Li, Yang, Xie, Jiaohong, Kim, Amy M., El-Basyouny, Karim.

Investigating trade-offs between optimal mobile photo enforcement programme plans.

AUTHOR POST PRINT VERSION

Li, Y., Xie, J., Kim, A. M., & El-Basyouny, K. (2019). Investigating trade-offs between optimal mobile photo enforcement programme plans. Journal of Multi-Criteria Decision Analysis, 26(1-2), 51-61. https://doi.org/10.1002/mcda.1658

Investigating Tradeoffs between Optimal Mobile Photo Enforcement Program Plans

Yang Li[1], Jiaohong Xie[2], Amy M. Kim[3]*, Karim El-Basyouny[3]

[1] Ph.D. Candidate, Department of Civil and Environmental Engineering, University of Alberta
[2] MITACS Undergraduate Student, Department of Civil and Environmental Engineering, University of Alberta
[3] Associate Professor, Department of Civil and Environmental Engineering, University of Alberta

*Corresponding author. Address: 6-269 Donadeo Innovation Centre for Engineering, 9211 - 116 Street NW, Edmonton, Alberta, Canada T6G 1H9. Tel.: +17804929203.

E-mail addresses: li18@ualberta.ca (Y. Li), jiaohongxie@gmail.com (J. Xie), amy.kim@ualberta.ca (A. Kim), karim.el-basyouny@ualberta.ca (K. El-Basyouny)

Abstract

Agencies that manage mobile photo enforcement (MPE) programs must decide where and when to send their limited resources to monitor compliance with speed limits. Usually, the goal is to select locations based on a number of concerns (i.e., high collision sites, high speed violation sites, school zones, etc.) which, in most cases, is conflicting. If certain locations are given more MPE resources, then by definition, other locations will receive less attention, and vice versa. This paper aims to provide insights about such MPE program tradeoffs. We present a systematic procedure for interpreting the results of a multi-objective MPE resource allocation problem. The procedure consists of three steps: 1) Pareto front (PF) generation, 2) front representation, and 3) tradeoff analysis. First, in generating a PF, we sequentially apply two well-known scalar optimization methods to obtain a comprehensive set of Pareto-optimal solutions. Second, the K-medoids clustering algorithm and the silhouette index are adopted to partition the generated PF into similar-sized clusters, in order to help MPE program agencies choose from a reduced set of solutions on the PF. Third, we use the response surface method to determine tradeoff patterns on the PF. The results of the front generation analysis showed that applying two optimization methods together resulted in a nearly complete PF with a relatively uniform and dense spread of solutions. Consequently, the identified set of solutions (i.e., 13,210 cases) was further partitioned into 12 clusters by silhouette index and K-medoids. With the aim of reducing decision fatigue for agencies, each cluster's representative solution is considered a possible MPE resource allocation candidate. The tradeoff analysis indicated how much one must sacrifice in the other objectives in order to increase attainment of one particular objective. Finally, the tradeoff rate and elasticity were used to explore the quantitative relationship between the considered objectives.

Keywords: mobile photo enforcement program planning, multi-objective optimization, tradeoff analysis, Pareto front analysis, resource allocation.

1. INTRODUCTION

Agencies that manage mobile photo enforcement (MPE) programs must decide where and when to send their limited resources for conducting speed enforcement. Locations include those with known traffic safety issues, those perceived by the public to be of concern, those with higher numbers of vulnerable road users. To this end, agencies may be aiming to achieve multiple, but often conflicting, goals through MPE. If certain (types of) locations are given more MPE resources, then by definition, other locations will receive less attention, and vice versa; both these resource allocation scenarios may be optimal in a multi-objective setting (Li, Kim, & El-Basyouny, 2016). For instance, managers may decide that greater MPE presence in school zones is warranted through September (i.e., usually start of the school term), which is achieved at the expense of presence at all other priority location types. The question is thus: what is the cost of achieving more school zone presence, in terms of reduced presence at other types of priority sites (high collision, high speed violation, etc.)?

This paper aims to provide insights about such MPE program tradeoffs (also referred to as elasticities in transportation planning), by exploring the resource allocations generated through multi-objective optimization. Resource allocation solutions from a multi-objective optimization problem make up a Pareto front (PF). In a PF, no solution is absolutely superior over any other; instead, in comparing two solutions, we observe tradeoffs between the (two or more) objectives. By generating and understanding the PF of MPE resource allocation solutions, agencies can understand what they are gaining (in one objective) but losing (in another) by choosing a particular resource allocation over another.

We present a systematic procedure for interpreting the results of a multi-objective MPE resource allocation problem (MPE-RAP). The procedure consists of three steps: 1) front generation, 2) front representation, and 3) tradeoff analysis. First, in generating a PF, we sequentially apply two well-known scalar optimization methods – weighted sum and epsilon constraint – to obtain a set of Pareto-optimal solutions approximating the PF of the MPE-RAP. The two methods can be used together to efficiently find a satisfactory solution set. Second, the K-medoids clustering algorithm is adopted to partition the generated PF into k similar-sized clusters. K-medoids was chosen because of its ease of implementation. In addition, it is well known for its efficiency in processing large amounts of data, so it was used to handle a large number of PF solutions quickly. In addition, K-medoids uses an existing solution on the PF to represent a cluster. In the third step we use the response surface method to determine tradeoff patterns on the PF. A quadratic polynomial model is estimated on the Pareto-optimal solutions (generated in step 1) to construct a continuous surface. From this, we can examine the tradeoffs between Pareto-optimal MPE resource allocation solutions.

We demonstrate our procedure using an example of the MPE program in the City of Edmonton, Canada, in September 2014. We aim to maximize enforcement coverage of high collision sites, high speed violation sites, and school zones (Li, Kim, & El-Basyouny, 2016). A set of Pareto-optimal solutions in the three-dimensional objective space were generated and further partitioned into clusters. Each cluster's representative solution (12 clusters in total) is considered a possible MPE resource allocation candidate, with the aim of reducing decision fatigue for MPE decision makers. Moreover, we take one cluster representative (candidate solution) as an example to explore the tradeoffs between the three objectives under consideration. By moving among solutions on the optimal tradeoff surface (estimated using a quadratic polynomial function), we are able to observe, for instance, how much collision and speeding site coverage would be sacrificed to achieve more enforcement presence in school zones.

This paper can help agencies managing MPE programs access and choose resource allocation strategies through a better understanding of the solutions generated, and the relationships (tradeoffs) between these solutions. It provides an evidence-based, methodologically sound, and ultimately traceable MPE resource allocation decision support system, in sharp contrast to existing MPE programs that rely on black box (i.e. qualitative, expert run) decision making. This paper adds to the literature on systematic methods of MPE resource allocation, in particular providing a method of better understanding the relationships between different allocation solutions.

2. LITERATURE REVIEW

A mobile photo enforcement MPE program requires radar and camera systems installed in vehicles to perform speed enforcement at various locations within a roadway network. MPE programs in Alberta, Canada, have six common enforcement goals: to provide presence at high collision sites, high speed violation sites, school zones, construction zones, high pedestrian volume sites, and sites with community speeding complaints (Li, Kim, El-Basyouny, & Li, 2016). Other MPE programs throughout the world have similar objectives.

To address multiple MPE deployment goals simultaneously, Kim et al. (2016) proposed a model for the MPE resource allocation problem MPE-RAP that combined multiple deployment goals together in a single weighted function. The proposed model measures the degree of achievement towards several deployment goals at a site, and combines these into a single numerical index for the site by the pre-determined weights assigned to each deployment goal. The authors select a pre-determined number of sites with the highest index rankings, to which a pre-set amount of enforcement resources are allocated randomly. Despite the model being easy to implement, in practice, it may be challenging for MPE program managers to specify appropriate weight values between the enforcement goals considered in the model.

The MPE-RAP was then further studied and solved in two stages: 1) generate candidate MPE resource allocation (operators & vehicles) plans for city neighborhoods while accounting for multiple goals (Li, Kim, & El-Basyouny, 2016), and 2) schedule neighborhood-level MPE resources for individual enforcement sites where the goals set in the 1st stage can be attained (Li, Kim, & El-Basyouny, 2017). In Stage 1, Li et al. (2016) constructed a multi-objective optimization model to simultaneously account for the multiple goals that an MPE program might aim for. These goals – previously only qualitatively defined by the Province of Alberta – had been quantified in previous work (Li, Kim, El-Basyouny, et al., 2016). The authors used an illustrative example of the City of Edmonton's MPE program operations from September 2014 to demonstrate the optimized MPE deployment plans identified by the proposed model. Three frequently addressed deployment priorities were considered: maximizing enforcement presence at 1) high collision sites, 2) high speed violation sites, and 3) school zones. These three deployment priorities were quantified into objectives using the measures of equivalent property-damage-only collision frequency per kilometer (EPK), speed violation indicator (SVI), and school zone density (SZD), respectively. The values of the three metrics were taken from three years of data (2012-2014) from the City of Edmonton. Correspondingly, the objective functions are: maximize 1) $\sum_{i=1}^{n} EPK_i \cdot x_i$, 2) $\sum_{i=1}^{n} SVI_i \cdot x_i$, and 3) $\sum_{i=1}^{n} SZD_i \cdot x_i$, where x_i is the number of enforcement shifts (i.e., a shift is a 10-hour duty span for officers to conduct speed enforcement) assigned to each neighborhood $i \in [1, \ldots, 388]$. The objective functions are subject to two constraints: 1) $\sum_i x_i$ equals the total number of shifts ($P = 458$) during the studied month and 2) there are minimum (L_i) and maximum (U_i) bounds on the number of times a neighborhood can be enforced (x_i). The generalized differential evolution 3 algorithm (GDE3) (Kukkonen & Lampinen, 2005) was used to solve the model. The algorithm yielded a set of 200 optimal solutions, called a Pareto front (PF) of the MPE-RAP example. The purpose of this model is to present MPE program managers with different optimal allocation solutions, and allow them to choose solutions from month to month that address their changing priorities.

In Stage 2, Li et al. (2017) developed a binary integer linear programming model to specify the daily sequence of enforcement shifts that are allocated to pre-determined enforcement sites within each neighborhood over one month. The model determines the shift sequence by minimizing the conflict between the shift assignment and the enforcement time halo effect. Enforcement time halo is the deterrent effect that an MPE program yields for a time period after its operation (Hauer, Ahlin, & Bowser, 1982), and therefore reducing shift assignment in the time halo period can achieve efficient resource utilization. In addition, the model assigns a neighborhood's shifts (x_i) to the pre-determined enforcement sites of that neighborhood based on the sites' weights in relation to the attainment of the desired goals set in Stage 1. An optimal solution of the PF generated for the MPE-RAP example in the 1st stage study was input to the scheduling model, which produced a diverse shift schedule for the City of Edmonton's MPE program operation in September 2014.

Although the MPE-RAP has been systematically addressed, program managers can face difficulties when presented with a Pareto front (PF) in the 1st stage of the proposed systematic approach. First, in many real-life multi-objective optimization problems, the PF can be very large or can even contain an infinite number of solutions; the greater the number of considered objectives, the larger the expected size is of the PF. It is therefore difficult to make a choice from a very large PF. Second, although each solution on the PF informs the value given to each objective, the exchange of the objective values between solutions is not directly revealed. This creates inconvenience for MPE program managers when they compare a large number of solutions and choose the desired tradeoff.

There are two main approaches to reducing the number of solutions to represent a PF: 1) define objective preferences and establish utility functions, and 2) cluster analysis. In the first approach, Taboada et al. (2007) proposed a non-numerical preferences ranking method, where the authors proposed a weighted utility function of objectives. The weights were based on decision makers' ranking of the importance of each objective. Pareto solutions were assessed by the utility function, and a subset of solutions having function values greater than a pre-defined pruning threshold can be identified. Branke et al. (2004) focus on the solutions in the center of the PF (referred to as knee solutions) in the absence of decision makers' preferences. Solutions are evaluated by either an angular measure or a marginal utility measure, where the solutions with highest measure values are the preferred knee solutions. Mattson et al. (2004) propose an insignificant tradeoff region where the difference between any two objective values is less than a user-specified threshold. Solutions that are positioned within the insignificant tradeoff region of reference solutions are removed from the PF. However, the above first-category approach requires multiple iterative calculations and most also require a-priori determinations and estimates of preferences between objectives. Conversely, clustering techniques (the second approach to generating a PF representation) do not require significant computational efforts and prior preference information. Incorporating a clustering procedure in analyzing Pareto results can be found in many studies (Morse, 1980; Rosenman & Gero, 1985; Zitzler & Thiele, 1999; Taboada & Coit, 2007; Taboada et al., 2007). Despite the use of different clustering algorithms, all studies group solutions on a PF into a pre-defined range of clusters consisting of similar solutions; only the solutions that represent the clusters are chosen to stand in for the entire PF.

Approaches to analyze tradeoffs among conflicting objectives are mainly focused on plotting results on two axes (with 2 objectives) or a hypersurface (3 or more objectives) in a discrete PF. The tradeoff between any two objective functions when moving from one solution to another along a PF is the slope of the line connecting the two solutions in the two-objective space (Miettinen, 1999). Hence, by connecting the solutions on the PF with smooth curves or surfaces, the objective tradeoff implied can be analyzed in an efficient manner in a PF with a large number of data points. For instance, Bai et al. (2011) used polynomial regression to generate pairwise tradeoff curves for five performance objectives considered in a highway asset management program. Goel et al. (2007) applied the response surface method to simultaneously analyze the tradeoffs of three goals related to a rocket injector design program. The authors constructed a polynomial model to build an (optimized) tradeoff surface for the three goals considered. However, tradeoffs were analyzed on a 2D contour map of the surface for simplicity. Note that the higher the objective dimension is, the more complex and difficult it is to interpret tradeoffs on a hypersurface. Therefore, when there are more than two objectives to consider, the easiest method is to perform pairwise comparisons of objective tradeoffs in 2D while keeping other dimensions constant (Bai et al., 2011).

Despite the rich literature on PF result analysis, there has been no application of this within MPE-RAP, which ideally benefits from PF tradeoff analysis. This paper is designed to help MPE agencies better understand and use the Pareto optimal set of resource allocation solutions obtained by multi-objective programming. First, to identify a representative subset of MPE optimal allocation solutions from a PF, we incorporate a clustering process that can be easily implemented without the user-specified preferences. Then, we use the response surface method to fit an optimal tradeoff surface that (typically) involves more than two enforcement objectives. Objective tradeoffs are evaluated in pairs, to provide easily-understood guidance to MPE program managers looking for resource allocation solutions.

Section 3 describes the methods to generate PF results for analysis. Section 4 explains how to cluster the PF results, and Section 5 presents the tradeoff analysis of the PF results. We close with concluding remarks in Section 6.

3. PARETO FRONT GENERATION

This paper uses the same City of Edmonton MPE-RAP example and data as previously introduced (in Section 2) by Li et al. (2016). Section 3.1 describes a method to generate a PF of the MPE-RAP example. Section 3.2 shows the generated PF results for the MPE-RAP example.

3.1 Pareto Front Generation Method

As the MPE-RAP example is large (recall that the variable vector is $n = 388$), the GDE3 algorithm used in previous work (Li, Kim, & El-Basyouny, 2016) is not efficient for searching a large number of Pareto-optimal solutions in a reasonable time. Therefore, we employ traditional scalar optimization techniques, which allow for a much lower computational time compared to evolutionary algorithms (Chiandussi, Codegone, Ferrero, & Varesio, 2012). The advantage of using an evolutionary algorithm is its ability to generate a representative subset of Pareto optimal solutions. From the representative solutions, decision makers can interactively choose answers based on their specific needs and preferences. However, this paper focuses on how to explore the relationships (tradeoffs) after these solutions are found. Therefore, considering the computation effort, we chose the scalar optimization method that can yield a PF faster than GDE3.

The weighted sum method (Miettinen, 1999), one of the most well-known and simplest scalar optimization techniques, is first employed to solve the MPE-RAP example. The formulation of the weighted sum method is shown in the following Problem P1.

Problem P1: Weighted Sum Method Formulation

As shown in Eqn. 1, the weighted sum method formulates the three-objective model of Li et al. (2016) as a single objective consisting of the weighted sum of the three individual objectives. Eqn. 2 normalizes the weights α_g, β_g, and γ_g assigned to each of the three metrics such that they sum to 1. The subscript g represents the algorithm iteration number (to a maximum of G). Eqns. 3 and 4 are the constraints from the original model on resources x_i (introduced in Section 2). Eqns. 1-4 are repeatedly evaluated for each g; each evaluation yields a Pareto-optimal solution.

$$max\ Z_g = \alpha_g \cdot \sum_{i=1}^{n} EPK_i \cdot x_i \ + \beta_g \cdot \sum_{i=1}^{n} SVI_i \cdot x_i \ + \gamma_g \cdot \sum_{i=1}^{n} SZD_i \cdot x_i \tag{1}$$

Subject to:

$$\alpha_g + \beta_g + \gamma_g = 1, \qquad \forall\ g \in [1, \dots, G] \tag{2}$$

$$\sum_{i=0}^{n} x_i = P \tag{3}$$

$$L_i \le x_i \le U_i, \qquad \forall\, i \in [1, \dots, n] \tag{4}$$

Note that the weighted sum method has a well-known drawback: it only searches for corner solutions in the feasible region of the weighted sum problem (Eqns. 1-4). Therefore, using various weight combinations is likely to also produce corner solutions (Branke, Deb, & Miettinen, 2008; Mavrotas, 2009). To identify intermediate (non-corner) solutions, we adopted another well-known scalar optimization approach, the ε-constraint method (Haimes, 1971). The ε-constraint method formulation for the MPE-RAP example is described in the following Problem 2.

Problem P2: ε-Constraint Method Formulation

The ε-constraint method described in Eqns. 5-8 optimizes EPK (equivalent property-damage-only collision frequency per kilometer) and transforms the remaining two measures (SVI, speed violation indicator, and SZD, school zone density) into inequality constraints that are greater than or equal to the pre-set values of ε_g^1 and ε_g^2. The choice to optimize one particular measure and set the others as constraints is arbitrary; we would expect any configuration to yield the same results because this three-objective problem is convex. By changing the ε values of Eqns. 6 and 7, the ε-constraint method is able to generate a different Pareto-optimal solution at every iteration (g).

$$\max_{x \in \Omega} Z_g = \sum_{i=1}^{n} EPK_i \cdot x_i \tag{5}$$

Subject to:

$$\sum_{i=1}^{n} SVI_i \cdot x_i \ \ge \varepsilon_g^1 \tag{6}$$

$$\sum_{i=1}^{n} SZD_i \cdot x_i \ge \varepsilon_g^2 \tag{7}$$

$$\Omega = \{x_i \,\Big|\, \sum_{i=0}^{n} x_i = P \ and\ L_i \le x_i \le U_i, \forall\, i \in [1, \dots, n]\} \tag{8}$$

The major disadvantage of the ε-constraint method is that it can be difficult to specify the values of ε_g^1 and ε_g^2 without knowing the bounds of objectives SVI and SZD for the PF

(Miettinen, 1999). However, the results from the weighted sum method in the previous step can be used to address this issue – we can define ε_g^1 and ε_g^2 values using the range of corresponding objective function values from the weighted sum solutions found in the previous step.

3.2 Results

To construct the weight combinations required for the weighted sum method (Eqns. 1-4), we first set each of α_g , β_g and γ_g to values at 0.05 increments between 0 and 1, to generate a total of 9,260 weight value combinations. Then, we normalized these weight values (i.e., such that they summed to 1) by dividing by the total weight of each combination. By removing the duplicated weight combinations and a zero-valued combination, a total of 7,758 different remaining weight combinations were input to Eqn. 1. The 7,758 optimizations of Eqns. 1-4 were implemented by CPLEX in the MATLAB environment on a PC with Intel Core i7-3770 CPU (3.4GHz) and 16GB RAM. A total of 244 unique solutions were found in 23 seconds.

Of the 244 weighted sum solutions, the objective function values for SVI and SZD are observed within the ranges of [208, 294] and [377, 1007], respectively. Therefore, these two ranges are used to limit the values of ε_g^1 and ε_g^2 used in Eqns. 6 and 7 of the ε-constraint method. We created 20,000 random numbers for ε_g^1 and ε_g^2 in a uniform sequence within the specified range. Eqns. 5-8 were implemented by the MATLAB CPLEX toolbox repeatedly at each of the 20,000 sets of ε_g^1 and ε_g^2 values. To build a dense PF, we limited the solution space of SVI and SZD examined per iteration to the neighborhood of (ε_g^1, ε_g^2). Specifically, each implementation is constrained in a search area where the SVI-axis step size was set to two and SZD-axis step size ten. The step size (of 2×10) accounts for 2% of the corresponding objective function interval; other sizes can be determined as needed. A total of 16,544 unique solutions were obtained in 27 seconds on the same PC described above.

The solutions found by the weighted sum and ε-constraint methods are then put together and compared with each other. Although various weights were used in both methods, 97% of weighted sum solutions and 22% of ε-constraint solutions have the same objective values as the solutions in the combined set. Therefore, we eliminated these repeated solutions, and a total of 243 weighted sum solutions (black circles in Fig. 1) and 12,967 ε-constraint solutions (grey points in Fig. 1) are finally considered in the PF.

Each point shown in Fig. 1 is the result of optimizing all metrics (represented on each of the three axes shown) simultaneously. The three red asterisks shown in Fig. 1 represent the extreme (corner) points of the PF; each corner point represents the maximization of one of the three objectives. They are generated from the weighted sum method using weight combinations (1, 0, 0), (0, 1, 0) and (0, 0, 1) for the measures (EPK, SVI, SZD). The values of these three extreme solutions maximizing EPK, SVI, and SZD are (4599, 239, 470), (3059, 294, 377), and (2449, 237, 1007) respectively.

Fig. 1 Pareto front identified for the MPE-RAP example.

As seen in Fig. 1, the 243 weighted sum solutions are not evenly distributed on the Pareto front despite the evenly spaced weights. This is because, in the weighted sum method, the relationship between the objective function weights and the objective function values of the Pareto solution (based on those weights) is nonlinear. Using geometry, Das & Dennis (1997) demonstrated that the weight used in a bi-objective weighted sum method to find a Pareto solution is the reciprocal of one minus the slope (i.e., the ratio of change between the two objective functions) of the PF at a given solution point. Thus, considering a uniform distribution of weights in Eqn. 1 is unlikely to result in uniformly distributed Pareto solutions.

In addition, the identified weighted sum solutions comprise only 3% of the utilized weight combinations. This demonstrates the drawback of the weighted sum method discussed earlier: the method ignores non-corner solutions, rendering the usage of a large portion of the weight combinations redundant (Branke et al., 2008; Mavrotas, 2009). However, the solutions generated by the ε-constraint method fill the empty spaces left by the weighted sum solutions on the PF, as illustrated in Fig. 1. Applying both solution methods results in a nearly complete PF with a relatively uniform and dense spread of solutions. We will use the solutions generated by using both methods as the basis for the following two post-Pareto analyses in Sections 4 and 5.

4. PARETO FRONT CLUSTERING

To be able to analyze the most important and salient features of the PF generated as per Fig. 1, we adopt the *K*-medoids algorithm (Kaufman & Rousseeuw, 1987) to group similar solutions into clusters and identify a representative solution for each cluster. *K*-medoids is a modification of the well-known *K*-means clustering algorithm (MacQueen & others, 1967), where existing data points are recognized as cluster centers (medoids) rather than creating new cluster centers. To use the existing Pareto-optimal solutions as candidates, we select *K*-medoids to conduct the clustering analysis.

Use of *K*-medoids requires a prior determination of how many data clusters should be created. A common tool for determining the optimal number of clusters is the silhouette index (Rousseeuw, 1987). The silhouette index evaluates the average distance, $a(i)$, between any data point i in cluster a and all other points in the same cluster. It also compares $a(i)$ with the average distance of the point i to all the points of a neighboring cluster b, $b(i)$. The silhouette index for point i is close to one if $b(i)$ is much larger than $a(i)$ (Rousseeuw, 1987) – meaning, cluster a points are more close to one another than points in cluster b. The optimal number of clusters is found by maximizing the average silhouette index for all data points.

We implement the R Package 'NbClust' (Charrad, Ghazzali, Boiteau, Niknafs, & Charrad, 2014) to compute the silhouette index for a pre-set range of clusters between 10 and 20 in the data set as illustrated in Fig. 1. The maximum silhouette index is observed when the number of clusters is set to 12; thus, we took 12 as the best cluster count for our data set. Then clustering is done by a *K*-medoids algorithm in MATLAB with these 12 clusters. The *K*-medoids algorithm identifies 12 medoids and partitions all other solutions around the 12 identified cluster medoids.

Note that the scales of the three metrics (axes) shown in Fig. 1 are not the same. This is likely to cause the Euclidean distance measure used in the computation of silhouette index and K-medoids clustering to be dominated by metrics with large values. Therefore, to avoid biases in results due to metric domination, we normalized metric values prior to the computation of silhouette index and K-medoids. The performance of a normalization method that takes the variable range (i.e. the difference between the minimum and maximum values of the variable) as the divisor has been proven to be superior over other normalization methods in cluster analysis (Milligan & Cooper, 1988); therefore, we selected this min-max normalization method to transform the objective vectors of the solutions in Fig. 1 into the range 0 to 1.

Fig. 2 shows the rescaled data with the clustering result. The crosses in Fig. 2 represent each of the 12 cluster medoids. Fig. 2 also differentiates solution clusters by colors. Table 1 summarizes the descriptive statistics for the 12 clusters. The size of each cluster is given. The range of objective vectors for the solutions in each cluster and the objective vectors of the 12 medoids are also indicated in Table 1.

Fig. 2 Clustering analysis of the Pareto front of the MPE-RAP example.

From Table 1, one can observe an average of about 1,100 solutions per cluster. In each cluster, the objective function values of the solutions with respect to EPK, SVI, and SZD vary within ranges of 688, 28, and 190, respectively. It is observed that the average differences in the three objective function values of EPK, SVI, and SZD between the medoid and the farthest solution in the same cluster is 344, 14, and 95, respectively. These ranges are about half that of the ranges of the objective vectors within each cluster. Thus, we can conclude that the 12 medoids are located in the relative center of each cluster, and are a reasonable representation of their respective cluster. An agency managing an MPE program could take these medoids as the initial deployment candidate options.

Table 1 shows that candidate options (or medoids) #1, #6, #7, and #9 have relatively low objective function values on the SVI axis but high values on the EPK axis. Therefore, these solutions lie on the left-hand side of the PF as illustrated in Fig. 2. The average objective function values of these solutions in SVI and EPK are 231 and 3,987, which fall in the first third of the SVI-axis scale (208 to 294) and in the last third of the EPK-axis scale (2,432 to 4,599), respectively. By choosing these solutions, one gives priority to the EPK objective (enforcing high collision sites) while largely ignoring SVI (enforcing high speed violation sites).

Conversely, candidate solutions occupying the PF's right-hand side on Fig. 2, such as solutions #4, #5, and #12, show low objective function values for EPK but high values for SVI, with average values of 3,296 (two-fifths of the EPK scale) and 279 (four-fifths of the SVI scale), respectively. Therefore, the solutions on the right-hand side of Fig. 2 give enforcement attention to high speed violation sites, regardless of the locations with high collision frequencies.

Other solutions in the middle of the PF have relatively average values in both EPK and SVI objectives (3757 and 259), indicating a balance of the two deployment goals. Specifically, solution #3 (in the middle top of the PF in Fig. 2) shows the highest SZD objective value (957) among the 12 medoids. This solution assigns school zone enforcement the greatest priority, while maintaining relatively average enforcement intensity at high collision sites and high speed violation sites. In contrast, solution #2 (which lies in the middle bottom of the PF) presents the lowest SZD value (477) among all the medoids; therefore, it gives school zone enforcement the lowest priority of the three objectives. Solutions #8, #10, and #11 are in the relative center of the PF. Their average values in EPK, SVI, and SZD are 3,840, 259, and 714, which lie at the midpoints of the corresponding axis' intervals. These types of solutions represent a balance of the three conflicting enforcement deployment goals; they almost reach the optimal value for each. When MPE managing agencies have no preference among the three enforcement priorities, these solutions may be of most interest.

5. PARETO FRONT TRADEOFF ANALYSIS

After making an initial selection from the clustering result, if MPE agencies want to move from the initial selection to another solution on the PF that better suits their requirements, it would be helpful to understand what the tradeoffs are (with respect to the three objectives) in moving to another solution on the PF. In other words, if one wanted to increase attainment of one objective, how much would one need to sacrifice in the other two objectives to achieve this? In this section we present a function fit to the Pareto data points found in Section 3. Note that the data points discussed in this section are not normalized. The (continuous) PF fitting function can be used to quantitatively evaluate the tradeoffs between the deployment objectives as one moves between (discrete) solutions on the PF.

5.1 Pareto Front Fitting Function

A (continuous) polynomial function is estimated on the discrete multi-objective optimization solutions comprising the PF of Fig. 1. A polynomial functional form is chosen because of its simple implementation and its ability to approximate the true PF (Fang, Rais-Rohani, Liu, & Horstemeyer, 2005; Goel et al., 2007). However, other fitting techniques, such as exponential and translog functions, may also be suitable given different distributions of optimized solutions.

In a multi-objective problem such as the MPE-RAP example (with three objectives EPK, SVI, and SZD), one objective should be chosen as the dependent variable of the PF fitting function while the remaining objectives are independent variables (Goel et al., 2007). To facilitate this decision, we created three quadratic polynomial functions for the three possible variable configurations using the response surface method. The quadratic polynomial is one of the most commonly used models for the response surface method, to describe the relationship between dependent and independent variables (Myers, Montgomery, & Anderson-Cook, 2016). The model is useful for generating a response surface that is reasonably close to the fitted data points (Box & Wilson, 1992), and such a model is easy to estimate and apply. Table 2 shows the R^2 values for each of these three functions, indicating the goodness of fit of each function to the Pareto data points shown in Fig. 1. Because polynomial regression is a special case of linear regression, in that it is linear in the coefficients on the independent variables, it is appropriate to use R^2 to determine model goodness-of-fit (Ostertagová, 2012).

In Table 2, the R^2 value of $SZD = f(EPK, SVI)$ is the highest (at 0.968) among all the three fitted functions. This suggests that the function taking the SZD metric as dependent variable is the best-fitting function for the MPE-RAP Pareto points generated, compared to the functions generated by the other two variable configurations. Furthermore, the R^2 value of this best-fitting function is close to one, suggesting that a quadratic polynomial function is an appropriate fit for the identified Pareto data points. Therefore, the function $SZD = f(EPK, SVI)$ is selected to represent the PF of the MPE-RAP example, and its estimated form is shown in Eqn. 9. All estimated parameters are statistically significant at the 95% confidence level.

$$SZD = -9.63e^{-5} \cdot EPK^2 + 0.001 \cdot EPK \cdot SVI + 0.149 \cdot EPK - 0.093 \cdot SVI^2 + 38.408 \cdot SVI - 3360.688 \quad (9)$$

Fig. 3 shows a plot of the PF fitting function of Eqn. 9 as a grey surface, and compares it against the set of Pareto-optimal solutions (first shown in Fig. 1) used to fit the function. It is observed that the function fits the plotted Pareto data points closely (as the R^2 value would indicate).

Fig. 3 Pareto-optimal solutions and the fitted Pareto surface, for the MPE-RAP example.

We observe a downward bend at the top of the fitted Pareto surface in Fig. 3(b). This bending is attributed to the fact that the graph of a quadratic polynomial is a parabola. Since the coefficients of the function terms with the highest degree in Eqn. 9 are negative, the function graph will always decrease exponentially at its edges. Therefore, it is important to note that the PF fitting function should only be used within the range of Pareto data point values taken to fit the function.

5.2 Illustrative Example of the Objective Tradeoff Analysis

Suppose Edmonton's MPE program manager (i.e. the managing agency) had chosen solution 3 from the 12 medoids identified in Fig. 2, for the month of September 2014. According to Table 1, among the 12 medoids, this solution reflects an elevated priority to have enforcement presence in school zones during the start of the school year, while also maintaining some coverage of high collision and high speed violation sites. This solution has the highest SZD value (at 957) of the 12 medoids, with relatively average values of EPK and SVI at 2,850 and 248, respectively. The three metric values of solution 3 represent an initial decision, which assigns 957, 2850, and 248 enforcement coverage units in school zones, high collision sites and high speed violation sites, respectively.

We assess the tradeoffs between each pair of the three metrics in solution 3, where the pairwise tradeoff results are benchmarked against set values of the 3rd metric. Fig. 4 illustrates the three cross sections of the PF fitting function along each of the three axes, viewed in solution 3. Curves in Fig. 4(a), (b), (c) are function contour lines at the solution 3's objective values ($SVI = 248$, $EPK = 2{,}850$, and $SZD = 957$). Hence, these curves depict how a change in one objective function value of solution 3 impacts the other function values.

Fig. 4 Contours of the Pareto fitting function at MPE-RAP example Solution 3.

Note that as discussed in Section 5.1, the polynomial function is only valid over the range of Pareto data points used to fit it. Therefore, based on the range of the found Pareto data points in the three axes (2,432-4,599 EPK, 208-294 SVI, 377-1,007 SZD), we found two endpoints on each curve, between which the curve is considered a valid description of objective tradeoff. The two endpoints of each curve are marked as crosses and labeled A and B, D and E, H and I, in Figs. 4(a), (b), and (c) respectively. The three components of the objective vectors for these endpoints are shown in Table 3.

The three plots in Fig. 4 show that as one objective decreases, the other two objectives increase. The average tradeoff rate (i.e., the slope of the curves) between SZD and EPK, SZD and SVI, and SVI and EPK is -0.2, -7.8, and -0.02, respectively. This means that for every one unit decrease in SZD (i.e., one less enforcement coverage unit in school zones), EPK increases by 0.2 (or, 0.2 more enforcement coverage units at high collision sites) when SVI is fixed at 248 (enforcement coverage units at high speed violation sites). However, if EPK remains at 2850, SVI increases by 7.8 for each reduced unit in SZD. Additionally, a one unit decrease in SVI leads to a 0.02 increase in EPK when $SZD = 957$. It is difficult for MPE program decision makers to intuitively interpret the tradeoff between more than two objectives. Therefore, these obtained pairwise tradeoff values provide useful information and support for multi-objective decision-making about MPE resource allocation. Specifically, decision makers can learn the result of changing a decision (choose a new PF solution), that is, when the expectations for the 3rd objective are met, how they adjust the resource allocation between the remaining two objectives.

As the ranges of the three objective values are different, the concept of curve elasticity is introduced to further understand how responsive (in a proportional manner) one objective is to a change in another objective. Elasticity is a measure that evaluates the proportional change of the abscissa divided by the proportional change of the ordinate. The metric on the abscissa is classified as being ordinate metric elastic if elasticity is greater than one, unit ordinate metric elastic if elasticity is equal to one, or ordinate metric inelastic if elasticity is less than one. The ordinate metric elasticity of the abscissa metric at a specific point (x^*, y^*) on the curve is expressed by Eqn. 10, which computes the reciprocal of the curve's derivative at that point multiplied by the ratio of y^* to x^*.

$$e = \frac{dx^*}{dy^*} \cdot \frac{y^*}{x^*} \tag{10}$$

As can be seen from Fig. 4, as abscissa values increase, the slopes of each of the three curves become steeper, indicating a continuously decreasing elasticity of each curve in the abscissa direction. Thus, by manipulating Eqn. 10 and the curve function (i.e., the PF fitting function holding one variable fixed), we found the location of the unit elastic point on each curve. Specifically, the unit-elastic points in Figs. 4(a), (b), (c) are marked as asterisks and labeled as C, F, G, respectively. Table 3 shows their objective function values. These unit-elastic points help divide the curve between two endpoints into two parts. The curve that lies to the left of the unit-elastic point is elastic, whereas the curve to the right side of the unit-elastic point is inelastic.

In Fig. 4(a) where SVI is fixed at 248, it is observed that EPK is elastic to the changes in SZD along the part AC of the curve where EPK is in the range of 2,432-3,561, and SZD is in the range of 836-978 as shown in Table 3. Conversely, on the curve between points C and B, where EPK lies between 3,561 and 4,599 and SZD is between 489 and 836, EPK is inelastic regardless of whether SZD changes. The average elasticities of curve segments AC and CB are -3.2 and -0.7, respectively. This indicates that EPK changes at 3.2 times the rate of SZD change on the AC curve segment, but the rate of change of EPK is reduced to 0.7 times the rate of change of SZD on the CB curve segment. Since solution 3 is positioned on the AC curve segment, reducing solution 3's SZD value by a small quantity, say 10%, may yield an approximately 32% increase in EPK. This tradeoff could be highly attractive to MPE program managers that are looking to reduce traffic collisions.

In Fig. 4(b), the unit-elastic point F also specifies the elastic and inelastic regions for the tradeoff of the SVI and SZD when EPK equals 2850. It is worth noting that the endpoint D and unit-elastic point F are very close together. As seen in Table 3, they have a difference of 5 and 18 respectively on the SVI and SZD axes, accounting for only 6% and 3% of the corresponding axis interval. Therefore, the curve portion DF may be negligible, leading to the fact that SVI is almost always inelastic in SZD. The reason that SVI almost always responds weakly to changes in SZD may be due to the fact that neighborhoods with relatively high school zone densities also experience more speed violations (in turn, due to lowered area speed limits). According to data from 2012 to 2014, Edmonton neighborhoods with the top 10% of school densities (the average SZD for these neighborhoods is 2.74) exhibited an average of 43% of speeding violations, 30% higher than the average of all city neighborhoods. Reducing MPE resource allocations to neighborhoods with more school zones may also have the secondary effect of reducing MPE resources to neighborhoods with high speed violations (because these neighborhoods are the same). It can be seen that solution 3 is on the FE curve segment with an average elasticity of -0.47. This suggests that a 10% decrease in SZD results in only about a 5% rise in SVI. Consequently, at solution 3, it may not be productive to reduce SZD value to gain the expected increase in SVI.

Contrary to the inelasticity of SVI and SZD, in Fig. 4(c) that sets SZD fixed, it is observed from the position of the unit-elastic point G (located to the right of the curve segment HI) that EPK is always SVI elastic. The average elasticity between points H and I on the curve is -6.8, indicating that a 10% drop in SVI will result in a 68% increase in EPK. This high elasticity implies the strong responsiveness of EPK to decreases in SVI. However, between solution 3 and endpoint I, the SVI value that can be traded-off for EPK is very limited, only 4 in SVI (as shown in Table 3).

Based on the tradeoff and elasticity results observed above, MPE program managers may want to locate a more suitable solution (other than solution 3) with the desired level for all objectives on the fitted Pareto surface. Suppose the manager decides to relax SZD by 100 (an approximately 10% decrease) and SVI by 4 (approximately 2% decrease) to improve EPK. By entering the (relaxed) values of SZD (854) and SVI (244) to the PF fitting function, we can obtain an EPK value of 3,522, which is an approximately 25% increase in EPK compared to that of solution 3 (we will refer this point on the fitted Pareto surface as solution X). This procedure can be continued until the program manager obtains a solution that best satisfies their needs. Then, one of the Pareto-optimal solutions (Fig. 1) closest to their preferred solution determined on the Pareto surface can be selected. A search of the nearest Pareto-optimal solution to solution X is conducted and the objective values of that Pareto-optimal solution for EPK, SVI, SZD are obtained (3477, 243, 839 respectively). The Pareto-optimal solution found nearest to X deviates from X by values of 45 (EPK), 1 (SVI), and 15 (SZD). Choosing this Pareto-optimal solution will generate a tradeoff that is slightly different from the tradeoff result estimated by the PF fitting function; however, the new tradeoff may not satisfy the desired tradeoff level. Therefore, it is recommended that in future research, we should further investigate how to identify the actual Pareto-optimal solution that best matches the solution found on the tradeoff surface.

6. CONCLUSIONS

Agencies that manage MPE programs must decide where and when to send their limited resources to monitor compliance with speed limits. Usually, the goal is to select locations based on multiple objectives (i.e., high collision sites, high speed violation sites, school zones, construction zones, high pedestrian volume sites, etc.) which, in most cases, is conflicting. If certain (types of) locations are given more MPE resources, then by definition, other locations will receive less attention, and vice versa. This paper aims to provide insights about such MPE program tradeoffs, by exploring the resource allocations generated through multi-objective optimization. We present a systematic procedure for interpreting the results of a multi-objective MPE-RAP. The procedure consists of three steps: 1) front generation, 2) front representation, and 3) tradeoff analysis. A case study is used to demonstrate the procedure while simultaneously optimizing three metrics: equivalent property-damage-only collision frequency per kilometer (EPK), speed violation indicator (SVI), and school zone density (SZD).

The procedure first generated an initial PF over the three-dimensional objective space by the weighted sum method. Then, based on the extent of the front obtained by the weighted sum method, the ε-constraint method was adopted to fill the vacant areas on PF that are not covered by the weighted sum solutions. The combination of the two scalar optimization methods were applicable to generate a PF with a large set of relatively uniform and dense distributed solutions over multi-dimensions in a time-efficient manner.

The second part of the analysis procedure employed K-medoids clustering to partition all the PF solutions generated in Part 1 of the procedure into 12 clusters. K-medoids clustering was used to choose 12 existing Pareto-optimal solutions as representative solutions for each cluster. Each of the 12 representative solutions was relatively located at the center of each cluster, representing about 1,100 solutions with an average objective value interval of 688, 28, and 190 in the three objectives EPK, SVI, and SZD, respectively. Based on these objective values (representing the clusters), the MPE agency can initially choose a solution from these representative solutions that suits their preference. Selecting representative solutions from each cluster greatly simplifies the challenging task of making decisions from a large Pareto-optimal set (recall that the initial PF was made out of 13,210 solutions).

The last part of the procedure created a tradeoff surface approximating the shape of the identified PF in Part 1. This surface was built using a quadratic polynomial regression, which models the relationship between the dependent variable (SZD) and the independent variables (EPK and SVI). The determination of the dependent and independent variables was done by examining which modeling formulation best fit the Pareto data points. The tradeoff analysis using the PF fitting function is illustrated by a case example, which studies the objective relationship at the representative solution for cluster 3. Specifically, the tradeoff relationship between any pair of the three considered objectives at cluster 3 is visualized by the contour curves of the function and is quantified by two measures: tradeoff rate and elasticity. Pairwise comparison results (contingent on the 3rd objective's set value) show that a one-unit reduction in SZD will result in an average increase of 0.2 units in EPK, or an average increase of 7.8 units in SVI, whereas reducing SVI by one unit can lead to an increase in EPK by 0.02. This means that one less enforcement coverage unit in school zones will compensate for 0.2 more enforcement coverage units at high collision sites or 7.8 more coverage units at high speed violation sites. Additionally, every one enforcement coverage unit reduction at high speed violation sites increases enforcement coverage at high collision sites by 0.02 units. Moreover, the elasticity measures show that EPK is almost always elastic in SVI with an average elasticity of -6.8, while SVI is almost SZD inelastic (-0.47 average elasticity). Additionally, when SZD is in the range of 836-978, EPK is SZD elastic; when SZD lies in the range 489-836, EPK is SZD inelastic. Specially, if SZD drops by 10% within its elastic region, EPK rises by 32%; however, the same percentage decrease in SZD will only increase EPK by 7% in the inelastic region. The ability to quantitatively assess pairwise tradeoffs among objectives, allows MPE agencies to understand the relationships between any pair of objective values, when moving along the PF, and how responsive a change in one objective is in relation to the other.

This paper presents a set of procedures that can lead to traceable and informed decisions on MPE deployment strategies. It specifically addresses a major limitation in current deployment plans by allowing an MPE agency to examine the effects of often conflicting enforcement objectives and quantitatively assessing their tradeoffs. Future work may include the study of more objectives (only three were considered in this paper), such as prioritizing enforcement in construction zones, high pedestrian volume sites and sites with community speeding complaints. In addition, future research will explore methods to help decision makers choose the final solution by applying clustering techniques to the tradeoff analysis.

7. ACKNOWLEDGEMENT

The authors thank the Office of Traffic Safety at the City of Edmonton for providing data and financial support. Prof. Amy Kim would also like to thank the Mitacs Globalink Research Internship program for sponsoring Jiaohong Xie.

References

Bai, Q., Labi, S., & Sinha, K. C. (2011). Trade-off analysis for multiobjective optimization in transportation asset management by generating Pareto frontiers using extreme points nondominated sorting genetic algorithm II. *Journal of Transportation Engineering*, *138*(6), 798–808.

Box, G. E., & Wilson, K. B. (1992). On the experimental attainment of optimum conditions. In *Breakthroughs in statistics* (pp. 270–310). Springer.

Branke, J., Deb, K., Dierolf, H., Osswald, M., & others. (2004). Finding knees in multi-objective optimization. In *PPSN* (Vol. 3242, pp. 722–731). Retrieved from http://repository.ias.ac.in/83511/1/15-a.pdf

Branke, J., Deb, K., & Miettinen, K. (2008). *Multiobjective optimization: Interactive and evolutionary approaches* (Vol. 5252). Springer Science & Business Media. Retrieved from https://books.google.ca/books?hl=en&lr=&id=N-1hWMNUa2EC&oi=fnd&pg=PA1&dq=multi+objective+optimization+interactive+and+evolutionary+approaches&ots=eDBBYx0K9R&sig=3fr86o7wvF8JbPCQMl2biJvR90A

Charrad, M., Ghazzali, N., Boiteau, V., Niknafs, A., & Charrad, M. M. (2014). Package 'NbClust.' *J. Stat. Soft*, *61*, 1–36.

Chiandussi, G., Codegone, M., Ferrero, S., & Varesio, F. E. (2012). Comparison of multi-objective optimization methodologies for engineering applications. *Computers & Mathematics with Applications*, *63*(5), 912–942.

Das, I., & Dennis, J. E. (1997). A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems. *Structural and Multidisciplinary Optimization*, *14*(1), 63–69.

Fang, H., Rais-Rohani, M., Liu, Z., & Horstemeyer, M. F. (2005). A comparative study of metamodeling methods for multiobjective crashworthiness optimization. *Computers & Structures*, *83*(25), 2121–2136.

Goel, T., Vaidyanathan, R., Haftka, R. T., Shyy, W., Queipo, N. V., & Tucker, K. (2007). Response surface approximation of Pareto optimal front in multi-objective optimization. *Computer Methods in Applied Mechanics and Engineering*, *196*(4), 879–893.

Haimes, Y. Y. (1971). On a bicriterion formulation of the problems of integrated system identification and system optimization. *IEEE Transactions on Systems, Man, and Cybernetics*, *1*(3), 296–297.

Hauer, E., Ahlin, F. J., & Bowser, J. S. (1982). Speed enforcement and speed choice. *Accident Analysis & Prevention*, *14*(4), 267–278.

Kaufman, L., & Rousseeuw, P. (1987). *Clustering by means of medoids*. North-Holland. Retrieved from https://lirias.kuleuven.be/handle/123456789/426382

Kim, A. M., Wang, X., El-Basyouny, K., & Fu, Q. (2016). Operating a mobile photo radar enforcement program: A framework for site selection, resource allocation, scheduling, and evaluation. *Case Studies on Transport Policy*. Retrieved from http://www.sciencedirect.com/science/article/pii/S2213624X16300104

Kukkonen, S., & Lampinen, J. (2005). GDE3: The third evolution step of generalized differential evolution. In *Evolutionary Computation, 2005. The 2005 IEEE Congress on* (Vol. 1, pp. 443–450). IEEE.

Li, Y., Kim, A., & El-Basyouny, K. (2016). A multi-objective resource allocation model for a mobile photo enforcement (MPE) program. *Submitted for Journal Publication*.

Li, Y., Kim, A., & El-Basyouny, K. (2017). Scheduling resources in a mobile photo enforcement program. In *Transportation Information and Safety (ICTIS), 2017 4th International Conference on* (pp. 645–652). IEEE.

Li, Y., Kim, A. M., El-Basyouny, K., & Li, R. (2016). Using GIS to interpret automated speed enforcement guidelines and guide deployment decisions in mobile photo enforcement programs. *Transportation Research Part A: Policy and Practice*, *86*, 141–158.

MacQueen, J., & others. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability* (Vol. 1, pp. 281–297). Oakland, CA, USA. Retrieved from https://books.google.ca/books?hl=en&lr=&id=IC4Ku_7dBFUC&oi=fnd&pg=PA281&dq=some+methods+for+classification+and+analysis+of+multivariate+observations&ots=nOSgD0IftM&sig=a5ROjuhGomjpKTeWahT03YSAA08

Mattson, C. A., Mullur, A. A., & Messac, A. (2004). Smart Pareto filter: Obtaining a minimal representation of multiobjective design space. *Engineering Optimization*, *36*(6), 721–740.

Mavrotas, G. (2009). Effective implementation of the ε-constraint method in multi-objective mathematical programming problems. *Applied Mathematics and Computation*, *213*(2), 455–465.

Miettinen, K. (1999). *Nonlinear Multiobjective Optimization, volume 12 of International Series in Operations Research and Management Science*. Kluwer Academic Publishers, Dordrecht.

Milligan, G. W., & Cooper, M. C. (1988). A study of standardization of variables in cluster analysis. *Journal of Classification*, *5*(2), 181–204.

Morse, J. N. (1980). Reducing the size of the nondominated set: Pruning by clustering. *Computers & Operations Research*, *7*(1–2), 55–66.

Myers, R. H., Montgomery, D. C., & Anderson-Cook, C. M. (2016). *Response surface methodology: process and product optimization using designed experiments*. John Wiley & Sons. Retrieved from https://books.google.ca/books?hl=en&lr=&id=T-BbCwAAQBAJ&oi=fnd&pg=PR13&dq=R.H.+Myers,+D.C.+Montgomery,+Response+Surface+Methodology&ots=O1fePof63N&sig=nH6uxG4jwv0uQQm6Zzwy6Hf22_I

Ostertagová, E. (2012). Modelling using polynomial regression. *Procedia Engineering*, *48*, 500–506.

Rosenman, M. A., & Gero, J. S. (1985). Reducing the Pareto optimal set in multicriteria optimization (with applications to Pareto optimal dynamic programming). *Engineering Optimization*, *8*(3), 189–206.

Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, *20*, 53–65.

Taboada, H. A., Baheranwala, F., Coit, D. W., & Wattanapongsakorn, N. (2007). Practical solutions for multi-objective optimization: An application to system reliability design problems. *Reliability Engineering & System Safety*, *92*(3), 314–322.

Taboada, H. A., & Coit, D. W. (2007). Data clustering of solutions for multiple objective system reliability optimization problems. *Quality Technology & Quantitative Management*, *4*(2), 191–210.

Zitzler, E., & Thiele, L. (1999). Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. *IEEE Transactions on Evolutionary Computation*, *3*(4), 257–271.

Table 1 Statistical Summary of the Objective Vectors in the Twelve Partitioned Clusters

Clusters	No. of Solutions	EPK			SVI			SZD		
		Medoid	Min	Max	Medoid	Min	Max	Medoid	Min	Max
1	1175	3865	3555	4172	219	208	233	774	683	859
2	1381	4416	4128	4572	270	255	284	477	402	557
3	1028	2850	2432	3249	248	228	269	957	882	1007
4	762	3813	3228	4184	286	278	294	491	378	623
5	1008	3131	2440	3467	264	246	279	857	740	922
6	1268	4209	3943	4438	233	215	248	663	581	742
7	1186	3371	2928	3651	227	210	246	893	820	960
8	1001	3665	3225	3902	270	257	283	736	624	803
9	1502	4503	4282	4599	243	228	258	530	440	618
10	1201	3727	3351	4016	245	231	258	779	693	857
11	1313	4128	3877	4360	263	248	279	626	546	706
12	385	2944	2433	3463	286	279	294	601	377	785

Table 2 R-squared Values of the Quadratic Pareto Front Fitting Function, by Variable Configurations

	$\boldsymbol{EPK = f(SZD, SVI)}$	$\boldsymbol{SVI = f(EPK, SZD)}$	$\boldsymbol{SZD = f(EPK, SVI)}$
R^2	0.939	0.899	0.968

Table 3 Objective Vectors of the Endpoints, the Unit-Elastic Points, and the Illustrative Candidate Point on the Tradeoff Curves.

Point Labels	Endpoint, unit-elastic point, or illustrative point	EPK	SVI	SZD
A	Endpoint	2432	248	978
B	Endpoint	4599	248	489
C	Unit-Elastic Point	3561	248	836
D	Endpoint	2850	241	979
E	Endpoint	2850	294	564
F	Unit-Elastic Point	2850	246	961
H	Endpoint	2432	252	957
I	Endpoint	2956	244	957
G	Unit-Elastic Point	3081	238	957
Solution 3	Illustrative Point	2850	248	954