



SPENSER: Towards a NeuroEvolutionary Approach for Convolutional Spiking Neural Networks

Henrique Branquinho
branquinho@dei.uc.pt

University of Coimbra, CISUC, DEI
Coimbra, Portugal

Nuno Lourenço
naml@dei.uc.pt

University of Coimbra, CISUC, DEI
Coimbra, Portugal

Ernesto Costa
ernesto@dei.uc.pt

University of Coimbra, CISUC, DEI
Coimbra, Portugal

ABSTRACT

Spiking Neural Networks (SNNs) have attracted recent interest due to their energy efficiency and biological plausibility. However, the performance of SNNs still lags behind traditional Artificial Neural Networks (ANNs), as there is no consensus on the best learning algorithm for SNNs. Best-performing SNNs are based on ANN to SNN conversion or learning with spike-based backpropagation through surrogate gradients. The focus of recent research has been on developing and testing different learning strategies, with hand-tailored architectures and parameter tuning. Neuroevolution (NE), has proven successful as a way to automatically design ANNs and tune parameters, but its applications to SNNs are still at an early stage. DENSER is a NE framework for the automatic design and parametrization of ANNs, based on the principles of Genetic Algorithms (GA) and Structured Grammatical Evolution (SGE). In this paper, we propose SPENSER, a NE framework for SNN generation based on DENSER, for image classification on the MNIST and Fashion-MNIST datasets. SPENSER generates competitive performing networks with a test accuracy of 99.42% and 91.65% respectively.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**; *Computer vision*; • **Theory of computation** → **Evolutionary algorithms**; **Grammars and context-free languages**.

KEYWORDS

spiking neural networks, neuroevolution, DENSER, computer vision

ACM Reference Format:

Henrique Branquinho, Nuno Lourenço, and Ernesto Costa. 2023. SPENSER: Towards a NeuroEvolutionary Approach for Convolutional Spiking Neural Networks. In *Genetic and Evolutionary Computation Conference Companion (GECCO '23 Companion)*, July 15–19, 2023, Lisbon, Portugal. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3583133.3596399>

1 INTRODUCTION

The advent of Artificial Neural Networks (ANNs) and Deep Learning (DL) has revolutionized the field of Machine Learning (ML) over

the last decade, allowing for the development of high-performing models for computer vision, speech recognition, and natural language processing [23]. However, the success of ANNs is highly dependable not only on the availability of annotated data but mostly on computationally powerful hardware such as Graphical Processing Units (GPU). This hardware dependency foresees an unsustainable future for Artificial Intelligence (AI), as the state-of-the-art models have millions of floating point parameters and require large pipelines of power-hungry hardware for training resulting in large carbon footprints [35].

Spiking Neural Networks (SNNs), often called the third generation of neural networks, are biologically inspired neural network models built with spiking neurons, where information is encoded in discrete binary events over time called action potentials or spikes [31]. SNNs are innately sparse and highly parallelizable, which favors processing speed and energetic efficiency. Albeit still lagging behind ANNs in terms of performance, SNNs show great promise for the future of biologically plausible and sustainable AI. Current bottlenecks in SNN research include the lack of an established learning strategy, such as error backpropagation in ANNs, due to the non-differentiability of the spiking neuron's activation function, and high sensitivity to parameter tuning. Due to this, the focus of recent research, especially regarding image classification problems, has been on developing and testing different learning strategies, with hand-tailored architectures and parameter tuning usually based on successful ANN models [34]. However, it is unclear if these ANN architectures are suited for SNNs as well.

Evolutionary computation (EC) methods are known to be an effective optimization tool [3], and their application to the optimization of ANNs, known as neuroevolution (NE), has proven successful both as a learning strategy as well as a way to automatically design networks and tune parameters [4]. DENSER [1, 2] is a NE framework for the automatic design and parametrization of ANNs, based on the principles of Genetic Algorithms (GA) and Structured Grammatical Evolution (SGE). DENSER has attained impressive results on several benchmark problems, and due to its grammar-based engine, can easily be generalized to a multitude of domains.

In this paper, we propose SPENSER (SPiking Evolutionary Network StructurEd Representation), a NE framework for evolving Convolutional Spiking Neural Networks (CSNN) based on DENSER. This paper is a preliminary experimental study to validate SPENSER for image classification problems. In this study, we evolved the architecture and parameters of SNNs with SPENSER on the MNIST [24] and Fashion-MNIST [42] public datasets, using a fixed learning strategy (Backpropagation Through Time and surrogate gradients).



This work is licensed under a Creative Commons Attribution International 4.0 License. *GECCO '23 Companion*, July 15–19, 2023, Lisbon, Portugal
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0120-7/23/07.
<https://doi.org/10.1145/3583133.3596399>

To the best of our knowledge, this is the first work focusing on evolving SNNs trained with BPTT for image classification, including not only different architectures but different neuronal dynamics and optimizers in the search space. The main contribution of this paper is the preliminary validation of neuroevolution through SPENSER in the automatic generation of competitively performing CSNNs. The main focus of the paper is on the performance of the generated networks in terms of accuracy.

The remainder of this paper is structured as follows: Section 2 provides a review of important concepts regarding SNNs; Section 3 covers related work regarding evolutionary approaches for SNNs; Section 4 describes SPENSER; Section 5 describes the experimental setup; Section 6 analyses the experimental results, covering the evolutionary search and the testing performance of the generated models; Section 7 provides some final remarks and suggested guidelines for future research.

2 SPIKING NEURAL NETWORKS

Spiking Neural Networks (SNNs) are a class of neural network models built with spiking neurons where information is encoded in the timing and frequency of discrete events called spikes (or action potentials) over time [31]. Spiking neurons can be characterized by a membrane potential $V(t)$ and activation threshold V_{thresh} . The weighted sum of inputs of the neuron increases the membrane potential over time. When the membrane potential reaches its activation threshold, a spike is generated (fired) and propagated to subsequent connections. In a feed-forward network, inputs are presented to the network in the form of spike trains (timed sequences of spikes) over T time steps, during which time spikes are accumulated and propagated throughout the network up to the output neurons.

There are a number of spiking neuron models that vary in biological plausibility and computational cost, such as the more realistic and computationally expensive Hodgkin-Huxley [13], to the more simplistic and computationally lighter models such as the Izhikevich [17], Integrate-and-Fire (IF) [22] and Leaky Integrate-and-Fire (LIF) [9]. We refer to Long and Fang [26] for an in-depth review of existing spiking neuron models and their behaviour.

The LIF neuron is the most commonly used in the literature due to its simplicity and low computational cost. The LIF neuron can be modulated as a simple parallel Resistor-Capacitor (RC) circuit with a "leaky" resistor:

$$C \frac{dV}{dt} = -g_L(V(t) - E_L) + I(t) \quad (1)$$

In Eq. 1, C is a capacitor, g_L is the "leaky" resistor (conductor), E_L is the resting potential and $I(t)$ is the current source (synaptic input) that charges up the capacitor to increase the membrane potential $V(t)$. Solving this differential equation through Euler method (demonstration in [11]), we can calculate a neuron's membrane potential at a given timestep t as:

$$V[t] = \beta V[t - 1] + WX[t] - Act[t - 1]V_{thresh} \quad (2)$$

In Eq. 2, β is the decay rate of the membrane potential, $X[t]$ is the input vector (corresponding to $I(t)$), W is the vector of input weights, and $Act[t]$ is the activation function. The activation function can be defined as follows:

$$Act[t] = \begin{cases} 1, & \text{if } V[t] > V_{thresh} \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

A LIF neuron's membrane potential naturally decays to its resting state over time if no input is received ($\beta V[t - 1]$). The potential increases when a spike is received from incoming connections, proportionally to the connection's weight ($WX[t]$). When the membrane potential $V(t)$ surpasses the activation threshold V_{thresh} a spike is emitted and propagated to outgoing connections and the membrane's potential resets ($-Act[t - 1]V_{thresh}$). Resetting the membrane's potential can be done either by subtraction, as is done in the presented example, where V_{thresh} is subtracted at the onset of a spike; or to zero, where the membrane potential is set to 0 after a spike. A refractory period is usually taken into account where a neuron's potential remains at rest after spiking in spite of incoming spikes. The decay rate and threshold can be static or trainable.

Existing frameworks such as *snntorch* [11] allow for the development of SNNs by integration of spiking neuron layers in standard ANN architectures such as Convolutional Neural Networks, by simply replacing the activation layer with a spiking neuron layer.

2.1 Information Coding

Spiking systems rely on discrete events to propagate information, so the question arises as to how this information is encoded. We focus on two encoding strategies: rate coding and temporal coding. In **rate coding**, information is encoded in the frequency of firing rates. This is the case in the communication between photoreceptor cells and the visual cortex, where brighter inputs generate higher frequency firing rates as opposed to darker inputs and respectively lower frequency firing rates [14]. ANNs rely on rate coding of information, as each neuron's output is meant to represent an average firing rate. In **temporal coding**, information is encoded in the precise timing of spikes. A photoreceptor system with temporal coding would encode a bright input as an early spike and a dark input as a last spike. When considering the output of an SNN for a classification task, the predicted class would either be: the one with the highest firing frequency, using rate coding; the one that fires first, using temporal coding.

Temporal coding is advantageous in terms of speed and power consumption, as fewer spikes are needed to convey information, resulting in more sparse events which translate to fewer memory accesses and computation. On the other hand, rate coding is advantageous in terms of error tolerance, as the timing constraint is relaxed to the overall firing rate, and promoting learning, as the absence of spikes can lead to the "dead neuron" problem, where no learning takes place as there is no spike in the forward pass. Increased spiking activity prevents the "dead neuron" problem.

2.2 Learning

Learning in SNNs remains one of the biggest challenges in the community due to the non-differentiability of the activation function of spiking neurons (Eq. 3), which does not allow for the direct transposition of the error backpropagation algorithm.

Commonly used learning strategies include unsupervised learning through Spike-Timing-Dependent Plasticity (STDP) [8], offline conversion from trained ANNs to SNNs (also known as shadow

training) [7, 36], and supervised learning through backpropagation either using spike times [5] or adaptations of the activation function to a continuous-valued function [15, 16, 25, 33, 38]. In this work, we focus on the latter, by training SNNs using backpropagation through time (BPTT) and surrogate gradients.

BPTT is an application of the backpropagation algorithm to the unrolled computational graph over time, usually applied to Recurrent Neural Networks (RNNs) [41]. In order to bypass the non-differentiability of the spiking neuron’s activation function, one can use surrogate gradients, by approximating the activation function with continuous functions centered at the activation threshold during the backward pass of backpropagation [33].

In this experimental study, we considered two surrogate gradient functions available in *snntorch* [11]:

- **Fast-Sigmoid**

$$Act \approx \frac{V}{1 + k|V|} \quad (4)$$

- **ATan** - Shifted arc-tan function

$$Act \approx \frac{1}{\pi} \arctan(\pi V \frac{\alpha}{2}) \quad (5)$$

Regarding the loss function, there are a number of choices available depending on the output encoding of the network (rate vs temporal), that calculate the loss based on spikes or on membrane potential. For this experimental study, we considered rate encoding for inputs and outputs, and as such, chose the **Mean Square Error Spike Count Loss** (adapted from [38]). The spike counts of both correct and incorrect classes are specified as targets as a proportion of the total number of time steps (for example, the correct class should fire 80% of the time and the incorrect classes should only fire 10%). The target firing rates are not required to sum to 100%. After a complete forward pass, the mean square error between the actual ($\sum_{t=0}^T Act[t]$) and target (\hat{Act}) spike counts of each class C is calculated and summed together (Eq.6).

$$\mathcal{L} = \frac{1}{T} \sum_{j=0}^{C-1} \left(\sum_{t=0}^T Act_j[t] - \hat{Act}_j \right)^2 \quad (6)$$

3 RELATED WORK

Recent works blending EC and SNNs are mostly focused on evolving a network’s weights, using evolutionary approaches as a learning strategy [20, 29, 30].

Schuman et al. [37] proposed Evolutionary Optimization for Neuromorphic Systems, aiming to train spiking neural networks for classification and control tasks, to train under hardware constraints, to evolve a reservoir for a liquid state machine, and to evolve smaller networks using multi-objective optimization. However, they focus on simple machine learning classification tasks and scalability is unclear. Elbrecht and Schuman [10] used HyperNeat [40] to evolve SNNs focusing on the same classification tasks. Grammatical Evolution (GE) has also been used previously by López-Vázquez et al. [30] to evolve SNNs for simple classification tasks.

The current state of the art in the automatic design of CSNN architectures are the works of Kim et al. [19] and AutoSNN by Na et al. [32]. Both works focus on Neural Architecture Search

(NAS), with an evolutionary search component implemented in AutoSNN, and attain state-of-the-art performances in the CIFAR-10, CIFAR-100 [21], and TinyImageNet datasets. However, both works fix the generated networks’ hyperparameters such as LIF neuron parameters and learning optimizer. Our work differs from these works by incorporating these properties in the search space.

4 SPENSER

SPENSER (SPiking Evolutionary Network StructurEd Representation) is a general-purpose evolutionary-based framework for the automatic design of SNNs, based on DENSER [1, 2], combining the principles of Genetic Algorithms (GA) [39] and Dynamical Structured Grammatical Evolution (DSGE) [27, 28]. SPENSER works on a two-level basis, separating the GA and the DSGE level, which allows for the modeling of the overall network structure at the GA level while leaving the network layer’s specifications for the DSGE (Figure 1). The use of a grammar is what makes SPENSER a general-purpose framework, as one solely needs to change the grammar to handle different network and layer types, problems and parameters range.

The GA level encodes the macrostructure representing the sequence of evolutionary units that form the network. Each unit corresponds to a nonterminal from the grammar that is later expanded through DSGE. With this representation, we can encode not only the network’s layers as evolutionary units but also the optimizer and data augmentation. Furthermore, by assigning each evolutionary unit to a grammar nonterminal, we can encode prior knowledge and bound the overall network architecture.

The DSGE level is responsible for the specification of each layer’s type and parameters, working independently from the GA level. DSGE represents an individual’s genotype as a set of expansion choices for each expansion rule in the grammar. Starting from a nonterminal unit from the GA level, DSGE follows the expansions set in the individual’s genotype until all symbols in the phenotype are nonterminals. Rules for the layer types and parameters are represented as a Context-Free Grammar (CFG), making it easier to adapt the framework to different types of networks, layers and problem domains.

An example encoding to build CSNNs could be defined by Grammar 1 and the following GA macro structure:

$$[(features, 1, 10), (classification, 1, 3), (output, 1, 1), (learning, 1, 1)]$$

The numbers in each macro unit represent the minimum and maximum number of units that can be incorporated into the network. With this example, the *features* block encodes layers for feature extraction, and therefore we can generate networks with convolutional and pooling layers, followed by 1 to 3 fully connected layers from the *classification* units. The activation layers are restricted to LIF nodes with different surrogate gradient options. The *learning* unit represents the optimizer used for learning and its parameters. The *output* unit encodes the network’s output layer. Numeric parameters are defined by their type, the number of parameters to generate, and the range of possible values.

```

<features> ::= <convolution> <activation>|<convolution> <pooling>
              <activation>
<convolution> ::= layer : conv [num-filters, int, 1, 32, 256] [filter-shape, int, 1, 2, 5]
<classification> ::= <fully-connected> <activation>
<fully-connected> ::= layer : dense [num-units, int, 1, 16, 128]
<activation> ::= layer : LIF <surrogate-gradient>
<surrogate-gradient> ::= <ATan> |<FastSigmoid>
<output> ::= layer : dense num-units : 10 <activation>
<learning> ::= <Adam> |<SGD>
...
    
```

Grammar 1: Example of a Convolutional Spiking Neural Network grammar.

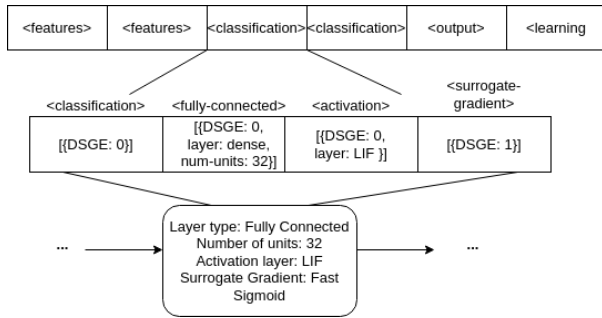


Figure 1: Individual generation by SPENSER. The first line represents the GA level where the macrostructure of the network is defined (this individual has 2 features units and 2 classification units). The second line represents the specification of a classification unit through DSGE. Each number in the DSGE level represents the index of the chosen expansion rule for the current non-terminal. The last line is the resulting phenotype of the layer in question [1].

Regarding variation operators, SPENSER relies on mutations on both levels. At the GA level, individuals can be mutated by adding, replicating, or removing genes i.e. layers. At the DSGE level, mutation changes the layers’ parameters by grammatical mutation (replacing grammatical expansions), integer mutation (replacing an integer parameter with a uniformly generated random one), and float mutation (modifying a float parameter through Gaussian perturbation). SPENSER follows a $(1 + \lambda)$ evolutionary strategy where the parent individual for the next generation is chosen by highest fitness and mutated to generate the offspring. This evolutionary strategy was chosen due to the computational demands of the network training process, which limits the population size in regard to execution time.

5 EXPERIMENTAL SETUP

For this experimental study, we evolved and tested networks on the MNIST [24] and Fashion-MNIST [42] datasets, available through the Torchvision library of Pytorch. All images were converted to grayscale and their original size was kept (28x28). In order to apply SNNs to these datasets, the images were converted to spike trains using rate coding. The pixel values are normalized between 0 and

1 and each pixel value is used as a probability in a Binomial distribution, which is then sampled from to generate spike trains of length T time steps. No data augmentation was used. We considered different time steps for each dataset according to their complexity.

Datasets were split in three subsets: EvoTrain, Fitness and Test. The Test split is the one provided by Torchvision. The EvoTrain and Fitness splits are a 70/30 split of the original Train split. Each independent run generates different EvoTrain and Fitness splits. Table 1 summarises the chosen time steps and the number of samples per split for each dataset.

Table 1: Time steps and number of samples per split for each dataset (MNIST and Fashion-MNIST).

	Time Steps (T)	Train		
		EvoTrain	Fitness	Test
MNIST	10	42000	18000	10000
F-MNIST	25			

As this is a preliminary study to validate SPENSER, we settled on one-pass training of individuals as a trade-off between speed and accuracy. During the evolutionary search, individuals are trained on the EvoTrain split for 1 epoch and tested against the Fitness split for fitness assignment. After the evolutionary search is complete, the best individual is further trained for 50 epochs on the entire Train set, and tested against the Test set for accuracy assessment.

We used *snntorch* [11] to assemble, train and evaluate SNNs based on rate coding. Individuals are trained using BPTT and the chosen loss function was the Mean Square Error Spike Count described in Section 2.2, with a target spiking proportion of 100% for the correct class and 0% for the incorrect class. The predicted class for a given instance is calculated based on the highest spike count of the output neurons. Accuracy is used as the fitness metric during the evolutionary search and as the final performance assessment of the best found individuals.

The macro structure of individuals for the GA level was set as:

$$[(features, 1, 6), (classification, 1, 4), (output, 1, 1), (learning, 1, 1)]$$

Because we are dealing with an image recognition problem, we defined a grammar that contains primitives allowing for the construction of CSNNs, as shown in Grammar 2. Following is a brief description of the grammar.

features units can be expanded to either Convolutional + Activation, Convolutional + Pooling + Activation, or Dropout layers. Convolutional layers are defined by the number of filters, filter shape, stride, padding and bias. Pooling layers are defined by the pooling type (max or average) and the kernel size. *classification* units can be expanded to either Fully-Connected + Activation or Dropout layers. Fully-Connected layers are defined by the number of units. Dropout layers are defined by the dropout rate. The *output* unit is set as a Fully-Connected + Activation where the number of units is fixed to the number of classes. Activation layers are currently limited to LIF neurons. LIF neurons are defined by the decay rate β , the activation threshold V_{thresh} , and the reset mechanism

(subtraction or zero). Furthermore, they are also defined by the surrogate gradient function, which in this case can be either the ATan or the Fast-Sigmoid functions described in Section 2.2. The *learning* unit encodes the optimizer and can be expanded to either Stochastic Gradient Descent, Adam, or RMSProp. We increased the probability of choosing feature extraction layers over dropout for *features* units (Grammar 2, line 1).

```

<features> ::= <aux-convolution> | <aux-convolution> | <aux-convolution>
| <dropout>
<aux-convolution> ::= <convolution> <pooling> <activation>
<activation> ::= layer : act <beta> <threshold> <surr-grad> <
reset-mechanism>
<reset-mechanism> ::= reset : subtract | reset : zero
<beta> ::= [beta, float, 1, 0, 1] <beta-trainable>
<threshold> ::= [threshold, float, 1, 0.5, 1.5] <threshold-trainable>
<beta-trainable> ::= beta-trainable : True | beta-trainable : False
<threshold-trainable> ::= threshold-trainable : True | threshold-trainable : False
<surr-grad> ::= surr-grad : atan | surr-grad : fast-sigmoid
<pooling> ::= <pool-type> [kernel-size, int, 1, 2, 4] | layer : no-op
<pool-type> ::= layer : pool-avg | layer : pool-max
<classification> ::= <fully-connected> <activation> | <dropout>
<convolution> ::= layer : conv [num-filters, int, 1, 32, 128] [filter-shape, int, 1, 2, 4]
[stride, int, 1, 1, 3] <padding> <bias>
<padding> ::= padding : same | padding : valid
<dropout> ::= layer : dropout [rate, float, 1, 0, 0.5]
<fully-connected> ::= layer : fc [num-units, int, 1, 32, 256] <bias>
<bias> ::= bias : True | bias : False
<output> ::= <fully-last> <activation>
<fully-last> ::= layer : fc num-units : 10 bias : True
<learning> ::= <gradient-descent> | <rmsprop> | <adam>
<gradient-descent> ::= learning : gradient-descent [lr, float, 1, 0.0001, 0.1]
[momentum, float, 1, 0.68, 0.99] [decay, float, 1, 0.000001, 0.001]
<nesterov>
<nesterov> ::= nesterov : True | nesterov : False
<adam> ::= learning : adam [lr, float, 1, 0.0001, 0.1]
[beta1, float, 1, 0.5, 0.9999] [beta2, float, 1, 0.5, 0.9999]
[decay, float, 1, 0.000001, 0.001] <amsgrad>
<amsgrad> ::= amsgrad : True | amsgrad : False
<rmsprop> ::= learning : rmsprop [lr, float, 1, 0.0001, 0.1] [rho, float, 1, 0.5, 1]
[decay, float, 1, 0.000001, 0.001]

```

Grammar 2: Convolutional Spiking Neural Network grammar.

Regarding SPENSER’s main hyper-parameters, we followed the recommendations of [1, 2], summarised in Table 2. The table is divided in two parts: i) evolutionary parameters, specifying the evolutionary engine properties such as number of generations, number of parents (μ), number of offspring (λ), mutation rates and fitness function; ii) training parameters, specifying the overall learning parameters fixed for all networks.

All the code, configuration files, grammar, and execution instructions for these experiments are publicly available at GitHub ¹.

¹<https://github.com/henriquejsb/spenser>

Table 2: Hyper-parameters for SPENSER.

Evolutionary Parameter	Value
Number of runs	5
Number of Generations	200
μ (#Parents)	1
λ (#Offspring)	10
Add Layer Rate	25%
Duplicate Layer Rate	15%
Remove Layer Rate	25%
Layer DSGE Rate	15%
Learning DSGE Rate	30%
Gaussian Perturbations	(0, 0.15)
Fitness Function	Accuracy

Training Parameters	Value
Number of epochs	1
Batch Size	64
Loss Function	Mean Square Error Spike Count
Correct Rate	1.0
Incorrect Rate	0.0

6 EXPERIMENTAL RESULTS

6.1 Evolutionary Search

The evolutionary results are promising and show that SPENSER is able to generate increasingly better-performing individuals. Figures 2 and 3 display the evolution of the best fitness and the average fitness of the population across 200 generations, and a violin plot of the fitness of the best found individuals in the MNIST and Fashion-MNIST datasets respectively. The more notable aspects of the evolutionary search are the constant increase in best fitness and the diminishing variance over generations (Fig. 2(a), 3(a)). These aspects showcase SPENSER’s ability to uncover new and better individuals, and its consistency over different runs in generating better performing individuals. Furthermore, the average fitness of the population also increases, particularly in the Fashion-MNIST dataset (Fig. 3(b)), which demonstrates SPENSER’s stability, as a random search would yield a constant average fitness.


In order to understand if there are any notably better design choices for CSNNs, we summarized the best individuals’ (both from MNIST and Fashion-MNIST) characteristics in Table 3. The most interesting result is the total absence of Average Pooling layers, as Kim et al. [19] had also stated that Average Pooling is not preferred for SNNs during their NAS and degraded performance. Furthermore, it is interesting to notice that the ATan surrogate gradient is preferred over the Fast-Sigmoid. The choice of Adam as a preferred optimizer is not surprising as it usually is the best-performing optimizer of the three.

6.2 Test Results


After evolving for 200 generations, the best individuals were trained further for another 50 epochs (totaling 51 epochs) and evaluated on the Test set. Violin plots of the test accuracy on the MNIST and Fashion-MNIST datasets are displayed in Fig. 4. Test results of

Table 3: Network characteristics (percentage) for the best 10 individuals from MNIST and Fashion-MNIST.


Layer Types		Reset Mechanism	
Convolutional	35%	Subtract	63%
Average Pooling	0%	Zero	37%
Max Pooling	19%		
Dropout	11%	Optimizers	
Fully-Connected	35%	Adam	70%
		SGD	20%
		RMSProp	10%
Surrogate Gradients			
ATan	76%		
Fast-Sigmoid	24%		



(a) Best Fitness




(b) Average Fitness




(c) Fitness of best found individuals


Figure 2: Evolutionary analysis of SPENSER on the MNIST dataset over 200 generations. The results are averaged over 5 runs.



(a) Best Fitness



(b) Average Fitness



(c) Fitness of best found individuals

Figure 3: Evolutionary analysis of SPENSER on the Fashion-MNIST dataset over 200 generations. The results are averaged over 5 runs.

different have small variations, showcasing SPENSER’s robustness in generating high-performing networks.

We compared the best attained test accuracy with other works that also trained hand-tailored networks through spike based back-propagation. A comparison of test results is presented in Tab. 4. Albeit not surpassing the state-of-the-art, networks generated by SPENSER are head-to-head with the best-performing networks in the literature.

In order to validate our choice of one epoch training for fitness assessment, we also trained the best networks found in the first generation of each run for another 50 epochs and tested their performance on the Test set. Fig. 5 displays violin plots for the test accuracy of the best individuals from generation 1 and generation




Figure 4: Violin plots of the Test accuracy of the best individuals after further training for 50 epochs.

Table 4: Test accuracy comparison of state of the art and our work.

	MNIST	Fashion-MNIST
Zhang et al. [43]	99.62%	90.13%
Cheng et al. [6]	99.50%	92.07%
Fang et al. [12]	99.72%	94.38%
Jiang et al. [18]	99.61%	94.35%
SPENSER (ours)	99.42%	91.65%

200. It is clear that the networks' performance is dependent on the architecture rather than training epochs and that the networks evolved by SPENSER perform better than random initialization.



Figure 5: Test accuracy on Fashion-MNIST for the best individuals from Generation 1 and Generation 200, after 50 epochs of training.

We hypothesize that a big limitation in this experimental study was the choice of the loss function's parameters, as it does not follow the literature's recommendations [34]. By setting the target firing rate of incorrect classes to 0%, we might be suppressing output activity which is important to distinguish between closely distanced inputs. Furthermore, this experimental setup is sluggish,

as training with BPTT is slower than in traditional ANNs and highly memory intensive. Kim et al. [19] have achieved impressive results without training the generated networks during the search phase, by estimating their future performance based on spike activation patterns across different data samples, and we believe this might be an important improvement to our framework. With faster experiments, we can focus on increasing diversity and coverage of the search space, so that SPENSER can yield better individuals.

7 FINAL REMARKS

In this paper we propose SPENSER, a NE framework to automatically design CSNNs. SPENSER is able to generate competitive performing networks for image classification at the level of the state of the art, without human parametrization of the network's architecture and parameters. SPENSER generated networks with competitive results, attaining 99.42% accuracy on the MNIST [24] and 91.65% accuracy on the Fashion-MNIST [42] datasets. Current limitations rely on the execution time, due to the computationally intensive BPTT learning algorithm and the memory requirements. Furthermore, we believe the configuration of the loss function played a role in suppressing output activity and potentially decreasing accuracy.

7.1 Future Work

In the future, we plan on:

- Experiment with different loss functions / encode the loss function as an evolvable macro parameter;
- Perform a more in-depth study of the preferred choices during evolution and observable patterns in the best-performing individuals. This could be relevant in uncovering novel optimal architectures and parameters;
- Experiment with different learning algorithms.
- Implement skip connections and back connections.
- Apply regularisation methods to prevent vanishing and exploding gradients.

ACKNOWLEDGMENTS

This research was supported by the Portuguese Recovery and Resilience Plan (PRR) through project C645008882-00000055, Center for Responsible AI, by the FCT - Foundation for Science and Technology, I.P./MCTES through national funds (PIDDAC), within the scope of CISUC R&D Unit - UIDB/00326/2020 or project code UIDP/00326/2020. The first author is partially funded by FCT - Foundation for Science and Technology, Portugal, under the grant 2022.11314.BD.

REFERENCES

- [1] Filipe Assunção, Nuno Lourenço, Penousal Machado, and Bernardete Ribeiro. 2019. DENSER: deep evolutionary network structured representation. *Genetic Programming and Evolvable Machines* 20, 1 (2019), 5–35.
- [2] Filipe Assunção, Nuno Lourenço, Bernardete Ribeiro, and Penousal Machado. 2021. Fast-DENSER: Fast Deep Evolutionary Network Structured Representation. *SoftwareX* 14 (2021), 100694.
- [3] Thomas Bäck and Hans-Paul Schwefel. 1993. An overview of evolutionary algorithms for parameter optimization. *Evolutionary computation* 1, 1 (1993), 1–23.
- [4] Alejandro Baldominos, Yago Saez, and Pedro Isasi. 2020. On the automated, evolutionary design of neural networks: past, present, and future. *Neural computing and applications* 32 (2020), 519–545.

- [5] Sander M Bohte, Joost N Kok, and Han La Poutre. 2002. Error-backpropagation in temporally encoded networks of spiking neurons. *Neurocomputing* 48, 1-4 (2002), 17–37.
- [6] Xiang Cheng, Yunzhe Hao, Jiaming Xu, and Bo Xu. 2021. LISNN: improving spiking neural networks with lateral interactions for robust object recognition. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*. 1519–1525.
- [7] Shikuang Deng and Shi Gu. 2021. Optimal Conversion of Conventional Artificial Neural Networks to Spiking Neural Networks. In *International Conference on Learning Representations*.
- [8] Peter U Diehl and Matthew Cook. 2015. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Frontiers in computational neuroscience* 9 (2015), 99.
- [9] Sangya Dutta, Vinay Kumar, Aditya Shukla, Nihar R Mohapatra, and Udayan Ganguly. 2017. Leaky integrate and fire neuron by charge-discharge dynamics in floating-body MOSFET. *Scientific reports* 7, 1 (2017), 8257.
- [10] Daniel Elbrecht and Catherine Schuman. 2020. Neuroevolution of Spiking Neural Networks Using Compositional Pattern Producing Networks. In *International Conference on Neuromorphic Systems 2020 (ICONS 2020)*. Association for Computing Machinery, New York, NY, USA, 1–5.
- [11] Jason K. Eshraghian, Max Ward, Emre Neftci, Xinxin Wang, Gregor Lenz, Girish Dwivedi, Mohammed Bannamoun, Doo Seok Jeong, and Wei D. Lu. 2022. Training Spiking Neural Networks Using Lessons From Deep Learning. arXiv:2109.12894 [cs].
- [12] Wei Fang, Zhaofei Yu, Yanqi Chen, Timothée Masquelier, Tiejun Huang, and Yonghong Tian. 2021. Incorporating learnable membrane time constant to enhance learning of spiking neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2661–2671.
- [13] A. L. Hodgkin and A. F. Huxley. 1952. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology* 117, 4 (Aug. 1952), 500–544.
- [14] David H Hubel and Torsten N Wiesel. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology* 160, 1 (1962), 106.
- [15] Dongsung Huh and Terrence J Sejnowski. 2018. Gradient descent for spiking neural networks. *Advances in neural information processing systems* 31 (2018).
- [16] Eric Hunsberger and Chris Eliasmith. 2015. Spiking deep networks with LIF neurons. arXiv preprint arXiv:1510.08829 (2015).
- [17] Eugene M Izhikevich. 2003. Simple model of spiking neurons. *IEEE Transactions on neural networks* 14, 6 (2003), 1569–1572.
- [18] Chunming Jiang and Yilei Zhang. 2023. KLIF: An optimized spiking neuron unit for tuning surrogate gradient slope and membrane potential. arXiv preprint arXiv:2302.09238 (2023).
- [19] Youngeun Kim, Yuhang Li, Hyoungseob Park, Yeshwanth Venkatesha, and Priyadarshini Panda. 2022. Neural Architecture Search for Spiking Neural Networks. <https://doi.org/10.48550/arXiv.2201.10355> arXiv:2201.10355 [cs, eess].
- [20] Katarzyna Kozdon and Peter Bentley. 2018. The Evolution of Training Parameters for Spiking Neural Networks with Hebbian Learning. In *ALIFE 2018: The 2018 Conference on Artificial Life*. MIT Press, 276–283.
- [21] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [22] Louis Lapicque. 1907. Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *Journal de physiologie et de pathologie générale* 9 (1907), 620–635.
- [23] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
- [24] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [25] Eimantas Ledinauskas, Julius Ruseckas, Alfonsas Juršėnas, and Giedrius Buračas. 2020. Training deep spiking neural networks. arXiv preprint arXiv:2006.04436 (2020).
- [26] Lyle Long and Guoliang Fang. 2010. A Review of Biologically Plausible Neuron Models for Spiking Neural Networks. In *AIAA Infotech@Aerospace 2010*. American Institute of Aeronautics and Astronautics, Atlanta, Georgia.
- [27] Nuno Lourenço, Filipe Assunção, Francisco B Pereira, Ernesto Costa, and Penousal Machado. 2018. Structured grammatical evolution: a dynamic approach. In *Handbook of Grammatical Evolution*. Springer, 137–161.
- [28] Nuno Lourenço, Francisco B Pereira, and Ernesto Costa. 2015. SGE: a structured representation for grammatical evolution. In *International Conference on Artificial Evolution (Evolution Artificielle)*. Springer, 136–148.
- [29] Sen Lu and Abhronil Sengupta. 2022. Neuroevolution Guided Hybrid Spiking Neural Network Training. *Frontiers in Neuroscience* 16 (April 2022), 838523.
- [30] G. López-Vázquez, M. Ornelas-Rodríguez, A. Espinal, J. A. Soria-Alcaraz, A. Rojas-Domínguez, H. J. Puga-Soberanes, J. M. Carpio, and H. Rostro-Gonzalez. 2019. Evolutionary Spiking Neural Networks for Solving Supervised Classification Problems. *Computational Intelligence and Neuroscience* 2019 (March 2019), e4182639. Publisher: Hindawi.
- [31] Wolfgang Maass. 1997. Networks of spiking neurons: The third generation of neural network models. *Neural Networks* 10, 9 (Dec. 1997), 1659–1671.
- [32] Byunggook Na, Jisoo Mok, Seongsik Park, Dongjin Lee, Hyeokjun Choe, and Sungroh Yoon. 2022. AutoSNN: Towards Energy-Efficient Spiking Neural Networks. In *Proceedings of the 39th International Conference on Machine Learning*. PMLR, 16253–16269.
- [33] Emre O Neftci, Hesham Mostafa, and Friedemann Zenke. 2019. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine* 36, 6 (2019), 51–63.
- [34] Joao D Nunes, Marcelo Carvalho, Diogo Carneiro, and Jaime S Cardoso. 2022. Spiking neural networks: A survey. *IEEE Access* 10 (2022), 60738–60764.
- [35] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. 2021. Carbon emissions and large neural network training. arXiv preprint arXiv:2104.10350 (2021).
- [36] Bodo Rueckauer, Iulia-Alexandra Lungu, Yuhuang Hu, Michael Pfeiffer, and Shih-Chii Liu. 2017. Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *Frontiers in neuroscience* 11 (2017), 682.
- [37] Catherine D Schuman, J Parker Mitchell, Robert M Patton, Thomas E Potok, and James S Plank. 2020. Evolutionary optimization for neuromorphic systems. In *Proceedings of the Neuro-inspired Computational Elements Workshop*. 1–9.
- [38] Sumit B Shrestha and Garrick Orchard. 2018. Slayer: Spike layer error reassignment in time. *Advances in neural information processing systems* 31 (2018).
- [39] S.N. Sivanandam and S. N. Deepa. 2007. *Introduction to Genetic Algorithms*.
- [40] Kenneth O Stanley, David B D'Ambrosio, and Jason Gauci. 2009. A hypercube-based encoding for evolving large-scale neural networks. *Artificial life* 15, 2 (2009), 185–212.
- [41] Paul J Werbos. 1990. Backpropagation through time: what it does and how to do it. *Proc. IEEE* 78, 10 (1990), 1550–1560.
- [42] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv:1708.07747 (2017).
- [43] Wenrui Zhang and Peng Li. 2019. Spike-train level backpropagation for training deep recurrent spiking neural networks. *Advances in neural information processing systems* 32 (2019).