Quality Diversity Optimization

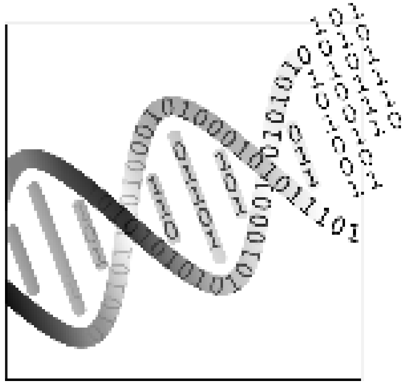

Evolution is diverse

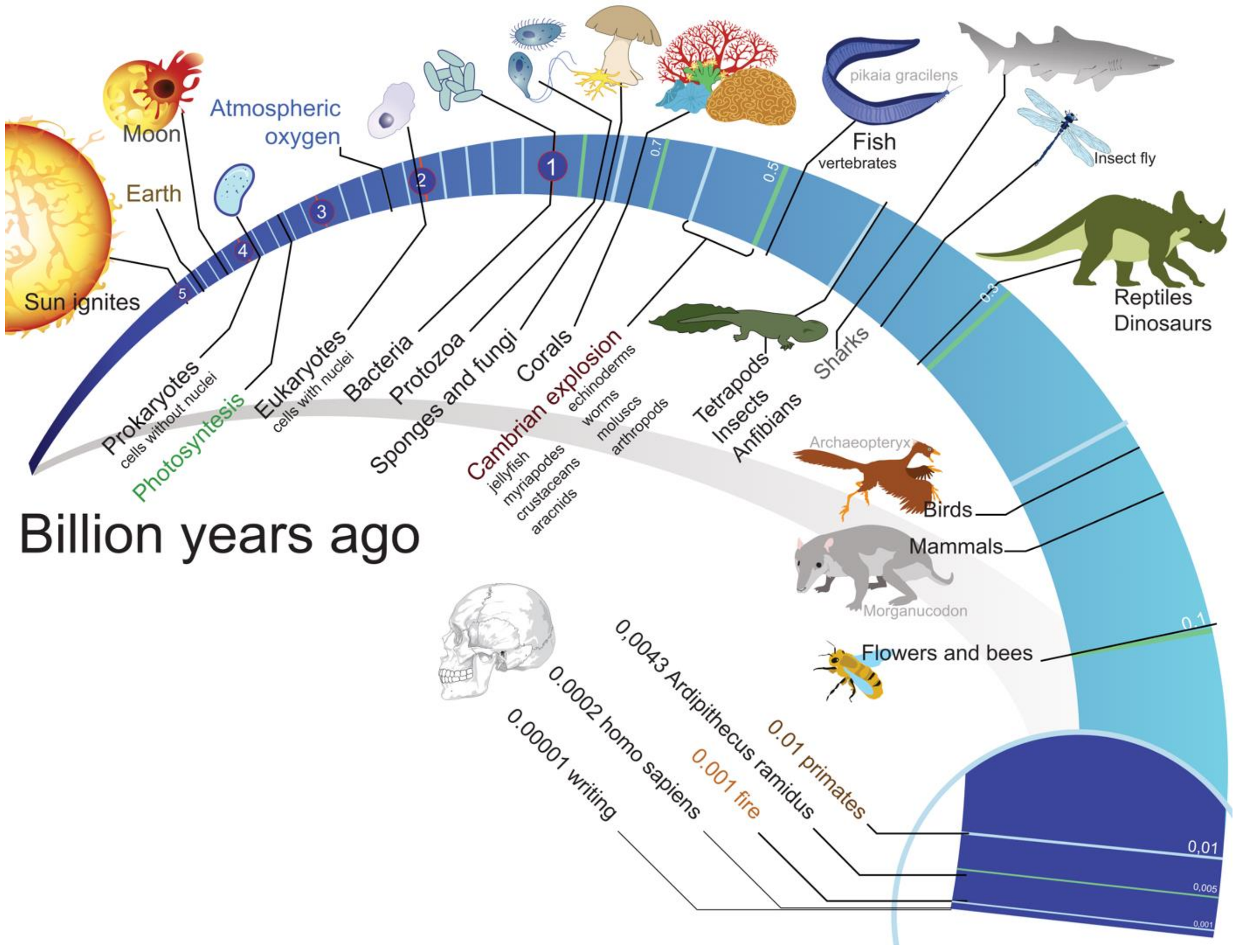

What you will learn in this lecture

- Concept of premature convergence
- Population diversity: genotype vs. phenotype
- Quality Diversity algorithms
  - Open-ended evolution
    - Novelty search
  - Map elites

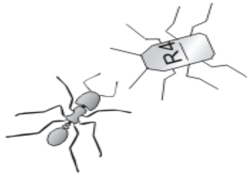

Examples of Evolutionary Algorithms

Genetic Algorithms (GA) - Holland, 1975

Binary genotypes, crossover and mutation

Genetic Programming (GP) - Koza, 1992

Tree-based genotypes, crossover and mutations

Steady-State GA (SSGA) – Whitley et al., 1988

Gradual replacement: Best individuals replace worst individuals

Differential Evolution (DE) – Storn & Prince, 1996

As SSGA, but with differential factor

Evolutionary Strategies (ES) - Rechenberg, 1973

Real-valued genotypes, mutation step(s) encoded in genotype

Covariance Matrix Adaptation ES (CMA-ES) – Hansen & Ostermeier, 2001

Evolutionary Strategies with correlated and adaptive mutations

Non-dominated Sorting GA (NSGA)– Srinivas, Deb, 1998

Multi-objective evolutionary optimization

Viability Evolution (ViE)– Maesani, Mattiussi, Floreano, 2014

Evolution without fitness ranking and diversity preservation

MAP Elites – Mouret and Clune, 2015

Preserve diversity by making similar solutions compete with each other

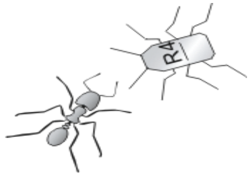

When to stop evolution?

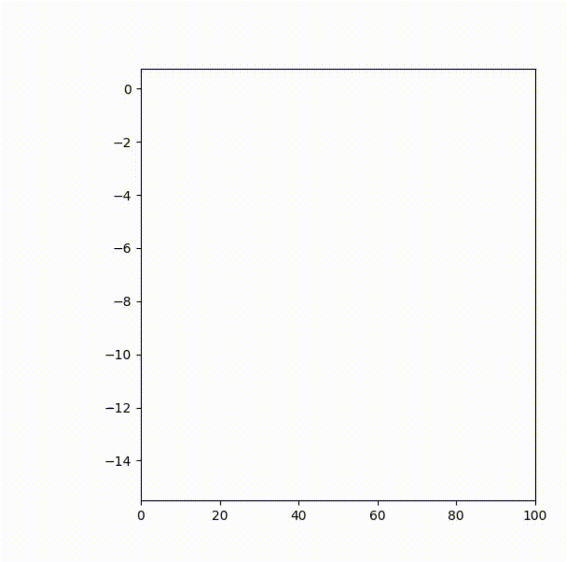

When to stop evolution?

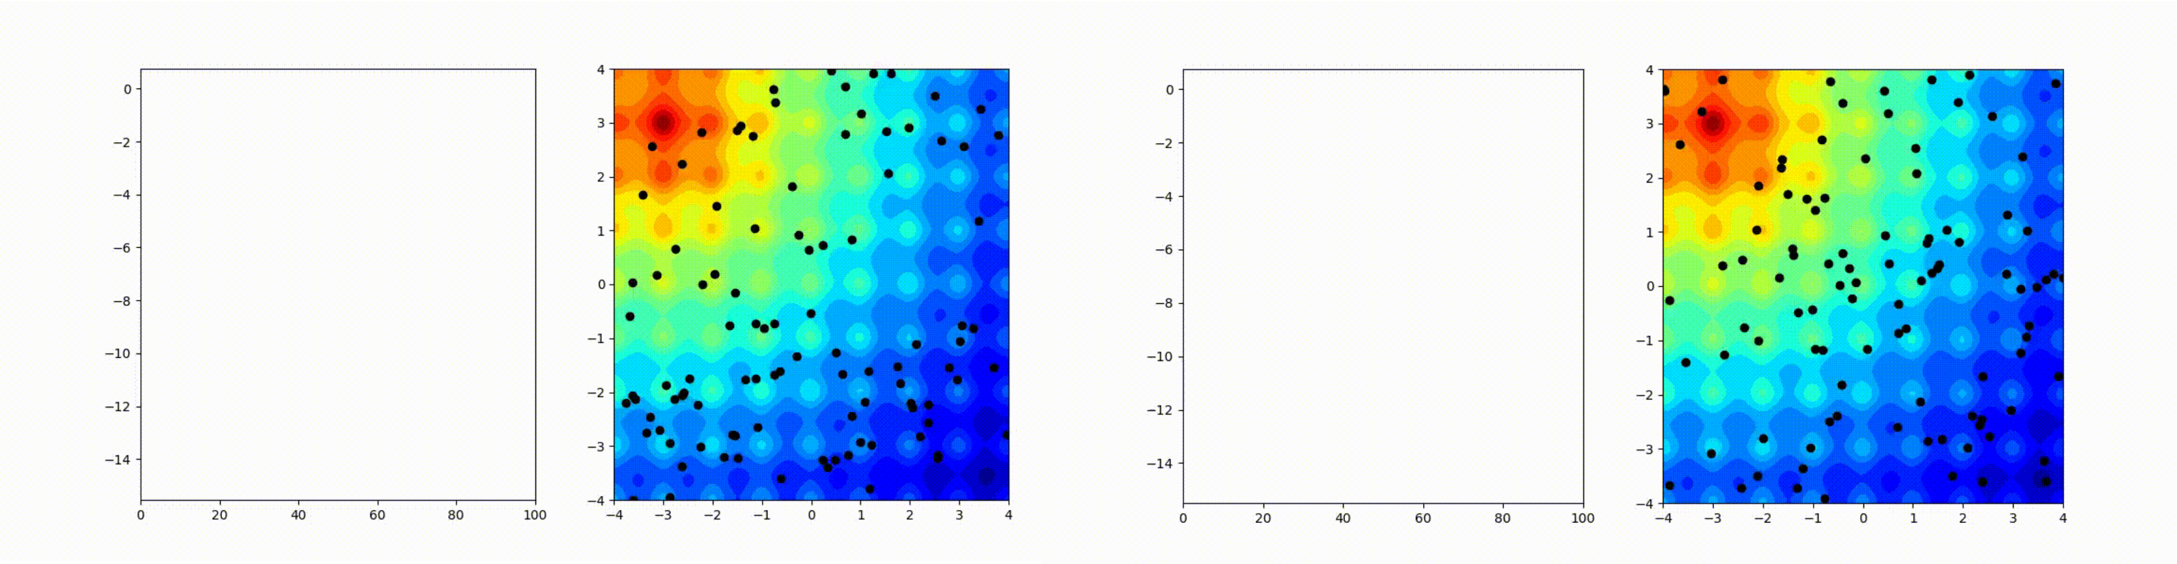

Premature convergence

Premature convergence is the situation where a population loses diversity too quickly and converges to a suboptimal solution (local optimum) before adequately exploring the search space.

- Quick reduction in spread of fitness (sign of convergence)

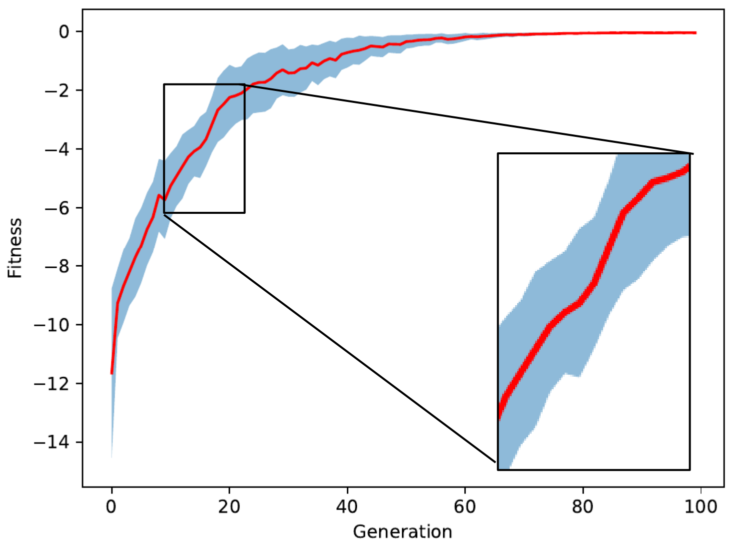

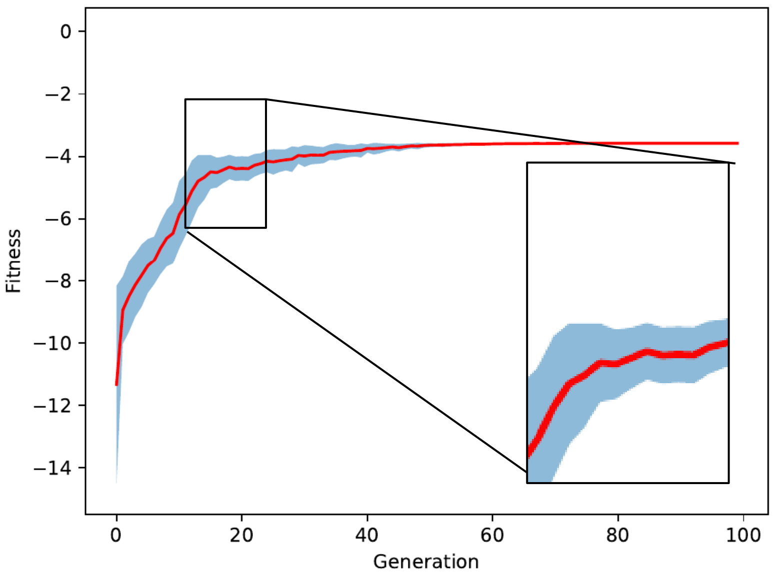

Sensitivity to genotype variation

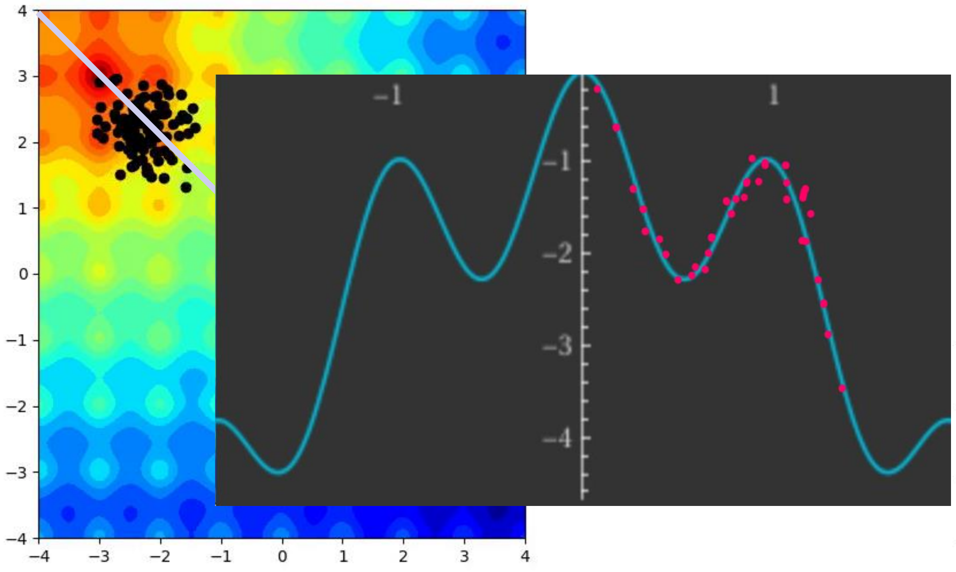

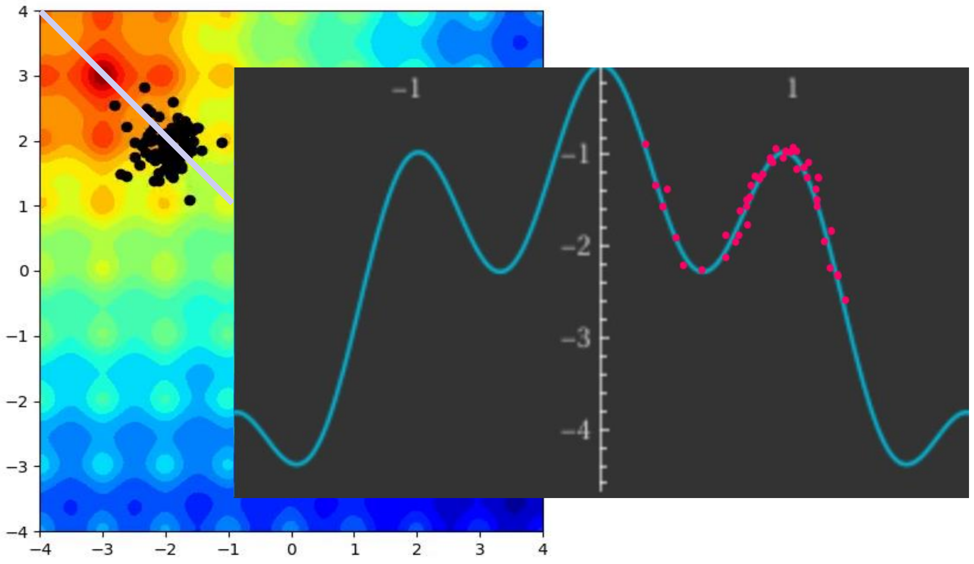

Exploitation vs. Exploration

The EA balancing act:

- Converge on best solution -> Selection
- Search for better genotypes -> mutation

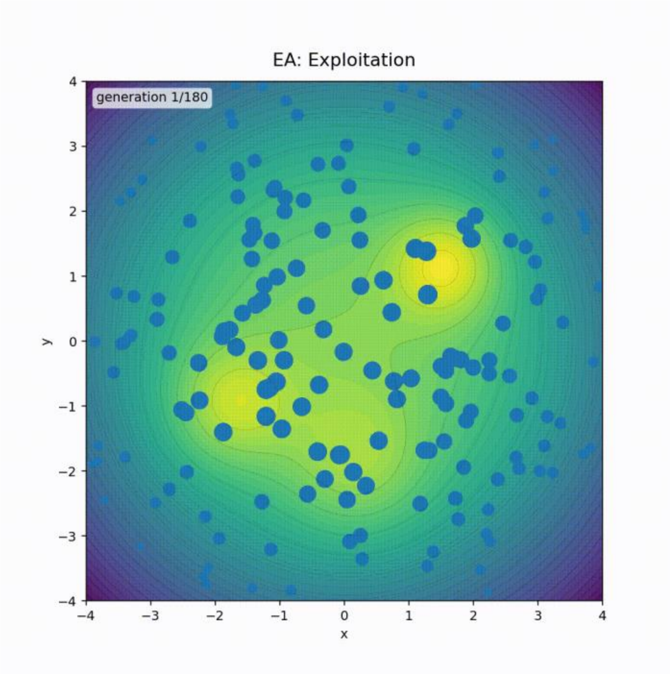

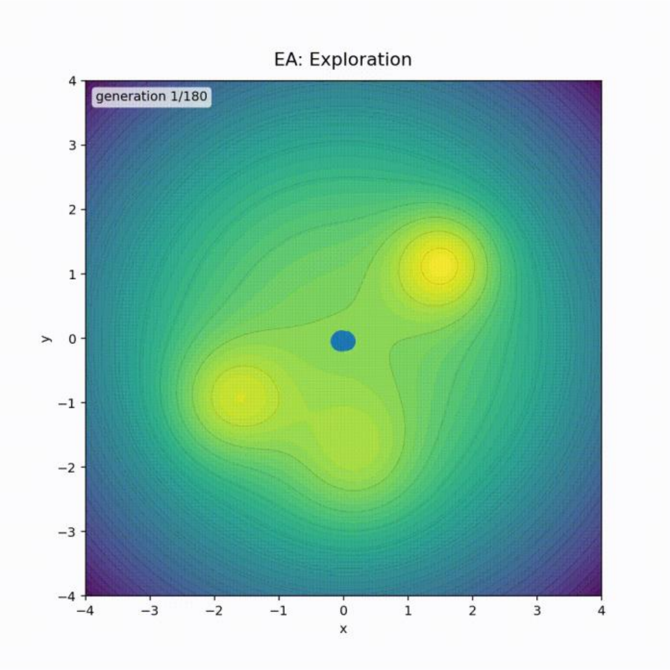

Genotype diversity

Premature convergence is the situation where a population loses diversity too quickly and converges to a suboptimal solution (local optimum) before adequately exploring the search space.

- Reduced spread in fitness (sign of convergence)
- Reduced spread in genotype
  - Adaptive step size (e.g. CMAES)
  - Co-evolve mutation rate

> MIT Technology Review
>
> Featured Topics
>
> scientists estimate we share about 50 percent of our DNA with the banana. This is the nature of DNA in evolution: that all organisms currently alive developed over the eons from common earlier organisms.

Maze problem

Maze problem:

- Robot actions {<,^,v,>}
- Genotype = T actions (e.g. [>,v,>])
- Fitness = dist(robot_pos$_T$, end_goal)

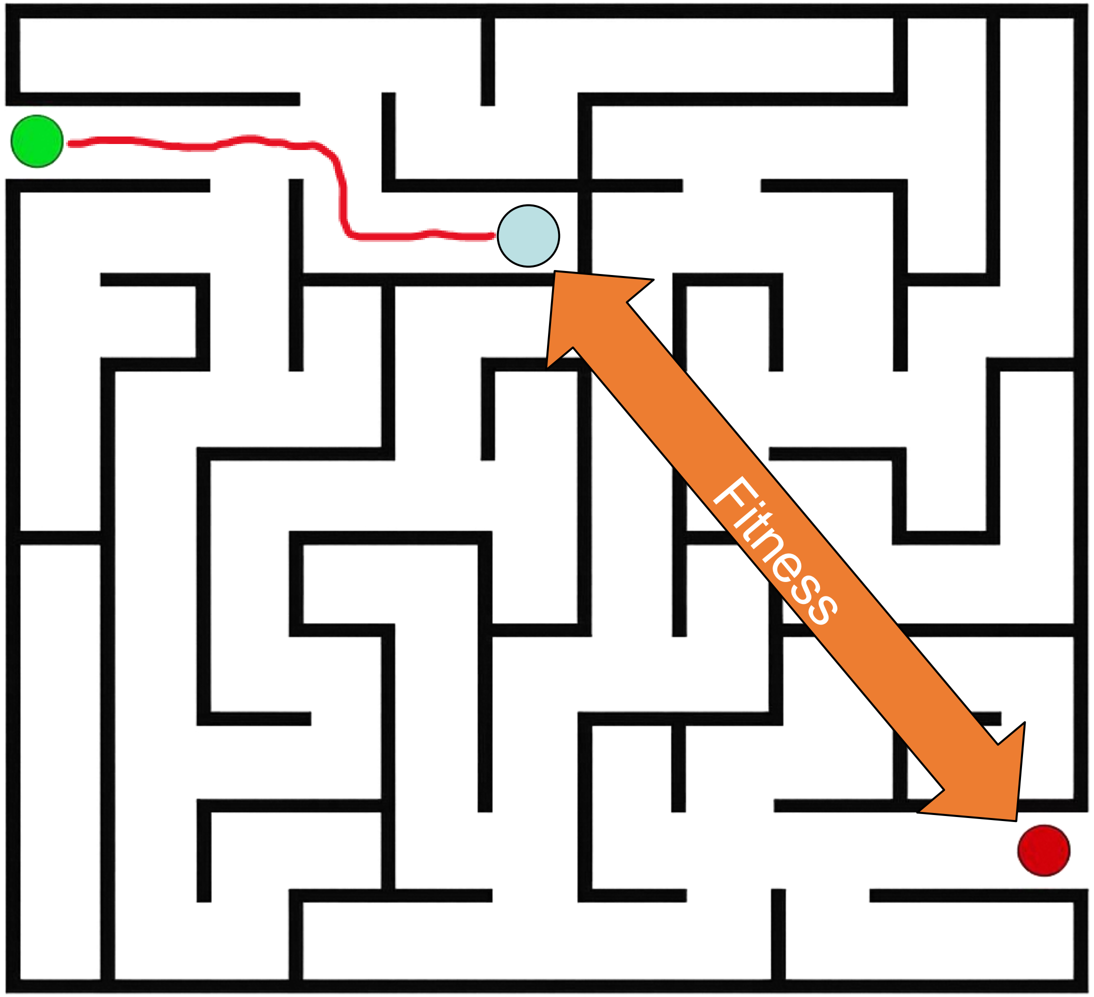

The deceptive trap

Optimizing the maze problem leads to a deceptive trap

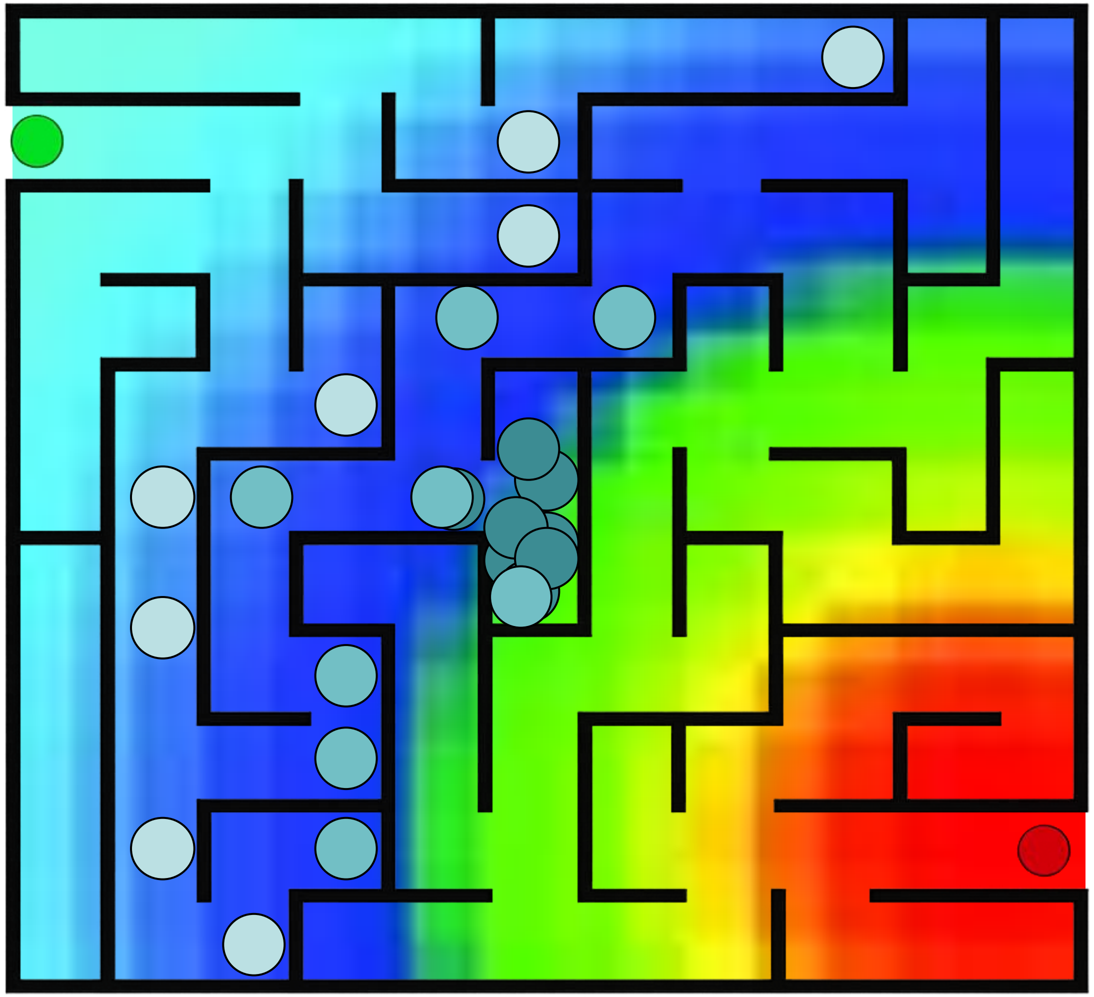

Genotype variation is not enough

Maze problem:

- Robot actions {<,^,v,>}

T1 = {>,^,<,>,v,^,<,v,<,v,<,v,>,^,<,^,>}

T2 = {>,v,<,v,>,^,<,^,>,^,>,^,v,<,v,>,v}

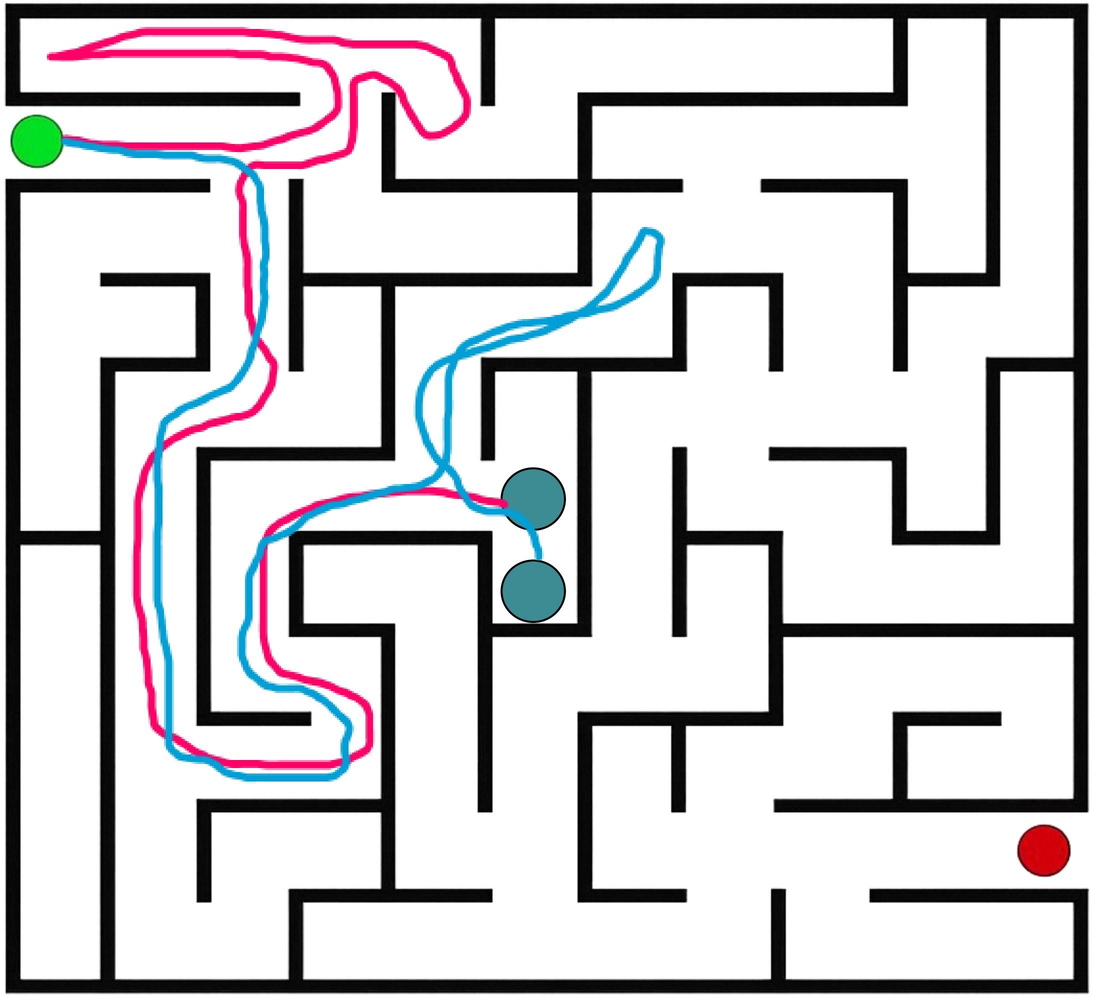

Genotype vs. Phenotype

Premature convergence is the situation where a population loses diversity too quickly and converges to a suboptimal solution (local optimum) before adequately exploring the search space.

- Reduced spread in fitness (sign of convergence)
- Reduced spread in genotype
  - Adaptive step size (e.g. CMAES)
  - Co-evolve mutation rate
- Reduced spread in phenotype

Deception in Objective-Based Search

- When Fitness is calculated as the distance between the robot's final position and the end goal.
- The more complex the maze, the more deceptive the fitness landscape becomes, trapping standard algorithms in local optima

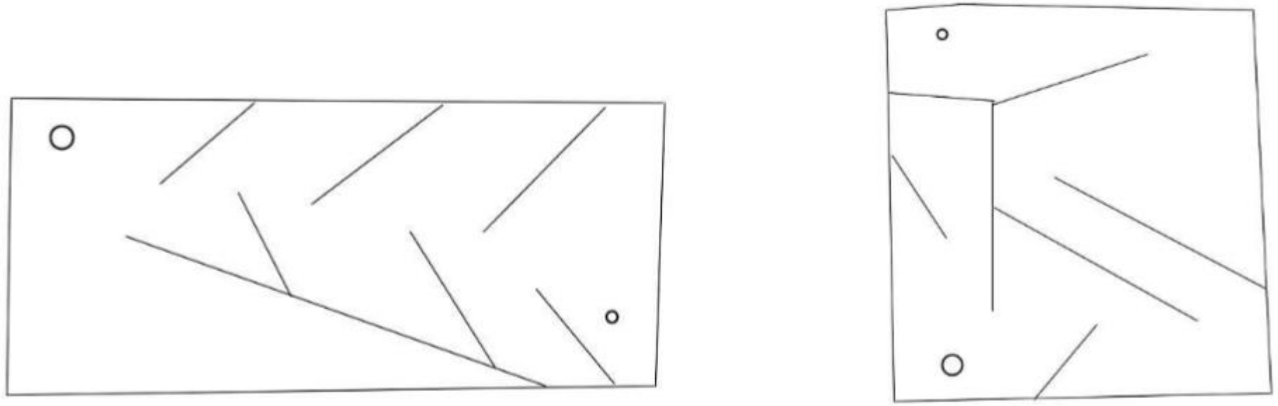

(a) Medium Map (b) Hard Map

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).
Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

Novelty search

The *novelty* of a newly generated individual is computed with respect to the *behaviors* (i.e. phenotype) of an *archive* of past individuals whose behaviors were highly novel when they originated.

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).

Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

Novelty search

The *novelty* of a newly generated individual is computed with respect to the *behaviors* (i.e. phenotype) of an *archive* of past individuals whose behaviors were highly novel when they originated.

1. Archive N individuals
2. Measure phenotype distance of each new individual with respect to k-nearest neighbors in the archive
3. If the measure exceeds a threshold, add new individual to the archive

$$\rho(x) = \frac{1}{k} \sum_{i=0}^{k} \mathrm{dist}(x, \mu_i)$$

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).
Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

Novelty search: 1) Archive N individuals

Archive:

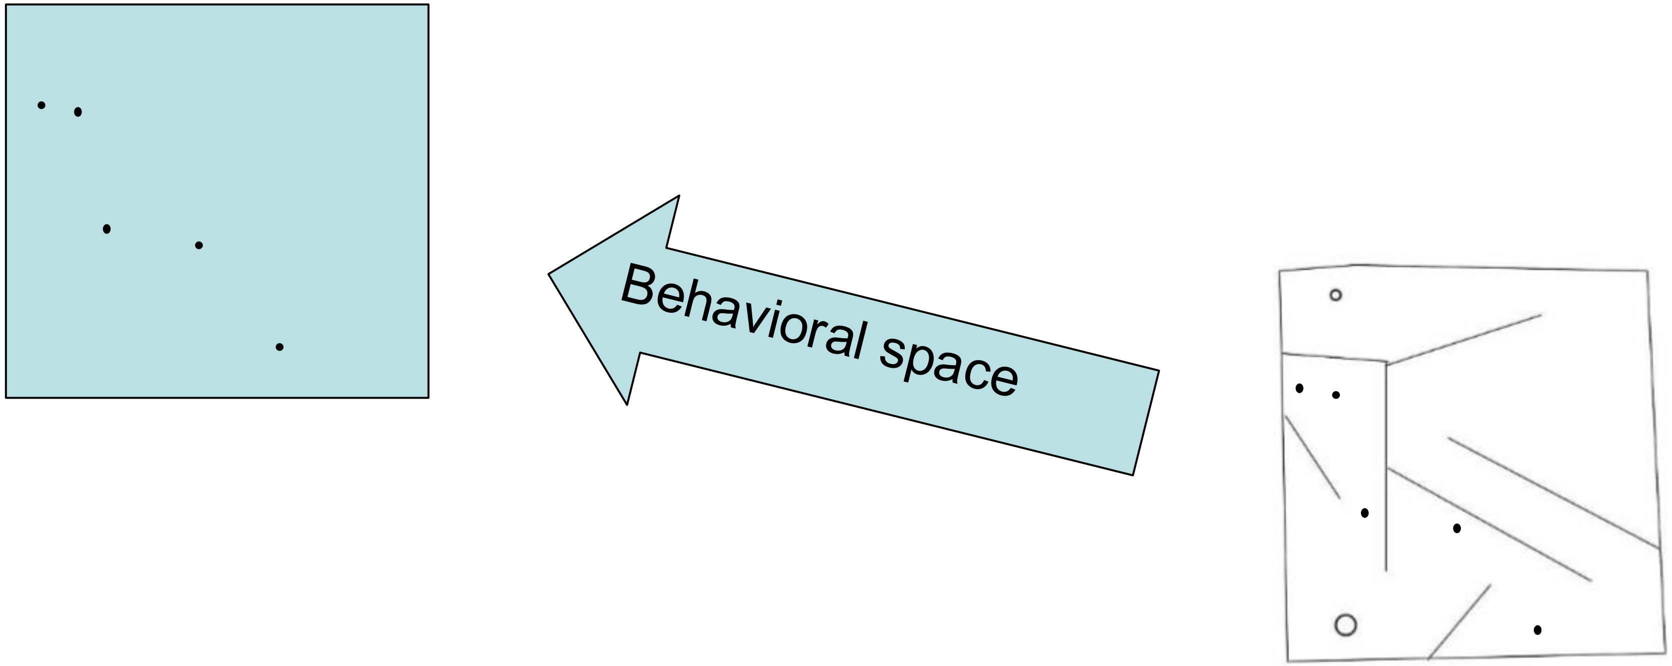

(b) Hard Map

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).
Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

Novelty search: 2) Measure phenotype distance

Archive:

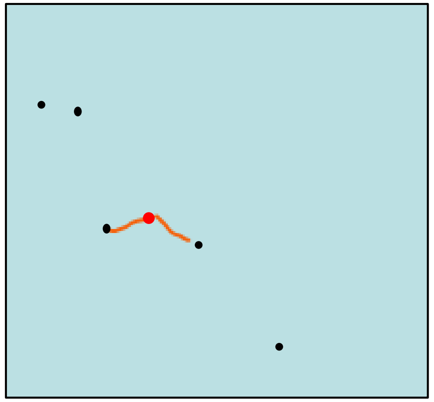

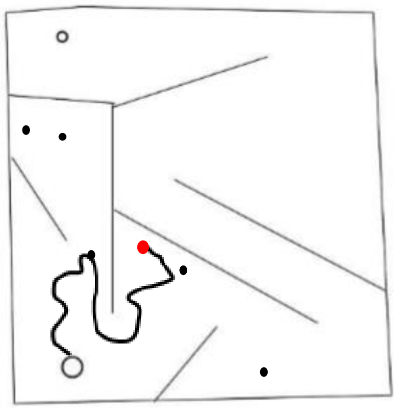

(b) Hard Map

$$\rho(x) = \frac{1}{k}\sum_{i=0}^{k} \mathrm{dist}(x, \mu_i)$$

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).

Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

Novelty search: 3) if measure exceeds threshold..

Archive:

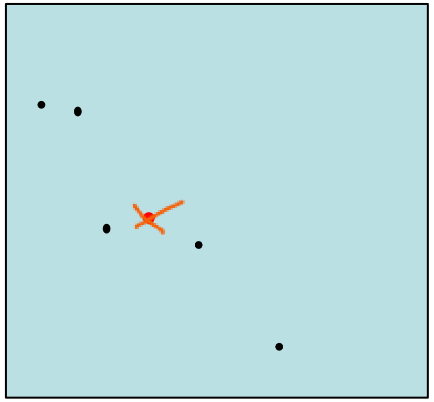

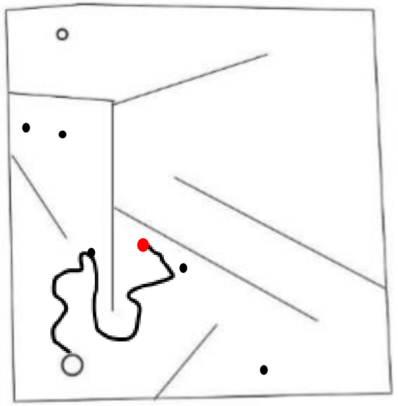

(b) Hard Map

$$\rho(x) = \frac{1}{k}\sum_{i=0}^{k} \text{dist}(x, \mu_i)$$

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).
Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

Novelty search: 3) if measure exceeds threshold..

Archive:

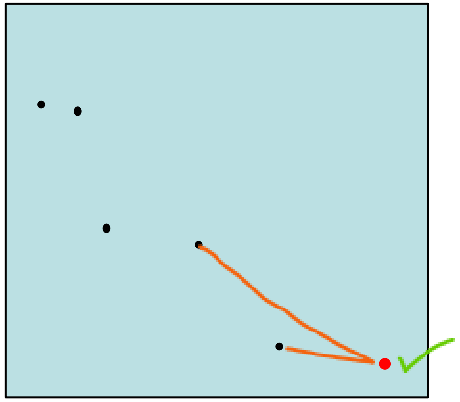

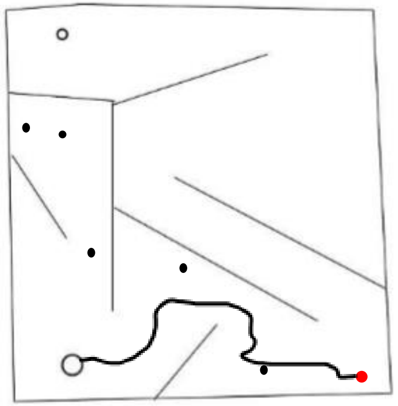

(b) Hard Map

$$\rho(x) = \frac{1}{k}\sum_{i=0}^{k} \mathrm{dist}(x, \mu_i)$$

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).

Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

New objective: fill archive as much as possible

Instead of goal-based optimization that can get trapped

We fill the behavioral space by optimizing for novelty

$$\rho(x) = \frac{1}{k}\sum_{i=0}^{k} \mathrm{dist}(x, \mu_i)$$

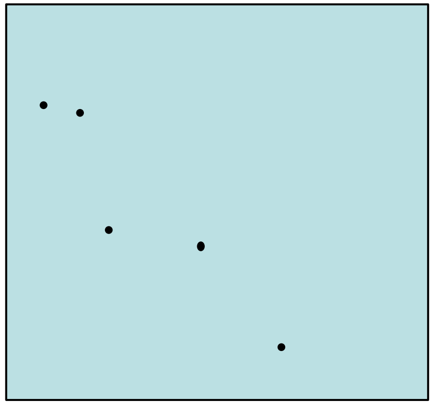

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).

Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

Novelty search outperforms distance-based fitness

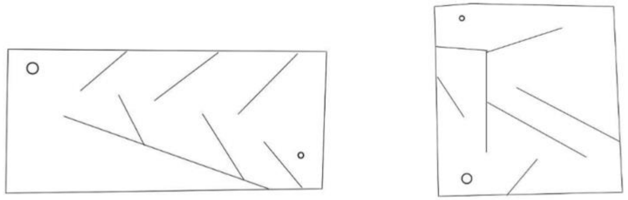

(a) Medium Map (b) Hard Map

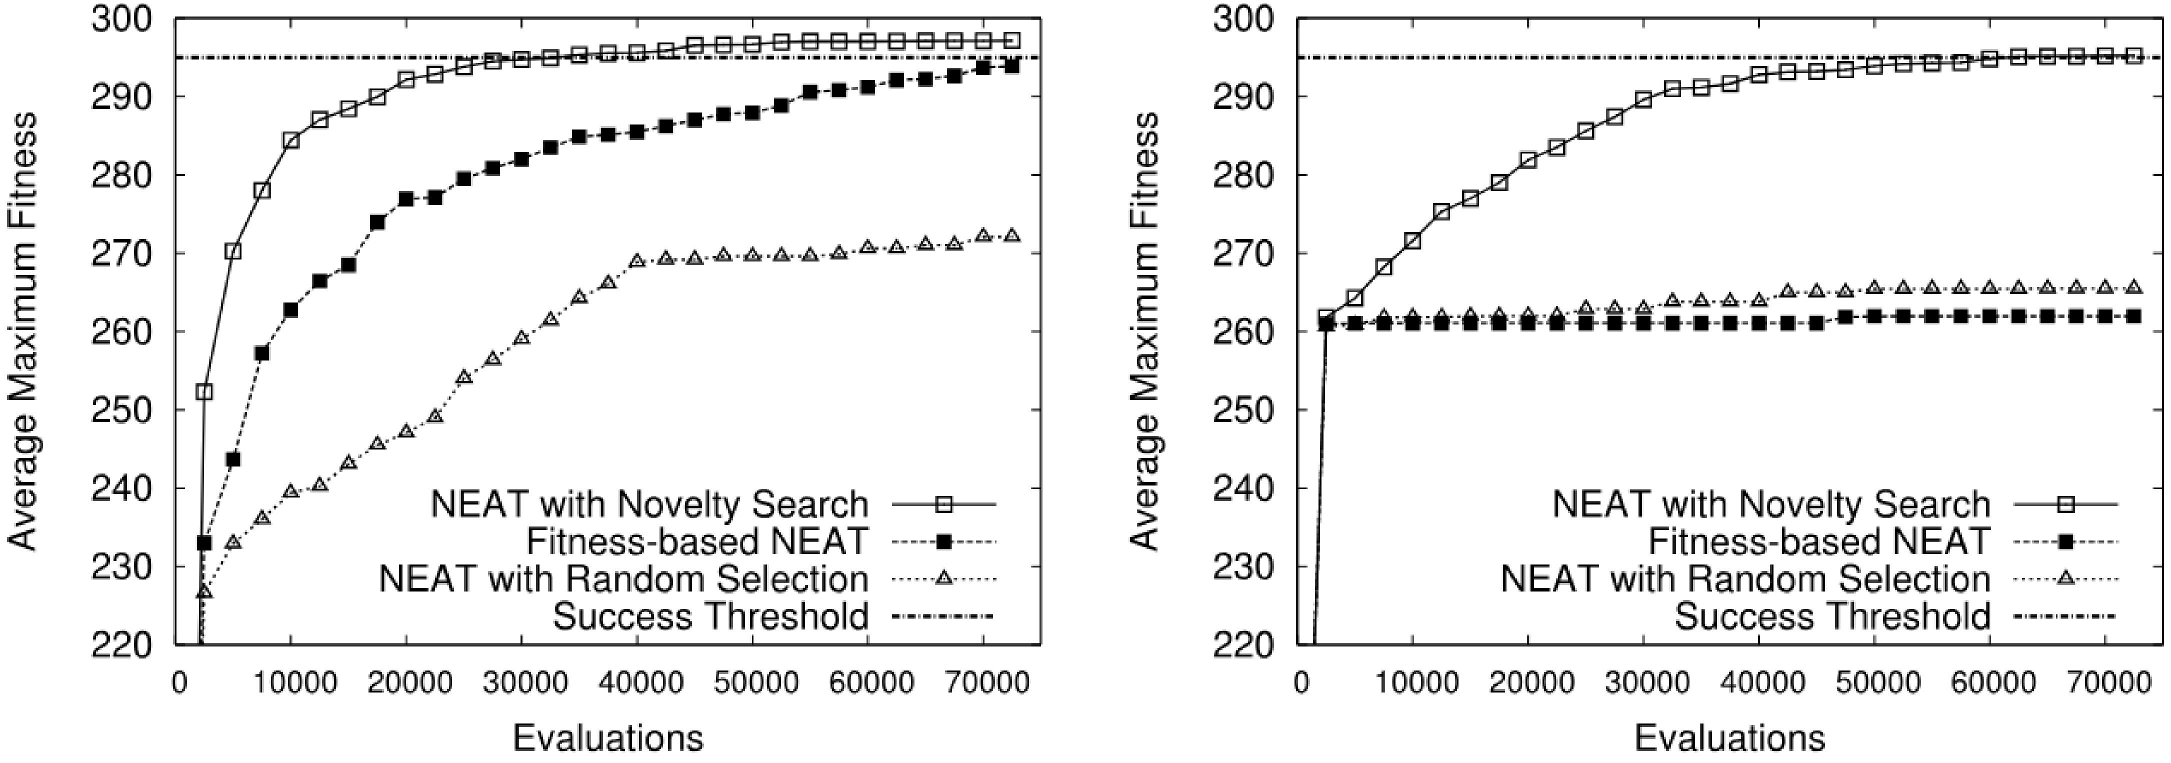

(a) Medium Map (b) Hard Map

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).

Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

Novelty search completes hard maze

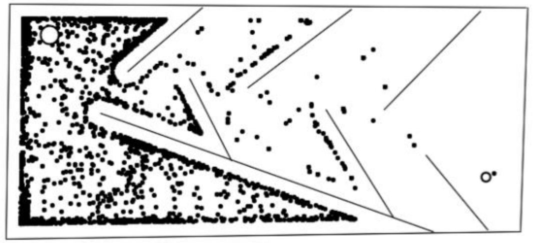

(a) Medium Map Novelty

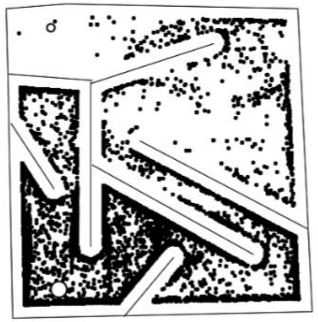

(b) Hard Map Novelty

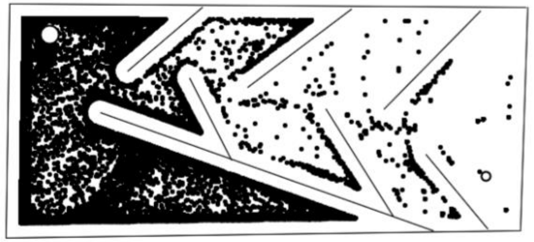

(c) Medium Map Fitness

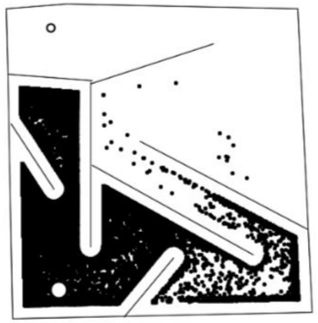

(d) Hard Map Fitness

Lehman, J., & Stanley, K. O. (2008, August). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE* (Vol. 65, pp. 329-336).

Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, *19*(2), 189-223.

Novelty search: key takeaways

- Instead of optimizing for an objective, Novelty Search rewards unique behaviors.
- The algorithm bypasses deceptive local optima, by optimizing phenotypical variation.

Open-ended evolution

- Premature convergence is caused by focusing on an objective.

Open-ended evolution redefines optimization as a never-ending process, rather than settling into a stable equilibrium.

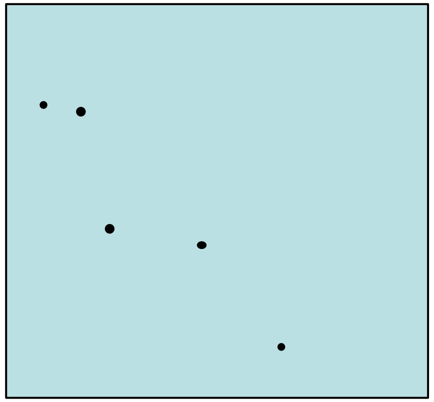

Open-ended evolution

- Premature convergence is caused by focusing on an objective.

Open-ended evolution redefines optimization as a never-ending process, rather than settling into a stable equilibrium.

- Relativism (e.g., peer-comparison)

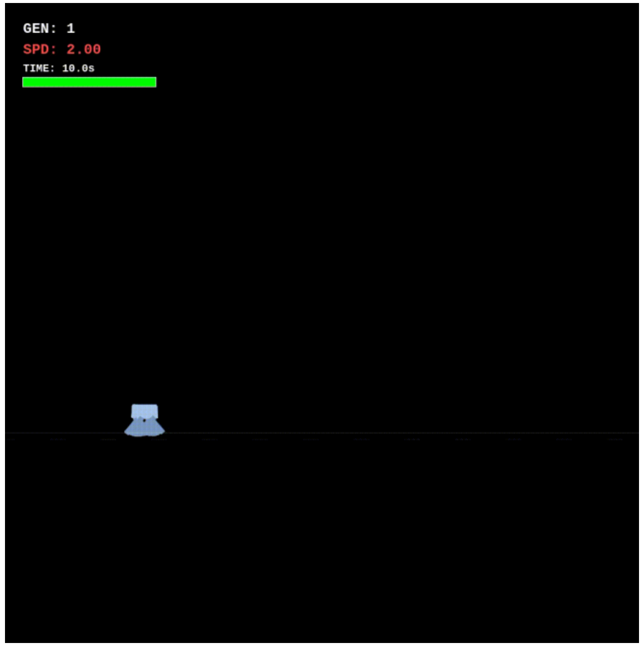

Open-ended evolution

- Premature convergence is caused by focusing on an objective.

Open-ended evolution redefines optimization as a never-ending process, rather than settling into a stable equilibrium.

- Relativism (e.g., peer-comparison)
- Artificial life (e.g. Cellular automata)

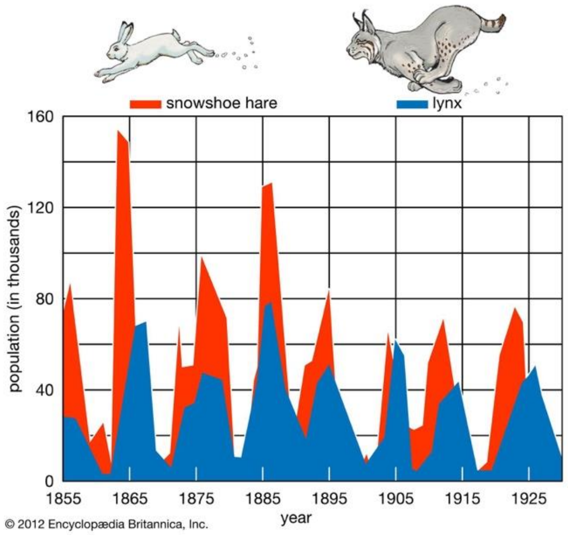

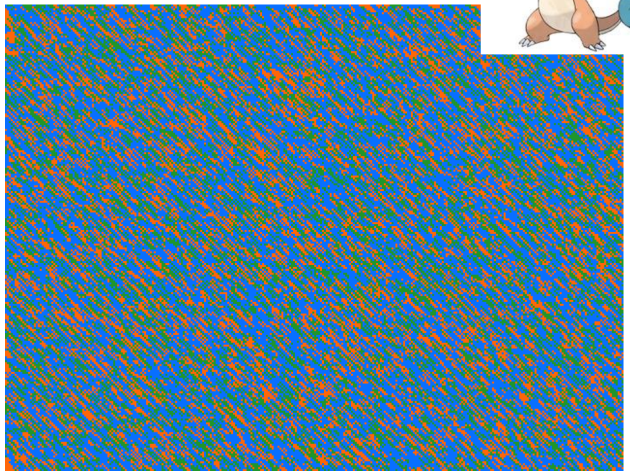

https://kennycason.com/posts/2016-11-25-cellular-automata-pokemon-types.html

Open-ended evolution

- Premature convergence is caused by focusing on an objective.

Open-ended evolution redefines optimization as a never-ending process, rather than settling into a stable equilibrium.

- Relativism (e.g., peer-comparison)
- Artificial life (e.g., Cellular automata)
- Viability Evolution

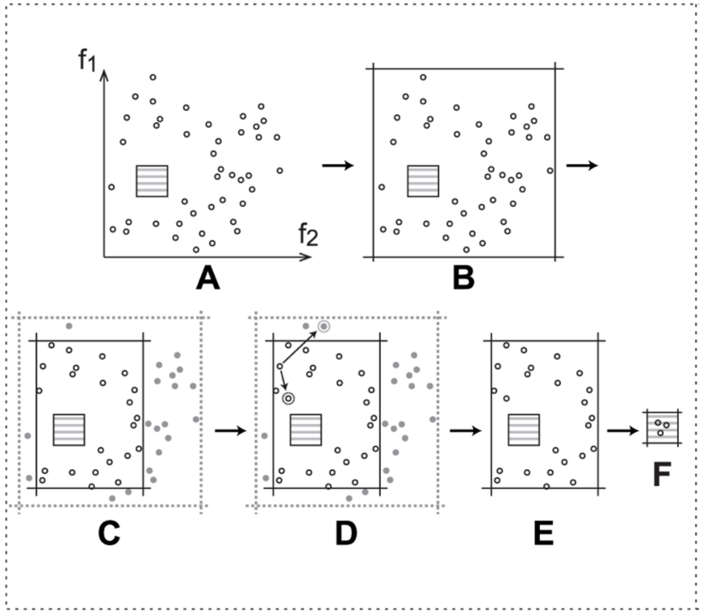

Open-ended evolution

- Premature convergence is caused by focusing on an objective.

Open-ended evolution redefines optimization as a never-ending process, rather than settling into a stable equilibrium.

- Relativism (e.g., peer-comparison)
- Artificial life (e.g., Cellular automata)
- Viability Evolution

Diversity does not guarantee quality!

MAP Elites

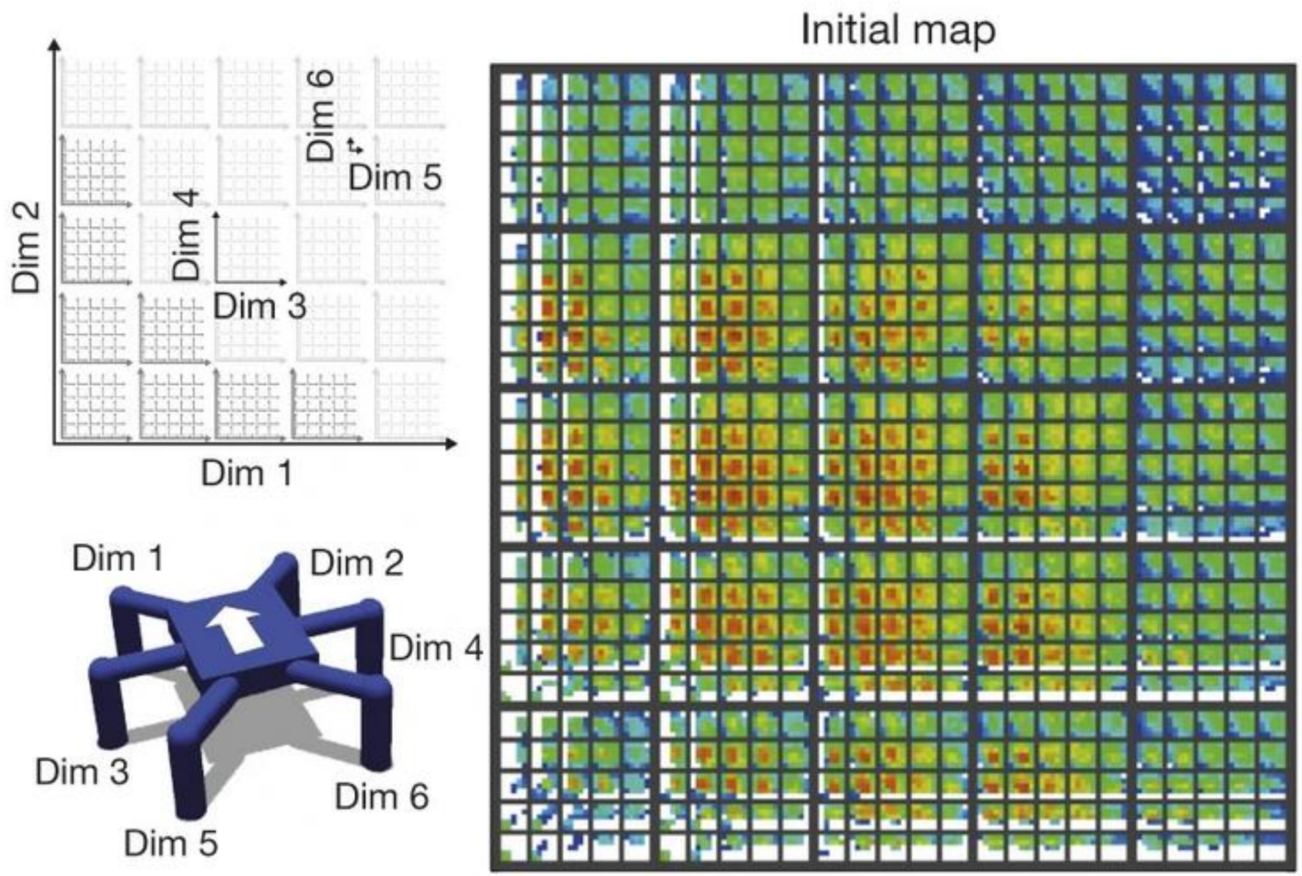
Initial map
Dim 2
Dim 1
Dim 6
Dim 5
Dim 4
Dim 3
Dim 1
Dim 2
Dim 4
Dim 3
Dim 6
Dim 5

Premature convergence in robot co-design

The co-design of control/body evolution results in different performances for similar designs, thereby worsening premature convergence

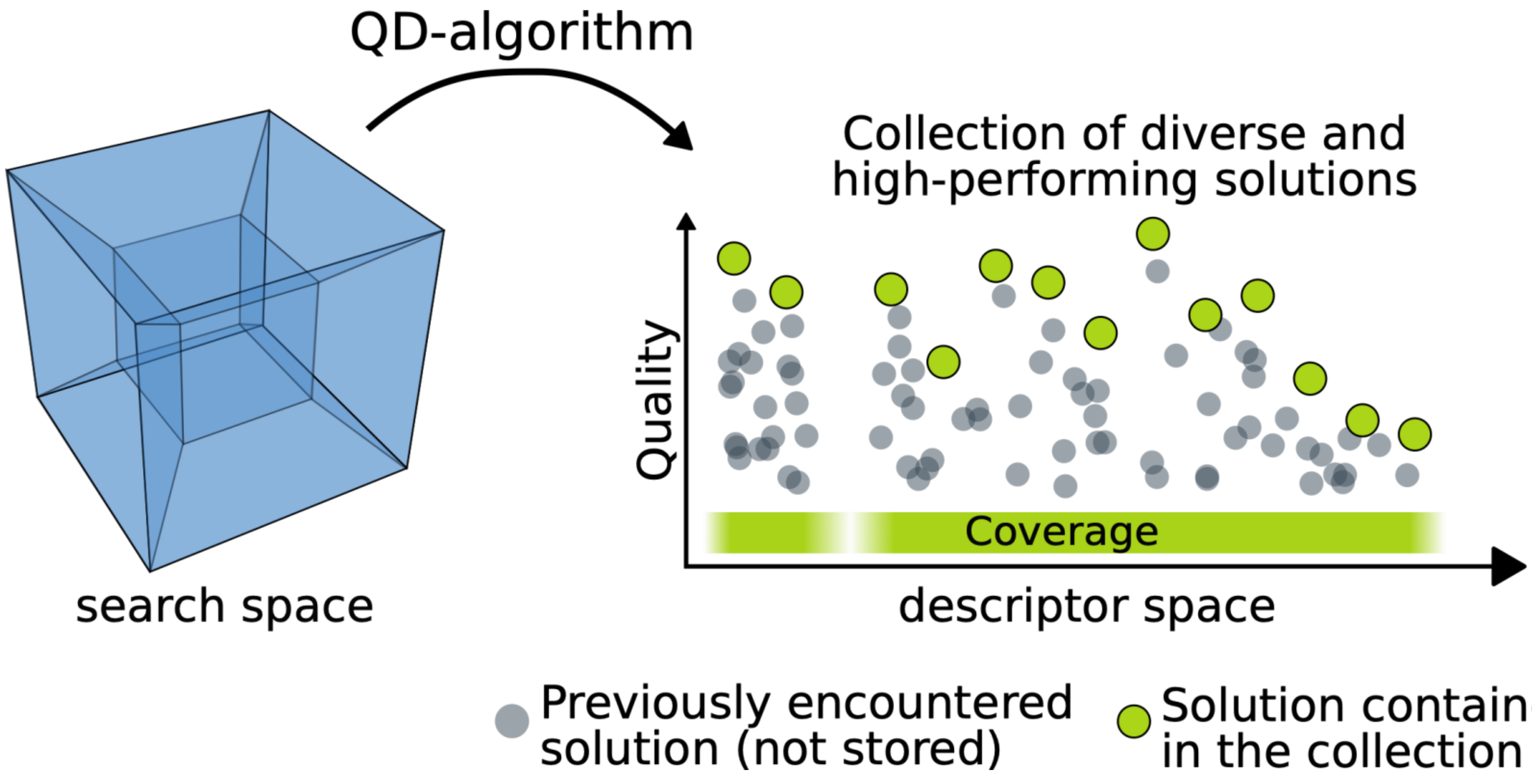

Mouret, J. B., & Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.

Quality Diversity algorithms

The co-design of control/body evolution results in different performances for similar designs, thereby worsening premature convergence

Quality-Diversity: Optimize both goal (quality) and phenotypical coverage (diversity)

Novelty search -> Maximize differences

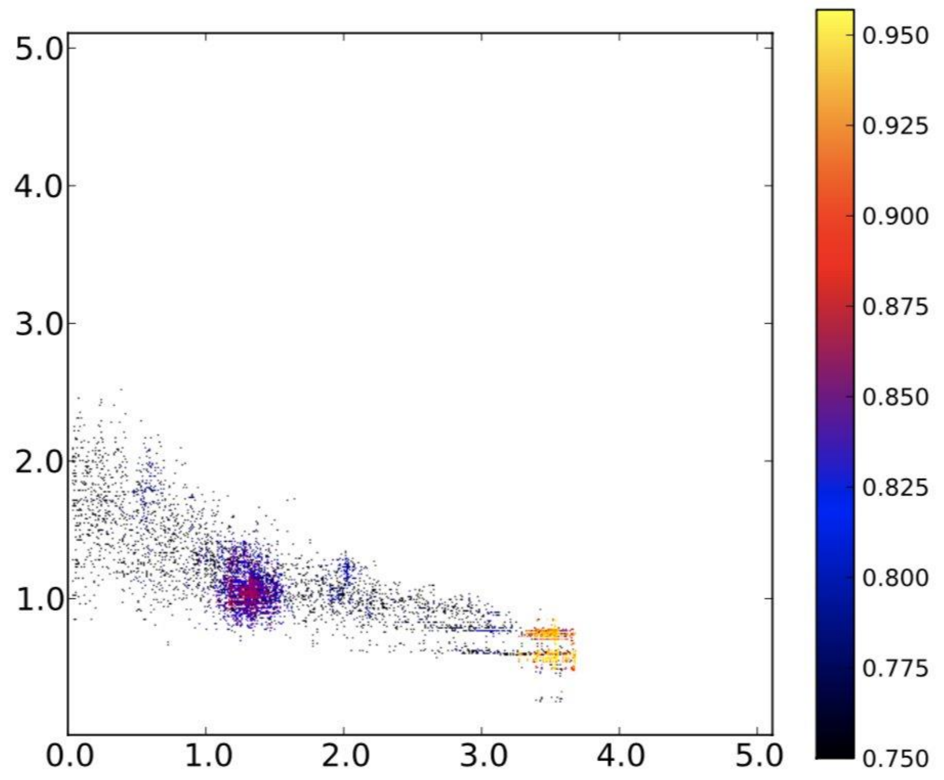

MAP-Elites -> Optimize within niche

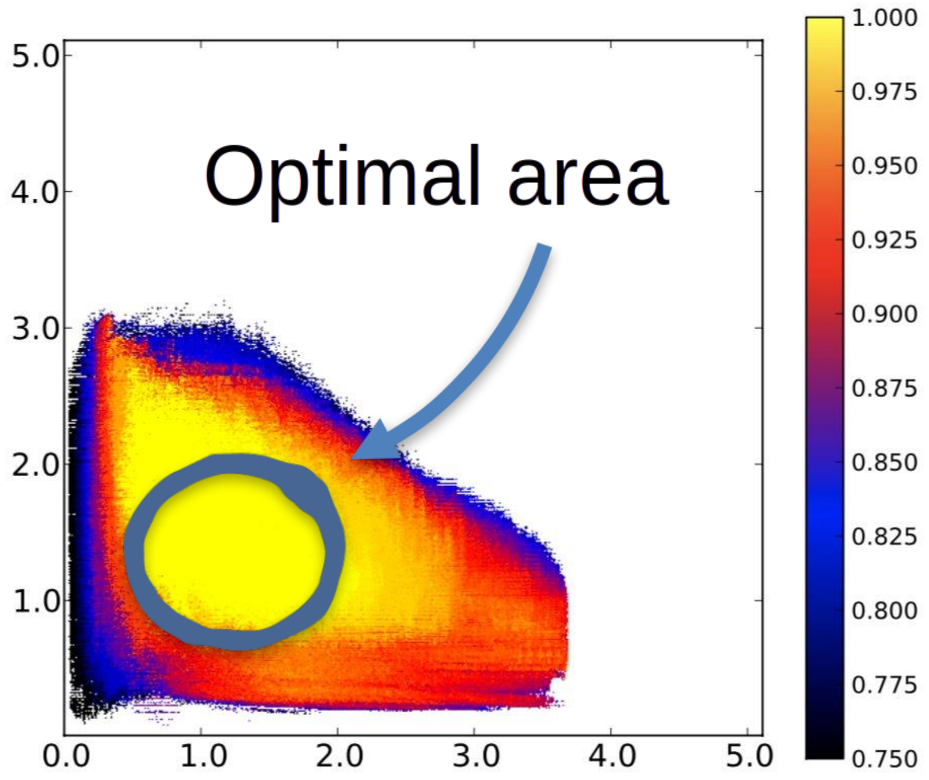

MAP Elites Archive

Define a phenotypical feature space (Archive)
Discretize feature space

Test individual fitness
Place individual in phenotype grid
- If individual fitness > current cell best; **replace**

procedure MAP-ELITES ALGORITHM (SIMPLE, DEFAULT VERSION)
    $(\mathcal{P} \leftarrow \emptyset, \mathcal{X} \leftarrow \emptyset)$
    for iter $= 1 \rightarrow I$ **do**
        if iter $< G$ **then**
            $\mathbf{x}' \leftarrow$ random_solution()
        else
            $\mathbf{x} \leftarrow$ random_selection($\mathcal{X}$)
            $\mathbf{x}' \leftarrow$ random_variation($\mathbf{x}$)
        $\mathbf{b}' \leftarrow$feature_descriptor($\mathbf{x}'$)
        $p' \leftarrow$performance($\mathbf{x}'$)
        if $\mathcal{P}(\mathbf{b}') = \emptyset$ or $\mathcal{P}(\mathbf{b}') < p'$ **then**
            $\mathcal{P}(\mathbf{b}') \leftarrow p'$
            $\mathcal{X}(\mathbf{b}') \leftarrow \mathbf{x}'$
    return feature-performance map ($\mathcal{P}$ and $\mathcal{X}$)

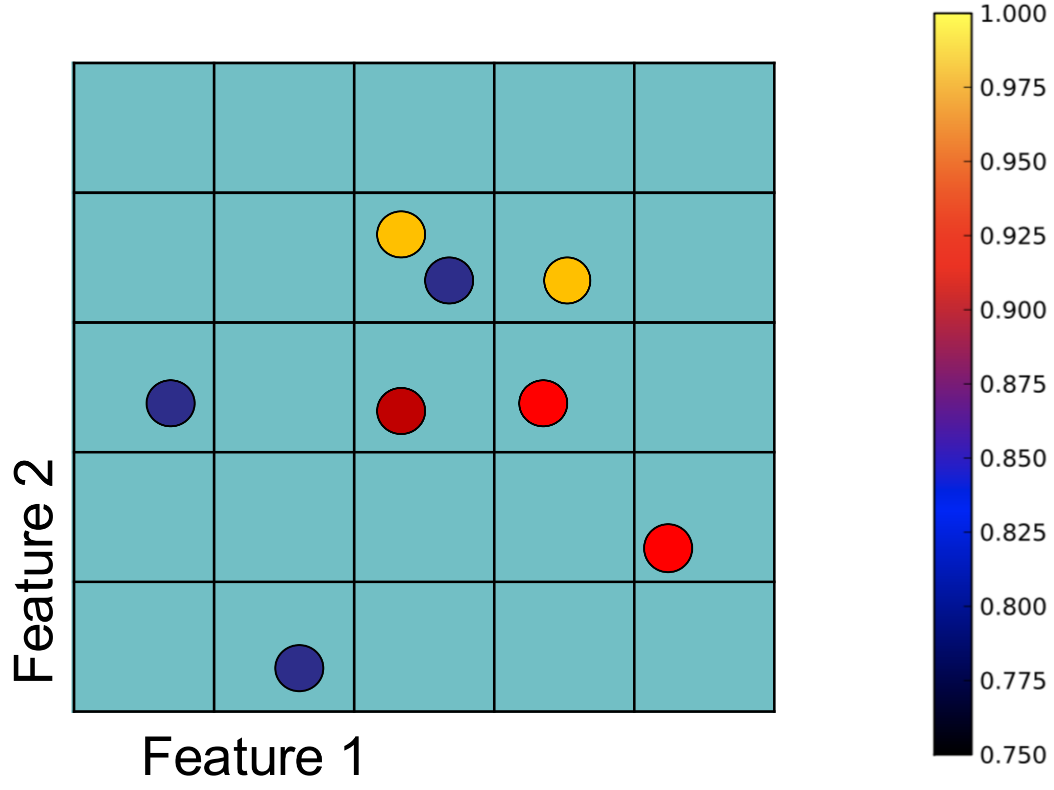

Mouret, J. B., & Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.

Niche construction in the archive

- MAP-Elites forces similar solutions to compete only with each other.
- Each map cell preserves a local elite, thus allowing for Phenotypical diversity

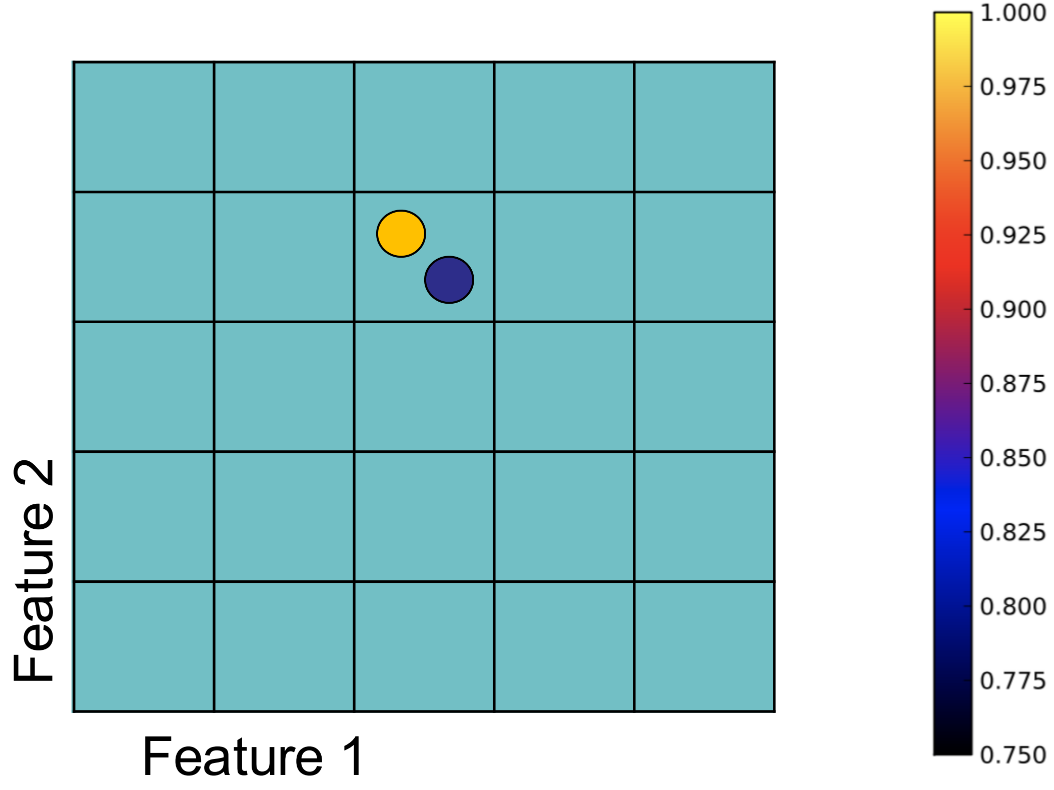

Mouret, J. B., & Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.

MAP Elites: soft-voxel robots

Example:

Evolve soft voxel robots walk as far as possible
Voxels: {empty, bone, soft, contract, expand}

Phenotype space:
- F1: Voxels not empty [%]
- F2: Voxels = bones [%]

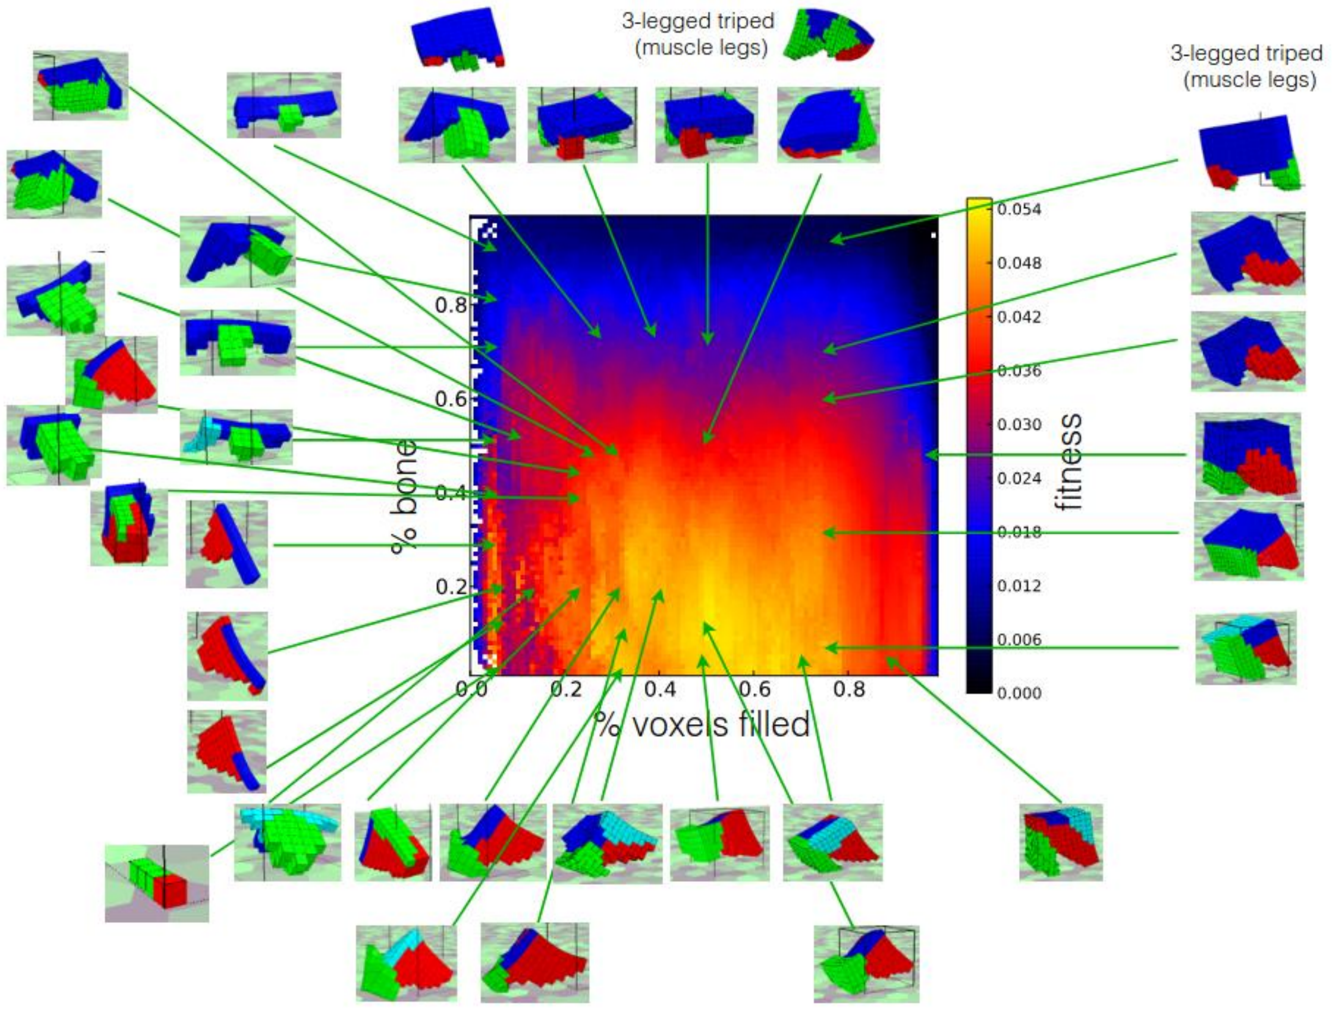

Mouret, J. B., & Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.

MAP Elites: 6-legged robot

Example:

Evolve 6-legged robot walk as far as possible 5s
Parameterized controller

Phenotype space Contact time per leg:

- F1: Dim 1 (leg 1)
- F2: Dim 2
- …
- F6: Dim 6

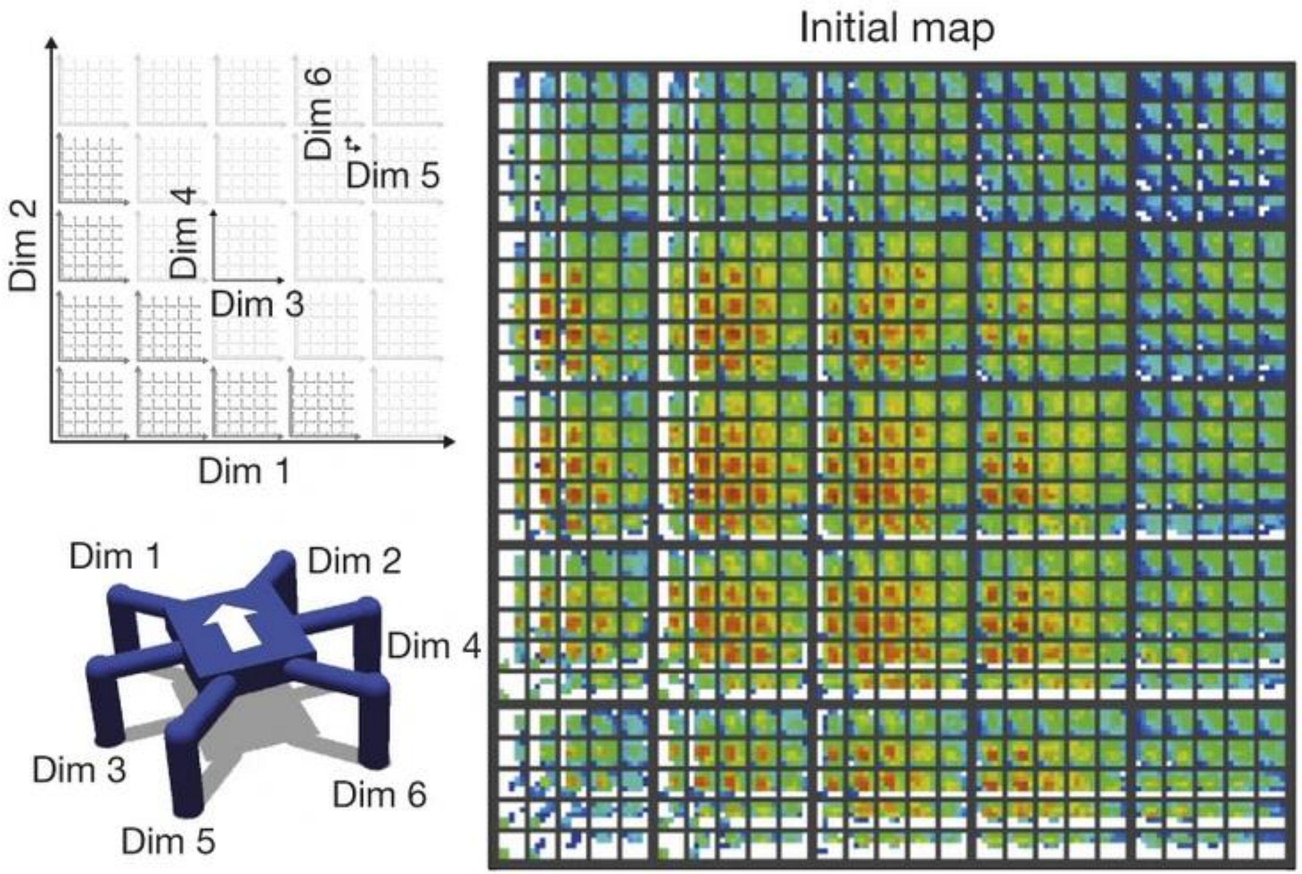

Cully, A., Clune, J., Tarapore, D., & Mouret, J. B. (2015). Robots that can adapt like animals. *Nature*, *521*(7553), 503-507.

Understanding the feature space

Analyzing feature space (time of ground contact)

E.G. right bottom robot
Dim 1 & Dim 2: low

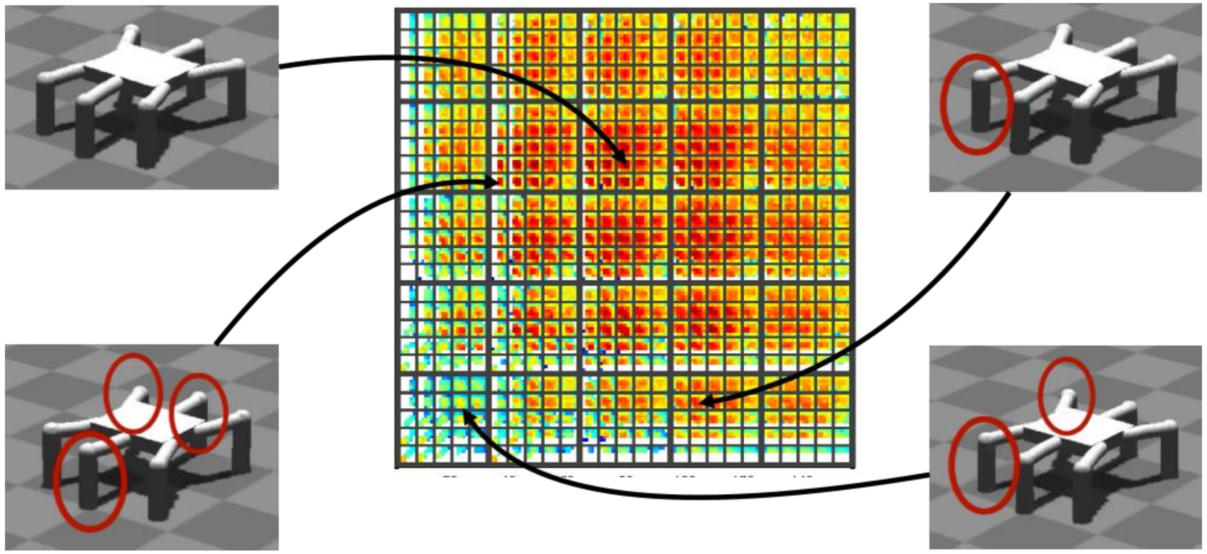

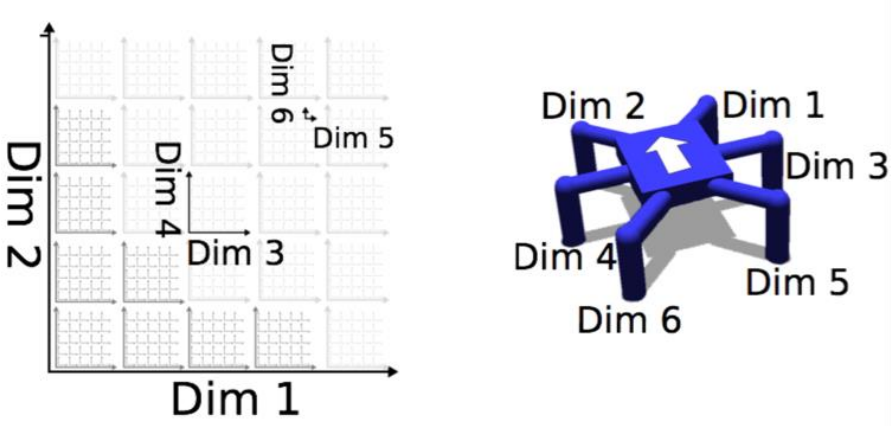

Cully, A., Clune, J., Tarapore, D., & Mouret, J. B. (2015). Robots that can adapt like animals. *Nature*, *521*(7553), 503-507.

Adaptation through archival motor primitives

Learning a wide range of behaviors allows for quick adaptation

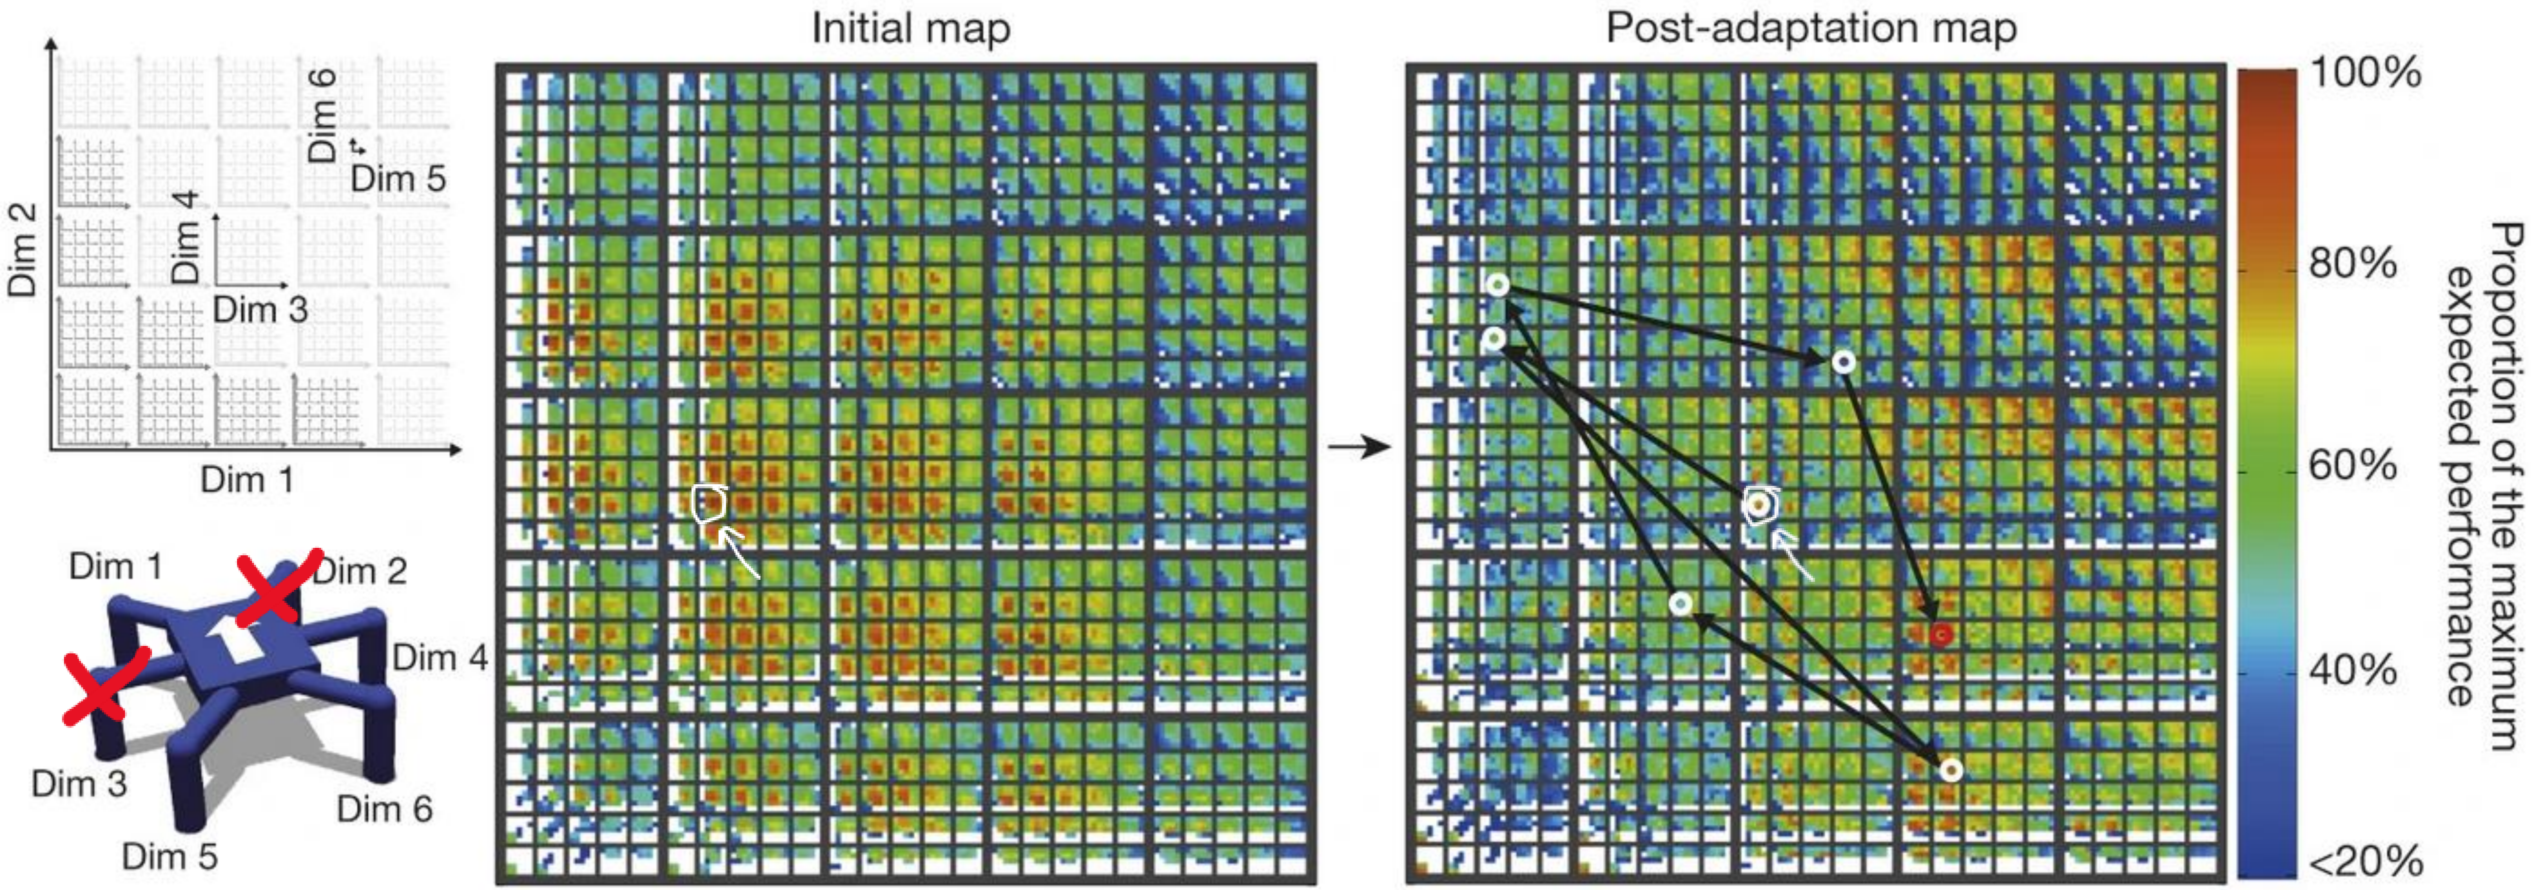

Cully, A., Clune, J., Tarapore, D., & Mouret, J. B. (2015). Robots that can adapt like animals. *Nature*, *521*(7553), 503-507.

Real world test

Learning a wide range of behaviors allows for quick adaptation

Cully, A., Clune, J., Tarapore, D., & Mouret, J. B. (2015). Robots that can adapt like animals. *Nature*, *521*(7553), 503-507.

MAP-Elites: key takeaways

- MAP-elites Optimize both goal (quality) and phenotypical coverage (diversity)
- Archive is a phenotype-performance space
- Unlike novelty-search MAP-elites do not reject similar phenotypes; but replaces individuals when performance improves (i.e. local competition/ niches)

Checkpoints

- Describe premature convergence
- Describe the differences between genotype/phenotype diversity
- What is open-ended evolution and how does it differ from goal-based evolution
- Describe the following Quality-Diversity algorithms:
  - Novelty search
  - MAP-Elites