

An Embedded Implementation of Bayesian Network Robot Programming Methods

By **Mark A. Post**

Lecturer, Space Mechatronic Systems Technology Laboratory.
Department of Design, Manufacture and Engineering Management.
University of Strathclyde, Glasgow, United Kingdom.

Abstract

A wide variety of approaches exist for dealing with uncertainty in robotic reasoning, but relatively few consider the propagation of statistical information throughout an entire robotic system. The concept of Bayesian Robot Programming (BRP) involves making decisions based on inference into probability distributions, but can be complex and difficult to implement due to the number of priors and random variables involved. In this work, we apply Bayesian network structures to a modified BRP paradigm to provide intuitive structure and simplify the programming process. The use of discrete random variables in the network can allow high inference speeds, and an efficient programming toolkit suitable for use on embedded platforms has been developed for use on mobile robots. A simple example of navigational reasoning for a small mobile robot is provided as an example of how such a network can be used for probabilistic decisional programming.

1. Introduction

One of the most pressing problems in mobile robotics is that of how to quantify and propagate certainty and uncertainty measures when making decisions for control. It is desirable for the rover itself to be able to deal with uncertainties probabilistically, since this gives a robot the ability to appropriately handle unexpected and uncertain circumstances. Bayesian Networks (BN) are well-suited for handling uncertainty in cause-effect relations, and handle dependence/independence relationships well provided that the network is constructed using valid relational assumptions. Some drawbacks of this method are that the variables, events, and values available must be well-defined from the beginning, and the causal relationships and conditional probabilities must be available initially Kjrulff, 2008.

We approach the problem of probabilistic robotics using the Bayesian Robot Programming (BRP) methodology developed by Lebeltel, Bessiere et al Bessiere et al., 2000 Lebeltel et al., 2000 Lebeltel et al., 2004, which provides a quite comprehensive framework for robotic decision-making using inference. We add to this method by formally using Bayesian networks as a knowledge representation structure for programming, which adds clarity of representation, a practical structure for constructing joint distributions dynamically, and reliance on a proven probabilistic methodology. Also, the use of recursive inference in a Bayesian network avoids the need to manually partition and decompose large joint distributions, which greatly simplifies the programming process. Our approach to using Bayesian networks for robotic programming is detailed in the dissertation by Post, 2014, and a summary of the approach is given here.

2. Probabilistic Concepts

The fundamental concept of a probability p is a real positive number between zero and 1 (denoting 100%) that is used to assign quantitative measures of confidence to a given result or set of results of that event from result space Ω . These principles are formalized in the concept of a random variable that assigns a value to each outcome $o \in \Omega$ and maps from the events $o \in O$ to real probability values p within a probability distribution P . In numerical terms, this means that for a countable number of events $o_1 \dots o_N$, the distribution P must at least provide N mappings $p_1 \dots p_N$. Each mapping is described using functional notation such as $p = P(x)$. We denote random variables with an uppercase letter such as X , while a specific value that the random variable can take is denoted in lowercase such as x like any other scalar. Hence, the statement $X = x$ refers to a specific value x within the variable X . We refer to a specific value $P(X = x)$ with the simple form $P(x)$, and the whole distribution by $P(X)$. Also, when dealing with several random variables, $P((X = x) \cap (Y = y))$ is written as $P(x, y)$.

In the field of robotics, conditional probabilities can be used to calculate conjunctions of events such as $(X = x) \cap (Y = y)$. For an example of an obstacle sensor, we can define the Boolean events $Y = y :=$ “An obstacle is detected” and $X = x :=$ “An obstacle is actually present”, so that $P((X = x) \cap (Y = y))$ represents the event that an obstacle is detected correctly, with priors such as $P(Y = y) = 0.3$ being necessary to calculate such probabilities. We quantify a conditional probability distribution that takes these priors into account, usually abbreviated to $P(x|y)$, as

$$P(X = x|Y = y) = \frac{P(Y = y \cap X = x)}{P(Y = y)}. \quad (2.1)$$

The chain rule for conditional probabilities generalizes the conditional probability $P(X_1 \cap X_2) = P(X_1)P(X_2|X_1)$ to incorporate n events as a product of conditional distributions

$$P(X_1 \cap X_2 \cap \dots \cap X_N) = P(X_1)P(X_2|X_1) \dots P(X_N|X_1 \cap \dots \cap X_{n-1}). \quad (2.2)$$

Because the conjunction operator is commutative, we can replace $P(Y = y \cap X = x)$ with $P(X = x \cap Y = y) = P(Y = y|X = x)P(X = x)$ in Equation 2.1 to express the conditional distribution entirely in terms of event distributions and conditionals. This is the form of *Bayes' Rule*, which is the foundation of all Bayesian inference systems, including Markov chains and Kalman filtering Engelbrecht, 2002.

$$P(Y|X) = \frac{P(X|Y)P(Y)}{P(X)} = \frac{(\textit{likelihood}) \cdot (\textit{prior})}{(\textit{evidence})}. \quad (2.3)$$

We can consider $P(Y = y|X = x)$ as the *likelihood* of getting a sensor reading if an object to be sensed is present, and $P(X = x)$ as the probability of an object being present *prior* to the reading being made. The sensor reading itself is the *evidence* given that points to object presence, and becomes a certainty of $P(Y = y) = 1$ after a reading is made, but is generally estimated as $0 < P(Y = y) < 1$ in the absence of a positive sensor reading. In this way, the process of finding $P(X = x|Y = y)$ is known as Bayesian Inference.

The use of joint probability distributions as a basis for modelling relationships and making predictions is a very powerful concept, but calculating all joint and conditional distributions for all but the simplest systems is an intractable problem. The key to solving this problem is the use of independence between random variables to greatly

decrease the number of values that must be stored in a joint distribution. Returning to the collision-detection sensor example, we can define X as the random variable of object detection with, and Y as the random variable of object actual presence. Rather than having to measure the specific distribution $P(X, Y)$ directly, only the probability of object detection and the accuracy of the sensor are needed. Assuming marginal independence $P \models (W \perp Y|X)$, We can write a conditional joint distribution with three random variables as

$$P(W, Y|X) = \frac{P(W|Y)P(X|Y)P(Y)}{P(X)}. \quad (2.4)$$

In “naive” Bayes models such as this, we can generalize the factorization of a joint distribution of M variables $X_1 \dots X_M$ that are marginally independent but dependent on a condition Z to

$$P(X_1, \dots, X_M, Z) = P(Z) \prod_{m=1}^M P(X_m|Z). \quad (2.5)$$

For us to be able to properly organize and represent a large set of joint distributions using factorization in this way, we need a method of clearly associating random variables that are dependent on each other, in the case of our sensor example the association of W and X with Y . A Bayesian network with random variables as nodes provides a compact way to encode both the structure of a conditional factorization of a joint distribution, and also (equivalently) the independence assumptions about the random variables included in the distribution. Independent random variables have no parents, while marginally dependent or conditional random variables have parents connected as in the case of $(W \perp X|Y)$. Hence, the Bayesian network can serve as a framework for formalizing expert knowledge about how the world works. Applying the chain rule to all nodes in a Bayesian network, the probability distribution over a given network or subnetwork of nodes $\varphi = \{X_1 \dots X_M\}$ can be said to factorize over φ according to the dependencies in the network if the distribution can be expressed as a product Koller and Friedman, 2009

$$P(\{X_1 \dots X_M\}) = \prod_{m=1}^M P(X_m|Pa(X_m)). \quad (2.6)$$

Due to the dependency structure of the network, the conditional posterior probability of a given node X depends on all its parents $Pa(X)$, so queries are a recursive operation through all parent nodes $Y \in Pa(X)$ to determine their probability distributions, then multiplying them by the probability distributions of each parent node Y such that

$$P(X = x) = \sum_{Y \in Pa(X)} P(X = x|Y = y)P(Y = y). \quad (2.7)$$

For single-parent random variables this is effectively a matrix multiplication, but for multiple parents, each combination of probabilities leading to a given value must be considered. Representing probability distributions by matrices, the distribution $P(X|Y)$ will be an $L+1$ -dimensional matrix for L parents, with the major dimension of size N for N random variable values in $V(X)$. A given parent Y is assumed to have a distribution of size M , so that the distribution $P(Y)$ is a $N \times 1$ matrix. Each element of the matrix represents the probability of a given combination of values from parents occurring, so to

calculate the conditional distribution for X , all probability products of each combination leading to a given outcome are added together, as in Equation 2.7. This results in an $N \times 1$ matrix that is stored as a temporary posterior distribution estimate for $P(X)$ which avoids frequent recalculation for determining the conditional distributions of children $Ch(X)$ while traversing the network. This process is graphically illustrated in Figure 1 for a set of single parents (a) and for multiple parents (b).

In this way, any node in a Bayesian network can be queried to obtain a probability distribution over its values. While exact inference as described into a Bayesian network is widely understood to be an NP-hard problem Wu and Butz, 2005, it is still much more efficient than the raw computation of a joint distribution, and provides an intuitive, graphical method of representing dependencies and independences. It is easy to see how any system that can be described as a set of independent but conditional random variables can be abstracted into a Bayesian network, and that a wealth of information regarding the probability distributions therein can be extracted relatively efficiently.

3. Bayesian Programming

A Bayesian program has been defined by Lebeltel et al. as a group of probability distributions selected so as to allow control of a robot to perform tasks related to those distributions. A “Program” is constructed from a “Question” that is posed to a “Description”. The “Description” in turn includes both “Data” represented by δ , and “Preliminary Knowledge” represented by π . This “Preliminary Knowledge π consists of the pertinent random variables, their joint decomposition by the chain rule, and “Forms” representing the actual form of the distribution over a specific random variable, which can either be parametric forms such as Gaussian distributions with a given mean and standard deviation, or programs for obtaining the distribution based on inputs Lebeltel et al., 2000.

It is assumed that the random variables used such as X are discrete with a countable number of values, and that a logical proposition of random variables $[X = x_i]$ is mutually exclusive such that $\forall i \neq j, \neg(X = x_i \wedge X = x_j)$ and exhaustive such that $\exists X, (X = x_i)$. All propositions represented by the random variable follow the Conjunction rule $P(X, Y|\pi) = P(X|\pi)P(Y|X, \pi)$, the Normalization rule $\sum_X P(X|\pi) = 1$, and the Marginalization rule $\sum_X P(X, Y|\pi) = P(Y|\pi)$ Bessiere et al., 2000. Rather than unstructured groups of variables, we apply these concepts to a Bayesian network of M random variables $\varphi = X_1, X_2, \dots, X_N \in \pi, \delta$, from which an arbitrary joint distribution can be computed using conjunctions. It is assumed that any conditional independence of random variables in π and δ (which must exist, though it was not explicitly mentioned by Lebeltel et al.) is represented appropriately by the Bayesian network, thus significantly simplifying the process of factorization for joint distributions. The general process we use for Bayesian programming, including changes from the original BRP, is as follows:

(a) **Define the set of relevant variables.** This involves identifying the random variables that are directly relevant to the program desired. In a Bayesian network, this is implicit in the edges between nodes that represent dependencies. Usually, a single child node is queried to include information from all related nodes.

(b) **Decompose the joint distribution.** The original BRP methodology explicitly partitioned a joint distribution of M variables $P(X_1, \dots, X_M|\delta, \pi)$ into subsets, each one a conjunction, and then used the product of the factorization of each subset, a process called decomposition Lebeltel et al., 2004. We make use of the properties of the Bayesian network for implicitly including information in parent nodes when queried and apply

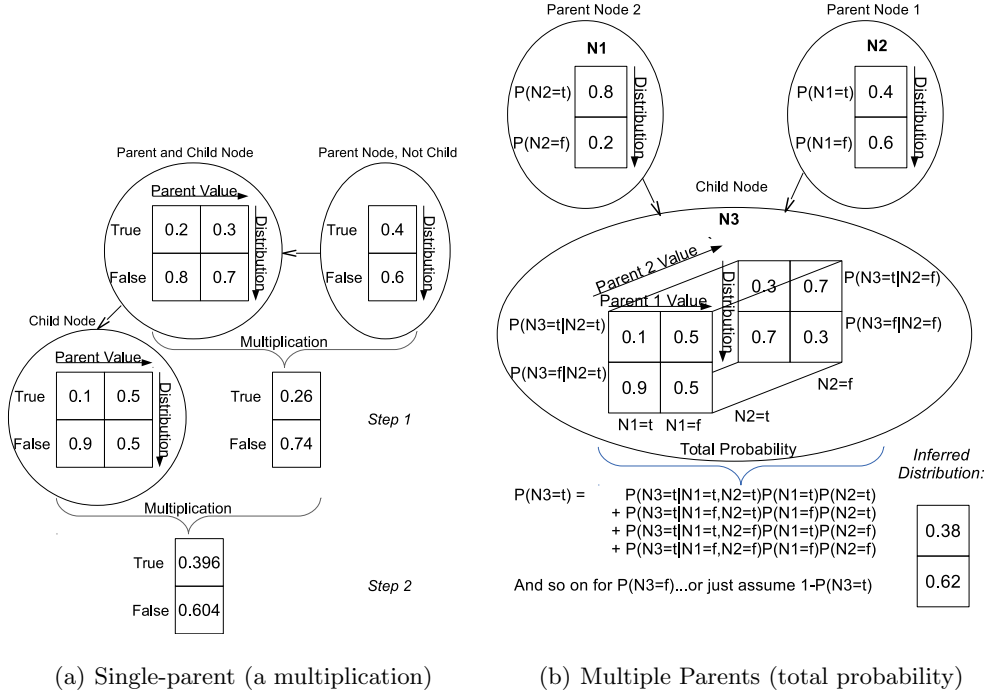


FIGURE 1. Matrix calculations for querying a discrete random variable

the factorization rules in Equation 2.5 to reduce $P(X_1, \dots, X_M | \delta, \pi)$ to a product of conditional distributions which are queried recursively

$$P(X) \prod_{m=1}^M P(X_m | \delta, \pi). \quad (3.1)$$

(c) **Define the forms.** For actual computations, the joint and dependent distributions must be numerically defined, which is done by inserting discrete values into each node. The most common function to be used, and the function used for continuous distributions in this work, is the Gaussian distribution with parameters mean \bar{x} and standard deviation σ that define the shape of the distribution, commonly formulated as

$$P(X = x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\bar{x})^2}{2\sigma^2}}. \quad (3.2)$$

A uniform distribution can also be set with $P(X = x) = k$, having the same probability regardless of value. Finally, because continuous distributions are programmed as function calls, a distribution that is a function into some arbitrary external process such as a map can be used.

(d) **Formulate the question.** Queries into a BRP system traditionally involve partitioning the random variables in φ into three sets: a searched set $Se \subset \varphi$ for variables that contain information we want to determine, a known set $Kn \subset \varphi$ for variables that contain an observed state so that $P(X = x) = 1$ for $X \in Kn$, and unknown variables $Un \subset \varphi$ that are only stochastically estimated. Under these sets, a “question” is formulated as the posterior probability distribution $P(Se | Kn, \pi)$, which makes intuitive sense because in any query we want to obtain the probability distribution over a given random variable based on all the distributions that affect it. Although many methodologies can

be used to determine the “final” answer to the question, namely a specific probability value in this distribution, we will generally use marginal MAP queries to obtain the value of highest probability taking into account the dependencies in Un , making the form of the question

$$\arg \max_{Se} \sum_{Un} P(Se, Un | Kn, \delta, \pi). \quad (3.3)$$

The use of a Bayesian network formalizes the relationships of these sets, so that a query into a common child node of Se incorporates information from parents Kn and Un of Se . It is important to note that a “question” is functionally another conditional distribution, and therefore operates in the same way as an additional node in the Bayesian network.

(e) **Perform Bayesian inference.** To perform inference into the joint distribution $P(X_1, \dots, X_M | \delta, \pi)$, the “Question” that has been formulated as a conjunction of the three sets Searched (Se), Known (Kn), and Unknown (Un) is posed to the system and solved as a Bayesian inference that includes all relevant information to the set Se . For our Bayesian network implementation. The “Answer” is obtained as a probability distribution, and a specific maximum or minimum value can be obtained using a value from the set Se and a Maximum A Posteriori (MAP) query as

$$\text{MAP}(X|Y = y) = \arg \max_x \sum_Z P(X \cap Z | Y). \quad (3.4)$$

The last step in Bayesian programming is the actual inference operation used to determine the probability distribution for the variable or set of variables in question. Obtaining the joint distribution $P(Se|Kn, \pi)$ is the goal, and requires information from all related random variables in $\{Kn, Un, \pi\}$, which in the Bayesian network are visualized as parents of Se . This distribution can always be obtained using Lebeltel and Bessière, 2008

$$P(Se|Kn, \delta, \pi) = \frac{1}{\Sigma} \sum_{Un} P(Se, Un, Kn | \delta, \pi) \quad (3.5)$$

where $\Sigma = \sum_{\{Se, Un\}} P(Se, Un, Kn | \delta, \pi)$ acts as a Normalization term. To complete the inference calculation, we only need to reduce the distribution $\sum_{Un} P(Se, Un, Kn | \delta, \pi)$ into marginally independent factors that can be determined. We assume that independence is denoted by the structure of the Bayesian network, so we only need be concerned with the ancestors of Se and do not need to scale by Σ . Given that inference into a Bayesian network typically involves querying a single node, we will assume that Se is the singleton $Se = \{X\}$. This can also be accomplished if Se is larger by making X a parent of all nodes in Se .

Applying the chain rule again to Bayesian networks, we can walk the Bayesian network backwards through the directed edges from X , determining the conditional distribution of each node from its parents as we go, and therefore breaking down the determination of the joint distribution into smaller, separate calculations. Considering Z to be the parents of each ancestor node Y and following the method of Equations 2.6, and Equation 2.7, and Equation 3.5, a general expression for the factorization of $P(Se|Kn, \delta, \pi)$ through the Bayesian network is

$$P(Se|Kn, \delta, \pi) = \sum_{Y \in \{X, \text{An}(X)\}} \left[\prod_{Z \in \text{Pa}(Y)} P(Y|Z)P(Z) \right]. \quad (3.6)$$

4. Bayesian Network Implementation

For our Bayesian network implementation, with the random variables in Se , Kn , and Un internally linked together as nodes and edges. Nodes associated with actions to be taken typically have conditional distributions that act as “questions” regarding their operational state. If for example, the question is asked what should the right-side motors do?, the network nodes related to obstacle presence and mapping, and in turn, prior sensor data, will have to be traversed to obtain the Answer, which is a posterior probability distribution that is used to select a particular value of motor speed. Unlike most Bayesian network implementations, our implementation is unique in that it uses fixed-point math for storage and calculation and is programmed in C for better portability and calculation efficiency on small-scale embedded systems that do not have a floating-point arithmetic unit Lauha, 2006.

At minimum, a random variable with N possible values will have a $1 \times N$ distribution matrix. A parent node with M values will create at most M additional distributions, and in general, if each parent has a distribution N_l values in size, and there are L parents, then the number of distributions M possible in the child node are $M = \prod_{l=1}^L N_l$, so the storage size of the node scales roughly as N^L . This can be mitigated to make processing more efficient by designing a deeper graph with more nodes and less parents per node. A parent node with an $M_l \times N_l$ distribution matrix, regardless of the number of parents and the size of M_l , will still only contribute N_l values to its child nodes. A given node X will then have to store a table of size $|V(X \cup Pa(X))|$. Total probability requires that all values in each row m sum to 1. The joint distribution P of probability values associated with a given random variable is the content actually stored.

Because the dimensionality of data stored in the node changes with number of parents, we use a single array for storing the distribution and index it with a linear index that is a function of the parent numbers of the node. To create the index, we model the array as an $L + 1$ -dimensional matrix for L parents and calculate an index i . For a two-dimensional row-major-order matrix with row (m) and column (n) indices, $i = n + m * columns$. By recognizing that each additional dimension must be indexed by multiplying past the sizes of all preceding dimensions, we can set index i using matrix indices m_1 for dimension 1 of size M_1 (we choose columns here for consistency), m_2 for dimension 2 of size M_1 (we choose rows here for consistency), and m_3, m_4, \dots and above for additional matrix dimensions of size M_3, M_4, \dots respectively, obtaining

$$i = m_1 + m_2 M_1 + m_3 M_2 M_1 + \dots + m_{L+1} \prod_{l=1}^L M_l = \sum_{n=1}^{L+1} \left(m_n \prod_{l=1}^{n-1} M_l \right). \quad (4.1)$$

5. Sensory Reasoning Example

To test the concepts and implementation detailed above, a simple six-node network was set up which is graphed with base probabilities as shown in Figure 2, consisting of three nodes representing obstacle sensors with states $\{Clear, Obstacle\}$ angled 30 degrees apart in front of a mobile robot, a sensor fusion node that uses Gaussian distributions to estimate the detection angles of the three sensors, and two motor nodes for control of left and right side motors. A MAP query is used to determine the most appropriate state of the motors from $\{Forward, Reverse\}$ given inference into the sensor data.

It is assumed that each obstacle sensor has only a 20% prior probability of seeing an obstacle in practice and 80% probability of no obstacles, but this is updated when the

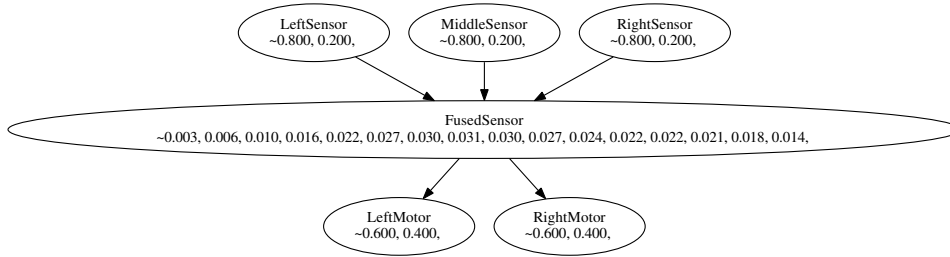


FIGURE 2. Simple Bayesian Network for Obstacle Avoidance Program

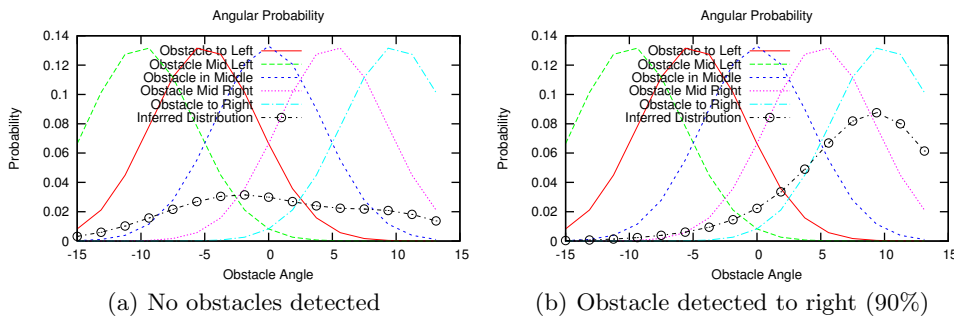


FIGURE 3. Sensor Fusion Probability Priors and resulting Inferred Distributions

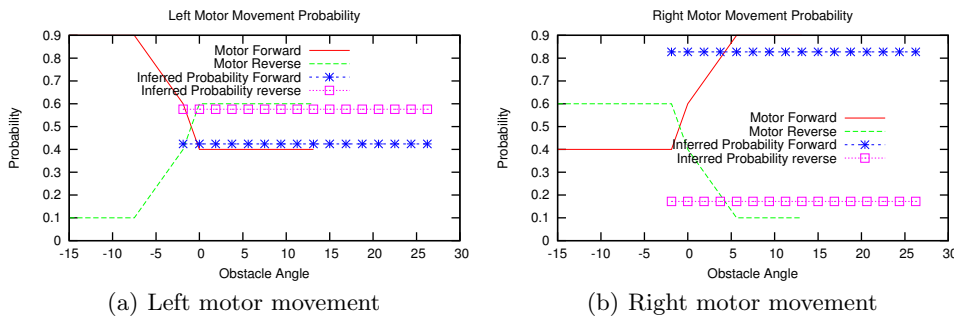


FIGURE 4. Priors for Motor Movement and Probabilities of Obstacle Avoidance for Obstacle Detected to Right (90%)

sensor actually detects an obstacle. The sensor fusion node is pre-loaded with a set of six Gaussian functions representing an approximate likelihood of the direction of an obstacle given the six possible states of the three two-state obstacle sensors. Figure 3 shows the angular probabilities of these Gaussian functions and the resulting inferred distribution (denoted with markers) for (a) no obstacles detected, and (b) an obstacle detected by the right-hand sensor with 90% probability.

The motor nodes are defined with a high probability of reverse movement for obstacles detected on the opposite side to cause obstacle avoidance behaviour, but generally a higher probability of forward movement to keep the robot moving forward if obstacles are unlikely. The functions used are shown in Figure 4 for the right and left motors

This function could be constructed in any way desired to reflect the response of the motors to likelihood of an obstacle, but a simple function is used in this case to illustrate that the system still works. The presence of an obstacle increases the inferred “prob-

ability” of the left motor reversing, which in this context would be contextualized as “probability that reversing will avoid an obstacle”, and the motor will reverse when a MAP query is applied. Optimal priors can be determined by expert knowledge or by statistical learning methods which will be a focus of future work. Graphing the priors used for inference at each node shows a notable similarity to fuzzy functions, and the process of inference into discrete values that of fuzzification, but there are still fundamental differences. A random variable is expected to have a single outcome while fuzzy sets represent multiple outcomes, or alternately imprecise degrees of belief. This imprecision is reflected in the extension to fuzzy sets of possibility theory, which can be seen as a form of upper probability theory Zadeh, 1999.

6. Conclusions

We have described a method for making decisions using the Bayesian Robot Programming paradigm using Bayesian networks for organization of data and efficient algorithms for storage and inference into discrete random variables. This method has been implemented as a novel, efficient, structured framework for practical use of Bayesian networks and probabilistic queries on an embedded system with fixed-point arithmetic, leading to an efficient and intuitive way of representing robotic knowledge and drawing conclusions based on statistical information encapsulated in the network.

The applicability of these methods to robotic programming and control is very wide, and many future applications are expected, as this method can be extended to much more complex systems simply by adding networked random variables for sensor and actuator quantities, so long as they can be probabilistically characterized. Future work includes implementations for new applications, determination of network structure and optimal priors from existing hardware data, and the implementation of statistical learning methods to improve performance of the system while it operates.

REFERENCES

- Bessiere, P., Lebeltel, O., Lebeltel, O., Diard, J., Diard, J., Mazer, E., and Mazer, E. (2000). Bayesian robots programming. In *Research Report 1, Les Cahiers du Laboratoire Leibniz, Grenoble (FR)*, pages 49–79.
- Engelbrecht, A. P. (2002). *Computational Intelligence: An Introduction*. John Wiley & Sons.
- Kjrluff, U. B. (2008). *Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis*. Springer Science.
- Koller, D. and Friedman, N. (2009). *Probabilistic Graphical Models, Principles and Techniques*. MIT Press, Cambridge, Massachusetts.
- Lauha, J. (2006). The neglected art of fixed point arithmetic. Presentation.
- Lebeltel, O. and Bessière, P. (2008). Basic Concepts of Bayesian Programming. In *Probabilistic Reasoning and Decision Making in Sensory-Motor Systems*, pages 19–48. Springer.
- Lebeltel, O., Bessière, P., Diard, J., and Mazer, E. (2004). Bayesian robot programming. *Autonomous Robots*, 16(1):49–79.
- Lebeltel, O., Diard, J., Bessiere, P., and Mazer, E. (2000). A bayesian framework for robotic programming. In *Twentieth International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering (MaxEnt 2000)*, Paris, France.
- Post, M. (2014). *Planetary Micro-Rovers with Bayesian Autonomy*. PhD thesis, York University.
- Wu, D. and Butz, C. (2005). On the complexity of probabilistic inference in singly connected bayesian networks. In Slezak, D., Wang, G., Szczuka, M., Dntsch, I., and Yao, Y., editors, *Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing*, volume 3641 of *Lecture Notes in Computer Science*, pages 581–590. Springer Berlin Heidelberg.
- Zadeh, L. A. (1999). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets Syst.*, 100:9–34.