/learning the latent distribution; (c) for which space reciprocity is achieved.

of the introspective training and propose a stable algorithm that is capable of high-quality image generation.

Table 1, originally devised by [44], compares between different autoencoding methods based on the following properties: (a) how does the model learn the distribution of the data (i.e., by similarity to the original data or in an adversarial manner?); (b) does the model impose a distribution type on the latent space (e.g., a Gaussian) or is the latent space prior learned? and (c) is the autoencoding reconstruction error measured (termed reciprocity) with respect to the data or the latent space? Soft-IntroVAE is a modification of IntroVAE, and shares similar high-level characteristics, as shown in the table.

### 3. Background

We formulate our problem under the variational inference (VI) setting [33]. Given some data  $x \sim p_{data}(x)$ , where  $p_{data}$  denotes the data distribution, one aims to fit the parameters  $\theta$  of a latent variable model  $p_\theta(x) = \mathbb{E}_{p(z)} [p_\theta(x|z)]$ , where the prior  $p(z)$  is given, and  $p_\theta(x|z)$  is learned. For general models, a typical objective is the maximum-likelihood  $\max_\theta \log p_\theta(x)$ , which is intractable, and can be approximated using VI methods.

**Evidence lower bound (ELBO)** For some model defined by  $p_\theta(x|z)$  and  $p(z)$ , let  $p_\theta(z|x)$  denote the posterior distribution that the model induces on the latent variable. The evidence lower bound (ELBO) states that for any approximate posterior distribution  $q(z|x)$ :

$$\begin{aligned} \log p_\theta(x) &\geq \mathbb{E}_{q(z|x)} [\log p_\theta(x|z)] - KL(q(z|x)||p(z)) \\ &\doteq ELBO(x), \end{aligned} \quad (1)$$

where the Kullback-Leibler (KL) divergence is  $KL(q(z|x)||p(z)) = \mathbb{E}_{q(z|x)} \left[ \log \frac{q(z|x)}{p(z)} \right]$ . In the variational autoencoder (VAE, [33]), the approximate posterior is represented as  $q_\phi(z|x) = \mathcal{N}(\mu_\phi(x), \Sigma_\phi(x))$  for some neural network with parameters  $\phi$ , the prior is  $p(z) = \mathcal{N}(\mu_0, \Sigma_0)$ , and the ELBO can be maximized using

the *reparameterization trick*. Since the resulting model resembles an autoencoder, the approximate posterior  $q_\phi(z|x)$  is also known as the *encoder*, while  $p_\theta(x|z)$  is termed the *decoder*. Typically,  $p_\theta(x|z)$  is modelled as a Gaussian distribution, and in this case  $\log p_\theta(x|z)$  is equivalent to the mean-squared error (MSE) between  $x$  and the mean of  $p_\theta(x|z)$ . In the following, the term reconstruction error refers to  $\log p_\theta(x|z)$ .

**Introspective VAE (IntroVAE)** The Introspective VAE adds an adversarial objective to the VAE training, where the intuition, following [25], is that since the VAE maximizes a lower bound of the data likelihood, there is no guarantee that points outside the support of  $p_{data}$  will not be assigned high likelihood. In [25], the KL term in the ELBO is seen as an ‘energy’ of a sample. Inspired by energy-based GANs (EBGAN, [58]), the encoder is encouraged to classify between the generated and real samples by *minimizing* the KL of latent distribution of real samples and the prior, and *maximizing* the KL of generated ones. The decoder, on the other hand, is trained to reconstruct real data samples using the standard ELBO, and to minimize the KL of generated samples that go through the encoder. Consider a real sample  $x$  and a generated one  $D_\theta(z) \sim p_\theta(x|z)$ , and let  $KL(E_\phi(D_\theta(z))) = KL(q_\phi(\cdot|D_\theta(z))||p(\cdot))$  denote the KL of the generated sample that goes through the encoder. The adversarial objective, which is to be *maximized*, for  $x$  and  $z$  is given by:

$$\begin{aligned} \mathcal{L}_{E_\phi}(x, z) &= ELBO(x) - [m - KL(E_\phi(D_\theta(z)))]^+, \\ \mathcal{L}_{D_\theta}(x, z) &= ELBO(x) - KL(E_\phi(D_\theta(z))), \end{aligned} \quad (2)$$

where  $\mathcal{L}_{E_\phi}$  is the objective function for the encoder,  $\mathcal{L}_{D_\theta}$  is the objective function for the decoder, and  $[\cdot]^+ = \max(0, \cdot)$ . The hard threshold  $m$  in  $\mathcal{L}_{E_\phi}$  limits the KL of generated samples in the objective, and is crucial for the analysis and practical implementation of IntroVAE. The encoder and decoder are trained simultaneously using stochastic gradient descent (SGD) on mini-batches of  $x$  and  $z$ .

The theoretical properties of IntroVAE were studied in [25] under a simplified objective that ignores the reconstruction terms in the ELBO, and omits the decoder loss for real samples, as given by:

$$\begin{aligned} \mathcal{L}_{E_\phi}(x, z) &= -KL(E_\phi(x)) - [m - KL(E_\phi(D_\theta(z)))]^+, \\ \mathcal{L}_{D_\theta}(z) &= -KL(E_\phi(D_\theta(z))). \end{aligned} \quad (3)$$

For this objective, [25] showed that a Nash equilibrium is obtained when the distribution  $p(x) = \mathbb{E}_z [p_\theta(x|z)]$  is exactly  $p_{data}(x)$ , showing the soundness of the model in terms of sample generation. However, the EBGAN-based analysis did not make claims about the inference capabilities of the model, and in particular, it is possible that at the Nash equilibrium the encoder distribution  $q_\phi(z|x)$  is different from the true posterior  $p_\theta(z|x)$ .<sup>1</sup>

<sup>1</sup>The theoretical results in [25], as the results in our work, are proved



## 4. Soft-IntroVAE

In this section, we propose a new introspective VAE model, termed *Soft-IntroVAE*, that mitigates two shortcomings of IntroVAE – the training instability due to the hard threshold function (cf. Eq. 2), and the difficulty in analysing the full optimization objective (cf. Eq. 3).

Recall that  $p_{data}(x)$  denotes the data distribution and  $p(z)$  represents some prior distribution over latent variables  $z$ . The objective functions for the encoder,  $E_\phi$ , and the decoder,  $D_\theta$ , for samples  $x$  and  $z$  in our model are given by:

$$\begin{aligned}\mathcal{L}_{E_\phi}(x, z) &= ELBO(x) - \frac{1}{\alpha} \exp(\alpha ELBO(D_\theta(z))), \\ \mathcal{L}_{D_\theta}(x, z) &= ELBO(x) + \gamma ELBO(D_\theta(z)),\end{aligned}\quad (4)$$

where  $ELBO(x)$  is defined in equation 1,  $ELBO(D_\theta(z))$  is defined similarly but with  $D_\theta(z) \sim p_\theta(x|z)$  replacing  $x$  in Eq. 1, and  $\alpha \geq 0$  and  $\gamma \geq 0$  are hyper-parameters. Note that Eq. 4 portrays a game between the encoder and the decoder: the encoder is induced to distinguish, through the ELBO value, between real samples (high ELBO) and generated samples (low ELBO), while the decoder is induced to generate samples the ‘fool’ the encoder. The complete Soft-IntroVAE objective takes an expectation of the losses above over real and generated samples:

$$\begin{aligned}\mathcal{L}_{E_\phi} &= \mathbb{E}_{x \sim p_{data}, z \sim p(z)} [\mathcal{L}_{E_\phi}(x, z)], \\ \mathcal{L}_{D_\theta} &= \mathbb{E}_{x \sim p_{data}, z \sim p(z)} [\mathcal{L}_{D_\theta}(x, z)].\end{aligned}\quad (5)$$

For a Gaussian posterior, these losses can be optimized effectively by SGD using the reparametrization trick [33].

There are two key differences between S-IntroVAE and IntroVAE. The first is that we utilize the complete ELBO term instead of just the KL, which will allow us to provide a complete variational inference-based analysis. The second difference is replacing the hard threshold in Eq. 2 with a soft exponential function over the ELBO, henceforth denoted *expELBO*. The effect of both of these functions is similar – they induce a separation between the posterior distributions over latent variables of real samples and generated ones. However, as we report in the sequel, the soft threshold in Eq. 4 is much easier to optimize, and results in improved training stability.

At this point, the reader may question what minimizing the ELBO for generated samples (through the *expELBO* in Eq. 4) means, as minimizing a lower bound on the log-likelihood does not imply that the likelihood of generated samples decreases [7]. In addition, for a decoder that produces near-perfect generated samples, it may seem that the *expELBO* term seeks to reduce sample quality. In the following, we answer these questions by analyzing the equilibrium of Eq. 5.

---

under the non-parametric setting. We nevertheless use the parametric setting notation to avoid introducing additional notation, and understand from the context that  $\theta$  and  $\phi$  do not carry meaning in the non-parametric setting.

## 4.1. Analysis

In this section we analyze the Nash equilibrium of the game in Eq. 5. We consider a non-parametric setting, where the encoder and decoder can represent any distribution. This is a typical setting for analyzing adversarial generative models [14, 25]. For simplicity, we focus on discrete distributions, but our analysis easily extends to the continuous case. Also, to ease the presentation, we focus on the case  $\alpha = 1$ . The analysis for general  $\alpha$  provides similar insights and is provided in Appendix 9.6 with detailed proofs.

We introduce the following notation. The encoder is represented by the approximate posterior distribution  $q \doteq q(z|x)$ . The decoder is represented using  $d \doteq p_d(x|z)$ . These are the controllable distributions in our generative model. The latent prior is denoted  $p(z)$  and is not controlled. Slightly abusing notation, we also denote  $p_d(x) = \mathbb{E}_{p(z)}[p_d(x|z)]$  as the distribution of generated samples. For some distribution  $p(x)$ , let  $H(p) = -\mathbb{E}[\log p(x)]$  denote its Shannon entropy.

We define  $d^*$  as follows:

$$d^* \in \arg \min_d \{KL(p_{data}||p_d) + \gamma H(p_d(x))\}.\quad (6)$$

Note that for  $\gamma = 0$ , we have that  $p_{d^*} = p_{data}$ . For  $\gamma > 0$ , however,  $p_{d^*}$  represents a balance between closeness to  $p_{data}$  and low entropy. We make the following assumption.

**Assumption 1.** *For all  $x$  such that  $p_{data}(x) > 0$  we have that  $p_{d^*}(x) \leq \sqrt{p_{data}(x)}$ .*

Assumption 1 can be seen as a condition of the closeness between  $p_{d^*}$  and  $p_{data}$ , and essentially requires that the effect of the entropy minimization term in Eq. 6 is limited. Intuitively, if  $\gamma$  is small enough, we should always be able to satisfy Assumption 1, as established in the next result.

**Proposition 2.** *For any  $p_{data}$ , there exists  $\gamma > 0$  such that  $p_{d^*}$ , as defined in equation 6, satisfies Assumption 1.*

We are now ready to state our main result – that  $p_{d^*}$  is an equilibrium point of the S-IntroVAE model.

**Theorem 3.** *Let  $d^*$  be defined as in equation 6. Denote  $q^* = p_{d^*}(z|x)$ . If Assumption 1 holds, then  $(q^*, d^*)$  is a Nash equilibrium of equation 5.*

Interestingly, Theorem 3 shows that the S-IntroVAE model does not converge to the data distribution, but to an entropy-regularized version of it. One should question the effect of such regularization, in light of the typical goal of generating samples that are similar to the data distribution. The experiments in Section 6, on various 2-dimensional datasets, illustrate that S-IntroVAE learns distributions with sharper supports than a standard VAE, but without negative effects such as mode dropping. Our image experiments further support this statement.

We now contrast our analysis with the analysis of [25]. First, we note that ignoring the reconstruction terms in the analysis, as done by [25], leads to significantly different insights. For example, if one removes the  $ELBO(x)$  term from  $\mathcal{L}_{D_\theta}$ , our analysis shows that the model will effectively only minimize  $H(p_d(x))$ , without any dependence on  $p_{data}$  (see Appendix 9.6 for more details). Indeed, the empirical results in [25] were only obtained with the  $ELBO(x)$  term in  $\mathcal{L}_{D_\theta}$ . Furthermore, our analysis, as detailed in the appendix, does not build on representing our model as a particular instance of the EBGAN, but explicitly builds on the variational properties of the ELBO. In this sense, our results more closely tie between variational inference principles and adversarial generative models.

It is important to note that at equilibrium, the encoder converges to the true posterior  $q^* = p_{d^*}(z|x)$ . Thus, the exponential penalty in  $\mathcal{L}_E$  does not harm the inference properties of the encoder. This important conclusion does not appear in the analysis of [25].

We finally remark on the parameter  $\gamma$ . While the analysis makes a strict assumption on  $\gamma$ , as evident in Assumption 1 and Proposition 2, in all our experiments we set  $\gamma = 1$ . Thus, it seems that in practice the requirements for obtaining decent results are much less stringent.

## 5. Implementation

In this section, we outline several implementation modifications of the S-IntroVAE algorithm that proved helpful in practice. Pseudo-code of our algorithm is depicted in Algorithm 1, and a more detailed version can be found in Appendix 9.1. In all our experiments, we set  $\alpha = 2$  (other values, such as  $\alpha = 1$ , work similarly) and  $\gamma = 1$  (cf. equation 4). Moreover, for the ELBO terms, we use the  $\beta$ -VAE [23] formulation and rewrite it as:  $ELBO(x) = \beta_{rec} \mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x|z)] - \beta_{kl} KL[q_\phi(z|x) || p(z)]$ . The hyperparameters  $\beta_{rec}$  and  $\beta_{kl}$  control the balance between inference and sampling quality respectively.<sup>2</sup> When  $\beta_{rec} > \beta_{kl}$ , the optimization is focused on good reconstructions, which may lead to less variability in the generated samples, as latent posteriors are allowed to be very different from the prior, and when  $\beta_{rec} < \beta_{kl}$ , there will be more varied samples, but reconstruction quality will degrade. Note that when  $\beta_{rec} \ll \beta_{kl}$ , the VAE is prone to *posterior collapse*, and we further discuss this issue in Appendix 9.5. For most cases, we found that values between 0.05 and 1 for  $\beta_{rec}$  and  $\beta_{kl}$  work well, and the best configuration depends on the architecture and data set, as we detail in Section 6.

Each ELBO term in equation 4 can be considered as an instance of  $\beta$ -VAE and can have different  $\beta_{rec}$  and  $\beta_{kl}$  pa-

<sup>2</sup>In principle, only the ratio between  $\beta_{rec}$  and  $\beta_{kl}$  affects the loss function. However, we found it easier to work with two parameters, as the balance between ELBO and expELBO is affected by both parameters.

rameters. However, we set them all to be the same, except for the ELBO inside the exponent. For this term,  $\beta_{kl}$  controls the repulsion force of the posterior for generated samples from the prior. We found that good results are obtained when  $\beta_{kl}$  is set to be on the order of magnitude of  $z_{dim}$  (e.g., for  $z_{dim} = 128$ ,  $\beta_{kl}$  should be set to values around 128). In Algorithm 1, this specific parameter is denoted as  $\beta_{neg}$ . Moreover, notice that the decoder tries to minimize the reconstruction error for generated data, which may slow down convergence, as at the beginning of the optimization the generated samples are of low quality. Thus, we introduce hyperparameter  $\gamma_r$  (not to be confused with  $\gamma$  from equation 4) that multiplies only the reconstruction term of the generated data in the ELBO term of the decoder in equation 4. In all our experiments we used a constant  $\gamma_r = 1e - 8$  independently of the data type. We remark that setting  $\gamma_r$  to zero had a significant detrimental effect on performance. Finally, we use a scaling constant  $s$  to balance between the ELBO and the expELBO terms in the loss, and we set  $s$  to be the inverse of the input dimensions. This scaling constant prevents the expELBO from vanishing for high-dimensional input. Our code is publicly available<sup>3</sup>.

## 6. Experiments

In this section, we detail our experiments with the following goals in mind: (1) Understanding the distributions that S-IntroVAE learns to generate; (2) Evaluating the robustness and stability of S-IntroVAE compared to IntroVAE; (3) Benchmarking S-IntroVAE on high-quality image synthesis; and (4) Demonstrating a practical application of S-IntroVAE to unsupervised image translation.

To answer (1) and (2), we investigate learning of 2D distributions, which are both easy to interpret, and enable a quantitative evaluation of quality and robustness. For (3) and (4), we experiment with standard image data sets.

### 6.1. 2D Toy Datasets

We evaluate our method on four 2D datasets: 8 Gaussians, Spiral, Checkerboard and Rings [8, 16], and compare with a standard VAE and IntroVAE.

We calculate 3 metrics to measure how well the methods learn the true data distribution: KL-divergence and Jensen-Shannon-divergence (JSD), and a custom metric, *grid-normalized ELBO* (gnELBO). The KL and JSD are calculated using empirical histograms of samples generated from the trained models and samples generated from the real data distribution. Thus, these metrics evaluate the generation capabilities of the learned model. The grid-normalized ELBO, on the other hand, treats the ELBO as an energy term. We normalize the ELBO of the learned model over a grid of points to produce a normalized energy

<sup>3</sup>[github.com/taldatech/soft-intro-vae-pytorch](https://github.com/taldatech/soft-intro-vae-pytorch)

**Algorithm 1** Training Soft-IntroVAE (pseudo-code)

---

**Require:**  $\beta_{rec}, \beta_{kl}, \beta_{neg}, \gamma_r$

- 1:  $\phi_E, \theta_D \leftarrow$  Initialize network parameters
- 2:  $s \leftarrow 1/\text{input dim}$  ▷ Scaling constant
- 3: **while** not converged **do**
- 4:    $X \leftarrow$  Random mini-batch from dataset
- 5:    $Z \leftarrow E(X)$  ▷ Encode
- 6:    $Z_f \leftarrow$  Samples from prior  $N(0, I)$
- 7:   **procedure** UPDATEENCODER( $\phi_E$ )
- 8:      $X_r \leftarrow D(Z), X_f \leftarrow D(Z_f)$  ▷ Decode
- 9:      $Z_{ff} \leftarrow E(X_f)$
- 10:      $X_{ff} \leftarrow D(Z_{ff})$
- 11:      $\text{ELBO} \leftarrow s \cdot \text{ELBO}(\beta_{rec}, \beta_{kl}, X, X_r, Z)$
- 12:      $\text{ELBO}_f \leftarrow \text{ELBO}(\beta_{rec}, \beta_{neg}, X_f, X_{ff}, Z_{ff})$
- 13:      $\text{expELBO}_f \leftarrow 0.5 \exp(2s \cdot \text{ELBO}_f)$
- 14:      $L_E \leftarrow \text{ELBO} - \text{expELBO}_f$  ▷ Eq. 4
- 15:      $\phi_E \leftarrow \phi_E + \eta \nabla_{\phi_E}(L_E)$  ▷ Adam update
- 16:   **end procedure**
- 17:   **procedure** UPDATEDECODER( $\theta_D$ )
- 18:      $X_r \leftarrow D(Z), X_f \leftarrow D(Z_f)$  ▷ Decode
- 19:      $Z_{ff} \leftarrow E(X_f)$
- 20:      $X_{ff} \leftarrow sg(D(Z_{ff}))$  ▷  $sg$ : stop-gradient
- 21:      $\text{ELBO} \leftarrow \beta_{rec} L_{rec}(X, X_r)$
- 22:      $\text{ELBO}_f \leftarrow \text{ELBO}(\gamma_r \cdot \beta_{rec}, \beta_{kl}, X_f, X_{ff}, Z_{ff})$
- 23:      $L_D \leftarrow s \cdot (\text{ELBO} + \text{ELBO}_f)$  ▷ Eq. 4
- 24:      $\theta_D \leftarrow \theta_D + \eta \nabla_{\theta_D}(L_D)$  ▷ Adam update
- 25:   **end procedure**
- 26: **end while**

---

term  $\widehat{ELBO}$ , and we measure  $\mathbb{E}_{x \sim p_{data}} [-\widehat{ELBO}(x)]$  by sampling from the data distribution. This metric is affected by the encoder and thus effectively measures the inference capabilities of the model – a good model should assign a high ELBO for likely samples and a low ELBO for unlikely ones. For all metrics, *lower is better*.

In Figure 2 we plot random samples from the models and a density estimation (for the rest of the datasets see Appendix 9.4.1), obtained by approximating  $p(x)$  with  $\exp(\widehat{ELBO})$ . In order to tune the algorithms for each dataset, we ran an extensive hyperparameter grid search of 81 runs for the standard VAE, 210 runs for S-IntroVAE, and 1260 runs for IntroVAE, due to the additional  $m$  parameter. The architecture for all methods is a simple 3-layer fully-connected network with 256 hidden units and ReLU activations, and the latent space dimension is 2. The complete set of hyperparameters is provided in Appendix 9.2.1.

Results for the different evaluation metrics are shown in Table 2. Evidently, Soft-IntroVAE outperforms IntroVAE both quantitatively and qualitatively, and both are superior to the standard VAE. Note that the standard VAE assigns low energy (high likelihood) to points outside the data support, as pointed out in Section 3. The adversarial loss in the

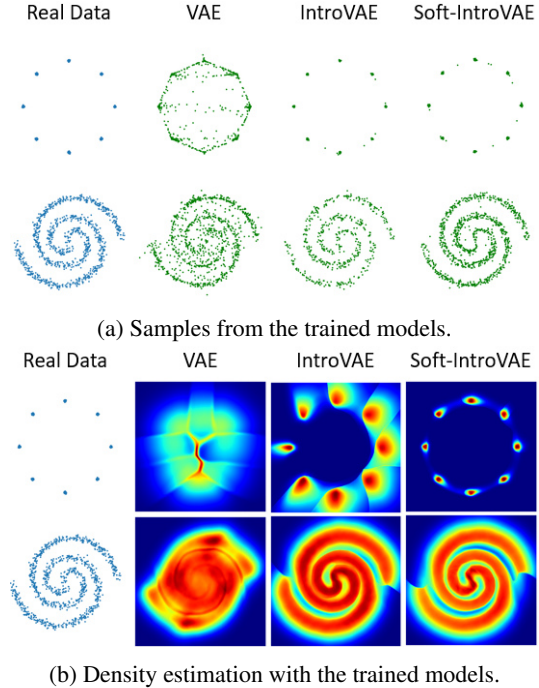


Figure 2: Unsupervised learning of 2D datasets.

introspective models, on the other hand, prohibits the model from generating such samples.

		VAE	IntroVAE	Soft-IntroVAE
8 Gaussians	gnELBO	7.42±0.07	1.29±0.76	<b>1.25±0.35</b>
	KL	6.72±0.46	2.53±1.07	<b>1.25±0.11</b>
	JSD	16.04±0.3	1.67±0.46	<b>0.96±0.15</b>
Spiral	gnELBO	6.19±0.06	5.87±0.03	<b>5.21±0.04</b>
	KL	9.8±0.48	8.38±0.45	<b>8.13±0.3</b>
	JSD	4.89±0.05	3.58±0.04	<b>3.37±0.04</b>
Checkerboard	gnELBO	8.53±0.1	8.54±0.11	<b>4.47±0.29</b>
	KL	20.91±0.45	<b>19.03±0.34</b>	20.27±0.21
	JSD	9.78±0.04	9.07±0.1	<b>9.06±0.15</b>
Rings	gnELBO	6.4±0.04	7.25±0.18	<b>6.3±0.08</b>
	KL	13.16±0.55	10.21±0.49	<b>9.18±0.33</b>
	JSD	7.26±0.07	4.24±0.11	<b>4.13±0.09</b>

Table 2: Results on 2D datasets. Grid-normalized ELBO is in  $1e^{-7}$  units. Results are averaged over 5 seeds.

## 6.2. Training Stability of Soft-IntroVAE

In practice, we found that training the original IntroVAE model was very difficult, and prone to instability. In fact, we were not able to reproduce the results reported in [25] on image datasets, even when using the authors’ published code [1] or other implementations, and countless parameter investigations. We suspect that the algorithm is very sensitive to the choice of the  $m$  parameter: at some point during training when the KL of generated samples is larger than the threshold  $m$ , there is no more adversarial signal for the encoder (as there are no gradients from the KL-divergence of generated samples), while the decoder still tries to ‘fool’ the encoder. Finding the right  $m$  where both the encoder




Figure 3: Stability investigation. KL divergence and JSD metrics for IntroVAE (I-VAE in the plot) with respect to  $m$  (see text for details). S-IntroVAE (S-VAE in the plot) does not require the  $m$  parameter, and we plot its performance with all other hyper-parameters the same as IntroVAE. Results are averaged over 5 seeds. Note that the constant lines of the KL and JSD of S-VAE overlap.

and decoder can maintain the adversarial signal and reach equilibrium is the key for IntroVAE success.

To demonstrate that the  $m$  hyperparameter in IntroVAE plays an important role in the stability and convergence of IntroVAE, we pick the best combination of hyperparameters from the extensive search we performed for the experiments in Section 6.1 on the 8-Gaussians 2D dataset, and plot the KL and JSD (lower is better) for varying values of  $m$ . For comparison, we plot the results of S-IntroVAE with the same hyperparameters (without  $m$ ) in Figure 3. Evidently, IntroVAE is very sensitive to this hyperparameter, while our method does not require it. We remark that for different datasets, significantly different values of  $m$  were required to obtain reasonable results. We found the soft expELBO term less sensitive, and we obtained good results on a wide range of values as described in Section 5.

### 6.3. Image Generation

In this section, we evaluate Soft-IntroVAE on image synthesis in terms of both inference (i.e., reconstruction) and sampling (generation). To measure the quality and variability of the images, we report the Fréchet inception distance (FID) based on 50,000 generated images from the model, and use the training samples as the reference images [29, 44]. Detailed hyperparameter settings and data set details are provided in the supplementary material.

**Architectures and Hyperparameters:** We experiment with two convolution-based architectures: (1) IntroVAE’s [25] encoder-decoder architecture with residual-based convolutional layers (see Figure 8a) and (2) ALAE’s [44] style-based autoencoder architecture, which adopted StyleGAN’s [30] style generator to a style-based encoder (see Figure 8c). For the style-based architecture, we also use progressive growing as in [28, 29, 44], where we start from low-resolution  $4 \times 4$  images and progressively increase the resolution by smoothly blending in new blocks in the en-

	CelebA-HQ	FFHQ
PGGAN [28])	<b>8.03</b>	-
BigGAN [4]	-	11.48
U-Net GAN [47]	-	<b>7.48</b>
GLOW [32]	68.93	-
Pioneer [21]	39.17	-
Balanced Pioneer [22]	25.25	-
StyleALAE [44]	19.21	-
SoftIntroVAE (Ours)	<b>18.63</b>	<b>17.55</b>

Table 3: Comparison of FID scores (lower is better) for CelebA-HQ and FFHQ datasets at a resolution of 256x256. Note the separation between GANs (top) and explicit density methods (bottom).



Figure 4: Interpolation in the latent space between two samples from a model trained on CelebA-HQ.

coder and decoder. The reconstruction loss is chosen to be the pixel-wise mean square error (MSE). For more details, see Appendix 9.2.2.

**CIFAR-10 dataset:** We evaluate both the class-conditional and unconditional settings using the architecture of the original IntroVAE [25]. For the unconditional setting we report FID of 4.6 and qualitative samples and reconstructions of unseen data are displayed in Figure 5. Evidently, our model is able to generate and reconstruct high-quality samples from the various classes of CIFAR-10. For the class-conditional setting, we report FID of 4.07 and qualitative samples can be found in Appendix 9.4.2 along with more samples from the unconditional model.


**CelebA-HQ and FFHQ datasets:** For both datasets we downscale the images to 256x256 resolution. We use the style-based architecture with a latent and style dimension of 512. Our model is capable of generating high-quality samples and faithfully reconstructing unseen samples, as demonstrated in Figure 1 and Figure 14 in the supplementary. We provide more samples in Appendix 9.4.2. Table 3 quantitatively compares our method’s performance to various GANs and explicit density methods. Evidently, S-IntroVAE outperforms all previous autoencoding-based models, further narrowing the gap to GANs.

**Interpolation in the latent space:** Figure 4 shows smooth interpolation between the latent vectors of two images from S-IntroVAE trained on the CelebA-HQ dataset. We provide additional interpolations in Appendix 9.4.3.

### 6.4. Image Translation

To demonstrate the advantage of our model’s inference capability, we evaluate S-IntroVAE on image translation –





(a) Generated samples (FID: 4.6). (b) Reconstructions on test data: Left: real, right: reconstruction.


Figure 5: Results for the CIFAR-10 dataset (unconditional).

learning disentangled representations for *class* and *content*, and transferring content between classes (e.g., given two images of cars from different visual classes, rotate the first car to be in the angle of the second car, without altering the car’s other visual properties). In a typical *class-supervised* image translation setting, class labels are used to encourage class-content separation. The current SOTA in this setting is LORD [12]. Our focus here, however, is on the *unsupervised* image translation setting, where no labels are used at any point. A recent study [38] showed that unsupervised learning of disentangled representations can only work by exploiting some form of inductive bias. Here, we claim that the encoder architecture in LORD effectively adds such a strong inductive bias – we show that when coupled with our training method, we can achieve unsupervised image translation results that come close to the supervised SOTA.

We adopt the two-encoder architecture proposed in LORD, where one encoder is for the class and the other for the content. LORD’s decoder uses adaptive Instance Normalization [26] to align the mean and variance of the content features with those of the class features, as depicted in Figure 8b. In our proposed architecture, the ELBO is the sum of the reconstruction error and two KL terms:  $KL[q_{\phi_{content}}(z | x) || p(z)] + KL[q_{\phi_{class}}(z | x) || p(z)]$ . The separation to two encoders imposes strong inductive bias, as the model explicitly learns different representations for the class and content. Following the implementation in [12], we replace the pixel-wise MSE reconstruction loss with the VGG perceptual loss [24].

We quantitatively evaluate our method on the Cars3D dataset [46], where the *class* corresponds to the car model and the *content* is the azimuth and elevation. We follow [12] and measure content transfer in terms of perceptual similarity by Learned Perceptual Image Patch Similarity (LPIPS) [57]. Unlike previous methods that use some kind of supervision signal (e.g., class label), our method does not require such signals. As demonstrated qualitatively in Figure 6a and quantitatively in Table 4, our method outperforms most of the supervised methods, and narrows the gap to the SOTA supervised approach. We present additional qualitative results on the KTH dataset [48] in Figure 6b. Further details can be found in Appendix 9.3.

Interestingly, [12] showed that even with the architec-



(a) Cars3D (b) KTH

Figure 6: Qualitative results for content transfer on test data from the Cars3D and KTH datasets. The  $4 \times 4$  bottom-right matrix of images is generated according to a class from images in the left column and content of images in the top row. It is recommended to zoom-in.

Szabó et al. [51]	Supervised	0.137
ML-VAE [2]	Supervised	0.132
Cycle-VAE [19]	Supervised	0.141
DrNet [9]	Supervised	0.095
LORD [12]	Supervised	<b>0.078</b>
SoftIntroVAE (Ours)	Unsupervised	<b>0.084</b>

Table 4: Content transfer reconstruction error (LPIPS, lower is better) on the Cars3D dataset.

ture described above, a standard VAE struggles with learning disentangled representations due to vanishing of the KL term (posterior collapse). Our study shows that when trained using our introspective manner, the model is able to overcome this issue, demonstrating an additional benefit of the introspective approach.

## 7. Conclusion

Introspective VAEs narrow the gap between VAEs and GANs in terms of sampling quality, while still enjoying the favorable traits of variational inference models, such as amortized inference. In this work, we proposed the SoftIntroVAE – a modification of IntroVAE that is stable to train and simpler to analyze. Our investigation resulted in new insights on introspective training, and our experiments demonstrate competitive image generation results.

We see great potential in introspective models, as they open the door for using high quality generative models in applications that also require fast and high-quality inference.

## 8. Acknowledgements

This work is partly funded by the Israel Science Foundation (ISF-759/19) and the Open Philanthropy Project Fund, an advised fund of Silicon Valley Community Foundation.



## References

- [1] <https://github.com/hhb072/IntroVAE>. 6
- [2] Diane Bouchacourt, Ryota Tomioka, and Sebastian Nowozin. Multi-level variational autoencoder: Learning disentangled representations from grouped observations. *arXiv preprint arXiv:1705.08841*, 2017. 1, 8
- [3] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, 2016. 20
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 7
- [5] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Neural photo editing with introspective adversarial networks. *arXiv preprint arXiv:1609.07093*, 2016. 2
- [6] Marek Capiński and Peter Ekkehard Kopp. *The Radon—Nikodym theorem*, pages 187–240. Springer London, London, 2004. 27
- [7] Tal Daniel, Thanard Kurutach, and Aviv Tamar. Deep variational semi-supervised novelty detection. *arXiv preprint arXiv:1911.04971*, 2019. 4
- [8] Nicola De Cao, Ivan Titov, and Wilker Aziz. Block neural autoregressive flow. *35th Conference on Uncertainty in Artificial Intelligence (UAI19)*, 2019. 5
- [9] Emily L Denton et al. Unsupervised learning of disentangled representations from video. In *Advances in neural information processing systems*, pages 4414–4423, 2017. 8
- [10] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016. 2, 3
- [11] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Olivier Mastropietro, Alex Lamb, Martin Arjovsky, and Aaron Courville. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016. 1, 2, 3
- [12] Aviv Gabbay and Yedid Hoshen. Demystifying inter-class disentanglement. In *International Conference on Learning Representations*, 2019. 8, 13, 15, 20
- [13] Partha Ghosh, Mehdi SM Sajjadi, Antonio Vergari, Michael Black, and Bernhard Schölkopf. From variational to deterministic autoencoders. *arXiv preprint arXiv:1903.12436*, 2019. 2
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1, 4
- [15] Anirudh Goyal, Alessandro Sordani, Marc-Alexandre Côté, Nan Rosemary Ke, and Yoshua Bengio. Z-forcing: Training stochastic recurrent networks, 2017. 20
- [16] Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018. 5
- [17] David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2018. 1
- [18] Tian Han, Erik Nijkamp, Linqi Zhou, Bo Pang, Song-Chun Zhu, and Ying Nian Wu. Joint training of variational auto-encoder and latent energy-based model. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2020. 2
- [19] Ananya Harsh Jha, Saket Anand, Maneesh Singh, and VSR Veeravasaru. Disentangling factors of variation with cycle-consistent variational auto-encoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 805–820, 2018. 8
- [20] Serhii Havrylov and Ivan Titov. Preventing posterior collapse with levenshtein variational autoencoder, 2020. 20
- [21] Ari Heljakka, Arno Solin, and Juho Kannala. Pioneer networks: Progressively growing generative autoencoder. *Lecture Notes in Computer Science*, page 22–38, 2019. 2, 3, 7
- [22] Ari Heljakka, Arno Solin, and Juho Kannala. Towards photographic image manipulation with balanced growing of generative autoencoders. *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, Mar 2020. 2, 7
- [23] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner.  $\beta$ -vae: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2(5):6, 2017. 5
- [24] Yedid Hoshen, Ke Li, and Jitendra Malik. Non-adversarial image synthesis with generative latent nearest neighbors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5811–5819, 2019. 8, 15
- [25] Huaibo Huang, Zhihang Li, Ran He, Zhenan Sun, and Tieniu Tan. Introvae: Introspective variational autoencoders for photographic image synthesis. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, page 52–63, Red Hook, NY, USA, 2018. Curran Associates Inc. 1, 2, 3, 4, 5, 6, 7, 11, 13
- [26] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. 8, 14, 15
- [27] Dimitris Kalatzis, David Eklund, Georgios Arvanitidis, and Søren Hauberg. Variational autoencoders with riemannian brownian motion priors, 2020. 2
- [28] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation, 2017. 7, 13, 15
- [29] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 7, 13, 14, 15
- [30] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 2, 7
- [31] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 12 2014. 12

- [32] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in neural information processing systems*, pages 10215–10224, 2018. [7](#)
- [33] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2014. [1](#), [3](#), [4](#)
- [34] N Kodali, J Abernethy, J Hays, and Z Kira. On convergence and stability of gans. arxiv 2017. *arXiv preprint arXiv:1705.07215*, 2017. [1](#)
- [35] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. [13](#)
- [36] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric, 2015. [1](#), [2](#), [3](#)
- [37] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. [13](#)
- [38] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schoelkopf, and Olivier Bachem. A sober look at the unsupervised learning of disentangled representations and their evaluation. *Journal of Machine Learning Research*, 21(209):1–62, 2020. [2](#), [8](#)
- [39] Teng Long, Yanshuai Cao, and Jackie Chi Kit Cheung. Preventing posterior collapse in sequence vaes with pooling, 2019. [20](#)
- [40] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015. [1](#), [2](#), [3](#)
- [41] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? *arXiv preprint arXiv:1801.04406*, 2018. [1](#)
- [42] Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002. [28](#)
- [43] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017. [12](#)
- [44] Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14104–14113, 2020. [1](#), [2](#), [3](#), [7](#), [13](#), [14](#), [15](#)
- [45] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in Neural Information Processing Systems*, pages 14866–14876, 2019. [2](#)
- [46] Scott E Reed, Yi Zhang, Yuting Zhang, and Honglak Lee. Deep visual analogy-making. In *Advances in neural information processing systems*, pages 1252–1260, 2015. [8](#), [15](#)
- [47] Edgar Schonfeld, Bernt Schiele, and Anna Khoreva. A u-net based discriminator for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8207–8216, 2020. [7](#)
- [48] Christian Schuldt, Ivan Laptev, and Barbara Caputo. Recognizing human actions: a local svm approach. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 3, pages 32–36. IEEE, 2004. [8](#), [15](#)
- [49] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. Ladder variational autoencoders. In *Advances in neural information processing systems*, pages 3738–3746, 2016. [2](#)
- [50] Akash Srivastava, Lazar Valkov, Chris Russell, Michael U Gutmann, and Charles Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. In *Advances in Neural Information Processing Systems*, pages 3308–3318, 2017. [2](#), [3](#)
- [51] Attila Szabó, Qiyang Hu, Tiziano Portenier, Matthias Zwicker, and Paolo Favaro. Challenges in disentangling independent factors of variation. *arXiv preprint arXiv:1711.02245*, 2017. [8](#)
- [52] Jakub Tomczak and Max Welling. Vae with a vampprior. In *International Conference on Artificial Intelligence and Statistics*, pages 1214–1223, 2018. [2](#), [3](#)
- [53] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. It takes (only) two: Adversarial generator-encoder networks. *arXiv preprint arXiv:1704.02304*, 2017. [2](#), [3](#)
- [54] Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder, 2020. [2](#)
- [55] Hongteng Xu, Dixin Luo, Ricardo Henao, Svati Shah, and Lawrence Carin. Learning autoencoders with relational regularization, 2020. [2](#)
- [56] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. [16](#)
- [57] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [8](#)
- [58] Junbo Zhao, Michael Mathieu, and Yann LeCun. Energy-based generative adversarial network, 2016. [1](#), [3](#)