

Efficient Path Routing Over Road Networks in the Presence of Ad-Hoc Obstacles

Ahmed Al-Baghdadi^{a,b}, Xiang Lian^{a,*}, En Cheng^c

^a Department of Computer Science, Kent State University, Kent, OH 44242, USA

^b Minister of Education in Al-Qadisiyah, Al-Diwaniyah 58001, Iraq

^c Department of Computer Science, the University of Akron, Akron, OH 44325, USA

ARTICLE INFO

Article history:

Received 5 March 2019

Received in revised form 13 July 2019

Accepted 2 September 2019

Available online 8 October 2019

Recommended by Yannis Manolopoulos

Keywords:

Ad-hoc weather-based obstacles

Ad-hoc keyword-based obstacles

Road networks

Path routing query over road networks with ad-hoc obstacles

ABSTRACT

Nowadays, the path routing over road networks has become increasingly important, yet challenging, in many real-world applications such as location-based services (LBS), logistics and supply chain management, transportation systems, map utilities, and so on. While many prior works aimed to find a path between a source and a destination with the smallest traveling distance/time, they do not take into account the quality constraints (e.g., obstacles) of the returned paths, such as uneven roads, roads under construction, and weather conditions on roads. Inspired by this, in this paper, we consider two types of ad-hoc obstacles, keyword-based and weather-based obstacles, on road networks, which can be used for modeling roads that the returned paths should not pass through. In the presence of such ad-hoc obstacles on roads, we formulate a *path routing query over road networks with ad-hoc obstacles* (PRAO), which retrieves paths from source to destination on road networks that do not pass ad-hoc keyword and weather obstacles and have the smallest traveling time. In order to efficiently answer PRAO queries, we design effective pruning methods and indexing mechanism to facilitate efficient PRAO query answering. Extensive experiments have demonstrated the efficiency and effectiveness of our approaches over real/synthetic data sets.

© 2019 Elsevier Ltd. All rights reserved.

1. Introduction

In many real-world applications [1–8] such as location-based services (LBS), supply chain management, map services, transportation systems, and so on, one classical and important problem is to find the best path (e.g., with the smallest traveling time or distance) between a source and a destination over road networks. Fig. 1 shows an example of a road network, where line segments (e.g., v_1v_2) represent roads, and nodes (e.g., v_1) denote intersection points of roads. In this example, given two points, *src* and *dst*, on road networks, the path routing task is to find a path from source *src* to destination *dst* with the smallest traveling time.

Different from prior works on the path routing problem, in this paper, we will take into account two specific types of obstacles on road networks, keyword-based and weather-based obstacles, and retrieve a path with the smallest traveling time that does not pass through some user-specified ad-hoc obstacles (or road segments).

Below, we give an example of path routing in the real applications of transferring chemical supplies.

Example 1 (Path Routing for Chemical Transportation). As illustrated in Fig. 1, assume that a company wants to deliver dangerous chemical compounds from a source, *src*, to a destination, *dst*, on road networks. Due to the potential spread of hazardous chemicals, this company would like to schedule an optimal path to transfer chemical supplies, which satisfies the following conditions: (1) the path should not pass by those roads close to areas with high population densities; (2) the wind speed on the path should not exceed 40 mph; and (3) the retrieved path has the smallest traveling time.

In this example, the company can online specify obstacle keywords such as “metropolitan” (i.e., ad-hoc keyword-based obstacles) and require wind speeds on roads smaller than or equal to 40 mph (i.e., ad-hoc weather-based obstacles). Given road networks in Fig. 1 and weather conditions in Table 1, we can obtain a path that does not visit roads with obstacle keyword “metropolitan”, with wind speeds always not exceeding 40 mph of high confidences, and having the smallest traveling time. In this example, one candidate path can be $src \rightarrow v_1 \rightarrow v_3 \rightarrow v_5 \rightarrow v_6 \rightarrow dst$. ■

In the example above, keyword-based obstacles indicate properties of roads, such as “uneven roads”, “roads under construction”, “roads close to downtown areas”, “roads with falling rocks”, “roads with deer”, and so on. In Fig. 1, edge $e (= v_4v_6)$ has

* Corresponding author.

E-mail addresses: aalbaghd@kent.edu (A. Al-Baghdadi), xliao@kent.edu (X. Lian), echeng@uakron.edu (E. Cheng).




Fig. 1. An example of a road network, illustrating wind speeds of vertex v_2 at 08:00 am and 09:00 am, and edge $e (= v_4v_6)$ passing through a metropolitan area.

Table 1
Weather forecasting in the example of Fig. 1.

V. ID	Wind Speed at 8am	Conf.	Wind Speed at 9am	Conf.
v_0	10 mph	0.9	15 mph	1
v_1	15 mph	0.8	20 mph	0.9
v_2	50 mph	0.9	65 mph	0.8
v_3	20 mph	0.9	25 mph	0.8
v_4	35 mph	0.9	40 mph	0.9
v_5	20 mph	1	25 mph	0.9
v_6	10 mph	0.8	15 mph	1
v_7	10 mph	0.9	10 mph	1

a keyword “metropolitan”. On the other hand, weather-based obstacles are highly related to the predicted weather conditions in the future, such as temperature, humidity, or wind speeds (as depicted in Table 1). Moreover, due to the accuracy of weather forecasting, each predicted weather value is often associated with a confidence. In Table 1, the wind speed at vertex v_2 is 50 mph at 8am with probability 0.9, and 65 mph at 9am with probability 0.8.

Example 2 (Supply Chain Management). In a recent study by the Atmospheric and Environmental Research Center [9], thirty percent (30%) of the U.S. gross domestic product is affected by the weather condition. One of the main weather effects comes from delivery issues. According to the article, weather conditions can make it hazardous and even impossible for delivery trucks to transport products to distributors and retailers. Weather conditions such as snow, ice, high temperature, high wind speed, and heavy rain can slow and stop transportation, making products such as groceries unavailable when they are most needed.

Therefore, in the real application of supply chain management, it is very important for many users (such as logistics companies or delivery drivers) to find an optimal path between source and destination for the product delivery, by considering weather conditions and road status. For example, a user can specify constraints (i.e., weather-based obstacle) with respect to temperature (e.g., temperature $< 45^\circ$), wind speed (e.g., wind speed < 40 mph), and/or any other weather-related obstacles at the query time, such that the retrieved path will avoid passing through roads that do not satisfy the specified weather conditions. Furthermore, the user can specify some keywords to be avoided in the retrieved path (i.e., keyword-based obstacles). For instance, the truck driver can specify “city road” as the obstacle keyword(s) to avoid roads passing through cities in the path finding process. ■

Note that, the aforementioned keyword-based and weather-based obstacles are ad-hocly given, in the sense that different path routing problems may specify different obstacle criteria, including an arbitrary set of user-specified obstacle keywords and domain-specific weather constraints for the paths to be returned. Thus, in the presence of such ad-hoc obstacles, it is very

challenging to effectively and efficiently compute the path over large-scale road networks.

Inspired by the motivation examples above, in this paper, we formalize a *path routing query with ad-hoc obstacles* (PRAO), which retrieves a path that does not pass through ad-hoc keyword-based and weather-based obstacles and achieves the smallest traveling time. As mentioned above, the PRAO problem has many potential users (e.g., logistics companies, delivery drivers, etc.) and high/frequent demands (e.g., map utility and services, mobile APP requests by drivers or logistics managers).

Due to the large amount of weather/keyword information over large-scale road networks (exponential number of possible worlds for the predicted weather data), it is rather challenging to efficiently process PRAO queries. One straightforward method to answer PRAO queries is to enumerate all possible paths between *src* and *dst* (e.g., via A^* algorithm [1]), check keyword-based and weather-based obstacle constraints on each road segment of paths, and return the one with the smallest traveling time. However, this straightforward method is not time-efficient, due to a large number of possible paths and the costly computations of checking weather predictions for any point on paths. Therefore, it cannot efficiently handle frequent query requests from users, and we need to design specific techniques (e.g., pruning/indexing/query processing) to tackle the PRAO problem, which are the main contributions of our work.

To tackle the PRAO problem, in this paper, we propose three effective pruning methods, keyword-based, weather-based, and traveling time pruning, to filter out false alarms of candidate paths and reduce the PRAO search space. We can ensure that our proposed method does not miss any optimal route nor produces false dismissals (i.e., we always prune those false alarms of paths containing ad-hoc keyword-based obstacles and/or weather-based obstacles with high confidences, and retrieve actual PRAO answers satisfying PRAO query predicates). Furthermore, we design a variant of R^* -tree index over road networks with obstacles, which supports dynamic updates of weather information. That is, we index road-network edges in a variant of R^* -tree where each tree node represents a group of spatially closed edges (roads), and all the information about weather predictions and keywords is summarized and stored in each node. Finally, we propose an efficient algorithm to facilitate PRAO query answering via the index. Our proposed pruning methods w.r.t. index nodes can guarantee no false dismissals.

In this paper, we make the following contributions.

1. We formulate the PRAO problem over road networks with ad-hoc obstacles in Section 2, and present the weather estimation for any location on roads in Section 3.
2. We present three effective pruning strategies to reduce the PRAO search space in Section 4.
3. We propose effective indexing and efficient PRAO query answering algorithms in Sections 5 and 6, respectively.

4. We demonstrate through extensive experiments the efficiency and effectiveness of our proposed PRAO query answering approach in Section 7.

Section 8 discusses related works on queries over road networks, probabilistic data management, and queries in the presence of obstacles. Section 9 concludes the paper.

2. Problem definition

2.1. Data model for road networks

First, we formally present the data model for road networks as follows.

Definition 1 (Road Networks). A road network is represented by a graph $G = (V, E, \phi)$, where V is a set of vertices v_1, v_2, \dots , and $v_{|V|}$, E is a set of edges e_1, e_2, \dots , and $e_{|E|}$, and ϕ is a mapping function: $V \times V \rightarrow E$, where each edge $e_i (= v_j v_k)$ is associated with a traveling time $e_i.w$ and a traveling distance $e_i.len (= dist(v_j, v_k))$.

In Definition 1, edges e_i (for $1 \leq i \leq |E|$) in a road network G represent roads, whereas vertices v_j (for $1 \leq j \leq |V|$) correspond to road intersection points. Fig. 1 shows an example of road networks, where vertices $v_0 \sim v_7$ represent intersection points on road networks. Edge $e = v_4 v_6$ is a road segment connecting two intersection points v_4 and v_6 , associated with: (1) the traveling time $e.w = 0.2$ hour, and (2) the edge length, $e.len = dist(v_4, v_6) = 30$ miles.

For simplicity, in this paper, we consider a *static* traveling time $e_i.w$ for each road segment e_i . For time-varying traveling times on roads, the data model is more complex, and the problem over such road networks is more challenging and worthy of investigation. We would like to leave the interesting topic of considering our problem over road networks with dynamic traveling times on roads as our future work.

2.2. Data model for ad-hoc obstacles

In this subsection, we provide formal definitions of two types of ad-hoc obstacles on road networks, ad-hoc keyword-based and weather-based obstacles.

Ad-Hoc Keyword-based Obstacles: We define ad-hoc keyword-based obstacles below.

Definition 2 (Ad-Hoc Keyword-based Obstacles). On road networks G , each edge e_i is associated with a set, $e_i.K$, of obstacle keywords. Given a set, S , of user-specified keywords, an edge e_i on road network G is said to be an *ad-hoc keyword-based obstacle*, if $e_i.K$ contains some obstacle keywords in S (i.e., $e_i.K \cap S \neq \emptyset$).

In particular, ad-hoc keyword-based obstacles can be on roads (edges) that pass through specific sites (e.g., metropolitan areas with high population densities), or on roads that travelers prefer not to pass through (e.g., uneven roads).

In Fig. 1, the edge $v_4 v_6$ (i.e., e) has a keyword set {metropolitan}. If a user-specified set, S , of obstacle keywords is {uneven, metropolitan}, then edge $v_4 v_6$ contains an ad-hoc keyword-based obstacle.

Ad-Hoc Weather-based Obstacles: Next, we give the definition of ad-hoc weather-based obstacles.

Definition 3 (Weather Predictions). Weather predictions for a point $o_i \in e_i$ are given by a set, $W(o_i)$, of quadruples, in the form $(W_type, W_val(o_i, t_i), p(o_i, t_i), t_i)$, where W_type represents the type of weather conditions, $W_val(o_i, t_i)$ shows the possible value of the weather condition at the timestamp t_i , and $p(o_i, t_i)$ denotes the confidence of weather value $W_val(o_i, t_i)$.

In Definition 3, the type, W_type , of weather conditions can be wind speed, temperature, and so on. As an example in Fig. 1, the quadruple, $(wind\ speed, 50\ mph, 0.9, 8:00)$, at vertex v_2 indicates that the predicted wind speed (i.e., W_type) 50 mph at 8am is accurate, with prediction confidence 0.9.

Note that, here we can obtain hourly weather prediction data from data sources such as Dark Sky [10]. For instance, highway I24 might have a 70% probability of wind speed 50 mph at timestamp 5am. Moreover, weather data are dynamically updated on an hourly basis. That is, the newly predicted weather data arrive every hour, whereas expired weather data can be removed from the system.

Furthermore, it is not feasible to predict the weather information at every point o_i of road-network edges e_i , thus, in this work, we assume that we can obtain the predicted weather information at the two endpoints v_j and v_k of each road-network edge $e_i (= v_j v_k)$. We can estimate weather conditions (as well as their confidences) at any point o_i on the edge e_i by using weather interpolation techniques. In particular, for a moving vehicle on the edge $e_i (= v_j v_k)$ departing from v_j at a timestamp t_{dep} toward v_k (with an expected arrival time t_{arr}), we can interpolate/estimate weather conditions on any points $o_i \in e_i$ at any timestamp t ($\in [t_{dep}, t_{arr}]$), according to the weather information known at v_j and v_k , which will be discussed later in Section 3.

Definition 4 (Ad-Hoc Weather-based Obstacles). Given weather predictions $W(o_i)$ on road network G , an ad-hoc weather threshold ϵ , and a probabilistic threshold α , any point o_i on edge e_i is said to be an *ad-hoc weather-based obstacle* at a timestamp t_i , if it holds that $Pr\{W_val(o_i, t_i) > \epsilon\} \geq \alpha$.

Intuitively, a point o_i on an edge e_i is an ad-hoc weather-based obstacle, if and only if the value of the weather condition (e.g., a wind speed 50 mph or a temperature 50° C) at o_i is greater than a threshold ϵ with high confidence (i.e., $Pr\{W_val(o_i, t_i) > \epsilon\} \geq \alpha$). Due to bad weather at o_i or the required weather conditions on the traveling path, those edges e_i containing weather-based obstacle o_i will be prohibited to be passed through.

2.3. The PRAO problem definition

In this subsection, we define the path routing over road networks in the presence of ad-hoc keyword-based and weather-based obstacles (i.e., the PRAO problem).

Definition 5 (Path Routing Over Road Networks with Ad-Hoc Obstacles, PRAO). Assume that we have a road network G with weather conditions $W(v_j) = \{(W_type, W_val(v_j, t_j), p(v_j, t_j), t_j)\}$ at any vertex $v_j \in V$ in G at future timestamps t_j , a query weather type W_type , a weather threshold ϵ , a probabilistic threshold α , a set, S , of obstacle keywords, a source vertex src , and a destination vertex dst . A *path routing query with ad-hoc obstacles* (PRAO) finds the best path, $Path$, from src to dst that satisfies the following three criteria:

- for any point o_i on path, $Path$, with arrival time t_i , it holds that:

$$Pr\{W_val(o_i, t_i) > \epsilon\} < \alpha; \quad (1)$$
- for any obstacle keyword $s \in S$ and edge $e_i \in Path$, it holds that $e_i.K$ does not contain s (i.e., $S \cap e_i.K = \emptyset$), and;
- the total traveling time of path, $Path$, that is, $\sum_{e_i \in Path} e_i.w$, is minimized.

In Definition 5, the PRAO problem searches for a path, $Path$, from src to dst that satisfies good weather conditions with high probabilities, avoids passing through edges with keyword obstacles, and has the minimum traveling time.

Table 2
Frequently used notations.

Symbol	Description
G	A road network (graph)
V	A set of vertices
E	A set of edges
$e_i.w$	The traveling time of the edge e_i
$e_i.len$	The length of road segment e_i
o_l	An object or an ad-hoc obstacle
$dist(v_j, v_k)$	The Euclidean distance between v_j and v_k
$dist_N(v_j, v_k)$	The shortest path distance between v_j and v_k
$e_i.K$	A set of obstacle keywords on edge e_i
S	A set of user-specified obstacle keywords
$W(v_j)$	The weather predictions at vertex v_j
$W.type$	The type of weather condition
$W.val(v_j, t_j)$	The weather value of vertex v_j at timestamp t_j
ϵ	An ad-hoc weather threshold
$p(v_j, t_j)$	The confidence of weather condition predictions at the intersection point v_j at timestamp t_j

2.4. Challenges

The main challenges of the PRAO problem are threefold. First, in the PRAO problem, there are many edges with keywords, which may potentially be ad-hoc obstacles on road networks (i.e., containing ad-hoc obstacle keywords). It is not trivial how to efficiently compute the best paths online that do not pass through these keyword-based obstacles. Second, the weather conditions are available only at road intersection points (i.e., vertices) with probabilistic confidences, and it is not trivial to accurately estimate weather conditions at any points on roads. Moreover, the weather-based obstacles are ad-hocly specified by users at the query time, and it is rather challenging to efficiently retrieve valid paths by avoiding ad-hoc weather-based obstacles. Third, road networks are usually of large scale, and there are a large number of possible paths between *src* and *dst*, which are not time-efficient to enumerate.

Inspired by the challenges above, in this paper, we will propose an efficient query processing approach to answer PRAO queries. The proposed approach follows a pruning procedure to filter out false alarms of paths and reduce PRAO search space. We will also design an efficient query answering algorithm to answer PRAO queries over road networks.

Table 2 depicts the commonly used symbols and their descriptions in this paper.

3. Weather condition estimation

As mentioned in Section 2.2, for each edge $e_i (= v_j v_k)$, we assume that the predicted weather conditions with confidences are available at vertices v_j and v_k . However, the weather conditions at other points on edge e_i (excluding v_j and v_k) are unknown. Therefore, in this section, we will discuss how to estimate the unknown weather condition (and its confidence as well) for any point on edges in road networks, based on known weather conditions at vertices.

Without loss of generality, consider a moving vehicle on an edge $e_i (= v_j v_k)$ departing from v_j at timestamp t_{dep} toward v_k with velocity $\frac{e_i.len}{e_i.w} (= \frac{dist(v_j, v_k)}{e_i.w})$. At timestamp t_l ($\in [t_{dep}, t_{arr}]$), we will estimate weather conditions at $o_l \in e_i$, inferred from the known (predicted) weather conditions, $W.val(v_j, t_l)$ and $W.val(v_k, t_l)$, at vertices v_j and v_k , respectively.

Since the predicted weather condition $W.val(v_j, t_l)$ at vertex v_j (or $W.val(v_k, t_l)$ at vertex v_k) is associated with a probability $p(v_j, t_l)$ (or $p(v_k, t_l)$), the predicted weather condition at v_j (or v_k) can be either accurate or inaccurate. Following the *possible worlds* semantics [11], we will consider four cases below (i.e., 4 possible worlds, based on whether or not the predicted weather data at

v_j and v_k appear in the real world), and estimate the weather condition at any location $o_l \in e_i$ in each of these four cases.

Case A (The predicted weather values at both vertices v_j and v_k are accurate). In this case, based on accurate weather predictions at vertices v_j and v_k , we can estimate/interpolate the weather information at $o_l \in e_i$ by using *Inverse Distance Weighting* (IDW) interpolation [12]. The IDW interpolates weather information at unknown points o_l from weighted average of the values available at the known points v_j and v_k as follows:

$$W.val(o_l, t_l) = \frac{dist(o_l, v_k) \cdot W.val(v_j, t_l) + dist(o_l, v_j) \cdot W.val(v_k, t_l)}{dist(o_l, v_j) + dist(o_l, v_k)},$$

where o_l is the point at which we are interested in finding its weather, v_j and v_k denote the vertices with known weather data, and $dist(x, y)$ is the Euclidean distance between x and y .

Intuitively, when point o_l becomes closer to v_j (or v_k), its estimated weather value $W.val(o_l, t_l)$ will be more similar to that of v_j (or v_k).

The confidence that point o_l has the weather value $W.val(o_l, t_l)$ can be derived from the confidence that both vertices v_j and v_k have accurate weather predictions. That is, we have:

$$p(o_l, t_l) = p(v_j, t_l) \cdot p(v_k, t_l),$$

where $p(o_l, t_l)$ is the confidence of the value of weather conditions of the point o_l at timestamp t_l .

Fig. 2 illustrates an example of the interpolation of weather information at point $o_l \in e_i$, where wind speeds at v_j and v_k are 30 mph and 20 mph, respectively. Based on the distance from point o_l to v_j or v_k (i.e., 2 miles and 8 miles, resp.), we can estimate the wind speed at o_l as $W.val(o_l, t_l) = \frac{8 \times 30 + 2 \times 20}{2+8} = 28$ mph.

CASE B (The predicted weather condition at vertex $v_j \in e_i$ is inaccurate, and that at vertex $v_k \in e_i$ is accurate). In such a case, we rely on known and accurate information at vertex v_k to estimate value of weather condition at o_l on edge e_i as follows:



$$W.val(o_l, t_l) = W.val(v_k, t_l).$$

The confidence of the estimated weather value at o_l can be derived as follows:

$$p(o_l, t_l) = (1 - p(v_j, t_l)) \cdot p(v_k, t_l).$$

CASE C (The predicted weather condition at vertex $v_j \in e_i$ is accurate, and that at vertex $v_k \in e_i$ is inaccurate). We estimate values of weather conditions at o_l by using the known and accurate weather information at v_j as follows:

$$W.val(o_l, t_l) = W.val(v_j, t_l).$$

(a) weather conditions at vertices v_j and v_k (b) the interpolation of weather conditions at o_l from known weather conditions at v_j and v_k **Fig. 2.** An example of the weather interpolation (Case ③).

Similarly, we derive the confidence of the weather value at o_l as follows:

$$p(o_l, t_l) = p(v_j, t_l) \cdot (1 - p(v_k, t_l)).$$

Case ④ (The predicted weather conditions at both v_j and v_k are inaccurate). In this case, the weather at any point o_l on edge e_i is unknown (i.e., it can be any value), with probability:

$$p(o_l, t_l) = (1 - p(v_j, t_l)) \cdot (1 - p(v_k, t_l)).$$

For any point o_l on the edge e_i at timestamp t_l , we can obtain different possible weather values, $W_val(o_l, t_l)$, corresponding to the four cases (①, ②, ③, and ④) above, each of which is associated with a confidence value $p(o_l, t_l)$.

4. Pruning strategies

In this section, we present three pruning strategies, *keyword-based*, *weather-based*, and *traveling time pruning*. The pruning strategies utilize the properties of the PRAO problem to improve the efficiency of PRAO query processing.

4.1. Keyword-based pruning

In this subsection, we present an effective *keyword-based pruning* method, which can reduce the PRAO search space, under the constraint of keyword-based obstacles.

Given a set, S , of obstacle keywords that should not appear on the path from src to dst , an edge, e_i , can be safely pruned, if there exists an obstacle keyword $s \in S$ in the keyword set, $e_i.K$, of edge e_i (i.e., $e_i.K \cap S \neq \emptyset$). Intuitively, if an edge e_i contains keyword-based obstacles $s \in S$, then we can safely prune this edge.

Lemma 1 (Keyword-based Pruning). Given an edge $e_i \in E$ and a set, S , of ad-hoc obstacle keywords, if $e_i.K \cap S \neq \emptyset$ holds, then edge e_i can be safely pruned.

Proof. For details, please refer to Appendix A in our technical report [13]. \square

In order to enable keyword-based pruning, we hash the keyword set $e_i.K$ of each edge e_i (or the set, S , of obstacle keywords) into a bitmap array [14], where each bit in the bitmap represents a hashed keyword and the bit value represents the existence of the keyword. If the bit value at a position in the bitmap is equal to 1, then the keyword mapped to that position exists; otherwise (i.e., the bit is 0), the corresponding keyword does not exist.

Therefore, we can test if a keyword $s \in S$ appears on the edge e_i , by applying the AND operation between the two bitmaps of $e_i.K$ and S . If the AND result (i.e., a bitmap) is nonzero, it indicates that sets $e_i.K$ and S share common obstacle keywords (i.e., $e_i.K \cap S \neq \emptyset$), and edge e_i can be safely pruned.

4.2. Weather-based pruning

In this subsection, we propose a *weather-based pruning* method to filter out edges, e_i , that do not satisfy the constraint of weather conditions in Definition 5. The basic idea of our weather-based pruning method is to prune those edges, e_i , with bad weather (or alternatively, with low chances of having good weather).

Lemma 2 (Weather-based Pruning). Given a point o_l on edge e_i , at timestamp t_l , we denote $UB_Pr\{W_val(o_l, t_l) \leq \epsilon\}$ as an upper bound of probability $Pr\{W_val(o_l, t_l) \leq \epsilon\}$. If it holds that $UB_Pr\{W_val(o_l, t_l) \leq \epsilon\} < 1 - \alpha$, then edge e_i can be safely pruned.

Proof. For details, please refer to Appendix B in our technical report [13]. \square

Discussions on How to Compute Probability Upper Bound: In order to compute probability upper bound, $UB_Pr\{W_val(o_l, t_l) \leq \epsilon\}$, for a point $o_l \in e_i$ from weather predictions provided at v_j and v_k , we may encounter three cases.

Case 1. First, we consider the case where values of weather conditions at both v_j and v_k are less than or equal to weather threshold ϵ at timestamp t_l (i.e., $W_val(v_j, t_l) \leq \epsilon$ and $W_val(v_k, t_l) \leq \epsilon$). In this case, the probability upper bound of weather values at any object $o_l \in e_i$ can be computed as:

$$UB_Pr\{W_val(o_l, t_l) \leq \epsilon\} = 1. \quad (2)$$

The reason for the probability upper bound in Eq. (2) is as follows. In Cases (A) ~ (C), as mentioned in Section 3, the estimated weather values $W_val(o_i, t_i)$ are definitely smaller than or equal to ϵ (due to the conditions of Case 1). For Case (D), the estimated weather value may be smaller than or equal to ϵ . Therefore, the upper bound, $UB_Pr\{\cdot\}$, of the probability $Pr\{W_val(o_i, t_i) \leq \epsilon\}$ is equal to 1 (i.e., Eq. (2) holds).

Fig. 3(a) shows an example of Case 1, where the weather threshold ϵ is set to 40, and weather predictions at v_j and v_k are both less than ϵ .

Case 2. When ϵ is between weather values at two endpoints, v_j and v_k , of edge e_i , we need to examine two sub-cases.

Case 2.1. In the case where it holds that $W_val(v_j, t_i) < \epsilon < W_val(v_k, t_i)$, the probability upper bound is calculated as follows:

$$\begin{aligned} UB_Pr\{W_val(o_i, t_i) \leq \epsilon\} &= p(v_k, t_i) \cdot p(v_j, t_i) \quad // \text{ Case (A)} \\ &+ p(v_j, t_i) \cdot (1 - p(v_k, t_i)) \quad // \text{ Case (C)} \\ &+ (1 - p(v_j, t_i)) \cdot (1 - p(v_k, t_i)) \quad // \text{ Case (D)} \\ &= 1 - p(v_k, t_i) \cdot (1 - p(v_j, t_i)), \end{aligned}$$

where $p(v_j, t_i)$ (or $p(v_k, t_i)$) is the probability that the predicted weather is accurate for vertex v_j (or v_k) at timestamp t_i .

In this scenario, as discussed in Section 3, the estimated weather values in Cases (A), (C), and (D) may be smaller than or equal to ϵ . Thus, we can overestimate their probabilities, and obtain the probability upper bound $UB_Pr\{\cdot\}$ by summing up appearance probabilities of these 3 cases. Please refer to an example of Case 2.1 in Fig. 3(b).

Case 2.2. Similarly, in the case that $W_val(v_k, t_i) < \epsilon < W_val(v_j, t_i)$ holds, the probability upper bound for any point o_i on edge e_i can be given by:

$$UB_Pr\{W_val(o_i, t_i) \leq \epsilon\} = 1 - p(v_j, t_i) \cdot (1 - p(v_k, t_i)). \quad (3)$$

Case 3. In the case that both predicted weather values at v_j and v_k are greater than the weather threshold ϵ (i.e., $W_val(v_j, t_i) > \epsilon$ and $W_val(v_k, t_i) > \epsilon$), the probability upper bound for any point o_i on edge $e_i (= v_j v_k)$ at timestamp t_i can be computed as follows:

$$\begin{aligned} UB_Pr\{W_val(o_i, t_i) \leq \epsilon\} &= (1 - p(v_j, t_i)) \cdot (1 - p(v_k, t_i)). \quad // \text{ Case (D)} \end{aligned} \quad (4)$$

This is because the estimated weather values $W_val(o_i, t_i)$ only has the chance to be greater than the weather threshold ϵ in Case (D) (i.e., unknown value, as discussed in Section 3). Fig. 3(c) shows an example of Case 3, where we set ϵ to 40, and weather values at both v_j and v_k are greater than threshold ϵ .

4.3. Traveling time pruning

In this subsection, we discuss the rationale of the *traveling time pruning* method, which rules out paths, $Path$, with the traveling time greater than that of paths we have seen so far. For any path, $Path'$, we denoted $LB_T(Path')$ and $UB_T(Path')$ as lower and upper bounds of its traveling time, respectively, and *best-path-so-far* as a valid path we have seen so far from src to dst (satisfying all obstacle constraints) with the smallest traveling time. Then, our traveling time pruning method filters out paths, $Path$, whose lower bounds of the traveling times, $LB_T(Path)$, are greater than the time upper bound $UB_T(best-path-so-far)$ for path *best-path-so-far*. Intuitively, $Path$ cannot be the PRAO answer, if there exists a path, *best-path-so-far*, that has the traveling time smaller than that of $Path$.

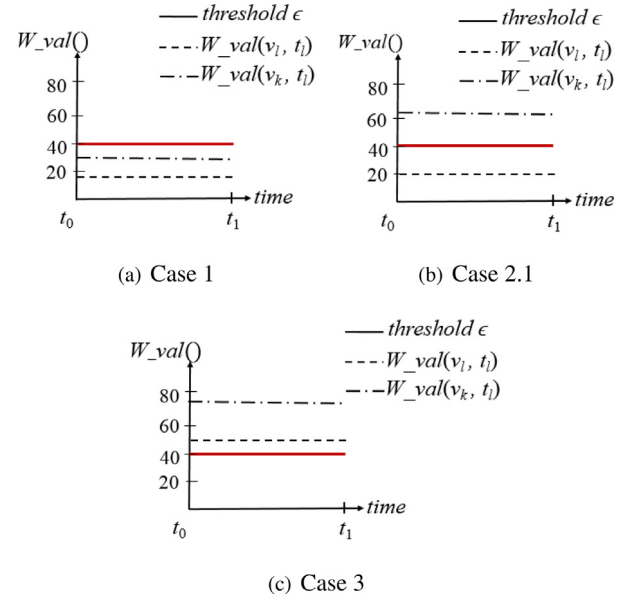


Fig. 3. Illustration of weather-based pruning with wind speed values at v_j and v_k w.r.t. threshold ϵ .

Lemma 3 (Traveling Time Pruning). Let *best-path-so-far* be the best path from src to dst that we have seen so far with the smallest traveling time. Then, any path, $Path$, can be safely pruned, if $UB_T(best-path-so-far) < LB_T(Path)$ holds.

Proof. For details, please refer to Appendix C in our technical report [13]. \square

In order to enable the traveling time pruning, lower and upper bounds of the traveling time have to be efficiently obtained. In the sequel, we will discuss how to infer lower and upper bounds of the traveling time for a path, $Path'$ (based on the traveling times on edges $e_i \in Path'$).

Discussion on How to Compute the Lower Bound of the Traveling Time $LB_T(Path)$: On road networks, the triangle inequality may not hold for the traveling time (due to variable vehicle speeds on roads), however, it holds for the shortest path distance.

As an example on road networks in Fig. 1, it holds that $dist_N(v_2, v_4) \leq dist_N(v_2, v_5) + dist_N(v_5, v_4)$, but it may not be the case for the traveling time, that is, $T(v_2, v_4) > T(v_2, v_5) + T(v_5, v_4)$, which violates the triangle inequality. This may be due to the traffic jams on the road $v_2 v_4$.

The traveling time on road networks can be obtained as the network distance divided by velocity. Thus, the lower bound, $LB_T(Path)$, of the traveling time on any path, $Path (= src \rightsquigarrow dst)$, can be given by the lower bound of the shortest path distance, $LB_dist_N(src, dst)$, divided by the maximum velocity on that path, that is,

$$LB_T(Path) = \frac{LB_dist_N(src, dst)}{\max_vel(src \rightsquigarrow dst)}, \quad (5)$$

where $\max_vel(src \rightsquigarrow dst)$ is the maximum velocity on the shortest path $src \rightsquigarrow dst$. Note that, here the velocity of each edge e_i on $Path$ can be given by $\frac{e_i.len}{e_i.w}$.

The Computation of a Lower Bound of the Shortest Path Distance $LB_dist_N(src, dst)$: We can obtain lower bounds of the shortest path distances on road networks, by utilizing the *triangle inequality* property, as described in [15]. In [15], the distance bound is computed by choosing a small number of pivots, then computing and storing shortest path distance between all vertices




Fig. 4. An example of computing the distance lower bound of the shortest path $src \rightsquigarrow dst$ on road networks.

and each of these pivots. By utilizing pivots, the lower bounds can be computed in constant time with the triangle inequality. In Fig. 4, assume that we have pre-computed $dist_N(src, v_l)$ and $dist_N(v_l, dst)$. Then, by the triangle inequality, we can obtain the lower bound of the shortest distance between src and dst as follows:

$$\begin{aligned} dist_N(src, dst) &\geq |dist_N(src, v_l) - dist_N(v_l, dst)| \\ &= LB_dist_N(src, dst). \end{aligned} \quad (6)$$

This way, we can utilize d pivots (i.e., vertices on road networks, such as v_l in the example of Fig. 4), and then offline pre-compute the shortest path distances from each vertex to these d pivots. By using these d pivots, we can apply Eq. (6) to compute a lower bound, $LB_dist_N(src, dst)$, of the shortest path distance between any two vertices src and dst . We will provide a cost model in Section 5.1 to select good pivots on road networks.

Discussion on How to Calculate the Upper Bound of the Traveling Time UB_T (best-path-so-far): In order to compute an upper bound of the traveling time, we only need to find a valid path, $Path$, from src to dst (satisfying all obstacle constraints), and treat the traveling time of this valid path as a time upper bound $UB_T(best-path-so-far)$, which can be used for the traveling time pruning.

To achieve high pruning power, the selected path $Path$ ($= src \rightsquigarrow dst$) has to meet two requirements: (1) $Path$ has to be valid (i.e., $Path$ should not contain edges with ad-hoc keyword-based and weather-based obstacles), and (2) the traveling time on $Path$ should be close to the traveling time of the optimal path.

Due to the existence of ad-hoc obstacles, it is challenging to find a valid path from src to dst . Prior works [16] usually used pivots and the triangle inequality to compute time upper bounds, however, they did not consider ad-hoc given obstacles on roads. In contrast, some roads in our PRAO problem may not be available online due to ad-hoc obstacles, which may be passed by some offline pre-computed paths to pivots. Thus, we cannot utilize offline pre-computed distances via pivots and velocity bounds to calculate the traveling time upper bounds. Another challenge is that, due to many possible valid paths that connect src and dst , it is not trivial how to efficiently choose a valid path with the traveling time close to that of the optimal path.

Inspired by the challenges above, in this paper, we will follow a greedy traversal algorithm to find a valid path, $Path$, that connects src and dst . Then, the traveling time upper bound of the best-so-far path, $best-path-so-far$, from src to dst is equal to the minimum traveling time of $Path$ that we have encountered, that is,

$$UB_T(best-path-so-far) = \min_{\forall \text{ valid path } Path} T(Path), \quad (7)$$

where $T(Path)$ is the traveling time of a valid path $Path$.




Fig. 5. An example of computing the distance lower bound, with the help of 2 pivots, piv_1 and piv_2 .

5. Offline pre-computations and indexing

In this section, we will illustrate offline pre-computation and indexing mechanisms over weather and keyword information on road networks, which can enable our proposed pruning methods and facilitate online PRAO query processing.

5.1. Pivot selection

Next, we discuss our pivot selection algorithm in details including our proposed cost model and the difference with the landmark selection algorithm in [15].

Cost Model for Pivot Selection. As discussed in Section 4.3, we can use the triangle inequality to estimate the lower bound of network distances between two vertices v_j and v_k via pivots. As an example, Fig. 5 shows two vertices, v_j and v_k , and two pivots, piv_1 and piv_2 , on road networks. We can offline pre-compute the shortest path distances from vertices v_j and v_k to these two pivots, that is, $dist_N(v_j, v_k) \geq |dist_N(v_j, piv_t) - dist_N(v_k, piv_t)|$, where $t = 1, 2$.

Intuitively, those pivots that result in large distance lower bound will give better pruning power. Thus, in Fig. 5, piv_2 gives a tighter distance lower bound (i.e., $13 = |7 - 20|$) than that of piv_1 (i.e., $3 = |13 - 10|$).

From the discussion above, we will design a cost model below to evaluate the “goodness” of the selected pivots, based on the distance differences between pivots and pairs of vertices on road networks.

$$C = \sum_{\forall (v_i, v_j) \in V \times V} \max_{\forall piv_t \in S_{piv}} |dist_N(v_i, piv_t) - dist_N(v_j, piv_t)|. \quad (8)$$

where S_{piv} is a set of d selected pivots.

Eq. (8) computes the tightness, C , of distance lower bounds (via d pivots in a pivot set S_{piv}), for all pairs of vertices (v_i, v_j) . Intuitively, larger C value will lead to higher pruning power. Therefore, our pivot selection algorithm aims to maximize the C value (i.e., based on our cost model) during the process of choosing pivots.

Pivot Selection Algorithm. Algorithm 1, namely Obtain_Pivots, chooses a set, S_{piv} , of d pivots, in light of our proposed cost model (i.e., maximizing C in Eq. (8)). We first initialize two parameters, $global_cost$ and S_{piv} , which store the global C value and pivot set, respectively (line 1). Initially, our algorithm selects d random pivots, forming a pivot set \mathcal{P} , and then evaluates the C value (recorded by $local_cost$) with respect to \mathcal{P} (lines 3–5). Next, for $swap_iter$ times, we will randomly pick a pivot piv_i from \mathcal{P} and a non-pivot $v_j \in V - \mathcal{P}$, and swap their roles so that a new pivot set \mathcal{P}^{new} is obtained (lines 6–8). We evaluate the new cost, C^{new} , over the updated pivot set \mathcal{P}^{new} (line 9), and compare it with

Algorithm 1: Obtain_Pivots

Input: road network $G = (V, E)$ and the number of pivots d
Output: the set of pivots S_{piv}

```

1 set  $global\_cost = -\infty$ ,  $S_{piv} = \emptyset$ 
2 for  $a = 1$  to  $global\_iter$  do
3   randomly select  $d$  pivots from  $V$  and form a pivot set  $\mathcal{P}$ 
4   evaluate the cost function  $\mathcal{C}$  of  $\mathcal{P}$ 
5   set  $local\_cost = \mathcal{C}$ 
6   for  $b = 1$  to  $swap\_iter$  do
7     select a random pivot  $piv_i \in \mathcal{P}$ 
8     randomly choose a non-pivot  $v_j \in (V - \mathcal{P})$ 
9      $\mathcal{P}^{new} = \mathcal{P} - \{piv_i\} + \{v_j\}$ 
10    evaluate the cost function  $\mathcal{C}^{new}$  w.r.t.  $\mathcal{P}^{new}$ 
11    if  $\mathcal{C}^{new} < local\_cost$  then
12       $local\_cost = \mathcal{C}^{new}$ 
13       $\mathcal{P} = \mathcal{P}^{new}$ 
14  if  $local\_cost < global\_cost$  then
15     $S_{piv} = \mathcal{P}$ 
16     $global\_cost = local\_cost$ 
17 return  $S_{piv}$ 

```

$local_cost$ (lines 10). If it holds that $\mathcal{C}^{new} < local_cost$, then we will accept new pivot set \mathcal{P}^{new} and update $local_cost$ (lines 11–12). The process of finding better pivots repeats for $swap_iter$ times (lines 6–12).

In order to avoid the locally optimal solution, we will execute the pseudo code above for $global_iter$ times, by selecting different initial pivot sets (line 2). For each round, we will update the global pivot set S_{piv} , if it holds that $local_cost < global_cost$ (lines 13–15). Finally, the algorithm returns the selected pivot set S_{piv} (line 16).

Discussions on the Number, d , of Pivots. In order to decide an appropriate number of pivots, we will execute our pivot selection algorithm with different numbers of pivots, where $d = 1, 2, \dots$. Intuitively, with more and more pivots, the increase of the pruning power (or the \mathcal{C} value) will become smaller. We will set parameter d , such that the difference of \mathcal{C} values with d and $(d+1)$ pivots is below a user-specified threshold.

Discussions on the Differences Between Pivot Selection and Landmark Selection. Note that, pivots used in our work are similar to the landmarks in [15]. Nevertheless, the criterion of selecting pivots in our PRAO algorithm is based on our proposed cost model (i.e., Eq. (8)), which is different from the heuristic-based landmark selection algorithm. Specifically, landmarks in [15] are selected in order to provide tight distance bounds in graphs. The landmark selection algorithms are based on different heuristics such as greedy landmark selection, using landmarks geometrically lying behind the destination, and so on. With optimization techniques, [15] removes a landmark and replaces it with the best candidate landmark (with the highest score w.r.t. distance lower bounds). In contrast, our work designs a specific cost model (as given in Eq. (8)), which has a different goal (i.e., score) from the one for the landmark selection. Furthermore, our pivot selection algorithm is in light of our proposed cost model, and aims to obtain the best pivot set with the highest \mathcal{C} value (in Eq. (8)) by avoiding local optimality. Thus, our cost-model-based pivot selection algorithm is different from heuristic-based landmark selection algorithm.

5.2. Pre-computations

In this subsection, we discuss how to choose pivots and the data pre-computation.

Pre-Computation of the Shortest Path Distances with Pivots: Once we select pivots, we can pre-compute distances between

all vertices $v_j \in V$ and pivots piv_i (for $1 \leq i \leq d$).

$$dist_N(v_j, piv_i) = \sum_{\forall e_i \in Path(v_j \rightsquigarrow piv_i)} e_i.len. \quad (9)$$

The pre-computed distances with pivots are used to compute lower bounds of the traveling time, as discussed in Section 4.3.

Pre-Computation of the Probability Upper Bounds of Weather Conditions: We pre-compute probability upper bounds of weather conditions, as well as weather values at endpoints of edges, as mentioned in Section 4. The pre-computed data will be indexed in a tree index for facilitating online PRAO query processing. Specifically, given a user-specified weather threshold ϵ and a probabilistic threshold α , we will utilize the pre-computed data to online compute the probability upper bound, based on Cases 1sim3 discussed in Section 4.2, which can be used to enable the weather-based pruning.

5.3. Indexing

In this subsection, we discuss the indexing mechanism over road networks, which can facilitate efficient PRAO query processing. In particular, we use an R^* -tree [17], denoted as \mathcal{I} , to index road network G and its pre-computed data. To construct the R^* -tree, we first partition road networks (graphs) G into subgraphs (each with edges spatially close to each other), and treat each subgraph as a leaf node, N , of the R^* -tree, which contains edges e_i in the subgraph, represented by *minimum bounding rectangles* (MBRs). Then, we recursively group MBRs of leaf or non-leaf nodes until we obtain a root of the R^* -tree.

Leaf Nodes: Each edge $e_i (= v_j v_k)$ in the leaf node N of index \mathcal{I} is represented by an MBR, which minimally bounds all spatial locations in edge e_i . Moreover, edge e_i is associated with two vectors of the pre-computed shortest path distances from the two endpoints of the edge e_i to pivots, that is, $(dist_N(v_j, piv_1), dist_N(v_j, piv_2), \dots, dist_N(v_j, piv_d))$ and $(dist_N(v_k, piv_1), dist_N(v_k, piv_2), \dots, dist_N(v_k, piv_d))$, the traveling time, $e_i.w$, on edge e_i , and a bitmap, $e_i.K$, that stores keywords associated with e_i .

We also store the functions of weather predictions at endpoints v_j and v_k of edge e_i , that is, functions $W_val(v_j, t)$ and $W_val(v_k, t)$ at time t , respectively. In addition, for any edge $e_i (= v_j v_k)$, we dynamically keep lower/upper bounds of weather predictions at edges e_i , denoted as $lb_W_val(e_i, t)$ and $ub_W_val(e_i, t)$, respectively, at timestamp t in a future period of time. That is, we have:

$$lb_W_val(e_i, t) = \min\{W_val(v_j, t), W_val(v_k, t)\}, \quad (10)$$

$$ub_W_val(e_i, t) = \max\{W_val(v_j, t), W_val(v_k, t)\}. \quad (11)$$

Non-Leaf Nodes: Each non-leaf node, N_c , of index \mathcal{I} contains multiple entries, denoted as N_c , each of which is represented by an MBR and minimally bounds edges in its subtree. Moreover, for each entry N_c , we also store minimum/maximum traveling times for any edge e_i under N_c .

$$lb_t(N_c) = \min_{\forall e_i \in N_c} e_i.w, \quad (12)$$

$$ub_t(N_c) = \max_{\forall e_i \in N_c} e_i.w. \quad (13)$$

Furthermore, each entry $N_c \in \mathcal{I}$ is associated with a summary, $N_c.K$, of keywords shared by all edges under entry N_c , that is,

$$N_c.K = \bigcap_{\forall e_i \in N_c} e_i.K. \quad (14)$$

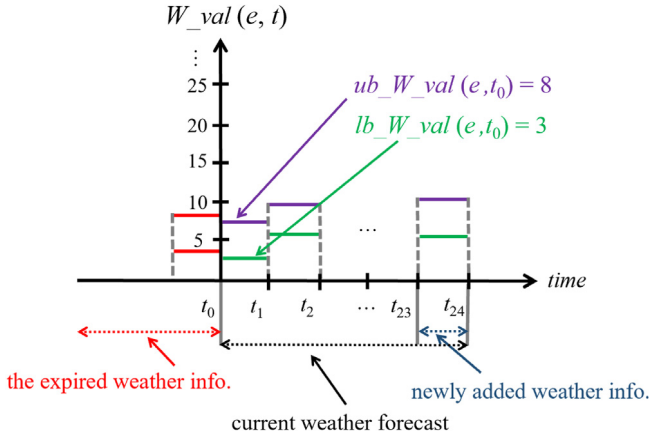


Fig. 6. Dynamic update of weather forecast information for an edge on road networks (or an index entry).

We also dynamically maintain the upper/lower bounds of weather values over time for edges e_i under node N_c , that is,

$$lb_W_val(N_c, t) = \min_{e_i \in N_c} \{lb_W_val(e_i, t)\},$$

$$ub_W_val(N_c, t) = \max_{e_i \in N_c} \{ub_W_val(e_i, t)\};$$

and probability upper bound of weather conditions for the 4 possible cases (as mentioned in Section 3):

$$UB_Pr\{N_c\} = \max_{e_i \in N_c \wedge o_i \in e_i} \{UB_Pr\{W_val(o_i, t_i)\}\}. \quad (15)$$

Dynamic Maintenance of Weather Data: In order to support efficient updates of weather data, for each entry of leaf/non-leaf nodes in the R*-tree, we store a pointer pointing to a (space-efficient) circular array that contains a sliding window of weather prediction data in a future period of time (e.g., in the next 24 h). Circular arrays are organized by a tree structure, the same as the structure of R*-tree, which can be used for weather data maintenance.

As shown in Fig. 6, we maintain weather forecast information (e.g., weather lower/upper bounds and probability upper bounds) on edges or entries in the index in a streaming manner. When new weather prediction data on edges arrive, we will replace the expired weather data on edges with the new ones in circular arrays, and aggregate/update the weather information for circular arrays of index entries in a bottom-up manner.

Note that, the R*-tree index only needs to be offline constructed once, whereas the weather forecasting data are periodically maintained (e.g., every hour). We will show later in Section 7.2 that the time cost to dynamically update the weather data is low (e.g., about 0.01 s, as will be confirmed in Fig. 9).

Auxiliary Synopses: For the space efficiency, we use bitmap to store a summary of keywords in leaf or non-leaf nodes. On the leaf level, all possible keywords in $e_i.K$ are hashed into a bitmap. If some keyword on an edge e_i is hashed to a bit in the bitmap, then this bit is set to 1; otherwise, the bit is set to 0. For leaf node N , we use the AND operator to aggregate all bitmaps of edges under N and obtain a bitmap for leaf node N . Similarly, for non-leaf node N_p , we can also compute a bitmap by applying the AND operator over bitmaps of its children N_c .

To facilitate the path expansion during the index traversal, on each level of the R*-tree, we will also maintain a *connection graph*, where each vertex of this graph corresponds to a node entry, N_i , on this level, and each edge, $N_i N_j$, of the connection graph indicates that there are two road segments under two entries, N_i

and N_j , respectively, sharing the same endpoint (i.e., connected via an intersection point).

In addition, in order to quickly look up index nodes containing a specific edge (e.g., the one containing source src) in road networks, we encode edges, as well as index nodes containing these edges, into a hash file, *Hash*.

5.4. Index node pruning

In this subsection, we discuss the rational of index node pruning and illustrate the pruning of candidate paths that involve tree nodes N_p in the index \mathcal{I} using pruning methods discussed in Section 4. In real-world scenarios, road-network obstacles (e.g., high temperature, high wind speed, etc.) may have an impact on an entire region. In such a case, all roads in this region that violate the weather constraints can be safely pruned. In other words, paths passing by these roads (with ad-hoc weather-based obstacles) can be safely pruned. Similarly, in the R*-tree index, each tree node N can be considered as a region containing multiple spatially close roads, associated with a summary of weather and keyword information in this region (as discussed in Section 5.3). Node N can be directly pruned, if all edges under this node can be safely pruned (w.r.t. weather or keywords). This way, we can effectively prune an index node representing a spatial region that does not satisfy the constraints of weather conditions or keywords.

Keyword-based Pruning: When we access a node $N_p \in \mathcal{I}$, we can check whether or not its child entries N_c contain keyword-based obstacles. Specifically, each child entry N_c is associated with a bitmap, $N_c.K$, which contains keyword information under node N_c . To check if node N_c contains any obstacle keyword $s \in S$, we only need to verify if the bit corresponding to s in bitmap $N_c.K$ is "1". If the answer is yes, it indicates that all edges under node N_c contain the obstacle keyword s . Thus, none of edges under node N_c can appear in the path of our PRAO query answer (due to the existence of the obstacle keyword), and entry N_c can be safely pruned.

Weather-based Pruning: Each node entry $N_c \in \mathcal{I}$ stores an probability upper bound of the weather condition $UB_Pr\{N_c\}$ (as given in Eq. (15)). Intuitively, if this probability upper bound is less than $1 - \alpha$ (i.e., $UB_Pr\{N_c\} < 1 - \alpha$), then node N_c can be safely pruned, where α is the probability threshold. This is because any object o_i on edges e_i under node N_c satisfies the following inequality:

$$Pr\{W_val(o_i, t_i) < \epsilon\} \leq UB_Pr\{N_c\}.$$

Thus, any edge under N_c cannot appear in our PRAO query answer, and node N_c can be safely pruned.

Traveling Time Pruning: We compute lower/upper bounds of traveling time for any edge under a node $N_c \in \mathcal{I}$. The bounds of traveling time of nodes can be used to compute lower/upper bounds of traveling time of paths, $Path = N_{src} \rightarrow N_1 \rightarrow \dots \rightarrow N_{dst}$ in the index \mathcal{I} .

Specifically, for any node N_c in path $Path$, we can obtain the time lower/upper bounds, $LB_T(N_c)$ and $UB_T(N_c)$, of an edge under N_c , given by Eqs. (12) and (13), respectively.

The only issue we need to address is the time lower/upper bounds for the first or last edge in a node N_{src} of $Path$ (i.e., the one containing src or dst). Without loss of generality, consider the first road segment e_{src} in node N_{src} containing src on $Path$. Then, $LB_T(N_{src}) = 0$ and $UB_T(N_{src}) = ub_t(N_{src})$. The case of node N_{dst} is similar and thus omitted.

Based on Eqs. (12) and (13), we can obtain the time lower bound, $LB_T(Path)$, of a path $Path$ (even if this path does not reach the destination):

$$LB_T(Path) = \sum_{N_c \in Path} lb_t(N_c). \quad (16)$$

Furthermore, when a path $Path$ is valid (i.e., satisfying keyword and weather constraints) and complete (i.e., reaching dst), we can obtain its time upper bound, $UB_T(Path)$, as follows:

$$UB_T(Path) = \sum_{\forall N_c \in Path} ub_t(N_c). \quad (17)$$

Algorithm 2: PRAO-QP

Input: a road network $G = (V, E)$, an index \mathcal{I} over G , a source src , a destination dst , a set, S , of obstacle keywords, a weather threshold ϵ , and a probabilistic threshold α

Output: a path, $Path$, that satisfies Definition 5

```

1 initialize an empty queue,  $\mathcal{Q}$ , with entries in the form  $(Path, level)$ 
2 initialize an empty min-heap,  $\mathcal{H}$ , accepting  $(key, Path)$  pairs
3  $\lambda \leftarrow +\infty$ ;  $S_{cand} \leftarrow \emptyset$ ;  $level \leftarrow height(\mathcal{I})$ 
4 obtain the entry  $N_{src} \in root(\mathcal{I})$  that contains  $src$ 
5 for each path,  $Path = N_{src} \rightarrow N_i (\in root(\mathcal{I}))$  do
6   apply keyword-based and weather-based pruning
7   if  $Path$  cannot be pruned then
8      $\mathcal{Q} \leftarrow (Path, level)$ 
9 for each path,  $(Path, level)$ , in  $\mathcal{Q}$  do
10  if  $level = -1$  then
11    terminate the loop
12   $(Path, level) = pop\_out(\mathcal{Q})$ ;
13  substitute node entries in  $Path$  with their children, and obtain paths  $Path_c$ 
14  for each  $Path_c$  do
15    expand  $Path_c$  by increasing the length by 1, and obtain  $Path'_c$ 
16    apply keyword-based and weather-based pruning
17    if  $Path'_c$  cannot be pruned then
18       $\mathcal{Q} \leftarrow (Path'_c, level - 1)$ 
19 if  $\mathcal{Q}$  is empty then
20   return  $\emptyset$ 
21 for each  $Path \in \mathcal{Q}$  do
22   if  $Path$  reaches  $dst$  then
23     if all edges,  $e_i \in Path$ , satisfy the weather condition, and  $Path$  cannot
       be pruned by traveling time pruning then
24        $S_{cand} \leftarrow S_{cand} \cup \{Path\}$ 
25       if  $\lambda > UB\_T(Path)$  then
26          $\lambda \leftarrow UB\_T(Path)$ 
27   else
28     insert the entry  $(LB\_T(Path), Path)$ , into  $\mathcal{H}$ 
29 while  $\mathcal{H}$  is not empty do
30    $(key, Path) = pop\_out(\mathcal{H})$ 
31   if  $key > \lambda$  then
32     terminate the loop
33   expand the path,  $Path$ , by increasing the length by 1, and obtain path,  $Path'$ 
34   if  $Path'$  cannot be pruned by keyword-based, weather-based, and traveling
     time pruning then
35     if  $Path'$  reaches  $dst$  then
36        $S_{cand} \leftarrow S_{cand} \cup \{Path'\}$ 
37       if  $\lambda > T(Path')$  then
38          $\lambda \leftarrow T(Path')$ 
39     else
40       insert the entry,  $(T(Path'), Path')$ , into  $\mathcal{H}$ 
41 refine candidate paths in  $S_{cand}$  and return the actual PRAO answer

```

6. Processing of PRAO queries

6.1. The PRAO query procedure

In this subsection, we will propose an efficient query processing approach to answer PRAO queries. Specifically, Algorithm 2 illustrates the pseudo code of our PRAO query answering algorithm, namely PRAO-QP, which follows the filter-and-refine framework. That is, PRAO-QP first obtains candidate paths of PRAO answers by traversing the tree index \mathcal{I} and applying our three pruning methods (keyword-based, weather-based, and traveling time pruning), stated in Section 4. Then, we refine the

retrieved candidate paths (via the index traversal), and return the actual PRAO answers.

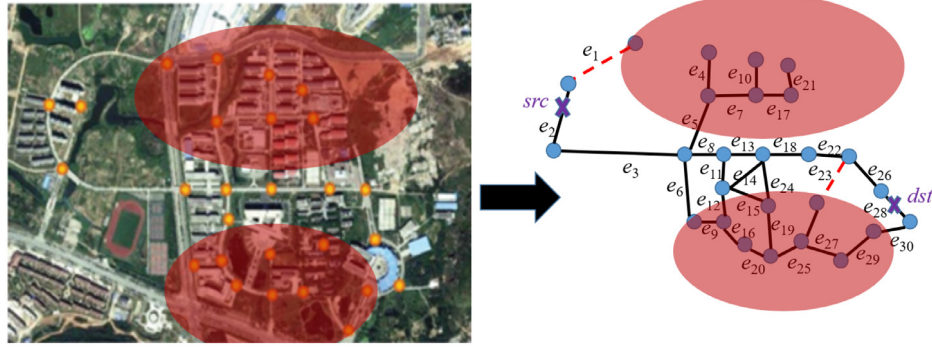
Initialization: Specifically, in Algorithm 2, we first initialize an empty queue \mathcal{Q} for traversing index \mathcal{I} , which accepts entries in the form of $(Path, level)$, where $Path$ is a path containing index nodes, and $level$ is the level of nodes in path $Path$ (line 1). We also keep a minimum heap, \mathcal{H} , with entries in the form $(key, Path, t)$ for path expansion after the index traversal, where key stores the lower bound of the traveling time on a path $Path$ and t is the timestamp that is the time spend on $Path$ added to the query time (line 2). In addition, we also set a traveling time threshold λ to $+\infty$, maintain an empty candidate set S_{cand} , and let initial traversing level, $level$, be the height, $height(\mathcal{I})$, of the R^* -tree (line 3).

Index Traversal: Next, we will obtain an entry, N_{src} in root, $root(\mathcal{I})$, that contains the source src , and obtain all paths, $Path = N_{src} \rightarrow N_i$, of length 2 (note: each node in $Path$ represents one edge), which start from an edge in entry N_{src} and pass the keyword-based and weather-based pruning, where $N_i \in root(\mathcal{I})$. We will then add these paths in the form $(Path, level)$ to queue \mathcal{Q} (lines 4–8).


In the sequel, we will use queue \mathcal{Q} to traverse the R^* -tree in a breadth-first manner and retrieve candidate paths (lines 9–28). In particular, if the candidate path, $Path$, from queue \mathcal{Q} contains entries on the edge level (i.e., $level = -1$), then we will terminate the loop (lines 9–11); otherwise, we will pop out a candidate path $(Path, level)$ from queue \mathcal{Q} , and substitute node entries in $Path$ with their children, which leads to paths, $Path_c$, on lower level (lines 12–13). Furthermore, for these paths $Path_c$, we will try to expand the path $Path_c$ (via connection graphs mentioned in Section 5.3) by including one more (connected) node entry on level $(level - 1)$, and obtain paths $Path'_c$ (lines 14–16). If the expanded path $Path'_c$ cannot be ruled out by keyword-based and weather-based pruning, then we will insert it into queue \mathcal{Q} for further checking (lines 17–18). The loop terminates when queue \mathcal{Q} is empty (line 9) or the index traversal reaches the edge level (lines 10–11).

When queue \mathcal{Q} is empty, it indicates that we cannot find candidate paths, and thus return an empty query answer set \emptyset (lines 19–20). Otherwise, we will prepare for expanding the length of candidate paths in \mathcal{Q} (lines 21–28). If paths $Path$ in \mathcal{Q} have reached destination dst , then we will check whether or not they satisfy the weather and traveling time pruning (lines 22–23). Those objects that cannot be pruned will be added to the candidate set S_{cand} , and threshold λ will be updated with the smallest upper bound of the traveling time for candidate paths we have seen so far (discussed in Section 4.3; lines 24–26). If the paths $Path \in \mathcal{Q}$ have not reached dst , then we will insert them into heap \mathcal{H} in the form $(LB_T(Path), Path)$, where $LB_T(Path)$ is the time lower bound for (partial) path $Path$ (as the path expansion is not completed; lines 27–28).

Path Expansion: Finally, we will use the heap \mathcal{H} to further expand (partial) paths in it, and obtain complete paths from src to dst (lines 29–40). Specifically, each time we obtain an entry $(key, Path)$ from heap \mathcal{H} with the minimum key (i.e., time lower bound; lines 29–30). Then, we try to expand the path $Path$ by increasing its length by 1, and obtain a path $Path'$ (lines 33). If path $Path'$ can pass the weather-based, keyword-based, and traveling time pruning and reaches destination dst , we will add it to the candidate set S_{cand} and update threshold λ (lines 34–38); otherwise (i.e., $Path'$ cannot be pruned and does not reach dst), we insert $(T(Path'), Path')$ into heap \mathcal{H} for further expansion (lines 39–40). The path expansion loop repeats until heap \mathcal{H} is empty (line 29) or the popped entry has key value above threshold λ



(a) Illustration of road network with its corresponding graph representation.



(b) The R*-tree representation of road network in (a).

Fig. 7. Illustration of road network and its corresponding graph representation and R*-tree index.

(i.e., all the remaining paths in \mathcal{H} cannot have smaller traveling time; lines 31–32).

Candidate Path Refinement: After the index traversal and path expansion, we can obtain a number of candidate paths in set S_{cand} . Then, we can refine these paths by checking constraints of keywords, weather, and traveling times, and return actual PRAO answers (line 41).

Next, we give a running example for the PRAO-QP algorithm in Algorithm 2. Fig. 7(a) shows an example of road networks and its corresponding graph representation, where the areas within red ovals correspond to those regions with bad weather conditions and the dashed (red) edges indicate roads with ad-hoc keyword obstacles. Moreover, Fig. 7(b) shows the corresponding R*-tree index of road networks in Fig. 7(a). Below, we give the detailed steps for this running example.

- The PRAO-QP procedure starts by obtaining an entry N_{src} that contains the src edge, N_1 ,
- obtain all paths, $Paths$, of length 2 that starts from N_1 , that is, $\{\{N_1N_1\}, \{N_1N_2\}\}$
- apply weather-based and keyword-based pruning on $Paths$
- insert remaining paths into the queue \mathcal{Q} in the form of $(Path, level)$, $\mathcal{Q}(\{(\{N_1N_1\}, 1), (\{N_1N_2\}, 1)\})$
- for all paths, $Path \in \mathcal{Q}$
 - extract $Path$ for \mathcal{Q}
 - substitute node entries in $Path$ with their children and expand the length of the path by 1,
 - resulting path is: $\{N_3N_4N_4\}$
 - apply weather-based and keyword-based pruning on the node level
 - $Path \{N_3N_4N_4\}$ will be directly pruned due to the existence of an unsafe node N_4
 - **note that no path exists, to continue, we extract another path from the queue \mathcal{Q}
 - extract $Path$ from \mathcal{Q} , $Path = \{N_1N_2\}$

- substitute node entries in $Path$ with their children and expand the length of the path by 1,
 - resulting paths are: $\{\{N_3N_4N_5\}, \{N_3N_4N_6\}, \{N_3N_4N_7\}, \{N_3N_5N_6\}, \{N_3N_5N_7\}, \{N_3N_6N_7\}\}$
 - apply weather-based and keyword-based pruning on the node level
 - $Paths \{N_3N_4N_5\}, \{N_3N_4N_6\}, \{N_3N_4N_7\}, \{N_3N_5N_6\}, \{N_3N_5N_7\}$ will be directly pruned due to the existence of unsafe nodes N_4 and N_5
 - insert the remaining path into the queue $\mathcal{Q}(\{N_3N_6N_7\}, 0)$
 - extract paths from the queue, $Path = \{N_3N_6N_7\}$
 - substitute node entries in $Path$ with their children and expand the length of the path by 1
 - compute upper bound of traveling time $UB_T(Path)$
 - resulting paths are: $\{e_2, e_3, e_{13}, e_{26}\}, UB_T(Path)$
 - insert the resulting paths into the heap \mathcal{H}
 - for all $Paths$ in the heap \mathcal{H}
 - further expand partial paths by increase length by 1 and obtain $Path'$, apply weather-based, keyword-based, and traveling time pruning
 - insert pair $(Path', T(Path'))$ in \mathcal{H}
 - continue until either the heap is empty, or a valid path from src to dst is formed with traveling time greater than the upper bound of traveling time of all paths in the heap \mathcal{H}

6.2. Complexity analysis

In this subsection, we discuss the time complexity of the PRAO-QP algorithm in Algorithm 2. The major time cost of Algorithm 2 consists of two portions: index traversal (lines 1–28) and candidate path expansion (lines 29–41).

Without loss of generality, assume that the *connection graph* (mentioned in Section 5.3) for the i th level of the R*-tree index

has the average degree deg_i . Then, during the index traversal, on the i th level, the number of candidate paths $Path$ (starting from N_{src} that contains src) can be given by $\sum_{j=1}^{h-i+1} (deg_i + 1)^j (= \frac{(deg_i + 1)^{h-i+1} - 1}{deg_i})$, where h is the height $height(\mathcal{I})$ of the R^* -tree. Moreover, in the worst case, we need to expand these candidate paths by including one more node. Thus, the traversal cost on the i th level is given by $\frac{(deg_i + 1)^{h-i+1} - 1}{deg_i} \cdot (deg_i + 1)$, and the total index traversal cost, C_{index} , is given by $\sum_{i=0}^h (\frac{(deg_i + 1)^{h-i+1} - 1}{deg_i} \cdot (deg_i + 1))$.

Furthermore, for candidate path expansion, there are totally about $(deg_0 + 1)^{h+1}$ candidate paths on the leaf level after the index traversal. Assume that the average degree of road network G is denoted by deg_G , and the length of valid path from src to dst is represented by σ . Then, the total cost of path expansions, $C_{expansion}$, can be given by $(deg_0 + 1)^{h+1} \cdot deg_G^{\sigma-h-1}$ in the worst case.

To summarize, in the worst case (i.e., candidate paths cannot be pruned by any pruning strategies), the worst-case time complexity of Algorithm 2 is given by $C_{index} + C_{expansion} = \sum_{i=0}^h (\frac{(deg_i + 1)^{h-i+1} - 1}{deg_i} \cdot (deg_i + 1)) + (deg_0 + 1)^{h+1} \cdot deg_G^{\sigma-h-1}$, or $O(\sum_{i=0}^h (deg_i + 1)^{h-i+1} + deg_0^{h+1} \cdot deg_G^{\sigma-h-1})$. In practice, degrees deg_i (or deg_G) of connection graphs (or road networks), and height, h , of the R^* -tree are usually small constants (e.g., 3-4, or smaller), and our pruning methods (via keyword, weather, and time constraints) are effective to prune candidate paths. Therefore, as confirmed by empirical studies in Section 7, the PRAO-QP algorithm can achieve good performance in the amortized case, in terms of the CPU time and I/O cost.

7. Experimental evaluation

7.1. Experimental settings

Real/synthetic Data Sets: We test our PRAO query answering approaches on both real and synthetic data sets. Specifically, for real graph data, we use the California Road Network [18], denoted as CA, which contains 21,048 road intersection points and 21,693 road segments. CA is originally obtained from Digital Chart of the World Server and U.S. Geological Survey. Each vertex in CA data set is represented by (longitude, latitude), and we obtain real weather forecast of each vertex from Dark Sky [10]. Moreover, we produce random integers within interval [0, 15] to simulate keywords associated with roads (edges), which may correspond to highway, city, uneven road, and so on. We also obtain the average traveling times on roads that are proportional to lengths of roads.

For synthetic data, we generate road networks (i.e., graphs) as follows. Specifically, we first generate random vertices in the 2D data space following either Uniform or Gaussian distribution. Then, we randomly connect vertices nearby through edges, such that all vertices are reachable in one single connected graph and the average degree of vertices is within [3, 4]. This way, with Uniform and Gaussian distributions of vertices, we can obtain two types of graphs, denoted as *Uni* and *Gau*, respectively. For each vertex, we also generate random weather data, for example, wind speed within interval [0, 100] and its forecasting accuracy within (0, 1]. Then, we also associate each edge e_i with keywords in $e_i.K$, and the traveling time $e_i.w$.

We build indexes (including a variant of R^* -tree [17] and auxiliary data structures) over each of real/synthetic road networks above (as mentioned in Section 5.3), where the page size is set to 4 K. For the pivot selection, we set the number of pivots to 5, based on the cost model as discussed in Section 5.1.

To evaluate the PRAO query performance, we generate 20 queries with source src and destination dst , where the average

Table 3
Parameter settings.

Parameter	Value
α	0.1, 0.3, 0.5 , 0.7, 0.9
ϵ	20, 30, 50 , 70, 90
σ	5, 8, 10 , 15, 20, 25, 30
$ S $	1, 3, 5 , 8, 10
d	2, 5 , 8, 10, 15
$ V $	10K, 20K, 30K , 40K, 50K, 200K, 500K, 1M, 1.2M

length, σ , of valid paths between src and dst is set to 5, 8, 10 (default value), 15, 20, 25, or 30.

Measures: We will report the CPU time and the I/O cost of our PRAO query processing approach, where the CPU time is the average time cost of conducting PRAO queries, and the I/O cost is defined by the number of node accesses during the index traversal.

Competitor: To the best of our knowledge, there are no prior works that answer path routing queries by considering both ad-hoc keyword-based and weather-based obstacles. In this paper, we compare our PRAO algorithm with two baseline algorithms, denoted as A^* and *filterfirst*. Specifically, the first baseline algorithm, A^* , is a variant of the A^* algorithm [1]. Note that, we cannot directly apply the A^* algorithm, due to the existence of the ad-hoc keyword-based and weather-based obstacles. Thus, we adapt the A^* algorithm to handle weather-based and keyword-based obstacles as follows. The A^* algorithm traverses the graph in a breadth-first manner. Whenever encountering an edge, we check the edge to see if it contains weather-based and/or keyword-based obstacles. In case the edge violates the weather/keyword constraints, it will be pruned; otherwise, it will be inserted into the queue for further checking.

The second baseline algorithm, *filterfirst*, starts by filtering out all edges in road networks that contain ad-hoc over the remaining road networks (excluding those with obstacles) weather-based and/or keyword-based obstacles. After that, *filterfirst* applies the conventional A^* algorithm [1] over the remaining road networks (excluding those with obstacles) to find the shortest path between the source, src , and the destination, dst .

Parameter Settings: Table 3 depicts the parameter settings in our experiments, where bold numbers are default parameter values. In each set of our subsequent experiments, we will vary one parameter while setting other parameters to their default values. We ran our experiments on a machine with Intel(R) Core(TM) i7-6700 CPU 3.40 GHz (8 CPUs) and 64 GB memory. All algorithms were implemented by C++.

7.2. The PRAO performance

The PRAO Performance vs. Real/Synthetic Data Sets: First, as illustrated in Fig. 8, we test the performance of our PRAO approach against two baseline algorithms namely A^* and *filterfirst* with real/synthetic data sets CA, *Uni*, and *Gau*. From figures, we can see that PRAO outperforms both A^* and *filterfirst* for all the three data sets, in terms of both CPU time and I/O cost. This is because PRAO applies effective pruning methods with the help of the index. The experimental results confirm the effectiveness of our pruning methods, and efficiency of our PRAO approach.

In Fig. 8, we can see that PRAO outperforms both baselines, A^* and *filterfirst*. Specifically, our PRAO algorithm utilizes a variant of the R^* -tree index to facilitate the query processing. On the node level of the index, we can quickly filter out many false alarms of path candidates (which violate the constraints of keyword or weather conditions) with low costs, and obtain a small set of

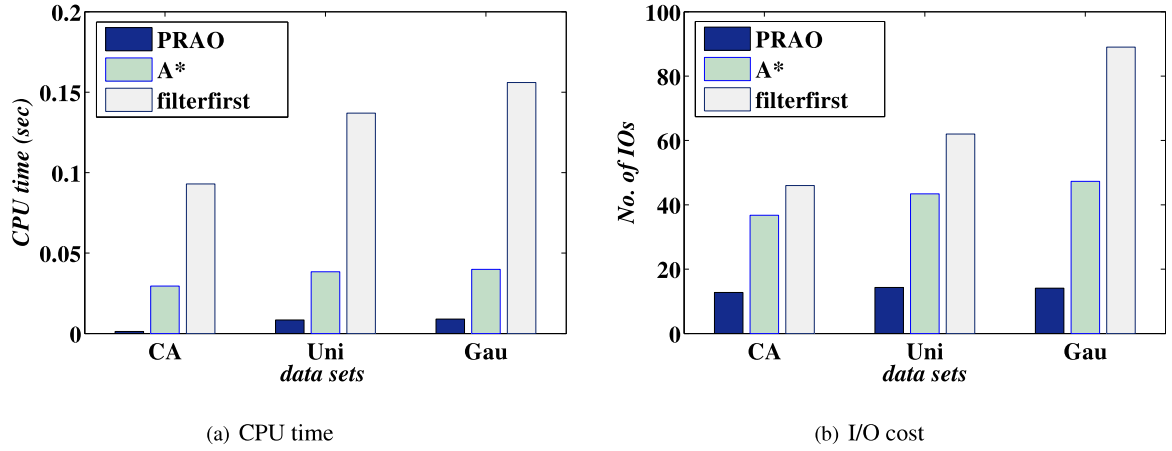


Fig. 8. The PRAO efficiency vs. real/synthetic data sets.

path candidates after the index traversal (i.e., many false alarms are pruned on higher levels with low costs). In contrast, A^* is based on a variant of the Dijkstra's algorithm starting from the source point to search for paths toward all directions, and many searching paths may not even be toward the destination point, which may consume extra computation costs for searching and refinement. Similarly, the second baseline, *filterfirst*, takes much computational time in pruning unnecessary edges at the filtering step. Similar to A^* , *filterfirst* struggles from exploring paths that may not be toward the destination. Therefore, our PRAO algorithm is more efficient than both A^* and *filterfirst* algorithms, as confirmed by Fig. 8.

Index Construction and Update Time vs. Real/Synthetic Data Sets: Fig. 9 presents the efficiency of the index construction and dynamic updates (as discussed in Section 5.3) over CA, Uni, and Gau data sets. From experimental results, the index construction takes about 13.9 s for CA data set, 31.3 s for Uni data, and 33.2 s for Gau data. As we discussed earlier, PRAO builds the R^* -tree index offline only once. Moreover, we can see that the cost of updating the index is low, for example, 0.011 s for CA (with $|V| = 21,048$ and $|E| = 21,693$), and about 0.017 s for Uni and Gau (with $|V| = 30$ K and $|E| = 40$ K).

In the sequel, to evaluate the robustness of our PRAO approach, we will vary different parameter values over synthetic data sets and report the experimental results.

Effect of Probabilistic Weather Confidence Threshold α : Fig. 10 shows the effect of the probabilistic threshold α for weather conditions on the PRAO performance over Uni and Gau data sets, where α varies from 0.1 to 0.9, and other parameters are set to their default values (see Table 3). From figures, for different α values, we can see that the CPU time and I/O cost remain low (e.g., about 0.01 s and less than 15 I/O accesses). This indicates the efficiency of our PRAO approach against different probabilistic weather threshold α .

Effect of Weather Obstacle Threshold ϵ : Fig. 11 reports the CPU time and I/O cost of our PRAO approach, by varying the weather obstacle threshold ϵ from 20 to 90. When threshold ϵ is small, many edges (in turn, paths) will be considered as containing weather-based obstacles, and thus high pruning power can be achieved, which incurs low CPU time and I/O cost. From the experimental results, for different ϵ values, the PRAO approach is efficient, with low CPU time (i.e., 0.005~0.015 s) and I/O cost (i.e., 10~15 I/Os), which confirms the efficiency of our PRAO approach with respect to different ϵ value.

Effect of Valid Path Length σ : Fig. 12 varies the average length, σ , of valid query paths from 5 to 30, where other parameters

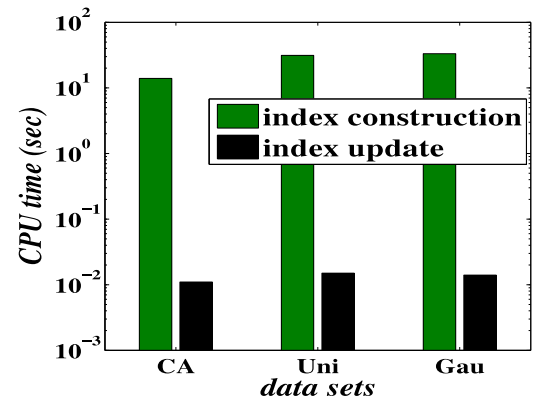


Fig. 9. The index construction and update times vs. real/synthetic data sets.

are set to default values. Intuitively, longer paths from PRAO answer set lead to higher search cost. Thus, when σ increases, as shown in Fig. 12(a), the CPU time also becomes higher. On the other hand, the I/O cost is not very sensitive to the σ value. From experimental results, the CPU time and I/O cost remain low (i.e., less than 0.03 s and 10~15 I/Os, respectively) (see Fig. 13).

Effect of the Size, $|S|$, of Keyword Obstacle Set: Fig. 13 examines the effect of the number, $|S|$, of ad-hoc obstacle keywords in set S , where $|S| = 1, 3, 5, 8$, and 10, and default values are used for other parameters. In figures, the CPU time and I/O cost are not very sensitive to the value of $|S|$. The CPU time for different $|S|$ remains low (i.e., around 0.01 s), whereas the number of I/Os is about 10~15, which indicates the efficiency of our proposed PRAO approach against $|S|$.

Effect of the Number, $|V|$, of Vertices: Fig. 14 demonstrates the scalability of our PRAO approach, by changing the number, $|V|$, of vertices in road networks G from 10 K to 1.2M, where other parameters are set to their default values. With larger road networks (i.e., larger $|V|$), there are more candidate paths between source src and destination dst . Therefore, as shown in Fig. 14(a), the CPU time tends to smoothly increase for larger $|V|$. Both CPU time and I/O cost remain low (i.e., 0.005~0.033 s and 11~51 I/Os, respectively). This confirms the scalability of our proposed PRAO approach against large road networks.

Effect of the Number, d , of pivots: Fig. 15 illustrates the performance of our PRAO algorithm varying the number, d , of pivots from 2 to 15, where other parameters are set to their default values. From figures, we can see that more pivots lead to better




Fig. 10. The PRAO efficiency vs. probabilistic weather confidence threshold α .




Fig. 11. The PRAO efficiency vs. weather obstacle threshold ϵ .

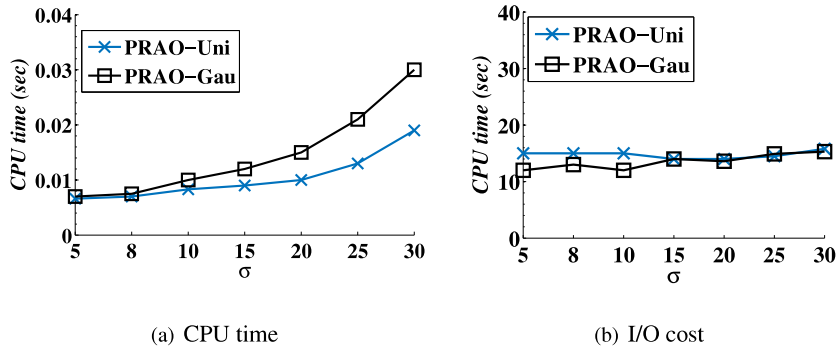


Fig. 12. The PRAO efficiency vs. valid path length σ .

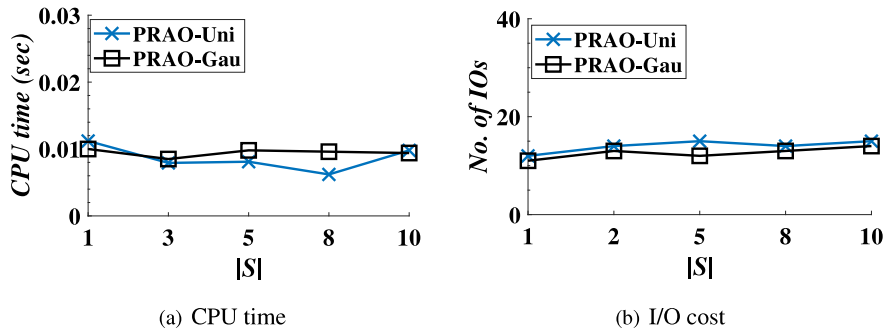





Fig. 13. The PRAO efficiency vs. the size, $|S|$, of keyword obstacle set.

PRAO performance (including CPU time and I/O cost). This is because more pivots will incur tighter bounds for the traveling time pruning (as discussed in Section 4.3). From experimental results, The CPU time and I/O cost with different numbers of

pivots remain low (i.e., below 0.013 s and below 13 page accesses, respectively).

Size Evaluation of Queue \mathcal{Q} and Heap \mathcal{H} : Next, we study the behaviors of the priority queue \mathcal{Q} and the heap \mathcal{H} during the processing of the PRAO-QP procedure. Figs. 16(a) and 16(b) report

Fig. 14. The PRAO efficiency vs. the number, $|V|$, of vertices.Fig. 15. The PRAO efficiency vs. the number, d , of pivots.Fig. 16. The maximum sizes of the queue Q and the heap H of PRAO-QP procedure while processing queries.

the maximum sizes of queue Q and heap H , respectively, during PRAO query processing, where $|V|$ varies from 10K to 1.2M, and other parameters are set to their default values. From figures, we can see that, when the size of the data set increases, the maximum sizes of both Q and H increase. Nevertheless, even for large data sets, the sizes of Q and H remain small. From our experimental results, even with a large road network with 1.2M vertices, the maximum sizes of Q and H are less than 211 and 600, respectively. This confirms the space efficiency of our PRAO query algorithm that utilizes data structures Q and H to enable efficient PRAO query processing.

We also tested other parameters (e.g., the number, $|E|$, of edges) and data distributions (e.g., skewed vertex, or obstacle keyword distributions), and do not report similar experimental results. In summary, our PRAO approach can achieve good performance, in terms of both CPU time and I/O cost.

8. Related work

Query Processing on Road Networks: There are many existing works on the data model of spatial road networks (or graphs), whose road segments (edges) are associated with deterministic

weights. The *shortest path query* is a classical query over road networks, which retrieves the shortest path between source and destination on road networks. Dijkstra [1] proposed the well-known Dijkstra's algorithm to solve such a problem. In order to improve the query efficiency, several variants have been proposed to heuristically prune the search space [2] or materialize some paths [3,4,15]. Huang et al. [5] studied the shortest path search, by verifying spatial constraints (e.g., altitude) of the passing areas through the join operator. Li et al. [6] explored the shortest path queries between source and destination that pass through some types of interesting data points. Song et al. [19] studied the repairing of inconsistent timestamps that do not conform to the required temporal constraints. Neighborhood constraints have been studied in [20,21], where the conflict in roads is detected. Wang et al. [22,23], considered constraints with AND/XOR semantics, in addition to the pairwise constraints. This work filled the missing events referring to the network constraints. Another graph repair approach was proposed in [24] that considers not only the constraints in network structure, but also the names (labeling) of events. Different from prior works, this paper studies obstacle-based shortest path search problem (i.e., PRAO), by considering both ad-hoc keyword-based and weather-based obstacles in the future during the path search,

which have not been studied before. Thus, we cannot borrow previous techniques to answer PRAO queries.

Apart from the shortest path query, many other query types have been studied in spatial road networks, for example, range queries [25,26], k -nearest neighbor (k NN) queries [25,27], k -nearest neighbor (k NN) queries with the incorporation of social influence [28], aggregate nearest neighbor query [29,30], reverse nearest neighbor queries [31], multi-source skyline queries [32], and so on. Furthermore, the similarity search over uncertain/certain graph databases has been extensively studied in [33–36], where the goal of this query is to find a set of subgraphs from the graph database that are similar to the subgraph query. Prior works on uncertain graphs follow *filtering-and-verification* framework. In the filtering step, different pruning techniques have been proposed to reduce the search space. In particular, they usually designed tight upper/lower bounds of the subgraph similarity that can be utilized to prune unmatched graphs. The verification step is used to compute the actual similarity between a candidate probabilistic graph g and the query graph q . However, pruning methods proposed in [33,35,36] cannot be applied to PRAO problem, since our PRAO query is a path routing problem (rather than a probabilistic (sub)graph matching problem). Thus, the pruning methods that prune (sub)graphs (not matching with a query graph) in [33,35,36] cannot be directly used for pruning a path from a source to a destination. Similarly, [34] retrieves the matching subgraphs over certain graph database. In this work, authors aim to scale the subgraph matching problem to very large datasets using the parallelism. The proposed techniques are for certain datasets under the distributed settings, and cannot be applied to solve our PRAO problem over road networks with ad-hoc probabilistic weather information.

There is another query type, keyword search, over uncertain/certain graph databases, which retrieves a set of nodes with certain keywords from the graph. Yuan et al. [37] proposed solutions to the keyword search following the *filtering-and-verification* framework. The filtering process uses three pruning techniques, *existence probability pruning*, *path-based probability pruning*, and *tree-based probability pruning*. In fact, these pruning techniques cannot be applied to our PRAO problem, since they do not consider ad-hoc weather constraints, and the returned answers are subgraphs (instead of a path) containing keywords (rather than excluding obstacle keywords). Moreover, in PRAO, the value of weather condition of an edge changes over time, for example, an edge might not satisfy weather condition at 9:00am, but it may be safe to be taken at 9:15am. Similarly, Yuan et al. [38] considered certain data sets over the distributed environment, whereas our PRAO problem is considered over a centralized machine and it involves the probabilistic weather condition (e.g., the temperature is 50° at 8am with probability 0.6)."

In these works, road networks are usually modeled by weighted graphs, and the traveling distance/time is often used to measure the path length between two points on road networks. In addition to weights on edges, in this paper, we also model real scenarios of road networks, by including the properties of road segments (i.e., keywords) and the predicted weather conditions on road networks, which integrate many heterogeneous road-network data. Therefore, with different query types and data models for road networks, we cannot directly apply previous solutions to solve our PRAO problem.

Probabilistic Graph Management: In the literature, there are some query types, such as path queries [8,39], reachability queries [40], and trip planning [41], over road networks, by considering uncertain traffic conditions or the availability of road segments (e.g., roads under construction). In these works, road networks are modeled by probabilistic graphs, where either edges are associated with uncertain velocity samples of vehicles, or

edges have existence probabilities (i.e., edges may exist or not exist). In contrast, in this paper, we consider uncertain weather conditions at vertices of road networks (e.g., wind speed 10 mph with the forecasting accuracy or probability 0.8), and utilize the known weather values at vertices to estimate (unknown) weather values at some points on edges (roads), associated with confidences. Moreover, our PRAO query predicates include avoiding ad-hoc weather-based obstacles on the path, which takes into account uncertain weather conditions, and has not been studied before. Our work thus needs to design novel weather-based pruning, specific for PRAO, to enable efficient query answering.

Queries in the Presence of Obstacles: There are some works [42–45] on queries taking into account physical obstacles (e.g., buildings and lakes) in the Euclidean space, so that the traveling paths should be detoured. For example, Zhang et al. [43] studied queries such as the range search, nearest neighbors, e -distance joins, and closest pairs in the Euclidean space and with physical obstacles. Due to such obstacles, visible nearest neighbor queries [44] or visible reverse nearest neighbor queries [45] are proposed, which retrieve nearest neighbors (NNs) or reverse nearest neighbors (RNNs) of a query point Q that are visible to q (i.e., without obstacles between q and NNs/RNNs). However, existing works usually considered static physical obstacles and/or obstacle-aware queries in the Euclidean space. In contrast, in our PRAO problem, users can arbitrarily specify *ad-hoc* obstacles on roads (including both keyword and weather obstacles) rather than *static* ones, and furthermore our PRAO queries are performed in graphs (road networks) to retrieve obstacle-aware shortest paths, instead of the Euclidean space. Therefore, we cannot directly use previous methods to tackle our PRAO problem.

9. Conclusion

In this paper, we formalize and tackle the problem of *path routing over road networks with ad-hoc obstacles* (PRAO), which retrieves paths from source to destination with the smallest traveling time, by avoiding ad-hocly specified keyword-based and weather-based obstacles. To efficiently process the PRAO query, in this paper, we design effective pruning methods (w.r.t. keywords, weather conditions, and traveling times) to filter out false alarms, and propose efficient indexing and query processing algorithms to answer PRAO queries. Extensive experiments have demonstrated the efficiency and effectiveness of our proposed PRAO approach over both real and synthetic data sets.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Xiang Lian is supported by NSF OAC, USA No. 1739491 and Lian Startup No. 220981, Kent State University, USA. En Cheng is supported by UA Startup No. 207993, The University of Akron, USA.

References

- [1] E.W. Dijkstra, A note on two problems in connexion with graphs, *Numer. Math.* (1959).
- [2] R.-M. Kung, E. Hanson, Y. Ioannidis, T. Sellis, L. Shapiro, M. Stonebraker, Heuristic search in data base systems, in: *Expert Database Systems*, 1986.

- [3] R. Agrawal, S. Dar, H. Jagadish, Direct transitive closure algorithms: Design and performance evaluation, *ACM Trans. Database Syst.* 15 (3) (1990) 427–458.
- [4] Y. Ioannidis, R. Ramakrishnan, L. Winger, Transitive closure algorithms based on graph traversal, *ACM Trans. Database Syst.* 18 (1993), <http://dx.doi.org/10.1145/155271.155273>.
- [5] Y.-W. Huang, N. Jing, E.A. Rundensteiner, Integrated query processing strategies for spatial path queries, in: *Proceedings 13th International Conference on Data Engineering*, IEEE, 1997, pp. 477–486.
- [6] F. Li, D. Cheng, M. Hadjieleftheriou, G. Kollios, S.-H. Teng, On trip planning queries in spatial databases, in: *International Symposium on Spatial and Temporal Databases*, Springer, 2005, pp. 273–290.
- [7] B. Ding, J.X. Yu, L. Qin, Finding time-dependent shortest paths over large graphs, in: *Proceedings of the 11th International Conference on Extending Database Technology: Advances in Database Technology*, vol. 2008, 2008, <http://dx.doi.org/10.1145/1353343.1353371>.
- [8] M. Hua, J. Pei, Probabilistic path queries in road networks: traffic uncertainty aware path selection, in: *Proceedings of the 13th International Conference on Extending Database Technology*, ACM, 2010, pp. 347–358.
- [9] U.T. NETWORK, The arizona Republic, 2017, URL <https://yourbusiness.azcentral.com/weather-affect-supply-chain-27783.html>.
- [10] The Dark Sky Team, Dark Sky, <https://darksky.net/dev/>.
- [11] N. Dalvi, D. Suciu, Efficient query evaluation on probabilistic databases, *VLDB J.* 16 (4) (2007).
- [12] D. Shepard, A two-dimensional interpolation function for irregularly-spaced data, in: *Proceedings of the 1968 23rd ACM National Conference*, ACM, 1968, pp. 517–524.
- [13] A. Al-Baghdadi, X. Lian, E. Cheng, Efficient Path Routing Over Road Networks in the Presence of Ad-Hoc Obstacles, *Tech. Rep.*, 2019, [arXiv: 1910.04786](https://arxiv.org/abs/1910.04786).
- [14] C.-Y. Chan, Y.E. Ioannidis, Bitmap index design and evaluation, *ACM SIGMOD Rec.* 27 (2) (1998) 355–366.
- [15] A.V. Goldberg, C. Harrelson, Computing the shortest path: A search meets graph theory, in: *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, 2005, pp. 156–165.
- [16] C. Shahabi, M.R. Kolahdouzan, M. Sharifzadeh, A road network embedding technique for k-nearest neighbor search in moving object databases, in: *Proceedings of the 10th ACM International Symposium on Advances in Geographic Information Systems*, 2002.
- [17] N. Beckmann, H.-P. Kriegel, R. Schneider, B. Seeger, The R*-tree: an efficient and robust access method for points and rectangles, *ACM SIGMOD Rec.* 19 (2) (1990).
- [18] Real datasets for spatial databases: Road networks and points of interest, <https://www.cs.utah.edu/~lifeifei/SpatialDataset.htm>.
- [19] S. Song, Y. Cao, J. Wang, Cleaning timestamps with temporal constraints, *Proc. VLDB Endow.* 9 (10) (2016) 708–719.
- [20] S. Song, B. Liu, H. Cheng, J.X. Yu, L. Chen, Graph repairing under neighborhood constraints, *VLDB J.* 26 (5) (2017) 611–635.
- [21] S. Song, H. Cheng, J.X. Yu, L. Chen, Repairing vertex labels under neighborhood constraints, *Proc. VLDB Endow.* 7 (11) (2014) 987–998.
- [22] J. Wang, S. Song, X. Zhu, X. Lin, Efficient recovery of missing events, *Proc. VLDB Endow.* 6 (10) (2013) 841–852.
- [23] J. Wang, S. Song, X. Zhu, X. Lin, J. Sun, Efficient recovery of missing events, *IEEE Trans. Knowl. Data Eng.* 28 (11) (2016) 2943–2957.
- [24] J. Wang, S. Song, X. Lin, X. Zhu, J. Pei, Cleaning structured event logs: A graph repair approach, in: *2015 IEEE 31st International Conference on Data Engineering*, IEEE, 2015, pp. 30–41.
- [25] D. Papadias, J. Zhang, N. Mamoulis, Y. Tao, Query processing in spatial network databases, in: *VLDB*, 2003, URL <http://dl.acm.org/citation.cfm?id=1315451.1315520>.
- [26] H. Jeung, M.L. Yiu, X. Zhou, C.S. Jensen, Path prediction and predictive range querying in road network databases, *VLDB J.* 19 (4) (2010) 585–602.
- [27] C. Shahabi, M.R. Kolahdouzan, M. Sharifzadeh, A road network embedding technique for k-nearest neighbor search in moving object databases, *Geoinformatica* 7 (3) (2003) 255–273.
- [28] Y. Yuan, X. Lian, L. Chen, Y. Sun, G. Wang, RSKNN: kNN search on road networks by incorporating social influence, *IEEE Trans. Knowl. Data Eng.* 28 (6) (2016) 1575–1588.
- [29] M.L. Yiu, N. Mamoulis, D. Papadias, Aggregate nearest neighbor queries in road networks, *IEEE Trans. Knowl. Data Eng.* 17 (6) (2005) 820–833.
- [30] D. Papadias, Y. Tao, K. Mouratidis, C.K. Hui, Aggregate nearest neighbor queries in spatial databases, *ACM Trans. Database Syst.* 30 (2) (2005) 529–576.
- [31] M.L. Yiu, D. Papadias, N. Mamoulis, Y. Tao, Reverse nearest neighbors in large graphs, *IEEE Trans. Knowl. Data Eng.* 18 (4) (2006) 540–553.
- [32] K. Deng, X. Zhou, H. Tao, Multi-source skyline query processing in road networks, in: *2007 IEEE 23rd International Conference on Data Engineering*, IEEE, 2007, pp. 796–805.
- [33] Y. Yuan, G. Wang, L. Chen, H. Wang, Graph similarity search on large uncertain graph databases, *VLDB J.* 24 (2) (2015) 271–296.
- [34] Y. Yuan, G. Wang, J.Y. Xu, L. Chen, Efficient distributed subgraph similarity matching, *VLDB J.* 24 (3) (2015) 369–394.
- [35] Y. Yuan, G. Wang, H. Wang, L. Chen, Efficient subgraph search over large uncertain graphs, *Proc. VLDB Endow.* 4 (11) (2011) 876–886.
- [36] Y. Yuan, G. Wang, L. Chen, H. Wang, Efficient subgraph similarity search on large probabilistic graph databases, *Proc. VLDB Endow.* 5 (9) (2012) 800–811.
- [37] Y. Yuan, G. Wang, L. Chen, H. Wang, Efficient keyword search on uncertain graph data, *IEEE Trans. Knowl. Data Eng.* 25 (12) (2013) 2767–2779.
- [38] Y. Yuan, X. Lian, L. Chen, J.X. Yu, G. Wang, Y. Sun, Keyword search over distributed graphs with compressed signature, *IEEE Trans. Knowl. Data Eng.* 29 (6) (2017) 1212–1225.
- [39] Y. Nie, X. Wu, Shortest path problem considering on-time arrival probability, *Transp. Res. B* 43 (6) (2009).
- [40] R. Jin, L. Liu, B. Ding, H. Wang, Distance-constraint reachability computation in uncertain graphs, *Proc. VLDB Endow.* 4 (2011), URL <http://dl.acm.org/citation.cfm?id=2002938.2002941>.
- [41] X. Lian, L. Chen, Trip planner over probabilistic time-dependent road networks, *IEEE Trans. Knowl. Data Eng.* 26 (8) (2014), <http://dx.doi.org/10.1109/TKDE.2013.159>.
- [42] T. Lozano-Pérez, M.A. Wesley, An algorithm for planning collision-free paths among polyhedral obstacles, *Commun. ACM* 22 (10) (1979), <http://dx.doi.org/10.1145/359156.359164>.
- [43] J. Zhang, D. Papadias, K. Mouratidis, M. Zhu, Spatial queries in the presence of obstacles, in: *International Conference on Extending Database Technology*, Springer, 2004, pp. 366–384.
- [44] S. Nutanong, E. Tanin, R. Zhang, Incremental evaluation of visible nearest neighbor queries, *IEEE Trans. Knowl. Data Eng.* 22 (5) (2010), <http://dx.doi.org/10.1109/TKDE.2009.158>.
- [45] Y. Gao, B. Zheng, G. Chen, W.-C. Lee, K.C. Lee, Q. Li, Visible reverse k-nearest neighbor queries, in: *2009 IEEE 25th International Conference on Data Engineering*, IEEE, 2009, pp. 1203–1206.