


Sheaf-Theoretic Stratification Learning


Adam Brown¹

Department of Mathematics, University of Utah
Salt Lake City, UT, USA
abrown@math.utah.edu

 <https://orcid.org/0000-0002-0955-0119>

Bei Wang²

School of Computing, Scientific Computing and Imaging Institute, University of Utah
Salt Lake City, UT, USA
beiwang@sci.utah.edu

 <https://orcid.org/0000-0002-9240-0700>

Abstract

In this paper, we investigate a sheaf-theoretic interpretation of stratification learning. Motivated by the work of Alexandroff (1937) and McCord (1978), we aim to redirect efforts in the computational topology of triangulated compact polyhedra to the much more computable realm of sheaves on partially ordered sets. Our main result is the construction of stratification learning algorithms framed in terms of a sheaf on a partially ordered set with the Alexandroff topology. We prove that the resulting decomposition is the unique minimal stratification for which the strata are homogeneous and the given sheaf is constructible. In particular, when we choose to work with the local homology sheaf, our algorithm gives an alternative to the local homology transfer algorithm given in Bendich et al. (2012), and the cohomology stratification algorithm given in Nanda (2017). We envision that our sheaf-theoretic algorithm could give rise to a larger class of stratification beyond homology-based stratification. This approach also points toward future applications of sheaf theory in the study of topological data analysis by illustrating the utility of the language of sheaf theory in generalizing existing algorithms.

2012 ACM Subject Classification Mathematics of computing → Algebraic topology; Computing methodologies → Dimensionality reduction and manifold learning

Keywords and phrases Sheaf theory, stratification learning, topological data analysis, stratification

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.14

Related Version A full version is available at <https://arxiv.org/abs/1712.07734>.

1 Introduction

Our work is motivated by the following question: Given potentially high-dimensional point cloud samples, can we infer the structures of the underlying data? In the classic setting of *manifold learning*, we often assume the support for the data is from a low-dimensional space with manifold structure. However, in practice, a significant amount of interesting data contains mixed dimensionality and singularities. To deal with this more general scenario, we assume the data are sampled from a mixture of possibly intersecting manifolds; the objective is to recover the different pieces, often treated as clusters, of the data associated with different

¹ NSF IIS-1513616

² NSF IIS-1513616 and NSF ABI-1661375



© Adam Brown and Bei Wang;

licensed under Creative Commons License CC-BY


34th International Symposium on Computational Geometry (SoCG 2018).


Editors: Bettina Speckmann and Csaba D. Tóth; Article No. 14; pp. 14:1–14:14

Leibniz International Proceedings in Informatics




LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany





■ **Figure 1** Example of a topological stratification of a *pinched torus*.



■ **Figure 2** Example of a homological stratification of a *sundial*.

manifolds of varying dimensions. Such an objective gives rise to a problem of particular interest in the field of *stratification learning*. Here, we use the word “stratification learning” loosely to mean an unsupervised, exploratory, clustering process that infers a decomposition of data into disjoint subsets that capture recognizable and meaningful structural information.

Previous work in mathematics has focused on the study of stratified spaces under smooth and continuous settings [12, 25] without computational considerations of noisy and discrete datasets. Statistical approaches that rely on inferences of mixture models and local dimension estimation require strict geometric assumptions such as linearity [13, 16, 24], and may not handle general scenarios with complex singularities. Recently, approaches from topological data analysis [3, 5, 23], which rely heavily on ingredients from computational [9] and intersection homology [2, 4, 10], are gaining momentum in stratification learning.

Topological approaches transform the smooth and continuous setting favored by topologists to the noisy and discrete setting familiar to computational topologists in practice. In particular, the local structure of a point cloud (sampled from a stratified space) can be described by a multi-scale notion of local homology [3]; and the point cloud data could be clustered based on how the local homology of nearby sampled points map into one another [5]. Philosophically, our main goal is to find a stratification where two points in the same strata (or cluster) can not be distinguished by homological methods, and two points in different strata (different clusters) can be distinguished by homological methods. The majority of the paper will be spent developing a rigorous and computable interpretation of the purposely vague statement “distinguished by homological methods”. Furthermore, we will see that our approach to computing the above stratification applies equally well to sheaves other than those based on local homology. As our work is an interplay between sheaf theory and stratification, we briefly review various notions of stratification before describing our results.

1.1 Stratifications

Given a topological space X , a *topological stratification* of X is a finite filtration, that is, an increasing sequence of closed subspaces $\emptyset = X_{-1} \subset X_0 \subset \dots \subset X_d = X$, such that for each i , $X_i - X_{i-1}$ is a (possibly empty) open i -dimensional topological manifold. See Figure 1 for an example of a *pinched torus*, that is, a torus with points along a geodesic with fixed longitude identified, and a spanning disc glued along the equator.

Ideally, we would like to compute a topological stratification for a given space. However, if we are restricted to using only homological methods, this is a dubious task. Topological

invariants like homology are too rough to detect when a space such as $X_i - X_{i-1}$ is an open i -manifold. A well known example of a homological manifold which is not a topological manifold can be constructed from a homology 3-sphere with nontrivial fundamental group. The suspension of such a space is a homological manifold, but not a topological manifold, since the links of the suspension points have nontrivial fundamental groups [18, page 326]. In this paper, we will avoid the difficult problem of computing topological stratifications, and instead aim to investigate stratifications which can be computed using homological methods. Therefore, we must first consider a definition of stratification which does not rely on topological conditions which are not distinguished by homology. We begin with an extremely loose definition of stratification (Definition 1) which only requires the properties necessary to discuss the constructibility of sheaves (defined in Section 2.2). We will then refine our definition of stratification by placing requirements on the constructibility of certain sheaves (Definition 3).

► **Definition 1.** Given a topological space X , a *stratification* \mathfrak{X} of X is a finite filtration of X by closed subsets X_i :

$$\emptyset = X_{-1} \subset X_0 \subset \cdots \subset X_d = X.$$

We refer to the space $X_i - X_{i-1}$ as *stratum*, denoted by S_i , and a connected component of S_i as a *stratum piece*.

Suppose we have two stratifications of the topological space X , denoted \mathfrak{X} and \mathfrak{X}' . We say that \mathfrak{X} is equivalent to \mathfrak{X}' if each stratum piece of \mathfrak{X} is equal to a stratum piece of \mathfrak{X}' .

► **Definition 2.** Given two inequivalent stratifications of X , \mathfrak{X} and \mathfrak{X}' , we say \mathfrak{X} is *coarser* than \mathfrak{X}' if each stratum piece of \mathfrak{X}' is contained in a stratum piece of \mathfrak{X} .

Homological stratification. There have been several approaches in the topology literature to define homological stratifications. While proving the topological invariance of intersection homology, Goresky and MacPherson defined a type of homological stratification which they call a \bar{p} -stratification [11, Section 4]. There have been several approaches for building on the ideas of Goresky and MacPherson, with applications to computational geometry and topology in mind ([4], [20]). In this paper, we choose to adopt the perspective of homological stratifications found in [21], with a view toward sheaf theoretic generalizations and applications in topological data analysis.

Consider a filtration $\emptyset = X_{-1} \subset X_0 \subset \cdots \subset X_n = X$ of a topological space X . Let \mathcal{L}_i denote the local homology sheaf on the space X_i (see Section 5.1 for the definition of the local homology sheaf). We say that a stratification is a *homological stratification* if \mathcal{L}_n is locally constant (see Section 2.2 for the definition of locally constant) when restricted to $X_i - X_{i-1}$, for each i . A stratification is a *strong homological stratification* if for each i both \mathcal{L}_i and \mathcal{L}_n are locally constant when restricted to $X_i - X_{i-1}$. Finally, a stratification is a *very strong homological stratification* if for each i and every $k \geq i$, \mathcal{L}_k is locally constant when restricted to $X_i - X_{i-1}$. As mentioned in [21], it would be interesting to study the relationship between these definitions of homological stratification. We plan to pursue this in future work. The cohomological stratification given in [20] can be considered as a cohomological analogue of the very strong homological stratification defined above. The utility of Nanda's definition is the extent to which it lends itself to the study of topological properties of individual strata. For example, it can be easily shown that the strata of such a stratification are R -(co)homology manifolds (R being the ring with which the local cohomology is computed). The trade off for using the very strong homological stratification is in the number of local (co)homology

groups which need to be computed. This is by far the most computationally expensive aspect of the algorithm, and the very strong homological stratification requires one to compute new homology groups for each sheaf \mathcal{L}_i . By contrast, our homological stratification only requires the computation of local homology groups corresponding to the sheaf \mathcal{L}_n . In this paper, we choose to study the less rigid (and more computable) notion of homological stratification (see Section 4 for more details on computing homological stratifications).

Sheaf-theoretic stratification. The definition of homological stratification naturally lends itself to generalizations, which we now introduce (while delaying formal definition of constructible sheaves to Section 2.2).

► **Definition 3.** Suppose \mathcal{F} is a sheaf on a topological space X . An \mathcal{F} -stratification (“sheaf-stratification”) of X is a stratification such that \mathcal{F} is constructible with respect to $X = \coprod S_i$. A *coarsest* \mathcal{F} -stratification is an \mathcal{F} -stratification such that \mathcal{F} is not constructible with respect to any coarser stratification.

For general topological spaces, a coarsest \mathcal{F} -stratification may not exist, and may not be unique if it does exist. The main focus of this paper will be proving existence and uniqueness results for certain coarsest \mathcal{F} -stratifications.

1.2 Our contribution

In this paper, we study stratification learning using the tool of constructible sheaves. As a sheaf is designed to systematically track locally defined data attached to the open sets of a topological space, it seems to be a natural tool in the study of stratification based on local structure of the data. Our contributions are three-fold:

1. We prove the existence of coarsest \mathcal{F} -stratifications and the existence and uniqueness of the minimal homogeneous \mathcal{F} -stratification for finite T_0 -spaces (Section 3).
2. We give an algorithm for computing each of the above stratification of a finite T_0 -space based on a sheaf-theoretic language (Section 4).
3. In particular, when applying the local homology sheaf in our algorithm, we obtain a coarsest homological stratification (Section 5.2).

In addition, we envision that our abstraction could give rise to a larger class of stratification beyond homological stratification. For instance, we give an example of a “maximal element-stratification” when the sheaf is defined by considering maximal elements of an open set (see the full version [6]).

Comparison to prior work. This paper can be viewed as a continuation of previous works that aim to adapt the stratification and homology theory of Goresky and MacPherson to the realm of topological data analysis. In [21], Rourke and Sanderson give a proof of the topological invariance of intersection homology on PL homology stratifications, and give an recursive process for identifying a homological stratification (defined in Section 5 of [21]). In [4], Bendich and Harer introduce a persistent version of intersection homology that can be applied to simplicial complexes. In [5], Bendich, Wang, and Mukherjee provide computational approach that yields a stratification of point clouds by computing transfer maps between local homology groups of an open covering of the point cloud. In [20], Nanda uses the machinery of derived categories to study cohomological stratifications based on local cohomology.

Motivated by the results of [20] and [5], we aim to develop a computational approach to the stratifications studied in [21]. Our main results can be summarized as the generalization of homological stratifications of [21] to \mathcal{F} -stratifications, and a proof of existence and uniqueness

of the minimal homogeneous \mathcal{F} -stratification of a finite simplicial complex. When \mathcal{F} is the local homology sheaf, we recover the homological stratification described by [21]. While admitting a similar flavor as [20], our work differs from [20] in several important ways. The most obvious difference is our choice to work with homology and sheaves rather than cohomology and cosheaves. More importantly however, we reduce the number of local homology groups which need to be computed. We will investigate the differences between homological, strong homological, and very strong homological stratifications in future work. Additionally, we plan to build on the results of [5] and [23], and extend the sheaf-theoretic stratification learning perspective described in this paper to the study of stratifications of point cloud data using persistent local homology.

2 Preliminaries

2.1 Compact polyhedra, finite T_0 -spaces and posets

Our broader aim is to compute a clustering of a finite set of points sampled from a compact polyhedron, based on the coarsest \mathcal{F} -stratification of a finite T_0 -space built from the point set. In this paper, we avoid discussion of sampling theory, and assume the finite point set forms the vertex set of a triangulated compact polyhedron. The finite T_0 -space is the set of simplices of the triangulation, with the corresponding partial order. To describe this correspondence in more detail, we first consider the connection between compact polyhedra and finite simplicial complexes. We then consider the correspondence between simplicial complexes and T_0 -topological spaces.

Compact polyhedra and triangulations. A *compact polyhedron* is a topological space which is homeomorphic to a finite simplicial complex. A *triangulation* of a compact polyhedron is a finite simplicial complex K and a homeomorphism from K to the polyhedron.

T_0 -spaces. A T_0 -*space* is a topological space such that for each pair of distinct points, there exists an open set containing one but not the other. Its correspondence with simplicial complex is detailed in [17]:

1. For each finite T_0 -space X there exists a (finite) simplicial complex K and a weak homotopy equivalence $f : |K| \rightarrow X$.
2. For each finite simplicial complex K there exists a finite T_0 -space X and a weak homotopy equivalence $f : |K| \rightarrow X$.

Here, weak homotopy equivalence is a continuous map which induces isomorphisms on all homotopy groups.

T_0 -spaces have a natural partial order. In this paper, we study certain topological properties of a compact polyhedron by considering its corresponding finite T_0 -space. The last ingredient, developed in [1], is a natural partial order defined on a given finite T_0 -space. We can define this partial ordering on a finite T_0 -space X by considering minimal open neighborhoods of each point (i.e. element) $x \in X$. Let X be a finite T_0 -space. Each point $x \in X$ has a minimal open neighborhood, denoted B_x , which is equal to the intersection of all open sets containing x .

$$B_x = \bigcap_{U \in \mathcal{N}_x} U,$$

where \mathcal{N}_x denotes the set of open sets containing x . Since X is a finite space, there are only finitely many open sets. In particular, \mathcal{N}_x is a finite set. So B_x is defined to be the

intersection of finitely many open sets, which implies that B_x is an open neighborhood of x . Moreover, any other open neighborhood V of x must contain B_x as a subset. We can define the partial ordering on X by setting $x \leq y$ if $B_y \subseteq B_x$.

Conversely, we can endow any poset X with the Alexandroff topology as follows. For each element $\tau \in X$, we define a minimal open neighborhood containing τ by $B_\tau := \{\gamma \in X : \gamma \geq \tau\}$. The collection of minimal open neighborhoods for each $\tau \in X$ forms a basis for a topology on X . We call this topology the Alexandroff topology. Moreover, a finite T_0 -space X is naturally equal (as topological spaces) to X viewed as a poset with the Alexandroff topology. Therefore, we see that each partially ordered set is naturally a T_0 -space, and each finite T_0 -space is naturally a partially ordered set. The purpose for reviewing this correspondence here is to give the abstractly defined finite T_0 -spaces a concrete and familiar realization.

Given a finite T_0 -space X with the above partial order, we say $x_0 \leq x_1 \leq \dots \leq x_n$ (where $x_i \in X$) is a maximal chain in X if there is no totally ordered subset $Y \subset X$ consisting of elements $y_j \in Y$ such that $y_0 \leq \dots \leq y_j \leq \dots \leq y_k$ and $\cup_{i=0}^n \{x_i\} \subsetneq Y$. The cardinality of a chain $x_0 \leq x_1 \leq \dots \leq x_n$ is $n + 1$. We say that a finite T_0 -space has dimension m if the maximal cardinality of maximal chains is $m + 1$. An m -dimensional simplicial complex is called *homogeneous* if each simplex of dimension less than m is a face of a simplex of dimension m . Motivated by this definition and the correspondence between simplicial complexes and T_0 -spaces, we say an m -dimensional finite T_0 -space is *homogeneous* if each maximal chain has cardinality $m + 1$.

The correspondences allow us to study certain topological properties of compact polyhedra by using the combinatorial theory of partially ordered sets. In particular, instead of using the more complicated theory of sheaves on the geometric realization $|K|$ of a simplicial complex K , we will continue by studying sheaves on the corresponding finite T_0 -space, denoted by X .

2.2 Constructible sheaves

Intuitively, a sheaf assigns some piece of data to each open set in a topological space X , in a way that allows us to glue together data to recover some information about the larger space. This process can be described as the mathematics behind understanding global structure by studying local properties of a space. In this paper, we are primarily interested in sheaves on finite T_0 -spaces, which are closely related to the cellular sheaves studied in [22], [8], and [20].

Sheaves. Suppose X is a topological space. Let $\mathbf{Top}(X)$ denote the category consisting of objects which are open sets in X with morphisms given by inclusion. Let \mathcal{F} be a contravariant functor from $\mathbf{Top}(X)$ to \mathcal{S} , the category of sets. For open sets $U \subset V$ in X , we refer to the morphism $\mathcal{F}(U \subset V) : \mathcal{F}(V) \rightarrow \mathcal{F}(U)$ induced by \mathcal{F} and the inclusion $U \subset V$, as a *restriction* map from V to U . We say that \mathcal{F} is a *sheaf*³ on X if \mathcal{F} satisfies the following conditions 1-4; a *presheaf* is a functor \mathcal{E} (as above) which satisfies conditions 1-3:

1. $\mathcal{F}(\emptyset) = 0$;
2. $\mathcal{F}(U \subset U) = \text{id}_U$.
3. If $U \subset V \subset W$, then $\mathcal{F}(U \subset W) = \mathcal{F}(U \subset V) \circ \mathcal{F}(V \subset W)$.
4. If $\{V_i\}$ is an open cover of U , and $s_i \in \mathcal{F}(V_i)$ has the property that $\forall i, j, \mathcal{F}((V_i \cap V_j) \subset V_i)(s_i) = \mathcal{F}((V_j \cap V_i) \subset V_j)(s_j)$, then there exists a unique $s \in \mathcal{F}(U)$ such that $\forall i, \mathcal{F}(V_i \subset U)(s) = s_i$.

³ See [8, Chapter 2], [15, Chapter 3] for various introductions to sheaf theory.

There is a useful process known as *sheafification*, which allows us to transform any presheaf into a sheaf. In the setting of finite T_0 -spaces, sheafification takes on a relatively simple form. Let \mathcal{E} be a presheaf on a finite T_0 -space X . Then the sheafification of \mathcal{E} , denoted \mathcal{E}^+ , is given by

$$\mathcal{E}^+(U) = \left\{ f : U \rightarrow \prod_{x \in U} \mathcal{E}(B_x) \mid f(x) \in \mathcal{E}(B_x) \text{ and } f(y) = \mathcal{F}(B_y \subset B_x)(f(x)) \text{ for all } y \geq x \right\}$$

For any presheaf \mathcal{E} , it can be seen that \mathcal{E}^+ is necessarily a sheaf. We only need to know the values $\mathcal{E}(B_x)$ for minimal open neighborhoods B_x , and the corresponding restriction maps between minimal open neighborhoods $\mathcal{E}(B_x \subset B_y)$, in order to define the sheafification of \mathcal{E} . The result is that two presheaves will sheafify to the same sheaf if they agree on all minimal open neighborhoods. We will use this fact several times in Section 3. Unless otherwise specified, for the remaining of this paper, we use X to denote a T_0 -space.

Pull back of a sheaf. For notational convenience, define for each subset $Y \subset X$ the *star* of Y by $\text{St}(Y) := \cup_{y \in Y} B_y$, where B_y is the minimal open neighborhood of $y \in X$. We can think of the star of Y as the smallest open set containing Y . Let X and Y be two finite T_0 -spaces. The following property can be thought of as a way to transfer a sheaf on Y to a sheaf on X through a continuous map $f : X \rightarrow Y$. Let \mathcal{F} be a sheaf on Y . Then the *pull back* of \mathcal{F} , denoted $f^{-1}\mathcal{F}$, is defined to be the sheafification of the presheaf \mathcal{E} which maps an open set $U \subset X$ to $\mathcal{E}(U) := \mathcal{F}(\text{St}(f(U)))$. We can avoid using direct limits in our definition of pull back because each point in a finite T_0 -space has a minimal open neighborhood [8, Chapter 5]. The pull back of \mathcal{F} along an inclusion map $\iota : U \hookrightarrow X$ is called the *restriction* of \mathcal{F} to U , and is denoted $\mathcal{F}|_U$.

Constant and locally constant sheaves. Now we can define classes of well-behaved sheaves, constant and locally constant ones, which we can think of intuitively as analogues of constant functions based on definitions common to algebraic geometry and topology [15]. A sheaf \mathcal{F} is a *constant sheaf* if \mathcal{F} is isomorphic to the pull back of a sheaf \mathcal{G} on a single point space $\{x\}$, along the projection map $p : X \rightarrow x$. A sheaf \mathcal{F} is *locally constant* if for all $x \in X$, there is a neighborhood U of x such that $\mathcal{F}|_U$ (the restriction of \mathcal{F} to U), is a constant sheaf.

► **Definition 4.** A sheaf \mathcal{F} on a finite T_0 -space X is *constructible* with respect to the decomposition $X = \coprod S_i$ of X into finitely many disjoint locally closed subsets, if $\mathcal{F}|_{S_i}$ is locally constant for each i .

3 Main results

In this section we state three of our main results, namely, the existence of \mathcal{F} -stratifications (Proposition 5), the existence of coarsest \mathcal{F} -stratifications (Theorem 6), and the existence and uniqueness of minimal homogeneous \mathcal{F} -stratifications (Theorem 9). Of course, Theorem 6 immediately implies Proposition 5. We choose to include a separate statement of Proposition 5 however, as we wish to illustrate the existence of \mathcal{F} -stratifications which are not necessarily the coarsest. We include proof sketches here and refer to the full version for technical details.

► **Proposition 5.** *Let \mathcal{F} be a sheaf on a finite T_0 -space X . There exists an \mathcal{F} -stratification of X (see Definition 3 and Definition 4).*

Proof sketch. \mathcal{F} is constructible with respect to the decomposition $X = \coprod_{x \in X} x$. ◀

► **Theorem 6.** *Let \mathcal{F} be a sheaf on a finite T_0 -space X . There exists a coarsest \mathcal{F} -stratification of X .*

Proof sketch. We can prove Theorem 6 easily as follows. There are only finitely many stratifications of our space X , which implies that there must be an \mathcal{F} -stratification with a minimal number of strata pieces. Such a stratification must be a coarsest stratification, since any coarser stratification would have fewer strata pieces.

However, the above proof is rather unenlightening if we are interested in computing the coarsest \mathcal{F} -stratification. Therefore we include a constructive proof of the existence of a coarsest \mathcal{F} -stratification which we sketch here. We can proceed iteratively, by defining the top-dimensional stratum to be the collection of points (i.e. elements) so that the sheaf is constant when restricted to the minimal open neighborhoods of the said points. We then remove the top-dimensional stratum from our space, and pull back the sheaf to the remaining points. We proceed until all the points in our space have been assigned to a stratum. We can see that this is a coarsest \mathcal{F} -stratification by arguing that this algorithm, in some sense, maximizes the size of each stratum piece, and thus any coarser \mathcal{F} -stratification is actually equivalent to the one constructed above. We refer the reader to the full version for the details of the above argument. ◀

To uniquely identify a stratification by its properties, we will need to introduce a minimal homogeneous \mathcal{F} -stratification.

► **Definition 7.** Suppose \mathcal{F} is a sheaf on a finite T_0 -space X . A *homogeneous \mathcal{F} -stratification* is an \mathcal{F} -stratification such that for each i , the closure of the stratum S_i in X_i is homogeneous of dimension i (defined in Section 2.1).

We will introduce a lexicographical preorder on the set of homogeneous \mathcal{F} -stratifications of a finite T_0 -space X . Let \mathfrak{X} be a homogeneous \mathcal{F} -stratification of X given by


$$\emptyset = X_{-1} \subset X_0 \subset X_1 \subset \cdots \subset X_n = X$$

We define a sequence $A^{\mathfrak{X}} := \{|X_n|, \dots, |X_i|, \dots, |X_0|\}$, where $|X_i|$ denotes the cardinality of the set X_i . Given two stratifications \mathfrak{X} and \mathfrak{X}' , we say that $\mathfrak{X}' > \mathfrak{X}$ if the first non-zero term of the sequence $A^{\mathfrak{X}'} - A^{\mathfrak{X}} = \{|X'_i| - |X_i|\}$ is positive; $\mathfrak{X} \leq \mathfrak{X}'$ if there are no non-zero terms. Notice that if \mathfrak{X} and \mathfrak{X}' are homogeneous stratifications such that \mathfrak{X} is coarser than \mathfrak{X}' , then we necessarily have that $\mathfrak{X} \leq \mathfrak{X}'$. We say that a stratification \mathfrak{X} is a *minimal homogeneous \mathcal{F} -stratification* if $\mathfrak{X} \leq \mathfrak{X}'$ for every other homogeneous \mathcal{F} -stratification \mathfrak{X}' .

► **Definition 8.** A homogeneous \mathcal{F} -stratification is *minimal* if it is minimal with respect to the lexicographic order on homogeneous \mathcal{F} -stratifications.

There are several examples which illustrate the necessity of introducing *minimal* homogeneous \mathcal{F} -stratifications, rather than studying only coarsest homogeneous \mathcal{F} -stratifications. Consider the simplicial complex K illustrated in Figure 3. Let \mathcal{C} denote the constant sheaf on the corresponding T_0 -space X by assigning the one dimensional vector space k to the star of each simplex σ , $\mathcal{C}(\text{St}\sigma) = k$. For the restriction maps, $\mathcal{C}(\text{St}(\tau) \subset \text{St}(\sigma))$ is an isomorphism for each pair $\sigma < \tau$. Figure 3 shows that there are two coarsest homogeneous \mathcal{C} -stratifications of X . The stratification on the right side of Figure 3 is the minimal homogeneous \mathcal{C} -stratification.

► **Theorem 9.** *Let K be a finite simplicial complex, and X be a finite T_0 -space consisting of the simplices of K endowed with the Alexandroff topology. Let \mathcal{F} be a sheaf on X . There exists a unique minimal homogeneous \mathcal{F} -stratification of X . Moreover, the unique minimal homogeneous \mathcal{F} -stratification is a coarsest homogeneous \mathcal{F} -stratification.*



■ **Figure 3** An example of two inequivalent coarsest homogeneous \mathcal{C} -stratifications, where \mathcal{C} is a constant sheaf on X . The stratification given on the right is the minimal homogeneous \mathcal{C} -stratification.

Proof sketch. The idea for this proof is very similar to that of the Theorem 6. We construct a stratification in a very similar way, with the only difference being that we must be careful to only construct homogeneous strata. The argument for the uniqueness of the resulting stratification uses the observation that this iterative process maximizes the size of the current stratum (starting with the top-dimensional stratum) before moving on to define lower-dimensional strata. Thus the resulting stratification is minimal in the lexicographic order. The top-dimensional stratum of any other minimal homogeneous \mathcal{F} -stratification then must equal the top stratum constructed above, since these must both include the set of top-dimensional simplices, and have maximal size. An inductive argument then shows the stratifications are equivalent. The comments preceding Definition 8 imply that the minimal homogeneous \mathcal{F} -stratification is a coarsest homogeneous \mathcal{F} -stratification. Again, we refer readers to the full version for the remaining details. ◀

4 A sheaf-theoretic stratification learning algorithm

We outline an explicit algorithm for computing a coarsest \mathcal{F} -stratification of a space X given a particular sheaf \mathcal{F} . We give two examples of stratification learning using the local homology sheaf (Section 5) and the sheaf of maximal elements in the full version.

Let X be a finite T_0 -space, equipped with a partial ordering. Instead of using the sheaf-theoretic language of Theorem 9, we frame the computation in terms of X and an “indicator function” δ . For every $x, y \in X$ with a relation $x \leq y$, δ assigns a binary value to the relation. That is, $\delta(x \leq y) = 1$ if the restriction map $\mathcal{F}(B_y \subset B_x) : \mathcal{F}(B_x) \rightarrow \mathcal{F}(B_y)$ is an isomorphism, and $\delta(x \leq y) = 0$ otherwise. We say a pair $w \leq y$ is *adjacent* if $w \leq z \leq y$ implies $z = w$ or $z = y$ (in other words, there are no elements in between w and y). Due to condition 3 in the definition of a sheaf (Section 2.2), δ is fully determined by the values $\delta(w \leq y)$ assigned to each adjacent pair (w, y) . If $a_1 \leq a_2 \leq \dots \leq a_k$ is a chain of adjacent elements (a_i is adjacent to a_{i+1} for each i), we have that $\delta(a_1 \leq a_k) = \delta(a_1 \leq a_2) \cdot \delta(a_2 \leq a_3) \cdot \dots \cdot \delta(a_{k-1} \leq a_k)$. As X is equipped with a finite partially ordering, computing δ can be interpreted as assigning a binary label to the edges of a Hasse diagram associated with the partial ordering (see Section 5 for an example).

For simplicity, we assume that δ is pre-computed, with a complexity of $O(m)$ where m denotes the number of adjacent relations in X . When X corresponds to a simplicial complex K , m is the number of nonzero terms in the boundary matrices of K . δ can, of course, be processed on-the-fly, which may lead to more efficient algorithm. In addition, determining the value of δ is a local computation for each $x \in X$, therefore it is easily parallelizable.

Computing a coarsest \mathcal{F} -stratification. If we are only concerned with calculating a coarsest \mathcal{F} -stratification as described in Theorem 6, we may use the algorithm below.

1. Set $i = 0$, $d_0 = \dim X$, $X_{d_0} = X$, and initialize $S_j = \emptyset$, for all $0 \leq j \leq d_0$.
2. While $d_i \geq 0$, do
 - a. For each $x \in X_{d_i}$, set $S_{d_i} = S_{d_i} \cup x$ if $\delta(w \leq y) = 1$, \forall adjacent pairs $w \leq y$ in $B_x \cap X_{d_i}$
 - b. Set $d_{i+1} = \dim(X_{d_i} - S_{d_i})$
 - c. Define $X_{d_{i+1}} = X_{d_i} - S_{d_i}$
 - d. Set $i = i + 1$
3. Return S

Here, i is the step counter; d_i is the dimension of the current strata of interest; the set S_{d_i} is the stratum of dimension d_i . d_i decreases from $\dim(X)$ to 0. To include an element x to the current stratum S_{d_i} , we need to check δ for adjacent relations among all x 's cofaces.

Computing the unique minimal homogeneous \mathcal{F} -stratification. If we would like to obtain the unique minimal homogeneous \mathcal{F} -stratification, then we need to modify step 2a. Let $c(x, i) = 1$ if all maximal chains in X_{d_i} containing x have cardinality d_i , and $c(x, i) = 0$ otherwise. Then the modified version of 2.a. is:

- 2.a. For each $x \in X_{d_i}$, set $S_{d_i} = S_{d_i} \cup x$ if $\delta(w \leq y) = 1, \forall$ adj. pairs $w \leq y$ in $B_x \cap X_{d_i}$ and $c(x, i) = 1$

5 Stratification learning with local homology sheaf

5.1 Local homology sheaf

For a finite T_0 -space X , consider the chain complex $C_\bullet(X)$, where $C_p(X)$ denotes a free R -module generated by $(p + 1)$ -chains in X , with chain maps $\partial_p : C_p(X) \rightarrow C_{p-1}(X)$ given by

$$\partial_p(a_0 \leq \dots \leq a_p) = \sum (-1)^i (a_0 \leq \dots \leq \hat{a}_i \leq \dots \leq a_p)$$

where \hat{a}_i means that the element a_i is to be removed from the chain.

We would like to remark on the decision to refer to this sheaf as the local homology sheaf. If X is a more general topological space (CW space, simplicial complex, manifold, etc), then the *local homology* of X at $x \in X$ is defined to be the direct limit of relative homology $H_\bullet(X, X - x) := \varinjlim H_\bullet(X, X - U)$, where the direct limit is taken over all open neighborhoods U of x with the inclusion partial order [19, page 196]. In our setting, the local homology of X (a finite T_0 -space) at a point $x \in X$ is given by $H_\bullet(X, X - B_x)$. Here we avoid using notions of direct limit by working with topological spaces that have minimal open neighborhoods. This motivates us to refer to the sheaf defined by relative homology $H_\bullet(X, X - U)$ for each open set U (see Theorem 10), as the *local homology sheaf*⁴. The following theorem, though straightforward, provides justification for applying the results of Section 4 to local homology computations (see the full version for its proof).

► **Theorem 10.** *The functor \mathcal{L} from the category of open sets of a finite T_0 -space to the category of graded R -modules, defined by*


$$\mathcal{L}(U) := H_\bullet(X, X - U)$$

where R is the ring of coefficients of the relative homology, is a sheaf on X .

⁴ See [7] for an interesting approach to the computation of homology groups of finite T_0 -spaces.

5.2 An example of stratification learning using local homology sheaf

If X is a T_0 -space corresponding to a simplicial complex K , then the local homology groups in Section 5.1 are isomorphic to the simplicial homology groups of K . We now give a detailed example of stratification learning using local homology sheaf for the sundial example from Figure 2. We will abuse notation slightly, and use K to denote the finite T_0 -space consisting of elements which are open simplices corresponding to the triangulated sundial (Figure 4). We choose this notation so that we can describe our T_0 -space using the more familiar language of simplicial complexes. For a simplex $\sigma \in K$, its minimal open neighborhood B_σ is its *star* consisting of all cofaces of σ , $\text{St}(\sigma) = \{\tau \in K \mid \sigma \leq \tau\}$. The *closed star*, $\overline{\text{St}}(\sigma)$, is the smallest subcomplex that contains the star. The *link* consists of all simplices in the closed star that are disjoint from the star, $\text{Lk}(\sigma) = \{\tau \in \overline{\text{St}}(\sigma) \mid \tau \cap \text{St}(\sigma) = \emptyset\}$. K is equipped with a partial order based on face relations, where $x < y$ if x is a proper face of y . This partial order gives rise to a Hasse diagram illustrated in Figure 5.




■ **Figure 4** A triangulated sundial and its stratification based on the local homology sheaf.

A sheaf on K can be considered as a labeling of each vertex in the Hasse diagram with a set and each edge with a morphism between the corresponding sets. Consider the local homology sheaf \mathcal{L} on K which takes each open set ⁵ $U \subset K$ to $H_\bullet(K, K - U) \cong \tilde{H}_\bullet(K/(K - U)) \cong \tilde{H}_\bullet(\text{Cl}(U)/\text{Lk}(U)) \cong H_\bullet(\text{Cl}(U), \text{Lk}(U))$, where $\text{Lk}(U) := \text{Cl}(U) - U$. The above isomorphisms follow from excision and the observation that $K - U$ (resp. $\text{Lk}(U)$) is a closed subcomplex of K (resp. $\text{Cl}(U)$), and therefore $(K, K - U)$ and $(\text{Cl}(U), \text{Lk}(U))$ form good pairs (see [14, page 124]). Our algorithm described in Section 4 can then be interpreted as computing local homology sheaf associated with each vertex in the Hasse diagram, and determining whether each edge in the diagram is an isomorphism. Our algorithm works by considering an element σ in the Hasse diagram to be in the top-dimensional strata if all of the edges above σ are isomorphisms, that is, if $\mathcal{L}(\sigma < \tau)$ is an isomorphism for all pairs $\sigma < \tau$.


As illustrated in Figure 5, first, we start with the 2-simplexes. Automatically, we have that \mathcal{L} is constant when restricted to any 2-simplex, and gives homology groups isomorphic to the reduced homology of a 2-sphere. For instance, the local homology groups of the 2-simplex $\sigma = [0, 1, 3]$ is isomorphic to the reduced homology of a 2-sphere, $H_\bullet(\overline{\text{St}}(\sigma), \text{Lk}(\sigma)) \cong \tilde{H}_\bullet(S^2)$.

Second, we consider the restriction of \mathcal{L} to the minimal open neighborhood of a 1-simplex. For instance, consider the 1-simplex $[1, 3]$; $B_{[1,3]} = [1, 3] \cup [0, 1, 3]$. It can be seen that $\text{Lk}(B)_{[1,3]} = [0] \cup [3] \cup [1] \cup [0, 3] \cup [0, 1]$, and $H_\bullet(\text{Cl}(B_{[1,3]}), \text{Lk}(B_{[1,3]}))$ is isomorphic to the reduced homology of a single point space. Therefore the restriction map $\mathcal{L}(B_{[1,3]}) \rightarrow \mathcal{L}(B_{[0,1,3]})$ is not an isomorphism (illustrated as a dotted blue line in Figure 5). On the other hand, let us consider the 1-simplex $[0, 3]$, where $B_{[0,3]} = [0, 3] \cup [0, 1, 3] \cup [0, 2, 3]$. We have that

⁵ In the finite simplicial setting, U is the support of a union of open simplices in K .



■ **Figure 5** The Hasse diagram of the triangulated sundial. For any two adjacent simplices $\tau < \sigma$, an edge between τ and σ in the diagram is solid red if $\mathcal{L}(B_\tau) \rightarrow \mathcal{L}(B_\sigma)$ is an isomorphism; otherwise it is in dotted blue. On the right of each simplex τ is either a point, a sphere, or the wedge of two spheres, chosen so that $\mathcal{L}(B_\tau)$ is isomorphic to the reduced homology of the associated space.



■ **Figure 6** The Hasse diagram after the top dimensional stratum has been removed. We can consider this the beginning of the second iteration of the algorithm in Section 4.

$\text{Lk}(B_{[0,3]}) = [0] \cup [1] \cup [2] \cup [3] \cup [0, 1] \cup [0, 2] \cup [1, 3] \cup [2, 3] \cup [1, 3]$. Therefore $\mathcal{L}(B_{[0,3]})$ is isomorphic to the reduced homology of a 2-sphere. Moreover, both of the restriction maps corresponding to $B_{[0,1,3]} \subset B_{[0,3]}$ and $B_{[0,2,3]} \subset B_{[0,3]}$ are isomorphisms (illustrated as solid red lines in Figure 5). This implies that $[0, 3] \in S_2 = X_2 - X_1$. Alternatively, we can consider the simplex $[0, 1]$ and see that $\mathcal{L}(B_{[0,1]}) \cong \tilde{H}_\bullet(\overline{\text{St}}([0, 1])/\text{Lk}([0, 1])) \cong \tilde{H}_\bullet(S^2 \vee S^2)$, the reduced homology of the wedge of two spheres. Therefore, for any 2-simplex τ , the restriction map $\mathcal{L}([0, 1] < \tau)$ can not be an isomorphism. We conclude that $[0, 1]$ is not contained in the top dimensional stratum. If we continue, we see that the top dimensional stratum is given by $S_2 = [0, 1, 3] \cup [0, 1, 2] \cup [0, 2, 3] \cup [0, 1, 4] \cup [0, 2] \cup [0, 3]$, see Figure 4.

Next, we can calculate the stratum $S_1 = X_1 - X_0$ by only considering restriction maps whose codomain is not contained in S_2 (see Figure 6). We get $S_1 = [0, 1] \cup [1, 3] \cup [1, 2] \cup [2, 3] \cup [0, 4] \cup [1, 4] \cup [2] \cup [3] \cup [4]$, which is visualized in Figure 4. Finally, the stratum $S_0 = X_0$ consists of the vertices which have not been assigned to any strata. So $S_0 = [0] \cup [1]$.

We observe that (for this example) the coarsest \mathcal{L} -stratification we calculated is actually the unique minimal homogeneous \mathcal{L} -stratification. We will investigate this coincidence for \mathcal{L} -stratifications elsewhere, to verify if a coarsest \mathcal{L} -stratification is automatically homogeneous or minimal. For low-dimensional examples, we observe that the local homology based stratification we recover is actually a topological stratification. For instance, the coarsest \mathcal{L} -stratification we obtain from our algorithm coincides with the stratification given in Figure 1 (for a suitable triangulation of the pinched torus). In general, local homology does not carry enough information to recover a stratification into manifold pieces, and examples exist in higher dimensions where \mathcal{L} -stratification are not topological stratifications.

6 Discussion

A triangulation of the pinched torus in Figure 1 can be thought of as a stratification of the space, by defining the d -dimensional stratum to be the collection of d -dimensional simplices. However, this stratification is too “fine” in a sense, as it breaks up the underlying space into too many pieces, resulting in each stratum piece retaining relatively little information about the structure of the total space. The results of this paper can be interpreted as a method for computing a coarsening of the stratification obtained from the triangulation of our underlying space, using sheaf-theoretic techniques (in particular, homological techniques) to determine when two simplices should belong to the same coarser stratum.

There are two key features of the sheaf-theoretic stratification learning algorithm which should be highlighted. The first feature is that we avoid computations which require the sheafification process. At first glance this may be surprising to those not familiar with cellular sheaves, since constructible sheaves can not be defined without referencing sheafification, and our algorithm builds a stratification for which a given sheaf is constructible. In other words, each time we want to determine the restriction of a sheaf to a subspace, we need to compute the sheafification of the presheaf referenced in the definition of the pull back of a sheaf (Section 2.2). We can avoid this by noticing two facts. Suppose \mathcal{E} is a presheaf and \mathcal{E}^+ is the sheafification of \mathcal{E} . First, in the setting of finite T_0 -spaces, we can deduce if \mathcal{E}^+ is constant by considering how it behaves on minimal open neighborhoods. Second, the behavior of \mathcal{E}^+ will agree with the behavior of the presheaf \mathcal{E} on minimal open neighborhoods. Symbolically, this is represented by the equalities $\mathcal{E}^+(B_x) = \mathcal{E}(B_x)$ and $\mathcal{E}^+(B_w \subset B_x) = \mathcal{E}(B_w \subset B_x)$ for all pairs of minimal open neighborhoods $B_w \subset B_x$ (where B_x is a minimal open neighborhood of x , and B_w is a minimal open neighborhood of an element $w \in B_x$). Therefore, we can determine if \mathcal{E}^+ is constant, locally constant, or constructible, while only using computations involving the presheaf \mathcal{E} applied to minimal open neighborhoods.

The second feature of our algorithm (which is made possible by the first) is that the only sheaf-theoretic computation required is checking if $\mathcal{F}(B_w \subset B_x)$ is an isomorphism for each pair $B_w \subset B_x$ in our space. This is extremely relevant for implementations of the algorithm, as it minimizes the number of expensive computations required to build an \mathcal{F} -stratification. For example, if our sheaf is the local homology sheaf, we will only need to compute the restriction maps between local homology groups of minimal open neighborhoods. The computation of local homology groups can therefore be distributed and computed independently. Additionally, once we have determined whether the local homology restriction maps are isomorphisms, we can quickly compute a coarsest \mathcal{F} -stratification, or a minimal homogeneous \mathcal{F} -stratification, without requiring any local homology groups to be recomputed.

There are several interesting questions related to \mathcal{F} -stratifications that we will investigate in the future. We are interested in studying \mathcal{F} -stratifications for natural sheaves other than the local homology sheaf, on finite T_0 -spaces. The primary objective would be to find easily computable sheaves that yield intuitive stratifications relevant for manifold learning algorithms and of interest to data analysts. We will also study the stability of \mathcal{F} -stratifications under refinements of triangulations of polyhedra. In this direction, it would be interesting to view \mathcal{F} -stratifications from the perspective of persistent homology. If we are given a point cloud sampled from a compact polyhedron, it would be natural to ask about the convergence of \mathcal{F} -stratifications and the properties of the strata under a filtration of the simplicial complex. Finally, we are also intrigued by the results of [7], and possible implementations of our algorithm using spectral sequences.

References

- 1 Pavel Sergeevich Alexandroff. Diskrete Räume. *Mathematicheskii Sbornik*, 2:501–518, 1937.
- 2 Paul Bendich. *Analyzing Stratified Spaces Using Persistent Versions of Intersection and Local Homology*. PhD thesis, Duke University, 2008.
- 3 Paul Bendich, David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Dmitriy Morozov. Inferring local homology from sampled stratified spaces. *IEEE Symposium on Foundations of Computer Science*, pages 536–546, 2007.
- 4 Paul Bendich and John Harer. Persistent intersection homology. *Foundations of Computational Mathematics*, 11:305–336, 2011.
- 5 Paul Bendich, Bei Wang, and Sayan Mukherjee. Local homology transfer and stratification learning. *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1355–1370, 2012.
- 6 Adam Brown and Bei Wang. Sheaf-theoretic stratification learning. *arXiv:1712.07734*, 2017.
- 7 Nicolás Cianci and Miguel Ottina. A new spectral sequence for homology of posets. *Topology and its Applications*, 217:1 – 19, 2017.
- 8 Justin Curry. *Sheaves, Cosheaves and Applications*. PhD thesis, University of Pennsylvania, 2014.
- 9 Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010.
- 10 Mark Goresky and Robert MacPherson. Intersection homology I. *Topology*, 19:135–162, 1982.
- 11 Mark Goresky and Robert MacPherson. Intersection homology II. *Inventiones Mathematicae*, 71:77–129, 1983.
- 12 Mark Goresky and Robert MacPherson. *Stratified Morse Theory*. Springer-Verlag, 1988.
- 13 Gloria Haro, Gregory Randall, and Guillermo Sapiro. Stratification learning: Detecting mixed density and dimensionality in high dimensional point clouds. *Advances in Neural Information Processing Systems (NIPS)*, 17, 2005.
- 14 Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- 15 Frances Clare Kirwan. *An introduction to intersection homology theory*. Chapman & Hall/CRC, 2006.
- 16 Gilad Lerman and Teng Zhang. Probabilistic recovery of multiple subspaces in point clouds by geometric lp minimization. *Annals of Statistics*, 39(5):2686–2715, 2010.
- 17 Michael C. McCord. Singular homology groups and homotopy groups of finite topological spaces. *Duke Mathematical Journal*, 33:465–474, 1978.
- 18 Washington Mio. Homology manifolds. *Annals of Mathematics Studies (AM-145)*, 1:323–343, 2000.
- 19 James R. Munkres. *Elements of algebraic topology*. Addison-Wesley, Redwood City, CA, USA, 1984.
- 20 Vidit Nanda. Local cohomology and stratification. *ArXiv:1707.00354*, 2017.
- 21 Colin Rourke and Brian Sanderson. Homology stratifications and intersection homology. *Geometry and Topology Monographs*, 2:455–472, 1999.
- 22 Allen Dudley Shepard. *A Cellular Description Of The Derived Category Of A Stratified Space*. PhD thesis, Brown University, 1985.
- 23 Primoz Skraba and Bei Wang. Approximating local homology from samples. *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 174–192, 2014.
- 24 René Vidal, Yi Ma, and Shankar Sastry. Generalized principal component analysis (GPCA). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27:1945–1959, 2005.
- 25 Shmuel Weinberger. *The topological classification of stratified spaces*. University of Chicago Press, Chicago, IL, 1994.