Towards a Behavioural Traffic Monitoring System

Marco Rigolli*
marco@robots.ox.ac.uk

Michael Brady
jmb@robots.ox.ac.uk

Robotics Research Group
Department of Engineering Science
University of Oxford
Oxford, UK

ABSTRACT

In recent years, traffic video surveillance has increased significantly. However, most of the footage is reviewed by humans or not at all. Therefore, tools capable of analysing traffic video sequences and autonomously extracting information are required. In this paper, we present an agent-based approach to analysing driver behaviour. Our work differs from normal road monitoring systems in that we are interested in inferences about driver behaviour and in learning 'normal' driving modes, rather than specific instances of driver actions. Our system provides a behavioural description of traffic scenes. First, we present a kinematic traffic simulator designed to test driving agents. The simulator supports multiple lanes, obstructions and different environmental conditions. Second, we specify the agent's perception and reasoning models. Contrary to current autonomous driving systems, our behavioural models primarily influence agent perception. This approach is supported by recent psychological studies carried out on human drivers. Furthermore it simplifies the system implementation, increasing the ease of testing alternative models. By embedding the agents in the simulator, we observe classical traffic behaviour. Finally, we suggest ways to use the system's results directly or within higher level tools.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems, Intelligent agents*; I.2.10 [**Artificial Intelligence**]: Vision and Scene Understanding—*Perceptual reasoning, Video analysis*

General Terms

Design, Human Factors, Theory

*The primary author of this paper is a graduate student.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.

Keywords

Behavioural Traffic Monitoring, Multi-Agent Systems

1. INTRODUCTION

In recent years, traffic video surveillance has increased significantly, especially on motorways. It already plays a significant role in monitoring the M25 London Orbital, the strategic hub of Britain's motorway network, through over 70 cameras in constant use 24 hours a day, 365 days a year [11]. At the moment, most of the footage is analysed by human operators, who assess whether traffic conditions are normal or are requiring intervention, or is not used at all. Current advances in vehicle tracking [2, 8] suggest that in the near future reliable real-time information about vehicles' location, speed and acceleration will be available. However, existing users of Intelligent Transportation Systems (ITS) already feel a need for technology capable of providing higher-value information. Particularly, as shown in [10], they are interested in systems which will help them to:

- effectively investigate collisions and accidents through modelling software and simulator tools;
- ensure drivers are fit to drive by detecting drunk drivers or drivers under the influence of drugs;
- improve road safety, road user's behaviour and road design through traffic system monitoring and direct intervention.

This paper presents a behavioural, multi-agent approach to modelling human drivers. Section 3 describes the kinematic simulator developed as an environment for our agents. Section 4 introduces our driving agent model with particular attention to its reasoning and perception modules. The macroscopic and microscopic performance of the system is qualitatively assessed in Section 5, while possible real-world uses are presented in Section 6.

2. RELATED WORK

Classical traffic systems have concentrated on estimating macroscopic behavioural parameters [4]. Their main purpose is to provide users with statistical information about the likely effect of structural (e.g. road works), normative (e.g. speed limit modifications) and environmental changes (e.g. bank holidays, peak hours) on the traffic flow. The classical approaches have been either to use fluid dynamics equations, or to model the relationship between a vehicle's

acceleration and the distance to the vehicle ahead (leader) under the assumptions of one-lane, one-way roads and that all drivers follow the same base model. Recently the focus has started to shift towards multi-lane [5], multi-class scenarios [7] and multi-agent systems [1] in order to provide higher quality flow predictions.

The classical approaches are a useful tool for traffic engineers and for journey planning, as they models global traffic flow; however, they are unable to provide a meaningful analysis of the local situations filmed by the surveillance cameras.

In the past decade, development has started on traffic monitoring systems aimed at analysing specific traffic situations. Buxton and Howarth [3] investigated a system that provides basic analysis of traffic situations produced at a roundabout junction: such as overtaking, following and waiting to enter the junction. Their work introduces the idea that a deictic approach is needed to successfully interpret traffic events, i.e. that traffic situations should be analysed from the viewpoint of the agents involved rather than from an 'omniscient observer'. In order to achieve task recognition the system relies on several layers of situated knowledge-bases in the form of a heavily segmented background, condition-actions rules and Bayesian networks with explicitly provided conditional probabilities. More recently, in order to reduce the amount of human provided knowledge required, Fernyhough et al. [6] have developed a system that attempts to automatically segment the background of video sequences into regions of interest using statistical analysis of trajectory information. When this is complemented by prior knowledge about object interactions and events composition, the system is able to detect overtaking manoeuvres on motorways.

The main limitation of these monitoring systems is that, although capable of identifying the occurrence of specific events, they are unable to provide a description of why those events have occurred. Essentially they can provide ITS users with a description of *what* has happened in a video but not *why* it happened as they lack a model of the drivers internal motivations. Therefore they are currently unsuitable for monitoring road users' behaviour (such as aggressive vs safe) or to assess whether the same action would be carried out by different driver types (such as sober vs drunk), which are two of the three needs identified in Section 1.

3. THE SIMULATOR

One of the most fundamental choices in any traffic system is which motion model should be used. This choice is important because it influences not only the realism of simulations, but also the agents' reasoning capabilities. In order to capture the driving agents' behaviours we implemented a kinematic simulation. This allows the modelling of realistic lane-changes, while being relatively easy to code and fast to execute. The main limitation is that the model doesn't allow for sliding or realistic collisions; however, we believe that such considerations play little role in human decision making processes while driving normally.

The simulator models a highway with no entry or exit lanes, desirable in order to be able to separate 'goal oriented' driving, such as moving towards slower lanes in order to exit at a junction, from purely 'behavioural oriented' driving. It allows modelling of road obstructions, through the placement of unmovable agents; different environmental condition such as rain and ice by changing the braking capacity of the agents; and fog by modifying the visibility range of the agents. Figure 1 shows a screenshot of the 12–lane simulator: each lane is 4m wide, cars are 4m long and 2m wide, while trucks are 7.5m by 3m.

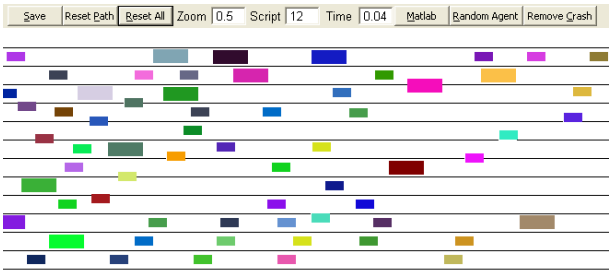

Figure 1: ***Screenshot of a simulated 12-lane highway.***

At each time step, which was set to 0.04s for the results shown in this paper, each agent updates its world model by assessing the location and speed of its neighbours. At the moment, the reasoning limits itself to an 8-neighbourhood. Then, based on the world model and its behavioural parameter, each agent can perform one of six actions: *accelerate*, *constant*, *decelerate a little* (the agents brakes at its preferred rate), *decelerate a lot* (the agent brakes at the vehicle maximum capacity), *overtake* or *get-back* (the agent moves towards a slower lane). Depending on which action the agent has decided to carry out, the simulator will then update its position and velocity while ensuring that speeds and acceleration are within the vehicles capabilities.

In order to test different scenarios we have developed a scripting language which allows us to either generate specific agents at given time intervals or to generate random agents according to different distribution functions such as the Poisson distribution.

4. THE DRIVING AGENT

4.1 The Model

Developing driving agents capable of interacting with real traffic, either human or artificial, has been the subject of autonomous driving research. In autonomous driving it is assumed that the environment can be accessed through onboard sensors and that the main task faced by the agent is fast action selection. The classical approach (Figure 2) is to build agents with different specialised reasoning modules such as lane tracking, car following, obstacle avoidance and lane changing. The optimal action is then chosen by weighting the suggestions of the different reasoning modules by a set of behavioural parameters, usually found through brute force or genetic programming. This approach, although capable of emulating specific driver types, is not helpful in providing a description of traffic situations as the action selection method and the physical and psychological meaning of the behavioural parameters often bear little resemblance to human thinking.

To develop a more 'human-like' driving agent we made the following assumptions regarding human behaviour:

- drivers are taught 'good' vs 'bad' driving in a stan-

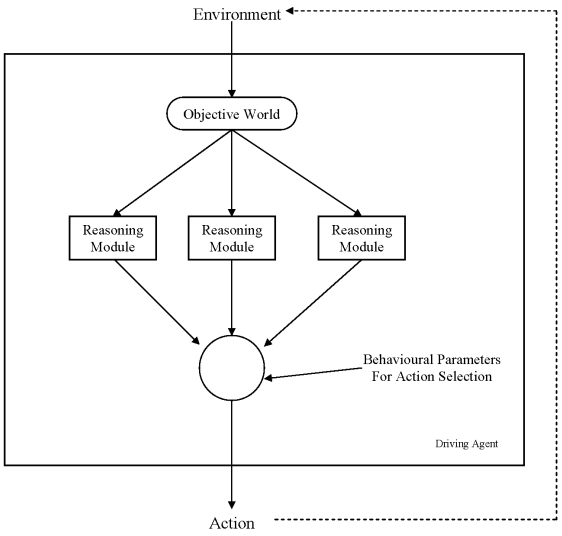

Figure 2: ***Classical autonomous driving agent system.***

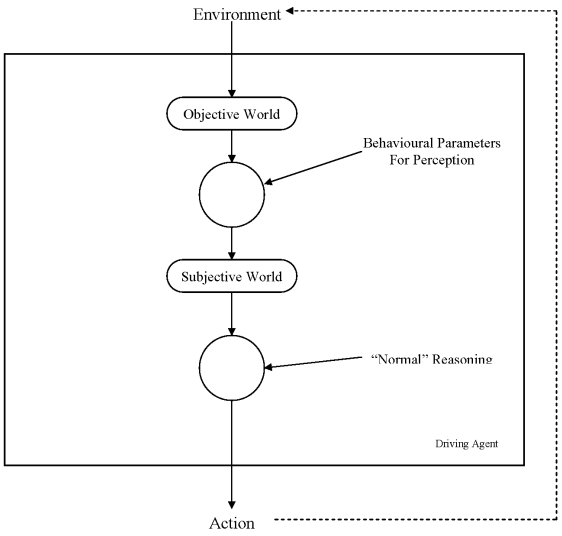

Figure 3: ***The driving agent implemented in this paper.***

dardised way as part of the process required to obtain a driving licence,

- most drivers believe that they have carried out the 'right' action given the situation they were in,
- drivers tend to describe other drivers behaviours in terms of their own driving style.

These assumptions are supported by the research [9] carried out by the Driver Behaviour Research Group (DBRG) for the UK Department of Transport which has identified three phases of driving:

1. Technical Mastery: driver learns to control, position and manoeuvre the vehicle.
2. Reading the Road: in this phase the driver learns to interpret, anticipate and interact with other drivers.
3. Expressive Phase: after having acquired sufficient competence in the first two phases, drivers use the manner in which they drive to give expression to personality, attitudinal and motivational characteristics.

In the UK, the practical driving test makes a through assessment of technical mastery and requires demonstration of a reasonable level of road literacy.

We believe that, when unconstrained by specific goals, drivers perform different actions not because of different reasoning, but because they perceive the world differently; as stated in [9] "Safe and risky drivers almost certainly differ in their interpretations of situations". Our agent therefore uses a set of parameters (section 4.3) to generate a qualitative description (*subjective world*) of the sensed environment (*objective world*). The reasoning module (section 4.2) is then implemented on top of the subjective environment in a straightforward way (Figure 3). This allows coding 'normal' behaviour and local knowledge while accounting for the variety of driving styles observed through perceptual differences.

4.2 Reasoning Module

The driving agent reasoning module is based on the agent's qualitative perception of the environment. The idea is that all drivers try to achieve a comfortable state which provides an acceptable compromise between safety and travelling at a desired speed. Therefore an agent satisfied with its current state will not change its course of action without some external influence. The reasoning module is composed of two parts:

- *action proposing*: during this phase a suitable action space is suggested in reaction to perceived external influences.
- *action filtering*: in this phase the action space generated above is filtered by assessing secondary environmental conditions.

In the case of a highway, the *action proposing* environmental circumstance is the distance between the agent and the vehicle ahead, as no driver would consciously choose to be part of a traffic accident. For example, a vehicle suddenly pulling in front of an agent requires a response. *Action filtering* conditions are: the difference between current speed and desired speed, the gap in the overtaking and get-back lanes and the presence of a close follower. For example, a very small gap on the faster lane precludes the possibility of overtaking, but it does not require a response from the agent.

A more detailed example follows: when approaching a vehicle (i.e. getting too close to be safe), the action space generated by an agent would be: *overtake*, *decelerate-little*. The agent then chooses an appropriate action after checking two things:

i) Is it safe (by checking that the gap in the overtaking lane is large enough)

ii) Is it desirable (by comparing its current speed with its desired velocity and deciding if it would like to travel faster).

Table 1: Behavioural Parameters

Currently Used	
Name	Description
$V_{desired}$	The desired speed on an empty road.
$T_{critical}$	The stopping time below which the agent considers itself in an emergency situation with respect to the leading vehicle.
$T_{braking}$	The stopping time below which the agent considers itself as approaching the leading vehicle.
$T_{accelerate}$	The stopping time above which the agent considers itself as free-flowing i.e. unconstrained by the leading vehicle.
T_{pushed}	The stopping time below which the follower vehicle is perceived as tailgating.
T_{leader}	The stopping time above which it is considered safe to perform a lane change with respect to the new leader.
$T_{follower}$	The stopping time above which it is considered safe to perform a lane change with respect to the new follower.

The reasoning rules have been implemented within the functional paradigm as this allows us to exhaustively test the core reasoning module. Also by not being hard-coded into the rest of the system it can be to easily changed for different situations, such as junctions and roundabouts.

4.3 Perception Module

A 'human-like' reasoning system needs to be complemented by a believable mapping between 'objective' and 'subjective' descriptions of the world. Our agents' perception module relies on the following two assumptions:

- Human drivers describe situations and reason in qualitative terms such as: 'too slow', 'not far enough' etc.
- Drivers are mostly aware of the capabilities of their vehicle and can generally judge their distance from other vehicles. This appears to be a fair assumptions as traffic accidents, although a serious problem, do not happen as frequently as they would otherwise do.

The assumptions above allow moving from a continuous description of the world to a discrete one. Equation 1 and the agent's behavioural parameters (Table 1) are used to translate objective distances between the agent and its neighbours into a set of qualitative descriptions.

$$D_{thres} = \frac{{V_s}^2}{2D_s} - \frac{{V_n}^2}{2D_n} + T_s V_s + X_0 \quad (1)$$

Where V is velocity in m/s, D is braking capacity in m/s^2, X_0 is the desired space between halted vehicles and T is a behavioural parameter expressed in seconds. The subscript s and n refer to subject agent and neighbour respectively.

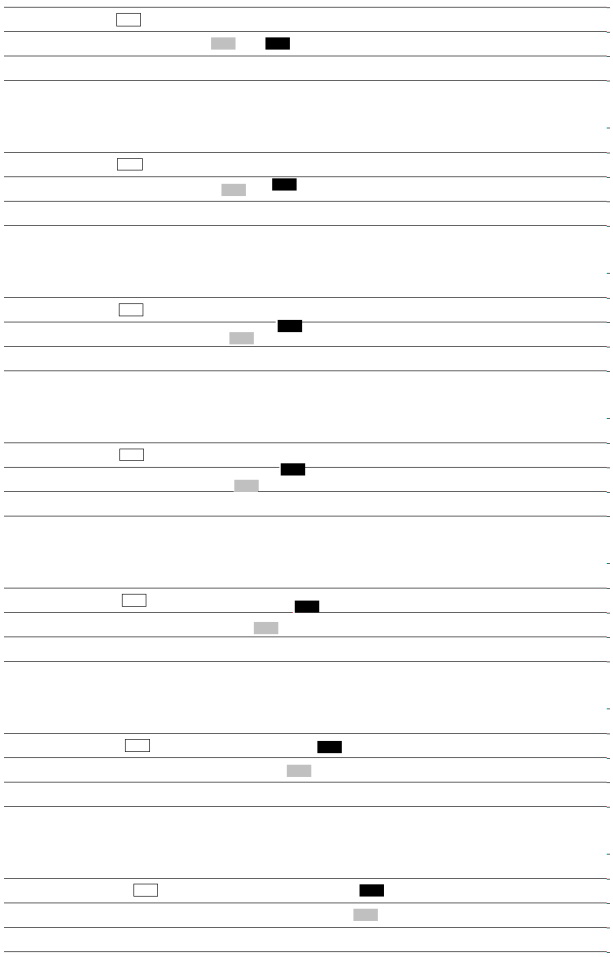

Figure 4: ***Sequence showing an aborted lane-change. The light grey vehicle starts to change lane in order to overtake the black vehicle — when it realises the black vehicle is moving towards the slower lane it stops the lane change.***

Finally D_{thres} is the distance which acts as a threshold between different qualitative interpretations. The advantage of using Equation 1 is that it has a clear physical meaning, i.e. it represents a stopping distance, and that behavioural parameters (T_s) also have a clear meaning i.e. they represent 'time to impact'. The agent can then assess if it is either: critically close to the car ahead; approaching it; following it; or free-driving. It can also assess whether the gap in another lane is safe or unsafe to move into or if the vehicle behind is tailgating. For example, the "two-second rule" taught in many British driving schools can be simulated by setting $T_s = 2$ and assuming $V_s = V_n$ and $D_s = D_n$.

5. QUALITATIVE SYSTEM ASSESSMENT

5.1 Microscopic Performance

The microscopic performance of the agents was considered by a number of experienced drivers as realistic both with respect to car-following and lane-changing decision making. The simple reasoning system implemented was capable of modelling realistic behaviours such as reduced number of

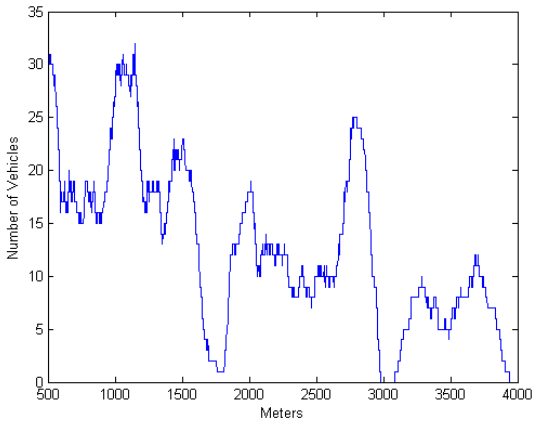

Figure 5: ***Compression Waves: alternating zones of high and low density (window size of 100m).***

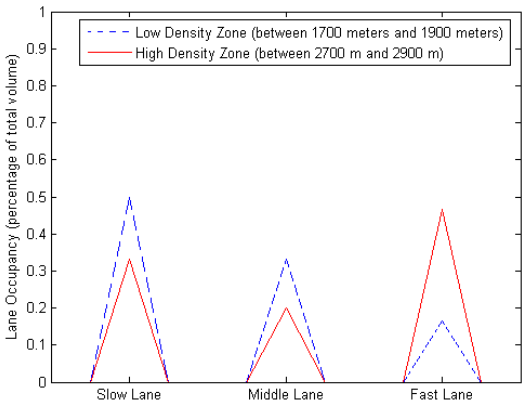

Figure 6: ***Lane Occupancy: the locations refer to Figure 5.***

unnecessary lane changes, weak interaction with a follower agent or no inside lane overtaking. Furthermore the agents showed behaviours they had not been programmed for such as aborted lane-changes (Figure 4) i.e. the agents are able to correctly respond to unforseen events (such as a pinch condition) and abort an already initiated lane change.

The natural occurrence of aborted overtaking is due to the fact that the system, although showing planning and anticipation thanks to the reasoning module, is essentially reactive. There was no need to impose a strict 'persistence of intentions' due to the overall smoothly varying nature of the environment and the qualitative approach to action selection; this allows the agent to react in a timely fashion to those environmental changes it perceives as significant while not showing erratic changes between different action plans.

The system was also successful in representing different driving styles, as human observers were able to clearly identify agents acting 'unsafe'. This was particularly interesting as currently acceleration and deceleration rates are the same across all agents, which seems to suggest that what distinguishes a safe driver from a dangerous driver is 'when' they choose to carry out a certain action.

5.2 Macroscopic Performance

Although the system had not been designed for macroscopic analysis, we decided to verify whether credible microscopic behaviours would lead to the emergence of realistic macroscopic data. This was also true as the system showed the occurrence of classical results such as compression waves (Figure 5) i.e. the presence of zones of higher and lower vehicular density.

Furthermore, the simulated agents displayed the lane occupancy trends seen in real life traffic (Figure 6) i.e. when traffic volume is low the majority of drivers are well behaved and prefer to travel on the slow lane, while in areas of high traffic volume, drivers tend to accumulate on the faster lanes. These findings provide further confidence in the driving agent developed.

The data used to produce the figures shown was obtained by simulating 330 agents with randomly generated behavioural parameters over 10Km of a 3-lane motorway and a time span of 15 minutes.

6. DISCUSSION ON POSSIBLE USAGE

6.1 Behaviour Clustering

An area which is not currently addressed by traffic monitoring systems is the clustering of drivers by behaviour types such as 'safe', 'unsafe', 'intoxicated' or 'learner'; information that ITS users would find highly-valuable. By being able to summarise the performance of a driver across time in seven parameters (Table 1) our driving agents offer a starting point to address this need. By providing our system with a set of behaviour types, it is possible to cluster previously unseen drivers as either belonging to one of the standard types or displaying new behaviour that requires immediate human attention. However, the task of learning standard behaviours from raw data of real human drivers has not yet been tackled and will be an area of future research.

6.2 Higher-level Tools Integration

Once a driver has been assigned a behaviour type it is possible to generate an automatic commentary of its actions and supposed motivations (Figure 7). This is useful as a way to automatically log video data. However, it can provide even greater benefits to some of the traffic monitoring systems introduced in Section 2 as it can significantly simplify their reasoning system by removing the noise generated by different driving styles and by providing higher-level information upon which to reason. For example, the white agents shown in Figure 8 and Figure 9 present similar spacing and velocities between each other; however a human observer would claim that Figure 8 shows a queue while Figure 9 shows five friends travelling together. To be able to capture the difference between the two situations we must be able to identify that, if Figure 9 was showing a queue, the normal behaviour would be for one of the agents to overtake the others. To obtain such understanding from just a rule-based system or background segmentation would require a significantly complex and fragile system. Conversely, if provided with the information that the agents in Figure 9 are satisfied with their current velocity, while the agents in Figure 8 do not have the option to overtake, it is easy to recognise the difference between these two scenarios.

```
At frame 6, special-agent-0 thinks:
 it is FREE-FLOWING;
 it is safe to move to NEITHER lanes;
 it would like to be travelling FASTER;
 the vehicle behind is NOT-PUSHING;
 Hence it decides to: ACCELERATE.
At frame 14, special-agent-0 thinks:
 it is FREE-FLOWING;
 it is safe to move to NEITHER lanes;
 it would like to be travelling AS-IT-IS;
 the vehicle behind is NOT-PUSHING;
 Hence it decides to: CONSTANT.
At frame 46, special-agent-0 thinks:
 it is APPROACHING the vehicle in front;
 it is safe to move to RIGHT lane;
 it would like to be travelling AS-IT-IS;
 the vehicle behind is NOT-PUSHING;
 Hence it decides to: DECELERATE-LITTLE.
At frame 59, special-agent-0 thinks:
 it is APPROACHING the vehicle in front;
 it is safe to move to RIGHT lane;
 it would like to be travelling FASTER;
 the vehicle behind is NOT-PUSHING;
 Hence it decides to: OVERTAKE.
```

Figure 7: ***Automatic Commentary: agent-0 is confronted by a vehicle that has pulled out in front of it.***

Figure 8: ***The white agents are in a queue***

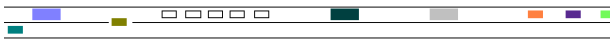

Figure 9: ***The white agents are not in a queue***

7. CONCLUSION

This paper has presented a novel approach to modelling human drivers, based on the idea that different behaviours lead to different interpretations of the surrounding environment rather than different types of reasoning. We have shown that both the microscopic and macroscopic performance of our driving agents is realistic and has a number of possible uses. Further work is required in developing efficient ways to learn behavioural parameters from real traffic data and to assess the complexity of using the current model in other traffic situations such as two-way traffic, road junctions and roundabouts.

8. ACKNOWLEDGMENTS

Support of this research by the *European Union (CogViSys Project, IST-2000-294004)* is acknowledged.

We are grateful to *Dr. Mark Gooding* for his assistance in the experimental work.

9. REFERENCES

[1] M. Balmer, N. Cetin, K. Nagel, and B. Raney. Towards truly agent-based traffic and mobility simulations. In *AAMAS*, 2004.

[2] D. Beymer, P. McLauchlan, B. Coifman, and J. Malik. A real-time computer vision system for measuring traffic parameters. In *Conference on Computer Vision and Pattern Recognition*, pages 495–501, 1997.

[3] H. Buxton and R. J. Howarth. Behavioural descriptions from image sequences. In *Workshop on Integration of Natural and Vision Processing Language*. AAAI, 1994.

[4] C. F. Daganzo. Estimation of gap acceptance parameters within and across the population from direct roadside observation. *Transportation Research Part B*, 15B:1–15, 1981.

[5] C. F. Daganzo. A behavioural theory of multi-lane traffic flow. Part I and II. *Transportation Research Part B*, 36:131–158,159–169, 2002.

[6] J. H. Fernyhough, A. G. Cohn, and D. Hogg. Constructing qualitative event models automatically from video input. *Image and Vision Computing*, 18:81–103, 2000.

[7] S. P. Hoogendoorn and P. H. L. Bovy. Continuum modeling of multiclass traffic flow. *Transportation rsearch Part B*, 34:123–146, 2000.

[8] J. Kim, C. Lee, K. Lee, T. Yun, and H. Kim. Wavelet-based vehicle tracking for automatic traffic surveillance. In *Proceedings of IEEE Region 10 International Conference on Electrical and Electronic Technology*, volume 1, pages 313–316. IEEE, 2001.

[9] D. Parker and S. Stradling. Influencing driver attitudes and behaviour. Technical report, DBRG, 2001.

[10] H. O. Science Policy Unit. The police science and technology strategy: 2004-2009. Technical report, Home Office, UK, 2004.

[11] R. Vanhamme. Telindus works with uk highways agency to bring the M25 CCTV system up to speed. Technical report, 2002.