

Laser-Radar Data Fusion with Gaussian Process Implicit Surfaces

Marcos P. Gerardo-Castro, Thierry Peynot and Fabio Ramos

Abstract This work considers the problem of building high-fidelity 3D representations of the environment from sensor data acquired by mobile robots. Multi-sensor data fusion allows for more complete and accurate representations, and for more reliable perception, especially when different sensing modalities are used. In this paper, we propose a thorough experimental analysis of the performance of 3D surface reconstruction from laser and mm-wave radar data using Gaussian Process Implicit Surfaces (GPIS), in a realistic field robotics scenario. We first analyse the performance of GPIS using raw laser data alone and raw radar data alone, respectively, with different choices of covariance matrices and different resolutions of the input data. We then evaluate and compare the performance of two different GPIS fusion approaches. The first, state-of-the-art approach directly fuses raw data from laser and radar. The alternative approach proposed in this paper first computes an initial estimate of the surface from each single source of data, and then fuses these two estimates. We show that this method outperforms the state of the art, especially in situations where the sensors react differently to the targets they perceive.

1 Introduction

The ability to build high-fidelity representations of the environment is critical for autonomous robots [14]. Range scanners such as laser range finders and radars are widely used in field robotics [10]. However, like any sensor, they suffer from limitations, e.g. in terms of field of view, resolution and noise [1]. Consequently, robots need techniques that can estimate accurate representations of the environment with incomplete data and uncertainty.

Gaussian Processes (GP) have become a popular technique in recent literature of robotic perception, due to their ability to learn spatial representations from noisy data in a non-parametric Bayesian fashion [12]. Gaussian Process Implicit Surfaces (GPIS) [19] is a mechanism to estimate the surface of an object with uncertainty, within a GP framework, by representing the geometry of the object as an Implicit Surface [15]. GPIS applied on range data offers a number of benefits to overcome

Australian Centre for Field Robotics (ACFR), The University of Sydney NSW 2006, Australia.
{m.castro, t.peynot, f.ramos}@acfr.usyd.edu.au

the aforementioned problems. Firstly, the generated model is fully predictive as it is able to predict a surface at arbitrary regions of an object that were not entirely observed by the range sensor [5]. Secondly, the model also yields the uncertainty of the estimates, at any point of the surface. By capturing the correlations between points using parametrized covariance functions, only a limited number of points are required to learn an accurate model. In addition, the GP can automatically handle the model selection (parameter estimation) efficiently.

The use of data from multiple sensing modalities can help to obtain a more complete and accurate representation of an object. As these modalities sense the environment using different physical processes, they also respond differently to some materials, textures, or environmental conditions [1, 4]. Consider the example of a car with windows perceived by a laser and a radar. Lasers operate at near-infrared frequencies of the electromagnetic spectrum, which are close to visible wavelengths. Therefore, a window appears mostly transparent to their sensing. In contrast, a mm-wave radar operates at lower frequencies (e.g. around $94GHz$), and is getting more returns back from the surface of the windows. On the other hand, on the rest of the car, the laser sensing is more accurate [1]. Therefore, a more complete and more accurate representation of the car can be obtained by fusing the data from the two sensor modalities.

In this paper we perform the fusion of data from two distinct sensing modalities (a laser and a radar) within a GPIS framework, in the context of field robotics. We propose a thorough experimental analysis of the performance of 3D surface reconstruction from laser and radar data in a realistic field robotics scenario: an unmanned ground vehicle (UGV) scanning an outdoor environment. We first analyse the performance of the GPIS approach using raw laser data alone, and using raw radar data alone, with 4 different choices of covariance matrices and different resolutions of the input data, on a total of 8 objects with different geometries. We then evaluate and compare the performance of two different fusion approaches within the GPIS framework. The first approach directly fuses raw data from the two different sensor modalities in the GPIS, as in [2]. In the second approach, we fuse points extracted from two initial estimates of the object surface that were built using raw laser data only and raw radar data only, respectively. We show that this novel approach outperforms the first state-of-the-art approach, especially in cases where the sensing modalities react differently, perceiving different targets. Although we put special emphasis on the implementation of data fusion for laser and radar data in this paper, the frameworks for GPIS fusion may be used with different range sensors.

The paper is organised as follows. Sec. 2 discusses related work on surface reconstruction with uncertainty representations, GPIS, and multi-sensor data fusion. Sec. 3 develops the GPIS framework used in this work and the fusion methods considered. Sec. 4 describes the experimental setup and Sec. 5 presents the experimental evaluation and analysis. Finally, Sec. 6 proposes conclusions and elements of future work.

2 Related Work

The problem of estimating continuous representations of data from range sensors has been extensively studied in the recent literature [13, 8]. To build continuous representations from range sensor data while accounting for uncertainties, different variations of Gaussian Processes have been implemented in the robotics community. Gaussian Beam Processes [11] give an independent treatment of the noise using heteroscedasticity on the beams, however, this approach is limited to 2D scenarios. Other applications using GP in 3D scenarios are related to mapping and terrain modelling [17, 7]. These approaches adopt the same parametrisation problem by associating a single elevation value z with any given position (x, y) in 2D Euclidean space. While this way of mapping is effective for applications such as terrain modelling, it is not suitable for applications that consider the full 3D mapping case, e.g. full 3D modelling of an object, where there can be multiple elevation values for a given (x, y) . Implicit Surfaces (IS) is a representation that is appropriate for this case [15]. It takes advantage of the geometry and topology of the objects. GPIS is a framework to estimate IS surfaces with uncertainties using Gaussian Process [19]. This approach has been applied to range sensor data for robotics applications in different contexts, such as change detection [18], active learning [5], and grasping [2].

Gaussian Process has been shown to be a powerful tool for multi-sensor data fusion when considering noisy input data [7, 3]. This concept has been adopted in [6], which proposes a sensor fusion framework based on a mixture of GPs. The application focussed on affect recognition, rather than object representations. GP data fusion was also explored in [17], where the author fuses raw data from laser scanner and Global Positioning System, with distinct noise models for each data source. However, this approach uses a representation that only allows for a single elevation value at a given (x, y) location, which is not appropriate for representing objects, as mentioned above. A similar fusion approach was used in [2] within a GPIS framework. Raw data from lasers and tactile sensors are directly fused in a GPIS with multi-variance noise in the input dimensions. However, the paper shows limited experimental results and no error analysis is provided. In this paper we propose an extensive experimental analysis of the performance of sensor data fusion within a GPIS framework, in the context of field robotics.

3 Gaussian Process Implicit Surfaces

3.1 Implicit Surfaces

Consider a set of points $X = [x_1, x_2, \dots, x_n]$ in Euclidean space corresponding to observations of the object. In order to model an object represented by X , an Implicit Surface is defined as the 0-level set of real-valued function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$, where the function f specifies whether a point x is inside the surface ($f(x) > 0$), outside the

surface ($f(x) < 0$), or on the surface ($f(x) = 0$). Such constraints values are assigned to the variable Y , so that $Y = [f(x_1), f(x_2) \dots f(x_n)]$. Direct observations made by range scanners are usually assumed to be on the surface of the object, therefore, *zero-value constraints* ($f(x) = 0$) are assigned to the sensor measurements. Additional observations of points known to be inside or outside the object may be added to help the estimation process.

3.2 Gaussian Process Implicit Surfaces

Gaussian process regression can be used to provide the surface estimate $f_*(x_*)$, with variance $\mathbb{V}(f_*(x_*))$ based on observation data from X and constraints Y (i.e. targets for the GP), defined as training data under the GPIS framework. This can be formulated as:

$$\mathbb{P}(f_*(x_*) | X, Y, \theta, x_*) = \mathcal{N}(\bar{f}_*, \mathbb{V}[f_*]), \quad (1)$$

where θ is a set of hyper-parameters. The mean \bar{f}_* and variance $\mathbb{V}[f_*]$ at a selected point x_* given the measured data X are:

$$\bar{f}_* = k(x_*, X)^T (K + \sigma_n^2 I)^{-1} Y \quad (2)$$

$$\mathbb{V}[f_*] = k(x_*, x_*) - k(x_*, X)^T (K + \sigma_n^2 I)^{-1} k(x_*, X), \quad (3)$$

where K is a covariance matrix. The noise variance of the observed data is represented by $\sigma_n^2 I$. Note that σ_n^2 can be learnt along with the other GP hyper-parameters.

In this paper, we implement different stationary kernels: the exponential covariance function,

$$k = k(x_i, x_j) = \sigma_f^2 \sum_{k=1}^d \exp(-(\frac{\Delta_k}{\ell})^\gamma), \quad (4)$$

and two variants of the Matérn type [12],

$$k_{M_{3/2}} = \sigma_f^2 \sum_{k=1}^d (1 + \frac{\sqrt{3}\Delta_k}{\ell}) \exp(-\frac{\sqrt{3}\Delta_k}{\ell}),$$

$$k_{M_{5/2}} = \sigma_f^2 \sum_{k=1}^d (1 + \frac{\sqrt{5}\Delta_k}{\ell} + \frac{5\Delta_k^2}{3\ell^2}) \exp(-\frac{\sqrt{5}\Delta_k}{\ell}).$$

The hyper-parameter σ_f^2 represents the signal variance and the length-scale is represented by ℓ , $\Delta_k = |x_i - x_j|$. $k_{M_{3/2}}$ and $k_{M_{5/2}}$ are stationary covariance functions used to amplify the sensitivity between the correlations of the points compared to the widely used square exponential (Eq. (4) with $\gamma = 2$), which produces a smooth kernel that drops off with distance. We extend our analysis to the exponential covariance function ($\gamma = 1$), which is even more sensitive to changes.

An important aspect of the GP is the optimisation of the hyper-parameters $\theta = (\sigma_f, \ell, \sigma_n)$. In this paper this was done by maximising the log-marginal likeli-

hood. The Cholesky decomposition was used to obtain the predictors (\tilde{f}_* and $\mathbb{V}[f_*]$) and the log-marginal likelihood [12]. Once an estimate f_* of the surface has been obtained, 3D surface points and corresponding variances are then computed for values of $f_* = 0$ in Eq. (2) and Eq. (3) by querying in a region pre-defined by x_* .

3.3 Single-Sensing-Modality GPIS

We define a *single-sensing-modality GPIS* as a GPIS whose input is a set of data provided by a single sensor type (in this paper, laser or radar), as illustrated in Fig. 1. We name the process $GPIS_i$, The input data X_i and Y_i , the estimated surface \tilde{f}_{i*} , and

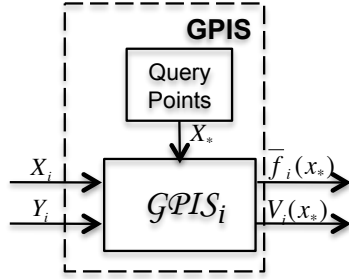


Fig. 1 $GPIS_i$ process, using laser ($i = L$) or radar ($i = R$) input data.



Fig. 2 Argo UGV equipped with the laser and radar sensors used in this study.

the variance \mathbb{V}_{i*} , where the index i specifies the nature of the input data: $i = L$ if the input data is provided by a laser, and $i = R$ if the data is from a radar. A global noise parameter σ_n^2 is used for all the input points in each $GPIS_i$ (see Eqs. (2) and (3)).

3.4 Multi-Sensor Data Fusion: $GPIS_{LR}$

The first fusion method, $GPIS_{LR}$, fuses two sets of raw data, X_L and X_R , acquired by laser and radar, respectively, in a single GPIS. The approach is similar to the one in [2]. The input data of $GPIS_{LR}$ is composed of all training points from each sensing modalities put together: $X = [X_L \ X_R]$ and $Y = [Y_L \ Y_R]$. A diagram of the process of $GPIS_{LR}$ is illustrated in Fig. 3(a). $GPIS_{LR}$ accounts for different noise parameters for each sensing modality by implementing an input-dependent noise process, i.e. heteroscedastic, similar to [11]. Let $\sigma^2 \in \mathbb{R}^n$ be the noise variances for n given sensing modalities. The predicted distributions become:

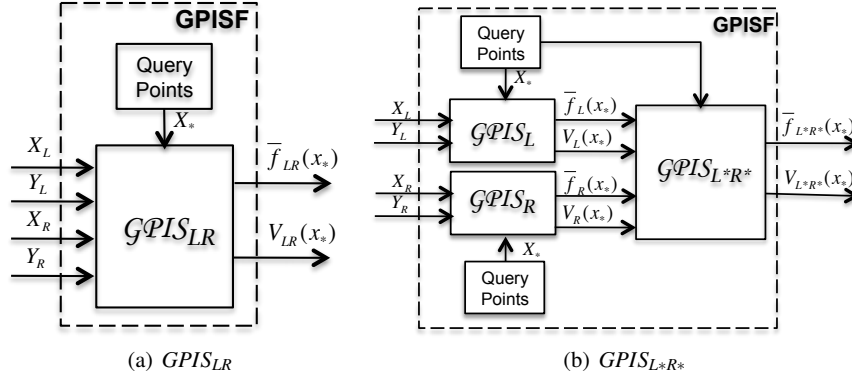


Fig. 3 The two alternative fusion processes: $GPIS_{LR}$ (a) and $GPIS_{L^*R^*}$ (b).

$$\bar{f}_* = k_*^T (K + H)^{-1} Y \quad (5)$$

$$\mathbb{V}[f_*] = k(x_*, x_*) - k_*^T (K + H)^{-1} k_* \quad (6)$$

where $H = \text{diag}(\sigma_1^2(X_1), \sigma_2^2(X_2) \dots \sigma_n^2(X_n))$ is a non-fixed noise matrix.

3.5 Alternative Fusion Method: $GPIS_{L^*R^*}$

In the alternative fusion approach we propose, $GPIS_{L^*R^*}$ (see Fig. 3(b)), we first estimate the object surface from raw laser points and from raw radar points separately, using two independent GPIS (i.e. $GPIS_L$ and $GPIS_R$). We then query a set of points, X_{L^*} and X_{R^*} , which are randomly sampled from the points where $\bar{f}_{L^*} = 0$ and $\bar{f}_{R^*} = 0$ respectively, along with the corresponding variances, \mathbb{V}_{L^*} and \mathbb{V}_{R^*} . The number of points in X_{L^*} and X_{R^*} is two times the original number of input points (i.e. in X_L and X_R). Associated constraints Y_{L^*} and Y_{R^*} are computed from the estimated surfaces \bar{f}_{L^*} and \bar{f}_{R^*} . The next step is to compute the final estimate using $GPIS_{L^*R^*}$, with inputs: $X = [X_{L^*} \ X_{R^*}]$ and $Y = [Y_{L^*} \ Y_{R^*}]$. The predicted uncertainties of \mathbb{V}_{L^*} and \mathbb{V}_{R^*} are integrated in the final GPIS as fixed noise parameters. We substitute H in Eqs. (5) and (6), with $H = \text{diag}(H_{L^*}, H_{R^*})$, where H_{L^*} is a fixed noise matrix, defined by the variances \mathbb{V}_{L^*} as $H_{L^*} = \text{diag}(\mathbb{V}_{L^*1}, \mathbb{V}_{L^*2} \dots \mathbb{V}_{L^*m})$, and, similarly, $H_{R^*} = \text{diag}(\mathbb{V}_{R^*1}, \mathbb{V}_{R^*2} \dots \mathbb{V}_{R^*p})$. m and p represent the number of points in X_{L^*} and X_{R^*} , respectively.

$GPIS_{L^*R^*}$ can be used to fuse two continuous surface estimates with uncertainties. Therefore, it may allow for a *consistency check* between the laser and radar perception prior to fusion. The data that passes this test should be fused to obtain a refined estimate, while the inconsistent data should not be fused. The implementation of this consistency check is left to future work.

4 Experimental Setup

4.1 Data Collection

Experiments were conducted with the Argo UGV (Fig. 2), used as a data collection platform. The Argo is equipped with a laser range finder, a mm-wave radar, and a *cm*-accuracy 6-DOF dGPS/INS localisation unit. The laser sensor is a Sick LMS-291 (range resolution: $0.01m$ and angular resolution: 0.25°) and the radar is a 94 GHz Frequency Modulated Continuous Wave (FCMW) radar, custom-built at ACFR (range resolution: $0.2m$ and angular resolution: 2°). The laser and radar were directed at the front of the vehicle with a constant nodding angle, so that the center of the beam intersected the ground at a look-ahead distance of approximately $11.4m$. The sensor data, along with the navigation data, were collected by the platform while it was moving around a rural environment. Consequently, the errors in the resulting 3D points were the result of the combination of 3 error sources: sensor noise, calibration and localisation. A detailed description of the platform, sensors and the datasets can be found in [9]. In particular, objects of different geometries (listed in Table 1) were partially scanned by the sensors on the UGV from distances varying from $2m$ up to $30m$ and used to evaluate the performance of the surface reconstruction techniques.

Table 1 List of objects, sizes and number of ground truth points.

Object	Comp.	Car	Wall	Wall2	Trailer	Pole	Pole2	Fence
Dim. (m^3)	3.3x1.7x1.4	2.9x2.8x2.0	14x3.1x1.7	9.0x2.5x9.0	4.6x4.6x0.4	0.4x0.4x0.4	0.4x0.4x0.4	4.5x4.5x4.5
Nb. Pts.	7,958	22,736	16,738	15,584	18,784	2,578	1,032	6,470

4.2 Data Pre-processing

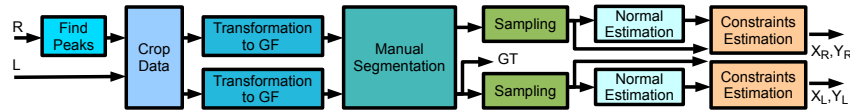


Fig. 4 Data Preparation Process.

The data pre-processing follows the process described in Fig. 4. The data provided by the laser sensor consist of a single range value for each bearing angle in its scan, which are the result of the target extraction developed by the sensor manufacturer. The raw radar data consist of multiple range values, with corresponding

intensities, for each bearing angle in its scan (note that the radar beam is much larger than the laser's). The extraction of the targets from the noise in the data is achieved using the approach presented by the authors in [4]. We first extract peaks (i.e. local maxima) for each bearing angle. Most robotics applications only consider the highest peak (global maxima) as a target detected by the radar, thereby using the radar as a laser. However, this leads to the loss of useful information contained in the rest of the radar beam. Therefore, along with the highest peaks, we also extract local maxima that correspond to secondary targets by using an adaptive intensity threshold [4]. The result is a set of 3D points per scan, similar to the data provided by a multi-echo laser sensor.

Laser and radar raw scans are then cropped to only keep data where the two sensors' FOVs overlap. Laser and radar points are then transformed into a common global navigation frame (GF). This transformation is obtained by combining the output of a prior extrinsic sensor calibration (using the technique in [16]) with the localisation of the Argo. The object of interest is then manually segmented from the full point cloud obtained with each sensing modality. A segmented object is a dense 3D point cloud, scanned from different perspectives. To evaluate the performance of the object reconstruction techniques, a small set of 5% of the data points is randomly sampled from the object. The rest of the *laser* data (i.e. 95% of the point cloud), constituting a dense point cloud, is used to build a ground truth (GT) (see Sec. 4.3).

For each sampled observation provided by the range sensor, the normal to the surface of the object at that point, N_i , is approximated by the perpendicular to the segment between this point and the closest point in the same scan. Points inside and outside the surface are computed using the *normal value constraints* [15]. We place two points on N_i : one outside the object at a given distance $d = -0.5m$ from the surface, and one inside at a distance $d = 0.2m$ (see Fig. 5). In the GPIS, these points constitute positive ($f(x) = 1$) and negative ($f(x) = -1$) constraints, respectively, and populate X_L, Y_L and X_R, Y_R which are used as input data of the GPIS, for the laser and radar, respectively.

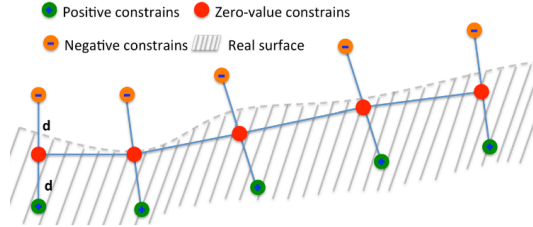


Fig. 5 The positive and negative constraints, given range sensor observations (red dots). The noise-free observations are assumed to be on the surface, represented by the dashed line.

4.3 Ground Truth (GT)

To quantify the errors made in the reconstructions, we used 95% of the full-resolution laser point cloud to obtain ground truth data, since the laser is the most accurate sensor available in our system. In this paper, the mean distance between any point in the full resolution point cloud and its closest neighbour was $3cm$. In some cases, due to the limitations of laser sensing, this dense point cloud of the object has to be corrected or completed to better reflect the actual surface of the object. For example, the side windows of the car shown in Fig. 8(a) were poorly represented (see Fig. 8(b) before correction), since the laser did not provide many returns from the windows surface. Therefore, points on the windows were manually added to complete the ground truth, using a resolution and noise level comparable to those of the original point cloud.

5 Experimental Results

5.1 Single-Sensing-Modality GPIS

Surfaces were estimated from laser (X_L) and radar (X_R) data separately, using the $GPIS_L$ and $GPIS_R$ processes described in Sec. 3.3. We evaluated the surfaces obtained with different resolutions of input data, and with the different covariance functions mentioned in Sec. 3.2. For all evaluations herein we queried the GPIS using a 3D grid of points (see x_* in Fig. 1) uniformly spaced at a resolution of $1cm$. In each of the 3D grid cells, values where $f(x_*) = 0$ were selected to represent the surface estimate. An example of a 3D surface reconstruction generated using GPIS with 376 input points is showed in Fig. 6. From a limited number of points, the general geometry of the object was recovered. A quantitative evaluation of the accuracy of surface estimates was performed by computing the 3D distances between points extracted from the estimated surfaces and our ground truth (the full resolution laser point cloud). Error statistics for each object were then obtained by calculating the root mean square (RMS) of these distances, providing an RMS error (RMSE) for each object. The analysis was performed for all 8 objects we considered, but for conciseness, we only show the RMSE for the different surface estimates of the compressor, in Fig. 7.

With an RMSE of $0.069m$ and $0.125m$ for laser and radar GPIS surfaces respectively, the exponential covariance function outperforms ($\gamma = 1$ in Eq. (4)) all the other kernels. Furthermore, the RMSE decreases as more sampled points are used, until the error does not change significantly.

Surface estimates obtained with $GPIS_L$ are more accurate than with $GPIS_R$. This was expected considering the higher accuracy and resolution of the laser sensor (see Sec. 4.1). However, in some cases radar surfaces showed an extended coverage of the object compared with laser surfaces. For example, Fig. 8(c) shows that most

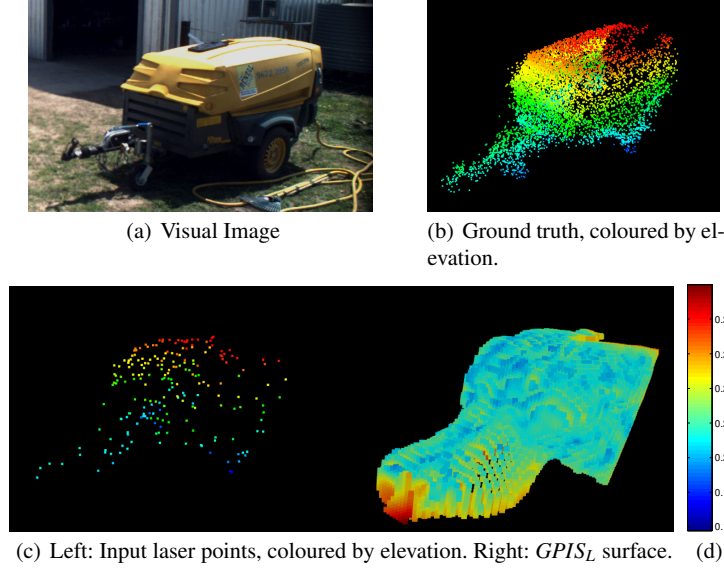


Fig. 6 Surface reconstruction of the compressor from sparse laser data. (a) Visual image of the compressor. (b) Full resolution 3D laser point cloud, used as ground truth only. (c) Surface estimated by $GPIS_L$ (right) using 376 input laser points (left). The surface is coloured by uncertainty (variance), using the colour scale shown in (d).

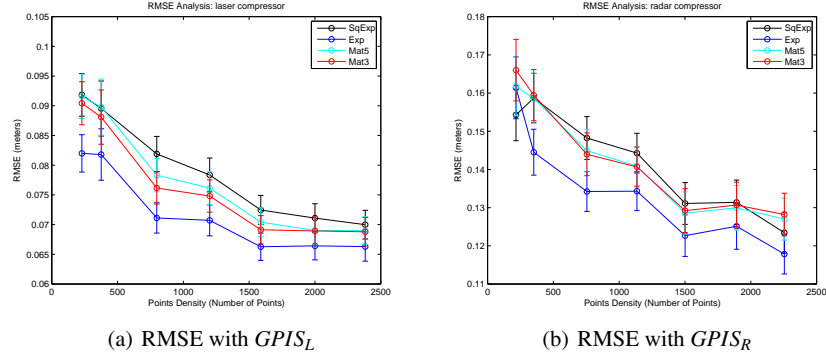


Fig. 7 RMSE of the compressor surface estimates obtained with $GPIS_L$ (a) and $GPIS_R$ (b) for different kernels and input point densities. Note the different scales of the errors obtained (larger errors with radar data).

parts of the car are quite well modelled by $GPIS_L$, but the surfaces of the windows are incomplete, since they are partially transparent to laser sensing. On the other hand, the radar perception of these windows is more complete, resulting in a more complete representation using ($GPIS_R$) (Fig. 8(e)).

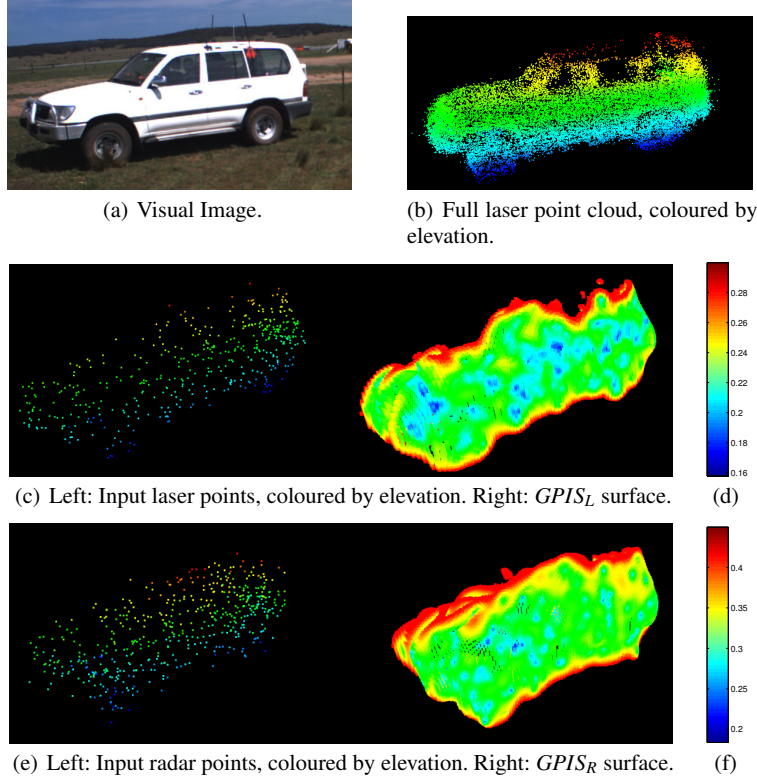


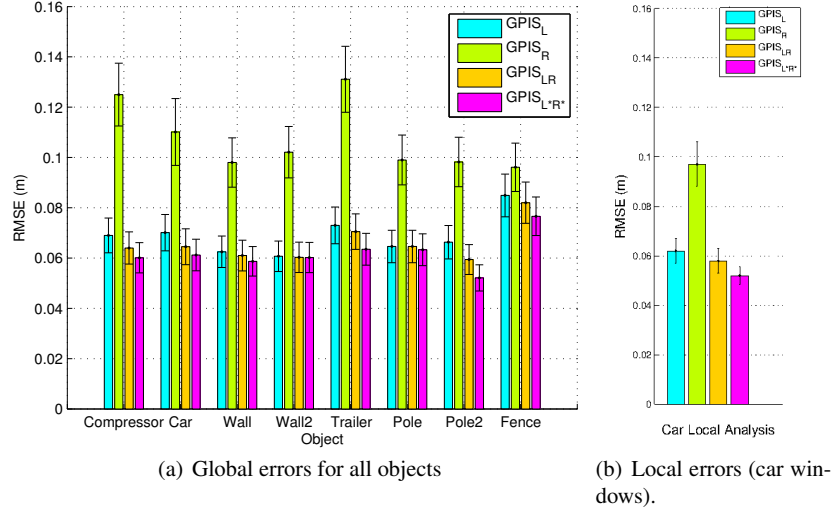
Fig. 8 Surface reconstruction of the car (a) using sparse laser or radar data. (b) Full resolution laser point cloud, used to build the ground truth data (shown *before* correction of the windows). (c) Surface estimated by $GPIS_L$ and (d) Surface estimated by $GPIS_R$. 1136 input points are used in each case. The surfaces are coloured by variance. Note the difference in scale between the variances for $GPIS_L$ and $GPIS_R$, the latter showing higher uncertainty overall. The edges of the car show high uncertainty since the car was only partially observed by both sensors.

5.2 Laser-Radar Data Fusion

Laser and radar data were fused using the $GPIS_{LR}$ and $GPIS_{L \times R}$ fusion methods described above. Table 2 shows the RMSE obtained using different covariance functions to estimate the surface of the compressor. The most accurate results seem to be obtained with the exponential kernel. However, the improvement is relatively insignificant, in particular for the $GPIS_{L \times R}$ method, whose accuracy is consistently the highest. We evaluated the performance of surface estimation on the 8 different objects listed in Table 1 using each GPIS method considered in this paper. Fig. 9(a) shows the RMSE for each surface estimate obtained using the exponential covariance function. As an example, Fig 10 shows the estimated surface of the car using $GPIS_{L \times R}$. The estimation of the surfaces of the compressor, the car and the second pole (Pole2) were significantly improved by the fusion ($GPIS_{LR}$). In addition, fur-

Table 2 RMSE (in m) obtained for the compressor using different covariance functions.

	$GPIS_L$	$GPIS_R$	$GPIS_{LR}$	$GPIS_{L^*R^*}$
SqExp	0.074	0.134	0.071	0.063
Exp	0.069	0.125	0.064	0.060
Mat3	0.071	0.132	0.069	0.063
Mat5	0.074	0.134	0.071	0.063

**Fig. 9** (a) RMSE (in m) obtained using the different $GPIS$ methods to estimate the surfaces of all 8 objects. (b) Local analysis of the RMSE in the area of the car windows, shown in Fig. 10. The black error bars represent the standard deviation.

ther improvement was obtained using the alternative fusion method ($GPIS_{L^*R^*}$). The trailer was not significantly better represented by the state-of-the-art fusion method ($GPIS_{LR}$), however, the improvement is clearer when using ($GPIS_{L^*R^*}$). For the other objects, such as the walls, the improvements are less significant. To better illustrate some of the benefits of the laser-radar fusion, Fig.9(b) proposes a local analysis of the errors obtained with the car, focussing on the area of the windows (see white box in Fig. 10). These results show that due to the poor perception of the windows by the car, the accuracy of the surface estimation was particularly improved by the laser-radar fusion process.

6 Conclusion

In this paper, we proposed an experimental analysis of the performance of continuous 3D surface reconstruction from laser and mm-wave radar data using Gaussian Process Implicit Surfaces, in a realistic field robotics scenario. We evaluated the

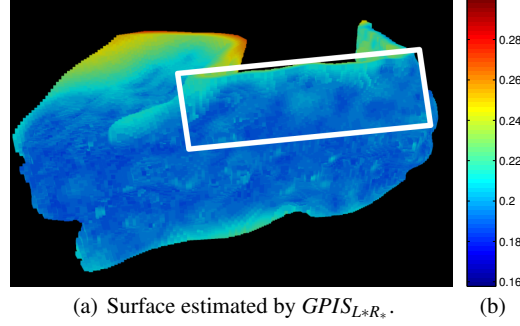


Fig. 10 Surface of the car (seen in Fig. 8(a)), reconstructed using sparse laser and radar data (1136 points) fused with the $GPIS_{L \star R_\star}$ method. The surface is coloured by uncertainty, with low uncertainty shown in blue (see corresponding colour bar in (b)). Note that the windows of the car were correctly reconstructed, with low uncertainty and no holes. The white box delimitates the area used to perform the local analysis of errors in the area of the windows.

performance of single-sensor approaches with different resolutions of input data and different kernels. We also compared the performance of these approaches with a state-of-the-art fusion approach and a new alternative method to multi-sensor data fusion. The GPIS fusion showed a significant improvement of the surface representations, especially when taking advantage of the complementarity of the two sensor modalities (e.g. in the case of the car windows, consistently detected by the radar but not by the laser). The proposed fusion process $GPIS_{L \star R_\star}$ outperformed the state-of-the-art fusion method.

The next step of this work will be to implement a test within the $GPIS_{L \star R_\star}$ framework to validate the consistency between the estimates obtained using laser and radar perception separately, prior to fusion. This will allow for higher resilience in challenging conditions, when laser and radar may not detect the same targets (e.g. in the presence of airborne dust, as in [4]). In this paper, the 3D data were manually segmented before applying the GPIS approaches. A direct extension will be the integration of automatic data segmentation in the same framework, or by analysing large scale environmental datasets rather than pre-segmented objects. In addition, the heterocedasticity treatment can be extended to sets of scans rather using a single noise parameter for all the points provided by each sensing modality.

Acknowledgements The authors thank Simon O’Callaghan for his input. This work was supported in part by the ACFR and the NSW State Government. This material is based on research sponsored in part by the Air Force Research Laboratory, under agreement number FA2386-10-1-4153. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

References

1. G. Brooker. *Sensors for Ranging and Imaging*. SciTech Publishing, Inc., 2009.
2. S. Dragiev, M. Toussaint, and M. Gienger. Gaussian process implicit surfaces for shape estimation and grasping. In *IEEE International Conference on Robotics and Automation*, 2011.
3. M. A. El-Beltagy and W. A. Wright. Gaussian processes for model fusion. In G. Dorffner, H. Bischof, and K. Hornik, editors, *Artificial Neural Networks ICANN*. Springer Berlin/Heidelberg, 2001.
4. M.P. Gerardo C. and T. Peynot. Laser-to-radar sensing redundancy for resilient perception in adverse environmental conditions. In *ARAA Australasian Conference on Robotics and Automation*, 2012.
5. G. A Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme. Active planning for underwater inspection and the benefit of adaptivity. *International Journal of Robotics Research*, 32(1), 2013.
6. A. Kapoor, H. Ahn, and R. Picard. Mixture of gaussian processes for combining multiple modalities. In *International Workshop on Multiple Classifier Systems*, 2005.
7. S. T. O’Callaghan, F. T Ramos, and H. Durrant-Whyte. Contextual occupancy maps incorporating sensor and location uncertainty. In *IEEE International Conference on Robotics and Automation*, 2010.
8. Y. Ohtake, A. Belyaev, M. Alexa, G. Turk, and H. Seidel. Multi-level partition of unity implicits. In *ACM SIGGRAPH Courses*, 2005.
9. T. Peynot, S. Scheduling, and S. Terho. The Marulan Data Sets: Multi-Sensor Perception in Natural Environment with Challenging Conditions. *International Journal of Robotics Research*, 29(13), 2010.
10. T. Peynot, J. Underwood, and S. Scheduling. Towards reliable perception for unmanned ground vehicles in challenging conditions. In *IEEE/RSJ International Conference on Robotics and Intelligent Systems*, 2009.
11. C. Plagemann, K. Kersting, P. Pfaff, and W. Burgard. Gaussian beam processes: A nonparametric bayesian measurement model for range finders. In *Robotics: Science and Systems III*, 2007.
12. C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
13. T. Whelan et al. Robust tracking for real-time dense rgb-d mapping with kintinuous. Technical Report MIT-CSAIL-TR-2012-031, CSAIL, MIT, 2012.
14. S. Thrun, W. Burgard, and D. Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3d mapping. In *IEEE International Conference on Robotics and Automation*, 2000.
15. G. Turk and J.F. O’Brien. Variational implicit surfaces. Technical Report GIT-GVU-99-15, Georgia Institute of Technology, 1999.
16. J. P. Underwood, A. Hill, T. Peynot, and S. J. Scheduling. Error modeling and calibration of exteroceptive sensors for accurate mapping applications. *Journal of Field Robotics*, 27(1), 2010.
17. S. Vasudevan. Data fusion with gaussian processes. *Robotics and Autonomous Systems*, 60(12), 2012.
18. A. W Vieira, PLJ Drews, and M. FM Campos. Efficient change detection in 3d environment for autonomous surveillance robots based on implicit volume. In *IEEE International Conference on Robotics and Automation*, 2012.
19. O. Williams and A. Fitzgibbon. Gaussian process implicit surfaces. In *Gaussian Processes in Practice Workshop*, 2007.