

Learning Spatial Correlations for Bayesian Fusion in Pipe Thickness Mapping

Teresa Vidal-Calleja, Daobilige Su, Freek De Bruijn and Jaime Valls Miro

Abstract—Pipe thickness maps are used to assess the condition in pipelines. Thickness maps are a 2.5D representation similar to elevation maps in robotics. Probabilistic frameworks, however, have barely been used in this context. This paper presents a general approach for generating probabilistic maps from heterogeneous sensor data. The key idea is to learn the spatial correlation of a sensor through Gaussian Process models and use it as priors for Bayesian fusion. This approach is applied to the novel application of pipe thickness mapping. Data from a 3D laser scanner on the outer surface of the pipe and thickness measurements from a contact ultrasonic sensor are fused into a single thickness map with associated uncertainty. Moreover, a dedicated algorithm to model the ultrasonic sensor using kernel density estimation is also proposed. The overall approach is evaluated using the full 3D profile (outer and inner surfaces) of the pipe section as ground truth.

I. INTRODUCTION

Deterioration of pipelines used in water, gas or oil transportation systems can lead to leaks and breakages that can result in significant damages to the infrastructure and adjacent properties, expensive repair costs and cause major inconvenience to the public. Detailed assessment of this critical infrastructure to monitor for corrosion and general pipeline wall thickness reduction is therefore a top priority for utilities.

Pipe thickness mapping can be considered as a 2.5D mapping problem similar to terrain modelling. In robotics, elevation maps are a common approach to map the surface of the terrain. In these maps each cell in a Cartesian grid is assigned with a height estimate. This is a compact representation useful for navigation tasks that is amenable to computationally fast implementations in terms of storage, access and evaluation. This map representation is equivalent to the thickness maps produced by the condition assessment industry in pipelines. Thickness maps in pipes are also represented as a Cartesian grid, where each cell is assigned the thickness estimate. These thickness maps are used to decide whether a pipe section needs to be replaced due to its poor condition (thinning of the wall thickness, or pitting due to corrosion) or not. An example of a thickness map associated to a pipe 3D profile is presented in Figure 1.

Probabilistic terrain mapping is a well-studied problem in robotics, however in pipeline condition assessment, the use of probabilistic tools for thickness mapping has rarely been tackled. The main contribution of this paper is a novel probabilistic framework for fusing multiple information from




Fig. 1. Example of a thickness map associated to a pipe 3D surface profile.

heterogeneous sensors, which can be applied in both, robotic mapping and thickness mapping for condition of pipelines.

In the context of the approach proposed in this paper, data fusion is performed through Bayesian fusion. The key idea is to produce a prior map by using a Gaussian Process model that learns the spatial correlation of the profile, further new sensor measurements (sparse or dense) with associated uncertainty can be directly fused to produce the updated map via a maximum *a posteriori* estimation. The aim is to build a map that is amenable to sensor fusion, at the map level, of any number and variety of sensors from which sensor models can be estimated. The purpose behind correlating spatially the initial set of measurements, is that a single observation at a given position will impact the surrounding neighbours. The main advantages of this approach are the generation of more complete, accurate and robust estimates, along with a spatial measure of the uncertainty in the estimate.

We applied this general fusion/mapping framework to a pipe thickness mapping application. For this novel application, we propose a dedicated algorithm to assign confidence values to ultrasonic thickness measurements. Kernel density estimation is used for this purpose. We also implemented an algorithm to produce thickness measurements from external 3D surfaces. Because of the high resolution of the latter measurements, we chose them as prior, where the spatial correlation is learned. The ultrasound measurements are then the ones used to update the priors via Bayesian fusion. Experimental results are validated using the thickness extracted from the external and internal 3D surfaces as ground-truth.

The remainder of the paper is as follow. The literature review in terrain modelling is discussed in Section II. Section III presents the overview of the proposed approach. Section IV gives the background on Gaussian Processes and describes how they are used in the context of this paper. Section V details the fusion approach. The implementation to produced probabilistic thickness maps fusing 3D laser profile information and ultrasonic sensors is described in Section VI. Finally, Section VII presents the experimental results and Section VIII the conclusions.

II. LITERATURE REVIEW

Several approaches in the literature to estimate the elevation of the terrain have been proposed. They typically fall within one of two categories in terms of their treatment of uncertainty. The first approach neglects explicit account of uncertainty by averaging over the multiple measurements that fall within a given cell, and/or discarding outliers that fit poorly the other data associated with the corresponding terrain. The second approach typically constitutes a Bayesian formulation that constructs an expression for the probability of a 3D surface, conditional on the collection of range measurements [1]. This approach, while most appealing because of its mathematical rigour, suffers a few drawbacks for our application. Its formulation requires an a priori parametric model (considering given shapes of the surface) to which the measurements are compared. Pipe surfaces and therefore thickness maps are quite irregular in the presence of corrosion pits (see Figure 1), making difficult to assume a particular shape (parametric model) for both outer and inner surfaces.

Other related approaches, in this respect, include [2], which uses the concept of a “scatter matrix” to represent the local geometric uncertainty in a grid, the approach in [3] to accommodate sensor resolution dependency on range, and in [4] overhanging objects are tackled but fusion is still done by the sum of weighted variances of the uncorrelated data. Our method is related to the latter in the sense that we also use Bayesian fusion, however, we consider the prior map as spatially correlated, therefore a single measurement of a cell makes an impact on its neighbours. The spatial correlation in our case is learned through Gaussian Process models.

Gaussian processes (GPs) [5] are powerful non-parametric Bayesian learning techniques that can fit a spatial correlation model to the data. As stated in [6], GPs incorporate and handle uncertainty in a statistically sound manner and represent structure appropriately. In an estimation sense, GPs provide the best linear unbiased estimate [7] based on the underlying stochastic model of the spatial correlation between the data points. GPs thus handle both uncertainty and incompleteness effectively. In terrain modelling GPs have been used before to model and use the spatial correlation of the given data to estimate the elevation values for other unknown points of interest [8], [9] and more recently [10]. A recent approach for occupancy mapping with GPs has also been presented in [11].

Sensor data fusion using Gaussian processes has also received some attention, examples include [12] and [13]. The former includes, as part of the input vector, one of the values to be fused, therefore the fusion is intrinsically done. The latter work attempts to generalise arbitrary functions on GP priors through linear transformations. It introduces heteroscedasticity (random variables with non-constant variance) and how information from different sources can be fused. The work in [14] integrates heterogeneous information within a classifier based on Gaussian process for a protein fold recognition application. Each feature representation is represented by a separate GP. The fusion uses the fact that individual feature representations are considered independent

and hence a composite covariance function would be defined in terms of a linear sum of Gaussian process priors. In a similar manner [6] fuses multiple, heterogeneous sources of correlated information. The latter work exploits the use of dependent Gaussian process to perform the fusion, learning a different independent GP model for each sensor data and then correlate them. Our approach, in the other hand, only learns spatial correlation for a single sensor data with a GP model, and all the other sources can be fused without the need of learning another GP model. This allows us to handle the fusion even with very sparse data.

III. APPROACH OVERVIEW

In this section we will give an overview of the framework to fuse multiple sensor data into a probabilistic map.

Let us consider the Bayesian fusion problem for mapping, where prior estimates are updated with new measurements from other sensor sources. The general methodology is as follow:

- Prior map generation,
- sensor modelling, and
- Bayesian fusion.

In order to generate prior maps, we proposed to learn and infer the spatial correlation of the measurement through GPs. In practice the sensor with the higher coverage will be the best candidate to be employed for building the prior map. Let $\Psi_1 = \{(\mathbf{x}_0, y_0), (\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ denote a data set of n referenced sensor readings, where $\mathbf{x}_i \in X$ is the position from which the sensor reading $y_i \in Y$ was taken. The data set Ψ_1 is drawn from a noisy process,

$$y_i = f(\mathbf{x}_i) + \varepsilon = \xi_i + \varepsilon, \quad (1)$$

where ε is independent identically distributed Gaussian with variance σ_n^2 and $\xi_i = f(\mathbf{x}_i) \in \xi$. The density $p(\xi|X)$ must then be estimated from Ψ_1 . Through GPs we can also obtain the predictive distribution $p(\xi|X^*, \Psi)$ for arbitrary points X^* .

Let us rename $\xi = [\xi, \xi^*]^\top$, then $\xi \sim \mathcal{GP}(\mu^-, P^-)$ is the jointly Gaussian distributed map, where μ^- and P^- are the posterior mean and covariance. Given this joint distribution, it is straightforward to use it as the prior $p(\xi|X)$ for Bayesian fusion. Because the correlations in position are already considered, a single measurement can update the full map. This Gaussian Process prior give us some more advantages such as the associated uncertainty and the ability to increase or decrease map resolution (inferring more or less points) as required by the new sensor data.

With this model as a prior, a new data set of sensor measurements $\Psi_2 = \{(\mathbf{x}_0, z_0), (\mathbf{x}_1, z_1), \dots, (\mathbf{x}_m, z_m)\}$ with m sensor readings \mathbf{z} taken from the $\mathbf{x}_i \in X^\circ$ positions, can be fused (given a sensor model and the location) using a Bayesian framework. Note that $X^\circ \subset X$ is considered as known.

A model is required for the new sensor data (measurements with associated uncertainty). Sensor modelling is highly dependent on the type of sensor and the application. In general for our approach a Gaussian sensor is considered and modelled as $p(\mathbf{z}|\xi, X) = \mathcal{N}(\mu, R)$, where R is the uncorrelated covariance function of the new measurements.

Note this approach can handle intrinsically sensor models with variable noise.

The maximum *a posteriori* estimator will fuse the prior map $p(\xi|X^*)$ with the new measurements $p(\mathbf{z}|\xi, X)$ to get the posterior $p(\xi|\mathbf{z}, X) \propto p(\mathbf{z}|\xi, X) \times p(\xi|X)^1$.

The output of the fusion is a probabilistic map $\{\mu^+, P^+\}$ with an updated mean and covariance function correlated based on location. This fusion framework allows to integrate multiple data sets from same or heterogeneous sensors. The advantage of our framework, over other general approaches such as [6], is that it allows sparse or dense measurement to be incorporated without the need of learning prior maps for each sensor data and then fuse them together.

IV. PRIOR MAPS

Gaussian Process models are used in this work to model the spatial correlation of the data to be described as a probabilistic map.

A. Gaussian Process

Gaussian Processes are a non-parametric approach in the sense that they do not explicitly specify a functional model between input and output [5]. GPs can be thought of as a Gaussian prior over the function space mapping inputs and outputs and are characterised by a mean function $\mu = E[\xi]$ and the covariance function $E[(\xi - \mu)(\xi - \mu)^\top]$ as $\xi \sim \mathcal{GP}(\mu, P)$. Although the functions are infinitely dimensional, GPs are used to infer, or predict, function values at a finite set of prediction points X^* from the observed data Ψ . The basic GP regression equations for a vector as given in [15] are,

$$\mu = K(X^*, X)[K(X, X) + \sigma_n^2 I]^{-1} Y \quad (2)$$

$$P = K(X^*, X^*) - K(X^*, X)[K(X, X) + \sigma_n^2 I]^{-1} \times K(X, X^*)^\top, \quad (3)$$

where I is the identity matrix, μ the posterior function mean and P is the posterior covariance². The matrix $K(X, X)$ denotes the joint prior distribution covariance of the function at inputs X . This covariance matrix has elements,

$$K(\mathbf{x}_i, \mathbf{x}_j) = P(\xi_i, \xi_j). \quad (4)$$

The matrix $K(X^*, X)$, obtained from the kernel K , is the covariance between the function at the prediction points X^* and the training inputs X . Many kernels have been designed to capture various properties of the modelled phenomenon including smoothness, dimensionality and stationarity. The following section presents the chosen kernel for the problem at hand.

B. Learning and inferring Spatial Correlations

As stated in [16], the problem of learning a prior “map” is one of regression: estimating a function mapping from inputs to output. In the context of the present problem, the inputs are locations and the outputs are sensor measurements. To apply a GP framework to this regression problem, one must first

¹In this approach it is assumed the location from where all the measurements were taken as known.

²A prior mean can also be considered in (2) as shown in [15].

select the kernel K . This specifies the kinds of functions that are expected, before any data have been seen. Technically, the kernel places a prior likelihood on all possible functions. Next, the parameters of the kernel are learned from the training data. Finally, the combination of the training data and the kernel induces not only the most likely state, but also a full posterior probability distribution over the whole map.

From the commonly used kernels we have chosen the Matérn class of functions,

$$K_{\text{Matern}}(X, X^*) = \sigma_f^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}r}{l} \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}r}{l} \right), \quad (5)$$

where r is the Euclidean distance between inputs X and X^* while the trainable hyperparameters l and σ_f represent the lengthscale and amplitude respectively. K_ν is a modified Bessel function and ν is a shape parameter that regulates the smoothness of the interpolation. We note that a number of other kernels, including non-stationary [9], could be adapted and substitute for (5).

All hyperparameters are learned simultaneously from the data Ψ_1 maximising the log marginal likelihood using a gradient decent as shown in [5]. The learning process implicitly trades off the various parameters such as sensor noise and input’s ability to generalise spatially. Note that there are relatively few free parameters that need to be tuned by hand: only the choice of the covariance function and the priors over σ_f and l . A disadvantage of this process is that it is relatively computationally intensive, since each step of gradient descent requires the inversion of the $n \times n$ matrix K_{Matern} . While exact methods were sufficient for the data sets addressed in this paper, Rasmussen and Williams describe a number of approximate methods [5].

Once the hyperparameters have been learned, producing the prior map involves deciding the spatial resolution X^* of the map and inferring the posterior distribution $p(\xi|X^*)$ for the whole map using (2) and (3). The GP fits a spatial correlation model to the data, where the correlation of the response between two observations decreases as the values of the independent variables become more distant.

V. BAYESIAN FUSION

Given the prior information $p(\xi|X) \sim \mathcal{N}(\mu^-, P^-)$, output of the GP inference (2) and (3), as the spatially correlated map,

$$\mu^- = \begin{pmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_n \end{pmatrix} \quad P^- = \begin{pmatrix} P_{11} & P_{12} & \cdots & P_{1n} \\ P_{21}^2 & P_{22} & \cdots & P_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ P_{n1}^2 & P_{n2}^2 & \cdots & P_{nn} \end{pmatrix} \quad (6)$$

and the new m independent measurements $\mathbf{z} = [z_1 \cdots z_m]^\top$ at positions $[\mathbf{x}_1 \cdots \mathbf{x}_m]^\top \subset X$ modelled as $p(z_i|\xi_i, \mathbf{x}_i) = \mathcal{N}(\mu_i, \sigma_{g,i})$,

$$\mathbf{z} = \begin{pmatrix} z_1 \\ z_2 \\ \vdots \\ z_m \end{pmatrix} \quad R = \begin{pmatrix} \sigma_{g,1}^2 & 0 & \cdots & 0 \\ 0 & \sigma_{g,2}^2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 \cdots & \sigma_{g,m}^2 \end{pmatrix} \quad (7)$$

with $\sigma_{g,i}^2$ as the noise variance for each sensor measurement, it is straightforward to perform the fusion. The maximum *a posteriori* estimator is used to integrate the prior maps together with the new independent measurements (with constant or non-constant noise). It is formulated as

$$\underset{\xi}{\operatorname{argmax}} p(\xi|\mathbf{z}, X). \quad (8)$$

Then the posterior density is $p(\xi|\mathbf{z}, X) \propto p(\mathbf{z}|\xi, X) \times p(\xi|X)$ computed as

$$\begin{aligned} P^+ &= ((P^-)^{-1} + (H^\top RH)^{-1})^{-1} \\ \mu^+ &= \mu^- + P^+((P^-)^{-1}\mu^- + (H^\top RH)^{-1}\mathbf{z}), \end{aligned} \quad (9)$$

where H is a $m \times n$ matrix that intrinsically selects part of the state $\xi_1 \dots \xi_m$ that is observed through \mathbf{z} . H is equal to $I_{n \times n}$ if the full state ξ is observed by \mathbf{z} .

Note that (9) and (10) are equivalent to the Kalman (or Information) filter update.

VI. APPLICATION TO THICKNESS MAPPING

In this section we present the approach described in Sections III, IV and V applied to pipe thickness mapping problem for assessing the condition on metal pipes.

The sensors currently used in assessing the condition of metal pipes may vary from acoustic (ultrasonic and hydrophones) that measure time-of-flight [17], magnetic flux leakage (MFL) that measure variations in magnetic fields [18], and remote Eddy currents that measure both time-of-flight and signal strength of a varying electromagnetic field [19].

3D range sensors such as lasers or cameras are also employed to obtain the external 3D profile of the pipe sections. The use of these sensors depends on the material and whether the corrosion or thinning can be exposed (*e.g.* grit-blasting the pipe). The resulting assessment is only external (*i.e.* from the outer surface), with no information given about the condition in the inner surface. Nevertheless, 3D profiling is a widely used alternative because of the high resolution and the ability to be done in-house without engaging expensive tools, which use MFL or Eddy currents sensors.

Acoustic sensor such as ultrasonic testing (UT) contact probes are also widely used because of the low-cost and easy interpretation. The disadvantages of UT sensors are that usually only spot measurements can be taken, and that they are sometimes unreliable in certain materials such as cast-iron. Combining 3D profiling with UT measurements, however, seems to be a viable option to get reliable thickness maps, but it has rarely been exploited by the condition assessment industry.

The aim in here is to integrate, in a probabilistic manner, information from the external 3D profile obtained by a laser range scanner and UT measurements. A probabilistic thickness map, in the form of an estimated thickness value and its associated uncertainty, will be the output of our approach. First, from an external 3D surface and assuming a regular inner surface, a prior thickness map with its associated covariance matrix that considers the spatial correlation will be generated. As described in previous sections, a GP will be used to learn and infer this spatial correlation.

UT independent thickness measurements and its associated variance will be integrated with the prior map through the Bayesian estimator to produce the posterior thickness map.

The following sections will detail sensor models and implementations for this particular application.

A. From 3D Surfaces to Thickness

A high-resolution 3D laser scanner from Creaform has been used to produce a 3D profile of a 660mm diameter cast-iron water pipe section. Figure 2 shows a picture of the pipe section. Both outer and inner surfaces of this pipe have been scanned to produce what we call “ground truth” to evaluate the proposed approach. The full 3D profile of a section of the pipe is shown in Figure 1a).

In real condition assessment scenarios, scanning the inner surface is not possible unless the pipe has a relatively large diameter (such in this case) and, more importantly, it is exhumed from the pipeline. For this reason, only the outer surface is commonly used to produce inaccurate thickness measurements, based on virtual gauges and nominal thicknesses. This motivates the use of other sensor modalities to obtain better thickness estimates.

In this work, in order to produce a measurement of thickness from only the outer surface, we have opted for fitting a concentric mesh of a cylinder using the nominal thickness value to decide its diameter. The nominal thickness value is in general well-known for any pipeline.

Given two concentric surfaces, we have developed an algorithm to compute the distance between the surfaces, *i.e.* the thickness of a pipe section. The algorithm is based on the ray tracing method with octrees developed in [20] implemented in the Point Cloud Library [21]. The mesh surface is converted to a point cloud and further into an octree representation. The pipe is ray traced in cylindrical coordinates from the centre point to the outer surface. The voxel resolution used for our purposes is 2mm. The output of the algorithm is a thickness measurement y_i at $\mathbf{x}_i = (x_i, c_i)$ location, where x is the axial position and c circumferential position in mm. An example plot of the output of the algorithm is presented in Figure 1b).

This algorithm works for both cases, when having only the external surface scanned with a inner concentric cylinder fitted, and when having the outer and the inner surfaces available.

B. Prior thickness map

Following the methodology presented in Section III the data set $\Psi_1 = \{(\mathbf{x}_0, y_0), (\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ is generated using the ray tracing algorithm. We have used 3674 thickness measurements from a 22×167 grid of 12.5mm cell size that covers the whole circumference in 275mm. Note that the high resolution achieved by the ray tracing algorithm has been deliberately reduced by averaging all the measurements within the cell size. This has been done to match the cell size of the UT sensor.

The mean μ^- and the Covariance P^- for the full grid are inferred following (2) and (3), after having learned the hyperameters of the Matér kernel (5). Cross validation led to the conclusion that $\nu = \frac{3}{2}$ produced the best results for our




Fig. 2. Picture of the cast-iron water pipe used in the experiments. On the left side, a grid of 22×167 cells has been painted on the surface. This grid shows the section measured using UT.

dataset. Note that $X^* = X$ in this particular data set, because we are only after the spatial correlation of the training points captured by P and there is no need to interpolate. The inferred P turns out to be almost block diagonal as expected, showing that each position is only correlated with its close neighbours.

Having the ground-truth, *i.e.* thickness obtained from the both inner and outer surfaces at high resolution, allows us to characterised the deviation associated to the thickness measurements. The sensor noise is set to $\sigma_n = 3.8\text{mm}$.

C. UT Modelling

Ultrasonic testing contact sensors are used for obtaining measurements of the thickness of pipes. UT contact probes emit a high frequency acoustic wave, which will go through a thin film (coupling) penetrating into the pipe. This wave will bounce back from the back wall of the pipe and come back to probe in the end. Once the time of flight of the reflected wave, the speed of sound in the pipe and the probe delay, and the coupling between the probe and the pipe are measured, the thickness of the pipe can be calculated with $d = c_s * \frac{t}{2}$, where c_s is speed of the sound in the material and t is the measured round-trip transit time.

In our experiments to obtain the measurements of Ψ_2 data set, a dual element direct contact UT probe with 1 MHz frequency has been used. The field-of-view of the ultrasonic probe was determined to be around 12.5 mm, giving the averaged thickness under this area. The cell-size of the grid is therefore chosen 12.5×12.5 mm. Figure 2 shows the grid used in the actual pipe.

Due to the particular material of this pipe (cast iron) and the roughness of the surface the UT signals can be unreliable. In order to capture this unreliability and associate an uncertainty to it, we propose to use a kernel density estimation (KDE) [22]. KDE is employed to estimate the probability of the time of arrival of the reflected wave and hence the thickness of the pipe. The possible thickness range for this pipe varies from 0 mm to 40 mm. Thus, a probability density function (PDF) built by the KDE is constructed for a certain range of possible arrival times of the reflected wave, which correspond to that interval. The mean value and standard deviation of the PDF, after being converted into thickness, are the cell's thickness and uncertainty value

respectively. The noise associated to each UT measurements is therefore non-constant.

In cast iron pipes the signal of the reflected wave is much weaker than that of steel, which is a ultrasonic's widely analysed material that reflects back most of the emitted signal. Unlike the steel plates, which reflected the wave including the initial bounced pulse followed by the attenuated fluctuating tail, the majority of reflected waves of cast iron only return a few short pulses as shown in Figure 3. In these circumstances, it is reasonable to assume that the PDF follows the normal distribution, whose mean value is near the back wall reflected pulse. Under this assumption, the kernel is chosen as Gaussian and its bandwidth determined by

$$h = \left(\frac{4\hat{\sigma}^5}{3n} \right)^{\frac{1}{5}} \approx 1.06\hat{\sigma}n^{-1/5}, \quad (11)$$

as described by [22]. The power of each time sample signal, which reflects the energy of the bounced wave, is used as weighting function for that time index during the construction of the PDF using KDE. One example of the constructed PDF is shown in Figure 4.




Fig. 3. Ultrasound wave signal in a cast iron pipe.




Fig. 4. PDF constructed by KDE on a certain range of time sample, which correspond to 0-40 mm possible pipe thickness.

In order to validate the proposed PDF function, 50 random measurements were taken from the pipe at a location with a good UT signal and another 50 at a location with a weak UT signal. A traditional threshold based method was used to determine the pipe wall thickness and a histogram was generated for representing the probability of the thickness at the two different location as shown in Figures 5 and 7. Then, one measurement from each of the above two locations was chosen randomly and the PDF was constructed using

the KDE based method as shown in Figures 6 and 8. From these plots, it is observed that the proposed PDF function based on KDE models the UT sensor in a reasonable manner. Note that high uncertainty values usually come from the pipe wall with irregular surface resulting in very poor or even no bounced wave. This type of measurements should, in a fusion framework, contribute very little to the estimate due to their lack of reliability.




Fig. 5. Histogram of a pipe wall thickness at a location where reflected ultrasound signal is clear, based on 50 random measurements.




Fig. 6. PDF of a pipe wall thickness at a location where reflected ultrasound signal is clear. Raw signal is shown in blue and PDF in green.




Fig. 7. Histogram of a pipe wall thickness at a location where reflected ultrasound signal is weak, based on 50 random measurements.

UT measurements from the pipe section were obtained using the KDE approach for the a total length of 275 mm by the full circumference (2087.5mm). As in the case of the 3D profile data set Ψ_1 , a total of 3674 independent measurements with uncertainty were generated to cover the same grid forming the data set Ψ_2 . The plots of these data are shown in the results section.

An extra set of independent measurements using this approach were generated for a small section of the circumference on a grid of 22×45 .

VII. RESULTS

As described in Section III, after having learned the prior maps and with the UT measurements and their associated




Fig. 8. PDF of a pipe wall thickness at a location where reflected ultrasound signal is weak. Raw signal is shown in blue and PDF in green.

TABLE I
COMPARISON RESULTS FROM THICKNESS ESTIMATION BETWEEN
CORRELATED P AND UNCORRELATED P FUSION.

	Naïve-2 Sensors	Correlated-2 Sensors	Correlated-3 Sensors
mse	6.774	3.476	2.122
max error	5.98	7.53	7.692

uncertainty available, it is straightforward to perform the Bayesian fusion using (9) and (10).

The results are presented in Figure 9 as thickness maps, where the prior, measurements and posterior, both mean and covariances, together with the ground-truth are plotted.

We have compared the results of the proposed approach with a naïve fusion approach, *i.e* considering independent variances as a prior. The fusion of this approach is done through the sum of weighted variances as follows:

$$\mu_i^+ = \frac{\sigma_{n,i}^{-2} y_i + \sigma_{g,i} z_i}{\sigma_{n,i}^{-2} + \sigma_{g,i}^{-2}} \quad (12)$$


$$\sigma_i^{2+} = \frac{1}{(\sigma_{n,i}^{-2} + \sigma_{g,i}^{-2})}. \quad (13)$$

The results of the naïve fusion approach are displayed in Figure 10. Table I presents the mean square error (mse) for both fusion approaches using the ground-truth, *i.e* thickness values resulting from ray tracing the inner and outer 3D surfaces of the pipe section, for evaluation.


The resultant posterior mean, from integrating UT measurements with the thickness extracted of the laser 3D profile using the proposed approach (third plot in Figure 9a)), shows an improvement from the fusion using the naïve approach (third plot in Figure 10a)) when compared with the ground-truth (fourth plot in Figure 9a)) although the variances look the same. This is because only the elements of the diagonal of P^+ were extracted to plot the 2σ values for the prior and posterior. The cross-correlation terms, however, are used in (10) to compute μ^+ getting a better thickness estimate as shown in Table I.

In order to show the simplicity of the proposed approach to fuse multiple sensor data, we have integrated another set of UT measurements³ to the previous data sets. For fusing this new set of measurements the output of the previous fusion is

³This data set was taken only in a small area of the pipe section with a different coupling. The area covers only 45 measurements on the circumference along the same 22 in the axial direction, corresponding in location to the first 45 cell starting from the bottom for the previous grid.



(a) Mean



(b) 2σ


Fig. 9. Estimated thickness maps with the proposed approach for a 275mm pipe section around the full circumference. a) plots are the mean prior, UT measurements \mathbf{z} , mean posterior and ground truth from left to right. b) plots are 2σ of prior P^- , measurements R and posterior P^+ from left to right.

used as a prior and, (9) and (10) are used to update this prior. The results are presented in Figure 11 and Table I showing a clear improvement, especially in the area covered by the new measurements.


The new thickness estimate after fusing the 3 data sets shows a clear improvement in mean (third plot in Figure 11a)) and reduced covariance (third plot in Figure 11b)).

VIII. CONCLUSIONS

This paper proposes a probabilistic framework to integrate heterogeneous sources of information into a single map. The fusion technique is generic and applicable as a general 2.5D mapping methodology. Gaussian process are used to incorporate and handle uncertainty in a statistically sound manner and the represent spatially correlated data. This correlation poses dependencies between points on the map which are not accounted for by many common 2.5D mapping techniques. The Gaussian process model learned from one data source is used as a prior of a standard Bayesian fusion. Our approach is an algorithm capable of generating accurate estimates at arbitrary resolutions to suit many applications. It



(a) Mean



(b) 2σ

Fig. 10. Estimated thickness maps with the naïve approach for a 275mm pipe section around the full circumference. a) plots are measurements Y , measurements \mathbf{z} , mean posterior and ground truth from left to right. b) plots are $2\sigma_f$, $2\sigma_{g,i}$ and $2\sigma^+$ computed from the uncorrelated variances from left to right.

also provides inferences with associated variances. Crucially, the technique can handle noisy data from multiple sources, and fuse it into a robust common probabilistic representation.

The proposed framework has been applied to the novel application of pipe thickness mapping. A dedicated algorithm for extracting the thickness from concentric 3D surfaces has been implemented. Moreover, a new technique based on KDE has been proposed to obtain measurements and uncertainty for contact UT sensors. The thickness information from a 3D surface and a fitted cylinder (at the nominal thickness) has been fused with UT thickness measurement to produce a common thickness map with mean and uncertainty estimation.

The results have been evaluated using the ground-truth extracted from the inner and outer 3D profiles and compared with the naïve approach that does not consider the spatial correlations.

We have chosen the thickness data from the external 3D surface as a prior, however one can choose the UT information instead. Choosing UT to learn the spatial correlation is

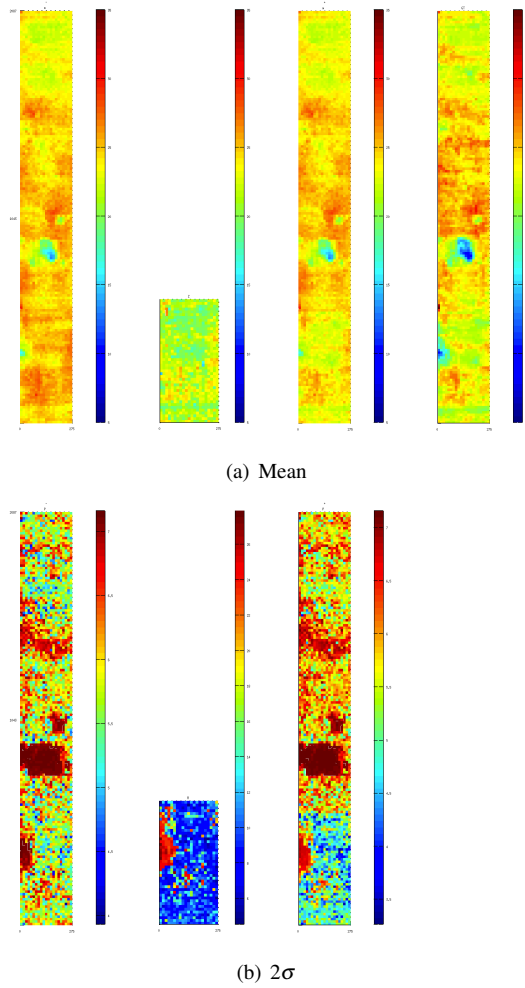


Fig. 11. Estimated thickness maps with the proposed approach (3 data sets) for a 275mm pipe section around the full circumference. a) plots are the mean prior (posterior of the first fusion as shown in Figure 9), new UT measurements, mean posterior and ground truth from left to right. b) plots are 2σ of prior P^+ (posterior of the first fusion as shown in Figure 9), measurements noise R and posterior from left to right.

possible in principle, however the fact that the measurement noise varies should be handled by the GP. Approaches such as in the one presented in [23] could be used.

A more efficient representation to allocate large covariance matrices and the use of submaps (where the cross-correlation over distant points is disregarded) should be considered in the case of large data sets to avoid the high computational complexity of the fusion algorithm due the inversion of a large covariance matrix. As depicted in the results the learned covariance matrix tends to be block diagonal for thickness mapping, therefore disregarding distant points is a reasonable approximation.

ACKNOWLEDGMENT

This publication is an outcome from the Critical Pipes Project funded by Sydney Water Corporation, Water Research Foundation of the USA, Melbourne Water, Water Corporation (WA), UK Water Industry Research Ltd, South Australia Water Corporation, South East Water, Hunter Water Corporation, City West Water, Monash University, University

of Technology Sydney and University of Newcastle. The research partners are Monash University (lead), University of Technology Sydney and University of Newcastle.

REFERENCES

- [1] R. T. Whitaker and J. Gregor, "A maximum-likelihood surface estimator for dense range data," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 10, pp. 1372–1387, 2002.
- [2] A. Kelly and A. Stentz, "An analysis of requirements for rough terrain autonomous mobility," in *Part II: Resolution and Accuracy, IEEE International Conference of Robotics and Automation*, 1997.
- [3] M. Montemerlo and S. Thrun, "A multi-resolution pyramid for outdoor robot terrain perception," in *In Proc. of the National Conference on Artificial Intelligence (AAAI)*, 2004.
- [4] P. Pfaff, T. R., and W. Burgard, "An efficient extension to elevation maps for outdoor terrain mapping and loop closing," *Int. J. Rob. Res.*, vol. 26, no. 2, pp. 217–230, Feb. 2007.
- [5] C. Rasmussen and C. Williams, *Gaussian Process for Machine Learning*. Cambridge: MA:MIT press, 2006.
- [6] S. Vasudevan, "Data fusion with gaussian processes," *Robotics and Autonomous Systems*, vol. 60, no. 12, pp. 1528–1544, 2012.
- [7] P. Kitanidis, *Introduction to geostatistics: applications to hydrogeology*, ser. Stanford-Cambridge Program. Cambridge University Press, 1997.
- [8] T. Lang, C. Plagemann, and W. Burgard, "Adaptive non-stationary kernel regression for terrain modeling," in *Robotics: Science and Systems (RSS)*, Atlanta, Georgia, USA, June 2007.
- [9] S. Vasudevan, F. Ramos, E. Nettleton, H. Durrant-Whyte, and A. Blair, "Gaussian process modeling of large scale terrain," in *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, 2009, pp. 1047–1053.
- [10] M. Smith, I. Posner, and P. Newman, "Efficient non-parametric surface representations using active sampling for push broom laser data," in *Proceedings of Robotics: Science and Systems VI*, Zaragoza, Spain, June 2010, 06.
- [11] S. T. O'Callaghan and F. T. Ramos, "Gaussian process occupancy maps," *I. J. Robotic Res.*, vol. 31, no. 1, pp. 42–62, 2012.
- [12] M. El-Beltagy and W. A. Wright, "Gaussian processes for model fusion," in *ICANN*, ser. Lecture Notes in Computer Science, G. Dorffner, H. Bischof, and K. Hornik, Eds., vol. 2130. Springer, 2001, pp. 376–383.
- [13] R. Murray-Smith and B. Pearlmuter, "Transformations of gaussian process priors," in *Deterministic and Statistical Methods in Machine Learning, First International Workshop*, Sheffield, UK, 2004. Springer.
- [14] M. Girolami, "Bayesian data fusion with gaussian process priors: An application to protein fold recognition," 2006.
- [15] S. Reece and S. J. Roberts, "An introduction to gaussian processes for the kalman filter expert," in *FUSION*. IEEE, 2010, pp. 1–9.
- [16] A. Brooks, A. Makarenko, and B. Upcroft, "Gaussian Process Models for Indoor and Outdoor Sensor-Centric Robot Localization," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1341–1351, Dec. 2008.
- [17] M. Bracken and D. Johnston, "Acoustic methods for determining remaining pipe wall thickness in asbestos cement, and ferrous pipes," in *Procs of Pipelines 2009: Infrastructure's Hidden Assets.*, 2009, pp. 271–281.
- [18] C. Edwards and S. B. Palmer, "The magnetic leakage field of surface-breaking cracks," *Journal of Physics D: Applied Physics*, vol. 19, no. 4, p. 657, 1986. [Online]. Available: <http://stacks.iop.org/0022-3727/19/i=4/a=018>
- [19] D. Atherton, "Remote field eddy current inspection," *IEEE Transactions on Magnetics*, vol. 31, no. 6, pp. 4142–4147, 1995.
- [20] J. Revelles, C. Urea, and M. Lastra, "An efficient parametric algorithm for octree traversal," in *Journal of WSCG*, 2000, pp. 212–219.
- [21] R. B. Rusu and S. Cousins, "3D is here: Point Cloud Library (PCL)," in *IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9–13 2011.
- [22] B. Silverman, *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall, 1998.
- [23] K. Kersting, C. Plagemann, P. Pfaff, and W. Burgard, "Most likely heteroscedastic gaussian process regression," in *International Conference on Machine Learning (ICML)*, Corvallis, Oregon, USA, March 2007.