Transportation Research Part C 108 (2019) 357–377

Contents lists available at ScienceDirect

Transportation Research Part C

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

Tradeoffs between safety and time: A routing view

Daniel R. Carmody[a], Richard B. Sowers[a,b,*]

[a] Department of Mathematics, University of Illinois at Urbana–Champaign, Urbana, IL 61801, United States
[b] Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana–Champaign, Urbana, IL 61801, United States

ARTICLE INFO

Keywords:
Intelligent urban navigation
Smart city
Safety
Optimization
Accidents
Routing

ABSTRACT

This article proposes a data-driven combination of travel times, distance, and collision counts in urban mobility datasets, with the goal of quantifying how intertwined traffic accidents are in the road network of a city. We devise a bi-attribute routing problem to capture the tradeoff between travel time and accidents. We apply this to a dataset from New York City. By visualizing the results of this computation in a normalized way, we provide a comparative tool for studies of urban traffic.

1. Introduction

1.1. Motivation

Global, mobility-related technological advances, e.g., Google maps, ride-hailing apps of Didi, Lyft, and Uber, and real-time flight status data, have enormously influenced travel. This might in fact be called an “age of transportation”. Building on a strong foundation of software and computation, we are at a point of generational change.

Mathematically, the problem of efficiently getting from an origin to a destination is a classical one, leading directly to Riemannian geometry, the calculus of variations, and later to Dijkstra’s efficient routing algorithm. With the advent of ubiquitous GPS-enabled cell phones, classical routing problems can be connected to real-time traffic models and estimators.

The purpose of this paper is to rethink routing problems to include safety. We want to: (*a*) redefine the notion of best path to include safety, and (*b*) develop broad measures of accident risk which will allow us to compare cities (cf. Dai et al., 2016). A key goal is to determine how intrinsic accidents are to traffic patterns. Our work develops a data-driven methodology for combining data sets to quantify tradeoffs between level of service and safety.

1.2. Problem statement

We aim to study the tradeoff between trip time and accident exposure. Motivated by a constrained optimization (routing)

The authors would like to thank Arnab Chakraborty for some helpful insights on risk. The authors would also like to thank A. Madhusudan and H. Athi, both Department of Industrial and Enterprise Systems Engineering at the University of Illinois at Urbana-Champaign, Urbana, Illinois, for their work on prior versions of this project. This material is based upon work supported by the National Science Foundation under Grant Nos. CMMI 1727785 and DMS 1345032. The authors would like to thank the High Performance Computing Center of the MSFE program at the University of Illinois for computational support. The code for this work is at https://github.com/dcarmod2/VisionZeroTheory_Submission.

* Corresponding author at: Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana–Champaign, Urbana, IL 61801, United States.
E-mail addresses: dcarmod2@illinois.edu (D.R. Carmody), r-sowers@illinois.edu (R.B. Sowers).

https://doi.org/10.1016/j.trc.2019.09.020
Received 20 November 2018; Received in revised form 19 September 2019; Accepted 21 September 2019
Available online 17 October 2019

problem, we consider the problem of optimally routing a vehicle from an origin to a destination, where optimality in our case is a combination of minimizing trip duration and maximizing trip safety. Our core question is: *how much time would you have to give up to achieve a certain level of safety?*. By asking this in a global way over many trips, we propose a descriptive way of quantifying how intertwined accidents are in traffic patterns. By further using time-minimizing trips as a point of comparison in computing relative tradeoffs between trip duration and accidents, we propose a normalized framework for studying and comparing different cities. As cities of different size and character are experimenting and innovating in mobility, we believe that measures of comparison are crucial to identifying best practices in urban management and intelligent transportation systems.

1.3. Datasets

Our methods combine three datasets. In this paper, we will apply our calculations to Manhattan Island weekday traffic data. We require

- speed data on most, if not all the links (directed edges) in a city; we will use Donovan et al. (2016).
- accident data in a city, annotated by time and location; we will use NYPD (n.d.).
- trip data (with some approximation allowed). The dataset of Donovan et al. (2016) is used as the basis of the estimates of Donovan et al. (2016), Donovan and Work (2016).

An important aspect of our approach is that trip data gives us broad statistical properties of travellers.

A similar approach in (Liu et al., 2019) combines multiple accessible urban datasets to provide insights into emission hotspots in traffic networks. The authors similarly use taxi data as a proxy for all traffic, but they utilize kriging with additional images of roads to infer vehicle volume from the taxi data. As more technologies give rise to new and diverse datasets, new insights will become available; see O'Keeffe et al. (2019), Massaro et al. (2016).

In the context of ride-hailing, Alemi et al. (2019) combines multiple data sources to look for factors (i.e. neighborhood type) which strongly influence the frequency of ride-hailing service usage. Another data-driven analysis, using unsupervised learning and visualization, was developed in Dong et al. (2018) to identify and understand spatio-temporal hotspots of ride-sharing activity.

1.4. Related work and contribution

In 1997, the Swedish Parliament passed a bill on traffic safety which stated: "Vision Zero means that eventually no one will be killed or seriously injured within the road transport system" (Belin et al., 2012; Haq, 2006; Johansson, 2009; Tingvall and Haworth, 2000). As of March 2017, 33 U.S. cities have committed to the Vision Zero effort[1] for the design and functioning of the road transport system (Belin et al., 2012; Johansson, 2009; Tingvall and Haworth, 2000).

By some estimates, more than 50% of the world's population currently lives in cities, a percentage which is projected to increase to 70% by the year 2050 (United Nations-World, 2011). The notion of "smart cities" has thus attracted significant attention (Kumar et al., 2018). Personalized location/direction recommendations based on the behavior of city dwellers on various social-network profiles have been developed in Benouaret and Lenne (2016), Berjani and Strufe (2011), Dai et al. (2015), Fu et al. (2014), Noulas et al. (2012), Yang et al. (2015). In the orienteering problem (Gionis et al., 2014; Vansteenwegen et al., 2011), routing problems are augmented by a set of user benefits along different links. Other routing problems have focused on cost of exposure to crime (Galbrun et al., 2016) and environmental costs such as fuel consumption (Boriboonsomsin et al., 2012; Dhaou, 2011; Guo et al., 2015a,b; Minett et al., 2011; Nie and Li, 2013; Rakha et al., 2012; Yang et al., 2018; Yang et al., 2014).

In Bruglieri et al. (2019) the authors study a constrained optimization problem for optimally routing alternative fuel vehicles where the constraints are related to the fact that vehicles must not run out of fuel. That paper focussed on efficient methods for solving a fixed optimization problem; our focus is on understanding the relationships among solutions to a parametric family of optimization problems for an entire network.

Multi-objective optimization techniques are used in Zero et al. (2017), Zografos and Davis (1989) to take travel cost and travel risk into account when routing hazardous materials. Those works consider routing for a single origin-destination pair; our work considers an ensemble of origin-destination pairs as a means to understand network-wide properties (and unlike the fuzzy cost analysis of Zero et al. (2017), we focus on a deterministic analysis).

Our work here is based on historical data; our algorithm is thus *descriptive*, and not *predictive*; we make no claim about predicting accidents. Accidents are affected by a complex combination of road and human factors (Caliendo et al., 2007; Lee et al., 2003; Lord and Persaud, 2000). The intent here is to provide city-wide measures by which cities (and traffic safety policy implementations) can be assessed. We claim no contributions to questions of realtime routing.

Routing problems also occur in computer science applications; see Chen and Chin (1990).

1.5. Outline

We start out in Section 2 (Section 2.1) by mathematically formulating a data-related, bi-attribute optimization problem (see also

[1] see https://visionzeronetwork.org/resources/vision-zero-cities/.

Wang et al., 2018; Yang et al., 2014). The corresponding objective function can be defined for a single trip (i.e., an origin-destination pair) using data consisting of traffic speed estimates and accident counts. We then extend this to a citywide optimization problem by considering a statistical sample of trips. We also construct (Sections 2.1 and 2.2) several measures on the output of these optimization problems. These measures are designed to capture the tradeoff between time and accident risk. In Section 3, we apply this analysis to the dataset mentioned in Section 1.3, starting with a broad description of the data in Section 3.1. Several different perspectives on the results are in Sections 3.2 and 3.3. An illustration of the comparative aspects of this work is developed in Section 3.4, where the calculations are specialized to Upper/Central Manhattan. Computational considerations are addressed in Section 3.5. Finally future work is described in Section 4. Appendix A connects the optimization problem of Section 2.1 to the theory of bi-objective costs. Some of our results will be parametric plots of tradeoffs; Appendix B reviews some numerical aspects of such plots. Finally, Appendix C contains an empirical study of desirable monotonicity properties.

2. Setup and theory

2.1. Optimization problem

We start out with a data-related optimization problem, with $\mathcal{T}$ being a collection of origin-destination (O, D) pairs (i.e., trips). As in Donovan and Work (2016), we want $\mathcal{T}$ to reflect typical travel statistics. For each $(O, D) \in \mathcal{T}$, let $\mathcal{P}(O, D)$ be the collection of paths which start at O and end at D. Let $\mathcal{P} \stackrel{\text{def}}{=} \bigcup_{(O,D)\in\mathcal{T}} \mathcal{P}(O, D)$.

We are interested in several measures of trips. For $p \in \mathcal{P}$, define

$$\begin{aligned} \mathbf{A}(p) &\stackrel{\text{def}}{=} \text{historical number of accidents/hr along } p, \\ \mathbf{L}(p) &\stackrel{\text{def}}{=} \text{trip length of } p, \\ \mathbf{T}(p) &\stackrel{\text{def}}{=} \text{trip duration of } p. \end{aligned} \tag{1}$$

The normalization by hour in $\mathbf{A}$ reflects risk which is stationary over some subset of a day (see Section 3.2). For $\alpha \in [0, 1]$, define a cost function

$$\mathbf{c}_\alpha(p) \stackrel{\text{def}}{=} (1 - \alpha)\mathbf{T}(p) + \alpha\mathbf{A}(p). \tag{2}$$

For $\alpha = 0$, this cost gives the trip duration; for $\alpha = 1$, it gives the accidents per hour along the path. We will use α as a parameter to capture the tradeoff between risk and travel time (see the figures of Section 3).

For each $(O, D) \in \mathcal{T}$, we can compute

$$\mathbf{c}_\alpha(O, D) \stackrel{\text{def}}{=} \min\{\mathbf{c}_\alpha(p): p \in \mathcal{P}(O, D)\}$$

and let $p^*_\alpha(O, D)$ be an argmin of this minimization problem; i.e.,

$$\mathbf{c}_\alpha(p^*_\alpha(O, D)) = \mathbf{c}_\alpha(O, D). \tag{3}$$

Then $p^*_\alpha(O, D)$ is a geodesic from O to D in the metric defined by $\mathbf{c}_\alpha$. Since $\mathbf{c}_\alpha$ is additive along links in a path, we can use standard off-the-shelf routing algorithms (e.g., Dikjstra's algorithm) to find $p^*_\alpha(O, D)$.

Noting that accidents are integer-valued, the path with the least accidents (minimizing c_1) may fail to be unique (two paths from origin O to destination D may both have 5 accidents). Given that trip times are floating point numbers (e.g., one path from O to D may take 14.6 minutes and the other may take 14.99 minutes according to our data), we can usually break any tie in selecting $p^*_1(O, D)$ (the path with fewest accidents) in favor of the faster path (the path which takes 14.6 minutes in our example). We do this by lexicographically ordering the (accident,time) pair (with $(5, 14.6) < (5, 14.99)$; see also Appendix A).

In cases where the data of floating-point travel times still does not lead to a unique solution of (3), we take any of the minimizers (which may depend on software implementations).

As α increases, our cost function gives more weight to accidents and less to travel time. See Fig. 1 for some examples of geodesics between a fixed origin and destination as α varies. The path $p^*_1(O, D)$ passes the fewest accidents, but may take the longest. In fact, the cost function 2 is connected to solutions to

$$\min\{\mathbf{T}(p): \mathbf{A}(p) \leqslant \widehat{A},\ \ p \text{ a path from } O \text{ to } D\}$$

where $\widehat{A}$ is a prespecified upper limit on the accidents on a path; see Miettinen (1998), Theorem 3.2.5.

From the standpoint of a driver seeking to balance between accident risk and trip time, we can assess the pair (O, D) in several ways. Define:

$$\tau_\alpha(O, D) \stackrel{\text{def}}{=} \frac{\mathbf{T}(p^*_\alpha(O, D))}{\mathbf{T}(p^*_0(O, D))},\ \ \sigma_\alpha(O, D) \stackrel{\text{def}}{=} \frac{\mathbf{A}(p^*_\alpha(O, D))/\mathbf{L}(p^*_\alpha(O, D))}{\mathbf{A}(p^*_0(O, D))/\mathbf{L}(p^*_0(O, D))}. \tag{4}$$

Informally, $\tau_\alpha(O, D)$ tells us how much longer (in time) a path is if we try to avoid accidents with strength α, and $\sigma_\alpha(O, D)$ tells us the accidents per meter (APM) along p^*_α, where we normalize the accidents per meter by that of the time minimizing path.

The more accidents are penalized, the more $p^*_\alpha(O, D)$ should deviate from $p^*_0(O, D)$ (e.g., the absolute safest path between two

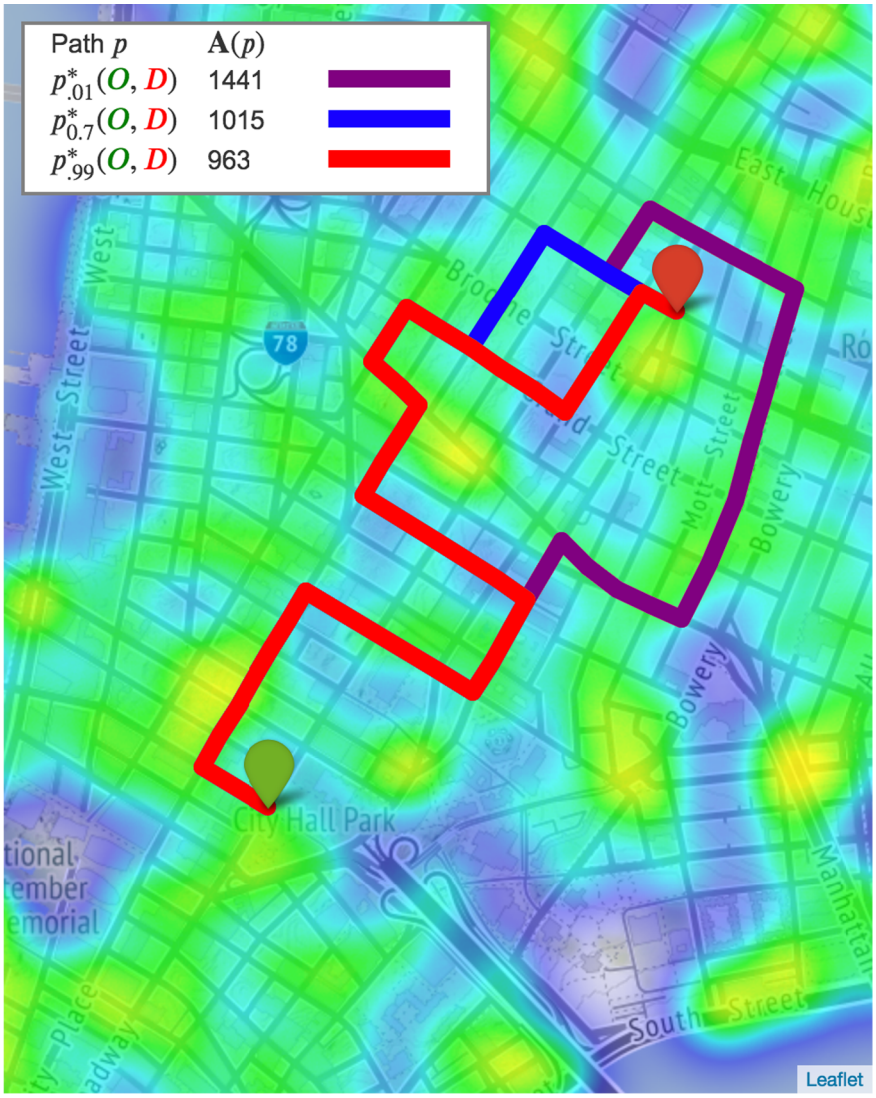

Fig. 1. Paths for different values of α from node 2821304137 (green marker) to node 42442247 (red marker). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

nearby intersections might conceivably follow a long detour to avoid historical accidents). In other words, one should seek safer paths at the cost of longer (in time) trips. Thus the map $\alpha \mapsto \mathbf{T}(p^*_\alpha(O, D))$ and thus $\alpha \mapsto \tau_\alpha(O, D)$ should be increasing.[2] Similarly, $\alpha \mapsto \mathbf{A}(p^*_\alpha(O, D))$ is decreasing, but we have no theoretical assurance that when we divide by path length, as in (4), the resulting ratios are also decreasing. For a fixed $(O, D) \in \mathcal{T}$, $\mathbf{A}(p^*_1(O, D)) \leqslant \mathbf{A}(p^*_0(O, D))$ (accidents are decreasing in α), but if $\mathbf{L}(p^*_1(O, D))$ is sufficiently small compared to $\mathbf{L}(p^*_0(O, D))$, then $\sigma_1(O, D) > \sigma_0(O, D)$.

Remark 2.1. Since (3) reflects a discrete optimization problem parametrized by the continuous parameter α, the map $\alpha \mapsto p^*_\alpha(O, D)$ will generically be constant over an interval (of α values). For each $\alpha \in [0, 1]$, define

$$\Gamma_\alpha \stackrel{\text{def}}{=} \{\alpha' \in [0, 1]: p^*_{\alpha'}(O, D) = p^*_\alpha(O, D)\},$$

Also define γ: $[0, 1] \to \mathbb{R}^2$ as

$$\gamma(\alpha) \stackrel{\text{def}}{=} (\bar{\tau}(\alpha), \bar{\sigma}(\alpha)), \quad \alpha \in [0, 1].$$

Then

[2] for simplicity of exposition, when considering monotonicity, "increasing" should more accurately be interpreted as nondecreasing, and "decreasing" should more accurately be interpreted as nonincreasing.

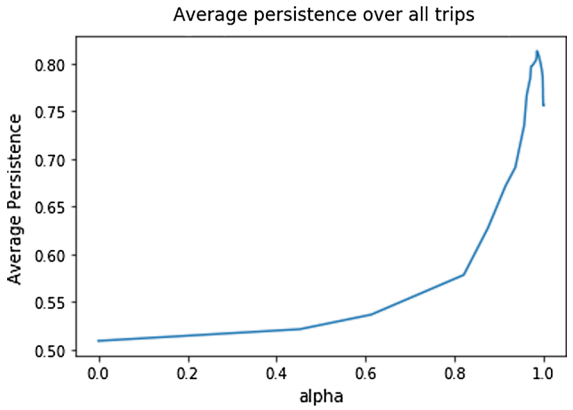

Fig. 2. Average persistence over all (O, D) pairs of various $\alpha \in [0, 1]$.

$$\text{pers}(\alpha) \overset{\text{def}}{=} \frac{arc\ length(\gamma(\Gamma_\alpha))}{arc\ length(\gamma([0,1]))}$$

is the percent coverage of the $(\bar{\tau}, \bar{\sigma})$-Pareto Front by Γ_α (see Appendix A) and measures the normalized "size" of Γ_α. This gives a measure of robustness of the best path.

Numerically, we approximate γ by a piecewise-linear function, in which case arc length is a finite sum of distances in $\mathbb{R}^2$ (see Appendix B). Theoretically, if γ is differentiable, arc length is defined as an integral against an arc length form. Most generally, arc length is given by one-dimensional Hausdorff measure.

Fig. 1, for example, shows several paths from node $O = 2821304137$ to node $D = 42442247$. Some numerical exploration (using results of Appendix B) shows that $0.91 < \text{pers}(1) < 0.96$, $0.04 < \text{pers}(0) < 0.09$, and $0 < \text{pers}(0.7) < 0.05$. A traveller who is unsure of her or his relative preference for travel time compared to historical accidents per meter might prefer the path $p_1^*(O, D)$ from O to D, as compared to $p_0^*(O, D)$ or $p_{0.7}^*(O, D)$;$p_1^*(O, D)$ is the best path for a wider range of α's.

Fig. 2 shows a plot of the average persistence over all trips of various values of $\alpha \in [0, 1]$. The global maximum persistence is achieved at $\alpha = 0.985$, indicating that $p_{0.985}^*(O, D)$ is on average a good choice of path for a decision maker who is unsure of their preference for travel time or safety. Since we are effectively excluding links for which we do not have travel times, the accidents along $p_1^*(O, D)$ may in fact be smaller (and travel times larger) in reality than for our model; accident-free paths may significantly meander in the larger graph corresponding to real travel times. Informally, more missing links may increase this effect, meaning larger values of persistence at $\alpha = 1$.

Remark 2.2. In Fig. 12 (which is described in (13)), there are a few points which lie above the line $y = 1$, corresponding to (O, D) trips for which $\sigma_1(O, D) > 1 = \sigma_0(O, D)$. For these few (O, D) trips, $\alpha \mapsto \sigma_\alpha(O, D)$ thus fails to be decreasing in α; the safest path is relatively short, even though it may take a longer time. Congestion being one reason why a short path may take a long time, it makes sense that congested paths which are in fact the safest are also rare.

See Appendix C for a detailed discussion of the monotonicity of σ_α.

We are interested in large citywide statistics and datasets; a crucial computational tool is efficient shortest-path algorithms. For any given trip, time, accidents, and length of trip are extensive quantities which can be written as sums of intensive quantities along pieces of the path. Accidents per meter, which is a ratio of such extensive quantities, cannot in general be written as such a sum of intensive quantities along the path (informally, a ratio of sums is not a sum of ratios). For computational efficiency, we thus compute σ_α from the extremals $p_\alpha^*(O, D)$ rather than trying to directly minimize accidents/meter.

We also note that the normalization of (4) means that the units of length (meter vs mile) cancel out.

We can capture the overall statistics of the city by averaging over all $(O, D) \in \mathcal{T}$; define

$$\bar{\tau}_\alpha \overset{\text{def}}{=} \frac{1}{|\mathcal{T}|} \sum_{(O,D)\in\mathcal{T}} \tau_\alpha(O, D), \quad \bar{\sigma}_\alpha \overset{\text{def}}{=} \frac{1}{|\mathcal{T}|} \sum_{(O,D)\in\mathcal{T}} \sigma_\alpha(O, D). \tag{5}$$

Note that $\alpha \mapsto \bar{\tau}_\alpha$ inherits the appropriate monotonicity in α from $\alpha \mapsto \tau_\alpha(O, D)$: an average of monotonic increasing functions is monotonic increasing. However, since the $\sigma_\alpha(O, D)$'s themselves are not assured to be monotone, we cannot guarantee monotonicity of $\alpha \mapsto \bar{\sigma}_\alpha$ (again, see Appendix C for empirical results).

Tradeoff Plot 1 (*APM vs Time*). To quantify the macroscopic tradeoff between safety and time from a driver's perspective, we parametrically plot $\{(\bar{\tau}_\alpha, \bar{\sigma}_\alpha)\colon\ 0 \leqslant \alpha < 1\}$. This tells us how much time we would have to give up to achieve a certain reduction in accidents/meter.

By basing all of our analysis on the cost function (2), we take advantage of a number of relatively nice properties. Given a cost function which is additive over links, cheapest paths have the property that subpaths of the cheapest path are themselves cheapest paths. This allows for efficient calculation of various cheapest-path algorithms (e.g., Dijkstra algorithms). Accidents/meter, on the

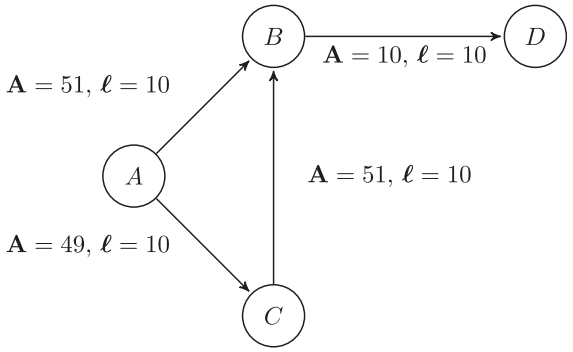

Fig. 3. A graph with a non-optimal subpath of an optimal path.

other hand, is not additive; Fig. 3 shows an example where this subpath property thus fails for accidents/meter. The cheapest path from A to D is the path $A \to B \to D$ with an accident/meter cost of 61 accidents/20 m = 3.05 accidents/meter. On the other hand, the cheapest path from A to B is the path $A \to C \to B$ with an accident/meter cost of 100 accidents/20 m = 5 accidents/meter (compared to 5.1 accidents/meter going directly from A to B); $A \to C \to B$ is clearly not a subpath of $A \to B \to D$. Any greedy routing algorithm will assign to node B the data from the path $A \to C \to B$, which will lead to the algorithm erroneously returning $A \to C \to B \to D$ as a cheapest path.

2.2. Accidents per vehicle-meter

A city planner might be interested in normalizing by traffic counts as an aid to studying how to best address large-scale mobility questions. From that standpoint, risk along a road might be proportional to the length of the road and the number of vehicles on the road; the ratio

$$\frac{\text{accidents}}{\text{vehicle count} \times \text{length}} \tag{6}$$

quantifies this risk (in units of accidents per vehicle-meter or vehicle-mile).

We can study the risk factor (6) using calculations like in Section 2.1. For any link ℓ, define

$$\mathbf{N}(\ell) \stackrel{\text{def}}{=} \text{traffic count along } \ell;$$

for any path p, we can then define

$$\mathbf{VM}(p) \stackrel{\text{def}}{=} \sum_{\ell \in p} \mathbf{N}(\ell) \times \mathbf{L}(\ell), \tag{7}$$

the sum being over the disjoint links ℓ in p. This is the accumulated vehicle-meters along a path p, and should informally be proportional to accident risk (which should increase with both distance travelled and traffic) along the path. It is additive along paths as are the quantities of (1). To within fluctuations in $\mathbf{N}$ and $\mathbf{L}$, $\mathbf{VM}$ of (7) scales linearly with the number of links in a path (whereas multiplying total vehicle count along p and total length of p would lead to scaling which is approximately quadratic in the number of links). For each $(O, D) \in \mathcal{T}$, we can also compute

$$\text{APVM}_\alpha\left(O, D\right) \stackrel{\text{def}}{=} \frac{\sum_{\ell \in p_\alpha^*(O,D)} \mathbf{A}(\ell)/\mathbf{VM}(\ell)}{\sum_{\ell \in p_0^*(O,D)} \mathbf{A}(\ell)/\mathbf{VM}(\ell)}$$

which measures the accidents per vehicle-meter for $p_\alpha^*(O, D)$, normalized by that of the travel-time minimizing path $p_0^*(O, D)$. Finally,

$$\overline{\text{APVM}}_\alpha \stackrel{\text{def}}{=} \frac{1}{|\mathcal{T}|} \sum_{(O,D)\in\mathcal{T}} \text{APVM}_\alpha\left(O, D\right) \tag{8}$$

averages this normalized risk over $(O, D) \in \mathcal{T}$. Normalization of risk by $p_0^*(O, D)$ means that quantities like (8) are 1 if $\alpha = 0$.

Tradeoff Plot 2 (*APVM vs Time*). To quantify the macroscopic tradeoff between safety and time from an urban planning perspective, we parametrically plot $\{(\bar{\tau}_\alpha, \overline{\text{APVM}}_\alpha)\colon\ 0 \leqslant \alpha < 1\}$. This captures the tradeoff between time and accidents per vehicle-meter for a statistical sample of drivers.

2.3. Traffic density

In order to get a rough comparison between plots involving accidents per meter and accidents per vehicle-meter, we can also consider traffic density. For any path p, define

$$\mathbf{D}^{\circ}(p) = \frac{\sum_{\ell \in p} \mathbf{N}(\ell)}{\text{number of links in } p}. \tag{9}$$

Set

$$\mathbf{D}_{\alpha}(O, D) \overset{\text{def}}{=} \frac{\mathbf{D}^{\circ}(p_{\alpha}^{*}(O, D))}{\mathbf{D}^{\circ}(p_{0}^{*}(O, D))}, \ (O, D) \in \mathcal{T}; \quad \bar{\mathbf{D}}_{\alpha} \overset{\text{def}}{=} \frac{1}{|\mathcal{T}|} \sum_{(O,D)\in\mathcal{T}} \mathbf{D}_{\alpha}(O, D) \tag{10}$$

If in fact every road would have the same number of cars, the accidents per vehicle-meter would be proportional to accidents per meter, and

$$\overline{\text{APVM}}_{\alpha} = \frac{\bar{\sigma}_{\alpha}}{\bar{\mathbf{D}}_{\alpha}};$$

even in the absence of this equality, the average density function of (10) will allow us to make some comparisons as α varies. We also note that (9) depends on how links are defined; extraneous link structures (e.g., where a railroad crosses a road) can affect both the numerator and denominator (and thus the ratio). As our interest in $\mathbf{D}$ is to suggest comparisons between σ and APVM, we do not process the digraph to find and remove such extraneous link structures.

Tradeoff Plot 3. To quantify the macroscopic tradeoff between safety and traffic density, we parametrically plot $\{(\bar{\sigma}_{\alpha}, \bar{\mathbf{D}}_{\alpha}): 0 \leqslant \alpha < 1\}$. This quantifies the relation between accidents/meter and traffic density.

3. Manhattan: calculations, results, and discussion

3.1. The datasets (mentioned in Section 1.3)

Let's apply these ideas to the datasets listed in Section 1.3. Our model of the street structure of Manhattan Island will be that of Open Street Map. According to Open Street Map, there are 4468 nodes (intersections) and 9683 links.

The data of Donovan et al. (2016) gives us hourly estimates of the travel times along different links in New York City. To reduce computational time, we restrict the dataset by randomly selecting one day out of each workweek in 2012. The reduced dataset has a datetime in the interval

[2012-1-6 00: 00: 00, 2012-12-26 23: 00: 00]

and contains 14,243,122 total traffic speed estimates. Further restricting this data to only those links which begin and end in Manhattan, we obtain data for 5,633 (roughly 60%) of the 9,683 links in our graph. Links with no travel times are removed from the graph, so those links are not used by Dijkstra's algorithm. When computing averages as in (5), we average only over those (O, D) pairs joined by a finite cost path. In our dataset, 67,613 (roughly 50%) out of 138,631 total (O, D) pairs are not joined by a finite cost path.

There may be several realistic reasons why (Donovan et al., 2016) does not have a complete set of estimates; either there was no traffic on the link, or any "optimizing" taxi driver would have avoided it (due to, e.g., congestion). Assuming that travel times are stationary on a daily basis, we can nevertheless average across days to get an hour-of-day estimate on each link. In other words, if the dataset has estimates of the travel time of some link at 8 AM on two different days as 50.84 s and 30.0 s, we will define the travel time

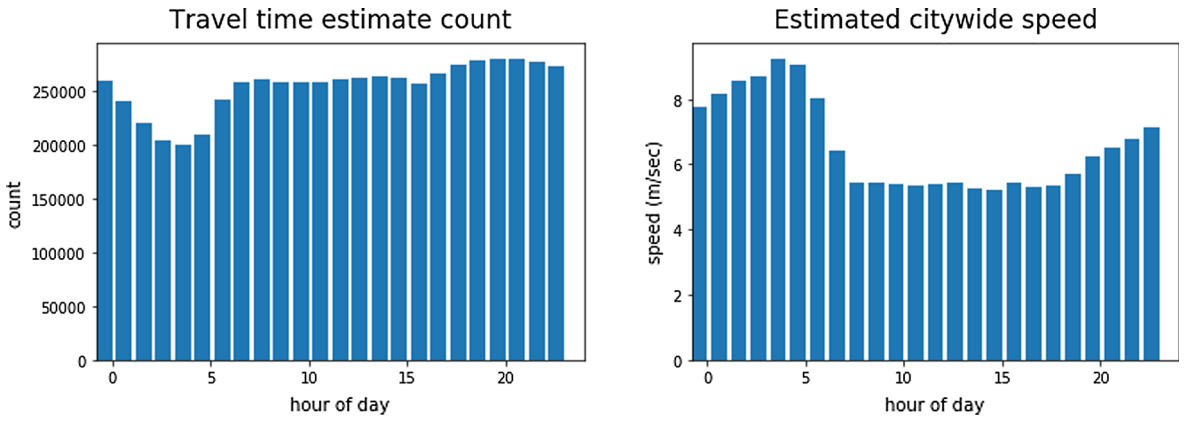

Fig. 4. Estimated travel times (left) and speeds (right).

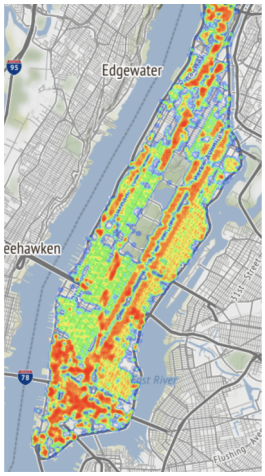

Fig. 5. Pace (slower links have a higher pace value and are in red). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

on that link to be 40.42 s.

Fig. 4 shows how the estimated travel times and speeds (i.e., link distance, as given by Open Street Map, divided by estimated time) vary throughout the hours of the day. The geographic distribution of the pace (the reciprocal of speed) is given in Fig. 5.

The accident data of NYPD (n.d.) contains 109, 093 accident reports within the time interval

$$[2014\text{-}01\text{-}01\ 00{:}\ 01{:}\ 00,\ 2017\text{-}12\text{-}29\ 23{:}\ 29{:}\ 00] \tag{11}$$

We here take all the accident counts from NYPD (n.d.); one could alternately choose accidents of only a certain severity. The location of these accidents is snapped to the nearest Open Street Map node. There are 108,658 weekday accidents in Manhattan in this file; this amounts to 3.1 accidents/hour. There were only 3,732 unique accident nodes. Fig. 6 shows the count of accidents in the dataset as a function of hour of day. Fig. 7 gives the geographic distribution of accidents. We note that works such as (Kieu et al., 2018) allow one to statistically augment datasets such as accidents in ways which preserve observed features.

To assign accidents to links (as opposed to nodes), we say that the accidents at each node are assigned to all links ending in that node (i.e., to exit a link, one has to pass through the accidents at the exit node). While this synthetically enlarges the accident count in any snapshot of the data, the routed paths will see each link and node at most once, so the effective accident count is not changed from the perspective of any driver.

The trip data of Donovan and Work (2016) contains 153,767 trips within the time interval

[2016-02-29 23: 00: 00, 2016-03-30 23: 58: 00].

There are 153,142 weekday trips (based on origin datetime) which begin and end in Manhattan; there are 138,632 unique origin-destination pairs. Since Manhattan is formally an island, we will assume that it is a geodesically (for all $\mathbf{c}_\alpha$) convex subset of the entire world traffic graph (A directed subgraph H of a directed graph G is said to be geodesically convex with respect to a cost function c if any c-shortest path in G which begins and ends in H in fact is entirely contained in H.).

In our analysis, we use taxi counts (available from Donovan et al., 2016) as an illustrative proxy for true traffic counts, but

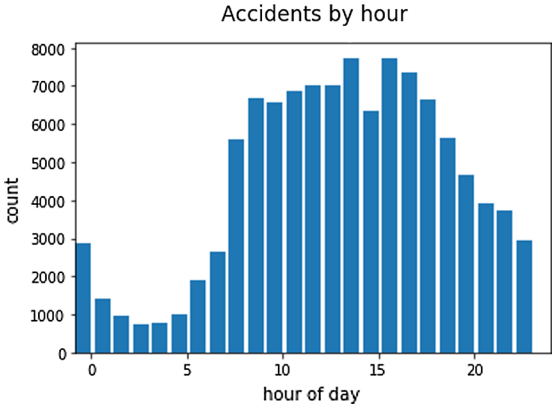

Fig. 6. Accident counts.

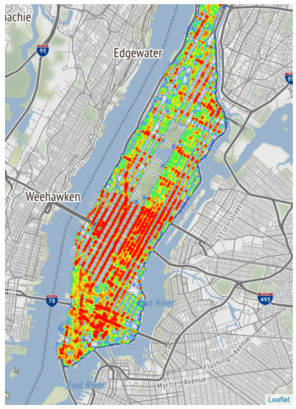

Fig. 7. Accidents (nodes in red have higher accident count). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

recognize that it is a potentially biased dataset. The function $\overline{\text{APVM}}_\alpha$ could be readily modified when non-biased counts (possibly available in private companies or from dedicated traffic counting sensors installed by municipalities) are available.

Fig. 8 contains a histogram of the duration of these trips and the distribution of trip origination by hour of day. Fig. 9 gives geographic distributions of origins and destinations.

Note that the travel time data, the accident data and the trip data are from different years; we have assumed yearly stationarity of traffic patterns.

In our parametric plots, we restrict α to the finite set

$$\mathcal{A} \overset{\text{def}}{=} \{0, 0.45, 0.61, 0.82, 0.87, 0.91, 0.94, 0.956, 0.967, 0.970, 0.978, 0.984, 0.988, 0.991, 0.996, 1\} \subset [0, 1]. \tag{12}$$

This collection of α is chosen so that the points on Fig. 10 will be relatively evenly spaced.

To show the comparative aspect of our analysis, we compare tradeoffs in normalized accidents/meter and trip time by windowing our three datasets (trips, travel times, and accidents) to morning rush hour, midday, evening rush hour, and evening/morning. See Table 1 for precise definitions. We assume that risk is stationary in time over each of these subintervals of the day, so we normalize the accident counts by the length of the relevant time interval (i.e., 2 accidents at a node during morning rush hour becomes 1 accident/hour) to approximate the accident rate (recall the definition of **A** in (1)).

3.2. Driver perspective

Plot 1 for this dataset is in Fig. 10. Normalization by the case $\alpha = 0$ in (4) means that every curve in Fig. 10 passes through $(1, 1)$, when $\alpha = 0$.

As expected, the curves are decreasing; if we are willing to increase time, the accidents/meter decrease. For most origins and destinations there is a path which provides a safer (historically fewer accidents per meter) alternative to the time minimizing route. Quantitatively, a 10% decrease in the number of accidents along a path can be achieved, on average, by only a 4% increase in travel time compared to the time minimizing path. A curve with a steeper negative slope, such as the evening and morning curve, indicates that there is a larger gain to be made in safety of a given path with relatively less given up in terms of travel time.

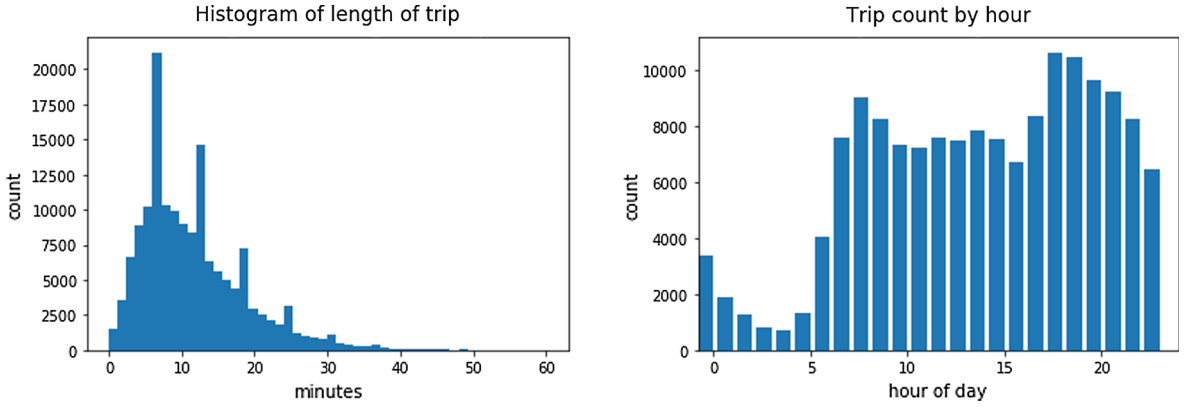

Fig. 8. Histogram of trip duration (left) and trip count by hour of trip origination (right).

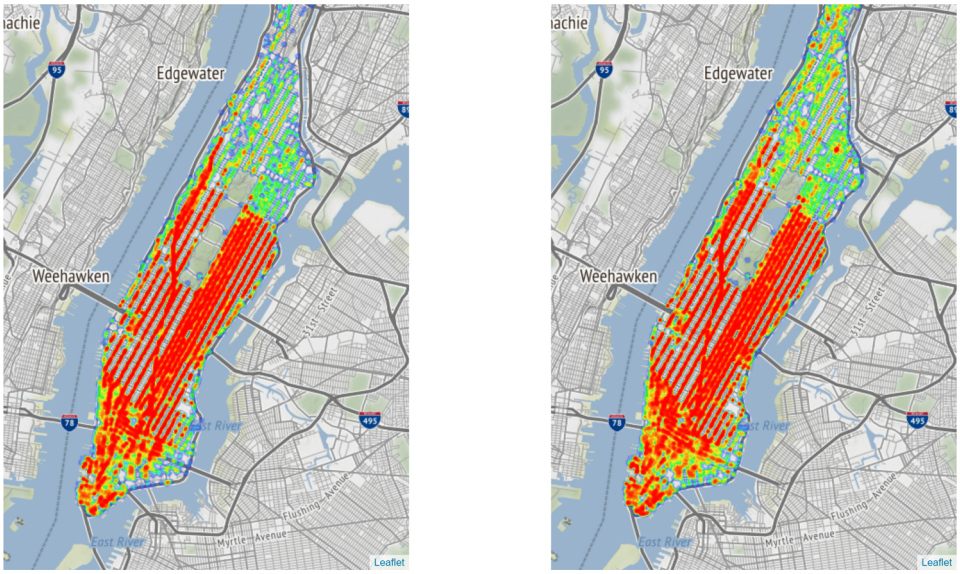

Fig. 9. Spatial distribution of origins (left) and destinations (right); nodes which are more frequent origins or destinations are in red. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

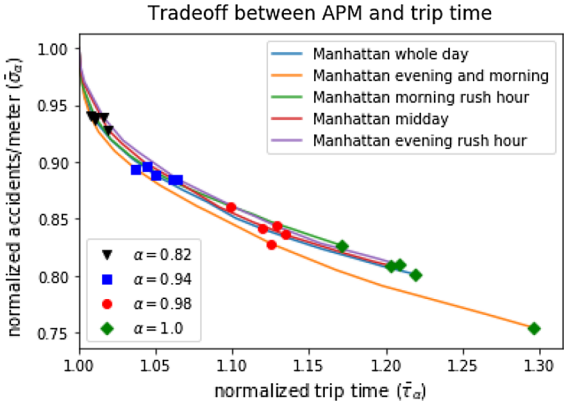

Fig. 10. Plot 1 for different time windows in Manhattan. Selected values of α are highlighted.

Table 1
Time windows.

Label	Time window
Whole day	00:00:00–23:59:59
Evening and morning	07:00:00–18:00:00
Morning rush hour	07:00:00–09:00:00
Midday	09:00:00–16:00:00
Evening rush hour	16:00:00–18:00:00

Fig. 11 gives a bit more insight into the monotonicity in Fig. 10. By definition of $p^*_\alpha(O, D)$ as the argmin of (3), we know that the trip time increases and the accident count decreases as we increase α. Thus the ratio of accidents/trip time must decrease as α increases. Fig. 11 tells us that, on average, trip length (positively) covaries with trip time (both increase together). This suggests that accidents/meter, as opposed to accidents/minute, will also decrease as α is increased.

Combining Figs. 10 and 11, we also note that safer trips in general end up being longer; as α increases, the points in Fig. 10 move down (implying safer routing), and the points in Fig. 11 move up (implying longer trips). It follows that the safer trips will be longer; this is natural as safer trips need to route around accident-prone areas (see Fig. 1).

Fig. 12 is a scatterplot of

$$\{(\bar{\tau}_1(O, D), \sigma_1(O, D)) : (O, D) \in \mathcal{T}\}. \tag{13}$$

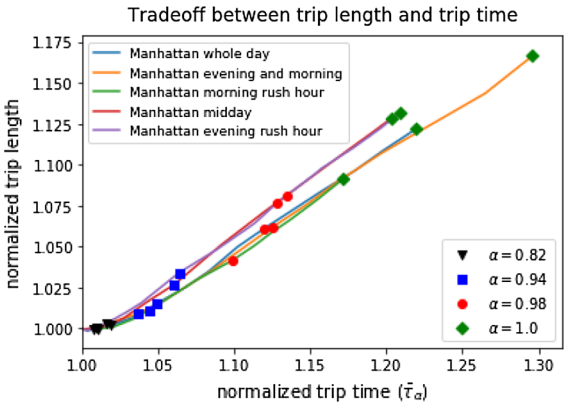

Fig. 11. Normalized trip length as a function of normalized trip time.

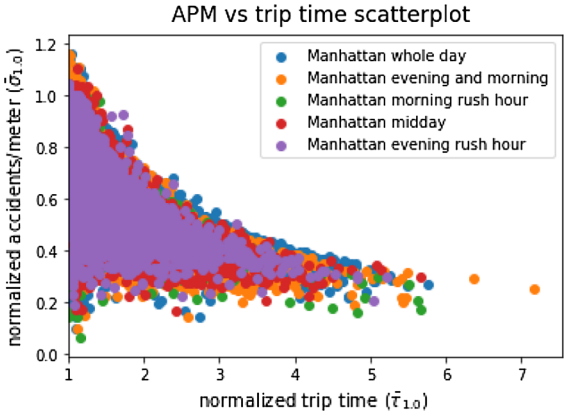

Fig. 12. Scatterplot of $(\tau_{.99}(O, D), \sigma_{.99}(O, D))$ for all $(O, D) \in \mathcal{T}$.

Each point on the plot indicates the travel time (relative to that of the quickest path) when taking the safest (least number of historical accidents/hr) path. The centroid of this scatterplot gives the corresponding point in Fig. 10 for $\alpha = 1$. This plot shows the variability in the conclusion that, for most origin-destination pairs, the $\mathbf{c}_1$-cheapest path provides a reasonably fast alternative for drivers interested in a safer route. As noted in Remark 2.2, a few points in Fig. 12 lie above the horizontal line $y = 1$, meaning that for the corresponding $(O, D) \in \mathcal{T}$, the safest path is short enough that the accidents/meter (as opposed to the accident count itself) is larger than that of the time minimizing path. These paths might be very congested as both the travel time is longer (since for large α we are de-emphasizing travel time in constructing $p^*_\alpha(O, D)$) and the path length is shorter. Assumedly, these trips would also show an exception to the generally increasing trend of Fig. 11.

3.3. City planner perspective

Plot 2 for this dataset is in Fig. 13. Again, normalization means that all curves pass through (1, 1) for $\alpha = 0$. These curves, however, are in general increasing, suggesting that the safest paths have low traffic counts (or at least taxi counts); low enough that if we normalize the points leading to Fig. 10 by traffic counts, the curve becomes increasing (even though the curves of Fig. 10 were decreasing). The decrease in traffic density along safer routes outweighs the decrease in accidents per meter. If the normalized traffic counts on the safest paths are low, the traffic counts on the time minimizing paths should be high; most travellers indeed seek the time minimizing path.

Plot 3, i.e., Fig. 14, gives more insight into the interaction of accidents/meter and traffic density. As usual, the curves are normalized to pass through (1, 1) at $\alpha = 0$. Both accidents/meter and traffic density decrease as we increase α, thus the plots in Fig. 14 carve out a curve to the left and down as α increases. The parametric plot is thus increasing, confirming that traffic densities and risk (in accidents/meter) positively covary; more accidents/meter are seen with higher traffic counts. There are more people on the quicker paths and less on the safer ones. Fig. 15 gives the same conclusion; safer trips take longer and encounter less traffic.

In fact, Fig. 14 suggests a linear relationship between the average traffic density function (10) and accidents/meter; the best fit line is

$$\text{normalized density} = 0.991\{\text{normalized accidents/meter} - 1\} + 1$$

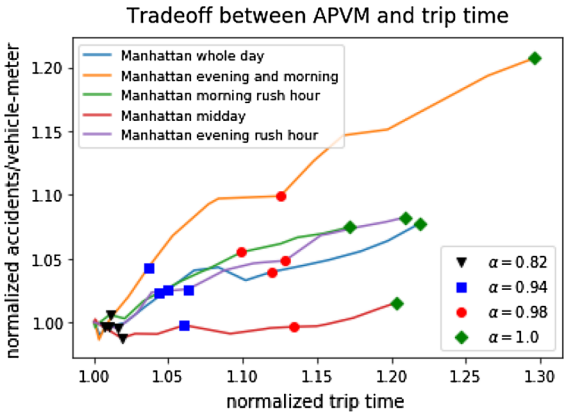

Fig. 13. Plot 2 for different time windows in Manhattan. Selected values of α are highlighted.

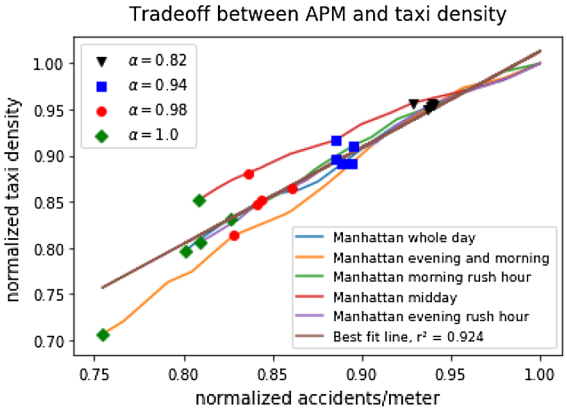

Fig. 14. Plot 3 for different time windows in Manhattan. Selected values of α are highlighted.

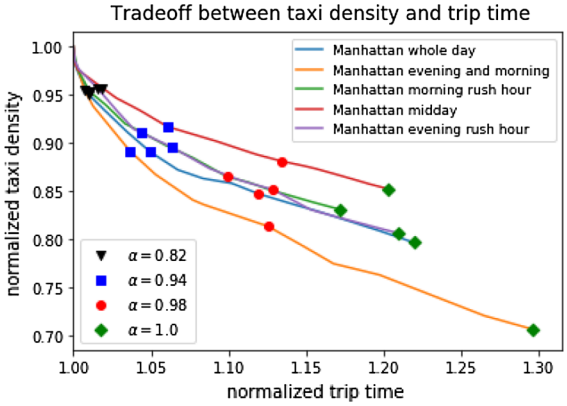

Fig. 15. Plot of $(\bar{\tau}_\alpha, \bar{\mathbf{D}}_\alpha)$ for $\alpha \in \mathcal{A}$ in each time window.

(recall that normalization requires that all curves go through (1, 1)) with correlation determination coefficient $r^2 = 0.954$. This gives us the sensitivity of normalized accidents/meter with respect to normalized taxi density;

$$\frac{\text{change in normalized accidents/meter}}{\text{change in normalized taxi density}} = \frac{1}{0.991} = 1.01. \tag{14}$$

Fig. 16. Line delineating Upper/Central Manhattan.

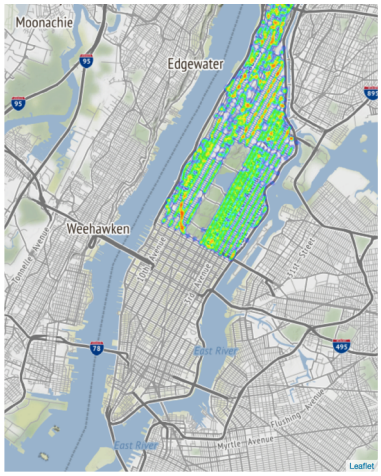

Fig. 17. Heat map of accidents in Upper/Central Manhattan during the morning rush hour.

3.4. Upper and central manhattan

The normalized nature of our analysis allows us to carry out comparative analyses on geographically different datasets. As an indication of this, consider trips which are above the lower end of Central Park—specifically trips which are contained in the part of Manhattan above the line determined by the (latitude, longitude) points (40.772619, −73.993321), (40.758382 −73.959375) (i.e., 59th St.); see Fig. 16. The space of paths contained in the subgraph corresponding to Upper/Central Manhattan (our abbreviation for Upper and Central Manhattan) is smaller than the space of paths in all of Manhattan. With this reduction in allowable paths, one might expect less of an ability to reduce accidents per meter as α varies in the Upper/Central Manhattan data. To explain this further and set up notation for this section, let G_M be the directed graph corresponding to our model of all of Manhattan, and let $G_{U/C} \subset G_M$ be the directed subgraph of vertices and edges which lie above the line in Fig. 16. Figs. 10 and 11 tell us that safer paths tend to be longer, and hence meander more than faster paths. For a given pair of vertices (O, D) in $G_{U/C}$, one might then expect that the faster paths remain in $G_{U/C}$, while the safer paths do not. For a fixed $(O, D) \in G_{U/C}$, the cheapest path in G_M may in fact be cheaper than the cheapest path restricted to $G_{U/C}$, so restricting paths to lie in $G_{U/C}$ might increase $\bar{\sigma}_1$ (Fig. 19 will support this).

About 20% (41,235 of 209,746) of the trips in our dataset start and end in Upper/Central Manhattan. In Figs. [3] 19 and 20, we plot various tradeoff curves during the morning rush hour for both Manhattan as a whole and for Upper/Central Manhattan. In Figs. 21 and 22 we plot various tradeoff curves during midday hours for Manhattan as a whole and for Upper/Central Manhattan.

[3] To save space in the plots, we refer to Upper/Central Manhattan as “Upper Manhattan”.

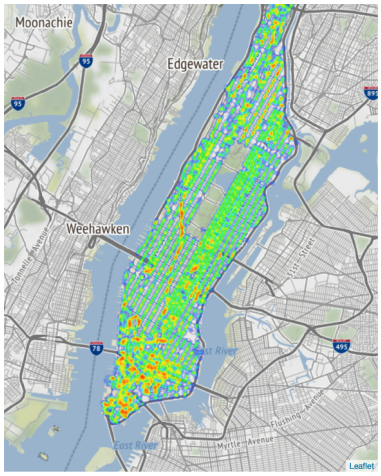

Fig. 18. Heat map of accidents in Manhattan during the morning rush hour.

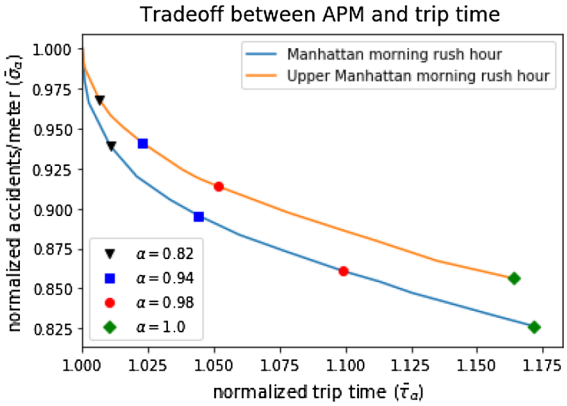

Fig. 19. Plot 1 for morning rush hour in Upper/Central Manhattan vs. all of Manhattan. Selected values of α are highlighted.

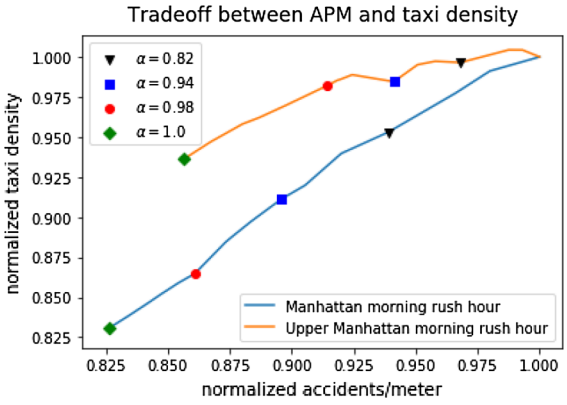

Fig. 20. Plot 3 for morning rush hour in Upper/Central Manhattan vs. all of Manhattan.

In particular, in Fig. 19, the curve for Upper/Central Manhattan is above the curve for Manhattan as a whole. For example, if one is willing to increase trip time by 6%, one can only decrease accidents/meter by about 8% in Upper/Central Manhattan, while one can decrease accidents/meter by about 11% in Manhattan as a whole. This suggests that accidents are more entwined with Upper/Central Manhattan morning traffic than with traffic in Manhattan as a whole. This is supported by Fig. 20, where the slope of the Upper/Central Manhattan curve is smaller than the slope of the overall Manhattan curve; any change in the number of taxis along a path increases the number of accidents per meter along that path in more Upper/Central Manhattan than along a typical path in Manhattan as a whole. Several potential reasons might be:

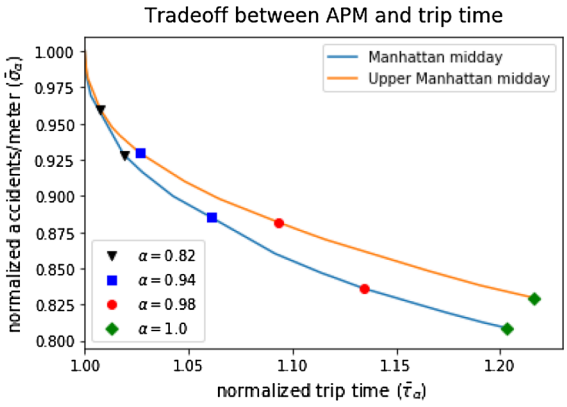

Fig. 21. Plot 1 for midday in Upper/Central Manhattan vs. all of Manhattan. Selected values of α are highlighted.

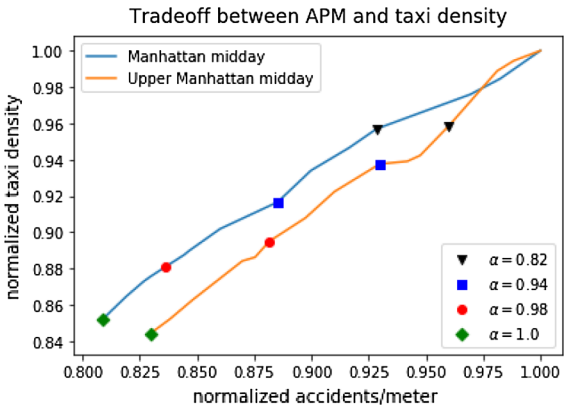

Fig. 22. Plot 3 for midday in Upper/Central Manhattan vs. all of Manhattan.

- Accidents are more evenly distributed over paths in Upper/Central Manhattan than Manhattan as a whole during the morning rush hour. This can be visually inferred from the heat maps in Figs. 17 and 18. If all paths pass a similar number of accidents, it may be difficult to find a safer route.
- Upper/Central Manhattan is not geodesically convex (as a subgraph of all of Manhattan; the definition of geodesic convexity was given in Section 3.1). For 35% of the (O, D) pairs in the graph $G_{U/C}$ of Upper/Central Manhattan, the c_1-minimizing path in $G_{U/C}$ is different than the c_1-minimizing path in the graph G_M of Manhattan as a whole. To break 35% up in detail, 13% of the (O, D) pairs have no finite cost path in $G_{U/C}$, but do have a finite cost path in G_M. The other 22% have a cheaper (with respect to c_1) path in G_M than $G_{U/C}$. Because we only average over finite cost paths when computing $\bar{\sigma}_\alpha$, these 22% of (O, D) pairs are what contribute to the difference in $\bar{\sigma}_1$ for the two curves in Fig. 19.
  On the other hand, there are no finite-cost c_0-minimizing paths in $G_{U/C}$ which are longer than their counterparts in G_M (though for 13% of the (O, D) pairs there is a finite cost path minimizing c_0 in G_M, but no such optimal path in $G_{U/C}$), confirming our intuition that c_0-minimizing paths between vertices in $G_{U/C}$ tend to remain in $G_{U/C}$.

For the Upper/Central Manhattan midday dataset, we expect the normalized tradeoff between accidents and travel time to be less than the corresponding tradeoff in Manhattan (as a whole) for the same reason as for the morning rush hour data; there are strictly fewer route choices in Upper/Central Manhattan. Fig. 22 confirms this. As α increases from 0 to 1, the curve for Upper/Central Manhattan is (end-to-end) steeper than that of Manhattan as a whole; the same overall (vertical) variation of normalized traffic density equates to a smaller change in accidents/meter in Upper/Central Manhattan than in Manhattan as a whole. One potential reason for the difference between morning and midday behavior in Upper/Central Manhattan in Figs. 20 and 22 would be an increase in route choice for the midday data in Upper/Central Manhattan. Because our travel time data has more entries in the midday time window as compared to the morning rush hour window, more links in our graph are assigned a finite cost, and hence are viable choices for Dijkstra's algorithm. Indeed, there are 2.2 times as many finite cost trips in the midday data for Upper/Central Manhattan as compared to the morning data for Upper/Central Manhattan. If this is indeed the reason, we would expect the analogue of Fig. 22 for the entire day's data to look similar. Fig. 23 supports this.

3.5. Computational considerations

We carried this work out on an NVIDIA Geforce 1070 FE processor with 8 GB GDDR5 (256-bit) on-board memory plus 1920 CUDA

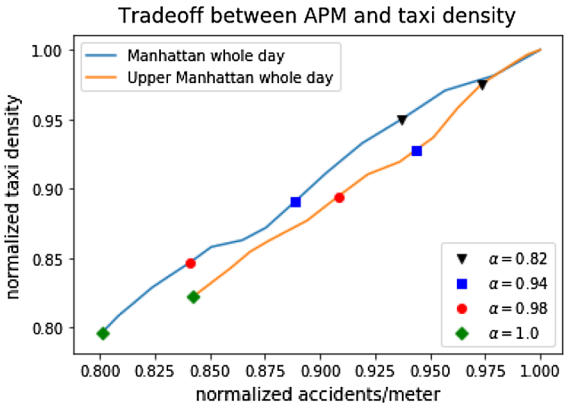

Fig. 23. Plot of $(\bar{\sigma}_\alpha, \bar{\mathbf{D}}_\alpha)$ for $\alpha \in \mathcal{A}$ in Manhattan and Upper/Central Manhattan for non time-windowed data.

cores. The traffic network was modeled as a directed graph with 4,468 nodes and 9,683 links. The computationally intensive part of the work is finding the cheapest paths (according to (2)) for each (O,D) pair (see Sabran et al., 2014).

For each origin-destination pair, we use Dijkstra's shortest path algorithm to compute a shortest path. Because the trips in any given window comprise less than 1% of the trip space (the 4468^2 possible origin destination pairs), this is much faster than using the Floyd-Warshall algorithm to compute all shortest paths in the graph then specializing to the trips in our window. More specifically, we use the Python package networkx's Dijkstra algorithm, which uses a min-priority queue implemented via a binary heap. One can achieve slightly better amortized runtime using a min-priority queue based on the Fibonacci heaps as developed by Fredman and Tarjan (1987); for graphs of our size, however, bookkeeping considerations make a binary heap faster.

As we vary α, we can parallelize the calculations by assigning the computations for each α to a different processor. We use the Python multiprocessing.Pool.map function to simultaneously compute a dictionary of cheapest paths with respect to the cost function $\mathbf{c}_\alpha$ for $\alpha \in \mathcal{A}$. Since our server has 16 processors and we compute for 16 values of α, this reduces the computation time by a factor of 16.

4. Future work

The analysis developed here could be adapted to a number of datasets. With datasets from two different cities, we could build comparative analytics (along the lines of Section 3.4). We could also investigate our various stationary assumptions. As weather patterns change, transients in driving patterns may make the various tradeoffs become more pronounced. The accident data of NYPD (n.d.) contains additional information detailing the cause of the accident; one could carry out comparative analyses for accidents of different severity.

We might also investigate Fig. 14 for other cities; is a linear relation generally applicable, and if so, what would be the sensitivity as in (14)? This would give another point of comparison.

We could also investigate why monotonicity fails, more particularly if there are any commonalities in the points above the line $y = 1$ in Fig. 12. If so, this might suggest specific situations against which to test new safety policies.

It would also be interesting to see if " best" paths become self-correcting. Drivers may avoid intersections known to be dangerous, eventually leading to less accidents along those paths (the perspective of Section 2.2). A related direction would be to carry out an analysis before and after safety policies (stop signs, traffic lights) have been implemented.

Appendix A. Bi-objective decision making

We here recall some aspects of a weighted-sum model as applied to the problem of choosing a "best" path which minimizes both travel time and accidents/meter.

For a fixed $(O, D) \in \mathcal{T}$, define the bi-objective $\mathbb{R}^2$-valued cost function

$$\mathrm{Obj}_{(O,D)}(p) \stackrel{\text{def}}{=} (\mathbf{T}(p), \mathbf{A}(p))$$

for all $p \in \mathcal{P}(O, D)$. Defining also the linear map

$$b_\alpha(t, a) \stackrel{\text{def}}{=} (1 - \alpha)t + \alpha a, \qquad (t, a) \in \mathbb{R}^2$$

from $\mathbb{R}^2 \to \mathbb{R}$, we have that

$$c_\alpha(p) = b_\alpha(\mathrm{Obj}_{(O,D)}(p)). \qquad p \in \mathcal{P}(O, D)$$

for all $(O, D) \in \mathcal{T}$ and $p \in \mathcal{P}(O, D)$. Standard results Miettinen (1998) then imply that any paths which minimize c_α over any fixed

$\mathcal{P}(O, D)$ must occur on the *Pareto optmal front* of $\text{Obj}_{(O,D)}(\mathcal{P}(O, D))$. This simplifies some aspects of our problem.

To develop these thoughts, recall the partial order on $\mathbb{R}^2$ defined by *Pareto dominance*. For (t, a) and (t', a') in $\mathbb{R}^2$, $(t, a) \preceq (t', a')$ if $t \leqslant t'$ and $a \leqslant a'$, while $(t, a) \prec (t', a')$ if $(t, a) \preceq (t', a')$ but either $t < t'$ or $a < a'$.

For $(O, D) \in \mathcal{T}$, the set

$$\text{PF}^*(O, D) \stackrel{\text{def}}{=} \{(t, a) \in \text{Obj}_{(O,D)}(\mathcal{P}(O, D)) \text{: there does not exist } (t', a') \in \text{Obj}_{(O,D)}(\mathcal{P}(O, D)) \text{ with } (t', a') \prec (t, a)\}.$$

of non-minorized points in $\text{Obj}_{(O,D)}(\mathcal{P}(O, D))$ is the Pareto optimal front of $\text{Obj}_{(O,D)}(\mathcal{P}(O, D))$; the preimage of the optimal front is the Pareto optimal set

$$\text{PS}^*\Big(O, D\Big) \stackrel{\text{def}}{=} \text{Obj}^{-1}_{(O,D)}\Big(\text{PF}^*\Big(O, D\Big)\Big)$$

Extrema of (the linear function) b_α, when restricted to $\text{Obj}_{(O,D)}(\mathcal{P}(O, D))$, must then (Miettinen, 1998) occur at

$$\text{Obj}_{(O,D)}(\mathcal{P}(O, D)) \cap \text{ext}(\overline{\text{conv}}(\text{Obj}_{(O,D)}(\mathcal{P}(O, D)))), \tag{15}$$

where, for any subset S of $\mathbb{R}^2$, $\text{conv}(S)$ denotes its convex hull, $\overline{S}$ denotes its closure, and $\text{ext}(S)$ denotes the collection of extremal points if S is convex. It follows that extrema of the composite function c_α must then occur on those parts of the Pareto optimal sets which are in the pre-image of (15).

Averaging allows us to consider all trips simultaneously; an average is minimal only when the terms in the average are minimal. Set

$$\text{Obj}(\mathbf{p}) \stackrel{\text{def}}{=} \frac{1}{|\mathcal{T}|} \sum_{(O,D)\in\mathcal{T}} \text{Obj}_{(O,D)}(p_{(O,D)})$$

for all $\mathbf{p} = (p_{(O,D)})_{(O,D)\in\mathcal{T}}$ in the set

$$\mathcal{P}^\times \stackrel{\text{def}}{=} \times_{(O,D)\in\mathcal{T}} \mathcal{P}(O, D).$$

This has Pareto optimal front

$$\text{PF}^* \stackrel{\text{def}}{=} \{(t, a) \in \text{Obj}(\mathcal{P}^\times): \text{ there does not exist } (t', a') \in \text{Obj}(\mathcal{P}^\times) \text{ with } (t', a') \prec (t, a)\}.$$

and Pareto optimal set

$$\text{PS}^* \stackrel{\text{def}}{=} \text{Obj}^{-1}(\text{PF}^*).$$

We then have that

$$\text{PS}^* = \times_{(O,D)\in\mathcal{T}} \text{PS}^*(O, D),$$

meaning that the average is indeed minimal when each path is minimal. Then $(p^*_\alpha(O, D))_{(O,D)\in\mathcal{T}}$ must be somewhere in the Obj-preimage of

$$\text{Obj}(\mathcal{P}^\times) \cap \text{ext}(\overline{\text{conv}}(\text{Obj}(\mathcal{P}^\times))),$$

The objective function Obj allows us to combine the optimizations of Section 2 over all (O, D)'s in $\mathcal{T}$ into a single calculation.

Appendix B. Parametric plot

Assume that we have two functions ϕ_1 and ϕ_2 in $C([0, 1])$, and we wish to numerically approximate the parametric plot

$$G \stackrel{\text{def}}{=} \{(\phi_1(\alpha), \phi_2(\alpha)): \alpha \in (0, 1]\}.$$

We wish to do so under the further assumption that ϕ_1 and ϕ_2 are computationally expensive to evaluate.

Let's start with a sequence

$$\mathcal{I} \stackrel{\text{def}}{=} (I_1, I_2...I_N) \tag{16}$$

of subintervals of $(0, 1]$ such that

$$I_n \stackrel{\text{def}}{=} (\alpha_{n-1}, \alpha_n]$$

for $n \in \{1, 2...N\}$, where $0 = \alpha_0 < \alpha_1 < ...\alpha_N = 1$. Define

$$\Phi^{\mathcal{I}}_i(\alpha) \stackrel{\text{def}}{=} \phi_i(\alpha_n)\frac{\alpha - \alpha_{n-1}}{\alpha_n - \alpha_{n-1}} + \phi_i(\alpha_{n-1})\frac{\alpha_n - \alpha}{\alpha_n - \alpha_{n-1}} \qquad \alpha_{n-1} < \alpha \leqslant \alpha_n$$

for each $n \in \{1, 2...N\}$; we can then set

$$G(\mathcal{I}) \overset{\text{def}}{=} \{(\Phi_1^{\mathcal{I}}(\alpha), \Phi_2^{\mathcal{I}}(\alpha)) : \alpha \in (0, 1]\}. \tag{17}$$

With intervals as in (16), we can create a new collection of intervals by splitting the one which corresponds to the largest segment of the graph of (17). Define

$$n^* \overset{\text{def}}{=} \underset{1 \leqslant n \leqslant N}{\text{argmax}} \sqrt{\sum_{i \in \{1,2\}} |\phi_i(\alpha_n) - \phi_i(\alpha_{n-1})|^2};$$

if the argmax is not unique, we randomly choose one of the elements of the set-valued argmax. We then define

$$\alpha' \overset{\text{def}}{=} \frac{\alpha_{n^*} + \alpha_{n^*-1}}{2}$$

and consider a new sequence of intervals

$$T(\mathcal{I}) \overset{\text{def}}{=} (I_1, I_2 ... I_{n^*-1}, I_{n*}^-, I_{n^*}^+, I_{n^*+1} ... I_N) \tag{18}$$

where

$$I_{n*}^- \overset{\text{def}}{=} (\alpha_{n^*-1}, \alpha'] \quad \text{and} \quad I_{n^*}^+ \overset{\text{def}}{=} (\alpha', \alpha_{n^*}].$$

Setting $\mathcal{I}^0 \overset{\text{def}}{=} \mathcal{I}$ ($\mathcal{I}$ given by (16)), we can then iteratively define $\mathcal{I}^{(d)} \overset{\text{def}}{=} T(\mathcal{I}^{(d-1)})$ for $d \in \{1, 2...\}$
We have that $G(\mathcal{I}^{(d)})$ converges to G.

Theorem 4.1. *We have that*

$$\lim_{d \to \infty} \sup_{\alpha \in (0,1]} \sqrt{\sum_{i \in \{1,2\}} |\Phi_i^{\mathcal{I}^{(d)}}(\alpha) - \phi_i(\alpha)|^2} = 0.$$

Proof. Let ω be the modulus of continuity of $\alpha \mapsto (\phi_1(\alpha), \phi_2(\alpha))$;

$$\omega(\delta) \overset{\text{def}}{=} \sup_{\substack{\alpha, \alpha' \in [0,1] \\ |\alpha - \alpha'| < \delta}} \sqrt{\sum_{i \in \{1,2\}} |\phi_i(\alpha) - \phi_i(\alpha')|^2}.$$

For intervals such as (16), define the approximate modulus of continuity

$$\Omega(\mathcal{I}) \overset{\text{def}}{=} \max_{1 \leqslant n \leqslant N-1} \sqrt{\sum_{i \in \{1,2\}} |\phi_i(\alpha_n) - \phi_i(\alpha_{n-1})|^2}.$$

The widths of the I_n's of (16) are

$$\{\alpha_1 - \alpha_0, \alpha_2 - \alpha_1, ... \alpha_N - \alpha_{N-1}\}.$$

The widths of the intervals of (18) is

$$\left\{\alpha_1 - \alpha_0, \alpha_2 - \alpha_1, ..., \alpha_{n^*-1} - \alpha_{n^*-2}, \frac{\alpha_{n^*} - \alpha_{n^*-1}}{2}, \alpha_{n^*+1} - \alpha_{n^*-1}, ... \alpha_N - \alpha_{N-1}\right\}.$$

We further have that

$$\Omega(\mathcal{I}) = \max_{1 \leqslant n \leqslant N} \sqrt{\sum_{i \in \{1,2\}} |\phi_i(\alpha_n) - \phi_i(\alpha_{n-1})|^2} = \sqrt{\sum_{i \in \{1,2\}} |\phi_i(\alpha_{n^*}) - \phi_i(\alpha_{n*-1})|^2} \leqslant \omega(\alpha_{n^*} - \alpha_{n^*-1}) \tag{19}$$

Thinking of the intervals of (16) as the first generation of parents, only one parent has a child in each generation, where (18) is the next generation after (16). The d-th generation will then have widths

$$\left\{\frac{\alpha_1 - \alpha_0}{2^{p_1(d)}}, \frac{\alpha_2 - \alpha_1}{2^{p_2(d)}}, ... \frac{\alpha_N - \alpha_{N-1}}{2^{p_N(d)}}\right\},$$

where the $p_n(d)$'s are nonnegative integers and $\sum_{n=1}^N p_n(d) = d$. Since N is fixed (the number of intervals in $\mathcal{I}^0$), we must have that

$$\lim_{d \nearrow \infty} p_{n'}(d) = \infty \tag{20}$$

for at least one $n' \in \{1...N\}$.

Assume now that

$$p_{n'}(d+1) = p_{n'}(d) + 1 \tag{21}$$

(i.e., at generation d, one of the descendants of $I_{n'}$ has a child) then, analogous to (19),

$$\Omega(\mathcal{I}^{(d)}) \leqslant \omega\left(\frac{\alpha_{n'} - \alpha_{n'-1}}{2^{p_{n'}(d)}}\right). \tag{22}$$

Let $\mathcal{D}(n')$ be the set of $d \in \mathbb{N}$ such that (21) holds. From (20), $|\mathcal{D}(n')| = \infty$ and

$$\lim_{\substack{d \nearrow \infty \\ d \in \mathcal{D}(n')}} p_{n'}(d) = \infty.$$

Thus

$$\lim_{\substack{d \nearrow \infty \\ d \in \mathcal{D}(n')}} \frac{\alpha_{n'} - \alpha_{n'-1}}{2^{p_{n'}(d)}} = 0.$$

The continuity of $\alpha \mapsto (\phi_1(\alpha), \phi_2(\alpha))$ implies that $\omega(0+) = 0$, so

$$\lim_{\substack{d \nearrow \infty \\ d \in \mathcal{D}(n')}} \omega\left(\frac{\alpha_{n'} - \alpha_{n'-1}}{2^{p_{n'}(d)}}\right) = 0.$$

Combining this with (22), we then have that

$$\lim_{\substack{d \nearrow \infty \\ d \in \mathcal{D}(n')}} \Omega(\mathcal{I}^{(d)}) = 0. \tag{23}$$

Let's now look at other parent intervals I_n, starting with the initial generation. From (23), we have that

$$\Omega(\mathcal{I}^{(d)}) < \min_{\substack{1 \leqslant n \leqslant N \\ n \neg = n'}} \left\{ \sum_{i \in \{1,2\}} |\phi_i(\alpha_n) - \phi_i(\alpha_{n-1})|^2 \right\}^{1/2},$$

for some sufficiently large generation d in $\mathcal{D}(n')$. In other words, $\Omega(\mathcal{I}^{(d)})$ is eventually smaller than the variation of the ϕ_i's between other parent endpoints. On or before that generation, one of the other I_n's must thus have had a child. Successive repetitions imply that at least one of the other I_n's must thus have an infinite number of descendents. Repeating this again, we see that all I_n's must have an infinite number of descendents. Consequently,

$$\lim_{\substack{d \nearrow \infty \\ d \in \mathcal{D}(n)}} \frac{\alpha_n - \alpha_{n-1}}{2^{p_n(d)}} = 0$$

for all $n \in \{1, 2...N\}$, and thus

$$\lim_{d \nearrow \infty} \max_{1 \leqslant n \leqslant N} \frac{\alpha_n - \alpha_{n-1}}{2^{p_n(d)}} = 0. \tag{24}$$

If $\mathcal{I}$ is as in (16) and $\alpha \in (\alpha_{n-1}, \alpha_n]$ for some $n \in \{1, 2...N\}$, then

$$\begin{aligned}
&\left\{ \sum_{i \in \{1,2\}} |\Phi_i^{\mathcal{I}}(\alpha) - \phi_i(\alpha)|^2 \right\}^{1/2} \\
&= \left\{ \sum_{i \in \{1,2\}} \left| \phi_i(\alpha_n) \tfrac{\alpha - \alpha_{n-1}}{\alpha_n - \alpha_{n-1}} + \phi_i(\alpha_{n-1}) \tfrac{\alpha_n - \alpha}{\alpha_n - \alpha_{n-1}} \right.\right. \\
&\qquad \left.\left. - \phi_i(\alpha) \tfrac{\alpha - \alpha_{n-1}}{\alpha_n - \alpha_{n-1}} - \phi_i(\alpha) \tfrac{\alpha_n - \alpha}{\alpha_n - \alpha_{n-1}} \right|^2 \right\}^{1/2} \\
&\qquad \leqslant \tfrac{\alpha - \alpha_{n-1}}{\alpha_n - \alpha_{n-1}} \sqrt{\sum_{i \in \{1,2\}} |\phi_i(\alpha_n) - \phi_i(\alpha)|^2} \\
&\qquad + \tfrac{\alpha_n - \alpha}{\alpha_n - \alpha_{n-1}} \sqrt{\sum_{i \in \{1,2\}} |\phi_i(\alpha_{n-1}) - \phi_i(\alpha)|^2} \\
&\leqslant \tfrac{\alpha - \alpha_{n-1}}{\alpha_n - \alpha_{n-1}} \omega\left(\left|\alpha_n - \alpha\right|\right) + \tfrac{\alpha_n - \alpha}{\alpha_n - \alpha_{n-1}} \omega\left(\left|\alpha - \alpha_{n-1}\right|\right) \\
&\qquad \leqslant \omega(\alpha_n - \alpha_{n-1}).
\end{aligned}$$

where we have used Minkowski's inequality and the fact that ω is nondecreasing. Thus

$$\sup_{0 < \alpha \leqslant 1} \left\{ \sum_{i \in \{1,2\}} |\Phi_i^{\mathcal{I}}(\alpha) - \phi_i(\alpha)|^2 \right\}^{1/2} \leqslant \max_{1 \leqslant n \leqslant N} \omega\left(\alpha_n - \alpha_{n-1}\right).$$

In the d-th generation, this becomes

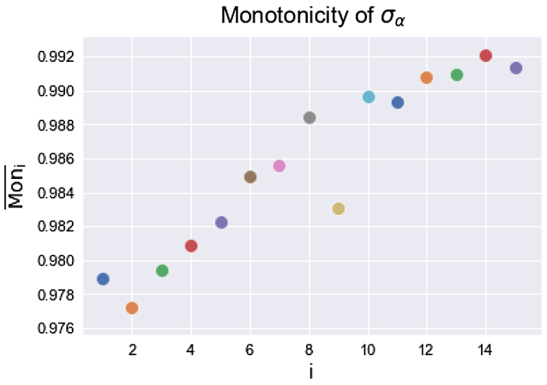

Fig. 24. Plot of $\overline{\mathbf{Mon}}_i$ as a function of i to demonstrate monotonicity on average of $\sigma_\alpha(O, D)$.

$$\sup_{0<\alpha\leqslant 1}\left\{\sum_{i\in\{1,2\}}|\Phi_i^{\mathcal{I}^{(d)}}(\alpha)-\phi_i(\alpha)|^2\right\}^{1/2}\leqslant\max_{1\leqslant n\leqslant N}\omega\left(\frac{|\alpha_n-\alpha_{n-1}|}{2^{p_n(d)}}\right).$$

The claim then follows from (24). □

To terminate the iterative algorithm in finite time with some intended accuracy, let's fix a threshold τ and say that the graph (17) is τ-approximate if

$$\max_{1\leqslant n\leqslant N}\left\{\sum_{i\in\{1,2\}}|\phi_i(\alpha_i(\alpha_n)-\phi_i(\alpha_{n-1})|^2\right\}^{1/2}\leqslant\tau\left\{\sum_{n=1}^{N}|\phi_i(\alpha_i(\alpha_n)-\phi_i(\alpha_{n-1})|^2\right\}^{1/2};$$

i.e., if the distance between succesive vertices in (17) is τ (percent) of the length of (17). We terminate the iterative procedure of the $\mathcal{I}^{(d)}$'s when $\mathcal{I}^{(d)}$ is τ-approximate.

To approximate persistence as defined in 2.1, we find a 0.05-approximate collection $\mathcal{I}$ of α, and use $\mathcal{I}$ to find a piecewise linear approximation of

$$\{(\bar{\tau}(\alpha),\bar{\sigma}(\alpha)):\alpha\in[0,1]\}.$$

Using this piecewise linear approximation to compute approximate arclength, we can give upper and lower bounds on the persistence of a given α.

Appendix C. Monotonicity of accidents per meter

As mentioned in Section 2.1, the function $\alpha\mapsto\sigma_\alpha(O,D)$ of (4) is not theoretically guaranteed to be decreasing for any given $(O,D)\in\mathcal{T}$. For a fixed $(O,D)\in\mathcal{T}$, we are guaranteed that $\alpha\mapsto\mathbf{T}(p_\alpha^*(O,D))$ is increasing and $\alpha\mapsto\mathbf{A}(p_\alpha^*(O,D))$ is decreasing. We are of course thus assured that $\alpha\mapsto\mathbf{A}(p_\alpha^*(O,D))/\mathbf{T}(p_\alpha^*(O,D))$ is decreasing. If $\mathbf{L}(p_\alpha^*(O,D))$ increases with $\mathbf{T}(p_\alpha^*(O,D))$, then $\alpha\mapsto\sigma_\alpha(O,D)$ must also be decreasing. Fig. 11 empirically does show that $\mathbf{T}$ and $\mathbf{L}$ both vary in the same direction. We also note that Fig. 10 suggests that $\alpha\mapsto\bar{\sigma}_\alpha$ is decreasing; as α is parametrically increased along each curve, the points fall lower and lower in the plot. To more directly study monotonicity of $\alpha\mapsto\sigma_\alpha(O,D)$, enumerate $\mathcal{A}$ of (12) in increasing order. We then define

$$\mathbf{Mon}_i\left(O,D\right)\overset{\text{def}}{=}\frac{\sigma_{\alpha_i}(O,D)}{\sigma_{\alpha_{i-1}}(O,D)}$$

for each (O,D) pair and then define

$$\overline{\mathbf{Mon}}_i\overset{\text{def}}{=}\frac{1}{|\mathcal{T}|}\sum_{(O,D)\in\mathcal{T}}\mathbf{Mon}_i\left(O,D\right)=\frac{1}{|\mathcal{T}|}\sum_{(O,D)\in\mathcal{T}}\frac{\sigma_{\alpha_i}(O,D)}{\sigma_{\alpha_{i-1}}(O,D)}$$

If $\mathbf{Mon}_i(O,D)\leqslant 1$ for a given trip (O,D), then $\alpha\mapsto\sigma_\alpha(O,D)$ is nonincreasing on $\{\alpha_{i-1},\alpha_i\}$. If $\overline{\mathbf{Mon}}_i\leqslant 1$, then $\mathbf{Mon}_i(O,D)\leqslant 1$ on average. Fig. 24 shows that $\max_{1\leqslant i\leqslant 15}\overline{\mathbf{Mon}}_i\leqslant 1$ and thus $\alpha\mapsto\sigma_\alpha(O,D)$ should on average be decreasing for each $(O,D)\in\mathcal{T}$.

Appendix D. Supplementary material

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.trc.2019.09.020.

References

Alemi, F., Circella, G., Mokhtarian, P., Handy, S., 2019. What drives the use of ridehailing in California? Ordered probit models of the usage frequency of uber and lyft. Transp. Res. Part C: Emerg. Technol. 102, 233–248. URL. http://www.sciencedirect.com/science/article/pii/S0968090X18318849.

Belin, M.-Å., Tillgren, P., Vedung, E., 2012. Vision zero–a road safety policy innovation. Int. J. Injury Control Safety Promot. 19 (2), 171–179.

Benouaret, I., Lenne, D., 2016. A composite recommendation system for planning tourist visits. In: 2016 IEEE/WIC/ACM International Conference on Web Intelligence (WI), pp. 626–631.

Berjani, B., Strufe, T., 2011. A recommendation system for spots in location-based online social networks. In: Proceedings of the 4th Workshop on Social Network Systems', SNS '11. ACM, New York, NY, USA, pp. 4:1–4:6. https://doi.org/10.1145/1989656.1989660.

Boriboonsomsin, K., Barth, M.J., Zhu, W., Vu, A., 2012. Eco-routing navigation system based on multisource historical and real-time traffic information. IEEE Trans. Intell. Transp. Syst. 13 (4), 1694–1704.

Bruglieri, M., Mancini, S., Pisacane, O., 2019. More efficient formulations and valid inequalities for the green vehicle routing problem. Transp. Res. Part C: Emerg. Technol. 105, 283–296.

Caliendo, C., Guida, M., Parisi, A., 2007. A crash-prediction model for multilane roads. Accident Anal. Prevent. 39 (4), 657–670.

Chen, Y., Chin, Y., 1990. The quickest path problem. Comput. Oper. Res. 17 (2), 153–161.

Dai, J., Yang, B., Guo, C., Ding, Z., 2015. Personalized route recommendation using big trajectory data. In: Data Engineering (ICDE), 2015 IEEE 31st International Conference on. IEEE, pp. 543–554.

Dai, J., Yang, B., Guo, C., Jensen, C.S., Hu, J., 2016. Path cost distribution estimation using trajectory data. Proc. VLDB Endowment 10 (3), 85–96.

Dhaou, I.B., 2011. Fuel estimation model for eco-driving and eco-routing. In: Intelligent Vehicles Symposium (IV), 2011 IEEE. IEEE, pp. 37–42.

Dong, Y., Wang, S., Li, L., Zhang, Z., 2018. An empirical study on travel patterns of internet based ride-sharing. Transp. Res. Part C: Emerg. Technol. 86, 1–22. URL. http://www.sciencedirect.com/science/article/pii/S0968090X17302954.

Donovan, B., Mori, A., Agrawal, N., Meng, Y., Lee, J., Work, D., 2016a. New York City hourly traffic estimates (2010–2013). https://doi.org/10.13012/B2IDB-4900670_V1.

Fredman, M.L., Tarjan, R.E., 1987. Fibonacci heaps and their uses in improved network optimization algorithms. J. ACM 34 (3), 596–615. URL. http://doi.acm.org/10.1145/28869.28874.

Fu, K., Lu, Y.-C., Lu, C.-T., 2014. Treads: a safe route recommender using social media mining and text summarization. In: Proceedings of the 22Nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems', SIGSPATIAL '14. ACM, New York, NY, USA, pp. 557–560. URL: http://doi.acm.org/10.1145/2666310.2666368.

Galbrun, E., Pelechrinis, K., Terzi, E., 2016. Urban navigation beyond shortest route: the case of safe paths. Inform. Syst. 57, 160–171. URL. http://www.sciencedirect.com/science/article/pii/S0306437915001854.

Gionis, A., Lappas, T., Pelechrinis, K., Terzi, E., 2014. Customized tour recommendations in urban areas. In: Proceedings of the 7th ACM International Conference on Web Search and Data Mining', WSDM '14. ACM, New York, NY, USA, pp. 313–322. URL: http://doi.acm.org/10.1145/2556195.2559893.

Guo, C., Yang, B., Andersen, O., Jensen, C.S., Torp, K., 2015a. Ecomark 2.0: empowering eco-routing with vehicular environmental models and actual vehicle fuel consumption data. GeoInformatica 19 (3), 567–599.

Guo, C., Yang, B., Andersen, O., Jensen, C.S., Torp, K., 2015b. Ecosky: Reducing vehicular environmental impact through eco-routing. In: Data Engineering (ICDE), 2015 IEEE 31st International Conference on. IEEE, pp. 1412–1415.

Donovan, Brian, Work, D., 2016b. New York City taxi trip data (2010–2013). https://doi.org/10.13012/J8PN93H8.

Haq, G., 2006. Vision zero: Adopting a target of zero for road traffic fatalities and serious injuries.

Johansson, R., 2009. Vision zero–implementing a policy for traffic safety. Saf. Sci. 47 (6), 826–831.

Kieu, T., Yang, B., Jensen, C.S., 2018. Outlier detection for multidimensional time series using deep neural networks, in '2018 19th. In: 2018 19th IEEE International Conference on Mobile Data Management (MDM). IEEE, pp. 125–134.

Kumar, H., Singh, M.K., Gupta, M., Madaan, J., 2018. Moving towards smart cities: Solutions that lead to the smart city transformation framework. Technological Forecasting and Social Change. URL: http://www.sciencedirect.com/science/article/pii/S004016251731394X.

Lee, C., Hellinga, B., Saccomanno, F., 2003. Real-time crash prediction model for application to crash prevention in freeway traffic. Transp. Res. Rec. J. Transp. Res. Board 1840, 67–77.

Liu, J., Han, K., Chen, X.M., Ong, G.P., 2019. Spatial-temporal inference of urban traffic emissions based on taxi trajectories and multi-source urban data. Transp. Res. Part C: Emerg. Technol. 106, 145–165.

Lord, D., Persaud, B., 2000. Accident prediction models with and without trend: application of the generalized estimating equations procedure. Transp. Res. Rec. J. Transp. Res. Board 1717, 102–108.

Massaro, E., Ahn, C., Ratti, C., Santi, P., Stahlmann, R., Lamprecht, A., Roehder, M., Huber, M., 2016. The car as an ambient sensing platform [point of view]. Proc. IEEE 105 (1), 3–7.

Miettinen, K., 1998. Nonlinear Multiobjective Optimization. Springer, US, Boston, MA.

Minett, C.F., Salomons, A.M., Daamen, W., Van Arem, B., Kuijpers, S., 2011. Eco-routing: comparing the fuel consumption of different routes between an origin and destination using field test speed profiles and synthetic speed profiles. In: Integrated and Sustainable Transportation System (FISTS), 2011 IEEE Forum on. IEEE, pp. 32–39.

Nie, Y.M., Li, Q., 2013. An eco-routing model considering microscopic vehicle operating conditions. Transp. Res. Part B: Methodol. 55, 154–170.

Noulas, A., Scellato, S., Lathia, N., Mascolo, C., 2012. A random walk around the city: New venue recommendation in location-based social networks. In: 2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing, pp. 144–153.

NYPD, P.D., n.d. Nypd motor vehicle collisions. https://data.cityofnewyork.us/Public-Safety/NYPD-Motor-Vehicle-Collisions/h9gi-nx95.

O'Keeffe, K.P., Anjomshoaa, A., Strogatz, S.H., Santi, P., Ratti, C., 2019. Quantifying the sensing power of vehicle fleets. Proc. Nat. Acad. Sci. 116 (26), 12752–12757.

Rakha, H.A., Ahn, K., Moran, K., 2012. Integration framework for modeling eco-routing strategies: logic and preliminary results. Int. J. Transp. Sci. Technol. 1 (3), 259–274.

Sabran, G., Samaranayake, S., Bayen, A., 2014. Precomputation techniques for the stochastic on-time arrival problem, in '2014. In: Proceedings of the Sixteenth Workshop on Algorithm Engineering and Experiments (ALENEX). SIAM, pp. 138–146.

Tingvall, C., Haworth, N., 2000. Vision zero: an ethical approach to safety and mobility. In: 6th ITE International Conference Road Safety & Traffic Enforcement: Beyond', vol. 1999, pp. 6–7.

United Nations-World Urbanization Prospects: The 2011 revision, 2011. http://www.un.org/en/development/desa/population/publications/pdf/urbanization/WUP2011_Report.pdf.

Vansteenwegen, P., Souffriau, W., Oudheusden, D.V., 2011. The orienteering problem: a survey. Eur. J. Oper. Res. 209, 1–10.

Wang, J.Y., Dirks, K.N., Ehrgott, M., Pearce, J., Cheung, A.K., 2018. Supporting healthy route choice for commuter cyclists: the trade-off between travel time and pollutant dose. Oper. Res. Health Care.

Yang, B., Dai, J., Guo, C., Jensen, C.S., Hu, J., 2018. Pace: a path-centric paradigm for stochastic path finding. VLDB J. Int. J. Very Large Data Bases 27 (2), 153–178.

Yang, B., Guo, C., Jensen, C.S., Kaul, M., Shang, S., 2014. Stochastic skyline route planning under time-varying uncertainty, in '2014. In: 2014 IEEE 30th International Conference on Data Engineering (ICDE). IEEE, pp. 136–147.

Yang, B., Guo, C., Ma, Y., Jensen, C.S., 2015. Toward personalized, context-aware routing. VLDB J. Int. J. Very Large Data Bases 24 (2), 297–318.

Yang, B., Kaul, M., Jensen, C.S., 2014. Using incomplete information for complete weight annotation of road networks. IEEE Trans. Knowl. Data Eng. 26 (5), 1267–1279.

Zero, L., Bersani, C., Paolucci, M., Sacile, R., 2017. Multi-objective shortest path problem with deterministic and fuzzy cost functions applied to hazmat transportation on a road network, in '2017 5th. In: IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)', pp. 238–243.

Zografos, K., Davis, C., 1989. Multi-objective programming approach for routing hazardous materials. J. Transp. Eng. 115.