Learning Spatio-temporal Context for Vehicle Reidentification

Ahmed Y. Tawfik **Aidong Peng**

School of Computer Science
University of Windsor
Windsor, Ontario, Canada
<atawfik|peng2>@uwindsor.ca

Seyed M. Tabib **Baher Abdulhai**

Intelligent Transportation Research Centre
University of Toronto
Toronto, Ontario, Canada
<baher|mtabib>@ecf.utoronto.ca

Abstract

Vehicle reidentification is the process of tracking a vehicle along a highway as it crosses inductive loop detectors. The present work uses decision tree induction to generate a specific decision tree for tracking vehicles along specific high sections. Initial experimental results show that this approach performs well specially when coupled with signature matching techniques.

Keywords: Decision tree learning, Vehicle reidentification, Intelligent highway systems

Introduction

Despite recent advent of an array of new detection systems, such as video, ultrasound, infrared, and microwave based detectors, Inductive Loop Detectors (ILD) still remain among the most widely used, and the most widely invested-in detection technologies. The ILDs provide the Intelligent Transportation Systems (ITS) with a constant source of information on traffic system conditions. ILDs detect vehicles as they pass over their electromagnetic field and provide counts, occupancies, and speeds. These measurements however are 'point' measures and do not necessarily reflect wide-area traffic conditions. For this reason, ITS researchers have been searching for 'better' detection systems. If these detectors could be used in a "smarter" way, more useful information or section measures of traffic system performance such as travel time and density could be obtained. This in turn translates into better traffic management and information systems via the use of accurate section measures. One critical application using section measures is dynamic origin/destination demand estimation. This application is a vital component of other intelligent traffic management strategies such as traveller information, traffic assignment, and route guidance. The way to "squeeze out" more information from loop detectors is by using the vehicle waveforms that are produced when a vehicle passes over the loop detector. In order to produce meaningful traffic information such as travel time and density, the signature of a vehicle needs to be reidentified at different sites.

Recently, research investigating the possibility of more 'intelligent' usage of ILD, has focused on deriving 'wide-area'/'section-related' measures from their outputs, as opposed to the limited conventional point measurements [3]. Such new approaches rely on matching vehicle features at two successive ILD stations in order to uniquely identify individual vehicles and hence derive true travel times and densities over a highway section. The significant implication of these emerging efforts is that the already existing vehicle detection infrastructure can be revived and used for numerous ITS applications without the need for additional expensive sensor technology, and without the loss of functionality.

Section-related tracking aims at detecting and re-matching vehicle signatures from successive loop detectors. As an individual vehicle passes over a loop sensor location for instance, a spatiotemporal signature is generated reflecting the levels of inductance unique to the vehicle as it spatially progresses over the detector in time. The uniqueness of the distribution of the metal mass along the body of a given vehicle results in a unique signature.

A general transportation network can be modelled using individual links where each link has only one ingress point and one egress point. Some examples of network links include multilane freeway sections, arterial sections, and ramps. If both the beginning and the end of each link is appropriately instrumented with a detector station, say a loop station with vehicle signature output, then vehicle waveforms can be obtained from the beginning (upstream) and ending (downstream) detectors. A vehicle waveform pair can then be formed using one downstream waveform and one upstream waveform. The set of vehicle waveform pairs is increasing over time as more vehicles enter and exit the link. If approached in a sequential fashion, the vehicle reidentification problem is to find the matching upstream vehicle waveform given a downstream vehicle waveform. Stated more formally, the vehicle reidentification problem is as follows:

Given a set of vehicle waveform pairs x_i, where x_i ? S, $i= 1,...,N_c$, S is the set of vehicle waveform pairs, and N_c is the number of waveform pair combinations. Find the waveform pair which is produced by the same vehicle.

The single solution of the reidentification problem produces the origin/destination tracking of a single while the sequential solution of a stream of vehicles will produce measures such as link density and group travel time. Both kinds of solutions produce valuable data that are of interest to operating agencies, researchers, and individual travellers.

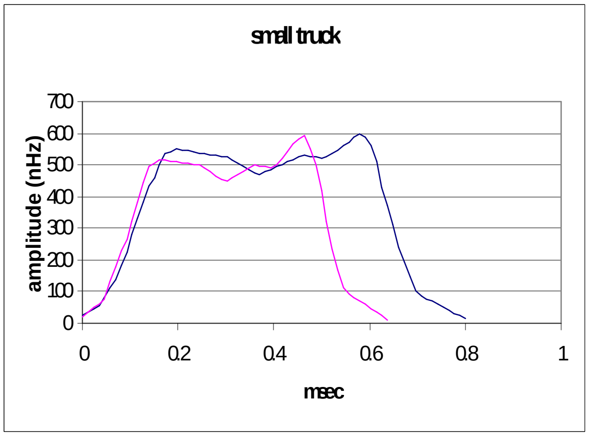

Figure 1. Vehicle Signature for a small truck

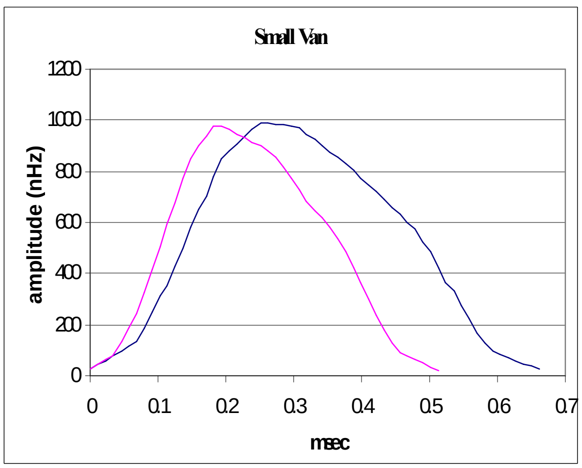

Figure 2. Vehicle Signature for a small van

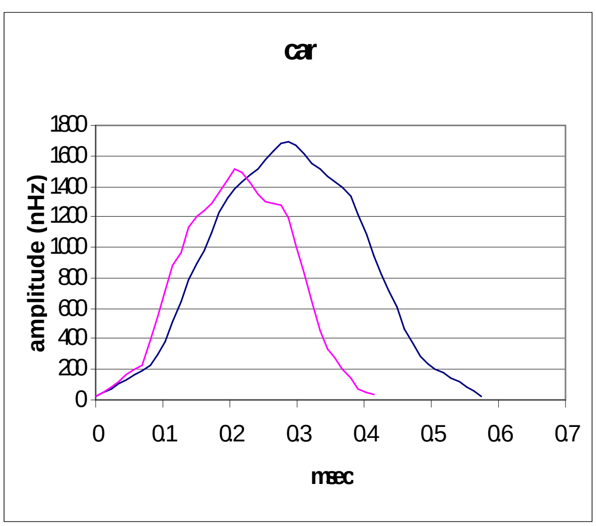

Figure 3. Vehicle Signature for a car

The inductive loop detector (ILD) operates on the principle of mutual inductance between the loop and the equivalent conducting plate of a vehicle. When a metallic mass passes through the magnetic field generated by the inductive loop, the disturbance produces a net reduction in the loop inductance or frequency, and the resonance circuit properties are altered. A motorcycle could, for example, produce a frequency shift of up to 0.08% (80 nHz, while an automobile could cause a shift of up to 3% (3500 nHz). The metallic component of the vehicle is what disturbs the loop inductance. As a result, double-axle trucks produce a twin-peaked vehicle signature (Figure 1) when the resolution of the detector is adequate. Thus, in concept this method can easily be used for vehicle-type identification purposes as the signature of a van (Figure 2) or that of a car (Figure 3) will be significantly different.

Vehicle identification is based on a set of feature vectors extracted from the raw vehicle waveforms, and other information. Such additional information can include vehicle speed computed from speed trap, known lane number, and location geometry. The main assumption in this approach is that the feature vectors would be less sensitive to the error disturbances than to the input, otherwise noise would dominate the system making classification impossible.

Raw data are obtained from sampling the inductive loop magnitude at intervals of every 11-14 ms (i.e. the scan rate of the detectors cards used during SR-24 data collection). The signal magnitude is normalized with respect to it's maximum amplitude to eliminate upstream and downstream variations. Further normalization adjusted for speed difference and re-sampling of a spline interpolation allowed point to point comparison despite differences in initial signal sampling.

Even after the described waveform transformation, some variability remained between upstream and downstream waveforms due to the following reasons:

?? Vehicle entrance angle into inductance field.

?? Vehicle offset from the physical loop center, even straddling (e.g. vehicle's inductance if covering half the loop is reduced from 3500 nHz to 2000 nHz).

?? Height from loop due to the suspension system.

?? Inaccuracies in the measured speed from the speed trap configuration .

?? Errors in the interpolation.

?? Aliasing due to sampling.

?? Quantization error in the A/D process.

?? Variability between upstream and downstream inductive loop systems including the physical loop installation and the loop energizing circuitry.

Lexicographic Optimization Problem.

The vehicle re-identification problem can be formulated as a lexicographical optimization problem [3,4].
The first objective constrains the set of vehicle pairs to those who had travel times within a lower and higher limit (i.e. a time window).

The second objective selects the pairs of vehicles that contain a candidate upstream vehicle magnitude that was within a certain percentage of the downstream vehicle that is being reidentified. This results from the fact that the inductance magnitude is inversely proportional to the height of the vehicle. For example, trucks would have much smaller magnitudes than passenger vehicles, so the candidate set is limited to trucks by using this constraint.

The third goal limits the candidate vehicle's length to be within a certain percentage of the downstream vehicle.
The fourth level lexicographic optimization objective is
to minimize the discrepancies between the upstream feature vector and the downstream feature vector in a vehicle waveform pair. Five different distance measures are used in the research: Euclidean, Correlation, Similarity, Lebesque, and Neural networks [5].

The last level objective is tries to optimize the distance measure with the maximum posterior probability. Equivalently, this resulted in the choosing of the "most likely" distance measure that would identify the correct match between a vehicle pair. The selection of the optimum objective function is therefore done by Bayesian analysis.About 75% correct matching was reported using the above described method for the passenger cars.

Decision Trees

Lexicographic methods provide an order of preferences in which options are compared on the most important criterion, and the best option is chosen unless other options tie for first place. In that case, evaluations on the second most important criterion are considered to break the tie. If that is not possible then the third most important criterion is consulted, and so on until one option can be chosen.

Decision tree induction algorithms [2] rely on information theoretic measure known as entropy to prioritize the most informative attributes. The resulting decision diagram shows the outcomes that may occur for a series of interdependent decisions. The actual outcome of each of the individual decisions at each stage is not known with certainty.

The attributes available for reidentification include vehicle specific ones such as the vehicle length and its signature, as well as other attributes that rely on the spatiotemporal context, the behavior of the driver and traffic/road conditions. These attributes include lane change, speed variation, and the time taken to cross a highway section.

Therefore, it is realistic to expect that for each highway section, a different decision tree should be used. Moreover, we expect that the induced trees may vary significantly with traffic conditions.

To induce decision trees, it is necessary to perform visual matching for some training data and apply a decision tree induction algorithm. The decision trees generated are reviewed by the user for soundness and tested on a validation set before adoption for practical use.

Experimental Evaluation

The test data were obtained from a field site on the westbound SR-24 freeway in Lafayatte, California in December 1996. Two data acquisition stations were instrumented with video, loop waveform dataloggers, and speed trap dataloggers. Standard 6'x 6' (1.82 m) loops were used at both stations. Several hours of data were collected, but a smaller portion of the data was reduced into two data sets and used for this initial investigation in vehicle reidentification. The reduced data set contains the waveforms of the upstream vehicles along with their speed, electrical length (derived from occupancy time), arrival time at stations, and the vehicle identification number. One data set was composed of moderate flow traffic (1000 Veh/h.l) and contained approximately 2000 vehicles. This dataset was recorded on December 6 at about 12:00 pm. Another dataset was composed of congested flow traffic (1800 veh/h.l) and contained approximately 3000 vehicles. The second dataset was recorded on December 12 at 8:00 am during the morning rush hour. Both datasets were divided into training and testing datasets. Figure 4 shows the data collection set-up for the SR-24 site.

The test data were stored in SIG (signature) files and ground truthed by video correlation. The unprocessed signature files and PVR (Per Vehicle Record) files were collected in the field. The video correlation database was produced in the laboratory from the video footage taken during data collection. Manual correlation of the downstream with the upstream vehicle using video was required for the development of the video correlation database.

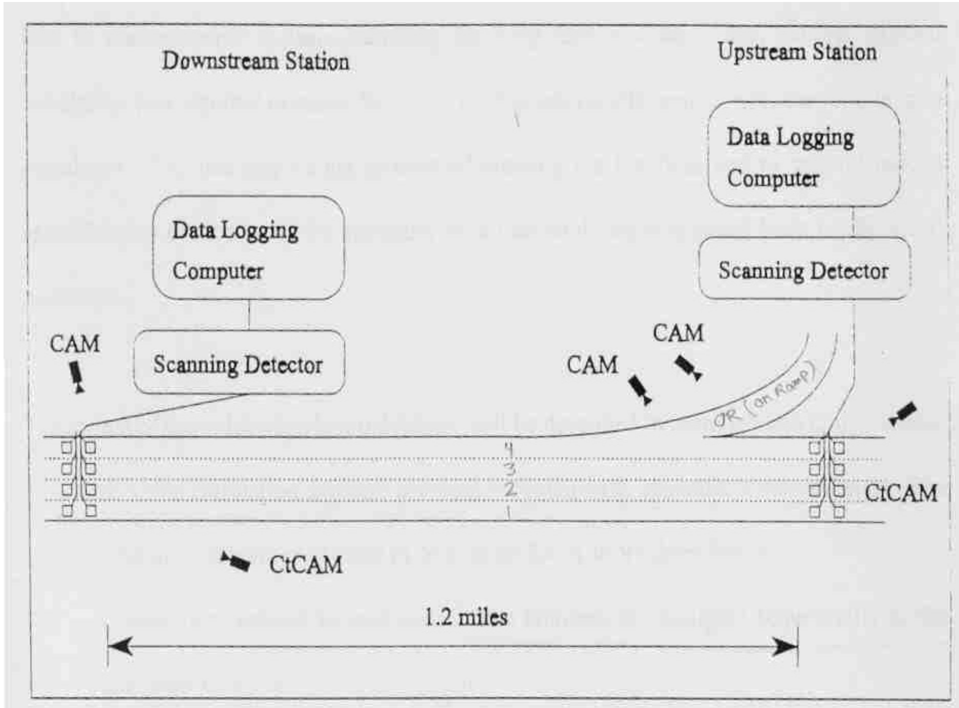

Figure 4. Lafayette data collection site

The dataset used for learning decision trees consists of the following attributes for an upstream and a downstream vehicles: absolute length difference, time difference between crossing the two IDLs, absolute speed difference, and aboslute lane difference. The training / test datasets have an additional attribute indicating if the record represents a match.

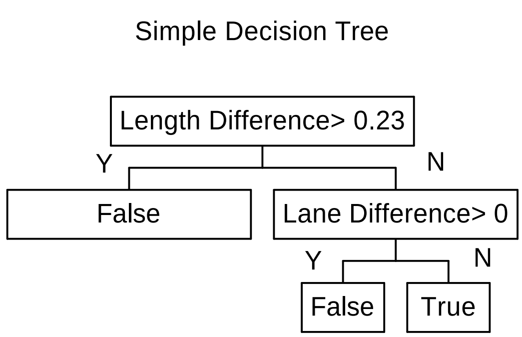

Figure 5. Simple Vehicle Reindentification Decision Tree.

The dataset is highly imbalanced as there are a lot more mismatches that matches. However, this imbalance is artificial to some extent because it is not necessry to compare each vehicle with all others. Spatiotemporal constraints would limit the comparison to vehicles that have crossed one ILD and not the other as described laster in this Section. Many techniques have been developed to allow effective learning from imbalanced data [1]. In this work, we use random sampling from the larger class to introduce some balance and avoid trivial decision trees that simply assume no matching. In addition, the ability to reliably determine the car type or category (car, truck, van, ...etc) eliminates many candidates. In order to eliminate the dependence on the actual traffic intensity in the experimental dataset we considered the following criteria to limit the number of vehicles to be compared to a given vehicle:

?? randomly selecting the comparison set
?? forcing a fixed size comparison set, and
?? limiting the comparison set by a fixed time duration

Initially, we have also limited the test set to a small set of vehicles containing 180 to 360 vehicles.

The decision trees generated differ somewhat in their complexity and performance depending on the constraints imposed on the training set. A simple decision tree, shown in Figure 5, that considers length and lane differences only properly matches 85% of the training set and 75% of the test set.

However, in all decision trees generated, the length difference attribute and the lane difference attribute proved to be important in identifying matches. Other attributes such as time difference and speed difference appear in some trees. A rather complex tree involving these attributes has been successful in learning 95% of the training cases and correctly matching 85% of test cases.

In a subsequent set of experiments, we used the travel time distribution shown in Figure 6 to limit the set of vehicles in the comparison group. About 95% of the vehicles arrive at the second loop detector within 160 seconds from crossing the first one but not before 88 seconds. This condition reduced the average number of vehicles to compare to 80 vehicles during the rush hour dataset and 46 vehicles for the other dataset.

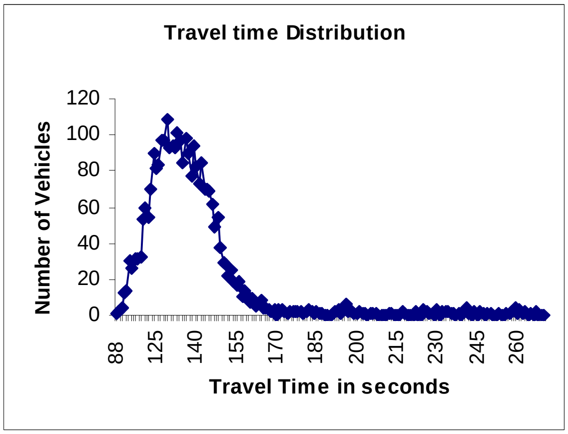

Figure 6. Travel time between the loop detectors

```
lengthDifference <= 0.23 :
|  laneDifference <= 0 : true
|  laneDifferemce > 0 :
|  |  lengthDiffference > 0.16 : false
|  |  lengthDifference <= 0.16 :
|  |  |  laneDifference <= 1 : true
|  |  |  laneDifference > 1 : false
lengthDifference > 0.23 :
|  laneDif fernce> 0 : false
|  laneDifference <= 0 :
|  |  timeDifference<= 0.79 :
|  |  |  lengthDifference <= 0.38 :
|  |  |  |  lengthDiffence <= 0.29
|  |  |  |  lengthDifference > 0.29
|  |  |  lengthDifference > 0.38 :
|  |  |  |  lengthDifference > 0.51 : false
|  |  |  |  lengthDifference <= 0.51 :
|  |  |  |  |  timeDifference <= 0.552 : false
|  |  |  |  |  timeDifference > 0.552 :
|  |  |  |  |  |  speedDifference<= 9.7 : true
|  |  |  |  |  |  speedDifferemce > 9.7 : false
|  |  timeDifference > 0.79 :
|  |  |  lengthDifference <= 1.59 : true
|  |  |  lengthDifference > 1.59 : false
```

Figure 7. A more complex decision tree

The more complex decision tree shown in Figure 7 has been used to classify vehicles in the rush hour test set after imposing the travel time constraints. The decision tree correctly identified 89% of the matches, and 78% of the mismatches. However, about 20% of the records were incorrectly labelled as matching. In fact, less than 10% of the reported matches are actual matches. To address this problem, it is necessary to compare the electromagnetic signatures for all reported matches as described in [5] to eliminate the majority of incorrect matches. This brings the percentage of correctly classified records to 90%.

Although the training set used to generate the tree in Figure 7 has been extracted from the rush hour dataset, the same tree performed well when tested on the lower traffic dataset. About 82% of the matches are correctly labelled as well as 82% of the mismatches. However, true matches continued to represent 10% of the records labelled as matching.

It is therefore necessary to consider the signature data to get reliable reidentification of vehicles. The decision tree can successfully reduce the number of records requiring signature verification by more than 75% making it possible to achieve 90% correct classification. It is also worth noting that correct identification of each vehicle may be impossible without using additional visual clues.

The following table compares the performance of different approaches to vehicle reidentification of passenger cars

Approach	Correct matching
Lexicographic Optimization[3,4]	75%
Signature matching [5]	61%
Decision Tree	82%
Decision tree + signature	90% (estimated)

Table 1. Accuracy of Vehicle Reidentification

The accuracy result reported in the above table for signature matching is based on the percentage of correctly matched passenger cars in heavy traffic. The accuracy of signature matching varies for different types of vehicles ranging from 42% for passenger cars in moderate traffic to 100% for trucks. The estimated rate of correct matching in the last row of the above table is based on some initial results.

Conclusion and Future Work

Comparing these results to those reported in [3], [4], and [5], the decision tree achieves comparable results even without comparing actual signatures. Performing signature matching as described in [5] helps further reducing the error rate by an additional 5% to 10%. We are currently experimenting with applying signature matching algorithms to branches of the decision tree with a high error rate, this error is decreased.

It is clear form the decision tree generated that spatio-temporal properties such as persistence of lane, speed can provide useful clues for vehicle reidentification.

Acknowledgments

The authors acknowledge the support of the Natural Sciences and Engineering Research Council (NSERC), Canada Foundation for Innovation, The University of Toronto and the University of California Irvine.

REFERENCES

[1]. Japkowicz, N. (2000), Learning from imbalanced data sets: A comparison of various strategies, In Papers form the AAAI Workshop on Learning from Imbalanced Data Sets, AAAI Press.

[2]. Quinlan, J.R. (1986) Induction of decision trees. Machine Learning, 1, 81-106.

[3]. Sun, Carlos, Ritchie, S. G., Tsai, Kevin, and Jayakrishnan, R. (1999). "Use of Vehicle Signature Analysis and Lexicographic Optimization for Vehicle Reidentification on Freeways." TRB Annual meeting.

[4]. Sun, Carlos (1998). *Use of Vehicle Signature Analysis and Lexicographic Optimization for Vehicle Reidentification on Freeways.* PhD Dissertation, University of California, Irvine.

[5]. Tabib, Seyed (2001). Vehicle re-identification based on inductance signature matching, M. Sc. Thesis, University of Toronto, Toronto, ON, Canada, 131 pages.