

UNSUPERVISED LEARNING OF SLOW FEATURES FOR DATA EFFICIENT REGRESSION

Anonymous authors

Paper under double-blind review

ABSTRACT

Research in computational neuroscience suggests that the human brain’s unparallelled data efficiency is a result of highly efficient mechanisms to extract and organize slowly changing high level features from continuous sensory inputs. In this paper, we apply this *slowness principle* to a state of the art representation learning method with the goal of performing data efficient learning of down-stream regression tasks. To this end, we propose the *slow variational autoencoder* (S-VAE), an extension to the β -VAE which applies a temporal similarity constraint to the latent representations. We empirically compare our method to the β -VAE and the Temporal Difference VAE (TD-VAE), a state-of-the-art method for next frame prediction in latent space with temporal abstraction. We evaluate the three methods against their data-efficiency on down-stream tasks using a synthetic 2D ball tracking dataset and a dataset generated using the DeepMind Lab environment. In both tasks, the proposed method outperformed the baselines both with dense and sparse labeled data. Furthermore, the S-VAE achieved similar performance compared to the baselines with 1/5 to 1/11 of data.

1 INTRODUCTION

Neuroscience suggests that a major difference between state of the art deep learning architectures and the human brain is that cells in the brain do not react to single stimuli, but instead extract invariant features from sequences of fast changing sensory input signals (Bengio & Bergstra, 2009). Evidence found in the hierarchical organization of simple and complex vision cells shows that time-invariance is the principle after which the cortex extracts the underlying generative factors of these sequences and that these factors usually change slower than the observed signal (Wiskott & Sejnowski, 2002; Berkes & Wiskott, 2005; Bengio & Bergstra, 2009). Computational neuroscientists have named this paradigm the *slowness principle* wherein individual measurements of a signal may vary quickly, but the underlying generative features vary slowly. For example, individual pixel values in a video change rapidly during short periods of time, but the scene itself changes slowly. This principle has found application in *Slow Feature Analysis* (SFA) (Wiskott & Sejnowski, 2002), transformation- and time-invariant object detection (Franzius et al., 2011; Zou et al., 2011), and neural network pre-training (Bengio & Bergstra, 2009).

In this paper, we apply the *slowness principle* to a state-of-the-art representation learning method, the β -Variational Autoencoder (Higgins et al., 2017), by adding a similarity loss term to the evidence lower bound (ELBO) that encourages similarity between latent representations based on their temporal proximity. We show that the slow representations that hence emerge improve the task performance and data efficiency of down-stream few-shot regression tasks. We compare the proposed method to state-of-the-art representation learning methods, the β -VAE and the Temporal Difference Variational Autoencoder (TD-VAE) (Gregor et al., 2019). Furthermore, we investigate the structure of the slow latent space and its influence on the task performance with respect to bias and variance of the down-stream model. The key **contributions** of this paper are:

- We propose the slow variational autoencoder (S-VAE) that extends the β -VAE with a similarity constraint in the latent space which imposes the slowness property on the latent representations.
- We empirically show that slow representations lead to more data-efficient learning of down-stream regression tasks for two datasets. We show that the task performance of the S-VAE

is between 2.38% and 17.6% better depending on the dataset and requires 4 to 11 times less data to achieve similar performance than other state of the art methods.

- We analyze the structure of the resulting latent spaces and show that slow latent spaces reduce the bias and variance of down-stream models.

2 SLOW VARIATIONAL AUTOENCODER

The variational autoencoder (VAE) (Kingma & Welling, 2013) is a representation learning method for dimensionality reduction using a loss on the reconstruction quality of observations decoded from the low dimensional representations. The VAE consists of an encoder q_θ parameterized by θ that returns a lower dimensional approximate posterior with unit Gaussian prior and a decoder p_ϕ parameterized by ϕ that reconstructs the input from the latent posterior. Higgins et al. (2017) introduced the β -VAE by adding a parameter β to adjust the weight of the KL divergence between the prior and the posterior to allows a trade-off between disentanglement of the latent factors and reconstruction quality. We introduce the **Slow Variational Autoencoder** (S-VAE) which extends the β -VAE formulation by a regularization term based on the slowness principle. To that end we extend the ELBO of the β -VAE with a *similarity loss term* (\mathcal{L}^{sim}), which enforces the similarity of latent representations for temporally close observations.

Let $D = (\tilde{o}_1, \dots, \tilde{o}_T)$ be a long sequence of unlabeled observations $\tilde{o} = (\mathbf{o}, t)$ that consist of an observation (\mathbf{o}) and the time index (t) that are used to determine the metric Δt that determines how far apart a pair of observations at times i, j are, where $i < j$. Specifically,

$$\Delta t(\tilde{\mathbf{o}}_i, \tilde{\mathbf{o}}_j) = j - i. \quad (1)$$

Let q_θ be the variational approximate posterior distribution obtained by an encoder network with parameters θ and \mathbf{z} be the latent vector such that $\mathbf{z} \sim q_\theta(\mathbf{z}|\mathbf{o})$. Considering two distinct yet sequential observations $\mathbf{o}_i, \mathbf{o}_j \in D \mid i < j$, the difference of the corresponding latent representations is given by the **approximate difference distribution**,

$$q_\theta(\mathbf{z}_j - \mathbf{z}_i | \mathbf{o}_j, \mathbf{o}_i) = \mathcal{N}(\mu_j - \mu_i, \Sigma_j + \Sigma_i) \equiv q_\theta(\Delta \mathbf{z} | \mathbf{o}_j, \mathbf{o}_i). \quad (2)$$

To express the decaying similarity with growing temporal separation Δt , we assume a prior that the latent vector exhibits Brownian motion. Denoting by Z_i, Z_j two increments of the Brownian motion at times i and j , respectively, the **prior** distribution $p(\Delta \mathbf{z})$ is given by

$$Z_j - Z_i = \sqrt{\Delta t} \cdot N \sim \mathcal{N}(0, \Delta t I) \equiv p(\Delta \mathbf{z}), \quad (3)$$

where $N \sim \mathcal{N}(0, I)$. This prior also encodes the time, resulting in observations further apart in time to have a prior distribution with a larger covariance.

The similarity of two observations can be computed as the Kullback-Leibler (KL) divergence between the approximate posterior and the prior distributions as

$$\mathcal{L}^{\text{sim}}(\tilde{\mathbf{o}}_i, \tilde{\mathbf{o}}_j) = D_{KL}(q_\theta(\Delta \mathbf{z} | \mathbf{o}_j, \mathbf{o}_i) || p(\Delta \mathbf{z})). \quad (4)$$

Δt can be considered a scaling factor for the variance of the Brownian motion. Thus, when considering two consequent elements in the sequence, we want to constrain \mathcal{L}^{sim} to be smaller than a certain bound. We consider only pairs of consecutive elements because the bound (scaled according to Δt) becomes weaker as the temporal distance between elements increases.

Combining the constraint term with the β -VAE, we can write the constrained optimization problem

$$\begin{aligned} \max_{\theta, \phi} \mathbb{E}_{(\tilde{\mathbf{o}}_i, \tilde{\mathbf{o}}_{i+1}) \sim D} [\mathbb{E}_{q_\theta(\mathbf{z}_i | \mathbf{o}_i)} [\log p_\phi(\mathbf{o}_i | \mathbf{z}_i)]] \text{ subject to } D_{KL}(q_\theta(\mathbf{z}_i | \mathbf{o}_i) || p(\mathbf{z})) < \epsilon_1, \\ D_{KL}(q_\theta(\Delta \mathbf{z} | \mathbf{o}_{i+1}, \mathbf{o}_i) || p(\Delta \mathbf{z})) < \epsilon_2 \end{aligned} \quad (5)$$

with a prior $p(\mathbf{z}) = \mathcal{N}(0, I)$ and the decoder network p_ϕ . The parameter ϵ_1 describes the strength of the latent bottleneck as in (Higgins et al., 2017) while ϵ_2 describes the strength of the temporal similarity constraint. Rewriting the above in Lagrange form we get,

$$\begin{aligned} \mathcal{F}(\theta, \phi, \beta, \lambda, \mathbf{o}_i, \mathbf{o}_{i+1}) = & \mathbb{E}_{q_\theta(\mathbf{z} | \mathbf{o}_i)} [\log p_\phi(\mathbf{o}_i | \mathbf{z})] \\ & - \beta (D_{KL}(q_\theta(\mathbf{z} | \mathbf{o}_i) || p(\mathbf{z})) - \epsilon_1) \\ & - \lambda (D_{KL}(q_\theta(\Delta \mathbf{z} | \mathbf{o}_{i+1}, \mathbf{o}_i) || p(\Delta \mathbf{z})) - \epsilon_2) \end{aligned} \quad (6)$$

Since $\epsilon_1, \epsilon_2 \geq 0$ we can rewrite the Lagrangian to arrive at the S-VAE loss function

$$\begin{aligned} \mathcal{F}(\theta, \phi, \beta, \lambda, \mathbf{o}_i, \mathbf{o}_{i+1}) \geq \mathcal{L}(\theta, \phi, \beta, \lambda, \mathbf{o}_i, \mathbf{o}_{i+1}) = & \mathbb{E}_{q_\theta(\mathbf{z}|\mathbf{o}_i)}[\log p_\phi(\mathbf{o}_i|\mathbf{z})] \\ & - \beta D_{KL}(q_\theta(\mathbf{z}|\mathbf{o}_i)||p(\mathbf{z})) \\ & - \lambda D_{KL}(q_\theta(\Delta\mathbf{z})|\mathbf{o}_{i+1}, \mathbf{o}_i||p(\Delta\mathbf{z})). \end{aligned} \quad (7)$$

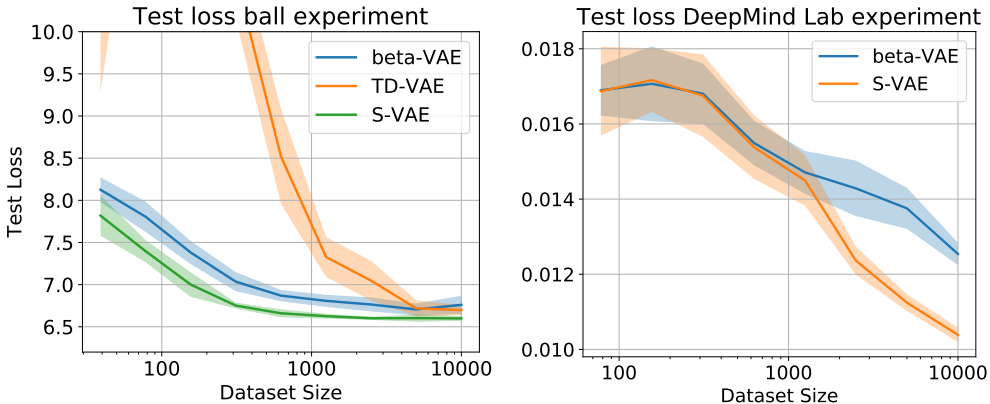
where β remains the same parameter as used in (Higgins et al., 2017). Our contribution, the additional similarity loss term scaled by the parameter λ consists of the KL-divergence between the approximate difference distribution of two observations and a random walk based prior. This formulation of the similarity loss allows us to ensure temporal consistency in both means and variances of the encoder as opposed to for example taking the $L1$ norm of the means μ_i and μ_j .

3 EXPERIMENTS

In this section, we will compare the performance of the following three methods: a β -VAE, the TD-VAE, and our method, the S-VAE. We used two datasets: a synthetic dataset consisting of a ball bouncing within the bounds of a 2D arena and a 3D dataset generated using the DeepMind Lab environment (Beattie et al., 2016).

3.1 SYNTHETIC DATASET

In the synthetic dataset experiment we compare the S-VAE, the TD-VAE and the β -VAE regarding the downstream few-shot learning performance when trying to learn the ball velocity from two consecutive frames. We generated sequences of 20 frames of a ball bouncing in a 100×100 uni-color 2D environment. For each sequence the ball is placed in a random position in the environment and initialized with a random direction and random but limited velocity vector. Then the environment performs 20 update steps and stores the frames and the ball’s x/y-velocity. Overall 10000 labeled sequences consisting of 20 datapoints each were generated.



(a) Test loss vs. dataset size Ball experiment.

(b) Test loss vs. dataset size DeepMind Lab experiment.

Figure 1: Average MSE loss with standard deviation on unseen test dataset over 12 runs with different random seeds compared to the labeled dataset size used during the few-shot learning.

During the **unsupervised representation learning step**, we use the full dataset without labels to train a S-VAE, a TD-VAE and a β -VAE model. Both S-VAE and β -VAE share the same architecture, a neural network with 4 fully connected layers with 300 hidden nodes that outputs means and logvariances of the 2D latent distributions. The β parameter of both VAE methods is 0.001 while λ is 0.001 for the S-VAE. The TD-VAE is trained using the same architecture as the moving MNIST experiment described in Appendix D in the paper (Gregor et al., 2019) with an 8D latent space.

In the supervised **few-shot down-stream task** we use subsets of $(1, 1/2, 1/4, \dots, 1/256)$ of the full labeled dataset with their labels to train the down-stream task of predicting the ball velocity. The down-stream task is trained by freezing the encoder networks trained in the previous step and feeding the latent representations into a fully connected neural network with 4 layers of 50 hidden nodes. The latent representations for the S-VAE and β -VAE are extracted from the latent bottleneck of the autoencoder structure. To obtain a representation from the TD-VAE we concatenate the sampled representation obtained from the belief states at t_1 and t_2 . Figure 1a shows the average performances (MSE between predicted and true labels) on the downstream task on an unseen test dataset across 12 runs. We can see that the S-VAE out-performs both competing methods for all subset sizes.

3.2 DEEPMIND LAB DATASET

In this experiment a dataset of 10000 sequences with 20 images each was generated using a random walker exploring a DeepMind Lab environment (Beattie et al., 2016). The training procedure is the same as for the bouncing ball experiment.

In the **unsupervised representation learning step** we trained a β -VAE and a S-VAE on the full dataset without labels. The network architectures of the S-VAE and β -VAE were adapted from the Ball experiment with 4 convolutional layers in both encoder and decoder as well as a $100D$ latent space. The β parameter for both methods was 0.00001 and λ was 10.0. The representations were extracted in the same way as for the Bouncing Ball experiment. We do not include the TD-VAE in the DeepMind Lab experiment as we were not able to recreate the results from (Gregor et al., 2019) with the more complex architecture and the given implementation instructions.

Table 1: Data efficiency improvement for both experiments. The table lists the amount of labeled data needed to achieve the best performance of the β -VAE, the second best performing method.

Method	Ball Exp.	DeepMind Lab Exp.
	$\beta = 0.001, \lambda = 0.001$	$\beta = 0.00001, \lambda = 10.0$
TD-VAE	– images	– images
β -VAE	5000 images	10000 images
S-VAE	446 images	2337 images
Improvement S-VAE:	91.08% less data	76.63% less data

Table 2: Best few-shot performance across subsets of the unlabeled training data in both experiments.

Method	Ball experiment			DeepMind Lab Exp.		
	$\beta = 0.001, \lambda = 0.001$			$\beta = 0.00001, \lambda = 10.0$		
	Full dataset	1/4 th	1/32	Full dataset	1/4 th	1/32
TD-VAE	6.70	7.05	10.66	–	–	–
β -VAE	6.76	6.76	7.03	0.0125	0.014	0.0168
S-VAE	6.56	6.60	6.75	0.0103	0.012	0.0167
Improvement S-VAE:	2.1%	2.4%	4.0%	17.6%	13.83%	0.60%

The **few-shot down-stream task** is then learned from subsets of $(1, 1/2, 1/4, \dots, 1/128)$ of the full dataset. The final performance on an unseen test dataset of each method for all subset sizes is shown in Fig. 1b. The S-VAE out-performs the β -VAE for medium and large subsets of labeled data.

3.3 DISCUSSION

In this section, we first discuss the results obtained from the experiments on both datasets. Then, we look at the structure of $2D$ slow latent spaces and how this influences the bias and variance of the downstream task across multiple runs.

3.3.1 DATA EFFICIENCY AND PERFORMANCE

To quantify the down-stream task data efficiency and performance of the S-VAE we compare the amount of data needed to achieve similar performance when compared to the best performing competing method.

From the results summarized in Table 1 we can see that the S-VAE requires 91.08% less training data for the bouncing ball experiment and 76.63% less in the DeepMind Lab experiment to achieve the best performance of the β -VAE of 6.71 and 0.0125 respectively. From Table 2 we can see that in the case of the Ball experiment, the S-VAE achieves a 2.1% better performance for the full labeled dataset when compared to the β -VAE and about 4% for $1/32$ and smaller subsets of the full dataset. While the TD-VAE performs similarly to the other methods on the full dataset, its performance degrades heavily for smaller portions of the labeled data. In the DeepMind Lab experiment the improvements over the β -VAE are 17.6% and 13.83% for the full labeled dataset and $1/4^{\text{th}}$ respectively. In this experiment it is also noteworthy that for very small subsets ($1/10^{\text{th}}$ of the dataset) both S-VAE and β -VAE perform similarly. This is due to the fact that the more complex down-stream training was prone to overfitting for small subsets and thus, we focus on analyzing the performance on larger subsets.

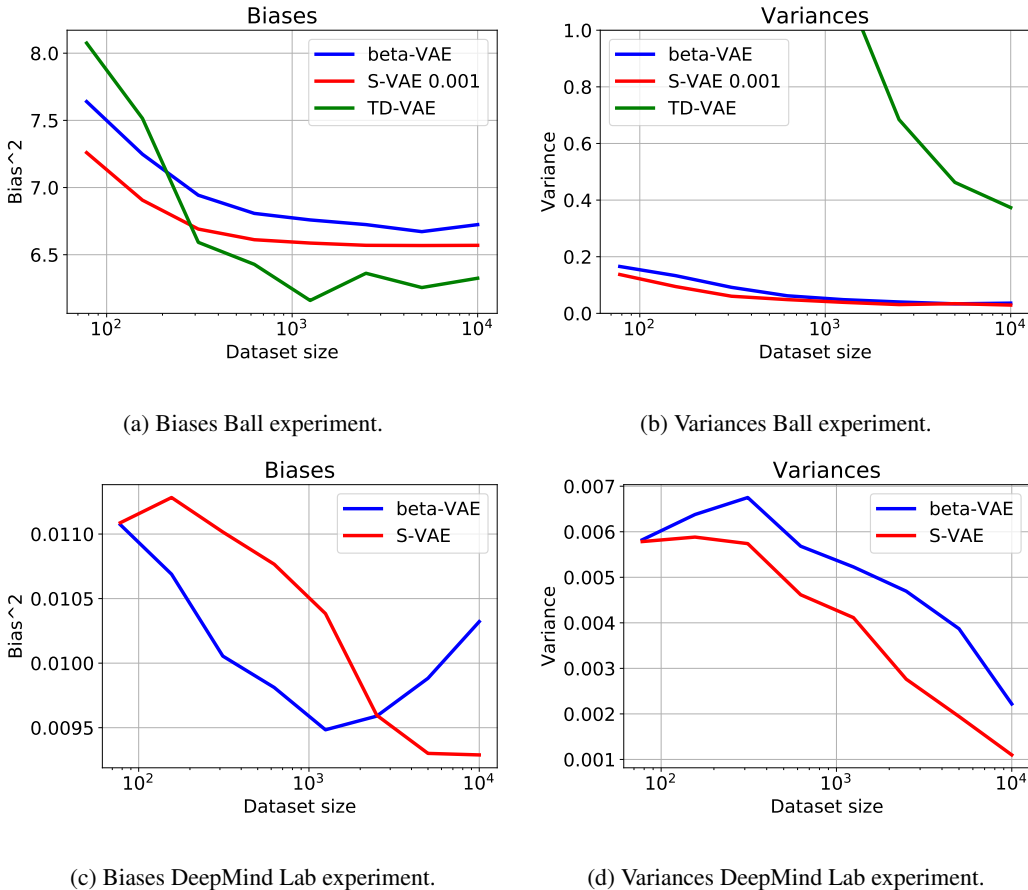


Figure 2: Bias variance decomposition for both experiments across all dataset sizes.

3.3.2 BIAS-VARIANCE DECOMPOSITION

To give perspective on the performance differences in down-stream task performance we decompose the bias and variances for each subset size across 12 runs with different random seeds, shown in Fig. 2. In the Ball experiment, the TD-VAE displays the lowest bias given more labeled data but also much larger variance and overall worse performance, confirming the initial observation of overfitting for small amounts of labeled data. The S-VAE displays both lower bias and variance compared to the β -VAE resulting in the improved performance and data efficiency observed before.

In the DeepMind Lab experiment the variance contributed more to the overall test error and was lower for the S-VAE through all subset sizes (Fig. 2d). The results on the bias are favoring the S-VAE for subsets with more than $1/5^{\text{th}}$ of the labeled data, for smaller subsets we assume the results are not reliable since the down-stream task training overfitted strongly.

Overall we conclude that the S-VAE displays lower bias for all considered subset sizes and lower variance across all subset sizes leading to less overfitting and overall better down-stream task performance.

3.3.3 LATENT SPACE STRUCTURE

Lastly we visualize the $2D$ latent space of the Ball experiment in Fig. 3 to demonstrate how the parameter λ influences its structure. Increasing the constraint on the similarity term in Eq. 7 increases the continuity of the latent space from β -VAE (Fig. 3a) where $\lambda = 0$ to S-VAE with $\lambda = 0.01$ and $\lambda = 1.0$ (Fig. 3b and 3c). Note that with increasing λ also the range of the latents decreases. We conclude that the similarity loss term structures the latent space according to the slowness principle, observations where the ball is in a similar position are also similar in latent space.

The TD-VAE has been visualized in two dimensions by applying PCA to the $8D$ latent space which exhibits high discontinuity. This is likely a result of the TD-VAE being trained for a specific task and not enforcing a Gaussian prior on the latent space. As we can see from Fig. 2b, this leads to large variances, lower bias and ultimately to extreme overfitting for small datasets.

Taking the results of the bias variance decomposition into account, we conclude that a more continuous latent space leads to lower bias and variance compared to the β -VAE, and at the same time the discontinuous latent space of the TD-VAE displays strong overfitting especially for smaller subsets. Ultimately this results in better data efficiency and better overall performance. These results indicate that continuous latent spaces might be beneficial for the optimization process of down-stream tasks which we would like to expand on in future work.

4 RELATED WORK

Unsupervised learning of invariant features from observations is an efficient way to extract higher level features about a scene without the need for human labels. However, it is still not clear what invariances lead to the most descriptive features and which features work best for down-stream tasks (Saunshi et al., 2019; Locatello et al., 2019; Bengio et al., 2013).

4.1 SLOWNESS PRINCIPLE

Research from neuroscience (Berkes & Wiskott, 2005) suggests that cell structures in the visual cortex emerge based on the underlying principle of extracting slowly varying features from the environment. This principle termed *slowness principle* follows the assumption that the underlying features of an environment and the internal representations vary on a different time scale than the sensory signals. More intuitively, we would like to extract invariant scene information which changes slowly over time, e.g. a car passing by, from a video whose individual pixel values change quickly assuming that the slowly changing factors are good higher level representations of the observations.

The most well known application of the slowness principle is the slow feature analysis method (SFA) introduced in (Wiskott & Sejnowski, 2002). SFA is an unsupervised learning algorithm capable to extract linearly decorrelated features by expanding and transforming the input signal such that it can be optimized for finding the most slowly varying features from an input signal (Wiskott & Sejnowski, 2002). Extending the SFA method to nonlinear features has shown that the learned

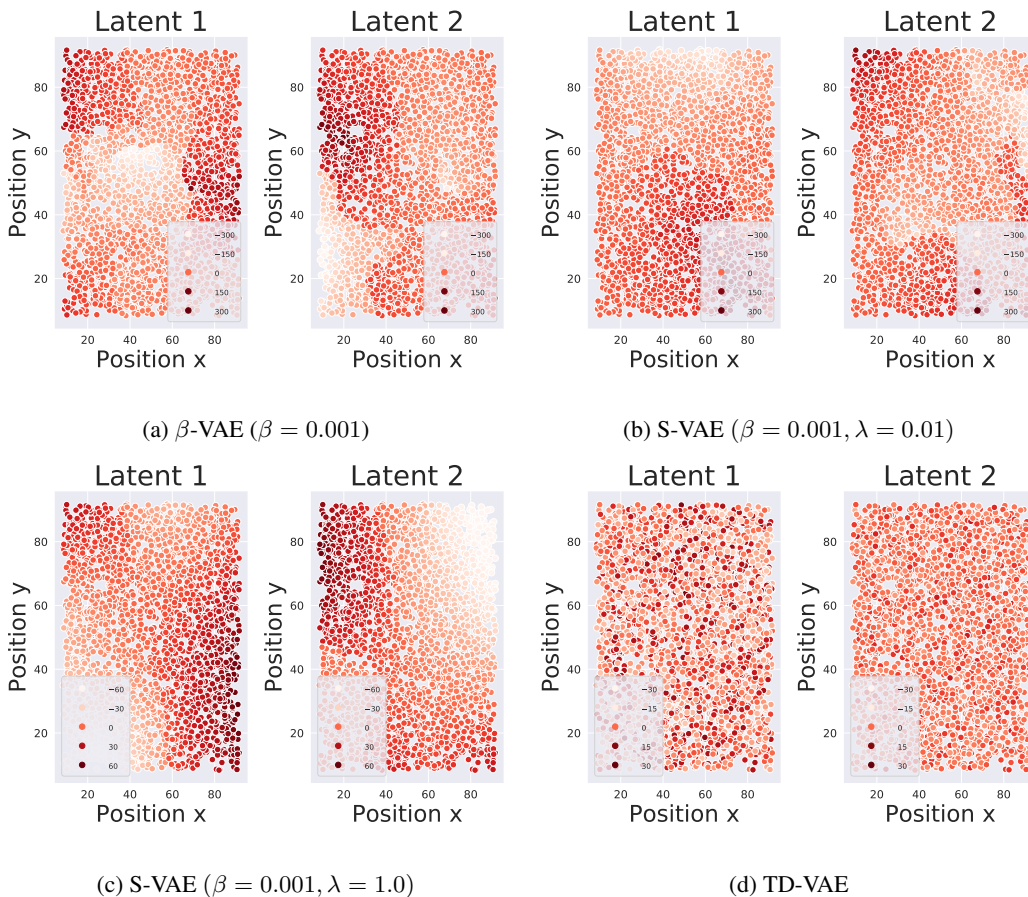


Figure 3: Scatter plot visualization of 3000 samples in the 2D latent space of the Ball experiment. X- and Y-Axis are the position of the ball in the 100×100 pixel large environment and the color scale represents the value of the corresponding latent dimension.

features share many characteristics with those of complex cells in the V1 cortex (Berkes & Wiskott, 2005). Further applications of the slowness principle include object detection invariant of spatial transformations (Franzius et al., 2011), pre-training neural networks for improved performance on the MNIST dataset (Bengio & Bergstra, 2009) and the self organization of grid cells, structures in the rodent brain used for navigation (Franzius et al., 2007b;a).

4.2 CONTRASTIVE LEARNING AND THE SLOWNESS PRINCIPLE

The equivalent in state of the art machine learning that could be considered related is contrastive learning. The objective of contrastive learning is to encode observations and place them in a latent space using a metric score that allows to express (dis-)similarity of the observations. Contrastive learning has been successfully applied in reinforcement learning (Laskin et al., 2020) and most recently for object classification in SimCLR (Chen et al., 2020a;b). These methods use a contrastive loss on augmented versions of the same observation, effectively learning transformation invariant features from images, and show that these representations benefit reinforcement learning and image classification tasks.

When using the time as the contrastive metric, similar to the slowness principle, we speak about *time-contrastive* learning. Time-contrastive learning has been applied successfully to learning view-point-invariant representations for learning from demonstration with a robot (Sermanet et al., 2018). Related to our work, Mobahi et al. (Mobahi et al., 2009) used the coherence in video material to train a CNN for a variety of specific tasks. While training two CNNs in parallel with shared parameters,

in alternating fashion a labeled pair of images was used to perform a gradient update minimizing training loss followed by selecting two unlabeled images from a large video dataset to minimize a time-contrastive loss based on the $L1$ norm of the representations at each individual layer. The experiments showed that supervised tasks can benefit from the additional pseudo-supervisory signal and that features invariant to pose, illumination or clutter can be learned.

Compared to the work by Mobahi et al. (2009), our method is focused on learning task-agnostic representations that encode uncertainty and facilitate data-efficient learning of downstream tasks. These goals are achieved by extending the state of the art β -VAE by an additional similarity loss term based on the Kullback-Leibler divergence.

As a comparison in this paper we use the β -VAE (Higgins et al., 2017; Kingma & Welling, 2013) and the temporal difference variational autoencoder (TD-VAE) (Gregor et al., 2019). We chose the TD-VAE as it learns representations that include temporal abstraction capabilities, encode an uncertain belief state and is not based on the variational autoencoder. The TD-VAE is trained on sequences from a video trying to predict a time step in the future from information that is encoded in a belief code at the current step.

5 CONCLUSION

In this paper, we introduced the Slow Variational Autoencoder (S-VAE) which applies the slowness principle to the state-of-the-art β -VAE by enforcing similarity in latent space based on temporal similarity in an observation sequence. To this end we derived a similarity loss term constrained by a parameter λ and added it to the ELBO of the β -VAE. We show empirically that unsupervised pre-training using time correlated data with this new loss term leads to improved downstream task performance and data efficiency. Qualitative analysis of the latent space structure and the bias-variance decomposition of the downstream task shows that the similarity loss enforces the slowness property on the latent space. This leads to more continuous latent spaces which facilitate more data efficient learning of downstream tasks. In future works, we would like to investigate how the structure of the latent space influences the optimization of downstream tasks and if these improvements are applicable to a wider variety of tasks.

REFERENCES

- Charles Beattie, Joel Z Leibo, Denis Teplyashin, Tom Ward, Marcus Wainwright, Heinrich Küttler, Andrew LeFrancq, Simon Green, Víctor Valdés, Amir Sadik, et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 2016.
- Yoshua Bengio and James S Bergstra. Slow, decorrelated features for pretraining complex cell-like networks. In *Advances in neural information processing systems*, pp. 99–107, 2009.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- Pietro Berkes and Laurenz Wiskott. Slow feature analysis yields a rich repertoire of complex cell properties. *Journal of vision*, 5(6):9–9, 2005.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020a.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020b.
- Mathias Franzius, Henning Sprekeler, and Laurenz Wiskott. Slowness and sparseness lead to place, head-direction, and spatial-view cells. *PLoS Comput Biol*, 3(8):e166, 2007a.
- Mathias Franzius, Roland Vollgraf, and Laurenz Wiskott. From grids to places. *Journal of computational neuroscience*, 22(3):297–299, 2007b.
- Mathias Franzius, Niko Wilbert, and Laurenz Wiskott. Invariant object recognition and pose estimation with slow feature analysis. *Neural computation*, 23(9):2289–2323, 2011.

- Karol Gregor, George Papamakarios, Frederic Besse, Lars Buesing, and Theophane Weber. Temporal difference variational auto-encoder. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=S1x4ghC9tQ>.
- 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. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=Sy2fzU9gl>.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Michael Laskin, Aravind Srinivas, and Pieter Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. In *Proceedings of the 37th Annual International Conference on Machine Learning (ICML)*, 2020.
- Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, pp. 4114–4124, 2019.
- Hossein Mobahi, Ronan Collobert, and Jason Weston. Deep learning from temporal coherence in video. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 737–744, 2009.
- Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*, pp. 5628–5637, 2019.
- Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain. Time-contrastive networks: Self-supervised learning from video. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1134–1141. IEEE, 2018.
- Laurenz Wiskott and Terrence J Sejnowski. Slow feature analysis: Unsupervised learning of invariances. *Neural computation*, 14(4):715–770, 2002.
- Will Y Zou, Andrew Y Ng, and Kai Yu. Unsupervised learning of visual invariance with temporal coherence. In *NIPS 2011 workshop on deep learning and unsupervised feature learning*, volume 3, 2011.