ARTICLE IN PRESS

European Journal of Operational Research xxx (2014) xxx–xxx

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

GP-DEMO: Differential Evolution for Multiobjective Optimization based on Gaussian Process models

Miha Mlakar *, Dejan Petelin, Tea Tušar, Bogdan Filipič

Jožef Stefan Institute, Jamova cesta 39, SI-1000 Ljubljana, Slovenia
Jožef Stefan International Postgraduate School, Jamova cesta 39, SI-1000 Ljubljana, Slovenia

ARTICLE INFO

Article history:
Received 14 June 2013
Accepted 7 April 2014
Available online xxxx

Keywords:
Multiple objective programming
Evolutionary algorithms
Surrogate models
Gaussian Process modeling
Probable Pareto dominance

ABSTRACT

This paper proposes a novel surrogate-model-based multiobjective evolutionary algorithm called Differential Evolution for Multiobjective Optimization based on Gaussian Process models (GP-DEMO). The algorithm is based on the newly defined relations for comparing solutions under uncertainty. These relations minimize the possibility of wrongly performed comparisons of solutions due to inaccurate surrogate model approximations. The GP-DEMO algorithm was tested on several benchmark problems and two computationally expensive real-world problems. To be able to assess the results we compared them with another surrogate-model-based algorithm called Generational Evolution Control (GEC) and with the Differential Evolution for Multiobjective Optimization (DEMO). The quality of the results obtained with GP-DEMO was similar to the results obtained with DEMO, but with significantly fewer exactly evaluated solutions during the optimization process. The quality of the results obtained with GEC was lower compared to the quality gained with GP-DEMO and DEMO, mainly due to wrongly performed comparisons of the inaccurately approximated solutions.

1. Introduction

Optimization problems are present in our everyday life and come in a variety of forms, e.g. the task to optimize certain properties of a system by correctly choosing the system parameters. Many of these optimization problems require the simultaneous optimization of multiple, often conflicting, criteria (or objectives). These problems are called multiobjective optimization problems. The solution to such problems is not a single point, but a family of points, known as the Pareto-optimal set. This set of solutions gives the decision maker an insight into the characteristics of the problem before a single solution is chosen.

One of the most effective ways to solve problems with more objectives is to use multiobjective evolutionary algorithms (MOEAs). MOEAs are population-based algorithms that draw inspiration from optimization processes that occur in nature. During the optimization process, in order to find a Pareto-optimal set, a lot of different solutions have to be assessed (evaluated). These solution evaluations can be costly, dangerous or computationally expensive. In such cases the goal is to minimize the number of exactly evaluated solutions, but still find the best solutions. In this paper we focus on computationally expensive problems, where one solution evaluation takes a lot of time.

In order to obtain the results of such an optimization problem more quickly (or even be able to obtain them in reasonable amount of time), we can use surrogate models in the optimization process to approximate the objective functions of the problem. To evaluate a solution, instead of using a time-consuming exact evaluation, a solution can be approximated with the surrogate model. Since one solution approximation is (much) faster, the whole optimization process can be accelerated. However, note that the time needed to create and update the surrogate models during the optimization process has to be considered and included in the whole duration of the optimization process. So, in the case where the exact solution evaluations are quick, it can happen that the surrogate-model-based optimization takes longer than the optimization without surrogates.

In surrogate-model-based multiobjective optimization, approximated values are often inappropriately used in the solution comparison. As a consequence, exactly evaluated good solutions can be discarded from the population because they appear to be dominated by the inaccurate and over-optimistic approximations. This can slow the optimization process or even prevent the algorithm from finding the best solutions.

* Corresponding author at: Jožef Stefan Institute, Jamova cesta 39, SI-1000 Ljubljana, Slovenia. Tel.: +386 14773633.

E-mail addresses: miha.mlakar@ijs.si (M. Mlakar), dejan.petelin@ijs.si (D. Petelin), tea.tusar@ijs.si (T. Tušar), bogdan.filipic@ijs.si (B. Filipič).

http://dx.doi.org/10.1016/j.ejor.2014.04.011
0377-2217/

Some surrogate models provide a distribution, from which the approximated value and also the confidence interval of the approximation can be calculated. Using this confidence interval, we define new dominance relations that take into account this uncertainty and propose a new concept for comparing solutions under uncertainty that requires exact evaluations only in cases where more certainty is needed. This minimizes the mistakes made in comparisons of inaccurately approximated solutions.

Based on this concept we propose a new surrogate-model-based multiobjective evolutionary algorithm, called Differential Evolution for Multiobjective Optimization based on Gaussian Process modeling (GP-DEMO). This algorithm is an extension of the Differential Evolution for Multiobjective Optimization (DEMO) algorithm (Robič & Filipič, 2005), which uses differential evolution to effectively solve numerical multiobjective optimization problems and, in addition, emphasizes the variation operators. In GP-DEMO, Gaussian Process (GP) modeling is employed to find approximate solution values together with their confidence intervals. Then, instead of comparing the solutions using the Pareto dominance relation, GP-DEMO uses the new uncertainty-based dominance relations, requiring exact evaluations of solutions as needed. The efficiency of GP-DEMO is assessed on several benchmark and two real-world optimization problems.

The structure of this paper is as follows. In Section 2, we overview the work done in the field of surrogate-model-based optimization, especially in multiobjective optimization. In Section 3, we describe the Gaussian Process modeling that is used to build the surrogate models in GP-DEMO. Then, in Section 4, we describe the new relations and methods for comparing solutions (presented with and without uncertainty). The GP-DEMO algorithm is presented in Section 5. In Section 6, we test and compare GP-DEMO with two other algorithms on benchmark and real-world multiobjective optimization problems. Finally, Section 7 concludes the paper with a summary of the work done and our ideas for future work.

2. Related work

In the literature the term surrogate model (sometimes also meta-model) based optimization is used where, during the optimization processes, some solutions are not evaluated with the original objective function, but are approximated using a model of this function. Different modeling methods are used to build the surrogate models. For single and multiobjective optimization similar methods are used. These methods typically return only one approximated value, which is why in multiobjective problems several models have to be used, so that every model approximates one objective. Some of the most commonly used methods are the Response Surface Method (Myers & Montgomery, 1995), Radial Basis Function (Hardy, 1971), Neural Network (Specht, 1990), Kriging (Stein, 1999) and Gaussian Process modeling (MacKay, 1998; Rasmussen & Williams, 2006; Seeger, 2004).

In single-objective optimization, the usage of surrogate models is well established and has proven to be successful. In the literature many different algorithms and various modeling techniques are used to solve benchmark and real-world problems (Emmerich, Giotis, Özdemir, Bäck, & Giannakoglou, 2002; Zhang & Sanderson, 2007). The results typically show that the surrogate-model-based optimization in comparison with optimization without surrogates provides comparable results in fewer objective function evaluations (Jin, Olhofer, & Sendhoff, 2001; Zhou, Ong, Nair, Keane, & Lum, 2007). The use of differential evolution in combination with surrogate models is mentioned in Zhang and Sanderson (2007). The authors presented an algorithm based on differential evolution that generates multiple offspring for each parent and chooses the promising one based on the confidence and the approximation of the current surrogate model.

In the field of surrogate-model-based multiobjective optimization, where the result is not just one solution but a nondominated front of solutions, the problem of finding these solutions is even more challenging. There are many approaches that differ in terms of which solutions are approximated and how they use the approximations. Surrogate models can aim at either a global approximation of the objective function, or a local one, focusing on the neighborhood of the best current individuals. In Zhou et al. (2007), the authors used a combination of local and global surrogate models for solving optimization problem of Aerodynamic Shape Design.

Within surrogate-model-based optimization algorithms a mechanism is needed to find a balance between the exact and approximate evaluations. In evolutionary algorithms this mechanism is called evolution control (Jin, 2003) and can be either fixed or adaptive.

In fixed evolution control, the surrogate model is trained from previously exactly evaluated solutions and then used directly instead of expensive objective function evaluations. In this approach the number of exact function evaluations that will be performed during the optimization is known in advance. Fixed evolution control can be further divided into generation-based control, where in some generations all solutions are approximated and in the others they are exactly evaluated (Deb & Nain, 2007), and individual based control, where in every generation some (usually the best) solutions are exactly evaluated and others approximated (Grierson & Pak, 1993).

In adaptive evolution control, the number of exactly evaluated solutions is not known in advance, but depends on the accuracy of the model for the given problem. Adaptive evolution control can be used in one of two ways: as a part of a memetic search or to pre-select the promising individuals which are then exactly evaluated (Pilat & Neruda, 2012).

In a memetic algorithm, an additional algorithm (e.g., a gradient-based or an evolutionary algorithm) is used to find the optimal solutions using the surrogate model. Once this optimum is found, the best solutions are exactly evaluated and used for updating the model. In Pilat and Neruda (2011), aggregated surrogate models are used in a memetic algorithm. The model is based on the distance to the currently known, nondominated set and is used to find new, nondominated individuals using local search. In memetic algorithms, especially if the surrogate model is not very accurate, a local optimum is often found instead of the global optimum.

In the case of pre-selecting the promising individuals, the surrogate model is used to find the promising or drop the low-quality individuals even before they are exactly evaluated, thus reducing the number of exact evaluations. For example, OEGADO (Chafekar, Shi, Rasheed, & Xuan, 2005) creates a surrogate model for each of the objectives. The best solutions in every objective get also approximated on other objectives, which helps with finding trade-off individuals. The best individuals are then exactly evaluated and used to update the models. ParEGO (Knowles, 2006) uses the weighted sum of the objective functions to perform a local search. The weights are generated randomly for each iteration. When a different model is used for each of the functions, the conversion from the multiobjective problem to the single-objective one has to be performed (or a multiobjective optimizer has to be used on the models). Moreover, if there are more models, their errors can add up, as well as the time needed to train the models.

Surrogate models are also used to rank and filter out offspring according to Pareto-related indicators like the hypervolume (Emmerich, Giannakoglou, & Naujoks, 2006), or a weighted sum of the objectives (Taboada, Baheranwala, Coit, & Wattanapongsakorn, 2007). The problem with the methods that use hypervolume as a