Are Quality Diversity Algorithms Better at Generating Stepping Stones than Objective-based Search?

Adam Gaier
Bonn-Rhein-Sieg University of Applied Sciences
Inria / CNRS / Université de Lorraine
adam.gaier@h-brs.de

Alexander Asteroth
Bonn-Rhein-Sieg University of Applied Sciences
Sankt Augustin, Germany
alexander.asteroth@h-brs.de

Jean-Baptiste Mouret
Inria / CNRS /
Université de Lorraine
Nancy, France
jean-baptiste.mouret@inria.fr

ABSTRACT

The route to the solution of complex design problems often lies through intermediate "stepping stones" which bear little resemblance to the final solution. By greedily following the path of greatest fitness improvement, objective-based search overlooks and discards stepping stones which might be critical to solving the problem. Here, we hypothesize that Quality Diversity (QD) algorithms are a better way to generate stepping stones than objective-based search: by maintaining a large set of solutions which are of high-quality, but phenotypically different, these algorithms collect promising stepping stones while protecting them in their own "ecological niche". To demonstrate the capabilities of QD we revisit the challenge of recreating images produced by user-driven evolution, a classic challenge which spurred work in novelty search and illustrated the limits of objective-based search. We show that QD far outperforms objective-based search in matching user-evolved images. Further, our results suggest some intriguing possibilities for leveraging the diversity of solutions created by QD.

KEYWORDS

Quality Diversity, Neuroevolution, Indirect Encodings, MAP-Elites

ACM Reference Format:
Adam Gaier, Alexander Asteroth, and Jean-Baptiste Mouret. 2019. Are Quality Diversity Algorithms Better at Generating Stepping Stones than Objective-based Search?. In *Genetic and Evolutionary Computation Conference Companion (GECCO '19 Companion), July 13–17, 2019, Prague, Czech Republic.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3319619.3321897

1 INTRODUCTION

In complex problems it is often not possible to follow incremental improvements to an optimal solution. These deceptive cases, where the path to high-performing regions leads through poor-performing regions, can cause difficulty or outright failure of algorithms which only consider optimization of the objective function.

GECCO '19 Companion, July 13–17, 2019, Prague, Czech Republic

ACM ISBN 978-1-4503-6748-6/19/07…$15.00
https://doi.org/10.1145/3319619.3321897

The deficiencies of objective search was made explicit by Woolley and Stanley [9] with an illustrative example drawn from the collaborative evolution project Picbreeder [6]. In Picbreeder users are presented with a set of images, and prompted to select one or more to act as parents for a next generation of evolved images. The parents are combined and mutated to produce a new set of images, which are once again presented to the user for selection. In this way the users evolve increasingly complex images, which can be shared with a community of other users to use as starting points for their own explorations. The results range from abstract shapes to images resembling butterflies, cars, dolphins, and skulls.

In [9] the authors attempted to find these same images using automated evolution, in place of human selection. Interestingly, even though the same method of representing and varying images was used, the algorithm failed dramatically at reproducing all but the simplest images. The automated algorithm chose as parents the closest pixel-by-pixel match to the target image, and it was this "single-minded approach" which the authors blamed for the failure.

This finding echoes those obtained a few years before with novelty search [4], a thought-provoking alternative to objective-driven search. In its purest form, novelty search abandons traditional objectives in favor of exploration of phenotypic space, such as robot behaviors or design features. By exploring the space of possible behaviors selection pressure is biased toward "novel"solutions rather than only those which bring the algorithm closer to a stated goal. Novelty-based methods operate by collecting a variety of "stepping stones" which may lead to better solutions, rather than searching for better solutions explicitly.

Looking for "interesting" results is an approach which mirrors the one taken by users of Picbreeder, up to a point: users explore the space of images without a particular goal until they find an image which resembles something they recognize, and then purposefully hone that resemblance. A directed component is similarly necessary to attack this picture matching task. Combinations of novelty search and objective search, known as Quality Diversity (QD) algorithms [2] borrow from both fields to produce a large number of high-quality, but qualitatively varied solutions to a problem.

Here we show that a QD algorithm can perform significantly better at the Picbreeder images matching task.

2 EXPERIMENTS

The evolution of images in Picbreeder is performed by evolving directed graphs called Compositional Pattern Producing Networks (CPPNs) [7]. CPPNs operate like neural networks; they take as inputs the coordinates of a pixel (e.g. x,y) and output a value to color this pixel (e.g. a gray scale intensity). Nodes of a CPPN have a

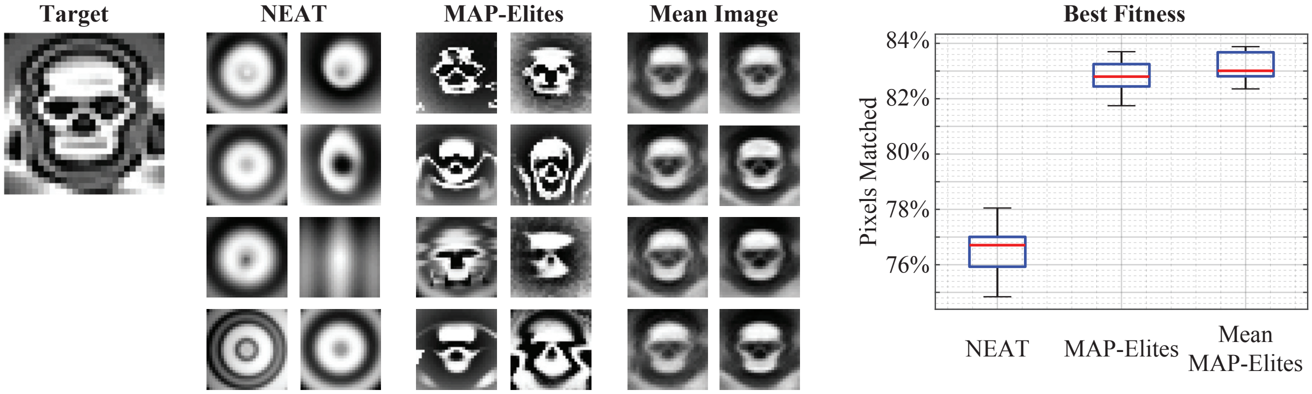

Figure 1: *Picbreeder Pixel Matching with NEAT and MAP-Elites.* From left to right: the target images, a sample of the highest performing images found by NEAT, MAP-Elites, and the mean of all images at the end of the run of MAP-Elites. Images shown are the end result of distinct runs. Far right: distribution of highest fitness found by each approach over 16 runs.

variety of activation functions, such as sin and Gaussian, allowing for the emergence of repetitive and symmetric patterns. CPPNs are evolved with the NEAT [8] algorithm, which evolves networks of increasingly complexity by adding nodes and connections through mutation. This CPPN-NEAT approach was used in both Picbreeder and Wooley's attempts to match the resulting images [9].

To illustrate the capabilities of QD in this problem we maintain the same representation and variation operators, but guide evolution using the QD algorithm MAP-Elites [1, 5]. MAP-Elites maintains an archive of solutions, known as elites, stored in "niches". New solutions are created by mutating these elites, which act as the parent population. Child solutions are then evaluated and assigned a niche based on their features. If this niche is already occupied by an elite, the individual with the higher fitness is placed in the niche and the other discarded. By repeating this process a set of increasingly optimal solutions are discovered for every combination of features. We refer to the archive of elites as a whole as a "map".

We explore two features which define the dimensions of the map. The first, complexity, is simply the number of connections in the network. For the second, we introduce a generic notion of novelty. Periodically all of the images in the map are compiled and used to train an autoencoder [3]; when a new image is produced we compress and reconstruct the image with the autoencoder, examining the pixel-by-pixel error of the reconstruction. Images similar to those which were used to train the autoencoder will be reconstructed well, those different poorly. The reconstruction error is used as generic approximation of "novelty" or "interestingness".

The ability of NEAT and MAP-Elites to match the target images is then compared. Each algorithm was run 16 times, and the closest pixel match recorded. The distribution of the best fitness found over all runs is shown above (Figure 1, right). To illustrate the meaning of this difference qualitatively we show a selection of high fitness images from each of the two approaches, each drawn from a different run (Figure 1, left).

NEAT is unable to find more than a basic circular shape, as reported in previous experiments [9]. Deviating from this shape decreases fitness, and these solutions are less likely to continue to the next generation. In contrast, MAP-Elites collects a variety of shapes, so possible stepping stones are kept, and the resulting images are much closer to the target.

QD produces varied solutions to a problem, and in this diversity differing aspects of the problem are solved. This variety can be combined to powerful effect: if the mean pixel value of all elites in the map is taken this "mean image" is a closer match then the best single image, and more recognizable to the human eye.

3 CONCLUSION

In this picture synthesis experiment, a QD algorithm (here, MAP-Elites) clearly outperforms objective-based search, which suggests that QD algorithms can be better at generating stepping stones than objective-based search. Nevertheless, the images produced by MAP-Elites are not a perfect match with the target image. Intriguingly, in every run the mean of all found solutions in the map produced by MAP-Elites improved on the performance of any individual solution. This points to the potential usefulness of the body of solutions produced by QD when taken as a whole, that is, as an "ensemble". That even the most naive ensemble approach yields results suggests promise in pursuing a more sophisticated way of combining the solutions found by QD.

REFERENCES

[1] A Cully, J Clune, D Tarapore, and J-B Mouret. 2015. Robots that can adapt like animals. *Nature* (2015).
[2] A Cully and Y Demiris. 2018. Quality and diversity optimization: A unifying modular framework. *IEEE Trans. on Evolutionary Computation* (2018).
[3] G Hinton and R Salakhutdinov. 2006. Reducing the dimensionality of data with neural networks. *Science* (2006).
[4] J Lehman and K O Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* (2011).
[5] J-B Mouret and J Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
[6] J Secretan, N Beato, D B D'Ambrosio, A Rodriguez, A Campbell, and K O Stanley. 2008. Picbreeder: evolving pictures collaboratively online. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*. ACM.
[7] K Stanley. 2007. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines* (2007).
[8] K Stanley and R Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary computation* (2002).
[9] B G Woolley and K O Stanley. 2011. On the deleterious effects of a priori objectives on evolution and representation. In *Proc. of GECCO*. ACM.