Investigating Collision Patterns to Support Autonomous Driving Safety

Carmen Kar Hang Lee*

School of Business, Singapore University of Social Sciences, Singapore.

Email: carmenleekh@suss.edu.sg

Eric Ka Ho Leung

Management School, University of Liverpool, United Kingdom.

Email: k.h.leung@liverpool.ac.uk

Ying Kei Tse

Cardiff Business School, Cardiff University, United Kingdom.

Email: tsem1@cardiff.ac.uk

Yu-Chung Tsao

Department of Industrial Management, National Taiwan University of Science and Technology, Taipei, Taiwan, ROC.

Email: yctsao@mail.ntust.edu.tw

*Correspondence Address

Carmen Kar Hang Lee

School of Business

Singapore University of Social Sciences

463 Clementi Road, Singapore 599494

Tel: +65 6240 8817

Email: carmenleekh@suss.edu.sg

Abstract

There is a debate on the importance of autonomous vehicles (AVs) and the methods for ensuring that AVs are safe to be driven on roads open to traffic, which has become the priority for engineers, artificial intelligence experts, and data scientists who wish to enhance autonomous driving safety. This paper proposes a machine learning framework for investigating AV collisions, thereby advancing the knowledge on the risk factors of AV collisions. We used the California Department of Motor Vehicles' AV collision reports from January, 2019 to October, 2021 to determine the association between risk factors and the level of damage to an AV due to collisions. Association rule mining was used to develop methodologies that can advance result interpretability, which is crucial in the transportation field as it will lead to the development of evidence-based policies. A total of twenty-one rules were determined and used to reveal the unique safety patterns of AVs to understand the factors that co-occur with AV damage. This study demonstrates that collision data, when analysed using appropriate machine learning algorithms, can generate useful insights that complement current AV-related policies and provide practical information that can be used to enhance autonomous driving safety in the long term.

Keywords*: Autonomous driving, autonomous vehicle collision, traffic accident analysis, machine learning, association rule mining*

Paper type: Research paper

1. Introduction

Many countries have announced their strategic plans to develop autonomous vehicles (AVs) as a means to improve their cities' sustainability (Taeihagh and Lim 2018). One of the challenges that these cities have to overcome is promoting consumer trust in AVs, which has been significantly decreasing in recent years because of the ongoing media coverage of AV crashes. For example, autonomous driving technology caused a fatal crash in Arizona, United States of America in 2018 when an AV operated by Uber struck and killed a pedestrian (Driggs and Wakabayashi 2018), which was the first pedestrian death associated with AVs. Although the COVID-19 pandemic has hastened the adoption of AVs (Lee and Wu 2021), over 75% of respondents in Euromonitor's Mobility survey conducted in February, 2020 indicated that they would not feel comfortable in an AV (Liuima 2020). In addition, Morning Consult, a consultancy firm in the United States, reported a continuous drop in consumer safety perception of AVs (Whalen 2022). In March, 2022, only 19% of consumers perceived AVs as safer than traditional vehicles, which was 8% lower than in 2018. In order to ease public fear, AV manufacturing firms need to safeguard AV passengers and other road users (e.g., motorcyclists, cyclists, conventional car drivers, pedestrians). Any accidents involving AVs will seriously influence consumer trust, eventually resulting in a lack of public acceptance. Therefore, reducing AV systems' flaws and enhancing road safety has become the priority for engineers, artificial intelligence experts, and data scientists.

Analysing the factors of traffic accidents can provide critical knowledge for enhancing AV safety and designing roadway infrastructure that accommodates AVs. Traditional statistical methods remain the most popular approach to analysing traffic accidents. However, these methods usually have pre-defined underlying relationships between the dependent and independent variables that may significantly hamper the determination of complex associations between the contributing factors (Xu et al. 2018). In addition, the assumptions made in these studies may affect the reliability of the results (Chang and Chen 2005). Machine learning approaches can be used to make up for these deficiencies as they go beyond traditional statistical methods when explaining traffic accidents. However, certain machine learning algorithms, such as artificial neutral networks, are widely criticised as being 'black boxes', which denotes decision makers not being able to explain how a specific solution is derived. This leads to a barrier for implementation in the transportation field in which searching only for high accuracy is not sufficient and the developed methodologies are expected to help drive evidence-based policies. On the other hand, rule-based approaches can generate results that can be understood by humans. Thus, studies have harnessed the features of association rules to enhance interpretation (Ghanem et al. 2021), which are a form of IF-THEN rules that can be easily understood by humans. In the transportation field, there are many risk factors for accidents, but by using the IF-THEN rules, decision makers can effectively identify the set of antecedents associated with the consequence (e.g., damage level, injury level). Studies show that traffic accidents are the result of complex interactions among multiple risk factors, such as road users' reactions, roadways' characteristics, and environmental conditions (Montella 2011; Xu et al. 2018). Therefore, association rule mining is a suitable technique as it shifts the focus from individual factors to a combination of factors (Jiang et al. 2020). This paper aims to contribute to the knowledge on the risk factors that contribute to traffic accidents in the context of AVs by proposing a machine learning framework that incorporates association rule mining to explain AV collisions. Accordingly, the first research question is as follows: *What are the key factors associated with collisions in the context of AVs?*

Traffic safety is imperative in autonomous transport systems, and there are two approaches to testing the safety of an AV. The first approach comprises an on-road test in which an AV operates in real driving situations and interacts with other road users. However, according to Kalra and Paddock (2016), it may take hundreds of years for an AV to travel enough miles to prove its reliability. The second approach consists of simulation, which is more cost-effective than on-road testing from a practical point of view (Makkula et al. 2012). It is also less time-consuming, and manufacturers do not need to go through a lot of administrative work to gain permits for this type of testing (Rosique et al.

2019). Nevertheless, various factors can affect the reliability of simulation results, such as sensor accuracy and the vehicle and environment models (Huang et al. 2016). Furthermore, simulation takes important test-scenario considerations into account, such as the need for established definitions of relevant traffic scenarios in which AVs can be tested, which are currently lacking in transport policies (Coelingh et al. 2010; Schöner 2018). To fill this gap, we propose a machine learning framework that can serve as a decision-support tool when determining crucial traffic scenarios for AV testing. Accordingly, the second research question is as follows: *What risk factors can be incorporated into traffic scenarios for AV testing?*

To answer these research questions, a machine learning approach is used to investigate AV collisions and safety patterns. We collected all of the AV collision reports from California's Department of Motor Vehicles (DMV) for collisions that occurred between January, 2019 and October, 2021. The DMV has granted over 50 permits to AV manufacturers, thereby allowing them to test their AVs on public roads in California. Regulations require manufacturers to report any collision that causes property damage, physical injury, or mortality within ten days. These AV collision reports are valuable data sources for identifying risk factors in AV collisions. Most importantly, we analysed a comprehensive dataset containing all the AV collision cases (n = 231). The findings of this study can aid in establishing transport policies for safe autonomous transport systems. This study contributes to the current policies on AV safety that do not discuss the interactions between the contributing factors of AV collisions. Unlike previous studies that investigate individual factors independently, this study analyses the factors in relation to one another. Therefore, the association rules derived from this investigation are different to the simple explanatory models implicit in conventional statistical methods.

The paper is structured as follows: Section 2 reviews the relevant literature on AVs and traffic accident analysis; Section 3 proposes a machine learning framework that can extract useful knowledge from AV collision reports in order to establish transport policies; Section 4 presents the results and discussion; and Section 5 concludes this study and highlights future research directions.

2. Literature Review

2.1 AVs

The rapid growth in information and communication technologies is enabling cities to undergo digital transformations and streamline smart services (Anthony 2021; Suman et al. 2020). Smart transport in particular is considered effective in addressing undesirable effects caused by conventional vehicles, such as greenhouse gas emissions and traffic congestion (Yigitcanlar 2018). AVs, also referred as self-driving or driverless cars, play an important role in recent urban development concepts that address problems arising from traffic, pollution, and energy consumption (Seuwou et al. 2020; Zhu et al. 2018). As such, AV development has been a crucial topic in contemporary urban policy agendas (Creutzig et al. 2015; Perveen et al. 2017). AVs utilise cutting-edge technologies and sensors to achieve self-driving capabilities without human involvement (Rosique et al. 2019). Despite AVs' potential to shape urban mobility, the adoption of AVs requires changing individuals' travel behaviours as, if users do not perceive AVs as safe, they will be less likely to switch from conventional vehicles to AVs. There are two dominant research streams related to AV safety: The first stream consists of studies investigating AV safety issues based on simulations, and the second stream comprises AV-involved accident analysis using conventional statistical methods, such as regression analysis.

In the literature, AV safety is mainly discussed according to simulation results rather than real AV-driving data. Ye and Yamamoto (2019) used simulations to examine the impact of connected autonomous vehicles on traffic safety, and their analysis reveals that traffic safety can be significantly improved when connected AVs' penetration rates on roads are increased. A similar conclusion is drawn by Papadoulis et al. (2019), who developed a connected AV control algorithm that was implemented in a simulated motorway segment. Their results indicate that the presence of connected AVs can ensure efficient traffic flow. Some simulations focus on roadway infrastructure, such as roundabouts and unsignalised intersections. Deluka Tibljaš et al. (2018) analysed roundabout safety levels in circumstances in which various AVs were mixed with conventional vehicles. These simulations were run in conjunction with the acquisition of field data on speed and traffic volumes from existing roundabouts in Croatia, and the results suggest that the introduction of AVs could change both operational and safety parameters at roundabouts. Another simulation model developed by Maryam Mousavi et al. (2020) was used to assess the safety impacts of AVs at urban unsignalised intersections and in the proximity of urban driveways. The results of both studies demonstrate the need to justify AVs' safety levels prior to their actual implementation in future.

On the other hand, several studies largely rely on traditional statistical methods for traffic accident analysis. However, the empirical findings from these studies are not univocal. For instance, Teoh and Kidd (2017) calculated the crash rates per million vehicle miles travelled based on Google data, and they concluded that Google's self-driving cars had much lower rates of police-reportable crashes per million vehicle miles travelled than human drivers. Contrastingly, Schoettle and Sivak (2015) and Favarò et al. (2017) discovered that AV-involved crash rates are much higher than normal crash rates. In light of a potentially higher rate of AV-involved crashes, Xu et al. (2019) used regression methods to discover key factors contributing to the severity level of crashes. However, they estimated the net effect of each factor independent from the others, thereby failing to address the interactions between the factors. This is a drawback because traffic accidents are complex and can be influenced by a combination of multiple factors. Consequently, it is imperative to understand the co-occurrence of multiple factors that could result in an accident.

These studies show that the majority of AV safety performance analyses have been statistical in nature or have been conducted using simulations. Currently, no artificial intelligence tools nor machine learning algorithms have been used to investigate the contributing factors for AV collisions and the interactions among those factors. Therefore, association rule mining, a machine learning algorithm capable of generating human-explainable association rules, is proposed under this study's framework. This study departs from prior studies in two ways: Firstly, unlike most studies, it analyses a set of real AV collision reports but not simulated data, and secondly, by using association rule mining, this study seeks association patterns in what causes AV collisions rather than treating each cause separately.

2.2 Association Rule Mining

One of the most widely used algorithms for association rule mining is the Apriori algorithm (Agrawal and Srikant 1994), which is used to identify co-occurrence patterns from data in the form of association rules. An association rule, denoted as $A \rightarrow C$, consists of an *antecedent* (i.e., A) and a *consequent* (i.e., C). Every possible value of a variable is considered an *item* in the Apriori algorithm. Association rules indicate the relationships among the items. In an AV collision analysis, A represents the set of risk factors for the collision, while C represents the outcomes of the collision.

Two hyperparameters of the Apriori algorithm are the *support* and *confidence threshold* values. The support of A is the percentage of instances containing A in the dataset, and it is the probability of A occurring. Therefore, $\textbf{Support}(A) = \textbf{P}(A)$. Its value has to be greater than or equal to the support threshold in order for A to be included in an association rule. The confidence of a rule is the percentage of instances containing A that also contain C; it is the conditional probability of C occurring given A. Thus, $\textbf{Confidence}(A \rightarrow C) = \textbf{P}(A \cup C)/\textbf{P}(A)$. It is used to measure the strength of the association between A and C. A rule is considered interesting if its confidence is equal to or larger than the confidence threshold. Before model construction, users must specify the support and confidence thresholds using trial-and-error approaches. The number of association rules decreases with the threshold values. In addition, *lift* is another measure that can be used to filter association rules as it is measures the number of times *A* and *C* co-occur more often than expected if *A* and *C* are statistically independent. Thus, **Lift(*A*→*C*) = Confidence(*A*→*C*)/Support(*C*)** (Das 2021). It indicates the independence between *A* and *C*. In general, a lift greater than one implies that *A* and *C* depend on each other.

Although many studies select association rules based on rule confidence (Lee and Wu 2021; Montella 2011), there is a drawback to this approach: Rules with a high confidence may not always be interesting. For example, in our dataset, the support of 'Damage Level = Minor' was 73%, so the rule '∅ → Damage Level = Minor' had a confidence of 73%, where '∅' represented an empty antecedent. Even if the fictional rule of 'Lighting = Daylight → Damage Level = Minor' with a confidence of 75% fulfilled the confidence threshold, it is not interesting because the presence of daylight does not have a significant influence on the consequent item as the confidence only increases slightly by 2%, from 73% to 75%. Hence, simply using the confidence threshold alone cannot suppress the second rule. On the other hand, the fictional rule 'AV Mode = Autonomous → Damage Level = Minor' with a low confidence of 30% may be interesting because it reveals that the chance of minor damage becomes significantly lower when the AV operates in an autonomous mode, from 73% to 30%. This example shows that the confidence of a potentially interesting rule should be significantly different from the confidence of the rule that has the same consequent but a simpler antecedent. Thus, adding an item to the antecedent is helpful only when it causes a significant change in the confidence of the rule. Therefore, another measure is employed herein in addition to confidence to evaluate association rules: confidence ratio, in that **Confidence ratio = 1 – min ($C_{posterior}/C_{prior}$, $C_{prior}/C_{posterior}$).** The $C_{posterior}$ represents the confidence of the association rule that is under evaluation, and the C_{prior} represents the confidence of the association rule with the same consequent item but an empty antecedent. Thus, the confidence ratio is higher when the difference between $C_{posterior}$ and C_{prior} is greater.

To summarise, the existing studies in the conventional vehicle domain are insufficient for advancing our knowledge of AV collision and AV safety. Therefore, this study uses association rules to determine collision patterns from AV collision reports. The obtained rules can be used to help AV manufacturers enhance AV safety, and they will also enable transport authorities or urban planners to design roadway infrastructure that accommodates future transport means, including AVs.

3. The Machine Learning Framework

This study proposes a machine learning framework for investigating AV collisions. The framework involves three phases (see fig. 1). Phase 1 comprises the data collection and preparation processes, which, in this study, started with extracting data from the DMV's AV collision reports and preparing the data for collision analysis. Phase 2 consists of applying the Apriori algorithm to discover the association rules and using k-fold cross-validation to construct *k* models, with each generating a set of rules from the training data that are validated using the testing data. Thus, the only valid rules are those that can be validated by the *k* models. This process iterates multiple times to further remove the resampling bias. Lastly, Phase 3 is the evaluation of the results. It consolidates all of the valid rules from Phase 2 upon each iteration, with the final valid rule set comprising the valid rules that exist in all the iterations. The quality of each of these rules is evaluated using the metrics' average values. Each phase is explained in detail below.

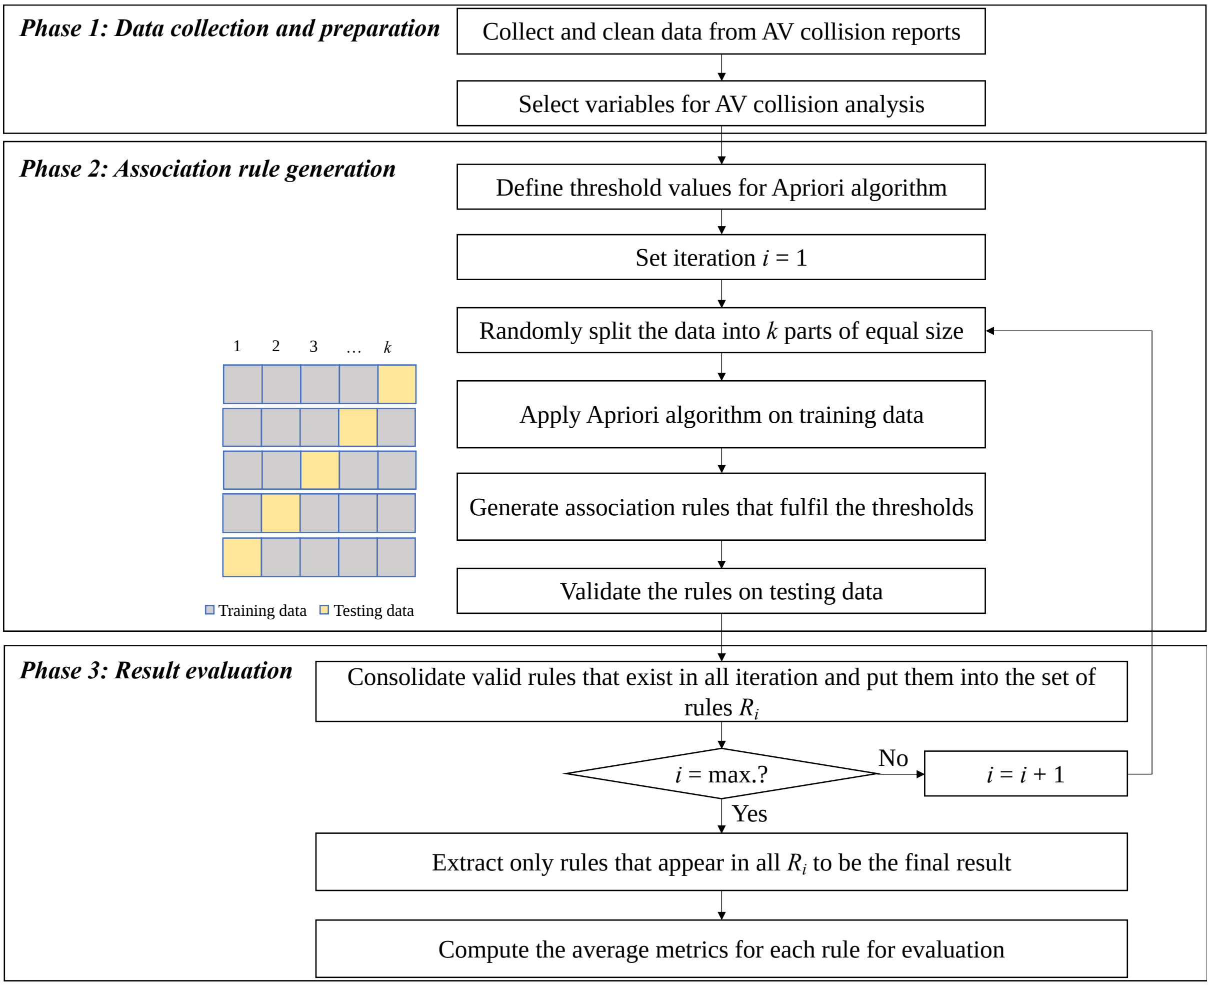

Figure 1. A machine learning framework for AV collision analysis

3.1 Data Collection and Preparation

Phase 1 started with collecting the AV collision reports from the DMV, which is a leader in promoting AV development. From January, 2019 to October, 2021, it received 231 AV collision reports from AV manufacturers within ten days of the collisions that occurred during AV testing. In each report, the areas of an AV that were damaged in a collision were recorded (see fig. 2). The DMV uploads the collision reports to its website in PDF format, for which a Python program was developed to process the data from the collision reports. Each record represents a collision instance. A careful examination of the dataset demonstrates that there was a contradiction in the values for some variables in certain records. For example, some reports stated that the AV was ‘moving’ in one variable, but the AV was indicated as being ‘stopped’ in another variable. To resolve these data inconsistencies, the textual field that contained the narrative of the collision was checked to correct the values of the fields accordingly. Table 1 describes the key variables used for this study after the data cleaning process.

Table 1. This study's variables

Variable	Description	Possible values: Distribution (%)
Time	The period of day at the time of the collision	AM (42.42%) PM (57.58%)
Object Collided	The object that collided with the AV	Passenger vehicle (83.55%) Bicyclist (4.33%) Object (4.33%) Truck (2.60%) Scooterist (2.16%) Skateboarder (1.30%) Bus (0.87%) Motorcyclist (0.43%)
AV's Movement	The movement of the AV prior to the collision	Stopped (38.96%) Proceeding straight (27.27%) Slowing/stopping (8.66%) Making right turn (7.79%) Making left turn (7.36%) Backing (3.90%) Changing lanes (2.60%) Parked (2.60%) Parking manoeuvre (2.16%) Merging (0.87%) Passing another vehicle (0.43%) Entering traffic (0.43%)
Colliding Object's Movement	The movement of the object prior to it colliding with the AV (vehicle-type objects only)	Proceeding straight (44.16%) Changing lanes (9.96%) Making right turn (6.93%) Backing (3.90%) Making left turn (3.46%) Parking manoeuvre (3.46%) Passing another vehicle (4.33%) Stopped (2.60%) Parked (2.60%) Other unsafe turning (2.16%) Entering traffic (1.73%) Slowing/stopping (1.73%) Other (1.30%) Crossing into opposing lane (0.87%) Merging (0.43%) Traveling the wrong way (0.43%)
AV Mode	Whether the AV is in autonomous mode or conventional mode (i.e., disengaged) at the time of the collision	Autonomous (43.29%) Disengaged (56.71%)
Weather	The weather at the time of the collision	Clear (87.88%) Cloudy (9.09%) Raining (2.16%) Fog/limited visibility (0.87%)
Lighting	The light conditions at the time of the collision	Daylight (72.29%) Dark with streetlights (24.24%) Dusk/dawn (3.46%)
Roadway Surface	The roadway's surface where the collision occurred	Dry (92.77%) Wet (4.33%) Unknown (3.90%)
Roadway Condition	The roadway's condition where the collision occurred	No unusual condition (87.45%) Unknown (10.39%) Construction repair zone (0.87%) Reduced roadway width (0.87%) Obstruction on roadway (0.43%)
Collision Type	The type of collision	Rear end (57.58%)

		Side swipe (21.65%) Broadside (9.09%) Hit object (4.33%) Head-on (3.46%) Other (3.46%) Unknown (0.43%)
Left Rear Corner Damaged	The AV's left rear corner was damaged in the collision	Yes (28.57%) No (71.43%)
Rear Bumper Damaged	The AV's rear bumper was damaged in the collision	Yes (31.60%) No (68.40)
Right Rear Corner Damaged	The AV's right rear corner was damaged in the collision	Yes (26.41%) No (73.59%)
Left Rear Passenger Side Damaged	The AV's left rear passenger side was damaged in the collision	Yes (2.60%) No (97.40%)
Right Rear Passenger Side Damaged	the AV's right rear passenger side was damaged in the collision	Yes (5.63%) No (94.37%)
Front Driver Side Damaged	the AV's front driver side was damaged in the collision	Yes (8.23%) No (91.77%)
Left Front Corner Damaged	The AV's left front corner was damaged in the collision	Yes (13.42%) No (86.58%)
Front Bumper Damaged	The AV's front bumper was damaged in the collision	Yes (6.49%) No (93.51%)
Front Passenger Side Damaged	The AV's front passenger side was damaged in the collision	Yes (7.36%) No (92.64%)
Right Front Corner Damaged	the AV's right front corner was damaged in the collision	Yes (9.09%) No (90.91%)
Damage Level	The AV's overall damage level from the collision	Minor (72.73%) Moderate (15.15%) None (10.82%) Major (1.30%)

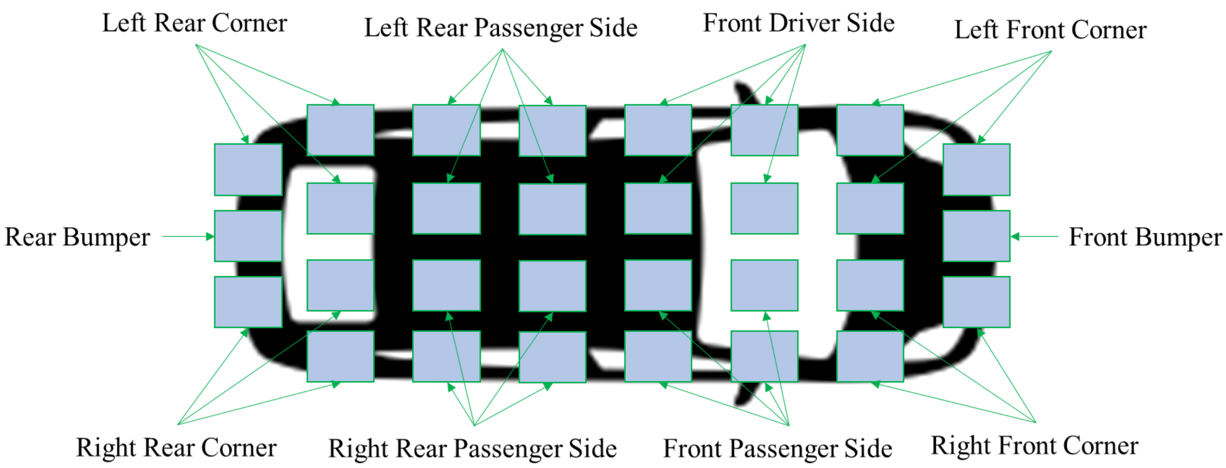

Figure 2. The areas of the AV that could be damaged in collisions

3.2 Association Rule Generation

The framework developed herein uses the Apriori algorithm to generate association rules. The AV's damage level was set as the consequent part of the rule (i.e., *C*), while the other variables formed the antecedent part of the rule (i.e., *A*). In this study, the thresholds for support, confidence, and confidence ratio were set to 5%, 10%, and 10%, respectively. A low confidence threshold value was required when the confidence ratio was used to measure rule interestingness, and a rule could be potentially interesting if adding an item to its antecedent caused a significant change in the confidence. Consequently, it was possible for an interesting rule to have a low confidence if adding an item drastically lowered the confidence. To allow these rules to be generated by the algorithm, a low confidence threshold value was used. Furthermore, the number of antecedent items was limited to three to ensure ease of interpretability, and the association rules with a lift smaller than one were excluded from the framework.

Cross-validation was necessary for ensuring the generalisability of the association rules. The data could have been partitioned into two subsets (i.e., train and test) or three (i.e., train, test, and validation). We chose the former approach because the latter would have reduced the size of each subset and is suitable only when the dataset is huge. The k-fold cross-validation with k=5 was used for this study. It started with randomly splitting the data into five subsets of equal size. Then, the Apriori algorithm was used to generate a set of association rules using the training data (i.e., k-1 subsets) and then evaluate the rules using the testing set (i.e., one subset). Each subset acted as the testing set only once, resulting in a total of five association rule sets.

3.3 Results

The five-fold cross-validation was repeated three times (i.e., i = 3) to remove the resampling bias. For each value of k, the number of training rules and the subsequent valid rules in the test set were obtained (see table 2). In each iteration, the valid rules were those that existed in all k values. In the first iteration, seventy-seven valid rules were obtained. In the second and the third iterations, fifty-seven and sixty-seven valid rules were obtained, respectively, with the final valid rules being those that existed in all three iterations, of which there were twenty-one. The performance measures of each of these rules were determined as the averages of the values. In particular, we calculated each rule’s average support, average confidence, average lift, and average confidence ratio.

Table 2. The number of validated association rules in the five-fold cross-validation with three iterations

	k = 1	k = 2	k = 3	k = 4	k = 5	Valid rules that exist in all k	Valid rules that exist in all i
i = 1	409 out of 912	447 out of 797	507 out of 887	405 out of 888	586 out of 912	77	21
i = 2	389 out of 894	447 out of 870	371 out of 857	341 out of 915	603 out of 874	57	
i = 3	438 out of 894	453 out of 962	464 out of 748	475 out of 898	377 out of 869	67	

4. Results and Discussion

4.1 Association Rules

Tables 3 to 5 show the final framework that contains a total of twenty-one association rules. Each table lists the rules from the most important to the least important based on confidence. The dataset contains very few (< 10%) AV collisions that resulted in major damage. ‘Damage Level = Major’ had a support value that was smaller than the threshold value, resulting in it not being included in any association rule. As a result, no rules with ‘Damage Level = Major’ were generated in the study. Therefore, this discussion focuses on the rules related to the other damage levels.

‘Damage Level = Minor’ was present in ten rules (see Table 3), with ‘Right Rear Corner Damaged = Yes’ being the most common occurrence (seven out of ten rules), followed by ‘AM=Yes’ (five out of ten rules). Interestingly, when ‘Time = AM’ was present, ‘Right Rear Corner Damaged = Yes’ was also present (i.e., Rules 1, 4, 7, 8, and 10). When these two items co-occurred, the confidence of the rules was at least 85.6%. Therefore, the probability of minor damage is high when the collision happens in the morning and the damage area is the AV’s right-rear corner. The probability increases to 95.83% if the lighting condition is daylight (i.e., Rule 1). In addition, ‘AV Mode = Autonomous’ and ‘Right Rear Corner Damaged = Yes’ occurred together and formed part of the antecedent of two association rules (i.e., Rules 6 and 9), with the rule confidences being 87.78% and 86.33% in the presence of ‘Lighting = Daylight’ and ‘Weather = Clear’, respectively. This small difference in the confidence value means that ‘Lighting = Daylight’ and ‘Weather = Clear’ similarly impact the probability of minor damage in cases in which the AV is in autonomous mode at the time of the collision and in which the right rear corner of the AV is damaged because of the collision.

‘Damage Level = Moderate’ was present in ten rules (see Table 4). All the rules had low confidence, but they are potentially interesting because of the high confidence ratio. When the antecedent was empty and the consequence was ‘Damage Level = Moderate’, the rule’s confidence was equivalent to the support of ‘Damage Level = Moderate’, which was low because of the small number of collisions with moderate damage. However, when additional items were added to the rule’s antecedent, the rule confidence differed significantly, as indicated by the confidence ratio. Among the ten rules, ‘Lighting = Daylight’ and ‘Left Rear Corner Damaged = Yes’ co-occurred five times (i.e., Rules 11, 12, 13, 15, and 17). Furthermore, when the collision type was a side swipe, the damage level was moderate, with a confidence of 23.90% (i.e., Rule 18). The confidence remained the same regardless of the presence of dry roadways (i.e., Rule 19). In addition, there were two rules that included ‘AV’s Movement = Proceeding straight’ as part of the antecedent (i.e., Rules 14 and 16). If the AV was proceeding straight on a dry roadway prior to the collision, the resultant damage level was moderate, with a confidence of 25.39%, and if the AV was proceeding straight prior to the collision when the weather was clear, then the resultant damage level was moderate, with a confidence of 23.98%.

Only one rule included ‘Damage Level = None’ as the consequent (see Table 5). If the roadway condition was unknown and the collision type was a rear end, then the damage level was zero, with a confidence of 12.78%. The high support (i.e., 52.77%) indicates that many rear-end collisions happened on roadways with unknown conditions. As collisions with no damage were a minority, the probability of having ‘Damage Level = None’ with an empty antecedent was low. Yet, the probability differed significantly when both ‘Roadway Condition = No unusual condition’ and ‘Collision Type = Rear end’ were present.

Consistent with previous studies (e.g., Montella 2011; Xu et al. 2018), this study determined that traffic accidents, such as AV collisions, are the result of complex interactions among multiple factors, including the weather, lighting conditions, and the mode in which the AV is operating at the time of the accident. It is therefore appropriate to focus on the association patterns among factors instead of the impact of individual factors when investigating AV safety. Although there are existing studies using association analysis to explain traffic accidents involving conventional vehicles (Lee and Wu 2021), their findings are not applicable to the AV context due to the fundamental differences between conventional vehicles and AVs. For instance, drivers performing secondary tasks (e.g., talking on their phone) is a factor leading to traffic accidents in the context of conventional vehicles (Kong et al. 2021), but it will not have the same effect in the context of AVs. In addition, there are other factors (e.g., whether the vehicle operates in autonomous mode or not) that are unique to AVs.

Table 3. Association rules with the consequent item of 'Damage Level = Minor'

No.	Antecedent	Support (%)	Confidence (%)	Confidence ratio (%)	Lift
1	'Right Rear Corner Damaged = Yes' and 'AM = Yes' and 'Lighting = Daylight'	8.09	95.83	23.55	1.32
2	'Roadway Condition = Unknown' and 'Object Collided = Passenger vehicle'	7.80	91.22	20.27	1.26
3	'Roadway Condition = Unknown' and 'Lighting = Daylight' and 'Weather = Clear'	7.38	91.21	21.74	1.26
4	'Right Rear Corner Damaged = Yes' and 'AM = Yes' and 'Weather = Clear'	9.36	89.59	20.41	1.24
5	'Left Front Corner Damaged = Yes' and 'Lighting = Daylight' and 'Object Collided = Passenger vehicle'	7.52	88.11	23.46	1.21
6	'Right Rear Corner Damaged = Yes' and 'AV Mode = Autonomous' and 'Lighting = Daylight'	7.23	87.78	20.20	1.21
7	'Right Rear Corner Damaged = Yes' and 'AM = Yes'	11.06	86.93	15.91	1.20
8	'Right Rear Corner Damaged = Yes' and 'AM = Yes' and 'Object Collided = Passenger vehicle'	10.21	86.37	15.18	1.19
9	'Right Rear Corner Damaged = Yes' and 'AV Mode = Autonomous' and 'Weather = Clear'	10.21	86.33	20.15	1.19
10	'Right Rear Corner Damaged = Yes' and 'AM = Yes' and 'Roadway Condition = No unusual condition'	10.21	85.60	14.23	1.18

Table 4. Association rules with the consequent item of 'Damage Level = Moderate'

No.	Antecedent	Support (%)	Confidence (%)	Confidence ratio (%)	Lift
11	'Left Rear Corner Damaged = Yes' and 'Lighting = Daylight' and 'Roadway Condition = No unusual condition'	20.57	25.85	38.57	1.64
12	'Left Rear Corner Damaged = Yes' and 'Lighting = Daylight' and 'Object Collided = Passenger vehicle'	20.57	25.84	39.50	1.66
13	'Left Rear Corner Damaged = Yes' and 'Lighting = Daylight' and 'Roadway Surface = Dry'	20.57	25.49	39.43	1.65
14	'AV's Movement = Proceeding straight' and 'Roadway Surface = Dry'	23.83	25.39	35.98	1.65
15	'Left Rear Corner Damaged = Yes' and 'Lighting = Daylight' and 'Weather = Clear'	19.72	24.92	37.85	1.63
16	'AV's Movement = Proceeding straight' and 'Weather = Clear'	23.55	23.98	35.49	1.58
17	'Left Rear Corner Damaged = Yes' and 'Lighting = Daylight'	22.27	23.90	36.73	1.53
18	'Collision Type = Side swipe'	21.56	20.76	28.13	1.38
19	'Collision Type = Side swipe' and 'Roadway Surface = Dry'	21.56	20.76	28.13	1.38
20	'AV Mode = Disengaged' and 'Roadway Surface = Dry'	53.05	19.69	22.13	1.25

Table 5. Association rules with the consequent item of 'Damage Level = None'

No.	Antecedent	Support (%)	Confidence (%)	Confidence ratio (%)	Lift
21	'Collision Type = Rear end' and 'Roadway Condition = No unusual condition'	52.77	12.78	23.46	1.25

4.2 Practical Implications

Policymaking on AV development requires a well-rounded perspective of the societal conditions from implementation to all of the possible consequences that a smart transport system may bring (Marchau et al. 2019). Additionally, a comprehensive transport policy that details a set of actions for governing AV systems is required. Thus, this study provides practical implications for reviewing and establishing AV-related policies on AV data management, traffic scenario definitions for testing, transport infrastructure, and legal issues surrounding AV collisions.

Firstly, our findings shed light on the standardisation of AV collision records for traffic accident analysis. Currently, AV development is in its infancy, so there is a lack of rich and robust AV collision data. Even though AV manufacturers may have their own databases for collision reports during AV testing, the factors considered by different manufacturers are likely different (Lee and Wu 2021). This study successfully identified the important factors in AV collisions that can be used to standardise AV crash records for accident analysis. Specifically, many of the generated rules indicate that daylight and clear weather co-occur in AV collisions. This could be due to imbalanced samples as the majority of the collisions we analysed occurred during the daytime and in clear weather. This could be due to the current on-road tests that are largely conducted during a certain period of time, such as AV test drivers' working hours and when the weather is good. Consequently, AV performance in some scenarios, such as during bad weather or at night, is still unknown. Therefore, AV manufacturers should develop a more comprehensive strategy that allows AV testing to be done in a wider range of situations, such as at different times and in various weather conditions. This is important because an AV system should be able to respond to all types of situations, including unanticipated ones, in real driving conditions (Zavala et al. 2021). Moreover, the policymakers should revisit the existing procedures for on-road testing as necessary safety measures must currently be adopted for on-road testing, but when the test is done in higher-risk scenarios, such as at night with low visibility, the existing safety measures may need to be enhanced.

In addition, our findings provide supplementary information for AV manufacturers when they are designing traffic scenarios for AV testing and for transport authorities when they are determining benchmark traffic scenarios for assessing newly manufactured AVs' driving performance. Specifically, the antecedents of the rules denote the traffic scenarios that commonly occur when AVs are driving on roads and involved in collisions. Since our study has determined twenty-one rules, the same number of traffic scenarios can be established. Moreover, the consequents of the rules can serve as a baseline of AV performance. For instance, if the damage level being moderate is the baseline, then the AV's performance can be considered better if the resultant damage level is less than this in the simulation results. Thus, authorities can use AV performance as the basis when designing an AV safety assurance program. This can result in AVs that demonstrate a good safety track record and that have capabilities to ensure road safety being considered eligible for longer license durations and higher award tiers (e.g., Gold, Silver, and Bronze). There are two potential benefits to this: When AVs are widely adopted in future, customers can refer to the award when choosing an AV for use, and programmes that include grading assessments can motivate AV manufacturers to maintain safety standards and provide better assurance to customers and other road users.

Amongst the twenty-one rules generated from the framework, twelve indicated a damaged rear corner, and one indicated a damaged rear end from collisions. Rear-end collisions can be attributed to improper manoeuvres or short following distances between vehicles. To avoid these rear-end collisions, AV manufacturers can perform more scenario-based AV tests to enhance an AV's ability to manoeuvre as well as to improve the AV sensors' sensitivity in terms of detecting the speed and distance of incoming vehicles or objects. A universal or national standard for the minimum safe following distance between an AV and another vehicle can be established by transport authorities. Furthermore, regarding future road infrastructure that accommodates AVs, policymakers can consider setting a minimum road width that allows AVs to manoeuvre safely as well as restricting the number of vehicles that can access roads that are open to AVs during peak hours.

Rule 20 indicates that AV collisions happen when the AV disengages, which is the rule with the highest support among all the generated rules, revealing the high possibility of an AV operating in manual mode on a dry roadway. Based on the narratives of the collisions contained in the reports, there are at least two possible situations in which AVs operate in a manual mode: The first scenario is the driver having driven the AV manually for a while before the collision, and the second scenario is the driver disengaging the autonomous mode when they are trying to avoid a collision. Whether the AV driver is legally responsible for the collision remains debatable. For instance, in the second scenario, it is unclear whether the driver might have done the right thing (e.g., applying the brake manually) or whether the accident could have been avoided if the driver did not do so. This happened in one of Google's AVs, when it was hit from behind after the Google employee manually braked when the AV had already started slowing for a pedestrian. It is doubtful whether this rear-ending could have been prevented if the employee did not apply brakes as it may have given the car behind more time to stop (Richtel and Dougherty 2015). Therefore, transport authorities can establish a governance framework that defines the legal responsibility in AV collisions, with the disengagement time prior to the AV collision being an important factor to consider.

5. Conclusion

Extensive efforts have been made by researchers, industrial practitioners, and public authorities to make AVs drive more like humans. However, the potential of machine learning approaches in AV collision analysis has rarely been discussed. This paper creates a new, promising direction on which future studies related to AV can focus as machine learning approaches can be used to improve not only the driving capability of AVs but also the evaluation of AV safety. They can also be used as a knowledge discovery tool to generate useful insights that can be incorporated into the ongoing creation of AV-related policies.

The machine learning framework proposed in this study was used to analyse a set of real AV collision reports and extract collision patterns using association rule mining. The resulting explainable association rules provide valuable insights for AV policymakers and stakeholders when attempting to understand the key factors associated with AV collisions. Additionally, based on the rules' antecedents, AV developers can establish a set of traffic scenarios for AV testing. Unlike previous studies, this study used confidence ratio instead of confidence to ensure the interestingness of the rules that are beneficial to the development of AVs.

This study has two limitations regarding data collection and analysis. Firstly, this study did not consider the behaviour of AV users upon collision. Thus, the action that AV users should and should not engage in to maintain safety when the AVs are in motion remains unexplored. Therefore, user data, such as the user choosing where they sit in the AV, whether the user has a driving license, and the tasks that the user is performing at the time of collision, should be included in future collision reports to ensure more comprehensive analysis that can be used to drive evidence-based policies for road safety. Additionally, the extent of users' freedom in terms of multitasking safely while on the move is another interesting question that remains to be explored. Secondly, the dataset used in this study contained only a few collision cases that resulted in major damage. Consequently, our results contain no rules that can be used to identify the factors that commonly occur when an AV is majorly damaged in a collision. From a safety perspective, it would be useful to understand collisions that cause major damage. Thus, future research could analyse the minority of AV collisions in which the damage is severe. Furthermore, analytical tools that are suitable for small sample sizes, such as qualitative comparative analysis, can be used to investigate the complex relationships among risk factors.

Disclosure Statement

There are no relevant financial or non-financial competing interests to report.

References

Acheampong, R. A., F. Cugurullo, M. Gueriau, and I. Dusparic. 2021. "Can Autonomous Vehicles Enable Sustainable Mobility in Future Cities? Insights and Policy Challenges from User Preferences over Different Urban Transport Options." *Cities* 112, 103134. doi:10.1016/j.cities.2021.103134.

Agrawal, R., and R. Srikant. 1994. "Fast Algorithms for Mining Association Rules in Large Databases." Proceedings of *20th international conference on very large databases*, Santiago de Chile (pp. 487–489).

Anthony Jnr, B. 2021. "Managing Digital Transformation of Smart Cities through Enterprise Architecture – A Review and Research Agenda." *Enterprise Information Systems* 15 (3): 299-331. doi:10.1080/17517575.2020.1812006.

Chang, L.-Y., and W.-C. Chen. 2005. "Data Mining of Tree-based Models to Analyze Freeway Accident Frequency." *Journal of Safety Research* 36 (4): 365–375. doi:10.1016/j.jsr.2005.06.013.

Coelingh, E., A. Eidehall, and M. Bengtsson. 2010, September. "Collision Warning with Full Auto Brake and Pedestrian Detection: A Practical Example of Automatic Emergency Braking." 13th International IEEE Conference on Intelligent Transportation Systems (pp. 155–160). Madeira Island, Portugal: IEEE.

Creutzig, F., P. Jochem, O. Y. Edelenbosch, L. Mattauch, D. P. van Vuuren, D. McCollum, and J. Minx. 2015. "Transport: A Roadblock to Climate Change Mitigation?" *Science* 350 (6263): 911–912. doi: 10.1126/science.aac8033.

Das, S. 2021. "Autonomous Vehicle Safety: Understanding Perceptions of Pedestrians and Bicyclists." *Transportation Research Part F: Traffic Psychology and Behaviour* 81: 31-54. doi:10.1016/j.trf.2021.04.018.

Deluka Tibljaš, A., T. Giuffrè, S. Surdonja, and S. Trubia. 2018. "Introduction of Autonomous Vehicles: Roundabouts Design and Safety Performance Evaluation." *Sustainability* 10 (4): 1060. doi:10.3390/su10041060.

Favarò, F. M., N. Nader, S. O. Eurich, M. Tripp, and N. Varadaraju. 2017. "Examining Accident Reports Involving Autonomous Vehicles in California." *PLoS one* 12 (9): e0184952. doi: 10.1371/journal.pone.0184952.

Ghanem, S., R. Couturier, and P. Gregori. 2021. "An Accurate and Easy to Interpret Binary Classifier Based on Association Rules Using Implication Intensity and Majority Vote." *Mathematics* 9 (12): 1315. doi: 10.3390/math9121315.

Griggs, T., and D. Wakabayashi. 2018. "How a Self-Driving Uber Killed a Pedestrian in Arizona." *The New York Times*, Accessed 13 March 2023. https://www.nytimes.com/interactive/2018/03/20/us/

self-driving-uber-pedestrian-killed.html?mtrref=www.google.com&gwh=223F48EE20B4FA42F5D4E5BABDE4E2A4&gwt=regi&assetType=REGIWALL

Huang, W., K. Wang, Y. Lv, and F. Zhu. 2016, November. "Autonomous Vehicle Testing Methods Review." *IEEE 19th International Conference on Intelligent Transportation Systems* (pp. 163–168). IEEE. Rio de Janeiro, Brazil.

Jiang, F., K. K. R. Yuen, and E. W. M. Lee. 2020. "Analysis of Motorcycle Accidents Using Association Rule Mining-based Framework with Parameter Optimization and GIS Technology." *Journal of Safety Research* 75: 292–309. doi:10.1016/j.jsr.2020.09.004.

Kalra, N., and S.M. Paddock. 2016. "Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?" *Transportation Research Part A: Policy and Practice* 94: 182–193. doi:10.1016/j.tra.2016.09.010.

Kong, X., S. Das, and Y. Zhang. 2021. "Mining Patterns of Near-crash Events with and without Secondary Tasks." *Accident Analysis and Prevention* 157: 106–162. doi:10.1016/j.aap.2021.106162.

Lee, C.K.H., and K.Y.K. Wu. 2021. "Making Autonomous Vehicle Systems Human-like: Lessons Learned from Accident Experiences in Traffic." *Enterprise Information Systems*. doi:10.1080/17517575.2021.1998641.

Liuima, J. 2020. "Coronavirus Increases Demand for Autonomous Vehicles." Accessed 13 March 2023. https://blog.euromonitor.com/coronavirus-increases-demand-for-autonomous-vehicles

Marchau, V., J. Zmud, and N. Kalra. 2019. "Editorial for the Special Issue–Autonomous Vehicle Policy." *Transportation Research Part A: Policy and Practice* 122: 120-124. doi: 10.1016/j.tra.2018.04.017

Markkula, G., O. Benderius, K. Wolff, and M. Wahde. 2012. "A Review of Near-Collision Driver Behavior Models." *Human Factors* 54 (6): 1117–1143. doi:10.1177/0018720812448474.

Maryam Mousavi, S., D. Lord, B. Dadashova, and S. Reza Mousavi, S. 2020. "Can Autonomous Vehicles Enhance Traffic Safety at Unsignalized Intersections?" In *International Conference on Transportation and Development 2020* (pp. 194-206). Reston, VA: American Society of Civil Engineers.

Montella, A. 2011. "Identifying Crash Contributory Factors at Urban Roundabouts and Using Association Rules to Explore Their Relationships to Different Crash Types." *Accident Analysis and Prevention* 43 (4): 1451–1463. doi:10.1016/j.aap.2011.02.023.

Papadoulis, A., M. Quddus, and M. Imprialou. 2019. "Evaluating the Safety Impact of Connected and Autonomous Vehicles on Motorways." *Accident Analysis & Prevention* 124: 12–22. doi:10.1016/j.aap.2018.12.019

Perveen, S., T. Yigitcanlar, M. Kamruzzaman, and J. Hayes. 2017. "Evaluating Transport Externalities of Urban Growth: A Critical Review of Scenario-based Planning Methods." *International Journal of Environmental Science and Technology* 14: 663–678. doi:10.1007/s13762-016-1144-7.

Richtel, M., and C. Dougherty. 2015. "Google's Driverless Cars Run Into Problems: Cars With Drivers." *The New York Times*. Accessed 13 March 2023. https://www.nytimes.com/2015/09/02/technology/personaltech/google-says-its-not-the-driverless-cars-fault-its-other-drivers.html#:~:text=MOUNTAIN%20VIEW%2C%20Calif.,-%E2%80%94%20Google%2C%20a%20leader&text=Researchers%20in%20the%20fledgling%20field,t%20behave%20by%20the%20book

Rosique, F., P.J. Navarro, C. Fernández, and A. Padilla. 2019. "A Systematic Review of Perception System and Simulators for Autonomous Vehicles Research." *Sensors* 19 (3): 168. doi:10.3390/s19030648.

Schoettle, B., and M. Sivak. 2015. "A Preliminary Analysis of Real-world Crashes Involving Self-driving Vehicles." *University of Michigan Transportation Research Institute*.

Seuwou, P., E. Banissi, and G. Ubakanma. 2020. "The Future of Mobility with Connected and Autonomous Vehicles in Smart Cities." In *Digital Twin Technologies and Smart Cities* (pp. 37-52). Springer, Cham.

Suman, M. N. H., F. A. Chyon, and M. S. Ahmmed. 2020. "Business Strategy in Bangladesh – Electric Vehicle SWOT-AHP Analysis: Case Study." *International Journal of Engineering Business Management* 12: 1–10. doi:10.1177/1847979020941487.

Taeihagh, A., and H. S. M. Lim. 2019. "Governing Autonomous Vehicles: Emerging Responses for Safety, Liability, Privacy, Cybersecurity, and Industry Risks." *Transport Reviews* 39 (1): 103–128. doi:10.1080/01441647.2018.1494640.

Teoh, E. R., and D. G. Kidd. 2017. "Rage against the Machine? Google's Self-driving Cars versus Human Drivers." *Journal of Safety Research* 63: 57-60. doi: 10.1016/j.jsr.2017.08.008.

Whalen, L. 2022. "For Autonomous Vehicles, Safety and Trust Remain Major Obstacles." Accessed 25 August 2022. https://morningconsult.com/2022/04/06/autonomous-vehicles-safety-and-trust/

Xu, C., J. Bao, C. Wang, and P. Liu. 2018. "Association Rule Analysis of Factors Contributing to Extraordinarily Severe Traffic Crashes in China." *Journal of Safety Research* 12: 65–75. doi:10.1016/j. jsr.2018.09.013.

Xu, C., Z. Ding, C. Wang, and Z. Li. 2019. "Statistical Analysis of the Patterns and Characteristics of Connected and Autonomous Vehicle Involved Crashes." *Journal of Safety Research* 71: 41-47. doi: 10.1016/j.jsr.2019.09.001.

Ye, L., and T. Yamamoto. 2019. "Evaluating the Impact of Connected and Autonomous Vehicles on Traffic Safety." *Physica A: Statistical Mechanics and its Applications* 526: 121009. doi:10.1016/j.physa.2019.04.245.

Yigitcanlar, T., M. Kamruzzaman, L. Buys, G. Ioppolo, J. Sabatini-Marques, E. Costa, and J. Yun. 2018. "Understanding 'Smart Cities': Intertwining Development Drivers with Desired Outcomes in a Multidimensional Framework." *Cities* 81: 145–160. doi: 10.1016/j.cities.2018.04.003.

Zavala, E., X. Fanch, J. Marco, and C. Berger. 2021. "Adaptive Monitoring for Autonomous Vehicles Using the HAFLoop Architecture." *Enterprise Information Systems* 15 (2): 270–298. doi:10.1080/ 17517575.2020.1844305.

Zhang, T., D. Tao, X. Qu, X. Zhang, R. Lin, and W. Zhang. 2019. "The Roles of Initial Trust and Perceived Risk in Public's Acceptance of Automated Vehicles." *Transportation Research Part C: Emerging Technologies* 98: 207–220. doi:10.1016/j.trc.2018.11.018.

Zhu, W., D. Gao, W. Zhao, H. Zhang, and H.-P. Chiang. 2018. "SDN-enabled Hybrid Emergency Message Transmission Architecture in Internet-of-vehicles." *Enterprise Information Systems* 12 (4): 471–491. doi:10.1080/17517575.2017.1304578.