Cultivating Open-Earedness with Sound Objects discovered by Open-Ended Evolutionary Systems

Björn Þór Jónsson[1], Çağrı Erdem[2], Stefano Fasciani[3], and Kyrre Glette[4]

[1,2,4]Department of Informatics, University of Oslo, Norway
[3]Department of Musicology, University of Oslo, Norway

Abstract

Interaction with generative systems can face the choice of generalising towards a middle ground or diverging towards novelty. Efforts have been made in the domain of sounds to enable divergent exploration in search of interesting discoveries. Those efforts have been confined by pre-trained models and single environments. We are building on those efforts to enable autonomous discovery of sonic landscapes. Furthermore, we draw inspiration from research on open-ended evolution to continuously provide evolutionary processes with new opportunities for sonic discoveries. Exposure to autonomously discovered sound objects can elevate openness to sonic experiences, which in turn offers inspiring opportunities for creative work involving sounds.

Introduction

To generalise towards a middle ground or to diverge towards novelty, that is the question. A question worth considering when interacting with generative systems to create artefacts, such as sounds. There is often a desire to generate something unique or distinct from existing creations. That desire may not be satisfied by typing requests into a system that can come up with anything you can imagine. What you can imagine may be limited by what you have already been exposed to, and what generative systems can come up with may be limited by what they have been trained on. Curiosity and the desire to create unexpected outcomes may be better satisfied by a generative system that can come up with novel and interesting discoveries.

Exposure to the right stepping stones is crucial on the path to interesting discoveries and enables a shift away from systems that merely offer combinatorial creativity, which is essential to overcome such impediments to unleashing human creativity. Many of the right stepping stones may seem unlikely, yet they could be necessary components of the pathways that lead to unexpected findings. Moreover, the process of transformational creativity is idiosyncratic; the paths and stepping stones that culminate in transformative discoveries for one individual can diverge significantly from those that ignite inventive creativity in another. Abundant examples from both natural evolution and human cultural history illuminate the diversity of routes that can lead to significant innovations (Stanley and Lehman 2015).

This suggests that to foster the continuous flourishing of human creativity, we should advocate for a diversity of exposure to artefacts that have transformative influences on divergent creative processes. Experiments like Picbreeder demonstrate the power of such diverse exposure in revealing a multitude of paths that lead to compelling discoveries (Secretan et al. 2011). In an era where we spin the wheels of generative systems trained on the vast corpus of current human digital output (Dingemanse 2023), it's more important than ever to cultivate environments that emphasise genuine creation over mere imitation. The focus should be on systems that break the cycle of replication and "produce outputs that are novel or surprising and which yield unexpected value" (Veale, Amílcar Cardoso, and Pérez Y Pérez 2019). Open-Endedness (OE) can illuminate and steer Computational Creativity (CC) through paths towards transformative discoveries by encouraging exploration and the generation of diverse and innovative artifacts in a more dynamic and less constrained manner (Soros et al. 2024).

Sound synthesisers were initially used to mimic the sound of existing acoustic instruments (Mathews and Pierce 1987; Smith III 1991). Their significant impact was realized when their generative capabilities were further investigated to produce distinctive timbres, rendering them as valuable tools in musical exploration. It is interesting to compare this evolution of synthesiser applications to the application of generative AI systems to produce realistic artefacts. Musicians, especially those working with electroacoustic sounds, can emphasise defining their own signature sounds. The process of discovering new sounds with sound synthesis has, though, required dedication and expertise to some degree. Some may possess the determination and dexterity to acquire deep knowledge on the inner workings of their instruments and skill in applying it to shape new sounds. Others may rather be inclined to be pleasantly surprised by new discoveries offered to them.

Here we report on steps we have taken towards enabling such discoveries by exploration with evolutionary algorithms. Furthermore we share our ambitions to open up the evolutionary processes enabled by our current findings and guided by insights from research on open-ended evolution.

Sound Innovation

Our initial explorations into applying Interactive Evolutionary Computation (IEC) (Takagi 2001) to the discovery of sounds (Jónsson, Hoover, and Risi 2015) revealed how challenging that domain is for human evaluators. The limited capacity of human evaluation (McCormack 2005) is particularly noticeable when considering the temporal dimension, as well as the potential mental fatigue that arises from assessing a series of diverse sounds. This can be compared with the palate fatigue that can be experienced in large wine tasting scenarios, where all wines start to taste the same.

To relieve humans from excessive evaluation and accelerate discovery, while still avoiding generalisation towards a middle ground of existing training data, it can be useful to automate exploration through the space of sounds in search of interesting discoveries. To that end, we have implemented a system of modular and reusable components that can be applied in distributed computation configurations (Jónsson, Erdem, and Glette 2024). Different facets of the system focus on generating sounds, measuring the diversity and quality, and driving search through the space defined by those measures.

The main approach offered by the system for generating sounds is based on Compositional Pattern Producing Networks (CPPNs) (Stanley 2007) for generating audio and control signals which are optionally fed through a Digital Signal Processing (DSP) network. CPPNs have been used as an abstract representation of the unfolding development in evolutionary processes, which build a phenotype over time. This can be compared with the process of timbral development, where musical expression depends on changes and nuances over time. The CPPN network outputs have been combined with a DSP network in an effort to broaden the sonic palette of the generative system, and the effect of this coupling was investigated specifically in (Jónsson et al. 2024). Both networks are evolved with NeuroEvolution of Augmenting Topologies (NEAT) (Stanley and Miikkulainen 2002) according to configurable rates. Evolving those networks with NEAT allows the evolutionary processes to start without any pre-training from the simplest possible network configuration, with only input and output nodes, which gradually complexifies during evolution. Exploration is based on Quality Diversity (QD) optimisation algorithms (Chatzilygeroudis et al. 2021), Multi-dimensional Archive of Phenotypic Elites (MAP-Elites) in particular (Mouret and Clune 2015), which keep track of many different classes of solutions and check the performance of offspring from one class in other classes, which may lead to the discovery of stepping stones through many classes on path to interesting discoveries.

An initial iteration of applying the system for automated sound discovery is inspired by the Innovation Engines approach (Nguyen, Yosinski, and Clune 2015b). To automate the exploration, Innovation Engines combine QD algorithms with a model that is capable of evaluating whether new solutions are interestingly new. A long-term goal of Innovation Engines is to learn classifying data they have seen so far and seek to produce new examples, unsupervised, without labeled data. To start our Innovation Engine explorations with sound (Jónsson et al. 2024), we started with a pre-trained model, a Deep Neural Network (DNN) classifier, to define the measurement space, replicating a setup from previous evaluations of the Innovation Engine algorithm in the visual domain (Nguyen, Yosinski, and Clune 2016). With this approach, the DNN classification is used to define diversity, and the classification confidence levels for each class are used as a measure of quality.

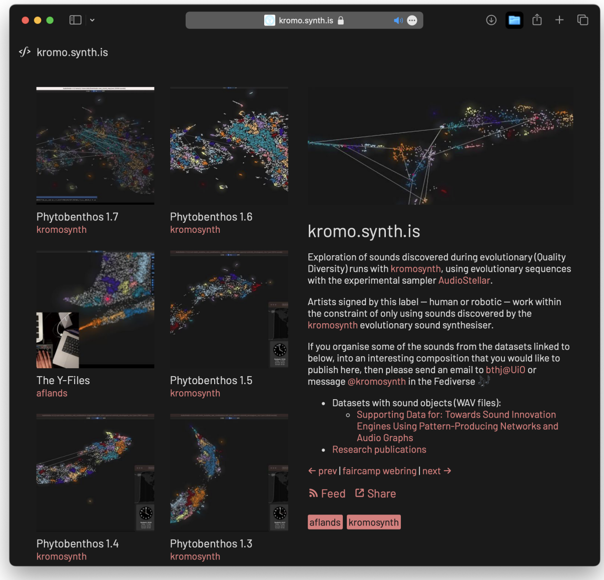

Figure 1: Independent record store at kromo.synth.is, dedicated to organised sounds from our evolutionary system.

Applying this approach and the aforementioned method for generating sounds led to the discovery of a diverse set of high-quality solutions, according to the confidence levels of the pre-trained classifier employed. The solutions are often not representative of their corresponding class, even though they have been assigned to it with high levels of confidence, which may be due to the fact that DNNs are easily fooled (Nguyen, Yosinski, and Clune 2015a). This has not been a drawback but rather a trait of our experiments, as the goal is not to imitate existing artefacts with state-of-the-art classification but rather to encourage diverse and interesting discoveries.

Empirical Tests of Artefacts While structured human evaluation of the discovered artefacts has not been conducted, their applicability has been put to the test in different contexts. We offer an online explorer which provides access to sonic renderings of the discovered artefacts from all our evolutionary runs[1]. Furthermore, we have applied the rendered sound objects in automated, evolutionary sound organisation and published several recordings of livestreams

[1] Evolution runs explorer:
`https://synth.is/exploring-evoruns`

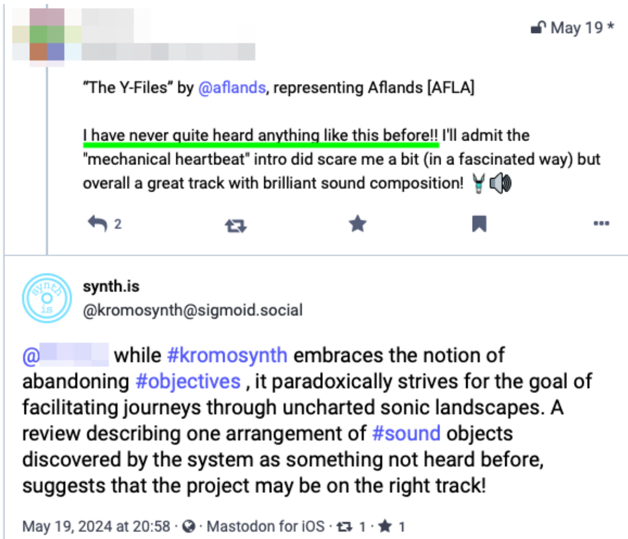

Figure 2: A comment on one manual composition of sounds from our Sound Innovation Engine, and our response.

from those compositions[2]. An independent distribution site—online record store or a digital garden (Basu 2020) — dedicated to sounds organised from those evolutionary processes, has been established at `kromo.synth.is` (Figure 1). Both stochastic and manual organisations of the discovered sound objects, arranged with the experimental sampler AudioStellar (Garber, Ciccola, and Amusategui 2021), can be obtained at the site. One of the manually organised pieces was entered into Fedivision[3], an annual song contest held in the Fediverse (Mansoux and Roscam Abbing 2020). The entry received several interesting comments, such as: "I have never quite heard anything like this before!!"[4] (Figure 2), which indicate a potential for transformation.

The diversity, quality, and applicability of the sound objects discovered with this evolutionary configuration suggest that further explorations based on this system would be worthwhile.

Autonomous Exploration of Sonic Landscapes

Currently, we are experimenting with enabling autonomous exploration through the space of sounds by augmenting this approach for evolutionary sound discovery. To that end we have separated the functionality of the previously employed pre-trained classifier, which served to define diversity and quality in one model, into finer grained components. Those components serve to extract audio features from rendered sound objects, project them into a behaviour space and evaluate their quality. We are experimenting with approaches for extracting features, such as with sets of manually selected features, perceptually aligned MFCC features and features representing high and low level structures of audio with deep models such as VGGish (Hershey et al. 2017). When defining the measurement space from the extracted features, we have drawn inspiration from previous work on autonomously discovering the behaviour space (Grillotti and Cully 2022), where unsupervised dimensionality reduction (DR) is applied and periodically retrained. Currently, we have applied linear projection with Principal Component Analysis (PCA) and discretized it into a regular grid, on which we apply the QD algorithm MAP-Elites in a similar fashion to its application on predefined classes in earlier experiments. A diagram of our experimental setup can be seen in Figure 3.

While comprehensive simulations and thorough analysis is yet to be conducted, preliminary results indicate that our approach still achieves high levels of quality and diversity. The trajectories of quality and coverage discovered with such newer configurations of sonic search simulations are similar to those we previously published in (Jónsson et al. 2024), although they are obtained by applying unsupervised approaches to discover the measurement space autonomously. A likely explanation is that the environment of each individual evolutionary run remains static throughout, though a diverse set of behaviours is explored within it. Eventually our simulations run out of opportunities for new solutions. This motivates us to turn our attention to the Open-Ended Evolution (OEE) research field.

Open-Ended Discovery of Sounds

The different constraints set by the idiosyncratic nature of transformational creativity, discussed in the Introduction, can be considered as different problems to solve or opportunities to evolve (Soros, Lehman, and Stanley 2017). To emulate these idiosyncrasies and continuously evolve towards increasingly diverse and complex sounds, we are looking towards the transfer of discovered sound objects between various environments, an approach inspired by the POET algorithm. (Wang et al. 2019) and subsequent works (Stensby, Ellefsen, and Glette 2021; Norstein, Ellefsen, and Glette 2022). The dynamic switching between feature spaces proposed by (Usui, Suzuki, and Arita 2023) aligns well with our previous experimental setups. To try a similar approach within our experimental environment, we have defined two or more containers for MAP-Elites to switch between dynamically. As an initial approach for defining each container as a unique terrain to navigate, we have chosen different sets of PCA components for each container, for example, components 0 and 2 for the first map, 1 and 3 for the second, and so on. Furthermore, each map can have different approaches for fitness evaluation, such as different sets of reference sounds for each. To further differentiate the map terrains—representing diverse and idiosyncratic constrains—we are investigating the effect of applying other non-linear projections, such as (Parametric-)Uniform Mani-

---

[2] Playlist with evolutionary sequences through sounds discovered by QD runs: `https://youtube.com/playlist?list=PLSYAaR-xYhEXk0czfHYKJSWmZ8vG35xEN&feature=shared`

[3] Fedivision: `https://fedivision.party`

[4] Social media correspondence on manually organised `kromosynth` sounds: `https://sigmoid.social/@kromosynth/112469287837850218`

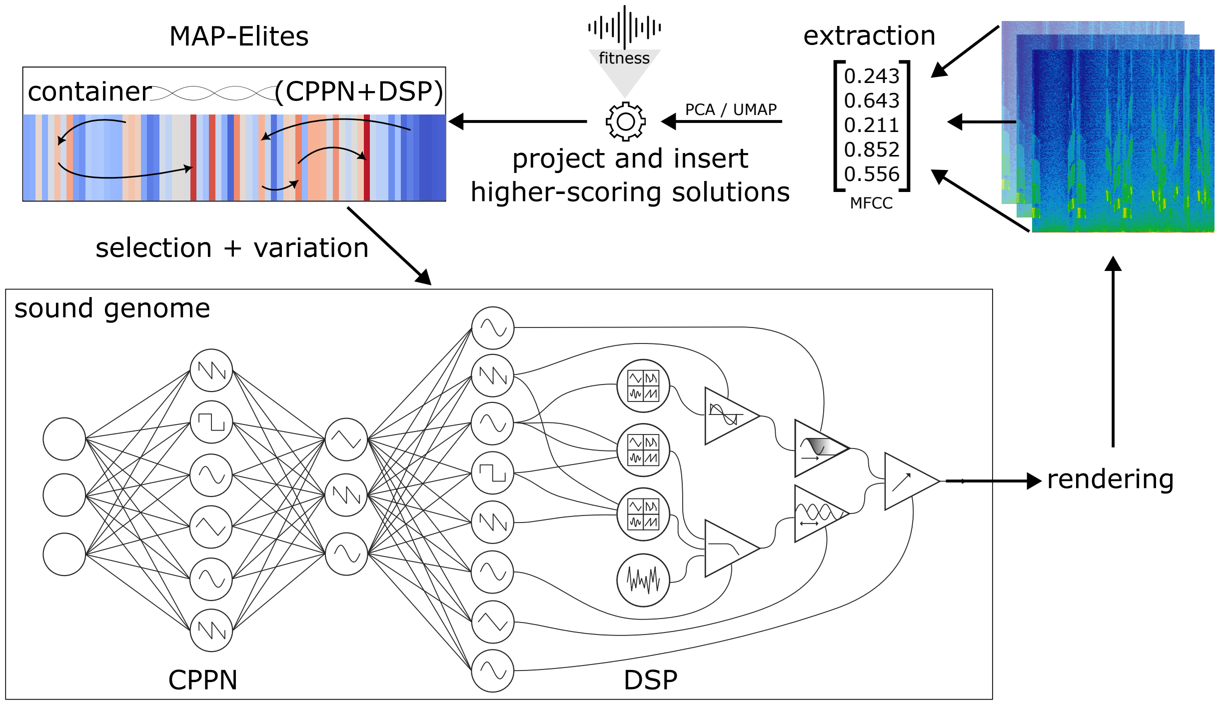

Figure 3: Diagram of the data flow in our experimental setup, showing how the genome fits within the data pipeline. This figure differs from a comparable diagram in a previous publication (Jónsson et al. 2024) in its illustration of exchangeable components for audio feature extraction, projection, and fitness evaluation.

fold Approximation and Projection (UMAP), hypothesising that different DR approaches can facilitate diverse competition dynamics between the cells in the grid, thus creating new problems to solve and opportunities to evolve. When switching environments, the whole population of elites in the current environment is mapped to the new environment, as in (Usui, Suzuki, and Arita 2023). Simulations of those and further combinations of such an OEE configuration are underway.

Emulation of dispersal through natural means, such as that of plant seeds in the forests, could be interesting to compare with the mass migration of elites described above. Could the dispersal of individual sound objects discovered in one environment to another be better suited to offer new opportunities for discovery? In that regard, different dispersal conditions (Maley 1993) could be investigated.

Fostering Openness to new Sonic Experiences

Open-earedness, as formulated by David J. Hargreaves (Hargreaves 1982; Hargreaves and Bonneville-Roussy 2018) and other researchers (Louven 2016) in the field of music psychology, revolves around the psychological openness to unfamiliar musical styles. Considering sounds broadly, and how they influence musical styles, open-earedness could be reinterpreted to encompass receptivity to novel sounds, such as those discovered by electronic synthesisers. Challenging listener's expectations and inviting them to appreciate unfamiliar or unconventional timbral textures can be conducive to elevating their openness to new sonic experiences. Computational aid from a system autonomously exploring sonic landscapes, such as the one we work towards and present in this paper, may foster the ability to appreciate novel sound discoveries. Improving such listener tolerance may in turn lead to exposure that provides inspiration for new sonic compositions.

References

Basu, T. 2020. Digital gardens let you cultivate your own little bit of the internet.

Chatzilygeroudis, K.; Cully, A.; Vassiliades, V.; and Mouret, J.-B. 2021. Quality-Diversity Optimization: A Novel Branch of Stochastic Optimization. In Pardalos, P. M.; Rasskazova, V.; and Vrahatis, M. N., eds., *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*, Springer Optimization and Its Applications. Cham: Springer International Publishing. 109–135.

Dingemanse, M. 2023. Human creativity in the age of text generators.

Garber, L.; Ciccola, T.; and Amusategui, J. 2021. AudioStellar, an open source corpus-based musical instrument for latent sound structure discovery and sonic experimentation. In *Proceedings of the International Computer Music Conference*, 62–67.

Grillotti, L., and Cully, A. 2022. Unsupervised Behavior Discovery With Quality-Diversity Optimization. *IEEE Transactions on Evolutionary Computation* 26(6):1539–1552.

Hargreaves, D. J., and Bonneville-Roussy, A. 2018. What is ‘open-earedness’, and how can it be measured? *Musicae Scientiae* 22(2):161–174.

Hargreaves, D. J. 1982. The development of aesthetic reactions to music. *Psychology of Music*. Publisher: Sage Publications.

Hershey, S.; Chaudhuri, S.; Ellis, D. P. W.; Gemmeke, J. F.; Jansen, A.; Moore, R. C.; Plakal, M.; Platt, D.; Saurous, R. A.; Seybold, B.; Slaney, M.; Weiss, R. J.; and Wilson, K. 2017. CNN Architectures for Large-Scale Audio Classification. arXiv:1609.09430 [cs, stat].

Jónsson, B. T.; Erdem, C.; Fasciani, S.; and Glette, K. 2024. Towards Sound Innovation Engines Using Pattern-Producing Networks and Audio Graphs. In Johnson, C.; Rebelo, S. M.; and Santos, I., eds., *Artificial Intelligence in Music, Sound, Art and Design*, volume 14633. Cham: Springer Nature Switzerland. 211–227. Series Title: Lecture Notes in Computer Science.

Jónsson, B. T.; Erdem, C.; and Glette, K. 2024. A System for Sonic Explorations With Evolutionary Algorithms. *Journal of the Audio Engineering Society* 72(4):257–266.

Jónsson, B. T.; Hoover, A. K.; and Risi, S. 2015. Interactively Evolving Compositional Sound Synthesis Networks. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, GECCO ’15, 321–328. New York, NY, USA: Association for Computing Machinery.

Louven, C. 2016. Hargreaves’ “open-earedness”: A critical discussion and new approach on the concept of musical tolerance and curiosity. *Musicae Scientiae* 20(2):235–247.

Maley, C. C. 1993. The Effects of Dispersal on the Evolution of Artificial Parasites. In *Proceedings of the Second European Conference on Artificial Life (ECAL ’93)*.

Mansoux, A., and Roscam Abbing, R. 2020. Seven theses on the fediverse and the becoming of FLOSS. 124–140. Publisher: Institute for Network Cultures and Transmediale.

Mathews, M. V., and Pierce, J. R. 1987. The Computer as a Musical Instrument. *Scientific American* 256(2):126–133. Publisher: Scientific American, a division of Nature America, Inc.

McCormack, J. 2005. Open Problems in Evolutionary Music and Art. In Hutchison, D.; Kanade, T.; Kittler, J.; Kleinberg, J. M.; Mattern, F.; Mitchell, J. C.; Naor, M.; Nierstrasz, O.; Pandu Rangan, C.; Steffen, B.; Sudan, M.; Terzopoulos, D.; Tygar, D.; Vardi, M. Y.; Weikum, G.; Rothlauf, F.; Branke, J.; Cagnoni, S.; Corne, D. W.; Drechsler, R.; Jin, Y.; Machado, P.; Marchiori, E.; Romero, J.; Smith, G. D.; and Squillero, G., eds., *Applications of Evolutionary Computing*, volume 3449. Berlin, Heidelberg: Springer Berlin Heidelberg. 428–436. Series Title: Lecture Notes in Computer Science.

Mouret, J.-B., and Clune, J. 2015. Illuminating search spaces by mapping elites. arXiv:1504.04909 [cs, q-bio].

Nguyen, A.; Yosinski, J.; and Clune, J. 2015a. Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images. arXiv. arXiv:1412.1897 [cs].

Nguyen, A. M.; Yosinski, J.; and Clune, J. 2015b. Innovation Engines: Automated Creativity and Improved Stochastic Optimization via Deep Learning. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, GECCO ’15, 959–966. New York, NY, USA: Association for Computing Machinery.

Nguyen, A.; Yosinski, J.; and Clune, J. 2016. Understanding innovation engines: Automated creativity and improved stochastic optimization via deep learning. *Evolutionary Computation* 24(3):545–572.

Norstein, E. S.; Ellefsen, K. O.; and Glette, K. 2022. Open-Ended Search for Environments and Adapted Agents Using MAP-Elites. In Jiménez Laredo, J. L.; Hidalgo, J. I.; and Babaagba, K. O., eds., *Applications of Evolutionary Computation*, Lecture Notes in Computer Science, 651–666. Cham: Springer International Publishing.

Secretan, J.; Beato, N.; D’Ambrosio, D. B.; Rodriguez, A.; Campbell, A.; Folsom-Kovarik, J. T.; and Stanley, K. O. 2011. Picbreeder: a case study in collaborative evolutionary exploration of design space. *Evolutionary Computation* 19(3):373–403.

Smith III, J. O. 1991. Viewpoints on the history of digital synthesis. In *Proceedings of the International Computer Music Conference*, 1–1. INTERNATIONAL COMPUTER MUSIC ACCOCIATION.

Soros, L.; Adams, A.; Kalonaris, S.; Witkowski, O.; and Guckelsberger, C. 2024. On Creativity and Open-Endedness. arXiv:2405.18016 [cs] version: 1.

Soros, L. B.; Lehman, J.; and Stanley, K. O. 2017. Open-Endedness: The Last Grand Challenge You’ve Never Heard Of.

Stanley, K. O., and Lehman, J. 2015. *Why Greatness Cannot Be Planned*. Cham: Springer International Publishing.

Stanley, K. O., and Miikkulainen, R. 2002. Evolving Neural Networks through Augmenting Topologies. *Evolutionary Computation* 10(2):99–127.

Stanley, K. O. 2007. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines* 8(2):131–162.

Stensby, E. H.; Ellefsen, K. O.; and Glette, K. 2021. Co-optimising Robot Morphology and Controller in a Simulated Open-Ended Environment. In Castillo, P. A., and Jiménez Laredo, J. L., eds., *Applications of Evolutionary Computation*, Lecture Notes in Computer Science, 34–49. Cham: Springer International Publishing.

Takagi, H. 2001. Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation. *Proceedings of the IEEE* 89(9):1275–1296. Conference Name: Proceedings of the IEEE.

Usui, K.; Suzuki, R.; and Arita, T. 2023. Towards open-ended evolution based on CVT-MAP-Elites with dynamic switching between feature spaces. MIT Press.

Veale, T.; Amílcar Cardoso, F.; and Pérez Y Pérez, R. 2019. Systematizing Creativity: A Computational View. In Veale, T., and Cardoso, F. A., eds., *Computational Creativ-*

ity. Cham: Springer International Publishing. 1–19. Series Title: Computational Synthesis and Creative Systems.

Wang, R.; Lehman, J.; Clune, J.; and Stanley, K. O. 2019. POET: open-ended coevolution of environments and their optimized solutions. In *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO '19, 142–151. New York, NY, USA: Association for Computing Machinery.