

Accepted Manuscript

Algorithms and Uncertainty Sets for Data-Driven Robust Shortest Path Problems

André Chassein, Trivikram Dokka, Marc Goerigk

PII: S0377-2217(18)30857-9
DOI: <https://doi.org/10.1016/j.ejor.2018.10.006>
Reference: EOR 15398



To appear in: *European Journal of Operational Research*

Received date: 4 December 2017
Revised date: 4 October 2018
Accepted date: 5 October 2018

Please cite this article as: André Chassein, Trivikram Dokka, Marc Goerigk, Algorithms and Uncertainty Sets for Data-Driven Robust Shortest Path Problems, *European Journal of Operational Research* (2018), doi: <https://doi.org/10.1016/j.ejor.2018.10.006>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Highlights

- We consider robust shortest path problems with real-world street data.
- Different uncertainty sets for data-driven robust optimisation are compared.
- Trade-offs between uncertainty sets are identified.
- A new branch-and-bound algorithm for ellipsoidal uncertainty is proposed.

ACCEPTED MANUSCRIPT

Algorithms and Uncertainty Sets for Data-Driven Robust Shortest Path Problems

André Chassein^a, Trivikram Dokka^b, Marc Goerigk^{c,*}

^a*Data Analytics Center of Excellence, Deutsche Post DHL Group, Bonn, Germany*

^b*Department of Management Science, Lancaster University, Lancaster, United Kingdom*

^c*Network and Data Science Management, University of Siegen, Germany*

Abstract

We consider robust shortest path problems, where the aim is to find a path that optimizes the worst-case performance over an uncertainty set containing all relevant scenarios for arc costs. The usual approach for such problems is to assume this uncertainty set given by an expert who can advise on the shape and size of the set.

Following the idea of data-driven robust optimization, we instead construct a range of uncertainty sets from the current literature based on real-world traffic measurements provided by the City of Chicago. We then compare the performance of the resulting robust paths within and outside the sample, which allows us to draw conclusions on the suitability of uncertainty sets.

Based on our experiments, we then focus on ellipsoidal uncertainty sets, and develop a new solution algorithm that significantly outperforms a state-of-the-art solver.

Keywords: robustness and sensitivity analysis, robust shortest paths, uncertainty sets, data-driven robust optimization

1. Introduction

For classic shortest path problems in street networks, considerable speed-ups over a standard Dijkstra's algorithm have been achieved thanks to algorithm engineering techniques [2], which makes the use of real-time information even in large networks possible. Most types of robust shortest path problems, on the other hand, are NP-hard (see [22]), and real-time re-optimization has not been an option.

To formulate a robust problem, it is necessary to have a description of all possible and relevant scenarios that the solution should prepare against, the so-called uncertainty set. We refer to the surveys [1, 18, 19] for a general overview

*Corresponding author

Email addresses: chassein@mathematik.uni-kl.de (André Chassein),
t.dokka@lancaster.ac.uk (Trivikram Dokka), m.goerigk@lancaster.ac.uk (Marc Goerigk)

on the topic. The current literature on robust shortest paths usually assumes this set to be given, by some mixture of data-preprocessing and expert knowledge that is not part of the study. This means that different types of sets have been studied (compare, e.g., [20, 9]), but it has been impossible to address the question which would be the “right” choice.

A recent paradigm shift is data-driven robust optimization (see [6]), where building the uncertainty set from raw observations is part of the robust optimization problem. This paper is the first to follow such a perspective for shortest path problems. Based on real-world observations from the City of Chicago, we build a range of uncertainty sets, calculate the corresponding robust solutions, and perform an in-depth analysis of their performance. This allows us to give an indication which set is actually suitable for our application, and which are not.

In the second part of this paper, we then focus on the case of ellipsoidal uncertainty, and provide a branch-and-bound algorithm that is able to solve instances considerably faster than an off-the-shelf solver.

Parts of this paper were previously published as a conference paper in [17]. In comparison, we provide a completely new set of experimental results based on an observation period of 46 days (instead of one single day), which leads to a more detailed insight into the performance of different uncertainty sets. Furthermore, we provide a new analysis for axis-parallel ellipsoidal uncertainty sets, including an efficient branch-and-bound algorithm that is able to outperform CPLEX [15] by several orders of magnitude, pushing robust shortest paths towards applicability in real-time navigation systems.

The remainder of the paper is structured as follows. In Section 2 we briefly introduce the robust shortest path problem along with six uncertainty sets used in this study. The experimental setup and results on real-world data are then presented in Section 3. Section 4 presents our algorithm for ellipsoidal uncertainty sets, and includes additional computational results. The paper is concluded in Section 5.

2. Uncertainty Sets for the Shortest Path Problem

In this section, we briefly introduce the robust shortest path problem and different approaches to model uncertainty sets that are used in the current literature. In the classic shortest path problem, we are given a directed graph $G = (V, A)$ where V denotes the set of nodes, and A denotes the set of arcs. For every arc $e \in A$, we know its traversal time $c_e \geq 0$. For a start node s and a target node t , the aim is to find a path minimizing the total travel time given as the sum of times over all arcs that are part of the path.

More formally, we denote this problem as

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

$\mathcal{X} \subseteq \{0, 1\}^n$ denotes the set of s - t -paths, and $n = |A|$ is the number of variables.

In our setting, we assume that travel times \mathbf{c} are not known exactly. Instead, we are provided with a set \mathcal{R} of travel time observations, where $\mathcal{R} = \{\mathbf{c}^1, \dots, \mathbf{c}^N\}$ with $\mathbf{c}^i \in \mathbb{R}^n$. We also refer to \mathcal{R} as the available raw data. Based on this data, an uncertainty set \mathcal{U} is generated which is then used within the robust shortest path problem

$$\min \left\{ \max_{\mathbf{c} \in \mathcal{U}} \mathbf{c}^t \mathbf{x} : \mathbf{x} \in \mathcal{X} \right\}$$

that is, we search for a path that minimizes the worst-case costs over all costs in \mathcal{U} . We briefly sketch approaches to generate \mathcal{U} in the following. Each is equipped with a scaling parameter to control its size (see also [14] on the problem of choosing the size of an uncertainty set with a given shape). A visual example using four data points in two dimensions is provided for each approach in Figure 1. We use the notation $[N] = \{1, \dots, N\}$ and denote by $\hat{\mathbf{c}}$ 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 more details on the resulting models, we refer to the conference version of this paper [17].

- Convex hull uncertainty (see, e.g., [19, 22]): We set

$$\mathcal{U}^{CH} = \mathcal{R}.$$

Note that this is equivalent to using the convex hull of raw data. To scale this set, we substitute each point \mathbf{c}^i with $\hat{\mathbf{c}} + \lambda(\mathbf{c}^i - \hat{\mathbf{c}})$ for a given $\lambda \geq 0$, and take the convex hull of the scaled data points.

- Interval uncertainty (see, e.g., [11]): We set

$$\mathcal{U}^I = \prod_{i \in [n]} [\hat{c}_i + \lambda(\underline{c}_i - \hat{c}_i), \hat{c}_i + \lambda(\bar{c}_i - \hat{c}_i)]$$

for some $\lambda \geq 0$, where $\bar{c}_i := \max_{j \in [N]} c_i^j$ and $\underline{c}_i := \min_{j \in [N]} c_i^j$. By \times , we denote the Cartesian product. Note that $\max_{\mathbf{c} \in \mathcal{U}^I} \mathbf{c}^t \mathbf{x} = \sum_{i \in [n]} (\hat{c}_i x_i + \lambda(\bar{c}_i - \hat{c}_i) x_i)$ for every $\mathbf{x} \in \mathcal{X}$, and the robust optimization problem can be solved by solving a nominal (non-robust) problem. In Figure 1(b), only the top right corner of the boxes becomes part of the robust problem.

- Ellipsoidal uncertainty (see, e.g., [3, 4]): We set

$$\mathcal{U}^E = \{\mathbf{c} : (\mathbf{c} - \hat{\mathbf{c}})^t \boldsymbol{\Sigma}^{-1} (\mathbf{c} - \hat{\mathbf{c}}) \leq \lambda\}$$

with $\boldsymbol{\Sigma} = \frac{1}{N} \sum_{i \in [N]} (\mathbf{c}^i - \hat{\mathbf{c}})(\mathbf{c}^i - \hat{\mathbf{c}})^t$ derived from a maximum-likelihood fit of a normal distribution. The resulting sets are ellipsoids, see Figure 1(c), containing the highest-density region of a fitted multivariate normal distribution.

- Budgeted uncertainty (see, e.g., [7, 8, 18]): We set

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

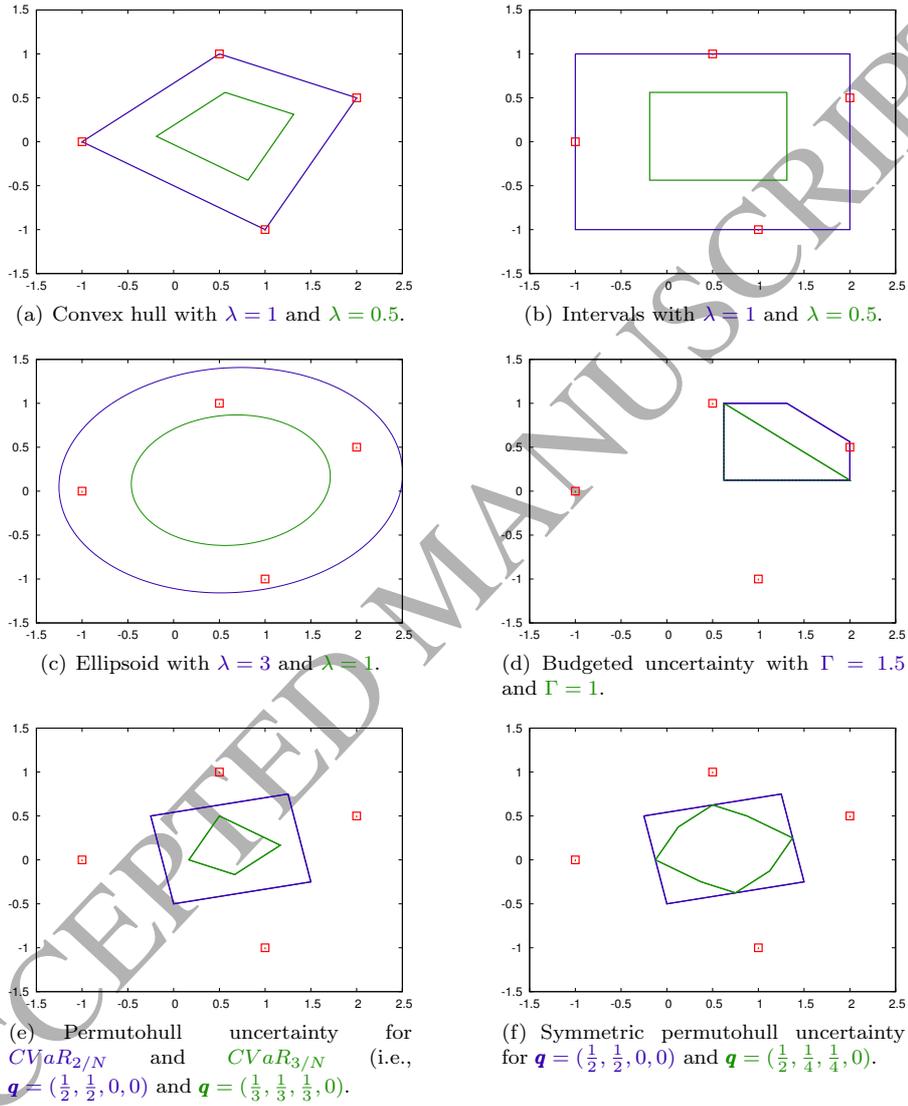


Figure 1: Example uncertainty sets.

where the parameter Γ controls the size of \mathcal{U}^B . This approach can be seen as a variant of interval uncertainty, where the top right corner is cut off to decrease conservatism. The size of Γ controls the size of the region that is cut off from the box.

- Permutohull uncertainty (see [5]): We set

$$\mathcal{U}^{PH} = \text{conv} \left(\left\{ \sum_{i \in [N]} q_{\sigma(i)} \mathbf{c}^i : \sigma \in S_N \right\} \right)$$

where S_N denotes the set of permutations on $[N]$, and \mathbf{q} is a column of the matrix

$$Q_N := \begin{pmatrix} 1 & \dots & \frac{1}{N-2} & \frac{1}{N-1} & \frac{1}{N} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \frac{1}{N-2} & \frac{1}{N-1} & \frac{1}{N} \\ 0 & \dots & 0 & \frac{1}{N-1} & \frac{1}{N} \\ 0 & \dots & 0 & 0 & \frac{1}{N} \end{pmatrix} \in \mathbb{R}^{N \times N}$$

Scaling is included via the choice of the column, where using the j th column of Q_N corresponds to using the conditional value at risk *CVaR* criterion with respect to risk level j/N . Using $\mathbf{q} = (\frac{1}{2}, \frac{1}{2}, 0, 0)$ as shown in Figure 1(e) means that we take the convex hull of the average of each pair of original points; using $\mathbf{q} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0)$ amounts to taking the convex hull of the average of each triplet of original points.

- Symmetric permutohull uncertainty (see [5]): As in the above case, but we generate \mathcal{U}^{SPH} using columns of the matrix $\tilde{Q} \in \mathbb{R}^{N \times (\lfloor N/2 \rfloor + 1)}$ defined by

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

instead, i.e., in the first column, all entries are $1/N$; in the second column, the first entry is $2/N$ and the last entry is 0, etc. The resulting sets are symmetric with respect to $\hat{\mathbf{c}}$.

In total we use six methods to generate uncertainty set \mathcal{U} based on the raw data \mathcal{R} . The resulting optimization model and its complexity are summarized in Table 1. Here, “(M)IP” stands for (mixed-)integer linear program, “LP” for linear program, and “MISOCP” for mixed-integer second order cone program. Furthermore, “NPH” denotes that the problem is NP-hard, while “P” means

	\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	MISOCP	MIP	MIP	MIP
Additional constraints	N	0	1	$n + 1$	N^2	N^2
Additional variables	1	0	1	n	$2n$	$2n$

Table 1: Uncertainty sets in this study.

that it is solvable in polynomial time. 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.

3. Real-World Experiments

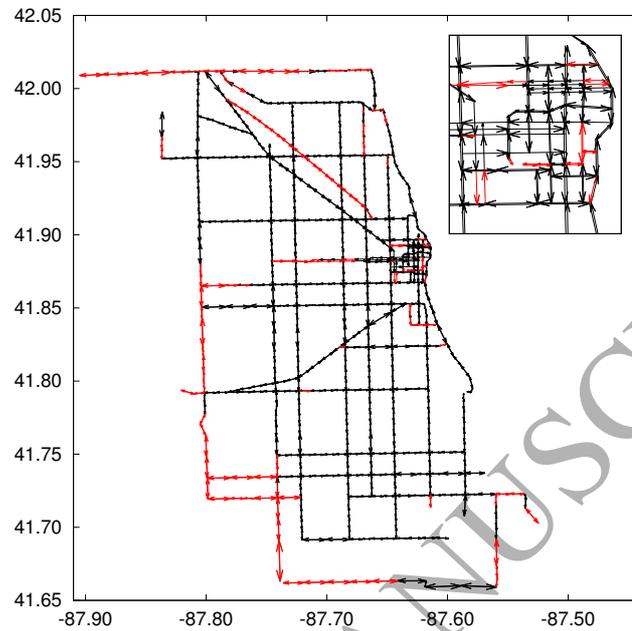
3.1. Data Collection and Cleaning

We used data provided by the City of Chicago¹, which provides a live traffic data interface. The data set consists of traffic updates for every 15-minute interval over a time horizon of 46 days spanning Tuesday morning, March 28, 2017 to Friday evening, May 12, 2017.

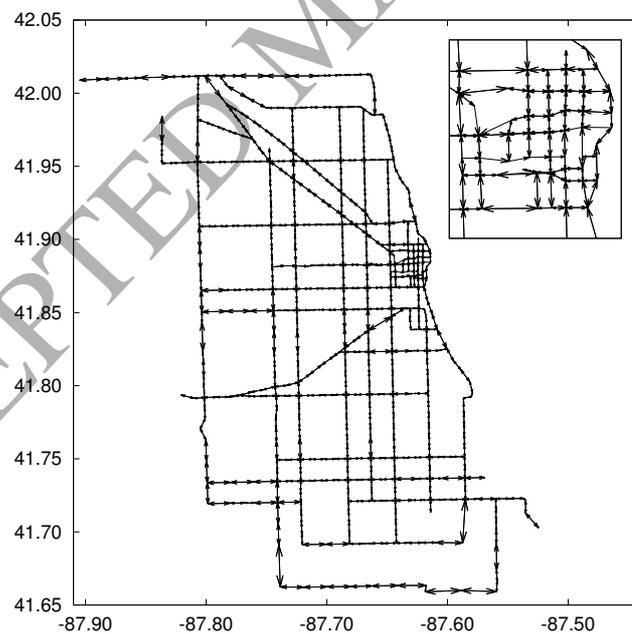
Out of all $46 \cdot 96 = 4,416$ potential observations, only 4,363 had usable data or were recorded due to server downtimes. Every data point contains the traffic speed for a subset of a total of 1,257 segments (that is, a part of a street given through its geographical coordinates). The resulting plot in Figure 2(a) with a zoom-in for the city center. The complete travel speed data set contains a total of 3,891,396 records. There were 1,045 segments where the data was recorded at least once in the 4,363 data points. For nearly 55% of the segments, at least 1,340 data points were recorded, and more than 90% of them have at least 450 data points. For almost all segments at least 400 data points were recorded. We used linear interpolation to fill the missing records keeping in mind that data was collected over time. For segments that did not have any data, we set the travel speed to 20 miles per hour (which is slightly slower than the average speed in the network). Any speed record below 3 miles per hour was set to 3 miles per hour to ensure reasonable travel times. 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 2(b); note that this process slightly simplified the network, but kept its structure intact. The final graph contains 538 nodes and 1,308 arcs. Using arc length and speed, we calculate their respective traversal time for each of the

¹<https://data.cityofchicago.org>



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



(b) Resulting graph model with zoom-in for the city center, longitude versus latitude.

Figure 2: Chicago instance.

4,363 data points. In the following, we refer to 4,363 *scenarios* generated this way.

We then used this full dataset to derive the following subsets aimed at providing robust solutions in different contexts:

1. Find a path that is robust when driving during morning rush hours. We only use scenarios sampled on weekdays from 8am to 10am. These are 271 such scenarios (“mornings dataset”).
2. Find a path that is robust when no additional information is given. We use all 4,363 scenarios (“complete dataset”).

In the following, we present results only for the mornings dataset. Results for the complete dataset can be found in Appendix A. Additionally, we also carried out experiments on an evenings dataset (272 scenarios for weekdays from 4pm to 6pm), a Tuesdays dataset (671 scenarios for all Tuesdays), and a weekends dataset (1,141 scenarios for all Saturdays and Sundays). As results were found to be very similar, they are not presented here. All dataset are available from the corresponding author upon request.

3.2. Data Exploration

Figure 3 visualizes the travel time data used in these experiments plotted against the time of one week, where each point represents the average travel speed over all segments in the network in one observation. The red line shows the hourly average travel speed, and the shaded area represents the corresponding 95% confidence band.

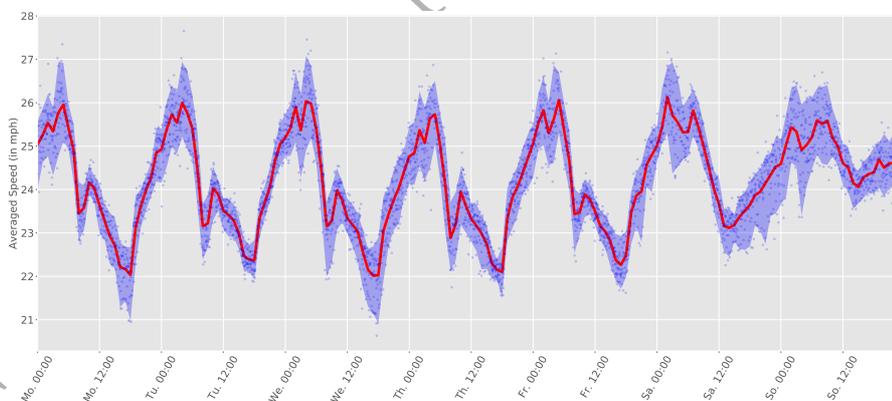


Figure 3: Averaged recorded travel speed data.

We find that there is little correlation between edges (for 99.36% of pairs of edges, the absolute value of correlation is ≤ 0.3), but a strong correlation between points in time (for 87.52% pairs of times, the absolute value of correlation is > 0.6). This gives rise to the question whether optimal paths between fixed pairs of nodes change over time (if not, a robust optimization approach

was not necessary). To consider this aspect, we generated 200 random $s - t$ pairs and calculated an optimal path for each scenario and pair. For only 6 out of 200 pairs, there was one path which was optimal in every scenario. For 128 pairs, there were at least 10 unique shortest paths; and for 63 pairs, there were at least 50 such paths. On average, there are 64.15 different shortest paths per $s - t$ pair.

3.3. Setup

Each uncertainty set can be tuned 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 , $\lambda \in \{0.05, 0.10, \dots, 1.00\}$.
- For \mathcal{U}^E , $\lambda \in \{0.5, 1.0, \dots, 10.0\}$.
- For \mathcal{U}^B , $\Gamma \in \{1, 2, \dots, 20\}$.
- For \mathcal{U}^{PH} , we used columns $\mathbf{q}_1, \mathbf{q}_3, \dots, \mathbf{q}_{39}$.
- For \mathcal{U}^{SPH} , we used columns $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{20}$.

Additionally, we calculate a solution to the average-case scenario $\hat{\mathbf{c}}$. Note that this is a special case of all uncertainty sets presented here if the scaling parameter is sufficiently small. Each uncertainty set is generated using 75% of scenarios sampled uniformly (e.g., 203 out of 271), and we evaluate solutions in-sample and out-sample separately. Furthermore, we generated 200 random $s - t$ pairs uniformly over the node set, and used each of the 6 · 20 methods on the same 200 pairs. Each of our 120 methods, hence, generates $200 \cdot 271 = 54,200$ objective values for the mornings set.

It is non-trivial to assess the quality of these robust solutions, see [13]. If one just uses the average objective value, as an example, then one could as well calculate the solution optimizing 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. Let $\mathcal{X}' = \{\mathbf{x}^1, \dots, \mathbf{x}^P\}$ with $P = 200$ be the set of paths generated by some solution approach.

- The average objective value over all $s - t$ pairs and all scenarios, i.e., the value

$$AV = \frac{1}{P} \sum_{i \in [P]} \left(\frac{1}{N} \sum_{j \in [N]} (\mathbf{c}^j)^t \mathbf{x}^i \right)$$

- The average of the worst-case objective value for each $s - t$ pair, i.e., the value

$$MAX = \frac{1}{P} \sum_{i \in [P]} \left(\max_{j \in [N]} (\mathbf{c}^j)^t \mathbf{x}^i \right)$$

- The average value of the worst 5% of objective values for each $s - t$ pair (as in the *CVaR* measure). Formally, let $C_j^i = \sum_{j \in [N]} (\mathbf{c}^j)^t \mathbf{x}^i$, and let $C_{(j)}^i$ be a sorting of indices such that $C_{(j)}^i \geq C_{(j+1)}^i$ for all $j \in [N - 1]$. Then we use

$$CVaR = \frac{1}{P} \sum_{i \in [P]} \left(\frac{20}{N} \sum_{j \in [N/20]} C_{(j)}^i \right)$$

Note that many more criteria would be possible to use.

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 specialized combinatorial algorithms are available for some problems).

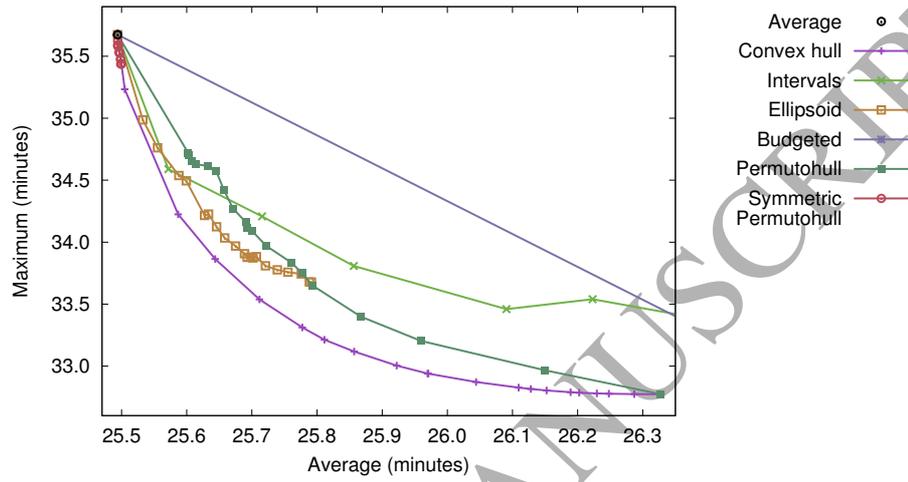
3.4. Results

We present the performance of solutions in Figures 4 and 5. In each plot, the 20 parameter settings that belong to the same uncertainty set are connected by a line, including the average case as a 21st point. They are complemented with Figure 6 showing the total computation times for the methods over all 20 shortest path calculations. For better visibility, Figures 4 and 5 do not show outlying data; plots showing all data can be found in Appendix A.

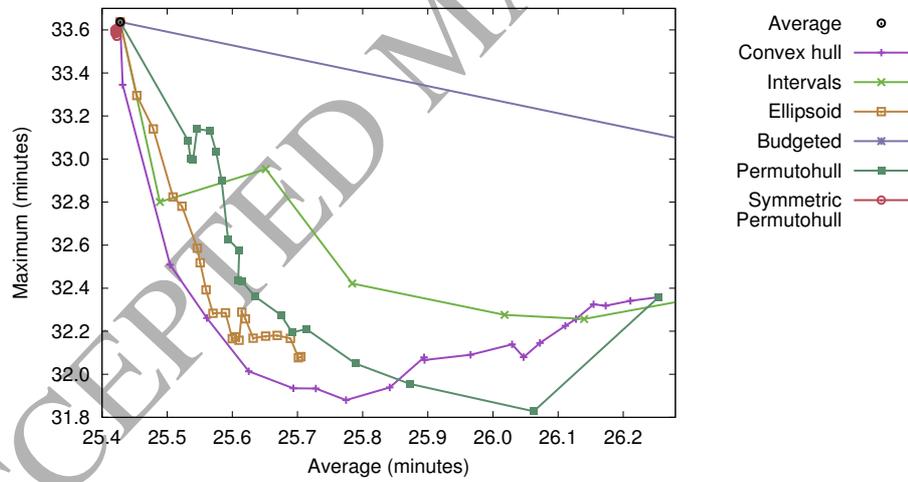
The first set of plots in Figure 4 shows the trade-off between the average and the maximum objective value; the second set of plots in Figure 5 shows the trade-off between the average and the average of the 5% worst objective values. For each case, the in-sample and out-sample performance is shown (i.e., the performance on the set of scenarios that were used in building the uncertainty set, and the performance on the set of scenarios which were not used). All values are in minutes of travel time. 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 are 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 the in-sample performance in Figure 4(a). In general, we find that most methods indeed present a trade-off between average performance and robustness through their scaling parameter (i.e., curves run from top left to bottom right). Solutions calculated using the convex hull dominate the others. Symmetric permutohull solutions tend to focus on a good average performance, while ellipsoid and permutohull solutions show a broader front over the two criteria. Interval and budgeted uncertainty solutions tend to perform worse for larger scaling values, without the desired trade-off property, which confirms previous findings in [12].

Comparing these findings with the out-sample results in Figure 4(b), we see that a general ranking of concepts is kept intact. The solutions generated by the convex hull lose their trade-off property, as they are apparently over-fitted

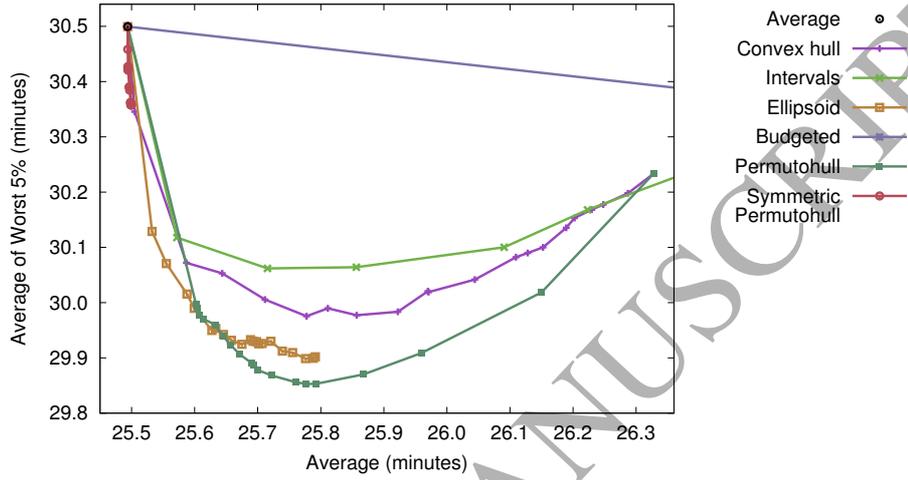


(a) in-sample

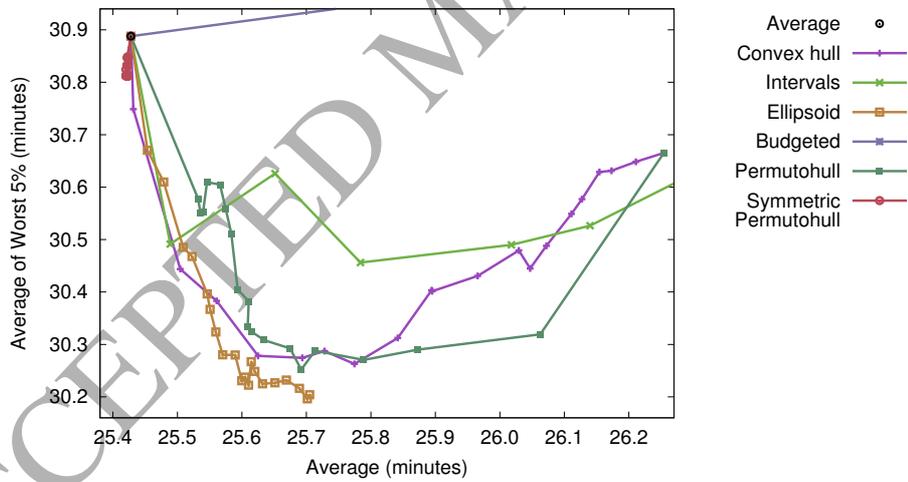


(b) out-sample

Figure 4: Average vs worst-case performance.



(a) in-sample



(b) out-sample

Figure 5: Average vs CVaR performance.

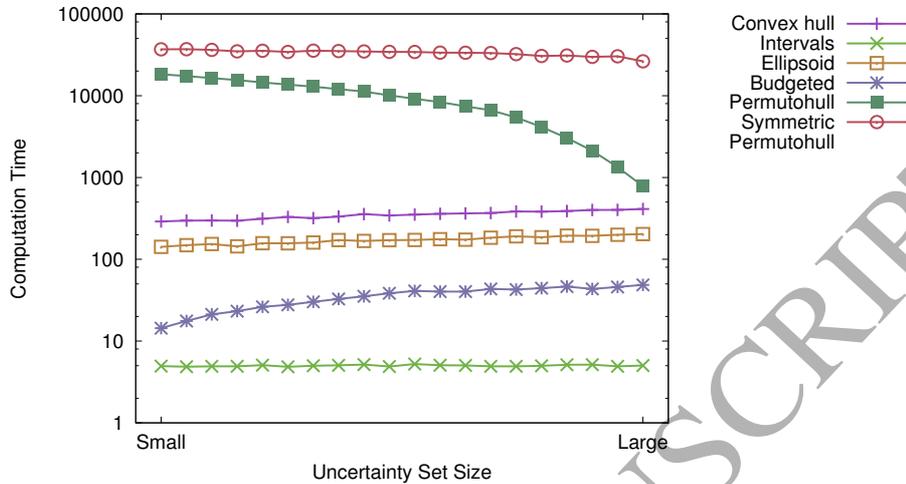


Figure 6: Computation times in seconds.

to the data used in the sample (see also the results in Appendix A). We also find that the interval uncertainty outperforms budgeted uncertainty here, while they showed similar performance in-sample.

We now consider the results presented in Figure 5. Here the average is plotted against the average performance of the 5% worst performing scenarios, averaged over all $s - t$ pairs. While the convex hull solutions showed the best trade-off in Figure 4(a), we find that the permutohull solutions are prominent among the non-dominated points in this case. As before, symmetric permutohull solutions tend to remain on one end of the spectrum with good average performance. Solutions obtained from ellipsoidal uncertainty are still among the best, and stable when considered out-sample (compare Figures 5(a) and 5(b)). Interval uncertainty and in particular budgeted uncertainty do not perform well in comparison with the other approaches.

Regarding computation times (see Figure 6), note that the two polynomially solvable approaches (intervals and budgeted) are also the fastest when using CPLEX; these computation times can be further improved using specialized algorithms. Using ellipsoids is faster than using the convex hull, which is in turn faster than using the symmetric permutohull. For the standard permutohull, the computation times are sensitive to the uncertainty set size; if the \mathbf{q} vector that is used in the model has only few entries, computation times are smaller. This is in line with the intuition that the problem becomes easier if fewer scenarios need to be considered.

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

- Convex hull solutions show good in-sample performance, but are not stable when facing scenarios out of sample.

- Interval solutions do not perform well in general, but are easy and fast to compute, which makes them a reasonable approach, in particular for smaller scalings.
- Budgeted uncertainty does not seem an adequate choice for robust shortest path problems. Scaling interval uncertainty sets gives better results and is easier to use and to solve.
- Ellipsoidal uncertainty solutions have good and stable overall performance and represent a large part of the non-dominated points in our results.
- 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, permutohull and ellipsoidal uncertainty tend to produce solutions with the best trade-off, while being computationally more challenging than most of the other approaches. The algorithmic research for robust shortest path problems with such structure should therefore be studied further. Results on additional experiments leading to the same conclusions can be found in the appendix. In the following section, we consider a variant of ellipsoidal uncertainty where correlation between arcs is ignored.

4. Ellipsoidal Uncertainty Sets

4.1. From General to Axis-Parallel Ellipsoids

Since our experiments have shown that ellipsoidal uncertainty sets are a reasonable choice, we devote this section to these sets. First, we show that changing the general ellipsoid to an axis-parallel ellipsoid by setting all non-diagonal entries of Σ to 0, has almost no effect on the solutions found. Second, we derive a specialized branch-and-bound algorithm for such axis-parallel ellipsoidal uncertainty sets which clearly outperforms the standard approach of using a generic solver. Problems with the same structure were previously considered in [21], where a heuristic method was proposed.

Comparing general and axis-parallel ellipsoids for the experiments presented in Section 3, we find that the maximum deviation over all plotted datapoints is less than 0.002% for average travel times, less than 0.005% for average worst-case values, and less than 0.003% for average CVaR values. That is, plotted in our figures, general and axis-parallel ellipsoids would look indistinguishable. On the other hand, using axis-parallel ellipsoidal uncertainty sets in the robust model instead of general ellipsoids decreases the computation time significantly, see Figure 7. The computation time can be further reduced by the use of specialized algorithms, as shown in the next section.

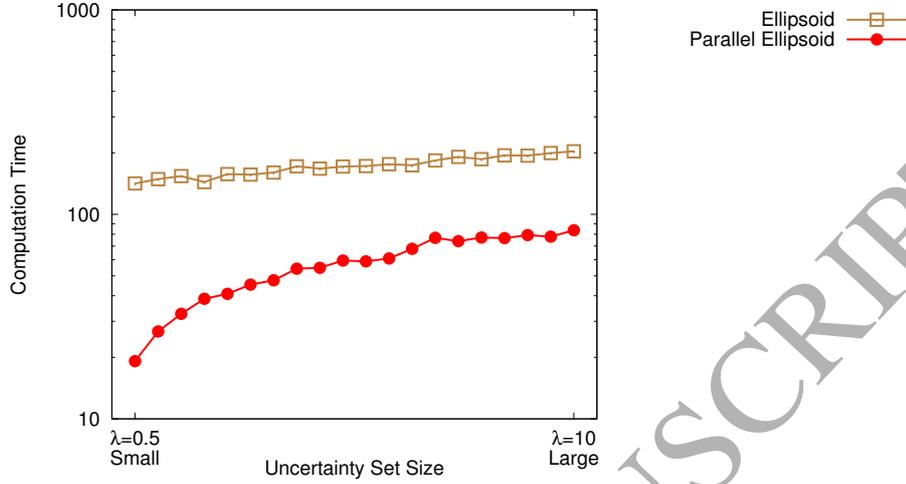


Figure 7: Comparison between general and axis-parallel ellipsoids, computation times in seconds.

4.2. Efficient Algorithm for Axis-Parallel Ellipsoids

4.2.1. A Bicriteria Perspective

In this section, we describe an efficient branch-and-bound algorithm to solve the robust shortest path problem if the uncertainty set is given as an axis parallel ellipsoid. Recall that the mathematical formulation of the problem is

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

where $\boldsymbol{\Sigma}$ is a diagonal matrix specifying the shape and the size of the ellipsoid. Let \mathbf{d} be the diagonal of $\boldsymbol{\Sigma}$. Since \mathbf{x} is a binary vector, we have $x_i^2 = x_i$ and we can simplify the quadratic expression $\mathbf{x}^t \boldsymbol{\Sigma} \mathbf{x}$ to a linear expression $\mathbf{d}^t \mathbf{x}$. Hence the problem can be reduced to

$$\begin{aligned} \min \quad & \hat{\mathbf{c}}^t \mathbf{x} + \sqrt{\mathbf{d}^t \mathbf{x}} \\ \text{s.t.} \quad & \mathbf{x} \in \mathcal{X} \end{aligned}$$

As pointed out in [21], this problem can be transformed to the following bicriteria optimization problem.

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

There it is shown that each optimal solution of the robust problem is an efficient extreme solution of this bicriteria optimization problem. We call a solution \mathbf{x}^*

of this bicriteria optimization problem *efficient extreme* if there exists α_0 and α_1 with $0 \leq \alpha_0 < \alpha_1 \leq 1$ such that for all $\alpha \in [\alpha_0, \alpha_1]$ it holds that there exists no other solution \mathbf{x}' with $(\alpha \mathbf{c} + (1 - \alpha) \mathbf{d})^t \mathbf{x}' < (\alpha \mathbf{c} + (1 - \alpha) \mathbf{d})^t \mathbf{x}^*$. This means that we can find efficient extreme solution by solving the following weighted sum problem which corresponds to a classic shortest path problem.

$$\begin{aligned} \min & (\alpha \hat{\mathbf{c}} + (1 - \alpha) \mathbf{d})^t \mathbf{x} \\ \text{s.t. } & \mathbf{x} \in \mathcal{X} \end{aligned}$$

Hence, the robust solution can be found by computing all efficient extreme solutions of the bicriteria problem. Unfortunately, there is no polynomial bound for the number of efficient extreme solutions of a bicriteria shortest path problem. In fact, it has been shown in [10] that there exist instances of the bicriteria shortest path problem with a subexponential number of efficient extreme solutions. Hence, [21] proposes a heuristic to compute only a subset of all efficient extreme solutions of the bicriteria problem. Among the found solutions, the best is chosen with respect to the robust objective function.

In the following, we present an exact algorithm which is guaranteed to find an efficient extreme solution which is optimal for the robust problem without computing all efficient extreme solutions. Unfortunately, we cannot prove that the number of computed solutions by the exact algorithm is polynomially bounded. However, for real-world or randomly generated instances the number of computed solutions is so small that it can be assumed to be constant. We verify this claim in computational experiments. For convenience, we first present a naive algorithm to compute the complete set of all extreme efficient solutions.

4.2.2. Naive Algorithm

Algorithm 1 Naive Algorithm to Compute all Efficient Extreme Solutions

- 1: Compute $\mathbf{x}_l = \text{arglexmin}_{\mathbf{x} \in \mathcal{X}}(\hat{\mathbf{c}}^t \mathbf{x}, \mathbf{d}^t \mathbf{x})$.
 - 2: Compute $\mathbf{x}_r = \text{arglexmin}_{\mathbf{x} \in \mathcal{X}}(\mathbf{d}^t \mathbf{x}, \hat{\mathbf{c}}^t \mathbf{x})$.
 - 3: **return** EXPLORE($\mathbf{x}_l, \mathbf{x}_r$)
-

Note that the lexmin in Step 1 (or Step 2, respectively) of Algorithm 1 can be found by solving a problem of the form $\min_{\mathbf{x} \in \mathcal{X}}((1 - \epsilon) \hat{\mathbf{c}} + \epsilon \mathbf{d})^t \mathbf{x}$ for a sufficiently small chosen ϵ . The subroutine described as Algorithm 2 recursively finds all efficient extreme solutions. The recursion halts if the found solution is not efficient extreme anymore.

We remark that for each efficient extreme solution \mathbf{x}^* , it is guaranteed that Algorithm 1 either finds \mathbf{x}^* or an alternative solution \mathbf{x}_{alt} with $\hat{\mathbf{c}}^t \mathbf{x}_{\text{alt}} = \hat{\mathbf{c}}^t \mathbf{x}^*$ and $\mathbf{d}^t \mathbf{x}_{\text{alt}} = \mathbf{d}^t \mathbf{x}^*$. Further, it is not guaranteed that all solutions returned by Algorithm 1 are efficient extreme, as the indifference curves $(\hat{\mathbf{c}}^t \mathbf{x}, \sqrt{\mathbf{d}^t \mathbf{x}})$ in the objective space are parabolas. However, all non efficient extreme solutions could be easily removed.

Algorithm 2 Recursive Subroutine for the Naive Algorithm

```

1: procedure EXPLORE( $\mathbf{x}_0, \mathbf{x}_1$ )
2:   Set  $\alpha_m$  such that  $(\alpha_m \hat{\mathbf{c}} + (1 - \alpha_m) \mathbf{d})^t \mathbf{x}_0 = (\alpha_m \hat{\mathbf{c}} + (1 - \alpha_m) \mathbf{d})^t \mathbf{x}_1$ .
3:   Compute  $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}} (\alpha_m \hat{\mathbf{c}} + (1 - \alpha_m) \mathbf{d})^t \mathbf{x}$ 
4:   if  $(\alpha_m \hat{\mathbf{c}} + (1 - \alpha_m) \mathbf{d})^t \mathbf{x}^* < (\alpha_m \hat{\mathbf{c}} + (1 - \alpha_m) \mathbf{d})^t \mathbf{x}_0$  then
5:     return EXPLORE( $\mathbf{x}_0, \mathbf{x}^*$ )  $\cup \{\mathbf{x}^*\} \cup$  EXPLORE( $\mathbf{x}^*, \mathbf{x}_1$ ).
6:   else
7:     return  $\emptyset$ 

```

4.2.3. Improved Algorithm

For the improved algorithm, we first find the two lexicographic minimal solutions \mathbf{x}_l and \mathbf{x}_r as in the naive method (see Figure 8(a)). We denote by $p : \mathcal{X} \rightarrow \mathbb{R}^2$, $p(\mathbf{x}) = (p_1(\mathbf{x}), p_2(\mathbf{x})) = (\hat{\mathbf{c}}^t \mathbf{x}, \mathbf{d}^t \mathbf{x})$ the map from the solution space to the two-dimensional objective space of the bicriteria optimization problem.

From the definition of efficient extreme solutions, it follows that for each efficient extreme solution \mathbf{x}^* , $p(\mathbf{x}^*)$ is contained in the triangle Δ_{unexp} with the vertices $p(\mathbf{x}_l)$, $(p_1(\mathbf{x}_l), p_2(\mathbf{x}_r))$, and $p(\mathbf{x}_r)$ (see Figure 8(a)).

Denote by $\mathbf{x}_{\text{rob}}^* = \operatorname{argmin}(\hat{\mathbf{c}}^t \mathbf{x}_l + \sqrt{\mathbf{d}^t \mathbf{x}_l}, \hat{\mathbf{c}}^t \mathbf{x}_r + \sqrt{\mathbf{d}^t \mathbf{x}_r})$ the current best solution for the robust problem and by OBJ its corresponding objective value. Assume to have a solution \mathbf{x} which improves the current best solution. It must hold that

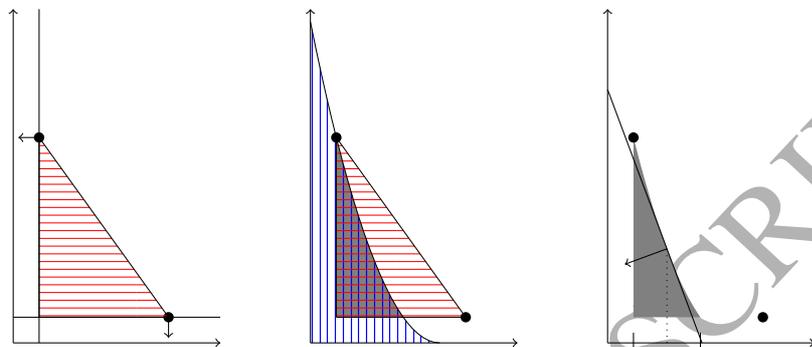
$$\begin{aligned}
& \hat{\mathbf{c}}^t \mathbf{x} + \sqrt{\mathbf{d}^t \mathbf{x}} < \hat{\mathbf{c}}^t \mathbf{x}_{\text{rob}}^* + \sqrt{\mathbf{d}^t \mathbf{x}_{\text{rob}}^*} \\
\Leftrightarrow & \mathbf{d}^t \mathbf{x} < \left(\hat{\mathbf{c}}^t \mathbf{x}_{\text{rob}}^* + \sqrt{\mathbf{d}^t \mathbf{x}_{\text{rob}}^*} - \hat{\mathbf{c}}^t \mathbf{x} \right)^2 \\
\Leftrightarrow & p_2(\mathbf{x}) < (OBJ - p_1(\mathbf{x}))^2
\end{aligned}$$

This means $p(\mathbf{x})$ must be contained in the region $\mathcal{R}_{\text{imp}} = \{z \in \mathbb{R}^2 \mid z_2 < (OBJ - z_1)^2\}$ (see Figure 8(b)).

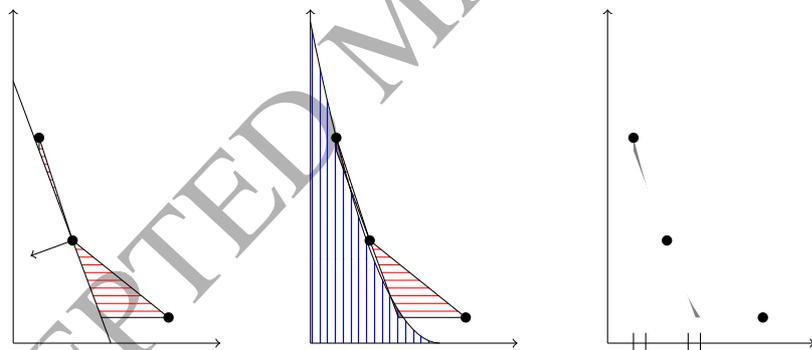
Intuitively, we always have two regions in which we are interested during the algorithm. First, the unexplored region, which may contain efficient extreme solutions which we have not found yet. At the beginning this region corresponds to Δ_{unexp} . Second, the improving region, corresponding to \mathcal{R}_{imp} , which could contain solutions that improve our current best solution. We call the intersection of these two regions \mathcal{R}_{int} in the following. During the algorithm we keep track of \mathcal{R}_{int} . The algorithm then solves, for different values of α , shortest path problems with the objective function $\alpha \cdot \hat{\mathbf{c}}^t \mathbf{x} + (1 - \alpha) \cdot \mathbf{d}^t \mathbf{x}$. After each solution of such a problem, \mathcal{R}_{int} can be shrunk. If \mathcal{R}_{int} is empty, we have proven to have an optimal solution.

Next we have to describe how we generate from \mathcal{R}_{int} the next α for which we solve the weighted sum problem.

At the beginning, we intersect the triangle Δ_{unexp} and \mathcal{R}_{imp} and project to the first axis. This results in the interval which we use to initialize a list of intervals \mathcal{L} (see Figure 8(c)). The idea of the improved algorithm is then to shrink, split or remove intervals from \mathcal{L} until \mathcal{L} is empty.



- (a) To initialize the algorithm we compute the two solutions which minimize the first and second objective function. The unexplored region which might contain efficient extreme solutions is marked with horizontal lines.
- (b) The region of improvement which might contain solutions that improve the actual best solution is marked with vertical lines. The intersection of both regions, shown in gray, defines the area in which we try to find solutions.
- (c) Projecting the gray area to the first axis defines the first interval the algorithm tries to shrink. The algorithm computes the midpoint of the interval and projects it to the parabola to obtain the next direction of optimization



- (d) A new solution is found. This solution intersects the previous unexplored region with one half-space, which results in two smaller triangles.
- (e) Intersecting the region of improvement with the new unexplored region leads to two small areas.
- (f) Projecting the two gray areas to the first axis defines two smaller intervals. The algorithm proceeds by shrinking, splitting or removing these intervals until the gray area corresponds to the empty set

Figure 8: Visualization of the improved algorithm.

To do so, the improved algorithm picks an interval $I = [a, b]$ from \mathcal{L} , computes its midpoint $m = 0.5(a + b)$ and defines a local approximation of the boundary of \mathcal{R}_{imp} at $(m, (OBJ - m)^2)$ (see Figure 8(c)). The slope of the obtained line l is $2(m - OBJ)$. We set $\alpha_m = 2(m - OBJ)/(2(m - OBJ) - 1)$ to optimize in the direction which is perpendicular to l (see Figure 8(c)).

We then compute $\mathbf{x}_{\text{new}} = \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}} (\alpha_m \hat{\mathbf{c}} + (1 - \alpha_m) \mathbf{d})^t \mathbf{x}$. We know for each other solution \mathbf{x} that

$$\begin{aligned} & \alpha_m \hat{\mathbf{c}}^t \mathbf{x}_{\text{new}} + (1 - \alpha_m) \mathbf{d}^t \mathbf{x}_{\text{new}} \leq \alpha_m \hat{\mathbf{c}}^t \mathbf{x} + (1 - \alpha_m) \mathbf{d}^t \mathbf{x} \\ \Leftrightarrow & \frac{\alpha_m}{1 - \alpha_m} (\hat{\mathbf{c}}^t \mathbf{x}_{\text{new}} - \hat{\mathbf{c}}^t \mathbf{x}) + \mathbf{d}^t \mathbf{x}_{\text{new}} \leq \mathbf{d}^t \mathbf{x} \\ \Leftrightarrow & \frac{\frac{2(m - OBJ)}{2(m - OBJ) - 1}}{1 - \frac{2(m - OBJ)}{2(m - OBJ) - 1}} (\hat{\mathbf{c}}^t \mathbf{x}_{\text{new}} - \hat{\mathbf{c}}^t \mathbf{x}) + \mathbf{d}^t \mathbf{x}_{\text{new}} \leq \mathbf{d}^t \mathbf{x} \\ \Leftrightarrow & -2(m - OBJ)(\hat{\mathbf{c}}^t \mathbf{x}_{\text{new}} - \hat{\mathbf{c}}^t \mathbf{x}) + \mathbf{d}^t \mathbf{x}_{\text{new}} \leq \mathbf{d}^t \mathbf{x} \\ \Leftrightarrow & 2(m - OBJ)p_1(\mathbf{x}) - 2(m - OBJ)\hat{\mathbf{c}}^t \mathbf{x}_{\text{new}} + \mathbf{d}^t \mathbf{x}_{\text{new}} \leq p_2(\mathbf{x}) \end{aligned}$$

Hence, we conclude that for all efficient extreme solutions \mathbf{x}^* , $p(\mathbf{x}^*)$ must lie above the line through $p(\mathbf{x}_{\text{new}})$ with slope $2(m - OBJ)$. Hence, we can exclude a half space from the unexplored region (see Figure 8(d)) and shrink, split or remove intervals contained in \mathcal{L} (see Figure 8(e)). Further, it might happen that \mathbf{x}_{new} improves the actual best solution in this case we update OBJ and \mathcal{R}_{imp} and consequently all intervals in \mathcal{L} .

If the algorithm has reduced \mathcal{L} to the empty set the current best solution is indeed the optimal solution of the robust optimization problem.

4.2.4. Computational Experiments for the Improved Algorithm

We test the performance of the branch-and-bound algorithm on grid graphs, using the same computational environment as in Section 3. We used the LEMON graph library v.1.3.1 [16] to solve the classic shortest path problems that needs to be solved during the branch-and-bound algorithm. The goal is to find a path from the upper left corner to the lower right corner of the grid. For each arc we chose 50 values uniform at random from $\{1, \dots, 100\}$. Further, we fit an axis-parallel ellipsoidal uncertainty set to these points as described previously. We vary the grid size from a 5×5 grid to a 20×20 grid. For each grid size we create 100 instances and solve them in two ways: By using CPLEX to solve the resulting MISOCP and by the proposed branch-and-bound algorithm. The averaged computation times are shown in Figure 9.

It can be seen that our approach outperforms CPLEX by several orders of magnitude (note the logarithmic vertical scale). While the computation times for CPLEX scale exponentially with the graph size, this is not observed for our method. We show the average number of shortest path calculations required by our method in Table 2, where we increase the grid size to 100×100 . It can be seen that on average only very few calls are required, and the increase is

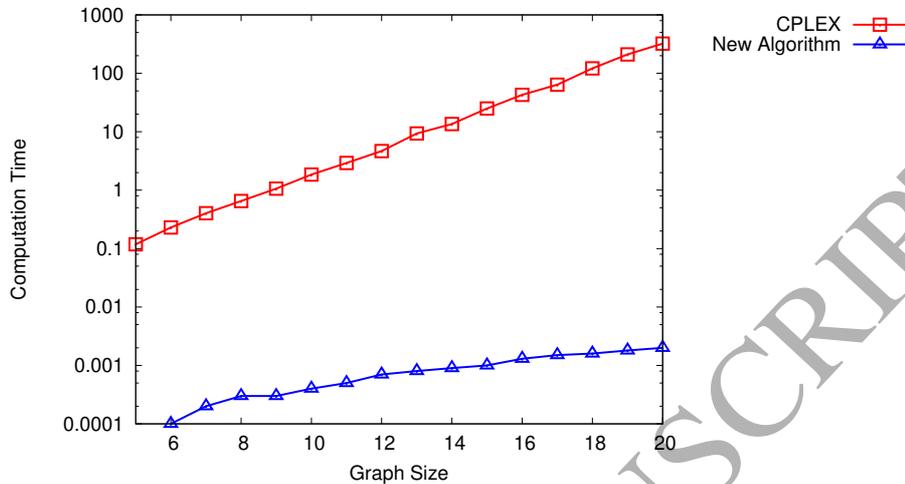


Figure 9: Comparison of average computation times for grid graphs.

slow. By using further improved algorithms for shortest path calculation in road networks (see [2]), the application of our method in real-time route planning is within reach.

Instance size	SP comp.
5×5	3.625
10×10	3.929
20×20	4.251
50×50	4.778
100×100	5.127

Table 2: Average number of shortest path computations for different grid sizes.

Finally, we revisit the computation times for the real-world instance from our previous experiment (Figure 7). Figure 10 shows the performance of our method for comparison.

5. Conclusions

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 in-sample and out-sample performance using different performance indicators. Naturally, conclusions can only be drawn within the reach of the available data. It remains to be seen how the considered uncertainty sets perform on other datasets for robust shortest paths.

We have observed that using ellipsoidal uncertainty sets provides high-quality solutions with less computational effort than for the permutohull. If one uses

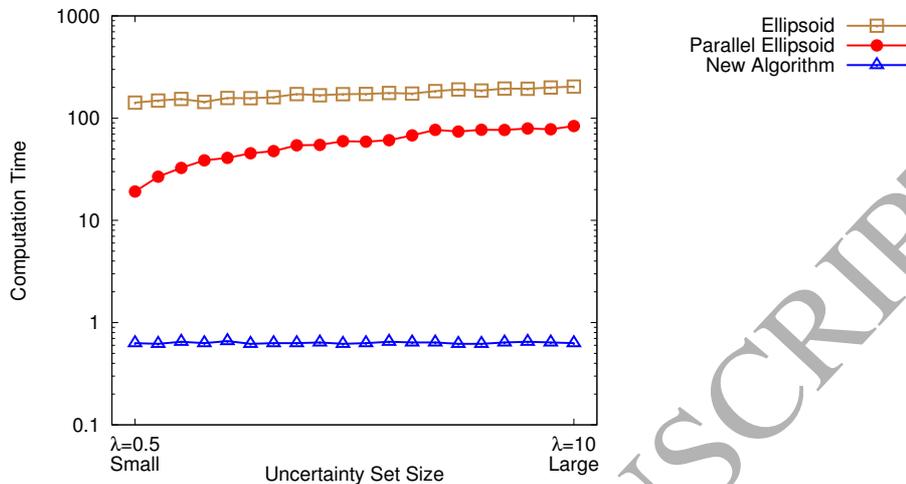


Figure 10: Comparison between general and axis-parallel ellipsoids, computation times in seconds.

only the diagonal entries of the matrix Σ , then one ignores the data correlation in the network, but the solution quality remains roughly the same. For the resulting problem, a specialized branch-and-bound algorithm was developed that is able to reduce computation times considerably compared to CPLEX. In fact, the computational effort to solve this problem is comparable to the complexity of solving a few classic shortest path problems, which even makes the application on real-time navigation devices a possibility.

References

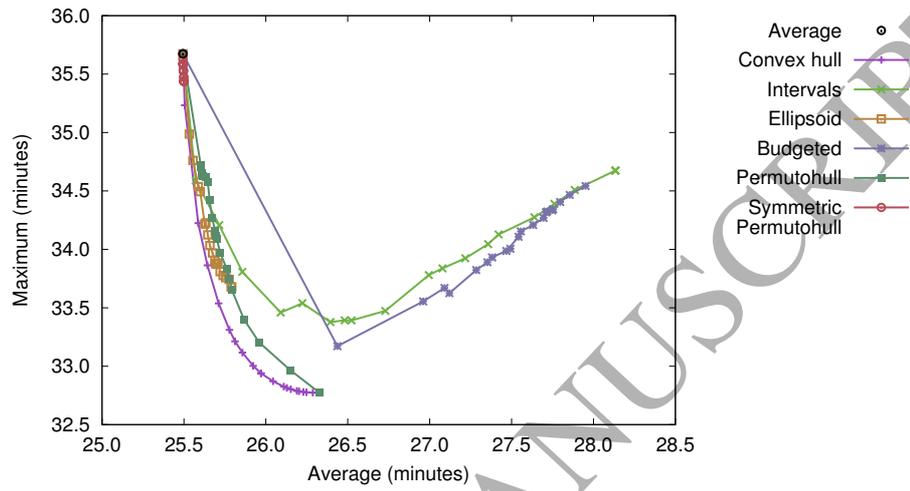
References

- [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.

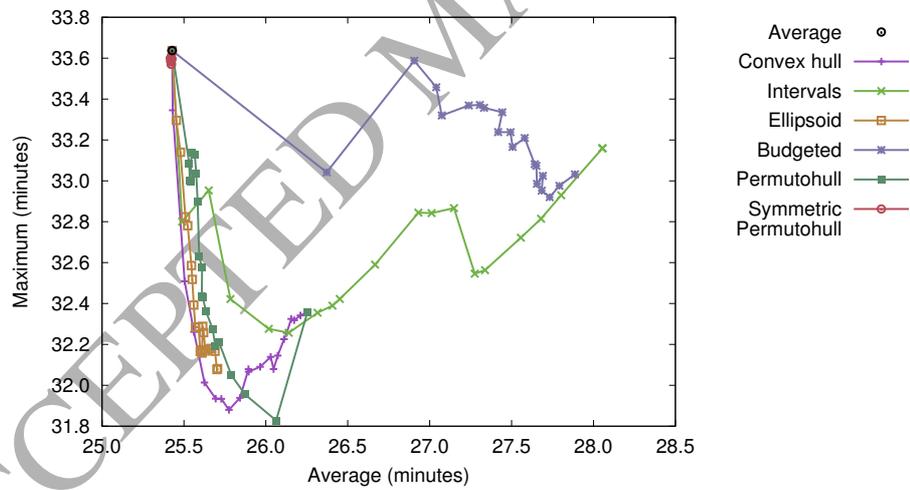
- [5] Dimitris Bertsimas and David B Brown. Constructing uncertainty sets for robust linear optimization. *Operations research*, 57(6):1483–1495, 2009.
- [6] Dimitris Bertsimas, Vishal Gupta, and Nathan Kallus. Data-driven robust optimization. *Mathematical Programming*, pages 1–58, 2017. Online first.
- [7] Dimitris Bertsimas and Melvyn Sim. Robust discrete optimization and network flows. *Mathematical programming*, 98(1):49–71, 2003.
- [8] Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations research*, 52(1):35–53, 2004.
- [9] Christina Büsing. Recoverable robust shortest path problems. *Networks*, 59(1):181–189, 2012.
- [10] Patricia J. Carstensen. *The Complexity of Some Problems in Parametric Linear and Combinatorial Programming*. PhD thesis, University of Michigan, 1983.
- [11] 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.
- [12] André Chassein and Marc Goerigk. A bicriteria approach to robust optimization. *Computers & Operations Research*, 66:181–189, 2016.
- [13] André Chassein and Marc Goerigk. Performance analysis in robust optimization. In *Robustness Analysis in Decision Aiding, Optimization, and Analytics*, pages 145–170. Springer, 2016.
- [14] André Chassein and Marc Goerigk. Variable-sized uncertainty and inverse problems in robust optimization. *European Journal of Operational Research*, 2017. Available online, to appear.
- [15] IBM ILOG CPLEX. 12.6. *CPLEX Users Manual*, 2014.
- [16] Balázs Dezsó, Alpár Jüttner, and Péter Kovács. LEMON – an open source C++ graph template library. *Electronic Notes in Theoretical Computer Science*, 264(5):23–45, 2011.
- [17] Trivikram Dokka and Marc Goerigk. An Experimental Comparison of Uncertainty Sets for Robust Shortest Path Problems. In Gianlorenzo D’Angelo and Twan Dollevoet, editors, *17th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2017)*, volume 59 of *OpenAccess Series in Informatics (OASICs)*, pages 16:1–16:13, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [18] Marc Goerigk and Anita Schöbel. Algorithm engineering in robust optimization. In *Algorithm engineering*, pages 245–279. Springer, 2016.

- [19] 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.
- [20] 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.
- [21] 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.
- [22] Gang Yu and Jian Yang. On the robust shortest path problem. *Computers & Operations Research*, 25(6):457–468, 1998.

Appendix A. Additional experimental results



(a) in-sample



(b) out-sample

Figure A.11: Average vs worst-case performance, mornings dataset.

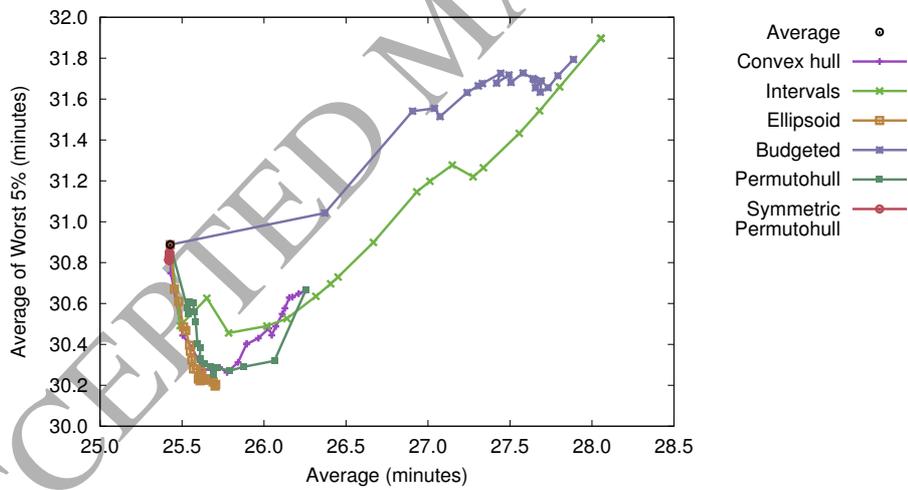
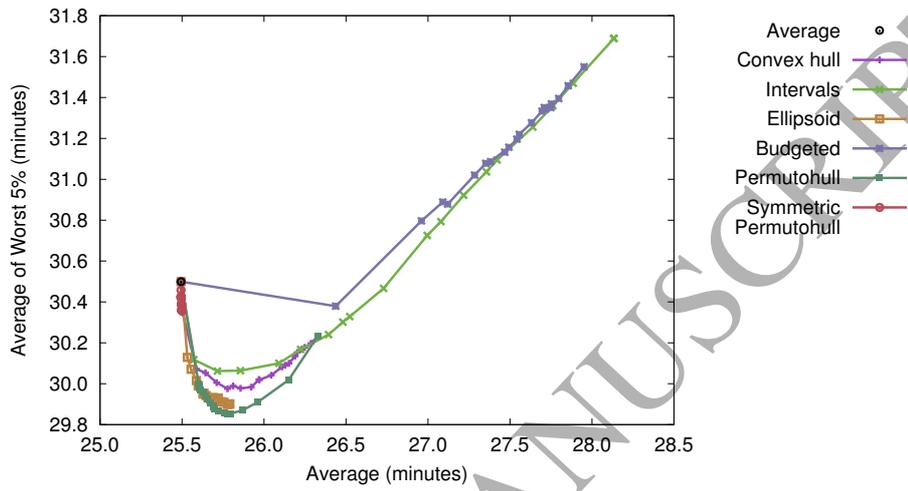
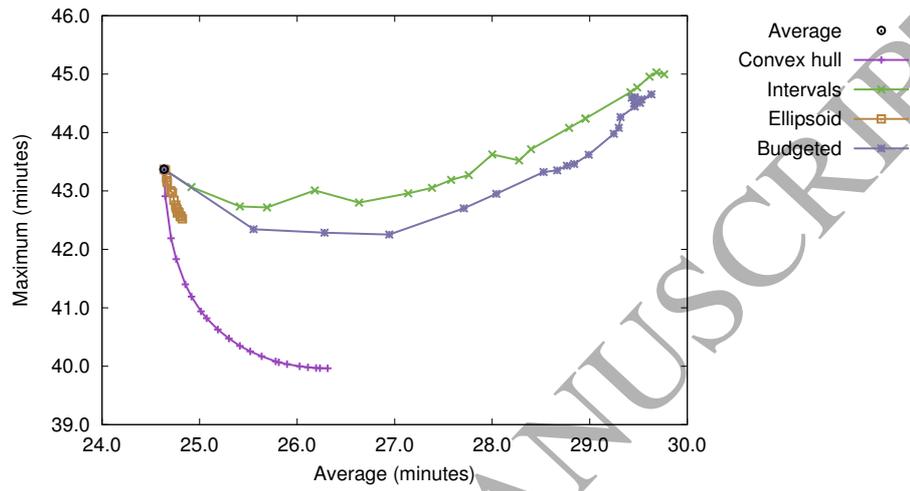
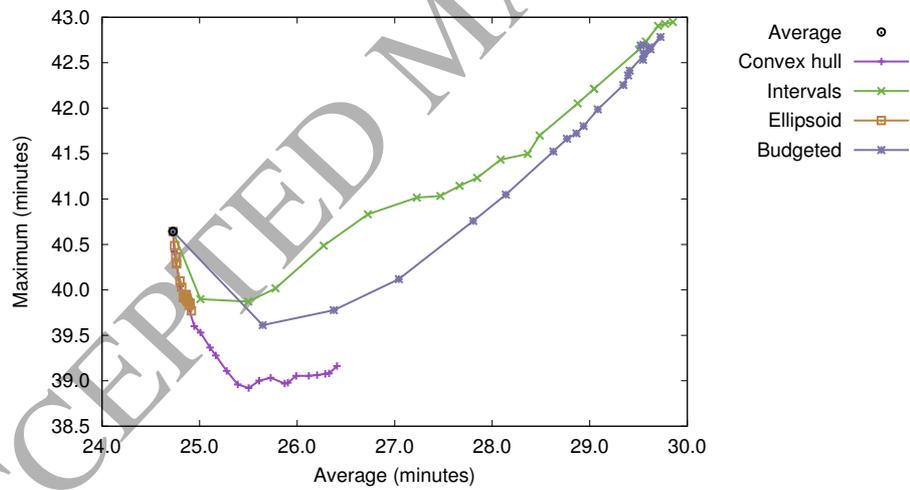


Figure A.12: Average vs CVaR performance, mornings dataset.

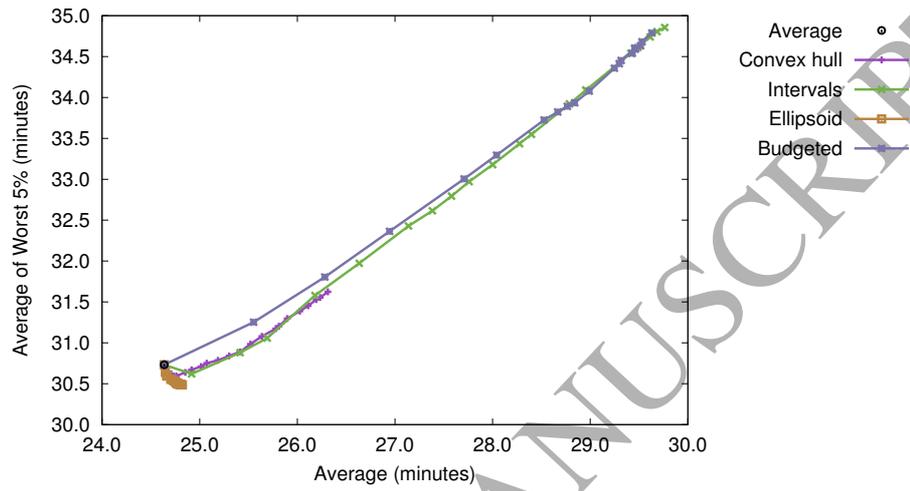


(a) in-sample

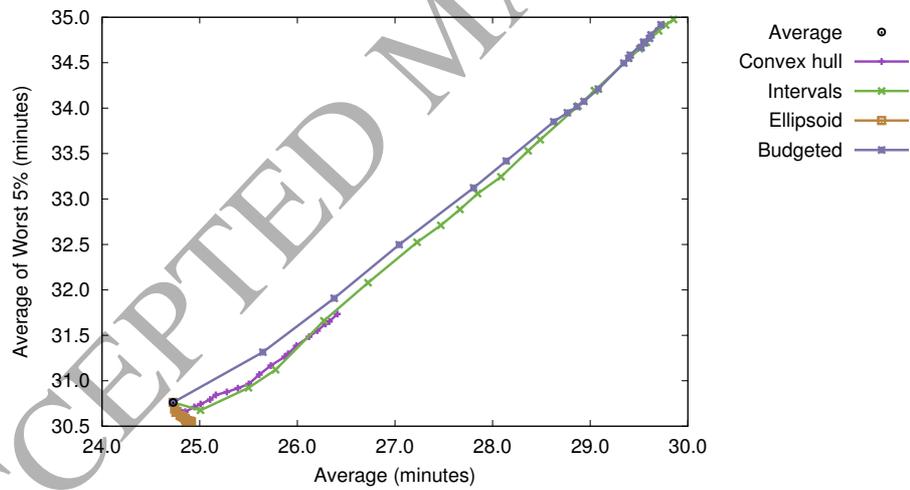


(b) out-sample

Figure A.13: Average vs worst-case performance, complete dataset.



(a) in-sample



(b) out-sample

Figure A.14: Average vs CVaR performance, complete dataset.