

Bayesian optimization for sparse neural networks with trainable activation functions

Mohamed Fakhfakh^{1,2} and Lotfi Chaari²

¹ University of Sfax, MIRACL, Tunisia
mohamed.fakhfakh@toulouse-inp.fr

² Toulouse INP, IRIT, University of Toulouse, France
lotfi.chaari@toulouse-inp.fr

Abstract—In the literature on deep neural networks, there is considerable interest in developing activation functions that can enhance neural network performance. In recent years, there has been renewed scientific interest in proposing activation functions that can be trained throughout the learning process, as they appear to improve network performance, especially by reducing overfitting. In this paper, we propose a trainable activation function whose parameters need to be estimated. A fully Bayesian model is developed to automatically estimate from the learning data both the model weights and activation function parameters. An MCMC-based optimization scheme is developed to build the inference. The proposed method aims to solve the aforementioned problems and improve convergence time by using an efficient sampling scheme that guarantees convergence to the global maximum. The proposed scheme has been tested across a diverse datasets, encompassing both classification and regression tasks, and implemented in various CNN architectures to demonstrate its versatility and effectiveness. Promising results demonstrate the usefulness of our proposed approach in improving models accuracy due to the proposed activation function and Bayesian estimation of the parameters.

Index Terms—Activation function, Deep neural networks, Optimization, MCMC, Hamiltonian dynamics

I. INTRODUCTION

Classification is a machine-learning task that identifies which objects are present in an image or video. It is critical for various applications such as computer vision [1]–[3], medical diagnostics [4], signal processing [5], and others. The process involves learning significant or nontrivial relationships from a set of training data and extending these relationships to interpret new test data [6]. The task involves categorizing elements into one of the finite set of classes by comparing the measured attributes of a given object with the known properties of objects to determine whether the object belongs to a specific category.

Convolutional Neural Networks (CNNs) [7]–[11] have become the industry standard in numerous applications over the past two decades, as they can process complex high-dimensional input data into simple low-dimensional concepts through a series of nonlinear transformations. Each feature layer in CNNs comprises features from the layer below, creating a hierarchical organization of ever-more-abstract concepts. They are particularly effective at capturing high-level abstractions in real-world observations, making them a popular choice for image classification tasks. CNNs can learn features from raw image pixels, reducing the need for manual feature engineering. This makes them well-suited for tasks such as

object recognition, where the goal is to identify the presence and location of specific objects in an image. In this context, the activation function plays a critical role in learning representative features. The Rectified Linear Unit (ReLU) [12] is currently the most widely used activation function for neural networks. When provided with a positive argument, the ReLU activation function is either zero or the identity. ReLUs have the additional benefit of alleviating the vanishing gradient problem, in addition to providing sparse codes [13]. According to the state of the art, activation functions can be fixed or trainable during a learning phase [14]. In most cases, gradient descent is the most widely used method in the literature for parameter estimation.

From another side, Bayesian methods have advanced significantly in many domains over time and have numerous useful applications. The primary idea is to represent all uncertainties in the model using probabilities. Bayesian techniques are distinctive in that they treat the problem as an inference problem [15]. One of the most significant advantages is the ability to incorporate prior information about the model parameters and hyperparameters. Recent advancements in Markov Chain Monte Carlo (MCMC) methods [16]–[19] make it easier to use Bayesian analyses in complex datasets with missing observations and to handle multidimensional outcomes. Recent studies have demonstrated that using a Bayesian framework in CNNs for the optimization process leads to more promising performances than standard gradient descent [20], [21].

In this study, we introduce a new trainable activation function. The parameters of the proposed function are automatically estimated from the data. For doing so, a Bayesian framework is used where these parameters as well as the network weights are assumed to be realizations of random variables. With an adequate likelihood, a hierarchical Bayesian model is built with priors and hyperpriors. An MCMC-based inference is then used, specifically a Gibbs sampler, to derive estimators from the target distributions. Our method is an extension of our previous work described in [22], where we employed non-smooth Hamiltonian methods to fit sparse artificial neural networks. Our main objective with the Bayesian scheme is to minimize the target cost function of the learning model. The use of non-smooth Hamiltonian techniques enables us to perform efficient and fast sampling, even when dealing with non-differentiable energy functions that arise due to the use of sparse regularization functions.

The contribution of this paper is therefore twofold: *i*) proposing a new trainable activation function, and *ii*) a more general

Bayesian formulation than in [22] to integrate the estimation of all parameters from the data.

The rest of this paper is organized as follows. The state of the art is the focus of Section II. Then, the Problem statement is in Section III. In section IV we detail the adopted hierarchical Bayesian model. The proposed Bayesian inference scheme is developed in Section V and validated in Section VI. Finally, The conclusion and future work are drawn in Section VII.

II. RELATED WORK

Finding the best activation function to integrate into an architecture is a challenging task. This section proposes a taxonomy of different activation functions described in the literature. A primary established classification is based on the ability to modify the shape of the activation function during the training phase, resulting in two major categories that can be distinguished [14].

A. Fixed-shape activation functions

This category pertains to the use of activation functions in neural network research, including sigmoid [23], hyperbolic tangent (tanh) [24], and ReLU, all of which have a defined form. The introduction of rectified functions, particularly ReLU, has led to a marked enhancement in neural network performance and heightened scientific interest. Consequently, this category can be subdivided into subcategories based on their distinctive characteristics.

Classic activation functions:

The findings presented in [25] demonstrated that a feed-forward and shallow network can efficiently manage any continuous function defined on a compact subset. For many years, bounded activation functions like Sigmoid and tanh were the preferred choices for neural networks, with researchers demonstrating their efficacy, particularly in shallow network architectures [26]. Although Sigmoid, Bipolar sigmoid [27], Hyperbolic tangent, Absolute value [28], and other bounded activation functions are commonly used, their efficacy is limited when training multi-layer neural networks due to the vanishing gradient problem [29].

Rectifier-based activation functions:

The primary advantage of using rectified activation functions is to alleviate the problem of vanishing gradient. The success of ReLU [30] has inspired the development of many new activation functions during the last years [31], [32]. While ReLU has numerous benefits such as solving the vanishing gradient issue and making sparse coding easier [33], it is not without flaws. The "dying" ReLU problem [34] and non-differentiability at zero are the main concerns.

To address these issues, several variations of ReLU have been developed, such as Leaky ReLU (LReLU) [34], Truncated rectified [35], softplus [36], Exponential linear unit (ELU) [37], E-swish [38], and Flatten-T Swish [39].

B. Trainable activation functions

The concept of using trainable activation functions is not new in the field of neural network research, with many studies published on this topic as early as the 1990s [40]–[42]. However, the growing interest in neural networks in recent years has led researchers to reconsider the potential benefits of trainable activation functions in improving network performance.

In this section, we discuss the main strategies proposed in the literature for learning activation functions from data. Based on their primary characteristics, these strategies can be classified into three families.

Parameterized standard activation functions:

In [40], a generalized hyperbolic tangent function was proposed by introducing two trainable parameters α and β , which adjust the saturation level and slope, respectively. Similarly, a sigmoid function with two trainable parameters was used in [41] to modify the activation function's shape. These parameters are learned along with the network weights using the backpropagation algorithm. More recently, the work in [43] aimed to avoid manually setting the parameter of the ELU unit by proposing an alternative based on two trainable parameters. The proposed activation function, called PELU, is defined as follows:

$$PELU(x) = \begin{cases} \frac{\beta}{\gamma}x & \text{if } x \geq 0 \\ \beta \times (\exp(\frac{x}{\gamma}) - 1) & \text{otherwise,} \end{cases} \quad (1)$$

where β is a trainable parameter.

A flexible ReLU function has been proposed in [44]:

$$frelu(x) = ReLU(x + \alpha) + \beta, \quad (2)$$

where α and β are parameters learned from data. This is done to capture negative information that is lost in the classic ReLU function [12]. The activation function introduced in [45] is another type of ReLU function that partially learns its shape from the training set. Indeed, it can modify the negative part of the data via the parameter α . This function is called Parametric ReLU (PReLU) and can be defined as follows:

$$PReLU(x) = \begin{cases} x & \text{if } x > 0 \\ \alpha \times x & \text{otherwise} \end{cases} \quad (3)$$

where the parameter α is learned jointly with the model using a gradient method.

Functions based on ensemble methods:

Functions based on ensemble methods involve combining multiple basic activation functions to form a more complex function.

In [46], an innovative method using genetic programming to optimize neural network activation functions is presented. This technique differs from standard activation functions by exploring combinations and tree structures to engineer novel and potentially superior functions. The goal is to enhance neural network performance through task-specific and adaptive

activation functions, providing a compelling alternative to traditional approaches. In [47], a method for investigating activation functions built as compositions of various basic activation functions are proposed. Similarly, [48] introduces a similar method using a genetic algorithm composed of two new activation functions: Exponential Linear Sigmoid Squashing (EliSH) and HardELiSH. EliSH is defined as follows:

$$EliSH(x) = \begin{cases} x/(1 + e^{-x}) & \text{if } x \geq 0 \\ (e^x - 1)/(1 + e^{-x}) & \text{otherwise.} \end{cases} \quad (4)$$

The negative part of EliSH is a multiplication of two functions, ELU and Sigmoid, while the positive part is shared with Swish [49]. HardELiSH is defined as a multiplication of HardSigmoid and ELU in the negative portion and HardSigmoid and Linear in the positive part. In [50], another interesting activation function is proposed, called the Mexican Hat Linear Unit (MeLU). This activation function solves the problems of unstable learning related to trainable parameters. Unstable learning can cause a decrease in accuracy and an increase in generalization error when the model's performance varies significantly in response to slight changes in the data or parameters. Mexican hat-type functions have a smoother curve than ReLU, which prevents saturation and allows for optimal performance.

Let f be the function defined by

$$f_{\gamma,\lambda}(x) = \max(\lambda - |x - \gamma|, 0), \quad (5)$$

where λ, γ are real numbers. This function returns zero when $|x - \gamma| > \lambda$. Moreover, it increases with a derivative of 1 between $\gamma - \lambda$ and γ , then decreases using a derivative of -1 between γ and $\gamma + \lambda$. MeLU is defined for each layer as

$$MeLU(x) = PReLU(x) + \sum_{j=1}^{k-1} c_j f_{\gamma_j, \lambda_j}(x), \quad (6)$$

where k represents the number of parameters that can be learned in each neuron. The c_j parameters are learnable real numbers, while γ_j, λ_j are fixed parameters.

In [51], the authors introduced a new function called Adaptive Blending Units (ABU) as a trainable linear combination of a set of activation functions

$$ABU(x) = \sum_{i=1}^k \alpha_i \times f_i(x), \quad (7)$$

where $(\alpha_1, \alpha_2, \dots, \alpha_k)$ are parameters to be learned, and $\{f_1(\cdot), f_2(\cdot), \dots, f_k(\cdot)\}$ is a set of activation functions. The parameters α_i are all initialized to the value $\frac{1}{k}$ and are trained using the gradient descent method.

Recent innovations also encompass methods like Kernel-based activation function [52], [53] and Trained activation function [54], which further broaden the array of tools available for neural network customization.

Building upon these advances, the Trainable Compound Activation Functions (TCAs) introduced in [55] represent

an innovation in activation functions design. By integrating multiple basis functions, each with individual trainable parameters, TCAs offer an interesting adaptability and precision. This approach allows the activation function to dynamically tune to the probability density functions (PDFs) of the data distribution during training. Such alignment can significantly enhance the functionality of complex neural network structures like Restricted Boltzmann Machines (RBM), Deep Belief Networks (DBN), and Projected Belief Networks (PBN). In generative networks, TCAs may play a special role as they effectively estimate the marginal distributions of each data dimension using a mixture model. This reduces the modality of the data and makes linear dimension reduction techniques more effective, thus enhancing the network's ability to capture and represent complex data distributions.

Furthering this innovation, Baggenstoss' subsequent work [55] applied TCAs to Projected Belief Networks (PBN and D-PBN), elevating auto-encoders' capabilities. By implementing PDF projection, these models achieve sophisticated data representation and efficient dimensionality reduction. This represents a major improvement over traditional approaches by integrating adaptable network structures with advanced activation functions, setting new standards for neural network efficiency and performance.

Activation functions based on other techniques:

Beyond traditional and standard functions, researchers are exploring innovative techniques for designing activation functions, being inspired by diverse areas such as evolutionary computation and fuzzy logic.

In [56], the authors introduce a novel tool named EvoNN, which uses evolutionary computation to simultaneously optimize the weights and activation functions of hidden layers in neural networks. Unlike traditional methods, EvoNN provides users with the flexibility to add their own activation functions to the set of available functions, thereby enriching the "alphabet" of functions usable by the optimizer. Comparative tests with standard neural networks show that, although EvoNN requires a longer execution time, it delivers superior performance, which could justify its use in certain demanding applications. In [42], a method using polynomial functions with adjustable coefficients is proposed. Similarly, in the context of fuzzy activation functions, a neural unit based on Type-2 fuzzy logic [57] was developed. Other works have proposed functions using interpolation and spline approaches [58]. However, depending on the chosen technique, these strategies may require additional input.

C. Comparison and analysis

Given the advancements in machine learning, the optimization of activation functions remains a central topic of study. The correct choice of activation functions is crucial, as it can significantly enhance learning architecture performance. While traditional fixed activation functions like

ReLU or sigmoid have been standard, there is a growing trend towards trainable activation functions, which offer modifications in shape driven by trainable parameters. Empirical studies have shown that these trainable functions lead to marked improvements in performance compared to their fixed counterparts.

In this sense, TCAs emerge as a notable innovation, providing a fresh perspective on managing activation functions compared to conventional approaches. Despite their benefits, TCAs introduce complex challenges. They exhibit a delicate and highly non-linear dependence on parameter variations, complicating the tuning process and increasing the risk of overfitting. Future research focused on re-parameterization strategies could mitigate these issues, potentially leading to enhanced model generalization and performance.

From another side, the refinement of various neural network components becomes increasingly interconnected with the pursuit of automated model optimization strategies. The trend towards automated model configuration, like Auto-Net [59], reflects the growing necessity to minimize human intervention in increasingly complex neural network architectures. Indeed, the incorporation of Bayesian optimization and HyperBand (BOHB) within Auto-Net 2.0 represents a significant step towards self-optimized neural networks, potentially competitive with manually tuned models.

While the potential of advanced activation functions to improve neural network performance is significant, their implementation can be more complex and require more time and resources for learning.

Despite encouraging results, it is still challenging to identify a strategy for automatic learning of an activation function that would solve different problems and significantly improve performance. Most trainable activation functions use the gradient descent method for hyperparameter estimation. The main limitations of these techniques lie in the computation time and gradient vanishing. This process prevents the network from learning deep features and may even lead to excessive processing capacity during training.

Indeed, neural networks can get stuck in local minima [60], which can harm the model's performance. This phenomenon is partly due to the gradient vanishing that occurs as derivatives decrease as the model deepens. This gradient decrease makes it harder to optimize the model, leading to a decrease in performance.

III. PROBLEM STATEMENT

In the previous Section, we have examined different activation functions presented in the literature. In this paper, we introduce a modification of the MeLU activation function [61], while integrating the estimation of its parameters into a global Bayesian optimization framework. The choice of this function is mainly justified by its promising performance, as well as its form which promotes non-linearity and sparsity. However, the limits of the MeLU function are mainly related

to memory requirements. This limit indicates that using this function may require a larger amount of memory compared to other activation functions, which can affect computational efficiency.

As a first contribution of this paper, the Modified Mexican ReLU (MMeLU) activation function is proposed to solve the complexity problem and improve model performance. The specificity of this new activation function is that it requires fewer parameters to estimate than MeLU. The second contribution is related to integrating the parameters estimation of the MMeLU function into a Bayesian optimizer [62], rather than using a standard optimization procedure such as the ADAM optimizer.

To define MMeLU, let

$$f_{\gamma,b}(x) = \max(b - |x - \gamma|, 0), \quad (8)$$

where γ and b are real numbers.

Using $f_{\gamma,b}$, the proposed activation function MMeLU can be defined as

$$MMeLU(x) = c \times f_{\gamma,b}(x) + (1 - c) \times ReLU(x) \quad (9)$$

where c is a real number belonging to the interval $[0, 1]$. For the proposed MMeLU function, c , b , and γ are the parameters to be estimated.

The shape of the proposed MMeLU function, as well as those of competing functions ReLU, FReLU, PReLU, and MeLU, are illustrated in Figure 1. The activation function proposed in this paper is made up of a mixture of the ReLU and the Mexican hat functions. The curves in Figure 1[(a)-(d)] clearly show the flexibility and non-linearity of our MMeLU function with different configurations of the parameters γ , b , and c .

The Mexican hat function is often used as an activation function in neural networks due to its advantages. Mexican hat functions are continuous, which means that changes in inputs produce continuous changes in outputs. This is important in neural networks because it allows a gradual update of weights. Additionally, they have high representation capacity, which means that they can accurately model complex functions. This allows neural networks to learn non-linear relationships between inputs and outputs.

The Mexican hat function looks like a bell but with a peak in the center that makes it more pronounced for small values, as shown by the MMeLU curves in Figures 1[a] and [b]. This means that for small input values, the Mexican hat function will have a stronger response than other activation functions, such as ReLU or FReLU. This behavior can be useful in certain situations, for example when the input data has a restricted value range or when the neural network's response needs to be more sensitive to small variations in the input data.

Let us now consider the convolutional neural network in Figure 2. The MMeLU activation function is applied after each convolutional layer, replacing the *ReLU* function. This strategy has the potential to increase the non-linearity of

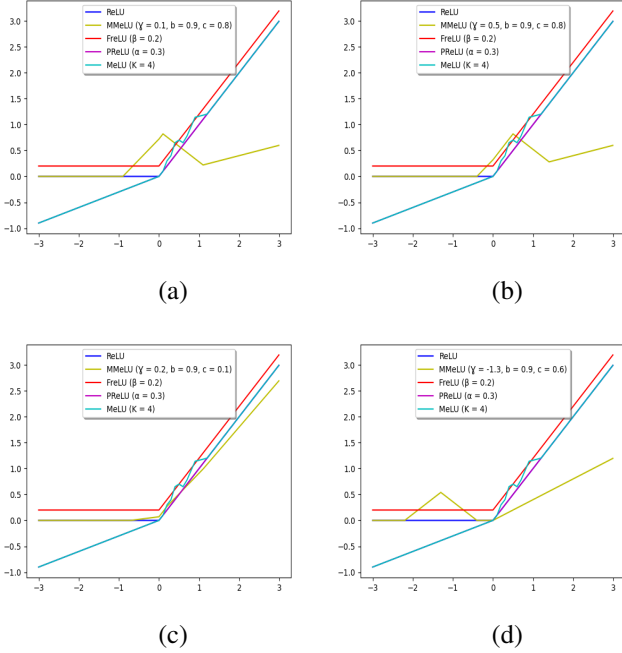


Fig. 1. Illustrations of ReLU, FreLU, PReLU, MeLU, and MMeLU curves with different configurations.

the model and improve its ability to represent features in images. It is worth noting that when c is estimated close to 0 ($c \sim 0$), the MMeLU function tends to the ReLU behavior.

Regarding the model fitting, let us assume that the estimated label (or numerical value) is obtained by applying the proposed activation function $MMeLU(\mathbf{x}, W)$, where \mathbf{x} is the input data and $W \in \mathbb{R}^N$ denotes the weights vector. The model parameters (weight vector and parameters of the activation function) can be determined during the training phase using a generic error function D (Euclidean, Minkowski, etc.). Akin to [21], the target CNN is assumed to be sparse. We also use the same Bayesian formulation of the optimization problem. For the M input data, we can write

$$\begin{aligned} \widehat{W} &= \arg \min_W \mathcal{L}(W) \\ &= \arg \min_{W, b, \gamma, c} \sum_{m=1}^M D(MMeLU(x^m; W) - y^m) + \sum_{l=1}^L \lambda_l \|W^l\|_1, \end{aligned} \quad (10)$$

where y^m is the ground truth for input data x^m , L is the number of layers in the network, and λ_l is a regularization parameter to be estimated for layer l that balances the solution between the data attachment term and the ℓ_1 sparse regularization terms.

In the following Section, we formulate the adopted hierarchical Bayesian model to conduct the inference and fit the model weights and the activation function parameters.

IV. HIERARCHICAL BAYESIAN MODEL

The problem of estimating the parameters of the MMeLU activation function is formulated within a Bayesian framework. In this sense, all parameters and hyperparameters are supposed to follow probability distributions. A likelihood is defined to model the relationship between the target weight vector, the activation function parameters, and the data. A prior distribution is defined to model the prior knowledge about the target weights and all activation function parameters.

A. Likelihood

Following the principle of minimizing the error between the reference vector \mathbf{y} (labels or continuous values) and its estimate $\widehat{\mathbf{y}}$, we define the likelihood distribution as

$$f(\mathbf{y}, \mathbf{x}; W, c, \gamma, b) \propto \prod_{m=1}^M \exp[-D; (MMeLU(x^m; W, c, \gamma, b) - y^m)]. \quad (11)$$

It is worth noting that when a Euclidean distance is used for D , the adopted likelihood is nothing but a Gaussian distribution.

B. Priors

In our model, the unknown parameters are grouped in the unknown vector $\theta = \{W, c, \gamma, b, \lambda\}$, where $\lambda = \{\lambda_1, \dots, \lambda_L\}$.

Prior for W :

To promote sparsity in the neural network, we use a Laplace distribution for the weight vector W akin to [21]:

$$f(W; \lambda) \propto \prod_{l=1}^L \prod_{k=1}^{K_l} \left[\frac{1}{\lambda_l} \exp\left(-\frac{|W_{kl}^l|}{\lambda_l}\right) \right], \quad (12)$$

where K_l is the number of weights in layer l of the network and λ_l is a parameter to be estimated.

Prior for λ_l :

Since $\lambda_l \in \mathbb{R}_+$, we chose to use an inverse gamma (IG) distribution:

$$f(\lambda_l; \delta, \mu) = IG(\lambda_l; \delta, \mu) \propto (\lambda_l)^{-1-\delta} \exp\left(-\frac{\mu}{\lambda_l}\right), \quad (13)$$

where δ and μ are positive parameters that were fixed at 10^{-3} to have a non-informative prior.

Prior for c :

Regarding the parameter c , we consider a uniform distribution over the interval $[0, 1]$, denoted as

$$c \sim U_{[0,1]}(c). \quad (14)$$

Prior for γ :

Since γ is a real value, a Gaussian distribution is used as follows

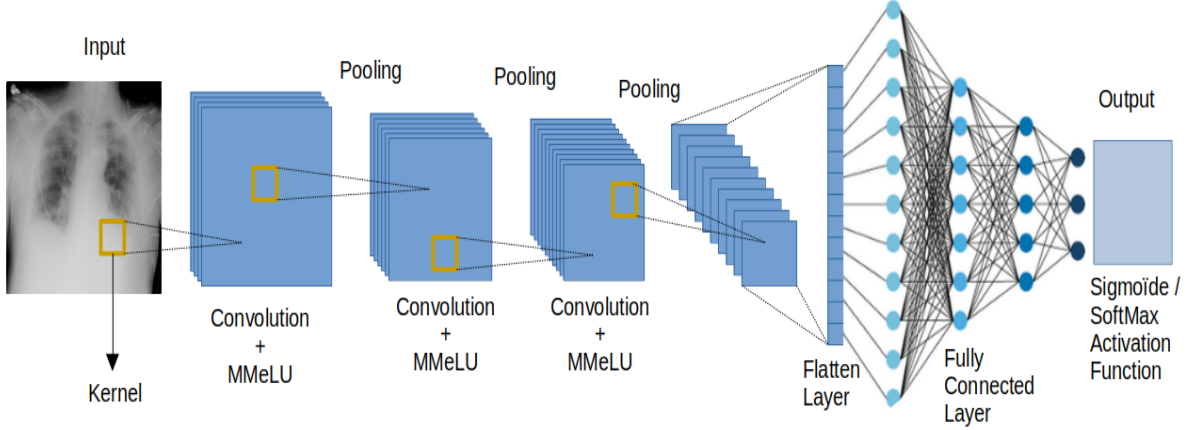


Fig. 2. A general diagram of a convolutional neural network with the MMELU activation function.

$$f(\gamma; \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma^2} \exp\left(-\frac{\gamma^2}{2\sigma^2}\right), \quad (15)$$

where σ^2 is a hyperparameter to be estimated.

Priori for b :

Since b is a positive real number, an exponential distribution is used as follows:

$$f(b; \lambda_b) \propto \begin{cases} \frac{1}{\lambda_b} \exp\left(-\frac{b}{\lambda_b}\right); & \text{if } b \geq 0 \\ 0; & \text{otherwise.} \end{cases} \quad (16)$$

where λ_b is a hyperparameter to be estimated. This prior penalizes large values of b .

C. Hyperpriors

Since λ_b and σ^2 are positive real numbers, an inverse gamma (IG) distribution was used as a hyper-*a priori*:

$$f(\lambda_b; \delta, \mu) = IG(\lambda_b; \delta, \mu) \propto (\lambda_b)^{-1-\delta} \exp\left(-\frac{\mu}{\lambda_b}\right) \quad (17)$$

and

$$f(\sigma^2; \delta, \mu) = IG(\sigma^2; \delta, \mu) \propto (\sigma^2)^{-1-\delta} \exp\left(-\frac{\mu}{\sigma^2}\right), \quad (18)$$

where δ and μ are positive parameters that were fixed at 10^{-3} .

V. INFERENCE SCHEME

By adopting a Maximum *a Posteriori* (MAP) approach, we first need to express the posterior distribution. Let Φ_e be the hyperparameters to be estimated, represented by $\Phi_e = \{\sigma^2, \lambda_b\}$, and Φ_m be the hyperparameters to be fixed, $\Phi_m = \{\delta, \mu\}$. Using the likelihood, the prior distributions, and the defined hyperpriors, we can write the posterior distribution as:

$$f(\theta, \Phi_e; y, \Phi_m) \propto f(y; \theta) f(\theta; \Phi_e) f(\Phi_e; \Phi_m) \quad (19)$$

which can be reformulated in a detailed version as

$$\begin{aligned} f(\theta, \Phi_e; \mathbf{y}, \mathbf{x}, \Phi_m) \propto & \prod_{m=1}^M \exp[-D(MMeLU(x^m; W, c, \gamma, b) - y^m)] \times \\ & \prod_{l=1}^L \frac{1}{\lambda_l^{K_l}} \prod_{k=1}^{K_l} \left[\exp\left(-\frac{|W_{k,l}^l|}{\lambda_l}\right) \right] \times (\lambda_l)^{-1-\delta} \exp\left(-\frac{\mu}{\lambda_l}\right) \times \\ & \exp\left(-\frac{\gamma^2}{2\sigma^2}\right) \times \frac{1}{\lambda_b} \exp\left(-\frac{b}{\lambda_b}\right) 1_{\mathbb{R}_+}(b) \times 1_{[0,1]}(c) \times \\ & (\lambda_b)^{-1-\delta} \exp\left(-\frac{\mu}{\lambda_b}\right) \times (\sigma^2)^{-1-\delta} \exp\left(-\frac{\mu}{\sigma^2}\right). \end{aligned} \quad (20)$$

One can clearly notice that the posterior in (20) is complicated to deal with in order to derive close-form estimators. We, therefore, resort to numerical approximations using a Markov Chain Monte Carlo technique (MCMC) [17], [63]. Specifically, we use a Gibbs sampler to sequentially sample according to the conditional posteriors. To calculate the conditional distributions associated with each parameter of the model, one needs to integrate the joint posterior distribution in (20) with respect to all the other parameters.

Regarding the parameter W , calculations based on (20) lead to the following form:

$$\begin{aligned} f(W; c, \gamma, b, \lambda) \propto & \exp\left[-\sum_{l=1}^L \sum_{k=1}^{K_l} \frac{|W_{k,l}^l|}{\lambda_l}\right] \times \\ & \exp\left[-\sum_{m=1}^M (D; (MMeLU(x^m; W, c, \gamma, b) - y^m))\right]. \end{aligned} \quad (21)$$

The conditional distribution for the parameter c is given by:

$$f(c; W, b, \gamma) \propto 1_{[0,1]}(c) \times \exp \left[- \sum_{m=1}^M (D; (MMeLU(x^m; W, c, \gamma, b) - y^m)) \right]. \quad (22)$$

For the parameter b , the condition distribution is given by:

$$f(b; W, c, \gamma, \lambda_b) \propto \exp \left(- \frac{b}{\lambda_b} \right) \times \exp \left[- \sum_{m=1}^M (D; (MMeLU(x^m; W, c, \gamma, b) - y^m)) \right]. \quad (23)$$

As regards γ , the conditional distribution writes:

$$f(\gamma; W, b, c, \sigma^2) \propto \exp \left(- \frac{\gamma^2}{2\sigma^2} \right) \times \exp \left[- \sum_{m=1}^M (D; (MMeLU(x^m; W, c, \gamma, b) - y^m)) \right]. \quad (24)$$

The conditional distribution for the parameter λ_l is given by:

$$f(\lambda_l; \delta, \mu) \propto \lambda_l^{-1-(\delta+K_l)} \exp \left(- \frac{\mu}{\lambda_l} \right) \propto IG(\delta + K_l, \mu). \quad (25)$$

For the hyperparameter vector Φ_e , it is necessary to calculate the conditional distributions from which it is possible to sample based on the likelihood and adopted priors.

The conditional distribution for the hyperparameter λ_b is given by:

$$f(\lambda_b; b, \mu, \delta) \propto \lambda_b^{-2-\delta} \exp \left(- \frac{b + \mu}{\lambda_b} \right) \propto IG(\delta + 1, b + \mu) \quad (26)$$

The conditional distribution for the hyperparameter σ^2 is given by:

$$f(\sigma^2; \mu, \gamma, \delta) \propto (\sigma^2)^{-1-\delta} \exp \left(- \frac{\gamma^2 + 2\mu}{2\sigma^2} \right) \propto IG(\delta, \gamma + 2\mu). \quad (27)$$

The sampling scheme is summarized in Algorithm 1, where the model weights W and the parameters of the proposed MMeLU function are sampled.

In Algorithm 1, S denotes the number of MCMC sampling iterations. After the burn-in period, the sampled coefficients are used to calculate the estimators \widehat{W} , \widehat{c} , \widehat{b} , $\widehat{\gamma}$, in addition to $\widehat{\sigma^2}$, $\widehat{\lambda}$ and $\widehat{\lambda}_b$.

VI. EXPERIMENTAL VALIDATION

To validate the proposed method, our experimental approach encompasses both classification and regression [64] tasks, leveraging a variety of datasets and application scenarios. We conduct three image classification experiments using distinct datasets: a COVID-19 dataset consisting of Computed tomography (CT) images for challenging classification [65], Blood Cell Images classification [66] and CIFAR-10 dataset

Algorithm 1: Gibbs sampler for the proposed method.

Fix the hyperparameters Φ_m ;

for $r = 1, \dots, S$ **do**

- * Sample c according to $f(c; W, b, \gamma)$;
- * Sample γ according to $f(\gamma; W, b, c, \sigma^2)$;
- * Sample b according to $f(b; W, c, \gamma, \lambda_b)$;
- * Sample σ^2 according to $f(\sigma^2; \mu, \gamma, \delta)$;
- * Sample λ_b according to $f(\lambda_b; b, \mu, \delta)$;
- * Sample λ_l according to $f(\lambda_l; \delta, \mu)$;
- $\forall l \in \{1, \dots, L\}$;
- * Sample W as described in [22] ;

end

[67]. Furthermore, our method is evaluated on three regression challenges involving the prediction of diabetes, head pose estimation, and age group detection from facial images, using the Biwi Head Pose Dataset [68], Faces - Age Detection Dataset [69], and a Diabetes Dataset [70]. These varied applications demonstrate the versatility and robustness of our algorithm in different contexts.

In our experimental setup, we assess the performance of various activation functions using two distinct optimizers: the widely used Adam optimizer and our recently developed non-smooth Hamiltonian Monte Carlo (ns-HMC) optimizer, which is set within a Markov Chain Monte Carlo framework [22]. The ns-HMC, paired with the conventional ReLU activation function, provides a Bayesian alternative to gradient descent methods for optimizing deep neural networks (DNNs) subjected to non-differentiable regularizations, such as the ℓ_1 norm that encourages sparsity in the weights distribution. By weaving Hamiltonian dynamics into the sampling process, the ns-HMC method circumvents the limitations of non-smooth regularizations, thereby improving the efficiency of optimization in scenarios where traditional backpropagation is not optimal.

For benchmarking, we employ two contrasting optimization strategies: "ReLU+ns-HMC," which combines the standard ReLU with ns-HMC in a Bayesian framework, and the Adam optimizer, set at a learning rate of 10^{-3} , alongside eight prevalent activation functions: ReLU, LReLU, ELU, PReLU, SeLU [71], swish, FReLU, and MeLU. As regards coding, we used python programming language with Keras and Tensorflow libraries on an Intel(R) Core(TM) i7-2720QM CPU 2.20GHZ architecture with 16 Go memory.

A. ConvNet models

In this work, four CNN architectures are utilized. Similar to the LeNet model [72], the first one has two fully-connected and three convolutional (Conv-32, Conv-64, and Conv-128) layers (FC-64 and FC-softmax). The second architecture has nine convolutional (3XConv-32, 3XConv-64, and 3XConv-128) and three FC layers (FC-128, FC-64 and FC-softmax) which are organized similarly to VGG-Net [73]. These architectures are shown in Table I. The other two CNNs are designed for regression tasks, as discussed in section VI-F.

Each incorporates convolutional layers with a stride size of 1, using 3×3 filters, and includes 2×2 max-pooling layers. Deep neural networks are expanded with three regularizing strategies since they can easily overfit when trained on small datasets: Batch Normalization [74], ℓ_1 Regularization [75] and Dropout [76].

We chose to use different architecture depths in the experiments, mainly to test the ability of our proposed method to achieve better performance in the shallow model. In this sense, training with large and complex data can be expensive.

TABLE I
CONVNET WITH REGULARIZATION TECHNIQUES.

CNN_1	CNN_2
Conv3x3-32:stride=1 BatchNormalization MaxPool 2x2 Dropout(0.2)	3 X Conv3x3-32:stride=1 BatchNormalization MaxPool 2x2 Dropout(0.3)
Conv3x3-64:stride=1 BatchNormalization MaxPool 2x2 Dropout(0.3)	3 X Conv3x3-64:stride=1 BatchNormalization MaxPool 2x2 Dropout(0.3)
Conv3x3-128:stride=1 BatchNormalization MaxPool 2x2 Dropout(0.4)	3 X Conv3x3-128:stride=1 BatchNormalization MaxPool 2x2 Dropout(0.4)
Flattening FC-64 Dropout(0.3)	Flattening FC-128 Dropout(0.35) FC-64 Dropout(0.35)
FC-softmax	FC-softmax

B. Sampling Results

After using the proposed Bayesian optimization method to train the CNN models detailed above for the classification of Covid-19 CT images, we analyzed the convergence behavior. Figure 3 presents the sampling chains for the γ , b , and c parameters of the proposed MMeLU function (a-c), as well as the histograms of the corresponding samples (d-f). The sampling chains and histograms of the sampled coefficients confirm the good convergence properties of the designed Gibbs sampler. After a burn-in period of 350 iterations, the algorithm achieves stable convergence and exhibits a good mixing rate of the sampled chains.

C. Experiment 1: COVID-19 classification using CT images

This section examines the effectiveness of our approach in classifying Covid-19 infections from other pneumonia in CT data. The COVID-CT dataset ¹ includes 397 images that are negative for COVID-19 and 349 images that are positive for COVID-19, and belong to 216 patients. We used 566 images for the train and 180 images for the test.

Table II presents compelling evidence that the proposed

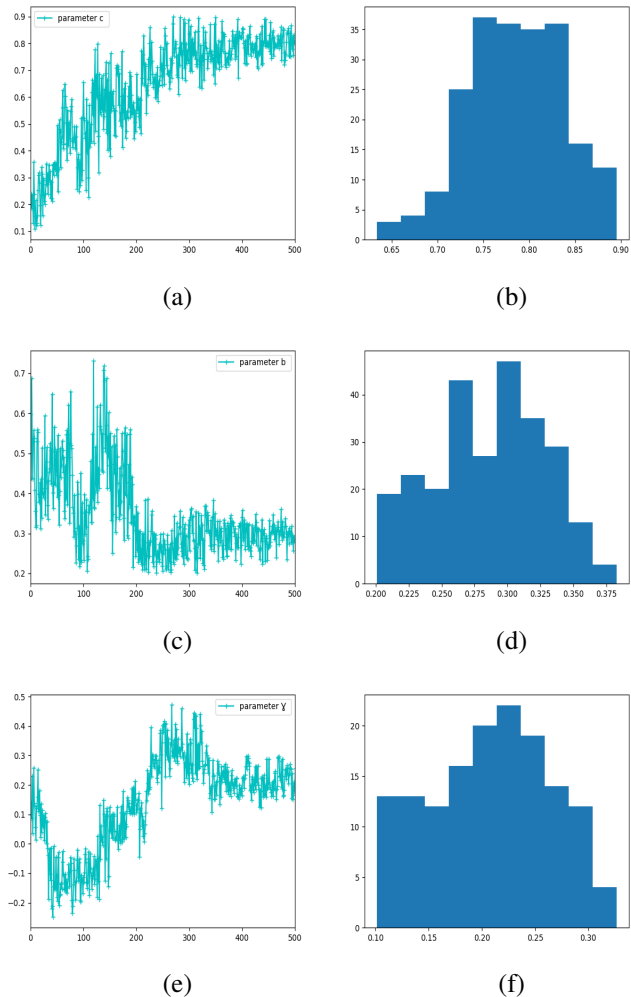


Fig. 3. Sampling of parameters c (a,b), b (c,d), and γ (e,f): chains and histograms.

Bayesian method outperforms all other activation functions for both the CNN_1 and CNN_2 architectures. Moreover, except with respect to "ReLU + ns-HMC", the computational time for convergence is shorter than all the other activation functions. With respect to our previous work ("ReLU + ns-HMC"), it is worth noting that the proposed method performs better in terms of accuracy and loss value, which confirms the usefulness of the used trainable activation function. However, due to the use of additional parameters, the proposed method is approximately 15% slower. These conclusions hold for both CNN_1 and CNN_2 .

Notably, when regularization is employed, a considerable drop in performance is observed for all competing activation functions, which is largely attributable to the inherent difficulty of classifying CT images due to their content richness and the similarities between Covid-19 infection and other types of pneumonia.

Learning and test curves (accuracy and loss) illustrated in Figures 4 and 5 confirm the good behavior of the proposed method, which is not necessarily the case of the other compet-

¹<https://www.kaggle.com/luisblanche/covidct>

TABLE II

EXPERIMENT 1: CT CLASSIFICATION RESULTS WITH CNN_1 AND CNN_2 (ACTIVATION FUNCTIONS (ACT FCTS), COMPUTATION TIME IN MINUTES, ACCURACY (ACC), LOSS, SENSITIVITY (SENS) AND SPECIFICITY (SPEC)).

Act. Fcts	CNN_1					CNN_2				
	Time	Acc.	Loss	Sens.	Spec.	Time	Acc.	Loss	Sens.	Spec.
MMeLU	46.21	0.90	0.23	0.87	0.86	61.92	0.91	0.21	0.87	0.87
ReLU+ns-HMC	40	0.84	0.26	0.82	0.80	53	0.88	0.24	0.86	0.85
ReLU	58	0.73	0.43	0.69	0.68	81	0.77	0.39	0.74	0.72
LReLU	65.18	0.73	0.52	0.71	0.69	105	0.78	0.44	0.76	0.75
ELU	63	0.75	0.47	0.75	0.74	97	0.76	0.46	0.75	0.75
PReLU	71.28	0.68	0.72	0.64	0.62	119	0.70	0.76	0.68	0.67
SeLU	64.75	0.77	0.78	0.74	0.72	107	0.76	0.69	0.74	0.73
Swish	83.41	0.68	0.58	0.65	0.62	132	0.73	0.55	0.71	0.70
FReLU	77.8	0.76	0.59	0.76	0.75	123	0.77	0.52	0.76	0.75
MeLU	95.89	0.77	0.43	0.77	0.76	146	0.80	0.38	0.80	0.80

ing models where a marked difference between the precision and loss curves can be noticed. For example, while the LReLU function introduces a negative bias that suppresses excessive activations, an inappropriate bias value can lead to underfitting. Similarly, although the ELU function allows negative activation, its exponential form can lead to an explosion of the activation value for large values of x .

The Swish function is known to accelerate learning convergence, but it can also lead to overfitting by being more sensitive to outliers. Likewise, although the FReLU function can capture complex data patterns, it can also suffer from overfitting if the parameters are not well chosen.

These remarkable differences confirm the interest and efficiency of our MMeLU function, which outperforms all competing activation functions in terms of accuracy and robustness to regularization.

D. Experiment 2: Blood Cell Images classification

In this experiment, we employed the BCCD (Blood Cell Count Dataset)² to measure the performance of our proposed activation function MMeLU. The BCCD database contains 12,447 enhanced blood cell images (JPEG), each labeled with the corresponding cell type (CSV). This collection includes 9957 images for training and 2490 images for testing. The images are divided into four distinct categories of blood cells: eosinophils, lymphocytes, monocytes, and neutrophils. In the original training image set, there are 2497 eosinophils, 2483 lymphocytes, 2478 monocytes, and 2499 neutrophils, while the test set contains 623 eosinophils, 623 lymphocytes, 620 monocytes, and 624 neutrophils.

Table III presents the results obtained using the BCCD dataset, where our proposed method significantly outperforms all competing activation functions. It achieves a high accuracy of up to 96% in the second CNN model. Furthermore, the table highlights a notable efficiency in processing time; while competing methods exceed 400 minutes for CNN_1 , our method requires only about 330 minutes. This efficiency is consistently observed in CNN_2 as well.

Additionally, the MMeLU function demonstrates superior global performance when compared to the "ReLU+ns-HMC" across both CNN models. A key advantage of MMeLU is

its fully automatic parameter estimation, which contrasts sharply with the ReLU+ns-HMC, where manual setting of some parameters is necessary. This aspect enhances the user-friendliness and practical applicability of MMeLU. Moreover, the ability of our method to deliver remarkable results in challenging scenarios is further evidenced in the first experiment, underlining its robustness and versatility in diverse applications.

E. Experiment 3: CIFAR-10 image classification

In this experiment, the learning performance of the proposed model (made up of the trainable activation function and the Bayesian sparse optimization) is evaluated using the standard CIFAR-10 dataset. The CIFAR-10 dataset contains 60000 32×32 color images divided into 10 classes with 6000 images per class. There are 50,000 training and 10,000 test images. Table IV presents the classification results for the CIFAR-10 dataset. It can be observed that the proposed Bayesian model performed well overall, even when multiple classes were used. In contrast, all competing activation functions (with Adam optimizer) had similar accuracy, around 88%, with a loss rate almost double that of the proposed model for both architectures.

Moreover, the MMeLU activation function takes significantly less time for training than other activation functions + Adam, which often take a long time before convergence. For the competing functions when used with Adam, we noticed similar performance between the trainable and standard activation functions, with a slight superiority observed with the trainable function ELU.

The same conclusions can be drawn by examining the results of "ReLU + ns-HMC" on this dataset. These findings suggest that the proposed MMeLU function provides better flexibility for the activation task, even for the multi-class case.

F. Experiment 4: Regression

In this section, we evaluate the performance of our algorithm on three regression challenges: diabetes prediction, head pose estimation, and age group detection from facial images.

Biwi Head Pose Dataset: This dataset contains over 15,000

²<https://www.kaggle.com/datasets/paultimothymooney/blood-cells>

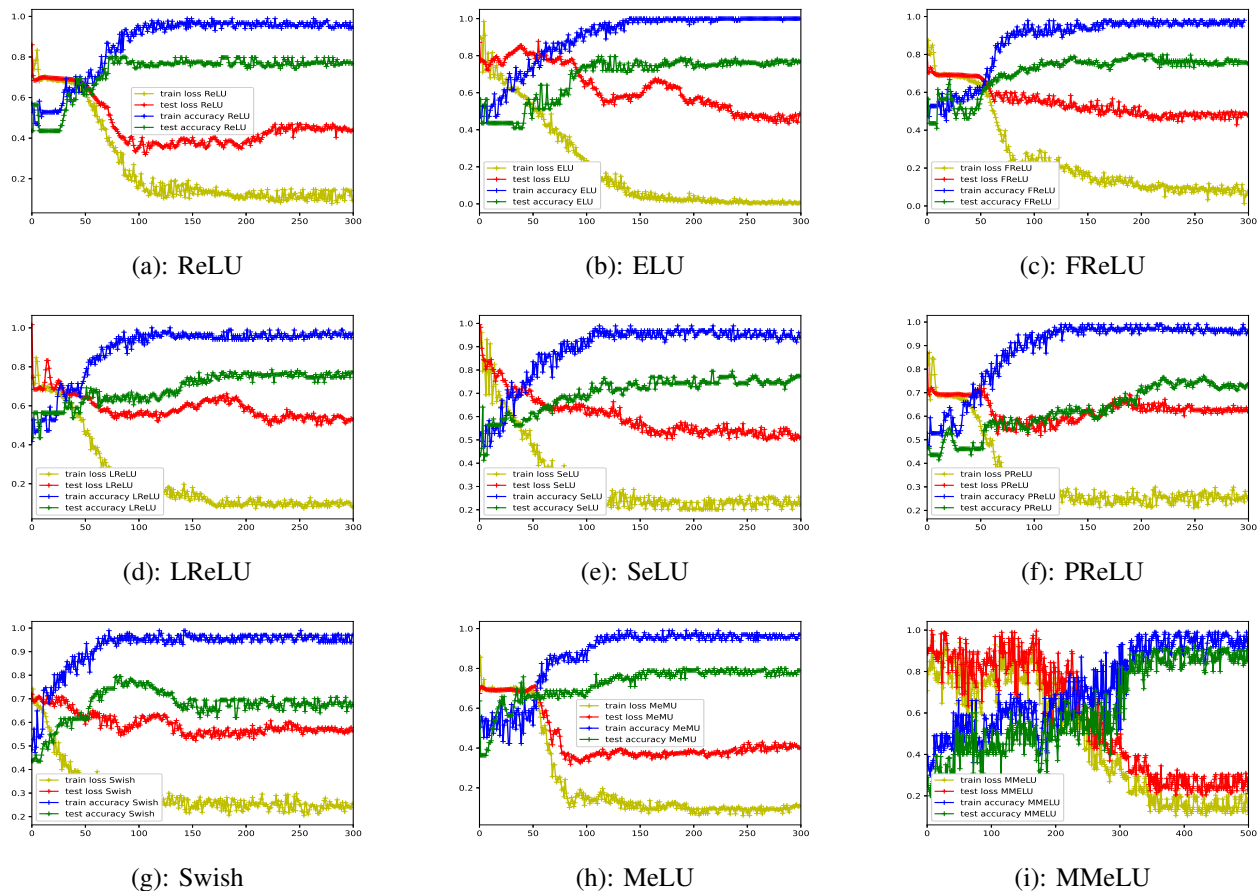


Fig. 4. Experiment 1: Train and test curves using CNN_1 for all competing activation functions.

TABLE III

EXPERIMENT 2: BLOOD CELL IMAGES CLASSIFICATION RESULTS WITH CNN_1 AND CNN_2 (ACTIVATION FUNCTIONS (ACT FCTS), COMPUTATION TIME IN MINUTES, ACCURACY (ACC), LOSS, SENSITIVITY (SENS) AND SPECIFICITY (SPEC)).

Act. Fcts	CNN_1					CNN_2				
	Time	Acc.	Loss	Sens.	Spec.	Time	Acc.	Loss	Sens.	Spec.
MMeLU	329.5	0.93	0.23	0.93	0.93	510.4	0.96	0.20	0.96	0.96
ReLU+ns-HMC	335.6	0.89	0.31	0.89	0.89	549.3	0.90	0.29	0.90	0.90
ReLU	406.2	0.79	0.64	0.78	0.78	642.4	0.84	0.43	0.84	0.83
LReLU	424.1	0.63	0.88	0.63	0.60	675.5	0.71	0.65	0.69	0.63
ELU	427.7	0.73	0.72	0.70	0.66	683.1	0.75	0.57	0.74	0.72
PReLU	453.9	0.72	0.69	0.71	0.67	697.5	0.70	0.66	0.68	0.66
SeLU	441.2	0.69	0.75	0.60	0.62	688.9	0.79	0.65	0.76	0.71
Swish	483.8	0.78	0.64	0.77	0.76	709.4	0.78	0.60	0.76	0.74
FReLU	425.7	0.71	0.71	0.69	0.67	670.8	0.81	0.47	0.81	0.79
MeLU	498	0.81	0.60	0.81	0.80	712.1	0.85	0.40	0.85	0.84

RGB-D images from videos of 20 individuals (16 men and 4 women), captured by a Kinect camera. It is used for the assessment of head orientations, with natural movements covering ranges of $[-60^\circ, 60^\circ]$ in pitch, $[-75^\circ, 75^\circ]$ in yaw, and $[-20^\circ, 20^\circ]$ in roll.

Faces - Age Detection Dataset: Comprising 19,906 images, this database is used for facial feature analysis and age group estimation. Each image is categorized into a specific age group, providing a valuable resource for applications in biometrics and entertainment.

Diabetes Dataset: Originating from the BRFS 2015, this

set includes 441,455 individuals for studying diabetes. It focuses on the accuracy of diabetes prediction from BRFS responses, identifying the most predictive risk factors, and the possibility of predicting diabetes with a subset of these factors.

These three datasets offer a wide range of applications for testing the versatility and robustness of our algorithm in various and challenging regression tasks. To effectively analyze these datasets, we employ two distinct CNN architectures. The first one, used for the Diabetes Dataset, comprises 5 fully connected layers (512, 256, 128, 64, and




Fig. 5. Experiment 1: Train and test curves using CNN_2 for all competing activation functions.

TABLE IV
EXPERIMENT 3: CIFAR-10 CLASSIFICATION RESULTS WITH CNN_1 AND CNN_2 (ACTIVATION FUNCTIONS (ACT FCTS), COMPUTATION TIME IN MINUTES, ACCURACY (ACC), LOSS, SENSITIVITY (SENS) AND SPECIFICITY (SPEC)).

Act. Fcts	CNN_1					CNN_2				
	Time	Acc.	Loss	Sens.	Spec.	Time	Acc.	Loss	Sens.	Spec.
MMeLU	120.7	0.91	0.21	0.89	0.88	332.5	0.93	0.20	0.90	0.88
ReLU+ns-HMC	100.7	0.90	0.25	0.89	0.87	331	0.92	0.21	0.90	0.87
ReLU	161	0.87	0.42	0.83	0.81	429	0.90	0.36	0.87	0.86
LReLU	170	0.87	0.52	0.85	0.84	437.9	0.85	0.55	0.83	0.81
ELU	165	0.88	0.41	0.86	0.85	438	0.88	0.39	0.86	0.85
PReLU	198.2	0.82	0.48	0.84	0.82	466.8	0.84	0.47	0.81	0.79
SeLU	173.6	0.84	0.43	0.85	0.84	454.3	0.86	0.39	0.87	0.86
Swish	209.7	0.83	0.49	0.82	0.79	478.3	0.85	0.46	0.83	0.81
FReLU	172.3	0.84	0.41	0.82	0.81	431.7	0.86	0.38	0.83	0.81
MeLU	219.9	0.86	0.40	0.83	0.83	485.6	0.90	0.34	0.88	0.87

32). The second CNN, applied to the other two datasets (Biwi Head Pose and Faces: Age Detection), includes 25 convolutional layers ($5 \times \text{Conv}3 \times 3-32$, $5 \times \text{Conv}3 \times 3-64$, $5 \times \text{Conv}3 \times 3-128$, $5 \times \text{Conv}3 \times 3-256$, and $5 \times \text{Conv}3 \times 3-512$) followed by three fully connected layers (FC-256, FC-128, and FC-64). All layers utilize convolutional layers with 3×3 kernel filters, alongside 2×2 max-pooling and a stride size of 1.

The detailed analysis of the results obtained on the Biwi Kinect Head Pose, Faces - Age Detection, and CDC Diabetes datasets, as presented in tables V, VI, and VII, respectively,

clearly demonstrates the outstanding performance of the MMeLU activation function. This performance is not only illustrated by lower scores in terms of Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE), but also by the significant impact of these scores on the quality of the results and their practical applications.

For the Biwi Kinect Head Pose Dataset, as shown in table V, the low MSE indicates that MMeLU provides a reliable precision throughout both the training and testing phases, a key factor for applications like augmented reality where

precision is paramount. Such a reduction in MSE underscores the model’s ability to generalize effectively, steering clear of overfitting.

For the Age Detection Dataset, as reported in Table VI, MMeLU demonstrated excellent performance, evidenced by low scores in MSE, MAE and RMSE. The low MAE value is particularly significant, indicating high average precision in predictions, crucial for age group classification in facial recognition applications. This implies that while errors are present, they are not significant on average, ensuring greater reliability in the correct identification of age brackets. Moreover, a low RMSE further suggests that errors in predictions are generally of low magnitude, critical for maintaining the integrity of biometric systems. The consistent performance of MMeLU, particularly in terms of MAE and RMSE, underscores its effectiveness in scenarios where prediction accuracy is directly linked to user experience and system reliability.

The same conclusion holds for the CDC Diabetes Dataset, as illustrated in table VII, which demonstrates low scores in MSE, MAE, and RMSE. Such performance is critical in the medical field, where accurate predictions can greatly impact treatment decisions and patient management. The consistency of MMeLU’s results across these metrics underscores its robustness and suitability in clinical settings, where the precision of predictions can have substantial implications for patient health.

These analyses reveal that the remarkable performance of MMeLU is not just limited to surpassing other activation functions in statistical scores, but also encompasses significantly enhancing the reliability and applicability of predictions across various practical contexts. The ability of MMeLU to consistently provide precise and reliable results in diverse regression scenarios underscores its potential as a versatile tool in the field of machine learning. This makes it a valuable asset for a broad spectrum of practical applications.

TABLE V

RESULTS OF CDC DATASET REGRESSION: ACTIVATION FUNCTIONS (ACT FCTS), MEAN SQUARED ERROR ON TRAINING SET (ACT FCTS), MEAN SQUARED ERROR ON TEST SET (MSE_{train}), MEAN ABSOLUTE ERROR ON TEST SET (MSE_{test}), MEAN ABSOLUTE ERROR WITH TEST SET (MAE_{test}), ROOT MEAN SQUARED ERROR ($RMSE$)).

Act. Fcts	MSE_{Train}	MSE_{Test}	MAE_{Test}	$RMSE$
MMeLU	0.20	0.38	0.07	0.61
ReLU+ns-HMC	0.26	0.46	0.10	0.68
ReLU	0.33	0.62	0.15	0.78
LReLU	0.38	0.69	0.18	0.83
ELU	0.34	0.64	0.16	0.80
PReLU	0.37	0.66	0.18	0.81
SeLU	0.33	0.65	0.14	0.80
Swish	0.41	0.71	0.20	0.85
FReLU	0.39	0.70	0.19	0.84
MeLU	0.34	0.63	0.17	0.79

VII. CONCLUSION

In this paper, we proposed a Bayesian approach for training sparse deep neural networks with trainable activation func-

TABLE VI

REGRESSION RESULTS FOR BIWI KINECT HEAD POSE DATASET: ACTIVATION FUNCTIONS (ACT FCTS), MEAN SQUARED ERROR ON TRAINING SET (MSE_{train}), MEAN SQUARED ERROR ON TEST SET (MSE_{test}), MEAN ABSOLUTE ERROR ON TEST SET (MAE_{test}), ROOT MEAN SQUARED ERROR ($RMSE$)).

Act. Fcts	MSE_{Train}	MSE_{Test}	MAE_{Test}	$RMSE$
MMeLU	2.81	15.33	2.90	3.91
ReLU+ns-HMC	4.91	25.95	4.89	5.09
ReLU	7.09	38.45	6.67	6.20
LReLU	5.18	31.05	7.59	5.57
ELU	6.21	36.86	8.14	6.07
PReLU	5.16	34.73	5.04	5.89
SeLU	5.05	33.47	6.99	5.87
Swish	9.75	43.83	8.51	6.62
FReLU	5.69	30.85	7.02	5.55
MeLU	6.19	33.92	5.67	5.82

TABLE VII

REGRESSION RESULTS FOR FACES - AGE DETECTION DATASET: ACTIVATION FUNCTIONS (ACT FCTS), MEAN SQUARED ERROR ON TRAINING SET (MSE_{train}), MEAN SQUARED ERROR ON TEST SET (MSE_{test}), MEAN ABSOLUTE ERROR ON TEST SET (MAE_{test}), ROOT MEAN SQUARED ERROR ($RMSE$)).

Act. Fcts	MSE_{Train}	MSE_{Test}	MAE_{Test}	$RMSE$
MMeLU	52.26	54.82	3.79	7.40
ReLU+ns-HMC	65.51	67.02	5.78	8.25
ReLU	80.71	83.87	7.63	9.15
LReLU	79.30	81.64	8.08	9.03
ELU	79.11	82.52	7.71	9.08
PReLU	83.03	88.42	7.26	9.40
SeLU	80.04	84.87	7.41	9.21
Swish	89.14	95.01	8.13	9.74
FReLU	79.64	82.72	7.25	9.09
MeLU	76.01	80.53	7.04	8.97

tions. Our method learns the weights and parameters of the proposed activation function directly from the data without any user configuration. Compared to competing algorithms, our approach achieves promising results with better classification and regression accuracy and high generalization performance, while also being computationally efficient, particularly in challenging cases. Our experiments, encompassing a range of standard datasets including CIFAR-10, as well as a variety of regression challenges such as diabetes prediction, head pose estimation, and age group detection from facial images, demonstrate the precision, robustness, and wide-ranging efficacy of our approach. These diverse applications highlight not only the adaptability of our method across multiple domains, but also its reliability in delivering accurate and stable results in both classification and regression tasks.

Future work will focus on parallelizing the algorithm to enable GPU calculations and further reduce computational time.

REFERENCES

- [1] A. Alsarhan, M. Alauthman, E. Alshdaifat, A.-R. Al-Ghuwairi, A. Al-Dubai, Machine learning-driven optimization for svm-based intrusion detection system in vehicular ad hoc networks, *Journal of Ambient Intelligence and Humanized Computing* (2021) 1–10.
- [2] Z. Wang, J. Chen, S. C. Hoi, Deep learning for image super-resolution: A survey, *IEEE transactions on pattern analysis and machine intelligence* 43 (10) (2020) 3365–3387.

- [3] C. Dong, C. C. Loy, K. He, X. Tang, Image super-resolution using deep convolutional networks, *IEEE transactions on pattern analysis and machine intelligence* 38 (2) (2015) 295–307.
- [4] V. Sree, J. Mapes, S. Dua, O. S. Lih, J. E. Koh, E. J. Ciaccio, U. R. Acharya, et al., A novel machine learning framework for automated detection of arrhythmias in eeg segments, *Journal of Ambient Intelligence and Humanized Computing* (2021) 1–18.
- [5] S. N. B. Jaini, D. Lee, S. Lee, M. Kim, Y. Kwon, Tool monitoring of end milling based on gap sensor and machine learning, *Journal of Ambient Intelligence and Humanized Computing* (2021) 1–13.
- [6] R. S. Mitchell, R. A. Sherlock, L. A. Smith, An investigation into the use of machine learning for determining oestrus in cows, *Computers and electronics in agriculture* 15 (3) (1996) 195–213.
- [7] A. Drewek-Ossowicka, M. Pietrolaj, J. Rumiński, A survey of neural networks usage for intrusion detection systems, *Journal of Ambient Intelligence and Humanized Computing* 12 (1) (2021) 497–514.
- [8] T. K. Sajja, H. K. Kalluri, Image classification using regularized convolutional neural network design with dimensionality reduction modules: Rcnndrm, *Journal of Ambient Intelligence and Humanized Computing* (2021) 1–12.
- [9] M. Fakhfakh, B. Bouaziz, F. Gargouri, L. Chaari, Prognostic: Covid-19 prognosis using recurrent and convolutional neural networks, *The Open Medical Imaging Journal* 12 (1) (2020).
- [10] Y. Li, H. Zhang, X. Xue, Y. Jiang, Q. Shen, Deep learning for remote sensing image classification: A survey, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 8 (6) (2018) e1264.
- [11] S. Ji, W. Xu, M. Yang, K. Yu, 3d convolutional neural networks for human action recognition, *IEEE transactions on pattern analysis and machine intelligence* 35 (1) (2012) 221–231.
- [12] K. Hara, D. Saito, H. Shouno, Analysis of function of rectified linear unit used in deep learning, in: 2015 international joint conference on neural networks (IJCNN), IEEE, 2015, pp. 1–8.
- [13] M. M. Lau, K. H. Lim, Review of adaptive activation function in deep neural network, in: 2018 IEEE-EMBS Conference on Biomedical Engineering and Sciences (IECBES), IEEE, 2018, pp. 686–690.
- [14] A. Apicella, F. Donnarumma, F. Isgrò, R. Prevete, A survey on modern trainable activation functions, *Neural Networks* 138 (2021) 14–32.
- [15] C. T. Butts, Network inference, error, and informant (in) accuracy: a bayesian approach, *social networks* 25 (2) (2003) 103–140.
- [16] C. Andrieu, A. Doucet, R. Holenstein, Particle markov chain monte carlo methods, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72 (3) (2010) 269–342.
- [17] C. Robert, G. Casella, *Monte Carlo statistical methods*, Springer Science & Business Media, 2013.
- [18] L. Chaari, H. Batatia, N. Dobigeon, J.-Y. Tourneret, A hierarchical sparsity-smoothness bayesian model for l0+l1+l2 regularization, in: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2014, pp. 1901–1905.
- [19] L. Chaari, A bayesian grouplet transform, *Signal, Image and Video Processing* 13 (2019) 871–878.
- [20] M. Fakhfakh, B. Bouaziz, L. Chaari, F. Gargouri, Efficient bayesian learning of sparse deep artificial neural networks, in: International Symposium on Intelligent Data Analysis, Springer, 2022, pp. 78–88.
- [21] M. Fakhfakh, B. Bouaziz, F. Gargouri, L. Chaari, Bayesian optimization using hamiltonian dynamics for sparse artificial neural networks, in: 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI), IEEE, 2022, pp. 1–4.
- [22] M. Fakhfakh, L. Chaari, B. Bouaziz, F. Gargouri, Non-smooth bayesian learning for artificial neural networks, *Journal of Ambient Intelligence and Humanized Computing* (2022).
- [23] U. of Illinois at Urbana-Champaign. Center for Supercomputing Research, Development, G. Cybenko, Continuous valued neural networks with two hidden layers are sufficient, 1988.
- [24] F. Xiao, Y. Honma, T. Kono, A simple algebraic interface capturing scheme using hyperbolic tangent function, *International journal for numerical methods in fluids* 48 (9) (2005) 1023–1040.
- [25] G. Cybenko, Approximation by superpositions of a sigmoidal function, *Mathematics of control, signals and systems* 2 (4) (1989) 303–314.
- [26] B. Ding, H. Qian, J. Zhou, Activation functions and their characteristics in deep neural networks, in: 2018 Chinese control and decision conference (CCDC), IEEE, 2018, pp. 1836–1841.
- [27] P. d. B. Harrington, Sigmoid transfer functions in backpropagation neural networks, *Analytical Chemistry* 65 (15) (1993) 2167–2168.
- [28] S. Chadha, Fractional programming with absolute-value functions, *European Journal of Operational Research* 141 (1) (2002) 233–238.
- [29] Y. Bengio, P. Simard, P. Frasconi, Learning long-term dependencies with gradient descent is difficult, *IEEE transactions on neural networks* 5 (2) (1994) 157–166.
- [30] X. Glorot, A. Bordes, Y. Bengio, Deep sparse rectifier neural networks, in: Proceedings of the fourteenth international conference on artificial intelligence and statistics, JMLR Workshop and Conference Proceedings, 2011, pp. 315–323.
- [31] C. Gulcehre, M. Moczulski, M. Denil, Y. Bengio, Noisy activation functions, in: International conference on machine learning, PMLR, 2016, pp. 3059–3068.
- [32] X. Nie, J. Cao, Multistability of second-order competitive neural networks with nondecreasing saturated activation functions, *IEEE Transactions on Neural Networks* 22 (11) (2011) 1694–1708.
- [33] A. Montalto, G. Tessoro, R. Prevete, A linear approach for sparse coding by a two-layer neural network, *Neurocomputing* 149 (2015) 1315–1323.
- [34] A. L. Maas, A. Y. Hannun, A. Y. Ng, et al., Rectifier nonlinearities improve neural network acoustic models, in: Proc. icml, Vol. 30, Atlanta, Georgia, USA, 2013, p. 3.
- [35] K. Konda, R. Memisevic, D. Krueger, Zero-bias autoencoders and the benefits of co-adapting features, *arXiv preprint arXiv:1402.3337* (2014).
- [36] C. Dugas, Y. Bengio, F. Bélisle, C. Nadeau, R. Garcia, Incorporating second-order functional knowledge for better option pricing, *Advances in neural information processing systems* 13 (2000).
- [37] A. Shah, E. Kadam, H. Shah, S. Shinde, S. Shingade, Deep residual networks with exponential linear unit, in: Proceedings of the third international symposium on computer vision and the internet, 2016, pp. 59–65.
- [38] M. Çelebi, M. Ceylan, New complex valued activation functions: Complex modifiedswish, complex e-swish and complex flatten-tswish., *International Journal of Advanced Research in Computer Science* 11 (2) (2020).
- [39] H. H. Chieng, N. Wahid, O. Pauline, S. R. K. Perla, Flatten-t swish: a thresholded relu-swish-like activation function for deep learning, *International Journal of Advances in Intelligent Informatics* 4 (2) (2018) 76–86.
- [40] C.-T. Chen, W.-D. Chang, A feedforward neural network with function shape autotuning, *Neural networks* 9 (4) (1996) 627–641.
- [41] S. Guarnieri, Multilayer neural networks with adaptive spline-based activation functions, in: Proceedings of Word Congress on Neural Networks WCNN’95, Washington, DC, July, 1995, pp. 17–21.
- [42] F. Piazza, A. Uncini, M. Zenobi, Artificial neural networks with adaptive polynomial activation function (1992).
- [43] L. Trottier, P. Giguere, B. Chaib-Draa, Parametric exponential linear unit for deep convolutional neural networks, in: 2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA), IEEE, 2017, pp. 207–214.
- [44] S. Qiu, X. Xu, B. Cai, Frelu: flexible rectified linear units for improving convolutional neural networks, in: 2018 24th international conference on pattern recognition (icpr), IEEE, 2018, pp. 1223–1228.
- [45] K. He, X. Zhang, S. Ren, J. Sun, Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, in: Proceedings of the IEEE international conference on computer vision, 2015, pp. 1026–1034.
- [46] A. Nader, D. Azar, Evolution of activation functions: An empirical investigation, *ACM Transactions on Evolutionary Learning and Optimization* 1 (2) (2021) 1–36.
- [47] A. Nader, D. Azar, Searching for activation functions using a self-adaptive evolutionary algorithm, in: Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion, 2020, pp. 145–146.
- [48] M. Basirat, P. M. Roth, The quest for the golden activation function, *Proceedings of the ARW & OAGM Workshop* (2019) 1–16.
- [49] P. Ramachandran, B. Zoph, Q. V. Le, Searching for activation functions, *arXiv preprint arXiv:1710.05941* (2017).
- [50] G. Maguolo, L. Nanni, S. Ghidoni, Ensemble of convolutional neural networks trained with different activation functions, *Expert Systems with Applications* 166 (2021) 114048.
- [51] L. R. Sutfield, F. Brieger, H. Finger, S. Fullhase, G. Pipa, Adaptive blending units: Trainable activation functions for deep neural networks, in: Science and Information Conference, Springer, 2020, pp. 37–50.
- [52] S. Scardapane, S. Van Vaerenbergh, S. Totaro, A. Uncini, Kafnets: Kernel-based non-parametric activation functions for neural networks, *Neural Networks* 110 (2019) 19–32.
- [53] S. Scardapane, S. Van Vaerenbergh, A. Hussain, A. Uncini, Complex-valued neural networks with nonparametric activation functions, *IEEE*

- Transactions on Emerging Topics in Computational Intelligence 4 (2) (2018) 140–150.
- [54] Ö. F. Ertuğrul, A novel type of activation function in artificial neural networks: Trained activation function, *Neural Networks* 99 (2018) 148–157.
- [55] P. M. Baggenstoss, Improved auto-encoding using deterministic projected belief networks, *arXiv preprint arXiv:2309.07481* (2023).
- [56] B. Shabash, K. C. Wiese, Evonn: a customizable evolutionary neural network with heterogeneous activation functions, in: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 2018, pp. 1449–1456.
- [57] O. Castillo, P. Melin, A review on interval type-2 fuzzy logic applications in intelligent control, *Information Sciences* 279 (2014) 615–631.
- [58] S. Scardapane, M. Scarpiniti, D. Comminiello, A. Uncini, Learning activation functions from data using cubic spline interpolation, in: *Italian Workshop on Neural Nets*, Springer, 2017, pp. 73–83.
- [59] M. Hector, K. Aaron, F. Matthias, S. Jost Tobias, U. Matthias, B. Michael, D. Maximilian, L. Marius, H. Frank, Towards automatically-tuned deep neural networks, *Automated Machine Learning*. Springer (2019) 135–149.
- [60] X. Wang, Y. Qin, Y. Wang, S. Xiang, H. Chen, Reltanh: An activation function with vanishing gradient resistance for sae-based dnns and its application to rotating machinery fault diagnosis, *Neurocomputing* 363 (2019) 88–98.
- [61] G. Maguolo, L. Nanni, S. Ghidoni, Ensemble of convolutional neural networks trained with different activation functions, *Expert Systems with Applications* 166 (2021) 114048.
- [62] M. Fakhfakh, B. Bouaziz, H. Batatia, L. Chaari, Bayesian optimization for sparse artificial neural networks: Application to change detection in remote sensing, in: *Proceedings of International Conference on Information Technology and Applications*, Springer, 2022, pp. 39–49.
- [63] L. Chaari, J.-Y. Tournet, C. Chau, H. Batatia, A Hamiltonian Monte Carlo method for non-smooth energy sampling, *IEEE Trans. on Signal Process.* 64 (21) (2016) 5585 – 5594.
- [64] S. Lathuilière, P. Mesejo, X. Alameda-Pineda, R. Horaud, A comprehensive analysis of deep regression, *IEEE transactions on pattern analysis and machine intelligence* 42 (9) (2019) 2065–2081.
- [65] P. Afshar, S. Heidarian, N. Enshaei, F. Naderkhani, M. J. Rafiee, A. Oikonomou, F. B. Fard, K. Samimi, K. N. Plataniotis, A. Mohammadi, Covid-ct-md, covid-19 computed tomography scan dataset applicable in machine learning and deep learning, *Scientific Data* 8 (1) (2021) 121.
- [66] A. Acevedo, A. Merino, S. Alférez, Á. Molina, L. Boldú, J. Rodellar, A dataset of microscopic peripheral blood cell images for development of automatic recognition systems, *Data in brief* 30 (2020).
- [67] R. C. Çalik, M. F. Demirci, Cifar-10 image classification with convolutional neural networks for embedded systems, in: *2018 IEEE/ACS 15th International Conference on Computer Systems and Applications (AICCSA)*, IEEE, 2018, pp. 1–2.
- [68] G. Fanelli, M. Dantone, J. Gall, A. Fossati, L. Van Gool, Random forests for real time 3d face analysis, *International journal of computer vision* 101 (2013) 437–458.
- [69] C. Fernández, I. Huerta, A. Prati, A comparative evaluation of regression learning algorithms for facial age estimation, in: *International Workshop on Face and Facial Expression Recognition from Real World Videos*, Springer, 2014, pp. 133–144.
- [70] V. Chang, M. A. Ganatra, K. Hall, L. Golightly, Q. A. Xu, An assessment of machine learning models and algorithms for early prediction and diagnosis of diabetes using health indicators, *Healthcare Analytics* 2 (2022) 100118.
- [71] G. Klambauer, T. Unterthiner, A. Mayr, S. Hochreiter, Self-normalizing neural networks, *Advances in neural information processing systems* 30 (2017).
- [72] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, *Proceedings of the IEEE* 86 (11) (1998) 2278–2324.
- [73] U. Muhammad, W. Wang, S. P. Chattha, S. Ali, Pre-trained vggnet architecture for remote-sensing image scene classification, in: *24th International Conference on Pattern Recognition (ICPR)*, 2018, pp. 1622–1627.
- [74] S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: *International conference on machine learning*, PMLR, 2015, pp. 448–456.
- [75] Z. Xu, H. Zhang, Y. Wang, X. Chang, Y. Liang, L 1/2 regularization, *Science China Information Sciences* 53 (6) (2010) 1159–1169.
- [76] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, *The journal of machine learning research* 15 (1) (2014) 1929–1958.