IDEA PROJECT FINAL REPORT
Contract ITS-9

IDEA Program
Transportation Research Board
National Research Council

October 2, 1995

DECISION-THEORETIC REASONING FOR TRAFFIC MONITORING AND VEHICLE CONTROL

Michael P. Wellman, University of Michigan
and Stuart J. Russell,
University of California at Berkeley

The ITS-IDEA program is jointly funded by the U.S. Department of Transportation's Federal Highway Administration, National Highway Traffic Safety Administration, and Federal Railroad Administration. For information on the IDEA Program contact Dr. K Thirumalai, IDEA Program Manager, Transportation Research Board, 2101 Constitution Avenue N.W., Washington, DC 20418 (phone 202-334-3568, fax 202-334-3471).

INNOVATIONS DESERVING EXPLORATORY ANALYSIS (IDEA) PROGRAMS MANAGED BY THE TRANSPORTATION RESEARCH BOARD (TRB)

This investigation was completed as part of the ITS-IDEA Program which is one of three IDEA programs managed by the Transportation Research Board (TRB) to foster innovations in surface transportation. It focuses on products and result for the development and deployment of intelligent transportation systems (ITS), in support of the U.S. Department of Transportation's national ITS program plan. The other two IDEA programs areas are Transit-IDEA, which focuses on products and results for transit practice in support of the Transit Cooperative Research Program (TCRP), and NCHRP-IDEA, which focuses on products and results for highway construction, operation, and maintenance in support of the National Cooperative Highway Research Program (NCHRP). The three IDEA program areas are integrated to achieve the development and testing of nontraditional and innovative concepts, methods and technologies, including conversion technologies from the defense, aerospace, computer, and communication sectors that are new to highway, transit, intelligent, and intermodal surface transportation systems.

1. Executive Summary

This report describes results from a joint Michigan/Berkeley project to develop technology for robust traffic monitoring and automated vehicle control using decision theory and probability. We have shown that high-level traffic monitoring situations can be modeled using modern techniques, and that solving such models in real time is computationally feasible.

Specifically, we have focused on models supporting the task of *pIan recognition in* highway *environments*. These models are used to infer the intended behavior of vehicles in traffic based on movement patterns and highway tactics (e.g., lane changes). In addition to their use in traffic monitoring (e.g., for gathering statistics of driver behavior), plan recognition models are indispensable for in-vehicle applications of several sorts---for any task in which it is crucial to anticipate movements of other vehicles.

Our models have been developed and encoded using off-the-shelf software, in particular the Hugin™ system. Using existing algorithms and some enhancements developed as part of this project, all of our models can be solved without prohibitive computational resources. For example, real-time performance is easily achievable using hardware based on Pentium-class processors. We have demonstrated this capability using Hugin™ in a standalone setup as well as in a driving simulator integrated with the SmartPATH animation system for real-time visualization of traffic scenarios.

1.1. Overview of Results

In the course of this project, we have designed and implemented a range of probabilistic models for *sifuafion assessment* in highway environments. Situation assessment can be applied toward such tasks as traffic management and design, emergency response, near-accident detection for intersection safety analysis, and intelligent traffic signals. Decision-theoretic models for situation assessment can also play a role in the intelligent control of vehicles. This project explored several of these areas, focusing primarily on general traffic interpretation through monitoring.

Our main conclusion is that current probabilistic reasoning methods based on Bayesian-network technology 'is indeed capable of supporting a range of traffic interpretation tasks, including those based on in-vehicle and extra-vehicle observation. Because the methods employ probabilistic reasoning, they are robust to sensor error and other forms of uncertainty. Successful deployment of this technology will depend on decreasing the cost of sensor technology (both video processing and in-vehicle sensor use).

These conclusions are based on our experience over the past year in constructing, testing, and integrating a variety of Bayesian-network models for a range of traffic-interpretation tasks. These tasks included both low-level interpretation (e.g., lane-change detection), and high-level situation assessment (e.g., driver plan recognition). Multilevel models can integrate these functions into a comprehensive traffic-interpretation system.

Developing these models required a few technical enhancements to the currently available modeling methods and probabilistic reasoning algorithms. Our most significant technical advances were in the following areas (described in more detail below):

- Methods for modifying the Bayesian network over time to support dynamic probabilistic reasoning (Forbes et al., 1995).
- A framework for plan recognition that properly accounts for the context in which the observed agent generated the plan, in addition to the planned activity itself (Pynadath and Wellman, 1995).
- New anytime approximation methods for networks with continuous variables, based on stochastic simulation (Kanazawa et al., 1995) and state-space abstraction (Liu and Wellman, 1995).

1.2. Specific Developments

We have developed models and algorithms for deriving high-level descriptions of traffic conditions, as well as the maneuvers and intentions of individual vehicles, from visual observation of a traffic scene. Our probabilistic network models represent both individual vehicles (their position, velocity, etc.) and aggregate variables concerned with the interaction of vehicles (flow, travel time, etc.). These models fuse information from a variety of sensors, which may be noisy and error-prone. The models have been successfully integrated with real-time visual processing on actual stationary highway monitoring video footage to identify and track individual vehicles accurately.

This framework has been extended to handle reasoning over time, by adding model elements to represent state variables at progressive time slices. At the next level of abstraction, our plan recognition framework can capture the beliefs and intentions of individual drivers and represent the dependence of their actions. This dependence encapsulates the driver's decision-making process in choosing a maneuver, based on its goals and information state. With the resulting probabilistic model, we use partial observations (e.g., lane changes, signals) of a vehicle to efficiently infer the driver's plan (e.g., passing, exiting) and project this plan to predict future actions.

In addition to demonstrating traffic monitoring on stationary highway video, we have demonstrated the use of similar integrated sensing and probabilistic network reasoning in dynamic vehicle traffic simulation. Through integration with a simple decision-tree based decision-making system, we have demonstrated autonomous intelligent driving. Our driving system is able to maintain a sensible estimate of the current traffic situation from simulated sensor inputs and negotiate a variety of challenging traffic contingencies. Our simulator is integrated with the SmartPATH animation system for real-time visualization of the traffic scenarios.

Complex probabilistic networks often require prohibitive computational resources for practical real-time traffic monitoring. We have designed, implemented, and demonstrated new approximation algorithms for probabilistic network inference especially well-suited to continual state updating and predicting, as often required for traffic monitoring and control. These algorithms produce the most accurate predictions possible within the time available for inference. As more computation time is allocated, the algorithms become increasingly accurate. In particular, we have implemented and tested algorithms based on (1) stochastic simulation and (2) abstraction of state spaces on versions of the models presented above.

2. Problem Statement

The purpose of this project was to investigate the feasibility of applying real-time, decision-theoretic reasoning technology to the intelligent monitoring of urban and freeway traffic and control of automated vehicles. The project built on recent, successful work on visual processing of real and simulated traffic images that is capable of identifying and tracking individual vehicles accurately. We aimed to show that such inputs can be processed further using artificial intelligence techniques, including probabilistic networks, to provide high-level descriptions of traffic conditions and individual vehicle maneuvers and intentions.

The problem of traffic interpretation is to derive a high-level description of what is going on in a traffic situation based on low-level information about vehicle movements. Low-level information includes vehicle tracks generated from visual monitoring or road-based or vehicle-centered sensors. It can also include turn signals and other indicators of drivers' intentions-information that is typically ignored in traffic monitoring schemes. High-level descriptions include summaries of flow levels, driver actions, and patterns of vehicle configurations. These descriptions can potentially be applied for congestion and flow analysis, near-accident detection (for safety analysis and/or directing emergency response), signal control, and development of high-fidelity driver models. Note that these applications include both on-line traffic control functions and off-line analysis tasks.

In particular, the automation of near-accident detection could provide crucial data for allocating safety resources, without requiring that we accumulate data on actual accidents (with the enormous social costs this entails). Visual monitoring coupled with intelligent interpretation can detect near-accidents, giving reliable measures of likely actual accidents within days rather than months or years.

To achieve our overall objectives for the project, our specific goals included:

- to develop a suite of probabilistic models for traffic situation assessment,
- to integrate these models with our existing software for highway driving simulation, assessment of high-level visual image inputs, and probabilistic model evaluation,
- to enhance available probabilistic inference technology to enable successful processing of these models,
- to test these models on simulated highway situations, and
- to use the results of this experience to evaluate the overall feasibility of probabilistic reasoning methods for traffic interpretation.

3. Research Approach

Currently available facilities for traffic assessment focus on relatively narrow contexts and low-level sensor fusion. We believe that by deriving high-level assessments of traffic situations and driver behavior, it can be possible to recognize and even anticipate significant traffic events, thereby improving safety and incident response. However, traffic monitoring at this level of abstraction is quite difficult, largely due to the inherent uncertainty in driver behavior, traffic dynamics, and interpretation of intended traffic movements. Standard stochastic methods from control theory cannot easily capture the structural uncertainty in alternate hypotheses about traffic situations.

Recent advances in probabilistic modeling technology by Artificial Intelligence researchers have led to significant improvements in the flexibility of specifications of probabilistic knowledge. Specifically, formalisms based on *Bayesian networks* support the representation of arbitrary patterns of probabilistic interdependence, and algorithms for exploiting the structure of relationships in the model. This standard probabilistic network framework has been extended to handle reasoning over time, by adding nodes to represent state variables at progressive time slices. This project is the first to apply such dynamic probabilistic networks to problems in traffic monitoring.

Our survey of the current state-of-the-art led us to conclude that existing methods for traffic monitoring and situation assessment will prove too inflexible for deployment in uncontrolled and dynamically changing environments. The generality and normative status of the decision-theoretic approach, along with recent computational advances in decision-theoretic reasoning, suggested that the time was ripe for its application to traffic monitoring.

The following enumerates some of the technical issues faced in the project. Specific accomplishments in tackling these issues are presented in the Results section.

3.1. Vehicle-Centered Models

In order to make appropriate control decisions, a driver must have accurate information about its own vehicle state and the state of its surrounding traffic environment. For example, a vehicle controller must know its own position, velocity, and intentions, and it must monitor those of its neighboring vehicles. It must also monitor road and weather conditions, since they may significantly affect driving ability.

The state of a vehicle's environment is only partially observable. Sensor information for variables such as vehicle positions and velocities may be incomplete and noisy, while driver intentions and road conditions may not be directly measurable at all. Thus, a controller cannot make decisions based merely upon the latest sensor readings. Rather, it must maintain estimates for the random variables that together represent the state of the world, and it must make its decisions based upon the joint probability distribution over all those variables.

3.2. Highway Models

In addition to models of individual vehicles, it is also important to model aggregate traffic variables, such as flow on road segments, travel time across various locations, etc. Such a model must accommodate information observable from in-vehicle sensors, in-highway sensors, road monitoring sites, and other sources of traffic data.

3.3. Dynamic Model Construction

Because traffic patterns evolve unpredictably, it is impractical to generate a probabilistic model in advance that fits the actual situation encountered during monitoring. Therefore, we must implement automated facilities to generate a model structure on-line, exploiting the specific features of interest for the given situation. In this project, we apply recent techniques for *knowledge-based model* construction to produce customized Bayesian networks that address the special features of traffic patterns observed dynamically by the visual input system.

Moreover, the networks themselves must be *dynamic,* that is, they must capture the evolution of uncertain state variables over time. We employ *dynamic Bayesian networks,* and *rollup* techniques to update the uncertain belief state at each time stage.

3.4. Maneuver Recognition

Since driver actions are normally limited to an enumerable set of maneuvers (e.g., lane changes, passing, exiting), it is reasonable to categorize driver actions at this high level. Recognizing the high-level maneuver being carried out can help in assessing an overall situation and in predicting the future behavior of drivers. To support maneuver recognition (or more generally, driver *plan recognition),* we require a probabilistic model relating the maneuver of a single car to observable features, from which we can categorize maneuvers and predict future behavior given partial information.

3.5. Anytime Approximation

In some traffic monitoring contexts, solving Bayesian networks at high fidelity may not be feasible or necessary. An alternate strategy is to employ simplification and approximation methods to produce the most accurate predictions possible given the available time for inference. As more computation time is allocated, the algorithms become increasingly accurate. This approach is called *anytime inference* in AI. The flexibility of anytime inference is especially important when random variables are continuous, as exact solutions can require an unbounded amount of computation.

4. Results

In the course of this project we have developed several probabilistic models for traffic interpretation tasks. The models (developed at both Berkeley and Michigan) have been integrated with the Hugin™ probabilistic reasoning tool and the SmartPATH traffic simulator. Both systems are installed and running at both project sites.

4.1. Vehicle-Centered Models

To maintain the vehicle controller's belief state, we employ *dynamic Bayesian networks* (DBNs). DBNs are an extension of Bayesian networks that allow variables to take on different values over time. Figure 1 shows the general structure of a DBN. Typically, observations are taken at regular "time slices", and a given network structure is replicated for each slice. DBNs model their domains as partially observable Markov processes, so nodes can be connected not only to other nodes within the same time slice but also (and only) to nodes in the immediately vehicle's representation of the world conforms to this property, we need not maintain the history of percepts to predict the next state since the accumulated effect of its observations is captured in the current belief state.

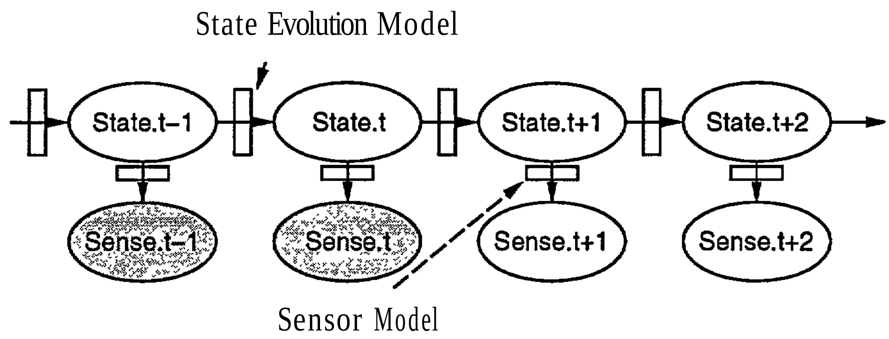

Figure 1: The structure of a dynamic probabilistic network. The ovals denote sets of state nodes or sensor nodes. The arcs going from one slice to the next form the state evolution model, and the arcs going into the sensor nodes form the sensor model. The shaded ovals denote observations available when predicting the state at time t+l.

As implemented, the system monitors each vehicle tracked by the sensor system with a separate DBN. Each network contains nodes for sensor observations, such as vehicle position and velocity, as well as nodes for predicting driver intentions, such as whether the driver intends to make a lane change or to slow down.

Like a Kalman filter, each network computes probability distributions for a vehicle's position and velocity based on both its latest observations and its previous state estimate (which reflects the influence of all previously observed evidence). Unlike a Kalman filter, which is limited to Gaussian distributions, the network predictions can be arbitrarily distributed. For example, if a vehicle were approaching some debris directly in front of it, the network could predict that the vehicle would move either to the right or to the left (but not straight) in order to avoid the debris. Also, the network could easily incorporate additional sensor information. If the sensor system recognized that a vehicle was flashing its right turn signal, the network could make predictions that biased the vehicle's position towards the right.

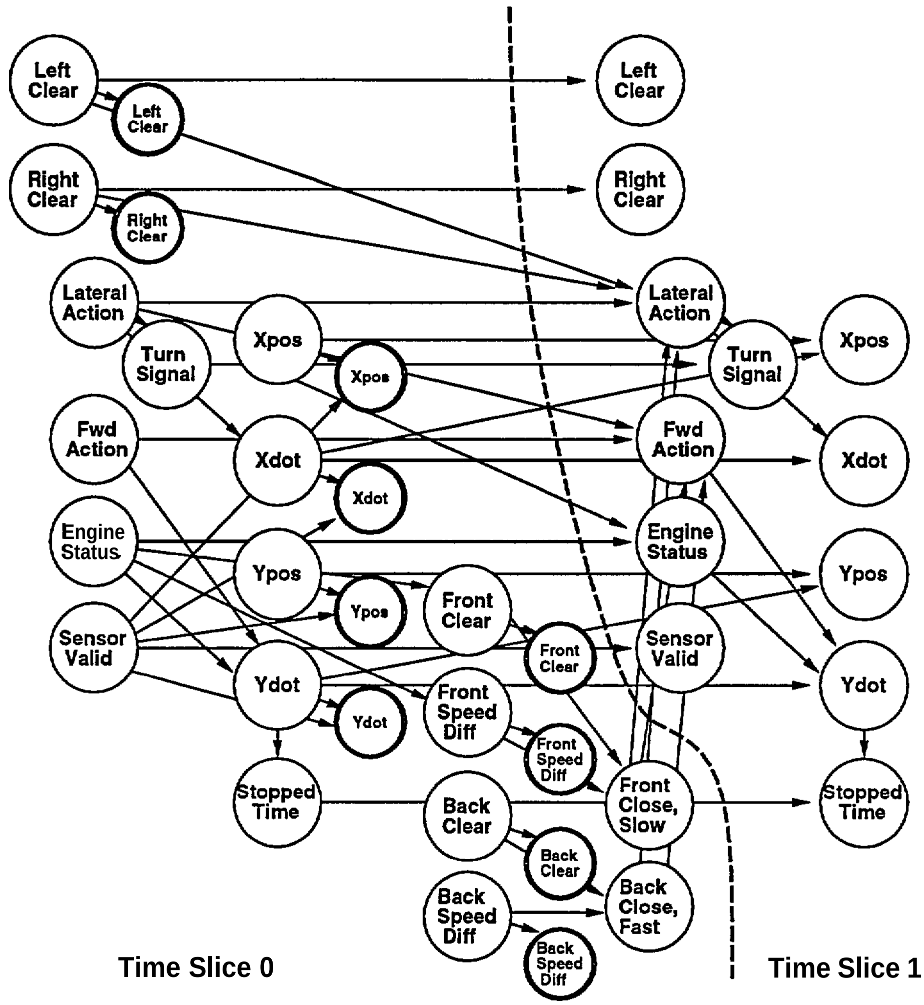

Figure 2: Dynamic Bayesian network for one vehicle, including inter-slice arcs. The smaller nodes with thicker outlines denote sensor observations.

To incorporate the influence of nearby vehicles, each network contains nodes corresponding to those vehicles. For example, the **Front Clear** and **Front Speed Diff** nodes in Figure 2 refer to "the space between this vehicle and the vehicle in front", and "the speed difference between this vehicle and the vehicle in front", respectively. Since the vehicle in front of or behind a given vehicle may change, *these indexical* nodes do not correspond to a specific vehicle. Instead, a preprocessing step using sensor data determines the spatial relationships among the vehicles and then sets the node states accordingly. Figure 2 shows an example vehicle network for one time slice, along with the inter-slice links to the next time slice.

4.2. Highway Models

Starting from standard deterministic models of traffic flow from the literature, we have developed Bayesian networks that capture the inherent uncertainty in congestion and the dynamics of traffic interactions. A schematic of the highestfidelity model we have explored is presented in Figure 3. We have also developed simpler, more tractable versions of this model that retain the two-dimensional space-time structure but avoid the complex interdependencies that make this model difficult to solve.

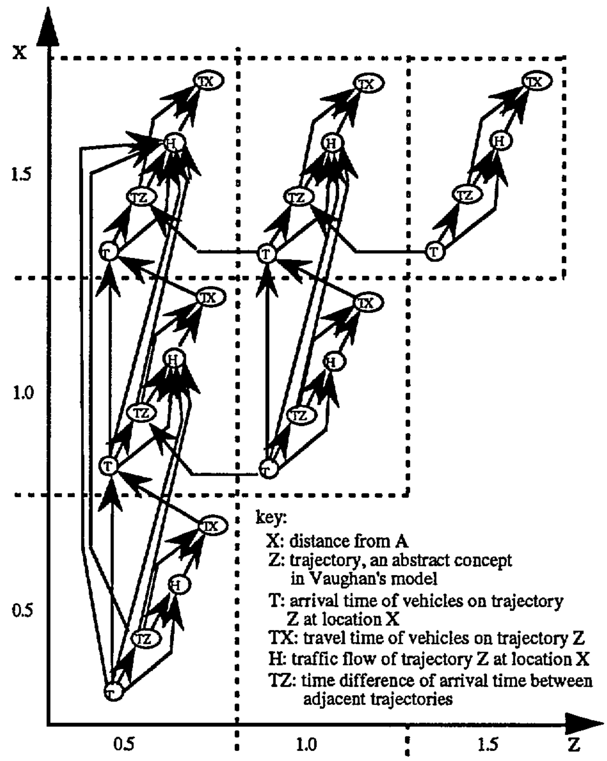

Figure 3: High-fidelity traffic model.

4.3. Maneuver Recognition

As an illustration of our plan recognition model, consider a driver on the highway, trying to predict the actions of the other drivers. Since these actions are normally limited to a small set of maneuvers (e.g., lane changes, passing, exiting), recognition of a driver's maneuvering *plan* would greatly assist in the prediction of future actions. To this end, we have worked on a probabilistic model of the maneuvers of a single car. We can then use this model to identify the current maneuver of an observed car and/or predict future actions, given only partial information.

We begin with a description of the plan execution model, since the assumptions made here motivated the choices made in the other portions of the model. We can classify the maneuvers according to the lane changes involved. The simplest plan is to simply continue driving in the same lane. At the next level of complexity, a driver can shift one lane to the left or right. We consider entering and exiting the highway as specific instances of these one-lane shifts. The driver could also shift two lanes to the left or right, where this could again involve entering or exiting the highway. As a final option, the driver may choose a passing maneuver, which we view as two successive lane shifts of opposite direction. In our Bayesian network, the variable **gen maneuver** takes on a value corresponding to the chosen plan.

We can also classify driving plans according to the acceleration. A driver may decide to speed up, slow down, or maintain current speed, depending on the current speed and the driver's desired traveling speed. The variable **acc maneuver** can take on one of the three values. The acceleration maneuver depends on the lane maneuver if we do not consider the plan selection mechanism. For instance, a deceleration is more likely with a right lane change plan than with a plan to pass.

The variable **spec pass** represents more specific passing plans. If the driver decides to pass, there are the options of passing on the left and passing on the right. And even if the driver chooses to pass, there may be cars blocking both lanes, forcing the driver to wait for another opportunity to pass. This variable clearly depends on **gen maneuver,** since if a passing maneuver is not chosen, then **spec pass** will be neither pass on left nor pass on right.

At this point we must model the initial world state. The current position and speed of the car are important factors in the driver's decision-making procedure, and we assume that both are observable, to the driver as well as to us. We also assume perfect sensors, but an extension to incorporate sensor noise is straightforward, as we have shown in some of our other models.

The random variables **x position** and **y position** represent the car's lane position and distance from the highway's start, respectively. The driver can be in one of three lanes or may be off the highway, either preparing to enter or having just exited. The y **speed** node, containing the car's speed, initially depends on the current node, since the farther left the lane, the faster its cars are usually traveling.

We can also observe the presence of other cars around the driver of interest, who must consider them in choosing a maneuver. For instance, if there is a car blocking the driver's front, then a passing maneuver is more likely. We can observe any cars to the driver's immediate front, back, left, and right, as well as in the four diagonal directions. In the Bayesian network, the Boolean random variable dir **clr?** represents the presence of any car immediately next to the driver in the *dir* direction.

The model of agent formulation in this case is greatly simplified by the assumption of perfect sensors, since the driver's beliefs about the world corresponds to the actual values in our simplified model. The only other factor in the driver's decision-making procedure is the driver's mental state. In most cases, the driver has the explicit goal of getting from one exit to another, although we may not always which specific exits. The random variable **exit position** represents the driver's desired exit. All of the possible exit positions are farther along the highway than the values of **y position.** **If** this were not the case, then the current position would provide evidence that the desired exit is probably not one that has been passed. Therefore, there would be a dependency, but to simplify the network, we make the sets of *y* and exit positions disjoint.

Also, there may be some constraint on the travel time between these exits, or the driver might have some target speed which is preferred for the duration of travel. The random variable **target y speed** represents this preference, with its values clustered around the speed limit. If the car has been on the highway for enough time, then its current speed should provide some clue as to the driver's target speed. We could model this with a link from *y* **speed.** On the other hand, if we have been observing the car and its maneuvers for some time, then these past observations should provide more conclusive evidence as to its target speed. If this is the case, we can make the target speed independent of current speed and encode our past observations in the prior probabilities. Our subnetwork in Figure 4 makes this assumption.

This network also contains the intermediate belief random variables, **at exit?** and **at driver?.** These reflect the driver's belief about the proximity of the desired exit and the desirability of the current speed, respectively. The **at exit?** variable depends only on the current position and the preferred exit, and is true only when the former is immediately before the latter. The **at target?** variable depends only on the current and preferred speeds, and its value indicates whether the current speed is too slow, too fast, or just right, with respect to the target.

The agent's beliefs about its capabilities are not represented explicitly. Instead, the driver is assumed to know all of the possible plans. The planning process also assumes that the driver has complete knowledge of how the plans can best satisfy its preferences in the current context. Thus the plan selection mechanism implicitly represents the driver's beliefs about its capabilities.

We can now model plan selection with some reliability. In our Bayesian network, the conditional probability table must specify the likelihood of certain maneuvers under every possible combination of world situation and driver mental state. Under most situations, there will be one maneuver that is clearly preferable. For example, suppose that the driver is currently traveling below its target speed and that there is another car directly in front while the lane to the left is clear. Then it is likely that driver will pass the car on the left. The complete plan selection subnetwork is shown in Figure 4. This model is closely related to the vehicle-centered model, especially in its effort to model a driver's decision-making procedure.

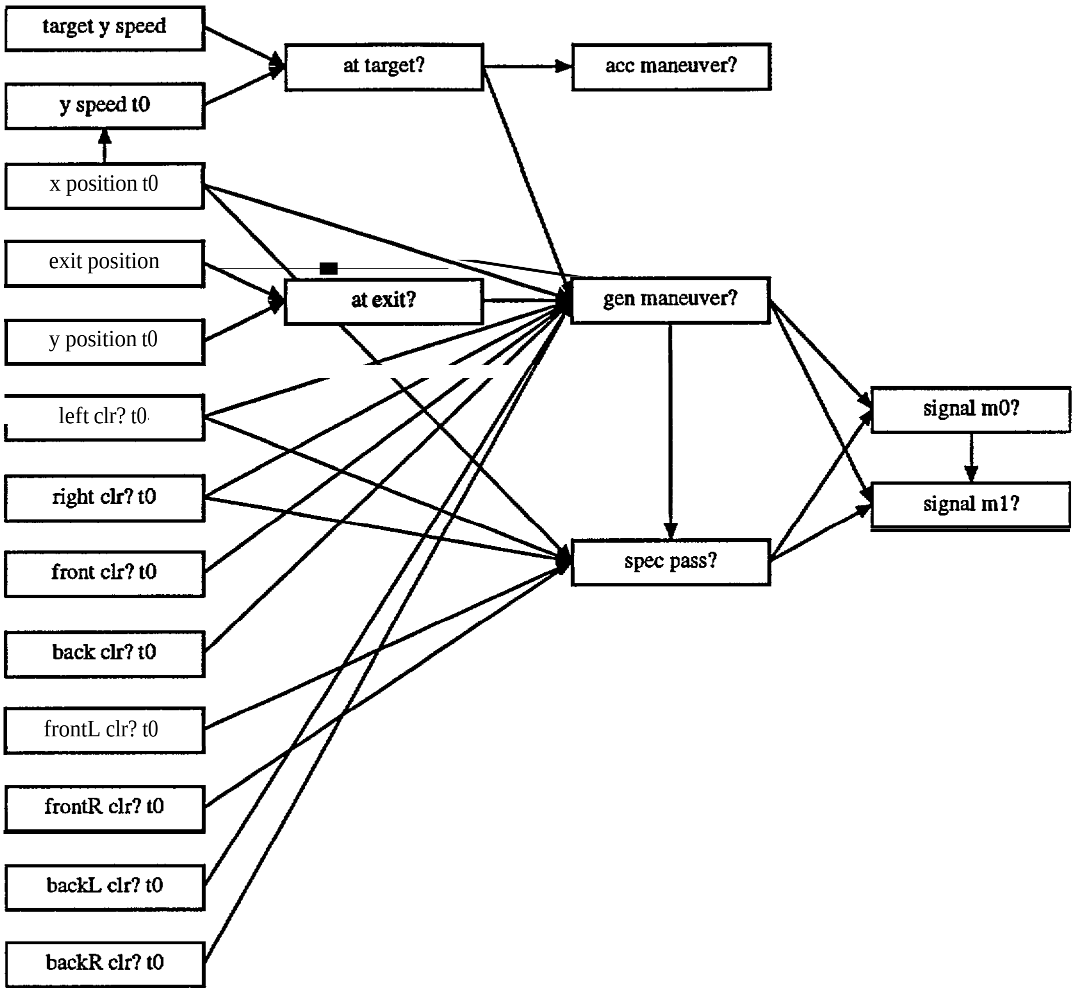

Figure 4: Planning process subnetwork.

The acceleration maneuver depends only on the desirability of the current speed. **Thus** the sole **link** to **acc maneuver** is from **at target?.** If the driver is at the target speed, then the current speed will be maintained. If the current speed is too low, then the driver will choose an acceleration maneuver. Likewise, if the current speed is too fast, then a deceleration maneuver will be chosen.

The lane change maneuver also depends on the desirability of the current speed. For instance, a car traveling at its target speed is unlikely to change lanes. However, there are other factors in the initial world state to consider. Obviously, the current lane is important, since a car in the leftmost lane can not change lanes to the left. In addition, the driver will consider any cars to the front or back. If there is a car blocking the front and the driver's current speed is too low, then a simple acceleration could cause a collision. The driver may instead choose to change lanes to the left. But a decision to change lanes must also consider the presence of cars to the driver's left or right, or any cars coming up from the back left or right. The links to the **gen maneuver** node represent these dependencies.

If the driver decides to pass, a direction must be chosen. Passing on the left is preferable to passing on the right, but the current situation may not allow it. For instance, any cars to the driver's left or to the front left could block the passing attempt. The same is true on the right side. If enough passing avenues are blocked, then the driver may decide to delay the passing attempt or to perform the initial lane change and wait to complete the pass.

There is no separate model of agent communication, but the extension is straightforward. For instance, the car's turn signal provides a simple mechanism for a driver to announce the intended lane change. Currently, any communication from the driver can only be modeled by directly instantiating any variables of the mental state which are announced.

Since there is no observed activity, most of our inference will come from observed effects. We must now model the dynamics of the traffic world, beginning with the changes in the position and speed of the car. We can view the actions of the driver to be transitions between world states. To simplify the model, we ignore observations taking place while the driver is performing an action. Thus, evidence is available only at the completion of a component action, and there are three stages of observable variables, including the context, as can be seen in Figure 5.

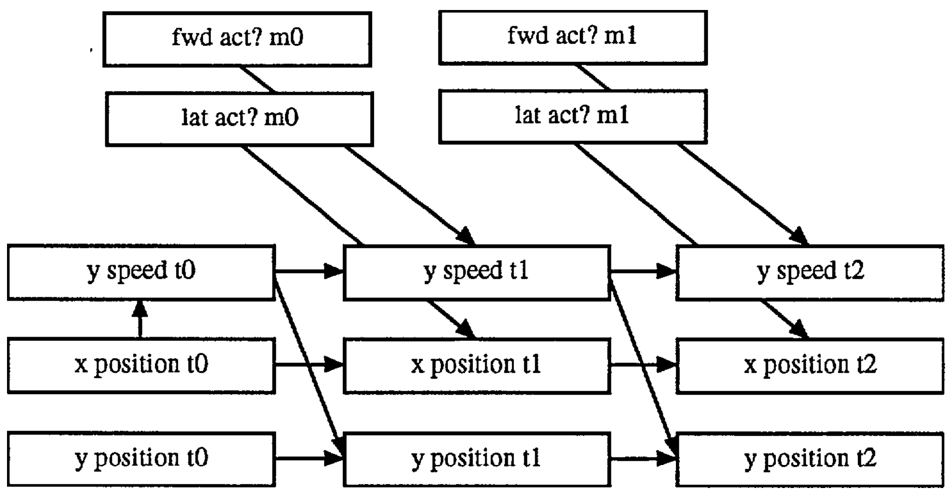

Figure 5 : Evidence subnetwork.

The evidence subnetwork includes the individual transitions in lane and speed, which are completely unobservable. At each step, the driver can change one lane to the left or right, or remain in the same lane. The driver can also increase, decrease, or maintain speed. All of the plans we consider produce a two-step action sequence. For instance, a plan to shift one lane to the left produces a left lane change followed by a "remain in lane" act. The **lat accm**x variables represent the lane changes at step x, while fwd **accm**x represents the acceleration at step x.

Our definition of the lane maneuvers completely determines the lane changes of the action sequences. The individual shifts depend on the general lane maneuver, as well as on the specific passing plan, but not on the acceleration maneuver. Likewise, the individual accelerations are independent of the general lane changes and the specific passing maneuvers if given the overall acceleration plan.

Finally, we must define the dependencies of these effects. Most of the observable variables depend on the driver's previous action, as well as their own previous values. For instance, the driver's lane is completely determined if we know what lane change just took place, as well as the lane value just before the change. Likewise, the driver's speed depends on the previous speed and whatever acceleration action took place, although this is clearly not a deterministic relationship.

The presence of other cars is a bit more complex, due to the driver's movements. For instance, after a left change, a car that was to the front and left is now probably directly in front. But if the driver stays in the same lane, then we must check whether there was a car blocking the front in the previous world state. Therefore, each clearance variable depends on the previous action, as well as all relevant clearance variables from the previous state. To simplify the network, we ignore the presence of other cars in the evidence. We do consider them when modeling plan selection, but since the driver's actions do not directly affect the other drivers' positions, we ignore these effects. As with the context, we assume perfect sensors, so there is no distinction between the actual and observed effects.

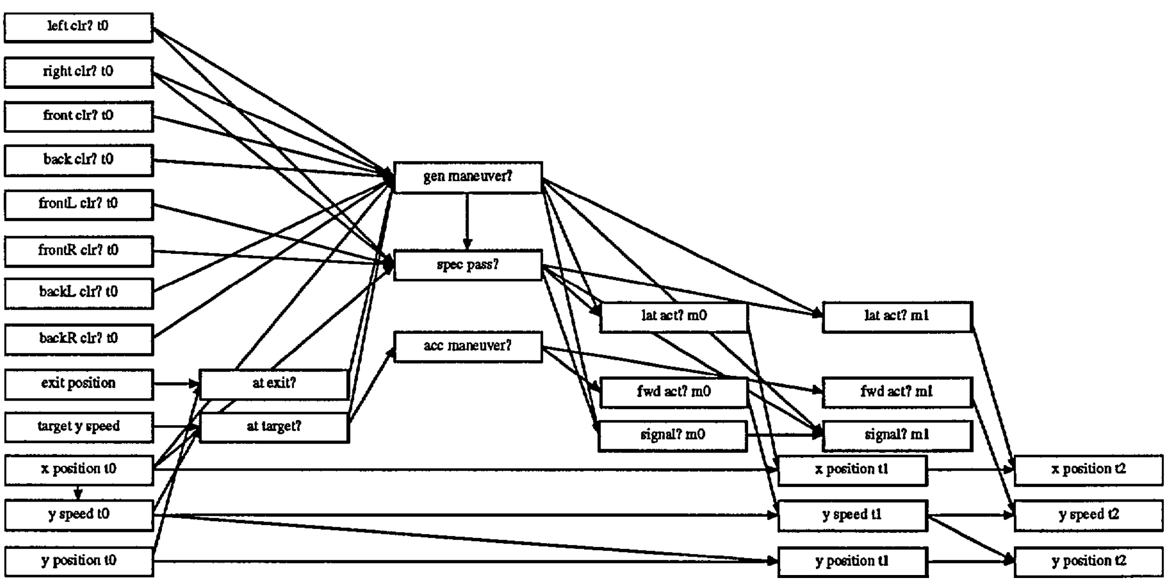

Figure 6 : Complete Bayesian network for maneuver recognition.

Once we have constructed the entire network, shown in Figure 6, we can handle plan recognition in a wide range of useful driving situations. For instance, suppose we are trying to predict the behavior of the car behind us as we are driving in the middle lane of a three-lane highway. We observe the car move into the rightmost lane, and we want to determine if it is passing us, or preparing to exit, or perhaps simply moving into the slower-moving lane.

Thus, in the context, we have observed **front clr? t0** to be false and **x position t0** to be the middle lane. The only observed effect is that **x position t1** is the right lane. If we want to infer the driver's plan, we can examine the **gen maneuver?** node to see that the posterior probability of a one-lane right shift is 0.64, while that of a pass is 0.35. The former is more plausible since we assume that drivers prefer to pass on the left-hand side, so passing on the right has a relatively low prior probability. There is a small probability (<0.01) that the car is about to exit, since we have assumed no knowledge about the location of exits beyond the prior probabilities.

If we are not interested in the driver's plan, but only in the future lane position, then we **can** examine the **x position t2** node. The posterior probability that the car will still be in the right lane is 0.65, while the probability that it will move to the middle lane is 0.34. The difference between these beliefs and that of the maneuvers arises from the nature of the passing maneuver. Even if the car decides to pass, it may not be able to do so immediately do to surrounding cars. In such a case, it will remain in its current lane until it can complete the maneuver. Thus, there is a slight probability that the car will stay in the right lane even if the driver has decided to pass.

Given no other contextual observations, it is reasonable to predict that the car will remain in the right lane. However, if we also observed that there was another car to our left, thus blocking the car behind us from passing on the left, we can instantiate the **frontL clr? t0** variable to be false. Repeating our observation of the nodes of interest, we find that the posterior probability that the car is passing has increased to 0.53, while that for the car simply shifting one lane to the right has dropped to 0.46. The probabilities for **x position t2** have changed as well, to 0.51 and 0.48 respectively.

Thus, we are able to perform valuable inference with only a limited subset of the possible observations. If we were to also observe that there were no other cars nearby, other than those already considered, then we could instantiate the remaining clearance context variables to be true. Doing so increases the posterior probability that the maneuver is a pass to 0.61, while decreasing that for a one-lane right shift to 0.39.

4.4. Anytime Approximation

4.4.1. State-Space Abstraction

In practical uncertain reasoning situations, computational resource limitations often preclude exact solution of probabilistic models. A common strategy in such cases is to fall back on inexact, *approximate* solutions, which trade accuracy and fidelity for computational time and/or space. The most typical approximation approaches involve some form of *abstraction*-suppression or neglect of detail for the purpose of simplification.

Abstraction and other approximation techniques are of particular interest for inference in Bayesian networks. That computing exact conditional probabilities in Bayesian networks is NP-hard suggests that ultimately, only approximate solutions will be accessible. Although it has been shown that approximating a conditional probability to a fixed degree of accuracy is also NP-hard, this does not seriously damage the case for approximation. First, even without guarantees of fixed **degrees** of accuracy, approximate methods offer reasonable prospects of significant accuracy, which is a lot better than many alternatives. Second, approximation offers us the opportunity to consider problems much larger than we could otherwise, which may compensate substantially for a loss of accuracy. And third, approximation methods often lend themselves well to *anytime algorithms*, where the quality of solution grows smoothly with available computation time, thus supporting robust performance over a range of time-stressed situations.

Approximation by evaluating abstracted Bayesian networks has attracted considerable attention in recent years. A Bayesian network may be abstracted by neglecting dependencies or state distinctions between or within variables. In our approach, we abstract the network by coarsening the granularity of the state spaces of selected nodes. The granularity is then refined in successive iterations as the computation progresses. Evaluation is terminated when the solution becomes accurate enough, or the allotted computation time is exhausted.

procedure Abstract-Iter(OBN, evidence)

1. Generate an initial Abstract Bayesian Network with one superstate per abstracted node.

2. Evaluate the probability distribution for each node given the evidence.

3. If all states for all nodes are elementary, return.

4. Split a superstate in an abstracted node.

5. Go to step 2.

Figure 7: The iterative abstraction procedure.

4.4.2. Sampling

A second anytime approach we have developed in the course of this project is a new version of stochastic simulation designed particularly for DBNs (Kanazawa et al., 1995). Specifically, we have introduced two new techniques, *evidence* *reversal* and *survival of fittest sampling.* Downstream evidence (i.e., observed nodes with parents) can degrade performance of stochastic simulation. Evidence reversal transforms a given network so that all evidence nodes are at the root. This is a provably correct operation that yields significant improvement in performance, especially in the type of DBNs used in our vehicle-centered model. Survival of fittest sampling is based on the use of a fixed-size population of samples. After each decision epoch, we extend the sample population by one time slice. We randomize the repopulation process based on the likelihood of the evidence given a particular sample (thus this algorithm is related to genetic algorithms although there is no crossover in our approach).

5. Follow-on Work

The development of improved methods for traffic scene analysis has contributed significantly to the success of a recent proposal submitted by UC Berkeley (including S. Russell as co-PI) in response to the RFP for New Sensor Technology for traffic surveillance from JPL (Ql0-1066-025). The project builds directly on some of the Bayesian methods we have developed for tracking, and a follow-on study will use the incident analysis networks. The effort will develop and test a prototype vision-based surveillance system for wide-area deployment in freeway and urban settings in the next 12 months. If successful, the system could be deployed at upwards of 100,000 sites in the United States.

At Michigan, we are extending our research in Bayesian methods for pattern and plan recognition, under funding from the Air Force Office of Scientific Research. We will enhance our traffic maneuver model as necessary to test out new techniques.

6. Conclusions

Our main findings to date confirm the feasibility of real-time probabilistic reasoning in support of various traffic assessment applications. We have built a variety of Bayesian network models, which have provided robust probabilistic estimates of traffic states given partial information and limited inference time.

Previous to this project, there had been little or no overlap between research in IVHS and decision-theoretic reasoning. One important contribution of this project was to bring probabilistic and decision-theoretic reasoning methods to bear on traffic sensing and control problems. We hope that our reports describing this work to the respective research communities will increase cross-fertilization between the fields.

Although we believe that the techniques demonstrated in this project will ultimately prove valuable in collision-avoidance systems, it appears that the sensors necessary to support vehicle-centered probabilistic reasoning will not be widely available in the near term. In contrast, the requirements for visual monitoring of traffic situations from a fixed location appear to be relatively reasonable. Traffic monitoring stations using probabilistic reasoning techniques can be deployed to collect data on traffic maneuvers, driving behavior, and incidents, for use in safety analysis, traffic engineering, long-term planning, and improvement of driver models.

7. Bibliography

Forbes, J., T. Huang, K. Kanazawa, et al. (1995). The BATmobile: Towards a Bayesian automated taxi. *Fourteenth International Joint Conference on Artificial Intelligence* Montreal.

Kanazawa, K., D. Koller, and S. Russell (1995). Stochastic simulation algorithms for dynamic probabilistic networks. *Eleventh Conference on Uncertainty in Artificial Intelligence,* Montreal.

Liu, C.-L., and M. I?. Wellman (1995). State-space abstraction for anytime evaluation of Bayesian networks. *IJCAI-95 Workshop on Anytime Algorithms and Deliberation Scheduling,* Montreal.

Pynadath, D. V., and M. P. Wellman (1995). Accounting for context in plan recognition, with application to traffic monitoring. *Eleventh Conference on Uncertainty in Artificial Intelligence,* Montreal.

Wellman, M. I?., C.-L. Liu, D. Pynadath, S. Russell, J. Forbes, T. Huang, and K. Kanazawa (1995). Decision-theoretic reasoning for traffic monitoring and vehicle control. *Intelligent Vehicles '95,* Detroit.