

# Hierarchies of probabilistic models of space for mobile robots: the Bayesian Map and the Abstraction operator

Julien Diard\* and Pierre Bessière and Emmanuel Mazer

Laboratoire GRAVIR / IMAG – CNRS  
INRIA Rhône-Alpes, 655 avenue de l’Europe  
38330 Montbonnot Saint Martin FRANCE

## Abstract

This paper presents a new method for probabilistic modelling of space, called the Bayesian Map formalism. It offers a generalization of some common approaches found in the literature, as it does not constrain the dependency structure of the probabilistic model. The formalism allows incremental building of hierarchies of models, by the use of the Abstraction Operator. In the resulting hierarchy, localization in the high level model is based on probabilistic competition of the lower level models. Experimental results validate the concept, and hint at its usefulness for large scale scenarios.

## 1 Introduction and related work

In robotics, modelling the environment that a robot has to face in a navigation task is a crucial problem, that has received a lot of attention in the community. The most promising approaches rely on the probability calculus, especially for its capacity to handle incomplete models and uncertain information. These approaches include – but are far from limited to – Kalman Filters [Leonard *et al.*, 1992], Markov Localization models [Thrun, 2000], (partially or fully) observable Markov Decision Processes [Boutillier *et al.*, 1999; Kaelbling *et al.*, 1998], and Hidden Markov Models [Rabiner and Juang, 1993]. For references that present several of them at once, giving unifying pictures, see [Bessière *et al.*, 2003; Murphy, 2002; Roweis and Ghahramani, 1999; Smyth *et al.*, 1997]. We will here assume that the reader has some familiarity with these approaches.

In this domain of probabilistic modelling for robotics, hierarchical solutions are currently flourishing – while still representing a very small part of the literature. The more active domain in this regard is decision theoretic planning: one can find variants of MDPs that accommodate hierarchies or that select automatically the partition of the state-space (see for instance [Hauskrecht *et al.*, 1998; Lane and Kaelbling, 2001], or browse through the references in [Pineau and

Thrun, 2002]). More exceptionally, one can find hierarchical POMDPs, as in [Pineau and Thrun, 2002], which is arguably the work that bears the most resemblance to the one presented here, although we do not use reward functions in this work. The current work can also be related to Thrun’s object mapping paradigm [Thrun, 2002], in particular concerning the aim of transferring some of the knowledge the programmer has about the task, to the robot. Some hierarchical approaches outside of the MDP community include Hierarchical HMMs and their variants (see [Murphy, 2002] and references therein), which, unfortunately, rely on the notion of final state of the automata. Another class of approaches relies on the extraction of a graph from a probabilistic model, like for example a Markov Localization model [Thrun, 1998], or a MDP [Lane and Kaelbling, 2002]. Using such deterministic notions is inconvenient in a purely probabilistic approach, as we are pursuing here. Indeed, the current work uses probabilities in all layers of the hierarchy of representations, allowing us to propagate and handle uncertainties in a uniform and formally coherent manner.

Moreover, the main philosophy used by all the previous approaches is to try to extract, from a very complex but intractable model, a hierarchy of smaller models (structural decomposition, see [Pineau and Thrun, 2002]). Of course, *automatically* selecting the relevant decomposition of a problem into sub-problems is quite a challenge – this challenge being far from restricted to the domain of navigation for robots facing uncertainties.

We pursue here an alternate route, investigating how, starting from a set of simple models, one can combine them for building more complex models. The goal of this paper is therefore to present a new formalism for building models of the space in which a robot has to navigate (the *Bayesian Map* model), and a method for combining such maps together in a hierarchical manner (the *Abstraction operator*).

This formalism allows for a new representation of space, in which the final program is built upon many imbricated models, each of them deeply rooted into lower level sensorimotor relationships. Hierarchies of sensorimotor models seem also relevant to biologically inspired models, as it appears that no single metric model can account alone for large scale navigation capacities of animals (see [Kuipers, 2000; Trullier *et al.*, 1997; Franz and Mallot, 2000]).

We will also argue that our approach draws away from the

---

\*Julien Diard is currently with the Laboratoire de Physiologie de la Perception et de l’Action, Collège de France, Paris and the National University of Singapore. He can be reached at Julien.Diard@free.fr.

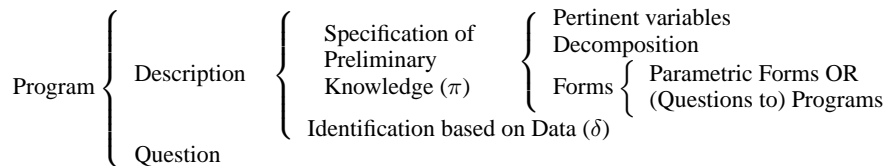


Figure 1: Structure of a Bayesian Robotic Program.

usual characteristics of the common models of space (Section 2.3), and that it is also more general than these models (Sections 3 and 6). For brevity, this paper will discuss neither of the learning methods that can be included into Bayesian Maps (mapping process), nor of another operator for merging Bayesian Maps (the Superposition operator). Preliminary work about these issues and all the details missing in the current paper can be found in Diard’s Ph.D. thesis [Diard, 2003].

The rest of this paper is organized as follows. Section 2 presents the basics of the Bayesian Robot Programming methodology, and discusses some of its characteristics. Sections 3 and 4 will quickly define our notion of Bayesian Map, and the Abstraction operator, respectively. Experimental results are presented Section 5, and the paper concludes on some research perspectives.

## 2 Bayesian Robot Programming

The work we present here is based on BRP, a Bayesian Robot Programming methodology. We briefly summarize it here, but still invite the interested reader to refer to [Lebeltel *et al.*, 2003]<sup>1</sup> for all the details about this methodology.

As this formalism is only based on the inference rules needed for probability calculus, it is very general. Indeed, a very wide class of probabilistic models found in the literature can be rewritten in the BRP framework, as is shown in great length in [Bessière *et al.*, 2003; Diard, 2003]<sup>2</sup>. For example, we can rewrite the Markov Localization model into the BRP formalism. We will use this particular example to illustrate both the BRP formalism itself and its generality.

### 2.1 Definition

In the BRP formalism, a bayesian robotic program is a structure (see Figure 1) made of two components.

The first is a *declarative* component, where the user defines a **description**. The purpose of a description is to specify a method to compute a joint distribution over a set of relevant variables  $\{X_1, X_2, \dots, X_n\}$ , given a set of experimental data  $\delta$  and preliminary knowledge  $\pi$ . This joint distribution is denoted  $P(X_1 X_2 \dots X_n | \delta \pi)$ . To specify this distribution, the programmer first lists the pertinent variables (and defines their domains), then decomposes the joint distribution as a

product of simpler terms (possibly stating conditional independence hypotheses so as to simplify the model and/or the computations), and finally, assigns forms to each term of the selected product (these forms can be parametric forms, or recursive questions to other bayesian programs). If there are free parameters in the parametric forms, they have to be assessed. They can be given by the programmer (*a priori* programming) or computed on the basis of a learning mechanism defined by the programmer and some experimental data  $\delta$ .

The second component is of a *procedural* nature, and consists of using the previously defined description with a **question**, *i.e.* computing a probability distribution of the form  $P(\text{Searched} | \text{Known})$ . Answering a “question” consists in deciding a value for the variable *Searched* according to  $P(\text{Searched} | \text{Known})$ . Different decision policies are possible, in our robotic experiments we usually choose to draw a value at random according to that distribution. It is well known that general Bayesian inference is a very difficult problem, which may be practically intractable. But, as this paper is mainly concerned with modelling issues, we will assume that the inference problems are solved and implemented in an efficient manner by the programmer<sup>3</sup>.

### 2.2 Example

As an example of a Bayesian Robot Program, we will rewrite the Markov Localization model in the BRP formalism. It can be seen as an extension of the Hidden Markov Model, where an action variable is added. Outside of the field of robotics, it is sometimes called input-output HMM [Bengio and Frasconi, 1995; Cacciatore and Nowlan, 1994]. Let us recall that a HMM is basically the decomposition  $P(O_t S_t S_{t-1}) = P(S_{t-1})P(S_t | S_{t-1})P(O_t | S_t)$ , where  $O_t$  is a perception variable,  $S_t$  and  $S_{t-1}$  are location variables at time  $t$  and  $t-1$ . Starting from this structure, the action variable  $A_t$  is used to refine the transition model  $P(S_t | S_{t-1})$  into  $P(S_t | A_t S_{t-1})$ . The resulting BRP model for Markov Localization is shown Figure 2.

### 2.3 BRP vs. other models

Let us now develop some remarks that arise from the comparison between the use of the BRP formalism and some aspects of the more common models of the representation of space (see Section 1). In particular, we now focus on solving navigation tasks using BRP programs.

The first remark relies on the fact that, in BRP, a form appearing in a description  $c^1$  can be a question to another

<sup>1</sup>[Lebeltel *et al.*, 2003] is currently in press. Please refer in the meantime to a preliminary version, at <http://www-laplace.imag.fr/publications/Rayons/Lebeltel2000.pdf>.

<sup>2</sup>[Bessière *et al.*, 2003] is downloadable from <http://www-laplace.imag.fr/publications/Rayons/RR-4730.pdf>.

<sup>3</sup>The inference engine we use to tackle these problems has been described elsewhere [Bessière *et al.*, 2003].

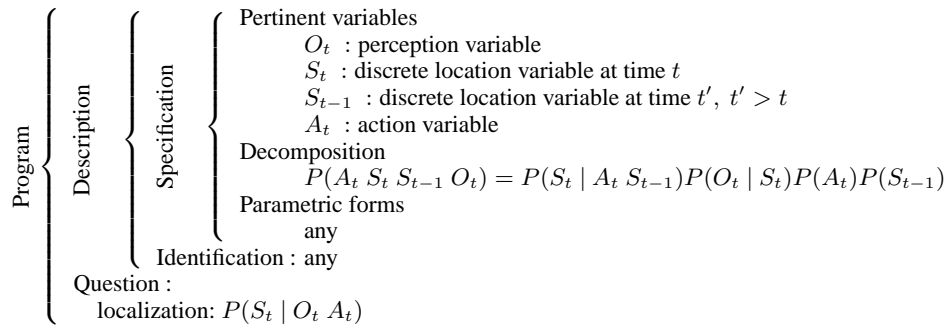


Figure 2: The Markov Localization definition expressed in the BRP formalism.

description  $c^2$ . This allows the programmer to decompose a robotic program into sub-programs, as in structured computer programming. Therefore, the first step for solving a navigation task is to imagine, or to copy from living beings (see [Kuipers, 2000; Trullier *et al.*, 1997; Franz and Mallot, 2000]), intermediary levels of descriptions or skills, that are relevant. This is somewhat different from most probabilistic models of space, that only rely on one level of description, *i.e.* that try to represent the environment using only one type of features. Forms being questions to other descriptions is a key feature of our Abstraction Operator (see Section 4).

The second remark is that the first step when designing a BRP description is the choice of variables. When dealing with the representation of space, one usually selects a perception variable, an action variable, and a location (or state) variable. Therefore, the programmer has to choose a set of locations that are *relevant for solving the task at hand, in the class of environments the robot will likely face*. The choice of the nature of these locations (metric or topologic, or dense or sparse, for instance) should come *as a consequence of these considerations*. This, again, somewhat differs from existing approaches, where the choice of model (Markov Localization or Kalman Filter, for instance), is rather a choice of a *dependency structure* or *form definition*, that implies properties on the choice of variables (Kalman Filters are well suited to continuous variables, for instance). In contrast, in the Bayesian Map formalism, we will not put constraints on the choice of decomposition or forms: the programmer will have all latitude left for choosing the semantic of the location variable that solves his navigation task (the constraints on the choice of variables will merely be syntactic).

The third and final remark is that, in BRP, the description phase is considered <sup>4</sup> independent of the utilization phase. This contrasts with most probabilistic models, where the terms appearing in the decomposition are usually chosen for a particular inference (when they are not said to be the only correct choice of decomposition, see [Boutilier *et al.*, 1999]). For example, action or transition models, which can be difficult to assess when the variables are not chosen well, are still very common because they are easily integrated into the lo-

calization estimation. In our Bayesian Map formalism, we will constraint what maps are used for (the questions), but not how the knowledge necessary for using the map is structured (the decomposition).

### 3 Bayesian Maps

#### 3.1 Definition

A Bayesian Map  $c$  is a description that defines a joint distribution  $P(P L_t L_{t'} A)$ , where:

- $P$  is a perception variable (the robot reads its values from physical sensors or lower level variables),
- $L_t$  is a *location* variable at time  $t$ ,
- $L_{t'}$  is a variable having the same domain than  $L_t$ , but at time  $t'$  (without loss of generality, let us assume  $t' > t$ ),
- and  $A$  is an action variable (the robot writes commands on this variable).

The choice of decomposition is not constrained. Any probabilistic dependency structure can therefore be chosen here: see the recent [Attias, 2003] for an example of how this leverage can lead to interesting new models. Finally, the definition of forms and the learning mechanism (if any) are not constrained, either.

For a Bayesian Map to be useable in practice, we need the description to be rich enough to generate *behaviors*. We call *elementary behavior* any question of the form  $P(A^i | X)$ , where  $A^i$  is a subset of  $A$ , and  $X$  a subset of the other variables of the map (*i.e.*, not in  $A^i$ ). A behavior can be not elementary, for example if it is a sequence of elementary behaviors, or, in more general terms, if it is based on elementary behaviors and some other knowledge (which need not be expressed in terms of maps).

For a Bayesian Map to be interesting, we will also require that it generates *several* behaviors – otherwise, defining just a single behavior instead of a map is enough. Such a map is therefore a resource, based on a location variable relevant enough to solve a class of tasks: this internal model of the world can be reified.

A “guide” one can use to “make sure” that a given map will generate useful behaviors, is to check if the map answers in a relevant manner the three questions  $P(L_t | P)$  (localization),  $P(L_{t'} | A L_t)$  (prediction) and  $P(A | L_t L_{t'})$  (control).

<sup>4</sup>Assuming the inference space and time requirements issues are taken care of by the programmer – usually by choosing a reasonable model, or decomposing further into sub-models.

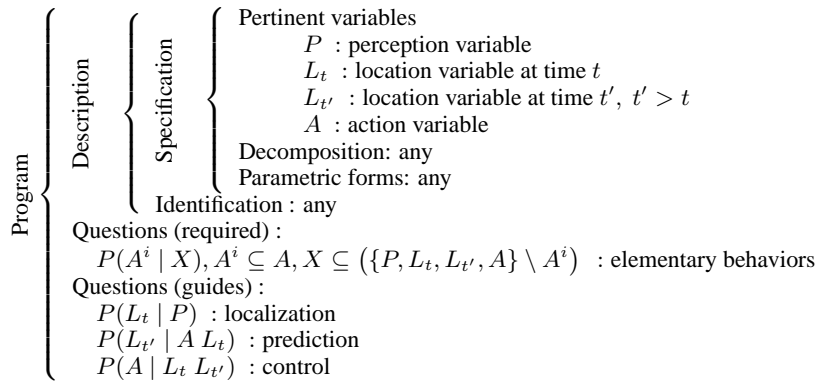


Figure 3: The Bayesian Map model definition expressed in the BRP formalism.

By “relevant manner”, we mean that these distributions have to be informative, in the sense that their entropy is “far enough” of its maximum (*i. e.* the distribution is different from a uniform distribution). This constraint is not formally well defined, but it seems intuitive to focus on these three questions. Indeed, the skills of localization, prediction and control are well identified in the literature as means to generate behaviors. Checking that the answers to these questions are informative is a first step to evaluate the quality of a Bayesian Map with respect to solving a given task.

Figure 3 is a summary of the definition of the Bayesian Map formalism.

### 3.2 Generality of the Bayesian Map formalism

We now invite the reader to verify that the Markov Localization model is indeed a special case of the Bayesian Map model by comparing Figures 2 and 3. Recall that Kalman Filters and Particle Filters are special cases of Markov Localization, as they add hypotheses over the choice of dependency structure made by the Markov Localization model. This implies that Kalman Filters and Particle Filters also are special cases of Bayesian Maps.

Bayesian Maps can therefore accomodate many different forms, depending on the needs or information at hand: for example, one Bayesian Map can be structured like a real valued Kalman Filter for tracking the angle and distance to some feature when it is available. If that feature is not present, or in cases where the linearity hypotheses fail, we can use another Bayesian Map, which need not be a Kalman Filter (for example, based on a symbolic variable).

Hierarchies built of Bayesian Maps (via the abstraction operator) can thus be hierarchies of Markov Localization models, hierarchies of Kalman Filters, etc. Moreover, heterogeneous hierarchies of these models can be imagined: ML over KFs, or even  $n$  KFs and one ML model, which, in our view, would be a formally satisfying alternative to [Tomatis *et al.*, 2001].

## 4 Abstraction of Bayesian Maps

Having defined the Bayesian Map concept, we now turn to defining operators for putting Bayesian Maps together. The

one we present here is called the abstraction of maps, it is defined Figure 4, and commented in the rest of this section.

As stressed above, in a Bayesian Map, the semantics of the location variable can be very diverse. The main idea behind the abstraction operator is to build a *Bayesian Map*  $c$  whose different locations are other Bayesian Maps  $c^1, c^2, \dots, c^n$ . The location variable of the abstract map will therefore take  $n$  possible symbolic values, one for each underlying map  $c^i$ . Each of these maps will be “nested” in the higher level abstract map, which justifies the use of the term “hierarchy” in our work. Recall that Bayesian Maps are designed for generating behaviors. In the abstract map, the lower level behaviors  $a^1, a^2, \dots, a^k$  can be used for linking the locations  $c^i$ . The action variable of the abstract map will therefore take  $k$  possible symbolic values, one for each behavior of the underlying maps. In order to build an abstract map having  $n$  locations, the programmer will have to have previously defined  $n$  lower level maps, which generate  $k$  behaviors. The numbers  $n$  and  $k$  are therefore small, and so the abstract map deals with a small internal space, having retained of each underlying map only a symbol, and having “forgotten” all their details. This justifies the use of the name “abstraction” for this operator. But this “summary mechanism” has yet to be described: that is what the perception variable  $P$  of the abstract map will be used for, as it will be the list of all the variables appearing in the underlying maps.

Given the four variables of the abstract map, we define its joint distribution  $P(P L_t L_{t'} A)$  with the following decomposition:

$$\begin{aligned}
 P(P L_t L_{t'} A) &= P(P^1 L_t^1 L_{t'}^1 A^1 \dots P^n L_t^n L_{t'}^n A^n L_t L_{t'} A) \\
 &= P(L_t) \prod_{i=1}^n P(P^i L_t^i L_{t'}^i A^i | L_t) P(L_{t'}) P(A | L_t L_{t'}).
 \end{aligned}$$

In this decomposition,  $P(L_t)$  and  $P(L_{t'})$  are defined as uniform distributions. All the terms of the form  $P(P^i L_t^i L_{t'}^i A^i | [L_t = c])$  are defined as follows: when  $c \neq c^i$ , the probabilistic dependency between the variables  $P^i, L_t^i, L_{t'}^i, A^i$  of the map  $c^i$  is supposed unknown, therefore defined by a uniform distribution. Whereas when  $c = c^i$ ,

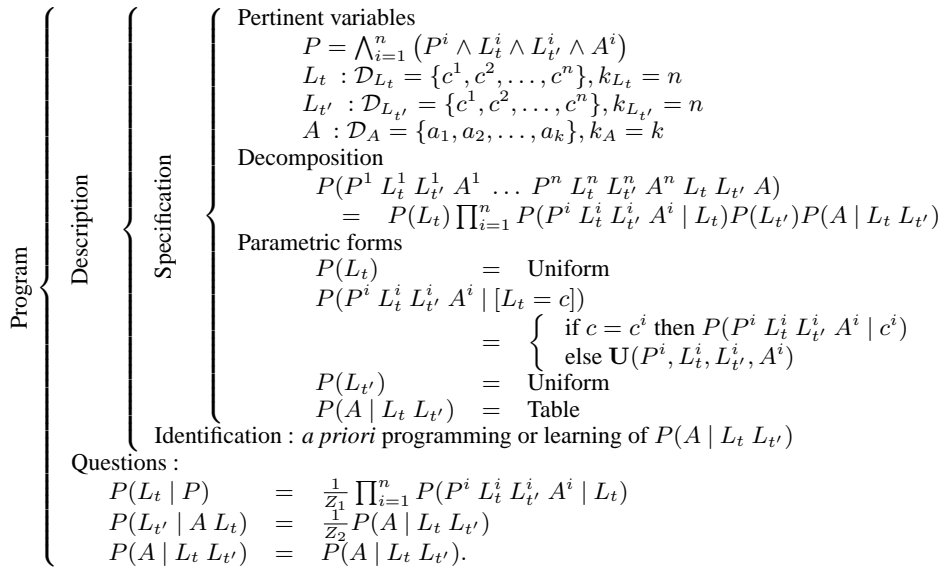


Figure 4: The abstraction operator definition expressed as a Bayesian Map.

this dependency is exactly what the map  $c^i$  defines. Therefore this term is a question to the description  $c^i$ , but a question that includes the whole sub-description by asking for the joint distribution it defines<sup>5</sup>. Since the last term,  $P(A | L_t L_{t'})$ , only includes symbolic variables that have a small number of values, it makes sense to define it as a table, which can be easily *a priori* programmed or learned experimentally.

The abstract Bayesian Map is now fully defined, and, given  $n$  underlying maps, can be automatically built. The last step is to verify that it generates useful behaviors. Here, we will examine the three guide questions of localization, prediction and control.

The localization question leads to the following inference (derivation omitted):

$$\begin{aligned}
 P(L_t | P) &= P(L_t | P^1 L_t^1 L_{t'}^1 A^1 \dots P^n L_t^n L_{t'}^n A^n) \\
 &= \frac{1}{Z} \prod_{i=1}^n P(P^i L_t^i L_{t'}^i A^i | L_t).
 \end{aligned}$$

The interpretation of this result is best explained with an example (see Section 5).

As for the prediction and control questions, they are easily computed: the control question  $P(A | [L_t = l] [L_{t'} = l'])$  appears in the decomposition, therefore it requires no inference. The prediction question is an “inversion”:  $P(L_{t'} | [A = a] [L_t = l]) \propto P([A = a] | [L_t = l] L_{t'})$ .

Recall that the final goal of any Bayesian Map is to provide behaviors. In the abstract map, this is done by answering a

<sup>5</sup>Of course, using a joint distribution as a question is dangerous with respect to the assumption that the inference is dealt with by the programmer. Special care has to be taken for this quantity to be used in the inference, as the dimension of the space of these distribution makes them prone to degenerate quickly to numerical epsilons. Discussing the possible solutions to this problem is beyond the scope of this paper.

question like  $P(A | [L_{t'} = c] [P = p])$ : what is the probability distribution over lower level behaviors, knowing all values  $p$  of the variables of the lower level, and knowing that we want to “go to map  $c$ ”? Answering this question thus allows to select the most relevant underlying behavior to reach a given high level goal. The computation is as follows:

$$\begin{aligned}
 P(A | L_{t'} P) &= \frac{1}{Z} \sum_{L_t} \left( \prod_{i=1}^n P(P^i L_t^i L_{t'}^i A^i | L_t) \right) P(A | L_t L_{t'}).
 \end{aligned}$$

This computation includes the localization question, to weigh the probabilities given by the control model  $P(A | L_t L_{t'})$ . In other words, the distribution over the action variable  $A$  includes all localization uncertainties. Each underlying model is used, even when the robot is located at a physical location that this model is not made for. As a direct consequence, there is no need to *decide* what map the robot is in, or to *switch* from map to map: the computation considers all possibilities and weighs them according to their (localization) probabilities. Therefore the underlying maps need not be “mutually exclusive” in a geographical sense.

## 5 Experimental validation

We report here an experiment made on the well-known Koala mobile robot platform<sup>6</sup>. In order to keep as much control as possible over our experiments and the different effects we observe, we simplify the sensorimotor system and its environment. We only use the 16 proximeters  $Px = Px_0 \wedge \dots \wedge Px_{15}$  of our robot, and keep two degrees of freedom of motor control, via the rotation and translation speed  $V_{rot}$  and  $V_{trans}$ . The environment we use is a 5 m  $\times$  5 m area made of movable planks, that allow us to easily set up any configuration (see a

<sup>6</sup>See <http://www.K-team.com>.

typical shape we use, Figure 5). The goal of this experiment is to solve a navigation task: we want the robot to be able to go hide in any corner, as if the empty space in the middle of the area were dangerous.

The first programming step is to analyse this task into sub-tasks. We particularize three situations that are relevant for solving the task: the robot can either be near a wall, and it should follow it in order to reach the nearest corner, or the robot can be in a corner, and it should stop, or finally it could be in empty space, and should therefore go straight, so as to leave the exposed area as quickly as possible.

## 5.1 Low level Bayesian Maps

Given this analysis, the second programming step is to define one Bayesian Map for each of the three situations. They all use the same perception variable  $P = Px$  and the same action variable  $A = Vrot \wedge Vtrans$ .

The first map,  $c^{wall}$  describes how to navigate in presence of a single wall, using a location variable  $L_t = \theta \wedge Dist$ : the phenomenon “wall” is summed up by an angle and a distance. Therefore,  $c^{wall}$  defines  $P(Px \theta_t Dist_t \theta_{t'} Dist_{t'} Vrot Vtrans | c^{wall})$ . We have implemented this map using 12 possible angle values, and 3 different distances. This lead to a compact model, yet accurate enough to solve the sub-tasks we wanted to solve. The dependency structure we choose is ( $c^{wall}$  on right hand sides omitted):

$$\begin{aligned} & P(Px \theta_t Dist_t \theta_{t'} Dist_{t'} Vrot Vtrans) \\ &= P(\theta_t Dist_t) \prod_i P(Px_i | \theta_t Dist_t) P(\theta_{t'} Dist_{t'}) \\ & \quad P(Vrot | \theta_t Dist_t \theta_{t'} Dist_{t'}) \\ & \quad P(Vtrans | \theta_t Dist_t \theta_{t'} Dist_{t'}). \end{aligned}$$

$P(\theta_t Dist_t)$  and  $P(\theta_{t'} Dist_{t'})$  are uniform probability distributions. Each term of the form  $P(Px_i | \theta_t Dist_t)$  is a set of Gaussians, that was identified experimentally, by a supervised learning phase: we physically put the robot in all 36 possible situations with respect to the wall, and recorded proximeter values so as to compute experimental means and standard deviations. Finally, the two control terms  $P(Vrot | \theta_t Dist_t \theta_{t'} Dist_{t'})$  and  $P(Vtrans | \theta_t Dist_t \theta_{t'} Dist_{t'})$  were programmed “by hand”: given the current angle and distance, and the angle and distance to be reached, what should be the motor commands?

This map successfully solves navigation tasks like “follow-wall-right”, “follow-wall-left”, “go-away-from-wall”, “stop”, using behaviors of the same name. For example, “follow-wall-right” is defined by the probabilistic question  $P(Vrot Vtrans | Px [L_{t'} = \langle 90, 1 \rangle])$ : compute the probability distribution on motor variables knowing the sensory input and knowing that the location to reach is  $\theta = 90^\circ$ ,  $Dist = 1$  (wall on the right at medium distance).

This map is an instance where a Kalman Filter could have been used instead. For example, if we had required a more accurate computation of the angle and distance to the wall, the supervised learning method we used would not have been manageable anymore, as the time to experimentally visit all

configurations would have grown. We could then have replaced the  $c^{wall}$  Bayesian Map by a Kalman Filter based Bayesian Map. In the case of our experiments, that was not necessary.

The two other Bayesian Maps we define are:

- $c^{corner}$  describes how to navigate in a corner, using a symbolic location variable that can take 4 values: *FrontLeft*, *FrontRight*, *RearLeft* and *RearRight*. This is enough for solving tasks like “quit-corner-and-follow-right”, “away-from-both-walls”, “stop”.
- $c^{empty-space}$ , which is very simple, describes how to navigate in empty space, *i.e.* when the sensors do not see anything. The behaviors defined here are “straight-ahead” and “stop”.

## 5.2 Abstract Bayesian Map

Given these three maps, the third and final programming step is to apply the abstraction operator on them. We obtain a map  $c$ , whose location variable is  $L_t = \{c^{wall}, c^{corner}, c^{empty-space}\}$ . The action variable lists the behaviors defined by the low level maps:  $A = \{\text{follow-wall-right, go-away-from-wall, } \dots\}$ . The rest of the abstract map is according to the schema given Section 4.

We want here to discuss the localization question. Let us assume that the robot is in empty space: all its sensors read 0. Let us also assume that the robot is currently applying the “straight-ahead” behavior, that sets  $Vrot$  and  $Vtrans$  near 0 (no rotation) and 40 (fast forward movement), respectively, using sharp Gaussian distributions.

Let us consider the probability to be in location  $c^{empty-space}$  (with  $w$  standing for *wall*,  $c$  for *corner* and  $e$  for *empty-space*):

$$\begin{aligned} & P([L_t = c^{empty-space}] | P) \\ & \propto \left( \begin{array}{l} P(P^w L_t^w L_{t'}^w A^w | [L_t = c^{empty-space}]) \\ P(P^c L_t^c L_{t'}^c A^c | [L_t = c^{empty-space}]) \\ P(P^e L_t^e L_{t'}^e A^e | [L_t = c^{empty-space}]) \end{array} \right). \end{aligned}$$

Of the three terms of the product, two are uniforms, and one is the joint distribution given by  $c^{empty-space}$ . That joint distribution gives very high probability for the current situation, as describing the phenomenon “going straight ahead in empty space” basically amounts to favoring sensory readings of 0 and motor commands near 0 and 40 for  $Vrot$  and  $Vtrans$ , respectively. The situation is quite the opposite for  $P([L_t = c^{wall}] | P)$ : for example,  $c^{wall}$  does not favor at all this sensory situation. Indeed, the phenomenon “I am near a wall” is closely related to the fact that the sensors actually sense something. The probability of seeing nothing on the sensors knowing that the robot is near a wall is very low:  $P([L_t = c^{wall}] | P)$  will be very low. The reasoning is similar for  $P([L_t = c^{corner}] | P)$ .

This computation can thus be interpreted as the *recognition of the most pertinent underlying map* for a given sensory situation  $P$ . Alternatively, it can be seen as a *measure of the coherence of the variables* of each underlying map, or even as a *Bayesian comparison of the relevance of models*, as assessed by the numerical value of the joint distributions of each lower level model. Since these distributions include (lower

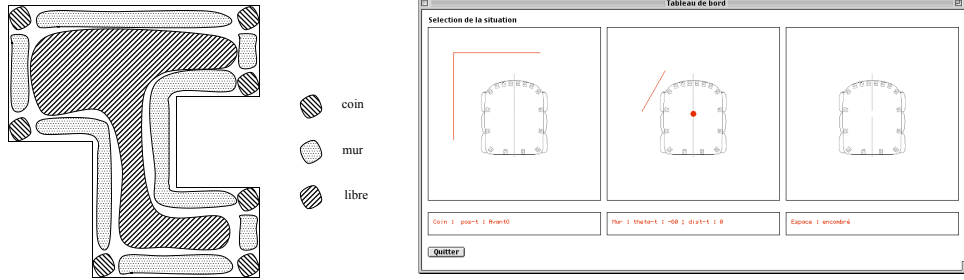


Figure 5: 2D projection of the estimated “validity zones” of the maps  $c^{wall}$ ,  $c^{corner}$  et  $c^{empty-space}$ . The right part of the figure is a screenshot of the localization module of the abstract map, that shows the “comparison” and competition between the underlying models. The winner is marked by the central dot: in this case, the robot was near a wall.

level) location and action variables, the maps are not only recognized by sensory patterns, but also by what the robot is currently doing.

The localization question can therefore be used to assess the “validity zones” of the underlying maps, *i.e.* the places of the environment where the hypotheses of each model hold. Experimentally, we have the robot navigate in the environment, and ask at each time step the localization question. We can summarize visually the answer, for example by drawing values for  $L_t$ , and report the drawn value on a cartesian map of the environment. A (simplified but readable) result is shown Figure 5. As can be seen, the robot correctly recognizes each situation that it has a model for. As a conclusion, let us note that the resulting zones are not contiguous in the environment: for example, all the corners of the environment are associated with the same symbol, namely,  $c^{corner}$ . This effect is known as *perceptual aliasing*. But this very simple representation is sufficient for solving the task that was given to the robot: we report here that the behavior “go-hide-in-corner” is indeed generated by the abstract map.

A typical trajectory for the robot, when we start it in the middle of the arena, is to start by going straight ahead. As soon as a couple of forward sensors sense something, the “empty-space” situation is not relevant anymore, and the robot applies the next best model it has, depending on the correlation between what the sensors see: if it looks like a wall and moves like a wall, then the probability for the “wall” model will be high. On the other hand, if it rather feels like a corner, then the corner model will win the probabilistic competition. Suppose it was near a wall, then it starts to follow it, until a corner is reached. In our first version, because the “corner” model was designed independently of the wall model, the validity zone of the  $c^{corner}$  map was too small, and seldom visited by the robot as it passed the corner using the “follow-wall-right” behavior. The robot would then miss the first corner, and stop at another one. This shows that the decomposition of the task gives independent sub-tasks only as a first approximation. We solved the problem by modifying the “corner” model, so that it would recognize a corner on a typical “follow-wall-right” trajectory.

## 6 Conclusion

We have presented a formalism called the Bayesian Map, which is a generalization of most probabilistic models of space found in the literature. Indeed, it drops the usual constraints on the decomposition, on the choice of forms, even on the implementation of the probability distributions. We have presented an operator, called the Abstraction operator, for building hierarchies of Bayesian Maps.

The experiments we presented are of course to be regarded only as “proofs of concept”. Their simplicity also serves didactic purposes, while on the other hand, the task considered could have been achieved in many other ways. However, these experiments, in our view, are a successful preliminary step toward applying our formalism. Part of the current work is aimed at enriching the Bayesian Maps presented here for testing the *scaling up* capacity of our formalism when confronted to a large scale scenario.

Moreover, since each map of the hierarchy is a full probabilistic model over the four variables of perception, action and locations at two different time-steps, it is potentially very rich. Possible computations based on these maps include questions like the prediction question  $P(L_{t'} | A L_t)$ , which forms the basis of *planning* processes. Hierarchies of Bayesian Maps are therefore to be placed alongside model based approaches, instead of pure reactive approaches. Exploiting such knowledge by integrating a planning process in our Bayesian Map formalism is also part of the ongoing work, and also requires a full scale experiment to be validated.

Although we are only at the beginning of this research track, we do believe it is a promising one, and hope it will spark some interest in the community as well.

## Acknowledgments

The authors wish to thank the anonymous reviewers for their comments.

## References

[Attias, 2003] Hagai Attias. Planning by probabilistic inference. In *Ninth International Workshop on Artificial Intelligence and Statistics Proceedings*, 2003.

- [Bengio and Frasconi, 1995] Y. Bengio and P. Frasconi. An input/output HMM architecture. In G. Tesauro, D.S. Touretzky, and T.K. Leen, editors, *Advances in Neural Information Processing Systems 7*, pages 427–434. MIT Press, Cambridge, MA, 1995.
- [Bessière *et al.*, 2003] Pierre Bessière, Juan-Manuel Ahuactzin, Olivier Aycard, David Bellot, Francis Colas, Christophe Coué, Julien Diard, Ruben Garcia, Carla Koike, Olivier Lebeltel, Ronan LeHy, Olivier Malrait, Emmanuel Mazer, Kamel Mekhnacha, Cédric Pradalier, and Anne Spalanzani. Survey: Probabilistic methodology and techniques for artefact conception and development. Technical Report RR-4730, INRIA Rhône-Alpes, Montbonnot, France, 2003.
- [Boutilier *et al.*, 1999] Craig Boutilier, Thomas Dean, and Steve Hanks. Decision theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 10:1–94, 1999.
- [Cacciatore and Nowlan, 1994] Timothy W. Cacciatore and Steven J. Nowlan. Mixtures of controllers for jump linear and non-linear plants. In Jack D. Cowan, Gerald Tesauro, and Joshua Alspector, editors, *Advances in Neural Information Processing Systems*, volume 6, pages 719–726. Morgan Kaufmann Publishers, Inc., 1994.
- [Diard, 2003] Julien Diard. *La carte bayésienne – Un modèle probabiliste hiérarchique pour la navigation en robotique mobile*. Thèse de doctorat, Institut National Polytechnique de Grenoble, Grenoble, France, Janvier 2003.
- [Franz and Mallot, 2000] Matthias Franz and Hanspeter Mallot. Biomimetic robot navigation. *Robotics and Autonomous Systems*, 30:133–153, 2000.
- [Hauskrecht *et al.*, 1998] Milos Hauskrecht, Nicolas Meuleau, Leslie Pack Kaelbling, Thomas Dean, and Craig Boutilier. Hierarchical solution of Markov decision processes using macro-actions. In Gregory F. Cooper and Serafín Moral, editors, *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pages 220–229, San Francisco, July, 24–26 1998. Morgan Kaufmann.
- [Kaelbling *et al.*, 1998] L.P. Kaelbling, M.L. Littman, and A.R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, 1998.
- [Kuipers, 2000] Benjamin J. Kuipers. The spatial semantic hierarchy. *Artificial Intelligence*, 119(1–2):191–233, 2000.
- [Lane and Kaelbling, 2001] Terran Lane and Leslie Pack Kaelbling. Toward hierarchical decomposition for planning in uncertain environments. In *Proceedings of the 2001 IJCAI Workshop on Planning under Uncertainty and Incomplete Information*, Seattle, WA, August 2001. AAAI Press.
- [Lane and Kaelbling, 2002] Terran Lane and Leslie Pack Kaelbling. Nearly deterministic abstractions of markov decision processes. In *Eighteenth National Conference on Artificial Intelligence (AAAI-2002)*, 2002.
- [Lebeltel *et al.*, 2003] Olivier Lebeltel, Pierre Bessière, Julien Diard, and Emmanuel Mazer. Bayesian robot programming. *Autonomous Robots (accepted for publication)*, 2003.
- [Leonard *et al.*, 1992] John Leonard, Hugh Durrant-Whyte, and Ingemar Cox. Dynamic map-building for an autonomous mobile robot. *The International Journal of Robotics Research*, 11(4):286–298, 1992.
- [Murphy, 2002] Kevin Murphy. *Dynamic Bayesian Networks : Representation, Inference and Learning*. Ph. D. thesis, University of California, Berkeley, Berkeley, CA, July 2002.
- [Pineau and Thrun, 2002] Joelle Pineau and Sebastian Thrun. An integrated approach to hierarchy and abstraction for POMDPs. Technical Report CMU-RI-TR-02-21, Carnegie Mellon University, August 2002.
- [Rabiner and Juang, 1993] Lawrence R. Rabiner and Biing-Hwang Juang. *Fundamentals of Speech Recognition*, chapter Theory and implementation of Hidden Markov Models, pages 321–389. Prentice Hall, Englewood Cliffs, New Jersey, 1993.
- [Roweis and Ghahramani, 1999] Sam Roweis and Zoubin Ghahramani. A unifying review of linear gaussian models. *Neural Computation*, 11(2):305–345, February 1999.
- [Smyth *et al.*, 1997] Padhraic Smyth, David Heckerman, and Michael I. Jordan. Probabilistic independence networks for hidden markov probability models. *Neural Computation*, 9(2):227–269, 1997.
- [Thrun, 1998] Sebastian Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71, 1998.
- [Thrun, 2000] Sebastian Thrun. Probabilistic algorithms in robotics. *AI Magazine*, 21(4):93–109, 2000.
- [Thrun, 2002] Sebastian Thrun. Robotic mapping : A survey. Technical Report CMU-CS-02-111, Carnegie Mellon University, February 2002.
- [Tomatis *et al.*, 2001] Nicola Tomatis, Illah Nourbakhsh, Kai Arras, and Roland Siegwart. A hybrid approach for robust and precise mobile robot navigation with compact environment modeling. In *Proceedings of the 2001 IEEE International Conference on Robotics and Automation (ICRA'01)*, pages 1111–1116, 2001.
- [Trullier *et al.*, 1997] Olivier Trullier, Sidney Wiener, Alain Berthoz, and Jean-Arcady Meyer. Biologically-based artificial navigation systems : Review and prospects. *Progress in Neurobiology*, 51:483–544, 1997.