A LITTLE TAXONOMY OF OPEN-ENDEDNESS

Asiiah Song
Department of Computational Media
University of California, Santa Cruz
Santa Cruz, CA 95064, USA
julinas@ucsc.edu

ABSTRACT

This paper aims to provide a partial taxonomy of the ways that the term open-endedness is used in Artificial Life, the ways that open-endedness is used outside of Artificial Life, open-endedness referred to via other terms. The conceptual definitions for open-endedness are rich, while the operational definitions—and the actual measurements—tend to be simplistic. Related concepts that we describe include meta-learning, compositionality (in language), and creativity.

1 INTRODUCTION

What is "open-endedness"? Bedau et al. (2000)'s open problems in Artificial Life (ALife) refers to "open-ended evolution", implying that open-endedness is a feature of evolutionary processes—or at least that it is the open-ended kind of evolutionary process that ALife is interested in. Packard et al. (2019a;b) describes how "open-endedness" over time becomes an important idea in its own right: although biology is believed to be open-ended and "creative production of ongoing novelty" tends to be called open-ended evolution, open-endedness is not necessarily biological or evolutionary. Stepney (2021); Stanley (2019) describes open-endedness as a major goal and grand challenge in ALife.

Two distinct types of definitions for open-endedness currently exist in ALife: the conceptual and the operational. The conceptual definitions in some instances describe relative levels of open-endedness (Banzhaf et al., 2016; Taylor, 2019), and in other instances place open-endedness in categories depending on the context in which open-endedness is considered (Packard et al., 2019a;b; Hintze, 2019). Packard et al. (2019a) notes that the term "open-endedness" is subjective and that it is best to identify which definitions of open-endedness—possibly multiple at once—is being used before diving into any discussion; it provides a partial list of definitions to start with. This plurality of definitions for open-endedness is not only in ALife, but also in robotics, deep learning, computational creativity, and so on. The operational definitions are commonly described in the evaluation of open-ended systems; these operationalizations often refer to the same concepts of novelty, diversity, and complexity without further discussion of conceptual frameworks of open-endedness.

We aim to provide a partial taxonomy of the ways that open-endedness is used in ALife, the ways that open-endedness is used outside of ALife, open-endedness referred to via other terms, as well as related but different concepts. We are not attempting to provide a genealogy or a history. Most of the papers cited are found via a forward-and-backward bibliography walk from Taylor (2019). The intended contribution of this paper is to help researchers answer questions such as the following:

- **Is my goal open-endedness for open-endedness' sake?** Which definition of open-endedness? Can this be implemented given certain computational resources? If the open-endedness of an implementation is contingent, is it measurable? How good of an approximation are my best measurements?
- **Is my goal the implications of open-endedness?** Does the implication logically or otherwise provably hold, or is just assumed by convention? What definition of open-endedness is required? Can the "implied characteristics" be directly measured? Can this be implemented given my computational resources? If the open-endedness of an implementation is contingent, is it measurable? Does the approximation of open-endedness using my best measurements imply those other desired characteristics?

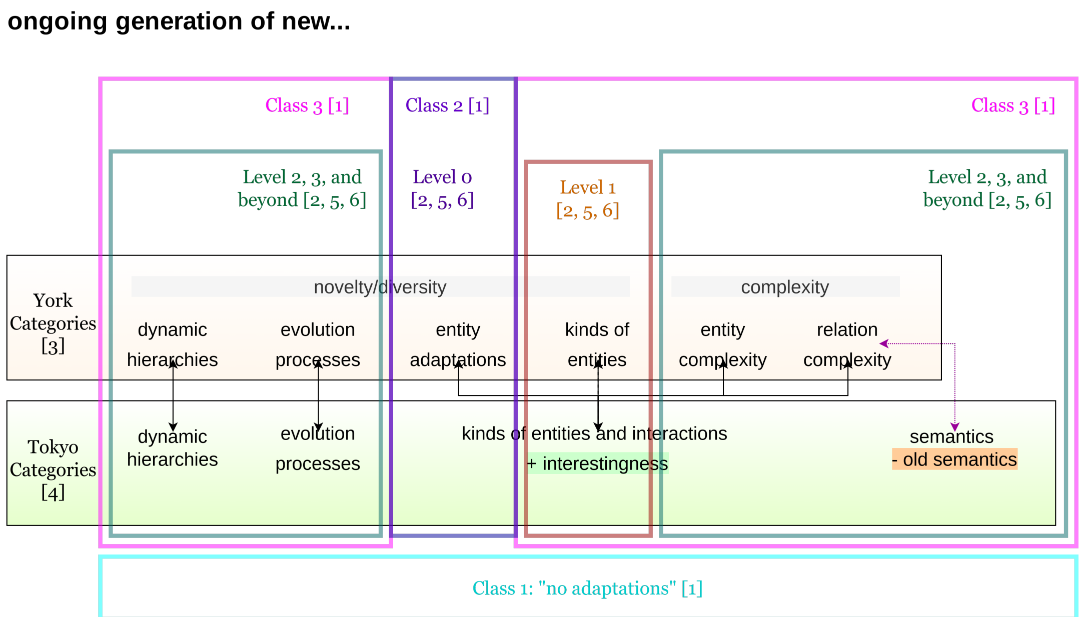

Figure 1: Several different conceptual frameworks of open-endedness. (1) Bedau et al. (1998); (2) Banzhaf et al. (2016); (3) Packard et al. (2019a); (4) Packard et al. (2019b); (5) Taylor (2019); (6) Stepney (2021)

- Is my goal actually one or more of the proxy characteristics used to measure open-endedness – novelty, diversity, complexity, or some other characteristic? Would the metrics used to measure novelty itself be the same as the metrics used to measure open-endedness-by-way-of-ongoing-novelty? Why not directly measure novelty, diversity, complexity, etc.?

2 Definitions of open-endedness

2.1 Open-endedness in artificial life: conceptual

See Figure 1 for a diagram of the conceptual definitions discussed in the following paragraphs.

Open-ended processes are ones that are ongoing and do not have a specific objective. This ongoing change should be unbounded rather than asymptotic. This open-endedness applies to the phase of a dynamical system before it converges, but does not require the dynamical system to never converge. (Banzhaf et al., 2016; Taylor et al., 2016) Without hypothesizing a theoretical dynamical system that never converges, it is obvious that such a system is ultimately bounded by resources in the physical or simulated world. (Ray, 1991) When considering simulated open-ended processes, it should not matter if the processes start simple and grow into an environment that allows for complex processes, or if the processes and the environment co-evolve. However, the environment should be complex enough to allow for the contained processes to be open-ended. (Taylor, 2015). In practice, simulations are ultimately limited by the programming language used.

Bedau et al. (1998) describes three classes of evolutionary processes: Class 1. no adaptations; Class 2. adaptations with bounded diversity (the number of atomic innovations in the system); Class 3. adaptations with unbounded diversity. Class 2 corresponds to Banzhaf et al. (2016)'s variation, and Class 3 corresponds to Banzhaf et al. (2016)'s innovation and emergence. The diversity of a system could simply be thought of as the novelty of each entity in relation to the rest of the entities in the system aggregated in some fashion.

Banzhaf et al. (2016) finds that various definitions of open-endedness in the literature boil down to novelty and complexity. Further, they find that novelty is a necessary but not sufficient requirement, while complexity is a sufficient but not necessary requirement. They extend the novelty requirement by adding the concept of the model and the meta-model, which results in three levels of novelty in relation to the most recent state of the model and meta-model: novelty within the model (variation), novelty that changes the model (innovation), and novelty that changes the meta-model (emergence). An open-ended system is one that can continually generate innovation and emergence events.

Taylor (2019) adapts the model in Banzhaf et al. (2016) and also describes three levels of novelty, but in relation to the initialization state of the model and meta-model: novelty within the model (exploratory), novelty that changes the model (expansive), and novelty that changes the meta-model (transformational). All three levels indicate some kind of open-endedness in this system, and the boundary between the levels can be fuzzy. The examples, given in a simulated creature evolution domain, are: from short limbs to long limbs (exploratory), from no limbs to limbs (expansive), and from running legs to flying wings (transformational). Additionally, the three levels may be evaluated in different contexts: the novelty of an event may be different in the genotype state space compared the the novelty of the same event in the phenotype state space. The question of whether open-endedness is an evaluation of process or outcome can be answered by the plurality of state spaces: the outcome is merely the mapping of the process in another state space. This also helps absorb some definitions of open-endedness into the categories discussed here. For example, Channon (2019) uses three classifications: unbounded diversity of components but bounded adaptive success; bounded diversity of components but unbounded adaptive success; unbounded diversity of components and unbounded adaptive success. The components (often the genome) and adaptive success (behaviors resulting from phenotypes) can be interpreted using the plurality of state spaces.

Stepney (2021), also adapting Banzhaf et al. (2016) and taking inspiration from Kleppe et al. (2003), suggests having levels beyond emergence—novelty that changes the (meta-)meta-meta model—as long as it's possible to identify such meta-ness. They suggest using a discontinuity in novelty value (or derivatives thereof) to identify the boundary between change to a model and change to a meta-model, making an analogy to phase change in physics. Sommerer & Mignonneau (2000); Sayama (2019); Meyerson et al. (2022) similarly propose the ideas "phase transition", "cardinality leap", and "miracle jump", respectively.

Taylor et al. (2016); Packard et al. (2019a) report a set of categories for open-ended evolution (the York Categories), the result of a workshop in York, UK on this topic. According to their survey, definitions of open-endedness are usually some combinations of these categories, which fall in two meta-categories: the ongoing generation of novelty and the ongoing generation of complexity. In their survey, definitions can have labels from both categories. The *novelty* aspect includes: (a) ongoing generation of new adaptations (within entities); (b) ongoing generation of new kinds of entities; (c) ongoing major transitions (Szathmáry & Smith, 1995); (d) evolution of evolvability (Pattee & Sayama, 2019). The first two types of novelty-based open-endedness (a, b) could map to the first two levels in Banzhaf et al. (2016); Taylor (2019); Stepney (2021). The third (c) can be interpreted as aggregating models and meta-models and valuing the both the whole and its internal layers and transitions all together. The fourth (d) refers to evolution of the programmed evolution process, but can also be interpreted conceptually that there is always the possibility of wrapping existing aggregates and designating novelty one layer above that—depending on how one reads the "programmed evolution process". The *complexity* aspect includes: (a) ongoing growth of entity complexity and (b) ongoing growth of interaction complexity. One could also extend this to include the ongoing growth of complexity of the aggregate in dynamic hierarchies, in terms of additional levels of sub-entities in the aggregate as wells as more complex intra-entity interactions.

Packard et al. (2019b) reports a revised set of categories (the Tokyo Categories), the result of a workshop in Tokyo, Japan on this topic. Those are the *ongoing generation* of: (a) interesting new kinds of entities and interactions, (b) evolution of evolvability, (c) major transitions, and (d) semantic evolution (Ikegami et al., 2019). The first three categories are supposed to subsume the ones in the York Categories. (d) refers to the shifting semantic information in a system, and has been identified from analyzing human-produced evolution processes. However, taking Kockelman (2020)'s interpretation of Saussurean meaning, that the meaning of a sign is "constituted by the relation (of combination or substitution) between relations (qua signifier-signified pairings)", (d) could be restated as "ongoing change in relation complexity", which is not far removed from the "ongoing growth in relation complexity" used in the York Categories. With (a), the authors account for the implication in a general understanding of open-endedness for *interestingness* rather than just novelty and complexity, while acknowledging that interestingness is a vague term that needs to be further specified in future work. It makes sense to collapse several categories from the York Categories,

novelty
• Euclidean k-nearest neighbor in feature space (Lehman et al., 2008) • count of immediate neighbors in a discrete feature space (Cully & Demiris, 2017) • count of new adaptive innovations divided by the count of all current adaptive innovations (Bedau et al., 1998)
diversity
• Euclidean k-nearest neighbors in feature space (Lehman et al., 2008; Lehman & Stanley, 2011b; Pugh et al., 2016; Mouret & Clune, 2015; Brant & Stanley, 2017) • Shannon complexity of the distribution of component size in the system (Bedau et al., 1998) • total count of persisting components in the system (Bedau et al., 1998) • total count of types of components in the system (Bedau et al., 1998) • pair-wise Levenstein distance (Hintze, 2019) • human eye (Hintze, 2019)
complexity
• grammar-based compression (Mondol & Brown, 2021) • sliding-window compression (Hintze, 2019) • jpeg compression (Earle et al., 2021) • (requires approximation) Kolmogorov complexity—modification: 0 if sequence is random (Standish, 2003) • (requires approximation) statistical complexity (Stepney, 2021) • Shannon complexity (Standish, 2003) • the count of atomic functions used to compose a more complex function (where the complex function has been generated by the program to start with) (Wang et al., 2019) • human eye (Hintze, 2019)

Table 1: metrics used to calculate properties contributive to open-endedness

because in practice it is difficult to determine the line between the adaptation of an entity and a new kind of entity, etc.

Conceptual treatment of open-endedness appears to come either in terms of mutually-exclusive categories (where individual evolutionary systems can straddle multiple categories) or levels (where each subsequent model level indicates some additional meta-quality compared to the previous level). Within the mutually-exclusive categories, each conceptual category could be associated with novelty, diversity, complexity, or some combination thereof; novelty, diversity, and complexity are in turn fuzzy terms that map to "new-ness" of entity characteristics and entity relations. See Figure 1.

2.2 Open-endedness in artificial life: measurements

See Table 1 for a list of algorithms used to calculate novelty, diversity, and complexity. In the papers reviewed, projects that involve actual calculations turn out to be one of the three York Categories: new entity adaptations, new kinds of entities, and entity complexity.

Bedau et al. (1998) suggests that diversity could be measured by a) Shannon complexity (Shannon & Weaver, 1949) of the distribution of sizes of components in the system, b) the count of persisting components in the system, c) the count of types of components in the system, or d) the count of new adaptive innovations divided by the count of all current adaptive innovations. Using the diversity of the system at different points of time, calculated using the second metric, one could then calculate the adaptive evolutionary activity statistics, or ongoing diversity, of the system. The weakness of the system is that what is considered a component in the system has to be decided before hand, in order to count their activity. This weakness may be resolved by having an automated algorithm that takes as input the observation of a system and outputs deciding criteria for what should be considered a component in that system.

Standish (2003) highlights the context-dependence of calculating complexity (e.g., whether it is based on genotypes or phenotypes) and proposes to measure complexity in general terms by using Shannon's information complexity (Shannon & Weaver, 1949) or Kolmogorov's algorithmic complexity (Kolmogorov, 1965). One problem with the latter is that random sequences have maximum Kolmogorov complexity, which is contrary to the human intuition that random sequences are not interesting, therefore they do not feel complex. To solve this problem, Standish (2003) changes the complexity of random sequences to 0, while keeping other parts of the Kolmogorov complexity the same. This is a simpler version of statistical complexity in Crutchfield (1994).

Also consider the humble cockroach, which is said to be able to adapt to many more environmental events than humans can; consider the thought experiment of monkeys on typewriters having a miniscule but non-zero chance of generating Shakespeare's poetry. Both are examples of a seemingly simple descriptive/genotypical space that project into a seemingly richer behavioral/phenotypical space. If the Shakespearean poetry is possible from the random typing, should random typing really have a complexity of 0? Perhaps complexity could be divided into potential complexity and structural complexity.

Lehman & Stanley (2012) discusses impressiveness in open-ended evolutionary systems. Impressiveness is closely related to interestingness and beauty; interestingness decreases over time, while impressiveness and beauty do not. They consider interestingness and beauty orthogonal to compressibility and make an analogy to P/NP, where an impressive product would be "easier to appreciate than to create", correlating to an NP-complete problem in complexity theory. To measure the open-endedness of their system, the authors use the criterion of novelty, operationalized as sparsity. More specifically, they calculate the novelty of an individual as the average Euclidean distance to the k-nearest neighbor of where that individual is located in the n-dimensional space of relevant features, also known as the behavioral characterization, or solution descriptors. (Lehman et al., 2008) In their experiment, they use the following eight features: brightness, BZip2 compression, wavelet compression, color variety, x-axis symmetry, y-axis symmetry, and choppiness. Several quality diversity algorithms Lehman & Stanley (2011b); Pugh et al. (2016); Mouret & Clune (2015) also use sparsity in the n-dimensional feature space to determine diversity. Brant & Stanley (2017) evaluates the minimum criterion coevolution algorithm on diversity, again via sparsity in the feature space where a relevant feature is the point in each agent's trajectory at a given time step.

Another quality diversity algorithm using the MAP-Elite multi-dimensional grid, Cully & Demiris (2017), calculates novelty using a simplification of sparsity: the count of the filled cells surrounding the location of the individual. That is, the more filled cells there are, the less novel it is, which sounds a lot like a gravitational search algorithm Rashedi et al. (2009) without attraction forces.

Hintze (2019)'s open-endedness includes diversity and complexity. They measure diversity with pair-wise Levenstein distance (Levenshtein et al., 1966) and to measure complexity with Kolmogorov complexity (Kolmogorov, 1965) by proxy of compression (Ziv & Lempel, 1978). Levenstein's distance is a rough approximation of diversity as well as a rough approximation of novelty, so one could interpret this model as, again, reducing to novelty and complexity. While they use sliding-window compression to approximate Kolmogorov complexity, some kind of grammar-based compression (Mondol & Brown, 2021) might be a better approximator of Kolmogorov complexity. Stepney (2021) complains that Kolmogorov complexity, even when accurately calculated, would be a measure of randomness rather than complexity in the general sense, and that it would be better to use something like statistical complexity to measure complexity. (Stepney, 2021) Statistical complexity is similar to Standish (2003)'s modified Kolmogorov complexity, but more specifically states that $Kolmogorov\_complexity(L) \approx statistical\_complexity(L) + Shannon\_complexity(L)$ for a string L. (Crutchfield, 1994)

3 Open-endedness in other fields, by other names

Lehman et al. (2008); Lehman & Stanley (2011a) come from the point of view of ALife and apply novelty to solve maze path-finding problems. For the authors, complexity comes necessarily after a continuous novelty search, because there are a limited number of simple novel behaviors in this domain; and the complexity allows the path-finding problem to be solved. They measure novelty using sparsity.

Earle et al. (2021) uses jpeg compression as a proxy for complexity in the domain of simulated artificial life in video games. Wang et al. (2019) uses co-evolved environments and agents to solve terrain walker problems. They observe that the after some evolution, the agents are able to walk relatively gracefully (compared to algorithms) in more complex domains. In an evolved environment like this, the count of smaller modular functions used to compose a certain terrain could be a good proxy for Kolmogorov complexity, although it is unclear if this is what the authors used.

In deep learning, approaches such as meta-learning Finn et al. (2017) and one-shot learning Vinyals et al. (2016) appear to have comparable goals to open-endedness of being able to embody changes to the meta-model in a system without explicitly modeling the meta-model, although the term open-endedness is not used, and the criteria for success are very different. It might be said that the deep learning varieties are more focused on the product of their systems, while open-endedness as a research goal is an idea that focuses on the correctness of the architecture and process.

Boden (2004)'s criteria for human creativity include surprise, novelty, and usefulness; couple this with Schmidhuber (2008)'s theory that usefulness and beauty correlate to compressibility, and the fact that compression is often used to estimate complexity, we arrive at a loose equivalence between creativity and open-endedness. Similar to Bedau et al. (1998); Banzhaf et al. (2016); Taylor (2019), Boden (2004)'s creativity has a triad of increasing meta-creativity: exploratory, combinatorial, and transformative creativity. The claims of open-endedness are weaker than that of Boden (2004)'s human creativity, and requires one of novelty, diversity, and complexity, while having the contingency for the other qualities.

Cellular automata (Von Neumann et al., 1966) such as the Game of Life correspond to exploratory or expansive open-endedness. Relatedly, Mange et al. (2004) implements Lindenmayer systems (L-systems) which are basically the expansion of context-free grammars. L-systems might exhibit exploratory open-endedness. A well-specified grammar (such the grammar in a natural language) could have compositionality, which is a desired goal in emergent language systems (Lowe et al., 2019). The idea of compositionality in emergent language systems is quite similar to open-endedness in evolutionary systems: both are initially identified as co-existing with natural systems (language and biology, respectively) and deemed important, if not essential to those natural systems. The desires to research compositionality and open-endedness are both sometimes motivated by the idea that perhaps they would, in the simulated systems, lead to desirable properties from the natural systems.

4 Conclusion

This paper presented a taxonomy of the ways that "open-endedness" is used in artificial life/open-ended evolution. We also make comparisons to similar concepts used in other areas. Some of those other areas, such as compositionality in emergent language research, seem to follow similar disjunction between the conceptual definitions, operational definitions, and the actual implementation of measurements. While it does not make an extensive survey, it is hoped that the taxonomy itself is relatively complete and provides illumination for future research.

References

Wolfgang Banzhaf, Bert Baumgaertner, Guillaume Beslon, René Doursat, James A Foster, Barry McMullin, Vinicius Veloso De Melo, Thomas Miconi, Lee Spector, Susan Stepney, et al. Defining and simulating open-ended novelty: requirements, guidelines, and challenges. *Theory in Biosciences*, 135(3):131–161, 2016.

Mark A Bedau, Emile Snyder, and Norman H Packard. A classification of long-term evolutionary dynamics. In *Artificial life VI*, pp. 228–237, 1998.

Mark A Bedau, John S McCaskill, Norman H Packard, Steen Rasmussen, Chris Adami, David G Green, Takashi Ikegami, Kunihiko Kaneko, and Thomas S Ray. Open problems in artificial life. *Artificial life*, 6(4):363–376, 2000.

Margaret A Boden. *The creative mind: Myths and mechanisms*. Routledge, 2004.

Jonathan C Brant and Kenneth O Stanley. Minimal criterion coevolution: a new approach to open-ended search. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 67–74, 2017.

Alastair Channon. Maximum individual complexity is indefinitely scalable in geb. *Artificial life*, 25 (2):134–144, 2019.

James P Crutchfield. The calculi of emergence: computation, dynamics and induction. *Physica D: Nonlinear Phenomena*, 75(1-3):11–54, 1994.

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

Sam Earle, Julian Togelius, and LB Soros. Video games as a testbed for open-ended phenomena. In *2021 IEEE Conference on Games (CoG)*, pp. 1–9. IEEE, 2021.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.

Arend Hintze. Open-endedness for the sake of open-endedness. *Artificial life*, 25(2):198–206, 2019.

Takashi Ikegami, Yasuhiro Hashimoto, and Mizuki Oka. Open-ended evolution and a mechanism of novelties in web services. *Artificial life*, 25(2):168–177, 2019.

Anneke G Kleppe, Jos B Warmer, Jos Warmer, and Wim Bast. *MDA explained: the model driven architecture: practice and promise*. Addison-Wesley Professional, 2003.

Paul Kockelman. *3. Sassurean Values*. Prickly Paradigm Press, 2020.

Andrei N Kolmogorov. Three approaches to the quantitative definition ofinformation'. *Problems of information transmission*, 1(1):1–7, 1965.

Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, 2011a.

Joel Lehman and Kenneth O Stanley. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pp. 211–218, 2011b.

Joel Lehman and Kenneth O Stanley. Beyond open-endedness: Quantifying impressiveness. In *ALIFE 2012: The Thirteenth International Conference on the Synthesis and Simulation of Living Systems*, pp. 75–82. MIT Press, 2012.

Joel Lehman, Kenneth O Stanley, et al. Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, pp. 329–336. Citeseer, 2008.

Vladimir I Levenshtein et al. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pp. 707–710. Soviet Union, 1966.

Ryan Lowe, Jakob Foerster, Y-Lan Boureau, Joelle Pineau, and Yann Dauphin. On the pitfalls of measuring emergent communication. *arXiv preprint arXiv:1903.05168*, 2019.

Daniel Mange, Andre Stauffer, Leonardo Peparaolo, and Gianluca Tempesti. A macroscopic view of self-replication. *Proceedings of the IEEE*, 92(12):1929–1945, 2004.

Elliot Meyerson, Xin Qiu, and Risto Miikkulainen. Simple genetic operators are universal approximators of probability distributions (and other advantages of expressive encodings). *arXiv preprint arXiv:2202.09679*, 2022.

Tiasa Mondol and Daniel G Brown. Grammar-based compression and its use in symbolic music analysis. *Journal of Mathematics and Music*, pp. 1–18, 2021.

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

Norman Packard, Mark A Bedau, Alastair Channon, Takashi Ikegami, Steen Rasmussen, Kenneth Stanley, and Tim Taylor. Open-ended evolution and open-endedness: Editorial introduction to the open-ended evolution i special issue. *Artificial Life*, 25(1):1–3, 2019a.

Norman Packard, Mark A Bedau, Alastair Channon, Takashi Ikegami, Steen Rasmussen, Kenneth O Stanley, and Tim Taylor. An overview of open-ended evolution: Editorial introduction to the open-ended evolution ii special issue. *Artificial life*, 25(2):93–103, 2019b.

Howard H Pattee and Hiroki Sayama. Evolved open-endedness, not open-ended evolution. *Artificial Life*, 25(1):4–8, 2019.

Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.

Esmat Rashedi, Hossein Nezamabadi-Pour, and Saeid Saryazdi. Gsa: a gravitational search algorithm. *Information sciences*, 179(13):2232–2248, 2009.

Thomas S Ray. An approach to the synthesis of life. *Artificial life II*, pp. 371–408, 1991.

Hiroki Sayama. Cardinality leap for open-ended evolution: Theoretical consideration and demonstration by hash chemistry. *Artificial life*, 25(2):104–116, 2019.

Jürgen Schmidhuber. Driven by compression progress: A simple principle explains essential aspects of subjective beauty, novelty, surprise, interestingness, attention, curiosity, creativity, art, science, music, jokes. In *Workshop on anticipatory behavior in adaptive learning systems*, pp. 48–76. Springer, 2008.

Claude E Shannon and Warren Weaver. The mathematical theory of communication. university of illinois. *Urbana*, 117, 1949.

Christa Sommerer and Laurent Mignonneau. Modeling emergence of complexity: The application of complex system and origin of life theory to interactive art on the internet. In *Artificial life VII: Proceedings of the Seventh International Conference on Artificial Life*, pp. 547–554. Citeseer, 2000.

Russell K Standish. Open-ended artificial evolution. *International Journal of Computational Intelligence and Applications*, 3(02):167–175, 2003.

Kenneth O Stanley. Why open-endedness matters. *Artificial life*, 25(3):232–235, 2019.

Susan Stepney. Modelling and measuring open-endedness. *Artificial Life*, 25(1):9, 2021.

Eörs Szathmáry and John Maynard Smith. *The major transitions in evolution*. WH Freeman Spektrum Oxford, UK:, 1995.

Tim Taylor. Requirements for open-ended evolution in natural and artificial systems. *arXiv preprint arXiv:1507.07403*, 2015.

Tim Taylor. Evolutionary innovations and where to find them: Routes to open-ended evolution in natural and artificial systems. *Artificial life*, 25(2):207–224, 2019.

Tim Taylor, Mark Bedau, Alastair Channon, David Ackley, Wolfgang Banzhaf, Guillaume Beslon, Emily Dolson, Tom Froese, Simon Hickinbotham, Takashi Ikegami, et al. Open-ended evolution: Perspectives from the oee workshop in york. *Artificial life*, 22(3):408–423, 2016.

Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016.

John Von Neumann, Arthur W Burks, et al. Theory of self-reproducing automata. *IEEE Transactions on Neural Networks*, 5(1):3–14, 1966.

Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. Paired open-ended trailblazer (poet): Endlessly generating increasingly complex and diverse learning environments and their solutions. *arXiv preprint arXiv:1901.01753*, 2019.

Jacob Ziv and Abraham Lempel. Compression of individual sequences via variable-rate coding. *IEEE transactions on Information Theory*, 24(5):530–536, 1978.