
On Implicit Regularization in β -VAEs

Abhishek Kumar¹ Ben Poole¹

Abstract

While the impact of variational inference (VI) on posterior inference in a fixed generative model is well-characterized, its role in regularizing a learned generative model when used in variational autoencoders (VAEs) is poorly understood. We study the regularizing effects of variational distributions on learning in generative models from two perspectives. First, we analyze the role that the choice of variational family plays in imparting uniqueness to the learned model by restricting the set of optimal generative models. Second, we study the regularization effect of the variational family on the local geometry of the decoding model. This analysis uncovers the regularizer implicit in the β -VAE objective, and leads to an approximation consisting of a deterministic autoencoding objective plus analytic regularizers that depend on the Hessian or Jacobian of the decoding model, unifying VAEs with recent heuristics proposed for training regularized autoencoders. We empirically verify these findings, observing that the proposed deterministic objective exhibits similar behavior to the β -VAE in terms of objective value and sample quality.

1. Introduction

Variational autoencoders (VAEs) (Kingma & Welling, 2014; Rezende et al., 2014) and variants such as β -VAEs (Higgins et al., 2017; Alemi et al., 2018) have been widely used for density estimation and learning generative samplers of data where latent variables in the model may correspond to factors of variation in the underlying data-generating process (Bengio et al., 2013). However, without additional assumptions it may not be possible to recover or identify the correct model (Hyvärinen & Pajunen, 1999; Locatello et al., 2019).

¹Google Research, Brain Team. Correspondence to: Abhishek Kumar <abhishk@google.com>, Ben Poole <pooleb@google.com>.

We study the regularizing effects of variational distributions used to approximate the posterior on the learned generative model from two perspectives. First, we discuss the role of *the choice of variational family* in imparting uniqueness properties to the learned generative model by restricting the set of possible solutions. We argue that restricting the variational family in specific ways can rule out non-uniqueness to a corresponding set of transformations on the generative model. Next, we study the impact of the variational family on the geometry of the decoding model ($p_\theta(x|z)$) in β -VAEs. Assuming that the first two moments exist for the variational family, we can relate the structure of the variational distribution’s covariance to the Jacobian matrix of the decoding model. For a Gaussian prior and variational distribution, we further use this relation to derive a deterministic objective that closely approximates the β -VAE objective.

We perform experiments to validate the correctness of our theory and accuracy of our approximations. On the MNIST dataset, we empirically verify the relationship between variational distribution covariance and the Jacobian of the decoder. In particular, we observe that block-diagonal variational distribution covariance encourages a block diagonal structure on the matrices $J_g(z)^\top J_g(z)$ where $J_g(z)$ is the decoder Jacobian matrix. This generalizes the results from recent work (Rolinek et al., 2019), which showed (for Gaussian observation models) that *diagonal* Gaussian variational distributions encourage orthogonal Jacobians, to general observation models and arbitrary structures on the covariance matrix. We also compare the objective values of the original β -VAE objective with our derived deterministic objectives. While our deterministic approximation holds for small β , we empirically find that the two objectives remain in agreement for a reasonably wide range of β , verifying the accuracy of our regularizers. We also train models using a tractable lower bound on the deterministic objectives for CelebA and find that trained models behave similar to β -VAE in terms of sample quality.

2. Latent-variable models and identifiability

We begin by discussing inherent identifiability issues with latent-variable models trained using maximum log-likelihood alone.

Latent-variable models. We consider latent-variable generative models that consist of a fixed prior $p(z)$, and a conditional generative model or *decoder*, $p_\theta(x|z)$. In maximum-likelihood estimation, we aim to maximize the marginal log-likelihood of the training data, $\log p_\theta(x) = \log \int dz p_\theta(z)p_\theta(x|z)$ in terms of the parameters θ . A desirable goal in learning latent-variable models is that the learned latent variables, z , correspond simply to “ground-truth factors of variation” that generated the dataset (Bengio et al., 2013; Higgins et al., 2017). For example, in a dataset of shapes we aim to learn latent variables that are in one-to-one correspondence with attributes of the shape like size, color, and position. Learning such “disentangled” representations may prove useful for downstream tasks such as interpretability and few-shot learning (Gilpin et al., 2018; van Steenkiste et al., 2019).

Identifiability and non-uniqueness. Unfortunately, there is an inherent identifiability issue with latent-variable models trained with MLE: there are many models that have the same marginal data and prior distributions, $p_\theta(x)$ and $p(z)$, but entirely different latent variables z . We can find transformations r of the latent variable z that leave its marginal distribution $p(z)$ the same, and thus the marginal data distribution $p_\theta(x)$ remains the same for a fixed decoder $p_\theta(x|z)$. The transformation r is not necessarily just a simple permutation or nonlinear transformation of individual latents, but *mixes* components of the latent representation. This leads to a *set* of solutions instead of a *unique* solution, all of which are equal according to their marginal log-likelihood, but that have distinct representations. Thus identifying the “right” model whose latents are in one-to-one correspondence with factors of variation is not feasible with MLE alone. In our work, we adopt the definition of identifiability from the ICA literature that ignores permutations and nonlinear transformations that act separately on each latent dimension (Hyvärinen & Pajunen, 1999). This identifiability issue is well-known in the ICA literature (Hyvärinen & Pajunen, 1999) and causal inference literature (Peters et al., 2017), and has been highlighted more recently in the context of learning disentangled representations with factorized priors (Locatello et al., 2019). In Appendix A, we summarize the construction from Locatello et al. (2019) for factorized priors, and extend their non-uniqueness result to correlated priors.

3. Uniqueness via variational family

In the previous section, we highlighted how there can be an infinite family of latent-variable models that all achieve the same marginal likelihood. Here, we show how additional constraints present in the learning algorithm can restrict the space of solutions leading to uniqueness.

We focus on the impact of posterior regularization (Zhu

et al., 2014; Shu et al., 2018) via the choice of variational family on the set of solutions to the β -VAE objective. In variational inference, we optimize not only the generative model, $p(x, z) = p(x|z)p(z)$, but also a variational distribution $q(z|x)$ to approximate the intractable posterior $p(z|x)$. When optimizing over $q(z|x)$, we constrain the distributions to come from some family \mathcal{Q} .

We first consider the case when the true posterior lies in the variational family \mathcal{Q} . Although there are many latent variable generative models with the same marginal data distribution $p(x)$, this equivalence class of models shrinks when we constrain ourselves to models with posteriors $p(z|x) \in \mathcal{Q}$. Consider applying a mixing transform r , with $|\det J_r(z)| = 1$, that preserves the prior distribution of z before feeding it to the decoder: $z' = r(z)$. The transformed posterior is given by:

$$\begin{aligned} p'(z|x) &= \frac{p(x|r(z))p(z)}{p(x)} = \frac{p(x|z')p(r^{-1}(z'))}{p(x)} \\ &= \frac{p(x|z')p(z')}{p(x)} = p(z'|x) = p(r(z)|x), \end{aligned} \quad (1)$$

where we used the facts that $p(z) = p(r(z)) = p(r^{-1}(z))$ and $|\det J_r(z)| = 1$. This implies $p(z|x) = p'(r^{-1}(z)|x)$, *i.e.*, the posterior r.v. is transformed by r^{-1} yielding $p'_{z|x} = r^{-1}\#p_{z|x}$ ¹. If $p_{z|x} \in \mathcal{Q}$ but $p'_{z|x} \notin \mathcal{Q}$, then the best ELBO we can achieve in the transformed model p' is strictly less than the untransformed model p . Thus even though r leaves the marginal data likelihood unchanged, it yields a worse ELBO and thus can no longer be considered an equivalent model under the training objective. The question remains that for a distribution family \mathcal{Q} and a $q \in \mathcal{Q}$, what class of transformations R guarantee that $r^{-1}\#q \notin \mathcal{Q}$ for all $r \in R$? This has to be answered on a case-by-case basis, but we consider an example below.

3.1. Example: Isotropic Gaussian prior and orthogonal transforms

An isotropic Gaussian prior is invariant under orthonormal transformations (characterized by orthogonal matrices with unit norm columns). For a mean-field family of distributions \mathcal{Q} , the necessary and sufficient conditions for $r^{-1}\#p_{z|x} \in \mathcal{Q}$, given $p_{z|x} \in \mathcal{Q}$, for all r belonging to the set of orthonormal transformations, are (Skitovitch, 1953; Darrois, 1953; Lukacs et al., 1954): (i) $p_{z|x}$ is a factorized Gaussian, and (ii) variances of the individual latent dimensions of the posterior are all equal. If either of these two conditions are not satisfied, the orthogonal symmetry in the generative model, which arises from using the isotropic Gaussian prior, is broken. When all latent variances are

¹For clarity, we denote the transformed distribution using push-forward notation: where $f\#p$ denotes the distribution created by transforming samples from p using the deterministic function f .

unequal, the only orthogonal transforms r allowed on the generative model will be the permutations. This Darmois-Skitovitch characterization of the Gaussian random vectors has also been used to prove identifiability results for linear ICA (Hyvärinen & Pajunen, 1999). However, in our context, we do not need any linearity assumption on the conditional generative model $p(x|z)$.

Probabilistic PCA (Tipping & Bishop, 1999) is one example in this model class with $p(z) = \mathcal{N}(0, I)$, $p(x|z) = \mathcal{N}(Az, \sigma^2 I)$ and $p(x) = \mathcal{N}(0, AA^\top + \sigma^2 I)$, that has a Gaussian posterior which is factored when A has orthogonal columns. However, there are equivalent models $p(x|z) = (AVz, \sigma^2 I)$ with V orthonormal that yield the same $p(x)$ and still have a factorized posterior if the columns of A have equal norms. This symmetry vanishes only when all columns of A have different norms (condition (ii) mentioned earlier).

3.2. Uniqueness in β -VAE

The choice of variational family can lead to uniqueness even when the true posterior is not in the variational family. Consider the β -VAE objective (Higgins et al., 2017):

$$L(q, p) := \mathbb{E}_{q(z|x)} \log p_{x|z}(x|z) - \beta \text{KL}(q(z|x) \| p_z(z)). \quad (2)$$

Note that we maximize the objective w.r.t. only the decoder $p_{x|z}$ (not the prior p_z) but use the shorthand $L(q, p)$ in place of $L(q, p_{x|z})$ for brevity. Let \mathcal{P} denote the set of all generative models with fixed prior p_z but arbitrary decoder $p_{x|z}$. Let R be the class of (deterministic) mixing transforms that leave the prior invariant, i.e., $p_z = r \# p_z$. Note that $r \in R \Rightarrow r^{-1} \in R$. Given a solution (q^*, p^*) , s.t., $(q^*, p^*) = \max_{q \in \mathcal{Q}, p \in \mathcal{P}} L(q, p)$, we would like to know if there is a transformational non-uniqueness associated with it, i.e., if it is possible to transform the generative model by r and get the same marginal $p(x)$ as well as same objective value of (2). The following result gives a condition under which this non-uniqueness is ruled out.

Theorem 1 (Proof in Appendix B). *Let R be the class of mixing transforms that leave the prior invariant, and p_r be the generative model obtained by transforming the latents of the model p by $r \in R$. Let $L(q^*, p^*) = \max_{q \in \mathcal{Q}, p \in \mathcal{P}} L(q, p)$. Let $\tilde{\mathcal{Q}}$ be the completion of \mathcal{Q} by R , i.e., $\tilde{\mathcal{Q}} = \{r \# q : q \in \mathcal{Q}, r \in R\} \cup \mathcal{Q}$. If \mathcal{Q} is s.t. $q \in \mathcal{Q} \Rightarrow r \# q \notin \mathcal{Q}$ for all $r \in R$, and $\arg \max_{q \in \tilde{\mathcal{Q}}} L(q, p^*)$ is unique, then $\max_{q \in \mathcal{Q}} L(q, p_r^*) < L(q^*, p^*)$.*

Corollary 1 (Proof in Appendix B). *If the set $\tilde{\mathcal{Q}}$ defined in Theorem 1 is convex, then $\arg \max_{q \in \tilde{\mathcal{Q}}} L(q, p^*)$ is unique. When all other assumptions of Theorem 1 are met, we have $\max_{q \in \mathcal{Q}} L(q, p_r^*) < L(q^*, p^*)$.*

Hence, for a choice of variational family \mathcal{Q} such that the

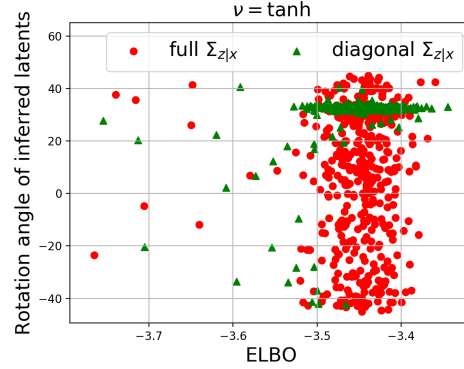


Figure 1. Rotation angle between true latents and inferred latents for VAEs (multiple runs): diagonal Σ encourages unique solution in terms of inferred latents as the ELBO increases. Please refer to text for details. More experiments are in Appendix C.

tuple (\mathcal{Q}, R) satisfies the condition of Theorem 1, transforming the latents by $r \in R$ will result in decreasing the optimal value of the β -VAE objective, thus breaking the non-uniqueness. From Section 3.1, we know that using a Gaussian prior, and restricting \mathcal{Q} to factorized *non-Gaussian* distributions is one such choice that ensures a unique (q^*, p^*) when R is the class of orthogonal transformations. Thus the choice of prior and variational family can help to guarantee uniqueness of models trained with the β -VAE objective w.r.t. transformations in R . We conduct an experiment highlighting this empirically. We generate $4k$ samples in two dimensions using $x = \nu(Wz) + \epsilon$, with $z \sim \mathcal{N}(0, I)$, $\epsilon \sim \mathcal{N}(0, .05^2 I)$. Entries of $W \in \mathbb{R}^{2 \times 2}$ are sampled from $\mathcal{N}(0, 1)$ and normalized such that columns have different ℓ_2 norms (2 and 1, respectively). ν is either Identity (Probabilistic PCA) or a nonlinearity (tanh, sigmoid, elu). A VAE is trained on this data, with decoder of the form $x = \nu(Az) + b$ and variational posterior of the form $q(z|x) = \mathcal{N}(Cx + d, \Sigma)$ with A, b, C, d and Σ learned. We do multiple random runs of this experiment and measure the rotation angle between true latents and inferred posterior means (after ignoring permutation and reflection of axes so the angle ranges in $(-45, 45]$). These runs are trained with a varying number of iterations so they reach solutions with varying degrees of suboptimality (ELBO). Fig. 1 shows a scatter plot of rotation angle vs ELBO for $\nu = \text{tanh}$. When Σ is diagonal (mean-field $q(z|x)$), there is a unique solution in terms of inferred latents as the ELBO increases (ignoring the permutations and reflections of latent axes). In Appendix C, we show this uniqueness is not due to additional inductive biases in the optimization algorithm, and provide results for unamortized VI and other activations ν .

Disentanglement. Locatello et al. (2019); Mathieu et al. (2019) highlight that learning disentangled representations using only unlabeled data is impossible without additional

inductive biases due to this non-uniqueness issue of multiple generative models resulting in same $p(x)$. While Mathieu et al. (2019) detail how the β -VAE objective is invariant to rotation of the latents with a flexible variational distribution, they do not consider the role mean-field plays in limiting the set of solutions. Based on the discussion in this section, it can be seen that the choice of variational family can be one such inductive bias, *i.e.*, it can be chosen in tandem with the prior to rule out issues of non-uniqueness due to certain transforms. Most existing works on unsupervised disentanglement using β -VAE's (Higgins et al., 2017; Kumar et al., 2018; Chen et al., 2018b; Kim & Mnih, 2018; Burgess et al., 2018) use a standard normal prior and factored Gaussian variational family which should not have this orthogonal non-uniqueness issue (*i.e.*, the construction in the proof of Theorem 1 of (Locatello et al., 2019) will change the ELBO), unless two or more variational distribution covariances are equal. Note that this observation is limited to the theoretical uniqueness aspect of the models and does not contradict the empirical findings of Locatello et al. (2019) who observe that β -VAEs converge to solutions with different disentanglement scores depending on the random seeds. There could be multiple reasons for this, such as models differing in their training objective values (*i.e.*, suboptimal solutions where our theoretical results do not apply), disentanglement metrics being sensitive to non-mixing transforms or the finite dataset size. We leave further investigation of these empirical aspects to future work.

4. Deterministic approximations of β -VAE

Next, we analyze how the covariance structure of the variational family can regularize the geometry of the learned generative model. We start with the β -VAE objective of Eq. (2), which reduces to the ELBO for $\beta = 1$. We use $p_\theta(x|z)$ to denote the decoding model parameterized by θ , and $q_\phi(z|x)$ to denote the amortized inference model (encoding model) parameterized by ϕ . We will sometimes omit the parameter subscripts for brevity. We explicitly consider the case where the first two moments exist for the variational distribution $q_\phi(z|x)$. We denote the latent dimensionality by d and the observation dimensionality by D .

Let us denote $f_x(z) = \log p(x|z)$ and view $\log p(x|z)$ as a scalar function of z for a given x . The second-order Taylor series expansion of $f_x(z)$ around the first moment (mean) $\mu_{z|x} = \mathbb{E}_{q_\phi(\cdot|x)}[z] = h_\phi(x)$ is given by:

$$f_x(z) \approx \log p(x|h_\phi(x)) + J_{f_x}(h_\phi(x))(z - h_\phi(x)) + \frac{1}{2}(z - h_\phi(x))^\top H_{f_x}(h_\phi(x))(z - h_\phi(x)), \quad (3)$$

where $J_{f_x}(z) \in \mathbb{R}^{1 \times d}$ and $H_{f_x}(z) \in \mathbb{R}^{d \times d}$ are the Jacobian and Hessian of f_x evaluated at z . Note that the expectation of the second term under $q_\phi(z|x)$ is zero as $\mathbb{E}_{q_\phi(\cdot|x)}[z] =$

$h_\phi(x)$. Substituting the approximation (3) in the original β -VAE objective (2), we obtain the following approximation

$$L(p, q) \approx \log p_\theta(x|h_\phi(x)) + \frac{1}{2}\text{tr}(H_{f_x}(h_\phi(x))\Sigma_{z|x}) - \beta \text{KL}(q_\phi(z|x)||p(z)), \quad (4)$$

where $\Sigma_{z|x} = \mathbb{E}_{q_\phi(z|x)}[(z - \mu_{z|x})(z - \mu_{z|x})^\top]$ is the covariance of the variational distribution. For this approximation to be accurate, we either need the higher central moments of the variational distribution to be small or the higher-order derivatives $\nabla_z^n \log p(x|z)|_{z=h_\phi(x)}$ ($n \geq 3$) to be small. For small values of β , the variational distribution will be fairly concentrated around the mean, resulting in small higher central moments (if any). While this objective no longer requires sampling, it involves a Hessian of the decoder, which can be computationally expensive to evaluate. Next we show how this Hessian can be computed tractably given certain network architectures.

Let $g_\theta : Z \rightarrow \mathbb{R}^D$ denote the deterministic component of the decoder that maps from a latent vector z to the parameters of the observation model (*e.g.*, to means of the Gaussian pixel observations, to probabilities of the Bernoulli distribution for binary images, etc.). We can further expand the Hessian $H_{f_x}(z)$ to decompose it as the sum of two terms, one involving the Jacobian of the decoder mapping $J_g(z) \in \mathbb{R}^{D \times d}$, and the other involving the Hessian of the decoder mapping, as follows:

$$\begin{aligned} H_{f_x}(z) &= \nabla_z^2 \log p(x|z) = \nabla_z^2 \log p(x; g(z)) \\ &= \nabla_z [(\nabla_{g(z)} \log p(x; g(z)))J_g(z)] \\ &= J_g(z)^\top (\nabla_{g(z)}^2 \log p(x; g(z)))J_g(z) + \\ &\quad (\nabla_{g(z)} \log p(x; g(z)))\nabla_z^2 g(z) \end{aligned} \quad (5)$$

It can be noted that $\nabla_z^2 g(z)$ in the second term is almost always zero when the decoder network uses only piecewise-linear activations (*e.g.*, relu, leaky-relu), which is the case in many implementations. Denoting $H_{p_x}(g(z)) = \nabla_{g(z)}^2 \log p(x; g(z))$, this leaves us with

$$H_{f_x}(z) = J_g(z)^\top H_{p_x}(g(z))J_g(z) \quad (6)$$

In the case of other nonlinearities, we can make the approximation that $\nabla_z^2 g(z)$ is small and can be neglected. Substituting the Hessian $H_{f_x}(z)$ from Eq. (6) into the approximate ELBO (4), we obtain

$$\begin{aligned} \max_{g, h} \log p(x|h(x)) - \beta \text{KL}(q_\phi(z|x)||p(z)) + \\ \frac{1}{2}\text{tr}(J_g(h(x))^\top H_{p_x}(g(h(x)))J_g(h(x))\Sigma_{z|x}) \end{aligned} \quad (7)$$

The above deductions in (4) and (7), obtained using the second order approximation of $\log p(x|z)$, imply that stochasticity of the variational distribution translates to regularizers that encourage the *alignment* of $H_{f_x}(h(x))$ and

$J_g(h(x))^\top H_{p_x}(g(h(x)))J_g(h(x))$ with covariance $\Sigma_{z|x}$ by maximizing their inner product with it. Note that for observation models that independently model stochasticity at individual pixels (e.g., pixel-wise independent Gaussian, Bernoulli, etc.), the matrix $H_{p_x}(g(h(x)))$ is diagonal and thus efficient to compute. In particular, for independent standard Gaussian observations (i.e., $p(x; g(z)) = \mathcal{N}(g(z), I)$), $H_{p_x}(g(h(x))) = -I$, the identity matrix of size D .

5. Impact of covariance structure of q

The deterministic approximation to the β -VAE objective allows us to study the relationship between the decoder Jacobian and the covariance of the variational distribution. Lets consider the specific case of a standard Gaussian prior, $p(z) = \mathcal{N}(0, I)$, and Gaussian variational family, $q_\phi(z|x) = \mathcal{N}(h_\phi(x), \Sigma_{z|x})$. The objective in (4) reduces to

$$\log p_\theta(x|h_\phi(x)) + \frac{1}{2}\text{tr}(H_{f_x}(h_\phi(x))\Sigma_{z|x}) - \frac{\beta}{2}(\|h_\phi(x)\|^2 + \text{tr}(\Sigma_{z|x}) - \log|\Sigma_{z|x}| - d), \quad (8)$$

where $|\Sigma_{z|x}|$ denotes the absolute value of the determinant of $\Sigma_{z|x}$. We consider the case when the mean of the variational distribution is parameterized by a neural network but the covariance is not amortized. The objective (8) is concave in $\Sigma_{z|x}$, and maximizing it w.r.t. $\Sigma_{z|x}$ gives

$$\begin{aligned} \frac{1}{2}H_{f_x}(h_\phi(x)) - \frac{\beta}{2}(I - \Sigma_{z|x}^{-1}) &= 0 \\ \implies \Sigma_{z|x} &= \left(I - \frac{1}{\beta}H_{f_x}(h_\phi(x))\right)^{-1} \end{aligned} \quad (9)$$

Substituting Eq. (9) into (8), we get the maximization objective of

$$\max \log p_\theta(x|h_\phi(x)) - \frac{\beta}{2}\|h_\phi(x)\|^2 - \frac{\beta}{2} \log \left| I - \frac{1}{\beta}H_{f_x}(h_\phi(x)) \right|. \quad (10)$$

Following similar steps for the objective (7) gives the optimal covariance of

$$\Sigma_{z|x} = \left(I - \frac{1}{\beta}J_g(h(x))^\top H_{p_x}(g(h(x)))J_g(h(x)) \right)^{-1}, \quad (11)$$

and the maximization objective of

$$\log p_\theta(x|h_\phi(x)) - \frac{\beta}{2}\|h_\phi(x)\|^2 - \frac{\beta}{2} \log \left| I - \frac{1}{\beta}J_g(h_\phi(x))^\top H_{p_x}(g(h_\phi(x)))J_g(h_\phi(x)) \right| \quad (12)$$

This analysis yields an interpretation of the β -VAE objective as a deterministic reconstruction term ($p_\theta(x|h_\phi(x))$), a regularizer on the encodings ($\|h_\phi(x)\|^2$), and a regularizer involving the Jacobian of the decoder mapping.

5.1. Regularization effect of the covariance

The relation between the optimal $\Sigma_{z|x}$ and the Jacobian of the decoder $J_g(z)$ in Eq. (11) uncovers how the choice of the structure of $\Sigma_{z|x}$ influences the Jacobian of the decoder. For example, having a diagonal structure on the covariance matrix (mean-field variational distribution) should drive the matrix $(J_g(h(x))^\top H_{p_x}(g(h(x)))J_g(h(x)))$ towards a diagonal matrix. As mentioned earlier, for observation models that independently model the individual pixels given the decoding distribution parameters $g(z)$, the matrix $H_{p_x}(g(z))$ is diagonal which implies that the matrix $|H_{p_x}(g(h(x)))|_{\text{abs}}^{1/2}J_g(h(x))$ is encouraged to have orthogonal columns, where $|H_{p_x}(g(h(x)))|_{\text{abs}}$ denotes taking elementwise absolute values of the matrix $H_{p_x}(g(h(x)))$.

For standard Gaussian observations, we have $H_{p_x}(g(z)) = -I$ for all z , implying that the decoder Jacobian $J_g(h(x))$ will be driven towards having orthogonal columns. For Bernoulli distributed pixel observations, the i 'th diagonal element of matrix $H_{p_x}(g(h(x)))$ is given by $-1/(1 - x_i - [g(h(x))]_i)^2$, where $x_i \in \{0, 1\}$ is the i 'th pixel and $[g(h(x))]_i$ is the predicted probability of i 'th pixel being 1. When the decoding network has the capacity to fit the training data well, i.e., $(1 - x_i - [g(h(x))]_i)^2$ is close to 1 for all i , the Jacobian matrix $J_g(h(x))$ will still be driven towards having orthogonal columns. Orthogonality of the Jacobian has been linked with semantic disentanglement and regularizers to encourage orthogonality have been proposed to that end (Ramesh et al., 2019). Our result shows that diagonal covariance of the variational distribution naturally encourages orthogonal Jacobians. As is evident from Eq. (11), more interesting structures on the decoder Jacobian matrix can be encouraged by placing appropriate sparsity patterns on $\Sigma_{z|x}^{-1}$. Often it may be expensive to encode constraints on the decoder Jacobian directly, as that would typically require computing the Jacobian to apply the constraint. Instead, this relationship highlights that we can impose soft constraints on the decoder Jacobian by restricting the covariance structure for the variational family.

5.2. Connection with Riemannian metric

For Gaussian observations, Eq. (11) reduces to $\Sigma_{z|x} = \left(I + \frac{1}{\beta}J_g(h(x))^\top J_g(h(x)) \right)^{-1}$. The decoder, $g : Z \rightarrow \mathbb{R}^D$, outputs the mean of the distribution and its image $M = g(Z) \subset \mathbb{R}^D$ is an embedded differentiable manifold² of dimension d , when $J_g(z)$ is full rank for all z . If we inherit the Euclidean geometric structure from \mathbb{R}^D , re-

²This assumes that the decoder uses differentiable nonlinearities which does not hold for piecewise-linear activations such as relu, leaky-relu. However, in principle, they can always be approximated arbitrarily closely by a smooth function for the sake of this argument.

restricting the inner product to the tangent spaces $T_x M$ gives us a Riemannian metric on M . The pullback of this metric to Z gives the corresponding Riemannian metric on Z , which is given by a symmetric positive-definite matrix field $G(z) = J_g(z)^\top J_g(z)$. Recent work has used this metric to study the geometry of deep generative models (Shao et al., 2018; Chen et al., 2018a; Arvanitidis et al., 2018; Kuhnel et al., 2018). Given the relation in Eq. (11), we can think of variational distribution covariance as indirectly inducing a metric on the latent space for Gaussian observations as $G(z) = J_g(z)^\top J_g(z) = \beta(\Sigma_{z|x}^{-1} - I)$ for $z = h(x)$. It will be interesting to investigate the metric properties of $\Sigma_{z|x}^{-1}$ in the context of deep generative models, *e.g.*, for geodesic traversals, curvature calculations, etc. (Shao et al., 2018; Chen et al., 2018a).

5.3. Training deterministic approximations of β -VAE

Beyond the theoretical understanding of the implicit regularization in β -VAE, it should also be possible to use the proposed objectives for training the models, particularly when the observation dimensions are conditionally independent as the Hessian $H_{p_x}(\cdot)$ is diagonal. This will also help in investigating the accuracy and utility of our deterministic approximations to the β -VAE objective. We consider the case of Gaussian observation model for simplicity as the matrix $H_{p_x}(g(z)) = -I$, reducing the objective to:

$$\min_{g,h} \frac{1}{2} \|x - g(h(x))\|^2 + \frac{\beta}{2} \|h(x)\|^2 + \frac{\beta}{2} \log \left| I + \frac{1}{\beta} J_g(h(x))^\top J_g(h(x)) \right|. \quad (13)$$

We refer to the objective (13) as Gaussian Regularized AutoEncoder (**GRAE**). However, for large latent dimensions, computing full Jacobian matrices can substantially increase the training cost. To reduce this cost, we propose training with a stochastic approximation of an upper bound of the regularizers in objectives (12) and (13). We use Hadamard inequality for determinants of positive definite matrices A , *i.e.*, $\det(A) \leq \prod_i A_{ii}$ with equality for diagonal A , and approximate the regularizer of (13) as:

$$\log \left| I + \frac{1}{\beta} J_g(h(x))^\top J_g(h(x)) \right| \leq \sum_i \log \left(1 + \frac{1}{\beta} \|[J_g(h(x))]_{:i}\|_2^2 \right) \quad (14)$$

Let p_c be a discrete distribution on the column indices of the Jacobian matrix (*e.g.*, uniform distribution). We can approximate the summation by sampling k column indices $\{c_i\}_1^k$ from the distribution p_c and using $\frac{1}{k} \sum_{i=1}^k \frac{1}{p_c(c_i)} \log \left(1 + \frac{1}{\beta} \|[J_g(h(x))]_{:c_i}\|_2^2 \right)$. Lacking any more information, we simply use a uniform distribution for

p_c , and $k = 1$ in all our experiments. We refer to the objective with this more tractable regularizer as **GRAE** $_{\approx}$. Similar steps can be followed for the regularizer in (12). The upper bound in (14) should be tight for mean-field variational distributions by virtue of the relation in Eq. (11). For such cases (*i.e.*, sparse precision matrices for variational distributions, including, diagonal), we can impose an additional orthogonal penalty on the corresponding Jacobian columns (in line with Eq. (11)) by penalizing their normalized dot products. This can also be implemented in stochastic form by sampling a pair of Jacobian columns.

It is possible to relate the regularizer in (13) with earlier work on regularized autoencoders. We can use the Taylor series for $\log(1+x)$ around $x=0$ to approximate the regularizer in (13), which converges when the singular values of $J_g(h(x))$ are smaller than $\sqrt{\beta}$. Restricting the Taylor series to just first order term yields the simplified regularizer of $\frac{1}{2} \|J_g(h(x))\|_F^2$ (note that this approximation is crude: β cancels out and does not influence the decoder regularizer in any way). Similar approximations can be obtained for (12) which will yield the regularizer of $\|J_g(h(x))[-H_{p_x}(g(h(x)))]^{1/2}\|_F^2$. This simplified regularizer of $\frac{1}{2} \|J_g(h(x))\|_F^2$ weighted by a free hyperparameter has been used in a few earlier works with autoencoders, including recently in (Ghosh et al., 2020), where it was motivated in a rather heuristic manner for encouraging smoothness. On the other hand, we are specifically motivated by uncovering the regularization implicit in β -VAE and the regularizer in (13) emerges as part of this analysis.

6. Related Work

Our work resembles existing work on marginalizing out noise to obtain deterministic regularizers. Maaten et al. (2013) proposed deterministic regularizers obtained with marginalizing feature noise belonging to exponential family distributions. Chen et al. (2012) marginalize noise in denoising autoencoders to learn robust representations. Marginalization of dropout noise has also been explored in the context of linear models as well as deep neural networks (Srivastava, 2013; Wang & Manning, 2013; Srivastava et al., 2014; Poole et al., 2014). Recently, Ghosh et al. (2020) argued for replacing stochasticity in VAEs with deterministic regularizers resulting in deterministic regularized autoencoder objectives. However, the regularizers considered there were motivated from a heuristic perspective of encouraging smoothness in the decoding model. Another recent work (Kumar et al., 2020) uses injective probability flow to derive autoencoding objectives with Jacobian-based regularizers. Different from these works, our goal in Sec. 4 is to uncover the regularizers implicit in the β -VAE objective by marginalizing out noise coming from the variational distribution, and solving for its optimal covariance, yielding the regularizers in (4), (7), (12) and (13).




Figure 2. MNIST: Precision matrices $\Sigma_{z|x}^{-1}$ of the variational distributions (left) and the corresponding $J_g(h(x))^\top J_g(h(x))$ (right), for different values of β and latent dimensionality d . Block size b for the covariance was taken to be 2 or 3. More plots are shown in Appendix E.

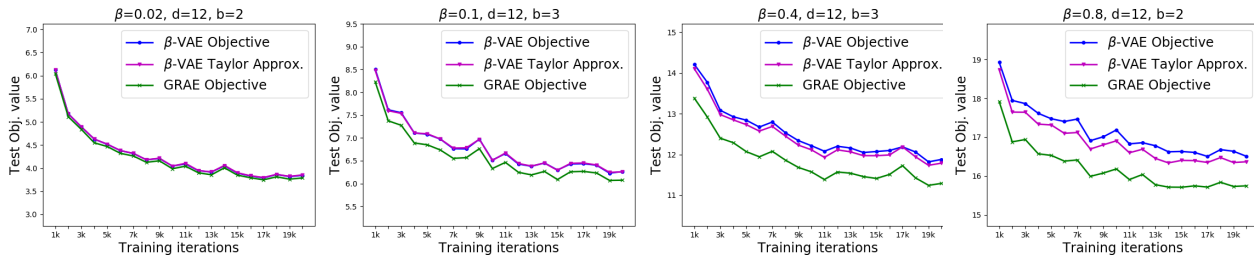


Figure 3. MNIST: Comparison of objective values of β -VAE, its Taylor approximation (Eq. (7)), and GRAE (Eq. (13)) for different values of β , and different values of block size b for covariance matrix of variational distributions. Latent dimensionality d is fixed to 12. More plots can be viewed in Appendix F.

Recent work by Rolinek et al. (2019) considered the case of VAEs with a Gaussian prior, Gaussian variational distribution with diagonal covariance and Gaussian observations, and showed that diagonal posterior covariance encourages the decoder Jacobian to have orthogonal columns. In our work, we explicitly characterize the regularization effect of variational distribution covariance, obtaining insights into the effect of arbitrary sparsity structure of the variational distribution’s precision matrix $\Sigma_{z|x}^{-1}$ on the decoder Jacobian (Eq. (11)). This leads to regularized autoencoder objectives with specific regularizers (e.g., Eq. (12)). Our analysis also generalizes to non-Gaussian observation models. Park et al. (2019) address the amortization gap and limited expressivity of diagonal variational posteriors in VAEs by using Laplace approximation around the *mode of the posterior*. On the other hand, our analysis starts with the Taylor approximation of the conditional $p(x|z)$ around the *mean of the variational posterior* and then further diverges from (Park et al., 2019) towards different goals. Stühmer et al. (2020) recently studied structured non-Gaussian priors for unsupervised disentanglement, motivated by rotational non-uniqueness issues associated with isotropic Gaussian prior, however there will still be mixing transforms that leave these priors invariant. Burgess et al. (2018) discuss how disentangled representations may emerge when using a mean field variational family, but do not present a rigorous argument connected to uniqueness. Recently, Khemakhem et al. (2020) provided uniqueness results for conditional VAEs with *conditionally* factorized priors.


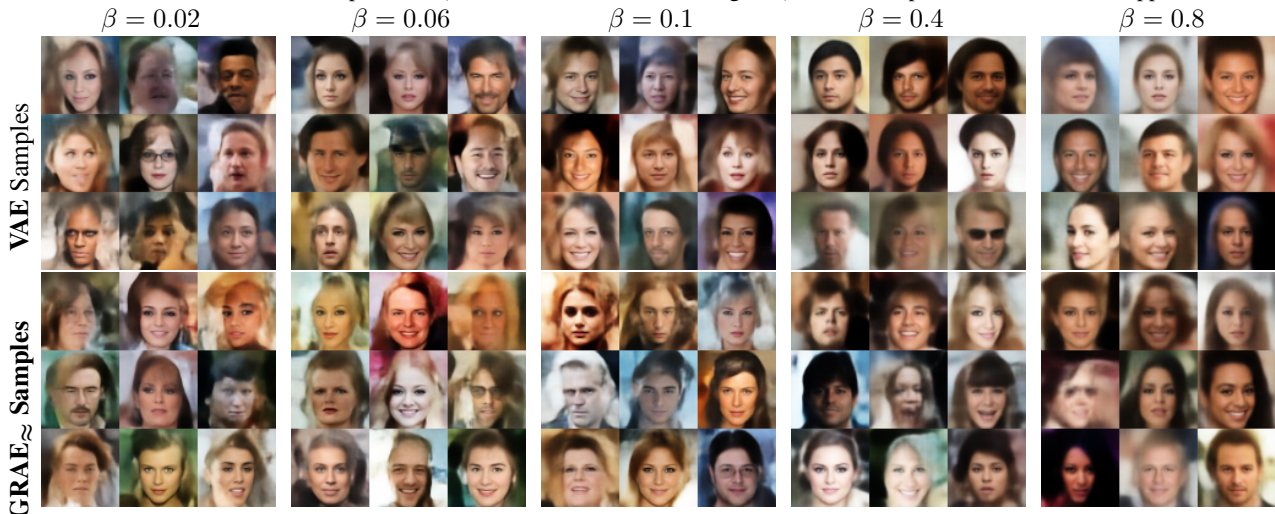


Figure 4. FID scores for CelebA generated samples using β -VAE and GRAE_{\approx} : averaged over 5 runs for each β . All standard deviations are less than 1. Samples from the two models stay close in terms of FID scores, further substantiating the validity of our derived deterministic approximation.

7. Experiments

We conduct experiments on MNIST (Lecun et al., 1998) and CelebA (Liu et al., 2015) to test the closeness of the proposed deterministic regularized objectives to the original β -VAE objective and the relation between the covariance of the variational distribution and the Jacobian of the decoder. Standard train-test splits are used in all experiments for both datasets. We use Gaussian distributions for the prior, variational distribution, and observation model in all our experiments for both β -VAE and the proposed deterministic objectives (which reduces to the GRAE objective of Eq. (13))

Table 1. CelebA samples using the standard normal prior for β -VAE and GRAE_{\approx} (Eq. (13) with the regularizer approximated by stochastic approximation of the upper bound of Eq. (14)) for different values of β . Samples from models trained using the GRAE_{\approx} objective have similar smoothness/blurriness as samples from β -VAE models for a wide range of β . More samples can be viewed in Appendix G.



for the Gaussian case). Note that for the isotropic Gaussian observation model we use, multiplication by β for the KL term in β -VAE is equivalent to assuming an isotropic decoding covariance of $\frac{1}{\beta}I$ in a VAE. We use 64×64 cropped images for CelebA faces as in several earlier works. We use 5 layer CNN architectures for both the encoder and decoder, with elu activations (Clevert et al., 2016) in all hidden layers. We use elu to test the approximation quality of GRAE objective when Eq. (6) is not exact. For piecewise-linear activations such as relu, Eq. (6) is exact and the approximation is better compared to elu. The same architecture is used for both β -VAE and GRAE, except an additional fully-connected (fc) output layer in the encoder for β -VAE that produces the standard deviations of the amortized mean field variational distributions. For non-factorized Gaussian variational distribution, we mainly work with block-diagonal covariance matrices, which restricts the corresponding precision matrices to also be block-diagonal. This is implemented by an additional fc layer in the encoder outputting the entries of the appropriate Cholesky factors $C_{z|x}$ such that $\Sigma_{z|x} = C_{z|x}C_{z|x}^{\top}$ is block diagonal with desired block sizes. All MNIST and CelebA models are trained for 20k and 50k iterations respectively, using the Adam optimizer (Kingma & Ba, 2015).

Relation between covariance of the variational distribution and Jacobian structure. Based on the relation in Eq. (11), we expect to see similar block diagonal structures in $\Sigma_{z|x}^{-1}$ (or $\Sigma_{z|x}$) and $J_g(h(x))^{\top}J_g(h(x))$. Figure 2 visualizes these two matrices for a few test points for different values of β , latent dimensionality d , and block size b . In many cases, these two matrices stay close to each other, providing evidence for our theoretical observation.

Comparing the β -VAE and GRAE objectives. We empirically compare the values of the β -VAE (minimization) objective (2), its Taylor approximation (7), and the GRAE objective (13). We train the same architecture as above on the β -VAE objective while using a single sample approximation for the expectation term $\mathbb{E}_{q(z|x)} \log p(x|z)$. We evaluate all three objectives on a fixed held-out test set of 5000 examples on checkpoints from the stochastic β -VAE model. Fig. 3 shows a comparison of the three objectives for MNIST. The objectives remain close over a wide range of β values, with the gaps among them decreasing for smaller β values. The gap between the β -VAE Taylor approximation and GRAE objectives increases more rapidly with β than the gap between β -VAE and its Taylor approximation. This gap is analytically given by $\frac{\beta}{2} [\text{tr}(I + \frac{1}{\beta}J_g(h(x))^{\top}J_g(h(x))\Sigma_{z|x}) - \log |(I + \frac{1}{\beta}J_g(h(x))^{\top}J_g(h(x))\Sigma_{z|x}) - d|]$, which vanishes when Eq. (11) holds exactly. Although these gaps increase with larger β where the variance of the variational distribution increases, we observe that the trend of the three objectives is still strikingly similar (ignoring constant offsets). We provide more comparison plots in Appendix F.

Training models using the GRAE_{\approx} objective. Next, we test how the GRAE objective behaves when used for training. As discussed earlier, we use a stochastic approximation of Eq. (14) for the decoder regularizer to train the model (referred as GRAE_{\approx} objective). We work with CelebA 64×64 images and use standard Gaussian prior with latent dimensionality of 128, Gaussian observations, and factored Gaussian variational distribution. Table 1 shows the decoded samples from standard normal prior for both GRAE_{\approx} and β -VAE trained with different values of β . It can be noted that the generated samples follow a similar trend qual-

itatively for both β -VAE and GRAE_{\approx} , with large β values resulting in more blurry samples. This provides an evidence for validity of GRAE_{\approx} as a training objective that behaves similar to β -VAE. More samples for a wider range of β values are shown in Appendix G. As an attempt to quantify the similarity, we also compare the FID scores of the generated samples for both models as a function of β in Fig. 4. The FID scores stay close for both models, particularly for $\beta \geq 0.01$ where sample quality is better (*i.e.*, lower FID scores). This further empirically substantiates the faithfulness of our derived deterministic approximation. As future work, it would be interesting to explore other tractable approximations to the GRAE objective and different weights for the encoder regularizer and decoder regularizer (which are both currently fixed to $\beta/2$ in the GRAE objective).

8. Conclusion

We studied regularization effects in β -VAE from two perspectives: (i) analyzing the role that the choice of variational family can play in influencing the uniqueness properties of the solutions, and (ii) characterizing the regularization effect of the variational distribution on the decoding model by integrating out the noise from the variational distribution in an approximate objective. The second perspective leads to deterministic regularized autoencoding objectives. Empirical results confirm that the deterministic objectives are close to the original β -VAE in terms of objective value and sample quality, further validating the analysis. Our analysis helps in connecting β -VAEs and regularized autoencoders in a more principled manner, and we hope that this will motivate novel regularizers for improved sample quality and diversity in autoencoders.

Acknowledgements. We would like to thank Alex Alemi and Andrey Zhmoginov for providing helpful comments on the manuscript. We also thank Matt Hoffman and Kevin Murphy for insightful discussions.

References

- Adamčík, M. The information geometry of bregman divergences and some applications in multi-expert reasoning. *Entropy*, 16(12), 2014.
- Alemi, A. A., Poole, B., Fischer, I., Dillon, J. V., Saurous, R. A., and Murphy, K. Fixing a broken elbo. In *International Conference on Machine Learning*, 2018.
- Arvanitidis, G., Hansen, L. K., and Hauberg, S. Latent space oddity: on the curvature of deep generative models. In *ICLR*, 2018.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8): 1798–1828, 2013.
- Burgess, C. P., Higgins, I., Pal, A., Matthey, L., Watters, N., Desjardins, G., and Lerchner, A. Understanding disentangling in β -vae. *arXiv preprint arXiv:1804.03599*, 2018.
- Chen, M., Xu, Z., Weinberger, K., and Sha, F. Marginalized denoising autoencoders for domain adaptation. In *ICML*, 2012.
- Chen, N., Klushyn, A., Kurle, R., Jiang, X., Bayer, J., and van der Smagt, P. Metrics for deep generative models. In *AISTATS*, 2018a.
- Chen, T. Q., Li, X., Grosse, R. B., and Duvenaud, D. K. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems*, pp. 2610–2620, 2018b.
- Clevert, D.-A., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (elus). In *ICLR*, 2016.
- Darmois, G. Analyse des liaisons de probabilité. In *Proc. Int. Stat. Conferences 1947*, pp. 231, 1951.
- Darmois, G. Analyse générale des liaisons stochastiques: etude particulière de l’analyse factorielle linéaire. *Revue de l’Institut international de statistique*, pp. 2–8, 1953.
- Ghosh, P., Sajjadi, M. S. M., Vergari, A., Black, M., and Schölkopf, B. From variational to deterministic autoencoders. In *ICLR*, 2020.
- Gilpin, L. H., Bau, D., Yuan, B. Z., Bajwa, A., Specter, M., and Kagal, L. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*, pp. 80–89. IEEE, 2018.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.
- Hyvärinen, A. and Pajunen, P. Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 12(3):429–439, 1999.
- Khemakhem, I., Kingma, D. P., and Hyvärinen, A. Variational autoencoders and nonlinear ica: A unifying framework. In *AISTATS*, 2020.
- Kim, H. and Mnih, A. Disentangling by factorising. In *ICML*, 2018.

- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *ICLR*, 2014.
- Kuhnel, L., Fletcher, T., Joshi, S., and Sommer, S. Latent space non-linear statistics. *arXiv preprint arXiv:1805.07632*, 2018.
- Kumar, A., Sattigeri, P., and Balakrishnan, A. Variational inference of disentangled latent concepts from unlabeled observations. In *ICLR*, 2018.
- Kumar, A., Poole, B., and Murphy, K. Regularized autoencoders via relaxed injective probability flow. In *AISTATS*, 2020.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pp. 2278–2324, 1998.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3730–3738, 2015.
- Locatello, F., Bauer, S., Lucic, M., Gelly, S., Schölkopf, B., and Bachem, O. Challenging common assumptions in the unsupervised learning of disentangled representations. In *ICML*, 2019.
- Lukacs, E., King, E. P., et al. A property of the normal distribution. *The Annals of Mathematical Statistics*, 25(2):389–394, 1954.
- Maaten, L., Chen, M., Tyree, S., and Weinberger, K. Learning with marginalized corrupted features. In *International Conference on Machine Learning*, pp. 410–418, 2013.
- Mathieu, E., Rainforth, T., Siddharth, N., and Teh, Y. W. Disentangling disentanglement in variational autoencoders. In *International Conference on Machine Learning*, pp. 4402–4412, 2019.
- Park, Y., Kim, C., and Kim, G. Variational laplace autoencoders. In *International Conference on Machine Learning*, 2019.
- Peters, J., Janzing, D., and Schölkopf, B. *Elements of causal inference: foundations and learning algorithms*. MIT press, 2017.
- Poole, B., Sohl-Dickstein, J., and Ganguli, S. Analyzing noise in autoencoders and deep networks, 2014.
- Ramesh, A., Choi, Y., and LeCun, Y. A spectral regularizer for unsupervised disentanglement. In *International Conference on Machine Learning*, 2019.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, 2014.
- Rolinek, M., Zietlow, D., and Martius, G. Variational autoencoders pursue pca directions (by accident). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 12406–12415, 2019.
- Sason, I. and Verdú, S. f -divergence inequalities. *IEEE Transactions on Information Theory*, 62(11):5973–6006, 2016.
- Shao, H., Kumar, A., and Thomas Fletcher, P. The riemannian geometry of deep generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 315–323, 2018.
- Shu, R., Bui, H. H., Zhao, S., Kochenderfer, M. J., and Ermon, S. Amortized inference regularization. In *Advances in Neural Information Processing Systems*, pp. 4393–4402, 2018.
- Skitovitch, V. On a property of the normal distribution. *DAN SSSR*, 89:217–219, 1953.
- Srivastava, N. Improving neural networks with dropout. *University of Toronto*, 182:566, 2013.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- Stühmer, J., Turner, R. E., and Nowozin, S. Independent subspace analysis for unsupervised learning of disentangled representations. In *AISTATS*, 2020.
- Tipping, M. E. and Bishop, C. M. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- van Steenkiste, S., Locatello, F., Schmidhuber, J., and Bachem, O. Are disentangled representations helpful for abstract visual reasoning? In *Advances in Neural Information Processing Systems*, 2019.
- Wang, S. and Manning, C. Fast dropout training. In *international conference on machine learning*, pp. 118–126, 2013.
- Zhu, J., Chen, N., and Xing, E. P. Bayesian inference with posterior regularization and applications to infinite latent svms. *Journal of Machine Learning Research*, 2014.

Appendix.

Here we provide more details, in particular:

Appendix A: non-uniqueness results for latent-variable models.

Appendix B: proof of Theorem 1

Appendix C: experiments to empirically verify the uniqueness results

Appendix D: architecture details for encoder and decoder

Appendix E: plots visualizing the approximate posterior precision matrix $\Sigma_{z|x}^{-1}$ and $J_g(h(x))^\top J_g(h(x))$ to verify the relation in Eq. (11)

Appendix F: plots comparing objective values of β -VAE and GRAE

Appendix G: comparison of samples from β -VAE and GRAE_{\approx}

A. Non-uniqueness of latent-variable models

Non-uniqueness for factorized priors. For models with continuous factorized priors, $p(z) = \prod_i p_i(z_i)$, [Locatello et al. \(2019\)](#) present a construction of transformations r that mix latents but leave the marginal data distribution unchanged. We summarize it here for completeness and also extend it to the case of non-factored priors which was conjectured but not covered in ([Locatello et al., 2019](#)). The construction first transforms all marginal distributions $p_i(z_i)$ to uniform distributions by the mapping $u_i = F_i(z_i)$ where $F_i(\cdot)$ is the c.d.f. corresponding to the density $p_i(\cdot)$. This independent uniform random vector u is then transformed to independent standard normal random vector by the mapping $n_i = \psi^{-1}(z_i)$ where $\psi(\cdot)$ is the c.d.f. of the standard normal. As the r.v. n is independent standard normal distributed and thus spherically symmetric, for any orthogonal transformation U , the transformed r.v. $z' = r(z) = (F^{-1} \circ \psi \circ U \circ \psi^{-1} \circ F)(z)$ has same density as the prior $p(z)$ while completely mixing the latent variables z (i.e., the transform has Jacobian $J_r(z)$ with nonzero off-diagonals, with $|\det J_r(z)| = 1$). The marginal observation density $p(x)$ remains unchanged as well, since $p'(x) = \int p(z)p(x|r(z))dz = \int p(r^{-1}(z'))p(x|z')dz' = \int p(z')p(x|z')dz' = p(x)$.

Non-uniqueness for correlated priors. Although not discussed in [Locatello et al. \(2019\)](#), it is also possible to generalize this construction to *non-factorized* priors by replacing transformations $u_i = F_i(z_i)$ in the first step with $u_i = F_{i|1...(i-1)}(z_i, z_1, \dots, z_{i-1})$ (and its inverse in the last step), where $F_{i|1...(i-1)}$ is the c.d.f. of the conditional distribution of z_i given *earlier* latents (with arbitrary ordering on the latents) ([Darmois, 1951](#); [Hyvärinen & Pajunen, 1999](#)).

B. Uniqueness in β -VAE

We provide the proof of Theorem 1 which states that any β -VAE solution obtained by transforming the generative model of the optimal solution while keeping the prior and the marginal log-likelihood same, will result in decreasing the objective value.

Theorem 1. *Let R be the class of mixing transforms that leave the prior invariant, and p_r be the generative model obtained by transforming the latents of the model p by $r \in R$. Let $L(q^*, p^*) = \max_{q \in \mathcal{Q}, p \in \mathcal{P}} L(q, p)$. Let $\tilde{\mathcal{Q}}$ be the completion of \mathcal{Q} by R , i.e., $\tilde{\mathcal{Q}} = \{r\#q : q \in \mathcal{Q}, r \in R\} \cup \mathcal{Q}$. If \mathcal{Q} is s.t. $q \in \mathcal{Q} \Rightarrow r\#q \notin \mathcal{Q}$ for all $r \in R$, and $\arg \max_{q \in \tilde{\mathcal{Q}}} L(q, p^*)$ is unique, then $\max_{q \in \mathcal{Q}} L(q, p_r^*) < L(q^*, p^*)$.*

Proof. The objective after transforming a model p with $r \in R$ yields:

$$\begin{aligned} L(q, p_r) &:= \mathbb{E}_{q(z|x)} \log p_{x|z}(x|r(z)) - \beta \text{KL}(q(z|x)||p_z(z)) \\ &= \int \log \left(p_{x|z}(x|r(z)) \left(\frac{p_z(z)}{q(z|x)} \right)^\beta \right) q(z|x) dz \\ &= \int \log \left(p_{x|z}(x|z') \left(\frac{p_z(r^{-1}(z'))}{q(r^{-1}(z')|x)} \right)^\beta \right) q(r^{-1}(z')|x) dz' \\ &= \int \log \left(p_{x|z}(x|z) \left(\frac{p_z(z)}{q(r^{-1}(z)|x)} \right)^\beta \right) q(r^{-1}(z)|x) dz \\ &= \mathbb{E}_{q_r} \log p_{x|z}(x|z) - \beta \text{KL}(q_r||p_z) = L(q_r, p). \end{aligned} \tag{15}$$

where we used the shorthand q_r to denote the transformed variational density $r\#q_{z|x}$. This also implies that $L(q, p) = L(q_{r^{-1}}, p_r)$. Hence, the optimum value of the β -VAE objective $L(q^*, p^*) = L(q_{r^{-1}}^*, p_r^*) \forall r \in R$. Choosing \mathcal{Q} such that $q \in \mathcal{Q} \Rightarrow r\#q \notin \mathcal{Q} \forall r \in R$ will ensure that $q_{r^{-1}}^* \notin \mathcal{Q}$. Note that optimizing the β -VAE objective over $\tilde{\mathcal{Q}} = \{r\#q : q \in \mathcal{Q}, r \in R\}$ does not change the optimal value, i.e., $L(q^*, p^*) = \max_{q \in \mathcal{Q}, p \in \mathcal{P}} L(q, p) = \max_{q \in \tilde{\mathcal{Q}}, p \in \mathcal{P}} L(q, p)$. The statement of the theorem can now be proven by contradiction. Suppose $\exists q' \in \mathcal{Q}$ s.t. $L(q', p_r^*) = L(q_r', p_r^*) \geq L(q^*, p^*)$. However, this will contradict the assumption that $q^* = \arg \max_{q \in \tilde{\mathcal{Q}}} L(q, p^*)$ and that it is unique. Hence $\forall q' \in \mathcal{Q}, L(q', p_r^*) < L(q^*, p^*)$. \square

Corollary 2. *If the set $\tilde{\mathcal{Q}}$ defined in Theorem 1 is convex, then $\arg \max_{q \in \tilde{\mathcal{Q}}} L(q, p^*)$ is unique. When all other assumptions of Theorem 1 are met, we have $\max_{q \in \mathcal{Q}} L(q, p_r^*) < L(q^*, p^*)$.*

Proof. As the β -VAE objective is strongly concave in q (first term is linear in q and the KL-divergence term is strongly convex in q ([Adamčík, 2014](#); [Sason & Verdu,](#)

2016)), optimizing it over a convex set $\tilde{\mathcal{Q}}$ for a fixed p will have a unique maximizer. Hence, the result of Theorem 1 holds when all other assumptions are met. \square

C. Uniqueness in β -VAE: empirical results

Here we provide more details on simulations we perform to verify that appropriately choosing the variational family \mathcal{Q} imparts uniqueness to the β -VAE (or plain VI) solution with respect to the corresponding class of transformations R . We consider the case of \mathcal{Q} being mean-field distributions and R being rotations. We work with two dimensional data so that rotations can be represented by a scalar (angle of planar rotation) which can be plotted conveniently.

We generate 4k samples in two dimensions by the generative model $x = \nu(Wz) + \epsilon$, with $z \sim \mathcal{N}(0, I_{2 \times 2})$, $\epsilon \sim \mathcal{N}(0, .05^2 I_{2 \times 2})$. Entries of $W \in \mathbb{R}^{2 \times 2}$ are sampled from $\mathcal{N}(0, I_{2 \times 2})$ and normalized such that columns have different ℓ_2 norms (2 and 1, respectively). The function ν acts elementwise on the vector and is taken to be either Identity (reducing it to Probabilistic PCA) or a nonlinearity (tanh, sigmoid, elu). A VAE is learned on this data, with the decoder of the same form as the generative model, *i.e.*, $x = \nu(Az) + b$, and variational posterior of the form $q(z|x) = \mathcal{N}(Cx + d, \Sigma)$ with A, b, C, d and Σ learned.

We also do unamortized variational inference with the variational distribution $q_i(z) = \mathcal{N}(\mu_i, \Sigma_i)$ for the i 'th data point ($\mu_i \in \mathbb{R}^2, \Sigma_i \in S_+^2$). For the case of mean-field variational inference, we learn only diagonal elements of Σ or Σ_i that appear in the objective (ELBO), leaving the off-diagonal elements free.

We do multiple runs by randomizing over W, ϵ, z for the true generative model of data and randomizing over the initializations of the VAE parameters. For each run, we measure the rotation angle between the true latents and the inferred posterior means. Rotation angle is computed after ignoring permutation and reflection of axes so the angle ranges in $(-45, 45]$ degrees with -45 degree rotation being same as 45 degree rotation. The rotation angle is estimated by solving the problem $\min_O \|Z - OM\|_F^2$ where O is an orthogonal matrix, $Z \in \mathbb{R}^{2 \times 4000}$ is the matrix of original sampled latents used for generating the data, $M \in \mathbb{R}^{2 \times 4000}$ is the matrix of inferred latents (the means of $q(z|x_i)$). This is known as Orthogonal Procrustes Problem whose solution is given by $O = UV^T$ where U and V are left and right singular vectors of ZM^T . We then extract axes permutation, axes reflection and the rotation angle from the matrix O .

We train the model parameters by gradient descent (use all data samples for a gradient update). The learning rate is set to 0.1 in the beginning and is reduced in a stepwise manner after every n iterations by a factor of 10, till the learning rate reaches 10^{-5} (*i.e.*, each model is trained for $5n$ iterations).

To train models up to varying degree of optimality, we take n to be $\{100, 500, 1k, 2k, 3k, 4k, 5k, 6k, 7k, 8k, 9k, 10k, 20k\}$. We do 20 random trials for each value of n from 100 to 9k, and 100 random trials for $n = 10k$ and $n = 20k$ (total 420 random trials per experiment).

We also take special care in how we parameterize the models so that the implicit bias of gradient descent does not affect the solution in any particular way. For amortized inference with VAE, we parameterize encoder parameters as $C = O_1 C'$ and $\Sigma = O_2 D' D'^T O_2^T$, where O_1, O_2 are random 2×2 rotation matrices that are fixed at the beginning of training (different for each random trial), and C' and D' are the parameters learned by gradient descent. For unamortized variational inference, we again parameterize $\Sigma_i = O D'_i D'_i{}^T O^T$ with O being a random 2×2 rotation matrix and D'_i being the trainable parameters. We notice that without these reparameterizations, implicit bias of gradient descent can impart some degree of uniqueness to the model, which we particularly noticed in the case of unamortized inference.

Fig. 5 shows the scatter plots of rotation angle vs ELBO for the amortized case (VAE) with $\nu = \text{identity, tanh, elu, sigmoid}$. Having diagonal Σ (mean-field $q(z|x)$) encourages unique solution in terms of inferred latents as the ELBO increases (ignoring the permutations and reflections of latent axes). Fig. 6 shows the scatter plots of rotation angle vs ELBO for the unamortized case, with similar observations regarding uniqueness.

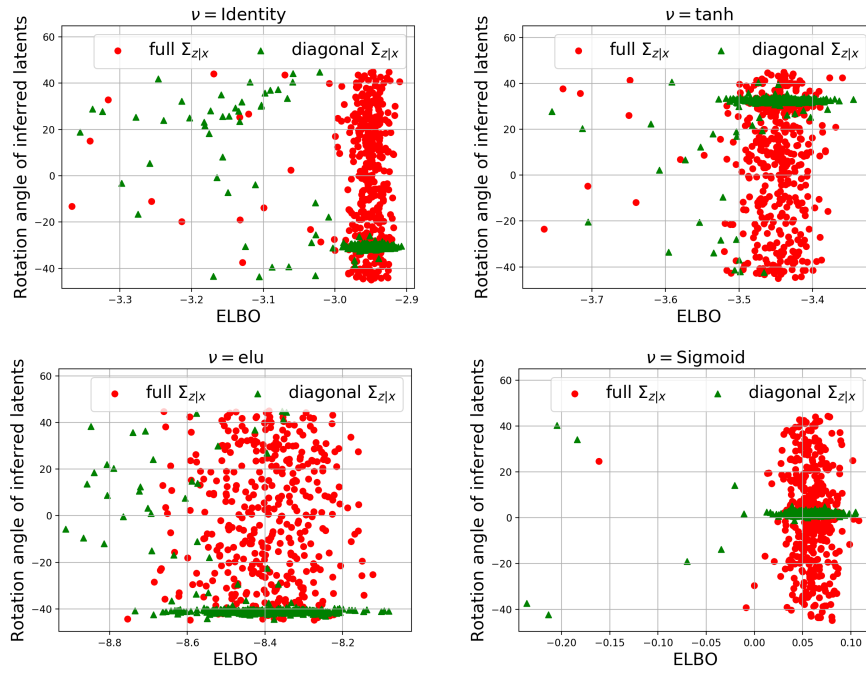


Figure 5. Rotation angle between true latents and inferred latents with VAEs (multiple runs): The true generative model for the data is $x = \nu(Wz) + \epsilon$, with $z \sim \mathcal{N}(0, I_{2 \times 2})$, $\epsilon \sim \mathcal{N}(0, .05^2 I_{2 \times 2})$. Diagonal Σ encourages unique solution in terms of inferred latents as the ELBO increases. Please refer to the text for more details.

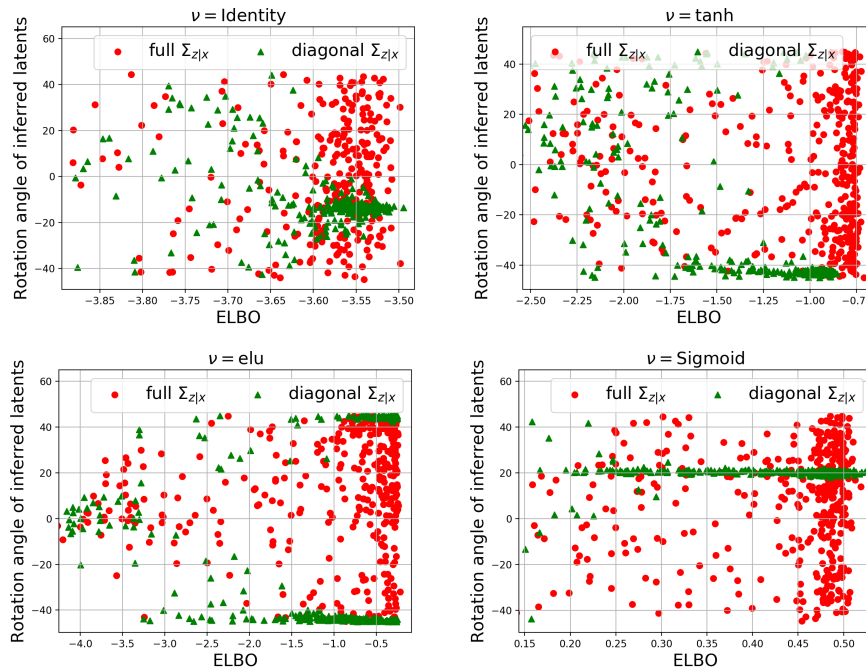


Figure 6. Rotation angle between true latents and inferred latents with unamortized variational inference (multiple runs): The true generative model for the data is $x = \nu(Wz) + \epsilon$, with $z \sim \mathcal{N}(0, I_{2 \times 2})$, $\epsilon \sim \mathcal{N}(0, .05^2 I_{2 \times 2})$. Diagonal Σ encourages unique solution in terms of inferred latents as the ELBO increases. Please refer to the text for more details.


D. Architectures

Here we describe the architecture for our MNIST and CelebA models. We list Conv (convolutional) and ConvT (transposed convolution) layers with their number of filters, kernel size, and stride. Latent dimension for CelebA is fixed to be 128 in all experiments, and the dimension for MNIST experiments is varied in the set $\{8, 12, 16, 20\}$. The architecture is shown for the GRAE model. For β -VAE, the encoder has an extra output through an fully-connected layer for standard deviations of the posterior factors in the case of mean-field posterior, or for the elements of the Cholesky factor in the case of block-diagonal posterior covariance matrices.


MNIST	CelebA
Encoder	
$x \in \mathcal{R}^{28 \times 28}$	$x \in \mathcal{R}^{64 \times 64 \times 3}$
\rightarrow Conv _{64,4,1} \rightarrow BN \rightarrow ELU	\rightarrow Conv _{128,5,1} \rightarrow BN \rightarrow ELU
\rightarrow Conv _{128,4,2} \rightarrow BN \rightarrow ELU	\rightarrow Conv _{256,5,2} \rightarrow BN \rightarrow ELU
\rightarrow Conv _{256,4,2} \rightarrow BN \rightarrow ELU	\rightarrow Conv _{512,5,2} \rightarrow BN \rightarrow ELU
\rightarrow Conv _{512,4,2} \rightarrow BN \rightarrow ELU	\rightarrow Conv _{1024,5,2} \rightarrow BN \rightarrow ELU
\rightarrow Conv _{512,4,1} \rightarrow BN \rightarrow ELU	\rightarrow Conv _{1024,5,2} \rightarrow BN \rightarrow ELU
\rightarrow Flatten \rightarrow FC _{d}	\rightarrow Flatten \rightarrow FC ₁₂₈
Decoder	
$z \in \mathcal{R}^d \rightarrow$ FC _{7\times7\times256}	$z \in \mathcal{R}^{128} \rightarrow$ FC _{16\times16\times512}
\rightarrow BN \rightarrow ELU	\rightarrow BN \rightarrow ELU
\rightarrow ConvT _{512,4,1} \rightarrow BN \rightarrow ELU	\rightarrow ConvT _{1024,5,1} \rightarrow BN \rightarrow ELU
\rightarrow ConvT _{256,4,1} \rightarrow BN \rightarrow ELU	\rightarrow ConvT _{512,5,2} \rightarrow BN \rightarrow ELU
\rightarrow ConvT _{128,4,2} \rightarrow BN \rightarrow ELU	\rightarrow ConvT _{256,5,2} \rightarrow BN \rightarrow ELU
\rightarrow ConvT _{64,4,2} \rightarrow ELU	\rightarrow ConvT _{128,5,2} \rightarrow ELU
\rightarrow Conv _{1,4,1} \rightarrow Sigmoid	\rightarrow Conv _{3,5,1} \rightarrow Sigmoid

E. Relation between covariance of the variational distribution and Jacobian structure


We provide more plots visualizing the variational distribution precision matrix $\Sigma_{z|x}^{-1}$ and $J_g(h(x))^\top J_g(h(x))$ in Figure 7. Based on the relation in Eq. (11), we expect to see similar block diagonal structures in $\Sigma_{z|x}^{-1}$ (or $\Sigma_{z|x}$) and $J_g(h(x))^\top J_g(h(x))$.




(a) $\beta = 0.2, d = 8, b = 2$




(b) $\beta = 0.4, d = 8, b = 2$




(c) $\beta = 0.6, d = 12, b = 2$




(d) $\beta = 1, d = 8, b = 2$



(e) $\beta = 0.4, d = 12, b = 3$



(f) $\beta = 0.6, d = 8, b = 3$




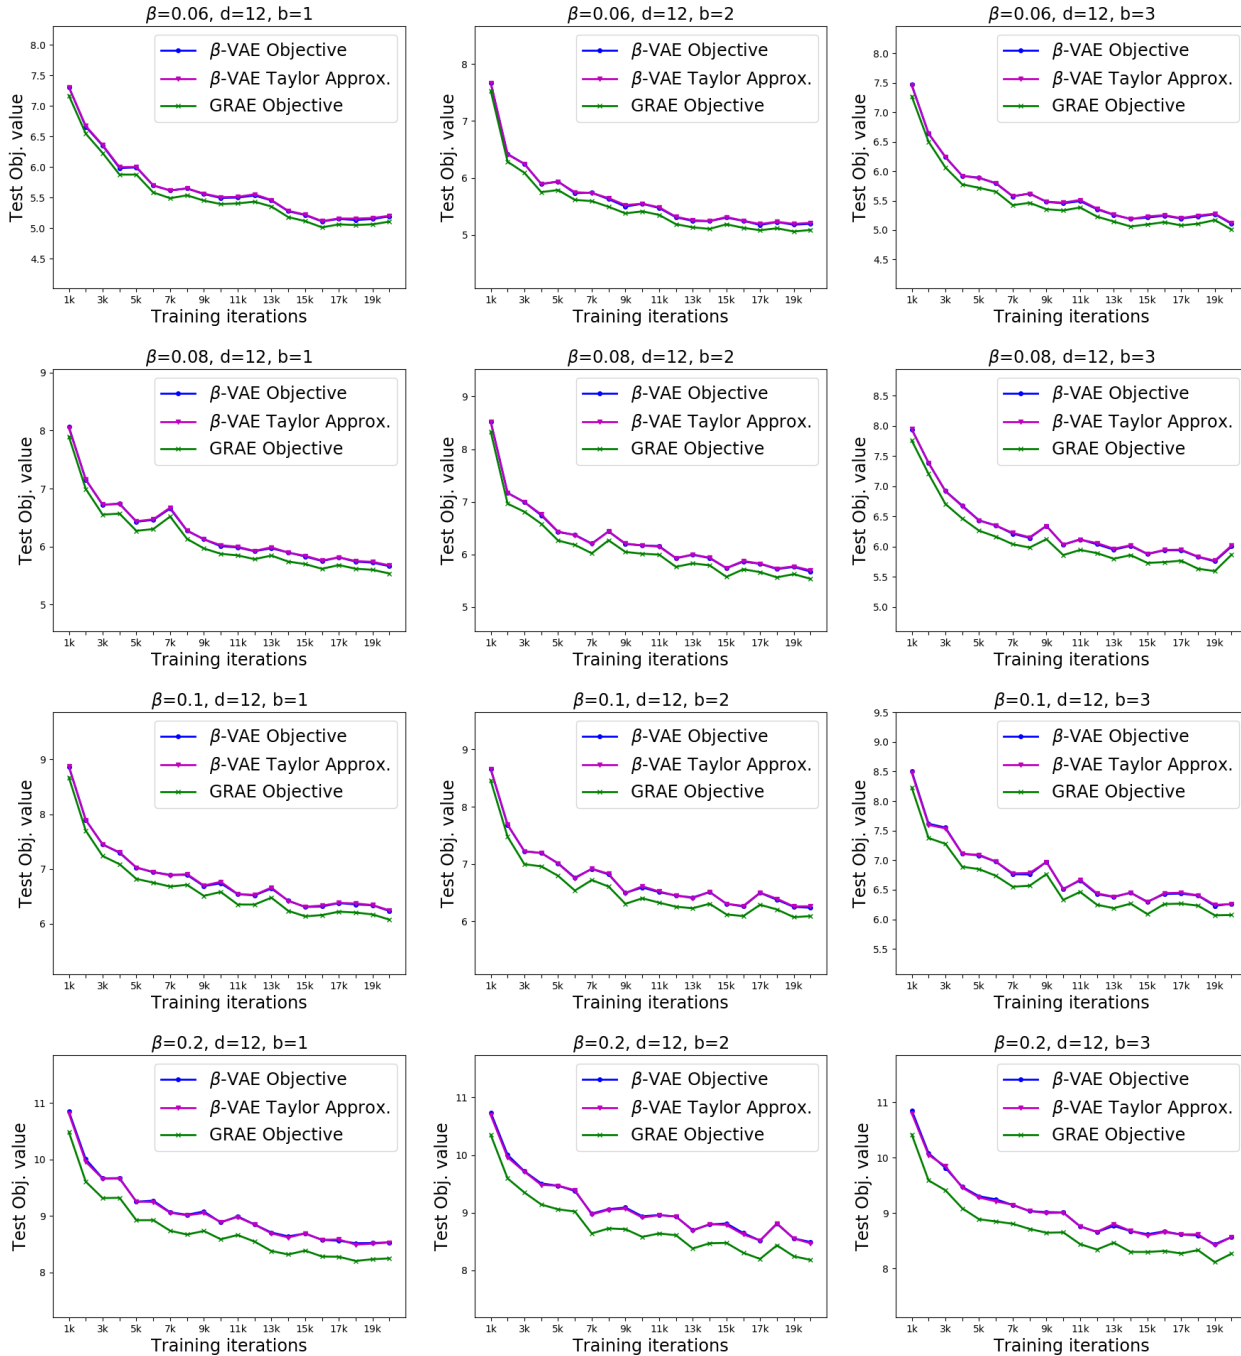
(g) $\beta = 0.8, d = 20, b = 3$

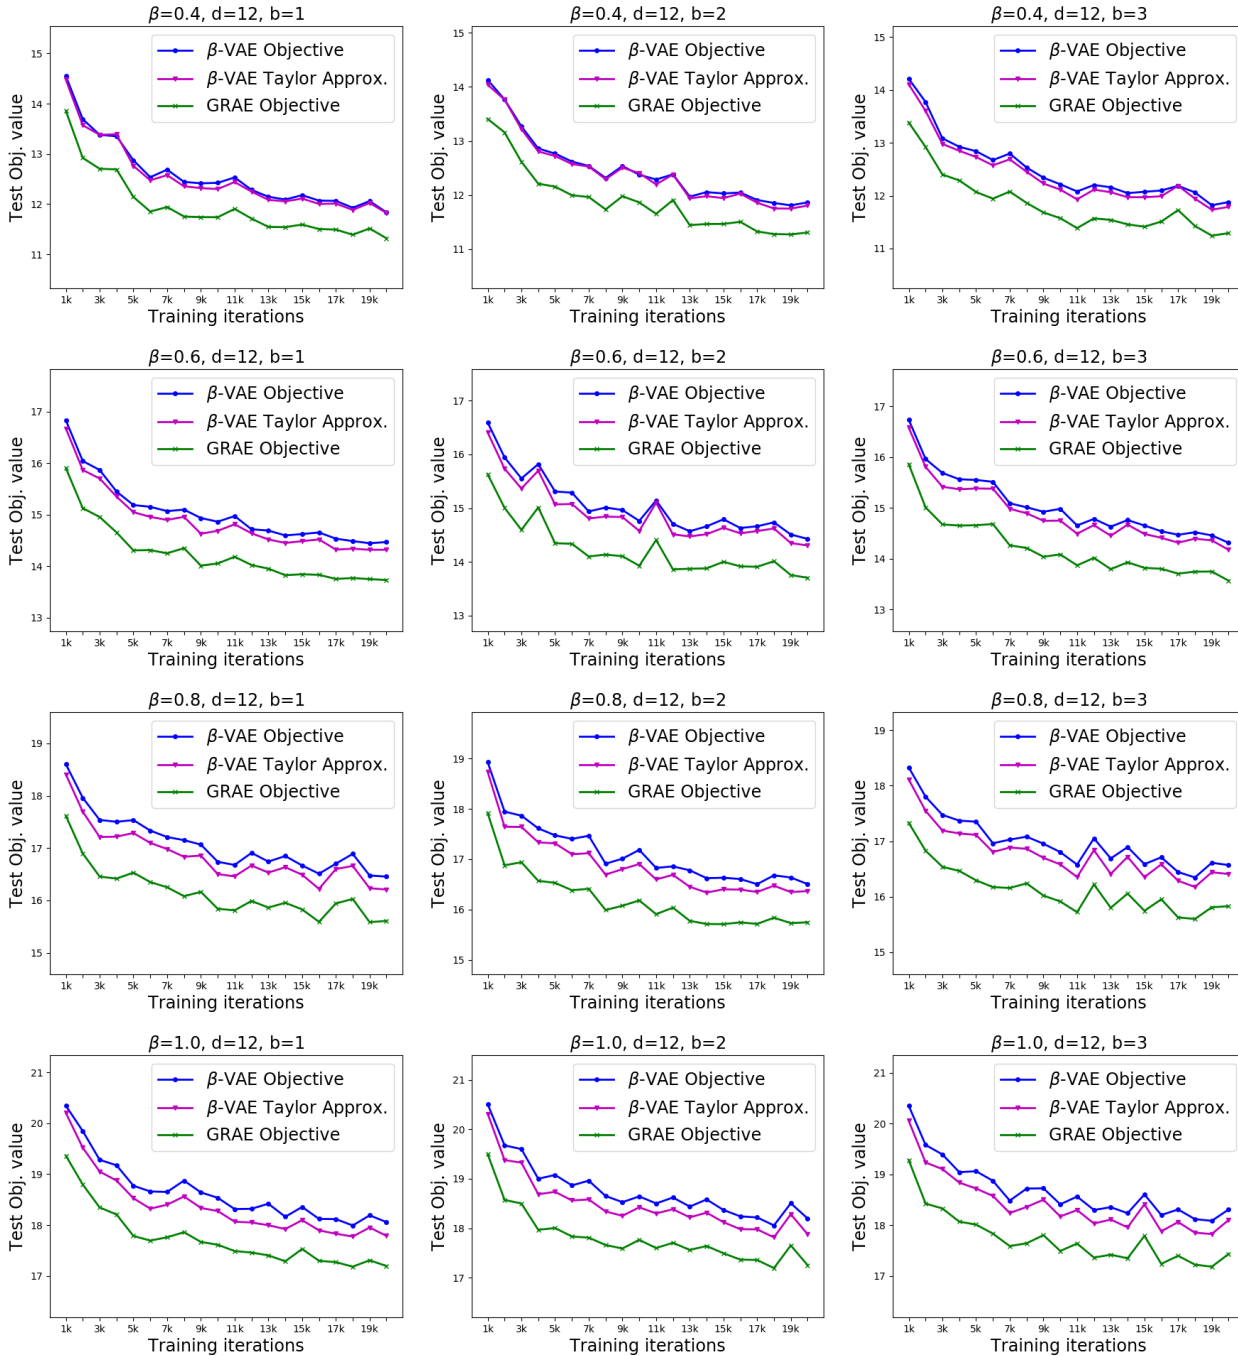
Figure 7. MNIST: Precision matrices $\Sigma_{z|x}^{-1}$ of variational distributions (left) and the corresponding $J_g(h(x))^\top J_g(h(x))$ (right), for different values of β and latent dimensionality d . Block size b for the posterior covariance was taken to be 2 (top row) or 3 (bottom row).

F. Comparison of objective values

We show more plots comparing objective values of β -VAE, its Taylor series approximation, and the GRAE objective for various values of β on MNIST. The latent dimension d is fixed to 12, and the block-size b (for block-diagonal posterior covariance) is varied over $\{1, 2, 3\}$. The model is trained using the β -VAE objective and is evaluated on all objectives using a held-out test set of 5000 examples. As β increases, the gap between the β -VAE objective and its Taylor approximation increases. The gap between the Taylor approximation and the GRAE objective also increases with β . This gap is analytically given by $\frac{\beta}{2} [\text{tr}(I + \frac{1}{\beta} J_g(h(x))^T J_g(h(x)) \Sigma_{z|x}) - \log |(I + \frac{1}{\beta} J_g(h(x))^T J_g(h(x)) \Sigma_{z|x}) - d|]$.







G. Samples from VAE and GRAE

We visualize random CelebA samples using the standard normal prior for β -VAE and Gaussian-RAE for different values of β below. Note that for training the GRAE model, we actually minimize the stochastic approximation of the upper bound of Eq. (17) for the decoder regularizer. Hence we will refer to it as **GRAE_≈**. Samples from models trained using the GRAE objective have similar smoothness/blurriness as samples from β -VAE models, particularly for small β values.



On Implicit Regularization in β -VAEs



On Implicit Regularization in β -VAEs



On Implicit Regularization in β -VAEs



On Implicit Regularization in β -VAEs

