

This is a repository copy of *Artificial Neural Microcircuits for use in Neuromorphic System Design*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/202697/>

Version: Accepted Version

Proceedings Paper:

Walter, Andrew, Wu, Shimeng, Tyrrell, Andy orcid.org/0000-0002-8533-2404 et al. (5 more authors) (2023) Artificial Neural Microcircuits for use in Neuromorphic System Design. In: ALife 2023: Ghost in the Machine: Proceedings. ALife 2023, 24-28 Jul 2023, University of Hokkaido. MIT Press, JPN, pp. 51-58.

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Artificial Neural Microcircuits for use in Neuromorphic System Design

Andrew Walter,¹ Shimeng Wu¹, Andy M. Tyrrell¹, Liam McDaid², Malachy McElholm²,
Nidhin Thandassery Sumithran², Jim Harkin², Martin A. Trefzer¹

¹ School of Physics, Engineering & Technology; University of York

² School of Computing, Engineering and Intelligent Systems, Ulster University

Abstract

Artificial Neural Networks (ANNs) are one of the most widely employed forms of biomorphic computation. However (unlike the biological nervous systems they draw inspiration from) the current trend is for ANNs to be structurally homogeneous. Furthermore, this structural homogeneity requires the application of complex training & learning tools that produce application specific ANNs, susceptible to pitfalls like overfitting. In this paper, an alternative approach is suggested, inspired by the role played in biology by Neural Microcircuits, the so called “fundamental processing elements” of organic nervous systems. How large neural networks can be assembled using Artificial Neural Microcircuits, intended as off-the-shelf components, is articulated; before showing the results of initial work to produce a catalogue of such Microcircuits through the use of Novelty Search.

Introduction

ANNs, in their various permutations, currently form the backbone of many non-standard computing tools (Prieto, et al., 2016). Typically, the state of the art is focused in one of two directions: large, topologically homogeneous feed forward (figure 1a) & recurrent networks (figure 1b) whose connection weights are adjusted by either increasingly complex machine learning (Schmidhuber, 2015) or genetic algorithms (Shifei, Li, Su, Yu, & Jin, 2013); or iterative evolutionary methodologies designed to produce complex topologies from scratch (Stanley & Miikkulainen, 2002) (figure 1c). With either of these

approaches, the resultant networks are often application specific, which both limits flexibility and raises the specter of issues such as overfitting as a consequence of the limited breadth of training data available.

However, advances in neuroscience are painting a picture of biological nervous system that possess much more nuanced architectures; ones built up of various computational subunits, or Neural Microcircuits, which work together both hierarchically and in parallel (Luo, 2016). Furthermore, these biological systems are more flexible, being able to carry out multiple different tasks, adapt to new ones, or apply approximate information to new situations.

In this paper, this more biological partitioned architecture is used as inspiration, employing novelty search to create a catalogue of Artificial Neural Microcircuits, which can then be used as off-the-shelf components to fashion larger application specific networks, an idea illustrated in figure 1. It is envisaged that this approach will improve the robustness of network behaviours by breaking that behaviour into sub-behaviours that can be the focus of individual specialised Microcircuits or groups of Microcircuits; as well as bringing other advantages such as reducing network development and training overheads by allowing them to be build out of off-the-shelf components and allowing for the updating or alteration of a network’s overall behaviour through swapping of Microcircuits.




Figure 1: An illustration of the various approaches to neural network topologies, including the approach proposed in this paper: a) Feed Forward; b) Recurrent; c) Evolved from scratch; d) Microcircuit based.




Figure 2: the crayfish ventral nerve cord (a), taken from (Demyanenko, Dzreyan, & Uzdensky, 2019); with a block illustration of the ganglia of the abdominal section (b), adapted from (Smarandache-Wellmann & Grätsch, 2014)

Neural Microcircuits

Neuroscientists refer to the multitude of specialised subunits within biological nervous systems as Neural Microcircuits. Commonly referenced examples of such include motor neuron sequencing circuits; the hazard amelioration reflexes of some invertebrates; sensory processing circuits; and the columns of the neocortex. Though varied in form and function, Neural Microcircuits share a common role as the “elementary processing units” of the nervous system (Grillner & Graybiel, 2006).

By way of an example, consider two elements of the locomotion system of *P. leniusculus*, the Signal crayfish. This organism’s primary form of propulsion through the water is via a set of appendages called Swimmerets, which beat back and forth in coordination with one another to produce motion. Control of these appendages is performed by a portion of the crayfish’s nervous system called the Swimmer System, found within four of the six ganglia of the abdominal nerve cord. Each of these ganglia, labeled as A2 through A5 in figure 2, controls the motion of one pair of Swimmerets via motor neurons RS & PS.




Figure 3: Illustration of the Microcircuit within the left hemiganglion of segment A4, counterparts of which are repeated in the segments A2 through A5. Arrows are connections to other microcircuits, dot lines are biochemical synapses, & the resistor represents an electrical synapse. Adapted from (Schneider, Blumenthal, & Smarandache-Wellmann, 2018).

Within each of the relevant hemiganglion is an identical neural Microcircuit, illustrated in figure 3, consisting of a pair of motor neurons and associated inhibitory neurons: Power Stroke (PS), Inhibits Power Stroke (IPS), Return Stroke (RS) & Inhibits Return Stroke (IRS); and three neurons that facilitate the interconnection of the Microcircuits in the different ganglia.

In isolation, the pattern generating and motor neurons produce the rhythmic beating motion of the Swimmerets; but via the three interconnection neurons that motion is modulated, such that all the Swimmerets move in concert and thus efficient movement results (Smarandache-Wellmann & Grätsch, 2014) (Schneider, Blumenthal, & Smarandache-Wellmann, 2018).

However, in situations where the crayfish is threatened, other locomotive behaviours are employed to facilitate fast escape. Collectively referred to Escape Reflexes, they activate the musculature that flexes the whole of the crayfish’s tail to rapidly propel the creature away from danger. Of interest to this work is the Lateral Giant Escape reflex, named because of the role played by the Lateral Giant neurons within the nerve cord.

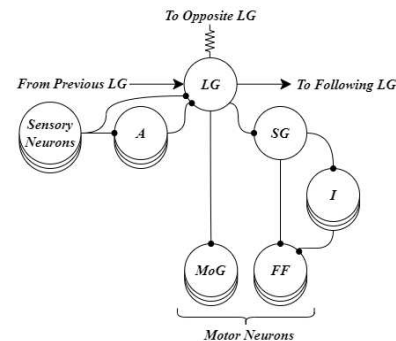


Figure 4: Illustration of the Lateral Giant Escape reflex Microcircuits within the left hemiganglion of segment A2, counterparts of which are repeated in segments A1 through A5. Adapted from (Vu, Berkowitz, & Krasne, 1997).

When stimulus is applied to the sensory neurons on the crayfish’s abdomen, signals are sent to the Lateral Giant (LG) neuron of that segment, either directly or through Interneuron A. As the LG neurons are connected in series down the abdomen and also linked to their counterparts on the opposite side of the animal, sufficient sensory input passes the threshold which is deemed to mean hostile action. The LG neuron then in turn activates the muscles of the tail, causing them to rapidly contract and push the crayfish out of danger. This motor action is caused by the LG directly signaling the Motor Giant neuron (MoG) in each segment, for the immediate motion, and the segments Segmental Giant neuron (SG), with in turn triggers the Fast Flexor neurons (FF) and their interneurons I, which extend the motion beyond the initial impulse (Vu, Berkowitz, & Krasne, 1997) (Edwards, 2017).

Artificial Neural Microcircuits

From this example, it is easy to see an argument for adopting a similar approach in neuromorphic systems; after all, such compartmentalization of different functionality, but with interconnection to allow for coordination of more macro behaviours is seen as a positive side effect if it emerges within neural networks during training or evolution. Furthermore, as individual neural Microcircuits' can be repeated throughout a nervous system wherever similar behaviours are needed, this paper proposes that it should be possible to create a general-purpose catalogue of Artificial Neural Microcircuits, which can be used as building blocks of a wide variety of different more complex neuromorphic systems, while assisting in scalable hardware implementations.

Methodology

In order to assemble a catalogue of Microcircuits that can be used, the following steps need to be taken: (i) generation of candidate Microcircuits; (ii) assessment of them based on some criteria; and (iii) the use of the assessment to select which Microcircuits would be useful components.

Generating Prospective Microcircuits

A Microcircuit is represented using a set of connection matrices: (i) one for external inputs; (ii) one for internal connections; and (iii) a final array for outputs. This approach would allow for the production of arbitrary topologies simply by randomising the values of these matrices, however in this paper a somewhat different direction has been taken.

One of the goals of the proposed Microcircuits-as-components method is to provide a degree of substrate agnosticism. That is to say, the exact spiking neuron model that is used to implement the Microcircuits should be divorced from their functionality. Hence, it is desirable to avoid using individual neurons as the building blocks, and instead employ something "one step up" from this.

This can be achieved while also injecting some biological domain knowledge into the methodology through the use of neural Circuit Motifs. Neurobiologist Liqun Luo identifies ten recurring patterns of connectivity within biological neurocircuits, termed Motifs (Luo, 2016). Each of these motifs consists of two to seven neurons and they exhibit distinct patterns of behaviour.

Initial investigations of the Microcircuit approach employ simplified versions of six of the ten motifs (Byrne, 1997), selected because they represent a breadth of functionalities without using a large number of neurons and connections: FeedForward Excitation (FFE); FeedBack Excitation (FBE); FeedBack Inhibition (FBI); ReCurrent Excitation (RCE); ReCurrent Inhibition (RCI); and LaTeral Inhibition (LTI).




Figure 5: The FeedForward Excitation (FFE) Motif




Figure 6: The FeedBack Excitation (FBE) Motif




Figure 7: The FeedBack Inhibition (FBI) Motif




Figure 8: The ReCurrent Excitation (RCE) Motif




Figure 9: The ReCurrent Inhibition (RCI) Motif




Figure 10: The LaTeral Inhibition (LTI) Motif

In addition to these, an additional “motif” is suggested, in the form of a Central Pattern Generator (CPG). This could be argued to be a Microcircuit in its own right, but here it is included as a “motif” due to its small size (three neurons), and the useful behaviour it adds (the ability to generate a consistently spiking output of a fixed frequency).




Figure 11: The Central Pattern Generator (CPG) Motifs

Each of the above Motifs has an associated connection matrix, which can be tiled together, and connections between the Motifs can then be randomly initialised to produce a potential Microcircuit. Note that, as the Motifs have defined input & output connections, not all possible connections

between Motifs are allowed: e.g. a connection could not end at neuron 2 of the FBE Motif, nor start at neuron 1.

Thus, the matrices of two input, one output Microcircuit consisting of two Motifs might be as illustrated in figure 12: with the associated topology shown in figure 13.

$$\begin{aligned}
 a) & \begin{pmatrix} 0.48 & 0.0 & 0.0 & 0.0 & n/a \\ 0.0 & 0.39 & 0.0 & 0.0 & n/a \end{pmatrix} \\
 b) & \begin{pmatrix} 0.0 & 1.0 & 1.0 & 0.69 & n/a \\ 1.0 & 0.0 & 1.0 & 0.82 & n/a \\ 1.0 & 1.0 & 0.0 & 0.14 & n/a \\ n/a & n/a & n/a & 0.0 & 1.0 \\ 0.07 & 0.22 & 0.14 & 1.0 & 0.0 \end{pmatrix} \\
 c) & (4)
 \end{aligned}$$

Figure 12: Example input (a); internal (b); and output connection (c) matrices.

Behavioral Assessment

Once a selection of prospective Microcircuits exists, it is necessary to assess their behaviour. This step is conceptually straightforward, as all that is required is to record the response of the various Microcircuits to a given stimulus or set of stimuli. Ideally, there would be a standard stimulus or set of stimuli which could provide a sufficiently comprehensive set of responses from the Microcircuit so as to consider its behaviour fully mapped, i.e., all possible forms of output from that Microcircuit have been solicited and recorded. However, the creation of such a test stimulus or set of stimuli has not yet been undertaken, hence preliminary experiments employed a less generic stimulus pattern, described later in this paper.




Figure 13: An example Microcircuit, consisting of a CPG (neurons 0, 1 & 2) and a FBE (neurons 3 & 4)




Figure 14: An illustration of the complete Microcircuit catalogue generation methodology

Catalogue Selection

Regardless of the exact test stimuli used, the result is a set of Microcircuits, each with one or more output spike trains which illustrate their response to the stimuli. Using these, it is then necessary to determine whether a Microcircuit is suitable for addition to the component catalogue. To increase the coverage of Microcircuits, Novelty Search is used to produce a catalogue that encompasses a sufficiently large breadth of the behavioral space.

Novelty Search is an alternative to fitness function based evolutionary algorithms. Instead of assessing individuals based on how successful their behaviours are at meeting some explicitly defined criteria, they are instead compared on the basis of how different their behaviours are from one another. Individuals surpassing some threshold of difference are then added to an archive, which forms the output of the search (Lehman & Stanley, 2008). Originally this archive was then searched to locate a final “best” individual, but in this use case, the archive constitutes the desired final product: a catalogue of Microcircuits with varying behaviours. Unlike fitness-based evolution, this approach allows for a catalogue to be created without needing to characterise any of the desired behaviours or even without knowing what those behaviours might be.

To employ Novelty Search, a method to compare the output spike trains of the various Microcircuits is required, ideally one that will provide a concise numerical value related to the difference between them. Neuroscientists have compared the recorded spike trains of biological neurons, and have produced a wide range of methods that can be applied in this work (Kreuz, Measures of spike train synchrony, 2011). Bivariate SPIKE-Distance has been chosen, a measure that has the advantages of taking a wide range of spike train metrics into consideration, while producing a single value. This value is between 0.0 and 1.0 and corresponds to the “distance” between the two spike trains, with a value of 0.0 only being returned if the two spike trains are identical (Kreuz, Chicharro, Greschner, & Andrzejak, 2010). By applying this measure to each unique pairing of Microcircuits, a distance matrix can be generated, with the average of the distances between one individual and all others being the metric used by the Novelty Search.

All that remains is to iterate the population of prospective Microcircuits in the standard fashion (using crossover & mutation as part of selection), before repeating the process to produce a suitable catalogue, as illustrated in figure 12.

Simulated Data Bus Experiment

To establish if the previously detailed methodology is sufficient to produce a selection of useful Microcircuits, it is necessary to put it to the test. For this purpose, an emulation of an 8-bit bus is used as the input stimulus, with the intended output being Microcircuits that displayed clear and consistent responses to different parts of the data streamed.

Experiment Stimulus

To provide an input stimulus that is meaningful in both quantity and quality, this experiment used a sample of text. This sample consisted of 2030 characters, including spaces, complete with punctuation. Each character of the text sample was translated to a single byte, following the UTF-8 encoding scheme, with these binary values then used to produce a set of eight input spike trains. A simple encoding scheme was used, where a binary 1 produced a 25ms long burst of spikes on the associated input channel (25ms was selected as it was the same as τ of the example spiking neuron being used).




Figure 15: A section of the input stimulus. Note that there are only 7 input spike trains shown, as UTF-8 does not use the Least Significant Bit (LSB) for all characters.

It should be noted that, while in the context of this experiment it is specific characters that are being considered, any Microcircuits that display the desired kind of behaviours will be doing so because they have latched onto the pattern of input spike trains those characters have been mapped to. Those

patterns could correspond to other stimuli of interest, for example: control codes on a microprocessor's instruction bus. Furthermore, validating the ability of this approach to reliably produce Microcircuits that latch onto patterns in a useful fashion would indicate its broad applicability, as not only could any conceivable input of interest be encoded thus, but the outputs of the Microcircuits themselves are spike train patterns, meaning that there is nothing to prevent them being assembled sequentially.

Experiment Parameters

The population of Microcircuits consisted of 100 individuals and was iterated 50 times. All Microcircuits had eight inputs and one output, with the starting pool of Microcircuits all having two Motifs. Input and internal connections had a 25% chance of existing (tuned to allow for suitably sparse connectivity) and a weight value between -1.0 and 1.0.

Each iteration of the population produced new Microcircuits by combining two Microcircuits from the previous iteration: one with a high average Bivariate SPIKE-Distance and one with a low average. This avoided the possibility of self-crossover, while also aiding diversity. Crossover ratio was 4:6 in favor of the low SPIKE-Distance individual. Following crossover, mutation was applied with the following probabilities: 35% chance to add a new Motif; 60% chance to replace a Motif; 60% chance to alter the weight of an input or internal connection; and a 60% chance of changing an output connection. These values were deliberately set quite high so as to fuel diversity within the population.

To be added to the catalogue, a Microcircuit needed to possess an average Bivariate SPIKE-Distance of at least 0.5, though this value was dynamic, as suggested in (Lehman & Stanley, 2008). If 10 iterations of Microcircuits passed without an addition being made to the catalogue, the threshold was reduced by 5%. Alternatively, if 10 individuals were added to the catalogue in the space of one iteration, the threshold was raised by 20%.

Results

From a total pool of 5000 Microcircuits created over the course of the experiment, 50 were deemed to display sufficiently novel behaviour to warrant transfer to the catalogue, with average Bivariate SPIKE-Distance values of between 0.44 and 0.63. As is to be expected from the increase in complexity that occurs in novelty searches (shown in figure 16), the number of motifs, and as such neurons, within the Microcircuits increased steadily, with the largest instance coming from generation 45: possessing 61 neurons across 20 motifs.




Figure 16: Average number of neurons (blue) and motifs (orange) of Microcircuits added to the catalogue, showing the trend towards increasing complexity.

The breakdown of which motifs are present in those Microcircuits added to the catalogue at which point in time is shown in figure 17. This data suggests there is no strong inclination towards a specific ratio of motifs for novel microcircuits, nor that the presence or absence of a given motif is of any specific importance.




Figure 17: Percentage breakdown of the motif makeup of Microcircuits added to the catalogue in a given generation

To determine if any of the Microcircuits in the catalogue demonstrated clearly useful behaviour, a correlation was determined between the output spike trains of the Microcircuits and the character the input spike trains were presenting. To do this, the average number of output spikes across each of the input characters was calculated. If this value was zero, the Microcircuit was not responding to a given character; values less than one would indicate that a character was only responded to sometimes (a weak correlation); while any value greater than one would indicate the Microcircuit spiked more than once on each instance of a character (a strong correlation)

Across all 50 Microcircuits in the catalogue, 8 Microcircuits exhibited strong correlations for all characters, which on closer investigation corresponded to different forms of consistent spiking. Of the others: 20 Microcircuits displayed weak correlations to one or more characters; 15 displayed a mixture of strong and weak correlations, while the remaining 7 displayed strong correlations with specific subsets of characters. Of these 7, two are of particular note: Microcircuits 392 & 466 (figures 18 and 19).

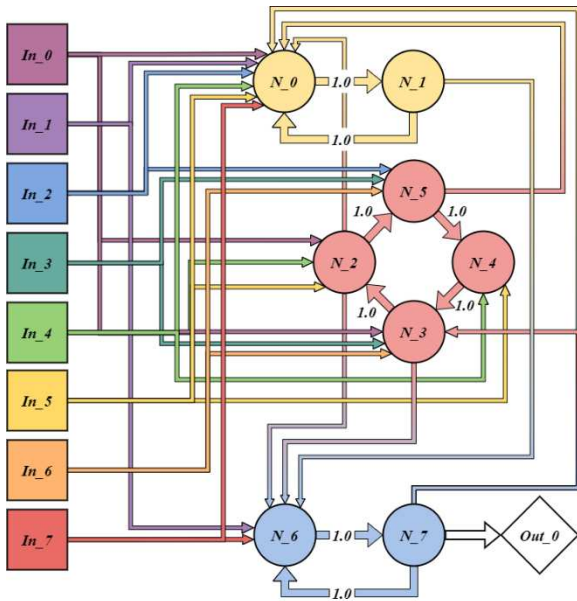


Figure 18: Illustration of Microcircuit 392

Microcircuit 392, produced in the fourth generation with 8 neurons across 3 motifs (FBE, RCI & FBE), displayed a strong correlation (average of 1.20 spikes) with the dashes in the sample text.

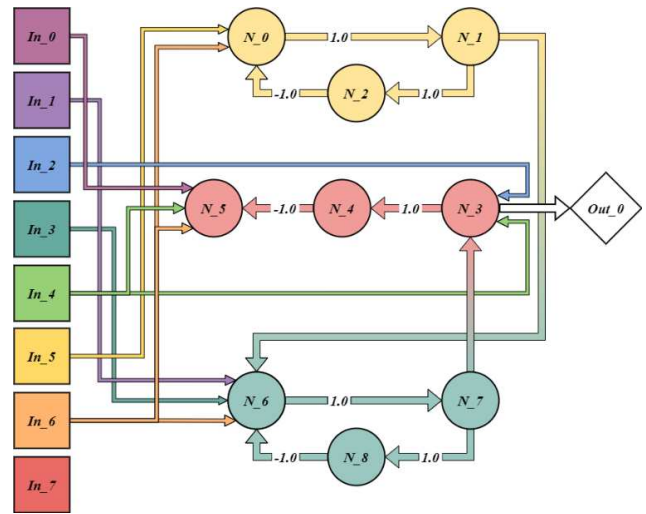


Figure 19: Illustration of Microcircuit 466

Microcircuit 466, produced in the fifth generation with 9 neurons across 3 motifs (FBI, LTI & FBI), displayed an even stronger correlation with all four non-letter characters in the sample text (Space: 4.31; Comma: 4.50; Dash: 3.80; Stop: 4.33).

Conclusions

Both of the two Microcircuits highlighted have clear uses within a larger neural network, indicating that the proposed approach to building networks has merit. It is easy to see how a more expansive version of the Microcircuit catalogue from an experiment like this could be used to assemble a network able to completely parse UTF-8 text streams, but of course it should be possible to go further.

The exact same Microcircuits could be used as part of a network that takes in handwritten characters, with other kinds of Microcircuit forming the interface between the character images and the character recognisers; or in a completely different application where the 8-bit patterns they recognise are some other stimuli of interest. In addition, the local connections within a Microcircuit support more scalable neuromorphic implementations as inter-microcircuit connectivity does not follow the traditional ANN arrangement of one-to-many neurons.

With the eventual goal being the creation of a usefully expansive catalogue of Microcircuit components, there are obviously steps to take before that point. A number of areas for improvement in the generation methodology have been identified, including updates to the selection of motifs; refinement of the novelty metric; and alterations to the Microcircuit population iterator to avoid unconstrained growth.

It should also be noted that these initial experiments have been focused on the generation of Microcircuits that react to spatial input patterns (i.e. which inputs are or are not spiking). A vital next step is to expand these efforts to investigate temporal patterns (i.e. sequences of spikes on both the same or multiple inputs).

Acknowledgements

This work was carried out as part of the NERVOUS Systems project, a joint effort of the University of York & Ulster University; funded by the UK's Engineering & Physical Sciences Research Council (EPSRC) under grants EP/W003759/1 & EP/W003783/1; working in collaboration with Xilinx Inc, ARM Ltd & Thales Alenia Space.

References

- Byrne, J. H. (1997). *Introduction to Neurons and Neuronal Networks*. Retrieved from Neuroscience Online: <https://nba.uth.tmc.edu/neuroscience/s1/introduction.html>
- Demyanenko, S. V., Dzrean, V. A., & Uzdensky, A. (2019). Axotomy-Induced Changes of the Protein Profile in the Crayfish Ventral Cord Ganglia. *Journal of Molecular Neuroscience*, 667–678.
- Edwards, D. (2017). *Crayfish Escape*. Retrieved from Oxford Research Encyclopedia of Neuroscience: <https://oxfordre.com/neuroscience/view/10.1093/acrefore/9780190264086.001.0001/acrefore-9780190264086-e-158>
- Grillner, S., & Graybiel, A. M. (2006). *Microcircuits: the Interface Between Neurons and Global Brain Function*. MIT Press.
- Kreuz, T. (2011). *Measures of spike train synchrony*. Retrieved from Scholarpedia: http://www.scholarpedia.org/article/Measures_of_spike_train_synchrony
- Kreuz, T., Chicharro, D., Greschner, M., & Andrzejak, R. G. (2010). Time-resolved and time-scale adaptive measures of spike train synchrony. *Journal of Neuroscience Methods*, 92-106.
- Lehman, J., & Stanley, K. O. (2008). Exploiting Open-Endedness to Solve Problems Through the Search for Novelty. *Proceedings of the Eleventh International Conference on Artificial Life* (pp. 329-336). Cambridge: MIT Press.
- Luo, L. (2016). *Principles of Neurobiology*. Abingdon: Garland Science.
- Prieto, A., Prieto, B., Ortigosa, E. M., Ros, E., Pelayo, F., Ortega, J., & Rojas, I. (2016). Neural Networks: an Overview of Early Research, Current Frameworks and New Challenges. *Neurocomputing*, 242-268.
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 85-117.
- Schneider, A., Blumenthal, F., & Smarandache-Wellmann, C. (2018). Adaptive Encoding of Coordinating Information in the Crayfish Central Nervous System. *Preprint*. Retrieved from bioRxiv.
- Shifei, D., Li, H., Su, C., Yu, J., & Jin, F. (2013). Evolutionary Artificial Neural Networks: a Review. *Artificial Intelligence Review*, 251-260.
- Smarandache-Wellmann, C., & Grätsch, S. (2014). Mechanisms of Coordination in Distributed Neural Circuits: Encoding Coordinating Information. *Journal of Neuroscience*, 5627–5639.
- Stanley, K. O., & Miikkulainen, R. (2002). Evolving Neural Networks through Augmenting Topologies. *Evolutionary Computation*, 99-127.
- Vu, E. T., Berkowitz, A., & Krasne, F. B. (1997). Postexcitatory Inhibition of the Crayfish Lateral Giant Neuron: A Mechanism for Sensory Temporal Filtering. *Journal of Neuroscience*, 8867-8879.