

Improved Auto-Encoding using Deterministic Projected Belief Networks and Compound Activation Functions

Paul M Baggenstoss

Fraunhofer FKIE

Fraunhoferstraße 20, 53343 Wachtberg, Germany

p.m.baggenstoss@ieee.org

Abstract—In this paper, we exploit the unique properties of a deterministic projected belief network (D-PBN) to take full advantage of trainable compound activation functions (TCAs). A D-PBN is a type of auto-encoder that operates by “backing up” through a feed-forward neural network. TCAs are activation functions with complex monotonic-increasing shapes that change the distribution of the data so that the linear transformation that follows is more effective. Because a D-PBN operates by “backing up”, the TCAs are inverted in the reconstruction process, restoring the original distribution of the data, thus taking advantage of a given TCA in both analysis and reconstruction. In this paper, we show that a D-PBN auto-encoder with TCAs can significantly out-perform standard auto-encoders including variational auto-encoders.

I. INTRODUCTION

A. Motivation

Auto-encoders are an important type of generative network used in current tasks such as unsupervised learning and open set classification [1], and dimension-reduction is the main task involved. It has been understood since the advent of support vector machines that non-linear operations can make the classification problem linearly separable [2], [3]. In the context of auto-encoding, non-linear operations can make the distribution representable in a lower-dimensional linear subspace. It is no surprise then, that the analysis task in auto-encoding is implemented by alternating applications of non-linear transformation (i.e. activation functions), followed by linear projection onto a lower-dimensional space. While this concept is clear for the analysis task, it is fair to ask the question: how can we construct an auto-encoder with non-linear transformations in such a way that the return path (synthesis task) is more effective? If non-linear transformations could be inverted during synthesis, then the data distribution at the input would be restored. The idea of inverting the non-linear transformation seems good at first, because the data distribution is changed in the forward path, and inverting it restores the original distribution. But this makes less sense in an auto-encoder with separate analysis and synthesis networks. But, a new type of auto-encoder based on the projected belief network (PBN) has recently been proposed, in which the

synthesis task is implemented by backing up through the analysis network [4]–[7]. Therefore, inverting the non-linear transformation is inherently part of the synthesis task. We seek to exploit properties of the PBN together with activation functions having complex shapes, to further improve the performance of an autoencoder.

B. Main Concept

A projected belief network (PBN) is unique among layered generative networks because instead of using an explicit generative scheme, it operates implicitly by backing-up through a feed-forward neural network (FFNN). The FFNN serves as the optimal feature extractor [4], and is the dual network for the generative process [7]. A PBN can also be sampled deterministically by choosing the conditional mean [4], resulting in the deterministic PBN (D-PBN). By passing forward through the FFNN, then backing up using D-PBN, a type of auto-encoder is formed [4]–[7]. By working forward, and then backward, an auto-encoder based on D-PBN is in a position to take full advantage of trainable compound activation functions (TCAs), which are activation functions with complex monotonic-increasing shapes. Like any non-linearity, TCAs can change the distribution of the data so that the linear transformation that follows is more effective [8]. By inverting the TCAs in the sampling process (synthesis task), the original distribution of the data is restored, realizing the advantage of a TCA in both analysis and reconstruction.

C. Previous Work

There is existing work on parameterizing activation functions, but these are either of low complexity, or are not necessarily invertible [9], [10]. Work on invertible activation functions exists in conjunction with normalizing flows (NF) [11], which is a method related to PBNs that has recently gained popularity. In NF, the network consists of a series of bijective (1:1) transformations, and the final output is usually regarded as being a set of independent identically-distributed (*iid*) random variables (RVs). The data generation process, is therefore the reverse of this: a set of *iid* RVs are subjected to the inverse series of transformations. Activation functions with complex shapes based on splines have been proposed for NF [12], [13]. In NFs, however, there is no

This work was supported jointly by the Office of Naval Research Global and the Defense Advanced Research Projects Agency under Research Grant - N62909-21-1-2024

dimension reduction, so this body of work does not encompass auto-encoders. In order to get around the requirement of dimension-preserving transformations, a way to combine NF with variational autoencoders (VAEs) has been described, [14], however this is a re-invention of PDF projection.

D. Motivation to use TCAs

TCAs are activation functions with complex monotonic increasing shapes that can make the linear transformation to a lower-dimensional space more effective, reducing the dimension required to represent the information. This is analogous to the motivation for non-linear projections in support vector machines (SVMs), making data linearly separable [2], [3]. A TCA can transform data with modal clusters into data with a uni-modal distribution if the activation function approximates the cumulative distribution of the input data [8]. Therefore, the first derivative of the transfer function should approximate the distribution of the input RV. The usefulness of this property is obvious, especially when the operation is inverted, allowing synthesizing multi-modal data from uni-modal data.

E. Novelty and Contributions

Both D-PBN and TCA have been previously separately introduced. The main novelty and contribution of this paper is the combination of D-PBN with TCA. The advantage of using TCA in D-PBN is that the TCA is inverted in the data synthesis process, thus restoring the original data distribution, even if it is multi-modal. Because D-PBN is the only kind of auto-encoder based on back-projection, it is the only auto-encoder where a given TCA can be used to advantage in both the analysis and reconstruction processes.

II. MATHEMATICAL APPROACH

A. Review of PDF Projection

The PBN and D-PBN are based on PDF projection, which we now review. Subject to mild constraints, any fixed dimension-reducing transformation, $\mathbf{y} = T(\mathbf{x})$, together with the known or assumed feature distribution $g(\mathbf{y})$, corresponds to a probability density function (PDF) on the input data [15]–[17] given by $G(\mathbf{x}) = \frac{p_{0,x}(\mathbf{x}) p_{0,y}(\mathbf{y})}{p_{0,x}(\mathbf{y})} g(\mathbf{y})$, where $p_{0,x}(\mathbf{x})$ is a prior distribution and $p_{0,x}(\mathbf{y})$ is its mapping to \mathbf{y} . In the notation “g” represents the *given* feature distribution, “G” is its projection to the higher-dimensional range of \mathbf{x} , and subscript “0” is a reminder that $p_{0,x}(\mathbf{x})$ can be seen as a *reference distribution*. In our simplified notation, the argument of a distribution defines its region of support, while the subscript defines the space where the prior distribution was originally applied. Therefore $p_{0,x}(\mathbf{y})$ has region of support on the range of \mathbf{y} , but is derived from the reference distribution on \mathbf{x} . On the other hand, $p_{0,y}(\mathbf{y})$, which we will introduce below, is a new reference distribution defined for \mathbf{y} . Note that if $p_{0,x}(\mathbf{x})$ is selected for maximum entropy, then $G(\mathbf{x})$ is unique.

The PDF of the input data can be estimated by training the parameters of the transformation to maximize the mean of $\log G(\mathbf{x})$, resulting in a transformation that extracts sufficient statistics and maximizes information [18]. And, when $g(\mathbf{y})$ is

fixed, then the true distribution under the available training data will be driven toward $g(\mathbf{y})$ [18]. We say that $G(\mathbf{x})$ is the “projection” of $g(\mathbf{y})$ back to the input data.

To draw a sample of $G(\mathbf{x})$, we first draw a sample \mathbf{y} from $g(\mathbf{y})$, then draw \mathbf{x} randomly from the set $\{\mathbf{x} : T(\mathbf{x}) = \mathbf{y}\}$, weighted by $p_{0,x}(\mathbf{x})$. PDF projection admits a chain-rule by applying the idea recursively to stages of a transformation. Consider the cascade of two transformations $\mathbf{y} = T_1(\mathbf{x})$, and $\mathbf{z} = T_2(\mathbf{y})$. Applying the above idea recursively, we have $G(\mathbf{x}) = \frac{p_{0,x}(\mathbf{x}) p_{0,y}(\mathbf{y})}{p_{0,x}(\mathbf{y}) p_{0,y}(\mathbf{z})} g(\mathbf{z})$. Note that $p_{0,x}(\mathbf{y})$ and $p_{0,y}(\mathbf{z})$ are two different distributions with support on the range of \mathbf{y} . While $p_{0,x}(\mathbf{y}) = T_1[p_{0,x}(\mathbf{x})]$, $p_{0,y}(\mathbf{z})$ is a canonical MaxEnt prior distribution for \mathbf{y} . The sampling process is also recursive, working backwards through the cascade of transformations.

B. Review of D-PBN

When PDF projection is applied layer by layer to a feed-forward neural network (FFNN), the result is a projected belief network (PBN) [7]. The PBN is stochastically sampled by “backing up” through the FFNN, but can be deterministically sampled by choosing the conditional mean estimate in each layer, resulting in a deterministic PBN (D-PBN). In any given layer of a FFNN, the hidden variable $\mathbf{y} \in \mathbb{R}^M$ is computed from the layer input $\mathbf{x} \in \mathbb{R}^N$, where we assume $N > M$, i.e. a dimension-reducing layer. Let $\mathbf{y} = \lambda(\mathbf{b} + \mathbf{W}'\mathbf{x})$, where $\lambda(\cdot)$ is the element-wise activation function, assumed to be strictly monotonic increasing (SMI). In the D-PBN, we seek to reconstruct \mathbf{x} from \mathbf{y} . Because the SMI activation function and bias are invertible, we prefer to work with the output of the linear transformation, denoted by $\mathbf{z} = \mathbf{W}'\mathbf{x}$. Specifically, \mathbf{x} is drawn randomly from the manifold $\mathcal{M}(\mathbf{z}) = \{\mathbf{x} : \mathbf{W}'\mathbf{x} = \mathbf{z}\}$, i.e. we draw \mathbf{x} from the set of samples that map to \mathbf{z} . For the PBN, we draw \mathbf{x} randomly from $\mathcal{M}(\mathbf{z})$ proportional to a prior distribution $p_{0,x}(\mathbf{x})$, and for the D-PBN, we select the conditional mean. We select $p_{0,x}(\mathbf{x})$ according to the principle of maximum entropy [19] (MaxEnt). Accordingly, $p_{0,x}(\mathbf{x})$ has the highest entropy subject to any constraints which include the range of the variable \mathbf{x} , denoted by \mathbb{X}^N , determined by which activation function was used in the up-stream layer. We consider three canonical ranges, $\mathbb{R}^N : x_i \in (-\infty, \infty), \forall i$, $\mathbb{P}^N : x_i \in [0, \infty), \forall i$, and $\mathbb{U}^N : x_i \in [0, 1], \forall i$. For a MaxEnt distribution to exist in the first two cases, we need constraints on the variance¹.

The canonical MaxEnt priors for the ranges \mathbb{R}^N , \mathbb{P}^N , and \mathbb{U}^N , are Gaussian, truncated Gaussian (TG), uniform, respectively, which are listed in Table I. The uniform distribution is a special case of the truncated exponential distribution (TED). Additional mathematical details can be found in the references [4], [5]. As estimate of \mathbf{x} , we use the conditional mean $\hat{\mathbf{x}} = \mathbb{E}_0(\mathbf{x}|\mathbf{z})$, i.e. the conditional mean under prior $p_{0,x}(\mathbf{x})$. In [4], we show that this is given by $\hat{\mathbf{x}} = \mathbb{E}_0(\mathbf{x}|\mathbf{z}) = \lambda(\mathbf{W}\mathbf{h})$, where $\lambda(\cdot)$ is a special MaxEnt activation function corresponding to

¹In the case of \mathbb{P}^N , a constraint on the mean is adequate, but this leads to an uninteresting solution.

\mathbb{X}^N	MaxEnt Prior $p_{0,x}(\mathbf{x})$	$\lambda(\alpha)$
\mathbb{R}^N	$\prod_{i=1}^N \mathcal{N}(x_i)$ (Gaussian)	α (Linear)
\mathbb{P}^N	$\prod_{i=1}^N 2\mathcal{N}(x_i), 0 < x_i$ (TG)	$\alpha + \frac{\mathcal{N}(\alpha)}{\Phi(\alpha)}$ (TG)
\mathbb{U}^N	1, $0 < x_i < 1$ (Uniform)	$\frac{e^\alpha}{e^\alpha - 1} - \frac{1}{\alpha}$ (TED)

TABLE I

MAXENT PRIORS AND ACTIVATION FUNCTIONS AS A FUNCTION OF INPUT DATA RANGE. TG="TRUNC. GAUSS.". TED="TRUNC. EXPON. DISTR".

$$\mathcal{N}(x) \triangleq \frac{e^{-x^2/2}}{\sqrt{2\pi}}, \text{ AND } \Phi(x) \triangleq \int_{-\infty}^x \mathcal{N}(x).$$

the MaxEnt prior (listed in Table I), and \mathbf{h} is the saddle point, which solves the equation

$$\mathbf{W}'\lambda(\mathbf{W}\mathbf{h}) = \mathbf{z}. \quad (1)$$

Sampling a D-PBN, therefore consists of the following steps: (a) Assume that hidden variable \mathbf{y} is given. (b) Invert the activation function and bias to obtain the linear transformation output variable \mathbf{z} . (c) Solve (1) to obtain \mathbf{h} . (d) Let $\hat{\mathbf{x}} = \lambda(\mathbf{W}\mathbf{h})$. (e) Now set the variable \mathbf{y} at the output of the previous layer to $\hat{\mathbf{x}}$, and repeat for the next layer (working backwards). As an aside, it was pointed out in [4], [6] that the inversion of the activation function in step 2 can be avoided if the FFNN uses MaxEnt activations in all layers (but not applicable when if TCAs are used).

D-PBN as Auto-Encoder. The D-PBN forms an auto-encoder if the starting feature is obtained by passing input data through the FFNN. In each layer, solving (1) presents a computational challenge due to the need to repeatedly invert matrices of the size $M \times M$, where M is the output dimension of a layer. The problem can be avoided or mitigated by various methods [5], some of which will be employed in the experiments.

Sampling failure. The solution to (1) is not guaranteed to exist unless $\mathbf{z} = \mathbf{W}'\mathbf{x}$ for some \mathbf{x} in the range of \mathbf{X} , denoted by \mathbb{X}^N [4]. When backing up through a FFNN, this condition is only strictly met for the last layer (first layer when working backwards), meaning that some input samples cannot be auto-encoded. Luckily, the *sampling efficiency*, which is the fraction of samples that succeed, can be driven to 1.0 with proper initialization and training [4], [6], [20]. Interestingly, a network can be initialized as a stacked RBM to have a sampling efficiency of 1 using the up-down algorithm, a multi-layer version of the contrastive divergence (CD) algorithm for fine-tuning stacked RBMs [21]. To use TCAs in the initial stacked RBM, the TCAs are initialized to a neutral state so that operates similar to a simple activation function (the base activation function, See Section II-C), and simple activation functions are used in the return (synthesis) path.

C. Review of TCAs

In Section I-D, we provided a motivation for TCAs. We now define the TCA precicely. Consider the compound activation function $f(x)$ given by

$$f(x) = \frac{\sum_{k=1}^K e^{a_k} f_k(e^{w_k} x + b_k)}{\sum_{k=1}^K e^{a_k}}, \quad (2)$$

where the functions $f_k(x)$ are simple activation functions, and $\mathbf{w} = \{w_k\}$, $\mathbf{a} = \{a_k\}$ and $\mathbf{b} = \{b_k\}$ are scale, weights, and bias parameters, respectively. The use of exponential functions is to insure positivity of the weights and scale factors. Note that if $f_k(x)$ are SMI, then the TCA is SMI. As a convention, all activation functions except the first activation function $f_1(x)$ are sigmoid-like activations (see TED activation in Table I). Because these are limited to the range $[0, 1]$, the ultimate data range of the TCA depends only on $f_1(x)$ which is called the *base activation*. The TCA behaves like the base activation, but with $K - 1$ "wiggles". The form of (2) is motivated by approximation of distributions with mixtures [8]. More to the point, (2) has a first derivative closely approximated by a Gaussian mixture [8], which is know to be a universal approximator of distributions [22]. It is the first derivative that is propagated as PDF estimate in a PBN [18].

D. D-PBN + TCA

We propose to use TCAs in the FFNN forward path, and TCA inversion in the D-PBN reconstruction path. A diagram of a D-PBN auto-encoder employing TCAs is shown in Figure 1. No explicit bias is needed because it is included in the

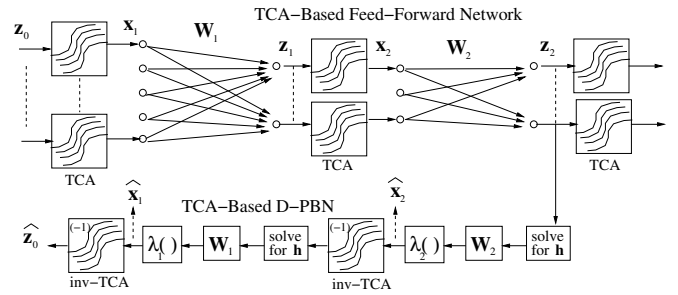


Fig. 1. Block diagram of a 2-layer D-PBN auto-encoder employing TCAs. Layer number is shown as subscript.

TCAs. Each D-PBN layer must choose one of the canonical data ranges, MaxEnt priors, and activation functions (Table III), which in the case of the TCA is identified with the base activation. Inversion of the TCA (i.e. find the x for which $f(x) = y$) is done by bisection search, then a few Newton-Raphson iterations. In a D-PBN, the last TCA in the forward path by itself does nothing, because it gets immediately inverted in the sampling process.

III. EXPERIMENT 1

A. Data Set and Goal

The the first data set was a subset of the MNIST handwritten data corpus, just three characters "3", "8", and "9". The data consists of sample images of 28×28 , or a data dimension of 784. We used 500 training samples from each character, for a total of 1500 samples. Since MNIST pixel data is coarsely quantized in the range $[0,1]$, a dither was applied to the pixel values². After applying dither, the data was "gaussianified",

²For pixel values above 0.5, a small exponential-distributed random value was subtracted, but for pixel values below 0.5, a similar random value was added.

Algorithm	Layers	MSE (train)	MSE (test)
AEC	3	3.05	3.15
AEC	5	2.86	3.22
VAE	3	3.87	4.18
UPDN	3	3.41	3.46
D-PBN	3	3.42	3.45
D-PBN/TCA	3	2.84	2.96

TABLE II
MEAN SQUARED ERROR (MSE) FOR EXPERIMENT 1.

i.e. transformed from range $[0, 1]$ to the real line by applying the logit (inverse sigmoid) function. The resulting pixel values had a smooth amplitude distribution in the range -10 and 10 . Using gaussianified data made the use of mean-squared error more meaningful. To facilitate duplicating the experimental results, the data we used is publicly available [20]. The goal is to auto-encode into a 16-dimensional bottleneck feature, then reconstruct the data with minimum mean-squared error. Training was unsupervised (no label information used).

B. Network

The network had three dense layers with 64, 32, and 16 nodes, respectively. Both TCA and non-TCA networks were tested. For the outputs of the first and second layer, we used 3-component TCAs. The base activation function is the TG (see Table I), which is very similar to softplus. We employed a 2-component TCA between the input data and first layer which used the linear base activation. A second, deeper network with 5 layers of 64, 48, 32, 24, and 16 nodes was also tested for the standard auto-encoder to dispell the notion that the TCA is nothing more than a type of additional layer.

C. Algorithms and Training

The following algorithms were tested: **AEC**: Standard auto-encoder with TG activation in all layers and untied weights (decoder network is the transpose of the analysis network, but with independent weights). **VAE**: Variational auto-encoder. The VAE has a special output layer, so the final TG activation is not used. **UPDN**: Stacked-RBM fine-tuned using the up-down algorithm [21]. **D-PBN**: Deterministic PBN using TG activation functions. **D-PBN/TCA**: Deterministic PBN using TCAs with TG base activation functions employing 3 components (base + 2). For data augmentation we used random shifts of a maximum of 1 pixel in both vertical and horizontal directions. Shifts were circular and calculated using the 2D FFT to attain vernier translation.

D. Results

As can be seen in Table II, the performance of D-PBN/TCA was superior to all other algorithms. The closest competitor for MSE on testing data was the 3-layer AEC, which due to the independent decoder network, has significantly more parameters. We also see with the 5-layer network that the additional layers do not help, but rather worsen overtraining.

IV. EXPERIMENT 2

A. Data set and Goal

Because MNIST data is widely seen to be over-simplified, we conducted a second experiment using more realistic data consisting of Google speech commands [23]. We trained unsupervised on two difficult to distinguish words: “three, and “tree”, sampled at 16 kHz and segmented into into 48 ms Hanning-weighted windows shifted by 16 ms. We used log-MEL band energy features with 20 MEL-spaced frequency bands and 45 time steps, representing a frequency span of 8 kHz and a time span of 0.72 seconds. The input dimension was therefore $N = 45 \times 20 = 900$. The number of training samples were 834, and 840, and testing samples were 430, and 449 for the two classes, respectively. The networks were trained unsupervised, taking all the data at once. The goal is to auto-encode into a 32-dimensional feature, then reconstruct the data with minimum mean-squared error.

B. Network and Training

We used a network with two convolutional layers, followed by two dense layers of 64 and 32 neurons. To mitigate the high computational load of D-PBN, the first 2 layers (convolutional) used linear activation, forming a Gaussian group, resulting in a much more efficient implementation of D-PBN [5]. A TCA with 2 components and linear base activation was applied to the input data, TCAs with 3 components and TG base activation were used before each dense layer. For all non-D-PBN algorithms, the same network structure was used with TG activation (similar to softplus) in all layers.

The PBN network (with or without TCA) was trained using the Theano framework with PBN Toolkit [20] with a mini-batch size of 288 samples and weight-decay. No data augmentation was used. Training a PBN with TCA requires using a very small learning rate due to the weak and highly non-linear dependence on the parameters of the TCA. Training took a week using a Quadro-P6000 GPU.

To give AEC the best chance, we asked a colleague to develop a Tensorflow/Keras-based implementation with no constrains (gloves off). The best-performing network, listed as “Unconstrained” in the results, had 21 layers and used ReLu activation and batch normalization (our networks did not). Data and D-PBN code is available for reader verification at [20].

C. Results

Results of Experiment 2 are provided in Table III. The validated mean squared error for D-PBN/TCA is significantly lower than all other algorithms. We provide examples of reconstructed spectrograms in Figure 2 in which it is seen that D-PBN/TCA visibly preserves the details better than the best of the other algorithms. We have outlined details in the Figure that show notable improvement. Details for DPBN/TCA (bottom row) are closer to the original (top row) that AEC (middle row).

Algorithm	L2 reg	MSE (Train)	MSE (Test)
AEC/Tied (*)	1e-3	2.500	3.093(**)
AEC/Untied (*)	1e-3	2.405	3.0558(**)
Unconstrained	0	2.38	3.02
PBN	0	2.777	3.030
PBN/TCA	0	2.4102	2.8092

TABLE III

MEAN SQUARED RECONSTRUCTION ERROR (MSE) FOR VARIOUS ALGORITHMS. (*) TIED: THE ANALYSIS AND RECONSTRUCTION WEIGHTS ARE THE SAME EXCEPT FOR A TRAINABLE SCALE FACTOR. UNTIED: WEIGHTS ARE INDEPENDENT. (**) OPTIMIZED OVER L2 REGULARIZATION FACTOR.




Fig. 2. Comparison of reconstructions for testing data. Top row: original spectrograms. Center: AEC (Untied). Bottom: PBN/TCA. Examples of areas where PBN/TCA shows better detail are circled.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we demonstrated that TCAs and D-PBN complement each other so that together, they provide a significant advantage in validated mean-squared error over other types of auto-encoders. Future work is needed to make the training more efficient. The weak and highly non-linear dependence on the TCA parameters might be mitigated by re-parameterizing the TCAs. Data, software, and parameters to recreate these results are provided at [20].

REFERENCES

- [1] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep Learning*. Cambridge, MA: MIT press, 2016.
- [2] D. Sebald, "Support vector machines and the multiple hypothesis test problem," *IEEE Trans. Signal Processing*, vol. 49, pp. 2865–2872, November 2001.
- [3] T. Joachims, "Making large-scale SVM learning practical," in *Advances in Kernel Methods - Support Vector Learning* (B. Schölkopf, C. Burges, and A. Smola, eds.), ch. 11, pp. 169–184, Cambridge, MA: MIT Press, 1999.
- [4] P. M. Baggenstoss, "A neural network based on first principles," in *ICASSP 2020, Barcelona (virtual)*, (Barcelona, Spain), Sep 2020.
- [5] P. M. Baggenstoss, "Using the projected belief network at high dimensions," *Proceedings of EUSIPCO 2022, Belgrade*, 2022.
- [6] P. M. Baggenstoss, "Applications of projected belief networks (PBN)," *Proceedings of EUSIPCO, A Coruña, Spain*, 2019.
- [7] P. M. Baggenstoss, "On the duality between belief networks and feed-forward neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–11, 2018.
- [8] P. M. Baggenstoss, "Trainable compound activation functions for machine learning," *Proceedings of EUSIPCO 2022, Belgrade*, 2022.
- [9] A. Molina, P. Schramowski, and K. Kersting, "Padé activation units: End-to-end learning of flexible activation functions in deep networks," 2020.

- [10] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," 2015.
- [11] I. Kobyzev, S. J. Prince, and M. A. Brubaker, "Normalizing flows: An introduction and review of current methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 11, pp. 3964–3979, 2021.
- [12] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, "Cubic-spline flows," in *Workshop on Invertible Neural Networks and Normalizing Flows, ICML*, 2019.
- [13] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, "Neural spline flows," in *Advances in Neural Information Processing Systems*, 2019.
- [14] D. Nielsen, P. Jaiji, E. Hooeboom, O. Winther, and M. Welling, "Survae flows: Surjections to bridge the gap between vaes and flows," in *Advances in Neural Information Processing Systems* (H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, eds.), vol. 33, pp. 12685–12696, Curran Associates, Inc., 2020.
- [15] P. M. Baggenstoss, "The PDF projection theorem and the class-specific method," *IEEE Trans Signal Processing*, pp. 672–685, March 2003.
- [16] P. M. Baggenstoss, "Beyond moments: Extending the maximum entropy principle to feature distribution constraints," *Entropy*, vol. 20, no. 9, 2018.
- [17] P. M. Baggenstoss, "Maximum entropy PDF design using feature density constraints: Applications in signal processing," *IEEE Trans. Signal Processing*, vol. 63, June 2015.
- [18] P. M. Baggenstoss and S. Kay, "Nonlinear dimension reduction by pdf estimation," *IEEE Transactions on Signal Processing*, 2022.
- [19] E. T. Jaynes, "Information theory and statistical mechanics i," *Physical Review*, p. 171–190, 1957.
- [20] P. Baggenstoss, "PBN Toolkit," {<http://class-specific.com/pbntk>}. Accessed: 2022-02-28.
- [21] G. E. Hinton, S. Osindero, and Y.-W. Teh, "A fast learning algorithm for deep belief nets," in *Neural Computation 2006*, 2006.
- [22] R. A. Redner and H. F. Walker, "Mixture densities maximum likelihood, and the EM algorithm," *SIAM Review*, vol. 26, April 1984.
- [23] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," *arXiv:1804.03209*, 2018.