

Homological Sensing for Mobile Robot Localization[†]

Jason Derenick¹, Alberto Speranzon² and Robert Ghrist³

Abstract—In this paper, we consider a minimalistic approach to mobile robot localization that constrains the robot’s ability to sense its environment to a binary detection of uniquely identifiable landmarks having unknown position (e.g., a commodity WiFi transceiver detecting network SSIDs). Central to the proposed solution are dual landmark and observation complexes (*simplicial nerve complexes*), which can be iteratively built through local observations without any metric or time-sequenced information. We have shown that these complexes approximate the topology of the underlying physical environment. Specifically, the notion of a “hole” within these complexes naturally represents a physical structure (e.g., a building) that limits landmark visibility/communication with respect to the robot’s location. Taking advantage of this property, we formulate a *homological sensing model* that operates on these constructs enabling the robot to “count” the number of structures in its vicinity using local homology computations as a metric-free surrogate sensor. Our homological sensor model is presented in the context of a Monte-Carlo localization algorithm that resolves robot location by correlating the measured number of holes with an unlabeled, metric map location.

I. INTRODUCTION

Imagine an autonomous robot, equipped with only a WiFi transceiver and an IMU¹, traveling through an urban canyon. Although the robot is given a map revealing building locations (perhaps via Google Maps), it cannot identify these structures nor any other potential anchors within the map. The robot has only available its IMU measurements, its unlabeled map, and the location-free SSIDs that it has seen with its transceiver (in a binary fashion) throughout the duration of its travels. Given such limited sensing information, can the robot recover its location within the map?

Key to enabling localization (albeit coarse) in this context is the realization that the assumed sensor model, although primitive, still contains enough information to generate a topological map (in the form of a simplicial nerve complex) of the robot’s workspace. Since this metric-free map is

constructed from a visibility-based sensor, under sufficient observation density, its topological invariants (e.g., holes) naturally correspond to physical structures in the robot’s workspace that limit its visibility (or, alternatively, communications) to landmarks. Taking advantage of this feature, we formulate a localization pipeline in which the agent first uses its binary sensing capabilities to build a simplicial representation of its environment to enable local homology computation revealing the number of proximal holes. In this case, we define proximity via a coarse pseudo-metric. Using the number of holes as a surrogate for the number of physical structures nearby, the agent can correlate its measurement with an unlabeled metric map thereby closing the loop.

It should be also noted that, in the context of RF-based localization, quite a few algorithms have been developed and demonstrated in the literature that use ranging and/or bearing, see [1], [2], [3] and references therein. However, to use such information, an estimate of the range and/or angle of arrival needs to be determined from measurements of, e.g., power, time of flight, phase difference, etc. For such measurements, statistical models can be very complex as they depend on the level of interference from other transmitters, materials within the environment, type and number of obstacles between receiver and transmitter, etc. The approach taken in this paper becomes very relevant when such statistical models are difficult to obtain. Typical examples are those corresponding to indoor or urban environments.

Although the robot has a map of its physical workspace, the position of the landmarks it uses to build its *landmark complex* is not known to the robot (and thus they are not provided in the map). Additionally, despite knowing that holes in its landmark complex correspond to physical structures in its environment, the robot cannot uniquely identify these holes and it has no knowledge of which hole maps to which structure. It can only enumerate these topological features in its local landmark homology.

II. RELATED WORK

The tools employed herein are rooted in algebraic topology and have, but recently, gained attention from the robotics community. Seminal results have used such tools in various application contexts that include coverage verification [4], [5], [6], [7], motion planning [8], [9], [10], [11], [12], metric-free target tracking and enumeration [13], and, even, the characterization of robot gaits [14]. Topological localization via signals of opportunity is considered in [15], although in a different context when compared to this paper. Specifically, it does not address the case of a moving robot and the associated (dynamic) filtering problem to estimate its

[†] The authors gratefully acknowledge support from DARPA, contract number FA8650-11-C-7139, and the United Technology Research Center. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. Approved for Public Release, Distribution Unlimited.

¹J. Derenick is a Senior Research Scientist in the Autonomous and Intelligent Robotics Laboratory, Systems Department, United Technologies Research Center, Hartford, CT, 06108, USA. Email: jason.derenick@utrc.utc.com

²A. Speranzon is a Staff Research Scientist in the Autonomous and Intelligent Robotics Laboratory, Systems Department, United Technologies Research Center, Hartford, CT, 06108, USA. Email: alberto.speranzon@utrc.utc.com

³Robert Ghrist is the Andrea Mitchell University Professor of Mathematics and Electrical/Systems Engineering at the University of Pennsylvania, Philadelphia, PA, 19104, USA. Email: ghrist@seas.upenn.edu

¹The use of an IMU can be eliminated provided the agent has a suitable motion model for the prediction update.

position. Others have also considered topological algorithms for estimation. However, in these works, topology is largely defined as a graphical embedding representing environmental free space. Recent results in this area include [16], [17], [18].

Due to the coarseness of the available information, in this paper we will rely on Monte-Carlo based filtering. These methods are a standard tool for non-parametric, state-based estimation that have been widely adopted by the robotics community. They appear with various levels of sophistication [19], [20], [21]. Additionally, localization using low-grade, commodity sensing modalities (e.g., WiFi, acoustics, depth cameras, etc.) has also been studied in a variety of contexts [22],[23],[24],[25]; however, central to many such approaches is a precise calibration procedure for correlating the signal of interest with some metric quantity. Finally, results with visibility-based simultaneous localization and mapping (SLAM) [26], [27] (and references therein) have a natural kinship with this research.

III. MATHEMATICAL BACKGROUND

Given the centrality of algebraic topology to the forthcoming results, we provide a brief index of key definitions and constructs. The reader is referred to [28] for a comprehensive treatment of the subject.

A. Simplices, Simplicial Complexes and Nerves

We make liberal use of simplicial complexes. An *abstract simplicial complex* based on a vertex set V is a collection \mathcal{K} of finite subsets of V (called *simplices* closed under the operation of taking subsets: if $\sigma \in \mathcal{K}$ and $\tau \subset \sigma$, then $\tau \in \mathcal{K}$ as well. In this case, τ is called a *face* of σ . The dimension of a simplex is given by $\dim \sigma = |\sigma| - 1$: e.g., a 2-simplex is an unordered 3-tuple in V , and a k -simplex is an unordered $(k+1)$ -tuple in V . The k -skeleton of a complex is defined as the union of all simplices of dimension k and below: $\mathcal{K}^{(k)} := \{\sigma \in \mathcal{K} : \dim \sigma \leq k\}$.

As with the more familiar case of graphs (1-skeleta), abstract simplicial complexes always possess a geometric realization, obtained by gluing together geometric simplices [28]: the illustrations in this paper represent simplicial complexes in this manner. This furthermore permits thinking of simplicial complexes as topological spaces, complete with topological invariants.

We make particular use of simplicial complexes that are *nerves* associated to a cover.

Nerve Complex: Given a collection of sets $\mathcal{U} = \{\mathcal{U}_\alpha\}$, the *nerve complex*, $\mathcal{N}_\mathcal{U}$, is the abstract simplicial complex whose k -simplices are unordered collections of $k+1$ elements of \mathcal{U} having non-empty intersection.

By the Nerve Theorem [29], when \mathcal{U} consists of, e.g., convex sets in Euclidean space, $\mathcal{N}_\mathcal{U}$ is homotopic (topologically equivalent) to the union $\cup_\alpha \mathcal{U}_\alpha$. Convexity is sufficient but not necessary.

B. Homology Groups

Homology provides an algebraic means to characterize the holes in a topological space. At its core lies the notion of a boundary homomorphism, which for simplicial homology encodes how simplices are attached to their lower-dimensional faces. To define this for a complex \mathcal{K} , we choose an ordering (in the manner of directed graphs) for each simplex. An ordering on a k -simplex, $\sigma \subset V$, is a literal ordering of its vertices $[v_0, \dots, v_k]$, up to the equivalence generated by even permutations. Given such a choice, consider the \mathbb{R} -vector space $C_k(\mathcal{K})$ with basis the oriented k -simplices in \mathcal{K} . Multiplication by -1 reverses the orientation. A *boundary homomorphism* is defined as the linear mapping $\partial_k: C_k(\mathcal{K}) \rightarrow C_{k-1}(\mathcal{K})$ given by sending each basis element of $C_k(\mathcal{K})$ to the formal sum of its (oriented) faces of dimension $k-1$.

With the boundary operator $\partial = \{\partial_k\}$ encoding the assembly instructions of \mathcal{K} , certain global topological features can be deduced from the resulting linear algebra. The k^{th} *homology group* of \mathcal{K} , $H_k(\mathcal{K})$, is, in this setting,

$$H_k(\mathcal{K}) = \ker \partial_k / \text{im } \partial_{k+1}. \quad (1)$$

This is the quotient vector space whose generators are k -cycles (corresponding to k -dimensional boundaryless subcomplexes surrounding a hole) modulo the equivalence relation that says two such k -cycles are the same (or *homologous*) whenever they are the (oriented) boundary of a $(k+1)$ -dimensional subcomplex. Among the key properties of $H_k(\mathcal{K})$ is that its dimension (i.e., number of generators) corresponds to the number of “ k -dimensional holes” in the simplicial complex.

IV. PROBLEM STATEMENT

Let x_t denote the position at time t of a mobile robot operating in a Euclidean domain \mathcal{D} , which is assumed tame (i.e., semi-algebraic). Let u_t and z_t respectively denote the robot’s noisy control inputs and measurements. Adopting the standard Bayesian approach, our objective is to compute the probability of $x_t \in \mathcal{D}$ conditioned upon the prior control and current measurement. Recursively, it is expressed

$$P(x_t|z_t, u_t) = \eta P(z_t|x_t)P(x_t|u_t, x_{t-1}) \quad (2)$$

where $\eta \in \mathbb{R}_+$ is the standard normalization factor and $P(x_t|u_t, x_{t-1})$ represents the prior which is assumed initially uniform over the domain \mathcal{D} .

Our objective in this research is to alleviate the need for precise metric measurements to close the loop. As such, the focus of our attention turns to the measurement likelihood $P(z_t|x_t)$ and how it may be computed without the agent having anything other than the ability to sense/detect (in a binary fashion) uniquely identifiable landmarks not appearing in a metric map. Assuming such limited sensing capabilities, the challenge in determining $P(z_t|x_t)$ reduces to reinterpreting z_t in such a way that it can be used to measure some features embedded within the map’s structure. In this case, we will capture these features by computing local homology within a simplicial nerve complex approximating our domain \mathcal{D} .

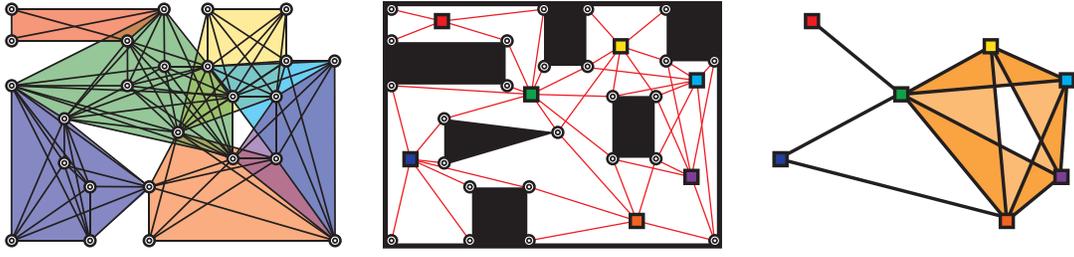


Fig. 1. A collection of observations (center, squares) register the presence of uniquely identifiable landmarks (center, circles) in a planar domain \mathcal{D} via our visibility-based model. The resulting dual complexes, the landmark complex \mathcal{K}_L (left) and the observation complex \mathcal{K}_O (right) approximate \mathcal{D} using two different vertex sets. Note the presence of two holes in the topological structures capture the presence of two large structures in the map (at its center). Although the figures are drawn geometrically, no coordinate information is assumed or used in their construction.

Towards this end, we adopt the following assumptions:

- (A1): The robot is given an unlabeled metric map which reflects the physical layout of \mathcal{D} (e.g., building footprints)
- (A2): The robot is capable of measuring its own motions (alternatively, it has knowledge of its motion model).
- (A3): The robot can *locally sense/detect* from a collection of uniquely identifiable landmarks \mathcal{L} . The landmarks' position is unknown to the robot.
- (A4): A collection of labeled observations \mathcal{O} are taken by the robot as it moves through the environment. The position at which observations are taken is unknown to the robot.
- (A5): The ability of the robot to detect landmarks is encoded as a relation $\mathcal{R} \subset \mathcal{O} \times \mathcal{L}$, in which $(\alpha, i) \in \mathcal{R}$ iff observation $\alpha \in \mathcal{O}$ detects landmark $i \in \mathcal{L}$.

A point of discussion should be assumptions (A3) and (A5), which defines our sensor model in terms of a sensing relationship between observations and the landmarks they see. The generality of this model should not go overlooked as it can be extended to include directional modalities as well. The interested reader is referred to our technical article [30] for additional details.

V. A MULTI-PHASE SOLUTION

Our solution to the problem outlined in IV comes in the form of a multi-phase (pipelined) localization algorithm defined by two key processes: 1) Building a metric-free topological map, \mathcal{K}_L , which we represent using our simplicial nerve constructs to approximate the structure of \mathcal{D} ; and 2) Coupling \mathcal{K}_L with a metric map of the robot's operating environment to enable Monte Carlo Localization (MCL) with local homological sensing.

Algorithm 1 presents our localization pipeline. As stated, the first step (Line 2) in this process is to iteratively build a landmark complex by taking a number of local observations as the robot navigates the environment's free space. This step is metric free and does not require time sequencing. During the second phase, the robot initializes a fixed-size set of random particles and begins the main filtering loop using \mathcal{K}_L and \mathcal{M} . The key components of the measurement update are given by Lines 8 and 11, which make use of two separate radii: r_{pseudo} and r_{geo} . The latter is used to define our local

homology neighborhood within the metric-map \mathcal{M} . As a result it is used for generating the hypothetical measurements for each particle (Line 11, to be discussed shortly). The former is used to loosely couple the simplicial complex with a pseudo-metric approximation to the geodetic bound r_{geo} to enable the robot to identify its local neighborhood in \mathcal{K}_L .

We now elaborate upon each of these points as we next examine the key phases of Algorithm 1. Notably, the algorithms stated in Lines 2, 8 and 11.

Algorithm 1: runTopologicalMCL

Input : Unlabeled metric map, \mathcal{M}
 Num. of particles, $numParts \in \mathbb{Z}_+$
 Num. of observations, $numObs \in \mathbb{Z}_+$
 Pseudo-metric radius, $r_{pseudo} \in \mathbb{R}_+$
 Geodetic radius, $r_{geo} \in \mathbb{R}_+$

```

1 begin
  // Phase I: Generate Map
2  $\mathcal{K}_L \leftarrow buildLandmarkComplex(numObs)$ ;
  // Phase II: Localize Robot
3  $x_{t-1} \leftarrow x_0$ ;
4  $\mathcal{P} \leftarrow initRandParticles(\mathcal{M}, numParts)$ ;
5 for  $t \leftarrow 1$  to  $\infty$  do
  // Motion Update:
6  $x_t \leftarrow moveRobot(x_{t-1})$ ;
7  $\hat{u}_t \leftarrow senseRobotMotion(x_t, x_{t-1})$ ;
  // Measurement Update:
8  $\chi \leftarrow senseLocalHomology(\mathcal{K}_L, r_{pseudo})$ ;
9 foreach Particle  $p_i = (x_{p_i}, w_{p_i}) \in \mathcal{P}$  do
10  $x_{p_i} \leftarrow moveParticle(\hat{u}_t, x_{p_i})$ ;
11  $\chi_{p_i} \leftarrow getLocHomology(p_i, \mathcal{M}, r_{geo})$ ;
12  $w_{p_i} \leftarrow getParticleWeight(\chi, \chi_{p_i})$ ;
13  $\mathcal{P} \leftarrow resampleParticlesByWeight(\mathcal{P})$ ;

```

A. Phase I: Building a Topological Map

To approximate our domain \mathcal{D} , we leverage the notion of *dual landmark and observation complexes* (see Figure 1), which we first presented in [30]. Although we may use either nerve construct as the simplicial representation of our map, for the purposes of this discussion, we focus upon landmark complexes with the understanding that both are

homologically equivalent and refer the interested reader to [30]. We define the landmark complex:

Landmark Complex: The nerve of the cover of \mathcal{O} by subsets $\mathcal{R}^i = \{\alpha \in \mathcal{O} : (\alpha, i) \in \mathcal{R}\}$, whose k -simplices are all collections of $(k + 1)$ distinct landmarks witnessed by some observation α . It is denoted $\mathcal{K}_{\mathcal{L}}$.

According to this definition, only those landmarks that *are observed* participate in the landmark complex; hence it is incorrect to say that, *e.g.*, $\mathcal{K}_{\mathcal{L}}$ is a simplicial complex with vertex set \mathcal{L} , though it is the case when each landmark is observed. Such are the caveats of working with abstract simplicial complexes: we emphasize the difference between the abstract landmarks and their geometric representations, $|\mathcal{L}| \subset \mathcal{D}$, within the workspace. Figure 1 illustrates this construct. Note that there is a corresponding geometric realization of the nerve complex $\mathcal{K}_{\mathcal{L}}$ which has a more direct connection to the problems of mapping and navigation, which is beyond the scope of this effort.

Towards solidifying this definition, refer to Figure 1 (Left) which depicts $\mathcal{K}_{\mathcal{L}}$ for the environment illustrated in Figure 1 (Center). The observation points (denoted as colored squares) can be interpreted as the robot’s location at some point in time where it has identified some set of landmarks. We see that at the lower-leftmost observation point in the center figure (*i.e.*, the blue square), the robot detects six landmarks as indicated by the edges between said observation point and those landmarks (circles). As such, a 5-simplex is introduced into $\mathcal{K}_{\mathcal{L}}$. This 5-simplex is geometrically depicted as the simplex on the lower-left (purple) in Figure 1 (Left).

As it turns out, the definition of the landmark complex readily lends itself to constructing a topological approximation of the free-space under a sufficient number of observations. Intuitively, as the robot drives around via a random walk or some other motion model, it uses its visibility-based, binary, sensing capability to make observations of visible landmarks. If the robot sees $k + 1$ landmarks, denoted as the set ℓ herein, at a particular observation point, it simply introduces a k -simplex into $\mathcal{K}_{\mathcal{L}}$ with vertices (*i.e.*, 0-simplices) corresponding to those landmark identifiers.

Algorithm 2: buildLandmarkComplex

Input : $numObs \in \mathbb{Z}_+$

Output: Landmark Complex $\mathcal{K}_{\mathcal{L}}$

```

1 begin
2    $x_t \leftarrow x_0$ ;
3    $\mathcal{K}_{\mathcal{L}} = \emptyset$ ;
4   for  $i \leftarrow 1$  to  $numObs$  do
5      $\ell \leftarrow getVisibleLandmarkIDs(x_t)$ ;
6      $\sigma_{\ell} \leftarrow createNewSimplex(\ell)$ ;
7      $\mathcal{K}_{\mathcal{L}} \leftarrow \mathcal{K}_{\mathcal{L}} \cup \sigma_{\ell}$ ;
8      $x_t \leftarrow moveToNewLocation(x_t)$ ;

```

Algorithm 2 presents a formal statement of the iterative process. Although the algorithm references the robot’s state,

x_t , the robot is assumed to have no knowledge of x_t , and its inclusion is solely for the purposes of clarity. Note that this Algorithm *does not require the robot to have knowledge of metric information regarding itself or the landmarks*. Furthermore, it *does not require any reference to time* to sequence the robot’s observations. In fact, with only slight modification, Algorithm 2 may be made fully distributed provided the robots involved assign consistent identifiers between viewed landmarks (note that in the SSID case, no consistency check, *i.e.* data association, is necessary since landmarks broadcast their unique identifiers). Such are the advantages of leveraging topological tools.

B. Phase II: MCL with Homological Measurements

Since the prediction step in our formulation (*i.e.*, Lines 6 and 7 in Algorithm 1) is standard, being either driven by an onboard IMU or a known motion model, we opt to focus instead on enabling the correction (*i.e.* the measurement update) for loop closure. The measurement update is driven by two key components: 1) the actual measurement taken by the robot (Line 8) and 2) the hypothetical (or expected) measurements for each of the particles reflecting a feasible state (Line 11). For the sake of brevity, we restrict our discussion of Phase II to the process of generating both the actual and hypothetical measurements and refer the interested reader to [19] for a proper treatment of particle filtering.

Herein, we consider a standard SIR (Sample Importance Resampling) particle filter over the robot’s state space (*i.e.*, its location in \mathcal{D}). Our choice of a non-parametric Bayesian filter is driven by the realization that homological sensing is inherently coarse and, thus, the posterior over robot positions may exhibit multiple modes, and, in general, will be non-Gaussian. Note that the structure of the posterior will heavily depend upon the topological structure of the environment as well as the density of landmarks.

A key advantage to our approach is that each hypothetical measurement (see Line 11 in Algorithm 1) can be obtained in $O(1)$ time by using a simple lookup table generated via a geometric preprocessing phase on the metric map, \mathcal{M} . Furthermore, the homological computations (Line 8, Algorithm 1) required by the robot may also be efficiently done (in-practice, under normal observation density) by doing a simple summation over the simplices comprising the subcomplex representing the robot’s local neighborhood when $x_t \in \mathbb{R}^2$.

1) *Taking Local Homological Measurements:* Given that $\mathcal{K}_{\mathcal{L}}$ is built during Phase I (see Line 2 in Algorithm 1), the robot may perform local homology computations by first extracting the subcomplex of $\mathcal{K}_{\mathcal{L}}$ corresponding to its local neighborhood. Since $\mathcal{K}_{\mathcal{L}}$ is constructed without explicit metric measurements, the notion of neighborhood is not readily defined on its structure. Fortunately, in our case, the local sensing relationship (**A3**) reflecting our underlying hardware capabilities enables us to exploit the fact that the agent can only sense within some maximal range. As such, we may couple the edges (*i.e.*, 1-simplices) within $\mathcal{K}_{\mathcal{L}}$ loosely with a metric scale by associating with each

edge the maximum sensing range, $s_{max} \in \mathbb{R}_+$ of the robot. Note that s_{max} need not be precise given the coarse nature of our homological sensor model. Assuming the robot has computed the shortest pairwise path between each landmark pair² in the 1-skeleton of $\mathcal{K}_{\mathcal{L}}$, it may simply lookup and “trace” the vertices along all the paths beginning in ℓ until r_{pseudo} is reached along each. The result of this process will be a collection of vertices that may be used as “indices” to extract the appropriate subcomplex of $\mathcal{K}_{\mathcal{L}}$. We denote the subcomplex capturing the robot’s local homology as \mathcal{K}_{ℓ} .

This allows us to embed a coarse notion of locality within $\mathcal{K}_{\mathcal{L}}$ without the need for precise landmark positions or high-fidelity sensor calibrations. Ultimately though, it enables us to use the k th-homology over local neighborhoods in $\mathcal{K}_{\mathcal{L}}$ to count (or sense) the number of topological invariants within the robot’s vicinity. Recall that computing the k th-homology corresponds to computing the generators representing non-reducible k -cycles within the local region of interest. Fortunately, the field of computational homology provides a number of methods to enable this computation [31]. Additional recent techniques include spectral methods (using Laplacians) [6] and discrete Morse theory [32].

2) *Euler Characteristic*: Though computable, $H_k(\mathcal{K})$ can be problematic, as the size of the complex can grow very quickly where landmark density is sufficiently high. Furthermore, the full computation of homology provides very rich information: more than is needed for our purposes. We can more easily compute a reduced topological invariant — the *Euler characteristic* — given by:

$$\chi(\mathcal{K}) = \sum_{\sigma} (-1)^{\dim \sigma} = \sum_{k=0}^{\infty} (-1)^k \dim C_k.$$

The following well-known result is a trivial but useful consequence of the fact the the Euler characteristic is also the alternating sum of the dimensions of the homology groups:

Lemma 5.1 ([31]): For \mathcal{K} a connected simplicial complex that is homotopic to a subcomplex of the plane \mathbb{R}^2 ,

$$\chi(\mathcal{K}) = 1 - \dim H_1(\mathcal{K}).$$

Use of the Euler characteristic is especially compelling for connected planar subsets, as it completely characterizes topological type via a simple simplex count. In what follows, we assume the use of the Euler characteristic to classify local landmark complexes to which, because they approximate local planar regions, Lemma 5.1 applies.

Algorithm 3 formalizes the process of computing the local homology within $\mathcal{K}_{\mathcal{L}}$ as an implicit function of the robot’s location in \mathcal{D} . The algorithm returns the number of holes or invariants, $\chi_{\ell} \in \mathbb{Z}_{\geq 0}$, sensed in the local homology \mathcal{K}_{ℓ} . Line 3 of the algorithm recovers the robot’s local subcomplex. Finally, in Line 4, the number of local homological invariants are computed giving the robot its measurement, χ_{ℓ} .

It should be noted that it is natural to restrict ourselves to the Euler characteristic whenever operating in \mathbb{R}^2 ; however,

²This process may be readily done prior to starting the localization filter in Algorithm 1

Algorithm 3: senseLocalHomology

Input : $\mathcal{K}_{\mathcal{L}}, r_{pseudo} \in \mathbb{R}_+$

Output: $\chi_{\ell} \in \mathbb{Z}_{\geq 0}$

1 **begin**

2 $\ell \leftarrow getVisibleLandmarkIDs();$

3 $\mathcal{K}_{\ell} \leftarrow getLocalSubcomplex(\mathcal{K}_{\mathcal{L}}, \ell, r_{pseudo});$

4 $\chi_{\ell} \leftarrow \chi(\mathcal{K}_{\ell});$

one may consider an alternate approach for higher dimensional localization where Lemma 5.1 does not hold. In such a case, local homology may be computed by analyzing the spectrum of the combinatorial Laplacian associated with a given simplicial complex \mathcal{K} :

$$L_k = \partial_k^T \partial_k + \partial_{k+1} \partial_{k+1}^T.$$

It is well-known that the dimensionality of the kernel of L_k corresponds to the number of non-trivial generators in the k th homology group [6]. The disadvantage of this approach is the need to do a spectral decomposition, which, in general, is computationally more expensive.

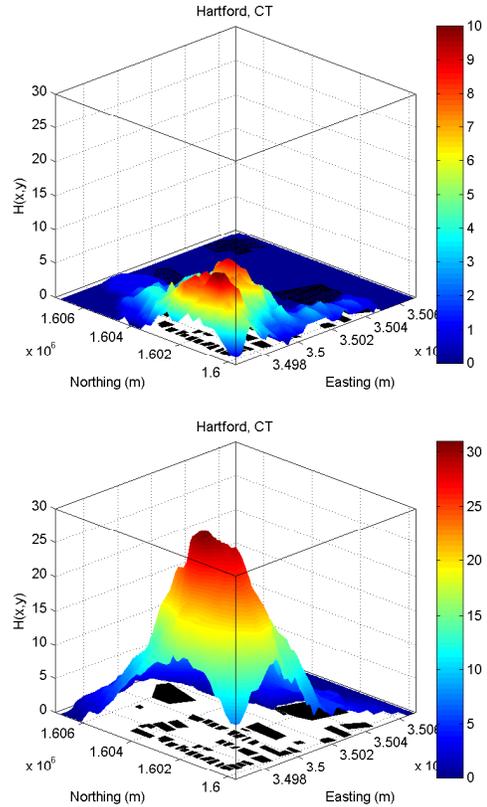


Fig. 2. Height maps corresponding to “near” (top) and “far” (bottom) for downtown region of Hartford, CT comprised of 44 buildings/structures. The maps were generated using our geometric preprocessing algorithm for computing the number of expected holes as a function of map position by exploiting the robot’s visibility-based sensing model. Note that since the maps are computed *a priori*, each particle in our filter can obtain its hypothetical measurements using an $O(1)$ table lookup.

3) *Generating Hypothetical Measurements*: Accordingly, we now turn our attention to the task of generating hypothetical measurements for our set of particles. Assuming our map, \mathcal{M} , is comprised of a finite collection of polygons (e.g., perhaps corresponding to building structure footprints), the task of computing a hypothetical measurement (Algorithm 1, Line 11) given a particular hypothesis (i.e., a particle) reduces to identifying the polygons that lie fully within the geodetic radii, r_{geo} defining the robot’s homological neighborhood. These structures represent expected holes in the landmark complex $\mathcal{K}_{\mathcal{L}}$ generated during Phase I. Exploiting this realization, we can employ a geometric preprocessing phase over \mathcal{M} in which we consider a tessellation (or, alternatively, random samples) of the map’s free space. If we map each such observation or sample to the expected number of holes/structures seen at that point, we can construct a “height” (or elevation) map over \mathcal{M} that reflects the local homology of the hypothesis. The advantage of computing this map *a priori* is that it can be used as a lookup table enabling each particle to obtain its hypothetical measurements in $O(1)$ time (e.g., by employing nearest neighbor interpolation). Given the straightforward nature of this algorithm, we omit a formal statement.

From an implementation standpoint, however, it should be noted that we assume a polygon is represented as a set of points given by a surface mesh over its comprising faces. For example, in \mathbb{R}^2 this mesh corresponds to discretizing the line segments comprising the edges of the polygon into smaller segments. Given this data representation, determining whether a polygon represents a likely topological invariant in $\mathcal{K}_{\mathcal{L}}$ reduces to verifying that all such mesh points satisfy the geodetic distance constraint.

VI. SIMULATION RESULTS

The proposed localization pipeline was implemented in a Matlab-based simulation. In the scenario, a region of downtown Hartford, Connecticut was selected. The metric map was comprised of 44 buildings with various geometric polygonal footprints. A total of 197 WiFi access points serving as landmarks were placed within the environment. The robot was charged with following a specified path along the roadways in the map. In lieu of a motion model, the robot was assumed outfitted with an IMU for its prediction step. Its noise model was assumed additive Gaussian with $\sigma_{linear} = 10.0(m/s^2)$ and $\sigma_{angular} = 0.435(rads/sec^2)$. A total of 8000 particles represented the prior/posterior distributions. It should be noted that we have obtained similar results with fewer samples.

Employing Algorithm 2 for Phase I results in the simplicial nerve complex seen in Figure 3 (Center). Although not necessary, in this particular instance, we assume $\mathcal{O} = \mathcal{L}$ thereby reducing the complexity of certain aspects of our implementation. Given the sufficient number of samples, we see that in this case $\mathcal{K}_{\mathcal{L}}$ faithfully captures the topology of \mathcal{D} . In addition to the metric map, which corresponds to the unlabeled map depicting building footprints, $\mathcal{K}_{\mathcal{L}}$ is fed as input into Phase II.

Given the coarse nature of the homological sensing modality, we decomposed the homological measurement z into two separate components, z_{near} and z_{far} , corresponding to “near” and “far” local homologies. Given the choice of sensor range, the number of scaled hops in the 1-skeleton of $\mathcal{K}_{\mathcal{L}}$ defining the near neighborhood was 2, while the far neighborhood was 3.

Since the agent is analyzing the local homology of $\mathcal{K}_{\mathcal{L}}$, it is possible for “phantom” holes to occur once the desired neighborhood (a subcomplex) is extracted from the global homological structure (i.e., a hole may be introduced that does not correspond to a physical structure in the map). Conversely, some structures in the map (depending upon the visibility-based sensing modality of the agent) may not necessarily introduce topological features in $\mathcal{K}_{\mathcal{L}}$. Such artifacts may be considered a form of topological noise.

For the purposes of discussion, we consider a standard additive model $\hat{z}_t = z_t + \epsilon_t$ to capture this noise and define our likelihood function as the product of terms

$$P(z_t|x_t) = P(z_t^{near}|x_t)P(z_t^{far}|x_t)$$

where, again, $P(z_t|x_t)$ reflects the probability of measurement z_t at location x_t at time t . Although a discrete distribution may be used, for the purposes of our simulations, we chose $\epsilon_t = N(\mu_t, \sigma_t^2)$ with inflated covariance to mitigate the disparity introduced by measuring only values of Z_+ . A complete characterization of a noise model is the focus of ongoing research; however, even with this simple model, results were very satisfactory.

Figure 3 shows the evolution of our particle filter with homological sensing for loop closure. Going from top to bottom, each row of images represents a snapshot of our filter with respect to time. The left most figure shows the geometric realization of the environmental landmarks. Landmarks sensed at a particular instance are indicated via squares with landmarks respectively belonging to the “near” and “far” neighborhoods, as determined by our pseudo-metric, being color-coded. The coordinate-free landmark complex is shown center with near and far local homology regions color-coded to match the geometric realization of the landmarks appearing in the left most figure. The right figure shows the approximated posterior given by the particle filter.

Figure 4 shows the absolute error of the homological measurements for the duration of the simulation shown in Figure 3. Note that despite substantial homological errors (on the order of 8 holes at one point), the algorithm still converges as shown in Figure 3 (Right).

VII. CONCLUSION

In this paper we have proposed a novel filtering approach that enables a mobile robot to localize leveraging binary information, in the form of SSIDs from WiFi access points, and fuse it with information from an IMU and a unlabeled map of the environment. We show that, despite the fact the uniquely identifiable features may have unknown positions (such as WiFi access points), it is possible to leverage such binary information for Monte-Carlo localization. In particular

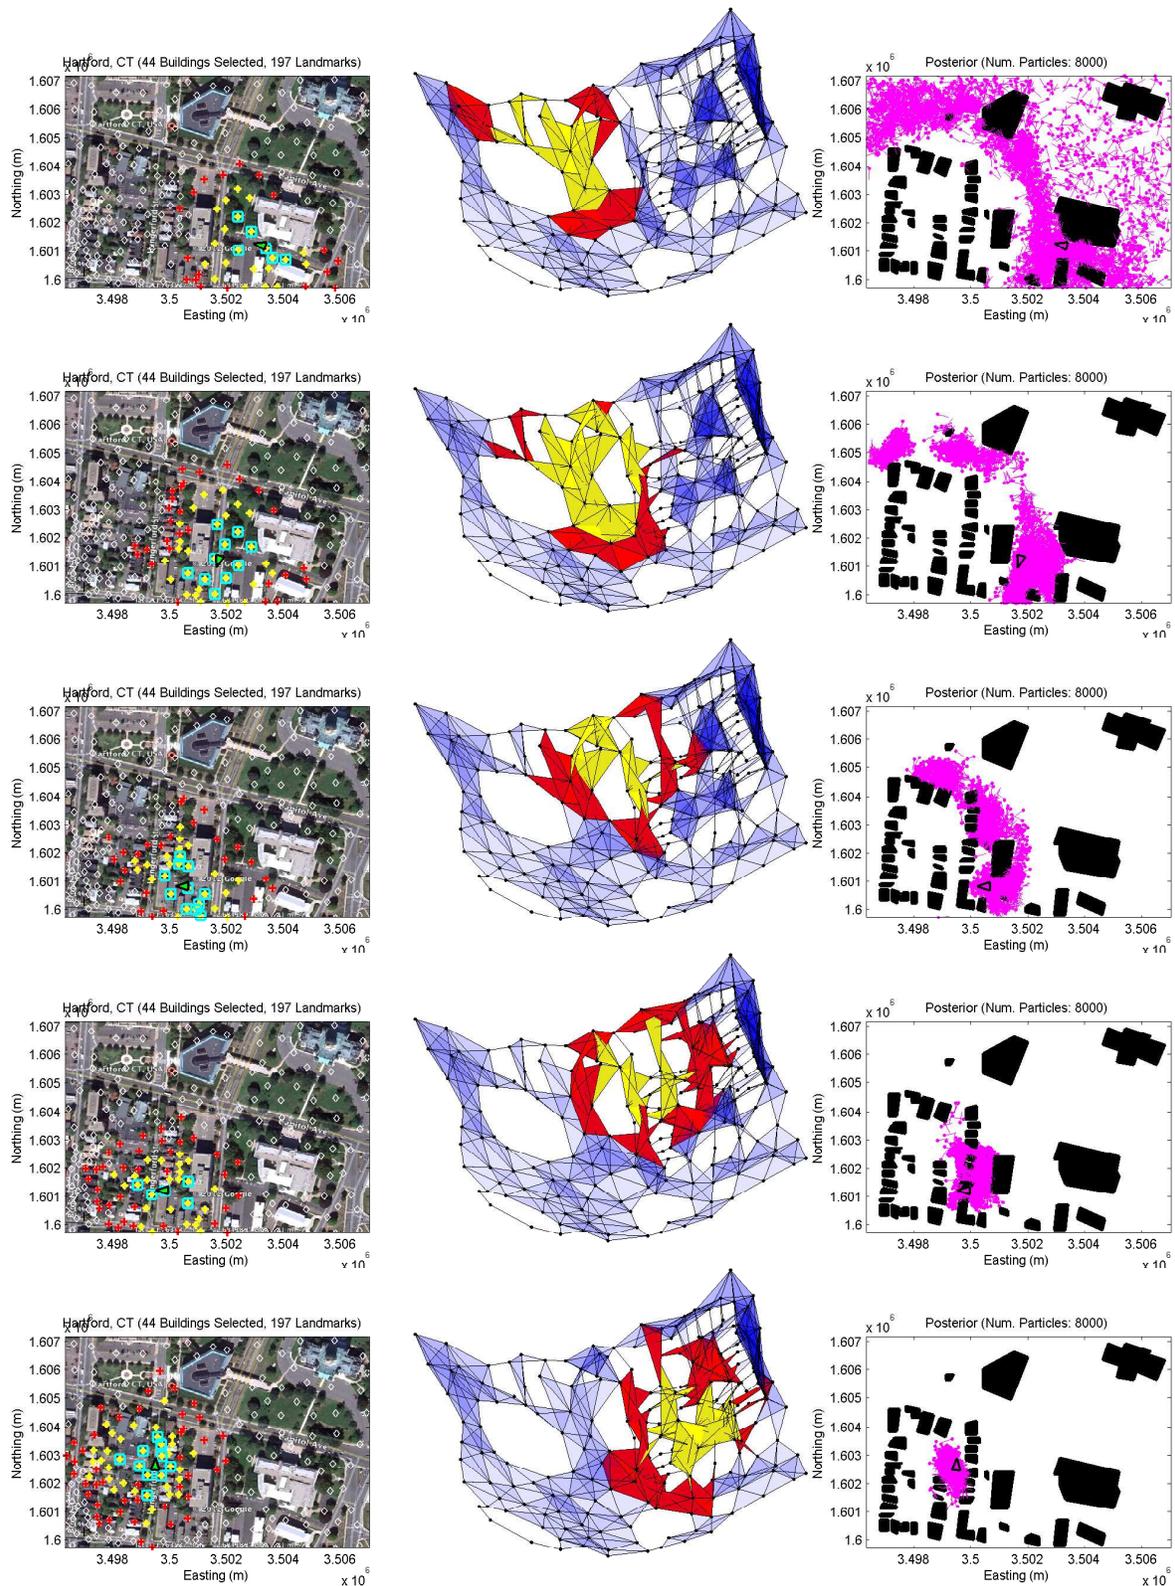


Fig. 3. MCL with homological sensing over downtown Hartford, CT [33] with 44 buildings and 197 simulated WiFi APs serving as landmarks (Top-to-Bottom, Left) - Squares: Visible landmarks used by robot to interrogate local homology within $\mathcal{K}_{\mathcal{L}}$; Red +': Landmarks in "far" neighborhood; Yellow +': Landmarks in "near" neighborhood; (Top-to-Bottom, Center) - Landmark complex/topological map shown with near/far homology regions as determined by visible landmarks in (Top-to-Bottom, Left); (Top-to-Bottom, Right) The evolution of the posterior over the robot's state space shown with environmental structures as polygons. Loop closure was achieved using only the coordinate-free, homological sensing of said structures. The robot has no knowledge of landmark coordinates, and it uses only a binary, visibility-based model for uniquely identifying landmarks before interrogating the local homology of $\mathcal{K}_{\mathcal{L}}$.

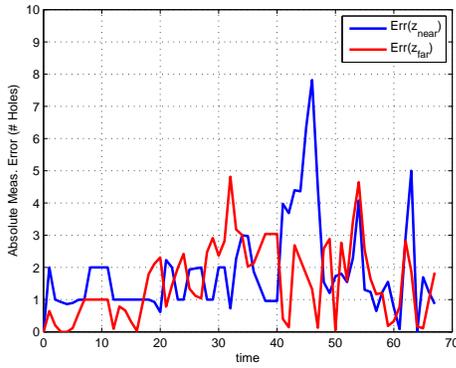


Fig. 4. Absolute error for homological measurements ($z = [z_{near}, z_{far}]$) corresponding to the simulation in Figure 3. The ground truth for a particular location was determined using the results of geometric preprocessing algorithm with linear interpolation between sample points. As shown in Fig. 3 (Right), the filter still converges to a reasonable estimate of the robot's position.

we show that the binary information can be fused to build a topological map that is consistent with the underlying geometric environment map. The mobile robot can interrogate the topological map through local homology computation and obtain the number of holes in the topological map. This can be efficiently correlated with the underlying geometric map to provide the likelihood position of the mobile robot.

The main advantage of the proposed approach is two-fold. On one side, it enables minimalistic localization and on the other it shows how it is possible to fuse efficiently highly coarse/quantized information within a Monte-Carlo filter.

As future work we plan to pursue a more quantitative characterization of the properties of such methodology. It is clear that the accuracy of the topological map (in term of hole-to-building correspondence) strongly depend upon the density of the landmarks (e.g., WiFi access points) and the size of the holes/buildings. In this context some recent results developed at the interface between signal processing and algebraic topology [15] will be instrumental to pursue this analysis. We also plan to study more precisely the effect of noise on the performance. Although we already showed that persistent homology can be used to remove/reduce the effect of noise at the topological level, see [30], its effect on Monte-Carlo localization requires further investigation.

REFERENCES

- [1] D. Kurth, G. Kantor, and S. Singh, "Experimental results in range-only localization with radio," in *IROS*, 2003.
- [2] D. Moore, J. Leonard, D. Rus, and S. Teller, "Robust distributed network localization with noisy range measurements," in *International conference on Embedded Networked Sensor Systems (SenSys)*, 2004.
- [3] E. Olson, J. Leonard, and S. Teller, "Range-only beacon localization," *Journal of Oceanic Engineering*, vol. 31, no. 4, 2006.
- [4] V. De Silva and R. Ghrist, "Coordinate-free coverage in sensor networks with controlled boundaries via homology," *Int. J. Rob. Res.*, vol. 25, no. 12, pp. 1205–1222, 2006.
- [5] A. Muhammad and M. Egerstedt, "Control using higher order Laplacians in network topologies," in *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems*, 2006, pp. 1024–1038.

- [6] A. Tahbaz-Salehi and A. Jadbabaie, "Distributed coverage verification in sensor networks without location information," *IEEE Trans. Automat. Contr.*, vol. 55, no. 8, pp. 1837–1849, 2010.
- [7] A. Muhammad and A. Jadbabaie, "Dynamic coverage verification in mobile sensor networks via switched higher order laplacians," in *Robotics: Science and Systems III*. The MIT Press, 2007.
- [8] M. Erdmann, "On the topology of discrete strategies," *Int. J. Rob. Res.*, vol. 29, no. 7, pp. 855–896, 2010.
- [9] S. Bhattacharya, M. Likhachev, and V. Kumar, "Topological constraints in search-based robot path planning," *Autonomous Robots*, vol. 33, pp. 273–290, 2012.
- [10] S. Bhattacharya, D. Lipsky, R. Ghrist, and V. Kumar, "Invariants for homology classes with application to optimal search and planning problem in robotics," *Electronic pre-print*, Aug 2012.
- [11] S. Kim, K. Sreenath, S. Bhattacharya, and V. Kumar, "Optimal trajectory generation under homology class constraints," in *51st IEEE Conference on Decision and Control*, 10-13 Dec 2012.
- [12] M. Farber and M. Grant, "Topological complexity of configuration spaces," *ArXiv e-prints*, 2008.
- [13] Y. Baryshnikov and R. Ghrist, "Target enumeration via euler characteristic integrals," *SIAM Journal of Applied Mathematics*, vol. 70, no. 3, pp. 825–844, 2009.
- [14] G. C. Haynes, F. R. Cohen, and D. E. Koditschek, "Gait transitions for quasi-static hexapedal locomotion on level ground," in *International Symposium of Robotics Research*, 2009.
- [15] M. Robinson and R. Ghrist, "Topological localization via signals of opportunity," *IEEE Transactions on Signal Processing*, pp. 2362–2373, 2012.
- [16] A. Ranganathan and F. Dellaert, "Online probabilistic topological mapping," *The International Journal of Robotics Research*, 2011.
- [17] W. H. Huang and K. R. Beevers, "Topological map merging," *Int. J. Rob. Res.*, vol. 24, no. 8, pp. 601–613, 2005.
- [18] D. Marinakis and G. Dudek, "Pure topological mapping in mobile robotics," *Trans. Rob.*, vol. 26, no. 6, pp. 1051–1064, 2010.
- [19] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*, ser. Intelligent robotics and autonomous agents. MIT Press, September 2005.
- [20] S. Thrun, D. Fox, W. Burgard, and F. Dellaert, "Robust monte carlo localization for mobile robots," 2001.
- [21] D. Frank, "A sample of Monte Carlo methods in robotics and vision," in *Proceedings of the ISM International Symposium on the Science of Modeling - The 30th Anniversary of the Information Criterion (AIC)*, 2003, pp. 211–222.
- [22] J. Biswas and M. M. Veloso, "Wifi localization and navigation for autonomous indoor mobile robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 4379–4384.
- [23] A. Howard, S. Siddiqi, and G. S. Sukhatme, "An experimental study of localization using wireless ethernet," in *Proceedings of the International Conference on Field and Service Robotics*, Jul 2003.
- [24] Y. Lin, P. Vernaza, J. Ham, and D. D. Lee, "Cooperative relative robot localization with audible acoustic sensing," in *IROS'05*, 2005, pp. 3764–3769.
- [25] B. Ferris, D. Hähnel, and D. Fox, "Gaussian processes for signal strength-based location estimation," in *Robotics: Science and Systems*, 2006.
- [26] F. Dellaert, F. Alegre, and M. E. B., "Intrinsic localization and mapping with 2 applications: Diffusion mapping and Marco Polo localization," in *ICRA*, May 2003.
- [27] K. E. Bekris, M. Glick, and L. E. Kavraki, "Evaluation of algorithms for bearing-only slam," in *Proceedings of the 2006 IEEE International Conference on Robotics and Automation*, 2006, pp. 1937–1943.
- [28] A. Hatcher, *Algebraic Topology*. Cambridge University Press, 2002.
- [29] A. Björner, "Topological methods," in *Handbook of combinatorics (vol. 2)*, R. L. Graham, M. Grötschel, and L. Lovász, Eds. Cambridge, MA, USA: MIT Press, 1995, pp. 1819–1872.
- [30] R. Ghrist, D. Lipsky, J. Derenick, and A. Speranzon, "Topological landmark-based navigation and mapping," University of Pennsylvania, Department of Mathematics, Tech. Rep., 08 2012.
- [31] T. Kaczynski, K. Mischaikow, and M. Mrozek, *Computational Homology*, ser. Applied Mathematical Sciences. Springer, 2004.
- [32] S. Harker, K. Mischaikow, M. Mrozek, and V. Nanda, "Discrete Morse theoretic algorithms for computing homology of complexes and maps," 2012, preprint.
- [33] Google, <http://maps.google.com>.