

---

# Recognizing Places with Weak Evidence

---

Edwin Olson

EOLSON@MIT.EDU

MIT Computer Science and Artificial Intelligence Laboratory, 32 Vassar Street, Cambridge MA, 02139 USA

## 1. Introduction

A central problem in robot navigation is recognizing when a robot is somewhere that it has been before. Without “loop closing”, the robot’s position uncertainty increases without bound; consequently navigation, map-building, and other common robot tasks become impossible.

In the presence of globally-unique features (RFID tags embedded in the environment, for example), closing the loop is trivial. More commonly, however, evidence of a loop closure is weak. Planar laser scanners are a typical robotic sensor, producing cross-sectional maps of environments (see Fig. 1). Unfortunately, indoor environments tend to be composed of similar-looking elements (corners, walls, etc.), leading to many false matches.




Figure 1. Three laser scans from CSAIL-G7. Some environments present rich alignment cues (left; an elevator lobby). Incorrect matches arise from both spartan areas (middle; a corridor) and from repetitive/cluttered areas (right; a cubical farm).

We present an algorithm that considers groups of several dozen loop closure hypotheses and robustly rejects the incorrect hypotheses. This paper’s central contribution is showing how to map the loop-closing problem onto the Single Cluster Graph Partitioning (SCGP) problem, which has an efficient solution (Olson et al., 2005).

We present results using laser data, including a data set collected at CSAIL and a standard mapping benchmark data set. We have only explored laser data, but our approach may be relevant to other sensing modalities as well.

## 2. Approach

The motion of a robot is well-represented by a graph whose nodes represent positions where the environment was ob-




Figure 2. Place recognition. A robot determines that it is near a previously-visited location and attempts to align laser scans from nearby poses, generating a number of hypotheses (left). Any two hypotheses form a loop whose cumulative rigid-body transformation *should* be the identity matrix (right); this allows us to test two hypotheses for consistency.

served, and whose edges represent the physical motion between nodes (i.e., rigid-body transformations *with covariances*); see Fig. 2a. As the robot moves around, it builds an acyclic chain graph by connecting sensor observations (nodes) with edges according to the robot’s dead-reckoning sensors (e.g., wheel odometry).

Recognizing a place is equivalent to adding a new edge to the graph; this makes the graph cyclic. Cycles make the position of the nodes *over-determined*, which prevents position uncertainty from increasing without bound and allows SLAM algorithms to produce good maps (Olson et al., 2006).

Using a conservative model of the robot’s dead-reckoning error, we can compute the set of nodes that *might* be within sensor range of any other node. Our approach attempts to align the laser scans of all such pairs of nodes, producing a large set of graph edges—many of which are incorrect.

The essential idea of the approach is that we can compute the consistency of any two hypotheses by considering the loop that they form (in conjunction with portions of the trajectory prior), as in Fig. 2b. The product of the rigid-body transformations around *any* loop should be the identity transform. Consequently, we can compute a pairwise

consistency metric: the probability that the loop is the identity matrix assuming that both hypotheses are true.

True hypotheses tend to be mutually consistent. In contrast, false hypotheses are not generally consistent, since the particular failure mode affecting one matching operation tends to change from one hypothesis to another. Consequently, we can look for correct hypotheses by searching for a subset of hypotheses that are maximally consistent.

One reasonable metric for the mutual consistency of a set of hypotheses is the *average consistency*, the sum of all the pairwise consistencies of the subset divided by the number of hypotheses in that set. Let  $A$  be a matrix such that  $A_{i,j}$  is the pairwise consistency of hypotheses  $i$  and  $j$ , and let  $u$  denote an indicator vector such that if  $u_i = 1$  then hypothesis  $i$  is in the subset. Average consistency can then be written:

$$\bar{c} = \frac{u^T A u}{u^T u} \quad (1)$$

We have previously shown how to compute an indicator vector  $u$  which approximately maximizes this expression (Olson et al., 2005). Briefly,  $u$  is a discretization of the dominant eigenvector of  $A$ .

### 3. Results

Once a set of true hypotheses has been identified, the hypotheses can be added to the graph and a SLAM algorithm used to generate an improved map.




Figure 3. Place recognition example. The same physical place is visited several times; before (top) and after (bottom) automatic loop closure.

We ran our algorithm on a data set collected from CSAIL-G7 (Fig. 3 and 4), as well as the benchmark SLAM data set collected at the Intel Research Center (Fig. 5). The quality of the maps are obviously improved. Animations of our algorithms are available at: <http://rvsn.csail.mit.edu/eolson/loopclosing>.

### 4. Related Work

Loop-consistencies were used to validate hypotheses in (Bosse, 2004), but without the batch outlier rejection presented here.



Figure 4. Map of CSAIL G7. Our algorithm enables high-quality maps of environments without human assistance.



Figure 5. Intel research center. A standard robotics data set before and after automatic loop closures.

Random Sample Consensus (RANSAC) could be employed in this domain. RANSAC repeatedly selects a hypothesis and counts the number of other hypotheses that are consistent with the model, using a user-tuned consistency threshold. The hypothesis with the most votes “wins”. This is equivalent to picking only one “true” hypothesis by picking the one with the greatest row-sum in matrix  $A$ .

In contrast, our approach returns a set of multiple “true” hypotheses (each hypothesis with a relative confidence indicator), and considers the entire network of pairwise consistencies (rather than just one row at a time) in order to produce a better answer. Despite these advantages, our approach has the same run-time complexity as RANSAC.

### References

- Bosse, M. C. (2004). *ATLAS: a framework for large scale automated mapping and localization*. Doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA.
- Olson, E., Leonard, J., & Teller, S. (2006). Fast iterative optimization of pose graphs with poor initial estimates. *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2262–2269). Orlando, FL, USA.

Olson, E., Walter, M., Leonard, J., & Teller, S. (2005). Single cluster graph partitioning for robotics applications. *Proceedings of Robotics Science and Systems* (pp. 265–272).