Available online at www.sciencedirect.com
ScienceDirect
European Journal of Operational Research 175 (2006) 806–820
EUROPEAN JOURNAL OF OPERATIONAL RESEARCH
www.elsevier.com/locate/ejor

Computing, Artificial Intelligence and Information Management

Learning genetic algorithm parameters using hidden Markov models

Jackie Rees [a,*], Gary J. Koehler [b,1]

[a] *Krannert Graduate School of Management, 403 West State Street, Purdue University, West Lafayette, IN 47907-2056, USA*
[b] *Department of Decision and Information Sciences, Warrington College of Business, University of Florida, P.O. Box 117169, Gainesville FL 32611-7169, USA*

Received 7 October 2003; accepted 14 April 2005
Available online 15 August 2005

Abstract

Genetic algorithms (GAs) are routinely used to search problem spaces of interest. A lesser known but growing group of applications of GAs is the modeling of so-called ''evolutionary processes'', for example, organizational learning and group decision-making. Given such an application, we show it is possible to compute the likely GA parameter settings given observed populations of such an evolutionary process. We examine the parameter estimation process using estimation procedures for learning hidden Markov models, with mathematical models that exactly capture expected GA behavior. We then explore the sampling distributions relevant to this estimation problem using an experimental approach.

Keywords: Artificial intelligence; Evolutionary computations; Evolutionary process; Adaptive agents; Genetic algorithms; Heuristics; Markov processes

1. Introduction

The area of experimental economics (e.g., Kagel and Roth, 1995) studies complex, multi-agent systems within a computer-simulated environment. Often it is desirable to have the artificial agents adapt to various events and pressures within their environment. One popular type of adaptive behavior is modeled after natural evolutionary processes. A simple, yet powerful, form of this behavior is captured in genetic algorithms (GAs) (see Goldberg, 1989).

* Corresponding author. Tel.: +1 765 494 0320; fax: +1 765 494 9658.
E-mail addresses: jrees@purdue.edu (J. Rees), koehler@ufl.edu (G.J. Koehler).
[1] Tel.: +1 352 846 2090; fax: +1 352 392 5438.

0377-2217/$ - see front matter
doi:10.1016/j.ejor.2005.04.045

Several authors have studied the evolutionary characteristics of systems by creating a simulation environment where a GA is used to mimic the adaptive behavior of some agent or group of agents. For example, Marks et al. (1995) examined oligopoly behavior in an adaptive framework and used a GA to simulate market-pricing movements in the coffee industry. An evolutionary model for electronic commerce was presented in Oliver (1996). Using a GA for learning, automated agents learned strategies for business negotiations within an electronic commerce framework. This investigation focused on a series of simulations of automated negotiation tasks using a genetic algorithm on automated agents. In Sikora and Shaw (1996), a GA was used to simulate group interaction leading to a solution among groups of agents. Genetic programming, (see Koza, 1992), another evolutionary computing technique, was used by Dworman et al. (1996) to automate the discovery of game theory models. A computational model of the organization based on the simple genetic algorithm was created by Bruderer and Singh (1996). Their research is unique in that it specifically uses a genetic algorithm as the model for organizational evolution, as opposed to simply simulating behavior using a genetic algorithm. In the above examples using evolutionary techniques to simulate or model particular processes, the authors have had to provide very rough estimates of parameters for the evolutionary technique used, for example, crossover or mutation rates, typically without much guidance of which values might be appropriate. Clearly, these simulations and models could progress further if more were known about the appropriate settings for the evolutionary parameters.

In Rees and Koehler (2000, 2002) the process is reversed. Rather than simulate behavior, the authors started with experimental data from an actual adaptive process (in this case a group decision-making case) and used the data to find a best-fit GA to mimic the evolutionary path. This fitting process was used to determine parameters required by the GA. They formulated and tested various hypotheses for these parameters (Rees and Koehler, 2000).

In traditional GA simulation, the GA parameters, χ, the crossover rate, and μ, the mutation rate, are either determined experimentally by running a series of preliminary simulations or are chosen based on previous results or standard values used in the GA community. As discussed above, there are many real-world phenomena that appear to possess evolutionary characteristics, not unlike a GA. Following the lead given in Rees and Koehler (2000, 2002), we seek to investigate further whether there is a way to characterize these processes in terms of GA instances? In other words, can we reliably learn GA parameters, such as crossover and mutation rates, from real-world processes?

The expected behavior of a GA process can be modeled exactly using a Markov chain (Nix and Vose, 1992). If we know or assume a real-world process is a GA process, then we have a hidden Markov model (HMM)—we know it can be modeled by a Markov chain, we just do not know the specific mutation and crossover rates that generate the transition probabilities. The objective of this research is to use HMM methods to compute the likely GA parameter settings given observed populations of such an evolutionary process. In this paper we study the process of learning or "fitting" GA parameters to such evolutionary processes in more detail than done in Rees and Koehler (2000, 2002). We examine this issue by using mathematical models that exactly capture expected GA behavior and explore the sampling distributions relevant to this estimation problem using an experimental approach.

The value of this research is to provide researchers with a tool to more accurately simulate real-world evolutionary behavior. Such simulations are becoming commonplace, especially in applications such as information security, artificial markets and retail management. In such simulations, there is no theoretical guidance on how to set GA parameters such as crossover and mutation rates. The technique described in this paper allows researchers to determine these parameter values from existing real-world data. This ability should lead to more accurate and useful simulation studies.

The remainder of this paper is organized as follows: Section 2 presents relevant background on genetic algorithms, on the applications of hidden Markov models to this problem, and specifically on the use of maximum likelihood estimates for computing best-fit genetic algorithm parameters, namely crossover and mutation rates, from experimental data. Section 3 describes the experimental study comparing known