Automated Road Safety Analysis Using Video Data

Paper Submission to the

2007 TRB Annual Meeting

Nicolas Saunier

Post Doctoral Fellow

Department of Civil Engineering

The University of British Columbia

6250 Applied Science Lane

Vancouver, BC, Canada V6T 1Z4

Tarek Sayed

Professor and Distinguished University Scholar

Department of Civil Engineering

The University of British Columbia

6250 Applied Science Lane

Vancouver, BC, Canada V6T 1Z4

Word Count: (5741 Text) +(1 Tables = 250)+ (5 Figure = 1250) = 7241

Automated Road Safety Analysis Using Video Data

Nicolas Saunier, Tarek Sayed

Abstract: Traffic safety analysis has often been undertaken using historical collision data. However, there are well-recognized availability and quality problems associated with collision data. In addition, the use of collision records for safety analysis is reactive: a significant number of collisions has to be recorded before action is taken. Therefore, the observation of traffic conflicts has been advocated as a complementary approach to analyze traffic safety. However, incomplete conceptualization and the cost of training observers and collecting conflict data have been factors inhibiting extensive application of the technique. The goal of this research is to develop a method for automated road safety analysis using video sensors in order to address the problem of a dependency on the deteriorating collision data. The method will automate the extraction of traffic conflicts from video sensors data. This should address the main shortcomings of the traffic conflict technique. This paper describes a comprehensive system for traffic conflict detection in video data. The system is composed of a feature-based vehicle tracking algorithm adapted for intersections and a traffic conflict detection method based on the clustering of vehicle trajectories. The clustering uses a K-means approach with Hidden Markov Models and a simple heuristic to find the number of clusters automatically. Traffic conflicts can then be detected by identifying and adapting pairs of models of conflicting trajectories. The technique is demonstrated on real world video sequences of traffic conflicts.

INTRODUCTION

Given the magnitude of the traffic safety problem, road authorities around the world are working hard to improve road safety in an effort to reduce the economic and societal costs associated with traffic collisions. Traffic safety analysis has been traditionally undertaken using historical collision data. However, there are well-recognized availability and quality problems associated with collision data. In many jurisdictions, the quantity and quality of collision data has been degrading for several years. In addition, 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. Because of these problems, 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 (*1,2,3,4*). The Traffic Conflict Technique involves observing and evaluating the frequency and severity of traffic conflicts at an intersection by a team of trained observers. Traffic conflicts are more frequent than collisions, and their study can give detailed information about safety. The technique therefore provides a means for the analysts to immediately observe and evaluate unsafe driving maneuvers at an intersection. However, incomplete conceptualization and the cost of training observers and collecting conflict data have been factors inhibiting extensive application of the technique. Therefore, the successful automation of extracting conflicts from video sensors data using computer vision techniques appears to have practical benefits for traffic safety analysis.

This research aims at implementing a complete system to interpret vehicle interactions and detect traffic conflicts in real world video data (provided by one stationary camera). This system should be generic, robust and low-cost, for regular use by traffic safety practitioners. The emphasis is on intersections which are crucial parts of the road networks for safety. This paper presents an approach to build such a traffic conflict detection system, and demonstrates its feasibility using experimental data.

PREVIOUS WORK

Traffic Conflicts

The concept of traffic conflicts was first proposed by Perkins and Harris (*5*). As an alternative to collision data, which in many cases are scarce, unreliable, or unsatisfactory. Their objective was to define traffic events or incidences that occur frequently, can be clearly observed, and are related to collisions. An internationally accepted definition of a traffic conflict is "an observable situation in which two or more road users approach each other in space and time for such an extent that there is a risk of collision if their movements remain unchanged." (*6*). A traffic conflict between two road users includes two components: a collision course and an emergency evasive action. Deciding if two road users are on a collision course depends on the extrapolation hypotheses. Common hypotheses are extrapolation with constant direction and speed. A traffic conflict is an interaction, defined as an observational situation in which two or more vehicles are close enough in space and time.

The relationship between traffic conflicts and collisions need to be established to use traffic conflicts as surrogates to collisions for safety analysis. Many researchers (*7,8*), 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 are often called quasi-collisions. The interactions can thus be recursively ranked in the safety hierarchy. 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 location in the hierarchy, i.e. measuring the proximity to the potential occurrence of a collision, which is related to the probability of collision. The main elements that constitute the severity are the distance in space and time between the drivers involved and their evasive action(s). Many severity indicators were developed (*8*,*9*), e.g. the Time-To-Collision (TTC), defined for two road users on a collision course as the extrapolated time for the collision to occur. The similarity between collisions and quasi-collisions is "striking" enough (*9*) to allow their use for safety diagnosis (*3*). Even if the relationship between traffic conflicts and collisions is not simple (*1*,*3*), traffic conflict techniques (TCTs) offer a powerful framework to investigate the safety of a location and provide detailed safety information, in the same way as costly in-depth collision analysis, since the processes are very similar.

Vision-Based Traffic Conflict Detection

The vision-based monitoring of road networks is a complex task. Monitoring intersections faces more problems than highways. These problems are related to the highly variable structure of the intersections, the presence of multiple flows of the vehicles with turning movements, the mixed traffic that ranges from pedestrians to trucks and vehicles that stop at traffic lights. Specific classification and occlusion management techniques are required. Common problems to highways and intersections include global illumination variations, multiple object tracking and shadow handling.

Despite the potential benefits of automated traffic safety analysis based on video sensors, limited computer vision research is directly applied to road safety, and especially the detection of traffic conflicts (*10*,*11*,*12*,*13*). Maurin et al. (*14*) state that "despite significant advances in traffic sensors and algorithms, modern monitoring systems cannot effectively handle busy intersections". Such a system requires a high level understanding of the scene and is traditionally composed of two modules (see Figure 1):

1. a video processing module for vehicle detection and tracking,
2. an interpretation module for traffic conflicts detection.

The input to the system is a stream of images. The detail level of the traffic conflict detection can vary. Detecting traffic conflicts, with details about the involved vehicles and their trajectories, is a difficult task and may not be achieved with enough accuracy for practical use. An alternative system can yield the sequences that contain traffic conflicts. Such a system is successful if all traffic conflicts are detected and the number of false detection is kept to the minimum. This system would act as a filter of video data that would summarize hours of video data in minutes containing the relevant incidents for further human review and analysis.

Vehicle Tracking

The goal of the tracking process is to maintain the identity of a moving object over the sequence of frames. Object tracking methods are usually classified into at least four groups (*12*,*15*). *Model-based* tracking exploits the a priori knowledge of typical objects in a given scene, which is matched to the image data. *Region/blob-based* tracking identifies connected regions of the image, blobs, associated with each vehicle. Regions are often obtained through background subtraction and then tracked over time. *Contour-based* tracking has a representation of the contour of a moving object that is updated dynamically. These first three methods have limitations. Model-based tracking requires models for all types of moving objects, and is computationally intensive. Region- and contour-based tracking don't cope well with congested traffic conditions involving partial occlusion. Many implementations of these methods also require a special initialization phase, sometimes manual.

Feature-based tracking abandons the idea of tracking objects as a whole, but instead tracks features such as distinguishable points or lines on the object. This is achieved through well known methods such as the Kanade-Lucas-Tomasi Feature Tracker. Since a vehicle can have multiple features, it must be determined what set of features belong to the same vehicle. The features are grouped using spatial proximity and common motion. Feature-based tracking algorithms have distinct advantages over other methods: they are robust to partial occlusions, they don't require any initialization, and can adapt successfully and rapidly to variable lighting conditions, allowing real-time processing and tracking of multiple objects in dense traffic. However, feature-based tracking may have difficulty in delineating occluding objects with similar motion vectors

Traffic Conflict Detection

Traffic conflict detection based on video sensors can be achieved with various methods. One would first try the direct approach of extrapolating vehicle trajectories. The approach described by Atev et al. (*13*) detects pairs of vehicles that would collide if they maintain their current speeds and directions. It employs a region-based tracking method with background subtraction, and a novel method for three-dimensional vehicle size estimation. Tracking and size estimation appear to be accurate, but there is no validation of the collision prediction results. However, this type of direct approach is likely to have unreliable performance because the tracking data are imperfect and noisy. Such explicit systems, using rules adjusted by trial and error, rarely provide generic and robust solutions, especially when the environment and the event patterns change. This would for example mean to redevelop the system for every intersection. A better model may be learnt from observation.

There are two approaches involving learning for traffic conflict detection: 1) supervised learning for interaction classification, and 2) unsupervised learning of vehicle dynamics and movements for prediction. The first approach casts traffic conflict detection as a binary classification problem between traffic conflicts and non-conflict interactions. A classifier can be learnt on instances of each class. Kamijo et al. (*11*) learn "ordinary" traffic vehicle interactions with Hidden Markov Models (HMM). An HMM is learnt for each type of activity that needs to be recognized. The results are very limited as there are very few training instances and the system is tested only on this training set.

The second approach aims at improving the movement extrapolation of the simple extrapolation approach by learning the recurrent vehicle movements. Hu et al. (*10*) use Self-Organizing Maps (SOM) to quantify the vehicle observations and predict the vehicle movements. They present a method that computes the probability for two vehicles to collide, which is then applied to traffic accident prediction. While this work seems promising, the traffic accident prediction system is validated only on two accident instances in a toy car experience. Robotics is another field that deals with movement learning and prediction, for motion planning. For example, Vasquez and Fraichard (*16*) use the statistical learning of motion patterns for the motion prediction of moving objects.

Messelodi and Modena (*12*) describe an alternative approach based on occupation rates to evaluate the average probability of accident for a passing vehicle. However, there is no direct detection of actual events. The collision probability is computed even for isolated vehicles without any other interacting vehicle in the intersection.

THE PROPOSED APPROACH

Tracking Vehicles in Intersections

For all the advantages stated in the previous section, the vehicle detection and tracking approach used in the system presented in this paper relies on a feature-based tracking method that extends the method described by Beymer et al.(*15*) to intersections.. Instead of grouping features independently in each frame, Beymer et al.(15) group the whole features tracks with common motion, by integrating this information over as many image frames as possible. However, Beymer et al.(*15*) deal only with highways, whereas trajectories in intersections are more variable, feature tracks are often disrupted by occlusions and pose change. In the system used in this paper, no assumption is made with respect to entrance or exit regions in the scene, and the algorithm is modified in order to make use of partial feature tracks.

The algorithm relies on world coordinates through the estimation of the homography matrix. The tracking accuracy has been measured between 84.7 % and 94.4 % on three different sets of sequences (see Figure 2 for examples of difficult trajectory reconstitution). Despite the problems in an open environment such as an intersection, the vast majority of vehicles are detected and their movement can be estimated with sufficient accuracy for the desired application, the detection of traffic conflicts. A detailed description of the tracking algorithm is presented in Saunier and Sayed (*17*).

Trajectory Clustering

The unsupervised learning approach is favored because it is difficult to gather sufficient data for training and testing the system. Traffic conflicts, although more frequent than collisions, are rare traffic events. On the other hand, vehicle movements can be learnt in an unsupervised way.

There are three categories of strategy for the clustering of sequential data regardless whether the sequence elements have numerical or nominal values (*18*):

1. In sequence similarity, the comparison between two sequences is viewed as a process to transform a given sequence into another. Many approaches with practical use follow that principle including the Dynamic Time Warping algorithm (DTW) widely used in speech recognition tasks and the Longest Common Subsequence model (LCSS) applied to the clustering of motion trajectories (*19*).
2. In indirect sequence clustering, a set of features is extracted from the sequences. In this space, traditional vector space-based clustering algorithms can be used. However, the process causes loss of information and needs additional knowledge. For example, Khalid and Naftel (*20*) use the leading Fourier coefficients, which they describe as being inherently unsuited for the representation of highly complex trajectories.
3. In statistical sequence clustering, statistical models like HMMs and more general Dynamic Bayesian Networks (DBNs) are built to describe the dynamics of each group of sequences. Observations are defined to be similar in terms of common similarity to a model, expressed through the likelihood function $P(Observation|Model)$.

Clustering trajectories reconstituted from video data is difficult as the trajectories are complex, noisy and often incomplete. Indirect sequence clustering is inadequate. Statistical sequence clustering is naturally suited to handle variable length trajectories. Among the sequence similarity approaches, methods based on the Euclidean distance are too simple to accommodate noisy and partial trajectories, although it is used by Vasquez and Fraichard (*16*) for a special indoor tracking application where all trajectories start from the same area.

Some studies have already explored statistical sequence clustering, mainly using HMMs. Most adopt a K-means formulation which is extended with soft memberships and applied to video data by Alon et al. (*21*). In HMM-based clustering, a set of HMMs, one HMM per cluster, is learnt iteratively. At each iteration, instances are assigned to the HMM that maximizes its probability or likelihood, and a HMM is trained on each cluster (see Figure 3). A distinct advantage is the availability of a HMM per cluster, that gives a simple and fast way to match a new trajectory in the detection phase.

The Proposed HMM-based Clustering Algorithm

The approach used in this paper, adopts the K-means formulation, adding a simple heuristics to determine the number of clusters. The HMM notation is first defined. For a detailed overview of HMMs, readers are directed to (Rabiner 1989). A complete specification of a first-order HMM with M states $[S_1, S_2, \ldots, S_M]$ and a simple Gaussian observation density is formally given by:

- a set of prior probabilities $\Pi = \{\pi_i\}$ where $\pi_i = P(q_1 = S_i),\ 1 \le i \le M$.
- a set of state transition probabilities $H = \{h_{i,j}\}$ where $h_{i,j} = P(q_{t+1} = S_j | q_t = S_i), 1 \le i, j \le M$.
- a set of output distributions $B = \{b_j\}$ where $b_j(o_t) = P(O_t = o_t | q_t = S_j) = N(O_t, \mu_j, \Sigma_j), 1 \le j \le M$, where μ_j and Σ_j are the mean and covariance of the Gaussian of state S_j .

where q_t and O_t are respectively the state and observation at time *t*. It is common to denote a mixture of *K* HMMs by $\lambda_m = (H_m, B_m, \Pi_m), 1 \le m \le K$. Algorithms exist to:

- compute the probability of observing a sequence, given a model.
- find the state sequence that maximizes the probability of the given sequence, when the model is known (the Viterbi algorithm).
- induce the HMM that maximizes (locally) the probability of the given sequence (the Baum-Welch algorithm, an expectation-maximization (EM) algorithm).

Algorithm 1: Algorithm for the clustering of variable length sequences with HMMs.

Input: Let $\lambda_m = (H_m, B_m, \Pi_m), 1 \le m \le K$ be a mixture of K HMMs

Let K_i be the initial number of clusters

Let $O = \{O^i\}, 1 \le i \le N$ be a set of N observations, where O^i is a sequence of T_i vector observations

Let $N_{\min}$ be the minimum number of observations assigned to a cluster to keep and retrain it

Let n be the maximum number of iterations without a change in the number of clusters

Output: A mixture of K_f HMMs.

begin

Initialize randomly $K = K_i$ HMMs λ_m, with a simple Gaussian observation density adjusted to the observation set O

$i = 1$

while $i \le n$ **do**

Assign each observation O^i to the HMMs λ_{m_i} that maximizes $P(O^i | \lambda_m)$

Let O_{λ_m} be the set of N_m observations assigned to λ_m

Keep only the HMMs λ_m such that $N_m \ge N_{\min}$, reset and train λ_m with O_{λ_m}

if No HMM was discarded **then**

$i = 1 + 1$

end

The HMMs are initialized in a standard way: uniform priors and transition matrices are used, the means and variances of the Gaussian observation density are estimated using K-means on the observation data. The number of clusters K is automatically adjusted through the use of a minimum number of assigned observations per cluster, in order to ensure a minimum reliability of the parameters of the HMM. The HMMs that do not satisfy this condition are discarded. The initial number K_i of HMMs should be larger than the "natural" number of clusters of the data. The learning process stops when the number of clusters doesn't change for a number n of iterations. A major drawback of HMM-based clustering is the instability of the resulting clustering, as the initialization of the HMMs is random.

Detecting Traffic Conflicts

Vehicles involved in traffic conflicts are on a collision course, i.e. vehicles will collide if their movements continue unchanged. To decide if two vehicles are on a collision course, their positions, speeds and movement directions are required. That is why velocity (speed and direction vector) estimations are included in the trajectories descriptions, whereas most other research uses the actual vehicle displacements. The clustered observations are therefore sequences of four-dimensions vectors, composed of the vehicle coordinates (x, y) and velocity (s_x, s_y) at each time step. The estimated size of the vehicles should be also useful, but it didn't improve the results in the experiments. Since the trajectories obtained through the vehicle tracking algorithm are noisy, they are smoothed using a moving average filter (*21*).

Even in the unsupervised approach, traffic conflict detection requires some supervision. Knowing typical vehicle movements, a new trajectory can be matched against them, and the best matches used for extrapolation. One approach is to combine the other probable simultaneous vehicle movements, and to estimate a collision probability. As experimented by Hu et al. (*10*), this measure is expected to reach values for traffic conflicts that are significantly different from normal interactions.

This paper presents a more simple approach introduced by Saunier and Sayed (*22*) that identifies the conjunction of movements that may lead to collisions. For that purpose, HMMs learnt on each obtained clusters are adapted into more specific models of conflicting trajectories by using a few traffic conflict instances. The resulting "conflicting clusters" are better models of conflicting trajectories than what could be learnt with the available training instances. In speech recognition, it is a traditional method to adapt general HMM to each speaker. This was used recently for unusual event detection by Zhang et al. (*23*). Readers are directed to Saunier and Sayed (*22*) for the HMM adaptation formulas. A weighting factor α controls the balance between the original model and the new estimates on the trajectories involved in the training traffic conflicts. Both the original model and the adapted model are kept in the set of HMMs used for detection.

When adapting the HMMs, the learning includes the memorization of the pairs of models to which the pairs of trajectories involved in the training traffic conflict instances were assigned. Consequently, the learning phase, before any detection, yields a set of HMMs and pairs of conflicting models. A traffic conflict is detected when two simultaneous trajectories match one of the memorized pairs of conflicting HMMs. The detection process proceeds as follows:

1. Vehicles are tracked.
2. If two vehicles are close enough (threshold on their distance) and nearing each other (their distance decreases), they are considered to be in interaction.
3. Each interacting vehicle trajectory is assigned to a HMM, say *A* and *B*.
4. If the two HMMs *A* and *B* of both interacting trajectories were both memorized as conflicting, a traffic conflict between these two vehicles is detected.

EXPERIMENTAL RESULTS

The evaluation is based on a set of ten traffic sequences on the same location (figure 4), initially used for the training of traffic conflict observers in the 1980s. Their length ranges from 10 seconds to 60 seconds. Despite the videotape aging and the approximate alignment of the field of view between sequences, this data was readily available and could be digitized to test the method. The tracking results on these sequences are reported by Saunier and Sayed (17) to be 86 %.

All *K* HMMs of a mixture have the same structure parameter values (number of states, simple Gaussian observation density), which are tuned for each cluster of observations with the algorithm. The trajectories are measured in the world coordinates and are smoothed with an average moving filter with support 5. HMMs are clustered with a set of 560 trajectories from eight traffic sequences. The trajectories of two sequences are not used for learning since there are many detection and tracking errors caused by camera jitter. They are however used to test the generalization ability of the method. The set is large and representative of all trajectories, but small enough to maintain reasonable computation times. The parameters *n* (number of iterations once the number of HMMs doesn't change anymore) and $N_{\min}$ (minimum size of the clusters) are tuned by trial and error respectively to 3 and 5.

Figure 5 shows that for different settings of the number of states M for all HMMs, the final number K_f of HMMs increases as a function of the initial number K_i. When K_i is superior to a certain value, K_f increases very slowly, and this value can be seen as the "natural" number of clusters to describe the set of trajectories. A value of M=3 states seems a good tradeoff for accurate and fast learning. An observation of the learnt HMMs indicates that the 3 states (the means of the Gaussian observation density) often correspond to the approach of the intersection, the conflict zone and the leaving of the intersection. Besides, the stopping criterion of the algorithm is quickly reached, most discarding of HMMs occurring at the first iteration of the algorithm, followed by n more iterations.

An example of the result of the clustering algorithm is presented in the Figure 6. The figures display the trajectories assigned to each cluster. The main vehicle streams are represented sometimes mixed in the same cluster. Most adapted HMMs have fewer assigned trajectories than non-adapted models. However, 66 trajectories are assigned to the largest cluster. The examination of these clusters suggests a plausible division of the trajectories space, covering different paths with differing velocities.

The training program highlights nine traffic conflicts in nine sequences (there is no traffic conflict in the last sequence). Among these, only five traffic conflicts are used as training instances. Another one (in the fourth sequence) is not used because the vehicles are not detected anymore at the time they are really in conflict, and the rest of the traffic conflicts involves pedestrians and cyclists that are very difficult to distinguish on account of the quality of the videotape. Result validation is difficult as there is no ground truth, apart from the traffic conflicts used for training, and this relies only on observers' judgment, not on objective measures such as the Time To Collision (TTC).

The detection results are presented in terms of interactions. They are computed on all the available data, the trajectories detected in the ten sequences. This is a two-class classification problem, traffic conflicts against non-conflict interactions, with two types of errors: false alarms (FA), when the system accepts a non-conflict interaction, and false rejections (FR), when the system rejects a traffic conflict. Designed as the system is, the traffic conflicts used to identify and adapt the HMMs will always be detected. It is sometimes difficult to decide upon the other detected traffic conflicts: some are clearly FA, but some could be real traffic conflicts, and are counted in their own category. When these uncertain traffic conflicts are not detected, no FR is counted. This allows a more flexible result evaluation.

1501 interactions are detected. The numbers are quite high because the tracking is sometimes lost, which increases the number of trajectories measured, and therefore the number of interactions. This is also why the five training traffic conflicts instances are detected as ten interactions. The tradeoff between adaptation and generalization is illustrated by studying the performance of a mixture of HMMs, with different levels of adaptation, controlled by the weighting factor α. In Table 1, the detection results are presented for the mixture of HMMs for which the clustered trajectories are displayed in Figure 6. The most striking is the improvement of results between the non-adapted HMMs and the smallest level of adaptation (α=0.05). It must be noted that most FA are caused by detection and tracking errors in one of the test traffic sequences, and not by the traffic conflict detection method. The system can generalize as there are no FA for all levels of adaptation on the other test sequence without traffic conflicts.

There seems to be a good tradeoff between adaptation and generalization for α=0.15 or 0.2. As the adaptation is more important, the system becomes very specific to the training traffic conflicts instances, and is not able anymore to generalize to unknown data, as it is showed by the overall decrease of the number of uncertain traffic conflicts. For example, the traffic conflict of sequence 4 is detected for the non-adapted HMMs and for $\alpha \leq 0.15$, but not anymore when the adaptation increases. Such overfitting should be avoided.

It should be noted that the limited data available for testing make this evaluation difficult, but these results demonstrate the possibility of detecting traffic conflicts among all interactions with the presented method, and even generalize to some unseen data. The system can be made more or less specific according to the application.

CONCLUSIONS AND FURTHER RESEARCH

This paper shows the feasibility of automated traffic conflict detection based on video sensors. A complete system is presented, including an improved feature-based tracking method and a first traffic conflict detection approach using the clustering of trajectories with HMMs in an unsupervised way, their adaptation and identification for further detection. First results on limited real data show the potential of this work. Further work includes the collection of more data for testing. The use of simulated traffic data is contemplated to evaluate on a varied set of controlled situations how generic the technique is. Improvements of the technique can be made at different levels. K-means clustering is only guaranteed to converge to a local optimum, and the results are somewhat dependent on the initialization of the algorithm. Other clustering techniques based on the sequence similarity approach (LCSS), possibly in combination with HMM-based clustering, can make the process more robust. Collision probability estimation methods similar to (Hu et al. 2004) will also be tried as it can provide severity measures such as the TTC and avoids casting traffic conflict detection in a strict classification framework, where finding a boundary between traffic conflicts and non-conflict interactions can only be arbitrary

Currently, a detailed validation plan is being developed to test the method presented in this paper. The first component of the validation will focus on comparing the conflicts identified by this method and human experts' identification and ranking of conflicts. As well, the validation will test the ability of using the extracted conflicts to produce statistically-significant differences in the safety for different intersection designs. The second component will compare the predictive safety performance capabilities of using the extracted conflicts with actual collision experience at signalized intersections.

REFERENCES

1. Brown, G. R. (1994). Traffic conflict for road user safety studies. Canadian Journal of Civil Engineering, 21:1–15.
2. Sayed, T., Brown, G. R., and Navin, F. (1994). Simulation of Traffic Conflicts at Unsignalised Intersections with TSC-Sim. Accident Analysis & Prevention, 26(5):593–607.
3. Sayed, T. and Zein, S. (1999). Traffic conflict standards for intersections. Transportation Planning and Technology, 22:309–323.
4. Svensson, A. and Hydén, C. (2006). Estimating the severity of safety related behaviour. Accident Analysis & Prevention, 38(2):379–385.

5. Perkins, S., and Harris, J. (1967). “Criteria for Traffic Conflict Characteristics”, Report GMR 632. General Motors Corporation.
6. Amundson, F., and Hyden, C. (1977). Proceedings of First Workshop on traffic Conflicts, Institute of Economics, Oslo.
7. Hyden, C. (1987). “The development of a method for traffic safety evaluation: The Swedish Traffic Conflicts Technique” PhD thesis, Lund University of Technology, Lund, Sweden, Bulletin 70.
8. Svensson, A. (1998). “*A Method for Analyzing the Traffic Process in a Safety Perspective*”, PhD thesis, University of Lund, 1998. Bulletin 166.
9. van der Horst, R. (1990), “A time-based analysis of road user behavior in normal and critical encounter”, PhD thesis, Delft University of Technology.
10. Hu, W., Xiao, X., Xie, D., Tan, T., and Maybank, S. (2004). Traffic accident prediction using 3d model based vehicle tracking. IEEE Transactions on Vehicular Technology, 53(3):677–694.
11. Kamijo, S., Matsushita, Y., Ikeuchi, K., and Sakauchi, M. (2000). Traffic monitoring and accident detection at intersections. IEEE Transactions on Intelligent Transportation Systems, 1(2):108–118.
12. Messelodi, S. and Modena, C. M. (2005). A computer vision system for traffic accident risk measurement: A case study. Technical Report ITC-irst T05-06-07, ITC.
13. Atev, S., Arumugam, H., Masoud, O., Janardan, R., and Papanikolopoulos, N. P. (2005). A vision-based approach to collision prediction at traffic intersections. IEEE Transactions on Intelligent Transportation Systems, 6(4):416– 423.
14. Maurin, B., Masoud, O., and Papanikolopoulos, N. P. (2005). Tracking all traffic: computer vision algorithms for monitoring vehicles, individuals, and crowds. Robotics & Automation Magazine, IEEE, 12(1):29–36.
15. Beymer, D., McLauchlan, P., Coifman, B., and Malik, J. (1997). 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. IEEE Computer Society.
16. Vasquez, A. D. and Fraichard, T. (2004). Motion prediction for moving objects: a statistical approach. In Proc. of the IEEE Int. Conf. on Robotics and Automation, pages 3931–3936, New Orleans, LA (US).
17. Saunier, N. and Sayed, T. (2006) “A Feature-based Tracking Algorithm for Vehicles in Intersections”, Third Canadian Conference on Computer and Robot Vision, IEEE, Quebec City, June 7-9.
18. Xu, R. and Wunsch, D. (2005). Survey of clustering algorithms. Neural Networks, IEEE Transactions on, 16(3):645–678.
19. Buzan, D., Sclaroff, S., and Kollios, G. (2004). Extraction and clustering of motion trajectories in video. In Proceedings of the 17th International Conference on Pattern Recognition, volume 2, pages 521–524.
20. Khalid, S. and Naftel, A. (2005). Classifying spatiotemporal object trajectories using unsupervised learning of basis function coefficients. In VSSN’05: Proceedings of the third ACM international workshop on Video surveillance & sensor networks, pages 45–52, New York, NY, USA. ACM Press.
21. Alon, J., Sclaroff, S., Kollios, G., and Pavlovic, V. (2003). Discovering clusters in motion time-series data. In IEEE Conference on Computer Vision and Pattern Recognition.

22. Saunier, N. and Sayed, T. (2006) “Clustering Vehicle Trajectories with Hidden Markov Models. Application to Automated Traffic Safety Analysis”, IEEE World Congress on Computational Intelligence, Vancouver, July 16-21.
23. Zhang, D., Gatica-Perez, D., Bengio, S., and McCowan, I. (2005). Semi-supervised adapted hmms for unusual event detection. In IEEE Conference on Computer Vision and Pattern Recognition, San Diego.

LIST OF TABLES

LIST OF FIGURES

TABLE 1 Detection results as a function of the weighting factor α, which controls the adaptation of the HMMs to the traffic conflict instances.

	Number of Traffic Conflicts		
α	Correctly Detected	Uncertain	False Alarms
No adaptation	10	17	38
0.05	10	13	6
0.10	10	13	10
0.15	10	12	6
0.20	10	3	3
0.25	10	5	2
0.30	10	5	2
0.35	10	4	1
0.40	10	4	0
0.45	10	4	0
0.50	10	3	0

FIGURE 1 Generic overview of a traffic conflict detection system, composed of two modules.

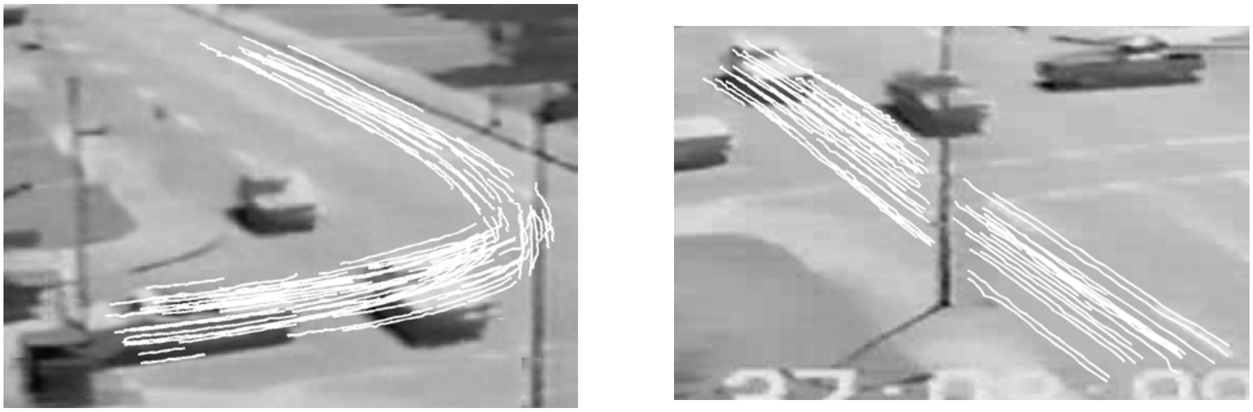

FIGURE 2 Example of disrupted feature tracks that are still successfully grouped.

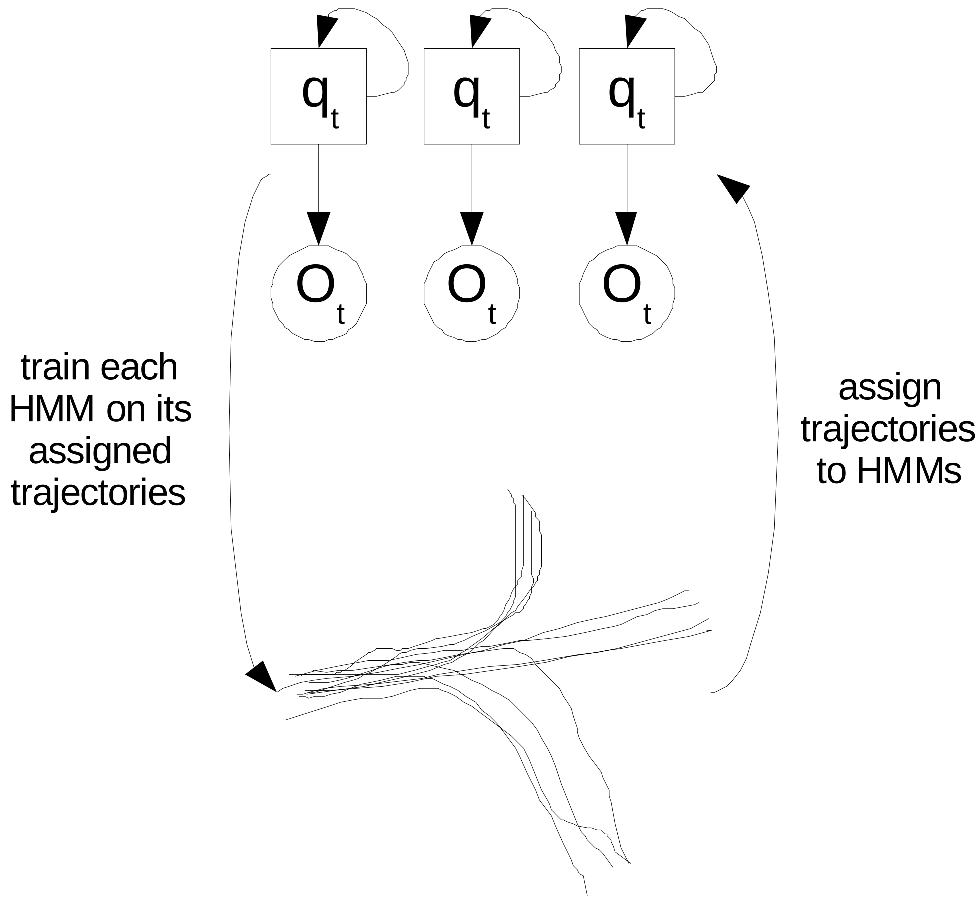

FIGURE 3 HMM-based clustering.

FIGURE 4 An image of the traffic sequences used for experiments.

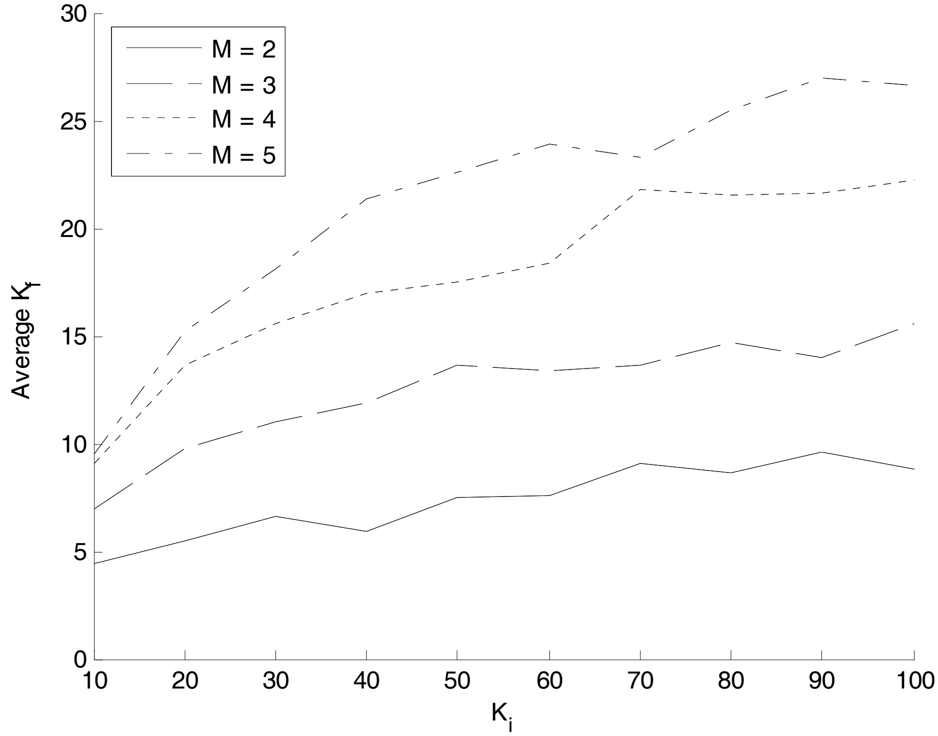

FIGURE 5 The average final number K_f of HMMs, over 10 runs of the algorithm, as a function of the initial number K_i, for different number of states M.

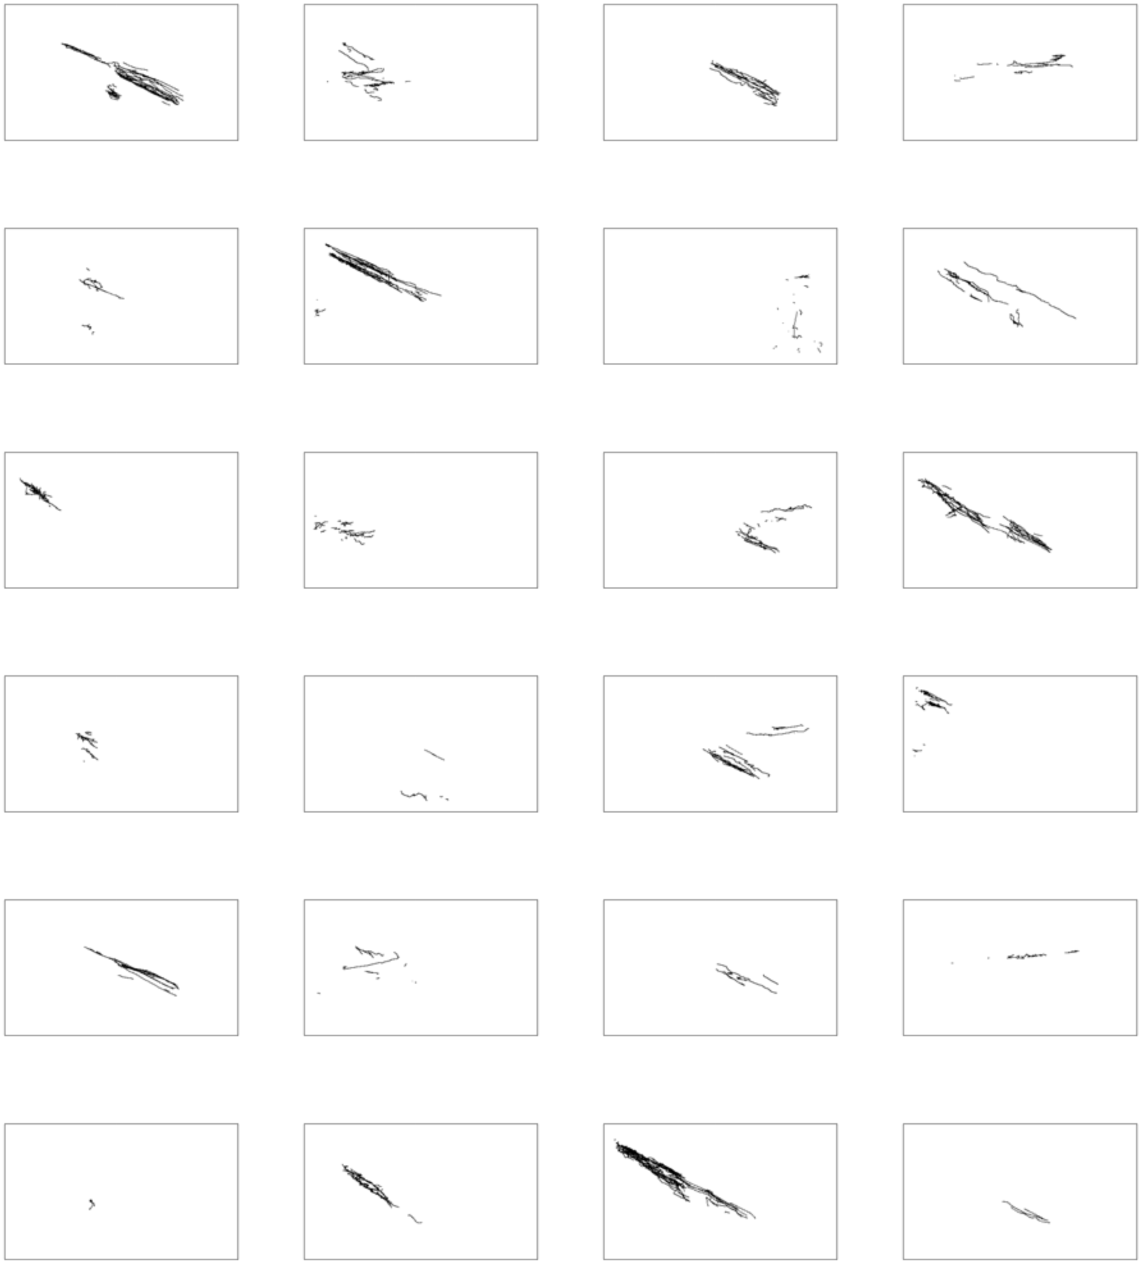

FIGURE 6 Each image represents the trajectories assigned to each cluster (the trajectories assigned to adapted HMMs are displayed in the last two rows).