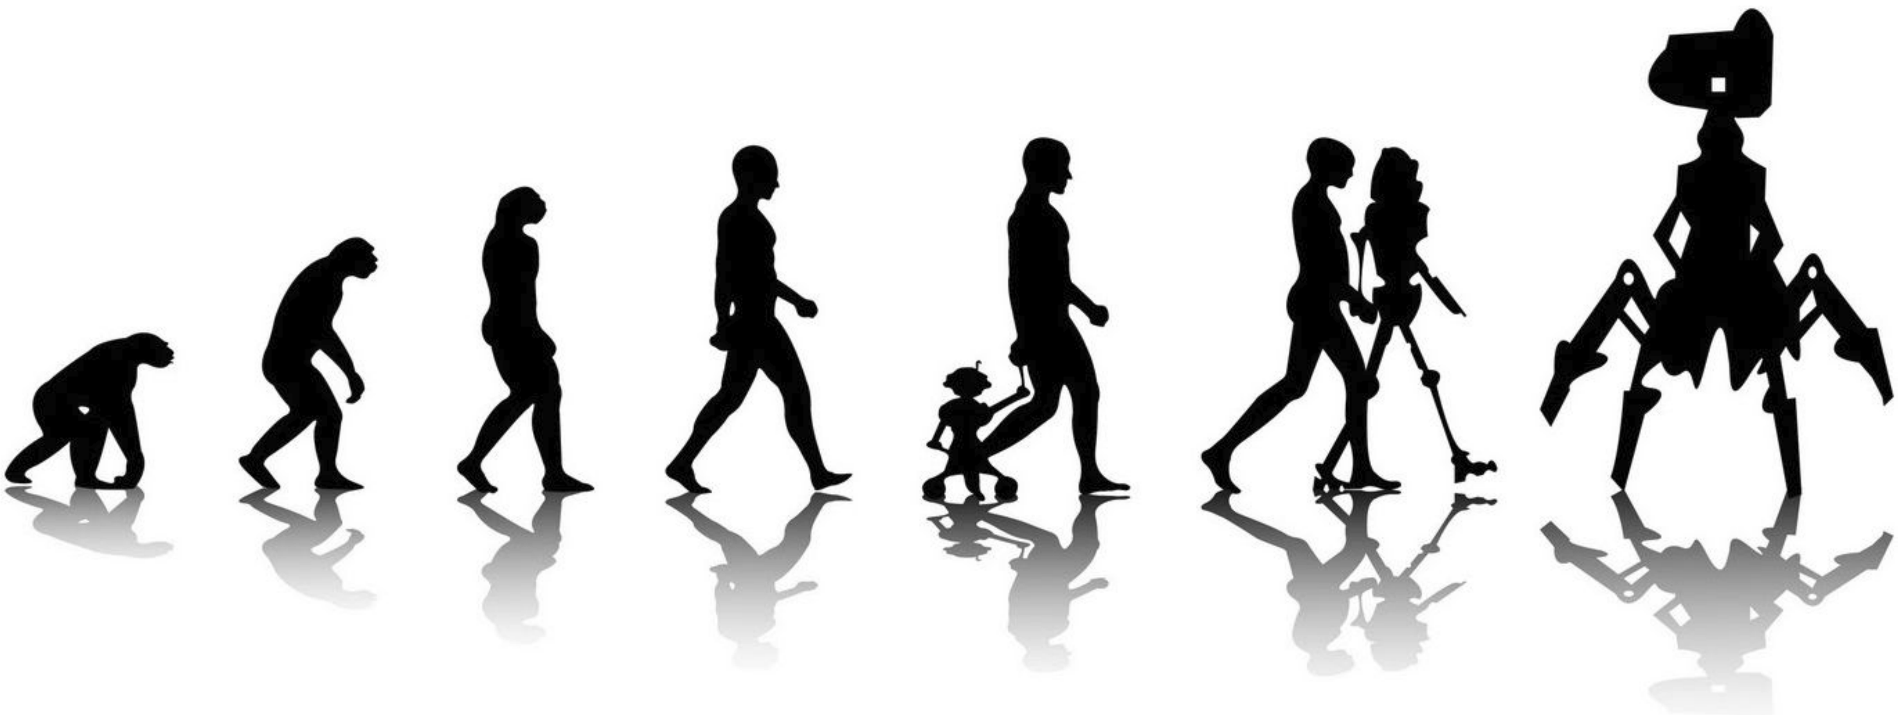

Modern Robotics: Evolutionary Robotics

COSC 4560 / COSC 5560

Professor Cheney

3/5/18

No class March 26th and 28th
(and still no class on the 30th)

Project proposals due April 1st

Behavioral Characteristics

The last paper used the endpoint
of the robot’s behavior
to calculate the distance
of it’s behavior from other robots’

What else could we use?

Trajectory of robot’s path
(it’s location throughout the simulation)

The robot’s heading through the simulation

The robot’s sensor/motor readings over time

For embodied robots?

The footprint plot of the robot
(when it’s touching the ground)

Volume

Height, Width, Depth

Convex Hull / Branching / Shape Entropy

What if we care more about
how many different ways
we can perform a task well

instead of just how well we can perform
it with our best solution

(optimizing for both diversity/novelty and fitness)

Confronting the Challenge of Quality Diversity

Justin K. Pugh, L. B. Soros, Paul A. Szerlip, and Kenneth O. Stanley

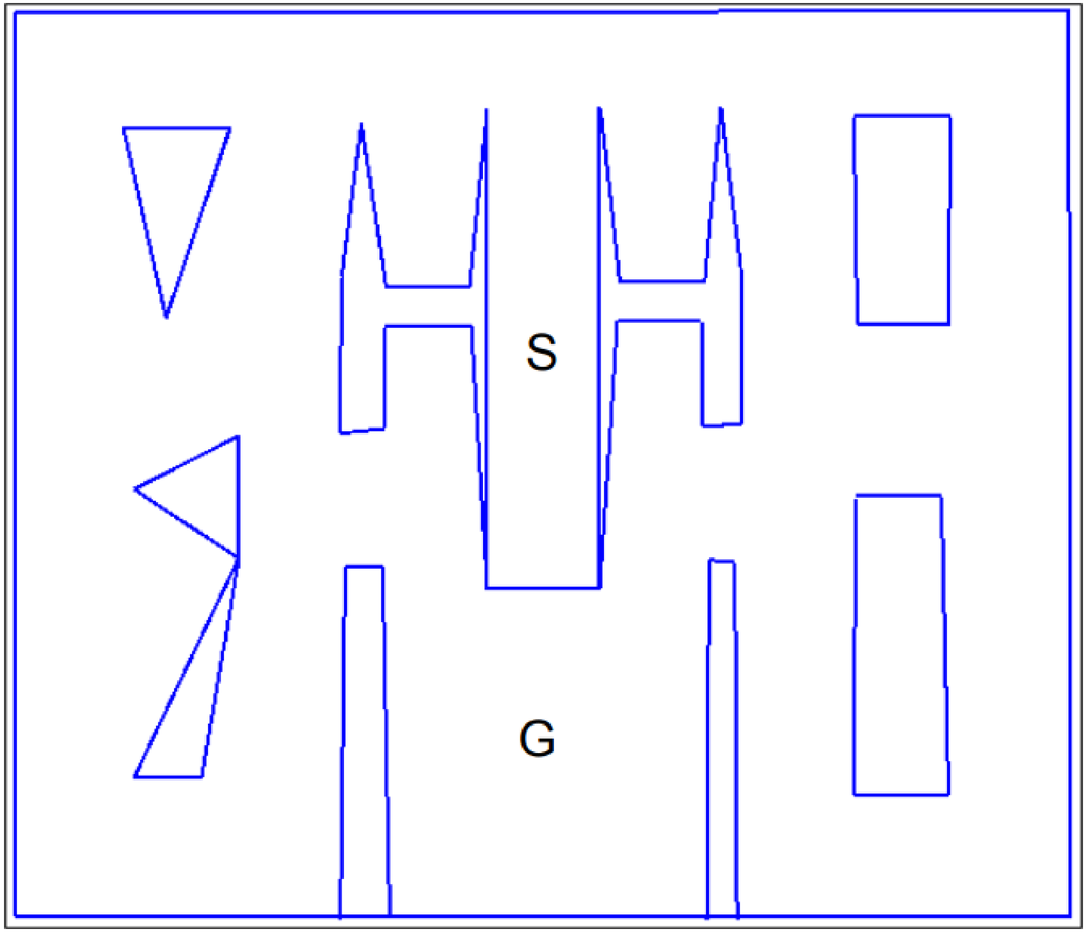

Figure 1: Quality diversity maze (QD-Maze). Individuals start at the point S and the highest quality solutions navigate to point G. While the maze is deceptive enough to challenge objective-based algorithms, it also contains a variety of solutions and thus serves as a benchmark for combining quality with diversity.

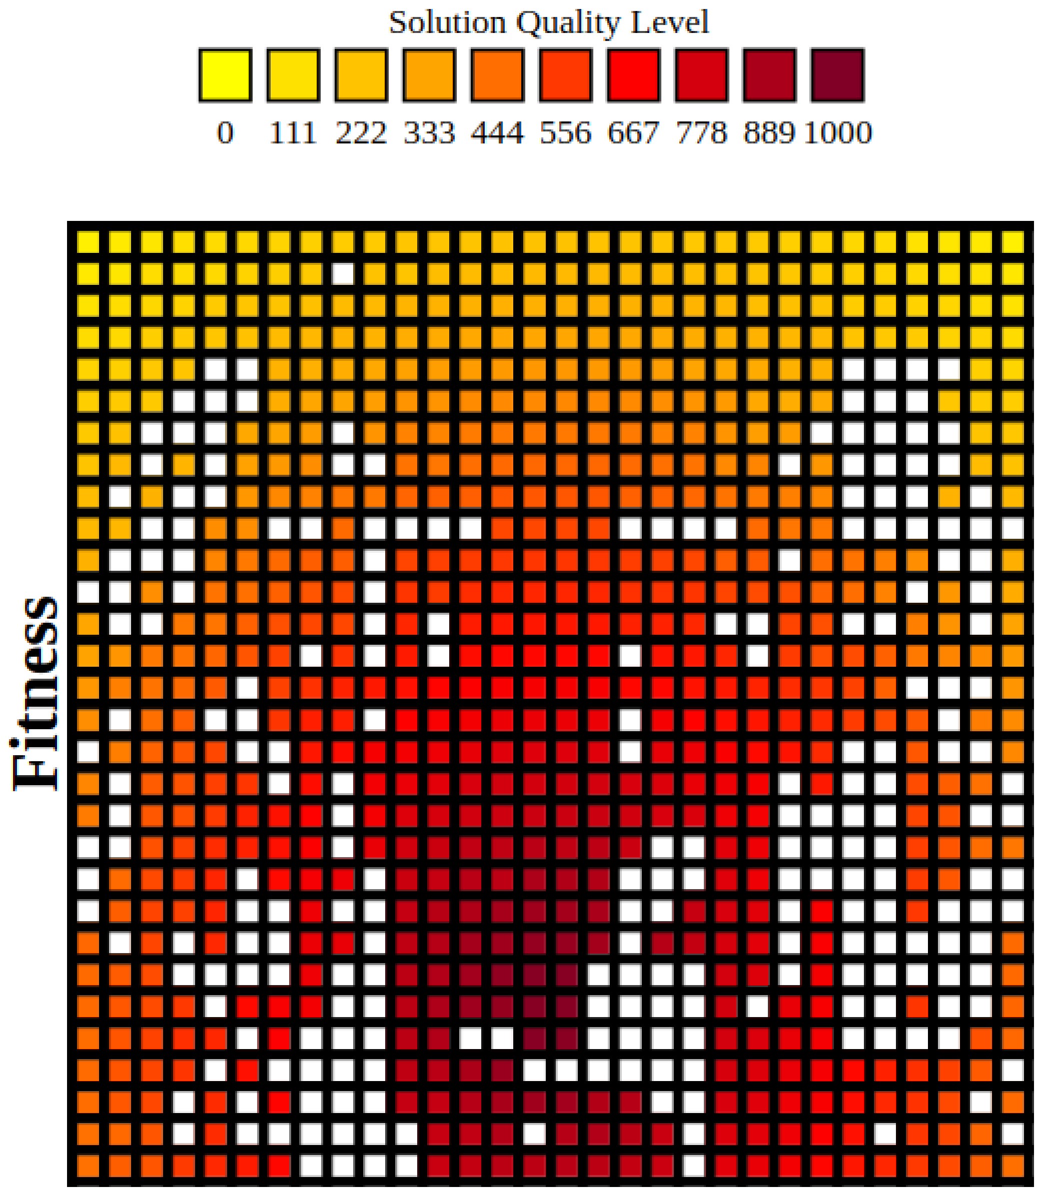

Solution Quality Level
0 111 222 333 444 556 667 778 889 1000
Fitness

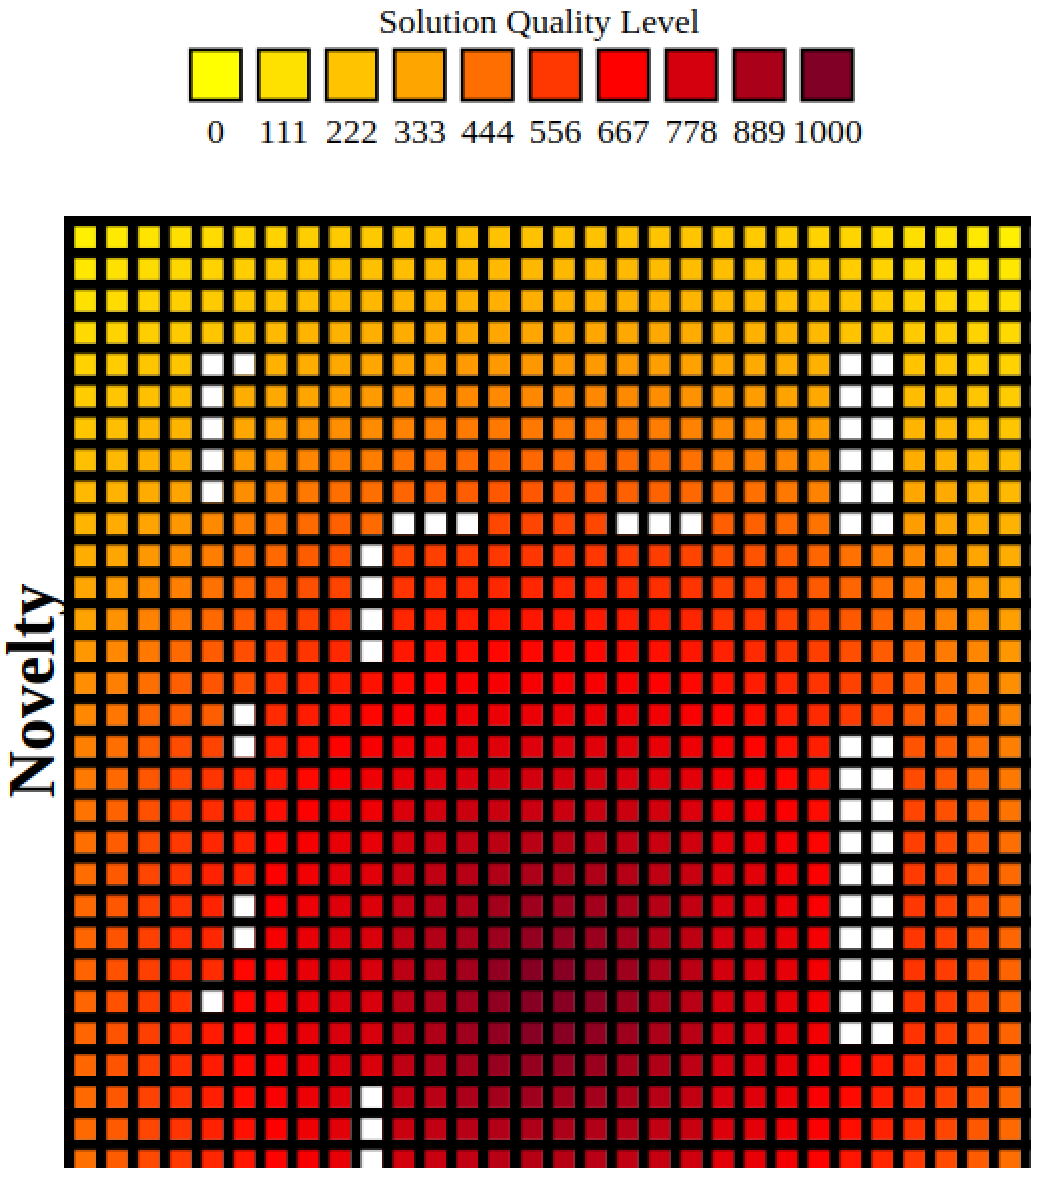
Solution Quality Level
0
111
222
333
444
556
667
778
889
1000
Novelty

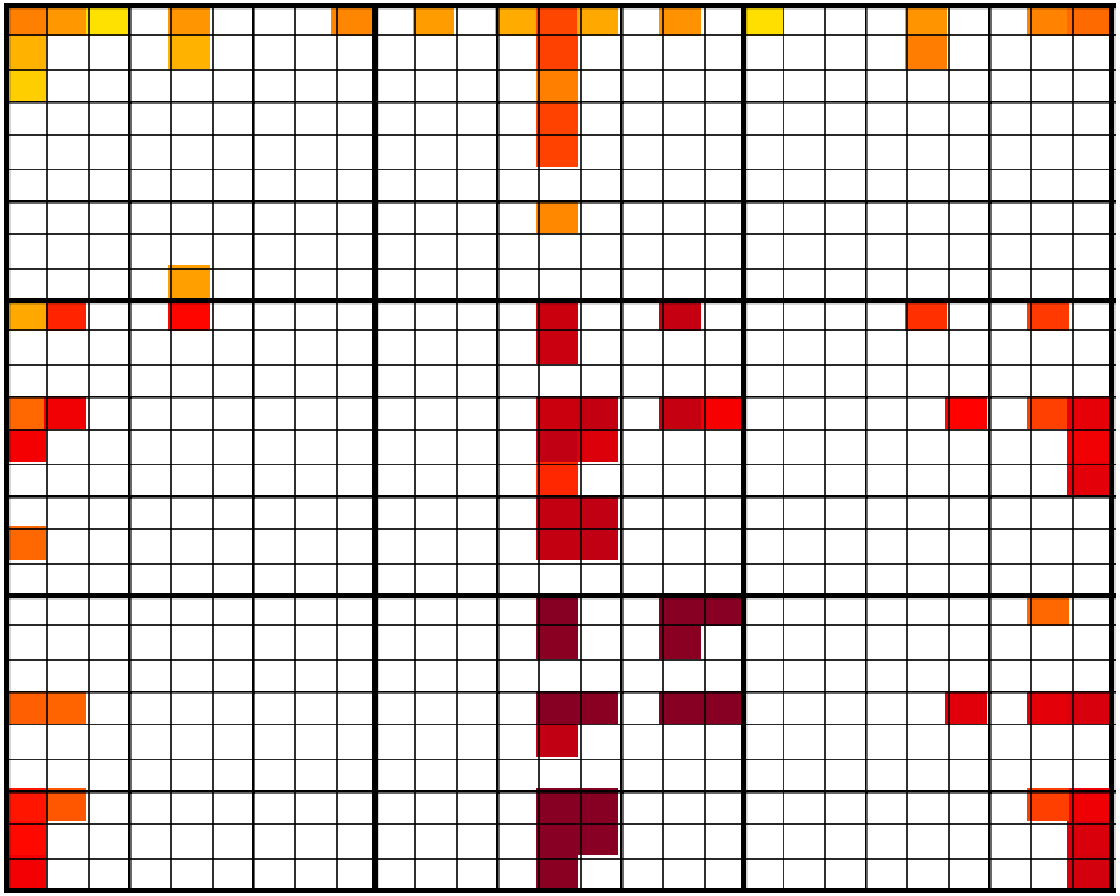

Figure 6: Fitness in the FullTrajectoryBC. In this example grid, the six-dimensional behavior space (FullTrajectoryBC) (discretized into three bins per dimension for a total of 729 bins) is visually depicted as a series of nested two-dimensional grids (each of which are 3×3). The color of each grid box corresponds to the quality of the solution found by the search algorithm after 250,000 evaluations: yellow corresponds to low quality, dark red to high quality, and white to unfilled bins. Fitness finds very few of the possible behaviors for this BC.

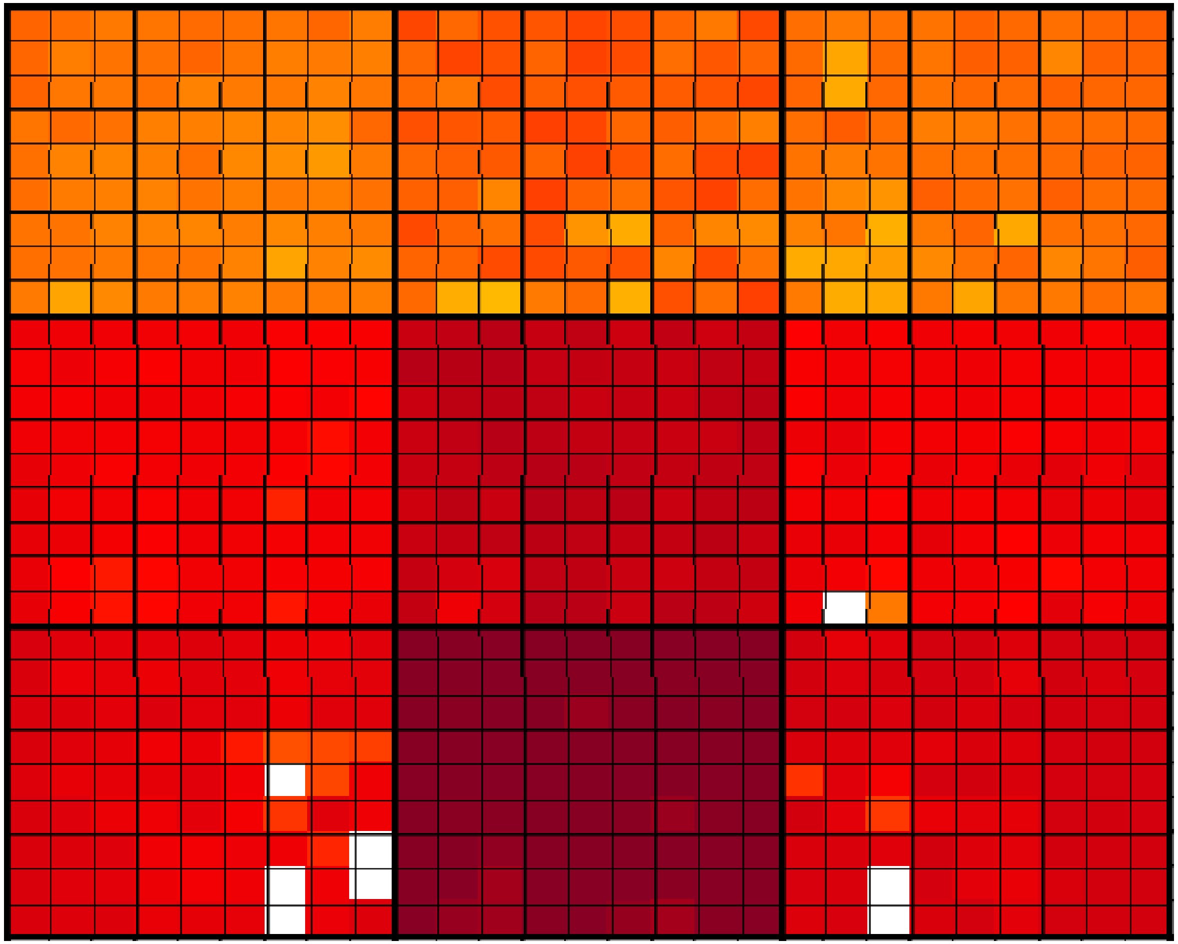

Figure 8: Novelty Search in the FullTrajectoryBC. Of the five compared algorithms, NS performs the best under the FullTrajectoryBC (featuring relatively high alignment between the BC and the objective) because it focuses exclusively on pursuing diversity. This conclusion is supported in the QD collection grid by almost all bins being filled (i.e. non-white).

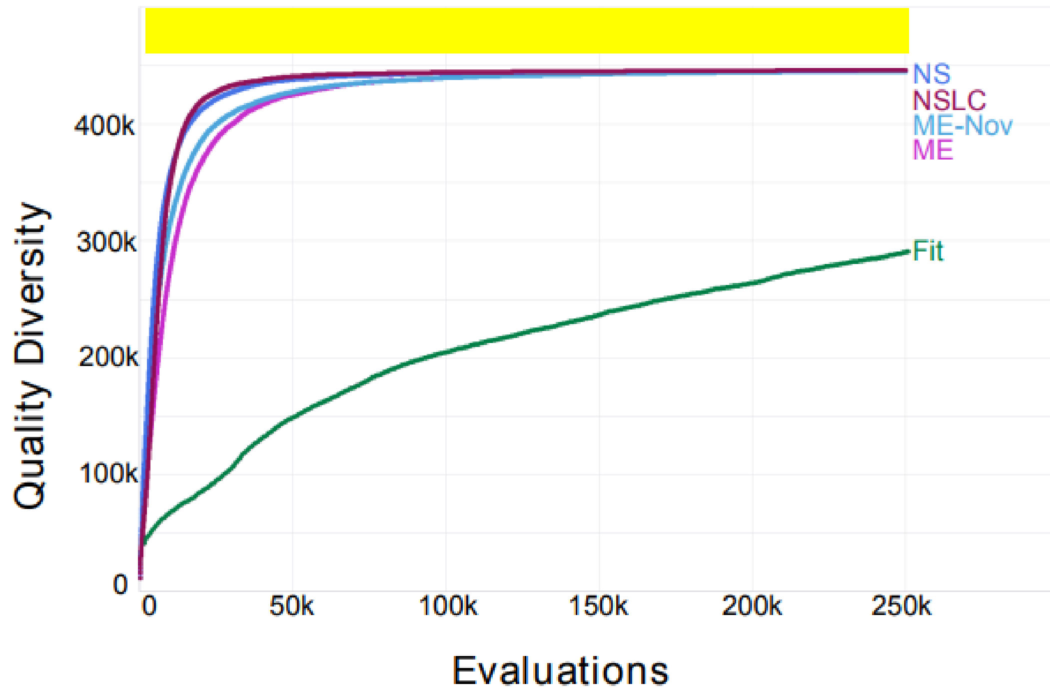

Figure 2: EndpointBC (very high alignment). In this performance comparison and others in this paper, the average QD (taken over 20 runs) for each variant method is shown across a run. The yellow strip at the top indicates the period during which there is a significant difference at least between the top and bottom method (exlcuding fitness, which is always significantly worse than all other methods). The method labels are color coded to match with their respective curves and are shown from top to bottom in the their rank order during the period of significance. Note that because of the large QD scale, sometimes significant differences exist even when not visually apparent. For EndpointBC, NS performs best, followed by NSLC.

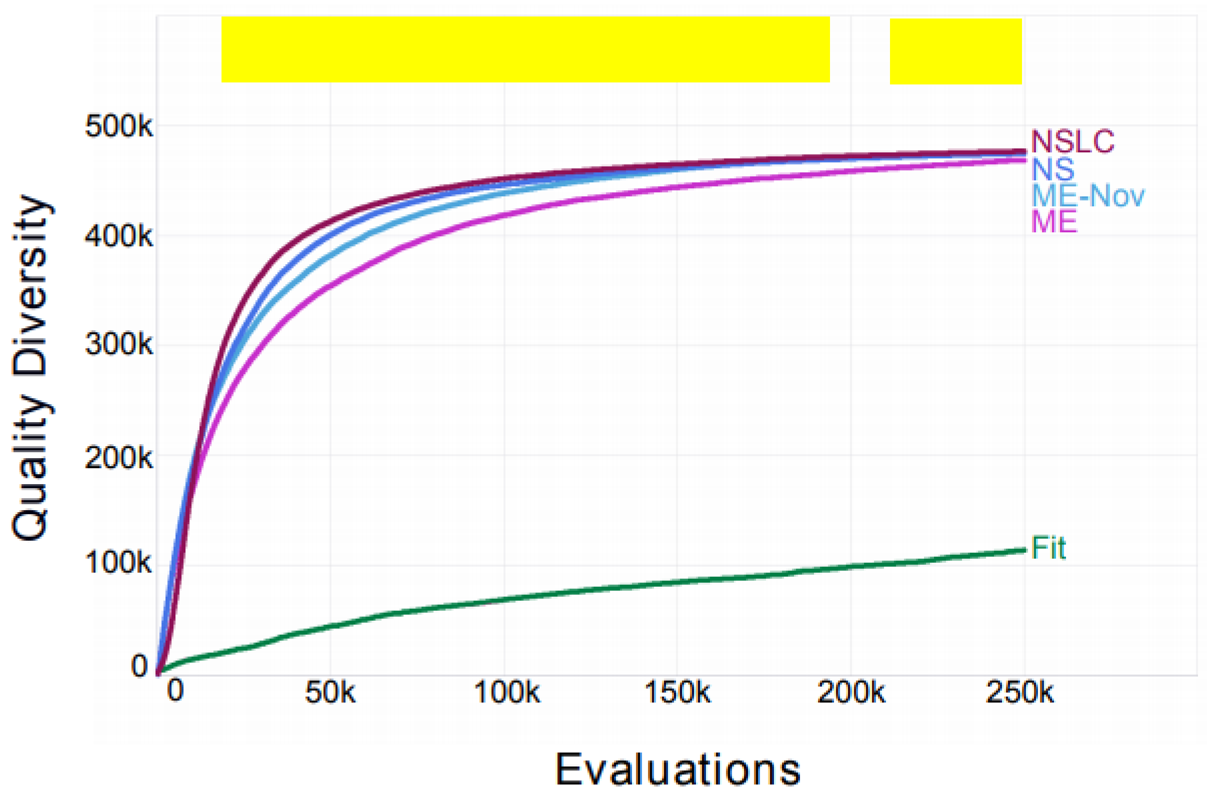

Figure 3: FullTrajectoryBC (high alignment). For this BC, which is slightly less aligned with quality than EndpointBC, NS, NSLC, and ME-Nov are effectively tied, with ME significantly behind.

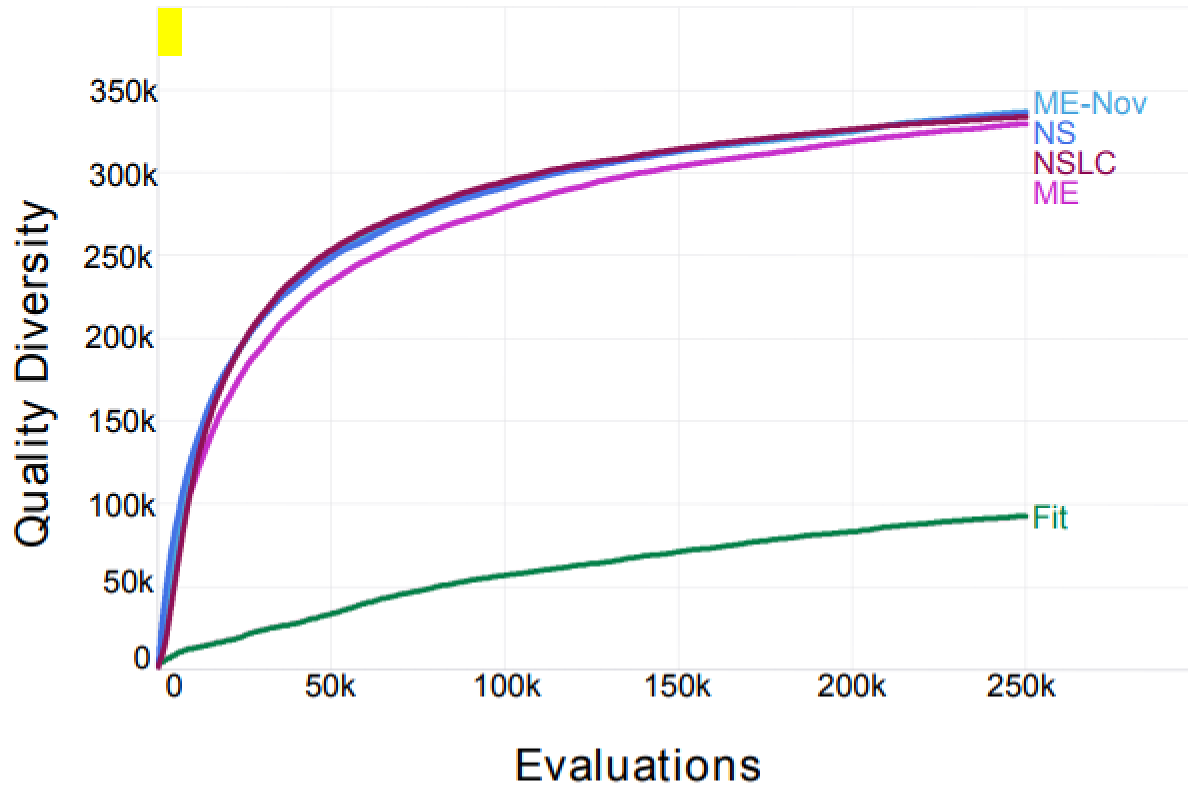

Figure 4: HalfTrajectoryBC (modest alignment). All the methods except fitness are tied for the vast majority of the run for this BC, which is even less aligned than FullTrajectoryBC.

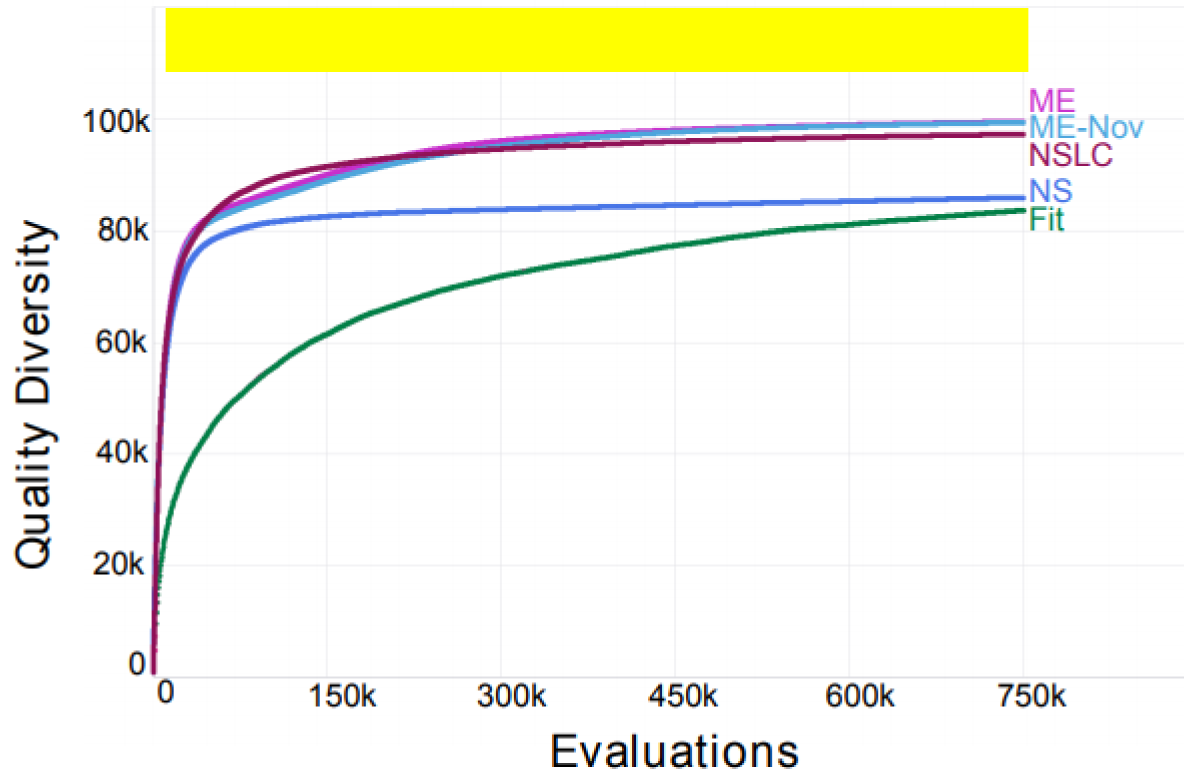

Figure 5: DirectionBC (low alignment). With an almost complete lack of alignment in this BC, ME and ME-Nov tie for first place, and NS trails far behind.

http://eplex.cs.ucf.edu/QD/GECCO-15/compare.html

Combining Novelty and Fitness

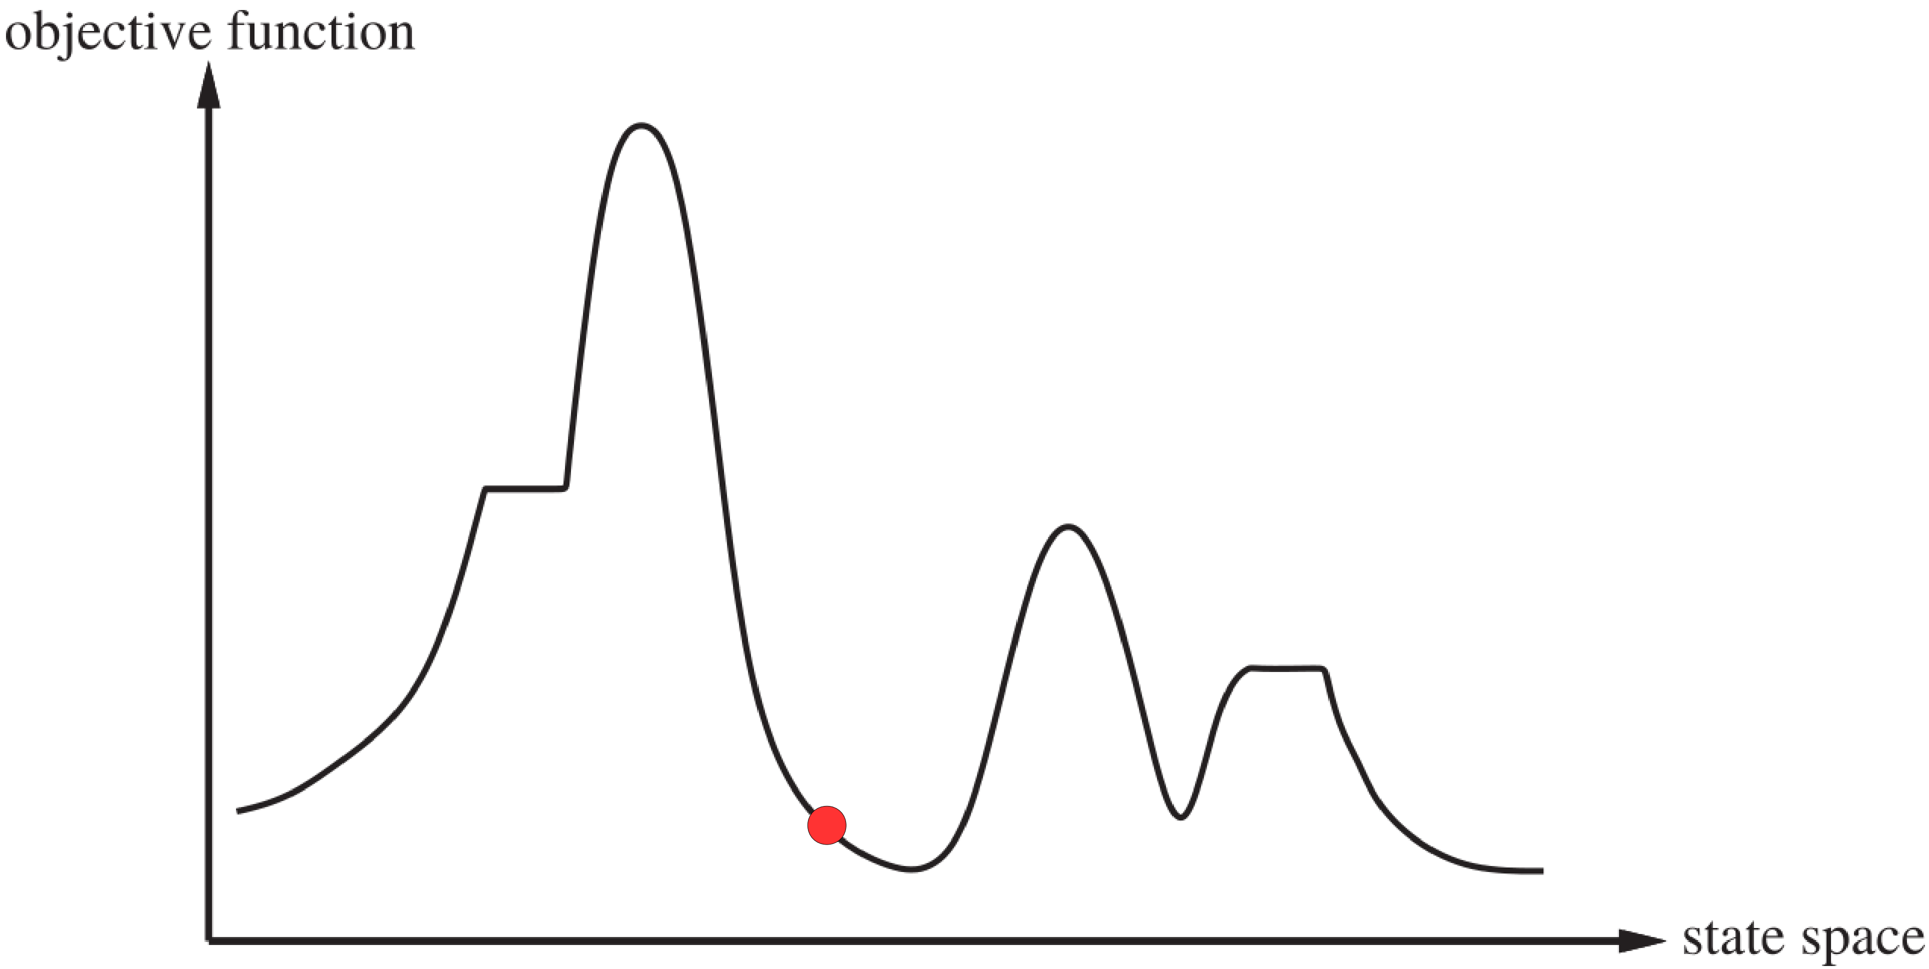

objective function
state space

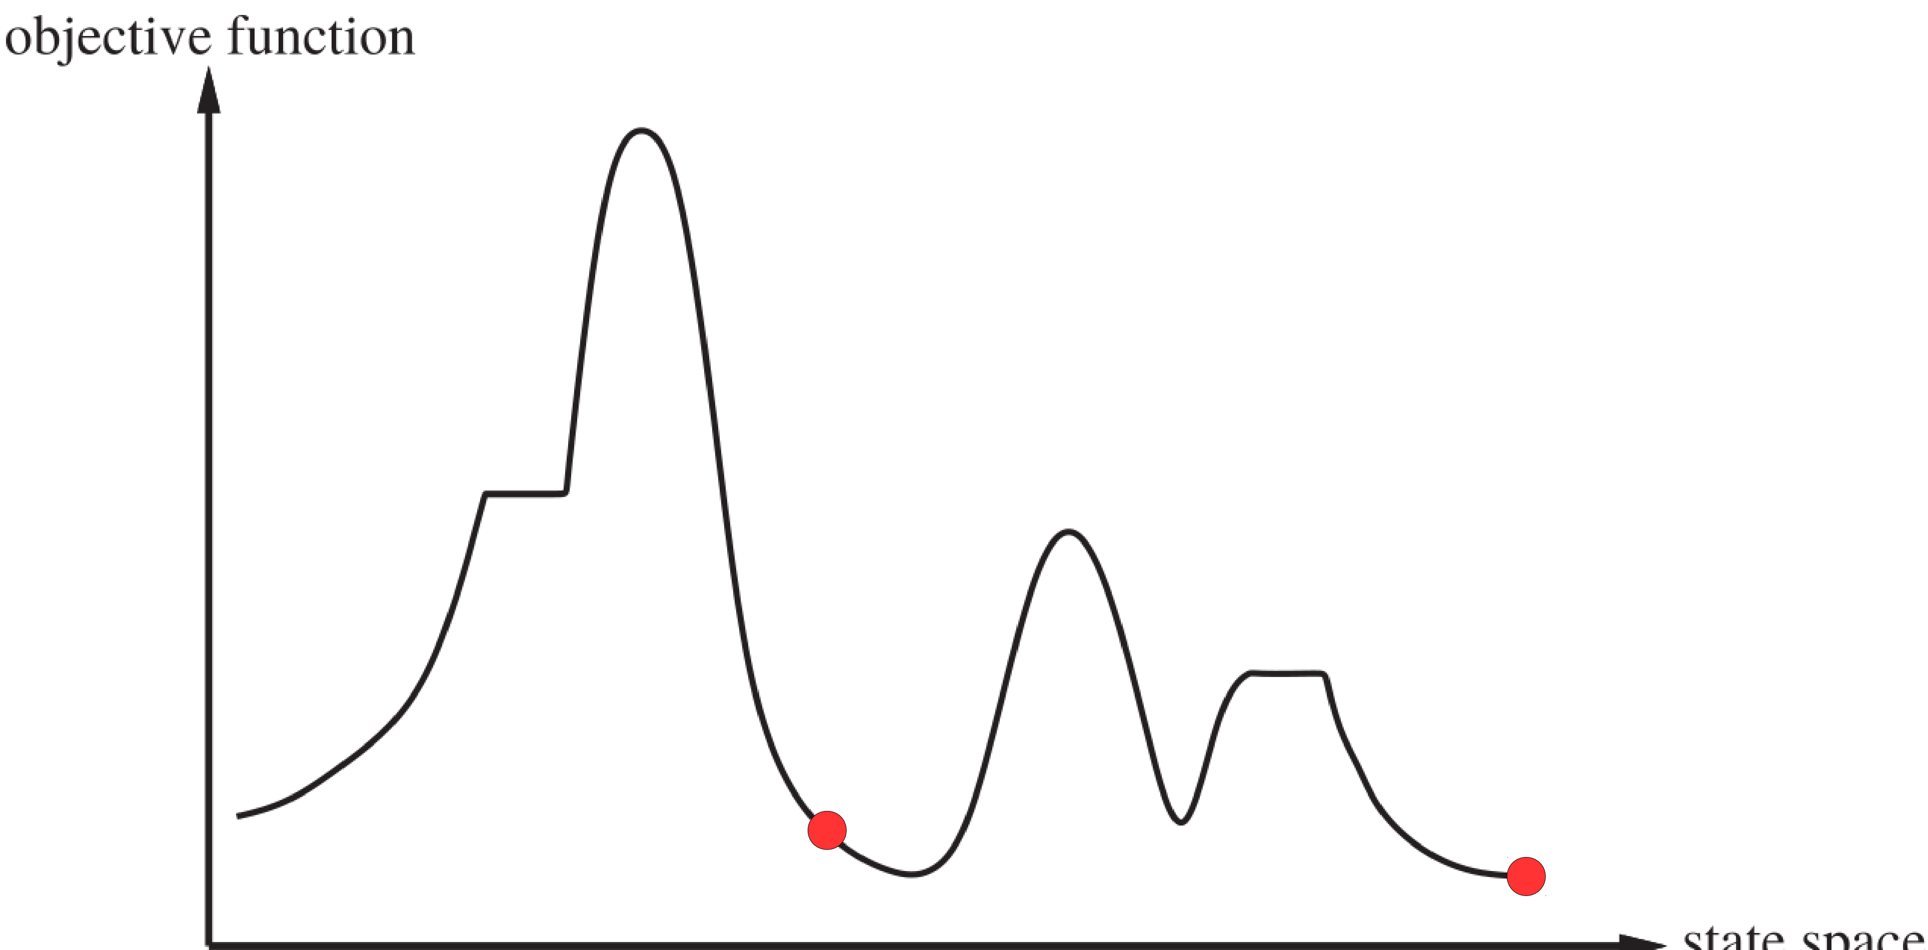
objective function
state space

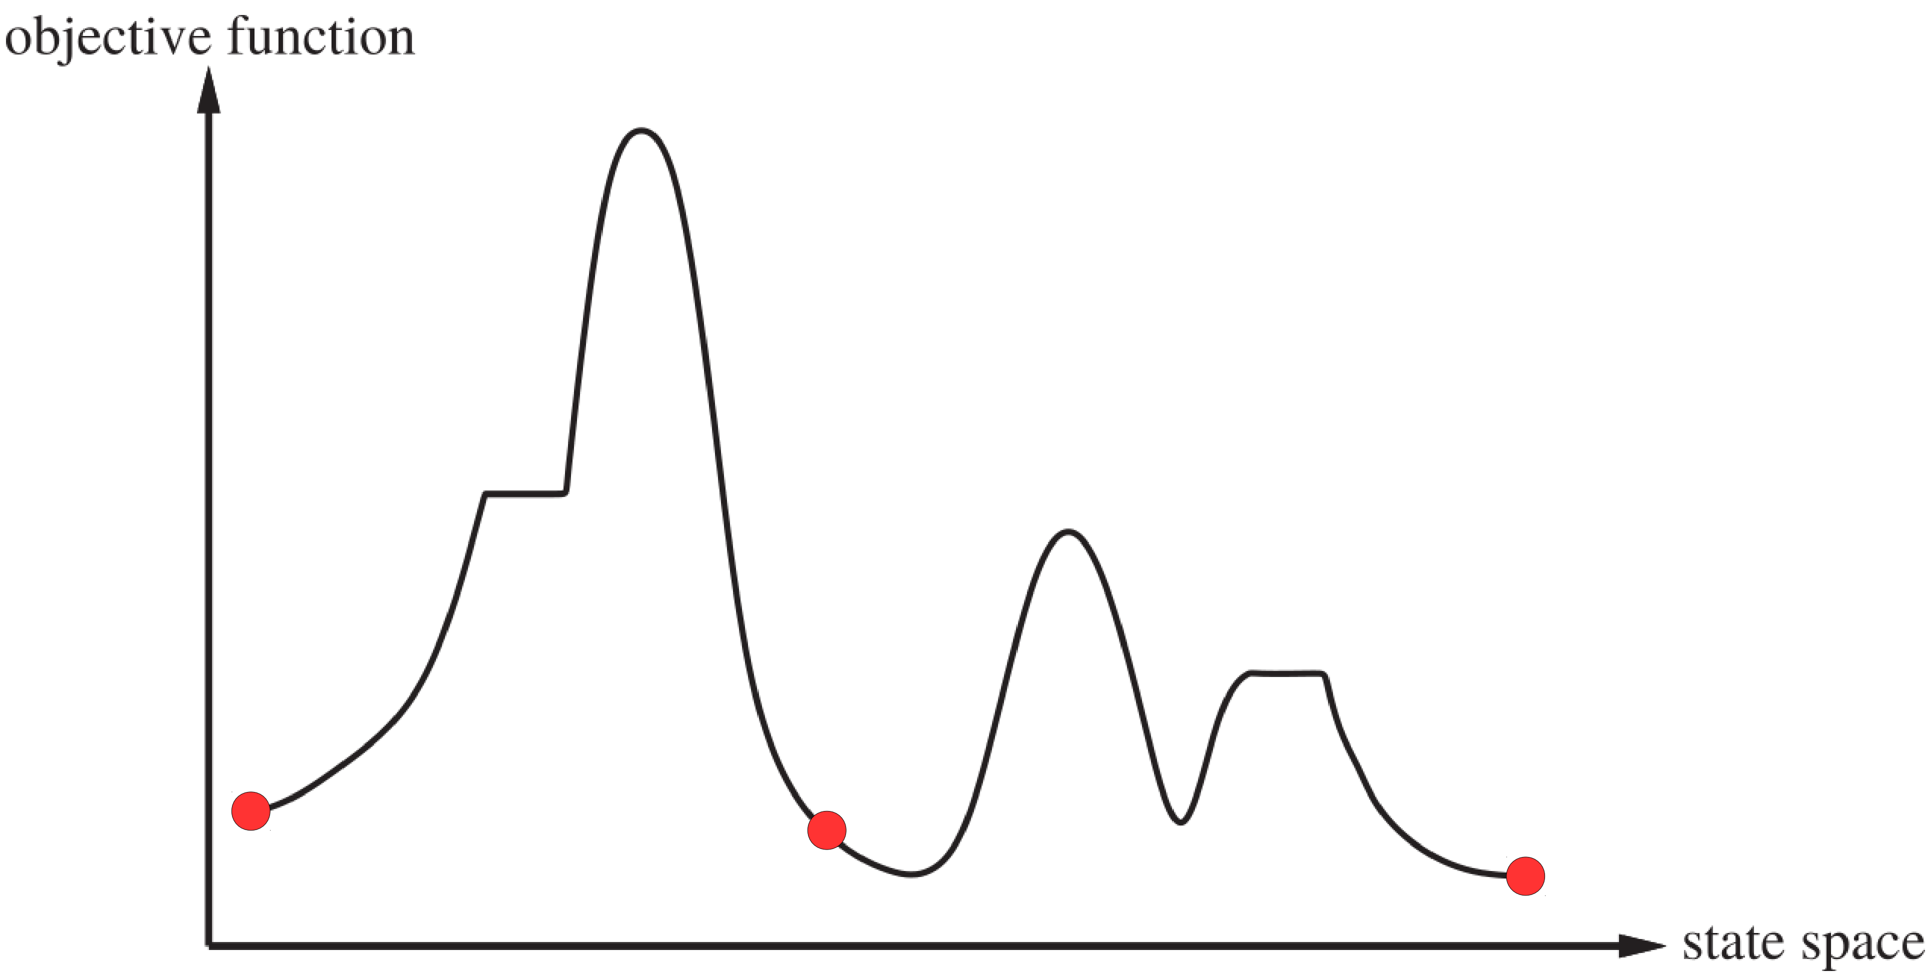
objective function
state space

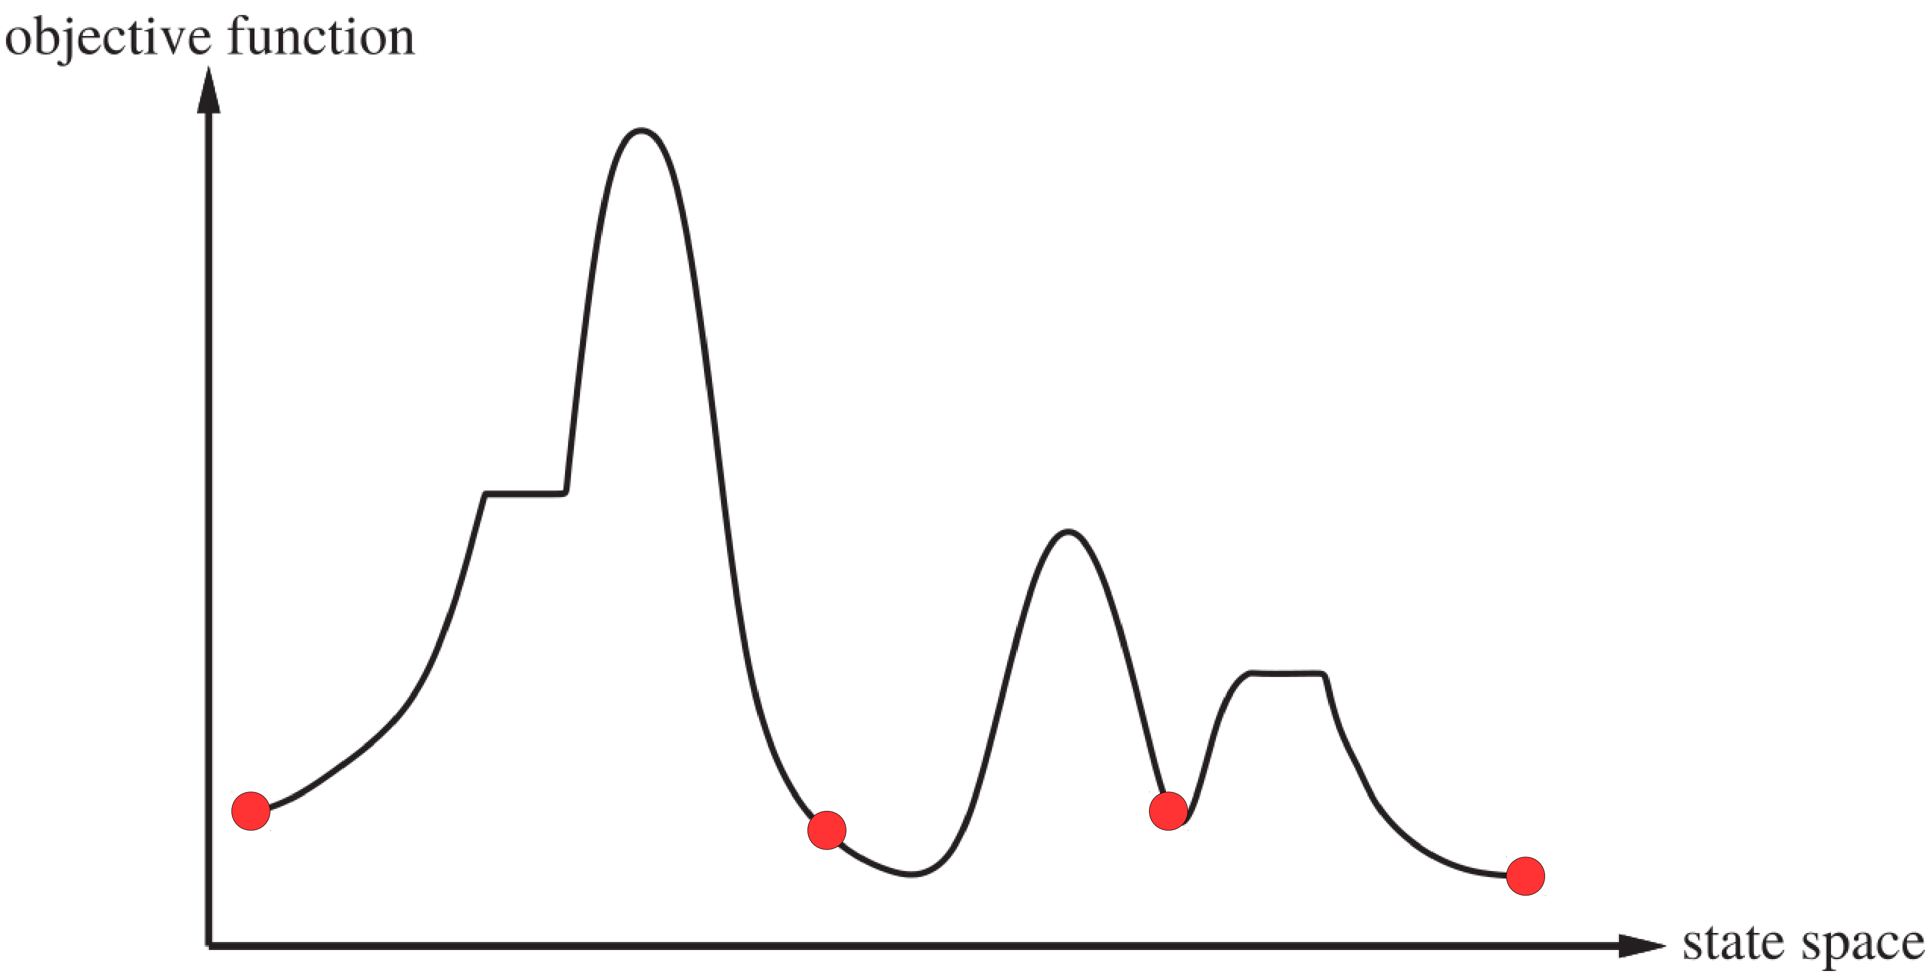
objective function
state space

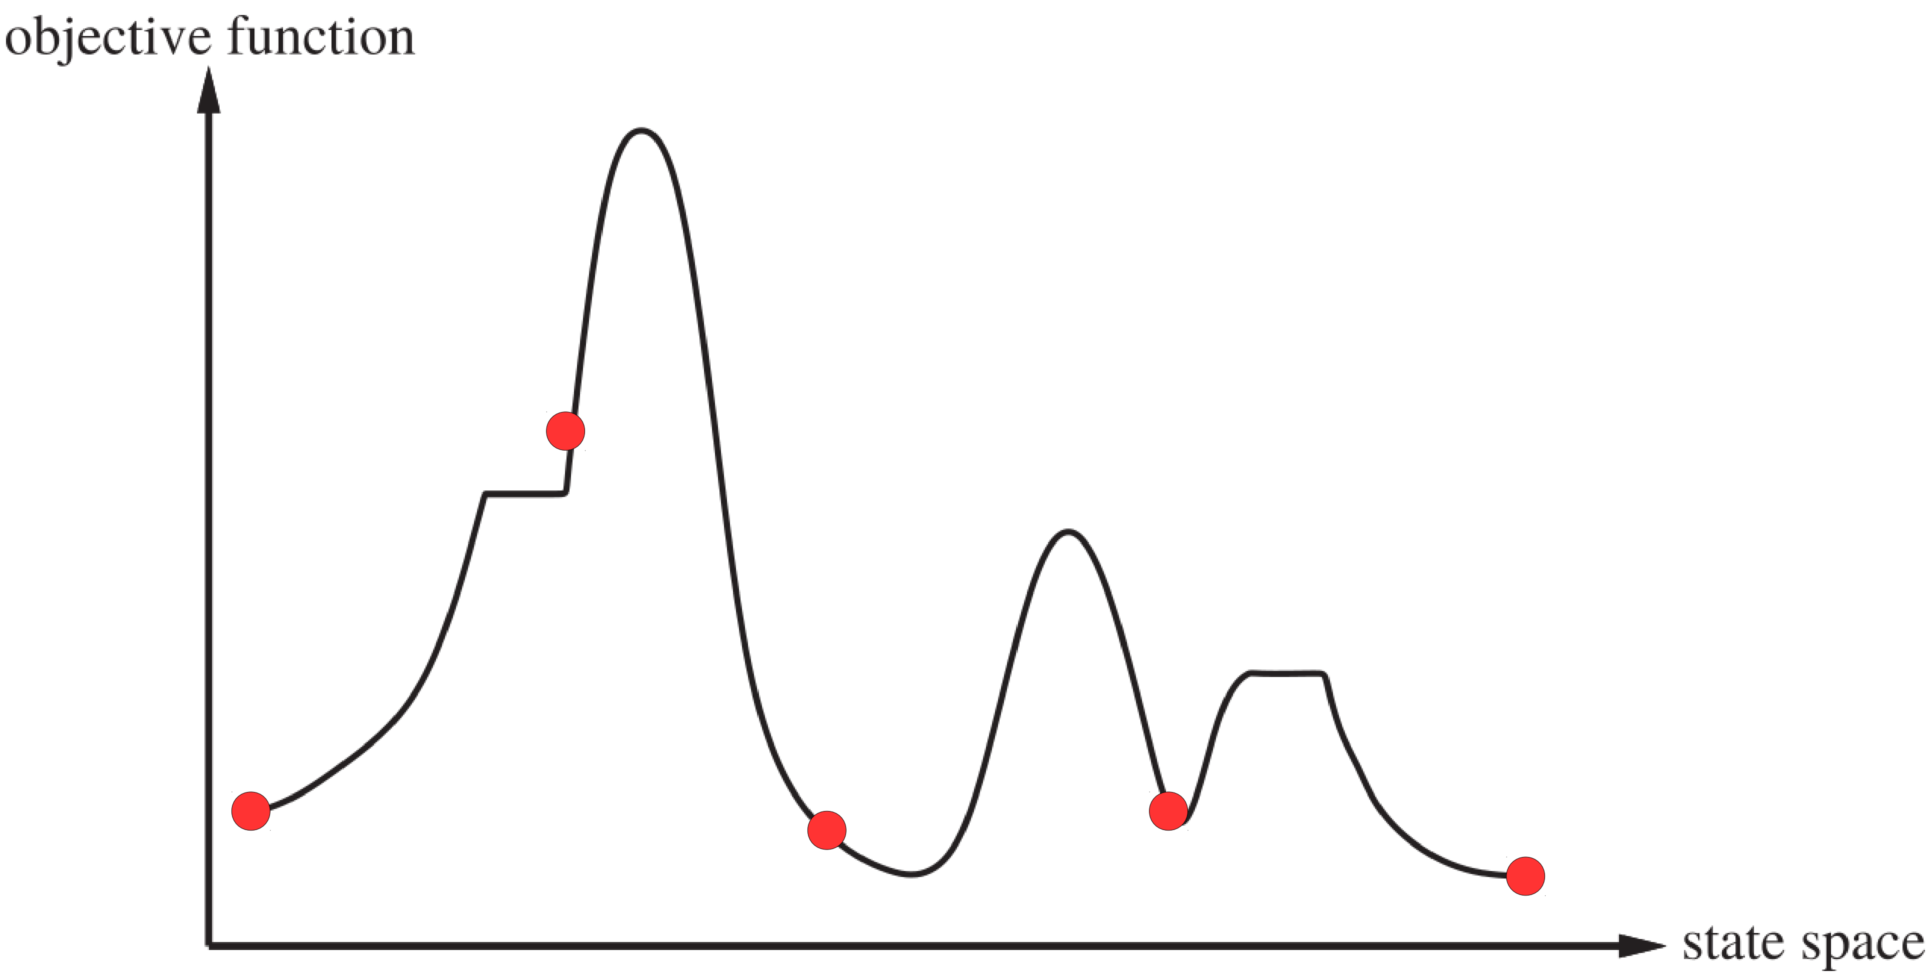
objective function
state space

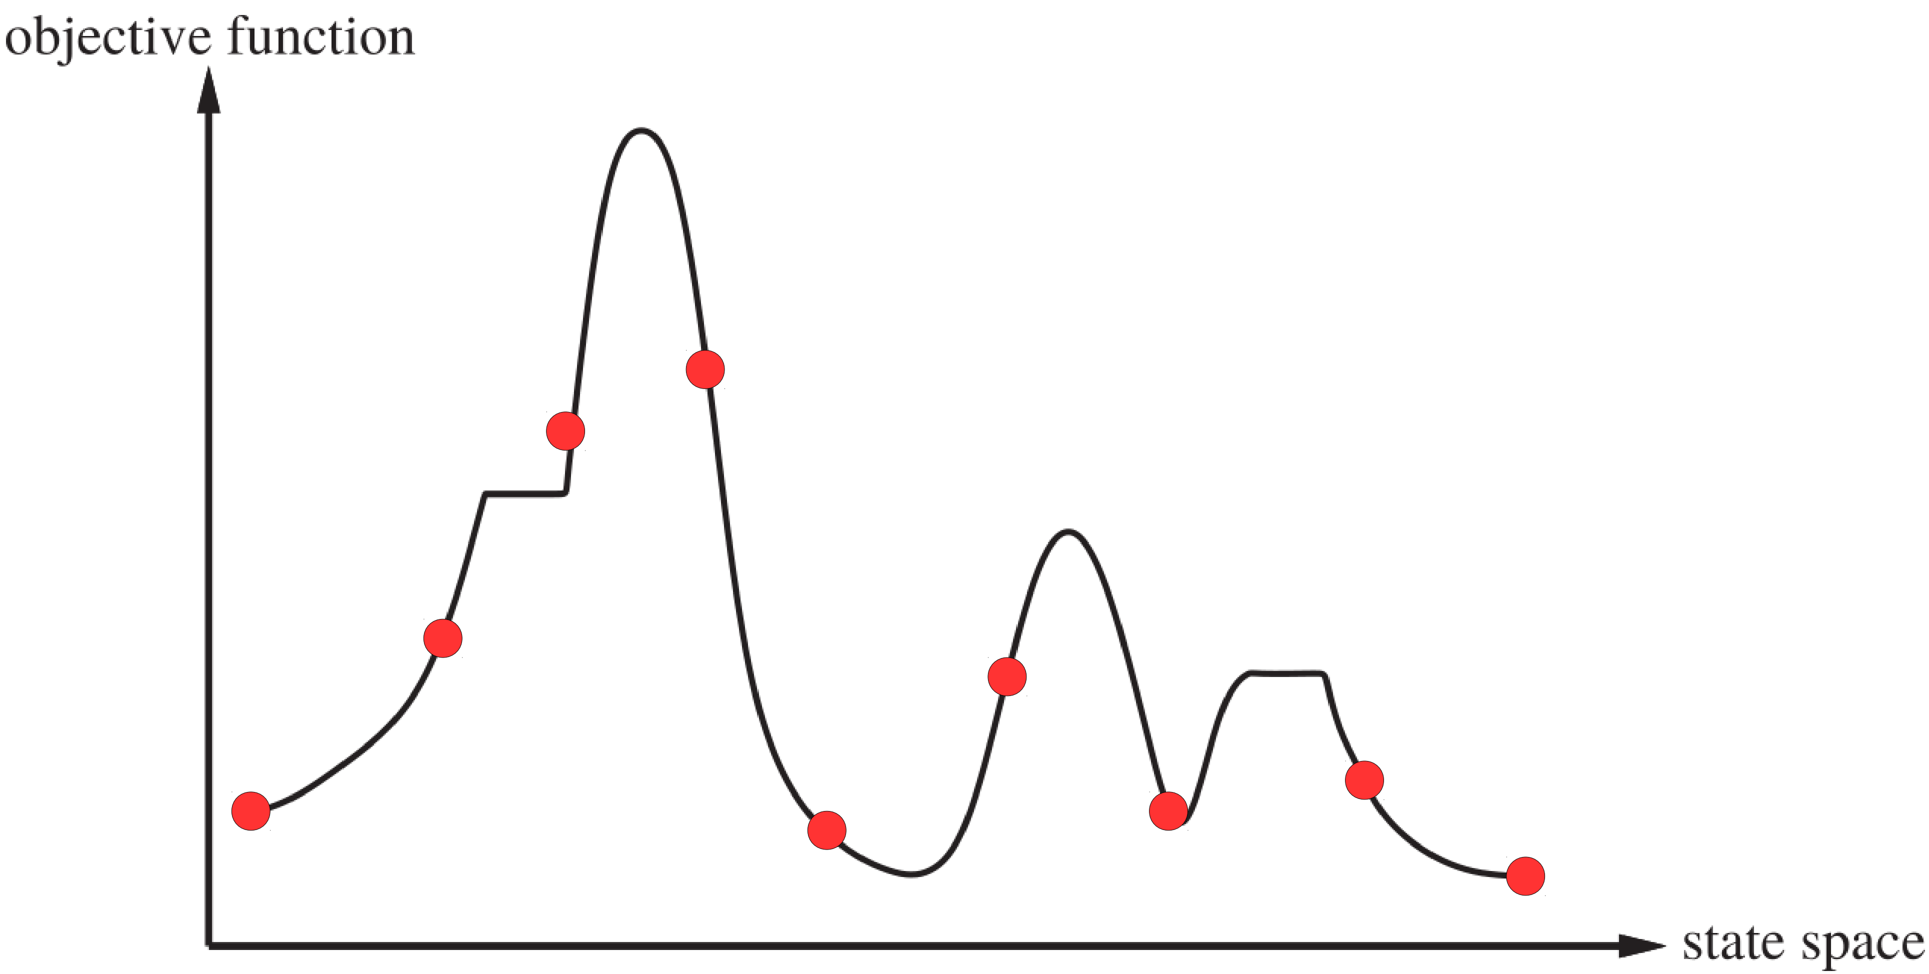
objective function
state space

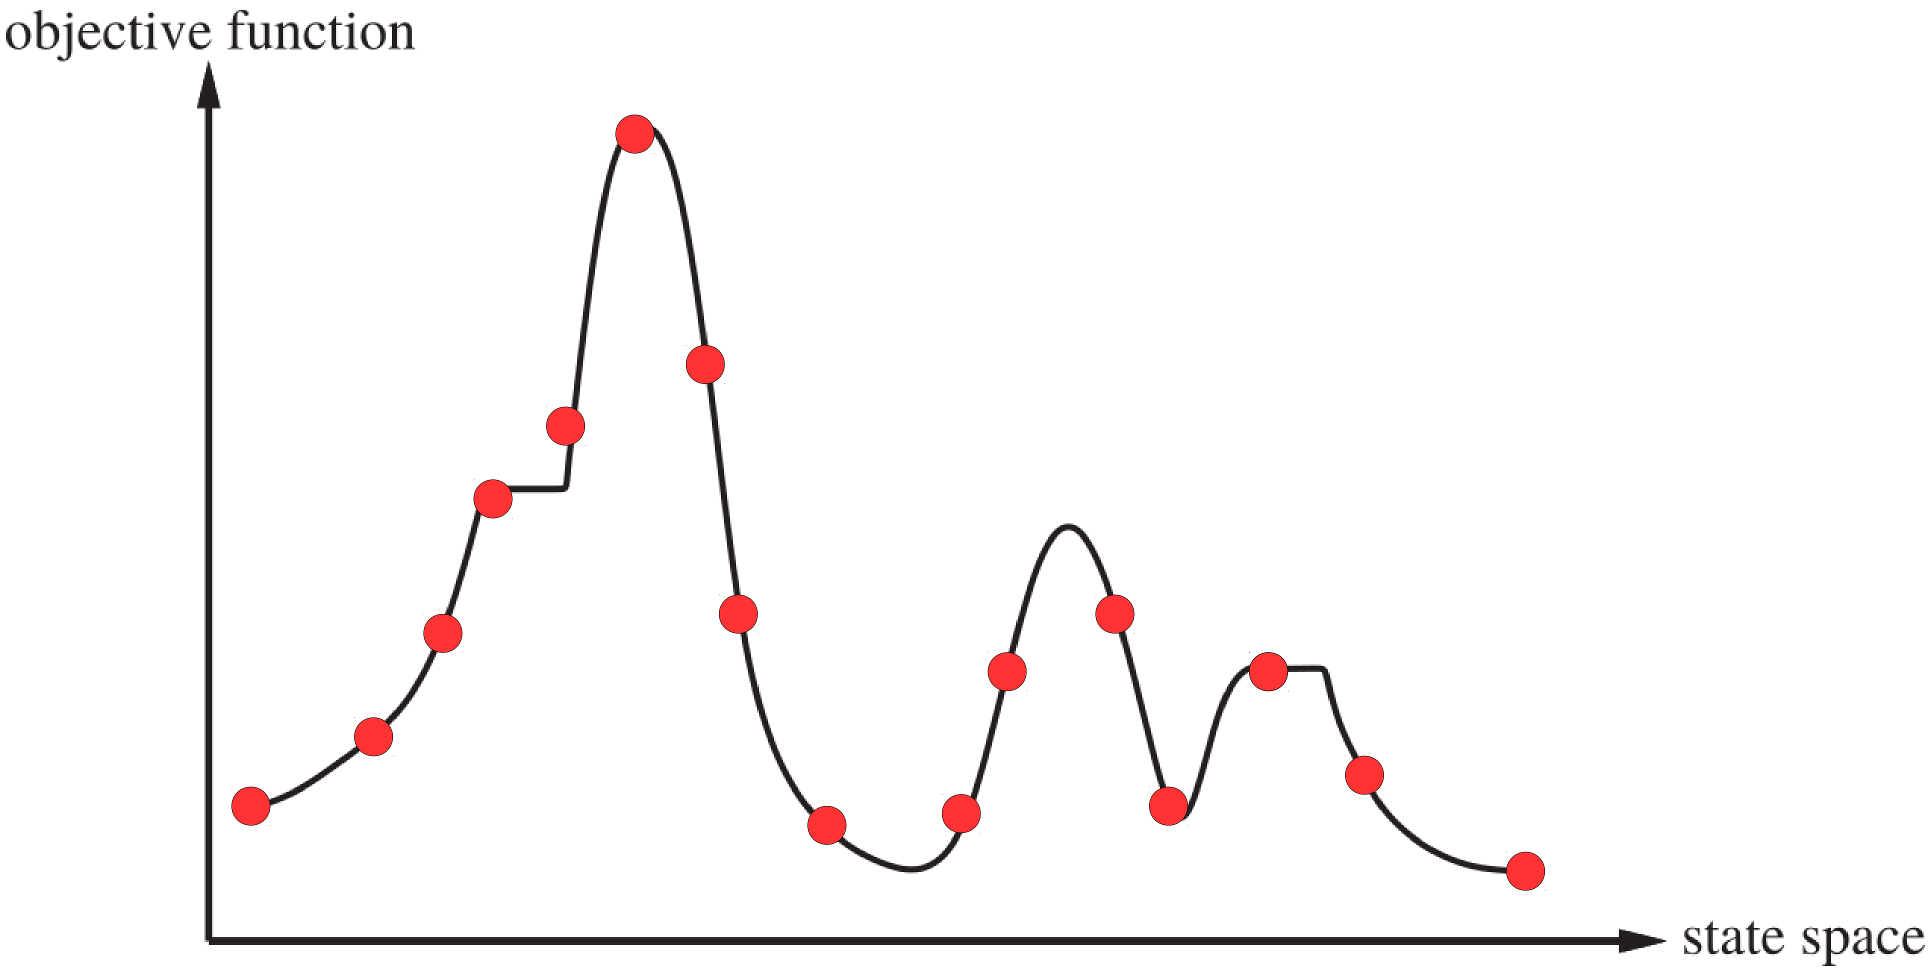
objective function
state space

Medium Map Novelty

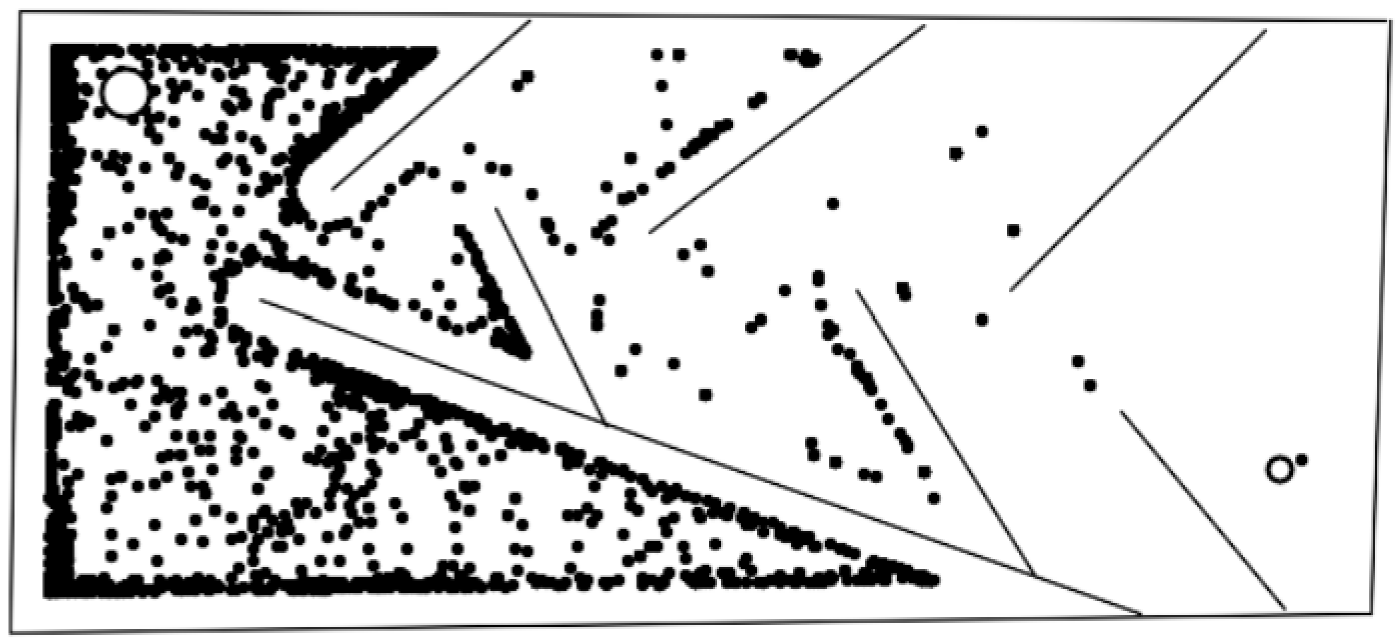

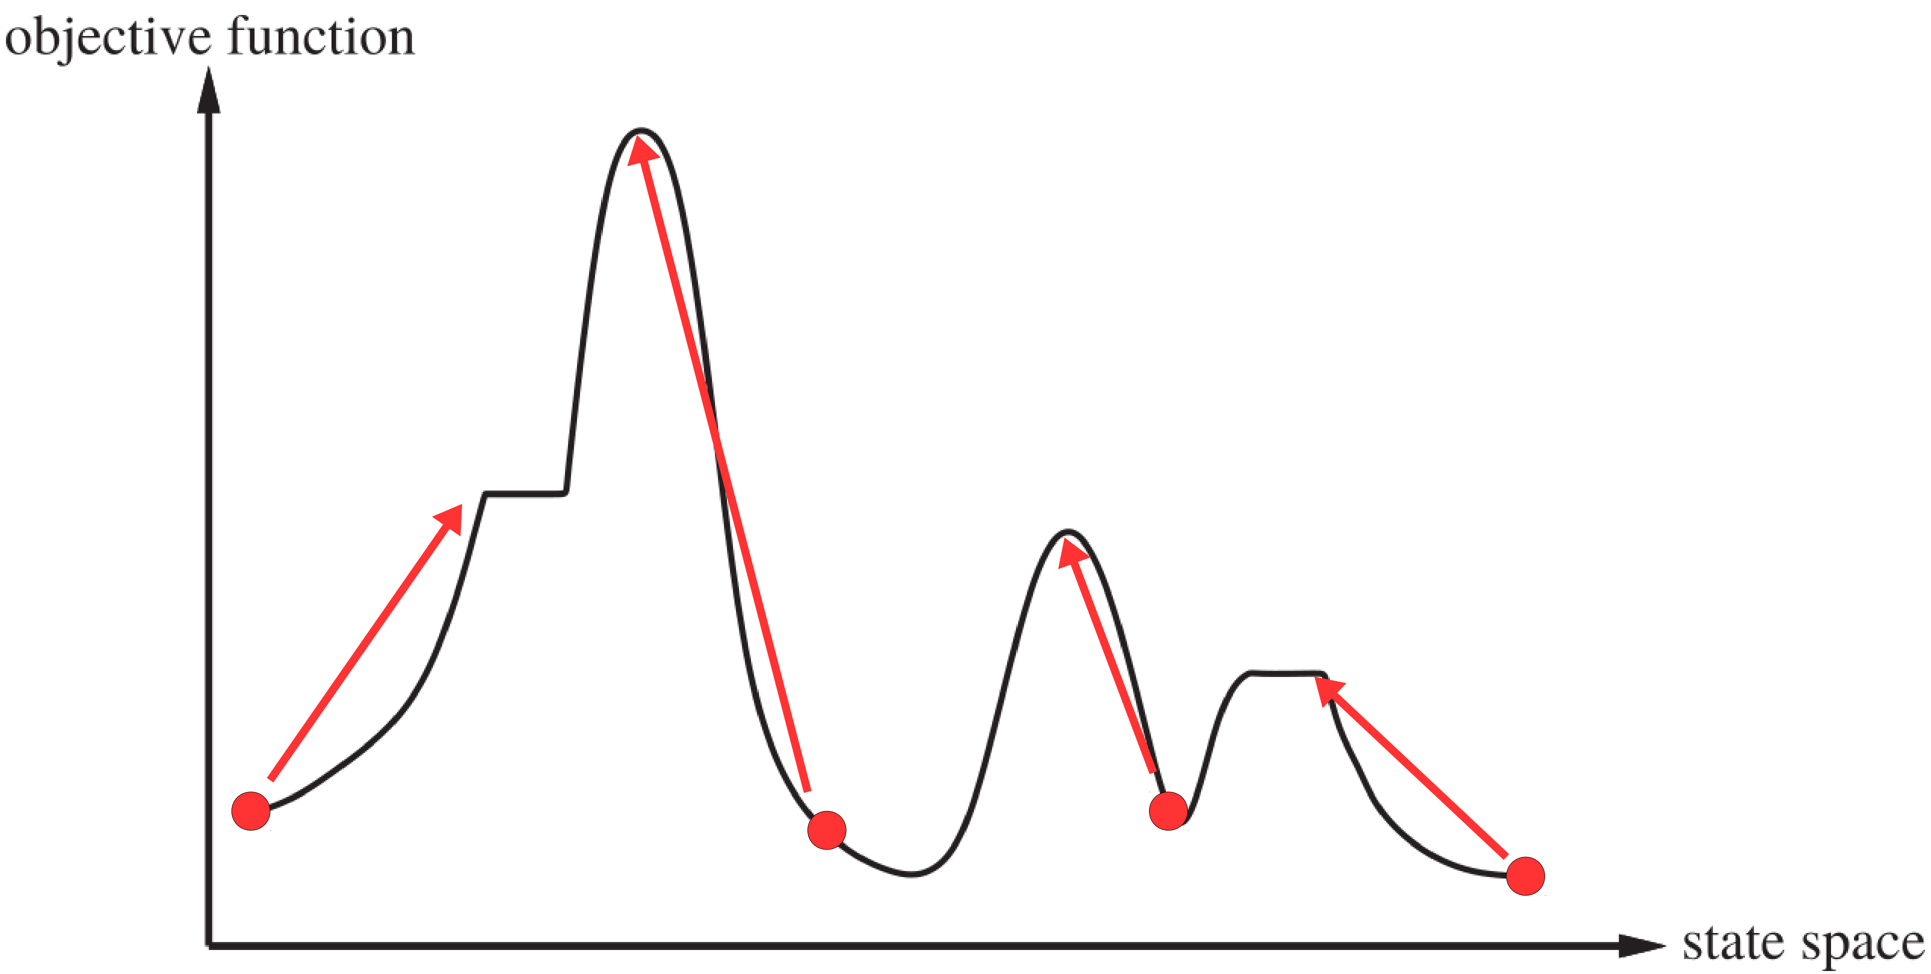
objective function
state space

Evolving a Diversity of Creatures through Novelty Search and Local Competition

In: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2011).* New York, NY:ACM

Joel Lehman

Kenneth O. Stanley

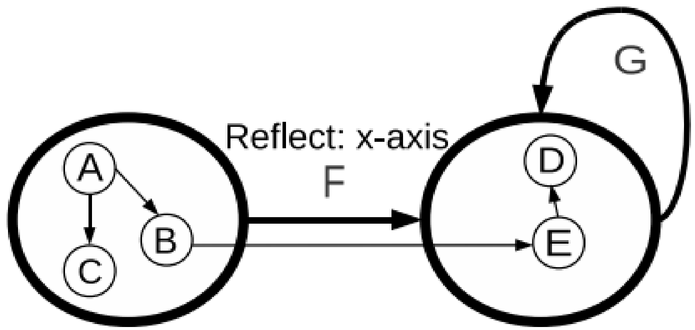

(a) Genotype

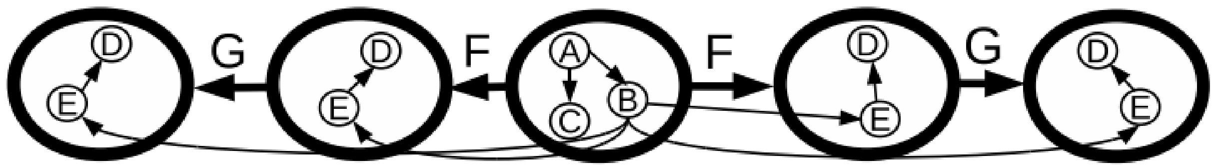

(b) Phenotype Blueprint

Figure 1: **ERO Encoding.** A nested graph genotype is shown in (a) that consists of outer nodes (i.e. the larger circles) that represent morphological parts, connections among outer nodes (e.g. F and G) that represent joints, inner nodes (i.e. the smaller circles) that represent neurons in an ANN, and neural connections among inner nodes (e.g. the connection between B and E). The phenotype shown in (b) is unrolled from the genotype (a) according to flags associated with the connections between outer nodes that can encode symmetry (e.g. the joint F and its children are reflected over the x-axis) and repetition (e.g. joints can be repeated to create articulated appendages like fingers), yielding a hierarchical repeating structure.

a two-dimensional morphology space can be constructed by considering the height and mass of a virtual creature. A search for novelty within this space will effectively explore the space of morphologies spanning those that are short and light to tall and heavy.

Balancing Achievement with Novelty

Multi-objective optimization is a popular paradigm within EC that addresses how to optimize more than one objective at the same time in a principled way [3]. Such multi-objective search suggests a simple way to combine the drive to optimize performance with the search for novelty: Reward both performance and novelty at the same time by making them separate objectives in a multi-objective search [16].

Local Competition

In practice, transforming a global competition score (e.g. the fitness function) into a local competition score requires a comparison of an individual's performance to that of its nearest neighbors in niche space. The more neighbors it outperforms, the higher its local competition score.

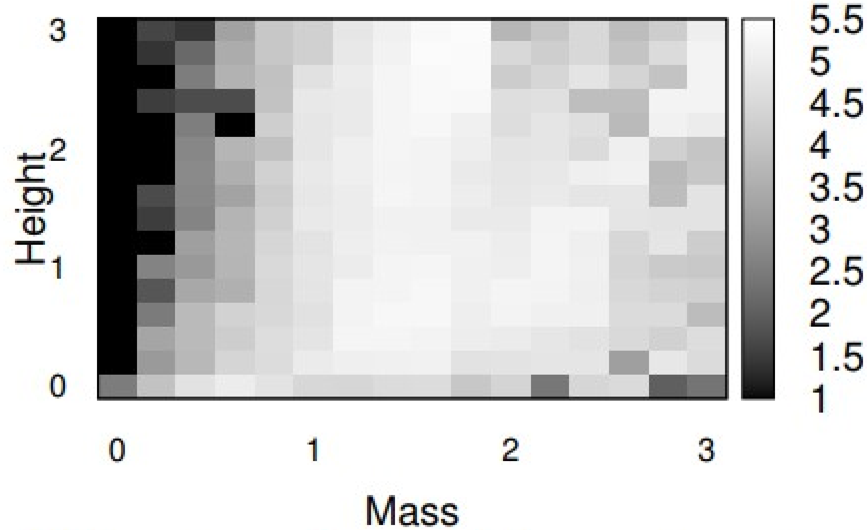

Figure 4: **Niche Capacity.** The capacity of evolution to exploit different morphological niches is illustrated above. Each square represents a segment of morphology space (only two out of three morphological dimensions are visualized), and its darkness is proportional to the logarithm of the highest fitness value found within that segment of morphology space (i.e. darker means more fit), over all runs of all variants; thus it is an estimate of the upper bound of fitness that the niche supports.

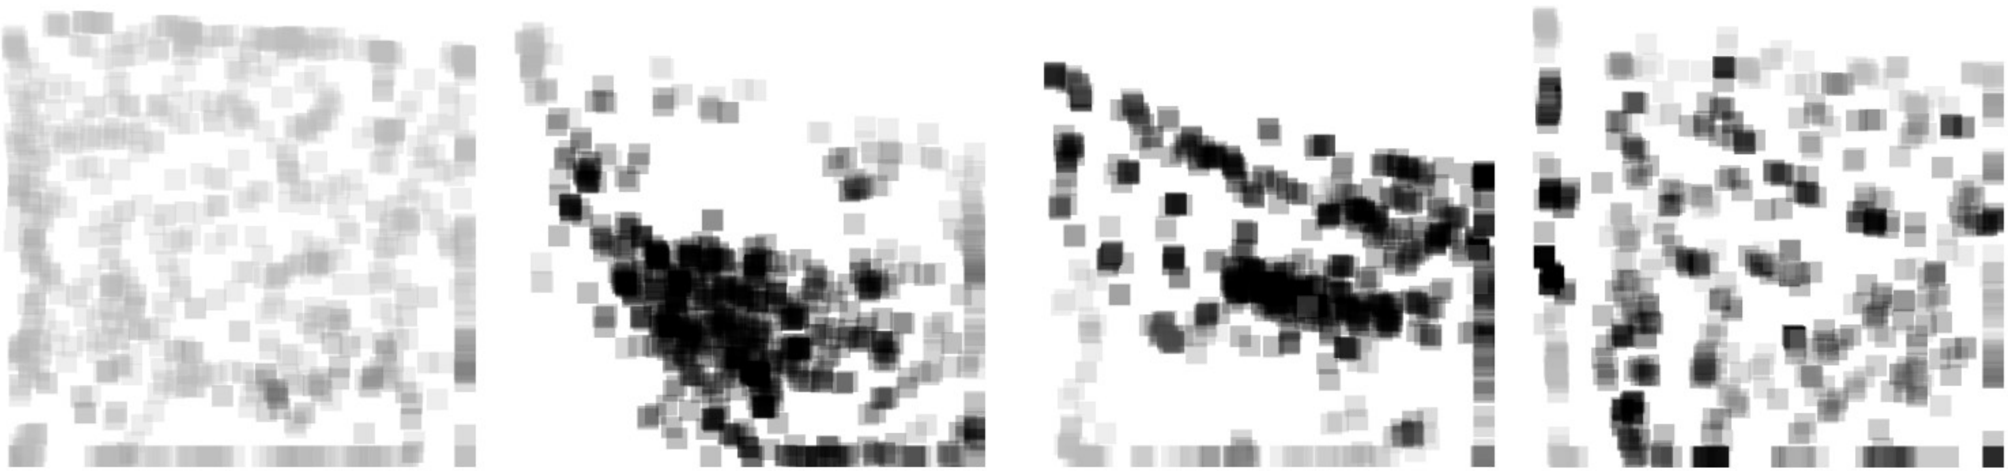

(a) Novelty Only (b) Fitness Only (c) Global Competition (d) Local Competition

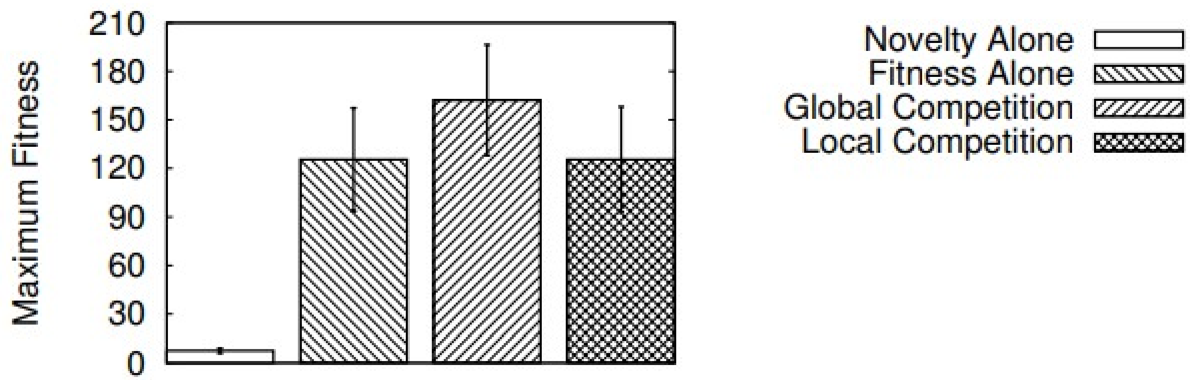

Figure 2: **Absolute Performance Comparison (larger is better).** For each setup, the maximum fitness discovered in a particular run is shown (averaged over 15 runs). The main result is that novelty search with global competition discovers the most fit individuals ($p < 0.001$).

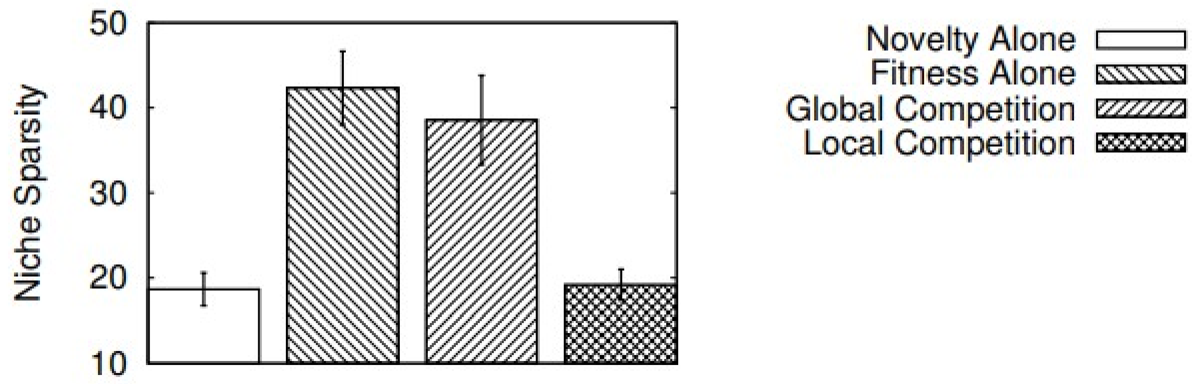

Figure 3: **Niche Sparsity (larger is worse).** For each setup, the average coverage of morphology space of the final population of a particular run is shown (lower is better; averaged over 15 runs). The main result is that novelty with local competition and novelty search alone cover the niche space the best ($p < 0.001$).

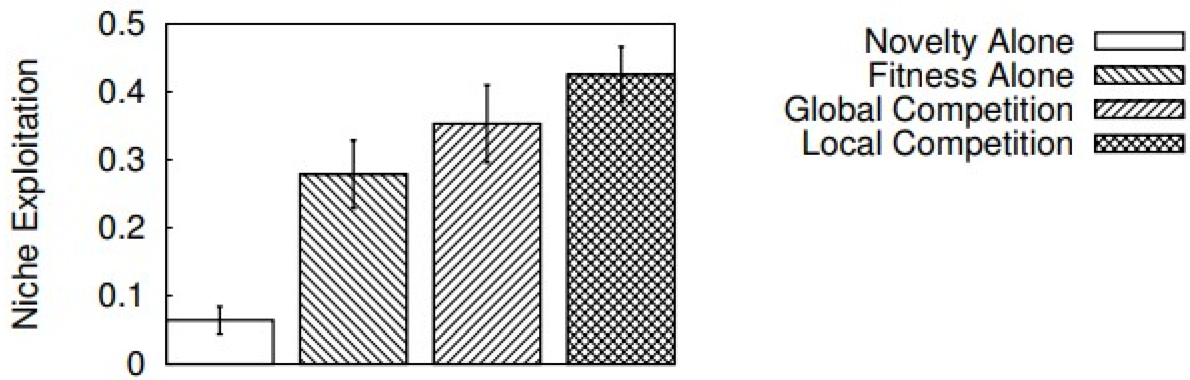

Figure 5: **Niche Exploitation (larger is better).** For each setup, the average niche exploitation is shown (averaged over 15 runs). The main result is that novelty with local competition exploits niches significantly better on average than the other setups ($p < 0.001$).

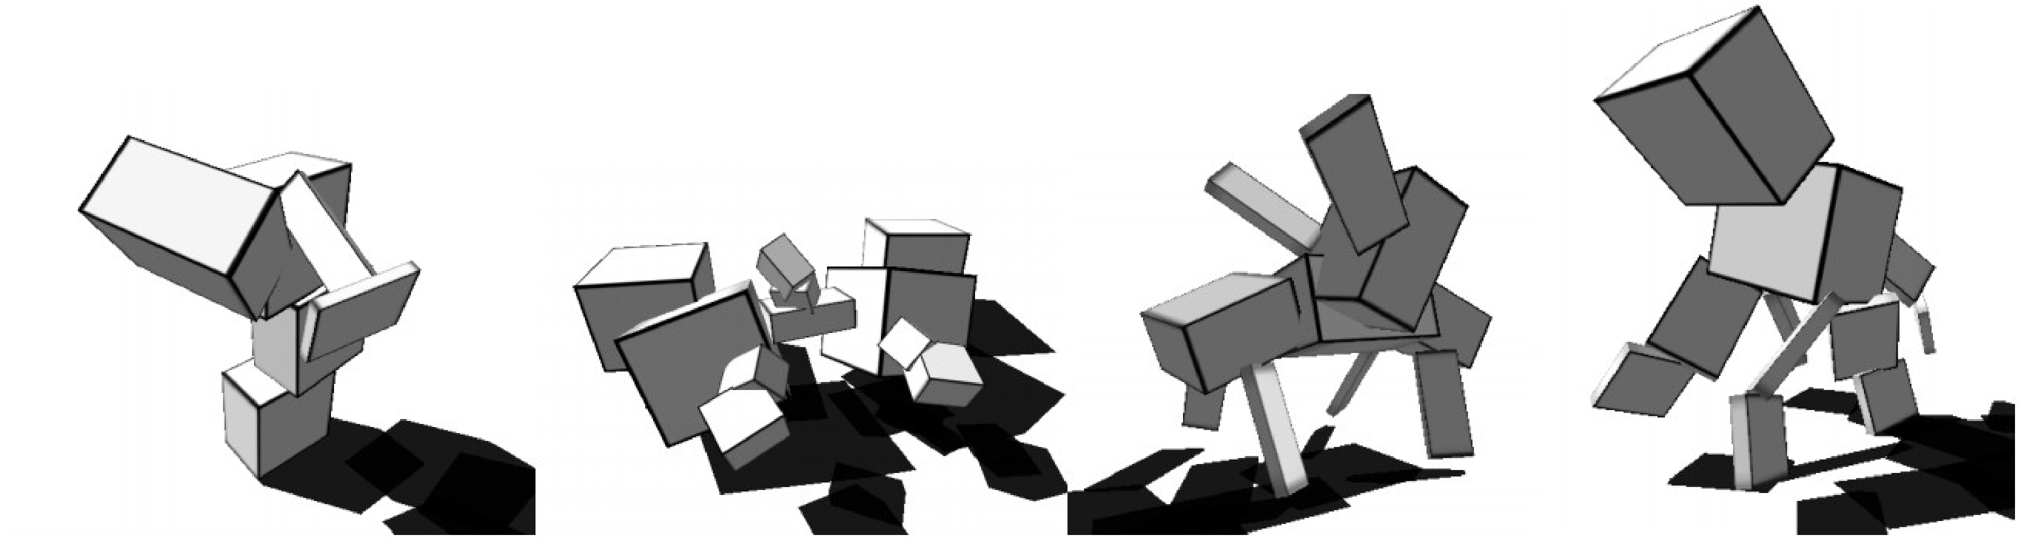

(a) Hopper (b) Crab (c) Quadruped (d) Tailed Quadruped

Figure 7: **Diverse competent morphologies discovered within a typical *single* run of local competition.** Various creatures are shown that have specialized to effectively exploit particular niches of morphology space. These creatures were all found in the final population of a typical run of local competition. The hopper (a) is a unipedal hopper that is very tall, (b) is a heavy short crab-like creature, and (c) and (d) are distinct quadrupeds. Creature (c) drives a large protrusion on its back to generate momentum, and (d) has a tail for balance.