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

Inferring Motion Uncertainty from Shape-Matching

Zuolei Sun, Joop van de Ven, Fabio Ramos, Xuchu Mao and Hugh Durrant-Whyte

Abstract— **This paper proposes a novel method for computing robot motion uncertainty from ranging sensor data. The method utilises the recently proposed CRF-Matching procedure which matches laser scans based on shape descriptors. Motion estimates are computed in a probabilistic framework by performing inference on a probabilistic graphical model. We propose an efficient sampling procedure for obtaining probable association hypothesis of the probabilistic graphical model. The hypothesis are used to generate estimates on the uncertainty of translational and rotational movements of the robot. Experiments demonstrate the benefits of the approach on simulated data sets and on laser scans from an urban environment. The approach is also combined with the well-established delayed-state information filter for a large-scale outdoor simultaneous localisation and mapping task.**

I. INTRODUCTION

Reliable navigation in mobile robotics requires the computation of robust motion estimates. While solutions based on inertial measurement units or GPS can provide an estimate and corresponding uncertainties directly, for ranging sensors this task is significantly more complex. The difficulty lies in obtaining robust point correspondences between consecutive scans from which the motion estimation is computed. Solutions based on the Iterative Closest Point (ICP) [1], [2] can, in general, provide reasonable motion estimates. However, for reliable navigation, deterministic estimates of the motion are not enough; it is also necessary to quantify the uncertainty on these estimates.

The motion estimation uncertainty computed from range sensors arrives from two main sources: 1) uncertainty in the point associations; 2) uncertainty in the range and bearing measurements from the sensor. For the particular case of laser range finders, range and bearing estimates are very accurate and almost insignificant compared to the uncertainty from the point association. The computation of uncertainty from the point association is however, much more challenging as it involves the evaluation of a potentially enormous set of possible associations. For example, in a conventional data association problem with a pair of laser scans with 361 laser points each, the number of possible associations is 361^{361}. The probability of a particular scan association is then computed by evaluating the likelihood of the association divided by the sum of the likelihoods of all possible associations. This computation quickly becomes infeasible even for laser scans with a relatively small number of points.

The recently proposed CRF-Matching algorithm [3] tackles the problem of laser scan association by formulating it as a probabilistic inference procedure in a graphical model. This allows the incorporation of general features that take into account shape descriptors to match the scans. With such a formulation, the laser point association problem can be addressed in a integrated fashion, jointly reasoning over the space of associations. This paper extends the former by providing a procedure to compute the motion uncertainty from the scan association. In particular, the main contribution presented here is the development of a novel inference algorithm for probabilistic networks that efficiently seeks for probable configurations of the network in the space of laser point associations.

II. RELATED WORK

Determining the uncertainty of scan matching is a significant problem in mobile robotics using laser range finders. The problem has received much attention since Lu [1] first addressed it. Most solutions focus on ICP like algorithms and can be categorised according to the uncertainty sources they model; measurement noise, lack of pairwise constrains, local minima and improper laser point associations. Related work for each of the four categories will be discussed next.

1) Measurement Noise: Intuitively, measurement noise from imperfect sensors adds uncertainty. The general approach to dealing with isotropic or non-isotropic sensor noise is by means of the *likelihood* function. In [4, section VI.A], the uncertainty is modelled as a multi-Gaussian likelihood distribution. The individual Gaussians are parameterised by the covariance of range measurements (after linearisation). Analogously, [5] attempts to capture the relationship between measurement noise and the uncertainty estimate from which a weighted matching algorithm is proposed. Wang [6, Section 3.3] again employs the likelihood function but for a grid-based approach. Other approaches, such as [7], [8], [1, Section 6.2], take advantage of *Laplace*'s method to approximate the sensor noise distribution by a Gaussian.

2) Lack of Pairwise Constrains: Non-Gaussian likelihood functions are often approximated by their first two moments; the mean and variance. This is feasible, and reasonable, only when sufficient associated laser point pairs are obtained at consecutive poses. Then the true uncertainty is highly peaked and the approximation from non-Gaussian to Gaussian is accurate. However, it is possible that there are very few correspondences due to a small overlap between scans, or

Zuolei Sun and Xuchu Mao are with the Navigation & Control Lab, Department of Instrumentation Engineering, School of Electronic, Information and Electrical Engineering, Shanghai Jiao Tong University, Shanghai, 200240, China sunzuolei@gmail.com, maoxc@sjtu.edu.cn

Joop van de Ven, Fabio Ramos and Hugh Durrant-Whyte are with the ARC Centre of Excellence for Autonomous Systems, Australian Centre for Field Robotics, University of Sydney 2006, NSW, Australia {j.vandeven, f.ramos, hugh}@acfr.usyd.edu.au

a lack of orthogonal constraints. In such cases the likelihood is ill-represented by a Gaussian [4]. Furthermore, in case when only 1 or 2 associations are found, minimising the error function of ICP (or its variations) is a high risk venture. This is shown in [8] for typical under-constrained situations such as the corridor and circular environments.

3) Local Minima: ICP optimises a concave error function for any given association. The ICP algorithm generally uses nearest neighbours to determine the association. However, depending on the association, the resulting minima may only be a local minima of the overall association problem. As shown in [4, Section VI.C], the only way to avoid such local minima is to adequately search the relevant space of possible associations. One way is to draw a set of samples from the prior distribution of the relative transformation. Each sample is used to initialise a separate run of ICP, each converging to a (perhaps) different local minima. The covariance estimate is then computed from the combination of these solutions. One such application is explored in [6].

4) Improper Laser Point Associations: It is generally accepted that for high accuracy sensors, such as scanning lasers, improper correspondences are much more critical compared to the above three uncertainty sources. Unfortunately, most solutions are based on the assumption that a perfect correspondence is available. Correspondence uncertainty modelling is seldom treated thoroughly. The Laplace method does not account for correspondence uncertainty [4]. The weighted method relies too much on parameters [5]. While the sampling-based "off-line" approach proposed in [7] is somewhat *ad hoc*.

CRF-Matching [3] defines the scan matching problem in a sound statistical framework; as an inference procedure in a graphical model. This allows for a natural extension to the expression of the uncertainty on correspondences. Investigating this issue is the main contribution of this paper.

III. Conditional Random Fields

A Conditional Random Field (CRF) [9] is a *probabilistic discriminative framework*. Typically the probability distribution of a CRF is represented using an undirected graphical model; the vertices of the graph index the distribution's random variables, while the edges of the graph capture relationships between variables. In the case of a CRF the hidden variables $\mathbf{x} = \langle x_1, x_2, ..., x_N \rangle$ are *globally* conditioned on the observations $\mathbf{z}$; $p(\mathbf{x} \mid \mathbf{z})$. The distribution must factorise according to the cliques of the graph; cliques are fully connected sub-graphs. Let $\mathscr{C}$ be the set of all cliques of the graph. The distribution must then factorise as a product of clique potentials $\phi_c(x_c, \mathbf{z})$ as follows:

$$p(\mathbf{x} \mid \mathbf{z}) = \frac{1}{Z(\mathbf{z})} \prod_{c \in \mathscr{C}} \phi_c(x_c, \mathbf{z}), \tag{1}$$

where x_c are the hidden variables of the clique and $Z(\mathbf{z}) = \sum_x \prod_{c \in \mathscr{C}} \phi_c(x_c, \mathbf{z})$ is the partition function; it normalises the product of clique potentials.

Clique potentials are commonly expressed by a log-linear combination of feature functions, $\phi_c(x_c, \mathbf{z}) = \exp\left(\mathbf{w}_c^T \cdot \mathbf{f}_c(x_c, \mathbf{z})\right)$, resulting in the following definition of a CRF:

$$p(\mathbf{x} \mid \mathbf{z}) = \frac{1}{Z(\mathbf{z})} \exp\left(\sum_{c \in \mathscr{C}} \mathbf{w}_c^T \cdot \mathbf{f}_c(x_c, \mathbf{z}) \right); \tag{2}$$

where $Z(\mathbf{z}) = \sum_x \exp(\sum_{c \in \mathscr{C}} \mathbf{w}_c^T \cdot \mathbf{f}_c(x_c, \mathbf{z}))$ is the partition function expressed using the log-linear form. $\mathscr{C}$ is again the set of all cliques in the graph. $\mathbf{w}_c$ are the parameters (or weights). The feature functions $\mathbf{f}_c$ extract feature vectors given the value of the clique variables x_c and observations $\mathbf{z}$.

A. Parameter Learning

The weights $\mathbf{w}_c$ express the relative importance of the feature functions $\mathbf{f}_c$. As such they play an important role in determining the shape of the distribution. Learning the values of the weights is achieved through maximisation of the conditional likelihood (Equation 2) given labelled training data. In our case this is computationally intractable; the partition function $Z(\mathbf{z})$ sums over the (very large) space of all hidden variables. We therefore employ maximum pseudo-likelihood learning [10] to determine the weights.

B. Inference

A CRF (Equation 2) defines a distribution over *all* hidden variables. Typically one is more interested in the configuration of the hidden variables that generates the maximum *a-posteriori* (MAP) value of the distribution. Inference is the process by which this configuration may be obtained.

Belief Propagation (BP) [11] is a class of inference algorithms in which each node sends messages to each of its neighbours in the graph. The messages convey what a node *believes* its neighbours' state should be given its own state. The received messages together with a node's own belief are then used to compute the MAP configuration; the Max-Product BP algorithm. Message construction is defined as follows:

$$m_{ij}(x_i) = \max_{x_j} \left(\phi(x_j) \phi(x_i, x_j) \prod_{k \in \mathscr{N}(j) \setminus i} m_{kj}(x_j) \right). \tag{3}$$

Here $m_{ij}(x_i)$ is the message node j sends to node i. $\phi(x_j)$ is the local clique potential for node j. Local clique potentials represent what a node believes its state is given the observations. Pairwise potentials $\phi(x_i, x_j)$ relate nodes i and j on either end of an edge. They transform the belief of node j into a belief suitable for node i and vice versa. Lastly a product of all incoming messages, *except* from the node to which the message is being sent.

Messages are first passed *inbound* from the leaves to the root of the graph. With each message m_{ij} sent, the node also records the states of x_j responsible for the message: δ_{ij}. Once at the root, the MAP state of the root is computed as;

$$x_r^* \in \arg\max_{x_r} \left(\phi(x_r) \prod_{j \in \mathscr{N}(r)} m_{jr}(x_r) \right). \tag{4}$$

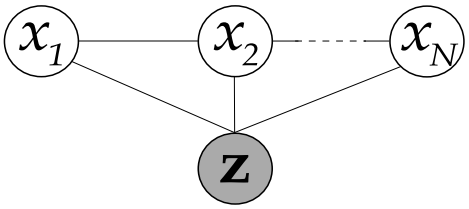

Fig. 1. Graphical model for scan matching.

Here x_r indicates the root node while x_r^* is its maximal configuration. As can be seen maximisation is performed on the root's own belief (local potential) combined with all incoming messages. The maximal configuration for the root node is then simply a state for which the combined belief is maximal. Using the maximal configuration of the root, the algorithm backtracks along the graph, using δ_{ij}, to find a consistent MAP configuration.

IV. MODEL DEFINITION

Data association is modelled using a conditional distribution; $p(\mathbf{x} \mid \mathbf{z})$. The hidden variables $\mathbf{x}$ represent the associations given an observation $\mathbf{z}$. The observation consists of two laser scans for which associations are to be determined.

The vector valued $\mathbf{x}$ consist of N multi-nomial random variables, where N is the number of points in the first scan. Each of these nodes is discrete with $M+1$ states; M being the number of points in the second scan. The states of a variable, say variable x_i, have the following interpretation. The first state indicates the likelihood that point i in the first scan associates to point 1 in the second scan. The second state is the likelihood of association to the second point in the second scan, etc. Finally, the $M+1$ state represents the likelihood that point i is an outlier.

The graphical model corresponding to this distribution is given in Figure 1. As can be seen, the hidden variables are modelled as a chain. The motivation for using a chain stems from the way scan data is obtained. The laser scanner obtains data in a single plane; the plane in which the beam scans. As such, points are acquired one after the other; a chain represents this acquisition. Correlations in the data are not lost in the graph as there is a path from any one node in a chain to any other node.

A. Feature Description

The popularity of CRFs is in no small measure due to the feature functions $\mathbf{f}_c$ in Equation 2. These encapsulate domain specific knowledge which can subsequently be used in the probabilistic framework of the CRF. The experiments employ a subset of the features used in [3]. Three local shape features *distance*, *angle* and *geodesic* and one pairwise feature *association* are used. The reader is referred to [3] for a more detailed description of these features.

In addition, one new local feature is used, the *ICP association* feature. ICP is able to compute good association results provided the scans are reasonably aligned. It is therefore beneficial to incorporate ICP estimates in the inference procedure. However, a mapping is required from the single valued ICP results to the multi-valued representation of the CRF. This is achieved by assigning 1 to the k-th state of a node, say node i, and zero to all its other states; where k is determined by the ICP association for point i.

Algorithm 1 Sample-Product message construction.

1: **for** $s = 1$ to $M+1$ **do**
2: $\quad r \leftarrow GenerateRandomNumber([0, x_{ij}^{c_{M+1}}(s)])$
3: $\quad l \leftarrow \arg\min_k x_{ij}^{c_k}(s) > r$
4: $\quad m_{ij}(s) \leftarrow x_{ij}(s,l)$
5: $\quad \delta_{ij}(s) \leftarrow l$
6: **end for**

V. SAMPLING INFERENCE

In order to construct uncertainty estimates over the robot pose, several association hypotheses are required. For our estimate construction we therefore wish to *efficiently* sample configurations that map to likely poses. From these samples the pose uncertainty can then be computed.

For the remainder of this section we shall assume that the mapping from association to pose is smooth; in the experiments a least squares mapping is used. With this assumption the problem is then one of sampling configurations that are *similar* to the MAP configuration. This can be achieved by replacing the max operator with a sampling operator as discussed next.

The $sample_{x_j}$ operator in Equation 5 samples on the product of incoming messages, pairwise potential and local potential; analogous to Max-Product. The values in the generated message m_{ij} are the values of the sampled states.

$$m_{ij}(x_i) \approx sample_{x_j}\left(\phi(x_j)\phi(x_i,x_j)\prod_{k\in\mathcal{N}(j)\setminus i} m_{kj}(x_j)\right) \quad (5)$$

In addition to generating the message, the sampled states are also recorded (analogous to Max-Product); this allows back-tracking to find a consistent configuration once message are propagated inbound to the root.

The sample operator is implemented as shown in Algorithm 1. Let $x_{ij} = \phi(x_j)\phi(x_i,x_j)\prod_{k\in\mathcal{N}(j)\setminus i} m_{kj}(x_j)$ be the product of local and pairwise potentials and incoming messages; a $M+1 \times M+1$ matrix. Let $x_{ij}(s)$ be the s-th row from this matrix, while $x_{ij}^{c_k}(s) = \sum_{l=1}^{k} x_{ij}^{l}(s)$ is the cumulative sum up to and including the k-th value of the s-th row.

The algorithm then proceeds, for each state of a node, as follows: On line 2 a random number is generated in the range 0 to $x_{ij}^{c_{M+1}}(s)$. The random number is then used to find the corresponding state l (line 3). The l-th value of $x_{ij}(s)$ is used in the construction of the message while l is recorded as the state responsible for $m_{ij}(s)$ on lines 4 and 5 respectively.

The commutative property does not hold for the sampling operator. Strictly speaking we are therefore not allowed to interchange max_{x_j} with $sample_{x_j}$ as done in Equation 5. The results in the next section however, show that good results are still achieved. Intuitively this can be explained

by the following observation. In highly peaked distributions (as association generally is) sampling is much like applying the max_{x_j} operator; most of the probability mass is located in only a few states. More importantly, the resulting algorithm is one that can very efficiently sample a configuration from the graph.

The above Sample-Product inference procedure is used to compute K association hypothesis; the experiments compute $K = 30$. For each hypothesis the solution (see also [2, Appendix C]) to translation $T = [x, y]$ and rotation $R = \theta$ is computed according to Equation 6;

$$R, T \leftarrow \min_{R,T} \sum_{i=1}^{N} (z_{1,i} R + T - z_{2,x_i^a})^2, \tag{6}$$

where N denotes the number of laser points in the first scan. The index x_i^a is the index of the point in the second scan, associated to the i-th point in the first scan. As such, association hypothesis are mapped to points in the robot motion space; $X = [\mathbf{x}, \mathbf{y}, \theta]$. The uncertainty is then computed as the covariance of the points in the robot motion space as follows:

$$\Sigma_{CRF} = E\left((X - E(X))(X - E(X))\right). \tag{7}$$

VI. EXPERIMENTS

To verify if the proposed algorithm is able to capture the underlying uncertainty of scan matching, we evaluate its performances on two data sets. The first data set consists of under-constrained simulated data in the form of corridor and circular environments. The second is a data set of the *University of Sydney* campus and it is used in a Simultaneous Localisation and Mapping (SLAM) application.

A. Simulated Data Set

As mentioned in Section II there are two under-constrained environments; the corridor and circular environments. These lead to large uncertainty in specific directions when aligning the pairwise scans. The expected shapes of the uncertainty ellipsis allow us to check the validity of the algorithm [7], [8].

1) Corridor Environment: In the simulated corridor environment the robot is located in a long corridor. The length of the corridor extends beyond the range of the range finder, the robot is thus only able to sense the walls.

Laser scans viewed from two consecutive poses, assuming the walls are sufficiently smooth, will appear identical to the robot. As a result, a lack of distinguishing features for consecutive scans will significantly impair scan matching. In particular along the direction of the walls, where most laser points are aligned. An extreme case is shown in Figure 2(a); here the robot moves in parallel with the walls. The laser range measurements are corrupted by Gaussian white noise with $\mu_r = 0$, $\delta_r = 0.06m$.

The 3-σ ellipsis are shown in Figure 2. The standard deviations of the relative motion estimates are shown in Table I. The results show that the relative motion estimate on X is more uncertain than on Y (see Figure 2(b)). In Figure 2(d) it can be seen that the uncertainty on X causes a large uncertainty on rotation as well.

	$\sigma_x(mm)$	$\sigma_y(mm)$	$\sigma_\theta(^\circ)$
Corridor	102.90	5.11	0.02
Circular	94.30	70.20	2.35

TABLE I
THE UNCERTAINTY ESTIMATES FOR THE SIMULATED (UNDER-CONSTRAINED) ENVIRONMENTS

2) Circular Environment: Analogous to the corridor case, when the robot stands in an ideal circular environment the walls will significantly impair the robot's ability to correctly determine its rotation estimate. A typical circular situation is pictured in Figure 3(a). During inference, it is difficult for the CRF to find optimal laser point associations as there are strong similarities between all points in the data. Compared with the corridor case, the uncertainties for laser point associations are expected to impact primarily on the rotation estimates. This is visualised by the long radius of the uncertainty ellipsis along the rotation axis as shown in Figures 3(c) and 3(d). The standard deviation of the motion estimates for the circular environment are given in Table I.

As shown above, the shape of the uncertainty ellipsis demonstrate that the proposed approach correctly models the underlying uncertainty in the under-constrained corridor and circular cases.

B. Large-Scale SLAM

The configuration of the mobile platform, used to collect the data, is a utility vehicle with SICK laser sensor. The laser range finder data was acquired around the main campus of the *University of Sydney* during the day time; it therefore contains a significant number of dynamic objects, such as cars and people.

To avoid difficulties in feature extraction and maintain tractability, a view-based delayed state SLAM framework is adopted. We do not explicitly model the environment by extracting features and estimating their positions; instead only a selection of key historical poses of the robot's trajectory are maintained in filtering. The raw laser scans are stored and registered in order to calculate pairwise constraints, detect potential loop closures and provide virtual observations of pose displacements. Furthermore, we utilise the extended information filter (EIF) rather than traditional extended Kalman filter (EKF) for practical reasons. The reader is referred to [12] for more details.

CRF-Matching functions, in some sense, as a *virtual odometry*. It can be used to generate estimates, based on consecutive laser scan alignments, for the control input which are then incorporated into the robot motion model. The uncertainty covariance matrix $\Sigma_{CRF}(t)$, computed from the hypothesis (see Section V), is incorporated into the system as the process noise.

GPS data from the mobile platform was collected for ground-truth evaluation. For comparison, ICP-based SLAM

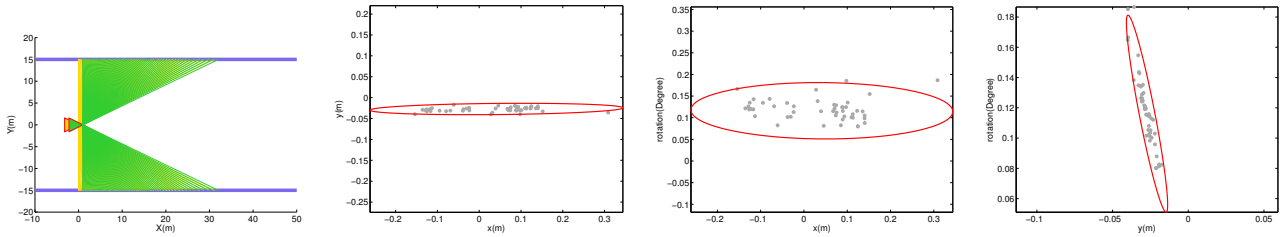

(a) (b) (c) (d)

Fig. 2. Laser scan alignment uncertainty estimates of the corridor environment. (a) Simulation scenario; Triangle indicates the robot. Yellow colour indicates the initial pose and range measurements. Green colour indicates the subsequent pose and measurements. (b) The X-Y 3-σ uncertainty ellipsis and samples computed from CRF-Matching with our approach. (c) X-Rotation 3-σ uncertainty ellipsis and samples. (d) Y-Rotation 3-σ uncertainty ellipsis and samples.

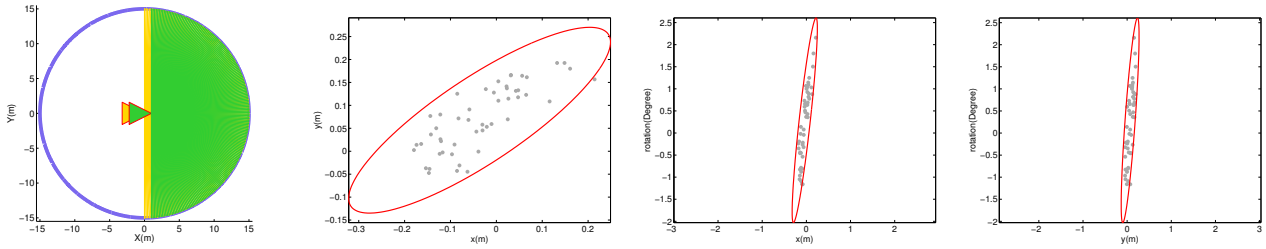

(a) (b) (c) (d)

Fig. 3. Laser scan alignment uncertainty estimate of the circular environment. (a) Simulation scenario; Triangle indicates the robot. Yellow colour indicates the initial pose and range measurements. Green colour indicates the subsequent pose and measurements. (b) The X-Y uncertainty ellipsis and samples computed from CRF-Matching with our approach. (c) X-Rotation uncertainty ellipsis and samples. (d) Y-Rotation uncertainty ellipsis and samples.

with uncertainty estimates from the *off-line* approach [7] are computed. The trajectories and maps are visualised in Figure 4 where the first loop closure corrects both the robot states and map representation.

Figure 4(a) shows the closed-loop paths of CRF-SLAM and ICP-SLAM and the uncertainty ellipses of CRF-SLAM. We can conclude that CRF-SLAM out-performs ICP-SLAM from two observations. First, the path of CRF-SLAM is much closer to the GPS trajectory than that of ICP-SLAM. The largest offset (30.5m) of the latter occurs after the robot has travelled around 420m while the offset of CRF-SLAM is 4.3m. Second, accounting for the inherent measurement error of GPS, we further check the results by superimposing the maps on an aerial photo. It can be seen from Figure 4(b) that the contour created by the laser map coincides with the real object quite well.

The uncertainty model discussed in this paper indicates the quality of the relative motion estimates. In other words, it tell us how much we can trust the estimates. This is crucial within the filtering framework, since proper quantifications of uncertainty can correct the estimates while performing time and measurement updates. Therefore, the CRF-SLAM results presented here show that our uncertainty inference algorithm can be integrated into a real SLAM framework and perform well.

VII. Discussions

Uncertainty estimation is indispensable in the construction of a filtering framework; the proposed Sample-Product inference procedure incorporates CRF laser scan matching into a practical SLAM application. To our knowledge, existing work tackling similar topics, such as [1], [7], [4], [5], [8], are based on iterative closest point (ICP) or its variants. [3] has demonstrated that CRF laser scan matching outperforms ICP-based approaches. Furthermore, in contrast to the ICP based techniques, the focus here is on the uncertainty caused by improper laser point associations rather than measurement noise. The proposed algorithm solves the task in a novel way, capturing laser point associations hypothesis with a shape-oriented graphical model.

The sampling approach is used to track the uncertainty distribution, like the *off-line* approach presented in [8]. The latter attempts to model the very large laser point correspondence space ($\mathcal{O}((M+1)^N)$) by a random brute force association configuration test. It essentially assigns an identical likelihood to the samples in the space, and randomly selects a small set to estimate the uncertainty. However, the correspondence space is too large for the quantity of samples. So the *off-line* approach can't guarantee to correctly capture the underlying distribution. With respect to Sample-Product inference, for a single node, the product of incoming messages, pairwise potential and local potential represents the *likelihood* of association candidates for that node. We perform sampling over this likelihood. It ensures entries with higher likelihood are more likely selected. As a result, the K most probable association estimates are selected generating better motion uncertainty estimates.

Instead of ranging over the entire association space, the proposed algorithm tracks the underlying uncertainty by focusing on the K most likely laser point associations. The

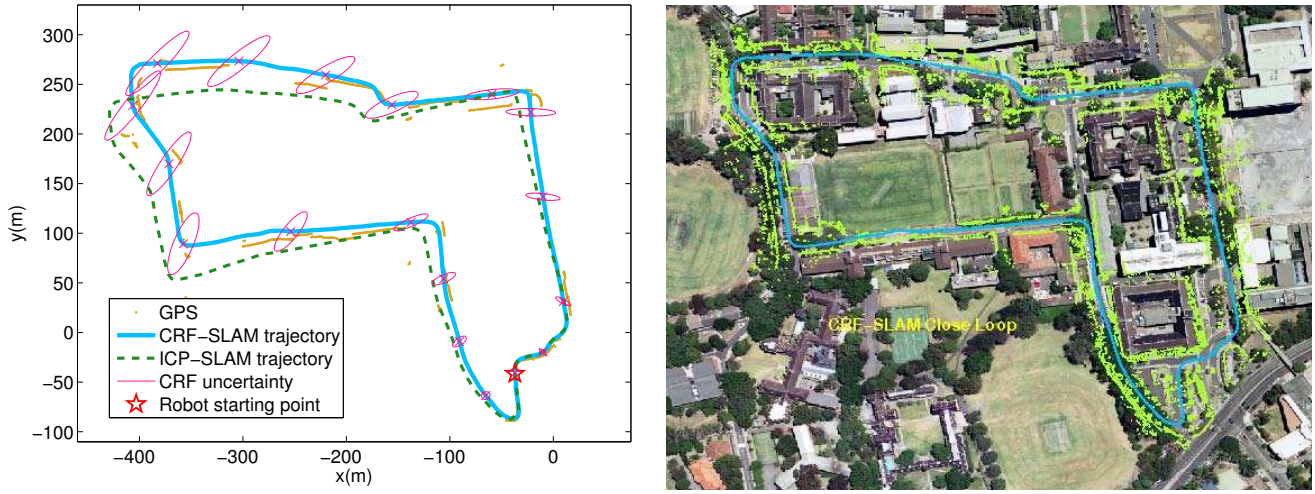

(a) (b)

Fig. 4. The large-scale SLAM experiments. (a) Robot trajectories after loop closure and the 1-σ uncertainty ellipsis of CRF-SLAM propagated by the filter for every 100 poses. (b) The close-loop map created from CRF-SLAM overlapped with the aerial photo.

computational cost is determined by the number of samples (K) drawn. In our current Matlab implementation, drawing $K = 30$ samples takes approximately 2.6 seconds. Scan matching and uncertainty quantification for 4391 laser scan pairs is processed within 15 minutes using Matlab's Parallel Computing toolbox.

VIII. CONCLUSIONS

This paper presents an efficient sampling inference procedure to estimate the motion uncertainty for laser scan registration with CRF-Matching. In contrast to existing techniques, our approach focuses on capturing the uncertainty from laser point associations. A sampling mechanism is employed during message propagation in a graphical model to produce probable association configuration hypothesis for two consecutive laser scans collected by a moving robot. This enables the computation of uncertainty for both translation and rotation in a probabilistic manner. The effectiveness of the proposed algorithm is demonstrated on a simulated data-set and in a large-scale SLAM for a challenging urban environment.

ACKNOWLEDGEMENTS

The authors wish to thank Tim Bailey, Karl Granstrom and Ian Mahon for the numerous helpful discussions.

This work would not have been possible without the assistance of the National Natural Science Foundation of China (No.40674002) and Professor Weifeng Tian for supporting Zuolei Sun's overseas research.

This work has been supported by the Rio Tinto Centre for Mine Automation and the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government.

REFERENCES

[1] F. Lu. *Shape registration using optimization for mobile robot navigation*. PhD thesis, University of Toronto, 1995.

[2] F. Lu and E. Milios. Robot pose estimation in unknown environments by matching 2D range scans. *Journal of Intelligent and Robotic Systems*, 18, 1997.

[3] F. Ramos, D. Fox, and H. Durrant-Whyte. CRF-matching: Conditional random fields for feature-based scan matching. In *Proc. of Robotics: Science and Systems*, 2007.

[4] T. Bailey and J. Nieto. Scan-slam: Recursive mapping and localisation with arbitrary-shaped landmarks. In *Workshop on Quantitative Performance Evaluation of Navigation Solutions for Mobile Robots, Robotics: Science and Systems Conference (RSS)*, Zurich, Switzerland, 2008.

[5] S. T. Pfister, K. L. Kriechbaum, S. I. Roumeliotis, and J. W. Burdick. Weighted range sensor matching algorithms for mobile robot displacement estimation. In *Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on*, volume 2, pages 1667–1674 vol.2, 2002.

[6] C. C. Wang. *Simultaneous Localization, Mapping and Moving Object Tracking*. PhD thesis, Carnegie Mellon University, 2004.

[7] O. Bengtsson and A. J. Baerveldt. Robot localization based on scan-matching–estimating the covariance matrix for the idc algorithm. *Robotics and Autonomous Systems*, 44(1):29–40, 2003.

[8] A. Censi. An accurate closed-form estimate of icp's covariance. In *Robotics and Automation, IEEE International Conference on*, pages 3167–3172, 2007.

[9] J. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proc. of the International Conference on Machine Learning (ICML)*, 2001.

[10] J. Besag. Statistical analysis of non-lattice data. *The Statistician*, 24, 1975.

[11] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., 1988.

[12] R. M. Eustice. *Large-area visually augmented navigation for autonomous underwater vehicles*. PhD thesis, Massachusetts Institute of Technology, 2005.