Playing Atari Space Invaders with Sparse Cosine Optimized Policy Evolution

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

Connecticut College

November 14, 2025

Motivations

- ➔ Large observations in robotics and game-playing demand large policies in order to classify observations to actions
- ➔ Evolutionary methods decrease in efficacy as dimensionality increases
- ➔ SCOPE aims to provide a method to compress inputs while retaining essential information for effective learning in large input spaces

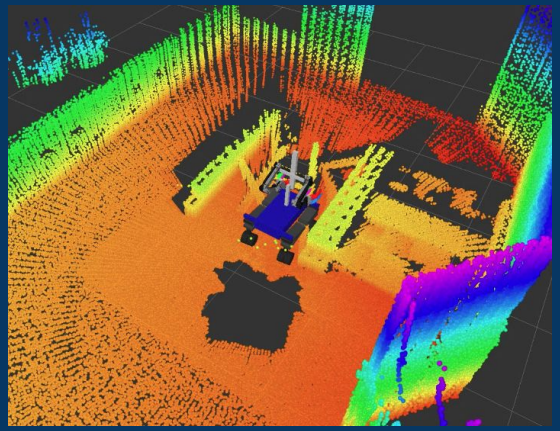

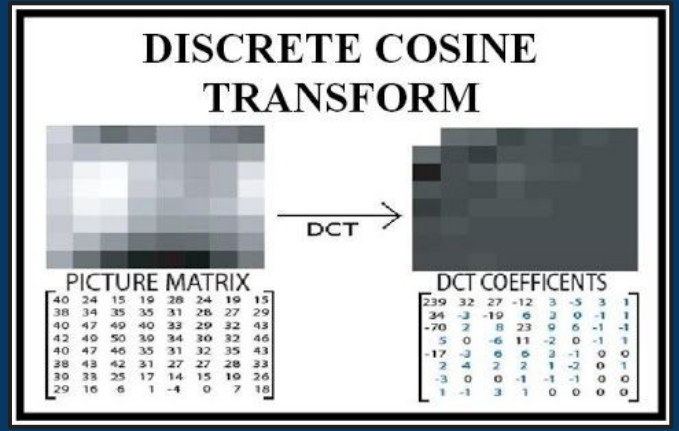

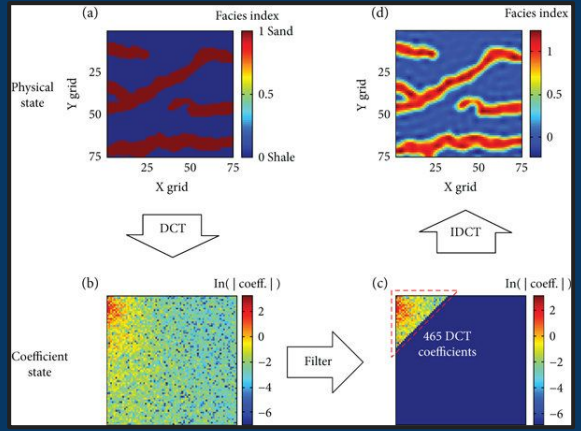

Sparse Cosine Optimized Policy Evolution (SCOPE)

- ➔ Uses the Discrete Cosine Transform (DCT) to convert high-dimensional inputs into the frequency domain
- ➔ Truncates and sparsifies the DCT coefficients, reducing input dimensionality
- ➔ Enables evolved policies to handle large input spaces with far fewer parameters compared to conventional solutions

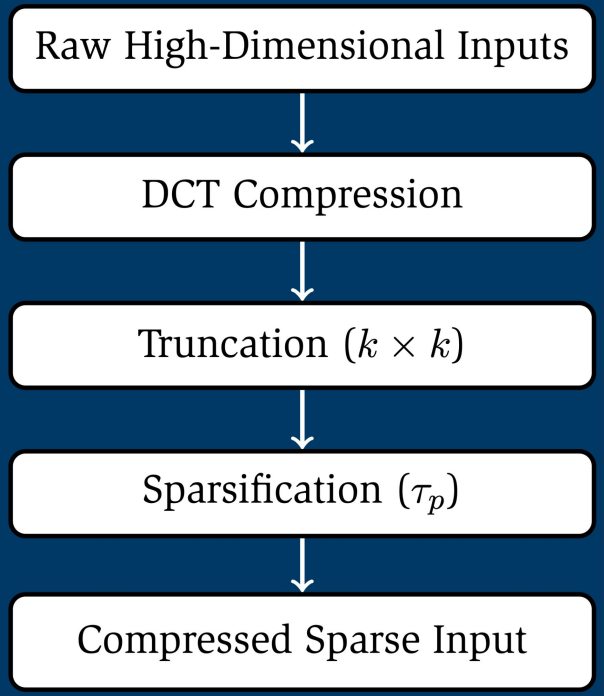

Effects of Truncating DCT Coefficients

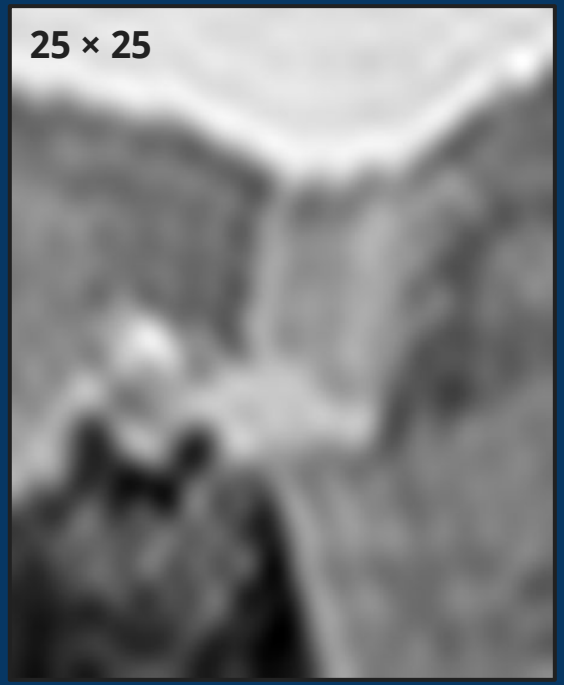

Effects of Enforcing Sparsity

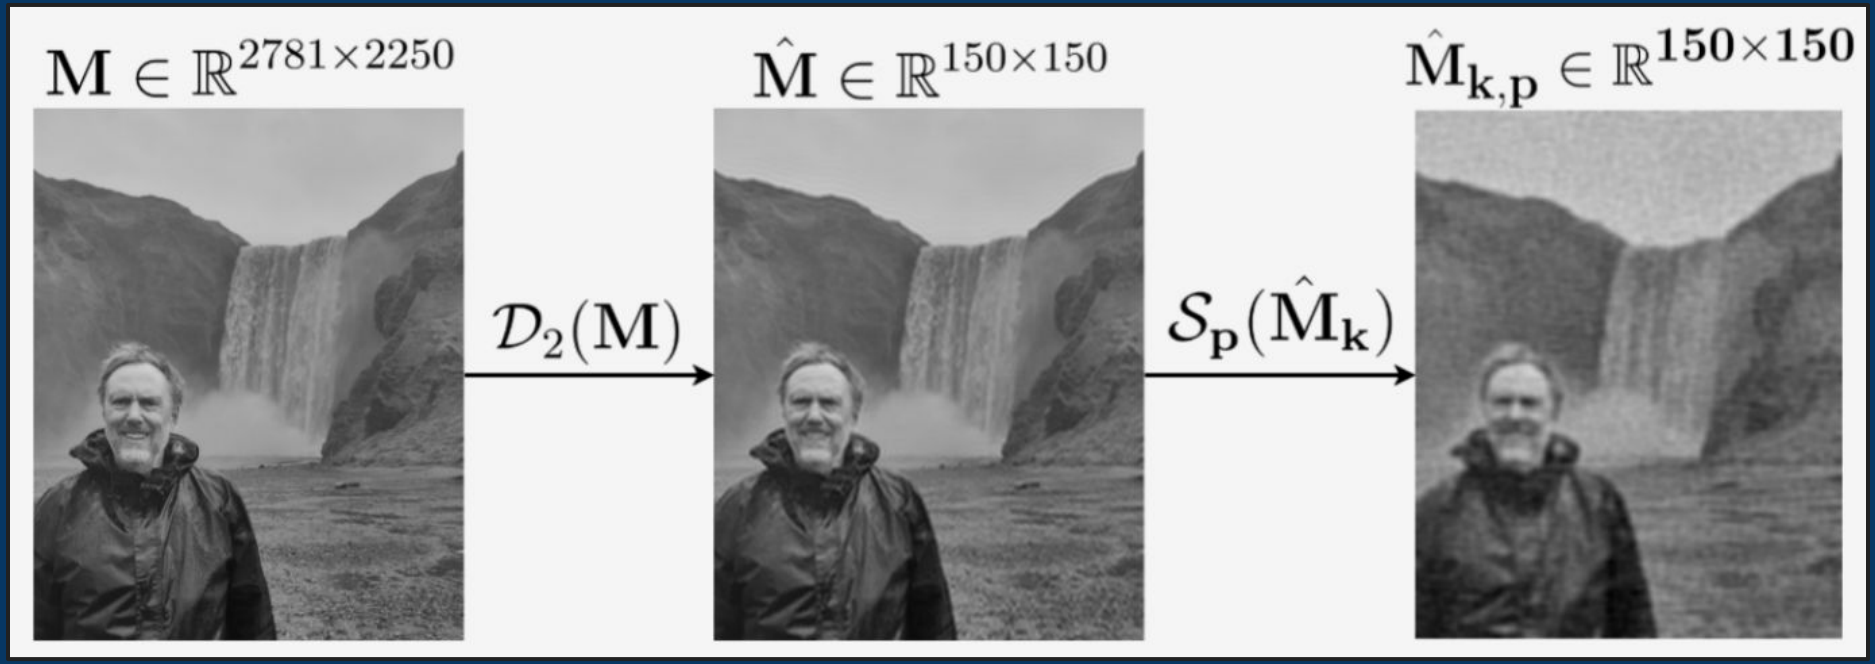

Evolution Pipeline

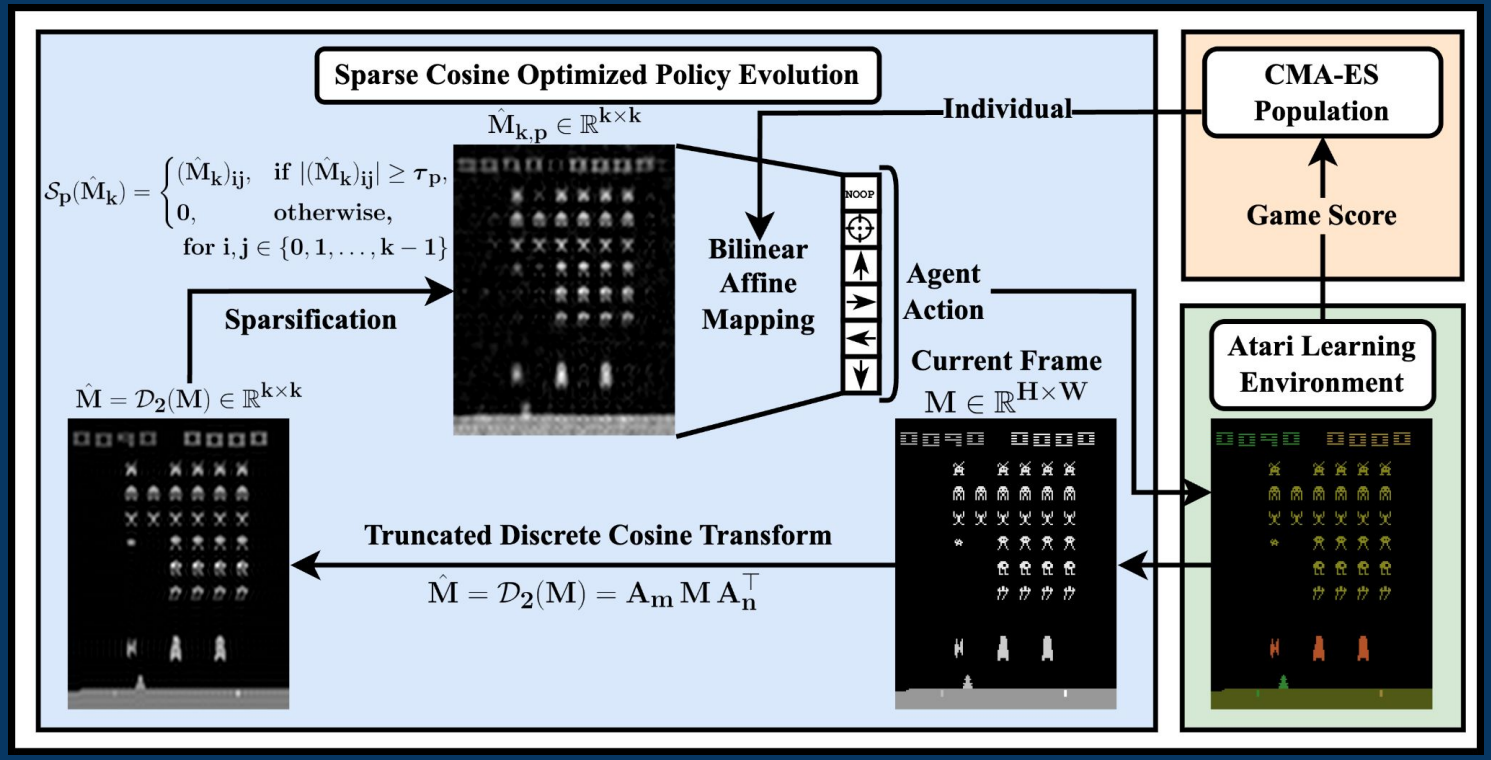

Atari Learning Environment & Space Invaders

- ➔ We use the Atari Learning Environment (ALE) as a standardized benchmark for evaluating game-playing agents
- ➔ Space Invaders provides patterns and localized high-frequency inputs
- ➔ Input frames are 210 x 160 grayscale images, compressed using SCOPE

Putting it all together…

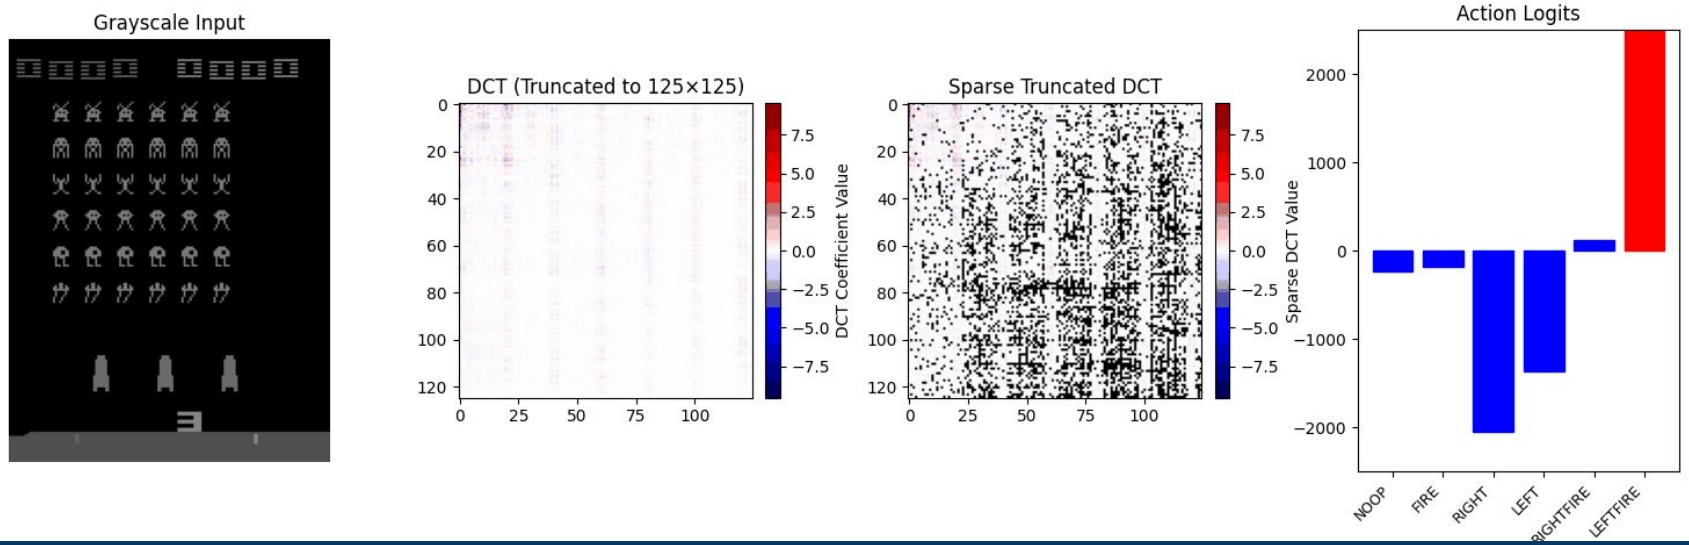

Results

- ➔ SCOPE reduced input size by 53%: from 33,600 to 15,625
- ➔ Best configuration achieved scores surpassing OpenAI-ES, HyperNEAT, and DQN baselines

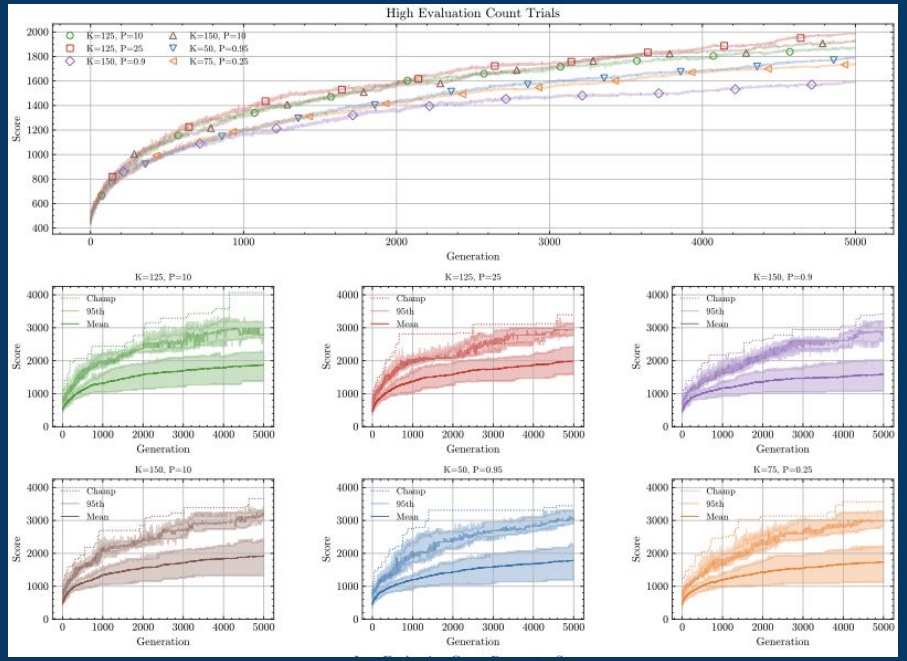

Method	K=125,P=10	K=125,P=25	K=150,P=10	OpenAI-ES	HyperNEAT	A2C FF	A3C FF (1 Day)	DQN	UCT
Mean	1,993.5	**2,109.1**	2,068.3	678.5	1,251	951.9	–	1,449.7	–
Top-5 Mean	3,261	3,175	**3,347**	–	–	–	2,214.7	–	–
Best	**4,065**	3,395	3,660	–	1,481	–	–	–	2,718
Parameter Count	**875**	**875**	1,050	680,448	116,602	680,448	680,448	1,692,672	N/A

Thank You!

Questions?