

Improving VAE generations of multimodal data through data-dependent conditional priors

Frantzeska Lavda^{*12} and Magda Gregorová^{*2} and Alexandros Kalousis²

Abstract. One of the major shortcomings of variational autoencoders is the inability to produce generations from the individual modalities of data originating from mixture distributions. This is primarily due to the use of a simple isotropic Gaussian as the prior for the latent code in the ancestral sampling procedure for the data generations. We propose a novel formulation of variational autoencoders, conditional prior VAE (CP-VAE), which learns to differentiate between the individual mixture components and therefore allows for generations from the distributional data clusters. We assume a two-level generative process with a continuous (Gaussian) latent variable sampled conditionally on a discrete (categorical) latent component. The new variational objective naturally couples the learning of the posterior and prior conditionals, and the learning of the latent categories encoding the multimodality of the original data in an unsupervised manner. The data-dependent conditional priors are then used to sample the continuous latent code when generating new samples from the individual mixture components corresponding to the multimodal structure of the original data. Our experimental results illustrate the generative performance of our new model comparing to multiple baselines.

1 Introduction

Variational autoencoders (VAEs) [20, 33] are deep generative models for learning complex data distributions. They consist of an encoding and a decoding network parametrizing the variational approximate posterior and the conditional data distributions in a latent variable generative model.

Though powerful and theoretically elegant, the VAEs in their basic form suffer from multiple deficiencies stemming from the mathematically convenient yet simplistic distributional assumptions. Multiple strategies have been proposed to increase the richness or the interpretability of the latent code [32, 3, 21, 4, 16, 30, 38, 1, 8, 6]. These mostly argue for more flexible posterior inference procedure or for the use of more complex approximate posterior distributions to facilitate the encoding of non-trivial data structures within the latent space.

In this paper we reason that for generating realistic samples of data originating from complex distributions, it is the prior which lacks expressiveness. Accordingly, we propose a new VAE formulation, conditional prior VAE (CP-VAE), with two-level hierarchical generative model combining a categorical and a continuous (Gaussian) latent variables.

The hierarchical conditioning of the continuous latent variable on the discrete latent component is particularly suitable for modelling multimodal data distributions such as distributional mixtures. Importantly, it also gives us better control of the procedure for generating new samples. Unlike in the standard VAE, we can sample data from specific mixture components at will. This is particularly critical if the generative power of VAEs shall be used in conjunction with methods requiring the identification of the distributional components such as in continual learning [31, 23].

As recently shown [6, 26], without supervision (as in our setting), enforcing independence factorization in the latent space does not guarantee recovering meaningful sources of variation in the original space. Therefore in our CP-VAE formulation, we let the model to fully utilize the capacity of the latent space by allowing for natural conditional decomposition in the generative and inference graphical models.

We formulate the corresponding variational lower bound on the data log-likelihood and use it as the optimization objective in the training. In the spirit of empirical Bayes, we propose to estimate the parameters of the conditional priors from the data together with the parameters of the variational posteriors in a joint learning procedure. This ensures that the inferred structure of the latent space can be exploited in the data generations.

2 Variational autoencoders

Variational autoencoders (VAEs) [20, 33] are deep Bayesian generative models relying on the principals of amortized variational inference to approximate the complex distributions $p(\mathbf{x})$ from which the observed data $\mathcal{S} = \{\mathbf{x}_i\}_{i=1}^n$ originate.

In their basic form, they model the unknown ground-truth $p(\mathbf{x})$ by a parametric distribution $p_\theta(\mathbf{x})$ with a latent variable generative process

$$p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z})d\mathbf{z} . \quad (1)$$

The typical assumptions are of a simple isotropic Gaussian prior $p(\mathbf{z})$ for the latent variable \mathbf{z} and, depending on the nature of the data \mathbf{x} , factorized Bernoulli or Gaussian distributions for the data conditionals $p_\theta(\mathbf{x}|\mathbf{z})$. These per-sample conditionals are parametrized by a deep neural network, a *decoder*. Once the decoder network is properly trained, we can sample new data examples from the learned data distribution $p_\theta(\mathbf{x})$ by ancestral sampling procedure: sample the latent \mathbf{z} from the prior $p(\mathbf{z})$ and pass it through the stochastic *decoder* $p_\theta(\mathbf{x}|\mathbf{z})$ to obtain the sample \mathbf{x} .

The VAEs employ the strategy of amortized variational inference. They approximate the intractable posteriors $p_\theta(\mathbf{z}|\mathbf{x})$ by factorized Gaussian distributions $q_\phi(\mathbf{z}|\mathbf{x})$ and infer the variational parameters

¹ University of Geneva, Switzerland, email: frantzeska.lavda@etu.unige.ch

² Geneva School of Business Administration (DMML Group), HES-SO, Switzerland, email: name.surname@hesge.ch,

* equal contribution

ϕ of the approximate per-sample posteriors through a deep neural network, an *encoder*.

The encoder and decoder networks are trained end-to-end by stochastic gradient-based optimization maximizing the sample estimate of lower bound $\mathcal{L}_{\theta,\phi} = \mathbb{E}_{p(\mathbf{x})} \mathcal{L}_{\theta,\phi}(\mathbf{x})$ on the data log-likelihood

$$\frac{1}{n} \sum_i \log p_{\theta}(\mathbf{x}_i) \geq \frac{1}{n} \sum_i \mathcal{L}_{\theta,\phi}(\mathbf{x}_i) \approx \mathcal{L}_{\theta,\phi}$$

$$\mathcal{L}_{\theta,\phi}(\mathbf{x}) = \underbrace{\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} \log p_{\theta}(\mathbf{x}|\mathbf{z})}_A - \underbrace{\text{KL}(q_{\phi}(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))}_B, \quad (2)$$

where $\text{KL}(\cdot||\cdot)$ is the Kullback-Leibler divergence. The first term A in equation (2) can be seen as a negative reconstruction cost, term B penalizes the deviations of the approximate posterior from the fixed prior $p(\mathbf{z})$ and has a regularizing effect on the model learning.

The gradients of the lower bound with respect to the model parameters θ can be obtained straightforwardly through Monte Carlo estimation. For the posterior parameters ϕ , the gradients are estimated by stochastic backpropagation using a location-scale transformation known as the reparametrization trick.

3 VAE with data-dependent conditional priors

The mathematically and practically convenient assumption of the factorial Gaussian approximate posterior $q_{\phi}(\mathbf{z}|\mathbf{x})$ has been previously contested as one of the major limitations of the basic VAE architecture. For complex data distributions $p(\mathbf{x})$ the simple Gaussian $q_{\phi}(\mathbf{z}|\mathbf{x})$ may not be flexible enough to approximate well the true posterior $p_{\theta}(\mathbf{z}|\mathbf{x})$. Scalable methods for enriching the posterior distributions through variable transformations have been proposed [32, 21, 36] alongside alternatives for improving the inference procedure [3, 38].

These methods improve rather dramatically the lower bound and log-likelihood maximization objectives by learning latent representations more appropriate for the complex data structures. However, these improvements cannot be translated into better generations without a closer match between the posterior and prior distributions used for sampling the latent variables during inference and data generations respectively.

We propose a new VAE formulation, conditional prior VAE (CP-VAE), with a conditionally structured latent representation which encourages a better match between prior and posterior distributions by jointly learning their parameters from the data.

3.1 Two-level generative process

We consider a two-level hierarchical generative process for the observed data

$$p_{\theta}(\mathbf{x}) = \sum_{\mathbf{c}} \int p_{\theta}(\mathbf{x}|\mathbf{z}, \mathbf{c}) p_{\varphi}(\mathbf{z}|\mathbf{c}) p(\mathbf{c}) d\mathbf{z}. \quad (3)$$

The latent space is composed of a continuous \mathbf{z} and a discrete \mathbf{c} component with a joint prior distribution

$$p(\mathbf{z}, \mathbf{c}) = p_{\varphi}(\mathbf{z}|\mathbf{c}) p(\mathbf{c}). \quad (4)$$

We assume a uniform categorical as a prior distribution for the discrete component \mathbf{c} so that for each of the K categories $p(c_k) =$

$1/K$, $k = 1, \dots, K$. The conditionals of the continuous component are factorised Gaussians³

$$p_{\varphi}(\mathbf{z}|\mathbf{c}_k) = \prod_i p_{\varphi}(z_i|\mathbf{c}_k) = \prod_i \mathcal{N}(z_i | \mu_{ik}, \sigma_{ik}^2), \quad k = 1, \dots, K. \quad (5)$$

The compositional prior we propose is well suited for generations of new samples from multimodal data distributions mixing multiple distributional components. In contrast to sampling from a simple isotropic Gaussian prior which concentrates symmetrically around the origin, we can sample the latent code from discontinuous parts of the latent space. These are expected to represent data clusters corresponding to the originating distributional mixing.

In addition, the variations encoded into the continuous part of the latent space are also sampled conditionally and therefore are better adapted to represent the important factors of data variations within the distributional clusters. This is in contrast to the single common continuous distribution of the basic VAE (section 2) or VAEs with similar continuous-discrete composition of the latents as ours which, however, assume independence between the two parts of the latent representation [10].

The data conditional $p_{\theta}(\mathbf{x}|\mathbf{z}, \mathbf{c})$ is parametrised by a *decoder* network $d_{\theta}(\mathbf{z}, \mathbf{c})$ as a Bernoulli($\mathbf{x} | d_{\theta}(\mathbf{z}, \mathbf{c})$) or a Gaussian $\mathcal{N}(\mathbf{x} | d_{\theta}(\mathbf{z}, \mathbf{c}), \sigma^2 I)$ distribution depending on the nature of the data \mathbf{x} .




Figure 1. To generate new examples from the learned data distribution $p_{\theta}(\mathbf{x})$, we sample the discrete and continuous latent variables from the two-level prior and pass those through the decoder.

3.1.1 Data-dependent conditional priors

There is no straightforward way to fix the parameters $\varphi = (\mu, \sigma)$ in the distributions (5) for each of the conditioning categories c_k a priori. Instead of placing hyper-priors on the parameters and expanding to full hierarchical Bayesian modelling, we estimate the prior parameters from the data through a relatively simple procedure resembling the empirical Bayes technique [29].

As explained in section 3.3, the conditional $p_{\varphi}(\mathbf{z}|\mathbf{c})$ enters our objective function through a KL divergence term. The prior parameters φ therefore can be optimized by backpropagation together with learning the encoder and decoder parameters ϕ and θ . Once the model is trained, all the parameters are fixed and the learned prior $p_{\varphi}(\mathbf{z}|\mathbf{c})$ can be used in the ancestral sampling procedure described above to generate new data samples.

3.2 Inference model

As in standard VAEs, we employ amortized variational inference to learn the unknown data distribution. We use the approximate posterior distribution

$$q_{\phi}(\mathbf{z}, \mathbf{c}|\mathbf{x}) = q_{\phi}(\mathbf{z}|\mathbf{x}, \mathbf{c}) q_{\phi}(\mathbf{c}|\mathbf{x}) \quad (6)$$

³ Full-covariance Gaussian could be parametrized through a lower-triangular matrix corresponding to the Cholesky decomposition at an increased computational cost [18, section 2.6.1].

in place of the intractable posterior $p_\theta(\mathbf{z}, \mathbf{c}|\mathbf{x})$. Our approximate posterior replicates the two-level hierarchical structure of the prior. In this way we ensure that the latent samples are structurally equivalent both during inference and new samples generations. This is not the case in other hierarchical latent models which rely on simplifying mean field assumptions for the posterior inference [16, 9].

We use *encoder* network with a gated layer $e_\phi(\mathbf{x}) = (\pi_\phi(\mathbf{x}), \mu_\phi(\mathbf{x}, \pi), \sigma_\phi(\mathbf{x}, \pi))$ for the amortized inference of the variational approximate posteriors

$$q_\phi(\mathbf{c}|\mathbf{x}) = \text{Cat}(\pi_\phi(\mathbf{x}))$$

$$q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c}) = \mathcal{N}(\mu_\phi(\mathbf{x}, \pi), \text{diag}(\sigma_\phi^2(\mathbf{x}, \pi))) .$$




Figure 2. The encoder infers the parameters of the discrete and continuous approximate posteriors using a gated layer for the hierarchical conditioning. At first it outputs the parameters of the discrete latent variable, π_ϕ . Then there is an extra layer that takes as input π_ϕ concatenated with the last hidden layer of the encoder and infers the parameters of the continuous latent variable.

3.3 Optimization objective

As customary in variational inference methods, our optimization objective is the maximization of the lower bound on the data log-likelihood

$$\mathcal{L}_{\theta, \phi}(\mathbf{x}) = \underbrace{\mathbb{E}_{q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x})} \log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})}_A - \underbrace{\text{KL}(q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x})||p_\varphi(\mathbf{z}, \mathbf{c}))}_B . \quad (7)$$

This is a straightforward adaptation of the bound from equation (2) to the compositional latent code (\mathbf{z}, \mathbf{c}) with similar interpretations for the A and B terms. Using the posterior and prior distribution decompositions from equations (4) and (6), the KL term in B can be rewritten as a sum of two KL divergences which are more amenable to practical implementation: $B1$ for the continuous conditional distributions and $B2$ for the discrete.

$$\begin{aligned} & \underbrace{\text{KL}(q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x})||p_\varphi(\mathbf{z}, \mathbf{c}))}_B \\ &= \underbrace{\text{KL}(q_\phi(\mathbf{c}|\mathbf{x})||p(\mathbf{c}))}_{B1} + \underbrace{\mathbb{E}_{q_\phi(\mathbf{c}|\mathbf{x})} \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c})||p_\varphi(\mathbf{z}|\mathbf{c}))}_{B2} \end{aligned} \quad (8)$$

The minimization of the KL divergence between the categorical posterior and the fixed uniform prior in the first term $B1$ is equivalent to maximizing the entropy of the categorical posterior $\mathbb{H}(q_\phi(\mathbf{c}|\mathbf{x}))$ (up to a constant).

$$\underbrace{\text{KL}(q_\phi(\mathbf{c}|\mathbf{x})||p(\mathbf{c}))}_{B1} = -\mathbb{H}(q_\phi(\mathbf{c}|\mathbf{x})) + \log K \quad (9)$$

The second term $B2$ can be seen as a weighted average of the KL divergences between the posterior and prior conditionals. The weights are the probabilities of the posterior categorical distribution so that

the two conditionals are pushed together more strongly for those observations \mathbf{x} and latent categories c_k to which the model assigns high probability. The KL s can be conveniently evaluated in a closed form as both the posterior and the prior conditionals are diagonal Gaussians.

We train the model by a stochastic gradient-based algorithm (Adam [19]). As the gradients of the variational lower bound $\mathcal{L}_{\theta, \phi}$ with respect to the model parameters are intractable, we use the usual well-established Monte-Carlo methods for their estimation.

For the decoder parameters θ , the gradient is estimated as the sample gradient of the conditional log-likelihood with the latent \mathbf{z} and \mathbf{c} sampled from the approximate posterior

$$\nabla_\theta \mathcal{L}_{\theta, \phi}(\mathbf{x}) \approx \nabla_\theta \log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c}), \quad (\mathbf{z}, \mathbf{c}) \sim q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x}) \quad (10)$$

For the encoder parameters ϕ , we use the pathwise gradient estimators [28] based on the standard location-scale $\mathbf{z} = f_\phi(\tilde{\mathbf{z}})$ and Gumbel-Softmax [15] $\mathbf{c} = g_\phi(\tilde{\mathbf{c}})$ reparametrizations with the auxiliary $\tilde{\mathbf{z}} \sim \mathcal{N}(0, 1)$ sampled from the standard normal and $\tilde{\mathbf{c}}$ sampled from the Gumbel(0,1) distribution.

$$\begin{aligned} \nabla_\phi \mathcal{L}_{\theta, \phi}(\mathbf{x}) &\approx \nabla_\phi \log p_\theta(\mathbf{x}, f_\phi(\tilde{\mathbf{z}}), g_\phi(\tilde{\mathbf{c}})) \\ &\quad - \nabla_\phi \log q_\phi(f_\phi(\tilde{\mathbf{z}}), g_\phi(\tilde{\mathbf{c}})|\mathbf{x}) \\ &\quad \tilde{\mathbf{z}} \sim \mathcal{N}(0, 1), \tilde{\mathbf{c}} \sim \text{Gumbel}(0, 1) \end{aligned} \quad (11)$$

Finally, the gradients with respect to the parameters φ of the conditional prior are estimated alongside the gradients of the decoder under the same sampling of the latents.

$$\nabla_\varphi \mathcal{L}_{\theta, \phi}(\mathbf{x}) \approx -\nabla_\varphi \log p_\varphi(\mathbf{z}|\mathbf{c}), \quad (\mathbf{z}, \mathbf{c}) \sim q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x}) \quad (12)$$

3.3.1 Analysis of the objective

The KL divergence in term B of the objective (7) has important regularization effects on the model learning. We expand on the discussion of these in the standard VAE objective (2) from [38] to analyse our more complex model formulation.

There are two major issues that optimizing the reconstruction term A of the objective (7) in isolation could cause. First, the model could completely ignore the categorical component of the latent representation \mathbf{c} by encoding all the data points \mathbf{x} into a single category with a probability $q_\phi(c_k|\mathbf{x}) = 1$. All the variation in the data \mathbf{x} would then be captured within the continuous component of the latent representation through the single continuous posterior $q_\phi(\mathbf{z}|\mathbf{x}, c_k)$. While this would not diminish the ability of the model to reconstruct the observed data and therefore would not decrease the reconstruction part of the objective A , it would degrade the generative properties of our model. Specifically, with all the data clusters pushed into a single categorical component and distributed within the continuous latent space, we would have no leverage for generating samples from the individual data distributional components which is one of the major requirements for our method. This pathological case is essentially equivalent to learning with the standard VAE.

Second, as discussed in [38], maximizing the log-likelihood in A naturally pushes the continuous posteriors to be concentrated around their means in disjoint parts of the continuous latent space with variances tending to zero. For such posteriors, the model could learn very specific decoding yielding very good reconstructions with very high log-likelihoods $p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})$. However, the generations would again suffer as the prior used for the ancestral sampling would not cover the same areas of the latent space as used during the inference.

To analyse the regularization effect of term B in the objective (7) on the learning, we decompose the expected KL divergence into three terms (and a constant):

$$\begin{aligned} \mathbb{E}_{p(\mathbf{x})} \text{KL}(q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x})||p_\varphi(\mathbf{z}, \mathbf{c})) &= \mathbb{I}_q((\mathbf{z}, \mathbf{c}), \mathbf{x}) \\ &+ \mathbb{E}_{q_\phi(\mathbf{c})} \text{KL}(q_\phi(\mathbf{z}|\mathbf{c})||p_\varphi(\mathbf{z}|\mathbf{c})) \\ &- \mathbb{H}(q_\phi(\mathbf{c})) + \log K. \end{aligned} \quad (13)$$

The first is the mutual information of the composite latent variable (\mathbf{z}, \mathbf{c}) and the data \mathbf{x} under the posterior distribution q . Minimization of the KL divergence in (7) pushes the mutual information between the two to be low and therefore prevents the overfitting of the latent representation to the training data described in the second point above.

The third term is the negative entropy of the marginal categorical posterior whose empirical evaluation over the data sample $\mathcal{S} = \{\mathbf{x}_i\}_{i=1}^n$ is often referred to as the *aggregated posterior* [27, 35].

$$q_\phi(\mathbf{c}) = \mathbb{E}_{p(\mathbf{x})} q_\phi(\mathbf{c}|\mathbf{x}) \approx \frac{1}{n} \sum_i^n q_\phi(\mathbf{c}|\mathbf{x}_i) \quad (14)$$

The regularizer maximizes the entropy of this distribution thus encouraging the model to use evenly all the categories of the discrete latent code counteracting the pathological case of the first point above.

Finally, the middle term pushes the marginalized conditional posteriors of the continuous latent variable \mathbf{z} to be close to the priors conditioned on the corresponding categories. It helps to distribute the variations in the data into the continuous component of the latent space in agreement between the inferential posteriors and the learned generative priors. It does so for each latent category c_k separately, putting more or less weights on the alignment as per the importance of the latent categories established through the categorical marginal posterior $q_\phi(\mathbf{c})$. It is this term in the objective of our VAE formulation which safeguards the generative properties of the model by matching the inferential posteriors and the learned generative priors used in the ancestral sampling procedure for new data examples.

4 Related work

Since their introduction in 2014 [20, 33] variational autoencoders have become one of the major workhorses for large-scale density estimation and unsupervised representation learning. Multitudes of variations on and enhancements of the original design have been proposed in the literature. These can broadly be categorized into 4 large groups (with significant overlaps as many methods mix multiple ideas to achieve the best possible performance).

First, it has been argued that optimizing the variational bound (2) instead of the intractable likelihood $p_\theta(\mathbf{x})$ inhibits the VAEs to learn useful latent representations both for data reconstructions and for downstream tasks. Methods using alternative objectives aim to encourage the learning towards representations that are better aligned with the data (measured by mutual information), e.g. InfoVAE [38], [1], or which separate important factors of variations in the data (disentangling), e.g. [14, 17, 10]. Though these methods report good results on occasions, there seem to be little evidence that breaking the variational bound brings systematical improvements [26, 6].

For our model, the analysis in section 3.3.1 suggests that our objective (which is a proper lower bound on the likelihood) encourages the encoding of the major source of variation, that of the originating mixture component, through the categorical variable without any extra alterations. At the same time, it should be noted that our goal is

not the interpretability of the learned representations nor their reuse outside the VAE model. Our focus is on generations reflecting the cluster structure of the original data space.

Second, the simplifying conditional independence assumptions for the data dimensions factored into the simple Gaussian decoder $p_\theta(\mathbf{x}|\mathbf{z})$ have been challenged in the context of modelling data with strong internal dependencies. More powerful decoders with autoregressive architectures have been proposed for modelling images, e.g. PixelVAE [13], or sequentially dependent data such as speech and sound, e.g. VRNN [5]. In our model, we use a hierarchical decoder $p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{c})$ corresponding to the cluster-like structure we assume for the data space. However, in this work we stick to the simple independence assumption for the data dimensions. Augmenting our method with stronger decoder should in principle be possible and is open for future investigation.

Third, the insufficient flexibility of the variational posterior $q_\phi(\mathbf{z}|\mathbf{x})$ to approximate the true posterior $p_\theta(\mathbf{z}|\mathbf{x})$ has led to proposals for more expressive posterior classes. For example, a rather successful approach is based on chaining invertible transformations of the latent variable [32, 21]. While increasing the flexibility of the approximate posterior improves the modelling objective through better reconstructions, without accompanied enhancements of the prior it does not guarantee better generations.

This has been recognised and addressed by the fourth group of improvements which focuses on the model prior and which our method pursues. For example, the authors in [37, 8] replace the distributional class of the prior (together with the posterior) by von Mises-Fisher distributions with potentially better characteristics for high-dimensional data with hyperspherical latent space.

More related to ours are methods that suggest to learn the prior. The VLVAE [4] uses the autoregressive flows in the prior which are equivalent to the inverse autoregressive flows in the posterior [21]. The increased richness of the encoding and prior distributions leads to higher quality generations, however, the prior cannot be used to generate from selected parts of the data space as our model can.

The VampPrior [35] proposes to construct the prior as a mixture of the variational posteriors over a learned set of pseudo-inputs. These could be interpreted as learned cluster prototypes of the data. However, the model does not learn the importance of the components in the mixture, and it does not align the prior and posteriors at an individual component level as our model does. Instead, it pushes the posteriors to align with the overall prior mixture which diminishes the models ability to correctly generate from the individual components of multimodal data.

The continuous-discrete decomposition of the latent space similar to ours have been used for data clustering through generative model in [16] and [9]. The first combines the VAE with a Gaussian mixture model through two stage procedure mimicking the independence assumptions in their inference model. The latter assumes (conditional) independence in the generative and inference models and extends to a full Bayesian formulation through the use of hyper-priors. Their complex model formulation exhibits some over-regularization issues that, as the authors acknowledge, are challenging to control.

Options for freeing the distributional class of the latent representations through Bayesian non-parametrics have been explored for example in [30, 12, 25]. The learned structures in the latent representations greatly increase the generative capabilities including also the (hierarchical) clustering ability. However, this comes at a cost of complex models that are tricky to train in a stable manner. In contrast, our model is elegantly simple and easy to train.

5 Empirical evaluation

We validate our new conditional prior (CP-VAE) model through experiments over synthetic data and two real-image datasets (MNIST [24] and Omniglot [22]). We compare the results with those produced by standard VAE (VAE), VAE with Mixture of Gaussian prior (MG) and VAE with VampPrior (VP) [35].

Not to obfuscate the analysis of the benefits of our method by various tweaks in the model architecture, we use the same structure of the encoder and decoder networks for all the methods in all our experiments. We set the dimensions of the continuous latent variable to 40, we use simple feed-forward networks with two hidden layers of 300 units each for both the encoder and the decoder, we initialise the weights according to Glorots’s method [11], and we utilize the gating mechanism of [7] as the element-wise non-linearity.

We trained all the models using ADAM optimizer [19] with learning rate 5×10^{-4} , mini-batches of size 100, and early stopping based on the stability of the objective over a validation-set. To avoid pathological local minima and numerical issues during training, we use a linear annealing/ warm-up scheme of 100 epochs [2], during which the KL regularization in the objective is annealed from 0 to 1 during training.

For generating new data examples, we use the ancestral sampling strategy with the latent variables sampled from the respective prior distributions of each method. In the simple VAE this is from the standard normal Gaussian $\mathbf{z} \sim p(\mathbf{z}) = \mathcal{N}(\mathbf{z}|0, I)$. In the MG model it is from the set of learned Gaussian components $\mathbf{z} \sim p_\varphi(\mathbf{z}) = \frac{1}{K} \sum_i^K \mathcal{N}(\mathbf{z}|\mu_k, \text{diag}(\sigma_k))$ with equal weighting. For VP it is from the mixture of variational posteriors $\mathbf{z} \sim p_\varphi(\mathbf{z}) = \frac{1}{K} \sum_i^K q_\phi(\mathbf{z}|\mathbf{u}_k)$ over the learned set of pseudo-inputs $\mathcal{U} = \{\mathbf{u}_k\}_{k=1}^K$ which first have to be passed through the encoder network. For our method, we follow the two level-generative process described in section 3.1 where we use the learned conditional priors for each of the categories for sampling the continuous latent component $\mathbf{z} \sim p_\varphi(\mathbf{z}|\mathbf{c} = c_k) = \mathcal{N}(\mathbf{z}|\mu_k, \text{diag}(\sigma_k))$ and the empirical aggregated posterior (14) for the discrete component $\mathbf{c} \sim q_\phi(\mathbf{c})$.

The implementation of our method together with the settings for replication of our experiments is available from our Bitbucket repository <https://bitbucket.org/dmmlgeneva/cp-vae/>.

5.1 Synthetic data experiments

In this section we demonstrate the effectiveness of the CP-VAE method through experiments over synthetic data. We use a toy dataset with 50000 examples $\mathbf{x} \in \mathbb{R}$ generated from a Gaussian mixture with two equally weighted components $\mathbf{x} \sim p(\mathbf{x}) = \frac{1}{2} (N(0.3, 0.05) + N(0.7, 0.05))$. We compare our method to the simple VAE (VAE), VAE with Mixture of Gaussians prior (MG) and VAE with VampPrior (VP). This simple set-up allows us to better understand the strengths and weaknesses of the method in terms of its density estimation performance for a known and rather simple ground-truth data distribution.

Because in real-life problems the number of distributional clusters in the data (the number of mixture components) may not be known or even easy to estimate, we use two experimental set-ups:

- **known** number of components: discrete latent variable \mathbf{c} with 2 categories (corresponding to the ground-truth two mixture components)
- **unknown** number of components: discrete latent variable \mathbf{c} with 150 categories

In Figure 3 we present histograms of data generated from the ground truth and the learned distributions. As we can see, our method (CP) correctly recovers the bi-modal structure of the data for both of the set-ups. This is important for practical utility of the method in situations where the domain knowledge does not provide an indication on the number of underlying generative clusters. With high enough number of categories within the discrete latent, our method can recover the correct multi-modal structure of the data.

Due to the simplicity of this set-up even methods which do not adjust their priors to the disjoint learned representation, such as the simple VAE, are able to recover the multimodal structure of the data at generation time. However, in contrast to CP-VAE, because of the nature of the model, these are not able to conditionally generate samples. MG and VP have difficulties to recover the structure of the data when a small number of components/ pseudo-inputs is used. This seems to be improved as the number of components/ pseudo-inputs is increased to 150 but at the expense of computational cost.




Figure 3. Histograms of the data generated from the ground-truth $\mathbf{x} \sim p(\mathbf{x}) = \frac{1}{2} (N(0.3, 0.05) + N(0.7, 0.05))$ and the learned distributions. Our CP method can recover the bi-modal structure of the data correctly irrespective of the number of categories used for the latent categorical component.

We further explore how our model handles the excess capacity within the categorical latent variable. For this we focus on the 150-category case and we generate data by sampling the discrete latent variable a) from the marginal posterior $\mathbf{c} \sim q_\phi(\mathbf{c})$, b) from the uniform prior $\mathbf{c} \sim p(\mathbf{c}) = \frac{1}{K}$.

Comparing the two in Figure 4, we see that unlike the generations sampled from the marginal posterior, the generations from the uniform prior display some mixing artifacts. This suggests that our model learns to ignore the excess capacity by assigning low marginal probabilities $q_\phi(c_k) \approx 0$ to some of the categories. The continuous latent representations corresponding to these parts of the disjoint latent space are irrelevant for both the reconstructions and the generations due to our weighted KL formulation in B^2 of equation (8).

5.2 Real data experiments

For the real-data experiments we use two image datasets, MNIST [24] and Omniglot [22], commonly used for evaluation of generative models. We use the dynamically binarized versions of the datasets, as in [34], with the following train-validation-test splits: for MNIST 50000-10000-10000, for Omniglot 23000-1345-8070.

We first demonstrate the benefits of learning the prior distribution p_φ on the generative performance of our CP-VAE method. For this




Figure 4. The CP-VAE learns to ignore the excess capacity of the disjoint latent space by assigning near-zero probability to some of the categories in the discrete latent space. These parts of the latent space are ignored for the reconstructions and by sampling the categorical variable from the marginal posterior $q_\phi(c)$ can correctly be ignored also for the generations.

we fix the number of categories within the discrete latent variable to 1. This is equivalent to learning the simple VAE with only the continuous latent representation but a learnable prior and the corresponding objective

$$\mathcal{L}_{\theta, \phi}(\mathbf{x}) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} \log p_{\theta}(\mathbf{x}|\mathbf{z}) - \text{KL}(q_{\phi}(\mathbf{z}|\mathbf{x})||p_{\varphi}(\mathbf{z})) \quad (15)$$

In Figures 5 and 6 we compare the reconstructions produced by our CP-VAE model with only a single category in the latent space to the standard VAE without a learnable prior. The visual quality of the data reconstructions for both of the models is very high. However, the quality of generations (Figure 6) seems to be improved when the model adapts the prior to the learned latent space (CP1).




Figure 5. Data reconstructions from models without a categorical structure in the latent space are similarly good with $p_{\varphi}(\mathbf{z})$ and without $p(\mathbf{z})$ using the learnable prior.




Figure 6. New data generations from models without a categorical structure in the latent space for the CP1 adapting the prior $p_{\varphi}(\mathbf{z})$ to the learned latent space (a) and for the standard VAE with fixed prior $p(\mathbf{z})$ (b).

Next, we examine the ability of our model to generate new examples from the underlying distributional clusters. For this we trained our model over the MNIST data with 10 categories in the discrete latent variable and we compared it to the VAE with VampPrior (VP) with 10 pseudo-inputs. In Figure 7 we show that our model learned to use the categorical variable for encoding the changes in the digits. For our model, the examples in each of the subplots were generated by fixing the discrete latent variable to one category and sampling the continuous latent from the corresponding learned prior. For VP the examples in each of the subplots were generated by using the same pseudo-input. We see that the learned discrete encoding accurately

captures the main categories of the data, while VP conditioned on the individual learned pseudo-inputs generates a mix of digits.

In fact, rather remarkably, though our model was trained in a purely unsupervised manner, if we use the learned discrete code as a classification output we achieve 94.6% classification accuracy. To compare the cluster assignments given by CP-VAE with the true image labels we follow the following procedure: We first manually match each category in the discrete latent variable to the corresponding label based on the conditional generations. Then we encode the data and we compare the assigned labels with the true labels to obtain the classification accuracy.




Figure 7. New variations of individual MNIST digits generated by our CP-VAE model (1st – 2nd row) with a latent discrete variable with 10 categories and by VP model (3rd – 4th row) with using 10 pseudo-inputs. Digits in the same subplot were generated from the same discrete category/ pseudo-input.

Finally, we compare the performance of our CP-VAE to a number of baselines varying also the size of the categorical variable. We experiment with $\{10, 150, 500\}$ categories for MNIST and $\{50, 500\}$ categories for Omniglot. The methods we compare to are the simple VAE (VAE) and the following methods from [35]: VAE with Mixture of Gaussians prior (MG), VAE with VampPrior (VP), hierarchical two-layered VAE with VampPrior (HVP), and hierarchical two-layered VAE with simple fixed prior (HVAE). For the VP and MG methods we use the same numbers of pseudo-inputs and mixture components as the number of the latent categories and for the two layers models we use 40 latent variables at each layer.

	MNIST			<i>Omniglot</i>	
	$c = 10$	$c = 150$	$c = 500$	$c = 50$	$c = 500$
CP	84.76	86.84	88.63	114.31	114.14
VAE	88.75	—	—	115.06	—
MG	89.43	88.96	88.85	116.31	116.12
VP	87.94	86.55	86.07	114.01	113.74
HVAE	86.7	—	—	110.81	—
HVP	85.90	85.09	85.01	110.50	110.16

Table 1. Comparison of negative variational lower bounds for the different methods over the test data sets.

We summarize the numerical results in terms of the negative variational lower bound calculated over the test data in Table 1. Our model achieves lower bounds when compared to VAE and MG in all the cases. It also performs the best of all the tested methods for the MNIST dataset with 10 latent categories. Otherwise, the VampPrior and especially the hierarchical VampPrior method seems to perform consistently the best. However, as explained in section 4, good values of the variational lower bound objective do not guarantee good generations and good control over the distributional clusters.

We present the new data examples generated by the various methods in Figure 8 and 9 for the MNIST and Omniglot data respectively. Our model is able to generate consistently good quality new samples for both the datasets irrespective of the number of latent categorical components. In contrast, all of the other methods (all VampPrior variations including the two-layer hierarchical, and the MG) fail to generate quality examples with only 10 components within the prior. They also seem to collapse to generating examples only from a few digits which suggest an important lack of flexibility available for the generations. As the number of components (pseudo-inputs) in the prior mixture increases, the VampPrior generations tend to improve, with the hierarchical version of the method systematically outperforming the simple VP version.




Figure 8. New data examples from the MNIST dataset generated by the various methods with increasing number of the prior components.

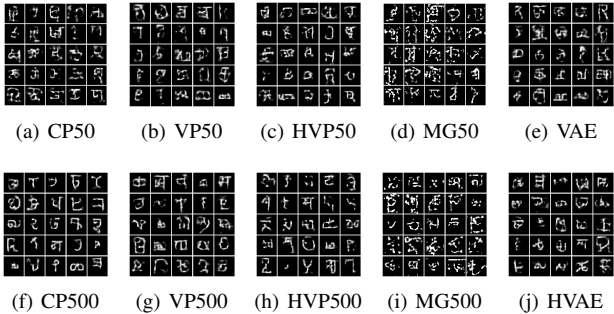


Figure 9. New data examples from the Omniglot dataset generated by the various methods with increasing number of the prior components.

6 Conclusion

In this paper, we have introduced CP-VAE, an unsupervised generative model in order to learn the multi-modal structure of the data. We propose a conditionally structured latent representation that enables our model to discover the modes of the training data. This is achieved through a better matching between prior and posterior distributions by jointly learning their parameters from the data. The experimental results demonstrate that our approach is able to recover the modes of the original data in an unsupervised manner, to conditionally generate new samples from the underlying mode distribution and has a high accuracy of cluster assignments. Moreover, CP-VAE is able to generate new samples irrespective of the number of latent categorical components. In contrast, all of the VP variations and the MG fail to generate good-quality examples with only 10 components within the prior.

Acknowledgements

This work has been in part supported by the Swiss National Foundation grant 177179: Modelling pathological gait resulting from motor impairments.

APPENDIX

Proof of equation (8)

$$\begin{aligned} \underbrace{\text{KL}(q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x})||p_\varphi(\mathbf{z}, \mathbf{c}))}_B &= \\ &= \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c})q_\phi(\mathbf{c}|\mathbf{x})} \log \frac{q_\phi(\mathbf{c}|\mathbf{x})}{p(\mathbf{c})} \frac{q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c})}{p_\varphi(\mathbf{z}|\mathbf{c})} = \\ &= \underbrace{\text{KL}(q_\phi(\mathbf{c}|\mathbf{x})||p(\mathbf{c}))}_{B1} + \underbrace{\mathbb{E}_{q_\phi(\mathbf{c}|\mathbf{x})} \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c})||p_\varphi(\mathbf{z}|\mathbf{c}))}_{B2} \end{aligned}$$

Proof of equation (9)

$$\begin{aligned} \underbrace{\text{KL}(q_\phi(\mathbf{c}|\mathbf{x})||p(\mathbf{c}))}_{B1} &= \mathbb{E}_{q_\phi(\mathbf{c}|\mathbf{x})} [\log q_\phi(\mathbf{c}|\mathbf{x}) - \log K^{-1}] \\ &= -\mathbb{H}(q_\phi(\mathbf{c}|\mathbf{x})) + \log K \end{aligned}$$

Proof of equation (13)

$$\begin{aligned} \mathbb{E}_{p(\mathbf{x})} \text{KL}(q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x})||p_\varphi(\mathbf{z}, \mathbf{c})) &= \\ &= \mathbb{E}_{p(\mathbf{x})} \mathbb{E}_{q_\phi(\mathbf{c}|\mathbf{x})} \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{c})||p_\varphi(\mathbf{z}|\mathbf{c})) = \\ &= \mathbb{E}_{q_\phi(\mathbf{x}, \mathbf{c})} \left[\log \frac{q_\phi(\mathbf{c}|\mathbf{x})}{q_\phi(\mathbf{c})} + \log \frac{q_\phi(\mathbf{c})}{p(\mathbf{c})} \right] \\ &+ \mathbb{E}_{q_\phi(\mathbf{x}, \mathbf{z}, \mathbf{c})} \left[\log \frac{q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x})}{q_\phi(\mathbf{c}|\mathbf{x})q_\phi(\mathbf{z}|\mathbf{c})} + \log \frac{q_\phi(\mathbf{z}|\mathbf{c})}{p_\varphi(\mathbf{z}|\mathbf{c})} \right] \\ &= \text{KL}(q_\phi(\mathbf{c})||p(\mathbf{c})) + \mathbb{E}_{q_\phi(\mathbf{c})} \text{KL}(q_\phi(\mathbf{z}|\mathbf{c})||p_\varphi(\mathbf{z}|\mathbf{c})) \\ &+ \mathbb{E}_{q_\phi(\mathbf{x}, \mathbf{z}, \mathbf{c})} \log \frac{q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x})}{q_\phi(\mathbf{z}|\mathbf{c})q_\phi(\mathbf{c})} \\ &= \log K - \mathbb{H}(q_\phi(\mathbf{c})) + \mathbb{E}_{q_\phi(\mathbf{c})} \text{KL}(q_\phi(\mathbf{z}|\mathbf{c})||p_\varphi(\mathbf{z}|\mathbf{c})) \\ &+ \mathbb{E}_{q_\phi(\mathbf{x}, \mathbf{z}, \mathbf{c})} \log \frac{q_\phi(\mathbf{z}, \mathbf{c}|\mathbf{x})}{q_\phi(\mathbf{z}, \mathbf{c})} \\ &= \log K - \mathbb{H}(q_\phi(\mathbf{c})) \\ &+ \mathbb{E}_{q_\phi(\mathbf{c})} \text{KL}(q_\phi(\mathbf{z}|\mathbf{c})||p_\varphi(\mathbf{z}|\mathbf{c})) + \mathbb{I}_{q_\phi}((\mathbf{z}, \mathbf{c}), \mathbf{x}) \end{aligned}$$

REFERENCES

- [1] Alexander A. Alemi, Ben Poole, Ian Fischer, Joshua V. Dillon, Rif A. Saurous, and Kevin Murphy, ‘Fixing a Broken ELBO’, in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, (2018).
- [2] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio, ‘Generating sentences from a continuous space’, in *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pp. 10–21, Berlin, Germany, (August 2016). Association for Computational Linguistics.
- [3] Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov, ‘Importance Weighted Autoencoders’, in *Proceedings of the 4th International Conference on Learning Representations (ICLR)*, (2016).
- [4] Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel, ‘Variational Lossy Autoencoders’, in *International Conference on Learning Representations*, p. 17, (2017).
- [5] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio, ‘A Recurrent Latent Variable Model for Sequential Data’, in *Advances in Neural Information Processing Systems 28*, eds., C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, 2980–2988, Curran Associates, Inc., (2015).
- [6] Bin Dai and David Wipf, ‘Diagnosing and Enhancing VAE Models’, in *International Conference on Learning Representations*, (2019).
- [7] Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier, ‘Language modeling with gated convolutional networks’, in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 933–941. JMLR. org, (2017).
- [8] Tim R. Davidson, Luca Falorsi, Nicola De Cao, Thomas Kipf, and Jakub M. Tomczak, ‘Hyperspherical Variational Auto-Encoders’, *arXiv:1804.00891 [cs, stat]*, (April 2018).
- [9] Nat Dilokthanakul, Pedro A. M. Mediano, Marta Garnelo, Matthew C. H. Lee, Hugh Salimbeni, Kai Arulkumaran, and Murray Shanahan, ‘Deep Unsupervised Clustering with Gaussian Mixture Variational Autoencoders’, *arXiv:1611.02648 [cs, stat]*, (January 2017).
- [10] Emilien Dupont, ‘Learning Disentangled Joint Continuous and Discrete Representations’, (March 2018).
- [11] Xavier Glorot and Yoshua Bengio, ‘Understanding the difficulty of training deep feedforward neural networks’, in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256, (2010).
- [12] Prasoon Goyal, Zhiting Hu, Xiaodan Liang, Chenyu Wang, and Eric Xing, ‘Nonparametric Variational Auto-encoders for Hierarchical Representation Learning’, *Proceedings of the IEEE International Conference on Computer Vision*, (August 2017).
- [13] Ishaan Gulrajani, Kundan Kumar, Faruk Ahmed, Adrien Ali Taiga, Francesco Visin, David Vazquez, and Aaron Courville, ‘PixelVAE: A Latent Variable Model for Natural Images’, *International Conference on Learning Representations (ICLR)*, (November 2017).
- [14] 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’, *International Conference on Learning Representations (ICLR)*, (November 2017).
- [15] Eric Jang, Shixiang Gu, and Ben Poole, ‘Categorical Reparameterization with Gumbel-Softmax’, in *International Conference on Learning Representations (ICLR)*, (2017).
- [16] Zhuxi Jiang, Yin Zheng, Huachun Tan, Bangsheng Tang, and Hanning Zhou, ‘Variational Deep Embedding: An Unsupervised and Generative Approach to Clustering’, in *IJCAI*, (2017).
- [17] Hyunjik Kim and Andriy Mnih, ‘Disentangling by Factorising’, (July 2018).
- [18] D. P. Kingma, *Variational Inference & Deep Learning: A New Synthesis*, Ph.D. dissertation, 2017.
- [19] Diederik P. Kingma and Jimmy Ba, ‘Adam: A Method for Stochastic Optimization’, *3rd International Conference for Learning Representations*, (January 2015).
- [20] Diederik P. Kingma and Max Welling, ‘Auto-Encoding Variational Bayes’, in *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, (2014).
- [21] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling, ‘Improved Variational Inference with Inverse Autoregressive Flow’, in *Advances in Neural Information Processing Systems 29*, eds., D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, (2016).
- [22] Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum, ‘Human-level concept learning through probabilistic program induction’, *Science*, **350**(6266), 1332–1338, (2015).
- [23] Frantzeska Lavda, Jason Ramapuram, Magda Gregorova, and Alexandros Kalousis, ‘Continual Classification Learning Using Generative Models’, *Continual learning workshop, Advances in Neural Information Processing Systems*, (December 2018).
- [24] Yann Lecun, Lon Bottou, Yoshua Bengio, and Patrick Haffner, ‘Gradient-based learning applied to document recognition’, in *Proceedings of the IEEE*, pp. 2278–2324, (1998).
- [25] Xiaopeng Li, Zhouong Chen, Leonard K. M. Poon, and Nevin L. Zhang, ‘Learning Latent Superstructures in Variational Autoencoders for Deep Multidimensional Clustering’, (September 2019).
- [26] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem, ‘Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations’, *Proceedings of the 36th International Conference on Machine Learning*, (June 2019).
- [27] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey, ‘Adversarial Autoencoders’, *arXiv:1511.05644 [cs]*, (May 2016).
- [28] Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih, ‘Monte Carlo Gradient Estimation in Machine Learning’, *arXiv:1906.10652 [cs, math, stat]*, (June 2019).
- [29] Kevin P. Murphy, *Machine Learning: A Probabilistic Perspective*, Adaptive Computation and Machine Learning Series, MIT Press, Cambridge, MA, 2012.
- [30] Eric Nalisnick and Padhraic Smyth, ‘Stick-Breaking Variational Autoencoders’, in *International Conference on Learning Representations*, (2017).
- [31] Jason Ramapuram, Magda Gregorova, and Alexandros Kalousis, ‘Life-long Generative Modeling’, *arXiv:1705.09847 [cs, stat]*, (November 2019).
- [32] Danilo Jimenez Rezende and Shakir Mohamed, ‘Variational Inference with Normalizing Flows’, in *Proceedings of the 32 International Conference on Machine Learning*, p. 9, (2015).
- [33] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra, ‘Stochastic Backpropagation and Approximate Inference in Deep Generative Models’, in *International Conference on Machine Learning*, pp. 1278–1286, (January 2014).
- [34] Ruslan Salakhutdinov and Iain Murray, ‘On the quantitative analysis of deep belief networks’, in *Proceedings of the 25th international conference on Machine learning*, pp. 872–879, (2008).
- [35] Jakub M. Tomczak and Max Welling, ‘VAE with a VampPrior’, *21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, (May 2018).
- [36] Rianne van den Berg, Leonard Hasenclever, Jakub M. Tomczak, and Max Welling, ‘Sylvester Normalizing Flows for Variational Inference’, *CoRR, abs/1803.05649*, (February 2019).
- [37] Jiacheng Xu and Greg Durrett, ‘Spherical Latent Spaces for Stable Variational Autoencoders’, in *Conference on Empirical Methods in Natural Language Processing*, (August 2018).
- [38] Shengjia Zhao, Jiaming Song, and Stefano Ermon, ‘InfoVAE: Information Maximizing Variational Autoencoders’, *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI)*, (June 2017).