

Generating Implicit Surfaces from Lidar Data

Mike Smith

Department of Engineering Science
Oxford University
Oxford, UK

Email: mike@robots.ox.ac.uk

Ingmar Posner

Department of Engineering Science
Oxford University
Oxford, UK

Email: ingmar@robots.ox.ac.uk

Paul Newman

Department of Engineering Science
Oxford University
Oxford, UK

Email: pneuman@robots.ox.ac.uk

Abstract—This paper is concerned with generating a continuous implicit representation of a robot’s workspace using sparse point cloud data. We adopt a Gaussian Process (GP) framework to model the underlying workspace surfaces and suggest a non-parametric formulation which allows us to capture the non-functional relation between ground plane and elevation that is not possible with, for example, terrain mapping algorithms. The point clouds we are processing are typically vast such that blind application of a Gaussian Process leads to computational intractability. We therefore use a mixture-of-GPs model where individual GP support sets are chosen via segmentation in both laser and appearance space. We provide results that highlight the robustness of our algorithm, and apply our framework to semantically guide resampling of the workspace.

I. INTRODUCTION

3D point clouds of outdoor scenes are now ubiquitous. They originate from diverse sources such as, for example, the aggregation of lidar data from a moving platform, a sequence of stereo pairs or, more recently, flash laser sensors. The significant increase in workspace coverage provided by these 3D range sensors over traditional systems offers advantages in problems such as localisation and planning. However, while this technology paves the way towards new vistas of operation, it also presents significant challenges to be overcome in terms of the amount and quality of the data provided. The sequential nature of the data acquisition has previously implied a trade-off between sparsity of the coverage and the time required to gather data from an entire workspace. Recent developments in sensor technology have led to systems providing rapid and relatively dense coverage, at rates in excess of two orders of magnitude beyond existing systems¹. This puts a significant strain on existing algorithms.

This work addresses the challenge of forming an efficient, continuous 3D workspace representation from a 3D laser point cloud. Our motivation lies in the value of obtaining richer workspace descriptions generated by a mobile robot. We especially consider the domain of assisted autonomy where a human operator relies on

¹This comparison is based on a SICK LMS200 model and the more recent Velodyne high definition laser range scanner.

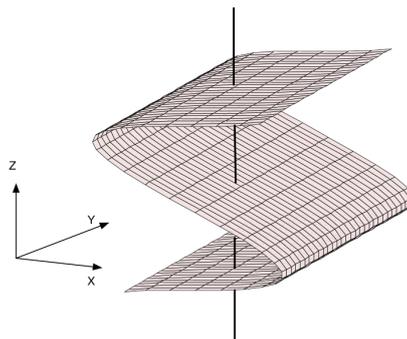


Figure 1: An illustration of a non-functional relation where multiple z values are associated with the same (x, y) pair.

a detailed and accurate representation of an otherwise inaccessible workspace. Concretely, based on the evidence of the original point cloud, our approach admits the calculation of every (x, y, z) point that lies on a sampled workspace surface. We adopt the formulation of [1] which is able to represent the *non-functional relation* between ground plane and elevation (see Figure 1) that is not possible with, for example, terrain mapping algorithms. This parameterisation allows us to model common structures such as walls, which require multiple elevation values to be associated with the same position on the Euclidean x - y -plane.

We recognise the value of uncertainty estimates in the surface predictions and therefore adopt a regression framework based on a mixture of Gaussian Processes. This Bayesian framework permits robustness to noisy and erroneous measurements (typical in highly reflective surfaces sampled by lasers). Uncertainty estimates also naturally reflect areas that have a low density of sampling and, together with the predictions, can be coupled with the corpus of Bayesian methods to provide a complete and principled processing chain. For example, one can imagine surface predictions from laser point clouds could be used to help refine vehicle trajectory estimates. As data becomes denser, we may also wish to leverage sparsification techniques found commonly in the GP literature (see [2]).

To enable richer workspace descriptions, and provide an example of the flexibility of this approach, we determine the subsets of the laser data used for the individual GPs by using a segmentation in appearance space. This choice is motivated by the common availability of classifiers which operate in this space and is envisioned to enable semantics-guided model-selection, filtering and resampling — e.g. as a precursor to CAD model generation — in the future. We note, however, that our framework is by no means bound to this form of segmentation, and demonstrate empirically that this approach to support-set selection does not adversely affect the GP regression performance. We illustrate an initial application for this framework using semantic information to resample a workspace.

The remainder of this paper is structured as follows. Section II provides an overview of related works. Our parameterisation of the workspace as well as the GP regression model adopted are introduced in Section III. We present results on real outdoor urban data in Section IV. Finally, we conclude in Section V.

II. RELATED WORKS

Within the robotics community, 3D surface reconstruction from sparsely sampled environments are commonly used to model an unknown terrain or in order to provide a faithful reconstruction of the workspace across scales. The latter is commonly achieved by meshing techniques (see, for example, [3]) whereby every datum of a point cloud forms the vertex of a polygon. Often such techniques are augmented by a mesh decimation step which aims to retain a faithful representation of the environment while at the same time reducing the amount of data to be stored. The creation of the mesh as well as the reduction of the number of vertices are based on sensor and environment dependent heuristics. Problems are encountered when the data present only non-homogeneous samples of a workspace and coverage is incomplete — which is commonly the case in the complex outdoor environments we consider.

An alternative approach considers a parametric representation based on geometric primitives [4, 5]. Planar representations are a common choice since they are ubiquitous in man-made environments. However, the reliance on geometric primitives places strong assumptions on the environment which are often not met in complex outdoor environments.

The requirement of robots to operate in ever larger, unstructured environments provides ample incentive for research into suitable models of the terrain encountered. Recent works in 3D reconstruction techniques trace their origin to this domain. Common approaches include discrete 2D, 2.5D and 3D methods such as [6] and [7, 8, 9, 10], which represent the world as a regular

grid of cells / voxels. Each cell stores the probability of it being occupied. As new sensor measurements arrive these probabilities are continuously updated. These representations suffer from the necessity to strike a trade-off between resolution (i.e. memory requirements) and volume covered. Like our algorithm, [6] also produces uncertainty estimates on surface predictions but through confidence maps which do not allow straightforward integration into a principled toolchain as we desire. As mentioned in [6], the technique is also prescriptive in the form of the surfaces permissible for a given point cloud. These algorithms also require further processing to produce a continuous surface representation.

In recent years there has been a further trend in the terrain modelling community to provide non-parametric continuous models of surfaces. In this context, GP models are particularly favoured for their ability to handle incomplete data in a principled probabilistic fashion. Examples of such approaches include [11, 12, 13]. While these approaches differ in their choice of (non-stationary) covariance function and GP sparsification method, they all adopt the same parameterisation of the problem — that is they model a *well defined function* $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ which associates a single elevation value z with any given position (x, y) in 3D Euclidean space. While this mapping is effective for terrain modelling it precludes an application in the full 3D mapping case where commonly encountered vertical structures such as walls require a *non-functional relation* representation. Our work remedies this shortcoming by using what we argue is a more natural parameterisation of the surface reconstruction problem.

While we share the common goal of faithful 3D surface reconstruction with the works discussed so far, the work bearing closest relation to the approach presented here is that of [14] and [1]. Gaussian Beam Processes [14] model laser data on a per-scan basis using a GP model to regress range on bearing, then applies this model to estimate 3DOF pose given a map. However, in this work we adopt the parameterisation proposed in [1] where the entire workspace of a moving laser scanner is expressed as a function of range, time and scan-angle (Section III). In doing so we gain the ability to represent vertical structures by way of a non-functional relation between range and position of a reflectance datum. The contribution of this paper consists of the exploration of the potential benefits of combining this parameterisation with appearance-based information from scene images. In particular, we develop an approach which employs image segmentation to select the support set of each one of an ensemble of GPs to allow the environment representation to scale to arbitrary point cloud sizes.

This approach naturally extends to the exploitation of semantic workspace classes to limit the extent and size of the support sets.

III. NON-FUNCTIONAL SURFACE REPRESENTATIONS VIA BEAM-SPACE PARAMETERISATION

Our aim is to form an implicit representation of workspace surfaces by processing data gathered from a laser sensor being moved through space at an arbitrary orientation. In particular, we are interested in answering range queries along arbitrary rays emanating from the sensor at any point along the sensor’s trajectory. To this end we adopt the formulation introduced in [1] which is summarised in the following. Let a beam from the laser be parameterised as a point $\mathbf{q} \in \mathbb{S} \times \mathbb{R}$ where $\mathbf{q} = [\theta, t]^T$. θ denotes the angular position of the laser beam and t denotes time. \mathbf{q} is a point in *beam space* \mathcal{Q} as shown in Figure 2 and refer to it as a particular configuration of the laser sensor.

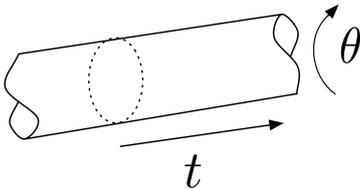


Figure 2: An illustration of beam space \mathcal{Q} . Every point \mathbf{q} on this manifold is parametrized by time t and the angular position of the laser beam θ .

Beam space provides a more natural domain for range regression since it is more closely related to the actual parameters of the laser sensor itself than the more commonly used 3D Euclidean space. Due to a laser sensor being unable to penetrate surfaces, the mapping from beam space to range is necessarily a well formed function — only a single range can be associated with any point in \mathcal{Q} . These considerations therefore lead to a functional mapping from a given time and angle to a scalar range r such that

$$G : \mathbb{S} \times \mathbb{R} \rightarrow \mathbb{R}, \mathbf{q} \mapsto r. \quad (1)$$

Consider now a mapping from beam space to the 3D Euclidean workspace \mathcal{W}

$$E'(\mathbf{q}) = E(V(\mathbf{q}), G(\mathbf{q})) \quad (2)$$

where:

$$V : \mathbb{S} \times \mathbb{R} \rightarrow \mathbb{S}^3 \times \mathbb{R}^3, \mathbf{q} \mapsto \mathbf{p} \quad (3)$$

$$E : (\mathbb{S}^3 \times \mathbb{R}^3) \times \mathbb{R} \rightarrow \mathbb{R}^3, (\mathbf{p}, r) \mapsto \mathbf{x} \quad (4)$$

Here, for every point \mathbf{q} in beam space $V(\mathbf{q})$ provides the 6DOF pose of the sensor $\mathbf{p} \in \mathbb{S}^3 \times \mathbb{R}^3$ — represented as roll, pitch, yaw and position — and $E(\mathbf{p}, r)$ associates a sensor pose \mathbf{p} and a scalar range

measurement r to a point \mathbf{x} in Euclidean space \mathcal{W} . Crucially, Equation 2 permits the functional mapping in Equation 1 to model the non functional relation between elevation z and (x, y) location (as indicated in Figure 1). As a consequence the beam-space parameterisation enables the modelling of vertical structures such as walls in Euclidean space.

In addition to this one-to-many correspondence in Euclidean space, we require that the mapping $G(\mathbf{q})$ provides confidence bounds on the range r , which can also be transformed through Equations 3 and 4 into 3D Euclidean space. Gaussian Process regression is chosen to provide $G(\mathbf{q})$ such that for a given set of measurements $\mathcal{D} = \{(\mathbf{q}_i, r_i)\}_{i=1}^N$ and any query point in beam space \mathbf{q}^* we obtain the predictive distribution $p(r^*|\mathbf{q}^*, \mathcal{D})$. In the following we present a brief summary of how this distribution is obtained.

A. Gaussian Process Regression

A Gaussian Process (GP) provides for non-parametric probabilistic regression. They assume that all measurements of the underlying latent functions are jointly normal distributed. In our particular case this latent function is the mapping $G(\mathbf{q})$. Thus, given known training ranges \mathbf{r} from different configurations \mathbf{q} and a query point \mathbf{q}^* with corresponding target range r^* we can write

$$\begin{bmatrix} \mathbf{r} \\ r^* \end{bmatrix} \sim \mathcal{N}(\mathbf{m}(\mathbf{q}), \mathbf{K}(\mathbf{q}, \mathbf{q}^*)), \quad (5)$$

where \mathbf{K} is a matrix of functions dependent, in the stationary case, only on a suitably chosen distance metric $d = \|\mathbf{q} - \mathbf{q}'\|$ between two points \mathbf{q} and \mathbf{q}' in beam space, and $\mathbf{m}(\mathbf{q})$ is the mean function. The derivation of the predictive distribution $p(r^*|\mathbf{q}^*, \mathcal{D})$ for mean and covariances are standard and can be found, for example, in [2]. Importantly, this formulation is able to explicitly account for noise in the training observations \mathbf{r} in terms of an additive white noise process of strength σ_n .

The workspaces in which we operate have heterogeneous geometric properties that are sampled at varying densities by the laser scanner. In addition, the quantities of data we consider render the application of a single monolithic GP infeasible — time complexity of a naïve implementation of GP regression is cubic in the number of training cases. In this work we therefore adopt a local approximation technique whereby a bank of GPs with stationary covariance functions and local support is used to provide a piecewise approximation to a single global non-stationary covariance function [15]. Our selection criteria for the local *support sets* is outlined in Section III-B.

Throughout this work we use a member of the Matérn class of covariance functions as advocated in [16]. This class of covariance functions provides

a shape parameter, ν , which regulates the smoothness of the interpolation and subsumes the more standard squared exponential covariance function as a special case when $\nu = \frac{1}{2}$. As suggested in [2] we explored several common choices of $\nu = \{\frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \infty\}$ over a number of workspaces varying in complexity. We found that $\nu = \frac{3}{2}$ consistently produced accurate surface reconstructions for a given support set size and density. Although the smoother covariance functions $\nu = \{\frac{5}{2}, \infty\}$ performed well for simple workspaces, they seemed over constrained in complex scenarios, and vice versa for the rough covariance function $\nu = \frac{1}{2}$. Thus

$$K(\mathbf{q}, \mathbf{q}') = \sigma_p^2 \left(1 + \frac{d\sqrt{3}}{l}\right) \exp\left(-\frac{d\sqrt{3}}{l}\right), \quad (6)$$

where l is the length scale, σ_p is the process noise, and $d = \|\mathbf{q} - \mathbf{q}'\|_{\mathcal{Q}}$ denotes the geodesic distance between \mathbf{q} and \mathbf{q}' in beam space. In order to estimate the hyperparameters $\phi = l, \sigma_n, \sigma_p$ we adopt the standard approach of maximising the logarithm of the marginal data likelihood [2].

B. Generating Support Sets

The unfavourable scaling of full GP inference as cubic in the number of training data has given rise to a plethora of approximation methods [17]. A common approach is based on a mixture-of-GP-experts [18, 19], which partitions the data domain into independent subsets of local *support sets*. These support sets are by design significantly smaller than the original training domain so that inference time is decreased. An additional advantage is provided since local regressors can learn local properties and so provide a more faithful representation of the overall data. The split of the input data can be performed according to various criteria. For example, [12] use a regular tiling of models while [20] perform clustering in the input space.

In accordance with the mixture-of-GP-experts paradigm we also split our range data into several support sets. In addition to a regular tiling in \mathbf{Q} -space we compare against a segmentation scheme that leverages another sensor modality: vision, which densely samples the workspace. Our reason for employing an additional vision sensor is two-fold. Firstly, it paves the way towards semantics-guided sampling of workspace regions. Workspace classes of interest can be explicitly included in or excluded from the surface representation, and can be provided by the user, or automatically by systems such as [21]. Secondly, we wish to exploit the common intuition that discontinuities in range are indicated by edges in an image. By using the dense sampling found in images we can more accurately delineate surfaces than if we applied segmentation directly on the laser data.



Figure 3: A typical scene from the New College data set and the corresponding image segmentation using [24] that gave the lowest RMS error in our framework.

This allows our segments to have more homogeneous surface properties which aids surface reconstruction, and will greatly aid the performance of sparsification techniques (see [2]) in future work.

IV. RESULTS

In this section we describe experimental results obtained using the New College data set [22]. Data were gathered using a Segway RMP 200 platform along a 2.2 km track spanning campus and parkland environments. The laser sensors used were LMS291-S14 mounted in vertical alignment on the sides of the vehicle. Image data were gathered using a PointGray Ladybug. Pose estimates were obtained using visual odometry from a forward facing stereoscopic camera [23].

A. Support set selection: a comparative analysis

This section establishes that support set selection via a segmentation in appearance space yields comparable reconstruction accuracy to the more conventional selection with regular tiling. To this end an arbitrary image from the New College data set was chosen (Figure 3a) together with the corresponding laser data from a window of time around the image timestamp. Calibrations between laser and vision cameras were made by hand, and refined based on resulting laser projections into images. We use a graph based [24] image segmentation algorithm due to the low number of required parameters, although other segmentations could equally be provided (Figure 3b). For a fair comparison, we performed a grid search to find parameters that produced the lowest root-mean-square (RMS) reconstruction error against a hold-out set for both segmentation strategies.

Given each segmentation, we used standard log-likelihood optimization (Section III-A) to find the best

Error Metric	Regular Tiled	Appearance Segmentation
RMS	0.25	0.27
Median abs. error	6.77×10^{-3}	8.75×10^{-3}
L_∞	5.58	4.65

Table I: Comparison of reconstruction error for both regular tiling and appearance based segmentation for the dataset shown in Fig. 3a.

hyperparameters for each segment using the average of all \mathbf{r} in the segment as $\mathbf{m}(\mathbf{q})$ (Equation 5). Resulting error statistics were recorded and compiled into Table I. In addition to RMS error the median absolute error and the L_∞ -norm are stated to provide a comparison with a more robust measure with respect to outliers as well as a measure of worst case performance. We note that the performance of the appearance based reconstruction method is slightly worse compared to the regular tiling when measured in RMS and median absolute error. The appearance based method has a slightly lower L_∞ -norm. This is an intuitive result since by design the image segmentation selects support sets that form visually and by correspondence, geometrically homogeneous regions. These are generally more amenable to stationary GP regression and may explain the lower worst case error. We note that although these results help highlight the robustness of our framework, errors may further be improved with a more precise calibration between laser and imagery.

Illustrations of both the regular tiled and the appearance-based approaches are provided in Figures 4 and 5 for the optimal (as found above) grid cell area and segmentation parameters, respectively. Depicted in both cases is the \mathcal{Q} -space representation of the segmentations as well as the projection of the corresponding laser data into 3D Euclidean space. The image segmentation giving rise to this transformation is shown in Figure 3b. We notice that in terms of geometrically uniform regions our framework naturally has a preference for over segmentation to ensure that each segment is as uniform as possible, as mirrored in the lowest overall RMS error.

To obtain a statistically more meaningful comparison between the two approaches and to validate that over-segmentation in the images is not a concern, the optimal segmentation parameters and grid cell area were fixed and the reconstruction errors were obtained for 100 randomly selected images in the dataset. The results are presented in Table II. Inspection of means and corresponding standard deviations suggest that we are not incurring any significant penalty in terms of error in using one segmentation over the other.

B. Semantics-guided Resampling

We have shown that we have not adversely affected the surface reconstruction performance by introducing

image based segmentation into our framework. We now present an initial application where semantic information from another sub-system is used to cue our actions.

We consider an oracle as given, such as [21] or hand-labelled segments of an image belonging to either one of the classes 'bush', 'wall', 'floor' or 'other'. Members of the class 'other' are removed from the data. Consider the scene depicted in Figure 6a. It was segmented using the optimal parameter set determined above. The mixture of GP regressors was trained accordingly. Having obtained a continuous workspace representation we are now able to use our GP framework to sample laser range data at arbitrary locations. Consider an application where the goal is to build an efficient mesh representation of the environment. Intuitively, a low sampling density of polygon vertices suffices in homogeneous regions - 'wall' and 'floor' - while geometrically heterogeneous regions - 'bush' - will benefit from a more dense sampling. As illustrated in Figure 6b, this is readily achievable with the method proposed here.

V. CONCLUSION

This paper introduces a non-parametric piecewise continuous representation of large 3D laser point clouds. Our approach is based on a novel parameterisation of a well defined function which allows a non-functional relation between positions on the ground plane and elevation values — thus enabling our approach to represent vertical structures, whilst employing a powerful Gaussian Processes regression framework. We use a mixture of GPs to ensure computational tractability. Although other segmentations are equally valid, in this paper we determine the local support sets for each GP through both a naive gridding in input space and a segmentation in appearance space.

The proposed method is applied to a publicly available outdoor, urban data set. An empirical investigation suggests that reconstruction performance as measured by RMS error, median of absolute error and L_∞ -norm are comparable for both segmentations, even under crude hand calibration of laser and vision modalities. A full statistically significant analysis is part of our future work.

We then illustrate how our approach also has the ability to be integrated into a larger system. In our example semantically significant objects in the workspace were resampled at a higher density than background objects to allow efficient meshing algorithms, but one can imagine many other uses, such as model selection, data filtering, CAD model generation and point cloud compression algorithms. Future work will build on this framework where the application of GP sparsification techniques could be particularly attractive due to their ability to select

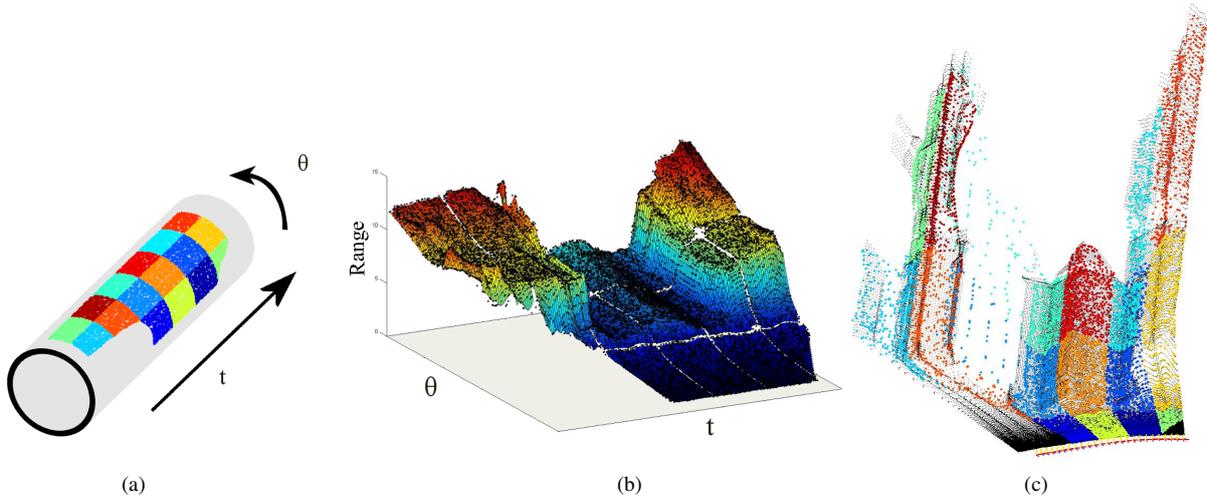


Figure 4: Support sets selected by regular tiling in \mathcal{Q} . (a) Beam space \mathcal{Q} . (b) Corresponding \mathcal{Q} (unwrapped) with t and θ of \mathcal{Q} space along the horizontal axes and laser range measurements up the vertical. (c) Laser range image of (b) projected into 3D Euclidean space via the robot trajectory (indicated by the coloured axes across the bottom of the image). Both (a) and (c) are coloured according to the support regions the predictions are formed from. For a fair comparison with the image based segmentation, only laser points (depicted in colour) that projected into the image (Figure 3a) were used. It is recommended to view this figure in colour.

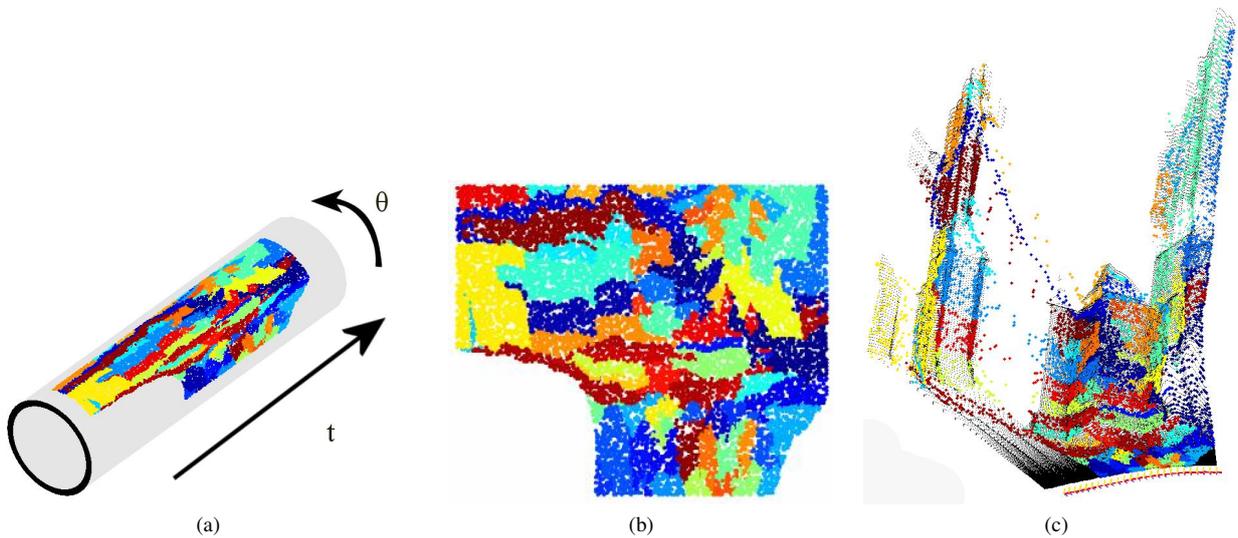


Figure 5: Support sets selected by appearance-based image segmentation in \mathcal{Q} . (a) Beam space \mathcal{Q} . (b) Corresponding \mathcal{Q} (unwrapped) with t horizontal, θ vertical. White sections in the image are as a result from segmenting in the image space and projecting into \mathcal{Q} . (c) 3D Euclidean plot of the laser data projected through the robot trajectory, with segmentation colouring. Only the laser data (depicted in colour) that projected into the image (Figure 3a) were used in the analysis. It is recommended to view this figure in colour.

Error Metric	Regular Tiled				Appearance Segmentation			
	Mean	Min	Max	Std. Dev.	Mean	Min	Max	Std. Dev.
RMS	0.36	64.46×10^{-3}	0.90	0.17	0.39	41.67×10^{-3}	0.89	0.19
Median abs. error	12.94×10^{-3}	4.59×10^{-3}	0.51×10^{-1}	8.1×10^{-3}	11.57×10^{-3}	4.34×10^{-3}	48.08×10^{-3}	6.89×10^{-3}
L_∞	6.55	1.27	21.46	3.84	7.47	1.71	21.09	4.23

Table II: Comparison of reconstruction error for both regular tiling and appearance based segmentation for a 100 images from the New College data set [22].

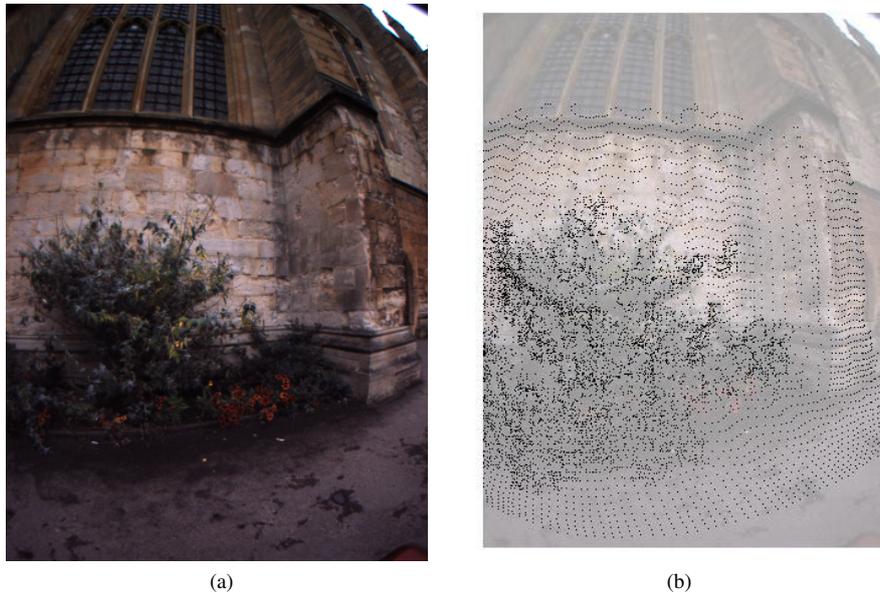


Figure 6: Illustration of semantics-guided resampling. a) Input image to our processing pipeline. Image segmentation was performed using [24] with the parameters values found in Section IV-A. Resulting segments were then hand labelled 'bush', 'wall', 'floor' or 'other'. b) Using our GP framework we could then sample geometrically homogeneous regions such as 'wall' and 'floor' sparsely, while geometrically complex regions such as 'bush' are sampled densely.

only the most pertinent data points to represent the underlying workspace surfaces. This could provide substantial data compression which is increasingly required for the vast amounts of data present in today's point clouds. Results so far, have provided enticing evidence to suggest that our image based segmentation shares the same preference for homogeneous surfaces (and therefore easily compressible segments) as that of sparsification techniques. Further work will also be to increase the accuracy of our system by optimising the laser-vision calibrations, and incorporating more complex mean and (non-stationary) covariance functions.

VI. ACKNOWLEDGMENTS

This work was funded through an Engineering and Physical Sciences Research Council CASE award with BAE Systems, and by the European Commission under grant agreement number FP7-231888-EUROPA.

REFERENCES

- [1] 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*, Zaragoza, Spain, 2010.
- [2] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [3] C. Früh and A. Zakhor, "Constructing 3d city models by merging aerial and ground views," *IEEE Comput. Graph. Appl.*, vol. 23, no. 6, pp. 52–61, 2003.
- [4] D. Hähnel, W. Burgard, and S. Thrun, "Learning compact 3D models of indoor and outdoor environment with a mobile robot," *Robotics and Autonomous Systems*, vol. 44, no. 1, pp. 15–27, 2003.
- [5] R. Triebel, W. Burgard, and F. Dellaert, "Using hierarchical em to extract planes from 3d range

- scans,” in *Int. Conf. Robotics and Automation (ICRA)*, 2005.
- [6] M. Pauly, N. Mitra, and L. Guibas, “Uncertainty and variability in point cloud surface data,” in *Eurographics symposium on point-based graphics*, vol. 84, 2004.
- [7] M. C. Martin and H. Moravec, “Robot evidence grids,” Robotics Institute, Tech. Rep. CMU-RI-TR-96-06, 1996.
- [8] J. Bares, M. Hebert, T. Kanade, E. Krotkov, T. Mitchell, R. Simmons, and W. R. L. Whittaker, “Ambler: An autonomous rover for planetary exploration,” *IEEE Computer*, vol. 22, no. 1, pp. 18–26, 1989.
- [9] P. Pfaff and W. Burgard, “An efficient extension of elevation maps for outdoor terrain mapping,” in *Intl. Conf. on Field and Service Robotics (FSR)*, 2005.
- [10] R. Triebel, P. Pfaff, and W. Burgard, “Multi-level surface maps for outdoor terrain mapping and loop closing,” in *Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2006.
- [11] T. Lang, C. Plagemann, and W. Burgard, “Adaptive non-stationary kernel regression for terrain modeling,” in *Robotics: Science and Systems (RSS)*, 2007.
- [12] C. Plagemann, S. Mischke, S. Prentice, K. Kersting, N. Roy, and W. Burgard, “Learning predictive terrain models for legged robot locomotion,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2008.
- [13] S. Vasudevan, F. Ramos, E. Nettleton, H. Durrant-Whyte, and A. Blair, “Gaussian process modeling of large scale terrain,” in *Int. Conf. Robotics and Automation (ICRA)*, 2009.
- [14] 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 (RSS)*, 2007.
- [15] E. Snelson and Z. Ghahramani, “Local and global sparse gaussian process approximations,” in *Artificial Intelligence and Statistics*, vol. 11, 2007.
- [16] M. L. Stein, *Interpolating Spatial Data*. Springer-Verlag, 1999.
- [17] J. Quinero-Candela, C. E. Rasmussen, and C. K. I. Williams, *Approximation Methods for Gaussian Process Regression*. MIT Press, 2007, ch. 9.
- [18] C. E. Rasmussen and Z. Ghahramani, “Infinite mixtures of gaussian process experts,” in *Advances in Neural Information Processing Systems*, vol. 14. MIT Press, 2002, p. 881:888.
- [19] V. Tresp, “Mixtures of gaussian processes,” in *Advances in Neural Information Processing Systems (NIPS)*, 2000, p. 654:660.
- [20] A. Schwaighofer and V. Tresp, “Transductive and inductive methods for approximate gaussian process regression,” in *Advances in Neural Information Processing Systems (NIPS)*. MIT Press, 2002.
- [21] I. Posner, M. Cummins, and P. Newman, “A generative framework for fast urban labeling using spatial and temporal context,” *Autonomous Robots*, vol. 26, pp. 153–170, 2009.
- [22] M. Smith, I. Baldwin, W. Churchill, R. Paul, and P. Newman, “The new college vision and laser data set,” *International Journal for Robotics Research (IJRR)*, vol. 28, no. 5, pp. 595–599, May 2009. [Online]. Available: <http://www.robots.ox.ac.uk/NewCollegeData/>
- [23] C. Mei, G. Sibley, M. Cummins, P. Newman, and I. Reid, “A constant time efficient stereo slam system,” in *BMVC*, 2009.
- [24] P. F. Felzenszwalb and D. P. Huttenlocher, “Efficient graph-based image segmentation,” *Int. J. Comput. Vision*, vol. 59, no. 2, pp. 167–181, 2004.