

A unified probabilistic approach to traffic conflict detection

Jiao, Yiru; Calvert, Simeon C.; van Cranenburgh, Sander; van Lint, Hans

DOI

[10.1016/j.amar.2024.100369](https://doi.org/10.1016/j.amar.2024.100369)

Publication date

2024

Document Version

Final published version

Published in

Analytic Methods in Accident Research

Citation (APA)

Jiao, Y., Calvert, S. C., van Cranenburgh, S., & van Lint, H. (2024). A unified probabilistic approach to traffic conflict detection. *Analytic Methods in Accident Research*, 45, Article 100369. <https://doi.org/10.1016/j.amar.2024.100369>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy


Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



ELSEVIER

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Analytic Methods in Accident Research

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

A unified probabilistic approach to traffic conflict detection

Yiru Jiao^{a,c,*}, Simeon C. Calvert^{a,c}, Sander van Cranenburgh^{b,c}, Hans van Lint^a^a Department of Transport & Planning, Delft University of Technology, Delft, the Netherlands^b Department of Engineering Systems and Services, Delft University of Technology, Delft, the Netherlands^c CityAI lab, Delft University of Technology, Delft, the Netherlands

ARTICLE INFO

Keywords:

Traffic safety
Autonomous driving safety
Conflict detection
Collision risk
Surrogate Measures of Safety

ABSTRACT

Traffic conflict detection is essential for proactive road safety by identifying potential collisions before they occur. Existing methods rely on surrogate safety measures tailored to specific interactions (e.g., car-following, side-swiping, or path-crossing) and require varying thresholds in different traffic conditions. This variation leads to inconsistencies and limited adaptability of conflict detection in evolving traffic environments, particularly as the integration of autonomous driving systems adds complexity. Consequently, there is an increasing need for consistent detection of traffic conflicts across interaction contexts. To address this need, we propose a unified probabilistic approach in this study. The proposed approach establishes a unified framework of traffic conflict detection, where traffic conflicts are formulated as context-dependent extreme events of road user interactions. The detection of conflicts is then decomposed into a series of statistical learning tasks: representing interaction contexts, inferring proximity distributions, and assessing extreme collision risk. The unified formulation accommodates diverse hypotheses of traffic conflicts and the learning tasks enable data-driven analysis of factors such as motion states of road users, environment conditions, and participant characteristics. Jointly, this approach supports consistent and comprehensive evaluation of the collision risk emerging in road user interactions. We demonstrate the proposed approach by experiments using real-world trajectory data. A unified metric for indicating conflicts is first trained with lane-change interactions on German highways, and then compared with existing metrics using near-crash events from the U.S. 100-Car Naturalistic Driving Study. Our results show that the unified metric provides effective collision warnings, generalises across distinct datasets and traffic environments, covers a broad range of conflict types, and captures a long-tailed distribution of conflict intensity. In summary, this study provides an explainable and generalisable approach that enables traffic conflict detection across varying interaction contexts. The findings highlight its potential to enhance the safety assessment of traffic infrastructures and policies, improve collision warning systems for autonomous driving, and deepen the understanding of road user behaviour in safety-critical interactions.

1. Introduction

Collision avoidance and accident prevention are key elements in efforts to improve traffic safety (Vahidi and Eskandarian, 2003; Zheng et al., 2021). Motivated by proactive prevention of accidents and not waiting for collisions to happen, Surrogate Measures of

* Corresponding author at: Department of Transport & Planning, Delft University of Technology, Delft, the Netherlands.
E-mail address: y.jiao-1@tudelft.nl (Y. Jiao).

<https://doi.org/10.1016/j.amar.2024.100369>

Received 4 September 2024; Received in revised form 14 November 2024; Accepted 17 December 2024

Available online 20 December 2024

2213-6657/© 2024 The Author(s).

Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Published by Elsevier Ltd. This is an open access article under the CC BY license

Safety were proposed as surrogates of real collisions for safety evaluation and improvement. Over the past decades, traffic conflicts have become one of the most comprehensive and prominent surrogates (Zheng et al., 2014b; Tarko, 2018). A traffic conflict is defined as “an observable situation in which two or more road users approach each other in space and time to such an extent that there is a risk of collision if their movements remain unchanged” (Cooper, 1977). Under this definition, every conflict is a potential collision; and every collision is a conflict until the moment when it becomes unavoidable. Despite the relative rareness when compared with safe daily interactions, successfully resolved conflicts that do not end in collisions offer opportunities to explore the emergence and resolution of collision risk (Tarko, 2018). This has been widely recognised in the traffic safety community, and is being adopted in emerging technologies such as driving assistance and autonomous driving.

Conflicts cannot always be directly measured due to the unclear boundary between safe and unsafe interactions, thus conflict detection often relies on surrogate metrics. These metrics are called “conflict/safety indicators” in the field of traffic safety (Vogel, 2003) or “criticality metrics” in the field of autonomous driving (Junietz et al., 2018). This difference in terminology implies different focuses on the scale of their respective research objective. Traffic safety studies concentrate more on reasoning about the causes of collisions and reducing accidents in general; while autonomous driving studies aim to ensure safe interactions between automated vehicles and other road users (Razi et al., 2022). As a result, conflict indicators are predominantly used for identifying contributing factors of collisions (Saunier et al., 2011; Wu and Jovanis, 2012) and evaluating the safety of transportation infrastructure, a traffic signal system, or a specific type of interaction (Saunier et al., 2007; Laureshyn et al., 2010; de Ceunynck, 2017). In contrast, criticality metrics are more used for predicting trajectories of other road users (Wang et al., 2022), developing autonomous driving strategies (Das and Won, 2021), and assessing individual collision risk (Broadhurst et al., 2005; Kim and Kum, 2018; Zhou and Zhong, 2020). In this paper, we use *surrogate metrics of conflicts* to uniformly refer to both conflict/safety indicators and criticality metrics.

A variety of surrogate metrics of conflicts have been developed for different types of interactions. For instance, Deceleration Rate to Avoid Collision (DRAC, Cooper and Ferguson, 1976) and its variants are primarily targeted at rear-end conflicts. Time-To-Collision (TTC, Hayward, 1972; Hydén, 1987) along with its variants can cover both rear-end and side-swipe conflicts. Time advantage (Laureshyn et al., 2010), also known as predicted Post-Encroachment-Time (PET), is specifically tailored for path-crossing conflicts. In addition, composite indices are designed by integrating multiple metrics to deal with more complicated conflicts such as those during lane changes (Uno et al., 2002; Li et al., 2017). Many summaries (to name a few, Arun et al., 2021; Lu et al., 2021; Wang et al., 2021; Arun et al., 2021; Westhofen et al., 2022) are available for an overview of these metrics, and new metrics are being actively proposed for interactions in a two-dimensional plane (Ward et al., 2015; Venthuruthiyil and Chunchu, 2022) and in traffic oscillations (Kuang et al., 2015; Wang et al., 2024).

The diversity of these metrics entails inconsistency as collision risk is heterogeneous in different traffic conditions, between different vehicles, and is subject to changes in road user behaviour (Manning et al., 2016; Zheng et al., 2021). Some experiments show that the perception of critical TTC can vary among drivers (Kusano et al., 2015) and in different traffic environments (Tageldin and Sayed, 2019; Arun et al., 2021; Chauhan et al., 2022). For example, a TTC of 3 s could be dangerous for vehicles rushing on highways, but not necessarily for vehicles making a cooperative lane-change, nor for vehicles decelerating to approach an urban intersection. Similarly, a 2-s PET could be accident-prone for cars crossing their paths at an intersection, but is not uncommon for cyclists (Beitel et al., 2018). In addition, with the increasing prevalence of driving assistance systems and electric vehicles, behavioural changes in interaction may gradually influence people’s perception of collision risk (Das et al., 2022; Wessels et al., 2022).

Increasing efforts are devoted to overcoming these inconsistencies. For example, there is a growing trend towards combining multiple metrics (e.g., Kuang et al., 2015; Nadimi et al., 2021; Mazaheri et al., 2023), particularly by applying deep learning methods (Formosa et al., 2020; Abdel-Aty et al., 2023). Such a combination builds understanding from existing metrics, and thus remains constrained by their underlying assumptions of conflicts. Many studies have also looked at determining robust thresholds to distinguish unsafe conflicts from safe interactions. If there is ground truth, the threshold should optimise the accuracy of conflict identification for, e.g., issuing crash warnings (Das and Won, 2021). When no label is available, the threshold selection often follows heuristic rules (Panou, 2018; Das and Maurya, 2020; Nadimi et al., 2022), or is guided by the extreme value theory modelling to satisfy certain hypotheses (Farah and Azevedo, 2017; Borsos et al., 2020; Niu et al., 2024).

However, as the development of autonomous driving rapidly advances, the inconsistency in conflict detection poses new challenges for traffic safety. Automated vehicles require a unified approach to estimating collision risk as they navigate in various road environments. Without consistent estimation, these vehicles may struggle to reliably interpret and respond to traffic conflicts, thereby jeopardising road safety. As different levels of vehicle automation increasingly share the road, traffic conditions will become more complex. Road user interactions will no longer be solely between humans, but will also involve driving algorithms developed by different manufacturers, creating new layers of complexity. Consequently, in the long run, the assessment and improvement of traffic safety will become more challenging.

To address the challenges, this study introduces a new approach for consistent and comprehensive conflict detection. First, we propose a unified framework of traffic conflict detection by formulating Bayesian collision risk. Second, we present a series of statistical learning tasks that apply the theoretical framework to practical conflict detection. Our approach considers a traffic conflict as an extreme event of normal interactions, and quantifies its context-dependent and proximity-characterised collision risk. Then the learning tasks break down conflict detection into interaction context representation, proximity distribution inference, and extreme event assessment. Theoretically, any existing surrogate metric of conflicts is a special case under the framework. With this approach, conflict detection can be consistent across different traffic environments, and traffic safety evaluation can involve more comprehensive considerations.

The rest of this paper is organised as follows. Section 2 first explains the unified framework of traffic conflict detection. Then Section 3 presents the statistical learning tasks to apply conflict detection, including probability estimation and conflict intensity

evaluation. In Section 4, demonstration experiments are designed, performed, and analysed to show the performance of the approach. Finally, Section 5 concludes this paper and envisions future research.

2. A unified framework of traffic conflict detection

This section introduces our unified framework of traffic conflict detection. We frame conflict detection as quantifying the risk of a potential collision c between two or more road users, based on their proximity s and other contextual observables X , i.e., $p(c|s, X)$. This risk quantification uses Bayes' theorem to consider the interaction situation of an event, the typical proximity behaviour of road users in the interaction context, and the probability variation with conflict intensity. These considerations are summarised in Eq. (1) for a preliminary overview, where the symbols are defined in Table 1.

$$p(c|s, X) = \iint p(c|s, \phi) p(\phi|\theta) p(\theta|X) d\phi d\theta, \text{ where } p(c|s, \phi) := C(n; s, \phi) \tag{1}$$

The Bayesian components of the integral in Eq. (1) represent a probabilistic breakdown of collision risk. Starting from the outermost component, the first term $p(\theta|X)$ denotes the probability of a specific interaction context θ given observable measurements X , such as road user motion states and environmental conditions. The second term $p(\phi|\theta)$ describes the probability of the parameters ϕ that characterise proximity patterns conditioned on the interaction context θ . Finally, the innermost term $p(c|s, \phi)$ expresses the probability of a potential collision occurring at a proximity s , based on the proximity distribution parameterised by ϕ in the specific context. This probability also accounts for conflict intensity n , which reflects the extremeness extent of the conflict within a spectrum of interactions in the same context.

The following subsections provide more detailed explanations and derivations to establish the framework. We begin with defining the probabilistic collision risk that depends on proximity and other situational variables. Then we assume a conflict hierarchy of risk perceptions and reactions in an interaction context, and describe the conflict hierarchy quantitatively by measuring how road user behaviour varies with proximity. Lastly, we consider conflicts as extreme events within a spectrum of interactions, and use extreme value theory to relate conflict probability with conflict intensity.

2.1. Context-dependent collision risk

In this subsection, we first present a general probabilistic description of conflict detection which depends on the context of interaction. A major branch of surrogate metrics of conflicts is based on proximity. Proximity means physically less space and time for reactions to prevent a potential collision; and due to the physical constraints, people perceive increased risk when the distance to an approaching object is closer (Schiff and Detwiler, 1979; Teigen, 2005). Therefore, a proximity-based conflict metric can cover both the objective aspect of collision risk and people's subjective perception of the collision risk.

Proximity-based conflict detection can be generally formulated as $p(c|s, X)$, which estimates the probability of a potential collision c based on proximity s ($s \geq 0$) and other observables X of the interaction context. The probabilistic formulation enables consistent evaluation across contexts, as probability has a normalised range between 0 and 1. This consideration has also been taken by some studies, for example, Saunier and Sayed (2008) and de Gelder et al. (2023). For an interaction involving two or more road users, s is a measure of the spatial-temporal closeness between the road users; and X can include, but is not limited to, motion states of the road users, traffic states, road layouts, and weather conditions.

Not all of the information in X is used, nor might it all be useful. Every existing surrogate metric of conflicts selects some key information, i.e., variables, and makes a hypothesis based on the selected information. For example, TTC uses the relative velocity between two approaching road users and assumes no movement change (constant acceleration) at the moment of conflict; DRAC also uses the relative speed but assumes an immediate brake. Here we use $\theta = \{\theta_1, \theta_2, \dots, \theta_k\}$ to denote k variables extracted from X , and θ are assumed to adequately represent the interaction context compressed out of the situation X . Considering all possible selections of θ , Eq. (2) holds according to Bayes' theorem.

$$p(c|s, X) = \int p(c|s, \theta) p(\theta|X) d\theta \tag{2}$$

If the selection of key information is deterministic as $\theta = \mathbb{R}(X)$, where \mathbb{R} refers to representation learning of the interaction context,

Table 1
Symbols and their definitions in the unified framework of traffic conflict detection.

Symbol	Definition
c	A potential collision, i.e., conflict
s	Proximity, the spatial-temporal closeness between road users
X	Observables that can be measured to describe interaction situations
θ	Representation for the interaction context involving key information selected from X
ϕ	Parameters used to characterise the proximity distribution of road users in a given θ
n	Conflict intensity, a conflict at intensity n occurs once per n times in the same interactions
C	Probability of a conflict with intensity n at proximity s in the context characterised by ϕ

$p(\theta|X)$ in Eq. (2) becomes a Dirac delta function of $\delta(\theta - \mathbb{R}(X))$. A Dirac delta function $\delta(x)$ has a value of 1 when $x = 0$ and 0 for any other values of x . Therefore, Eq. (2) can be approximated to Eq. (3).

$$p(c|s, X) = p(c|s, \theta) = p(c|s, \mathbb{R}(X)), \text{ if } p(\theta|X) = \delta(\theta - \mathbb{R}(X)) \tag{3}$$

To better explain how a surrogate metric of conflicts fits into our formulation, we provide two examples. The first is Post-Encroachment Time (PET), where s is the time interval between one vehicle leaving a conflict area and another vehicle arriving in the same conflict area. The θ for PET is the existence of a conflict area. As shown in Fig. 1 (a), $p(c|s, \theta)$ can be a step function with probability 1 when PET exceeds a threshold PET^* and 0 otherwise; or a continuous function based on a cumulative Gaussian probability that increases gradually from 0 to 1 and is 0.5 at the threshold PET^* . Then another example uses TTC = $s/\Delta v$ as shown in Fig. 1 (b). For two vehicles approaching each other, s is the net distance between them; while Δv is their relative speed, and serves as a univariate context θ . The probability $p(c|s, \theta)$ varies at different relative speeds. If there is a critical threshold TTC^* that differentiates safe and unsafe interactions, the threshold of proximity then follows $s^* = TTC^* \Delta v$.

2.2. Proximity-characterised conflict hierarchy

Proximity patterns in an interaction context reflect the aggregated behaviour of road users in this context, which is shaped by people’s perception of collision risk. Within the same context of interaction, a shorter spatial or temporal gap between road users consistently implies more risk of collision. This increase in collision risk with decreasing proximity motivates road users to maintain an acceptable distance from the others at varying levels (Camara and Fox, 2020). Therefore, the proximity behaviour of road users embodies a conflict hierarchy perceived by the road users. On the one hand, proximity pushes or prevents the road users from approaching each other; on the other hand, the road users adjust their behaviours to maintain comfortable proximity.

We characterise the proximity behaviour of road users with a conditional probability distribution. Seeing proximity as a random variable S , in a certain interaction context θ , the conditional probability distribution of S given a context θ is $p(s|\theta)$. We use $f_S(s|\theta; \phi)$ to describe the density function of $p(s|\theta)$, where ϕ is a set of parameters. To incorporate context-dependent proximity behaviour into conflict detection, we consider all possible sets of ϕ in a context θ , thus integrate $p(\phi|\theta)$ in Eq. (4) according to Bayes’ theorem.

$$p(c|s, \theta) = \int p(c|s, \phi)p(\phi|\theta)d\phi \tag{4}$$

As we parameterise the conditional probability $p(s|\theta)$ with $f_S(s|\theta; \phi)$, inferring $p(s|\theta)$ is to obtain ϕ , and we denote this inference as $\phi = \mathbb{I}(\theta)$. When the inference is deterministic, $p(\phi|\theta)$ in Eq. (4) becomes a Dirac delta function $\delta(\phi - \mathbb{I}(\theta))$. In the same way as explained when deriving Eq. (3), we can derive Eq. (5).

$$p(c|s, \theta) = p(c|s, \phi) = p(c|s, \mathbb{I}(\theta)), \text{ if } p(\phi|\theta) = \delta(\phi - \mathbb{I}(\theta)) \tag{5}$$

Numerous empirical studies have observed the transition from comfort to discomfort when a road user is approached by other vehicles (Taieb-Maimon and Shinar, 2001; Lewis-Evans et al., 2010; Siebert et al., 2014; Siebert et al., 2017). Therefore, we can generally assume that $p(c|s, \phi)$ monotonically increases while s is decreasing. In a conflict hierarchy characterised by ϕ , the closer the proximity between approaching road users, the less safe they are and the higher the probability of a potential collision. This monotonicity of $p(c|s, \phi)$ implies that Eq. (6) holds for any ϕ .




Fig. 1. Illustration examples of context-dependent proximity-based conflict probability.

$$\begin{cases} \lim_{s \rightarrow \infty} p(c|s, \phi) = 0 \\ \lim_{s \rightarrow 0} p(c|s, \phi) = 1 \end{cases} \tag{6}$$

We use Fig. 2 to help explain the proximity-characterised conflict hierarchy. On the left of Fig. 2, we present an adapted pyramid of conflict hierarchy (Hydén, 1987), which conceptually relates conflict intensity and frequency, and reflects road user behaviour in balancing safety and efficiency (Laureshyn et al., 2010). When the approaching between road users entails a potential collision, the closer they are, the higher the conflict intensity and the lower the frequency of such conflicts. When the proximity is too small for safe interaction, a conflict emerges entailing a potential collision. Furthermore, an accident will happen if there is no (successful) evasion or the proximity is too small to prevent a collision. This hierarchy of conflicts is not fixed, but has varying shapes in different interaction contexts. As illustrated in the right of Fig. 2, $f_S(s|\theta; \phi)$ is designed to capture such context-dependent proximity patterns. Consequently, the estimated probability $p(c|s, \phi = \mathbb{I}(\theta))$ may vary for the same proximity in different contexts, and evolve as the interaction context changes over time.

2.3. Extreme value theory-based interaction spectrum

Traffic interactions exist on a spectrum varying in conflict intensity and potential consequences. Within this spectrum, we assume that conflicts are extreme events at different levels of collision risk and collisions are the most extreme cases. Then we can establish the relation between conflict intensity and conflict probability utilising extreme value theory. In traffic safety research, the extreme value modelling has been traditionally used to treat crashes as the extreme events of conflicts and extrapolate average crash risk from traffic conflicts (Songchitruksa and Tarko, 2006; Zheng et al., 2014a; Farah and Azevedo, 2017; Tarko, 2021; Nazir et al., 2024; Niu et al., 2024). Here we zoom out the view and provide a derivation that collisions and conflicts are both extreme events of ordinary interactions, at varying levels of intensity.

Our derivation uses a similar logic to lifetime survival analysis, which is a branch of statistics to model the expected duration of time until one event, such as deaths or failures of medical treatment, occurs (Reiss and Thomas, 1997). The essential idea of lifetime survival analysis is to estimate the probability of events from the records of human lives or running systems, where the frequency of these events reduces when living time increases. Therefore, the survival function is $S(t) = \Pr(T > t)$, representing the cumulative probability that the duration of survival T is longer than some specified time t . The longer the t , the closer to death or failure.

For estimating the probability of conflicts, collisions are system failures, and our evaluation is based on the records of daily interactions. Contrary to lifetime survival analysis which measures survival duration and longer duration is closer to deaths, here conflict analysis measures proximity and shorter proximity is closer to collisions. The frequency of conflicts increases when proximity reduces. We thus define a ‘‘conflict function’’ $F(s) = \Pr(S < s)$, which represents the cumulative probability that the interaction proximity S is less than a certain s . Recall that in Section 2.2 we use $f_S(s|\theta; \phi)$ to describe the proximity distribution of S in an interaction context θ . Rewriting the function as $f_S(s; \phi)$ for convenience, given that $\phi = \mathbb{I}(\theta)$ and $\theta = \mathbb{I}^{-1}(\phi)$, then Eq. (7) shows the conflict function in an interaction context θ where the conflict hierarchy is characterised by ϕ .

$$F(s; \phi) = \Pr(S < s; \phi) = \int_0^s f_S(x; \phi) dx \tag{7}$$

The shorter the s , the more likely a collision is to occur, but the likelihood varies with conflict intensity. Aligned with the conflict hierarchy in Fig. 2, we define conflict intensity as the inverse of conflict frequency, i.e., a conflict at intensity n occurs once per n times in the same interaction context. Specifically, assuming an interaction where the proximity between road users is s ; if s is always smaller




Fig. 2. Proximity-characterised conflict hierarchy varies in different interaction contexts.

than the observed proximities for n times interactions in the same situation, this interaction is a conflict with a frequency of $1/n$, and we consider it an extreme event of intensity n . According to extreme value theory, we can use $(\Pr(S \geq s; \phi))^n$ to calculate the cumulative probability that a proximity s is the minima in n times of observations in the same interaction context. Therefore, for an interaction at proximity s and in the context θ that is characterised by ϕ , its probability of conflict is not a value, but a function regarding conflict intensity n . We denote this function by C as shown in Eq. (8).

$$p(c|s, \phi) := C(n; s, \phi) = (1 - F(s; \phi))^n = \left(\int_s^\infty f_s(x; \phi) dx \right)^n \tag{8}$$

Fig. 3(a) shows the probability distributions of Eq. (8) under different intensity n . As displayed in Fig. 3(b), this probability fulfils the requirements in Eqs. (6), monotonically increasing while s decreases. In Fig. 3(c), we present an extreme value theory-based interaction spectrum described by Eq. (8). It corresponds to the practical understanding of a conflict: the smaller the proximity and the lower the conflict intensity, the greater the conflict probability. Note that conflict probability does not equate to collision probability in our derivation. Collision probability specifically refers to the case when n is very large and approaches infinity, representing the most intense and least frequent interaction. In that case, Eq. (8) converges to either Gumbel, Frechet, or Weibull distribution. Here we do not particularly consider this convergence at an infinite n , as traffic and driving safety focuses not just on collisions, but on conflicts at varying levels of intensity.

3. Statistical learning tasks for conflict detection

The previous section explains the unified theoretical foundation of conflict detection. Now we continue by framing a series of statistical learning tasks for application in practice. Based on the derivation in Section 2, conflict detection estimates the probability of context-dependent and proximity-characterised extreme events, and we can decompose it into three tasks as shown in Fig. 4. The first task $\theta = \mathbb{R}(X)$ maps situational observables X to a compressed information space as θ , which represents interaction context. In a certain interaction context θ , the second task $\phi = \mathbb{I}(\theta)$ infers the conditional probability distribution of typical proximity behaviour of road users. Then the third task uses extreme value theory to determine the relations between proximity, conflict intensity, and conflict probability. When $\theta = \mathbb{R}(X)$ and $\phi = \mathbb{I}(\theta)$ are deterministic, we can rewrite Eq. (1) into Eq. (9).

$$p(c|s, X) = \iint p(c|s, \phi) p(\phi|\theta) p(\theta|X) d\phi d\theta = p(c|s, \phi) := C(n; s, \phi), \text{ if } p(\theta|X) = \delta(\theta - \mathbb{R}(X)) \text{ and } p(\phi|\theta) = \delta(\phi - \mathbb{I}(\theta)) \tag{9}$$

These three tasks can be integrated within a pipeline for end-to-end learning, and can also be performed separately. In the following subsections, we will explain the tasks further and present preliminary methods for application. We emphasise here that there are many useful learning methods, and the ones we use in this study are not necessarily the best. Future investigation on optimising suitable statistical learning techniques is required.

3.1. Representation of interaction context

The first task $\theta = \mathbb{R}(X)$ is to create an informative representation for interaction contexts by selecting and transforming relevant observables of interaction situations. Traditionally, key variables such as absolute speed, relative speed, and deceleration rate have been widely validated and used in existing surrogate measures of safety. However, there may be more factors at play than those hypothesised in the existing measures. As mentioned in Section 2.1, X may cover various aspects such as vehicle motion states, environment factors, and participant characteristics. Incorporating different factors can be particularly helpful for a comprehensive assessment of collision risk and for user-customised collision warnings.




Fig. 3. Interaction spectrum described with context-dependent proximity distribution and extreme value theory.




Fig. 4. The unified framework and statistical learning tasks for conflict detection.

If the observables of an interaction situation are numerous or use time-series and/or image formats, data-driven representation learning can be effectively employed. Such techniques can compress selected observables into a lower-dimensional space, where different interaction contexts are adequately encapsulated. One approach can be auto-encoders, of which the learning objective is to minimise the reconstruction error between compressed representation and original information. Another approach is contrastive learning, which aims to minimise the difference between the representations of similar samples while maximising the difference between dissimilar samples. Representation learning is an active and evolving field in machine learning, where many other methods are worth investigating. It is important to note that purely data-driven methods may not always capture the nuances of complex interaction situations. Therefore, integrating domain knowledge in traffic safety is necessary and can improve the robustness and reliability of context representation.

To efficiently present the theoretical contributions of this research, we do not perform data-driven representation learning in this study. However, we have some notes for readers on learning representations from variables relevant to motion states. Firstly, the output in the next task of inference is about the proximity between interacting road users. To avoid label leakage, we recommend not including a complete series of position or speed vectors for all involved road users. Secondly, we suggest transforming the variables into local coordinate systems. This transformation can be centred on every road user involved in an interaction situation, as the same situation may be perceived differently by different road users. Such view transformation can ensure the consistency and enhance the comparability of interaction context representations. Once transformed, an interaction context can be represented as instantaneous frames or over a continuous period.

3.2. Inference of conflict hierarchy

The second task $\phi = \mathbb{I}(\theta)$ needs to learn the conditional probability distribution of proximity in a certain interaction context, i.e., $p(s|\theta)$. Our previous study (Jiao et al., 2023) used conditional sampling to infer this probability. Such sampling requires a large amount of data to ensure enough samples for each condition range, while the condition ranges are finite samples of an infinite space. In this study, we use Gaussian Process Regression (GPR) to avoid discretising the condition space, and the rest of this subsection will explain more details. Note that there are other useful methods to learn the conditional probability, but we choose GPR in order to obtain the equation of $p(s|\theta)$ and thus its analytical cumulative probability.

Many studies have found that lognormal distribution best fits the distribution of spatial and temporal gaps between road users (Meng and Qu, 2012; Pawar and Patil, 2016; Anwari et al., 2023). Therefore, we assume that the proximity s in a certain interaction context θ follows a lognormal distribution. Recall that in Section 2.2 we use $f_s(s|\theta; \phi)$ to describe $p(s|\theta)$, and now we have $\phi = \{\mu, \sigma\}$ where μ and σ parameterise a lognormal distribution as shown in Eq. (10). This implies that the logarithm of proximity, $\ln(s)$, follows a Gaussian distribution that is parameterised by the same μ and σ .

$$f_s(s|\theta; \mu, \sigma) = \frac{1}{s\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln(s) - \mu)^2}{2\sigma^2}\right) \quad (10)$$

The Gaussian distribution of $\ln(s)$ can be utilised to learn $p(s|\theta)$ with GPR. First, we consider a mapping $g: \theta \rightarrow \ln(s)$ between proximity s and the interaction context θ that s is in. Assuming a normally distributed noise in this mapping, we denote that $\ln(s) = g(\theta) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Second, given that $\ln(s)$ is derived to be normally distributed, the distribution of function $g(\theta)$ is Gaussian. Then we can consider that $g(\theta)$ is drawn from a Gaussian Process, as shown in Eq. (11). This suggests that a $g(\theta)$ is one sample from the multivariate Gaussian distribution of all possible mappings.

$$g(\theta) \sim \mathcal{GP}(m(\theta), K(\theta, \theta)), \quad (11)$$

where $\mu = m(\theta)$ and $\sigma = K(\theta, \theta)$ are the mean function and covariance function that specify a Gaussian Process. For two similar interaction contexts θ and θ' , $g(\theta)$ is expected to be close to $g(\theta')$.

With GPR, the mean and covariance functions are learned from data. Considering that we may include many variables in the representation of interaction context and thus a fairly large amount of data is needed to serve training, we use the Scalable Variational Gaussian Process (SVGP) model with the python library GPyTorch (Gardner et al., 2018). As the cost to increase scalability, SVGP does not ensure an exact solution and may underestimate variance. To train this model effectively, we maximise a predictive log-likelihood

in Eq. (12), which is proposed in (Jankowiak et al., 2020) and we adapt it here. In addition to the symbols that we have consistently used, N is the number of samples; u denotes inducing variables that are introduced to perform sparse training and reduce computational load; $q(u)$ is the variational distribution of u ; $p(u)$ is the prior distribution of u ; $D_{\text{KL}}[q(u)||p(u)]$ computes the Kullback–Leibler divergence of $q(u)$ and $p(u)$; and β controls the regularisation effect of KL divergence.

$$L = \mathbb{E}_{p_{\text{data}}(\ln(s), \theta)} [\ln(p(\ln(s)|\theta))] - \beta D_{\text{KL}}[q(u)||p(u)] \approx \sum_{i=1}^N \ln \left\{ \mathbb{E}_{q(u)} \left[\int p(\ln(s_i) | g_i) p(g_i | u, \theta_i) dg_i \right] \right\} - \beta D_{\text{KL}}[q(u)||p(u)] \quad (12)$$

Theoretically, learning the mean and covariance functions can approximate any form of mapping g between interaction context θ and proximity s . As explained in Sections 2.1 and 2.2, a surrogate metric for traffic conflicts essentially assumes such a mapping, based on which a threshold is then determined to distinguish safe interactions and conflicts. Therefore, in theory, any metric based on spatial–temporal proximity is a particular case under our unified framework. Without assuming a specific format of g , the mapping embedded in data can be statistically derived. This suggests that, given the same interaction context, a data-driven metric using the unified probabilistic approach proposed in this study should be no less conflict-indicative than a pre-assumed metric.

3.3. Conflict probability estimation and intensity evaluation

In the third task, we use extreme value theory to estimate conflict probability and evaluate conflict intensity. This task plays a similar role to the selection of thresholds when using traditional metrics of conflicts. After the previous two tasks, we can learn parameters $\phi = \{\mu, \sigma\}$ that characterise the conflict hierarchy in different interaction contexts represented with θ . Based on the derivation in Section 2.3, then we can write conflict function $F(s; \phi)$ and obtain $C(n; s, \phi)$ that relates conflict intensity n and conflict probability.

Given that we infer $f_S(s|\theta; \phi)$ as the probability density function of the lognormal distribution, the conflict function $F(s; \phi)$ in Eq. (7) can be further derived as Eq. (13). Here $\text{erf}(z) = 2 \int_0^z e^{-x^2} dx / \sqrt{\pi}$ and is the Gaussian error function. Then conflict probability function is as shown in Eq. (14).

$$F(s; \mu, \sigma) = \int_0^s \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln(x) - \mu)^2}{2\sigma^2}\right) dx = \frac{1}{2} + \frac{1}{2} \text{erf}\left(\frac{\ln(s) - \mu}{\sigma\sqrt{2}}\right) \quad (13)$$

$$C(n; s, \mu, \sigma) = (1 - F(s; \mu, \sigma))^n = \left(\frac{1}{2} - \frac{1}{2} \text{erf}\left(\frac{\ln(s) - \mu}{\sigma\sqrt{2}}\right)\right)^n \quad (14)$$

There are two general purposes for conflict detection. The first is conflict probability estimation. For an interaction, given current proximity s and parameters ϕ that characterise the conflict hierarchy in the interaction context, we can use Eq. (15a) to estimate the probability of a potential collision at different levels of conflict intensity n . The second purpose is safety/conflict evaluation. Assuming that the probability of a potential collision is larger than the probability of no collision, i.e., conflict probability is larger than 0.5, we can use Eq. (15b) to evaluate the maximum possible conflict intensity.

$$\hat{p} = C(n; s, \mu, \sigma) = \left(\frac{1}{2} - \frac{1}{2} \text{erf}\left(\frac{\ln(s) - \mu}{\sigma\sqrt{2}}\right)\right)^n, \quad n \geq 1 \quad (15a)$$

$$\hat{n} = C^{-1}(p; s, \mu, \sigma) = \frac{\ln p}{\ln\left(\frac{1}{2} - \frac{1}{2} \text{erf}\left(\frac{\ln(s) - \mu}{\sigma\sqrt{2}}\right)\right)}, \quad 0.5 < p < 1 \quad (15b)$$

Eqs. (15a) and (15b) are useful in different practices. Conflict probability estimation can be used to issue collision warnings, which alert human drivers or driving assistance systems to prevent potential collisions. Conflict intensity evaluation can be used to identify varying levels of conflict cases in daily traffic. This helps to assess the impact of an infrastructure or traffic policies on traffic safety, and then make according improvements.

4. Demonstration

This section applies the proposed unified framework and statistical learning tasks on real-world trajectory data, to demonstrate the characteristics of this new conflict detection approach. Table 2 presents an overview of our experiment design. In Section 4.1, we introduce the datasets used and experiment details. Then respectively in Sections 4.2 and 4.3, we show the experiment results in conflict probability estimation and conflict intensity evaluation. For convenience, we use *approach* to collectively refer to the unified probabilistic approach proposed in this study; and use the term *unified metric* and its abbreviation *Unified* to refer to the surrogate metric trained with our approach. It is also important to emphasise that these experiments are designed primarily for demonstration. Further exploration is expected in future research.

4.1. Data and experiment details

We use two naturalistic trajectory datasets in this study. First is the highD dataset that was collected at 6 different locations on German highways using drones (Krajewski et al., 2018). It includes detailed information about vehicle types, sizes, and movements, with a positioning error typically less than 10 cm. Considering that a significant part of driving on highways is independent without interactions with other vehicles, we select trajectories that involve lane-changes in highD. In addition to the high-quality drone-collected data for model training, a dataset of real-world conflicts is necessary for demonstration. The second dataset we use is from the 100-Car Naturalistic Driving Study (NDS), which is an instrumented-vehicle study conducted in the U.S. over 2 years in the early 2000s (Dingus et al., 2006). An event database (Custer, 2018) resulted from the study compiles information on 68 crashes and 760 near-crashes. From the time-series sensor data of radars and accelerometers, we reconstruct bird’s eye view trajectories for these events¹. Due to missing values, inaccuracy of sensing, and the lack of ground truth, not all of the events can be reconstructed. We end up matching 180 events based on the constraint of insufficient space (distance less than 4.5 m) for undetected vehicles, including 11 crashes and 169 near-crashes as summarised in Table 3.

The first experiment estimates conflict probability at each time moment for the events in 100-Car NDS data. We utilise the estimation to issue collision warnings, and compare warning effectiveness with using 3 other commonly used surrogate metrics. Effective collision warning maximises true positives while minimising false positives. To reserve safe interactions within an event, we select events with a duration of at least 6 s, no hard braking (acceleration larger than -1.5 m/s^2) in the first 3 s, and speeds larger than 3 m/s at the first time moment for both vehicles involved. After the selection, 4 crashes and 66 near-crashes remained. Considering the relatively low reliability of data in crashes and for a consistent comparison, we use only the selected near-crashes, which are also shown in Table 3. The 3 broadly used metrics for comparison are Proportion of Stopping Distance (PSD), Deceleration Rate to Avoid a Crash (DRAC), and Time-to-Collision (TTC). Table 4 provides an overview of them, where Δs is the distance between vehicle bounding boxes and Δv is relative velocity. For PSD we use a braking rate of 5.5 m/s^2 , as emergency braking is typically between 1.6 and 5.5 m/s^2 (Fambro et al., 2000; Deligianni et al., 2017).

The second experiment evaluates conflict intensity at each time moment for lane-change interactions in the highD dataset. We use lane-changes to demonstrate the applicability of our unified metric to two-dimensional conflicts, which remain challenging to be integratively indicated by traditional metrics. We identify lane changes using the lane references provided in highD data. Then we determine the start and end moments of a lane-change based on the time when the vehicle deviates 1/3 of its vehicle width from the centerline of its current lane. In this way, overtaking is considered as two consecutive lane-changes. Seeing the vehicle making a lane-change as an ego vehicle, the lane-change is potentially interactive if there is a vehicle in front or at rear of the ego vehicle in either the original lane or the target lane. Then we use both the unified metric and two-dimensional TTC (2D-TTC) to evaluate conflict intensity (if any) at each time moment. Here 2D-TTC follows the typical definition of TTC (see Fig. 2, Tarko, 2018), assuming constant velocities for two approaching vehicles at the moment of evaluation². The smaller the TTC value, the higher the conflict intensity; when TTC is infinite, no potential collision is expected to occur.

In both experiments, we use a single unified metric trained with the trajectories involving lane-changes in highD data. In this way, the first experiment can additionally demonstrate the generalisability of our approach; and the metric thresholds calibrated with real-world near-crashes can be used in the second experiment to distinguish conflicts. For a fair comparison with existing metrics, we avoid complex representations for interaction context θ . Firstly, we include the information used in PSD, DRAC, and TTC, i.e., v_{follower}^2 , $\|\Delta v\|^2$, and $\|\Delta v\|$ as analysed in Table 4. Since the lead-follow relationship evolves during lane-changing, we use the squared speeds of both ego vehicles (which change lanes) and target vehicles (in the surrounding). Further, we consider the accelerations of ego vehicles³, heading directions of target vehicles relative to their ego vehicles, and lengths of both ego and target vehicles. In order to indicate two-dimensional conflicts, we define proximity s utilising the two-dimensional spacing proposed in (Jiao et al., 2023). This spacing is denoted by (x, y) in the real-time transformed relative coordinate systems, and we convert (x, y) into (ρ, s) into polar coordinate systems, where ρ is the angle between (x, y) and $(1, 0)$ within a range of $[-\pi, \pi]$. Then we add ρ as one more variable in θ , and use s as the proximity measure to reflect conflict hierarchy.

We train two sets of SVGP models under different settings of the hyper-parameter β in Eq. (12), with $\beta = 5$ and $\beta = 10$. Our training is on 60% of the selected interaction trajectories involving lane-changes; and the validation set and test set account for 20% each. During the training, we reduce the learning rate dynamically to avoid overfitting of the models, and stop training early when validation loss converges. Fig. 5 shows the training progress, where there are 271 batches per epoch and 2,048 samples per batch. Training in both settings converges well, and stops earlier as well as reaches lower loss values when $\beta = 5$. We then evaluate model performance by both loss and negative log-likelihood (NLL) on validation and test sets, as presented in Fig. 5. As a result, we select the model achieving the minimum test loss and the second minimum test NLL, i.e., the model with $\beta = 5$ after 52 epochs of training, to apply in the subsequent demonstration experiments.

Since we consider two-dimensional interactions, maintaining a consistent coordinate system is necessary. The highD dataset uses a coordinate system where the x-axis points from left to right and the y-axis points downwards. This is a mirrored system from the

¹ We open-source our trajectory reconstruction code at <https://github.com/Yiru-Jiao/Reconstruct100CarNDSData>

² We open-sourced an algorithm for fast computation of two-dimensional TTC at <https://github.com/Yiru-Jiao/Two-Dimensional-Time-To-Collision>. This algorithm is used for all experiments in this study.

³ We do not include target vehicle accelerations because this information is lacking in 100-Car NDS data and cannot be reliably derived due to speed fluctuations prior to potential collisions. This lack also prevents comparison with the metric Modified Time to Collision (MTTC).

Table 2

An overview of experiment design for performance demonstration.

Purpose	Training data	Application data	Characteristics: this approach
Conflict probability estimation $\hat{p} = C(n; s, \mu, \sigma)$	Trajectories involving lane-changes in highD	Near-crashes in 100-Car NDS	1) is no less conflict-indicative than any of PSD, DRAC, and TTC; 2) is generalisable across datasets.
Conflict intensity evaluation $\hat{n} = C^{-1}(p; s, \mu, \sigma)$	Trajectories involving lane-changes in highD	Lane-changes in highD	1) covers more diverse conflicts during lane-changes than TTC; 2) detects conflicts in a long-tailed distribution of intensity.

Table 3

Summary of matched and selected events in 100-Car NDS data.

Event happened with	Matched crashes	Matched near-crashes	Selected near-crashes
leading vehicle	5	119	47
following vehicle	6	30	17
vehicle in adjacent lane	0	13	2
vehicle turning across oncoming traffic	0	4	0
vehicle crossing through an intersection	0	2	0
vehicle in oncoming traffic	0	1	0
In total	11	169	66

Table 4

Existing surrogate metrics of conflicts that are used for warning comparison.

Metric	Calculation	Note	Reference(s)	Proximity s	Context θ
PSD	$\frac{\Delta s}{v_{\text{follower}}^2/2/\text{dec}}$	dec = 5.5 m/s ²	(Allen et al., 1978)	Δs	v_{follower}^2
DRAC	$\frac{\ \Delta v\ ^2}{2\Delta s}$	approaching only	(Cooper and Ferguson, 1976)	Δs	$\ \Delta v\ ^2$
TTC	$\frac{\Delta s}{\ \Delta v\ }$	approaching only	(Hayward, 1972; Hydén, 1987)	Δs	$\ \Delta v\ $

traditional engineering coordinate system where the y-axis points upwards. When applying a model that is trained on highD data to other datasets, such as the 100-Car dataset, it is essential to adjust for this inconsistency. We make adjustments by exchanging the x and y coordinates of positions, velocities, and heading directions in 100-Car NDS data before coordinate transformation. We hereby remind the readers to correct potential inconsistencies in the coordinate systems between training data and application data.

4.2. Collision warning compared with existing metrics

The first experiment based on conflict probability estimation is designed to demonstrate two characteristics of our approach. First, the unified metric performs at least as well as any among PSD, DRAC, and TTC in collision warning. This is determined in theory, but experimental results with real-world near-crashes can provide additional evidence. Second, we trained the unified metric with trajectories in the highD dataset while applying it to the events in 100-Car NDS data. This cross-validation involves not only the differences between German and American driving, but also the evolution of driving over more than 15 years. If the highD-trained metric performs well for 100-Car near-crashes, we can argue for the generalisability of this approach, which implies common principles in human road use interactions.

In order to systematically compare warning effectiveness, we define performance indicators as explained in Table 5. Collision warning classifies safe and unsafe interactions based on the values and specific thresholds of metrics. Seeing the metrics as different classification models, we can compare their ROC curves, and a metric is better if the area under its curve is larger. For each metric, we then select an optimal threshold such that the corresponding point on their ROC curve is closest to the ideal point with zero false positive and all true positives (0%, 100%). Based on the optimal threshold, we further calculate warning periods and warning timeliness. The warning period of an event should ideally be close to 100%, but good warning timeliness does not necessarily mean early warning. On average, people need 1 to 1.3 s in response to an obstacle by braking, and in emergencies this can be less than 1 s (Summala, 2000; Markkula et al., 2016). Therefore, good warning timeliness should not be too large, as it may distract people earlier than they need to be warned; neither should it be too small, as the best timing to prevent a potential collision may be missed.

Fig. 6 presents the comparison of collision warning effectiveness. In the plot of ROC curves, we use circles centred at the ideal point (0%, 100%) and crossing the optimal points of different metrics to facilitate comparing their closeness to the ideal point. The metrics of DRAC, TTC, and Unified have comparable areas under their ROC curves, while PSD is less effective. Zooming in to see more details at the top left corner, TTC and Unified are consistently better than DRAC, and are very close to the ideal point. More specifically, the




Fig. 5. Evaluation of SVGP training and selection of the model to apply.

Table 5
Definitions of the performance indicators for collision warning effectiveness comparison.

Performance indicator	Definition
True positive	Seeing the moment when conflicting vehicles reach the minimum distance as a critical moment, the 3 s prior to the moment are supposed to be dangerous; any warning in this period marks a true positive.
False positive	The first 3 s in each of the selected events are supposed to be safe, so any warning in this period marks a false positive.
True positive rate and false positive rate (%)	The rate of true positives and false positives, respectively, among all events. True positive rate ideally approximates 100% and false positive rate 0%.
ROC curve	Receiver Operating Characteristic curve that plots true positive rate against false positive rate at various thresholds.
Warning period (%)	For each event with warnings issued under the optimal threshold, the percentage of warned time moments within the annotated conflict period.
Warning timeliness (s)	For each event with warnings issued under the optimal threshold, the time interval from the last safe-unsafe shift of warning until the critical moment.

optimal threshold of Unified is $n^* = 17$, reaching a true positive rate of 95.45% and a false positive rate of 4.55%; $TTC^* = 4.2$ s and achieves 93.94% true positive rate and 0.00% false positive rate; $DRAC^* = 0.45$, with a true positive rate of 92.42% and a false positive rate of 7.58%. Much less comparably, the optimal threshold of PSD is 0.52, with a true positive rate of 57.58% only but a false positive rate of 25.76%.

These metrics are further compared at their optimal thresholds by the box plots of warning period and warning timeliness in Fig. 6. Given the weak effectiveness of PSD, its warning period and timeliness are not comparable with other metrics. Both DRAC and TTC have higher median warning periods than Unified. This suggests that the unified metric may not issue a warning at the beginning of the annotated conflicts. This is also verified in the plot of warning timeliness, where DRAC and TTC have similar distributions and their median timeliness is both earlier than that of Unified. In addition, the warning timeliness of Unified is less varied. From a positive perspective, this means that using the unified metric gives consistent warnings and does not distract drivers too much; whereas from a negative perspective, this means that the unified metric may miss the chance to prevent a potential collision.

For a more detailed analysis, Fig. 7 shows two frames in the trip indexed by 8332. This event is warned by DRAC only, while TTC is larger than 10.59 s throughout and the conflict probability estimated by Unified remains lower than 0.5. In this figure, we present the profiles of speed, acceleration, and metric values, with the annotated conflict period shaded. We also plot heatmaps of proximity density function $f(s; \phi)$ and conflict probability function $C(s; \phi, n)$, where the y-axis points to the heading direction of ego vehicle and the x-axis points from right to left to align with highD’s coordinate system. For each location around the ego vehicle, we assume the target vehicle is at that location and use actual interaction context θ for estimation. Fig. 7a is the frame at 4.4 s before the critical






Fig. 6. Collision warning effectiveness comparison between PSD, DRAC, TTC, and Unified. In the boxplots of warning period and timeliness, median values are marked below the labels of metrics.

moment; and Fig. 7b is 1.4 s before. Notably, in Fig. 7b, the proximity density distribution, i.e., probable positions of the target vehicle given the interaction context, does not point to the ego vehicle’s heading direction. This is due to the lateral interaction between these two vehicles, which is in line with the narrative data of this event: “the target vehicle was merging into the right lane ahead of the ego vehicle, causing the ego vehicle to brake to avoid a collision”. Comparing the two frames, the unified metric takes into account the ego vehicle’s deceleration as successful prevention of an immediate conflict, and thus the estimated conflict probability stays low during the annotated conflict period.

We summarise the success and failure cases of conflict warnings in Table 6. PSD is not included due to its lower effectiveness. For the events where at least one of Unified, TTC, and DRAC fails, as well as for the false warnings of Unified, we provide dynamic visualisations along with our open-sourced code. The readers are referred to the Acknowledgements to find links. Impressively, 58 out of the 61 near-crashes are correctly warned by all the metrics of Unified, TTC, and DRAC. This is because most events in the 100-Car dataset are rear-end conflicts during car following, for which TTC and DRAC have been proven to be useful. In the next experiment, TTC and our unified metric are challenged to handle two-dimensional traffic conflicts.



(a) Frame 4.4 seconds prior to the critical moment of trip 8332.



(b) Frame 1.4 seconds prior to the critical moment of trip 8332.

Fig. 7. Visualisation example for the collision warning of trip 8332.

Table 6
Summary of success and failure cases of collision warnings by Unified, TTC, and DRAC.

Unified	TTC	DRAC	Number of events	Trip ID
Succeed	Succeed	Succeed	58	
Succeed	Succeed	Fail	3	8622, 8854, 9101
Succeed	Fail	Succeed	0	
Succeed	Fail	Fail	2	8463, 8810
Fail	Succeed	Succeed	1	8761
Fail	Succeed	Fail	0	
Fail	Fail	Succeed	2	8332, 8702
Fail	Fail	Fail	0	

4.3. Conflict intensity evaluation for lane-changing interactions

The second experiment evaluates the conflict intensity of lane-changing interactions, with which we want to demonstrate two more characteristics of our approach. First, this approach can cover a more diverse range of conflicts during lane-changes than TTC. This stands in theory because TTC assumes constant movements (velocities in this study) of interacting vehicles, which identifies potential lane-change conflicts only if two vehicles have crossing velocity directions and could collide without movement change. In contrast, a unified metric trained with our approach on two-dimensional daily driving data considers conflicts in all directions and varying interaction situations. Second, the intensity of conflicts evaluated by this approach is expected to have a long-tailed distribution. More specifically, the number of detected conflicts should decrease according to a power-law as conflict intensity increases. A long tail in such a distribution indicates that, although high-intensity conflicts are extremely rare, they occur with a non-negligible probability. Not all surrogate metrics of conflicts have this characteristic, but it is crucial to indicate very rare conflicts.

There are 13,364 lane-changes in the highD dataset, and in total 713 (5.34%) of the lane-changes involve at least 1 s consecutively identified as a conflict by either TTC or the unified metric. This identification is under the optimal thresholds found in the first experiment, i.e., of 4.2 s for TTC and 17 for Unified. Among these lane-changes, in 679 the ego vehicle conflicts with one other vehicle, in 31 the ego vehicle conflicts with 2 other vehicles, and in 3 the ego vehicle conflicts with 3 other vehicles. These constitute 750 conflicts, among which 67 are indicated by both TTC and Unified; 18 are indicated by TTC only; and 665 are indicated by Unified only. In Fig. 8, we plot the scattered relative locations of target vehicles to their ego vehicles aggregated at (0,0). These locations are where a lane-changing interaction reaches the minimum TTC or the maximum intensity evaluated by Unified.

Fig. 8 clearly shows that conflict detection by TTC strongly relies on the assumption of constant movements, while training a unified metric can cover all directions around the ego vehicle and accounts for the heterogeneity of proximity behaviour in every different direction. There are significantly more conflicts between the ego vehicles making lane-changes and the target vehicles in the left lane than those in the right lane. This may be due to more lane-changes from right to left or higher speeds of vehicles on the left lane, but future investigation is needed for more precise reasons. Notably, most of the lane-changing interactions have intensities lower than 100, which means one such conflict on average occurs in every 100 or fewer interactions. In comparison, conflicts with intensities higher than 100 are fewer, and those with intensities higher than 1000 are significantly fewer. However, despite their seeming rarity, they are far from unlikely to occur. In fact, foreseeing and preventing these safety-critical events remain key challenges for safe autonomous driving (Liu and Feng, 2024).

We present a closer look at the intensity distribution with Fig. 9. In the left half of the figure, we plot the histograms of TTC values and Unified intensities respectively, for those at each individual moment and for averaged during each lane-change conflict case. The averaged values consider only lane-changing process, while the individual moments include car-following periods before and after lane-changes. The histograms show a decreasing frequency when intensity increases (TTC decreases), which is aligned with the assumption of conflict hierarchy. To further see if the distributions are long-tailed, in the right half of Fig. 9, we make log-log scatter plots of the distributions, as well as their dashed trend lines and calibrated functions. All of the logarithm relationships are linear, therefore, all distributions to different extents have the characteristic of power-law. However, the trend line of TTC values averaged during lane-change conflicts has a higher slope than the line of TTC values at individual moments. This suggests that TTC detects fewer




Fig. 8. Spatial distributions of conflict moments with minimum TTC or maximum intensity in every lane-change. Target vehicle positions are transformed to a coordinate system centred at the ego vehicle position and with the longitudinal axis pointing to the ego vehicle heading direction. The dashed lines mark lanes at the average lane width in highD data.




Fig. 9. Intensity distributions during lane-changing interactions and at individual moments in the highD dataset.

lane-changing conflicts than car-following conflicts. In contrast, the unified metric evaluates conflicts in the same theoretical framework, and thus has more consistent trend lines between averaging during lane-changes and at individual moments.

To provide more intuitive information, we make dynamic visualisations for the detected conflicts in different ranges of intensity indicated by TTC and Unified. A link can be found in the Acknowledgements. Here in Fig. 10 we present an example at location 1 (indexed in the highD dataset) with vehicle track indices of 1860 and 1858. The ego vehicle in red makes two sequential lane-changes, and conflicts with the target vehicle in the intermediate lane. Similar to the visualisation for near-crashes in Section 4.2, we plot the profiles of speed, acceleration, evaluated conflict intensity, and TTC values, as well as a real-time heatmap of proximity distribution. In the plots of profiles, lane-change periods are shaded. In addition, we visualise the interaction spectrum described by $C(n, s; \phi)$, which relates conflict intensity n , current proximity s , and conflict probability p .

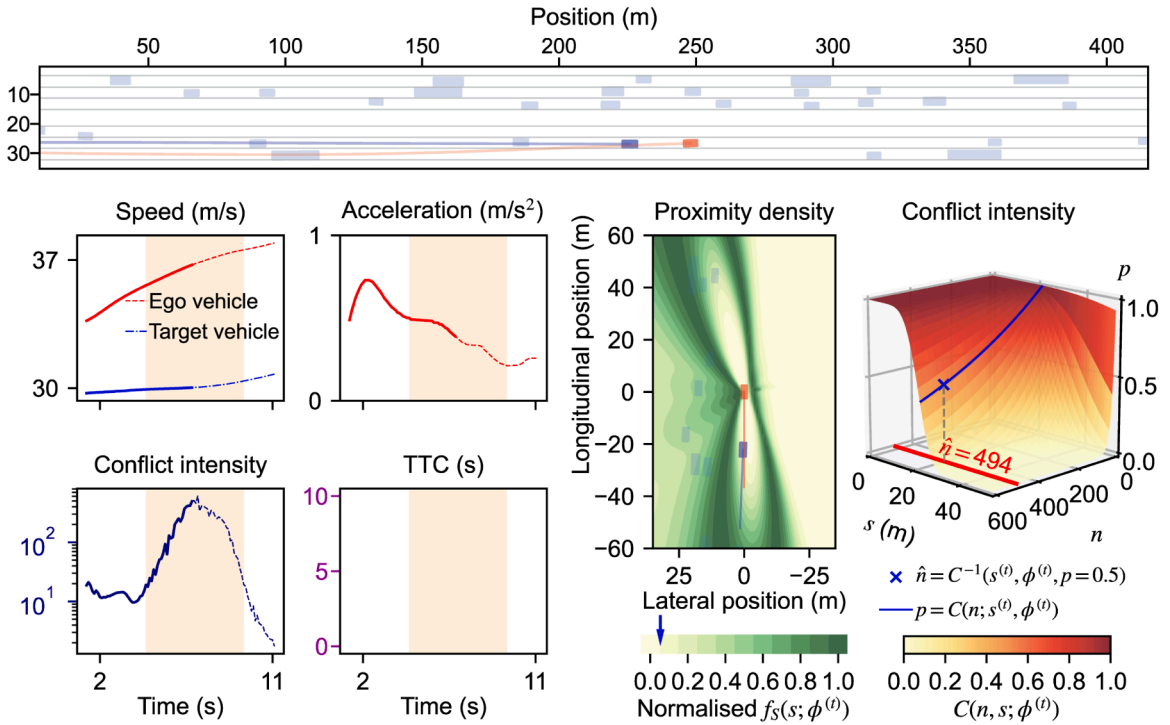
The first frame in Fig. 10a is when the ego vehicle finishes its first lane-change from the right lane to the middle lane. In this frame, the conflict intensity is evaluated to reach the highest level in the whole process of sequential lane-change. On the surface of the interaction spectrum, we plot a line showing the relation between conflict intensity and probability at real-time proximity. This moment is at a very high probability to be considered as a minor conflict, and the maximum intensity to be considered as a conflict, i.e., the maximum intensity if conflict probability is larger than 0.5, reaches 494. This can be considered a serious conflict that occurs once per 494 times in the same interaction situations. Then the evaluated intensity decreases when the ego vehicle leaves the middle lane and continues moving to the left lane. In the frame shown in Fig. 10b, at the end of these two sequential lane changes, the evaluated intensity returns to safe levels.

5. Conclusion and discussion

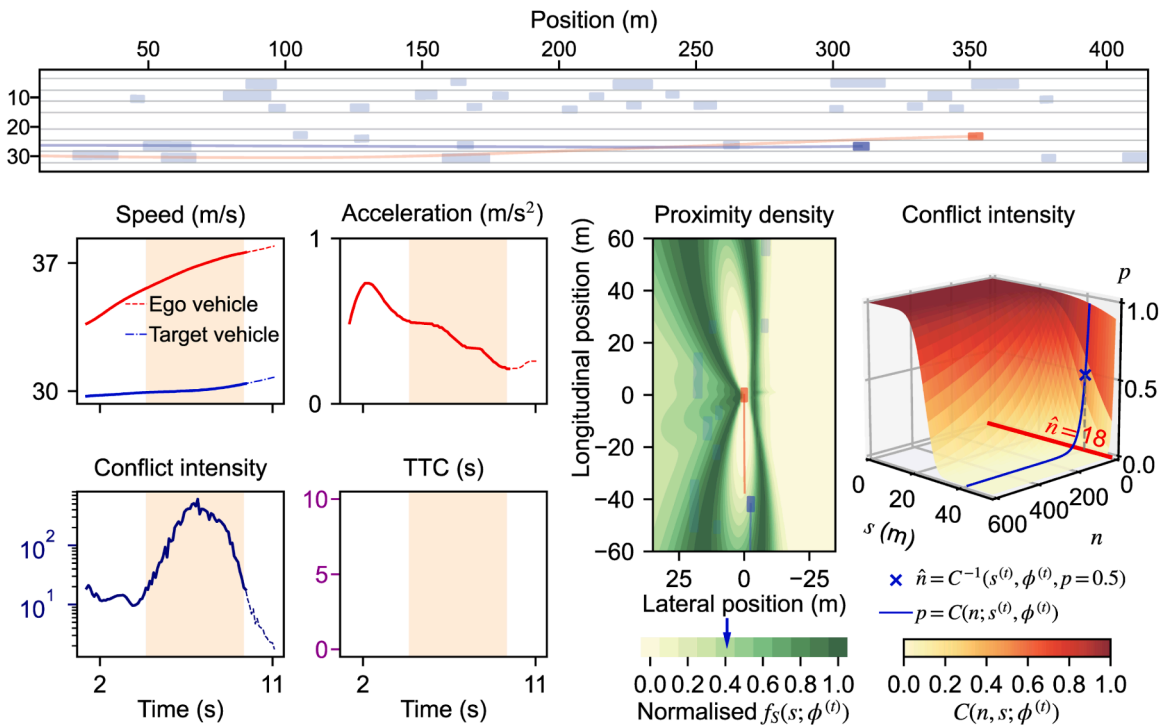
Conflicts do not arise out of nowhere, and every conflict is an extreme continuation of preceding safe interactions. Based on this assumption, this study presents a unified probabilistic approach to detecting traffic conflicts. The unified framework of traffic conflict detection models conflicts as context-dependent extreme events in ordinary interactions. Under this framework, any existing surrogate safety measure that captures certain aspects of conflicts is a special case. Then the statistical learning tasks allow for data-driven hypotheses of traffic conflicts, making the unified detection at least as effective as using a pre-hypothesised surrogate measure. Preliminary experiments with real-world trajectory data demonstrate that the proposed approach provides effective collision warnings, generalises well across datasets, captures various conflicts consistently, and detects conflicts in a long-tailed distribution of intensity.

These features enable consistent and comprehensive conflict detection, which can support scalable and reliable road safety research in the future. For example, more comprehensive surrogate measures of safety can be learned from naturalistic data. In the same way, an integrated estimation of collision risk becomes feasible for automated vehicles. Data-driven conflict detection also allows for the analysis of a broader range of factors underlying unsafe interactions in various situations. When the risk of conflicts is consistently evaluated across interaction contexts, it becomes possible to compare cross-modal road user interactions and provide consistent safety assessment of road infrastructures. Furthermore, safety-critical events that are rarely observed in data may be effectively identified, reinforced, and generated, contributing to training and testing automated vehicles. Pressing societal concerns, such as the impact of vehicle automation on road safety, could also be explored.

The current study has several limitations that should be investigated in future research. First, the demonstration in this study uses spatial proximity and trajectories in car-following and lane-changing scenarios; however, verification in more diverse traffic environments is needed. The framework could also be applied to temporal proximity, and involve other road users such as pedestrians and cyclists. Future research on cross-modal interaction safety is thus promising. Second, the proposed approach does not directly predict when a collision will occur. This could be addressed by incorporating trajectory prediction models and forming a more complete methodology for safe interaction planning. Third, the theoretical framework is based on proximity and does not indicate the severity of a potential collision. Collision severity depends on the energy released by a collision, which is more physics driven than extreme value theory-based. As severity is a key aspect of collision avoidance, further exploration is necessary. Lastly, the uncertainty in conflict probability estimation, which could inform the reliability of detection, is not quantified in this study. This quantification can be



(a) Frame 9719 at which the ego vehicle finishes first lane-change from the right lane to the middle lane.



(b) Frame 9747 at which the ego vehicle changes lane again to the left lane.

Fig. 10. Example dynamic visualisation of a lane-changing conflict detected by Unified only.

achieved by inferring the distributions of proximity characteristic parameters, given that the parameters are learned using Gaussian models.

An important consideration regarding data should be noted. Although learning the representation of interaction context is identified as a key task, this study does not conduct representation learning. In theory, the use of learning methods can include almost all the information one can collect to describe the interactions between road users. This can go beyond movements and include road layouts, weather, individual characteristics of road users, etc. It is necessary to note, however, that the more information gets involved, the more diverse data is required to train an effective unified metric. Since the framework is well-suited for learning from accumulative evidence, future research could expand into continual learning as more data becomes available, allowing for increasingly robust and adaptive detection of traffic conflicts over time.

Declaration of Competing Interest

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

Data availability

We have shared data and code via links in acknowledgements.

Acknowledgments

This work is supported by the TU Delft AI Labs programme. We thank Dr. Haneen Farah for her insightful discussion and internal review. We thank Dr. Guopeng Li for his discussion about technical details. We also would like to thank the editor and anonymous reviewers for their valuable comments and advice.

We open-source our code for applying the approach and repeating our experiments at <https://github.com/Yiru-Jiao/UnifiedConflictDetection>. Dynamic visualisations of near-crashes in 100-Car NDS data are in the subfolder `./Data/DynamicFigures/ProbabilityEstimation/gifs/`; of lane-changing conflicts in highD data are in the subfolder `./Data/DynamicFigures/IntensityEvaluation/gifs/`.

References

- Abdel-Aty, M., Wang, Z., Zheng, O., Abdelraouf, A., 2023. Advances and applications of computer vision techniques in vehicle trajectory generation and surrogate traffic safety indicators. *Accid. Anal. Prev.* 191, 107191.
- Allen, B., Shin, B., Cooper, P., 1978. Analysis of traffic conflicts and collisions. *Transp. Res. Rec.* 667, 67–74.
- Anwari, N., Abdel-Aty, M., Goswamy, A., Zheng, O., 2023. Investigating surrogate safety measures at midblock pedestrian crossings using multivariate models with roadside camera data. *Accid. Anal. Prev.* 192, 107233.
- Arun, A., Haque, M., Bhaskar, A., Washington, S., Sayed, T., 2021. A systematic mapping review of surrogate safety assessment using traffic conflict techniques. *Accid. Anal. Prev.* 153, 106016.
- Arun, A., Haque, M., Washington, S., Sayed, T., Mannering, F., 2021. A systematic review of traffic conflict-based safety measures with a focus on application context. *Analytic Methods in Accident Research* 32, 100185.
- Beitel, D., Stipanovic, J., Manaugh, K., Miranda-Moreno, L., 2018. Assessing safety of shared space using cyclist-pedestrian interactions and automated video conflict analysis. *Transp. Res. Part D* 65, 710–724.
- Borsos, A., Farah, H., Laureshyn, A., Hagenzieker, M., 2020. Are collision and crossing course surrogate safety indicators transferable? A probability based approach using extreme value theory. *Accid. Anal. Prev.* 143, 105517.
- Broadhurst, A., Baker, S., Kanade, T., 2005. Monte Carlo road safety reasoning. In: *Proceedings of IEEE Intelligent Vehicles Symposium*. Las Vegas, NV, USA.
- Camara, F., Fox, C., 2020. Space invaders: pedestrian proxemic utility functions and trust zones for autonomous vehicle interactions. *International Journal of Social Robotics* 13, 1929–1949.
- Chauhan, R., Dhamaniya, A., Arkatkar, S., 2022. Challenges in rear-end conflict-based safety assessment of highly disordered traffic conditions. *Transp. Res. Rec.* 2677 (2), 624–634.
- Cooper, D., Ferguson, N., 1976. Traffic studies at T-junctions – A conflict simulation record. *Traffic Engineering and Control* 17 (7), 306–309.
- Cooper, P., 1977. Reports from group discussions. *Proceedings of the first workshop on traffic conflicts*. Oslo, Norway.
- Custer, K., 2018. 100-Car Data.
- Das, L., Won, M., 2021. SAINT-ACC: safety-aware intelligent adaptive cruise control for autonomous vehicles using deep reinforcement learning. *Proceedings of the 38th International Conference on Machine Learning* 139, 2445–2455.
- Das, S., Maurya, A., 2020. Defining Time-to-Collision thresholds by the type of lead vehicle in non-lane-based traffic environments. *IEEE Trans. Intell. Transp. Syst.* 21 (12), 4972–4982.
- Das, T., Samandar, M., Roupail, N., 2022. Longitudinal traffic conflict analysis of autonomous and traditional vehicle platoons in field tests via surrogate safety measures. *Accid. Anal. Prev.* 177, 106822.
- de Ceunynck, T., 2017. Defining and applying surrogate safety measures and behavioural indicators through site-based observations. PhD thesis Transport and Roads, Hasselt University.
- de Gelder, E., Adjenughwure, K., Manders, J., Snijders, R., Paardekooper, J.-P., Op den Camp, O., Tejada, A., De Schutter, B., 2023. PRISMA: A novel approach for deriving probabilistic surrogate safety measures for risk evaluation. *Accid. Anal. Prev.* 192, 107273.
- Deligianni, S., Quddus, M., Morris, A., Anvuur, A., Reed, S., 2017. Analyzing and modeling drivers deceleration behavior from normal driving. *Transp. Res. Rec.* 2663, 134–141.
- Dingus, T., Klauer, S., Neale, V., Petersen, A., Lee, S., Sudweeks, J., Perez, M., Hankey, J., Ramsey, D., Gupta, S., Bucher, C., 2006. The 100-car naturalistic driving study, Phase II-Results of the 100-car field experiment DOT-HS-810-593. United States. Department of Transportation. National Highway Traffic Safety Administration.
- Fambro, D., Koppa, R., Picha, D., Fitzpatrick, K., 2000. Driver braking performance in stopping sight distance situations. *Transp. Res. Rec.* 1701, 9–16.
- Farah, H., Azevedo, C., 2017. Safety analysis of passing maneuvers using extreme value theory. *IATSS Research* 41 (1), 12–21.
- Formosa, N., Quddus, M., Ison, S., Abdel-Aty, M., Yuan, J., 2020. Predicting real-time traffic conflicts using deep learning. *Accid. Anal. Prev.* 136, 105429.

- Gardner, J., Pleiss, G., Weinberger, K., Bindel, D., Wilson, A., 2018. GPyTorch: blackbox matrix-matrix Gaussian process inference with GPU acceleration. *Advances in Neural Information Processing Systems* 31, 587–7597.
- Hayward, J., 1972. Near miss determination through use of a scale of danger. 51st Annual Meeting of the Highway Research Board 384, 24–34.
- Hydén, C., 1987. The development of a method for traffic safety evaluation: The Swedish Traffic Conflicts Technique. PhD thesis. Lund University.
- Jankowiak, M., Pleiss, G., Gardner, J., 2020. Parametric Gaussian process regressors. *Proceedings of the 37th International Conference on Machine Learning* 119, 4702–4712.
- Jiao, Y., Calvert, S., van Cranenburgh, S., van Lint, H., 2023. Inferring vehicle spacing in urban traffic from trajectory data. *Transp. Res. Part C* 155, 104289.
- Junietz, P., Bonakdar, F., Klamann, B., Winner, H., 2018. Criticality metric for the safety validation of automated driving using model predictive trajectory optimization. In: *IEEE 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 60–65.
- Kim, J., Kum, D., 2018. Collision risk assessment algorithm via lane-based probabilistic motion prediction of surrounding vehicles. *IEEE Trans. Intell. Transp. Syst.* 19 (9), 2965–2976.
- Krajewski, R., Bock, J., Kloeker, L., Eckstein, L., 2018. The highD Dataset: A drone dataset of naturalistic vehicle trajectories on German highways for validation of highly automated driving systems. *IEEE 21st International Conference on Intelligent Transportation Systems (ITSC)*. Maui, HI, United States, pp. 2118–2125.
- Kuang, Y., Qu, X., Wang, S., 2015. A tree-structured crash surrogate measure for freeways. *Accid. Anal. Prev.* 77, 137–148.
- Kusano, K., Chen, R., Montgomery, J., Gabler, H., 2015. Population distributions of time to collision at brake application during car following from naturalistic driving data. *Journal of Safety Research* 54, 95.e29–104.
- Laureshyn, A., Svensson, Å., Hydén, C., 2010. Evaluation of traffic safety, based on micro-level behavioural data: Theoretical framework and first implementation. *Accid. Anal. Prev.* 42 (6), 1637–1646.
- Lewis-Evans, B., De Waard, D., Brookhuis, K., 2010. That's close enough: A threshold effect of time headway on the experience of risk, task difficulty, effort, and comfort. *Accid. Anal. Prev.* 42 (6), 1926–1933.
- Li, Y., Lu, J., Xu, K., 2017. Crash risk prediction model of lane-change behavior on approaching intersections. *Discrete Dynamics in Nature and Society* 2017 (1), 7328562.
- Liu, H., Feng, S., 2024. Curse of rarity for autonomous vehicles. *Nature Communications* 15 (1), 4808.
- Lu, C., He, X., van Lint, H., Tu, H., Happee, R., Wang, M., 2021. Performance evaluation of surrogate measures of safety with naturalistic driving data. *Accid. Anal. Prev.* 162, 106403.
- Mannering, F., Shankar, V., Bhat, C., 2016. Unobserved heterogeneity and the statistical analysis of highway accident data. *Analytic methods in accident research* 11, 1–16.
- Markkula, G., Engström, J., Lodin, J., Bärgrman, J., Victor, T., 2016. A farewell to brake reaction times? Kinematics-dependent brake response in naturalistic rear-end emergencies. *Accid. Anal. Prev.* 95, 209–226.
- Mazaheri, A., Saffarzadeh, M., Nadimi, N., NaserAlavi, S., 2023. A revise on using surrogate safety measures for rear-end crashes. *IATSS Research* 47 (1), 105–120.
- Meng, Q., Qu, X., 2012. Estimation of rear-end vehicle crash frequencies in urban road tunnels. *Accid. Anal. Prev.* 48, 254–263.
- Nadimi, N., Amiri, A., Sadri, A., 2021. Introducing novel statistical-based method of screening and combining currently well-known surrogate safety measures. *Transportation Letters* 14 (4), 385–395.
- Nadimi, N., NaserAlavi, S., Asadamraji, M., 2022. Calculating dynamic thresholds for critical time to collision as a safety measure. *Proceedings of the Institution of Civil Engineers - Transport* 175 (7), 403–412.
- Nazir, F., Ali, Y., Haque, M., 2024. Effects of sample size on pedestrian crash risk estimation from traffic conflicts using extreme value models. *Analytic Methods in Accident Research* 44, 100353.
- Niu, D., Sayed, T., Fu, C., Mannering, F., 2024. A cross-comparison of different extreme value modeling techniques for traffic conflict-based crash risk estimation. *Analytic Methods in Accident Research* 44, 100352.
- Panou, M., 2018. Intelligent personalized ADAS warnings. *European Transport Research Review* 10 (2), 59.
- Pawar, D., Patil, G., 2016. Critical gap estimation for pedestrians at uncontrolled mid-block crossings on high-speed arterials. *Saf. Sci.* 86, 295–303.
- Razi, A., Chen, X., Li, H., Wang, H., Russo, B., Chen, Y., Yu, H., 2022. Deep learning serves traffic safety analysis: A forwardlooking review. *IET Intel. Transport Syst.* 17 (1), 22–71.
- Reiss, R.-D., Thomas, M., 1997. *Statistical analysis of extreme values*. Birkhäuser, Basel, Switzerland.
- Saunier, N., Mourji, N., Agard, B., 2011. Mining microscopic data of vehicle conflicts and collisions to investigate collision factors. *Transp. Res. Rec.* 2237, 41–50.
- Saunier, N., Sayed, T., 2008. Probabilistic Framework for Automated Analysis of Exposure to Road Collisions. *Transp. Res. Rec.* 2083, 96–104.
- Saunier, N., Sayed, T., Lim, C., 2007. Probabilistic collision prediction for vision-based automated road safety analysis. *IEEE Intelligent Transportation Systems Conference*. Bellevue, WA, USA.
- Schiff, W., Detwiler, M., 1979. Information used in judging impending collision. *Perception* 8 (6), 647–658.
- Siebert, F., Oehl, M., Bersch, F., Pfister, H.-R., 2017. The exact determination of subjective risk and comfort thresholds in car following. *Transp. Res. Part F* 46, 1–13.
- Siebert, F., Oehl, M., Pfister, H.-R., 2014. The influence of time headway on subjective driver states in adaptive cruise control. *Transp. Res. Part F* 25, 65–73.
- Songchitruksa, P., Tarko, A., 2006. The extreme value theory approach to safety estimation. *Accid. Anal. Prev.* 38 (4), 811–822.
- Summala, H., 2000. Brake reaction times and driver behavior analysis. *Transportation Human Factors* 2 (3), 217–226.
- Tageldin, A., Sayed, T., 2019. Models to evaluate the severity of pedestrian-vehicle conflicts in five cities. *Transport-metrica A: Transport Science* 15 (2), 354–375.
- Taieb-Maimon, M., Shinar, D., 2001. Minimum and comfortable driving headways: reality versus perception. *Hum. Factors* 43 (1), 159–172.
- Tarko, A., 2018. *Surrogate measures of safety. Safe mobility: challenges, methodology and solutions*, 11. Emerald Publishing Limited, Leeds, pp. 383–405.
- Tarko, A., 2021. A unifying view on traffic conflicts and their connection with crashes. *Accid. Anal. Prev.* 158, 106187.
- Teigen, K., 2005. The proximity heuristic in judgments of accident probabilities. *Br. J. Psychol.* 96 (4), 423–440.
- Uno, N., Iida, Y., Itsubo, S., Yasuhara, S., 2002. A microscopic analysis of traffic conflict caused by lane-changing vehicle at weaving section. In: *Proceedings of the 13th mini-EURO conference-handling uncertainty in the analysis of traffic and transportation systems*. Bari, Italy, pp. 10–13.
- Vahidi, A., Eskandarian, A., 2003. Research advances in intelligent collision avoidance and adaptive cruise control. *IEEE Trans. Intell. Transp. Syst.* 4 (3), 143–153.
- Venturuthiyil, S., Chunchu, M., 2022. Anticipated Collision Time (ACT): A two-dimensional surrogate safety indicator for trajectory-based proactive safety assessment. *Transp. Res. Part C* 139, 103655.
- Vogel, K., 2003. A comparison of headway and time to collision as safety indicators. *Accid. Anal. Prev.* 35 (3), 427–433.
- Wang, C., Xie, Y., Huang, H., Liu, P., 2021. A review of surrogate safety measures and their applications in connected and automated vehicles safety modeling. *Accid. Anal. Prev.* 157, 106157.
- Wang, X., Alonso-Mora, J., Wang, M., 2022. Probabilistic risk metric for highway driving leveraging multi-modal trajectory predictions. *IEEE Trans. Intell. Transp. Syst.* 23 (10), 19399–19412.
- Wang, Y., Li, Z., Liu, P., Xu, C., Chen, K., 2024. Surrogate safety measures for traffic oscillations based on empirical vehicle trajectories prior to crashes. *Transp. Res. Part C* 161, 104543.
- Ward, J., Agamennoni, G., Worrall, S., Bender, A., Nebot, E., 2015. Extending Time to Collision for probabilistic reasoning in general traffic scenarios. *Transp. Res. Part C* 51, 66–82.
- Wessels, M., Kröling, S., Oberfeld, D., 2022. Audiovisual time-to-collision estimation for accelerating vehicles: The acoustic signature of electric vehicles impairs pedestrians' judgments. *Transportation research part F* 91, 191–212.
- Westhofen, L., Neurohr, C., Koopmann, T., Butz, M., Schütt, B., Utesch, F., Neurohr, B., Gutenkunst, C., Böde, E., 2022. Criticality metrics for automated driving: A review and suitability analysis of the state of the art. *Archives of Computational Methods in Engineering* 30 (1), 1–35.
- Wu, K., Jovanis, P., 2012. Crashes and crash-surrogate events: Exploratory modeling with naturalistic driving data. *Accid. Anal. Prev.* 45, 507–516.
- Zheng, L., Ismail, K., Meng, X., 2014a. Freeway safety estimation using extreme value theory approaches: A comparative study. *Accid. Anal. Prev.* 62, 32–41.

- Zheng, L., Ismail, K., Meng, X., 2014b. Traffic conflict techniques for road safety analysis: open questions and some insights. *Can. J. Civ. Eng.* 41 (7), 633–641.
- Zheng, L., Sayed, T., Mannering, F., 2021. Modeling traffic conflicts for use in road safety analysis: a review of analytic methods and future directions. *Anal. Methods Accid. Res.* 29, 100142.
- Zhou, H., Zhong, Z., 2020. Evasive behavior-based method for threat assessment in different scenarios: A novel framework for intelligent vehicle. *Accid. Anal. Prev.* 148, 105798.