Robust Real-Time Local Laser Scanner Registration with Uncertainty Estimation

Justin Carlson, Charles Thorpe, David L. Duke

► To cite this version:

Justin Carlson, Charles Thorpe, David L. Duke. Robust Real-Time Local Laser Scanner Registration with Uncertainty Estimation. 6th International Conference on Field and Service Robotics - FSR 2007, Jul 2007, Chamonix, France. inria-00195850

HAL Id: inria-00195850

<https://hal.inria.fr/inria-00195850>

Submitted on 11 Dec 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Robust Real-Time Local Laser Scanner Registration with Uncertainty Estimation

Justin Carlson   Charles Thorpe   David L. Duke
Carnegie Mellon University - Qatar Campus
{justinca|cet|dl Duke}@ri.cmu.edu

Abstract

We present a fast, robust method for registering successive laser rangefinder scans. Correspondences between the current scan and previous scans are determined. Gaussian uncertainties of the correspondences are generated from the data, and are used to fuse the data together into a unified egomotion estimate using a Kalman process. Robustness is increased by using a RANSAC variant to avoid invalid point correspondences. The algorithm is very fast; computational and memory requirements are $O(n \log n)$ where n is the number of points in a scan. Additionally, a covariance suitable for use in SLAM and filter techniques is cogenerated with the egomotion estimate. Results in large indoor environments are presented.

1 Introduction and Motivation

Maps are extremely useful in mobile robotics; they are of immense value to a robot tasked with localization, planning, and navigation. At present, it is common to use dead reckoning with range sensing in simultaneous localization and mapping (SLAM) algorithms to bound global error in unknown environments. The effectiveness of such algorithms depends in large part on the quality of dead reckoning. As dead reckoning degrades, the search for landmark correspondences grows, and becomes the dominating factor in most SLAM approaches. Although high-quality dead reckoning solutions exist, their components, such as fiber-optic gyroscopes, tend to be expensive.

At the same time, laser rangefinder registration is capable of providing significant improvements to dead reckoning estimates; however, to be useful in a larger context, such registrations must provide useful error metrics to higher level systems. This work incorporates data-derived uncertainties in scanner registrations to both improve registration results and generate accurate uncertainty estimates to higher level systems.

2 Related Work

Most of the popular methods for laser scanner registration find their roots in the Iterative Closest Point algorithm presented in [3]. [10] and [11] present a modified ICP-based approach to estimate robot egomotion using different correspondence algorithms to estimate rotation and translation, but the work doesn't take into account the quality of individual point matches nor does it generate an estimate of the quality of the registration. In [14] and [13], a Bayesian framework is added to separate moving objects from stationary landmarks; this work is orthogonal to and compatible with that framework. [5] and [4] use a grid to divide previous scans into localized bins. Covariances are calculated for these bins, and new scans are registered into old ones by numerical maximization of a pseudo probability density function. However, the artificial division of the world and required selection of representation is not straightforward.

[2] presents two methods for generating covariances from the Iterative Dual Correspondence algorithm; a real-time approach is based on examining the minimized error function, and an offline approach shifts registration around a minimum and uses a correlation metric to estimate registration uncertainty. The offline approach generates very good results, but is not suitable for a real-time system.

[12] uses Gaussian noise models to improve estimation of linear surfaces in laser scans. [16] uses line segments to generate localization matches. However, these methods require specific structure in the environment; this method does not.

3 Problem Definition

Consider a mobile robot which is capable of translation and rotation on a flat surface. This robot is equipped with a rigidly fixed range sensor which generates sets of range returns from a plane parallel with the surface of motion. At periodic intervals while moving, the robot takes a set of readings. The sensor and environment are assumed to provide a significant number of returns, but the environment is specifically not constrained to be linear or smooth.

In order to estimate the robot's path, we wish to generate an estimate of the rigid transformation needed to most accurately register successive sets of measurements. To enable incorporation into a higher level framework, we wish to simultaneously generate a corresponding metric of the uncertainty of the estimate.

Figure 1: Sample Match and Covariance Ellipse

4 Algorithm

We deal with two sets of returns: $\mathbf{A}$ is collected, the robot moves, then $\mathbf{B}$ is collected. The motion of the robot between scans is $m = (m_x, m_y, m_\theta)^T$. We seek an estimate $\hat{m}$ of this motion as well as a covariance Q for that estimate.

Each return in a set is parameterized as a Cartesian coordinate in a frame with its origin at the robot's center of rotation at the time $\mathbf{A}$ was gathered. Generating this parameterization for $\mathbf{B}$ requires some initial value for $\hat{m}$. If no information is available to seed the registration process, the initial estimate can be 0. Alternatively, the initial estimate can come from odometry or a motion model.

For each point a_i in $\mathbf{A}$, we find b_i^* , the closest point on the polyline defined by the bearing-sorted points of $\mathbf{B}$. All returns in $\mathbf{B}$ within a fixed bearing window of b_i^* are used to calculate a 2×2 covariance matrix, R_i .

Figure 1 illustrates the various aspects of a single generated point match.

The next step is best discussed in the framework of an update step of a Kalman Filter. One popular formulation of the filter, from [15], can be summarized as follows:

Given a state estimate $\hat{x}_k$ at time k with associated covariance matrix P , if we have an observation of the form:

$$z_k = H_k x_k + v_k$$

where

$$p(v_k) \sim N(0, R_k)$$

$\hat{x}$ and P are updated as follows:

$$K_k = P_k H_k^T (H_k P_k H_k^T + R)^{-1}$$

$$\hat{x}_{k+1} = \hat{x}_k + K_k (z_k - H \hat{x}_k)$$

$$P_{k+1} = (I - K_k H_k) P_k$$

We initialize the filter with a $\hat{x}$ of 0 and we let P_0 be an arbitrarily large diagonal matrix. We consider each point match $z_i = b_i^* - a_i$ to be an independent observation of the state to be estimated, with associated covariance R_i . The mapping from state to observation is nonlinear, so we substitute the Jacobian:

$$H = \begin{pmatrix} 1 & 0 & -d_i \sin(\phi_i) \\ 0 & 1 & d_i \cos(\phi_i) \end{pmatrix}$$

where $d_i = \|a_i\|$ and $\phi_i = \text{atan2}(a_{iy}, a_{ix})$.

After processing all the point correspondences, the filter contains an incremental refinement to the overall transformation:

$$\hat{m}_{k+1} = \hat{m}_k + \hat{x}$$

As with ICP, as the estimate is refined, better correspondences are found between scans. Thus, we generate new matches with the refined estimate and iterate until convergence. When converged, Q is simply the calculated P from the final iteration.

5 Robustness

This algorithm is a modified least squares optimization. As such, it inherits the sensitivity of such methods to outliers. The most common cause of such outliers is moving objects in the scene.

The resulting scan registrations are greatly improved by adding a RANSAC step to the point selection [6]. The process for doing so is straightforward. n point matches are drawn from the registration,

and used to generate an estimate and covariance. Inliers from the rest of the set are determined from the Mahalanobis distance of the generated estimate. If a sufficient proportion of the points are inliers, the overall estimate is recalculated using the inliers, otherwise a new set of n points is drawn. Empirically, generating a sample with 4 single point motion estimates does an excellent job of rejecting outliers while not being computationally burdensome.

6 Test Platform

Data is presented from two different robotic platforms.

Indoor results are presented using data gathered from a heavily modified Segway Robotic Mobility Platform (RMP). The RMP carries a SICK LMS mounted 45 cm from the ground. Due to the inverted pendulum nature of the RMP motion, the SICK is mounted on a servo which allows the robot to keep it nominally parallel to the ground. The RMP provides state updates, including odometry and kinematic estimates, at 100 Hz. The SICK provides scans of 181 returns spaced 1 degree apart at 75 Hz.

The RMP data sets presented here were gathered at speeds of up to 2.5 meters per second, and angular rates of up to 2 radians per second.

Additional testing used data sets collected at Stanford and graciously provided online at [8]. The Stanford data sets were collected in the Gates building using a SICK LMS on a Pioneer XT.

7 Accuracy

For this paper, we call our algorithm “Kalreg”, and compare our results against both dead reckoning and an implementation of the Iterative Dual Correspondence algorithm [10], a widely used scan-matching system.

To make comparisons more valid, results shown use odometry to seed the registration process in all cases. The “Kalreg” results do not fuse the odometry measurements into the final registration.

Figure 2 shows a run consisting of about 31,000 registered scans over a path of approximate 750 meters. Both registration algorithms encounter some difficulties with an open area in the middle of the run; the heading estimate in both cases is noticeably skewed.

Figure 3 uses the first 50,000 scans of a traversal of the Gates building at Stanford. This data set is intended for SLAM-type appli-

Dead Reckoning

IDC

Kalreg

Figure 2: Result: CMU-Qatar building, 31000 scans

Dead Reckoning

IDC

Kalreg

Figure 3: Result: Stanford Gates building, 50,000 scans

cations, and contains many loops and backtracks. Despite the lack of loop closure in these results, our algorithm generates an overall registration which is visibly more consistent than the IDC result.

8 Complexity and Performance

The algorithm is designed to be very efficient and fast. With the exception of finding correspondences, the entire algorithm is linear in the number of points in a scan.

Given a bearing-ordered sensor, several shortcuts can be used to find reasonable candidates for matches in constant time. In the general case, points can be sorted into bearing order, making approximate correspondence search an $O(n \log(n))$ proposition.

Other computational factors that are significant are the number of RANSAC trials needed to find a consistent estimate and the number of iterations needed to converge. In the environments tested, where most of the scene is static, neither factor appears to be onerous.

Interestingly, the uncertainty estimates for matches reduce the iterations required for convergence, often significantly. One problem in many ICP approaches is the tendency of linear scene segments correspondences with near-zero offsets to cause a “drag” effect on the convergence. The original ICP paper proposes a form of gradient descent to bypass the problem. In this method, such segments are associated with appropriate covariances, so the “drag” effect is dramatically lessened, obviating such a requirement.

The authors’ most recent C++ implementation runs comfortably in real-time on a 2.4 Ghz Pentium 4 workstation. It is possible to perform the bulk of necessary computations as highly pipelinable parallelized vector operations using SIMD instructions. The authors did write one mildly optimized implementation which was capable of registering scans at 320 Hz. With more effort, this likely could be improved by an order of magnitude.

9 Conclusion

This paper put forth a new method for estimating robot egomotion with explicit uncertainty reasoning. The algorithm provides good registration results comfortably in real-time, and generates uncertainty estimates concurrently.

10 Future Work

One of the goals in development of this algorithm was to generate useful but pessimistic error estimates for registration. Ongoing work is being carried out on validation of the generated error estimates using loop closures. Additionally, there are no obvious obstacles to extending the developed framework into registering points in three dimensions.

References

- [1] Y. Bar-Shalom and X.-R. Li. *Estimation and Tracking: Principles, Techniques and Software*. Artech House, Inc., 1993.
- [2] O. Bengtsson and A.-J. Baerveldt. Robot localization based on scan-matching - estimating the covariance matrix for the icd algorithm. *Robotics and Autonomous Systems*, 44:29–40, 2003.
- [3] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):239–256, 1992.
- [4] P. Biber, S. Fleck, and W. Straßer. A probabilistic framework for robust and accurate matching of point clouds. In *26th Pattern Recognition Symposium (DAGM 04)*, 2004.
- [5] P. Biber and W. Straßer. The normal distributions transform: A new approach to laser scan matching. In *IEEE/RJS International Conference on Intelligent Robots and Systems*, 2003.
- [6] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981.
- [7] J. Guivant and E. Nebot. Optimization of the simultaneous localization and map building algorithm for real time implementation. *IEEE Transactions on Robotics and Automation*, May 2001.
- [8] A. Howard and N. Roy. The robotics data set repository (radish), 2007.

- [9] J. K. Lindsey. *Introduction to Applied Statistics: A Modelling Approach*. Oxford University Press, second edition, 2004.
- [10] F. Lu and E. Miliotis. Robot pose estimation in unknown environments by matching 2d range scans. In *CVPR94*, pages 935–938, 1994.
- [11] F. Lu and E. Miliotis. Globally consistent range scan alignment for environment mapping. *Auton. Robots*, 4(4):333–349, 1997.
- [12] S. T. Pfister, S. I. Roumeliotis, and J. W. Burdick. Weighted line fitting algorithms for mobile robot map building and efficient data representation. In *Proceedings of the 2003 IEEE International Conference on Robotics and Automation*, pages 1304–1311, 2003.
- [13] C.-C. Wang. *Simultaneous Localization, Mapping and Moving Object Tracking*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, April 2004.
- [14] C.-C. Wang, C. Thorpe, and S. Thrun. Online simultaneous localization and mapping with detection and tracking of moving objects: Theory and results from a ground vehicle in crowded urban areas. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Taipei, Taiwan, September 2003.
- [15] G. Welch and G. Bishop. An introduction to the Kalman filter. Technical report, University of North Carolina at Chapel Hill, 2006.
- [16] X. Zezhong, L. Jilin, and X. Zhiyu. Scan matching based on cls relationships. In *IEEE International Conference on Robotics, Intelligent Systems, and Signal Processing*, pages 99–104, October 2003.