
Christophe Coué
Cédric Pradalier
Christian Laugier
Thierry Fraichard
Pierre Bessière

INRIA
Rhône-Alpes and Gravier-CNRS
<http://www.inrialpes.fr/sharp>
cedric.pradalier@inrialpes.fr

Bayesian Occupancy Filtering for Multitarget Tracking: An Automotive Application

Abstract

Reliable and efficient perception and reasoning in dynamic and densely cluttered environments are still major challenges for driver assistance systems. Most of today's systems use target tracking algorithms based on object models. They work quite well in simple environments such as freeways, where few potential obstacles have to be considered. However, these approaches usually fail in more complex environments featuring a large variety of potential obstacles, as is usually the case in urban driving situations. In this paper, we propose a new approach for robust perception and risk assessment in highly dynamic environments. This approach is called Bayesian occupancy filtering; it basically combines a four-dimensional occupancy grid representation of the obstacle state space with Bayesian filtering techniques.

KEY WORDS—multitarget tracking, Bayesian state estimation, occupancy grid

1. Introduction

1.1. Advanced Driver Assistance Systems Context

Unlike regular cruise control systems, adaptive cruise control (ACC) systems use a range sensor to regulate the speed of the car while ensuring collision avoidance with the vehicle in front. ACC systems were introduced on the automotive market in 1999. Since then, surveys and experimental assessments have demonstrated the interest for this type of system. They are the first step towards the design of future advanced driver assistance systems (ADASs) that should help the driver in increasingly complex driving tasks. The use of today's commercially available ACC systems is generally limited to mo-

torways or urban expressways without crossings. The traffic situations encountered are rather simple and attention can be focused on a few, well-defined detected objects (cars and trucks). None the less, even in these relatively simple situations, these systems display a number of limitations. They are not very good at handling fixed obstacles and because of this may generate false alarms. Moreover, in “cut-in” situations, i.e., when the intrusion of another vehicle or a pedestrian in the detection beam is too close to the vehicle, they may be unable to react appropriately.

A wider use of such systems requires an extension of their range of operation to more complex situations in dense traffic environments, around or inside urban areas. In such areas, traffic is characterized by lower speeds, tight curves, traffic signs, crossings, and “fragile” traffic participants such as motorbikes, bicycles, or pedestrians.

1.2. Problem

Prerequisites to a reliable ADAS in such complex traffic situations are the following.

- Robust and accurate sensing of the environment. In particular, dynamic characteristics of the traffic participants, such as position and velocity, have to be correctly estimated.
- Appropriate sensing representation that allows both a good understanding of the traffic situation and selection of the most appropriate driving decisions.

1.3. Related Work

1.3.1. Multitarget Tracking

The estimation of the dynamic characteristics of the traffic participants is basically a multitarget tracking problem. The objective is to collect observations, i.e., sensor data, on one or

more potential obstacles in the environment of the vehicle, and then to estimate at each time step (and as robustly as possible) the obstacles positions and velocities. The classical approach is to track the different objects independently, by maintaining a list of tracks, i.e., a list of currently known objects. The main difficulty of multitarget tracking is known as the data association problem. It includes observation-to-track association and track management problems. The goal of observation-to-track association is to decide whether a new sensor observation corresponds to an existing track or not. Then track management includes deciding whether existing tracks should be maintained or deleted, and whether new tracks should be created. Numerous methods exist to perform data association (Bar-Shalom and Li 1995; Streit and Luginbuhl 1995; Gauvrit, Le Cadre, and Jauffret 1997). The reader is referred to Blackman and Popoli (2000) for a complete review of the existing tracking methods with one or more sensors.

Urban traffic scenarios are still a challenge in the multitarget tracking area: the traditional data association problem is intractable in situations involving numerous appearances, disappearances and occlusions of a large number of rapidly maneuvering targets.

In Wang, Thorpes, and Thrun (2003), a classical multiple hypothesis tracking technique is used to track moving objects while stationary objects are used for simultaneous localization and mapping (SLAM). Unfortunately, the authors did not explicitly address the problem of the interaction between tracked and stationary objects, e.g., when a pedestrian is temporarily hidden by a parked car. It is one of the purposes of our approach to solve this problem.

1.3.2. Grid Representation of the Environment

The occupancy grid framework (Moravec 1988; Elfes 1989) is a classical way to describe the environment of a mobile robot. It has been extensively used for static indoor mapping using a two-dimensional grid (Thrun 2002). The goal is to compute from the sensor observations the probability that each cell is full or empty. To avoid a combinatorial explosion of grid configuration, the cell states are estimated as independent random variables.

More recently, occupancy grids have been adapted to track multiple moving objects (Prassler, Scholz, and Elfes 2000). In this approach, spatio-temporal clustering applied to temporal maps is used to perform motion detection and tracking. A major drawback of this work, relatively to the ADAS context, is that a moving object may be lost due to occlusion effects.

1.4. Contribution

The objective of this paper is to propose a new approach for a robust perception, representation and analysis of highly dynamic environments. Four main motivations were taken into account in the design of this approach, as follows.

- Taking explicitly into account the uncertainty (which is inherently present in any model of a real phenomenon) when estimating the state of the environment.
- Avoiding the “data association problem”, which usually fails to solve the complex scenarios we would like to address, i.e., scenarios involving numerous appearances, disappearances and occlusions of several rapidly maneuvering targets.
- Increasing the robustness of the system relatively to object occlusions, appearances and disappearances, by exploiting at any time all the relevant information on the vehicle environment; this information includes the description of the occupied areas, of the unoccupied areas, and of the hidden areas (i.e., areas of the environment that are temporarily hidden from sensors by an obstacle).
- Designing a method that could be implemented later on a dedicated hardware, in order to both reach high performances and decrease the costs of the final system.

Our approach is based on a probabilistic grid representation of the obstacle state space. This approach allows us to meet the four previous objectives.

- Uncertainty is explicitly taken into account thanks to the probabilistic reasoning paradigm, which is becoming a key paradigm in robotics: various approaches based on this paradigm have already been successfully used to address several robotic problems, such as CAD modeling (Mekhnacha, Mazer, and Bessière 2001) or map building and localization (SLAM; Thrun 1998; Kaelbling, Littman, and Cassandra 1998; Arras, Tomatis, and Siegwart 2001).
- The data association problem is avoided by reasoning on a probabilistic grid representation of the dynamic environment. In such a model, concepts such as objects or tracks do not explicitly exist; they are replaced by more useful properties such as occupancy or risk, which are directly estimated for each cell of the grid using both sensor observations and some prior knowledge. Furthermore, when estimating occupancy probability thanks to an adequate sensor model, the hidden parts of the environments can also be explicitly characterized. Since we consider both the positions and the velocities of the potential obstacles with respect to our vehicle, this grid is four-dimensional and is called the obstacle state space (OSS) grid .
- The dynamic nature of the environment and robustness to object occlusion is addressed using a novel two-step mechanism allowing us to take into account the sensor observation history and the temporal consistency of

the scene. This mechanism estimates, at each time step, the state of the occupancy grid by combining a prediction step (history) and an estimation step (new measurements). This approach is derived from the Bayes filter approach (Jazwinsky 1970); our filter is called the Bayesian occupancy filter (BOF).

- The BOF has been designed in order to be highly parallelizable. So, a hardware implementation on a dedicated chip is possible, which will lead to a very efficient way to represent the environment of an automotive vehicle. This problem (i.e., SoC for the *BOF*) is currently addressed within the scope of the Safemove France–Korea project.

1.5. Outline of the Paper

In the next section we present the estimation of the occupancy grid in a static case, i.e., taking into account only the last sensor observation. In Section 3 we describe the BOF itself, i.e., how we take into account the sensor observation history. Experimental results are shown in Sections 2 and 3. In Section 4 we describe a collision avoidance application based on the BOF. This approach has been implemented on the CyCab, an electric vehicle (see <http://www.robosoft.fr>) equipped with a Sick laser range finder, allowing the system to estimate target position and velocity (Figure 1).

2. Static Estimation of the Occupancy Probability

In this section, we introduce first the Bayesian formalism (called Bayesian programming) which has been used for developing our approach. Then, we present the Bayesian program, which has been developed for estimating the state of the occupancy grid of the vehicle environment, by using only the latest sensor observations (i.e., static estimation). Finally, experimental results are presented and discussed.



Fig. 1. The CyCab experimental vehicle. It is equipped with a Sick laser range finder.

2.1. Bayesian Formalism and Related Computational Tools

The Bayesian programming framework has initially been developed by our research team, for designing robust robot control programs (Lebeltel 1999; Coué and Bessière 2001; Lebeltel et al. 2004). Today, it is used for addressing various problems involving uncertain or incomplete knowledge (e.g., Mekhnacha, Mazer, and Bessière 2001; Le Hy et al. 2003). This framework is based on a well-defined mathematical theory, and it provides both formal and computational tools for designing applications in a systematic way.

From the formal point of view, a Bayesian program is made up of two main parts: a description and a question.

- The description. This part of the Bayesian program can be viewed as a knowledge base containing the a priori information available on the problem at hand; it mainly represents a joint probability distribution. This description is made up of three components:
 1. a set of relevant variables on which the joint distribution is defined, e.g., motor, sensory, or internal state variables;
 2. a decomposition of the joint distribution as a product of simpler terms, which is obtained by applying Bayesian rules and by taking advantages of the conditional independences that may exist between the variables;
 3. the parametric forms assigned to each of the terms appearing in the decomposition (they are required to compute the joint distribution).
- The probabilistic questions. Given a distribution, it is possible to ask probabilistic questions. Basically, a question can be expressed by first partitioning the set of variables into three subsets (S, K, F) (representing respectively the “searched variables”, the “known variables”, and the “free variables”), and by writing a probabilistic expression of the type:

$$P(S | K)? \quad (1)$$

Given the previous description, it is always possible to answer such a question, i.e., to compute the probability distribution $P(S | K)$. This can be done using the following general inference:

$$\begin{aligned} P(S | K) &= \frac{\sum_F P(S F K)}{P(K)} \\ &= \frac{1}{\alpha} \times \sum_F P(S F K), \end{aligned} \quad (2)$$

where α is a normalization term.

From the computational point of view, it is well known that a brute force approach cannot be applied in practice for solving the previous inference problem (Bayesian inference has

been shown to be NP-hard; Cooper 1990), and that a previous symbolic simplification phase can drastically reduce the number of sums necessary to compute a given distribution. This is why we have developed an Application Programming Interface (API), called ProBT[®], which allows us to easily express Bayesian programs and to perform in an efficient way the related symbolic and numerical operations. This engine operates in two complementary stages:

- a symbolic simplification stage, which reduces the complexity of the probability distribution to be computed;
- a numerical stage that actually computes the related distributions.

This engine is now commercialized by the Probayes company (see <http://www.probayes.com>). The approach described in this paper has been implemented using this framework.

2.2. Bayesian Program for Static Estimation

Our goal is to estimate the occupancy probability of each cell of the grid, using the last set of sensor observations. These observations represent preprocessed information given by a sensor. At each time step, the sensor is able to return a list of detected objects, along with their associated positions and velocities in the sensor reference frame. In practice, this set of observations could also contain two types of false measurements: the false alarms (i.e., when the sensor detects a non-existing object) and the missed detection (i.e., when the sensor does not detect an existing object).

Then, solving the previous static estimation problem can be done by building the following Bayesian program.

1. Choosing the relevant variables and decomposition.

- C : the cell itself; this variable is four-dimensional and represents a position and a speed relative to the vehicle.
- E_C : the state of the cell C , occupied or not.
- \mathcal{Z} : the sensor observation set; one observation is denoted Z_s ; the number of observations is denoted S ; each variable Z_s is four-dimensional.
- M : the “matching” variable; its goal is to specify which observation of the sensor is currently used to estimate the state of a cell.

Then, the following decomposition of the joint distribution determined by these variables can be chosen:

$$P(C E_C \mathcal{Z} M) = \left(\begin{array}{l} P(C)P(E_C | C)P(M) \\ \times \prod_{s=1}^S P(Z_s | C E_C M) \end{array} \right). \quad (3)$$

2. Assigning the parametric forms. According to our knowledge of the problem to be solved, we can assign the following parametric forms to each of the terms of the previous decomposition:

- $P(C)$ represents the information on the cell itself. As we always know the cell for which we are currently estimating the state, this distribution does not need to be specified.
- $P(E_C | C)$ represents the a priori information on the occupancy of the cell. If available, a prior distribution could be used to specify it. Otherwise, a uniform distribution has to be selected. In the next section we explain how the prior distribution may be obtained from passed estimation.
- $P(M)$ is chosen uniform.
- The shape of $P(Z_s | C E_C M)$ depends on the value the matching variable.
 - If $M \neq s$, the observation is not due to the cell C . Consequently, we cannot say anything on this observation. $P(Z_s | C E_C M)$ is defined by a uniform distribution.
 - If $M = s$, the form of $P(Z_s | C E_C M)$ is given by the sensor model. Its goal is to model the sensor response knowing the cell state. Details on this model can be found in Elfes (1989).

3. Solution of the problem. It is now possible to ask the Bayesian question corresponding to the searched solution.¹ Since the problem to solve consists in finding a good estimate of the cell occupancy, the question can be stated as follows:

$$P(E_C | \mathcal{Z} C)? \quad (4)$$

Following the general inference mechanism given by eq. (2), the result of the inference can be written as follows:

$$P(E_C | \mathcal{Z} C) \propto \sum_{M=1}^S \left(\prod_{s=1}^S P(Z_s | E_C C M) \right). \quad (5)$$

The result of this inference is computed by our inference engine. During this inference, the sum on the previous variables allows us to consider every sensor observation when updating the state of one cell. It should be noticed that the estimation step is performed without any explicit association between cells and observations; this problematic operation is replaced by the integration of all the possible values of M .

1. The answer to this question is given by our inference engine. It represents the searched probability distribution.




Fig. 2. Example of static grid estimation. (a) The situation and the sensor observations, given in (x, y, \dot{x}, \dot{y}) . (b), (c) Values of occupancy probability for two “slices” of the four-dimensional grid, i.e., for all possible positions at a given speed.

2.3. Experimental Result

We have tested the previous approach for the static estimation of an occupancy grid, using both a simulator and a real vehicle equipped with a laser sensor (the CyCab).

Figure 2 shows some resulting grid estimations. Figure 2(a) depicts the situation: the sensor mounted on the CyCab is located at $(x = 0, y = 0)$; the part of the environment covered by the grid is represented by the light gray rectangle; the sensor field of view is modeled by the dark gray area. In this situation, three obstacles (black disks) are present in front of the CyCab; two of these are stationary objects, and the third is moving from the left to the right, at a speed of 0.8 m s^{-1} represented by a black arrow. In this first experiment, the CyCab is not moving.

As mentioned earlier, we use a four-dimensional grid. Thus, only two-dimensional “slices” of the grid are depicted by Figures 2(b) and (c). In both figures, the occupancy probability value is given by the gray level of the cell. We can see the correspondence between probability value and gray level on the right side of the figure.

Figure 2(b) depicts the occupancy probability of each cell corresponding to a null relative velocity (i.e., $c = [x, y, 0, 0]$), which is the speed associated with the two previous sensor observations. As expected, two areas with high occupancy probabilities are visible. These probability values depend on several factors attached to the sensor model: the probability of true detection, the probability of false alarm, and the sensor accuracy. Since the measured speed for the third obstacle is not null, any area of high occupancy probability corresponding to this observation is only represented in the related slices of the grid (i.e., the slice corresponding to $c = [x, y, 0, 0.8]$ in this case; see Figure 2(c)). It should be noticed that the cells located outside the sensor field of view, or the cells hidden by one of the three sensor observations (i.e., the cells located

behind the three detected obstacles) cannot be observed; consequently, nothing really consistent can be said about these cells, and the system has given an occupancy probability value of 0.5 for these cells (red areas). Finally, all the cells located in the observable area and not related to any sensor observation are associated with a low occupancy probability value (purple areas).

3. Bayesian Occupancy Filter

3.1. Problem Addressed and Approach

We are now interested in taking into account the sensor observation history, in order to be able to make more robust estimations in changing environments (i.e., in order to be able to process temporary objects occlusions and detection problems). Our approach for solving this problem is to make use of an appropriate Bayesian filtering technique called the BOF.

Bayes filters (Jazwinsky 1970) address the general problem of estimating the state sequence x^k , $k \in \mathbb{N}$ of a system given by

$$x^k = f^k(x^{k-1}, u^{k-1}, w^k), \quad (6)$$

where f^k is a possibly nonlinear transition function, u^{k-1} is a “control” variable (e.g., speed or acceleration) for the sensor which allows to estimate its ego-movement between time $k - 1$ and time k , and w^k is the process noise. This equation describes a Markov process of order one.

Let z^k be the sensor observation of the system at time k . The objective of the filtering is to recursively estimate x^k from the sensor measurements:

$$z^k = h^k(x^k, v^k). \quad (7)$$

Here, h^k is a possibly nonlinear function and v^k is the measurement noise. This function models the uncertainty of the measurement z^k of the system's state x^k .

In other words, the goal of the filtering is to recursively estimate the probability distribution $P(X^k | Z^k)$, known as the posterior distribution. In general, this estimation is done in two stages: prediction and estimation. The goal of the prediction stage is to compute an a priori estimate of the target's state known as the prior distribution. The goal of the estimation stage is to compute the posterior distribution, using this a priori estimate and the current measurement of the sensor.

Exact solutions to this recursive propagation of the posterior density do exist in a restrictive set of cases. In particular, the Kalman filter (Kalman 1960; see <http://www.cs.unc.edu/~welch/kalman/index.html>) is an optimal solution when the functions f^k and h^k are linear and the noises w^k and v^k are Gaussian. However, in general, solutions cannot be determined analytically, and an approximate solution has to be computed.

In our case, the state of the system is given by the occupancy state of each cell of the grid, and the required conditions for being able to apply an exact solution such as the Kalman filter are not always verified. Moreover, the particular structure of the model (occupancy grid) and the real-time constraint coming from the ADAS application have led us to develop the new concept of the BOF. This filter consists in estimating the occupancy state in two steps, as depicted in Figure 3.

3.2. Bayesian Occupancy Filter Estimation Step

In this loop, the estimation step is similar to the static estimation of the grid depicted in the previous section, except that the a priori on the cell occupancy $P(E_C^k | C^k)$ is not given by a uniform distribution, but by the result of the previous prediction step.

3.3. Bayesian Occupancy Filter Prediction Step

The goal of this processing step is to estimate an a priori model of the occupancy probability at time k of a cell using the latest estimation of the occupancy grid, i.e., the estimation at time $k - 1$. The variables that are relevant here are the following:




Fig. 3. BOF as a recursive loop.

- C^k : cell C considered at time k ;
- E_C^k : state of cell C at time k ;
- C^{k-1} : cell C at time $k - 1$;
- E_C^{k-1} : state of cell C at time $k - 1$;
- U^{k-1} : “control” input of the CyCab at time $k - 1$, e.g., it could be a measurement of its instantaneous velocity at time $k - 1$.

Using these variables, we can define the following decomposition of the joint distribution:

$$P(C_k E_C^k C^{k-1} E_C^{k-1} U^{k-1}) = \left(\begin{array}{l} P(U^{k-1})P(C^{k-1}) \\ \times P(E_C^{k-1} | C^{k-1}) \\ \times P(C^k | C^{k-1} U^{k-1}) \\ \times P(E_C^k | E_C^{k-1} C^{k-1} C^k) \end{array} \right). \quad (8)$$

Then, we can assign the following parametric to each of the previous decomposition terms.

- $P(C^{k-1})$ and $P(U^{k-1})$ are chosen as uniform distributions.
- $P(E_C^{k-1} | C^{k-1})$ is given by the result of the estimation step at time $k - 1$.
- $P(C^k | C^{k-1} U^{k-1})$ is given by the dynamic model. It represents the probability that an object has moved from the cell C^{k-1} to the cell C^k . This movement is due to the object itself and to the CyCab movement between times $k - 1$ and k . To define this model, we suppose a constant velocity model subject to zero mean Gaussian errors for the moving objects.
- $P(E_C^k | E_C^{k-1} C^{k-1} C^k)$ represents the probability that an existing object at time $k - 1$ (i.e., $[E_C^{k-1} = 1]$) still exists at time k (i.e., $[E_C^k = 1]$). As we consider that objects cannot disappear, Dirac are chosen for these distributions.

The problem to be solved is to find an estimate of the occupancy probability for each cell of the grid. Solving this problem can be done by asking the following question to our inference engine:

$$P(E_C^k | C^k U^{k-1})? \quad (9)$$

The result of the inference can be expressed as follows:

$$P(E_C^k | C^k U^{k-1}) \propto \sum_{\substack{C^{k-1} \\ E_C^{k-1}}} \left(P(C^k | C^{k-1} U^{k-1}) \times P(E_C^{k-1} | C^{k-1}) \right). \quad (10)$$

Unfortunately, most of the time this expression cannot be expressed analytically, and it cannot be computed in real time. This is why an approximate solution of the integral term has to be computed. Our approach for making this computation is based on the fact that only a few points are needed to approximate the integral. Thus, for each cell of the grid at time $k - 1$, we can compute the probability distribution $P(C^k | C^{k-1})$; then, a cell c^k is drawn according to this probability distribution; finally, the cell C^{k-1} is used to update the predicted state of the cell c^k . It should be noticed that the complexity of the previous algorithm increases linearly with the number of cells in our grid, and ensures that the most informative points are used to compute the sum appearing in eq. (10).

Using the BOF approach, the estimation of the occupancy grid at time k is carried out in two steps. (1) The prediction step makes use of both the result of the estimation step at time $k - 1$ and a dynamical model, for computing an a priori estimate of the grid. (2) Then, the estimation step makes use of both this prediction result and the sensor observations at time k , to compute the grid.

3.4. Experimental Results

Figure 4 shows a short sequence of successive prediction and estimation results, for a dynamic scene involving two moving obstacles. The objective of this example is to experimentally demonstrate the robustness of our approach to objects occlusions. The first row describes the experimental conditions: the CyCab is immobile, and its sensors can observe two moving objects $O1$ and $O2$ ($O1$ is moving from left to right, and $O2$ is moving from right to left). In the situation depicted by Figure 4(c1), $O1$ is temporary hidden by $O2$ (and thus $O1$ is not detected by the Sick laser range finder).

The second and third rows show the results of the prediction and estimation steps, respectively. We have chosen to only represent the cells of the grid corresponding to a relative speed equal to $(\dot{x} = 0.0 \text{ m s}^{-1}, \dot{y} = 1.0 \text{ m s}^{-1})$, which is close to the speed of $O1$. The occupancy probabilities of the related cells are represented by several shades.

In this example, an area of “high occupancy probability”, which corresponds to the moving objects, is well characterized in Figures 4(a2) and (a3). We can also see that the areas hidden by the moving objects have an occupancy probability value equal to 0.5. Similar results can be seen in Figures 4(b2) and (b3). Figure 4(c2) shows the result of the prediction step, based on the grid of Figure 4(b3) and on the used dynamic model; this prediction shows that an object is probably located in the area hidden by $O2$ (i.e., an area of high occupancy probability is found in Figure 4c3). Of course, confidence in the presence of a hidden object (i.e., the values of the occupancy probability in the grid) progressively decreases when this object is not observed by the sensor during the next time steps. In the example depicted by Figure 4(d3), the object is no longer hidden by $O2$; it is detected by the laser, and the related occupancy probability values increase.

Figure 5 shows another sequence of successive prediction and estimation results. The first row describes the experimental conditions: the CyCab is moving forward at a constant speed equal to 2.0 m s^{-1} ; a static object is present in the sensor field of view of the CyCab. Since the CyCab is moving forward, this object finally goes out of the sensor field of view, as depicted in Figure 5(d1). The second and third rows show the results of the prediction and estimation steps, respectively. We have chosen to only represent the cells of the grid corresponding to a relative speed equal to $(\dot{x} = -2.0 \text{ m s}^{-1}, \dot{y} = 0.0 \text{ m s}^{-1})$, which is close to the relative object/CyCab. As in the case of the example of Figure 4, the prediction step allows us to infer knowledge on the current occupancy state of the CyCab environment, even if the object is no longer observed by the sensor. This is the situation depicted by Figure 5(d3), where an area of high occupancy probability still exists when the object is going out of the sensor field of view. In some sense, our prediction step can be seen as a “short-term memory”, which allows us to combine in an evolutive way past and current observations.

3.5. Performance

As mentioned earlier, thanks to our approximation algorithm, both the prediction step and the estimation step complexities increase linearly with the number of cells of the grid. This makes the approach tractable in real situations involving reasonable grid sizes. This is the case for the experimental examples described in this section and in the next section. Typically, the characteristics of the related grid models are the following:

- 0 to 10 m in the X direction, with a 0.5 m resolution;
- -5 to 5 m in the Y direction, with a 0.5 m resolution;
- -3 to 1 m s^{-1} in the \dot{X} direction, with a 0.4 m s^{-1} resolution;
- -3.2 to 3.2 m s^{-1} in the \dot{Y} direction, with a 0.4 m s^{-1} resolution.

Using such a grid of 64,000 cells, the computation time for both prediction and estimation steps is about 100 ms on a 1 GHz computer. This is fast enough to control the CyCab at a maximum speed of 2 m s^{-1} .

However, this grid size is not fine enough for an automotive application involving higher speeds. In this case, the number of cells increases quickly, and the required computational time becomes too high for satisfying the real-time constraints. For instance, doubling the size of the grid in all directions will result in a global size of 1,024,000 cells, and in a required computation time of about 2400 ms (using a single 1 GHz processor).

Hopefully, our BOF approach has been designed in order to be highly parallelizable: thanks to the hypothesis that each




Fig. 4. A short sequence of a dynamic scene. The first row describes the situation: a moving object is temporary hidden by a second object. The second row shows the predicted occupancy grids, and the third row shows the result of the estimation step. The grids show $P([E_C^k = 1] | x \ y [\dot{x} = 0.0] [\dot{y} = 1.0])$.

cell is independent, the state of each cell can be computed independently. Current work deals with the development of a dedicated hardware that will exploit this characteristic (and consequently allows a real-time application of the BOF in large-scale applications).

In this section we have described the way the system operates for continuously interpreting the state of the environment using the BOF. In the next section we explain how we have exploited this information for implementing safe basic behaviors on an autonomous robot (and more precisely on the CyCab).

4. Bayesian Occupancy Filter Based Collision Avoidance

The goal of this section is to show how the BOF can be used for developing a collision avoidance function on an autonomous vehicle. This function has been implemented and experimented on the CyCab experimental vehicle; the goal of the implemented function is to continuously select the forward speed values, in order to move safely (i.e., while avoiding moving obstacles such as pedestrians or other cars) along a given road lane.




Fig. 5. A short sequence of a dynamic scene. The CyCab is moving forward at a constant speed. The grids show $P(\{E_C^k = 1\} | x \ y \ [\dot{x} = -2.0] \ [\dot{y} = 0.0])$.

As mentioned earlier, the cell state can be used to encode some relevant properties of the robot environment (e.g., occupancy, observability, reachability, etc.). In the previous sections, only the occupancy characteristic was stored; in the case of the vehicle application, we also encode the danger property. This leads us to control the vehicle by combining occupancy and danger criteria.

4.1. Danger Estimation

For each cell of the grid, the probability that this cell is hazardous is estimated; this estimation is done independently of the occupancy probability property. Let $P(D_C^k | C^k)$ be the probability distribution associated with the cell C^k of the ve-

hicle environment, where D_X^k is a Boolean variable that indicates whether this cell is hazardous or not.

Basically, both “time to collision” and “safe traveling distance” may be seen as two complementary relevant criteria to be used for estimating the danger to associate to a given cell. In our current implementation, we are using the following related criteria which can easily be computed. (1) The closest point of approach (CPA), which defines the relative positions of the pair (vehicle, obstacle) corresponding to the “closest admissible distance” (i.e., safety distance). (2) The time to the closest point of approach (TCPA), which is the time required for reaching the CPA. (3) The distance at the closest point of approach (DCPA), which is the distance separating the vehicle and the obstacle when the CPA has been reached.

In some sense, these criteria give an assessment of the future relative trajectories of any pair of environment components of the type (vehicle, potential obstacle).

The previous criteria are evaluated for each cell, and at each time step k , by taking into account the dynamics characteristics of both the vehicle and the potential obstacles. In practice, both TCPA and DCPA are estimated under the hypothesis that the related velocities at time k remain constant. This computation can easily be one using some classical geometrical algorithms (see, for instance, <http://softsurfer.com/algorithms.htm>).

As previously mentioned, our goal is to estimate the “danger probability” to associate to each cell of the grid (or in other terms, the probability for each cell C^k that a collision will occur in the near future between the CyCab and a potential obstacle in C^k). Since each cell C^k represents a pair (position, velocity) defined relatively to the CyCab, it is easy to compute the TCPA and DCPA factors, and in a second step to estimate the associated danger probability using given intuitive user knowledge. In the current implementation, this knowledge roughly states that when the DCPA and the TCPA decrease, the related probability of collision increases. In a future version of the system, we could expect that such knowledge could be acquired using a learning phase.

Figure 6 shows the cells for which the danger probability is greater than 0.7 in our CyCab application. In the picture, each cell is represented by an arrow: the tail of the arrow indicates the position, and its length and direction indicate the associated relative speed. This figure exhibits quite intuitive data: any cell located in the vicinity of the front part of the CyCab are considered as having a high danger probability for any relative velocity (the arrows are pointing in all directions); the other cells having a high “oriented” danger probability are those having a relative speed vector oriented towards the CyCab. Since we only consider relative speeds for constructing the danger grid, the content of this grid does not depend on the actual CyCab velocity.

4.2. Collision Avoidance Behaviors

In this section we describe how we can control the longitudinal speed of the autonomous vehicle (the CyCab), for avoiding partially observed moving obstacles having a high probability of collision with the vehicle. The implemented behavior consists of braking or accelerating, in order to adapt the velocity of the vehicle to the level of risk estimated by the system.

As mentioned earlier, this behavior derives from the combination of two criteria defined on the grid: the danger probability associated with each cell C^k of the grid (characterized by the distribution $P(D_C^k | C^k)$; see Section 4.1), and the occupancy probability of this cell (characterized by the posterior distribution $P(E_C^k | Z^k C^k)$; see Section 3). In practice, we search, at each time step, for the most hazardous cell that is considered as probably occupied. This can be done using the




Fig. 6. Cells of high danger probabilities. For each position, arrows model the speed.

following equation:

$$\max_{C^k} \{P(D_C^k | C^k), \text{ with } P(E_C^k | C^k) > 0.5\}.$$

Then the longitudinal acceleration/deceleration to apply to the CyCab controller can be decided according to the obtained level of danger and to the actual velocity of the CyCab.

Figure 7 depicts the scenario used for experimentally validating the previous collision avoidance behavior on the CyCab. In this scenario, the CyCab is moving forward, the pedestrian is moving from right to left, and during a small period of time the pedestrian is temporarily hidden by a parked car.




Fig. 7. Scenario description: the pedestrian is temporary hidden by a parked car.



Fig. 8. Snapshots of the experimental pedestrian avoidance scenario (see Extension 1 for the video).




Fig. 9. Velocity of the CyCab during the experiment involving a pedestrian occlusion.

Figure 8 shows some snapshots of the experiment (see also Extension 1, which shows the entire video). The CyCab brakes to avoid the pedestrian, who is temporarily hidden by the parked car, and then it accelerates as soon as the pedestrian has crossed the road.

Figure 9 shows the velocity of the CyCab during this experiment. From $t = 0$ s to $t = 7$ s, the CyCab accelerates, up to 2 m s^{-1} . At $t = 7$ s, the pedestrian is detected; as a collision could possibly occur, the CyCab decelerates. From $t = 8.2$ s to $t = 9.4$ s, the pedestrian is hidden by the parked car; thanks to the BOF results, the hazardous cells of the grid are still considered as probably occupied and, as a consequence, the CyCab still brakes. When the pedestrian reappears at $t = 9.4$ s, there is no longer a risk of collision, and the CyCab can accelerate.

5. Conclusions

In this paper we have addressed the problem of designing a new approach for robust perception and danger assessment of highly dynamic environments. The proposed approach is called Bayesian occupancy filtering, which basically combines a four-dimensional occupancy grid representation of the obstacle state space with Bayesian filtering techniques. This approach can be seen as an alternative to complex multitarget tracking algorithms, which usually fail in situations involving

numerous appearances, disappearances and occlusions of a large number of rapidly maneuvering targets. It also brings a significant improvement to traditional approaches, by including a prediction step which allows us to make more robust estimation relatively to temporary occlusions.

This approach has experimentally been validated on our experimental vehicle (the CyCab), for avoiding partially observed moving obstacles. A scenario involving the CyCab, a moving pedestrian, and a parked car which temporarily hides the pedestrian from the sensors of the CyCab, has successfully been executed. In this experiment, the avoidance behavior has been obtained by combining the occupancy probability and the danger probability of each cell of the grid.

Current and future work deal with three major points.

1. Improvement of the approximation algorithm used for the prediction step; this improvement should allow to estimate bigger grids (which are required for dealing with complex urban traffic situations).
2. Development of a dedicated hardware that will exploit the parallelizable property of the BOF algorithm, for being able to meet the real-time constraint in large-scale applications.
3. Fusion of the occupancy grid with higher-level information, such as GPS maps, to better estimate the danger of the situation.

Appendix: Index to Multimedia Extensions

The multimedia extension page is found at <http://www.ijrr.org>.

Table of Multimedia Extensions

Extension	Type	Description
1	Video	Complete experimentation video for the scenario described in section IV.B. The CyCab detect and track a temporarily occluded obstacle using the BOF, while adjusting its speed to minimise collision risk.

Acknowledgments

This work was partially supported by the European project IST-1999-12224 "Sensing of Car Environment at Low Speed Driving" (<http://www.carsense.org>).

References

- Arras, K. O., Tomatis, N., and Siegwart, R. 2001. Multisensor on-the-fly localization: precision and reliability for applications. *Robotics and Autonomous Systems* 44:131–143.
- Bar-Shalom, Y. and Li, X. 1995. *Multitarget Multisensor Tracking: Principles and Techniques*, YBS Publishing, Storrs, CT.
- Blackman, S. and Popoli, R. 2000. *Design and Analysis of Modern Tracking Systems*, Artech House, Boston, MA.
- Cooper, G. 1990. The computational complexity of probabilistic inference using Bayesian belief network. *Artificial Intelligence* 42(2–3):393–405.
- Coué, C. and Bessière, P. 2001. Chasing an elusive target with a mobile robot. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Maui, HI, October.
- Elfes, A. 1989. Using occupancy grids for mobile robot perception and navigation. *IEEE Computer* 22(6):46–57.
- Gauvrit, H., Le Cadre, J. P., and Jauffret, C. 1997. A formulation of multitarget tracking as an incomplete data problem. *IEEE Transactions on Aerospace and Electronic Systems* 33(4):1242–1257.
- Jazwinsky, A. H. 1970. *Stochastic Processes and Filtering Theory*, Academic Press, New York.
- Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101:99–134.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Transactions of the ASME, Journal of Basic Engineering* 82(D):35–45.
- Lebeltel, O. 1999. *Programmation Bayésienne des Robots*. Thèse de doctorat, Institut National Polytechnique de Grenoble, Grenoble, France, September.
- Lebeltel, O., Bessière, P., Diard, J., and Mazer, E. 2004. Bayesian robot programming. *Autonomous Robots* 16:49–79.
- Le Hy, R., Arrigoni, A., Bessière, P., and Lebeltel, O. 2003. Teaching Bayesian behaviors to video game characters. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, NV, October.
- Mekhnacha, K., Mazer, E., and Bessière, P. 2001. The design and implementation of a Bayesian CAD modeler for robotic applications. *Advanced Robotics* 15(1):45–70.
- Moravec, H. P. 1988. Sensor fusion in certainty grids for mobile robots. *AI Magazine* 9(2):61–74.
- Prassler, E., Scholz, J., and Elfes, A. 2000. Tracking multiple moving objects for real-time robot navigation. *Autonomous Robots* 8(2):105–116.
- Streit, R. L. and Luginbuhl, T. E. 1995. Probabilistic multi-hypothesis tracking. Technical Report 10,428, Naval Undersea Warfare Center Division, Newport, RI.
- Thrun, S. 1998. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence* 99(1):21–71.
- Thrun, S. 2002. Robotic mapping: a survey. *Exploring Artificial Intelligence in the New Millennium*, Morgan Kaufmann, San Mateo, CA.
- Wang, C-C., Thorpes, C., and Thrun, S. 2003. On-line simultaneous localization and mapping with detection and tracking of moving objects: theory and results from a ground vehicle in crowded urban areas. *Proceedings of the IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, September, pp. 842–849.