
Preface

Evolutionary computation is a class of problem optimization methodology with the inspiration from the natural evolution of species. In nature, the population of a species evolves by means of selection and variation. These two principles of natural evolution form the fundamental of evolutionary algorithms (EAs). During the past several decades, EAs have been extensively studied by the computer science and artificial intelligence communities. As a class of stochastic optimization techniques, EAs can often outperform classical optimization techniques for difficult real world problems.

Due to the ease of use and robustness, EAs have been applied to a wide variety of optimization problems. Most of these optimization problems tackled are stationary and deterministic. However, many real-world optimization problems are subjected to dynamic and uncertain environments that are often impossible to avoid in practice. For example, the fitness function is uncertain or noisy as a result of simulation errors, measurement errors or approximation errors. In addition, the design variables or environmental conditions may also perturb or change over time. For these dynamic and uncertain optimization problems, the objective of the EA is no longer to simply locate the global optimum solution, but to continuously track the optimum in dynamic environments, or to find a robust solution that operates optimally in the presence of uncertainties. This poses serious challenges to classical optimization techniques and conventional EAs as well. However, conventional EAs with proper enhancements are still good tools of choice for optimization problems in dynamic and uncertain environments. This is because EAs are inspired by principles of natural evolution, which takes place in the ever-changing dynamic and uncertain environment in nature.

Handling dynamic and uncertain optimization problems has been a topic since the early days of evolutionary computation and has received increasing research interests over recent years due to its challenge and its importance in practice. Several events, e.g., journal special issues, workshops and conference special sessions, have taken place in recent years in the field of evolutionary computation in dynamic and uncertain environments. A variety of methods

have been reported across a broad range of application backgrounds in recent years. This motivated the project of this book. This book aims to timely reflect the most recent advances, present sophisticated real-world applications, and explore future research directions in the field.

We have a total of 26 chapters in this book, which cover a broad range of topics relevant to evolutionary computation in dynamic and uncertain environments. Further, the chapters in this book are presented as the following four categories:

- Part I: Optimum Tracking in Dynamic Environments
- Part II: Approximation of Fitness Functions
- Part III: Handling Noisy Fitness Functions
- Part IV: Search for Robust Solutions

Part I: Optimum Tracking in Dynamic Environments

Most problems studied by the evolutionary computation community are stationary optimization problems where no change occurs over time. For stationary optimization problems, the goal is to design EAs that can quickly and precisely locate the optimal solution(s) to the problem at hand. However, for dynamic optimization problems (DOPs) where change occurs over time, the main task is not to find one optimal solution but to track the moving optimum as soon and narrow as possible. This poses a serious challenge to conventional EAs due to the convergence problem. For stationary optimization problems, convergence at a proper pace other than premature convergence is exactly what is expected for EAs to locate the optimal solution. However, convergence becomes a big problem for DOPs because once converged, it is difficult for conventional EAs to adapt to the changing environment. DOPs usually require EAs to maintain certain level of diversity in the population. In order to deal with this problem, several approaches have been developed in recent years to enhance the performance of EAs in dynamic environments. Part I of the book encapsulates nine chapters that reflect the state-of-the-art research on EAs for problem optimization in dynamic environments and their application to real world dynamic problems.

The first six chapters of Part I present advanced EA approaches for general DOPs. In Chapter 1, Yang investigates the application of two kinds of explicit memory schemes, direct memory and associative memory, for genetic algorithms (GAs) and univariate marginal distribution algorithms (UMDAs) for DOPs. Based on a series of systematically constructed dynamic test environments, experiments are carried out to compare the direct and associative memory schemes for GAs and UMDAs. Blackwell in Chapter 2 studies the use of charged swarms in the particle swarm optimization (PSO) algorithm for DOPs. A self-adapting multi-swarm approach with an exclusion operator that provides effective repulsion between swarms is advocated in this chapter.

A simple rule for swarm birth and death is proposed so that the multi-swarm may adjust its size dynamically and in relation to the number of peaks in the dynamic environments. Chapter 3 by Schönemann experimentally investigates evolution strategies (ESs) for dynamic numerical optimization problems. The results demonstrate that self-adaptive ESs are powerful methods for dynamic environments. To avoid the handicaps of existing performance measures, a new measurement, called average best function value (ABFV), is developed to compare EAs for DOPs. This chapter also discusses the choice for different strategy parameters, e.g., the optimal number of mutation step sizes, for ESs for practical application. An orthogonal dynamic hill-climbing algorithm (ODHC) is presented by Zeng et al. in Chapter 4 for continuous DOPs. In ODHC, the local peak climber is not a solution, but a “niche” (a small hyper-rectangle). An orthogonal design method is employed on the niche in order to seek a potential peak more quickly. An archive is also used to store the latest found higher peaks, so the ODHC algorithm can learn from the past search. Chapter 5 by Tinós and Yang presents a self-organizing random immigrants scheme for GAs to address DOPs. In this scheme, the worst individual and its neighbours are replaced by random immigrants, which are placed in a sub-population to protect them from being replaced by fitter individuals in the main population. In this way, when the fitness of the individuals are close, one single replacement of an individual can affect a large number of individuals of the population in a chain reaction. This simple approach can take the system to a self-organization behaviour, which is useful for GAs in dynamic environments. Bosman in Chapter 6 investigates the use of learning and anticipation for EAs for online DOPs. The time-linkage property, i.e., decisions taken now may influence the score in the future, has been identified as an important source of problem-difficulty. A means to address time-linkage is to predict the future (i.e. anticipation) by learning from the past. This is formalized into an algorithmic framework. Experimental results show that in the presence of time-linkage EAs based on this algorithmic framework outperform conventional EAs.

The last three chapters of Part I present work on the application of EAs for real world dynamic problems. In Chapter 7, Dam et al. investigates XCS, a genetics-based learning classifier system, for online dynamic data mining problems with different degrees of concept changes. In order to reduce the recovery time of XCS after concept changes, three strategies are proposed to force the system to learn quickly after severe changes. The effect of noise on the recovery time after a concept change is also experimentally investigated. Chapter 8 by Michalewicz et al. discusses the prediction and optimization issues in dynamic environments and suggests a system architecture, called Adaptive Business Intelligence, to handle a kind of real world problems where the evaluation functions are based on the prediction of the future values of some variables. Three diverse case studies in dynamic environments: pollution control, ship navigation, and car distribution, are presented. All these problems require some level of prediction and optimization for recommending

the best course of action. Quintão et al. in Chapter 9 present the application of EAs to the area coverage and node connectivity problems in wireless sensor networks (WSNs), a kind of ad-hoc networks with distributed communication, sensing, and processing capacities. EAs are provided to support the network manager with the concern of controlling the energy consumption in the network and the quality of service.

Part II: Approximation of Fitness Functions

A continuing trend in science and engineering is the use of increasingly accurate simulation codes in the design and analysis process so as to produce ever more reliable and high quality products. Such technologies now play a central role in aiding scientists validate crucial designs and to study the effects of altering key design parameters on product performance. Nonetheless, the use of accurate simulation methods can be very timing consuming, leading to possibly unrealistic design cycle. Further, it poses a serious impediment to the practical application of existing optimization methods for automatically establishing the critical design parameters present in real world problems in science and engineering. Particularly, EAs typically require many thousands of function calls to the simulation codes in order to locate a near optimal solution. One promising way to significantly reduce the computational cost of EAs by employing computationally cheap approximation models or surrogates in place of the original computationally expensive fitness functions during evolutionary optimization. The five chapters showcased in Part II of the book reflect the recent state-of-the-art research on single and multi-objective evolutionary frameworks for tackling problems with computationally expensive optimization functions in the context of real world applications.

To reduce the number of expensive fitness function evaluations in evolutionary optimization, Graning et al. in Chapter 10 present a study on several individual-based and generation-based adaptive strategies for neural network metamodel management. In their preliminary study, it was reported that some of adaptation mechanisms proposed do not perform well as expected. The individual-based meta-model management was found to be most promising among all and subsequently applied to real world 3D blade design optimization problem. Song in Chapter 11 considers the use of approximation models based on Gaussian Processes for structural shape optimization. Application examples of the proposed surrogate-assisted evolutionary approaches are given in areas of firtree shape optimization using finite element method and engine nacelle optimization using computational fluid dynamics.

The next three chapters contributed by Reyes-Sierra and Coello in Chapter 12, Deb and Nain in Chapter 13, and Mack et al. in Chapter 14 present three independent studies on using approximation models in the context of multi-objective optimization. In particular, Reyes-Sierra and Coello present an empirical study on using fitness inheritance over approximation

models in the context of PSO and multi-objective optimization for enhancing evolutionary search. Deb and Nain, on the other hand, present a successive fitness landscape modelling for reducing the exact function evaluation calls while retaining the basic search capability of NSGA-II. Using a case study in space propulsion, Mack et al. show that besides obtaining substantial improvements in the efficiency of the evolutionary search, surrogate-based optimization is also useful for novel or exploratory design tasks by offering a global view of the characteristics of the design space, thus enabling one to define previously unknown feasible design space boundaries and to reveal important physics in the design.

Part III: Handling Noisy Fitness Functions

It rarely happens that the fitness of real-world problems can be calculated by a deterministic analytical function. In most cases, the quality of a candidate solution has either to be measured by sensors or estimated using a numerical method. The sensory measurements are usually contaminated with noise in the environment, and the estimations are often subject to randomness. Though EAs are more robust against noise compared to derivative-dependent optimization methods, special attention needs to be paid in many cases. This part of the book presents four interesting chapters describing various approaches to handling noise in fitness evaluations.

Chapter 15 by Neri and Mäkinen describes a hierarchical EA for optimal design of an electrical grounding grid and an elastic structure. In this hierarchical algorithm, the fitness of a population depends on the results from another population, which is therefore noisy. To achieve reliable results, counter measures including population sizing, sampling sizing and survivor selection are taken. Evolution of multi-rover systems in noisy environments has been discussed in Chapter 16 by Tumer and Agogino. Since it is impractical to evaluate the fitness of rovers in collective, the authors presented different methods for designing the fitness function for individual rovers without degrading the performance. Noise introduced by sensors are also considered. A memetic algorithm combining a trust-region based local search with evolutionary global search is presented in Chapter 17 where a trust-region method is combined with an evolutionary search. It is shown that on the one hand, evolutionary algorithms are inherently more robust against noise due to their derivative-free characteristics, the quadratic model used for fitness estimation also contributes to reducing the influence of noise. Chapter 18 by Tezuka et al deals with a financial optimization problem where the fitness values are based on a Monte Carlo method. The explicit sampling method is adopted for reducing the influence of noise. To reduce the computational costs, a selection efficiency index is proposed and the sampling size is adapted in such a way that the selection efficiency is maximized.

Part IV: Search for Robust Solutions

Solving optimization problems using EAs has always been perceived as finding the optimal solution over the entire search space. However, the global optima may not always be the most desirable solution in many real world engineering design problems. In practice, if the global optimal solution is very sensitive to uncertainties, for example, small changes in design variables or operating conditions, then it may not be appropriate to use this highly sensitive solution. Part III showcases eight chapters primarily on new methodologies of EAs for robust search.

Lim et al. in Chapter 19 report a study on several single and multi-objective inverse robust evolutionary optimization schemes that make little assumption on the uncertainty structure. The inverse approach searches for solutions that guarantee a certain degree of maximum uncertainty and, at the same time, satisfy the desired nominal performance of the final design solution. A multi-objective algorithm is also proposed in Chapter 20 by Goh and Tan for robust optimization. Their method incorporates the features of micro-GA (as a local search) to locate a worst case scenario of the candidate solution, a memory-based feature of tabu restriction to guide the evolutionary process and periodic re-evaluation of archived solutions to reduce uncertainty of evolved solutions. In the context of real world robust design applications, Hu et al. in Chapter 21 describe a robust design approach that exploits the open-ended topological synthesis capability of genetic programming and bond graph modelling (GPBG) for evolving robust lowpass and highpass analog filters with respect to parameter perturbations. Handa et al. on the other hand, describes a novel route planning memetic optimization system for a fleet of salting trucks that remains robust under different road temperatures and different temperature distributions in a road network in Chapter 22. Fan et al. in Chapter 23 report a method for robust layout synthesis of micro-electromechanical resonators subjected to inherent geometric uncertainties such as the fabrication error on the sidewall of the structure. An alternative technique that hybridizes EAs and Interval Arithmetic is also described in Chapter 24 by Rocco et al. Barrico and Antunes in Chapter 25 present the concept of degree of robustness in a multi-objective evolutionary approach. The information on the degree of robustness of solutions can then be used to support the decision maker in the selection of a robust compromise solution. Finally, Ling et al. report a study on the effect of the sampling number of Monte-Carlo simulation method used in a standard crowding genetic algorithm for robust optimal design of varied-line-spacing holographic grating in recording optics in Chapter 26.

Generally speaking, this book fulfils the original aims quite well. The four parts represent a great variety of work in the area of evolutionary computation in dynamic and uncertain environments. We hope that the publication of this book will further promote this emerging research field.

Finally, we would like to thank Dr. Janusz Kacprzyk for inviting us to edit this book in the Springer book series “Studies in Computational Intelligence”. We acknowledge the authors for their fine contributions and cooperation during the book preparation. We are grateful to Thomas Ditzinger and Heather King of Springer for their kind support for this book.

Shengxiang Yang
Yew-Soon Ong
Yaochu Jin
November 2006