Generalized Coverage in Homological Sensor Networks

Nicholas J. Cavanna Kirk Gardner
Donald R. Sheehy

1 Introduction

The homological sensor network (HSN) coverage problem is to verify that a collection of sensors cover an unknown, bounded domain. The sensors can be interpreted as a collection of discs with uniform radius in $\mathbb{R}^2$, or more generally, uniform d-balls centered at a set of points $\mathbb{R}^d$. Moreover, it is assumed that the network is known only through neighbor relationships, and the coordinates of the points are unknown.

Over a series of papers de Silva and Ghrist developed a theory of coordinate-free coverage which used the relative homology of a sample to a sampled boundary to verify coverage of a domain. Assuming the network is sampled from domain with a smooth boundary it is shown that coverage of the domain by a sensor network can be verified by looking at the persistent homology of the network sensor's radius relative to a sampled boundary. The work culminated in a simple, computable topological coverage criterion (TCC), a necessary (but not sufficient) condition to guarantee a collection of sensors cover the unknown domain [4].

This paper provides a new proof of the TCC that allows us to extend results about HSNs in several different directions. The main contributions are as follows:

1. We give a new proof of correctness for the TCC that replaces the assumption that the boundary of the domain is a smooth manifold with a very weak topological condition.

2. We extend the TCC to weighted k-nearest neighbor distance in order to allow verification of coverage by k weighted sample points. The weights allow for different radii.

3. We give an algorithm for *certified homology inference*, i.e. we compute the homology of the domain and provide some guarantee if there was sufficient data.

2 Background

Throughout the domain $\mathcal{D}$ is a compact, locally contractible set with boundary $\mathcal{B}$ forming a pair $(\mathcal{D}, \mathcal{B})$. The pair (P, Q) is composed of a sample $P \subset \mathcal{D}$ and $Q = P \cap \mathcal{B}^\alpha$ for a fixed scale $\alpha \geq 0$. In order to generalize coverage in the TCC we consider verifying coverage of the domain by k sensors with distinct radii. This is achieved by defining a distance function from points in $\mathbb{R}^d$ to subsets of k weighted sample points. For a compact point set $A \subset \mathbb{R}^d$ let $\binom{A}{k}$ be the set of k-element subsets of A. The weighted distance from a point $x \in R^d$ to a weighted point $y \in A$ is defined as the power distance $\rho_y(x)^2 := \|x - y\|^2 + w_y^2$. The **weighted k-nearest neighbor distance** from a point x to k points in a weighted compact set A is defined as

$$d_k(x, A)^2 := \inf_{K \in \binom{A}{k}} \max_{y \in K} \rho_y(x).$$

We define the **weighted (k, ε)-offsets** of a point set $A \subset \mathbb{R}^d$ to be

$$A_k^\varepsilon := \{x \in X \mid d_k(x, A) \leq \varepsilon\}.$$

The k-barycentric decomposition of a simplicial complex $\mathcal{K}$ is a new simplicial complex that has a vertex for every simplex of $\mathcal{K}$ of dimension at least $k-1$. The simplices are flags of these simplices, i.e., sequences ordered by inclusion. By the Nerve Theorem and [3], the weighted (k, ε)-offsets are homotopy equivalent to the k-barycentric decomposition of the Nerve of the sample point's weighted offsets, or Čech complex at scale ε denoted $\mathcal{C}_\varepsilon^k(P)$.

Although the Čech complex is homotopy equivalent to the offsets, a related complex, the (Vietoris-)Rips complex can be computed directly from the network information. This complex has a simplex for every clique of the intersection graph of the balls. By [3] the k-barycentric decomposition of the Rips complex, denoted $\mathcal{R}_\varepsilon^k(A)$, is a ϑ_d-approximation to the (k, ε)-offsets A_k^ε for a weighted set A at scale ε. This generalizes an important result about the relationship of Čech and Rips complexes following from Jung's Theorem [2], relating the diameter of a point set P and the radius of the minimum enclosing ball as follows:

$$\mathcal{C}_\varepsilon^k(P) \subseteq \mathcal{R}_\varepsilon^k(P) \subseteq \mathcal{C}_{\vartheta_d \varepsilon}^k(P) \tag{1}$$

where $\vartheta_d = \sqrt{\frac{2d}{d+1}}$. Our new TCC proof is greatly simplified by working directly with the union of balls $(P_k^{\vartheta_d \varepsilon}, Q_k^{\vartheta_d \varepsilon})$ to first give the geometric underpinnings and then using the relationship from Jung's Theorem to give an algorithm in terms of a sampled pair of weighted (k, ε)-Rips complexes, denoted $\mathrm{R}^k_\varepsilon = (\mathcal{R}^k_\varepsilon(P), \mathcal{R}^k_\varepsilon(Q))$.

The Alexander Duality is used to turn questions about coverage into questions about connectivity. The d-dimensional cohomology of the pair $(\mathcal{D}, \mathcal{B})$ is isomorphic to the 0-dimensional homology of its complement. Under the duality this is equivalent to the number of connected components in the shrunken domain that are disconnected from the complement of the domain.

$$H^d(\mathcal{D}, \mathcal{B}) \cong H_0(\overline{\mathcal{B}}, \overline{D}) \cong H_0(\mathcal{D} \setminus \mathcal{B}).$$

3 The TCC

To certify coverage without coordinates we need to make assumptions on the shape of our bounded domain as a subset of $\mathbb{R}^d$. In our full paper, the new TCC posed in Theorem 1 is proven first in terms of the union of balls. The computable condition about the homology of Rips coplexes then follows via the the Rips-Čech interleaving and the so-called Persistent Nerve Lemma. The proof is based on the following diagrams. Since $Q \subseteq \mathcal{B}^\alpha$, for $\varepsilon \geq 0$, $Q^\varepsilon_k \subseteq \mathcal{B}^{\alpha+\varepsilon}$ we have that Diagram 2 commutes where the maps are inclusions on the pairs of spaces.

$$\begin{array}{ccc} (P^\alpha_k, Q^\alpha_k) & \hookrightarrow & (P^\beta_k, Q^\beta_k) \\ \downarrow & & \downarrow \\ (\mathcal{D}^{2\alpha}, \mathcal{B}^{2\alpha}) & \hookrightarrow & (\mathcal{D}^{\alpha+\beta}, \mathcal{B}^{\alpha+\beta}) \end{array} \qquad (2)$$

We can then use the Alexander duality to construct the following commutative diagram which allows us to state our assumptions in terms of connectivity.

$$\begin{array}{ccc} \mathrm{H}_0(\overline{\mathcal{B}^{\alpha+\beta}}, \overline{\mathcal{D}^{\alpha+\beta}}) & \xrightarrow{j_*} & \mathrm{H}_0(\overline{\mathcal{B}^{2\alpha}}, \overline{\mathcal{D}^{2\alpha}}) \\ \downarrow & & \downarrow \\ \mathrm{H}_0(\overline{Q^\beta_k}, \overline{P^\beta_k}) & \xrightarrow{i_*} & \mathrm{H}_0(\overline{Q^\alpha_k}, \overline{P^\alpha_k}) \end{array} \qquad (3)$$

Theorem 1 poses the TCC in terms of the number of connected components of the domain relative to its boundary.

Theorem 1 (Generalized TCC)**.** *Let $\mathcal{D} \subset \mathbb{R}^d$ be a set with boundary $\mathcal{B}$, a weighted sample P and $Q = P \cap B^\alpha$ where $\beta \geq 3\alpha$. If*

$$j_* = H_0((\mathcal{D} \setminus \mathcal{B}^{\alpha+\beta}) \hookrightarrow (\mathcal{D} \setminus \mathcal{B}^{2\alpha}))$$

is an isomorphism and

$$\mathrm{H}_d\left(\mathrm{R}^k_{\alpha/\sqrt{2}} \hookrightarrow \mathrm{R}^k_\beta\right) = \mathrm{rank}\ j_*$$

then $\mathcal{D} \setminus B^{2\alpha} \subseteq P^\alpha_k$

4 Algorithm

Algorithm 1 is used to check coverage of the shrunken domain i.e., that $\mathcal{D} \setminus \mathcal{B}^{2\alpha} \subseteq P^\alpha_k$, assuming we can calculate or know a priori the number of connected components m of $\mathcal{D} \setminus \mathcal{B}^{2\alpha}$. In order to apply the criterion

Algorithm 1 Check if $\mathcal{D} \setminus B^{2\alpha} \subseteq P^\alpha_k$

```
procedure COVERAGE(α, β, P, Q, m)
    construct R^k_{α/ϑ_d} and R^k_β
    let r := rank H_d(R^k_{α/ϑ_d} ↪ R^k_β)
    if m = r then return True
    else return False
```

to cases where the number of connected components is not known we generalize Algorithm 1 to give a certification of coverage on those components that have been sampled. Then, even if we do not know the number of connected components of D_0, as long as we know which components have been sampled we can provide a condition which certifies weighted k-coverage of the sampled subdomain.

Certified Coverage

By throwing out points too close to the boundary we can do homology inference under the assumption we only have partial coverage. We can use the topological coverage criterion to certify coverage assuming the number of connected components is known by taking a subsample of P of points far enough from the boundary. Under the stronger shape hypothesis that the so-called weak feature size of $\mathcal{B}$ is sufficiently large as in [1], all of the Betti numbers can be computed from the homology of subsamples inclusions of Rips complexes. This connects results on homology inference to results on HSNs.

References

[1] Frédéric Chazal and André Lieutier. Weak feature size and persistent homology: Computing homology of solids in R^n from noisy data samples. In *Proceedings of the 21st ACM Symposium on Computational Geometry*, pages 255–262, 2005.

[2] Heinrich Jung. Über die kleinste kugel, die eine räumliche figur einschließt. *J. Reine Angew. Math*, 123:214–257, 1901.

[3] Donald R. Sheehy. A multicover nerve for geometric inference. In *The Canadian Conference in Computational Geometry*, pages 309–314, 2012.

[4] Vin De Silva and Robert Ghrist. Coordinate-free coverage in sensor networks with controlled boundaries via homology. *International Journal of Robotics Research*, 25:1205–1222, 2006.