Risk-Aware Planning: Methods and Case Study on Safe Driving Routes

John Krumm and Eric Horvitz
Microsoft Research
Redmond, WA USA
jckrumm@microsoft.com, horvitz@microsoft.com

Abstract

Vehicle crashes account for over one million fatalities and many more millions of injuries annually worldwide. Some roads are safer than others, so driving routes optimized for safety may reduce the number of crashes. We have developed a method to estimate the probability of a crash on any road as a function of the traffic volume, road characteristics, and environmental conditions. We trained a regression model to estimate traffic volume and a binary classifier to estimate crash probability on road segments. Modeling a route's crash probability as a series of Bernoulli trials, we employ the Dijkstra routing algorithm to compute the safest route between two locations. We find that, compared to the fastest route, the safest route is approximately 1.7 times as long in duration and about half as dangerous. We also show how to smoothly trade off safety for travel time, and demonstrate how drivers could be offered several route options, each with different crash probabilities and durations.

Introduction

Annual deaths worldwide due to vehicle crashes are estimated at over 1.2 million, along with 20-50 million non-fatal injuries (WHO, 2009). The World Health Organization predicts that road crashes will rise to the fifth leading cause of death by 2030, at about 3.6% of the total (WHO, 2009).

The U.S. National Highway Traffic Safety Administration blames outside conditions, along with drivers and vehicles, as one of the three broad causes of vehicle crashes (NHTSA, 2016). Choosing less crash-prone roads may reduce accident rates, even for careful drivers. However, very little work to date has aimed at finding driving routes that promise to reduce the probability of a vehicle crash. As an example, Figure 1 shows three computed routes. The green route is the fastest, identified by minimizing the driving time. The red route is the safest, according to our analysis, computed by minimizing the crash probability. The black route is a compromise between the two extremes.

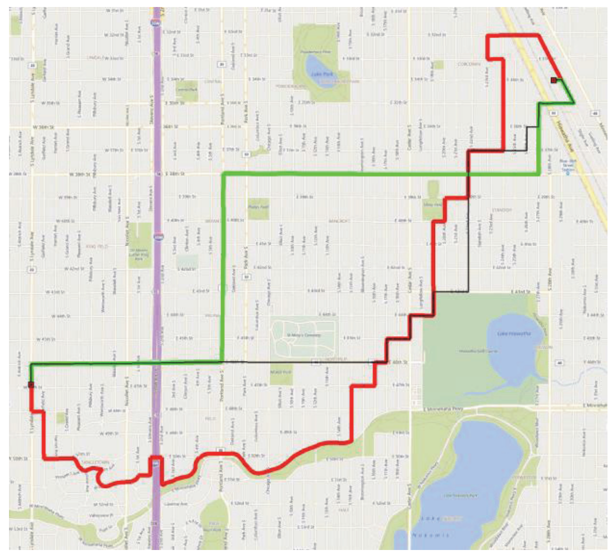

Figure 1: Green path is the fastest route between the two locations, while red route is the safest. Red route takes 1.5 times as long to drive as green route, but probability of crashing on red route is half that of green route. Black route is a compromise that is 30 seconds longer and 34% safer than fastest route.

We couple route planning and predictive models to identify the safety of routes. At the core of the approach, we seek to estimate vehicle counts and crash probabilities on individual road segments using publicly available data from road sensors and crash reports. We develop a regression model to interpolate from spatially sparse vehicle count measurements to estimate counts on all roads. These vehicle counts, along with other features, serve as inputs to a crash classifier that provides estimates of crash risk on any road. We use a Bernoulli probability model to compute the crash probability of any route, and we show how to use Dijkstra's algorithm to compute the route that minimizes the crash probability. We test our algorithm on 100 sample routes, showing how the fastest and safest routes compare in terms of driving time and crash probability. Finally, we demonstrate how to trade off safety and driving time, giving several different route options with different durations and crash probabilities.

Related to our work are studies of specific road configurations and how they affect crash risk, such as traffic signal timing (Retting, Chapline, & Williams, 2002), roundabouts (Daniels, Brijs, Nuyts, & Wets, 2010), and other conditions (Othman, Thomson, & Lannér, 2009). Other related work includes efforts to understand alternate routing criteria such as risks for transporting hazardous materials (Zografos & Davis, 1989) and reducing fuel consumption (Ericsson, Larsson, & Brundell-Freij, 2006). Nadi and Delavar introduced a route planner that simultaneously attempts to optimize for ten different criteria, although safety is not one of them (Nadi & Delavar, 2011). Thus, previous studies include methods for estimating crash risk and on considering alternate routing criteria, but not on routing based on estimated crash risk.

Crash and Traffic Data

We used datasets on crashes and on vehicle counts to build a predictive model for the per car probability of a crash along route segments. We obtained both crash and vehicle count data from the Traffic Data Management System of the City of Minneapolis, MN, USA (Minneapolis). The dataset includes 15,401 vehicle crashes over 30 months, spanning January 1, 2013 to June 30, 2015. Each crash report includes a date/time and latitude/longitude. As the reporting entity is the city of Minneapolis, the data does not include crashes on federal highways passing through the city, and thus we ignored federal highways in the routes we consider.

The vehicle count dataset includes hourly vehicle count data for 939 different roads in Minneapolis over years 2012-2015. Significantly, the data does not come with detailed time stamps for the counts. Instead, the counts are reported as the number of vehicles traversing the road for a given hour on a given day of the week and year. A map of the reported crashes and vehicles counts is shown in Figure 2.

Finally, we leverage a routable road network drawn from the Bing Maps U.S. database. Connectivity is represented as a directed graph, with intersections as vertices and road segments as edges between the vertices. We matched each crash and traffic count to the geographically nearest road.

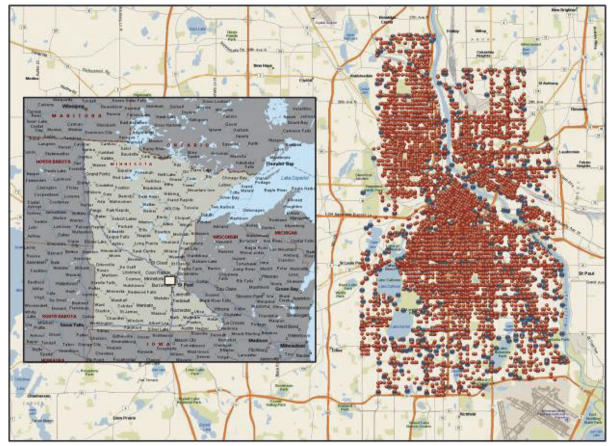

Figure 2: Red dots show locations of 15,401 vehicle crashes over 30 months in Minneapolis, MN. Black dots show 939 vehicle count measurement locations over 4 years.

Learning Crash Risk

We estimate the crash probability of road segments by first interpolating to infer the hourly vehicle count and then classifying to infer the crash probability.

Inferring Traffic Counts

For a given date/time on a road segment, we must estimate the hourly vehicle count. This is first used as a feature to estimate crash risk, where it proved to be the most important feature among those we used. The estimate is also used in the arithmetic of computing crash probabilities: if a road segment will host one crash and N vehicles in an hour, then each vehicle on that road segment has a $1/N$ chance of crashing over that hour.

When we seek a traffic count estimate on a given road segment for a given date/time, sometimes that road segment has an actual measurement. In this case, our traffic count query first looks for a count taken in the year nearest the given date/time. Given this, it then looks for a count in that year taken on the nearest day of week. (Recall that our traffic count data does not come with absolute dates, just a year and a day of week.) With this nearest year and day of week, we then look up the traffic count for the desired hour.

For identifying safe driving routes across an entire metropolitan area, we need to do inferences on all roads of the city. Thus, we need to estimate the hourly traffic counts on roads that do not have measurements. We make this estimation via interpolation from traffic counts from roads that are nearby in space and time. We perform regression where the dependent variable is the desired traffic count, and the independent variables are features including nearby traffic counts. We consider 65 independent variables in the regression. Five of the independent variables characterize the date/time and the road segment in question: hour of day, day of week, road type from {major road, arterial, street}, number of lanes, and speed limit. We also include independent variables from nearby roads that have measured vehicle counts. Specifically, we looked for the nearest five measured roads for each road type in {major road, arterial, street}, giving data on 15 total nearby roads. For each of these 15, the independent variables are:

- Hourly vehicle count per the vehicle count estimation described above
- Straight line distance from road segment in question
- Difference in number of lanes from road segment in question
- Difference in speed limit from road segment in question

In summary, we use 65 independent variables to estimate the hourly traffic count on the road segment in question for the specified date and time.

We employed as the regression function a boosted forest of decision trees, as described in (Friedman, 2001). We swept through parameters for the learning rate and forest parameters to find the most accurate function in terms of l_1 error. In particular, the optimal forest consisted of 500 trees with 74 leaves per tree and a minimum of 10 instances in each leaf.

We tested our estimates by using our learned regression function to estimate traffic counts on held-out roads where we had actual traffic count measurements. Specifically, for each measured road segment, we extracted the actual vehicle count for each hour from our data. Using ten-fold cross validation, we recorded regression estimates for each measured instance. The results are shown in Figure 3, which shows a good correlation between measured and estimated vehicle counts, with a linear least squares fit of R^2=0.9772.

With this regression function in place, we can compute a vehicle count estimate for any road segment in the region for any given date/time.

Predicting Crashes

Based on our crash data, we built a binary classifier that identifies the risk of a crash. We trained the classifier with positive examples from our crash data and negative examples generated randomly and uniformly in space and time. Each example is an hour-long instance on a road segment that either did or did not host a crash. The positive examples are simply the 15,401 crashes in our data, with the date/time truncated to the previous integer hour and the location represented as the nearest road segment. It was exceedingly rare to have more than one crash on the same road segment during the same hour, so all our crash examples implicitly represent a single reported crash incident. We generated an equal number of negative crash examples in hour-long intervals and road segments that were not

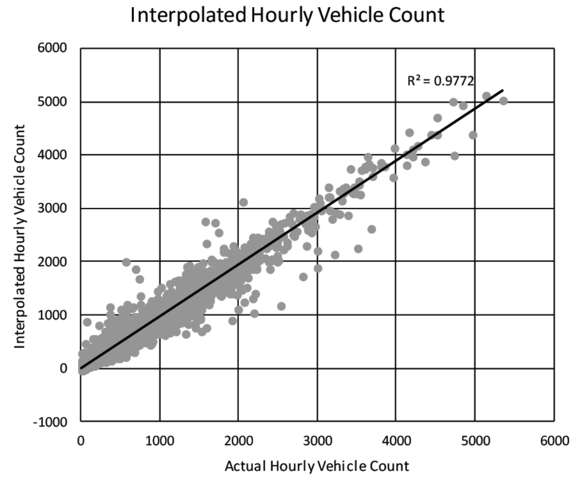

Figure 3: Estimated vehicle counts vs. actual vehicle counts.

reported as crash occurrences. As some crashes may go unreported, we assume that there may have been false negatives in the data.

The crash classifier is based on a set of 29 features from each example, including:

Vehicle Count: vehicle count estimate

Date/Time: day of week, hour of day

Road Segment Layout: road type from {major road, arterial, street}, number of lanes, speed limit, divider (binary), length of road segment, mean slope of road segment, minimum, maximum & mean radii of curvature

Sun Angle: sun above or below horizon (binary), elevation, azimuth, subtended angles between azimuth and road heading (sum to 180^o)

Weather: temperature, wind speed, snow depth, visibility (miles), cloud ceiling (feet), precipitation in last hour, last 6 hours, and last 24 hours

Nearby Structures: number and density of residences along road segment, number and density of businesses along road segment

Most of these features are self-explanatory. The road segment layout features, residences, and businesses were drawn from the Bing Maps road database. The sun angle features were motivated by reports of the role of glare in car accidents (Choi & Singh, 2005; Mitri & Washington, 2012). The feature definitions came from equations presented in (Honsberg & Bowden, 2016). The features involving the sun's azimuth angle capture the propensity for sun glare in drivers' eyes. We took our weather features for the Minneapolis, MN area from the National Oceanic and Atmospheric Administration which gives recorded weather conditions in approximately one-hour increments (NOAA, 2016).

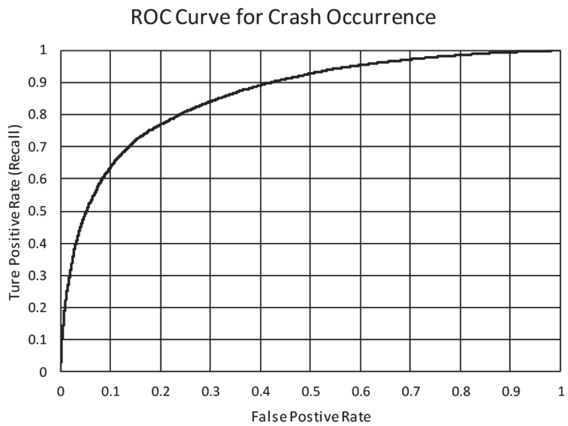

Figure 4: Receiver operating characteristic curve for classifying crash occurrences.

Traffic count estimate	1.00
Length of road segment	0.74
Mean slope of road segment	0.71
Number of businesses along road segment	0.64
Number of residences per meter along road segment	0.60
Sun elevation angle	0.53
Number of businesses per meter along road segment	0.50
Temperature	0.49
Sun azimuth angle	0.48
Number of residences along road segment	0.46
Road to sun azimuth angle min	0.38
Minimum radius of curvature	0.37
Wind speed	0.37
Cloud ceiling (feet)	0.35
Day of week	0.34
Road to sun azimuth angle max	0.33
Hour of day	0.32
Speed limit	0.31
Maximum radius of curvature	0.26
Number of lanes	0.26
Visibility (miles)	0.24
Road classification	0.23
Mean radius of curvature	0.23
Precipitation in last hour (inches)	0.15
Snow depth (inches)	0.14
Precipitation in last 24 hours (inches)	0.08
Precipitation in last 6 hours (inches)	0.08
Divider	0.08
Sun above or below horizon (binary)	0.07

Figure 5: Relative importance of features in anticipating a crash.

Given the features, we predicted the instance as either a crash or not. As with the regression function, we used a boosted forest of decision trees, but for classification instead of regression. We swept through different parameters with 10-fold cross validation, settling on a forest of 500 trees, 296 leaves per tree, and a minimum of 1 instance in each leaf.

Ten-fold cross validation gave an overall classification accuracy of 78.4%. Positive precision and recall were 79% and 77.5%, respectively, and negative precision and recall were 77.9% and 79.4%. The receiver-operator characteristic (ROC) curve is shown in Figure 4, and the area under this curve is 0.863. The classifier can provide insights about the conditions under which a crash is likely to occur.

The importance of a feature can be measured by how often it is used in the forest of decision trees. The relative importance of the 29 crash classification features is shown in Figure 5. Notably, the vehicle count estimate is the most important feature. This is despite the fact that this is an estimate from an imperfect interpolation. The next two most used features are the length and mean slope of the road segment. These are followed by two "Nearby Structures" features that indicate the number of businesses and density of residences along the road. The sixth most important feature is the sun's elevation angle. The minimum radius of curvature (maximum turn sharpness) of the road is in the top half of features. Both (Othman et al., 2009) and (Fink & Krammes, 1995) showed the importance of curvature in assessing crash risk.

Computing Individual Crash Probability

We used the crash occurrence classifier to estimate the relative risk of a crash occurring on each road segment for a given hour. We now describe how we compute the crash probability for an individual traversing a road segment and then how we compute the probability of a crash over a trip traversing multiple road segments.

Crash Probability on Road Segment

We seek to compute the probability p_i of a single vehicle crashing on a road segment i over a given hour. From the crash occurrence classifier above, we infer a class probability c_i that predicts crash risk given environmental features over the hour in question. This class probability pertains to the occurrence of a crash among all the vehicles on the road segment over the hour. In that hour, there are v_i vehicles traversing the road segment, where v_i comes from the road count estimate as described previously. With p_i and v_i for the relevant hour, the class probability for any individual vehicle crash in that hour is $c_i' = c_i,/v_i$.

Neither the class probability c_i, nor the individual class probability c_i', are calibrated to consider the overall expected number of crashes in the entire region, *i.e.* the city of Minneapolis in our case. Both have an unrealistically inflated view of crash occurrences, because the classifier was trained on an equal number of positive and negative crash examples. In actuality, there are far fewer positive instances of crashes than negative instances. Thus we must calibrate these class probabilities by scaling such that the expected number of crashes over all the region's road segments is the same as the historical number of crashes over the given hour. The scale factor is ρ, and the calibrated probability of a single vehicle crashing is $p_i = \rho c_i'$. If A is the random variable representing the total number of hourly crashes in the region over an hour-long period, and $E[A] = a$, then

$$E[A] = \sum_{i=1}^{S} p_i = \sum_{i=1}^{S} \rho c_i' = a \quad (1)$$

Here S is the total number of road segments in the region. From Equation (1), we have $\rho = a / \sum_{i=1}^{S} c_i'$. This scaling by ρ ensures that the total number of predicted crashes matches the historical number of crashes, a, in the given hour. In our formulation, a is the mean number of crashes observed on the day of week and hour of day corresponding to the hour in question.

In numerical terms, the crash classifier produces class probabilities c_i in the range [0,1]. Hourly road counts v_i are typically $O(100)$, giving $c'_i = c_i / v_i$ a typical range of $[0, O(0.01)]$. Calibrating with $p_i = \rho c'_i$ gives individual crash probabilities of roughly $[0, O(10^{-6})]$.

Crash Probability on Route

A route consists of a sequence of connected road segments. Since the order of traversal does not matter for our computations, we will designate a route as a set of road segments R, where each element of R is an index of a road segment.

We model the crash probability p_R of a route as a set of independent Bernoulli trials, with one trial for each road segment. The probability of traversing a single road segment i without crashing is $1 - p_i$. The probability of traversing the entire route R without crashing is then $\prod_{i \in R}(1 - p_i)$. Thus, the probability of crashing anywhere along the route is then

$$p_R = 1 - \prod_{i \in R}(1 - p_i) \tag{2}$$

Route Length vs. Safety Tradeoff

Intuitively, a safe route could be one that tediously weaves a long path around unsafe road segments. However, even with smaller crash probabilities, a longer route with more road segments leads to potentially more crash exposure. The Bernoulli probabilities make it easy to analyze this tradeoff. As a simple example, suppose that a shorter route consists of n_{short} road segments, each with an equal crash probability of p_{short}. From Equation (2), the probability of experiencing a crash on this route is $p_s = 1 - (1 - p_{short})^{n_{short}}$. An alternate longer route has a lower crash probability $p_{long} = \gamma p_{short}$ on each of its segments, with $0 < \gamma < 1$. The longer route also has more road segments $n_{long} = s n_{short}$, with $s > 1$. The probability of experiencing a crash on the longer route is

$$p_l = 1 - \left(1 - p_{long}\right)^{n_{long}} = 1 - (1 - \gamma p_{short})^{s n_{short}}.$$

Setting p_s=p_l gives the equivalence point where the two routes are equally safe:

$$1 - (1 - p_{short})^{n_{short}} \quad = \quad 1 - (1 - \gamma p_{short})^{s n_{short}} \tag{3}$$

For a given value of p_{short}, it is easy to compute values of γ and s that satisfy this equation. Plotted as ordered pairs

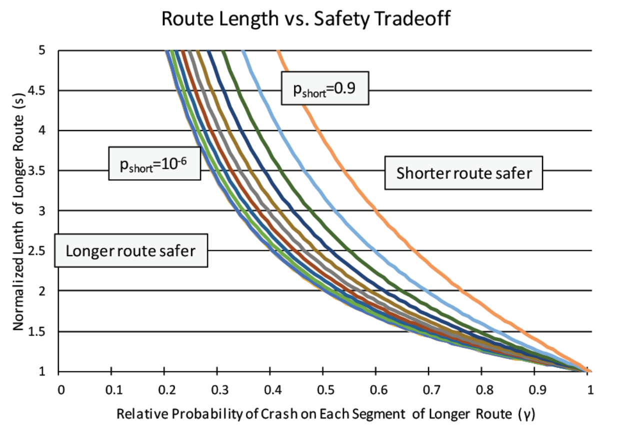

Figure 6: A longer route with safer road segments can be more or less dangerous than a shorter route with more dangerous road segments. (Note that vertical axis starts at 1.0.)

(γ, s), points on the lower left side of this curve indicate the longer route is safer. Values on the other side indicate the shorter route is safer. This is illustrated for several values of p_{short} in Figure 6. The lower left region of the (γ, s) space represents longer routes that are not very much longer than the corresponding shorter route and whose road segment crash probabilities are significantly less than those of the corresponding shorter route. In reality, the values of p_{short} are small, $O(10^{-5})$ or $O(10^{-6})$, meaning the realistic tradeoff curve is toward the left of the illustrated curves.

As an example, we examine the case with $p_{short} = 10^{-6}$, which pertains to the left-most curve in Figure 6. Examining the horizontal axis at 0.5, this represents a longer route with $p_{long} = 0.5 * p_{short}$, which might be considered safer. Looking vertically from 0.5, the curve is at about $s = 2$. This means that the longer route will be safer if it is less than about twice the length of the shorter route. If it is over twice as long, then the longer route will be ultimately less safe, even though its constituent road segments are safer.

Safer Driving Routes

Driving routes in navigation systems are almost always computed by minimizing a sum of costs. For computing safe routes, we used the probability of a crash as the cost to minimize. From the Bernoulli trials formulation, the probability of a crash along a route is given by Equation (2). Even though this is not a sum, even in log space, a simple Dijkstra algorithm can still be used to compute the safest route.

More specifically, Dijkstra's algorithm maintains a list that gives the minimum cost for traveling to each node visited on a potential route. Normally these costs are simply the sum of individual edge costs leading to that node. In

Hour of Day	Fastest Route Duration (minutes)	Fastest Route Crash Probability	Safest Route Duration (minutes)	Safest Route Crash Probability	Fastest to Safest Driving Time Multiple	Fastest to Safest Crash Probability Multiple
4 a.m.	13.14	0.0000098	20.68	0.0000047	1.56	0.52
8 a.m.	13.14	0.0000108	22.04	0.0000060	1.66	0.57
6 p.m.	13.14	0.0000062	24.83	0.0000031	1.85	0.51
Mean	13.14	0.0000089	22.52	0.0000046	1.69	0.53

Table 1: Means of route duration and crash probability for 100 routes evaluated at three different times of day. Safer routes have longer durations and smaller crash probabilities.

our case, these costs are instead the partial route crash probabilities computed from Equation (2). For computational efficiency and numerical stability, we maintained a parallel list giving $\sum log(1-p_i)$ for the partial route to each visited node.

We picked 100 arbitrary start and end pair locations in our study area for testing, for which we could compute the fastest and safest routes. The fastest (shortest time) route used only the lengths and speed limits of the road segments to estimate traversal time. Note that the fastest route computation did not account for turns, traffic lights, traffic congestion, or other delays as might be considered in navigation systems. Thus the computations of speed are independent of the time of day. Figure 1 shows one example pair of routes between the same endpoints.

To demonstrate the results, we looked at the hours of 4 a.m., 8 a.m., and 6 p.m. on January 20, 2015, which was a Tuesday. For each time of day, Table 1 shows the mean driving times and mean crash probabilities for the 100 test routes. It also shows the relative multiples between the driving times and crash probabilities between the fastest and safest routes. On average, the safest route takes 1.69 times as long to drive and has 0.53 times the crash probability, illustrating the tradeoff between driving time and safety. The mean crash probability of the fastest routes is $8.9\text{x}10^{-6}$, and the mean crash probability of the safest routes is $4.6\text{x}10^{-6}$.

It is possible to find routes that are between the safest and fastest, incurring a certain driving time penalty for a certain safety benefit. We computed these routes by using a cost function that includes both driving time and crash probability. Specifically, the cost function is

$$\alpha\beta p_R + (1-\alpha)\sum_{i\in R} T_i \tag{4}$$

Here p_R is the crash probability along a route whose edge indices are in set R, as given by Equation (2). The T_i are the costs of the route's edges in terms of driving time, and they sum in the usual way. The variable $0 \le \alpha \le 1$ is a blending parameter that controls the tradeoff between crash probability and driving time. Finally, β is an optional parameter that helps equalize the numerical magnitude of p_R and $\sum_{i\in R} T_i$. The crash probability p_R is usually $O(10^{-5})$ or (10^{-6}), and the driving time in seconds is usually

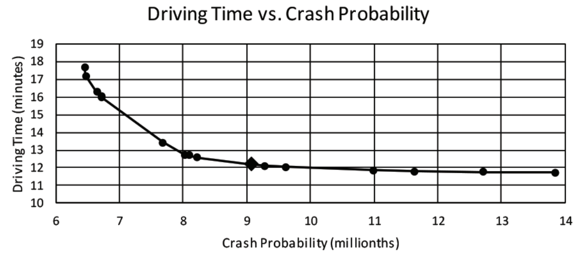

Figure 7: Smooth tradeoff between driving time and crash probability for the two routes shown in Figure 1. Each marker on the curve represents one route option. The route option marked with a diamond is displayed in black in Figure 1.

$O(10^2)$ or $O(10^3)$. Without β, the driving time dominates for almost the whole range of α. If we set $\beta = \sum_{i\in R} T_i / p_R$, then the full range of sampled α has a better chance of producing different routes. However, the method still holds for $\beta = 1$. Note also that the driving time part of the cost function could be any available cost function, including those that take into account delays due to traffic, turns, and other factors.

We used the cost function in Equation (4) to compute routes for the endpoints shown in Figure 1 over $0 \le \alpha \le 1$. The resulting tradeoff between driving time and crash probability is shown in Figure 7. Here it is apparent that lower crash probabilities are accompanied by longer driving times, as seen on the left side of the plot. The curve flattens as the driving time drops toward its minimum and crash probability rises. From Figure 7, there are many compromise routes whose driving time is approximately 12 minutes, but whose crash probability varies considerably. The fastest route, represented as the point farthest to the right in the plot, has a crash probability of $13.8x10^{-6}$ and a driving time of 11.7 minutes. Moving left from this point, the diamond-shaped point represents a compromise route with a driving time of 12.2 minutes (30 seconds longer than the fastest route) and a crash probability of $9.1x10^{-6}$, which is a 34% reduction. Many drivers would likely consider this to be a worthwhile tradeoff. This compromise route is shown in black in Figure 1. The compromise route in black partially overlaps both the fastest and safest routes.

Discussion and Conclusion

We described the use of data captured from vehicle crashes and counts to build predictive models that enable us to compute the probability of a crash along a driving route. The analytical pipeline involves estimating hourly traffic counts on unmeasured roads using a learned regression function and estimating crash risk as a function of environmental conditions using a learned classifier. We showed how to use results from the classifier to estimate Bernoulli crash probabilities along road segments and how to combine these probabilities to compute the crash probability of a route. These crash probabilities can be used in a Dijkstra-based route planner to compute the safest route between two points as well as to find compromise routes that trade off driving time for safety. We envision a routing application that presents a user with this tradeoff, giving drivers the flexibility to make decisions about travel time and safety. Although the probabilities of crash on any one trip are small, risks for drivers could accrue over years of driving, such as the overall risk associated with years of recurrent commuting trips.

On future directions, there is an opportunity to consider the severity of crashes and to study the implications of minimizing different kinds of safety-related outcomes, such as minimizing the probability of any crash versus minimizing the probability of fatalities or incapacitating injuries. Gaining access to severity of each crash and the number of vehicles involved would help us to refine crash probability inferences. On another direction, there is an opportunity to identify the likelihood of crashes associated with different components of routes such as for such connectives among segments as right and left hand turns and merges. Such probabilities of crashes could be integrated into the assembly of segments into routes.

Beyond use in decision-support for individuals, there is opportunity for analyses of safety being incorporated into large-scale directions services offering routing recommendations to populations of drivers. Potential influences of such recommendations on the likelihood of crashes based on unmodeled influences of higher usage of specific sets of segments would have to be monitored. Finally, we believe the information gain associated with sets of features in predictive models for crashes could be used to frame insight building and causal studies of factors influencing the safety of routes. Such efforts could lead to road revisions and design guidance that enhance the safety of driving.

References

NHTSA, National Highway Traffic Safety Administration. (2016). 2014 Motor Vehicle Crashes: Overview: U.S. Department of Transportation.

NOAA, National Oceanic and Atmospheric Administration. (2016). Integrated Surface Global Hourly Data. Retrieved June 21, 2016, 2016, from https://gis.ncdc.noaa.gov/geoportal/catalog/search/resource/details.page?id=gov.noaa.ncdc:C00532

WHO, World Health Organization. (2009). Global Status Report on Road Safety: Time for Action (D. o. V. I. P. D. (VIP), Trans.).

Choi, Eun-Ha, & Singh, Santokh. (2005). Statistical assessment of the glare issue-human and natural elements. Paper presented at the Federal Committee on Statistical Methodology Workshop.

Daniels, Stijn, Brijs, Tom, Nuyts, Erik, & Wets, Geert. (2010). Externality of risk and crash severity at roundabouts. Accident Analysis & Prevention, 42(6), 1966-1973.

Ericsson, Eva, Larsson, Hanna, & Brundell-Freij, Karin. (2006). Optimizing route choice for lowest fuel consumption–potential effects of a new driver support tool. Transportation Research Part C: Emerging Technologies, 14(6), 369-383.

Fink, Kenneth L, & Krammes, Raymond A. (1995). Tangent length and sight distance effects on accident rates at horizontal curves on rural two-lane highways. Transportation Research Record(1500), 162-168.

Friedman, Jerome H. (2001). Greedy Function Approximation: A Gradient Boosting Machine. Annals of Statistics, 29(5), 1189-1232.

Honsberg, Christina, & Bowden, Stuart. (2016). PV Education. Retrieved June 21, 2016, 2016, from http://pveducation.org/pvcdrom/properties-of-sunlight/declination-angle

Minneapolis, City of. Transportation Data Management System. Retrieved June 8, 2016, 2016, from http://minneapolis.ms2soft.com/tcds/tsearch.asp?loc=Minneapolis

Mitri, Sudeshna, & Washington, Simon. (2012). On the significance of omitted variables in intersection crash modeling. Accident Analysis & Prevention 49, 439-448.

Nadi, Saeed, & Delavar, Mahmood Reza. (2011). Multi-criteria, personalized route planning using quantifier-guided ordered weighted averaging operators. International Journal of Applied Earth Observation and Geoinformation, 13(3), 322-335.

Othman, Sarbaz, Thomson, Robert, & Lannér, Gunnar. (2009). Identifying critical road geometry parameters affecting crash rate and crash type. Annals of advances in automotive medicine, 53.

Retting, Richard A, Chapline, Janella F, & Williams, Allan F. (2002). Changes in crash risk following re-timing of traffic signal change intervals. Accident Analysis & Prevention, 34(2), 215-220.

Zografos, Konstantinos G, & Davis, Christian F. (1989). Multi-objective programming approach for routing hazardous materials. Journal of Transportation engineering, 115(6), 661-673.