

A Probabilistic Framework
for the Automated Analysis of the Exposure to Road Collision

Paper Submission to the
2008 TRB Annual Meeting

Nicolas Saunier
Postdoctoral Research Associate
Department of Civil Engineering
The University of British Columbia
6250 Applied Science Lane
Vancouver, BC, Canada V6T 1Z4

Tarek Sayed
Professor
Department of Civil Engineering
The University of British Columbia
6250 Applied Science Lane
Vancouver, BC, Canada V6T 1Z4

Word Count: (5572 Text) +(0 Tables = 0)+ (6 Figures = 1500) = 7072

A Probabilistic Framework for the Automated Analysis of the Exposure to Road Collision

Nicolas Saunier, Tarek Sayed

Abstract: The advent of powerful sensing technologies, especially video sensors and computer vision techniques, has allowed for the collection of large quantities of detailed traffic data. They allow us to further advance towards completely automated systems for road safety analysis. This paper presents a comprehensive probabilistic framework for automated road safety analysis. Building upon traffic conflict techniques and the concept of the safety hierarchy, it provides computational definitions of the probability of collision for road users involved in an interaction. It proposes new definitions for individual road users and aggregated measures over time. This allows the interpretation of traffic from a safety perspective, studying all interactions and their relationship to safety. New and more relevant exposure measures can be derived from this work, and traffic conflicts can be detected. A complete vision-based system is implemented to demonstrate the approach, providing experimental results on real world video data.

INTRODUCTION

Road safety is characterized by the absence of accidents, i.e. collisions between road users. The safety is traditionally measured by the number of collisions, or rather its expected number at a given time. Traffic safety diagnosis has been traditionally undertaken using historical collision data. However, there are well-recognized problems of availability and quality associated with collision data. Additionally, the use of collision records for safety analysis is a reactive approach: a significant number of collisions has to be recorded before action is taken.

Therefore there has been considerable interest in research dealing with surrogate safety measures (1). The observation of traffic conflicts has been advocated as an alternative or complementary approach to analyze traffic safety from a broader perspective than collision statistics alone (2, 3, 4, 5, 6). Traffic conflicts are interactions with very similar processes to collisions, but without a collision. A conflict is defined as “an observational situation in which two or more road users approach each other in space and time to such an extent that a collision is imminent if their movements remain unchanged” (7). The concept of collision course is derived from this widely accepted definition of traffic conflicts. In (8), users are defined to be on a collision course when, “unless the speed and/or the direction of the road users changes, they will collide”. Deciding if two road users are on a collision course thus depends on extrapolation hypotheses. The definition of (8) uses the common hypothesis of extrapolation with constant velocity, i.e. speed and direction. Some definitions of traffic conflicts also include that at least one of the road users involved takes a successful evasive action, often in emergency.

The relationship between traffic conflicts and collisions must be established to use traffic conflicts as surrogates to collisions for safety analysis. Most road users move freely in the traffic, without having to take into account other road users. Otherwise, road users interact with each other. An interaction is defined as a situation in which two or more road users are close enough in space and time, and their distance is decreasing. Many researchers, especially in Scandinavian countries (2, 6), assume that all interactions can be ranked in a safety hierarchy, with collisions at the top. The interactions located next to the collisions in the safety hierarchy, very similar to collisions, but without an actual collision, can be called quasi-collisions. The interactions can thus be recursively ranked in the safety hierarchy. One can imagine “quasi-quasi-collisions”, and so on... The pyramid shape of the hierarchy stands for the frequency of events. The further from collisions, the less severe and more frequent the events.

For this concept to be operational, the safety hierarchy is transferred into measurable parameters based on certain assumptions. For each interaction in the hierarchy, a severity can be estimated, matching its position in the hierarchy, i.e. measuring the proximity to the potential occurrence of a collision, which can be interpreted as the probability of collision. Many severity indicators, such as the Time-To-Collision (TTC) and the Post-Encroachment Time (PET), have been developed to evaluate the distance in space and time between the vehicles involved and their evasive action(s) (9, 10). Traffic Conflict Techniques (TCTs) involve observing and evaluating the frequency and severity of traffic conflicts at an intersection by a team of trained observers.

Several automated systems, using mostly video sensors, have been and are being developed for traffic monitoring (11, 12, 13, 14, 15). In a previous paper (16), the authors have shown that traffic conflicts can be detected in video sequences using HMM-based semi-supervised machine learning techniques. This paper presents an extension of the earlier work (16, 17): a comprehensive probabilistic framework relying on the concept of the safety hierarchy. It provides a computational definition of severity as the probability of collision that is suited for

an automated system. Being able to compute the probability for any road user to collide at a given time also allows to detect traffic conflicts and estimate detailed exposure measures.

The rest of the paper proceeds as follows. The next section presents related work in traffic safety. The third section will present the probabilistic framework and the computation of the collision probability. The fourth section summarizes the developed vision-based system that provides real traffic data. In the fifth section, experimental results will illustrate the approach on this real traffic data.

RELATED WORK

The basis for this work draws heavily from the traffic conflict literature, and the concept of safety hierarchy, as shown in the introduction. Considerable work has been accomplished to validate the TCTs, involving especially the development of severity indicators, in order to obtain more objective judgments from observers. The relationship between traffic conflicts and collisions is not simple (3). TCTs try to identify the subset of the most serious traffic conflicts that are the closest to collisions. Calibration conferences (Malmö, 1983 and Trautenfels 1985) compared most of the different TCTs developed so far, and found them to mostly agree (although the definitions of the considered traffic conflicts differ). Researchers and practitioners use various combinations of severity indicators to detect traffic conflicts and estimate their severity.

Among TCTs, The Swedish TCT is one of the best-known, and is still being actively used for everyday safety assessments. It relies on the Time to Accident (TA), defined as “the time that is remaining from when the evasive action is taken until the collision would have occurred if the road users had continued with unchanged speeds and directions. Its value can be calculated based on the estimates of the distance to the potential point of collision and the speed when the evasive action is taken” (18). This speed and the TA are used to determine the border between “serious conflicts” and “non-serious conflicts”. This TA indicator is the value at a special instant of the interaction of the general TTC severity indicator, which is defined as long as a collision course exists. Another severity indicator is the PET, defined as the time measured from the moment the first road user leaves the potential collision point to the moment the other road user enters this point. Other indicators include speed and its derivatives, and subjective elements of the observers’ judgment of the chance of a collision, incorporating “speed, proximity of vehicles, apparent control of the driving task and environmental conditions such as visibility or a wet road surface” (3).

The various severity indicators provide cues to the estimation of the severity of interactions. Yet the quest for the right “border” between traffic conflicts and non-conflict interactions doesn’t appear so relevant when considered into the general framework of the safety hierarchy which places all interactions on a continuum along the severity dimension. The potential for the use of all interactions was investigated in (8, 6). They studied the shape of the distribution of interactions according to their severity, and concluded that different shapes stand for different safety situations and that there is a severity threshold, under which interactions indicate normal road users’ interactions, and above which interactions are hints of safety issues.

If one considers the whole continuum of traffic events, the concept of exposure comes into play. Exposure is typically defined as a “measure of spatial or temporal duration in the traffic system in relation to the number of dynamic system objects, road-users, vehicles (axles), etc” (19, 10). Common measures are a number of inhabitants or some amount of travel, either in distance traveled (road user-kilometers), or in time traveled (road user-hours) that takes into account the speed of road users. Exposure was introduced to make comparisons more fair

between different situations, for example annual numbers of collisions for countries with different number of inhabitants, or different car ownership levels. Dividing the number of collisions by the corresponding exposure yields a collision rate, relative to that measure of exposure, that is routinely used to compare varied situations (countries, period of observation, group of users...). However, this should be done with care since there is no reason to believe a priori that the relationship between the safety and the exposure is linear (20). More generally, any situation that is necessary for a collision to happen can be considered as exposure to collision. All interactions in the safety hierarchy are exposure to collision, more or less close according to their severity. Estimating the severity of interactions provides exposure measures that are more detailed and relevant than mere volume counts, giving more insight into the processes that lead to collisions.

There have been a few attempts at building a system for automated road safety analysis (21, 22, 23, 24, 25). To our knowledge, the system presented in (16) is the first automated system for traffic conflict detection. There is limited research in automated systems that can provide severity indicators, except for specific situations and type of traffic conflicts. Only (22) describes a system that can measure the probability of collision for any two interacting road users, but it is validated only on a few experiments with toy cars. The formulas to compute the collision probability for two interacting road users are taken from this work. This paper presents extensions for the collision probability of one road user and aggregated measures over time.

THE PROBABILISTIC FRAMEWORK

Re-thinking the Collision Course

Any situation that is necessary for a collision to happen can be considered as exposure to collision. If there is no exposure, e.g. one stays at home, or is the only road user, no collision with another road user is possible. There are obviously physical limits to possible road user movements (bounded acceleration, maximum angle of turning movement, road users' reaction times...), so that it is not possible for all road users to collide in the future. If a collision between two road users is possible, i.e. there is a possible chain of events that can lead to a collision, the probability of collision can be considered. The collision probability must take into account all the possible movements of the road users, which will return the severity, or proximity to a potential collision.

The definition of a collision course (8) takes into account only "unchanged road user movements", i.e. movements without road users' intervention. This raises many questions and is difficult to properly define. It is often used experimentally with simplified hypotheses for road users' movements. However a collision course can be better defined as an interaction in which the collision probability is non-zero at a given time. The severity of the collision course at a given instant is the collision probability, summing the probability of all chain of possible events that can lead to a collision. This implies the existence of a probability distribution over all traffic events, more precisely over road users' movements, and requires a practical way to estimate this distribution for real-world use.

The Collision Probability for Two Road Users

The formulas presented in this part are based on (22), and to a lesser extent on (23). The collision probability for a given interaction between two road users can be computed at a given instant by summing the collision probability over all possible motions that lead to a collision, given the road users' states. This requires the ability to generate for each road user at any instant

a distribution over its possible future positions given its previous positions. A possible future motion, i.e. a temporal series of predicted positions, defines an extrapolation hypothesis. The collision probability computation is approximated by a discrete sum when taking into account a finite number of the most probable extrapolation hypotheses.

First the collision probability at time t_0 for two road users A_1 and A_2 with respective observed trajectories $Q_{1,t \leq t_0}$ and $Q_{2,t \leq t_0}$ (before t_0) is defined when considering only one extrapolation hypothesis for each, respectively H_i and H_j . The predicted positions according to the hypotheses H_i and H_j are computed for a number of time steps: the predicted time of the collision $t_{i,j}$ is the first instant at which the road users would be in contact. The larger $\Delta_{i,j} = t_{i,j} - t_0$, the more likely the road users can react and avoid the collision. This time takes into account speed and distance and is directly measurable against the road users' reaction times. The formula of the probability of collision given hypotheses H_i and H_j is taken from (22)

$$P(\text{Collision}(A_1, A_2) | H_i, H_j) = e^{-\frac{\Delta_{i,j}^2}{2\sigma^2}} \quad (1)$$

where σ is a normalizing constant. It is estimated in (22) that this probability should change when the elapsed time $\Delta_{i,j}$ is close to the road user reaction time. Therefore σ is chosen to be equal to an average user reaction time¹. The number of predicted positions can be limited to 3σ , as the resulting probability is very close to zero when $\Delta_{i,j}$ reaches that value. Based on (22), the collision probability for two road users A_1 and A_2 at t_0 is

$$P(\text{Collision}(A_1, A_2) | Q_{1,t \leq t_0}, Q_{2,t \leq t_0}) = \sum_{i,j} P(H_i | Q_{1,t \leq t_0}) P(H_j | Q_{2,t \leq t_0}) e^{-\frac{\Delta_{i,j}^2}{2\sigma^2}} \quad (2)$$

where $P(H_i | Q_{1,t \leq t_0})$ is the probability of road user A_1 to move according to extrapolation hypothesis H_i (same for A_2 and H_j). The sum is done over a variety of extrapolation hypotheses, although this number must be limited to maintain reasonable computation times. This formula is illustrated in a simplified example in Figure 1. In a traditional TCT, one could choose a threshold on the collision probability and other indicators to define traffic conflicts. In the new approach described in this paper, road safety can be automatically analyzed in detail by computing continuously the collision probability of all interactions.

The Collision Probability for One Road User

When considering the collision probability for only one road user, the formulas have to be adapted. It is not possible to directly sum the collision probabilities of the interactions in which the road user is involved, as only one collision can happen for each extrapolation hypothesis. The predicted positions according to hypothesis H_i and all hypotheses that the other interacting road user may follow are computed for a number of time steps. If the road user follows the motion hypothesis H_i , the predicted time of the collision t_i is the first instant at which the road user following motion hypothesis H_i would be in contact with another road user ($\Delta_i = t_i - t_0$). Let $Q_{k(i)}$ be the observed trajectory of this road user and $H_{j(i)}$ the hypothesis that leads this road user to a collision. The collision probability of the road user A_1 with n other road users at t_0 is

¹A value of 1.5 seconds is chosen for the experiments described in this paper.

$$P(\text{Collision}(A_1) | Q_{1,t \leq t_0}, Q_{2,t \leq t_0}, \dots, Q_{n,t \leq t_0}) = \sum_{i,j} P(H_i | Q_{1,t \leq t_0}) P(H_{j(i)} | Q_{k(i),t \leq t_0}) e^{-\frac{\Delta_i^2}{2\sigma^2}} \quad (3)$$

Aggregating over Time

The previous definitions deal only with one road user or one interaction between two road users at a given instant. The collision probability for two road users in interaction can be used for the detection of traffic events relevant to safety. However, to characterize a given period of time at a location, one needs a method to accumulate the indicators over all interactions that occurred in the monitored area during this period of time, or over all road users that went through the monitored area during this period of time.

The first aggregation level is the interaction or the road user. This indicator should reflect the highest collision probability over time, but also the amount of time during which this collision probability was high. This should therefore be similar to an integral of the instantaneous collision probability over time. However, issues arise when dealing with real data, e.g. collected after automated road user tracking using video sensors: tracking errors and noise produce measures of collision probability over time which may be randomly truncated and noisy. Hu et al. report similar observations in (22). This would make it difficult to compare fairly the interactions. Consequently, to improve robustness, it is preferred to use the average of a small number of largest values taken by the collision probability over time. Let n be that number. Let $SeverityIndex(A_1, A_2)$ and $SeverityIndex(A_1)$ be the averages of the n largest values taken respectively by the collision probability $P(\text{Collision}(A_1, A_2) | Q_{1,t \leq t_0}, Q_{2,t \leq t_0})$ over the time that the two road users A_1 and A_2 interacted in the monitored area, and by the collision probability $P(\text{Collision}(A_1) | Q_{1,t \leq t_0}, Q_{2,t \leq t_0}, \dots, Q_{n,t \leq t_0})$ over the time that the road user A_1 has spent in the monitored area. The values can subsequently be summed over time for all interactions or road users. The severity indices for the time interval $[t_1 t_2]$ are

$$InteractionSeverityIndex([t_1 t_2]) = \sum_{\substack{(i,j) \text{ such that } A_i \text{ and } A_j \text{ are observed} \\ \text{in Interaction during } [t_1 t_2]}} SeverityIndex(A_i, A_j) \quad (4)$$

$$UserSeverityIndex([t_1 t_2]) = \sum_{\substack{i \text{ such that } A_i \text{ is observed} \\ \text{during } [t_1 t_2]}} SeverityIndex(A_i) \quad (5)$$

OVERVIEW OF A VISION-BASED AUTOMATED SYSTEM

This framework is used in a complete automated vision-based system for road safety analysis. Such a system requires a high level understanding of the scene and is traditionally composed of two levels of modules (see Figure 2):

1. a video processing module for road user detection and tracking,
2. interpretation modules for interaction analysis and traffic conflict detection.

For road safety applications, the approach relies on the building of two databases: a trajectory database, where the results of the video processing module are stored, and an interaction database, where all interactions between road users within a given distance are considered, and for which various indicators, including collision probability and other severity indicators, are automatically computed. Identifying traffic conflicts and measuring other traffic parameters becomes the problem of mining these databases.

The road user detection and tracking module used in the system described in this paper relies on a feature-based tracking method that extends to intersections the method described in

(11). In this approach, distinguishable points or lines in the image are tracked: a moving object may have multiple features, which must be grouped for each object. A detailed description of the tracking algorithm is presented in (17). The algorithm relies on world coordinates through the estimation of the homography matrix. The tracking accuracy for motor vehicles has been measured between 84.7% and 94.4% on three different sets of sequences (pedestrians and two-wheels may also be tracked, but less reliably). This means that most trajectories are detected by the system, although overgrouping and oversegmentation still happens and creates some problems. The most important limitation for traffic conflict detection is the inaccuracy in the estimation of road user sizes.

Trajectories provided by the first module are used in subsequent modules to extract relevant information. A first system was developed to detect directly traffic conflicts using HMM-based semi-supervised machine learning techniques (16). To apply the probabilistic safety framework presented in this paper, it is necessary to be able to predict road users' future positions. Motion patterns, represented by actual prototype trajectories without any special pre-processing, are learnt incrementally using the Longest Common Sub-sequence Similarity (LCSS) (26). The description of the motion pattern learning algorithm is beyond the scope of this paper and is described in detail in (27). The motion pattern probabilities are computed by matching all trajectories over a given period of time using LCSS, and can be updated continuously in a real-time application, as traffic patterns change in time. When computing the collision probability, the partial trajectories of each considered road user at each time are matched against the set of learnt prototypes using the LCSS.

EXPERIMENTAL RESULTS

The core architecture of the system has been implemented, using the Intel OpenCV library². On the contrary to (22) which uses toy cars, the present work is tested on real traffic video data, and a few traffic conflict instances identified by trained traffic conflict observers³.

Three sets of data are used. The first is a set of traffic sequences on the same location initially used for the training of traffic conflict observers in the 1980s. Their length ranges from 10 seconds to 60 seconds. This "Conflict" set contains 2941 feature trajectories of a minimum length of 40 frames, and 327 road user trajectories. The second dataset is composed of two long sequences, each close to one hour long, recorded at an intersection in the Twin Cities (United States), in Minnesota. This "Minnesota" set contains 88255 feature tracks of a minimum length of 40 frames, and 11734 road user trajectories. The third dataset is composed of 6 sequences, each about 20 minutes long, recorded in Reggio di Calabria (south Italy). This "Italy" dataset contains 138009 feature tracks of a minimum length of 40 frames, and 9849 road user trajectories..

First the motion patterns are learnt from the feature trajectories, which are smoothed using a Kalman filter beforehand. It is difficult to evaluate such an unsupervised task. The learnt prototypes for the datasets are presented in Figure 3. The visual examination of the motion patterns suggests a plausible division of the trajectory space. Traffic patterns are well identified, and the traffic volumes are consistent with observation.

Traffic Conflict Study

²<http://sourceforge.net/projects/opencvlibrary/>

³Additional experimental results are available at the address

<http://www.confins.net/saunier/data/saunier07trb.html>.

Since only a few traffic conflict instances are available in the Conflict dataset, only preliminary results obtained for the three detectable traffic conflict instances are reported in this paper (these three traffic conflict instances belong to three sequences of the Conflict dataset). It appears that the prototype trajectories are well suited for the computation of the collision probability. An example of movement prediction is presented for one conflict in Figure 4.

The curves of the collision probability as a function of time, computed using formula 2,, are displayed for the three traffic conflicts in Figure 5. For each of these instances, one vehicle is over-segmented, resulting in two trajectories, and thus two traffic events (and two curves). It appears that the collision probability shows an expected evolution over time, starting with low values, increasing until the probability of collision reaches a maximum, to decrease afterward, often truncated due to tracking errors and disrupted trajectories.

Over-segmentation of tracked road users can cause major problems. The same road user detected twice can entail the detection of an interaction between two very close “imaginary” road users, often with very high computed collision probability. Fortunately, these interactions are mostly filtered out by testing for the similarity between the trajectories of interacting road users using the LCSS distance. In two of the three sequences containing traffic conflicts, querying interactions for which the severity index is superior to 0.1 returns only the traffic conflicts. For the third sequence, it returns the traffic conflict and some interactions between road users in traffic moving in opposite directions. Querying the other sequences that contain no detectable traffic conflicts also return these “normal” interactions that can be easily identified. This shows that traffic conflict detection can be achieved by computing the collision probability. Adding other severity indicators will further improve the detection results.

Severity Indices

Using formula 5, the severity indices of all interactions are computed for the sequences of the Minnesota and Italy datasets, which are both more than one hour long. The distributions of the interaction according to their severity indices are represented individually for each sequence of the two datasets in the Figure 6. As expected, the distributions exhibit the shape of the safety hierarchy, with the frequency of events decreasing as the severity increases. The different sequences in each dataset exhibit different distributions. For example, more interactions for all level of severity are observed in the sequence 2 in the Minnesota dataset. This type of analysis could be performed to compare different situations, for example in before and after studies. It is also possible to study interactions by their locations, by building severity maps, and therefore analyze particular problems in the intersection.

CONCLUSIONS AND FUTURE WORK

This paper presents a comprehensive probabilistic framework for automated road safety analysis. It provides computation definitions of the probability of collisions for road users involved in an interaction, extending the work of (22) to an individual collision probability and aggregated measures over time. By integrating this framework into a complete system for vision-based road safety analysis, it is shown that these definitions are suitable for an automated system. This provides detailed severity measures, exposure estimates and a method to detect and study traffic conflicts. The system is demonstrated using real traffic data, including some traffic conflict instances, illustrating the approach and its usefulness.

New data is currently being collected to expand the results and validate the computed measures. Further research is needed to investigate and validate the relationship of collision probability to safety.

ACKNOWLEDGEMENT

The authors thank Stefan Atev from the University of Minnesota in the Twin Cities for providing us with the video sequences of the Minnesota dataset, Jean-Michel Auberlet from INRETS for providing the video sequences of the Italian intersection, Karim Ismail and Clark Lim for their helpful comments.

REFERENCES

1. Gettman D. and L. Head. Surrogate safety measures from traffic simulation models, final report. Technical Report FHWA-RD-03-050, Federal Highway Administration, 2003.
2. Brown G. R. Traffic conflict for road user safety studies. *Canadian Journal of Civil Engineering*, 21:1–15, 1994.
3. Hydén C. The development of a method for traffic safety evaluation: The Swedish Traffic Conflicts Technique. PhD thesis, Lund University of Technology, Lund, Sweden, 1987. Bulletin 70.
4. Sayed T., G. R. Brown, and F. Navin. Simulation of Traffic Conflicts at Unsignalised Intersections with TSC-Sim. *Accident Analysis & Prevention*, 26(5): 593–607, 1994.
5. Sayed T. and S. Zein. Traffic conflict standards for intersections. *Transportation Planning and Technology*, 22:309–323, 1999.
6. Svensson Å. and C. Hydén. Estimating the severity of safety related behaviour. *Accident Analysis & Prevention*, 38(2): 379–385, March 2006.
7. Amundsen F. and C. Hydén, editors. Proceedings of the first workshop on traffic conflicts, Oslo, Norway, 1977. Institute of Transport Economics.
8. Svensson Å. A Method for Analyzing the Traffic Process in a Safety Perspective. PhD thesis, University of Lund, 1998. Bulletin 166.
9. Archer J. Methods for the Assessment and Prediction of Traffic Safety at Urban Intersections and their Application in Micro-simulation Modelling. Academic thesis, Royal Institute of Technology, Stockholm, Sweden, December 2004.
10. van der Horst R. A time-based analysis of road user behavior in normal and critical encounter. PhD thesis, Delft University of Technology, 1990.
11. Beymer D., P. McLauchlan, B. Coifman, and J. Malik. A real-time computer vision system for measuring traffic parameters. In Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97), pages 495–501, Washington, DC, USA, 1997. IEEE Computer Society. ISBN 0-8186-7822-4.
12. Hu W., T. Tan, L. Wang, and S. Maybank. A survey on visual surveillance of object motion and behaviors. *Systems, Man and Cybernetics, Part C, IEEE Transactions on*, 34(3):334– 352, August 2004a.
13. Kastrinaki V., M. Zervakis, and K. Kalaitzakis. A survey of video processing techniques for traffic applications. *Image and Vision Computing*, 21(4): 359–381, April 2003.
14. Laureshyn A. and H. Ardö. Straight to video? *Traffic Technology International*, pages 107–109, 2007.
15. Maurin B., O. Masoud, and N. P. Papanikolopoulos. Tracking all traffic: computer vision algorithms for monitoring vehicles, individuals, and crowds. *Robotics & Automation Magazine, IEEE*, 12 (1):29–36, March 2005.
16. Saunier N. and T. Sayed. Automated Road Safety Analysis Using Video Data. In *Transportation Research Board Annual Meeting Compendium of Papers*, Washington, D.C., January 2007.

17. Saunier N. and T. Sayed. A feature-based tracking algorithm for vehicles in intersections. In Third Canadian Conference on Computer and Robot Vision, Québec, June 2006. IEEE.
18. Anonymous. The swedish traffic conflict technique. <http://www.tft.lth.se/rapporter/Conflict1.pdf>, 2005.
19. Hakkert A.S. and L. Braimaister. The uses of exposure and risk in road safety studies. Technical Report R-2002-12, SWOV, Leidschendam, 2002.
20. Hauer E. On exposure and accident rate. *Traffic Engineering + Control*, 36(3): 134–138, 1995.
21. Atev S., H. Arumugam, O. Masoud, R. Janardan, and N. P. Papanikolopoulos. A vision-based approach to collision prediction at traffic intersections. *IEEE Transactions on Intelligent Transportation Systems*, 6(4):416– 423, December 2005.
22. Hu W., X. Xiao, D. Xie, T. Tan, and S.J. Maybank. Traffic accident prediction using 3d model based vehicle tracking. *IEEE Transactions on Vehicular Technology*, 53 (3):677–694, May 2004b.
23. Kamijo S., Y. Matsushita, K. Ikeuchi, and M. Sakauchi. Traffic monitoring and accident detection at intersections. *IEEE Transactions on Intelligent Transportation Systems*, 1(2):108–118, June 2000.
24. Laureshyn A. and H. Ardö. Automated video analysis as a tool for analysing road safety behaviour. In ITS World Congress, London, 2006.
25. Messelodi S. and C. M. Modena. A computer vision system for traffic accident risk measurement: A case study. Technical Report ITC-irst T05-06-07, ITC, 2005.
26. Vlachos M., G. Kollios, and D. Gunopulos. Elastic translation invariant matching of trajectories. *Machine Learning*, 58(2-3):301–334, February 2005.
27. Saunier N., T. Sayed, and C. Lim. Probabilistic collision prediction for vision-based automated road safety analysis. In The 10th International IEEE Conference on Intelligent Transportation Systems, Seattle, October 2007.

LIST OF FIGURES

FIGURE 1 In this simplified situation, two vehicles approach a T intersection at time t_0 . Only two extrapolation hypotheses are considered for each vehicle. Vehicle 1 is expected to turn left or right, with respective probabilities 0.4 and 0.6. Vehicle 2 is expected to go straight or turn left, with respective probabilities 0.7 and 0.3. There are two potential collision points, that can happen at times t_1 and t_2 . The collision probability at time t_0 is

computed as $P(\text{Collision}) = 0.4 \times 0.7 \times e^{-\frac{(t_1-t_0)^2}{2\sigma^2}} + 0.4 \times 0.3 \times e^{-\frac{(t_2-t_0)^2}{2\sigma^2}}$.

FIGURE 2 Overview of a modular system for vision-based automated road safety analysis.

FIGURE 3 Motion patterns learnt respectively on sequences of the Conflict dataset (top), the Minnesota dataset (bottom left) and the Italy dataset (bottom right), resulting respectively in 58, 128 and 58 prototype trajectories. The tracks are displayed in color, from white to red indicating the number of matched trajectories in the sequence for each pattern, i.e. the traffic volume along these patterns.

FIGURE 4 An example of movement prediction in a real traffic conflict situation (Sequence 1, See top plot in Figure 5). The vehicle trajectories are red and blue, with a dot marking their position, and the future positions are respectively cyan and yellow.

FIGURE 5 Graphs of the collision probability for the three traffic conflicts (collected in three separate sequences), as a function of time (counted in frame numbers). In all sequences, vehicle 1 travels south-bound through the intersection and vehicle 2 comes from an opposing approach. Vehicle 2 turns left in sequence 1 (top) (See Figure 4), right in sequence 2 (middle) and stops in sequence 3 (bottom).

FIGURE 6 Distribution of the interactions according to their severity indices (with a zoom on the higher severities), quantified by 0.1 (the point at severity index x stands for the number of interactions with severity index between $x-0.1$ and x), for the sequences of the Minnesota dataset (top) and the Italy dataset (bottom).




FIGURE 1 In this simplified situation, two vehicles approach a T intersection at time t_0 . Only two extrapolation hypotheses are considered for each vehicle. Vehicle 1 is expected to turn left or right, with respective probabilities 0.4 and 0.6. Vehicle 2 is expected to go straight or turn left, with respective probabilities 0.7 and 0.3. There are two potential collision points, that can happen at times t_1 and t_2 . The collision probability at time t_0 is computed as $P(\text{Collision}) = 0.4 \times 0.7 \times e^{-\frac{(t_1 - t_0)^2}{2\sigma^2}} + 0.4 \times 0.3 \times e^{-\frac{(t_2 - t_0)^2}{2\sigma^2}}$.

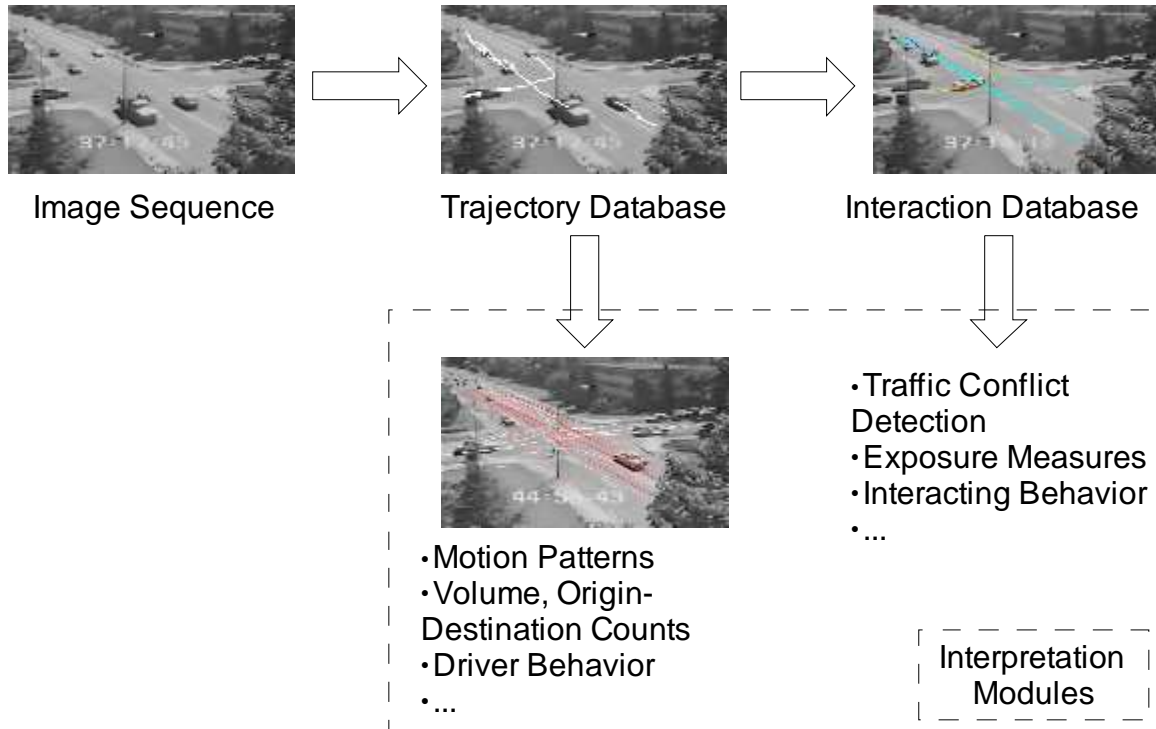


FIGURE 2 Overview of a modular system for vision-based automated road safety analysis.




FIGURE 3 Motion patterns learnt respectively on sequences of the Conflict dataset (top), the Minnesota dataset (bottom left) and the Italy dataset (bottom right), resulting respectively in 58, 128 and 58 prototype trajectories. The tracks are displayed in color, from white to red indicating the number of matched trajectories in the sequence for each pattern, i.e. the traffic volume along these patterns.




FIGURE 4 An example of movement prediction in a real traffic conflict situation (Sequence 1, See top plot in Figure 5). The vehicle trajectories are red and blue, with a dot marking their position, and the future positions are respectively cyan and yellow.




FIGURE 5 Graphs of the collision probability for the three traffic conflicts (collected in three separate sequences), as a function of time (counted in frame numbers). In all sequences, vehicle 1 travels south-bound through the intersection and vehicle 2 comes from an opposing approach. Vehicle 2 turns left in sequence 1 (top) (See Figure 4), right in sequence 2 (middle) and stops in sequence 3 (bottom).


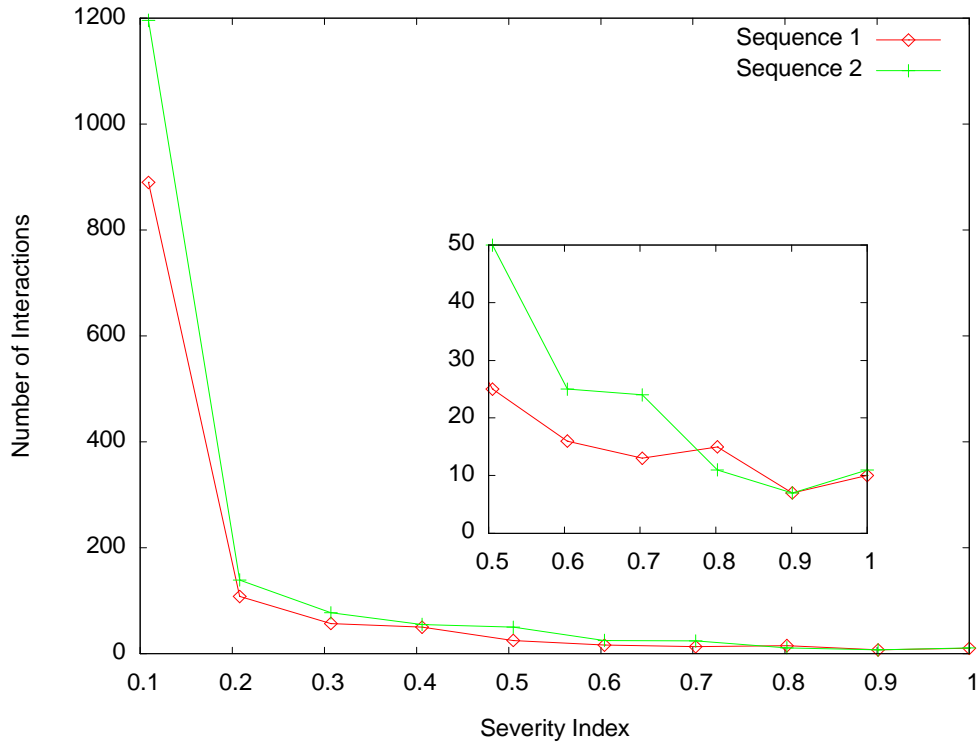


FIGURE 6 Distribution of the interactions according to their severity indices (with a zoom on the higher severities), quantified by 0.1 (the point at severity index x stands for the number of interactions with severity index between $x-0.1$ and x), for the sequences of the Minnesota dataset (top) and the Italy dataset (bottom).