Negative Slope Coefficient: A Measure to Characterize Genetic Programming Fitness Landscapes

Leonardo Vanneschi[1], Marco Tomassini[2], Philippe Collard[3], and Sébastien Vérel[3]

[1] Dipartimento di Informatica, Sistemistica e Comunicazione (D.I.S.Co.), University of Milan-Bicocca, Milan, Italy
[2] Computer Systems Department, University of Lausanne, Lausanne, Switzerland
[3] I3S Laboratory, University of Nice, Sophia Antipolis, France

Abstract. Negative slope coefficient has been recently introduced and empirically proven a suitable hardness indicator for some well known genetic programming benchmarks, such as the even parity problem, the binomial-3 and the artificial ant on the Santa Fe trail. Nevertheless, the original definition of this measure contains several limitations. This paper points out some of those limitations, presents a new and more relevant definition of the negative slope coefficient and empirically shows the suitability of this new definition as a hardness measure for some genetic programming benchmarks, including the multiplexer, the intertwined spirals problem and the royal trees.

1 Introduction

What makes a problem easy or hard for Evolutionary Algorithms? A first effort to answer this challenging question has been done by Goldberg and coworkers (e.g., see [3, 5]) in the field of Genetic Algorithms (GAs). Their approach consisted in constructing functions that should *a priori* be easy or hard for GAs to solve. These ideas have been followed by many others (e.g. [12, 4]) and have been the source of many hypotheses as to what makes a problem easy or difficult for GAs. One concept that underlies many of these approaches is the notion of *fitness landscape* [15]. A fitness landscape is a plot where the points in the horizontal subspace represent the different individual genotypes in a search space and the points in the vertical direction represent their fitness [9]. Individual genotypes are usually placed on the horizontal subspace according to a certain neighborhood structure. If genotypes can be visualized in two dimensions, the plot can be seen as a three-dimensional surface, which may contain peaks and valleys. The task of finding the best solution to the problem is equivalent to finding the highest peak (for maximization problems). The problem solver is seen as a short-sighted explorer searching for it. The fitness landscape plot can be helpful to understand the difficulty of a problem, i.e. the ability of a searcher to find the optimal solution for that problem (see for instance [17] for a deep analysis). Nevertheless, fitness landscapes are impossible to be plotted in practice, given the generally huge size of the space of solutions and the multi-dimentionality and complexity of the possible neighborhood structures. For this reason, in the last few years researchers have been looking for an algebraic measure able to capture some of the interesting properties of fitness landscapes. Early attemps

P. Collet et al. (Eds.): EuroGP 2006, LNCS 3905, pp. 178–189, 2006.
© Springer-Verlag Berlin Heidelberg 2006

are represented by [22, 11, 7]. A signicant contribution to this field has been given by Jones [6] with the introduction of an hardness measure for GAs called *fitness distance correlation* (*fdc*). This measure has been extended to tree-based Genetic Programming (GP) and proven a suitable hardness indicator in [17, 16, 19, 20, 2]. Nevertheless, these contributions have also shown that *fdc* has some flaws, the most important one being the fact that *fdc* is not predictive, i.e. the optimal solution (or solutions) must be known beforehand, which is almost unrealistic in applied search and optimization problems. Thus, it is important to investigate other approaches based on quantities that can be measured without any explicit knowledge of the genotype of optimal solutions. Preliminary results of this enquiry can be found in [18], where a new measure called *negative slope coefficient* (*nsc*) has been introduced.

This paper aims at extending and generalizing the study of *nsc* for tree-based GP. It is structured as follows: section 2 introduces the concept of *fitness cloud* on which *nsc* is based. Section 3 takes up the original definition of the *nsc* as it has been presented in [18] and section 4 points out its main limitations. Section 5 proposes a new method for calculating the *nsc* and shows some experimental results pointing out that this method enables to overcome some drawbacks of the original definition of the *nsc*. Finally, section 6 offers our conclusions and hints for future research.

2 Fitness Clouds

Evolvability is a feature that is intuitively related, although not exactly identical, to problem difficulty. It has been defined as the capability of genetic operators to improve fitness quality [1]. The most natural way to study evolvability is, probably, to plot the fitness values of individuals against the fitness values of their neighbors, where a neighbor is obtained by applying one step of a genetic operator to the individual. Such a plot has been first introduced for binary landscapes by Vérel and coworkers [21] and called by them *fitness cloud*. In this paper, the genetic operator used to generate fitness clouds is standard subtree mutation [8], i.e. mutation obtained by replacing a subtree of the selected individual with a randomly generated tree.

2.1 Definition

Let $\Gamma = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$ be the whole search space of a GP problem and let $V(\gamma_j) = \{v_1^j, v_2^j, \dots, v_{m_j}^j\}$ be the set of all the neighbors of individual $\gamma_j, \forall j \in [1, n]$. Now let f be the fitness function of the problem at hand. The following set of points can be defined: $S = \{(f(\gamma_j), f(v_k^j)),\ \forall j \in [1, n],\ \forall k \in [1, m_j]\}$. The graphical representation of S on a bidimentional plane, or fitness cloud, is the scatterplot of the fitness of all the individuals belonging to the search space against the fitness of all their neighbors. The main idea is that the shape of this scatterplot can give an indication of the evolvability of the genetic operators used and thus some hints about the difficulty of the problem at hand. The fitness cloud also implicitly gives some insight on the genotype to phenotype map: the set of genotypes that all have equal fitness is a *neutral set*. Such a set corresponds to one abscissa in the fitness/fitness plane; according to this abscissa, a vertical slice from the cloud represents the set of fitnesses that could be reached from this set of neutrality.