2007 IEEE International Conference on
Robotics and Automation
Roma, Italy, 10-14 April 2007

Histogram Matching and Global Initialization for Laser-only SLAM in Large Unstructured Environments

Michael Bosse and Jonathan Roberts
Autonomous Systems Laboratory, CSIRO ICT Centre
PO Box 883, Kenmore, Qld 4069, Australia
Email: Firstname.Lastname@csiro.au

***Abstract*— This paper presents an enhanced algorithm for matching laser scan maps using histogram correlations. The histogram representation effectively summarizes a map's salient features such that pairs of maps can be matched efficiently without any prior guess as to their alignment. The histogram matching algorithm has been enhanced in order to work well in outdoor unstructured environments by using entropy metrics, weighted histograms and proper thresholding of quality metrics. Thus our large-scale scan-matching SLAM implementation has a vastly improved ability to close large loops in real-time even when odometry is not available. Our experimental results have demonstrated a successful mapping of the largest area ever mapped to date using only a single laser scanner. We also demonstrate our ability to solve the lost robot problem by localizing a robot to a previously built map without any prior initialization.**

***Index Terms*— Scan matching, closing the loop, SLAM.**

I. INTRODUCTION AND PREVIOUS WORK

One of the major difficulties that large scale SLAM algorithms encounter is the closing of large loops. One common obstacle is the difficulty of recognizing when a robot has returned to a previously mapped area; this problem is exacerbated when an environment contains repeated structures, in which case multiple ambiguous matches may exist. Additionally, the unbounded growth of open loop uncertainties makes it necessary to check increasingly larger areas for revisits. In worst case, the unbounded uncertainties result in a situation equivalent to the lost robot problem, a case where there is no prior knowledge of the robot's location with respect to an existing map.

The ready availability of 2D laser scanners has resulted in the development of many scan matching algorithms for map-making, the majority of which require an initial guess to converge on the alignment between scans [1], [2], [3], [4]. While these algorithms are suitable for pose tracking and incremental map building, they are unacceptable for global localization because they will fail to converge. Those algorithms which are appropriate for global localization generally require polygonal environments as they depend upon the extraction of geometric features such as lines and corners [5], [6], [7], [8]. These approaches are likely to fail in unstructured outdoor scenes.

The localization and map-making algorithms utilized in this paper run within the *Atlas* framework [9], [10], [11]. The *Atlas* framework was developed to be a hybrid metrical/topological approach to SLAM that achieves efficient mapping of a variety of large-scale environments. The framework is maintained as a graph of coordinate frames: each vertex in the graph contains a local map, and each edge represents the transformation between adjacent map frames. Loops are detected and closed by matching maps and adding the alignment transformation to the graph. The framework's iterative scan match implementation works well for local mapping of outdoor environments; however, the same iterative approach when applied to map matching was inefficient in finding loops when the prior uncertainties are large.

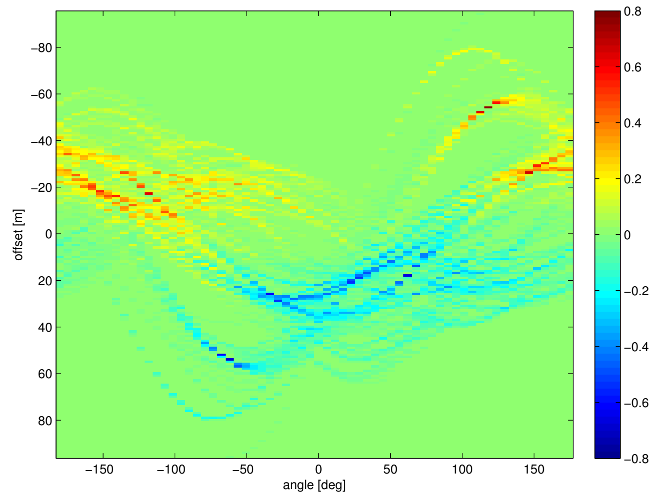

Fig. 1. A sample set of weighted projection histograms displayed as an image. The negative weights help to distinguish contributions from scan points with opposite orientations.

The improvements to the scan map matching algorithm described in this paper employ a histogram correlation approach very similar to [6], but which has been enhanced to work well in unstructured outdoor environments.

The paper is organized as follows. In Section II, the histogram correlation algorithm and its enhancements are described. Section III demonstrates how the approach makes mapping a large outdoor area and solving the lost robot problem feasible. Finally we conclude in Section IV.

II. THE ALGORITHM

The algorithm in this paper is described in parts: first, the generation of the histograms for each local map, and second, the correlation of local maps' histograms to determine whether there is a match, and if so, to find the ensuing transformation between their coordinate frames.

A. Histogram Map Representation

Each local map requires a compact representation of its salient characteristics; this representation can be used to distinguish the local map from other maps and to determine its transformation to another map of the same area. This representation consists of an *orientation histogram* of the scan normals plus a set of weighted *projection histograms* created from the orthogonal projections of scan points onto lines of varying orientation.

The *orientation histogram* is used to compute the rotational offset between pairs of local maps irrespective of any translational offsets. The peaks of the orientation histogram, most apparent when there are numerous flat surfaces visible in the laser scans, represent the dominant surface orientations. We have determined empirically that an angle bin size of about 3 degrees works well in our environment of industrial buildings and residential neighborhoods. In general, the histogram's bin size should be commensurate with the noise and certainty of the scans in the local map.

The set of weighted *projection histograms* are used to determine the translational offsets between pairs of maps once their rotational offset has been determined. Each projection histogram $H_p(\theta_p, d)$ is generated by orthogonally projecting every scan point, (x_i, y_i), on to a line with angle θ_p and creating a histogram of the offsets d_i of the points which have been weighted by the dot product of the scan normals (n_{x_i}, n_{y_i}) with the line:

$$d_i = x_i \cos(\theta_p) + y_i \sin(\theta_p) \quad (1)$$

$$H_p(\theta_p, d) = \sum_{\|d_i - d\| < \frac{\triangle}{2}} n_{x_i} \cos(\theta_p) + n_{y_i} \sin(\theta_p) \quad (2)$$

The dynamic range of the projection histograms is enhanced by weighting each point according to its surface orientation. Points with orientations parallel to the projection line are de-weighted so that they do not blur the histogram, whereas points with surfaces perpendicular to the line are given more weight. Furthermore, since the weights can be negative, it is possible to distinguish contributions from scan points with opposite normal orientations. This improves the dynamic range and saliency of the projection histograms, because the cumulative contribution from long walls will not wash out parallel projections, and perpendicular walls will only match with walls of the same orientation.

The orientation of each projection line and ultimately the number of projection histograms generated are determined by the number of angle bins. For the projection histograms, the size of the offset bins should be small enough that the details of the environment's structure are captured, but not so small that there are too few points in the bins. We have determined empirically that a bin size of $1\,\mathrm{m}$ works well for our environment.

An example of the weighted projection histograms is depicted as an image in Fig. 1.

B. Histogram Matching

The goal of the histogram matching algorithm is to determine quickly whether or not any pair of maps match, and

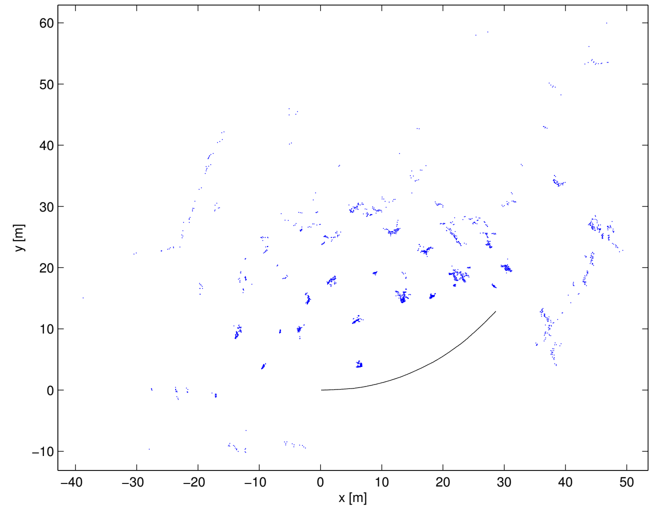

(a)

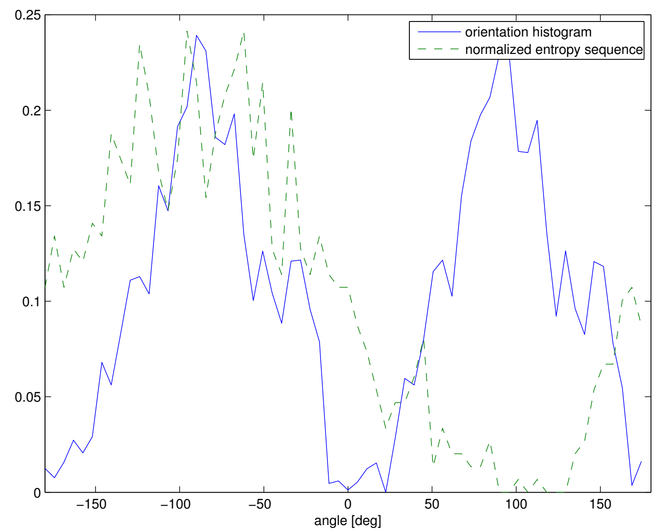

(b)

Fig. 2. (a) An example environment where the entropy sequence gives better results than the orientation histogram when computing rotational alignment. (b) The peaks in the orientation histogram are a result of noise; correlations cannot be used for reliable matching. The normalized entropy sequence of projection histograms, however, is more likely to produce a correct match since it is not dominated by noise.

if so, to find the transformation between their coordinate frames.

The first step of the matching process is to compute the correlation between the maps' orientation histograms in order to determine possible rotational offsets. In order to avoid boundary effects, the correlation is computed by circular convolution, where the histograms have been normalized by their Frobenius norms such that the circular convolution of the histogram with itself has a maximum value of 1.0, see Fig. 3(c).

As an alternative to using orientation histograms, it is also possible to compute the rotational offset using the sequence of entropy measures $E(\theta_p)$ for each projection histogram $H(\theta_p, d)$.

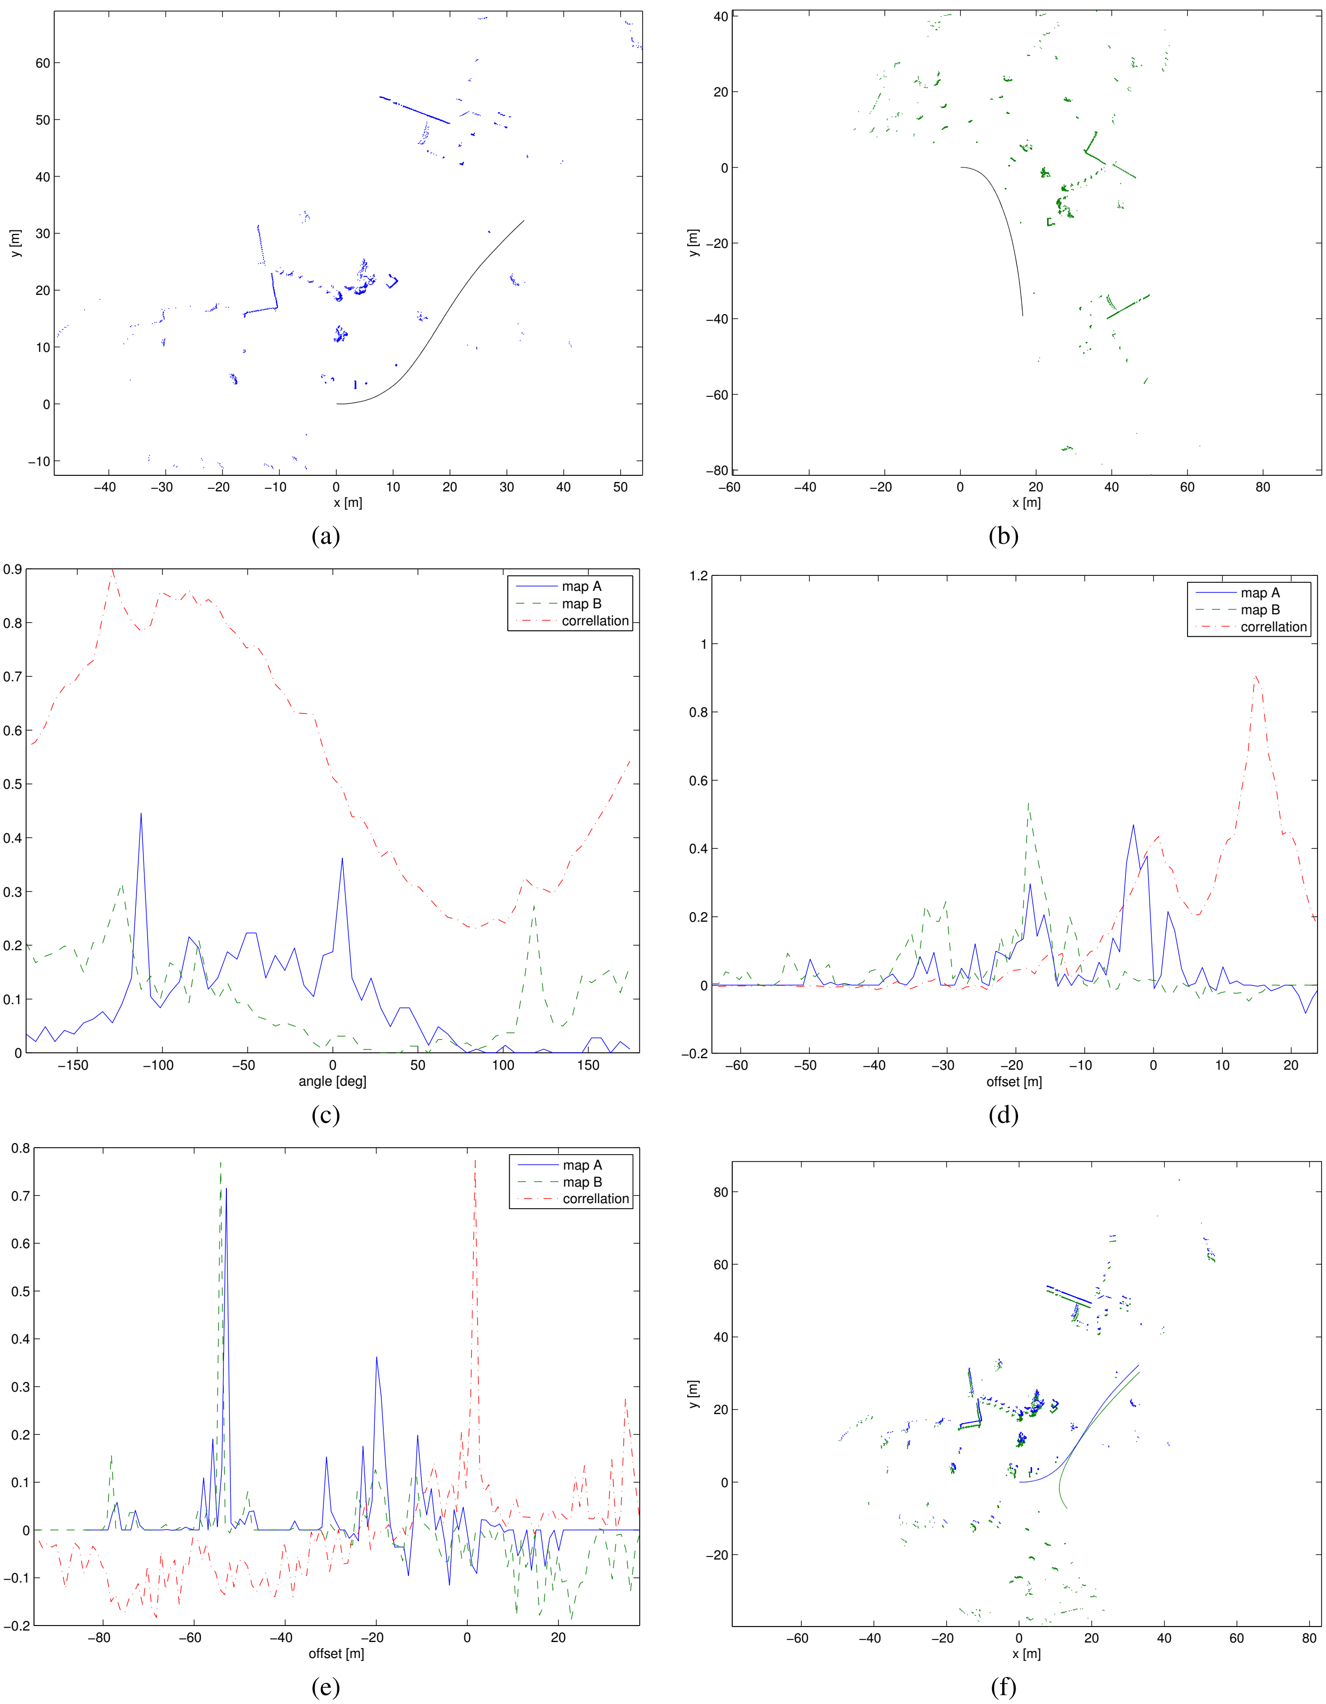

Fig. 3. (a) and (b) show two example local scan-match maps, where the dots are the scan points, and the solid line is the trajectory of the vehicle. (c) The orientation histograms and their correlation. (d) and (e) The two orthogonal projection histograms used to score the match. (f) The maps brought into alignment using the peaks of the match.

$$\hat{H}(\theta_p, d) = \frac{\|H(\theta_p, d)\|}{\sum_d \|H(\theta_p, d)\|} \quad (3)$$

$$E(\theta_p) = \sum_d \hat{H}(\theta_p, d) \log\left(\hat{H}(\theta_p, d)\right) \quad (4)$$

The absolute value of each projection histogram element is necessary since the weights may be negative.

Prior to matching, the entropy sequence is normalized to make it commensurate with an orientation histogram by shifting such that the minimum is zero and dividing by the Frobenius norm.

$$\hat{E}(\theta_p) = -\frac{E(\theta_p) - \max(E(\theta_p))}{\sqrt{\sum_{\theta_p}(E(\theta_p) - \max(E(\theta_p)))^2}} \quad (5)$$

The entropy sequence is negated because minimum entropy projections correspond to angles where the points are tightly packed.

Since the sequence of entropy measures repeats every $180°$, each peak in the correlation of entropy sequences produces two potential orientation offsets. This ambiguity, however, is often resolved during the next step due to the fact that the false offset will have a low correlation in the projection histograms.

This approach works better in unstructured outdoor environments where there are few dominant peaks in the orientation histograms but there is nevertheless a strong signal in the normalized sequence of projection entropies, see Fig. 2. In such environments, orientation histograms are too noisy for reliable matching.

Once the rotational offset has been computed, two perpendicular projection histograms are selected to compute the translational offsets. It is best to use the projection histogram with the least entropy (typically the one corresponding to the largest peak in the orientation histogram) and its orthogonal partner, as this choice slightly improves the saliency of the match. The peaks in the correlation between each chosen histogram and its correspondent in the other map represent the offsets that are in the direction of the projection lines, see Fig. 3(d) and (e).

The translation vector between the maps' coordinate frames is calculated by solving the simple linear system for t_x and t_y:

$$\begin{pmatrix} \cos(\theta_p) & -\sin(\theta_p) \\ \sin(\theta_p) & \cos(\theta_p) \end{pmatrix} \begin{pmatrix} t_x \\ t_y \end{pmatrix} = \begin{pmatrix} d_{\theta_p} \\ d_{\theta_p + \frac{\pi}{2}} \end{pmatrix} \quad (6)$$

where θ_p is the angle that the first projection line makes with the x-axis, d_{θ_p} is the offset of the maximum peak of the correlation of the first projection histogram, and $d_{\theta_p+\frac{\pi}{2}}$ is the offset computed from the correlation of the second projection histogram.

The accuracy of the resulting transformation is dependant upon the bin sizes. Even though the alignment may be coarse, as in Fig. 3(f), the accuracy should be sufficient for the standard iterative scan matching algorithm [1] to converge on the precise match.

While it is possible to compute a transformation for any two maps, that transformation will only make sense

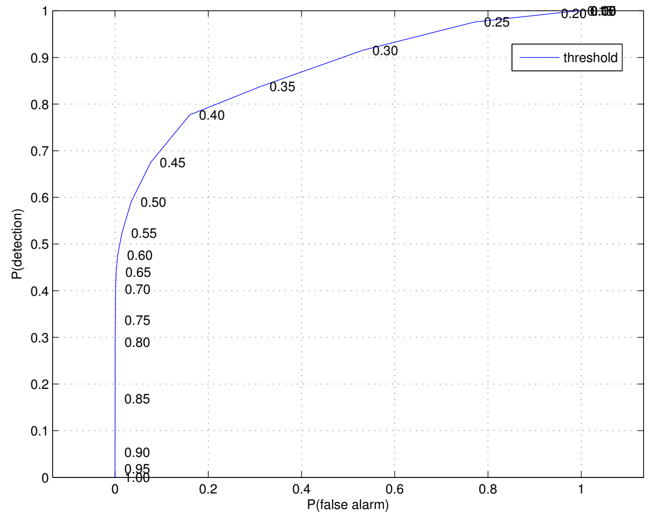

(a)

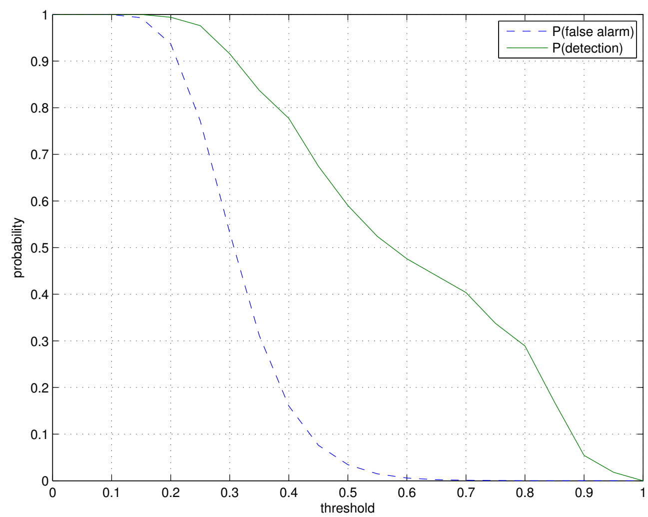

(b)

Fig. 4. The probability of detection vs. probability of false alarms for various match correlation threshold settings. Since false matches are problematic, we have chosen a threshold of .6 which had a P(FA) of less than 1 percent and a detection probability of 48%.

if the maps actually overlap. Therefore, it is important to assess the match quality before accepting a match. Match quality is assessed by multiplying the peak value of the orientation histogram correlation with the peak values of the two projection histogram correlations.

The threshold on the quality metric for accepting matches is determined empirically. The probabilities of detection, $P(D)$, and false alarms, $P(FA)$, for each threshold value are estimated by analyzing a large set of quality metrics of known correct and false matches, see Fig. 4. Using these probability functions, we have chosen a threshold value of 0.6 which has a $P(D)$ of 0.48, and a $P(FA)$ of less than 1%.

In any case, all transformations computed with valid quality metrics are subsequently verified by the iterative scan matching alignment. Furthermore, any false matches from truly ambiguous environments are filtered by the cycle

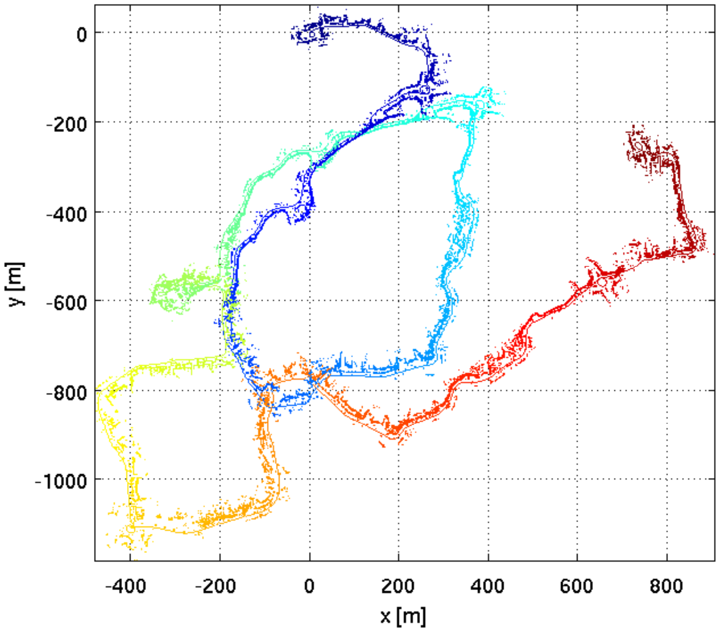

(a)

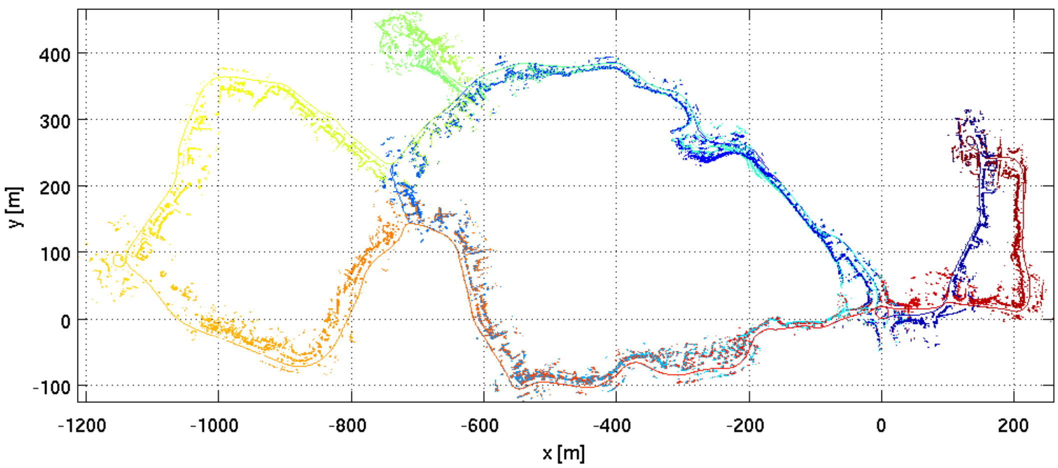

(b)

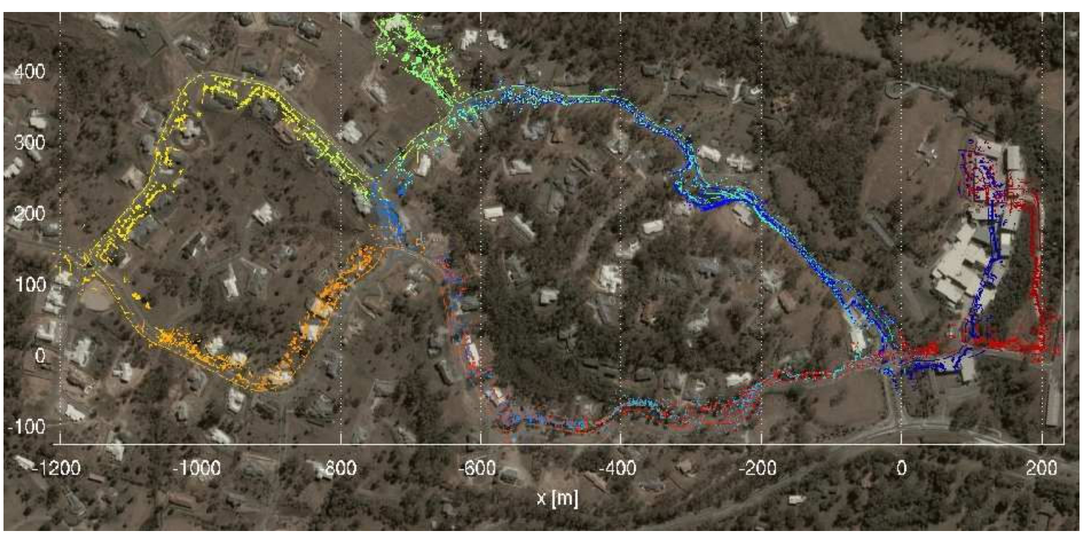

(c)

Fig. 5. (a) The map showing no loop closures. (b) The optimized map after distributing the errors around the loops. (c) The map overlaid onto an aerial image.

Fig. 6. The laser is mounted sideways to the front of the vehicle.

verification procedure in the *Atlas* framework.

III. RESULTS

This section presents the results of a mapping of our industrial compound and adjacent neighborhood. These results demonstrate the effectiveness of the histogram matching approach for closing loops in large scale environments: of particular note is the fact that we did not use odometry in generating the local maps.

Furthermore, since histogram matching allows us to recognize places we have been before without any prior transformation, we are able to use the technique to solve the lost robot problem.

A. Large Scale Mapping without Odometry

The experiment utilized a single SICK LMS laser scanner mounted on the front of a car, positioned such that its $180°$ field-of-view is oriented to the left, see Fig. 6. Each laser scan consists of 361 points at $40\,\mathrm{Hz}$. GPS was recorded for ground truth, but was not used in any of the map-making. The car was driven around our industrial compound and adjacent neighborhood at 25 $\mathrm{km/h}$ speeds for a total path length of 6.7 km.

During the run, 231 local maps were generated, each map containing 15 saved scans spaced approximately 2 m apart. Maps were processed with the Scan-Match implementation of the *Atlas* framework which had been modified slightly to suppress map reuse in order that the network structure be more apparent.

Fig. 5(a) illustrates what the global map looks like when histogram matching is disabled. Fig. 5(b) depicts the globally optimized map with histogram matching enabled, and Fig. 5(c) overlays the map onto an aerial photo, demonstrating the correctness of the map's topology.

Repeated loops are indicated by off diagonal entries in the graph adjacency matrix, depicted in Fig. 7(a). The same network structure is visible in the image of histogram match quality metrics, Fig. 7(b). The map connectivity percentage (as described in Chapter 7 of [9]) is 97%.

Without histogram matching, it would not have been possible to recognize loop closures of this size.

B. Lost Robot Problem

The lost robot problem occurs when, on startup, an existing map has been loaded but we do not initialize the robot's location in the map. In such a situation, the robot begins by generating a disjoint *Atlas* graph while continually attempting to match the current map to any previously existing map. Once a successful match has been found and verified, the two graphs will be connected and the robot is found.

To demonstrate solving the lost robot problem, we reprocessed the same data set in two parts. The process was initialized with a graph of the first large loop generated from the initial 11 minutes of data. The next 4 minutes of data was discarded to ensure that the robot would be completely lost in a new area, and then the map-making process was resumed.

Since this location was not in the loaded maps a disjoint graph was established. After 2 minutes of travel, the vehicle first re-entered an area covered by the previously loaded maps, and the histogram matching algorithm immediately discovered the first link. Verification of the link was suspended until the discovery of the next link to the existing area. At this point, the cycle verification procedure verified both links and the graphs were connected.

Please refer to the attached video for an animation of this experiment.

IV. CONCLUSION AND FUTURE WORK

This paper has presented an enhanced algorithm for matching laser scan maps using histogram correlations. The histogram representation effectively summarizes a map's salient features such that pairs of maps can be matched efficiently without any prior guess as to their alignment. The algorithm has been enhanced in order to work well in outdoor unstructured environments in several ways: A sequence of entropy measures of projection histograms can be utilized in place of orientation histograms, the projection histograms are made more salient by weighting them by the scan point's orientations, and an empirical method is used to determine a threshold on the quality metric that yields few false alarms.

This technique vastly improves the *Atlas* scan match implementation's ability to close large loops in real-time even when odometry is not available. Moreover, the ability to discern matches quickly and without the need of an initial guess as to the alignment allows the lost robot problem to be solved in real-time. Our experimental results have demonstrated a successful mapping of the largest area ever mapped to date using only a single laser scanner. We have also demonstrated our ability to solve the lost robot problem by localizing a robot to a previously built map without any prior initialization.

Since the overlaps between disjoint maps can now be easily discovered, we would like to investigate the possibility of automatically combining the maps generated from multiple disjoint runs or even dividing the mapping task among multiple robots. Further future work will involve extending

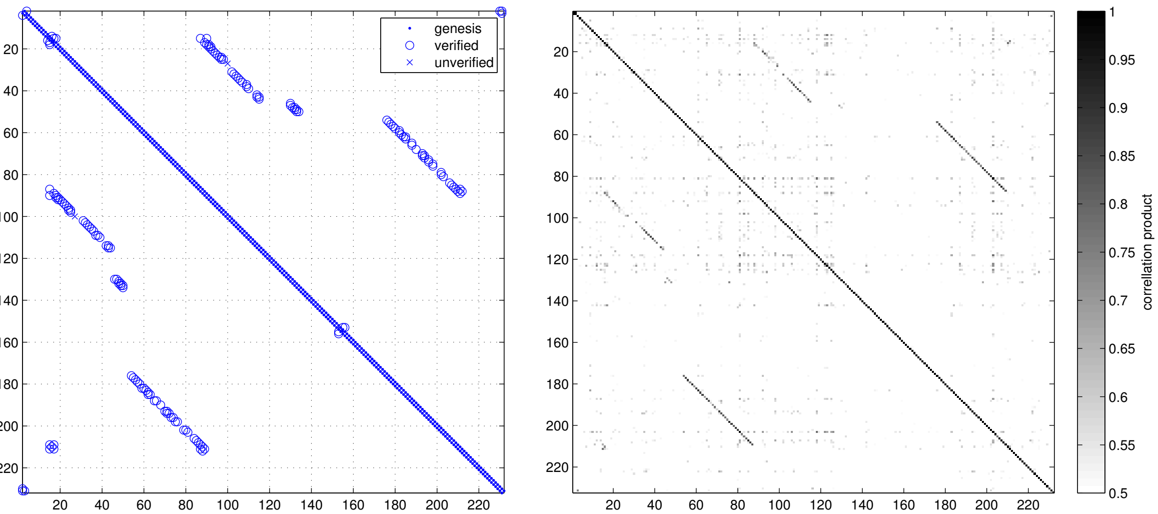

(a) (b)

Fig. 7. (a) The *Atlas* graph's adjacency matrix. (b) The product of the match correlations for each pair of maps. The repeated loops of the environment are indicated by the off diagonal strips. Highly ambiguous maps are indicated by maps that match many columns in the matrix, i.e. the maps 75-80 are of a strip of road with trees planted at regular intervals.

the histogram matching approach to 3D laser maps and other map representations such as those generated by sonar, radar, or stereo vision.

ACKNOWLEDGMENTS

This work was funded by the CSIRO ICT Centre under the ROVER and Dependable Field Robotics projects. The authors would like to thank the Autonomous Systems Laboratory team for their support of this work. Special thanks go to Peter Corke, Paul Flick, Navid Nourani, Cedric Pradalier, Ashley Tews and Kane Usher, who all contributed to this work.

REFERENCES

[1] F. Lu and E. Milios, "Robot pose estimation in unknown environments by matching 2D range scans," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Seattle, USA, June 1994, pp. 935–938.

[2] I. Cox, "Blanche – an experiment in guidance and navigation of an autonomous robot vehicle," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 2, pp. 193–204, 1991.

[3] J. Gutmann and K. Konolige, "Incremental mapping of large cyclic environments," in *Proceedingsof the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, Monterey, California, 1999, pp. 318–325. [Online]. Available: citeseer.ist.psu.edu/gutmann00incremental.html

[4] K. Lingermann, H. Surmann, A. Nuchter, and J. Hertzberg, "Indoor and outdoor localization for fast mobile robots," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sendai, Japan, October 2004, pp. 2185–2190.

[5] X. Zezhong, L. Jilin, and X. Zhiyu, "Scan matching based on CLS relationships," in *Proceedings of the IEEE International Conference on Robotics, Intelligent Systems and Signal Processing*, Chansha, China, October 2003, pp. 99–104.

[6] G. Weiss, C. Wetzler, and E. von Puttkamer, "Keeping track of position and orientation of moving indoor systems by correlation of range-finder scans," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Munich, Germany, September 1994, pp. 595–601.

[7] T. Röfer, "Using histogram correlation to create consistent laser scan maps," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, October 2002, pp. 625–630.

[8] J.-S. Gutmann, T. Weigel, and B. Nebel, "A fast, accurate, and robust method for self-localization in polygonal environments using laser-range-finders," *Advanced Robotics*, vol. 14, no. 8, pp. 651–668, 2001.

[9] M. C. Bosse, "ATLAS: A Framework for Large Scale Automated Mapping and Localization," Ph.D. dissertation, Massachusetts Institute of Technology, Feb 2004.

[10] M. Bosse, P. Newman, J. Leonard, M. Soika, W. Feiten, and S. Teller, "An Atlas Framework for scalable mapping," in *International Conference on Robotics and Automation*, Taipei, Taiwan, September 2003, pp. 1899–1906.

[11] M. Bosse, P. Newman, J. Leonard, and S. Teller, "Simultaneous localization and map building in large-scale cyclic environments using the Atlas Framework," *International Journal of Robotics Research*, vol. 23, no. 12, pp. 1113–1139, December 2004.