

# Evolving Spiking Neural Networks for Nonlinear Control Problems

Huanneng Qiu

*School of Engineering and Information Technology  
The University of New South Wales at  
the Australian Defence Force Academy  
Canberra, Australia  
huanneng.qiu@student.unsw.edu.au*

Matthew Garratt

*School of Engineering and Information Technology  
The University of New South Wales at  
the Australian Defence Force Academy  
Canberra, Australia  
m.garratt@adfa.edu.au*

David Howard

*Robotics and Autonomous Systems Group  
DATA61 at CSIRO  
Brisbane, Australia  
david.howard@data61.csiro.au*

Sreenatha Anavatti

*School of Engineering and Information Technology  
The University of New South Wales at  
the Australian Defence Force Academy  
Canberra, Australia  
s.anavatti@adfa.edu.au*

**Abstract**—Spiking Neural Networks are powerful computational modelling tools that have attracted much interest because of the bioinspired modelling of synaptic interactions between neurons. Most of the research employing spiking neurons has been non-behavioural and discontinuous. Comparatively, this paper presents a recurrent spiking controller that is capable of solving nonlinear control problems in continuous domains using a popular topology evolution algorithm as the learning mechanism. We propose two mechanisms necessary to the decoding of continuous signals from discrete spike transmission: (i) a background current component to maintain frequency sufficiency for spike rate decoding, and (ii) a general network structure that derives strength from topology evolution. We demonstrate that the proposed spiking controller can learn significantly faster to discover functional solutions than sigmoidal neural networks in solving a classic nonlinear control problem.

**Index Terms**—Spiking Neural Networks, neuroevolution, recurrent networks

## I. INTRODUCTION

Artificial Neural Networks (ANNs) are learning based computational systems inspired by the structure of animal brains. The current generation of neural networks which are dominant in the computational intelligence community are sigmoidal neural networks, whose neurons consist of two components: a weighted sum of inputs and a sigmoidal activation function generating the output accordingly. Despite their exceptional performance in a variety of tasks, they are vaguely related to their biological counterparts and do not approximate the signal transmission in biological neural systems.

The idea of Spiking Neural Networks (SNNs) [1], therefore, is to bridge the gap between neuroscience and computational intelligence, by employing biologically realistic neuron models to carry out computation. However, over the past decade most of the work analysing information processing in SNNs has

been focused on non-behavioural functionality, e.g., character recognition [2] and approximation [3]. Meanwhile, there is also a rising interest in behaviourally functional SNNs, which addresses neural activities in closed-loop interaction with the environment [4]–[6].

The most apparent advantage of *neurocontrollers* is that neural networks can learn to perform satisfactory tasks without explicit models of the plants. This is highly preferred in situations where accurate models are difficult to obtain. SNNs can be more suitable controllers than non-spiking neural systems because: (i) they are bioinspired plastic learning architectures and can provide faster information processing as observed in biological neural systems [7]; (ii) they are fundamentally computationally more powerful than non-spiking neural systems [1]; and (iii) massive network implementation on hardware has already been realised [8], [9], as spiking neuron models can be represented using simple electric circuits.

One issue of SNN implementation is that the *learning* process is challenging using gradient based methods, partly due to the difficulty of extracting gradient information from discrete events, but also because of the recurrent architecture that is essential to preserve memory through internal connections. In this paper, we use the popular NEAT algorithm [10] as the learning mechanism to generate action selection policies and consequently to seek functional network compositions. NEAT is an ideal neuroevolution strategy due to (i) the efficacy of automatic topology altering along with connection weight; and (ii) compatibility with the discrete nature of SNNs.

The main purpose of this work is to derive continuous motion actions from discrete spike trains in a reinforcement learning task. We claim this has not yet been fully resolved because of two problems: decrease of spike activities through synapses and low resolution of conventional decoding meth-

ods. Therefore, we construct a recurrent spiking network with two essential settings: a background current to deal with frequency loss, and a decoding method based on a weighted firing rate. We show that the proposed spiking controller has better performance than its sigmoidal counterpart in solving the classic pole balancing problem.

Organization of the rest of this paper is as follows. Section II will briefly describe the spiking neuron model and the mechanism to preserve spike activities, followed by a review of the NEAT algorithm in Section III. Section IV presents the experimental setup to test the controller learning algorithm and results of both spiking and non-spiking approaches. Finally, Discussion and Conclusion will be covered in Section V and Section VI.

## II. COMPUTATION WITH SPIKES

Unlike traditional neural networks, the information transmission in SNNs relies on electrical pulses. These so-called *spikes* are discrete events that occur at certain points in time. The information carried by spikes is not by the amplitude or form, but rather by the number and the timing of spikes.

A spike is fired when the *membrane potential* of the *presynaptic* neuron exceeds its threshold; then it will travel through the synapses and arrive at all forward-connected *postsynaptic* neurons. One incoming spike would increase the membrane potential of the neuron, which will decrease gradually until it reaches the resting potential if no other arriving spike is observed within a certain period. Therefore, a cluster of incoming spikes would be necessary for a neuron to generate a spike.

### A. Neural Model

Several spiking neuron models have been proposed over past decades [11]. The two-dimensional Izhikevich model [12] is used in this paper, because of its simplicity whilst being able to produce rich firing patterns to simulate biological neurons by adjusting only a few parameters. This model is formulated by two ordinary differential equations:

$$\dot{v} = 0.04v^2 + 5v + 140 - u + I \quad (1)$$

$$\dot{u} = a(bv - u) \quad (2)$$

with after-spike resetting following:

$$\text{if } v \geq v_t, \text{ then } \begin{cases} v = c \\ u = u + d \end{cases} \quad (3)$$

where  $v$  represents the membrane potential of the neuron;  $u$  represents a recovery variable;  $I$  represents the synaptic current injected into the neuron;  $v_t$  is a threshold value; and  $a, b, c$  and  $d$  are dimensionless parameters to form different spike patterns [12]. Fig. 1 shows the voltage response of an Izhikevich neuron when injected with a square wave current signal.

Firing time  $t^{(f)}$  is defined as the moment of  $v$  crossing threshold  $v_t$  from below:

$$t^{(f)} : v(t^{(f)}) = v_t \text{ and } \dot{v}(t^{(f)}) > 0 \quad (4)$$




Fig. 1. Voltage response of an Izhikevich neuron to a square wave current signal  $I$  corresponds to  $a = 0.02$ ;  $b = 0.2$ ;  $c = -65$ ;  $d = 2$  and  $v_t = 30$  mV.

A spike train is then denoted as the sequence of spike times:

$$s(t) = \sum_f \delta(t - t^{(f)}) \quad (5)$$

where  $\delta(t)$  is the Dirac  $\delta$  function<sup>1</sup>.

### B. Spike Transmission

When interfacing the spiking controller with its plant, one critical problem we should consider is how to decode spikes into output commands and encode sensing data into spikes.

The information representation of spiking neural systems can be distinguished between *rate coding* and *temporal coding* schemes [13]. In a rate coding scheme, neural information is encoded into the number of spikes occurring during a given time window, aka. firing *rate* of spikes. In a temporal coding scheme, the context is encoded into the exact *timing* between presynaptic and postsynaptic spikes.

In this paper the rate coding scheme is adopted, as it is easier to implement. Mean firing rate is averaged during a given time window, which will be further decoded to generate the output signal.

Encoding of sensing data will go through two stages. Input variables are first normalized within the range of [0,1], after which, the rescaled signal will be linearly converted into a current value that will be injected into forward-connected neurons through synapses. This current encoding method is commonly used in literature, and is compatible with different coding schemes, as the neuron receiving larger input currents will not only fire at higher rate, but also be earlier to spike [14].

In a rate coding framework, the neuron will fire at a steady-state firing rate  $f$ , given the injected current  $I$  is fixed in time. Therefore, the spike train frequency can be defined as a function of the magnitude of the input stimulus. Fig. 2 shows the curve of mean firing rate of an Izhikevich neuron when the input current varies from zero to 200 nA.

<sup>1</sup>[https://en.wikipedia.org/wiki/Dirac\\_delta\\_function](https://en.wikipedia.org/wiki/Dirac_delta_function)




Fig. 2.  $f$ - $I$  curve of an Izhikevich neuron with parameters as:  $a = 0.02$ ;  $b = 0.2$ ;  $c = -50$ ;  $d = 2$  and  $v_t = 30$  mV.

It is critical to maintain sufficient spike transmission in order to calculate a smooth firing rate. However, as mentioned, a cluster of incoming spikes are generally required to force a neuron to spike. Therefore, there tends to be a transmission loss as spikes travel through the synapses. In this paper, we try to resolve this problem by simply adding a background current, with which neurons will fire at a certain frequency even if there is no other input stimulus. Such an effect can also be found in biology, where synaptic background noise can significantly affect the neuron characteristics [15].

### III. NEUROEVOLUTION

Despite the success of gradient methods in training traditional multilayer neural networks, their implementation on SNNs is still problematic when extracting gradient information from output spike times. Instead, population-based neuroevolution (NE) is an ideal learning architecture for our settings – recurrent network learning and evolution of network topology.

#### A. Evolutionary Algorithms (EAs)

EAs [16] are an abstraction of the natural evolution process. A general scheme of EAs is given in Algorithm 1. One of the cornerstones of EAs is competition based selection, which is widely known as *survival of the fittest* – individuals that are more adapted to the environment have a higher chance to create offspring or survive. Another basis lies in phenotype variation. To generate different individuals, variation operations (i.e. *recombination* (also termed *crossover*) and *mutation*) are applied, producing new individuals that loosely resemble their parents. The combination of variation and selection thereby leads to a population that is better adapted to the environment or better able to complete a given task.

However, EAs are likely to converge around one single solution as evolution goes on – a phenomenon known as *genetic drift* [16], which can lead to premature convergence and consequently getting stuck at local optimum, because the population may quickly converge on whatever network that happens to perform best in the initial population. Therefore,

---

#### Algorithm 1 General Scheme of Evolutionary Algorithms

---

- 1: INITIALISE a population of candidate solutions
  - 2: EVALUATE each individual & ASSIGN fitness values
  - 3: **while** (terminate condition not met) **do**
  - 4:     SELECT parents
  - 5:     RECOMBINATION parents
  - 6:     MUTATE offspring
  - 7:     EVALUATE new candidates
  - 8:     SELECT individuals for the next generation
  - 9: **end while**
- 

the NEAT algorithm [10] is introduced in this paper to overcome this problem.

#### B. NeuroEvolution of Augmenting Topology (NEAT)

NEAT is a powerful evolutionary approach for neural network learning which evolves network topologies along with connection weights. The efficacy of NEAT is guaranteed by: (i) historical markings to solve the variable-length genome problem; (ii) speciation to protect innovation and preserve network diversity, to avoid premature convergence; and (iii) incremental structural growth to avoid troublesome hand design of network topology.

Typical EAs are difficult to genetically crossover neural networks with variant topologies, because they can only operate within fixed-sized genome space. Recombination of divergent genomes tends to produce damaged offspring. However, similar network solutions sharing similar functionalities can be encoded using completely different genomes – a phenomenon known as the *Competing Convention Problem* [10]. To address this problem, NEAT uses historical markings which act as artificial evidence to track the origin of genes. When two genes share the same historical marking, they are categorised as alleles. Therefore, NEAT can match up genomes representing similar network structures and allow mating in a rationale manner.

NEAT also uses an explicit fitness sharing scheme [16] as a population management approach to preserve network multimodality. Historical markings are used as a measurement of the genetic similarity of network topologies, based on which, genomes are speciated into different *niches* (also termed *species*). Individuals clustered into the same species will share their fitness score together [16]. The fitness of each individual is scaled according to the number of individuals falling in the niche. As competition within the species becomes more intense, solutions will have a lower fitness. Thus, the sparsely populated species will become more attractive, avoiding any one taking over. Therefore, innovations will be protected within niches to have time to optimize.

Finally, an incremental growth mechanism is used in NEAT to discover the least complex effective neural topology, by beginning searching minimal network structure and gradually expanding to more complex networks during evolution.

The NEAT algorithm has been successfully applied to different control problems [10], [17], [18]. However, imple-

mentation of NEAT has been restricted to traditional neural systems, with one exception [19] targeting spiking networks.

### C. Speciation Measurement

NEAT use a compatibility distance function  $\delta$  to determine the similarity of network solutions. When the distance between any two individuals is smaller than a threshold  $\delta_t$ , they are categorized into the same species. The compatibility distance  $\delta$  is defined as:

$$\delta = \frac{c_1 E}{N} + \frac{c_2 D}{N} + c_3 \overline{W} \quad (6)$$

where  $E$  and  $D$  denote the number of *excess* and *disjoint* genes;  $\overline{W}$  denotes the average connection weight difference;  $N$  is the total number of genes;  $c_1$ ,  $c_2$  and  $c_3$  are user-defined coefficients for altering the significance of these factors.

In our second task illustrated later, we will apply a slight modification to this function by taking account of the influence of the decoding method. There will be a more detailed description in Section IV-D.

## IV. EXPERIMENTS

We evaluate the system's performance using a classic nonlinear control benchmark – the pole balancing problem (also known as the inverted pendulum problem). This problem is not only inherently unstable, but also capable of varying degrees of complexity by limiting the state variables provided to the controller, which makes it ideal for designing and testing nonlinear control methods.

Previous attempts to solve this problem using SNNs with fixed-topology can be found in [20], [21]. In this paper, we take a different approach using NEAT. Results of the proposed spiking controller are benchmarked against the original sigmoidal counterpart. The original NEAT C++ source code is available publicly<sup>2</sup>, which is tailored to be amenable to our network model. All the experiments are programmed in C++ and performance analysis is carried out using MATLAB.

### A. Benchmark Problem

Fig. 3 shows the cart-pole system to be controlled, which consists of a cart that can move left or right within a bounded one-dimensional track, and a pole that is hinged to the cart. The problem is to balance the pole upright for as long as possible by applying a force  $F_t$  to the cart parallel to the track. The system has four state variables:

- $x$  – cart position
- $\theta$  – pole angle
- $\dot{x}$  – cart velocity
- $\dot{\theta}$  – pole angular velocity

For simplicity, we neglect the friction of cart on track and that of pole on cart. The system is then formulated by two nonlinear differential equations [22]:




Fig. 3. Cart-pole system, taken from [22].

$$\ddot{\theta}_t = \frac{g \sin \theta + \cos \theta \left( \frac{-F_t - m_p l \dot{\theta}^2 \sin \theta}{m_c + m_p} \right)}{l \left( \frac{4}{3} - \frac{m_p \cos^2 \theta}{m_c + m_p} \right)} \quad (7)$$

$$\ddot{x}_t = \frac{F_t + m_p l (\dot{\theta}^2 \sin \theta - \ddot{\theta} \cos \theta)}{m_c + m_p} \quad (8)$$

where  $g = 9.8 \text{ m/s}^2$  denotes the gravitational acceleration;  $m_c = 1.0 \text{ kg}$  denotes the mass of cart and  $m_p = 0.1 \text{ kg}$  denotes the mass of pole;  $l = 0.5 \text{ m}$  is half length of the pole.

The discrete form of state variables are updated following:

$$\begin{cases} x_{t+1} = x_t + \tau \dot{x}_t \\ \dot{x}_{t+1} = \dot{x}_t + \tau \ddot{x}_t \\ \theta_{t+1} = \theta_t + \tau \dot{\theta}_t \\ \dot{\theta}_{t+1} = \dot{\theta}_t + \tau \ddot{\theta}_t \end{cases} \quad (9)$$

with  $\tau$  representing the time step.

A force generated by the spiking controller at each time step will be used to update the state variables following (7), (8) and (9). A failure signal is generated when the cart reaches the track boundary, which is  $\pm 2.4$  meters from the track centre, or if the pole tilts beyond the failure angle, which is  $\pm 12$  degrees (or about 0.21 radian) from the vertical.

### B. Experimental Setup

We first start with the basic balancing task with complete state variables. This Markovian problem can act as a base performance measurement before we go to the more challenging non-Markovian version without velocity information. In both tasks the cart-pole system model is unknown to the spiking controller.

The controller contains a population of 150 networks, which will be evolved using a combined EA, including a (150, 150) Evolution Strategy [16] plus species-based Elitism. Per epoch, the champion of each species is duplicated when the number of networks in that species is larger than 5. The best 20% of networks in each species are allowed to reproduce, after which, all parents are discarded and the remaining 150 offspring will form the next generation.

<sup>2</sup><http://nn.cs.utexas.edu/?neat-c>

Each network will be evaluated and assigned a fitness value. Fitness is defined as the number of time steps that the balanced criteria is not violated. Otherwise a failure signal is generated and evaluation is moved on to the next network.

### C. Pole Balancing with Velocity

The first task is to balance the pole with velocity information. At initialization, each network consists of 5 spiking neurons – 4 input nodes each receiving one state variable and one output node generating a force applied to the cart; and 4 connections each connecting one input node to the output node, respectively. We use a probabilistic rate coding method here. The normalized input variables are encoded as a probability to generate a spike. At each time step, this probability will be used to determine the firing status of the input neuron. Firing rate is then calculated based on the output node. A binary force ( $F_t = \pm 10$  Newtons) is generated each time to be applied to the cart.

During evolution, hidden nodes are allowed to be added with a probability of 0.03. Connections are added with a probability of 0.1. The time step  $\tau$  in (9) is set to 0.02 seconds. A successful solution is identified when it is able to balance the pole for 100,000 time steps, which is equivalent to around 30 minutes of simulated time.

We apply NEAT to evolve both SNNs and sigmoidal networks. Each test is run for 60 episodes. Table I shows a summary of fewest generations needed to complete the task. A failure run means the controller fails to find a solution within 100 generations. Both approaches failed once in 60 runs. For those successful runs, it is interesting to notice that the spiking controller takes fewer generations to solve the task.

Example runs of both tests are shown in Fig. 4 and Fig. 5. The spiking approach has more jerky outputs due to the probability coding method we use to encode sensing data. Smaller output errors can be observed when the magnitude of the force  $F_t$  is set smaller.

### D. Pole Balancing without Velocity

Our main task is to balance the pole without velocity information. Instead of using bang-bang control, we use a continuous output force within  $[-10, 10]$  Newtons for this more challenging problem. A normalized force  $F_n$  is first calculated based on a weighted sum of firing rates from connected neurons following a modified sigmoid function, which is then scaled and shifted to generate the force  $F_t$ :

$$F_n = \frac{1}{1 + \exp(-\sigma \sum w_i r_i)} \quad (10)$$

$$F_t = 10(2F_n - 1) \quad (11)$$

where  $\sigma$  is a positive decay variable, which will be automatically tuned during evolution;  $r_i$  denotes the firing rate of the  $i^{\text{th}}$  connected neuron and  $w_i$  denotes the corresponding connection weight.

The aforementioned compatibility distance function (6) is modified by considering  $\sigma$  in (10). A component of  $\sigma$  differ-

TABLE I  
FEWEST GENERATIONS REQUIRED TO FIND A SUCCESSFUL SOLUTION FOR MARKOVIAN POLE BALANCING PROBLEM

	Best	Worst	Median <sup>a</sup>	Mean <sup>a</sup>	Failure Rate
NEAT original	3	55	17.5	21.68	1/60
NEAT-SNN	1	92	9.5	16.55	1/60

<sup>a</sup>Values are calculated assuming failure runs take 101 generations.




Fig. 4. NEAT original: cart position  $x$  and pole angle  $\theta$  of the pole balancing task with velocity information.




Fig. 5. NEAT-SNN: cart position  $x$  and pole angle  $\theta$  of the pole balancing task with velocity information.

ence (denoted as  $D_\sigma$ ) is added to the original compatibility distance:

$$\delta = \frac{c_1 E}{N} + \frac{c_2 D}{N} + c_3 \bar{W} + c_4 D_\sigma \quad (12)$$

Fig. 6 shows a possible network topology during evolution. Inputs are cart position and pole angle. Connections and nodes will be added based on a probability of 0.1 and 0.03. Recurrence is allowed within spiking neurons, facilitating internal calculation of derivatives.

The initial state variables are set to 0, except the pole angle is set to 3 degrees. Time step  $\tau$  is set to 0.01 seconds. Successful solutions are dictated if the pole is balanced for 5000 time steps.




Fig. 6. Spiking network topology that allows internal recurrence. Network inputs consist of cart position  $x_t$  and pole angle  $\theta_t$ . Output force  $F_t$  is calculated based on a weighted sum of firing rates  $\sum w_i r_i$ , in which  $r_i$  denotes the firing rate of the  $i^{\text{th}}$  connected neuron and  $w_i$  denotes the corresponding connection weight.

TABLE II  
FEWEST GENERATIONS REQUIRED TO FIND A SUCCESSFUL SOLUTION  
FOR NON-MARKOVIAN POLE BALANCING PROBLEM

	Best	Worst	Median <sup>a</sup>	Mean <sup>a</sup>	Failure Rate
NEAT original	8	96	44	50.35	7/60
NEAT-SNN	5	65	21.5	24.70	0/60

<sup>a</sup>Values are calculated assuming failure runs take 101 generations.

Similarly, we apply NEAT to the spiking controller and compare the results against the original NEAT. A summary is shown in Table II. The proposed spiking controller is essentially better than sigmoidal networks in solving this problem. It requires fewer generations to find a functional solution. It also has a lower failure rate over 60 runs, showing the potential to be more adaptive. Further, the Mann-Whitney  $U$ -test is used to assess the statistic difference between the two sets of samples. The  $p$  value is smaller than 0.01, showing that the spiking controller has significantly better performance.

To visualize the evolution progress, we average the best networks' fitness values over 60 runs. Fig. 7 shows the mean and standard deviation of fitness values of both tests at successive generations. In the beginning, only some of the networks can find a path to optimization, thus introducing a large fitness deviation. As evolution goes on, individuals with higher fitness values will gradually take over the entire population.

Fig. 8 and Fig. 9 show the status of cart position, pole angle and the output force applied to the cart during a successful run. In both approaches, all these three variables are oscillating around a certain point.

## V. DISCUSSION

The design of functional SNNs is considered to be difficult, because SNNs behave as complex systems with transient dynamics [7]. Therefore, parameter setting of spiking network models to solve a given task is non-trivial and still not entirely resolved. Apart from the neuron model, synaptic dynamics with transmission delay also acts as a significant component




Fig. 7. Best networks' mean fitness values in progress over 60 runs, with error bars denoting the standard deviation.




Fig. 8. NEAT original: cart position  $x$ , pole angle  $\theta$  and force  $F_t$  applied to the cart of the non-Markovian pole balancing task (without velocity information).




Fig. 9. NEAT-SNN: cart position  $x$ , pole angle  $\theta$  and force  $F_t$  applied to the cart of the non-Markovian pole balancing task (without velocity information).

to the computation power of SNNs. We argue evolution can be beneficial to the automatic tuning of these network parameters.

Another active implementation problem is to develop a meaningful approach to transform continuous variables into a spike representation, and vice versa. Rate coding methods either require a relatively long time or large group of homogeneous neurons to calculate a smooth firing rate. On the other hand, temporal coding is considered to be more biologically plausible. Although not yet fully understood, this coding scheme is evidentially able to provide faster signal processing and hence less reaction time.

## VI. CONCLUSION

NEAT is a performance-guaranteed neuroevolution method. This algorithm is tailored in this work to be amenable to SNNs. Through the experiments, we have demonstrated that SNNs can solve continuous control problems by maintaining sufficient spike activities and decoding from weighted spike frequencies. The proposed spiking controller has shown to be more rapid than sigmoidal networks in finding functional solutions. Our work is a first step toward robot control and the results have encouraged further experimentation on more challenging dynamic systems. We expect ensemble-based implementations [23] would provide redundancy to neuroevolution and hence achieve robust performance. We are also looking at integration of synaptic plasticity as a means to online-offline hybrid learning [24]. Our end target is to implement the spiking controller on flight systems such as flapping wing Micro Air Vehicles.

## REFERENCES

- [1] W. Maass, "Networks of spiking neurons: The third generation of neural network models," *Neural Networks*, vol. 10, no. 9, pp. 1659 – 1671, 1997.
- [2] K. L. Rice, M. A. Bhuiyan, T. M. Taha, C. N. Vutsinas, and M. C. Smith, "FPGA implementation of Izhikevich spiking neural networks for character recognition," in *2009 International Conference on Reconfigurable Computing and FPGAs*, Dec 2009, pp. 451–456.
- [3] L. F. Abbott, B. DePasquale, and R.-M. Memmesheimer, "Building functional networks of spiking model neurons," *Nature Neuroscience*, vol. 19, pp. 350 EP –, Feb 2016, perspective.
- [4] D. Floreano, N. Schoeni, G. Caprari, and J. Blynel, "Evolutionary bits'n'spikes," in *Proceedings of the eighth international conference on Artificial life*, 2003, pp. 335–344.
- [5] R. Batllori, C. Laramee, W. Land, and J. Schaffer, "Evolving spiking neural networks for robot control," *Procedia Computer Science*, vol. 6, pp. 329 – 334, 2011, complex adaptive systems.
- [6] D. Howard and A. Elfes, "Evolving spiking networks for turbulence-tolerant quadrotor control," *Artificial Life 14*, 2014.
- [7] H. Paugam-Moisy and S. Bohte, *Computing with Spiking Neuron Networks*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 335–376.
- [8] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, B. Brezzo, I. Vo, S. K. Esser, R. Appuswamy, B. Taba, A. Amir, M. D. Flickner, W. P. Risk, R. Manohar, and D. S. Modha, "A million spiking-neuron integrated circuit with a scalable communication network and interface," *Science*, vol. 345, no. 6197, pp. 668–673, 2014.
- [9] J. S. Seo, B. Brezzo, Y. Liu, B. D. Parker, S. K. Esser, R. K. Montoye, B. Rajendran, J. A. Tierno, L. Chang, D. S. Modha, and D. J. Friedman, "A 45nm CMOS neuromorphic chip with a scalable architecture for learning in networks of spiking neurons," in *2011 IEEE Custom Integrated Circuits Conference (CICC)*, Sept 2011, pp. 1–4.
- [10] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary Computation*, vol. 10, no. 2, pp. 99–127, 2002.
- [11] W. Gerstner and W. Kistler, *Spiking Neuron Models: An Introduction*. New York, NY, USA: Cambridge University Press, 2002.
- [12] E. M. Izhikevich, "Simple model of spiking neurons," *IEEE Transactions on Neural Networks*, vol. 14, no. 6, pp. 1569–1572, Nov 2003.
- [13] F. Theunissen and J. P. Miller, "Temporal encoding in nervous systems: A rigorous definition," *Journal of Computational Neuroscience*, vol. 2, no. 2, pp. 149–162, Jun 1995.
- [14] D. Gamez, A. K. Fidjeland, and E. Lazdins, "iSpike: a spiking neural interface for the iCub robot," *Bioinspiration & Biomimetics*, vol. 7, no. 2, p. 025008, 2012.
- [15] J.-M. Fellous, M. Rudolph, A. Destexhe, and T. Sejnowski, "Synaptic background noise controls the input/output characteristics of single cells in an in vitro model of in vivo activity," *Neuroscience*, vol. 122, no. 3, pp. 811 – 829, 2003.
- [16] A. E. Eiben, J. E. Smith *et al.*, *Introduction to Evolutionary Computing*, 2nd ed., ser. Natural Computing Series. Springer, Berlin, Heidelberg, 2015.
- [17] D. Pardoe, M. Ryoo, and R. Miikkulainen, "Evolving neural network ensembles for control problems," in *Proceedings of the 7th Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO '05. New York, NY, USA: ACM, 2005, pp. 1379–1384.
- [18] J. F. Shepherd III and K. Tumer, "Robust neuro-control for a micro quadrotor," in *Proceedings of the Genetic and Evolutionary Computation Conference*, Portland, OR, July 2010, pp. 1131–1138.
- [19] A. Vandesompele, F. Walter, and F. Rhrbein, "Neuro-evolution of spiking neural networks on SpiNNaker neuromorphic hardware," in *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, Dec 2016, pp. 1–6.
- [20] M. C. Vasu and E. J. Izquierdo, "Information bottleneck in control tasks with recurrent spiking neural networks," in *Artificial Neural Networks and Machine Learning – ICANN 2017*, A. Lintas, S. Rovetta, P. F. Verschure, and A. E. Villa, Eds. Cham: Springer International Publishing, 2017, pp. 236–244.
- [21] T. S. Kang and A. Banerjee, "Learning deterministic spiking neuron feedback controllers," in *2017 International Joint Conference on Neural Networks (IJCNN)*, May 2017, pp. 2443–2450.
- [22] A. G. Barto, R. S. Sutton, and C. W. Anderson, "Neuronlike adaptive elements that can solve difficult learning control problems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-13, no. 5, pp. 834–846, Sept 1983.
- [23] D. Howard, L. Bull, and P.-L. Lanzi, "A cognitive architecture based on a learning classifier system with spiking classifiers," *Neural Processing Letters*, vol. 44, no. 1, pp. 125–147, Aug 2016.
- [24] G. Howard, E. Gale, L. Bull, B. de Lacy Costello, and A. Adamatzky, "Evolution of plastic learning in spiking networks via memristive connections," *IEEE Transactions on Evolutionary Computation*, vol. 16, no. 5, pp. 711–729, Oct 2012.