

---

# Simple and Effective VAE Training with Calibrated Decoders

---

Oleh Rybkin<sup>1</sup> Kostas Daniilidis<sup>1</sup> Sergey Levine<sup>2</sup>

## Abstract

Variational autoencoders (VAEs) provide an effective and simple method for modeling complex distributions. However, training VAEs often requires considerable hyperparameter tuning to determine the optimal amount of information retained by the latent variable. We study the impact of calibrated decoders, which learn the uncertainty of the decoding distribution and can determine this amount of information automatically, on the VAE performance. While many methods for learning calibrated decoders have been proposed, many of the recent papers that employ VAEs rely on heuristic hyperparameters and ad-hoc modifications instead. We perform the first comprehensive comparative analysis of calibrated decoder and provide recommendations for simple and effective VAE training. Our analysis covers a range of image and video datasets and several single-image and sequential VAE models. We further propose a simple but novel modification to the commonly used Gaussian decoder, which computes the prediction variance analytically. We observe empirically that using heuristic modifications is not necessary with our method. Project website: <https://orybkin.github.io/sigma-vae/>.

## 1. Introduction

Deep density models based on the variational autoencoder (VAE) (Kingma & Welling, 2014; Rezende et al., 2014) have found ubiquitous use in probabilistic modeling and representation learning as they are both conceptually simple and are able to scale to very complex distributions and large datasets. These VAE techniques are used for tasks such as future frame prediction (Castrejon et al., 2019), image segmentation (Kohl et al., 2018), generating speech (Chung et al., 2015) and music (Dhariwal et al., 2020), as well as model-based reinforcement learning (Hafner et al., 2019a).

<sup>1</sup>University of Pennsylvania <sup>2</sup>UC Berkeley. Correspondence to: Oleh Rybkin <oleh@seas.upenn.edu>.

However, in practice, many of these approaches require careful manual tuning of the balance between two terms that correspond to distortion and rate from information theory (Alemi et al., 2017). This balance trades off fidelity of reconstruction and quality of samples from the model: a model with low rate would not contain enough information to reconstruct the data, while allowing the model to have high rate might lead to unrealistic samples from the prior as the KL-divergence constraint becomes weaker (Alemi et al., 2017; Higgins et al., 2017). While a proper variational lower bound does not expose any free parameters to control this tradeoff, many prior works heuristically introduce a weight on the prior KL-divergence term, often denoted  $\beta$ . Usually,  $\beta$  needs to be tuned for every dataset and model variant as a hyperparameter, which slows down development and can lead to poor performance as finding the optimal value is often prohibitively computationally expensive. Moreover, using  $\beta \neq 1$  precludes the appealing interpretation of the VAE objective as a principled bound on the data likelihood, and is undesirable for applications like density modeling.

While many architectures for calibrating decoders have been proposed in the literature (Kingma & Welling, 2014; Kingma et al., 2016; Dai & Wipf, 2019), more applied work typically employs VAEs with uncalibrated decoding distributions, such as Gaussian distributions without a learned variance, where the decoder only outputs the mean parameter (Castrejon et al., 2019; Denton & Fergus, 2018; Lee et al., 2019; Babaeizadeh et al., 2018; Lee et al., 2018; Hafner et al., 2019b; Pong et al., 2019; Zhu et al., 2017; Pavlakos et al., 2019), or uses other ad-hoc modifications to the objective (Sohn et al., 2015; Henaff et al., 2019). Indeed, it is well known that naively attempting to learn the variance in a Gaussian decoder leads to poor results as the variance may converge to zero (Rezende & Viola, 2018; Dai & Wipf, 2019). As a result, it remains unclear whether practical empirical performance of VAEs actually benefits from calibrated decoders or not.

To rectify this, our first contribution is a comparative analysis of various calibrated decoder architectures and practical recommendations for simple and effective VAE training. We find that, while naïve calibrated decoders often lead to worse results, a careful choice of the decoder distribution can work very well, and removes the need to tune the additional parameter  $\beta$ . Indeed, we note that the entropy of

the decoding distribution controls the mutual information  $I(x; z)$ . Calibrated decoders allow the model to control  $I(x; z)$  automatically, instead of relying on manual tuning. Our second contribution is a simple but novel technique for optimizing the decoder variance analytically, without requiring the decoder network to produce it as an additional output. We call the resulting approach to learning the Gaussian variance the  $\sigma$ -VAE. In our experiments, the  $\sigma$ -VAE outperforms the alternative of learning the variance through gradient descent, while being simpler to implement and extend. We validate our results on several VAE and sequence VAE models and a range of image and video datasets.

## 2. Related Work

Prior work on variational autoencoders has studied a number of different decoder parameterizations. Kingma & Welling (2014); Rezende et al. (2014) use the Bernoulli distribution for the binary MNIST data and Kingma & Welling (2014) use Gaussian distributions with learned variance parameter for grayscale images. However, modeling images with continuous distributions is prone to instability as the variance can converge to zero (Rezende & Viola, 2018; Mattei & Frellsen, 2018; Dai & Wipf, 2019). Some work has attempted to rectify this problem by using dequantization (Gregor et al., 2016), which is theoretically appealing as it is tightly related to the log-likelihood of the original discrete data (Theis et al., 2016), optimizing the variance in a two-stage procedure (Arvanitidis et al., 2017), or training a post-hoc prior (Engel et al., 2017; Ghosh et al., 2019). (Takahashi et al., 2018; Barron, 2019) proposed more expressive distributions. Additionally, different choices for representing such variance exist, including diagonal covariance (Kingma & Welling, 2014; Sønderby et al., 2016; Rolfe, 2016), or a single shared parameter (Kingma et al., 2016; Dai & Wipf, 2019; Edwards & Storkey, 2016; Rezende & Viola, 2018). We analyze these and notice that learning a single variance parameter shared across images leads to stable training and good performance, without the use of dequantization or even clipping the variance, although these techniques can be used with our decoders; and further improve the estimation of this variance with an analytic solution.

Early work on discrete VAE decoders for color images modeled them with the Bernoulli distribution, treating the color intensities as probabilities (Gregor et al., 2015). Further work has explored various parameterizations based on discretized continuous distributions, such as discretized logistic (Kingma et al., 2016). More recent work has improved expressivity of the decoder with a mixture of discretized logistics (Chen et al., 2016; Maaløe et al., 2019). However, these models also employ powerful autoregressive decoders (Chen et al., 2016; Gulrajani et al., 2016; Maaløe et al., 2019), and the latent variables in these models may not rep-

resent all of the significant factors of variation in the data, as some factors can instead be modeled internally by the autoregressive decoder (Alemi et al., 2017).<sup>1</sup>

While many calibrated decoders have been proposed, outside the core generative modeling community uncalibrated decoders are ubiquitous. They are used in works on video prediction (Denton & Fergus, 2018; Castrejon et al., 2019; Lee et al., 2018; Babaeizadeh et al., 2018), image segmentation (Kohl et al., 2018), image-to-image translation (Zhu et al., 2017), 3D human pose (Pavlakos et al., 2019), as well as model-based reinforcement learning (Henaff et al., 2019; Hafner et al., 2019b;a), and representation learning (Lee et al., 2019; Watter et al., 2015; Pong et al., 2019). Most of these works utilize the heuristic hyperparameter  $\beta$  instead, which is undesirable both as the resulting objective is no longer a bound on the likelihood, and as  $\beta$  usually requires extensive tuning. In this work, we analyze the common pitfalls of using calibrated decoders that may have prevented the practitioners from using them, propose a simple and effective analytic way of learning such calibrated distribution, and provide a comprehensive experimental evaluation of different decoding distributions.

Alternative discussions of the hyperparameter  $\beta$  are presented by Zhao et al. (2017); Higgins et al. (2017); Alemi et al. (2017); Achille & Soatto (2018), who show that it controls the amount of information in the latent variable,  $I(x; z)$ . Peng et al. (2018); Rezende & Viola (2018) further discuss constrained optimization objectives for VAEs, which also yield a similar hyperparameter. Here, we focus on  $\beta$ -VAEs with Gaussian decoders with constant variance, as commonly used in recent work, and show that the hyperparameter  $\beta$  can be incorporated in the decoding likelihood for these models.

## 3. Analysing Decoding Distributions

The generative model of a VAE (Kingma & Welling, 2014; Rezende et al., 2014) with parameters  $\theta$  is specified with a prior distribution over the latent variable  $p_\theta(z)$ , commonly unit Gaussian, and a decoding distribution  $p_\theta(x|z)$ , which for color images is commonly a conditional Gaussian parameterized with a neural network. We would like to fit this generative model to a given dataset by maximizing the evidence lower bound (ELBO, (Neal & Hinton, 1998; Jordan et al., 1999; Kingma & Welling, 2014; Rezende et al., 2014)), which uses an approximate posterior distribution  $q_\phi(z|x)$ , also commonly a conditional Gaussian specified

<sup>1</sup>BIVA (Maaløe et al., 2019) uses the Mixture of Logistics decoder proposed in (Salimans et al., 2017) that produces the channels for each pixel autoregressively, see also App E.

with a neural network:

$$\ln p_\theta(x) \geq \mathbb{E}_{q_\phi(z|x)} [\ln p_\theta(x|z)] + \mathbf{D}_{\text{KL}}(q_\phi(z|x)||p_\theta(z)) \quad (1)$$

In this work, we focus on the form of the decoding distribution  $p_\theta(x|z)$ . To achieve the best results, we want a decoding distribution that represents the required probability  $p(x|z)$  accurately and is *calibrated* in the statistical sense. In this section, we will review and analyze various choices of decoding distributions that enable better decoder calibration, including expressive decoding distributions that can represent both the prediction of the image and the uncertainty about such prediction, or even multimodal predictions.

### 3.1. Gaussian Decoders

We first analyse the commonly used Gaussian decoders. We note that the commonly used MSE reconstruction loss between the reconstruction  $\hat{x}$  and ground truth data  $x$  is equivalent to the negative log-likelihood objective with a Gaussian decoding distribution with constant variance:

$$\begin{aligned} -\ln p(x|z) &= \frac{1}{2} \|\hat{x} - x\|^2 + D \ln \sqrt{2\pi} = \frac{1}{2} \|\hat{x} - x\|^2 + c \\ &= \frac{D}{2} \text{MSE}(\hat{x}, x) + c, \quad (2) \end{aligned}$$

where  $p(x|z) \sim \mathcal{N}(\hat{x}, I)$ ,  $\hat{x}$  is produced with a neural network  $\hat{x} = \mu_\theta(z)$ , and  $D$  is the dimensionality of  $x$ .

This demonstrates a drawback of methods that rely simply on the MSE loss (Castrejon et al., 2019; Denton & Fergus, 2018; Lee et al., 2019; Hafner et al., 2019b; Pong et al., 2019; Zhu et al., 2017; Henaff et al., 2019), as it is equivalent to assuming a particular, constant variance of the Gaussian decoding distribution. By learning this variance, we can achieve much better performance due to better calibration of the decoder. There are several ways in which we can specify this variance. An expressive way to specify the variance is to specify a diagonal covariance matrix for the image, with one value per pixel (Kingma & Welling, 2014; Sønderby et al., 2016; Rolfe, 2016). This can be done, for example, by letting a neural network  $\sigma_\theta$  output the diagonal entries of the covariance matrix given a latent sample  $z$ :

$$p_\theta(x|z) = \mathcal{N}(\mu_\theta(z), \sigma_\theta(z)^2). \quad (3)$$

This parameterization of the decoding distribution outputs one variance value per each pixel and channel. While powerful, we observe in Section 5.3 that this approach attains suboptimal performance, and is moreover prone to numerical instability. Instead, we will find experimentally that a simpler parameterization, in which the covariance matrix is specified with a single *shared* (Kingma et al., 2016; Dai & Wipf, 2019; Edwards & Storkey, 2016; Rezende & Viola, 2018) parameter  $\sigma$  as  $\Sigma = \sigma I$  often works better in practice:

$$p_{\theta,\sigma}(x|z) = \mathcal{N}(\mu_\theta(z), \sigma^2 I). \quad (4)$$

The parameter  $\sigma$  can be optimized together with parameters of the neural network  $\theta$  with gradient descent. Of particular interest is the interpretation of this parameter. Writing out the expression for the decoding likelihood, we obtain

$$\begin{aligned} -\ln p(x|z) &= \frac{1}{2\sigma^2} \|\hat{x} - x\|^2 + D \ln \sigma \sqrt{2\pi} = \\ \frac{1}{2\sigma^2} \|\hat{x} - x\|^2 + D \ln \sigma + c &= D \ln \sigma + \frac{D}{2\sigma^2} \text{MSE}(\hat{x}, x) + c. \quad (5) \end{aligned}$$

The full objective of the resulting Gaussian  $\sigma$ -VAE is:

$$\mathcal{L}_{\theta,\phi,\sigma} = D \ln \sigma + \frac{D}{2\sigma^2} \text{MSE}(\hat{x}, x) + D_{\text{KL}}(q(z|x)||p(z)). \quad (6)$$

Note that  $\sigma$  may be viewed as a weighting parameter between the MSE reconstruction term and the KL-divergence term in the objective. Moreover, this objective explicitly specifies how to select the optimal variance: the variance should be selected to minimize the (weighted) MSE loss while also minimizing the logarithm of the variance.

**Decoder Calibration** It is important that the decoder distribution be calibrated in the statistical sense, that is, the predicted probabilities correspond to the frequencies of seeing a particular value of  $x$  given that prediction (DeGroot & Fienberg, 1983; Dawid, 1982). The calibration of a neural network can be usually improved by estimating the prediction uncertainty (Guo et al., 2017), such as the variance of a Gaussian (Kendall & Gal, 2017). Since the naive MSE loss assumes a constant variance, it does not effectively represent the uncertainty of the prediction, and is often poorly calibrated. Instead, learning the variance as in Eq. 6 leads to better uncertainty estimation and better calibration. In Sec 5.1, we show that learning a good estimate of this uncertainty is crucial for the quality of the VAE generations.

**Connection to  $\beta$ -VAE.** The  $\beta$ -VAE objective (Higgins et al., 2017) for a Gaussian decoder with unit variance is:

$$\mathcal{L}^\beta = \frac{D}{2} \text{MSE}(\hat{x}, x) + \beta D_{\text{KL}}(q(z|x)||p(z)). \quad (7)$$

We see that it can be interpreted as a particular case of the objective (6), where the variance is constant and the term  $D \ln \sigma$  can be ignored during optimization. The  $\beta$ -VAE objective is then equivalent to a  $\sigma$ -VAE with a constant variance  $\sigma^2 = \beta/2$  (for a particular learning rate setting). In recent work (Zhu et al., 2017; Denton & Fergus, 2018; Lee et al., 2019),  $\beta$ -VAE models are often used in this exact regime. By tuning the  $\beta$  term, practitioners are able to tune the variance of the decoder, manually producing a more calibrated decoder. However, by re-interpreting the  $\beta$ -VAE objective as a special case of the VAE and introducing the missing  $D \ln \sigma$  term, we can both obtain a valid evidence




Figure 1. Different types of calibrated decoders for Gaussian VAE, model parameters are denoted with enclosing squares. Left: both the mean  $\mu$  and the variance  $\sigma$  are output by a neural network with parameters  $\theta$ . Center:  $\sigma$ -VAE with shared variance, the mean is output by a neural network with parameters  $\theta$ , but the variance is itself a global parameter. Right: the proposed optimal  $\sigma$ -VAE, the mean is output by a neural network with parameters  $\theta$ , and the variance is computed analytically from the training data  $D$ .

lower bound, and remove the need to manually select  $\beta$ . Instead, the variance  $\sigma$  can instead simply be learned end-to-end, reducing the need for hyperparameter tuning.

An alternative discussion of this connection in the context of linear VAEs is also presented by Lucas et al. (2019). While the  $\beta$  term is not necessary for good performance if the decoder is calibrated, it can still be employed if desired, such as when the aim is to attain better disentanglement (Higgins et al., 2017) or a particular rate-distortion tradeoff (Alemi et al., 2017). However, we found that with calibrated decoders, the best sample quality is obtained when  $\beta = 1$ .

**Loss implementation details.** For the correct evidence lower bound computation, it is necessary to add the values of the MSE loss and the KL divergence across the dimensions. We observe that common implementations of these losses (Denton & Fergus, 2018; Abadi et al., 2016; Paszke et al., 2019) use averaging instead, which will lead to poor results if the number of image dimensions is significantly different from the number of the latent dimensions. While this can be conveniently ignored in the  $\beta$ -VAE regime, where the balance term is tuned manually anyway, for the  $\sigma$ -VAE it is essential to compute the objective value correctly.

**Variance implementation details.** Since the variance is non-negative, we parameterize it logarithmically as  $\sigma^2 = e^{2\lambda}$ , where  $\lambda$  is the logarithm of the standard deviation. For some models, such as per-pixel variance decoders, we observed that it is necessary to restrict the variance range for numerical stability. We do so by using the soft clipping operations proposed by Chua et al. (2018):

$$\begin{aligned}\lambda &:= \lambda_{\max} - \text{softplus}(\lambda_{\max} - \lambda), \\ \lambda &:= \lambda_{\min} + \text{softplus}(\lambda - \lambda_{\min}).\end{aligned}$$

We observe that setting  $\lambda_{\min} = -6$  to lower bound the standard deviation to be at least half of the distance between

allowed color values works well in practice. We also observe this clipping is unnecessary when learning a shared  $\sigma$  value.

### 3.2. Discrete Decoders

It is possible to use discrete decoding distributions to generate images, as color values are commonly restricted to a fixed set of integer pixel intensities (e.g. 0..255). Indeed, for discrete color values, discrete distributions are arguably more appropriate. In the most general case, a discrete decoding distribution factorized per each pixel and channel would be specified by a categorical distribution with a probability mass vector  $\hat{x}$  with 256 entries, one per each possible intensity value, similarly to a per-pixel classifier of the intensity value. We can implement it with a soft-max layer, yielding the following log-likelihood loss (sometimes called the cross-entropy loss) for a true pixel with intensity  $i$ :

$$-\ln p(x|z) = -\ln \frac{\exp(\hat{x}_i)}{\sum_j \exp(\hat{x}_j)},$$

We will evaluate these and further choices of discrete decoders, described in Appendix E.

## 4. Optimal Variance Estimation for Calibrated Gaussian Decoders

In this section, we propose a simple but novel analytic way of obtaining a calibrated decoder for continuous distributions that further improves performance. The Gaussian decoders with learned variance described in Section 3.1 are calibrated and work better than naïve unit variance decoders. However, for  $\sigma$ -VAE optimized with gradient descent or Adam (Kingma & Ba, 2015), we observe that careful learning rate tuning can yield significantly better performance, which is in line with prior work that reported poor performance of gradient descent for optimizing Gaussian distributions (Amari, 1998; Peters & Schaal, 2008). A smaller




Figure 2. Images or videos (right) sampled from the proposed optimal  $\sigma$ -VAE and a  $\beta$ -VAE models. The parameter  $\beta$  is tuned manually in  $\beta$ -VAE and learned in  $\sigma$ -VAE. Higher values of  $\beta$  cause the images to lose detail, while lower values of  $\beta$  might make samples unrealistic. The proposed optimal  $\sigma$ -VAE is able to learn the balance end-to-end on all datasets, without requiring tuning the  $\beta$  hyperparameter

learning rate often produces better performance, but slows down the training, as the likelihood values  $p(x|z)$  will be very suboptimal in the beginning. Instead, here we propose an analytic solution for the value of  $\sigma$ , which computes it analytically and does not require gradient descent.

The maximum likelihood estimate of the variance given a known mean is the average squared distance from the mean:

$$\sigma^{*2} = \arg \max_{\sigma^2} \mathcal{N}(x|\mu, \sigma^2 I) = \text{MSE}(x, \mu), \quad (8)$$

where  $\text{MSE}(x, \mu) = \frac{1}{D} \sum_i (x_i - \mu_i)^2$ . Eq. 8 can be easily shown using manual differentiation, and is a generalization of the fact that the MLE estimate of the variance is the sample variance.

The optimal variance for the decoder distribution under the maximum likelihood criterion is then simply the average MSE loss over the data and the encoder distribution. We leverage this to create an optimal analytic solution for the variance. In the batch setting, the optimal variance would be simply the MSE loss, and can be updated after every gradient update for the other parameters of the decoder. In the mini-batch setting, we use a batchwise estimate of the variance computed for the current minibatch. We analyze these approximations in Appendix C. At test time, a running average of the variance over the training data is used. This method, which we call *optimal  $\sigma$ -VAE*, allows us to learn very efficiently as we use the optimal variance estimate at every training step. It is also easier to implement, the

variance parameter can simply be computed from the current batch and there is often no need to separately track it. If the variance is not needed at test time, it can also be simply discarded after training.

**Per-image optimal  $\sigma$ -VAE.** Optimal  $\sigma$ -VAE uses a single variance value shared across all data points. However, the optimal  $\sigma$ -VAE also allows more powerful variance estimates, such as learning a variance value per each pixel, or even a variance value per each image, the difference in implementation simply being the dimensions across which the averaging in Equation 8 operates. This approach can be interpreted as variational variance prediction in the framework of (Stirn & Knowles, 2020), details in App. ??.

## 5. Experimental Results

We now provide an empirical analysis of different decoding distributions, and validate the benefits of our  $\sigma$ -VAE approach. We use a small convolutional VAE model on SVHN (Netzer et al., 2011), a larger hierarchical HVAE model (Maaløe et al., 2019) on the CelebA (Liu et al., 2015) and CIFAR (Krizhevsky et al., 2009) datasets, and a sequence VAE model called SVG (Denton & Fergus, 2018) on the BAIR Pushing dataset (Finn & Levine, 2017). We evaluate the ELBO values as well as visual quality measured by the Fréchet Inception Distance (FID, (Heusel et al., 2017)). Images are  $28 \times 28$  for SVHN and  $32 \times 32$  for CelebA and CIFAR, while video experiments were performed on

Table 1. Analysis of learned variance on SVHN. The parameter  $\beta$  is tuned manually in  $\beta$ -VAE and learned in  $\sigma$ -VAE.  $\sigma$ -VAE achieves better performance, while the value of  $\beta$  (implicitly defined via the decoder variance) automatically converges close the value found by manual tuning.

	$\beta$	$-\log p \downarrow$	FID $\downarrow$
$\beta$ -VAE	0.001	< 21.43	44.54
$\beta$ -VAE	0.01	< -3186	27.93
$\beta$ -VAE	0.1	< -1223	28.3
$\beta$ -VAE	1	< 1381	70.39
$\beta$ -VAE	10	< 4056	219.3
$\sigma$ -VAE	0.006	< -3333	<b>22.25</b>

64  $\times$  64 frames following (Denton & Fergus, 2018). We do not use KL annealing as it did not improve the results in our experiments. Further experimental details are in App. B.

### 5.1. Do calibrated decoders balance the VAE objective without tuning $\beta$ ?

As detailed in Section 3.1, a  $\beta$ -VAE with a unit variance Gaussian decoder commonly used in prior work is equivalent to a  $\sigma$ -VAE with constant, manually tuned variance. There is a simple relationship between beta and the variance:  $\sigma^2 = \beta/2$ . To compare the variance that the  $\sigma$ -VAE learns to the manually tuned variance in the case of the  $\beta$ -VAE, we compare the ELBO values and the corresponding values of  $\beta$  in Table 1. We find that learning the variance produces similar values of  $\beta$  to the manually tuned values in the  $\beta$ -VAE case, indicating that the  $\sigma$ -VAE is able to learn the balance between the two objective terms in a single training run, without hyperparameter tuning. Moreover, the  $\sigma$ -VAE outperforms the best  $\beta$ -VAE run. This is because end-to-end learning produces better estimates of the variance than is possible with manual search, improving the likelihood (as measured by the lower bound) and the visual quality. Figure 2 shows the qualitative results from this experiment.

We further validate our results on both single-image and sequential VAE models on a range of datasets in Table 2 and Figure 2. Single-sample ELBO values are reported, and ELBO values on discretized data are reported for discrete distributions. We see that learning a shared variance in a Gaussian decoders (shared  $\sigma$ -VAE) outperforms the naïve unit variance decoder (Gaussian VAE) as well as tuning the  $\beta$  constant for the Gaussian VAE manually. We also see that calibrated discrete decoders, such as full categorical distribution or mixture of discretized logistics, perform better than the naïve Gaussian VAE. Using Bernoulli distribution by treating the color intensities as probabilities (Gregor et al., 2015; Watter et al., 2015) performs poorly. Our results further improve upon the sequence VAE method of Denton & Fergus (2018), which uses a unit variance Gaussian with the

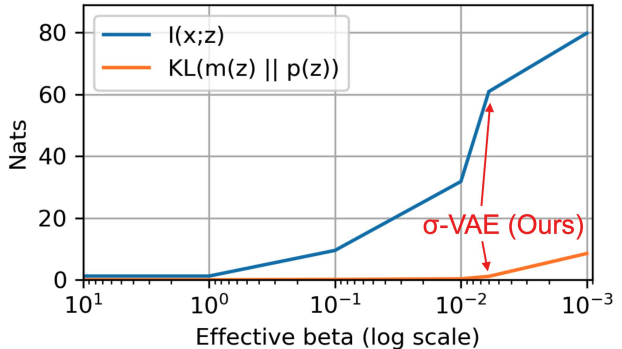


Figure 3. Comparison of  $\beta$ -VAE and  $\sigma$ -VAE on SVHN in terms of mutual information  $I_e(x; z)$  and marginal KL divergence  $KL(m(z)||p(z))$  (see Sec. 5.2).  $I_e(x; z)$  increases with lower  $\beta$ , yielding expressive representations and better reconstruction. However, after a certain point, lowering  $\beta$  leads to a rapid increase in the marginal KL, yielding poor samples from the prior. The  $\sigma$ -VAE is able to automatically find the inflection point after which the marginal KL begins to increase, capturing as much information as possible while still producing good samples.

$\beta$ -VAE objective.

### 5.2. How does learning calibrated decoders impact the latent variable information content?

We saw above that calibrated decoders result in higher log-likelihood bounds. Are calibrated decoders also beneficial for representation learning? We evaluate the mutual information  $I_e(x; z)$  between the data  $p_d(x)$  and encoder samples  $q(z|x)$ , as well as the mismatch between the prior  $p(z)$  and the marginal distribution  $m(z) = E_{p_d(x)}q(z|x)$ , measured by the marginal KL  $D_{KL}(m(z)||p(z))$ . These terms are related to the rate term of the VAE objective as follows (Alemi et al., 2017; Bozkurt et al., 2019):

$$\begin{aligned}
 & E_{p_d(x)} [D_{KL}(q(z|x)||p(z))] \\
 &= E_{p_d(x)} [D_{KL}(q(z|x)||m(z))] + D_{KL}(m(z)||p(z)) \\
 &= I_e(x; z) + D_{KL}(m(z)||p(z)).
 \end{aligned} \tag{9}$$

That is, the rate term decomposes into the true mutual information and the marginal KL term. We want to learn expressive latent variables with high mutual information. However, doing so by tuning the  $\beta$  value relaxes the constraint that the encoder and the prior distributions match, and leads to degraded quality of samples from the prior, which creates a trade-off between expressive representations and ability to generate good samples. To compare the  $\beta$ -VAE and  $\sigma$ -VAE in terms of these quantities, we estimate the marginal KL term via Monte Carlo sampling, as proposed by Rosca et al. (2018), and plot the results in Figure 3. As expected, we see that lower  $\beta$  values lead to higher mutual



Table 2. Generative modeling performance of the proposed  $\sigma$ -VAE on different models and datasets. For SVG, we compare with the original method (Denton & Fergus, 2018), which uses  $\beta$ -VAE. We see that uncalibrated decoders such as mean-only Gaussian perform poorly.  $\beta$ -VAE allows to calibrate the decoder but needs careful hyperparameter tuning. Calibrated decoders such as categorical or  $\sigma$ -VAE perform best. [1] (Gregor et al., 2015), [2] (Takahashi et al., 2018), [3] (Higgins et al., 2017).

	CelebA HVAE		SVHN VAE		CIFAR HVAE		BAIR SVG	
	$-\log p \downarrow$	FID $\downarrow$	$-\log p \downarrow$	FID $\downarrow$	$-\log p \downarrow$	FID $\downarrow$	$-\log p \downarrow$	FID $\downarrow$
Bernoulli VAE [1]		177.6		43.26		284.5		122.6
Categorical VAE	< <b>6359</b>	71.5	< 9179	46.13	< <b>7179</b>	<b>101.7</b>	N/A	N/A
Bitwise-categorical VAE	< 9067	66.61	< 10800	33.84	< 9390	<b>91.2</b>	< 48744	46.13
Logistic mixture VAE	< 7932	65.3	< <b>9085</b>	43.19	< 8443	143.1	< <b>40616</b>	42.94
Gaussian VAE	< 7173	186.5	< 2184	112.5	< 7186	293.7	< -10379	35.64
Per-pixel $\sigma$ -VAE	< -7814	159.3	< -3592	114.7	< -7222	131	< -14051	41.98
Student-t VAE [2]	< -8401	71.06	< - <b>3659</b>	70.4	< - <b>7419</b>	123.6	-	-
$\beta$ -VAE [3]	< -2713	<b>61.6</b>	< -3186	27.93	< -331	<b>103</b>	< -13472	34.64
Shared $\sigma$ -VAE	< -6374	<b>60.7</b>	< -3349	<b>22.25</b>	< -5435	116.1	< -13974	34.24
Optimal $\sigma$ -VAE	< - <b>8446</b>	<b>60.3</b>	< -3333	27.25	< -5677	<b>101.4</b>	< - <b>14173</b>	34.13
Opt. per-image $\sigma$ -VAE		66.01		26.28		<b>104.0</b>		<b>33.21</b>

information. However, after a certain point, lower values of  $\beta$  also cause a significant mismatch between the marginal and the prior distributions. By calculating the “effective”  $\beta$  for the  $\sigma$ -VAE, as per Section 4, we can see that the  $\sigma$ -VAE captures an inflection point in the  $D_{KL}(m(z)||p(z))$  term, learning a representation with the highest possible MI, but without degrading sample quality. This explains the high visual quality of the optimal  $\sigma$ -VAE samples: since the marginal and the prior distributions match, the samples from the prior look similar to reconstructions, while for a  $\beta$ -VAE with low  $\beta$ , the samples from the prior are poor. We see that, in contrast to the  $\beta$ -VAE, where the mutual information is controlled by a hyperparameter, the  $\sigma$ -VAE can adjust the appropriate amount of information automatically and is able to find the setting that produces both informative latents and high quality samples.

An alternative discussion of tuning  $\beta$  is presented by Alemi et al. (2017); Rezende & Viola (2018), who show that  $\beta$  controls the rate-distortion trade-off. Here, we show that the crucial trade-off also controlled by  $\beta$  is the trade-off between two components of the rate itself, which control expressivity of representations and the match between the variational and the prior distributions, respectively.

### 5.3. What are the common challenges in learning the variance that prevent practitioners from using it, and how to rectify them?

If learning the decoder variance improves generation, why are learned variances not used more often? In this section, we discuss how the naïve approach to learning variances, where the decoder outputs a variance for each pixel along with the mean, leads to poor results. First, we find that

this method often diverges very quickly due to numerical instability, as the network is able to predict certain pixels with very high certainty, leading to degenerate variances. In contrast, learning a shared variance is always numerically stable in our experiments. We can rectify this instability by bounding the output variance (Section 3.1). However, even with bounded variance, we observe that learning per-pixel variances leads to poor results in Table 2. While the per-pixel variance achieves a good ELBO value, it produces very poor samples, as measured by FID and visual inspection.

We see that the specific form of learned variance: a shared variance, a per-image variance, or a per-pixel variance, can lead to very different performance in practice. We hypothesize the per-pixel decoder performs poorly as it incentivizes the model to focus on particular pixels that can be predicted well, instead of focusing equally on all parts of the image. This is consistent with prior work on denoising diffusion models which noted that likelihood-based models place too much focus on imperceptible details, which leads to deteriorated results (Ho et al., 2020). The shared and per-image variance models mitigate this issue at the cost of introducing more bias (conf. Shu et al. (2018)), and work better in practice.

To evaluate why per-pixel sigma VAE performs poorly, we analyze the marginal KL and the visual quality of the generations. Further, we produce a series of models with partially shared  $\sigma$  by only sharing some dimensions. Figure 4 shows the marginal KL and MI with respect to number of variance parameters. We observe that with more expressive variance decoders, the visual quality and the marginal KL, which measures the quality of the prior distribution, both decrease and the models start producing unrealistic samples.

We note that while the marginal KL only grows up to 4 nats, this means that only 1 generated sample in 50 will be constrained to the prior (Theis et al., 2016). Given this and the results in Table 2, we observe that sharing the variance parameter across the image is a useful inductive bias which leads to better quality of generated images. We hypothesize that sharing the variance is effective because it prevents the model from focusing on particular pixels and forces it to focus on the whole image equally. In contrast, the per-pixel sigma-VAE may sacrifice image quality for improved results on just a few particular pixels.

This analysis confirms the result of prior work that focused on the case of multimodal decoders (Alemi et al., 2017; Rezende & Viola, 2018), and shows that expressive decoders can lead to poor learned representations even when they are unimodal, like a Gaussian decoder, but with expressive variance prediction. While with current methods learning a single shared variance works best, the ability to produce more fine-grained uncertainty will be desirable on more complex data with non-uniform dimensions. Our analysis reveals a shortcoming of current methods for producing calibrated uncertainties, wherein a more expressive uncertainty prediction may hinder performance, and highlights the need for more research in this direction.

#### 5.4. Can an analytic solution for optimal variance further improve learning?

We evaluate the *optimal*  $\sigma$ -VAE which uses an analytic solution for the variance (Section 4). Table 2 shows that it achieves superior results in terms of log-likelihood. We also note that the optimal  $\sigma$ -VAE converges to a good variance estimate instantaneously, which speeds up learning (highlighted in Figure 9 in the Appendix). In addition, we evaluate the per-image optimal  $\sigma$ -VAE, in which a single variance is computed per image. This model achieves significantly higher visual quality. While producing this per-image variance with a neural network would require additional architecture tuning, optimal  $\sigma$ -VAE is very simple to implement (it can be implemented simply as changing the axes of summation), not requiring any new tunable parameters.

We further analyse the the quality of the variance estimate of the optimal  $\sigma$ -VAE numerically. The optimal  $\sigma$ -VAE numerically requires computing the following estimate of the variance

$$\begin{aligned} \sigma^* &= \arg \max_{\sigma} \mathbb{E}_{x \sim \text{Data}} \mathbb{E}_{q(z|x)} [\ln p(x|\mu_{\theta}(z), \sigma^2 I)] \\ &= \mathbb{E}_{x \sim \text{Data}} \mathbb{E}_{q(z|x)} \text{MSE}(x, \mu_{\theta}(z)). \end{aligned} \quad (10)$$

This requires computing two expectations, with respect to the data in the dataset, and with respect to the encoder distribution. We use MC sampling to approximate both expectations. Inspired by common practices in VAEs, we use

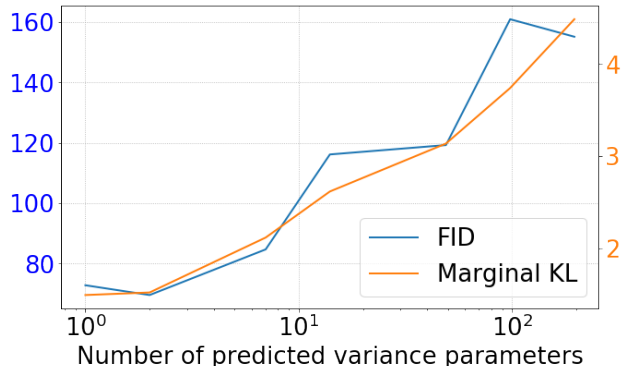


Figure 4. An analysis of number of predicted variance parameters on MNIST. Shared  $\sigma$ -VAE contains one parameter and is on one end of the spectrum, while the per-pixel  $\sigma$ -VAE has the full number of parameters and is on the other end of the spectrum. We see that with more expressive uncertainty prediction, visual quality and the quality of the prior measured by marginal KL quickly degrade. This observation highlights the good performance of shared  $\sigma$ -VAE and a shortcoming of current VAE methods that are not able to handle expressive uncertainty estimates well.

one sample per data point to approximate the inner expectation. On SVHN, the standard error of this approximation is 0.26% of the value of sigma. We further approximate the outer expectation with a single batch instead of the entire dataset. On SVHN, the standard error of this approximation is 2% of the value of sigma. We see that both approximations are accurate in practice. The second approximation yields a biased estimate of the evidence lower bound because the same batch is used to approximate the variance and compute the lower bound estimate. However, this bias can be corrected by using a different batch, or with a running average of the variance with an appropriate decay. This running average can also be used to reduce the variance of the estimate and to achieve convergence guarantees, but we did not find it necessary in our experiments.

## 6. Conclusion

We presented a simple and effective method for learning calibrated decoders, as well as an evaluation of different decoding distributions with several VAE and sequential VAE models. The proposed method outperforms methods that use naïve unit variance Gaussian decoders and tune a heuristic weight  $\beta$  on the KL-divergence loss, as commonly done in prior work. Moreover, it does not use the heuristic weight  $\beta$ , making it easier to train than this prior work. We expect that the simple techniques for learning calibrated decoders can allow practitioners to speed up the development cycle, obtain better results, and reduce the need for manual hyperparameter tuning.



## 7. Acknowledgements

We thank Vagelis Chatzipantazis, Marvin Zhang, Vitychyr Pong, Rahul Ramesh and Danijar Hafner for fruitful discussions, and Karl Pertsch for feedback on the Sequential VAE setting. We also thank Rahul Ramesh, Coline Devin, Michael Janner, Justin Fu, and Jorge Mendez on feedback on an earlier version of the paper. This work was supported through the following grants: ARL DCIST CRA W911NF-17-2-0181, NSF TRIPODS 1934960, and by Honda Research Institute.

## References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pp. 265–283, 2016.
- Achille, A. and Soatto, S. Information dropout: Learning optimal representations through noisy computation. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2897–2905, 2018.
- Alemi, A. A., Poole, B., Fischer, I., Dillon, J. V., Saurous, R. A., and Murphy, K. Fixing a broken elbow. *arXiv preprint arXiv:1711.00464*, 2017.
- Amari, S.-I. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276, 1998.
- Arvanitidis, G., Hansen, L. K., and Hauberg, S. Latent space oddity: on the curvature of deep generative models. *arXiv preprint arXiv:1710.11379*, 2017.
- Babaeizadeh, M., Finn, C., Erhan, D., Campbell, R. H., and Levine, S. Stochastic variational video prediction. 2018.
- Barron, J. T. A general and adaptive robust loss function. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4331–4339, 2019.
- Bozkurt, A., Esmaili, B., Tristan, J.-B., Brooks, D. H., Dy, J. G., and van de Meent, J.-W. Rate-regularization and generalization in vaes. *arXiv preprint arXiv:1911.04594*, 2019.
- Castrejon, L., Ballas, N., and Courville, A. Improved conditional vrns for video prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 7608–7617, 2019.
- Chen, X., Kingma, D. P., Salimans, T., Duan, Y., Dhariwal, P., Schulman, J., Sutskever, I., and Abbeel, P. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pp. 4754–4765, 2018.
- Chung, J., Kastner, K., Dinh, L., Goel, K., Courville, A. C., and Bengio, Y. A recurrent latent variable model for sequential data. 2015.
- Dai, B. and Wipf, D. Diagnosing and enhancing vae models. *arXiv preprint arXiv:1903.05789*, 2019.
- Dawid, A. P. The well-calibrated bayesian. *Journal of the American Statistical Association*, 77(379):605–610, 1982.
- DeGroot, M. H. and Fienberg, S. E. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 32(1-2):12–22, 1983.
- Denton, E. and Fergus, R. Stochastic video generation with a learned prior. 2018.
- Dhariwal, P., Jun, H., Payne, C., Kim, J. W., Radford, A., and Sutskever, I. Jukebox: A generative model for music. *arXiv preprint arXiv:[TODO]*, 2020.
- Edwards, H. and Storkey, A. Towards a neural statistician. *arXiv preprint arXiv:1606.02185*, 2016.
- Engel, J., Hoffman, M., and Roberts, A. Latent constraints: Learning to generate conditionally from unconditional generative models. *arXiv preprint arXiv:1711.05772*, 2017.
- Finn, C. and Levine, S. Deep visual foresight for planning robot motion. 2017.
- Ghosh, P., Sajjadi, M. S., Vergari, A., Black, M., and Schölkopf, B. From variational to deterministic autoencoders. *arXiv preprint arXiv:1903.12436*, 2019.
- Gregor, K., Danihelka, I., Graves, A., Rezende, D. J., and Wierstra, D. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015.
- Gregor, K., Besse, F., Rezende, D. J., Danihelka, I., and Wierstra, D. Towards conceptual compression. 2016.
- Gulrajani, I., Kumar, K., Ahmed, F., Taiga, A. A., Visin, F., Vazquez, D., and Courville, A. Pixelvae: A latent variable model for natural images. *arXiv preprint arXiv:1611.05013*, 2016.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. *arXiv preprint arXiv:1706.04599*, 2017.

- Hafner, D., Lillicrap, T., Ba, J., and Norouzi, M. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019a.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. 2019b.
- Henaff, M., Canziani, A., and LeCun, Y. Model-predictive policy learning with uncertainty regularization for driving in dense traffic. *arXiv preprint arXiv:1901.02705*, 2019.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pp. 6626–6637, 2017.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-VAE: Learning basic visual concepts with a constrained variational framework. 2017.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33, 2020.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- Kendall, A. and Gal, Y. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pp. 5574–5584, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. 2014.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pp. 4743–4751, 2016.
- Kohl, S., Romera-Paredes, B., Meyer, C., De Fauw, J., Ledsam, J. R., Maier-Hein, K., Eslami, S. A., Rezende, D. J., and Ronneberger, O. A probabilistic u-net for segmentation of ambiguous images. In *Advances in Neural Information Processing Systems*, pp. 6965–6975, 2018.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Lee, A. X., Zhang, R., Ebert, F., Abbeel, P., Finn, C., and Levine, S. Stochastic adversarial video prediction. *arXiv:1804.01523*, abs/1804.01523, 2018.
- Lee, A. X., Nagabandi, A., Abbeel, P., and Levine, S. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. *arXiv preprint arXiv:1907.00953*, 2019.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Lucas, J., Tucker, G., Grosse, R. B., and Norouzi, M. Don’t blame the elbow! a linear vae perspective on posterior collapse. In *Advances in Neural Information Processing Systems*, pp. 9403–9413, 2019.
- Maaløe, L., Fraccaro, M., Liévin, V., and Winther, O. Biva: A very deep hierarchy of latent variables for generative modeling. In *Advances in neural information processing systems*, pp. 6548–6558, 2019.
- Mattei, P.-A. and Frellsen, J. Leveraging the exact likelihood of deep latent variable models. In *Advances in Neural Information Processing Systems*, pp. 3855–3866, 2018.
- Neal, R. M. and Hinton, G. E. A view of the em algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*, pp. 355–368. Springer, 1998.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pp. 8024–8035, 2019.
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A. A., Tzionas, D., and Black, M. J. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 10975–10985, 2019.
- Peng, X. B., Kanazawa, A., Toyer, S., Abbeel, P., and Levine, S. Variational discriminator bottleneck: Improving imitation learning, inverse rl, and gans by constraining information flow. *arXiv preprint arXiv:1810.00821*, 2018.
- Peters, J. and Schaal, S. Reinforcement learning of motor skills with policy gradients. *Neural networks*, 21(4):682–697, 2008.

- Pong, V. H., Dalal, M., Lin, S., Nair, A., Bahl, S., and Levine, S. Skew-fit: State-covering self-supervised reinforcement learning. *arXiv preprint arXiv:1903.03698*, 2019.
- Rezende, D. J. and Viola, F. Taming vaes. *arXiv preprint arXiv:1810.00597*, 2018.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. 2014.
- Rolfe, J. T. Discrete variational autoencoders. *arXiv preprint arXiv:1609.02200*, 2016.
- Rosca, M., Lakshminarayanan, B., and Mohamed, S. Distribution matching in variational inference. *arXiv preprint arXiv:1802.06847*, 2018.
- Salimans, T., Karpathy, A., Chen, X., and Kingma, D. P. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. *arXiv preprint arXiv:1701.05517*, 2017.
- Shu, R., Bui, H. H., Zhao, S., Kochenderfer, M. J., and Ermon, S. Amortized inference regularization. *arXiv preprint arXiv:1805.08913*, 2018.
- Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. 2015.
- Sønderby, C. K., Raiko, T., Maaløe, L., Sønderby, S. K., and Winther, O. Ladder variational autoencoders. In *Advances in neural information processing systems*, pp. 3738–3746, 2016.
- Stirn, A. and Knowles, D. A. Variational variance: Simple and reliable predictive variance parameterization. *arXiv preprint arXiv:2006.04910*, 2020.
- Takahashi, H., Iwata, T., Yamanaka, Y., Yamada, M., and Yagi, S. Student-t variational autoencoder for robust density estimation. In *IJCAI*, pp. 2696–2702, 2018.
- Theis, L., Oord, A. v. d., and Bethge, M. A note on the evaluation of generative models. *ICLR*, 2016.
- Watter, M., Springenberg, J., Boedecker, J., and Riedmiller, M. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in neural information processing systems*, pp. 2746–2754, 2015.
- Zhao, S., Song, J., and Ermon, S. Infovae: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*, 2017.
- Zhu, J.-Y., Zhang, R., Pathak, D., Darrell, T., Efros, A. A., Wang, O., and Shechtman, E. Toward multimodal image-to-image translation. In *Advances in neural information processing systems*, pp. 465–476, 2017.