2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

Contextual Occupancy Maps Incorporating Sensor and Location Uncertainty

Simon T. O'Callaghan, Fabio T. Ramos and Hugh Durrant-Whyte

Abstract— **This paper describes a method of incorporating sensor and localisation uncertainty into contextual occupancy maps to provide for robust mapping. This paper builds on a recently proposed application of the Gaussian process (GP) to occupancy mapping. An extension of GPs is employed which incorporates uncertain inputs into the covariance function. In turn, this allows statistically consistent, multi-resolution maps to be constructed which exploit the spatial inference properties of GPs while correctly accounting for sensor and localisation errors. Experiments are described, with both synthetic and real data, which show the benefits of complete uncertainty modeling and how contextual occupancy maps may be constructed by fusing data from different sensors on different robots in a common probabilistic representation.**

I. INTRODUCTION

This paper 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.

A popular approach to mapping is to model the hypothesis of occupancy across the space. Occupancy grids, where the occupancy hypothesis is computed on a regular spatial grid, is one of the most widely used methods of mapping within the robotics community. Developed by Elfes and Moravec in the late 1980's, [4] and [9], this mapping technique has appeared in various forms employing sensors including sonar, laser and stereo vision [11], [20], [2]. By segmenting the environment into a field of discrete and independent cells, the occupancy grid decomposes the high dimensional mapping problem into many single dimensional binary classification tasks.

This 'independence between cells' assumption made in the occupancy grid approach to mapping results in substantial simplification in implementation and much improved real-time performance. However, the assumption also ignores the fact that in the real-world, cells of occupancy are not distributed independently at random over the environment, rather there exists a spatial correlation between cells due to the physical structure of objects and environment. Use of this assumption frequently results in cells of high uncertainty in regions where spatial context could assist in classifying the state of a cell. This is perhaps most clearly seen in occluded areas or segments between sensor beams where there is absence of sensor information to make a specific classification. Furthermore, the classical approach of representing the environment as a rigid grid leads to discretisation errors and fixes the resolution of the map which ultimately impedes scalability.

Recent work [12] proposes an occupancy mapping method that employs Gaussian processes (GPs) to describe and exploit spatial dependencies inherently introduced by structure in real-world environments. 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. This approach enables accurate maps to be generated from relatively sparse sensor information and allows inference of occupancy state and associated variance in unscanned regions of the environment. The continuous nature GP estimation means that scale now becomes arbitrary. Large coarse resolution maps or detailed reconstructions of specific areas of interest can be simply sampled from the underlying model.

However, the mapping technique described in [12] does not incorporate or account for sensor measurement or localisation uncertainty which may confuse the training process and seriously degrade the quality and consistency of the resulting map. This paper proposes an approach that addresses uncertainty in these *training input locations* by defining GP covariance functions over distributions rather than deterministic locations.

The principle contributions of this paper are:

1) A major extension to the contextual occupancy maps proposed in [12] to incorporate the effects of uncertainty in sensor readings and platform localisation;
2) The derivation of a general method to approximate covariance functions defined over probability distributions.

This paper is organised as follows. Section II summarises relevant papers literature in occupancy mapping and discusses how uncertainty is incorporated into some of the proposed techniques. Section III describes the fundamental principles behind GP-based contextual occupancy maps and proposes an approach to incorporating sensor measurement and localisation uncertainty into this model. Section IV presents experimental results for both simulated and real data sets which demonstrates the efficacy of the proposed approach. Section V discusses the principle conclusions.

S. T. O'Callaghan, F. T. Ramos and H. Durrant-Whyte are with the Australian Centre for Field Robotics, University of Sydney, J04, NSW 2006 Australia. This work is supported in part by the Australian Research Council (ARC) Centre of Excellence program. Correspondence should be addressed to s.ocallaghan@cas.edu.au

II. RELATED WORK

There is a considerable literature on the development of spatial occupancy representations that account for uncertainty in sensor measurements and localisation. The conventional approach with occupancy grids is to transfer this uncertainty into the sensor model where a Bayesian update procedure is then employed to update uncertainty in the map, [8]. While the addition of uncertainty increases the spread of the sensor model and, consequently, the number of grid cells whose hypothesis of occupancy is updated from the prior, it is important to note that this smearing effect does not lead to any notion of spatial correlation. The independence between cells assumption still exists.

Over the years, several papers have proposed alternative methods of fusing uncertain sensor data into a map. Notable is the histogram grid approach [1], and the use of the Dempster-Shafer formulation, a generalisation of Bayes rule, to estimate the evidence supporting hypotheses that a region is occupied or unoccupied [13]. While these approaches improve on traditional occupancy grids in some respects, ultimately, inference in regions beyond the coverage of the sensor model is not possible as there is no notion of spatial correlation.

More recently, a number of authors have taken advantage of the intrinsic structure in an environment to develop mapping techniques that remove this independence between cells assumption. Veeck and Burgard [18] train polylines to form a continuous representation of the environment's boundaries based on discrete range samples. This approach greatly compacts the size of the map using a set of heuristically derived rules to iteratively optimise the polylines. A drawback to this method however is that uncertainty in the sensor measurements is not handled in a probabilistic framework and noisy data that does not conform to the user-defined list of optimisation criteria can lead to maps converging to incorrect representations.

A more robust mapping technique was proposed by Paskin in [14] using polygonal random fields to probabilistically reason about occupancy, rather than the boundaries, of the environment. The maps generated are continuous and allow for inference to be made in unscanned regions. While this approach can not iteratively add new data, the use of a probabilistic model for occupancy enables this approach to produce accurate maps with both laser and sonar datasets.

The work described in [12] extends this idea using GPs with non-stationary neural network covariance functions to model environment occupancy. This enables predictive representations of the environment to be statistically inferred from measured data by exploiting the inherent spatial correlation structure of the GP. This was demonstrated to significantly improve the accuracy and completeness of the occupancy map. This paper further develops the GP model approach. It incorporates the idea of modeling noisy inputs to a GP [6], such as sensor measurements and location uncertainty, through modification of the GP kernel function to naturally fuse information from noisy sources into a common probabilistic model.

III. 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 [16] is combined with a probabilistic least squares classifier [15] to represent the environment as a probability distribution and label it into regions of occupancy and unoccupancy.

Uncertainties in the physical system are integrated into the map initially 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. Modifications to the conventional Gaussian process are made which redefine the covariance function as a measure of correlation between distributions rather than point locations. This enables information from noisy inputs to be integrated into the probabilistic spatial representation of the environment. Finally, the Gauss-Hermite quadrature is employed to generalise these modifications to all categories of covariance functions regardless of whether a closed form solution exists or not.

Section III is divided into four main parts. The theory behind Contextual Occupancy Maps is detailed in [12] however a brief summary of the principal components is presented in Subsection A. Subsection B introduces the key equations by which uncertainty is incorporated into the mapping technique. Subsection C provides an overview of the Gauss-Hermite quadrature and Subsection D discusses the manner in which probability distributions are estimated using the unscented transform.

A. Contextual Occupancy Map

1) Gaussian Process: The proposed contextual occupancy map is based upon the GP's ability to predict $p(O|\mathbf{x})$, where O is the hypothesis of occupancy and $\mathbf{x}$ 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 $\left\{\mathbf{x}_i, y_i\right\}_{i=1\rightarrow n}$ where n is the number of training points and y_i, the training output or target data, represents occupancy $(+1)$ or unoccupancy (-1) at a scanned location. The resulting continuous function can then be used to interpolate between data points and after applying a sigmoid function prediction can be made on the probability of occupancy in unscanned and occluded regions using the well understood Bayesian statistical framework.

The Gaussian process itself can be viewed as a distribution over an infinite number of possible functions and inference takes place directly in the space of functions. By assuming that all occupancy hypotheses, indexed by their corresponding location in the environment, are jointly Gaussian, we obtain

$$f(\mathbf{x}_*) = \mathcal{N}(\mu, \sigma^2), \tag{1}$$

where

$$\mu = k^\top(\mathbf{x}_*, X)^\top \left[K(X,X) + \sigma_n^2 I\right]^{-1}\mathbf{y}, \tag{2}$$

$$\sigma^2 = k(\mathbf{x}_*, \mathbf{x}_*) - k(\mathbf{x}_*, X)\left[K(X,X) + \sigma_n^2 I\right]^{-1} k(X, \mathbf{x}_*). \tag{3}$$

Here, $\mathbf{x}_*$ refers to a query or test location, X the training inputs, σ_n^2 the variance of the global noise and K is the covariance matrix. The elements of the covariance matrix $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ are defined depending on a covariance function k parameterised by hyperparameters θ. In this application, the hyperparameters' optimal values for the datasets are derived by maximising the log marginal likelihood using a simulated annealing algorithm followed by quasi-Newton gradient ascent. An extensive explanation and derivation of the Gaussian process can be found in [16].

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 non-stationary and is capable of modeling the sharp shifts in the trend of the underlying function, $f(\cdot)$.

Williams in [19] following the work of Neal [10] derived the following expression for the covariance function:

$$K(\mathbf{x}, \mathbf{x}_*) = \sigma_f^2 \arcsin\left(\frac{2\tilde{\mathbf{x}}^\top \Sigma \tilde{\mathbf{x}}_*}{\sqrt{\left(1 + 2\tilde{\mathbf{x}}^\top \Sigma \tilde{\mathbf{x}}\right)\left(1 + 2\tilde{\mathbf{x}}_*^\top \Sigma \tilde{\mathbf{x}}_*\right)}}\right). \tag{4}$$

Here $\tilde{\mathbf{x}} = (1, x_1, ..., x_d)^\top$ is an augmented vector and σ_f^2 is a hyperparameter signal variance used to scale the correlation between points. Σ in this case, is a diagonal matrix with a component for each dimension of the training inputs and a bias term. During inference, nearest neighbour approximations are used. Translational symmetry is preserved by centring the test point and local training data at the origin. The hyperparameters of the covariance function are thus $\Theta = [\sigma_f^2, \Sigma_{11}, ..., \Sigma_{(D+1)(D+1)}, \sigma_n^2]$.

A post-processing classifier stage described in [15] is used to represent the regressor's mean and variance predictions as a distribution over probabilities of occupancy in space.

Using the outputted probabilistic model, 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.

A more in-depth description of the contextual occupancy map algorithm including the sensor model employed, local approximations made to handle large datasets and the training method can be found in [12].

B. Propagation of Uncertainty in the Gaussian Process Model

At its core, the Gaussian Process is a regression technique. The technique assumes that the training outputs may be noisy and accounts for this by the inclusion of the hyperparameter σ_n^2 in Equation 2. However, the classical GP does not account for the possibility of uncertain training inputs, i.e. noise in the X domain.

To illustrate the effects that uncertain training inputs have on the Gaussian process, a dataset containing ten measurements, one of which has a degree of uncertainty associated

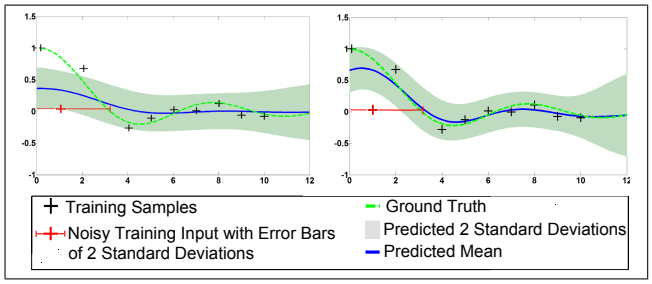

(a) Effects of an uncertain training input on a Gaussian process. Classic GP (left), noisy input GP (right).

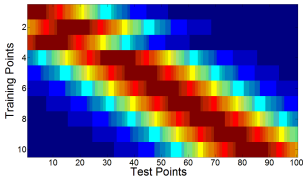

(b) Classic squared exponential covariance matrix. Noisy input's covariance shown on third row

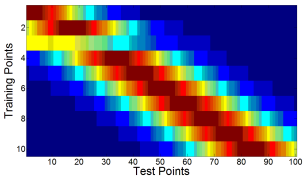

(c) Squared exponential covariance matrix defined over distributions. Noisy input's covariance shown on third row.

Fig. 1. Comparison of standard GP with noisy input GP.

with its location, is used to estimate the underlying ground truth. Fig. 1(a) compares the output of the traditional GP and a noisy-input GP. By assuming that each training input is a deterministic point, the standard method incorrectly deduces very noisy *observations* are the most likely explanation for the dataset. This results in the importance of relatively accurate data points being underestimated. The modified GP, on the other hand, which learns dependencies between distributions rather than single values, closely matches the ground truth.

Examining the associated covariance matrix of the classical GP, Fig. 1(b), reveals that the noisy training point is wrongly influencing test points far from its true position. Ideally, by accounting for this uncertainty in the X domain, the influence of noisier training points should be dispersed in proportion to the magnitude of the associated variance.

Redefining the training inputs as observations corrupted with some Gaussian noise, $\boldsymbol{\epsilon}_{\mathbf{x}_i} = \mathcal{N}(0, \Sigma_{x_i})$,

$$\mathbf{x}_i = \mathbf{u}_i + \boldsymbol{\epsilon}_{\mathbf{x}_i}. \tag{5}$$

The training inputs can now be expressed as $x_i \sim \mathcal{N}(u_i, \Sigma_{x_i})$. Although this addition of uncertainty to X results in the process no longer being Gaussian, Girard details in [6] how it is still possible to calculate the predictive mean and variance of $f(\cdot)$. This approach can be extended to define kernel functions over probability distributions;

$$K_n(\mathbf{u}_i, \mathbf{u}_j) = \iint_{-\infty}^{+\infty} K(\mathbf{x}_i, \mathbf{x}_j) p(\mathbf{x}_i, \mathbf{x}_j), d\mathbf{x}_i d\mathbf{x}_j \tag{6}$$

where $p(\mathbf{x}_i) = \mathcal{N}_{\mathbf{x}_i}(\mathbf{u}_i, \mathbf{\Sigma}_{\mathbf{x}_i})$, $p(\mathbf{x}_j) = \mathcal{N}_{\mathbf{x}_j}(\mathbf{u}_j, \mathbf{\Sigma}_{\mathbf{x}_j})$ and K_n denotes the noisy covariance matrix.

In the case of sensor inaccuracies, it is reasonable to assume that the noise is largely independent and so Equation 6 can be written as

$$K_n(\mathbf{u}_i, \mathbf{u}_j) = \iint_{-\infty}^{+\infty} K(\mathbf{x}_i, \mathbf{x}_j) p(\mathbf{x}_i) p(\mathbf{x}_j), d\mathbf{x}_i d\mathbf{x}_j. \quad (7)$$

It is possible to derive an exact expression for $K_n(\mathbf{u}_i, \mathbf{u}_j)$ when the chosen covariance function is Gaussian. In [6], Girard formulates an exact representation for the commonly used squared exponential covariance function:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f \exp\Big[\frac{1}{2}(\mathbf{x}_i - \mathbf{x}_j)^\top L^{-1}(\mathbf{x}_i - \mathbf{x}_j)\Big] \quad (8)$$

$$k_n(\mathbf{u}_i, \mathbf{u}_j) = \sigma_f \mid I + L^{-1}(\Sigma_{\mathbf{x}_i} + \Sigma_{\mathbf{x}_j}) \mid^{-1/2} \exp\Big[\frac{1}{2}(\mathbf{u}_i - \mathbf{u}_j)^\top (L + \Sigma_{\mathbf{x}_i} + \Sigma_{\mathbf{x}_j})^{-1}(\mathbf{u}_i - \mathbf{u}_j)\Big] \quad (9)$$

Examining Equation (9) shows that the inclusion of uncertainty has the effect of reducing the covariance's magnitude (σ_f) while simultaneously increasing the lengthscale ($\mathbf{L}$) to extend the influence of the training point to additional neighbouring points. This diffusion of the training point's covariance over a larger area enables accurate fusing of the uncertainty's effects into the system. Fig. 1(c) demonstrates the dissolution of the noisy inputs influence across neighbouring test points.

However, in the case where the covariance function is non-Gaussian, e.g. the Neural Network, or dependencies exist between the probability distributions of training points, such as correlation resulting from uncertainties in the vehicle's location, then a closed form solution does not exist and approximations must be performed.

C. Gauss-Hermite Quadrature

The Gauss-Hermite quadrature is a modification of the Gauss quadrature which approximates the integration of a function between limits -1 and $+1$ as a weighted sum of function values at specified points within the domain of integration. It introduces a decaying function, e^{-x}, to extend the limits to $-\infty$ and $+\infty$.

$$\int_{-\infty}^{+\infty} e^{-x} f(x) dx \approx \sum_{m=1}^{n} w_m f(x_m) \quad (10)$$

n is the number of samples depending on the user-defined level of approximation while x_m refers to the roots of the Hermite polynomial, $H_n(x)(i = 1, 2, ..., n)$. The corresponding weights for each sample point are given by:

$$w_m = \frac{2^{n-1} n \sqrt{\pi}}{n^2 [H_{n-1}(x_m)]^2}. \quad (11)$$

Equation 6 can be reformulated to resemble Equation 10 by first expressing $p(\mathbf{x}_i, \mathbf{x}_j)$ as a multivariate Gaussian distribution:

$$p(\mathbf{x}_i, \mathbf{x}_j) \sim \mathcal{N}\Big(\mathbf{x}; \boldsymbol{\mu}, \mathbf{\Sigma}\Big) \quad (12)$$

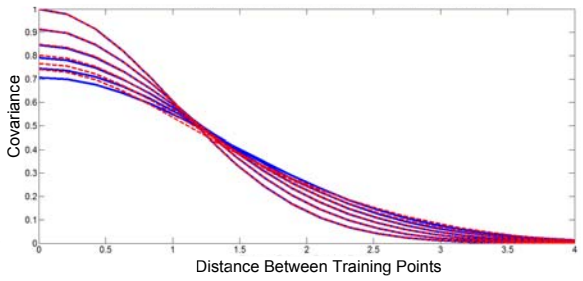

(a) Level 1

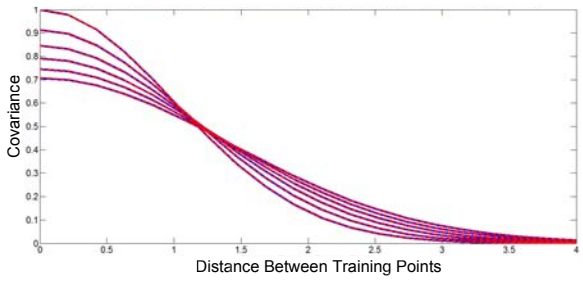

(b) Level 2

Fig. 2. Gauss Hermite quadrature (dashed red line) approximating the modified squared exponential covariance function (Eq. 9) (blue line) evaluated over several distributions of increasing variance. ($\sigma = 0 \rightarrow 1$)

where $\mathbf{x} = [\mathbf{x}_i, \mathbf{x}_j]^\top$, $\boldsymbol{\mu} = [\boldsymbol{\mu}_i, \boldsymbol{\mu}_j]^\top$ and $\mathbf{\Sigma} = \begin{bmatrix} \boldsymbol{\sigma}_i^2 & \boldsymbol{\sigma}_{ij} \\ \boldsymbol{\sigma}_{ji} & \boldsymbol{\sigma}_j^2 \end{bmatrix}$.

Following from this, the Gauss-Hermite quadrature can be used to accurately approximate the effects of noisy training inputs on the performance of the GP by expressing the covariance function as follows:

$$K_n(\mathbf{x}_i, \mathbf{x}_j, \boldsymbol{\mu}_i, \boldsymbol{\mu}_j, \boldsymbol{\sigma}_i, \boldsymbol{\sigma}_j) \approx \frac{1}{2\pi} \sum_{m=1}^{n} w_m K(\mathbf{z}_{i_m}, \mathbf{z}_{j_m}) \quad (13)$$

through direct substitution by letting $\mathbf{z}_i = \sqrt{2\mathbf{\Sigma}_{ij}}\mathbf{x}_i + \boldsymbol{\mu}_i$.

Figure 2 compares a Gauss-Hermite quadrature implementation of the squared exponential covariance function with the performance of the closed form solution derived by Girard (Equation 9). A level 1 approximation (11 sample points) accurately matches the closed form case for low degrees of uncertainty (Fig. 2(a)). However for significantly large variances, the quadrature undercompensates for their influence. Increasing the level of the quadrature considerably improves its ability to match the output of the exact solution. Fig. 2(b) illustrates the output of a level 2 approximation consisting of 49 samples from the kernel function.

D. Unscented Transform

Equation 12 assumes the training inputs to have Gaussian distributions. The variance associated with vehicle pose (location and orientation) and sensor readings (range and bearing) manifests itself as a non-trivial pdf when represented in a the global coordinate frame. An unscented transform [7] is

used to estimate the first two moments of these distributions and represent them as Gaussians. A preferred alternative to employing Jacobians when the transition model is highly-non linear, the unscented transform uses a deterministic sampling technique to pick a minimal set of sample points ($2n+1$ for an n-dimensional space) around the mean.

These sample or 'sigma' points, $\boldsymbol{\chi}$, and their associated weights, w, can be obtained using the following equations:

$$\chi^{[0]} = \mu, \tag{14}$$

$$\chi^{[i]} = \mu + \left(\sqrt{(n+\lambda)\Sigma}\right)_i \quad \text{for } i = 1, ..., n, \tag{15}$$

$$\chi^{[i]} = \mu - \left(\sqrt{(n+\lambda)\Sigma}\right)_{i-n} \text{ for } i = n+1, ..., 2n. \tag{16}$$

Here $\lambda = \alpha^2(n+\kappa) - n$ where α and κ are scaling parameters that determine how far the sigma points are spread from the mean.

These test points are propagated through the true nonlinear system, $g(x)$; {range,bearing,location,orientation}$\Rightarrow$ global position, allow estimation of the posterior mean and covariance (μ' and Σ', respectively) accurate to the third order for any nonlinearity.

$$\gamma^{[i]} = g(\chi^{[i]}), \tag{17}$$

$$\mu' = \frac{1}{2n}\sum_{i=0}^{2n}\gamma^{[i]}, \tag{18}$$

$$\Sigma' = \frac{1}{2n}\sum_{i=0}^{2n}(\gamma^{[i]} - \mu')(\gamma^{[i]} - \mu')^T. \tag{19}$$

The weights and pattern scale of the sigma points may also be adjusted however these variations are omitted here due to space restrictions.

The resulting Gaussian functions approximate the training points' true probability distributions and can now be used to form the inputs to the GP via Equation 13.

IV. RESULTS

A. Simulated Dataset

Initial tests were carried out using simulated datasets. This approach enabled the levels of sensory and localisation uncertainty to be easily manipulated while also providing a known ground truth for quantitative comparisons. Fig. 3 is an example of one of these simulated environments.

In this experiment, two robotic platforms are employed to map the environment. 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. Table I compares the key characteristics of both platforms. Figure 3 illustrates the sensor data obtained during the simulation.

Inputting this dataset into the traditional occupancy grid, the previous version of the contextual occupancy mapping technique [12] and the method proposed in this paper provides an interesting comparison that highlights a number of the key benefits of the new approach.

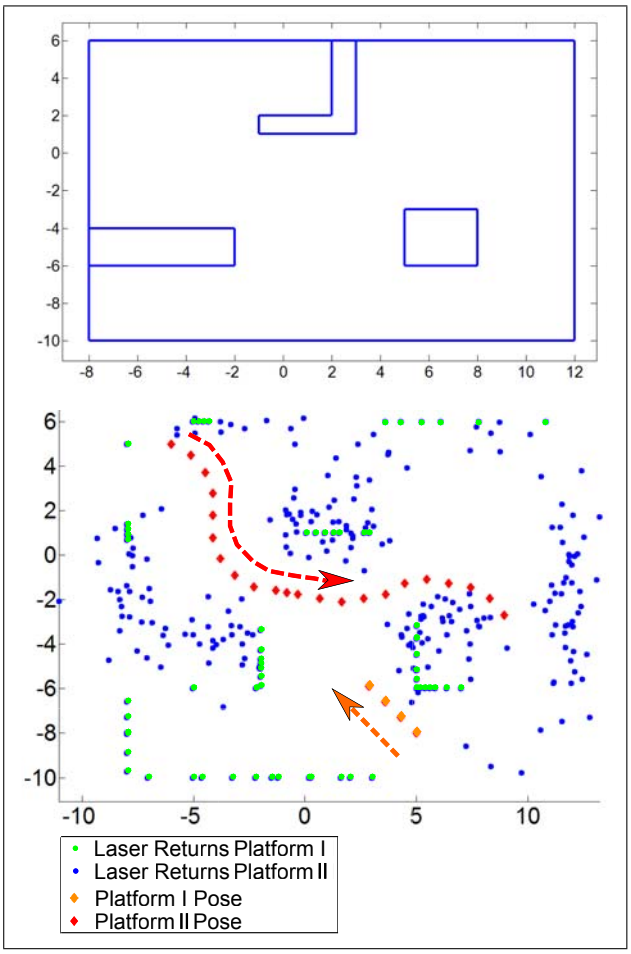

Fig. 3. Simulated ground truth (top). Platform poses and laser returns from experiment (bottom). Platform I poses (orange) and laser returns (green), Platform II poses (red) and laser returns (blue)

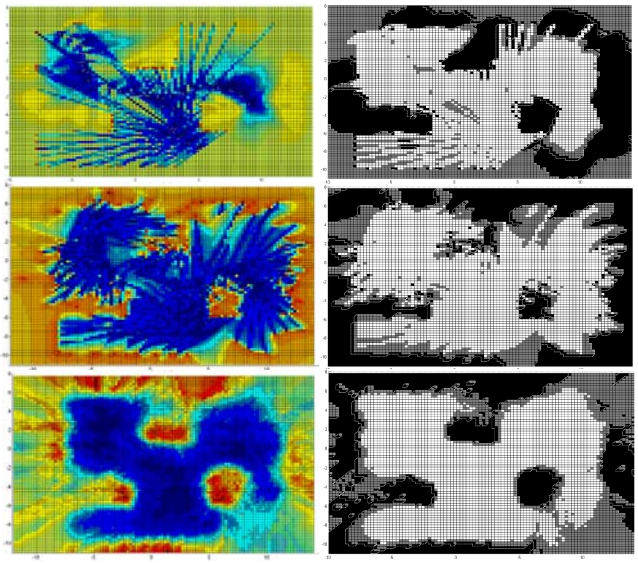

Fig. 4. Sequence of images comparing the performance of the traditional occupancy grid (first row), the GP mapping technique in [12] (second row) and the new proposed method which incorporates uncertainty in the observations (third row). Probability of occupancy versus location prior to thresholding are shown in the left column. Reddish areas indicate regions with high probability of occupancy while bluish regions suggest the area is most likely unoccupied. The right column illustrates their corresponding classified maps after applying thresholds. Classification labels: Black = Occupied; White = Unoccupied; Grey = Unsure.

TABLE I

COMPARISON OF SIMULATED ROBOTIC PLATFORMS

	Platform I	Platform II
Speed (m/s)	0.5	4
Firing Rate (scans/s)	0.5	4
Sensor Range (m)	20	8
Range Variance (m^2)	Negligible	0.8
Bearing Variance ($degrees^2$)	Negligible	0.5
Localisation Uncertainty (m^2)	0	0

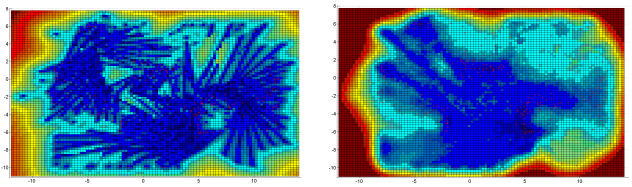

(a) (b)

Fig. 5. Predictive variance of previous, (a), and proposed, (b), approaches.

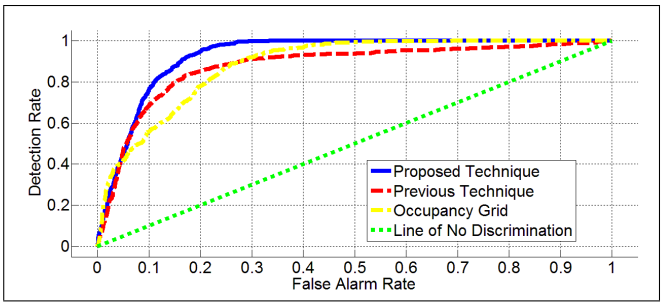

Fig. 6. ROC Curve comparison between proposed and previous mapping techniques.

TABLE II

QUANTITATIVE COMPARISON OF EXPERIMENTAL RESULTS

	Area under the curve	FP detection rate for TP detection rate of 0.97
Proposed Method	0.9303	0.2258
Previous Method	0.8715	0.7957
Occupancy Grid	0.8824	0.3405
No Discrimination	0.5	0.95

The independence-between-cells assumption made by the occupancy grid leads to large portions of the map in the first image of Fig.4 remaining relatively unaltered from the prior hypothesis of occupancy despite significant contextual information available to perform reasonably confident inference in those regions.

Alternatively, using contextual occupancy maps but ignoring the variance associated with the location of the training inputs, as in [12], has led to the learning process converging to incorrect hyperparameters. The outputted occupancy maps (second row of Fig. 4) are influenced heavily by the noisy data resulting in fragmented walls and objects as well as several phantom obstacles. Conflicting hypotheses of occupancy in a number of regions between the noisy and accurate sensors remain unresolved due to both sources being treated as equally relevant. Due to these conflicts, the classified map highlights a large number of areas where the label is unsure despite having been extensively scanned by the more reliable sensor.

In contrast, the proposed new approach generates maps (third row of Fig.4) that bear a much stronger resemblance to the ground truth. The boundaries of objects that were scanned by the accurate sensor are well defined. Additionally, it can be seen that occupancy estimates in regions scanned only by the noisier sensor, such as in the far right of the map, are correctly less certain.

The classified output (bottom right of Fig. 4) 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 too 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. 5(a) shows how the previous approach does not discern between varying levels of sensor noise and the resulting plot poorly represents the degree of confidence one could have in the associated predicted occupancy hypotheses for each region. Alternatively, Fig. 5(b) 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 to maximise the system's overall understanding of the environment.

A quantitative comparison between the outputs of both approaches and the traditional occupancy grid was carried out using a receiver operating curver (ROC), Fig. 6. By monitoring the true positive rate (TP) as the frequency of false positives (FP) is increased the benefits of the proposed approach become apparent. Ignoring the effects of noise in the training inputs (red curve) has led to several misclassified regions which can be catastrophic when considered in a navigation context and explains the relatively slow increase in detection rate. Table II compares key features of each curve.

B. Outdoor Dataset

To test whether the desirable characteristics of the new approach could be replicated in a real dataset, a SICK laser rangefinder was mounted onboard a vehicle which traveled across approximately 1.2 kilometres of the University Of Sydney's campus. The vehicle itself used a NovAtel IMU-HG1700 to perform dead-reckoning. Fig. 7 illustrates the path of the vehicle and laser returns while Fig. 8 shows the covariance for a small subset of the dataset's poses both prior to and following loop closure using a sparse extended information filter.

Fig. 7. Satellite view of scanned region with laser returns indicated

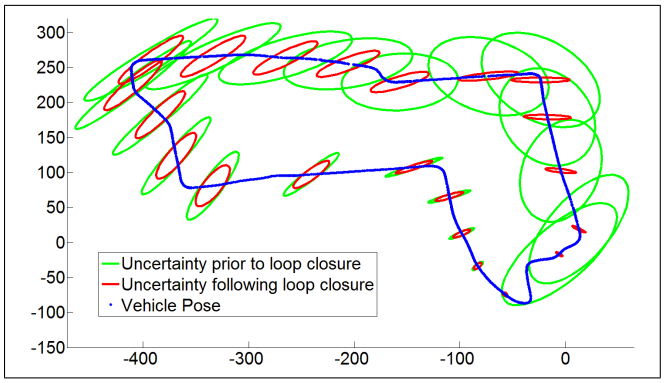

Fig. 8. Vehicle poses and covariance ellipses before and after loop closure

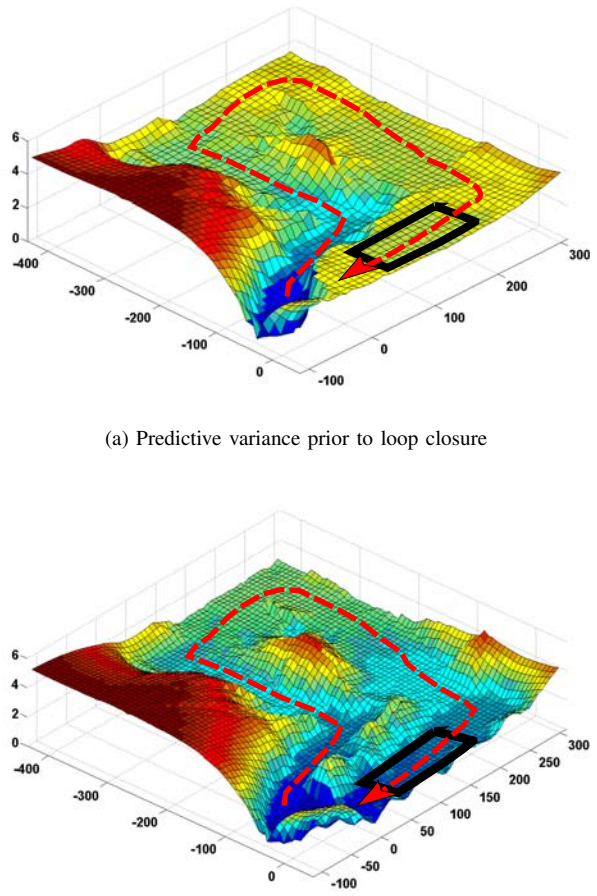

(a) Predictive variance prior to loop closure

(b) Predictive variance following loop closure

Fig. 9. Comparison of coarse predictive variance maps before (a) and after (b) loop closure using the real outdoor dataset. Arrows illustrate path of vehicle.

Initially, coarse maps depicting the outputted variance in occupancy estimates over the entire region were generated using the proposed approach. As was suggested by the simulated experiments, uncertainties associated with the location of the training data now influence the confidence in the predictions made by the mapper. Fig. 9(a) correctly shows a growth in variance as errors in the vehicle's dead-reckoning increases. Similarly, following loop closure detection and subsequent shrinking of the covariance ellipses, Fig. 9(b) indicates the improved belief in the accuracy of the predictions based on a higher degree of confidence in the location of the training inputs. In both maps, the variance is highest in regions where no measurements were taken.

Estimates made of the hypothesis of occupancy within a region should reflect the ambiguity of its training data. Fig. 10(a) and 10(b) plot the probability of occupancy focusing on an area highlighted by the black rectangle in 9(a) and 9(b), respectively. 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.

Once loop closure detection shrinks the covariances associated with vehicle pose, the unscented transform in turn reduces the uncertainty in the training inputs. The resulting map becomes sharp and more certain with predictions ranging from almost 0 to 1. The roadway and buildings are easily identifiable. Similar to the results in [12], Bayesian inference is carried out between training points and in occluded regions. As a result, a number of cars and side streets can also be identified. It is important to note also that the probability of occupancy correctly returns towards 0.5 for estimates made far from any scanned areas.

The proposed changes currently increase the computational cost of the procedure by an order of magnitude. However, there are several trade-off choices in the approximation methods such as the number of neighbouring points selected during inference or the level of the Gauss-Hermite quadrature that can be tuned by the user to suit the intended application. Additionally, performing inference with GPs lends itself to distributed computing and refining the algorithm to work online efficiently will be the focus of future work.

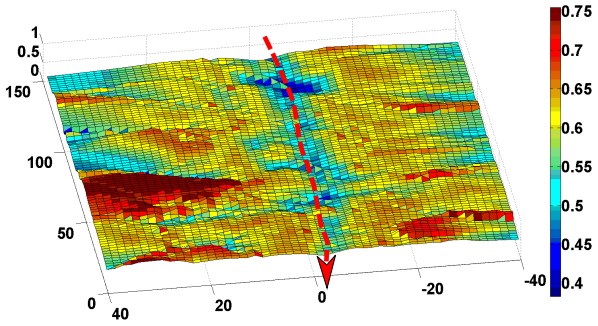

(a) Occupancy map of area highlighted by black rectangle in Fig.7 and Fig.9(a)

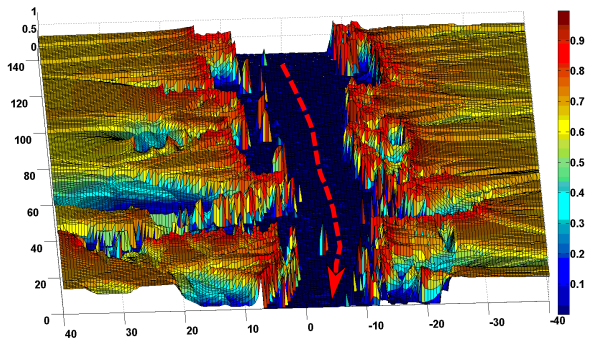

(b) Occupancy map of area highlighted by black rectangle in Fig.7 and Fig.9(b)

Fig. 10. Comparison of contextual occupancy maps generated for a street segment before (a) and after (b) loop closure using the real outdoor dataset.Arrows illustrate path of vehicle. As localisation uncertainty decreases, contextual occupancy maps referenced to the global coordinate frame become sharper and more certain about the hypotheses of occupancy within its regions.

V. CONCLUSIONS

Contextual 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.

This paper proposes several crucial extensions to the existing technique that make it 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.

Similarly, the outputted variance no longer depends entirely on the learnt hyperparameters and the sparsity of the data as is the case with the classical GP but rather now accounts for the ambiguity in the training inputs. The resulting maps provide predictions of occupancy with corresponding variance that can be used to optimise path planning algorithms to maximise the robot's understanding of its environment.

In recent years, Gaussian processes have been applied to a wide range of problems in the robotics community. The approach proposed in this paper to incorporating the uncertainty in training inputs into any generic GP kernel function is not limited exclusively to occupancy maps. Applications such as terrain modeling [17], Wi-Fi SLAM [5] and reinforcement learning [3] may also benefit from directly modeling the effects of noisy measurements in the system.

ACKNOWLEDGMENTS

This work has been supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government. Finally, thanks to Sun Zuolei for generating the outdoor dataset used in this paper.

REFERENCES

[1] J. Borenstein and Y. Koren. Real-time obstacle avoidance for fast mobile robots in cluttered environments. *IEEE Transactions on Systems, Man, and Cybernetics*, 19:1179–1187, 1990.

[2] C. Braillon, C. Pradalier, K. Usher, J. Crowley, and C. Laugier. Occupancy grids from stereo and optical flow data. In *Proc. of the Int. Symp. on Experimental Robotics*, Rio de Janeiro (BR) France.

[3] M. P. Deisenroth, C. E. Rasmussen, and J. Peters. Gaussian process dynamic programming. *Neurocomput.*, 72(7-9):1508–1524, 2009.

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

[5] B. Ferris, D. Fox, and N. D. Lawrence. Wifi-slam using gaussian process latent variable models. pages 2480–2485, 2007.

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

[7] S. J. Julier and J. K. Uhlmann. A new extension of the kalman filter to nonlinear systems. pages 182–193, 1997.

[8] M. C. Martin and H. Moravec. Robot evidence grids. Technical Report CMU-RI-TR-96-06, Robotics Institute, Pittsburgh, PA, March 1996.

[9] H. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, 9(2):61–74, 1988.

[10] R. M. Neal. *Bayesian Learning for Neural Networks*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1996.

[11] S. Noykov and C. Roumenin. Occupancy grids building by sonar and mobile robot. *Robot. Auton. Syst.*, 55(2):162–175, 2007.

[12] S. O'Callaghan, F. Ramos, and H. Durrant-Whyte. Contextual Occupancy Maps using Gaussian Processes. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1054–1060, 2009.

[13] D. Pagac, E. M. Nebot, and H. Durrant-Whyte. An evidential approach to probabilistic map-building. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 1, pages 745–750, 1996.

[14] M. Paskin and S. Thrun. Robotic mapping with polygonal random fields. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 450–458, 2005.

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

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

[17] S. Vasudevan, F. Ramos, E. Nettleton, and H. Durrant-Whyte. Gaussian Process Modeling of Large Scale Terrain. *Journal of Field Robotics*, 26(10), 2009.

[18] M. Veeck and W. Burgard. Learning polyline maps from range scan data acquired with mobile robots. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2004.

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

[20] M. Yguel, O. Aycard, and C. Laugier. Efficient GPU-based Construction of Occupancy Grids Using several Laser Range-finders. In *Proc. of the IEEE-RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Beijing (CN) China, 2006.