

An Experimental Comparison of Uncertainty Sets for Robust Shortest Path Problems

Trivikram Dokka¹ and Marc Goerigk²

1 Department of Management Science, Lancaster University, Lancaster, UK
t.dokka@lancaster.ac.uk

2 Department of Management Science, Lancaster University, Lancaster, UK
m.goerigk@lancaster.ac.uk

Abstract

Through the development of efficient algorithms, data structures and preprocessing techniques, real-world shortest path problems in street networks are now very fast to solve. But in reality, the exact travel times along each arc in the network may not be known. This led to the development of robust shortest path problems, where all possible arc travel times are contained in a so-called uncertainty set of possible outcomes.

Research in robust shortest path problems typically assumes this set to be given, and provides complexity results as well as algorithms depending on its shape. However, what can actually be observed in real-world problems are only discrete raw data points. The shape of the uncertainty is already a modelling assumption. In this paper we test several of the most widely used assumptions on the uncertainty set using real-world traffic measurements provided by the City of Chicago. We calculate the resulting different robust solutions, and evaluate which uncertainty approach is actually reasonable for our data. This anchors theoretical research in a real-world application and gives an indicator which robust models should be the future focus of algorithmic development.

1998 ACM Subject Classification G.1.6 Optimization, G.2.2 Graph Theory, G.2.3 Applications

Keywords and phrases robust shortest paths, uncertainty sets, real-world data, experimental study

Digital Object Identifier 10.4230/OASIScs.ATMOS.2017.16

1 Introduction

The problem of finding shortest paths in real-world networks has seen considerable algorithmic improvements over the last decade [2]. In the typical problem setup, one assumes that all data is given exactly. But also robust shortest path problems have been considered, where travel times are assumed to be given by a set of possible scenarios. In [18], it was shown that the problem of finding a path that minimises the worst-case length over two scenarios is already weakly NP-hard. For general surveys on results in robust discrete optimisation, we refer to [1, 15].

There are many possibilities how to model the scenario set that is used for the robust optimisation process (see, e.g., [16, 8]), and it is not obvious which is "the right" one. Part of the current research ignores the problem by simply assuming that the uncertainty is "given" in some specific form, while this does not happen in reality.

In fact, the starting point for all uncertainty sets is raw data, given as a set of observations of travel times. This data is then processed to fit different assumptions on the shape and size of the uncertainty set, and preferences of the decision maker. So far, the discussion of these uncertainty sets has been led by theoretical properties, such as the computational tractability



© Trivikram Dokka and Marc Goerigk;
licensed under Creative Commons License CC-BY

17th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2017).

Editors: Gianlorenzo D'Angelo and Twan Dollevoet; Article No. 16; pp. 16:1–16:13

Open Access Series in Informatics



OASIS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of the resulting robust model. We believe that this leads to a gap in the literature, where models are not sufficiently underpinned by actual real-world data to verify results. The purpose of this paper is to close this gap. We use real-world traffic observations by the City of Chicago to create a selection of the best-known and most-used uncertainty sets from the research literature. Using these uncertainty sets, we calculate different robust solutions and compare their performance. This allows us to determine which uncertainty sets are actually valuable for the real-world robust shortest path application at hand. Our results give strong impulses for future research in the field by pointing out which problems are most worthy to solve more efficiently, assuming that a similar performance can be reproduced on other data sets.

In Section 2 we briefly introduce all six uncertainty sets used in this study, and discuss the complexity of the resulting robust problems. The experimental setup and results are then presented in Section 3, before concluding this paper in Section 4.

2 Uncertainty Sets for the Shortest Path Problem

Let a directed graph $G = (V, A)$ with nodes V and arcs A be given. In the classic shortest path problem, each arc e has some specific travel time $c_e \geq 0$. Given a start node s and a target node t , the aim is to find a path minimising the total travel time, i.e., to solve

$$\min\{\mathbf{c}\mathbf{x} : \mathbf{x} \in \mathcal{X}\}$$

where $\mathcal{X} \subseteq \{0, 1\}^n$ denotes the set of s - t -paths, and $n = |A|$. For our setting we assume instead that a set \mathcal{R} of travel time observations is given, $\mathcal{R} = \{\mathbf{c}^1, \dots, \mathbf{c}^N\}$ with $\mathbf{c}^i \in \mathbb{R}^n$. This is the available raw data. In the well-known robust shortest path problem we assume that an uncertainty set \mathcal{U} is produced based on this raw data, and solve the robustified problem

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{c} \in \mathcal{U}} \mathbf{c}\mathbf{x},$$

that is, we search for a path that minimises the worst-case costs over all scenarios. In the following sections we detail different possibilities from the current literature to generate the shape of \mathcal{U} based on \mathcal{R} . They are illustrated by an example in Figure 1. Each set is equipped with a scaling parameter to control its size (see also [13] on the problem of choosing the size of an uncertainty set with a given shape).

2.1 Convex Hull

In this approach, also known as discrete uncertainty (see [18, 15]), we set $\mathcal{U}^{CH} = \mathcal{R}$. The resulting robust problem can then be written as

$$\begin{aligned} \min z \\ \text{s.t. } z &\geq \mathbf{c}^i \mathbf{x} && \forall i \in [N] \\ &\mathbf{x} \in \mathcal{X} \end{aligned}$$

where we use the notation $[N] = \{1, \dots, N\}$. Note that this problem is equivalent to using $\mathcal{U}^{CH} = \text{conv}(\{\mathbf{c}^1, \dots, \mathbf{c}^N\})$, where conv denotes the convex hull of a set of points. The problem is known to be NP-hard already for two scenarios [18].

Scaling: Let $\hat{\mathbf{c}}$ be the average of $\{\mathbf{c}^1, \dots, \mathbf{c}^N\}$, i.e., $\hat{\mathbf{c}} = \frac{1}{N} \sum_{i \in [N]} \mathbf{c}^i$. For a given $\lambda \geq 0$, we substitute each point \mathbf{c}^i with $\hat{\mathbf{c}} + \lambda(\mathbf{c}^i - \hat{\mathbf{c}})$, and take the convex hull of the scaled data points.

2.2 Intervals

We set $\mathcal{U}^I = \times_{i \in [n]} [\min_{j \in [N]} c_i^j, \max_{j \in [N]} c_i^j]$ as the smallest hypercube containing all data. For ease of notation, we write $\bar{c}_i := \max_{j \in [N]} c_i^j$ and $\underline{c}_i := \min_{j \in [N]} c_i^j$. The resulting robust problem is then

$$\begin{aligned} \min \quad & \bar{\mathbf{c}}\mathbf{x} \\ \text{s.t.} \quad & \mathbf{x} \in \mathcal{X} \end{aligned}$$

which is a classic shortest path problem. Robust shortest path problems with interval uncertainty are therefore easy to solve, but frequently used, especially in the so-called min-max regret setting (see [10]).

Scaling: We use $\mathcal{U}^I = \times_{i \in [n]} \left[\frac{\bar{c}_i + \underline{c}_i}{2} - \lambda \frac{\bar{c}_i - \underline{c}_i}{2}, \frac{\bar{c}_i + \underline{c}_i}{2} + \lambda \frac{\bar{c}_i - \underline{c}_i}{2} \right]$ for some $\lambda \geq 0$.

2.3 Ellipsoid

Ellipsoidal uncertainty sets were first proposed in [3, 4] and stem from the observation that the iso-density locus of a multivariate normal distribution is an ellipse. We use an ellipsoid of the form $\mathcal{U}^E = \{\mathbf{c} : (\mathbf{c} - \boldsymbol{\mu})^t \boldsymbol{\Sigma}^{-1} (\mathbf{c} - \boldsymbol{\mu}) \leq \lambda\}$ with scaling parameter $\lambda \geq 0$ that is centred on $\hat{\mathbf{c}}$. We create it using a normal distribution found as a maximum-likelihood fit. Recall that the best fit of a multivariate normal distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with respect to data points $\mathbf{c}^1, \dots, \mathbf{c}^N$ is given by

$$\boldsymbol{\mu} = \hat{\mathbf{c}} = \frac{1}{N} (\mathbf{c}^1 + \dots + \mathbf{c}^N)$$

and

$$\boldsymbol{\Sigma} = \frac{1}{N} \sum_{i \in [N]} (\mathbf{c}^i - \boldsymbol{\mu})(\mathbf{c}^i - \boldsymbol{\mu})^t$$

The resulting problem can then be formulated as

$$\begin{aligned} \min \quad & \hat{\mathbf{c}}\mathbf{x} + z \\ \text{s.t.} \quad & z^2 \geq \lambda (\mathbf{x}^t \boldsymbol{\Sigma} \mathbf{x}) \\ & \mathbf{x} \in \mathcal{X} \end{aligned}$$

which is an integer second-order cone program (ISOCP), see [3] for details. Due to the convexity of the constraints, the problem can still be solved with little computational effort by standard solvers.

2.4 Budgeted Uncertainty

This approach was introduced in [6], and is based on interval uncertainty $\mathcal{U} = \times_{i \in [n]} [\hat{c}_i, \bar{c}_i]$. To reduce the conservatism of this approach one assumes that only at most $\Gamma \in \{0, \dots, n\}$ many values can be simultaneously higher than the midpoint $\hat{\mathbf{c}}$. Formally,

$$\mathcal{U}^B = \{\mathbf{c} : c_i = \hat{c}_i + (\bar{c}_i - \hat{c}_i)\delta_i \text{ for all } i \in [n], \mathbf{0} \leq \boldsymbol{\delta} \leq \mathbf{1}, \sum_{i \in [n]} \delta_i \leq \Gamma\}.$$

Using the dual of the inner worst-case problem, the following compact mixed-integer program can be found:

$$\begin{aligned}
\min \quad & \hat{\mathbf{c}}\mathbf{x} + \Gamma\pi + \sum_{i \in [n]} \rho_i \\
\text{s.t.} \quad & \pi + \rho_i \geq (\bar{c}_i - \hat{c}_i)x_i && \forall i \in [n] \\
& \pi, \boldsymbol{\rho} \geq 0 \\
& \mathbf{x} \in \mathcal{X}
\end{aligned}$$

This approach has the advantage that probability bounds can be found that compare favorably with those for ellipsoidal uncertainty [7], while this problem also remains polynomially solvable by enumerating possible values for the π variable. This means that $\mathcal{O}(n)$ many problems of the original type need to be solved. For these reasons, the budgeted uncertainty approach has been very popular in the literature (see, e.g., [14]).

2.5 Permutohull

The final two uncertainty sets we consider were proposed in [5]. The original inspiration comes from risk measures; the authors show that any so-called distortion risk measure leads to a polyhedral uncertainty set. A risk measure μ is a distortion risk measure if and only if there exists $\mathbf{q} \in \{\mathbf{q}' \in \Delta^N : q_1 \geq \dots \geq q_N\}$, where Δ^N denotes the N -dimensional simplex such that

$$\mu(\mathbf{x}) = - \sum_{i \in [N]} q_i(\mathbf{c}^{(i)}\mathbf{x})$$

where the sorting (i) is chosen such that $\mathbf{c}^{(1)}\mathbf{x} \geq \dots \geq \mathbf{c}^{(N)}\mathbf{x}$.

The conditional value at risk $CVaR_\alpha$ with $\alpha \in (0, 1]$ is a well-known distortion risk measure. Intuitively, it is the expected value amongst the α worst outcomes. Using the matrix

$$\mathbf{Q}_N := \begin{pmatrix} 1 & \dots & \frac{1}{N-2} & \frac{1}{N-1} & \frac{1}{N} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \frac{1}{N-2} & \frac{1}{N-1} & \frac{1}{N} \\ 0 & \dots & 0 & \frac{1}{N-1} & \frac{1}{N} \\ 0 & \dots & 0 & 0 & \frac{1}{N} \end{pmatrix}$$

the j th column of \mathbf{Q}_N induces the risk measure $CVaR_{j/N}$. The corresponding polyhedra are called the \mathbf{q} -permutohull and defined as

$$\Pi_{\mathbf{q}}(\mathbf{c}^1, \dots, \mathbf{c}^N) := \text{conv} \left(\left\{ \sum_{i \in [N]} q_{\sigma(i)} \mathbf{c}^i : \sigma \in S_N \right\} \right).$$

To find the resulting robust problem, we first consider the worst-case problem for fixed $\mathbf{x} \in \mathcal{X}$.

$$\begin{aligned}
\max \quad & \sum_{i, j \in [N]} q_i(\mathbf{c}^j\mathbf{x})p_{ij} \\
\text{s.t.} \quad & \sum_{i \in [N]} p_{ij} = 1 && \forall j \in [N] \\
& \sum_{j \in [N]} p_{ij} = 1 && \forall i \in [N] \\
& p_{ij} \geq 0 && \forall i, j \in [N]
\end{aligned}$$

■ **Table 1** Uncertainty sets in this study.

	\mathcal{U}^{CH}	\mathcal{U}^I	\mathcal{U}^E	\mathcal{U}^B	\mathcal{U}^{PH}	\mathcal{U}^{SPH}
Complexity	NPH	P	NPH	P	NPH	NPH
Model	IP	LP	ISOCP	MIP	MIP	MIP
Add. Const.	N	0	1	$n + 1$	N^2	N^2
Add. Var.	1	0	1	n	$2n$	$2n$

Dualising this problem then gives the robust counterpart

$$\begin{aligned}
 \min \quad & \sum_{i \in [n]} (v_i + w_i) \\
 \text{s.t.} \quad & v_i + w_j \geq q_i(\mathbf{c}^j \mathbf{x}) \quad \forall i, j \in [N] \\
 & \mathbf{v}, \mathbf{w} \geq 0 \\
 & \mathbf{x} \in \mathcal{X}
 \end{aligned}$$

which is a mixed-integer program (note that this approach is actually the same as the ordered weighted averaging method, see [9]). The problem is NP-hard, as it contains the convex hull of $\{\mathbf{c}^1, \dots, \mathbf{c}^N\}$ as a special case. Through the choice of \mathbf{q} , there are N possible sizes of this uncertainty.

2.6 Symmetric Permutohull

In the same setting as before, the symmetric permutohull was also introduced in [5]. By using the $\lfloor N/2 \rfloor + 1$ columns of the matrix

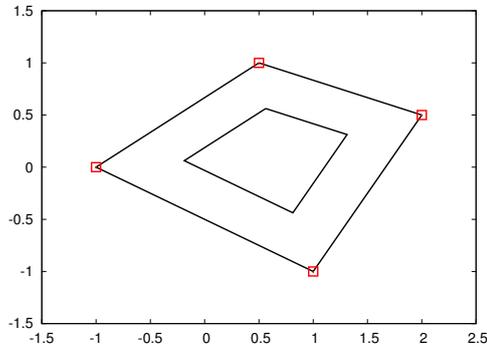
$$\tilde{\mathbf{Q}} := \frac{1}{N} \begin{pmatrix} 1 & 2 & 2 & \dots & 2 \\ 1 & 1 & 2 & \dots & 2 \\ 1 & 1 & 1 & \dots & 2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \dots & 0 \\ 1 & 1 & 0 & \dots & 0 \\ 1 & 0 & 0 & \dots & 0 \end{pmatrix}$$

it was shown that the resulting polyhedra are symmetric with respect to $\hat{\mathbf{c}}$. Note that these problems are also NP-hard, as $\tilde{\mathbf{Q}}$ contains the min-max approach for $N = 2$ as a special case.

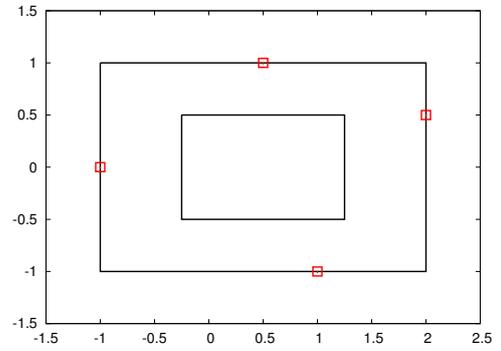
2.7 Summary of Uncertainty Sets

In total we described six methods to generate uncertainty set \mathcal{U} based on the raw data \mathcal{R} . Figure 1 illustrates these sets using a raw dataset with four observations (shown as red points). The complexity to solve the resulting robust models, as well as the type of program with the numbers of additional variables and constraints compared to the classic shortest path problem are shown in Table 1.

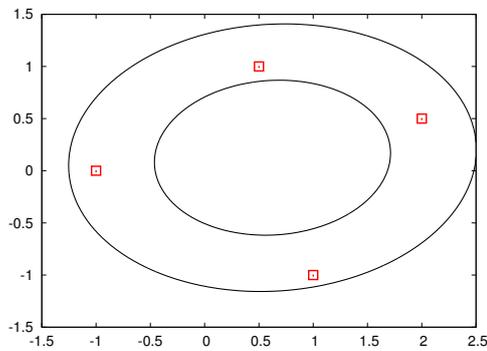
While the robust model with budgeted uncertainty sets can be solved in polynomial time using combinatorial algorithms, we still used the MIP formulation for our experiments, as it was sufficiently fast.



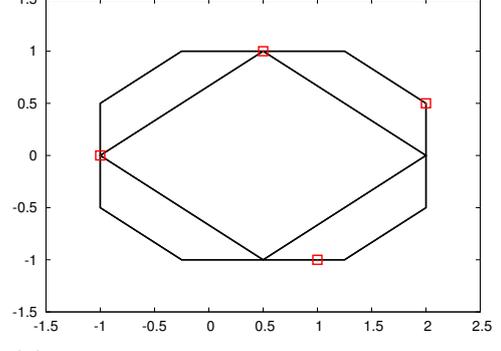
(a) Convex hull with $\lambda = 1$ and $\lambda = 0.5$.



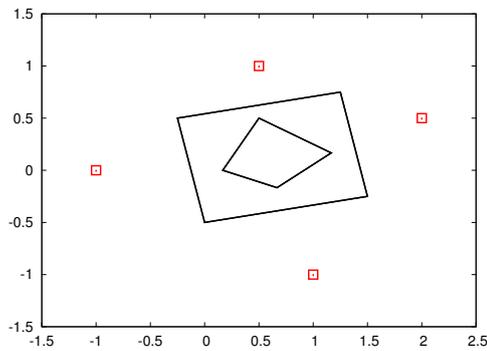
(b) Intervals with $\lambda = 1$ and $\lambda = 0.5$.



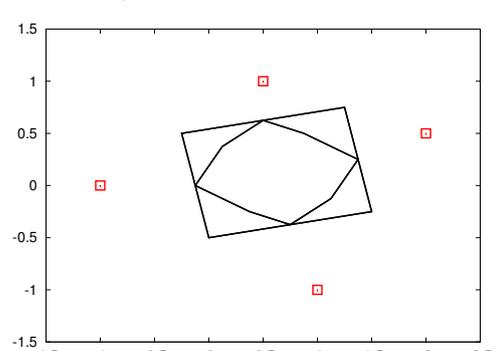
(c) Ellipsoid with $\lambda = 3$ and $\lambda = 1$.



(d) Budgeted uncertainty with $\lambda = 1$, $\Gamma = 1$ and $\lambda = 1$, $\Gamma = 1.5$.



(e) Permutohull uncertainty for $CVaR_{2/N}$ and $CVaR_{3/N}$ (i.e., $\mathbf{q} = (\frac{1}{2}, \frac{1}{2}, 0, 0)$ and $\mathbf{q} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$).



(f) Symmetric permutohull uncertainty for $\mathbf{q} = (\frac{1}{2}, \frac{1}{2}, 0, 0)$ and $\mathbf{q} = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4}, 0)$.

■ **Figure 1** Example uncertainty sets.

3 Real-World Experiments

3.1 Data Collection and Cleaning

We used data provided by the City of Chicago¹, which provides a live traffic data interface. We recorded traffic updates in a 15-minute interval over a time horizon of 24 hours spanning Monday March 27th 2017 morning to Tuesday March 28th 2017 morning. A total of 96 data points were thus used.

Every data point contains the traffic speed for a subset of a total of 1,257 segments. For each segment the geographical position is available, see the resulting plot in Figure 2a with a zoom-in for the city center. The complete travel speed data set contains a total of 54,295 records. There were 1,027 segments where the data was recorded at least once of the 96 data points. Nearly for 88% of the segments, speeds were recorded for at least 50 data points with only 1% (10 segments) where only one data point was recorded. We used linear interpolation to fill the missing records keeping in mind that data was collected over time. The data after removing missing records and filling missing values can be found at www.lancaster.ac.uk/~goerigk/robust-sp-data.zip. Segment lengths were given through longitude and latitude coordinates, and approximated using the Euclidean distance.

As segments are purely geographical objects without structure, we needed to create a graph for our experiments. To this end, segments were split when they crossed or nearly crossed, and start- and end-points that were sufficiently close to each other were identified as the same node. The resulting graph is shown in Figure 2b; note that this process slightly simplified the network, but kept its structure intact. The final graph contains 538 nodes and 1308 arcs. Using arc length and speed, we calculate their respective traversal time for each of the 96 data points. In the following, we refer to 96 *scenarios* generated this way.

3.2 Setup

Each uncertainty set is equipped with a scaling parameter. For each parameter we generated 20 possible values, reflecting a reasonable range of choices for a decision maker:

- For \mathcal{U}^{CH} and \mathcal{U}^I , $\lambda \in \{0.1, 0.2, \dots, 2\}$.
- For \mathcal{U}^E , $\lambda \in \{0.2, 0.4, \dots, 4\}$.
- For \mathcal{U}^B , $\Gamma \in \{5, 10, \dots, 100\}$.
- For \mathcal{U}^{PH} , we used columns $\mathbf{q}_1, \mathbf{q}_3, \dots, \mathbf{q}_{39}$.
- For \mathcal{U}^{SPH} , we used columns $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{20}$.

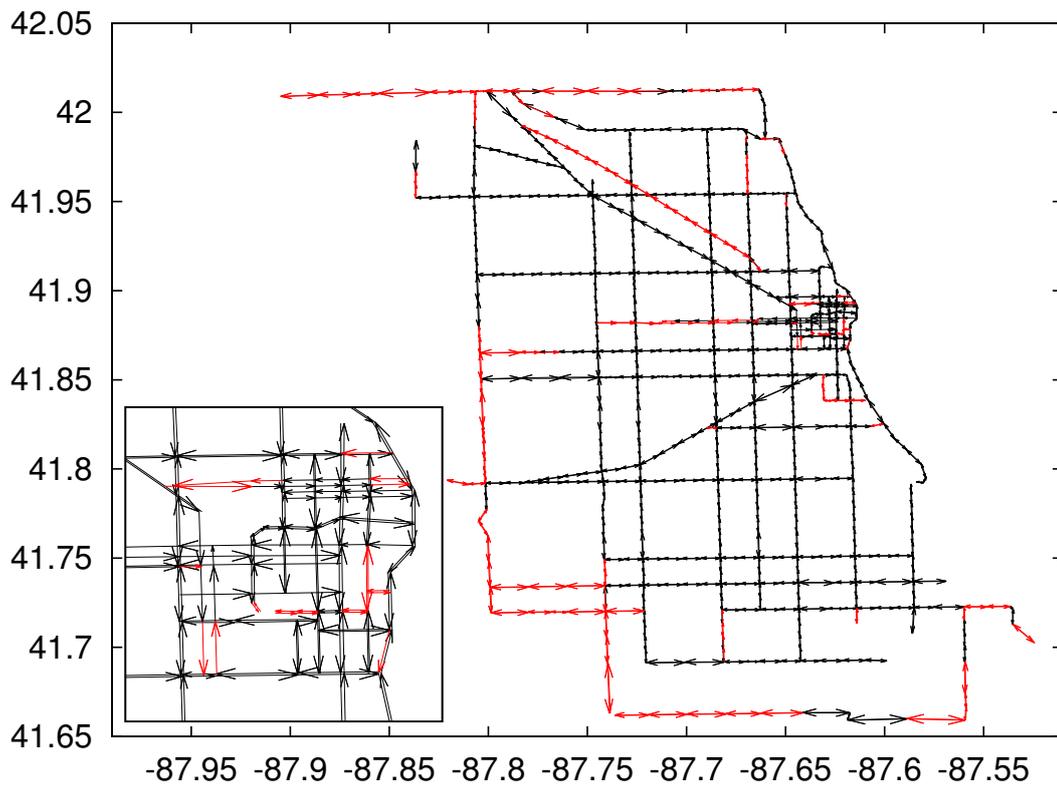
Each uncertainty set is generated using only every second scenario (i.e., 48 out of 96), but all 96 scenarios are then used to evaluate the solutions. Furthermore, we generated 200 random $s - t$ pairs uniformly, and used each of the $6 \cdot 20$ methods on the same 200 pairs. Each of our 120 methods hence generates $200 \cdot 96 = 19,200$ objective values.

It is highly non-trivial to assess the quality of these solutions, see [12]. If one just uses the average objective value, as an example, then one could as well calculate the solution optimising the average scenario case to find the best performance with respect to this measure. To find a balanced evaluation of all methods, we used three performance criteria:

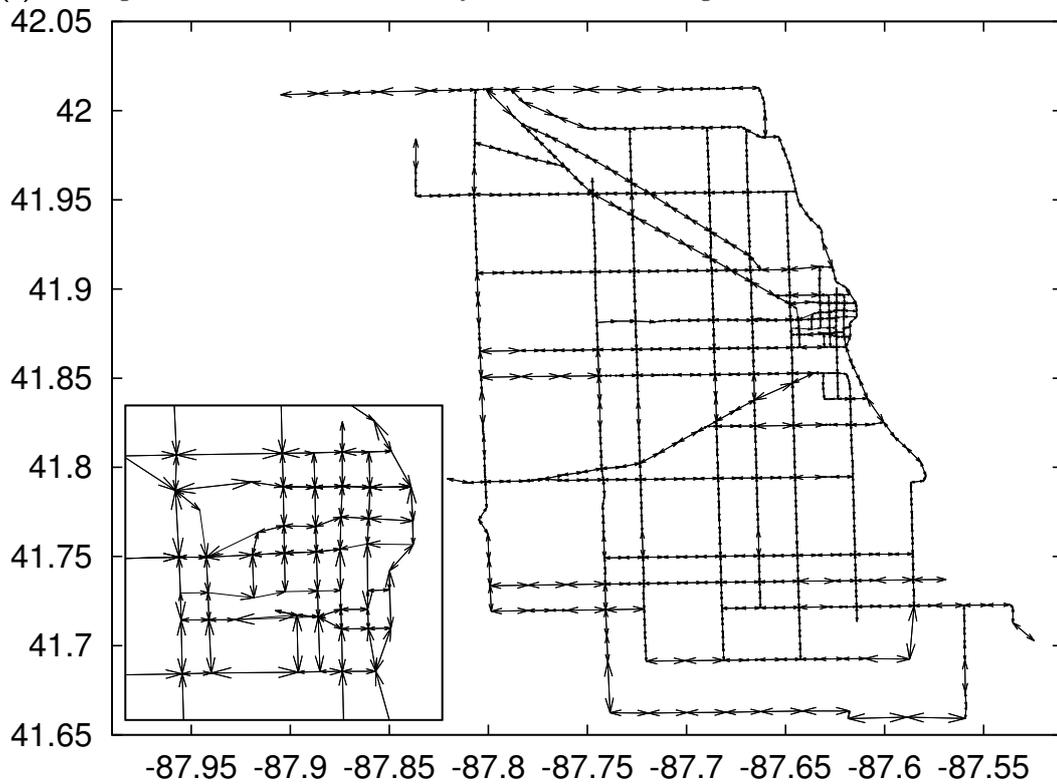
- the average objective value over all $s - t$ pairs and all scenarios,
- the average of the worst-case objective value for each $s - t$ pair, and
- the average value of the worst 5% of objective values for each $s - t$ pair (as in the CVaR measure)

Note that many more criteria would be possible to use.

¹ <https://data.cityofchicago.org>



(a) Raw segments with zoom-in for the city center. In red are segments without data.



(b) Resulting graph model with zoom-in for the city center.

■ Figure 2 Chicago instance.

For all experiments we used a computer with a 16-core Intel Xeon E5-2670 processor, running at 2.60 GHz with 20MB cache, and Ubuntu 12.04. Processes were pinned to one core. We used Cplex v.12.6 to solve all problem formulations (note that specialised combinatorial algorithms may be available for some problems).

3.3 Results

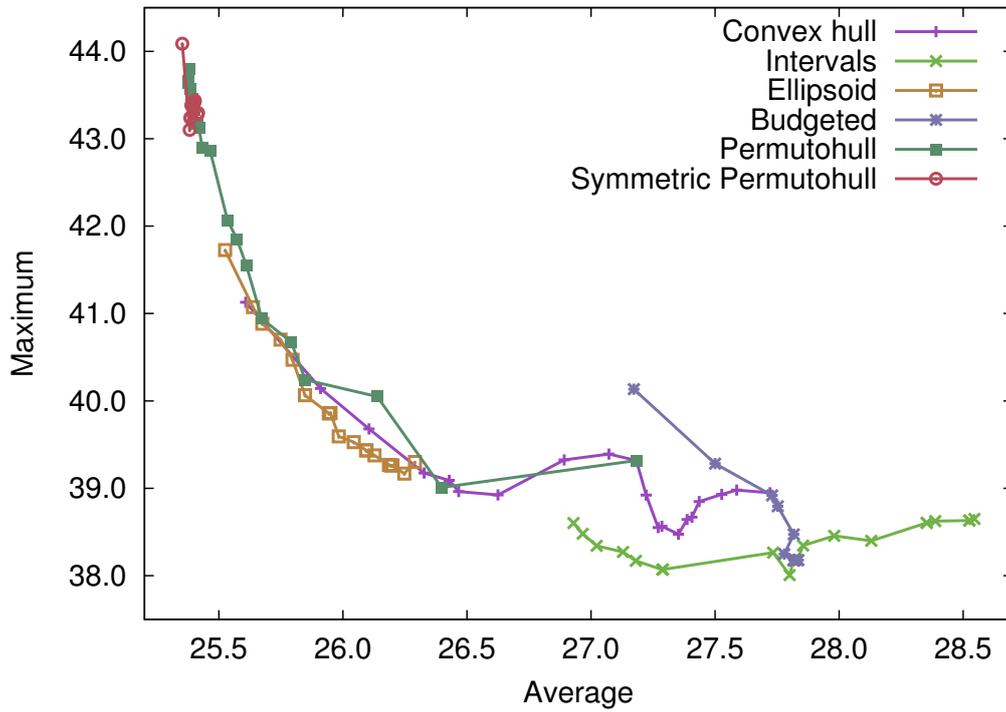
We present our findings in the two plots of Figure 3. In each plot, the 20 parameter settings that belong to the same uncertainty set are connected by a line. They are complemented with Figure 4 showing the total computation times for the methods over all 200 shortest path calculations.

The first plot in Figure 3a shows the trade-off between the average and the maximum objective value; the second plot in Figure 3b shows the trade-off between the average and the average of the 5% worst objective values. All values are in minutes. Note that for all performance measures, smaller values indicate a better performance – hence, good trade-off solutions should move from the top left to the bottom right of the plots. In general, the points corresponding to the parameter settings that give weight to the average performance can be on the left sides of the curves, while the more robust parameter settings are on the right sides, as would be expected.

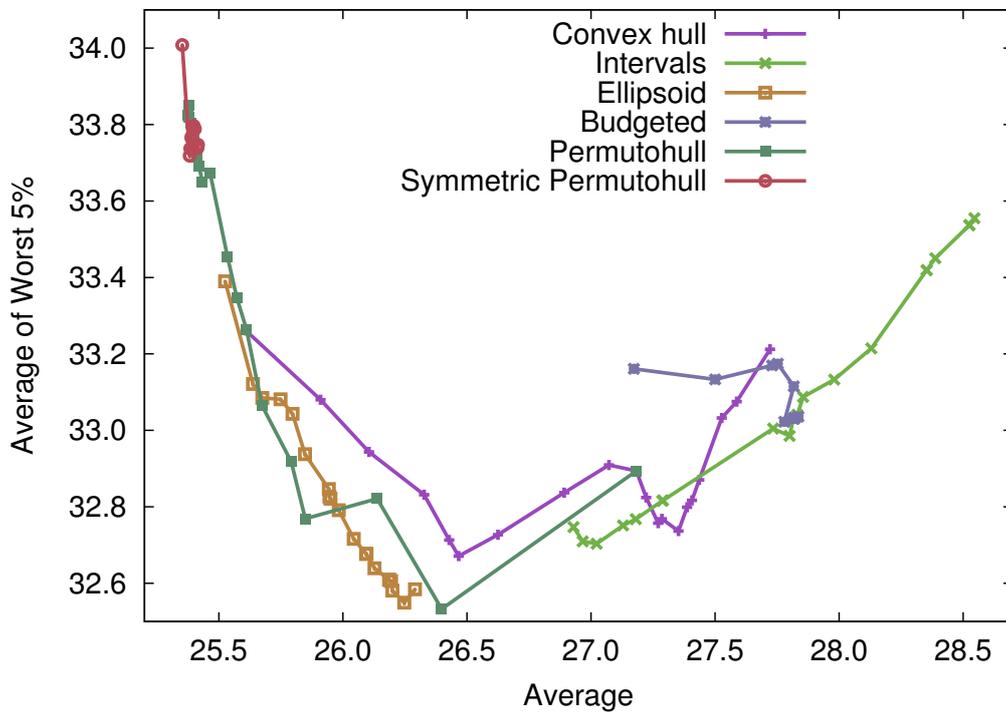
We first discuss Figure 3a. In general, we find that most concepts indeed present a trade-off between average performance and robustness through their scaling parameter. The symmetric permutohull solutions have the best average performance, while interval solutions are the most robust. Interestingly, that even holds for interval solutions where the scaling parameter is very small. The budgeted uncertainty does not give a good trade-off between worst-case and average-case performance, which confirms previous results on artificial data [11]. Scaling interval uncertainty sets achieves better results than using budgeted uncertainty. Solutions generated with ellipsoidal uncertainty sets slightly outperform (dominate in the Pareto sense) solutions generated with permutohull uncertainty. We also note that methods that are computationally more expensive tend to achieve better average performance at the cost of decreases robustness. The simplest and cheapest method, interval uncertainty, gives the most robust solutions. Solutions using the convex hull of raw data tend to be outperformed by the approaches that process data.

We now consider the results presented in Figure 3b. Here the average is plotted against the average performance of the 5% worst performing scenarios, averaged over all $s - t$ pairs. We note that for interval uncertainty, these two criteria are connected, with the best solutions for small scaling parameters dominating all solutions for larger scaling parameters. For the permutohull and the ellipsoidal uncertainty solutions, the order slightly changed with the former often dominating the latter. Permutohull solutions are designed to be efficient for the CVaR criterion, and the best-performing solution with respect to this aspect is indeed generated by this approach. However, also solutions with ellipsoidal, interval and convex hull uncertainty perform well.

Regarding computation times (see Figure 4), note that the two polynomially solvable approaches are also the fastest when using Cplex; these computation times can be further improved using specialised algorithms. Using the convex hull is faster than using ellipsoids, which are in turn faster than using the symmetric permutohull. For the standard permutohull, the computation times are sensitive to the uncertainty size; if the \mathbf{q} vector that is used in the model has only few entries, then computation times are smaller. This is in line with the intuition that the problem becomes easier if fewer scenarios need to be considered.

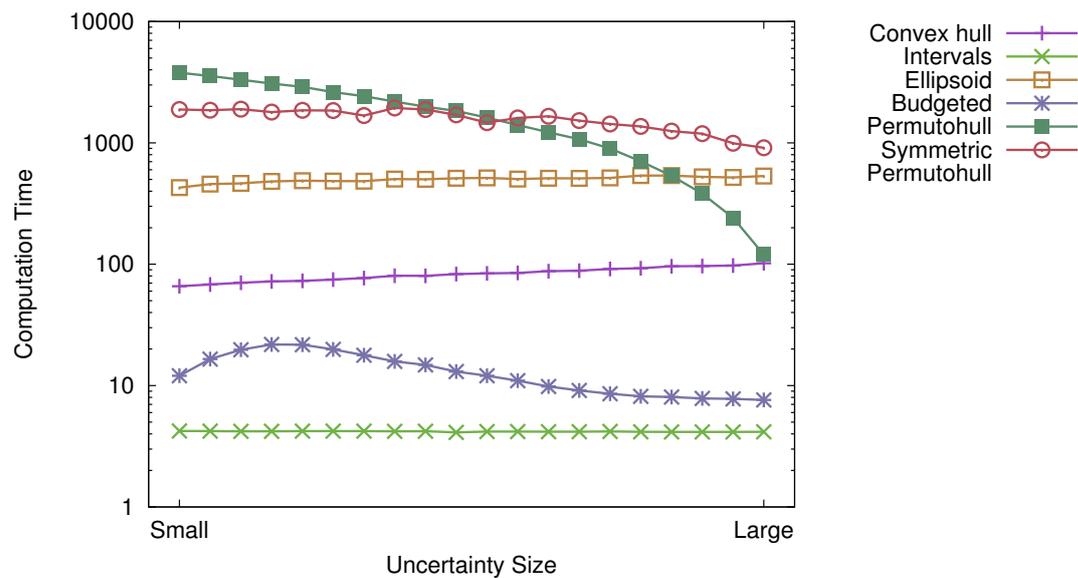


(a) Average vs worst-case performance.



(b) Average vs CVaR performance.

■ Figure 3 Performance results.



■ **Figure 4** Computation times in seconds.

To summarise our findings in our experiment on the robust shortest path problem with real-world data:

- Convex hull solutions are amongst the more robust solutions, but tend to be outperformed by the other approaches.
- Interval solutions perform bad on average, but are the most robust. Especially when the scaling is small they can give a decent trade-off, and are easy and fast to compute.
- Ellipsoidal uncertainty solutions have very good overall performance and represent a large part of the non-dominated points in our results.
- We do not encourage the use of budgeted uncertainty for robust shortest path problems. Scaling interval uncertainty sets gives better results and is easier to use and to solve.
- Permutohull solutions offer good trade-off solutions, whereas symmetric permutohull solutions tend to be less robust, but provide an excellent average performance. These methods also require most computational effort to find.

In the light of these findings, the interval and discrete (=convex hull) uncertainty sets that are widely used in robust combinatorial optimisation do warrant research attention, as they may not produce the best solutions, but are relatively fast to solve. However, permutohull and ellipsoidal uncertainty tend to produce solutions with a better trade-off, while being computationally more challenging. The algorithmic research for robust shortest path problems with such structure should therefore become a future focus.

4 Conclusion

In this paper we constructed uncertainty sets for the robust shortest path problem using real-world traffic observations for the City of Chicago. We evaluated the model suitability of these sets by finding the resulting robust paths, and comparing their performance using different performance indicators.

Naturally, conclusions can only be drawn within the reach of the available data. In our setting we considered solutions that are robust with respect to all possible travel times within a day. A use-case would be that a path needs to be computed for a specific day, but the

precise hour is not known. Using different sets of observations will result in solutions that are different in another sense, e.g., one could use observations over different days during the morning rush hours, or observations that span work days and a weekend. It is possible that these sets will provide different structure.

Finally, we have observed that using ellipsoidal uncertainty sets provides high-quality solutions with less computational effort than for the permutohull. If one uses only the diagonal entries of the matrix Σ , then one ignores the data correlation in the network. For the resulting problem specialised algorithms exist, see, e.g. [17]. In additional experiments we found that even by using Cplex, computation times were considerably reduced when only using the diagonal entries of Σ , but the solution quality remained roughly the same.

References

- 1 Hassene Aissi, Cristina Bazgan, and Daniel Vanderpooten. Min–max and min–max regret versions of combinatorial optimization problems: A survey. *European journal of operational research*, 197(2):427–438, 2009.
- 2 Hannah Bast, Daniel Delling, Andrew Goldberg, Matthias Müller-Hannemann, Thomas Pajor, Peter Sanders, Dorothea Wagner, and Renato F Werneck. Route planning in transportation networks. In *Algorithm Engineering*, pages 19–80. Springer, 2016.
- 3 Aharon Ben-Tal and Arkadi Nemirovski. Robust convex optimization. *Mathematics of operations research*, 23(4):769–805, 1998.
- 4 Aharon Ben-Tal and Arkadi Nemirovski. Robust solutions of uncertain linear programs. *Operations research letters*, 25(1):1–13, 1999.
- 5 Dimitris Bertsimas and David B. Brown. Constructing uncertainty sets for robust linear optimization. *Operations research*, 57(6):1483–1495, 2009.
- 6 Dimitris Bertsimas and Melvyn Sim. Robust discrete optimization and network flows. *Mathematical programming*, 98(1):49–71, 2003.
- 7 Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations research*, 52(1):35–53, 2004.
- 8 Christina Büsing. Recoverable robust shortest path problems. *Networks*, 59(1):181–189, 2012.
- 9 André Chassein and Marc Goerigk. Alternative formulations for the ordered weighted averaging objective. *Information Processing Letters*, 115(6):604–608, 2015.
- 10 André Chassein and Marc Goerigk. A new bound for the midpoint solution in minmax regret optimization with an application to the robust shortest path problem. *European Journal of Operational Research*, 244(3):739–747, 2015.
- 11 André Chassein and Marc Goerigk. A bicriteria approach to robust optimization. *Computers & Operations Research*, 66:181–189, 2016.
- 12 André Chassein and Marc Goerigk. Performance analysis in robust optimization. In *Robustness Analysis in Decision Aiding, Optimization, and Analytics*, pages 145–170. Springer, 2016.
- 13 André Chassein and Marc Goerigk. Variable-sized uncertainty and inverse problems in robust optimization. *European Journal of Operational Research*, 2017. Available online, to appear.
- 14 Marc Goerigk and Anita Schöbel. Algorithm engineering in robust optimization. In *Algorithm engineering*, pages 245–279. Springer, 2016.
- 15 Adam Kasperski and Paweł Zieliński. Robust discrete optimization under discrete and interval uncertainty: A survey. In *Robustness Analysis in Decision Aiding, Optimization, and Analytics*, pages 113–143. Springer, 2016.

- 16 Roberto Montemanni and Luca Maria Gambardella. An exact algorithm for the robust shortest path problem with interval data. *Computers & Operations Research*, 31(10):1667–1680, 2004.
- 17 Evdokia Nikolova. High-performance heuristics for optimization in stochastic traffic engineering problems. In *International Conference on Large-Scale Scientific Computing*, pages 352–360. Springer, 2009.
- 18 Gang Yu and Jian Yang. On the robust shortest path problem. *Computers & Operations Research*, 25(6):457–468, 1998.