Computers & Industrial Engineering 76 (2014) 253–267

Contents lists available at ScienceDirect

Computers & Industrial Engineering

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

Incorporating response variability and estimation uncertainty into Pareto front optimization

Jessica L. Chapman [a,*], Lu Lu [b], Christine M. Anderson-Cook [c]

[a] Department of Mathematics, Computer Science, and Statistics, St. Lawrence University, United States
[b] Department of Mathematics and Statistics, University of South Florida, United States
[c] Statistical Science Group, Los Alamos National Laboratory, United States

ARTICLE INFO

Article history:
Received 21 November 2013
Received in revised form 28 July 2014
Accepted 30 July 2014
Available online 7 August 2014

Keywords:
Response surface
Multiple response optimization
Incorporating estimation uncertainty
Trade-offs
Graphical summaries

ABSTRACT

Pareto front optimization has been commonly used for balancing trade-offs between different estimated responses. Using maximum likelihood or least squares point estimates or the worst case confidence bound values of the response surface, it is straightforward to find preferred locations in the input factor space that simultaneously perform well for the various responses. A new approach is proposed that directly incorporates model parameter estimation uncertainty into the Pareto front optimization. This step-by-step approach provides more realistic information about variability in the estimated Pareto front and how it affects our decisions about the potential best input factor locations. The method is illustrated with a manufacturing example involving three responses and two input factors.

1. Introduction

When optimizing several estimated responses, the two-stage Pareto front approach (Lu, Anderson-Cook, & Robinson, 2011) to identify promising candidate solutions and then select a best match to user priorities can add structure and rigor to decision-making. Traditionally, Pareto front (PF) optimization approaches involving multiple estimated responses have focused on the maximum likelihood or least squares point estimates of the response at a given set of inputs (referred to as the "mean model" throughout this paper). However, uncertainty in the model parameter estimates can have an impact on which input factor combinations are identified as best. Since the responses are likely to have different natural variability in the operational space, the precision with which the parameters are estimated differs, making it difficult to anticipate effects on the Pareto front solutions. Naively treating the estimated response surfaces as fixed can lead to overconfidence in the conclusions and potentially sub-optimal input factor level choices which do not perform well when implemented in practice.

Costa, Espirito Santo, and Oliveira (2011) and Mattson and Messac (2005) propose visualization approaches to understand how uncertainty impacts the construction of the PF. Martins and Lambe (2013) provide a survey of design optimization architectures, while Yao, Chen, Luo, van Tooren, and Guo (2011) review strategies for uncertainty-based optimization. Hu and Youn (2011a, 2011b), Wei, Cui, and Chen (2008) and Chowdhury, Rao, and Prasad (2009) consider strategies for summarizing the impacts of uncertainty on complex systems and their reliability.

Chapman, Lu, and Anderson-Cook (2014) propose using the worst case bounds of prediction intervals as a simple way of incorporating uncertainty into the decision-making process. In this paper, we propose an alternative approach for quantifying and characterizing the impact of estimation uncertainty on solution selection. The uncertainty impacts both which solutions are located on the PF, as well as which solutions are best for the particular priorities of the study as measured by a desirability function with user-specified weightings of the different criteria.

To illustrate the proposed methodology, we consider the optimization of a chemical process described in Myers, Montgomery, and Anderson-Cook (2009), [p. 253] where three responses (y_1 = yield, y_2 = viscosity, y_3 = number-average molecular weight) are of interest. Two input variables (time, $\xi_1 \in [77\text{ min}, 93\text{ min}]$ and temperature, $\xi_2 \in [167\,F, 183\,F]$) can be adjusted to influence the responses. To estimate the relationships between inputs and responses, a 13-run central composite design (Myers et al., 2009, p. 297) for a circular coded region with maximum radius of $\sqrt{2}$ was run and data were collected for each response. After fitting quadratic response surface models and removing non-significant terms, the estimated mean models for each response are as follows:

* Corresponding author.

http://dx.doi.org/10.1016/j.cie.2014.07.028
0360-8352/

$$\widehat{y_1} = 79.94 + 0.995x_1 + 0.52x_2 + 0.25x_1x_2 - 1.38x_1^2 - 1.00x_2^2$$

$$\widehat{y_2} = 70.0 - 0.16x_1 - 0.95x_2 - 1.25x_1x_2 - 0.69x_1^2 - 6.69x_2^2$$

$$\widehat{y_3} = 3386.2 + 205.1x_1 + 177.4x_2$$

The goal of the optimization is to simultaneously maximize yield, y_1, and minimize both the molecular weight, y_3, and the distance from the viscosity to a target value of 65, $|y_2 - 65|$. The ideal solution is a combination of time and temperature that performs well for all three objectives. Since the three criteria cannot all achieve their optimum simultaneously, which location is selected depends on the relative importance that is placed on the different responses' performance.

Because there are objective and subjective aspects to selecting a best solution to an optimization problem, the PF approach in Lu et al. (2011) considers decision-making in distinct stages. Stage 1 is objective, since it removes all poor candidates that are strictly inferior to others. A solution is inferior if at least one solution exists that has all criterion values at least as good as the inferior solution and at least one that is strictly better. Eliminating these inferior choices is rational and simplifies subsequent steps by removing non-contenders from further consideration. The PF is comprised of all non-inferior solutions. Stage 2 is subjective as it considers how important good performance on the different criteria is to the decision-maker. It examines solutions on the PF and determines how well they match the priorities of the study. Clearly there are different ways to consider the subjective aspect of the decision-making. Our approach is to quantify the desirability of different options subject to different priorities, and then provide methods to explore the robustness of the solutions to changes in priorities. Graphical summaries of the different alternatives and how they compare can help guide the selection of which individual solution best suits the needs of the decision-maker and facilitate discussion with quantitative measures if several stakeholders have different priorities for the solution. Methods are adapted from Lu and Anderson-Cook (2012) and Lu, Anderson-Cook, and Robinson (2012).

The process for selecting a best overall solution is further complicated when the estimated responses have associated uncertainty, which suggests a range of plausible values for the model parameters that are consistent with the data observed. To capture this uncertainty, we use the estimated models to simulate a large number of response surfaces all consistent with the observed data. This collection of alternative solutions becomes the basis for examining the impact of estimation uncertainty on our conclusions. The overall goal of the selection process is to highlight a small number of combinations of input factor levels that give optimal performance for the responses of interest, subject to how we have chosen to prioritize them. To help with the discussion of the subjective Stage 2 when uncertainty is present, we have broken this stage into several sub-steps (2a–2c), each with distinct goals and customized graphical summaries. We now provide an overview of the different steps in the decision-making process, before illustrating the methods in detail with the example.

Step 0: Generate Alternate Response Surfaces Consistent with Data: The goal of this step is to generate a large number of sets of model parameter values that are consistent with the observed data. These values can then be used to obtain response surfaces representative of the plausible relationship between the inputs and the responses. These response surfaces serve as the basis for our understanding of the impact of estimation uncertainty. The new response surfaces each lead to different Pareto fronts, with the points on the front having different criteria values. Understanding which points are on the Pareto fronts more frequently and the likely range of criteria values can inform the decision process.

Step 1: Characterize the Pareto Front (objective): The goal of this step is to summarize the uncertainty associated with the PFs that is propagated from the estimation uncertainty in the individual responses. Using the PFs for each simulated surface, we summarize the frequency with which input factor combinations appear on the front. Those combinations which do not appear on the front frequently can be eliminated from further consideration. At the conclusion of this step, the decision-maker should see how the PF changes across the spectrum of anticipated response values as well as which locations are commonly chosen on the PF.

Step 2a: Identify Promising General Solutions (subjective): The goal of this step is to gain understanding about which locations are frequently selected as best for different weight combinations for the user-specified desirability function and scaling. To summarize this information, we combine the criteria into a single measure, identify how frequently different locations are best across all simulated response PFs and the entire set of weight combinations as well as how robust they are to different weightings of the criteria. Examining trade-offs between the criteria allows the decision-maker to understand how much compromise is needed on some of the responses to improve others. At the end of this step, the decision-makers should have improved understanding of which regions of the input factor space perform well for different weighting combinations as well as how frequently locations are best for some weighting.

Step 2b: Find Promising Solutions for More Focused Priorities (subjective): As the decision-makers narrow their search for a best solution to match their priorities, this step focuses on how frequently different solutions are identified as the best choice for a particular set of weights. Initially, weights for the desirability function can be partitioned into larger regions, and then subsequently a particular set of weightings of interest can be explored. For the selected range of weightings, we examine how frequently different solutions are selected as best. As different prioritizations of the criteria are considered, different locations in the input factor space are highlighted as common choices of best solutions. At the end of this step, the decision-maker has information about which locations are common choices for best for the specific priorities of the study.

Step 2c: Make Final Performance-Based Selection (subjective): Since the optimization of the product or process often necessitates selecting a single input factor combination from which to operate, this step guides the users to a final decision. Numerical and graphical summaries allow comparisons between individual solutions which inform the decision-makers of the relative merits of the available choices. Once commonly identified best solutions in the range of interest have been highlighted, evaluating and comparing their performance to the best available alternative for each weighting combination provides understanding about the merits of a solution. At the end of this step, the decision-makers should understand what choices are available and how they perform at optimizing the responses for the weightings of interest.

In the remainder of the paper, we describe the details of Steps 0 through 2c for the chemical process example. We illustrate how the numerical and graphical summaries in each of the steps can be used to identify promising candidates and eliminate non-contenders until a final solution is selected. The descriptive summaries also provide a quantitative means of justifying the choice. In Section 2, the simulation step is described. Section 3 discusses the objective Step 1 for characterizing the PF. Sections 4–6 describe the decision-making process of Steps 2a–2c that incorporates the user-specified desirability function for combining the measures and the priorities of the study as summarized by the weighing combinations of interest. Section 7 provides some discussion of extensions to the methods, while conclusions are given in Section 8.