
Don't Blame the ELBO!

A Linear VAE Perspective on Posterior Collapse

James Lucas^{‡,*}, George Tucker[†], Roger Grosse[‡], Mohammad Norouzi[†]

[‡]University of Toronto

[†]Google Brain

Abstract

Posterior collapse in Variational Autoencoders (VAEs) arises when the variational posterior distribution closely matches the prior for a subset of latent variables. This paper presents a simple and intuitive explanation for posterior collapse through the analysis of linear VAEs and their direct correspondence with Probabilistic PCA (pPCA). We explain how posterior collapse may occur in pPCA due to local maxima in the log marginal likelihood. Unexpectedly, we prove that the ELBO objective for the linear VAE does not introduce additional spurious local maxima relative to log marginal likelihood. We show further that training a linear VAE with exact variational inference recovers an identifiable global maximum corresponding to the principal component directions. Empirically, we find that our linear analysis is predictive even for high-capacity, non-linear VAEs and helps explain the relationship between the observation noise, local maxima, and posterior collapse in deep Gaussian VAEs.

1 Introduction

The generative process of a deep latent variable model entails drawing a number of latent factors from the prior and using a neural network to convert such factors to real data points. Maximum likelihood estimation of the parameters requires marginalizing out the latent factors, which is intractable for deep latent variable models. The influential work of Kingma and Welling [24] and Rezende et al. [35] on Variational Autoencoders (VAEs) enables optimization of a tractable lower bound on the likelihood via a reparameterization of the Evidence Lower Bound (ELBO) [21, 6]. This has led to a surge of recent interest in automatic discovery of the latent factors of variation for a data distribution based on VAEs and principled probabilistic modeling [18, 7, 10, 16].

Unfortunately, the quality and the number of the latent factors learned is influenced by a phenomenon known as *posterior collapse*, where the generative model learns to ignore a subset of the latent variables. Most existing papers suggest that posterior collapse is caused by the KL-divergence term in the ELBO objective, which directly encourages the variational distribution to match the prior [7, 25, 38]. Thus, a wide range of heuristic approaches in the literature have attempted to diminish the effect of the KL term in the ELBO to alleviate posterior collapse [7, 33, 38, 20]. While holding the KL term responsible for posterior collapse makes intuitive sense, the mathematical mechanism of this phenomenon is not well understood. In this paper, we investigate the connection between posterior collapse and spurious local maxima in the ELBO objective through the analysis of linear VAEs. Unexpectedly, we show that spurious local maxima may arise even in the optimization of exact marginal likelihood, and such local maxima are linked with a collapsed posterior.

While linear autoencoders [37] have been studied extensively [4, 26], little attention has been given to their variational counterpart from a theoretical standpoint. A well-known relationship exists between linear autoencoders and PCA – the optimal solution of a linear autoencoder has decoder weight

*Intern at Google Brain

Code available at <https://sites.google.com/view/dont-blame-the-elbo>

columns that span the same subspace as the one defined by the principal components [4]. Similarly, the maximum likelihood solution of probabilistic PCA (pPCA) [39] recovers the subspace of principal components. In this work, we show that a linear variational autoencoder can recover the solution of pPCA. In particular, by specifying a diagonal covariance structure on the variational distribution, one can recover an identifiable autoencoder, which at the global maximum of the ELBO recovers the exact principal components as the columns of the decoder’s weights. Importantly, we show that the ELBO objective for a linear VAE does not introduce any local maxima beyond the log marginal likelihood.

The study of linear VAEs gives us new insights into the cause of posterior collapse and the difficulty of VAE optimization more generally. Following the analysis of Tipping and Bishop [39], we characterize the stationary points of pPCA and show that the *variance of the observation model* directly influences the stability of local stationary points corresponding to posterior collapse – it is only possible to escape these sub-optimal solutions by simultaneously reducing noise and learning better features. Our contributions include:

- We verify that linear VAEs can recover the true posterior of pPCA. Further, we prove that the global optimum of the linear VAE recovers the principal components (not just their spanning sub-space). More importantly, we prove that using ELBO to train linear VAEs does not introduce any additional spurious local maxima relative to log marginal likelihood training.
- While high-capacity decoders are often blamed for posterior collapse, we show that posterior collapse may occur when optimizing log marginal likelihood even without powerful decoders. Our experiments verify the analysis of the linear setting and show that these insights extend even to high-capacity non-linear VAEs. Specifically, we provide evidence that the observation noise in deep Gaussian VAEs plays a crucial role in overcoming local maxima corresponding to posterior collapse.

2 Preliminaries

Probabilistic PCA. The probabilistic PCA (pPCA) model is defined as follows. Suppose latent variables $\mathbf{z} \in \mathbb{R}^k$ generate data $\mathbf{x} \in \mathbb{R}^n$. A standard Gaussian prior is used for \mathbf{z} and a linear generative model with a spherical Gaussian observation model for \mathbf{x} :

$$\begin{aligned} p(\mathbf{z}) &= \mathcal{N}(\mathbf{0}, \mathbf{I}), \\ p(\mathbf{x} | \mathbf{z}) &= \mathcal{N}(\mathbf{W}\mathbf{z} + \boldsymbol{\mu}, \sigma^2\mathbf{I}). \end{aligned} \tag{1}$$

The pPCA model is a special case of factor analysis [5], which uses a spherical covariance $\sigma^2\mathbf{I}$ instead of a full covariance matrix. As pPCA is fully Gaussian, both the marginal distribution for \mathbf{x} and the posterior $p(\mathbf{z} | \mathbf{x})$ are Gaussian, and unlike factor analysis, the maximum likelihood estimates of \mathbf{W} and σ^2 are tractable [39].

Variational Autoencoders. Recently, amortized variational inference has gained popularity as a means to learn complicated latent variable models. In these models, the log marginal likelihood, $\log p(\mathbf{x})$, is intractable but a variational distribution, denoted $q(\mathbf{z} | \mathbf{x})$, is used to approximate the posterior $p(\mathbf{z} | \mathbf{x})$, allowing tractable approximate inference using the Evidence Lower Bound (ELBO):

$$\log p(\mathbf{x}) = \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x}, \mathbf{z}) - \log q(\mathbf{z} | \mathbf{x})] + D_{KL}(q(\mathbf{z} | \mathbf{x}) || p(\mathbf{z} | \mathbf{x})) \tag{2}$$

$$\geq \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x}, \mathbf{z}) - \log q(\mathbf{z} | \mathbf{x})] \tag{3}$$

$$= \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x} | \mathbf{z})] - D_{KL}(q(\mathbf{z} | \mathbf{x}) || p(\mathbf{z})) \quad (:= ELBO) \tag{4}$$

The ELBO [21, 6] consists of two terms, the KL divergence between the variational distribution, $q(\mathbf{z}|\mathbf{x})$, and prior, $p(\mathbf{z})$, and the expected conditional log-likelihood. The KL divergence forces the variational distribution towards the prior and so has reasonably been the focus of many attempts to alleviate posterior collapse. We hypothesize that the log marginal likelihood itself often encourages posterior collapse.

In Variational Autoencoders (VAEs), two neural networks are used to parameterize $q_\phi(\mathbf{z}|\mathbf{x})$ and $p_\theta(\mathbf{x}|\mathbf{z})$, where ϕ and θ denote two sets of neural network weights. The encoder maps an input \mathbf{x} to the parameters of the variational distribution, and then the decoder maps a sample from the variational distribution back to the inputs.




Figure 1: **Stationary points of pPCA.** Two zero-columns of \mathbf{W} are perturbed in the directions of two orthogonal principal components (μ_5 and μ_7) and the optimization landscape around zero-columns is shown, where the goal is to maximize log marginal likelihood. The stability of the stationary points depends critically on σ^2 (the observation noise). Left: σ^2 is too large to capture either principal component. Middle: σ^2 is too large to capture one of the principal components. Right: σ^2 is able to capture both principal components.

Posterior collapse. A dominant issue with VAE optimization is posterior collapse, in which the learned variational distribution is close to the prior. This reduces the capacity of the generative model, making it impossible for the decoder network to make use of the information content of all of the latent dimensions. While posterior collapse is widely acknowledged, formally defining it has remained a challenge. We introduce a formal definition in Section 6.2 which we use to measure posterior collapse in trained deep neural networks.

3 Related Work

Dai et al. [14] discuss the relationship between robust PCA methods [8] and VAEs. They show that at stationary points the VAE objective locally aligns with pPCA under certain assumptions. We study the pPCA objective explicitly and show a direct correspondence with linear VAEs. Dai et al. [14] showed that the covariance structure of the variational distribution may smooth out the loss landscape. This is an interesting result whose interactions with ours is an exciting direction for future research.

He et al. [17] motivate posterior collapse through an investigation of the learning dynamics of deep VAEs. They suggest that posterior collapse is caused by the inference network lagging behind the true posterior during the early stages of training. A related line of research studies issues arising from approximate inference causing a mismatch between the variational distribution and true posterior [12, 22, 19]. By contrast, we show that posterior collapse may exist even when the variational distribution matches the true posterior exactly.

Alemi et al. [2] used an information theoretic framework to study the representational properties of VAEs. They show that with infinite model capacity there are solutions with equal ELBO and log marginal likelihood which span a range of representations, including posterior collapse. We find that even with weak (linear) decoders, posterior collapse may occur. Moreover, we show that in the linear case this posterior collapse is due entirely to the log marginal likelihood.

The most common approach for dealing with posterior collapse is to anneal a weight on the KL term during training from 0 to 1 [7, 38, 30, 18, 20]. Unfortunately, this means that during the annealing process, one is no longer optimizing a bound on the log-likelihood. Also, it is difficult to design these annealing schedules and we have found that once regular ELBO training resumes the posterior will typically collapse again (Section 6.2).

Kingma et al. [25] propose a constraint on the KL term, termed "free-bits", where the gradient of the KL term per dimension is ignored if the KL is below a given threshold. Unfortunately, this method reportedly has some negative effects on training stability [33, 11]. Delta-VAEs [33] instead choose prior and variational distributions such that the variational distribution can never exactly recover the prior, allocating free-bits implicitly. Several other papers have studied alternative formulations of the VAE objective [34, 13, 2, 29, 41]. Dai and Wipf [13] analyzed the VAE objective to improve image fidelity under Gaussian observation models and also discuss the importance of the observation noise. Other approaches have explored changing the VAE network architecture to help alleviate posterior collapse; for example adding skip connections [30, 15]

Rolinek et al. [36] observed that the diagonal covariance used in the variational distribution of VAEs encourages orthogonal representations. They use linearizations of deep networks to prove their results under a modification of the objective function by explicitly ignoring latent dimensions with posterior collapse. Our formulation is distinct in focusing on linear VAEs without modifying the objective function and proving an exact correspondence between the global solution of linear VAEs and the principal components.

Kunin et al. [26] studied the optimization challenges in the linear autoencoder setting. They exposed an equivalence between pPCA and Bayesian autoencoders and point out that when σ^2 is too large information about the latent code is lost. A similar phenomenon is discussed in the supervised learning setting by Chechik et al. [9]. Kunin et al. [26] also showed that suitable regularization allows the linear autoencoder to recover the principal components up to rotations. We show that linear VAEs with a diagonal covariance structure recover the principal components *exactly*.

4 Analysis of linear VAE

This section compares and analyzes the loss landscapes of both pPCA and linear variational autoencoders. We first discuss the stationary points of pPCA and then show that a simple linear VAE can recover the global optimum of pPCA. Moreover, when the data covariance eigenvalues are distinct, the linear VAE identifies the individual principal components, unlike pPCA, which recovers only the PCA subspace. Finally, we prove that ELBO does not introduce any additional spurious maxima to the loss landscape.

4.1 Probabilistic PCA Revisited

The pPCA model (Eq. (1)) is a fully Gaussian linear model, thus we can compute both the marginal distribution for \mathbf{x} and the posterior $p(\mathbf{z} | \mathbf{x})$ in closed form:

$$p(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}, \mathbf{W}\mathbf{W}^\top + \sigma^2\mathbf{I}), \quad (5)$$

$$p(\mathbf{z} | \mathbf{x}) = \mathcal{N}(\mathbf{M}^{-1}\mathbf{W}^\top(\mathbf{x} - \boldsymbol{\mu}), \sigma^2\mathbf{M}^{-1}), \quad (6)$$

where $\mathbf{M} = \mathbf{W}^\top\mathbf{W} + \sigma^2\mathbf{I}$. This model is particularly interesting to analyze in the setting of variational inference, as the ELBO can also be computed in closed form (see Appendix C).

Stationary points of pPCA We now characterize the stationary points of pPCA, largely repeating the thorough analysis of Tipping and Bishop [39] (see Appendix A of their paper). The maximum likelihood estimate of $\boldsymbol{\mu}$ is the mean of the data. We can compute \mathbf{W}_{MLE} and σ_{MLE}^2 as follows:

$$\sigma_{\text{MLE}}^2 = \frac{1}{n-k} \sum_{j=k+1}^n \lambda_j, \quad (7)$$

$$\mathbf{W}_{\text{MLE}} = \mathbf{U}_k(\boldsymbol{\Lambda}_k - \sigma_{\text{MLE}}^2\mathbf{I})^{1/2}\mathbf{R}. \quad (8)$$

Here \mathbf{U}_k corresponds to the first k principal components of the data with the corresponding eigenvalues $\lambda_1, \dots, \lambda_k$ stored in the $k \times k$ diagonal matrix $\boldsymbol{\Lambda}_k$. The matrix \mathbf{R} is an arbitrary rotation matrix which accounts for weak identifiability in the model. We can interpret σ_{MLE}^2 as the average variance lost in the projection. The MLE solution is the global optimum. Other stationary points correspond to zeroing out columns of \mathbf{W}_{MLE} (posterior collapse).

Stability of \mathbf{W}_{MLE} In this section we consider σ^2 to be fixed and not necessarily equal to the MLE solution. Equation 8 remains a stationary point when the general σ^2 is swapped in. One surprising observation is that σ^2 directly controls the stability of the stationary points of the log marginal likelihood (see Appendix A). In Figure 1, we illustrate one such stationary point of pPCA for different values of σ^2 . We computed this stationary point by taking \mathbf{W} to have three principal component columns and zeros elsewhere. Each plot shows the same stationary point perturbed by two orthogonal vectors corresponding to other principal components.

The stability of the pPCA stationary points depends on the size of σ^2 — as σ^2 increases the stationary point tends towards a stable local maximum so that we cannot learn the additional components. Intuitively, the model prefers to explain deviations in the data with the larger observation noise. Fortunately, decreasing σ^2 will increase likelihood at these stationary points so that when learning σ^2 simultaneously these stationary points are saddle points [39]. Therefore, learning σ^2 is necessary for gaining a full latent representation.

4.2 Linear VAEs recover pPCA

We now show that linear VAEs can recover the globally optimal solution to Probabilistic PCA. We will consider the following VAE model,

$$\begin{aligned} p(\mathbf{x} | \mathbf{z}) &= \mathcal{N}(\mathbf{W}\mathbf{z} + \boldsymbol{\mu}, \sigma^2\mathbf{I}), \\ q(\mathbf{z} | \mathbf{x}) &= \mathcal{N}(\mathbf{V}(\mathbf{x} - \boldsymbol{\mu}), \mathbf{D}), \end{aligned} \tag{9}$$

where \mathbf{D} is a diagonal covariance matrix, used globally for all of the data points. While this is a significant restriction compared to typical VAE architectures, which define an amortized variance for each input point, this is sufficient to recover the global optimum of the probabilistic model.

Lemma 1. *The global maximum of the ELBO objective (Eq. (4)) for the linear VAE (Eq. (9)) is identical to the global maximum for the log marginal likelihood of pPCA (Eq. (5)).*

Proof. Note that the global optimum of pPCA is defined up to an orthogonal transformation of the columns of \mathbf{W} , *i.e.*, any rotation \mathbf{R} in Eq. (8) results in a matrix \mathbf{W}_{MLE} that given σ_{MLE}^2 attains maximum marginal likelihood. The linear VAE model defined in Eq. (9) is able to recover the global optimum of pPCA when $\mathbf{R} = \mathbf{I}$. Recall from Eq. (6) that $p(\mathbf{z} | \mathbf{x})$ is defined in terms of $\mathbf{M} = \mathbf{W}^\top \mathbf{W} + \sigma^2 \mathbf{I}$. When $\mathbf{R} = \mathbf{I}$, we obtain $\mathbf{M} = \mathbf{W}_{\text{MLE}}^\top \mathbf{W}_{\text{MLE}} + \sigma_{\text{MLE}}^2 \mathbf{I} = \boldsymbol{\Lambda}_k$, which is diagonal. Thus, setting $\mathbf{V} = \mathbf{M}^{-1} \mathbf{W}_{\text{MLE}}^\top$ and $\mathbf{D} = \sigma_{\text{MLE}}^2 \mathbf{M}^{-1} = \sigma_{\text{MLE}}^2 \boldsymbol{\Lambda}_k^{-1}$, recovers the true posterior with diagonal covariance at the global optimum. In this case, the ELBO equals the log marginal likelihood and is maximized when the decoder has weights $\mathbf{W} = \mathbf{W}_{\text{MLE}}$. Because the ELBO lower bounds log-likelihood, the global maximum of the ELBO for the linear VAE is the same as the global maximum of the marginal likelihood for pPCA. \square

The result of Lemma 1 is somewhat expected because the posterior of pPCA is Gaussian. Further details are given in Appendix C. In addition, we prove a more surprising result that suggests restricting the variational distribution to a Gaussian with a diagonal covariance structure allows one to *identify* the principal components at the global optimum of ELBO.

Corollary 1. *The global maximum of the ELBO objective (Eq. (4)) for the linear VAE (Eq. (9)) has the scaled principal components as the columns of the decoder network.*

Proof. Follows directly from the proof of Lemma 1 and Eq. (8). \square

We discuss this result in Appendix B. This full identifiability is non-trivial and is not achieved even with the regularized linear autoencoder [26].

So far, we have shown that at its global optimum the linear VAE recovers the pPCA solution, which enforces orthogonality of the decoder weight columns. However, the VAE is trained with the ELBO rather than the log marginal likelihood — often using SGD. The majority of existing work suggests that the KL term in the ELBO objective is responsible for posterior collapse. So, we should ask whether this term introduces additional spurious local maxima. Surprisingly, for the linear VAE model the ELBO objective *does not* introduce any additional spurious local maxima. We provide a sketch of the proof below with full details in Appendix C.

Theorem 1. *The ELBO objective for a linear VAE does not introduce any additional local maxima to the pPCA model.*

Proof. (Sketch) If the decoder has orthogonal columns, then the variational distribution recovers the true posterior at stationary points. Thus, the variational objective will exactly recover the log marginal likelihood. If the decoder does not have orthogonal columns then the variational distribution is no longer tight. However, the ELBO can always be increased by applying an infinitesimal rotation to the right-singular vectors of the decoder towards identity: $\mathbf{W}' \leftarrow \mathbf{W}\mathbf{R}_\epsilon$ (so that the decoder columns are closer to orthogonal). This works because the variational distribution can fit the posterior more closely while the log marginal likelihood is invariant to rotations of the weight columns. Thus, any additional stationary points in the ELBO objective must necessarily be saddle points. \square

The theoretical results presented in this section provide new intuition for posterior collapse in VAEs. In particular, the KL between the variational distribution and the prior is not entirely responsible for posterior collapse — log marginal likelihood has a role. The evidence for this is two-fold. We have shown that log marginal likelihood may have spurious local maxima but also that in the linear case the ELBO objective does not add any additional spurious local maxima. Rephrased, in the linear setting the problem lies entirely with the probabilistic model. We should then ask, to what extent do these results hold in the non-linear setting?

5 Deep Gaussian VAEs

The deep Gaussian VAE consists of a decoder D_θ and an encoder E_ϕ . The ELBO objective can be expressed as,

$$\mathcal{L}(\mathbf{x}; \theta, \phi) = -\text{KL}(q_\phi(\mathbf{z} | \mathbf{x}) \| p(\mathbf{z})) - \frac{1}{2\sigma^2} \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\|D_\theta(\mathbf{z}) - \mathbf{x}\|^2] - \frac{1}{2} \log(2\pi\sigma^2) \quad (10)$$

The role of σ^2 in this objective invites a natural comparison to the β -VAE objective [18], where the KL term is weighted by $\beta \in \mathbb{R}^+$. Alemi et al. [2] propose using small β values to force powerful decoders to utilize the latent variables, but this comes at the cost of poor ELBO. Practitioners must then use downstream task performance for model selection, thus sacrificing one of the primary benefits of likelihood-based models. However, for a given β , one can find a corresponding σ^2 (and a learning rate) such that the gradient updates to the network parameters are identical. Importantly, the Gaussian partition function for a Gaussian observation model (the last term on the RHS of Eq. (10)) prevents ELBO from deviating from the β -VAE’s objective with a β -weighted KL term while maintaining the benefits to representation learning when σ^2 is small. For the Gaussian VAE, this helps connect the dots between the role of local maxima and observation noise in posterior collapse vs. heuristic approaches that attempted to alleviate posterior collapse by diminishing the effect of the KL term [7, 33, 38, 20]. In the following section, we will study the nonlinear VAE empirically and explore connections to the linear theory.

6 Experiments

In this section, we present empirical evidence found from studying two distinct claims. First, we verify our theoretical analysis of the linear VAE model. Second, we explore to what extent these insights apply to deep nonlinear VAEs.

6.1 Linear VAEs

We ran two sets of experiments on 1000 randomly chosen MNIST images. First, we trained linear VAEs with learnable σ^2 for a range of hidden dimensions². For each model, we compared the final ELBO to the maximum-likelihood of pPCA finding them to be essentially indistinguishable (as predicted by Lemma 1 and Theorem 1). For the second set of experiments, we took the pPCA MLE solution for \mathbf{W} for each number of hidden dimensions and computed the likelihood under the observation noise which maximizes likelihood for 50 hidden dimensions. We observed that adding additional principal components (after 50) will initially improve likelihood but eventually adding more components (after 200) actually decreases the likelihood. In other words, the collapsed solution is actually preferred if the observation noise is not set correctly — we observe this theoretically through the stability of the stationary points (e.g. Figure 1).




Figure 2: The log marginal likelihood and optimal ELBO of MNIST pPCA solutions over increasing hidden dimension. Green represents the MLE solution (global maximum), the red dashed line is the optimal ELBO solution which matches the global optimum. The blue line shows the log marginal likelihood of the solutions using the full decoder weights when σ^2 is fixed to its MLE solution for 50 hidden dimensions.

²The VAEs were trained using the analytic ELBO (Appendix C.1) and without mini-batching gradients.




Figure 3: Stochastic vs analytic ELBO training: using the analytic gradient of the ELBO led to faster convergence and better final ELBO (950.7 vs. 939.3).




Figure 4: VAEs with linear decoders trained on real-valued MNIST with nonlinear preprocessing [31]. Final average ELBO on training set are (ordered by legend): -1098.2, -1108.7, -1112.1, -1119.6.

Effect of stochastic ELBO estimates In general, we are unable to compute the ELBO in closed form and so instead rely on unbiased Monte Carlo estimates using the reparameterization trick. These estimates add high-variance noise and can make optimization more challenging [24]. In the linear model, we can compare the solutions obtained using the stochastic ELBO gradients versus the analytic ELBO³ (Figure 3). Additional experimental details are in Appendix E. We found that stochastic optimization had slower convergence (when compared to analytic training with the same learning rate) and, unsurprisingly, reached a worse final training ELBO value (in other words, worse steady-state risk due to the gradient variance).

Nonlinear Encoders With a linear decoder and nonlinear encoder, Lemma 1 still holds, and the optimal variational distribution is the same as the true posterior has not changed. However, Corollary 1 and Theorem 1 no longer hold in general. Even a deep linear encoder will not have a unique global maximum and new stationary points (possibly maxima) may be introduced to ELBO in general. To investigate how deeper networks may impact optimization of the probabilistic model, we trained linear decoders with varying encoders using ELBO. We do not expect the linear encoder to be outperformed and indeed the empirical results support this (Figure 4).

6.2 Investigating posterior collapse in deep nonlinear VAEs

We explored how the analysis of the linear VAEs extends to deep nonlinear models. To do so, we trained VAEs with Gaussian observation models on the MNIST [27] and CelebA [28] datasets. We apply uniform dequantization as in Papamakarios et al. [31] in each case. We also adopt the nonlinear logit preprocessing transformation from Papamakarios et al. [31] to provide fair comparisons with existing work. We also report results of models trained directly in pixel space in the appendix (there is no significant difference for the hypotheses we test).

Measuring posterior collapse In order to measure the extent of posterior collapse, we introduce the following definition. We say that latent dimension i has (ϵ, δ) -collapsed if $\mathbb{P}_{\mathbf{x} \sim p}[KL(q(z_i|\mathbf{x})||p(z_i)) < \epsilon] \geq 1 - \delta$. Note that the linear VAE can suffer $(0, 0)$ -collapse. To estimate this practically, we compute the proportion of data samples which induce a variational distribution with KL divergence less than ϵ and finally report the percentage of dimensions which have (ϵ, δ) -collapsed. Throughout this work, we fix $\delta = 0.01$ and vary ϵ .

Investigating σ^2 We trained MNIST VAEs with 2 hidden layers in both the decoder and encoder, ReLU activations, and 200 latent dimensions. We first evaluated training with fixed values of the observation noise, σ^2 . This mirrors many public VAE implementations where σ^2 is fixed to 1 throughout training (also observed by Dai and Wipf [13]), however, our linear analysis suggests that this is suboptimal. Then, we consider the setting where the observation noise and VAE weights are learned simultaneously.

In Table 1 we report the final ELBO of nonlinear VAEs trained on real-valued MNIST. For fixed σ^2 , we found that the final models could have significant differences in ELBO which were maintained even after tuning σ^2 to the learned representations — the converged representations are less good when σ^2 is too large as predicted by the linear model. Additionally, we report the final ELBO

³We use 1000 MNIST images, as before, to enable full-batch training so that the only source of noise is from the reparameterization trick [24]

	Model		ELBO	σ^2 -tuned ELBO	Tuned σ^2	Posterior collapse (%)	KL Divergence
	Init σ^2	Final σ^2					
MNIST	10.0		-1450.3 ± 4.2	-1098.2 ± 28.3	1.797	89.88	28.8 ± 1.4
	1.0		-1022.1 ± 5.4	-1018.3 ± 5.3	1.145	27.38	125.4 ± 4.2
	0.1		-3697.3 ± 493.3	-1190.8 ± 37.4	0.968	3.25	368.7 ± 94.6
	0.01		-38612.5 ± 1189.8	-2090.8 ± 975.1	0.877	0.00	695.9 ± 118.1
	0.001		-504259.1 ± 49149.8	-1744.7 ± 48.4	0.810	0.00	756.2 ± 12.6
	10.0	1.320	-1022.2 ± 4.5	-1022.3 ± 4.6	1.318	73.75	73.8 ± 9.8
	1.0	1.183	-1011.1 ± 2.7	-1011.1 ± 2.8	1.182	47.88	106.3 ± 2.5
	0.1	1.194	-1025.4 ± 8.6	-1025.4 ± 8.6	1.195	29.25	116.1 ± 11.4
	0.01	1.194	-1030.6 ± 3.5	-1030.5 ± 3.5	1.191	23.00	121.9 ± 7.7
	0.001	1.208	-1038.7 ± 5.6	-1038.8 ± 5.6	1.209	27.00	124.9 ± 1.6
CELEBA 64	10.0		-73328.4 ± 0.49	-55186.7 ± 35.1	0.2040	80.56	56.12 ± 0.4
	1.0		-59841.8 ± 30.1	-51294.8 ± 333.7	0.1020	2.52	213.4 ± 6.3
	0.1		-50760.3 ± 353.4	-50698.5 ± 393.9	0.0883	32.72	483.8 ± 36.2
	0.01		-82478.7 ± 1823.3	-51373.9 ± 213.3	0.0817	0.00	1624.2 ± 8.8
	0.001		-531924.5 ± 17177.6	-57381.5 ± 512.6	0.0296	0.00	2680.2 ± 41.5
	10.0	0.0962	-51109.5 ± 408.2	-51109.5 ± 408.3	0.0963	53.32	364.5 ± 26.4
	1.0	0.0875	-50631.2 ± 163.4	-50631.0 ± 163.3	0.0875	54.76	462.2 ± 20.0
	0.1	0.0863	-50646.9 ± 269.0	-50645.9 ± 267.5	0.0869	28.84	520.9 ± 11.7
	0.01	0.0911	-51285.0 ± 708.1	-51284.8 ± 708.1	0.0963	5.64	557.0 ± 50.5
	0.001	0.1040	-51695.1 ± 322.4	-51694.8 ± 322.7	0.0974	0.00	537.5 ± 46.2

Table 1: Evaluation of deep Gaussian VAEs (averaged over 5 trials) on real-valued MNIST. We report the ELBO on the training set in all cases. Collapse percent gives the percentage of latent dimensions which are within 0.01 KL of the prior for at least 99% of the encoder inputs.

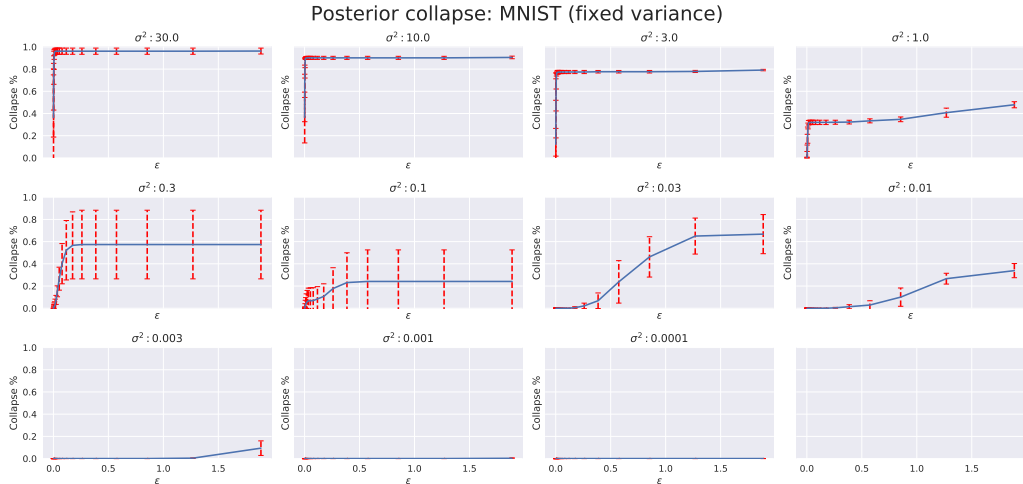


Figure 5: Posterior collapse percentage as a function of ϵ -threshold for a deep VAE trained on MNIST. We measure posterior collapse for trained networks as the proportion of latent dimensions that are within ϵ KL divergence of the prior for at least a $1 - \delta$ proportion of the training data points ($\delta = 0.01$ in the plots).

values when the model is trained while learning σ^2 with different initial values of σ^2 . The gap in performance across different initializations is smaller than for fixed σ^2 but is still significant. The linear VAE does not predict this gap which suggests that learning σ^2 correctly is more challenging in the nonlinear case.

Despite the large volume of work studying posterior collapse it has not been measured in a consistent way (or even defined so). In Figure 5 and Figure 6 we measure posterior collapse for trained networks as described above (we chose $\delta = 0.01$). By considering a range of ϵ values we found this was (moderately) robust to stochasticity in data preprocessing. We observed that for large choices of σ^2 initialization the variational distribution matches the prior closely. This was true even when σ^2 is learned — suggesting that local optima may contribute to posterior collapse in deep VAEs.

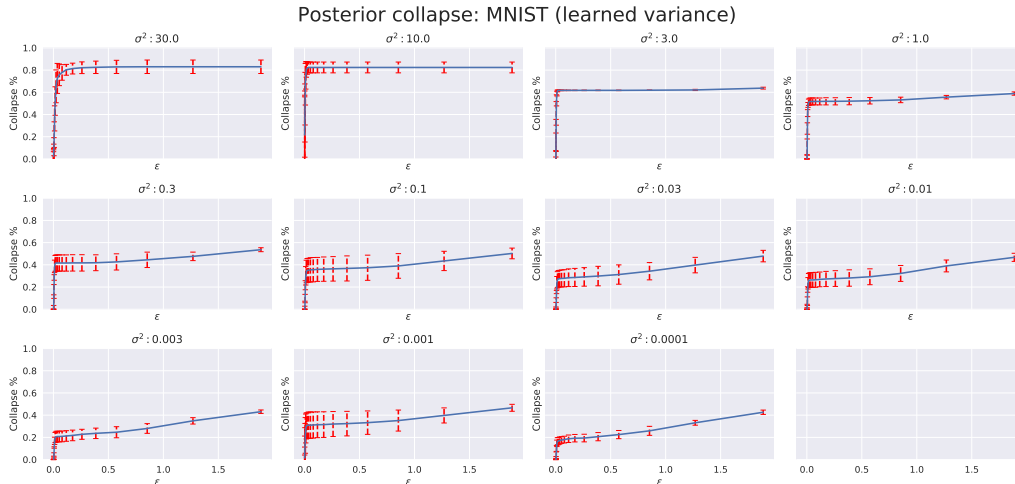


Figure 6: Posterior collapse percentage as a function of ϵ -threshold for a deep VAE trained on MNIST. We measure posterior collapse for trained networks as the proportion of latent dimensions that are within ϵ KL divergence of the prior for at least a $1 - \delta$ proportion of the training data points ($\delta = 0.01$ in the plots).

CelebA VAEs We trained deep convolutional VAEs with 500 hidden dimensions on images from the CelebA dataset (resized to 64x64). We trained the CelebA VAEs with different fixed values of σ^2 and compared the ELBO before and after tuning σ^2 to the learned representations (Table 1). Further, we explored training the CelebA VAE while learning σ^2 over varied initializations of the observation noise. The VAE is sensitive to the initialization of the observation noise even when σ^2 is learned (in particular, in terms of the number of collapsed dimensions).

7 Discussion

By analyzing the correspondence between linear VAEs and pPCA, this paper makes significant progress towards understanding the causes of posterior collapse. We show that for simple linear VAEs posterior collapse is caused by ill-conditioning of the stationary points in the log marginal likelihood objective. We demonstrate empirically that the same optimization issues play a role in deep non-linear VAEs. Finally, we find that linear VAEs are useful theoretical test-cases for evaluating existing hypotheses on VAEs and we encourage researchers to consider studying their hypotheses in the linear VAE setting.

8 Acknowledgements

This work was guided by many conversations with and feedback from our colleagues. In particular, we thank Durk Kingma, Alex Alemi, and Guodong Zhang for invaluable feedback on early versions of this work. RG acknowledges support from the CIFAR Canadian AI Chairs program.

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- [2] A. A. Alemi, B. Poole, I. Fischer, J. V. Dillon, R. A. Saurous, and K. Murphy. Fixing a broken ELBO. *arXiv preprint arXiv:1711.00464*, 2017.
- [3] J. Atchison and S. M. Shen. Logistic-normal distributions: Some properties and uses. *Biometrika*, 67(2):261–272, 1980.
- [4] P. Baldi and K. Hornik. Neural networks and principal component analysis: Learning from examples without local minima. *Neural networks*, 2(1):53–58, 1989.
- [5] D. J. Bartholomew. *Latent variable models and factors analysis*. Oxford University Press, Inc., 1987.
- [6] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 2017.
- [7] S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*, 2015.
- [8] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011.
- [9] G. Chechik, A. Globerson, N. Tishby, and Y. Weiss. Information bottleneck for gaussian variables. *Journal of machine learning research*, 6(Jan):165–188, 2005.
- [10] R. T. Q. Chen, X. Li, R. Grosse, and D. Duvenaud. Isolating sources of disentanglement in variational autoencoders. *Advances in Neural Information Processing Systems*, 2018.
- [11] X. Chen, D. P. Kingma, T. Salimans, Y. Duan, P. Dhariwal, J. Schulman, I. Sutskever, and P. Abbeel. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016.
- [12] C. Cremer, X. Li, and D. Duvenaud. Inference suboptimality in variational autoencoders. *arXiv preprint arXiv:1801.03558*, 2018.
- [13] B. Dai and D. Wipf. Diagnosing and enhancing VAE models. In *International Conference on Learning Representations*, 2019.
- [14] B. Dai, Y. Wang, J. Aston, G. Hua, and D. Wipf. Hidden talents of the variational autoencoder. *arXiv preprint arXiv:1706.05148*, 2017.
- [15] A. B. Dieng, Y. Kim, A. M. Rush, and D. M. Blei. Avoiding latent variable collapse with generative skip models. *arXiv preprint arXiv:1807.04863*, 2018.
- [16] R. Gomez-Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernandez-Lobato, B. Sanchez-Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *American Chemical Society Central Science*, 2018.
- [17] J. He, D. Spokoynny, G. Neubig, and T. Berg-Kirkpatrick. Lagging inference networks and posterior collapse in variational autoencoders. In *International Conference on Learning Representations*, 2019.
- [18] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2016.

- [19] D. Hjelm, R. R. Salakhutdinov, K. Cho, N. Jovic, V. Calhoun, and J. Chung. Iterative refinement of the approximate posterior for directed belief networks. In *Advances in Neural Information Processing Systems*, 2016.
- [20] C.-W. Huang, S. Tan, A. Lacoste, and A. C. Courville. Improving explorability in variational inference with annealed variational objectives. In *Advances in Neural Information Processing Systems*, 2018.
- [21] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An introduction to variational methods for graphical models. *Machine learning*, 1999.
- [22] Y. Kim, S. Wiseman, A. C. Miller, D. Sontag, and A. M. Rush. Semi-amortized variational autoencoders. *arXiv preprint arXiv:1802.02550*, 2018.
- [23] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [24] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [25] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751, 2016.
- [26] D. Kunin, J. M. Bloom, A. Goeva, and C. Seed. Loss landscapes of regularized linear autoencoders. *arXiv preprint arXiv:1901.08168*, 2019.
- [27] Y. LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [28] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- [29] X. Ma, C. Zhou, and E. Hovy. MAE: Mutual posterior-divergence regularization for variational autoencoders. In *International Conference on Learning Representations*, 2019.
- [30] L. Maaløe, M. Fraccaro, V. Liévin, and O. Winther. BIVA: A very deep hierarchy of latent variables for generative modeling. *arXiv preprint arXiv:1902.02102*, 2019.
- [31] G. Papamakarios, T. Pavlakou, and I. Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*. 2017.
- [32] K. B. Petersen et al. The matrix cookbook.
- [33] A. Razavi, A. van den Oord, B. Poole, and O. Vinyals. Preventing posterior collapse with delta-VAEs. In *International Conference on Learning Representations*, 2019.
- [34] D. J. Rezende and F. Viola. Taming VAEs. *arXiv preprint arXiv:1810.00597*, 2018.
- [35] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- [36] M. Rolinek, D. Zietlow, and G. Martius. Variational autoencoders pursue PCA directions (by accident). *arXiv preprint arXiv:1812.06775*, 2018.
- [37] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [38] C. K. Sønderby, T. Raiko, L. Maaløe, S. K. Sønderby, and O. Winther. Ladder variational autoencoders. In *Advances in neural information processing systems*, pages 3738–3746, 2016.
- [39] M. E. Tipping and C. M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- [40] J. M. Tomczak and M. Welling. Vae with a VampPrior. *arXiv preprint arXiv:1705.07120*, 2017.
- [41] S. Yeung, A. Kannan, Y. Dauphin, and L. Fei-Fei. Tackling over-pruning in variational autoencoders. *arXiv preprint arXiv:1706.03643*, 2017.