Autonomous Vacuum Cleaners Must Be Bayesian

R. Wade Brittain,
Bruce D'Ambrosio
Department of Computer Science
Oregon State University
Dearborn Hall 303
Corvallis, OR 97331-3202

Abstract

Due to limited power supplies, finite dust storage, poor location sensing, and random failures caused by sucking up numerous small objects, autonomous vacuum cleaners require accurate and timely information about environmental parameters and the state of internal components. The best methods for providing such estimates are Bayesian. The recent advent of compact, efficient probabilistic representations has made the implementation of these methods not only more tractable, but has also extended the range and scope of architectures in which they can be implemented. We explore the anomalies of the vacuuming domain and illustrate how Bayesian techniques can be applied to reduce their effects on agent performance.

1 Introduction

For long-range autonomous robots to be successful, they must be able to maintain accurate information about their location, available resources, and the state of critical components. Autonomous vacuuming agents are no exception. Under the assumptions that the physical architectures for such agents include an on-board, self contained power supply, (to avoid the cord-entanglement problem), and sensors which may be prone to failure and noise, the basic problem is one of tracking and resource management. An accurate and timely representation of critical parameters and system state must be maintained if the task is to be accomplished with any efficiency or reliability. Bayesian techniques provide the most effective and accurate means to maintain this information. Most of these techniques were developed in the early 1960's. However, many complex control algorithms based on them were infeasible to implement due to the exhaustive computation required. The advent the *belief network* in recent years has made probabilistic inference more tractable. Belief networks, (also called influence diagrams if a decision model is represented), represent a probabilistic problem model as a directed acyclic graph, thereby explicitly capturing much of the problem structure. This can dramatically reduce the size of the model representation, ease the assessment of model parameters and permit more efficient computation[12]. Until recently, belief network representations were strictly homogeneous, (i.e. either all variables were discrete or all were continuous.) However, several researchers have extended this ontology to allow for the representation of discrete and continuous variables within the same probabilistic model[10, 11, 7, 3]. This allows for a direct connection between the signal processing interface, the abstract symbolic representation of system state and continuous models of available resources. Since models of continuous dynamics of physical systems are themselves dependent upon system state , this greatly extends the representational capability and accuracy of Bayesian networks. Such accurate representation is crucial to the efficient performance of complex autonomous agents operating in uncertain domains.

In addition, the mixed discrete/continuous Bayesian network can provide both integration with low level controllers, (either by actually providing refined sensory input, or by simply observing system inputs and outputs to determine if the system is working correctly), and higher-level decision-theoretic reasoning [13, 14]. Bayesian methods can therefore be used in conjunction with a purely subsumptive architecture[4] , or one with more centralized processing, depending upon the level of integration desired. We follow with a description of the vacuuming task, an outline of Bayesian methods, and a discussion of how these methods might be applied in an autonomous vacuuming agent.

2 The Task

Given a rechargeable vacuum, ANY agent, (human or otherwise), must deal with the following problems:

1. Adequately removing the dirt from the floor.
2. Keeping the vacuum charged.
3. Keeping the vacuum inlets and turbines free from clogs
4. Emptying the dust bag when full.

Autonomous vacuums of this sort must deal with all of these and face the additional difficulties:

- Location sensors are significantly poorer than those of human agents.
- failure to address problem 2 above will render the agent completely unable to finish the vacuuming task or recharge.

At a minimum, the task of adequately removing the dirt from the living room floor requires the agent to insure that its path when vacuuming is such that all regions of the floor area have been covered at least once. This requires a knowledge of the rooms topography and continuous tracking of the vacuum's location in order to determine when the task is complete. Therefore, agents who do not condition their behavior on their beliefs about the past will perform poorly in this domain. Unfortunately, in the case of the automatic agent, the ability to track is made difficult by noisy, failure prone sensing devices. The task is further complicated by stationary obstacles and intermittent, dynamic clutter, (pets, people, and the junk they move around), – which can confound the agent by changing the nature of sensor returns in regions which have already been visited. While performing this task, the agent must also monitor it's internal resources, (remaining charge on power supply, space available in dust bag), and the state of internal components, (intakes, turbines, locomotive affectors etc.). These must be managed effectively in order to insure that the vacuum is continuously removing as much dirt as possible and that the system will remain functional so that the task can be completed. Sensor noise, failure, and task complexity make purely reactive behavior strategies ineffective – a more sophisticated approach is required.

3 Solution Method

The task is one in which objective accomplishment, tracking, and resource management all must occur simultaneously under uncertain conditions. We submit that Bayesian tracking methods and Bayesian decision theory are the best methods applicable for the following reasons:

1. Bayesian tracking methods are optimal in the linear case and can be easily extended to low-order, non-linear systems[1].
2. Bayesian tracking methods incorporate knowledge of system dynamics with sensor inputs, thus providing some robustness to noise and other sensor anomilies.
3. Bayesian decision theory is ideally suited for action selection when resource management/task accomplishment tradeoffs must be considered under conditions of uncertainty.
4. Bayesian decision theory has been successfully used in simple autonomous robots to construct internal representations of indoor environments with dynamic clutter[6, 9].
5. Mixed discrete/continuous Bayesian networks[10, 7, 11, 3] can provide a unified representation of Tracking, Fault-detection and Decision models and allows for more efficient, query-driven computation of agent beliefs.

We will now discuss specifically how these methods can be applied in the task of vacuuming the living room floor.

3.1 Fault Detection

Fault detection is a crucial element of the vacuum cleaning task. (e.g. Is the bag full? Is the beater bar working properly? Is the turbine clogged?). The development of efficient mixed discrete/continuous probabilistic representations allows for the conditioning of dynamical models on the corresponding assumptions about system state – thus enabling tracking and fault monitoring to proceed simultaneously. To detect faults, one simply makes a prediction about the value which an observable parameter should have, given the existing conditions and current state model of the system, and compares it with the actual observation via a statistical test. The prior probabilities for the state of the system assumed by each dynamic model can then be updated by taking the test results to be an observed posterior probability. For example, we can predict the expected rpm of the vacuum motor given it's consumption of electrical current. If the observed rpm is too high, we should consider the possibility that the intake is clogged, as the observed model is consistent with the prediction made by a dynamical model which also made this base assumption. If, on the other hand, the observed rpm is to low, we should instead consider that the turbine might be jammed, (again, because the consistent model made this assumption.) Integrating the continuous predictive models with discrete assumptions about system state allows the agent to update it's beliefs in it's state, provided that an effective test is used to compare predictions and observations. The models typically used to represent continuous parameters are Gaussian distributions. Therefore the test that's often used to compare the predicted and observed values is a likelihood ratio test. However, we cannot always detect faults as single valued observations. For example, how much electrical current does a shorted electric motor draw? The answer depends on the system in question and the severity and location of the short. Nonetheless, for any system, there is a *minimum* value for this parameter beyond which an electrical short may indeed be the cause. In such cases, we advocate the use of *the large sample test for the difference between two means.* This test allows for the specification of parameters as upper bounds, lower bounds and even intervals bounded by white noise. A more extensive discussion of the use of this test for fault detection can be found in[3].

The fault detection schemes we have described thus far are ideal for support of traditional adaptive control methods [1] since the parameter estimates for each plant model have already been computed by detection time. However, if adaptive control is not being used, the technique may be viewed as rather expensive and unnecessary. In many cases, it would suffice to simply recognize that a given parameter is beyond allowable limits and initiate some sort of recovery action. To facilitate this, Belief networks allow for efficient access of any parameter. This enables the monitoring of any critical value, measured, predicted or estimated, and also provides the probability that the estimate is valid. This latter information can be used by detection and recovery routines to determine if the situation requires further investigation or action.

3.2 Tracking

Kalman filtering is a parameter estimation technique used to provide better estimation and reduce uncertainty in the presence of sensor noise[8]. It works by using the plant model and a previous state estimate to produce a prediction about the next observed state, and then refining the actual observation with that estimate. When the plant model is linear, the method has been shown to be optimal[1]. It is, however, not restricted to linear applications, and has been successfully applied to such non-linear problems as aircraft control, submarine ballast system control, spacecraft control, tracking an airborne object from the ground and visually tracking objects on the ground from a moving aircraft. We feel that this method will be more than sufficient for estimating the position of a small robot in a living room. We do, however, propose an inhancement for the other tracking tasks involved in vacuuming: Namely, that the upper bound, lower bound, and interval types, (mentioned in the previous section), be included as possible parameter types. For example, predictive white noise models are derivable for the amount of air passing the turbine in a given instant, but how much dirt did the air contain? (i.e. What is a reasonable expectation for the level of amount of dirt entering the dust bag as a function of time?) We submit that this is better represented by an interval bounded by white noise than a single Gaussian. A subsequent filtering operation with a sensory observation on the level of dirt in the bag would only refine the observation to the extent that it exceeds the amount of dirt that the vacuum could possibly suck up in the time elapsed since the last estimate. An example of how such an enhancement could be implemented can be found in [3]. Other than this minor anomaly, we feel that Kalman filter-based tracking algorithms are already more than adequate for parameter estimation in autonomous vacuuming agents, and could easily scale to more complex tasks (mowing the lawn?)

3.3 Support for Control

3.3.1 Control Architectures

Bayesian parameter estimation techniques can be used with almost any agent architecture. For purely subsumptive architectures, this requires that each control level employ some means of stochasticly filtering it's input. Within the limitations of this architecture[2], this is feasible. Should centralized sensor processing be required, this too is directly supported by existing Bayesian methods, specifically belief networks, which allow efficient access to any parameter. A controller need therefore access only the data that it needs to perform it's task – and networks can be designed such that different parameters can be updated at different rates. Mixed discrete/continuous Bayesian networks extend this utility further by providing discrete, continuous, and symbolic data. They can therefore be used in architectures consisting of multiple heterogeneous controllers, provided some form of synchronization is implemented to insure timely distribution of the data and avoid resource contention deadlocks.

3.3.2 Decision Theoretic Control

Lastly, complete decision-theoretic control models can be embedded in Bayesian networks of the type we have described[5] . Decision theoretic control is ideally suited for action selection in agents operating in uncertain and risky environments with limited resources. This is exactly the situation in the autonomous vacuuming domain. The agent has a limited power supply, limited dust bag capacity, and limited sensing ability. Additionally, every action has associated has associated with it some sort of risk or utility. The agent must constantly decide whether to continue removing dirt from the floor, empty the bag, recharge it's batteries, or verify it's position. The situation is further complicated by the existence of objects whose position may be transient (People, pets, and any objects they might relocate), thus confounding the agents navigational ability. Getting lost means the agent will be unable to recharge or complete the task. However, maintaining good estimates of it's location often requires some extra activity to pinpoint landmarks and reference points. It might even require returning to the starting point and "recalibrating" position estimates.

This kind of activity is at odds with the task of actually removing the dirt from the floor, but is often necessary if the task is to be successfully completed. To operate efficiently, an agent must be able to weigh the costs of position orientation activities, (in terms of the work that could be done), against the resulting probable increase in position estimate accuracy. The problem is therefore one involving risks and uncertainty. We do not believe that planning is necessary or desirable in this domain, but we do believe that explicit treatment of uncertainty and risk is necessary. These are specifically embodied in Bayesian decision theory[15].

4 Summary

We have shown that autonomous vacuuming requires reliable tracking, state estimation and efficient resource management under uncertain conditions. We have enumerated the Bayesian approaches for performing or supporting each of these tasks. We have noted the proven optimality of some of these methods, and have described new methods which allow for the compact representation and efficient computation of probabilistic models. We have shown how each of these techniques can be applied towards the task of vacuuming the living room floor, and have shown that they can be used in almost any architecture an autonomous vacuuming agent may have.

5 Conclusion

For the successful and efficient completion of many real-world tasks, agents require an accurate description of physical parameters and system state. Optimal performance is dependent upon timely situation recognition – And timely recognition is dependent upon precise parameter estimation. This is particularly true for autonomous vacuuming agents, which must deal with obstacles, small power supplies, finite size dust bags, and different floor surfaces. Bayesian methods can provide timely compact, efficient, accurate (and sometimes optimal) estimates of these items. Furthermore, they can be implemented in almost any agent architecture. Therefore, whatever approach is employed in autonomous vacuum cleaner design, the best of such agents will be Bayesians.

References

[1] B Anderson and J Moore. *Optimal Filtering.* Prentice Hall, 1979.

[2] J. G. Bellingham, T. R. Consi, R. M. Beaton, and W. Hall. Keeping layered control simple. In *Proceedings of AUV-90*, pages 3 – 8. Symposium on Autonomous Underwater Vehicle Technology, June 1990.

[3] R. W. Brittain. Probabilistic tracking and real-time assessment of physical systems. Technical report, Technical Report, CS Dept., Oregon State University, 1993.

[4] R. Brooks. A layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, 1986.

[5] G. Cooper. A method for using belief networks as influence diagrams. In *Proceedings of the 1988 Workshop on Uncertainty in AI*, pages 55–63. Association for Uncertainty in AI, August 1988.

[6] Thomas Dean, Kenneth Basye, Robert Chekaluk, Seungseok Hyun, Moises Lejter, and Margaret Randazza. Coping with uncertainty in a control system for navigation and exploration. In *Proceedings AAAI-90*, pages 1010–1015. AAAI, 1990.

[7] Chang K. and Fung R. Symbolic probabilistic inference with both discrete and continuous variables. *Awaiting Publication*, 1992.

[8] R. E. Kalman. A new approach to linear filtering and prediction problems. *J. Basic Eng., Trans. ASME, Series D*, 82(1):35–45, 1960.

[9] J Kirman, K. Bayse, and T. Dean. Sensor abstractions for control of navigation. In *Procceedings 1991 Int. Conf. on Robotics and Automation*, pages 2812 –2817, 1991.

[10] S.L. Lauritzen. Mixed graphical association models. Technical report, Department of Mathematics and Computer Science, University of Aalborg, Denmark, July 1988.

[11] Kristian G. Olesen. Causual probabilistic networks with both discrete and continuous variables. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(2), March 1993.

[12] J. Pearl. *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufmann, Palo Alto, 1988.

[13] S. Russell. Principles of metareasoning. In *First International Conference on Principles of Knowledge Representation and Reasoning.* AAAI, 1989.

[14] S. Russell and E. Wefald. *Do the Right Thing.* MIT Press, 1991.

[15] L. Savage. *The Foundations of Statistics.* Dover Publications, 1972.