



Disentangling Generative Factors of Physical Fields Using Variational Autoencoders

Christian Jacobsen and Karthik Duraisamy*

Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI, United States

The ability to extract generative parameters from high-dimensional fields of data in an unsupervised manner is a highly desirable yet unrealized goal in computational physics. This work explores the use of variational autoencoders for non-linear dimension reduction with the specific aim of *disentangling* the low-dimensional latent variables to identify independent physical parameters that generated the data. A disentangled decomposition is interpretable, and can be transferred to a variety of tasks including generative modeling, design optimization, and probabilistic reduced order modelling. A major emphasis of this work is to characterize disentanglement using VAEs while minimally modifying the classic VAE loss function (i.e., the Evidence Lower Bound) to maintain high reconstruction accuracy. The loss landscape is characterized by over-regularized local minima which surround desirable solutions. We illustrate comparisons between disentangled and entangled representations by juxtaposing learned latent distributions and the true generative factors in a model porous flow problem. Hierarchical priors are shown to facilitate the learning of disentangled representations. The regularization loss is unaffected by latent rotation when training with rotationally-invariant priors, and thus learning non-rotationally-invariant priors aids in capturing the properties of generative factors, improving disentanglement. Finally, it is shown that semi-supervised learning - accomplished by labeling a small number of samples ($O(1\%)$)—results in accurate disentangled latent representations that can be consistently learned.

Keywords: generative modeling, unsupervised learning, variational autoencoders, scientific machine learning, disentangling

OPEN ACCESS

Edited by:

Traian Iliescu,
Virginia Tech, United States

Reviewed by:

Claire Heaney,
Imperial College London,
United Kingdom
Andrey Popov,
Virginia Tech, United States

*Correspondence:

Karthik Duraisamy
kdur@umich.edu

Specialty section:

This article was submitted to
Statistical and Computational Physics,
a section of the journal
Frontiers in Physics

Received: 07 March 2022

Accepted: 19 May 2022

Published: 30 June 2022

Citation:

Jacobsen C and Duraisamy K (2022)
Disentangling Generative Factors of
Physical Fields Using
Variational Autoencoders.
Front. Phys. 10:890910.
doi: 10.3389/fphy.2022.890910

1 INTRODUCTION

Unsupervised representation learning is a popular area of research because of the need for low-dimensional representations in unlabeled data. Low-dimensional *latent* representations of high-dimensional data have many applications ranging from facial image generation [1] and music generation [2] to autonomous controls [3] among many others. Generative adversarial networks (GANs) [4], variational autoencoders (VAEs) [5] and their variants [6–9], among other methods, aim to approximate an underlying distribution $p(y)$ of high-dimensional data through a two-step process. Compressed representations z are sampled from a low-dimensional—yet unknown—distribution $p(z)$. In the case of VAEs, which is the focus of this work, an encoding distribution $p(z|y)$ and a decoding distribution are learned simultaneously by maximizing a bound on the likelihood of the data (i.e., the evidence lower bound (ELBO) [5]). Thus, a mapping from the high-dimensional space to a low-dimensional space and the corresponding inverse mapping is

learned simultaneously, allowing approximations of both $p(y)$ and $p(z)$. Learning the lower-dimensional representation, or *latent space*, can facilitate computationally-efficient data generation and extract only the information necessary to reconstruct the data [10]. Modifications to the ELBO objective have been suggested in the literature, primarily with improved disentanglement in mind. The β -VAE [6] was introduced to improve disentanglement by adjusting the weight of regularization loss. FactorVAE [8] introduces a total correlation term (TC) to encourage learning a factorized latent representation. InfoVAE [9] augments the ELBO with a term to promote maximization of the mutual information between the data and the learned representation. Many other developments based on the VAE objective have been introduced in the literature. VAEs have been implemented in many applications including inverse problems [11], extracting physical parameters from spatio-temporal data [12], and constructing probabilistic reduced order models [13, 14], among others.

To illustrate the idea of disentanglement and its implications, consider a dataset consisting of images of teapots [15]. Each image is generated from 3 parameters indicating the color of the teapot (RGB) and 2 parameters corresponding to the angle the teapot is viewed from. Thus, even though the RGB image may be very high dimensional, the intrinsic dimensionality is just 5. Representation learning can be used to extract a low-dimensional latent model containing useful and meaningful representations of the high-dimensional images. Learned latent representations need not be *disentangled* to be useful in some sense, but disentanglement enhances interpretability of the representation. Disentanglement references a structure of the latent distribution in which changes in each parameter in the learned representation correspond directly to changes in a single yet different generative parameter. Humans tend to naturally and easily identify independent factors of variation, and thus a disentangled representation often corresponds to one which would be naturally identified by a human. A representation which is more naturally explained by a human observer is therefore one characterized by greater interpretability. In an unsupervised setting, one cannot guarantee that a disentangled representation can be learned.

The requirement for disentanglement depends on the task at hand, but a disentangled representation may be used in many tasks containing different objectives. Indeed, [10] state that “the most robust approach to feature learning is to disentangle as many factors as possible, discarding as little information about the data as is practical.” In the teapot example, changes in one of the learned latent dimensions may correspond to changes in the color red and one of the viewing angles, which would indicate an entangled representation. Another example, more relevant to our work, is that of fluid flow over an airfoil. Learning a disentangled representation of the flow conditions along with the shape parameters using VAEs can allow rapid prediction of the flow field with enhanced interpretability of the latent representation, facilitating efficient computation of the task at hand. The disentangled representation can be transferred to a variety of tasks easily such as design optimization, developing reduced order models in the latent space or parameter inference from

flow fields. It is the ability of disentangled representations to transfer across tasks with ease and interpretability which makes them so useful. In many practical physics problems, full knowledge regarding the underlying generative parameters of high-dimensional data may not exist, thus making it challenging to ascertain the quality of disentanglement.

Disentanglement using VAEs was first addressed in the literature [6]; [8] by modifying the strength of regularization in the ELBO loss, with the penalty of sub-optimal compression and reconstruction. FactorVAEs [6] encourage a factorized representation, which can be useful for disentanglement in the case of independent generative parameters, but undesirable when parameters are correlated. [16] suggest that the ability of the VAE to learn disentangled representations is not inherent to the framework itself, but an “accidental” byproduct of the typically assumed factorized form of the encoder. The prior distribution is of particular importance as the standard normal prior often assumed allows for rotation of the latent space with no effect on the ELBO loss. Disentangled representations are still often learned due to a factorized form of the encoding distribution with sufficiently large weight on regularization. Additional interpretations and insight into the disentanglement ability of VAEs are found in [17].

Our work on unsupervised representation learning is motivated from a computational-physics perspective. We focus on the application of VAEs for use with data generated by partial differential equations (PDEs). The central questions we seek to answer in this work are: 1) can we reliably disentangle parameters from data obtained from PDEs governing physical problems using VAEs, and 2) what are the characteristics of disentangled representations? Learning disentangled representations can be useful in many capacities: developing probabilistic reduced order models, design optimization, parameter extraction, and data interpolation, among others. Many of the applications of such representations, and the ability to transfer between them, rely heavily on the disentanglement of the latent space. Differences in disentangled and entangled representations are identified, and conclusions are drawn regarding the inconsistencies in learning such representations. Our goals are not to compare the available methods to promote disentanglement, as in [18], but rather to illustrate the use of VAEs without modifying the ELBO and to understand the phenomenon of disentanglement itself in this capacity. The use of hierarchical priors is shown to greatly improve the prospect of learning a disentangled representation in some cases without altering the standard VAE loss through the learning of non-rotationally-invariant priors. Along the way, we provide intuition on the objective of VAEs through connections to rate-distortion theory, illustrate some of the challenges of implementing and training VAEs, and provide potential methods to overcome some of these issues such as “vanishing KL” [19].

The outline of this paper is as follows: In **Section 2**, we introduce the VAE, connect it to rate-distortion (RD) theory, discuss disentanglement, and derive a bound on the classic VAE loss (the ELBO) using hierarchical priors (HP). In **Section 3**, we introduce a sample application of Darcy flow as the main

illustrative example of this work. In **Section 4**, we present challenges in training VAEs and include possible solutions, and investigate the ELBO loss landscape. We illustrate disentanglement of parameters on the Darcy flow problem, and provide insight into the phenomenon of disentanglement in **Section 5**. The use of a small amount of labeled data (semi-supervised learning) is considered in **Section 6**. In **Section 7**, conclusions and insights are drawn on the results of our work, and future directions are discussed.

The numerical experiments in this paper can be recreated using our code provided in <https://github.com/christian-jacobsen/Disentangling-Physical-Fields>.

2 VARIATIONAL AUTOENCODER FORMULATION

In many applications of representation learning, it is generally desirable that the latent representation be maximally compressed. In other words, the low dimensional representation contains only the information required to reconstruct the original data, discarding irrelevant information. The VAE framework used extensively in this work is a method of data compression with many ties to information theory [20]. In applications with little to no knowledge regarding the nature of obtained data, the latent factors extracted using VAEs can act as a set of features describing the generative parameters underlying the data. A direct correlation between the generative parameters and the compressed representation, or a disentangled representation, is sought such that the representation can be applied to a multitude of downstream tasks. Some example tasks include performing predictions on new generative parameters, interpreting the data in the case of unknown generative parameters, and computationally efficient design optimization.

Data snapshots obtained from some physical system or a model of that system is represented here by random variable $Y: \Omega \rightarrow \mathcal{Y}$ where Ω is a sample space and \mathcal{Y} is a measurable space ($\mathcal{Y} = \mathbb{R}^m$ will be assumed for the remainder of this work). Each realization of Y is generated from a function of Θ such that $\Theta: \Omega \rightarrow \mathbb{R}^s$ is a random variable representing generative parameters with distribution $p(\theta)$. With no prior knowledge of Θ , the random variable $Z: \Omega \rightarrow \mathbb{R}^n$ represents the latent parameters to be inferred from the data. A probabilistic relationship between Θ and Y is sought in an unsupervised manner using only samples from $p(y)$.

The VAE framework infers a latent-variable model by replacing the posterior $p(z|y)$ with a parameterized approximating posterior $q_\phi(z|y)$ [5], known as the encoding distribution. A parameterized decoding distribution $p_\psi(y|z)$ is also constructed to predict data samples given samples from the latent space. Only the encoding distribution and the decoding distribution are learned in the VAE framework, but the *aggregated posterior* $q_\phi(z)$ (to the best of our knowledge, first referred to in this way by [21]), is of particular importance in

disentanglement. It is defined as the marginal latent distribution induced by the encoder

$$q_\phi(z) \triangleq \int_{\mathcal{Y}} p(y)q_\phi(z|y)dy, \tag{1}$$

where the true data distribution is denoted by $p(y)$. The induced data distribution is the marginal output distribution induced by the decoder

$$p_\psi(y) \triangleq \int_{\mathbb{R}^n} p(z)p_\psi(y|z)dz. \tag{2}$$

It is noted that the true data distribution is typically unknown; only samples of data $\{y^{(i)}\}_{i=1}^N$ are available. The empirical data distribution is thus denoted $\hat{p}(y)$, and any expectation with respect to the empirical distribution is simply computed as an empirical average $\mathbb{E}_{\hat{p}(y)}[f(y)] \triangleq \frac{1}{N}\sum_{i=1}^N f(y^{(i)})$.

Learning the latent model is accomplished by simultaneously learning the encoding and decoding distributions through maximizing the evidence lower bound (ELBO), which is a lower bound on the log-likelihood [22]. To derive the ELBO loss, we begin by expanding the relative entropy between the data distribution and the *induced* data distribution

$$D_{KL}[p(y)||p_\psi(y)] = \mathbb{E}_{Y \sim p(y)}[\log p(y)] - \mathbb{E}_{Y \sim p(y)}[\log p_\psi(y)]$$

where the first term on the right hand side is the negative differential entropy $-H(Y)$. Noting that relative entropy D_{KL} —also often called the Kullback–Leibler divergence, which is a measure of the distance between two probability distributions—is always greater than or equal to zero and introducing Bayes’ rule as

$$p_\psi(y) = \frac{p_\psi(y|z)p(z)p_\phi(z|y)}{p(z|y)p_\phi(z|y)},$$

we arrive at the following inequality

$$H(Y) + \mathbb{E}_{Y \sim p(y)}[D_{KL}[p_\phi(z|y)||p(z|y)]] \leq \mathbb{E}_{p(y)}\left[\mathbb{E}_{q_\phi(z|y)}[\log p_\psi(y|z)]\right] - \mathbb{E}_{p(y)}D_{KL}[q_\phi(z|y)||p(z)].$$

Thus,

$$\mathbb{E}_{p(y)}[\log(p(y))] \geq \mathbb{E}_{p(y)}\left[\mathbb{E}_{q_\phi(z|y)}[\log p_\psi(y|z)]\right] - \mathbb{E}_{p(y)}D_{KL}[q_\phi(z|y)||p(z)], \tag{3}$$

where $p(z)$ is a prior distribution. The prior is specified by the user in the classic VAE framework. The right-hand side in **Eq. 3** is the well-known ELBO. Maximizing this lower bound on the log-likelihood of the data is done by minimizing the negative ELBO. The optimization is performed by learning the encoder and decoder parameterized as neural networks. The negative ELBO is defined as

$$-ELBO = \mathbb{E}_{p(y)}D_{KL}[q_\phi(z|y)||p(z)] + \mathbb{E}_{p(y)}\left[\mathbb{E}_{q_\phi(z|y)}[-\log p_\psi(y|z)]\right], \tag{4}$$

and we assume $\mathcal{L}_{VAE} \approx -ELBO$, where the difference results in the expectation being evaluated over the empirical data distribution in \mathcal{L}_{VAE} . The VAE loss function is defined as

$$\mathcal{L}_{VAE} = \mathbb{E}_{\hat{p}(y)} D_{KL} [q_{\phi}(z|y) || p(z)] + \mathbb{E}_{\hat{p}(y)} \left[\mathbb{E}_{q_{\phi}(z|y)} [-\log p_{\psi}(y|z)] \right], \quad (5)$$

where the first term on the right-hand side is the regularization loss \mathcal{L}_{REG} and drives the encoding distribution closer (in the sense of minimizing KL divergence) to the prior distribution. The second term on the right-hand side is the reconstruction error \mathcal{L}_{REC} and encourages accurate reconstruction of the data.

Selecting the prior distribution as well as the parametric form of the encoding and decoding distribution can allow closed form solutions to compute \mathcal{L}_{VAE} . The prior distribution is often conveniently chosen as a standard normal distribution $p(z) = \mathcal{N}(z; 0, I_{n \times n})$. The encoding and decoding distributions are also often chosen as factorized normal distributions $q_{\phi}(z|y) = \mathcal{N}(z; \mu_{\phi}(y), \text{diag}(\sigma_{\phi}(y)))$ and $p_{\psi}(y|z) = \mathcal{N}(y; \mu_{\psi}(z), \text{diag}(\sigma_{\psi}(y)))$, where the mean and log-variance of each distribution are functions parameterized by neural networks. Selecting the parameterized form of these distributions facilitates the reparameterization trick [5], allowing backpropagation through sampling operations during training. This selection of the prior, encoding, and decoding distributions allows a closed form solution to compute \mathcal{L}_{VAE} .

2.1 Disentanglement

Disentanglement is realized when variations in a single latent dimension correspond to variations in a single generative parameter. This allows the latent space to be interpretable by the user and improves transferability of representations between tasks. Disentanglement may not be required for some tasks which may not require knowledge on each parameter individually or perhaps only a subset of the generative parameters. Nevertheless, a disentangled representation can be leveraged across many tasks. [10] note that a disentangled representation captures each of the relevant features of the data, but downstream applications may only require a subset of these factors. We therefore hypothesize that disentangled representations lead to a more comprehensive range of downstream applications over non disentangled representations.

Many metrics of disentanglement exist in the literature [18], few of which take into account the generative parameter data. Often knowledge on the generative parameters is lacking, and these metrics can be used to evaluate disentanglement in that case (although there is no consensus on which metric is appropriate). In controlled experiments, however, knowledge on generative parameters is available, and correlation between the latent space and the generative parameter space can be directly determined. To evaluate disentanglement in a computationally efficient manner, we propose a disentanglement score

$$S_D = \frac{1}{n} \sum_i \frac{\max_j |\text{cov}(z_i, \theta_j)|}{\sum_j |\text{cov}(z_i, \theta_j)|}, \quad (6)$$

where z_i indicates the i th component of the latent vector $\forall i \in \{1, \dots, n\}$ and θ_j indicates the j th component of the generative parameter vector $\forall j \in \{1, \dots, s\}$. Noting that

$$\frac{\max_j |\text{cov}(z_i, \theta_j)|}{\sum_j |\text{cov}(z_i, \theta_j)|} \in [1/s, 1],$$

it is clear that $S_D \in [1/s, 1]$. It is noted that this score is not used during the training process. This score is created from the intuition that each latent parameter should be correlated to only a single generative parameter. One might note some issues with this disentanglement score. For instance, if multiple latent dimensions are correlated to the same generative parameter dimension, the score will be inaccurate. Similarly, if the latent dimension is greater than the generative parameter dimension, some latent dimensions may contain no information about the data and be uncorrelated to all dimensions, inaccurately reducing the score. For the cases presented here (we will use the score only when $n = s$), **Eq. 6** suffices as a reasonable measure of disentanglement. This score is used as an efficient means of scoring disentanglement when efficiency is important, but we propose another score based on comparisons between disentangled and entangled representations.

We observed empirically that disentanglement is highly correlated to a match in shape between the generative parameter distribution $p(\theta)$ and the aggregated posterior $q_{\phi}(z)$ (**Section 5**). A match in the scaled-and-translated shapes results in good disentanglement but an aggregated posterior which does not match the shape of the generative parameter distribution or contains incorrect correlations (“rotated”) relative to the generative parameter distribution does not. Using this knowledge, another disentanglement metric is postulated to compare these shapes by leveraging the KL Divergence (**Eq. 7**) where \circ denotes the Hadamard product. The disentanglement score is given by

$$S_{KL} = \min_{a,b} D_{KL} [p(\theta) || q_{\phi}(a \circ (z - b))]. \quad (7)$$

This metric compares the shapes of the two distributions by finding the minimum KL divergence between the generative parameter distribution and a scaled and translated version of the aggregated posterior. When $q_{\phi}(a \circ (z - b))$ is close to $p(\theta)$ for some vectors $a, b \in \mathbb{R}^n$, disentanglement is observed.

It is noted in [16] that rotation of the latent space certainly has a large effect on disentanglement, which is precisely what we observe (**Section 5**). Additionally, the ELBO loss is unaffected by rotations of the latent space when using rotationally-invariant priors such as standard normal (**Appendix A**).

2.1.1 β -VAE

The β -VAE objective gives greater weighting to the regularization loss,

$$\mathcal{L}_{\beta\text{-VAE}} = \beta \mathbb{E}_{\hat{p}(y)} D_{KL} [q_{\phi}(z|y) || p(z)] + \mathbb{E}_{\hat{p}(y)} \left[\mathbb{E}_{q_{\phi}(z|y)} [-\log p_{\psi}(y|z)] \right].$$

This encourages greater regularization, often leading to improved disentanglement over the standard VAE loss [6]. It is worth noting that when $\beta = 1$, with a perfect encoder and decoder, the VAE loss reduces to the Bayes rule [23]; [24]. More details on the β -VAE are provided in **Section 2.2**.

2.2 Connections to Rate-Distortion Theory

Rate-distortion theory [25, 26, 27] aids in a deeper understanding in the trade off and balance between the regularization and reconstruction losses. The general rate distortion problem is formulated before making these connections. Consider two random variables: data $Y: \Omega \rightarrow \mathbb{R}^m$ and a compressed representation of the data $Z: \Omega \rightarrow \mathbb{R}^n$. An encoder $p(z|y)$ is sought such that the compressed representation contains a minimal amount of information about the data subject to a bounded error in reconstructing the data. A model $\tilde{y}(z)$ is used to reconstruct Y from samples of Z , and a distortion metric $d(y, \tilde{y})$ is used as a measure of error in the reconstruction of Y with respect to the original data.

A rate-distortion problem thus takes the general form

$$R(D) = \min_{p(z|y)} I(Y; Z) \quad \text{s.t.} \quad \mathbb{E}_{Y,Z}[d(y, \tilde{y}(z))] \leq D, \quad (8)$$

where $D \in \mathbb{R}$ is an upper bound on the distortion. Solutions to **Eq. 8** consist of an encoder $p(z|y)$ which extracts as little information as possible from Y while maintaining a bounded distortion on the reconstruction of Y from Z through the model $\tilde{y}(z)$. Mutual information is minimized to obtain a maximally compressed representation of the data. Learning unnecessary information leads to “memorization” of some aspects of the data rather than extracting only the information relevant to the task at hand. This optimization problem formulated as the rate-distortion Lagrangian is

$$\min_{p(z|y)} \mathcal{J}(\beta) = \min_{p(z|y)} I(Y; Z) + \beta \mathbb{E}_{Y,Z} d(y, \tilde{y}(z)) - D. \quad (9)$$

Given an encoder and decoder, solutions to the rate-distortion problem lie on a convex curve referred to as the rate-distortion curve [20]. Points above this curve correspond to *realizable* yet sub-optimal solutions. Points below the RD curve correspond to solutions which are *not* realizable; no possible compression exists with distortion below the RD curve. As the RD curve is convex, optimal solutions found by varying β lie along the curve. Increasing β increases the tolerable distortion, decreasing the mutual information between the compressed representation and data, providing a more compressed representation. Conversely, decreasing β requires a more accurate reconstruction of the data, increasing the mutual information between compressed representation and data.

The β -VAE loss is tied to a rate-distortion problem. Rearranging the VAE regularization loss (\mathcal{L}_{REG}), we obtain

$$\mathcal{L}_{REG} = \mathbb{E}_{\hat{p}(y)} D_{KL}[q_\phi(z|y) || p(z)] = I_\phi(Y; Z)$$

which is equal to the mutual information between Y and Z according to the data and encoding distributions. Minimizing the β -VAE loss gives the optimization problem

$$\min_{\phi, \psi} \mathcal{L}_{\beta\text{-VAE}} = \min_{\phi, \psi} I_\phi(Y; Z) + \beta \mathbb{E}_{\hat{p}(y) p_\psi(z|y)} [-\log p_\phi(y|z)].$$

This optimization problem is similar to a rate-distortion problem with $d(y, \tilde{y}) = -\log p_\phi(y|z)$ and the mutual information $I_\phi(Y; Z)$ just an approximation to the true mutual information $I(Y; Z)$. Depending on β , solutions can be found at any location along the RD curve with each containing differing properties. RD curve for VAEs is simply an analogy: \mathcal{L}_{REG} is considered the rate R and \mathcal{L}_{REC} is considered the distortion D .

With increased β , the β -VAE minimizes the mutual information between the data and the latent parameters, limiting reconstruction accuracy. In [16], disentanglement is illustrated to be caused inadvertently through the assumed factored form of the encoding distribution even though rotations of the latent space have no effect on the ELBO. However, their proof relies on training in the “polarized” regime characterized by loss of information or “posterior collapse” [28]. Training in this regime often requires increasing the weight of the regularization loss, necessarily decreasing reconstruction performance in the process. In our work, we illustrate disentanglement through training VAEs with the ELBO loss ($\beta = 1$), keeping reconstruction accuracy high. [16] presents good insights into disentanglement.

2.3 Hierarchical Priors

Often the prior (in the case of classic VAEs, specified by the user) and generative parameter distributions (data dependent) may not be highly correlated. Hierarchical priors [7] (HP) can be implemented within the VAE network such that the prior is learned as a function of additional random variables, potentially leading to more expressive priors and aggregated posteriors. Hierarchical random variables ξ_i are introduced such that “sub-priors” can be assumed on each ξ_i (typically standard normal). In the case of a single hierarchical random variable

$$\begin{aligned} p(z) &= \int_{\Xi} p(z|\xi) p(\xi) d\xi = \int_{\Xi} \frac{p(\xi|z)}{p(\xi|z)} p(z|\xi) p(\xi) d\xi \\ &= \mathbb{E}_{\Xi \sim p(\xi|z)} \left[\frac{p(z|\xi) p(\xi)}{p(\xi|z)} \right]. \end{aligned}$$

The conditional distributions $p(\xi|z)$ and $p(z|\xi)$ are the *prior encoder* and *prior decoder*, respectively. These distributions can be approximated by parameterizing them with neural networks. The parameterized distributions are noted as $q_\gamma(\xi|z)$ and $p_\pi(z|\xi)$ where γ are the trainable parameters of the approximating prior encoder and π are the trainable parameters of the prior decoder. Thus, the VAE prior can be approximated through the prior encoding and decoding distributions

$$p(z) \approx \mathbb{E}_{\Xi \sim q_\gamma(\xi|z)} \left[\frac{p_\pi(z|\xi) p(\xi)}{q_\gamma(\xi|z)} \right]. \quad (10)$$

Rearranging the VAE regularization loss

$$\begin{aligned} \mathcal{L}_{REG} &= \int_{Y,Z} \hat{p}(y) q_\phi(z|y) \log \frac{q_\phi(z|y)}{p(z)} dy dz \\ &= \mathbb{E}_{Y,Z \sim \hat{p}(y) q_\phi(z|y)} [\log q_\phi(z|y)] \\ &\quad - \int_{Y,Z} \hat{p}(y) q_\phi(z|y) \log p(z) dy dz, \end{aligned} \quad (11)$$

and substituting the approximating hierarchical prior **Eq. 10** into **Eq. 11**, the final term on the right-hand side becomes

$$-\int_{Y,Z} \hat{p}(y)q_\phi(z|y)\log p(z)dydz = -\int_{Y,Z} \hat{p}(y)q_\phi(z|y)\log \left[\mathbb{E}_{\Xi \sim q_y(\xi|z)} \left[\frac{p_\pi(z|\xi)p(\xi)}{q_y(\xi|z)} \right] \right] dydz.$$

The logarithm function is strictly concave; therefore, by Jensen’s inequality the right-hand side is upper bounded by

$$-\int_{Y,Z} \hat{p}(y)q_\phi(z|y)\log \left[\mathbb{E}_{\Xi \sim q_y(\xi|z)} \left[\frac{p_\pi(z|\xi)p(\xi)}{q_y(\xi|z)} \right] \right] dydz \leq -\int_{Y,Z} \hat{p}(y)q_\phi(z|y)\mathbb{E}_{\Xi \sim q_y(\xi|z)} \left[\log \frac{p_\pi(z|\xi)p(\xi)}{q_y(\xi|z)} \right] dydz.$$

This bound is rearranged to the form

$$\mathbb{E}_{Y,Z \sim \hat{p}(y)q_\phi(z|y)} D_{KL} [q_y(\xi|z) || p(\xi)] - \mathbb{E}_{Y,Z \sim \hat{p}(y)q_\phi(z|y)} \left[\mathbb{E}_{q_y(\xi|z)} [\log p_\pi(z|\xi)] \right]. \tag{12}$$

Equation 12 takes the same form as the overall VAE loss, but applied to the prior network itself. Thus, the hierarchical prior can be thought of as a system of sub-VAEs within the main VAE. In summary, the VAE loss is upper bounded by

$$\mathcal{L}_{VAE} \leq \mathbb{E}_{Y,Z \sim \hat{p}(y)q_\phi(z|y)} [\log q_\phi(z|y)] + \mathbb{E}_{Y,Z \sim \hat{p}(y)q_\phi(z|y)} [D_{KL} [q_y(\xi|z) || p(\xi)]] \tag{13}$$

$$- \mathbb{E}_{Y,Z \sim \hat{p}(y)q_\phi(z|y)} \left[\mathbb{E}_{q_y(\xi|z)} [\log p_\pi(z|\xi)] \right] \tag{14}$$

$$- \mathbb{E}_{Y,Z \sim \hat{p}(y)q_\phi(z|y)} [\log p_\psi(y|z)]. \tag{15}$$

Implementing hierarchical priors can aid in learning non-rotationally-invariant priors, frequently inducing a learned disentangled representation, as shown below.

3 APPLICATION TO DARCY FLOW

To characterize the training process of the VAEs and to study disentanglement, we employ an application of flow through porous media. A two-dimensional steady-state Darcy flow problem in c spatial dimensions (our experiments employ $c = 2$) is governed by [29].

$$\begin{aligned} u(x) &= -K(x)\nabla p(x), \quad x \in \mathcal{X} \\ \nabla \cdot u(x) &= f(x), \quad x \in \mathcal{X} \\ u(x) \cdot \hat{n}(x) &= 0, \quad x \in \partial\mathcal{X} \\ \int_{\mathcal{X}} p(x)dx &= 0. \end{aligned} \tag{16}$$

Darcy’s law is an empirical law describing flow through porous media in which the permeability field is a function of the spatial coordinate $K(x): \mathbb{R}^c \rightarrow \mathbb{R}$. The pressure $p(x): \mathbb{R}^c \rightarrow \mathbb{R}$ and velocity $u(x): \mathbb{R}^c \rightarrow \mathbb{R}^c$ are found given the source term $f(x): \mathbb{R}^c \rightarrow \mathbb{R}$, permeability, and boundary conditions. The integral constraint is given to ensure a unique solution.

A no-flux boundary condition is specified, and the source term models an injection well in one corner of the domain and a production well in the other

$$f(x) = \begin{cases} r, & |x_i - \frac{1}{2}w| \leq \frac{1}{2}w, \quad i = 1, 2 \\ -r, & |x_i - 1 + \frac{1}{2}w| \leq \frac{1}{2}w, \quad i = 1, 2 \\ 0, & \text{otherwise} \end{cases} \tag{17}$$

where $w = \frac{1}{8}$ and $r = 10$. The computational domain considered is the unit square $\mathcal{X} = [0, 1]^2$.

3.1 Karhunen-Loeve Expansion Dataset

The dataset investigated uses a log-permeability field modeled by a Gaussian random field with covariance function k

$$K(x) = \exp(G(x)), \quad G(\cdot) \sim \mathcal{N}(\bar{\mu}, k(\cdot, \cdot)). \tag{18}$$

Generating the data first requires sampling from the permeability field (**Eq. 18**). We take the covariance function as

$$k(x, x') = \exp(-\|x - x'\|_2/l) \tag{19}$$

in our experiments, as in [29]. After sampling the permeability field, solving **Eq. 16** for the pressure and velocity fields produces data samples. We discretize the spatial domain on a 65×65 grid and use a second-order finite difference scheme to solve the system.

The intrinsic dimensionality of the data will be the total number of nodes in system (4,225 for our system) [29]. For dimensionality reduction, the intrinsic dimensionality s of the data is specified by leveraging the Karhunen-Loeve Expansion (KLE), retaining only the first s terms in

$$G(x) = \bar{\mu} + \sum_{i=1}^s \sqrt{\lambda_i} \theta_i \phi_i(x), \tag{20}$$

where λ_i and $\phi_i(x)$ are eigenvalues and eigenfunctions of the covariance function (**Eq. 19**) sorted by decreasing λ_i , and each θ_i are sampled according to some distribution $p(\theta)$, denoted the *generative parameter distribution*.

Each dataset contains some intrinsic dimensionality s , and we denote each dataset using the permeability field (**Eq. 18**) as KLEs. For example, a dataset with $s = 100$ is referred to as KLE100. Samples from datasets of various intrinsic dimension are illustrated in **Figure 1**. Variations on the KLE2 dataset are employed for our explorations in this work. The differences explored are related to varying the generative parameter distribution $p(\theta)$ in each set.

Each snapshot $y^{(i)}$ from a single dataset $\{y^{(i)}\}_{i=1}^N$ contains the pressure $p(x)$ and velocity fields $u(x)$ at each node in the computational domain. These are used as a 3-channel input to the VAE; the permeability field and KLE expansion coefficients (generative parameters) are saved and used only for evaluation purposes.




FIGURE 1 | Samples from datasets (top left) KLE2 (top right) KLE10 (bottom left) KLE100 (bottom right) KLE1000.

4 TRAINING SETUP AND LOSS LANDSCAPE

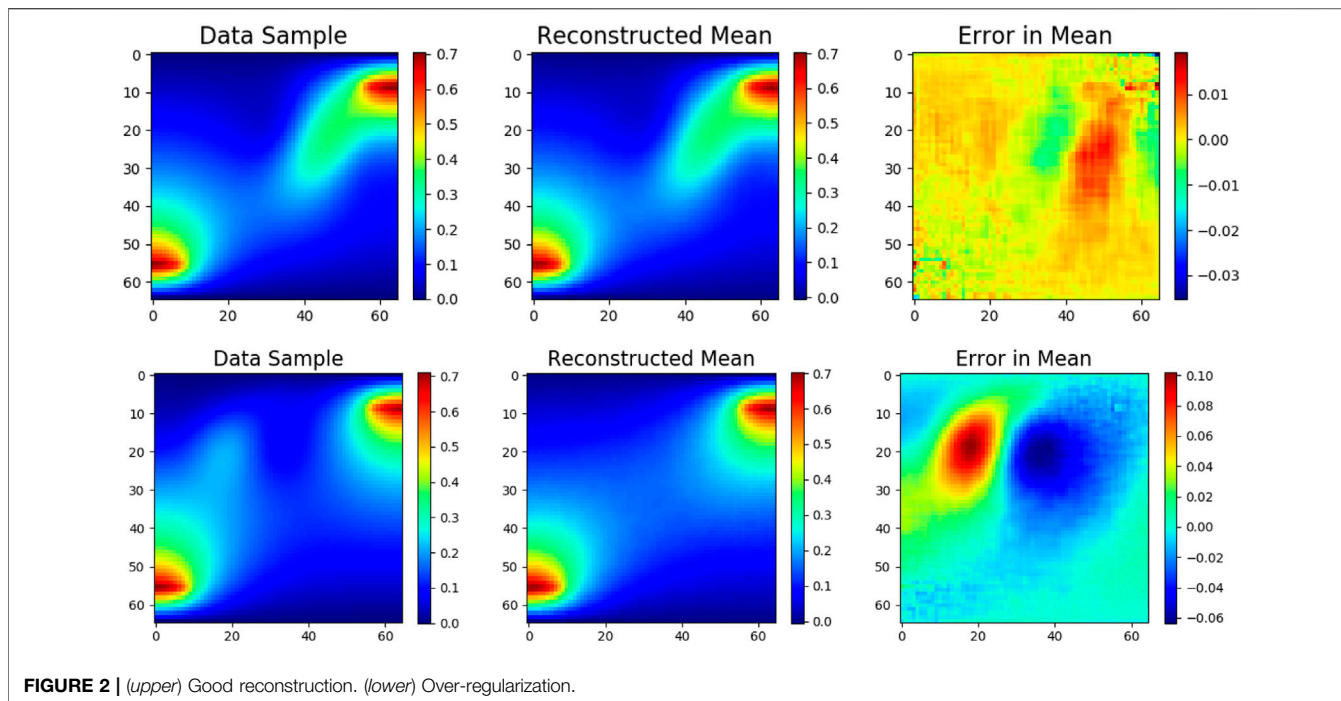
The process of training a VAE involves a number of challenges. For example, convergence of the optimizer to local minima can greatly hinder reconstruction accuracy and failure to converge altogether remains a possibility. A recurrent issue with VAE training in our experiments is that of over-regularization. Over-regularized solutions are characterized by disproportionately small regularization loss ($\mathcal{L}_{REG} \ll 1$). More information on this issue is detailed in **Section 4.2**.

To mitigate some of the issues inherent to training VAEs, we employ a training method tailored to avoid over-regularization. All experiments are performed using the Adam optimizer in Pytorch. We use $\mathcal{L}_{\beta\text{-VAE}}$ to train the models. The β value is varied with epochs, but at the end of training the model is converged with $\beta = 1$. The model is trained initially with $\beta_0 \ll 1$, typically around $\beta_0 = 10^{-7}$, for some number of epochs r_0 (depending on learning rates) until reconstruction accuracy is well below that of an over-regularized solution (**Section 3** illustrates this necessity). When β_0 is too small, the regularization loss can become too large, preventing convergence altogether. Training is continued by implementing a β scheduler [7] to slowly increase the weight of the regularization loss. The learning rate is then decreased to $lr_1 = c(lr_0)$ after some number of epochs r_1 to enhance reconstruction accuracy. This training method—in particular the heavily

weighted reconstruction phase and the β scheduler—result in much more stable training which avoids the local minima characterized by over-regularization and improves convergence consistency. Similar methods have been employed to avoid this issue. In particular, [30] refers to this issue as “KL vanishing” and uses a cyclical β schedule to avoid the issue. However, this can take far more training epochs and cycle iterations to converge than the method employed here.

4.1 Architecture

The primary architecture for the VAE is adapted from [29] and a more detailed description including architecture optimization is given in the included **Supplementary Material**. This architecture consists of a series of encoding blocks to form the encoder, and a series of decoding blocks to form the decoder. Each encoding/decoding block consists of a dense block followed by an encoding/decoding layer. Contrary to the name, dense blocks do not contain any dense layers, but rather a series of skip connections and convolutional layers. Encoding and decoding layers consist of convolutions. The architecture is called DenseVAE and is used for all VAEs trained in this work. The latent and output distributions are assumed to be Gaussian. We use the dense block based architecture to parameterize the encoder mean and log-variance separately, as well as the decoder mean. The decoding distribution log-variance is learned but constant as introducing a learned output log-



variance did not aid in reconstruction or improving disentanglement properties in our experiments but increased training time.

4.2 Over-Regularization

Over-regularization has been identified as a challenge in the training of VAEs [30]. This phenomenon is characterized by the latent space containing no information about the data; i.e., the regularization loss becomes zero. The output of the decoder becomes identical across all inputs. Thus, the output of the decoder is a constant distribution which does not depend on the latent representation. The constant distribution it learns becomes a normal distribution with mean and variance of the data. With zero regularization loss, the learned decoding distribution becomes $p_\psi(y|z) = \mathcal{N}(y; \hat{\mu}_y, \text{diag}(\hat{\sigma}_y^2)) \forall z$ where $\hat{\mu}_y = \frac{1}{N} \sum_{i=1}^N y^{(i)}$ and $\hat{\sigma}_y^2 = \frac{1}{N} \sum_{i=1}^N (y^{(i)} - \hat{\mu}_y)^2$. This is proven to minimize \mathcal{L}_{VAE} in Theorem 1. The solution is not shown to be unique, but our experiments indicate that this is the over-regularized solution found during training. As Theorem 1 illustrates validity for any encoder $q(z|y)$, this is the most robust solution for the VAE to converge to when over-regularization occurs. The decoder learns to predict as accurately as possible given nearly zero mutual information between the latent and data random variables. As the encoder and decoder are trained simultaneously, predicting a constant output regardless of z prevents the necessity of the decoder to adjust as the encoder changes. An empirical comparison between good reconstruction and over-regularization is shown in Figure 2.

Theorem 1 requires that the output variance is constant. Parameterizing the output variance with an additional network may aid in avoiding over-regularization.

THEOREM 1. : Given data $\{y^{(i)}\}_{i=1}^N$ and the VAE framework defined in Section 2, and assuming a decoding distribution of the form $p(y|z) = \mathcal{N}(y; \mu(z), \text{diag}(\sigma^2))$, if $\mathcal{L}_{REG} = 0$, then $\arg \min_{\mu(z), \sigma^2} \mathcal{L}_{VAE} = \{\hat{\mu}_y, \hat{\sigma}_y^2\}$, where $\hat{\mu}_y = \frac{1}{N} \sum_{i=1}^N y^{(i)}$ and $\hat{\sigma}_y^2 = \frac{1}{N} \sum_{i=1}^N (y^{(i)} - \hat{\mu}_y)^2$.

Proof: For any $q(z|y)$ s.t. $\mathcal{L}_{REG} = 0$:

$$\begin{aligned} \mathcal{L}_{VAE} &= \mathcal{L}_{REC} = \mathbb{E}_{\hat{p}(y|q(z|y))} [-\log(p(y|z))] \\ &= \mathbb{E}_{q(z|y)} \left[\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^m \frac{1}{2} \log(2\pi) + \log(\sigma_j) + \frac{1}{2\sigma_j^2} (y_j^{(i)} - \mu_j(z))^2 \right]. \end{aligned}$$

To minimize \mathcal{L}_{VAE} , take derivatives $\frac{\partial \mathcal{L}_{VAE}}{\partial \mu_j(z)}$ and $\frac{\partial \mathcal{L}_{VAE}}{\partial \sigma_j}$ (assuming derivative and expectation can be interchanged), where $j \in \{1, \dots, m\}$:

$$\frac{\partial \mathcal{L}_{VAE}}{\partial \mu_j(z)} = \mathbb{E}_{q(z|y)} \left[-\frac{1}{N} \sum_{i=1}^N \frac{1}{\sigma_j^2} (y_j^{(i)} - \mu_j(z)) \right] = 0.$$

Thus, $\mathbb{E}_{q(z|y)} [\sum_{i=1}^N y_j^{(i)} - \mu_j(z)] = 0$ and

$$\mathbb{E}_{q(z|y)} [\mu_j(z)] = \frac{1}{N} \sum_{i=1}^N y_j^{(i)}. \tag{21}$$

Eq. 21 holds $\forall z, j$ if

$$\mu_j(z) = \hat{\mu}_j = \frac{1}{N} \sum_{i=1}^N y_j^{(i)}. \tag{22}$$

Taking the derivative w.r.t. variance, we have $\frac{\partial \mathcal{L}_{VAE}}{\partial \sigma_j} = \mathbb{E}_{q(z|y)} [\frac{1}{N} \sum_{i=1}^N \frac{1}{\sigma_j} - \frac{1}{\sigma_j^3} (y_j^{(i)} - \mu_j(z))^2] = 0$, and rearranging, we have

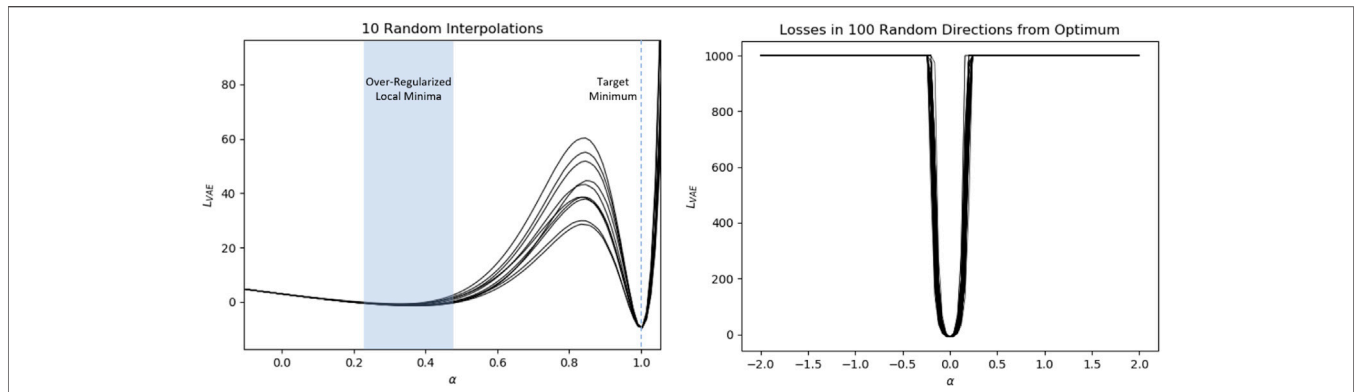


FIGURE 3 | (left) Loss along interpolated lines between 10 random weight initializations and a desirable converged solution. (right) Loss along 100 (of 1,000) random lines emanating from a desirable solution of the DenseVAE architecture. The parameter α indicates the distance along each random direction in parameter space and does not necessarily correspond to the same parameter α in the left figure. Note that the loss is limited to 1,000 for illustration purposes.

$$\sigma_j^2 = \frac{1}{N} \sum_{i=1}^N (y_j^{(i)} - \mu_j(z))^2 \quad \forall z, j. \quad (23)$$

Substituting Eq. 21 into Eq. 23 results in:

$$\hat{\sigma}_j^2 = \frac{1}{N} \sum_{i=1}^N (y_j^{(i)} - \hat{\mu}_j)^2 \quad \forall z, j. \quad (24)$$

With Eqs 21, 24 valid for all z and j , we can combine them into vector form and note that Eq. 25 minimizes \mathcal{L}_{VAE} as required.

$$\hat{\mu}_y = \begin{bmatrix} \hat{\mu}_1 \\ \vdots \\ \hat{\mu}_m \end{bmatrix}, \quad \hat{\sigma}_y^2 = \begin{bmatrix} \hat{\sigma}_1^2 \\ \vdots \\ \hat{\sigma}_m^2 \end{bmatrix}. \quad (25)$$

There exists a region in the trainable parameter loss landscape characterized by over-regularized local minimum solutions which partially surrounds the “desirable” solutions characterized by better reconstruction accuracy and latent properties. This local minima region is often avoided by employing the training method discussed previously, but random initialization of network parameters and changes in hyperparameters between training can render it difficult to avoid convergence to this region.

We illustrate the problem of over-regularization by training VAEs using the architecture described in Section 4.1 on the KLE2 Darcy flow dataset with $p(\theta)$ being standard normal.

A VAE is trained with 512 training samples (each sample is $65 \times 65 \times 3$), converging to a desirable solution with low reconstruction error and nearly perfect disentanglement. The parameters of this trained network are denoted P_T . After the VAE is trained and a “desirable” solution obtained, 10 additional VAEs with identical setup to the desirable solution are initialized randomly using the Xavier uniform weight initialization on all layers. Each of the 10 initializations contain parameters P_i . A line in the parameter space is constructed between the converged “desirable” solution and the initialized solutions as a function of α :

$$P(\alpha) = (1 - \alpha)P_i + \alpha P_T. \quad (26)$$

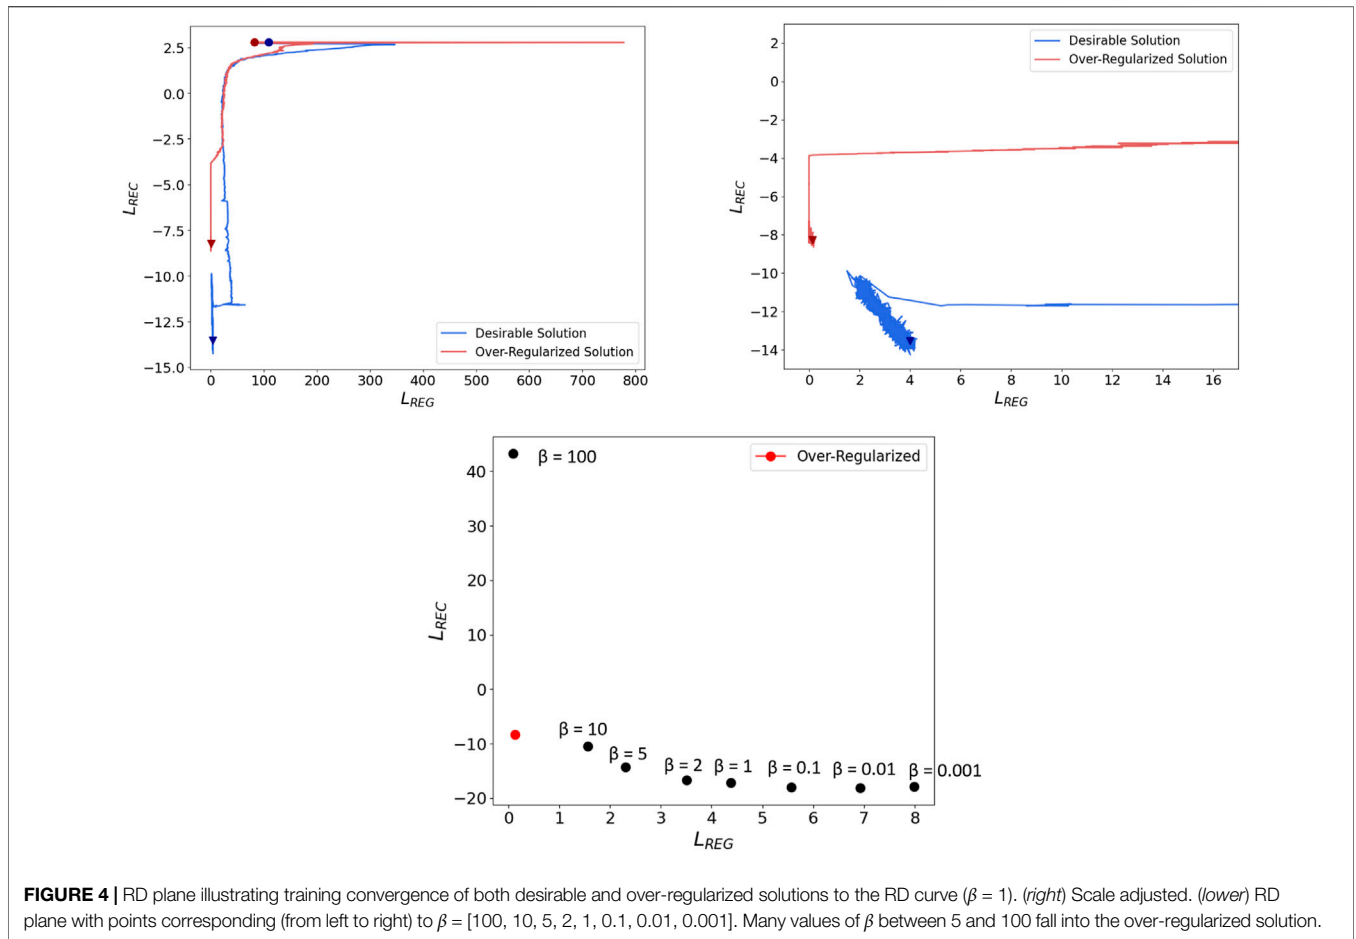
Losses are recorded along each of the 10 interpolated lines and plotted in Figure 3. Between the random initializations and “desirable” converged solutions there exists a region of local minima in the loss landscape, and these local minima are characterized by over-regularization. Losses illustrated are computed as an expectation over all training data and a Monte Carlo estimate of the reconstruction loss with 10 latent samples to limit errors due to randomness.

We include an illustration of the avoidance of these over-regularized local minima using our training method in the **Supplementary Material**.

Of interest is that this over-regularized local minima region does not fully surround the “desirable” region. Instead of interpolating in parameter space between random initializations and a converged solution, lines emanating away from the converged solution along 1,000 random directions in parameter space are created and the loss plotted along each. Figure 3 illustrates that indeed no local minima are found around the converged solution. We note that there are around 800,000 training parameters in this case, so 1,000 random directions may not completely encapsulate the loss landscape around this solution.

The Xavier uniform weight initialization scheme, and most other initialization schemes, limit the norm of the parameters in parameter space to near the origin. The local minima region exists only between the converged solution region and points in parameter space near the origin. In this case, there may be alternative initialization schemes which can greatly aid in the convergence of VAEs. This has been observed in [31] where the initialization scheme proposed greatly accelerates the speed of convergence and accuracy of reconstruction.

Over-regularized local minima follow a similar path during training as desirable solutions. A region of attraction exists in the loss landscape, and falling too close to this region will result in an over-regularized solution, illustrated in Figure 4. One VAE which obtains a desirable solution shares a similar initial path with an over-regularized solution. Plotted are the VAE losses computed



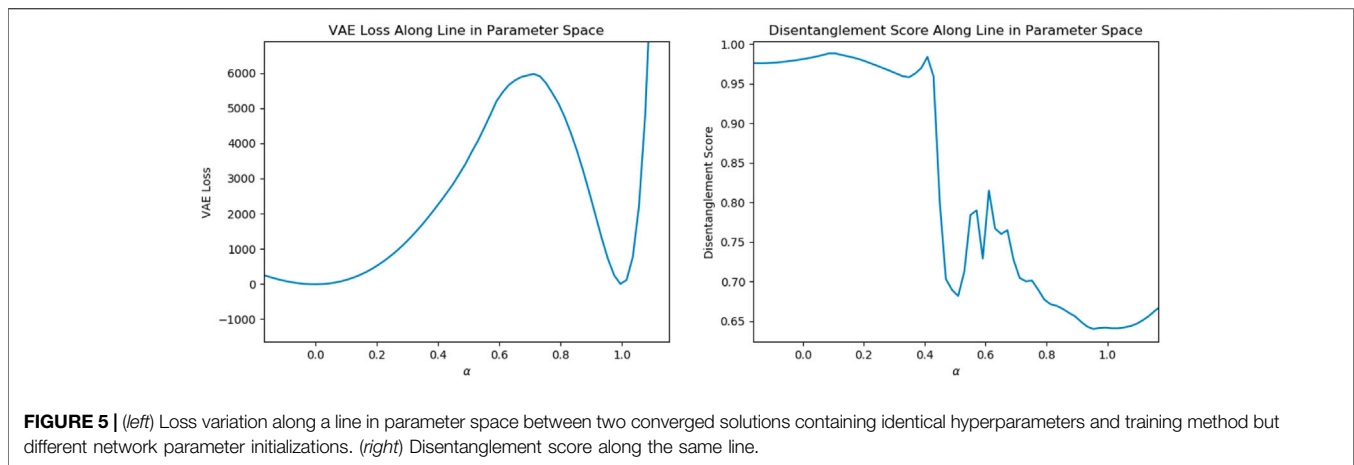
during training, not the training losses. The over-regularized solution breaks from the desired path too early, indicating a necessity for a longer reconstruction-heavy phase.

Training many VAEs with various β values facilitates a visualization of over-regularization in the RD plane. Each point in **Figure 4** shows the loss values of converged VAEs trained with different values of β . The over-regularized region of attraction prevents convergence to desirable solutions for many values of β . Interpolating in parameter space between each of these points (corresponding to a VAE with its own converged parameters) using the base VAE loss ($\beta = 1$), no other points on the RD curve are local minima of the VAE loss. In **Figure 4**, we observe that during training, the desirable solution reaches the RD curve but continues toward the final solution.

4.3 Properties of Desirable Solutions

Avoiding over-regularization aids in convergence to solutions characterized by low reconstruction error. Among solutions with similar final loss values, inconsistencies remain in latent properties. Two identical VAEs initialized separately often converge to similar loss values, but one may exhibit disentanglement while the other does not. This phenomenon is also explored in [18] and [16]. Two VAEs are trained with identical architectures, hyperparameters, and training method;

they differ only in the random initialization of network parameters P . We denote the optimal network parameters found from one initialization as P_1 and optimal network parameters found from a separate initialization P_2 . The losses for each converged solution are quite similar ($L_{VAE_1} \approx -9.50$, $L_{VAE_2} \approx -9.42$); however, disentanglement properties of each are dramatically different. We interpolate between these two solutions in parameter space (Eq. 26) and record losses and disentanglement scores along the line (Figure 5). The first network contains a nearly perfectly disentangled latent representation while the second network does not produce a disentangled representation. It is evident that multiple local minima exist in parameter space which converge to similar values in the loss landscape, but contain very different latent correlations. Local minima exist throughout the loss landscape, and with each initialization, a different local minimum may be found. Many such differing solutions are found throughout our experiments. This phenomenon is partially due to invariance of the ELBO to rotations of the latent space when using rotationally invariant priors. Disentanglement is heavily dependent on a factorized representation of the latent representation. With rotations not affecting the training loss, learning a disentangled representation seems to be somewhat random in this case.



This phenomenon exhibits the difficulties in disentangling generative parameters in an unsupervised manner; without prior knowledge of the factors of variation, conclusions cannot be drawn regarding disentanglement by observing loss values alone. In controlled experiments, knowledge of the underlying factors of variation is available, but when only data is available, full knowledge of such factors is often not. It is encouraging that the VAE does have the power to disentangle generative parameters in an unsupervised setting, but the nature of disentanglement must first be understood better to create identifying criterion.

5 CHARACTERIZING DISENTANGLEMENT

In this section, we explore the relationship between disentanglement, the aggregated posterior ($q_\phi(z)$), and the generative parameter distribution ($p(\theta)$) by incrementally increasing the complexity of $p(\theta)$. Disentanglement is first illustrated to be achievable but difficult using the classic VAE assumptions and loss due to a lack of enforcement of the rotation of the latent space caused by rotationally-invariant priors. Hierarchical priors are shown to aid greatly in disentangling the latent space by learning non-rotationally-invariant priors which enforce a particular rotation of the latent space through the regularization loss.

5.1 Standard Normal Generative Distributions

The intrinsic dimensionality of the data is set to $p = 2$ with a generative parameter distribution $p(\theta) = \mathcal{N}(\theta; 0, I_{2 \times 2})$, the standard normal distribution. Limiting p to 2 aids greatly in the visualization of the latent space and understanding of the ideas investigated. The standard latent prior is identical to the generative parameter in this case, creating a relatively simple problem for the VAE.

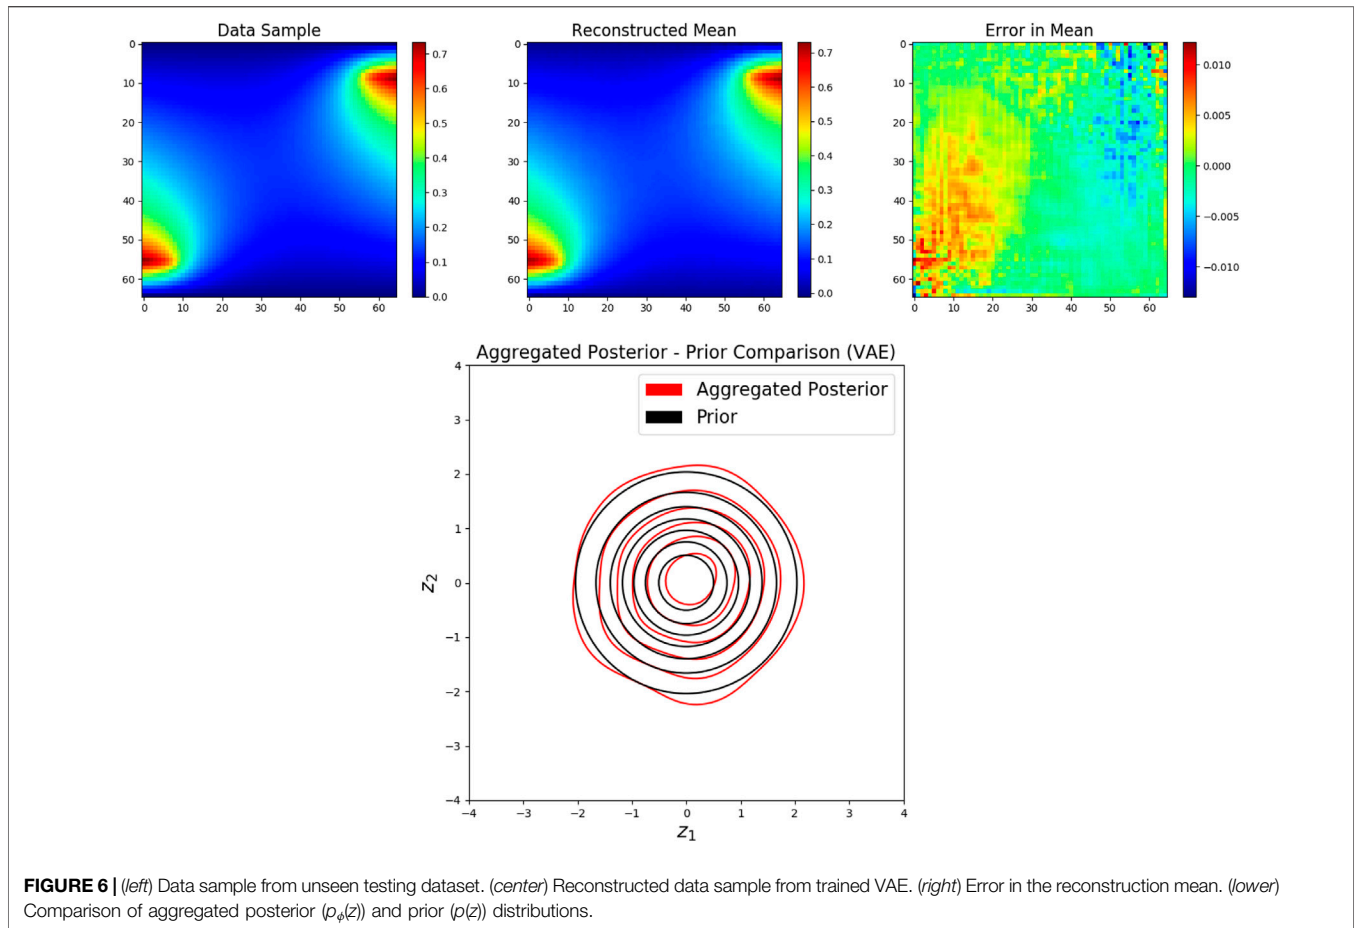
Using the architecture described in Section 4.1, the relationship between regularization, reconstruction, and disentanglement and the number of training samples is illustrated in the included **Supplementary Material**. A similar

study is performed in [18] with a greater sample size. Reconstruction losses continue to fall with the number of training data, indicating improved reconstruction of the data with increased number of samples; however, the regularization loss increases slightly with the number of training data. With too few samples, reconstruction performance is very poor and over-regularization (near zero regularization loss) seems unavoidable. Clear and consistent correlations exist among the loss values and number of training data, but disentanglement properties vary greatly among converged VAEs (Section 4.3). The compressed representations range from nearly perfect disentanglement to nearly completely entangled.

Although disentanglement properties are inconsistent between experiments, desirable properties of disentanglement are often observed. Training is performed using the maximum amount of available data (512 snapshots), and analysis included for 512 testing samples on the KLE2 dataset (regardless of $p(\theta)$). Regularization loss is large during the reconstruction phase in which $\beta_0 = 10^{-7}$, and the y-axis is truncated for clarity. A comparison between a test data sample and the reconstructed mean using the trained VAE is depicted in Figure 6, showing little error between the mean $\mu_\psi(z)$ of the decoding distribution and the input data sample. With small reconstruction error, a disentangled latent representation is learned. Figure 6 also illustrates the aggregated posterior matching the prior distribution in shape. This is unsurprising with a generative parameter and prior distribution match and an expressive network architecture. Finally, Figure 7 shows the correlation between the generative parameters of the training and testing data against the latent distribution as a qualitative measure of disentanglement. Each latent dimension is tightly correlated to a single but different generative parameter. Figure 7 also illustrates the uncertainty in the latent parameters, effectively $q_\phi(z|\theta)$. The latent representation is fully disentangled; each latent parameter contains only information about a single generative factor.

5.2 Non Standard Gaussian Generative Distributions

The generative parameter distribution and the prior are identical (independent standard normal) in the previous example. Most

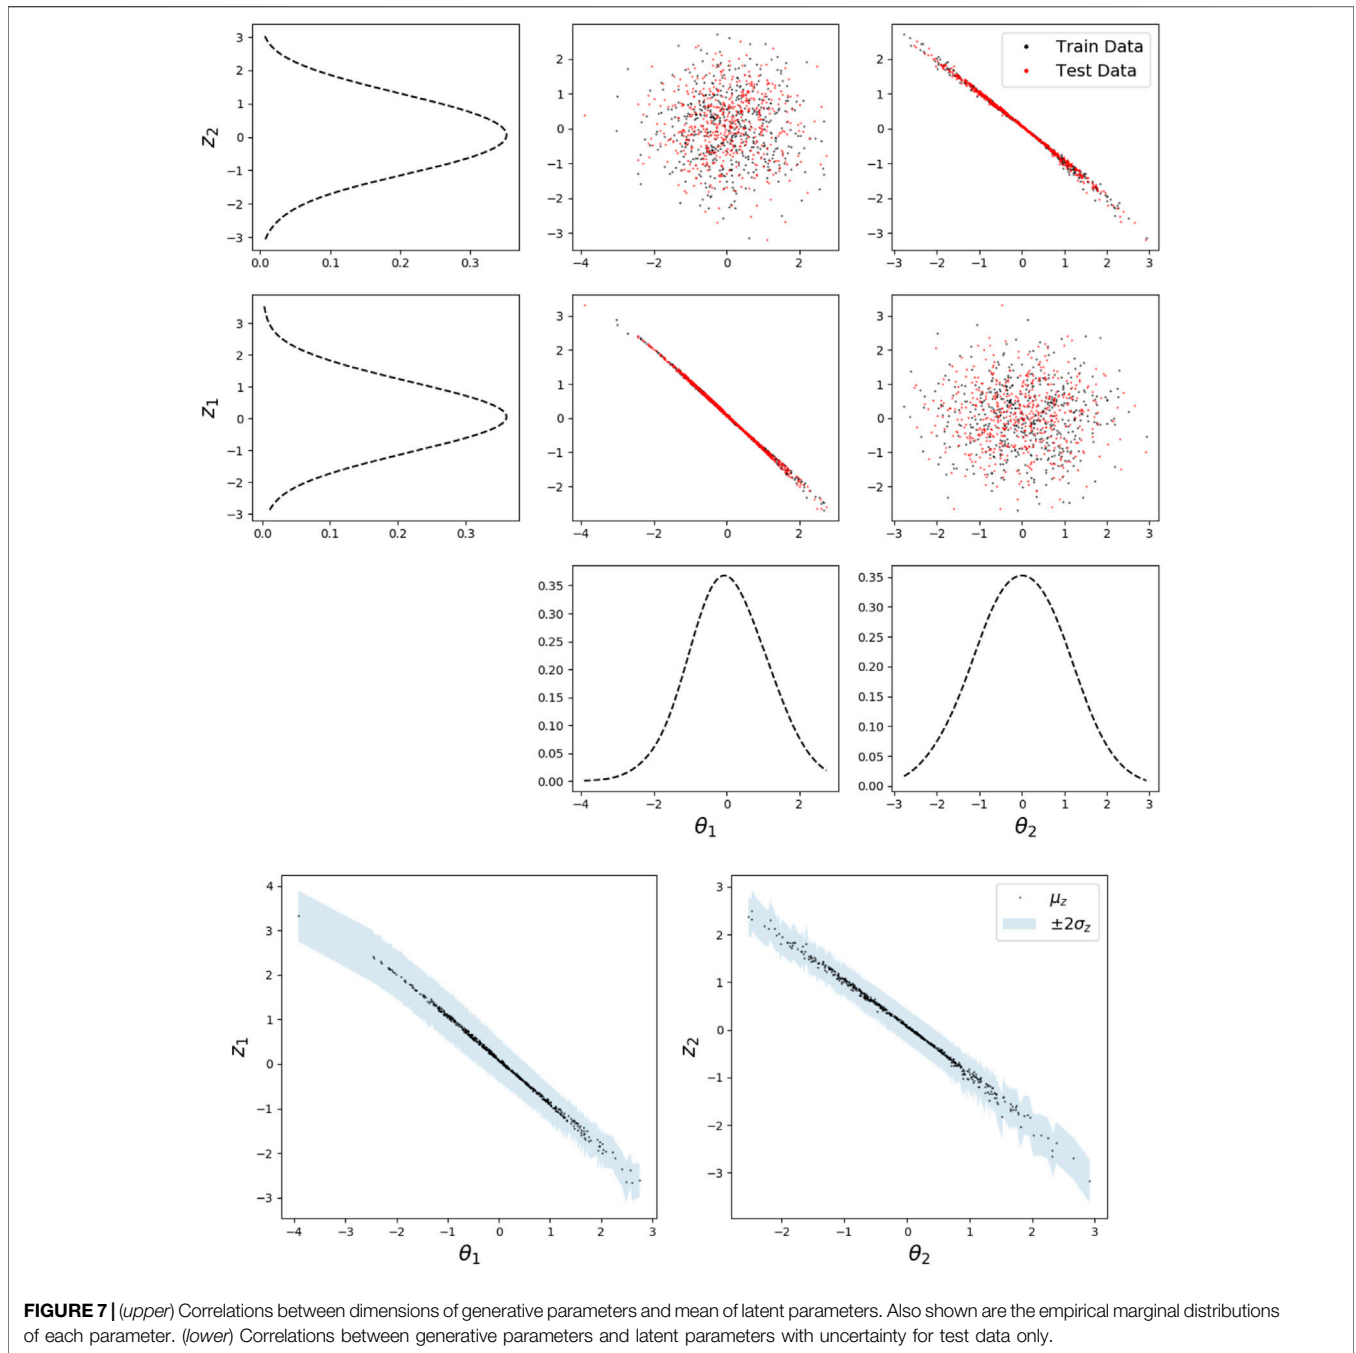


often, however, knowledge of the generative parameters is not possessed. The specified prior in this case is unlikely to match the generative parameter distribution. The next example illustrates the application of a VAE in which the generative parameter distribution and prior do not match. Another KLE2 dataset is generated with a non standard Gaussian generative parameter distribution. The generative parameter distribution is Gaussian, but scaled and translated relative to the previous example $p(\theta) = \mathcal{N}(\theta; [1; 1], [0.5, 0; 0, 0.5])$.

Training a standard VAE on this dataset results in high reconstruction accuracy, but undesirable disentanglement after many trials. With the use of an additional hierarchical prior network, good disentanglement can be achieved even with a mismatch in the prior and generative parameter distributions. The sub-prior (see Section 2.3) is the standard normal distribution, but the hierarchical network learns a non-standard normal prior. Still, the learned prior and generative parameter distributions do not match. Figures 8, 9 illustrate comparisons in results obtain from the VAE with and without the hierarchical prior network. When using hierarchical priors, the learned prior and aggregated posterior match reasonably well but do not match the generative parameter distribution. However, this does not matter as long as the latent representation is not rotated relative to the generative parameter distribution, as illustrated in the next example. Low reconstruction error and

disentanglement are observed using hierarchical priors, but disentanglement was never observed using the standard VAE after many experiments. This may be because β is not large enough to enforce a regularization loss large enough to produced an aggregated posterior aligned with the axes of the generative parameter distribution. Therefore, the rotation of the learned latent representation will be random and disentanglement is unlikely to be observed, even in two dimensions. The hierarchical network consistently enforces a factorized aggregated posterior, which is essential for disentanglement when generative parameters are independent. One potential cause of this is the learning of non-rotationally-invariant priors, such as a factorized Gaussian with independent scaling in each dimension. The ELBO loss in this case is affected by rotations of the latent space, aligning the latent representations to the axes of the generative parameters.

A latent rotation can be introduced such that the reconstruction loss is unaffected, but regularization loss changes with rotation. Introducing a rotation matrix A with angle of rotation ω to rotate the latent distribution, the encoding distribution becomes $q_\phi(z|y) = \mathcal{N}(z; A\mu_\phi(y), \text{Adiag}(\sigma_\phi(y))A^T)$. Reversing this rotation when computing the decoding distribution (i.e., $p_\psi(y|z) = \mathcal{N}(y; \mu_\phi(A^T z), \text{diag}(\sigma_\psi(A^T z)))$) preserves the reconstruction loss. However, the regularization loss can be



plotted as a function of the rotation angle (included in the **Supplementary Material**). When a rotationally-invariant prior is used to train the VAE, regularization loss is unaffected by latent rotation. However, when the prior is non-rotationally-invariant, the regularization loss is affected by latent rotation. Thus, rotation of the latent space is enforced by the prior during training.

Although the hierarchical prior adds some trainable parameters to the overall architecture, the increase is only 0.048%. This is negligible, and it is assumed that this is not the root cause of improved disentanglement. Rather, it is the ability of the additional hierarchical network to consistently

express a factorized aggregated posterior and learn non-rotationally-invariant priors which improves disentanglement. More insights are offered in the next example and **Section 7**.

5.3 Multimodal Generative Distributions

In this setup, disentanglement not only depends on a factorized $q_\phi(z)$, but the correlations in $p(\theta)$ must be preserved as well, i.e., rotations matter. The previous example illustrates a case in which the standard VAE fails in disentanglement but succeeds with the addition of hierarchical priors due to improved enforcement of a factorized $q_\phi(z)$ through learning non-

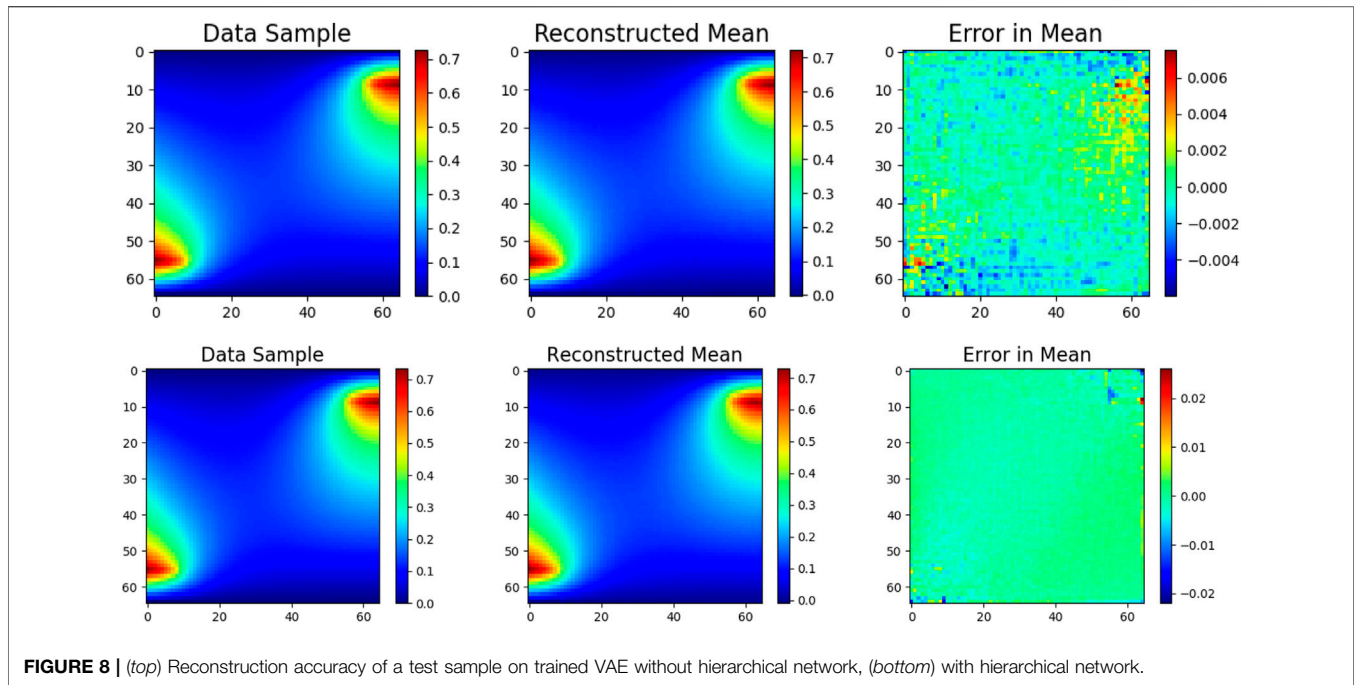


FIGURE 8 | (top) Reconstruction accuracy of a test sample on trained VAE without hierarchical network, (bottom) with hierarchical network.

rotationally-invariant priors. The generative parameter distribution is radially symmetric, thus visualization of rotations in $q_\phi(z)$ relative to $p(\theta)$ is difficult. To illustrate the benefits of using hierarchical priors for disentanglement, the final example uses data generated from a more complex generative parameter distribution with four lines of symmetry for better visualization. The generative parameter distribution is multimodal (a Gaussian mixture) and is more difficult to capture than a Gaussian distribution, but allows for better rotational visualization:

$$\begin{aligned}
 p(\theta) = & \frac{1}{4} \mathcal{N}(\theta; [-1; -1], [0.25, 0; 0, 0.25]) \\
 & + \frac{1}{4} \mathcal{N}(\theta; [1; 1], [0.25, 0; 0, 0.25]) \\
 & + \frac{1}{4} \mathcal{N}(\theta; [-1; 1], [0.25, 0; 0, 0.25]) \\
 & + \frac{1}{4} \mathcal{N}(\theta; [1; -1], [0.25, 0; 0, 0.25]).
 \end{aligned}$$

Training VAEs without hierarchical priors results in over-regularization more often than with the implementation of HP. Out of 50 trials, 10% trained without HP were unable to avoid over-regularization while all trials with HP successfully avoided over-regularization. More epochs are required in the reconstruction only phase (with and without HP) to avoid over-regularization than in previous examples. Disentanglement was never observed without the use of hierarchical priors. This again is due to rotation of the latent space relative to the generative parameter distribution due to rotationally invariant priors. To illustrate this concept, **Figure 10** illustrates the effects of rotation of the latent space on disentanglement. Clearly, rotation dramatically impacts disentanglement, and the

standard normal prior does not enforce any particular rotation of the latent space.

Implementing hierarchical priors, consistent observation of not only better reconstruction (avoiding over-regularization) but also reasonable disentanglement of the latent space in roughly half of all trained VAEs (out of 50) exemplifies the improved ability of hierarchical priors to produce a disentangled latent representation. Reconstruction of test samples is more accurate when implementing the hierarchical prior network, as illustrated in **Figure 11**. We hypothesize that disentanglement is observed in roughly half of our experiments due to local minima in the regularization loss corresponding to 45° rotations of the latent space, illustrated in the **Supplementary Material**. The learned priors using HP are often non-rotationally-invariant and aligned with the axes. However, the posterior is often rotated 45-degrees relative to this distribution, creating a non-factorized and therefore non-disentangled representation.

Comparing $p(\theta)$, $p(z)$, and $q_\phi(z)$ with and without HP (**Figure 12**), stark differences are noticeable. Without the HP network, the aggregated posterior often captures the multimodality of the generative parameter distribution, but it is rotated relative to $p(\theta)$, creating a non-factorized $q_\phi(z)$. Training the VAE with hierarchical priors, the learned prior becomes non-rotationally invariant. The rotation of the aggregated posterior is therefore controlled by the orientation of the prior through the regularization loss, but mimics the shape of the generative parameter distribution. It is clear that the prior plays a significant role in terms of disentanglement: it controls the rotational orientation of the aggregated posterior.

A qualitative measure of disentanglement is compared in **Figure 12**. Without HP, the latent parameters are entangled; they are each weakly correlated to both of the generative parameters. Adding HP to the VAE results in disentanglement

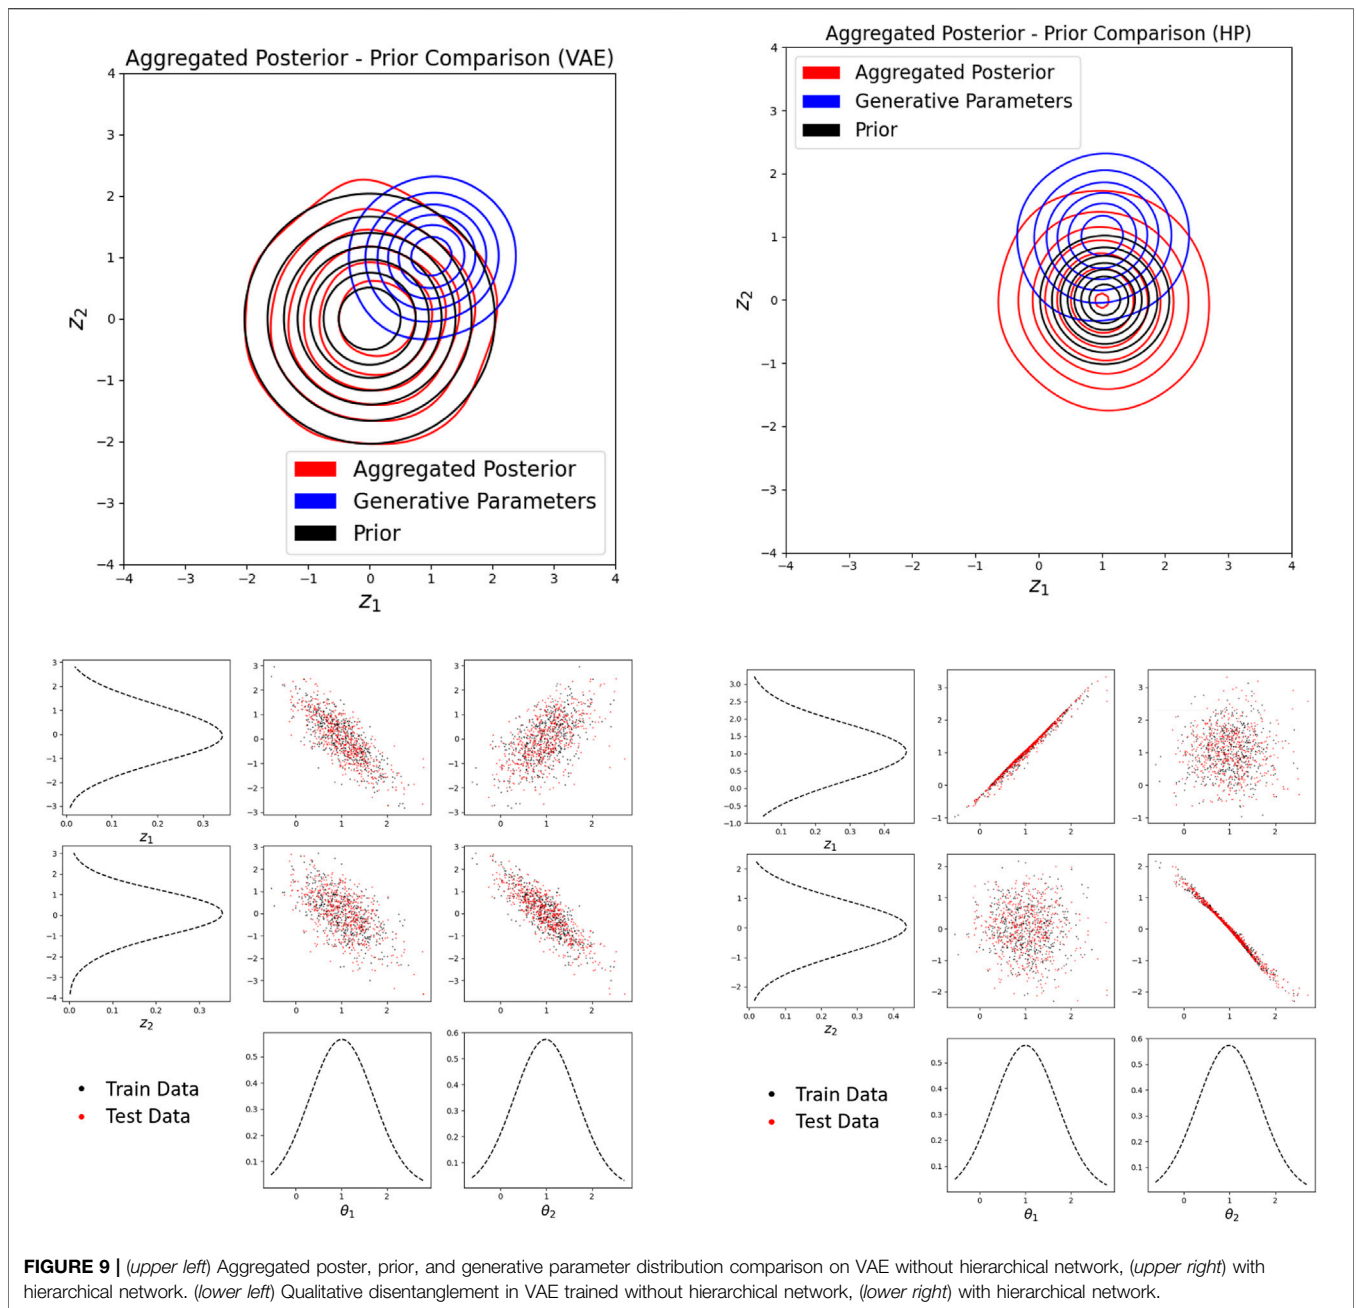


FIGURE 9 | (upper left) Aggregated poster, prior, and generative parameter distribution comparison on VAE without hierarchical network, (upper right) with hierarchical network. (lower left) Qualitative disentanglement in VAE trained without hierarchical network, (lower right) with hierarchical network.

in nearly half of our trials. When disentanglement does occur, each latent factor contains information on mostly a single but different generative factor. Through the course of our experiments, a relationship between disentanglement and the degree to which the aggregated posterior matches the generative parameter distribution is recognized. When disentanglement does not occur with the use of HP, the aggregated posterior is rotated relative to $p(\theta)$, or non-factorized (it has always been observed at around a 45-degree rotation). Only when $q_\phi(z)$ can be translated and scaled to better match $p(\theta)$, maintaining the correlations, does disentanglement occur. Thus, a quantitative measure of disentanglement (Eq. 7) is created from this idea. The

KL divergence is estimated through sampling using the k -nearest neighbors (k -NN) approach (version $\epsilon 1$) found in [19]. The optimization is performed using the gradient-free Nelder-Mead optimization algorithm [32].

In low-dimensional problems, humans are adept at determining disentanglement from qualitative measurements of disentanglement such as Figure 12. It is, however, more difficult to obtain quantitative measurements of these properties. Figure 13 shows the relationship between Eq. 7 and a qualitative measurement of disentanglement. Lower values of S_{KL} indicate better disentanglement. This measure of disentanglement and

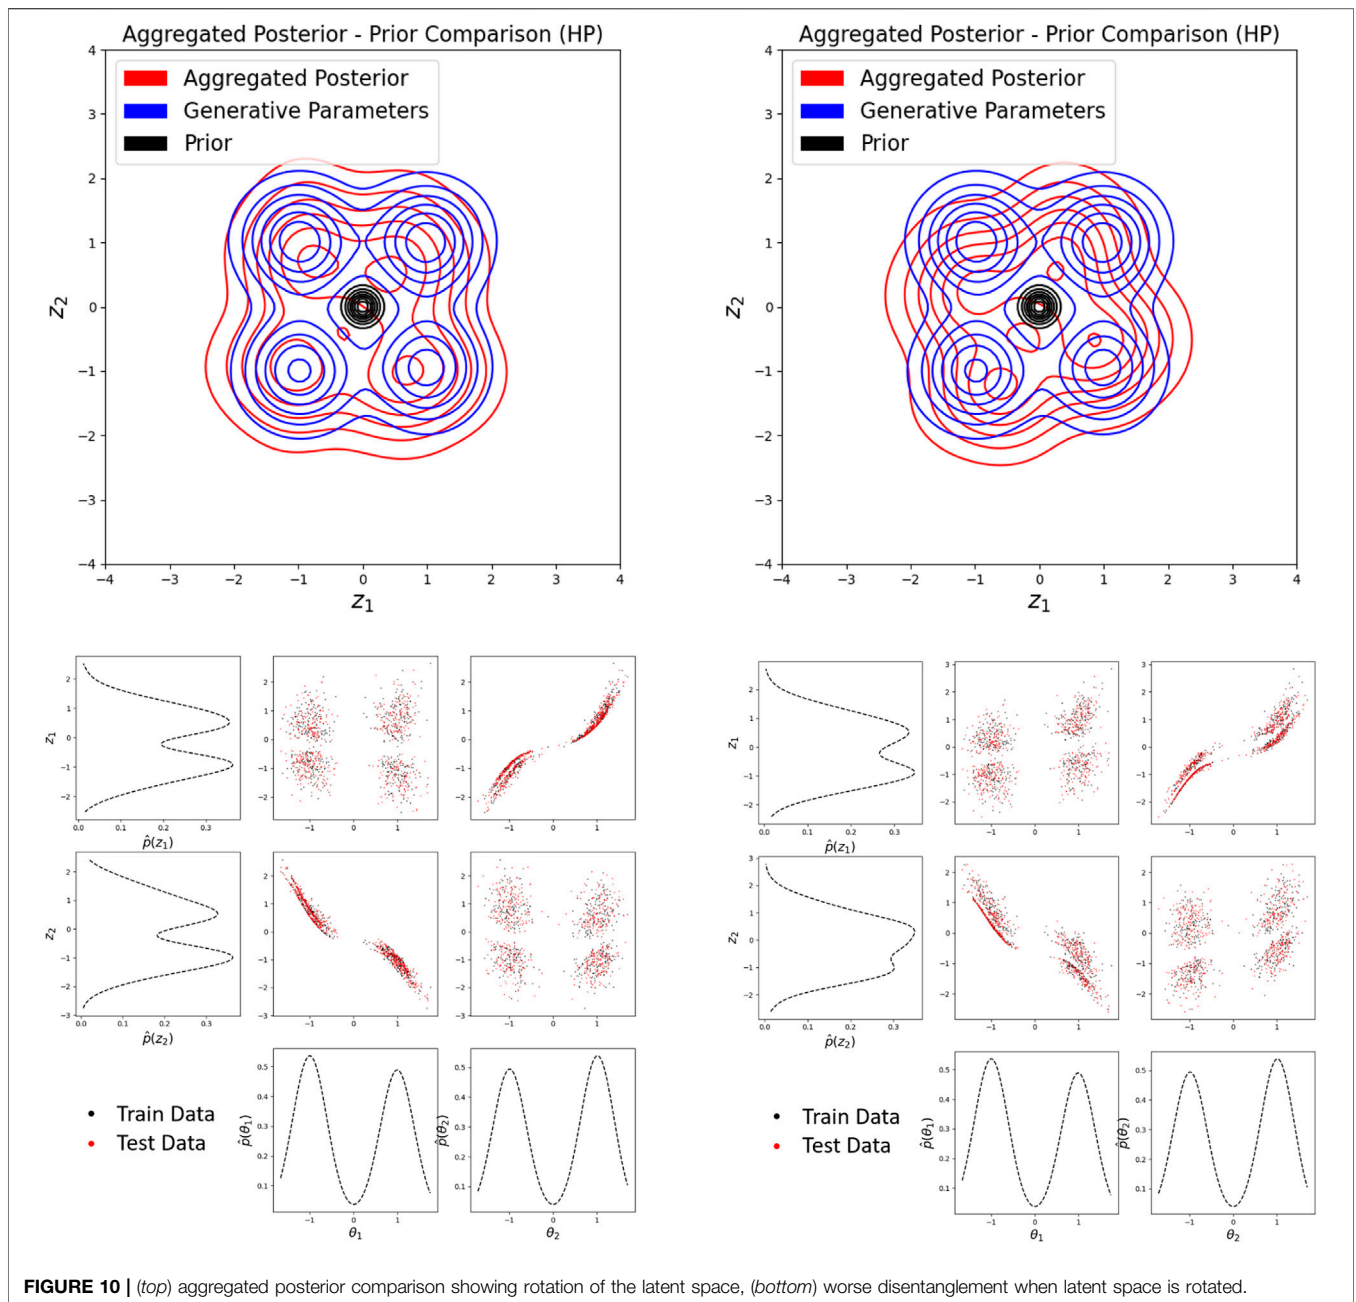


FIGURE 10 | (top) aggregated posterior comparison showing rotation of the latent space, (bottom) worse disentanglement when latent space is rotated.

the intuition behind it is discussed further in the conclusions section.

6 SEMI-SUPERVISED TRAINING

Difficulties with consistently disentangling generative parameters have been illustrated up to this point with an unsupervised VAE framework. In some cases, however, generative parameters may be known for some number of samples, suggesting the possibility of a semi-supervised approach. These labeled samples can be leveraged to further

improve the consistency of learning a disentangled representation. Consider data consisting of two partitions: labeled data $\{y^{(i)}, \theta^{(i)}\}_{i=1}^l$ and unlabeled data $\{y^{(i)}\}_{i=l+1}^{u+l}$. A one-to-one mapping between the generative parameters θ and the learned latent representation z is sought when disentanglement is desired. Thus, enforcing the latent representation to match the generative parameters for labeled data in a semi-supervised approach should aid in achieving our desired objective more consistently.

We begin the intuition behind a semi-supervised loss function by illustrating its connection to the standard ELBO VAE loss. One method of deriving the ELBO loss is to first expand the

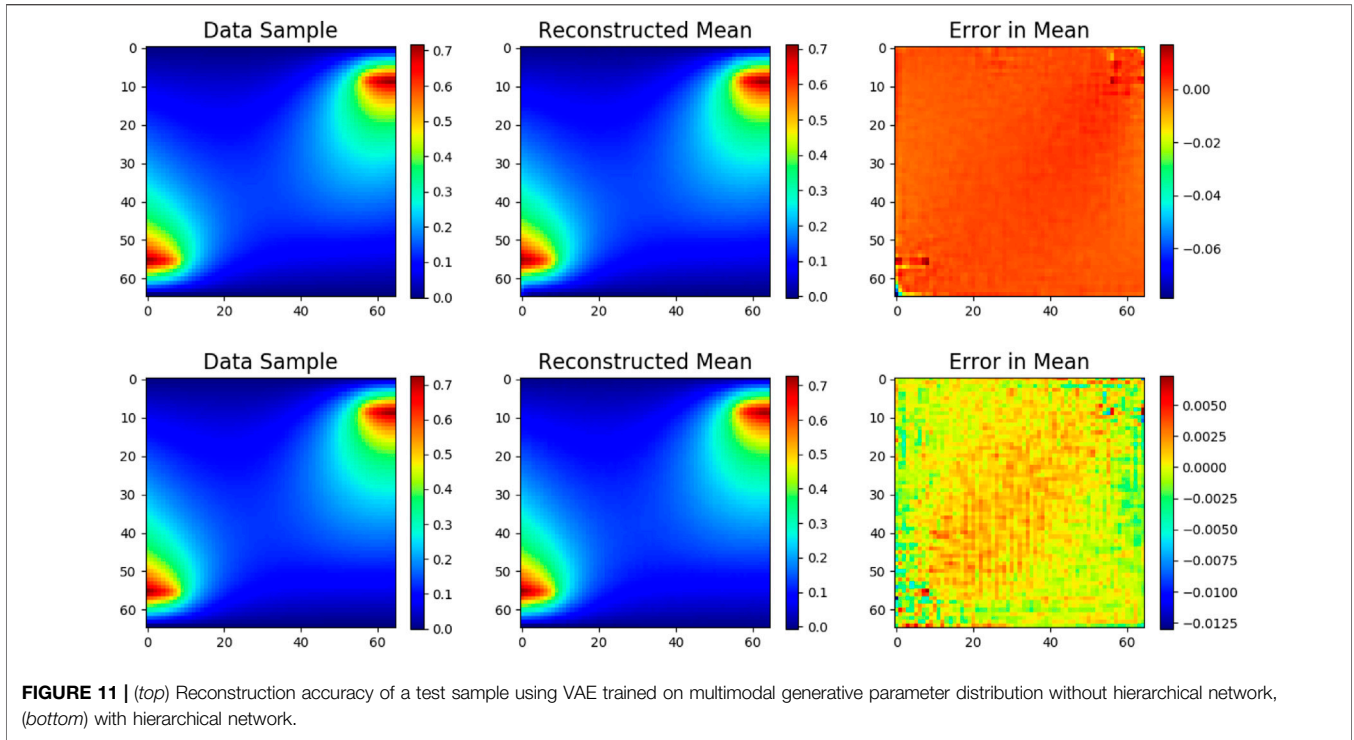


FIGURE 11 | (top) Reconstruction accuracy of a test sample using VAE trained on multimodal generative parameter distribution without hierarchical network, (bottom) with hierarchical network.

relative entropy between the data distribution and the induced data distribution to obtain

$$D_{KL}[p(y)||p_{\psi}(y)] = -H(Y) + \mathbb{E}_{p(y)}[D_{KL}[q_{\phi}(z|y)||p(z)]] - \mathbb{E}_{p(y)}[D_{KL}[q_{\phi}(z|y)||p(z|y)]] - \mathbb{E}_{p(y)q_{\phi}(z|y)}[\log p_{\psi}(y|z)]$$

where $-H(Y)$ is constant and the “true” encoder $p(z|y)$ is unknown. Therefore, the term

$$\mathbb{E}_{p(y)}[D_{KL}[q_{\phi}(z|y)||p(z|y)]]$$

is usually ignored and we arrive at the ELBO, which upper bounds the left hand side. However, a relationship between z and y is known for labeled samples. This relationship can be used to assign $p(z^{(i)}|y^{(i)})$ on the labeled partition. For unlabeled data, the standard ELBO loss is still used for training and the semi-supervised loss to be minimized becomes

$$\mathcal{L}_{VAE-SS}(\phi, \psi) = \mathbb{E}_{p(y)}[D_{KL}[q_{\phi}(z|y)||p(z)]] - \mathbb{E}_{p_l(y)}[D_{KL}[q_{\phi}(z|y)||p(z|y)]] - \mathbb{E}_{p(y)q_{\phi}(z|y)}[\log p_{\psi}(y|z)] \quad (27)$$

where $p_l(y)$ is the distribution of inputs with corresponding labels, $p(y)$ is the distribution of all inputs (labeled and unlabeled), and $\mathbb{E}_{p_l(y)}[D_{KL}[q_{\phi}(z|y)||p(z|y)]]$ is denoted \mathcal{L}_{SS} .

Note that the term $\mathbb{E}_{p(y)}[D_{KL}[q_{\phi}(z|y)||p(z)] - D_{KL}[q_{\phi}(z|y)||p(z|y)]]$ is minimized by $q_{\phi}(z|y) = p(z|y)$ at $I(Z; Y)$, the mutual information between the generative parameters and high dimensional data. In unsupervised VAEs, the regularization term $\mathbb{E}_{p(y)}[D_{KL}[q_{\phi}(z|y)||p(y)]]$ is minimized

when $q_{\phi}(z|y) = p(y)$. As observed in previous sections, disentanglement is observed when the aggregated posterior is “close” to the generative parameter distribution. With the semi-supervised loss being minimized when they are equivalent, the learned latent representations should be more easily and consistently disentangled.

However, empirically it is found that this loss is very sensitive to changes in network parameters and unreasonably small learning rates are required for stability. Additionally, there is no obvious way to determine the variance of $p(z^{(i)}|y^{(i)})$ for each sample, only the mean is easily identifiable. We therefore propose to train with $\mathcal{L}_{SS} = \mathbb{E}_{p_l(y)}[-\log q_{\phi}(z|y)]$ instead such that the loss function becomes

$$\mathcal{L}_{VAE-SS}(\phi, \psi) = \mathbb{E}_{p(y)}[D_{KL}[q_{\phi}(z|y)||p(z)]] - \mathbb{E}_{p_l(y)}[\log q_{\phi}(z|y)] - \mathbb{E}_{p(y)q_{\phi}(z|y)}[\log p_{\psi}(y|z)]. \quad (28)$$

Training with this loss achieves the desired outcome of consistently learning disentangled representations while being simple and efficient to implement.

Incorporating some labeled samples into training the VAE, a disentangled latent representation can be consistently learned. **Figure 14** illustrates the relationship between increasing the number of labeled samples and the disentanglement score of the learned latent representation. In each case, there are 512 unlabeled samples. Each trial varies in the number of labeled samples, and VAEs trained with the same number of labeled samples are trained with a different set of labeled samples. Ten VAEs are trained at each point, and the range illustrated

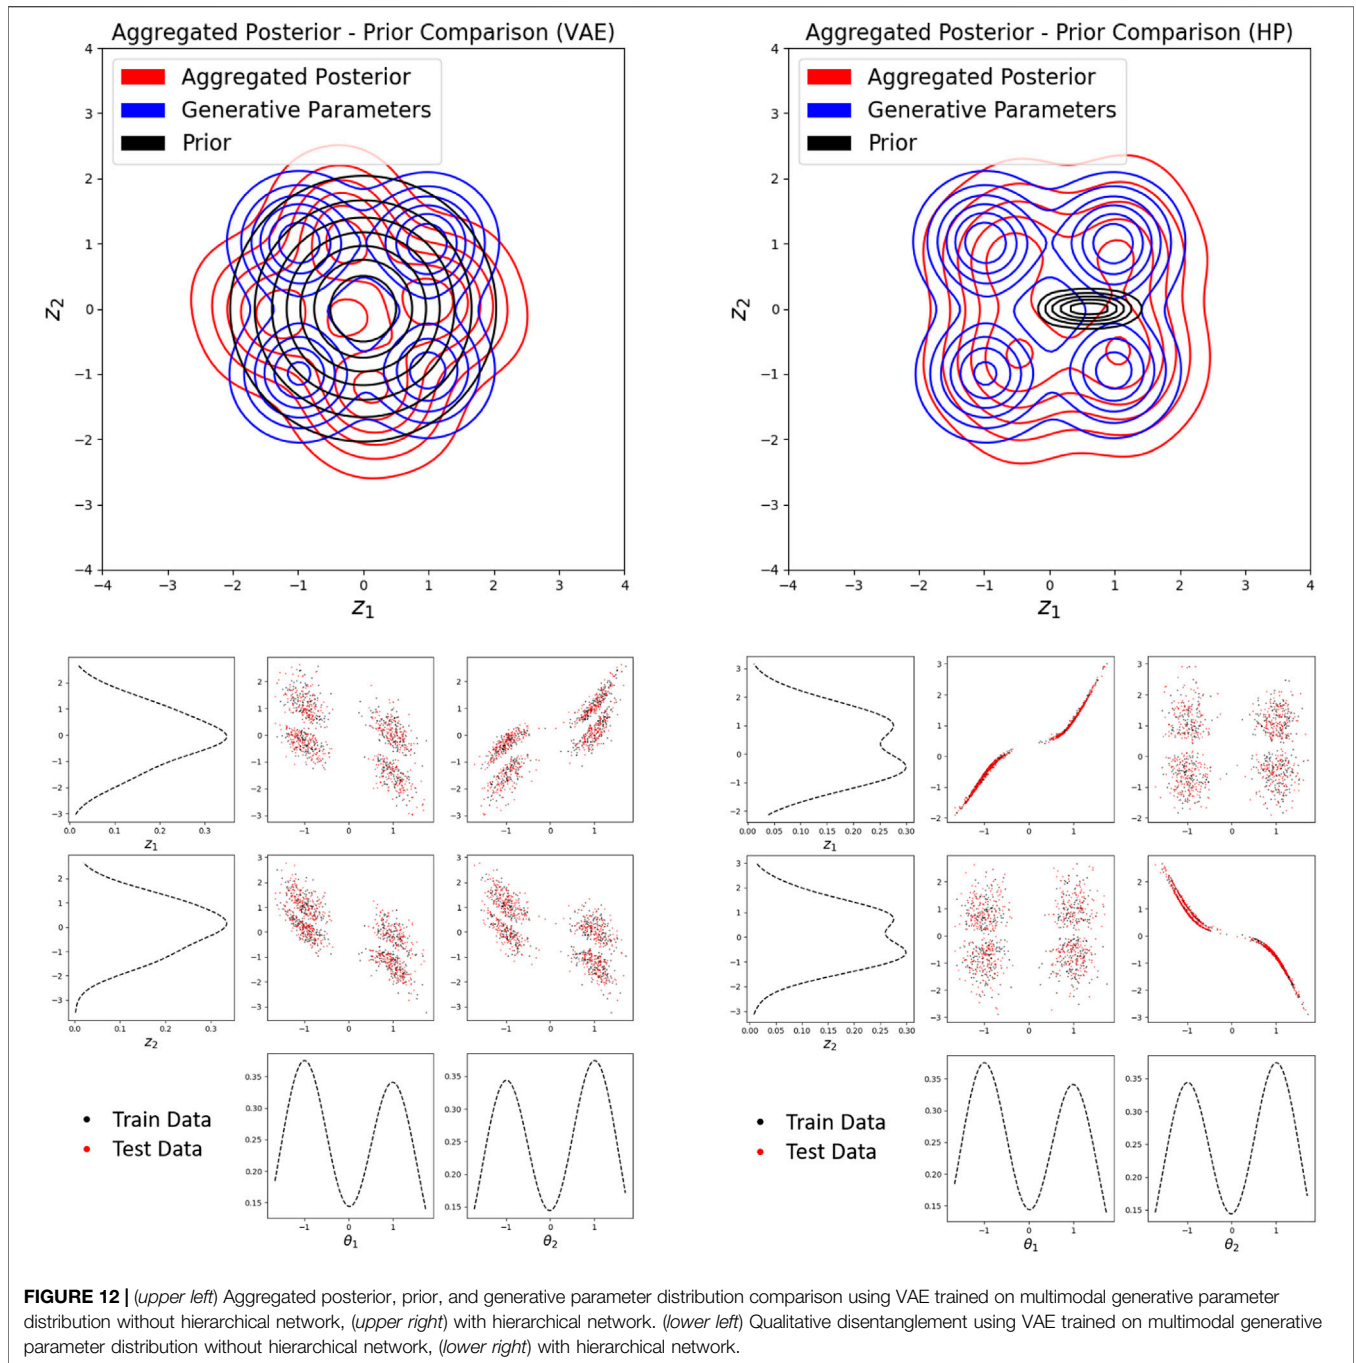


FIGURE 12 | (upper left) Aggregated posterior, prior, and generative parameter distribution comparison using VAE trained on multimodal generative parameter distribution without hierarchical network, (upper right) with hierarchical network. (lower left) Qualitative disentanglement using VAE trained on multimodal generative parameter distribution without hierarchical network, (lower right) with hierarchical network.

represents the maximum and minimum disentanglement score across the 10 trials.

The training losses do not seem to be effected by the number of labeled samples, only the disentanglement score is effected. With a low number of labeled samples, the semi-supervised VAE trains very similarly to the unsupervised VAE. That is, disentanglement is observed rather randomly, and the learned latent representation varies dramatically between trials. Labeling around 1% of the samples begins to result in consistently good disentanglement. Labeling between 3% and 8% results in learning

disentangled latent representations which are nearly identical between trials. It follows from these results that disentangled representations can be consistently learned when training with Eq. 28 when using a sufficient number of labeled samples (assuming a sufficiently expressive architecture).

Using a semi-supervised method also improves the ability of the VAE to predict data in regions of lower density. In Figure 15, we observe that the aggregated posterior matches the generative parameter distribution much better than the unsupervised case with just over 1% of the samples labeled. Additionally, regions of

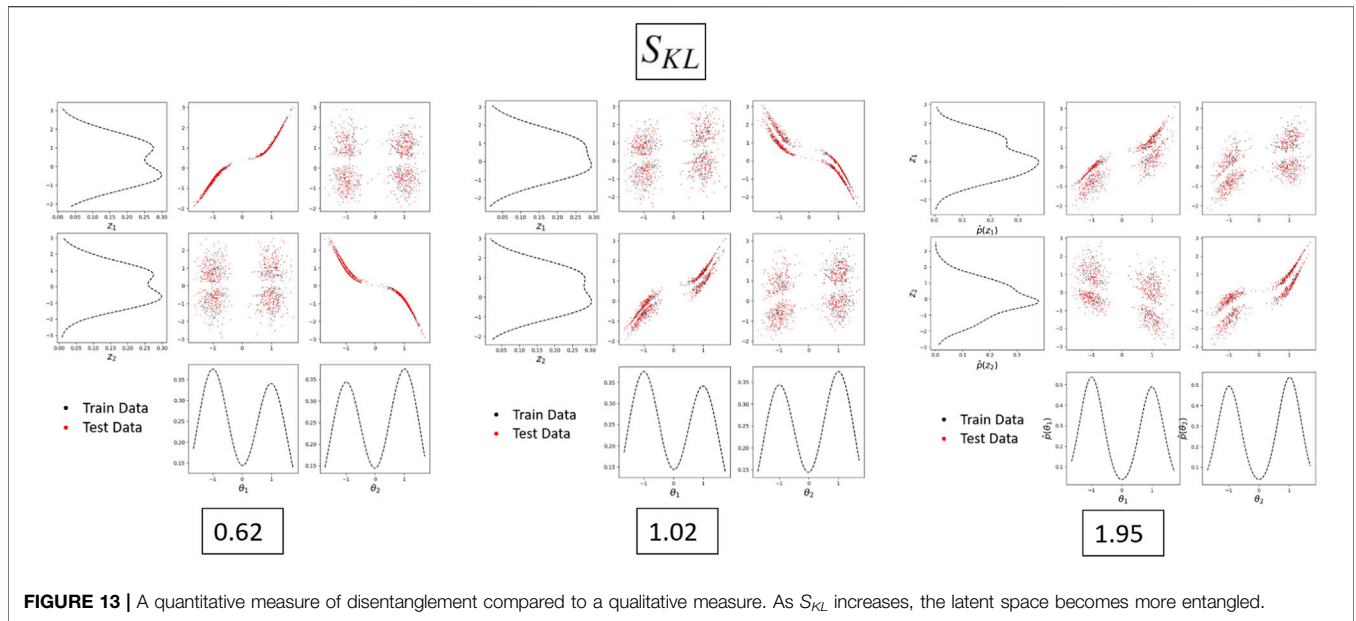


FIGURE 13 | A quantitative measure of disentanglement compared to a qualitative measure. As S_{KL} increases, the latent space becomes more entangled.

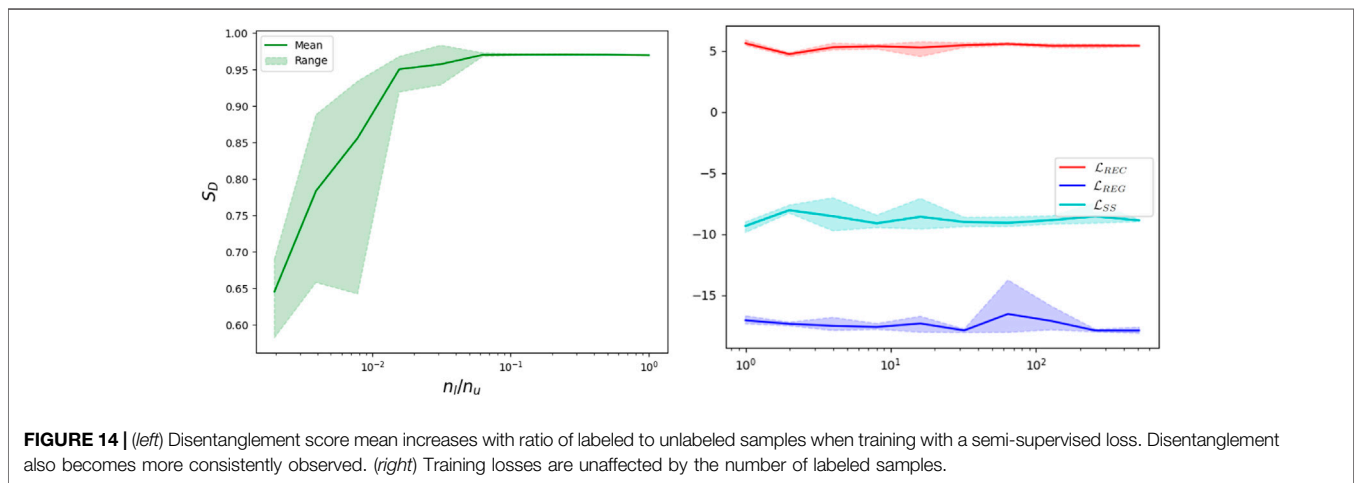


FIGURE 14 | (left) Disentanglement score mean increases with ratio of labeled to unlabeled samples when training with a semi-supervised loss. Disentanglement also becomes more consistently observed. (right) Training losses are unaffected by the number of labeled samples.

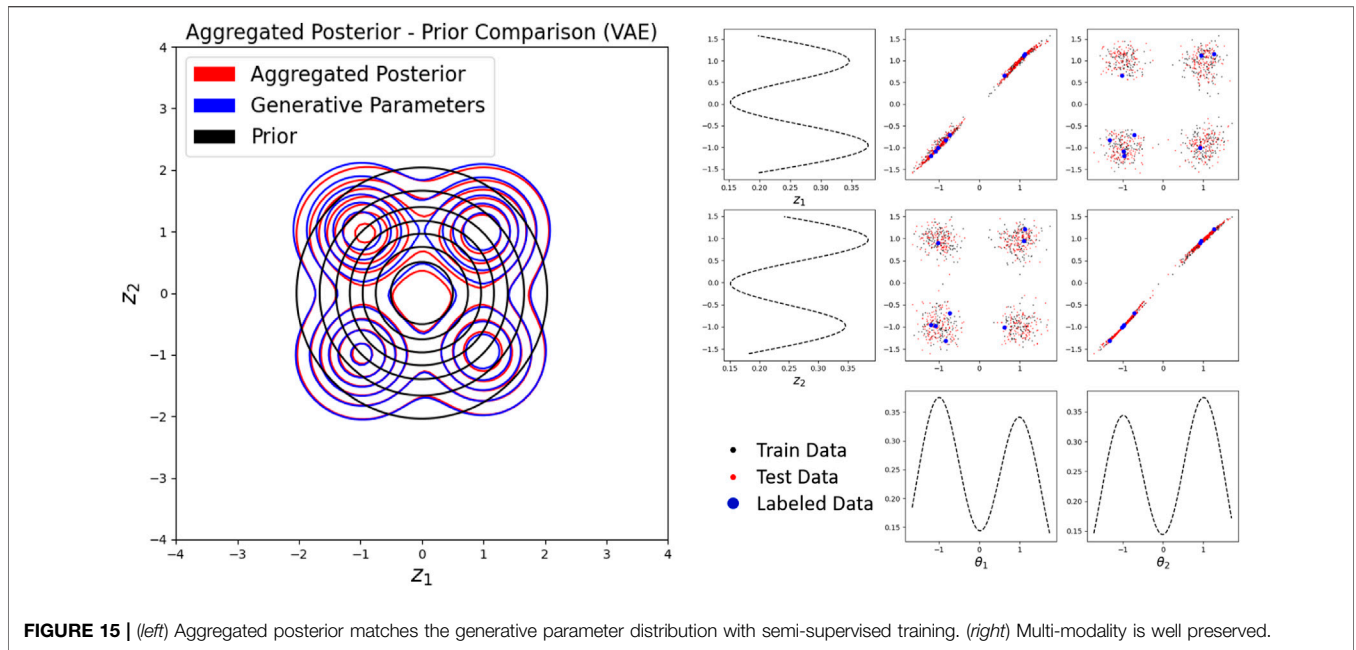
low density in the generative parameter distribution are better represented in the semi-supervised case over the unsupervised case; in other words, multimodality is better preserved (compare to **Figure 12**).

7 CONCLUDING REMARKS AND PERSPECTIVES

Learning representations such that each latent dimension corresponds to a single physical generative factor of variation is useful in many applications, particularly when learned in an unsupervised manner. Learning such *disentangled* representations using VAEs is dependent on many factors including network architecture, assumed form of distributions, prior selection, hyperparameters, and random seeding. The goal

of our work is to develop 1) a consistent unsupervised framework to learn disentangled representations of data obtained through physical experiments or PDE simulations, and 2) to comprehensively characterize the underlying training process, and to recommend strategies to avoid sub-optimal representations.

Accurate reconstruction is desirable from a variety of perspectives, including being necessary for consistent disentanglement. Given two samples near one another in data space, and an accurate decoder, those two samples will be encourage to be near one another in latent space. This is a result of the sampling operation when computing the reconstruction loss. The reconstruction loss is minimized if samples near one another in latent space correspond to samples near one another in data space. Thus, finding an architecture suitable for accurate prediction from latent



representations is of great importance to learn disentangled representations. In our experiments, different architectures were implemented before arriving at and refining the dense architecture (Section 4.1), which was found to accurately reconstruct the data from latent codes. Even with a suitable architecture, however, significant obstacles need to be overcome to arrive at a consistent framework for achieving disentangled representations. Over-regularization can often be difficult to avoid, especially when the variation among data samples is minor. This again emphasizes the necessity to accurately reconstruct the data first before attempting to learn meaningful representations. We have illustrated methods of avoiding over-regularization when training VAEs, but rotationally invariant priors can still create additional difficulties in the ability to disentangle parameters. We illustrated in Section 5 that the standard normal prior typically assumed (which is rotationally-invariant) does not enforce any particular rotation of the latent space, often leading to entangled representations. Rotation of the latent space matters greatly, and without rotational enforcement on the encoder, disentanglement is rarely, or rather randomly, achieved when training with the ELBO loss. We have also shown that the implementation of hierarchical priors allows one to learn non-rotationally-invariant priors such that the regularization loss enforces a rotational constraint on the encoding distribution. However, the regularization loss can contain local minima as the latent space rotates, enforcing a non-factorized and thus incorrectly rotated aggregated posterior. This indicates the need for better prior selection, especially in higher latent dimensions when rotations create more complex effects.

Matching the aggregated posterior to the generative parameter distribution can also be enforced by including labeled samples during training. Including some number of labeled samples in the

dataset and training with a semi-supervised loss, the aggregated posterior consistently matches the shape and orientation of the generative parameter distribution, effectively learning a disentangled representation. The multimodality of the data distribution is also better represented when using labeled data, indicating that the VAE can better predict data in regions of low density over the unsupervised version.

In reference to Section 5, the total correlation (TC) $D_{KL}[q_\phi(z) \parallel \prod_{i=1}^n q_\phi(z_i)]$ appears to be a useful and simpler measurement of disentanglement. When the generative parameters are completely independent (i.e., $p(\theta) = \prod_{i=1}^p p(\theta_i)$) and disentanglement occurs when a factorized $q_\phi(z)$ is learned (aligning the latent space axes with the generative parameter axes). This is the objective of the FactorVAE framework [8], which can successfully encourage a factorized $q_\phi(z)$ through the introduction of TC into the loss function, modifying to ELBO. However, considering a case in which the generative parameters are correlated, a factorized $q_\phi(z)$ is not necessarily desirable. It is in anticipation of a more correlated $p(\theta)$ that we use Eq. 7 as a measure of disentanglement. Additionally, in our work we do not modify the standard VAE objective to produce more accurate reconstruction of the data.

Complete disentanglement has not been observed when generative parameters are correlated in our experiments, but after many trials the same conclusions have been drawn as the uncorrelated case: for disentanglement to occur, the aggregated posterior must contain the same “shape” as the generative parameter distribution - this includes correlations up to permutations of the axes. The **Supplementary Material** further illustrates these ideas. Future work will include disentangling correlated generative parameters, which may be facilitated through learning correlated priors using HP.

In addition to disentangling correlated generative parameters, our broader aim is to extend our work to more complex problems to create a general framework for consistent unsupervised or semi-supervised representation learning. Through our observations here regarding non-rotationally invariant priors along with insights gained from [16], we hypothesize that such a framework will be largely focused on both prior selection and the structural form of the encoding and decoding distributions. Additionally, in a completely unsupervised setting, one must find an encoder and decoder which disentangle the generative parameters, but the dimension of the generative parameters may be unknown. The dimension of the latent space is always user-specified; if the dimension of the latent space is too small or too large, how does this effect the learned representation? Can one successfully and consistently disentangle generative parameters in higher dimensions? These are some of the open questions to be addressed in the future.

The issue of over-regularization often greatly hinders our ability to train VAEs (Section 4.2). Different initialization strategies may be investigated to increase training performance and avoid the issue altogether. It has been shown that principled selection of activation functions, architecture, and initialization can greatly improve not only the efficiency of training, but also facilitate greater performance in terms of reconstruction [31].

The greater scope of this work is to develop an unsupervised and interpretable representation learning framework to generate probabilistic reduced order models for physical problems and use learned representations for efficient design optimization.

REFERENCES

- Li D, Zhang M, Chen W, Feng G. Facial Attribute Editing by Latent Space Adversarial Variational Autoencoders. In: International Conference on Pattern Recognition (2018). p. 1337–42. doi:10.1109/icpr.2018.8545633
- Wang T, Liu J, Jin C, Li J, Ma S. An Intelligent Music Generation Based on Variational Autoencoder. In: International Conference on Culture-oriented Science Technology (2020). p. 394–8. doi:10.1109/iccst50977.2020.00082
- Amini A, Schwarting W, Rosman G, Araki B, Karaman S, Rus D. Variational Autoencoder for End-To-End Control of Autonomous Driving with Novelty Detection and Training De-biasing. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (2018). p. 568–75. doi:10.1109/iros.2018.8594386
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative Adversarial Networks. *Adv Neural Inf Process Syst* (2014) 3.
- Kingma D, Welling M. Auto-Encoding Variational Bayes. In: International Conference on Learning Representations (2013).
- Higgins I, Matthey L, Pal A, Burgess CP, Glorot X, Botvinick M, et al. Eta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In: International Conference on Learning Representations (2017).
- Klushyn A, Chen N, Kurlle R, Cseke B, Smagt PVD. Learning Hierarchical Priors in VAEs. In: Conference on Neural Information Processing Systems (2019).
- Kim H, Mnih A. Disentangling by Factorising. In: International Conference on Machine Learning (2018).
- Zhao S, Song J, Ermon S. InfoVAE: Information Maximizing Variational Autoencoders. *arXiv 1706.02262* (2017).
- Bengio Y, Courville A, Vincent P. Representation Learning: A Review and New Perspectives. *IEEE Trans Pattern Anal Mach Intell* (2013) 35:1798–828. doi:10.1109/tpami.2013.50

DATA AVAILABILITY STATEMENT

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found in the article/Supplementary Material.

AUTHOR CONTRIBUTIONS

CJ developed the methodology for generative modeling, wrote the code, and performed experiments. KD developed the problem statement, and templated ideas on simple problems. CJ was the primary author of the paper while KD provided the direction.

FUNDING

The authors acknowledge support from the Air Force under the grant FA9550-17-1-0195 (Program Managers: Dr. Mitat Birkan, and Dr. Fariba Fahroo).

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fphy.2022.890910/full#supplementary-material>

- Tait DJ, Damoulas T. Variational Autoencoding of PDE Inverse Problems. *arXiv 2006.15641* (2020).
- Lu PY, Kim S, Soljačić M. Extracting Interpretable Physical Parameters from Spatiotemporal Systems Using Unsupervised Learning. *Phys Rev* (2020) 10(3): 031056. doi:10.1103/physrevx.10.031056
- Lopez R, Atzberger P. Variational Autoencoders for Learning Nonlinear Dynamics of Physical Systems. *arXiv 2012.03448* (2020).
- Donà J, Franceschi JY, Sylvain Lamprier, Patrick Gallinari. PDE-Driven Spatiotemporal Disentanglement. In: International Conference on Learning Representations (2021).
- Eastwood C, Williams CKI. A Framework for the Quantitative Evaluation of Disentangled Representations. In: International Conference on Learning Representations (2018).
- Rolinek M, Zietlow D, Martius G. Variational Autoencoders Pursue PCA Directions (By Accident). *arXiv 1812.06775* (2019).
- Burgess CP, Higgins I, Pal A, Matthey L, Watters N, Desjardins G, et al. Understanding Disentangling in β -VAE (2018). *arXiv 1804.03599* (2018).
- Locatello F, Bauer S, Lucic M, Gelly S, Schölkopf B, Bachem O. Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations. *arXiv 1811.12359* (2019).
- Wang Q, Kulkarni SR, Verdu S. Divergence Estimation for Multidimensional Densities via k -Nearest-Neighbor Distances. *IEEE Trans Inform Theor* (2009) 55:2392–405. doi:10.1109/tit.2009.2016060
- Alemi A, Poole B, Fischer I, Dillon J, Saurous RA, Murphy K. Fixing a Broken ELBO. Editor Dy J, Krause A. Proceedings of the 35th International Conference on Machine Learning, July 10–15, 2008. PMLR, 80, 159–168 (2018). Available at: <http://proceedings.mlr.press/v80/alemi18a/alemi18a.pdf>
- Makhzani A, Shlens J, Jaitly N, Goodfellow I. Adversarial Autoencoders. In: International Conference on Learning Representations (2016).
- Odaibo SG. Tutorial: Deriving the Standard Variational Autoencoder (VAE) Loss Function. *arXiv 1907.08956* (2019).

23. Yu R. A Tutorial on VAEs: From Bayes' Rule to Lossless Compression. *arXiv* **2006.10273** (2020).
 24. Duraisamy K. Variational Encoders and Autoencoders : Information-Theoretic Inference and Closed-form Solutions. *arXiv* **2101.11428** (2021).
 25. Berger T. *Rate Distortion Theory. A Mathematical Basis for Data Compression*. Englewood Cliffs: Prentice-Hall (1971).
 26. Cover TM, Thomas JA. *Elements of Information Theory Wiley Series in Telecommunications and Signal Processing*. United States: Wiley-Interscience (2001).
 27. Gibson J. *Information Theory and Rate Distortion Theory for Communications and Compression*. Springer Cham. doi:10.1007/978-3-031-01680-6 (2013).
 28. Lucas J, Tucker G, Grosse RB, Norouzi M. Understanding Posterior Collapse in Generative Latent Variable Models. In: International Conference on Learning Representations (2019).
 29. Zhu Y, Zabarav N. Bayesian Deep Convolutional Encoder-Decoder Networks for Surrogate Modeling and Uncertainty Quantification. *J Comput Phys* (2018) 366:415–47. doi:10.1016/j.jcp.2018.04.018
 30. Fu H, Li C, Liu X, Gao J, Çelikyılmaz A, Carin L. Cyclical Annealing Schedule: A Simple Approach to Mitigating KL Vanishing. In: North American chapter of the Association for Computational Linguistics (2019).
 31. Sitzmann V, Martel JN, Bergman AW, Lindell DB, Wetzstein G. Implicit Neural Representations with Periodic Activation Functions. *arXiv* **2006.09661** (2020).
 32. Nelder JA, Mead R. A Simplex Method for Function Minimization. *Computer J* (1965) 7:308–13. doi:10.1093/comjnl/7.4.308
- Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.
- Publisher's Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.
- Copyright © 2022 Jacobsen and Duraisamy. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.*

APPENDIX A: ROTATIONALLY-INVARIANT DISTRIBUTIONS

A matrix $\mathbf{R} \in \mathbb{R}^{n \times n}$ is a rotation matrix if for all $z \in \mathbb{R}^n$, $\|\mathbf{R}z\|_2 = \|z\|_2$.

A probability distribution $p(z)$ is said to be *rotationally-invariant* if $p(z) = p(\mathbf{R}z)$ for all $z \in \mathbb{R}^n$ and for all rotation matrices $\mathbf{R} \in \mathbb{R}^{n \times n}$.

The ELBO loss is unaffected by rotations of the latent space when training with a rotationally-invariant prior. This is shown in detail in [16].