

Stochastic Bottleneck: Rateless Auto-Encoder for Flexible Dimensionality Reduction

Koike-Akino, Toshiaki; Wang, Ye

TR2020-075 June 24, 2020

Abstract

We propose a new concept of rateless auto-encoders (RL-AEs) that enable a flexible latent dimensionality, which can be seamlessly adjusted for varying distortion. In the proposed RL-AEs, instead of a deterministic bottleneck architecture, we use an over-complete representation that is stochastically regularized with weighted dropouts. Our RL-AEs employ monotonically increasing dropout rates across the latent representation nodes such that the latent variables become sorted by importance like in principal component analysis (PCA). This is motivated by the rateless property of conventional PCA, where the least important principal components can be discarded to realize variable rate dimensionality reduction that gracefully degrades the distortion. Our proposed stochastic bottleneck framework enables seamless rate adaptation with high reconstruction performance, without requiring predetermined latent dimensionality at training. We experimentally demonstrate that the proposed RL-AEs can achieve variable dimensionality reduction while retaining nearly optimal distortion compared to conventional AEs.

IEEE International Symposium on Information Theory (ISIT) 2020

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Stochastic Bottleneck: Rateless Auto-Encoder for Flexible Dimensionality Reduction

Toshiaki Koike-Akino and Ye Wang

Mitsubishi Electric Research Laboratories (MERL), 201 Broadway, Cambridge, MA 02139, USA.

Email: {koike, yewang}@merl.com

Abstract—We propose a new concept of rateless auto-encoders (RL-AEs) that enable a flexible latent dimensionality, which can be seamlessly adjusted for varying distortion. In the proposed RL-AEs, instead of a deterministic bottleneck architecture, we use an over-complete representation that is stochastically regularized with weighted dropouts. Our RL-AEs employ monotonically increasing dropout rates across the latent representation nodes such that the latent variables become sorted by importance like in principal component analysis (PCA). This is motivated by the rateless property of conventional PCA, where the least important principal components can be discarded to realize variable rate dimensionality reduction that gracefully degrades the distortion. Our proposed stochastic bottleneck framework enables seamless rate adaptation with high reconstruction performance, without requiring predetermined latent dimensionality at training. We experimentally demonstrate that the proposed RL-AEs can achieve variable dimensionality reduction while retaining nearly optimal distortion compared to conventional AEs.

I. INTRODUCTION

In many real-world applications, the raw data measurements (e.g., audio/speech, images, video, biological signals) often have very high dimensionality. Adequately handling high-dimensionality often requires the application of dimensionality reduction techniques [1] that transform the original data into meaningful feature representations of reduced dimensionality. Such feature representations should reduce the dimensionality to the minimum number required to capture the salient properties of the data. Dimensionality reduction is vital in many machine learning applications, since one needs to mitigate the so-called “curse of dimensionality” [2]. In the past few decades, latent representation learning based on auto-encoders (AEs) [3]–[10] has been widely used for dimensionality reduction, since this nonlinear technique has shown superior real-world performance compared to classical linear counterparts, such as principal component analysis (PCA).

One of the challenges in dimensionality reduction is to determine the optimal latent dimensionality that can sufficiently capture the data features required for particular applications. Although some regularization techniques, such as sparse AE (SAE) [7] and rate-distortion AE [11], may be useful to self-adjust the effective dimensionality, there are no existing methods that provide a *rateless* property [26] that allows for seamless adjustment of the latent dimensionality depending on varying distortion requirements for different downstream applications, without modification of the trained AE model. However, realizing a rateless AE is not straightforward, since traditional AEs typically learn nonlinear manifolds where

the latent variables are equally important, unlike the linear manifold models used for PCA.

In this paper, we introduce a novel AE framework which can universally achieve flexible dimensionality reduction while retaining excellent distortion performance. Motivated by the fact that the traditional PCA is readily adaptable to any dimension by just appending or dropping sorted principal components, we propose a stochastic bottleneck architecture to associate upper latent variables with higher-principal nonlinear features so that the user can freely discard the least-principal latent variables if desired. Our contributions are summarized:

- We introduce a new concept of rateless AEs designed for flexible dimensionality reduction.
- A stochastic bottleneck framework is proposed to prioritize the latent space non-uniformly.
- An extended regularization technique called TailDrop is considered to realize rateless AEs.
- We discuss dropout distribution optimization under the principle of multi-objective learning.
- We demonstrate that the proposed AEs achieve excellent distortion performance over the variable range of dimensionality in the standard MNIST and CIFAR-10 image datasets.
- We evaluate AE models trained for a perceptual distortion measure based on structural similarity (SSIM) [25] as well as the traditional mean-square error (MSE) metric.

II. RATELESS AUTO-ENCODER (RL-AE)

A. Dimensionality reduction

Due to the curse of dimensionality, representation learning to reduce the dimensionality is often of great importance to handle high-dimensional datasets in machine learning. To date, there have existed many algorithms for dimensionality reduction [1], e.g., PCA, Isomap, diffusion maps, locally linear embedding, Laplacian eigenmaps, Sammon mapping, and manifold charting along with AE. Among all, AE [3]–[10] has shown its high potential to learn lower-dimensional latent variables required in the nonlinear manifold underlying the datasets. AE is an artificial neural network having a bottleneck architecture as illustrated in Fig. 1(a), where N -dimensional data is transformed to M -dimensional latent representation (for $M \leq N$) via an encoder network. The latent variables should contain sufficient feature capable of reconstructing the original data through a decoder network.

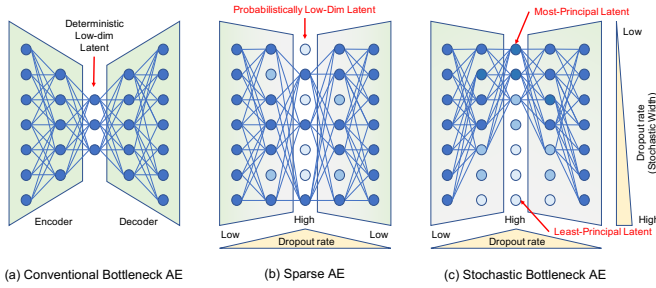


Fig. 1: Auto-encoder architectures: (a) conventional bottleneck network, (b) sparse AE regularized by dropout, having probabilistically lower-dimension representation, (c) stochastic bottleneck, having two-dimensional regularization with non-identical dropout rates in both depth and width directions to realize rateless property by ordered-principal latent variables.

From the original data $\mathbf{x} \in \mathbb{R}^N$, the corresponding latent representation $\mathbf{z} \in \mathbb{R}^M$, with a reduced dimensionality M is generated by the encoder network as $\mathbf{z} = f_\theta(\mathbf{x})$, where θ denotes the encoder network parameters. The latent variables should adequately capture the statistical geometry of the data manifold, such that the decoder network can reconstruct the data as $\mathbf{x}' = g_\phi(\mathbf{z})$, where ϕ denotes the decoder network parameters and $\mathbf{x}' \in \mathbb{R}^N$. The encoder and decoder pair (f_θ, g_ϕ) are jointly trained to minimize the reconstruction loss (i.e., distortion), as given by:

$$\min_{\theta, \phi} \mathbb{E}_{\mathbf{x} \sim \text{Pr}(\mathbf{x})} [\mathcal{L}(\mathbf{x}, g_\phi(f_\theta(\mathbf{x})))], \quad (1)$$

where the loss function $\mathcal{L}(\mathbf{x}, \mathbf{x}')$ is chosen to quantify the distortion (e.g., MSE) between \mathbf{x} and \mathbf{x}' .

B. Motivation: rateless property

AEs are also known as nonlinear PCA (NLPCA) [4]–[6]. If we consider a simplified case where there is no nonlinear activation in the AE model, then the encoder and decoder functions will reduce to simple affine transformations. Specifically, the encoder becomes $f_\theta(\mathbf{x}) = \mathbf{W}\mathbf{x} + \mathbf{b}$ where trainable parameters θ are the linear weight $\mathbf{W} \in \mathbb{R}^{M \times N}$ and the bias $\mathbf{b} \in \mathbb{R}^M$. Likewise, the decoder becomes $g_\phi(\mathbf{z}) = \mathbf{W}'\mathbf{z} + \mathbf{b}'$ with $\phi = \{\mathbf{W}', \mathbf{b}'\} \in \{\mathbb{R}^{N \times M}, \mathbb{R}^N\}$. If the distortion measure is MSE, the optimal linear AE coincides with the classical PCA when the data follows the multivariate Gaussian distribution according to the Karhunen–Loève theorem.

To illustrate, assume Gaussian data $\mathbf{x} \sim \mathcal{N}(\mathbf{m}, \mathbf{C})$ with mean $\mathbf{m} \in \mathbb{R}^N$ and covariance $\mathbf{C} \in \mathbb{R}^{N \times N}$, which has the eigen-decomposition: $\mathbf{C} = \mathbf{\Phi}\mathbf{\Lambda}\mathbf{\Phi}^T$, where $\mathbf{\Phi} \in \mathbb{R}^{N \times N}$ is the unitary eigenvectors matrix and $\mathbf{\Lambda} = \text{diag}[\lambda_1, \lambda_2, \dots, \lambda_N] \in \mathbb{R}^{N \times N}$ is the diagonal matrix of ordered eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N \geq 0$. For PCA, the encoder uses M principal eigenvectors $\mathbf{\Phi}\mathbf{I}_{N,M}$ to project the data onto an M -dimensional latent subspace with $\mathbf{W} = \mathbf{I}_{M,N}\mathbf{\Phi}^T$ and $\mathbf{b} = -\mathbf{W}\mathbf{m}$, where $\mathbf{I}_{M,N} \in \mathbb{R}^{M \times N}$ denotes the incomplete identity matrix with diagonal elements equal to one and zero

elsewhere. The decoder uses the transposed projection with $\mathbf{W}' = \mathbf{\Phi}\mathbf{I}_{N,M}$ and $\mathbf{b}' = \mathbf{m}$. The MSE distortion is given by

$$\bar{\mathcal{L}}_M = \mathbb{E}_{\mathbf{x}} [\|\mathbf{W}'(\mathbf{W}\mathbf{x} + \mathbf{b}) + \mathbf{b}' - \mathbf{x}\|^2] = \sum_{n=M+1}^N \lambda_n. \quad (2)$$

Since the eigenvalues are sorted, the distortion gracefully degrades as principal components are removed in the corresponding order. Of course, the MSE would be considerably worse if an improper ordering (e.g., reversed) is used.

One of the benefits of classical PCA is its graceful *rateless* property due to the ordering of principal components. Similar to rateless channel coding such as fountain codes [26], PCA does not require a pre-determined compression ratio M/N for dimensionality reduction, and the latent dimensionality can be later freely adjusted depending on the downstream application. More specifically, the PCA encoder and decoder learned for a dimensionality of M can be universally used for any lower-dimensional PCA of latent size $L \leq M$ without any modification of the PCA model but simply dropping the least-principal D components ($D = M - L$) in $\mathbf{z} = [z_1, z_2, \dots, z_M]^T$, i.e., nullifying variables as $z_m = 0$ for all $m \in \{L+1, \dots, M\}$.

The rateless property is greatly beneficial in practical applications since the optimal latent dimensionality is often not known beforehand. Instead of training multiple encoder and decoder pairs for different compression rates, one common PCA model can cover all rates L/N for $1 \leq L \leq M$ by simply dropping trailing components, while still attaining the optimal distortion as given by $\bar{\mathcal{L}}_L$. Nevertheless, traditional PCA often underperforms in comparison to nonlinear dimensionality reduction techniques on real-world datasets. Exploiting nonlinear activation functions such as rectified linear unit (ReLU), AEs can better learn inherent nonlinearities of the latent representations underlying the data. However, existing AEs do not readily achieve the rateless property, because the latent variables are generally learned to be equally important. This drawback still holds for the progressive dimensionality reduction approaches employed by stacked AEs [3] and hierarchical AEs [4], those of which require multiple training and re-tuning for different dimensionality. In this paper, we propose an effective technique employing a stochastic bottleneck to realize rateless AEs that are adaptable to any rates.

C. StochasticWidth bottleneck

Several variants of AE have been proposed, e.g., sparse AE (SAE) [7], variational AE (VAE) [8]–[10], rate-distortion AE [11], and compressive AE [12]. We introduce a new AE family which has no fixed bottleneck architecture to realize the rateless property for seamless dimensionality reduction. Our method can be viewed as an extended version of SAE, similar in its over-complete architecture, but also employing a varying dropout distribution across the width of the network. This aspect is a key for achieving excellent distortion while allowing a flexible dimensionality reduction.

Unlike a conventional AE with a deterministic bottleneck architecture, as shown in Fig. 1(a), the SAE employs a

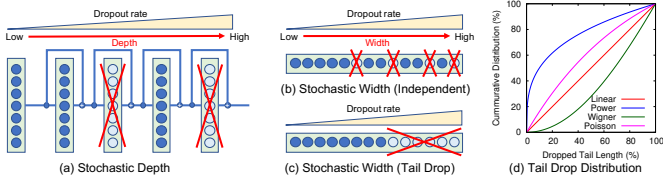


Fig. 2: Non-uniform dropout regularization: (a) StochasticDepth [17] to control depth by prioritizing shallower layers, (b) StochasticWidth to control width by prioritizing head neurons with independent and increasing dropout, (c) StochasticWidth to nullify consecutive burst neurons by TailDrop, (d) example of tail drop distributions.

probabilistic bottleneck with an effective dimensionality that is stochastically reduced by dropout, as depicted in Fig. 1(b). For example, the SAE encoder generates M -dimensional variables \mathbf{z} which are randomly dropped out at a probability of p , resulting in an effective latent dimensionality of $\bar{L} = (1-p)M$. Although the SAE has better adaptability than deterministic AE to further dimensionality reduction by dropping latent variables, the latent variables are still trained to be equally important for reconstruction of the data, and thus it is limited in achieving flexible ratelessness.

Our AE employs a stochastic bottleneck that imposes a specific dropout rate distribution that varies across both the width and depth of the network, as shown in Fig. 1(c). In particular, our StochasticWidth technique employs a monotonically increasing dropout rate from the head (upper) latent variable nodes to the tail (lower) nodes in order to encourage the latent variables to be ordered by importance, in a manner analogous to PCA. By concentrating more important features in the head nodes, we hope to enable adequate data reconstruction even when some of the least important dimensions (analogous to least-principal components) are later discarded.

D. TailDrop regularization

Dropout [13], [14] has been widely used to regularize over-parameterized deep neural networks. The role of dropout is to improve generalization performance by preventing activations from becoming strongly correlated, which in turn leads to over-training. A recent theory [24] provides a viable interpretation of dropout as a Bayesian inference approximation.

There are many related regularization methods proposed in literature; e.g., [15]–[23]. In order to facilitate the rateless property for stochastic bottleneck AE architectures, we introduce an additional regularization mechanism referred to as TailDrop, as one realization of StochasticWidth.

The stochastic bottleneck uses non-uniform dropout to adjust the importance of each neuron as explained in Fig. 1(c). This regularization technique is related to StochasticDepth [17] used in deep residual networks. As illustrated in Fig. 2(a), StochasticDepth drops out entire layers at a higher chance when dropping deeper layers so that an effective network depth is constrained and shallower layers are dominantly trained. Analogously, non-uniform dropouts are carried out

across the width direction for StochasticWidth as shown in Fig. 2(b), where independent dropouts at increasing rates are used for each neuron. The monotonically increasing dropout rates can be also realized by dropping consecutive nodes at the tail as shown in Fig. 2(c), which we call TailDrop. For TailDrop, the desired dropout rates can be achieved by adjusting the probability distribution of the tail drop length as depicted in Fig. 2(d). We focus on the use of this TailDrop regularization for rateless AE in this paper.

E. Multi-objective learning

Finding an appropriate dropout probability distribution is a key consideration in the design of high-performance rateless AEs. We now offer insights on how to do so. The objective function in (1) should be re-formulated to realize the rateless property. Our ultimate goal is to find AE model parameters θ and ϕ that simultaneously minimize distortion across multiple rates. Specifically, this problem is an M -ary multi-objective optimization as follows:

$$\min_{\theta, \phi} \left[\bar{\mathcal{L}}(\theta, \phi; 1), \bar{\mathcal{L}}(\theta, \phi; 2), \dots, \bar{\mathcal{L}}(\theta, \phi; M) \right], \quad (3)$$

where $\bar{\mathcal{L}}(\theta, \phi; L)$ denotes the expected distortion for the candidate AE model parameterized by θ and ϕ , given that the M -dimensional latent variables \mathbf{z} are further reduced to L -dimensional variables by dropping the last $D = M - L$ variables. In this multi-objective problem, a rateless AE model must account for the best balance across multiple dimensionalities in order to approach the Pareto-front solutions.

One commonly used naïve method in multi-objective optimization is a weighted sum optimization to reduce the problem to a single objective function as follows:

$$\min_{\theta, \phi} \sum_{L=1}^M \omega_L \bar{\mathcal{L}}(\theta, \phi; L), \quad (4)$$

with some weights $\omega_L \geq 0$. One may choose the weights to scale the distortion to a similar amplitude as $\omega_L \simeq 1/\bar{\mathcal{L}}^*(\theta, \phi; L)$ for positive distortions where $\bar{\mathcal{L}}^*(\theta, \phi; L)$ denotes the ground solution. As the expected distortion may depend on the eigenvalues as shown in (2), understanding the nonlinear eigenspectrum can facilitate in optimizing the weight distributions. The stochastic TailDrop regularization at training phase can be interpreted as a weight ω_L since the conventional single-objective optimization in (1) will effectively become the weighted sum optimization in (4). Accordingly, the weights will be the survivor length probability, i.e., the TailDrop distribution is $\Pr(D = M - L) = \omega_L$.

In this paper, we consider parametric eigenspectrum assumptions for simplicity. Under a model-based approach of nonlinear eigenspectrum assumptions, we evaluated several parametric distributions for TailDrop probability, e.g., Laplacian, and Wigner distributions, some of which are depicted in Fig. 2(d). Through a preliminary experiment, it was found that the power cumulative distribution function $\Pr(D < \tau M) = \tau^\beta$ for an order of $\beta \simeq 1$ (τ denotes a compression rate) performed well for most cases. Accordingly, we focus on the power distribution for TailDrop in the following experiments.

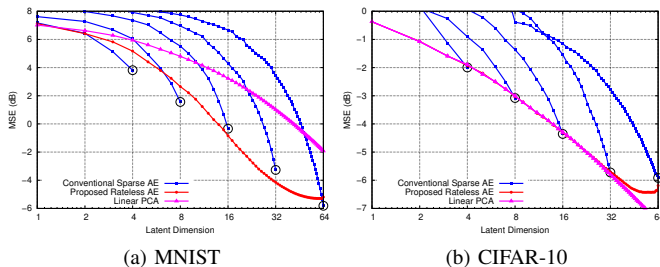


Fig. 3: MSE performance of SAE and RL-AE as a function of survivor latent dimensionality L .

III. EXPERIMENTS

To demonstrate the principle-of-concept benefits of our rateless AEs, we use standard image datasets of MNIST and CIFAR-10 [27]. MNIST contains handwritten 10-class gray-scale images of size 28-by-28, and thus the raw data dimensionality is $N = 28^2 = 784$. The dataset has 60,000 and 10,000 images for training and testing, respectively. CIFAR-10 is a dataset of 32-by-32 color images, representing 10 classes of natural scene objects. The raw data dimensionality is thus $N = 32^2 \times 3 = 3072$. The training set and test set contain 50,000 and 10,000 images, respectively.

For simplicity, we use fully-connected three-layer perceptron with ReLU activation functions for both encoder and decoder networks. Note that the concept of StochasticWidth regularization to realize ordered-principal feature can be applied to recurrent and convolutional networks in a straightforward manner. The number of nodes in the hidden layers is 1024 for MNIST and 2048 for CIFAR-10. For conventional SAE, we used 90% sparsity as a baseline to evaluate the robustness of flexible latent dimensionality. Model training was performed using the adaptive momentum (Adam) stochastic gradient descent method [28] with a learning rate of 0.001, and a mini-batch size of 100. The maximum number of epochs is 500 while early stopping with a patience of 20 was applied.

A. MSE measure

Figs. 3(a) and (b) show the MSE performance of the conventional SAE and proposed RL-AE for MNIST and CIFAR-10 datasets, respectively. For conventional SAE, multiple AE models are trained at the intended latent dimensionality of $M = 2^m$ for $m \in \{2, 3, 4, 5, 6\}$. The rateless AE is optimized at the dimensionality of $M = 64$ using TailDrop with a power distribution (with $\beta = 0.67$ for MNIST, and $\beta = 2.1$ for CIFAR-10). The latent dimensionality L used for image reconstruction is varied during testing evaluation by deterministically dropping tail variables.

As shown in Fig. 3(a), the conventional AE does not adapt well to variable dimensionality, with the MSE performance drastically degrading when the testing dimensionality L is reduced from the intended dimensionality M . For the SAE model trained for $M = 64$, dropping 50% of the latent variables to yield a reduced dimensionality of $L = 32$,

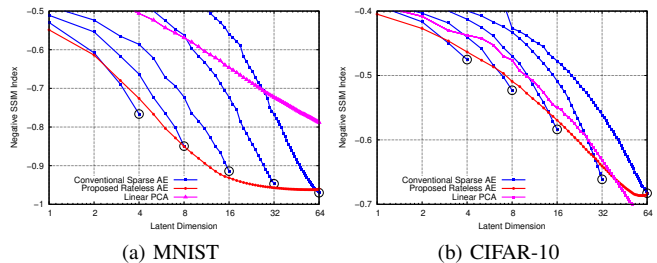


Fig. 4: SSIM performance of SAE and RL-AE as a function of survivor latent dimensionality L .

the MSE degrades to 3.5 dB from the -6 dB obtained at $L = 64$, which is significantly worse than an SAE model trained for $M = 32$ that obtains an MSE of -3.5 dB. This shows that the existing SAEs cannot be universally reused for flexible dimensionality reduction, and hence adaptive switching between multiple trained SAE models would be required depending on the desired dimensionality. However, our proposed RL-AE, which is trained once for dimensionality $M = 64$, flexibly operates over the wide range of further reduced dimensionalities $L \leq 64$, while achieving low MSE distortion close to the ideal MSEs obtained by SAE models trained for the specific dimensionality L .

Similar observations can be made in the results for the CIFAR-10 dataset, as shown in Fig. 3(b). It confirms that the nearly optimal distortion can be achieved by a single AE model for different compression rates by using the stochastic bottleneck regularization. This benefit comes from non-uniform dropout rates across neurons to concentrate the most-principal feature in upper nodes. Conventional uniform-rate dropout, as used in existing SAEs, still requires the target dimensionality to be known during training.

It should be noted that the linear PCA dimensionality reduction performs surprisingly well, competitive to the proposed nonlinear AE for CIFAR-10 datasets in Fig. 3(b). Because MNIST images are nearly binary bitmaps whose statistics are far from the Gaussian distribution, PCA did not work well as shown in Fig. 3(a). However, most natural images such as CIFAR-10 are often well-modeled by the Gauss–Markov process. This may be the primary reason why PCA works sufficiently well in particular for the MSE metric. Although it was unexpected that the nonlinear AEs could not improve the MSE performance over the linear PCA for CIFAR-10 datasets, the MSE curve of our AE perfectly agreed that of PCA for $L \leq 32$, which implies that our stochastic bottleneck approach could learn the ordered-principal components as intended.

B. SSIM measure

Since the classical MSE metric is known to be inconsistent with perceptual image quality, the structural similarity (SSIM) index [25] has been recently used as an alternative measure of perceptual distortion. The SSIM index ranges from -1 to 1 , indicating perceptual similarity between the original and

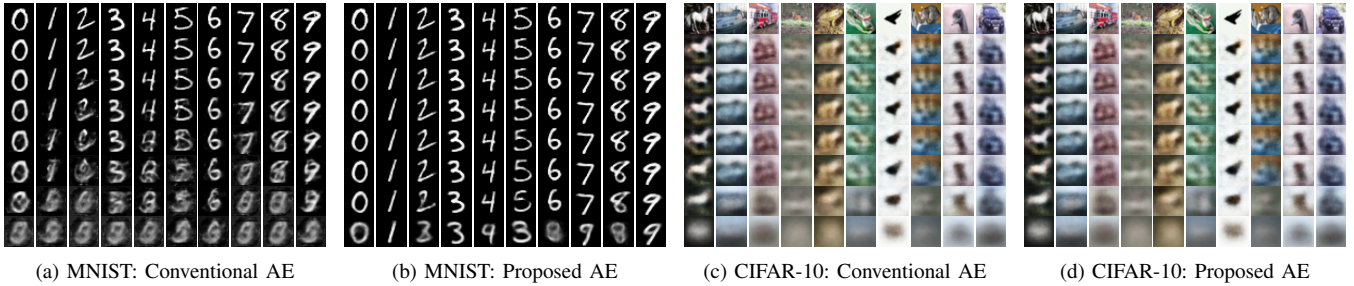


Fig. 5: Reconstructed snapshots varying the latent dimensionality L using AE model designed at dimensionality of $M = 64$. The top row is original, and subsequent rows are reconstructed images for a reduced dimensionality of $L = \{64, 54, 44, 34, 24, 14, 4\}$.

distorted images, from the worst to best quality, respectively. We use a negative SSIM index as a new loss function to fine-tune the AE models, which were pre-trained for the MSE metric, so as to improve the perceptual image quality.

Figs. 4(a) and (b) plot the negative SSIM index of the reconstructed images by the conventional SAE and proposed RL-AE for MNIST and CIFAR-10 datasets, respectively. It is confirmed in those figures that the conventional SAE cannot be universally used for flexibly varying dimensionality in the SSIM distortion metric. Although the proposed RL-AE may perform worse than the conventional SAEs at some dimensionalities, for which the SAE models were dedicatedly optimized, our RL-AE flexibly achieves SSIM performance closely comparable to the best SSIMs obtained by the ensemble of SAEs over the wide range of dimensionalities $L \leq 64$.

We can also see that the traditional PCA has a higher loss in the perceptual SSIM metric compared to the MSE metric. In particular for MNIST in Fig. 4(a), the SSIM degradation of the PCA over our RL-AE is noticeable over the whole range of dimensionalities, while the PCA worked well for lower dimensionality for the MSE metric, as seen in Fig. 3(a). More importantly, our AE can offer a perceptual performance benefit in the SSIM metric over PCA even for CIFAR-10 datasets, for which the AEs could not outperform the PCA in the MSE metric as discussed in Fig. 3(b). This makes sense because the PCA does not consider any perceptual quality.

C. Reconstructed images

Figs. 5 show visual samples randomly chosen from test datasets, respectively for SAE and RL-AE reconstructions. The top row displays the original images, and the subsequent rows are reconstructed images for a reduced dimensionality of $L = \{64, 54, 44, 34, 24, 14, 4\}$. Both types of models are trained at a latent dimensionality of $M = 64$ under the MSE measure. Our proposed RL-AE clearly exhibits improved visual quality for flexible dimensionality reduction versus the conventional SAE, without requiring retraining for each reduced dimensionality.

D. Latent representation

Finally we show a latent space geometry in Figs. 6(a) and (b), where the first two latent variables of all MNIST test

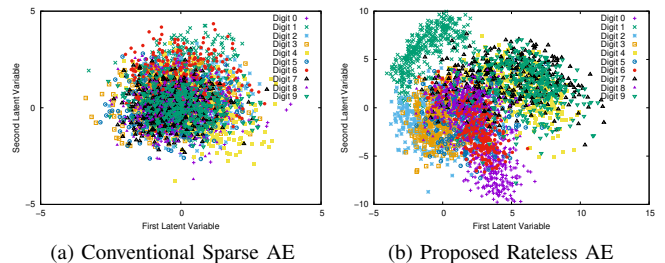


Fig. 6: The first two latent variables of MNIST test images encoded by SAE and RL-AE, which are trained at a latent dimensionality of $M = 64$ in terms of MSE measure.

images are plotted for the traditional SAE and proposed RL-AE, respectively. One can clearly see that the label-dependent distribution in our RL-AE is more clearly observable than the conventional AE, since the most-principal latent components are properly associated with the upper latent variables via the proposed stochastic bottleneck technique.

IV. CONCLUSIONS

We proposed new a type of auto-encoders employing a form of stochastic bottlenecking with non-uniform dropout rates for flexible dimensionality reduction. The proposed auto-encoders are rateless, i.e., the compression rate in dimensionality reduction is not pre-determined at the training phase and the user can freely change the dimensionality at testing phase without severely degrading quality. To realize rateless AEs, a simple regularization method called TailDrop was introduced to impose higher priority at upper neurons for learning the most-principal nonlinear features. This paper showed proof-of-concept results based on the standard MNIST and CIFAR-10 image datasets. Nearly optimal distortion performance was obtained with a single AE model irrespective of the flexible dimensionality reduction rate, which was obtained by simply dropping the least-principal latent dimensions. More rigorous analysis and theoretical optimization of dropout rate distributions for real-world data are left for future work. Multi-objective learning to account for various downstream applications is also an important open question to pursue.

REFERENCES

- [1] Van Der Maaten, L., Postma, E. & Van den Herik, J. (2009). Dimensionality reduction: A comparative review, *J Mach Learn Res.* **10**(66-71):13.
- [2] Jimenez, L.O. & Landgrebe, D.A. (1998). Supervised classification in high-dimensional space: Geometrical, statistical, and asymptotical properties of multivariate data. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews* **28**(1):39–54.
- [3] Hinton, G.E. & Salakhutdinov, R.R. (2006). Reducing the dimensionality of data with neural networks. *Science* **313**(5786):504–507.
- [4] Scholz, M., Fraunholz, M. & Selbig, J. (2008). Nonlinear principal component analysis: Neural network models and applications. In *Principal manifolds for data visualization and dimension reduction* (pp. 44–67). Springer, Berlin, Heidelberg.
- [5] Kramer, M.A. (1991). Nonlinear principal component analysis using auto-associative neural networks. *AIChE Journal* **37**(2):233–243.
- [6] DeMers, D., Cottrell, G.W. (1993). Nonlinear dimensionality reduction. *Advances in Neural Information Processing Systems 5*, San Mateo, CA, Morgan Kaufmann, 580–587.
- [7] Ng, A. (2011). Sparse autoencoder. *CS294A Lecture notes* **72**:1–19.
- [8] Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y. & Manzagol, P. A. (2010). Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research* **11**:3371–3408.
- [9] Doersch, C. (2016). Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*.
- [10] Sønderby, C.K., Raiko, T., Maaløe, L., Sønderby, S.K. & Winther, O. (2016). Ladder variational autoencoders. In *Advances in neural information processing systems* (pp. 3738–3746).
- [11] Giraldo, L.G.S. & Principe, J.C. (2013). Rate-distortion auto-encoders. *arXiv preprint arXiv:1312.7381*.
- [12] Theis, L., Shi, W., Cunningham, A. & Huszár, F. (2017). Lossy image compression with compressive autoencoders. *arXiv preprint arXiv:1703.00395*.
- [13] Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R.R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*.
- [14] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* **15**(1):1929–1958.
- [15] Wan, L., Zeiler, M., Zhang, S., Le Cun, Y. & Fergus, R. (2013). Regularization of neural networks using dropconnect. In *International conference on machine learning* (pp. 1058–1066).
- [16] Wu, Z., Nagarajan, T., Kumar, A., Rennie, S., Davis, L.S., Grauman, K. & Feris, R. (2018). Blockdrop: Dynamic inference paths in residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 8817–8826).
- [17] Huang, G., Sun, Y., Liu, Z., Sedra, D. & Weinberger, K.Q. (2016). Deep networks with stochastic depth. In *European conference on computer vision* (pp. 646–661). Springer, Cham.
- [18] Larsson, G., Maire, M. & Shakhnarovich, G. (2016). Fractalnet: Ultra-deep neural networks without residuals. *arXiv preprint arXiv:1605.07648*.
- [19] Yamada, Y., Iwamura, M., Akiba, T. & Kise, K. (2018). Shake-Drop Regularization for Deep Residual Learning. *arXiv preprint arXiv:1802.02375*.
- [20] Tompson, J., Goroshin, R., Jain, A., LeCun, Y. & Bregler, C. (2015). Efficient object localization using convolutional networks. *IEEE Conference on Computer Vision and Pattern Recognition* (pp. 648–656).
- [21] Krueger, D., Maharaj, T., Kramár, J., Pezeshki, M., Ballas, N., Ke, N.R. & Pal, C. (2016). Zoneout: Regularizing RNNs by randomly preserving hidden activations. *arXiv preprint arXiv:1606.01305*.
- [22] Gastaldi, X. (2017). Shake-shake regularization. *arXiv preprint arXiv:1705.07485*.
- [23] Huang, Z. & Wang, N. (2017). Data-driven sparse structure selection for deep neural networks. *arXiv preprint arXiv:1707.01213*.
- [24] Gal, Y. & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning* (pp. 1050–1059).
- [25] Wang, Z., Bovik, A.C., Sheikh, H.R. & Simoncelli, E.P. (2004). Image quality assessment: From error visibility to structural similarity. *IEEE transactions on image processing* **13**(4):600–612.
- [26] MacKay, D.J. (2005). Fountain codes. *IEE Proceedings-Communications* **152**(6): 1062–1068.
- [27] Krizhevsky, A. & Hinton, G. (2009). Learning multiple layers of features from tiny images. *Technical report, University of Toronto* **1**(4):7.
- [28] Kingma, D.P. & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [29] Xiao, H., Rasul, K. & Vollgraf, R. (2017). Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- [30] Clanuwat, T., Bober-Irizar, M., Kitamoto, A., Lamb, A., Yamamoto, K. & Ha, D. (2018). Deep learning for classical Japanese literature. *arXiv preprint arXiv:1812.01718*.
- [31] Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B. & Ng, A. Y. (2011). Reading digits in natural images with unsupervised feature learning. *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*.