

INTRODUCTION

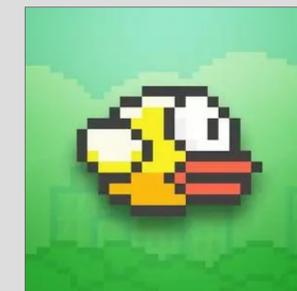
Machine learning is the process by which computers can learn to accomplish difficult tasks. Rather than requiring a programmer to create step-by-step instructions for a computer, machine learning instead allows the computer to learn by itself, without the need for much instruction from a programmer. Machine learning algorithms come in many forms, but they often try to imitate natural processes - the most common types of machine learning mimic the way brains learn or mimic the evolution of organisms.

NEAT is an example of one of these machine learning algorithms. It models evolution with neural networks acting as the "organisms" evolving (neural networks are a common object in machine learning - think of them as simplified digital brains). This research aims to determine the changes that can be made to this learning process to speed it up, improve the effectiveness of the final product, or in other ways improve the process.

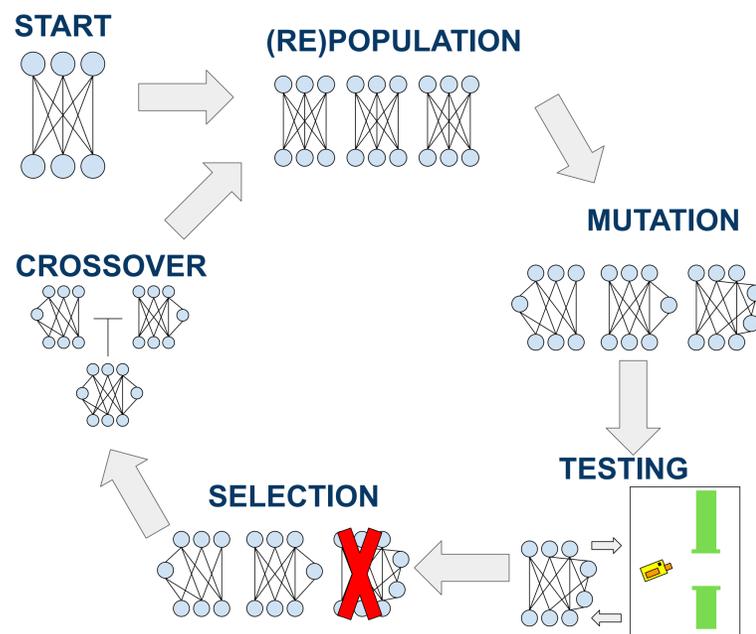
RESEARCH METHODOLOGIES

The experiment consisted of running **750** instances of the NEAT algorithm in order to teach it how to play Flappy Bird. Certain parameters (population size, mutation rate multiplier, and species elitism - basically, how different two NEAT neural nets can be before they are considered different species) were changed. 6 different mutation rates, 5 population sizes, and 5 values for species elitism were tested, with each combination having 5 trials done. Each trial consisted of 100 generations in which the agents were tested against a game of Flappy Bird. Data was collected on time each generation took to compute and the average score per generation. Score was calculated like this:

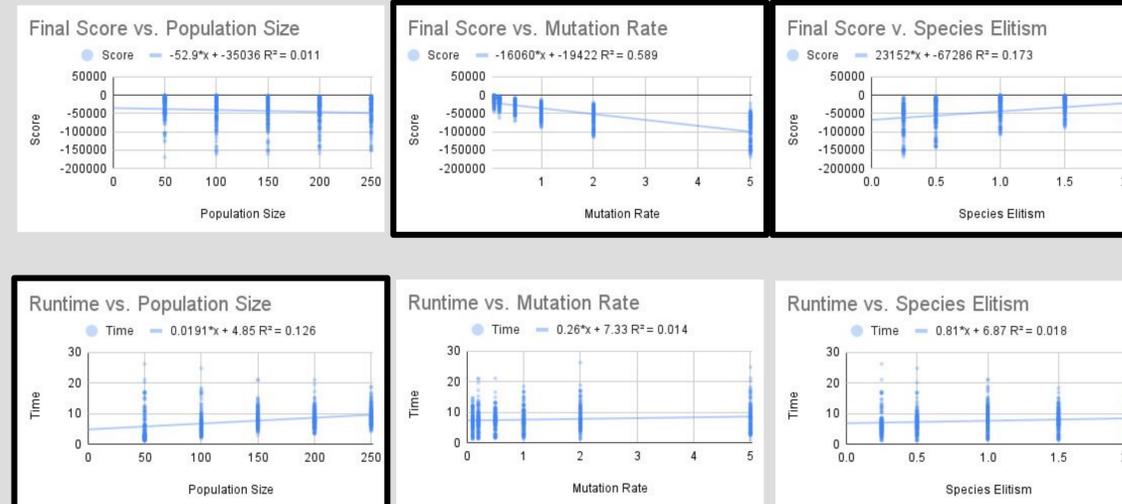
- 1 point per game tick survived (no active collisions)
- -50 for collision with a pipe
- -100 for each game tick spent out of bounds



The NEAT Algorithm, Visualized



DATA AND FINDINGS



CONCLUSIONS, IMPLICATIONS, AND NEXT STEPS

Conclusions

- The data supports the hypothesis that selecting good parameters is important for the performance of the NEAT algorithm in some cases

Implications

- Not many conclusions can be reliably extended to all uses of the NEAT algorithm, due to issues with the experiment described below

Next steps

- Conduct testing against a more complex task
 - The simplicity of Flappy Bird made it possible for NEAT to create successful players in the first generation
- Test against a task that does not require randomization
 - The random generation of Flappy Bird levels led to substantial variance in the score during individual training runs, which complicated analysis
- Test against a larger range of tasks
 - Testing only against Flappy Bird doesn't show how these parameters affect the NEAT algorithm in general
 - Other types of tasks (ie. competitive games like chess) may produce to different results

DISCUSSION, ANALYSIS, AND EVALUATION

Analysis of raw data consisted of determining the point at which each trial roughly stopped improving (found by taking the maximum of the data after smoothing and a slight downward shift increasing as time goes on in order to eliminate the effect of small improvements toward the end). This gave a good enough estimate for the "endpoint" of each instance of training.

Statistical analysis of the final scores at the endpoints and the time it took to reach those endpoints was conducted. T-tests for linear regression were performed.

The most significant trends shown by the data were a strong negative correlation between mutation rate and final score ($r^2=0.589$, $t>32$, $p\approx 0$), a weak positive correlation between population size and runtime ($r^2=0.126$, $t>10$, $p\approx 0$), and a weak positive correlation between species elitism and final score ($r^2=0.173$, $t>12$, $p\approx 0$).

The strengths and significance of these correlations (especially given the large sample size) makes them useful as evidence for forming conclusions.

FULL RESEARCH PAPER

***Special thanks to Shayan Doroudi and Jun Shen for helping make this project possible.

