

Efficient Mapping Through Exploitation of Spatial Dependencies

Yaron Rachlin[†], John M. Dolan[‡], and Pradeep Khosla^{† ‡}
Department of Electrical and Computer Engineering[†], The Robotics Institute[‡]
Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213, USA
rachlin@ece.cmu.edu, jmd@cs.cmu.edu, pkk@ece.cmu.edu

Abstract—Occupancy grid mapping algorithms assume that grid block values are independently distributed. However, most environments of interest contain spatial patterns that are better characterized by models that capture dependencies among grid blocks. To account for such dependencies, we model the environment as a pairwise Markov random field. We specify a belief propagation-based mapping algorithm that takes these dependencies into account when estimating a map. To demonstrate the potential benefits of this approach, we simulate a simple multi-robot minefield mapping scenario. Minefields contain spatial dependencies since some landmine configurations are more likely than others, and since clutter, which causes false alarms, can be concentrated in certain regions and completely absent in others. Our belief propagation-based approach outperforms conventional occupancy grid mapping algorithms in the sense that better maps can be obtained with significantly fewer robot measurements. The belief propagation algorithm requires a modest amount of increased computation, but we contend that in applications where significant energy and time expenditure is associated with robot movement and active sensing, the reduction in the required number of samples will justify the increased computation.

Index Terms—occupancy grids, belief propagation, markov random fields, demining

I. INTRODUCTION

Occupancy grid mapping algorithms [1] are widely used, in part due to their computational efficiency. This computational efficiency is achieved by breaking the high-dimensional task of estimating a grid into a set of one-dimensional problems, by assuming that all grid block values are independently distributed. While such an assumption is useful from a computational perspective, occupancy grid methods do not take advantage of spatial dependency in the environment. Most environments (e.g. hallways, natural terrain, mine shafts) contain spatial structure, and thus measuring one area in the map gives one information about adjacent areas in the map. By failing to exploit this structure, occupancy grid-based approaches have to explicitly sense the entire environment they seek to map. Spatial dependencies can be modeled, and in this paper we seek to do this through a Markov random field [2]. Given such an environment, where the content of a grid block is dependent on the contents of other grid blocks, the occupancy grid algorithm proves inefficient in the number of samples it requires to produce an accurate map. In this paper we demonstrate that in a structured environment, a belief propagation-based [3] algorithm that takes local

dependencies into account can outperform the occupancy grid algorithm in the sense that it will achieve more accurate maps with significantly fewer samples of the environment. We seek to demonstrate this potential in the context of mapping a minefield using multiple robots.

Unmapped minefields constitute an unresolved humanitarian threat. The amount of time required to clear a minefield is directly proportional to the false alarm rate [4]. High false alarm rates are also a primary cause of operator fatigue, which can result in deminer death and injury. Despite unresolved challenges, risk for human deminers and imprecise detection methods indicate the potential benefit of a robotic solution to minefield mapping. Given that demining robots operate outdoors, there exist localization solutions such as [5]. Thus in this paper we choose to focus on the problem of mapping assuming good estimates of robot position are available. To demonstrate our approach we introduce a spatial Markov random field model and apply it to a minefield mapping scenario in Section III. We introduce the belief propagation algorithm corresponding to this model in Section IV, and present results comparing the performance of occupancy grids and belief propagation in Section V. Our results show that belief propagation requires significantly fewer samples to construct maps of a similar quality to the ones produced by occupancy grids. This performance gain is achieved with a modest increase in computational burden, which we will characterize in Section V. In applications such as robotic minefield mapping where significant energy and time expenditures are associated with each robot sample (e.g., movement, active sensing), and assuming that fast, energy-efficient microprocessors are available, the significant reduction in the required number of robot samples will offset the additional processing required to run the belief propagation algorithm.

II. RELATED WORK

[6] used Markov random fields in order to combine local maps into a consistent global map. In this paper we apply Markov random fields to characterize the entire map, not only the interactions among local maps. [7] explores the problem of estimating a grid directly by taking into account spatial dependencies in the grid. However, the dependencies are introduced by sensors that observe multiple grid blocks at once, and are not due to the structure of the environment.

[8] develops efficient robot coverage strategies for minefields by estimating a mine emplacement pattern online in order to make exploration more efficient. This work focuses on estimating grid-like patterns of mine emplacement. By using Markov random fields to model the overall structure of an environment, we model both mine emplacement patterns and the pattern of clutter, which is the primary source of false alarms in the environment.

[9] proposed the use of the belief propagation algorithm for robotic exploration, but did not explore the issue of spatial dependence in the environment. In our work, we explicitly combine a model that captures environmental structure with a belief propagation algorithm.

III. SPATIAL MODEL

We denote random variables and functions by upper-case letters, and instantiations or constants by lower-case letters. Bold-font denotes vectors, and bold-font upper-case letters denote random vectors.

A. Markov Random Fields

We model the environment as a discrete grid where each grid block F_i takes a value from a finite alphabet \mathcal{F} . We assume that the grid \mathbf{f} occurs with probability $P_{\mathbf{F}}(\mathbf{f})$ and is assumed to be distributed as a pairwise Markov random field (also called an auto-model) [2]. This is a widely used model that is popular because it is the simplest form of Markov random field that still conveys spatial contextual information. Its simple form allows for efficient algorithms. This model differs from the equiprobable i.i.d. distribution associated with conventional occupancy grid models, and allows one to model environments with structure such as sparsity, likely configurations, and spatial contiguity. More complex Markov random field models could be specified to model environments with more complex dependencies, though such models would entail more complex algorithms.

A pairwise Markov random field is modeled as a graph. Each grid block F_i in grid \mathbf{F} is represented as a node and is connected to the grid blocks immediately adjacent to it as shown in Figure 1. The set of grid blocks connected to grid block F_i are written as \mathcal{N}_i . This field has two properties. First, we assume that all grid configurations have positive probability. Second, the pairwise Markov field implies that, conditioned on the grid blocks to which it is connected, the probability of a grid block is independent of the remaining grid blocks. According to the Hammersley-Clifford theorem, a Markov random field that obeys these two properties is distributed as a Gibbs distribution [2]. A Gibbs distribution is written as a normalized product of positive functions over the cliques in the graph of the Markov random field. In our pairwise Markov field model there are two types of cliques: single nodes (F_i) with associated function Ψ_i , and pairwise cliques ($F_i, F_j, j \in \mathcal{N}_i$) with associated function Ψ_{ij} . We have the following Gibbs distribution for \mathbf{F} ,

$$P_{\mathbf{F}}(\mathbf{f}) = \frac{1}{Z} \prod_i \Psi_i(f_i) \prod_{j \in \mathcal{N}_i} \Psi_{ij}(f_i, f_j) \quad (1)$$

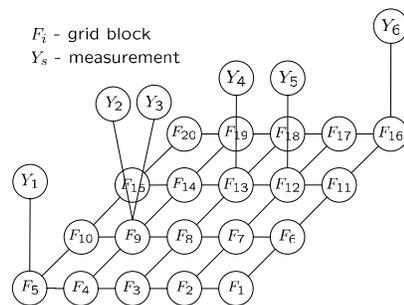


Fig. 1. Graph of Markov random field.

where Z is a normalization constant. The meaning of Ψ_i and Ψ_{ij} is application-dependent, but intuitively these functions indicate the likelihood of a grid block being in a particular state or the likelihood of a pair of adjacent grid blocks being in a particular configuration, respectively.

B. Minefield Model

We apply the above model to the task of mapping a minefield. In order to test the efficacy of mine detection technologies, the US Army maintains a test site at Fort A.P. Hill, VA consisting of a 20m by 49m test grid, where each grid square contains nothing, a landmine, or clutter [4]. Clutter consists of objects that confuse a mine detection technology and lead to higher false alarm rates. For Ground Penetrating Radars (GPR), a promising mine detection technology [4], clutter consists of objects that induce soil inhomogeneity such as rocks, roots, and pockets of wet soil. In light of this example, we let $\mathcal{F} = \{0, 1, 2\}$, whose values correspond to nothing, a landmine, and clutter, respectively. We specify $\Psi_i(f) = \frac{1}{|\mathcal{F}|}$, where $|\mathcal{F}|$ is the cardinality of \mathcal{F} , in order to highlight the effect of spatial dependency as opposed to an a priori state preference. We specify Ψ_{ij} as follows:

$$\Psi_{ij} = \begin{pmatrix} 0.6 & 0.3 & 0.1 \\ 0.475 & 0.05 & 0.475 \\ 0.1 & 0.3 & 0.6 \end{pmatrix} \quad (2)$$

Each row indicates the likelihood of adjacent grid block value configurations. For example, the preference of a landmine being adjacent to a landmine is 0.05. Figure 2 shows an example of two sample minefield grids corresponding to the occupancy grid assumption and Markov random field model, with Ψ_i and Ψ_{ij} specified above. The grid generated by sampling our Markov Random field model exhibits significant spatial structure, with clutter clustered together and landmines spatially separated. In non-simulation-based applications, the parameters of Ψ_i and Ψ_{ij} should be empirically estimated based on samples from the environment using algorithms such as the ones described in [2]. We chose values of Ψ_i and Ψ_{ij} that clearly demonstrate the effect of spatial structure on mapping.

In our experiments, robots are randomly placed in a grid world and sample the grid block where they are located. The robot reports a noisy measurement Y_s of the observed grid block F_i according to the noise model $P_{Y|F}$ as shown in

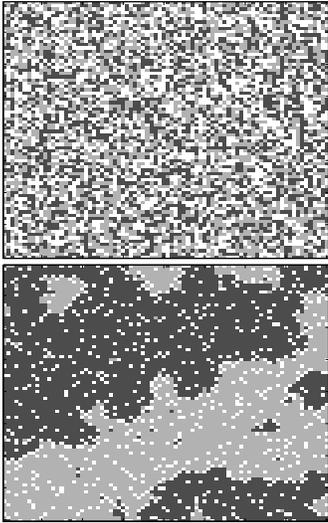


Fig. 2. Two randomly generated 80×80 grids. The top grid corresponds to an occupancy grid assumption about the environment, where the value of each grid block is chosen independently and all values are equiprobable. The bottom grid corresponds to the probabilistic assumption associated with a pairwise Markov random field. Dark gray corresponds to nothing, light gray corresponds to clutter, and white corresponds to landmines.

Figure 1. The subscript s indexes the measurements, and \mathcal{S}_i is the set of measurements associated with grid block F_i . After sampling, each robot moves one block (in a randomly chosen direction) and samples again. This process is continued for a fixed amount of time. We assume our robot is equipped with a ground penetrating radar. The GPR registers higher false alarms on clutter than on nothing (i.e., no clutter). After such an exploration process, each grid block has a finite number of sensor samples associated with it. We write the conditional probability of a grid \mathbf{f} given the measurements \mathbf{y} as follows,

$$P_{\mathbf{F}|\mathbf{Y}}(\mathbf{f}|\mathbf{y}) \propto \frac{1}{Z} \prod_i \Psi_i(f_i) \prod_{j \in \mathcal{N}_i} \Psi_{ij}(f_i, f_j) \cdot \prod_{s \in \mathcal{S}_i} P_{Y|F}(y_s|f_i) \quad (3)$$

where Z is a normalization constant, and $P_{Y|F}$ corresponds to the sensor noise model. Since in our demining scenario each robot observes a single block in the grid at a time and since we assume the world is static, the measurement conditional on the grid block to which it corresponds is independent of the other measurements. However, any two unconditioned measurements of separate grid blocks are generally not independent since the two grid blocks to which they correspond are dependent. This differs from a conventional occupancy grid approach where, according to the mathematical model used, while two unconditioned measurements are dependent if they correspond to the same grid block [10], two unconditioned measurements that correspond to two adjacent grid blocks are considered independent.

IV. MINEFIELD MAPPING ALGORITHMS

Given a vector of measurements \mathbf{y} , a mapping algorithm seeks to infer which map \mathbf{f} best corresponds to the envi-

ronment being mapped. We first review the occupancy grid approach to this problem, and then introduce a loopy belief propagation-based algorithm for addressing this problem.

A. Occupancy Grid

In the occupancy grid approach, the problem of finding the best map \mathbf{f} is broken down into a series of one-dimensional problems. The world is represented as a binary grid, where each grid block contains the log odds ratio of the probabilities of that grid block equalling one or zero given the measurements. The log odds ratio form is specialized for a binary grid. We write a more general form of the occupancy grid below that accommodates non-binary grids, while still maintaining the essential occupancy grid assumptions of independent grid blocks and conditional independence of measurements given the grid block to which they correspond. Each grid block F_i has the following a posteriori probabilities associated with it for each $f \in \mathcal{F}$,

$$\begin{aligned} P_{F_i|\mathbf{Y}}(f|\mathbf{y}) &= P(f|y_s \in \mathcal{S}_i) \\ &\propto P(y_s \in \mathcal{S}_i|f)P_{F_i}(f) \\ &= \prod_{s \in \mathcal{S}_i} P_{Y|F}(y_s|f)P_{F_i}(f) \end{aligned} \quad (4)$$

Typically these probabilities are written as a recursively updatable log odds ratio for each binary grid block. The above form is equivalent to the recursive form given the assumption of the occupancy grid model. Given the set of probabilities associated with a grid block, we can estimate its contents $\hat{f}_i = \arg \max_{f \in \mathcal{F}} P_{F_i|\mathbf{Y}}(f|\mathbf{y})$.

B. Loopy Belief Propagation

The belief propagation algorithm [3] allows one to efficiently compute marginal probabilities for a set of random variables with a tree dependency structure. In our case, we seek to compute the marginal probability for each grid block in our map, conditioned on all the robot measurements. All measurements are relevant for each grid block since our environment has spatial dependence among grid blocks. Thus observations of other grid blocks provide information about the grid block of interest. Because our Markov random field graph contains loops, computing marginal probabilities is difficult. To conduct inference on such a graph, we apply belief propagation to our graph as if it did not contain loops. This is a common approximation referred to as loopy belief propagation. Though this algorithm is only approximate for graphs with cycles such as our model, loopy belief propagation has proven extremely effective in conducting inference in other graphs with cycles, such as the graphs associated with high-performance error-correcting codes [11]. Belief propagation has other advantages for a distributed mapping scenario. By its structure belief propagation is amenable to distributed implementation. In addition, it provides a consistent probabilistic framework for integrating observations from multiple types of sensors. We describe the belief propagation algorithm for our minefield mapping application below.

TABLE I
NOISE MODEL FOR A GPR SENSOR IN A BINARY GRID.

$P_{Y F}$	$Y = 0$	$Y = 1$
$F = 0$	0.95	0.05
$F = 1$	0.05	0.95

The belief propagation algorithm is a message-passing algorithm. At each iteration, each node in the Markov random field sends different messages to its neighboring nodes based on the messages it received in the previous time step. The message $m_{ij}(f)$ from F_i to F_j consists of a probability distribution over the states of F_j and is thus specified as a function of $f \in \mathcal{F}$, as follows,

$$m_{ij}(f) \leftarrow \alpha \sum_{f' \in \mathcal{F}} \Psi_i(f') \Psi_{ij}(f', f) \prod_{s \in \mathcal{S}_i} P_{Y|F}(y_s | f') \cdot \prod_{k \in \mathcal{N}_i \setminus j} m_{ki}(f') \quad (5)$$

where α is a normalization constant, and $\mathcal{N}_i \setminus j$ corresponds to the set of neighbors of F_i excluding F_j . We run the algorithm until either the algorithm converges, or until the maximum number of iterations is exceeded. We determine convergence by computing the average change in each node's estimated probability distribution or belief $b_i(f)$. The beliefs are computed as follows,

$$b_i(f) \leftarrow \alpha \Psi_i(f) \prod_{j \in \mathcal{N}_i} m_{ji}(f) \prod_{s \in \mathcal{S}_i} P_{Y|F}(y_s | f) \quad (6)$$

where α is a normalization constant. Given the set of beliefs associated with a grid block, we can estimate its value $\hat{f}_i = \arg \max_{f \in \mathcal{F}} b_i(f)$. It is interesting to compare (6) and (4). Both expressions contain a product over the measurement probabilities and the a priori likelihood of the grid block being in a particular state, but (6) also contains a product over estimates from adjacent grid blocks.

V. SIMULATION RESULTS

Since we want to highlight the effect of spatial dependencies on the performance of mapping algorithms, we will assume an a priori probability for the occupancy grid algorithm of $P_{F_i}(f) = \frac{1}{|\mathcal{F}|}$. Similarly, we assume that in our Markov random field model $\Psi_i(f) = \frac{1}{|\mathcal{F}|}$.

In our experiments we use the following GPR sensor noise models. In the binary case, we use a model with a false positive rate of 0.05 and a false negative rate of 0.05. Table I shows this model. In the ternary case, we use a model with a false positive rate of 0.025 over nothing, a false positive rate of 0.05 over clutter, and a false negative probability of 0.05 over landmines. Table II shows the model we used. The specific values used in our noise models are meant to reflect the higher rates of false alarms when sensing clutter, and do not correspond to empirically estimated noise model parameters.

TABLE II
NOISE MODEL FOR A GPR SENSOR IN A TERNARY GRID.

$P_{Y F}$	$Y = 0$	$Y = 1$	$Y = 2$
$F = 0$	0.95	0.025	0.025
$F = 1$	0.025	0.95	0.025
$F = 2$	0.05	0.05	0.90

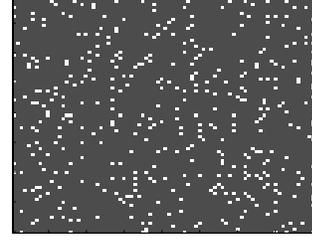


Fig. 3. A randomly generated 80×80 binary grid generated as a pairwise Markov random field. Dark gray corresponds to nothing and white corresponds to landmines.

A. Binary Maps

Before demonstrating simulation results on a model that includes clutter, we explore a simpler binary model that contains either landmines or nothing. A binary map also facilitates direct comparison with the occupancy grid algorithm. We generate the binary random fields for this experiment by sampling from a Markov random field with a distribution specified by the following functions.

$$\Psi_{ij} = \begin{pmatrix} 0.8 & 0.2 \\ 0.9 & 0.1 \end{pmatrix}$$

These values correspond to the likelihood of adjacent grid block values. For example, a landmine directly adjacent to a landmine has a value of 0.1, while a landmine adjacent to nothing has a value of 0.9. Such values model scattered non-contiguous landmines, as shown in Figure 3. We generated a set of 200 80×80 binary grids by randomly sampling the Markov random field distribution. For each grid we randomly deployed a fixed number r of robots and allowed them to collect measurements for 80 time steps. There were 200 such deployments per binary grid. For each deployment we obtained $80 \times r$ measurements. We processed the measurements using both occupancy grids and loopy belief propagation. The quality of the maps was assessed using various error metrics described below, and the average performance (across all deployments and grids) for each error metric was determined. We conducted a set of experiments to determine algorithm performance as we vary the number of robots (and therefore robot samples).

To evaluate the quality of the estimated minefield maps we used two types of error metric. The first is a map score metric of the type introduced by [10], which compares two occupancy grids on a block-by-block basis. Specifically, we use a sum of squares version of the map score metric [12] which is specified as follows,

$$\text{Score} = \sum_i (g_i - f_i)^2$$

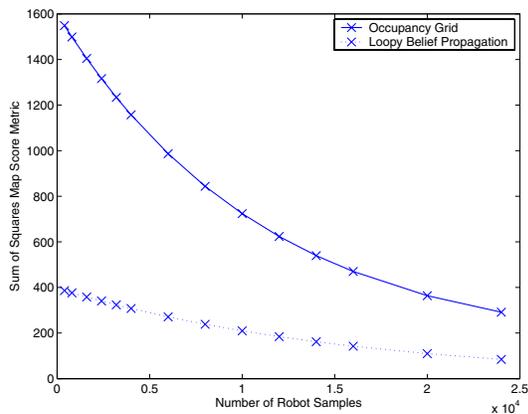


Fig. 4. Sum of squares map score metric.

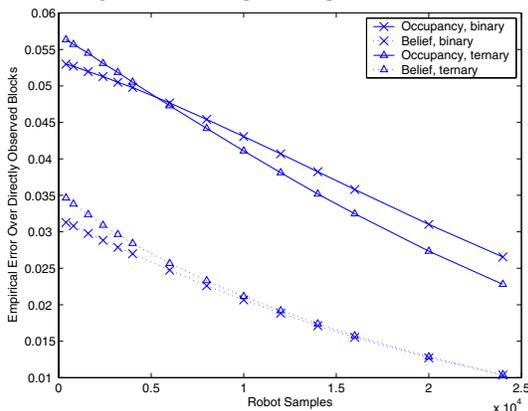


Fig. 5. Empirical error for occupancy grid and belief propagation for binary and ternary grids, excluding areas not observed directly.

where f_i corresponds to the true value of the binary grid, and g_i corresponds to the probability of grid block F_i equalling one. Thus $g_i = P_{F|Y}(1|\mathbf{y})$ for occupancy grids and $g_i = b_i(1)$ for loopy belief propagation. For unobserved blocks F_i , we set the occupancy grid value to its a priori probability $g_i = 0.5$. Loopy belief propagation produces an estimate for all grid blocks even if only a subset of them were directly observed, by taking advantage of spatial dependence to conduct inference. The performance according to this metric is shown in Figure 4.

The second type of error metric reflects the mine detection objective, which is to identify as many landmines as possible with as few false positives as possible. For each grid block that was directly observed, we choose a grid block value that corresponds to the highest probability or belief (i.e., we estimate the value of the directly observed grid block based on the robot samples and model of the environment). We count the number of times that this estimate indicates a landmine and the ideal map does not (i.e., false alarms), and the number of times that the estimate indicates the absence of a landmine while the true map contains a landmine (i.e. false negatives). This sum is normalized by the total number of directly observed grid blocks, and is referred to as the empirical error. The performance of the two algorithms according to this error metric is shown in Figure 5.

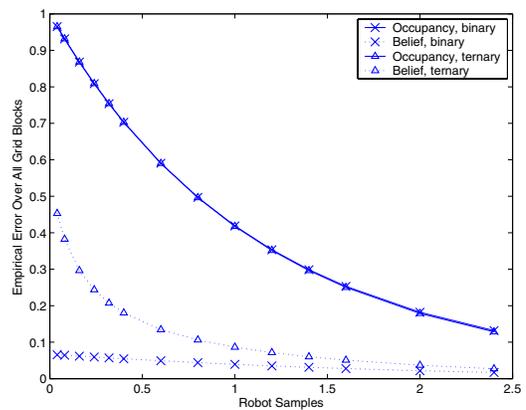


Fig. 6. Empirical error for occupancy grid and belief propagation algorithm for binary and ternary grids, for all grid blocks.

In the context of minefield mapping, a fair error metric should penalize the occupancy grid approach for its inability to conduct inference on areas that were not directly observed. Unmapped areas are unusable because their safety cannot be verified. Therefore in the following graph we modify the empirical error metric to count unobserved areas as an error. The results of our experiments according to this error metric are shown in Figure 6. This result highlights the ability of the belief propagation-based approach to infer parts of the map that were not directly observed by exploiting spatial dependencies. Figure 7 demonstrates this advantage for one particular map and robot deployment. As shown in the figure, the belief propagation algorithm provides reasonable estimates for grid blocks that were not directly observed.

Under the error metrics used above, loopy belief propagation infers maps of superior quality with significantly fewer robot samples. Its ability to do so requires more computation than occupancy grids. We timed the execution of both the occupancy grid and loopy belief propagation algorithms in our previous experiments in order to quantify this difference. Figure 8 summarizes our results.

The results were obtained on a Pentium IV 1.8GHz computer, with a belief propagation convergence criterion of 0.001, and a maximum number of allowable iterations of 80. The bump in the timing curve of the belief propagation algorithm recurred in other experiments and on other computers. Though belief propagation takes more time to run, it is faster on average than a tenth of a second in the experiments described above. When evaluating the relative merits of the two algorithms, one needs to examine the reduction in samples required to achieve the same map accuracy against the additional computational burden. For example, let us consider the number of robot samples required to achieve an empirical error of approximately 0.032 over the observed areas as shown in Figure 5. The occupancy grid-based approach requires 16000 samples, while the belief propagation-based approach achieves the same result in 1600 samples. For 16000 samples, the occupancy grid ran in 0.002 seconds, while for 1600 samples, belief propagation ran in 0.076 seconds. While

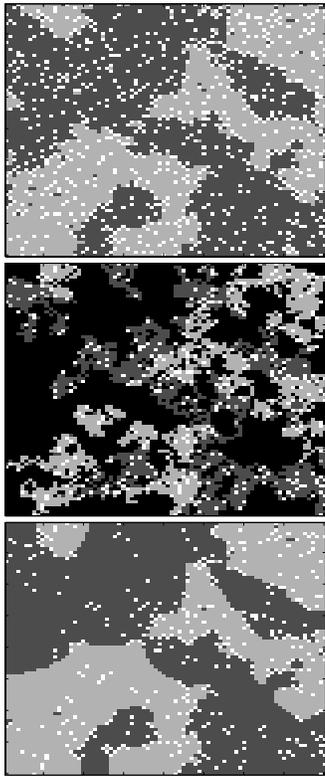


Fig. 7. The top 80×80 grid corresponds to ground truth. The grid immediately below it corresponds to the occupancy grid estimate based on 6000 robot samples (randomly gathered). The bottom grid corresponds to a loopy belief propagation-based estimate using the same 6000 samples. Black corresponds to areas not directly observed in the occupancy grid framework, dark gray corresponds to nothing, light gray corresponds to clutter, and white corresponds to landmines.

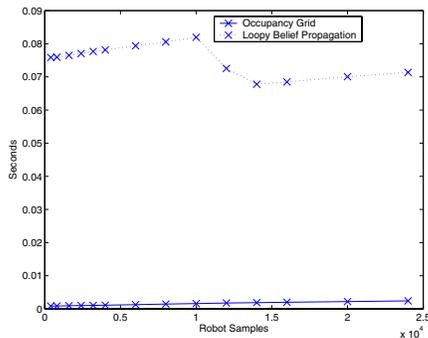


Fig. 8. Algorithm execution time as a function of robot samples.

the decision will vary according to application, we believe that there are a number of applications where 0.074 seconds of computation would be more efficient in terms of energy and time than gathering ten times the number of samples due to the significant energy and time expenditures associated with robot movement and active sensing. Of course we realize that some robots possess limited processing, but the potential sharp reduction in number of samples might still prove the more complex algorithm worthwhile, particularly in light of the increasing availability of fast, low-power processors. Further, such a disparity in the number of samples may suggest that one could build a more energy-efficient robot

by incorporating stronger computational capabilities under the assumption that as a result, the robot will be able to run more complex algorithms and sample the environment fewer times.

B. Ternary Maps

We conducted experiments identical to those conducted in the previous section for a ternary grid that includes clutter as well as landmines and nothing, using Ψ_{ij} as specified in (2). While the score metric is only meaningful for a binary grid, the empirical error metric applies directly. Our results were similar to the previous section. It is important to note that the empirical error does not penalize the algorithm for clutter mistaken for nothing, or nothing mistaken for clutter, since such errors are not relevant to the mapping goal. Figure 5 shows the empirical error over the directly observed areas of the map and Figure 6 shows the empirical error over the entire map. In both cases, the belief propagation algorithm achieves a lower average empirical error with fewer robot samples.

VI. CONCLUSIONS

The results of our experiments with a mapping algorithm that takes advantage of spatial dependencies in the environment indicate that a significant reduction in the number of required samples can be obtained at a modestly increased computational burden. In the future we seek to study the dependence of the belief propagation-based approach on the accuracy of the environmental model, and algorithms for online estimation of Markov random field parameters as a part of robot mapping.

REFERENCES

- [1] A. Elfes, *Occupancy grids: a probabilistic framework for mobile robot perception and navigation*, Ph.D. thesis, Electrical and Computer Eng. Dept., Carnegie Mellon University, 1989.
- [2] S.Z. Li, *Markov Random Field Modeling in Computer Vision*, Springer, 1995.
- [3] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 1988.
- [4] J. MacDonald, J. R. Lockwood, and J. McFee et. al., "Alternatives for landmine detection," Tech. Rep., RAND Corporation, 2003.
- [5] S. Sukkarieh, E.M. Nebot, and H.F. Durrant-Whyte, "A high integrity IMU/GPS navigation loop for autonomous land vehicle applications," *IEEE Trans. on Robotics and Automation*, 1999.
- [6] S. Thrun, S. Thayer, W. Whittaker, C. Baker, W. Burgard, D. Ferguson, D. Hähnel, M. Montemerlo, A. Morris, Z. Omohundro, C. Reverte, and W. Whittaker, "Autonomous exploration and mapping of abandoned mines," *IEEE Robotics and Automation Magazine*, vol. 11, no. 4, 2005.
- [7] S. Thrun, "Learning occupancy grids with forward sensor models," *Autonomous Robots*, vol. 15, pp. 111–127, 2003.
- [8] Yangang Zhang, Mark Schervish, and Howie Choset, "Probabilistic hierarchical spatial model for mine locations and its application in robotic landmine search," in *Proc. 2002 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, October 2002.
- [9] J. Thorpe and R. McEliece, "Data fusion algorithms for collaborative robotic exploration," in *IPN Progress Report 42-149*, J. Yuen, Ed. Jet Propulsion Laboratory, NASA, 2002.
- [10] M. Martin and H. Moravec, "Robot evidence grids," Tech. Rep., Carnegie Mellon University, 1996.
- [11] D. Mackay, *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, 2003.
- [12] J. Collins, M. Eaton, M. Masfield, D. Haskett, and S. O'Sullivan, "Developing a benchmarking framework for map building paradigms," in *Proc. 9th Int. Symp. Art. Life and Robotics*, 2004.