Quality Diversity: Treating Diversity as the Primary Goal

Justin K. Pugh and Kenneth O. Stanley

University of Central Florida, Department of Computer Science,
4328 Scorpius Street, Orlando, FL 32816-2362, USA
{jpugh,kstanley}@cs.ucf.edu

Abstract. Historically, evolutionary computation and machine learning have focused mainly on optimization problems, where the goal is to find the best-performing (global optimum) point in a fitness landscape. Multimodal optimization takes a step away from this paradigm by searching for a collection of local optima, employing techniques such as speciation and fitness sharing to encourage the search to explore multiple "niches" simultaneously. However, even when the goal is to return multiple solutions, the search strategies employed by multimodal optimization, like those of single-solution approaches, are inherently *convergent*: over time the search focuses only on the best-performing regions. This property of convergence is inconvenient for some domains, especially those in computational creativity and open-ended evolution, where it may be desirable to find the best performance at *all* parts of the search space (including regions of lower performance). Such applications are better suited to a *divergent* search designed to fully explore the vast space of possibilities. Addressing this need, a new class of *quality diversity* (QD) algorithms has begun to emerge where the primary goal is to explore the entire space (i.e. discover all possible niches) and optimization within each niche is constrained so that all niches are considered equally important.

Keywords: novelty search, non-objective search, quality diversity, behavioral diversity, neuroevolution

Acknowledgments. This work was supported by the National Science Foundation under grant no. IIS-1421925. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.