

Development of a Real-time Prediction Model of Driver Behavior at Intersections Using Kinematic Time Series Data.

Yaoyuan V. Tan* Michael R. Elliott† Carol A.C. Flannagan ‡

Abstract

As autonomous vehicles enter the fleet, there will be a long period when these vehicles will have to interact with human drivers. One of the challenges for autonomous vehicles is that human drivers do not communicate their decisions well. However, the kinematic behavior of a human-driven vehicle may be a good predictor of driver intent within a short time frame. We analyzed the kinematic time series data (e.g., speed) for a set of drivers making left turns at intersections to predict whether the driver would stop before executing the turn or not. We used Principal Components Analysis (PCA) to generate independent dimensions that explain the variation in vehicle speed before a turn. These dimensions remained relatively consistent throughout the maneuver, allowing us to compute independent scores on these dimensions for different time windows throughout the approach to the intersection. We then linked these PCA scores to whether a driver would stop before executing a left turn using the Bayesian Additive Regression Trees (BART). Our model achieved an Area Under the receiver operating characteristic Curve (AUC) of more than 0.90 by -25m away from the center of an intersection.

Key Words: Area Under the receiver operating characteristic Curve, Bayesian Additive Regression Trees, Naturalistic Driving Data, Principal Components Analysis

1. Introduction

An autonomous vehicle can be loosely defined as a vehicle where no human supervision or human controlled driving is needed. The National Highway Traffic Safety Administration (NHTSA) provides a more detailed definition with five levels of classification (National Highway and Traffic Safety Administration, 2013):

- Level 0: The driver completely controls the vehicle at all times.
- Level 1: Individual vehicle controls are automated, such as electronic stability control or automatic braking.
- Level 2: At least two controls can be automated in unison, such as adaptive cruise control in combination with lane keeping.
- Level 3: The driver can fully cede control of all safety-critical functions in certain conditions. The car senses when conditions require the driver to retake control and provides a “sufficiently comfortable transition time” for the driver to do so.
- Level 4: The vehicle performs all safety-critical functions for the entire trip, with the driver not expected to control the vehicle at any time. As this vehicle would

*University of Michigan Department of Biostatistics, 1415 Washington Heights Ann Arbor, MI 48109

†University of Michigan Department of Biostatistics, 1415 Washington Heights Ann Arbor, MI 48109

‡University of Michigan Transportation Research Institute, 2901 Baxter Rd, Ann Arbor, MI 48109



Figure 1: Google car.

control all functions from start to stop, including all parking functions, it could include unoccupied cars.

Society of Automotive Engineers (SAE) International provides an alternative classification system (SAE International, 2014), but in this study we utilize the NHTSA classification and viewed an autonomous vehicle as a Level 4 classification.

A good example of a Level 4 autonomous vehicle is a vehicle from the Google Self-Driving Car Project. These vehicles could either be a modified Lexus Sport Utility Vehicle (SUV) or a prototype vehicle (Figure 1) built by Google. In 2009, Google started testing these self-driven vehicles on the streets of Mountain View, California and Austin, Texas. As of August 2015, Google reported that they have self-driven these vehicles for more than 1 million miles (Google, 2015). In July 2015, Google stated in a CNN report that these self-driven vehicles have been involved in a total of 14 accidents since 2009 including 11 rear-end crashes (CNNMoney, 2015). In all these accidents, Google asserted that human error and inattention was the main cause.

Google's claim is not surprising since a consequence of these autonomous vehicles being on the streets is that they will have to interact with human drivers in other vehicles. Unfortunately, human drivers do not always communicate their decisions clearly, leading to near crashes and crashes. As such, these autonomous vehicles will have to learn how to predict human driver decisions using information conveyed by the human driver's vehicle.

In this study, we hypothesized that the kinematic behavior of a human driven vehicle would provide enough information to make a good prediction of driver intent within a short time frame. In particular, we studied the speed of a human driven vehicle. We focused on predicting whether a human driver would stop at an intersection before executing a left turn. We believe that once the prediction model of such a simple driving behavior is fine-tuned to produce satisfactory results, we can extend this model to other forms of driving behavior. Our ultimate goal is to develop a prediction model of human driving behavior using the vehicle speed from the human driven vehicle.

To build the prediction model, we used naturalistic driving data from about 100 licensed drivers in Michigan. We converted the time series data to a distance series and defined a new distance-varying outcome. Because we believe that recent speeds contain more information about the human driver's intention to stop compared to past speeds, we employed a moving window on the distance-varying speeds. We next used Principal Components Analysis (PCA) to reduce the number of variables

used in our prediction algorithm. To link our distance-varying outcomes to our Principal Component (PC) variables, we used the Bayesian Additive Regression Trees (BART). We evaluated our BART model's prediction performance at every meter away from the center of an intersection by using the Area Under the receiver operating characteristic Curve (AUC). Finally, to visually search for an optimal predicted probability cut-off level that would balance both unnecessary stops by the autonomous vehicle and a crash, we plotted the Capture Ratio (CR) and False Positive Ratio (FPR) profile.

We organized the rest of our paper as follows: In Section 2, we provide additional details on the dataset, data manipulation, and statistical methods. In Section 3, we present the results of our analysis and finally in Section 4, we discuss our methods and results.

2. Data and Methods

2.1 Data

We obtained our dataset from a previous study by Sayer et al. (2011). In brief, our naturalistic driving data was collected from 108 licensed drivers in Michigan between April 2009 and April 2010. Sixteen late-model Honda Accords were fitted with cameras, recording devices, and a collision warning system – the Integrated Vehicle Based Safety System (IVBSS) – to collect visual and kinematic data from the drivers for a total of 40 days – 12 days baseline period with IVBSS switched off followed by 28 days with IVBSS activated. We used the 12 days baseline unsupervised driving data for this analysis. Because information about road types and intersections outside Michigan was not available, we restricted our analysis to driving within Michigan in order to facilitate the accurate identification of an intersection and its associated road type. Accurate identification of an intersection allows us to determine a reference time to start extracting the information necessary for this analysis.

In this study, we had data from 108 drivers who made 3,795 turns. Of these 3,795 turns, 1,823 were left turns. We took the time at -100m away from the center of an intersection as the reference point for the start of data extraction and stopped extraction at the time the vehicle was beyond the center of an intersection. We extracted both the speed of the vehicle (in m/s) and the amount of distance traveled (in m) at 10 millisecond intervals starting from our reference point. We also defined a vehicle as stopped when its speed was $\leq 1\text{m/s}$.

2.2 Data manipulations

During the early stages of analysis, we quickly noticed that in order to build a practical algorithm, the time series format of vehicle speed cannot be used. This is because vehicles are approaching the intersection at varying speeds. Hence, each vehicle takes a different amount of time to “cross” the center of intersection. This implies that we are unable to set a common reference time to indicate that the vehicle has crossed the center of an intersection. As such, we converted our time series of vehicle speeds to a distance series starting from -100m away from the center of an intersection to -1m away from the center of an intersection at every 1 meter interval.

We illustrate this conversion using an example with Driver 40 Trip 34 Turn 1. The figure on the left of Figure 2 shows the speed profile of this particular turn. In

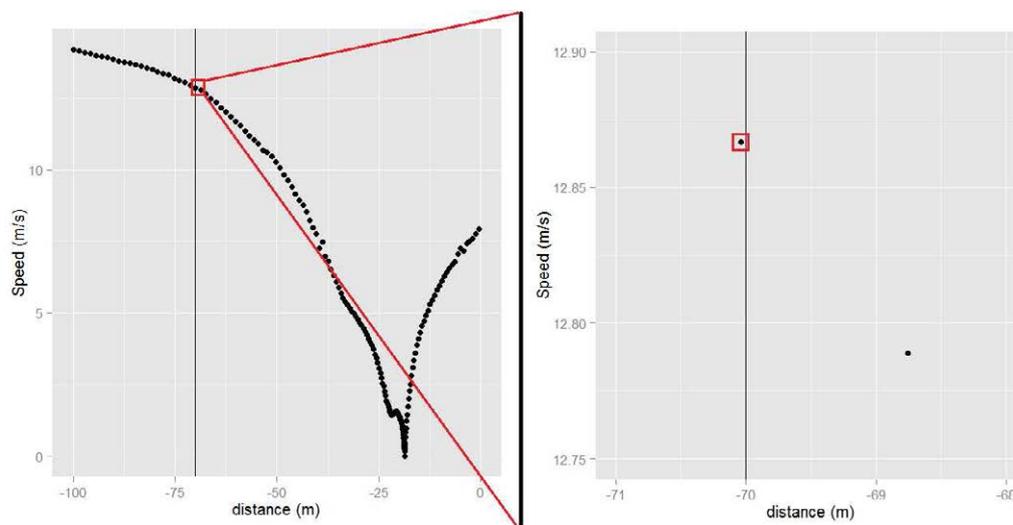


Figure 2: Original speed profile of Driver 40 Trip 34 Turn 1.

this example, our target is the vehicle speed at -70m . To obtain this speed, we first “draw” a line at -70m and focus on the speed sample points closest to this -70m line. The figure on the right shows the blow up of this focal point. To set the speed at -70m , we then compared which of the two speed sample points was closest to -70m . In our example, because the point on the left was closest, it was set as the speed at -70m for this turn. For the speeds of this turn from -100m to -1m at every 1m interval, we employed a similar approach. In the situation where more than one speed sample point was closest to the line, we took their average as the speed at that meter.

In addition to converting our time series to a distance series, we defined a new distance-varying outcome. This was done because we were interested in the question of “Will the human driven vehicle stop in the future?” at every meter away from the center of an intersection. An additional reason for defining a new distance-varying outcome was that we found turns where the vehicle stopped early. If we defined an overall outcome for each turn based on whether the vehicle stopped during the whole turn maneuver, we would find that at distances before the vehicle stopped, a prediction model would perform fairly well. Once the vehicle has stopped and starts to re-accelerate, the prediction model would begin estimating low probabilities of stopping; however, the observed outcome is still that the vehicle will stop, which is not accurate since the vehicle has already stopped and is now accelerating.

Hence, to define the new distance-varying outcome, we employed the following notation. Let i be the i^{th} turn and j be the j^{th} meter away from the center of intersection, $j = -100, \dots, -1$. Let s_{ij} be the new distance series of vehicle speed and y_{ij} be the distance-varying outcome (1 =stopped in future, 0 =will not stop in future) of the i^{th} turn at j be the j^{th} meter. Then, we defined y_{ij} as follows:

1. If $s_{ij} > 1\text{m/s} \forall j = -100, \dots, -1$, then set $y_{ij} = 0 \forall j$.
2. If $s_{ij} \leq 1\text{m/s}$ for some $j \in \{-100, \dots, -1\}$, let $c \in \{-100, \dots, -1\}$ be the index such that $\forall j > c, s_{ij} > 1\text{m/s}$. We set $y_{i,-100} = y_{i,-99} = \dots = y_{i,c} = 1$ and $y_{i,c+1} = y_{i,c+2} = \dots = y_{i,-1} = 0$.

Point 1 means that if the new distance series speed profile of a particular turn was more than 1m/s throughout, the distance-varying outcome would be set to 0

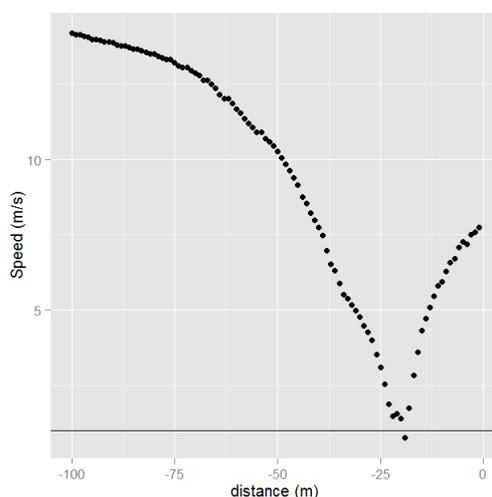


Figure 3: New distance series speed profile of Driver 40 Trip 34 Turn 1.

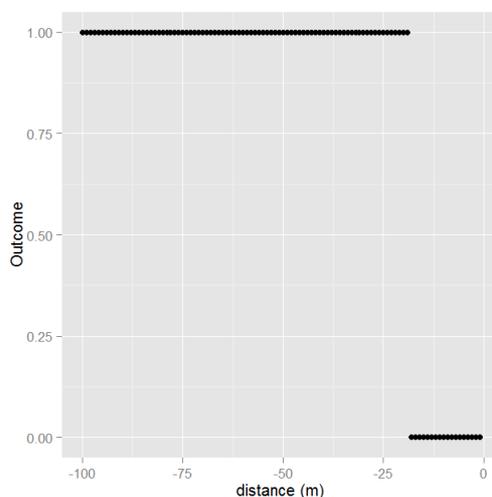


Figure 4: Distance-varying outcome for the speed profile of Driver 40 Trip 34 Turn 1.

throughout. Figure 3 and 4 clarifies point 2. Figure 3 corresponds to the new distance series of Driver 40 Trip 34 Turn 1. The horizontal line indicates 1m/s. We can see that for $j > -19$, the speed of the vehicle was more than 1m/s. Hence in Figure 4, the distance-varying outcome y_{ij} is set to 0 for $j = -18, \dots, -1$. On the other hand, because for $j = -100, \dots, -19$, the speeds s_{ij} could be less than or equal to 1m/s, we set their distance-varying outcome to 1.

2.3 Statistical methods

With the conversion and definition of the distance-varying outcome in place, we began developing our prediction model by first employing a moving window of speeds. This was done because as the vehicle approaches the center of intersection, recent vehicle speeds contain information on whether a human driver will decide to stop. The full profile of a vehicle's past speeds may include this information as well, but they may also contain irrelevant information making the full profile of a

vehicle’s past speeds “noisier” compared to a window of recent speeds. For every j^{th} meter $j = -100, \dots, -1$, we defined the moving window of speeds as,

$$M_{ij} = \{s_{i,j-w+1}, s_{i,j-w+2}, \dots, s_{ij}\}$$

where w is the size of the moving window.

Next, we used Principal Components Analysis (PCA) on these M_{ij} s to reduce the number of covariates in our prediction model. Here, the covariates are $s_{.,j-w+1}, s_{.,j-w+2}, \dots, s_{.j}$. We let

$$M_j = \begin{bmatrix} s_{1,j-w+1} & s_{1,j-w+2} & \dots & s_{1j} \\ s_{2,j-w+1} & s_{2,j-w+2} & \dots & s_{2j} \\ \vdots & \vdots & \vdots & \vdots \\ s_{n,j-w+1} & s_{n,j-w+2} & \dots & s_{nj} \end{bmatrix} \quad \text{and} \quad u(j) = \begin{bmatrix} u_{j-w+1} \\ u_{j-w+2} \\ \vdots \\ u_j \end{bmatrix}$$

where M_j is the matrix of moving windows with the first row being M_{1j} , second row being M_{2j}, \dots , and n^{th} row being M_{nj} . There are w orthogonal vectors $u(j)$ that decompose the variance of M_j into w parts under the condition that for each $u(j)$, $\|u(j)\| = 1$. To obtain the w decomposed variances, we used the formula: $PC_j = \text{Var}[M_j u(j)]$. If we let $PC_{j(q)}$ be the ordered statistic where $q = 1, \dots, w$ and $u(j)_{(q)}$ be the ordered vector corresponding to $PC_{j(q)}$, then the first Principal Component (PC) is $X_{j1} = M_j u(j)_{(w)}$, the second PC is $X_{j2} = M_j u(j)_{(w-1)}$, and so on. In our study, we found that the first 3 PCs explained at least 99% of the variation in M_j for all j (See Figure 6). Hence, the first 3 PCs were used as the predictors in our model.

In order to link our distance-varying outcomes to the first 3 PCs, we employed the Bayesian Additive Regression Trees (BARTs) model developed by Chipman et al. (2010). BART models the mean outcome of y_{ij} given covariates by a sum of regression trees and incorporates the additive effects of predictors. Formally, BART is written as

$$y_{ij} = \sum_{k=1}^m g(x_{i1}, x_{i2}, x_{i3}; T_k, V_k) + \epsilon_{ij}, \quad \epsilon_{ij} \sim N(0, \sigma^2). \tag{1}$$

The prior for equation (1) is decomposed as

$$\begin{aligned} p[(T_1, V_1), \dots, (T_m, V_m), \sigma] &= [\prod_k p(T_k, V_k)] p(\sigma) \\ &= [\prod_k p(V_k | T_k) p(T_k)] p(\sigma) \\ &= [\prod_k \{ \prod_l p(\mu_{lk} | T_k) \} p(T_k)] p(\sigma). \end{aligned}$$

The following distributions are then imposed on the priors

$$\begin{aligned} T_k &\sim \alpha(1 + d)^{-\beta}, \quad \alpha \in (0, 1), \beta \in [0, \infty), \\ \mu_{lk} | T_k &\sim N(\mu_\mu, \sigma_\mu^2), \text{ and} \\ \sigma^2 &\sim \nu\lambda / \chi_\nu^2. \end{aligned}$$

Gibbs-sampling is used compute the posterior distribution of σ and V_k while the Metropolis-Hastings algorithm is used to compute the posterior distribution of T_k .

Details of BART can be found in Chipman et al. (2010) while an explicit formulation of the Metropolis-Hastings algorithm can be found in Kapelner and Bleich (2014).

Figure 2 of Hill (2012) provides a good illustration of BART. In Hill (2012) Figure 2, the left panel shows an example of a single regression tree fit with z and x_1 as the covariates and y being the continuous dependent variable. A single regression tree fit is characterized by a splitting rule at each non-terminal node that would move the observed outcome y either to the left or to the right. Typically, the splitting rule would be based on some criterion for example, minimizing the variance, minimizing the Mean Squared Error (MSE), or maximizing AUC. If the covariates of the observed outcome y satisfy the splitting rule, it will be dropped to the left, else, it would be moved to the right. At the terminal nodes, the mean of the allocated y s would be calculated. The right panel of Hill (2012) Figure 2 compares visually the difference between a single tree fit and a BART fit. The blue discretized step-wise curve corresponds to the single tree fit in the left panel of Hill (2012) Figure 2. The red dashed smoother curve corresponds to the BART estimation. BART produces this smooth curve by first drawing many samples of m single tree fits and then taking a weighted sum of these samples. Here, the weights will be the probability that the draw is correct.

Because we had binary outcomes, we needed to modify the BART formulation slightly. Following Chipman et al. (2010)'s recommendation, we linked our distance-varying outcome to BART using a probit model:

$$p(x_{i1}, x_{i2}, x_{i3}) \equiv P[y_{ij} = 1 | x_{i1}, x_{i2}, x_{i3}] = \Phi[G(x_{i1}, x_{i2}, x_{i3})], \quad (2)$$

where $G(x_{i1}, x_{i2}, x_{i3}) \equiv \sum_{k=1}^m g(x_{i1}, x_{i2}, x_{i3}; T_k, V_k)$ and $\Phi[\cdot]$ is the c.d.f of a standard normal. In this probit model formulation, σ is fixed at 1 for identification purposes and the rest of the priors remain the same (Chipman et al., 2010). To calculate the posterior distribution, Chipman et al. (2010) suggested using a latent variable approach where z_{1j}, \dots, z_{nj} are independent identically distributed $N(G(x_{i1}, x_{i2}, x_{i3}), 1)$ such that $y_{ij} = 1$ if $z_{ij} > 0$ and $y_{ij} = 0$ if $z_{ij} \leq 0$. This formulation leads to

$$\begin{aligned} z_{ij} | y_{ij} = 1 &\sim \max\{N(g(x_{i1}, x_{i2}, x_{i3}), 1), 0\} \\ z_{ij} | y_{ij} = 0 &\sim \min\{N(g(x_{i1}, x_{i2}, x_{i3}), 1), 0\} \end{aligned}$$

The simulation from the posterior then becomes drawing $z_{ij} | y_{ij}, i = 1, \dots, n$ followed by drawing the posterior of T_k and V_k given z_{ij} instead of y_{ij} .

To evaluate our prediction model at every j^{th} meter away from the center of the intersection, we plotted the AUC value and its 95% confidence interval (CI) at every j^{th} meter. AUC calculates the proportion of observed outcomes that were ranked higher in terms of their predicted probability compared to the observed non-outcomes. Thus, a value close to 1 indicates that the prediction model is performing much better than chance while a value close to 0.5 indicates that the prediction model performs no better than chance. We computed the CI of the AUC using the method of Hanley and McNeil (1982), which uses a linear approximation of the AUC to the Somer's D statistic to obtain an estimate of the variance of AUC.

In addition to the AUC, we plotted the profile of the Capture Ratio (CR), the y-axis of the Receiver Operating Characteristic (ROC) curve and the profile of the False Positive Ratio (FPR), the x-axis of the ROC curve. For both profiles, we plotted them at nine different predicted probability cut-offs. Plotting the CR and

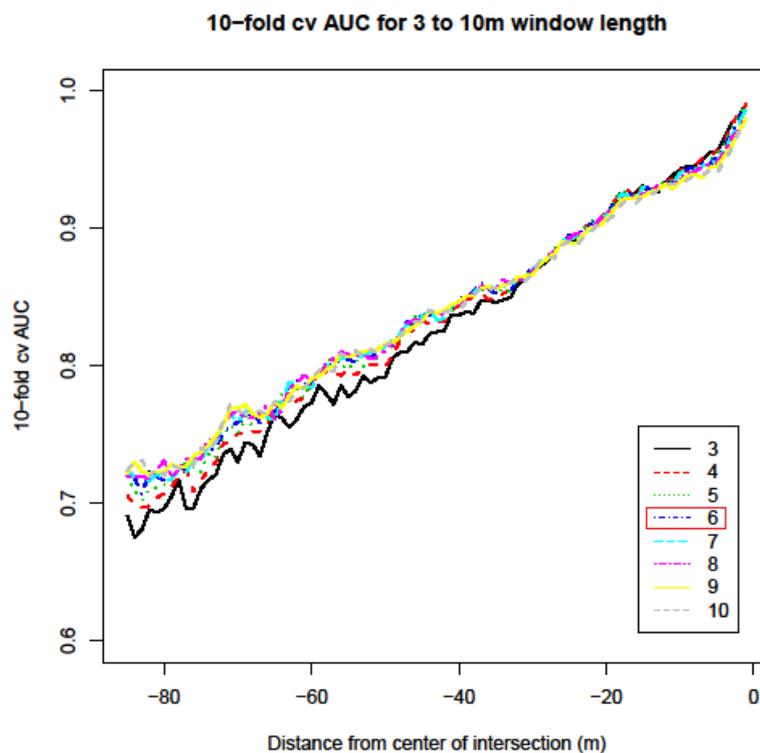


Figure 5: 10-fold cross validation AUC for window lengths 3 to 10.

FPR profile allows us to find the optimal predicted probability cut-off that will balance the probability of an unnecessary stop by the autonomous vehicle and the probability of a crash between the autonomous vehicle and a human driven vehicle.

3. Results

Our dataset contained 1,823 left turns with 894 of these turns started on major surface road types, 613 started on minor surface roads, and 316 were started on local roads. Major surface road types include roads supporting moderate travel within cities and quick travel between cities while minor surface roads include roads supporting moderate speed travel between neighborhoods. Local roads are defined as roads that support lower speed travel between neighborhoods. We also found 812 eventual stops defined as $s_{ij} \leq 1m/s$ for some $j \in \{-100, \dots, -1\}$. The average speed in all turns was 10.5 with a standard error of 4.2 and each driver took about 17 left turns (16.9, standard error 10.8).

We determined the length of our moving window w by using a 10-fold cross validation AUC (cvAUC) with the first 3 PCs as the variables and BART as the prediction model. We compared the cvAUC profiles with w from 3 to 50. Figure 5 shows the results of w from 3 to 10. We did not present the cvAUC profiles of w from 11 to 50 because they were all below the cvAUC profiles of w from 3 to 10. We chose a window length of 6 because the 10-fold cvAUC profile was higher compared to window lengths of 3 to 5 from -95m to -30m. Similarly, for distances more than -30m, the cvAUC of window length 6 was more than that of window lengths 7 to 10.

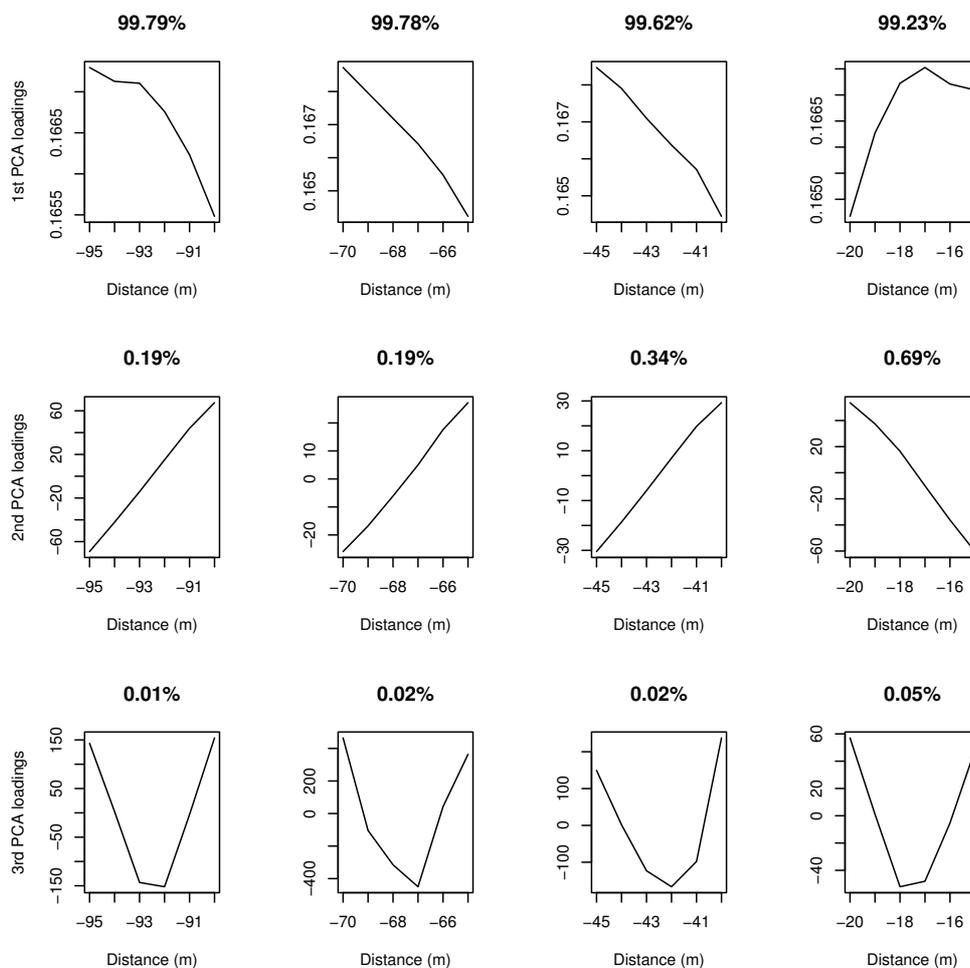


Figure 6: Principal Component loadings for the first, second, and third PC from -95m to -90m, -70m to -65m, -45m to -40m, and -20m to -15m (left to right).

Figure 6 shows the coefficients of the PCs allocated to the speeds from -95m to -90m, -70m to -65m, -45m to -40m, and -20m to -15m (left to right). Aside from the first 3 PCs explaining nearly 99% of the variation in M_j for all j , we also noted that the coefficient profile showed remarkable consistency throughout the approach to the center of intersection. What we found more surprising was the resemblance of the first PC to a form of average speed because of the small range of the coefficients. Similarly, the second PC resembled a form of acceleration or deceleration because of the linear profile of the coefficients and the wide range. We discuss some of the implications of this finding in Section 4.

Our BART model with $w = 6$ and using the first 3 PCs as predictors produced fairly good AUC results (Figure 7). The AUC profile together with its 95% CI were all above 0.7 throughout the left turn maneuver. Our AUC profile was 0.75 at -95m away from the center of intersection and steadily increased to over 0.80 by -60m out. It reached 0.90 by -25m out, and increased to 1 as the vehicle approached the center of intersection.

Figure 8 shows the CR and FPR profiles under nine different predicted probability cut-offs, 10%, 20%, ..., 90%. By a $x\%$ predicted probability cut-off we mean that for any predicted probability produced by our BART prediction model, those that were more than $x\%$ were labeled as stops and those that were less than or equal

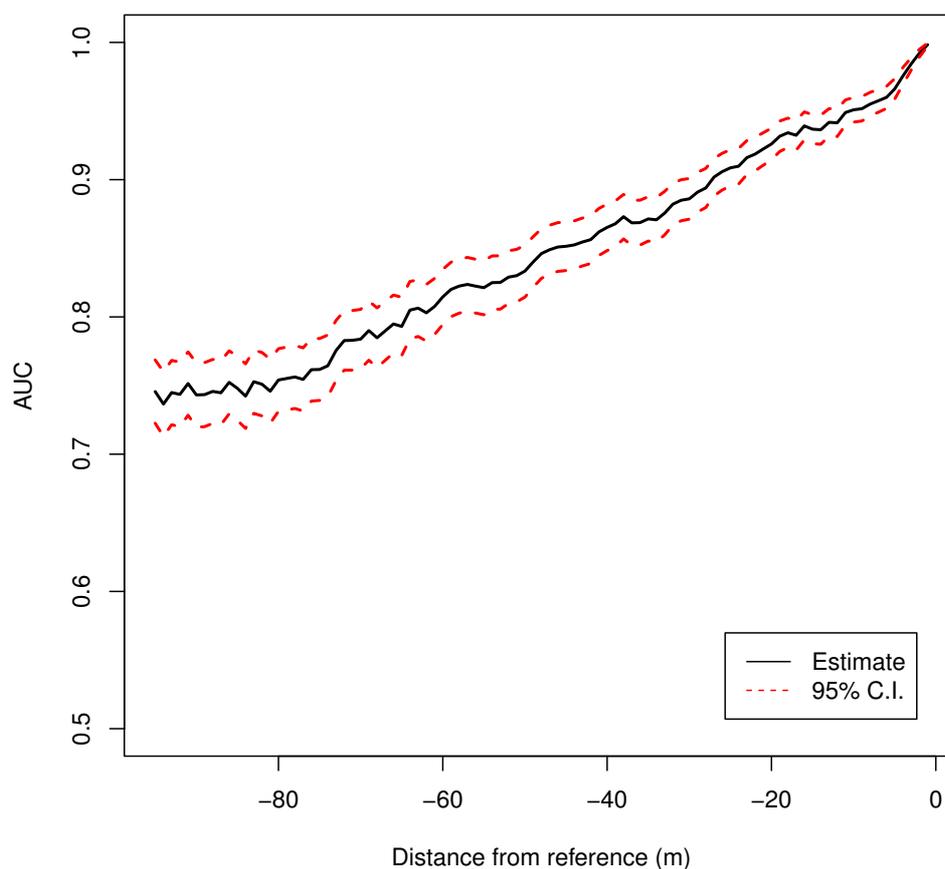


Figure 7: Area Under the receiver operating characteristic Curve (AUC) profile with 95% confidence interval (CI) of our BART prediction model.

to $x\%$ were labeled as non-stops. The CR then looks at the proportion of actual stops that were labeled correctly as stops and the FPR looks at the proportion of non-stops that were labeled incorrectly as stops. The solid lines in Figure 8 represent the CR and is equivalent to the autonomous vehicle correctly predicting that the human driven vehicle would stop using our BART model with the particular predicted probability cut-off. The dotted lines represent the FPR and is equivalent to the autonomous vehicle incorrectly predicting that the human driven vehicle would stop and hence a crash with the human driven vehicle would occur.

4. Discussion

In this study, we showed how we could use the kinematic behavior of speed from a human driven vehicle to predict the human driver's decision of stopping before executing a left turn. We employed a moving window of vehicle speeds to capture relevant information for prediction and used PCA to reduce the number of variables in our model. We then employed a recently developed model, BART, to link our distance-varying outcome to the PC variables. Finally we evaluated our prediction model by plotting the AUC, CR, and FPR profiles.

Six meters of speed data at each j^{th} meter away from the center of intersection gave us good cvAUC performance both near and far from the center of intersection. We used the first 3 PCs as the covariates in our prediction model because they explained at least 99% of the variation in M_j at each j^{th} meter. Our BART model

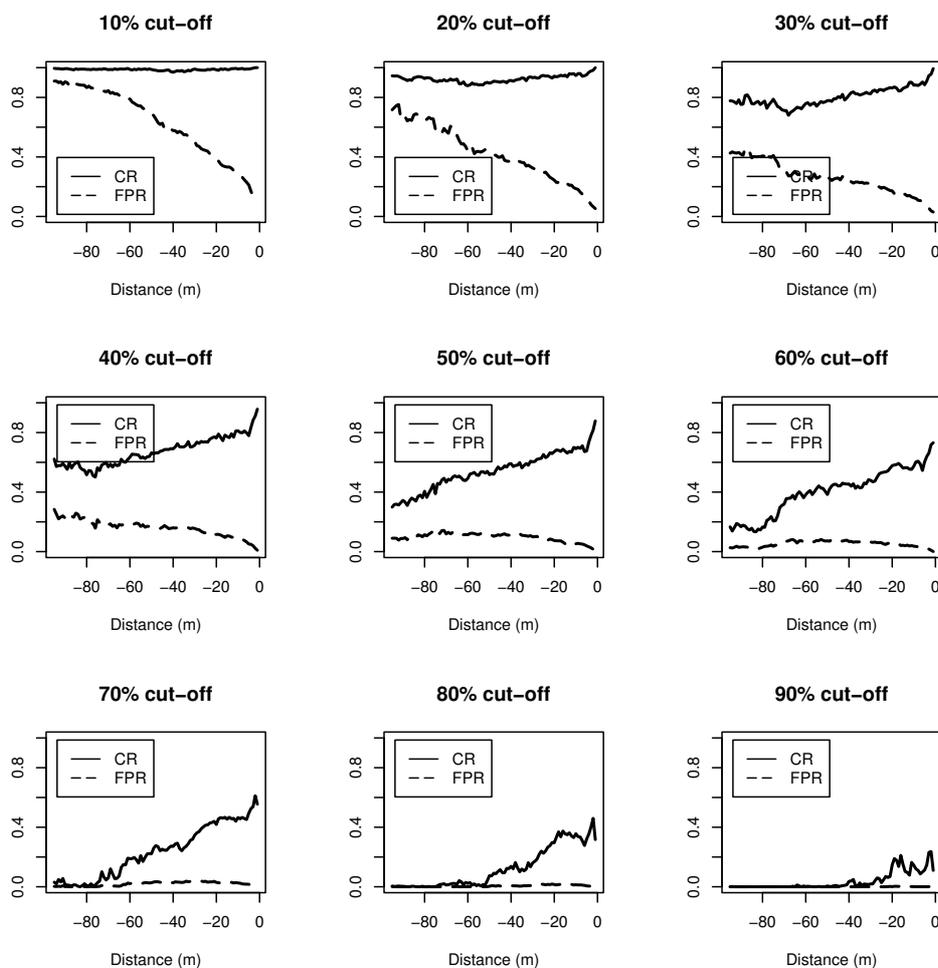


Figure 8: Capture Ratio (CR) and False Positive Ratio (FPR) profiles under nine different predicted probability cut-offs.

produced an AUC of 0.75 at -95m away from the center of intersection and this value increased steadily to 1 as the vehicle approaches the center of intersection.

When we converted the time series of vehicle speed to a distance series, we could have used a more sophisticated method to determine the vehicle speed. Some examples are linear interpolation, non-linear interpolation, and smoothing splines. However, we chose not to employ any of these methods because the likely loss in precision of estimating the speed would be at most 0.01 to 0.02 m/s. Such a small loss in the precision of the vehicle speed is unlikely to influence the final results. Hence, we preferred our current method.

We also considered many alternative statistical methods which we have not presented. For the use of a moving window of speeds, we originally employed a long window of vehicle speeds where at each meter, we kept increasing the window of speeds we considered. This corresponded to the definition of M_j as

$$M_j = \begin{bmatrix} s_{1,-100} & \cdots & s_{1j} \\ s_{2,-100} & \cdots & s_{2j} \\ \vdots & \vdots & \vdots \\ s_{n,-100} & \cdots & s_{nj} \end{bmatrix}.$$

We found that the AUC profile produced was not comparable to using a moving

window definition. We also tried to see if taking a weighted combination of moving window and long window definition would improve the AUC profile but it did not. Hence, we concluded that if we included past speeds in our prediction model, more “noise” would be introduced, reducing the model’s prediction performance.

We decided to use PCA as the method to reduce the number of variables in our model because of the surprising consistency we found in the profile of the PC coefficients (Figure 6). We did not base our choice of the first 3 PCs as the variables in our model only on the amount of variation in M_j explained. Since our ultimate goal is prediction, we investigated how much prediction performance would be added with the inclusion of the first 5 PCs in terms of AUC. We found huge increases in the AUC profile when the first 2 PCs were added and a substantial increase when the third PC was added. When the fourth and fifth PC were added, we found no improvement in the AUC profile. Moreover, when we plotted the PC values of the fourth PC and above, we found them to be rather inconsistent and difficult to interpret. We also considered using speed and acceleration in place of the first and second PC, given their resemblance to these quantities. However, we found that the resulting AUC profile was substantially lower compared to the AUC profile from using the first and second PC. Another alternative to using the first 3 PCs was the direct use of the 6 meters of speed as variables in the BART model. The rationale of this method was we can view the first 3 PCs as linear combinations of the 6 meters of speed since $X_{j(q)} = M_j u(j)_{(q)}$. So the use of the first 3 PCs and the 6 meters of speed data would be “similar”. In addition, PCA involves matrix multiplications which could slow down computation when the number of observations increase. Unfortunately, this alternative method does not produce an AUC profile better or comparable to the AUC profile produced using the first 3 PCs. We suspect the reason is PCA further extracts useful information from the 6 meters of speed data. And by using all the information from the 6 meters of speed data, some noise may have been added.

We also considered many prediction models as alternatives to the BART model including the linear logistic regression model with the first 3 PCs as covariates, the non-linear logistic regression model using cubic splines with a knot at the mean of each of the 3 PCs, and the Super-Learner (van der Laan and Polley, 2010). The Super-Learner is an ensemble method that combines the prediction result of any machine learning to obtain a better prediction model. In our Super-Learner implementation, we used the following machine learning methods: elastic-net regularization path for logistic regression (Friedman et al., 2010), logistic regression, K-Nearest Neighbor (KNN), Generalized Additive Models (Hastie and Tibshirani, 1990), mean of the outcomes, and BART. The AUC profile of the BART model was better compared to the linear and non-linear logistic regression model. Although the AUC profile of the Super-Learner was somewhat better compared to BART, the improvement was highly variable with various distances performing the same as BART. Therefore, we chose BART as our prediction model.

Although our BART model performed well in predicting a pre-left turn stop, there is still room for improvement. Firstly, we did not use other baseline covariates like presence of a lead vehicle, distance from the center of intersection the turn signal was first activated, and many others. Including these variables may improve the performance of our prediction model further away from the center of intersection. We were less concerned about the performance near the center of an intersection since the AUC of our model was already close to 1. We intend to investigate which covariates should be included to improve performance by using the BART variable

selection method proposed in Bleich et al. (2014). A point to note here is the inclusion of variables such as the gender and age of the driver may not be practical since it is unlikely that the sensors equipped on an autonomous vehicle would be able to capture such information.

Second, we are aware that our naturalistic driving data was collected from same drivers traveling on similar road types multiple times. This implies that our assumption that each turn is independent from the other may be violated since there could be some form of intra-driver or intra-road type correlation between turns. We believe this can be solved by extending the BART model to include a random intercept. Preliminary results by stratification showed promise and we are currently working on implementing a random intercept BART model to our data.

Finally, on closer inspection of our nine different CR and FPR profile plots, we can see that different predicted probability cut-offs could be proposed at different distances instead of one overall cut-off. This implies that different decisions could be made at different distances depending on the cost we decide to allocate to either correctly predicting a driver stop and hence avoid unnecessary stops in the autonomous vehicle, or incorrectly predicting a driver stop and hence a crash with the human driven vehicle would occur. To obtain the different optimal cut-offs that would balance the CR and FPR at each distance, we suggest attaching different costs to the CR and FPR at each j and then employ numerical methods to solve for the optimal cut-off.

References

- Bleich, J., Kapelner, A., George, E., and Jensen, S. (2014). Variable selection for BART: An application to gene regulation. *The Annals of Applied Statistics* **8**, 1750–1781.
- Chipman, H., George, E., and McCulloch, R. (2010). BART: Bayesian Additive Regression Trees. *The Annals of Applied Statistics* **4**, 266–298.
- CNNMoney (2015). Injuries in Google self-driving car accident, Retrieved August 26, 2015, from <http://money.cnn.com/2015/07/17/autos/google-self-driving-car-injury-accident/>
- Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software* **33**, 1–22.
- Google (2015). What were up to, Retrieved August 26, 2015, from <http://www.google.com/selfdrivingcar/>
- Hanley, J. and McNeil, B. (1982). The Meaning and Use of the Area under a Receiver Operating Characteristic (ROC) Curve. *Radiology* **143**, 29–36.
- Hastie, T. and Tibshirani, R. (1990). *Generalized Additive Models*. Chapman and Hall, London.
- Hill, J. (2012). Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics* **20**, 217–240.
- Kapelner, A. and Bleich, J. (2014). bartMachine: Machine Learning with Bayesian Additive Regression Trees. [arXiv:1312.2171v3 \[stat.ML\]](https://arxiv.org/abs/1312.2171v3) .

- National Highway and Traffic Safety Administration (2013). U.S. Department of Transportation Releases Policy on Automated Vehicle Development, Retrieved August 26, 2015, from <http://www.nhtsa.gov/About+NHTSA/Press+Releases/U.S.+Department+of+Transportation+Releases+Policy+on+Automated+Vehicle+Development>
- SAE International (2014). *Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems*. SAE International, Warrendale, PA.
- Sayer, J., Bogard, S., Buonarosa, M., LeBlanc, D., Funkhouser, D., Bao, S., Blankespoor, A., and Winkler, C. (2011). Integrated Vehicle-Based Safety Systems Light-Vehicle Field Operational Test Key Findings Report DOT HS 811 416, Retrieved August 26, 2015, from <http://www.nhtsa.gov/DOT/NHTSA/NVS/Crash%20Avoidance/Technical%20Publications/2011/811416.pdf>
- van der Laan, M. and Polley, E. C. (2010). Super Learner in Prediction. *U.C. Berkeley Division of Biostatistics Working Paper Series Working Paper 266*, <http://biostats.bepress.com/ucbbiostat/paper266>