

Resource Allocation in an Uncertain Environment: Application to Snowplowing Operations in Utah

Yinhu Wang,^a Ye Chen,^b Ilya O. Ryzhov,^{c,*} Xiaoyue Cathy Liu,^a Nikola Marković^a

^aDepartment of Civil and Environmental Engineering, University of Utah, Salt Lake City, Utah 84112; ^bDepartment of Statistical Sciences and Operations Research, Virginia Commonwealth University, Richmond, Virginia 23284; ^cRobert H. Smith School of Business, University of Maryland, College Park, Maryland 20742

*Corresponding author

Contact: yinhu.wang@utah.edu (YW); ychen24@vcu.edu (YC); iryzhov@umd.edu,  <https://orcid.org/0000-0002-4191-084X> (IOR); cathy.liu@utah.edu (XCL); nikola.markovic@utah.edu,  <https://orcid.org/0000-0003-0883-2701> (NM)

Received: January 19, 2023

Revised: November 17, 2023; March 20, 2024

Accepted: April 18, 2024

Published Online in Articles in Advance:
May 22, 2024

<https://doi.org/10.1287/trsc.2023.0024>

Copyright: © 2024 INFORMS

Abstract. We consider a two-stage planning problem where a fleet of snowplow trucks is divided among a set of independent regions, each of which then designs routes for efficient snow removal. The central authority wishes to allocate trucks to improve service quality across the regions. Stochasticity is introduced by uncertain weather conditions and unforeseen failures of snowplow trucks. We study two versions of this problem. The first aims to minimize the maximum turnaround time (across all regions) that can be sustained with a user-specified probability. The second seeks to minimize the total expected workload that has not been completed within a user-specified time frame. We develop algorithms that solve these problems effectively and demonstrate their practical value through a case application to snowplowing operations in Utah, obtaining solutions that significantly outperform the allocation currently used in practice.

Funding: Financial support from the Utah Department of Transportation [Grant 218138]; the National Science Foundation [Grant CMMI-2112758]; and the Mountain-Plains Consortium [Grant 637] is gratefully acknowledged.

Supplemental Material: The e-companion is available at <https://doi.org/10.1287/trsc.2023.0024>.

Keywords: resource allocation • stochastic optimization • arc routing • snow plowing • fleet management


1. Introduction

Winter road maintenance is essential for ensuring public mobility and safety, especially in areas with long winters and frequent snowfalls. According to the Federal Highway Administration, more than 70% of U.S. roads are located in regions affected by snowfalls (FHWA 2020). Furthermore, snow removal and ice control cost taxpayers over \$2.3 billion/year (Han and Yu 2017). Such high expenditures are unsurprising, given the amount of personnel and mechanization involved in winter road maintenance. For example, Utah experiences more than 25 snowstorms a year, and the Utah Department of Transportation (UDOT) employs hundreds of trucks to maintain state-owned roadway facilities, with costs averaging at about \$1 million/storm (UDOT 2021). In the presence of such large costs, departments of transportation (DOTs) are seeking to better manage limited resources and meet the public's rising level-of-service expectations.

Snow removal operations are too complex to be managed centrally across the entire state road network. Thus, state DOTs partition their road networks into regions that independently operate their designated fleets of vehicles (Perrier, Langevin, and Campbell 2007b;

Kinable, van Hoeve, and Smith 2021). Specifically, each region designs plowing routes such that all the road lanes within its jurisdiction are cleaned as efficiently as possible, an example of what is known as the arc-routing problem (Golden and Wong 1981). To give one example, UDOT divides northern Utah into 12 regions, each of which is equipped with snowplow trucks and crews for snow removal operations. Figure 1 shows the geographic partition and allocation of trucks that UDOT used in 2020. Although the state DOT does not coordinate activities within or between regions, it can use the allocation of trucks to regions to influence the efficiency of the system as a whole; for example, the performance metrics in Figure 1 indicate that region 2 is underequipped and that service quality could be improved by adjusting the fleet allocation.

Service quality in region r , given x_r allocated trucks, can be evaluated in two ways. Turnaround time, denoted by $Q_r(x_r)$, is a commonly adopted indicator (Salazar-Aguilar, Langevin, and Laporte 2012; Dowds and Sullivan 2021) of the efficiency of snow removal operations. A second metric is the number $M_r(x_r)$ of uncleared lane miles after some prespecified time T_r , the idea being that it is necessary to complete snow removal by this time to ensure safe

Figure 1. (Color online) Winter Road Maintenance in Northern Utah, Based on the Data for 2020 and Assuming All Trucks Are Operational

Notes. (a) Northern Utah is divided into $R = 12$ regions, and each region r is allocated x_r trucks. (b) The resulting turnaround time in each region, $Q_r(x_r)$, varies from 64 to 187 minutes. (c) After 90 minutes for all r , the number $M_r(x_r)$ of uncleaned lane miles in region r varies from 8 to 175. See an animation of the snowplowing routes: <https://youtu.be/Aqov3wbdWxI>. (a) Fleet allocation in 2020. (b) Turnaround times. (c) Uncleaned lane miles.

access to hospitals, fire stations, and key activity areas. Both metrics may be subject to uncertain weather conditions: for example, heavy snowfall may significantly impede the mobility of snowplow trucks (Schultz et al. 2022), thus increasing turnaround times. Snowplowing operations are also commonly affected (Li et al. 2019, Iyer et al. 2021) by randomly arising “failures” of trucks (e.g., mechanical issues or driver no-shows), whose likelihood may also be influenced by the weather. In this paper, we consider two fleet allocation problems that account for these sources of uncertainty in distinct ways. Both problems can be stated very concisely. The first allocates K trucks between R regions to minimize the maximum turnaround time that can be sustained with a user-specified probability:

$$\begin{aligned} \min_{x \in \mathbb{Z}_{++}^R, t \geq 0} \quad & t \\ \text{s.t.} \quad & \sum_{r=1}^R x_r = K, \\ & \mathbb{P}\left(\max_{r=1, \dots, R} Q_r(x_r - \xi_r) \leq t\right) \geq 1 - \alpha. \end{aligned} \quad (1)$$

The quantities $\xi_r \in \{0, 1, \dots, x_r\}$ denote truck breakdowns in each region. The second problem allocates K trucks between R regions to minimize the expected total (across all regions) uncleaned lane miles:

$$\begin{aligned} \min_{x \in \mathbb{Z}_{++}^R} \quad & \mathbb{E}\left[\sum_{r=1}^R M_r(x_r - \xi_r)\right] \\ \text{s.t.} \quad & \sum_{r=1}^R x_r = K. \end{aligned} \quad (2)$$

Thus, problem (1) is a risk-averse formulation that aims to guarantee a certain level of service quality in the presence of uncertainty (the parameter α controls the user-specified probability that the attained objective will not

be exceeded), whereas problem (2) is a risk-neutral formulation that simply optimizes average service quality. The probability in (1) and the expectation in (2) are taken over the joint distribution of weather and truck failures.

Before we discuss these two problems in more detail, it is important to explain how we handle the cost functions $Q_r(x_r)$ and $M_r(x_r)$. For any fixed r and x_r , both quantities can be computed by repeatedly solving arc-routing problems (for different weather scenarios) on the corresponding network data. Many papers on fleet management, such as Crainic et al. (2016), model the resource allocation decision jointly with all subsequent routing decisions using a two-stage formulation. We do *not* do this; rather, we compute the cost functions entirely offline, and problems (1) and (2) treat all of the $Q_r(x_r)$ and $M_r(x_r)$ values as given inputs.

At first glance, such an approach seems to incur significant computational overhead. However, by separating the routing problems from the resource allocation decision in this way, the offline computations become much easier. Under a fixed set of weather conditions, the values $Q_r(x_r)$ and $M_r(x_r)$ are completely unrelated across regions r (each region has its own network topology), and so the routing problems are naturally parallelizable. In the context of northern Utah, a single such problem can be solved in five seconds (for a small region) to 30 minutes (for a large one) on a laptop. It is therefore easy to solve multiple problems concurrently on different machines. With these computations done, the actual resource allocation decisions can also be computed much more efficiently (the approaches we develop require no more than five minutes). This approach is also highly modular—if we decide to tweak the performance metrics (e.g., by changing the threshold T_r) or to include additional relevant constraints in the routing problems, these changes can be made offline without rerunning a large two-stage model. The flexibility

thus afforded in integrating routing and fleet allocation is, in and of itself, a novel and attractive aspect of our work.

Thus, the methodological focus of our paper is not on the routing problems solved by individual regions, but on the central authority's resource allocation decision, as represented by problems (1) and (2). Both problems are highly nonlinear and nonconvex. In (1), the main difficulty is the last constraint, which is an example of a chance constraint (Ahmed and Xie 2018). Chance constraints can sometimes be transformed into their deterministic equivalents, but it is well known that these reformulations may be more difficult to solve (Tang et al. 2020). We find, however, that in the special case where weather is known and only truck failures are uncertain, the chance constraint becomes convex, and the problem can be solved to arbitrary precision using a fast binary search procedure that exploits a certain exchange property of the set of feasible allocations (a "jump system" in the sense of Murota 2021). The efficiency of this technique allows it to be used as a heuristic for the general case of stochastic weather conditions. Essentially, we find a separate optimal allocation for each individual weather scenario (treating that scenario as given), evaluate these allocations under uncertain weather, and take the best. Although this approach does not guarantee an optimal solution to (1), it proves to be quite effective in practice.

In Problem (2), the main difficulty is the objective. Problems of this form are often approached using the sample average approximation method (Shapiro, Dentcheva, and Ruszczyński 2021), but this technique typically assumes that the underlying probability distribution does not depend on the decision variable (as a result, the decision maker can evaluate any feasible decision under the same set of sampled scenarios). In our setting, however, the support of ξ_r depends on x_r , making it difficult, if not impossible, to pregenerate samples of truck failures. For this reason, we devise a customized branch-and-bound algorithm to optimize the nonconvex objective directly. At each node, the proposed method employs linear programming to generate a discrete relaxation, which can be solved efficiently without resorting to integer programming, using the theory of optimization over jump systems. The method provably finds the optimal solution, even in the presence of stochastic weather.

We demonstrate the practical value of the proposed modeling and algorithmic approaches in a case study based on real data from northern Utah. First, the solutions obtained by our methods substantially outperform UDOT's current fleet allocation in both the risk-averse and risk-neutral settings. The solutions to (1) improve the maximum turnaround time by 22.5%–34.1%. Second, the solutions to (2) outperform UDOT's fleet allocation, with the improvements ranging between 51.6%

and 70.5%. Both types of solutions also significantly outperform a heuristic that allocates trucks proportionally to lane-miles, showing that network topology plays a vital role in performance. In the risk-averse setting, we find that explicitly modeling truck failures helps to reduce the incidence of extremely unfavorable outcomes, which not only carry social costs, but also can attract negative media attention. All of these results are made interpretable by examining the Utah network in key regions: essentially, our proposed allocations are better adapted to certain characteristics of the network topology (such as long isolated routes) that can cause significant disruptions in service quality in the event of truck failure. We also find that these improvements are easy to implement, requiring only a minor repositioning of trucks relative to current practice.

2. Literature Review

The four-part survey by Perrier, Langevin, and Campbell (2006a, b, 2007a, b) provides a comprehensive summary of existing work on the planning, design, and operations of winter road maintenance. We may roughly categorize this body of work into three principal areas—namely, design of regions, fleet allocation, and snowplow routing. These three classes of problems are closely connected: the size and network structure of a region obviously impacts the fleet size needed to serve it, while routing ensures that the fleet is being utilized effectively. However, the holistic integration of all three aspects is very intricate, and many studies resort to simplifications. For example, Kandula and Wright (1997) and Labelle, Langevin, and Campbell (2002), both of which focus on the design of regions, use various methods to estimate deadheading miles without fully solving the routing problem. Similarly, many studies on fleet allocation bypass a comprehensive solution to the routing problem, either by working with a fixed set of routes (Fu, Trudel, and Kim 2009; Hajibabai and Ouyang 2016) or by analytically solving stylized models (Chien, Gao, and Meegoda 2013; Abdel-Malek et al. 2014). In our case, the region design is fixed, but we fully integrate fleet allocation and routing, with the optimal (or near-optimal) values of the routing problems serving as cost functions used to evaluate allocations.

Those studies that do investigate routing tend to look at it in isolation. One stream of research uses mathematical programming formulations, such as arc routing (Dror 2012; Corberán et al. 2021) or constraint programming (Kinable, van Hoeve, and Smith 2021), with various practical constraints, such as road priorities (Perrier, Langevin, and Amaya 2008; Quirion-Blais, Langevin, and Trépanier 2017), echelon routing (Salazar-Aguilar, Langevin, and Laporte 2012), and left-turn penalty (Corbett et al. 2020). An array of heuristic algorithms (Dusault et al. 2013; Hajibabai et al. 2014; Liu et al. 2014;

Quirion-Blais, Trépanier, and Langevin 2015; Holmberg 2019) is available for different variants of the problem. See also Campbell et al. (2015) for a discussion of recent modeling and algorithmic advances. In our paper, the routing problem (given a fixed fleet size) is embedded inside the computation of Q_r and M_r , and our research focus is on the efficient allocation of trucks under these cost functions. The literature on snowplow routing could be viewed as complementary to our work because one could potentially make use of any existing method to perform the required offline computations.

In practice, weather conditions are often uncertain (Berrocal et al. 2010), posing a significant challenge for snowplowing operations. There have been several efforts to incorporate stochasticity and/or dynamic environmental conditions into operational models. For example, Fu, Trudel, and Kim (2009) and Xu and Kwon (2021) model multiple time periods and allow various parameters to be updated using simulation models or, potentially, real-time data. Hajibabai and Ouyang (2016) consider stochastic factors, such as snowstorm severity, duration, and impacted area, as well as the possibility of truck failure. These studies focus on managing a fixed fleet of trucks within a single region; as such, their models are all potentially compatible with our modular framework. We use vehicle routing problems (VRPs) to compute cost functions because UDOT expressed strong interest in improving routes and has only limited access to real-time data. In principle, however, one could also use these alternate models to compute costs under a given fleet allocation.

The literature on fleet sizing is less abundant in comparison. Chien, Gao, and Meegoda (2013) and Abdel-Malek et al. (2014) account for the uncertainty of snowfall intensity and duration and develop analytical models to determine the snowplowing fleet that should be contracted for the upcoming winter. Li, RazaviAlavi, and AbouRizk (2021) use a simulation-based approach to assist winter maintenance agencies with designing short-term plans for fleet sizing and routing. These studies make simplifying assumptions—for example, they view trucks as an unlimited resource, which is quite unrealistic in practice, as DOTs are severely constrained by fleet sizes and plow driver shortage (MoDOT 2022). Miller et al. (2018) determine optimal fleet sizes for three Ohio Department of Transportation districts under several different objectives. In their work, each district is treated individually, and the tradeoffs involved in allocating trucks between regions are not considered. Lastly, some studies (Jang, Noble, and Hutsel 2010; Sullivan et al. 2019) attempt to jointly consider multiple aspects of winter maintenance, including districting, fleet allocation, and routing, but these approaches are completely heuristic in nature, though they may yield useful insights in numerical experiments.

3. Methodology: Risk-Averse Formulation

Table 1 summarizes notation used in this and subsequent sections. This section focuses on Problem (1), which is approached in two steps. First, in Section 3.1, we assume that truck failures are the only source of uncertainty—that is, the weather conditions are known. We show how the problem can be solved to arbitrary precision in this setting. Second, in Section 3.2, we describe our model of weather uncertainty and discuss how the technique developed in Section 3.1 can be used as a heuristic to solve this more general case.

3.1. Efficient Solution Under Known Weather

For the time being, we assume that the values $Q_r(x_r)$ are fixed and known for all $r = 1, \dots, R$ and $x_r = 1, \dots, K$. In words, given that k trucks are available in region r , we know the turnaround time that will be achieved in that region. Thus, the weather conditions are fixed, and the only random variables in (1) are the numbers ξ_r of truck failures.

We begin by deriving an explicit form for the chance constraint in (1). Because Q_r is nonincreasing with non-negative range, for $t \in [Q_r(K), Q_r(K-1), \dots, Q_r(0)]$ for any r , define

$$Q_r^{-1}(t) \triangleq \inf\{k \in \mathbb{Z}^+ : Q_r(k) \leq t\},$$

to be the smallest number of trucks whose allocation to region r would result in a turnaround time below t . For convenience, let $Q_r^{-1}(t) = \infty$ if $t < Q_r(K)$ and $Q_r^{-1}(t) = -\infty$ if $t > Q_r(0)$. Furthermore, denote by F_r the cumulative distribution function (CDF) of the stochastic failure ξ_r . Note that, given $x = (x_1, \dots, x_R)$, $Q_r(x_r - \xi_r)$ is a discrete random variable with support $\{Q_r(x_r), \dots, Q_r(0)\}$. Then, for $t \in \mathbb{R}$, we have

$$\mathbb{P}\left(\max_r Q_r(x_r - \xi_r) \leq t\right) = \prod_{r=1}^R \mathbb{P}(Q_r(x_r - \xi_r) \leq t) \quad (3)$$

$$= \prod_{r=1}^R \mathbb{P}(x_r - \xi_r \geq Q_r^{-1}(t)) \quad (4)$$

$$= \prod_{r=1}^R F_r(x_r - Q_r^{-1}(t)),$$

where (3) holds from the independence of ξ_r and (4) holds by the definition of Q_r^{-1} .

As mentioned before, if we assume that each truck within a region is equally likely to fail, then ξ_r will follow a binomial distribution. In that case,

$$\begin{aligned} & \mathbb{P}\left(\max_{r=1, \dots, R} Q_r(x_r - \xi_r) \leq t\right) \\ &= \prod_{r=1}^R \sum_{m=0}^{x_r - Q_r^{-1}(t)} \binom{x_r}{m} p_r^m (1 - p_r)^{x_r - m}, \end{aligned} \quad (5)$$

Table 1. Notation Used Throughout the Paper

Notation	Meaning
R	Number of regions
r	Region index
K	Total number of trucks that can be allocated between R regions
x_r	Number of trucks assigned to region r
ξ_r	Number of failures among x_r trucks assigned to region r
F_r	Cumulative distribution function (CDF) of the random variable ξ_r
U	Number of snowfall weather scenarios
u	Weather scenario index
q_u	Probability of weather scenario u occurring
$p_{r,u}$	Probability of truck failure in region r under weather scenario u
α	Probability threshold used in risk-averse formulation
$Q_{r,u}(x_r)$	Turnaround time in region r with x_r trucks under weather scenario u
$Q_r^{-1}(t)$	Smallest number of trucks assigned to region r attaining a turnaround time below t
$M_{r,u}(x_r)$	Uncleaned lane miles in region r with x_r trucks under weather scenario u
T_r	Time threshold for region r used in the computation of $M_{r,u}$
$\tau^{(u)}$	Optimized turnaround time across R regions under weather scenario u
$x^{(u)}$	Optimized truck allocation under weather scenario u
$t^{(u)}$	Turnaround time achieved by $x^{(u)}$ under stochastic weather
t^*	Best turnaround time achieved by any $x^{(u)}$
$\Gamma(x_r)$	Expected value of $M_r(x_r - \xi_r)$
$\bar{\Gamma}(x_r)$	Greatest convex minorant of $\Gamma(x_r)$ with $x_r \in [1, K]$
$\bar{\Gamma}^{[a_r, b_r]}(x_r)$	Greatest convex minorant of $\Gamma(x_r)$ restricted to $x_r \in [a_r, b_r]$
x^*	Optimized truck allocation obtained from our methods
x_{NF}^*	Optimized truck allocation obtained without considering truck failure
x_{UDOT}	Truck allocation currently used by UDOT
x_{LM}	Truck allocation made proportionally to the number of lane miles in each region

and the chance constraint in (1) becomes

$$\prod_{r=1}^R \left(\sum_{m=0}^{x_r - Q_r^{-1}(t)} \binom{x_r}{m} p_r^m (1 - p_r)^{x_r - m} \right) \geq 1 - \alpha, \quad (6)$$

where we must have $x_r \geq Q_r^{-1}(t)$ for $r = 1, \dots, R$ to make the left-hand side well-defined; otherwise, the left-hand side of the chance constraint will equal zero. Next, we apply a log transform to both sides of (6) to obtain the reformulation

$$\begin{aligned} \min_{x \in \mathbb{Z}_{++}^R, t \geq 0} \quad & t \\ \text{s.t.} \quad & \sum_{r=1}^R x_r = K, \\ & \sum_{r=1}^R \log \left(\sum_{m=0}^{x_r - Q_r^{-1}(t)} \binom{x_r}{m} p_r^m (1 - p_r)^{x_r - m} \right) \geq \log(1 - \alpha), \\ & x_r \geq Q_r^{-1}(t), \quad r = 1, \dots, R. \end{aligned} \quad (7)$$

It can be observed that the objective function t is involved in a nonlinear term $Q_r^{-1}(t)$, which sets the lower bound to x_r and also appears in the upper limit of the second sum in the chance constraint. For a fixed value of t , one can check whether there exists an allocation x that would satisfy the chance constraint, by

solving

$$\begin{aligned} \psi^* = \max_{x \in \mathbb{Z}_{++}^R} \quad & \sum_{r=1}^R \log \left(\sum_{m=0}^{x_r - Q_r^{-1}(t)} \binom{x_r}{m} p_r^m (1 - p_r)^{x_r - m} \right) \\ \text{s.t.} \quad & \sum_{r=1}^R x_r = K, \\ & x_r \geq Q_r^{-1}(t), \quad r = 1, \dots, R. \end{aligned} \quad (8)$$

If $\psi^* \geq \log(1 - \alpha)$, this means that the chance constraint can be feasibly satisfied. Given our ability to solve (8) efficiently, we can find the *smallest* t for which (6) holds by using the binary search (Williams 1976) procedure outlined in Algorithm 1.

Algorithm 1 (A Binary Search Procedure for Problem (7))

Input: Predefined α and optimality tolerance ϵ .

Output: The ϵ -efficient t^* and x^* .

Step 1: Let $t_l = \max_r Q_r(K)$, and $t_u = \max_r Q_r(0)$. Let $k = 1$ and $\tau_k \leftarrow \frac{t_l + t_u}{2}$.

Step 2: Solve

$$\begin{aligned} \psi^* = \max_{x \in \mathbb{Z}_{++}^R} \quad & \sum_{r=1}^R \log \left(\sum_{m=0}^{x_r - Q_r^{-1}(\tau_k)} \binom{x_r}{m} p_r^m (1 - p_r)^{x_r - m} \right) \\ \text{s.t.} \quad & \sum_{r=1}^R x_r = K, \\ & x_r \geq Q_r^{-1}(\tau_k), \quad r = 1, \dots, R. \end{aligned}$$

Step 3: Repeat

If $\psi^* \geq \log(1 - \alpha)$, update $t_u \leftarrow \tau_k, k \leftarrow k + 1, \tau_k \leftarrow \frac{t_l + t_u}{2}$ and go to Step 2.

Otherwise, update $t_l \leftarrow \tau_k, k \leftarrow k + 1, \tau_k \leftarrow \frac{t_l + t_u}{2}$ and go to Step 2.

Until $(\psi^* \geq \log(1 - \alpha) \text{ and } t_u - t_l \leq \epsilon)$.

Step 4: Return τ_k and corresponding x .

To solve (8), we first show that ψ^* is the optimal value of a separable concave function. It is known (An 1997) that, for fixed n and p , the mapping

$$x_r \mapsto \log \left(\sum_{m=0}^{x_r} \binom{n}{m} p_r^m (1 - p_r)^{n-m} \right),$$

is (piecewise linear) concave in x_r . The summand in the objective of (8) differs mainly in that n is replaced by x_r . However, we can prove a similar property—that is, the term of interest is concave for $x_r \geq Q_r^{-1}(\tau)$. It directly follows that the objective of (8) is separable and concave, being a sum of R concave functions.

Proposition 1. *The function*

$$x_r \mapsto \log \left(\sum_{m=0}^{x_r - Q_r^{-1}(\tau)} \binom{x_r}{m} p_r^m (1 - p_r)^{x_r - m} \right),$$

is piecewise-linear concave for $x_r \geq Q_r^{-1}(\tau)$.

The feasible region of (8) is described by a hyperplane bounded from below. We show below that this set represents a jump system (a concept introduced by Bouchet and Cunningham 1995), which is a set of integer points with a certain exchange property, formalized using the following definitions.

Definition 1. For $x, y \in \mathbb{Z}^n$, a step s from x to y is a vector $s \in \{0, \pm 1\}^n$ such that $x + s \in [x, y]$ and $\|s\|_1 = 1$.

Definition 2. A jump system $J \subseteq \mathbb{Z}^n$ is a nonempty set satisfying either of the following two conditions:

1. If s is a step from x to y , then $x + s \in J$, or
2. There exist steps s and s' from x to y such that $x + s + s' \in J$ while $x + s \notin J$.

Proposition 2. *The set*

$$J^\tau = \left\{ x \in \mathbb{Z}_{++}^R \mid \sum_{r=1}^R x_r = K, x_r \geq Q_r^{-1}(\tau) \right\},$$

is a jump system.

Ando, Fujishige, and Naitoh (1995) proposed a class of greedy algorithms that provably find the optimal solution to separable convex minimization on jump systems. Shioura and Tanaka (2007) subsequently proved that such algorithms terminate in pseudo-polynomial time. Algorithm 2 gives an instance of such an algorithm for the problem of optimizing (8). The correctness and

running time of this algorithm follow directly from the theory of optimization over jump systems.

Algorithm 2 (Greedy Algorithm to Find the Optimal Solution to (8))

Input: A separable concave function $w(x) = \sum_{r=1}^R \log \left(\sum_{m=0}^{x_r - Q_r^{-1}(\tau)} \binom{x_r}{m} p_r^m (1 - p_r)^{x_r - m} \right)$, a finite jump system J^τ and a feasible solution $x^0 \in J^\tau$.

Output: The optimal solution x^* to (8).

Step 1: $x \leftarrow x^0$.

Step 2: If neither of the two conditions below is satisfied, then stop. Otherwise, go to Step 3.

1. There exists a step s such that $x + s \in J^\tau$ and $w(x + s) > w(x)$.
2. There exists steps s, s' such that $x + s \notin J^\tau, x + s + s' \in J^\tau$ and $w(x + s + s') > w(x)$.

Step 3: Compute

$$w_1 \leftarrow \max \{ w(x + s) \mid x + s \in J^\tau, w(x + s) > w(x) \}, \quad (9)$$

$$w_2 \leftarrow \max \{ w(x + s) \mid x + s \notin J^\tau, x + s + s' \in J^\tau, w(x + s + s') > w(x) \}, \quad (10)$$

where the maximum over the empty set is defined to be $-\infty$.

Step 4: $\hat{w} \leftarrow \max \{ w_1, w_2 \}$. If $\hat{w} = w_1$, let $\hat{s}$ be the step s that attains the maximum of (9), $x \leftarrow x + \hat{s}$ and go to Step 2. If $\hat{w} = w_2$, let $\hat{s}$ and $\hat{s}'$ be the steps s and s' that attains the maximum of (10), $x \leftarrow x + \hat{s} + \hat{s}'$ and go to Step 2.

It then follows that Algorithm 1 solves (7) near-optimally in pseudo-polynomial time. It is known (Sabharwal 2019) that binary search has a computational complexity of $\mathcal{O}(\log(W/\epsilon))$, where W is the width of the initial interval and ϵ is the error tolerance. Each iteration of binary search uses Algorithm 2 to solve an instance of Problem (8) in pseudo-polynomial time. We state this result below for completeness.

Proposition 3. *Algorithm 1 returns an ϵ -optimal solution to (7) in pseudo-polynomial time.*

3.2. Heuristic Approach for Uncertain Weather

Suppose now that weather conditions are uncertain at the time when we make the fleet allocation decision. We assume, however, that this uncertainty is described by a finite set of scenarios, indexed by $u = 1, \dots, U$. Each scenario has a probability q_u of happening. In practice, we use historical data on past snowfalls to define the scenarios and probabilities; see Section 5 for a concrete example.

We now denote by $Q_{r,u}(x_r)$ the turnaround time achieved by x_r vehicles in region r under weather scenario u . This models the direct impact of weather on performance: for example, trucks may be forced to

reduce their movement speed under more severe weather conditions, reducing the turnaround time. The distribution of ξ_r may also depend on u through the failure probability $p_{r,u}$ (in other words, the number of truck failures is now conditionally binomial given the weather). As before, we assume that the values $Q_{r,u}(x_r)$, for all r, u , and x_r , are available to us.

With these modifications, (6) becomes

$$\sum_{u=1}^U q_u \prod_{r=1}^R \left(\sum_{m=0}^{x_r - Q_{r,u}^{-1}(t)} \binom{x_r}{m} p_{r,u}^m (1 - p_{r,u})^{x_r - m} \right) \geq 1 - \alpha, \quad (11)$$

that is, we first condition on the weather scenario and then apply the results of Section 3.1. Unfortunately, the left-hand side of (11) is no longer log-concave, meaning that (1) can no longer be formulated as a convex integer program. To provide intuition for why this is the case, we recall a fact from the theory of continuous optimization (Boyd and Vandenberghe 2004, p. 107) that a probability can be generally guaranteed to be log-concave in the decision variable only when the underlying probability density is log-concave (unimodal). Although our problem is not continuous, we see a similar issue: the binomial distribution is unimodal, helping to explain why we saw log-concavity in Section 3.1, but the joint distribution of weather and truck failures need not be.

At the same time, although (11) is intractable from the point of view of finding the optimal x_r , it becomes easy to optimize over t when x_r is fixed. That is, any allocation can be made to satisfy the chance constraint: once we fix x_r , a simple binary search will yield the smallest t for which (11) holds. Thus, an allocation obtained through other means can still be accurately evaluated in the setting of stochastic weather.

This suggests a heuristic approach that leverages the developments of Section 3.1 for the known-weather setting. For each $u = 1, \dots, U$, we run Algorithms 1 and 2 with $Q_{r,u}$ and $p_{r,u}$ as inputs. Essentially, we find an optimal allocation for each weather scenario individually. Letting $x^{(u)}$ and $\tau^{(u)}$ be the output of Algorithm 1 for scenario u , we simply discard $\tau^{(u)}$. Instead, we plug $x^{(u)}$ into (11) and compute a different value $t^{(u)}$ that achieves $(1 - \alpha)$ -level coverage under stochastic weather, rather than only under scenario u . We then let $u^* = \arg \min_u t^{(u)}$ and return $x^* = x^{(u^*)}$ as the recommended allocation. The quantity $t^* = t^{(u^*)}$ is precisely the turnaround time sustained by x^* with probability $1 - \alpha$ in the stochastic-weather setting.

The main idea is that, although the number $\tau^{(u)}$ returned by Algorithm 1 may be infeasible under stochastic weather, the actual resource allocation $x^{(u)}$ can still straightforwardly be made to satisfy the correct chance constraint. Depending on the specific set of scenarios that we are given, some u may be natural choices

for risk-averse optimization—for example, if there is a single scenario with the “most severe” weather. In general, however, we can simply solve for each u and take the best allocation. Although this approach can be computationally costly, the cost is also easy to distribute because we can run different u in parallel.

Our practical experience (described in Section 5) has been that many realistic weather scenarios are quite similar (for example, the weather conditions in two scenarios may differ in only one or two regions). Thus, an allocation that is optimal for scenario u will also perform very well for other scenarios that are similar to u . Our heuristic approach was thus able to obtain significant performance improvements over current practice.

4. Methodology: Risk-Neutral Formulation

We now consider Problem (2), assuming uncertain weather from the beginning. As in Section 3.2, we are given weather scenarios $u = 1, \dots, U$, each with a probability q_u of happening. We denote by $M_{r,u}(x_r)$ the lane miles that have not been cleaned by x_r vehicles in T_r time units, conditional on weather scenario u . As before, we let $p_{r,u}$ be the truck failure probability in each scenario.

The objective function may be computed more explicitly as

$$\mathbb{E} \left[\sum_{r=1}^R M_r(x_r - \xi_r) \right] = \sum_{r=1}^R \sum_{u=1}^U q_u \sum_{k=0}^{x_r} M_{r,u}(x_r - k) \cdot \binom{x_r}{k} p_{r,u}^k (1 - p_{r,u})^{x_r - k}, \quad (12)$$

where we condition on the weather scenario as in (11). Now, (2) is given as an integer program with a nonlinear and nonconvex objective. Such problems are commonly tackled using branch and bound (B&B) methods (Burer and Letchford 2012). The performance of a B&B algorithm depends on its ability to generate good lower bounds, typically done by solving *continuous* relaxations of the original problem. For example, lower bounds for mixed integer programs are generated by solving linear programming relaxations (Wolsey 2020), while similar continuous relaxations are also employed in spatial B&B methods to tackle integer programs with *nonlinear* objectives (Smith and Pantelides 1999; Gerard, Köppe, and Louveaux 2017). This idea, however, is not applicable to our setting. Because x_r appears in the binomial coefficients in (12), our objective function is not defined for noninteger x_r , which prevents us from leveraging continuous relaxations of the problem. Instead, we design custom *discrete* relaxations, which can be solved efficiently. These are discussed in Sections 4.1 and 4.2.

Section 4.3 gives a full statement of the algorithm, and Section 4.4 discusses computational cost.

4.1. Initial Lower Bound

Greatest convex minorants are a natural approach to obtaining a lower bound. For notational convenience, let

$$\Gamma_r(x_r) = \mathbb{E}[M_r(x_r - \xi_r)] = \sum_{u=1}^U q_u \sum_{k=0}^{x_r} M_{r,u}(x_r - k) \cdot \binom{x_r}{k} p_{r,u}^k (1 - p_{r,u})^{x_r - k},$$

denote the summand in the objective (12). Next, consider the greatest convex minorant of $\Gamma_r(x_r)$, defined as

$$\check{\Gamma}_r(x_r) = \sup\{\Phi_r(x_r) : \Phi_r \text{ is piecewise convex and } \Phi_r(x_r) \leq \Gamma_r(x_r) \text{ for } x_r \in [1, K]\}.$$

Because $\check{\Gamma}_r(x_r)$ is additionally nonincreasing and represents a K -vector (see illustration in Figure 2(a)), it can be easily computed as the solution to a linear program:

$$\begin{aligned} \min_{\check{\Gamma}_r \in \mathbb{R}_+^K} \quad & \sum_{k=1}^K (\Gamma_r(k) - \check{\Gamma}_r(k)) \\ \text{s.t.} \quad & \check{\Gamma}_r(k) \leq \Gamma_r(k), \quad k = 1, \dots, K, \\ & \check{\Gamma}_r(k) - \check{\Gamma}_r(k+1) \geq \check{\Gamma}_r(k+1) - \check{\Gamma}_r(k+2), \\ & \quad k = 1, \dots, K-2, \end{aligned} \quad (13)$$

where the last set of constraints ensures convexity. Repeating Proposition 2, we find that

$$\min_{x \in \mathbb{Z}_{++}^R} \left\{ \sum_{r=1}^R \check{\Gamma}_r(x_r) \mid \sum_{r=1}^R x_r = K \right\},$$

is an instance of separable convex minimization on a

jump system, and thus solvable in pseudo-polynomial time using a greedy algorithm similar to Algorithm 2. The solution of this problem provides a lower bound on the optimal value of (2).

4.2. Branching

Let $\bar{x}$ denote the optimal solution to the convex relaxation. We select the region

$$r^* = \arg \max_r \{\Gamma_r(\bar{x}_r) - \check{\Gamma}_r(\bar{x}_r)\},$$


with the greatest discrepancy between the original and relaxed objective, and branch on the variable x_{r^*} to create two subsets of the search space with tighter relaxations over more restricted intervals. To obtain such relaxations, let

$$\check{\Gamma}_r^{[a_r, b_r]}(x_r) = \begin{cases} \sup\{\Phi_r(x_r) : \Phi_r \text{ is piece-wise convex and } \Phi_r(x_r) \leq \Gamma_r(x_r)\}, & \text{if } x_r \in [a_r, b_r] \\ \Gamma_r(x_r), & \text{otherwise,} \end{cases}$$

denote the greatest minorant of $\Gamma_r(x_r)$, which is convex only over the domain $[a_r, b_r]$ of interest, and otherwise coincides with $\Gamma_r(x_r)$ (see Figure 2(b) for an illustration). This minorant can be computed with a linear program similar to (13), where the last set of convexity constraints will be specified only for $k \in [a_r, b_r]$. Because convexity constraints are relaxed over the remaining domain, the resulting minorant will provide a tighter relaxation over the interval $[a_r, b_r]$ of interest.


When branching on variable x_{r^*} , we create two nodes with search spaces restricted to $x_{r^*} \in [1, \bar{x}_{r^*}]$ and $x_{r^*} \in [\bar{x}_{r^*} + 1, K]$. At the former node, we tighten the convex relaxation by replacing $\check{\Gamma}_{r^*}^{[1, K]}(x_{r^*})$ with $\check{\Gamma}_{r^*}^{[1, \bar{x}_{r^*}]}(x_{r^*})$ in the objective. At the latter node, $\check{\Gamma}_{r^*}^{[1, K]}(x_{r^*})$ is replaced with $\check{\Gamma}_{r^*}^{[\bar{x}_{r^*} + 1, K]}(x_{r^*})$. The described branching procedure

Figure 2. (Color online) Examples of $\Gamma_r(x_r)$, $\check{\Gamma}_r(x_r)$, and $\check{\Gamma}_r^{[6, 12]}(x_r)$ for a Single Region



Notes. (a) $\Gamma_r(x_r)$ and $\check{\Gamma}_r(x_r)$ over the whole interval. (b) $\check{\Gamma}_r(x_r)$ over a restricted interval.

Figure 3. (Color online) An Illustrative Example of Branching on Variable x_{r^*}



Notes. (a) In this example, we obtain the optimal solution $\bar{x}$ to the relaxed problem and identify r^* as the index for which the difference between $\Gamma_r(\bar{x}_r)$ and $\check{\Gamma}_r(\bar{x}_r)$ is maximized. (b) and (c) We then branch variable x_{r^*} at point $\bar{x}_{r^*} = 5$ and generate two child nodes where the search space of x_{r^*} is restricted to $[1, \bar{x}_{r^*}]$ and $[\bar{x}_{r^*} + 1, 12]$, respectively. (a) $\Gamma_{r^*}(5) \neq \check{\Gamma}_{r^*}(5)$. (b) $x_{r^*} \in [1, 5]$. (c) $x_{r^*} \in [6, 12]$.

is illustrated in Figure 3. As we continue to branch and construct (partially) convex minorants over more and more restricted domains, we will obtain tighter and tighter relaxations. This ensures that the gap between the (partially) convex relaxation and the original non-convex function decreases monotonically.

Algorithm 3 (Branch-and-Bound Procedure for problem (2))

Initialization: Set the bounds for x_r to $a_r = 1, b_r = K$.

Denote the root node with a relaxation by $\mathcal{P}^0$.

Update the node set $\Omega \leftarrow \{\mathcal{P}^0\}$.

Solve the relaxation on $\mathcal{P}^0$ using Algorithm 4 in the e-companion to obtain the optimal solution $\bar{x}^0$ with LB^0 and UB^0 . Update $x^* \leftarrow \bar{x}^0, \hat{UB} \leftarrow UB^0$.

while $\Omega \neq \emptyset$ **do**

Select the node $\mathcal{P}^* \in \Omega$ with the smallest lower bound LB^* and solution $\bar{x}$ to its relaxation,

$$\min_{a_r \leq x_r \leq b_r} \left\{ \sum_r^R \check{\Gamma}_r^{[a_r, b_r]}(x_r) \mid \sum_{r=1}^R x_r = K \right\}, \quad (14)$$

and update $\Omega \leftarrow \Omega \setminus \{\mathcal{P}^*\}$.

if there exists r such that $\check{\Gamma}_r^{[a_r, b_r]}(\bar{x}_r) \neq \Gamma_r(\bar{x}_r)$ **then**

Find the region

$$r^* = \arg \max_r \{\Gamma_r(\bar{x}_r) - \check{\Gamma}_r(\bar{x}_r)\};$$

Branch on $\bar{x}_{r^*}$ and create two nodes:

$\mathcal{P}'$ with a relaxation that coincides with (14) with $b_{r^*} = \bar{x}_{r^*}$

$\mathcal{P}''$ with a relaxation that coincides with (14) with $a_{r^*} = \bar{x}_{r^*} + 1$

for $\mathcal{P}^s$ in $\{\mathcal{P}', \mathcal{P}''\}$ **do**

Solve the relaxation on $\mathcal{P}^s$ using Algorithm 4 in the e-companion to obtain the optimal solution $\bar{x}^s$ with its LB^s and UB^s .

if $LB^s < \hat{UB}$ **then**

Update $\Omega \leftarrow \Omega \cup \{\mathcal{P}^s\}, x^* \leftarrow \bar{x}^s, \hat{UB} \leftarrow \min\{\hat{UB}, UB^s\}$.

else

Discard $\mathcal{P}^s$.

Return the solution x^* as the optimal solution to (2).

With the proposed branching rule, the relaxation solved at each node of the B&B tree is of the form

$$\min_{a_r \leq x_r \leq b_r} \left\{ \sum_r^R \check{\Gamma}_r^{[a_r, b_r]}(x_r) \mid \sum_{r=1}^R x_r = K \right\}. \quad (15)$$

Every such problem represents a convex-separable minimization on a jump system, which is specified with a hyperplane, as well as upper and lower bounds on x_r . As a result, similarly to Problem (8), it can be solved in pseudo-polynomial time using a greedy algorithm, which is provided in the e-companion for completeness.

4.3. Outline and Correctness

Finally, we state the outline of the entire B&B method in Algorithm 3. Specifically, the relaxation at each node is solved with Algorithm 4 in the e-companion to obtain $\bar{x}$ and LB . Plugging $\bar{x}$ into (12) yields the upper bound UB for the node.

If the obtained LB is less than the best upper bound found thus far, $\hat{UB}$, then we look for the region r^* with the highest discrepancy between the original and relaxed objective, and we branch on x_{r^*} to create two new nodes. The procedure is repeated while pruning nodes that cannot produce a better solution than the incumbent (e.g., the nodes where $LB > \hat{UB}$). After all remaining nodes have met the termination condition and cannot be further split, the allocation with the least UB is returned. Thus, we are able to solve Problem (2)

exactly, even when both weather and truck failures are uncertain.

Proposition 4. *Algorithm 3 returns the optimal solution to Problem (2).*

4.4. Computational Performance of the Branch-and-Bound Method

In our experience, the proposed branch-and-bound method is often able to find the optimal solution in a very small number of iterations. We illustrate this behavior on an instance with $R = 12$ regions. The cost functions in this instance are based on our case study in Section 5, but for illustration purposes, we simplify the problem so that there is only one weather scenario, and $p_r = p$ for all regions r . Figure 4(a) reports the total number of convex relaxations solved by B&B for different values of p . For comparison, the figure also presents the number of convex programs that need to be solved via a simple brute-force partition, involving $\prod_{r=1}^R c_r$ problems of the form (15), with c_r being the number of subintervals on which Γ_r is locally convex. We see that, out of thousands of possible convex programs, the proposed method generally solves no more than 10. For example, when $p = 0.04$, the B&B method only solves 9 convex relaxations out of a possible 10,000. When $p \in \{0.10, 0.25\}$, B&B finds an optimal solution by solving only one convex relaxation in each instance.

Figure 4(b) shows the B&B tree generated by the algorithm for the instance when $p = 0.06$. In this case, there are only five nodes. The problem can be solved quickly because the greatest convex minorants provide tight relaxations. Figure 4(c) shows (for different choices of p) the gaps between the best lower and upper bounds at each level of the tree; even at the root nodes, the gaps range only between 0.5% and 7.1%. Thus, the nonconvexity of the cost functions does introduce some error into the solutions, but the relaxation captures enough of the

overall shape of the functions that we are able to make reasonably good tradeoffs between different curves.

5. Case Study: Allocation of Snowplow Trucks in Northern Utah


We demonstrate the practical value of our framework in a case study based on data from northern Utah. Section 5.1 describes the infrastructure and weather data used to compute cost functions in different scenarios. Sections 5.2 and 5.3 present results for problems (1) and (2), respectively. Section 5.4 compares the two solutions to each other, illustrating the tradeoff between risk-averse and risk-neutral allocation. Finally, Section 5.5 discusses practical implementation of the proposed allocations through fleet repositioning.

5.1. Infrastructure and Weather Data

In this application, we have $K = 61$ single-wing plows allocated over $R = 12$ regions. For each region, we extract detailed infrastructure data about the number of lanes and travel speed on each road. Taking the Wellsville region as an example (see Figure 5), the number of lanes on each road segment determines the minimum required number of traverses, which is the key input for optimizing snowplowing routes. We extract truck trajectories from automated vehicle location data to reconstruct UDOT's plowing routes and validate the required number of traverses for each road segment, as well as travel speeds.


Weather conditions, in each region, can be categorized into three levels (Ye 2009) of snowfall intensity: light (0.06–0.1 inches/hour (in/hr)), moderate (0.1–0.5 in/hr), and heavy (above 0.5 in/hr). The National Centers for Environmental Information furnish daily snowfall records across the United States through monitoring stations (NCEI 2023). From this database, we extracted snowfall data for the 12 Utah regions considered in our study, recorded between 2019 and 2022 during the

Figure 4. (Color online) Computational Efficiency of B&B Framework and the B&B Tree with $p = 0.06$, as well as Convergence of Lower Bound and Upper Bound



Notes. (a) # of convex programs solved. (b) The B&B tree with $p = 0.06$. (c) Convergence of UB and LB.

Figure 5. (Color online) An Illustrative Example of Detailed Infrastructure and Vehicle Location Data for Wellsville Region



Notes. (a) Roadways in Wellsville. (b) Travel speeds. (c) Trajectory data.

months of January, February, March, April, November, and December. After data cleaning and removal of missing entries, we were left with snowfall information for 195 distinct days. These 195 instances of snowfall can be systematically classified into 55 distinct scenarios, each of which represents a unique combination of weather conditions for each region (thus, weather is *not* independent across regions).

Figure 6 illustrates the snowfall levels across regions for each scenario, as well as the frequency with which each scenario was observed in the data. For ease of visualization, the scenarios are numbered in order of decreasing frequency; thus, for example, 8% of historical snowfall instances followed scenario 1, 4% of instances followed scenario 2, etc. We see that scenario 1 is one of the mildest in terms of snowfall intensity, whereas scenarios 2 and 3 are the most severe. Although there is clearly some variation, most scenarios involve a combination of light and moderate snowfall levels, and many of these combinations are quite similar: for example, scenarios 26–33 and 44–51 have similar conditions in many regions.

To compute $Q_{r,u}(x_r)$ and $M_{r,u}(x_r)$ for each region r and weather scenario u , we transform the snowplow

routing problem into a vehicle routing problem and run routing heuristics repeatedly to minimize the turn-around time and uncleaned lane miles, given varying fleet sizes. A general formulation of this problem is given in the e-companion for completeness, but a high-level understanding is sufficient for the present discussion. As shown in Figure 7, each arc to be cleaned is augmented with an additional node. The resulting vehicle routes, which visit all of these nodes, are guaranteed to clean all road lanes. Travel costs are calculated based on assumed vehicle speeds. Many studies have found that speed is affected by weather severity; for example, Agarwal, Maze, and Souleyrette (2005) estimated a 7%–9% reduction in speed (relative to the speed limit) under light snowfall and 11%–15% under heavy snowfall. A Utah-based study by Schultz et al. (2022) reported a broader range of reductions, from as low as 3% under light snowfall to as high as 40% under heavy snowfall. Taking these diverse results into account, we assume that plowing speed is reduced by 5%, 15%, and 30% of the speed limit under light, moderate, and heavy snowfall levels, respectively.

Thus, the values of $Q_{r,u}(x_r)$ and $M_{r,u}(x_r)$ are computed by solving the routing problem, for various fleet

Figure 6. (Color online) The Combination of Snowfall Intensity Across 12 Regions Within Each Individual Scenario (a) and the Likelihood of Each Snowfall Scenario Occurring (b)






Figure 7. (Color online) Example Transformation for Wellsville, UT

Notes. Based on the road network information (a), we create an instance of the arc-routing problem (b), in which the nodes represent intersections, while the arrows denote lanes between adjacent intersections. The corresponding node-routing problem is obtained using a standard transformation (Pearn, Assad, and Golden 1987; Hajibabai et al. 2014), where “delivery” nodes are inserted in the middle of each arc (c). (a) Wellsville network. (b) Arc routing problem. (c) Node routing problem.

sizes x_r , under the travel speeds prescribed by scenario u for region r (as seen in Figure 6). These cost functions are shown in Figure 8. For each r , M_r is the number of uncleaned lane miles after $T_r = 90$ minutes, which is considered an excellent level of service in providing access to important locations (Utah Winter Maintenance 2021). For most regions, $M_r(x_r)$ is zero for large enough x_r , but this is not always true because the network topology may make it impossible to clean every road by time T_r .

The total number of routing problems to be solved is fairly large: in theory, we would need to consider every possible combination of region, fleet size, and snowfall intensity. In practice, every cost curve in Figure 8 has

flattened out for fleet sizes of 12 or more trucks, so the actual number of problems to solve is smaller than $3 \cdot 12 \cdot R$. Furthermore, the difficulty of these problems varies widely between regions. Table 2 presents computation times for three representative regions (smallest, largest, and medium-sized) and three fleet sizes. We do not differentiate by weather scenario in this table because weather only affects the travel costs, not the problem size. Computational cost tends to decrease with the fleet size because it is easier to satisfy all of the demand with more trucks. Some of the individual problems can be solved in 5 seconds or less, whereas others require upward of 30 minutes. Because these problems

Figure 8. (Color online) Q_r and M_r Are Computed by Running Vehicle Routing Heuristics, Given Varying Fleet Sizes Under Light, Moderate, and Heavy Snowfalls

Notes. (a) Q_r for region r given x_r operational trucks light moderate heavy. (b) M_r for region r given x_r operational trucks.

Table 2. Computation Times for Several Representative (r, x_r) Combinations

Region (ID)	Fleet size	Computation time (s)
Smallest (5)	2	5.6
Smallest (5)	4	2.1
Smallest (5)	6	0.5
Medium (2)	2	103
Medium (2)	4	173
Medium (2)	6	85
Largest (3)	2	1,739
Largest (3)	4	1,021
Largest (3)	6	833

Note. All computation times were obtained on a machine with an Intel® Core™ i7-10510U CPU and 16 GB of RAM.

can be easily solved in parallel on different machines, all of the required offline computation could conceivably be done in one day.

5.2. Truck Allocation Based on Problem (1)

As described in Section 3.2, we obtain a solution x^* to (1) by running Algorithm 1 for each scenario u . For each of the 55 allocations thus obtained, we compute the shortest turnaround time that it can sustain with probability $1 - \alpha$, in the uncertain-weather setting, using (11). We then let x^* be the allocation with the shortest turnaround time among these.

We set the optimality tolerance in Algorithm 1 to $\epsilon = 0.001$ minutes. The failure probabilities $p_{r,u}$ are set to 0.03, 0.06, and 0.09 if the snowfall intensity in region r under scenario u is light, moderate, or heavy, respectively (thus, vehicles are more likely to fail in more severe weather). In the following, we consider a variety of α values. Although x^* changes depending on α , we found that scenarios 5 and 8 consistently produced the best-performing allocations. Referring back to Figure 6, we see that scenario 5 has moderate snowfall intensities in every region, whereas scenario 8 is mostly moderate,


with a few cases of light snowfall. Thus, the most robust allocations are obtained by assuming conditions that are somewhat worse than average, but not the worst possible (scenario 2, which has the most severe weather, produces overly conservative results).

Figure 9(a) reports the final allocations x^* and their corresponding turnaround times t for different α . Three benchmarks are considered: the allocation x_{UDOT} currently used by UDOT; a simple heuristic x_{LM} (“lane miles”) that allocates trucks proportionally to the number of lane miles in each region; and a modification of x^* denoted by x_{NF}^* (“no failures”) that ignores the possibility of truck failure. This last allocation is computed in the same manner as x^* , by individually optimizing with respect to each scenario and taking the best, but with the additional assumption that $p_{r,u} \equiv 0$. Thus, x_{NF}^* does account for weather uncertainty, but not for vehicle reliability, allowing us to assess the significance of considering both types of uncertainty together. Turnaround times for all three benchmarks are obtained using (11).

We find that x^* significantly reduces the maximum turnaround time compared with x_{UDOT} , for all α , with improvements ranging between 22.5% and 34.1%. For example, UDOT’s allocation can sustain a maximum turnaround time of 186.8 minutes with a probability of 0.95, whereas x^* can sustain 144.6 minutes (an improvement of 22.5%) with the same guarantee. Similarly, x^* surpasses x_{LM} consistently across the entire spectrum of α values, yielding enhancements ranging from 9.1% to 24.1%. The performance of x^* and x_{NF}^* is much closer, owing to the fact that the allocations themselves are often similar, but x^* still produces improvements for 9 out of 10 distinct values of α , ranging up to 8.8%. We will revisit this comparison further down, showing other ways in which x^* is preferable to x_{NF}^* .

It is illustrative to examine the overall performance distribution for each allocation, rather than just the objective value. Consider the case $\alpha = 0.10$. For each

Figure 9. (Color online) Performance of ϵ -Optimal Allocations Based on Problem (1) with Different α



Notes. (a) ϵ -optimal allocations based on (1). (b) Comparing different allocations.


snowfall scenario u , we generate $500,000 \cdot q_u$ truck breakdown situations from binomial distributions whose means depend on u and the allocations x_{UDOT} , x_{LM} , x_{NF}^* , and x^* . For every such simulated situation, we record the turnaround time if it exceeds $t^* = 128.9$ minutes, which is the objective value achieved by x^* under $\alpha = 0.10$. We then report the frequencies of such situations, out of the total sample of 500,000, in Figure 10. For the optimal allocation x^* , the total weight of these situations is close to 0.1, as guaranteed by the model. For x_{UDOT} , x_{LM} , and x_{NF}^* , it is, respectively, 1, 0.32, and 0.11. We can immediately see that x_{UDOT} has a severe shortage of vehicles in region 2, whereas x_{LM} has a shortage in regions 5 and 6. On the other hand, x^* and x_{NF}^* differ only by a single vehicle, which x^* diverts from region 8 to region 4. The overall improvement is small, because a reduction in turnaround time for region 4 is offset by an increase in region 8, but there is, nonetheless, a net gain.

To provide further insight into these comparisons, Figure 11, (a) and (b) show the roadways within regions 2 and 6, respectively. It is clear that both regional networks suffer from poor connectivity. A single truck failure will lead to a significant increase in turnaround time

because one of the remaining trucks will then have to navigate and clear multiple isolated “arms” of the network. For example, in region 2, the turnaround time for x_{UDOT} (under $\alpha = 0.10$) will jump up to 167.1 minutes, even under light snowfall, and 186.8 minutes under moderate snowfall. To mitigate this risk, x^* assigns three additional trucks to region 2. In this way, network topology can force us to allocate additional resources, even to relatively small regions. The benchmark x_{LM} , which is calculated solely based on lane miles, assigns only two trucks to region 6, but the lack of connectivity in the network makes this insufficient.

Figure 12 shows the roadways within regions 3, 4, and 8. As we saw in Figure 10(d), region 3 has the highest frequency of excessive turnaround times, despite the fact that (under $\alpha = 0.10$) x^* assigns a substantial fleet of eight trucks to it. This may be due to the fact that (from Figure 6) this region also has the highest probability of heavy snowfall (10%). At the same time, even under severe weather conditions, a single truck failure in region 3 will only increase the turnaround time by 10 minutes, whereas for region 4, there will be a more substantial escalation of 20.4 minutes. For this reason,

Figure 10. (Color online) Frequency of Each Region in Sampling Experiments for Problem (1) with $\alpha = 0.10$ and $(p_l, p_m, p_h) = (0.03, 0.06, 0.09)$



Notes. Resulting in $t > 128.9$ for (a) x_{UDOT} , (b) x_{LM} , and (c) x_{NF}^* , as well as (d) x^* . Note that the vertical axis of (a) and (b) has a different scale from the others.

Figure 11. (Color online) Responsible Roadways Within Regions 2 and 6



Notes. (a) Responsible roadways within region 2. (b) Responsible roadways within region 6.


when x^* moves one of the trucks assigned by x_{NF}^* to the reasonably well-connected region 8, it is preferable to divert this additional vehicle to region 4, rather than 3.

Let us now compare x^* and x_{NF}^* more closely in the case $\alpha = 0.05$, where the two allocations achieve nearly identical turnaround times. Despite this similarity, the two allocations produce visibly different performance distributions. Figure 13(a) shows the empirical distribution of maximum turnaround times greater than $t^* = 144.6$ minutes from a sample of 500,000 scenarios (in other words, this is the conditional distribution of turnaround time, given that it exceeded 144.6). It is clear that x^* is more robust, with a smaller median and lower concentration above it. The difference between the two allocations (as seen in Figure 13, (b) and (c)) is that x^* removes two trucks from region 8 and one truck from each of regions 1 and 3 and adds one truck to regions 6, 7, 9, and 11. We have already seen that region 6 is vulnerable to truck failures; as for regions 1, 7, 9, and 11, their road networks are shown in Figure 14, (a)–(d). Within region 7, a segment of Highway 91 contains five

lanes in both directions. If all three trucks allocated by x_{NF}^* are operational, one truck would travel several times along US-91 to clean all the lanes, leaving the remaining roadways to the other two trucks. A single failure would increase turnaround time by approximately 45 minutes; with the fourth truck added by x^* , the cost of one failure is reduced to 17 minutes. Similarly, both regions 9 and 11 contain at least one long roadway off to the side: thus, taking region 11 as an example, the cost of one truck failure is 53.8 minutes under x_{NF}^* , even under light snowfall, but only 17.6 minutes under x^* . Increasing the fleet in these regions does not completely eliminate the risk of outliers (because it is also possible for two trucks to fail), but it will reduce costs in “typical” situations.


Finally, we remark on the cost of computing x^* . The average cost of finding an optimal allocation for a single scenario u is 3.29 seconds across all α values. The problem appears to be easier for large α values, so the average cost per scenario is 2.63 seconds for $\alpha = 0.5$ and 5.61 seconds for $\alpha = 0.05$. Even in the worst case,

Figure 12. (Color online) Responsible Roadways Within Regions 3, 4, and 8



Notes. (a) Region 3. (b) Region 4. (c) Region 8.

Figure 13. (Color online) Robustness Comparison and Analysis Between x^* and x_{NF}^* Based on Problem (1) with $\alpha = 0.05$ and $(p_l, p_m, p_h) = (0.03, 0.06, 0.09)$



Notes. (a) Distribution of $t > 144.6$ for x_{NF}^* and x^* . (b) Frequencies of the seven regions resulting in turnaround time greater than 144.6 minutes under x_{NF}^* . (c) Frequencies of the seven regions resulting in turnaround time greater than 144.6 minutes under x^* .


however, we can find allocations for all 55 scenarios in approximately five minutes (or less, if we compute them in parallel).

5.3. Truck Allocation Based on Problem (2)

We implement the proposed B&B framework to solve (2). Unlike the risk-averse formulation, problem (2) does not have a tunable parameter, so we vary the level of risk by considering four settings for the truck failure


probabilities. Each setting consists of three possible values of $p_{r,u}$ corresponding to light, moderate, and heavy snowfall. These numbers are 0.02, 0.04, and 0.06 for setting 1; 0.03, 0.06, and 0.09 for setting 2; 0.04, 0.08, and 0.12 for setting 3; and 0.05, 0.10, and 0.15 for setting 4. Figure 15(a) provides optimal allocations x^* for different failure probability settings, whereas Figure 15(b) compares them against x_{UDOT} , x_{LM} , and x_{NF}^* , computed analogously to the previous setting. The expected

Figure 14. (Color online) Responsible Roadways Within Regions 1, 3, 9, and 11



Notes. (a) Region 1. (b) Region 7. (c) Region 9. (d) Region 11.

Figure 15. (Color online) Performance of Allocations Based on Problem (2) with Varying Failure Probability Settings



Notes. (a) Optimal allocations based on (2). (b) Comparing different allocations.

performance of each allocation is evaluated in the stochastic-weather setting using (12).

Most of the interesting characteristics of individual regions have already been discussed in Section 5.2, so we do not repeat them here, but we briefly summarize the results. As before, x^* consistently outperforms the three benchmarks. Service quality is improved by 51.6%–70.5% over x_{UDOT} and 24.8%–47.3% over x_{LM} . The improvement over x_{NF} ranges from 2.7% to 8.2% (and increases with the failure probability). Overall, however, the risk-neutral setting is less sensitive to truck failures than the risk-averse setting. In problem (2), any improvement in one region is partially offset by a reduction in another, whereas the risk-averse formulation rewards us for eliminating large outliers.

As in Section 5.2, we remark on computational cost. Problem (2) turns out to be easier to solve than (1), with all four settings requiring less than one second for the B&B method to return the optimal solution. In fact, under settings 1, 2, and 4, the optimal solution was found on the root node, whereas setting 3 required

solving three relaxations. Interestingly, the presence of weather uncertainty appears to have the effect of smoothing out some of the nonconvexity of the cost curves, as these instances require fewer iterations than the single-scenario example of Section 4.4.

5.4. Comparison of Optimal Solutions for (1) and (2)

It is also useful to compare the optimal allocations for Problems (1) and (2) by assessing the performance of one solution using the objective of the other model. Because of space considerations, we present results for two of the four settings (namely, settings 2 and 3) defined in Section 5.3 for the truck failure probabilities. These comparisons are shown in Figures 16 and 17. Note that (1) has a different solution for each value of α , whereas (2) has only a single optimal solution per failure setting.

We find that the risk-averse allocations, when evaluated in terms of expected uncleaned lane miles, consistently increase the objective by 25%–86% (under setting 2) and 22%–76% (under setting 3) relative to optimal.

Figure 16. (Color online) Evaluation of x^* to (1) and (2) in Uncleaned Lane Miles (a) and Turnaround Time (b) Under Setting 2 of the Truck Failure Probabilities





Figure 17. (Color online) Evaluation of x^* to (1) and (2) in Uncleaned Lane Miles (a) and Turnaround Time (b) Under Setting 3 of the Failure Probabilities


Because these ranges are fairly wide, one could potentially use such an analysis to select α if the decision maker wanted to partially accommodate the second objective. On the other hand, when the optimal solution to (2) is evaluated in terms of maximum turnaround time, the objective value is increased by 13%–25% (under setting 2) and 10%–20% (under setting 3) relative to the best risk-averse solution. There is thus a clear tradeoff between the two objectives. If the risk tolerance parameter α is high, the risk-neutral solution can perform reasonably well in both settings, but if the decision maker is greatly concerned about outliers, the risk-averse solution will be preferable.

5.5. Fleet Repositioning

In the following, we show that our improved truck allocations can be implemented at minimal cost to UDOT, simply by repositioning a very small number of trucks.

This is done by solving a transportation problem to minimize the total travel distance needed to transform UDOT's current allocation into the proposed allocations. The travel distance between any two maintenance stations is the network-based driving distance queried from Google Maps. Because of space considerations, we only consider setting 2 of the truck failure probabilities in this discussion.

Figure 18(a) presents the least expensive repositioning strategy to transform x_{UDOT} into allocation x^* based on Problem (1) with $\alpha = 0.10$. This is achieved by moving only 4 out of 61 trucks: region 2 will receive one and two trucks from regions 6 and 10, respectively, and region 6 will provide one truck to region 3. Similarly, Figure 18(b) presents the least expensive repositioning strategy to transform x_{UDOT} into allocation x^* based on Problem (2). This requires us to move 7 out of 61 trucks: region 10 will send one truck to region 12; region 6 will

Figure 18. (Color online) The Least-Expensive Truck Repositioning Strategies for Fleet Reassignment Based on Problem (1) with $\alpha = 0.10$ and Problem (2), Under Setting 2 of the Truck Failure Probabilities

Notes. (a) Transforming x_{UDOT} into x^* based on (1). (b) Transforming x_{UDOT} into x^* based on (2).

send two trucks to region 3; region 5 will send one truck to region 4 and another one to region 2; and region 1 will send two trucks to region 2. One additional conclusion that can be drawn from this discussion is that performance is quite sensitive to the resource allocation: very significant practical impact can be achieved by making very minimal adjustments to current practice.

However, once the fleet is repositioned, we expect that x^* will remain stable because the weather patterns, the road network, the partition of the network into regions, and the available fleet all change fairly slowly, and it is reasonable to suppose that they will remain constant, at least for the duration of a single winter season. Regions generally do not share or exchange vehicles on a short-term basis, as this would create a problem of accountability. UDOT has no protocols in place for shared decision making, and, in fact, this would defeat the purpose of partitioning the network in the first place, which is done to make operations more manageable and to distribute the responsibility between different teams.

6. Conclusions

We have considered the optimal allocation of snowplow trucks among a set of independent regions, while accounting for stochastic weather and truck failures. Both risk-averse and risk-neutral allocations can be efficiently computed. These allocations are shown to significantly outperform the UDOT's current allocation of trucks in northern Utah, as well as other benchmarks.

Although the developed methodology is applied to winter road maintenance in Utah, the same framework can be employed more broadly to allocate trucks for other practical problems in logistics where a large service area is subdivided into multiple regions or districts, such as waste collection or parking enforcement. Moreover, the developed framework would be applicable when allocating other types of resources that are subject to stochastic failures, such as the allocation of heavy machinery among construction sites with the objective to minimize the makespan. Our approach is most suitable when there is one type of resource with a clear priority over the others. For example, single-wing plow trucks represent 85% of UDOT's fleet. At the same time, the assignment of larger trucks (i.e., double-wing and tow-plow trucks) is fairly restricted to a few major highways. As a result, the allocation of single-wing trucks has a clear priority, as it is expected to have the highest impact on the efficiency of snow removal operations.

References

- Abdel-Malek LL, Chien SI, Meegoda JN, Yu H (2014) Approach to fleet contracting for snow plowing operations. *J. Infrastructure Systems* 20(2):04014006:1–04014006:7.
- Agarwal M, Maze TH, Souleyrette R (2005) Impacts of weather on urban freeway traffic flow characteristics and facility capacity. *Proc. 2005 Mid-Continent Transportation Res. Sympos.*, vol. 9 (Center for Transportation Research and Education, Iowa State University, Ames, IA), 15–28.
- Ahmed S, Xie W (2018) Relaxations and approximations of chance constraints under finite distributions. *Math. Program.* 170(1):43–65.
- An MY (1997) Log-concave probability distributions: Theory and statistical testing. Duke University Department of Economics Working Paper 95-03, Duke University, Durham, NC.
- Ando K, Fujishige S, Naitoh T (1995) A greedy algorithm for minimizing a separable convex function over a finite jump system. *J. Oper. Res. Soc. Japan* 38(3):362–375.
- Berrocal VJ, Raftery AE, Gneiting T, Steed RC (2010) Probabilistic weather forecasting for winter road maintenance. *J. Amer. Statist. Assoc.* 105(490):522–537.
- Bouchet A, Cunningham WH (1995) Delta-matroids, jump systems, and bisubmodular polyhedra. *SIAM J. Discrete Math.* 8(1):17–32.
- Boyd S, Vandenberghe L (2004) *Convex Optimization* (Cambridge University Press, Cambridge, UK).
- Burer S, Letchford AN (2012) Non-convex mixed-integer nonlinear programming: A survey. *Surv. Oper. Res. Management Sci.* 17(2): 97–106.
- Campbell JF, Langevin A, Perrier N (2015) Advances in vehicle routing for snow plowing. Corberán A, Laporte G, eds. *Arc Routing: Problems, Methods, and Applications* (SIAM, Philadelphia), 321–350.
- Chien SI, Gao S, Meegoda JN (2013) Fleet size estimation for snowplowing operation considering road geometry, weather, and traffic speed. *J. Transportation Engrg.* 139(9):903–912.
- Corberán Á, Eglese R, Hasle G, Plana I, Sanchis JM (2021) Arc routing problems: A review of the past, present, and future. *Networks* 77(1):88–115.
- Corbett J, Crocker M, Locke B, Show TJ (2020) Optimizing snow plow routes for the city of Bozeman. Report, Montana State University, Bozeman, MT.
- Crainic TG, Errico F, Rei W, Ricciardi N (2016) Modeling demand uncertainty in two-tier city logistics tactical planning. *Transportation Sci.* 50(2):559–578.
- Dowds J, Sullivan J (2021) Synthesis of technical requirements and considerations for automated snowplow route optimization. Report, Transportation Research Center, University of Vermont, Burlington, VT.
- Dror M (2012) *Arc Routing: Theory, Solutions and Applications* (Springer Science & Business Media, New York).
- Dussault B, Golden B, Groër C, Wasil E (2013) Plowing with precedence: A variant of the windy postman problem. *Comput. Oper. Res.* 40(4):1047–1059.
- FHWA (2020) Snow and ice. Accessed February 19, 2022, https://ops.fhwa.dot.gov/weather/weather_events/snow_ice.htm.
- Fu L, Trudel M, Kim V (2009) Optimizing winter road maintenance operations under real-time information. *Eur. J. Oper. Res.* 196(1): 332–341.
- Gerard D, Köppe M, Louveaux Q (2017) Guided dive for the spatial branch-and-bound. *J. Global Optim.* 68(4):685–711.
- Golden BL, Wong RT (1981) Capacitated arc routing problems. *Networks* 11(3):305–315.
- Hajibabai L, Ouyang Y (2016) Dynamic snow plow fleet management under uncertain demand and service disruption. *IEEE Trans. Intelligent Transportation Systems* 17(9):2574–2582.
- Hajibabai L, Nourbakhsh SM, Ouyang Y, Peng F (2014) Network routing of snowplow trucks with resource replenishment and plowing priorities: Formulation, algorithm, and application. *Transportation Res. Rec.* 2440(1):16–25.
- Han C, Yu XB (2017) Feasibility of geothermal heat exchanger pile-based bridge deck snow melting system: A simulation based analysis. *Renewable Energy* 101:214–224.
- Holmberg K (2019) The (over) zealous snow remover problem. *Transportation Sci.* 53(3):867–881.
- Iyer AV, Dunlop SR, Senicheva O, Thakkar DJ, Yan R, Subramanian K, Vasu S, Siddharthan G, Vasandani J, Saurabh S (2021) Improve

- and gain efficiency in winter operations. Joint Transportation Research Program Publication No. FHWA/IN/JTRP-2021/11, Purdue University, West Lafayette, IN.
- Jang W, Noble JS, Hutsel T (2010) An integrated model to solve the winter asset and road maintenance problem. *IIE Trans.* 42(9):675–689.
- Kandula P, Wright JR (1997) Designing network partitions to improve maintenance routing. *J. Infrastructure Systems* 3(4):160–168.
- Kinable J, van Hoeve W-J, Smith SF (2021) Snow plow route optimization: A constraint programming approach. *IIE Trans.* 53(6):685–703.
- Labelle A, Langevin A, Campbell JF (2002) Sector design for snow removal and disposal in urban areas. *Socio Econom. Planning Sci.* 36(3):183–202.
- Li Y, RazaviAlavi S, AbouRizk S (2021) Data-driven simulation approach for short-term planning of winter highway maintenance operations. *J. Comput. Civil Engrg.* 35(5):04021013.
- Li Y, Xu S, Wu L, AbouRizk S, Kwon TJ, Lei Z (2019) A generic simulation model for selecting fleet size in snow plowing operations. Lazarova-Molnar S., ed. *Proc. 2019 Winter Simul. Conf.* (IEEE, Piscataway, NJ), 2966–2977.
- Liu G, Ge Y, Qiu TZ, Soleymani HR (2014) Optimization of snow plowing cost and time in an urban environment: A case study for the city of Edmonton. *Canadian J. Civil Engrg.* 41(7):667–675.
- Miller T, Gleichert B, Crabtree H, Hendershot J, Nuveman R, Schneider W (2018) Role of route optimization in benefiting winter maintenance operations. *Transportation Res. Record* 2672(12):232–242.
- MoDOT (2022) Lack of snowplow drivers. Accessed February 25, 2022, https://www.komu.com/news/midmissourinews/lack-of-plow-drivers-will-cause-delays-modot-says/article_e968e3c8-82c2-11ec-ac89-27a2b1b0d151.html.
- Murota K (2021) A note on M-convex functions on jump systems. *Discrete Appl. Math.* 289:492–502.
- NCEI (2023) Daily U.S. snowfall. Accessed August 22, 2023, <https://www.ncei.noaa.gov/access/monitoring/daily-snow/>.
- Pearn W-L, Assad A, Golden BL (1987) Transforming arc routing into node routing problems. *Comput. Oper. Res.* 14(4):285–288.
- Perrier N, Langevin A, Amaya C-A (2008) Vehicle routing for urban snow plowing operations. *Transportation Sci.* 42(1):44–56.
- Perrier N, Langevin A, Campbell JF (2006a) A survey of models and algorithms for winter road maintenance. Part I: System design for spreading and plowing. *Comput. Oper. Res.* 33(1):209–238.
- Perrier N, Langevin A, Campbell JF (2006b) A survey of models and algorithms for winter road maintenance. Part II: System design for snow disposal. *Comput. Oper. Res.* 33(1):239–262.
- Perrier N, Langevin A, Campbell JF (2007a) A survey of models and algorithms for winter road maintenance. Part III: Vehicle routing and depot location for spreading. *Comput. Oper. Res.* 34(1):211–257.
- Perrier N, Langevin A, Campbell JF (2007b) A survey of models and algorithms for winter road maintenance. Part IV: Vehicle routing and fleet sizing for plowing and snow disposal. *Comput. Oper. Res.* 34(1):258–294.
- Quirion-Blais O, Langevin A, Trépanier M (2017) A case study of combined winter road snow plowing and de-icer spreading. *Canadian J. Civil Engrg.* 44(12):1005–1013.
- Quirion-Blais O, Trépanier M, Langevin A (2015) A case study of snow plow routing using an adaptive large hood search meta-heuristic. *Transportation Lett.* 7(4):201–209.
- Sabharwal CL (2019) Blended root finding algorithm outperforms bisection and regula falsi algorithms. *Mathematics* 7(11):1118.
- Salazar-Aguilar MA, Langevin A, Laporte G (2012) Synchronized arc routing for snow plowing operations. *Comput. Oper. Res.* 39(7):1432–1440.
- Schultz GG, Lau SK, Shoaf M, Bassett D, Eggett DL (2022) Analysis of using V2X DSRC-equipped snowplows to request signal preemption. Report, Utah Department of Transportation, Salt Lake City, UT.
- Shapiro A, Dentcheva D, Ruszczyński A (2021) *Lectures on Stochastic Programming: Modeling and Theory* (SIAM, Philadelphia).
- Shioura A, Tanaka K (2007) Polynomial-time algorithms for linear and convex optimization on jump systems. *SIAM J. Discrete Math.* 21(2):504–522.
- Smith EM, Pantelides CC (1999) A symbolic reformulation/spatial branch-and-bound algorithm for the global optimisation of nonconvex MINLPs. *Comput. Chem. Engrg.* 23(4–5):457–478.
- Sullivan JL, Dowds J, Novak DC, Scott DM, Ragsdale C (2019) Development and application of an iterative heuristic for roadway snow and ice control. *Transportation Res. Part A Policy Pract.* 127:18–31.
- Tang K, Dong S, Ma X, Lv L, Song Y (2020) Chance-constrained optimal power flow of integrated transmission and distribution networks with limited information interaction. *IEEE Trans. Smart Grid* 12(1):821–833.
- UDOT (2021) Utah snow removal. Accessed February 19, 2022, <https://www.udot.utah.gov/connect/public/snow-removal/>.
- Utah Winter Maintenance (2021) Winter weather maintenance and snow plow information. Accessed October 28, 2022, https://staging.udottraffic.utah.gov/PDFFiles/Snow_Plow_Priorities.pdf.
- Williams LF Jr (1976) A modification to the half-interval search (binary search) method. White MP, Kramer J, eds. *ACM-SE Proc. 14th Annu. Southeast Regional Conf.* (Association for Computing Machinery, New York), 95–101.
- Wolsey LA (2020) *Integer Programming* (John Wiley & Sons, Hoboken, NJ).
- Xu S, Kwon TJ (2021) Optimizing snowplow routes using all-new perspectives: Road users and winter road maintenance operators. *Canadian J. Civil Engrg.* 48(8):959–968.
- Ye Z (2009) Evaluation of the Utah DOT Weather Operations/RWIS Program on Traffic Operations. Technical report, Western Transportation Institute, Montana State University, Bozeman, MT.