Step size self-adaptation for SGD

Ilona Kulikovskikh and Tarzan Legović

Abstract—Convergence and generalization are two crucial aspects of performance in neural networks. When analyzed separately, these properties may lead to contradictory results. Optimizing a convergence rate yields fast training, but does not guarantee the best generalization error. To avoid the conflict, recent studies suggest adopting a moderately large step size for optimizers, but the added value on the performance remains unclear. We propose the LIGHT function with the four configurations which regulate explicitly an improvement in convergence and generalization on testing. This contribution allows to: 1) improve both convergence and generalization of neural networks with no need to guarantee their stability; 2) build more reliable and explainable network architectures with no need for overparameterization. We refer to it as step size self-adaptation.

Index Terms—step size, activation function, non-monotonicity, adaptive optimization, stochastic gradient descent

I. INTRODUCTION

NEURAL networks imitate signal transmission within neurons in the brain with units which are interconnected through weighted links and assembled in layers [1], [2]. Training a neural network implies updating the model weights to best map inputs to outputs. This process is framed as an optimization problem that involves minimizing the model errors on a training dataset. When training a network with gradient-based methods, accelerating convergence to the solution is of a high priority [3], [4], but not the only performance variable to optimize. Minimizing the difference between the model errors on a training and a testing dataset, which is called the generalization error, plays a fundamental role [5]–[11].

The iterative optimization schemes with an adaptive step size schedule converge faster [12]–[15], but generalize poorly [16]–[20]. They are often outperformed by non-adaptive stochastic gradient descent (SGD) [21] for over-parameterized neural networks, where the number of trainable parameters is much higher than the number of samples they are trained on. Exploring critical generalization capacity, several studies explained this phenomenon: overparametrization ensures faster convergence [22]–[28] while inducing implicit regularization of the original problem, which can potentially ease the minimization of the generalization error [29]–[35]. However, overparameterized models require an enormous number of units and layers to represent, process, and store data. This heavily reduces the transparency of neural networks, making them difficult to interpret.

I. Kulikovskikh is with the Department of Information Systems and Technologies, Samara University, Samara 443086, Russia (e-mail: kulikovskikh.im@ssau.ru).

T. Legović is with the Institute of Applied Ecology, Oikon Ltd., with Libertas International University, and with Ruđer Bošković Institute, Zagreb 10000, Croatia (email: tlegovic@oikon.hr)

Manuscript received xx xx, xxxx; revised xx xx, xxxx. (Corresponding author: Ilona Kulikovskikh).

What makes neural networks generalize well? Relying on the extensive empirical studies of SGD, it became evident that the step size maximizing the test accuracy is usually larger than the step size which minimizes the training loss [36], [37], [65]. The occurrence of an implicit regularizer demystifies this matter as well. For a small step size, SGD behaves similar to GD on the full batch loss function. When a step size increases, the regularizer starts penalizing the mean Euclidean norm of the minibatch gradients [33], [34], [37] that makes the training loss non-monotonic. Another explanation is that faster gradient descent methods naturally generate chaotic dynamical systems [38], which bring the optimizer to the edge of stability [36], [37] and, thus, yield a non-monotone decrease pattern in the loss function. This finding shares some similarities with the edge of chaos concept [39]. According to the concept, deep networks may be trained only sufficiently close to criticality, avoiding the regions of vanishing and exploding gradients, which correspond to the ordered and chaotic phases, respectively.

Answering the research question, we hypothesize that non-monotonicity of the loss function itself may establish faster convergence and better generalization bounds, even without adaptive optimization tricks. Our reasoning comes from the thorough review of literature on the positive effect of non-monotonicity [40], [41] in neural networks caused by discontinuity [42]–[44], delays [41], [45]–[49], differential inclusions, [50], [51], sliding modes [49], [52] on the global finite-time convergence and stability [53]–[56].

To examine the hypothesis, we simulate non-monotonicity in neural networks with the LIGHT (LogIstic Growth with HarvesTing) activation function, which originates from population dynamics [57], [58] and behaves as follows. It starts growing with the rate r by the logistic law. At the time $t = T$, it starts declining with the rate E . The y-intercepts at the moments 0 and T are specified. The default step size of SGD is modified with regard to r and E , respectively. For a diagnostic purpose, we suggest four configurations of the function to regulate explicitly an improvement in convergence and generalization on testing (see Fig. 1):

- default-: no improvement;
- r-: an improvement in convergence;
- E-: an improvement in generalization;
- Er-: both convergence and generalization are improved.

Increasing the rate r allows us to push a learning system towards the edge of stability. By increasing the rate E , we fix this edge by pushing the system towards the equilibrium point. The presence of sliding modes along discontinuity surfaces, which are modelled with T , establishes the control over the system behavior.

We show that the formulated hypothesis is valid and that the LIGHT function contributes to: 1) improving convergence and

Fig. 1. The balance between convergence and generalization for the LIGHT configurations

generalization by training neural networks with a moderately large step size with no need to ensure its stability; 2) building more reliable and explainable network architectures with no need for overparameterization. We refer to it as step size self-adaptation.

II. RELATED WORK

A. Adaptive optimization

SGD is one of the most dominant first-order optimization algorithms for training neural networks [21]. Albeit its popularity and simplicity, SGD scales the gradient equally in all directions that results in worse convergence than the adaptive methods, for example, Adam [14] and Adagrad [12] that scale the gradient using the information from past gradients [16]. Adopting optimization methods with variable step sizes leads to faster convergence and worse generalization compared to non-adaptive methods [15], [59], [60]: they train faster, but their performance plateaus on testing due to less stable and predictable behavior. Consequently, further development of optimization techniques is directed towards a better trade-off between convergence and generalization.

In order to address the above problem, the first group of studies is directed towards different SGD modifications such as SGD with adopting extended differentiators [61], random reshuffling [62], local changes in gradients [63], [64] and etc.

The analysis of SGD based optimization for overparameterized models has recently become another active area of research interest [4], [20], [22], [24], [34], [65]. Recent studies indicated that large step sizes can preserve good generalization and accelerate SGD convergence without any additional gradient scaling. While analyzing the effect of overparametrization, Wu et al. [34] pointed to the difference in directional biases for SGD and GD with a moderate and annealing step size. Vaswani et al. [66] explored line-search techniques and provided heuristics to automatically set larger step sizes. Li and Arora [65] carried out an analysis of an exponential step size schedule. They showed that using SGD with momentum [26] and an exponentially increasing step size, coupled with batch normalization, maintains good balance between convergence and generalization across all standard architectures. Nitanda and Suzuki [67] provided an analysis of a convergence rate for the averaged SGD in the Neural Tangent Kernel Regime. The authors disclosed the conditions on which the method can achieve the minimax optimal convergence rate, with the global convergence guarantee.

In parallel with more successful SGD adoption for overparameterized models, substantial progress has been achieved in optimization methods with adaptive step sizes. SGDP

and AdamP use effective step sizes without changing the update directions [19]. This allows to preserve the original convergence properties of GD optimizers. RAdam adopts the learning rate warmup heuristic to rectify the variance of adaptive step sizes [26] and, by that, stabilize training, accelerate convergence, and improve generalization. In an attempt to balance generalization and convergence on unstable and extreme step sizes, Luo et al. [16] put forward AdaBound and AMSBound which adopt dynamic bounds on step sizes to eliminate the generalization gap between adaptive methods and SGD and maintain higher learning rate early in the training. These methods were further developed with regard to a dynamic decay rate in [17]. Xie et al. [18] proved that the normalized Adagrad ensures robustness to the choice of hyper-parameters and achieves a linear convergence rate for a subset of either strongly convex functions or non-convex functions that satisfy the Polyak-Lojasiewicz (PL) inequality. Zhou et al. [20] proposed the SAGD method that leverages differential privacy to boost the generalization performance of adaptive gradient methods.

B. Adaptive activation

Using adaptive activation functions in neural networks is one more way to balance convergence and generalization. The first activation function presented “all-or-none” character of nervous activity with a step function [68], [69] to solve a binary classification problem. Wilson and Cowan [70] derived coupled nonlinear differential equations from the dynamics of spatially localized populations containing both excitatory and inhibitory model neurons. They investigated population responses to various types of stimuli and introduced an s-shaped monotonic function of stimulus intensity - a sigmoid function. Yamada and Yabuta [71] suggested an approach to optimally tune the shape of the sigmoid function in control systems. While comparing the approximation capabilities of activation functions, DasGupta and Schnitger [72] pinpointed that the standard sigmoid is more powerful than the binary threshold even when computing boolean functions [73]. Piazza et al. [74] proposed the adaptive polynomial activation function to address the issue of complexity in neural networks. Xu and Zhang [75] proposed another adaptive activation function to reduce a network size. Goh and Mandic [76] suggested to adapt the amplitude of activation functions, while reconsidering recurrent neural networks in terms of nonlinear adaptive filters. Bai et al. [77] showed that varying the slope of an activation function with different step sizes is more beneficial than using momentum and an adaptive step size in the backpropagation algorithm. Flennerhag [78] suggested simple drop-in replacements that learn to adapt their parameterization with regard to the network inputs. PPolyNets [79] are accurate and efficient parametric polynomial activations specifically developed for encryption schemes which support only polynomial operations. Goyal et al. [80] suggested to normalize polynomial activations to increase the stability of neural networks. Kunc and Klěma [81] proposed a novel transformative adaptive activation function that improves the gene expression inference by generalizing existing adaptive activation functions.

De Felice et al. [40] drew attention to the fact that the biological activation function has a more complicated behavior which reduces to the usual (step or sigmoid) function for some hyperparameters describing its shape and stated that the non-monotonicity of the function increases the capacity of neural networks. Baldi and Atiya [45] extended previously known results regarding the effects of delays on stability and convergence properties. Forti and Nistri [53] introduced a general class of neural networks, where the neuron activations are modeled by discontinuous functions. The authors discovered that the presence of sliding modes ensures global convergence in neural networks in finite time. Duan et al. [49] established the existence and global exponential stability of almost periodic solution for the delayed high-order Hopfield neural networks. The study [46] discussed the dynamics of a class of the delayed neural networks with discontinuous activation functions. The authors concluded that the solution of delayed neural networks with discontinuous activation functions can be regarded as a limit of the solutions of delayed neural networks with high-slope continuous activation functions. According to [82], [83], rectified linear units (ReLU) and their different modifications [84], [85] in the hidden layers of neural networks demonstrate better convergence and generalization in comparison with the continuous activations. Exploring monostability and multistability of almost-periodic solutions in the fractional-order neural networks, Wan et al. [86] indicated that the dynamics in neural networks with the unsaturating piecewise linear activation functions is more complex. Nie and Zheng [41] looked into the problem of coexistence and dynamical behaviors of multiple equilibrium points for neural networks with discontinuous non-monotonic piecewise linear activation functions and time-varying delays. The study revealed that discontinuous neural networks can have greater storage capacity than the continuous ones. Hayou et al. [87], however, mentioned that only a specific choice of hyperparameters such as initialization and activation with regard to the concept of chaos [39] improves convergence and generalization in neural networks.

III. PRELIMINARIES

For a dataset $\{x_i, y_i\}_{i=1}^m$ with $x_i \in \mathbb{R}^n$, $y_i \in \{-1, 1\}$, we minimize an empirical loss function for each mini-batch dataset $B(t) \subseteq \{1, \dots, m\}$ with a weight vector $\Theta \in \mathbb{R}^n$:

$$\mathcal{L}(\Theta) = \sum_{i \in B(t)} \ell(y_i \langle \Pi(\Theta), x_i \rangle), \quad (1)$$

where ℓ measures the discrepancy between the output y and the model prediction. The SGD optimizer finds the weight vector with a fixed step size η :

$$\Theta_l(t+1) = \Theta_l(t) - \eta \nabla_{\Theta_l} \mathcal{L}(\Theta), \quad (2)$$

where $\Pi(\Theta) = \Theta_1 \times \Theta_2 \times \dots \times \Theta_L$, $\Theta = \{\Theta_l \in \mathbb{R}^{d_{l-1} \times d_l} : l = 1, 2, \dots, L\}$, L is the number of layers, d_l is the number of nodes in the layer l .

IV. LIGHT

We built a diagnostic function $\ell^{r,E}(t)$ on the standard sigmoid by simulating different types of non-monotonicity with the growth rate r and decline rate E . The function grows with a constant rate r according to the logistic law. After time T , it declines with a constant rate E . We call the function LIGHT (LogIstic Growth with HarvesTing) as its behavior inherits the principles of population dynamics [57], [58]. Let us present the LIGHT function.

Definition (LIGHT) For any time $t \in \mathbb{R}$, time instant $T > 0$, growth rate $r > 0$ and decline rate $E \geq 0$, a non-monotonic function $\ell^{r,E}(t)$, such that

$$\lim_{t \rightarrow -\infty} \ell^{r,E}(t) = 0, \quad \lim_{t \rightarrow \infty} \ell^{r,E}(t) = \varepsilon,$$

where ε is the extent to which r is impacted by E , behaves as:

$$\ell^{r,E}(t) = \begin{cases} \varepsilon e^{(\ln_q N_T + \frac{E}{r})e^{-r(t-T)}} & \text{if } t \geq T, \\ \varepsilon e^{(\ln_q N_0)e^{-rt}} & \text{otherwise,} \end{cases} \quad (3)$$

with the derivative $\ell'^{r,E}(t)$:

$$\begin{cases} -\varepsilon r (\ln_q N_T + \frac{E}{r}) e^{-r(t-T) + (\ln_q N_T + \frac{E}{r})e^{-r(t-T)}} & \text{if } t \geq T, \\ -\varepsilon r \ln_q N_0 e^{-rt + \ln_q N_0 e^{-rt}} & \text{otherwise,} \end{cases}$$

where $N_0 = \ell^{r,E}(0)$, $N_T = \ell^{r,E}(T)$, $\ln_q(x)$ is the q -logarithm, where q is the rate with which the function grows when smaller.

By introducing q , we move from the infinitesimal calculus to quantum calculus [88]–[91] to avoid the concept of limits and, thus, simplify the definition. Another justification of this parameter is from the point of explainability: it generalizes the Verhulst ($q = 1$, **light-v**) [57], [92], [93] and Gompertz ($q \rightarrow 0$, **light-g**) [94]–[96] laws of population dynamics [58, see Section 7.2.3].

To introduce non-monotonicity in the LIGHT diagnostic function, we propose four different configurations (see Fig. 2). We augmented the -default- configuration, which reduces to **sigmoid** if $q = 1$ and $N_0 = 0.5$, with three more configurations which regulate an improvement in convergence and generalization:

- default**-: $r = 1$, $E = 0$; no improvement;
- r**-: $r \uparrow$, $E = 0$; an improvement in convergence;
- E**-: $r = 1$, $E \uparrow$; an improvement in generalization;
- Er**-: $r \uparrow$, $E \uparrow$; an improvement in both convergence and generalization.

The sign $\uparrow$ means that value of the parameter is significantly increased. The decline in growth, delayed by T , induces a discontinuity in the function. A simultaneous increase in r and E makes the discontinuity even more noticeable by scaling the function magnitude that results in a greater impact on a convergence/generalization trade-off. To diagnose a neural network capability with the LIGHT function, we modified a step size η in the SGD optimizer (2) by replacing $\ell(t)$ with $\ell^{r,E}(t)$ in the loss (1).

Fig. 2. The LIGHT configurations and their derivatives subject to $T = 0.75$. (a), (b) -default- configuration: $r = 1$, $E = 0$, $N_0 = 0.5$ (in blue). (c), (d) -r- configuration: $r = 3$, $E = 0$, $N_0 = 0.3$ (in red). (e), (f) -E- configuration: $r = 1$, $E = 3$, $N_0 = 0.3$ (in green). (g), (h) -Er- configuration: $r = 3$, $E = 4$, $N_0 = 0.3$ (in black).

V. EXPERIMENTAL RESULTS

A. Experimental Setup

As noted in the surveyed literature on adaptive optimization and activation, the success in balancing convergence and generalization is often attributed to the complexity and capacity of neural networks. To clearly observe the distinctive contribution of the LIGHT function to the step size adaptation, we focused primarily on creating the simplest network architecture - a single neuron $L = 0$ [97]–[103], the capacity of which recently drives the renewal interest in neural networks [104]–[107]. To investigate how a small increase in the network complexity, without overparameterization, may impact the proposed instrument, we also complemented the model with a hidden layer $L = 1$ of the ReLU neurons $d_l = 5$, where the LIGHT function is applied only to the output.

We compared four non-adaptive methods - SGD with the -default- configuration (**sigmoid-sgd**) and SGD with the -r-, -E-, -Er- configurations (**light-v-sgd** and **light-g-sgd**) - to two popular adaptive methods with the -default- configuration - Adam (**sigmoid-adam**) and AdaGrad (**sigmoid-adagrad**). For all the optimizers, we used the default parameters, batch size $|B(t)| = 75$, and $n_{\text{epoch}} = 1500$. The number of runs were equal to 10.

The light-v and light-g hyperparameters were optimized with a random search [108] with a 2.5% random pick of all possible combinations from the full grid space within the following ranges: $r \in [0.1, 20]$ with the number of points $n_r = 5$; $E \in [0, 20]$, $n_E = 5$; $T \in [0, 3]$, $n_T = 3$; $N_T \in [0.2, 0.8]$, $n_{N_T} = 5$. The number of epochs for hyperparameter search was equal to 1. As we optimized the hyperparameters with a small percentage of random combinations and one epoch, we added redundancy to the experimental setup by implementing both variations of the light function (light-v and light-g). This allowed us to validate the consistency of the chosen hyperparameters and to distinguish the examples where the rates r and E were properly balanced.

The LIGHT function in two variations was implemented as a custom output activation layer with Keras `class LIGHT(Layer)`. The layer controls a convergence and generalization trade-off with the LIGHT configurations. The

code is available at the repository: <https://github.com/yukinoi/light-diagnostic-function>.

B. Synthetic Data

We generated a set of synthetic linearly separable and non-separable datasets ($m = 1000$, $n = 2$) with lower and higher levels of variance (see Fig. 3). The datasets were randomly split into training (80%) and testing (20%) subsets.

Fig. 3. Synthetic datasets. (a) BLOBS: $cluster\ std = 0.25$ (lower variance), (b) BLOBS: $cluster\ std = 0.5$ (higher variance). (c) XOR: $cluster\ std = 0.45$ (lower variance), (d) XOR: $cluster\ std = 0.9$ (higher variance). (e) CIRCLES: $noise = 0.1$ (lower variance), (f) CIRCLES: $noise = 0.25$ (higher variance). (g) MOONS: $noise = 0.1$ (lower variance), (h) MOONS: $noise = 0.25$ (higher variance).

For conciseness, all the plots on synthetic datasets for different combinations of network architectures are deferred to Appendix A. Figures A1-A16 (a), (c), (e) depict the test accuracy curves for the -r-, -E-, -Er- configurations (light-v and light-g) on the synthetic datasets with lower and higher variance. For the sake of comparison, we added the -default- configuration of non-adaptive (sigmoid-sgd) and adaptive (sigmoid-adam and sigmoid-adagrad) optimizers to each plot. The presented results comply with the expected behavior of curves given in Figure 1. The curves on CIRCLES with one neuron A10-A10 seem different as they reach the maximum accuracy in a few epochs in comparison with the other datasets. However, they behave similarly within these epochs.

Tables I-IV summarize the maximum test accuracy and the number of epochs needed to reach it for each dataset. The best balance between the maximum accuracy and the number of epochs is highlighted in bold. We can observe that both light-v-sgd and light-g-sgd outperform other optimizers with the -default- configuration. The differences between the light variations point to the apparent discrepancy in the balance between r and E occurred due to the setting of hyperparameter optimization. When r and E are not fully balanced, we can also observe that the -r- or -E- configuration become slightly superior to the -Er- configuration (see Table I, light-g-sgd; Table II, light-v-sgd; Table III, light-g-sgd).

The optimal light-v and light-g hyperparameters for different configurations are demonstrated in Figures A1-A16 (b), (d), (f). By analyzing the boxplots, we see that a non-zero decline rate in the -Er- configuration increases the growth rate r compared to the value r in the -r- configuration. It also brings more stable results as the standard deviation of accuracy curves is substantially reduced.

TABLE I
MAXIMUM ACCURACY AND NUMBER OF EPOCHS ON BLOBS DATASET

Method	Configuration	L = 0				L = 1			
		Lower variance		Higher variance		Lower variance		Higher variance	
		Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}
sigmoid-adam	-default-	100	349	97.45	531	100	88	97.2	113
sigmoid-adagrad	-default-	99.7	1480	96.25	1488	100	196	97.2	472
sigmoid-sgd	-default-	100	112	97.5	330	100	69	97.4	232
light-v-sgd	-r-	95.65	1469	94.55	1478	100	726	96.8	1206
	-E-	100	127	97.5	311	100	150	97.4	262
	-Er-	100	11	97.5	47	100	67	96.8	818
light-g-sgd	-r-	99.9	769	96.7	1139	100	760	96.95	1206
	-E-	100	56	97.5	137	100	69	97.45	78
	-Er-	94.85	7	96.8	64	100	94	96.65	221

TABLE II
MAXIMUM ACCURACY AND NUMBER OF EPOCHS ON XOR DATASET

Method	Configuration	L = 0				L = 1			
		Lower variance		Higher variance		Lower variance		Higher variance	
		Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}
sigmoid-adam	-default-	60.5	401	55.2	240	99.95	1420	88.3	450
sigmoid-adagrad	-default-	59.1	1493	54.85	1358	97.15	1280	91.75	1486
sigmoid-sgd	-default-	60.5	283	55.4	198	99.6	296	90.05	1163
light-v-sgd	-r-	62.6	18	57.6	14	99.95	1420	92.4	425
	-E-	62.3	162	56.4	159	99.95	554	89.3	1256
	-Er-	68.25	8	57.85	14	99.7	58	92.1	462
light-g-sgd	-r-	68.7	6	58.35	10	94.55	1489	87.7	1106
	-E-	59.75	126	56.25	151	97.05	171	90.3	1417
	-Er-	70.75	6	60.85	11	89.15	147	91.55	580

TABLE III
MAXIMUM ACCURACY AND NUMBER OF EPOCHS ON CIRCLES DATASET

Method	Configuration	L = 0				L = 1			
		Lower variance		Higher variance		Lower variance		Higher variance	
		Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}
sigmoid-adam	-default-	51.45	119	52.6	95	92.75	625	87.15	443
sigmoid-adagrad	-default-	52.44	0	51.3	710	99.65	1423	85.7	1143
sigmoid-sgd	-default-	49.35	44	51.5	237	99.95	627	87.55	743
light-v-sgd	-r-	49.8	2	52.4	3	93.65	1023	77.55	1499
	-E-	46.65	245	51.55	251	100	1232	87	1177
	-Er-	54.1	1	50.7	0	99.95	388	85.6	88
light-g-sgd	-r-	60.7	1	56.1	0	100	102	82.65	514
	-E-	55.05	34	54.75	72	96.25	1259	86.75	1287
	-Er-	61.55	2	55.25	0	100	69	87.56	34

TABLE IV
MAXIMUM ACCURACY AND NUMBER OF EPOCHS ON MOONS DATASET

Method	Configuration	L = 0				L = 1			
		Lower variance		Higher variance		Lower variance		Higher variance	
		Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}	Accuracy, %	n_{epoch}
sigmoid-adam	-default-	88.5	639	85.55	617	97.2	849	92.75	1275
sigmoid-adagrad	-default-	85.45	1486	82.8	1477	91.65	1459	87.7	1469
sigmoid-sgd	-default-	88.5	436	85.75	428	94.1	1491	90.9	1496
light-v-sgd	-r-	84.35	1475	84	1461	95.95	1407	91.75	900
	-E-	88.5	458	85.75	444	93.6	1490	90	1493
	-Er-	88.5	47	85.75	36	97.95	270	94.45	624
light-g-sgd	-r-	86.7	1477	83.55	1481	99.9	1469	93.5	704
	-E-	88.5	297	84	46	99.45	1339	92.35	1495
	-Er-	87.2	31	83.65	1455	98.75	687	93.65	1276

When optimizing r on training (validation), the result does not deliver good generalization. This means that the system is pushed towards the edge of stability, which is not clearly defined. By increasing E , we fixed the edge of stability, improving generalization. In addition, it allowed us to shift the region for picking r to the right, allowing for more extreme values, and, thus, accelerating convergence.

To underline the LIGHT benefits in trading off convergence and generalization, we also analyzed the number of epochs at a test accuracy threshold (see Tables V-VIII). The lowest number of epochs needed to reach the accuracy threshold is strengthened with bold font. The hyphen indicates that the accuracy threshold is not reached in 1500 iterations. We can see that the light-based SGD greatly outperforms other optimization methods.

TABLE V
NUMBER OF EPOCHS AT ACCURACY THRESHOLD ON BLOBS DATASET

Method	Config.	Acc. $\geq 95\%$, $L = 0$		Acc. $\geq 95\%$, $L = 1$	
		low. var.	high. var.	low. var.	high. var.
sigmoid-adam	-default-	160	150	35	62
sigmoid-adagrad	-default-	372	875	21	61
sigmoid-sgd	-default-	27	45	21	43
light-v-sgd	-Er-	4	6	2	7
light-g-sgd	-Er-	–	4	1	8

TABLE VI
NUMBER OF EPOCHS AT ACCURACY THRESHOLD ON XOR DATASET

Method	Config.	Acc. $\geq 60\%$, $L = 0$		Acc. $\geq 90\%$, $L = 1$	
		low. var.	high. var.	low. var.	high. var.
sigmoid-adam	-default-	290	–	36	–
sigmoid-adagrad	-default-	–	–	70	368
sigmoid-sgd	-default-	243	–	63	1100
light-v-sgd	-Er-	1	–	2	53
light-g-sgd	-Er-	0	4	–	23

TABLE VII
NUMBER OF EPOCHS AT ACCURACY THRESHOLD ON CIRCLES DATASET

Method	Config.	Acc. $\geq 55\%$, $L = 0$		Acc. $\geq 85\%$, $L = 1$	
		low. var.	high. var.	low. var.	high. var.
sigmoid-adam	-default-	–	–	69	221
sigmoid-adagrad	-default-	–	–	67	466
sigmoid-sgd	-default-	–	–	77	315
light-v-sgd	-Er-	–	–	2	51
light-g-sgd	-Er-	0	0	3	13

C. Application

We validated the proposed step size adaptation approach with the LIGHT diagnostic function on MNIST, Fashion MNIST, and CIFAR10 datasets. The labels of the image classification datasets were binarized with the target class {5}. The samples were randomly extracted ($m = 1000$) from each of them and split into training (80%) and testing (20%) subsets. To classify the images, we used the one-layer network

TABLE VIII
NUMBER OF EPOCHS AT ACCURACY THRESHOLD ON MOONS DATASET

Method	Config.	Acc. $\geq 85\%$, $L = 0$		Acc. $\geq 90\%$, $L = 1$	
		low. var.	high. var.	low. var.	high. var.
sigmoid-adam	-default-	244	310	256	500
sigmoid-adagrad	-default-	1306	–	780	–
sigmoid-sgd	-default-	47	172	580	1379
light-v-sgd	-Er-	2	15	24	29
light-g-sgd	-Er-	5	–	5	58

architecture and the light-g variation with the pre-defined growth and decline rate: $r = 4.08$, $E = 6.4$.

Figure 4 shows the accuracy curves on testing for the -Er- configuration. The number of epochs needed to reach the maximum accuracy and the test accuracy threshold for each image dataset are summarized in Tables IX, X, and XI. By analogy, bold font indicates the lowest number of epochs needed to reach the accuracy threshold. The hyphen shows that the threshold is not reached. As we can see, the provided results disclose the benefits of the proposed step size adaptation in managing the trade-off between convergence and generalization with reliable and explainable network architectures.

Fig. 4. Test accuracy on MNIST, Fashion MNIST, and CIFAR10 datasets for $L = 1$, $d_l = 5$, $m = 1000$

TABLE IX
EXPERIMENTAL RESULTS ON MNIST

Method	Acc., %	n_{epoch}	Acc. $\geq 97.2\%$
sigmoid-adam	97.2	1312	1312
sigmoid-adagrad	97.2	405	405
sigmoid-sgd	97.1	952	–
light-g-sgd	97.6	358	102

TABLE X
EXPERIMENTAL RESULTS ON FASHION MNIST

Method	Acc., %	n_{epoch}	Acc. $\geq 97.6\%$
sigmoid-adam	97.5	94	–
sigmoid-adagrad	97.4	473	–
sigmoid-sgd	97.6	1319	1319
light-g-sgd	97.8	258	235

VI. CONCLUSION

We contributed to the direction of SGD based optimization with step size self-adaptation. This technique allows to increase a step size by some fixed growing rate r which is self-stabilized with some fixed declining rate E in order to

TABLE XI
EXPERIMENTAL RESULTS ON CIFAR10

Method	Acc., %	n_{epoch}	Acc. $\geq 90\%$
sigmoid-adam	88.1	43	—
sigmoid-adagrad	88.2	213	—
sigmoid-sgd	88.1	67	—
light-g-sgd	90.5	887	0

ensure best balance between convergence and generalization. It equips the optimizer with a simple instrument for explicit control over convergence/generalization trade-off which is the key to building reliable network architectures.

Rather than suggesting another adaptive and non-monotone activation function, we put forward this instrument as a simple diagnostic function. The function adopts sliding modes in line with some fixed growing r and declining E rates to simulate discontinuities and explicitly regulate their influence on convergence and generalization. In addition, the LIGHT function relies on the laws of population dynamics as an original s-shaped monotonic function [70] but exhibits more complex behavior as noted in [40]. This means that the proposed self-adaptation mechanism may open up new opportunities for building not only reliable but explainable neural network architectures [109] with greater capacity.

APPENDIX PLOTS ON SYNTHETIC DATASETS

Fig. A1. Test accuracy and optimal hyperparameters on BLOBS (lower variance) for $L = 0$, $d_l = 0$, $m = 1000$, $n = 2$, $\text{cluster std} = 0.25$

Fig. A2. Test accuracy and optimal hyperparameters on BLOBS (higher variance) for $L = 0$, $d_l = 0$, $m = 1000$, $n = 2$, $\text{cluster std} = 0.5$

Fig. A3. Test accuracy and optimal hyperparameters on BLOBS (lower variance) for $L = 1$, $d_l = 5$, $m = 1000$, $n = 2$, $\text{cluster std} = 0.25$

Fig. A4. Test accuracy and optimal hyperparameters on BLOBS (higher variance) for $L = 1$, $d_l = 5$, $m = 1000$, $n = 2$, $\text{cluster std} = 0.5$

Fig. A5. Test accuracy and optimal hyperparameters on XOR (lower variance) for $L = 0$, $d_l = 0$, $m = 1000$, $n = 2$, $\text{cluster std} = 0.45$

Fig. A6. Test accuracy and optimal hyperparameters on XOR (higher variance) for $L = 0$, $d_l = 0$, $m = 1000$, $n = 2$, $\text{cluster std} = 0.90$

Fig. A7. Test accuracy and optimal hyperparameters on XOR (lower variance) for $L = 1$, $d_l = 5$, $m = 1000$, $n = 2$, $cluster\ std = 0.45$

Fig. A11. Test accuracy and optimal hyperparameters on CIRCLES (lower variance) for $L = 1$, $d_l = 5$, $m = 1000$, $n = 2$, $noise = 0.1$

Fig. A8. Test accuracy and optimal hyperparameters on XOR (higher variance) for $L = 1$, $d_l = 5$, $m = 1000$, $n = 2$, $cluster\ std = 0.9$

Fig. A12. Test accuracy and optimal hyperparameters on CIRCLES (higher variance) for $L = 1$, $d_l = 5$, $m = 1000$, $n = 2$, $noise = 0.25$

Fig. A9. Test accuracy and optimal hyperparameters on CIRCLES (lower variance) for $L = 0$, $d_l = 0$, $m = 1000$, $n = 2$, $noise = 0.1$

Fig. A13. Test accuracy and optimal hyperparameters on MOONS (lower variance) for $L = 0$, $d_l = 0$, $m = 1000$, $n = 2$, $noise = 0.1$

Fig. A10. Test accuracy and optimal hyperparameters on CIRCLES (higher variance) for $L = 0$, $d_l = 0$, $m = 1000$, $n = 2$, $noise = 0.25$

Fig. A14. Test accuracy and optimal hyperparameters on MOONS (higher variance) for $L = 0$, $d_l = 0$, $m = 1000$, $n = 2$, $noise = 0.25$

Fig. A15. Test accuracy and optimal hyperparameters on MOONS (lower variance) for $L = 1$, $d_l = 5$, $m = 1000$, $n = 2$, $\text{noise} = 0.1$

Fig. A16. Test accuracy and optimal hyperparameters on MOONS (higher variance) for $L = 1$, $d_l = 5$, $m = 1000$, $n = 2$, $\text{noise} = 0.25$

REFERENCES

- [1] G. Marcus, Ed., *The algebraic mind: Integrating connectionism and cognitive science*. Cambridge, MA: MIT Press, 2001.
- [2] I. Goodfellow, Y. Bengio, and A. Courville, Eds., *Deep learning*. MIT Press, 2016.
- [3] A. Dieuleveut, N. Flammarion, and F. Bach, “Harder, better, faster, stronger convergence rates for least-squares regression,” *Journal of Machine Learning Research*, vol. 18, pp. 1–51, 2017.
- [4] S. Arora, N. Cohen, N. Golowich, and W. Hu, “A convergence analysis of gradient descent for deep linear neural networks,” in *ICLR*, 2019.
- [5] S. Amari, “A theory of adaptive pattern classifiers,” *IEEE Transactions on Electronic Computers*, vol. EC-16, no. 3, pp. 299–307, 1967.
- [6] C. Z. Tang and H. K. Kwan, “Convergence and generalization properties of multilayer feedforward neural networks,” in *1992 IEEE International Symposium on Circuits and Systems*, vol. 1, 1992, pp. 65–68.
- [7] M. Hardt, B. Recht, and Y. Singer, “Train faster, generalize better: Stability of stochastic gradient descent,” in *ICML*, 2016.
- [8] B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro, “Exploring generalization in deep learning,” in *NIPS*, 2017.
- [9] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding deep learning requires rethinking generalization,” in *ICLR*, 2017.
- [10] R. Giryes, Y. C. Eldar, A. M. Bronstein, and G. Sapiro, “Tradeoffs between convergence speed and reconstruction accuracy in inverse problems,” *IEEE Transactions on Signal Processing*, vol. 66, no. 7, pp. 1676–1690, 2018.
- [11] S. Lin, “Generalization and expressivity for deep nets,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 5, pp. 1392–1406, 2019.
- [12] J. Duchi, E. Hazan, and Y. Singer, “Adaptive subgradient methods for online learning and stochastic optimization,” *JMLR*, vol. 12, pp. 2121–2159, 2011.
- [13] M. Zeiler, “Adadelta: An adaptive learning rate method,” *arXiv preprint*, vol. arXiv:1212.5701, 2012.
- [14] D. Kingma and J. Ba, “Adam: A method for stochastic optimizations,” in *ICLR*, 2015.
- [15] A. Wilson, R. Roelofs, M. Stern, N. Srebro, and B. Recht, “The marginal value of adaptive gradient methods in machine learning,” in *NIPS*, 2017.
- [16] L. Luo, Y. Xiong, Y. Liu, and X. Sun, “Adaptive gradient methods with dynamic bound of learning rate,” in *ICLR*, 2019.
- [17] D. Liang, F. Ma, and W. Li, “New gradient-weighted adaptive gradient methods with dynamic constraints,” *IEEE Access*, vol. 8, pp. 110929–110942, 2020.
- [18] Y. Xie, X. Wu, and R. Ward, “Linear convergence of adaptive stochastic gradient descent,” in *AISTATS*, 2020.
- [19] B. Heo, S. Chun, S. J. Oh, D. Han, S. Yun, G. Kim, Y. Uh, and J.-W. Ha, “Adamp: Slowing down the slowdown for momentum optimizers on scale-invariant weights,” in *ICLR*, 2021.
- [20] Y. Zhou, B. Karimi, J. Yu, Z. Xu, and P. Li, “Towards better generalization of adaptive gradient methods,” in *NeurIPS*, 2021.
- [21] H. Robbins and S. Monro, “A stochastic approximation method,” *Annals of Mathematical Statistics*, vol. 22, no. 3, pp. 400–407, 1951.
- [22] S. Arora, N. Cohen, and E. Nazan, “On the optimization of deep networks: Implicit acceleration by overparameterization,” in *ICML*, 2018.
- [23] Y. Yuanzhi Li and Y. Liang, “Learning overparameterized neural networks via stochastic gradient descent on structured data,” in *NeurIPS*, 2018.
- [24] Z. Allen-Zhu, Y. Li, and Z. Song, “A convergence theory for deep learning via over-parameterization,” in *ICML*, 2019.
- [25] S. Oymak and M. Soltanolkotabi, “Overparameterized nonlinear learning: Gradient descent takes the shortest path,” in *ICML*, 2019.
- [26] C. Liu and M. Belkin, “Accelerating sgd with momentum for over-parameterized learning,” in *ICLR*, 2020.
- [27] S. Oymak and M. Soltanolkotabi, “Toward moderate overparameterization: Global convergence guarantees for training shallow neural networks,” *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 1, pp. 84–105, 2020.
- [28] Z. Chen, Y. Cao, D. Zou, and Q. Gu, “How much over-parameterization is sufficient to learn deep relu networks?” in *ICLR*, 2021.
- [29] B. Neyshabur, R. Tomioka, and N. Srebro, “In search of the real inductive bias: On the role of implicit regularization in deep learning,” in *ICLR*, 2014.
- [30] D. Soudry, E. Hoffer, M. Nacson, S. Gunasekar, and N. Srebro, “The implicit bias of gradient descent on separable data,” *JMLR*, vol. 19, pp. 1–57, 2018.
- [31] M. Nacson, N. Srebro, and D. Soudry, “Stochastic gradient descent on separable data: Exact convergence with a fixed learning rate,” in *AISTATS*, 2019.
- [32] L. Chizat and F. Bach, “Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss,” in *COLT*, 2020.
- [33] S. L. Smith, B. Dherin, D. Barrett, and S. De, “On the origin of implicit regularization in stochastic gradient descent,” in *ICLR*, 2021.
- [34] J. Wu, D. Zou, V. Braverman, and Q. Gu, “Direction matters: On the implicit regularization effect of stochastic gradient descent with moderate learning rate,” in *ICLR*, 2021.
- [35] C. Yun, S. Krishnan, and H. Mobahi, “A unifying view on implicit bias in training linear neural networks,” in *ICLR*, 2021.
- [36] V. D. Bortoli, A. Durmus, X. Fontaine, and U. Simsekli, “Quantitative propagation of chaos for sgd in wide neural networks,” in *NeurIPS*, 2020.
- [37] J. Cohen, S. Kaur, Y. Li, J. Z. Kolter, and A. Talwalkar, “Gradient descent on neural networks typically occurs at the edge of stability,” in *ICLR*, 2021.
- [38] K. van den Doel and U. Ascher, “The chaotic nature of faster gradient descent methods,” *Journal of Scientific Computing*, vol. 51, pp. 560–581, 2012.
- [39] S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein, “Deep information propagation,” in *ICLR*, 2017.
- [40] P. D. Felice, C. Marangi, G. Nardulli, G. Pasquariello, and L. Tedesco, “Dynamics of neural networks with non-monotone activation function,” *Network: Computation in Neural Systems*, vol. 4, no. 1, pp. 1–9, 1993.
- [41] X. Nie and W. Zheng, “Multistability of neural networks with discontinuous non-monotonic piecewise linear activation functions and time-varying delays,” *Neural Networks*, vol. 65, pp. 65–79, 2015.
- [42] D. A. Findlay, “Training networks with discontinuous activation functions,” in *1989 First IEEE International Conference on Artificial Neural Networks*, (Conf. Publ. No. 313), 1989, pp. 361–363.
- [43] H. Wu, “Global stability analysis of a general class of discontinuous neural networks with linear growth activation functions,” *Information Sciences*, vol. 179, no. 19, pp. 3432–3441, 2009.

- [44] X. Peng, H. Wu, and J. Cao, "Global nonfragile synchronization in finite time for fractional-order discontinuous neural networks with nonlinear growth activations," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 7, pp. 2123–2137, 2019.
- [45] P. Baldi and A. F. Atiya, "How delays affect neural dynamics and learning," *IEEE Transactions on Neural Networks*, vol. 5, no. 4, pp. 612–621, 1994.
- [46] W. Lu and T. Chen, "Dynamical behaviors of delayed neural network systems with discontinuous activation functions," *Neural Computation*, vol. 18, no. 3, pp. 683–708, 2006.
- [47] J. Liu, X. Liu, and W.-C. Xie, "Global convergence of neural networks with mixed time-varying delays and discontinuous neuron activations," *Information Sciences*, vol. 183, no. 1, pp. 92–105, 2012.
- [48] X. Nie and J. Cao, "Existence and global stability of equilibrium point for delayed competitive neural networks with discontinuous activation functions," *International Journal of Systems Science*, vol. 43, no. 3, pp. 459–474, 2012.
- [49] L. Duan, L. Huang, and Z. Guo, "Stability and almost periodicity for delayed high-order hopfield neural networks with discontinuous activations," *Nonlinear Dynamics*, vol. 77, p. 1469–1484, 2014.
- [50] Z. Cai, L. Huang, Z. Guo, and X. Chen, "On the periodic dynamics of a class of time-varying delayed neural networks via differential inclusions," *Neural Networks*, vol. 33, pp. 97–113, 2012.
- [51] O. Romero and M. Benosman, "Finite-time convergence in continuous-time optimization," in *ICML*, 2020.
- [52] A. Levant, "On fixed and finite time stability in sliding mode control," in *52nd IEEE Conference on Decision and Control*, 2013, pp. 4260–4265.
- [53] M. Forti and P. Nistri, "Global convergence of neural networks with discontinuous neuron activations," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 50, no. 11, pp. 1421–1435, 2003.
- [54] G. Huang and J. Cao, "Multistability of neural networks with discontinuous activation function," *Communications in Nonlinear Science and Numerical Simulation*, vol. 13, no. 10, pp. 2279–2289, 2008.
- [55] Z. Guo and L. Huang, "Global exponential convergence and global convergence in finite time of non-autonomous discontinuous neural networks," *Computational Brain & Behavior*, vol. 58, p. 349, 2009.
- [56] S. Qin and X. Xue, "Global exponential stability and global convergence in finite time of neural networks with discontinuous activations," *Neural Processing Letters*, vol. 29, pp. 189–204, 2009.
- [57] T. Legović, "Dynamic population models," in *Ecological model types*, S. Jorgensen, Ed. Elsevier, 2016, pp. 39–63.
- [58] W. Gray and G. Gray, Eds., *Introduction to Environmental Modeling*. Cambridge, UK: Cambridge University Press, 2017.
- [59] E. Hoffer, I. Hubara, and D. Soudry, "Train longer, generalize better: closing the generalization gap in large batch training of neural networks," in *NIPS*, 2017.
- [60] S. Gunasekar, J. Lee, D. Soudry, and N. Srebro, "Characterizing implicit bias in terms of optimization geometry," in *ICML*, 2018.
- [61] X. Yang, X. Zheng, and H. Gao, "Sgd-based adaptive nn control design for uncertain nonlinear systems," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 10, pp. 5071–5083, 2018.
- [62] B. Ying, K. Yuan, S. Vlaski, and A. H. Sayed, "Stochastic learning under random reshuffling with constant step-sizes," *IEEE Transactions on Signal Processing*, vol. 67, no. 2, pp. 474–489, 2019.
- [63] S. R. Dubey, S. Chakraborty, S. K. Roy, S. Mukherjee, S. K. Singh, and B. B. Chaudhuri, "diffgrad: An optimization method for convolutional neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 11, pp. 4500–4511, 2020.
- [64] H. Wang, Y. Luo, W. An, Q. Sun, J. Xu, and L. Zhang, "Pid controller-based stochastic optimization acceleration for deep neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 12, pp. 5079–5091, 2020.
- [65] Z. Li and S. Arora, "An exponential learning rate schedule for deep learning," in *ICLR*, 2020.
- [66] S. Vaswani, A. Mishkin, I. Laradji, M. Schmidt, G. Gidel, and S. Lacoste-Julien, "Painless stochastic gradient: Interpolation, line-search, and convergence rates," in *NeurIPS*, 2019.
- [67] A. Nitanda and T. Suzuki, "Optimal rates for averaged stochastic gradient descent under neural tangent kernel regime," in *ICLR*, 2021.
- [68] W. P. Warren S. McCulloch, "A logical calculus of the ideas immanent in nervous activity," *Bulletin of Mathematical Biophysics*, vol. 5, no. 4, pp. 115–133, 1943.
- [69] F. Rosenblatt, "The perceptron: A probabilistic model for information storage and organization in the brain," *Psychological Review*, vol. 65, no. 6, pp. 386–408, 1958.
- [70] H. R. Wilson and J. D. Cowan, "Excitatory and inhibitory interactions in localized populations of model neurons," *Biophysical Journal*, vol. 12, no. 1, pp. 1–24, 1972.
- [71] T. Yamada and T. Yabuta, "Neural network controller using autotuning method for nonlinear functions," *IEEE Transactions on Neural Networks*, vol. 3, no. 4, pp. 595–601, 1992.
- [72] B. Dasgupta and G. Schnitger, "The power of approximating: a comparison of activation functions," in *NIPS*, 1993, pp. 615–622.
- [73] E. D. Sontag, "Sigmoids distinguish more efficiently than heavisides," *Neural Computation*, vol. 1, no. 4, pp. 470–472, 1989.
- [74] F. Piazza, A. Uncini, and M. Zenobi, "Neural network complexity reduction using adaptive polynomial activation functions," in *ICANN*, 1993.
- [75] S. Xu and M. Zhang, "A novel adaptive activation function," in *Proceedings of the 2002 International Joint Conference on Neural Networks. IJCNN'01*, vol. 4, 2001, pp. 2779–2782 vol.4.
- [76] S. Goh and D. Mandic, "Recurrent neural networks with trainable amplitude of activation functions," *Neural Networks*, vol. 16, no. 8, pp. 1095–1100, 2003.
- [77] Y. Bai, H. Zhang, and Y. Hao, "The performance of the backpropagation algorithm with varying slope of the activation function," *Chaos, Solitons & Fractals*, vol. 40, no. 1, pp. 69–77, 2009.
- [78] S. Flennerhag, H. Yin, J. Keane, and M. Elliot, "Breaking the activation function bottleneck through adaptive parameterization," in *NeurIPS*, 2018.
- [79] W. Wu, J. Liu, H. Wang, F. Tang, and M. Xian, "Ppolynets: Achieving high prediction accuracy and efficiency with parametric polynomial activations," *IEEE Access*, vol. 6, pp. 72 814–72 823, 2018.
- [80] M. Goyal, R. Goyal, and B. Lall, "Improved polynomial neural networks with normalised activations," in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–8.
- [81] V. Kunc and J. Klěma, "On transformative adaptive activation functions in neural networks for gene expression inference," *Chaos, Solitons & Fractals*, vol. 16, no. 1, p. e0243915, 2021.
- [82] R. H. Hahnloser and H. S. Seung, "Permitted and forbidden sets in symmetric threshold-linear networks," 2000.
- [83] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," 2011.
- [84] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, "Fast and accurate deep network learning by exponential linear units (elus)," in *ICLR*, 2016.
- [85] M. Zhu, W. Min, Q. Wang, S. Zou, and X. Chen, "Pfluf and pfllu: Two novel non-monotonic activation functions in convolutional neural networks," *Neurocomputing*, vol. 429, pp. 110–117, 2021.
- [86] P. Wan, D. Sun, M. Zhao, and H. Zhao, "Monostability and multistability for almost-periodic solutions of fractional-order neural networks with unsaturating piecewise linear activation functions," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–15, 2020.
- [87] S. Hayou, A. Doucet, and J. Rousseau, "On the impact of the activation function on deep neural networks training," in *ICML*, 2019.
- [88] T. Ernst, "A method for q-calculus," *Journal of Nonlinear Mathematical Physics*, vol. 4, pp. 487–525, 2003.
- [89] F. Jackson, "On q-functions and a certain difference operator," *Trans. Roy. Soc. Edin.*, vol. 46, pp. 253–281, 1908.
- [90] C. Tsallis, "Possible generalization of boltzmann-gibbs statistics," *Journal of Statistical Physics*, vol. 52, no. 1/2, pp. 479–487, 1988.
- [91] —, "What are the numbers that experiments provide?" *Quimica Nova*, vol. 17, no. 468, pp. 468–471, 1994.
- [92] P. Verhulst, "Notice sur la loi que la population poursuit dans son accroissement," *Correspondance mathématique et physique*, vol. 10, pp. 113–121, 1838.
- [93] M. B. Schaefer, "Some aspects of the dynamics of populations important to the management of commercial marine fisheries," *Bulletin of the Inter-American Tropical Tuna Commission*, vol. 1, pp. 27–56, 1954.
- [94] B. Gompertz, "On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies," *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, vol. 182, pp. 513–585, 1825.
- [95] K. Tjørve and E. Tjørve, "The use of gompertz models in growth analyses, and new gompertz-model approach: An addition to the unified-richards family," *PLoS ONE*, vol. 12, no. 6, p. e0178691, 2017.
- [96] C. Winsort, "The gompertz curve as a growth curve," *Proc. Nat. Acad. Sci.*, vol. 18, no. 1, pp. 1–8, 1932.
- [97] M. Minsky and S. Papert, *Perceptrons: An Introduction to Computational Geometry*. MIT Press, 1969.

- [98] M. Aizerman, E. Braverman, and L. Rozonoer, "Theoretical foundations of the potential function method in pattern recognition learning," *Automation and Remote Control*, vol. 25, pp. 821–837, 1964.
- [99] C. Koch, T. Poggio, V. Torre, and B. B. Boycott, "Retinal ganglion cells: a functional interpretation of dendritic morphology," *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, vol. 298, no. 1090, pp. 227–263, 1982.
- [100] J. Hopfield, "Neurons with graded response have collective computational properties like those of two-state neurons," *Proceedings of the National Academy of Sciences*, vol. 81, no. 10, pp. 3088–3092, 1984.
- [101] P. Fromherz and V. Gaede, "Exclusive-or function of single arborized neuron," *Biological Cybernetics*, vol. 69, pp. 337–344, 1993.
- [102] D. R. Hush, "Training a sigmoidal node is hard," *Neural Computation*, vol. 11, no. 5, pp. 1249–1260, 1999.
- [103] J. Šíma, "Training a single sigmoidal neuron is hard," *Neural Computation*, vol. 14, no. 11, pp. 2709–2728, 2002.
- [104] S. Frei, Y. Cao, and Q. Gu, "Agnostic learning of a single neuron with gradient descent," in *NeurIPS*, 2020.
- [105] A. Gidon, T. A. Zolnik, P. Fidzinski, F. Bolduan, A. Papoutsis, P. Poirazi, M. Holtkamp, I. Vida, and M. E. Larkum, "Dendritic action potentials and computation in human layer 2/3 cortical neurons," *Science*, vol. 367, no. 6473, pp. 83–87, 2020.
- [106] I. S. Jones and K. Kording, "Can single neurons solve mnist? the computational power of biological dendritic trees," *arXiv preprint*, vol. arXiv:2009.01269v1, 2020.
- [107] G. Yehudai and O. Shamir, "Learning a single neuron with gradient methods," in *COLT*, 2020.
- [108] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *JMLR*, vol. 13, pp. 281–305, 2012.
- [109] E. Tjoa and C. Guan, "A survey on explainable artificial intelligence (xai): Toward medical xai," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–21, 2020.

Ilona Kulikovskikh received her Candidate of Sciences (Ph.D.) in 2011 and Doctor of Sciences (Dr. Sc.) in 2020 from the Higher Attestation Committee in Moscow. In 2018-2019, she was a Visiting Scholar at the University of Zagreb and Ruder Bošković Institute, where she worked on bio-inspired learning systems. She is currently with Samara University as a Full Professor and a Senior Researcher, leading projects and conducting transdisciplinary studies on explainable and reliable AI, evolutionary computation, and dynamic systems. She is a member of the

Academy of Navigation and Motion Control.

Tarzan Legović received his M.Sc. from the University of Toronto in 1976 and Ph.D. from the University of Zagreb in 1980. He serves as the president of the International Society for Ecological Modelling. He is employed as a professor at the Libertas International University and the president of the Scientific Council at the OIKON Ltd. At the OIKON he also heads the Data Science Laboratory which is concerned with methods of AI and their applications to ecology. As an external scientific advisor at the R. Bošković institute, he helps in the

development of predictive models for environmental management.