



Tuning metaheuristics: A data mining based approach for particle swarm optimization

Stefan Lessmann^{a,*}, Marco Caserta^{a,1}, Idel Montalvo Arango^{b,2}

^a Institute of Information Systems, University of Hamburg, Von-Melle-Park 5, D-20146 Hamburg, Germany

^b Centro Multidisciplinar de Modelación de Fluidos, Polytechnic University of Valencia, Spain

ARTICLE INFO

Keywords:

Metaheuristics
Particle swarm optimization
Forecasting
Data mining

ABSTRACT

The paper is concerned with practices for tuning the parameters of metaheuristics. Settings such as, e.g., the cooling factor in simulated annealing, may greatly affect a metaheuristic's efficiency as well as effectiveness in solving a given decision problem. However, procedures for organizing parameter calibration are scarce and commonly limited to particular metaheuristics. We argue that the parameter selection task can appropriately be addressed by means of a data mining based approach. In particular, a hybrid system is devised, which employs regression models to learn suitable parameter values from past moves of a metaheuristic in an online fashion. In order to identify a suitable regression method and, more generally, to demonstrate the feasibility of the proposed approach, a case study of particle swarm optimization is conducted. Empirical results suggest that characteristics of the decision problem as well as search history data indeed embody information that allows suitable parameter values to be determined, and that this type of information can successfully be extracted by means of nonlinear regression models.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

The support of managerial decision making is a key area of interest in many scientific domains and several qualitative and quantitative planning methods have been developed to support various types of operational, tactical and strategic planning problems. Especially well-structured and repetitive decision tasks are commonly approached by translating the problem into a mathematical program whose optimal solution is then determined by an optimization procedure like, e.g., linear programming.

Over the past decades, heuristic procedures from the field of intelligent search have received much attention and nowadays represent a well-established approach towards solving complex optimization problems. Offering high-quality solutions in a timely manner, they complement classical mathematical programming methods, especially in large, combinatorial settings where optimal solutions are usually unattainable. Respective techniques have been developed in several domains (e.g., Operational Research, Engineering, Soft Computing, etc.) and include algorithms like simulated annealing (Kirkpatrick, Gelatt, & Vecchi, 1983), tabu-search (Glover & Laguna, 1997), cross-entropy (Boer, Kroese, Mannor, &

Rubinstein, 2005) as well as nature-inspired methods like genetic algorithms (Goldberg, 1989; Holland, 1975), evolution strategies (Beyer & Schwefel, 2002) and, more recently, ant colony optimization (Dorigo & Stützle, 2005) and particle swarm optimization (PSO, Kennedy & Eberhart, 1995). The umbrella term *metaheuristics* is used throughout the paper to subsume respective approaches (see, e.g., Glover, 1986 for a formal definition of a metaheuristic).

A large body of literatures evidences the success of metaheuristics in various applications (see, e.g., Blum & Roli, 2003; Caserta & Voß, 2009; Glover & Kochenberger, 2003 for a survey). One key factor that seems to have a strong impact on the algorithmic performance is the fine tuning of the method; that is, the determination of the metaheuristic's parameters (e.g., the number of generations in genetic algorithms, the length of the tabu-list in tabu-search or the cooling factor in simulated annealing, etc.). According to Adenso-Diaz and Laguna (2006), there is evidence that 10% of the time required to develop a new metaheuristic is devoted to the actual development and that the remaining 90% is spent on fine tuning of algorithmic parameters. Consequently, it is of paramount importance to make a concerted effort in identifying and establishing a set of "standard" techniques to fine-tune a metaheuristic. One of the major achievements of such an effort would be to offset parameter specific issues in evaluating an algorithm. In addition, reproducibility of results would also be enhanced by such an approach, by making transparent the way in which parameter values should be set to tackle a given problem instance.

* Corresponding author. Tel.: +49 40 42838 4706; fax: +49 40 42838 5535.

E-mail addresses: lessmann@econ.uni-hamburg.de (S. Lessmann), caserta@econ.uni-hamburg.de (M. Caserta), imontalvo@gmmf.upv.es (I.M. Arango).

¹ Tel.: +49 40 42838 4294; fax: +49 40 42838 5535.

² Tel.: +34 963879890; fax: +34 963877981.

Some work has been carried out to develop elaborate parameterization approaches. For example, adaptations of individual metaheuristics have been proposed to tune parameters in a self-adaptive fashion (see, e.g., Battiti, Brunato, & Mascia, 2009; Beyer & Schwefel, 2002). Whereas the potential of such reactive metaheuristics is undoubted, the objective of this paper is to lay ground for an orthogonal approach towards parameter tuning. In particular, we examine whether a formal relationship between effective parameter values and characteristics of the underlying optimization problem exist and whether it is sufficiently strong to be exploited for an automated tuning approach.

Most metaheuristics operate in an iterative manner: Given a starting solution to an optimization problem, neighboring solutions are evaluated and used in some way to improve the present one, i.e., find better solutions to the problem. Consequently, data concerning the assessed candidate solutions and their appropriateness naturally becomes available during the execution of a metaheuristic. The particular way in which a solution is altered or a new solution is obtained on the basis of a present one, respectively, depends upon settings of the metaheuristic's parameters. Hence, we consider a parameter setting to be *effective*, if it facilitates improving a candidate solution's fitness. Employing the data produced, we aim at constructing a prediction model capable of estimating suitable parameter values for a subsequent iteration by means of regression. It is argued that if such a model shows high forecasting accuracy, it could be employed as a *tuning agent* to predict effective parameter settings on the basis of the metaheuristic's current status and problem-specific information. This differs from previous approaches to achieve self-adaptivity in the sense that it does not depend upon any particular search algorithm.

Therefore, the paper contributes to the literature by proposing a data mining based solution to the problem of tuning metaheuristics' parameters and providing empirical evidence of its efficacy. To that end, a case study is undertaken to systematically explore the feasibility of forecasting effective parameter settings for one particular metaheuristic, the PSO algorithm. The study comprises established as well as state-of-the-art regression methods to scrutinize the nature of the relationship between the employed independent variables and effective parameter values (e.g., linear versus nonlinear) and identifies promising candidate models. Since embedding a forecasting model as tuning agent into PSO – or any other metaheuristic – would involve processing batches of data that become available throughout successive PSO iterations, a *learning curve analysis* is conducted to explore the forecasting models' sensitivity towards data size and simulate *online learning*. This complements the assessment of different candidate models and lays ground for future research to develop and assess a hybrid metaheuristic with integrated tuning agent.

The paper is organized as follows: Related work concerning the tuning of metaheuristics as well as hybrid procedures at the interface of metaheuristics and data mining is reviewed in the following section, before the proposed tuning approach together with some background on PSO and regression modeling is elaborated in Section 3. Section 4 explains the design of the empirical study; respective results are reported in Section 5. Limitations and opportunities for future research are discussed in Section 6, before conclusions are drawn in Section 7.

2. Related work

The paper proposes employing data-driven regression models to determine effective parameter settings for a particular metaheuristic, the PSO algorithm. Therefore, using the taxonomy of (Meisel & Mattfeld, 2010), the work can be described as an application of

“Data Mining (DM) to increase the effectiveness of Operational Research (OR) by refinement”.

Previous approaches at the interface of OR and DM have mainly considered a reverse approach. That is, OR techniques have been employed to support DM tasks by, e.g., automating some parts of the DM process. In particular, the use of metaheuristics in conjunction with predictive regression or classification models is widespread. First, building a prediction model inevitably involves an optimization step to, e.g., minimize some loss function of the model's forecasts and true outcomes over a set of example cases. Consequently, metaheuristics can be employed to perform this optimization, i.e., construct the forecasting model from empirical data. Examples include, e.g., the construction of classification models by means of genetic or ant algorithms (Martens et al., 2007; Sharpe & Glover, 1999) or the training of support vector machine (SVM) or neural network classifiers with evolutionary algorithms (Pendharkar, 2001; Pendharkar & Rodger, 2004; Stoean, Preuss, Stoean, El-Darzi, & Dumitrescu, 2008).

Furthermore, the ability of metaheuristics to solve complex combinatorial optimization tasks has facilitated modifications and extensions of existing data mining techniques and the development of novel algorithms, respectively, to better reflect the requirements of real-world decision making tasks. For example, (Bhattacharyya, 1999) proposes a customer targeting model for direct marketing applications. Contrary to standard approaches that first estimate customers' likelihood of responding by means of classification and subsequently decide upon the number of prospects to be contacted, this model explicitly considers the available marketing budget (i.e., how many customers can be solicited) to optimize the response rate among those customers who will eventually be included in the campaign. As a result, the model directly optimizes campaign profit, rather than response rate, and is in this respect more aligned with marketers' objective. However, incorporating such additional information complicates the construction of the forecasting model. Therefore, (Bhattacharyya, 1999) develops a genetic algorithm to perform this task. In a similar fashion, (Orsenigo & Vercellis, 2004) consider the SVM classifier, and argue that the optimization necessary to construct a SVM model exhibits inaccuracies. For mathematical convenience, classification errors are taken into account in an approximate way, whereas a discrete counting of errors is more aligned with the principles of statistical learning theory. However, building a discrete SVM requires solving a complex mixed-integer problem, which is accomplished by means of a tabu-search metaheuristic.

The most popular application of metaheuristics in predictive modeling is concerned with determining a model's *hyperparameters* (e.g., the number of hidden layers in a neural network) to adapt them to a particular forecasting task. This approach has received particular attention within in scope of SVM based classification (Chen & Hsiao, 2008; Friedrichs & Igel, 2005; Guo, Yang, Wu, Wang, & Liang, 2008; Hou & Li, 2009) and regression models (Fei, Liu, & Miao, 2009; Gijssberts, Metta, & Rothkrantz, 2010; Tang, Zhuang, & Jiang, 2009; Üstün, Melssen, Oudenhuijzen, & Buydens, 2005). Commonly, the model selection problem (i.e., hyperparameter tuning) is considered alongside the task of feature selection, which is concerned with discarding less informative independent variables to ease prediction. Due to the potential of metaheuristics to effectively identify high quality solutions in very large search spaces, they facilitate a joint optimization of model parameters and features. Respective approaches can be found in, e.g., Huang, Chen, and Wang (2006), Huang and Wang (2006), Lin, Shiue, Chen, and Cheng (2009), Min, Lee, and Han (2006), Wu, Tzeng, Goo, and Fang (2007). Dedicated feature selection algorithms that ground on the principles of heuristic search are developed in, e.g., Kim, Street, Russell, and Menczer (2005), Meiri and Zahavi (2006), Pacheco, Casado, and Núñez (2009), Tosun, Turhan, and Bener (2009).

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات