Amherst College

University of Massachusetts Amherst

The Emergence of Diversity

Ryan Boldi

University of Massachusetts Amherst
rbahlousbold@umass.edu

PUSH*lab
programs under selection + heredity

ALIFE 2023
GHOST IN THE MACHINE

There is a vast diversity of life

Why Care About Diversity?

In practice:

- Helps you solve new combinations of problems that were previously not possible

- Overcome deceptive landscapes

Lexicase Selection[1, 2]

Initially designed for “uncompromising problems”, where superior performance on one task does not compensate for bad performance on another.

Selects parents based on a large (>5) set of test cases.

Does not aggregate fitness or accuracy metrics.

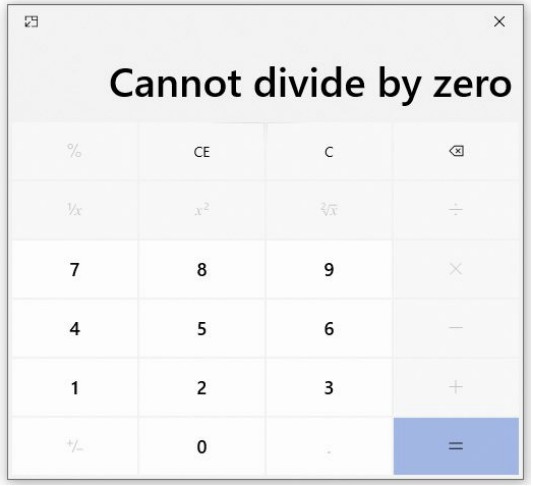

Instead, it selects a parent that is elite on a set of cases in random order.

[1] Spector, Lee. "Assessment of problem modality by differential performance of lexicase selection in genetic programming: a preliminary report." *Proceedings of the 14th annual conference companion on Genetic and evolutionary computation*. 2012.

[2] Thomas Helmuth, Lee Spector, and James Matheson. Solving uncompromising problems with lexicase selection. IEEE Transactions on Evolutionary Computation, 19(5):630–643, 2014.

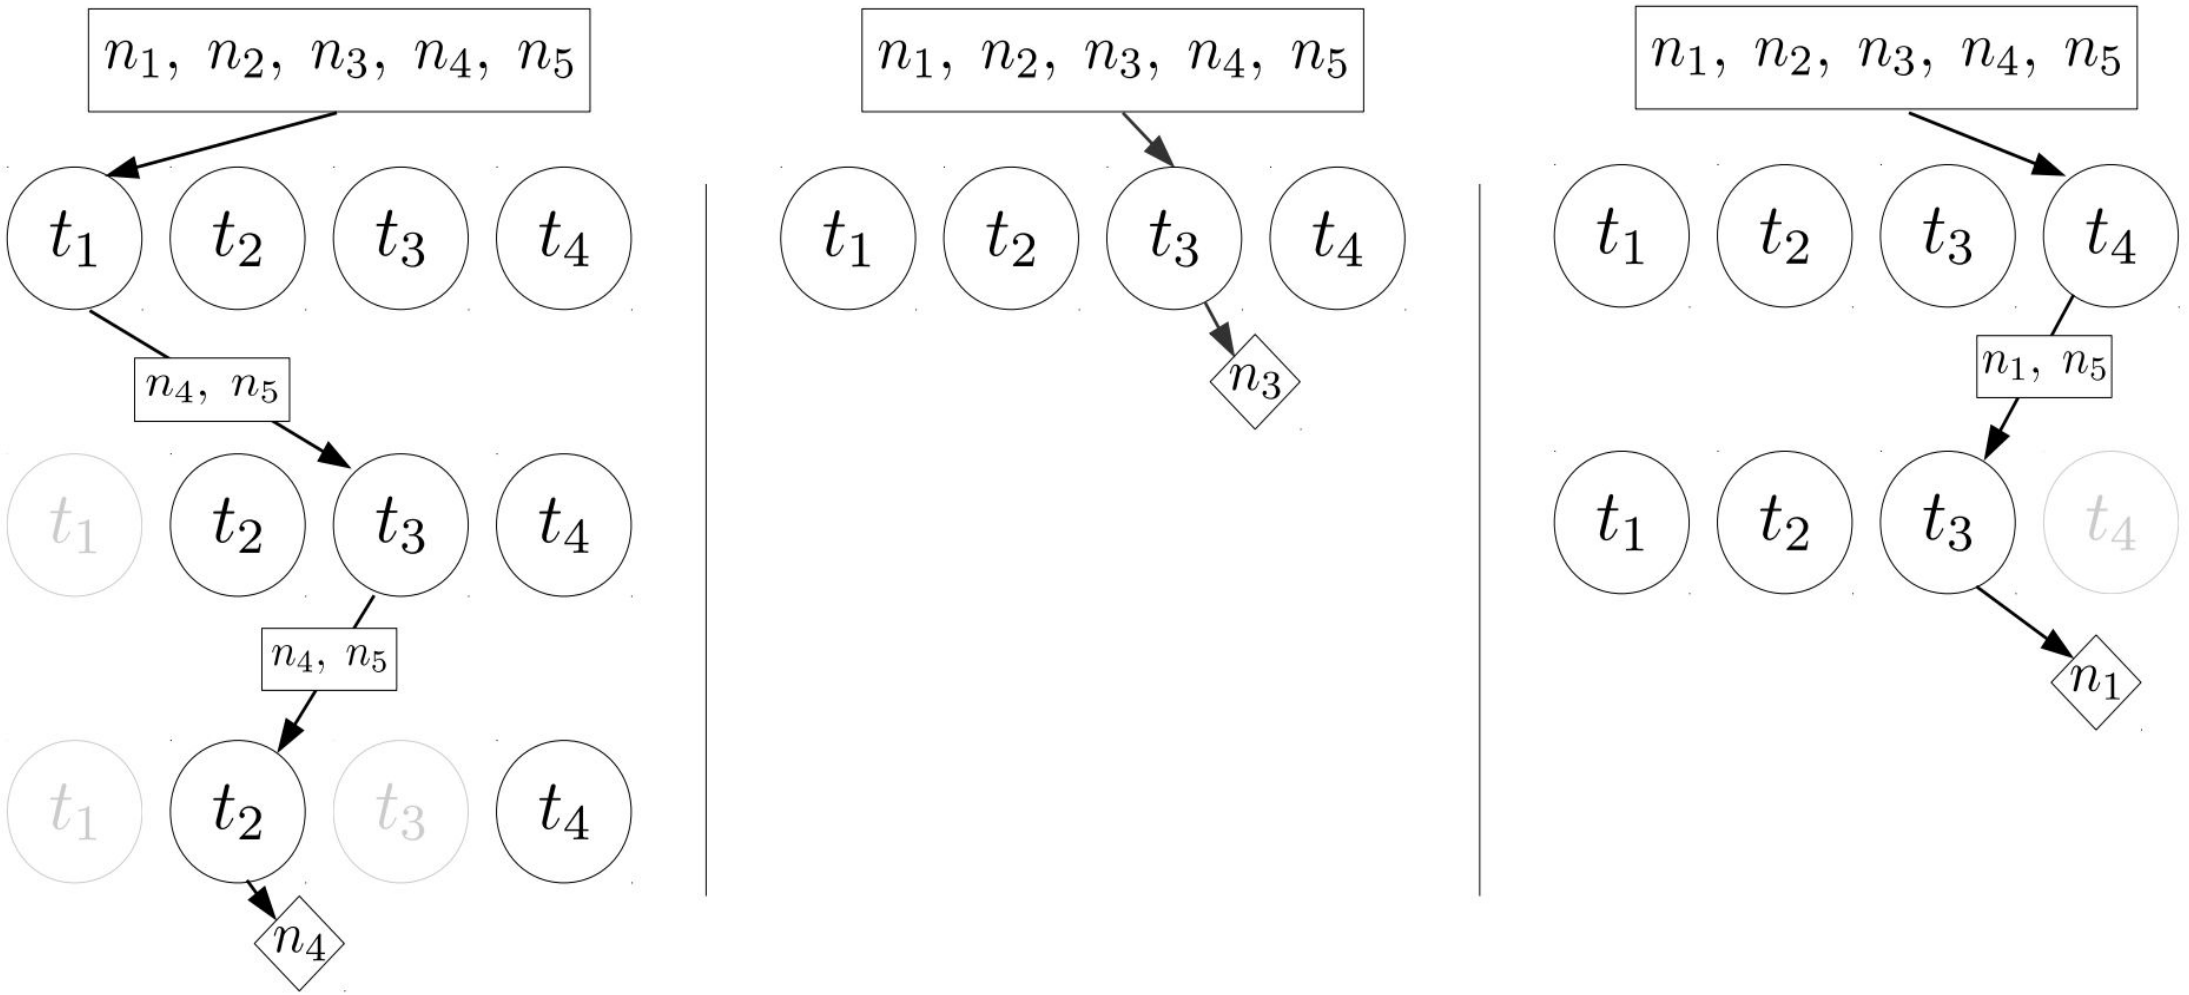

La Cava W, Helmuth T, Spector L, Moore JH. A Probabilistic and Multi-Objective Analysis of Lexicase Selection and ε-Lexicase Selection. Evol Comput. 2019 Fall;27(3):377-402. doi: 10.1162/evco_a_00224. Epub 2018 May 10. PMID: 29746157; PMCID: PMC9453780.

Environmental Interpretation

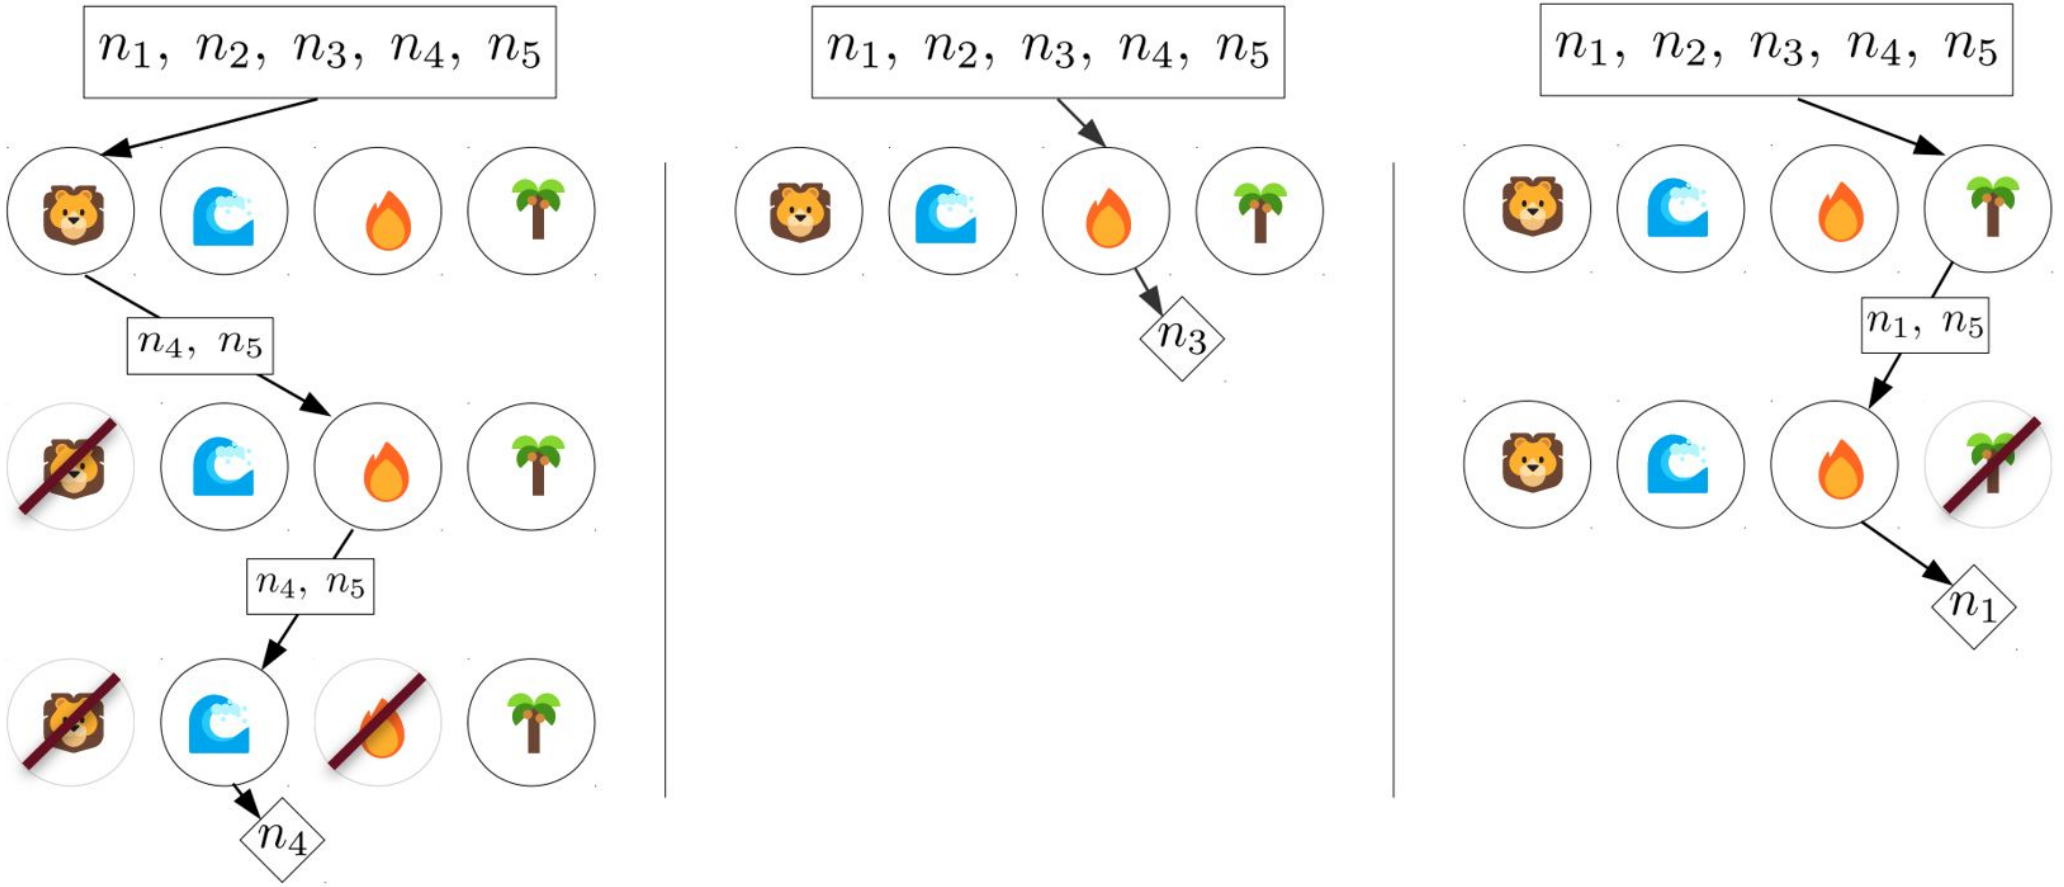

Ryan Boldi, Thomas Helmuth, and Lee Spector. 2022. The environmental discontinuity hypothesis for down-sampled lexicase selection. In The 2022 Conference on Artificial Life - Why it Didn't Work-Shop (ALIFE '22)

Quality Diversity

Novelty Search[1]

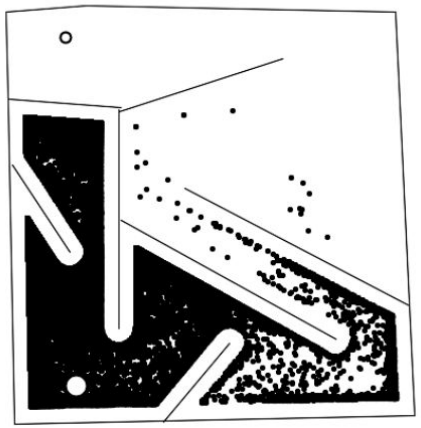

Hard Map Fitness

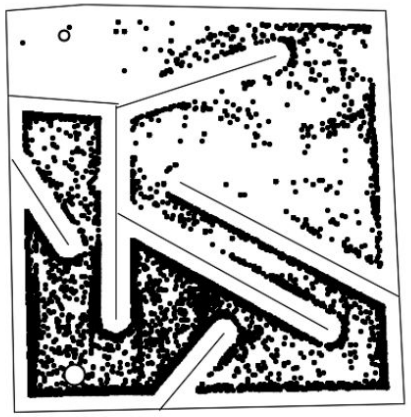

Hard Map Novelty

MAP-Elites[2]

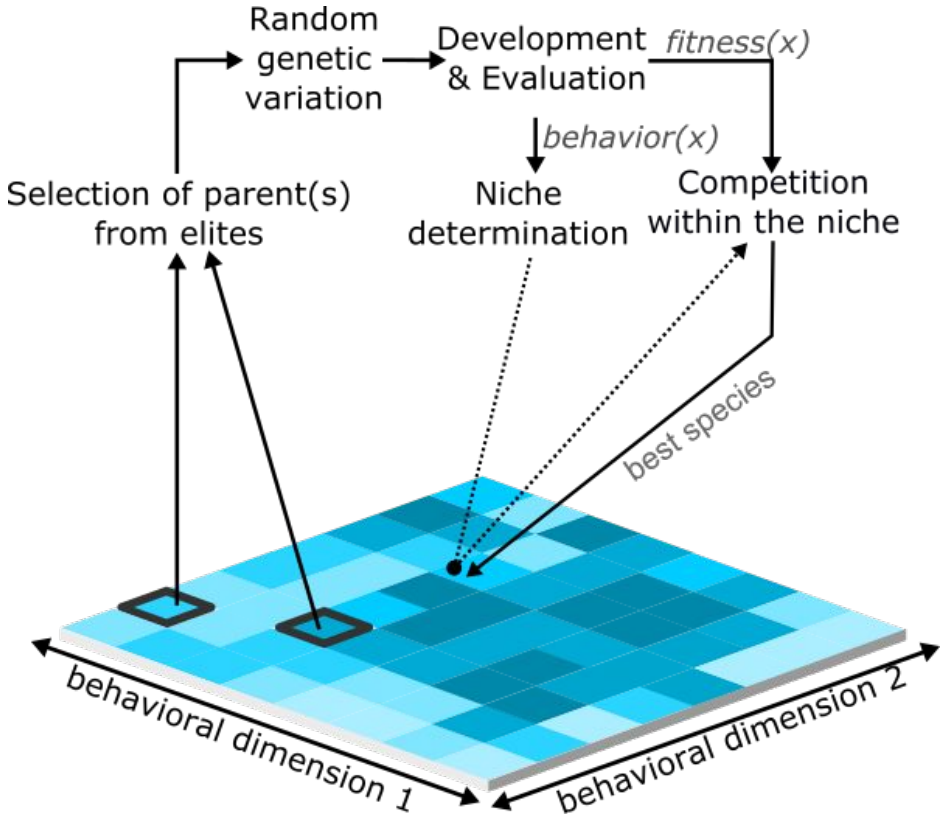

[1] J. Lehman and K. O. Stanley, "Abandoning Objectives: Evolution Through the Search for Novelty Alone," in *Evolutionary Computation*, vol. 19, no. 2, pp. 189-223, June 2011, doi: 10.1162/EVCO_a_00025.

[2] Mouret, Jean-Baptiste, and Jeff Clune. "Illuminating search spaces by mapping elites." *arXiv preprint arXiv:1504.04909* (2015).

Do you need explicit diversity maintenance?

- Quality diversity techniques **explicitly** maintain diversity

- Do we need this?
- Does nature do this?

Lexicase improves population diversity[1]

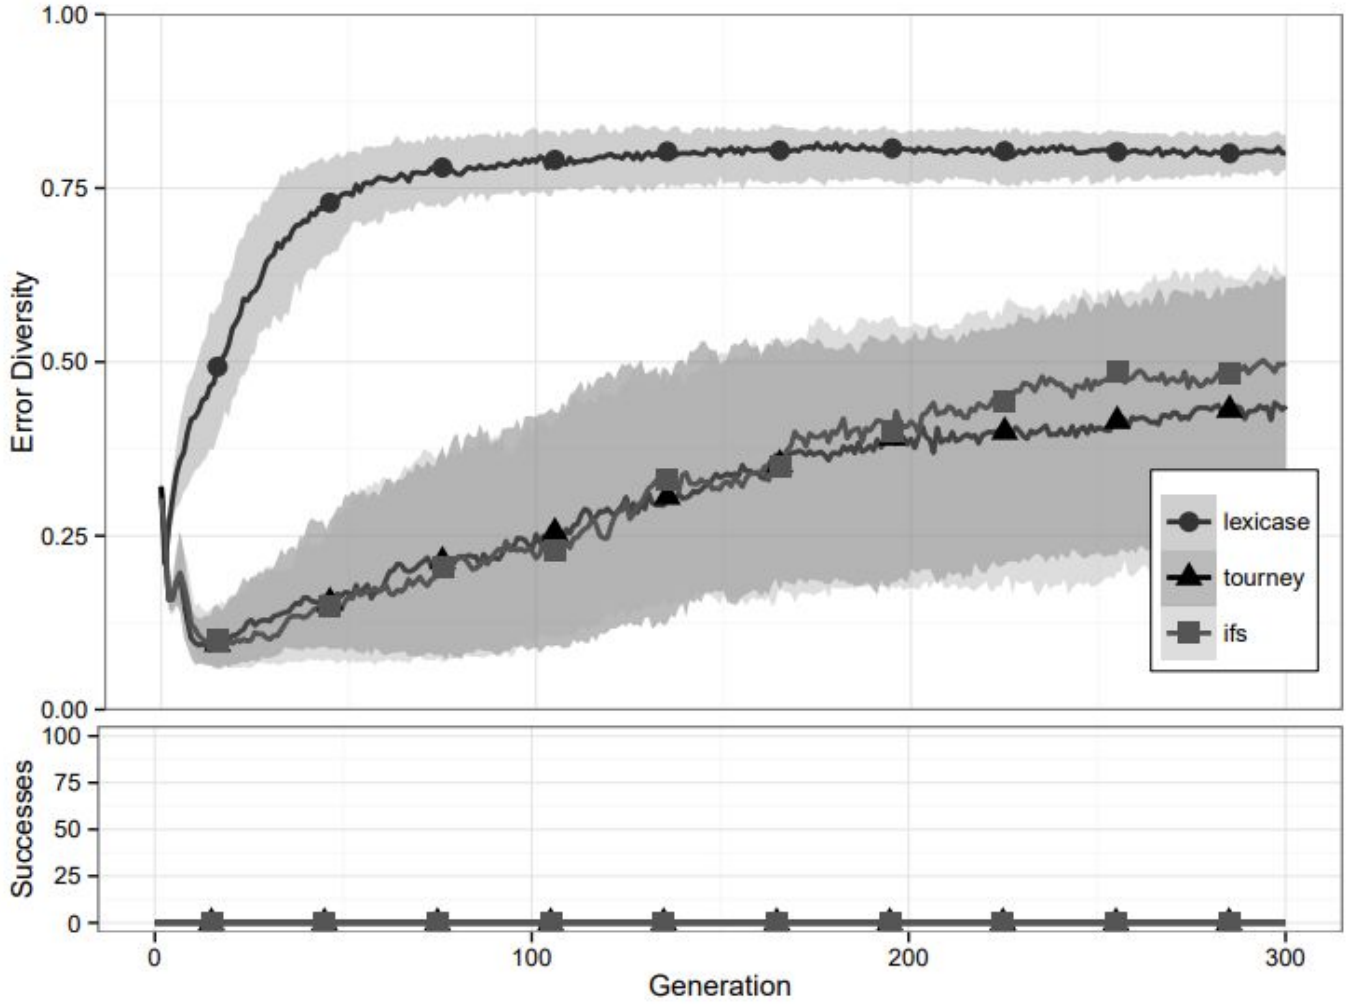

Fig. 13 Checksum – error diversity

[1] Helmuth, T., McPhee, N.F., Spector, L. (2016). Lexicase Selection for Program Synthesis: A Diversity Analysis. In: Riolo, R., Worzel, W., Kotanchek, M., Kordon, A. (eds) Genetic Programming Theory and Practice XIII. Genetic and Evolutionary Computation. Springer, Cham. https://doi.org/10.1007/978-3-319-34223-8_9

Lexicase improves population diversity[2]

When starting from low diversity

When starting from high diversity

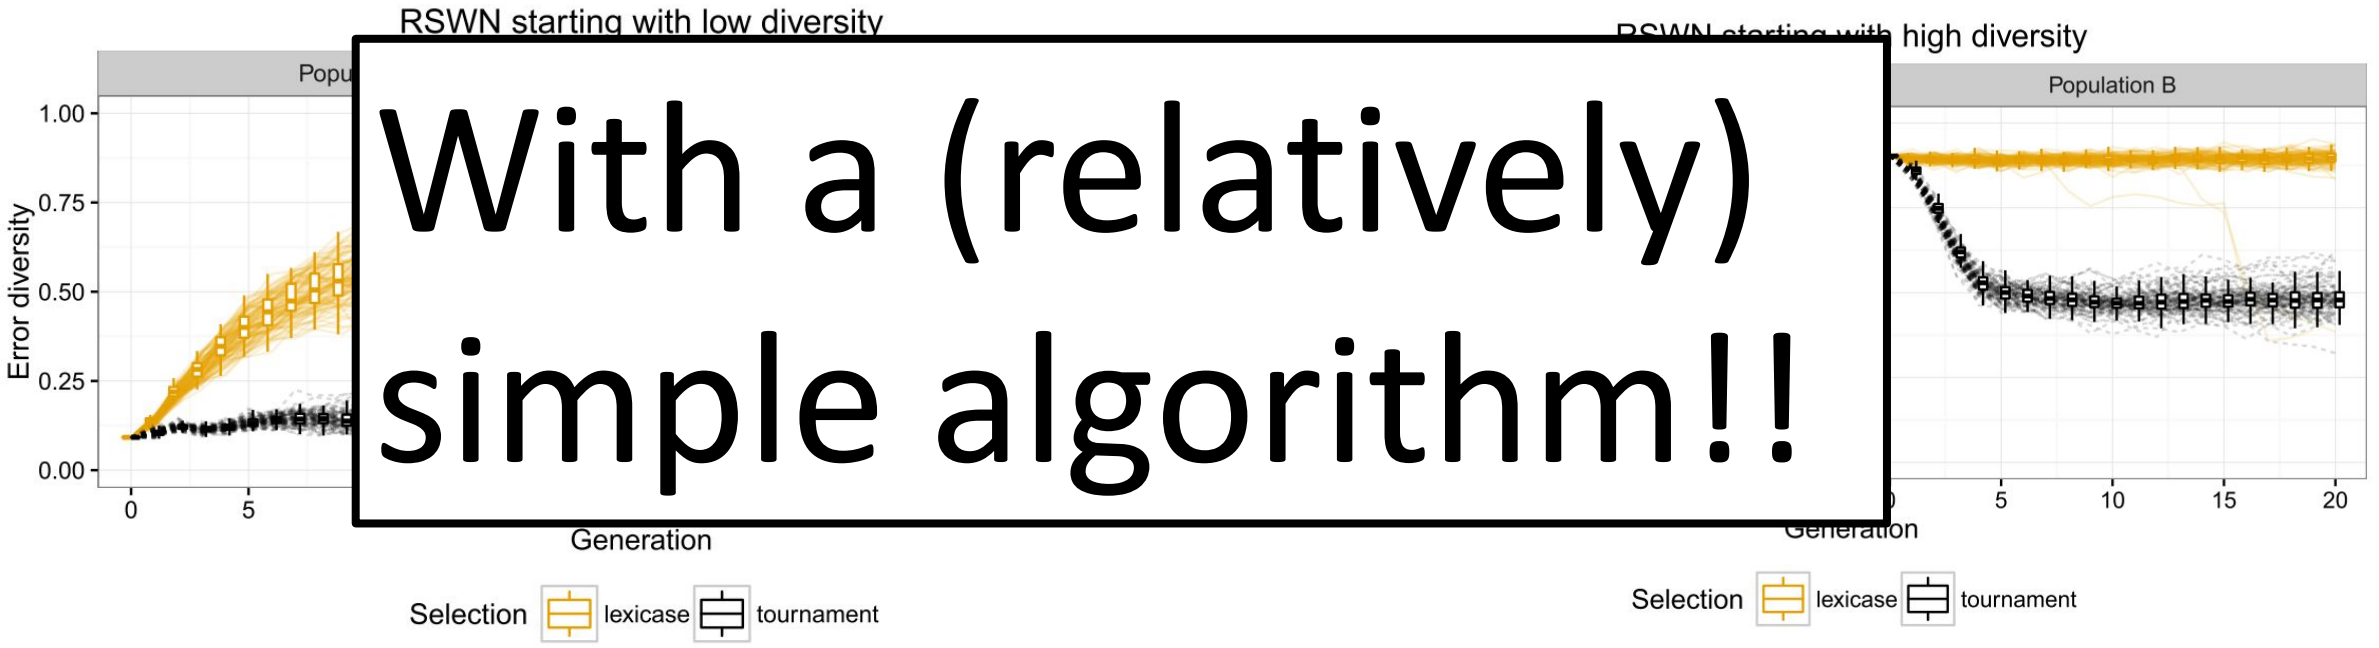

With a (relatively) simple algorithm!!

[2] Thomas Helmuth, Nicholas Freitag McPhee, and Lee Spector. 2016. Effects of Lexicase and Tournament Selection on Diversity Recovery and Maintenance. In Proceedings of the 2016 on Genetic and Evolutionary Computation Conference Companion (GECCO '16 Companion). Association for Computing Machinery, New York, NY, USA, 983–990. https://doi.org/10.1145/2908961.2931657

Many Objectives Facilitate the Emergence of Diversity

- Automatically
- When you treat them in a random order
- When you evaluate on them recursively

What Objectives Do We Use?

Genetic Programming – Training Cases

Reinforcement Learning – Environment Features

Deep Learning – Training Data

Particularities[1]

[1] Lee Spector, Li Ding, and Ryan Boldi. 2023. Particularity. In Genetic Programming Theory and Practice XX. New York: Springer. To appear

Downsampling

Downsampling: Using only a portion of the training cases for computational reasons.

What portion do we use?

- Random[1]
- Informed[2]

[1] Hernandez, Jose Guadalupe, et al. "Random subsampling improves performance in lexicase selection." *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 2019.

[2] Ryan Boldi, Martin Briesch, Dominik Sobania, Alexander Lalejini, Thomas Helmuth, Franz Rothlauf, Charles Ofria, and Lee Spector. 2023. Informed Down-Sampled Lexicase Selection: Identifying productive training cases for efficient problem solving. https://arxiv.org/abs/2301.01488

Informed Downsampling[1]

Intuition: Find cases that measure sufficiently different behaviors
- Build down-samples that maintain diverse cases

- Improved success rates of GP runs[1,2]

- Helped maintain better test coverage than random down-sampling[1]

- Improved problem-solving performance when using **lexicase selection** more than other techniques[2]

[1] Ryan Boldi, Alexander Lalejini, Thomas Helmuth, Lee Spector. 2023. A static analysis of informed down-samples. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '23).

[2] Ryan Boldi, Ashley Bao, Martin Briesch, Thomas Helmuth, Dominik Sobania, Lee Spector, Alexander Lalejini. 2023. The Problem Solving Benefits of Down-Sampling Vary by Selection Scheme. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '23).

Extracting Objectives Directly From Environments[1]

Train autoencoder on phenotypes or environmental interactions

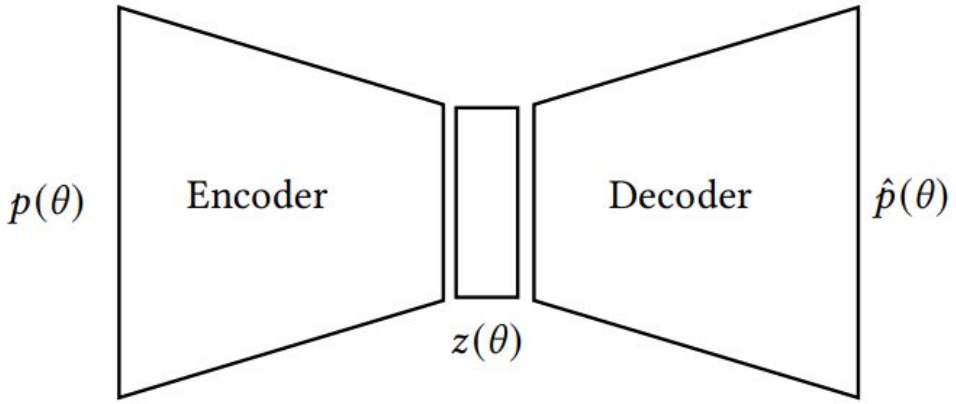

Learn linear combination of features that predicts fitness best

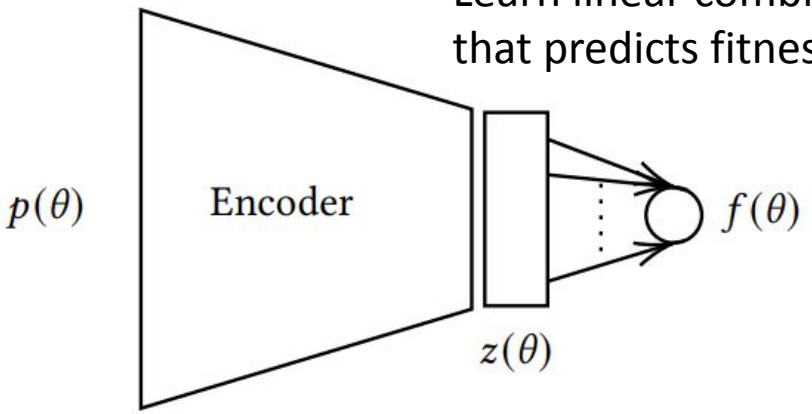

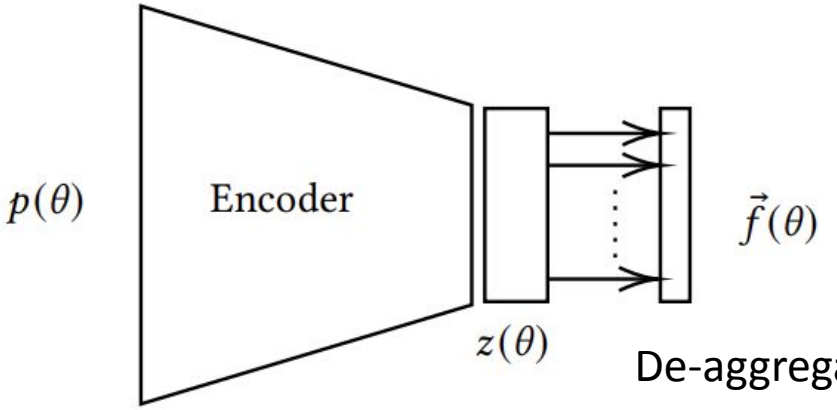

De-aggregate!

[1] Ryan Boldi and Lee Spector. 2023. Can the Problem-Solving Benefits of Quality Diversity Be Obtained Without Explicit Diversity Maintenance? In Genetic and Evolutionary Computation Conference Companion (GECCO `23).

Extracting Objectives Directly From Environments[1]

- Can be thought of as informed down-sampling
- Has been found to make lexicase selection competitive with some quality diversity techniques[1]
- Works to improve diversity of recommendations produced by recommendation engines[2]

[1] Ryan Boldi and Lee Spector. 2023. Can the Problem-Solving Benefits of Quality Diversity Be Obtained Without Explicit Diversity Maintenance? In Genetic and Evolutionary Computation Conference Companion (GECCO `23).

[2] Ryan Boldi, Aadam Lokhandwala, Edward Annatone, Yuval Schecter, Alexander Lavrenenko, Cooper Sigrist. 2023. Improving Recommendation System Serendipity Through Lexicase Selection. https://arxiv.org/abs/2305.11044

Conclusions

Lexicase selection is a powerful problem-solving technique

With the right particularities:

- It can be widely applicable and successful
- It can facilitate the spontaneous emergence of diversity

How we find the right particularities is still an open research question.

Thanks!

Alex Lalejini, Charles Ofria, Dominik Sobania, Emily Dolson, Franz Rothlauf, Lee Spector, Martin Briesch, Matt Fontaine, Stefanos Nikolaidis, Thomas Helmuth and the PUSH Lab at Amherst College.

This material is based upon work supported by the National Science Foundation under Grant No. 2117377. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation. This work was performed in part using high performance computing equipment obtained under a grant from the Collaborative R&D Fund managed by the Massachusetts Technology Collaborative.

ALIFE 2023
GHOST IN THE MACHINE

PUSH*lab
programs under selection + heredity

Ryan Boldi
University of Massachusetts Amherst
rbahlousbold@umass.edu

@RyanBoldi
https://ryanboldi.github.io/

Amherst College

University of Massachusetts Amherst