Computers & Operations Research 43 (2014) 248–260

Contents lists available at ScienceDirect

Computers & Operations Research

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

A two-stage approach to the orienteering problem with stochastic weights

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

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

ARTICLE INFO

Available online 5 October 2013

Keywords:
Uncertain orienteering problem
Stochastic programming
Stochastic heuristic

ABSTRACT

The Orienteering Problem (OP) is a routing problem which has many interesting applications in logistics, tourism and defense. The aim of the OP is to find a maximum profit path or tour, which is feasible with respect to a capacity constraint on the total weight of the selected arcs. In this paper we consider the Orienteering Problem with Stochastic Weights (OPSWs) to reflect uncertainty in real-life applications. We approach this problem by formulating a two-stage stochastic model with recourse for the OPSW where the capacity constraint is hard. The model takes into account the effect that stochastic weights have on the expected total profit value to be obtained, already in the modeling stage. Since the expected profit is in general non-linear, we introduce a linearization that models the total profit that can be obtained for a given tour and a given scenario of weight realizations. This linearization allows for the application of Sample Average Approximation (SAA). The SAA solution asymptotically converges to the optimal solution of the two-stage model, but is computationally expensive. Therefore, to solve large instances, we developed a heuristic that exploits the problem structure of the OPSW and explicitly takes the associated uncertainty into account. In our computational experiments, we evaluate the benefits of our approach to the OPSW, compared to both a standard deterministic approach, and a deterministic approach that is extended with utilization of real-time information.

1. Introduction

The Orienteering Problem (OP) is a routing problem which possesses characteristics of both the well-known Traveling Salesman Problem (TSP) and Knapsack Problem (KP). The OP considers a profit associated to each node and a weight associated to each arc. The profit is collected only if a node is visited. The aim is to find a maximum profit path or tour, satisfying a capacity constraint on the total weight of the arcs selected in the solution. Contrary to the TSP, in the OP not all nodes have to be visited. The selection of nodes in the most profitable way, such that the weight associated to this selection does not exceed a predefined capacity, relates to the KP. In the KP however, the ordering of the items in the selection does not influence the associated total weight, like in the OP. The weights in the OP represent for instance travel costs, travel (and service) time or fuel consumption and the profits model the importance of the locations.

The OP was introduced by Tsiligirides [1] and was shown to be NP-Hard by Golden et al. [2]. An exact algorithm for the OP was proposed by Fischetti et al. [3] and a recent survey on the OP is given by Vansteenwegen et al. [4]. The OP is also known as the selective traveling salesperson problem [5–8], the maximum collection problem [9,10], the time-constrained traveling salesman problem [11] and the bank robber problem [12]. 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 are always defined as a tour planning problem.

The OP has many interesting applications in logistics, tourism and defense. For example, a military application of the OP considers Unmanned Aerial Vehicle (UAV) mission planning to collect intelligence information about different locations in the area of operations (see for example [13–15]). The aim of these missions is to acquire as much information as possible during the flight, while the length of the flight is limited by the available fuel capacity of the UAV. Another application of the OP is the tourist tour planning problem. In this problem, a tourist wants to visit several different sightseeing locations, for each of which the

* Corresponding author.
E-mail address: lanah_evers@hotmail.com (L. Evers).

0305-0548/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cor.2013.09.011

tourist has a different preference level. The length of the tourist tour is restricted by the total time the tourist can spend on sightseeing. Research on this problem and similar applications was presented e.g. in [16–19]. Yet other applications of the OP include maintenance service planning [20], determining which suppliers to audit to maximize recovered claims [21] and planning of scouting tours to different high schools to scout talented sports players [10,22].

In reality, the weights of the OP are often uncertain. This uncertainty can be caused by several factors like weather circumstances, congestion, and other unforeseen events. Therefore, the solution to the deterministic OP might turn out to be infeasible or suboptimal in reality. In this paper, we focus on the problem of constructing an OP tour, while only the distribution of the uncertain weights is known beforehand. This problem can be classified as the Orienteering Problem with Stochastic Weights (OPSWs). While executing the tour, the realizations of the weights will be gradually revealed. We assume that the capacity constraint on the total weight is hard: the depot has to be reached before the total realized weight exceeds the capacity and no profit can be obtained for the unvisited nodes in the planned tour. This assumption is suitable to different applications of the OP. For example, in UAV mission planning, the UAV has to return to the depot as soon as the remaining fuel capacity drops below a certain level. The same applies to the tourist planning problem due to, for instance, tight tourist schedules.

1.1. Related literature

In spite of its practical relevance, literature on uncertain variants of the OP is limited. The existing approaches used in the literature to tackle the OPSW and variants thereof are stochastic programming and robust optimization.

The modeling approaches from stochastic programming that were applied to variants of the OPSW, are chance constrained programming and the use of two-stage stochastic recourse models. Tang and Miller-Hooks [8] propose a chance constrained model of a generalization of the OPSW, in which the total duration of the tour has to satisfy the capacity constraint with at least a predefined probability, based on discrete probability distributions. Recourse models for the OPSW were introduced by Teng et al. [11] and Campbell et al. [23]. In a two-stage recourse model of the OP, the first-stage decision is to select and sequence the nodes to be visited, based on knowledge about the distribution of the uncertain weights and the expected recourse cost in the second stage. This second-stage recourse cost describes a penalty for the effect that the actual realizations of the weights of the arcs might have on the tour that was chosen in the first stage. As second-stage cost, Teng et al. [11] consider a penalty that is proportional to the amount in which the time limit is exceeded. This penalty cost represents for example overtime that has to be paid to drivers for shifts that take longer than planned. The second-stage cost that is considered by Campbell et al. [23] involves a penalty for nodes in the first-stage solution that cannot be visited before the deadline, representing for example a loss in customer goodwill. Contrary to the problem that we consider, both recourse models treat the capacity constraint as a soft constraint: nodes can still be visited after the total weight realizations exceed the total capacity. Besides the penalty functions, an additional difference between the models by Teng et al. [11] and Campbell et al. [23] is the set of probability distributions that is considered for the stochastic weights: Teng et al. [11] consider only discrete probability distributions, while Campbell et al. [23] consider several continuous probability distributions. Furthermore, Teng et al. [11] consider the tour planning OP, while Campbell et al. [23] consider a path solution that starts at a predefined begin point, where the capacity constraint has to be satisfied only until the final node selected in the solution.

The first (and to our knowledge, so far only) robust optimization approach to the OP was introduced by Evers et al. [15]. Their approach produces solutions that are protected against given degrees of weight uncertainty, or in other words, that can cope with all weight realizations within a predefined uncertainty set. The size of the uncertainty set has two effects on the solution produced. A large uncertainty set will ensure that the solution will remain feasible for realizations of the uncertain parameters with a higher probability. On the other hand, the planned profit value of such a tour will be lower. Specific agile policies are incorporated to respond to the effect of the weight realizations during execution of the tour. These effects are, however, not explicitly taken into account beforehand in constructing the optimal solution.

Related to the OPSW, is the Orienteering Problem with Stochastic Profit (OPSP), introduced by Ilhan et al. [21]. In the OPSP the profits are stochastic, while the weights are deterministic and known beforehand. The aim of the OPSP is to maximize the probability of collecting more than a prespecified target profit level. Ilhan et al. [21] develop an exact solution scheme based on a parametric formulation of the problem. Another OP in which stochastic profits are considered, is the Stochastic Scouting Problem (SSP) [22]. In the SSP, not only the profits, but also the weights of the OP are assumed to be stochastic. A two-stage approach is used to maximize the expected profit under both these uncertainties.

1.2. Our approach

In this paper we will introduce a new two-stage recourse model to tackle the OPSW. We will consider stochasticity in the weights of the arcs of the OP and the effect thereof on the profit value to be obtained. The first-stage decision is to construct a tour, which may have to be aborted in the second stage before reaching the final node because of the weight realizations. This recourse action has some similarity to the recourse action used by Laporte et al. [24] for the capacitated vehicle routing problem with stochastic demands. Contrary to the stochastic recourse models for uncertain OPs from the existing literature, we assume that the capacity constraint on the total weight is hard. Consequently, we define the second-stage recourse cost as the sum of the profit of the nodes in the first-stage tour that cannot be reached due to specific weight realizations. Our approach allows for any predefined probability distribution.

We propose two approaches for solving the OPSW: Sample Average Approximation (SAA), a well-known technique from stochastic programming, which we base on the two-stage formulation, and a heuristic approach which takes advantage of the problem structure of the OPSW. To facilitate the SAA approach, we introduce a linearization which determines the profit to be obtained for a given tour and a given set of realizations of uncertain weights. The performance of both solution approaches is explored in our computational experiments. These experiments also illustrate the advantages of our approach compared to the deterministic OP solution and an extension thereof.

This paper is structured as follows. In Section 2 we provide a formulation of the deterministic OP and we will introduce the two-stage model. We will describe the SAA solution method applied to the two-stage model in Section 3 and we will describe our OPSW heuristic in Section 4. In Section 5, we discuss our computational experiments in which we compare our approach to the deterministic OP, to an improved policy based on the deterministic OP, and to upper bounds on the performance of our solution approaches. Concluding remarks and directions for further research are given in Section 6.