Adaptation and Ruggedness in an Evolvability Landscape

Terry Van Belle and David H. Ackley

Department of Computer Science
University of New Mexico
Albuquerque, New Mexico, USA
{vanbelle, ackley}@cs.unm.edu

Evolutionary processes depend on both *selection*—how fit any given individual may be, and on *evolvability*—how and how effectively new and fitter individuals are generated over time. While genetic algorithms typically represent the selection process explicitly by the fitness function and the information in the genomes, factors affecting evolvability are most often implicit in and distributed throughout the genetic algorithm itself, depending on the chosen genomic representation and genetic operators. In such cases, the genome itself has no direct control over evolvability except as determined by its fitness.

Researchers have explored mechanisms that allow the genome to affect not only fitness but also the distribution of offspring, thus opening up the potential of evolution to improve evolvability. In prior work [1] we demonstrated that effect with a simple model focusing on heritable evolvability in a changing environment.

In our current work [2], we introduce a simple evolvability model, similar in spirit to those of Evolution Strategies. In addition to genes that determine the fitness of the individual, in our model each individual contains a distinct set of 'evolvability genes' that determine the distribution of that individual's potential offspring. We also present a simple dynamic environment that provides a canonical 'evolvability opportunity' by varying in a partially predictable manner.

That evolution might lead to improved evolvability is far from obvious, because selection operates only on an individual's current fitness, but evolvability by definition only comes into play in subsequent generations. Two similarly-fit individuals will contribute about equally to the next generation, even if their evolvabilities vary drastically. Worse, if there is any fitness cost associated with evolvability, more evolvable individuals might get squeezed out before their advantages could pay off. The basic hope for increasing evolvability is circumstances where weak selective pressure allows diverse individuals to contribute offspring to the next generation, and then those individuals with better evolvability in the current generation will tend to produce offspring that will dominate in subsequent fitness competitions. In this way, evolvability advantages in the ancestors can lead to fitness advantages in the descendants, which then preserves the inherited evolvability mechanisms.

A common tool for imagining evolutionary processes is the fitness landscape, a function that maps the set of all genomes to a single-dimension real fitness value. Evolution is seen as the process of discovering peaks of higher fitness, while avoiding valleys of low fitness. If we can derive a scalar value that plau-

E. Cantú-Paz et al. (Eds.): GECCO 2003, LNCS 2723, pp. 150–151, 2003.

sibly captures the notion of evolvability, we can augment the fitness landscape conception with an analogous notion of an *evolvability landscape*. With our algorithm possessing variable and heritable evolvabilities, it is natural to wonder what the evolution of a population will look like on the evolvability landscape as well as the fitness landscape.

We adopt as an evolvability metric the *online fitness* of a population: The average fitness value of the best of population from the start of the run until a fixed number of generations have elapsed. The online fitness of a population with a fixed evolvability gives us the 'height' of the evolvability landscape at that point. In cases where evolvability is adaptive, we envision the population moving across the evolvability landscape as evolution proceeds, which in turn modifies the fitness landscape. Figures 1 and 2 show some of our results.

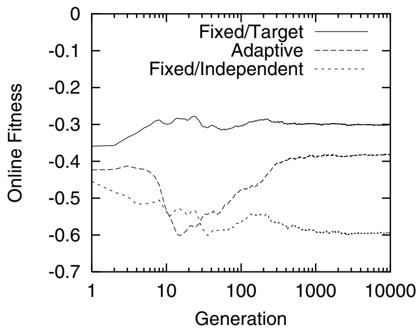

Fig. 1. *Fixed/Independent* is standard GA evolvability, in which all gene mutations are independent. *Fixed/Adaptive*, with an evolvable evolvability, does significantly better. *Fixed/Target* does best, but assumes advance knowledge of the environmental variation pattern.

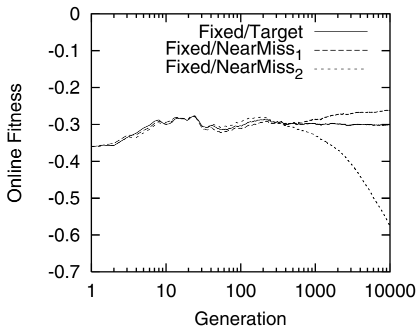

Fig. 2. Evidence of a 'cliff' in the evolvability landscape. Fixed evolvabilities that are close to optimal, but not exact, can produce extremely poor performance.

Acknowledgments. This research was supported in part by DARPA contract F30602-00-2-0584, and in part by NSF contract ANI 9986555.

References

[1] Terry Van Belle and David H. Ackley. Code factoring and the evolution of evolvability. In *Proceedings of GECCO-2002*, New York City, July 2002. AAAI Press.

[2] Terry Van Belle and David H. Ackley. Adaptation and ruggedness in an evolvability landscape. Technical Report TR-CS-2003-14, University of New Mexico, Department of Computer Science, 2003.
`http://www.cs.unm.edu/colloq-bin/tech_reports.cgi?ID=TR-CS-2003-14`.