European Journal of Operational Research 000 (2014) 1–7

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Cross entropy for multiobjective combinatorial optimization problems with linear relaxations

Rafael Caballero[a], Alfredo G. Hernández-Díaz[b], Manuel Laguna[c], Julián Molina[a,*]

[a] Department of Applied Economics (Mathematics), University of Malaga, Malaga, Spain
[b] Department of Economics, Quantitative Methods and Economic History, Pablo de Olavide University, Seville, Spain
[c] Leeds School of Business, University of Colorado, Boulder, CO 80309-0419, USA

ARTICLE INFO

Article history:
Received 18 February 2013
Accepted 29 July 2014
Available online xxx

Keywords:
Multiobjective combinatorial optimization
Cross entropy
EMO
Linear relaxation

ABSTRACT

While the cross entropy methodology has been applied to a fair number of combinatorial optimization problems with a single objective, its adaptation to multiobjective optimization has been sporadic. We develop a multiobjective optimization cross entropy (MOCE) procedure for combinatorial optimization problems for which there is a linear relaxation (obtained by ignoring the integrality restrictions) that can be solved in polynomial time. The presence of a relaxation that can be solved with modest computational time is an important characteristic of the problems under consideration because our procedure is designed to exploit relaxed solutions. This is done with a strategy that divides the objective function space into areas and a mechanism that seeds these areas with relaxed solutions. Our main interest is to tackle problems whose solutions are represented by binary variables and whose relaxation is a linear program. Our tests with multiobjective knapsack problems and multiobjective assignment problems show the merit of the proposed procedure.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Cross entropy (CE) has a relatively short history in the realm of optimization methodologies. Rubinstein (1997) developed CE as a method for estimating probabilities of rare event in complex stochastic networks. Two years later, CE was applied for the first time in the context of combinatorial optimization, triggering a number of other operation research applications that are now well documented (de Boer, Kroese, Mannor, & Rubinstein, 2005a). A fairly comprehensive list of CE applications and tutorial materials is found in http://www.cemethod.org. In its most basic form, CE consists of the repeated execution of the following two steps (de Boer, Kroese, Mannor, & Rubinstein, 2005b):

1. Generate a random sample from a pre-specified probability distribution function.
2. Use the sample to modify the parameters of the probability distribution in order to produce a "better" sample in the next iteration.

Performance of CE implementations varies according to the rules (and related parameter values) used to solve the so-called *associated stochastic problem* (ASP), which is the problem of estimating the probability that the objective function of a vector of random variables that follow a probability distribution function with parameters v exceeds a given value γ. The problem is solved by generating a sequence of (v, γ) values that are adaptively updated from iteration to iteration. If successful, the search converges to a small neighborhood around or precisely to the optimal values denoted by (v^*, γ^*). The CE implementations in the literature are variants of this process, where most of the changes consist of the specific rules to obtain (v_{t+1}, γ_{t+1}) in iteration $t+1$ from the values (v_t, γ_t) in iteration t. The exception is, perhaps, the adaptation introduced by Laguna, Duarte, and Martí (2009), where the method is hybridized with the addition of local search. The main difference between CE and the so-called estimation of distribution algorithms (EDAs) is the assumption of independence of the decision variables. In its simplest form, CE attempts to estimate the probability that at optimality a variable takes on the value of 1, and it does so by treating each variable independently. This is can be a disadvantage if variables are strongly correlated, but, on the other hand, it makes CE very easy to implement and to adjust. Applications to single objective combinatorial optimization problems have shown that the methodology can be quite effective even when not directly addressing the presence of covariance.

* Corresponding author.
E-mail addresses: rafael.caballero@uma.es (R. Caballero), agarher@upo.es (A.G. Hernández-Díaz), laguna@colorado.edu (M. Laguna), julian.molina@uma.es (J. Molina).

http://dx.doi.org/10.1016/j.ejor.2014.07.046
0377-2217/© 2014 Elsevier B.V. All rights reserved.

Please cite this article as: R. Caballero et al., Cross entropy for multiobjective combinatorial optimization problems with linear relaxations, European Journal of Operational Research (2014), http://dx.doi.org/10.1016/j.ejor.2014.07.046

We are aware of only three CE applications in the context of multiobjective optimization. The first one is the work by Ünveren and Acan (2007) that tackles the problem of finding efficient frontiers for problems with multiple multimodal objective functions of continuous variables. They "introduce the notion of clustered non-dominated solutions on the Pareto front to adapt the probability distribution parameters" within the CE. In single-objective optimization problem, the updating of the (v, γ) values is based on the best (elite) solutions found in the random sample of the current iteration. In multiobjective optimization, however, nondominated solutions are found along the Pareto front, rendering the use of a common set of (v, γ) values impractical. This is why Ünveren and Acan (2007) divide the set of nondominated solutions into clusters and essentially execute one CE procedure in each cluster. The number of clusters is set to be the number of objective functions in the problem plus one. The updating of each set of (v, γ) values is cluster-dependent and therefore each CE is responsible for finding the best possible set of nondominated solutions in its assigned region of the Pareto front (represented by the cluster).

The CE method applied to each cluster developed by Ünveren and Acan (2007) is practically identical to the CE Algorithm for Continuous Optimization described by Kroese, Porotsky and Rubinstein (2006). In particular, the procedure starts with an initial set of values for the mean and variance of the decision variables. A sample is taken from Normal distributions with the appropriate parameters for each variable. The mean and variances are updated using the elite solutions in the sample. Finally, a small (epsilon) value is used to compare the maximum variance and determine whether or not all variable values in the sample have converged to their corresponding mean. The method is compared to eight existing methods from the literature using seven well-known optimization problems with multiple multimodal functions, six with two objective functions and one with three. The key search parameters are set to a sample of 5000 solutions (per CE) and 1000 iterations. This means that for a bi-objective problem, the procedure evaluates $5000 \times 1000 \times 3 = 15$ million solutions. A total of 20 million solutions are evaluated in problems with three objective functions. The authors conclude that their procedure "performs better than its competitors for most of the tests cases." Connor (2008) applies the procedure by Ünveren and Acan (2007) to a real-world problem with two objective functions. In this application, ten clusters instead of the recommended three are employed.

The second multiobjective CE method that we have found in the literature is due to Perelman, Ostfeld, and Salomons (2008). Within the field of water resource management, Perelman, Ostfeld and Salomons address the problem of designing a water distribution system. Essentially, the problem consists of choosing where to locate pipes in a water distribution network and also select the size of the pipes. The single-objective version of the problem attempts to find a design that minimizes the overall construction cost while satisfying minimum water-pressure requirements at demand-nodes in the network. This is a classical problem in civil engineering that dates back to the work by Schaake and Lai (1969) where a linear programming model was used to find an approximate solution to the New York Tunnels Water Distribution System. Since then, numerous procedures have been applied to the single-objective function problem, including metaheuristics (e.g., Lippai, Heaney, & Laguna, 1999).

The multiobjective optimization version consists of making the water-pressure requirements a second objective as opposed to a constraint. This results in two extreme points, the zero-cost solutions, where the demands are not satisfied, and the highest cost solution that corresponds to satisfying all the demands (at a minimum cost). The first point is obtained trivially. The second requires the solution of the single-objective problem where the demand requirements are treated as hard constraints (as explained above). The CE implementation by Perelman et al. (2008) for the multiobjective optimization problem uses binary variables to select a pipe size in each location. Therefore, corresponding to each location, there is a binary vector of m variables, where m is the number of available pipe diameters (including the "zero" or "do nothing" alternative). A value of 1 indicates the selection of a particular pipe diameter.

In order to identify the elite solutions that will serve as the basis for updating the CE parameters, the authors invoke the concept of ranking, introduced by Fonseca and Fleming (1995). The ranking method assigns a rank value of one to all nondominated solutions. Dominated solutions are assigned higher rank values corresponding to the number of solutions that dominate them. In this sense, the rank method incorporates density information corresponding to different regions of the Pareto front. The rank values are used to determine the set of elite solutions in the sample, which in turn become the basis for updating the (v, γ) values.

Bekker and Aldrich (2011) propose a multiobjective CE for continuous problems with characteristics that are similar to the one developed by Perelman et al. (2008) for discrete optimization. The ranking of solutions is employed as a mechanism for selecting the set of elite solutions from the sample to update the (v, γ) values. In particular, the elite solutions become the basis for constructing empirical probability distribution functions, one for each variable in the problem. The size of the sample is set between 30 and 50 times the number of objective functions in the problem. In the computational experiments with eight standard problems from the literature the number of function evaluations is in the order of 10,000, resulting in high-quality fronts and computational times not exceeding 25 seconds. No comparisons against other methods are provided but convergence and dispersion metrics are used to assess the quality of the solutions. One of the main differences between Bekker and Aldrich (2011) and Perelman et al. (2008) is the selection of the "elite" solutions from the sample. While Perelman et al. (2008) use a fixed size throughout the search, Bekker and Aldrich (2011) use a size that grows over time and that is controlled by the ranking value (e.g., all solutions with ranking less than or equal to two may be selected).

2. Multiobjective combinatorial optimization (MOCO) problems

The MOCO problems that we are interested in solving have the following form:

$$\begin{array}{ll} \text{Maximize} & f_1(x), f_2(x), \ldots, f_m \\ \text{Subject to} & x \in \mathcal{X} \end{array}$$

where $\mathcal{X}$ represents the set of all feasible solutions and may constraint x to take on integer values. As customary in multiobjective optimization, solution x is said to dominate solution y if $f_i(x) \geq f_i(y)$ for all i and $f_i(x) > f_i(y)$ for at least one i. We focus on two specific MOCO problems in order to test our methodology: the multiobjective knapsack problem and the multiobjective assignment problem.

The multiobjective knapsack problem (MOKP) consists of selecting a subset of items from a set of n items in order to maximize the utility (profit) of m different knapsacks without violating their individual capacities. The profit for knapsack i associated with selecting item j is given by p_{ij} and the corresponding weight is denoted by w_{ij}. The capacity of the ith knapsack is represented by c_i. The problem may be formulated as follows:

$$\begin{array}{ll} \text{Maximize} & f_i(x) = \sum_{j=1}^{n} p_{ij} x_j \quad i = 1, \ldots, m \\ \text{Subject to} & \sum_{j=1}^{n} w_{ij} x_j \leq c_i \quad i = 1, \ldots, m \\ & x \in \{0,1\} \end{array}$$

Download English Version:

https://daneshyari.com/en/article/6896767

Download Persian Version:

https://daneshyari.com/article/6896767

Daneshyari.com