A Bayesian Approach to Occupancy Mapping With Uncertain Inputs

Simon T. O'Callaghan
Australian Centre for Field Robotics
University of Sydney
NSW 2006 Australia
s.ocallaghan@acfr.usyd.edu.au

Fabio T. Ramos
Australian Centre for Field Robotics
University of Sydney
NSW 2006 Australia
f.ramos@acfr.usyd.edu.au

1 Introduction

This work addresses the problem of occupancy mapping with uncertain measurements taken from one or more mobile robots. Appropriate modeling of sensor and localisation uncertainty is critical to obtaining consistent and robust maps which may subsequently be used in planning and motion control.

The proposed approach uses an occupancy mapping method that employs Gaussian processes (GPs) to describe and exploit spatial dependencies inherently introduced by structure in real-world environments. Traditional occupancy grids [1] decomposes the high dimensional mapping problem into many single dimensional, independent, binary classification tasks. This independence assumption ignores the fact that in the real-world, cells of occupancy are not distributed at random over the environment, rather there exists a spatial correlation between cells due to the physical structure of objects and environment.Rather than discretising the world into a grid, GPs describe the properties associated with a map, such as occupancy, in the form of a continuous non-parametric function.

Adopting a Bayesian approach enables accurate maps to be generated from relatively sparse sensor information and allows inference of occupancy state with associated variance in unscanned regions of the environment. The continuous nature of the GP means that the generated maps are not confined to a single scale. Large coarse resolution maps or detailed reconstructions of specific areas of interest can be simply sampled from the underlying model. Uncertainties in the vehicles' measurements and locations are incorporated into the hypothesis of occupancy estimates by defining covariance functions to represent correlations between distributions over the noisy inputs. This allows information from multiple sources with varying degrees of uncertainty to aid in the training of the GP and to be integrated into a common probabilistic representation of the environment.

The principle advantages of our mapping technique include:

1. The ability to generate accurate maps with relatively sparse and uncertain sensor information from multiple sources.
2. Producing an associated variance plot of the occupancy estimates which could be used to highlight unexplored regions and optimise path planning.
3. The incorporation of context in occupancy maps.
4. Multi-resolution maps.

Experimental results, using both synthetic and real outdoor data, have been performed that demonstrate the benefits of exploiting structural dependencies and merging information from various sources into the Bayesian occupancy maps. Quantitative comparisons show the Bayesian approach outperforming traditional occupancy grid on a number of key metrics, particularly when the datasets contain sparse measurements or numerous occluded regions. Furthermore, the related predictive variance maps correctly identify areas of uncertain occupancy hypotheses due to noisy observations or lack of information.

2 Theory

The task of mapping the robot's surroundings is considered as a classification problem. A trained, non-parametric, Bayesian regressor known as the Gaussian process [4] is combined with a probabilistic least squares classifier [3] to represent the environment as a probability distribution and label it into regions of occupancy and unoccupancy.
Bayesian occupancy maps are generated using the following procedure:

1. Initially, uncertainties in the physical system are integrated into the training inputs using an unscented transform to represent the training points as probability distributions governed by the associated noise in sensor range, bearing, vehicle position and orientation.
2. At its core, the Gaussian process is a regression technique. The technique assumes that the training outputs may be noisy however the classical GP does not account for the possibility of uncertain training inputs, **x**. Girard in [2] discusses a method of modifying the conventional GP to redefine the covariance function as a measure of correlation between distributions rather than deterministic point locations. Essentially, the influence of noisy training inputs is dispersed in proportion to their associated variance. This enables information from multiple sources with vary noise levels to be integrated into the probabilistic spatial representation of the environment.
3. The Bayesian approach to mapping is based upon the GP's ability to predict $p(O|\mathbf{x})$, where O is the hypothesis of occupancy and **x** represents a physical location within the map. O_i can be considered as a class, either occupied or unoccupied, referenced by its corresponding location, $\mathbf{x}_i$. The Gaussian process is used to fit a likelihood function to the training data $\{\mathbf{x}_i, y_i\}_{i=1\rightarrow n}$ where n is the number of training points and y_i, the training output or target data, represents occupancy or unoccupancy at a scanned location. The resulting continuous function can then be used to interpolate between data points to predict the occupancy probability in unscanned and occluded regions using the theoretically appealing Bayesian statistical framework.
4. Due to the non-stationary behaviour of typical map datasets (sudden changes from non-occupied to occupied regions), the commonly used squared exponential covariance function with its smoothing properties is not suitable for this application. The neural network covariance function is employed due to its non-stationary characteristics and its ability to model the sharp shifts in the trend of the underlying function, $f(\cdot)$. [5]
5. The Gauss-Hermite quadrature is used to generalise GP inference for uncertain inputs to all categories of covariance functions when a closed form solution does not exist.
6. Training the hyperparameters is performed using densely sampled points from sensor observations in a subsection of the scanned environment.
7. The large nature of the datasets associated with this problem forces local approximations during inference to reduce the required computational time. Thus, a KD-Tree data structure is employed to extract only the observations in each test point's neighbourhood.
8. A probabilistic least squares classifier described in [3] is used to represent the regressor's mean and variance predictions as probabilities of occupancy in space. Using the outputted probabilistic distribution, the environment can be categorised into occupied, unoccupied and unsure regions using user-defined thresholds which can be tuned to match the desired level of greediness.

3 Results

Initial tests were carried out using simulated datasets that provided known ground truths for quantitative analysis. Performing Bayesian inference into unscanned regions yields more complete maps with an associate variance for each occupancy estimate. An early experiment (Fig. 1) with no uncertainty in vehicle position or sensor readings compared the Bayesian approach to the traditional occupancy grid. Of the two mapping techniques, the resulting Bayesian map most closely resembled the ground truth. All of the roads and side-streets can be identified despite the fact that the sensor data did not fully map them. Using local observations to perform inference meant that the shapes of the buildings are also comparable to those of the actual environment. The parked vehicles are also easily identifiable despite the fact that the areas behind them were reasonably occluded from the robot's sensor. ROC tests showed the the false positive rate of the Bayesian occupancy map for a true positive rate of 0.95 was 0.026 compared to 0.1913 for the occupancy grid.

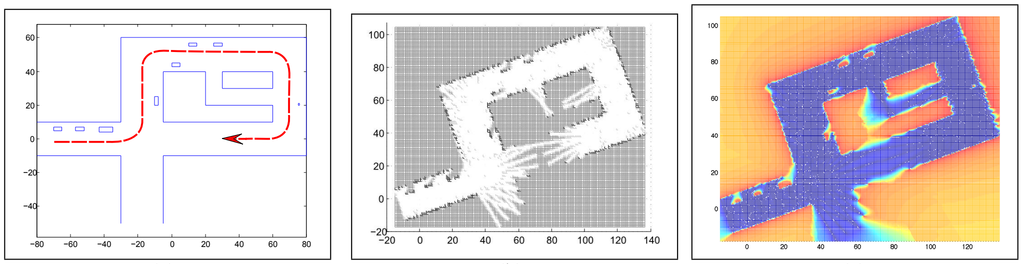
(a) (b) (c)

Figure 1: (a) Simulated street ground truth and vehicle path. (b) Resulting occupancy grid. (c) Bayesian occupancy map. Reddish areas indicate regions with high probability of occupancy while bluish regions suggest the area is most likely unoccupied.

A second study produced a more challenging dataset involving two mobile robotic platforms possessing sensors with varying levels of uncertainty in their range measurements. The first platform is relatively slow moving but possesses a highly accurate range finder. The second platform, in contrast, is fast moving but the sensory information it gathers is extremely noisy. Fig. 2 illustrates the data gathered from both platforms. Information from both sources can be merged using the proposed mapping technique to produce a common probabilistic representation of the occupancy hypothesis across the environment. The outputs of the algorithm are presented in Fig 3. By correctly handling the uncertainty in the dataset, the Bayesian occupancy map, Fig.3(a), reconstructs the general layout of the environment despite the extremely noisy data. The boundaries of objects that were scanned by the accurate sensor are well defined. Although information from the second platform is distorted, it can still be used to make rough estimates of the occupancy hypothesis in regions where Platform I has not scanned such as the far right of the map.

The classified output (Fig. 3(b)) supports this observation with the left section of the map (scanned by both robots) contains significant portions of confident and accurate classification. Conversely, the right half of the map becomes unsure in areas where only a few noisy readings occur (upper right corner) and in regions where neither platform have scanned (lower right corner).

Handling training input uncertainty appropriately also yields benefits in the associated predictive variance output. Fig. 3(c) illustrates how the predictive variance in the proposed method is lowest in areas that have been accurately mapped by Platform I and increases in regions scanned by the less reliable Platform II. As expected, the variance is highest where estimates of occupancy are predicted in the largely unexplored lower right quadrant. Crucially, this output of the algorithm could be combined with a subsequent path planner to optimise the trajectories of the vehicles with the goal of maximising the system's overall understanding of the environment.

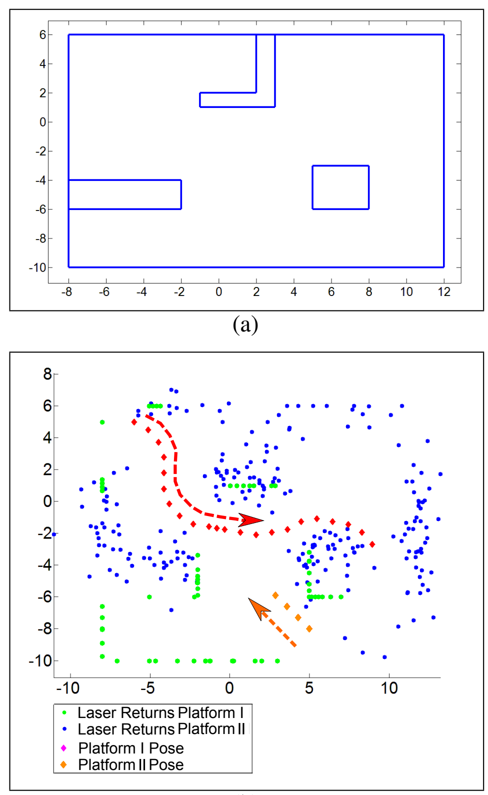

(b)

Figure 2: (a) Simulated dataset ground truth. (b) Platform I poses (orange) and laser returns (green), Platform II poses (red) and laser returns (blue).

The mapping technique has also been tested on several outdoor datasets such as that illustrated in Fig. 4(a). Using only dead reckoning to localise the platform, positional uncertainty continuous to

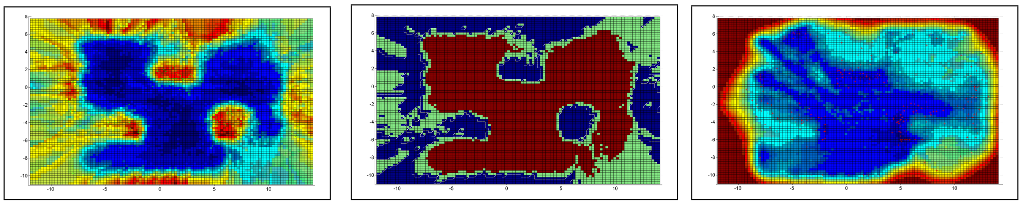

(a) (b) (c)

Figure 3: (a) Continuous probability of occupancy representations. (b) Classified map after applying thresholds. Classification labels: Blue = Occupied; Red = Unoccupied; Green = Unsure. (c) Predictive variance.

rise over time. This is reflected in the global representation of the map by probability of occupancy estimates becoming less and less certain (Fig.4(b)). Akin to the simulated tests, predictions using training inputs with distributions of higher variance results in less well defined boundaries however rough estimates of larger objects such as the road and buildings are still discernable. As a result of this poorly localised data, the majority of occupancy probability predictions within the map range from 0.4 to 0.65. Following loop closure using a sparse extended information filter, the variance in position with respect to a global reference frame shrinks. Consequently, the resulting map becomes sharp and more certain with predictions ranging from almost 0 to 1 (Fig. 4(c)). The roadway, buildings and smaller objects such as parked cars are now much more identifiable.

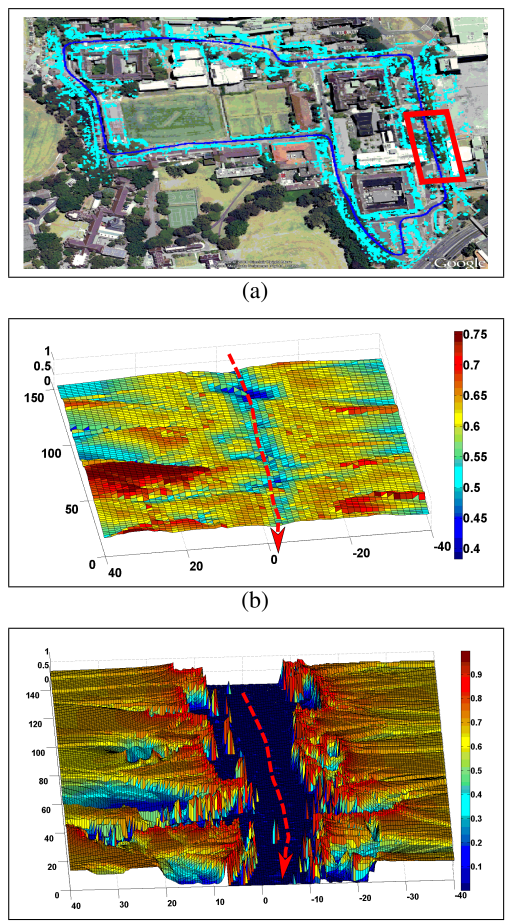

(a) (b) (c)

Figure 4: Results for outdoor dataset. (a) Satellite few of traveled route. (b & c) Probability of occupancy using Bayesian map for area highlighted by the red rectangle in (a) both before and after loop closure, respectively. Red arrow = Path of vehicle.

4 Discussion

Bayesian occupancy maps offer several important benefits when compared to other mapping techniques being employed by the robotics community today. Using a Gaussian processes with non-stationary neural network covariance functions to model occupancy in real-world environments allows Bayesian inferences to be performed to produce continuous probabilistic representations of occupancy estimates with associate variance plots.

The modifications to the GP make the maps robust to the inescapable effects of uncertainty present in measurements and localisation. Sensor readings from multiple sources of differing noise levels can now be naturally integrated into the learning and inference procedures to create an accurate common probabilistic model of the system's surroundings.

Currently, this technique is being revised to handle 3-D datasets. A future extension of this work will focus on including semantic labeling into the occupancy maps.

References

[1] A. Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, 1989.

[2] A. Girard. *Approximate Methods for Propagation of Uncertainty with Gaussian Process*. PhD thesis, 2004.

[3] J. C. Platt. Probabilities for SV machines. In *Advances in Large Margin Classifiers*, pages 61–74. MIT Press, 2000.

[4] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[5] C. K. I. Williams. Neural computation with infinite neural networks. *Neural Computation*, 10:1203–1216.