SCOPE for Hexapod Gait Generation

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

Connecticut College

IJCCI Conference on Evolutionary Computation and Theory and Applications

Mathematical Foundations of SCOPE

The Discrete Cosine Transform (DCT)

- The DCT expresses a sequence of real numbers as a sum of cosine functions oscillating at different frequencies
- Unlike the Fourier Transform, the DCT uses only real-valued basis functions
- Widely used in signal compression for its energy compaction properties such as JPEG compression

1D Discrete Cosine Transform

Given input vector $\mathbf{x} = (x_0, x_1, \dots, x_{N-1})$:

$$X_k = \alpha_k \sum_{n=0}^{N-1} x_n \cos\left[\frac{\pi}{N}\left(n + \tfrac{1}{2}\right) k\right],$$

where the normalization constant is

$$\alpha_k = \begin{cases} \sqrt{\frac{1}{N}}, & k = 0, \\ \sqrt{\frac{2}{N}}, & k > 0. \end{cases}$$

- Produces real-valued frequency coefficients X_k
- Low-frequency coefficients (small k) capture coarse structure; high-frequency coefficients capture fine details or noise

2D Discrete Cosine Transform

- Extends the 1D DCT by applying it along rows and columns of a matrix
- Widely used for image compression

For input matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$:

$$\hat{\mathbf{M}} = \mathcal{D}_2(\mathbf{M}) = \mathbf{A}_m \mathbf{M} \mathbf{A}_n^\top,$$

where $\mathbf{A}_m$ and $\mathbf{A}_n$ are the DCT basis matrices.

Energy Compaction Property of DCT

- The DCT concentrates most of the signal's energy into a few low-frequency coefficients
- Truncating high-frequency coefficients reduces dimensionality with minimal information loss

2D DCT Example

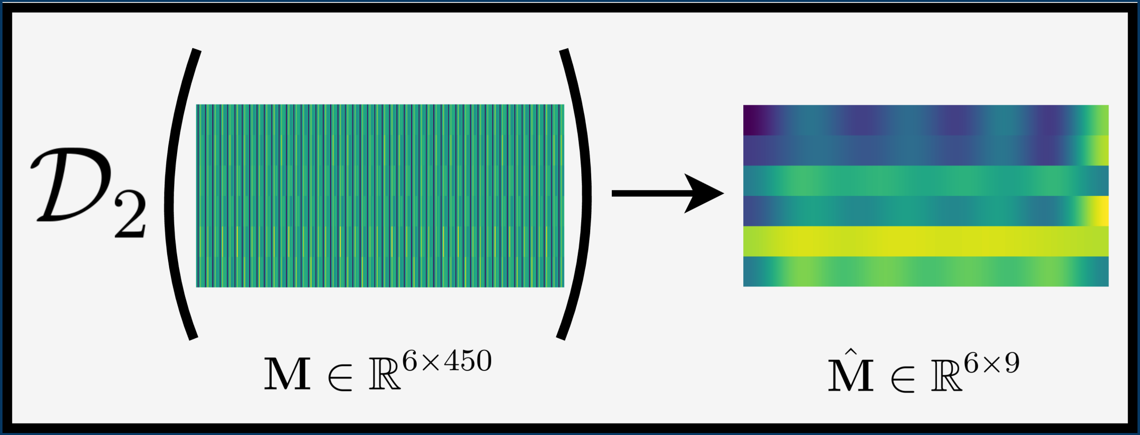

Why Use DCT in SCOPE?

- Enables compression of high-dimensional inputs by retaining only high-energy (low-frequency) components
- Reduces the number of parameters needed for policy optimization
- Acts as a pseudo-attention mechanism, focusing on relevant patterns while discarding noise

Sparsification of the DCT Output

- After truncating the $k \times k$ DCT matrix $\hat{\mathbf{M}}_k$, we apply a sparsification operator $\mathcal{S}_p$:

$$\mathcal{S}_p(\hat{\mathbf{M}}_k)_{ij} = \begin{cases} \hat{\mathbf{M}}_{ij}, & \text{if } |\hat{\mathbf{M}}_{ij}| \geq \tau_p, \\ 0, & \text{otherwise}, \end{cases}$$

  where τ_p is the p-th percentile threshold of absolute coefficient magnitudes.
- This keeps only high-energy coefficients while zeroing out low-energy components
- The resulting sparse matrix $\hat{\mathbf{M}}_{k,p}$ is passed to the policy's bilinear affine map

Sparsification Example

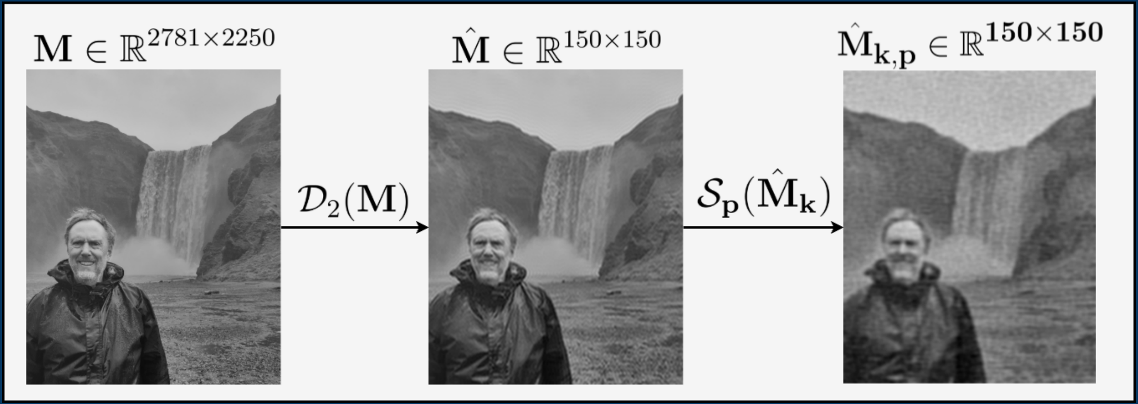

Why Sparsification Helps SCOPE

- High-energy DCT coefficients capture low-frequency structures; low-energy coefficients often represent noise or irrelevant detail
- Sparsification reduces the number of active input features, effectively shrinking the policy's parameter search space
- By dynamically adjusting which coefficients are retained, the policy can adapt its "focus" without explicit attention
- This allows SCOPE to evolve efficient policies with fewer parameters, faster convergence, and improved robustness to noise

SCOPE for Hexapod Gait Generation

Hexapod Robot Design

- We use a simulated 6-legged hexapod robot in the Webots simulator
- Each leg has 3 joints: coxa, femur, tibia
- The robot's movement is controlled via motors in each joint
- Designed and built by Matt Denton

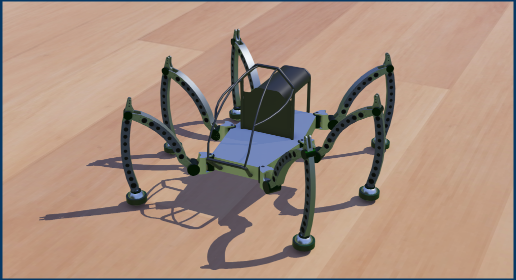

Evolution Pipeline

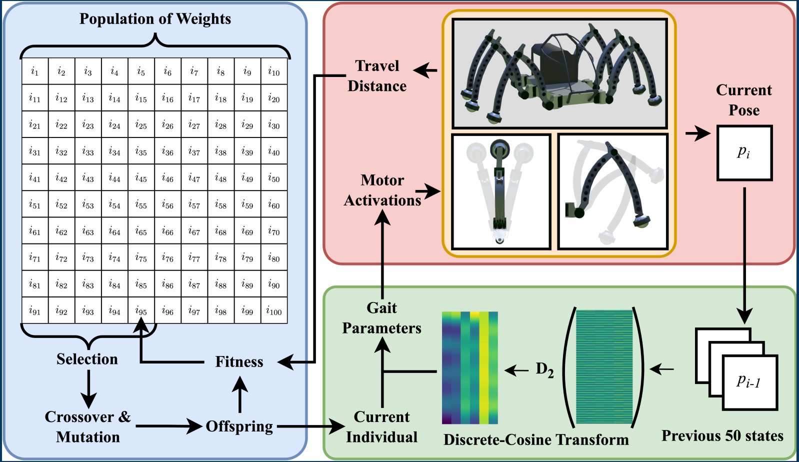

Results

- SCOPE provides a 98% reduction in input size: from 2700 to 54 features
- SCOPE reached average fitness of 14.24, compared to 11.88 for baseline across 500 individual trials, representing a 20% improvement in walking distance on average

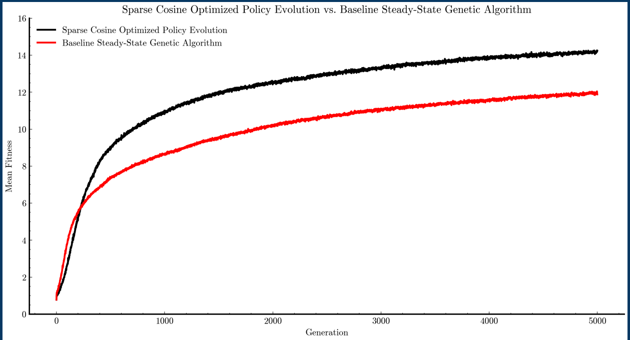

Thank You!

Questions?