Chapter 1.18
On Measuring the Attributes of Evolutionary Algorithms:
A Comparison of Algorithms Used for Information Retrieval

J. L. Fernández-Villacañas Martín
Universidad Carlos III, Spain

P. Marrow and M. Shackleton
Intelligent Systems Laboratory, BTextract Technologies, UK

ABSTRACT

In this chapter we compare the performance of two contrasting evolutionary algorithms addressing a similar problem, of information retrieval. The first, BTGP, is based upon genetic programming, while the second, MGA, is a genetic algorithm. We analyze the performance of these evolutionary algorithms through aspects of the evolutionary process they undergo while filtering information. We measure aspects of the variation existing in the population undergoing evolution, as well as properties of the selection process. We also measure properties of the adaptive landscape in each algorithm, and quantify the importance of neutral evolution for each algorithm. We choose measures of these properties because they appear generally important in evolution. Our results indicate why each algorithm is effective at information retrieval, however they do not provide a means of quantifying the relative effectiveness of each algorithm. We attribute this difficulty to the lack of appropriate measures available to measure properties of evolutionary algorithms, and suggest some criteria for useful evolutionary measures to be developed in the future.

INTRODUCTION

Evolutionary methods have been the focus of much attention in computer science, principally because of their potential for performing a partially directed search in very large combinatorial spaces (Sloman, 1998). Evolutionary algorithms (EAs) have the potential to balance exploration of the search space with exploitation of useful features of that search space. However the correct balance is difficult to achieve and places limits on what can be predicted about the algorithm's behaviour. In addition, EAs are often implemented in system-specific ways, making it very difficult to predict and evaluate performance on different implementations. This makes the need for measures to evaluate and compare different algorithms all the more urgent.

In this chapter we focus upon the comparison between algorithms for information retrieval. This is one of the tasks at which evolutionary algorithms have been found particularly effective. Such algorithms deal with the situation where a relevant sub-set of documents or records must be isolated from a larger pool. This chapter considers two such algorithms which were developed for the task of information filtering in a telecommunications context. The BTGP is a genetic programming system where the programs produced execute Boolean searches through keywords (Fernández-Villacañas & Exell, 1996). The MGA is a genetic algorithm which also uses a Boolean tree representation, through a relatively complicated mapping between genotype and phenotype.

We compare the performance of these algorithms using a collection of measures chosen from consideration of evolutionary processes. Such measures have been developed within an evolutionary computation context and also within evolutionary biology. To understand why such measures might be useful, we first consider the evolutionary process itself.

Evolution can be described as "*...any net directional change or any cumulative change in the characteristics of organisms or populations over many generations ...*" (Endler, 1986).

But this evolutionary change may occur as the consequence of a number of different processes, acting to differing extent. Comparison of biological and computational evolution shows the importance of three classes of phenomena in making natural and artificial evolutionary systems evolvable. These are variation, selection and adaptive landscape structure.

The existence of variation is crucially important for evolutionary processes because there would otherwise be no possibility for the selection scheme to exploit the search space. Measuring the amount of variation gives an indication of the potential of the population to be selected, although it does not of course tell about the potential of the population to vary in the future. Ideally we need to know about the propensity of the population to vary in the future in order to get a full picture of the evolvability of the system. This distinction between the amount of variation and the variability of a population has been emphasized, in the context of evolvability by Wagner and Altenberg (1996).

Variation may be measured through genetic variance, which can be calculated provided it is possible to set values on the different genetic variants present (Falconer, 1989; Lynch & Walsh, 1998). Depending on the evolutionary algorithm under consideration, it may be more appropriate to take a phenotypic variance measure, as the representation of the genotype in the phenotype may crucially affect the way in which available variation influences the selective process. The method of measurement of phenotypic variation will depend upon the representation used.

Mutation is an important means of generating further variation, and acts in part to counteract the loss of variation through selection. It must therefore be important for evolvability. It is with this in mind that Wagner and Altenberg (1996) have proposed mutational variance, the variation in effect of possible mutants that can arise in a population, as a measure of the evolvability or evolutionary performance of a system. While mutation variance may be very difficult to calculate in natural populations, it is at least in principle derivable for a given evolutionary algorithm.