

Learning activation functions from data using cubic spline interpolation

Simone Scardapane, Michele Scarpiniti, Danilo Comminiello, and Aurelio Uncini

Department of Information Engineering, Electronics and Telecommunications (DIET),
“Sapienza” University of Rome,
Via Eudossiana 18, 00184, Rome.

Email: {simone.scardapane, michele.scarpiniti, danilo.comminiello}@uniroma1.it;
aurel@ieee.org

Abstract. Neural networks require a careful design in order to perform properly on a given task. In particular, selecting a good activation function (possibly in a data-dependent fashion) is a crucial step, which remains an open problem in the research community. Despite a large amount of investigations, most current implementations simply select one *fixed* function from a small set of candidates, which is not adapted during training, and is shared among all neurons throughout the different layers. However, neither two of these assumptions can be supposed optimal in practice. In this paper, we present a principled way to have data-dependent adaptation of the activation functions, which is performed *independently* for each neuron. This is achieved by leveraging over past and present advances on cubic spline interpolation, allowing for local adaptation of the functions around their regions of use. The resulting algorithm is relatively cheap to implement, and overfitting is counterbalanced by the inclusion of a novel damping criterion, which penalizes unwanted oscillations from a predefined shape. Preliminary experimental results validate the proposal.

Keywords: Neural network, activation function, spline interpolation

1 Introduction

Neural networks (NNs) are extremely powerful tools for approximating complex nonlinear functions [7]. The nonlinear behavior is introduced in the NN architecture by the elementwise application of a given nonlinearity, called the activation function (AF), at every layer. Since AFs are crucial to the dynamics and computational power of NNs, the history of the two over the last decades is deeply connected [15]. As an example, the use of *differentiable* AFs was one of the major breakthroughs in NNs, leading directly to the back-propagation algorithm. More recently, progress on piecewise linear functions was shown to facilitate backward flow of information for training very deep networks [4]. At the same time, it is somewhat surprising that the vast majority of NNs only use a small handful of fixed functions, to be hand-chosen by the practitioner before the learning process. Worse, there is no principled reason to believe that a ‘good’ nonlinearity might be the same across all layers of the network, or even across neurons in the same layer.

This is shown clearly in a recent work by Agostinelli *et al.* [1], where every neuron in a deep network was endowed with an adaptable piecewise linear function with possibly different parameters, concluding that “*the standard one-activation-function-fits-all approach may be suboptimal*” in current practice. Experiments in AF adaptation have a long history, but they have never met a wide applicability in the field. The simplest approach is to parameterize each sigmoid function in the network by one or more ‘shape’ parameters to be optimized, such as in the seminal 1996 paper by Chen and Chang [3] or the later work by Trentin [16]. Along a similar line, one may consider the use of polynomial AFs, wherein each coefficient of the polynomial is adapted by gradient descent [11]. Additional investigations can be found in [20, 5, 2, 10, 9]. One strong drawback of these approaches is that the parameters involved affect the AF *globally*, such that a change in one region of the function may be counterproductive on a different, possibly faraway, region.

Several years ago, an alternative approach was introduced by using spline interpolating functions as AFs [17, 6], resulting in what was called a spline AF (SAF). Splines are an attractive choice for interpolating unknown functions, since they can be described by a small amount of parameters, yet each parameter has a local effect, and only a fixed number of them is involved every time an output value is computed [18]. The original works in [17, 6] had two main drawbacks that prevented a wider use of the underlying theory. First, SAFs were only investigated in an online setting, where updates are computed one sample at a time. Whether an efficient implementation is possible (and feasible) also for batch (or mini-batch) settings was not shown. Secondly, the obtained SAFs had a tendency to overfit training data, resulting in oscillatory behaviors which hindered performance. Inspired by recent successes in the field of nonlinear adaptive filtering [13, 14], our aim in this paper is two-fold. On one hand, we provide a modern introduction to the use of SAFs in neural networks, with a particular emphasis on their efficient implementation in the case of batch (or mini-batch) training. Our treatment clearly shows that the major problem in their implementation, which is evident from the discussion above, is the design of an efficient way to regularize their control points. In this sense, as a second contribution we provide a simple (yet effective) ‘damping’ criterion to prevent unwanted oscillations in the testing phase, which penalizes deviations from the original points in terms of ℓ_2 norm. A restricted set of experiments shows that the resulting formulation is able to achieve a lower test error than a standard NN with fixed AFs, while at the same time learning non-trivial activations with different shapes across different neurons.

The rest of the paper is organized as follows. Section 2 presents the basic theory of SAFs for the case of a single neuron. Section 3 extends the treatment to the case of a NN with one hidden layer, by deriving the gradient equations for the SAFs parameters in the internal layer. Then, Section 4 goes over the experimental results, while we conclude with some final remarks in Section 5.

2 The spline activation function

We begin our treatment of SAFs with the simplest case of a single neuron endowed with a flexible AF (see [17, 13] for additional details). Given a generic input $\mathbf{x} \in \mathbb{R}^D$,

the output of the SAF is computed as:

$$s = \mathbf{w}^T \mathbf{x}, \quad (1)$$

$$y = \varphi(s; \mathbf{q}), \quad (2)$$

where $\mathbf{w} \in \mathbb{R}^D$ (we suppose that an eventual bias term is added directly to the input vector), and the AF $\varphi(\cdot)$ is parameterized by a vector $\mathbf{q} \in \mathbb{R}^Q$ of internal parameters, called *knots*. The knots are a sampling of the AF values over Q representative points spanning the overall function. In particular, we suppose the knots to be uniformly spaced, i.e. $q_{i+1} = q_i + \Delta x$, for a fixed $\Delta x \in \mathbb{R}$, and symmetrically spaced around the origin. Given s , the output is computed by spline interpolation over the closest knot and its P right-most neighbors. The common choice $P = 3$, which we adopt in this paper, corresponds to *cubic* interpolation, and it is generally a good trade-off between locality of the output and interpolating precision.

Given the index i of the closest knot, we can define the normalized abscissa value between q_i and q_{i+1} as:

$$u = \frac{s}{\Delta x} - \left\lfloor \frac{s}{\Delta x} \right\rfloor. \quad (3)$$

where $\lfloor \cdot \rfloor$ is the floor operator. From u we can compute the normalized reference vector $\mathbf{u} = [u^P \ u^{P-1} \ \dots \ u \ 1]^T$, while from i we can extract the relevant control points $\mathbf{q}_i = [q_i \ q_{i+1} \ \dots \ q_{i+P}]^T$. We refer to the vector \mathbf{q}_i as the *ith span*. The output (2) is then computed as:

$$y = \varphi(s) = \mathbf{u}^T \mathbf{B} \mathbf{q}_i, \quad (4)$$

where $\mathbf{B} \in \mathbb{R}^{(P+1) \times (P+1)}$ is called the spline basis matrix. In this work, we use the Catmull-Rom (CR) spline with $P = 3$, given by:

$$\mathbf{B} = \frac{1}{2} \begin{bmatrix} -1 & 3 & -3 & 1 \\ 2 & -5 & 4 & -1 \\ -1 & 0 & 1 & 0 \\ 0 & 2 & 0 & 0 \end{bmatrix}. \quad (5)$$

Different bases give rise to alternative interpolation schemes, e.g. a spline defined by a CR basis passes through all the control points, but its second derivative is not continuous.

Apart from the locality of the output, SAFs have two additional interesting properties. First, the output in (4) is extremely efficient to compute, involving only vector-matrix products of very small dimensionality. Secondly, derivatives with respect to the internal parameters are equivalently simple and can be written down in closed form. In particular, the derivative of the nonlinearity $\varphi(s)$ with respect to the input s is given by:

$$\frac{\partial \varphi(s)}{\partial s} = \varphi'(s) = \frac{\partial \varphi(s)}{\partial u} \cdot \frac{\partial u}{\partial s} = \left(\frac{1}{\Delta x} \right) \dot{\mathbf{u}} \mathbf{B} \mathbf{q}_i, \quad (6)$$

where:

$$\dot{\mathbf{u}} = \frac{\partial \mathbf{u}}{\partial u} = [Pu^{P-1} \ (P-1)u^{P-2} \ \dots \ 1 \ 0]^T. \quad (7)$$

Given this, the derivative of the SAF output y with respect to \mathbf{w} is straightforward:

$$\frac{\partial \varphi(s)}{\partial \mathbf{w}} = \varphi'(s) \cdot \frac{\partial s}{\partial \mathbf{w}} = \varphi'(s) \mathbf{x}, \quad (8)$$

Similarly, for \mathbf{q}_i we obtain:

$$\frac{\partial \varphi(s)}{\partial \mathbf{q}_i} = \mathbf{B}^T \mathbf{u}. \quad (9)$$

while we have $\frac{\partial \varphi(s)}{\partial q_k} = 0$ for any element q_k outside the current span \mathbf{q}_i .

3 Designing networks with SAF neurons

3.1 Computing outputs and inner derivatives

Now we consider the more elaborate case of a single hidden layer NN, with a D -dimensional input, H neurons in the hidden layer, and O output neurons.¹ Every neuron in the network uses a SAF with possibly different adaptive control points, which are set independently during the training process. For easiness of computation, we suppose that the sampling set of the splines is the same for every neuron (i.e., each neuron has Q points equispaced according to the same Δx), and we also have a single shared basis matrix \mathbf{B} . The forward phase of the network is similar to that of a standard network. In particular, given the input \mathbf{x} , we first compute the output of the i th hidden neuron, $i = 1, \dots, H$, as:

$$h_i = \varphi(\mathbf{w}_{h_i}^T \mathbf{x}; \mathbf{q}_{h_i}). \quad (10)$$

These are concatenated in a single vector $\mathbf{h} = [h_1, \dots, h_H, 1]^T$, and the i th output of the network, $i = 1, \dots, O$, is given by:

$$f_i(\mathbf{x}) = y_i = \varphi(\mathbf{w}_{y_i}^T \mathbf{h}; \mathbf{q}_{y_i}). \quad (11)$$

The derivatives with respect to the parameters $\{\mathbf{w}_{y_i}, \mathbf{q}_{y_i}\}$, $i = 1, \dots, O$ can be computed directly with (8)-(9), substituting \mathbf{x} with \mathbf{h} . By back-propagation, the derivative of the i th output with respect to the j th (inner) weight vector \mathbf{w}_{h_j} is similar to a standard NN:

$$\frac{\partial y_i}{\partial \mathbf{w}_{h_j}} = \varphi'(s_{y_i}) \cdot \varphi'(s_{h_j}) \cdot [\mathbf{w}_{h_j}]_i \cdot \mathbf{x}, \quad (12)$$

where with a slight abuse of notation we let s_{y_i} denote the activation of the i th output (and similarly for s_{h_j}), $[\cdot]_i$ extracts the i th element of its input vector, and the two $\varphi'(\cdot)$ are given by (6). For the derivative of the control points of the j th hidden neuron, denote by $\mathbf{q}_{h_j,k}$ the currently active span, and by \mathbf{u}_{h_j} the corresponding reference vector. The derivative with respect to the i th output is then given by:

$$\frac{\partial y_i}{\partial \mathbf{q}_{h_j,k}} = \varphi'(s_{y_i}) \cdot [\mathbf{w}_{h_j}]_i \cdot \mathbf{B}^T \mathbf{u}_{h_j}. \quad (13)$$

¹We note that the following treatment can be extended easily to the case of a network with more than one hidden layer. However, restricting it to a single layer allow us to keep the discussion focused on the problems/advantages arising in the use of SAFs. We leave this extension to a future work.

3.2 Initialization of the control points

An important aspect that we have not discussed yet is how to properly initialize the control points. One immediate choice is to sample their values from an AF which is known to work well on the given problem, e.g. a hyperbolic tangent. In this way, the network is guaranteed to work similarly to a standard NN in the initial phase of learning. Additionally, we have found good improvements in error by adding Gaussian noise $\mathcal{N}(0, \sigma^2)$ with small variance σ^2 to a randomly chosen subset of control points (around 5% in our experiments). This provides a good variability in the beginning, similarly to how connections are set close to (but not identically equal to) zero during initialization.

3.3 Choosing a training criterion

Suppose we are provided with a training set of N input/output pairs in the form $\{\mathbf{x}_i, \mathbf{d}_i\}_{i=1}^N$. For simplicity of notation, we denote by \mathbf{w} the concatenation of all weight vectors $\{\mathbf{w}_{h_i}\}$ and $\{\mathbf{w}_{y_i}\}$, and by \mathbf{q} a similar concatenation of all control points. Training can be formulated as the minimization of the following cost function:

$$J(\mathbf{w}, \mathbf{q}) = \frac{1}{N} \sum_{i=1}^N L(\mathbf{d}_i, \mathbf{f}(\mathbf{x}_i)) + \lambda_w R_w(\mathbf{w}) + \lambda_q R_q(\mathbf{q}), \quad (14)$$

where $L(\cdot, \cdot)$ is an error function, while $R_w(\cdot)$ and $R_q(\cdot)$ provide meaningful regularization on the two set of parameters. The first two terms are well-known in the neural network literature [7], and they can be set accordingly. Particularly, in our experiments we consider a squared error term $L(\mathbf{d}_i, \mathbf{f}(\mathbf{x}_i)) = \|\mathbf{d}_i - \mathbf{f}(\mathbf{x}_i)\|_2^2$, and ℓ_2 regularization on the weights $R_w(\mathbf{w}) = \|\mathbf{w}\|_2^2$. The derivatives of $L(\cdot, \cdot)$ can be computed straightforwardly with the formulas presented in Section 3.1.

The term $R_q(\mathbf{q})$ is used to avoid overfitted solutions for the control points. In fact, its presence is the major difference with respect to previous attempts at implementing SAFs in neural networks [17], wherein overfitting was counterbalanced by choosing a large value for Δx , which in a way goes outside the philosophy of spline interpolation itself. At the same time, choosing a proper form for the regularization term is non-trivial, as the term should be cheap to compute, and it should introduce just as much *a priori* information as needed, without hindering the training process. Most of the literature on regularizing \mathbf{w} cannot be used here, as the corresponding formulations do not make sense in the context of spline interpolation. As an example, simply penalizing the ℓ_2 norm of \mathbf{q} leads to functions close to the zero function, while imposing sparsity is also meaningless.

For the purpose of this paper, we consider the following ‘damping’ criterion:

$$R_q(\mathbf{q}) = \|\mathbf{q} - \mathbf{q}_o\|_2^2, \quad (15)$$

where \mathbf{q}_o represents the initial values for the control points, as discussed in the previous section (without considering additional noise). The criterion makes intuitive sense as follows: while for \mathbf{w} we wish to penalize unwanted deviations from very small weights (which can be justified with arguments from learning theory), in the case of \mathbf{q} we

are interested in penalizing changes with respect to a ‘good’ function parameterized by the initial control points \mathbf{q}_0 , namely one of the standard AFs used in NN training. In fact, setting a value for λ_q very high essentially deactivates the adaptation of the control points. Clearly, other choices are possible, and in this sense this paper serves as a starting point for further investigations towards this objective. As an example, we may wish to penalize first (or second) order derivatives of the splines in order to force a desired level of smoothness [18].

3.4 Remarks on the implementation

In order to be usable in practice, SAFs require an efficient implementation to compute outputs and derivatives concurrently for the entire training dataset or, alternatively, for a properly chosen mini-batch (in the case of stochastic optimization algorithms). To begin with, we underline that the equations for the reference vector (see (3)) do not depend on the specific neuron, and for this reason they can easily be vectorized layer-wise on most numerical algebra libraries to obtain all vectors concurrently. Additionally, the indexes and relative terms $\mathbf{B}\mathbf{q}_i$ in (4) can be cached during the forward pass, to be reused during the computation of the derivatives. In this sense, the outputs of a layer *and* its derivatives can be computed by one 4×4 matrix-vector computation, and three 4-dimensional inner products, which have to be repeated for every pair input/neuron. In our experience, the cost of a relatively well-optimized implementation does not exceed twice that of a standard network for medium-sized batches, where the most onerous operation is the reshaping of the gradients in (9) and (13) into a single vector of gradients relative to the global vector \mathbf{q} .

4 Experimental results

4.1 Experimental setup

To evaluate the preliminary proposal, we consider two simple regression benchmarks for neural networks, the ‘chemical’ dataset (included among MATLAB’s testbeds for function fitting), and the ‘California Housing’.² They have respectively 498 and 20640 examples, and 8 numerical features. Inputs are normalized in the $[-1, +1]$ range, while outputs are normalized in the $[-0.5, +0.5]$ range. We compare a NN with 5 hidden neurons and $\tanh(\cdot)$ AFs (denoted as ‘Standard’ in the results), and a NN with the same number of neurons and SAF nonlinearities. The weight vector \mathbf{w} is initialized with the method described in [4]. Each SAF is initialized from a $\tanh(\cdot)$ nonlinearity, and control points are defined in the $[-2, +2]$ range with $\Delta x = 0.2$, which is a good compromise between locality of the SAFs and the overall number of adaptable parameters. For the first scenario, λ_q is kept to a small value of 10^{-5} . For each experiment, a random 30% of the dataset is kept for testing, and results are averaged over 15 different splits to average out statistical effects. Error is computed with the Normalized Root Mean-Squared Error (NRMSE). The optimization problems are solved using a freely available MATLAB implementation of the Polack-Ribiere variant of the nonlinear conjugate

²http://www.dcc.fc.up.pt/~ltorgo/Regression/cal_housing.html

Table 1. Average results for scenario 1 ($\lambda_w = 1$), together with one standard deviation.

Dataset	Nonlinearity	Tr. RMSE	T.st NRMSE
Chemical	Standard	1.00 ± 0.00	1.00 ± 0.01
	SAF	0.29 ± 0.02	0.31 ± 0.02
Calhousing	Standard	1.02 ± 0.00	1.01 ± 0.01
	SAF	0.56 ± 0.01	0.57 ± 0.02

gradient optimization algorithm by C.E. Rasmussen. [12].³ The optimization process is allowed 1500 maximum iterations. MATLAB code for the experiments is also available on the web.⁴ We briefly remark that the MATLAB library, apart from repeating the experiments presented here, is also designed to handle networks with more than a single hidden layer, and implements the ADAM algorithm [8] for stochastic training in case of a larger dataset.

4.2 Scenario 1: strong underfitting

As a first example, we consider a scenario of strong underfitting, wherein the user has misleadingly selected a very large value of $\lambda_w = 1$, leading in turn to extremely small values for the elements of \mathbf{w} after training. Results in terms of training and test RMSE are provided in Tab. 1. Since the activations of the NN tend to be very close to 0 (where the hyperbolic tangent operates in an almost-linear regime), standard NNs have a constant zero output, leading to a RMSE of 1. Nonetheless, SAF networks are able to reach a very satisfactory level of performance, which in the first case is almost comparable to that of a fully optimized network (see the following section).

To show the reasons for this, we have plotted four representative nonlinearities after training in Fig. 1. It is easy to see that the nonlinearities have adapted to act as ‘amplifiers’ for the activations in their operating regime, with mild and strong peaks around 0. Of particular interest is the fact that the resulting SAFs need not be perfectly centered around 0 (e.g. Fig. 1c), or even symmetrical around the y -axis (e.g. Fig. 1d). In fact, the splines are able to efficiently counterbalance a bad setting for the weights, with behaviors which would be very hard (or close to impossible) using standard setups with fixed, shared, mild nonlinearities.

4.3 Scenario 2: well-optimized parameters

In our second scenario, we consider a similar comparison with respect to before, but we fine-tune the parameters of the two methods using a grid-search with a 3-fold cross-validation on the training data as performance measure. Both λ_w and λ_q (only for the

³<http://learning.eng.cam.ac.uk/car1/code/minimize/>

⁴<https://bitbucket.org/ispamm/spline-nn>

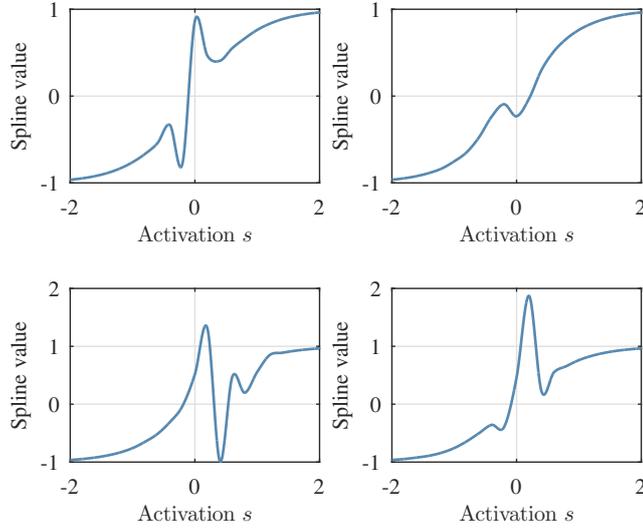


Fig. 1. Non-trivial representative SAFs after training for scenario 1.

Table 2. Optimal parameters (averaged over the runs) found by the grid-search procedure for scenario 2.

Dataset	Nonlinearity	λ_w	λ_q
Chemical	Standard	10^{-3}	—
	SAF	10^{-2}	10^{-4}
Calhousing	Standard	10^{-4}	—
	SAF	10^{-3}	10^{-4}

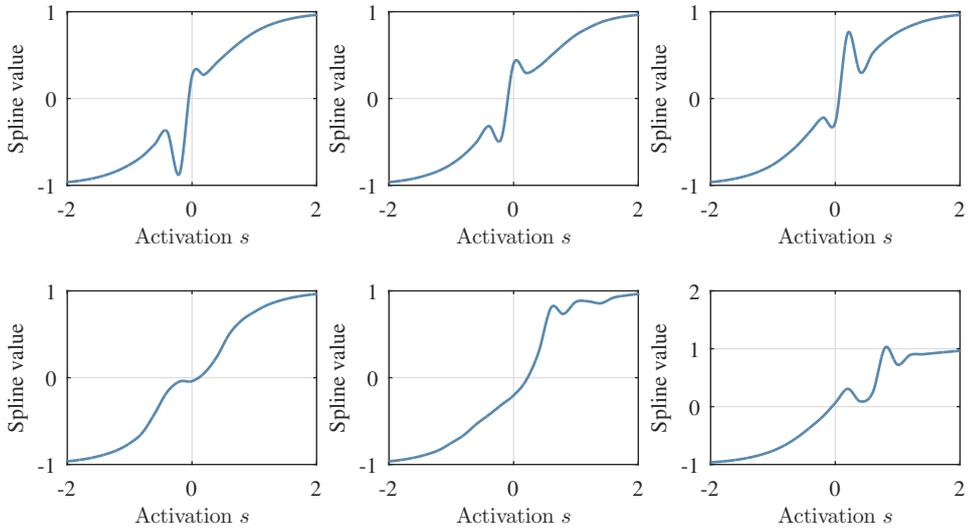
proposed algorithm) are searched in an exponential interval 2^j , with $j = -10, \dots, 5$. Optimal parameters found by the grid-search are listed in Table 2, while results in terms of training and test NRMSE are collected in Table 3.

Overall, we see that the NNs endowed with the SAF nonlinearities are able to surpass by a large margin a standard NN, and the results from the previous scenario. The only minor drawback evidenced in Table 3 is that the SAF network has some overfitting occurring in the ‘chemical’ dataset (around 2 points of NRMSE), showing that there is still some room for improvement in terms of spline optimal regularization.

Also in this case, we plot some representatives SAFs after training (taken among those which are not trivially identical to the tanh nonlinearity) in Fig. 2. As before, in general SAFs tend to provide an amplification (with a possible change of sign) of their activation around some region of operation. It is interesting to observe that, also in this case, the optimal shape need not be symmetric (e.g. Fig. 2a), and might even be far

Table 3. Average results for scenario 2 (fine-tuning for parameters), together with one standard deviation.

Dataset	Nonlinearity	Tr. RMSE	T.st NRMSE
Chemical	Standard	0.32 ± 0.01	0.32 ± 0.02
	SAF	0.26 ± 0.01	0.28 ± 0.02
Calhousing	Standard	0.55 ± 0.01	0.55 ± 0.01
	SAF	0.51 ± 0.02	0.51 ± 0.02

**Fig. 2.** Non-trivial representative SAFs after training for scenario 2.

from centered around 0 (e.g. Fig. 2c). Resulting nonlinearities can also present some additional non-trivial behaviors, such as a small region of insensitivity around 0 (e.g. Fig. 2d), or a region of pre-saturation before the actual tanh saturation (e.g. Fig.s 2e-2f).

5 Conclusion

In this paper, we have presented a principled way to adapt the activation functions of a neural network from training data, locally and independently for each neuron. Particularly, each nonlinearity is implemented with cubic spline interpolation, whose control points are adapted in the optimization phase. Overfitting is controlled by a novel ℓ_2 regularization criterion avoiding unwanted oscillations. Albeit efficient, this criterion does constrain the shapes of the resulting functions by a certain degree. In this sense,

the design of more advanced regularization terms is a promising line of research. Additionally, we plan on exploring the application of SAFs to deeper networks, where it is expected that the statistics of the neurons' activations can change significantly layer-wise [4].

References

1. Agostinelli, F., Hoffman, M., Sadowski, P., Baldi, P.: Learning activation functions to improve deep neural networks. arXiv preprint arXiv:1412.6830 (2014)
2. Chandra, P., Singh, Y.: An activation function adapting training algorithm for sigmoidal feed-forward networks. *Neurocomputing* 61, 429–437 (2004)
3. Chen, C.T., Chang, W.D.: A feedforward neural network with function shape autotuning. *Neural networks* 9(4), 627–641 (1996)
4. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: *Int. conf. on artificial intell. and stat.* pp. 249–256 (2010)
5. Goh, S., Mandic, D.: Recurrent neural networks with trainable amplitude of activation functions. *Neural Networks* 16(8), 1095–1100 (2003)
6. Guarnieri, S., Piazza, F., Uncini, A.: Multilayer feedforward networks with adaptive spline activation function. *IEEE Trans. Neural Netw.* 10(3), 672–683 (1999)
7. Haykin, S.: *Neural networks and learning machines*. Pearson Education, 3rd edn. (2009)
8. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. In: *3rd International Conference for Learning Representations* (2015), arXiv preprint arXiv:1412.6980
9. Lin, M., Chen, Q., Yan, S.: Network in network. arXiv preprint arXiv:1312.4400 (2013)
10. Ma, L., Khorasani, K.: Constructive feedforward neural networks using hermite polynomial activation functions. *IEEE Trans. Neural Netw.* 16(4), 821–833 (2005)
11. Piazza, F., Uncini, A., Zenobi, M.: Artificial neural networks with adaptive polynomial activation function. In: *Int. Joint Conf. on Neural Networks*. vol. 2, pp. II–343. *IEEE/INNS* (1992)
12. Rasmussen, C.: *Gaussian processes for machine learning*. MIT Press (2006)
13. Scarpiniti, M., Comminiello, D., Parisi, R., Uncini, A.: Nonlinear spline adaptive filtering. *Signal Process.* 93(4), 772–783 (2013)
14. Scarpiniti, M., Comminiello, D., Scarano, G., Parisi, R., Uncini, A.: Steady-state performance of spline adaptive filters. *IEEE Trans. Signal Process.* 64(4), 816–828 (2016)
15. Schmidhuber, J.: Deep learning in neural networks: An overview. *Neural Networks* 61, 85–117 (2015)
16. Trentin, E.: Networks with trainable amplitude of activation functions. *Neural Networks* 14(4), 471–493 (2001)
17. Vecci, L., Piazza, F., Uncini, A.: Learning and approximation capabilities of adaptive spline activation function neural networks. *Neural Networks* 11(2), 259–270 (1998)
18. Wahba, G.: *Spline models for observational data*. Siam (1990)
19. Wang, Y., Shen, D., Teoh, E.: Lane detection using catmull-rom spline. In: *IEEE Int. Conf. on Intell. Vehicles*. pp. 51–57 (1998)
20. Zhang, M., Xu, S., Fulcher, J.: Neuron-adaptive higher order neural-network models for automated financial data modeling. *IEEE Trans. Neural Netw.* 13(1), 188–204 (2002)