ARTICLE IN PRESS
JID: EOR [m5G;March 14, 2018;16:33]
European Journal of Operational Research 000 (2018) 1–15
Contents lists available at ScienceDirect
European Journal of Operational Research
journal homepage: www.elsevier.com/locate/ejor

Continuous Optimization

Trade-off preservation in inverse multi-objective convex optimization

Timothy C. Y. Chan[a], Taewoo Lee[b,*]

[a] Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, ON M5S 3G8, Canada
[b] Industrial Engineering, University of Houston, 4722 Calhoun Road, Houston, TX 77204, USA

ARTICLE INFO

Article history:
Received 21 June 2017
Accepted 20 February 2018
Available online xxx

Keywords:
Multiple objective programming
Convex programming
Inverse optimization
OR in cancer therapy

ABSTRACT

Given an input solution that may not be Pareto optimal, we present a new inverse optimization methodology for multi-objective convex optimization that determines a weight vector producing a weakly Pareto optimal solution that preserves the decision maker's trade-off intention encoded in the input solution. We introduce a notion of trade-off preservation, which we use as a measure of similarity for approximating the input solution, and show its connection with minimizing an optimality gap. We propose a linear approximation to the inverse model and a successive linear programming algorithm that balance between trade-off preservation and computational efficiency, and show that our model encompasses many of the existing inverse optimization models from the literature. We demonstrate the proposed method using clinical data from prostate cancer radiation therapy.

1. Introduction

Given a feasible solution to an (forward) optimization problem, the inverse optimization problem aims to determine parameter values – typically objective function parameters – that make the given solution optimal. Classical inverse approaches leverage duality to ensure optimality of a given solution (e.g., Ahuja & Orlin, 2001; Iyengar & Kang, 2005; Lamperski & Schaefer, 2015; Schaefer, 2009; Zhang & Xu, 2010). However, these inverse models return a trivial solution (e.g., a coefficient vector of all zeros for inverse linear optimization) if the given feasible solution is not a candidate to be optimal for the forward problem.

In general, there is no guarantee that a given solution is exactly optimal for the assumed forward problem. An original solution might have been adjusted post-optimization for implementability reasons, or the assumed forward model is itself a simplification of a complex system where an observed solution is near-optimal, which the decision maker wishes to use in the future. When inverse problems involve such uncertainty around the model and data, it is important to determine an objective function that captures the intention of the decision maker who "implemented" the given solution so that it replicates the implemented solution as closely as possible in future decision making. Recent studies have generalized the classical inverse models to overcome the issue of given solutions not being exactly optimal. For example, Troutt, Sohn, and Brandyberry (2005), Troutt, Pang, and Hou (2006), Keshavarz, Wang, and Boyd (2011), Chow and Recker (2012), Chan, Craig, Lee, and Sharpe (2014), Bertsimas, Gupta, and Paschalidis (2015), Aswani, Shen, and Siddiq (2018), and Chan, Lee, and Terekhov (2018) developed approximate inverse optimization models that impute model parameters that make the given solutions minimally suboptimal. In this paper, we bring together the ideas of Keshavarz et al. (2011) and Chan et al. (2014) and develop a new inverse optimization model for multi-objective convex optimization where the given solutions may not be Pareto optimal.

In multi-objective optimization, decision making is typically based in the objective space, i.e., the space where the vector of objective values resides. Deciding between different solutions on the Pareto frontier in this space involves examining the trade-off in objective values. With a weighted objective function, as is common in convex multi-objective problems, solutions are generated by solving the forward problem with different nonzero weight vectors, which explicitly quantify the trade-offs in the objectives deemed acceptable by the decision maker. Conversely, without access to the weights, and only observing a solution on the Pareto frontier, it is possible using classical inverse optimization methods to reverse engineer the weight vector that generated the solution and therefore determine the decision maker's intention with respect to trade-offs. If the observed solution is not on the Pareto frontier, on the other hand, approximate inverse optimization methods may be applied, which return a weight vector that generates a solution on the Pareto frontier.

Whether such a weight vector truly reflects the trade-offs intended in the given solution can be characterized by the relationship between the given solution and the newly generated (weakly)

* Corresponding author.
E-mail addresses: tcychan@mie.utoronto.ca (T.C.Y. Chan), tlee6@uh.edu (T. Lee).

https://doi.org/10.1016/j.ejor.2018.02.045
0377-2217/

Pareto optimal solution. In multiobjective optimization, trade-offs are represented by how much a decision maker is willing to compromise on some objective values to achieve improvements in other objectives (Eskelinen & Miettinen, 2012). Formally, given two solutions, a trade-off between any given two objectives is measured by the ratio of changes in the objective values (Miettinen, 1999). In other words, when moving from one solution to another, if the ratio of the changes in the objective values equals 1 across all the objectives, the solutions are considered to share (hence *preserve*) the same trade-off preference (Chankong & Haimes, 1983; Eskelinen & Miettinen, 2012; van Haveren et al., 2017). By incorporating this definition of trade-off into the inverse optimization context, the inferred weight vector then should generate a solution on the Pareto frontier that adjusts the objective values of the given solution by the same amount across all the objectives. In this paper, we generalize the above definition of trade-off preservation by considering a ratio of *weighted* differences in the objective values of two solutions, and develop an inverse optimization model that infers a weight vector that preserves the trade-off encoded in a given solution. We show that under some special cases of the trade-off preservation definition, finding a trade-off-preserving weight vector leads to minimizing the optimality gap associated with the given solution. Our assumption is that although a given solution may not be Pareto optimal, it is a desirable solution that reflects the decision maker's implicit preferences, from which our inverse model seeks to infer the weights.

Our work generalizes the approach of Chan et al. (2014) by considering convex multi-objective optimization problems. Also, our notion of trade-off preservation is general enough to represent various ways to characterize trade-offs across multiple objectives. For example, our model can be specialized to the duality gap minimization approaches in Chan et al. (2014) and the single-objective approximate inverse convex model in Keshavarz et al. (2011). We show that existing inverse optimization models designed for single-objective optimization which could potentially be used for multi-objective optimization, e.g., Keshavarz et al. (2011), may not take into consideration trade-offs across multiple objectives encoded in the given solution. As in Keshavarz et al. (2011) and Chan et al. (2014), we assume that a set of objectives is pre-specified.

We provide geometric interpretation of inverse optimization concerning Pareto optimality by relating the inversion process to projection of a given solution to the Pareto surface, which in turn is related to the reference point method in the multiobjective optimization literature (Miettinen, 1999; Ogryczak & Kozlowski, 2011; Wierzbicki, 1986). Our results elucidate the connection between the two areas by bringing the concept of trade-off preservation into the framework of inverse optimization. A related paper proposed a method to quantify the relative importance of multiple objectives given a prioritized order of the objectives and a Pareto optimal solution (Breedveld, Storchi, & Heijmen, 2009). Our work can be considered an extension since it applies to non-Pareto optimal solutions. Our contributions are summarized below.

(1) We generalize previous inverse optimization approaches and develop a new inverse convex multi-objective optimization model that accommodates any input solution and determines a nonzero weight vector that *preserves the trade-off encoded in the input solution*. We introduce a notion of trade-off preservation that is generally applicable to multi-objective optimization and prove that some special cases of trade-off preservation are equivalent to the concept of duality gap minimization that has been widely used in the inverse optimization literature. We elucidate the new relationship between the existing inverse optimization and multiobjective optimization techniques.
(2) We propose an efficient linear approximation of the proposed inverse problem as well as a successive linear programming algorithm that bridges the exact and approximate methods. Based on the linear approximation, we show that our proposed inverse optimization method is general and encompasses many of the inverse models from the literature.
(3) We demonstrate the application of our inverse optimization model to a clinical treatment planning problem using real prostate cancer radiation therapy data. Assuming given treatments are clinically desirable and a set of objective functions is provided, we show that weights that preserve the trade-off encoded in the given objective values produce treatments that are similar to the original treatments clinically. We show that inverse models that are not trade-off-preserving may lead to treatments that deviate substantially from the original treatments and violate clinical acceptability criteria.

2. Background

We first define a canonical multi-objective convex optimization problem as the forward problem. Then, we briefly review the inverse optimization models from Iyengar and Kang (2005), Keshavarz et al. (2011), and Chan et al. (2014).

2.1. Forward optimization problem

Let $f_k : \mathbb{R}^n \to \mathbb{R}, k = 1, \ldots, K$ and $g_l : \mathbb{R}^n \to \mathbb{R}, l = 1, \ldots, L$ be convex functions. Let $\mathbf{x} \in \mathbb{R}^n, \mathbf{A} \in \mathbb{R}^{m\times n}$, and $\mathbf{b} \in \mathbb{R}^m$. We define the forward optimization problem (FOP) as

$$\text{FOP}(\boldsymbol{\alpha}): \quad \underset{\mathbf{x}}{\text{minimize}} \quad \sum_{k=1}^{K} \alpha_k f_k(\mathbf{x}) \tag{1a}$$

$$\text{subject to} \quad g_l(\mathbf{x}) \le 0, \quad l = 1, \ldots, L, \tag{1b}$$

$$\mathbf{Ax} = \mathbf{b}, \tag{1c}$$

where α_k is the weight for the kth objective function. Let $\mathbf{X}$ be the feasible region of (1). We assume $\boldsymbol{\alpha} \in \mathbb{R}^K_+ \backslash \{\mathbf{0}\}$, $f_k(\mathbf{x}) > 0, k = 1, \ldots, K$ for $\mathbf{x} \in \mathbf{X}$, and $\mathbf{A}$ has full rank. We also assume that Slater's condition holds (Boyd & Vandenberghe, 2004). We define $\Omega(\boldsymbol{\alpha})$ to be the set of optimal solutions to FOP($\boldsymbol{\alpha}$) and $\Omega := \bigcup_{\boldsymbol{\alpha}\in\mathbb{R}^K_+\backslash\{\mathbf{0}\}} \Omega(\boldsymbol{\alpha})$. A solution $\mathbf{x} \in \mathbf{X}$ is weakly Pareto optimal if there is no other $\mathbf{y} \in \mathbf{X}$ such that $f_k(\mathbf{y}) < f_k(\mathbf{x})$, for all $k = 1, \ldots, K$; a solution $\mathbf{x} \in \mathbf{X}$ is Pareto optimal if there is no other $\mathbf{y} \in \mathbf{X}$ satisfying $f_k(\mathbf{y}) \le f_k(\mathbf{x})$ for all $k = 1, \ldots, K$ with at least one k such that $f_k(\mathbf{y}) < f_k(\mathbf{x})$. It is known that for a convex multi-objective optimization problem, the set Ω consists of all weakly Pareto optimal solutions. For any $S \subseteq \mathbf{X}$, we write $\mathbf{f}(S) = \{(f_1(\mathbf{x}), \ldots, f_K(\mathbf{x})) \mid \mathbf{x} \in S\}$. We denote $\mathbf{f}(\mathbf{X})$ as the feasible region in the objective space and the set $\mathbf{f}(\Omega)$ as the Pareto surface.

2.2. Inverse conic optimization

We begin by illustrating the approach of Iyengar and Kang (2005) using our forward problem (1). Given K pre-specified objectives and a solution $\hat{\mathbf{x}} \in \mathbf{X}$, assumed to be a regular point (Bazaraa, Sherali, & Shetty, 2006), a weight vector that makes $\hat{\mathbf{x}}$ optimal can be found by solving the following problem:

$$\underset{\boldsymbol{\alpha},\boldsymbol{\sigma},\boldsymbol{\pi}}{\text{minimize}} \quad 0 \tag{2a}$$

$$\text{subject to} \quad \sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}}) + \sum_{l=1}^{L} \sigma_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \mathbf{A}'\boldsymbol{\pi} = \mathbf{0}, \tag{2b}$$