

Introduction

The paper "Evolving Neural Networks through Augmenting Topologies" aims to propose a method, NeuroEvolution of Augmenting Topologies (NEAT), that is able to gain an advantage over other NeuroEvolution strategies by evolving its network topology along with the weights. This is in contrast to traditional NeuroEvolution approaches where the topology is chosen and fixed before the experiment begins. There has been much academic discussion around the effectiveness of evolving topologies. One such argument supporting the evolution of both topology and weights was put forward by Gruau et al. (1996), where a method that solved a pole-balancing benchmark problem was provided. However, it was later shown that the structure was not necessary and a fixed-topology method that convincingly outclassed this method was subsequently released. With the inconclusive nature of these results, the paper aims to demonstrate that evolving topologies and weights can significantly enhance the performance of a NeuralEvolution approach. According to the paper, such performance enhancements are achieved because of several key factors. The first of which is using crossover to combine different topologies. Next, structural innovation is protected using speciation, and lastly, the topologies are only grown incrementally and minimized throughout the entire evolution process. In addition, considering a method that uses evolving topologies also raises several concerns. For example, how do we crossover topologies that are significantly different from each other and how do we prevent the premature termination of potentially successful topologies? The paper also presents solutions to these concerns.

We chose this paper as all of our group members found the concept of applying evolutionary strategies to deep learning interesting. We were intrigued by the idea of solving computer science problems using an interdisciplinary approach that draws inspiration from the natural world. Additionally, we were very impressed by past implementations of this method.

Lastly, this approach is best suited to solve reinforcement learning problems. That is also why we have chosen to implement the methods of this paper to see how it performs on a video game, such as Flappy Bird.

Challenges

It was challenging to gain a full understanding of the research paper as it uses concepts and vocabulary from the field of evolutionary biology. As such, we had to refer to other forms of academic literature to gain the contextual knowledge required to understand and implement the research paper.

The conversion of the genotype (the list of connections and weights in the network) to the network itself has proved a bit more challenging than anticipated (a graph traversal needs to be done to find the number of hidden layers and I've run into some problems when there's cycles), but these challenges are mostly handled and it should be ready to run on our sample problems by the end of the week.

Insights

Currently we have all the major functions (breeding, mutation, initialization, and speciation) complete (and tested) that we will need for the project. Given our tests work for these, we don't

anticipate having to go back and change them. We do not have any results yet as we are still completing the last components of the project, but we are on track to have results as soon as our next components are complete

Plan

Moving forward, we are in the process of implementing several aspects of the project, including converting our genome data structure into a network, integrating the network with the Flappy Bird game, and the training process. As of now, we do not foresee any changes to our plans, but we will adapt if necessary.

We have three major components of our project left to complete:

- 1) The conversion of the genotype (list of nodes and connections) to a network
- 2) The loop that performs the evolutionary functions in order
- 3) The integration with flappy bird.

As of now, we do not foresee any changes to our plans, but we will adapt if necessary.