Network of Spiking Neurons Driven by Compression

Alexander Gain: Department of Computer Science, Tulane University
Lawrence Holder: School of EE and Computer Science, Washington State University

Our work aims to design an intelligent agent that chooses its actions based on compression as a reward signal. The design falls into the category of spiking neural network. In the scenarios tested, the goal of the agent is to compress an input stream of bytes. Neurons are organized by layers and connected to other neurons at adjacent layers. A neuron receives an increase in voltage based on how well its associated action compresses the input stream. The agent performs actions, based on the neuron, according to a probability that's calculated based on how successful the associated action has been at compressing the input stream and the strength of the neuron's connections to neighboring neurons that recently fired. The agent has been successful at learning to perform an action out of a subset of actions that leads to the most compression, and at learning sequences of actions that lead to significant or optimal compression. In some specific scenarios, the agent performs close to optimal actions that lead to significant compression rather than the optimal sequence of actions. There are limitations on the length of sequences of actions that can be learned, and some specific types of sequences could not be learned given the current structure, e.g., sequences that change with each iteration could not be learned. This agent compresses data in a novel and effective way and shows promise for simultaneously displaying intelligent behavior and compressing data intelligently.

The model may be conceptualized as comprised of the following primary components: Neurons, structure, performing actions, reward signal, spiking, voltage, and connections. Neurons can be thought of as nodes which are connected to each other. The neurons are iterated over periodically to perform actions. Each neuron performs an action with probability p, which is dependent on the number of times the neuron has spiked and the strength of its connection to its input neuron. The voltage of a neuron is a real number that starts at 0 and decays over time when positive. It increases if the neuron's action leads to an increase in reward signal. A neuron spikes if its voltage crosses a certain threshold. Then, the voltage resets to a value of 0. Connections between neurons are strengthened if a neuron previously fired and another neuron fires soon after.

This agent is able to learn sequences of actions that maximize, or come close to maximizing, the reward signal in several scenarios. Given that the input data to the neurons is a stream of bits and that the reward signal is the compression of this stream of bits, the agent was able to learn sequences of actions that came close maximizing compression in several scenarios. In general, it was able to learn a sequence of n actions given an $n \times n$ structure of neurons.

Using compression as a reward signal, this agent was able to choose actions that led to compression of data in a variety of scenarios. In the majority of cases, the agent chose the optimal set or sequence of actions that led to the most compression.

Although the agent has limitations, it performed optimally in the majority of scenarios tested. Also, the spiking neural network variation is novel. Taking these facts into consideration, further development of this model is promising.

One direction to take is to create separate networks that are connected to each other and serve different purposes. For example, there could be a lower-level network and a higher level network. The lower-level could handle basic compression operations such as replacing recurring substrings or bits, whereas the higher level could perform higher-level operations such as moving a camera, which consequently would change the input data. How well the lower level compresses after a certain duration of time could be the reward signal for the higher level. Thus, if moving the camera to a certain position caused the lower level to compress better, then the higher level would continue keeping the camera at that position.

If this direction is taken and the results are positive then the agent would compress data in a novel and effective manner while possibly displaying intelligent behavior. Thus, continuing further development of this model appears to be a promising course of action.

1068-0314/16 $31.00 © 2016 IEEE
DOI 10.1109/DCC.2016.116