Computers & Industrial Engineering 111 (2017) 29–38

Contents lists available at ScienceDirect

Computers & Industrial Engineering

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

Partitioning methods for pruning the Pareto set with application to multiobjective allocation of a cross-trained workforce

Michael J. Brusco

Department of Business Analytics, Information Systems, and Supply Chain, College of Business, Florida State University, Tallahassee, FL 32306-1110, United States

ARTICLE INFO

Article history:
Received 1 December 2016
Received in revised form 10 May 2017
Accepted 29 June 2017
Available online 1 July 2017

Keywords:
Partitioning
Pruned Pareto sets
Multiobjective programming
Workforce allocation

ABSTRACT

The Pareto (or nondominated set) for a multiobjective optimization problem is often of nontrivial size, and the decision maker may have a difficult time establishing objective criterion weights to select a solution. In light of these issues, clustering or partitioning methods can be of considerable value for pruning the Pareto set and limiting the decision to a few choice exemplars. A three-stage approach is proposed. In stage one, a variance-to-range measure is used to normalize the criterion function values. In stage two, maximum split partitioning and p-median partitioning are each applied to the normalized measures, thus producing two partitions of the Pareto set and two sets of exemplars. Finally, in stage three, the union of the exemplars obtained by the two partitioning methods is accepted as the final set of exemplars. The partitioning methods are compared within the context of multiobjective allocation of a cross-trained workforce to achieve both operational and human resource objectives.

1. Introduction

The Pareto set of a multiobjective optimization problem consists of all *nondominated* solutions. Denoting ψ as a solution to a multiobjective problem, P as the number of objective criteria, $f_p(\psi)$ as the objective function value of criterion p (for $1 \leq p \leq P$), and assuming that each criterion should be maximized, a solution ψ' dominates ψ if $[[f_p(\psi') \geq f_p(\psi)\ \forall\ 1 \leq p \leq P] \wedge [f_p(\psi') \neq f_p(\psi) \forall\ 1 \leq p \leq P]]$. In other words, for ψ' to dominate ψ, the criterion values for ψ' must be as good or better than those of ψ for each criterion, and strictly better for at least one criterion. If ψ is not dominated by any $\psi' \neq \psi$, then it is nondominated and a member of the Pareto set.

The number of solutions in the Pareto set of a multiobjective programming problem can be large and, therefore, difficult for decision makers to process. In such cases, it is desirable to *prune* the Pareto set to extract several solutions as *exemplars* for further consideration. During the past 10 years, there has been a resurgence of interest in the development of methods for pruning Pareto sets (Aguirre, Taboada, Coit, & Wattanapongsakorn, 2011; Eusebio, Figueira, & Ehrgott, 2014; Guo, Wong, Li, & Ren, 2013; Jornada & Leon, 2016; Kulturel-Konak, Coit, & Baheranwala, 2008; Taboada & Coit, 2007, 2008; Vaz, Paquete, Fonseca, Klamroth, & Stiglmayr, 2015). Jornada and Leon (2016) classify these Pareto set reduction (PSR) methods as primarily falling into one of two categories: (1) ranking methods and (2) clustering methods. Taboada and Coit (2007, 2008) have indicated that clustering methods are perhaps more appropriate when decision makers do not have an a priori idea for a suitable set of decision weights. We focus on clustering methods for PSR throughout the remainder of this paper.

The use of clustering methods for PSR dates back (at least) to the work of Morse (1980), Steuer and Harris (1980), and Törn (1980). Morse (1980), for example, evaluated several hierarchical clustering algorithms, as well as block clustering methods for the Pareto set. More recently, Taboada and Coit (2007, 2008) applied K-means partitioning to the standardized (or normalized) objective function values of multiobjective optimization problems. Aguirre et al. (2011) described the use of tree-based clustering methods for PSR. Vaz et al. (2015) presented an approach that used three location-based models for PSR in a biobjective context.

Our approach in this paper most closely corresponds to the work of Taboada and Coit (2007, 2008) because we also partition standardized objective function values using K-means, as well as other partitioning methods. Succinctly, there are four dimensions to our contribution to the literature. First, we begin with a careful description of the process of standardizing the competing objective functions and obtaining a proximity representation among the solutions in the Pareto set. Second, we discuss four partitioning methods for PSR: (i) maximum split partitioning, (ii) minimum diameter partitioning, (iii) K-means, and (iv) p-median partitioning. Our discussion highlights important advantages and disadvantages of each of these methods. Third, we present a new

E-mail address: mbrusco@business.fsu.edu

http://dx.doi.org/10.1016/j.cie.2017.06.035
0360-8352/

three-stage partitioning process for PSR. Fourth, we present a comparison of alternative partitioning approaches within the framework of a multiobjective workforce allocation problem.

The first stage of the proposed partitioning approach for PSR uses a variance-to-range standardization procedure to normalize the criterion function values. This procedure has been shown to substantially outperform z-score and range transformations with respect to recovery of cluster structure (Steinley & Brusco, 2008). In the second stage, maximum split partitioning and p-median partitioning are applied to the normalized measures to establish two partitions and two sets of exemplars. In the third stage, the union of the exemplars obtained by the two partitioning methods is selected as the final set of exemplars. Our introduction of p-median partitioning for PSR is noteworthy because, relative to the other three methods, p-median partitioning has the distinct advantage of directly providing an exemplar for each cluster. Moreover, exact solution of p-median problems is more computationally feasible than exact solution of K-means problems. The decision to augment p-median partitioning with maximum split partitioning is predicated on the fact that, unlike the other three methods considered herein, split is a between-cluster measure. Therefore, it tends to be useful for isolating well-separated solutions on the Pareto frontier.

Although the partitioning method for PSR can be used for many different types of multiobjective optimization problems, we focus on applications pertaining to the allocation of a cross-trained workforce. Brusco (2015) recently developed an algorithm for generating the entire Pareto set for a biobjective workforce allocation problem associated with the competing objectives of service utility and assignment desirability. However, his method did not provide formal guidance for the decision maker in terms of choosing a solution from the Pareto set, which contained 100 or more solutions in some instances. In light of this problem, as well the fact that decision makers would tend not to have an appropriate set of decision weights in advance, a partitioning approach is apt to be particularly useful for PSR in this context.

Section 2 focuses on data partitioning for the Pareto set. This includes the pre-processing of the data, descriptions of the various methods, choosing the number of clusters, and choosing an exemplar from each cluster. Section 3 presents a brief review of the workforce allocation literature and defines the underlying multiobjective optimization problem. Section 4 provides a biobjective workforce allocation example to illustrate the results obtained from each of the methods in Section 2, and proposes the three-stage approach to capitalize on the strengths and mitigate the weaknesses of methods. An application to a triobjective problem is provided in Section 5. The paper concludes in Section 6 with a summary, discussion of limitations, and suggestions for future research.

2. Partitioning the Pareto set

2.1. Standardizing the data to be clustered

A description of a data clustering approach for the Pareto set begins with the definition of $\mathbf{X} = [x_{ip}]$ as an $N \times P$ matrix of objective function values, where N is the number of solutions in the Pareto set, indexed $S = \{1, 2, \ldots, N\}$, and P is the number of objective criteria. Although clustering methods could be applied directly to $\mathbf{X}$, this is generally not advised. If two or more of the objective criteria are measured on markedly different scales, then those criteria with larger measurements are apt to dominate the clustering solution, effectively rendering the criteria with smaller measures as superfluous to the clustering process. For this reason, the importance of standardizing clustering variables to assure proper recovery of cluster structure is well recognized in the clustering literature (Steinley & Brusco, 2008), and standardization has also been advised when using data clustering to prune the Pareto set (Taboada & Coit, 2007, 2008). In our application, we standardize $\mathbf{X}$ using a measure developed by Steinley and Brusco (2008, pp. 83–84) that defines the relative clusterability of a variable based on it variance, $Var(p)$, conditioned by its range, $Range(p)$. The specific measure for each variable is:

$$\varsigma_p = \frac{[12 \times Var(p)]}{(Range(p))^2}, \quad \forall 1 \leqslant p \leqslant P, \tag{1}$$

and the relative clusterability index, $RC(p)$, for each variable is obtained as follows:

$$RC(p) = \frac{\varsigma_p}{\min_{1 \leqslant p \leqslant P}\{\varsigma_p\}}, \quad \forall 1 \leqslant p \leqslant P. \tag{2}$$

The $RC(p)$ measures are subsequently used to transform the columns of $\mathbf{X}$ using the process:

Step 1. Transform each column of $\mathbf{X}$ to z-scores by differencing each variable from its mean and dividing by the standard deviation. Define this $N \times P$ matrix as $\mathbf{Z} = [z_{ip}]$.
Step 2. Let $Range(z_p)$ indicate the range of the z-scores for variable p. Also, define $Range(z_{\min}) = [Range(z_p)$:$p$: $\varsigma_p = \min_{1 \leqslant p \leqslant P}\{\varsigma_p\}]$.
Step 3. Compute the *variance-to-range* standardized data as the $N \times P$ matrix, $\mathbf{Y} = [y_{ip}]$, from the $RC(p)$ indices as follows:

$$y_{ip} = z_{ip}\sqrt{\frac{RC(p)[Range(z_{\min})]^2}{[Range(z_p)]^2}}, \quad \forall i \in S, \quad \forall 1 \leqslant p \leqslant P. \tag{3}$$

The next step is to establish an $N \times N$ dissimilarity matrix, $\mathbf{D} = [d_{ij}]$, among the N solutions in the Pareto set based on $\mathbf{Y}$. This is accomplished using squared Euclidean distance:

$$d_{ij} = \sum_{p=1}^{P}(y_{ip} - y_{jp})^2, \quad \forall i \in S, \quad \forall j \in S. \tag{4}$$

Squared Euclidean distance was selected because it is the foundation for popular clustering methods, namely Ward's (1963) method and K-means (Forgy, 1965; Steinhaus, 1956). Although this is the measure of dissimilarity we selected to implement herein, alternative measures such as Euclidean (not squared), Manhattan, or Mahalanobis distance would also be viable.

2.2. Selecting a partitioning method

The next step is to partition the Pareto set based on the information in $\mathbf{D}$. Two desirable properties of a partition are: (i) that each pair of clusters is well-separated, and (ii) that the solutions within each cluster are homogeneous. Although there are many possibilities for constructing partitions, we limit our discussion to four methods: (1) maximum-split partitioning (MSP), (2) minimum-diameter partitioning (MDP), (3) K-means, and (4) p-median partitioning.

2.2.1. Maximum-split partitioning (MSP)

The goal of MSP is to partition the indices of the Pareto set into K subsets $(S_1, S_2, \ldots, S_K)$ such that the smallest pairwise dissimilarity measure between two solutions that are not in the same subset is maximized. Mathematically, the MSP can be stated as follows:

$$\max : \Gamma_1 = \min_{1 \leqslant k < l \leqslant K}\left[\min_{(i \in S_k, j \in S_l)}(d_{ij})\right], \tag{5}$$