

Hyperprior Induced Unsupervised Disentanglement of Latent Representations

Abdul Fatir Ansari, Harold Soh

Department of Computer Science, National University of Singapore
 {afatir, harold}@comp.nus.edu.sg

Abstract

We address the problem of unsupervised disentanglement of latent representations learnt via deep generative models. In contrast to current approaches that operate on the evidence lower bound (ELBO), we argue that statistical independence in the latent space of VAEs can be enforced in a principled hierarchical Bayesian manner. To this effect, we augment the standard VAE with an inverse-Wishart (IW) prior on the covariance matrix of the latent code. By tuning the IW parameters, we are able to encourage (or discourage) independence in the learnt latent dimensions. Extensive experimental results on a range of datasets (2DShapes, 3DChairs, 3DFaces and CelebA) show our approach to outperform the β -VAE and is competitive with the state-of-the-art FactorVAE. Our approach achieves significantly better disentanglement and reconstruction on a new dataset (CorrelatedEllipses) which introduces correlations between the factors of variation.

1 Introduction

Learning semantically interpretable representations of data remains an important open problem in artificial intelligence. In particular, there has been considerable attention on learning *disentangled* representations—equivariant codes that exhibit predictable changes in a *single* associated dimension when a factor of variation is altered (Bengio, Courville, and Vincent 2013). Disentangled representations are beneficial for a variety of tasks including exploratory data analysis (EDA), transfer learning, and generative modeling. For example, one may seek to change a single aspect of a generated face (e.g., lighting, orientation, or hair color). With an appropriate disentangled representation, only one dimension of the latent code needs to be modified to obtain the required change. By analogy to inverse graphics, the disentangled representation can be regarded as independent parameters fed to a rendering engine to synthesize an image.

In this work, we focus on pure *unsupervised* learning of disentangled representations with deep generative models. Alternative (semi-) supervised approaches have been explored in recent work (Reed et al. 2014; Kulkarni et al. 2015; Siddharth et al. 2017; Mathieu et al. 2016), but these methods require labels that can be costly to obtain. Moreover, in

certain applications such as EDA, the factors of variation are unknown and precisely the information we seek to uncover.

In the unsupervised setting, a prior notion of disentanglement is required: we adopt a current standard assumption that the data is generated from a fixed number of *statistically independent* factors. A popular generative model under this assumption is the β -VAE (Higgins et al. 2017), which learns good disentangled representations whilst being easy to train. However, the mechanism employed to encourage disentanglement—by increasing the weight on the KL divergence between the variational posterior and prior—sacrifices reconstruction fidelity.

Recognizing this deficiency, very recent models—the FactorVAE (Kim and Mnih 2018) and β -TCVAE (Chen et al. 2018)—have improved upon the β -VAE by augmenting the VAE loss with an extra penalty term that encourages independence in the latent codes. Although this penalty term is well-motivated via *total correlation*, the drive to maximize statistical independence in this manner may not be robust to factor correlations in the data that exist due to biased sampling. Furthermore, the need to add additional weighted terms to the variational lower bound is unsatisfying from a probabilistic modeling perspective; it points to an inadequacy in the underlying model formulation.

This paper takes a step back and asks whether trading-off reconstruction and disentanglement can be handled in a more principled fashion. Rather than introducing additional terms or weights to the VAE evidence lower bound (ELBO), we focus instead on the generative model and specifically, its *latent representation prior* $p(\mathbf{z})$. In the standard VAE, $p(\mathbf{z})$ is a standard multivariate Gaussian $p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \Sigma)$ where $\Sigma = \mathbf{I}$. By introducing a suitable *hyperprior* on the covariance matrix—e.g., the inverse-Wishart (IW) used in this study—one imagines that we can encourage (or discourage) independence in the learnt latent dimensions via the hyperprior’s parameters. When trained via variational inference using an approximate distribution $q(\mathbf{z}, \Sigma)$, the hyperprior’s effect naturally manifests as additional terms in the ELBO, rather than having to be inserted post-hoc as in previous studies. This approach is very natural from a Bayesian perspective, but surprisingly, has yet to be explored in the literature.

Unlike previous work, our model formulation entails learning a full covariance matrix Σ ; this allows the model

to capture possible correlations in the dataset, but requires additional treatment to ensure stable training. We employ a *structured variational posterior* and present approximation techniques to enable efficient and stable inference. We term the resulting model and inference scheme as the Covariance Hyperprior VAE (CHyVAE).

Experiments on a range of image datasets—2DShapes, CelebA, 3DFaces, and 3DChairs—show that CHyVAE outperforms β -VAE both in terms of disentanglement and reconstruction error and is competitive with the state-of-the-art FactorVAE. We also compare the three models on a novel dataset (CorrelatedEllipses) which introduces strong correlations between the factors of variation. Here, CHyVAE outperforms both β -VAE and FactorVAE by a significant margin. These results indicate that disentanglement and reconstruction can be traded-off in an alternative manner, i.e., at the model specification level, compared to existing approaches that operate on the ELBO.

In summary, this paper makes the following key contributions:

- A hierarchical Bayesian approach for learning disentangled latent space representations in an unsupervised manner;
- A specific generative model with an inverse-Wishart hyperprior and an efficient inference scheme with a structured variational posterior, which results in the CHyVAE;
- Extensive empirical results and analyses comparing CHyVAE to β -VAE and the FactorVAE on a range of datasets, which show that disentanglement can be achieved without resorting to “ELBO surgery” (Hoffman and Johnson 2016; Kim and Mnih 2018; Chen et al. 2018).

2 Generation and Disentanglement with the Variational Autoencoder

To begin, we give a brief overview of the Variational Autoencoder (VAE) (Kingma and Welling 2013)¹ and the alterations used to encourage disentanglement. We first consider the standard generative scheme, where our objective is to find parameters θ that maximize the expected log probability of the dataset under the data distribution,

$$\operatorname{argmax}_{\theta} \mathbb{E}_{\text{data}}[\log p_{\theta}(\mathbf{x})] \quad (1)$$

and $\mathbf{x} \in \mathcal{X}$ is an observed data item of interest. For real world data, $p_{\theta}(\mathbf{x})$ may be highly complex and non-trivial to generate samples from. Furthermore, in unsupervised learning, we may wish to obtain representations of \mathbf{x} that are more amenable to downstream analysis. One approach for achieving these aims is to further specify the log-distribution within Eq. (1),

$$\log p_{\theta}(\mathbf{x}) = \log \mathbb{E}_{p(\mathbf{z})}[p_{\theta}(\mathbf{x}|\mathbf{z})] \quad (2)$$

where we introduce the conditional distribution $p_{\theta}(\mathbf{x}|\mathbf{z})$ and variables $\mathbf{z} \in \mathcal{Z}$ with prior $p(\mathbf{z})$. Intuitively, each \mathbf{z} is a latent representation or code associated with \mathbf{x} . By choosing

¹We refer readers wanting more detail on the VAE and an alternative derivation to related work (Doersch 2016; Kingma and Welling 2013).

an appropriate condition and setting the prior $p(\mathbf{z})$ to be a simple distribution, e.g., $\mathcal{N}(\mathbf{0}, \mathbf{I})$, we can easily generate \mathbf{x} by sampling from $p(\mathbf{z})$. In addition, the diagonal covariance \mathbf{I} indicates a prior expectation that underlying data representation comprises statistically independent Gaussians (one for each latent dimension), and is therefore disentangled.

Computing $\log p_{\theta}(\mathbf{x})$ requires marginalizing out the latent variables \mathbf{z} which is generally intractable, e.g., when $p_{\theta}(\mathbf{x}|\mathbf{z}) = p(\mathbf{x}|f(\mathbf{z}))$ and f is a nonlinear neural network. To perform approximate inference, the VAE employs a *recognition* or *inference model*, $q_{\phi}(\mathbf{z}|\mathbf{x})$, and maximizes the variational or evidence lower bound (ELBO),

$$\mathcal{L}_{\text{ELBO}}^{\text{VAE}} = \frac{1}{N} \sum_{n=1}^N \mathbb{E}_{q_{\phi}}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] - D_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))$$

The above can be seen as the expectation of the data likelihood under the inference model with a KL divergence term that measures how different $q_{\phi}(\mathbf{z}|\mathbf{x})$ is from the prior $p(\mathbf{z})$.

2.1 Encouraging Disentanglement in VAEs

In our context, the second term in the ELBO is of particular interest: when $p(\mathbf{z})$ is intentionally chosen to factorize across the dimensions, minimizing the KL divergence encourages independence in the learned latent representations. The β -VAE (Higgins et al. 2017) takes advantage of this observation and encourages disentanglement by emphasizing the KL divergence with a weight β :

$$\mathcal{L}^{\beta} = \frac{1}{N} \sum_{n=1}^N \mathbb{E}_q[\log p(\mathbf{x}|\mathbf{z})] - \beta D_{\text{KL}}[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]$$

With larger β and a factorized prior, maximizing this modified objective favors latent representations possessing greater independence across the dimensions.

However, disentanglement gains in the β -VAE are often off-set by a decrease in reconstruction performance. Recent work (Kim and Mnih 2018; Chen et al. 2018) has argued that increasing β has the undesirable side-effect of inadvertently penalizing the mutual information between \mathbf{x} and \mathbf{z} . This can be seen by decomposing the KL divergence term:

$$\mathbb{E}_{\text{data}}[D_{\text{KL}}[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]] = I(\mathbf{x}; \mathbf{z}) + D_{\text{KL}}[q(\mathbf{z})||p(\mathbf{z})]$$

where $I(\mathbf{x}; \mathbf{z})$ is the mutual information between \mathbf{x} and \mathbf{z} (Hoffman and Johnson 2016). Penalizing the first term decreases the informativeness of \mathbf{z} about \mathbf{x} and hence, reduces reconstruction quality. To overcome this problem, both Kim and Mnih and Chen et al. optimize an augmented objective:

$$\mathcal{L}^{\gamma} = \mathcal{L}_{\text{ELBO}}^{\text{VAE}} - \gamma D_{\text{KL}} \left[q(\mathbf{z}) \parallel \prod_i q(z_i) \right] \quad (3)$$

where $D_{\text{KL}}[q(\mathbf{z})||\prod_i q(z_i)]$ is the *total correlation* (TC). Computing the TC is generally intractable and has to be approximated; the FactorVAE (Kim and Mnih 2018) estimates TC by means of a discriminator using the density-ratio trick, where else β -TCVAE (Chen et al. 2018) uses a minibatch-based alternative. Empirical results show that optimizing \mathcal{L}^{γ} leads to better reconstruction with similar disentanglement compared to the β -VAE.

3 A Hyperprior Approach for Learning Disentangled Representations

We observed in the previous section that current state-of-the-art methods attempt to procure disentanglement by augmenting the ELBO. In this section, we describe an alternative approach by further expanding upon the log-distribution within Eq. (1). At a high-level, we desire a means to “regularize” the latent codes towards a disentangled form, yet preserve sufficient *flexibility* to achieve good reconstruction. A natural Bayesian approach to achieve these aims is to place a hyperprior $p(\Sigma)$ on the covariance parameter of $p(\mathbf{z}|\Sigma)$:

$$\log p_\theta(\mathbf{x}) = \log \mathbb{E}_{p(\mathbf{z}|\Sigma)p(\Sigma)}[p(\mathbf{x}|\mathbf{z})]. \quad (4)$$

Our proposed model denotes a modified generative process relative to the standard VAE. In particular, an observed sample $\mathbf{x} \in \mathbb{R}^D$ is generated by first sampling a covariance matrix $\Sigma \sim p(\Sigma)$, followed by the latent representation $\mathbf{z} \sim p(\mathbf{z}|\Sigma)$ and finally, $\mathbf{x} \sim p_\theta(\mathbf{x}|\mathbf{z})$. As such, the joint probability factorizes as

$$p(\mathbf{x}, \mathbf{z}, \Sigma) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z}|\Sigma)p(\Sigma) \quad (5)$$

since \mathbf{x} and Σ are independent conditioned on \mathbf{z} . Notice that Σ is no longer constrained to be simply \mathbf{I} , but disentanglement can be encouraged in a straight-forward manner simply by placing greater weight on independence between the latent dimensions. In other words, tuning the strength or *informativeness* of the hyperprior would then allow us to naturally vary of level of disentanglement desired. By allowing some deviation from strict independence, the model recognizes that individual latent representations *may* have correlated factors of variation; this is potentially a more accurate reflection of real world data where different sub-populations (e.g., dog breeds) often have correlated factors of variation (e.g., color, size).

ELBO Decomposition Akin to the VAE, inference can be achieved via variational approximation. Let $q(\mathbf{z}, \Sigma|\mathbf{x})$ be the variational posterior distribution. Using Jensen’s inequality, the log-likelihood can be written as

$$\log p(\mathbf{x}) \geq \mathbb{E}_{q(\mathbf{z}, \Sigma|\mathbf{x})} \left[\log \frac{p(\mathbf{z}, \Sigma, \mathbf{x})}{q(\mathbf{z}, \Sigma|\mathbf{x})} \right] = \mathcal{L}_{\text{ELBO}} \quad (6)$$

Consider a structured variational distribution $q(\mathbf{z}, \Sigma|\mathbf{x}) = q(\mathbf{z}|\mathbf{x})q(\Sigma|\mathbf{z})$. Then, the ELBO in Eq. (6) decomposes into the following terms:

$$\begin{aligned} \mathcal{L}_{\text{ELBO}} = & \left[\frac{1}{N} \sum_{n=1}^N \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\log p(\mathbf{x}|\mathbf{z})] \right] - \underset{\text{(index-code MI)}}{I(\mathbf{x}; \mathbf{z})} \\ & - \underset{\text{(marginal KL to prior)}}{D_{\text{KL}}(q(\mathbf{z})||p(\mathbf{z}))} - \underset{\text{(covariance penalty)}}{\mathbb{E}_{q(\mathbf{z})} [D_{\text{KL}}(q(\Sigma|\mathbf{z})||p(\Sigma|\mathbf{z}))]} \end{aligned} \quad (7)$$

We can combine the first three terms in Eq. (7) to obtain the standard VAE ELBO (Hoffman and Johnson 2016):

$$\mathcal{L}_{\text{ELBO}} = \mathcal{L}_{\text{ELBO}}^{\text{VAE}} - \mathbb{E}_{q(\mathbf{z})} [D_{\text{KL}}(q(\Sigma|\mathbf{z})||p(\Sigma|\mathbf{z}))]$$

revealing $\mathbb{E}_{q(\mathbf{z})} [D_{\text{KL}}(q(\Sigma|\mathbf{z})||p(\Sigma|\mathbf{z}))]$ as an additional quantity minimized in our model. Intuitively, this term




Figure 1: Desired matrix Σ_0 and a random sample from an inverse-Wishart distribution with $\nu = 10$, $\nu = 100$, and $\nu = 1000$. As ν increases, random samples become closer to the desired matrix.

matches the approximate covariance distribution to the covariance prior (across the latent code distribution). Hence, a factorized $p(\Sigma|\mathbf{z})$ can encourage disentangled representations. A similar term was introduced directly into the ELBO in the DIP-VAE (Kumar, Sattigeri, and Balakrishnan 2017), which constrains the covariance matrix (averaged over the mini-batch) to be near \mathbf{I} . In contrast, our hyperprior approach enables this term to emerge naturally in the ELBO.

3.1 Model Specification and Inference

In this subsection, we derive a specific model—the Covariance Hyperprior VAE (CHyVAE)—under the hyperprior framework outlined above. Specifically, we set a Gaussian prior over the latent code, and an inverse-Wishart prior over its covariance matrix. Note that the overarching framework is not constrained to these specific distributions, i.e., alternative hyperpriors and variational distributions can be used without significant changes to the overall methodology.

An Inverse-Wishart Hyperprior The inverse-Wishart distribution $\mathcal{W}_p^{-1}(\Psi, \nu)$ is a popular distribution that has support over real-valued positive-definite matrices. It is parameterized by a positive-definite scale matrix $\Psi \in \mathbb{R}^{p \times p}$ and degrees of freedom (DoF) $\nu > p - 1$. An inverse-Wishart distributed random matrix $\mathbf{X} \in \mathbb{R}^{p \times p}$ has probability density function

$$p(\mathbf{X}) = \frac{|\Psi|^{\nu/2}}{2^{\nu p/2} \Gamma_p(\frac{\nu}{2})} |\mathbf{X}|^{-(\nu+p+1)/2} e^{-\frac{1}{2} \text{tr}(\Psi \mathbf{X}^{-1})}$$

where $\Gamma_p(\cdot)$ is the multivariate gamma function. The mean of an inverse-Wishart random variable is given by $(\nu - p - 1)^{-1} \Psi$. For a desired specification Σ_0 of the covariance matrix, a reasonable choice of Ψ would be $(\nu - p - 1) \Sigma_0$.

With a desired covariance matrix \mathbf{I} , the DoF parameter ν can be varied to control the desired statistical independence (Fig. 1). Intuitively, ν can be regarded as pseudo-observations and thus, controls the strength/informativeness of the prior; high values ($\nu \gg p$) indicate a strong prior, while $\nu = p + 1$ is the least informative setting.

Approximate Inference In our model, we have

$$p(\Sigma) = \mathcal{W}_p^{-1}(\Sigma|\Psi, \nu) \quad (8)$$

$$p(\mathbf{z}|\Sigma) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \Sigma) \quad (9)$$

and maximizing the log probability of the data (Eq. 4) is intractable. Our initial approach was to completely employ a mean-field variational approximation, i.e., factorize

$q(\mathbf{z}, \Sigma|\mathbf{x})$ into $q(\mathbf{z}|\mathbf{x})q(\Sigma|\mathbf{x})$. This factorization can be realized using shared neural networks that output both the prior and the hyperprior parameters. While simple and tractable, our preliminary experiments with this factorized form were unsuccessful; training was unstable and results were poor. One potential reason is that the decoupling renders the hyperprior ineffective.

An alternative approach is to factorize $q(\mathbf{z}, \Sigma|\mathbf{x}) = q(\mathbf{z}|\mathbf{x}, \Sigma)q(\Sigma|\mathbf{x})$, but explicit reparameterization is not applicable to the inverse-Wishart². It is also possible to employ a single variational distribution $q(\mathbf{z}|\mathbf{x})$ by recognizing that marginalization of the prior $p(\mathbf{z}|\Sigma)$ under an inverse-Wishart hyperprior leads to a multivariate Student’s t -distribution. However, explicit reparameterization is also not applicable in this case and analytic marginalization may not be possible with arbitrary hyperprior specifications.

In the following, we describe approximate inference using a structured variational distribution $q(\mathbf{z}, \Sigma|\mathbf{x}) = q(\mathbf{z}|\mathbf{x})q(\Sigma|\mathbf{z})$. Unlike the VAE, we learn a full covariance matrix and set the conditional $q(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}|\tilde{\boldsymbol{\mu}}, \tilde{\Sigma} = \tilde{\mathbf{L}}\tilde{\mathbf{L}}^\top)$ where $\tilde{\boldsymbol{\mu}}$ is the mean vector and $\tilde{\mathbf{L}}$ is the Cholesky factor of the covariance matrix, $\tilde{\Sigma} = \tilde{\mathbf{L}}\tilde{\mathbf{L}}^\top$. Both $\tilde{\boldsymbol{\mu}}$ and $\tilde{\mathbf{L}}$ are obtained via a neural network $f_\phi(\mathbf{x})$. Samples of \mathbf{z} from $q(\mathbf{z}|\mathbf{x})$ are obtained via explicit reparameterization,

$$\mathbf{z} = \tilde{\boldsymbol{\mu}} + \tilde{\mathbf{L}}\boldsymbol{\varepsilon} \quad (10)$$

where $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Next, we describe how $q(\Sigma|\mathbf{z})$ can be estimated efficiently. $\mathcal{L}_{\text{ELBO}}$ in Eq. (6) can be written as

$$\begin{aligned} \mathcal{L}_{\text{ELBO}} &= \mathbb{E}_{q(\mathbf{z}|\mathbf{x})q(\Sigma|\mathbf{z})} \left[\log \frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{q(\mathbf{z}|\mathbf{x})} + \log \frac{p(\Sigma|\mathbf{z})}{q(\Sigma|\mathbf{z})} \right] \\ &= \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \left[\log \frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{q(\mathbf{z}|\mathbf{x})} \right] \\ &\quad - \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [D_{\text{KL}}(q(\Sigma|\mathbf{z})||p(\Sigma|\mathbf{z}))]. \end{aligned} \quad (11)$$

The first term in Eq. (11) is independent of Σ and the second term is non-negative. As such, $\mathcal{L}_{\text{ELBO}}$ is maximized when $q(\Sigma|\mathbf{z})$ matches $p(\Sigma|\mathbf{z})$. With this in mind,

$$q(\Sigma|\mathbf{z}) \approx p(\Sigma|\mathbf{z}) = \frac{p(\mathbf{z}|\Sigma)p(\Sigma)}{\int_{\Sigma'} p(\mathbf{z}|\Sigma')p(\Sigma')d\Sigma'} \quad (12)$$

For $p(\mathbf{z}|\Sigma) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \Sigma)$, $p(\Sigma) = \mathcal{W}_p^{-1}(\Sigma|\Psi, \nu)$, and a sample $\mathbf{z}_i \sim p(\mathbf{z}|\Sigma)$, we can exploit the fact that the inverse-Wishart is a conjugate prior for the multivariate normal. We marginalize Σ' from the denominator in Eq. (12) and obtain $p(\Sigma|\mathbf{z}) = \mathcal{W}_p^{-1}(\Psi + \mathbf{z}_i\mathbf{z}_i^\top, \nu + 1)$. Using this distribution for $q(\Sigma|\mathbf{z})$, we now write the ELBO as

$$\begin{aligned} \mathcal{L}_{\text{ELBO}} &= \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\log p(\mathbf{x}|\mathbf{z})] \\ &\quad - \mathbb{E}_{p(\Sigma|\mathbf{z})} [D_{\text{KL}}(\mathcal{N}(\tilde{\boldsymbol{\mu}}, \tilde{\Sigma})||\mathcal{N}(\mathbf{0}, \Sigma))] \\ &\quad - \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [D_{\text{KL}}(\mathcal{W}_p^{-1}(\Phi, \lambda)||\mathcal{W}_p^{-1}(\Psi, \nu))] \end{aligned} \quad (13)$$

where $\Phi = \Psi + \mathbf{z}_i\mathbf{z}_i^\top$ and $\lambda = \nu + 1$.

²Very recent work (Figurnov, Mohamed, and Mnih 2018) may alleviate this issue and applying this technique is future work.

All three terms in the lower bound above have closed-form expressions and can be computed in a straight-forward manner (please refer to the supplementary material³ for detailed expressions). The first term is the reconstruction error, similar to other VAE based models. The second term represents the distance from the prior and discourages the latent codes from being too far away from the zero mean prior (this enables sampling and ensures that CHyVAE remains a valid generative model). The third term is an additional penalty on the covariance matrix; to encourage disentanglement, the prior is set as the identity matrix \mathbf{I} . As previously mentioned, when ν is increased, independence in the latent dimensions is more enforced, leading to disentangled representations.

Sample Generation Generally, we use Bartlett decomposition to obtain samples from the inverse-Wishart distribution (more details in the supplementary material). For models trained with large values of ν , we found that directly sampling from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ also generates good images.

4 Related Work

Early works that have demonstrated disentanglement in limited settings include (Schmidhuber 1992; Glorot, Bordes, and Bengio 2011; Desjardins, Courville, and Bengio 2012), and several prior research has addressed the problem of disentanglement in *supervised* or *semi-supervised* settings (Kulkarni et al. 2015; Kingma et al. 2014; Reed et al. 2014; Siddharth et al. 2017). In this work, we focus on unsupervised learning of disentangled features. Unsupervised generative models such as (Kingma and Welling 2013; Makhzani et al. 2015; Radford, Metz, and Chintala 2015) have been also shown to learn disentangled representations, although this was not the main motivation of these works. Recent work has also sought to disentangle factors of variation in sequential data in an unsupervised manner (Denton and Birodkar 2017; Hsu, Zhang, and Glass 2017).

VAE-based Models Our work builds upon the VAE and is inspired by recent work on learning disentangled factors (Higgins et al. 2017; Kim and Mnih 2018; Chen et al. 2018). In addition to these papers (covered in Sec. 2), there has been further work in uncovering the principles behind disentanglement in VAEs. Burgess et al. (2017) argue from the information bottleneck principle that penalizing mutual information in β -VAE results in a compact and disentangled representation. Based on their analyses, Kumar, Sattigeri, and Balakrishnan (2017) add an additional penalty to the VAE ELBO based on how much the covariance of $q(\mathbf{z})$ deviates from \mathbf{I} . In contrast to these previous work, we attempt to introduce disentanglement at the model specification stage through the covariance hyperprior.

InfoGAN An alternative approach towards deep generative modeling is the Generative Adversarial Network (GAN) (Goodfellow et al. 2014). Chen et al. (2016) have

³Supplementary material for this paper is available at <https://arxiv.org/abs/1809.04497>

argued that maximizing mutual information between the observed sample and a subset of latent codes encourages disentanglement and capitalized on this idea to develop the InfoGAN. Kim and Mnih (2018) evaluated the disentanglement performance of InfoWGAN-GP, a version of InfoGAN that uses WGAN (Arjovsky, Chintala, and Bottou 2017) and gradient penalty (Gulrajani et al. 2017).

Other priors Previous work has explored different priors for the latent space in VAEs including mixture of Gaussians (Jiang et al. 2016), Dirichlet process (Nalisnick and Smyth 2016), Beta, Gamma, and von Mises (Figurnov, Mohamed, and Mnih 2018). With advancements in reparameterization for discrete distributions, recent work (Esmaili et al. 2018; Pineau and Lelarge 2018) have proposed adding different priors to different subsets of the latent code to separately model discrete and continuous factors of variation in the data. However, to the best of our knowledge, we are the first to apply hierarchical priors towards learning disentangled representations.

5 Experiments

In this section, we report on experiments comparing CHyVAE to two VAE-based unsupervised disentangling models: the β -VAE and state-of-the-art FactorVAE⁴. Due to space constraints, we briefly describe the experimental setup and focus on the main findings; details are available in the supplementary material and our code base is available for download at <https://github.com/crslab/CHyVAE>.

5.1 Experimental Setup

Model Implementation and Training To ease comparisons between the methods and prior work, we use the same network architecture across all the compared methods. Specifically, we follow the model in (Kim and Mnih 2018): a convolutional neural network (CNN) for the encoder and a deconvolutional NN for the decoder. We normalize all datasets to $[0, 1]$ and use sigmoid cross-entropy as the reconstruction loss function. For training, we use Adam optimizer (Kingma and Ba 2014) with a learning rate of 10^{-4} . For the discriminator in FactorVAE, we use the parameters recommended by Kim and Mnih (2018).

Datasets Our experiments were conducted using five datasets, including four standard benchmarks:

- 1) Datasets with *known* generative factors:
 - (a) **2DShapes** (or *dSprites*) (Matthey et al. 2017): 737,280 binary 64×64 images of 2D shapes (heart, square, ellipse) with five ground truth factors of variation namely x-position, y-position, scale, orientation, and shape. All factors except shape are continuous.
 - (b) **CorrelatedEllipses**: 200,000 grayscale 64×64 images of ellipses with dependent ground truth factors

⁴We exclude comparisons with InfoGAN as the VAE-based models have been shown to obtain better disentanglement performance relative to InfoGAN in previous work (Higgins et al. 2017; Kim and Mnih 2018).

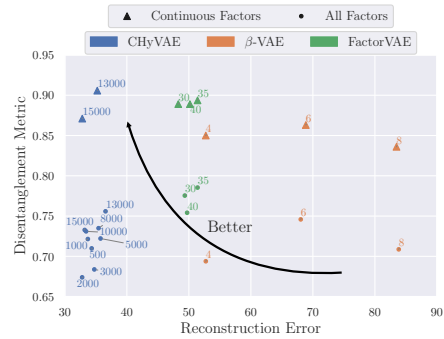


Figure 2: Disentanglement Metric plotted against Reconstruction error for CHyVAE, β -VAE, and FactorVAE on 2DShapes dataset averaged over 10 random restarts.



Figure 3: Disentanglement Metric plotted against Reconstruction error for CHyVAE, β -VAE, and FactorVAE on CorrelatedEllipses dataset averaged over 10 random restarts.

x-position correlated with y-position and scale correlated with orientation. This dataset embodies the “adversarial” case where the factors may be correlated or the dataset was obtained with sampling bias (arguably common in many real-world data). Dataset construction details are in the supplementary material.

- 2) Datasets with unknown generative factors:
 - (a) **3DFaces** (Paysan et al. 2009): 239,840 greyscale 64×64 images of 3D Faces.
 - (b) **3DChairs** (Aubry et al. 2014): 86,366 RGB 64×64 images of CAD chair models.
 - (c) **CelebA** (Liu et al. 2015): 202,599 RGB images of celebrity faces center-cropped to dimensions 64×64 .

Disentanglement Metric and Latent Traversals Suitable evaluation criteria for disentanglement remains an area of active research. Several metrics have been recently proposed based on linear mappings from latent codes to generative factors (Higgins et al. 2017; Eastwood and Williams 2018) and mutual information (Chen et al. 2018).

We use the metric proposed by Kim and Mnih (2018) for evaluating the models primarily because of its interpretability and computational efficiency. The metric uses a majority vote classifier matrix $\mathbf{V}^{p \times K}$ that maps each latent dimension to only one ground truth factor where p is the dimension of

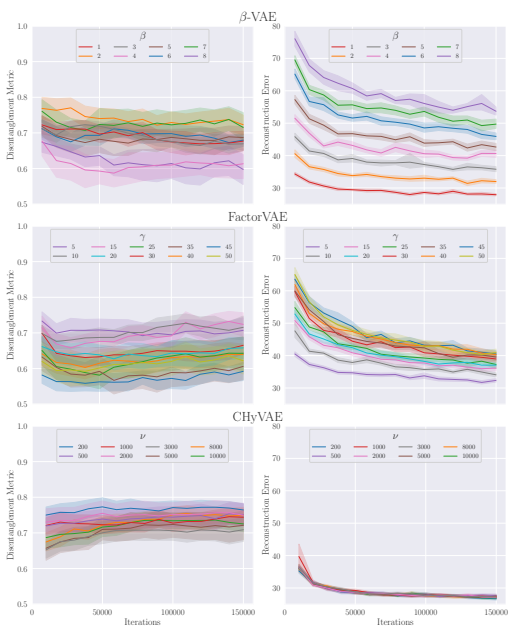


Figure 4: Disentanglement metric scores (left) and reconstruction error (right) with iterations for β -VAE, FactorVAE, and CHyVAE on CorrelatedEllipses dataset averaged over 10 random restarts.

the latent code and K is the number of ground truth factors. Each element \mathbf{V}_{ij} for $i \in \{1 \dots p\}, j \in \{1 \dots K\}$ is a count of the number of batches with a fixed factor j that have minimum variance in the dimension i of the latent code. Using the vote matrix in the metric each latent dimension can be mapped to a ground truth factor and the dimensions can be annotated. Note that quantitative evaluation can only be performed on datasets with *known* factors of variation.

When the factors of variation are *unknown*, it is common to examine latent traversals. These traversals are obtained by fixing all latent dimensions and varying only one. Inspection of latent traversals tells us little about the robustness of a model but is currently the only available method of comparing disentanglement performance on datasets with unknown factors of variation.

5.2 Quantitative Evaluation

Our main results on the 2DShapes dataset are summarized in Figure 2; it shows the disentanglement metric plotted against the reconstruction error (scores averaged over 10 random restarts) for varying values of β , γ , and ν^5 . Better performing methods fall on the top left of the graph (high disentanglement with low reconstruction error).

Figure 2 clearly shows that CHyVAE outperforms both β -VAE and FactorVAE on the continuous factors, achieving far better reconstruction error at similar—if not, slightly better—disentangling performance. Higher values of ν tend to produce better disentanglement, while preserving recon-

⁵For β -VAE and FactorVAE, we show results using the best performing hyperparameter values reported in Kim and Mnih (2018).

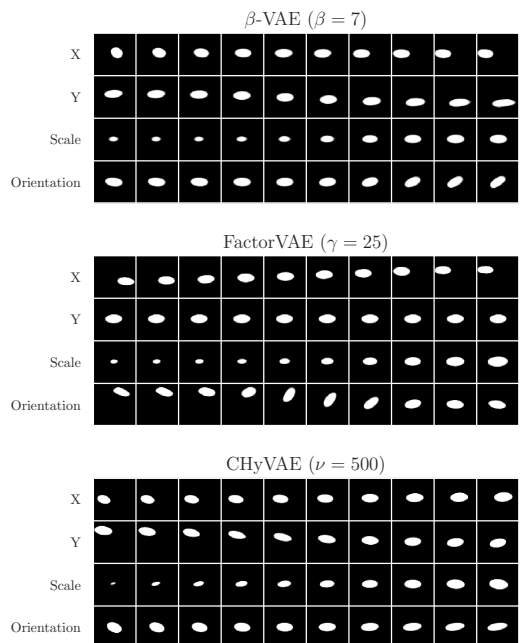


Figure 5: Traversals across latent dimensions on CorrelatedEllipses dataset annotated with the factor of variation assigned by the majority vote matrix for best performing β -VAE, FactorVAE, and CHyVAE.

struction capability. When the discrete factor (shape) is included, the FactorVAE achieves slightly higher disentanglement on average, but the best performing models are comparable—0.909 ($\gamma = 35$) and 0.905 ($\nu = 13000$) for FactorVAE and CHyVAE respectively. However, the latent traversals show that all the models struggle with the discrete factor (in supplementary material). For the CHyVAE, it is unsurprising that the disentanglement would be poorer with discrete latent factors: CHyVAE enforces a hierarchical prior on a *continuous* latent space. Discrete factors can be handled in a principled manner within our framework by using a suitable prior (Esmaili et al. 2018; Pineau and Lelarge 2018) with associated hyperprior.

Results for CorrelatedEllipses are summarized in Fig. 3. CHyVAE starkly outperforms β -VAE and FactorVAE both in terms of the metric and the reconstruction error across the different parameters. We posit that this was due to the extra flexibility afforded by the prior and hyperprior and learning a full covariance matrix; lower values of ν , which allow more deviation from the identity covariance, achieve better disentanglement and reconstruction. Figure 4 shows the the disentanglement metric and the reconstruction error as training progressed for different parameter values; compared to FactorVAE and β -VAE, CHyVAE achieves lower reconstruction error and better disentanglement. Figure 5 shows the latent traversals for best performing models on the disentanglement metric. Interestingly, both β -VAE and CHyVAE learn a slightly-entangled representation for the y-position but the FactorVAE fails to capture this factor.




Figure 6: Traversals across latent dimensions for β -VAE, FactorVAE, and CHyVAE on 3DFaces dataset annotated with the factor of variation.

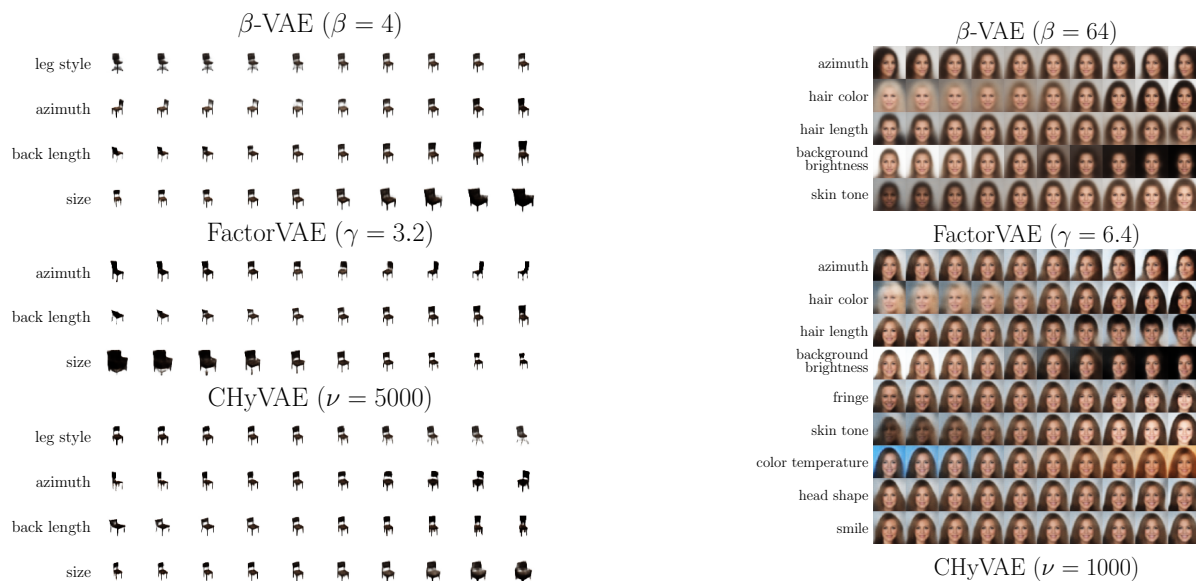


Figure 7: Traversals across latent dimensions for β -VAE, FactorVAE, and CHyVAE on 3DChairs dataset annotated with the factor of variation.

5.3 Qualitative Evaluation

In absence of a metric for comparison of the disentangling performance of different models on datasets with unknown generative factors, the only evaluation method available is inspecting latent traversals. Figs. 6 and 8 show that CHyVAE is able to learn semantically reasonable factors of variation for 3DFaces and CelebA. For 3DChairs (Fig. 7) CHyVAE is able to learn the *leg-style* factor which is missed by FactorVAE but learnt by the β -VAE. In terms of reconstruction, CHyVAE achieves superior performance relative to β -VAE and comparable to FactorVAE; see Fig. 9 and refer the supplementary material for plots for 3DFaces and 3DChairs.

6 Conclusion

State-of-the-art methods for learning disentangled representations in VAEs have focussed primarily on manipulating the ELBO. In contrast, we pursued an alternative principled approach by placing a hyperprior on the covariance matrix of the VAE prior. The inverse-Wishart used in our study exposes its degrees-of-freedom parameter which can be tuned to control the informativeness of a desired independent covariance \mathbf{I} and thus, encourage disentanglement.

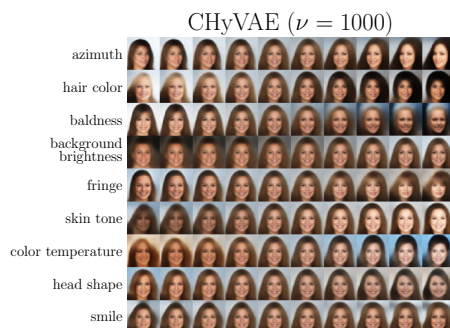


Figure 8: Traversals across latent dimensions for β -VAE, FactorVAE, and CHyVAE on CelebA dataset annotated with the most prominently varying factor of variation.




Figure 9: Reconstruction error vs iterations for β -VAE, FactorVAE, and CHyVAE on CelebA averaged over 5 random restarts.

Extensive experiments on a variety of datasets show that our model, CHyVAE, outperforms the β -VAE and is comparable to the FactorVAE in terms of disentanglement, while achieving better reconstruction. Our experimental results with a new dataset also demonstrate that encouraging factorial codes may not learn suitable disentangled representations when correlations are present; instead, a more flexible model such as CHyVAE may disentangle better.

While we have focussed on the inverse-Wishart hyperprior in this work, our key idea of using a hierarchical model can be extended to alternative distributions. As future work, we plan to examine the effects of different hyperpriors, and extend the approach towards learning disentangled representations with both discrete and continuous latent variables. We also plan to explore technical improvements, e.g., the structured factorization explored in this work is an improvement over standard mean-field, but it may lose information due to the choice of simplified distributions and the optimality assumption underlying the approximation of $q(\Sigma|\mathbf{z})$.

References

- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *ICML*.
- Aubry, M.; Maturana, D.; Efros, A. A.; Russell, B. C.; and Sivic, J. 2014. Seeing 3D Chairs: Exemplar Part-Based 2D-3D Alignment Using a Large Dataset of CAD Models. In *CVPR*.
- Bengio, Y.; Courville, A. C.; and Vincent, P. 2013. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35:1798–1828.
- Burgess, C. P.; Higgins, I.; Pal, A.; Matthey, L.; Watters, N.; Desjardins, G.; and Lerchner, A. 2017. Understanding disentangling in β -VAE. In *NIPS Workshop on Learning Disentangled Representations: from Perception to Control*.
- Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS*.
- Chen, T. Q.; Li, X.; Grosse, R. B.; and Duvenaud, D. K. 2018. Isolating sources of disentanglement in variational autoencoders. *CoRR* abs/1802.04942.
- Denton, E. L., and Birodkar, V. 2017. Unsupervised learning of disentangled representations from video. In *NIPS*.
- Desjardins, G.; Courville, A. C.; and Bengio, Y. 2012. Disentangling factors of variation via generative entangling. *CoRR* abs/1210.5474.
- Doersch, C. 2016. Tutorial on Variational Autoencoders. 1–23.
- Eastwood, C., and Williams, C. K. I. 2018. A framework for the quantitative evaluation of disentangled representations. In *ICLR*.
- Esmaeili, B.; Wu, H.; Jain, S.; Bozkurt, A.; Siddharth, N.; Paige, B.; Brooks, D. H.; Dy, J.; and van de Meent, J.-W. 2018. Structured Disentangled Representations. *ArXiv e-prints*.
- Figurnov, M.; Mohamed, S.; and Mnih, A. 2018. Implicit Reparameterization Gradients. *ArXiv e-prints*.
- Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NIPS*.
- Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. C. 2017. Improved training of wasserstein gans. In *NIPS*.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2017. β -VAE: Learning basic visual concepts with a constrained variational framework. In *ICLR*.
- Hoffman, M. D., and Johnson, M. J. 2016. Elbo surgery: yet another way to carve up the variational evidence lower bound. In *NIPS Workshop in Advances in Approximate Bayesian Inference*.
- Hsu, W.-N.; Zhang, Y.; and Glass, J. R. 2017. Unsupervised learning of disentangled and interpretable representations from sequential data. In *NIPS*.
- Jiang, Z.; Zheng, Y.; Tan, H.; Tang, B.; and Zhou, H. 2016. Variational Deep Embedding: An Unsupervised and Generative Approach to Clustering. In *IJCAI*.
- Kim, H., and Mnih, A. 2018. Disentangling by factorising. In *ICML*.
- Kingma, D. P., and Ba, J. 2014. Adam: A Method for Stochastic Optimization. In *ICLR*.
- Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes. *CoRR* abs/1312.6114.
- Kingma, D. P.; Rezende, D. J.; Mohamed, S.; and Welling, M. 2014. Semi-supervised learning with deep generative models. In *NIPS*.
- Kulkarni, T. D.; Whitney, W. F.; Kohli, P.; and Tenenbaum, J. B. 2015. Deep convolutional inverse graphics network. In *NIPS*.
- Kumar, A.; Sattigeri, P.; and Balakrishnan, A. 2017. Variational inference of disentangled latent concepts from unlabeled observations. *CoRR* abs/1711.00848.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep learning face attributes in the wild. In *ICCV*.
- Makhzani, A.; Shlens, J.; Jaitly, N.; and Goodfellow, I. J. 2015. Adversarial autoencoders. *CoRR* abs/1511.05644.
- Mathieu, M.; Zhao, J. J.; Sprechmann, P.; Ramesh, A.; and LeCun, Y. 2016. Disentangling factors of variation in deep representations using adversarial training. In *NIPS*.
- Matthey, L.; Higgins, I.; Hassabis, D.; and Lerchner, A. 2017. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>.
- Nalisnick, E., and Smyth, P. 2016. Stick-Breaking Variational Autoencoders. In *ICLR*.
- Paysan, P.; Knothe, R.; Amberg, B.; Romdhani, S.; and Vetter, T. 2009. A 3d face model for pose and illumination invariant face recognition. *2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance* 296–301.
- Pineau, E., and Lelarge, M. 2018. InfoCatVAE: Representation learning with categorical variational autoencoders. *CoRR* abs/1806.08240.
- Radford, A.; Metz, L.; and Chintala, S. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *CoRR* abs/1511.06434.
- Reed, S. E.; Sohn, K.; Zhang, Y.; and Lee, H. 2014. Learning to disentangle factors of variation with manifold interaction. In *ICML*.
- Schmidhuber, J. 1992. Learning factorial codes by predictability minimization. *Neural Computation* 4.
- Siddharth, N.; Paige, B.; van de Meent, J.-W.; Desmaison, A.; Wood, F. D.; Goodman, N. D.; Kohli, P.; and Torr, P. H. S. 2017. Learning disentangled representations with semi-supervised deep generative models. In *NIPS*.