The Orienteering Problem under Uncertainty
Stochastic Programming and Robust Optimization compared

Lanah Evers[a,b,c,*], Kristiaan Glorie[c], Suzanne van der Ster[d], Ana Barros[a,b], Herman Monsuur[b]

[a] *TNO*
[b] *Netherlands Defence Academy*
[c] *Econometric Institute, Erasmus University Rotterdam*
[d] *Department of Operations Research, VU University Amsterdam*

Econometric Institute Report EI 2012-21

Abstract

The Orienteering Problem (OP) is a generalization of the well-known traveling salesman problem and has many interesting applications in logistics, tourism and defense. To reflect real-life situations, we focus on an uncertain variant of the OP. Two main approaches that deal with optimization under uncertainty are stochastic programming and robust optimization. We will explore the potentialities and bottlenecks of these two approaches applied to the uncertain OP. We will compare the known robust approach for the uncertain OP (the robust orienteering problem) to the new stochastic programming counterpart (the two-stage orienteering problem). The application of both approaches will be explored in terms of their suitability in practice.

Keywords: uncertain orienteering problem, stochastic programming, robust optimization

1. Introduction

A large class of real-life problems encountered for example in industrial and logistic settings can be translated into a mathematical model. In traditional mathematical models, the input parameters are assumed to be known beforehand. In reality however, some or all of these parameters might be uncertain. Therefore, the solution to the deterministic model might turn out to be infeasible or suboptimal in reality. To this end, models that incorporate uncertainty already in the modeling stage have gained increasing attention, especially during the last decade.

Two areas of research dealing with optimization under uncertainty are stochastic programming and robust optimization. Each approach has its own optimization objective and operates under its own assumptions about the data uncertainty. Therefore, for a spe-

*Corresponding author: Lanah.Evers@tno.nl

cific optimization problem and specific characteristics of the uncertain data, one approach might be preferred over the other.

In this paper we will focus on an uncertain variant of the Orienteering Problem (OP). The OP has a number of challenging applications in logistics, tourism and defense, as we will address later in this paper. The OP is a generalization of the Traveling Salesman Problem (TSP), which does not require all nodes to be visited and considers a profit associated with each node. The profit is collected only if a node is visited. The goal is to find a feasible tour of maximum profit, starting and ending at the depot. A tour is feasible if the total weight of the arcs that are included in the tour satisfies a capacity constraint. In practice, the weights may for example model travel costs, travel time or fuel consumption. The non-deterministic version of the OP under our consideration has uncertain weights on the arcs.

We will apply techniques from both stochastic programming and robust optimization to the uncertain OP. To this end, we introduce the Two-Stage Orienteering Problem (TSOP) by applying techniques from stochastic programming and we will compare this model to the known robust approach to the uncertain OP: the Robust Orienteering Problem (ROP) [1]. We will present the advantages and disadvantages of both approaches and we will discuss why specific characteristics of different applications of the OP might make either one of the approaches more suitable for these applications. Although a lot of research has been done on applying either stochastic programming or robust optimization to a specific application, research on comparing both approaches is scarce. To the best of our knowledge, no comparison between a stochastic programming approach and robust optimization approach to the OP has been made yet.

This paper is structured as follows. In the next section we will provide background on both stochastic programming and robust optimization and we will mention how our uncertain OP relates to the available modeling approaches. In Section 3 we will provide some background on the OP. We will formulate the nominal OP in which all parameters are assumed to be deterministic and we will discuss the literature on uncertain variants of the OP, among which the ROP. Next, we will introduce the TSOP in Section 4 and we will describe how we use Sample Average Approximation (SAA) to solve the TSOP. In Section 5 we conduct a case study where we use both the ROP approach and the TSOP approach to solve the same problem instance. We will compare the results of both models and we will discuss their suitability for two different applications of the OP in Section 6. Concluding remarks and directions for further research will be stated in Section 7.

2. Optimization under uncertainty

In this section we will provide some background on stochastic programming and robust optimization. We will first briefly summarize the main characteristics of both approaches and the way in which they differ. In stochastic programming the probability distribution of the uncertain parameters is assumed to be known or close to some estimated distribution. In robust optimization on the other hand, the data is assumed to be 'unknown but bounded'. With respect to the optimization objective, in stochastic programming the goal is to optimize the expected value of the objective function, based on the specific probability distribution of the uncertain parameters. In robust optimization the goal is to optimize the worst-case value of the objective function when considering all possible realizations of the uncertain parameters within a so-called uncertainty set.

The basic foundations of stochastic programming were proposed already in the 1950s by Dantzig [2], Beale [3] and Charnes and Cooper [4]. For an introduction to stochastic programming we refer to Birge and Louveaux [5] and more advanced theory is provided by Ruszczinski and Shapiro [6]. A bibliography on theory and applications is provided by van der Vlerk in [7].

The first step in the area of robust optimization was made in 1973 by Soyster [8]. One of the characteristics of the Soyster approach is that it might provide very conservative solutions, since resulting solutions can deal with all parameters reaching their worst-case value at the same time. However, the probability of realization of this worst-case scenario is very low in practice. In the late nineties, new robust optimization frameworks that deal with this issue were developed for Integer Programming (Kouvelis and Yu [9]) and Convex Programming (Ben-Tal and Nemirovski [10, 11] and El Ghaoui et al. [12, 13]).

Both stochastic programming and robust optimization provide models and tools for static as well as dynamic optimization problems under uncertainty. In a static uncertain optimization problem all decisions have to be made before the actual data reveal themselves. The main approach in stochastic programming to static problems is the use of chance constraints. Chance constraints define a lower bound on the probability that a constraint has to be satisfied, based on the predefined probability distributions. The idea of chance constraints was introduced by Charnes and Cooper [4, 14]. For theory, background and illustrative examples of the use of chance constraints we refer to Prékopa [15]. In static robust optimization adjusting the size and shape of the uncertainty set allows the decision maker to find a balance between the feasibility and the worst-case objective value that can be obtained for all realizations in the uncertainty set.

In a dynamic uncertain optimization problem however, the effect of prior decisions combined with realizations that are revealed later on, can be corrected by decisions (recourse actions) in a future stage. These problems are referred to as two-stage or multi-stage problems, depending on the number of stages in which the decisions can be taken. Stochastic programming deals with these problems by minimizing the sum of the first-stage cost and the expected cost in the future stages, based on the assumed probability distributions of the uncertain parameters. When no such strict assumptions on the uncertain parameters can be made, the use of robust optimization seems more appropriate. The first extension of robust optimization to dynamic settings was proposed by Ben-Tal et al. [16]. They show that two-stage robust linear programming is computationally intractable and propose a tractable alternative referred to as affinely adjustable robust linear programming. Affinely adjustable robustness requires the recourse decision variables to be an affine function of the realizations of the uncertain parameters. Later, related approaches to dynamic robust optimization have been proposed (see. e.g. [17, 18, 19]).

The amount of research focussing on comparing robust optimization and stochastic programming is limited. Regarding the class of static uncertain problems, both Ben-Tal and Nemirovski [20] and Bertsimas and Sim [21] derived bounds on constraint violation, based on the uncertainty set that was chosen in a robust optimization problem. These bounds can be related to the chance constraints that are used in stochastic programming. Since the assumptions on the distribution of the data uncertainty in robust optimization are less restrictive than the ones used in stochastic programming, the robust optimization solutions associated to the bounds on constraint violation, might turn out to be relatively conservative. A drawback of chance constraint modeling on the other hand, is that the associated models are generally intractable, which is often not the case for robust optimization models. Chen et al. [22] also address the relation between the bounds suggested in robust optimization and chance constraints. They extend the ideas proposed in [20] and [21] by defining new uncertainty sets that can capture asymmetry in the underlying random variables. Bertsimas and Thiele [23] present a framework of robust optimization techniques for both static and dynamic uncertain optimization problems. Throughout their tutorial they mention some relations to stochastic programming approaches.

The non-deterministic version of the OP that we consider in this paper has uncertain weights on the arcs. This problem was solved by means of robust optimization in [1], where the Robust Orienteering Problem (ROP) was introduced. The ROP finds a solution to a static problem: finding an initial tour. In applications of this uncertain OP however, the solution should be implemented in a dynamic environment. Due to its problem structure, it is not possible to apply aforementioned robust dynamic approaches to our non-deterministic OP. Therefore an alternative approach that extends the static solution approach with agility principles was suggested in [1]. In the remainder of this paper we will refer to this approach, including its dynamic extension as the 'Agile ROP'. In order to compare the Agile ROP to its stochastic dynamic counterpart, we introduce the Two-Stage Orienteering Problem (TSOP), which is a dynamic model to begin with.

3. Background on the orienteering problem

The OP was introduced by Tsiligirides [24] and was shown to be NP-Hard by Golden et al. [25]. An exact algorithm for the OP was proposed by Fischetti et al. [26] and a recent survey on the OP is given by Vansteenwegen et al. [27]. The OP is also known as the selective traveling salesperson problem [28, 29, 30], the maximum collection problem [31, 32], the time constrained traveling salesman problem [33] and the bank robber problem [34]. The general OP is defined as a path planning problem, but in most applications the aim is to find a tour. The selective traveling salesperson problem and the time constrained traveling salesman problem however, are always defined as a tour planning problem.

The OP belongs to the class of Traveling Salesman Problems with profits (TSPs with profits), of which an overview is given by Feillet et al. [35]. The TSP with profits consists of three problem variants: the OP, the Profitable Tour Problem (PTP) and the Price Collecting TSP (PCTSP). These variants differ in the way the weights and the profits are handled in the optimization problem. They are modeled either in the objective function or in a constraint. In the PTP both the weights and the profits are contained in the objective function (see [36]). The objective of the PCTSP however is to minimize the total weight of the tour, such that a constraint on a minimum profit level is met, while the OP deals with the weights and the profits exactly in the opposite way. In the original definition of the PCTSP by Balas [37], penalty terms are also added to the objective function for unvisited nodes. In most recent literature on the PCTSP however, the penalties are set to zero.

3.1. The nominal orienteering problem

Consider a set of nodes N and let $|N|$ be its cardinality. Denote the depot location by node $0 \notin N$. For notational convenience, define $N^+ = N \cup \{0\}$. To each node $i \in N$ we associate a profit value p_i. We formulate the OP on a complete graph $G = (N^+, A)$ with $|N| + 1$ vertices. To each arc $(i, j) \in A$ we associate a value f_{ij} representing the expected weight of arc (i, j). The total capacity of the arcs that can be selected in any tour is denoted by F. We introduce a binary decision variable x_{ij} for every arc $(i, j) \in A$. x_{ij} is set to 1 if arc (i, j) is used in the tour and 0 otherwise. An auxiliary variable u_i is introduced to denote the position of node i in the tour. The goal is to find a tour of maximum profit, that is feasible with respect to the capacity constraint, starting and ending at the depot. Based on these definitions, the formulation of the nominal OP is the following.

$$\text{(OP)} \qquad \max \sum_{i \in N} p_i \sum_{j \in N^+ \setminus \{i\}} x_{ij}, \tag{1}$$

subject to

$$\sum_{i \in N} x_{0i} = \sum_{i \in N} x_{i0} = 1, \tag{2}$$

$$\sum_{i \in N^+ \setminus \{k\}} x_{ik} = \sum_{i \in N^+ \setminus \{k\}} x_{ki} \leq 1 \qquad \forall k \in N, \tag{3}$$

$$\sum_{(i,j) \in A} f_{ij} x_{ij} \leq F, \tag{4}$$

$$u_i - u_j + 1 \leq (1 - x_{ij})|N| \qquad \forall i, j \in N, \tag{5}$$

$$1 \leq u_i \leq |N| \qquad \forall i \in N, \tag{6}$$

$$x_{ij} \in \{0, 1\} \qquad \forall (i, j) \in A. \tag{7}$$

Constraint (2) guarantees that the tour starts and ends at the depot. Constraints (3) are the flow conservation constraints and ensure that a node is visited at most once. Constraint (4) is the capacity constraint. Finally, Constraints (5) prevent the construction of subtours. This formulation is consistent with the one used in [1].

3.2. Uncertain orienteering problems

Although the deterministic OP has been well studied in the literature, research on non-deterministic variants of the OP is limited. This section reviews the literature that does exist on this topic. In all of these variants (part of) the uncertainty is in the weights on the arcs. The assumptions on the distribution of the uncertain weights however, differ among the variants. Some variants also contain additional uncertain parameters. Different problem formulations and solution approaches are used. All approaches aim at finding an initial tour by taking uncertainty into account.

Teng et al. [33] introduce and solve the Time-Constrained Traveling Salesman Problem with stochastic travel and service times (TCTSP). For tours that exceed the time constraint, a penalty is imposed, proportional to the amount in excess of the cost limit. The authors consider a two stage stochastic program with recourse and use an Integer-L-Shaped algorithm to solve this problem. The probability distributions are assumed to be discrete in the TCTSP.

A related problem is addressed by Tang and Miller-Hooks [38]. They introduce the Stochastic Selective Traveling Salesperson Problem (SSTSP) where the aim is to find a tour with a maximum objective value consisting of total reward minus total travel cost. The total duration of the tour has to satisfy a chance constraint, based on discrete probability distributions. Since travel costs also occur in the objective function, this variant is a generalization of the OP.

Campbell at al. [39] introduce yet another stochastic variant of the OP where a penalty is incurred for the nodes in the planned ordering that cannot be reached, which they call the Orienteering Problem with Stochastic Travel and Service Times (OPSTS). The objective is to maximize the total expected profit, consisting of both rewards from the nodes that can be reached and penalties for the nodes in the planned order that cannot be reached. Contrary to the TCTSP and the SSTSP, the OPSTS does not demand the solution to be a tour. The deadline is only considered up to the last node in the planned order (or tour) and therefore there is no constraint on the total weight of the entire tour. In [39] exact solutions for some specific instances are provided and a variable neighborhood search heuristic is used to solve the general case.

3.2.1. The agile robust orienteering problem

The Robust Orienteering Problem with its dynamic extension (the Agile ROP) was introduced by Evers et al. [1]. We will devote special attention to the Agile ROP, since this modeling strategy will be used to compare to the TSOP strategy. The goal of the Agile ROP is to find a maximum profit tour that remains feasible for all realizations of the uncertain parameters in a predefined uncertainty set. In the Agile ROP, no penalty costs are added to the problem instance. By varying the uncertainty set, a balance can be achieved between the probability that a tour remains feasible and the objective value of such a feasible tour. Based on the choice of the uncertainty set, problem specific theoretical bounds on constraint feasibility were derived. Before further addressing the Agile ROP we will briefly go into more detail regarding the robust optimization framework provided by Ben-Tal et al. [40].

Ben-Tal at al. define a solution to be 'robust feasible' if the solution remains feasible for all realizations of the uncertain parameters within a predefined uncertainty set. The uncertainty set can for example be defined using norms. Special cases of such uncertainty sets may lead to intuitively easy to understand requirements on the robust feasible solutions. The use of the L_∞-norm for example, where the uncertainty set has a certain size (say ρ), leads to a solution that remains feasible if all uncertain parameters differ at most ρ times their maximum deviation from their expected value.

Evers et al. [1] applied this framework by Ben-Tal at al. to introduce and analyze the ROP. Contrary to the nominal OP, in the ROP the weights of the arcs are assumed to be uncertain. The expected weight of arc (i, j) is denoted by $\overline{f_{ij}}$. Define a random variable $\zeta_{ij} \in [-1, 1]$. The realizations of the weights f_{ij} are assumed to lie within a certain predefined interval, such that

$$f_{ij} = \overline{f_{ij}} + \zeta_{ij}\sigma_{ij}. \tag{8}$$

Recall that in robust optimization an uncertainty set describes the set of realizations of the uncertain parameters against which we want the solution to be protected. The larger this uncertainty set is chosen, the larger the probability will be that the solution will remain feasible under the realizations of the uncertain parameters. But a large uncertainty set might on the other hand lead to a relatively conservative solution. As such, by varying the size and shape of the uncertainty set, a balance can be found between feasibility and the objective value that will be obtained in case the solution is feasible. In the ROP, the uncertainty set containing all realizations of the weights against which the solution should be protected is denoted by Z, which is defined as

$$Z = \{\zeta \in \mathbf{R}^{|A|} : \|\zeta\|_s \leq \rho_s \forall s \in S\} = \bigcap_{s \in S} B_s^{|A|}(\rho_s). \tag{9}$$

Thus, Z is allowed to be an arbitrary intersection of balls, where, to simplify notation,

$$B_s^n(\rho_s) = \{\zeta \in \mathbf{R}^n : \|\zeta\|_s \leq \rho_s\}. \tag{10}$$

A solution to the ROP is robust against the uncertainty in the weights of the arcs if it satisfies Constraint (4) for all realizations $\zeta \in Z$. That is, if it holds that

$$\sum_{(i,j) \in A} (\overline{f_{ij}} + \sigma_{ij}\zeta_{ij})x_{ij} \leq F \quad \forall \zeta \in Z. \tag{11}$$

Note that Equation (11) contains infinitely many constraints, but it is shown in [1] that the problem can be rewritten such that the associated robust counterpart only contains a finite number of constraints.

As just described, by the use of uncertainty sets the robust counterpart takes uncertainty into account already in the modeling stage. The resulting solution is an initial tour plan, i.e. a solution to a static problem. Note though, that at the execution phase the actual realizations of the weights will determine whether or not this tour can be completed. In other words, if the weights turn out to be higher than accounted for, the tour has to be aborted and potential profit of the unvisited nodes will not be obtained. In such a case, the sum of the profit of the nodes that cannot be visited is referred to as the *profit shortage*. On the other hand, when the realizations of the weights turn out to be relatively low, real-time information on the remaining capacity is exploited at the end of the tour to add nodes to the tour, if possible. The sum of the profit of the additional nodes that can be visited in such a case, is referred to as the *profit surplus*. To this end, it is discussed in [1] how the initial ROP tour can be executed based on agility principles. The agility principles describe policies on when to abort the tour and how to extend the tour for relatively high or low realizations of the weights, respectively. For details on how to define such agility principles we refer to the paper by Evers et al. [1]. A disadvantage of the Agile ROP though, is that it does not take the consequence of the realizations and these policies on the total obtained profit into account in advance. In constructing the initial tour, it only balances feasibility against the planned objective value. We will now discuss how stochastic programming can help overcome this issue.

4. The two-stage orienteering problem

In this section we will introduce the TSOP and we will suggest a solution method to solve the TSOP. Regarding the assumptions about the uncertain parameters, in the TSOP we assume that the weights of the arcs are stochastic and follow a certain probability distribution. Note that this assumption is more specific than the assumption about the uncertain parameters used in robust optimization. Contrary to what is done in the ROP, in the TSOP we aim to take into account the effect of the realizations of the weights on the profit value that will actually be obtained.

Consider a certain initial tour and an associated ordering of nodes. The further ahead a node is scheduled in the initial tour, the higher the probability will be that the node cannot be reached as a result of the uncertainty in the weights. During the execution of the tour, at some point a decision to return to the depot might be required, based on the actual realizations of the weights. In that case all profit of the remaining nodes is lost with respect to the profit value that was associated to the initial tour. By means of a two stage recourse model all possible moments of returning to the depot and the associated loss in profit can be taken into account. We will now go into more detail on how this idea is incorporated in the TSOP.

In the TSOP the fist stage decision is to find an initial tour. Then, before the second-stage we assume the realizations of the weights on the arcs are revealed. Based on these realizations we can easily determine until what point the initial tour can be executed. Of course, in reality the realizations will be revealed during the execution of the tour, but in terms of modeling this problem in two stages, the realizations are used in the second-stage to find the point at which the tour will be aborted during the actual execution. The following recourse action is applied: we will return to the depot as soon as the remaining capacity drops below the expected weight required to return from the current location to the depot. We assume that a certain amount of extra capacity is available to fulfill the function of safety stock. This amount should cover the maximum deviation from the expected weight on any of the arcs to the depot. Note that this safety stock is available to cover a potentially unbeneficial deviation from the expected value in the realization of the weight required to return to the depot, but it is not part of the capacity F which is used in the model. Associated to this recourse action, in the TSOP we define the recourse cost of returning to the depot as the sum of the profits of the nodes in the preplanned tour that cannot be visited. Like in the ROP, we will refer to this loss in profit with respect to the first-stage planned profit, as the *profit shortage*. The aim of the TSOP is then to find a first-stage tour such that the sum of the first-stage profit, corrected by the expected second-stage profit shortage for that specific tour, is maximized.

This way of approaching the problem has some similarity to the capacitated vehicle routing problem with stochastic demands, studied by Laporte et al. [41]. In the model by Laporte et al. a recourse action is required when the realizations of the stochastic demands of the customers that were visited up to customer j exceed the vehicle capacity. The recourse action is a return trip to the depot in order for the vehicle to (un)load, after which the initially planned tour is continued. Consequently, the recourse cost is modeled as the travel cost from customer j to the depot and back. Note that in the TSOP we do not continue the initial tour after returning to the depot.

4.1. Model

We will now give a formal problem definition of the TSOP. Consider the definition of all parameters and decision variables described by the nominal OP. Assume that the weight f_{ij} of arc (i, j) is a random variable which follows a certain distribution. For notational convenience, we will also use f_{ij} to denote the realizations of the weights. We denote the average value by $\overline{f_{ij}}$. Now, the following two-stage recourse model maximizes the expected profit, taking into account the possibilities of shortages with respect to the planned tour.

$$\text{(TSOP)} \qquad \max \sum_{i \in N} p_i \sum_{j \in N^+ \setminus \{i\}} x_{ij} + Q(x), \tag{12}$$

subject to Constraints (2) and (3) and Constraints (5) to (7), where the function $Q(x)$ expresses the recourse cost of decision x. $Q(x)$ denotes the expected value of the recourse actions over all realizations of f. The function $Q(x)$ can be specified as follows:

$$Q(x) = \mathbb{E}_f \left(v(x, f) \right), \tag{13}$$

where $v(x, f)$ is a function that expresses the profit shortage for a tour described by x and a realization of weights described by f. Note that this formulation will produce long initial tours, since a tour will not be optimal as long as nodes exist that have a positive probability of being reached, and since adding one of those nodes to the tour increases the objective value (12).

4.1.1. Linearization

One way to approach this two-stage optimization problem is by first finding a way to express v used in Formula (13) in terms of decision variables and linear relationships. Since in the second-stage x was already decided upon and the realization f is known, $v(x, f)$ should express the cost of nodes in the tour described by x that cannot be reached for that specific f. We can express these second-stage costs by means of an optimization problem. In order to do so, we introduce the following decision variables:

$$x_{ijk} = \begin{cases} 1, & \text{if arc } (i,j) \text{ is the } k\text{th arc in the tour described by the first-stage variable } x \\ 0, & \text{otherwise.} \end{cases}$$

$$z_k^f = \begin{cases} 1, & \text{if the } k\text{th target in the tour cannot be reached with repect to realization } f \\ 0, & \text{otherwise.} \end{cases}$$

$$y_i^f = \begin{cases} 1, & \text{if target } i \text{ was in the preplanned tour, but cannot be reached with repect to } f \\ 0, & \text{otherwise.} \end{cases}$$

The resulting second-stage can be formulated as the following optimization problem:

$$v(x,f) = \max\left(-\sum_{i\in N} p_i y_i^f\right), \tag{14}$$

subject to

$$x_{0j1} \geq x_{0j} \qquad \forall j \in N, \tag{15}$$

$$x_{ijk} \geq x_{ij} + \sum_{l\in N^+} x_{li(k-1)} - 1 \qquad \forall i,j \in N, k = 1,\ldots,|N|, \tag{16}$$

$$\left(\sum_{k=1}^{K}\sum_{(i,j)\in A} f_{ij}x_{ijk}\right) + \sum_{(i,j)\in A} \overline{f_{j0}} x_{ijK} \leq F + M z_K^f \qquad \forall K = 1,\ldots,|N|, \tag{17}$$

$$z_k^f \geq z_{k-1}^f \qquad \forall k = 2,\ldots,|N|, \tag{18}$$

$$y_j^f \geq \sum_{i\in N^+} x_{ijk} + z_k^f - 1 \qquad \forall j \in N, k = 1,\ldots,|N|, \tag{19}$$

$$x_{ijk} \in \{0,1\} \qquad \forall i,j \in N^+, k = 1,\ldots,|N|, \tag{20}$$

$$z_k^f \in \{0,1\} \qquad \forall k = 1,\ldots,|N|, \tag{21}$$

$$y_i^f \in \{0,1\} \qquad \forall i \in N. \tag{22}$$

The objective given by Formula (14) is to minimize the loss in profit with respect to the preplanned tour. Constraints (15) identify the first arc in the tour. The first-stage variable x_{0j} is equal to 1 for the arc that connects node j to the depot location 0. Since this is the first arc in the tour, x_{0j1} will be set to one. Constraints (16) identify the location of the other arcs in the tour. For arc (i,j) to be the kth arc in the tour, two conditions have to be satisfied: 1) arc (i,j) has to be part of the tour and 2) node i has to be the $(k-1)$th node in the tour. Thus, if both x_{ij} and $\sum_{l\in N^+} x_{li(k-1)}$ are equal to 1, this constraint will set x_{ijk} to 1. Constraints (17) determine the number of nodes of the initial tour defined by x that can and cannot be reached, based on the realizations f_{ij}. These constraints model the previously defined recourse action: as long as the remaining capacity exceeds the expected weight required to return to the depot, the execution of the tour is continued. That is, all nodes up to the Kth node in the tour can be reached in case $\left(\sum_{k=1}^{K} \sum_{(i,j)\in A} f_{ij} x_{ijk}\right) + \sum_{(i,j)\in A} \overline{f_{j0}} x_{ijK} \leq F$. If the Kth node cannot be reached, z_K^f will have to be set to 1 in order for Constraints (17) to be satisfied, where M is a sufficiently large number. Constraints (18) make sure that all the nodes in the tour after the first node that cannot be reached, cannot be reached either. This constraint is necessary since the triangle inequality might not hold due to the randomness of the weights, e.g. it could be that for node j and its successor j' the realization between these nodes is such that $\overline{f_{j0}} \geq f_{jj'} + \overline{f_{j'0}}$. Constraints (19) identify the nodes in the preplanned tour that cannot be reached, based on their index. Two conditions are satisfied for such a target: 1) the incoming arc to node j is the kth arc in the tour and 2) the kth node in the tour cannot be reached with respect to the realizations. Thus, if both $\sum_{i\in N^+} x_{ijk}$ and z_k^f are equal to 1, this constraint will set y_j^f to 1.

Now define the relaxed problem $v_R(x, f)$ as the problem $v(x, f)$ in which (20) and (22) are relaxed to $0 \leq x_{ijk} \leq 1$ and $0 \leq y_i^f \leq 1$ respectively. We have the following result.

Theorem 1. *The objective values of $v_R(x, f)$ and $v(x, f)$ are equal.*

Proof. Since $v_R(x, f)$ is a relaxation of $v(x, f)$ and both are minimization problems we have $v_R(x, f) \leq v(x, f)$. To prove $v_R(x, f) \geq v(x, f)$, let $(x_{ij}, x_{ijk}, y_j, z_k)$ be a solution to $v_R(x, f)$. We will construct a feasible solution to $v(x, f)$ with equal or lower objective value. To do so, define $x'_{ijk} = \lfloor x_{ijk} \rfloor$ and $y'_j = \lfloor y_j \rfloor$ and consider the solution $(x_{ij}, x'_{ijk}, y'_j, z_k)$. As $y'_j \leq y_j$, this solution has an equal or lower objective value. It remains to show that $(x_{ij}, x'_{ijk}, y'_j, z_k)$ is feasible for $v(x, f)$. Constraints (15), (17) and (18) are trivially satisfied. In order to show that Constraints (16) are also satisfied, rewrite these constraints as $x_{ijk} + 1 \geq x_{ij} + \sum_{l\in N^+} x_{li(k-1)}$. In case $x_{ijk} = 1$ the constraint is satisfied in $v(x, f)$ since then only the right hand side is possibly rounded down to obtain the solution to $v(x, f)$. In case $x_{ijk} < 1$ we have $2 > x_{ijk} + 1 \geq x_{ij} + \sum_{l\in N^+} x_{li(k-1)}$. The right hand side of this expression contains $|N^+| + 1$ terms, of which at most one may be equal to 1. After rounding down these terms, the right hand side is therefore at most 1 in $v(x, f)$. Hence, the constraint is satisfied. A similar reasoning applies to Constraints (19). □

4.2. Solution method: sample average approximation

In order to solve the TSOP we use Sample Average Approximation (SAA). SAA is a well-known solution technique within stochastic programming for solving stochastic optimization problems through the use of Monte Carlo simulation. In this technique the objective function of the stochastic problem is approximated by a sample average estimate derived from a random sample. The stochastic problem with this approximate objective function (called the SAA problem) can then be solved either by special purpose algorithms or by deterministic optimization techniques. By repeating the optimization with different samples, candidate solutions and statistical estimates of their optimality gaps can be obtained.

Here we will describe how SAA can be applied to the TSOP, where the recourse costs are modeled as a second-stage linear optimization problem, as just described in Section 4.1.1. We introduce a set of scenarios of weight realizations S. These scenarios are used to estimate $\mathbb{E}_f\left(v(x,f)\right)$ from Formula (13) by the average profit shortage over all scenarios

$$\frac{1}{|S|}\sum_{f\in S} v(x,f). \tag{23}$$

The resulting TSOP is a large mixed integer program, which aims at finding the maximum total estimated expected profit, consisting of the first-stage profit reduced by the average second-stage cost over all predefined scenarios. For each scenario $f \in S$ it should be determined for how long the preplanned tour can be continued, based on the weight realizations f. Thus, $v(x,f)$ has to be solved for each scenario. The SAA objective function of the TSOP becomes:

$$\text{(TSOP-SAA)} \qquad \max \sum_{i\in N} p_i \sum_{j\in N^+\setminus\{i\}} x_{ij} - \frac{1}{|S|}\sum_{f\in S}\sum_{i\in N} p_i y_i^f. \tag{24}$$

The constraints of the SAA problem are Constraints (2) and (3) and (5) to (7), which are related to the first-stage problem, as well as Constraints (15) to (22), where we relax Constraints (20) and (22), and where Constraints (17), (18), (19), (21) and (22) have to be satisfied for all $f \in S$.

Theory on SAA subscribes how to repeat this sampling process in order to obtain statistical estimates of upper and lower bounds on the original objective value (in our case Objective (12)), as well as estimates of the variances of these bounds. This is done by generating M independent samples and by solving the associated M separate SAA problems. Solving these SAA problems will provide M objective values, say $z^1, \ldots, z^M$, and candidate solutions $\hat{x}^1, \ldots, \hat{x}^M$. The average of the optimal objective values of the M SAA problems

$$\overline{z} = \frac{1}{M}\sum_{m=1}^{M} z^m \tag{25}$$

is a statistical estimate for an upper bound on the optimal value of the original problem. An estimate for a lower bound can be obtained by noting that, for any feasible point $\hat{x} \in X$, the value $p^T\hat{x} + Q(\hat{x})$ is a lower bound of the objective value of our TSOP. Using such a feasible point, this lower bound can be estimated using a set N' of newly generated scenarios, independently of the previously created scenarios used to solve the associated SAA problem. Thus,

$$\underline{z}(\hat{x}) = p^T\hat{x} + \frac{1}{|N'|}\sum_{f\in N'} v(\hat{x},f) \tag{26}$$

is a statistical estimate for this lower bound. Typically $|N'|$ is chosen much larger than $|S|$, since the scenarios in N' are only used to evaluate a candidate solution, as opposed to the set S which is used to optimize a large integer program. The variances of the estimators $\overline{z}$ and $\underline{z}(\hat{x})$ can be estimated by respectively

$$\hat{\sigma}^2_{\overline{z}} = \frac{1}{(M-1)\,M} \sum_{m=1}^{M} (z^m - \overline{z})^2, \text{ and} \tag{27}$$

$$\hat{\sigma}^2_{\underline{z}(\hat{x})} = \frac{1}{(|N'|-1)\,|N'|} \sum_{f \in N'} \left(p^T \hat{x} + v(\hat{x}, f) - \underline{z}(\hat{x})\right)^2. \tag{28}$$

Let $\hat{x}^* \in \arg\max \left\{\underline{z}(\hat{x}) : \hat{x} \in \left\{\hat{x}^1, \ldots, \hat{x}^M\right\}\right\}$ denote the solution that provides the highest lower bound. The quality of $\hat{x}^*$ can be evaluated by computing the optimality gap estimate and its associated variance:

$$\underline{z}\left(\hat{x}^*\right) - \overline{z}, \text{ and} \tag{29}$$

$$\hat{\sigma}^2_{\underline{z}(\hat{x}^*)-\overline{z}} = \hat{\sigma}^2_{\overline{z}} + \hat{\sigma}^2_{\underline{z}(\hat{x})}. \tag{30}$$

This method of applying statistical estimates in order to evaluate candidate solutions was suggested in [42] and developed in [43]. Convergence properties were studied in [44].

5. Case study

In this section we will illustrate the use of both the Agile ROP and the TSOP formulation by means of a case study. The TSOP is approached by solving the SAA formulation as just described in Section 4.2. We will first discuss the data for this case study. The same data was also used in the previous research on the Agile ROP. Next, we will describe the model settings like the uncertainty sets in the ROP and the scenario definitions in the TSOP, followed by the comparison of the results of both models. Both models were implemented in Java and optimized by CPLEX 12.1 on an Intel(R) Core(TM) 2 Duo CPU, 2.40 GHz, 1.95 GB of RAM.

5.1. Data

The survey of the OP by Vansteenwegen et al. [27] contains an overview of benchmark instances for the OP. These data sets were used in earlier research on the OP [45, 24, 26], but since the standard OP is deterministic, the data sets do not include uncertainty parameters. Therefore, after describing the nominal data based on one of the benchmark instances, we will discuss how uncertainty was added to the instance.

For our experiments we used the data set ‘Tsiligirides problem 2’ [24]. This data set contains 19 nodes, each with a profit of either 15, 20, 25, 30, 40 or 50. Because of the size of the MIP involved in the SAA solution approach (and the resulting large computation time), we only selected the first 8 nodes out of this data set. For the purpose of this paper, comparing modeling strategies, the size of the test problems is adequate and is quite common for the applications that will be discussed in the next section. The nodes are positioned on a rectangular area of size 15 by 15 units. We use the Euclidean distance $dist(i, j)$ to represent the first part of the average weight between node i and j. As a second part of the average weight, we add a fixed amount to represent the average time spent at a node, to model elements like service times in logistical applications. For each of the nodes we set this amount to 2 units. Summarizing, $\overline{f_{ij}} = dist(i,j) + 2$ for all $(i,j) \in A \setminus \{(i,j)|j = 0\}$. Note that the arcs $(i,0)$, which are the arcs returning to the depot, only have a weight based on the Euclidean distance. The capacity equals 20 units.

Since the benchmark instance does not contain any uncertainty, we need to make an assumption on the distribution of the weight parameters. In robust optimization, the uncertain parameters are assumed to lie within a certain predefined interval. More specifically, recall that in the ROP the realizations f_{ij} are assumed to lie within the interval $[\overline{f_{ij}} - \sigma_{ij}, \overline{f_{ij}} + \sigma_{ij}]$, where $\overline{f_{ij}}$ denotes the average weight and σ_{ij} defines the maximum deviation from the average weight. In stochastic programming on the other hand, we need to define a distribution which we assume the uncertain parameters to follow. Satisfying the requirements of both approaches, we assume in these experiments that the realizations of the weights on arc (i,j) are uniformly distributed in the interval $[\overline{f_{ij}} - \sigma_{ij}, \overline{f_{ij}} + \sigma_{ij}]$. We chose to use the uniform distribution since it covers the whole range of the interval, where all realizations are equally likely to appear. As such, the use of this distribution helps in illustrating the difference between both models for very uncertain situations with respect to the kind of realizations that one can expect. In order to define σ_{ij}, we assume this bound on the deviation from the expected value to consist of two parts. The first part relates to the Euclidean distance between the two nodes and we assume this part to be a fixed percentage α of $dist(i,j)$. For the second part, which is related to the time spent at the node, we assume a fixed maximum absolute deviation c that is equal for all nodes. Summarizing, we construct the interval of the weight between node i and node j by defining $\sigma_{ij} = \alpha dist(i,j) + c$. In this case study we use $\alpha = 0.15$ en $c = 0.5$.

5.2. Model settings

Both models require some settings defining the experimental design. For the ROP we will describe how the uncertainty sets were defined. Next, we will discuss how the scenarios were created and which parameters were chosen to allow for the SAA solution approach of the TSOP.

5.2.1. Agile ROP settings

Recall that in the Agile ROP it is assumed that the weights satisfy

$$f_{ij} = \overline{f_{ij}} + \sigma_{ij}\zeta_{ij}, \tag{31}$$

where $\zeta_{ij} \in [-1, 1]$. In previous research on the Agile ROP [1], several different uncertainty sets were used, each defining a convex set of realizations of the vector ζ against which any solution to ROP should be protected. The uncertainty sets were defined using an L_1, an L_2, an L_∞ norm, or as intersections between two of those norms. For comparison purposes in these experiments we will only focus on the use of the L_∞ norm with size ρ between 0 and 1, and step size 0.1. More specifically, the fuel uncertainty sets that we will use in these experiments are defined as:

$$Z_\infty(\rho) = \left\{\zeta \in \mathbf{R}^{|A|} : \|\zeta\|_\infty \leq \rho\right\} = \left\{\zeta \in \mathbf{R}^{|A|} : \max_{(i,j)\in A} |\zeta_{ij}| \leq \rho\right\}. \tag{32}$$

This setting implies the following: a solution that is protected against all realizations of $\zeta \in Z_\infty(\rho)$ still remains feasible if all realizations differ at most ρ times their maximum deviation from their expected value. Also, with high probability, the solution still remains feasible for realizations outside of this uncertainty set.

Recall that the ROP formulation, combined with these uncertainty sets, is used to find an initial tour. In our results we will include the average shortage and average surplus of the tour information value, in case the agility principles are applied to execute the tour.

5.2.2. TSOP settings

For solving the TSOP by means of the SAA method, we defined $N = 8$ scenarios to estimate the expected profit within each run. This was repeated for $M = 10$ runs in order to define an estimated upper bound on the expected profit. Each scenario includes $|A|$ realizations of weight parameters: one for each of the arcs of the network. We draw these realizations from the previously mentioned distribution: a uniform distribution within the interval $[\overline{f_{ij}} - \sigma_{ij}, \overline{f_{ij}} + \sigma_{ij}]$. To estimate the lower bound on the TSOP, we consider all individual solutions from the M runs, and evaluate them using $N' = 10000$ scenarios, independent from the scenarios previously used to construct the solutions of these M runs. The highest of the objective values evaluated over these N' runs, represents the estimated lower bound of our problem. For each individual run we have used the same 10000 scenarios to estimate the lower bound. As such, different runs resulting in the same tour will result in the same lower bound.

5.3. Results

In this subsection we will discuss the results of the case study for both models and we will show one of the resulting tours for each of the models.

5.3.1. Agile ROP results

Table 1 shows the results for the Agile ROP. The second column shows that increasing the size of the uncertainty set results in saving more planned slack of the available capacity, or equivalently in a lower nominal weight of the planned tour. Column 4 illustrates how the planned profit objective value decreases, when the size of the fuel uncertainty set is increased. This is the profit value that can be obtained for all realizations within the uncertainty set. As such, a tour obtained by a large uncertainty set has a lower profit objective value, but has a higher probability of remaining feasible. This is illustrated in Column 3, where the feasibility values were estimated based on 10000 scenarios. For each scenario, the weights on the arcs of the tour were randomly (uniformly) drawn from the associated interval. In doing so, the proportion of tours that remained feasible, given the weight realizations, was derived. At the same time, some scenarios resulted in a profit shortage and others resulted in a profit surplus. The surpluses were constructed by assuming an online greedy agile strategy as described in [1]. The averages of these shortages and surpluses are given in Columns 5 and 6. Finally, Column 7 contains the estimated expected profit based on these simulations: it is the sum of the profit objective value and the average (negative) profit shortage and the average profit surplus. Note that using the first uncertainty set, $Z_\infty(0)$, comes down to solving the nominal OP. Table 1 clearly shows that, although the nominal tour has the highest planned objective value, it not the best tour with respect to the expected profit. This is due to the fact that the nominal tour very often has to be aborted before completion and therefore part of the planned profit is not obtained. For our instance, the solution to the Agile ROP with $Z_\infty(0.5)$ gives the highest estimated expected profit.

uncertainty	nom.fuel	feas.	obj.	av.shortage	av.surplus	exp.profit
$Z_\infty(0)$	19.851	0.568	70	8.640	0.130	61.490
$Z_\infty(0.1)$	19.131	0.808	60	3.840	1.080	57.240
$Z_\infty(0.2)$	19.131	0.817	60	3.660	0.090	56.430
$Z_\infty(0.3)$	18.732	0.892	60	1.080	0.000	58.920
$Z_\infty(0.4)$	17.845	0.996	55	0.080	5.880	60.800
$Z_\infty(0.5)$	17.592	0.998	55	0.030	7.655	**62.625**
$Z_\infty(0.6)$	17.592	1.000	55	0.000	6.580	61.580
$Z_\infty(0.7)$	17.592	1.000	55	0.000	6.580	61.580
$Z_\infty(0.8)$	17.370	0.999	50	0.010	0.030	50.020
$Z_\infty(0.9)$	16.668	1.000	50	0.000	10.550	60.550
$Z_\infty(1)$	16.668	1.000	50	0.000	10.550	60.550

Table 1: results of the ROP

5.3.2. TSOP results

As explained before, we evaluate the tours of the ROP by correcting them for shortages and surpluses by means of simulation, which we referred to as the Agile ROP. Now, instead of evaluating the estimated expected profit afterwards, the TSOP aims to optimize this value beforehand. Table 2 shows the results of the 10 SAA runs of the TSOP. The first column shows the objective value of each run. This is the average of the profit value of the resulting tour, based on the 8 scenarios within each run, as defined by Objective (24). It is the sum of the first-stage profit and the average (negative) second-stage cost, which are given in Columns 2 and 3 respectively. Note that both the objective value and the second-stage cost are estimates, based only on the 8 scenarios within a specific run. In the same way as we did for the tours resulting from the ROP, we performed a simulation of 10000 scenarios in order to better estimate the expected profit. Columns 4 and 5 show the resulting estimated expected profit shortage and estimated expected profit respectively. Note that the value in Column 5 is again the sum of the profit of the planned tour and the estimated (negative) expected profit shortage. Since the expected profit was obtained in the same manner as the evaluation of the Agile ROP tours, we can compare these values with one another. Even though the ROP utilizes real-time information on the actual fuel realizations at the end of its preplanned tour according to an agile strategy, which the TSOP does not, the TSOP is able to find a tour with a higher estimated expected profit than all of the tours found by the Agile ROP. Notice that such a good solution is already achieved when only eight scenarios were used to solve the associated SAA problem.

obj.	first-stage profit	second-stage cost	av.shortage	exp.profit
63.75	110	46.25	46.363	**63.638**
66.25	110	43.75	46.719	63.282
65	120	55	58.992	61.008
62.5	80	17.5	18.667	61.333
67.5	80	12.5	16.363	**63.638**
68.75	100	31.25	38.667	61.333
66.25	110	43.75	46.719	63.282
63.75	135	71.25	77.646	57.354
65	105	40	47.038	57.962
63.75	120	56.25	56.363	**63.638**

Table 2: results of the TSOP

Table 3 shows the statistics that were derived by the SAA method. The upper bound is the average of the objective values in the first column of Table 2 over the 10 runs. The lower bound is the best solution out of the 10 runs, evaluated over 10000 scenarios. This value equals the simulation evaluation in Column 5 of Table 2 of the tours found in Run 1, 5 and 10. The gap and the variances were derived as described in Section 4.2.

upper bound	65.25
var. upper bound	0.375
lower bound	63.638
var. lower bound	0.007
gap	1.6125
var. gap	0.382

Table 3: SAA statistics of the TSOP

5.3.3. Topology of tours for both approaches

Table 5.3.3 contains the tour resulting from the ROP on the left side and the tour resulting from the TSOP on the right side. The values depicted above the nodes are the profit values that represent the importance of the nodes. The values between brackets depicted below the nodes are the indices of the nodes. Both tours are built for the same data instance, 'Tsiligirides problem 2'. We selected the Agile ROP tour that resulted from the choice of $Z_\infty(0.5)$ as the uncertainty set, since this choice resulted in the highest expected profit. The black lines, together with the dotted line represent the initial tour that resulted from this uncertainty set. The two thick arrows from node (5) to node (4) and from node (4) back to the depot are depicted to illustrate how the agile modeling works in practice, based on one possible outcome of the realizations of the weights. In the illustrated case, we encountered beneficial realizations of the weights, such that after the final node of the initial tour, the tour was extended based on greedy agility principles. The tour on the right side of Table 5.3.3 was the best TSOP tour found with respect to the estimated expected profit. In the next section we will elaborate on the interpretation of the results just presented and the consequence of selecting one of the tours illustrated in Table 5.3.3.

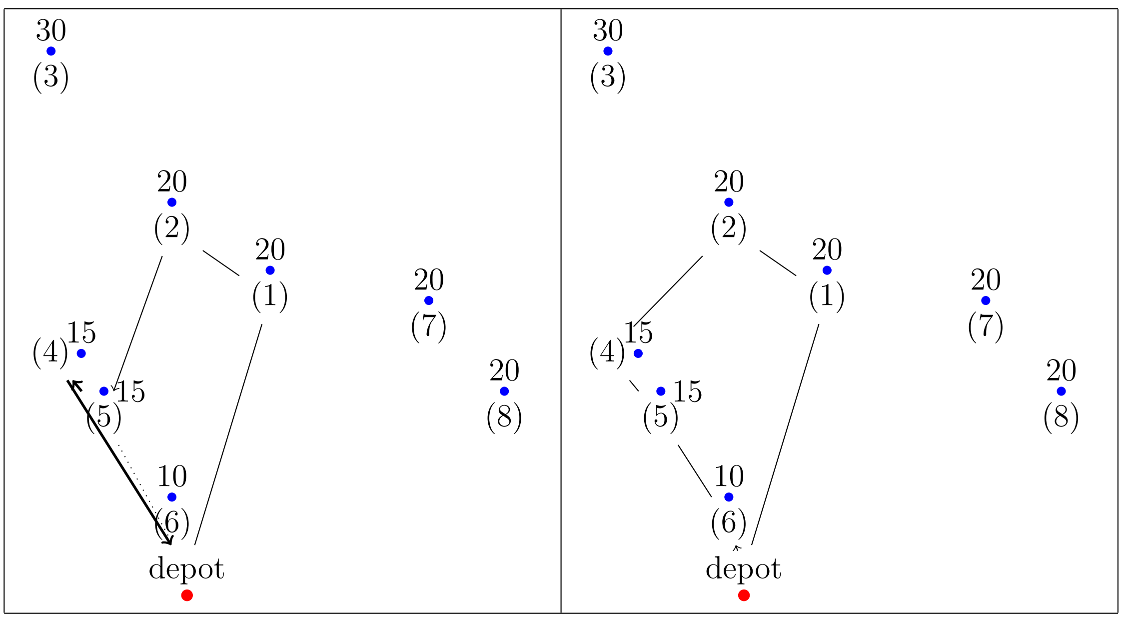

Table 4: Left: the ROP tour resulting from $Z_\infty(0.5)$ and the agile addition for a specific set of weight realizations (the thick arrows). Right: the TSOP tour with the highest estimated expected profit found.

6. Comparing modeling strategies

Based on the results of the case study, in this section we will discuss the suitability of the ROP and the TSOP for different applications of the OP. We will first describe two applications of the OP to illustrate which characteristics might influence the suitability of the use of the ROP or the TSOP. The first one is the Unmanned Aerial Vehicle (UAV) mission planning problem and the second one is the tourist tour planning problem.

6.1. Background on the applications

UAVs are important assets in requiring intelligence information for both military and civilian purposes. In military missions, they are used to record imagery of important 'target' locations. Civilian purposes of the use of UAVs are for example protection of wild parks against poachers or wildfires. In both situations, the tour of a UAV mission needs to be scheduled by taking into account the importance of all interesting locations and the fuel capacity of the UAV. Several papers discuss the use of the OP and extensions of the OP to optimize UAV tours (see for example [46, 47, 1]). In modeling the problem, the importance of the locations is represented by the profit values of the OP. The fuel requirement from location i to location j (possibly including the fuel requirement for recording target j) is modeled by the weights of the OP. The so-called 'UAV recovery point' is represented by the depot location of the OP instance.

The second application of the OP that we will consider is the tourist tour planning problem. In this problem a tourist aims to visit several different sightseeing locations, for each of which the tourist has a different preference level, represented by the profit values of the nodes in the OP. The length of the tourist tour is restricted by the total time the tourist wants to spend on sightseeing. In modeling this problem the time the tourist spends visiting a location is added to the travel time required to arrive at the location (similar to what is done for the recording time in the UAV problem). This total time is then represented by the weight of the arcs. Research on this problem and similar applications was presented e.g. in [48, 49, 50, 51].

6.2. Suitability of ROP and TSOP to the applications

In general, when deciding between a stochastic programming or a robust optimization approach, the characteristics of the uncertain parameters that are available to the decision maker play a role. As mentioned before, robust optimization requires less specific assumptions on the distribution of the uncertain parameters than stochastic programming. This yields an advantage when few information is known about the behavior of the uncertain parameters or when this behavior is difficult to approximate. Also, when the cause of the uncertainty is related to multiple factors, the behavior of the uncertain parameters can be hard to predict. This is for example the case for the fuel uncertainty in the UAV problem, since the cause of uncertainty might originate from different factors like deviations from the expected weather conditions, uncertain recording times and other unforeseen events. Additionally, it is very difficult to produce accurate estimates on the average fuel consumption due to the complex relations between the fuel usage and explanatory variables like weather conditions, air pressure and flight behavior of the UAV. A similar reasoning applies to the tourist planning problem. Here one could also argue that the distribution of the travel time between the sightseeing locations is hard to approximate and especially the time spent at the sightseeing location might be hard to predict beforehand. Thus, we assume that for both applications the behavior of the data uncertainty is hard to approximate. Consequently, we chose to use a uniform distribution in our case study on the TSOP. When more information on the distribution of the uncertain data would be available, the results of the TSOP will only improve further, since the likelihood of specific realizations within the interval is then tuned by the selected probability distribution. In that case, the advantage of stochastic programming is that it will explicitly exploit the information available on data uncertainty.

The important difference between the two applications mentioned is the consequence of not satisfying the capacity constraint. Since UAVs are expensive assets, it is of high importance that the UAV will actually return to its recovery point before it runs out of fuel. For a tourist on the other hand, violation of the capacity constraint would only imply that the time spent on the sightseeing tour will exceed the time limit that was decided upon in advance. Because of these different interpretations of the capacity constraint, the ROP is the best choice for the UAV mission planning problem, while the TSOP is the best strategy the tourist planning problem. More specifically, in military missions guaranteeing the feasibility of completion of a plan is of uttermost importance. The ROP formulation allows the mission planner to balance a specific feasibility against the planned objective value. In that case, for a certain predetermined feasibility, the best objective value is found. For the tourist problem, exceeding the time limit might mean not all of the sightseeing locations can be visited, because for example a museum may have closed already. Consequently, when the time limit is exceeded less profit can be obtained, but the tourist will in most cases still be able to return to the endpoint. Therefore, the tourist can afford to take a higher 'risk' of violating the capacity constraint than the UAV mission planner and as such, the TSOP is the preferred modeling strategy for the tourist planning problem.

Relating our applications to the results of our case study, the following observations can be made. The difference between the topological structure of the tours in Table 5.3.3 illustrates why the TSOP is able to produce a tour with a higher estimated expected total profit value. As just mentioned, in the military setting the ROP tour is preferred over the TSOP tour. When the UAV operator decides on the ROP tour illustrated in Table 5.3.3 to be the initial tour plan, our simulation results in Table 1 show that the operator has an estimated probability of 0.998 of actually being able to reach all targets in this tour. Only when the UAV operator has achieved its primary goal of reaching all targets of its initial tour, the possibility of exploiting the real-time knowledge on the remaining fuel will be examined. As the thick arrows in the ROP tour show, in case of beneficial fuel realizations additional targets will be added to the tour after target number 5. The figure shows that, if possible, target 4 is the first target to be added, which actually lies further ahead from the recovery point than the previous target, number 5. In case of very low fuel realizations the operator might then still have enough fuel to also visit target 6 before returning to the recovery point, but for the situation depicted in Table 5.3.3 this is not the case. The order of the targets in the TSOP tour on the other hand, might for the same fuel realizations have allowed for target 6 to be obtainable as well. This example thus illustrates the effect of the different optimization objectives combined with the topological characteristics of the data instance on the total obtainable profit value. While the UAV operator primarily wants to perform its initial flight plan, the tourist applying the TSOP tour in Table 5.3.3 just wants to be able to visit the combination of sightseeing locations with the highest possible expected value.

7. Conclusion

In order to compare modeling strategies we applied both stochastic programming and robust optimization to an uncertain version of the orienteering problem in which the weights of the arcs are uncertain. Each of the associated models deals with uncertainty in its own way. The first model is the known Agile Robust Orienteering Problem (Agile ROP) which ensures feasibility of a tour against all realizations of the uncertain weight parameters within a predefined uncertainty set and prescribes how to execute this tour in a dynamic environment. Since the initial tour construction does not take into account what happens to the total profit value in case of (un)beneficial weight realizations, we introduced the Two-Stage Orienteering Problem (TSOP). The TSOP optimizes the expected profit, which is modeled as the expected difference of the first-stage initial tour and the second-stage profit shortage. The profit shortage is the profit value of the last node(s) of the first-stage tour, that cannot be obtained as a result of the realizations of the weights. To solve the TSOP, we implemented one of the well-known solution methods in stochastic programming: the sample average approximation method. In a case study we illustrated how the use of the TSOP results in higher expected profit than the use of the ROP. This is due to the fact that the TSOP optimizes over the profit value that is corrected for the effects of the actual weight realizations on the total profit to be obtained, contrary to the ROP. We also discussed the suitability of both models for two applications of the OP: the UAV mission planning problem and the tourist tour planning problem.

The TSOP that was introduced in this paper is a two-stage recourse model where the recourse costs only consist of profit shortage. The resulting tours are relatively long and allow the decision maker to continue to the next node of the tour as long as the remaining capacity allows for it. Future research could focus on a multi-stage model, including a 'profit surplus'. The resulting first-stage solution is an initial tour that will be followed as long as possible, and additional nodes could be added to the tour according to a predefined online strategy. In that way, the expected profit could possibly be further increased, since the solution to the two-stage model would be part of the solution space of the multi-stage model. Also, further research could focus on implementing other solution approaches.

Concluding, both stochastic programming and robust optimization provide tools to appropriately take into account data uncertainty already in the modeling stage. Depending on the characteristics of the application the use of one of the approaches might be preferred over the other. In particular, the Agile ROP is the preferred modeling strategy feasibility if an initial plan is highly valued and/or the probabilistic distribution of the data is unknown. When the feasibility constraint is less important or when more information on the data uncertainty is available, the TSOP modeling strategy is preferred since this it optimizes the expected profit value in the uncertain OP.

References

[1] L. Evers, T. Dollevoet, A. Barros, H. Monsuur, Robust UAV mission planning, EI 2011-07, Econometric Institute, Erasmus University Rotterdam (2011).

[2] G. Dantzig, Linear Programming and Extensions, Princeton University Press, Princeton, NJ, 1963.

[3] E. Beale, On minimizing a convex function subject to linear inequalities, Journal of the Royal Statistic Society Series B 17 (1955) 173–184.

[4] A. Charnes, W. W. Cooper, Chance-constrained programming, Management Science 6(1) (1959) 73–79.

[5] J. R. Birge, F. Louveaux, Introduction to Stochastic Programming, Springer, 2000.

[6] A. Ruszczinski, A. Shapiro, Stochastic Programming, Elsevier, Amsterdam, 2003.

[7] M. H. van der Vlerk, Stochastic programming bibliography, http://www.eco.rug.nl/mally/spbib.html (1996-2007).

[8] A. Soyster, Convex programming with set-inclusive constraints and applications to inexact linear programming, Operations Research 21(5) (1973) 1154–1157.

[9] P. Kouvelis, G. Yu, Robust Discrete Optimization and its Applications, Kluwer Academic Publishers, London, 1997.

[10] A. Ben-Tal, A. Nemirovski, Stable truss topology design via semidefinite programming, SIAM J. on Optimization 7(4) (1997) 991–1016.

[11] A. Ben-Tal, A. Nemirovski, Robust convex optimization, Mathematics of Operations Research 23(4) (1998) 769–805.

[12] L. El Ghaoui, H. Lebret, Robust solution to least-squares problems with uncertain data, SIAM Journal on Matrix Analysis and Applications 18 (1997) 1035–1064.

[13] L. El Ghaoui, F. Oustry, H. Lebret, Robust solutions to uncertain semidefinite programs, SIAM J. on Optimization 9 (1998) 33–52.

[14] A. Charnes, W. W. Cooper, Deterministic equivalents for optimizing and satisficing under chance constraints, Operations Research 11 (1963) 18–39.

[15] A. Prékopa, Stochastic Programming, Kluwer, Dordrecht, 1995.

[16] A. Ben-Tal, A. Goryashko, E. Guslitser, A. Nemirovski, Adjustable robust solutions of uncertain linear programs, Mathematical Programming 99 (2004) 351–376.

[17] A. Atamturk, M. Zhang, Two-stage robust network flow and design under demand uncertainty, Operations Research 55 (2007) 662–673.

[18] X. Chen, M. Sim, P. Sun, J. Zhang, A linear-decision based approximation approach to stochastic programming, Operations Research 56(2) (2008) 344–357.

[19] D. Bertsimas, C. Caramanis, Finite adaptability in multistage linear optimization, IEEE Transactions on Automatic Control 55(12) (2010) 2751–2766.

[20] A. Ben-Tal, A. Nemirovski, Robust solutions of linear programming problems contaminated with uncertain data, Mathematical Programming 88 (2000) 411–424.

[21] D. Bertsimas, M. Sim, The price of robustness, Operations Research 52(1) (2004) 35–53.

[22] X. Chen, M. Sim, P. Sun, A robust optimization perspective on stochastic programming, Operations Research 55(6) (2007) 1058–1071.

[23] D. Bertsimas, A. Thiele, Robust and data-driven optimization modern decision-making under uncertainty, Tutorial (2006).

[24] T. Tsiligirides, Heuristic methods applied to orienteering, Journal of the Operational Research Society 35 (1984) 797–809.

[25] B. Golden, L. Levy, R. Vohra, The orienteering problem, Naval Research Logistics 34 (1987) 307–318.

[26] M. Fischetti, J. Salazar, P. Toth, Solving orienteering problem through branch-and-cut, INFORMS Journal on Computing 10 (1998) 133–148.

[27] P. Vansteenwegen, W. Souffria, D. van Oudheusden, The orienteering problem: A survey, European Journal of Operational Research 209(1) (2011) 1–10.

[28] G. Laporte, S. Martello, The selective traveling salesman problem., Discrete Applied Mathematics 26 (1990) 193–207.

[29] M. Gendreau, G. Laporte, F. Semet, A branch-and-cut algorithm for the undirected selective travelling salesman problem, Networks 32 (1998) 263–273.

[30] T. Thomadsen, T. Stidsen, The quadratic selective travelling salesman problem, Technical report 2003-17, Technical University of Denmark, Informatics and Mathematical Modelling (2003).

[31] S. Kataoka, S. Morito, An algorithm for the single constraint maximum collection problem, Journal of the Operations Research Society of Japan 31(4) (1988) 515–530.

[32] S. Butt, T. Cavalier, A heuristic for the multiple tour maximum collection problem, Computers and Operations Research 21 (1994) 101–111.

[33] H. L. Teng, S. Y.and Ong, H. C. Huang, An integer L-shaped algorithm for the time-constrained traveling salesman problem with stochastic travel times and service times, Asia-Pacific Journal of Operational Research 21 (2004) 241–257.

[34] E. Arkin, J. Mitchell, G. Narasimhan, Resource-constrained geometric network optimisation, Proc. 14th ACM Symposium on Computational Geometry June 6-10 (1998) 307–316.

[35] D. Feillet, P. Dejax, M. Gendreau, Traveling salesman problems with profits: an overview, Transportation Science 39 (2005) 188–205.

[36] M. Dell'Amico, F. Maffioli, P. Värbrand, On prize-collecting tours and the asymmetric travelling salesman problem, International Transactions in Operational Research 2(3) (1995) 297–308.

[37] E. Balas, The prize collecting traveling salesman problem, Networks 19(6) (1989) 621–636.

[38] H. Tang, E. Miller-Hooks, Algorithms for a stochastic selective travelling salesperson problem, The Journal of the Operational Research Society 56 (2005) 439–452.

[39] A. M. Campbell, M. Gendreau, B. W. Thomas, The orienteering problem with stochastic travel and service times, Annals of Operations Research 186 (2011) 61–81.

[40] A. Ben-Tal, L. E. Ghaoui, A. Nemirovski, Robust Optimization, Princeton University Press, Princeton and Oxford, 2009.

[41] G. Laporte, F. Louveaux, L. van Hamme, An integer L-shaped algorithm for the capacitated vehicle routing problem with stochastic demands, Operations Research 50 (2002) 415–423.

[42] V. Norkin, G. Pflug, A. Ruszcynski, A branch and bound method for stochastic global optimization, Mathematical Programming 83 (1998) 425–450.

[43] W. Mak, D. Morton, R. Wood, Monte Carlo bounding techniques for determining solution quality in stochastic programs, Operations Research Letters 24 (1999) 47–56.

[44] A. Kleywegt, A. Shapiro, T. H. de Mello, The sample average approximation method for stochastic discrete optimization, SIAM Journal on Optimization 12 (2001) 479–502.

[45] I. Chao, B. Golden, E. Wasil, Theory and methodology: a fast and effective heuristic for the orienteering problem, European Journal of Operational Research 88 (1996) 475–489.

[46] J. O. Royset, D. N. Reber, Optimized routing of unmanned aerial systems for the interdiction of improvised explosive devices, Military Operations Research 3(4) (2010) 2917–2922.

[47] F. Mufalli, R. Batta, R. Nagi, Simultaneous sensor selection and routing of unmanned aerial vehicles for complex mission plans, Working paper, Department of Industrial and Systems Engineering, University at Buffalo (State University of New York) (2010).

[48] P. Vansteenwegen, D. van Oudheusden, The mobile tourist guide: An OR opportunity, OR Insight 20 (3) (2007) 21–27.

[49] W. Souffriau, P. Vansteenwegen, J. Vertommen, G. Vanden Berghe, D. Van Oudheusden, A personalised tourist trip design algorithm for mobile tourist guides, Applied Artificial Intelligence 22 (2008) 964–985. doi:10.1080/08839510802379626.
URL `https://lirias.kuleuven.be/handle/123456789/221911`

[50] M. Schilde, K. F. Doerner, R. F. Hartl, G. Kiechle, Metaheuristics for the bi-objective orienteering problem., Swarm Intelligence (2009) 179–201.

[51] X. Wang, B. L. Golden, E. A. Wasil, Using a genetic algorithm to solve the generalized orienteering problem, In: Golden, B., Raghavan, S., Wasil, E. (Eds.) The veheicle routing problem: latest advances and new challanges (2008) 263–274.