

# PROGRESSIVE LEARNING AND DISENTANGLEMENT OF HIERARCHICAL REPRESENTATIONS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Learning rich representation from data is an important task for deep generative models such as variational auto-encoder (VAE). However, by extracting high-level abstractions in the bottom-up inference process, the goal of preserving all factors of variations for top-down generation is compromised. Motivated by the concept of “starting small”, we present a strategy to progressively learn independent hierarchical representations from high- to low-levels of abstractions. The model starts with learning the most abstract representation, and then progressively grow the network architecture to introduce new representations at different levels of abstraction. We quantitatively demonstrate the ability of the presented model to improve disentanglement in comparison to existing works on two benchmark data sets using three disentanglement metrics, including a new metric we proposed to complement the previously-presented metric of mutual information gap. We further present both qualitative and quantitative evidence on how the progression of learning improves disentangling of hierarchical representations. By drawing on the respective advantage of hierarchical representation learning and progressive learning, this is to our knowledge the first attempt to improve disentanglement by progressively growing the capacity of VAE to learn hierarchical representations.

## 1 INTRODUCTION

Variational auto-encoder (VAE), a popular deep generative model (DGM), has shown great promise in learning interpretable and semantically meaningful representations of data (Higgins et al. (2017); Chen et al. (2018); Kim & Mnih (2018)). However, VAE has not been able to fully utilize the depth of neural networks like its supervised counterparts, for which a fundamental cause lies in the inherent conflict between the bottom-up inference and top-down generation process (Zhao et al. (2017); Li et al. (2016)): while the bottom-up abstraction is able to extract high-level representations helpful for discriminative tasks, the goal of generation requires the preservation of all generative factors that are likely at different abstraction levels. This issue was addressed in recent works by allowing VAEs to generate from details added at different depths of the network, using either memory modules between top-down generation layers (Li et al. (2016)), or hierarchical latent representations extracted at different depths via a variational ladder autoencoder (VLAE, Zhao et al. (2017)).

However, it is difficult to learn to extract and disentangle all generative factors at once, especially at different abstraction levels. Inspired by human cognition system, Elman (1993) suggested the importance of “starting small” in two aspects of the learning process of neural networks: *incremental input* in which a network is trained with data and tasks of increasing complexity, and *incremental memory* in which the network capacity undergoes developmental changes given fixed external data and tasks — both pointing to an incremental learning strategy for simplifying a complex final task. Indeed, the former concept of *incremental input* has underpinned the success of curriculum learning (Bengio et al. (2015)). In the context of DGMs, various stacked versions of generative adversarial networks (GANs) have been proposed to decompose the final task of high-resolution image generation into progressive sub-tasks of generating small to large images (Denton et al. (2015); Zhang et al. (2018)). The latter aspect of “starting small” with incremental growth of network capacity is less explored, although recent works have demonstrated the advantage of progressively growing the depth of GANs for generating high-resolution images (Karras et al. (2018); Wang et al. (2018)). These works, so far, have focused on progressive learning as a strategy to improve image generation.

We are motivated to investigate the possibility to use progressive learning strategies to improve learning and disentangling of hierarchical representations. At a high level, the idea of progressively or

sequentially learning latent representations has been previously considered in VAE. In Gregor et al. (2015), the network learned to sequentially refine generated images through recurrent networks. In Lezama (2019), a teacher-student training strategy was used to progressively increase the number of latent dimensions in VAE to improve the generation of images while preserving the disentangling ability of the teacher model. However, these works primarily focus on progressively growing the capacity of VAE to generate, rather than to extract and disentangle hierarchical representations.

In comparison, in this work, we focus on 1) progressively growing the capacity of the network to extract hierarchical representations, and 2) these hierarchical representations are extracted and used in generation from different abstraction levels. We present a simple progressive training strategy that grows the hierarchical latent representations from different depths of the inference and generation model, learning from high- to low-levels of abstractions as the capacity of the model architecture grows. Because it can be viewed as a progressive strategy to train the VLAE presented in Zhao et al. (2017), we term the presented model *pro*-VLAE. We quantitatively demonstrate the ability of *pro*-VLAE to improve disentanglement on two benchmark data sets using three disentanglement metrics, including a new metric we proposed to complement the metric of mutual information gap (MIG) previously presented in Chen et al. (2018). These quantitative studies include comprehensive comparisons to  $\beta$ -VAE (Higgins et al. (2017)), VLAE (Zhao et al. (2017)), and the teacher-student strategy as presented in (Lezama (2019)) at different values of the hyperparameter  $\beta$ . We further present both qualitative and quantitative evidence that *pro*-VLAE is able to first learn the most abstract representations and then progressively disentangle existing factors or learn new factors at lower levels of abstraction, improving disentangling of hierarchical representations in the process.

## 2 RELATED WORKS

A hierarchy of feature maps can be naturally formed in stacked discriminative models (Zeiler & Fergus (2014)). Similarly, in DGM, many works have proposed stacked-VAEs as a common way to learn a hierarchy of latent variables and thereby improve image generation (Sønderby et al. (2016); Bachman (2016); Kingma et al. (2016)). However, this stacked hierarchy is not only difficult to train as the depths increases (Sønderby et al. (2016); Bachman (2016)), but also has an unclear benefit for learning either hierarchical or disentangled representations: as shown in Zhao et al. (2017), when fully optimized, it is equivalent to a model with a single layer of latent variables. Alternatively, instead of a hierarchy of latent variables, independent hierarchical representations at different abstraction levels can be extracted and used in generation from different depths of the network (Rezende et al. (2014); Zhao et al. (2017)). A similar idea was presented in Li et al. (2016) to generate lost details from memory and attention modules at different depths of the top-down generation process. The presented work aligns with existing works (Rezende et al. (2014); Zhao et al. (2017)) in learning independent hierarchical representation from different levels of abstraction, and we look to facilitate this learning by progressively learning the representations from high- to low-levels.

Progressive learning has been successful for high-quality image generation, mostly in the setting of GANs. Following the seminar work of Elman (1993), these progressive strategies can be loosely grouped into two categories. Mostly, in line with *incremental input*, several works have proposed to divide the final task of image generation into progressive tasks of generating low-resolution to high-resolution images with multi-scale supervision (Denton et al. (2015); Zhang et al. (2018)). Alternatively, in line with *incremental memory*, a small number of works have demonstrated the ability to simply grow the architecture of GANs from a shallow network with limited capacity for generating low-resolution images, to a deep network capable of generating super-resolution images (Karras et al. (2018); Wang et al. (2018)). This approach was also shown to be time-efficient since the early-stage small networks require less time to converge comparing to training a full network from the beginning. This latter group of works provided compelling evidence for the benefit of progressively growing the capacity of a network to generate images, although its extension for growing the capacity of a network to learn hierarchical representations has not been explored.

Limited work has considered incremental learning of representations in VAE. In Gregor et al. (2015), recurrent networks with attention mechanisms were used to sequentially refine the details in generated images. It however focused on the generation performance of VAE without considering the learned representations. In Lezama (2019), a teacher-student strategy was used to progressively grow the dimension of the latent representations in VAE. Its fundamental motivation was that, given a teacher model that has learned to effectively disentangle major factors of variations, progressively

learning additional nuisance variables will improve generation without compromising the disentangling ability of the teacher – the latter accomplished via a newly-proposed Jacobian supervision. The capacity of this model to grow, thus, is by design limited to the extraction of nuisance variables. In comparison, we are interested in a more significant growth of the VAE capacity to progressively learn and improve disentangling of important factors of variations which, as we will later demonstrate, is not what the model in Lezama (2019) is intended for. In addition, neither of these works considered learning different levels of abstractions at different depths of the network, and the presented *pro*-VLAE provides a simpler training strategy to achieve progressive representation learning.

Learning disentangled representation is a primary motivation of our work, and an important topic in VAE. Existing works mainly tackle this by promoting the independence among the learned latent factors in VAE (Higgins et al. (2017); Kim & Mnih (2018); Chen et al. (2018)). The presented progressive learning strategy provides an entirely different approach to improve disentangling that is orthogonal to these existing methods and a possibility to augment them in the future.

### 3 METHODS

#### 3.1 MODEL: VAE WITH HIERARCHICAL REPRESENTATIONS

We assume a generative model  $p(\mathbf{x}, \mathbf{z}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z})$  for observed  $\mathbf{x}$  and its latent variable  $\mathbf{z}$ . To learn hierarchical representations of  $\mathbf{x}$ , we decompose  $\mathbf{z}$  into  $\{z_1, z_2, \dots, z_L\}$  with  $z_l (l = 1, 2, 3, \dots, L)$  from different abstraction levels. We define the hierarchical generative model  $p_\theta$  as:

$$p(\mathbf{x}, \mathbf{z}) = p(\mathbf{x}|z_1, z_2, \dots, z_L) \prod_{l=1}^L p(z_l). \quad (1)$$

Note that there is no hierarchical dependence among the latent variables as in common hierarchical latent variable models. Rather, similar to that in Rezende et al. (2014) and Zhao et al. (2017),  $z_l$ 's are independent and each represents generative factors at an abstraction level not captured in other levels. We then define an inference model  $q_\phi$  to approximate the posterior as:

$$q(z_1, z_2, \dots, z_L|\mathbf{x}) = \prod_{l=1}^L q(z_l|\mathbf{h}_l(\mathbf{x})), \quad (2)$$

where  $\mathbf{h}_l(\mathbf{x})$  represents a particular level of bottom-up abstraction of  $\mathbf{x}$ . We parameterize  $p_\theta$  and  $q_\phi$  with an encoding-decoding structure and, as in Zhao et al. (2017), we approximate the abstraction level with the network depth. The full model is illustrated in Fig. 1(c), with a final goal to maximize a modified evidence lower bound (ELBO) of the marginal likelihood of data  $\mathbf{x}$ :

$$\log p(\mathbf{x}) \geq \mathcal{L} = \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x}|\mathbf{z})] - \beta KL(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})), \quad (3)$$

where  $KL$  denotes the Kullback-Leibler divergence, prior  $p(\mathbf{z})$  is set to isotropic Gaussian  $\mathcal{N}(0, \mathbf{I})$  according to standard practice, and  $\beta$  is a hyperparameter introduced in Higgins et al. (2017) to promote disentangling, defaulting to the standard ELBO objective when  $\beta = 1$ .

#### 3.2 PROGRESSIVE LEARNING OF HIERARCHICAL REPRESENTATION

We present a progressive learning strategy, as illustrated in Fig. 1, to achieve the final goal in equation (3) by learning the latent variables  $z_l$  progressively from the highest ( $l = L$ ) to the lowest ( $l = 1$ ) level of abstractions. We start by learning the most abstraction representations at layer  $L$  as show in Fig. 1(a). In this case, our model degenerates to a vanilla VAE with latent variables  $z_L$  at the deepest layer. We keep the dimension of  $z_L$  small to start small in terms of the capacity to learn latent representations, where we define the inference model at progressive step  $s = 0$  as:

$$z_L \sim \mathcal{N}(\mu_L(\mathbf{h}_L), \sigma_L(\mathbf{h}_L)), \mathbf{h}_l = f_l^e(\mathbf{h}_{l-1}), \text{ for } l = 1, 2, \dots, L, \text{ and } \mathbf{h}_0 \equiv \mathbf{x}, \quad (4)$$

and the generative model as:

$$\mathbf{g}_L = f_L^d(z_L), \mathbf{g}_l = f_l^d(\mathbf{g}_{l+1}), \mathbf{x} = D(\mathbf{x}; f_0^d(\mathbf{g}_0)), \quad (5)$$

where  $f_l^e$ ,  $\mu_L$ , and  $\sigma_L$  are parts of the encoder architecture,  $f_l^d$  are parts of the decoder architecture, and  $D$  is the distribution of  $\mathbf{x}$  parametrized by  $f_0^d(\mathbf{g}_0)$ , which can be either Bernoulli or Gaussian depending on the data. Next, as shown in Fig. 1, we progressively grow the model to learn

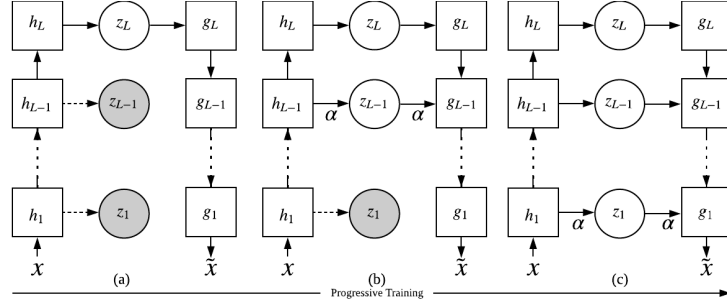


Figure 1: Progressive learning of hierarchical representations. White blocks and solid lines are VAE models at the current progression.  $\alpha$  is a fade-in coefficient for blending in the new network component. Gray circles and dash line represents (optional) constraining of the future latent variables.

$z_{L-1}, \dots, z_2, z_1$  from high to low abstraction levels. At each progressive step  $s = 1, 2, \dots, L - 1$ , we move down one abstraction level, and grow the inference model by introducing new latent code:

$$z_l \sim \mathcal{N}(\mu_l(\mathbf{h}_l), \sigma_l(\mathbf{h}_l)), \quad l = L - s. \quad (6)$$

Simultaneously, we grow the decoder such that it can generate with the new latent code as:

$$\mathbf{g}_l = f_l^d([m_l(z_l); \mathbf{g}_{l+1}]), \quad l = L - s, \quad (7)$$

where  $m_l$  includes transposed convolution layers outputting a feature map in the same shape as  $\mathbf{g}_{l+1}$ , and  $[\cdot; \cdot]$  denotes a concatenation operation. The training objective at progressive step  $s$  is then:

$$\mathcal{L}_{pro} = \mathbb{E}_{q(z_L, z_{L-1}, \dots, z_{L-s} | \mathbf{x})} [\log p(\mathbf{x} | z_L, z_{L-1}, \dots, z_{L-s})] - \beta \sum_{l=L-s}^L KL(q(z_l | \mathbf{x}) || p(z_l)), \quad (8)$$

By replacing the full objective in equation (3) with a sequence of the objectives in equation (8) as the training progresses, we incrementally learn to *extract* and *generate with* hierarchical latent representations  $z_l$ 's from high to low levels of abstractions. Once trained, the full model as shown in Fig. 1(c) will be used for inference and generation, and progressive processes are no longer needed.

### 3.3 IMPLEMENTATION STRATEGIES

Two important strategies are utilized to implement the proposed progressive representation learning. First, directly adding new components to a trained network often introduce a sudden shock to the gradient: in VAEs, this often leads to the explosion of the variance in the latent distributions. To avoid this shock, we adopt the popular method of "fade-in" (Karras et al. (2018)) to smoothly blend the new and existing network components. In specific, we introduce a "fade-in" coefficient  $\alpha$  to equations (6) and (7) when growing new components in the encoder and the decoder:

$$z_l \sim \mathcal{N}(\mu_l(\alpha \mathbf{h}_l), \sigma_l(\alpha \mathbf{h}_l)), \mathbf{g}_l = f_l^d([\alpha m_l(z_l); \mathbf{g}_{l+1}]), \quad (9)$$

where  $\alpha$  increases from 0 to 1 within a certain number of iterations (5000 in our experiments) since the addition of the new network components  $\mu_l, \sigma_l$ , and  $m_l$ .

Second, we further stabilize the training by weakly constraining the distribution of  $z_l$ 's before they are added to the network. This can be achieved by applying a KL penalty, modulated by a small coefficient  $\gamma$ , to all latent variables that have not been used in the generation at progressive step  $s$ :

$$\mathcal{L}_{pre-trained} = \gamma \sum_{l=1}^{L-s-1} [-KL(q(z_l | \mathbf{x}) || p(z_l))], \quad (10)$$

where  $\gamma$  is set to 0.5 in our experiments. The final training objective at step  $s$  then becomes:

$$\mathcal{L} = \mathcal{L}_{pro} + \mathcal{L}_{pre-trained} \quad (11)$$

Note that the latent variables at the hierarchy lower than  $L - s$  are neither meaningfully inferred nor used in generation at progressive step  $s$ , and  $\mathcal{L}_{pre-trained}$  merely intends to regularize the distribution of these latent variables before they are added to the network. In the experiments below, we use both "fade-in" and  $\mathcal{L}_{pre-trained}$  when implementing the progressive training strategy.

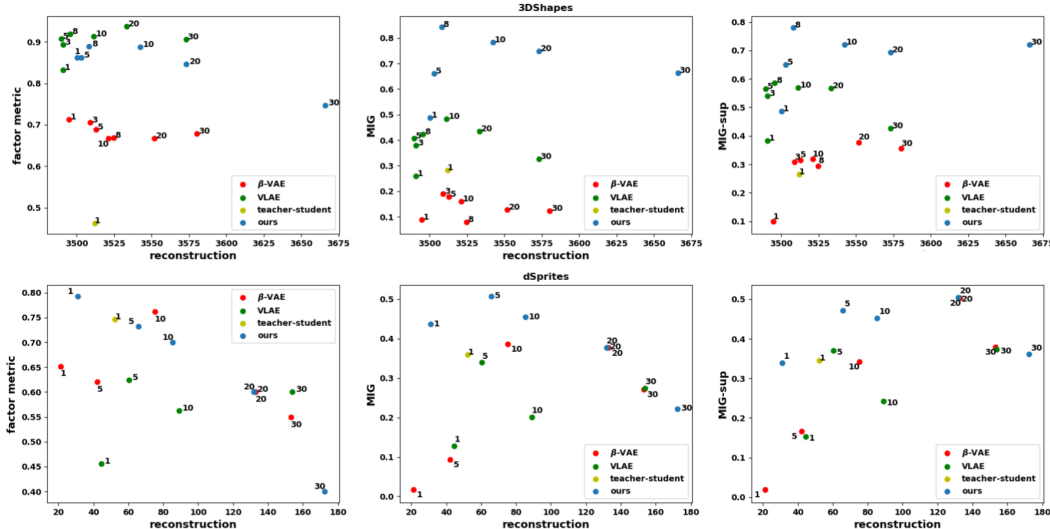


Figure 2: Quantitative comparison of disentanglement metrics. Each point is annotated by the  $\beta$  value and averaged over top three best random seeds for the given  $\beta$  on the give model. **Left to right:** reconstruction errors vs. disentanglement metrics of factor, MIG, and MIG-sup, a higher value indicating a better disentanglement in each metric.

### 3.4 DISENTANGLEMENT METRIC

Various quantitative metrics for measuring disentanglement have been proposed (Higgins et al. (2017); Kim & Mnih (2018); Chen et al. (2018)). However, to date, no metrics have been able to measure the extent to which a network can provide a one-to-one relationship between a latent dimension and a generative factor. For instance, the recently proposed MIG metrics (Chen et al. (2018)) measures the gap of mutual information between the top two latent dimensions that have the highest mutual information with a given generative factor. A low MIG score, therefore, suggests an undesired outcome that the same factor is split into multiple dimensions. However, if different generative factors are entangled into the same latent dimension, the MIG score will not be affected.

Therefore, we propose a new disentanglement metric to supplement MIG by recognizing the entanglement of multiple generative factors into the same latent dimension. We define MIG-sup as:

$$\frac{1}{J} \sum_1^J (I_{norm}(z_j; v_{k^{(j)}}) - \max_{k \neq k^{(j)}} I_{norm}(z_j; v_k)), \quad (12)$$

where  $k^{(j)} = \operatorname{argmax}_k I_n(z_j; v_k)$ ,  $J$  is the number of meaningful latent dimensions, and  $I_{norm}(z_j; v_k)$  is normalized mutual information  $I(z_j; v_k)/H(v_k)$ . Considering MIG and MIG-sup together will provide a more complete measure of disentanglement, accounting for both the splitting of one factor into multiple dimensions and the encoding of multiple factors into the same dimension. In an ideal disentanglement, both MIG and MIG-sup should be 1. Note that the factor metric (Kim & Mnih (2018)) has similar properties with MIG-sup, although MIG-sup is more strict on penalizing any amount of other minor factors in the same dimension.

## 4 EXPERIMENT

We tested the presented *pro*-VLAE on four benchmark data sets: dSprites (Matthey et al. (2017)), 3DShapes (Kim & Mnih (2018)), MNIST (LeCun et al. (1998)), and CelebA (Liu et al. (2015)), where the first two include ground-truth generative factors that allow us to carry out comprehensive quantitative comparisons of disentangling metrics with existing models. In the following, we first quantitatively compare the disentangling ability of *pro*-VLAE in comparison to three existing models using three disentanglement metrics. We then analyze *pro*-VLAE from the aspects of how it learns progressively, its ability to disentangle, and its ability to learn abstractions at different levels.

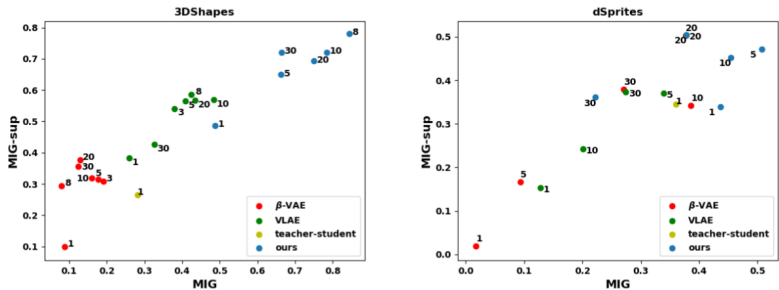


Figure 3: MIG vs. MIG-sup following a similar presentation in Fig. 2. A better disentanglement should have higher MIG and higher MIG-sup, locating at the top-right quadrant of the plot.

**Comparisons in quantitative disentanglement metrics:** For quantitative comparisons, we considered the factor metric in Kim & Mnih (2018), the MIG in Chen et al. (2018), and the MIG-sup presented in this work. We compared *pro*-VLAE (changing  $\beta$ ) with beta-VAE (Higgins et al. (2017)), VLAE (Zhao et al. (2017)) as a hierarchical baseline without progressive training, and the teacher-student model (Lezama (2019)) as the most related progressive VAE without hierarchical representations. All models were considered at different values of  $\beta$  except the teacher-student model: the comparison of  $\beta$ -VAE, VLAE, and the presented *pro*-VLAE thus also provides an ablation study on the effect of learning hierarchical representations and doing so in a progressive manner.

For fair comparisons, we strictly required all models to have the same number of latent variables and the same number of training iterations. For instance, if a hierarchical model has three layers that each has three latent dimensions, a non-hierarchical model will have nine latent dimensions; if a progressive method has three progressive steps with 15 epochs of training each, a non-progressive method will be trained for 45 epochs. Three to five experiments were conducted for each model at each  $\beta$  value, and the average of the top three is used for reporting the quantitative results in Fig. 2.

As shown, for MIG and MIG-sup, VLAE generally outperformed  $\beta$ -VAE at most  $\beta$  values, while *pro*-VLAE showed a clear margin of improvement over both methods. With the factor metric, *pro*-VLAE was still among the top performers, although with a smaller margin and a larger overlap with VLAE on 3DShapes, and with  $\beta$ -VAE ( $\beta = 10$ ) on dSprites. The teacher-student strategy with Jacobian supervision in general had a low to moderate disentangling score, especially on 3DShapes. This is consistent with the original motivation of the method for progressively learning nuisance variables after the teacher learns to disentangle effectively, rather than progressively disentangling hierarchical factors of variations as intended by *pro*-VLAE. Note that *pro*-VLAE in general performed better with a smaller value of  $\beta$  ( $\beta < 20$ ), suggesting that progressive learning already had an effect of promoting disentangling without the need of large  $\beta$  regularization on the KL penalty.

Fig. 3 shows MIG vs. MIG-sup scores among the tested models. As shown, results from *pro*-VLAE were well separated from the other three models at the right top quadrant of the plots, obtaining simultaneously high MIG and MIG-sup scores as a clear evidence for improved disentangling ability.

Fig. 4 provides images generated by traversing each latent dimension using the best *pro*-VLAE ( $\beta = 8$ ), the best VLAE ( $\beta = 10$ ), and the teacher-student model on 3DShapes data. As shown, *pro*-VLAE learned to disentangle the object, wall, and floor color in the deepest layer; the following hierarchy of representations then disentangled objective scale, orientation, and shape, while the lowest-level of abstractions ran out of meaningful generative factors to learn. In comparison, the VLAE distributed six generative factors over the nine latent dimensions, where color was split across the hierarchy and sometimes entangled with the object scale (in  $z_2$ ). The teacher-student model was much less disentangled, which we will delve into further in the following section.

**Information flow during progressive learning:** To further understand what happened during progressive learning, we use mutual information  $I(\mathbf{x}, \mathbf{z}_l)$  as a surrogate to track the amount of information learned in each hierarchy of latent variables  $\mathbf{z}_l$  during the progressive learning. We adopted the approach in Chen et al. (2018) to empirically estimate the mutual information by stratified sampling.

Fig. 5 shows an example from 3DShapes. At progressive step 0, *pro*-VAE was only learning the deepest latent variables in  $\mathbf{z}_3$ , discovering most of the generative factors including color, objective

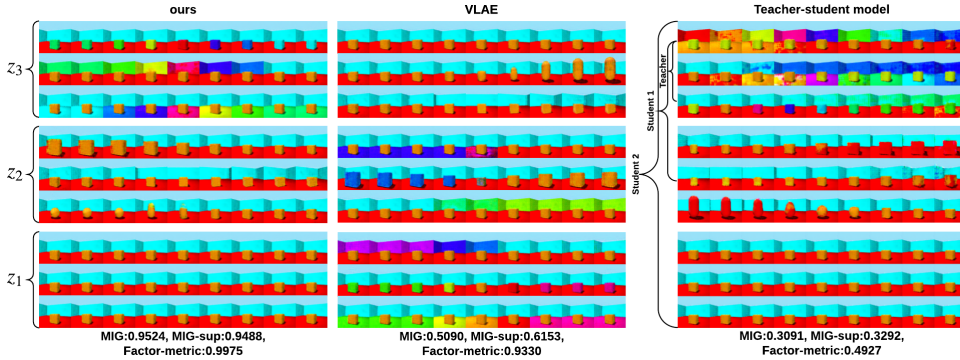


Figure 4: Traversing each latent dimension in *pro*-VLAE ( $\beta = 8$ ), VLAE ( $\beta = 10$ ), and teacher-student model. The hierarchy of the latent variables is noted by brackets on the side.

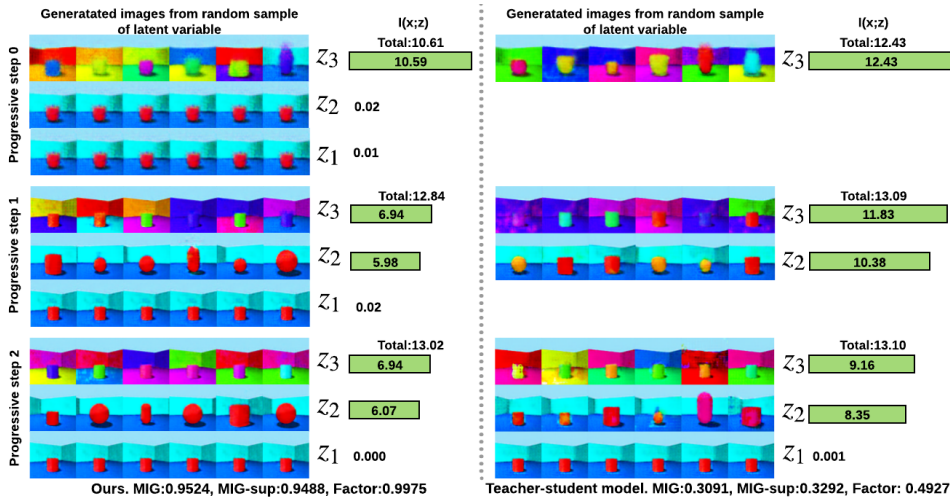


Figure 5: Progressive learning of hierarchical representations. At each progression and for each  $z_l$ , the row of images are generated by randomly sampling from its prior distributions while fixing the other latent variables (this is NOT traversing). The green bar at each row tracks the mutual information  $I(x; z_l)$ , while the total mutual information  $I(x; z)$  is labeled on top.

shape, and orientation entangled within  $z_3$ . At progressive step 1, interestingly, the model was able to “drag” out shape and rotation factors from  $z_3$  and disentangle them into  $z_2$  along with a new scale factor. Thus  $I(x; z_3)$  decreased from 10.59 to 6.94 while  $I(x; z_2)$  increased from 0.02 to 5.98 in this progression, while the total mutual information  $I(x; z)$  increased from 10.61 to 12.84, suggesting the overall learning of more detailed information. Since 3DShapes only has 6 factors, the lowest-level representation  $z_1$  had nothing to learn in progressive step 2, and the allocation of mutual information remained nearly unchanged. Note that the sum of  $I(x, z_l)$ ’s does not equal to  $I(x, z)$  and  $I_{over} = \sum_1^L I(x, z_l) - I(x, z)$  suggests the amount of information that is entangled.

In comparison, the teacher-student model was less effective in progressively dragging entangled representations to newly added latent dimensions, as suggested by the slowing changing of  $I(x, z_l)$ ’s during progression and the larger value of  $I_{over}$ . This suggests that, since the teacher-student model was motivated for progressively learning nuisance variables, the extent to which its capacity can grow for learning new representations is limited by two fundamental causes: 1) because it increases the dimension of the same latent vectors at the same depth, the growth of the network capacity is limited in comparison to *pro*-VLAE, and 2) the Jacobian supervision further restricts the student model to maintain the same disentangling ability of the teacher model.

**Disentangling hierarchical representations:** We also qualitatively examined *pro*-VLAE on data with both relatively simple (MNIST) and complex (CelebA) factors of variations, all done in unsupervised training. On MNIST (Figure 6), while the deepest latent representations encoded the




Figure 6: Visualization of hierarchical features learnt for MNIST data. Each sub-figure is generated by randomly sampling from the prior distribution of  $z_l$  at one abstraction level while fixing the others. The original latent code is inferred from a image with digit “0”. **From left to right:**  $z_3$  encodes the highest abstraction: digit identity;  $z_2$  encodes stroke width; and  $z_1$  encodes other digit styles.

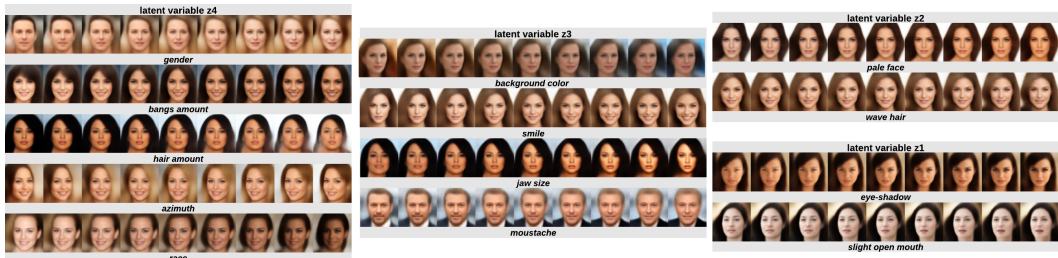


Figure 7: Visualization of hierarchical features learnt for CelebA data. Each subfigure is generated by traversing along a selected latent dimension in each row within each hierarchy of  $z_l$ 's. **From left to right:** latent variables  $z_4$  to  $z_1$  progressively learn major (e.g., gender in  $z_4$  and smile in  $z_3$ ) to minor representations (e.g. wavy-hair in  $z_2$  and eye-shadow in  $z_1$ ) in a disentangled manner.

highest-level features in terms of digit identity, the representations learned at shallower levels encoded changes in writing styles. In Figure 7, we show the latent representation progressively learned in CelebA from the highest to lowest levels of abstractions, along with disentangling within each level demonstrated by traversing one selected dimension at a time. These dimensions are selected as examples associated with clear semantic meanings. As shown, while the deepest latent representation  $z_4$  learned to disentangle high-level features such as gender and race, the shallowest representation  $z_1$  learned to disentangle low-level features such as eye-shadow. Moreover, the number of distinct representations learned decreased from deep to shallow layers. While demonstrating disentangling by traversing each individual latent dimension or by hierarchically-learned representations has been separately reported in previous works (Higgins et al. (2017); Zhao et al. (2017)), to our knowledge this is the first time the ability of a model to disentangle individual latent factors in a hierarchical manner has been demonstrated. This provides evidence that the presented progressive strategy of learning can improve the disentangling of first the most abstract representations followed by progressively lower levels of abstractions.

## 5 CONCLUSION

In this work, we present a progressive strategy for learning and disentangling hierarchical representations. Starting from a simple VAE, the model first learn the most abstract representation. Next, the model learn independent representations from high- to low-levels of abstraction by progressively growing the capacity of the VAE deep to shallow. Experiments on several benchmark data sets demonstrated the advantages of the presented method. An immediate future work is to include stronger guidance for allocating information across the hierarchy of abstraction levels, either through external multi-scale image supervision or internal information-theoretic regularization strategies.



## REFERENCES

- Philip Bachman. An architecture for deep, hierarchical generative models. In Advances in Neural Information Processing Systems, pp. 4826–4834, 2016.
- Yoshua Bengio, Jerome Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In International Conferences on Machine Learning, 2015.
- Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In Advances in Neural Information Processing Systems, pp. 2610–2620, 2018.
- Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In Advances in neural information processing systems, pp. 1486–1494, 2015.
- Jeffrey L Elman. Learning and development in neural networks: The importance of starting small. Cognition, 48(1):71–99, 1993.
- Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra. Draw: A recurrent neural network for image generation. In International Conference on Machine Learning, 2015.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. ICLR, 2(5):6, 2017.
- Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In International Conference on Learning Representations, 2018.
- Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In International Conference on Machine Learning, 2018.
- Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In Advances in neural information processing systems, pp. 4743–4751, 2016.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11):2278–2324, 1998.
- José Lezama. Overcoming the disentanglement vs reconstruction trade-off via jacobian supervision. In International Conference on Learning Representations, 2019.
- Chongxuan Li, Jun Zhu, and Bo Zhang. Learning to generate with memory. In International Conference on Machine Learning, pp. 1177–1186, 2016.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In Proceedings of the IEEE international conference on computer vision, pp. 3730–3738, 2015.
- Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In International Conference on Machine Learning, 2014.
- Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. Ladder variational autoencoders. In Advances in neural information processing systems, pp. 3738–3746, 2016.
- Yifan Wang, Federico Perazzi, Brian McWilliams, Alexander Sorkine-Hornung, Olga Sorkine-Hornung, and Christopher Schroers. A fully progressive approach to single-image super-resolution. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 864–873, 2018.


Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In European conference on computer vision, pp. 818–833. Springer, 2014.

Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan++: Realistic image synthesis with stacked generative adversarial networks. IEEE transactions on pattern analysis and machine intelligence, 41(8):1947–1962, 2018.

Shengjia Zhao, Jiaming Song, and Stefano Ermon. Learning hierarchical features from deep generative models. In Proceedings of the 34th International Conference on Machine Learning-Volume 70, pp. 4091–4099. JMLR. org, 2017.

## A APPENDIX

### A.1 EXAMPLES: PERFORMANCE OF DIFFERENT METRICS



**MIG: 0.1295, MIG-sup:0.4225, factor: 0.6653**

Figure 8: An example of one factor being encoded in multiple dimensions. Each row is a traverse for one dimension (dimension order adjusted for better visualization). Notice that both dim1 and dim2 are encoding floor-color, both dim3 and dim4 are encoding wall-color, and both dim5 and dim6 are encoding object color. Therefore, the MIG is very low since it penalizes splitting one factor to multiple dimensions. On the other hand, the MIG-sup and factor-metric is not too bad since one dimension mainly encodes one factor, even though there are some entanglement of color-vs-shape and color-vs-scale.

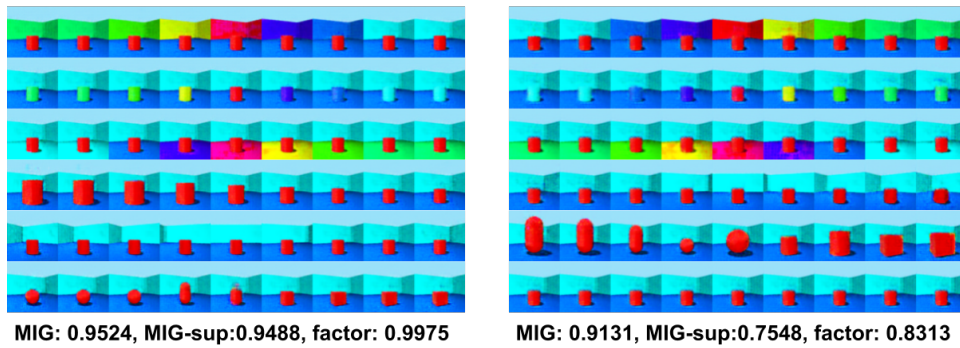


Figure 9: An example of one dimension containing multiple factors. Each row is a traverse for one dimension (dimension order adjusted for better visualization). Notice that both models achieve high and similar MIG because all 6 factors are encoded and no splitting to multiple dimensions. However, the right model has much lower MIG-sup and factor-metric because than left model. Because both scale and shape is encoded in dim5 and dim6 has no factor. And both MIG-sup and factor-metric penalizes encoding multiple factors in one dimensions. Besides, our MIG-sup is lower and drops more than factor-metric because MIG-sup is more strict about this case.