

How Far Ahead Should Autonomous Vehicles Start Resolving Predicted Conflicts? Exploring Uncertainty-Based Safety-Efficiency Trade-Off

Li, Guopeng; Li, Zirui; Knoop, Victor Lambert; van Lint, J. W.C.

DOI

[10.1109/TITS.2024.3393641](https://doi.org/10.1109/TITS.2024.3393641)

Publication date

2024

Document Version

Final published version

Published in

IEEE Transactions on Intelligent Transportation Systems

Citation (APA)

Li, G., Li, Z., Knoop, V. L., & van Lint, J. W. C. (2024). How Far Ahead Should Autonomous Vehicles Start Resolving Predicted Conflicts? Exploring Uncertainty-Based Safety-Efficiency Trade-Off. *IEEE Transactions on Intelligent Transportation Systems*, 25(10), 14183–14195. <https://doi.org/10.1109/TITS.2024.3393641>

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.

How Far Ahead Should Autonomous Vehicles Start Resolving Predicted Conflicts?

Exploring Uncertainty-Based Safety-Efficiency Trade-Off

Guopeng Li^{ID}, Zirui Li, Victor Lambert Knoop^{ID}, and J.W.C. van Lint^{ID}

Abstract—Resolving predicted conflicts is vital for safe and efficient autonomous vehicles (AV). In practice, vehicular motion prediction faces inherent uncertainty due to heterogeneous driving behaviours and environments. This spatial uncertainty increases non-linearly with prediction time horizons, leading AVs to perceive more road space occupied by conflicting vehicles. Reacting early to resolve predicted conflicts can ensure safety but may adversely affect traffic efficiency. Therefore, determining how far ahead AVs should start resolving predicted conflicts based on safety and traffic efficiency constraints is crucial. To answer this question, this study proposes a novel approach to explore the trade-off between safety and traffic efficiency considering prediction uncertainty. Firstly, a continuous-time motion prediction framework is proposed for estimating the spatial probability distribution of a vehicle's future position at any moment within the maximum time horizon. Subsequently, average driver space and the corresponding traffic flow are derived from the safety settings of AV and prediction uncertainty. As such, the safety-efficiency trade-off can be quantified. Experiments show that mandatory decision points, high speeds, and traffic state transitions usually cause fast-increasing prediction uncertainty. A case study of Intelligent Driver Models (IDM) shows that traffic efficiency drops rapidly when AVs resolve predicted conflicts longer than 1.5 seconds ahead. AVs can act earlier on motorways for efficiency concerns but must be myopic at urban intersections. Prediction uncertainty fundamentally constrains the safety-efficiency performance of AVs. These findings are instructive for designing traffic-compatible AVs.

Index Terms—Uncertainty quantification, driver space, traffic efficiency, traffic safety, autonomous vehicles.

I. INTRODUCTION

A. Background

MODELLING, predicting, and understanding human driving behaviours is one of the central topics in autonomous driving. In the context of mixed traffic, where

autonomous vehicles (AV) co-exist with human-driving vehicles (HV), AVs must interact with HVs and, to some extent, learn from or adjust to human behaviours. Therefore, motion prediction becomes an essential component of autonomous driving systems [1]. Accurate and reliable predictions regarding the motion of surrounding HVs can enhance the robustness of decision-making and motion planning [2]. These application values have attracted significant attention from researchers, leading to a wide range of prediction models. Recently, open large real-world AV datasets, such as Waymo [3], nuScenes [4], and Argoverse-2 [5], have greatly expedited the use of data-driven approaches, particularly deep-learning models. Relevant studies are numerous in the literature. We refer the reader to Mozaffari et al. [6] and Huang et al. [7] for comprehensive reviews.

However, besides accuracy, *uncertainty* holds equal importance for motion prediction [8]. Prediction uncertainty is a key factor connecting safety and traffic efficiency. The spatial prediction uncertainty increases with the temporal prediction horizon. Meanwhile, vehicles keep desired space and act ahead of time to guarantee safety in conflict resolution. The spatial-temporal resonance between prediction uncertainty and the safety setting of vehicles together determines traffic efficiency. Generally speaking, acting earlier to resolve predicted conflicts and keeping larger space can ensure safety but reduce road capacities [9]. High prediction uncertainty also damages traffic efficiency by enlarging the required driver space. Therefore, quantifying how uncertainty changes with the prediction time horizon is crucial for further study of when AVs should start resolving predicted conflicts. By exploring the interplay between safety settings, prediction uncertainty, and traffic efficiency, we aim to address one critical concern:

Given safety and traffic efficiency constraints, how far in advance should autonomous vehicles start resolving predicted conflicts?

To answer this question, it is first essential to distinguish between different sources of uncertainty and how uncertainty in motion prediction impacts safety and traffic efficiency.

B. Uncertainty, Safety, and Traffic Efficiency

Human driving behaviour is stochastic. Uncertainty arises principally from the lack of observability. In practice, the

Manuscript received 10 August 2023; revised 3 February 2024; accepted 21 April 2024. This work was supported by the Dutch Research Council (NWO)/Applied and Engineering Sciences (TTW) Project MiRRORS under Grand number 16270. The Associate Editor for this article was S. C. Wong. (Corresponding author: Guopeng Li.)

Guopeng Li, Victor Lambert Knoop, and J.W.C. van Lint are with the Faculty of Civil Engineering and Geosciences, Delft University of Technology, 2628 CN Delft, The Netherlands (e-mail: g.li-5@tudelft.nl).

Zirui Li is with the Chair of Traffic Process Automation, Faculty of Transport and Traffic Sciences, Dresden University of Technology, 01069 Dresden, Germany (e-mail: zirui.li@tu-dresden.de).

Digital Object Identifier 10.1109/TITS.2024.3393641

overall uncertainty has two sources: *aleatoric* uncertainty, which stems from the inherent randomness of the process, and *epistemic uncertainty*, which is caused by the limited available data [10].

Aleatoric uncertainty pertains to the unpredictability of a system due to the exclusion of certain unobservable state variables that may influence the outcome. In motion prediction, some environmental parameters and signals are usually ignored due to limited sensing methods, such as road conditions and vehicle signals (blinkers) [11]. Another source of aleatoric uncertainty is the human factor [12], [13], [14]. In-vehicle information and driving styles are not always observable. Conversely, epistemic uncertainty arises from the limited observations of similar cases, leading to unreliable outcomes due to the lack of knowledge. Near-collision behaviour is an example [15]. Most motion datasets have few collisions. Consequently, models trained on such datasets can provide neither reliable predictions nor accurate estimates of uncertainty in near-collision behaviour. Having clarified these concepts, we emphasize that this study focuses on aleatoric uncertainty only.

Uncertainty in motion prediction has significant implications for both safety and traffic efficiency [16]. A basic assumption is that a smaller microscopic driver space (at optimal speeds) leads to higher macroscopic traffic efficiency [17], but also to higher risk [18]. For stable traffic flows, such as one-dimensional stable car-following, the Fundamental Diagram (FD) derived from observed spacing and speed directly represents the relationship between traffic efficiency and spacing. This is a special case with perfect predictability (no uncertainty). However, this stability assumption breaks down when interactions arise to resolve conflicts. In conflict resolution, the anchor point for spacing must shift from the current state to the predicted future. In AVs' eyes, the future position of the conflicting vehicle becomes more uncertain with longer prediction time horizons so this conflicting vehicle occupies an increasingly larger road space. For example, the Waymo motion prediction open challenge¹ requires predictions up to 8 s of future trajectories. We note that the miss rate for predictions is approximately 15%, even with 6 predicted trajectories using state-of-the-art models. This observation implies that one vehicle can appear in a large area after 8 s due to high prediction uncertainty. If AVs retain a larger time buffer (act earlier), traffic efficiency will be lower due to the larger required road space. Prediction uncertainty makes the spatial and temporal safety settings of AVs entangled. Therefore, quantifying how spatial uncertainty varies over prediction time horizons is critical in deciding when to start resolving conflicts.

Although trading off safety and traffic efficiency concerns from the uncertainty perspective is important, this topic has not been discussed in depth due to two research gaps. Firstly, precisely modelling the continuous change of the spatial probability distribution of a vehicle's future position with prediction time horizons is challenging. Existing methods either accurately learn the spatial distribution at a certain predicted

time horizon only, such as non-parametric *goal-oriented* deep learning paradigms [19], [20], or consider simplified distributions for a series of discrete moments alone, such as Gaussian models [21]. The prediction time horizon is not explicitly considered a continuous variable in these models, but rather a discrete time step. Therefore, exploring horizon settings with an arbitrary resolution is not possible. Secondly, embedding microscopic prediction uncertainty into macroscopic traffic efficiency assessment requires novel methods. However, few studies directly and efficiently connect these two aspects at different scales. Embedding uncertainty-aware motion prediction models (particularly deep-learning models) into macroscopic traffic simulators will cause unacceptable computational complexity. On the other hand, although most traffic-flow-theory-based methods are easy to apply, they rely on the real-time spacing strategy instead of considering the impact of (unknown) prediction uncertainty. This gap also hinders engineers from considering traffic efficiency when tuning the key safety parameters of AVs.

In summary, the two research gaps in modelling the continuous evolution of prediction uncertainty and effectively linking this uncertainty to safety-efficiency trade-offs call for a novel methodological solution. Such a framework offers the opportunity to revisit the motion prediction task from a comprehensive perspective.

C. Contributions and Outline

This study proposes an extendable framework for quantifying the relationship between safety, traffic efficiency, and motion prediction uncertainty. Firstly, we propose a continuous-time probabilistic motion prediction model that can estimate the time-horizon-dependent spatial uncertainty of a vehicle. Next, a generic method that combines the prediction model, driver space, and simulation is proposed to estimate macroscopic traffic efficiency from the safety settings of AVs and prediction uncertainty. The primary contributions of this paper are summarized as follows:

- Propose a flexible motion prediction model that can evaluate the continuous change of spatial uncertainty in motion prediction.
- Connect the uncertainty in microscopic motion prediction and safety setting of AVs to the formation of driver space, and further to macroscopic traffic efficiency.
- Interpret that high speeds, mandatory decision points, and traffic condition transitions make predictive uncertainty increase rapidly.
- Give insights into how prediction uncertainty influences the trade-off between safety settings and traffic efficiency for AVs.

We first overview relevant studies in Section II. Next, the proposed motion prediction model and the training strategy are explained in Section III. Section IV presents the generic method of assessing traffic efficiency from predictive uncertainty and driving models. Then experiments are carried out in Section V and the results are shown and discussed in Section V. Finally, Section VII concludes this paper.

¹<https://waymo.com/open/challenges/2022/motion-prediction/>

II. RELATED STUDIES

This section provides a brief overview of pertinent research in the literature and then points out the novelty of this study. Given the scope of this paper, particular attention will be paid to two parts. The first part focuses on safety measures and traffic efficiency assessment. The second part is about uncertainty quantification in motion prediction.

How to measure safety depends on the interested scale. Surrogate safety measures, such as Time-to-Collision (TTC), Post-Encroachment Time (PET), and Proportion of Stopping Distance (PSW), are widely used to evaluate microscopic safety performance from time or space perspectives. Wang et al. [22] systematically review surrogate safety measures applied for connected and automated vehicles. As TTC is derived from projected trajectories, it is the most common safety setting in vehicles. For macroscopic or enforcement purposes, Vogel [23] show that a shorter average (time or space) headway generates more dangerous situations. Correspondingly, cruising headway and critical TTC are two key safety parameters for AVs. In contrast, traffic efficiency is a macroscopic concept that is usually measured by maximum traffic flow, which is determined by how the average spacing changes with speed. The relationship between averaged flow, space, and speed is the so-called Fundamental Diagram (FD). Next, we overview how to assess traffic efficiency in different scenarios.

For one-dimensional stable car-following, the impact of AVs has been intensively studied since the birth of traffic flow theory. For example, Bose and Ioannou [24] assume a constant time headway for AVs and derive AV-involved FDs and capacities. A shorter time headway means driving more aggressively but can potentially benefit traffic efficiency. In principle, any microscopic car-following model can yield an FD [25]. More complexity can be added to this fundamental model. For example, considering intra-platooning headway strategy [26], different types of AVs [27], heterogeneity of driving models [28], [29], name a few. Two-dimensional interactive scenarios, such as intersections, have drawn more attention recently. In general, the same efficiency-safety principle holds. Reducing the average gap in all directions leads to decreased safety but improved traffic efficiency [30], [31], [32]. However, traffic streams can cross at intersections and lanes have different passing priorities. Therefore, the use of FD and capacity also needs to be adapted. For example, Helbing [33] derive an FD from a simplified traffic flow model at intersections. Wu et al. [34] propose a so-called arterial FD to describe interrupted traffic flow. Different 2D spacing models, e.g., Zhao et al. [32], lead to various safety-efficiency relationships. It is useful to note that many 2D methods are initially proposed to describe pedestrian flows, such as the social force model [35] and the continuum model [36]. They can also be adapted to vehicles. We refer to the review of Vanumu et al. [37] for more information.

The studies above assess safety and traffic efficiency performances from driving models or traffic flow properties. The spacing strategy of vehicles is the key. Nevertheless, these studies consider measured spacing only instead of how the spacing (or driver space [38]) is generated in interactions

and conflict resolution. Particularly, how the uncertainty in motion prediction influences spacing, which is crucial for AVs, is ignored. Therefore, the second half of the review focuses on deep-learning-based uncertainty quantification methods applied for motion prediction. Three major approaches are identified from the literature. They are random sampling, parametric method, and non-parametric method.

Random sampling is a straightforward method for uncertainty estimation. The basic concept is generating multiple random predictions through Monte Carlo simulation, Bayesian networks, multi-modal predictions etc. The diversity of predictions, quantified by e.g. (co)variance, measures the uncertainty. For example, Schreier et al. [39] propose a prediction model containing inherent random parameters in how drivers execute particular manoeuvres. Monte Carlo simulation is used to estimate the collision risk from the output ensemble. Tran and Firl [40] use similar methods to recognize manoeuvres and predict driving behaviours at urban intersections. Instead of running the same simulation multiple times, running multiple models in parallel is also a common method. For instance, Abbas et al. [41] utilise an ensemble of Extended Kalman Filter (EKF) to predict a set of potential scenarios for the future location of a vehicle.

The parametric method describes the output distribution by a parametric prior, such as Gaussian [42] or mixture Gaussian [43]. The uncertainty can be measured by the covariance matrix or differential entropy. For instance, Deo and Trivedi [44] propose a manoeuvre-based LSTM to predict lane-changing behaviours. The output is a series of mixture Gaussian distributions. Huang et al. [21] employ a variational neural network equipped with a Gaussian Mixture Model (GMM) layer to predict trajectories, also providing a measure of confidence for various time horizons. Similar methods involve Graph Vehicle-Pedestrian Attention Network (GVAT) [45], Joint Time Series Model (JTSM) [46], to name a few. To consolidate the estimation of mixture density distributions, Zyner et al. [47] propose a clustering algorithm that recognizes the drivable paths and ranks them based on their possibility scores. Mercat et al. [48] further extend the GMM-based method to the joint prediction of multi-agent trajectories and show competitive performance. The parametric method can track how uncertainty evolves with the horizon using a smaller number of output parameters. However, assimilating the exact components of MDN is numerically difficult. In many real cases, the long-term positions of a vehicle cannot even be approximated by GMM due to map constraints and surrounding agents.

To overcome the limitations mentioned above, the non-parametric method is proposed. The primary difference is that the output distribution is not simplified to a parametric prior, but directly approximated by a two-dimensional histogram. This approach is more realistic and flexible. GoHome [19] and DenseTNT [20] are two representative models. Graph neural networks are used to extract input features and subsequently map them onto a 2D mesh grid. Gilles et al. [49] further apply this method to multi-agent trajectory prediction. However, these models are goal-oriented because of the large number of units in mesh grids. Only the spatial

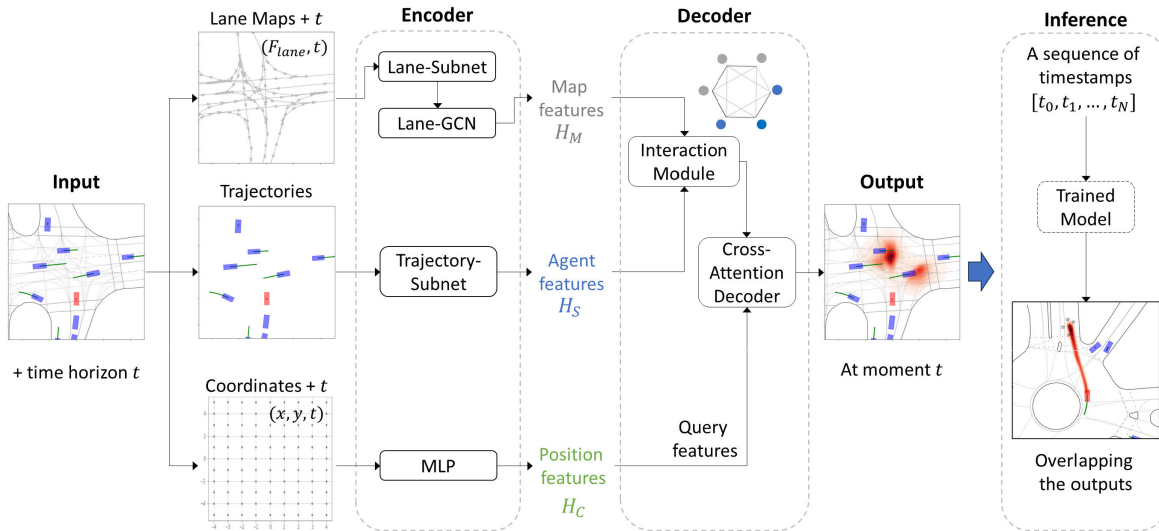


Fig. 1. The proposed model structure and the inference process. The red vehicle is the target vehicle to predict.

distribution at the last moment (e.g., after 3 s) is learnt. The uncertainty through the entire predicted trajectory is unknown. For safety and traffic efficiency assessment purposes, we need an efficient novel single model to learn the continuous change of spatial uncertainty with prediction time horizons.

To summarize, existing studies on traffic efficiency assessment do not consider how prediction uncertainty influences the spacing between vehicles. To include precise prediction uncertainty in arbitrary time budget settings, the current non-parametric uncertainty quantification approach needs to be extended to continuous prediction time horizons. This paper proposes a novel framework to fill these gaps and further clarify the relationship between prediction uncertainty, safety, and traffic efficiency.

III. MODEL

This section presents the proposed continuous-time motion prediction model and the training strategy. This model is fundamental for the following sections of traffic efficiency assessment and trade-off relationship analysis.

A. Continuous-Time Prediction Model

We first formulate the continuous-time probabilistic trajectory forecasting problem. Given the trajectories of all surrounding road agents in the past 1 sec (noted as \mathcal{S}) and the environment variables (high-definite maps, traffic rules, traffic lights, etc. noted as \mathcal{E}). We aim to predict the joint distribution of a target vehicle's position (x, y) at any moment t within the maximum prediction time horizon T :

$$p_t(x, y) = \mathcal{M}(\mathcal{S}, \mathcal{E}; t) \quad \forall t \in (0, T], \quad t \in \mathbb{R} \quad (1)$$

Notice that the output is the posterior distribution that marginalizes all possible interactions and intentions. The model \mathcal{M} is adapted from UQnet proposed in our previous study [50]. UQnet is a goal-oriented model that improves generalizability by adding a causation-based regularization module. Because only aleatoric uncertainty is considered in

this paper, the causal regularization part is removed. To imbue the model with complete awareness of uncertainty throughout the prediction horizon, the model structure and the training method are changed. Arbitrary prediction time horizon t is explicitly included.

The model structure is presented in Fig.1. It follows the compact vectorized representation proposed in VectorNet [51]. The model comprises an encoder and a decoder. The input to the model consists of the current scenario and the arbitrary prediction horizon t . The scenario is abstracted and decomposed into three parts, namely, the historical trajectories of surrounding agents \mathcal{S} , maps information \mathcal{E} , and the spatial mesh-grid \mathcal{C} . Each trajectory and lane centerline is abstracted into a series of end-to-head vectors (polylines). The value of t is concatenated with lane features and mesh-grid features. The three features are passed into three parallel sub-nets in the encoder. An MLP is used to convert \mathcal{C} into position features. The *Trajectory-subnet* and the *Lane-subnet* are two polyline-level graph neural networks [52] that encode the historical motion of each agent and the characteristics of each lane respectively. Lane features are processed by an additional *LaneGCN*-like subnet [53] to capture the connectivity among lanes. The outputs of the encoder are extracted map features \mathbf{H}_M , agent features \mathbf{H}_S and position features \mathbf{H}_C .

The decoder takes the encoded information to generate the predicted spatial distribution. The interaction module in the decoder concatenates map feature \mathbf{H}_M with agent features \mathbf{H}_S and learns the interaction among all the lanes and agents by graph attention layers. Next, a cross-attention decoder projects the output of the interaction module to each point of the mesh grid queried by \mathbf{H}_C (note that prediction horizon t has already been concatenated with \mathbf{H}_C). Finally, an MLP converts the feature dimension to 1 for each point. So the final output is a 2D histogram (heatmap) that approximates the distribution $p_t(x, y)$, denoted as $\hat{\mathbf{Y}}_t$.

In inference, users can pass a sequence of arbitrary time horizons at any desired resolution and intervals into the model. The model can return the predicted spatial distributions

at the corresponding moments in batches. Overlapping the heatmaps together gives a direct representation of how spatial uncertainty changes with time horizons.

Considering that the model structure is not the focus of this paper, we refer the readers to Li et al. [50] and the open-source code² for more details.

B. Training Strategy

The proposed model can predict evolving 2D distributions. However, the inclusion of the time horizon significantly increases the complexity of the output space. For example, if the spatial resolution is 1 m, the expected drivable space is a 50×50 m² area, the time resolution is 0.1 s. The maximum prediction horizon is 3 s. Then the output space (the spatiotemporal cube) has $50 \times 50 \times 30 = 75000$ mesh units (voxels) to query in the cross-attention decoder. To address this computational challenge, a random sampling training strategy is used.

First, the ground-truth trajectory (label) occupies a small percentage of the total voxels (take the same example above, the percentage is $30/75000 = 0.04\%$, we will use this example in this subsection). Learning such sparse features is difficult. To mitigate this sparsity, the label is enhanced by adding white noise. For example, if the ground-truth position is (x_t, y_t) at moment t (a pixel on the heatmap), then a Gaussian noise centred at the ground-truth with standard covariance matrix $\epsilon \mathbf{I}$ is added to this time slice. The voxel $Y_{i,j,t}$ at (x_i, y_j, t) in the enhanced label is:

$$Y_{i,j,t} = \frac{1}{2\pi\sigma_s^2} \exp\left[-\frac{(x_i - x_t)^2 + (y_j - y_t)^2}{2\sigma_s^2}\right] \quad (2)$$

This method results in a significantly larger proportion of non-zero voxels. In this study, $\sigma_s = 0.4$ m. Instead of considering all voxels in training, we use the following method. First, a moment t_s is randomly drawn between 0 and the maximum horizon for each training sample. Next, only voxels in a small spatiotemporal cube around the sampled ground truth are considered in the overall loss function. For example, if the sampled t_s and the label (x_{t_s}, y_{t_s}) are known, then only those voxels (x_c, y_c, t_c) fall in the following area are considered:

$$\begin{aligned} t_c &\in [t_s - \tau, t_s + \tau] \\ x_c &\in [x_{t_s} - \Delta, x_{t_s} + \Delta] \\ y_c &\in [y_{t_s} - \Delta, y_{t_s} + \Delta] \end{aligned} \quad (3)$$

Here τ and Δ control the spatiotemporal range. If we choose $\tau = 0.5$ s and $\Delta = 10$ m, then only $20 \times 20 \times 10 = 4000$ voxels are queried in the decoder, which significantly decreases the spatial complexity. For the validation of this training strategy, we refer the readers to Appendix A.

Focal loss [54] is the main loss function. Focal loss can address the issue of imbalanced classes by adding a modulating factor that reduces the contribution of easy cases but focuses more on rare examples. If the predicted probability

TABLE I
LIST OF NOTATIONS

Notation	Meaning
t	Prediction time horizon
T_c	Critical Time-to-Conflict setting of AV
q	Instantaneous prediction-based traffic flow
Q	Macroscopic traffic flow
v_a	Speed of the AV
v_b	Speed of the conflicting (target) vehicle
l	Length of vehicles
w_v	Width of vehicles
w_l	Width of lanes
d_{lon}	Longitudinal driver space
d_{lat}	Lateral driver space
S	Overall required driver space

at all selected voxels is denoted as a vector $\hat{\mathbf{Y}}$ and the noise-enhanced ground truth vector is \mathbf{Y} , then the objective function is (p is the indices of all selected voxels):

$$L_{\text{fl}}(\hat{\mathbf{Y}}, \mathbf{Y}) = -\frac{1}{P} \sum_p (Y_p - \hat{Y}_p)^2 f(Y_p, \hat{Y}_p) \quad (4)$$

$$f(Y_p, \hat{Y}_p) = \begin{cases} \ln \hat{Y}_p & \text{if } Y_p = 1 \\ (1 - Y_p)^4 \ln(1 - \hat{Y}_p) & \text{else} \end{cases} \quad (5)$$

As the current time and position of the predicted vehicles are arbitrary and the prediction time horizon is specified in the input, any desired temporal and spatial resolution can be used for inference, as shown in Fig. 1.

IV. TRAFFIC EFFICIENCY ASSESSMENT

Having the continuous-time motion prediction model, this section explains how to assess the traffic efficiency from the model output and the vehicle's safety settings. We first introduce the framework and then use IDM to showcase parameter calibration. Notations are listed in Table I.

A. Method Framework

The traffic efficiency assessment method is based on the concept of driver space. For safety and comfort reasons, drivers maintain a speed-dependent minimum 2D spacing around the vehicle [38]. Conflict is thus defined as one vehicle invading another vehicle's driver space in this context. To avoid potential collisions, AVs usually set a key parameter called critical Time-to-Conflict, denoted as T_c . When an AV's planned trajectory intersects with another vehicle's predicted driver space within T_c , this AV starts acting to resolve the predicted conflict. Note that the reaction time is also included. If the predicted position of this target vehicle is uncertain, more road space will be "occupied". This concept is illustrated by the example in Fig. 2. In the early stage of lane-changing, the intention is unclear, so the target vehicle's predicted position has bi-modality. If T_c is long, the target vehicle's predicted location will occupy both lanes and thus require double driver space at this moment. AVs must react to avoid invading the overall driver space. Or equivalently, one target vehicle creates two "virtual" conflicting vehicles in AVs' eyes.

For estimating the traffic efficiency conditioned on T_c , we first define the *instantaneous prediction-based traffic flow*.

²<https://github.com/RomainLITUD/Uncertainty-based-safety-efficiency-tradeoff>




Fig. 2. An example of lane-changing. The red spots are the predicted spatial distribution of the target vehicle after T_c and v_b is the predicted speed (assumed the same for both modes). Blue polygons are required driver space with a total area of S .

Given T_c , at moment τ , the predicted speed of the target vehicle is denoted as v_b and the required driver space is S , the definition of this flow is:

$$q(\tau|T_c) = \frac{v_b}{S} \quad (6)$$

The faster the speed is, and the smaller the required driver space is, the higher the instantaneous efficiency is. This definition is consistent with Edie's flow [55] (or flux). If one vehicle i is continuously tracked for a duration of T_{\max} and predictions are performed for every moment, then its average flow is:

$$\bar{q}_i(T_c) = \frac{1}{T_{\max}} \int_0^{T_{\max}} q_i(\tau|T_c) d\tau \quad (7)$$

which measures this vehicle's impact on traffic efficiency. For example, if the lane change is not decisive, the bi-modality and large driver space status will last longer so the average flow is lower. If enough data is collected, comprising different traffic conditions and diverse interactions (passing orders and driving styles), then the average flow among all recorded vehicles can give the estimate of macroscopic traffic efficiency. Here the lane width w_l converts the unit to conventional pcu/lane/h:

$$Q(T_c) = w_l \times \frac{1}{N} \sum_{i=1}^N \bar{q}_i(T_c) \quad (8)$$

Notice that Eq.(6) is calculated from the minimum required space in a short future, so Eq.(8) represents traffic efficiency in an imaginary extreme situation where all vehicles resolve conflicts according to the minimum spacing and the critical TTC setting. If there is no uncertainty and no conflict, Eq.(8) degrades to the basic stable car-following. This approach is consistent with the methodology proposed by Shi and Li [56], which constructs macroscopic traffic flows from microscopic AV data. We further decompose the real-time measures of spacing into its prediction-based generation process in interactions.

The key step is determining the total required driver space S at each moment and ensuring its consistency with measured flows. For simplicity, we assume that the driver space of a perfectly predictable vehicle is a square along the yaw direction. The longitudinal and lateral widths are denoted as d_{lon} and d_{lat} respectively. First, we define the so-called *preoccupied area*, which encompasses positions where probability density is higher than a threshold β . In practice, β is set to cover 99.9% overall probability:

$$A(T_c) = \{(x, y) \mid p_{T_c}(x, y) \geq \beta \text{ s.t. } \int_{p_{T_c}(x, y) \geq \beta} p_{T_c}(x, y) dx dy > 0.999\} \quad (9)$$

$A(T_c)$ represents where the target vehicle's centre may appear after T_c . For generality, we assume that AVs keep space buffer to a potential conflict point (x_c, y_c) based on speeds, probability density, and critical TTC:

$$\begin{aligned} d_{lon} &= f_{lon}(v_a, v_b, T_c, p_{T_c}(x_c, y_c)) \\ d_{lat} &= f_{lat}(v_a, v_b, T_c, p_{T_c}(x_c, y_c)) \end{aligned} \quad (10)$$

Their specific forms depend on the driving models and the interaction process. All points in $A(T_c)$ must be considered to ensure safety. Therefore, for all $(x_c, y_c) \in A(T_c)$, driver spaces are overlapped to generate the overall required driver space. This area comprises both prediction uncertainty and conflict resolution behaviours. Geometrically, this process equivalently creates a "space buffer" around $A(T_c)$.

For motorway corridors, this approach is straightforward. The relationship between $Q(T_c)$ and speed gives the fundamental diagram and road capacity. However, the road capacity does not hold for intersections because traffic streams cross and some lanes have the priority to pass. Thus, vehicles on some lanes are mostly located at the low-flow end. The maximum flow cannot correctly reflect the intersection's capacity. Many methods are proposed to solve this problem, such as the classical gap acceptance calculation [57], conflict techniques [58], and the fictive traffic light method [59]. In this section, we directly use the overall average flow given in Eq.(8) to measure traffic efficiency at intersections. This is reasonable because the distribution of speed has negligible shift when changing T_c on a given dataset (Time translation does not change the observed velocity distribution). Therefore, $Q(T_c)$ measures the expectation of extreme traffic efficiency under the same speed distribution (or traffic conditions). But one must note that this efficiency is not necessarily the conventional road "capacity".

The description above gives how to estimate traffic efficiency involving prediction uncertainty. The efficiency-safety trade-off can be quantified by changing critical TTC settings and estimating the corresponding traffic flow. A critical step is determining the size of fundamental driver space d_{lon} and d_{lat} . We emphasize that the objective of this study is not any specific control algorithms. Therefore, in the next subsection, Intelligent Driver Models (IDM) [60] is used as a case study to show how to calibrate driver space in different driving scenarios.

B. Case Study: IDM

First of all, the size of driver space depends on driving scenarios. All scenarios are categorized into two types, *crossing area*, where vehicles coming from and going towards different directions have intersected trajectories, and *non-crossing area*, where only car-following, lane-changing, merging, diverging, and over-taking are allowed. We further assume that d_{lon} and d_{lat} are independent of $p_{T_c}(x_c, y_c)$, which means that AVs use the same spacing strategy for all points in the preoccupied area. In other words, $A(T_c)$ is considered a solid obstacle.




Fig. 3. Calculating required driver space from prediction models in non-crossing areas. The longitudinal buffer is set as $d_{lat}/2$ because the space is shared by two vehicles.

The acceleration of an IDM (no reaction time) vehicle is determined by:

$$a = \alpha \left(1 - \left(\frac{v_a}{v_0} \right)^4 - \left(\frac{s^*(v_a, \Delta v)}{s} \right)^2 \right) \quad (11)$$

$$s^*(v_a, \Delta v) = s_0 + vT + \frac{v_a \Delta v}{2\sqrt{\alpha\beta}}$$

Here v_a is the speed of AV, Δv is the speed difference, s is the distance gap, and a is the output acceleration. IDM has 5 parameters. v_0 is desired velocity, s_0 is minimum gap, T is desired time headway, α is maximum acceleration, and β is comfortable deceleration.

For non-crossing areas, car-following is the fundamental driving behaviour. The end of lane-changing, merging, and diverging is also stable car-following on the target lane. Therefore, d_{lon} is set as the desired distance headway. Setting $\Delta v = 0$ and $a = 0$, we have:

$$d_{lon} = \sqrt{\frac{(s_0 + v_a T)^2 \cdot v_0^4}{v_0^4 - v_a^4}} + l \quad (12)$$

where l is the vehicle length. When $v_a \ll v_0$, the relationship is quasi-linear.

The lateral distance d_{lat} is the width of lanes w_l . All lanes intersecting with $A(T_c)$ are considered occupied, as shown in Fig.3. This approximation is reasonable because uncertain multi-modal intentions mainly cause the efficiency drop in non-crossing areas. The impact of lateral movement is also considered implicitly by the integration in Eq.(7). For example, a predictable and decisive lane-changing is highly efficient because the bi-modality and occupied dual lane status (both require larger driver space) last shorter. Conversely, hesitating and slow lane-changing can decrease traffic efficiency by enlarging the average required driver space.

In crossing areas, d_{lon} is set the same as Eq.(12). The lateral space buffer involves two conflicting streams with different speeds. Thus, d_{lat} must be calibrated from controlled simulations. Recall that d_{lat} is the fundamental driver space without prediction uncertainty. We hereby consider a simple and fully deterministic experiment depicted in Fig.4. A black AV and an orange vehicle are approaching a fixed crossing point. Their initial speeds are v_a and v_b , respectively.




Fig. 4. Simulation at a fixed conflicting point. The black vehicle is an AV. The conflict resolution process starts if one vehicle invades another vehicle's front driver space within T_c .

The passing order is determined immediately based on rules. If both vehicles keep their speeds, the vehicle that arrives earlier passes first. If they arrive at the same time, the vehicle with a higher speed passes first. If speeds are also the same, the AV always yields. Conflicts are defined as one vehicle invading another vehicle's front driver space within T_c . The conflict resolution process is also rule-based. The first-go vehicle passes at a constant speed. The yielding vehicle decelerates to avoid predicted conflicts. The conflict point is regarded as a standing object and the deceleration is determined by IDM. The duration for both vehicles passing the conflicting point, denoted as ΔT , can be obtained by simulation. Therefore, the flow for one conflict resolution process is (note that there are 2 vehicles and also two lanes):

$$q_i = \frac{2}{2 \times \Delta T \times w_l} \quad (13)$$

For each combination of v_a , v_b , and T_c , one can easily determine a range of initial positions that trigger conflict resolution. Assuming that the probability distribution within this range is uniform, we can quantify the average flow $\bar{q}(v_a, v_b, T_c)$ from Monte-Carlo simulations. Then using Eq.(6), we have:

$$d_{lat}(v_a, v_b, T_c) = \frac{S}{d_{lon}} = \frac{v_b}{d_{lon} \times \bar{q}} \quad (14)$$

where d_{lon} is given in Eq.(12). Similar to non-crossing areas, this driver space can be applied to all points in $A(T_c)$ and create a buffer around the preoccupied area $A(T_c)$ as shown in Fig.5.

This subsection shows an example of IDM. This can be extended to whatever driving models (e.g. IDM with reaction time, social force, deep-neural networks). We emphasize that driver space is an empirical concept that can always be calibrated. For instance, one can consider interactive models rather than this rule-based IDM, or use Monte-Carlo simulation against human drivers to include heterogeneous negotiation processes. If the driving model is unknown, one can also make assumptions about desired driver space like in traffic flow theory. This paper aims to present the concept. Studying more possibilities is beyond the scope of the paper.

In summary, the advantage of this traffic efficiency assessment method lies in two parts. First, uncertainty in motion prediction is considered in the generation of driver space. Second, simulations, traffic flow theory, and deep-learning prediction models are combined to balance the computational complexity and the utility of real-world data.




Fig. 5. Calculating the total required driver space S in crossing areas. The bottom figure shows that all driver spaces (rectangles) centred in the preoccupied area are overlapped to form the overall required space. The size of the rectangles is calibrated from simulation or computed from driving models.

V. EXPERIMENTS

The experiments require a dataset that provides smooth and precise trajectories. Moreover, the dataset should include labelled interactive scenarios without traffic lights (to have more crossing conflicts). The INTERACTION dataset [61] fulfils these requirements. This drone-based dataset gives high-quality trajectories of vehicles, cyclists, and pedestrians in different types of scenarios, including both motorways and intersections. Trajectories observed in the past 1 s are given and the maximum prediction time horizon is 3 s. The recording frequency is 10Hz.

To normalize the input features, all trajectories and the maps are re-oriented and centred to the target vehicle's last observed position (set as the origin of the coordinate system). The y-axis points in the yaw direction. In training, the spatial resolution of the mesh grid is set as 0.5 m. During inference, this resolution can be set to a smaller value. The details of the models, including the hyper-parameters setting, can be found in the open-source code.

Accuracy benchmarking is not the objective of this study. We refer the readers to Li et al. [50] for more details. To show the reliability of uncertainty quantification, we only compare the proposed model with the backbone UQnet. Three metrics are used. *Final Displacement Error* (FDE) is the L2-squared distance between the ground truth and the prediction at 3 s. *Average Displacement Error* (ADE), similarly, is the average L2-squared error along the predicted trajectory. *Missing Rate* (MR) originates from the concept of risk field [62]. If the predicted position is outside of this area around the ground truth, potential conflict may be triggered. This is called "missed". MR measures the percentage of missed predictions. The risk field is assumed to be a rectangle for the INTERACTION dataset. The lateral threshold is 1 m and the longitudinal

TABLE II
PERFORMANCE COMPARISON

Metrics	ADE (m)	FDE (m)	MR (%)
Regular			
Ours	0.387	0.615	1.56
UQnet [50]	0.392	0.621	1.96
Generalizability			
Ours	0.675	0.860	6.83
UQnet	0.679	0.861	6.86
All			
Ours	0.486	0.699	3.37
UQnet	0.490	0.704	3.64

threshold is a piece-wise function depending on the velocity:

$$\text{th}(v) = \begin{cases} 1 & v < 1.4 \text{ m s}^{-1} \\ 1 + \frac{v - 1.4}{11 - 1.4} & 1.4 \text{ m s}^{-1} \leq v \leq 11 \text{ m s}^{-1} \\ 2 & v > 11 \text{ m s}^{-1} \end{cases} \quad (15)$$

For each prediction, 6 trajectories are sampled and the result is measured as the one with the lowest error. Table II shows that using the proposed training strategy and extending the goal-oriented scheme to a continuous-time scheme do not damage the performance. The proposed model is on par with the state-of-the-art so the following analysis is reliable.

VI. ANALYSIS AND DISCUSSION

This section focuses on analyzing the results. We will present the uncertainty-horizon profile and discuss the trade-off between safety, traffic efficiency, and prediction time horizon.

A. Uncertainty-Horizon Profiles

This subsection first analyzes how spatial uncertainty changes with the prediction time horizon. Besides preoccupied area $A(t)$, *differential entropy* is also used to measure the amount of randomness in the random variable of predicted locations, which is related to the lower bound of prediction error [63]:

$$H_t(X, Y) = - \int_{\mathcal{X}, \mathcal{Y}} p_t(x, y) \ln p_t(x, y) dx dy \quad (16)$$

Fig. 6 shows how the average preoccupied area $\bar{A}(t)$ and the average entropy $\bar{H}(t)$ change with time horizons. The uncertainty is almost 0 in the short term but increases rapidly after 1-1.5 s. Fig. 7 further presents the long-tail distributions of $A(t = 3 \text{ s})$ and $H(t = 3 \text{ s})$. Most samples are located at the low uncertainty end, which means that most cases are highly predictable within 3 s.

According to the speed-threshold function in Eq. (15), we hereby identify the *critical moment* when the predicted occupied area surpasses the threshold (note that this is different from the critical TTC). After this moment, at least two trajectories must be given to ensure 0 MR. This can be regarded as a critical decision point. Similarly, we can




Fig. 6. Relationships between (left) average occupied space, (right) average differential entropy, and prediction horizon t .




Fig. 7. Distributions of preoccupied area $A(t = 3 \text{ s})$ and differential entropy $H(t = 3 \text{ s})$.




Fig. 8. The distribution of estimated critical moments. The high column at 3 s means equal to or higher than 3 s.

identify the critical position, defined as where the vehicle will most probably appear at the critical moment on the map. The distribution of critical moments is presented in Fig. 8. Around 20% of cases have critical moments that are equal to or longer than 3 s. In all cases, trajectories are highly predictable within 1 s. The distribution shows the diversity of uncertainty profiles from another aspect. Fig. 9 shows 3 representative examples. The upper two cases have strong multi-modality. Each mode represents an intended direction. The uncertainty increases fast after passing the decision point. In the bottom case, the red vehicle is approaching a congested merging area at a relatively high velocity. The vehicle will slow down soon. The future position is thus sensitive to deceleration. The uncertainty also increases fast after entering the merging zone. The 3 examples represent the major factors causing high uncertainty, namely the incoming decision point, braking with high initial speed, and transition of traffic conditions.

B. The Trade-off Between Safety and Efficiency for IDM

This subsection presents the trade-off between safety and efficiency using the model output and the IDM assumption. Huang et al. [64] calibrated IDMs from the Waymo open dataset [3] and we directly adapt the reported parameters. v_0 is

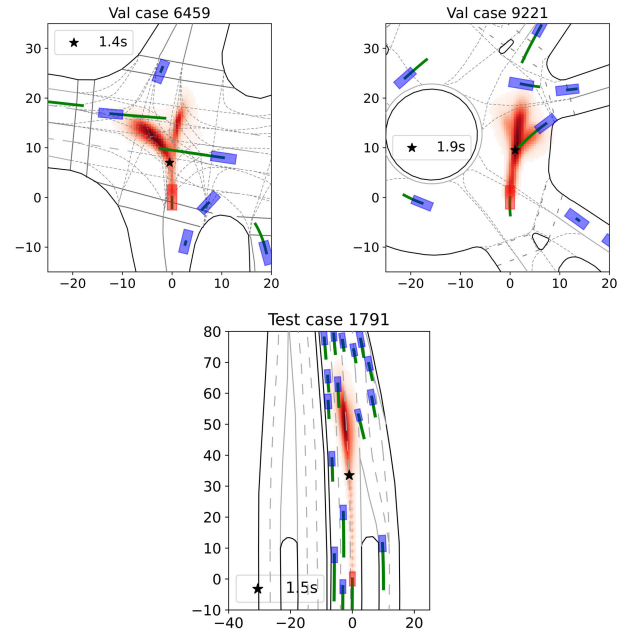


Fig. 9. Examples of trajectory predictions and uncertainty profiles. The black stars are identified as critical positions.

TABLE III
IDM PARAMETERS FOR WAYMO AV

IDM Parameter	Value	Other parameter	Value
$\alpha \text{ (m s}^{-2}\text{)}$	2.53	$l \text{ (m)}$	4.8
$\beta \text{ (m s}^{-2}\text{)}$	5.43	$w_v \text{ (m)}$	1.8
$T \text{ (s)}$	1.0	$w_l \text{ (m)}$	3.25
$s_0 \text{ (m)}$	3.49		
$v_0 \text{ (m s}^{-1}\text{), highway}$	33.3		
$v_0 \text{ (m s}^{-1}\text{), urban}$	20		

changed to the speed limit in the INTERACTION dataset. The used parameters are listed in Table III. Eq. (12) and (14) give d_{lon} and d_{lat} from v_a , v_b , and T_c via calculation or controlled simulations.

For non-crossing areas on motorways, derived fundamental diagrams and road capacities are shown in Fig. 10. In the left figure, the black line represents theoretical stable car-following with no interactions and perfect predictability. The capacity can reach around 2400 pcu/lane/h at 20.6 m s^{-1} . The blue scatters are derived from the INTERACTION dataset using the prediction model and predefined critical TTC, which is set as 2.5 s as an example. The blue line is the corresponding fitted fundamental diagram. We see that the capacity drops to around 1750 pcu/lane/h due to prediction uncertainty and conflict resolution. The corresponding optimal speed is also reduced to 16.6 m s^{-1} because prediction uncertainty increases more rapidly at high speeds. This means that both road capacity and average travel time decrease. The right figure directly shows the trade-off between road capacity and critical TTC settings. The capacity keeps horizontal when T_c is shorter than 1 s. When T_c is longer than 1.5 s, traffic efficiency drops rapidly.

In crossing areas at intersections, as explained in Sec. IV-A, average flow is used to measure traffic efficiency. Fig. 11 presents the trade-off relationship between flow and critical TTC. Similarly, traffic efficiency starts decreasing

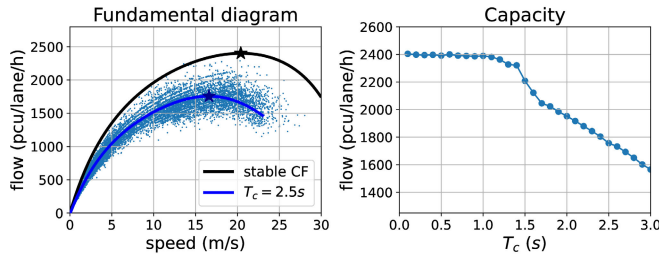


Fig. 10. Left: comparison between stable CF and an example of 2.5 s TTC setting. The efficiency-safety trade-off relationship in non-crossing areas, measured by critical TTC and road capacity.




Fig. 11. The efficiency-safety trade-off relationship in crossing areas, measured by critical TTC and average flow.

when the critical TTC setting is longer than 1.5 s. This 1.5 s law is universal in mixed traffic, no matter what driving models are used. The reason is that human behaviours are almost perfectly predictable within this horizon, as shown in Fig.6 and Fig.8. Nevertheless, compared to non-crossing areas, traffic flow drops faster at intersections. Compared to $T_c = 0$, when $T_c = 3$ s, traffic efficiency drops by 55% in the crossing area but only by 30% in the non-crossing area. This is because intersections involve more decision points so uncertainty increases significantly faster, causing lower traffic efficiency.

The results and analysis above confirm the significant role of motion prediction uncertainty in quantifying the safety-efficiency trade-off relationship. In summary, AVs cannot always improve traffic efficiency by resolving nearer conflicts. When the critical horizon is shorter than 1.5 s, driving more aggressively sacrifices safety but cannot gain traffic efficiency. Additionally, due to traffic efficiency concerns, AVs can forecast further ahead to resolve conflicts in non-crossing scenarios but have to shorten the critical horizon (be myopic) at intersections.

C. Implications for AV Design

In the previous section, the safety-efficiency trade-off is quantified using one specific IDM. Now we consider a generic inverse problem that is independent of the specific driving model:

Amongst critical TTC, space-keeping strategy, and macroscopic traffic efficiency, with given constraints on one of them, how to trade off the other two aspects when designing the key parameters of AVs?

As motorway corridors have already been intensively studied in the literature, we are particularly interested in crossing areas at unsignalized intersections.

Conceptually, we consider an *average* lateral space buffer to represent the average loss of efficiency in resolving conflicts




Fig. 12. The relationship between average traffic flow, critical TTC T_c , and an average lateral space buffer d_{lat} . The blue line represents that conflicts happen between two HVs; the red line represents HV-AV conflicts. These are just two examples.

between two crossing streams (under a certain speed distribution, such as in the used datasets). Such an average lateral space always exists, which makes the flow estimated from this average value the same as the one estimated using the spacing in the real case (mean value theorem). Its value depends on the driving model. Then the relationship between this average lateral space, critical TTC, and traffic flow can be quantified using the proposed method.

The heatmap in Fig.12 shows the result. The isolines are almost vertical when the average lateral buffer is large. It suggests that only excellent conflict resolution capability can significantly benefit the efficiency of crossing streams at intersections. This figure also shows the comprehensive trade-off relationship between the involved three aspects. For example, when traffic efficiency is constrained, the isolines quantify how to balance lateral space-keeping and critical TTC at intersections.

Li et al. [65] derive a conflict resolution dataset at unsignalized intersections from the open Argoverse-2 AV motion data [5]. The authors report that the traffic flows for HV-HV and AV-HV conflict resolution cases are 663 and 576 pcu/lane/h, respectively. We lend these results and also plot their isolines on the heatmap. Assume that the critical TTC for Argoverse AV is 2.5 s, as the red star marks. If we want to ensure that adding AVs to the traffic flow does not make traffic efficiency worse, then there are two major improvement directions. If critical TTC stays unchanged, then the average lateral space buffer must be reduced from 8.8 m to 5.8 m, which means that the crossing conflict resolution efficiency needs to be improved by 34% on average to increase the traffic flow from 576 to 663 pcu/lane/h (by 15% only, the difference is caused by prediction uncertainty). The other option is to sacrifice critical TTC, which requires faster and more decisive actions. The exact choice depends on the difficulty in both directions. This example illustrates how the proposed method helps engineers design and tune AV's key parameters and performances considering macroscopic traffic efficiency, which has been barely addressed in the literature.

To summarize, we come back to the primary research question about when to resolve predicted conflicts. The answer depends on three major factors, the desired traffic efficiency (social and traffic requirements), the capability of AVs, and different scenarios. Lower efficiency requirements and a more aggressive space-keeping strategy necessitate earlier actions (longer time horizon). Conversely, if we want to keep traffic

efficiency when integrating AVs into the current traffic system, Fig. 12 gives insights into how to tune the parameters in an AV control module for unsignalized intersections. Currently, AVs have flaws in the perception module and the control algorithms are safety-first. AVs hence prefer to drive more defensively and act earlier. For example, Hu et al. [66] report that Waymo AVs will potentially reduce traffic efficiency due to conservative driving styles. Deploying this type of AV in current urban traffic may reduce road capacity and cause more congestion.

VII. CONCLUSION AND PERSPECTIVE

In this paper, we propose a deep-learning model for estimating the continuous-time spatial uncertainty in motion prediction and reconsider how prediction uncertainty influences traffic efficiency and the safety setting of AVs. Experiments show that the model can effectively identify the time and the location of critical points and thus indicate the maximum horizon that the prediction stays reliable. In addition, we explored and built the quantitative trade-off relationships between safety, traffic efficiency, and prediction uncertainty by using prediction-based driver space and simulations. This study provides a fundamental and extendable approach for deciding when to act to resolve conflicts.

More realistic and detailed models can be implemented into this framework in further studies, for example, the impact of the penetration rate of AVs. Most results in this paper give the contribution of the AV-HV interaction component of the overall traffic efficiency. If we calibrate a set of parameters for interaction between different types of agents (AV-AV, HV-HV, HV-AV) respectively, and find the relationship between their ratios and the penetration rate of AVs, the overall traffic efficiency can be assessed by, e.g., weighted average.

Based on these conclusions, we hereby suggest several relevant research topics. First, the uncertainty quantification is based on the currently available information provided by the perception module of AVs. If AVs can effectively detect e.g. turn signals, the prediction uncertainty is expected to decrease so the efficiency-horizon trade-off can be improved. Studying what information is indeed needed to forecast trajectories is more important than modelling itself. Second, we admit that this study simplifies the problem and focuses on normal cases. How emergent cases and other dimensions, such as the comfort of AVs, the string stability of car-following behaviours, and the queuing at intersections may influence the relationship needs more investigation. Finally, how the proactive interactions between AVs and human drivers through communication technology influence the conclusion also needs deeper studies. They can be integrated into the proposed framework. We believe that these research topics are critical for helping integrate in-developing AVs into current urban traffic and mitigate the negative impact.

APPENDIX A

VALIDATION OF THE TRAINING STRATEGY

For validating the proposed training strategy, we consider the following methods and compare their prediction qualities. The concept is illustrated in Fig. 13.




Fig. 13. The proposed training strategy: a 2D projection. The colour map represents the probability density. Red means high value and blue means low value. σ_s is the standard variance of the added noise. r_s is the width of the selected area. t_s is the randomly sampled prediction horizon.

- S-1 The first strategy is that the moment t_s is randomly sampled between 0 and the maximum prediction time during training. We only consider the voxels located on this particular spatial slice. This method has the same complexity as the goal-oriented approach ($50 \times 50 = 2500$ voxels per sample). However, it is not efficient for two primary reasons. First, it cannot ensure the continuity of the distribution along the time axis. Second, due to the sparsity of the sampled moments, longer training periods are necessary to comprehensively cover every time step.
- S-2 The second strategy is exclusively considering those voxels located around the labels (the hatch area in Fig. 13). For example, if the range is $r_s = 5$ m, this method includes $5 \times 5 \times 30 = 750$ voxels per sample. Due to the reduced number of voxels that are all informative, this training strategy exhibits faster convergence and smoother distribution evolution. However, models trained by this method tend to be over-fitted because of the local vision.
- S-3 (Used) The third strategy involves a small spatiotemporal range around the sampled ground truth (the rectangle in Fig. 13). The added rectangle area includes local voxels from multiple adjacent timestamps and thus addresses the efficiency and continuity issues of S-1. The involved considerable number of zero-value voxels can mitigate the over-fitting problem of S-2.

In addition to accuracy, it is imperative to compare the continuity of the distribution series generated by various training strategies over time. Considering the smoothness of trajectories, the limited complexity of traffic rules and road connectivity, as well as the short prediction horizon, it is reasonable to presume that the mode series (positions that maximize probability density) ought to be continuous and smooth. So first, we sample the smooth mode series through tri-cubic interpolation [67] and the Bezier smoothing algorithm. Next, the acceleration of each of the generated mode series is computed, and a comparison is made against the predetermined limit (5.89 m s^{-2} from the training set) to ascertain whether the maximum value exceeds the feasible range. High acceleration means discontinuity. The percentage of the mode series exhibiting excessively high acceleration values serves as an indicator of the predictions' continuity. We call it the *Anomaly ratio* (noted as Anomaly%).

Table IV shows the MR and the Anomaly% of 3 models trained by different strategies after 10 epochs (the dataset

TABLE IV
COMPARISON OF THE TRAINING STRATEGIES

Strategy	S-1	S-2	S-3
Regular MR%	1.60	1.60	1.56
Generalizability MR%	7.19	11.46	6.83
All MR%	3.52	4.98	3.37
Anomaly% (validation)	4.91	3.08	3.21

is iterated 10 times). The MR is measured on the test set and the Anomaly% is measured on the validation set only. The regular MRs of the three methods are quite similar; however, the generalizability MR (for cases from scenarios that have not been seen in training) of S-2 is notably higher. While S-2 and S-3 have similar Abnormal%, S-1 shows relatively lower continuity. These findings indicate that S-3 can effectively address both the issues of over-fitting and discontinuity, simultaneously using relatively lower memory requirements.

REFERENCES

- [1] W. Schwarting, J. Alonso-Mora, and D. Rus, "Planning and decision-making for autonomous vehicles," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 1, no. 1, pp. 187–210, 2018.
- [2] Y. Kuwata, J. Teo, S. Karaman, G. Fiore, E. Frazzoli, and J. How, "Motion planning in complex environments using closed-loop prediction," in *Proc. AIAA Guid., Navigat. Control Conf. Exhib.*, Aug. 2008, p. 7166.
- [3] P. Sun et al., "Scalability in perception for autonomous driving: Waymo open dataset," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2020, pp. 2446–2454.
- [4] H. Caesar et al., "nuScenes: A multimodal dataset for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 11621–11631.
- [5] B. Wilson et al., "Argoverse 2: Next generation datasets for self-driving perception and forecasting," 2023, *arXiv:2301.00493*.
- [6] S. Mozaffari, O. Y. Al-Jarrah, M. Dianati, P. Jennings, and A. Mouzakitis, "Deep learning-based vehicle behavior prediction for autonomous driving applications: A review," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 1, pp. 33–47, Jan. 2020.
- [7] Y. Huang, J. Du, Z. Yang, Z. Zhou, L. Zhang, and H. Chen, "A survey on trajectory-prediction methods for autonomous driving," *IEEE Trans. Intell. Vehicles*, vol. 7, no. 3, pp. 652–674, Sep. 2022.
- [8] C. Tang, *Designing Explainable Autonomous Driving System for Trustworthy Interaction*. Berkeley, CA, USA: Univ. California, Berkeley, 2022.
- [9] A. Bazzi et al., "A Hardware-in-the-Loop evaluation of the impact of the V2X channel on the traffic-safety versus efficiency trade-offs," in *Proc. 14th Eur. Conf. Antennas Propag. (EuCAP)*, Mar. 2020, pp. 1–5.
- [10] A. D. Kiureghian and O. Ditlevsen, "Aleatory or epistemic? Does it matter?" *Structural Saf.*, vol. 31, no. 2, pp. 105–112, Mar. 2009.
- [11] J. Zhang, D. Kumor, and E. Bareinboim, "Causal imitation learning with unobserved confounders," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 12263–12274.
- [12] M. Saifuzzaman and Z. Zheng, "Incorporating human-factors in car-following models: A review of recent developments and research needs," *Transp. Res. C, Emerg. Technol.*, vol. 48, pp. 379–403, Nov. 2014.
- [13] J. W. C. van Lint and S. C. Calvert, "A generic multi-level framework for microscopic traffic simulation—Theory and an example case in modelling driver distraction," *Transp. Res. B, Methodol.*, vol. 117, pp. 63–86, Nov. 2018.
- [14] A. Sarker et al., "A review of sensing and communication, human factors, and controller aspects for information-aware connected and automated vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 1, pp. 7–29, Jan. 2020.
- [15] R. West, D. French, R. Kemp, and J. Elander, "Direct observation of driving, self reports of driver behaviour, and accident involvement," *Ergonomics*, vol. 36, no. 5, pp. 557–567, May 1993.
- [16] S. Khaitan, Q. Lin, and J. M. Dolan, "Safe planning and control under uncertainty for self-driving," *IEEE Trans. Veh. Technol.*, vol. 70, no. 10, pp. 9826–9837, Oct. 2021.
- [17] T. U. Yokota, "Evaluation of ahs effect on mean speed by static method," in *Proc. Towards New Horizon Together. 5th World Congr. Intell. Transp. Syst., Held*, Seoul South Korea, Oct. 1998, pp. 135–283, Paper 3201.
- [18] B. Lewis-Evans, D. De Waard, and K. A. Brookhuis, "That's close enough—A threshold effect of time headway on the experience of risk, task difficulty, effort, and comfort," *Accident Anal. Prevention*, vol. 42, no. 6, pp. 1926–1933, Nov. 2010.
- [19] T. Gilles, S. Sabatini, D. Tsishkou, B. Stanculescu, and F. Moutarde, "GOHOME: Graph-oriented heatmap output for future motion estimation," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2022, pp. 9107–9114.
- [20] J. Gu, C. Sun, and H. Zhao, "DenseTNT: End-to-end trajectory prediction from dense goal sets," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, Oct. 2021, pp. 15303–15312.
- [21] X. Huang, S. G. McGill, B. C. Williams, L. Fletcher, and G. Rosman, "Uncertainty-aware driver trajectory prediction at urban intersections," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2019, pp. 9718–9724.
- [22] C. Wang, Y. Xie, H. Huang, and P. Liu, "A review of surrogate safety measures and their applications in connected and automated vehicles safety modeling," *Accident Anal. Prevention*, vol. 157, Jul. 2021, Art. no. 106157.
- [23] K. Vogel, "A comparison of headway and time to collision as safety indicators," *Accident Anal. Prevention*, vol. 35, no. 3, pp. 427–433, May 2003.
- [24] A. Bose and P. Ioannou, "Mixed manual/semi-automated traffic: A macroscopic analysis," *Transp. Res. C, Emerg. Technol.*, vol. 11, no. 6, pp. 439–462, Dec. 2003.
- [25] Z. Vander Laan and P. Schonfeld, "Modeling heterogeneous traffic with cooperative adaptive cruise control vehicles: A first-order macroscopic perspective," *Transp. Planning Technol.*, vol. 43, no. 2, pp. 113–140, Feb. 2020.
- [26] J. B. Michael, D. N. Godbole, J. Lygeros, and R. Sengupta, "Capacity analysis of traffic flow over a single-lane automated highway system," *J. Intell. Transp. Syst.*, vol. 4, nos. 1–2, pp. 49–80, 1998.
- [27] C. Zhao, L. Li, X. Pei, Z. Li, F.-Y. Wang, and X. Wu, "A comparative study of state-of-the-art driving strategies for autonomous vehicles," *Accident Anal. Prevention*, vol. 150, Feb. 2021, Art. no. 105937.
- [28] Q. Wang, Z. Li, and L. Li, "Investigation of discretionary lane-change characteristics using next-generation simulation data sets," *J. Intell. Transp. Syst.*, vol. 18, no. 3, pp. 246–253, Jul. 2014.
- [29] A. Ghiasi, O. Hussain, Z. Qian, and X. Li, "A mixed traffic capacity analysis and lane management model for connected automated vehicles: A Markov chain method," *Transp. Res. B, Methodol.*, vol. 106, pp. 266–292, Dec. 2017.
- [30] H. Yang and K. Oguchi, "Intelligent vehicle control at signal-free intersection under mixed connected environment," *IET Intell. Transp. Syst.*, vol. 14, no. 2, pp. 82–90, 2020.
- [31] R. Hult, M. Zanon, S. Gros, H. Wymeersch, and P. Falcone, "Optimisation-based coordination of connected, automated vehicles at intersections," *Vehicle Syst. Dyn.*, vol. 58, no. 5, pp. 726–747, May 2020.
- [32] J. Zhao, V. L. Knoop, and M. Wang, "Microscopic traffic modeling inside intersections: Interactions between drivers," *Transp. Sci.*, vol. 57, no. 1, pp. 135–155, Jan. 2023.
- [33] D. Helbing, "Derivation of a fundamental diagram for urban traffic flow," *Eur. Phys. J. B*, vol. 70, no. 2, pp. 229–241, Jul. 2009.
- [34] X. Wu, H. X. Liu, and N. Geroliminis, "An empirical analysis on the arterial fundamental diagram," *Transp. Res. B, Methodol.*, vol. 45, no. 1, pp. 255–266, Jan. 2011.
- [35] X. Chen, M. Treiber, V. Kanagaraj, and H. Li, "Social force models for pedestrian traffic—state of the art," *Transp. Rev.*, vol. 38, no. 5, pp. 625–653, Sep. 2018.
- [36] S. P. Hoogendoorn, F. L. M. van Wageningen-Kessels, W. Daamen, and D. C. Duives, "Continuum modelling of pedestrian flows: From microscopic principles to self-organised macroscopic phenomena," *Phys. A, Stat. Mech. Appl.*, vol. 416, pp. 684–694, Dec. 2014.
- [37] L. D. Vanumu, K. R. Rao, and G. Tiwari, "Fundamental diagrams of pedestrian flow characteristics: A review," *Eur. Transp. Res. Rev.*, vol. 9, no. 4, pp. 1–13, Dec. 2017.
- [38] Y. Jiao, S. C. Calvert, S. van Cranenburgh, and H. van Lint, "Inferring vehicle spacing in urban traffic from trajectory data," *Transp. Res. C, Emerg. Technol.*, vol. 155, Oct. 2023, Art. no. 104289.

- [39] M. Schreier, V. Willert, and J. Adamy, "Bayesian, maneuver-based, long-term trajectory prediction and criticality assessment for driver assistance systems," in *Proc. 17th Int. IEEE Conf. Intell. Transp. Syst. (ITSC)*, Qingdao, China, Oct. 2014, pp. 334–341.
- [40] Q. Tran and J. Firl, "Online maneuver recognition and multimodal trajectory prediction for intersection assistance using non-parametric regression," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2014, pp. 918–923.
- [41] M. T. Abbas, M. A. Jibran, M. Afaq, and W. Song, "An adaptive approach to vehicle trajectory prediction using multimodel Kalman filter," *Trans. Emerg. Telecommun. Technol.*, vol. 31, no. 5, May 2020, Art. no. e3734.
- [42] L. Hewing, E. Arcari, L. P. Fröhlich, and M. N. Zeilinger, "On simulation and trajectory prediction with Gaussian process dynamics," 2019, *arXiv:1912.10900*.
- [43] J. Wiest, M. Höffken, U. Kresel, and K. Dietmayer, "Probabilistic trajectory prediction with Gaussian mixture models," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2012, pp. 141–146.
- [44] N. Deo and M. M. Trivedi, "Multi-modal trajectory prediction of surrounding vehicles with maneuver based LSTMs," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2018, pp. 1179–1184.
- [45] S. Eiffert, K. Li, M. Shan, S. Worrall, S. Sukkarieh, and E. Nebot, "Probabilistic crowd GAN: Multimodal pedestrian trajectory prediction using a graph vehicle-pedestrian attention network," *IEEE Robot. Autom. Lett.*, vol. 5, no. 4, pp. 5026–5033, Oct. 2020.
- [46] Y. Xing, C. Lv, and D. Cao, "Personalized vehicle trajectory prediction based on joint time-series modeling for connected vehicles," *IEEE Trans. Veh. Technol.*, vol. 69, no. 2, pp. 1341–1352, Feb. 2020.
- [47] A. Zyner, S. Worrall, and E. Nebot, "Naturalistic driver intention and path prediction using recurrent neural networks," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 4, pp. 1584–1594, Apr. 2020.
- [48] J. Mercat, T. Gilles, N. El Zoghby, G. Sandou, D. Beauvois, and G. P. Gil, "Multi-head attention for multi-modal joint vehicle motion forecasting," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2020, pp. 9638–9644.
- [49] T. Gilles, S. Sabatini, D. Tishkhou, B. Stanculescu, and F. Moutarde, "THOMAS: Trajectory heatmap output with learned multi-agent sampling," 2021, *arXiv:2110.06607*.
- [50] G. Li, Z. Li, V. Knoop, and H. van Lint, "UQNet: Quantifying uncertainty in trajectory prediction by a non-parametric and generalizable approach," 2022. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4241523
- [51] J. Gao et al., "VectorNet: Encoding HD maps and agent dynamics from vectorized representation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 11525–11533.
- [52] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," 2017, *arXiv:1710.10903*.
- [53] M. Liang et al., "Learning lane graph representations for motion forecasting," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 541–556.
- [54] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2980–2988.
- [55] L. C. Edie, "Discussion of traffic stream measurements and definitions," in *Proc. 2nd Int. Symp. Theory Traffic Flow*. Paris, France: OECD, 1956.
- [56] X. Shi and X. Li, "Constructing a fundamental diagram for traffic flow with automated vehicles: Methodology and demonstration," *Transp. Res. B, Methodol.*, vol. 150, pp. 279–292, Aug. 2021.
- [57] R. M. Kimber and R. D. Coombe, "The traffic capacity of major/minor priority junctions," *Transp. Res. Lab., Tech. Rep. SR 582 Monograph*, 1980.
- [58] W. Brilon and N. Wu, "Capacity at unsignalized intersections derived by conflict technique," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 1776, no. 1, pp. 82–90, Jan. 2001.
- [59] E. Chevallier and L. Leclercq, "A macroscopic theory for unsignalized intersections," *Transp. Res. B, Methodol.*, vol. 41, no. 10, pp. 1139–1150, Dec. 2007.
- [60] A. Kesting, M. Treiber, and D. Helbing, "Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity," *Philosophical Trans. Roy. Soc. A, Math., Phys. Eng. Sci.*, vol. 368, no. 1928, pp. 4585–4605, Oct. 2010.
- [61] W. Zhan et al., "INTERACTION dataset: An INTERNATIONAL, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps," 2019, *arXiv:1910.03088*.
- [62] S. Kolekar, J. de Winter, and D. Abbink, "Human-like driving behaviour emerges from a risk-based driver model," *Nature Commun.*, vol. 11, no. 1, pp. 1–13, Sep. 2020.
- [63] G. Li, V. L. Knoop, and H. van Lint, "Estimate the limit of predictability in short-term traffic forecasting: An entropy-based approach," *Transp. Res. C, Emerg. Technol.*, vol. 138, May 2022, Art. no. 103607.
- [64] Y. Huang, Y. Ye, J. Sun, and Y. Tian, "Characterizing the impact of autonomous vehicles on macroscopic fundamental diagrams," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 6, pp. 6530–6541, Jun. 2023.
- [65] G. Li, Y. Jiao, S. C. Calvert, and J. W. C. van Lint, "A comparative conflict resolution dataset derived from Argoverse-2: Scenarios with vs. Without autonomous vehicles," 2023, *arXiv:2308.13839*.
- [66] X. Hu, Z. Zheng, D. Chen, and J. Sun, "Autonomous vehicle's impact on traffic: Empirical evidence from waymo open dataset and implications from modelling," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 6, pp. 6711–6724, Jun. 2023.
- [67] F. Lekien and J. Marsden, "Tricubic interpolation in three dimensions," *Int. J. Numer. Methods Eng.*, vol. 63, no. 3, pp. 455–471, 2005.



Guopeng Li received the M.S. and Engineering Diploma degrees in applied mathematics from ENSTA-Paris in 2018 and the Ph.D. degree (cum laude) from Delft University of Technology (TU Delft) in 2023 on the data-driven uncertainty quantification in multiscale traffic systems. He was a Post-Doctoral Researcher with TU Delft from March 2023 to February 2024, working on safety and efficiency impact assessment of automated vehicles. Since March 2024, he has been a Post-Doctoral Researcher with the Faculty of Mechanical Engineering, TU Delft. His current research interests include developing epistemic AI methods for safe and trustworthy automated driving.



Zirui Li received the B.S. degree from Beijing Institute of Technology (BIT), Beijing, China, in 2019, where he is currently pursuing the Ph.D. degree in mechanical engineering. From June 2021 to July 2022, he was a Visiting Researcher with Delft University of Technology (TU Delft). Since August 2022, he has been a Visiting Researcher with the Chair of Traffic Process Automation, Faculty of Transportation and Traffic Sciences "Friedrich List," Dresden University of Technology (TU Dresden). His research interests include interactive behavior modeling, risk assessment, and motion planning of automated vehicles.



Victor Lambert Knoop was born in The Netherlands in 1981. He received the M.S. degree in physics from Leiden University in 2005 and the Ph.D. degree from Delft University of Technology in 2009 on the effects of incidents on driving behavior and traffic congestion. From 2009 to 2010, he was a Post-Doctoral Researcher with the University of Lyon on lane changing. He was with Imperial College London and the University of California at Berkeley. Since 2018, he has been a tenured Associate Professor with the Transport and Planning Department, Delft University of Technology. He is currently the Co-Director of the Traffic Dynamics Modelling and Control Laboratory (TDMaC-Lab). His main research interests include interaction between microscopic and macroscopic traffic flow phenomena.



J.W.C. van Lint received the M.Sc. degree in civil engineering informatics and the Ph.D. degree in transportation from Delft University of Technology (TU Delft) in 1997 and 2004, respectively. He was an Information Analyst and a Transport Engineer at various organizations. After receiving the Ph.D. degree, he was appointed as an Anthonie van Leeuwenhoek (AvL) Full Professor for the Executive Board of TU Delft in 2013 (an honor reserved for only a few young scientists and educators). He is currently the Director of the Data Analytics and Traffic Simulation Laboratory (DiTTlab) in the Faculty of Civil Engineering and Geosciences. He is active in many international projects and collaborations.