
Disentangling to Cluster: Gaussian Mixture Variational Ladder Autoencoders

Matthew Willetts, Stephen Roberts, Chris Holmes
University of Oxford
Alan Turing Institute
{mwilletts, sproberts, cholmes}@turing.ac.uk

Abstract

In clustering we normally output one cluster variable for each datapoint. However it is not necessarily the case that there is only one way to partition a given dataset into cluster components. For example, one could cluster objects by their colour, or by their type. Different attributes form a hierarchy, and we could wish to cluster in any of them. By disentangling the learnt latent representations of some dataset into different layers for different attributes we can then cluster in those latent spaces. We call this *disentangled clustering*. Extending Variational Ladder Autoencoders (Zhao et al., 2017), we propose a clustering algorithm, VLAC, that outperforms a Gaussian Mixture DGM in cluster accuracy over digit identity on the test set of SVHN. We also demonstrate learning clusters jointly over numerous layers of the hierarchy of latent variables for the data, and show component-wise generation from this hierarchical model.

1 Introduction

What do we mean when we talk of clustering a set of images? We have data $x \in \mathcal{D}$ available at train time. The model assigns some cluster variable $y \in 1, \dots, K$ to each datapoint x . We score our method using some existing ground-truth label information $t \in 1, \dots, T$ that was not used when we applied our clustering algorithm.

But plausibly there are different ways to cluster the same set of images, so it is limiting to insist *a priori* that there is only one cluster variable y per datapoint and that the clustering algorithm must successfully match that y to t over some dataset. Consider clustering images of digits. While it might make sense to cluster them against the ground truth classes corresponding to what digit they represent, one could conceivably cluster based on other aspects: What colour is the digit? What colour is the background? What is the style of the typeface?

As we analyse more complex data, intuitively we can expect to find an increasing number of different aspects on which one could plausibly cluster. To have one latent cluster variable capturing all the different aspects, the number of cluster components needed would be the product of the number of clusters for each aspect. For example, for digits, one might need a cluster component for each combination of digit identity and colour.

Because having such a large number of cluster components would be unwieldy and unparsimonious, we are interested in outputting a set of L cluster variables $\{y_\ell\}$, one of which might correspond to a particular given ground truth label and the others may capture other ways of clustering the data. This broadened conception of clustering, which we call *disentangled clustering*, requires us to learn sets of latent variables at different levels of the hierarchy of attributes so as to perform clustering over them.

So we wish for disentangled representations. Further we want these representations to be ordered in some way. The Variational Ladder Autoencoder (VLAE) [1] provides this. It separates out subsets of

latent variables of images via the degree of computation needed to map between each layer of latent variable and the image: ‘high level’ and ‘low level’ aspects of the image have their associated latent variables separated from each other by how expressive the mapping is between that latent variable and the data.

We augment VLAEs so they can perform clustering at each layer, by introducing mixture distributions for each subset of latent variables. We call this model *Variational Ladder Autoencoder for Clustering* or *VLAC*. Like other deep generative models that are trained using amortised variational inference, we produce recognition networks as inference artifacts that can be applied after training to new data. This enable us to perform *test set clustering*.

We find that we can learn a clustering variable in the hierarchy that corresponds to the ground truth for SVHN [2], and that we get better test set clustering accuracy than when using a single entangled layer of latent variables in a Gaussian Mixture DGM [3, 4, 5].

2 Related Work

Our approach has some links to multi-view clustering [6, 7, 8, 9, 10, 11], where like in multi-view learning one has a feature vector that is composed of distinct chunks of features each about some different aspect of that datapoint. These subsets of features can then be each be used to produce clustering assignments. Often the aim is to use these different sources of information to try to create the same clustering assignments [6, 7].

However, unlike multi-view clustering, we do not have access to the already-chunked feature vector that divides up the different aspects one could cluster over. And further, whereas in multi-view clustering the different views are used to bolster one overall clustering assignment for each datapoint, here want to cluster distinctly in each learnt set of latent variables.

Various deep learning-based algorithms have been proposed for clustering, including: Gaussian Mixture DGMs [3, 4, 5], GM-VAE [12], VaDE [13], IMSAT [14], DEC [15] and ACOL-GAR [16].

Many recent papers on learning disentangled representations are based around achieving statistical independence between the different dimensions of the latent variables in the aggregate posterior $q(z) = \sum_{i=1}^{|\mathcal{D}|} q_\phi(z|x_i)$. This then leads to simple generation of synthetic data by ancestral sampling, where each dimension in the learnt latent variable then controls a single (often human interpretable) aspect of the data. Examples of this include Factor VAE [17], β -TCVAE [18] and HFVAE [19]. These approaches do not learn hierarchies of disentangled factors, having only one stochastic layer, and in their approach are orthogonal to the method of disentangling by degree of computation that is used by VLAEs.

Recent theoretical work has studied definitions of disentangling around symmetry groups [20] and the effect of different priors in DGMs on their posterior representations [21]. In [21], different posited varieties of disentangling (such as: the above idea of axis alignment of interpretable generative factors, learning of sparse representations, or learning to cluster) arise from different priors in DGMs when the ELBO has been augmented to contain a divergence between the aggregate posterior $q(z)$ and the prior $p(z)$, that one aims to minimise.

While our model does not include this additional divergence in our objective, our approach is an example of trying to obtain a particular variety of disentangling – here a hierarchy of layers of variables that internally demonstrate clustering – through the interplay of a prior and a suitable variational posterior.

Other works closely similar to ours include DGMs that attempt to learn the structure of network of latent variables, given the data. For example, [22] learns a Bayesian non-parametric model, a nested Chinese Restaurant Process [23], as the generative model for a VAE.

The work philosophically most similar to ours is [24], where the authors aim to learn a set of clustering variables given the data, that like ours describe different aspects of the data.

Unlike these structure-learning approaches, we choose to constrain the structure of the hierarchy of latent variables in the generative model to be much simpler – our discrete latent variables are all independent in the generative model.

Further, work in this area has often focused on learning not on raw image data but on features extracted by some pre-trained method. A popular choice is the activations in the penultimate layer of a deep net trained on Imagenet [25]. [22, 24] are examples of this. Unlike that work, our method trains on raw image data directly. This is beneficial as it means the model has access to more information. We are not constrained only to have access to the aspects of the data that have been picked out so as to be useful for the task under which the feature extractor was trained. The features needed to classify an image can even be optimised by throwing away information we might care about.

For instance, in training a classifier on SVHN, one could imagine that dropping (most) colour information as quickly as possible might be worthwhile – if in the dataset there is no correlation between colour and digit identity. The classification training objective would be telling us to view colour information as noise on the signal we care about, and thus should play a minimal role in the embedding we get in the final layer. Thus, we are interested in learning directly on raw image data.

3 VLAC: Variational Ladder Autoencoders for Clustering

VLAEs To gain a more expressive model over a vanilla VAE that has a single set of latent variables z [26, 27], it is natural to consider having a hierarchy of latent variables $z \rightarrow \mathbf{z} = \{z_\ell\}$ for $\ell \in 1, \dots, L$ each with dimensionality d_ℓ^z . The simplest VAE with a hierarchy of conditional stochastic variables in the generative model is the Deep Latent Gaussian Model [27]. Here we have a Markov chain in the generative model: $p_\theta(x, \mathbf{z}) = p_\theta(x|z_1) \prod_{\ell=1}^{L-1} [p_\theta(z_\ell|z_{\ell+1})] p(z_L)$. Performing inference in this model is challenging. The latent variables further from the data can fail to learn anything informative [28, 1]: in the worst case a single-layer VAE can train in isolation within this hierarchical model: each $p_\theta(z_\ell|z_{\ell+1})$ distribution can become a fixed distribution not depending on z_{i+1} such that each KL divergence present in the objective between corresponding z_ℓ layers is driven to a local minima. [1] gives a proof of this separation for the case where the model is perfectly trained ($\text{KL}(q_\phi(z, x) || p_\theta(x, z)) = 0$).

The Variational Ladder Autoencoder (VLAE) [1] avoid this collapse in hierarchical VAEs. Here we have a ‘flat hierarchy’ in \mathbf{z} . Instead of having the set of z_ℓ variables conditioned on each other, the prior for $\mathbf{z} = \{z_\ell\}$ is a set of independent standard Gaussians: $p(\mathbf{z}) = \prod_{\ell=1}^L \mathcal{N}(z_\ell|0, \mathcal{I})$, $|z_\ell| = d_\ell^z$, and inside the conditional distribution $p_\theta(x|\mathbf{z})$ there is a ladder [29, 30, 28] over z_ℓ variables. This separates out aspects of the data by the degree of computation needed to map between their latent representation and x . Thus $p_\theta(x|\mathbf{z})$ is defined implicitly by:

$$\tilde{z}_L = f_L^\theta(z_L) \quad (1)$$

$$\tilde{z}_\ell = f_\ell^\theta(z_\ell, \tilde{z}_{\ell+1}) \quad (2)$$

$$x \sim p(x|f_0^\theta(\tilde{z}_1)) \quad (3)$$

for $\ell \in 1, \dots, L-1$. The posterior follows a similar structure, but in reverse:

$$h_\ell = g_\ell^\phi(h_{\ell-1}) \quad (4)$$

$$z_\ell \sim \mathcal{N}(z_\ell | \mu_\ell^\phi(h_\ell), \sigma_\ell^\phi(h_\ell)) \quad (5)$$

for $\ell \in 1, \dots, L$ and where $h_0 = x$.

VLAC: Variational Ladder Autoencoder for Clustering To enable us to cluster, we alter the generative model above so we have a mixture distribution in z_ℓ : $p_\theta(x, \mathbf{z}) \rightarrow p_\theta(x, \mathbf{z}, \mathbf{y}) = p(x|\mathbf{z}) \prod_{\ell=1}^L p_\theta(z_\ell|y_\ell) p(y_\ell)$. $p(z_\ell|y_\ell) = \mathcal{N}(z_\ell | \mu_{y_\ell}, \sigma_{y_\ell})$ and $p(y_\ell) = \text{Cat}(1/K_\ell)$. Where $\mathbf{K} = \{K_\ell\}$ is the vector of the dimensionalities of our discrete variables $\{y_\ell\}$. Our variational posterior is now $q_\phi(\mathbf{z}, \mathbf{y}|x)$. We choose to factorise this as $q_\phi(\mathbf{z}|\mathbf{y}, x) q_\phi(\mathbf{y}|x)$. Each of these is a product over our L layers. $q_\phi(\mathbf{y}|x) = \prod_{\ell=1}^L q_\phi(y_\ell|x)$, $q_\phi(y_\ell|x) = \text{Cat}(\pi_\ell^\phi(x))$. $q_\phi(\mathbf{z}|\mathbf{y}, x) = \prod_{\ell=1}^L q_\phi(z_\ell|x, \{y_i\}_{i \leq \ell})$, and so the new counterparts to Eqs (4-5) are:

$$h_1 = g_1^\phi(x) \quad (6)$$

$$h_\ell = g_\ell^\phi(h_{\ell-1}, y_{\ell-1}) \text{ for } \ell > 1 \quad (7)$$

$$z_\ell \sim \mathcal{N}(z_\ell | \mu_\ell^\phi(h_\ell, y_\ell), \sigma_\ell^\phi(h_\ell, y_\ell)) \quad (8)$$

See Figure 1 for a graphical representation of this model for $L = 2$. After training the $q_\phi(y_\ell|x)$ networks are inference artifacts that can be applied to new datapoints.

Thus the ELBO for our model is:

$$\mathcal{L}(x) = \mathbb{E}_{\mathbf{z} \sim q} \log p_\theta(x|\mathbf{z}) - \sum_{\ell=1}^L \mathbb{E}_{\mathbf{y} \sim q} \text{KL}(q_\phi(z_\ell|x, \{y_i\}_{i \leq \ell}) || p(z_\ell|y_\ell)) - \sum_{\ell=1}^L \text{KL}(q_\phi(y_\ell|x) || p(y_\ell)) \quad (9)$$

If all $K_\ell = 1$ then VLAC reduces to a VLAE. It is not necessary to have $K_\ell > 1$ for all layers in VLAC.

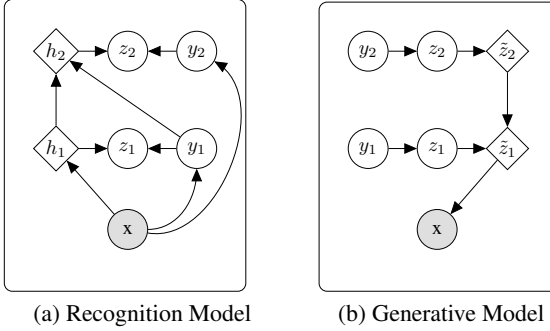


Figure 1: VLAC with $L = 2$

Evaluation Metric Following [13, 31, 16] we use cluster accuracy (ACC), also known as cluster purity, to evaluate our models:

$$\text{ACC} = \max_{P \in \mathcal{P}} \frac{\sum_{i=1}^{|\mathcal{D}|} \mathbb{I}[t_i = Py_i]}{|\mathcal{D}|} \quad (10)$$

where P is a $T \times K$ rectangular permutation matrix that attributes each y to a ground truth class t .

4 Experiments

We trained our model with $L = 4$ and convolutional g and deconvolutional f networks. $p(x|f_0^\theta(\tilde{z}_1))$ was a Gaussian distribution with fixed variance. For full imple-

mentation details, see the code at¹. We used CONCRETE sampling/the Gumbel-Softmax Trick [32, 33] to estimate stochastically the expectations over discrete variables, rather than exactly marginalise them out. This avoids us having to calculate numerous forward passes through the model.

We apply our model to SVHN [2], as it gives variations in style of one type of object while also having distinct ground-truth class structure (here t indexes digit identity) that we can benchmark against.

When running our implementation of a VLAE over SVHN we observed that the 3rd layer was associated most clearly with variation in digit identity. In our experiments we ran VLAC with $\mathbf{K} = \mathbf{K}^{\text{one}} = [1, 1, 50, 1]$ and $\mathbf{K}^{\text{two}} = [1, 5, 50, 1]$. We evaluate the cluster accuracy of y_3 over the test set, taking as our predictions the argmax of the posterior $q_\phi(y_3|x)$.

In addition to published baseline results, we also compared against a single- z -layer Gaussian mixture DGM [3, 4] with an encoder-decoder structure matching that of the sub-networks needed for the 3rd layer of VLAC with $K_3 = 50$.

We also perform class-conditional generation from the layers with $K_\ell > 1$, sampling from each cluster component $p(z_\ell|y_\ell)$. See Figure 3.

5 Discussion

Our model does not achieve state of the art clustering for SVHN. However, we can see that from the \mathbf{K}^{one} results that clustering inside a ladder of stochastic variables is better than an equivalent GM-DGM baseline. And as VLAC with \mathbf{K}^{two} gets better test set accuracy than VLAC with \mathbf{K}^{one} , we see that having a hierarchy of clusters increases performance further still.

From Figures 2 & 3 we can see that our model does separate out class information: class variation is mostly associated with one layer (as in a vanilla VLAE) and cluster components within that layer generally correspond to particular ground truth classes. Figure 3a shows that the model has also

¹<https://github.com/MatthewWilletts/VLAC>

Table 1: Test set cluster accuracy on SVHN for our approach and baselines.

Model	Cluster Accuracy
VLAC with \mathbf{K}^{one}	0.351 ± 0.038
VLAC with \mathbf{K}^{two}	0.378 ± 0.022
Equivalent GM-DGM with $ K = 50$	0.252 ± 0.004
IMSAT [14]	0.573 ± 0.040
DEC [15]	0.560 ± 0.016
ACOL-GAR [16]	0.768 ± 0.013




Figure 2: Decoder means when resampling the latent representation of a datapoint in each layer of VLAC with \mathbf{K}^{two} . For each layer, sample from the marginal $p_\theta(z_\ell) = \sum_i^{K_\ell} p_\theta(z_\ell | y_\ell^i) p(y_\ell^i)$ while keeping all other layers fixed.




Figure 3: Decoder means when sampling from VLAC with \mathbf{K}^{two} from layers with $K_\ell > 1$. Sample from $p_\theta(z_\ell | y_\ell)$ for each cluster component y_ℓ for one layer while keeping all other layers fixed with one sample from their priors. One column per cluster component.

discovered clusters describing the colour temperature of the image. Overall we are pleased to have demonstrated the benefits of clustering in disentangled spaces, and hope that this inspires more research both into how to cluster datasets over different aspects of the data and how disentangling can be used to improve performance of various classical machine learning tasks when working with images.

6 Acknowledgements

We thank Alexander Camuto for his helpful insights.

References

- [1] Shengjia Zhao, Jiaming Song, and Stefano Ermon. Learning Hierarchical Features from Generative Models. In *ICML*, 2017.
- [2] Yuval Netzer. Reading Digits in Natural Images with Unsupervised Feature Learning Yuval. In *NeurIPS Deep Learning and Unsupervised Feature Learning Workshop*, 2011.
- [3] Rui Shu. Gaussian Mixture VAE: Lessons in Variational Inference, Generative Models, and Deep Nets, 2016.
- [4] Eric Nalisnick, Lars Hertel, and Padhraic Smyth. Approximate Inference for Deep Latent Gaussian Mixtures. In *NeurIPS Bayesian Deep Learning Workshop*, 2016.
- [5] Matthew Willetts, Stephen J Roberts, and Christopher C Holmes. Semi-Unsupervised Learning with Deep Generative Models: Clustering and Classifying using Ultra-Sparse Labels. *CoRR*, 2019.
- [6] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the 11th Conference on Computational Learning Theory*, pages 92–100, 1998.
- [7] Steffen Bickel and Tobias Scheffer. Multi-View Clustering. In *ICDM*, 2004.
- [8] Chang Xu and Chao Xu. A Survey on Multi-view Learning. *CoRR*, 2013.
- [9] Hua Wang and Heng Huang. Multi-View Clustering and Feature Learning via Structured Sparsity. In *ICML*, 2013.
- [10] Hongchang Gao. Multi-View Subspace Clustering. In *ICCV*, 2015.
- [11] Handong Zhao, Zhengming Ding, and Yun Fu. Multi-view clustering via deep matrix factorization. In *AAAI*, 2017.
- [12] Nat Dilokthanakul, Pedro A M Mediano, Marta Garnelo, Matthew C H Lee, Hugh Salimbeni, Kai Arulkumaran, and Murray Shanahan. Deep Unsupervised Clustering with Gaussian Mixture VAE. *CoRR*, 2017.
- [13] Zhuxi Jiang, Yin Zheng, Huachun Tan, Bangsheng Tang, and Hanning Zhou. Variational Deep Embedding: An Unsupervised and Generative Approach to Clustering. In *IJCAI*, 2017.
- [14] Weihua Hu, Takeru Miyato, Seiya Tokui, Eiichi Matsumoto, and Masashi Sugiyama. Learning Discrete Representations via Information Maximizing Self-Augmented Training. In *ICML*, 2017.
- [15] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised Deep Embedding for Clustering Analysis. *ICML*, 2016.
- [16] Ozsel Kilinc and Ismail Uysal. Learning Latent Representations in Neural Networks for Clustering Through Pseudo Supervision and Graph-based activity Regularization. In *ICLR*, 2018.
- [17] Hyunjik Kim and Andriy Mnih. Disentangling by Factorising. In *NeurIPS*, 2018.
- [18] Ricky T Q Chen, Xuechen Li, Roger Grosse, and David Duvenaud. Isolating Sources of Disentanglement in Variational Autoencoders. In *NeurIPS*, 2018.
- [19] Babak Esmaeili, Hao Wu, Sarthak Jain, Alican Bozkurt, N Siddharth, Brooks Paige, Dana H Brooks, Jennifer Dy, and Jan-Willem van de Meent. Structured Disentangled Representations. In *AISTATS*, 2019.
- [20] Irina Higgins, David Amos, David Pfau, Sebastien Racaniere, Loic Matthey, Danilo Rezende, and Alexander Lerchner Deepmind. Towards a Definition of Disentangled Representations. *CoRR*, 2018.
- [21] Emile Mathieu, Tom Rainforth, N. Siddharth, and Yee Whye Teh. Disentangling Disentanglement in Variational Autoencoders. In *ICML*, 2019.

- [22] Prasoon Goyal, Zhiting Hu, Xiaodan Liang, Chenyu Wang, and Eric P Xing. Nonparametric Variational Auto-Encoders for Hierarchical Representation Learning. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2017-Octob, pages 5104–5112, 2017.
- [23] Yee Whye Teh, Michael I Jordan, Matthew J Beal, and David M Blei. Hierarchical Dirichlet processes. *Journal of the American Statistical Association*, 101(476):1566–1581, 2006.
- [24] Xiaopeng Li, Zhouong Chen, Leonard K.M. Poon, and Nevin L Zhang. Learning latent superstructures in variational autoencoders for deep multidimensional clustering. In *ICLR*, 2019.
- [25] J Deng, W Dong, R Socher, L.-J. Li, K Li, and L Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.
- [26] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. In *NeurIPS*, 2013.
- [27] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *ICML*, 2014.
- [28] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. Ladder Variational Autoencoders. In *NeurIPS*, 2016.
- [29] Antti Rasmus, Harri Valpola, Mikko Honkala, Mathias Berglund, and Tapani Raiko. Semi-Supervised Learning with Ladder Networks. In *NeurIPS*, 2015.
- [30] Mohammad Pezeshki, Linxi Fan, Philemon Brakel, Aaron Courville, and Yoshua Bengio. Deconstructing the ladder network architecture. In *ICML*, 2016.
- [31] Jianwei Yang, Devi Parikh, Dhruv Batra, and Virginia Tech. Joint Unsupervised Learning of Deep Representations and Image Clusters. In *CVPR*, 2016.
- [32] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *ICLR*, 2017.
- [33] Eric Jang, Shixiang Gu, and Ben Poole. Categorical Reparameterization with Gumbel-Softmax. In *ICLR*, 2017.

A Sampling from the prior of our GM-DGM

Comparing Figure 4 to Figure 3, we see that the GM-DGM encoding in z entangles class, colour and style information, unlike VLAC.



Figure 4: Decoder means when sampling from the latent layer of GM-DGM with $|K| = 50$, sampling from $p_\theta(z|y)$ for each cluster component y .