Quality diversity

From Novelty Search to the MAP-Elites algorithm

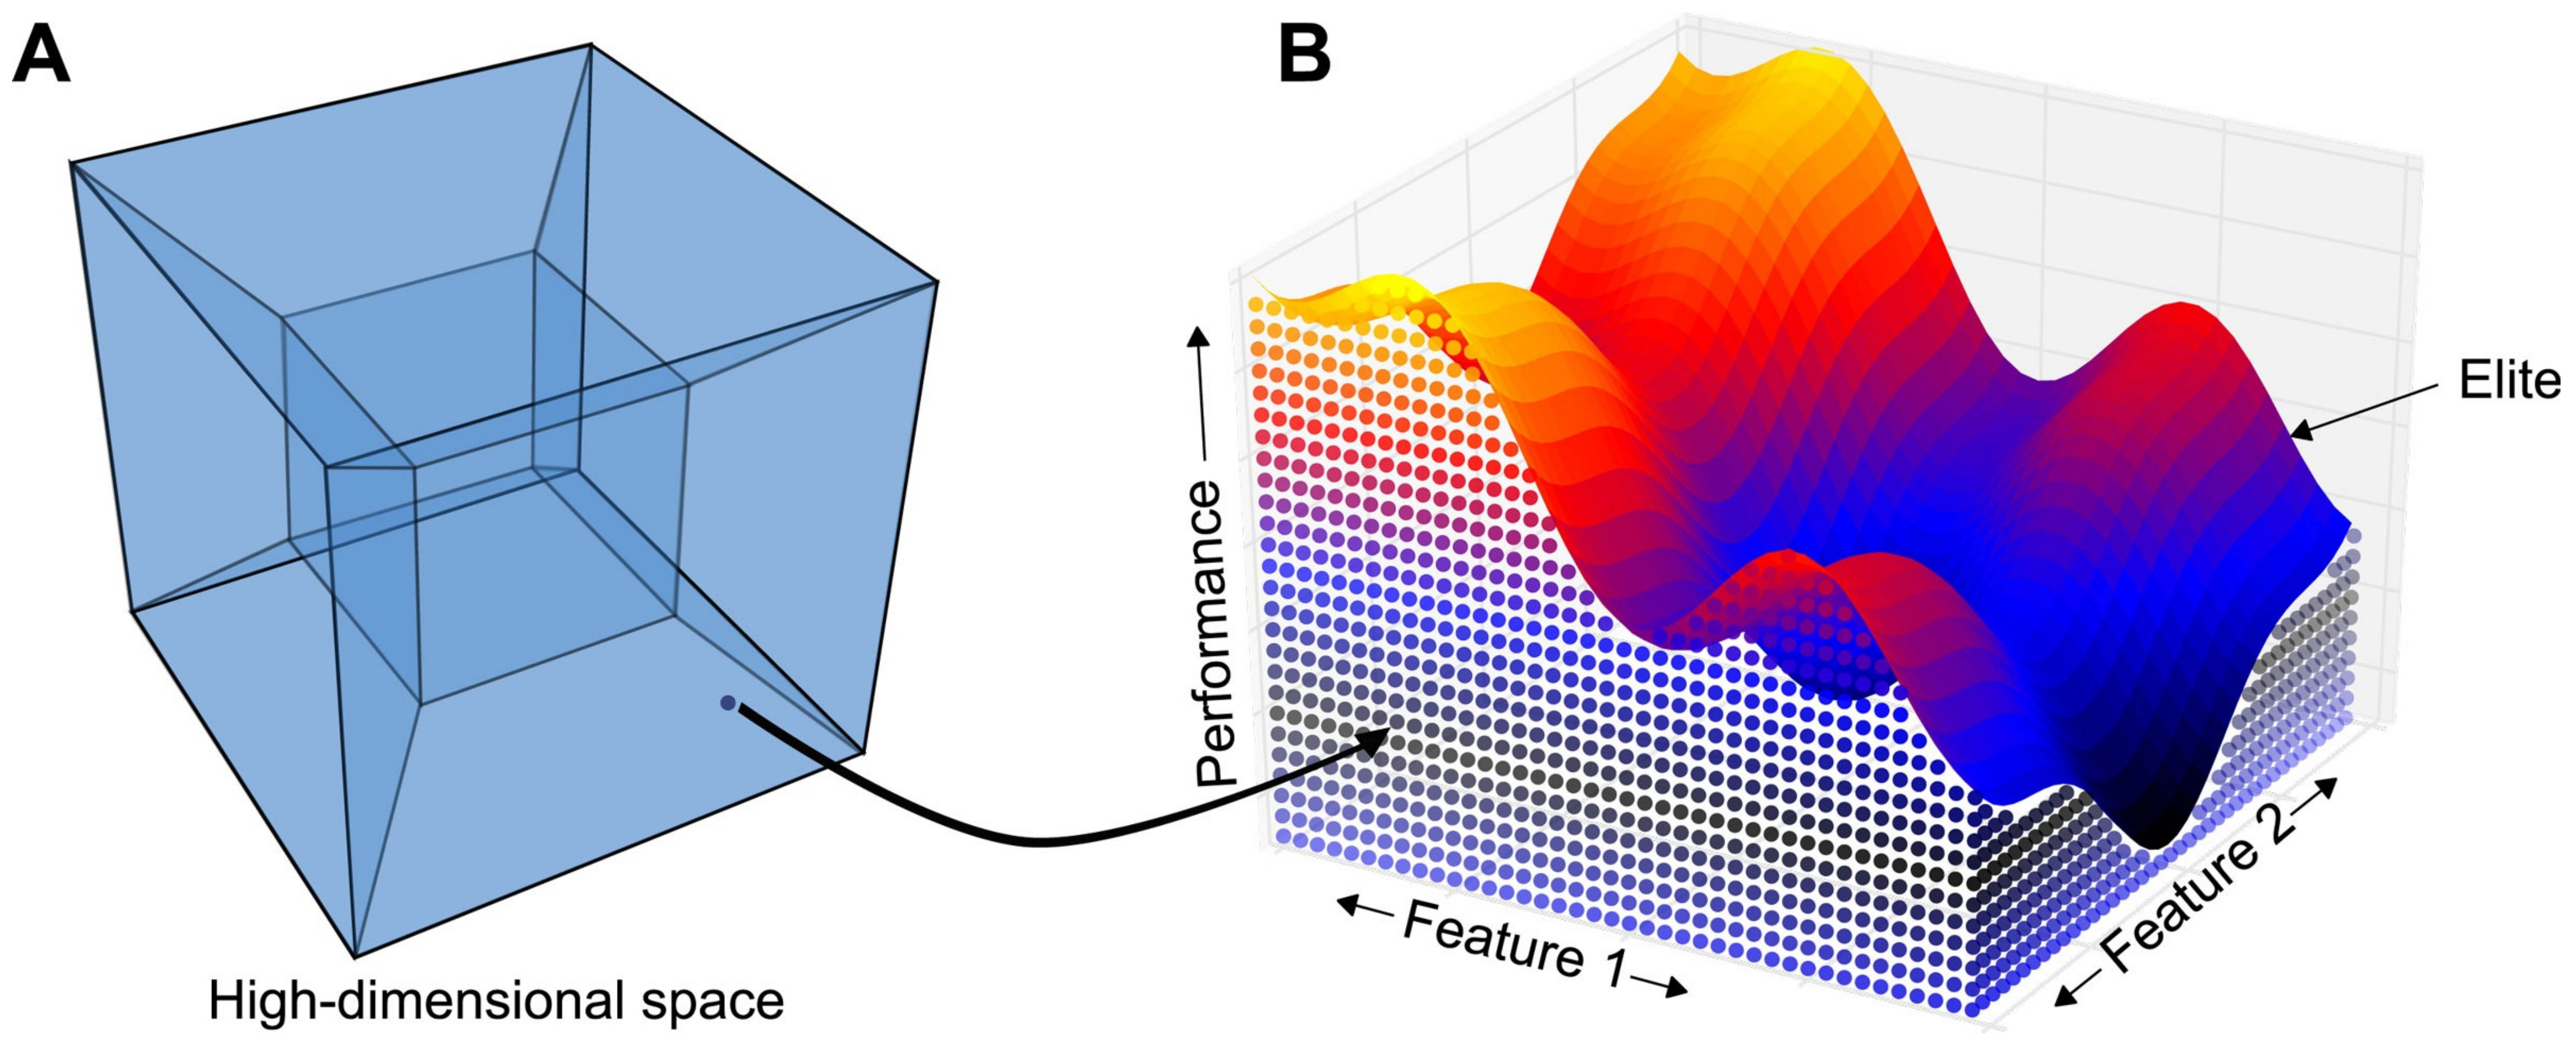

Jean-Baptiste Mouret — Inria

@jbmouret

@jb_mouret

The problem with artificial evolution

Where is the creativity?

- ES (e.g., CMA-ES) are good (global) black-box optimizers, inc. for RL problems
- We can evolve the weights of deep neural networks for RL (OpenAI-ES, etc.)
- MOEA (e.g., NSGA-II) are good multi-objective (black-box) optimizers, inc. for RL problems
- We can evolve the structure of neural networks (e.g., NEAT & HyperNEAT)

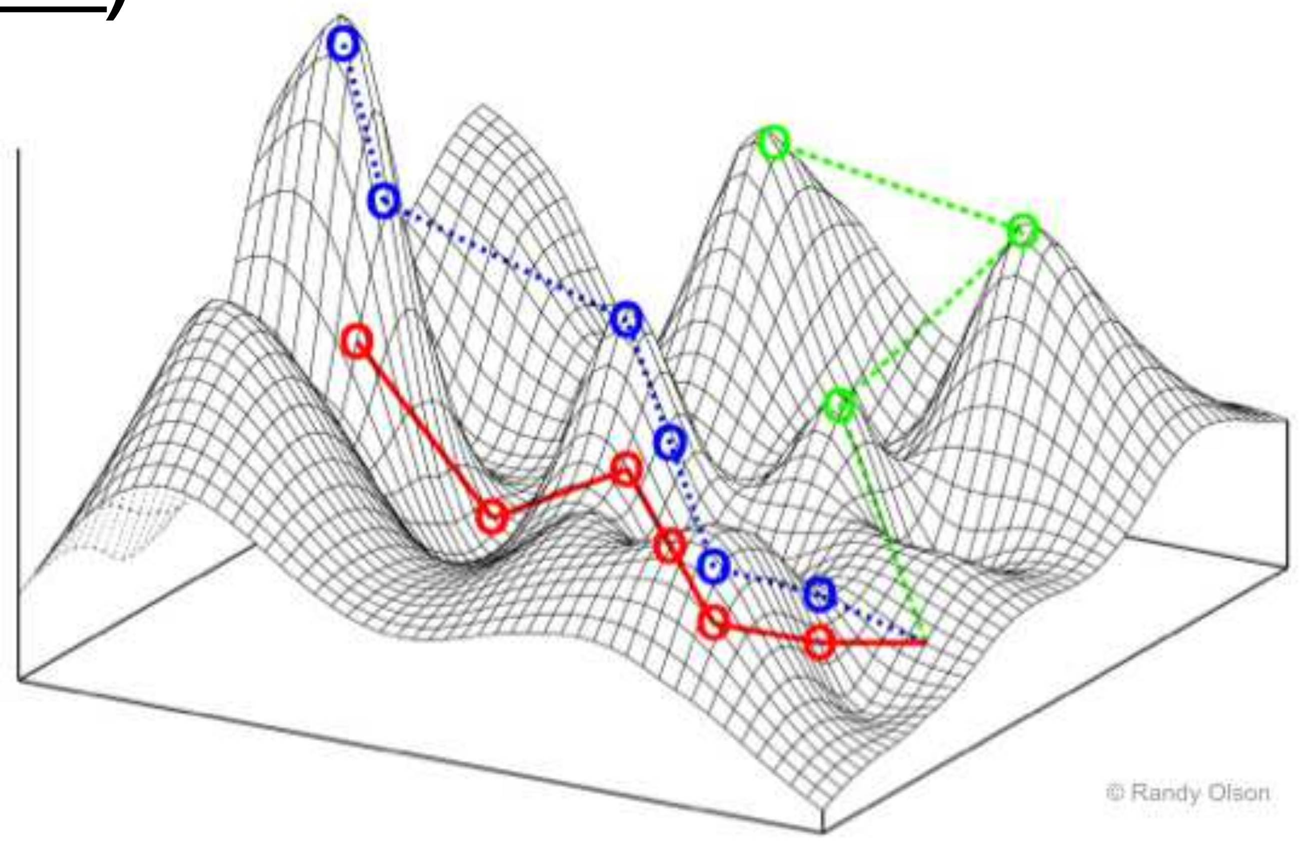

- **but…**
  - a lot of "tuning" and "fitness shaping" for the "success stories"
  - (yes, this is not highlighted in the videos / papers)

➔ **Where is open-ended, creative evolution?**

➔ **What is missing? how to do better?**





The Picbreeder Experiment

Interactive evolution

- Collaborative evolution of images (online)
- inspired by Dawkins' Biomorph
- No goal
- Encoding of images using CPPNs (see neuroevolution part)

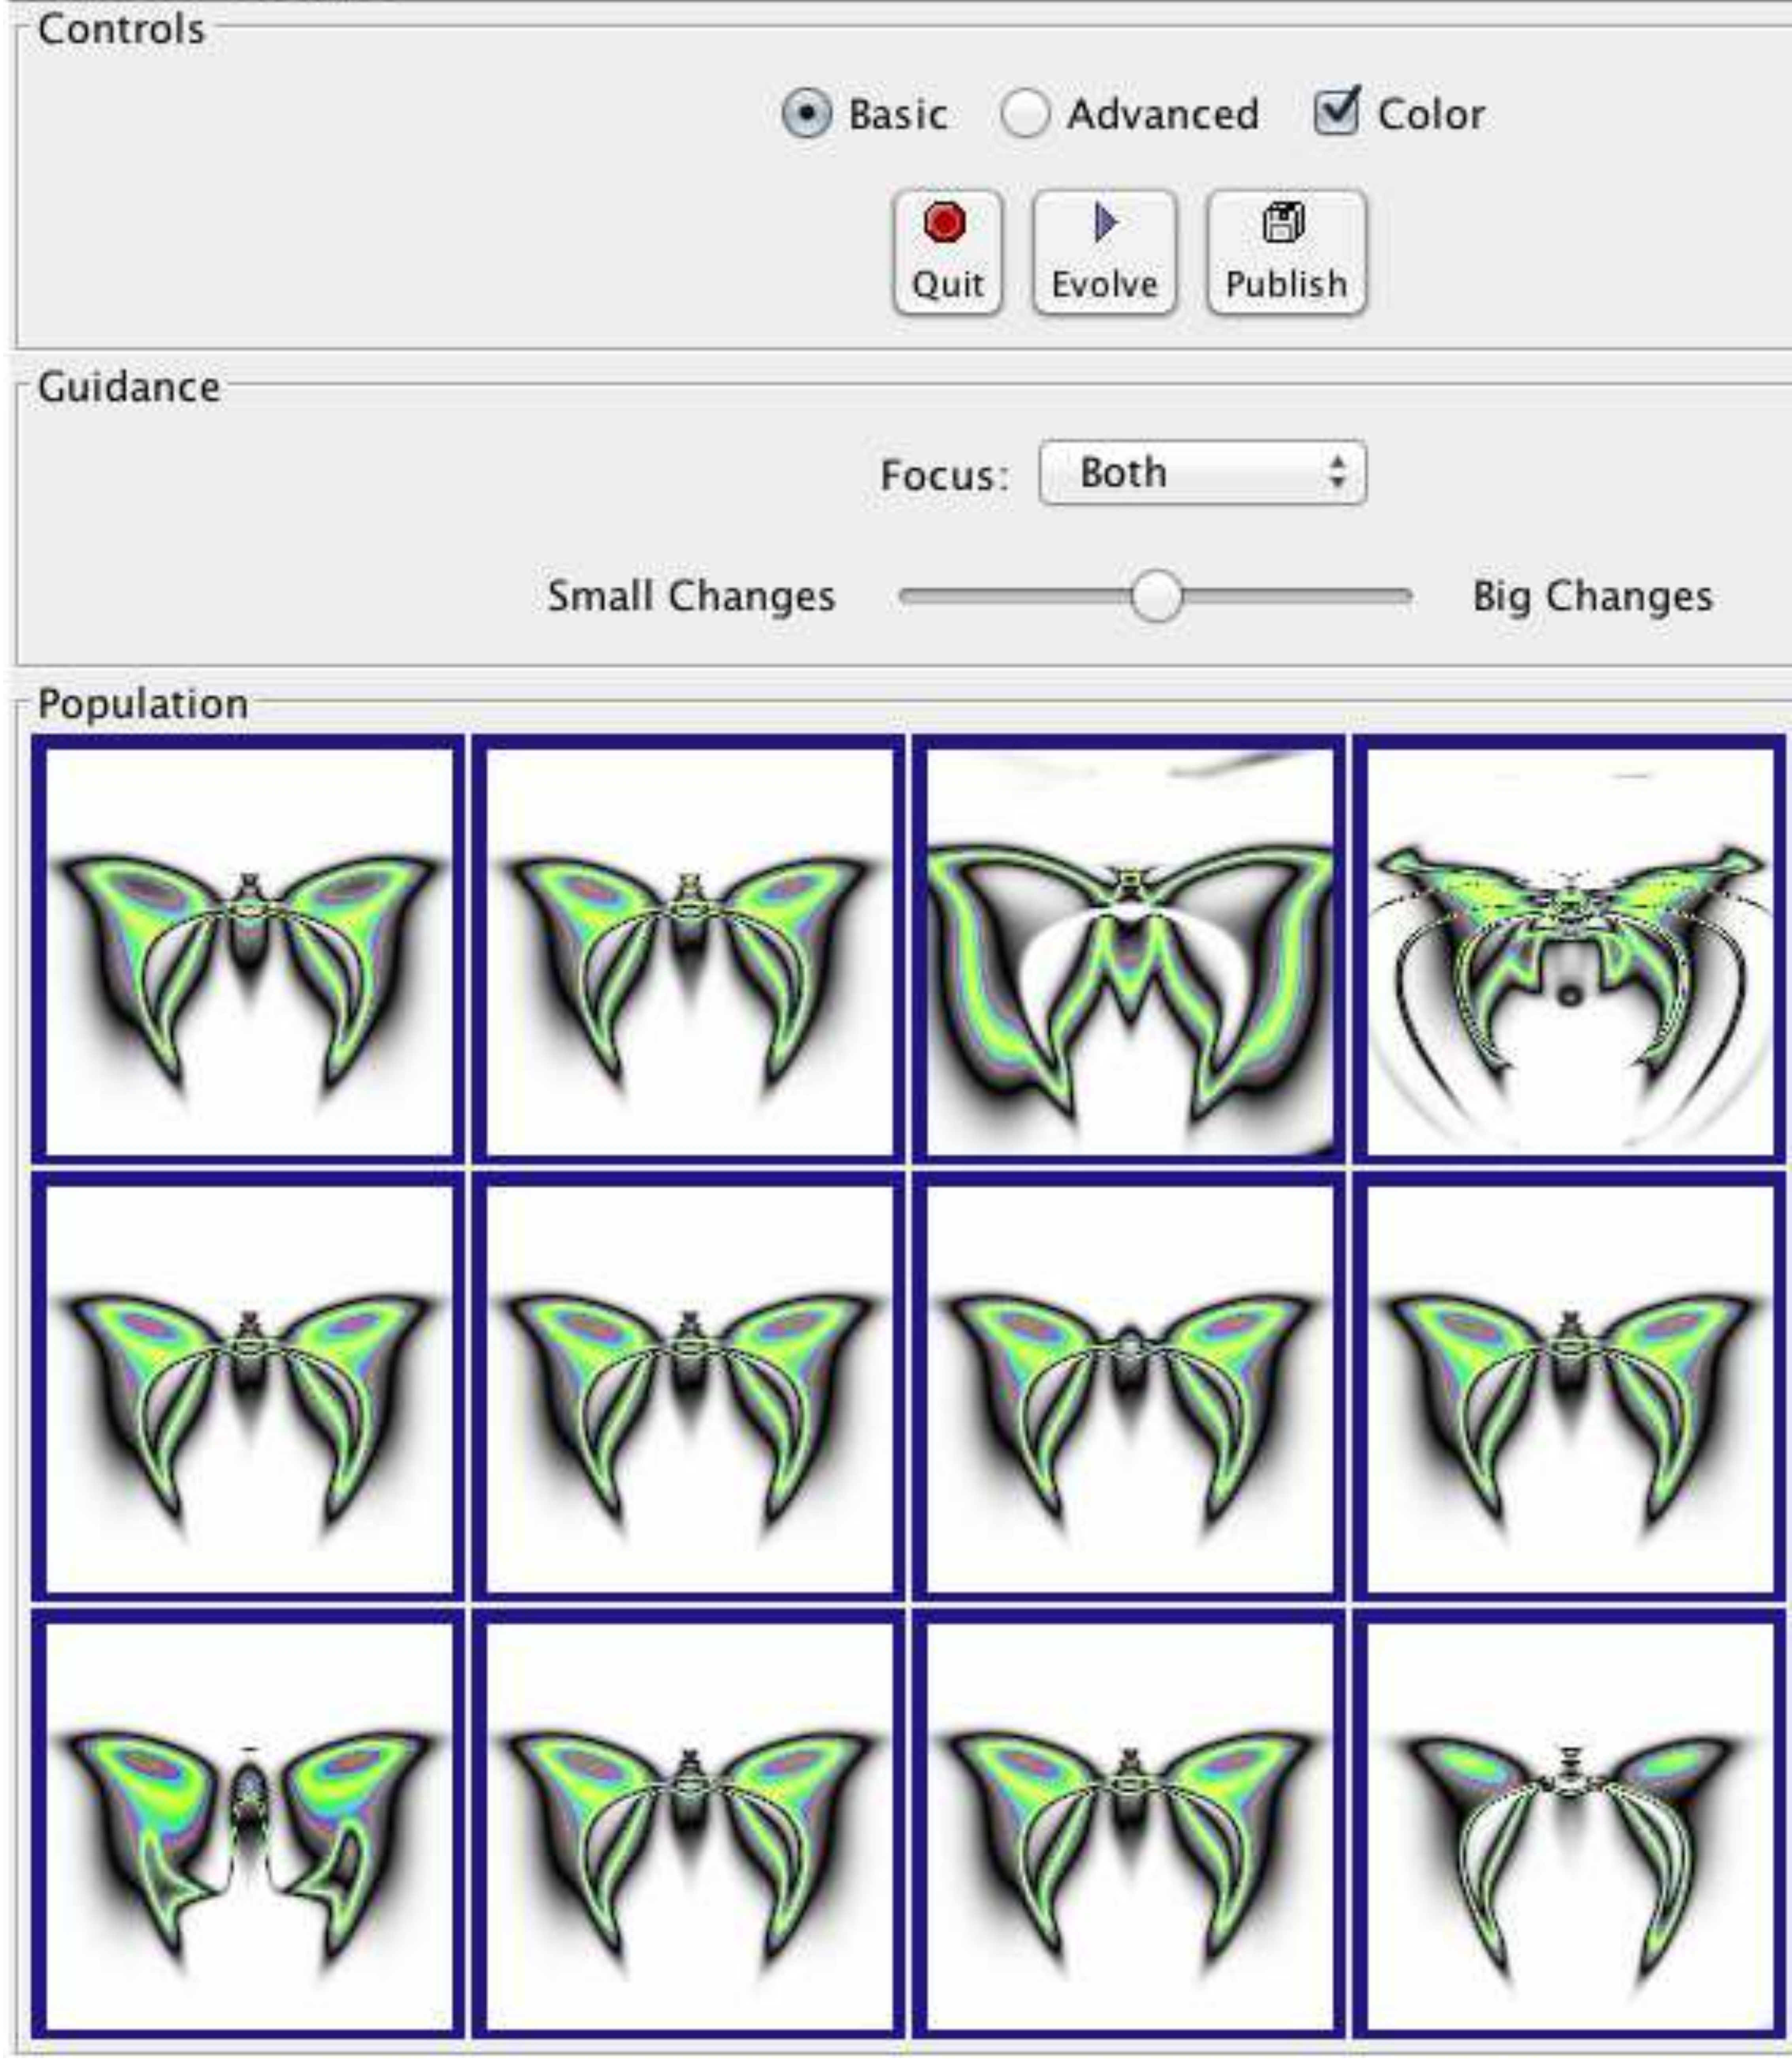

K. O. Stanley, J. Lehman (2015) - Why Greatness Cannot Be Planned,- Springer

Secretan J, Beato N, D'Ambrosio DB, Rodriguez A, Campbell A, Folsom-Kovarik JT, Stanley KO. Picbreeder: A case study in collaborative evolutionary exploration of design space. Evolutionary computation. 2011 Sep;19(3):373-403.

Results from Picbreeder (interactive evolution)

source: http://www.picbreeder.org

Interactive evolution (no goal):

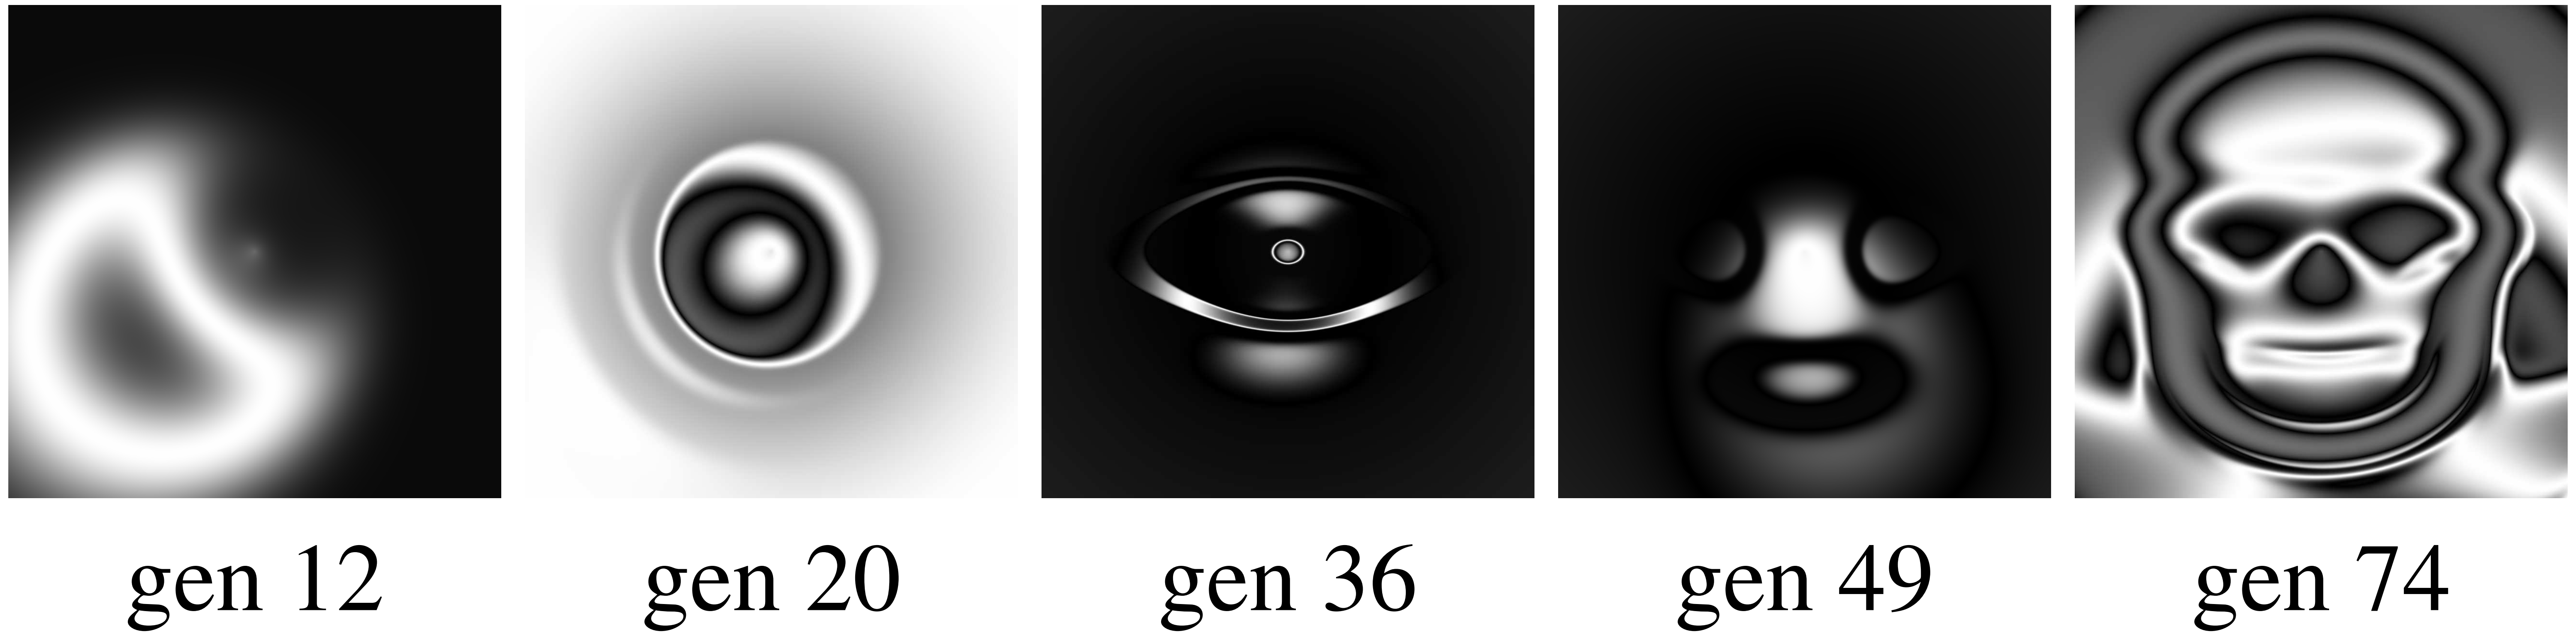

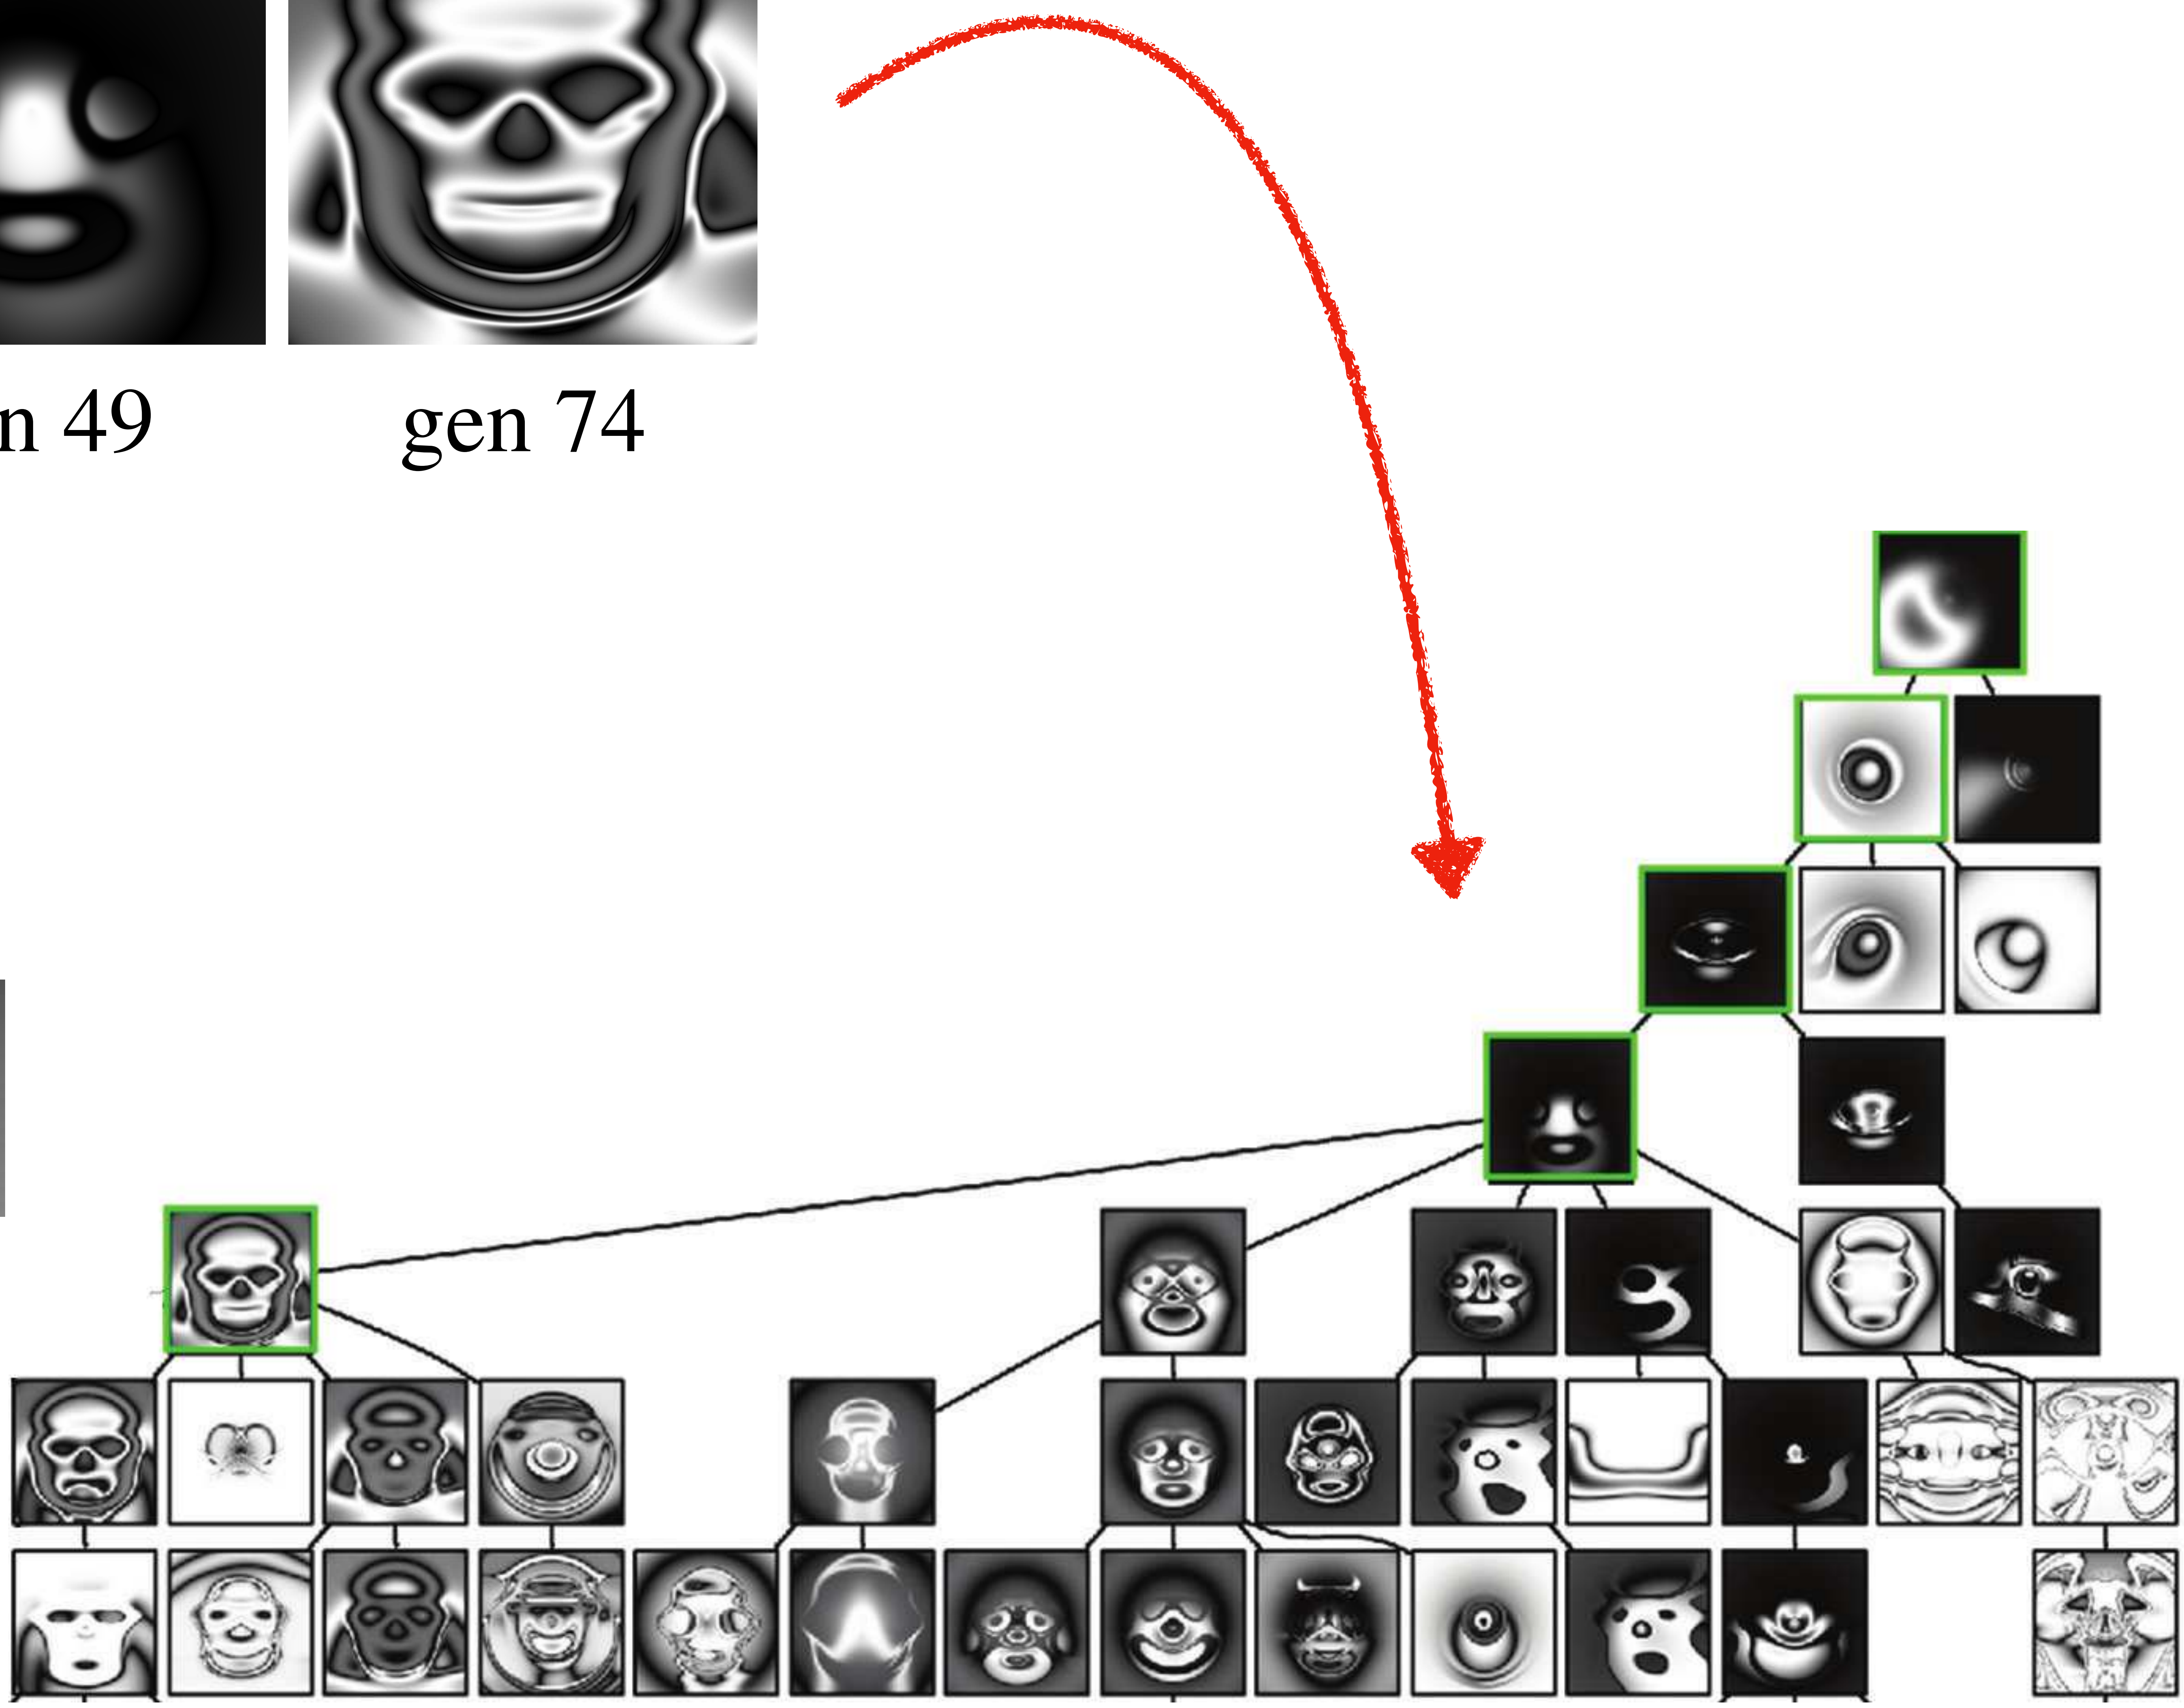

Objective-based evolution

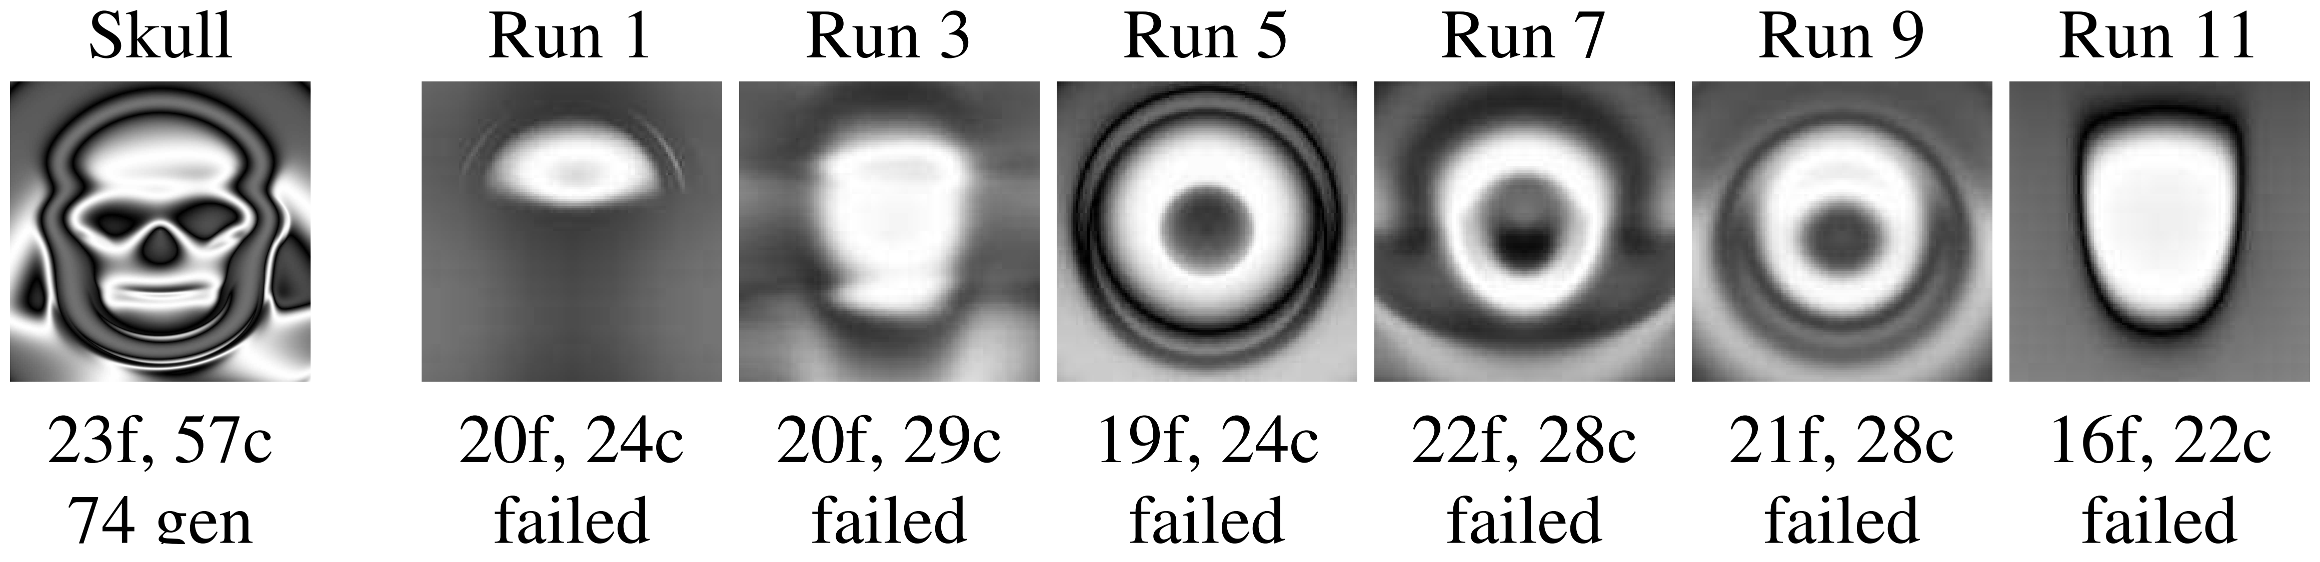

Woolley, B. G., & Stanley, K. O. (2011). On the deleterious effects of a priori objectives on evolution and representation. In Proceedings of the 13th annual conference on Genetic and evolutionary computation (pp. 957-964). ACM.

Deceptive search spaces

- The "stepping stones" to get to the solution are <u>NOT</u> like the final solution

… but this is the "basic" heuristic of most search algorithms: solutions that are closer (better reward) to the goal should be favored!

- Interpretation: the search space "<u>deceptive</u>" (attractive local minima)

➔We need more exploration

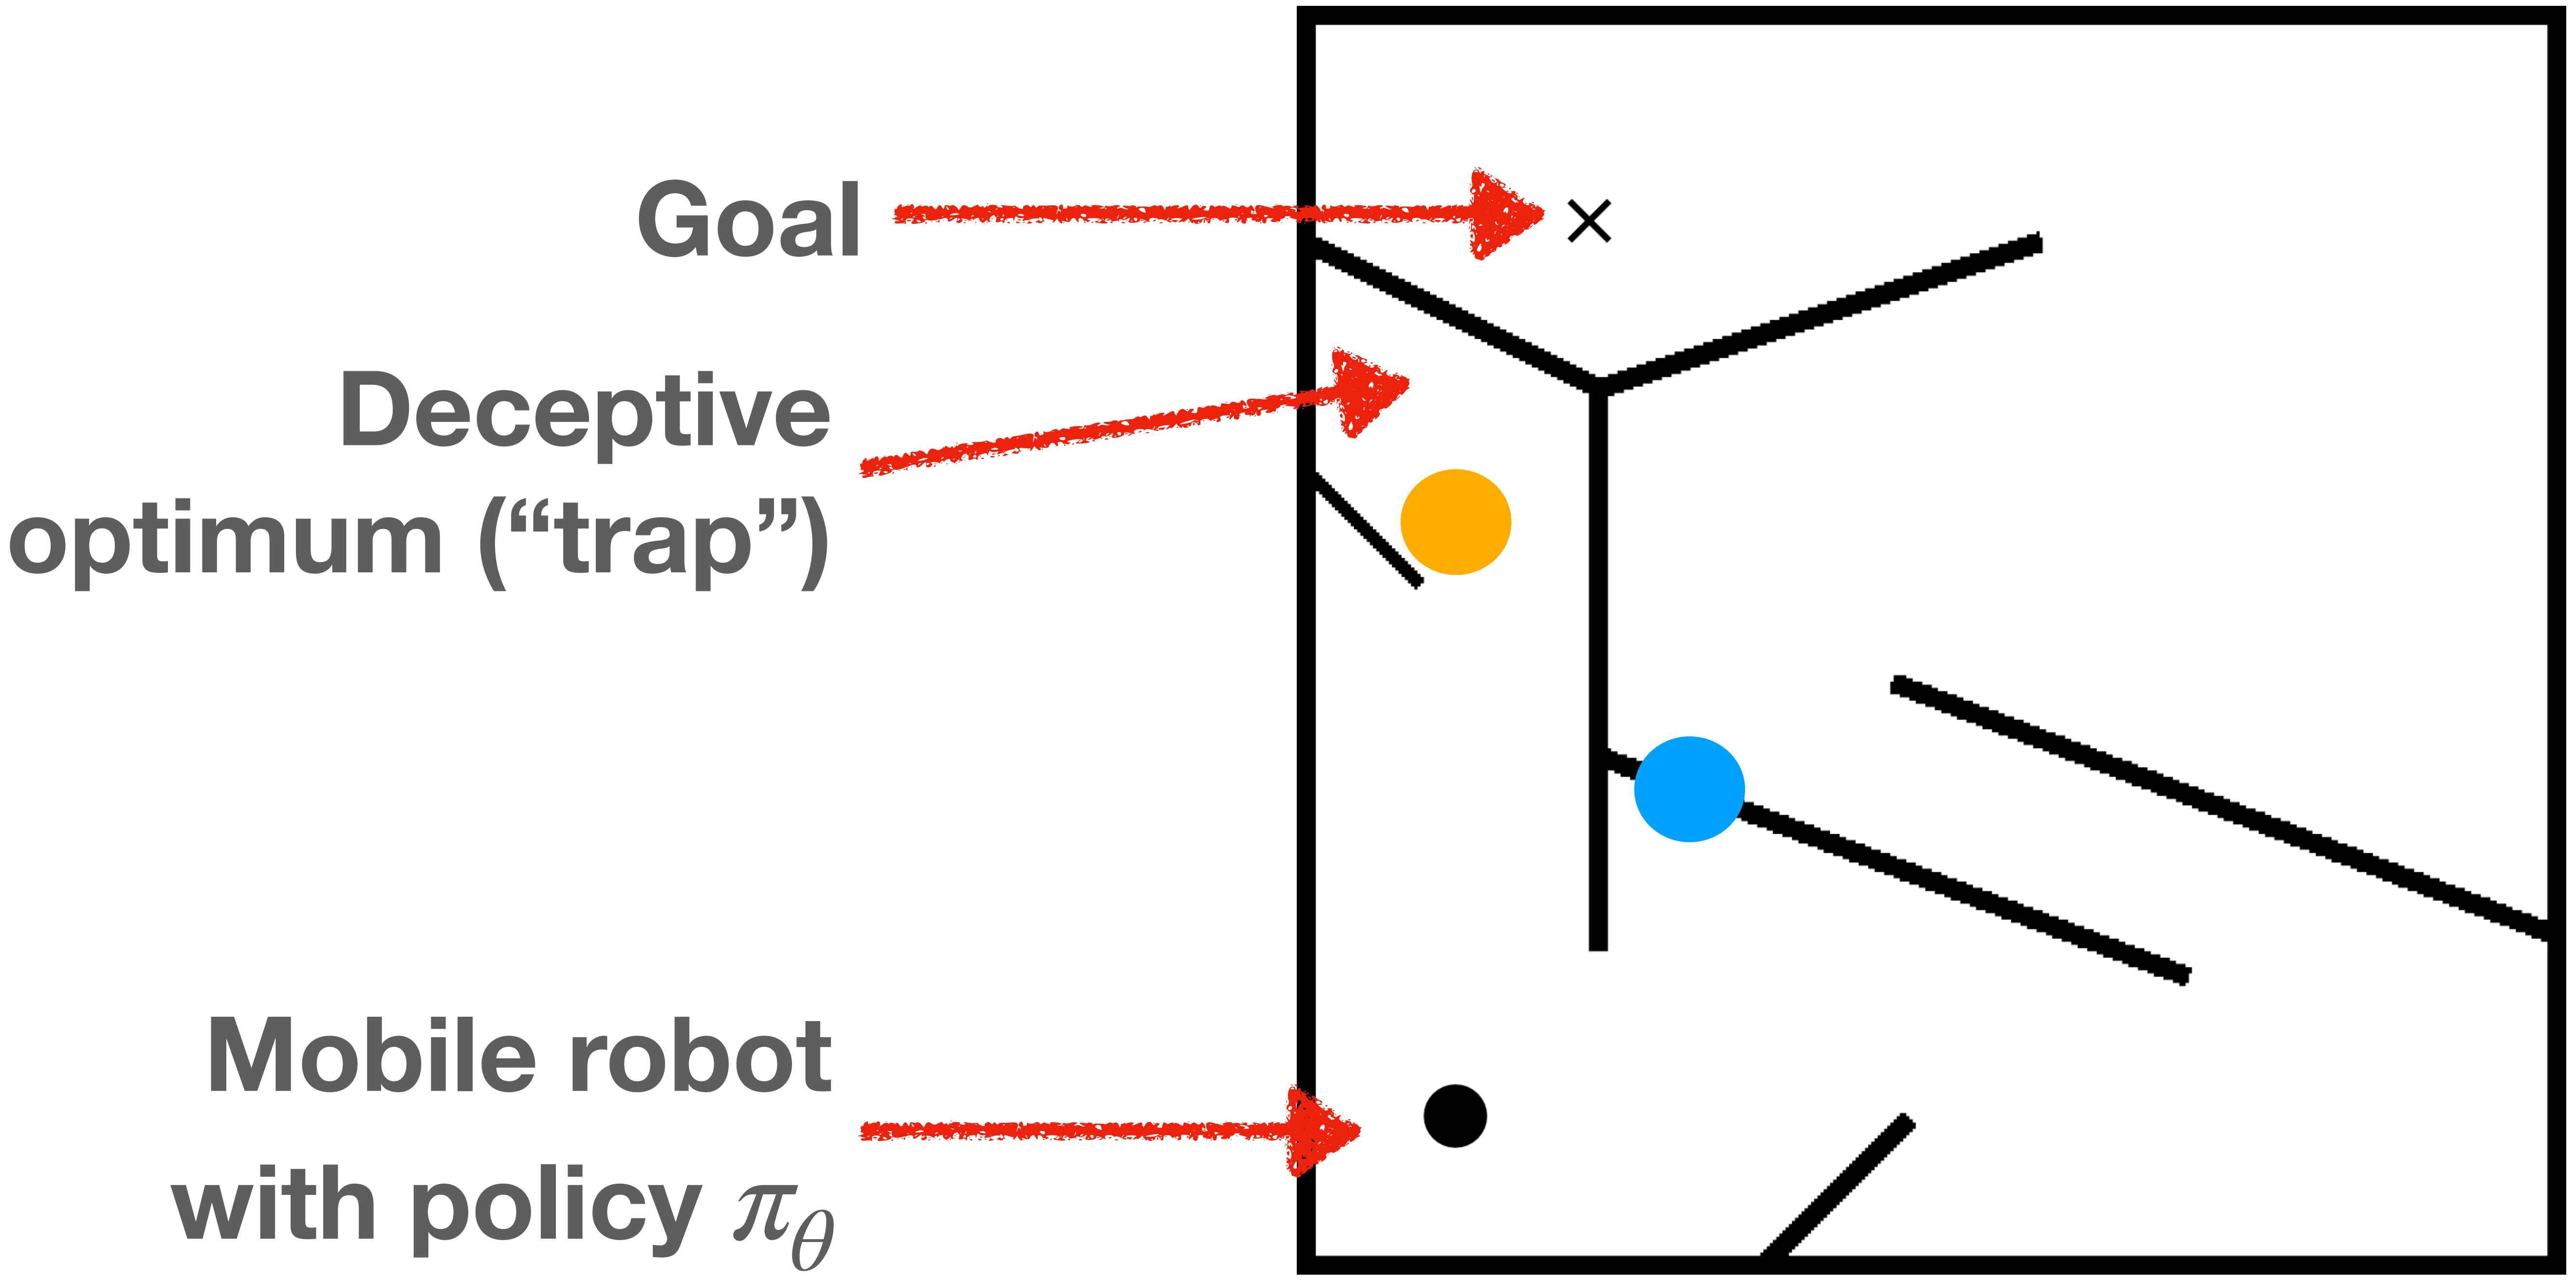

Fitness = distance to the goal at the end of the episode

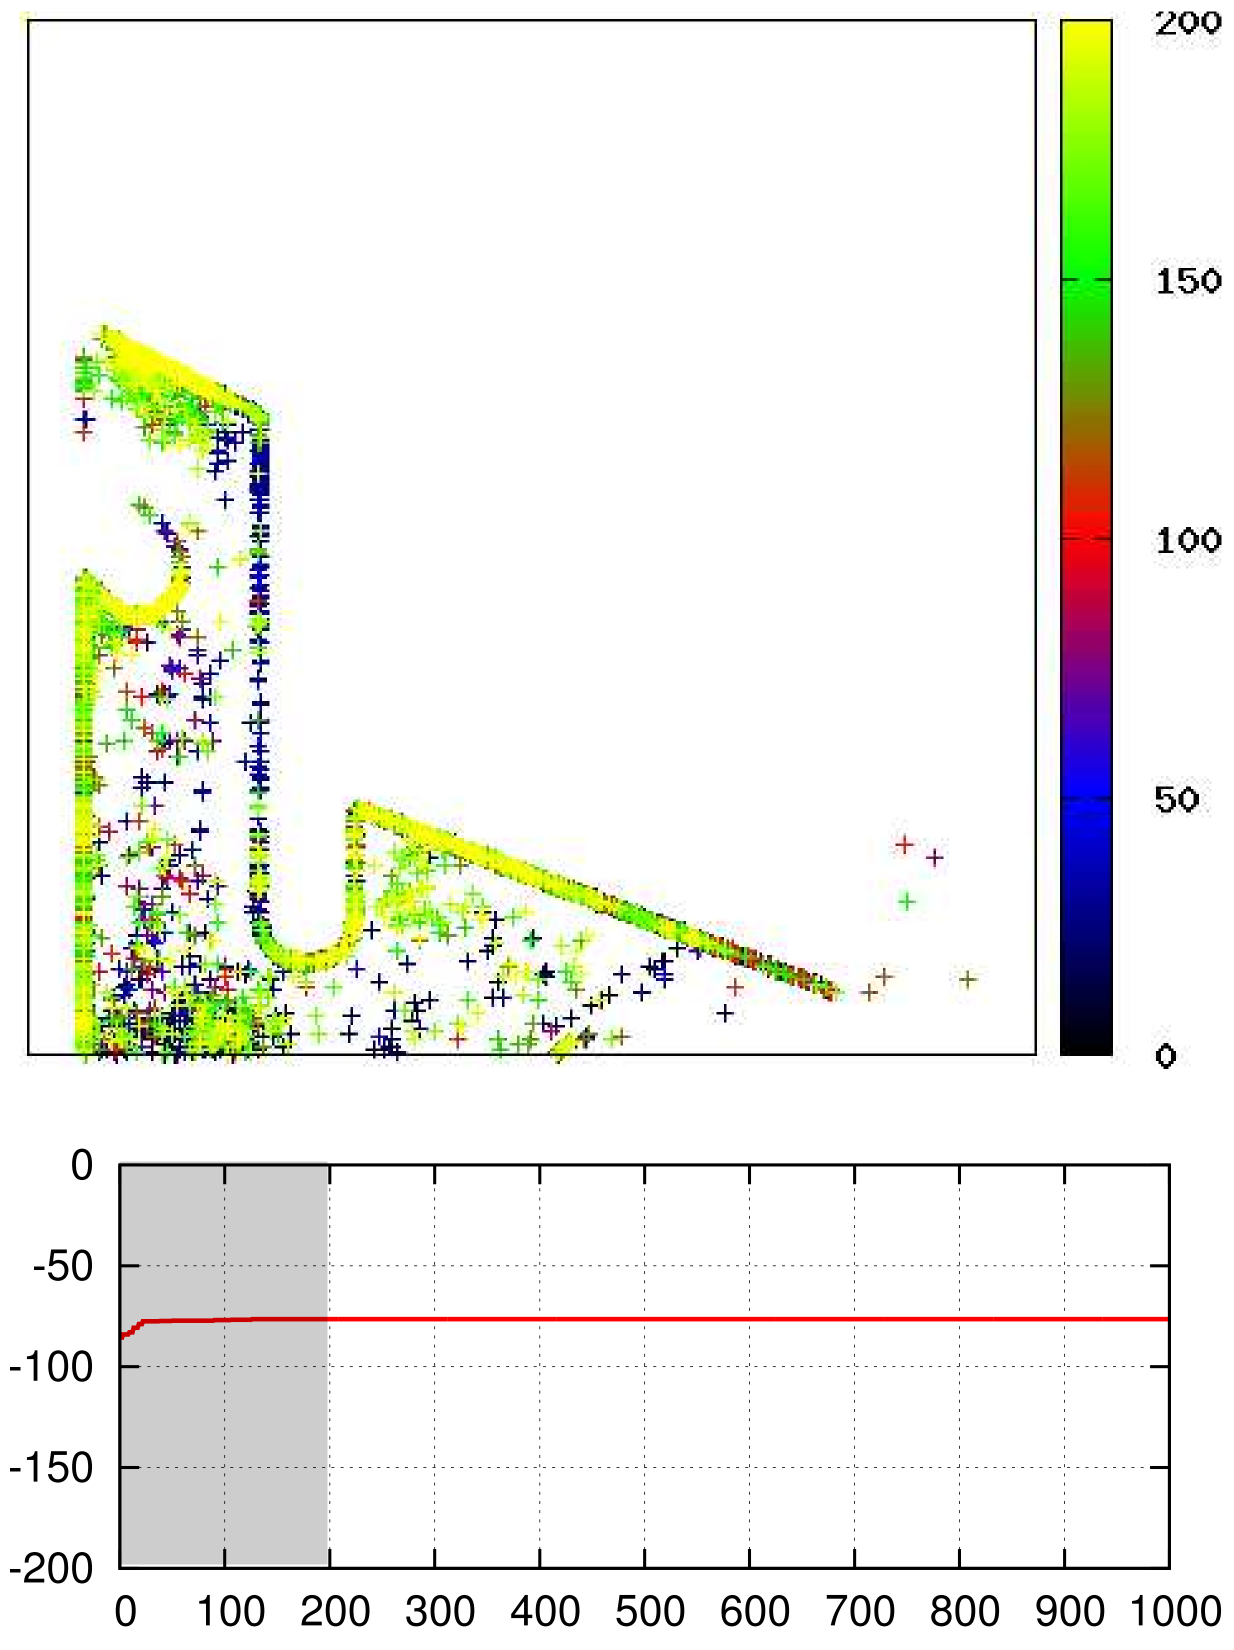

(a) Fitness only (single obj.)

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

Novelty Search concept

- Radical idea: what if we ignore the fitness function?

… and search for novel "things"

- In RL/Evo, "things" = behavior

➔ search for novel <u>behaviors</u>

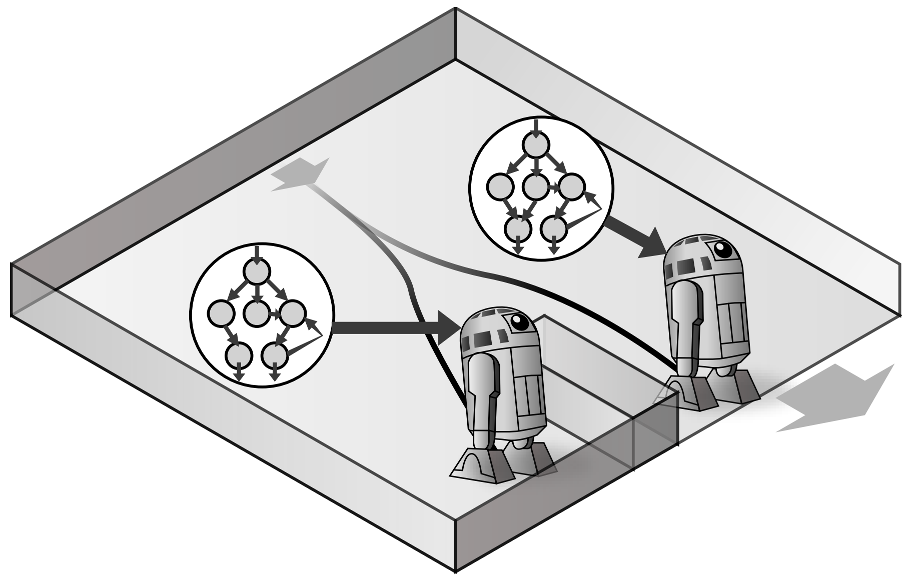

Close genotype / different behavior

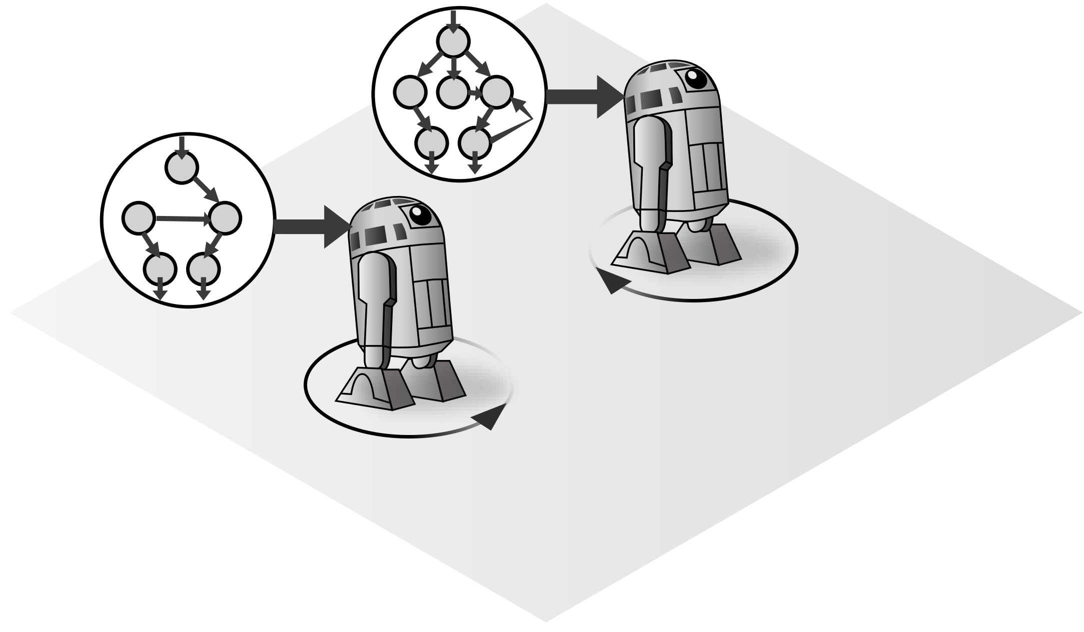

Different genotype / same behavior

Novelty search:

➔ characterize behavior (e.g., final position, not list) with a vector

➔ replace fitness by novelty

… computed by the behavioral distance to the archive & population

$$\rho(x) = \frac{1}{k}\sum_{j=0}^{k} d_b(x, \mu_j)$$

Novelty of x

Sum over the k nearest neighbors from archive and population

behavioral distance between x and μ_j

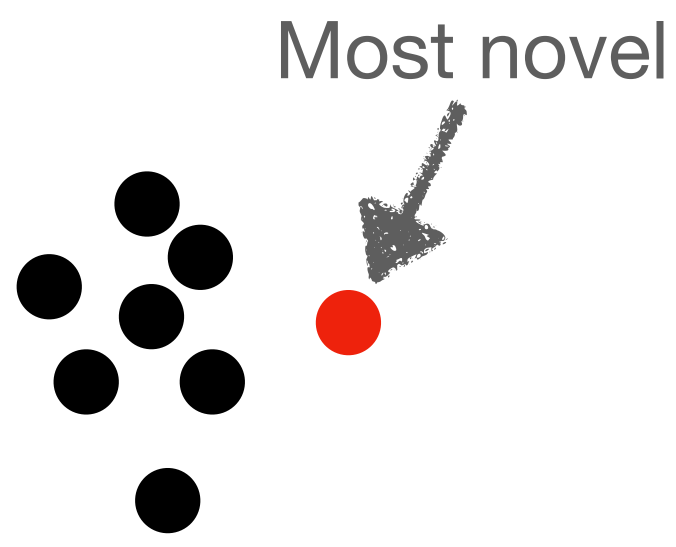

Novelty search

deceptive maze

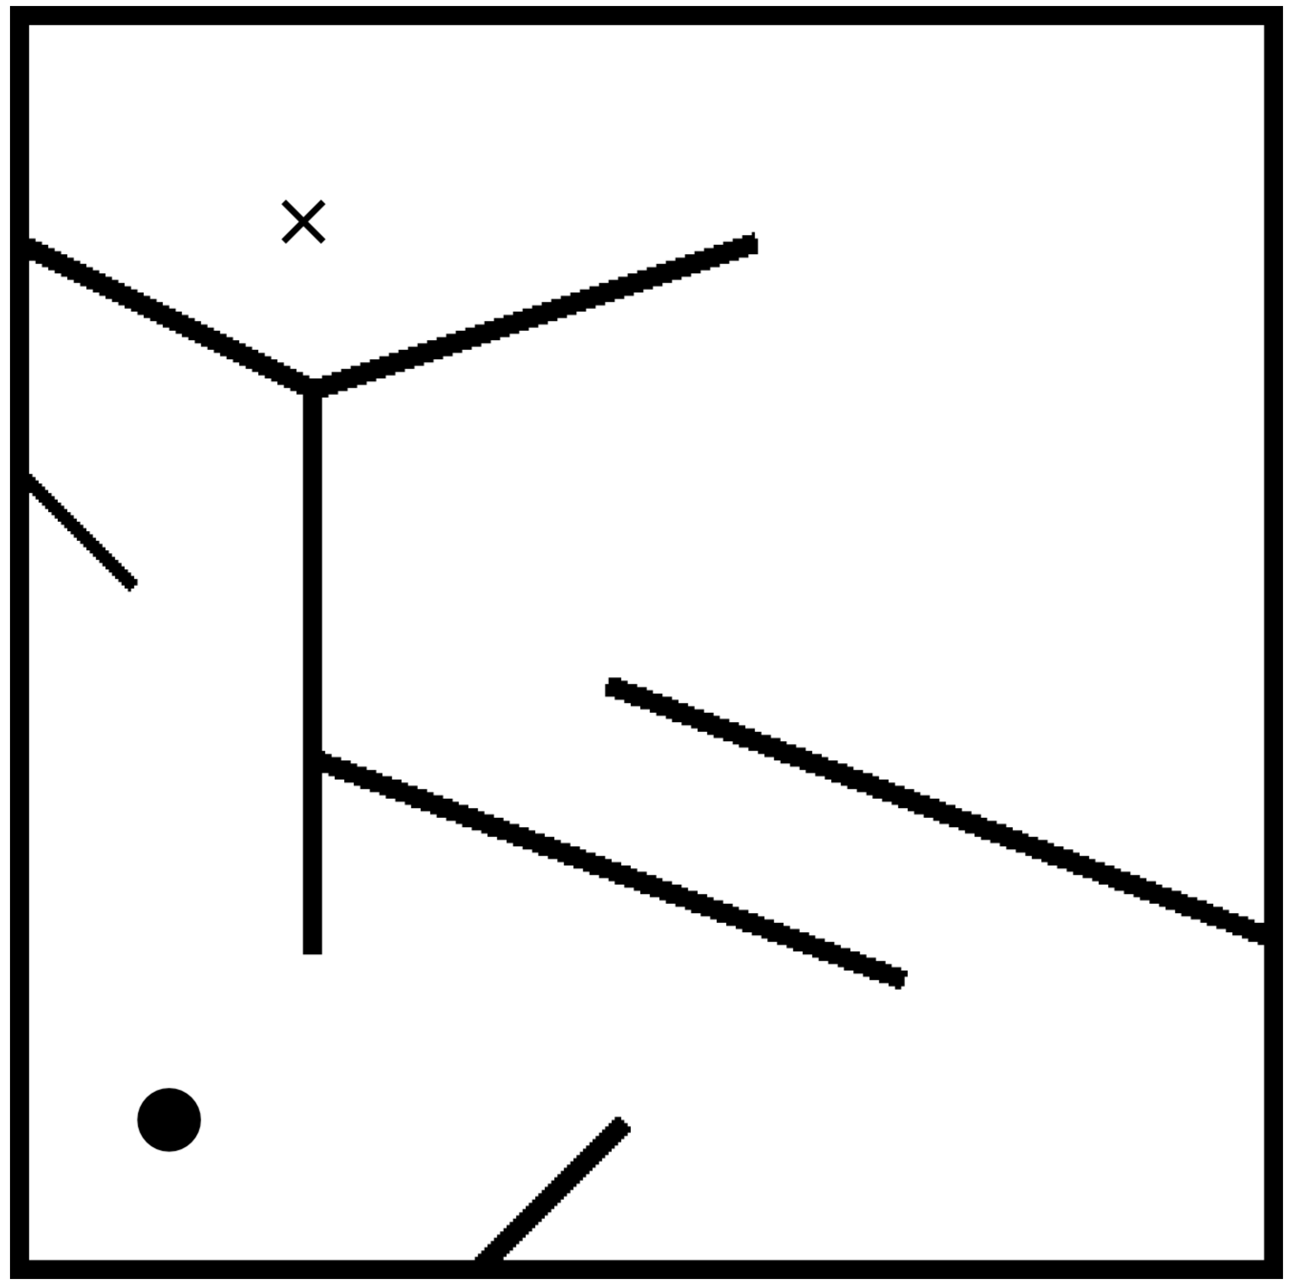

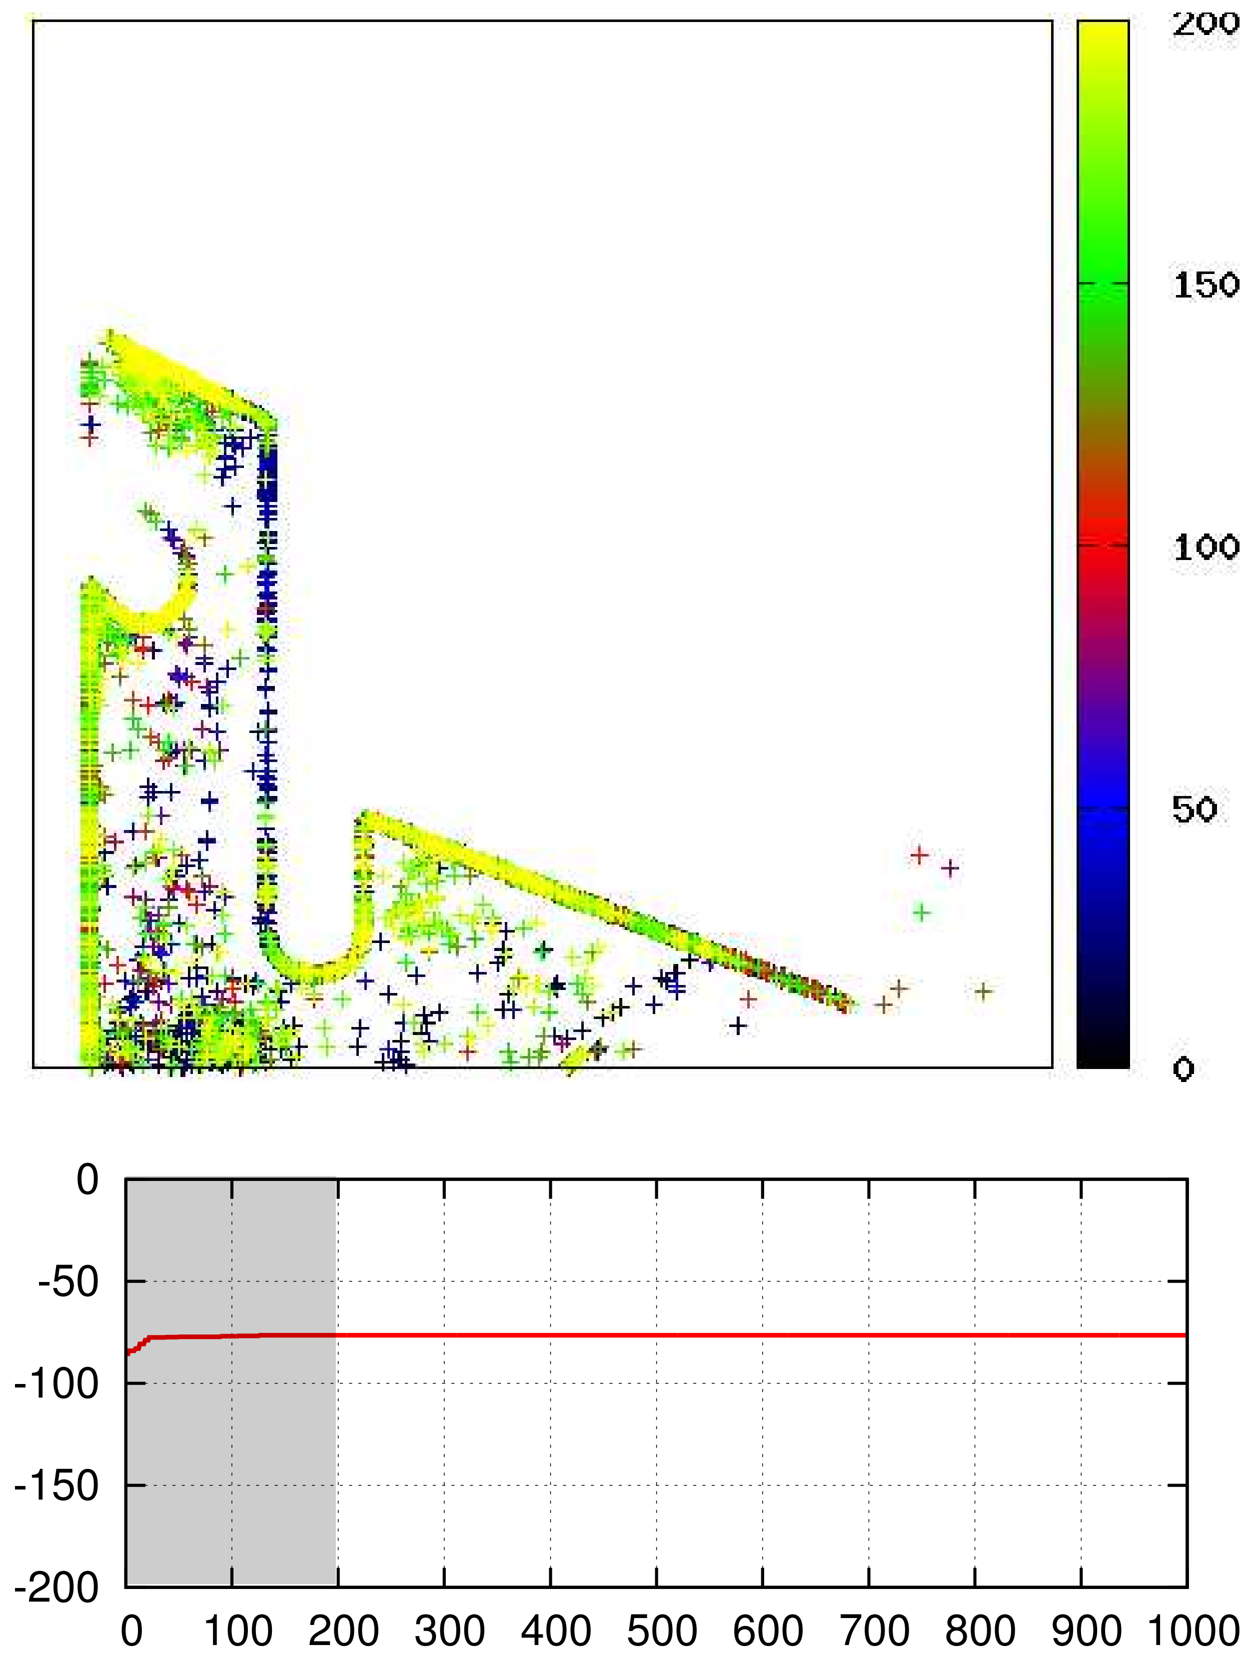

(a) Fitness only (single obj.)

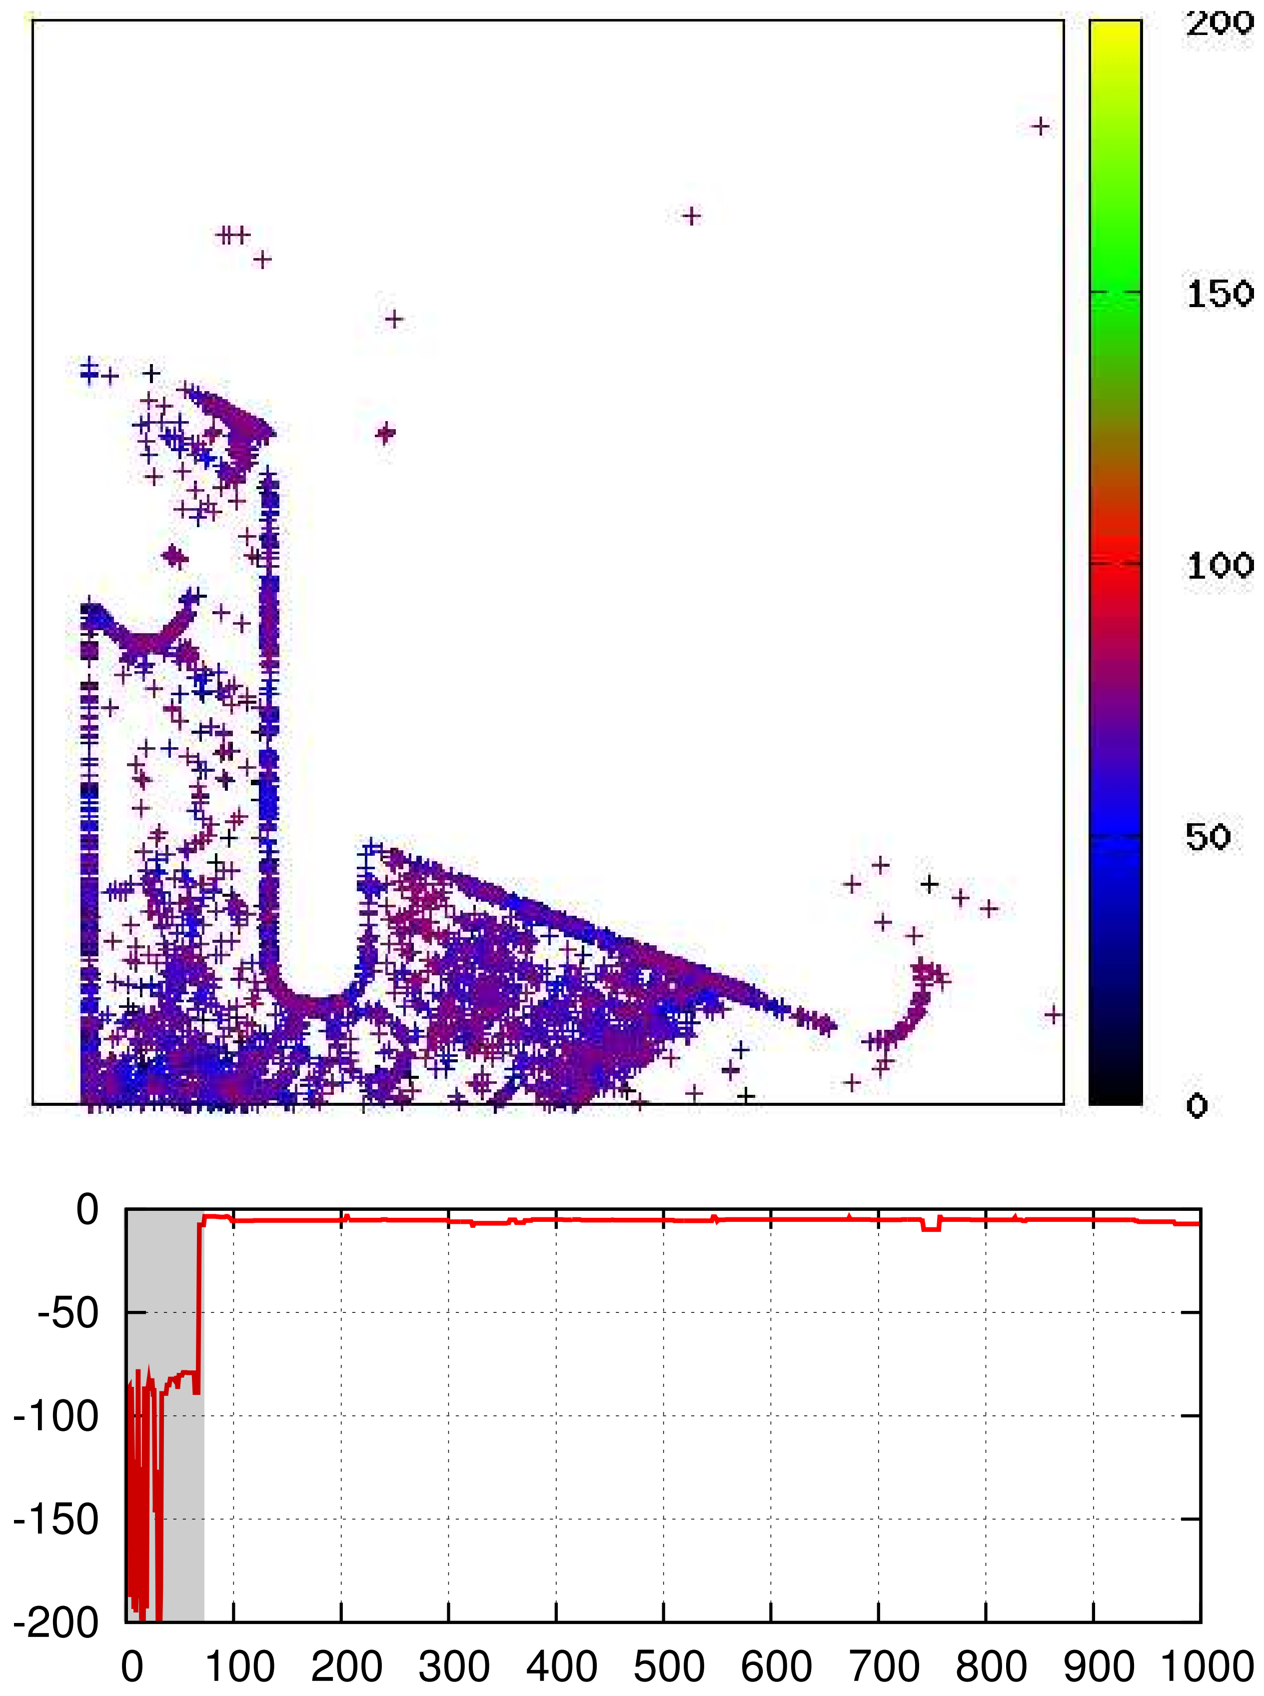

(e) Novelty (original, single obj.)

Behavior space = (x,y) position at the end of the evaluation period
Genotype = neural network (direct encoding, e.g. NEAT)

Novelty Search: demo

Another maze

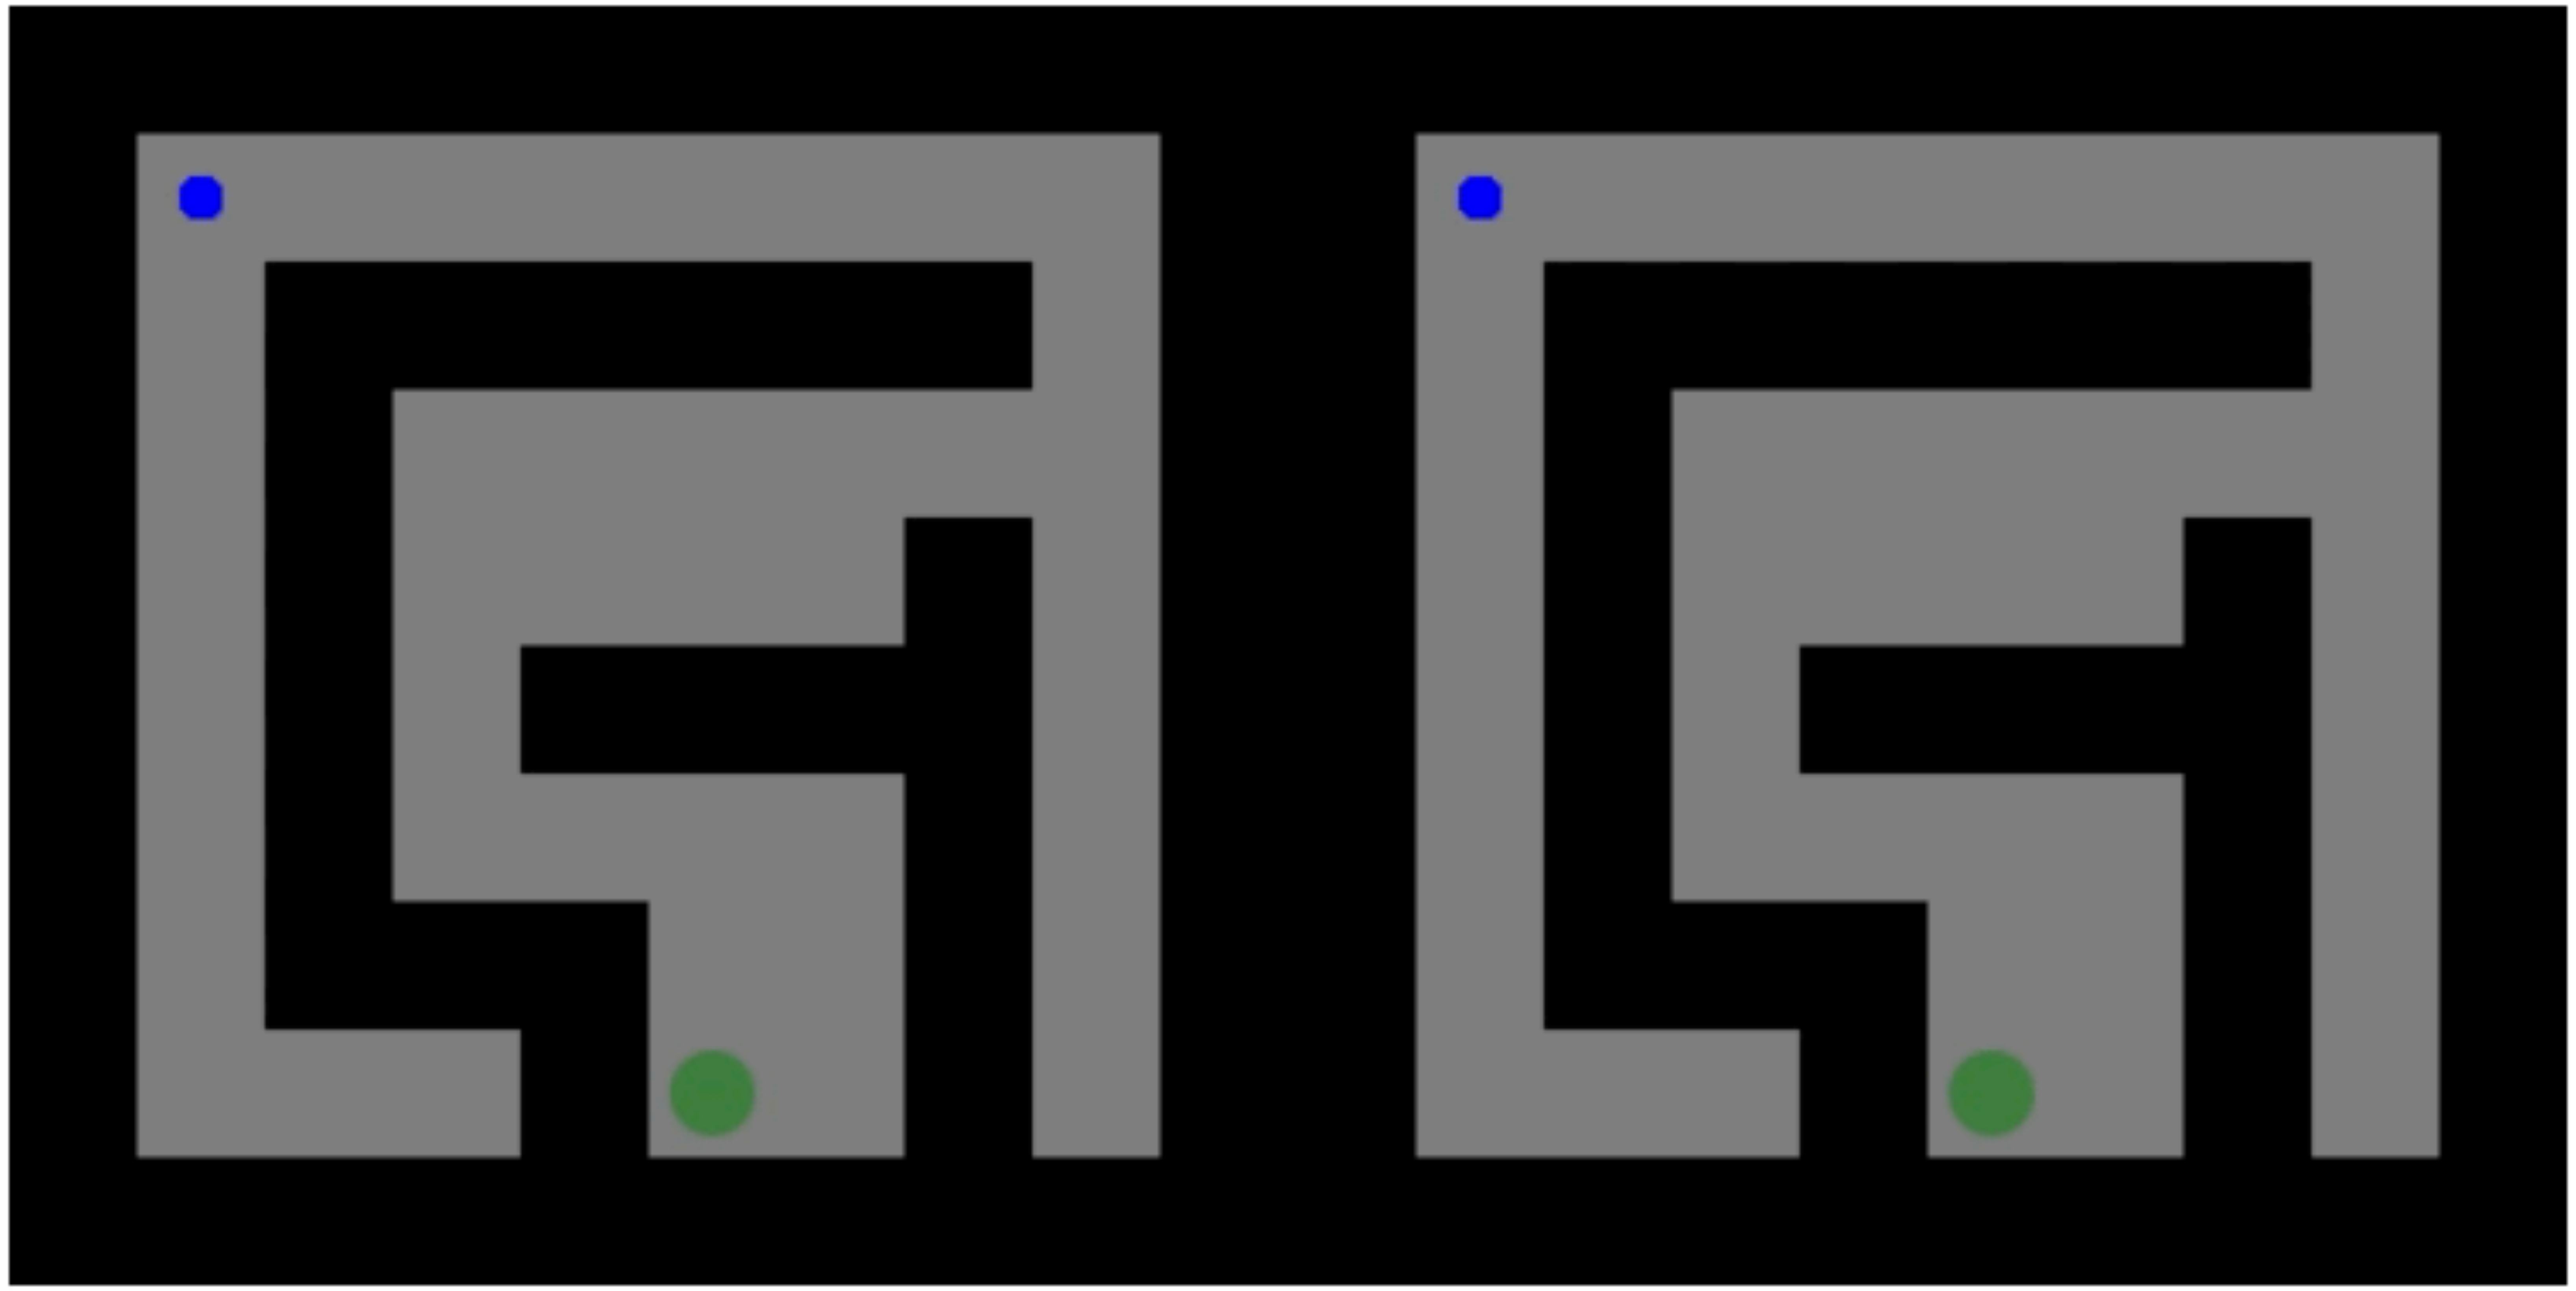

Fitness   Novelty

Source: http://eplex.cs.ucf.edu/noveltysearch/userspage/demo.html#d3

Novelty search

more complex example

- Once all the “easy behaviors” exist in the archive (e.g. falling)
- … the agents have to be creative! (e.g., walking)

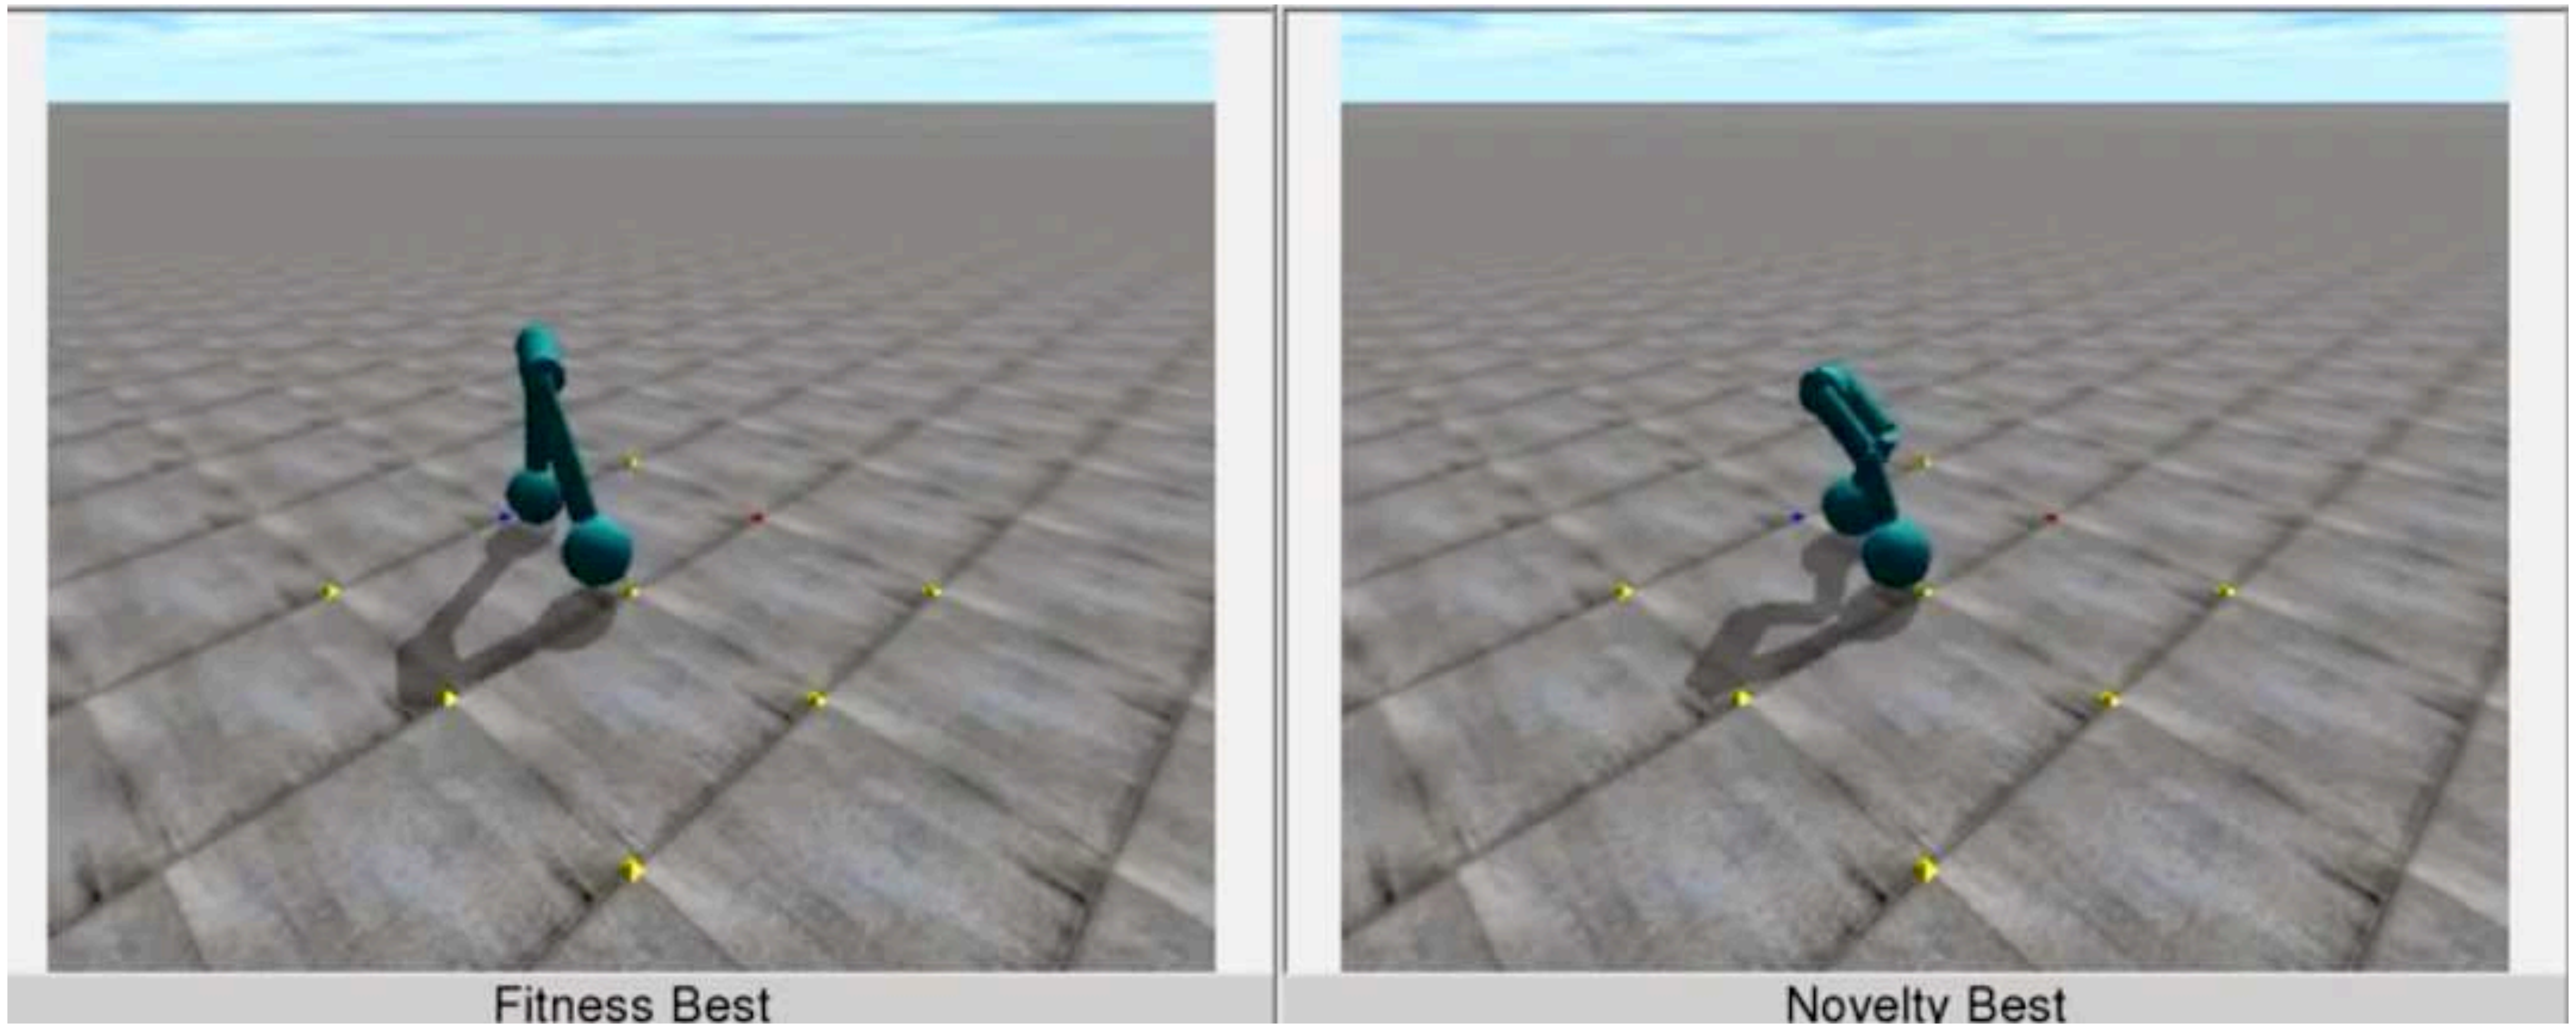

- behavior = position of the center of mass at the end of the episode

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

MAP-Elites

Multi-dimensional Archive of Phenotypic Elites

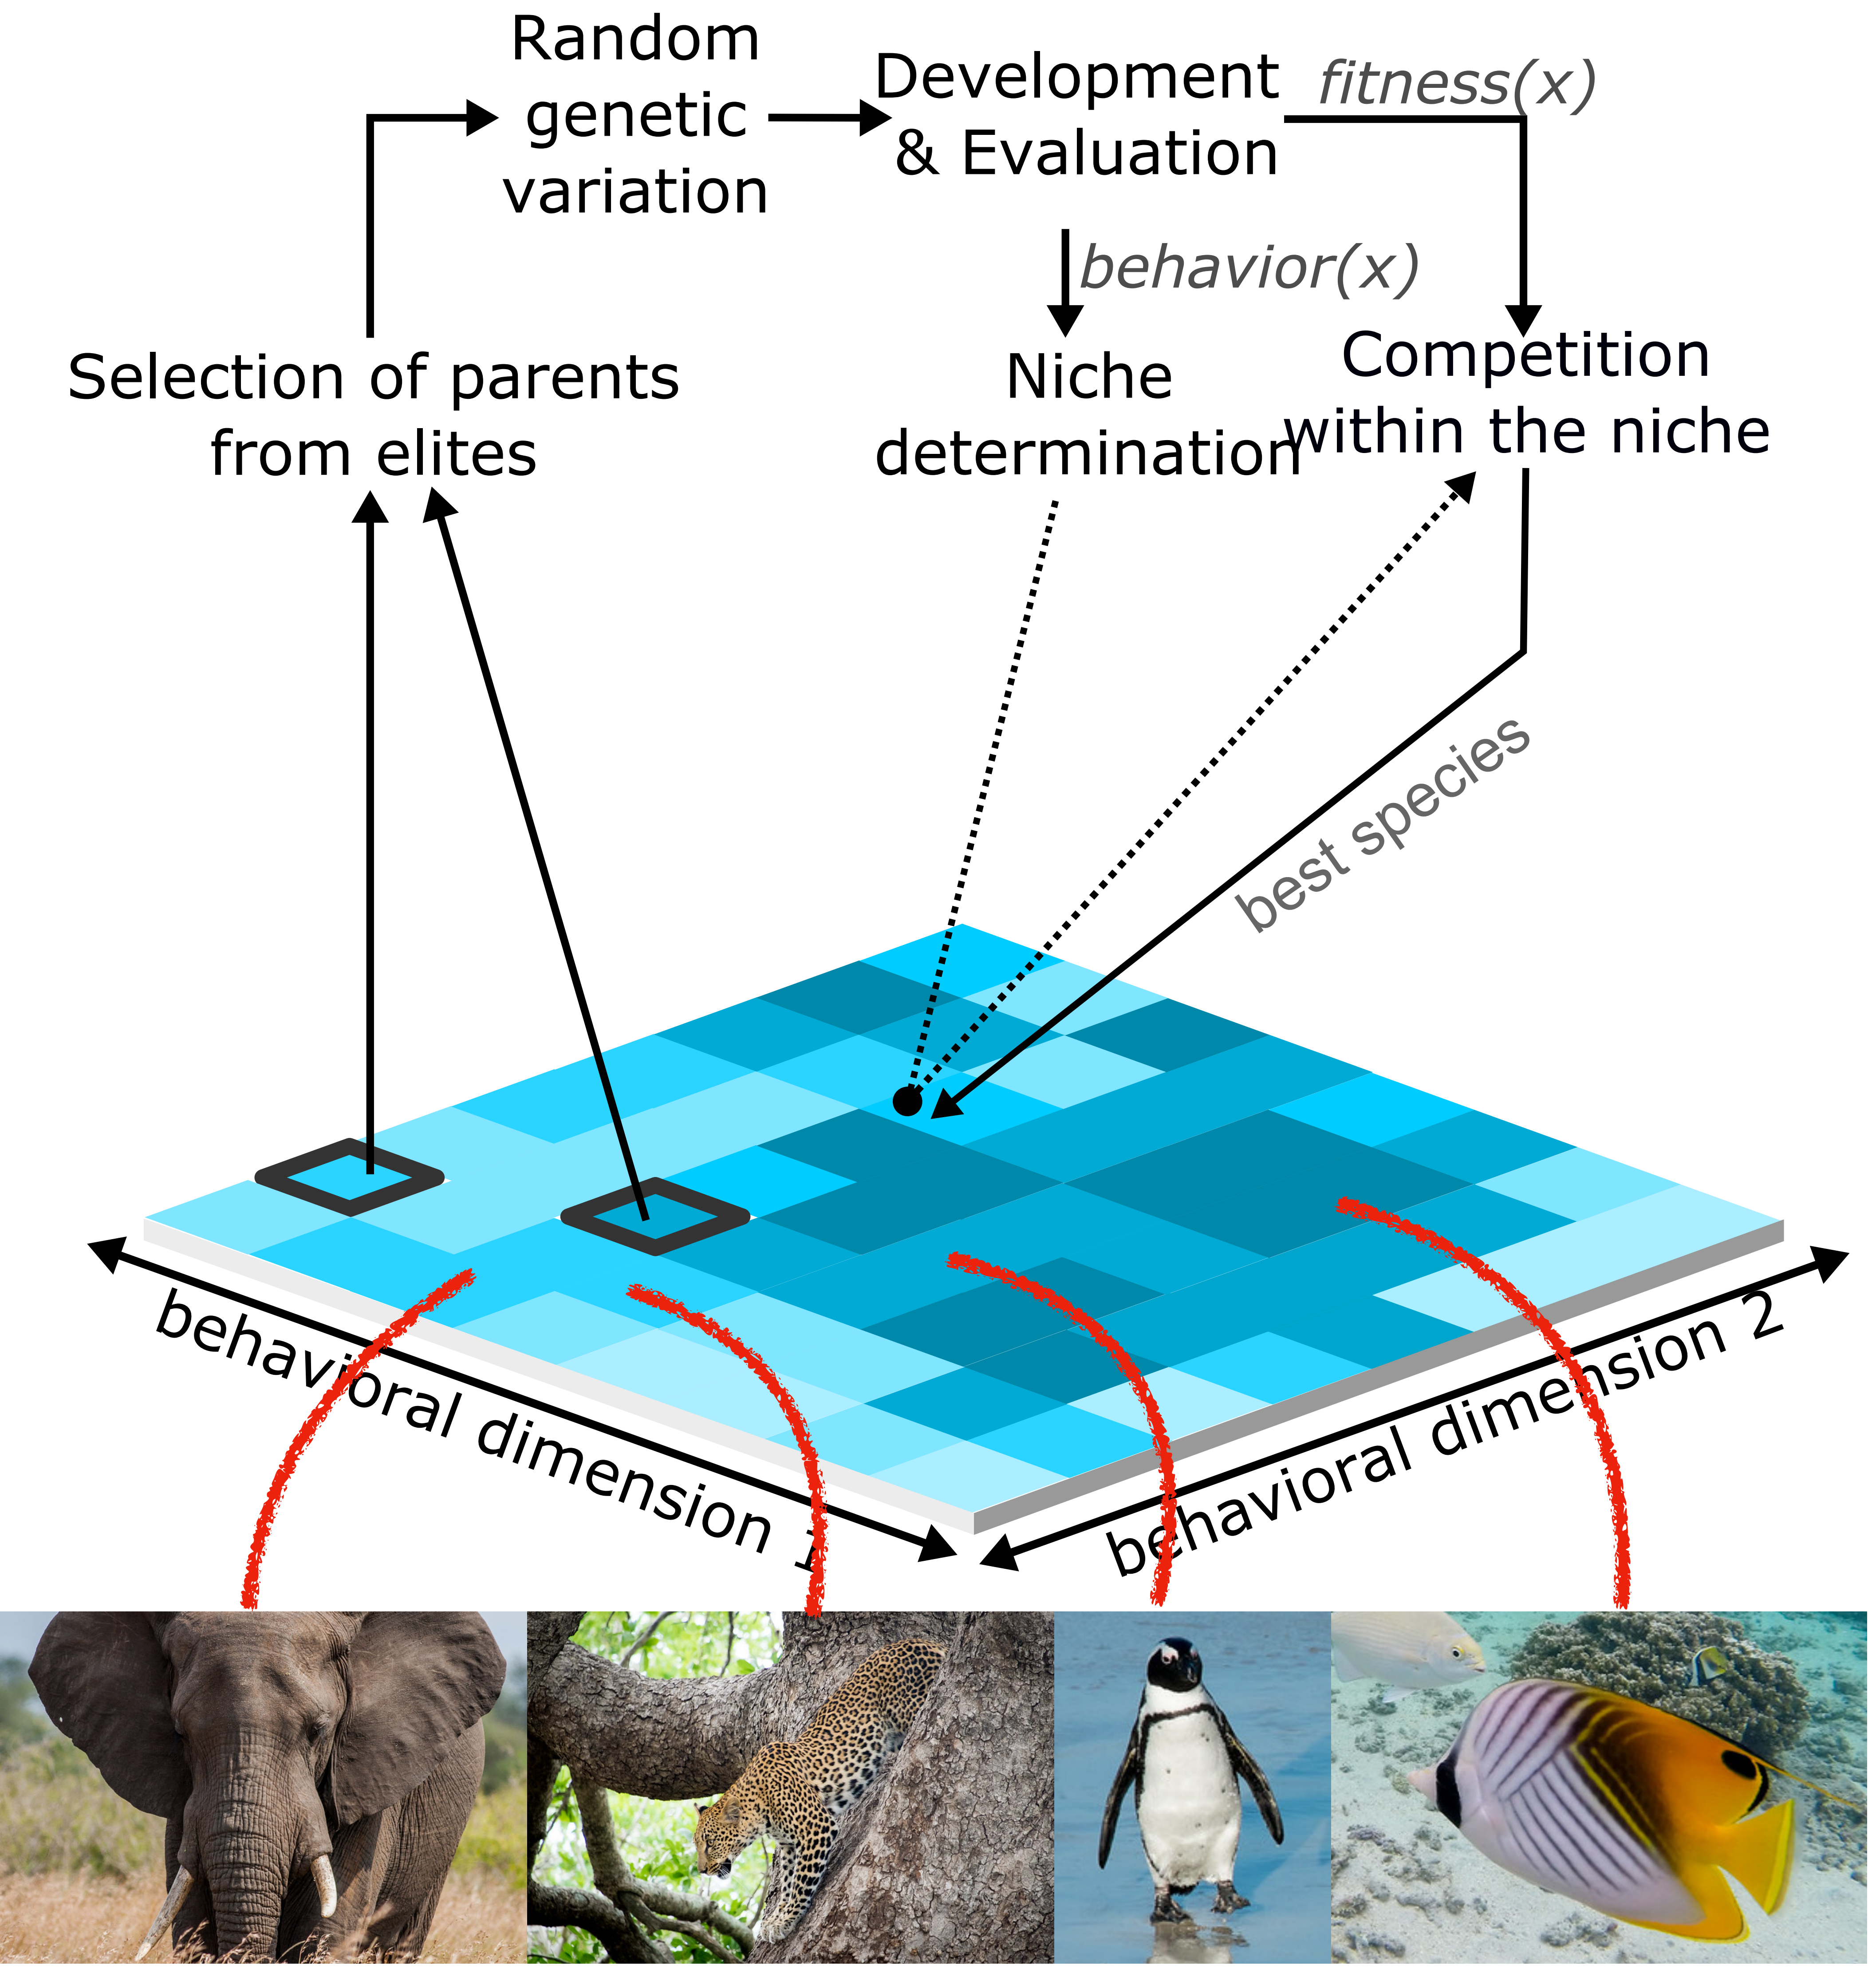

Objective: Find many good ways of solving a problem

Assumption: the fitness/reward function returns:
- a fitness/reward
- a behavioral vector (how is it solved)

$$f_\theta, b_\theta \leftarrow f(\theta)$$

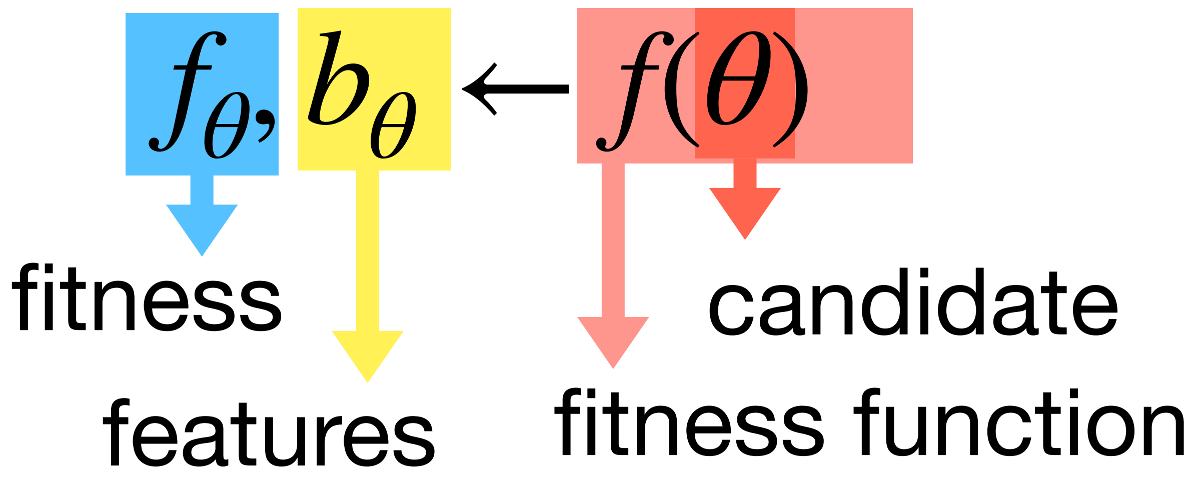

Underlying ideas:
- closer to natural evolution, emphasize diversity
- more creative process (not pure RL/optimization)
- less exhaustive than Novelty Search

Mouret, J.-B., and J. Clune. "Illuminating search spaces by mapping elites." arXiv preprint arXiv:1504.04909 (2015).
Mouret J.-B. Evolving the behavior of machines: from micro to macroevolution. iScience. 2020 Oct 28:101731.

Example: planar arm

see notebook

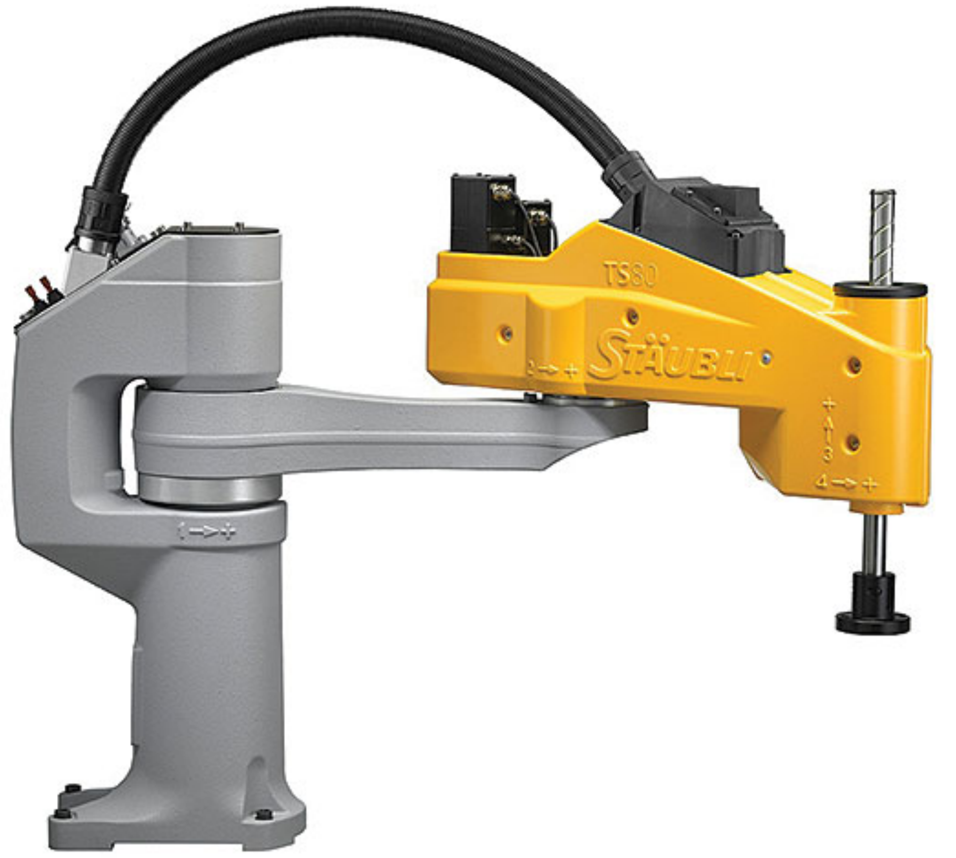

Image credit:
Jamit (Wikimedia commons))

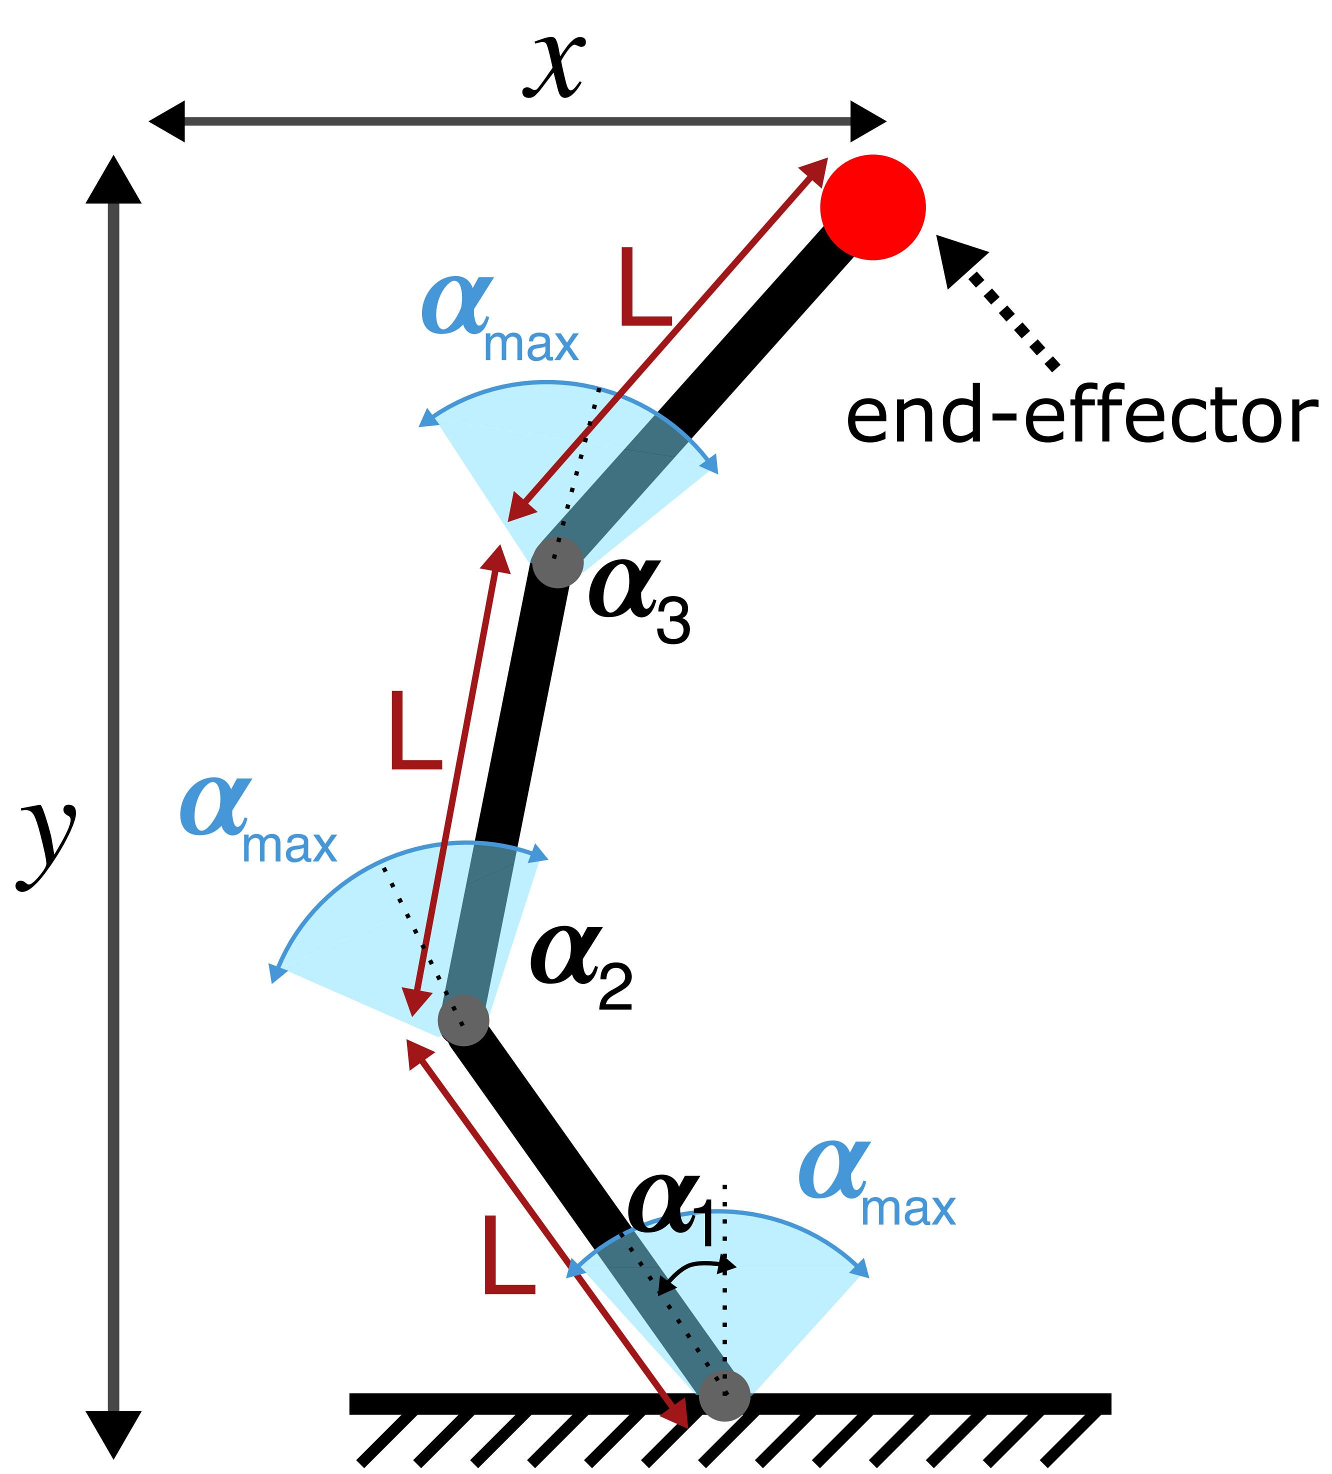

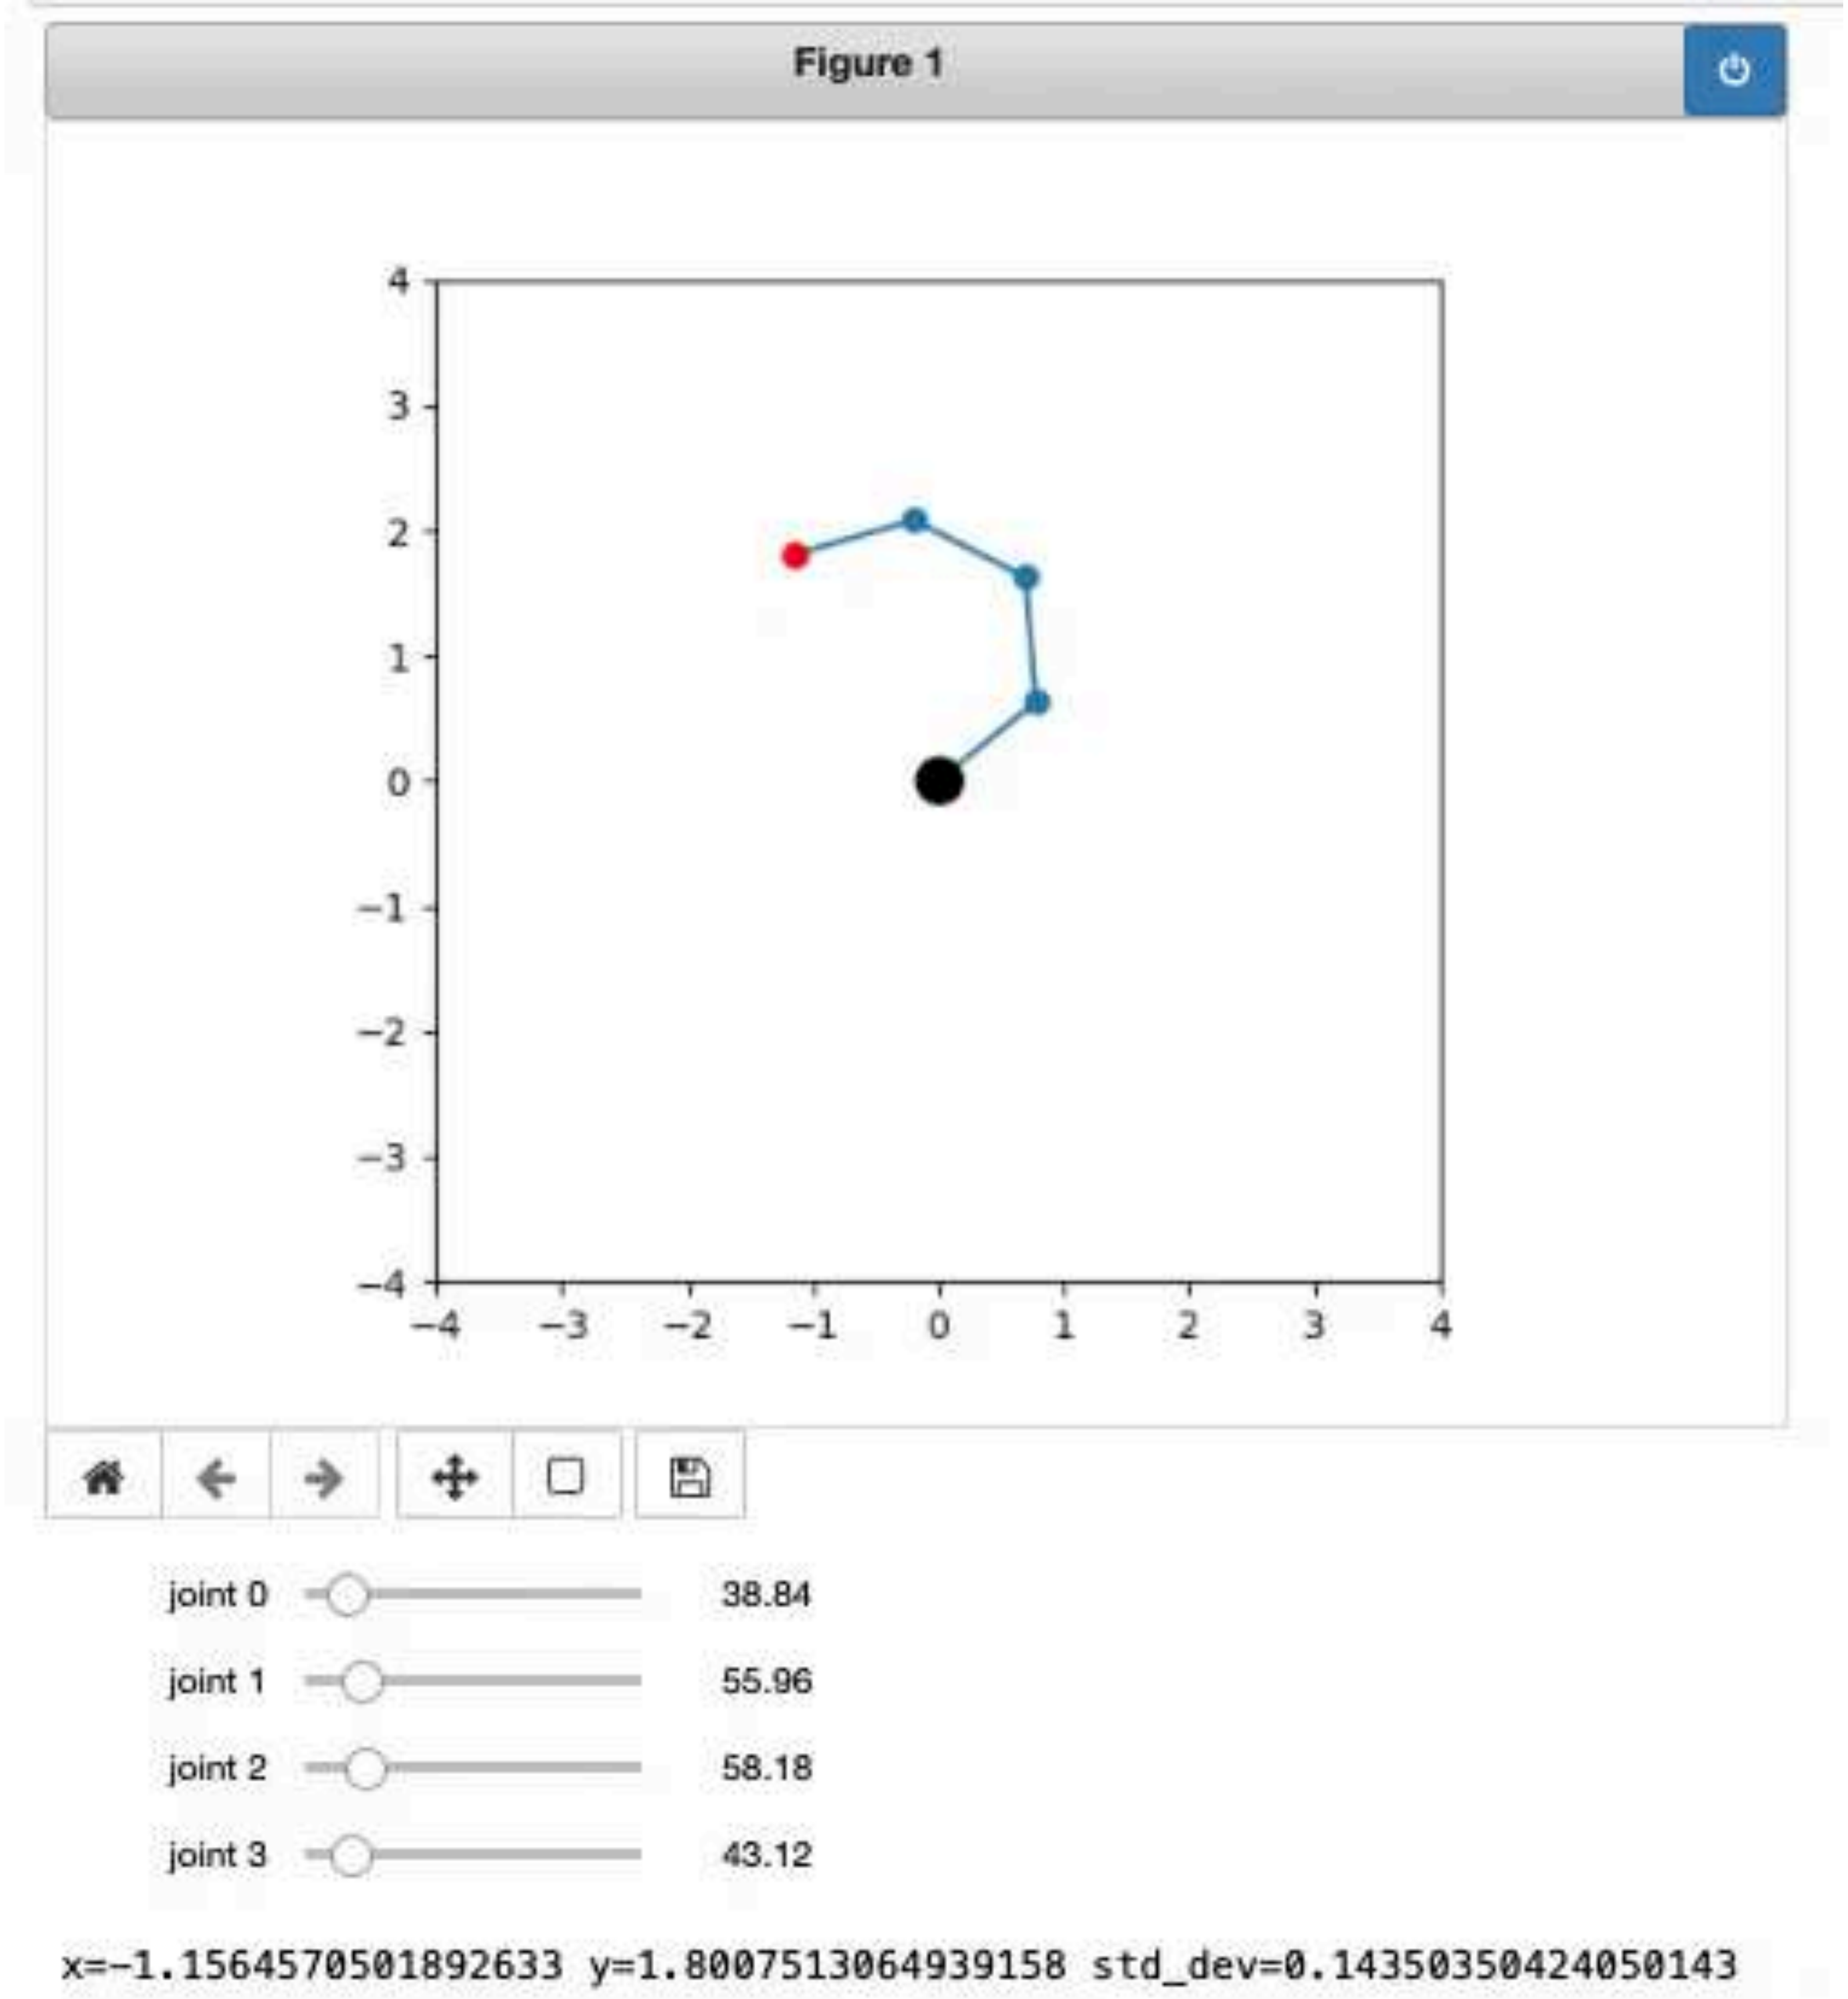

- **Search space:** $[\alpha_1, \cdots, \alpha_n]$ (n-dimensional)
- **Behavior space:** (x,y) (2-dimensional)

Example: planar arm

Fitness function (notebook)

```python
def fitness(genotype):
    # fitness is the standard deviation of joint angles (Smoothness)
    # (we want to minimize it)
    fit = 1 - np.std(genotype)

    # now compute the behavior
    #   scale to [0,2pi]
    g = np.interp(genotype, (0, 1), (0, 2 * math.pi))
    j = forward_kinematics(g, [1]*len(g))
    #  normalize behavior in [0,1]
    b = (j[-1,:]) / (2 * len(g)) + 0.5
    return fit, b# the fitness and the position of the last joint
```

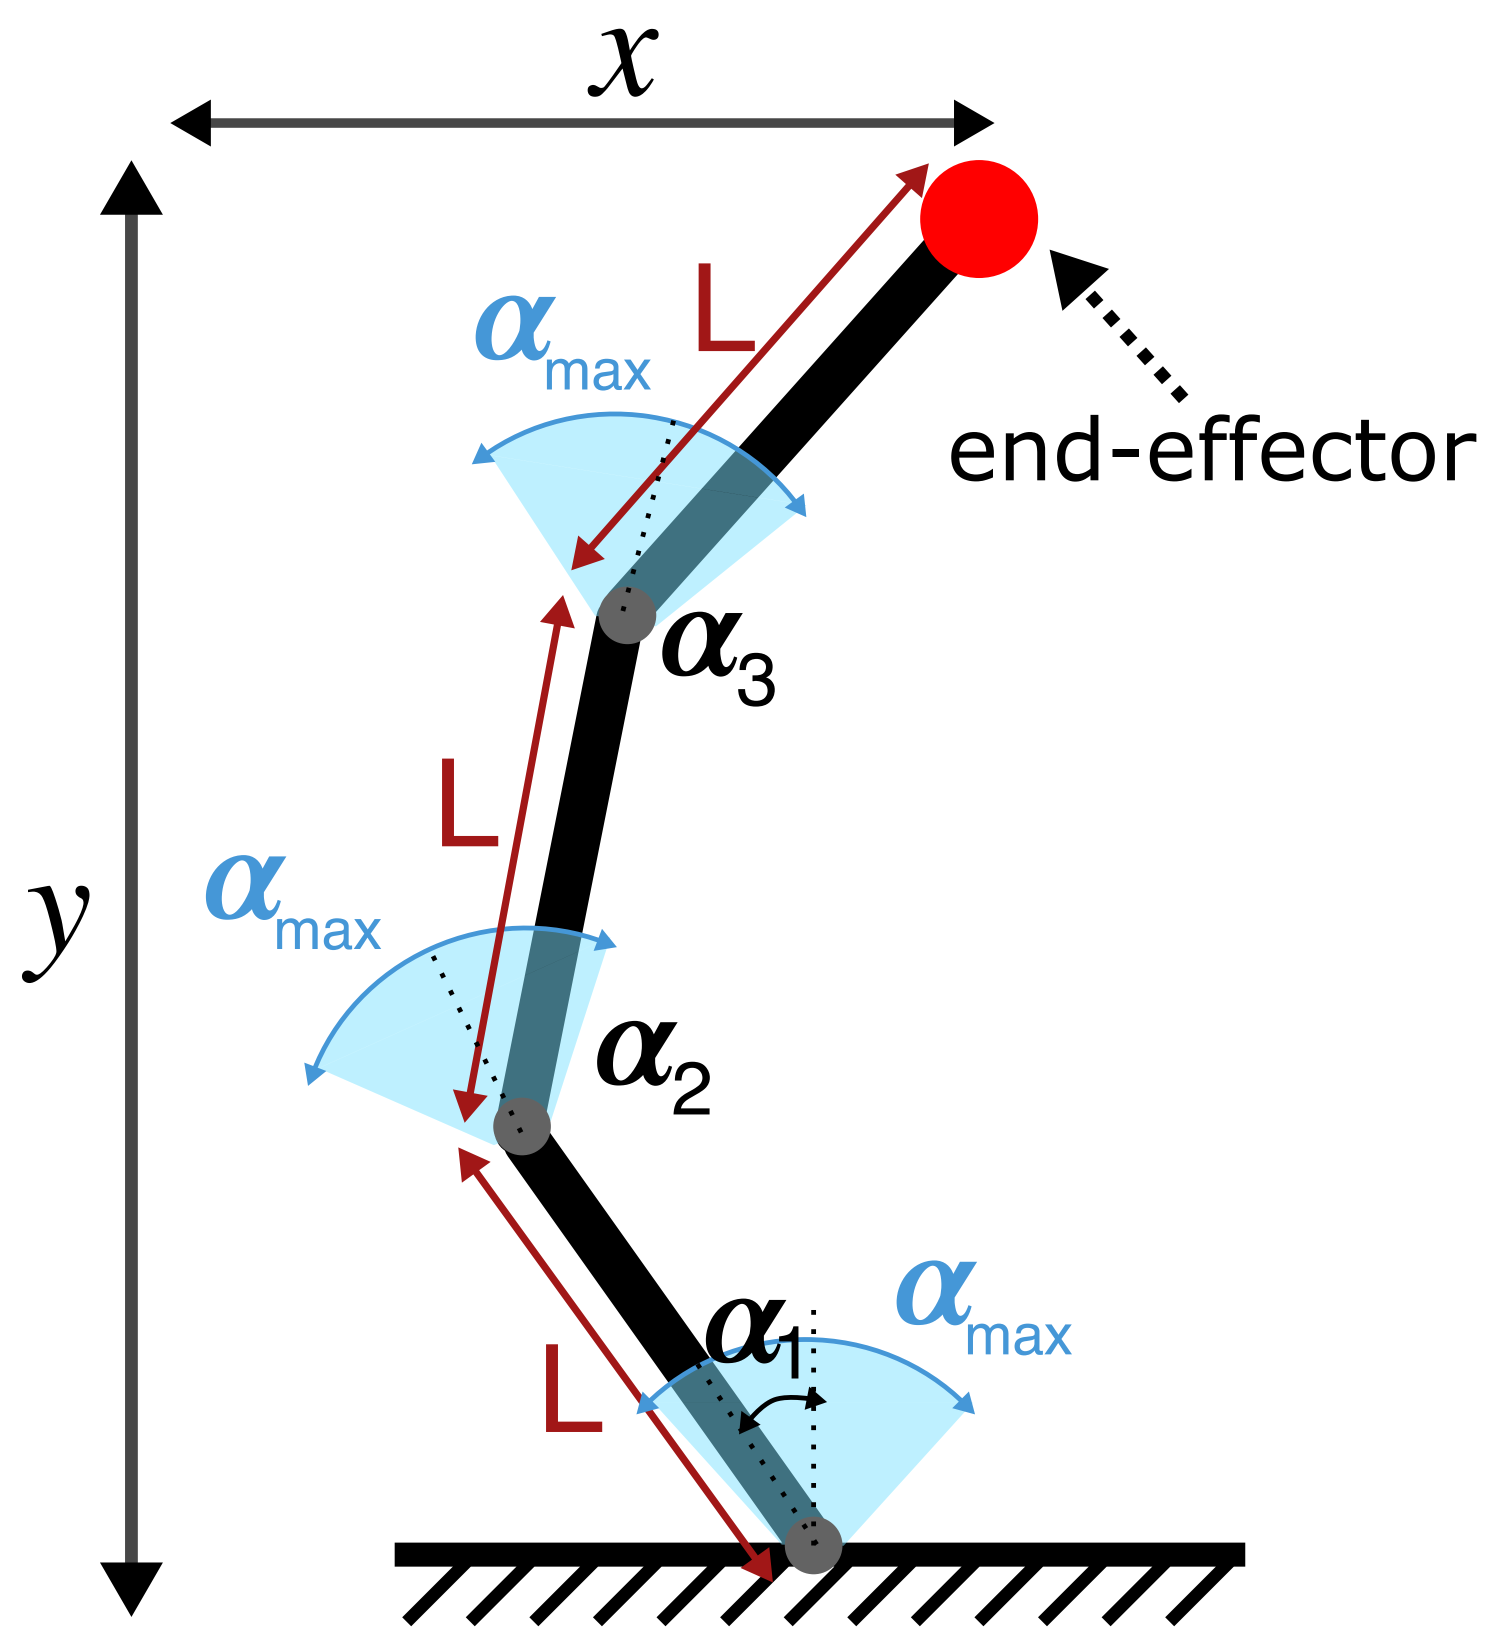

Example: planar arm

Archive management

```python
# for simplicity, this is 2-Dimensional MAP-Elites
cols = 30
rows = 30
num_random = 100
num_dofs = 4

# we should use a dataclass, but this 3.9+
class Species:
    def __init__(self, genotype, behavior, fitness, niche=[]):
        self.genotype = genotype
        self.behavior = behavior
        self.fitness = fitness
        self.niche = niche

def add_to_archive(archive, species):
    n = species.behavior * np.array([rows, cols])
    x,y = min(round(n[0]), rows-1), min(round(n[1]), cols-1)
    if (not (x,y) in archive) or (archive[(x,y)].fitness < species.fitness):
        archive[(x,y)] = species
        species.niche = (x,y)
```

Archive initialization

```python
def init_archive(num_random):# fill the archive with some random solutions
    # create an archive: a dictionnary indexed by coordinates
    archive = {}
    for i in range(0, num_random):
        g = np.random.rand(num_dofs)
        f, b = fitness(g)
        add_to_archive(archive, Species(g, b, f))
    return archive
```

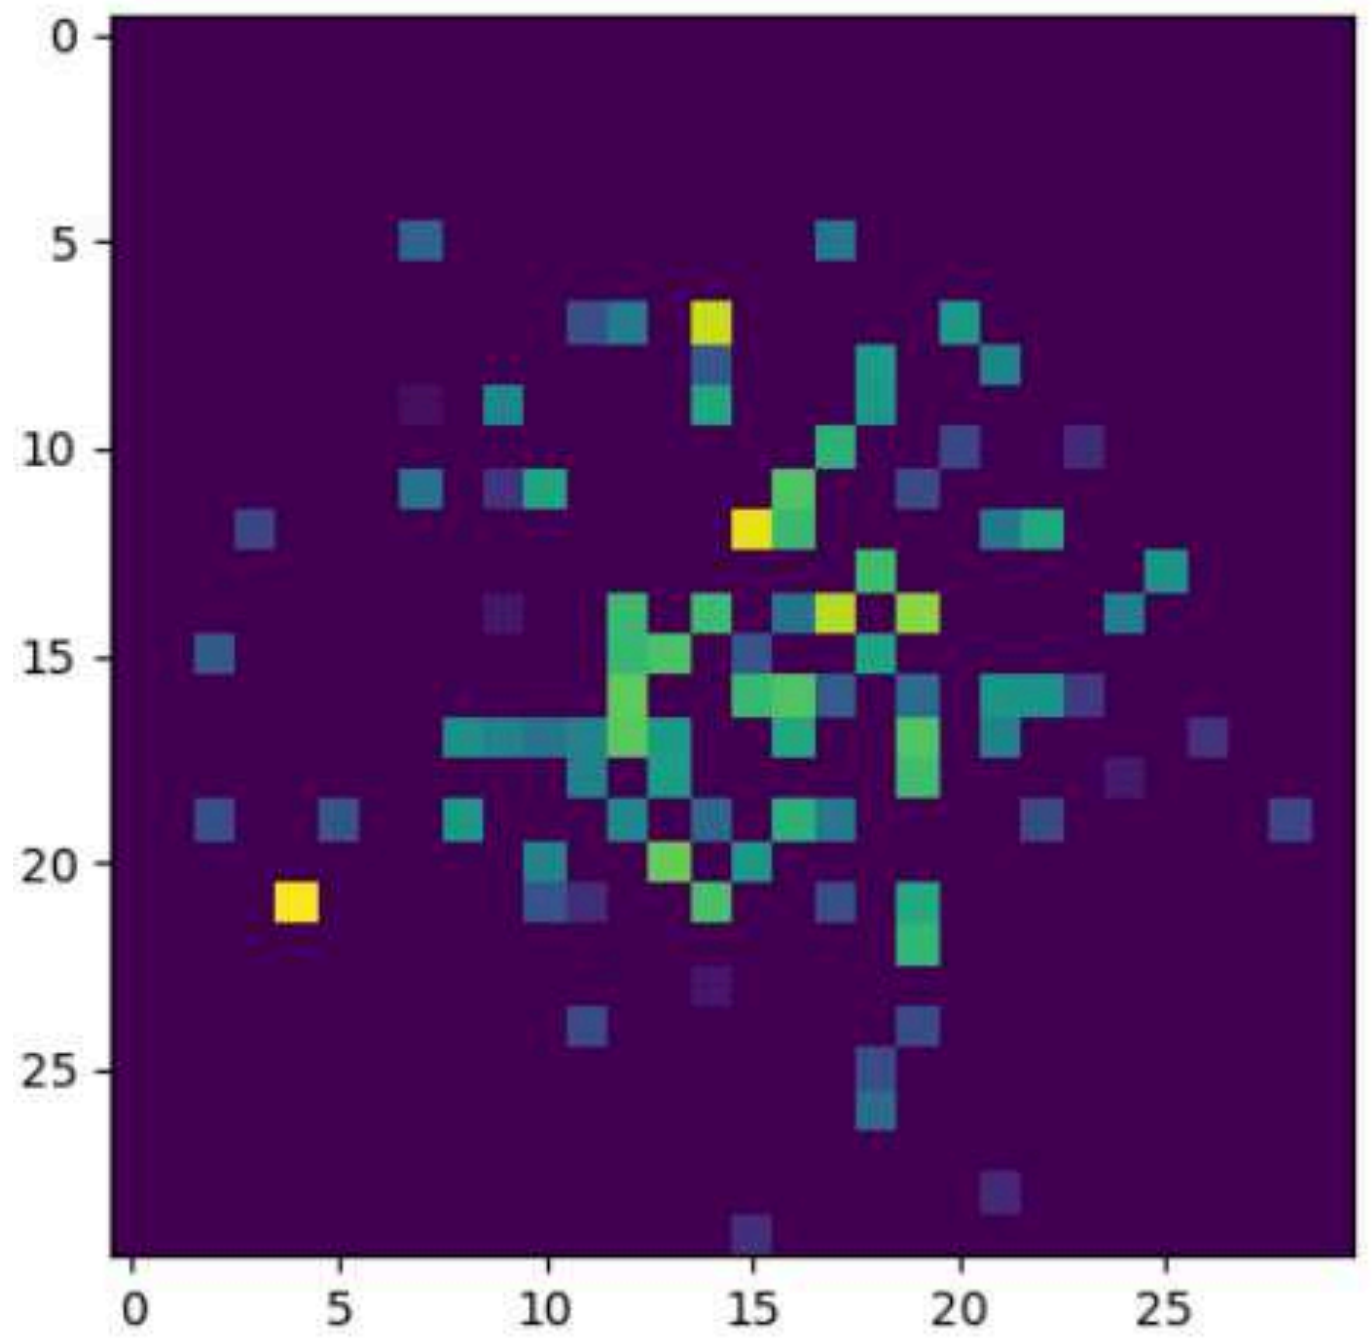

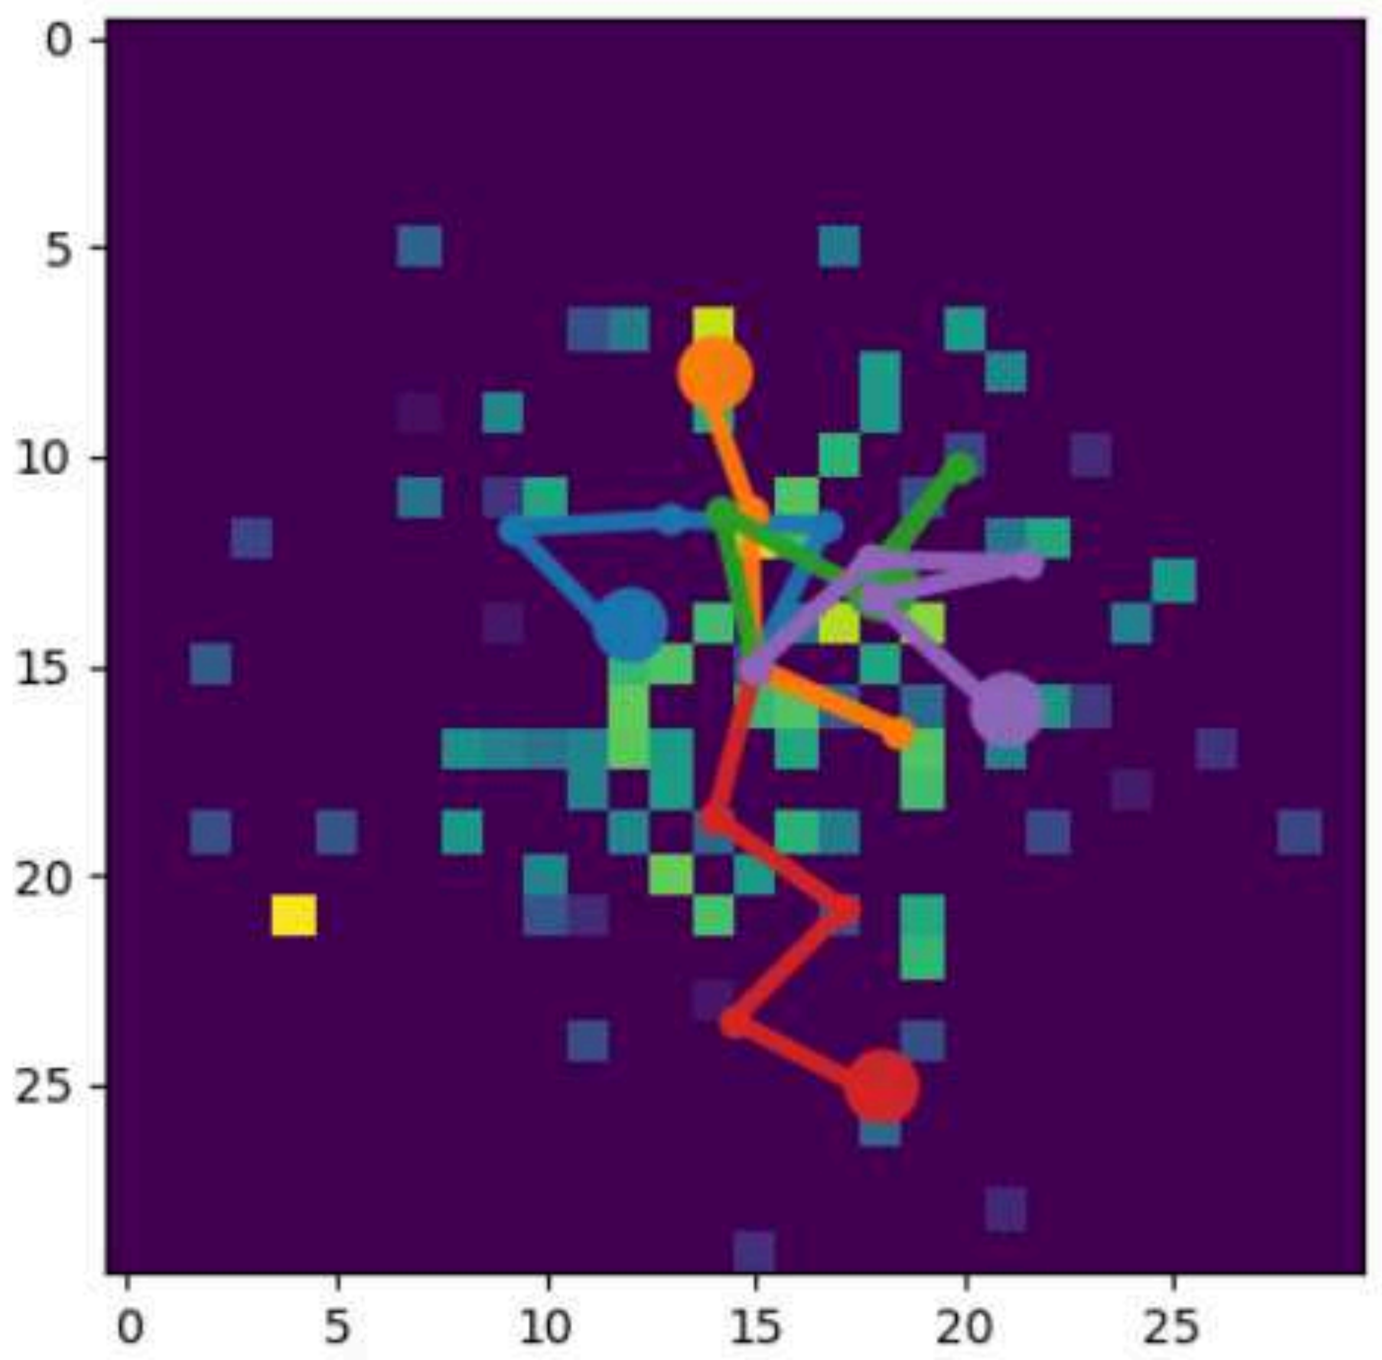

Example: planar arm

Main loop

```
num_iterations = 50
batch_size = 500

for j in tqdm(range(0, num_iterations)):
    display_archive(archive)
    for i in range(0, batch_size):
        # pick an existing random point in the archive
        x = random.choice(list(archive.values()))
        # mutate it (we can use more advanced variation techniques: NEAT, etc.)
        g = x.genotype + np.random.normal(0, 0.1, x.genotype.shape[0])
        # compute the fitness
        f, b = fitness(g)
        # add to archive
        add_to_archive(archive, Species(g, b, f))
```

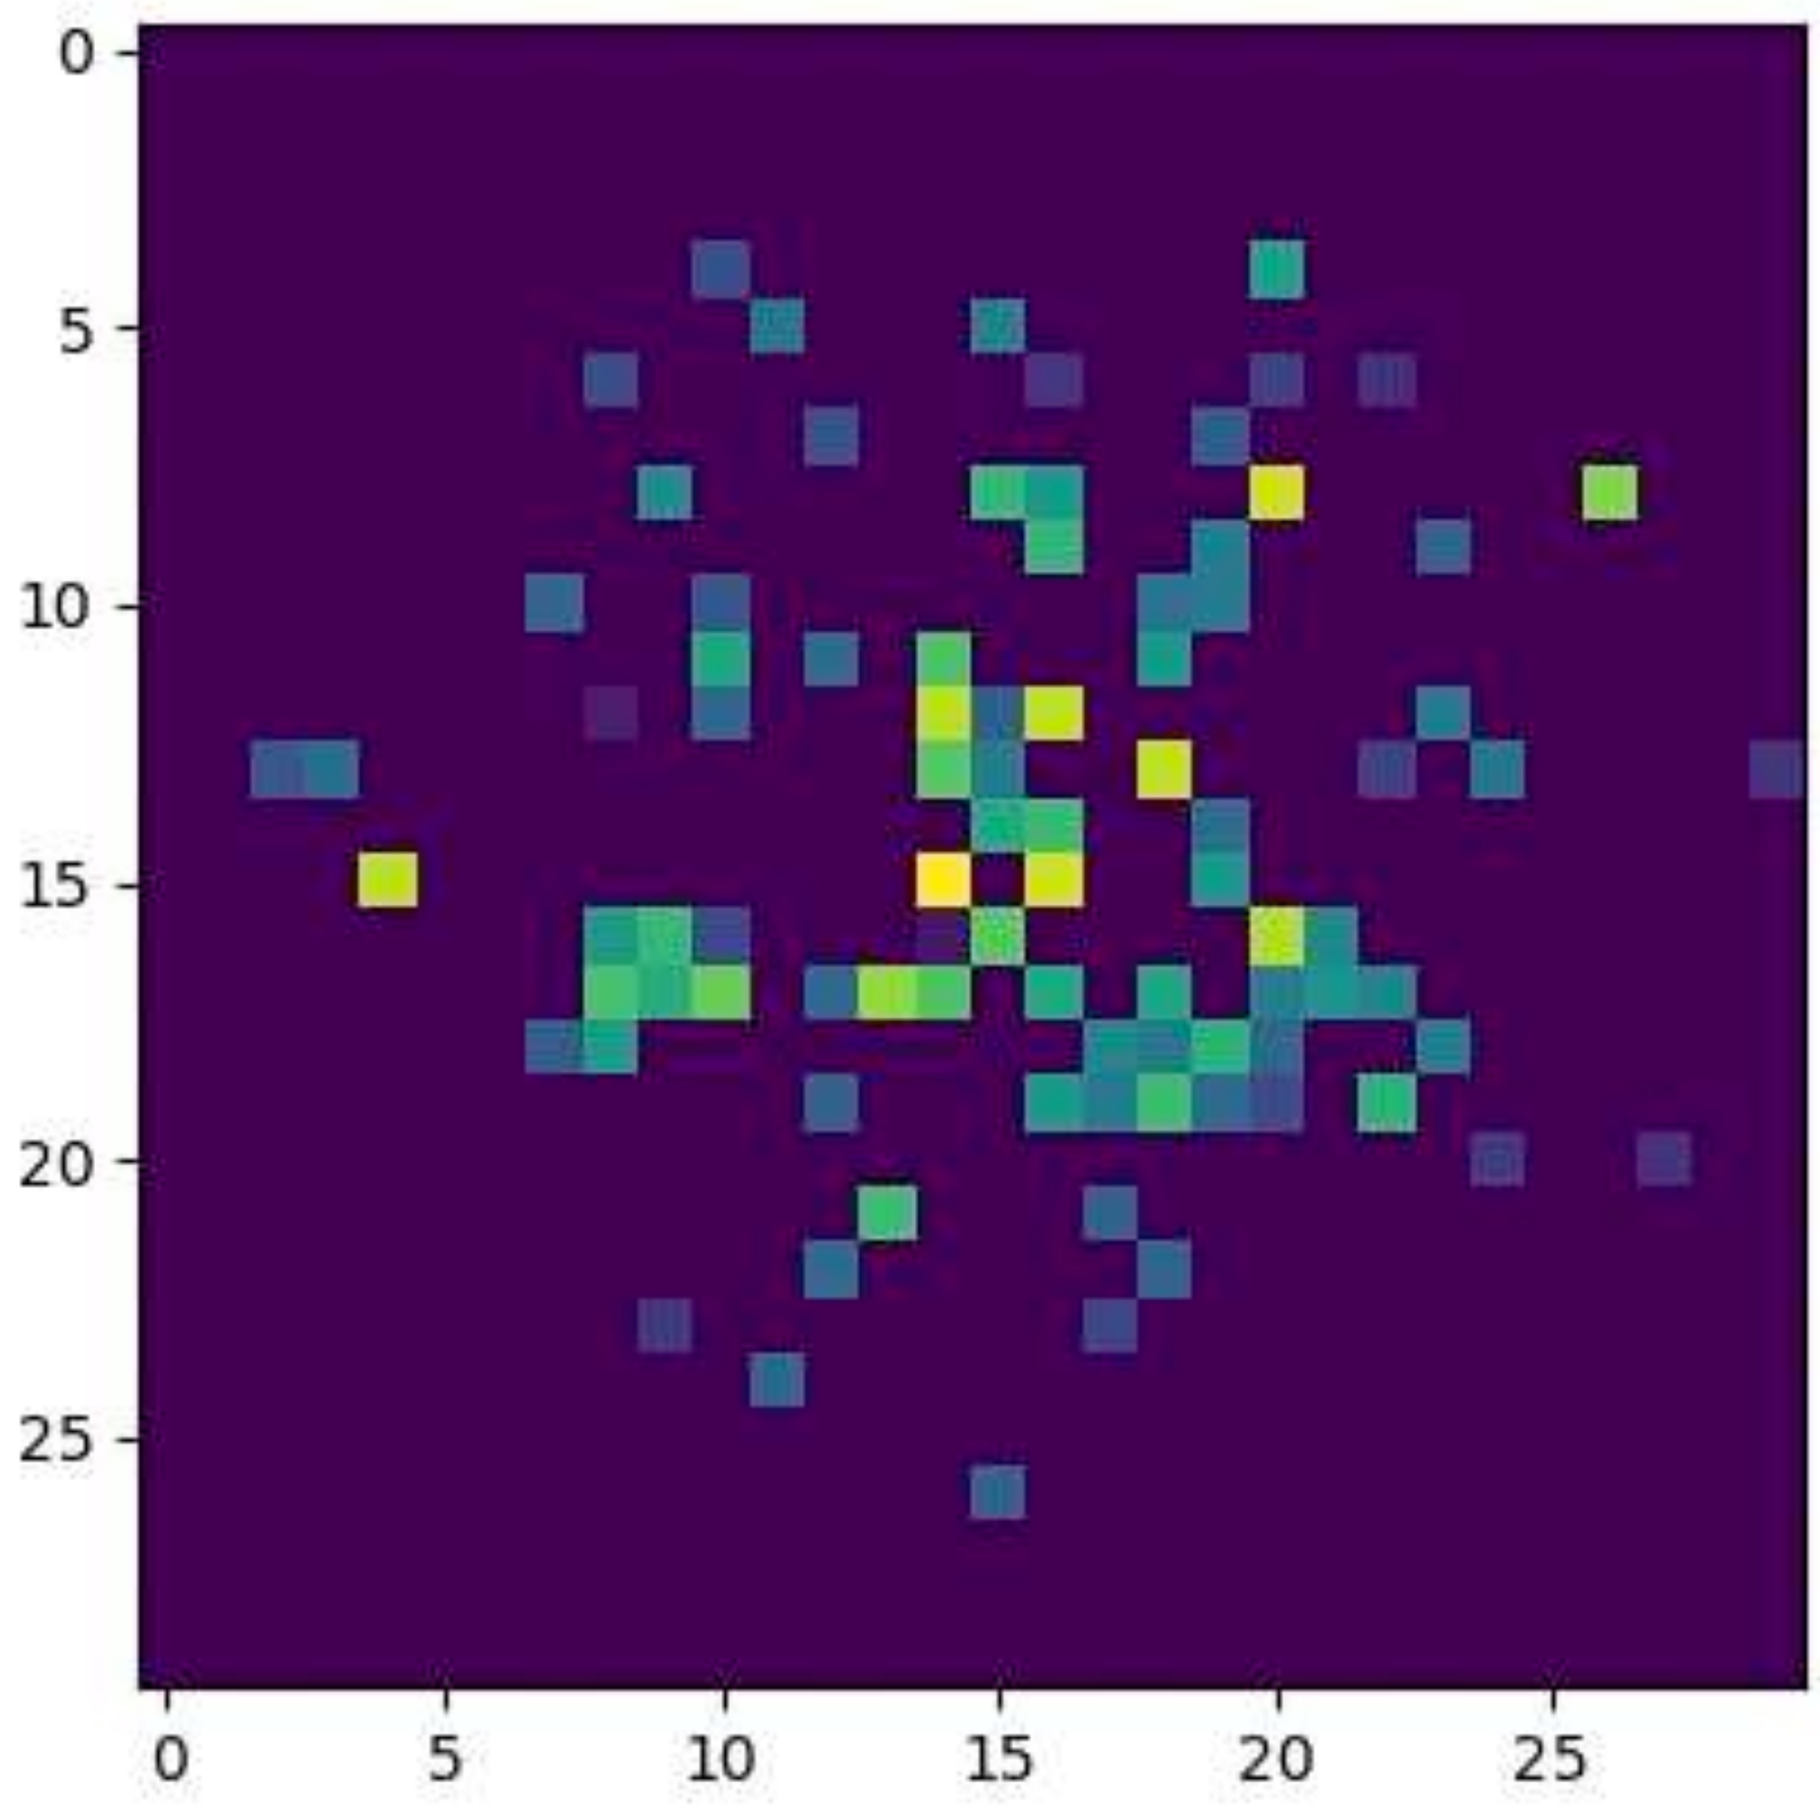

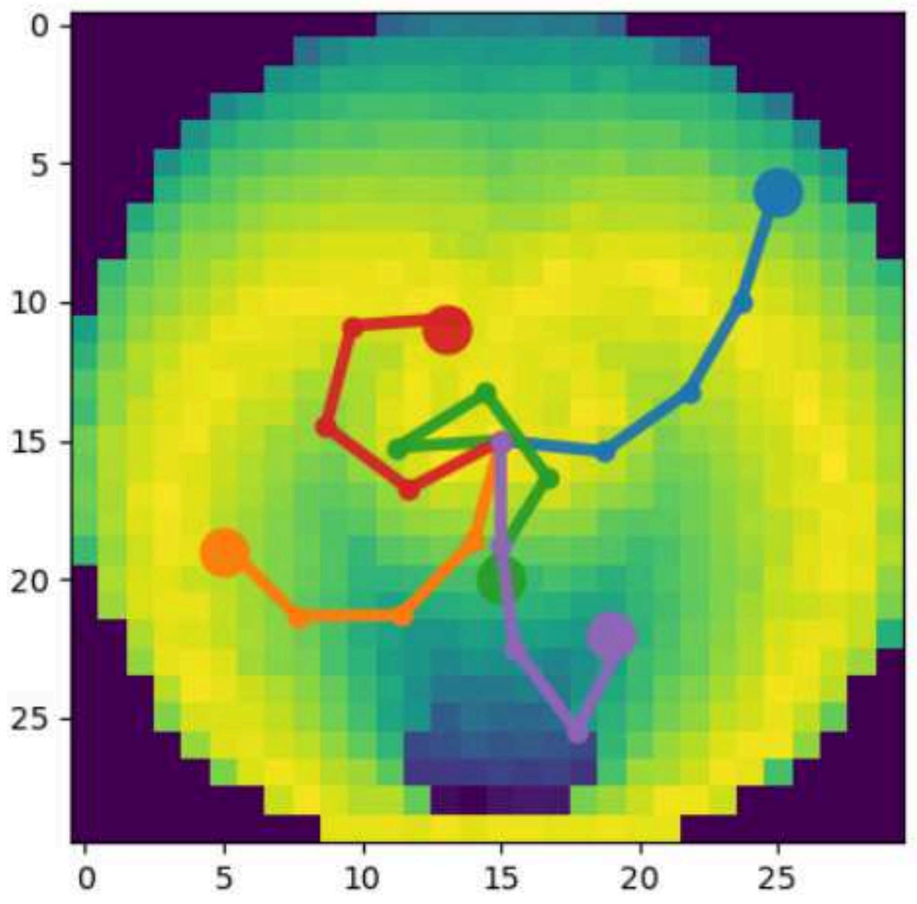

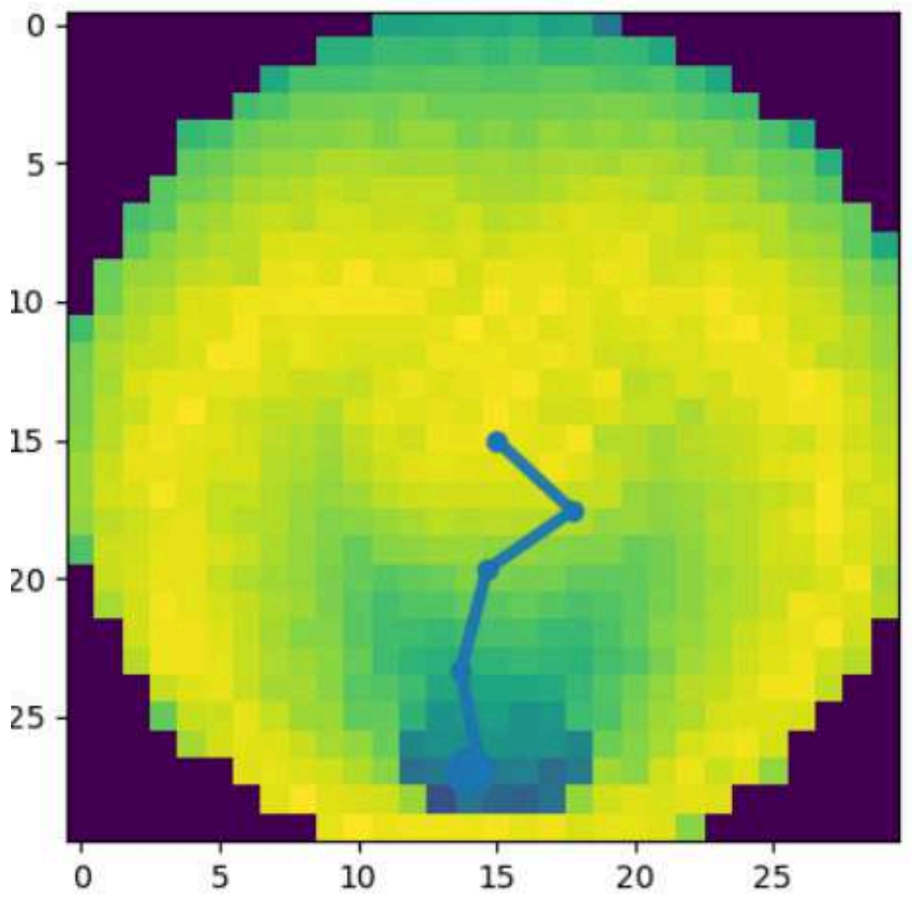

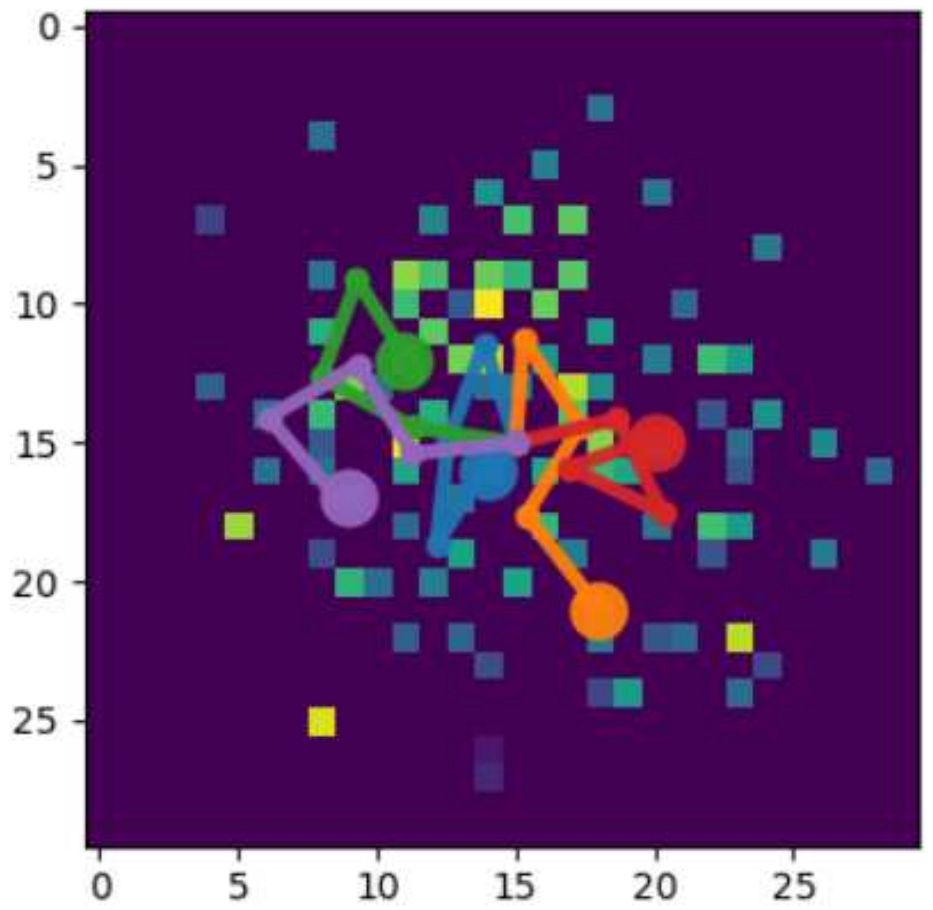

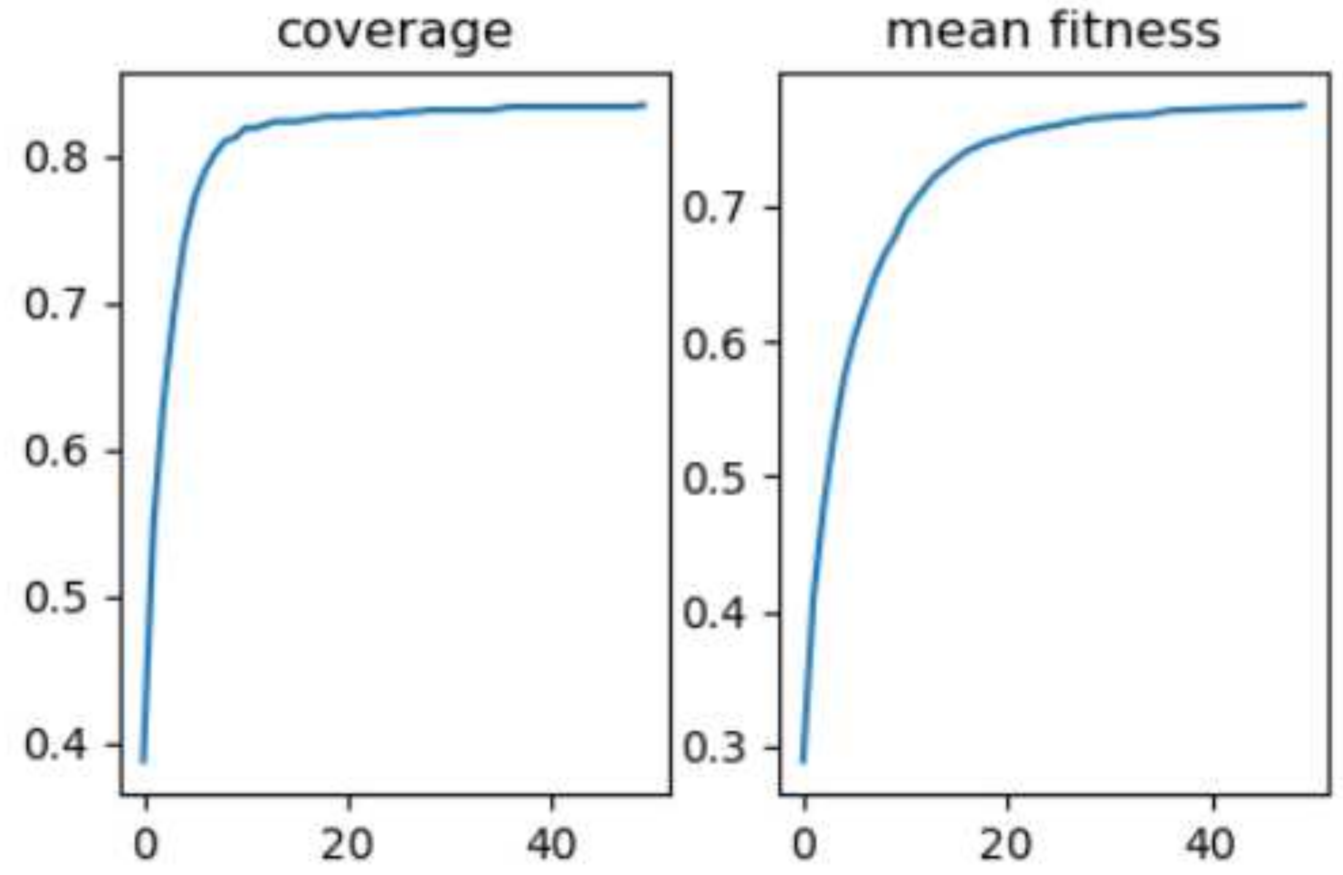

MAP-Elites: elite hypervolume

Good solutions share common “recipes"

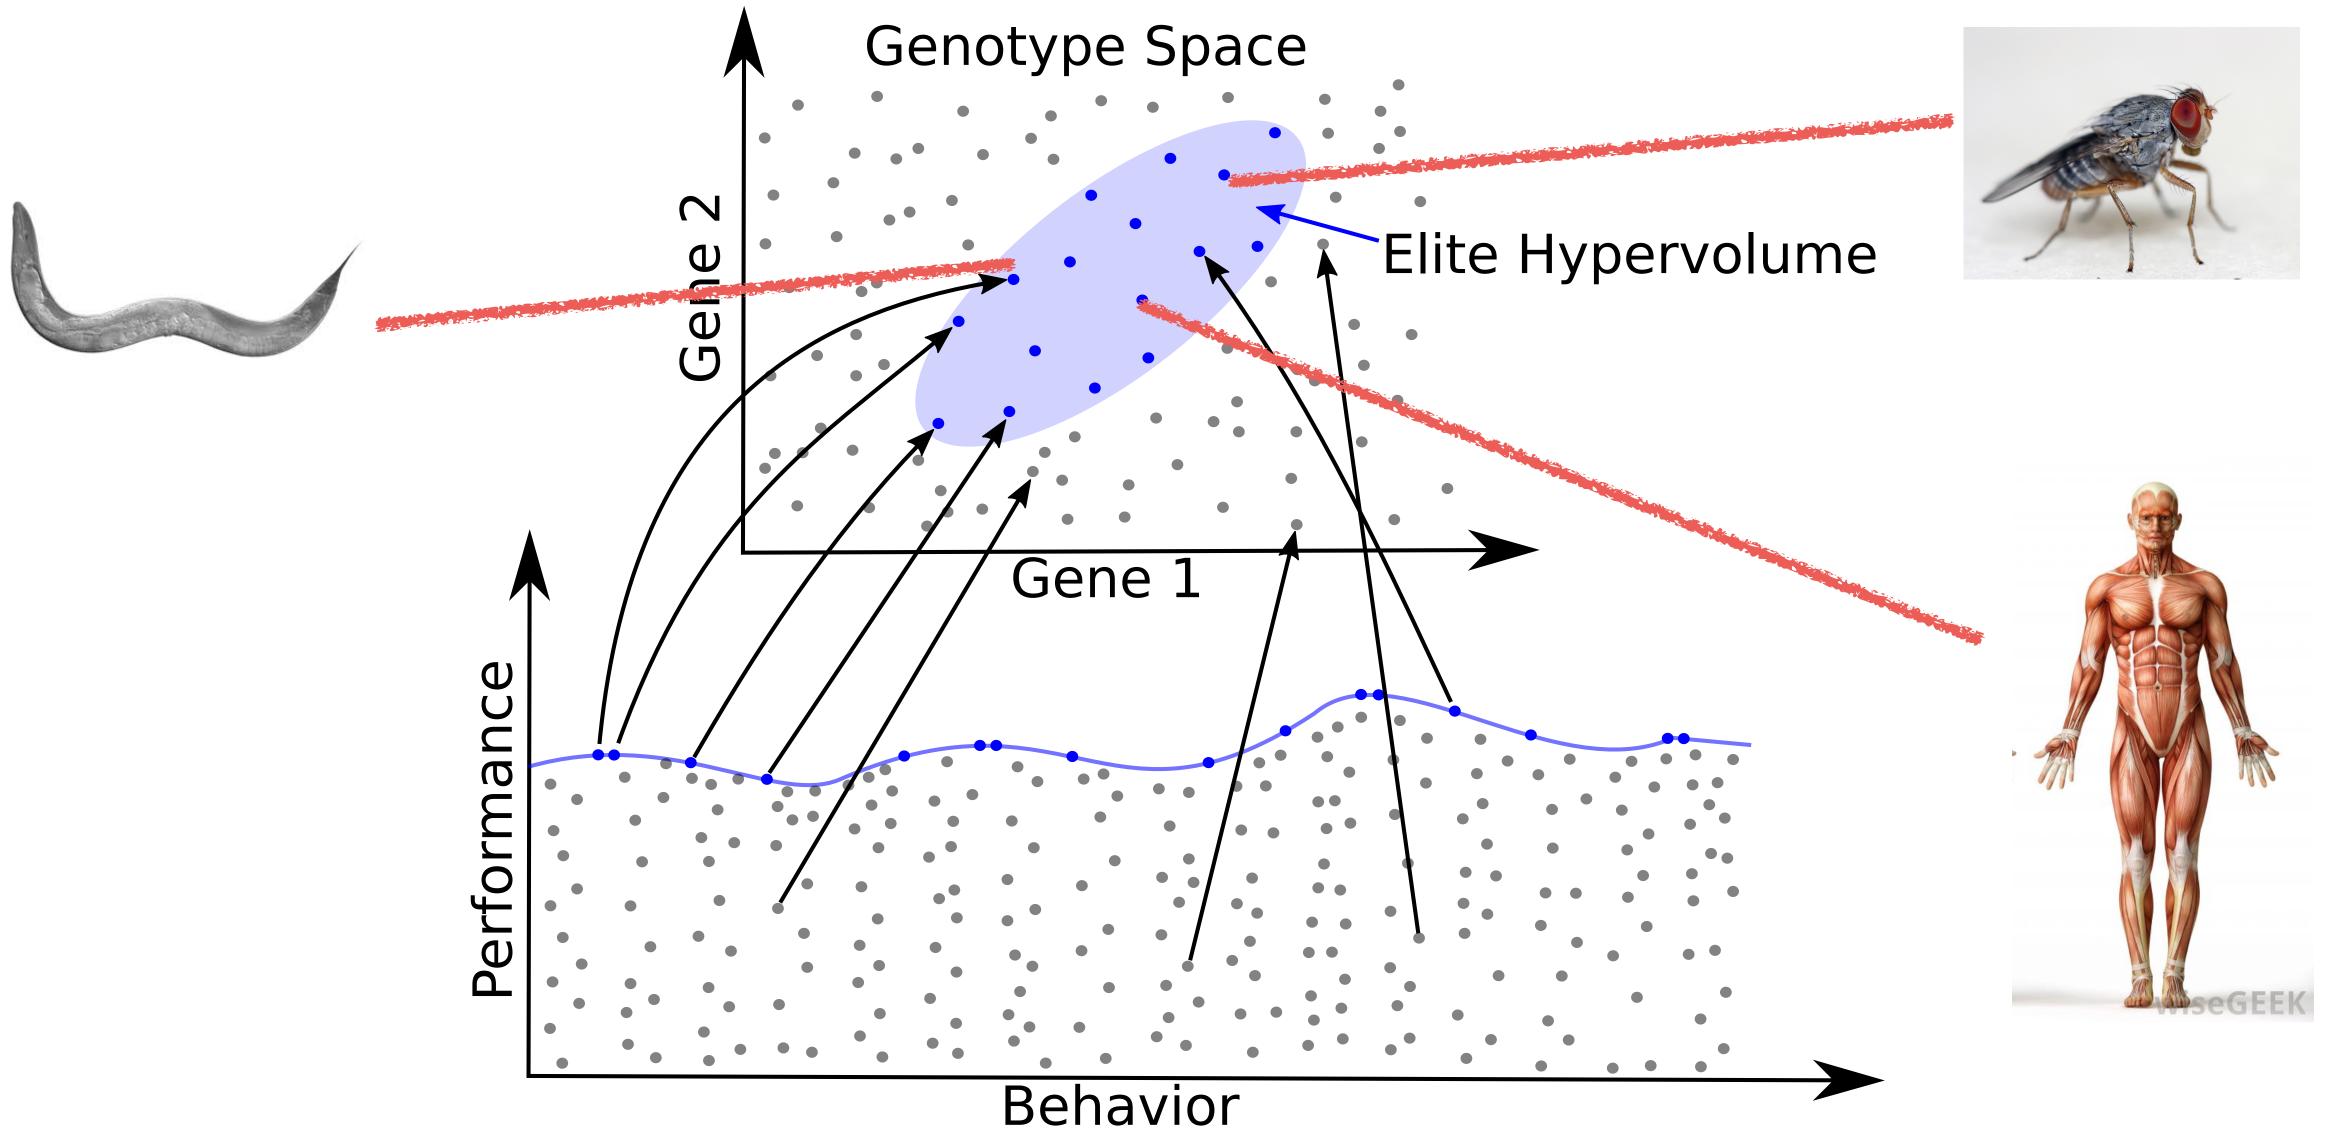

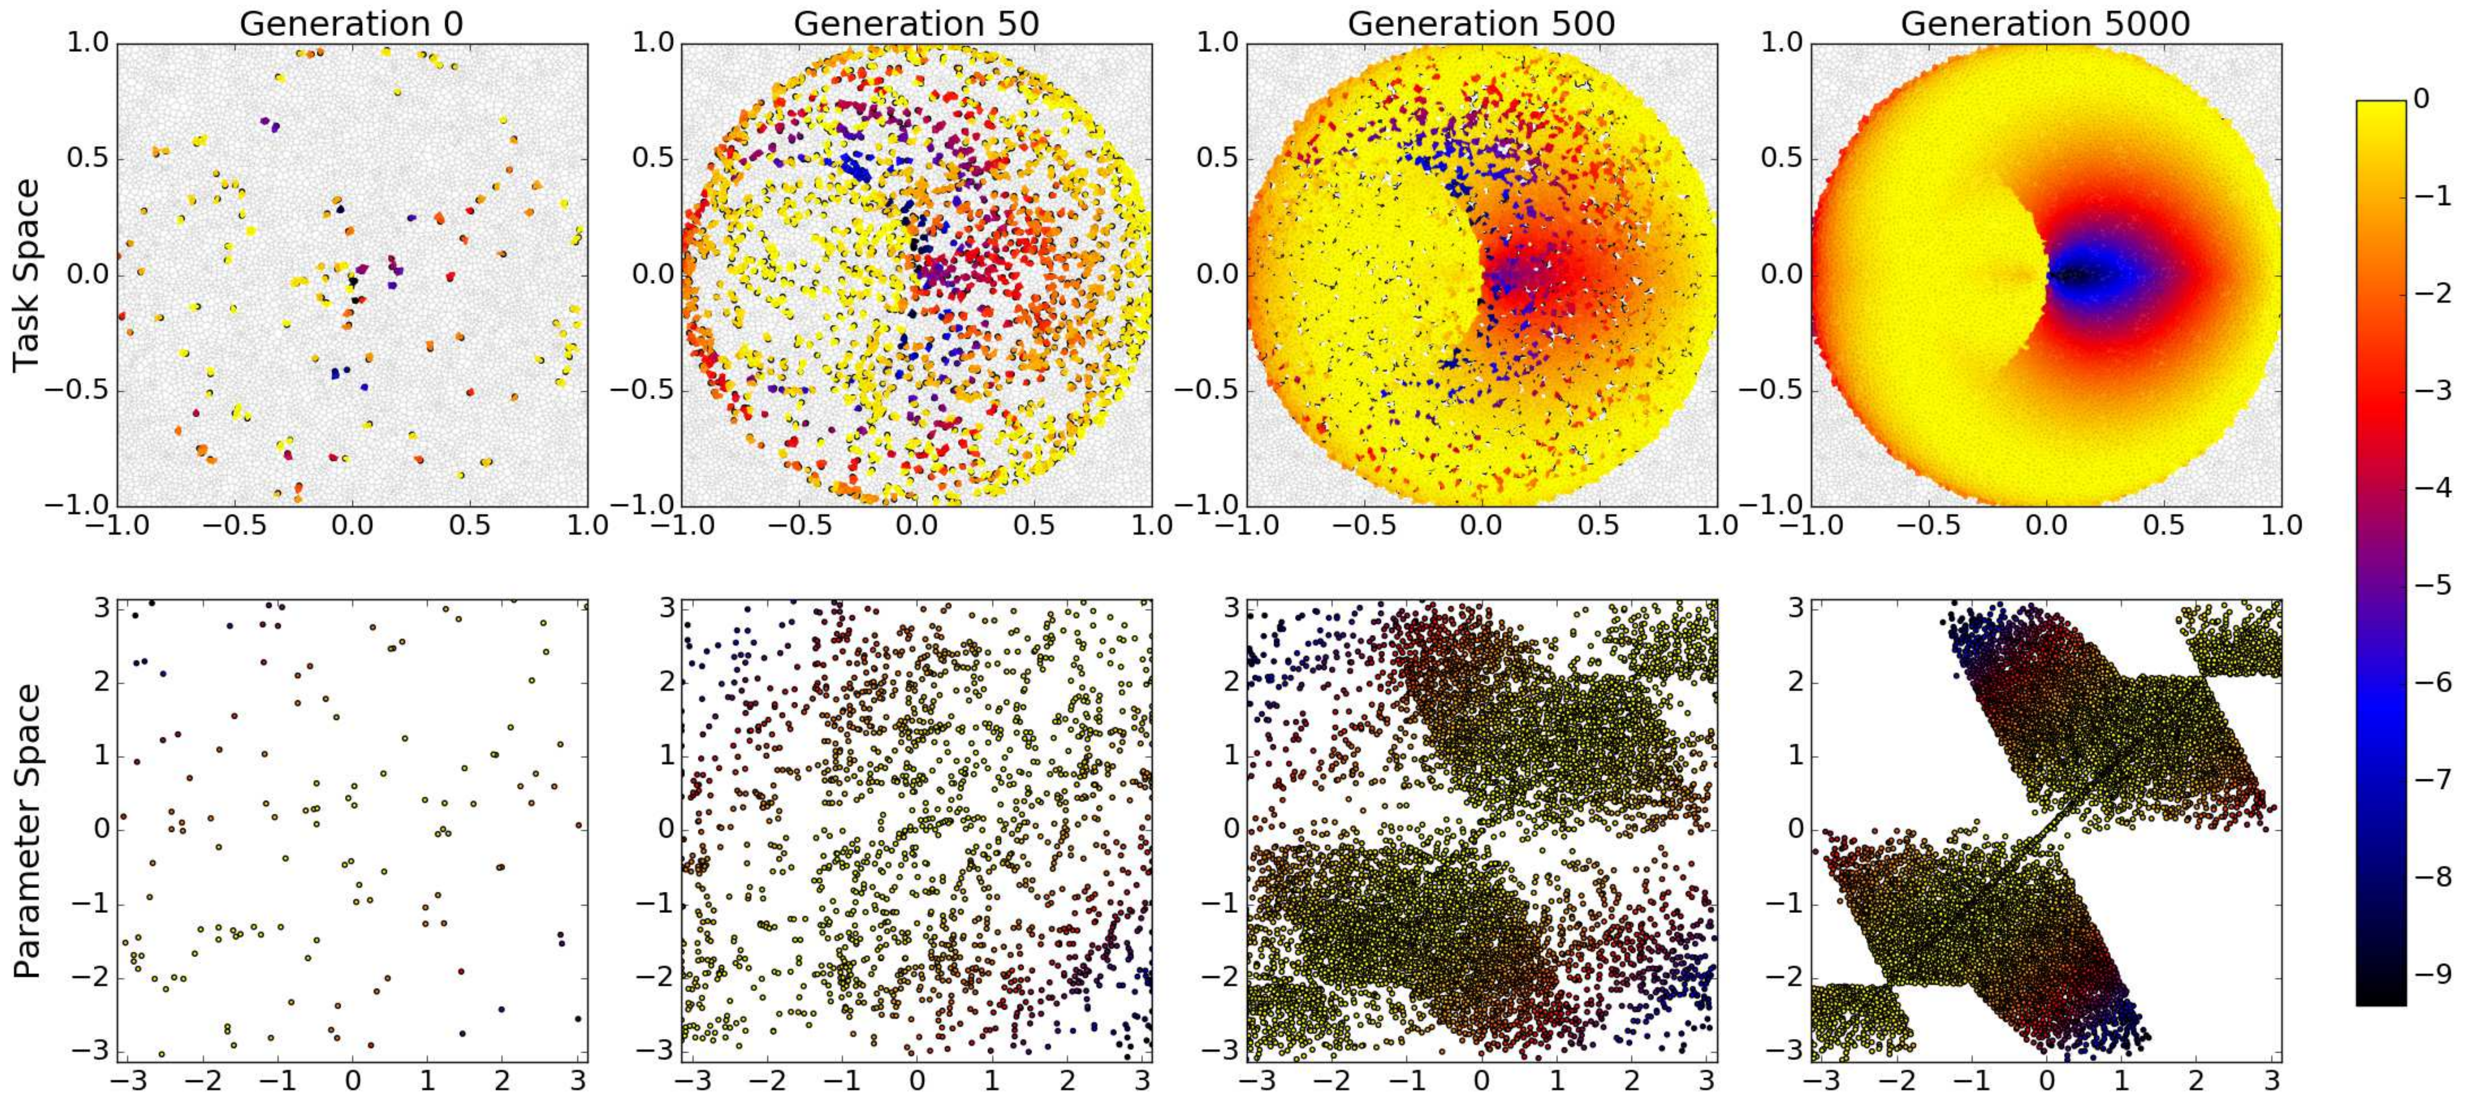

Human and fruit flies: 60% of common genes!

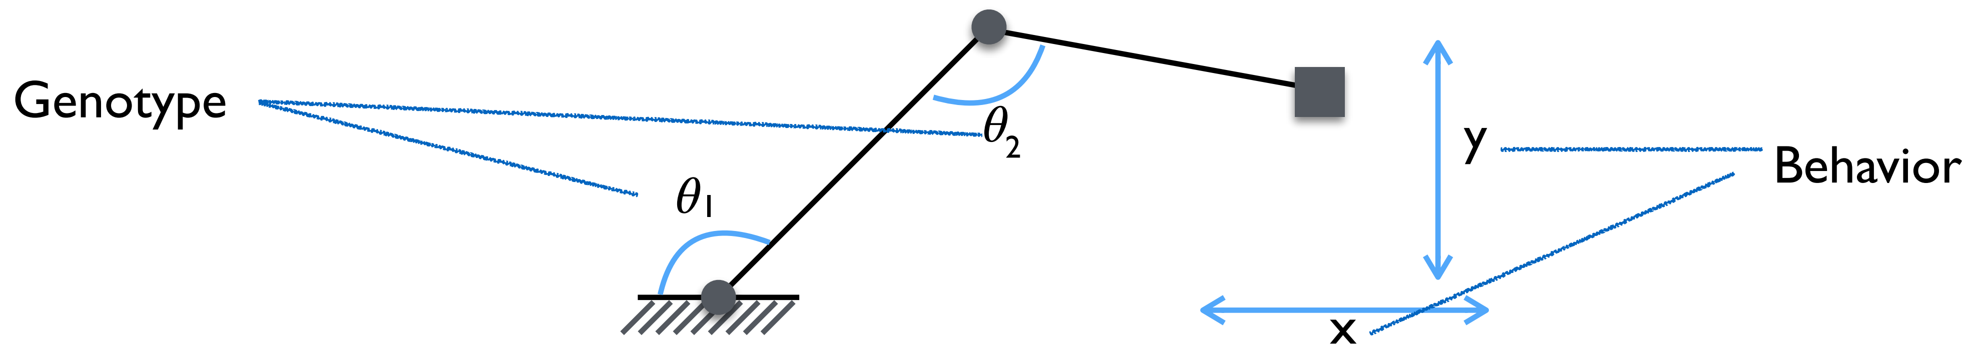

V. Vassiliades & J.-B. Mouret (2018). Discovering the Elite Hypervolume by Leveraging Interspecies Correlation. Proc. of GECCO.
Adams, M. D.,et al. (2000). The genome sequence of Drosophila melanogaster. *Science*, *287*(5461), 2185-2195.

MAP-Elites: elite hypervolume

Good solutions share common “recipes"

- What is a good variation operator?

➔ highly likely to generate an individual in the elite hypervolume

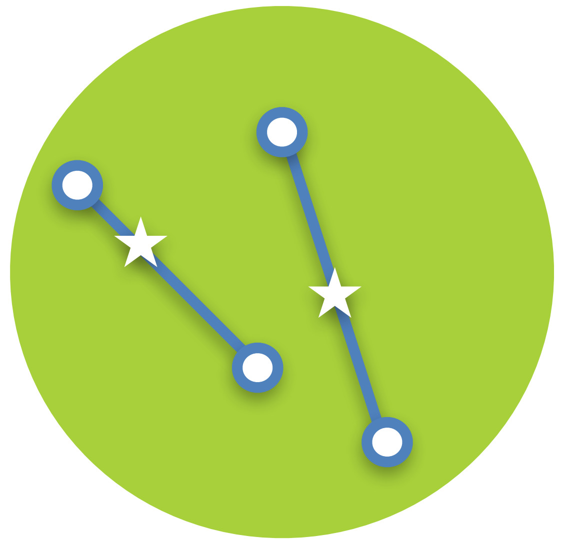

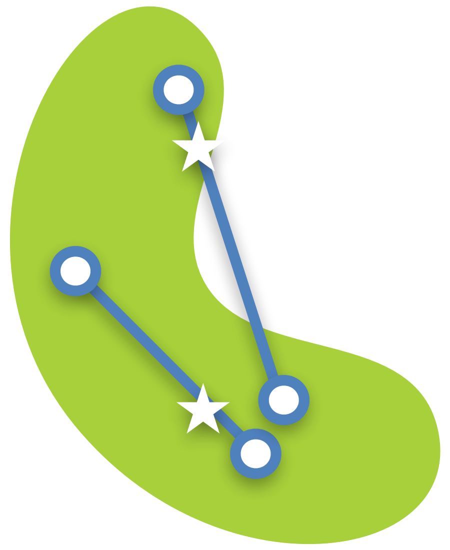

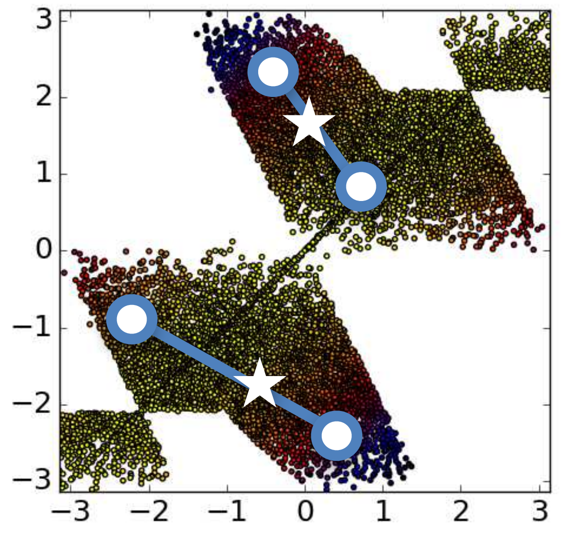

- if we take two points from a convex volume, any point on the segment is in the volume too

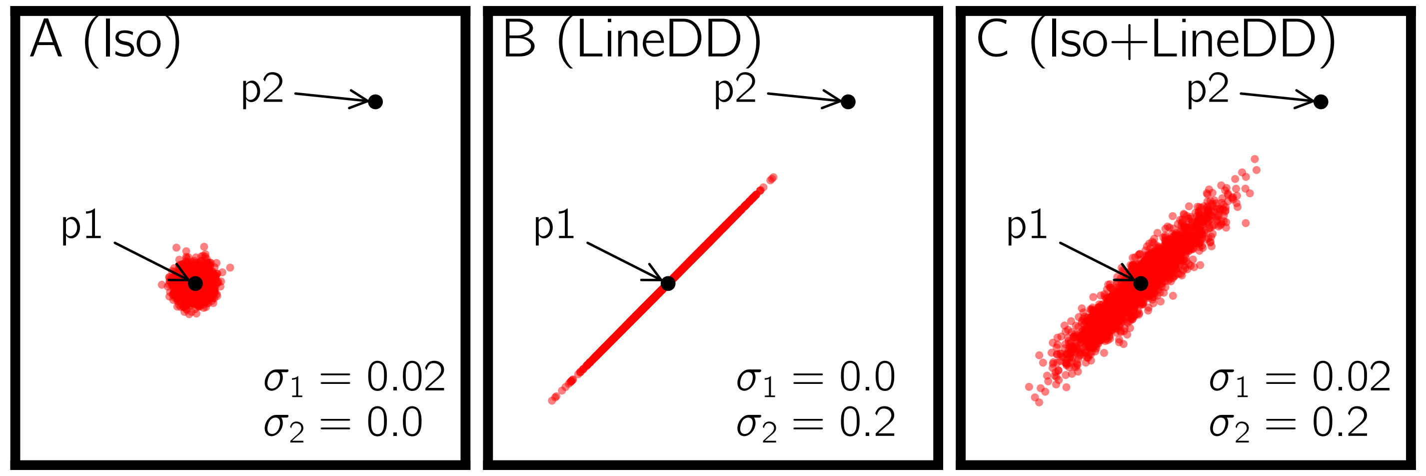

$$\mathbf{x}_i^{(t+1)} = \mathbf{x}_i^{(t)} + \sigma_1 \mathcal{N}(\mathbf{0}, \mathbf{I}) + \sigma_2(\mathbf{x}_j^{(t)} - \mathbf{x}_i^{(t)})\mathcal{N}(0, 1)$$

random perturbation for each dimension

weight a single perturbation by the difference between the parents

➔ Directional variation (~ cross-over)

➔ Adapts the step size

V. Vassiliades & J.-B. Mouret (2018). Discovering the Elite Hypervolume by Leveraging Interspecies Correlation. Proc. of GECCO.

Adams, M. D.,et al. (2000). The genome sequence of Drosophila melanogaster. *Science*, *287*(5461), 2185-2195.

MAP-Elites: elite hypervolume

Notebook

```
# random initialisation
archive = init_archive(num_random)

for j in tqdm(range(0, num_iterations)):
    display_archive(archive)
    for i in range(0, batch_size):
        # pick TWO points the archive
        x = random.choice(list(archive.values())).genotype
        y = random.choice(list(archive.values())).genotype
        # mutate it (we can use more advanced variation techniques: NEAT, etc.)
        # be careful: the first random.normal() is a 1-d normal
        g = x + (x - y) * np.random.normal(0, 0.1) + np.random.normal(0, 0.025, x.shape[0])
        # compute the fitness
        f, b = fitness(g)
        # add to archive
        add_to_archive(archive, Species(g, b, f))
```

$$\mathbf{x}_i^{(t+1)} = \mathbf{x}_i^{(t)} + \sigma_1 \mathcal{N}(\mathbf{0}, \mathbf{I}) + \sigma_2(\mathbf{x}_j^{(t)} - \mathbf{x}_i^{(t)})\mathcal{N}(0,1)$$

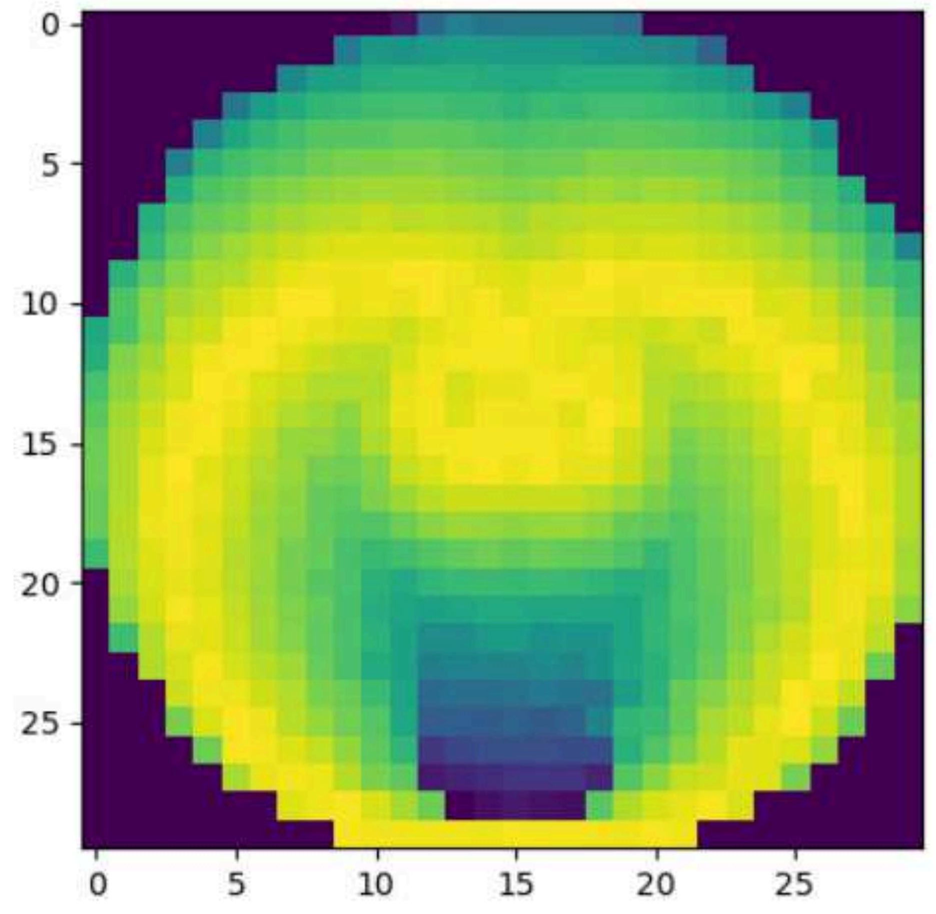

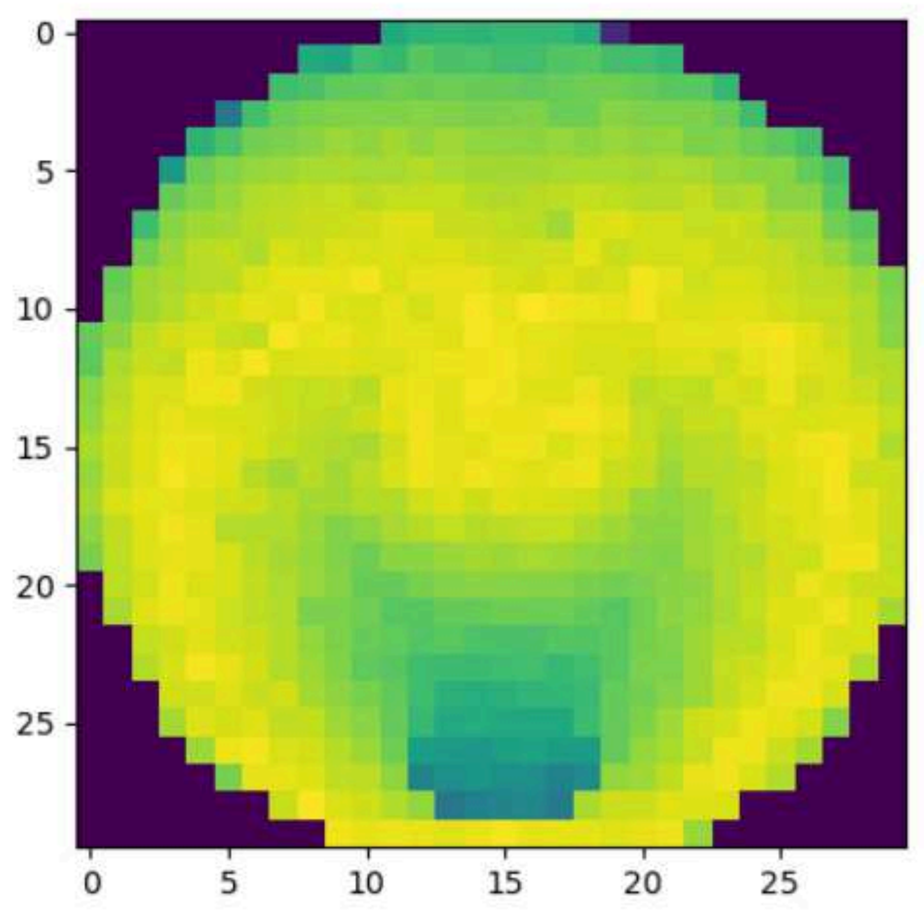

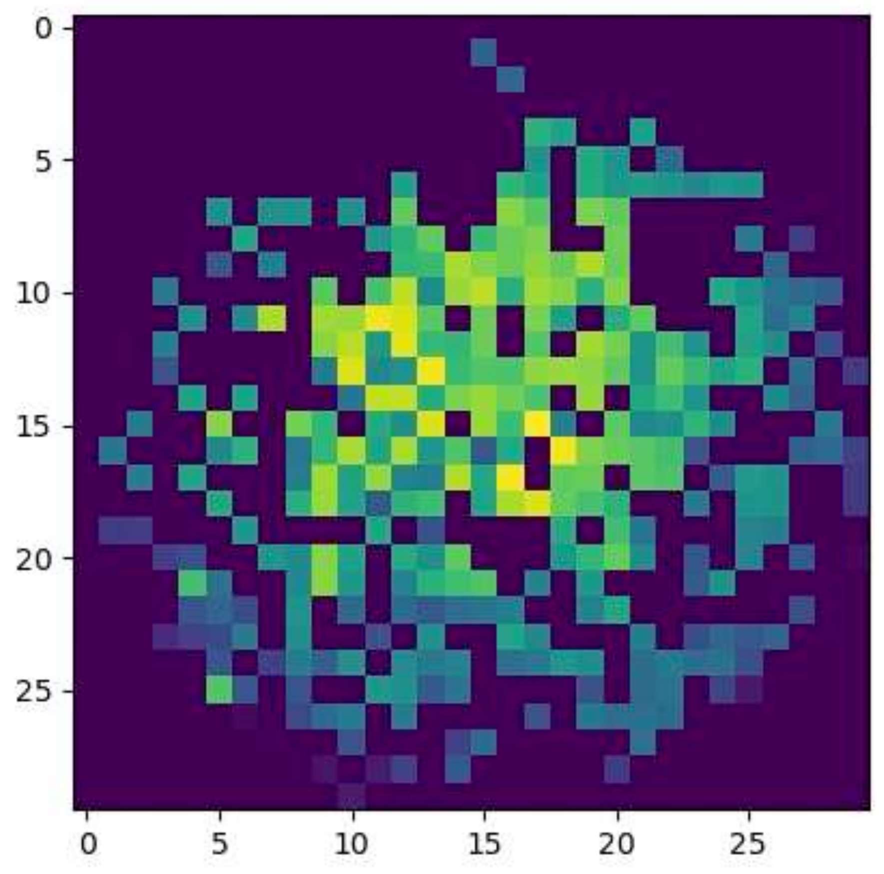

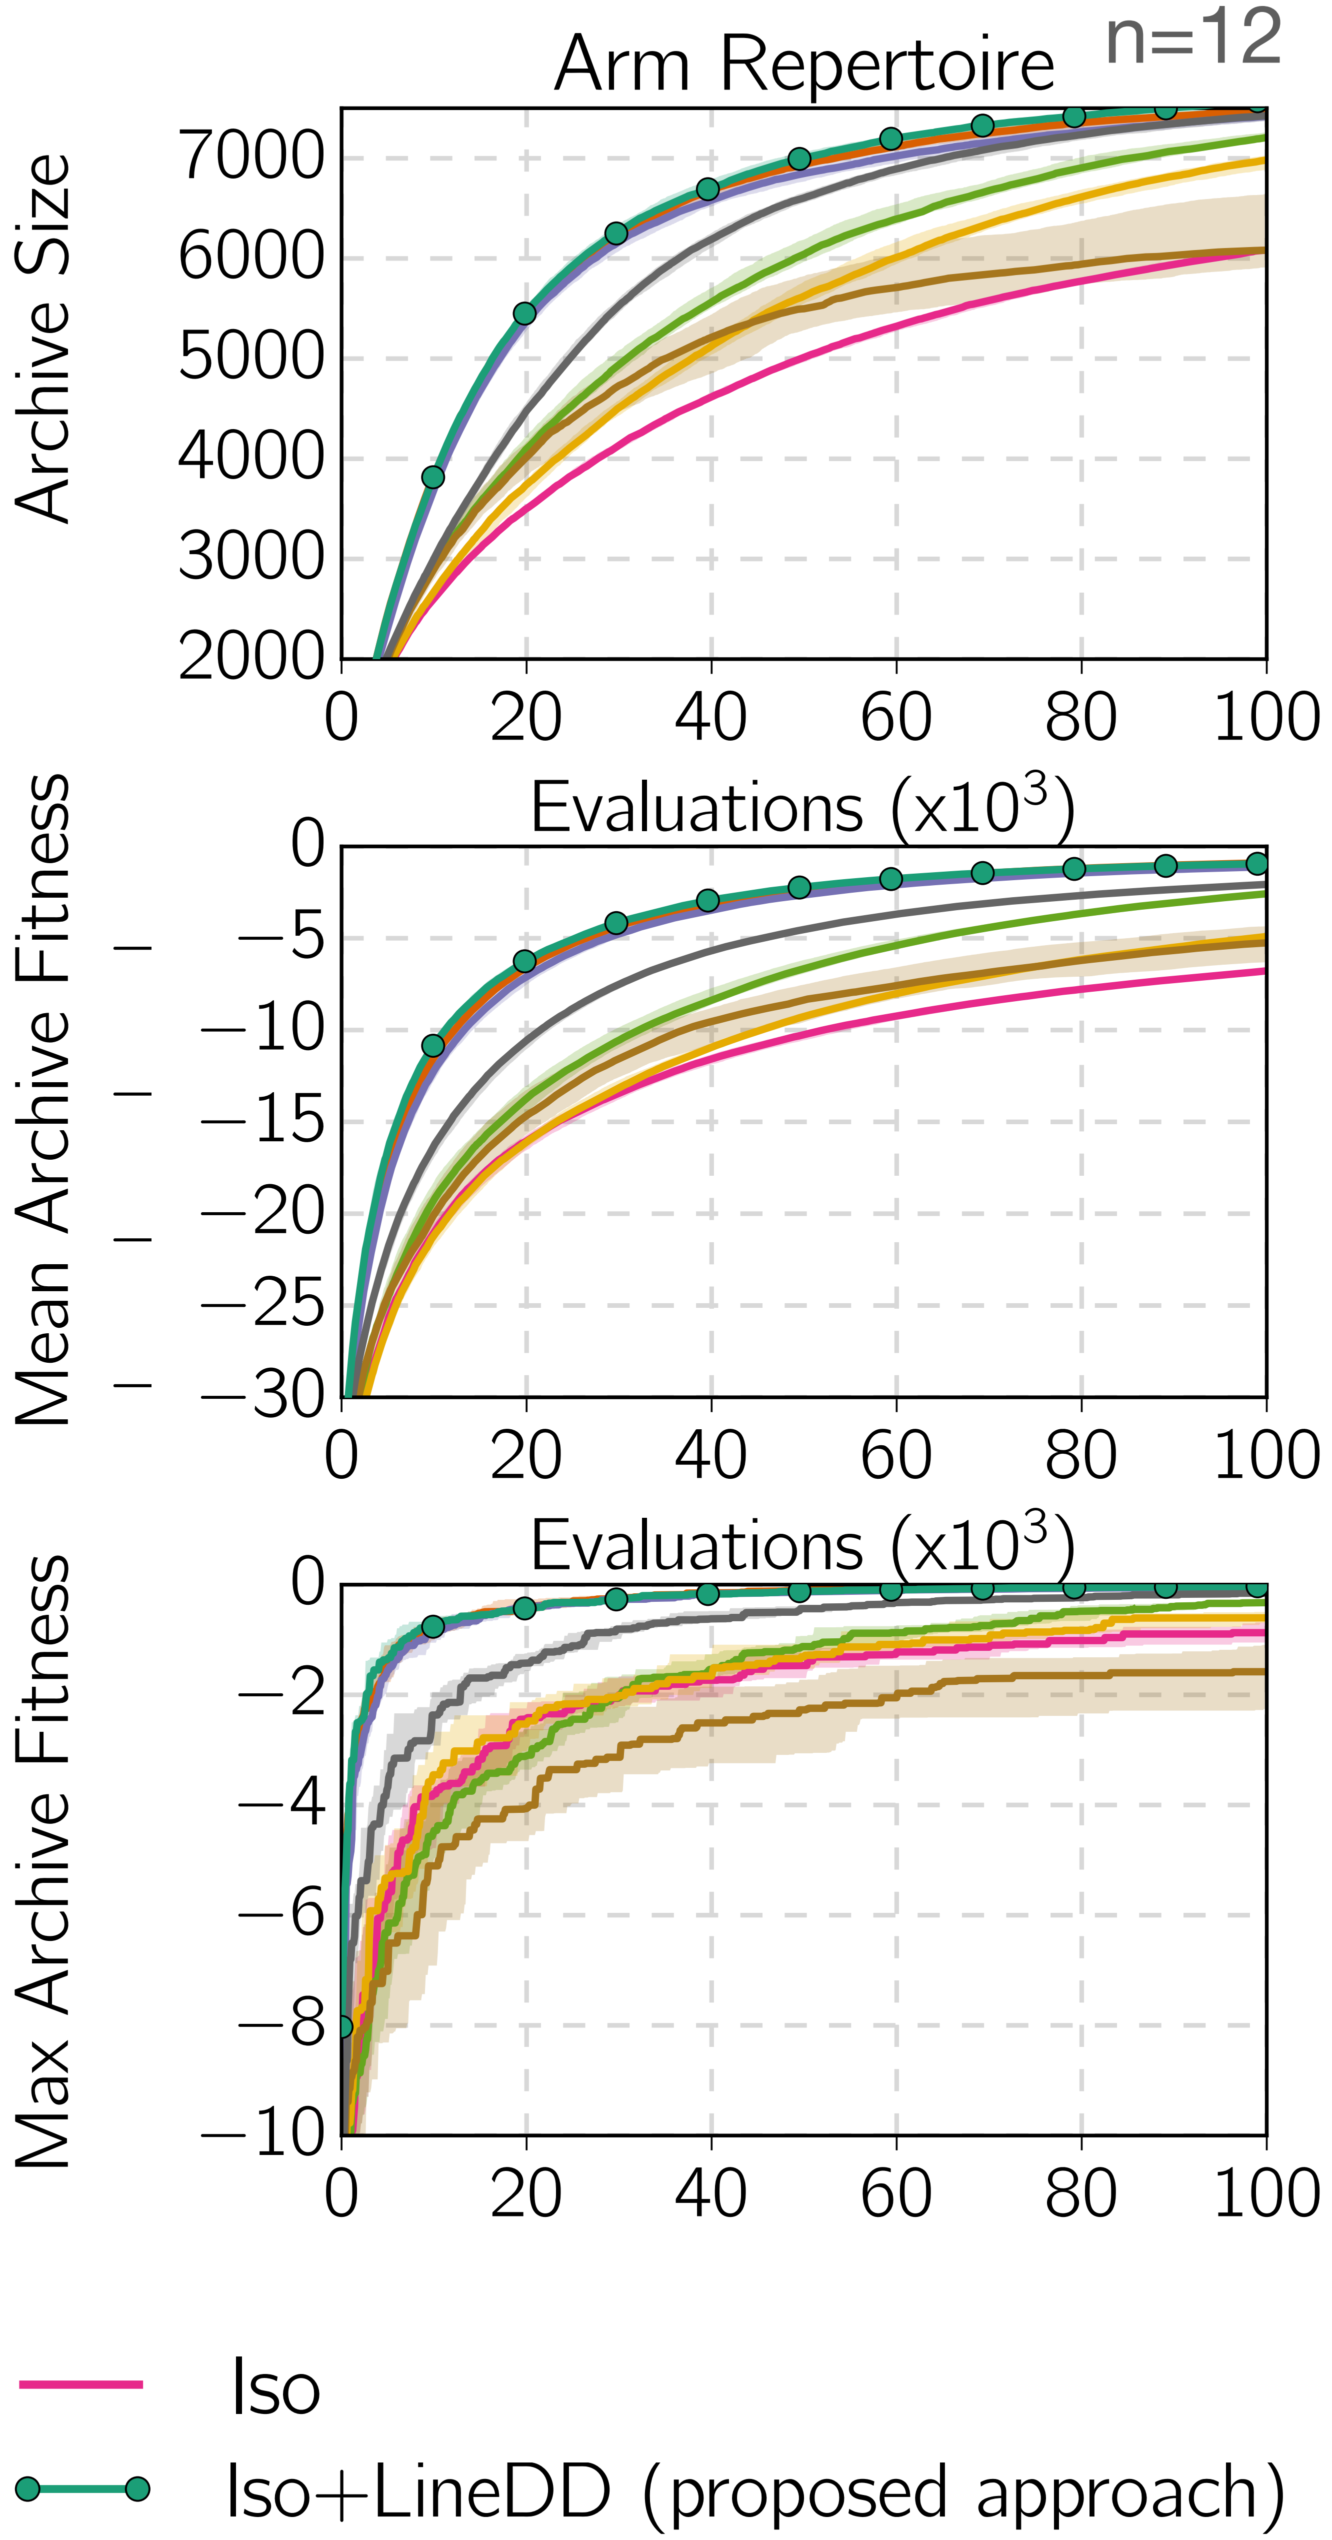

V. Vassiliades & J.-B. Mouret (2018). Discovering the Elite Hypervolume by Leveraging Interspecies Correlation. Proc. of GECCO.

Link with GO-Explore

- Exploration concept inspired by MAP-Elites
  - behavior = state traversed
  - keep fastest to reach the state in a archive/map

- Additions:
  - learn policies from the state sequence (robustification)
  - select cells from the map with weights + other heuristics

$$W = \frac{1}{\sqrt{C_{\text{seen}} + 1}},$$

- Best results in the hard games (Montezuma revenge)

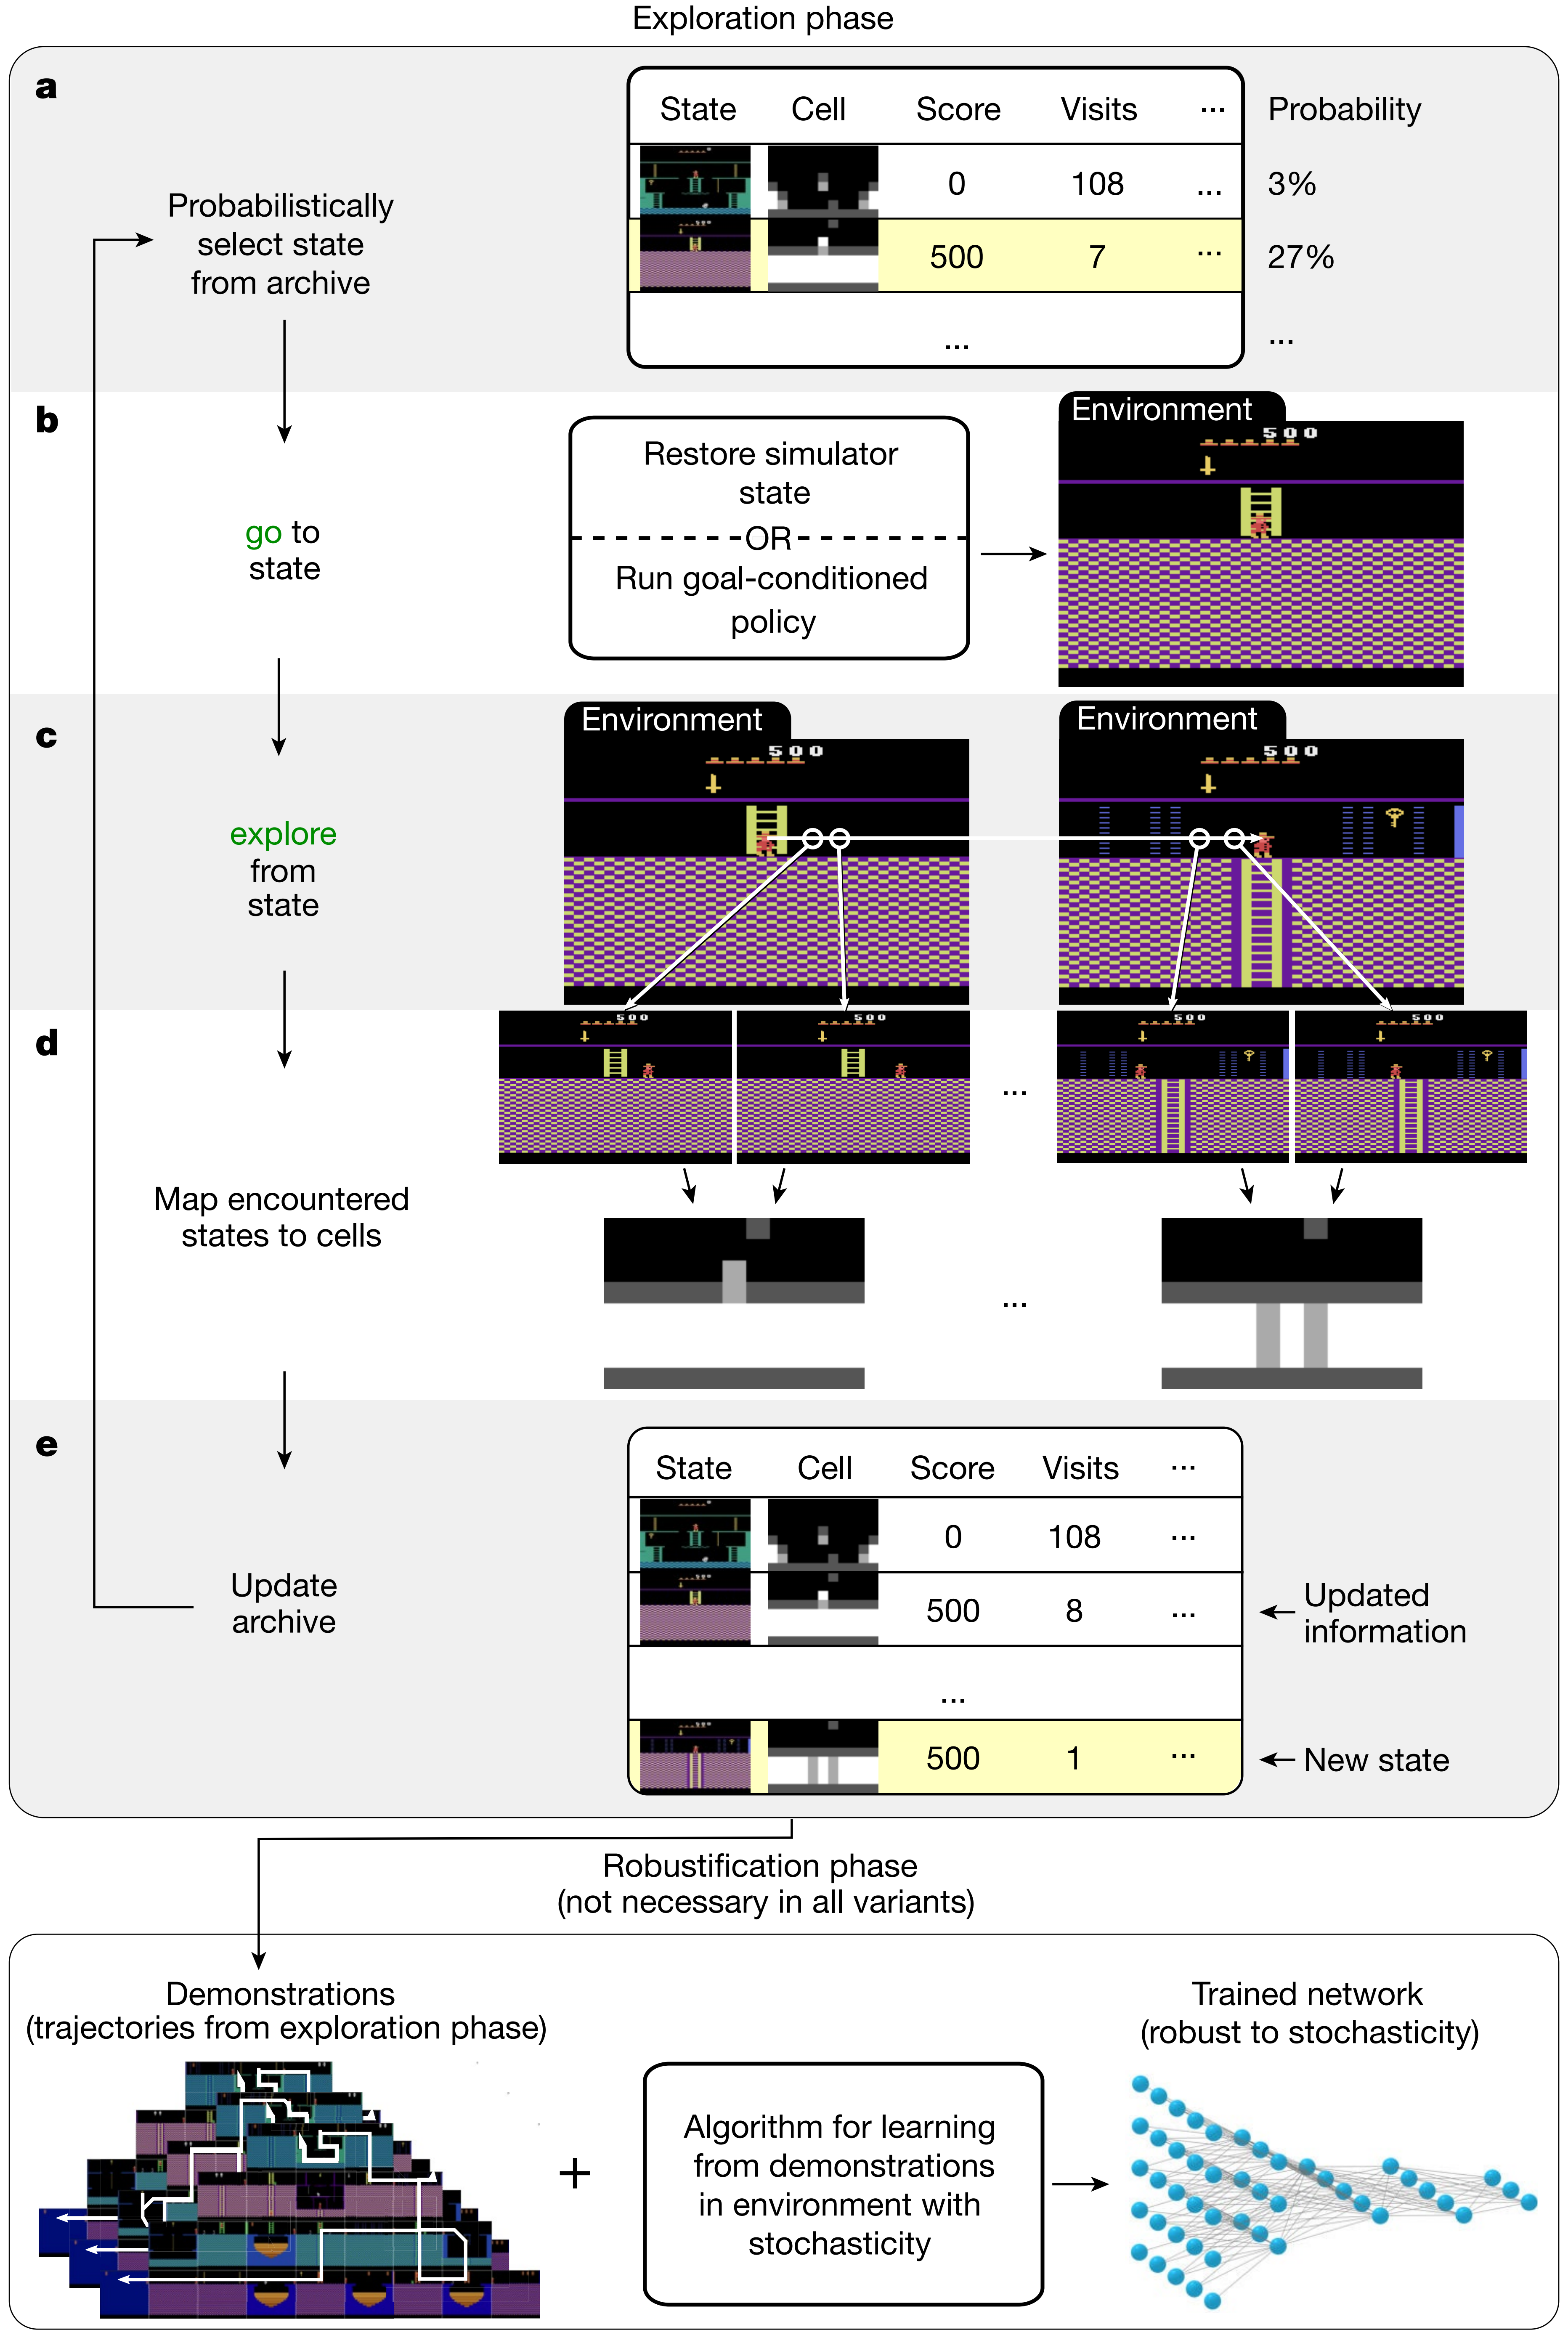

Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., & Clune, J. (2021). First return, then explore. *Nature*, *590*(7847), 580-586.

Scaling up to high-dimensional behavioral spaces:

- **centroidal Voronoi tessellation to split the volume in cells**
  - ➔ **Vassiliades, V., Chatzilygeroudis, K., & Mouret, J. B. (2017).** Using Centroidal Voronoi Tessellations to Scale Up the Multi-dimensional Archive of Phenotypic Elites Algorithm. *IEEE Transactions on Evolutionary Computation*, 9.
- **distance-based archive**
  - ➔ **Cully, A., & Demiris, Y. (2017).** Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation*, *22*(2), 245-259.
- **VAE-based dimensionality reduction**
  - ➔ **Cully, A. (2019).** Autonomous skill discovery with quality-diversity and unsupervised descriptors. In Proceedings of the Genetic and Evolutionary Computation Conference (pp. 81-89).

MAP-Elites

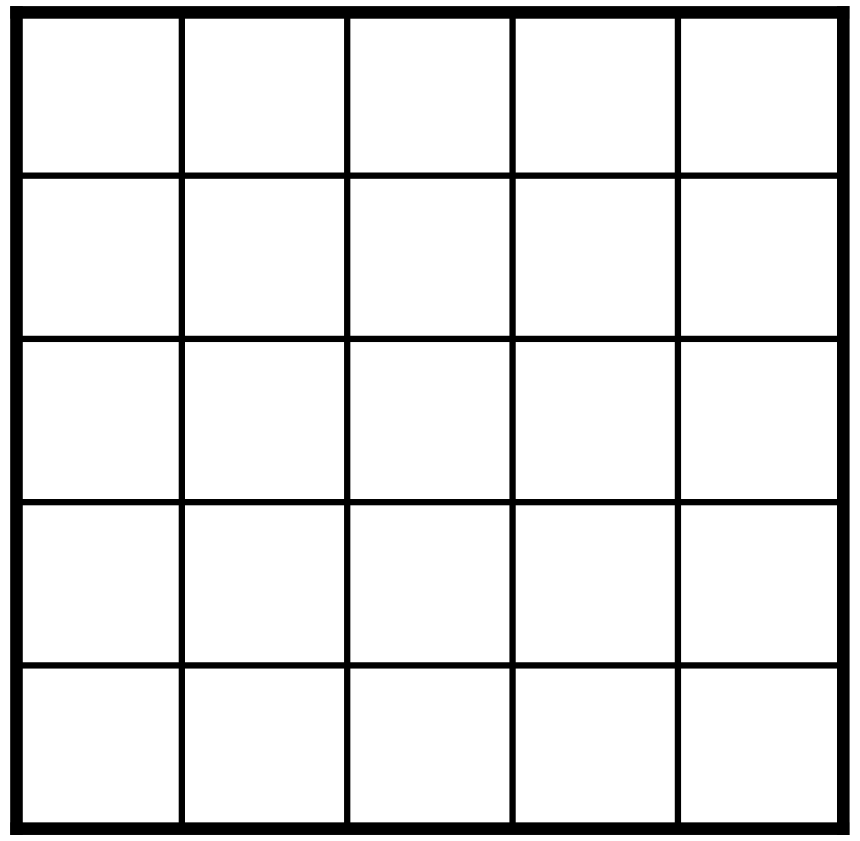

CVT-MAP-Elites

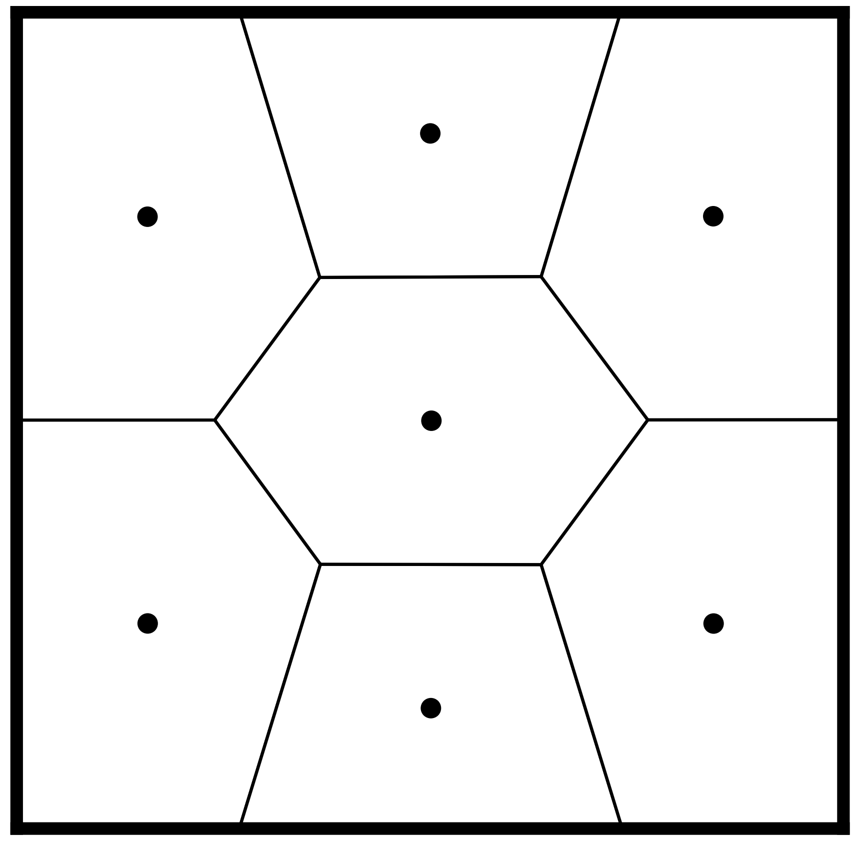

Scaling up to high-dimensional genotypes/search spaces

- **learn the hypervolume with a VAE**
  - ➔ **Gaier, A., Asteroth, A., & Mouret, J. B. (2020).** Discovering representations for black-box optimization. In Proceedings of the 2020 Genetic and Evolutionary Computation Conference (pp. 103-111).
- **take inspiration from OpenAI-ES**
  - ➔ **Colas, C., Madhavan, V., Huizinga, J., & Clune, J. (2020).** Scaling map-elites to deep neuroevolution. In Proceedings of the 2020 Genetic and Evolutionary Computation Conference (pp. 67-75).

Improve data-efficiency

- **with surrogate models**
  - ➔ **Gaier, A., Asteroth, A., & Mouret, J. B. (2018).** Data-efficient design exploration through surrogate-assisted illumination. *Evolutionary computation*, *26*(3), 381-410.
- **taking inspiration from CMA-ES**

➔ **Fontaine, M. C., Togelius, J., Nikolaidis, S., & Hoover, A. K. (2020)**. Covariance matrix adaptation for the rapid illumination of behavior space. In *Proceedings of the 2020 genetic and evolutionary computation conference* (pp. 94-102).

Example: many ways of walking

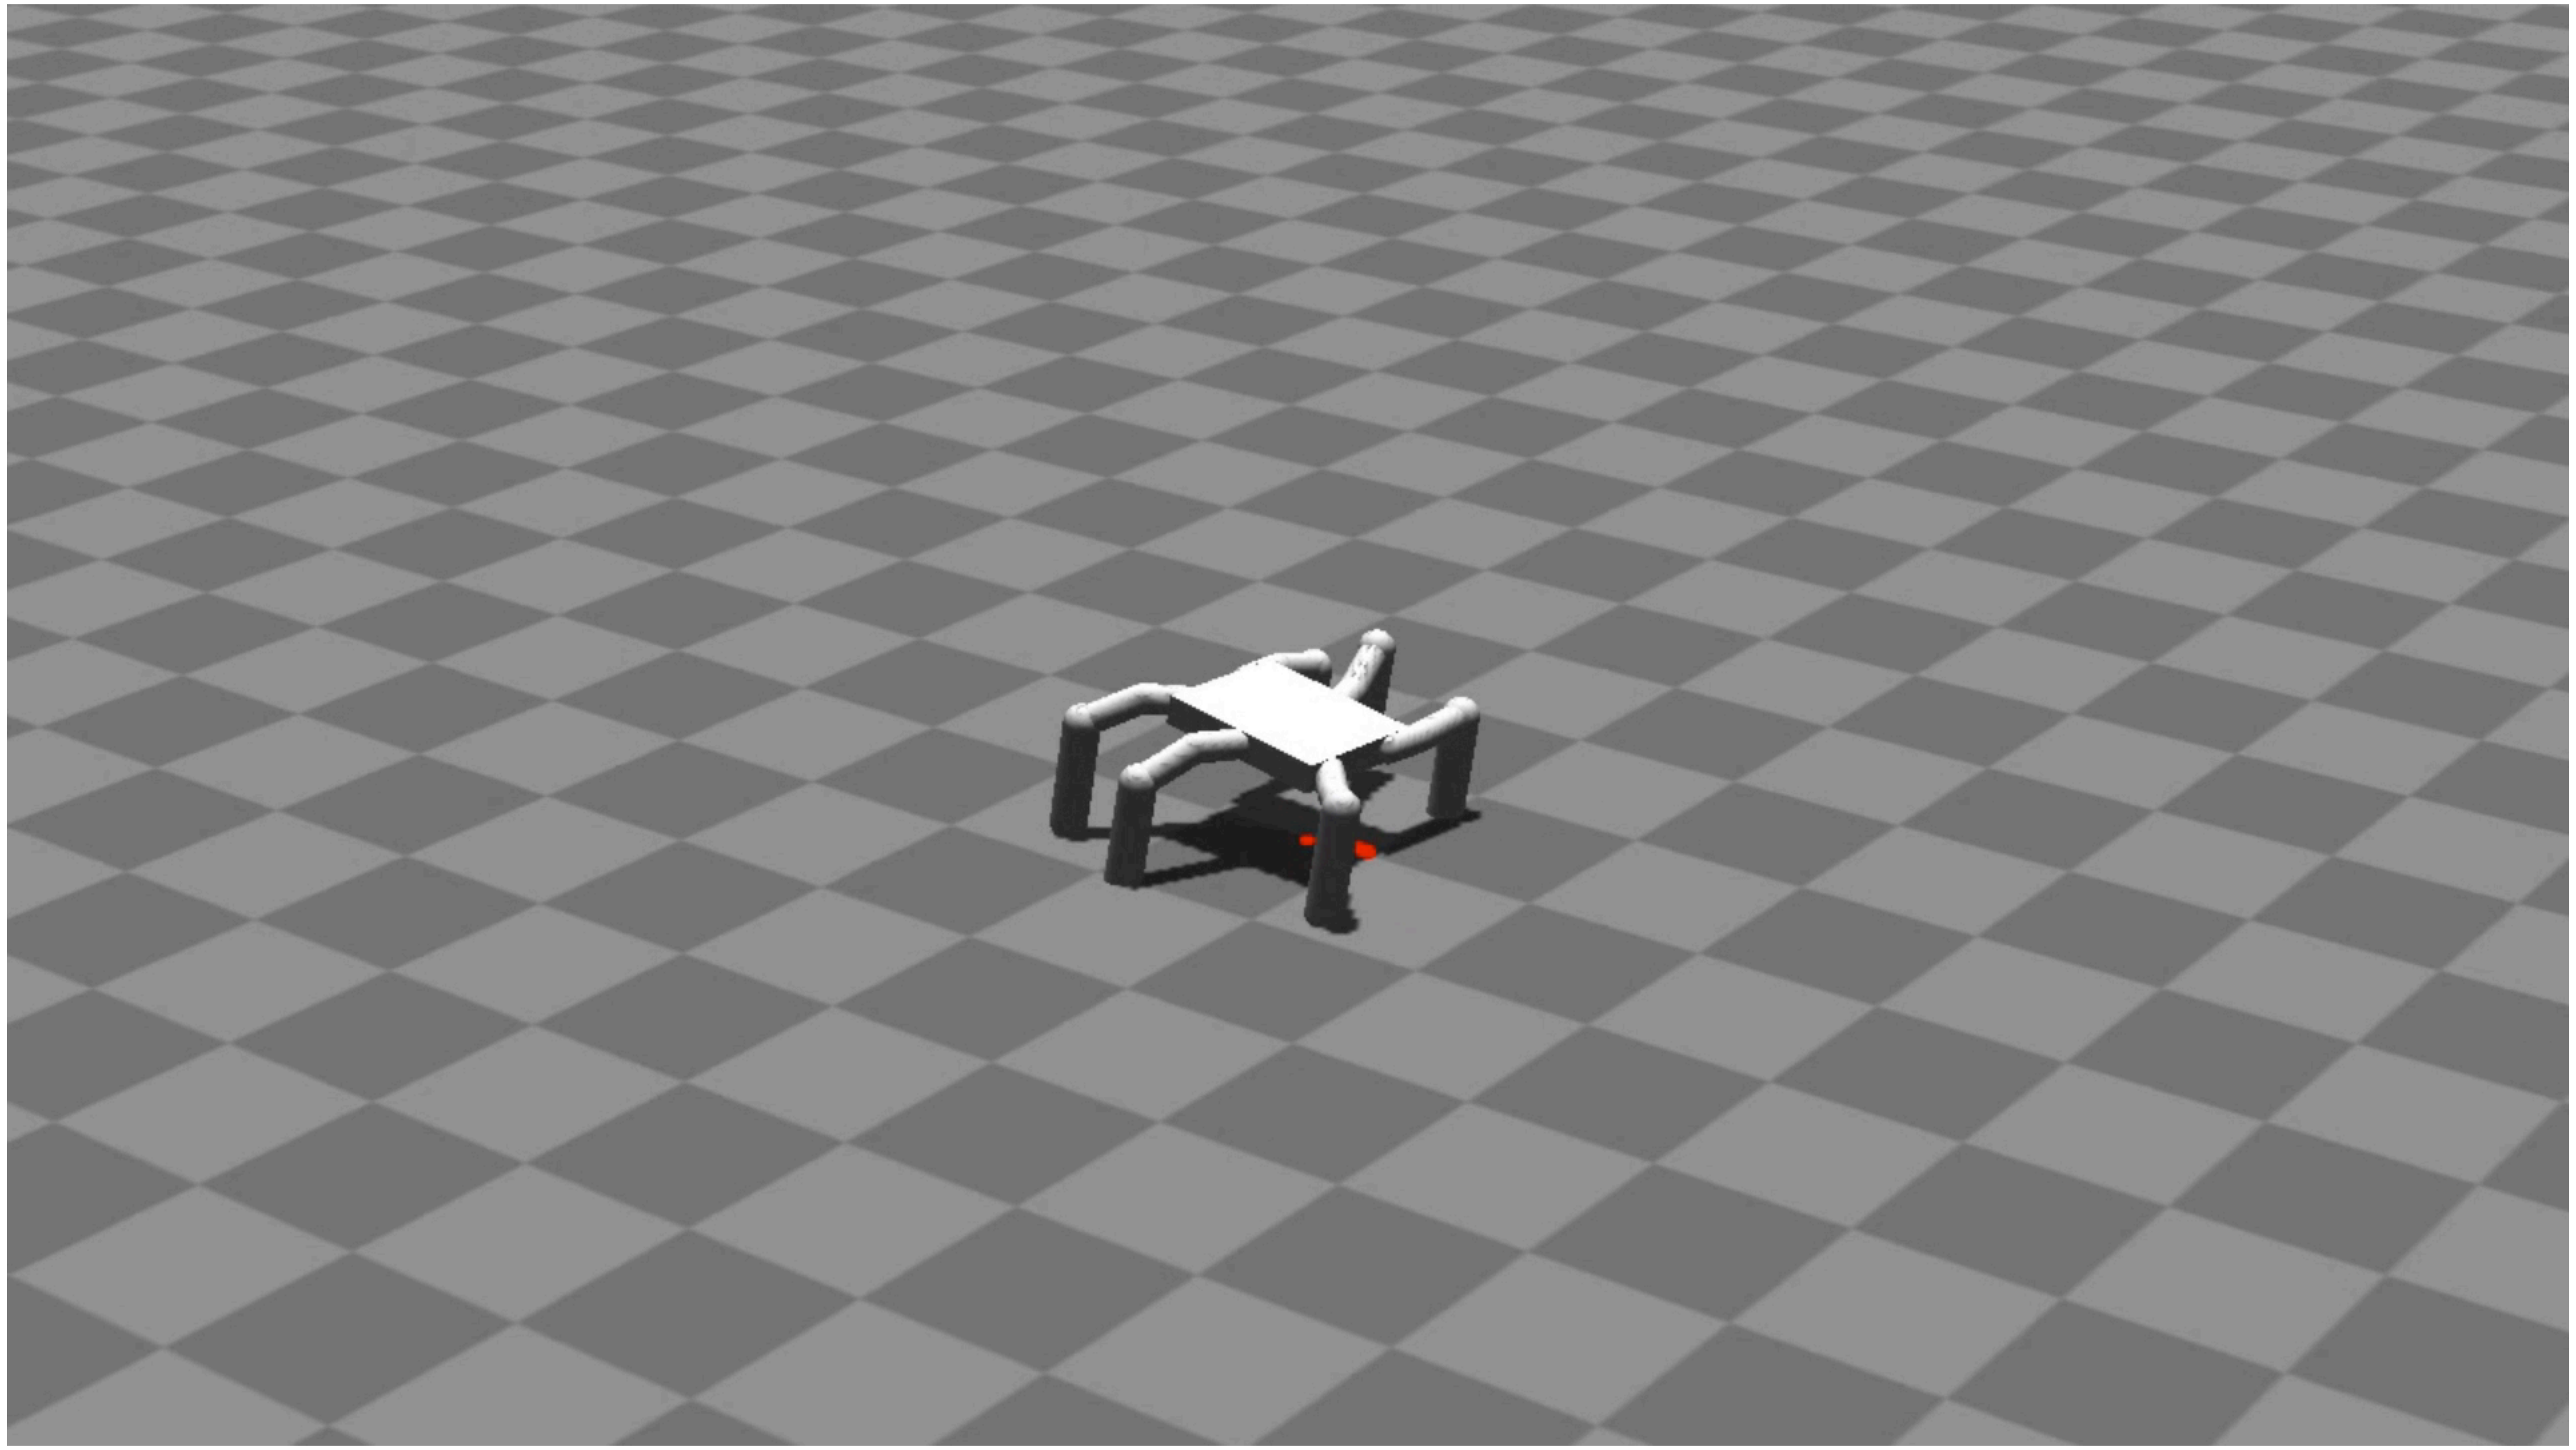

Search space: 36 parameters
(open-loop controller)

Behavior space: 6-D
(% of contact for each foot)

Cully A, Clune J, Tarapore D, Mouret JB. Robots that can adapt like animals. Nature. 2015 May;521(7553):503-7.

Example: many ways of walking

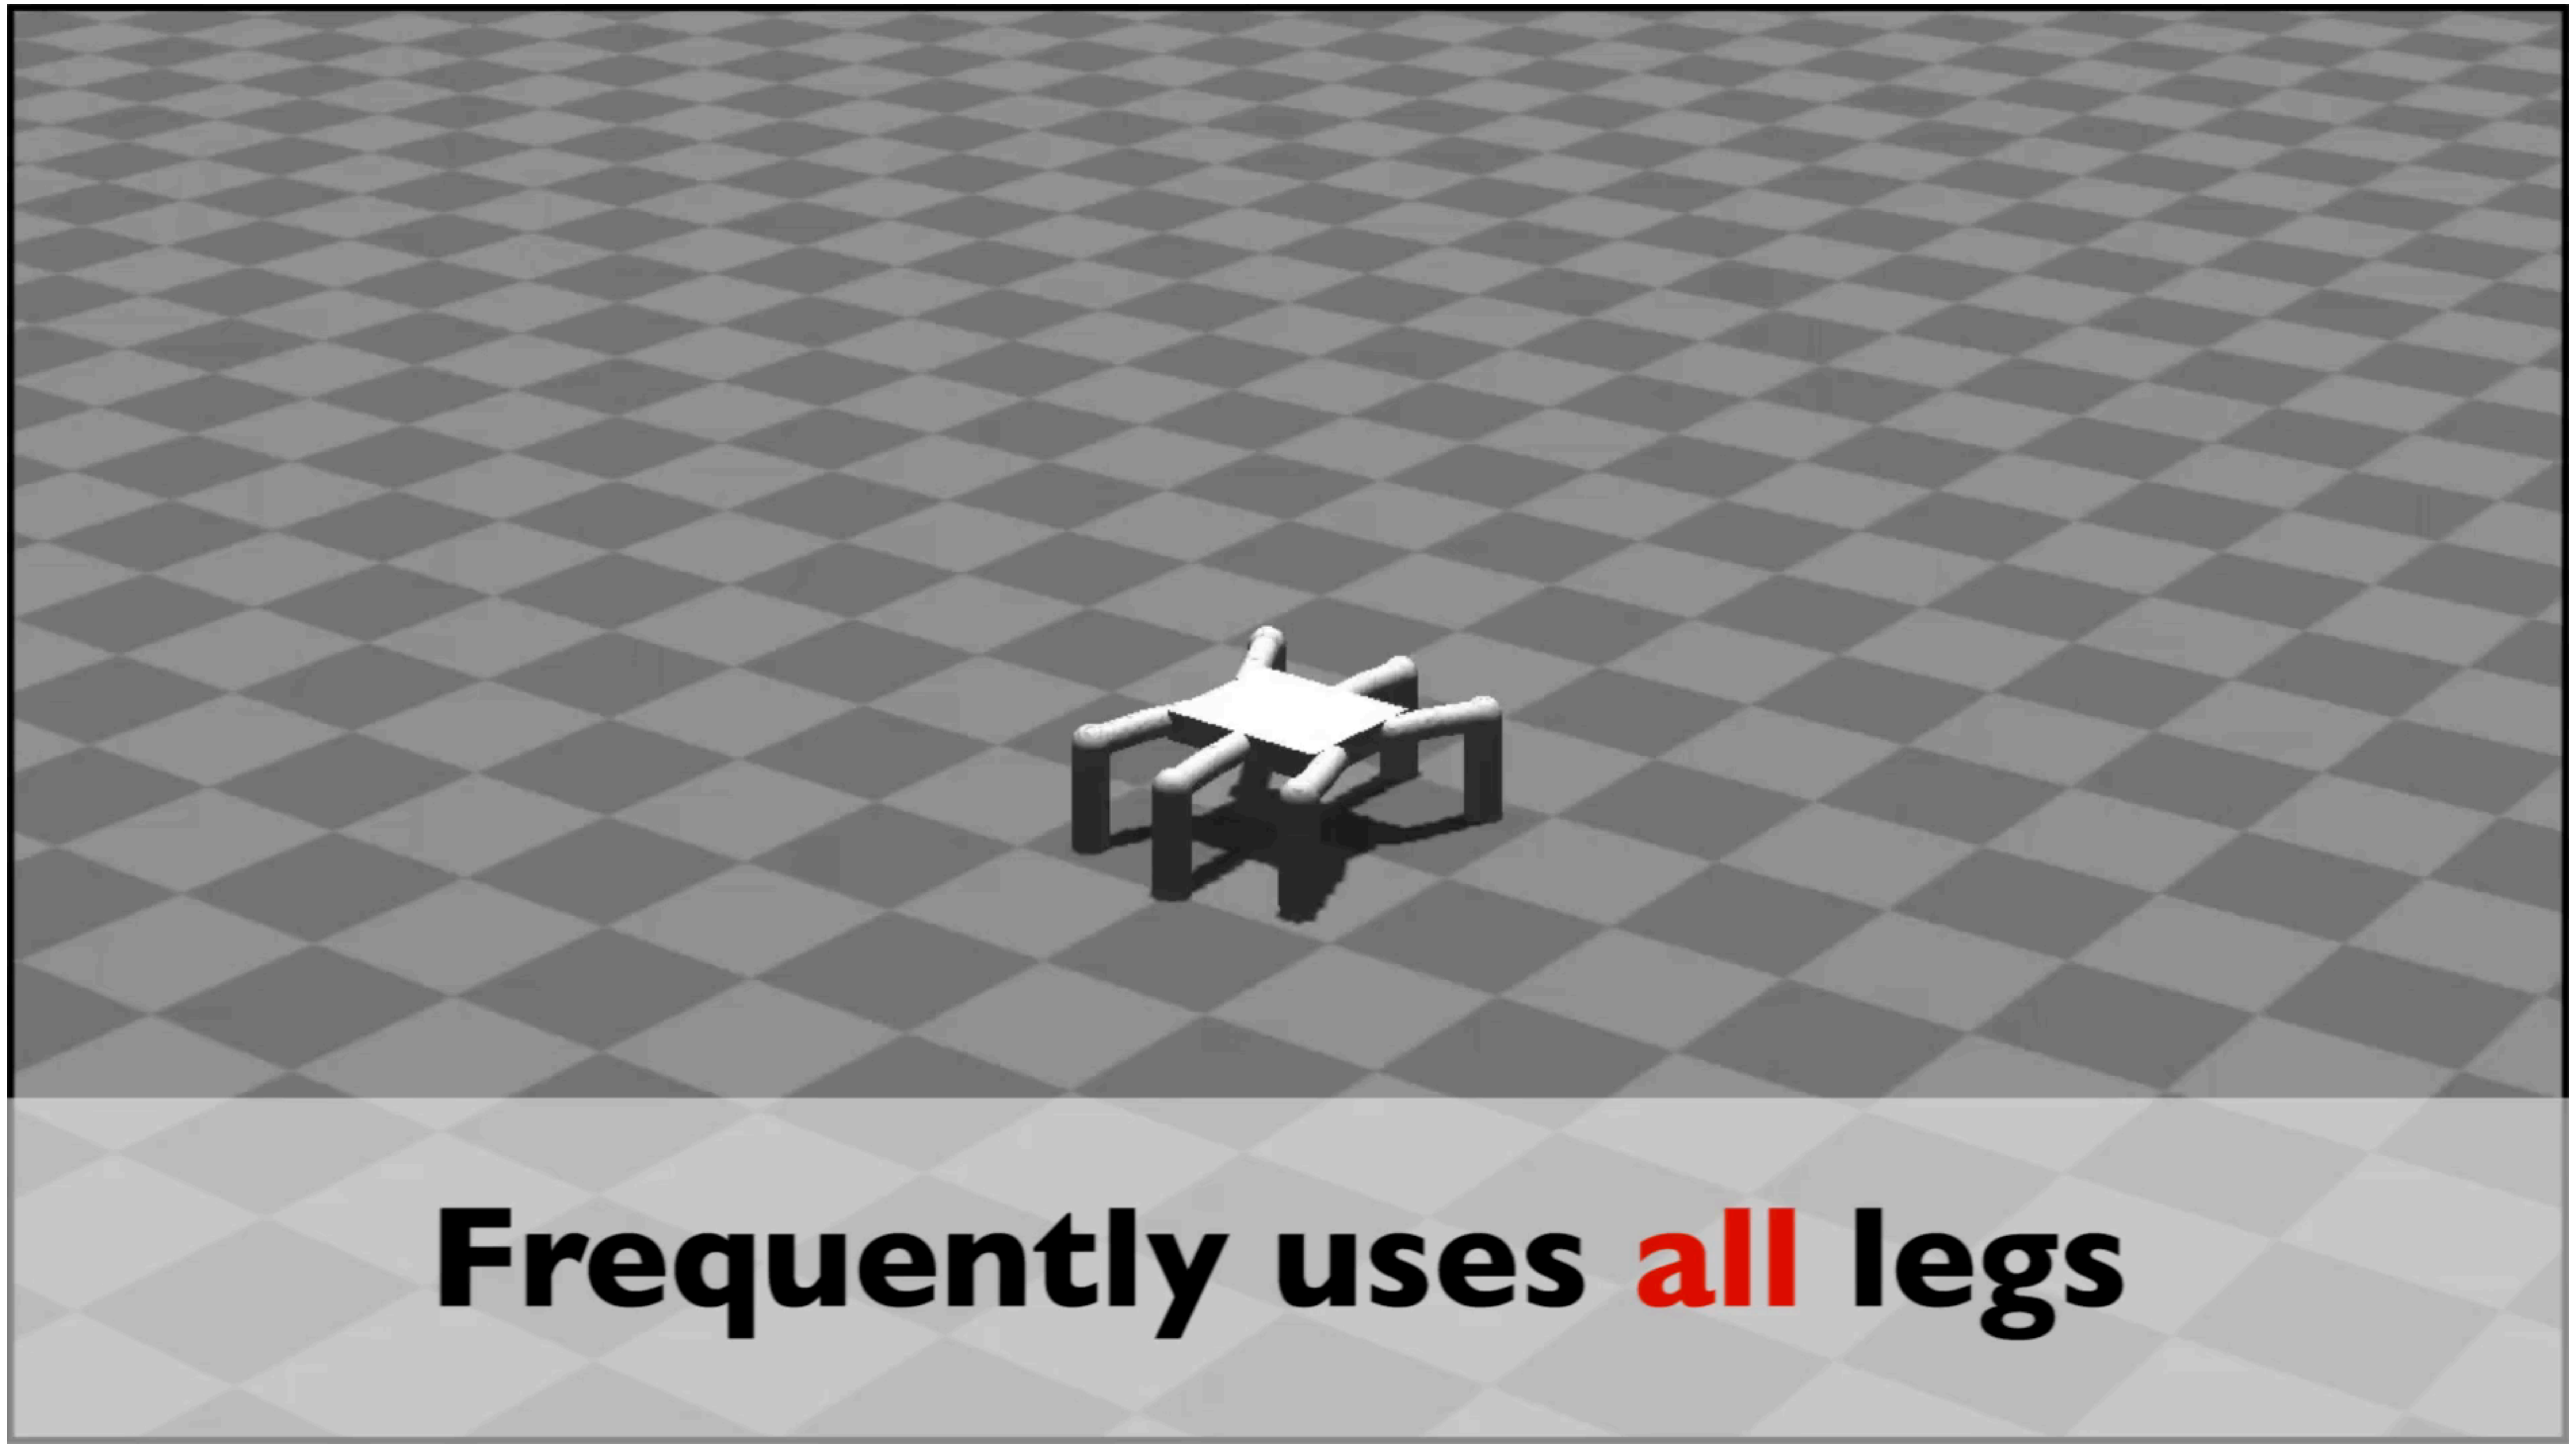

Search space: 36 parameters
(open-loop controller)

Behavior space: 6-D
(% of contact for each foot)

Cully A, Clune J, Tarapore D, Mouret JB. Robots that can adapt like animals. Nature. 2015 May;521(7553):503-7.

MAP-Elites vs PPO

Hexapod robot — neural network

- 2 hidden layers of 6 neurons
- 18 outputs (joint positions

- Not a bad optimizer (small space)!
- Keep in mind: different problem (diversity)

Open loop: input = time-modulo period

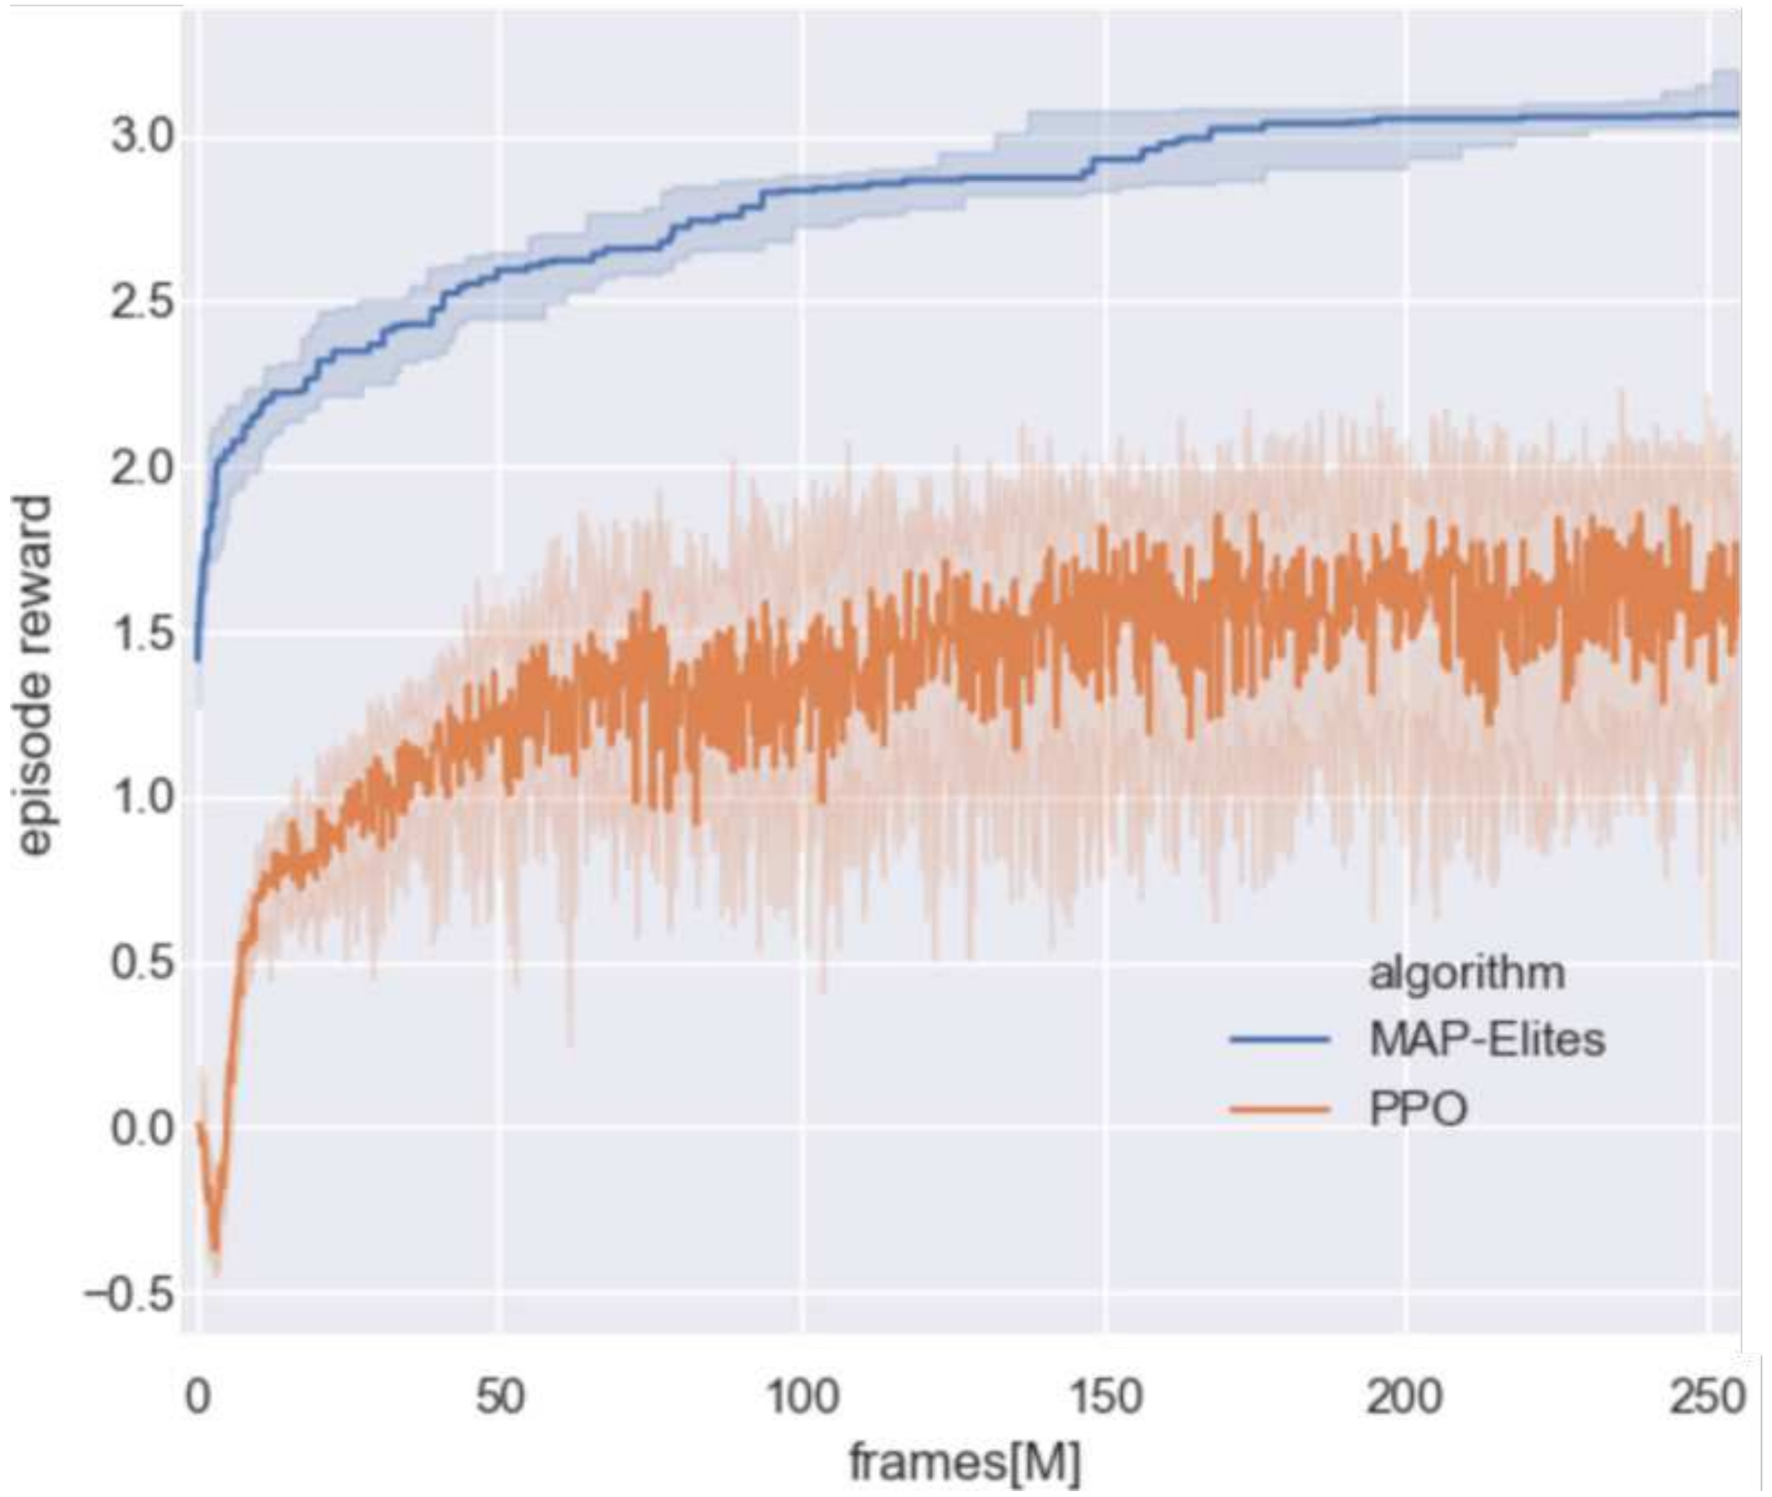

Closed-loop: input = robot state (98 to 282 weights)

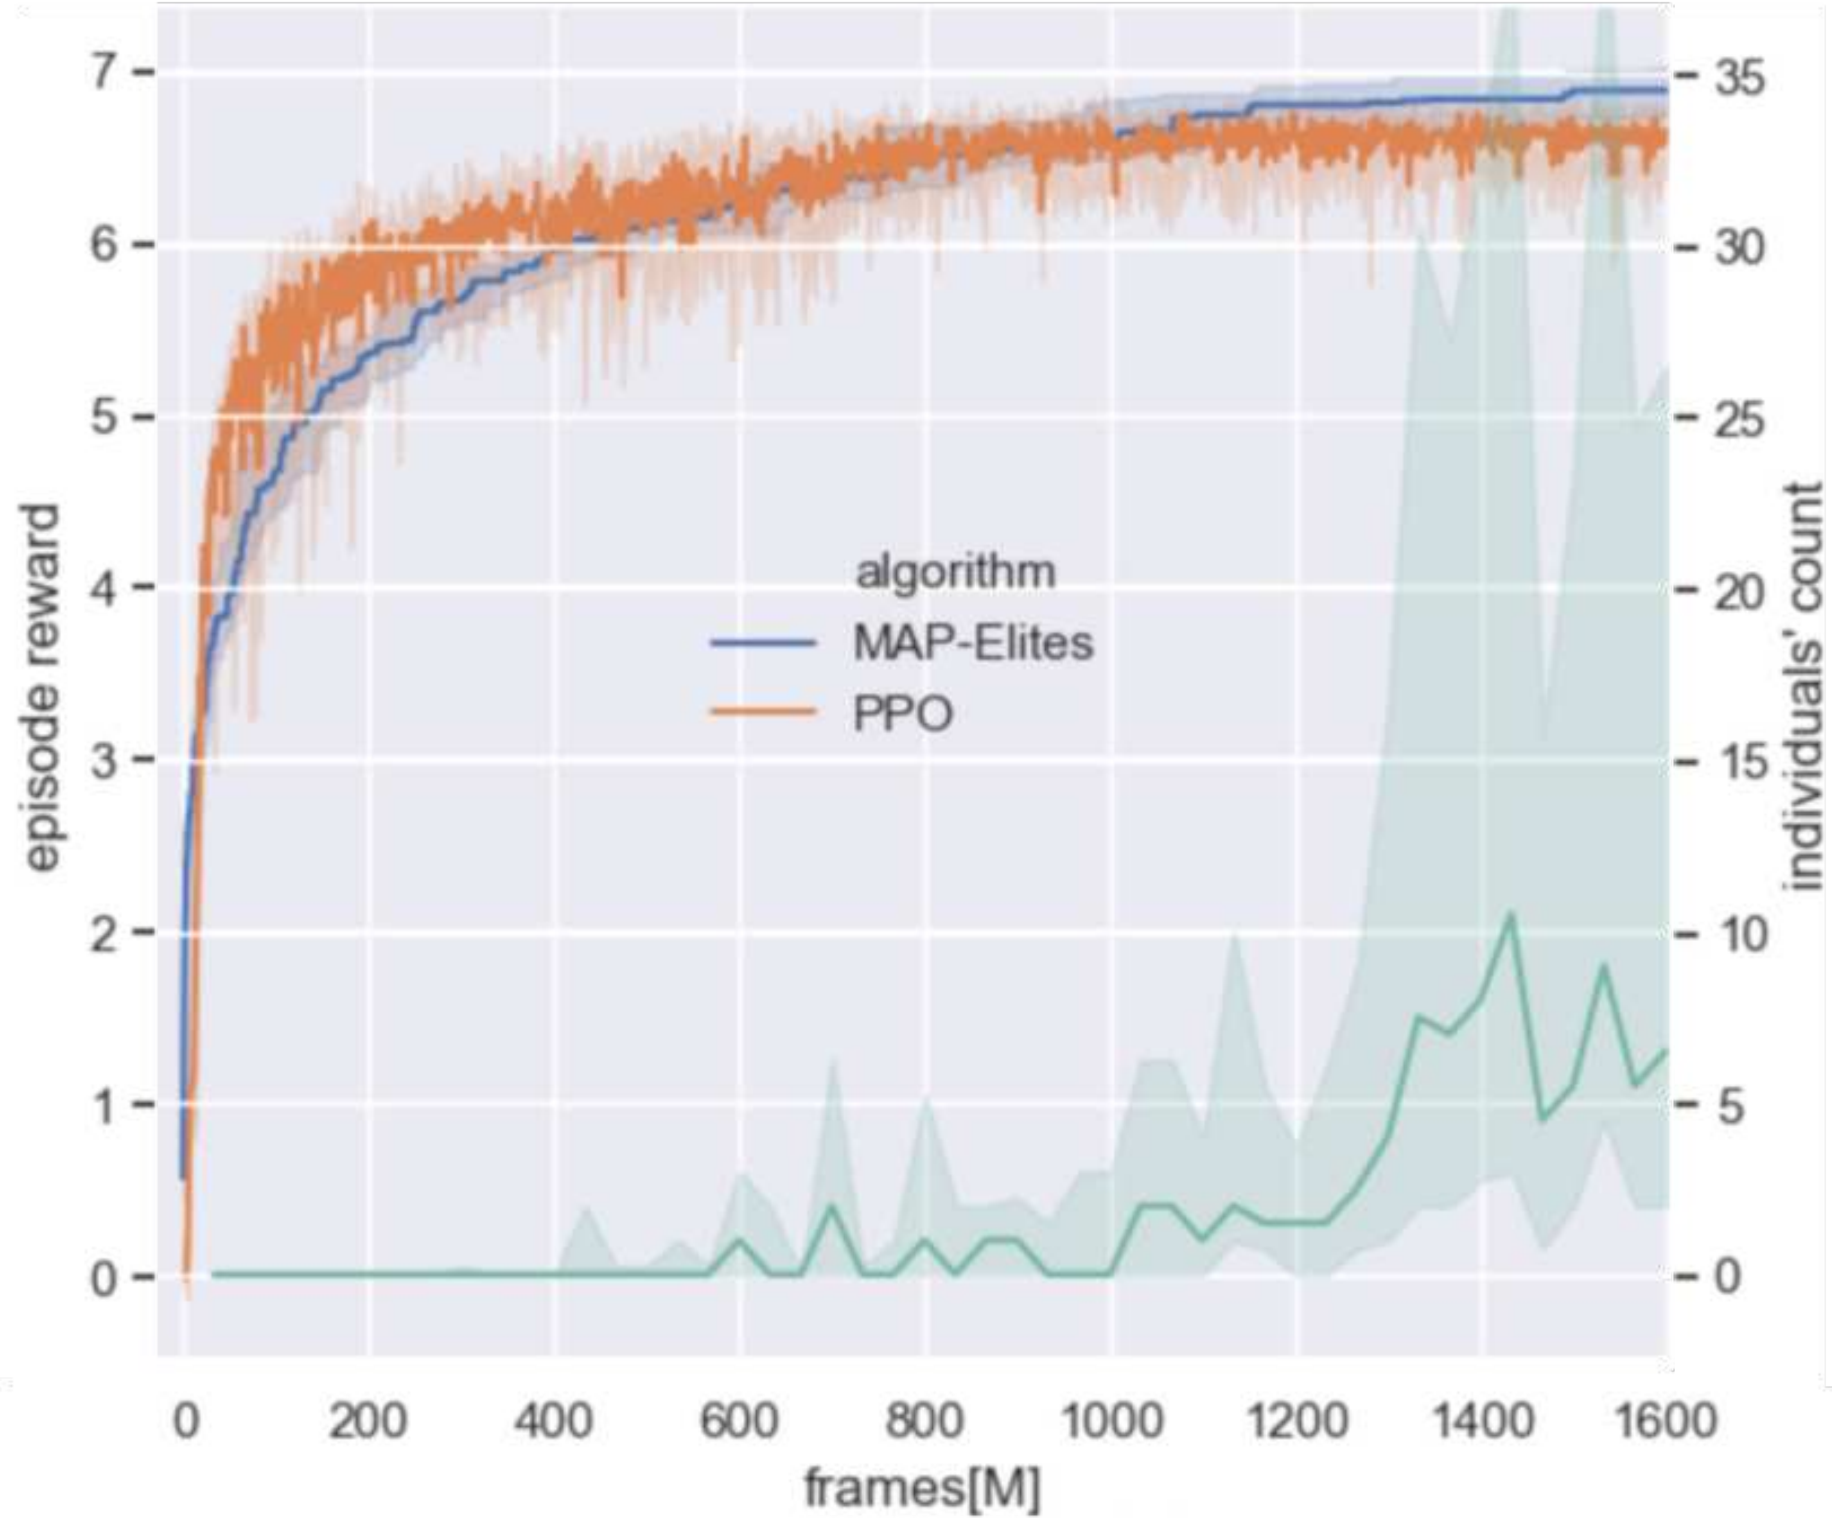

Brych, S., & Cully, A. (2020). Competitiveness of MAP-Elites against Proximal Policy Optimization on locomotion tasks in deterministic simulations. *arXiv preprint arXiv:2009.08438*.

Example: designing a dataset for grasping

EGAD!

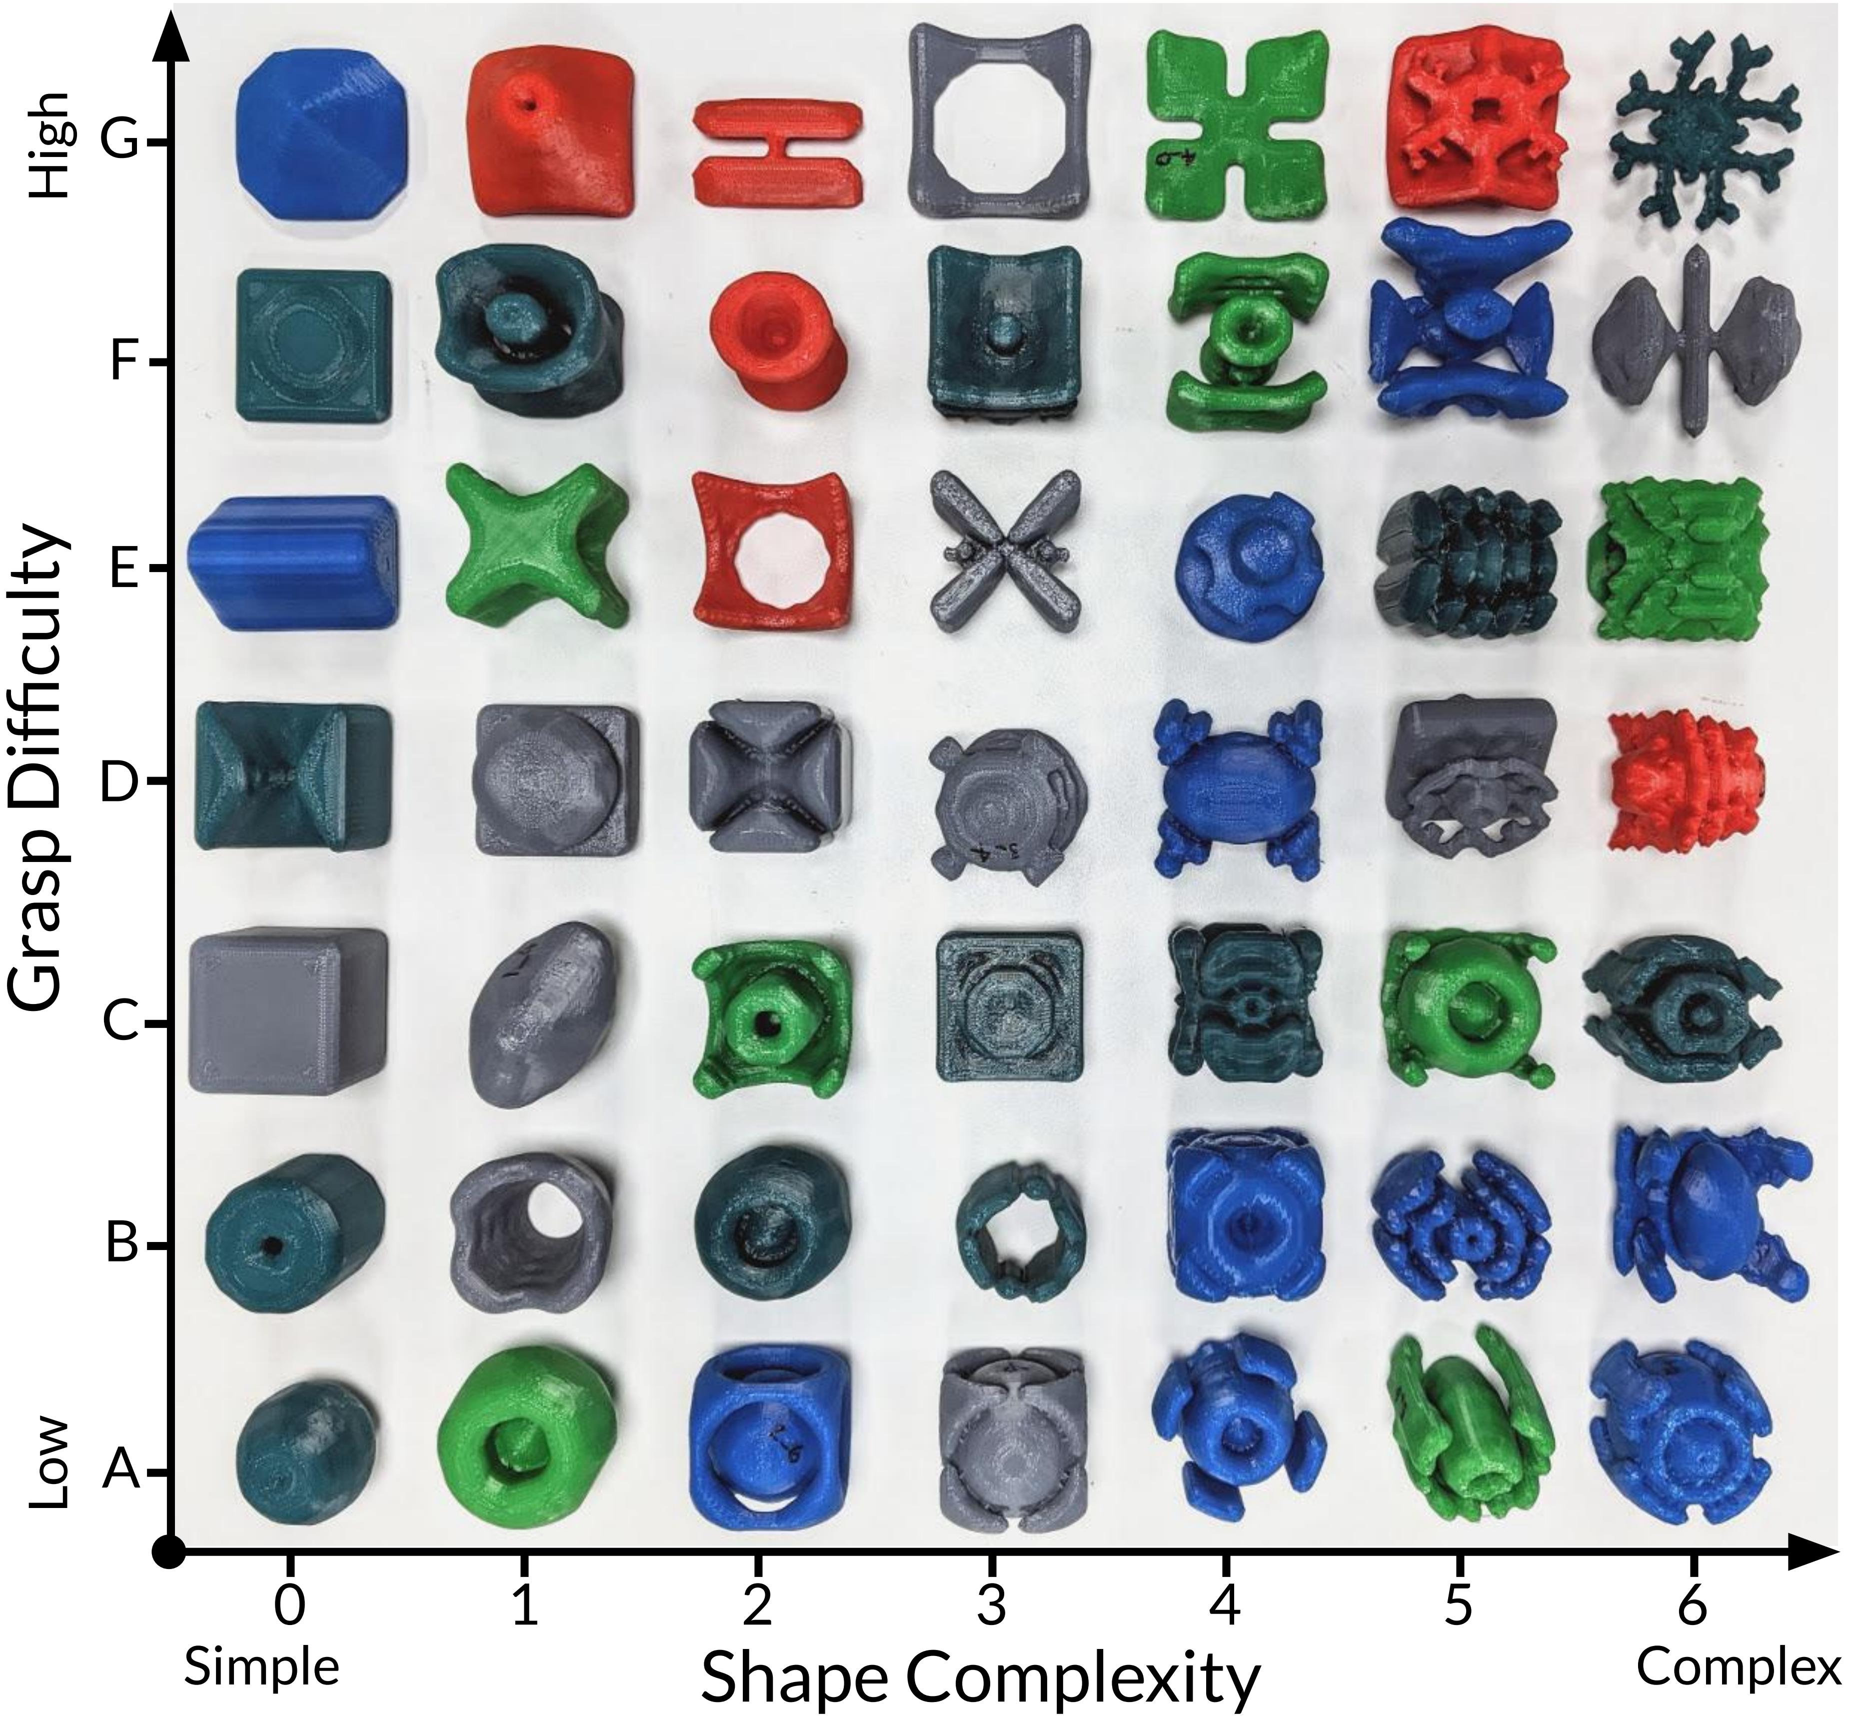

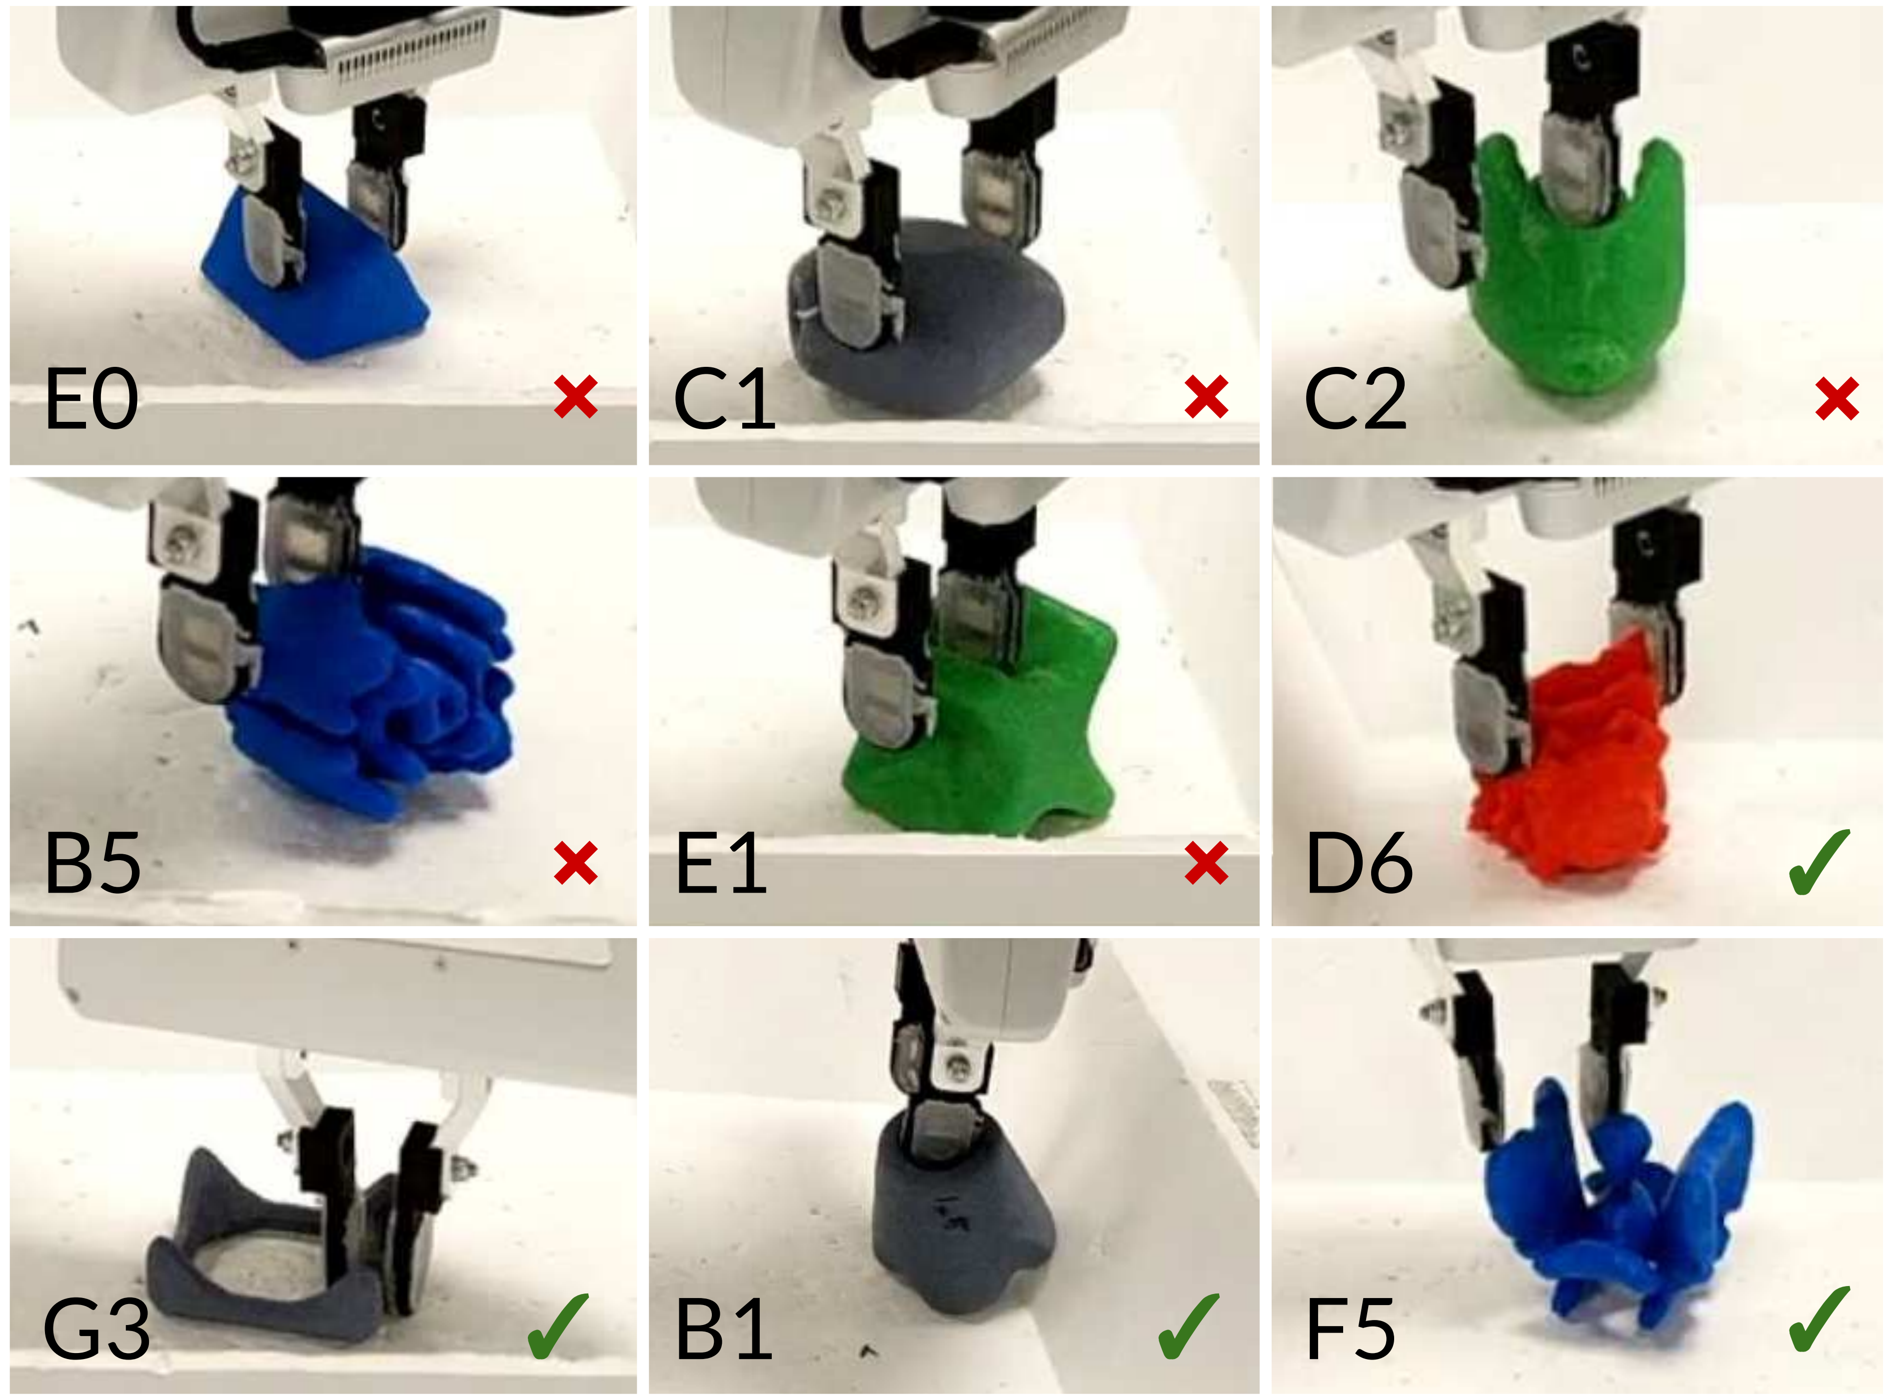

Morrison, D., Corke, P., & Leitner, J. (2020). Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation. IEEE Robotics and Automation Letters, 5(3), 4368-4375.

Example: designing airfoils

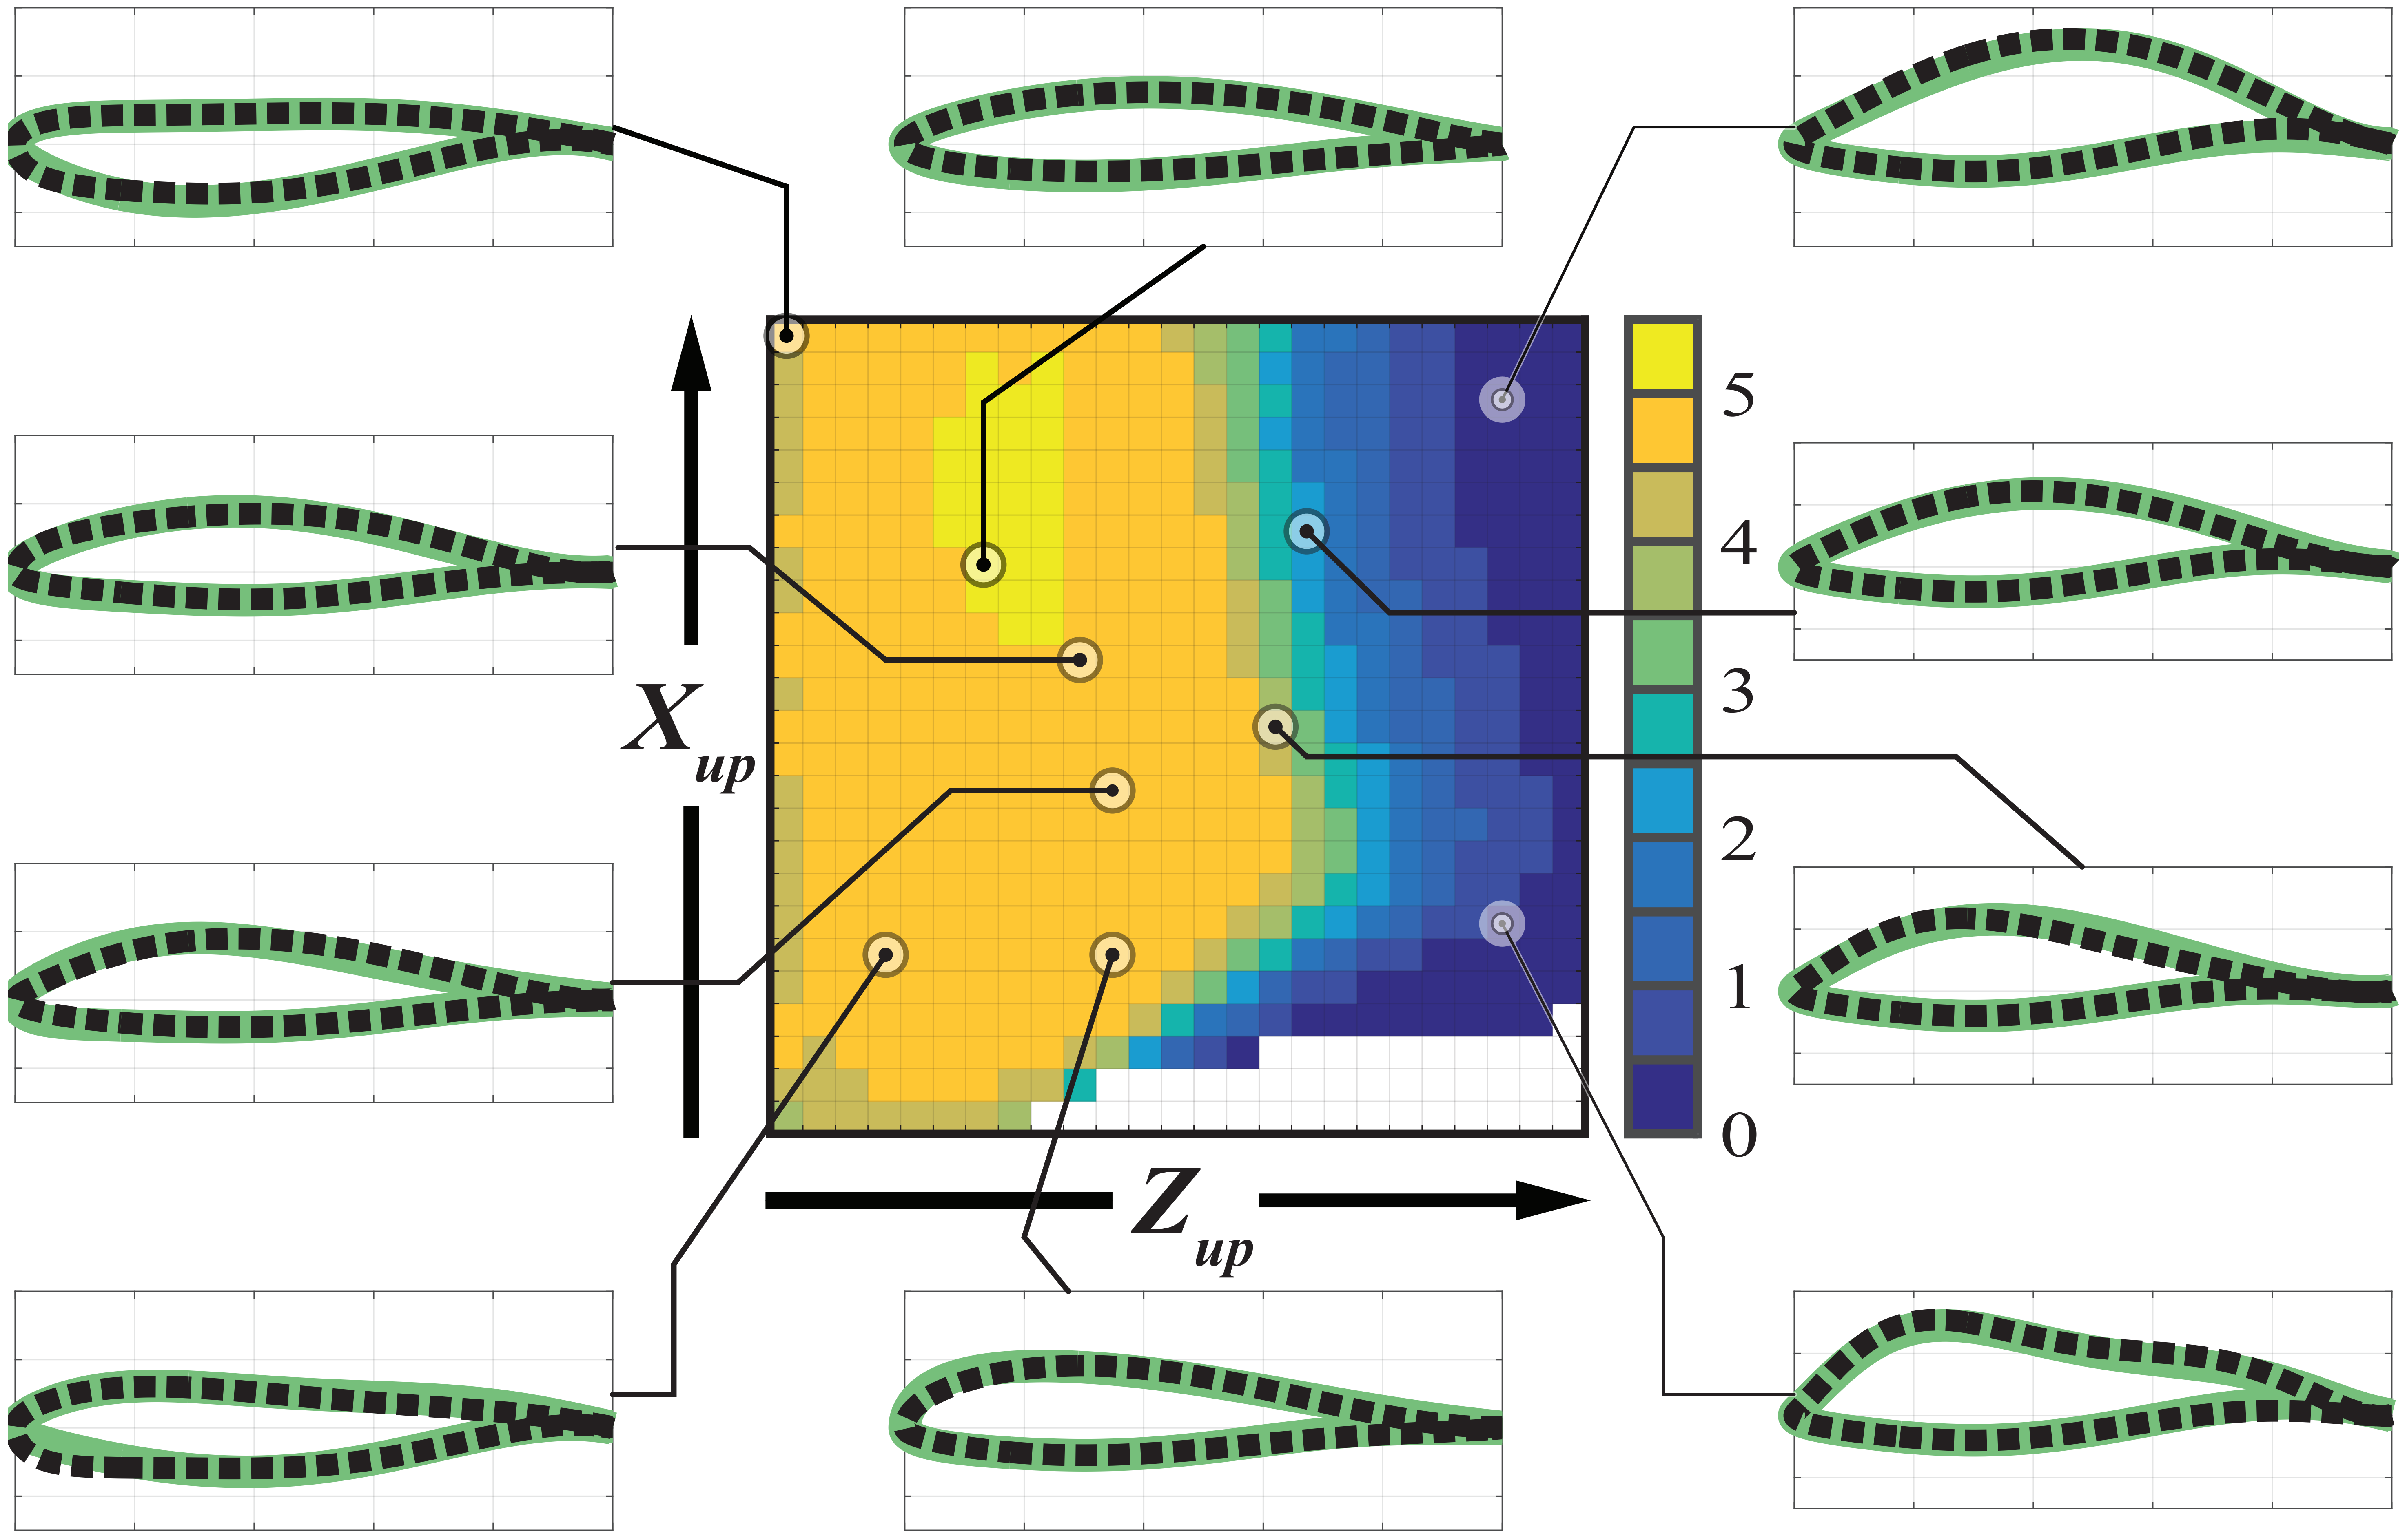

Gaier A, Asteroth A, Mouret JB. Data-efficient design exploration through surrogate-assisted illumination. Evolutionary computation. 2018

Trying to find the skull with MAP-Elites

gen 74

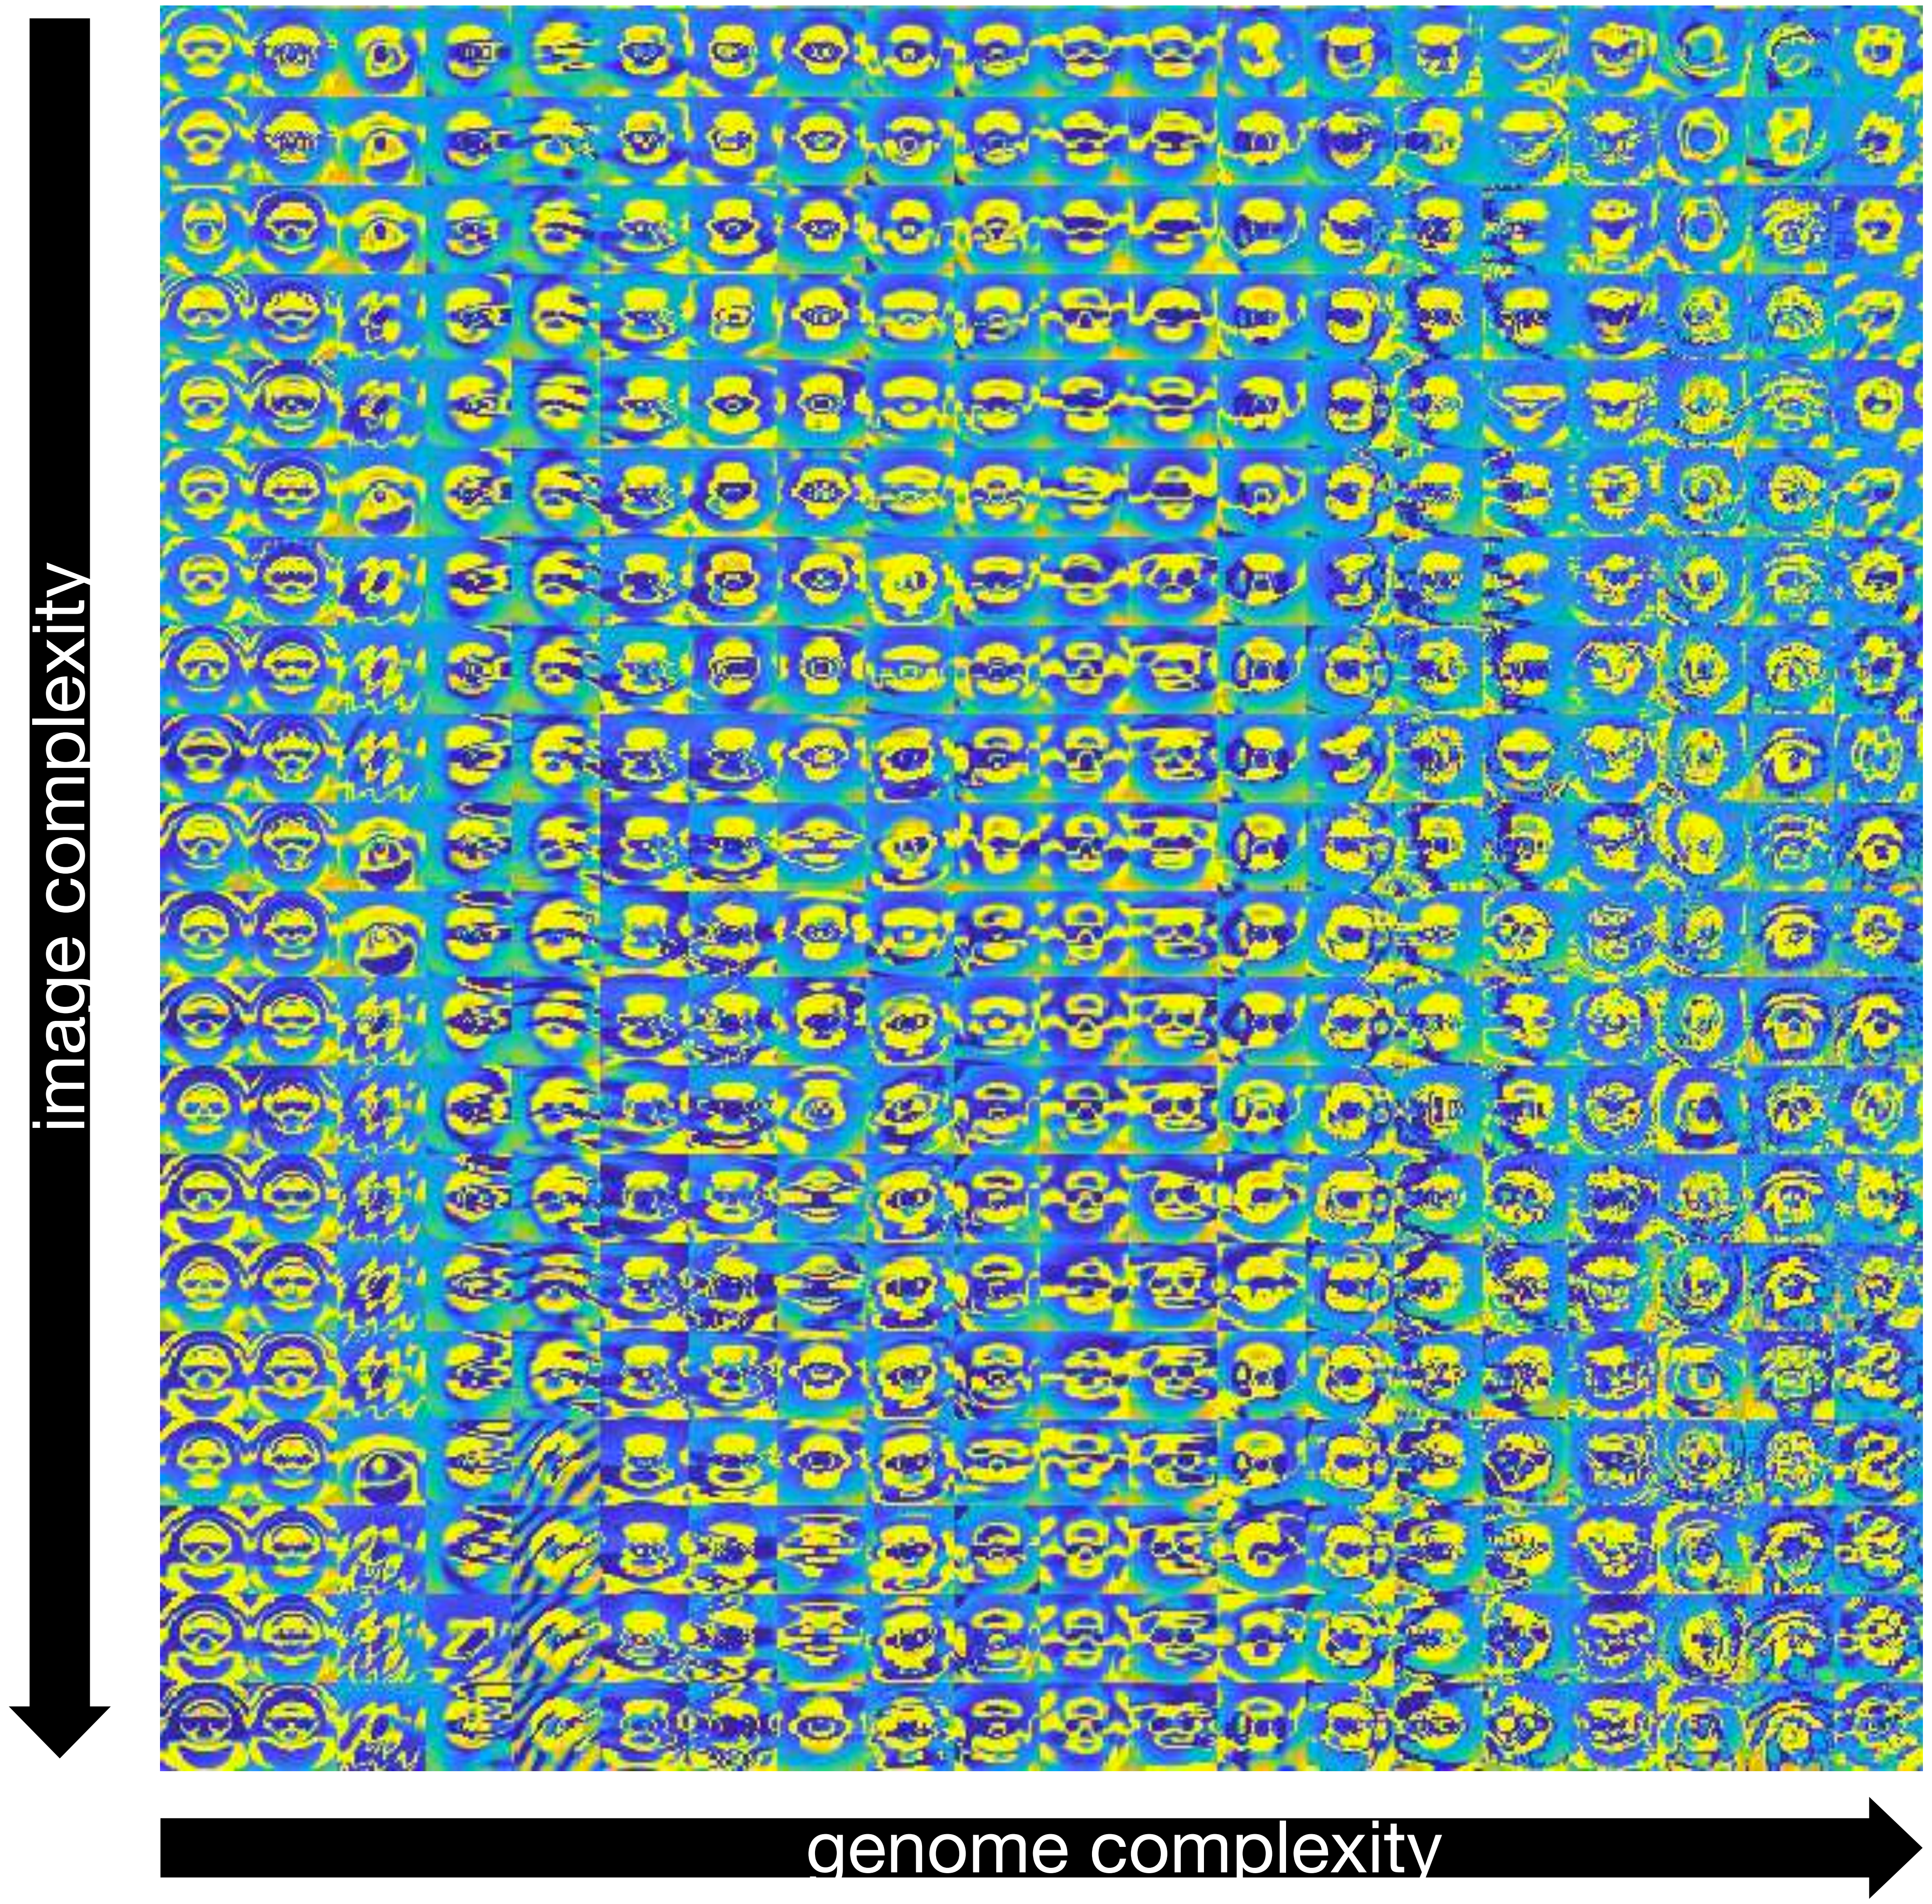

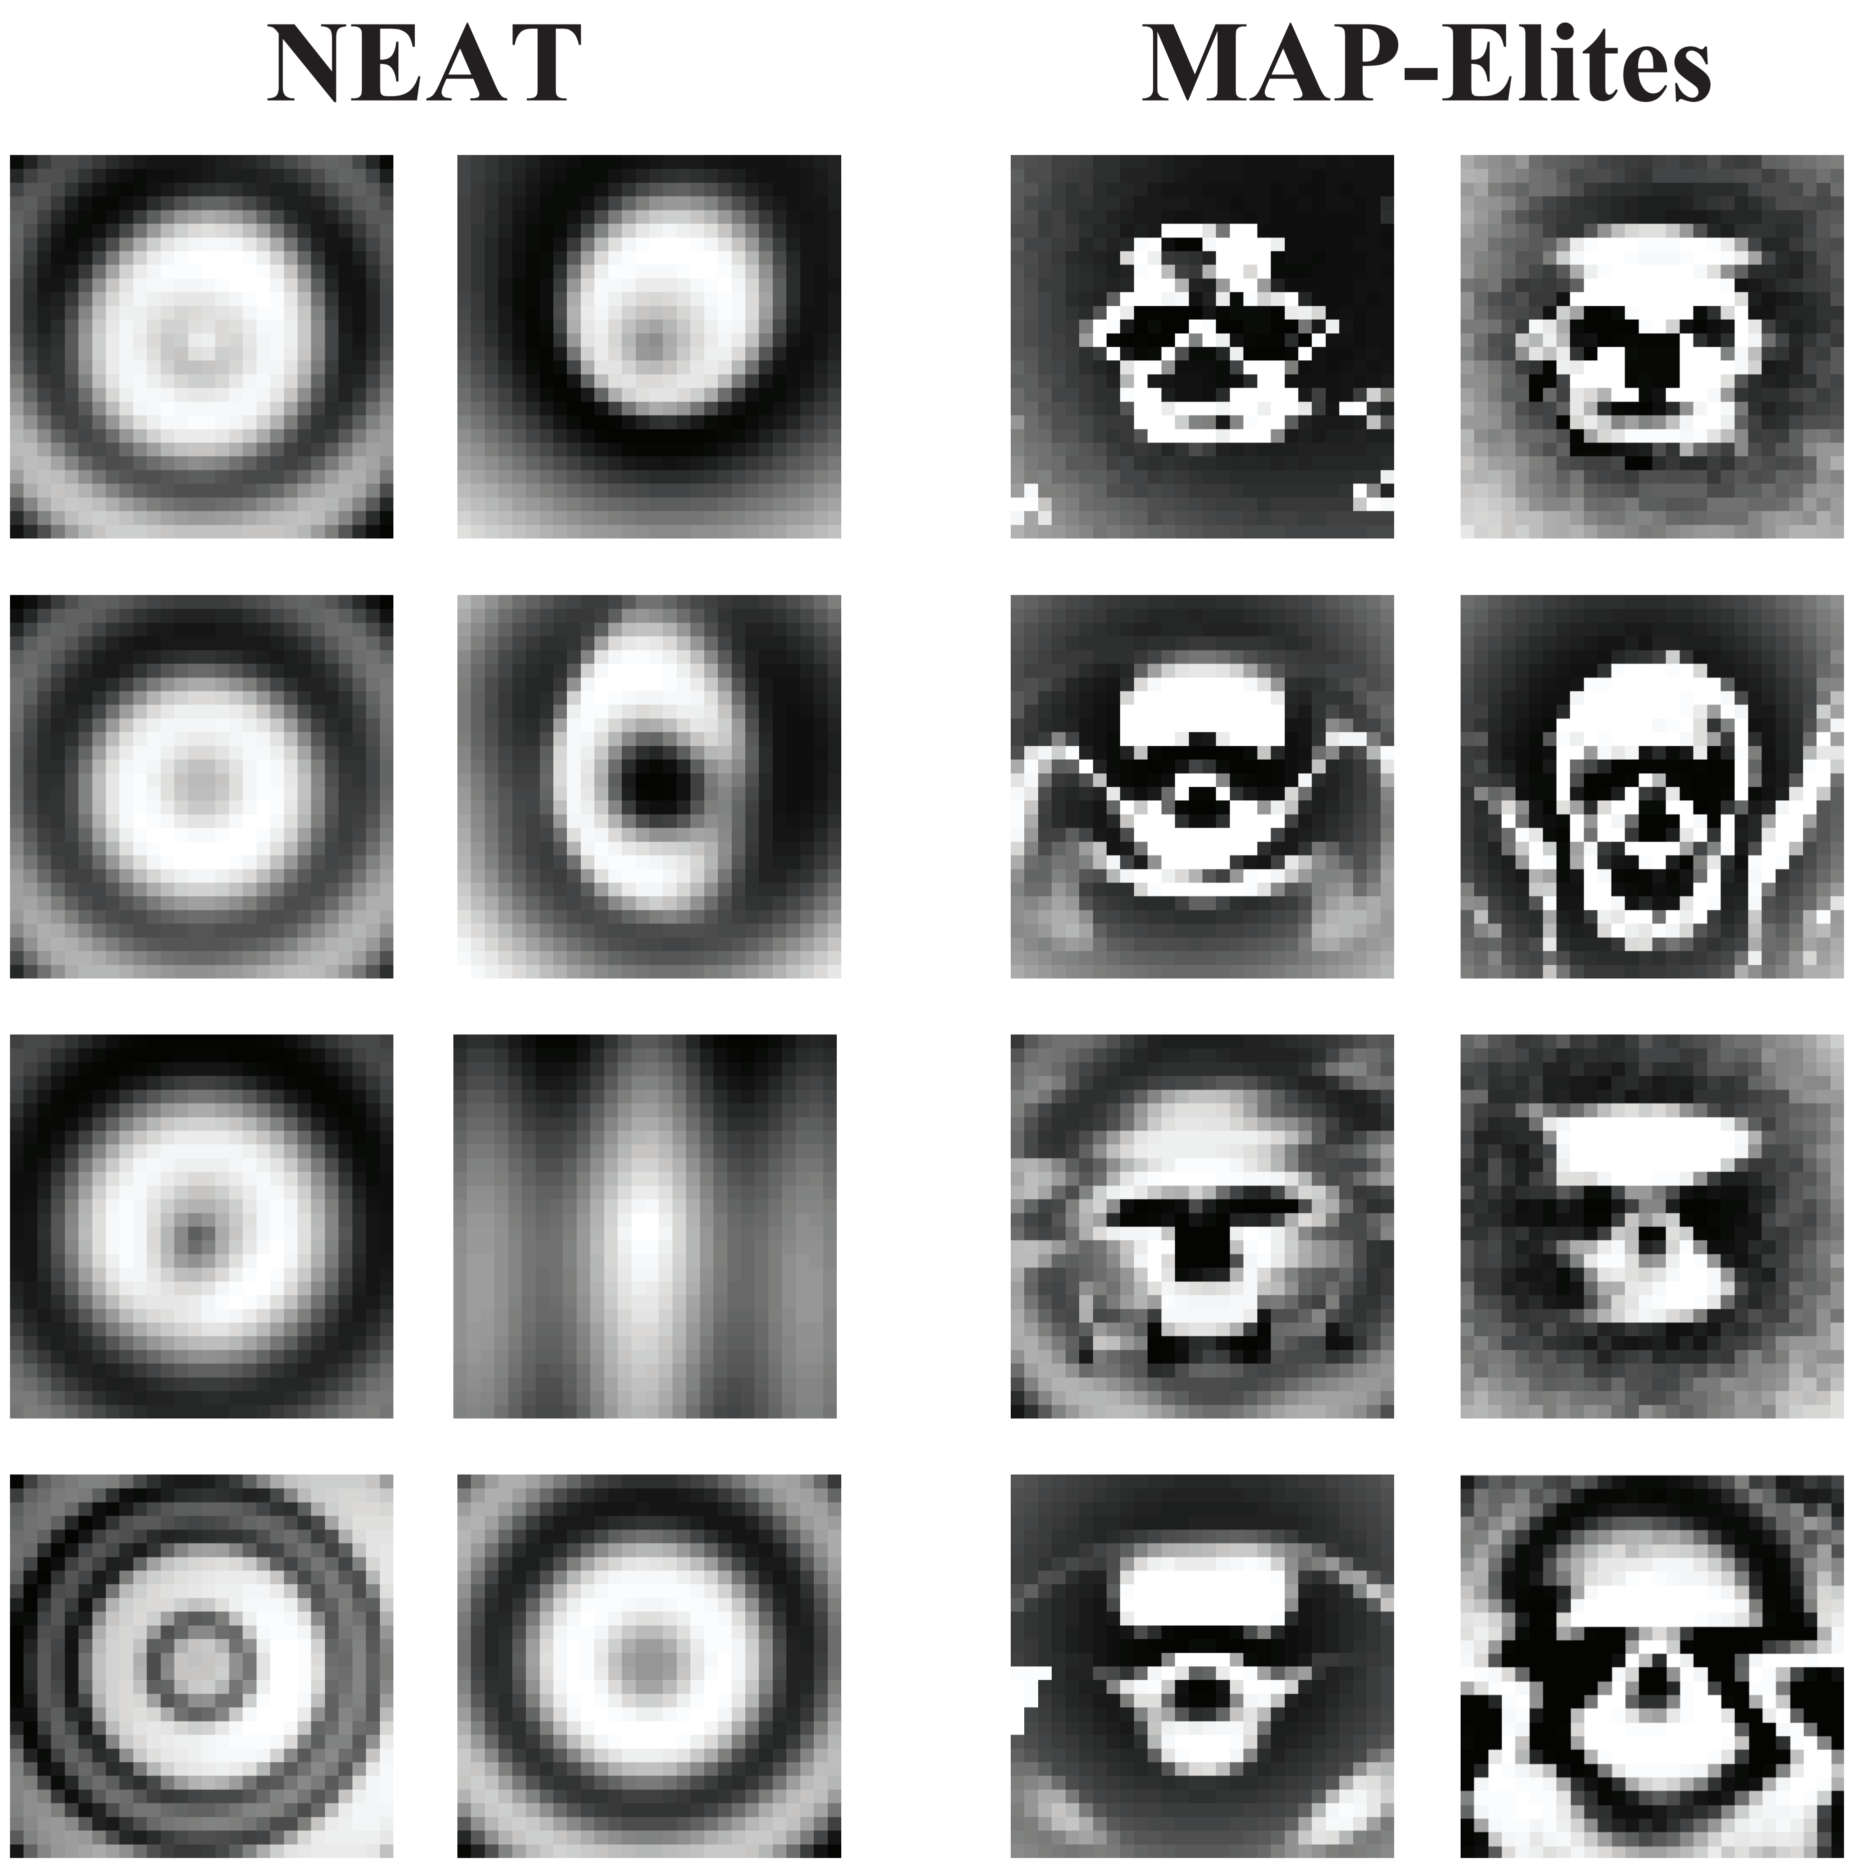

Gaier, A., Asteroth, A. & Mouret, J,-B. (2019). Does Quality Diversity Generate Better Stepping Stones than Objective-based Search? Proc. of GECCO

Trying to find the skull with MAP-Elites

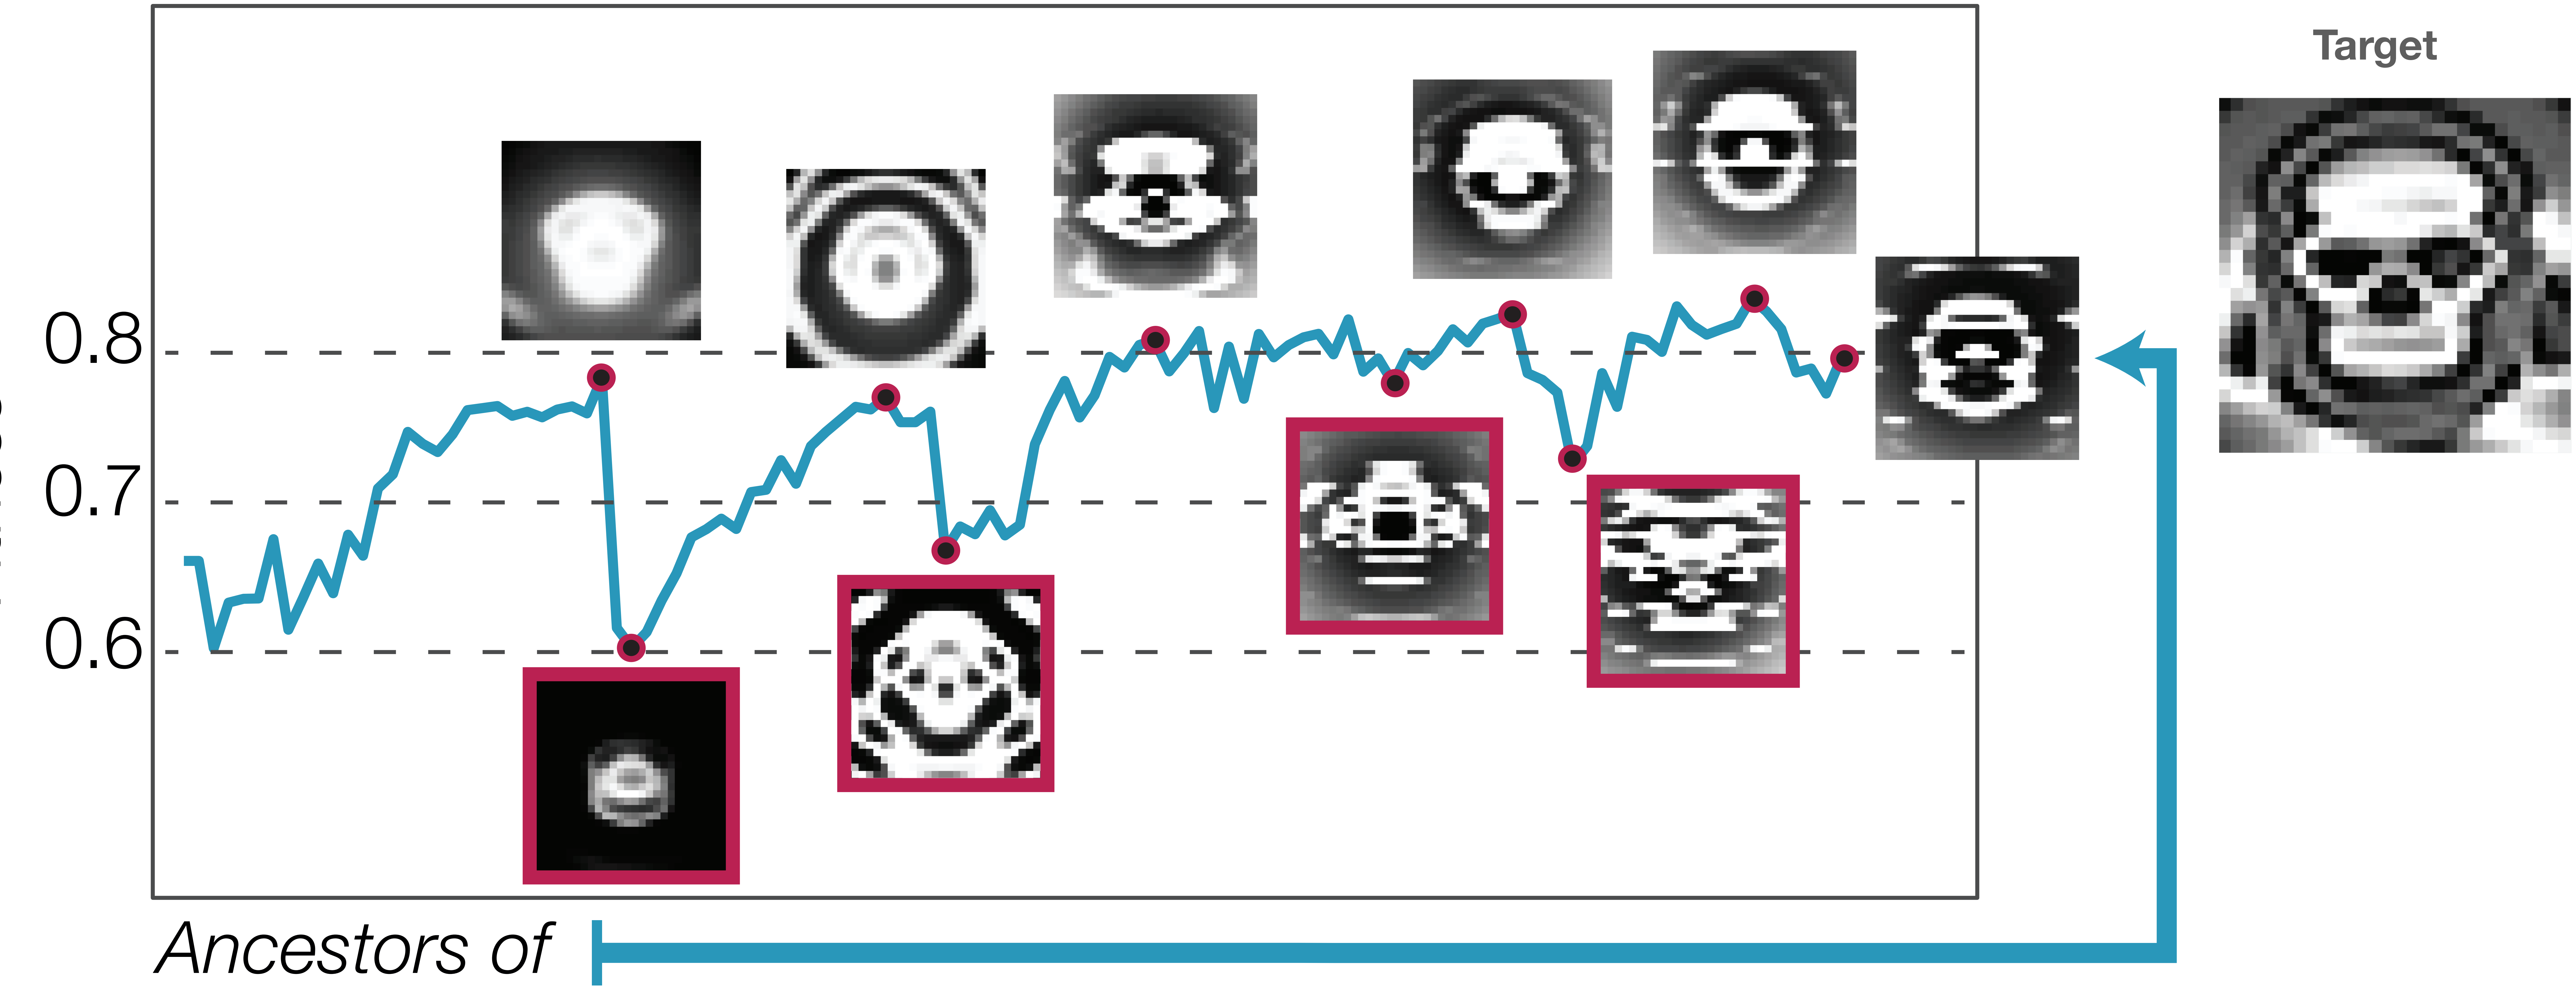

Gaier, A., Asteroth, A. & Mouret, J,-B. (2019). Does Quality Diversity Generate Better Stepping Stones than Objective-based Search? Proc. of GECCO

Further readings about quality diversity

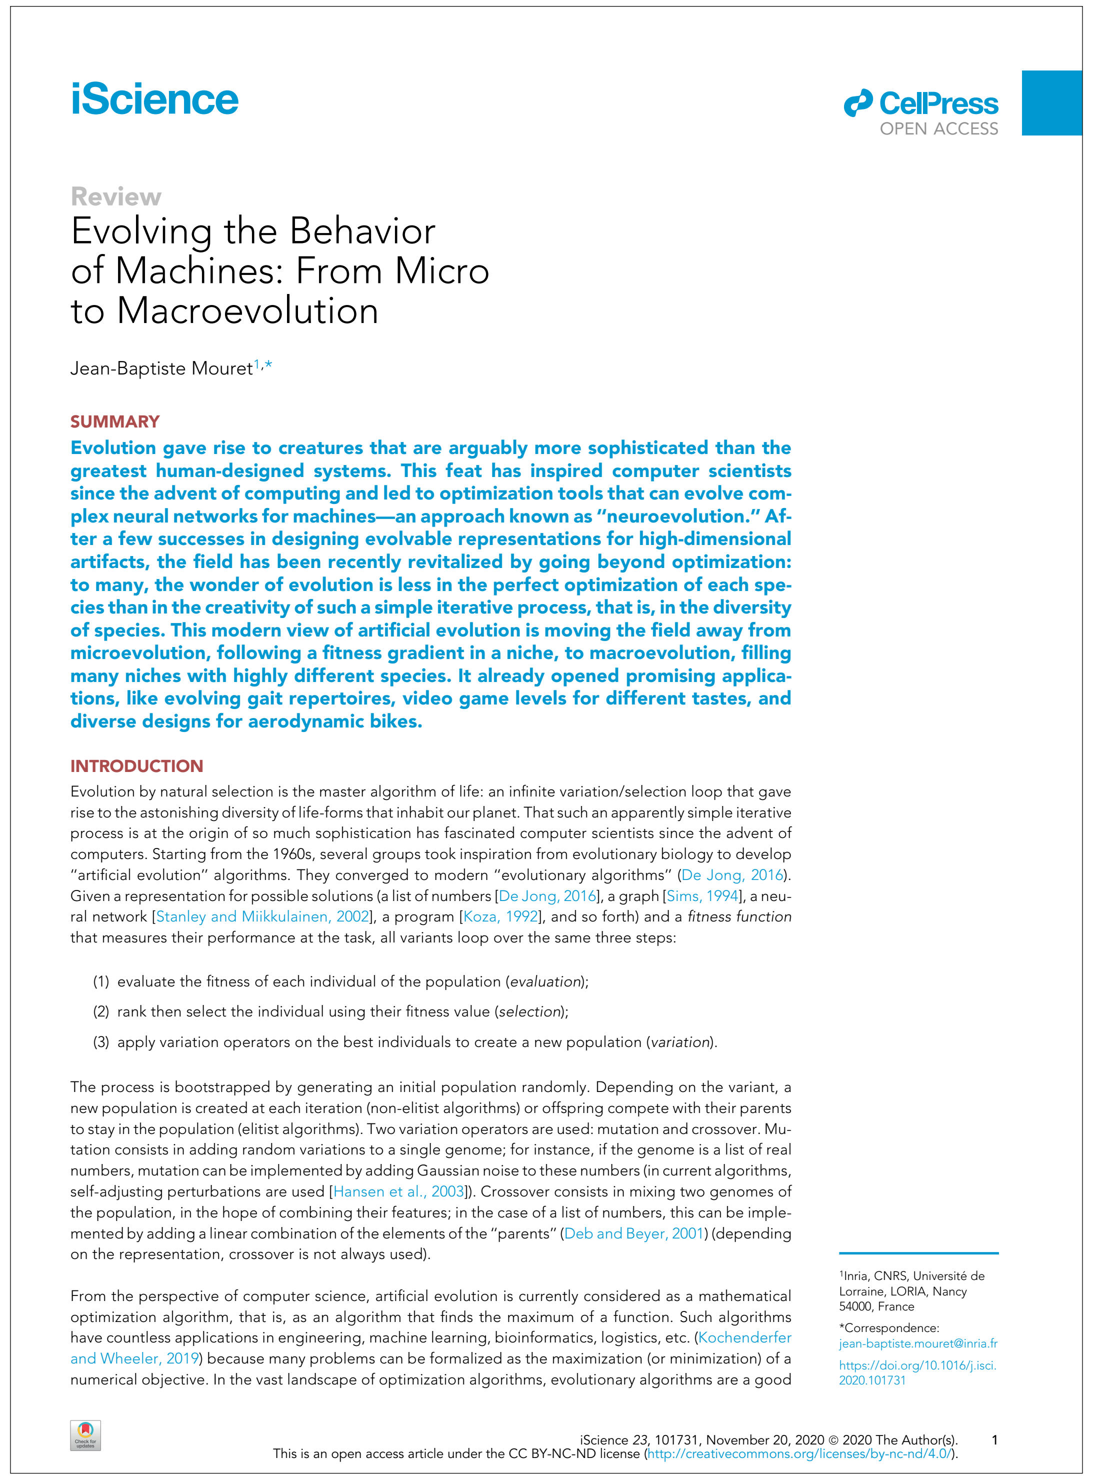

Mouret, J. B. (2020). Evolving the behavior of machines: from micro to macroevolution. *Iscience*, 101731.

Cully, A., & Demiris, Y. (2017). Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation*, *22*(2), 245-259.

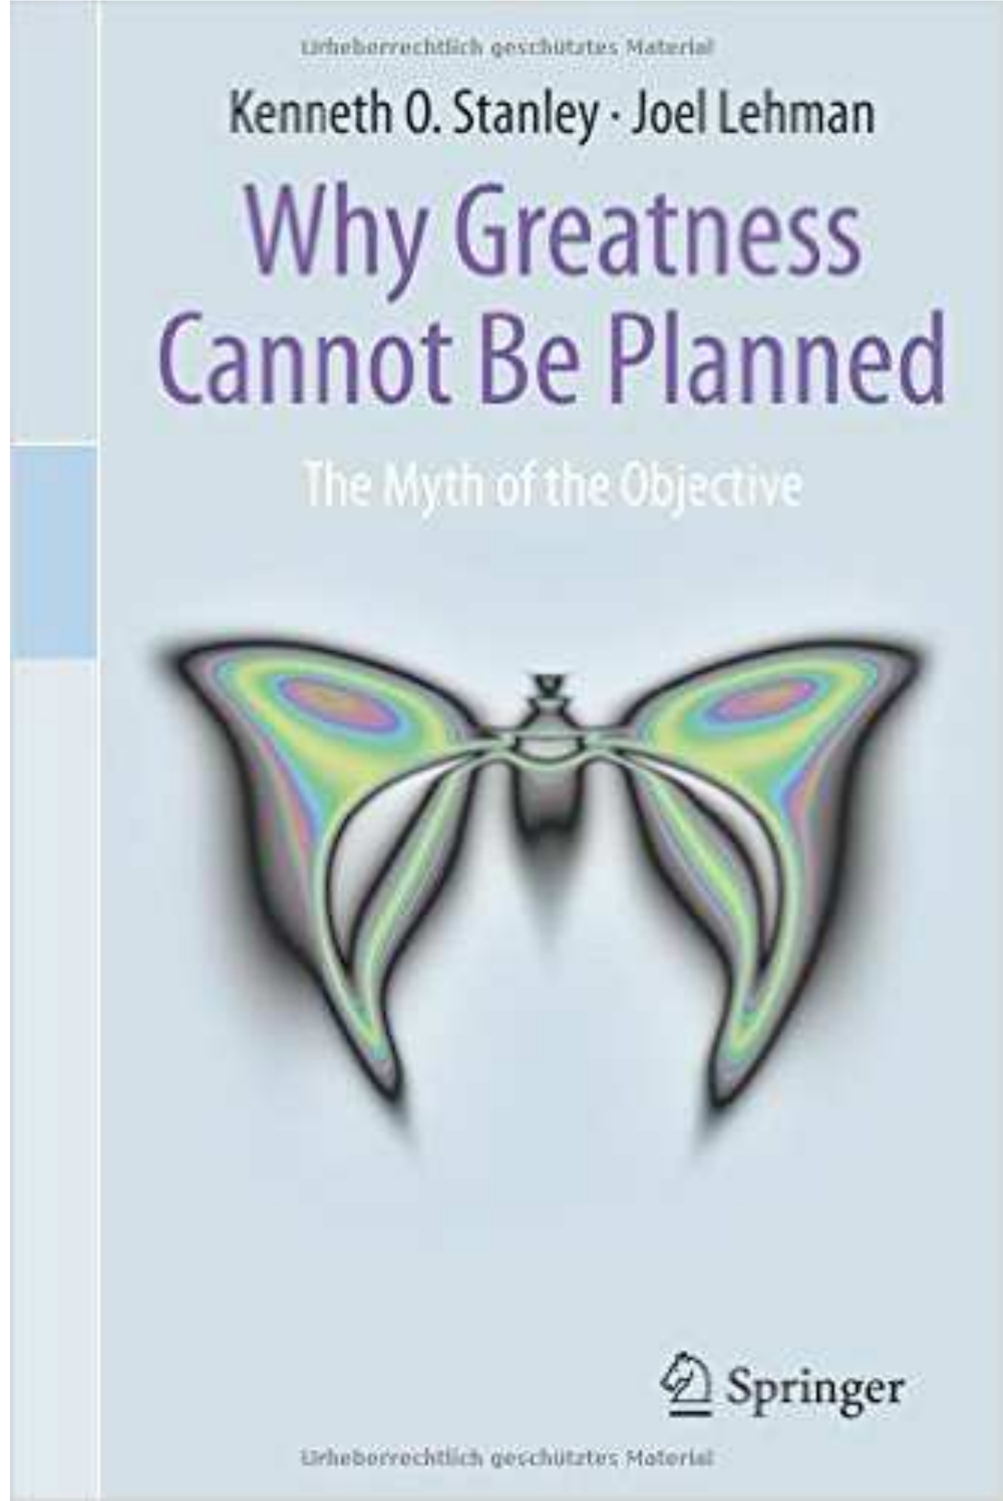

KO Stanley, J Lehman. *Why Greatness Cannot Be Planned* (2015) - Springer

https://quality-diversity.github.io