

Loop Closing in Topological Maps

Kristopher R. Beevers and Wesley H. Huang
Rensselaer Polytechnic Institute, Department of Computer Science
110 8th Street, Troy, New York 12180, U.S.A.
{beevek,whuang}@cs.rpi.edu

Abstract—In order to create consistent maps of unknown environments, a robot must be able to recognize when it has returned to a previously visited place. In this paper, we introduce an evidential approach to the loop-closing problem for topological maps, based on the Dempster-Shafer theory of evidence. In our approach, the robot makes a hypothesis whenever it may have revisited a place. It then attempts to verify hypotheses by continuing to traverse the environment, gathering evidence that supports (or refutes) the hypotheses. We describe methods for managing belief about multiple loop-closing hypotheses, and for determining a belief assignment given a piece of evidence. We also discuss methods for reducing the false alarm rate of our loop-closing algorithm, and provide simulated and real-world experimental results that verify the effectiveness of our approach.

I. INTRODUCTION

An important problem in mobile robotics is for an individual robot to enter an unknown environment and create a map that can subsequently be used for navigation. One approach to this problem is to create topological maps, generally represented as graphs where the nodes correspond to “places” in the environment and the edges are paths between two places.

The fundamental challenge in creating topological maps is for the robot to recognize when it has returned to a place it has previously visited. This problem is known as *closing the loop* since the robot usually returns to a previously visited place via a different path. Without this capability, a robot will have several nodes in its map that represent the same place, making the map inconsistent.

With sufficiently rich sensing, places can be recognized by recording a unique “sensing signature” [7] for each place. Our focus, however, has been on robots with much more limited sensing information. We envision applications where, for example, an inexpensive (and therefore potentially disposable) robot is sent into a contaminated building to perform hazard assessment. In order to make this robot inexpensive, a limited array of sensors will be available.

In this paper, we describe a new method for closing loops in topological maps that relies only upon odometry measurements. We have applied this method to a small robot that uses wall-following and other behaviors to travel through the environment. The nodes in our map are interior and exterior corners of walls; the edges represent a sequence of behaviors to move the robot from one node to another. The robot measures the distance between nodes (with error) using odometry, and the single “feature” of a node is whether it is an exterior or interior corner.

Our method maintains an estimate of the robot’s position with respect to its starting node. Using a model of odometry error, we compute confidence bounds on this estimate. When these bounds encompass a previously-visited node, the robot makes a hypothesis that it has closed the loop. The robot then continues its forward traversal of the loop, looking for evidence to confirm or reject a hypothesis. A hypothesis may be rejected by a “structural mismatch,” e.g., an interior corner encountered where there should be an exterior corner.

Otherwise, the length of each edge traversed is compared with the previous measurement of the corresponding edge under the hypothesis. The degree to which a new measurement matches the old measurements is interpreted as evidence for or against a loop-closing hypothesis. We use Dempster-Shafer theory [8] to combine evidence and maintain the current belief in each of the loop-closing hypotheses. When one hypothesis has garnered enough support, that hypothesis is accepted as correct.

One key element of our approach is a new method for modifying a Dempster-Shafer frame of discernment to accommodate discovery of new hypotheses. Another aspect is a belief function we have devised to compute a basic probability assignment that reflects the evidence a path-length measurement provides over all sets of hypotheses.

In the remainder of this section, we review related work and our assumptions. Then, after a brief overview of Dempster-Shafer theory in Section II, we describe the details of our method in Section III. Simulated and real-world experiments are presented in Section IV.

In the course of this work, we have found that environments with structural and metric self-similarity — particularly environments that “spiral” inward or outward and environments with repeating substructure — are the most difficult to map correctly. In these worlds, incorrect loop-closing hypotheses are often formed. This leads to “false alarms” — incorrect hypotheses being confirmed as correct because of the self-similarity. In Section V, we discuss methods for limiting the rate of loop-closing false alarms by measuring the self-similarity of an environment and responding accordingly.

A. Related work

The loop-closing problem has been studied in conjunction with mapping for well over two decades; however, we believe this work is the first to place this problem in a decision theoretic framework.

A simple (but undesirable) approach to closing the loop is to “drop a pebble.” Upon returning to the pebble, the robot knows that it has returned to the same location. Bender *et al.* [2] describe algorithms that use a pebble in order to explore graphs. Another approach is to use sensors that can record a distinctive “sensing signature” for each place. Kuipers and Beeson [6] use supervised learning to recognize the sensing signatures at nodes in their topological maps.

Kuipers’ “rehearsal procedure” [7] encapsulates the general idea of using the map topology to make a loop-closing decision. Choset and Nagatani [3], whose topological maps are based on the generalized Voronoi diagram of the environment, describe an approach where structural characteristics of the map (e.g., the degree of vertices and the order of incident edges) are the primary criteria for verification. Tomatis *et al.* [10] embed this comparison in a POMDP that should show a single peak when the loop has been closed.

Our approach, based on accumulating evidence about loop-closing hypotheses, is related to that of Cox and Leonard [4], who maintain multiple hypotheses about the state of a dynamic world. They assign and update the probability of each hypothesis using a Bayesian framework. The main advantage of a Dempster-Shafer based method is the ability to represent “ignorance” about which hypothesis is supported by a piece of evidence — particularly useful when, for example, evidence appears to support more than one hypothesis.

This work focuses on one specific aspect of the single-robot mapping problem (loop closing). We omit many of the details of the rest of our mapping approach. A complete discussion of the remainder of our mapping algorithm is available in our previous publications [1], [5], which include details about behaviors and our other extensions to topological maps (mostly for dealing with open spaces).

B. Assumptions

We consider a robot with known error models for movement, odometry and sensing. We assume errors are random and zero-mean and that we can merge and compound measurements and compute confidence bounds for a given confidence level. For this work, we assume that the robot perfectly detects features (i.e., there are no structural errors in the robot’s map — only metric errors). Additionally, we assume an enclosed, static environment.

II. DEMPSTER-SHAFER THEORY

Our loop-closing algorithm makes extensive use of the Dempster-Shafer theory of evidence [8]. As such, a brief overview of the central concepts of Dempster-Shafer theory is warranted.

Dempster-Shafer theory provides a framework for the mathematical representation of uncertainty that is based on modeling “belief” about a set of possibilities. Dempster-Shafer differs from traditional probability theory in several key ways:

- The allocation of belief mass to *sets* of mutually exclusive possibilities — not just individual possibilities — is allowed.
- For this reason, Dempster-Shafer theory can represent “ignorance,” as belief mass assigned to a set of multiple possibilities (reflecting lack of knowledge about which specific possibility a piece of evidence supports).
- Because of its ability to represent ignorance, Dempster-Shafer theory requires no *a priori* knowledge about the world; in the presence of no knowledge, all belief mass is assigned to ignorance.

In Dempster-Shafer theory, the set of mutually exclusive possibilities (or events), called the *frame of discernment*, is denoted Θ . In our loop-closing problem, the elements in the frame are hypotheses about loops in the environment. In Dempster-Shafer theory, probability is assigned over the *power set* of Θ , 2^Θ — rather than over Θ itself, as it would be in traditional probability theory. A *basic probability assignment* (b.p.a.) $m : 2^\Theta \rightarrow [0, 1]$ is a function for which $m(\emptyset) = 0$ and $\sum_{A \subseteq \Theta} m(A) = 1$. The quantity $m(A)$ is the basic probability assigned to A . This is the measure of belief committed exactly to A — but not the total belief committed to A . The total belief in A is the sum of all belief committed exactly to A and its subsets, i.e.

$$\text{Bel}(A) = \sum_{B \subseteq A} m(B) \quad (1)$$

The *plausibility* of A is defined as

$$\text{Pl}(A) = 1 - \text{Bel}(\overline{A}) \quad (2)$$

Plausibility is the probability mass that does not support A ’s negation.

It is useful to combine two b.p.a.’s into a single b.p.a. that reflects our belief given all of the evidence represented in each. Dempster-Shafer theory provides a mechanism for this, called *Dempster’s rule of combination*. We combine two b.p.a.’s m_1 and m_2 with:

$$m_1 \oplus m_2(A) = \frac{\sum_{B, C \in 2^\Theta: B \cap C = A} m_1(B)m_2(C)}{1 - \sum_{B, C \in 2^\Theta: B \cap C = \emptyset} m_1(B)m_2(C)} \quad (3)$$

(unless $A = \emptyset$, in which case $m_1 \oplus m_2(\emptyset) = 0$). The denominator in Dempster’s rule is a normalization coefficient; if it is zero, there is total conflict between the two b.p.a.’s.

III. ALGORITHM OVERVIEW

When the robot detects it has potentially returned to the start vertex, it adds a new hypothesis to its frame of discernment. As the robot continues to circumnavigate, each path-length measurement gives evidence about hypotheses that have been made so far. A belief function maps the evidence into a basic probability assignment over 2^Θ , which is merged with the global b.p.a.. When sufficient belief is concentrated in a single hypothesis, that hypothesis is accepted as the correct one.

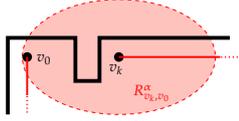


Fig. 1. A situation in which we hypothesize that $v_k = v_0$. The region R_{v_k, v_0}^α represents the positional uncertainty in the location of v_k with respect to v_0 . Since this region contains v_0 , it is possible that the two vertices represent the same feature.

A. Hypotheses

For each vertex v_i in the map, we use the model of the robot’s odometry error to maintain a probability density function U_{v_i, v_0} describing the robot’s location with respect to the start vertex v_0 . The uncertainty in this distribution grows monotonically as the map expands. For a confidence level α specified to our mapping algorithm, let R_{v_i, v_0}^α represent confidence bounds of U_{v_i, v_0} .

If the robot reaches a vertex v_k that is structurally the same as v_0 and R_{v_k, v_0}^α contains v_0 , then we create a new hypothesis $H_k \equiv v_k = v_0$ (i.e. we hypothesize that v_k closes the loop). For an example of such a situation, see Figure 1. The robot then continues to traverse the environment, attempting to prove H_k .

Belief in the correctness of hypotheses is maintained according to Dempster-Shafer theory over a frame of discernment where each hypothesis constitutes a singleton. This formulation reflects the fact that loop-closing hypotheses are not independent. A singleton representing “none of the current hypotheses” (N) is also included in the frame of discernment.

As an example: when a single loop-closing hypothesis H_1 has been made, the frame of discernment is $\Theta = \{H_1, N\}$ and b.p.a.’s are over the set $2^\Theta = \{\emptyset, \{H_1\}, \{N\}, \{H_1, N\}\}$. When two loop-closing hypotheses have been made, b.p.a.’s are made over the set

$$2^\Theta = \{\emptyset, \{H_1\}, \{H_2\}, \{N\}, \{H_1, H_2\}, \{H_1, N\}, \{H_2, N\}, \{H_1, H_2, N\}\} \quad (4)$$

Unless disproven by a structural mismatch encountered as the robot traverses, no hypothesis is ever removed from the frame of discernment until some hypothesis has been deemed correct.

B. Expanding the frame of discernment

When a new loop-closing hypothesis is made, the frame of discernment must be updated to include it, and the belief from the previous frame must be reassigned as the initial belief for the new frame. An existing frame of discernment Θ_{k-1} is updated when a new hypothesis H_k is discovered by adding H_k to the frame, i.e. $\Theta_k = \Theta_{k-1} \cup \{H_k\}$.

The b.p.a. over $2^{\Theta_{k-1}}$ cannot be transferred directly to the corresponding members of 2^{Θ_k} . Doing so would require that elements of 2^{Θ_k} that contain H_k be assigned zero probability mass, and Dempster’s rule would never assign probability mass to any element supporting H_k , regardless of the evidence found in the future supporting H_k .

Basic probabilities are reassigned in the new frame as follows: for all elements $\{A \in 2^{\Theta_{k-1}} | N \subseteq A\}$, evidence is reassigned into 2^{Θ_k} as:

$$m^k(A \cup \{H_k\}) = m^{k-1}(A) \quad (5)$$

So, any elements from the previous b.p.a. containing N (“none of the current hypotheses”) have their evidence reassigned to the element in the new b.p.a. consisting of the union of H_k and the previous element. Note that this also makes intuitive sense: evidence representing none of the previously-made hypotheses may, in fact, have been in support of the new hypothesis. This operation preserves the plausibility of all hypotheses (and N); the plausibility of H_k initially becomes the same as the plausibility of N . The belief in each element of $2^{\Theta_{k-1}}$ of which N is a subset is not preserved; for all other elements, it is preserved. (Belief in N itself is reassigned to $\{N, H_k\}$.)

The initial belief in H_k after this operation is zero. However, depending on the mapping strategy being used, we may have some initial evidence in support of H_k . Our particular mapping strategy defines a ball $B_{v_0}^{r^*}$ of nonzero radius around the start vertex within which no other structurally similar vertex can exist (see [5] for complete details). So, any evidence indicating the robot is within this radius of v_0 supports the new hypothesis. We compute a simple b.p.a. in which

$$m(H_k) = \iint_{B_{v_0}^{r^*}} U_{v_0, v_k}(x, y) dA \quad (6)$$

and the remainder of the probability mass is assigned to complete ignorance. Using Dempster’s rule, this b.p.a. is combined with our global b.p.a. to reflect the initial belief in H_k . Forms of initial evidence about H_k that are specific to particular mapping strategies may be integrated likewise.

C. Belief function

When a new path-length measurement is taken, it gives evidence about the hypotheses that have been made so far. If we have made k loop-closing hypotheses so far, there are k sets of measurements $L_1 \dots L_k$, one associated with each hypothesis, of previous measurements that should match the new measurement under the respective hypotheses. Based on these sets of measurements, we must assign basic probability to each element of 2^Θ . This b.p.a. constitutes our belief function given the new measurement. Once we have made this b.p.a., we combine it with our previous belief using Dempster’s rule. Note that our belief function is responsible for determining the amount of “ignorance” about which hypotheses are supported by a piece of evidence — Dempster’s rule does not do this for us.

Our particular belief function draws on methods of statistical inference. It essentially computes probabilities that the measurements in each L_i were taken from the same wall, given the robot’s error model, and then uses these probabilities to determine a b.p.a. over 2^Θ .

For each set L_i of measurements, we first compute a goodness-of-fit value based on the z -score of the measurements in L_i . (The z -score represents the deviation

of a sample from the expected mean, expressed in units of the standard deviation of the underlying distribution.) We use the square of the z -score, computed using a distribution with standard deviation $\hat{\sigma}$ generated according to the robot's error model for the current "best estimate" $\hat{\ell}$ (based on L_i) of the length of the wall:

$$z^2 = \sum_{\ell_j \in L_i} (\ell_j - \hat{\ell})^2 / \hat{\sigma}^2 \quad (7)$$

The value of z^2 follows a χ^2 distribution with $n = |L_i|$ degrees of freedom. So, the degree of support provided by L_i for H_i is computed as:

$$\Phi_i = \int_{z^2}^{\infty} \frac{y^{n/2-1} e^{-y/2}}{2^{n/2} \Gamma(n/2)} dy \quad (8)$$

(This is simply the χ^2 distribution with n degrees of freedom.) Φ_i is thus a basic probability measuring the likelihood that the measurements in L_i come from the same distribution. At the end of this step, we have a Φ_i for each L_i . We use the Φ_i 's to determine the b.p.a. over 2^Θ .

First, we find a basic probability to assign to N , the proposition that none of the hypotheses is correct. We label the basic probability associated with N as Φ_0 for convenience. Let $\Phi_0 = 1 - \max_{i>0} \Phi_i$. This value reflects the fact that the extent to which none of the hypotheses are supported by a piece of evidence is limited by the maximum support for any one hypothesis.

We then normalize the Φ_i 's (including Φ_0) so that they sum to one (i.e., so that they constitute a valid b.p.a.): $\Phi_i = \Phi_i / \sum_j \Phi_j$. Finally, we compute the b.p.a. for 2^Θ using the following procedure:

Algorithm 1 COMPUTE-BPA:

```

for  $j = |\Theta| \dots 2$  do
  Let  $\Lambda = \{A \in 2^\Theta \mid |A| = j\}$ 
  for all  $A \in \Lambda$  do // compute PIC
     $p_A \leftarrow \sum_i \Phi_i \mid_{H_i \in A}$ 
     $\eta_A \leftarrow 1 + \frac{\sum_i \Phi_i \log \frac{\Phi_i}{p_A}}{\log |A|}$ 
  for all  $H_i \in \Theta$  do // compute normalization constant
     $t_i \leftarrow \max \left( 1, \sum_{A \in \Lambda \mid H_i \in A} \eta_A \right)$ 
  for all  $A \in \Lambda$  do // compute b.p.a.
     $m(A) \leftarrow \sum_i \Phi_i \mid_{H_i \in A} \frac{1 - \eta_A}{t_i}$ 
  for all  $H_i \in \Theta$  do // "bleed off" probability mass
     $\Phi_i \leftarrow t_i \Phi_i$ 
for all  $H_i \in A \in 2^\Theta \mid |A| = 1$  do
   $m(A) = \Phi_i$  // assign remaining mass to singletons

```

The value $\eta_A \in [0, 1]$ is the "probability information content" [9] of the (renormalized) Φ_i 's for the hypotheses in element $A \in 2^\Theta$. This essentially measures the uniformity of the Φ_i 's for these hypotheses. By using the probability information content in this way, we profess more ignorance when evidence supports multiple hypotheses, or when it is unclear whether evidence supports or does not support a hypothesis.

This procedure, in essence, computes the b.p.a. by determining the amount of ignorance about each element in set A . In the first iteration, "overall" ignorance is determined (because $A = \Theta$). In later iterations (in which the b.p.a.'s for lower-cardinality elements of 2^Θ are computed), ignorance about the remaining evidence stored in the Φ_i 's belonging to each set is computed. By the time the procedure arrives at the singleton elements, the remaining probability mass is that which should be attributed directly to specific loop-closing hypotheses (or to N). The computations involving the renormalization constant t_i are necessary because without it, it is possible for more than the total probability mass remaining in a Φ_i to be "bled off" in a single iteration of the outer loop.

Note that this belief function meets the requirements stated in Section II.

D. Decision making

With our methods for modifying the frame of discernment and for determining a b.p.a. for a piece of evidence, we have all the tools necessary for accumulating belief about loop-closing hypotheses. As the robot continues to traverse the environment, belief about each new path-length measurement is combined with our global belief. Whenever a new loop-closing hypothesis is made, it is added to the frame of discernment. However, we must also provide a mechanism for making decisions about when to accept a hypothesis as the correct one.

A simple way to do this is to specify a threshold β to the algorithm. When the belief in any single hypothesis exceeds β , that hypothesis is accepted. In fact, this simple strategy works well in most cases — particularly in real-world situations — because after taking several new measurements, the belief in the correct hypothesis converges quickly to one. The difficulty arises when an incorrect loop-closing hypothesis is made, and it is not obvious after several measurements that it is a wrong hypothesis.

These kinds of hypotheses occur when the world is both structurally and metrically self-similar. Two ways in which a world might be self-similar are when the world: (1) consists of paths that spiral inward or outward; or (2) consists entirely or partially of repeating similar sequences of paths and features. Our evidential approach fares well in "spiral" worlds (by their nature, these worlds yield comparisons between different-length edges). Worlds with repeating subsequences are more difficult; we discuss strategies for effective loop-closing in such worlds in Section V.

IV. EXPERIMENTAL RESULTS

We have implemented our algorithm in simulation and on a mobile robot (Figure 2) that uses the behavior-based mapping strategy detailed in [5]. Simulated experiments were performed on a variety of environments, including simple worlds and others crafted to present hard loop-closing scenarios. Real-world experiments were performed in both hand-made and unmodified building environments. All experiments were performed using a fixed hypothesis-acceptance threshold requiring that belief in a hypothesis

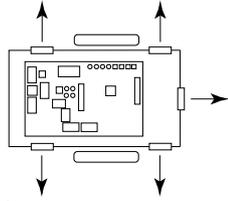
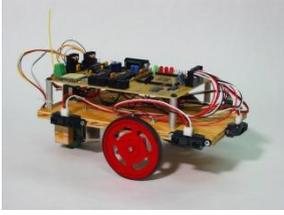


Fig. 2. Our prototype robot is a differential drive mobile robot approximately 20 cm long. It has five Sharp GP2D12 infrared range sensors, 256 CPR encoders on each wheel, and an Atmel ATMEGA64 microcontroller as the main processor.

$\sigma\%$	w1	w2	w4	w5	w9	w11	w12
4	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0
6	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0
8	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0	9/1/0
10	10/0/0	10/0/0	10/0/0	10/0/0	9/0/1	10/0/0	9/1/0
12	9/0/1	10/0/0	10/0/0	10/0/0	10/0/0	10/0/0	9/1/0
14	10/0/0	10/0/0	10/0/0	10/0/0	9/0/1	10/0/0	10/0/0
16	10/0/0	9/0/1	10/0/0	10/0/0	10/0/0	9/0/1	8/2/0

Fig. 3. Results of experiments performed in simulated environments, listed in order of complexity (self-similarity and size): w1 is the simplest and w12 the most complex. Each entry is of the form: “correct loop closings / incorrect hypotheses accepted / correct hypothesis not made.” The error model used in testing assumed standard deviation directly proportional to the length of a measurement; $\sigma\%$ is thus the percentage of the measurement length that was taken to be the standard deviation. For each world and each error percentage, ten simulations were run.

exceed 0.99. The robot used a confidence limit of 0.99 to compute confidence bounds for making hypotheses.

More than 700 simulations were performed in seven different simulated environments under various error conditions. Figure 3 presents the results. Each environment had a single loop; the robot either closed the loop correctly, accepted an incorrect hypothesis, or failed to make the correct loop-closing hypothesis. In 98.6% of these simulations, loops in the environments were closed correctly.

Of the failures, half occurred in a world specifically designed to defeat the algorithm by presenting an especially challenging degree of self-similarity (Figure 4). These failures were due to the acceptance of an incorrect hypothesis prior to the discovery of the correct one. (In our previous work, this world defeated our original loop-closing algorithm more than 44% of the time; the new algorithm failed only 7% of the time.) All other failures occurred

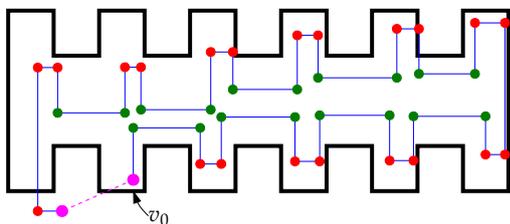


Fig. 4. A highly self-similar environment (w12) that sometimes fools our loop-closing algorithm under certain error conditions. The start vertex (v_0) is in the “well” that is second from left on the bottom (the worst case scenario for this environment). By the time the robot circumnavigates the environment to reach the leftmost well, it mistakes this for the start (due to error) and generates an incorrect hypothesis, shown as a dashed line between the matched vertices. Because of the self-similarity, this hypothesis is occasionally confirmed within four edge traversals, before the correct hypothesis is made.

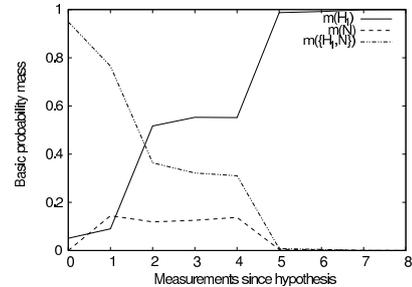
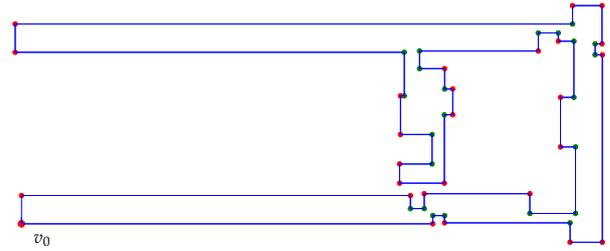


Fig. 5. Basic map, after loop-closing, of the corridors on the first floor of an academic building (Amos Eaton) at RPI, constructed using a wall-following behavior. Six measurements were needed to verify the correct loop-closing hypothesis, which was the only loop-closing hypothesis made (the start vertex is the large dot at bottom left). The convergence of the global belief function is shown in the plot.

when the correct hypothesis was not made because the start vertex fell outside the confidence bounds on the robot’s location. Note that this occurred less than 1% of the time — consistent with our choice of a 0.99 confidence limit for making hypotheses.

Our robots made real-world maps of several different environments of varying size. The largest map, shown in Figure 5, is of the first floor hallways of a 12 m \times 30 m academic building. In every real-world test, the robot closed the loops in its map correctly.

V. DISCUSSION

Continuing the forward traversal of the environment after making a loop-closing hypothesis presents several issues. An important concern is that one can never be completely certain that a hypothesis is correct. Nevertheless, we believe that the idea of building evidence in support of and against a hypothesis is both intuitive and reasonable for use in practical situations.

A. Repeating subsequences

Worlds with self-similarity in the form of repeating subsequences (such as that in Figure 4) provide a difficult challenge. In these types of worlds, an incorrect loop-closing hypothesis is sometimes accepted as correct before the correct hypothesis is even made. The situation in which an incorrect hypothesis is accepted is a “false alarm,” and the frequency at which this occurs is the *false alarm rate*. (Note that the frequency at which the correct hypothesis is made is the *probability of detection*, and is generally equal to the confidence limit α specified to the algorithm.)

The simplest way to reduce the false alarm rate is to increase the acceptance threshold β , requiring more evidence to accept any hypothesis. In order to guide the choice

of β , we can consult measures of the self-similarity of the environment. For example, we might compute the entropy of the structure and path lengths in the environment. The lower the entropy, the more self-similar the environment, and the higher we should make β in order to reduce the false alarm rate. Of course, we can only base our self-similarity measures on what has been encountered in the environment so far. We make the assumption that no subset of the environment is significantly more complex than the average complexity of the whole — i.e. that the environment is relatively uniform. This assumption has proven realistic in most practical scenarios.

A key advantage of our loop-closing formulation is the ability to represent ignorance about which hypotheses are supported by a piece of evidence. This means that when evidence matches several hypotheses equally, it affects our belief as we would expect — our direct commitment of belief to the hypothesis singletons is small because most of our belief is assigned to the set containing all of the closely-matching hypotheses. This is an important tool for closing loops in environments with repeating substructure in which incorrect hypotheses occur frequently, and often appear to be good matches initially.

B. Detecting “*n*th-lap” hypotheses

One tricky implementation detail merits brief discussion. In very simple environments, it is sometimes the case that the robot creates the correct hypothesis, but does not acquire enough evidence to accept it before it again returns to the start vertex and creates another, separate hypothesis. This can lead to problems because, for the two hypotheses, the world appears completely self-similar. We term this an “*n*th-lap” situation, because a new hypothesis is created like this on every “lap” around the loop to be closed. In practical situations, *n*th-lap hypotheses are typically not an issue, since most environments are extensive enough that loop-closing decisions are made during the second lap.

In general, detecting such a situation is reasonably straightforward. The number of paths that must be explored to make an *n*th-lap hypothesis must be a multiple of the number of paths explored before making the first-lap hypothesis. If two hypotheses H_a and H_b meet this criteria, and if $\eta_{\{H_a, H_b\}}$ is frequently close to zero (evidence frequently supports both H_a and H_b), the hypotheses are likely *n*th-lap hypotheses.

Upon detecting an *n*th-lap situation, the first of the two hypotheses should be assigned the belief in the second, and the second hypothesis should be removed from the frame of discernment.

VI. CONCLUSIONS

In this paper, we have presented an evidential approach to the problem of “closing the loop” in a topological map — recognizing when the robot has returned to a node it has already visited. Though our loop-closing algorithm was discussed primarily in the framework of a simple “circumnavigation” mapping strategy, it can be easily extended to different mapping scenarios. Our approach is based on the

Dempster-Shafer theory of evidence, which allows us to express “ignorance” about which loop-closing hypothesis is supported by a piece of evidence. This capability provides important benefits when faced with difficult loop-closing problems, particularly in highly structurally and metrically self-similar environments.

Our loop closing approach modifies the Dempster-Shafer frame of discernment whenever a new hypothesis is discovered. It also introduces a belief function that makes a basic probability assignment reflecting our belief in each hypothesis, given a certain piece of evidence. This function determines the “conflict” inherent in evidence and assigns ignorance accordingly.

We have also discussed methods for reducing the false alarm rate of our loop-closing algorithm in difficult worlds, based on measures of the self-similarity of the environment. Our approach shows promise in this respect, but more work must be done to clearly identify the relationship between self-similarity measures and false alarm rate. In particular, we are interested in providing guarantees on the false alarm rate, depending on the self-similarity of the environment.

Experiments show that our loop-closing algorithm is effective, even in difficult worlds. In simulated tests, the algorithm closed loops correctly more than 98% of the time — and more than 99% of the time in environments not specifically designed to defeat it. In real-world tests, the correct loop-closing hypothesis was chosen in every case.

ACKNOWLEDGMENT

This work was supported by the NSF through award IIS-9983642.

REFERENCES

- [1] K. R. Beevers. Topological mapping and map merging with sensing-limited robots. Master’s thesis, Rensselaer Polytechnic Institute, Troy, NY, April 2004.
- [2] M. Bender, A. Fernandez, D. Ron, A. Sahai, and S. Vadhan. The power of a pebble: Exploring and mapping directed graphs. In *Proc. Annual Symp. on Foundations of Computer Science*, pages 269–278, 1998.
- [3] H. Choset and K. Nagatani. Topological simultaneous localization and mapping (SLAM): Toward exact localization without explicit localization. *IEEE Transactions on Robotics & Automation*, 17(2):125–137, April 2001.
- [4] I. Cox and J. Leonard. Modeling a dynamic environment using a Bayesian multiple hypothesis approach. *Artificial Intelligence*, 66:311–344, 1994.
- [5] W. H. Huang and K. R. Beevers. Topological mapping with sensing-limited robots. In *6th International Workshop on the Algorithmic Foundations of Robotics (WAFR 2004)*, pages 367–382, 2004.
- [6] B. Kuipers and P. Beeson. Bootstrap learning for place recognition. In *Proc. 18th Natl. Conf. on Artificial Intelligence*, pages 174–180, Edmonton, Canada, 2002.
- [7] B. Kuipers and Y.-T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47–63, 1991.
- [8] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [9] J. Sudano. Pignistic probability transforms for mixes of low- and high-probability events. In *Proc. 4th International Conference on Information Fusion*, pages 23–27, Montreal, August 2001.
- [10] N. Tomatis, I. Nourbakhsh, and R. Siegwart. Hybrid simultaneous localization and map building: closing the loop with multi-hypothesis tracking. In *Proc. 2002 IEEE Intl. Conf. on Robotics & Automation*, May 2002.