# Open-Ended Dreamer: An Unsupervised Diversity-oriented Neurosymbolic Learner

Claire Glanois, Shyam Sudhakaran, Elias Najarro & Sebastian Risi

ITU, Copenhagen

ICLR 2023 Workshop NeSy-GeMs, 2023

# Motivating unsupervised open-ended program discovery

- *Open-ended (OE) search*: accumulating diverse stepping stones may help tackling more complex challenges (cf. Novelty Search, Quality Diversity).
- *Compositionality for OE*: Ability to re-use, combine, transfer or scale previous discoveries is crucial to both human learning and adaptation... and open-endedness.
- *Program Synthesis*: Programs as compelling encoding of both artefacts and behaviours, by their compositional, versatile, robust, expressive, and interpretable qualities.
  Promising Neuro-Symbolic models for Program Synthesis and Library Learning, such as DreamCoder [EWN+21].

# Proposition: Open-Ended Dreamer

*Research Question:*

*Could such a neurosymbolic model, growing a hierarchical library of programs, be leveraged for unsupervised open-ended program discovery ?*

*Proposition:*

*Open-Ended Dreamer, a first prototype of an unsupervised diversity-oriented neuro-symbolic learner, built upon DreamCoder to support open-ended program discovery.* $\rightsquigarrow$ *Supervision relaxation via Quality Diversity (MAP-Elites [MC15])*

# Iterative Learning Process

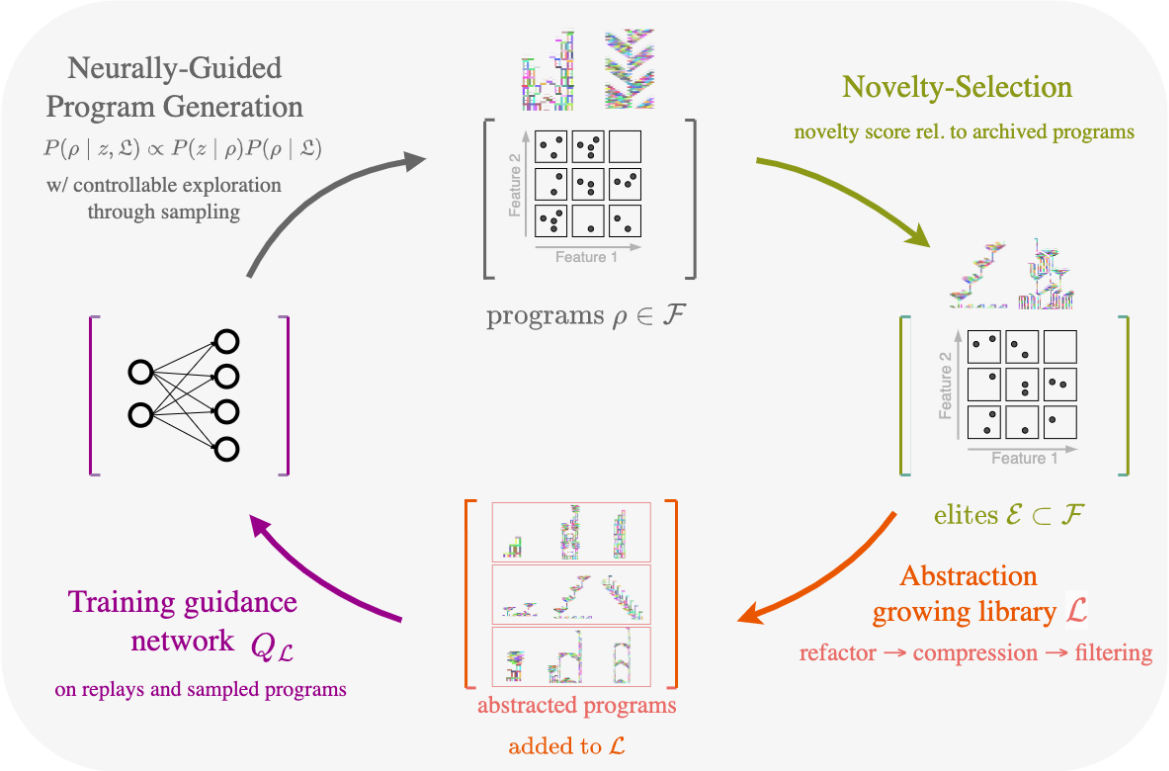


Figure: (1) *Neurally Guided Generation*: by auto-regressively querying the neural model, programs are generated for each niche; (2) *Novelty Selection*: most novel programs are selected; (3) *Abstraction* across niches: library learning stage; (4) *Neural Training*: consolidate the neural guidance on both replays and sampled programs.

# Experiments: Tower Domain

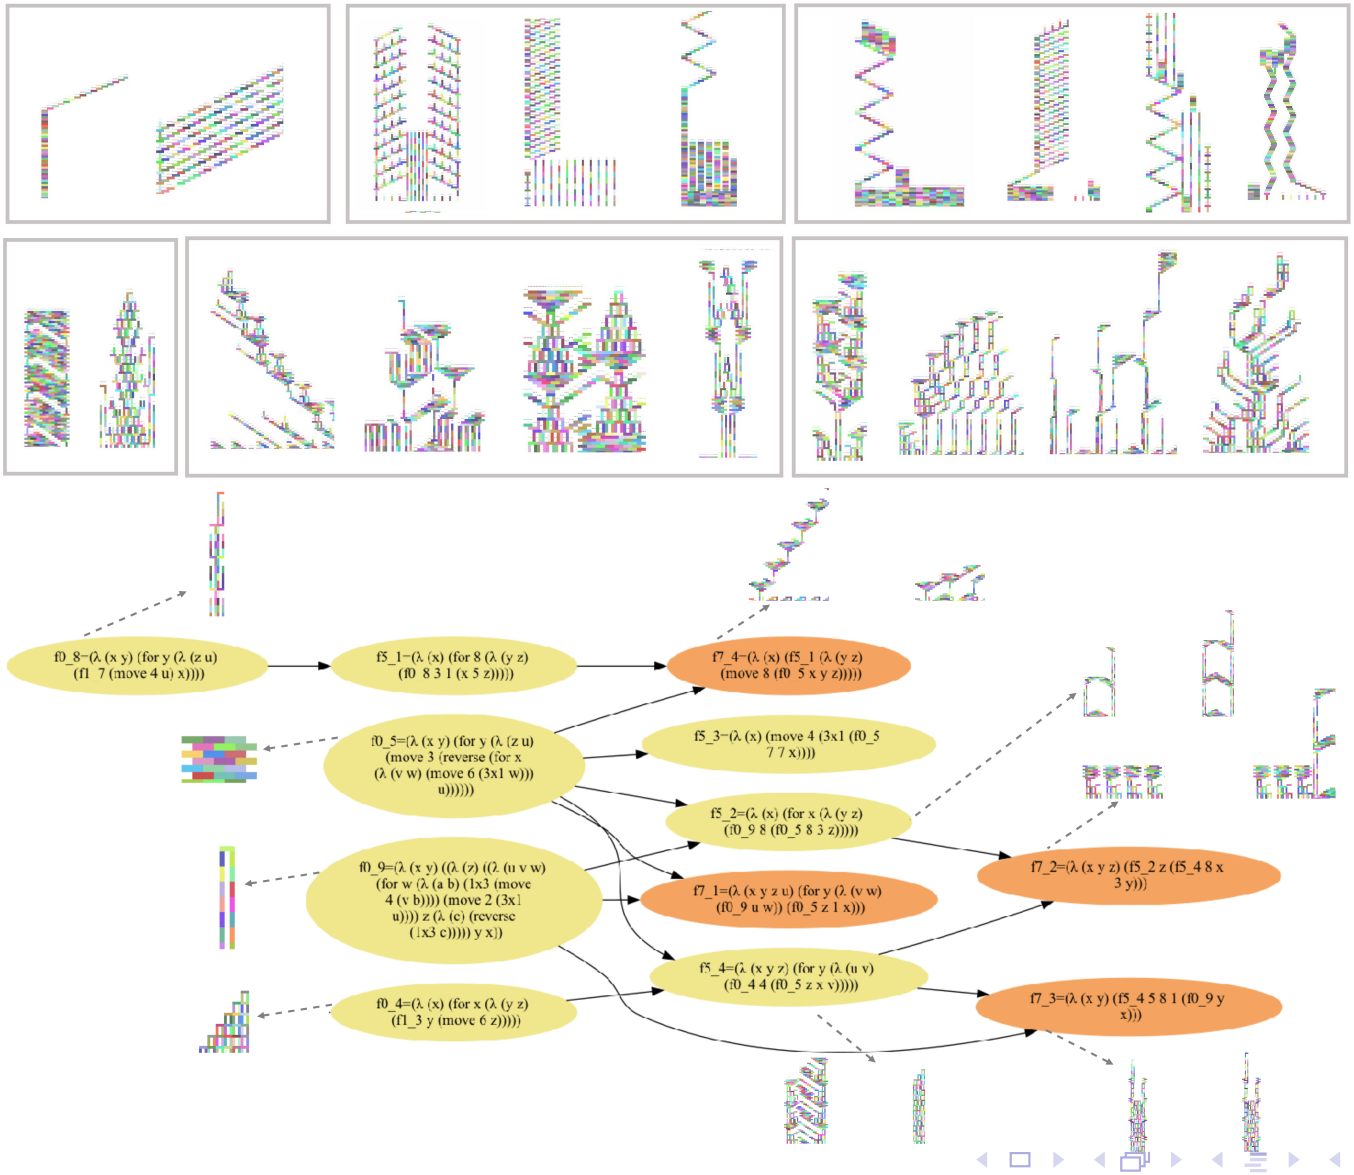

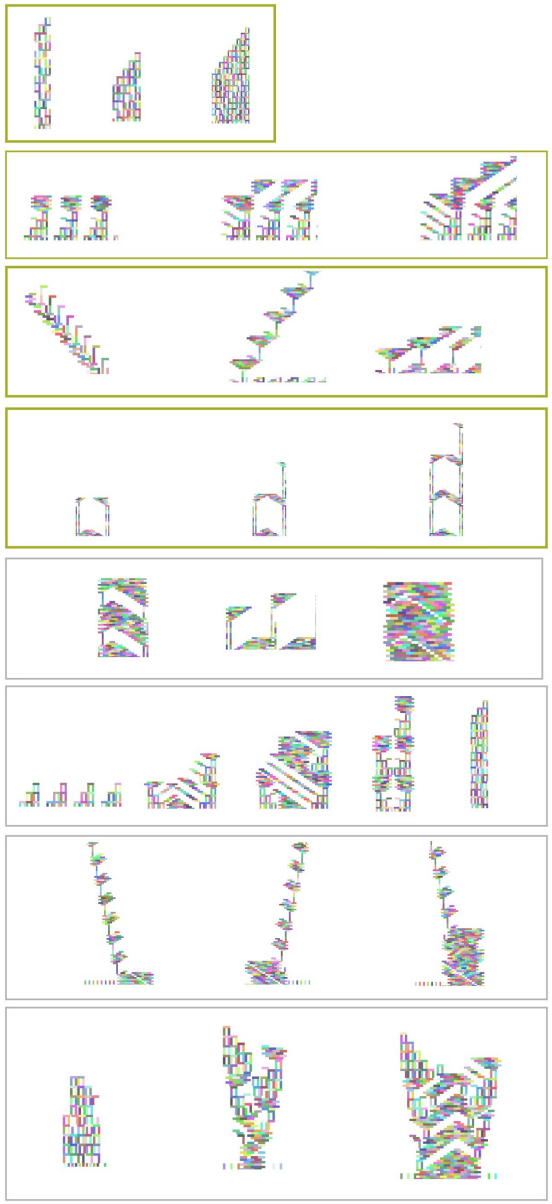

Figure: Library extract, from early evolution (top) to late evolution (bottom)

# Experiment I: Ablation Experiments

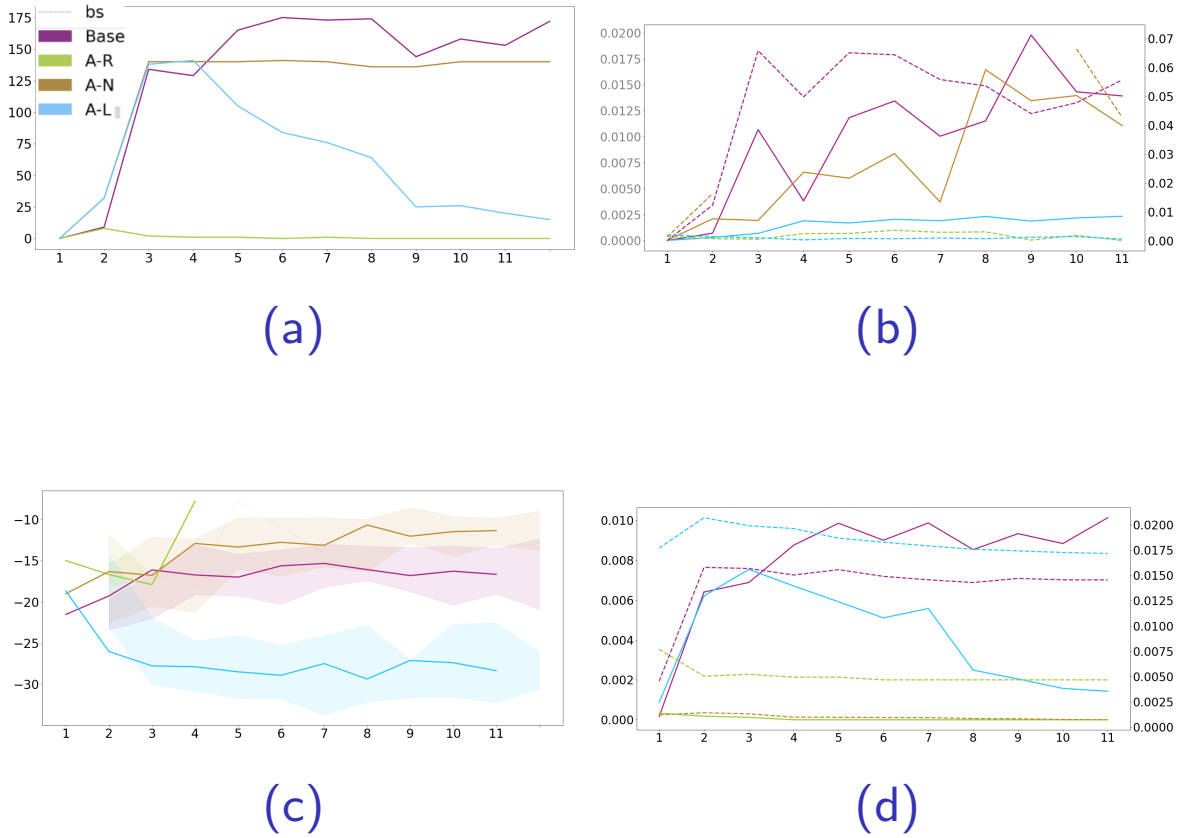


Figure: Ablation Experiments of (A-L) Library, (A-R) Neural guidance, (A-N) Novelty selection. (a) Number of Elites (b) Ratio Legit Programs found, either Bottom Up (left axis) or Top Down (dashed, right axis) (c) Log Priors (d) Novelty Scores.

# Experiment II: Biasing Symbol Learning

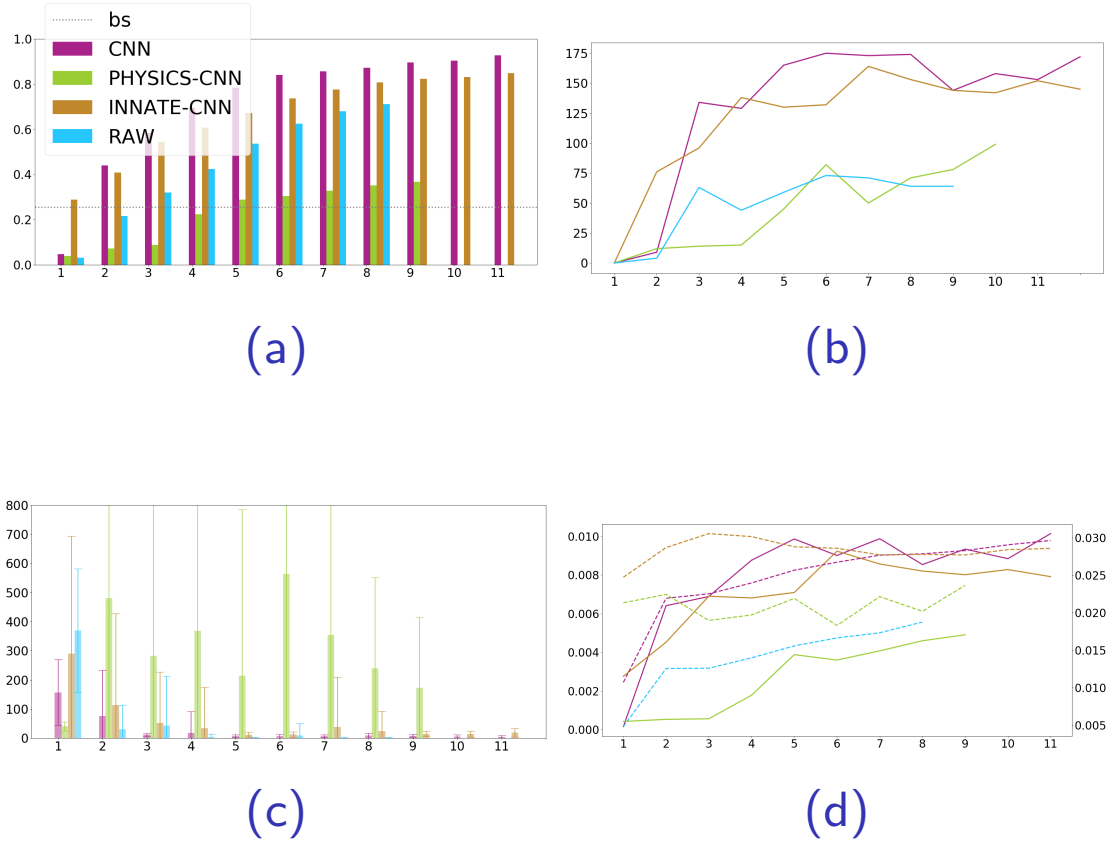


Figure: Bias Study, with (INNATE) bootstrapped library, (CNN) CNN-metric, (PHYSICS) gravity-aware environement, (RAW) downsample metric. (a) Niches Population (b) Number of Elites (c) Enumeration Search Times (d) Novelty Scores.

# A Novelty-Efficiency Trade Off

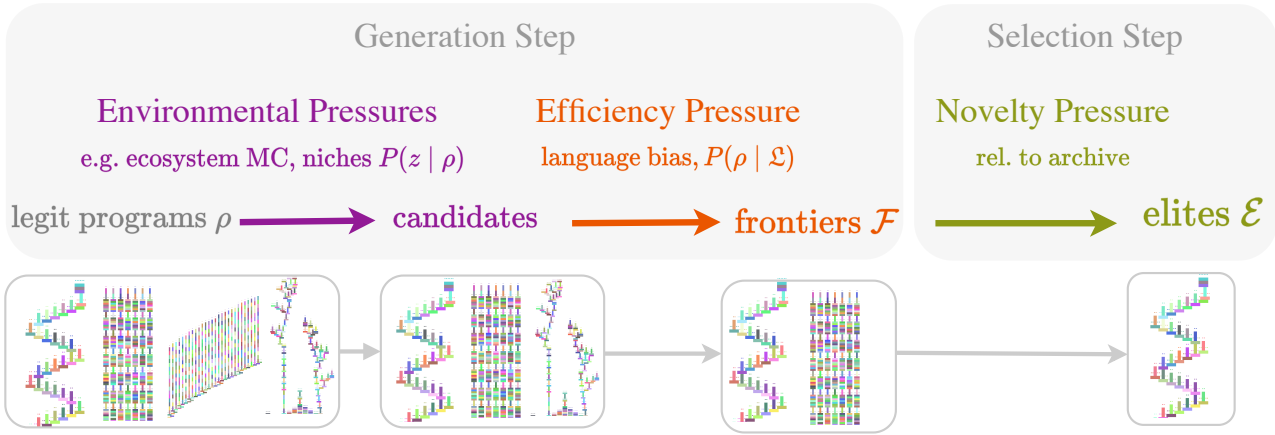


Figure: Different pressures at play in the programs in the diversity-enabling programs search.

- Learning language $\rightsquigarrow$ bias search space $\rightsquigarrow$ may impedes diversity
- Stronger efficiency pressures/novelty $\rightsquigarrow$ convergent evolution phenomenon
- Ways to balance out novelty and efficiency tradeoff with OED hyperparameter (regulating exploration)

# Conclusion

- Diversity-Oriented Search with (Neurally Guided) Library Learning promising framing for open-endedness; preserves regularities in a more potent and versatile way (as programs)
- May guide the development of symbolic knowledge using innate priors, biased metrics, and environmental constraints.
- Promoting greater exploration and stochasticity is crucial to offset the bias introduced by the growing language to adjust the novelty-efficiency trade-off
- Deployment in richer environments (e.g. Minecraft), with functional objectives, instead of aesthetic.

# References


Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sablé-Meyer, Lucas Morales, Luke Hewitt, Luc Cary, Armando Solar-Lezama, and Joshua B Tenenbaum.
Dreamcoder: Bootstrapping inductive program synthesis with wake-sleep library learning.
In *Proceedings of the 42nd acm sigplan international conference on programming language design and implementation*, pages 835–850, 2021.

Jean-Baptiste Mouret and Jeff Clune.
Illuminating search spaces by mapping elites.
*arXiv preprint arXiv:1504.04909*, 2015.