

Bidirectional Variational Autoencoders

Bart Kosko

Department of Electrical and Computer Engineering
University of Southern California
Los Angeles, CA, USA
kosko@usc.edu

Olaoluwa Adigun

Department of Electrical and Computer Engineering
University of Southern California
Los Angeles, CA, USA
adigun@usc.edu

Abstract—We present the new **bidirectional variational autoencoder (BVAE) network architecture**. The BVAE uses a single neural network both to encode and decode instead of an encoder-decoder network pair. The network encodes in the forward direction and decodes in the backward direction through the same synaptic web. Simulations compared BVAEs and ordinary VAEs on the four image tasks of image reconstruction, classification, interpolation, and generation. The image datasets included MNIST handwritten digits, Fashion-MNIST, CIFAR-10, and CelebA-64 face images. The bidirectional structure of BVAEs cut the parameter count by almost 50% and still slightly outperformed the unidirectional VAEs.

Index Terms—variational autoencoder, bidirectional backpropagation, directional likelihoods, ELBO, evidence lower bound

I. DIRECTIONAL VARIATIONAL AUTOENCODERS

This paper introduces the new *bidirectional* variational autoencoder (BVAE) network. This architecture uses a single parametrized network for encoding and decoding. It trains with the new bidirectional backpropagation algorithm that jointly optimizes the network’s bidirectional likelihood [1], [2]. The algorithm uses the same synaptic weights both to predict the target y given the input x and to predict the converse x given y . Ordinary or unidirectional VAEs use separate networks to encode and decode.

Unidirectional variational autoencoders (VAEs) are unsupervised machine-learning models that learn data representations [3], [4]. They both learn and infer with directed probability models that often use intractable probability density functions [5]. A VAE seeks the best estimate of the data likelihood $p(x|\theta)$ from samples $\{x^{(n)}\}_{n=1}^N$ if x depends on some observable feature z and if θ represents the system parameters. The intractability involves marginalizing out the random variable z to give the likelihood $p(x|\theta)$:

$$p(x|\theta) = \mathbb{E}_{z|\theta}[p(x|z, \theta)] = \int_z p(x|z, \theta) p(z|\theta) dz. \quad (1)$$

Kingma and Welling introduced VAEs to solve this computational problem [5]. The VAE includes a new recognition (or encoding) model $q(z|x, \phi)$ that approximates the intractable likelihood $q(z|x, \theta)$. The probability $q(z|x, \phi)$ represents a probabilistic *encoder* while $p(x|z, \theta)$ represents a probabilistic *decoder*. These probabilistic models use two neural networks

with different synaptic weights. Figure 1a shows the architecture of such a unidirectional VAE. The recognition model doubles the number of parameters and the computational cost of this solution.

The new bidirectional backpropagation (B-BP) algorithm trains a neural network to run forwards and backwards by jointly maximizing the respective directional probabilities. This among other things allows such a network to run backward from output code words to expected input patterns. Running a unidirectionally trained network backwards just produces noise. B-BP jointly maximizes the forward likelihood $q_f(z|x, \theta)$ and backward likelihood $p_b(x|z, \theta)$ or the equivalent sum of their respective log-likelihoods:

$$\theta^* = \arg \max_{\theta} q_f(z|x, \theta) p_b(x|z, \theta) \quad (2)$$


$$= \arg \max_{\theta} \underbrace{\ln q_f(z|x, \theta)}_{\text{Forward pass}} + \underbrace{\ln p_b(x|z, \theta)}_{\text{Backward pass}}. \quad (3)$$

A BVAE approximates the intractable $q(z|x, \theta)$ with the forward likelihood $q_f(z|x, \theta)$. Then the probabilistic *encoder* is $q_f(z|x, \theta)$ and the probabilistic *decoder* is $p_b(x|z, \theta)$. So the two densities share parameter θ and there is no need for a separate network. Figure 1b shows the architecture of a BVAE.

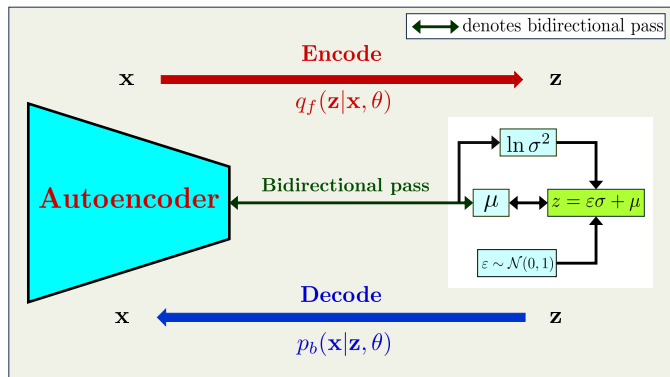
VAEs vary based on the choice of latent distribution, the method of training, and the use of joint modeling with other generative models, among other factors. The β -VAE introduced the adjustable hyperparameter β . It balances the latent channel capacity of the encoder network and the reconstruction error of the decoder network [6]. It trains on a weighted sum of the reconstruction error and the Kullback-Leibler divergence $D_{\text{KL}}(q(z|x, \phi)||p(z|\theta))$. The β -TCVAE (Total Correlation Variational Autoencoder) extends β -VAE to learning isolating sources of disentanglement [7]. A disentangled β -VAE modifies the β -VAE by progressively increasing the information capacity of the latent code while training [8].

Importance weighted autoencoders (IWAEs) use priority weights to derive a strictly tighter lower bound on the log-likelihood [9]. Variants of IWAE include the partially importance weighted auto-encoder (PIWAE), the multiply importance weighted auto-encoder (MIWAE), and the combined importance weighted auto-encoder (CIWAE) [10].

Hyperspherical VAEs use a non-Gaussian latent probability density. They use a von Mises-Fisher (vMF) latent density that



(a) Unidirectional VAE architecture



(b) Bidirectional VAE architecture

Fig. 1: Bidirectional vs. unidirectional variational autoencoders: Unidirectional VAEs use the forward passes of two separate networks for encoding and decoding. Bidirectional VAEs run their encoding on the forward pass and decoding on the backward pass with the same synaptic weights—weight matrices in both directions. This cuts the number of tunable parameters roughly in half. (a) The decoder network with parameter θ approximates $p(x|z, \theta)$ and the encoder network with parameter ϕ approximates $q(z|x, \theta)$. (b) Bidirectional VAEs use the forward pass of a network with parameter θ to approximate $q(z|x, \theta)$ and the backward pass of the network to approximate $p(x|z, \theta)$.

gives in turn a hyperspherical latent space [4]. Other VAEs include the Consistency Regularization for Variational Auto-Encoder (CRVAE) [3], the InfoVAE [11], and the Hamiltonian VAE [12] and so on. All these VAEs use separate networks to encode and decode.

Vincent et alia [13] suggests the use of tied weights in stacked autoencoder networks. This is a form of constraint that parallels the working of restricted Boltzmann machines RBMs [14] and thus a simple type of bidirectional associative memory or BAM [15]. It forces the weights to be symmetric using W^T on the backward pass. The building block here is a shallow network with no hidden layer [16], [17]. They further suggest that combining this constraint with a nonlinear activation would most likely lead to poor reconstruction error.

Bidirectional autoencoders BAEs [18] extend bidirectional neural representations to image compression and denoising. BAEs differ from autoencoders with tied weights because they relax the constraint by extending the bidirectional assumption over the depth of a deep network. BAEs differ from bidirectional VAEs because they do not require the joint optimization of the directional likelihoods. This limits the generative capability of BAEs.

The next sections review ordinary VAEs and introduce probabilistic BVAEs using the new B-BP algorithm. Section IV compares them on the four standard image test datasets: MNIST handwritten digits, Fashion-MNIST, CIFAR-10, and CelebA-64 datasets. We find that BVAEs cut the number of tunable parameters roughly in half while still performing slightly better than the unidirectional VAEs.

II. UNIDIRECTIONAL VARIATIONAL AUTOENCODERS

Let $p(x|\theta)$ denote the data likelihood and z denote the hidden variable. The data likelihood simplifies as

$$p(x|\theta) = \frac{p(x, z|\theta)}{q(z|x, \theta)} = \frac{p(x|z, \theta)p(z|\theta)}{q(z|x, \theta)}. \quad (4)$$

The likelihood $q(z|x, \theta)$ is intractable to solve. So unidirectional VAEs introduce a new likelihood that represents the *recognition* or *encoding* model. The term $q_f(z|x, \phi)$ represents the forward likelihood of the encoding network that approximates the intractable likelihood $q(z|x, \theta)$:

$$p(x|\theta) = \frac{p(x|z, \theta)p(z|\theta)}{q(z|x, \theta)} = \frac{p(x|z, \theta)p(z|\theta)q_f(z|x, \phi)}{q(z|x, \theta)q_f(z|x, \phi)}. \quad (5)$$

The corresponding data log-likelihood $\ln p(x|\theta)$ is

$$\ln p(x|\theta) = \ln \left[\frac{p(x|z, \theta)p(z|\theta)q_f(z|x, \phi)}{q(z|x, \theta)q_f(z|x, \phi)} \right] \quad (6)$$

$$= \ln p(x|z, \theta) + \ln \frac{p(z|\theta)}{q_f(z|x, \phi)} + \ln \frac{q_f(z|x, \phi)}{q(z|x, \theta)} \quad (7)$$

$$= \ln p(x|z, \theta) - \ln \frac{q_f(z|x, \phi)}{p(z|\theta)} + \ln \frac{q_f(z|x, \phi)}{q(z|x, \theta)}. \quad (8)$$

Now take the expectation of (8) with respect to $q_f(z|x, \phi)$:

$$\mathbb{E}_{z|x, \phi} [\ln p(x|\theta)] = \int_z q_f(z|x, \phi) \ln p(x|\theta) dz \quad (9)$$

$$= \ln p(x|\theta) \int_z q_f(z|x, \phi) dz \quad (10)$$

$$= \ln p(x|\theta) \quad (11)$$

because $q_f(z|x, \phi)$ is a probability density function and its integral over the domain of z equals 1. The expectation of the term on the right-hand side of (8) with respect $q_f(z|x, \phi)$ is

$$\begin{aligned} \mathbb{E}_{z|x, \phi} \left[\ln p(x|z, \theta) - \ln \frac{q_f(z|x, \phi)}{p(z|\theta)} + \ln \frac{q_f(z|x, \phi)}{q(z|x, \theta)} \right] \\ = \mathbb{E}_{z|x, \phi} [\ln p(x|z, \theta)] - D_{\text{KL}}(q_f(z|x, \phi) || p(z|\theta)) \\ + D_{\text{KL}}(q_f(z|x, \phi) || q(z|x, \theta)) \end{aligned} \quad (12)$$

where

$$D_{\text{KL}}(q_f(z|x, \phi) || p(z|\theta)) = \mathbb{E}_{z|x, \phi} \left[\ln \frac{q_f(z|x, \phi)}{p(z|\theta)} \right] \quad (13)$$

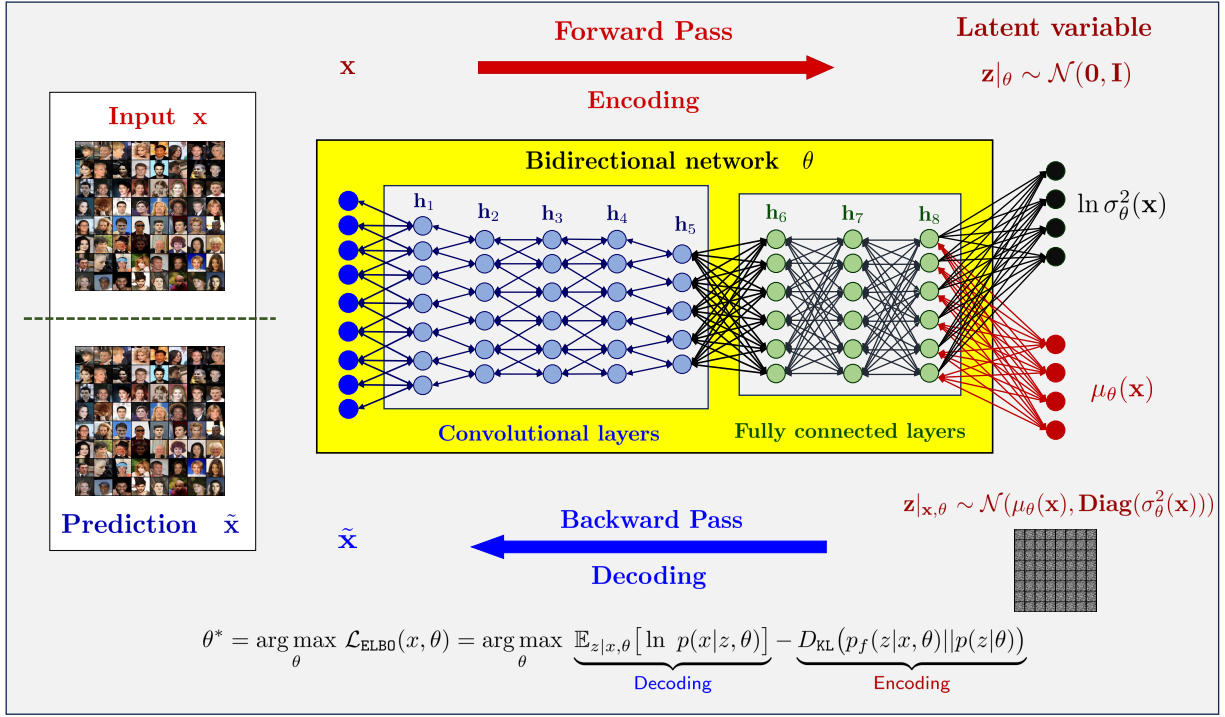


Fig. 2: Training a bidirectional variational autoencoder with bidirectional backpropagation algorithm: This framework uses a single network for encoding and decoding. The forward pass with likelihood $q_f(z|x, \theta)$ runs the encoding to the latent space. The backward pass with likelihood $p_b(x|z, \theta)$ decodes the latent features.

and

$$D_{\text{KL}}(q_f(z|x, \phi) || q(z|x, \theta)) = \mathbb{E}_{z|x, \phi} \left[\ln \frac{q_f(z|x, \phi)}{q(z|x, \theta)} \right]. \quad (14)$$

Combining (8), (9), and (12) gives

$$\begin{aligned} \ln p(x|\theta) &= \underbrace{\mathbb{E}_{z|x, \phi} [\ln p(x|z, \theta)]}_{\text{Decoding}} - \underbrace{D_{\text{KL}}(q_f(z|x, \phi) || p(z|\theta))}_{\text{Encoding}} \\ &\quad + \underbrace{D_{\text{KL}}(q_f(z|x, \phi) || q(z|x, \theta))}_{\text{Error-gap}}. \end{aligned} \quad (15)$$

The KL-divergence between $q_f(z|x, \phi)$ and $q(z|x, \theta)$ yields the following inequality because of Jensen's inequality:

$$D_{\text{KL}}(q_f(z|x, \phi) || q(z|x, \theta)) = \mathbb{E}_{z|x, \phi} \left[\ln \frac{q_f(z|x, \phi)}{q(z|x, \theta)} \right] \quad (16)$$

$$= \int_z q_f(z|x, \phi) \ln \frac{q_f(z|x, \phi)}{q(z|x, \theta)} dz \quad (17)$$

$$= - \int_z q_f(z|x, \phi) \ln \frac{q(z|x, \theta)}{q_f(z|x, \phi)} dz \quad (18)$$

$$\geq - \ln \int_z q_f(z|x, \phi) \frac{q(z|x, \theta)}{q_f(z|x, \phi)} dz \quad (19)$$

$$= - \ln \int_z q(z|x, \theta) dz \quad (20)$$

$$= - \ln 1 = 0 \quad (21)$$

because the negative of the natural logarithm is convex. So

$$\ln p(x|\theta) \geq \underbrace{\mathbb{E}_{z|x, \phi} [\ln p(x|z, \theta)] - D_{\text{KL}}(q_f(z|x, \phi) || p(z|\theta))}_{\mathcal{L}(x, \theta, \phi)} \quad (22)$$

and

$$\mathcal{L}(x, \theta, \phi) = \mathbb{E}_{z|x, \phi} [\ln p(x|z, \theta)] - D_{\text{KL}}(q_f(z|x, \phi) || p(z|\theta)) \quad (23)$$

where $\mathcal{L}(x, \theta, \phi)$ is the *evidence lower bound* (ELBO) on the data log-likelihood $p(x|\theta)$.

Unidirectional VAEs train on the estimate $\tilde{\mathcal{L}}_{\text{ELBO}}(x, \theta, \phi)$ of the ELBO using ordinary or unidirectional backpropagation (BP). This estimate involves using the forward pass $q_f(z|x, \phi)$ to approximate the intractable encoding model $q(z|x, \theta)$ and the forward pass $p_f(x|z, \theta)$ to approximate the encoding model. The gradient update rules for the encoder and decoder networks at the $(n+1)^{\text{th}}$ iteration or training epoch are

$$\theta^{(n+1)} = \theta^{(n)} + \eta \nabla_{\theta} \tilde{\mathcal{L}}_{\text{ELBO}}(x, \theta, \phi) \Bigg|_{\theta=\theta^{(n)}, \phi=\phi^{(n)}} \quad (24)$$

$$\phi^{(n+1)} = \phi^{(n)} + \eta \nabla_{\phi} \tilde{\mathcal{L}}_{\text{ELBO}}(x, \theta, \phi) \Bigg|_{\theta=\theta^{(n)}, \phi=\phi^{(n)}} \quad (25)$$

where η is the learning rate, $\phi^{(n)}$ is the encoder parameter, and $\theta^{(n)}$ is the decoder parameter after n training iterations.

III. BIDIRECTIONAL VARIATIONAL AUTOENCODERS

Bidirectional VAEs use the directional likelihoods of a network with parameter θ to approximate the data log-likelihood $\ln p(x|\theta)$. They use the same bidirectional associative network to model the encoding and decoding phases. The forward-pass likelihood $q_f(z|x, \theta)$ models the encoding and the backward-pass likelihood $p_b(x|z, \theta)$ models the decoding. So BVAEs do not need an extra likelihood $q(z|x, \phi)$ or an extra network with parameter ϕ .

The data log-likelihood is

$$p(x|\theta) = \frac{p(x|z, \theta)p(z|\theta)}{q(z|x, \theta)}. \quad (26)$$

Then

$$\ln p(x|\theta) = \ln \left[\frac{p(x|z, \theta)p(z|\theta)}{q(z|x, \theta)} \right] \quad (27)$$

$$= \ln \left[\frac{p(x|z, \theta)p(z|\theta)q_f(z|x, \theta)}{q(z|x, \theta)q_f(z|x, \theta)} \right] \quad (28)$$

$$= \ln [p(x|z, \theta)] - \ln \left[\frac{q_f(z|x, \theta)}{p(z|\theta)} \right] + \ln \left[\frac{q_f(z|x, \theta)}{q(z|x, \theta)} \right]. \quad (29)$$

Now take the expectation of (29) with respect to $q_f(z|x, \theta)$ and consider the left-hand side of (29):

$$\mathbb{E}_{z|x, \theta} [\ln p(x|\theta)] = \int_z q_f(z|x, \theta) \ln p(x|\theta) dz \quad (30)$$

$$= \ln p(x|\theta) \int_z q_f(z|x, \theta) dz \quad (31)$$

$$= \ln p(x|\theta). \quad (32)$$

The expectation of the right-hand term is

$$\begin{aligned} \mathbb{E}_{z|x, \theta} \left[\ln \frac{p(x|z, \theta)}{p(z|\theta)} - \ln \frac{q_f(z|x, \theta)}{q(z|x, \theta)} \right] \\ = \mathbb{E}_{z|x, \theta} [\ln p(x|z, \theta)] - D_{\text{KL}}(q_f(z|x, \theta)||p(z|\theta)) \\ + D_{\text{KL}}(q_f(z|x, \theta)||q(z|x, \theta)) \end{aligned} \quad (33)$$

where

$$\mathbb{E}_{z|x, \theta} \left[\ln \frac{q_f(z|x, \theta)}{p(z|\theta)} \right] = D_{\text{KL}}(q_f(z|x, \theta)||p(z|\theta)) \quad (34)$$

and

$$\mathbb{E}_{z|x, \theta} \left[\ln \frac{q_f(z|x, \theta)}{q(z|x, \theta)} \right] = D_{\text{KL}}(q_f(z|x, \theta)||q(z|x, \theta)). \quad (35)$$

The corresponding data log-likelihood of a BVAE with parameter θ is

$$\begin{aligned} \ln p(x|\theta) = \underbrace{\mathbb{E}_{z|x, \theta} [\ln p(x|z, \theta)]}_{\text{Decoding}} - \underbrace{D_{\text{KL}}(q_f(z|x, \theta)||p(z|\theta))}_{\text{Encoding}} \\ + \underbrace{D_{\text{KL}}(q_f(z|x, \theta)||q(z|x, \theta))}_{\text{Error-gap}}. \end{aligned} \quad (36)$$

The log-likelihood of the BVAE is such that

$$\ln p(x|\theta) \geq \underbrace{\mathbb{E}_{z|x, \theta} [\ln p(x|z, \theta)] - D_{\text{KL}}(q_f(z|x, \theta)||p(z|\theta))}_{\mathcal{L}(x, \theta)} \quad (37)$$

because the KL-divergence $D_{\text{KL}}(q_f(z|x, \theta)||p(z|\theta)) \geq 0$. So

$$\mathcal{L}(x, \theta) = \mathbb{E}_{z|x, \theta} [\ln p(x|z, \theta)] - D_{\text{KL}}(q_f(z|x, \theta)||p(z|\theta)) \quad (38)$$

where $\mathcal{L}(x, \theta)$ is the ELBO on $\ln p(x|\theta)$ and we take the expectation $\mathbb{E}_{z|x, \theta}$ with respect to $q_f(z|x, \theta)$.

Bidirectional VAEs train on the estimate $\tilde{\mathcal{L}}_{\text{ELBO}}(x, \theta)$ of the ELBO that uses bidirectional neural representation [2]. This estimate involves using the forward pass $q_f(z|x, \theta)$ to approximate the intractable encoding model $q(z|x, \theta)$ and the reverse pass $p_b(x|z, \theta)$ to approximate the decoding model. The update rule at the $(n+1)^{\text{th}}$ iteration or training epoch is

$$\theta^{(n+1)} = \theta^{(n)} + \eta \nabla_{\theta} \tilde{\mathcal{L}}_{\text{ELBO}}(x, \theta) \Big|_{\theta=\theta^{(n)}} \quad (39)$$

where η is the learning rate and $\theta^{(n)}$ is the autoencoder network parameter just after the n^{th} training iteration. Figure 2 shows the probabilistic approximation of a BVAE with the directional likelihoods of a bidirectional network.

IV. SIMULATIONS

We compared the performance of unidirectional VAEs and bidirectional VAEs using different tasks, datasets, network architectures, and loss functions. We first describe the image test sets for our experiments.

A. Datasets

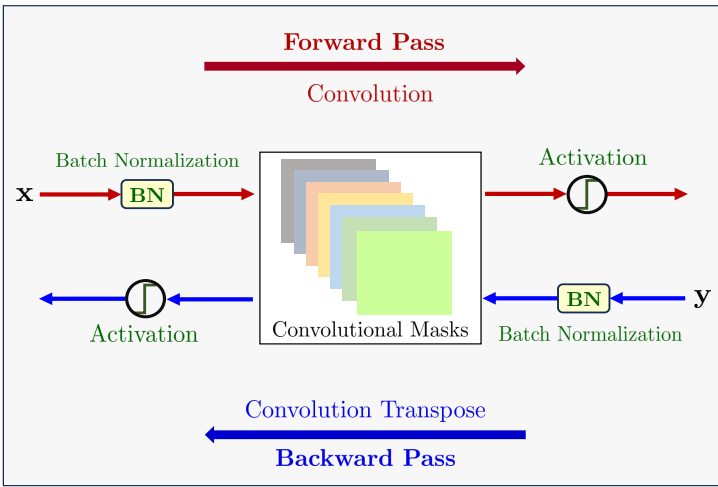
The simulations compared results on four standard image datasets: MNIST handwritten digits [19], Fashion-MNIST [20], CIFAR-10 [21], and CelebA [22] datasets.

The MNIST handwritten digit dataset contains 10 classes of handwritten digits $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$. This dataset consists of 60,000 training samples with 6,000 samples per class, and 10,000 test samples with 1,000 samples per class. Each image is a single-channel image with dimension 28×28 .

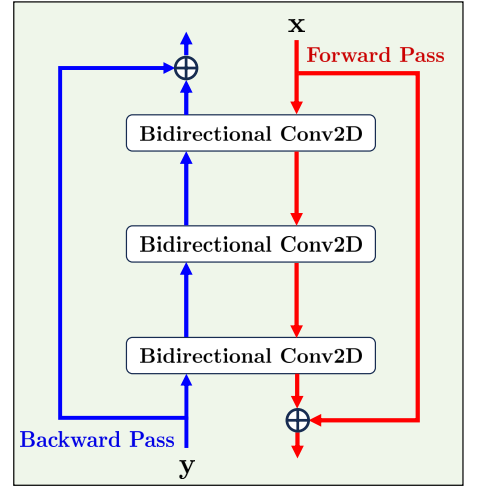
The Fashion-MNIST dataset is a database of fashion images. It is made of 10 classes namely ankle boot, bag, coat, dress, pullover, sandal, shirt, sneaker, trouser, and t-shirt / top. Each class has 6,000 training samples and 1,000 testing samples. Each image is also a single-channel image with dimension 28×28 .

The CIFAR-10 dataset consists of 60,000 color images from 10 categories. Each image has size $32 \times 32 \times 3$. The 10 pattern categories are airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck. Each class consists of 5,000 training samples and 1,000 testing samples.

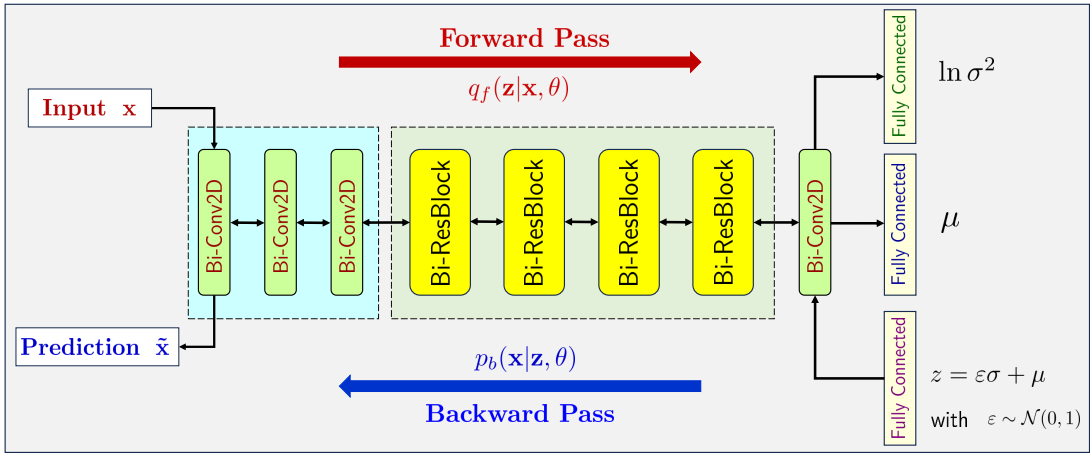
The CelebA dataset is a large-scale face dataset of 10,177 celebrities [22]. This dataset is made up of 202,599 color



(a) Bidirectional convolutional (Bi-Conv2D) layer



(b) Bidirectional residual block (Bi-ResBlock)



(c) Bidirectional residual VAE network

Fig. 3: Bidirectional VAE with residual network architecture: This cuts the tunable parameters roughly in half compared with unidirectional VAEs. (a) is the bidirectional convolutional layer. Convolution runs in the forward pass and convolution transpose runs in reverse with the same set of convolution masks. (b) is the architecture of a bidirectional residual block with bidirectional skip connections.

(three-channel) images. This is not a balanced dataset. The number of images per celebrity varies between 1 – 35. We divided the dataset into two splits of 9,160 celebrities for training and 1,017 celebrities for testing the VAEs. This resulted in 185,133 training samples and 17,466 testing samples. We resized each image to $64 \times 64 \times 3$.

B. Tasks

We compared the performance of bidirectional VAEs and unidirectional VAEs on the following four tasks.

1) *Image compression and reconstruction*: We tested the self-mapping of the image datasets with unidirectional and bidirectional VAEs. This involved the encoding of images with latent variable z and the subsequent decoding to reconstruct the image after the latent sampling. We evaluated the performance of VAEs on this task using the Peak Signal-to-Noise

Ratio (PSNR), Fréchet Inception Distance (FID) [23], [24], and Structural Similarity Index Measure (SSIM) [25].

2) *Downstream image classification*: We trained simple classifiers on the latent space features of VAEs. These VAEs compress the input images and a simple classifier maps the latent features to their corresponding classes. We evaluated the classification accuracy of this downstream classification.

The VAE-extracted features for the MNIST handwritten dataset trained on simple linear classifiers. We used a neural classifier with one hidden layer of 256 logistic hidden neurons to classify the VAE-extracted features from the Fashion-MNIST and CIFAR-10 datasets.

We used the t -distributed stochastic neighbor embedding (t -SNE) method to visualize the reduced features. This method uses a statistical approach to map the high-dimensional representation of data $\{\mathbf{x}_i\}_{i=1}^N$ to their respective low-dimensional

TABLE I: MNIST handwritten digits dataset with VAEs. We used the residual network architecture. The dimension of the latent variable is 64. The BVAEs each used 42.2MB storage memory and the VAEs each used 84.4MB storage memory.

Model	Parameters	Generative Task			Reconstruction and Classification			
		NLL ↓	AU	FID ↓	PSNR ↑	SSIM ↑	rFID ↓	Accuracy ↑
VAE	22.1M	86.72	18	4.340	20.59	0.8959	3.510	94.97%
Bidirectional VAE	11.1M	86.69	15	3.370	20.64	0.8988	3.050	95.26%
β -VAE ($\beta = 0.5$)	22.1M	88.12	23	2.021	22.19	0.9306	1.711	93.15%
Bidirectional β-VAE ($\beta = 0.5$)	11.1M	88.09	26	2.034	22.39	0.9332	1.702	94.32%
β -VAE ($\beta = 1.5$)	22.1M	86.99	15	4.855	19.75	0.8726	4.370	95.20%
Bidirectional β-VAE ($\beta = 1.5$)	11.1M	87.21	15	5.589	19.80	0.8736	5.011	96.19%
β -TCVAE ($\beta = 0.5$)	22.1M	88.18	24	1.975	22.21	0.9299	1.662	93.37%
Bidirectional β-TCVAE ($\beta = 0.5$)	11.1M	87.99	26	2.000	22.30	0.9325	1.760	94.17%
β -TCVAE ($\beta = 1.0$)	22.1M	86.87	18	3.544	20.68	0.8988	3.143	95.08%
Bidirectional β-TCVAE ($\beta = 1.0$)	11.1M	86.68	18	3.419	20.75	0.9007	3.085	95.66%
β -TCVAE ($\beta = 1.5$)	22.1M	87.20	14	5.623	19.64	0.8690	5.052	95.56%
Bidirectional β-TCVAE ($\beta = 1.5$)	11.1M	86.90	15	4.968	19.86	0.8756	4.548	96.54%
IWAE	22.1M	86.02	20	2.958	19.30	0.8627	2.774	96.40%
Bidirectional IWAE	11.1M	86.00	19	3.330	19.21	0.8587	3.030	96.29%

Algorithm 1: BVAE Training with bidirectional backpropagation

Input: Data $\{\mathbf{x}_n\}_{n=1}^N$ and latent space dimension J .
Initialize: Synaptic weights $\theta \in \{\mathbf{N}_\theta, \mathbf{V}_\theta, \mathbf{W}_\theta\}$, learning rate α , and other hyper-parameters.

- for** iteration $t = 1, 2, \dots$ **do**
- Pick a mini-batch $\{\mathbf{x}_m\}_{m=1}^B$ of B samples
- for** $m = 1, 2, \dots, B$ **do**
- Forward Pass (Encoding):** Predict the variational mean and log-covariance:
$$\hat{\mu}_m = \mathbf{W}_\theta(\mathbf{N}_\theta(\mathbf{x}_m))$$
and
$$\ln \hat{\sigma}_m^2 = \mathbf{V}_\theta(\mathbf{N}_\theta(\mathbf{x}_m))$$
- Sample the latent features \mathbf{z}_m from the variational Gaussian distribution with condition \mathbf{x}_m :
$$\varepsilon_m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad \text{and} \quad \mathbf{z}_m = \hat{\mu}_m + \varepsilon_m \cdot \hat{\sigma}_m$$
- Backward Pass (Decoding):** Map the latent variable back to the input space:
$$\hat{a}_m^{(x)} = \mathbf{N}_\theta^\top(\mathbf{W}_\theta^\top(\mathbf{z}_m))$$
- end for**
- Estimate the negative log-likelihood $\text{NLL}(\mathbf{x}, \theta)$:
$$\text{KLD}(\mathbf{x}, \theta) = \frac{1}{B} \sum_{m=1}^B D_{\text{KL}}(\mathcal{N}(\hat{\mu}_m, \mathbf{Diag}(\hat{\sigma}_m^2)) || \mathcal{N}(\mathbf{0}, \mathbf{I}))$$

$$\text{BCE}(\mathbf{x}, \theta) = -\frac{1}{B} \sum_{m=1}^B \ln p_b(\mathbf{x}_m | \mathbf{z}_m, \theta)$$

$$\tilde{\mathcal{L}}_{\text{ELBO}}(\mathbf{x}, \theta) = \text{BCE}(\mathbf{x}, \theta) + \text{KLD}(\mathbf{x}, \theta)$$
- Update θ by backpropagating $\tilde{\mathcal{L}}_{\text{ELBO}}(\mathbf{x}, \theta)$ through the weights.
- end for**
- return** θ

representation $\{\mathbf{y}_i\}_{i=1}^N$ based on the similarity of the datapoints [26]. This low-dimensional representation provides insight into the degree of separability among the classes.

3) *Image generation:* We compared the generative performance of bidirectional VAEs with their corresponding unidirectional ones. These VAEs trained with the Gaussian latent distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. We tested the VAEs using both Gaussian sampler and Gaussian mixture model samplers post-training [27]. We used the estimates of the negative of the data log-likelihood (NLL) and the number of active latent units (AU) [28] as the quantitative metrics for performance evaluation.

4) *Image interpolation:* We conducted linear interpolation of samples. The interpolations involve a convex combination of two images over 10 steps. The encoding step transforms the mixture of two images in the latent space. The decoding step reconstructs the interpolated samples.


C. Model Architecture

We used different neural network architectures for various datasets and tasks.

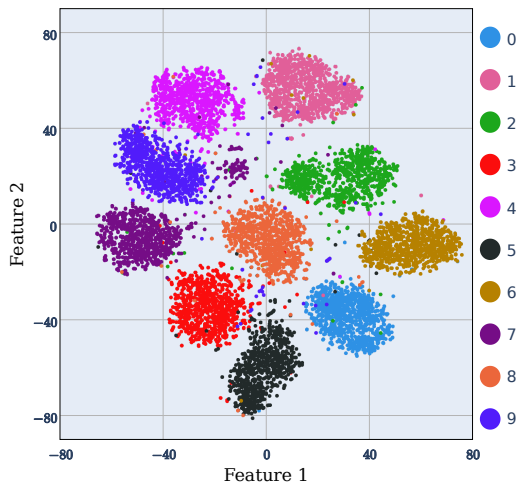
Variational Autoencoders: We used deep convolutional and residual neural network architectures. VAEs that trained on the MNIST handwritten and Fashion-MNIST used the residual architecture. Figure 3c shows the architecture of the bidirectional residual networks that trained on the MNIST datasets. The unidirectional VAEs with this architecture used two such networks each: one for encoding and the other for decoding. The BVAEs used just one of such network each. Encoding runs in the forward pass and decoding runs in the backward pass.

Each of the encoder and decoder networks that trained on the CIFAR-10 dataset used six convolutional layers and two fully connected layers. The corresponding BVAEs used only one network for encoding and decoding each. The dimension of the hidden convolutional layers is $\{64 \leftrightarrow 128 \leftrightarrow 256 \leftrightarrow 512 \leftrightarrow 1024 \leftrightarrow 2048\}$. The dimension of the fully connected layers is $\{2048 \leftrightarrow 1024 \leftrightarrow 64\}$.

The configuration of the VAEs that trained on the CelebA dataset differs slightly. The sub-networks each used nine convolutional layers and two fully connected layers. The dimension of the hidden convolutional layers is $\{128 \leftrightarrow$



(a) Unidirectional β -VAE



(b) Bidirectional β -VAE

Fig. 4: t -SNE embedding for the MNIST handwritten digit dataset: Latent space dimension is 128. (a) A simple linear classifier that trained on the unidirectional VAE-compressed features achieved a 95.2% accuracy. (b) The simple classifier achieved 97.32% accuracy when it trained on the BVAE-compressed features.

$128 \leftrightarrow 192 \leftrightarrow 256 \leftrightarrow 384 \leftrightarrow 512 \leftrightarrow 768 \leftrightarrow 1024 \leftrightarrow 1024$. The dimension of the fully connected layers is $\{4096 \leftrightarrow 2048 \leftrightarrow 256\}$.

The VAEs used *generalized nonvanishing* (G-NoVa) hidden neurons [29], [30]. The G-NoVa activation $a(x)$ of input x is

$$a(x) = \alpha x + x\sigma(\beta x) = \alpha x + \frac{x}{1 + e^{-\beta x}} \quad (40)$$

where $\alpha > 0$ and $\beta > 0$. Each layer of a BVAE performs probabilistic inference in both the forward and backward passes. The convolutional layers use bidirectional kernels. The kernels run convolution in the forward pass and transposed convolution in the backward pass. Transposed convolution projects feature maps to a higher-dimensional space [31].

Downstream Classification: We trained simple linear classifiers on VAE-extracted features from the MNIST digit dataset. We trained shallow neural classifiers with one hidden layer and 100 hidden neurons each on the extracted features from the Fashion-MNIST images. Similar neural classifiers with one hidden layer and 256 hidden neurons each trained on VAE-extracted features from the CIFAR-10 dataset.

D. Training

We considered four implementations of VAEs and compare them with their respective bidirectional versions. The four VAEs are vanilla VAE [5], β -VAE [6], β -TCVAE [7], and IWAE [9]. We trained these over the four datasets across the four tasks. We used the AdamW optimizer [32] with the OneCycleLR [33] learning rate scheduler. The optimizer trained on their respective ELBO estimates.

We designed a new framework for bidirectional VAEs and implemented unidirectional VAEs with the Pythae framework [34]. All the models trained on a single A100 GPU. Tables I - IV and Figures 4 - 6 present the results.

E. Evaluation Metrics

We considered the performance of the VAE models on generative and compression tasks. We used the following quantitative metrics:

- *Negative Log-Likelihood (NLL):* This is an estimate of the negative of $\ln p(x|\theta)$. This is computational intractable so we used the Monte Carlo method to estimate this. Lower value means the model generalizes well to unseen data.
- *Number of Active Latent Units (AU)* [28]: This reflects the number of latent variables with variance above a given threshold ϵ . We have

$$AU = \sum_{d=1}^D \mathbb{I} \left[Cov_{\mathbf{x}} \left[\mathbb{E}_{\mathbf{z}|\mathbf{x},\phi} [\mathbf{z}_d] \right] \geq \epsilon \right] \quad (41)$$

where \mathbb{I} is an indicator function, \mathbf{z}_d represents the d^{th} component of the latent variable, and $\epsilon = 0.01$. Higher AU means the model uses more features for the latent space representation. But having too many active units can lead to overfitting.

- *Peak Signal-to-Noise Ratio (PSNR):* This compares reconstructed images with their target images. Higher value implies better reconstruction from data compression.
- *Structural Similarity Index (SSIM):* This is a perceptual metric. It quantifies the degradation from data compression. Higher SSIM value implies better reconstruction from image compression.
- *Downstream Classification Accuracy:* This is the classification accuracy of simple classifiers that trained on the latent or VAE-extracted features. Higher accuracy means the compression extracts easy-to-classify features.
- *Fréchet Inception Distance (FID)* [23]: This metric evaluates the quality of generated images. It measures the similarity between the distribution of real images and the distribution of generated images. Lower value implies that the generated images are closer to the real images.




Fig. 5: MNIST handwritten image: Image interpolation with variational autoencoder networks.

TABLE II: Fashion-MNIST dataset with VAEs. We used the residual network architecture. The dimension of the latent variable is 64. The BVAEs used 42.2MB memory parameters and the unidirectional VAEs used 84MB memory parameters.

Model	Parameters	Generative Task			Reconstruction and Classification			
		NLL ↓	AU	FID ↓	PSNR ↑	SSIM ↑	rFID ↓	Accuracy ↑
VAE	22.1M	231.3	13	3.082	19.19	0.6964	2.521	87.17%
Bidirectional VAE	11.1M	231.1	15	3.045	19.21	0.6973	2.478	87.84%
β -VAE ($\beta = 0.5$)	22.1M	232.4	17	1.946	20.01	0.7282	1.569	87.16%
Bidirectional β-VAE ($\beta = 0.5$)	11.1M	232.4	20	1.965	20.05	0.7303	1.658	87.93%
β -VAE ($\beta = 1.5$)	22.1M	231.5	9	3.220	18.59	0.6714	2.947	86.89%
Bidirectional β-VAE ($\beta = 1.5$)	11.1M	231.5	9	3.346	18.60	0.6729	2.893	86.66%
β -TCVAE ($\beta = 0.5$)	22.1M	232.3	24	2.420	20.04	0.7238	1.996	87.71%
Bidirectional β-TCVAE ($\beta = 0.5$)	11.1M	232.3	19	2.026	20.09	0.7288	1.749	88.06%
β -TCVAE ($\beta = 1.0$)	22.1M	231.1	12	2.833	19.18	0.6941	2.427	87.21%
Bidirectional β-TCVAE ($\beta = 1.0$)	11.1M	231.2	12	2.891	19.21	0.6954	2.469	87.06%
β -TCVAE ($\beta = 1.5$)	22.1M	231.8	9	3.579	18.54	0.6678	3.234	86.62%
Bidirectional β-TCVAE ($\beta = 1.5$)	11.1M	231.4	10	3.273	18.58	0.6702	2.969	87.00%
IWAE	22.1M	230.3	17	2.489	17.78	0.6529	2.070	88.14%
Bidirectional IWAE	11.1M	230.4	14	2.881	18.10	0.6559	2.432	88.05%




Fig. 6: Image interpolation with VAEs on the Fashion-MNIST dataset.

TABLE III: CIFAR-10 dataset with VAEs. The dimension of the latent space is 256. The BVAEs each used 107MB memory parameters and the unidirectional VAEs each used 214MB memory parameters.

Model	Parameters	Generative Task			Reconstruction and Classification			
		NLL ↓	AU	FID ↓	PSNR ↑	SSIM ↑	rFID ↓	Accuracy ↑
VAE	56.6M	1817.0	51	2.483	18.67	0.4532	2.305	51.73%
Bidirectional VAE	28.3M	1814.1	37	2.442	18.93	0.4742	2.255	50.55%
β -VAE ($\beta = 0.5$)	56.6M	1816.7	85	2.225	19.48	0.5191	1.955	51.78%
Bidirectional β-VAE ($\beta = 0.5$)	28.3M	1815.1	65	2.122	19.83	0.5466	1.827	51.74%
β -VAE ($\beta = 1.5$)	56.6M	1820.9	32	2.625	18.12	0.4113	2.483	50.44%
Bidirectional β-VAE ($\beta = 1.5$)	28.3M	1816.1	40	2.535	18.39	0.4337	2.342	51.91%
β -TCVAE ($\beta = 0.5$)	56.6M	1816.8	50	2.225	19.26	0.5035	2.035	50.96%
Bidirectional β-TCVAE ($\beta = 0.5$)	28.8M	1814.0	59	2.101	19.63	0.5321	1.942	50.51%
β -TCVAE ($\beta = 1.0$)	56.6M	1819.6	27	2.485	18.47	0.4362	2.348	49.57%
Bidirectional β-TCVAE ($\beta = 1.0$)	28.3M	1814.0	42	2.380	18.91	0.4752	2.201	50.93%
β -TCVAE ($\beta = 1.5$)	56.6M	1824.1	33	2.673	17.93	0.3940	2.604	50.52%
Bidirectional β-TCVAE ($\beta = 1.5$)	28.3M	1816.8	35	2.576	18.36	0.4315	2.437	51.04%
IWAE	56.6M	1814.6	122	2.415	18.55	0.4540	2.173	51.94%
Bidirectional IWAE	28.3M	1811.5	128	2.309	18.84	0.4838	2.015	51.65%

V. CONCLUSION

Bidirectional VAEs encode and decode through the same synaptic web of a deep neural network. This bidirectional

flow captures the joint probabilistic structure of both directions during learning and recall. BVAEs cut the synaptic parameter

TABLE IV: CelebA-64 dataset with VAEs. The dimension of the latent variable is 256. The BVAEs each used 133.8MB memory parameters and the unidirectional VAEs each used 267.6MB memory parameters.

Model	Parameters	Generative Task		Reconstruction	
		NLL ↓	AU	PSNR ↑	SSIM↑
VAE	69.1M	6221.9	80	20.54	0.6528
Bidirectional VAE	34.6M	6217.7	52	20.66	0.6532
β -VAE ($\beta = 0.1$)	69.1M	6243.0	240	22.35	0.7225
Bidirectional β-VAE	34.6M	6261.9	90	22.60	0.7243
β -TCVAE ($\beta = 0.1$)	69.1M	6277.2	93	22.15	0.7189
Bidirectional β-TCVAE	34.6M	6275.2	145	23.55	0.7681
β -IWAE	69.1M	6221.1	77	20.44	0.6521
Bidirectional β-IWAE	34.6M	6217.5	204	20.81	0.6551

count roughly in half compared with unidirectional VAEs. The simulations on the four image test sets showed that the BVAEs still performed slightly better than the unidirectional VAEs.

REFERENCES

- [1] O. Adigun and B. Kosko, "Bidirectional representation and backpropagation learning," in *International Joint Conference on Advances in Big Data Analytics, ABDA 2016*. CSREA, 2016, pp. 3–9.
- [2] —, "Bidirectional backpropagation," *IEEE Trans. Syst. Man Cybern. Syst.*, vol. 50, no. 5, pp. 1982–1994, 2020.
- [3] S. Sinha and A. B. Dieng, "Consistency regularization for variational auto-encoders," in *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021*, pp. 12 943–12 954.
- [4] T. R. Davidson, L. Falorsi, N. D. Cao, T. Kipf, and J. M. Tomczak, "Hyperspherical variational auto-encoders," in *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018, Monterey, USA, August 6-10, 2018*. AUAI Press, 2018, pp. 856–865.
- [5] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014*, 2014.
- [6] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "beta-vae: Learning basic visual concepts with a constrained variational framework," in *International conference on learning representations*, 2016.
- [7] R. T. Chen, X. Li, R. B. Grosse, and D. K. Duvenaud, "Isolating sources of disentanglement in variational autoencoders," *Advances in neural information processing systems*, vol. 31, 2018.
- [8] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner, "Understanding disentangling in β -vae," *CoRR*, vol. abs/1804.03599, 2018. [Online]. Available: <http://arxiv.org/abs/1804.03599>
- [9] Y. Burda, R. Grosse, and R. Salakhutdinov, "Importance weighted autoencoders," *arXiv preprint arXiv:1509.00519*, 2015.
- [10] T. Rainforth, A. Kosiorek, T. A. Le, C. Maddison, M. Igl, F. Wood, and Y. W. Teh, "Tighter variational bounds are not necessarily better," in *International Conference on Machine Learning*. PMLR, 2018, pp. 4277–4285.
- [11] S. Zhao, J. Song, and S. Ermon, "Infovae: Information maximizing variational autoencoders," *arXiv preprint arXiv:1706.02262*, 2017.
- [12] A. L. Caterini, A. Doucet, and D. Sejdinovic, "Hamiltonian variational auto-encoder," in *Advances in Neural Information Processing Systems 31: NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 2018, pp. 8178–8188.
- [13] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *Journal of Machine Learning Research*, vol. 11, no. 110, pp. 3371–3408, 2010. [Online]. Available: <http://jmlr.org/papers/v11/vincent10a.html>
- [14] P. Smolensky *et al.*, "Information processing in dynamical systems: Foundations of harmony theory," 1986.
- [15] B. Kosko, "Bidirectional associative memories," *IEEE Transactions on Systems, man, and Cybernetics*, vol. 18, no. 1, pp. 49–60, 1988.
- [16] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [17] H. Larochelle and Y. Bengio, "Classification using discriminative restricted boltzmann machines," in *Proceedings of the 25th international conference on Machine learning*, 2008, pp. 536–543.
- [18] O. Adigun and B. Kosko, "Bidirectional backpropagation autoencoding networks for image compression and denoising," in *2023 International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 2023, pp. 730–737.
- [19] Y. LeCun, "The mnist database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, 1998.
- [20] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
- [21] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [22] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [23] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 2017, pp. 6626–6637.
- [24] M. Fréchet, "Sur la distance de deux lois de probabilité," in *Annales de l'ISUP*, vol. 6, no. 3, 1957, pp. 183–198.
- [25] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [26] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [27] P. Ghosh, M. S. M. Sajjadi, A. Vergari, M. J. Black, and B. Schölkopf, "From variational to deterministic autoencoders," in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020.
- [28] Y. Burda, R. B. Grosse, and R. Salakhutdinov, "Importance weighted autoencoders," in *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016*, 2016.
- [29] O. Adigun and B. Kosko, "Deeper bidirectional neural networks with generalized non-vanishing hidden neurons," in *21st IEEE International Conference on Machine Learning and Applications, ICMLA Nassau, Bahamas, December 12-14, 2022*. IEEE, 2022, pp. 69–76.
- [30] —, "Deeper neural networks with non-vanishing logistic hidden units: Nova vs. relu neurons," in *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 2021, pp. 1407–1412.
- [31] V. Dumoulin and F. Visin, "A guide to convolution arithmetic for deep learning," *arXiv preprint arXiv:1603.07285*, 2016.
- [32] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [33] L. N. Smith and N. Topin, "Super-convergence: Very fast training of neural networks using large learning rates," in *Artificial intelligence and machine learning for multi-domain operations applications*, vol. 11006. SPIE, 2019, pp. 369–386.
- [34] C. Chadebec, L. J. Vincent, and S. Allasonnière, "Pythae: Unifying generative autoencoders in python - A benchmarking use case," in *NeurIPS*, 2022.