



HAL
open science

Geodesics, Non-linearities and the Archive of Novelty Search

Achkan Salehi, Alexandre Coninx, Stephane Doncieux

► **To cite this version:**

Achkan Salehi, Alexandre Coninx, Stephane Doncieux. Geodesics, Non-linearities and the Archive of Novelty Search. GECCO '22: Genetic and Evolutionary Computation Conference, Jul 2022, Boston Massachusetts, United States. pp.156-159, 10.1145/3520304.3528944 . hal-03661188

HAL Id: hal-03661188

<https://hal.science/hal-03661188v1>

Submitted on 7 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Geodesics, Non-linearities and the Archive of Novelty Search

Achkan Salehi, Alexandre Coninx, Stephane Doncieux

Sorbonne Université, CNRS, ISIR, F-75005

Paris, France

{achkan.salehi,alexandre.coninx,stephane.doncieux}@sorbonne-universite.fr

Abstract

The Novelty Search (NS) algorithm was proposed more than a decade ago. However, the mechanisms behind its empirical success are still not well formalized/understood. This short note focuses on the effects of the archive on exploration. Experimental evidence from a few application domains suggests that archive-based NS performs in general better than when Novelty is solely computed with respect to the population. An argument that is often encountered in the literature is that the archive prevents exploration from backtracking or cycling, *i.e.* from revisiting previously encountered areas in the behavior space. We argue that this is not a complete or accurate explanation as backtracking—beside often being desirable—can actually be enabled by the archive. Through low-dimensional/analytical examples, we show that a key effect of the archive is that it counterbalances the exploration biases that result, among other factors, from the use of inadequate behavior metrics and the non-linearities of the behavior mapping. Our observations seem to hint that attributing a more active role to the archive in sampling can be beneficial.

CCS Concepts

• Computing methodologies → Search methodologies; Evolutionary robotics.

Keywords

Evolutionary Robotics, Complex Systems, Theory

ACM Reference Format:

Achkan Salehi, Alexandre Coninx, Stephane Doncieux. 2022. Geodesics, Non-linearities and the Archive of Novelty Search. In *Genetic and Evolutionary Computation Conference Companion (GECCO '22 Companion)*, July 9–13, 2022, Boston, MA, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3520304.3528944>

1 Introduction

In recent years, Novelty Search (NS) [3, 8] has increasingly been studied and applied in contexts ranging from robotic manipulation [7] and Reinforcement Learning [6] to swarm robotics [5] and games [9].

Traditionally, two different variations of NS have been used: the archive-free one, also called behavioral diversity [11], in which the Novelty score is computed with respect to the population, and the archive-based one, in which the aforementioned score is computed with respect to the archive and the population. Some (limited) evidence [2, 4, 10] suggests that the use of an archive is in general beneficial for exploration itself. However, very few explanations have been proposed for these observations, and they often do not go beyond speculation. An argument that is often encountered in the literature [8, 10, 13, 15] is that the archive prevents backtracking

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '22 Companion, July 9–13, 2022, Boston, MA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9268-6/22/07.

<https://doi.org/10.1145/3520304.3528944>

(also called "cycling") in the behavior space: that is, its role is to prevent the exploration from visiting areas that have been previously visited. This is inaccurate or at least incomplete, as while, in fact, avoiding recently visited areas can promote dynamic exploration in the short term, in the longer term, backtracking is often *desirable*: even in the ideal and simplified case where the mapping from genotype space to phenotypes is one-to-one, accessing some areas of behavior space might depend on rare events, whose realization might require to return frequently enough to a particular area of the behavior space. In practice, as illustrated by the low-dimensional analytical problems that we discuss, archives that are large enough to succeed on a given task in general encourage backtracking over longer time horizons.




Figure 1: (Left) An Archimedean spiral with parameter $a = 0.01$ and $t \in [0, 30\pi]$ embedded in \mathbb{R}^2 . Using the Euclidean distance and $k = 1$, points A and B will have higher Novelty scores than C and D. In contrast, using the geodesic distance, C and D would have much higher Novelty than A and B. In this figure, the red dot indicates the starting point of the agents. (Right) In practice, even when the reachable behavior space is a smooth manifold, Isotropic Gaussian mutations of a genotype x result in behavior exploration that is locally biased. In this figure, isotropic Gaussian mutations within one standard deviation are mapped to the interior of the bean curve $r = 0.7 \cos^3(t)^3 + \sin^3(t)$ which is then mapped to the surface $2x^2 - 4y^2$. In more practical cases, *e.g.* with neural network controllers, the resulting distribution will in general be much more complex.

In this paper, we study the advantages of using an archive for exploration by first inspecting how exploration can be biased by non-linearity and the use of the often inadequate Euclidean metric, which does not account for curvature or discontinuities. We build on low-dimensional, analytical examples, based on which we hypothesize that the usefulness of the archive for exploration stems, in part, from its ability to compensate those biases. While that bias mitigation can naturally occur in bounded behavior spaces with unbounded archives, we observe that this might not necessarily be the case for practical archives with manageable sizes without some improvements to the manner in which archives are defined and used in NS. In particular, those observations, coupled with the fact that individuals that are far in genotype space can in many situations be mapped to (nearly) identical behaviors, indicate that the archive should play a more active role in the sampling process and not be decoupled from the population.

Notations and reference algorithms. Given a genotype space \mathcal{G} , we assume that a hand-engineered or learned behavior space \mathfrak{B} has been defined,




Figure 2: The cumulative results of 20 Novelty search experiments in the four problem settings. Note that in (a) and (c), a Gaussian distribution centered at $g_b \in \mathcal{G}_b$ is mapped to a distribution with mean $\phi_u(g_b)$ but that is skewed towards the exterior of the spiral. In (a) and (b), the use of the Euclidean metric results in lower Novelty scores towards the interior of the spiral. In these 20 experiments, Novelty search was performed without an archive, and full behavior space coverage proved to be possible only in the case where the geodesic distance and non-biased mutations were used (d).

and denote $\phi : \mathcal{G} \rightarrow \mathfrak{B}$ the mapping between the two spaces. Considering an evolutionary optimization process running for G_{max} generations, the current population at generation g will be noted \mathcal{P}^g , and we will write \mathcal{A} the archive that is used with NS. Both unstructured and structured archives have been used in our experiments. In the unstructured case, archive growth was based on the selection of random elements from the population at each generation, and removals were also random. The structured archives that we used were simple grid-like partitionings as in the original MAP-elites algorithm [10], and as we focus on pure exploratory search (*i.e.* no fitness/reward), addition of individuals to each cell was based on its occupancy information: individuals falling in empty cells were immediately added. Otherwise, the individual already occupying the cell was replaced with a small probability ϵ . The bound imposed on the size of bounded archives will be noted \mathcal{A}_{max} .

2 On biased exploration

In this section, we place ourselves in the idealized case where the phenotype space is a Riemannian manifold embedded in \mathbb{R}^n . We define *locally biased* exploration at a point $b \in \mathfrak{B}$ as an exploration process that favors particular directions in behavior space, *i.e.* a sampling process that is not isotropic. We note that undesired biases might arise from these sources:

- The Novelty objective is often computed based on the Euclidean metric¹, which could not only be different but entirely contradictory with Novelty based on geodesic distances, as illustrated in figure 1 (left).
- The non-linearity of the genotype to phenotype mapping ϕ results in non-isotropic exploration of the behavior space (figure 1 (right)). Indeed, even if we assume isotropic gaussian mutations with sufficiently small variance σ (a condition that will not hold in practice), the resulting random variable $\phi(x)$ will have covariance $\sigma \frac{\partial \phi}{\partial x} \frac{\partial \phi}{\partial x}^T$, which is not isotropic.

We hypothesize that balanced exploration of the behavior space requires the mitigation of those biases, and illustrate that point by studying the following toy setting: let us assume a reachable phenotype space \mathfrak{B}_{toy} given by an archimedean spiral (figure 1, left) embedded in \mathbb{R}^2 whose equation in cartesian coordinates is given by (for $t \in \mathbb{R}^+$)

$$\gamma(t) = (at \cos(t), at \sin(t)) \quad (1)$$

and for which the geodesic distance naturally coincides with its arc-length, given by

$$S(t_1, t_2) = \frac{a}{2} (t\sqrt{t^2+1} + \log(t + \sqrt{t^2+1}))|_{t_1}^{t_2}. \quad (2)$$

Let us also define the two distinct bounded genotype spaces

$$\begin{cases} \mathcal{G}_b \triangleq [0, \alpha\pi] \\ \mathcal{G}_u \triangleq \{S(0, l) | l \in [0, \alpha\pi]\} \end{cases} \quad (3)$$

for some positive real α . Indeed, \mathcal{G}_b corresponds to the angle parametrization of the curve, and \mathcal{G}_u is the set of possible arc-lengths. Let us use an isotropic Gaussian mutation operator, and finally, let us define two mappings from those genotypes to \mathbb{R}^2 :

$$\begin{cases} \phi_b(g_b) \triangleq \gamma(g_b) & \forall g_b \in \mathcal{G}_b \\ \phi_u(g_u) \triangleq \gamma(S^{-1}(g_u)) & \forall g_u \in \mathcal{G}_u \end{cases} \quad (4)$$

The motivation behind this choices for ϕ_b and ϕ_u , which both have \mathfrak{B}_{toy} as their image, is that they behave differently under Gaussian mutations: it can easily be verified that an isotropic Gaussian distribution in \mathcal{G}_b will be mapped by ϕ_b to a distribution that is skewed towards the exterior of the spiral, while an isotropic Gaussian distribution in \mathcal{G}_u will be mapped by ϕ_u to an isotropic Gaussian on γ . Note that the mapping S^{-1} can not be expressed in closed form. As a result, whenever necessary, we approximate its values by solving the corresponding ordinary differential equation.

These settings define four problems, that result from pairing one of the two parametrizations \mathcal{G}_b or \mathcal{G}_u with one of the two metrics (Euclidean or Geodesic $S(t_1, t_2)$). We expect that using the geodesics distance in conjunction with \mathcal{G}_u should result in balanced exploration, and that all other cases NS will heavily favor the exploration to move towards the exterior of the spiral. We verified this hypothesis by running 20 experiments corresponding to each of those four problems, setting $\alpha = 30$, with population and offspring sizes of $M = N = 30$, mutations sampled from $\mathcal{N}(0, 0.3)$ and the consideration of $k = 10$ nearest neighbors. In all of the experiments, the number of generations was set to $G_{max} = 1000$, and all individuals were initialized with a particular value of s (the red dot in figure 1 (left)). The cumulated results, that are reported in figure 2 confirm our hypothesis: the only outcome in which the cumulative exploration results fully cover the behavior space (figure 2 (d)) corresponds to the case where the geodesic distance is used with \mathcal{G}_u . In figures 2(a) and (c), exploration seems to be trapped around the external borders of the behavior space. Figure 2 (b)

¹In very few instances of applying the Novelty objective, such as in the context of robotic manipulation [14], distances are measured in the tangent plane of $SE(N)$. However, the reachable behavior space is limited by the scene and the limits of the actuators, resulting in discontinuities and irregularities that no longer correspond to those metrics.

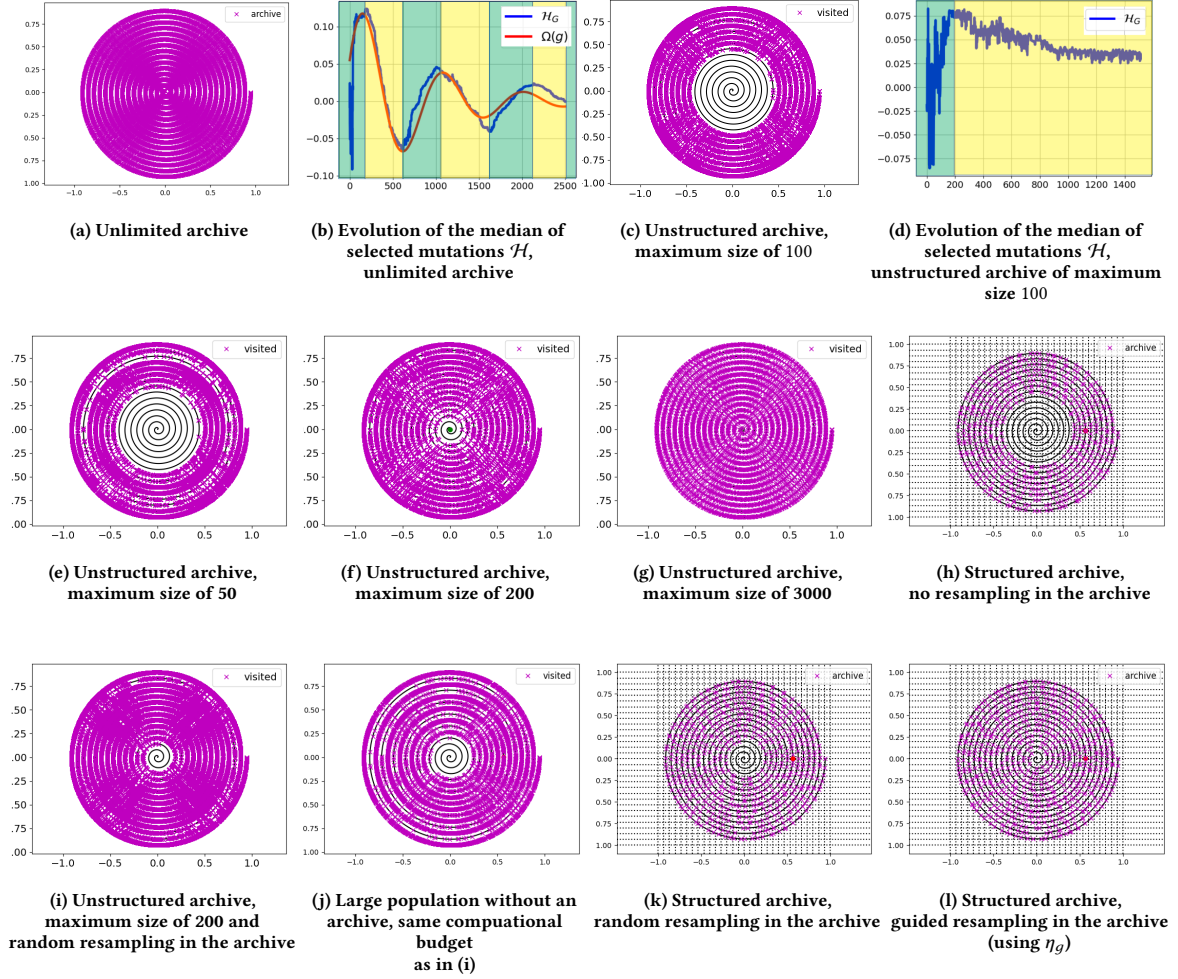


Figure 3: Cumulated results of experiments (with $G_{max} = 1000$) obtained with different archive configurations when the Euclidean metric and the ϕ_b parametrization were used.

seems to indicate that using a parametrization \mathcal{G}_u that does not induce bias is not sufficient to compensate the selection pressure that results from the use of the wrong metric.

This finally brings us to the archive. As we will discuss in the next section, the use of a large unstructured archive can compensate the aforementioned biases in the long term, precisely *because it encourages back-tracking*, without the need for learning the correct (*i.e.* geodesic) metric or designing mutation operators that do not induce bias (both of which are impractical in real-world problems).

3 Bias compensation using an archive

The experiments presented throughout the section use the bias-inducing parametrization \mathcal{G}_b , and the Euclidean metric. Unless explicitly stated, the same settings for α, M, N, k, G_{max} as in the previous section are used. The same holds for mutation standard deviations.

3.1 Large unstructured archives can overcome exploratory biases by enabling backtracking

With bounded behavior spaces, it is intuitively expected that an unbounded archive will eventually be able to overcome any biases in the selection pressures, and that is indeed the case for the spiral (3 (a)). To explicit the effects of the archive on selection pressures, let us consider the history of of all selected mutations, over G generations. More precisely, denote by $\psi(\cdot)$ the operator that maps an individual in the population \mathcal{P}^g to its parent in \mathcal{P}^{g-1} , and define

$$h_g \triangleq \{S(0, p) - S(0, \psi(p))\}_{p \in \mathcal{P}^g}. \quad (5)$$

In other terms, each element of h_g is the behavioral change, expressed as a change in arc length, that results from a mutation to some individual p at generation g . Now, consider the history of the medians of all such behavioral mutations over G generations $\mathcal{H}_G \triangleq \{\text{median}(h_i)\}_{i=0}^{G-1}$, which is displayed in figure 3 (b) (blue curve). We see the emergence of a pattern reminiscent of a *damped harmonic oscillatory motion*, as illustrated by the fitted curve that is displayed in red in figure 3(b), and which converges to 0 in term of

amplitude as g increases. Interestingly, each of its half-periods correspond to either an *expansion* or *retraction* stage in exploration: each expansion stage (shaded in green), corresponds to a set of consecutive generations during which the population consistently moves towards the exterior of the spiral. Likewise, each retraction stage, shaded in yellow, corresponds to a set of consecutive generations during which the population is pushed towards the center of the spiral. It can be seen that the amplitude (*i.e.* the median of the selected mutations) approaches 0 as g grows, indicating that asymptotically, the use of an unbounded archive does indeed lead to unbiased, balanced exploration.

This pattern seems to hint that in this settings —and using the population-based sampling of traditional Novelty search —, a bounded archive will fail to cover the entire space unless it is sufficiently large to push the exploration first through an expansion phase before *backtracking* towards the center. This is confirmed by the experiments with archive sizes of $\mathcal{A}_{max} = 50, 100, 200$ (figure 3(c,d,e,f)). In particular, it is clear from 3(d) that those archive sizes fail because they are not large enough to encourage a complete backtracking phase. In our experiments, it appeared that in order to ensure success over a single run consistently it was necessary that $\mathcal{A}_{max} > 1000$. An example for $\mathcal{A}_{max} = 3000$ is given in figure 3(g).

Structured archives fail to cover the behavior space in G_{max} generations without significant modification to the sampling process (figures 3(h, k)). This is expected, since in that case, sampling only from the population eliminates the possibility of backtracking (figure 3(h)) and allowing sampling from the archive as in MAP-elites (figure 3(k)) converges too slowly: backtracking requires the selection of a single particular cell, which happens too infrequently. We recognize that this might be a particularity of the low-dimensional problem at hand. Indeed, while only a single path between two points $b_0, b_1 \in \mathfrak{B}_{to y}$ exists on the spiral, in higher dimensions, unless we consider degenerate cases, an infinity of such paths would be available. However, we note the notorious difficulty of choosing the adequate partitioning and resolution of a structured archive which is also present in higher dimensions.

3.2 Discussion

We saw that non-linearities and discrepancies between the distance metric and the local structure of the behavior space were a source of bias in exploration, and showed that it could be compensated by using an unstructured archive that was large enough to enable backtracking past the initialization point. As the time complexity of NS with an unstructured archive is $O(|\mathcal{A}| \log |\mathcal{A}|)$ [10], in practice, bounds should be imposed on its size. Consequently, a logical step would be to modify the sampling strategy —which by default samples the population —to facilitate backtracking. However, as illustrated by figure 3(i,k), simply allowing resampling from the archive (as in classical MAP-elites) leads to unreasonably slow convergence as progress relies on the sampling of rare events. As the reader has probably noticed, the situation in which we allow resampling into an unstructured archive is to some extent similar, conceptually, to Novelty search without an archive but with a larger population (keeping the same computational budget). As is illustrated in figure 3(j), the latter performs worse than the former (results aggregated over 5 runs). This is not surprising from an intuitive point of view, as archive-less NS does not impose the additional local selection pressures that result from archive densifications.

We hypothesize that a more reliable sampling approach is to augment NS with a selection pressure that is *absolute*, in the sense that it is relative to the phenotype space \mathfrak{B} rather than relative to an ever-moving unstructured archive. For example, considering NS with a structured archive, let us assign a score η_g , that we dub *temporary discovery score*, to each individual p from the parent population \mathcal{P}^g :

$$\eta_g(p) = \tau \eta_{g-1}(p) + (1 - \tau) \frac{\sum_{q \in \xi_p} \kappa(q)}{\sum_{p \in \mathcal{P}^g} \sum_{q \in \xi_p} \kappa(q)}, \quad (6)$$

where ξ_p denotes the set of offsprings from p at generation g , and $\kappa(q)$ is simply set to whether the cell in which $\phi(q)$ falls is empty or not. The parameter $\tau \in [0, 1]$ sets the rate at which this score is updated with new information. As is shown in figure 3(l), combining the population-based sampling of NS with resampling in the archive guided by the η_g score consistently results in highly improved coverage of the behavior space in all executions.

Note that the suggested score can be thought of as a more flexible version of the Curiosity score proposed in [2] in the context of Quality-Diversity. Indeed, the latter is discrete and penalizes the lack of new discoveries by the same amount by which it rewards new behaviors. In contrast, the τ parameter of the η_g score allows for balancing between positive rewards and penalties.

We close this section with the remark that a possible way in which κ could be defined in the context of NS with unstructured archives is to use a loss based on Random Network Distillation [1]. However, that is conditional to whether or not one is able to mitigate the catastrophic inference issues [12] that are well-known in the context of continuous learning with encoders.

ACKNOWLEDGMENT

This work was supported by the European Union’s H2020-EU.1.2.2 Research and Innovation Program through FET Project VeriDream under Grant Agreement Number 951992.

References

- [1] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. 2018. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894* (2018).
- [2] Antoine Cully and Yiannis Demiris. 2017. Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation* 22, 2 (2017), 245–259.
- [3] Stéphane Doncieux, Alban Laflaquière, and Alexandre Coninx. 2019. Novelty search: a theoretical perspective. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 99–106.
- [4] Jorge Gomes, Pedro Mariano, and Anders Lyhne Christensen. 2015. Devising effective novelty search algorithms: A comprehensive empirical study. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*. 943–950.
- [5] Jorge Gomes, Paulo Urbano, and Anders Lyhne Christensen. 2013. Evolution of swarm robotics systems with novelty search. *Swarm Intelligence* 7, 2 (2013), 115–144.
- [6] Ethan C Jackson and Mark Daley. 2019. Novelty search for deep reinforcement learning policy network weights by action sequence edit metric distance. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 173–174.
- [7] Seungsu Kim, Alexandre Coninx, and Stéphane Doncieux. 2021. From exploration to control: learning object manipulation skills through novelty search and local adaptation. *Robotics and Autonomous Systems* 136 (2021), 103710.
- [8] Joel Lehman and Kenneth O Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19, 2 (2011), 189–223.
- [9] Antonios Liapis, Georgios N Yannakakis, and Julian Togelius. 2015. Constrained novelty search: A study on game content generation. *Evolutionary computation* 23, 1 (2015), 101–129.
- [10] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
- [11] J-B Mouret and Stéphane Doncieux. 2012. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary computation* 20, 1 (2012), 91–133.
- [12] Benedikt Pfüfl, Alexander Gepperth, S Abdullah, and A Kilian. 2018. Catastrophic forgetting: still a problem for DNNs. In *International conference on artificial neural networks*. Springer, 487–497.
- [13] Achkan Salehi, Alexandre Coninx, and Stéphane Doncieux. 2021. BR-NS: an archive-less approach to novelty search. *arXiv preprint arXiv:2104.03936* (2021).
- [14] Achkan Salehi, Alexandre Coninx, and Stéphane Doncieux. 2021. Few-shot Quality-Diversity Optimisation. *arXiv preprint arXiv:2109.06826* (2021).
- [15] Roby Velez and Jeff Clune. 2014. Novelty search creates robots with general skills for exploration. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*. 737–744.