Transformational Creativity Through the Lens of Quality-Diversity

Jonathan Demke[1], **Kazjon Grace**[2], **Francisco Ibarrola**[2] **and Dan Ventura**[1]
[1]Computer Science Department, Brigham Young University, Provo, UT 84602 USA
jpd0057@byu.edu, ventura@cs.byu.edu
[2]School of Architecture, Design and Planning, The University of Sydney, Sydney, Australia
kazjon.grace@sydney.edu.au, francisco.ibarrola@sydney.edu.au

Abstract

Quality-Diversity algorithms are a useful tool for creative search because they can evolve artefacts with different search strategies that focus on different criteria such as behavioral novelty or improving quality. These different strategies find diverse solutions with high quality. However, Quality-Diversity algorithms can exhibit inefficiency and a lack of control when exploring a genome space for diverse solutions. We propose two different approaches that connect the genome space of artefacts to their phenotypic behavior. The approaches may allow a computationally creative system not only to explore the space of solutions more quickly but also to guide its search towards underrepresented and surprising behaviors.

Introduction

Artefact-producing computational creative (CC) systems focus on generating artefacts that exhibit qualities like typicality, novelty, and quality. Developing metrics to evaluate these criteria is a challenging task, and so is finding the right balance for this multi-objective optimization problem: an insufficient focus on novelty would lead to artefacts that are merely good; an insufficient focus on typicality would lead to artefacts that are merely weird; an insufficient focus on quality would lead to artefacts that are merely artefacts. Accordingly, the subject of creative search has long been a focus of CC research (Boden 1992; Wiggins 2006). A key concept in this debate has been the contrast between exploratory and transformational creativity, with the former involving the search of a pre-defined space, and the latter being search that combines both redefinition and search of a space. Transformational creativity, viewed as the more significant of the two, is motivated by studies in the cognitive science of creativity showing iterative problem re-framing and the gradual emergence of both a problem and its solution (Maher and Poon 1996; Schon and Wiggins 1992). The basic nature of search algorithms implies a fixed space, and creative search has been operationalized as a kind of "meta-search"—a higher-level search over the space of search spaces permitting some level of transformational creativity in the original space.

Quality-Diversity (QD) algorithms are evolutionary algorithms (EA) that simultaneously optimize one fitness function (for "quality") while producing a set of solutions with good coverage of one or more other functions (or "behaviors", for "diversity"). QD algorithms only allow a solution to compete against its neighboring solutions; for example, it may not make sense to compare Edgar Allen Poe's horror poems to Shel Silverstein's comedic poems, but we may aptly compare Poe's poems with Emily Dickinson's horror poems. They enforce diversity by developing behavioral niches to stratify solutions while also finding quality solutions within those niches, leading to a collection of diverse solutions, each with high quality relative to their neighbors. The QD approach is not strictly a kind of meta-search, having a static genome space and being better understood as a kind of multi-objective search for a set of solutions. However, we argue that it exhibits some of the same properties, in that the diverse high-quality artefacts it tends to produce were evolved using different search strategies (i.e. combinations of the behavior and fitness functions), and tend to achieve their quality in very different ways (as the behavior space is likely nonlinear with respect to the genome space). QD algorithms are particularly useful in a co-creative CC context, where a set of diverse, high-quality options can be presented to a user, resulting in increased novelty in the final product and offering the potential for transformational creativity from the perspective of the user, who may have preconceptions about the space being searched.

Unfortunately, QD algorithms often struggle to search complex spaces thoroughly in reasonable amounts of time, as they make the assumption that the fitness and behavior evaluations can be run millions of times (Chatzilygeroudis et al. 2021). One approach to addressing this is introducing a surrogate evaluation—some rapid approximator of the more-expensive fitness and/or behavior functions (Ong, Nair, and Keane 2003). While surrogates have been applied to some CC domains (Zhang et al. 2022), for many domains the evaluation functions—as well as the expression from genotype to phenotype—are both expensive and challenging to approximate. As the dimensionality of the genome and the number of desired behaviors increases, many QD algorithms expend considerable resources exploring and re-exploring well-travelled and low-potential regions.

In this paper we propose two different approaches to improve QD search, particularly motivated by the CC context. Both approaches are based on attempting to predict how the

behavior space will change based on changes in the genome space. Our first approach is a low-level and local one: learning local gradients of individual behaviors and adapting individuals along them. Our second approach is a high-level and global one: trying to discover latent behavioral structure in the genome space.

Background

MAP-Elites (Mouret and Clune 2015) is one of the most influential QD algorithms to date. MAP-Elites maps an individual to a point in the behavior space $\mathbb{B}$, which specifies different behavioral qualities of the individual's phenotype. For example, we could map the space of stories to k behaviors, such as the amount of humor, horror, and romance contained within the story. MAP-Elites then searches the genome space to find quality solutions within behavioral niches such as the best story with a lot of humor and romance, but little horror.

The approach requires definition of a genome space $\mathbb{G}$; phenome space $\mathbb{P}$; genome to phenome map $T : \mathbb{G} \to \mathbb{P}$; behavior function $B : \mathbb{G} \to \mathbb{R}^k$, where k is the number of behavioral attributes; fitness function $P : \mathbb{G} \to \mathbb{R}$; and an archive of the elite solutions $(\mathbf{P}, \mathbf{X})$, where $\mathbf{P}$ stores the fittest score for each behavioral niche and $\mathbf{X}$ stores the fittest individual for each behavioral niche.[1] For example, let $\mathbb{P}$ be the domain of stories, $\mathbb{G}$ some genomic representation (such as a plot graph), T a story generator, and $k = 3$ behaviors: humor, horror, and romance. If B maps each behavior to a real value between 0 (e.g. not scary) and 1 (e.g. the scariest), then our behavior space $\mathbb{B} = [0, 1]^3$. $\mathbb{B}$ is then discretized along all k dimensions to form the phenome's behavioral niches. If horror and romance are discretized into bins, $[0, 0.5), [0.5, 1]$, meaning a story is considered either not scary (not romantic) or scary (romantic), and humor is discretized into bins, $[0, 0.33), [0.33, 0.66), [0.66, 1]$, so that a story can be mapped to low, medium, and high levels of humor, the result is twelve behavioral niches.

MAP-Elites randomly samples an individual $g \sim \mathbb{G}$, retrieves its behavioral niche $b \leftarrow B(g)$ and performance $p \leftarrow P(g)$, and checks if g is the fittest within its behavioral niche $p > \mathbf{P}[b]$. If it is, then the archive of elites is updated: $\mathbf{P}[b] \leftarrow p$ and $\mathbf{X}[b] \leftarrow g$. After enough random samples, MAP-Elites starts searching the genome space by performing genetic operations, e.g. crossover and mutation, among the elite solutions in $\mathbf{X}$.

Approximating Behavior Gradients

When a genome is mutated, the resulting child g usually exhibits a change in behavior Δb. Unfortunately, it is difficult to determine what specific change Δg led to Δb. Furthermore, it is unclear whether additional mutation in the direction of Δg leads to additional behavior change in the direction of Δb. This is a credit assignment problem. Our local, "low-level" approach is an attempt to alleviate this problem by approximating the gradient of our behavior function, ∇B, so that we can correctly assign credit to each gene.

[1] B and P measure *phenotypic* behavior and fitness, respectively, so each includes an implicit use of the mapping T.

By utilizing a differentiable regression model $f_\theta : \mathbb{G} \to \mathbb{B}$ as a surrogate for B, we can approximate ∇B by instead computing ∇f_θ. To compute f_θ, a genome $g \in \mathbb{G}$ can be mutated to generate neighbors g_i; their corresponding behaviors b_i retrieved; and f_θ trained to minimize error, i.e. $\min_\theta \|f_\theta(g_i) - b_i\|$. We can then either apply $\nabla f_\theta(g)$ directly: $g \leftarrow g \pm \nabla f_\theta(g)$ or increase the mutation rate in the direction of $\nabla f_\theta(g)$. Biasing mutation in the direction of maximal expected behavioral change could reduce the time an algorithm like MAP-Elites (which applies Gaussian noise as its mutation operator) spends searching regions of the genome space that have little chance of producing improvements in either quality or diversity.

The utility of this approach as an efficiency improvement would depend on the amount of data required to train a local behavioral regressor and the size of the region in genome space that said regressor could reasonably approximate gradients over. If a single global f_θ is accurate enough to resemble B, then utilizing ∇f_θ over the entire genome space would significantly speed up search. This might be possible if an underlying structure between $\mathbb{G}$ and $\mathbb{B}$ exists; that structure could be discovered with a neural network f_θ. However, it is unlikely that a single global f_θ will suffice, and therefore it may be necessary to employ multiple local regression models to approximate B in piecewise fashion.

A naïve first approach would define some radius around a genome and build a regressor on the mutations taken around the genome. If the search ever moves beyond the radius of the genome then we create a new regressor. Simple linear regressors could be trained with few data examples to give a quick approximation of a genome's local behavior gradient.

It might also be useful to utilize both global and local regressive models. Ensemble disagreement (Lakshminarayanan, Pritzel, and Blundell 2017) or randomized prior functions (Osband, Aslanides, and Cassirer 2018) can be built with regressive neural networks to simulate Bayesian uncertainty, which can allow EAs to exploit the global regressive model's gradient approximation when the global model's uncertainty is low and utilize a local regressive model when its uncertainty is high.

Our approach is comparable to natural evolution strategies (Wierstra et al. 2008) and covariance matrix adaptation evolution strategies (Hansen 2016), which also adapt mutation towards an approximate gradient; however, by uncoupling our gradient approximators from the current search we can use them for purposes other than finding the next population artefacts, such as backtracking or exploring the genome space on a different behavioral axis. Our approach also includes the possibility of using a global gradient approximator. Local gradient approximation is also analogous to "local explanation", an approach used as a form of explainable ML such as LIME (Ribeiro, Singh, and Guestrin 2016), which model the local environment around a datapoint using a simple, scrutable model, allowing the reasons for its classification to be made clear.

Even with gradient approximation, it can still be difficult to navigate the genome space to find some expected behavior—sometimes moving towards one behavior axis can move you away from another behavior axis.

Learning a Genome-Behaviour Latent Space

There may exist behaviorally-induced global latent structure within the genome space that may be discoverable during search. For example, a variational autoencoder (VAE) could be utilized to construct a latent space that is easier to explore (than genome space), because it clusters the genomes behaviorally, allowing sampling from the VAE's simple prior distribution to get genomes within each cluster—sampling and decoding from the latent space facilitates "intelligently" jumping around the genome space. To ensure the VAE's latent structure captures the desired behavior, the prior, encoder and decoder may be conditioned on that behavior (Sohn, Yan, and Lee 2015).

There are two challenges we see in this approach. First, VAEs commonly use continuous latent distributions to represent the data, most notably the multivariate Gaussian, which typically has a smoothness artefact that biases the mapping from similar latent values toward similar decoded outputs; however, dissimilar genomes may share behavioral features. They can also suffer from posterior collapse, where an overparameterized decoder will largely ignore most of the latent structure. In such cases, discretized latents, such as vector-quantized or categorical latents, may prove useful for alleviating these issues (van den Oord, Vinyals, and Kavukcuoglu 2017; Hafner et al. 2021).

The second and potentially more challenging issue is retrieving the necessary data to train the VAE; VAEs learn by maximizing the evidence lower bound, but what serves as the evidence for the genome space? If all genomes are considered equally likely, then the expected fitness of a VAE sample should approximate the expected fitness of the entire genome space, which may be extremely low in large genome spaces. The VAE likely wouldn't be a useful tool in this scenario. A possible solution to this problem might be to weight genomes based on their fitness, similar to how EAs perform parent selection; however, careful attention is required to ensure that the few high-performing genomes do not dominate as the evidence, since the VAE will likely overfit on the few samples and not generalize to other high-performers in the genome space. Similarly, we could weight genomes in underrepresented behavioral niches more heavily to enhance the coverage of behaviors.

Assuming a well-trained VAE model, with likelihood (encoder) $p_\theta(z \mid g, b)$, prior $p(z \mid b)$, and posterior (decoder) $q_\phi(g \mid z, b)$ distributions, one can sample from the model to find elites in the neighborhood of other elites: $g' \sim q_\phi(g \mid z, b)$ where $z \sim p_\theta(z \mid g, b)$ is a sample near the encoding of an elite $g \in \mathbf{X}$ with corresponding behavior niche b. We can also sample for a new elite in a behavorial niche b where no elite exists yet, by sampling our prior $z \sim p(z \mid b)$ and retrieving a genome from our posterior $g \sim q_\phi(g \mid z, b)$. It is important to note that although $g \sim q_\phi(g \mid z, b)$ is conditioned on behavior b, it does not guarantee that $B(g) = b$. However, this allows us to understand where behavior is not well understood within our models; we can use VAE samples to analyze whether their true behavior matches the given conditional behavior as a way to measure the information gap of behaviors in the genome space, e.g. $\mathbb{E}_{z \sim p(z|b)} \left[\| b - B(q_\phi(g \mid z, b)) \| \right]$.

Genome-Behavior Models and QD Curiosity

In complex genome spaces, the number of datapoints required to train the models in either approach may be large enough to eliminate any efficiency gains. Yet if either of these approaches is effective at connecting the genome and behavior spaces in QD applications—and we stress that *if*, because at present we haven't tested either approach beyond toy problems—then there may be more benefit to CC than any gains in efficiency. Maximizing coverage of one or more behavior functions is interesting, in that it offers a stepping stone to more CC-relevant concepts like novelty, but outside of the co-creative "offering diverse suggestions to a human" use case it is actually somewhat conceptually unsatisfying as a step towards creative search.

Creative search—and the transformational creativity it seeks to enable—are fundamentally motivated by the discovery of specific radically new solutions. The constant outward pressure of QD algorithms, however, values the entire behavior space equally at all times. Radically new solutions may emerge but are treated no differently than incrementally more-fit or more-diverse ones. This evokes 1960s "ideational fluency" notions of creativity (Torrance 1966), in the sense that QD algorithms produce the largest possible set of meaningfully different solutions to a problem. Being only a half-century behind the psychologists is still not bad for a CC algorithm, but it's possible we can do better by rethinking the problem definition for a CC context.

In classic QD algorithms the selection of where to search next is random—either by selection of a random existing elite in MAP-Elites and its differentiable derivatives like OMG-MEGA (Fontaine and Nikolaidis 2021), or by sampling from a distribution learned over the genome space in evolutionary strategy derived QD approaches like CMA-ME (Fontaine et al. 2020). This randomness seems unavoidable: QD algorithms are driven to explore the behavior space but cannot act directly within it—they must instead search the genome space and hope that doing so illuminates new behavior. The approaches we propose in this paper, however, offer an opportunity for a curiosity-motivated QD algorithm, grounded in a connection between genome and behavior. A "Curious Quality Diversity" algorithm might choose where to search next (within the genome space) based not on direct predictions of behavior, but on a drive to improve the quality of those predictions. This curiosity drive could be used to dynamically nudge QD search towards regions where behavior is not well understood—which may indicate potential for radically new (and potentially high-quality) artefacts.

The term "curiosity" has been used in QD algorithm selection before, with the "curiosity score" assigned to each elite in (Cully and Demiris 2017) being the expected probability that selecting that elite and mutating it would lead to offspring that are themselves elites (i.e. are either sufficiently different to all known elites or better than all elites they are similar to). In a sense, this is a model of what Berlyne would call "general curiosity", the drive towards any new stimulus (Berlyne 1960), which is consistent with the overall aim of QD algorithms. By contrast, "Curious QD" gives preference to new individuals that would improve the system's model of the behavior space, consistent

with Berlyne's "specific curiosity" and other similar "learning progress" notions (Oudeyer 2004; Schmidhuber 2010; Grace and Maher 2015).

While we admit to not having yet implemented any of these ideas, "Curious QD" could be implemented using either of our above approaches for connecting behavior and genome space combined with techniques from the field of Bayesian optimization (BO). BO techniques are active learning approaches that (when applied to learning a Bayesian ML model like a Gaussian Process) define an information-theoretic acquisition function over where to look next. Typical acquisition functions include upper confidence bounds (i.e. picking the spot that could theoretically be best, given uncertainty) and expected information gain (i.e. picking the spot that will reduce uncertainty the most). Applied to either of our proposed approaches, which would by necessity be learned in an active learning context, these BO techniques could produce the kind of medium-term search dynamics more recognizable as specific curiosity.

Conclusion

We have proposed two approaches that work in conjunction with the QD algorithm MAP-Elites. Our approaches focus on connecting a genome space to its phenotypic behavior space, either by approximating the local gradients of the behavior functions, or by finding a latent structure that correlates the genome space with the behavior space. These approaches not only may promote efficiency of creative search, but also, by modelling the behavior space, they offer an ability to control *how* a CC system explores that space. We also offer some initial thoughts about "Curious QD" and how a CC system could utilize these models of the behavior space to find specific, radically novel artefacts with high-quality. Although this work is still in the preliminary stages, it appears encouraging as a way to think about and operationalize the concept of transformational creativity.

References

Berlyne, D. E. 1960. *Conflict, Arousal, and Curiosity*. McGraw-Hill Book Company.

Boden, M. 1992. *The Creative Mind*. London: Abacus.

Chatzilygeroudis, K.; Cully, A.; Vassiliades, V.; and Mouret, J.-B. 2021. Quality-diversity optimization: A novel branch of stochastic optimization. In *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*. 109–135.

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

Fontaine, M., and Nikolaidis, S. 2021. Differentiable quality diversity. In *Advances in Neural Information Processing Systems*, 10040–10052.

Fontaine, M. C.; Togelius, J.; Nikolaidis, S.; and Hoover, A. K. 2020. Covariance matrix adaptation for the rapid illumination of behavior space. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 94–102.

Grace, K., and Maher, M. L. 2015. Specific curiosity as a cause and consequence of transformational creativity. In *Proceedings of the International Conference on Computational Creativity*, 260–267.

Hafner, D.; Lillicrap, T. P.; Norouzi, M.; and Ba, J. 2021. Mastering Atari with discrete world models. In *Proceedings of the International Conference on Learning Representation*.

Hansen, N. 2016. The CMA evolution strategy: A tutorial. *ArXiv* abs/1604.00772.

Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, 6405–6416.

Maher, M. L., and Poon, J. 1996. Modeling design exploration as co-evolution. *Computer-Aided Civil and Infrastructure Engineering* 11(3):195–209.

Mouret, J.-B., and Clune, J. 2015. Illuminating search spaces by mapping elites. *ArXiv* abs/1504.04909.

Ong, Y. S.; Nair, P. B.; and Keane, A. J. 2003. Evolutionary optimization of computationally expensive problems via surrogate modeling. *AIAA Journal* 41(4):687–696.

Osband, I.; Aslanides, J.; and Cassirer, A. 2018. Randomized prior functions for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, 8626–8638.

Oudeyer, P.-Y. 2004. Intelligent adaptive curiosity: A source of self-development. In *Proceedings of the Fourth International Workshop on Epigenetic Robotics*, 127–130.

Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. "Why should I trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144.

Schmidhuber, J. 2010. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development* 2(3):230–247.

Schon, D. A., and Wiggins, G. 1992. Kinds of seeing and their functions in designing. *Design Studies* 13(2):135–156.

Sohn, K.; Yan, X.; and Lee, H. 2015. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, 3483–3491.

Torrance, E. P. 1966. Torrance tests of creative thinking. *Educational and Psychological Measurement*.

van den Oord, A.; Vinyals, O.; and Kavukcuoglu, K. 2017. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, 6309–6318.

Wierstra, D.; Schaul, T.; Peters, J.; and Schmidhuber, J. 2008. Natural evolution strategies. In *Proceedings of the IEEE Congress on Evolutionary Computation*, 3381–3387.

Wiggins, G. A. 2006. Searching for computational creativity. *New Generation Computing* 24:209–222.

Zhang, Y.; Fontaine, M. C.; Hoover, A. K.; and Nikolaidis, S. 2022. Deep surrogate assisted MAP-elites for automated hearthstone deckbuilding. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 158–167.