Sparse Cosine Optimized Policy Evolution for Hexapod Gait Generation

Jim O'Connor, Jay B. Nash, Derin Gezgin, Prof. Gary B. Parker

Autonomous Agent Learning Lab, Department of Computer Science

Connecticut College

Introduction

Hexapods (6-legged robots) can learn to walk using evolutionary algorithms, but as input data grows, the learning becomes harder and slower due to the exponential growth in search-space. Our method, **Sparse Cosine Optimized Policy Evolution (SCOPE)**, helps solve this by compressing complex input data using the Discrete Cosine Transform (DCT). This keeps the most useful features while removing noise, allowing the robot to learn better gaits with fewer parameters.

Methodology

SCOPE helps a hexapod robot learn to walk by compressing its recent motion data. It starts with the last 50 poses, which include position, velocity, and acceleration for each motor, adding up to 2,700 values. Instead of sending all of these values to the controller, we apply the **Discrete Cosine Transform (DCT)** to shrink the input while keeping the important features, compressing the input down to just 54 values.

These 54 compressed numbers are passed to a simple controller that decides how the robot should move. The controller outputs parameters for sine wave patterns that drive each motor smoothly, producing natural, coordinated walking motions.

To improve the controller over time, we use a **Steady-State Genetic Algorithm (SSGA)**. The SSGA evolves the controller by trying small changes and keeping the ones that help the robot move farther. This controller runs continuously without resetting the simulation.

Together, the DCT compression and steady learning loop allow our robot to evolve strong walking behaviors, even with a huge amount of input data.

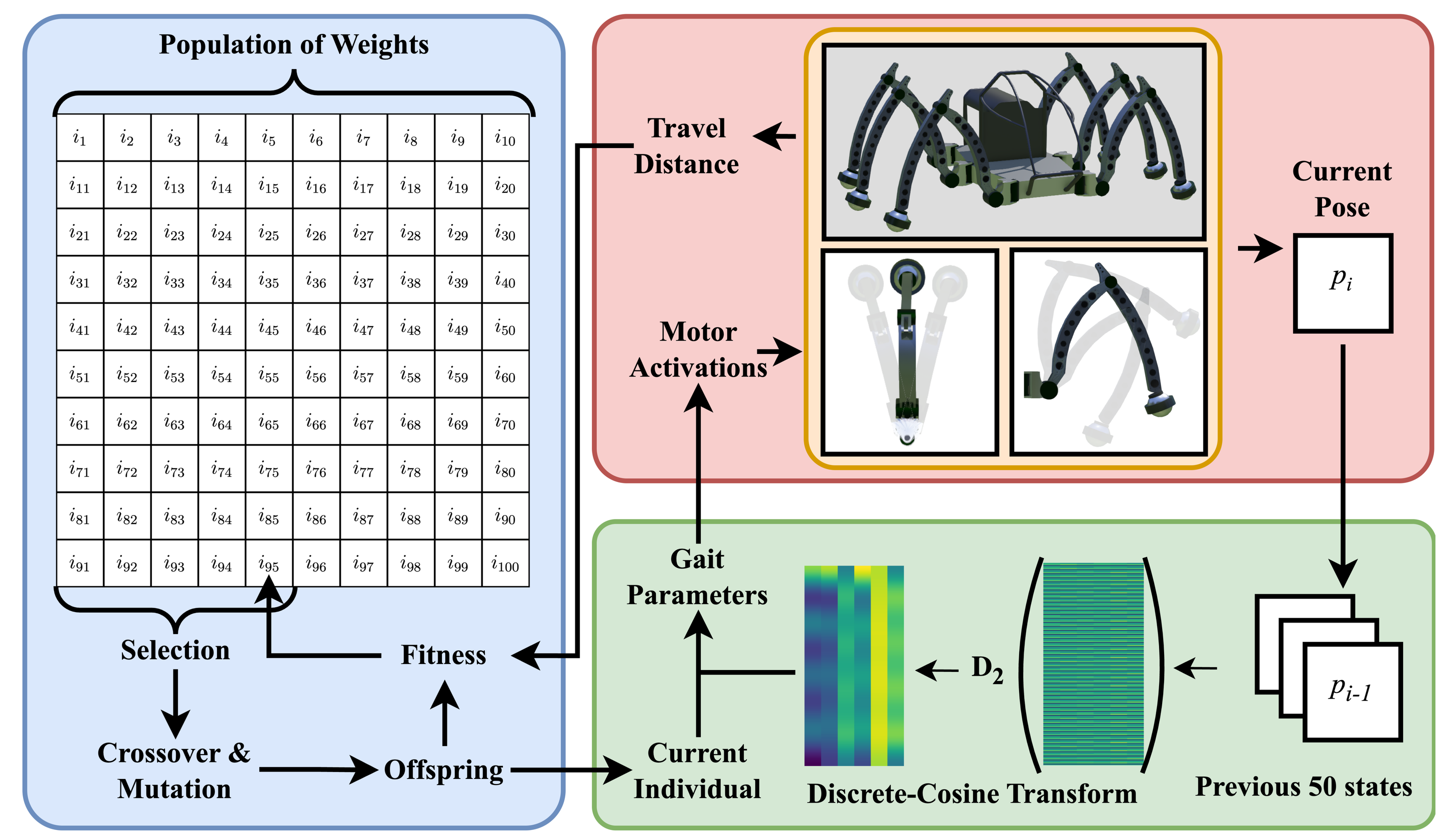

Figure 1. Overview of the complete methodology. The Steady-State Genetic Algorithm provides evolved policy weights for SCOPE, which decides the gait generation parameters for the Mantis robot based on the previous 50 states. A solution is assigned the distance traveled as its fitness, and is used to replace the worst individual in the tournament selection.

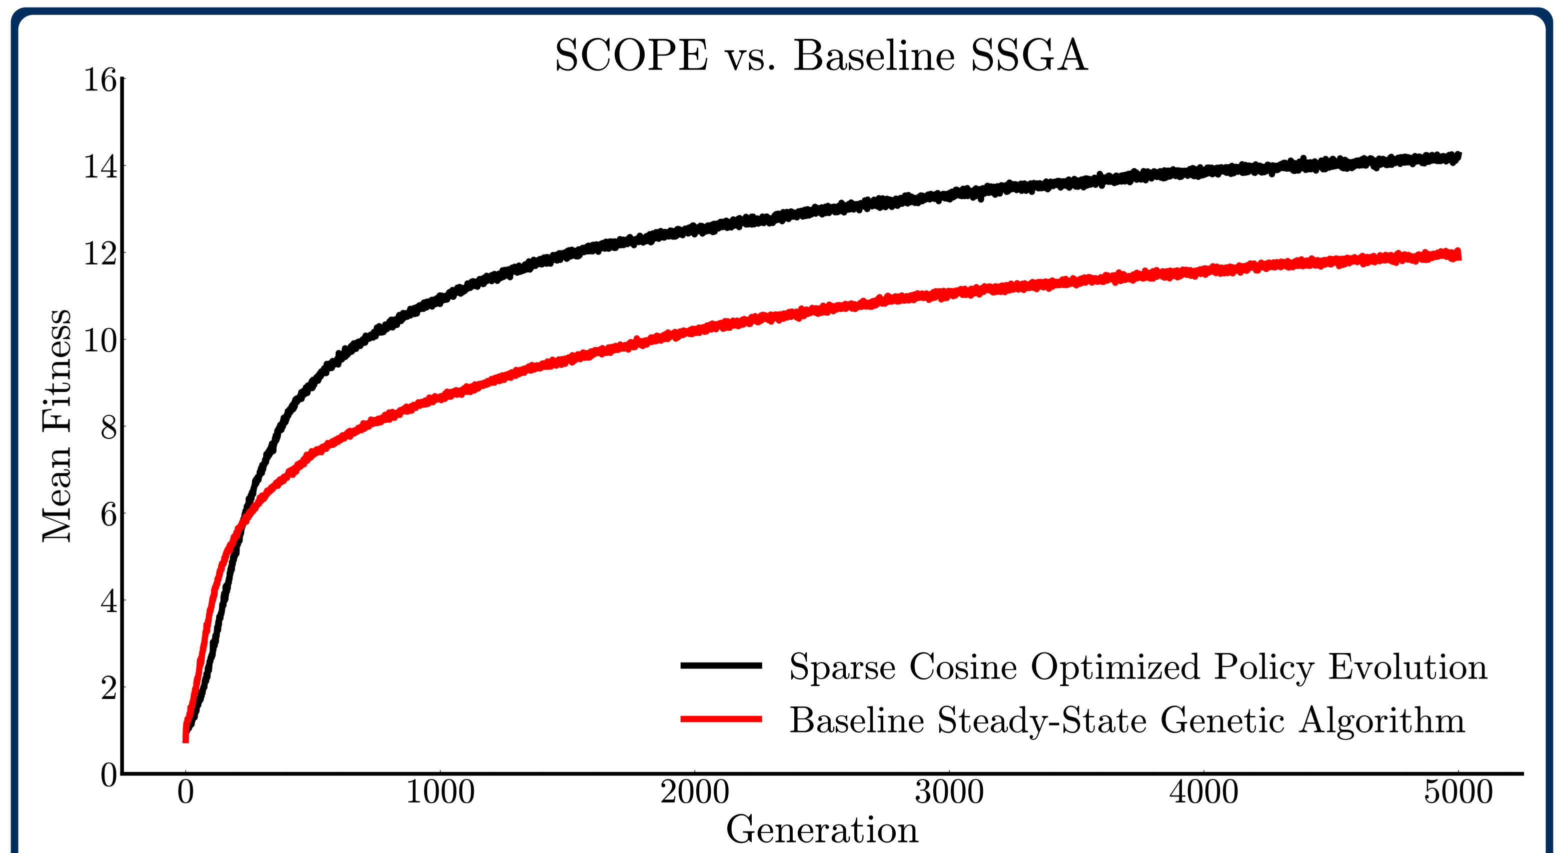

Figure 2. The mean fitness curve of a Steady-State Genetic Algorithm optimizing the weights of the gait generation policy with and without the use of SCOPE. Each algorithm was allowed to evolve over 5000 generations and the graph shows an average of 500 individual trials for both algorithms.

Results & Conclusion

Adding SCOPE to a baseline evolutionary controller resulted in a consistent improvement in learning performance. Across 500 trials, controllers using SCOPE achieved an average fitness of **14.24**, while the baseline averaged **11.88**, resulting in a **20% gain**. This was achieved despite compressing 2,700 input values down to only 54 using DCT, a **98%** reduction

By discarding noisy or redundant components and preserving key structure, SCOPE made it easier for the evolutionary algorithm to find useful solutions in a smaller search space. This shows that effective gait generation is possible even under significant input compression, without sacrificing quality.

Future Work

Following the successful applications of SCOPE in robotics and Atari *Space Invaders*, we are now expanding SCOPE to the entire Atari 2600 benchmark. This will test SCOPE's generalizability across a diverse set of games and visual environments.

Future research will also focus on developing adaptive sparsification methods to make the compression more dynamic. We will continue to integrate SCOPE with alternative evolutionary strategies, to evaluate its benefits across broader optimization frameworks.

Acknowledgments

- Summer Science Research Institute
- Office of the Dean of the College

An academic paper on this work was published in the 17th International Conference on Evolutionary Computation Theory and Applications