

A Closer Look at Disentangling in β -VAE

Harshvardhan Sikka^{*†}, Weishun Zhong^{*‡||}, Jun Yin[§], Cengiz Pehlevan[†]

[†]School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA

[‡]Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA

^{||}Center for Brain Sciences, Harvard University, Cambridge, MA, USA

[§]Department of Physics, Harvard University, Cambridge, MA, USA

Abstract—In many data analysis tasks, it is beneficial to learn representations where each dimension is statistically independent and thus disentangled from the others. If data generating factors are also statistically independent, disentangled representations can be formed by Bayesian inference of latent variables. We examine a generalization of the Variational Autoencoder (VAE), β -VAE, for learning such representations using variational inference. β -VAE enforces conditional independence of its bottleneck neurons controlled by its hyperparameter β . This condition is in general not compatible with the statistical independence of latents. By providing analytical and numerical arguments, we show that this incompatibility leads to a non-monotonic inference performance in β -VAE with a finite optimal β .

Index Terms—Autoencoder, Bayesian Inference

I. INTRODUCTION

Although data sampled from the natural world appear to be high-dimensional, their variations can usually be explained using a much smaller number of latent factors. Both biological and artificial information processing systems exploit such structure and learn explicit representations that are faithful to data generative factors, known commonly as disentangled representations [1]. For example, sparse coding, an influential model of the primary visual cortex, proposes that the visual cortex neurons are coding for latent variables of natural scenes: oriented edges [2]. A very popular method of extracting latent variables is by using the bottleneck neurons of deep autoencoders [3], [4]. In this paper, we examine unsupervised learning of disentangled representations in the context of variational inference and a generalization of the Variational Autoencoder (VAE) [5], β -VAE, developed specifically for disentangled representation learning [6].

We will adopt a probabilistic framework for latent-variable modeling of data [7], where a generative model $p_{\theta}(\mathbf{x}, \mathbf{z})$ for data \mathbf{x} and latent variables \mathbf{z} is assumed:

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

Here θ denotes the parameters of our model, $p_{\theta}(\mathbf{x}|\mathbf{z})$ models the stochastic process that generates the data given the latent variables, and $p(\mathbf{z})$ is the prior on the latent variables. An interpretable and common choice for $p(\mathbf{z})$, and the subject of our paper, is a factorized distribution $p(\mathbf{z}) = \prod_{i=1}^k p_i(z_i)$, which implies statistical independence. Examples of models with independent priors include popular methods such

as Independent Component Analysis [8], [9] and Principal Component Analysis [10].

While a common definition of learning disentangled representations has yet to be agreed upon [1], [5], [11], [12], extracting statistically independent latent factors is a natural choice [1], [8] and is the definition we will adopt. Such a representation is efficient in that it carries no redundant information [13], and at the same time sufficient information to generate the data.

In our probabilistic framework, the model posterior distribution $p_{\theta}(\mathbf{z}|\mathbf{x})$ allows inference of true latent variables. In principle, this could be used to form disentangled representations. However, model posterior is often intractable [7], and variational methods are used to estimate it.

We focus on a state-of-the-art variational inference method for learning disentangled representations, β -VAE [6]. The β -VAE training objective includes a hyperparameter, β , encapsulating the original VAE [5] as a special case with choice $\beta = 1$. When β is larger than unity, *conditional* independence of the learned representations at the bottleneck layer are enforced, corresponding to a conditional independence assumption on data generating latent variables, i.e. $p(\mathbf{z}|\mathbf{x}) = \prod_i p_i(z_i|\mathbf{x})$ [6]. However, as pointed above, a more natural assumption on latents is full statistical independence. Further, statistically independent latents are in general not conditionally independent. Given the popularity of VAEs in representation learning, it is important to understand the role of the β hyperparameter in learning disentangled (statistically independent) latent variables.

Our main contributions are as follows:

- 1) We provide general results about variational Bayesian inference in β -VAE. Specifically, we prove that the β -VAE objective is non-increasing with increasing β , leading to worse reconstruction performance but more conditionally independent representations. Further, we argue that latent variable inference performance generally tends to be non-monotonic in β .
- 2) We introduce an analytically tractable model for β -VAE, specializing to statistically independent latent generative factors. We analytically calculate the optimality conditions for this model, and numerically find that there is an optimal β for the best inference of latent variables.
- 3) We test our insights from the general theorems and the analytically tractable model using a realistic β -

*equal contribution, alphabetical order.

VAE architecture, using a synthetic MNIST dataset. Simulations agree well with our theory.

The rest of this paper is organized as follows. In Section II, we provide a review of variational inference and β -VAE. In Section III, we prove several theorems about variational inference in the context of β -VAE. In Section IV, we introduce our analytical results. In Section V, we test our insights from the general theorems and the tractable models using a β -VAE architecture on a synthetic MNIST dataset. Finally, in Section VI we discuss our results and present our conclusions.

II. VARIATIONAL INFERENCE AND β -VAE

Inference of latent variables in probabilistic models is often an intractable calculation [5], [7]. Variational methods instead optimize over a set of tractable distributions, $q_\phi(\mathbf{z}|\mathbf{x})$, that best approximates $p_\theta(\mathbf{z}|\mathbf{x})$. We will refer to $q_\phi(\mathbf{z}|\mathbf{x})$ as the inference model. The difference between the two distributions can be quantified using the Kullback-Leibler (KL) divergence, which we call Model Inference Error (MIE):

$$\text{MIE} \equiv \mathbb{E}_{p(\mathbf{x})} [D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x}))]. \quad (2)$$

We distinguish between MIE and the True Inference Error (TIE),

$$\text{TIE} \equiv \mathbb{E}_{p(\mathbf{x})} [D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_{g-t}(\mathbf{z}|\mathbf{x}))], \quad (3)$$

which can only be known when one has access to the underlying ‘ground-truth’ data generative process and the ground-truth posterior, $p_{g-t}(\mathbf{z}|\mathbf{x})$.

VAEs fit the parameters of the probabilistic model and the variational distribution simultaneously. A key identity in doing so is [14]

$$\begin{aligned} \ln p_\theta(\mathbf{x}) - D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x})) \\ = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\ln p_\theta(\mathbf{x}|\mathbf{z})] - D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})). \end{aligned} \quad (4)$$

Model fitting is done by maximizing the data log-likelihood, $\ln p_\theta(\mathbf{x})$, under model parameters. Because the KL divergence is non-negative, the right hand side of (4) serves as a lower bound for $\ln p_\theta(\mathbf{x})$ and is called the Evidence Lower Bound (ELBO)

$$\begin{aligned} \text{ELBO}(\boldsymbol{\theta}, \boldsymbol{\phi}) \\ \equiv \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z})] - D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})). \end{aligned} \quad (5)$$

VAE parameterizes the distributions $p_\theta(\mathbf{x}|\mathbf{z})$ and $q_\phi(\mathbf{z}|\mathbf{x})$ with neural networks, and maximizes ELBO as a proxy for maximizing the data likelihood.

The neural network realization of the $p_\theta(\mathbf{x}|\mathbf{z})$ is referred to as a decoder [5]. Once the VAE is trained, the decoder can be used as to generate new samples from the model data distribution [5], [15]. The term $\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z})]$ measures the reconstruction performance of the generative model. We will refer to it as the reconstruction objective.

The neural network realization of the inference model is referred to as an encoder [5]. Its outputs constitute a bottleneck layer and represent inferred latent variables. Note that the MIE

calculated from this representation appears on the left hand side of (4).

β -VAE is an extension of the traditional VAE, where an extra, adjustable hyperparameter β is placed in the training objective:

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}; \beta) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z})] - \beta D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})). \quad (6)$$

Specifically, when $\beta = 1$, the β -VAE is equivalent to VAE and $\text{ELBO}(\boldsymbol{\theta}, \boldsymbol{\phi}) = \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}; 1)$.

Higher values of β emphasizes the KL divergence between the inference model $q_\phi(\mathbf{z}|\mathbf{x})$ and the independent prior $p(\mathbf{z})$ in the objective (6). Smaller values of the KL divergence favor a conditionally independent inference model. This can be used to learn disentangled representations of conditionally independent latent variables, whose probability distributions factorize when conditioned on data [6], [16].

However, as alluded to in our introduction, in many cases of interest and application [17]–[19], latent variables are conditionally dependent while being independent [8], [10]. We will encounter an analytically tractable case in Section IV. In such cases, it is not clear if a β different than 1 helps learning a disentangled representation which extracts statistically independent latent factors. Our goal in the remaining of this paper is to examine this case analytically and numerically.

For convenience, we also attach a table of terms and corresponding mathematical expressions used throughout the paper (Table I).

Term	Mathematical Expression
Prior	$p(\mathbf{z})$
Model Posterior	$p_\theta(\mathbf{z} \mathbf{x})$
Ground-Truth Posterior	$p_{g-t}(\mathbf{z} \mathbf{x})$
Inference Model	$q_\phi(\mathbf{z} \mathbf{x})$
Data Log-Likelihood	$\log p_\theta(\mathbf{x})$
Reconstruction Objective	$\mathbb{E}_{q_\phi(\mathbf{z} \mathbf{x})} [\log p_\theta(\mathbf{x} \mathbf{z})]$
Conditional Independence Loss	$D_{KL}(q_\phi(\mathbf{z} \mathbf{x}) p(\mathbf{z}))$
MIE	$\mathbb{E}_{p(\mathbf{x})} [D_{KL}(q_\phi(\mathbf{z} \mathbf{x}) p_\theta(\mathbf{z} \mathbf{x}))]$
TIE	$\mathbb{E}_{p(\mathbf{x})} [D_{KL}(q_\phi(\mathbf{z} \mathbf{x}) p_{g-t}(\mathbf{z} \mathbf{x}))]$
Evidence Lower Bound (ELBO)	$\mathbb{E}_{q_\phi(\mathbf{z} \mathbf{x})} [\log p_\theta(\mathbf{x} \mathbf{z})] - D_{KL}(q_\phi(\mathbf{z} \mathbf{x}) p(\mathbf{z}))$

Table I. Table of terms and corresponding mathematical expressions.

III. HOW β AFFECTS MODEL PERFORMANCE AND INFERENCE OF LATENT VARIABLES

In this section, we provide general statements on the effect of the β parameter on the representation learning and the generative functions of β -VAE. We do this by proving propositions about how various terms in the identity (4) change as a function of β . Our first two propositions imply that

increasing β worsens the quality of reconstructed samples while improving conditional disentangling. While these points have been shown in simulations [6], [16], here we provide analytical statements. Our last proposition gives a handle on understanding behavior of MIE through ELBO.

In the following, we will denote optimal parameters of a β -VAE that maximizes the objective (6) by θ^* and ϕ^* . They are given as a solution to

$$\frac{\partial \mathcal{L}}{\partial \theta} = \mathbf{0}, \quad \frac{\partial \mathcal{L}}{\partial \phi} = \mathbf{0}. \quad (7)$$

We denote the value of the optimal objective by

$$\mathcal{L}^*(\beta) \equiv \mathcal{L}(\theta^*(\beta), \phi^*(\beta), \beta), \quad (8)$$

and the value of ELBO at the optimal point by

$$\text{ELBO}^*(\beta) \equiv \text{ELBO}(\theta^*(\beta), \phi^*(\beta)). \quad (9)$$

Our first proposition concerns the behavior of $\mathcal{L}^*(\beta)$ as a function of β .

Proposition 1. *The optimal value of the β -VAE objective, $\mathcal{L}^*(\beta)$, is non-increasing with increasing β :*

$$\frac{\partial \mathcal{L}^*(\beta)}{\partial \beta} = -D_{KL}(q_{\phi^*}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) \leq 0. \quad (10)$$

Proof. Follows from an application of the chain rule, the optimality conditions (7), and the nonnegativity of the KL-divergence:

$$\begin{aligned} \frac{\partial \mathcal{L}^*}{\partial \beta} &= \left(\frac{\partial \mathcal{L}}{\partial \theta} \cdot \frac{\partial \theta}{\partial \beta} + \frac{\partial \mathcal{L}}{\partial \phi} \cdot \frac{\partial \phi}{\partial \beta} + \frac{\partial \mathcal{L}}{\partial \beta} \right) \Big|_{\theta=\theta^*, \phi=\phi^*} \\ &= -D_{KL}(q_{\phi^*}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) \leq 0. \end{aligned} \quad (11)$$

□

The next proposition shows how the two terms in \mathcal{L}^* change with β .

Proposition 2. *The KL divergence between the inference model and the prior is non-increasing with increasing β :*

$$\frac{d}{d\beta} D_{KL}(q_{\phi^*}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) \leq 0. \quad (12)$$

Together with Proposition (1), this implies that

$$\frac{d \mathbb{E}_{q_{\phi^*}(\mathbf{z}|\mathbf{x})} [\log p_{\theta^*}(\mathbf{x}|\mathbf{z})]}{d\beta} \leq 0 \quad (13)$$

Proof. See Appendix A. □

The next proposition is about the behavior of ELBO^* .

Proposition 3. *ELBO* is maximized at $\beta = 1$.*

Proof. Note that by definition

$$\mathcal{L} = \text{ELBO} + (1 - \beta) D_{KL}(q_{\phi}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})). \quad (14)$$

By evaluating (14) at $\theta = \theta^*$ and $\phi = \phi^*$, and the chain rule, we get:

$$\begin{aligned} \frac{d \text{ELBO}^*(\beta)}{d\beta} &= \frac{d}{d\beta} [\mathcal{L}^* - (1 - \beta) D_{KL}(q_{\phi^*}(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))] \\ &= (\beta - 1) \frac{d}{d\beta} D_{KL}(q_{\phi^*}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})). \end{aligned} \quad (15)$$

The proposition follows from this result and (12). □

For simplicity of notation, we presented most of our formulas and propositions for a single data point. All our results generalize to the case where one averages over the data distribution $p(\mathbf{x})$, or a finite training set.

Inference of latent variables, measured by MIE, is affected by β as well. In the $\beta \rightarrow \infty$ limit the inference model becomes more and more conditionally independent, deviating from the model posterior. Is the behavior monotonic? While MIE is not explicitly calculable, we can get a hint of its behavior by rearranging (4), and evaluating it at the optimal β -VAE parameters:

$$\text{MIE}(\beta) = \mathbb{E}_{p(\mathbf{x})} [\ln p_{\theta^*(\beta)}(\mathbf{x}) - \text{ELBO}^*(\beta)]. \quad (16)$$

As reconstruction performance worsens with β , it is reasonable to expect that the data likelihood decreases with β . Because ELBO is non-monotonic with a maximum, even if the data log-likelihood was monotonic with β , we can expect a non-monotonic behavior of MIE with an optimal value. In the next section, we will see two specific examples of this.

IV. ANALYTICAL RESULTS

In this section we demonstrate our general theory for two different analytically tractable cases.

A. β -VAE with a fixed decoder does not lead to better disentangling

A simple case is when the decoder of the β -VAE is not trained. In our notation, this amounts to θ being fixed. Then the β -VAE objective (6) only trains the encoder network and the inference model, $q_{\phi}(\mathbf{z}|\mathbf{x})$. We can deduce the behavior of MIE as a function of β from (16). The data likelihood, $p_{\theta}(\mathbf{x})$, does not change as a function of training. ELBO^* is maximized at $\beta = 1$ from Proposition 3, which can be seen to apply to fixed θ . This means MIE is minimum at $\beta = 1$. In this case, $\beta = 1$, or the original VAE is best at learning the true latent variables.

B. Optimal β values in an analytically tractable model

Next, we present a tractable VAE model, in which we can explicitly calculate the β -dependence in every term in eq. (4).

We assume that our data \mathbf{x} comes from mixing of ground truth latent variables (or sources) $\mathbf{s} \in \mathbb{R}^k$ through a mixing matrix $\mathbf{A} \in \mathbb{R}^{N \times k}$, then corrupted by noise $\boldsymbol{\eta} \in \mathbb{R}^N$,

$$\mathbf{x} = \mathbf{A}\mathbf{s} + \boldsymbol{\eta}. \quad (17)$$

We assume $\mathbf{s} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$, $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N)$. The data distribution is found to be,

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{0}, \mathbf{A}\mathbf{A}^\top + \mathbf{I}_N) \equiv \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{x}}). \quad (18)$$

We denote a $d \times d$ identity matrix as \mathbf{I}_d . In this model we can calculate the ground-truth posterior exactly (see Appendix B-C for details):

$$\begin{aligned} p_{\mathbf{s}|\mathbf{x}}(\mathbf{s}|\mathbf{x}) &= \mathcal{N}(\boldsymbol{\mu}_{\mathbf{s}|\mathbf{x}}, \boldsymbol{\Sigma}_{\mathbf{s}|\mathbf{x}}), \\ \text{with } \boldsymbol{\mu}_{\mathbf{s}|\mathbf{x}} &= (\mathbf{A}^\top \mathbf{A} + \mathbf{I}_k)^{-1} \mathbf{A}^\top \mathbf{x} \\ \text{and } \boldsymbol{\Sigma}_{\mathbf{s}|\mathbf{x}} &= (\mathbf{A}^\top \mathbf{A} + \mathbf{I}_k)^{-1}. \end{aligned} \quad (19)$$

Note that the covariance matrix of the posterior is non-diagonal. Even though the latent factors are statistically independent, when conditioned on data they are dependent. Therefore, we expect a non-trivial dependence of MIE and TIE on the hyperparameter β .

Our encoder $q_\phi(\mathbf{z}|\mathbf{x})$ contains a fully-connected layer $\{\mathbf{W}^\mu, \mathbf{b}^\mu\}$ with linear activation that codes for the mean $\boldsymbol{\mu}_{\mathbf{z}}$ of the latent variables \mathbf{z} , and a fully-connected layer $\{\mathbf{W}^\sigma, \mathbf{b}^\sigma\}$ with exponential activation that codes for the diagonal part of the covariance matrix $\boldsymbol{\Sigma}_{\mathbf{z}}$. Given an input $\mathbf{x} \in \mathbb{R}^N$, we generate latent variables $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}}, \boldsymbol{\Sigma}_{\mathbf{z}}) \in \mathbb{R}^k$ by

$$\boldsymbol{\mu}_{\mathbf{z}} = \mathbf{W}^\mu \mathbf{x} + \mathbf{b}^\mu, \quad \boldsymbol{\Sigma}_{\mathbf{z}} = \text{diag}(\exp(\mathbf{W}^\sigma \mathbf{x} + \mathbf{b}^\sigma)), \quad (20)$$

where the diag operation maps vectors in \mathbb{R}^k to the diagonal of a diagonals matrix in $\mathbb{R}^{k \times k}$. The exponential nonlinearity in the definition of the covariance matrix acts elementwise and prevents negative covariances.

Our decoder consists of a single fully-connected layer $\{\mathbf{D}, \mathbf{b}^D\}$ with linear activations. We assume the output $\mathbf{y} \in \mathbb{R}^N$ is normally distributed, $\mathbf{y} \sim \mathcal{N}(\mathbf{D}\mathbf{z} + \mathbf{b}^D, \sigma_y^2 \mathbf{I}_N)$, where σ_y^2 is a hyperparameter. Without loss of generality, from now on we choose $\sigma_y^2 = 1$.

The decoder defines $p_\theta(\mathbf{x}|\mathbf{z})$. The full data likelihood can be calculated using the prior $p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$ through $p_\theta(\mathbf{x}) = \int d\mathbf{z} p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z})$. With this setup, our decoder is fully capable of modeling the data generative process (18), by choosing $\mathbf{D} = \mathbf{A}$, $\mathbf{b}^D = \mathbf{0}$ and $\sigma_y^2 = 1$. Any deviation from these parameters will be due to the encoder, or the inference model, deviating from the ground-truth distribution.

In order to solve this model, we integrate out data (i.e., performing $\mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})}[\cdot]$, using eq. (18)) in the β -VAE objective in eq. (5) to arrive at (see Appendix B-A for details)

$$\begin{aligned} \mathcal{L}_\beta &= -\frac{1}{2} \left\{ \text{Tr} \left[(\mathbf{D}\mathbf{W}^\mu - \mathbf{I}_N) \boldsymbol{\Sigma}_{\mathbf{x}} (\mathbf{D}\mathbf{W}^\mu - \mathbf{I}_N)^\top \right] \right. \\ &\quad + \beta \text{Tr} \left[\mathbf{W}^\mu \boldsymbol{\Sigma}_{\mathbf{x}} (\mathbf{W}^\mu)^\top \right] \\ &\quad + \sum_i^k \left([\mathbf{D}^\top \mathbf{D}]_{ii} + \beta \right) e^{\frac{1}{2} [\mathbf{W}^\sigma \boldsymbol{\Sigma}_{\mathbf{x}} (\mathbf{W}^\sigma)^\top]_{ii} + b_i^\sigma} \\ &\quad \left. + (\mathbf{D}\mathbf{b}^\mu + \mathbf{b}^D)^2 + \beta (\mathbf{b}^\mu)^2 - \beta \sum_i^k b_i^\sigma \right\}. \end{aligned} \quad (21)$$

We optimize over the network parameters, which amounts to setting the partial derivative of \mathcal{L}_β with respect to $\{\mathbf{W}^\mu, \mathbf{b}^\mu, \mathbf{W}^\sigma, \mathbf{b}^\sigma, \mathbf{D}, \mathbf{b}^D\}$ to zero. Upon simplifying, we find (see Appendix B-B for details)

$$\mathbf{b}^\mu = \mathbf{b}^D = 0, \quad (22)$$

and the remaining equations are ($a = 1, \dots, N$; $b = 1, \dots, k$):

$$\begin{aligned} 0 &= [(\mathbf{D}^\top (\mathbf{D}\mathbf{W}^\mu - \mathbf{I}_N) + \beta \mathbf{W}^\mu) \boldsymbol{\Sigma}_{\mathbf{x}}]_{ab}, \\ 0 &= [(\mathbf{D}\mathbf{W}^\mu - \mathbf{I}_N) \boldsymbol{\Sigma}_{\mathbf{x}} (\mathbf{W}^\mu)^\top]_{ab} + D_{ab} e^{\frac{1}{2} [\mathbf{W}^\sigma \boldsymbol{\Sigma}_{\mathbf{x}} (\mathbf{W}^\sigma)^\top]_{bb} + b_b^\sigma}, \\ 0 &= ([\mathbf{D}^\top \mathbf{D}]_{aa} + \beta) e^{\frac{1}{2} [\mathbf{W}^\sigma \boldsymbol{\Sigma}_{\mathbf{x}}]_{ab} W_{ab}^\sigma + b_a^\sigma} [\mathbf{W}^\sigma \boldsymbol{\Sigma}_{\mathbf{x}}]_{ab}, \\ 0 &= ([\mathbf{D}^\top \mathbf{D}]_{aa} + \beta) e^{\frac{1}{2} [\mathbf{W}^\sigma \boldsymbol{\Sigma}_{\mathbf{x}} (\mathbf{W}^\sigma)^\top]_{aa} + b_a^\sigma} - \beta. \end{aligned} \quad (23)$$

We can calculate the model posterior distribution $p_\theta(\mathbf{z}|\mathbf{x})$ at the network optimum, eqs. (22) and (23). Using Bayes' rule we find (see Appendix B-D)

$$\begin{aligned} p_\theta(\mathbf{z}|\mathbf{x}) &= \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}|\mathbf{x}}, \boldsymbol{\Sigma}_{\mathbf{z}|\mathbf{x}}), \\ \text{with } \boldsymbol{\mu}_{\mathbf{z}|\mathbf{x}} &= (\mathbf{D}^\top \mathbf{D} + \mathbf{I}_k)^{-1} \mathbf{D}^\top \mathbf{x} \\ \text{and } \boldsymbol{\Sigma}_{\mathbf{z}|\mathbf{x}} &= (\mathbf{D}^\top \mathbf{D} + \mathbf{I}_k)^{-1}. \end{aligned} \quad (24)$$

Note that when $\mathbf{A} = \mathbf{D}$, eq. (24) reduces to eq. (19), and the model posterior matches with the ground-truth posterior. We are interested in the inference errors MIE and TIE, eqs. (2) and (3). Upon integrating out the data, we find (see Appendix B-D for derivations)

$$\begin{aligned} \text{MIE/TIE} &= \frac{1}{2} \left\{ \sum_i^k E_{ii}^{-1} \exp \left[\frac{1}{2} \left(\mathbf{W}^\sigma \boldsymbol{\Sigma}_{\mathbf{x}} (\mathbf{W}^\sigma)^\top \right)_{ii} + b_i^\sigma \right] - \sum_i^k b_i^\sigma \right. \\ &\quad \left. + \text{Tr} \log \mathbf{E} + \text{Tr} \left[(\mathbf{F} - \mathbf{W}^\mu)^\top \mathbf{E}^{-1} (\mathbf{F} - \mathbf{W}^\mu) \boldsymbol{\Sigma}_{\mathbf{x}} \right] - k \right\}, \end{aligned} \quad (25)$$

where for MIE

$$\mathbf{E} = (\mathbf{D}^\top \mathbf{D} + \mathbf{I}_k)^{-1}, \quad \mathbf{F} = (\mathbf{D}^\top \mathbf{D} + \mathbf{I}_k)^{-1} \mathbf{D}^\top, \quad (26)$$

and for TIE

$$\mathbf{E} = (\mathbf{A}^\top \mathbf{A} + \mathbf{I}_k)^{-1}, \quad \mathbf{F} = (\mathbf{A}^\top \mathbf{A} + \mathbf{I}_k)^{-1} \mathbf{A}^\top. \quad (27)$$

As an example, we numerically solve eq. (23) for $N = 128, k = 2$, $A_{ij} = 1/2(1 + \delta_{ij})$, and use the optimal network parameters $\{\mathbf{W}^{\mu*}, \mathbf{b}^{\mu*}, \mathbf{W}^{\sigma*}, \mathbf{b}^{\sigma*}, \mathbf{D}^*, \mathbf{b}^{D*}\}$ to calculate ELBO (Fig. 1(A)) and inference errors (Fig. 1(B)). We see that ELBO is maximized at $\beta = 1$, while the inference error is not monotonically decreasing and has a minimum at some β . This confirms the theory we outlined earlier. Also, data log-likelihood is monotonically decreasing with β (not shown). We further calculate individual terms in the ELBO: the reconstruction objective (Fig. 1(C)), $\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z})]$, and the conditional Independence Loss (Fig. 1(D)), $D_{KL}(q_\phi(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z}))$. Indeed both terms are monotonically decreasing with β , confirming our propositions.

V. NUMERICAL SIMULATIONS

In this section, we examine a deep, nonlinear β -VAE on a synthetic dataset. The dataset is generated according to eq. (17) by mixing 10 MNIST digits, arranged as columns of a matrix \mathbf{A} , with ground truth sources, $\mathbf{s} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$, and subsequently adding a noise $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N)$. Other experimental setups and corresponding datasets that were explored are included in Appendix C (Fig. C.1).

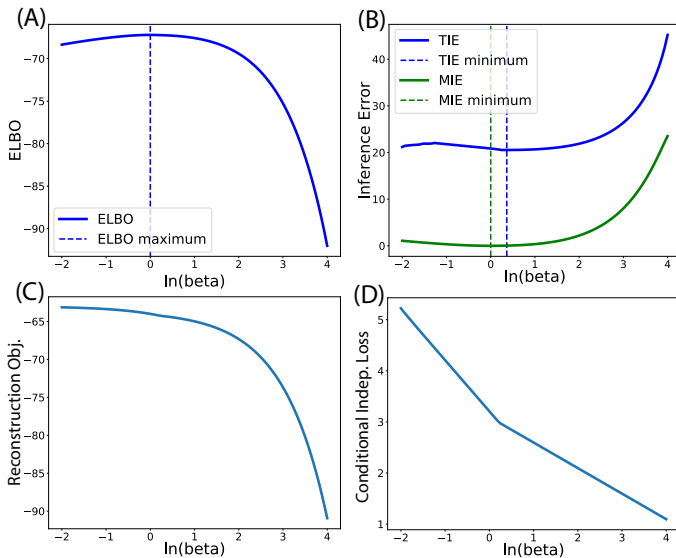


Fig. 1. β -dependence of various quantities at the optimal parameter configuration of β -VAE. (A) ELBO as a function of β . (B) MIE/TIE as a function of β . (C) Reconstruction objective as a function of β . (D) Conditional Independence Loss as a function of β . In these plots, we averaged the plotted quantities over the data distribution.

The encoder, $q_\phi(\mathbf{z}|\mathbf{x})$, consists of three feed-forward fully-connected layers with tanh activations, ending in two separate output layers encoding the mean of the latent variables \mathbf{z} , $\mu_{\mathbf{z}}$, and the variance, $\Sigma_{\mathbf{z}}$. These are each parameterized by k encoding units. The decoder, $p_\theta(\mathbf{x}|\mathbf{z})$, consists of three feed-forward fully-connected layers with tanh activation functions, which takes its input from the encoder, and outputs the reconstructed image. Model details are included in Appendix C.

After training, we calculate individual terms in the β -VAE objective and demonstrate their dependence on β . These terms correspond to the Reconstruction Objective, (Fig. 2(C)), and the conditional Independence Loss, (Fig. 2(D)). As we observed in the analytically tractable case, and predicted by our theory, these terms are decreasing with β . Correspondingly, after being maximized around $\beta = 1$ the entire ELBO term decreases with β (Fig. 2(A)). We also calculate the TIE for the β -VAE at various β , which follows a non-monotonic trend and has an optimal β (Fig. 2(B)).

VI. DISCUSSION AND CONCLUSION

In this paper, we examined the learning of disentangled representations by extracting statistically independent latent variables in β -VAE. We proved general theorems on variational Bayesian inference in the context of β -VAE and introduced an analytically tractable β -VAE model. We also performed experiments on synthetic datasets to test our insights from the general theorems and the tractable model, and found good agreements.

β -VAE enforces conditional independence of its units at the bottleneck layer. This preference is not compatible with




Fig. 2. Values for error terms across 100 random initializations of the network. Solid line represents the average. Dashed lines around the solid line represent the minimum and maximum values, and vertical dashed line represent the extremum. (A) ELBO as a function of β . (B) TIE as a function of β . (C) Reconstruction Objective as a function of β . (D) Conditional Independence Loss as a function of β .

independence of latent variables, and therefore may lead to an optimal value of β for latent variable inference.

There are other perspectives on what constitutes a disentangled representation not addressed in this paper [1], [16], including definitions not statistical in nature, instead taking into account the manifold structure and symmetry transformations in data [1], [12], [20]. Other deep learning approaches to disentangling include the adversarial setting [21]–[23]. Disentangled representations have also been studied in supervised and semi-supervised contexts [24].

REFERENCES

- [1] Y. Bengio, A. Courville, and P. Vincent, “Representation learning: A review and new perspectives,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [2] B. A. Olshausen and D. J. Field, “Sparse coding with an overcomplete basis set: A strategy employed by v1?” *Vision research*, vol. 37, no. 23, pp. 3311–3325, 1997.
- [3] G. E. Hinton and R. S. Zemel, “Autoencoders, minimum description length and helmholtz free energy,” in *Advances in neural information processing systems*, 1994, pp. 3–10.
- [4] A. Alemi, B. Poole, I. Fischer, J. Dillon, R. A. Saurus, and K. Murphy, “An information-theoretic analysis of deep latent-variable models,” 2018.
- [5] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [6] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, “beta-va: Learning basic visual concepts with a constrained variational framework.” *ICLR*, vol. 2, no. 5, p. 6, 2017.
- [7] D. P. Kingma and M. Welling, “An introduction to variational autoencoders,” *arXiv preprint arXiv:1906.02691*, 2019.
- [8] A. Hyvärinen and E. Oja, “Independent component analysis: algorithms and applications,” *Neural networks*, vol. 13, no. 4-5, pp. 411–430, 2000.
- [9] I. Khemakhem, D. P. Kingma, and A. Hyvärinen, “Variational autoencoders and nonlinear ica: A unifying framework,” *arXiv preprint arXiv:1907.04809*, 2019.

- [10] M. E. Tipping and C. M. Bishop, "Probabilistic principal component analysis," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 61, no. 3, pp. 611–622, 1999.
- [11] F. Locatello, S. Bauer, M. Lucic, S. Gelly, B. Schölkopf, and O. Bachem, "Challenging common assumptions in the unsupervised learning of disentangled representations," *arXiv preprint arXiv:1811.12359*, 2018.
- [12] I. Higgins, D. Amos, D. Pfau, S. Racaniere, L. Matthey, D. Rezende, and A. Lerchner, "Towards a definition of disentangled representations," *arXiv preprint arXiv:1812.02230*, 2018.
- [13] P. Dayan, L. F. Abbott *et al.*, *Theoretical neuroscience*. Cambridge, MA: MIT Press, 2001, vol. 806.
- [14] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Machine learning*, vol. 37, no. 2, pp. 183–233, 1999.
- [15] C. Doersch, "Tutorial on variational autoencoders," *arXiv preprint arXiv:1606.05908*, 2016.
- [16] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner, "Understanding disentangling in β -vae," *arXiv preprint arXiv:1804.03599*, 2018.
- [17] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz, "Multimodal unsupervised image-to-image translation," in *The European Conference on Computer Vision (ECCV)*, September 2018.
- [18] G. Lample, N. Zeghidour, N. Usunier, A. Bordes, L. Denoyer, and M. Ranzato, "Fader networks: Manipulating images by sliding attributes," in *Advances in Neural Information Processing Systems*, 2017, pp. 5967–5976.
- [19] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4401–4410.
- [20] J. J. DiCarlo and D. D. Cox, "Untangling invariant object recognition," *Trends in cognitive sciences*, vol. 11, no. 8, pp. 333–341, 2007.
- [21] E. L. Denton *et al.*, "Unsupervised learning of disentangled representations from video," in *Advances in neural information processing systems*, 2017, pp. 4414–4423.
- [22] L. Tran, X. Yin, and X. Liu, "Disentangled representation learning gan for pose-invariant face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1415–1424.
- [23] V. John, L. Mou, H. Bahuleyan, and O. Vechtomova, "Disentangled representation learning for text style transfer," *arXiv preprint arXiv:1808.04339*, 2018.
- [24] N. Siddharth, B. Paige, J.-W. Van de Meent, A. Desmaison, N. Goodman, P. Kohli, F. Wood, and P. Torr, "Learning disentangled representations with semi-supervised deep generative models," in *Advances in Neural Information Processing Systems*, 2017, pp. 5925–5935.
- [25] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

APPENDIX A PROOF OF PROPOSITION 2

We prove a more general version of eq. (12) given in Prop. 2. Eq. (13) follows from eq. (12) and Prop. 1.

Proposition 4. Consider an objective function given by a sum of two terms,

$$O(\boldsymbol{\kappa}; \beta) = A(\boldsymbol{\kappa}) - \beta B(\boldsymbol{\kappa}), \quad (28)$$

to be maximized over parameters $\boldsymbol{\kappa}$, and β is a hyperparameter. Let $\boldsymbol{\kappa}^*(\beta) = \arg \max_{\boldsymbol{\kappa}} O(\boldsymbol{\kappa}, \beta)$. As β increases $B(\boldsymbol{\kappa}^*(\beta))$ is nonincreasing.

Proof. The proof uses contradiction. Let $\beta_2 > \beta_1$ and

$$\boldsymbol{\kappa}_1 \equiv \boldsymbol{\kappa}^*(\beta_1), \quad \boldsymbol{\kappa}_2 \equiv \boldsymbol{\kappa}^*(\beta_2). \quad (29)$$

Then

$$\begin{aligned} O(\boldsymbol{\kappa}_1, \beta_1) &= O(\boldsymbol{\kappa}_1, \beta_2) + (\beta_2 - \beta_1)B(\boldsymbol{\kappa}_1) \\ &\leq O(\boldsymbol{\kappa}_2, \beta_2) + (\beta_2 - \beta_1)B(\boldsymbol{\kappa}_1), \end{aligned} \quad (30)$$

where the first line is an identity, and the second line follows from the optimality of $\boldsymbol{\kappa}_2$ at $\beta = \beta_2$.

Now we assume $B(\boldsymbol{\kappa}_2) > B(\boldsymbol{\kappa}_1)$, and see that this leads to a contradiction.

$$\begin{aligned} O(\boldsymbol{\kappa}_2, \beta_2) + (\beta_2 - \beta_1)B(\boldsymbol{\kappa}_1) \\ < O(\boldsymbol{\kappa}_2, \beta_2) + (\beta_2 - \beta_1)B(\boldsymbol{\kappa}_2) = O(\boldsymbol{\kappa}_2, \beta_1). \end{aligned} \quad (31)$$

The inequality follows from our assumption, and the equality from (28). Combined with (30), this implies

$$O(\boldsymbol{\kappa}_1, \beta_1) < O(\boldsymbol{\kappa}_2, \beta_1) \quad (32)$$

which contradicts (29). Therefore if $\beta_2 > \beta_1$, then $B(\boldsymbol{\kappa}_2) \leq B(\boldsymbol{\kappa}_1)$. \square

APPENDIX B

DETAILS OF THE ANALYTICALLY TRACTABLE β -VAE MODEL

A. Integrating out data from the objective

The full β -VAE objective is (6) averaged with respect to the data distribution $p(\mathbf{x})$:

$$\begin{aligned} L(\boldsymbol{\theta}, \boldsymbol{\phi}; \beta) &\equiv \mathbb{E}_{p(\mathbf{x})} [\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}; \beta)] \\ &= \mathbb{E}_{p(\mathbf{x})} [\mathbb{E}_{q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})} [\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})] - \beta D_{KL}(q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))]. \end{aligned} \quad (33)$$

We first calculate $\mathbb{E}_{q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})} [\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})]$. We use the reparametrization trick: For $\mathbf{z} \sim q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}}, \boldsymbol{\Sigma}_{\mathbf{z}})$, we can write $\mathbf{z} = \boldsymbol{\mu}_{\mathbf{z}} + \boldsymbol{\Sigma}_{\mathbf{z}}^{1/2} \boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$. Then,

$$\begin{aligned} &\mathbb{E}_{\mathbf{z} \sim q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})} [\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})] \\ &= \mathbb{E}_{\mathbf{z} \sim q_{\boldsymbol{\phi}}(\mathbf{z}|\mathbf{x})} \left[\log \mathcal{N}(\mathbf{x}; \mathbf{D}\mathbf{z} + \mathbf{b}^D, \mathbf{I}_N) \right] \\ &= \mathbb{E}_{\boldsymbol{\epsilon} \sim \mathcal{N}(0,1)} \left[\log \mathcal{N}(\mathbf{x}; \mathbf{D}(\boldsymbol{\mu}_{\mathbf{z}} + \boldsymbol{\Sigma}_{\mathbf{z}}^{1/2} \boldsymbol{\epsilon}) + \mathbf{b}^D, \mathbf{I}_N) \right] \\ &= -\frac{N}{2} \log(2\pi) - \frac{1}{2} (\mathbf{D}\boldsymbol{\mu}_{\mathbf{z}} + \mathbf{b}^D - \mathbf{x})^2 \\ &\quad - \frac{1}{2} \mathbb{E}_{\boldsymbol{\epsilon} \sim \mathcal{N}(0,1)} \left[\boldsymbol{\epsilon}^{\top} (\mathbf{D}\boldsymbol{\Sigma}_{\mathbf{z}}^{1/2})^{\top} (\mathbf{D}\boldsymbol{\Sigma}_{\mathbf{z}}^{1/2}) \boldsymbol{\epsilon} \right]. \end{aligned}$$

The last term can be calculated by the following useful trick. Let's introduce a source term \mathbf{J} into the generating functional,

$$Z[\mathbf{J}] = \int \frac{d\mathbf{z}}{(2\pi)^{n/2} \sqrt{\det \boldsymbol{\Sigma}_{\mathbf{z}}}} \exp \left(-\frac{1}{2} \mathbf{z}^{\top} \boldsymbol{\Sigma}_{\mathbf{z}}^{-1} \mathbf{z} + \mathbf{J}^{\top} \mathbf{A} \mathbf{z} \right), \quad (34)$$

then differentiating with respect to the source,

$$\left(\frac{\delta}{\delta \mathbf{J}} \right)^{\top} \left(\frac{\delta}{\delta \mathbf{J}} \right) Z[\mathbf{J}] \Big|_{\mathbf{J}=\mathbf{0}} = \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, \boldsymbol{\Sigma}_{\mathbf{z}})} [\mathbf{z}^{\top} \mathbf{A}^{\top} \mathbf{A} \mathbf{z}] \quad (35)$$

On the other hand, we can perform the Gaussian integral in $Z[\mathbf{J}]$ to obtain,

$$Z[\mathbf{J}] = \exp \left\{ \frac{1}{2} (\mathbf{J}\mathbf{A})^{\top} \boldsymbol{\Sigma}_{\mathbf{z}} (\mathbf{J}\mathbf{A}) \right\}. \quad (36)$$

Then we arrive at

$$\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, \boldsymbol{\Sigma}_{\mathbf{z}})} [\mathbf{z}^{\top} \mathbf{A}^{\top} \mathbf{A} \mathbf{z}] = \text{Tr}(\mathbf{A}\boldsymbol{\Sigma}_{\mathbf{z}}\mathbf{A}^{\top}) \quad (37)$$

Eq. (37) is central to the calculations of many results presented in the text.

Going back to the reconstruction objective, using eq. (37) we have (up to constants)

$$\mathbb{E}_{\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z})] = -(\mathbf{D}\boldsymbol{\mu}_z + \mathbf{b}^D - \mathbf{x})^2 - \text{Tr}(\mathbf{D}^\top \mathbf{D} \boldsymbol{\Sigma}_z). \quad (38)$$

Similarly we can calculate the conditional independence loss,

$$D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) = -\frac{1}{2} (k + \text{Tr} \log \boldsymbol{\Sigma}_z - \boldsymbol{\mu}_z^\top \boldsymbol{\mu}_z - \text{Tr} \boldsymbol{\Sigma}_z). \quad (39)$$

Putting everything together, the objective function we want to maximize is (neglecting constant terms)

$$L(\boldsymbol{\theta}, \phi; \beta) = \frac{1}{2} \mathbb{E}_{p(\mathbf{x})} [-(\mathbf{D}\boldsymbol{\mu}_z + \mathbf{b}^D - \mathbf{x})^\top (\mathbf{D}\boldsymbol{\mu}_z + \mathbf{b}^D - \mathbf{x}) - \beta \boldsymbol{\mu}_z^\top \boldsymbol{\mu}_z - \text{Tr}(\mathbf{D}^\top \mathbf{D} \boldsymbol{\Sigma}_z) + \beta \text{Tr} \log \boldsymbol{\Sigma}_z - \beta \text{Tr} \boldsymbol{\Sigma}_z]. \quad (40)$$

The expectation with respect to \mathbf{x} amounts to performing Gaussian integrals in \mathbf{x} , as $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_x)$, and thus can be done exactly. After plugging in the definition of $\boldsymbol{\mu}_z, \boldsymbol{\Sigma}_z$ from eq. (20), and performing the \mathbf{x} integrals, the result is given in eq. (21).

B. Taking derivatives of the objective

In order to take derivatives of eq. (21), we unpack the indices (to ease the notation, we denote $\boldsymbol{\Sigma}_x$ as $\boldsymbol{\Sigma}$, and follow the Einstein summation convention, repeated indices are to be summed over unless the summation is explicitly specified)

$$L = -\frac{1}{2} \left\{ (D_{ij} W_{jk}^\mu - \delta_{ik}) \Sigma_{kl} (W_{ml}^\mu D_{im} - \delta_{il}) - \beta \sum_i b_i^\sigma + \beta (W^\mu)_{ij} \Sigma_{jk} W_{ik}^\mu + (D_{ij} b_j^\mu + b_i^D)^2 + \beta (b_i^\mu)^2 + \sum_i \left(D_{ii}^2 + \beta \right) \exp \left(\frac{1}{2} W_{ij}^\sigma \Sigma_{jk} W_{ik}^\sigma + b_i^\sigma \right) \right\} \quad (41)$$

Then,

$$0 = \frac{\partial L}{\partial W_{ab}^\mu} = [(\mathbf{D}^\top (\mathbf{D} \mathbf{W}^\mu - \mathbf{I}_N) + \beta \mathbf{W}^\mu) \boldsymbol{\Sigma}]_{ab}, \quad (42)$$

$$0 = \frac{\partial L}{\partial b_a^\mu} = [(\mathbf{D} \mathbf{b}^\mu + \mathbf{b}^D) \mathbf{D} + \beta \mathbf{b}^\mu]_a, \quad (43)$$

$$0 = \frac{\partial L}{\partial D_{ab}} = [(\mathbf{D} \mathbf{W}^\mu - \mathbf{I}_N) \boldsymbol{\Sigma} (\mathbf{W}^\mu)^\top]_{ab} + [\mathbf{D} \mathbf{b}^\mu + \mathbf{b}^D]_a b_b^\mu + D_{ab} e^{\frac{1}{2} [\mathbf{W}^\sigma \boldsymbol{\Sigma} (\mathbf{W}^\sigma)^\top]_{bb} + b_b^\sigma}, \quad (44)$$

$$0 = \frac{\partial L}{\partial b_a^D} = [\mathbf{D} \mathbf{b}^\mu + \mathbf{b}^D]_a \quad (45)$$

$$0 = \frac{\partial L}{\partial W_{ab}^\sigma} = ([\mathbf{D}^\top \mathbf{D}]_{aa} + \beta) e^{\frac{1}{2} [\mathbf{W}^\sigma \boldsymbol{\Sigma}]_{ab} W_{ab}^\sigma + b_a^\sigma} [\mathbf{W}^\sigma \boldsymbol{\Sigma}]_{ab}, \quad (46)$$

$$0 = \frac{\partial L}{\partial b_a^\sigma} = ([\mathbf{D}^\top \mathbf{D}]_{aa} + \beta) e^{\frac{1}{2} [\mathbf{W}^\sigma \mathbf{C} (\mathbf{W}^\sigma)^\top]_{aa} + b_a^\sigma} - \beta. \quad (47)$$

From the b_a^μ and b_a^D equations we can immediately see $\mathbf{b}^\mu = \mathbf{b}^D = \mathbf{0}$.

C. Derivation of the ground-truth posterior

We observe that since both \mathbf{s} and $\boldsymbol{\eta}$ are independently normally distributed in (17), \mathbf{s} and $\boldsymbol{\eta}$ are jointly normal, i.e., $p(\mathbf{s}, \boldsymbol{\eta})$ is a normal distribution. However, note that $p(\mathbf{s}, \boldsymbol{\eta})$ is just $p(\mathbf{s}, \mathbf{x})$ up to a coordinate transformation, so $p(\mathbf{s}, \mathbf{x})$ is also normal. Also, as $\mathbf{s} \in \mathbb{R}^k$, $\mathbf{x} \in \mathbb{R}^N$, $(\mathbf{s}, \mathbf{x}) \in \mathbb{R}^{N+k}$. We can think of \mathbf{s} and \mathbf{x} partition a $(N+k)$ -dimensional normal distribution $p((\mathbf{s}, \mathbf{x}))$. Therefore, to find the conditional probability $p_{\mathbf{g}\text{-t}}(\mathbf{s}|\mathbf{x})$, we can just use the formula for conditioning multivariate normal distribution:

$$p_{\mathbf{g}\text{-t}}(\mathbf{s}|\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{s}|\mathbf{x}}, \boldsymbol{\Sigma}_{\mathbf{s}|\mathbf{x}}), \quad (48)$$

where

$$\begin{aligned} \boldsymbol{\mu}_{\mathbf{s}|\mathbf{x}} &= \boldsymbol{\mu}_s + \text{Cov}(\mathbf{x}, \mathbf{s})^\top (\boldsymbol{\Sigma}_x)^{-1} (\mathbf{x} - \boldsymbol{\mu}_x) \\ \boldsymbol{\Sigma}_{\mathbf{s}|\mathbf{x}} &= \boldsymbol{\Sigma}_s - \text{Cov}(\mathbf{x}, \mathbf{s})^\top (\boldsymbol{\Sigma}_x)^{-1} \text{Cov}(\mathbf{x}, \mathbf{s}). \end{aligned} \quad (49)$$

Now specializing to our case (17),

$$\text{Cov}(\mathbf{x}, \mathbf{s}) = \text{Cov}(\mathbf{A} \mathbf{s} + \boldsymbol{\eta}, \mathbf{s}) = \mathbf{A} \boldsymbol{\Sigma}_s + \text{Cov}(\mathbf{s}, \boldsymbol{\eta}) = \mathbf{A}. \quad (50)$$

Note that $\boldsymbol{\Sigma}_x = \mathbf{A} \mathbf{A}^\top + \mathbf{I}_N$, then

$$\boldsymbol{\mu}_{\mathbf{s}|\mathbf{x}} = \mathbf{A}^\top (\mathbf{A} \mathbf{A}^\top + \mathbf{I}_N)^{-1} \mathbf{x} = (\mathbf{A}^\top \mathbf{A} + \mathbf{I}_k)^{-1} \mathbf{A}^\top \mathbf{x}, \quad (51)$$

where in the second equality we have used the *matrix push-through identity*: For any matrices $\mathbf{U} \in \mathbb{R}^{N \times k}$, $\mathbf{V} \in \mathbb{R}^{k \times N}$,

$$(\mathbf{I}_N + \mathbf{U} \mathbf{V})^{-1} \mathbf{U} = \mathbf{U} (\mathbf{I}_k + \mathbf{V} \mathbf{U})^{-1}. \quad (52)$$

Now the covariance,

$$\begin{aligned} \boldsymbol{\Sigma}_{\mathbf{s}|\mathbf{x}} &= \mathbf{I}_k - \mathbf{A}^\top (\mathbf{A} \mathbf{A}^\top + \mathbf{I}_N)^{-1} \mathbf{A} \\ &= \mathbf{I}_k - \mathbf{A}^\top \mathbf{A} (\mathbf{A}^\top \mathbf{A} + \mathbf{I}_k)^{-1} \\ &= \mathbf{I}_k - \mathbf{A}^\top \mathbf{A} [(\mathbf{A}^\top \mathbf{A})^{-1} - (\mathbf{A}^\top \mathbf{A})^{-1} (\mathbf{A}^\top \mathbf{A} + \mathbf{I}_k)^{-1}] \\ &= (\mathbf{A}^\top \mathbf{A} + \mathbf{I}_k)^{-1}, \end{aligned} \quad (53)$$

where in the third equality we have used the *Woodbury matrix identity*: For any invertible matrix $\mathbf{B} \in \mathbb{R}^{N \times N}$ and size compatible matrices $\mathbf{U} \in \mathbb{R}^{N \times k}$ and $\mathbf{V} \in \mathbb{R}^{k \times N}$:

$$(\mathbf{B} + \mathbf{U} \mathbf{V})^{-1} = \mathbf{B}^{-1} - \mathbf{B}^{-1} \mathbf{U} (\mathbf{I}_k + \mathbf{V} \mathbf{B}^{-1} \mathbf{U})^{-1} \mathbf{V} \mathbf{B}^{-1}. \quad (54)$$

D. Derivation of the model posterior

Our goal is to use the Bayes rule to calculate the model posterior, $p_\theta(\mathbf{x}|\mathbf{z}) = p_\theta(\mathbf{x}|\mathbf{z}) p(\mathbf{z}) / p_\theta(\mathbf{x})$.

In order to do so, we first need to calculate the *evidence* $p_\theta(\mathbf{x})$,

$$p_\theta(\mathbf{x}) = \int_{\mathbb{R}^k} dz p_\theta(\mathbf{x}|\mathbf{z}) p(\mathbf{z}) \quad (55)$$

$$= \int_{\mathbb{R}^k} dz \mathcal{N}(\mathbf{D} \mathbf{z}, \mathbf{I}_N) \mathcal{N}(\mathbf{0}, \mathbf{I}_k) \quad (56)$$

$$= \mathcal{N}(\mathbf{0}, (\mathbf{D} \mathbf{D}^\top + \mathbf{I}_N)), \quad (57)$$

where in the third equality we have used eq.s (52) and (54) to simplify. Therefore,

$$p_\theta(\mathbf{x}|\mathbf{z}) = \frac{\mathcal{N}(\mathbf{D} \mathbf{z}, \mathbf{I}_N) \mathcal{N}(\mathbf{0}, \mathbf{I}_k)}{\mathcal{N}(\mathbf{0}, (\mathbf{D} \mathbf{D}^\top + \mathbf{I}_N))} \quad (58)$$

After some simplifications using eq.s (52) and (54), we arrived at

$$p_{\theta}(\mathbf{x}|\mathbf{z}) = \mathcal{N}((\mathbf{D}\mathbf{D}^{\top} + \mathbf{I}_N)^{-1}\mathbf{D}^{\top}, (\mathbf{D}\mathbf{D}^{\top} + \mathbf{I}_N)^{-1}) \\ \equiv \mathcal{N}(\boldsymbol{\mu}_{\mathbf{z}|\mathbf{x}}, \boldsymbol{\Sigma}_{\mathbf{z}|\mathbf{x}}). \quad (59)$$

E. Derivation of MIE/TIE

First let's consider MIE. Let

$$\boldsymbol{\mu}_{\mathbf{z}|\mathbf{x}} \equiv \mathbf{F}\mathbf{x}, \quad \boldsymbol{\Sigma}_{\mathbf{z}|\mathbf{x}} \equiv \mathbf{E}. \quad (60)$$

Then, we can write MIE as

$$\text{MIE} = \mathbb{E}_{p(\mathbf{x})} [D_{KL}(q_{\phi}(\mathbf{z}|\mathbf{x})||p_{\theta}(\mathbf{x}|\mathbf{z}))] \\ = \frac{1}{2} \mathbb{E}_{p(\mathbf{x})} \left[(\mathbf{F}\mathbf{x} - \boldsymbol{\mu}_{\mathbf{z}})^{\top} \mathbf{E}^{-1} (\mathbf{F}\mathbf{x} - \boldsymbol{\mu}_{\mathbf{z}}) \right. \\ \left. + \text{Tr}(\mathbf{E}^{-1} \boldsymbol{\Sigma}_{\mathbf{z}}) - \log \left(\frac{\det \boldsymbol{\Sigma}_{\mathbf{z}}}{\det \mathbf{E}} \right) - k \right]. \quad (61)$$

Plugging in eq. (20) and performing the \mathbf{x} Gaussian integrals as in Appendix B-A, we arrive at eq. (25).

Note that at network optimum, our model posterior $p_{\theta}(\mathbf{z}|\mathbf{x})$ equals to the ground-truth posterior $p(\mathbf{s}|\mathbf{x})$ upon changing \mathbf{D} to \mathbf{A} . Therefore, we just need to replace \mathbf{D} by \mathbf{A} in the above derivation to obtain the results for TIE.

APPENDIX C SIMULATION DETAILS

The deep neural network models used for the numerical experiments task used the same overall architecture. The encoder is a feed forward network with 3 hidden layers, with 256, 200, and 200 units. 2 parallel hidden layers with 2 neurons parameters the mean and variance for $k = 2$ latent variables. The decoder consists of 3 feed-forward hidden layers with 200, 200, and 256 units, then outputs the reconstructed image. The network was trained for 1000 epochs over the entire synthetic dataset, comprising of 1000 examples. We used a tanh activation function used along with Adam Optimization [25] with a learning rate of 1e-3. Experiments were repeated across 300 realizations for each β value. Results shown were averaged over the whole set of realizations.

The Reconstruction Objective was calculated for each trained model through generating 1000 samples from the encoder, passing them to the decoder to approximately calculate $\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} [\log p_{\theta}(\mathbf{x}|\mathbf{z})]$, and averaging over the data \mathbf{x} . The Conditional Independence Loss was calculated directly using the Tensorflow Distributions library's native KL Divergence method. The ELBO was calculated by numerically taking the difference of these two terms, and the β -VAE objective was an extension of this with the hyperparameter β included. The Inference Error was calculated numerically using the modelled $\boldsymbol{\mu}_{\mathbf{z}}$ and $\boldsymbol{\Sigma}_{\mathbf{z}}$ and estimating $p(\mathbf{x})$ from mini-batches.

In Fig. C.1, we show results on another simulation consistent with our findings.

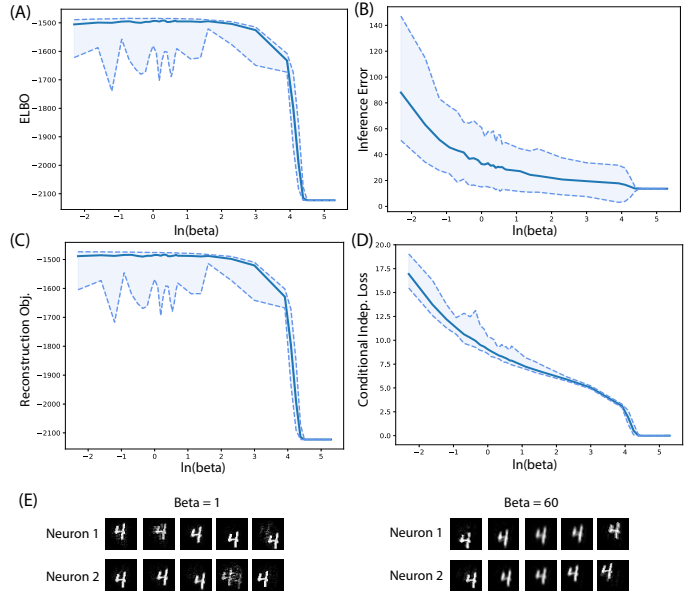


Fig. C.1. Values for error terms across 300 random initializations of the network for a synthetic dataset, which comprises of a single MNIST digit localized at different locations on a blank canvas. The cartesian coordinate of the digit in a sample from our data, \mathbf{x} , is determined by eq. (17), with $\mathbf{A} = \begin{pmatrix} 2.73 & 0.73 \\ 0.73 & 2.73 \end{pmatrix}$, $\mathbf{s} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$, $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_N)$, $k = 2$. Dashed lines represent the minimum and maximum values, and solid line represents the average. (A) ELBO as a function of β . (B) TIE as a function of β . Its minima over various random initialization follow a non-monotonic trend. (C) Reconstruction objective as a function of β . (D) Conditional independence loss as a function of β . (E) Traversal of latent encoding in bottleneck neurons for small and large β . One neuron is held fixed while the other is modulated to generate reconstructions. Reconstruction of the digits noticeable worsens with higher β , while units in the bottleneck encode for structured, orthogonal axes of motion.