Playing Atari Space Invaders with Sparse Cosine Optimized Policy Evolution

Jim O'Connor Jay B. Nash Derin Gezgin Gary B. Parker
Autonomous Agent Learning Lab, Dept. of Computer Science, Connecticut College

Abstract

Evolutionary algorithms often struggle with the large visual state spaces found in games such as those in the Atari Learning Environment. **Sparse Cosine Optimized Policy Evolution (SCOPE)** compresses high-dimensional inputs using the Discrete Cosine Transform (DCT) followed by sparsification, retaining only the most energetic coefficients. This frequency-domain compression enables the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) to evolve compact policies efficiently. When applied to *Space Invaders*, SCOPE outperforms both standard evolutionary methods (e.g., OpenAI-ES, HyperNEAT) and classic reinforcement learning baselines (DQN, A3C) using just 875 parameters.

Method Overview

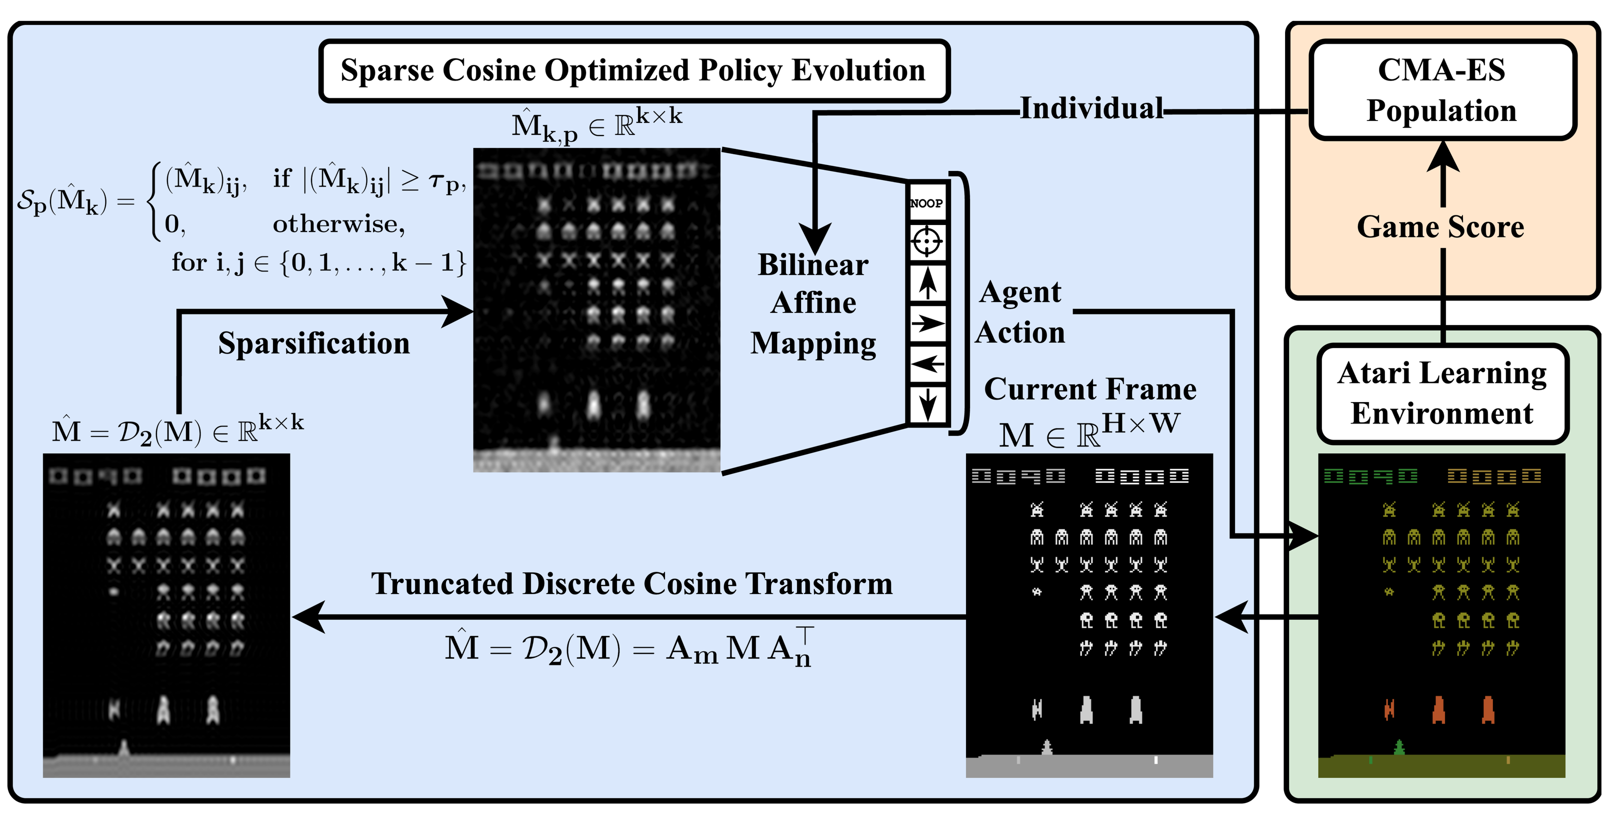

Figure 3. SCOPE pipeline: grayscale frame → 2D DCT → $K \times K$ truncation → percentile sparsification → affine mapping → CMA-ES optimization.

Experimental Setup

Environment: *ALE/SpaceInvaders-v5*, grayscale frames, frame skip 4, six discrete actions.

Training: CMA-ES with 5000 generations, 10,000 steps per episode. Parameter sweeps over truncation K and sparsity P determine optimal compression. An initial parameter sweep used only 500 generations with CMA-ES to find promising candidate parameter pairs.

Conclusions

SCOPE demonstrates that aggressive compression and sparsification in the frequency domain can yield highly efficient visual policies without deep architectures or backpropagation. Despite using under one thousand parameters, it achieves competitive scores in *Space Invaders*, confirming that compact policy search in a transformed space is both tractable and effective.

Results

We found that the parameters $K = 125$, $P = 10$ resulted in the best performance, achieving a peak score of 4,065 and a mean score of ≈ 2,000 over 50 runs. This outperforms many other methods with orders of magnitude more parameters.

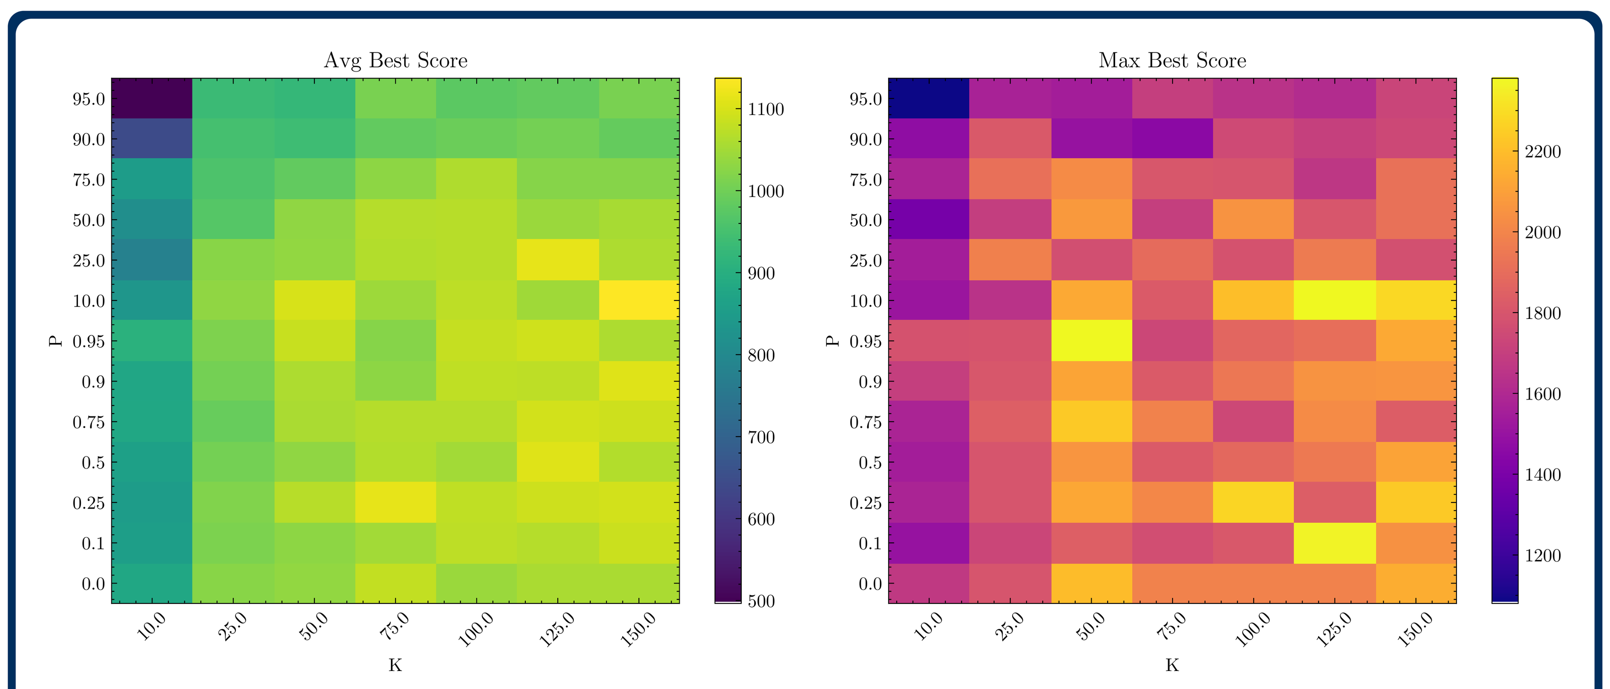

Figure 1. Heatmap of mean scores over 100 trials. Moderate truncation ($K \in [100, 150]$) and low-to-moderate sparsity ($P \in [5, 25]$) yield the best region.

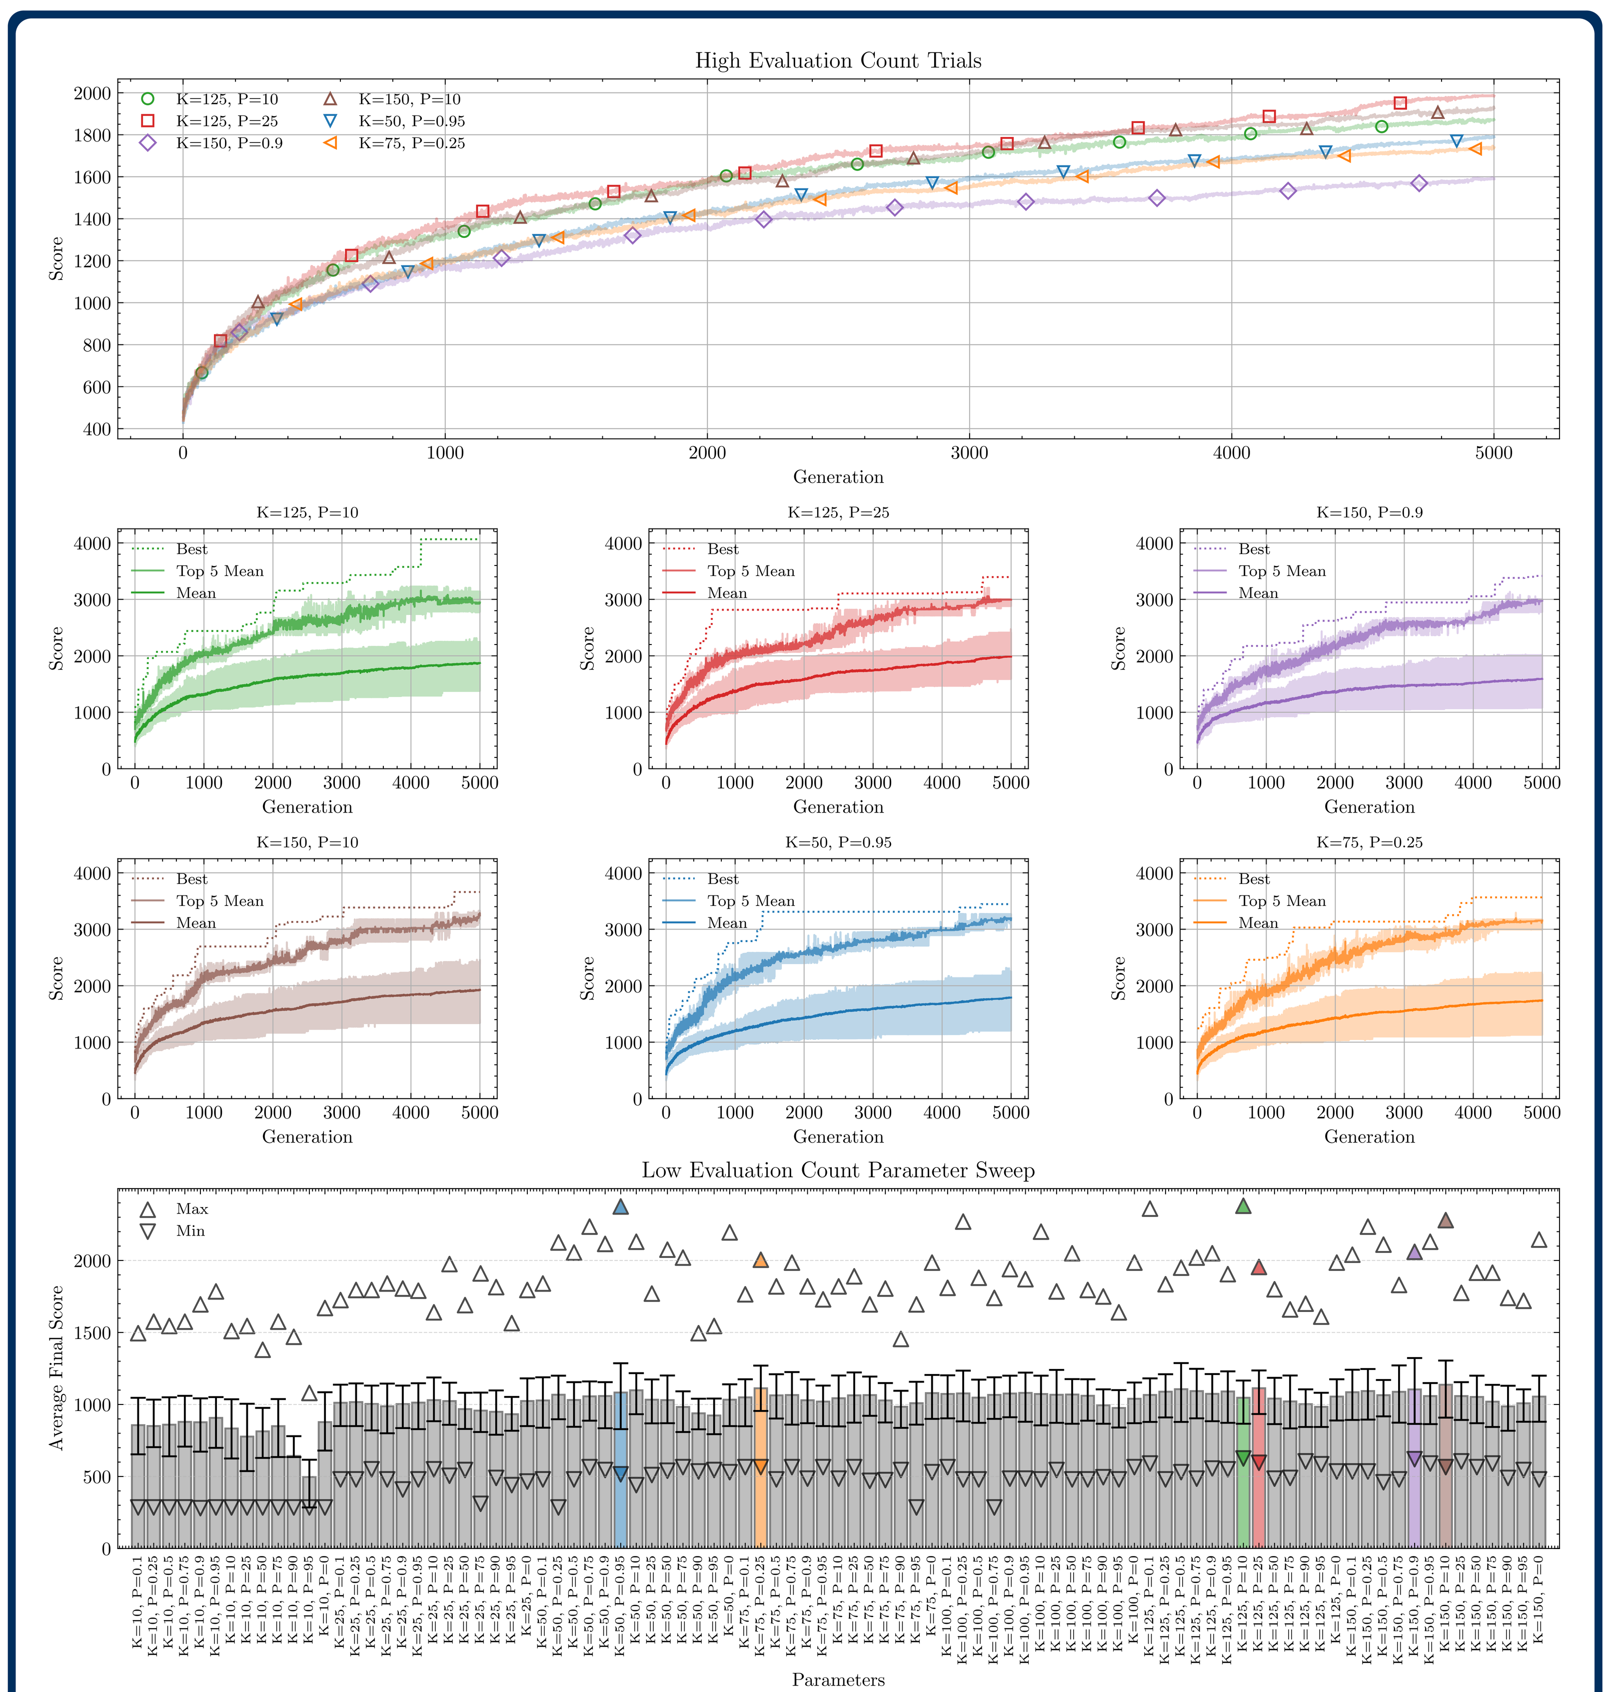

Figure 2. Full accounting of parameter sweep and extended trials. Top: six best (K, P) combinations averaged over 50 runs.

Future Work

- Extend SCOPE to continuous control and 3D domains.
- Evaluate across the full 50-game ALE suite.
- Explore integration with evolutionary model-based RL pipelines.