



Local Cohomology and Stratification

Vidit Nanda¹

Received: 14 August 2017 / Revised: 9 March 2019 / Accepted: 29 April 2019 / Published online: 13 June 2019
© The Author(s) 2019

Abstract

We outline an algorithm to recover the canonical (or, coarsest) stratification of a given finite-dimensional regular CW complex into cohomology manifolds, each of which is a union of cells. The construction proceeds by iteratively localizing the poset of cells about a family of subposets; these subposets are in turn determined by a collection of cosheaves which capture variations in cohomology of cellular neighborhoods across the underlying complex. The result is a nested sequence of categories, each containing all the cells as its set of objects, with the property that two cells are isomorphic in the last category if and only if they lie in the same canonical stratum. The entire process is amenable to efficient distributed computation.

Keywords Canonical stratification · Local cohomology

Mathematics Subject Classification 52S60 · 55N30 · 18E35

1 Introduction

Setting aside technicalities for the moment, an n -dimensional **stratification** of a given topological space X is a nested sequence of closed subspaces

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

where each successive difference $X_d - X_{d-1}$ resembles a (possibly empty) d -dimensional manifold whose connected components are the d -dimensional *strata*. The details which have been suppressed are designed to guarantee the uniformity of small neighborhoods around points in a given stratum. Stratified spaces are of fundamental importance in any branch of mathematics which seriously confronts singularities.


Communicated by Herbert Edelsbrunner.

✉ Vidit Nanda
nanda@maths.ox.ac.uk


¹ University of Oxford, Oxford, UK

Every topological manifold admits a straightforward stratification into its connected components; less trivially, the following spaces all admit, and are often analyzed using, stratifications: finite CW complexes, quotients of properly discontinuous Lie group actions on smooth manifolds, (semi)algebraic varieties and (sub)analytic sets.


The two-dimensional singular space below—let us call it Y —is built by pinching a torus along a meridian and attaching a disk across an equator:




Any regular CW structure, such as the illustrated decomposition into little squares, constitutes a stratification of Y where d -dimensional strata are precisely the d -cells; passing to a subdivision further refines this stratification in the sense that every new cell is entirely contained in the interior of an old cell. On the other hand, one can discover a much coarser stratification by examining the topology of small neighborhoods around points of Y . Up to homeomorphism, these fall into three different classes depending on whether the central point is at the pinch, on the singular equator, or on one of the manifold-like two-dimensional regions:



The neighborhoods above deformation retract onto their central vertices and are therefore contractible; however, their one-point compactifications (obtained by collapsing their boundaries to points) are new stratified spaces with potentially interesting topology. The compactified neighborhood around the pinch point is homeomorphic to two 2-spheres joined at their north and south poles with a spanning disk across the middle. The compactified neighborhood of any point in the singular equator resembles a 2-sphere whose interior has been partitioned into two by a disk. And finally, the compactified neighborhood around any point in either of the two-dimensional regions is homeomorphic to an ordinary 2-sphere:



Our main result here involves algorithmically recovering the coarsest stratification of a finite regular CW complex where all strata are unions of cells. This is called the **canonical stratification** of the complex; its existence and uniqueness for a special class of spaces (called *pseudomanifolds*) plays a central role in Goresky and MacPherson’s proof of the topological invariance of intersection homology [14, Sec 4]. Our argument, much like theirs, has an intuitive geometric core but invokes algebraic and categorical machinery. For the purposes of this introductory section, we focus on geometry and ask: Given a finite cellulation of Y , how might one identify the canonical strata and determine which cells lie in each canonical stratum, as shown below?



(Since the pinch point has a different compactified neighborhood than a generic point on the singular equator, it must constitute a separate stratum.)

In light of the discussion above, one hopes to recover canonical strata by clustering together cells whenever they exhibit similar compactified neighborhoods. But already in this example, we encounter two significant difficulties: First, the compactified neighborhoods do not distinguish cells in the two 2-strata from each other. The second difficulty is somewhat subtler—although the compactified neighborhoods of cells in the 0-stratum and 1-stratum are not homeomorphic, both are homotopy-equivalent to a wedge of two 2-spheres. Therefore, they cannot be distinguished by weaker, more computable topological invariants such as cohomology. We tackle the first problem by constructing a *complex of cosheaves* which encodes how local topology varies across cells, and we bypass the second problem by working exclusively in the category of *cohomologically stratified spaces*.

Consider the task of determining the canonical stratification of a finite-dimensional regular CW complex X into R -cohomology manifolds, where R is a fixed non-trivial commutative ring with unity. In this case, our complex of cosheaves assumes a particularly appealing form: It is a functor

$$L^\bullet : \mathbf{Fc}(X) \rightarrow \mathbf{Ch}(R)$$

from the poset of cells in X (where $x > y$ denotes that y is a face of x , or equivalently, that x is a co-face of y) to the category of lower-bounded R -module cochain complexes. To each cell x of X , it assigns a cochain complex $L^\bullet(x)$:

$$L^0(x) \xrightarrow{\beta_x^0} L^1(x) \xrightarrow{\beta_x^1} L^2(x) \xrightarrow{\beta_x^2} \dots ,$$

where $L^d(x)$ is the free R -module generated by all d -dimensional co-faces of x —in particular, the module $L^d(x)$ is trivial for $d < \dim x$ and $d > \dim X$. The co-differentials β_x^\bullet are inherited verbatim from incidence degrees among cells of X taking values in R . Thus, $L^\bullet(x)$ computes the (reduced) cohomology of a compactified small

open neighborhood around x in X . Given another cell y satisfying $x \geq y$, there is an inclusion map

$$L^\bullet(x \geq y) : L^\bullet(x) \hookrightarrow L^\bullet(y)$$

arising from the fact that co-faces of x are also co-faces of y . Here is a simple version of our main result. (The full statement has been recorded in Theorem 5.10.)

Theorem 1.1 *There is a category \mathcal{S} obtained from $\mathbf{Fc}(X)$ by formally inverting a particular subset of those face relations $x \geq y$ for which $L^\bullet(x \geq y)$ induces isomorphisms on cohomology; two cells lie in the same canonical stratum of X if and only if they are isomorphic in \mathcal{S} .*

Before proceeding to the L^\bullet -induced definition of \mathcal{S} and a proof of this theorem, we highlight three salient features. First, the construction of \mathcal{S} is algorithmic and isomorphism classes of its objects are explicitly (and efficiently) computable. Second, the non-isomorphisms in \mathcal{S} reveal when a canonical stratum contains another in its boundary. And third, our construction actually yields a nested sequence of intermediate categories

$$\mathcal{S}_0 \hookrightarrow \mathcal{S}_1 \hookrightarrow \cdots \hookrightarrow \mathcal{S}_{\dim X} = \mathcal{S},$$

each of which receives a functor from $\mathbf{Fc}(X)$ compatible with the previous ones, and all the canonical strata of dimension exceeding $(\dim X - d)$ may be recovered from certain isomorphism classes in \mathcal{S}_d .

This paper is organized as follows: Sect. 2 tersely collects relevant background material involving stratifications and cosheaves, while Sect. 3 describes the local cohomology complex of cosheaves L^\bullet associated with a regular CW complex X . In Sect. 4 we use L^\bullet to define an initial functor $\mathbf{Fc}(X) \rightarrow \mathcal{S}_0$ via categorical localization and prove that the top-dimensional canonical strata of X correspond bijectively with isomorphism classes of (images of) top-cells in \mathcal{S}_0 . Section 5 contains the heart of our argument: It concludes the proof of Theorem 1.1 by describing the inductive construction of $\mathbf{Fc}(X) \rightarrow \mathcal{S}_d$ from the previous $\mathbf{Fc}(X) \rightarrow \mathcal{S}_{d-1}$. The final Sect. 6 may be read independently of the preceding ones—it is devoted to computational matters and includes an efficient distributed algorithm for recovering canonical strata in practice.

Context

Our goal is to provide a principled topological pre-processor for various *manifold learning* algorithms. The fundamental aim of manifold learning is to automatically infer the intrinsic dimension of a compact submanifold M (often with boundary) of Euclidean space from a finite point sample $P \subset \mathbb{R}^n$ lying on, or at least near, M . An enormous amount of work has been done in this field, and we will make no attempt to even summarize it here—the curious reader is referred to [19] and its myriad citations. In most cases, one constructs a graph around the points in P by inserting all edges of Euclidean length bounded above by a parameter threshold $\epsilon > 0$. Shortest paths

along this graph then serve as proxies for intrinsic geodesics of M and are used to find distance-preserving projections to lower-dimensional subspaces of \mathbb{R}^n via spectral methods.

When M is not a single manifold, but rather a union of different manifolds across several dimensions, our standard manifold learning methods necessarily result in overfitting along regions where M has low intrinsic dimension. Theorem 1.1 provides a remedy by conferring the ability to partition the Delaunay triangulation [23, Ch 5] around P at radius ϵ into distinct clusters, each of which is guaranteed to lie on a cohomology manifold of known dimension.¹ These clusters may then be independently subjected to standard manifold learning techniques. Perhaps more important from a practical perspective is the fact that the computations involved in partitioning P are easily distributed across several processors. And as a straightforward by-product, one also gains the capacity to identify (those points of P which lie in) every anomalous, singular region of M .

There now exists a substantial body of work where filtered (co)homology groups of cell complexes built around data points play a central role [3,8,12]. With the knowledge of canonical strata of such cell complexes comes the prospect of efficiently computing far more refined topological invariants such as intersection cohomology [13] groups. Also accessible, thanks to suitable flavors of Morse theory [21,22], is the ability to cluster canonical d -strata by $(d + 1)$ -dimensional cobordisms internal to the ambient dataset.

2 Preliminaries

Our primary references for stratification theory are [15,18], and [30]; for categorical localization see [7] and [10]; and for cellular (co)sheaves see [4] and [25]. The interplay between cell structures and general stratifications on the same underlying space has been explored thoroughly in [26]. Here we will be concerned entirely with cohomologically stratified spaces, which the reader may have encountered before in [14, Sec 4.1], [11, Sec 3.3] or [24, Sec 5].

2.1 Cohomological Stratifications

Fix a non-trivial commutative ring R with identity, and recall that an **R -cohomology manifold** of dimension $d \geq 0$ is a locally compact Hausdorff topological space Z where each point z has an open neighborhood $U_z \subset Z$ whose compactly supported R -valued cohomology agrees with that of d -dimensional Euclidean space \mathbb{R}^d , i.e.,

$$H_c^\bullet(U_z; R) \simeq \begin{cases} R & \text{if } \bullet = d, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

It is essential to use compactly supported (rather than ordinary singular) cohomology to detect d -dimensionality, since the usual cohomology of a d -dimensional open ball

¹ The probabilistic version of such a framework is described in [1].

agrees with that of a point regardless of d . Readers unfamiliar with this notion may safely replace the compactly supported cohomology of any open set encountered henceforth by the standard cohomology of its one-point compactification relative to the additional point:

$$H^\bullet(U_z; R) \simeq H^\bullet(U_z \cup \{\infty\}, \{\infty\}; R).$$

Definition 2.1 An n -dimensional (**cohomological**) **stratification** of a locally compact Hausdorff topological space X is an ascending sequence of closed subspaces

$$\emptyset = X_{-1} \subset X_0 \subset \dots \subset X_{n-1} \subset X_n = X,$$

where the connected components of each difference $(X_d - X_{d-1})$, called d -dimensional **strata**, are required to obey the following axioms.

- *Frontier* if a stratum σ intersects the closure of another stratum τ , then in fact σ is completely contained in the closure of τ and $\dim \tau \geq \dim \sigma$ (with equality occurring if and only if $\tau = \sigma$). This relation, denoted by $\tau \succeq \sigma$ henceforth, forms a partial order on the set of all strata.
- *Link* for each d -stratum σ , there exists an $(n - d - 1)$ -dimensional stratified space $L = L(\sigma)$, called the *link* of σ :

$$\emptyset = L_{-1} \subset L_0 \subset \dots \subset L_{n-d-1} = L,$$

so that every open neighborhood around a point p in σ admits a *basic* open sub-neighborhood $U_p \subset X$ with the following structure. The intersections $U_p \cap X_i$ are empty for $i < d$, while for $d \leq i \leq n$ there are stratified quasi-isomorphisms of compactly supported singular cochain complexes (of R -modules):

$$C_c^\bullet(\mathcal{C}L_{i-d-1} \times \mathbb{R}^d) \xrightarrow{\simeq} C_c^\bullet(U_p \cap X_i), \tag{2}$$

where $\mathcal{C}L_\bullet$ denotes the open cone² on L_\bullet .

By quasi-isomorphism here (and elsewhere) we mean a cochain map which induces isomorphisms on all cohomology groups. The additional requirement that the maps from (2) be stratified means precisely that they behave naturally in three separate ways: They should be *filtered*, *stratum-preserving* and *refinable* as explained below.

- (1) They should respect (contravariant maps induced by) the inclusions $L_\bullet \subset L_{\bullet+1}$ in their domains and $X_\bullet \subset X_{\bullet+1}$ their codomains so that the following squares commute:

$$\begin{array}{ccc} C_c^\bullet(\mathcal{C}L_{i-d-1} \times \mathbb{R}^d) & \xrightarrow{\simeq} & C_c^\bullet(U_p \cap X_i) \\ \uparrow & & \uparrow \\ C_c^\bullet(\mathcal{C}L_{i-d} \times \mathbb{R}^d) & \xrightarrow{\simeq} & C_c^\bullet(U_p \cap X_{i+1}) \end{array}$$

² The open cone on Z is the quotient of $Z \times [0, 1)$ by the relation which identifies $(z, 0)$ with $(z', 0)$ for any z and z' in Z ; by convention, the open cone on the empty set is the one-point space.

(2) They should preserve strata in the sense that there exist surjective set-maps

$$\Phi_i : \{(i - d - 1) - \text{strata of } L\} \rightarrow \{i - \text{strata } \tau \succeq \sigma \text{ of } X\}$$

so that for each i -stratum $\tau \succeq \sigma$, the cochain maps from (2) restrict to quasi-isomorphisms

$$C_c^\bullet(\mathcal{C}\Phi_i^{-1}(\tau) \times \mathbb{R}^d) \xrightarrow{\simeq} C_c^\bullet(U_p \cap \tau).$$


(3) They should refine to smaller basic neighborhoods V_p so that the following triangle of cochain maps commutes:

$$\begin{array}{ccc} C_c^\bullet(\mathcal{C}L_{i-d-1} \times \mathbb{R}^d) & \xrightarrow{\simeq} & C_c^\bullet(U_p \cap X_i) \\ & \searrow \simeq & \downarrow \\ & & C_c^\bullet(V_p \cap X_i) \end{array}$$

Here the vertical map is induced by the inclusion $V_p \subset U_p$, and since the triangle commutes, this map is forced to also be a quasi-isomorphism.

The best way to acquire intuition for these constraints is to examine (pictures of) links associated with singular strata in low-dimensional spaces.

Example 2.2 The decomposition of the space Y from the Introduction into the pinch point, the equatorial circle, a disk 2-stratum and a toral 2-stratum is a cohomological stratification. The pinch point lies at the frontier of the equatorial circle, which in turn lies at the frontier of both 2-strata. The link L of the singular equator (which we will call σ here) is a three-point space. An open cone on L is therefore the union of three half-open intervals along a common boundary point. The product of this cone with the real line is homeomorphic to small neighborhoods around points in σ :



From such a homeomorphism, one obtains the desired cohomological equivalence (2). This equivalence is

- filtered, provided we stratify the open cone $\mathcal{C}L$ with one 0-stratum and three 1-strata (and the real line by a single 1-stratum),
- stratum-preserving, provided that the map Φ_2 sends two of the points in L to the toral 2-stratum and the third point to the disk 2-stratum, and
- refinable, because one can shrink the neighborhood around σ by a small amount while still preserving the homeomorphism.

Remark 2.3 To see a non-example of a cohomological stratification, try combining the pinch point p of Y with the singular equator σ into a single 1-stratum. The refinability constraint will not be satisfied in this case. In particular, given any point $q \in \sigma$ near p with two basic neighborhoods $V_q \subset U_q$ satisfying $p \in U_q - V_q$, the inclusion of V_q into U_q fails to induce an isomorphism on two-dimensional compactly supported cohomology. Instead, one obtains a rank one map $R^2 \rightarrow R^2$.

The following observations concern technical aspects of Definition 2.1.

- (1) By a theorem of Eilenberg and Zilber [9], links are well defined (only) up to filtered quasi-isomorphism type: We may replace L by a filtered cochain complex by substituting a tensor product for the left side of (2).
- (2) The link axiom (when $i = d$) guarantees that each d -stratum is indeed a d -dimensional R -cohomology manifold.
- (3) This definition, unlike similar notions which typically appear in the intersection homology literature, does not require X to be a **pseudomanifold** (see [14, Sec 1.1] for instance). In other words, X_{n-1} need not equal X_{n-2} .

Every finite-dimensional regular³ CW complex admits a cohomological stratification whose strata are the cells.

Definition 2.4 The **skeletal stratification** of a finite-dimensional regular CW complex X is defined as follows: Writing the face partial order among cells as \geq , we recall that each cell y of X has

- an open star $\mathbf{st}(y)$ containing all cells x which satisfy $x \geq y$, and
- a link $\mathbf{lk}(y)$, containing all cells x that share a co-face but no face with y .

The d -dimensional skeletal strata are the d -cells of X , and the frontier partial order coincides with \geq . Since $\mathbf{st}(y)$ of a d -cell y is simultaneously homeomorphic to all sufficiently small neighborhoods around points in y and to $\mathcal{C}\mathbf{lk}(y) \times \mathbb{R}^d$, the link $L(y)$ in the sense of Definition 2.1 is precisely $\mathbf{lk}(y)$.

We call one stratification a *coarsening* of another whenever each stratum of the former is a union of strata of the latter. All stratifications encountered henceforth will be coarsenings of the skeletal stratification for a fixed finite-dimensional regular CW complex X .

Definition 2.5 The **canonical** stratification of a finite-dimensional regular CW complex X is the coarsest stratification of X whose strata are all unions of cells.

It may not be immediately clear why canonical stratifications should exist at all. We will establish both their existence and uniqueness for regular CW complexes in the sequel via an explicit construction, which can be made efficiently algorithmic whenever the number of cells in X is finite. The core of our construction is heavily inspired by Goresky and MacPherson’s proof that *canonical \bar{p} -filtrations* of stratified pseudomanifolds exist and are unique [14, Sec 4.2].

³ Regularity in this context means that the characteristic map of every cell is an embedding.

Remark 2.6 For readers familiar with the Goresky–MacPherson construction, there are two differences between their canonical \bar{p} -filtration proof and the argument that we will present in the sequel here. Namely,

- (1) since our starting point is a regular CW complex rather than a stratified pseudo-manifold, we have direct recourse to the combinatorics of cell incidences which determine the topology of the underlying space, and
- (2) here we use the sheaf of compactly supported cellular cochains on the underlying space rather than the sheaf of intersection chains with respect to a perversity function $\bar{p} : \mathbb{Z}_{\geq 2} \rightarrow \mathbb{Z}$.

Aside from these modifications, our arguments below will follow theirs quite closely.

The following result is a direct consequence of the frontier axiom from Definition 2.1 for cells lying in top strata.

Proposition 2.7 *Let X be a regular CW complex of dimension n equipped with any stratification coarser than its skeletal stratification. If a cell y of X lies in an n -dimensional stratum σ , then so must every cell x which satisfies $x \geq y$.*

Proof Let τ be the unique coarse stratum containing x . Since y lies in the boundary of x by assumption, the closure of τ intersects σ non-trivially at y . The desired conclusion now follows from the axiom of the frontier in Definition 2.1 and the fact that there are no strata of dimension exceeding n , since $\dim \tau \geq \dim \sigma = n$ forces $\tau = \sigma$. \square

Thus, membership of cells in top-dimensional canonical strata is *upward-closed* with respect to the face partial order.

2.2 Localizations of Posets

The localization of a poset \mathbf{P} about a subcollection Σ of its order relations is the minimal category containing \mathbf{P} in which the relations of Σ have been rendered invertible [10, Ch I.1].

Definition 2.8 Let (\mathbf{P}, \geq) be a poset and $\Sigma = \{(x_\bullet \geq y_\bullet)\}$ a subset of its relations which is closed⁴ under composition. The **localization** of \mathbf{P} about Σ is a category $\mathbf{P}[\Sigma^{-1}]$ whose objects are precisely the elements of \mathbf{P} , while morphisms from p to q are given by equivalence classes of finite (but arbitrarily long) Σ -zigzags of order relations in \mathbf{P}

$$p \geq y_0 \leq x_0 \geq \cdots \geq y_k \leq x_k \geq q,$$

where

- (1) only relations in Σ and equalities can point backward (i.e., \leq),
- (2) composition is given by concatenation, and
- (3) the trivial zigzag $p = p$ represents the identity of each element p .

⁴ That is, if $(p \geq q)$ and $(q \geq r)$ both lie in Σ , then so does $(p \geq r)$.

The equivalence between Σ -zigzags is generated by the transitive closure of the following basic relations. Two such zigzags are related

- *horizontally* if one is obtained from the other by removal of intermediate equalities:

$$\begin{aligned}
 (\dots \leq x \geq y = y \geq x' \leq \dots) &\sim (\dots \leq x \geq x' \leq \dots), \\
 (\dots \geq y \leq x = x \leq y' \geq \dots) &\sim (\dots \geq y \leq y' \geq \dots),
 \end{aligned}$$

- or *vertically*, if they form the rows in a grid:

$$\begin{array}{ccccccccc}
 p & \geq & y_0 & \leq & x_0 & \geq & \cdots & \leq & x_k & \geq & q \\
 \parallel & & \text{IV} & & \text{IV} & & & & \text{IV} & & \parallel \\
 p & \geq & y'_0 & \leq & x'_0 & \geq & \cdots & \leq & x'_k & \geq & q
 \end{array}$$

where all vertical face relations (also) lie in Σ .

The category $P[\Sigma^{-1}]$ enjoys the following universal property [10, Lem I.1.3]—there is a canonical **localization functor** $P \rightarrow P[\Sigma^{-1}]$ which acts by inclusion. In particular, it sends each element of P to itself and each relation $x \geq y$ to its own equivalence class of Σ -zigzags. Moreover, any functor $P \rightarrow C$ which maps order relations in Σ to isomorphisms in the target category C factorizes uniquely across the localization functor.

2.3 Cellular Cosheaves

Sheaves (and their dual cosheaves) assign algebraic objects to open sets of topological spaces [2]. Our interest here is in a particularly tame and computable subclass of cosheaves, where the base space is always a finite-dimensional regular CW complex X and the assigned algebraic objects are cell-wise constant. We write $\mathbf{Fc}(X)$ to indicate the poset of cells with order relation $x > y$ denoting that y is a face of x .

Definition 2.9 A **cellular cosheaf** over X taking values in an Abelian category A is a functor $F : \mathbf{Fc}(X) \rightarrow A$. Thus, it assigns to each cell x an A -object $F(x)$ and to each face relation $x \geq y$ an A -morphism $F(x \geq y) : F(x) \rightarrow F(y)$ so that

- $F(x = x)$ is the identity on $F(x)$ for each cell x , and
- $F(y \geq z) \circ F(x \geq y)$ equals $F(x \geq z)$ across any triple of cells $x \geq y \geq z$.

We call $F(x)$ the **stalk** of F at x , and call $F(x \geq y)$ the **extension map** of F at $x \geq y$.

A **morphism** of cellular cosheaves $\eta : F \rightarrow G$ over X is a natural transformation; it assigns to each cell x an A -morphism $\eta_x : F(x) \rightarrow G(x)$ so that for each $x \geq y$ the relevant square in A commutes:

$$\begin{array}{ccc}
 F(x) & \xrightarrow{F(x \geq y)} & F(y) \\
 \eta_x \downarrow & & \downarrow \eta_y \\
 G(x) & \xrightarrow{G(x \geq y)} & G(y)
 \end{array}$$

Cosheaf morphisms may be composed stalkwise, and there is always a zero morphism $0 : F \rightarrow G$ which (as one might expect) assigns the zero map $0_x : F(x) \rightarrow G(x)$ in A to each cell x of X .

Definition 2.10 A (lower-bounded) **complex** of cellular cosheaves F^\bullet (over X , taking values in A) is a sequence of A -valued cellular cosheaves over X connected by cosheaf morphisms:

$$F^0 \xrightarrow{\eta^0} F^1 \xrightarrow{\eta^1} F^2 \xrightarrow{\eta^2} \dots$$

so that every successive composition $\eta^{n+1} \circ \eta^n$ yields the zero morphism.

Every complex of cosheaves F^\bullet on X may be reinterpreted as a single cosheaf which takes values in the category $\mathbf{Ch}(A)$ of lower-bounded cochain complexes in A —consider the collection of commuting squares in A that are assigned to face relations $x \geq y \geq z \geq \dots$ of X :

$$\begin{array}{ccccccc}
 F^0(x) & \longrightarrow & F^1(x) & \longrightarrow & F^2(x) & \longrightarrow & \dots \\
 \downarrow & & \downarrow & & \downarrow & & \\
 F^0(y) & \longrightarrow & F^1(y) & \longrightarrow & F^2(y) & \longrightarrow & \dots \\
 \downarrow & & \downarrow & & \downarrow & & \\
 F^0(z) & \longrightarrow & F^1(z) & \longrightarrow & F^2(z) & \longrightarrow & \dots \\
 \vdots & & \vdots & & \vdots & & \ddots
 \end{array}$$

The horizontal slices of this diagram confirm that a cochain complex in A is allocated to each cell, while vertical arrows between these slices prescribe the desired extension maps. Taking cohomology horizontally, one therefore obtains an ordinary cellular cosheaf of (graded) A -objects over X . We call this the **cohomology** of F^\bullet and denote it by $H^\bullet F$.

3 Local Cohomology of CW Complexes

Let X be a finite-dimensional regular CW complex with face poset $\mathbf{Fc}(X)$ and R a fixed non-trivial commutative ring with identity 1_R . We write $\mathbf{Mod}(R)$ to denote the category of R -modules, and $\mathbf{Ch}(R)$ —rather than the cumbersome $\mathbf{Ch}(\mathbf{Mod}(R))$ —to indicate the category of cochain complexes of R -modules indexed by the nonnegative natural numbers.

Definition 3.1 [5, Def 7.1] The **local cohomology** L^\bullet of X is a complex of cosheaves

$$L^0 \xrightarrow{\beta^0} L^1 \xrightarrow{\beta^1} L^2 \xrightarrow{\beta^2} \dots$$

over X taking values in $\mathbf{Mod}(R)$, prescribed by the following data.

- (1) For each dimension $d \geq 0$ and cell x of X , the cosheaf L^d has as its stalk $L^d(x)$ the free R -module with basis

$$\{z \in \mathbf{Fc}(X) \mid z \geq x \text{ and } \dim z = d\}.$$

When $x \geq y$, the extension $L^d(x \geq y) : L^d(x) \hookrightarrow L^d(y)$ is determined by the obvious inclusion of basis cells.

- (2) The cosheaf morphism β^d assigns to each cell x the map $\beta_x^d : L^d(x) \rightarrow L^{d+1}(x)$ defined by (linearly extending) the following action on basis cells. For each d -cell $z \geq x$, we have

$$\beta_x^d(z) = \sum_w \langle w, z \rangle_R \cdot w,$$

where the sum is taken over all $(d + 1)$ -cells $w \geq x$, and $\langle w, z \rangle_R$ is the R -valued degree of the attaching map in X from the boundary of w onto z . (Since we have assumed that X is regular, this number $\langle w, z \rangle_R$ takes values in $\{0, \pm 1_R\}$ for all cells w and z .)

Routine calculations confirm that $\beta^{d+1} \circ \beta^d$ is always zero and that the β_x^\bullet 's commute with all relevant extension maps. In light of the discussion following Definition 2.10, we will shift perspective and regard L^\bullet as a single cosheaf valued in $\mathbf{Ch}(R)$. In this setting, the stalk $L^\bullet(x)$ lying over each cell x is the entire cochain complex

$$L^0(x) \xrightarrow{\beta_x^0} L^1(x) \xrightarrow{\beta_x^1} L^2(x) \xrightarrow{\beta_x^2} \dots,$$

and $H^\bullet L(x)$ is the compactly supported cohomology of x 's open star in X . By excision, one may describe $H^\bullet L(x)$ as the ordinary relative cohomology of the pair $(\overline{\mathbf{st}}(x), \partial \overline{\mathbf{st}}(x))$ where $\overline{\mathbf{st}}(x)$ is the closure of x 's open star in X and ∂ denotes the topological boundary:

$$\partial \overline{\mathbf{st}}(x) = \overline{\mathbf{st}}(x) - \mathbf{st}(x).$$

When y is a zero-dimensional cell, we have $\partial \overline{\mathbf{st}}(y) = \mathbf{lk}(y)$.

Definition 3.2 For each $d \geq 0$, we write $R[d]^\bullet$ to indicate the special cochain complex in $\mathbf{Ch}(R)$ which is trivial except for a single copy of R in the d th position:

$$0 \rightarrow \dots \rightarrow 0 \rightarrow R \rightarrow 0 \rightarrow \dots$$

Thus, $H^\bullet R[d]$ is the cohomology of small neighborhoods in d -dimensional R -cohomology manifolds, as in (1).

By Definition 2.1, if a cell x lies in some top-dimensional stratum of X , then we must have an isomorphism $H^\bullet L(x) \simeq H^\bullet R[\dim X]$. For cells of high dimension, this requirement has strong consequences.

Proposition 3.3 *If $n = \dim X$, then every n -cell x has $L^\bullet(x) = R[n]^\bullet$. And moreover, every $(n - 1)$ -cell y with $H^\bullet L(y) \simeq H^\bullet R[n]$ lies in the boundary of exactly two n -cells.*

Proof The first assertion follows from Definition 3.1 and the fact that there are no cells of dimension $\geq (n + 1)$ in X . Turning to the second assumption, note that the local cohomology of an $(n - 1)$ -cell y is only supported in dimensions $(n - 1)$ and n :

$$L^\bullet(y) = 0 \rightarrow \dots \rightarrow 0 \rightarrow R \xrightarrow{\beta_y^{n-1}} R^k \rightarrow 0 \rightarrow \dots$$

where k counts the number of n -cells x satisfying $x > y$. If $L^\bullet(y)$ is quasi-isomorphic to $R[n]^\bullet$, then β_y^{n-1} must be injective (since $H^{n-1} L(y)$ is trivial), so $k > 0$. By definition, β_y^{n-1} acts by sending a basis element of R to a linear combination of basis elements in R^k , and by the regularity of X , all coefficients in this linear combination lie in $\{\pm 1_R\}$. The n th cohomology of $L^\bullet(y)$ is the quotient

$$H^n L(y) = \frac{R^k}{\text{im}(\beta_y^{n-1})},$$


and it must be isomorphic to $R = H^n R[n]$ by assumption. Since R is a commutative ring, it satisfies the invariant basis property—that is to say, R^k and R^ℓ are isomorphic as R -modules if and only if $k = \ell$ [29, Def 1.1]. Thus, we have $k = 2$ as desired. \square

In light of the preceding result and (1), one might hope to identify top-dimensional canonical strata by simply clustering together adjacent cells $x \geq y$ in X whenever there exists a top-dimensional cell w satisfying $w \geq x$ so that all extension maps below are quasi-isomorphisms:

$$R[\dim X]^\bullet = L^\bullet(w) \xrightarrow{\simeq} L^\bullet(x) \xrightarrow{\simeq} L^\bullet(y).$$

Our next example shows that this (necessary) condition for x and y to belong to the same top stratum is not sufficient when $\dim w - \dim y > 1$.

Example 3.4 Consider the two-dimensional simplicial complex depicted below: It is obtained by subdividing a parallelogram into four 2-simplices along diagonals and attaching an extra 2-simplex along one of the four resulting half-diagonals.



Let w, x and y , respectively, denote the extra 2-simplex, the shared half-diagonal and the central vertex. It is easy to confirm that $L^\bullet(y)$ is quasi-isomorphic to $R[2]^\bullet$ since its link has the cohomology of a circle. And moreover, given any sequence

$w' > x' > y$ of cells where w' and x' differ from w and x , respectively, the inclusion maps

$$R[2]^\bullet = L^\bullet(w') \hookrightarrow L^\bullet(x') \hookrightarrow L^\bullet(y)$$

are all quasi-isomorphisms. Note, however, that $L^\bullet(x)$ is not quasi-isomorphic to $R[2]^\bullet$ by Proposition 3.3 since there are three 2-simplices containing x as a face. Thus, x does not belong to any canonical 2-stratum, and hence (by Proposition 2.7), neither does y .

4 Extracting Top Strata

Throughout this section, X denotes an n -dimensional regular CW complex, R is a fixed non-trivial commutative ring with identity, and $L^\bullet : \mathbf{Fc}(X) \rightarrow \mathbf{Ch}(R)$ is the local cohomology complex of X from Definition 3.1. Consider the collection of face relations $(x \geq y)$ in $\mathbf{Fc}(X)$ sent by L^\bullet to quasi-isomorphisms, i.e.,

$$E_0 = \{(x \geq y) \text{ in } \mathbf{Fc}(X) \mid H^\bullet L(x \geq y) \text{ is an isomorphism,}\} \tag{3}$$

and define the subcollection

$$W_0 = \{(x \geq y) \text{ in } E_0 \mid H^\bullet L(y) \simeq H^\bullet R[n], \text{ and } (x' \geq y) \in E_0 \text{ for all } x' \geq y \text{ in } \mathbf{Fc}(X)\}. \tag{4}$$

In other words, W_0 is what remains of E_0 when we impose

- *dimensionality* retain only those face relations in which both cells have the local cohomology (isomorphic to that) of a top-dimensional cell, and
- *upward-closure* remove $(x \geq y)$ if there exists some $x' \geq y$ for which $L^\bullet(x' \geq y)$ is not a quasi-isomorphism. (Recall that cells lying in top strata are upward-closed by Proposition 2.7.)

It is clear that if W_0 contains both $(x \geq y)$ and $(y \geq z)$, then it also contains $(x \geq z)$.

Definition 4.1 The category S_0 is the localization of the face poset $\mathbf{Fc}(X)$ about W_0 .

Recall that S_0 has the cells of X as objects, while its morphisms are equivalence classes of W_0 -zigzags (as described in Definition 2.8). And, there is a canonical functor $\mathbf{Fc}(X) \rightarrow S_0$ which is universal with respect to rendering all the face relations from W_0 invertible.

Proposition 4.2 *Two cells $w \neq z$ are isomorphic in S_0 if and only if there is a W_0 -zigzag*

$$w \geq y_0 \leq x_0 \geq \cdots \geq y_k \leq x_k \geq z$$

whose last relation $(x_k \geq z)$ lies in W_0 .

Proof Given any W_0 -zigzag from w to z such as the one depicted above, recall that every $(x_\bullet \geq y_\bullet)$ is in W_0 . And by the upward-closure requirement, each preceding $(x_{\bullet-1} \geq y_\bullet)$ is automatically in W_0 (with the understanding that $x_{-1} = w$). Since the last relation $(x_k \geq z)$ is assumed to lie in W_0 as well, one can directly construct an inverse zigzag in S_0 :

$$z = z \leq x_k \geq y_k \leq \cdots \leq x_0 \geq y_0 \leq w = w.$$

Conversely, if w and z are isomorphic in S_0 , then there exist W_0 -zigzags from w to z and vice versa whose composites are (equivalent to) identity morphisms. We label some relevant cells in these zigzags:

$$w \geq \cdots \leq x \geq z \quad \text{and} \quad z \geq y \leq \cdots \geq w.$$

A brief examination of the horizontal and vertical relations from Definition 2.8 confirms that any zigzag in the equivalence class of the identity must have all its forward-pointing \geq 's lying in W_0 . As a consequence, when we compose the zigzags labeled above, we obtain $(x \geq y) \in W_0$. An appeal to the upward-closure requirement now forces $(x \geq z) \in W_0$, as desired. \square

Here is the main result of this section.

Proposition 4.3 *Two n -dimensional cells lie in the same canonical n -stratum of X if and only if they are isomorphic in S_0 .*

Proof To show that the first assertion implies the second, let σ be the n -stratum containing two distinct cells w and z . Since σ is connected by Definition 2.1, there exists a zigzag of cells lying entirely in σ from w to z , say

$$w \geq v_0 \leq v_1 \geq \cdots \geq v_k \leq z.$$

By the same definition and Proposition 2.7, it follows that every face relation appearing above lies in W_0 . By adding an identity to the end of this zigzag, we may therefore construct a morphism in S_0 from w to z ,

$$w \geq v_0 \leq v_1 \geq \cdots \geq v_k \leq z = z.$$

It is now evident from Proposition 4.2 that w and z are isomorphic in S_0 . In order to show the reverse implication, assume the existence of an isomorphism from w to z in S_0 , say

$$w \geq y_0 \leq x_0 \geq \cdots \geq y_k \leq x_k = z,$$

where the last equality follows from the fact that X contains no cells of dimension $> n$. Define the union

$$Z = \bigcup_{i=0}^k \text{st}(y_i)$$

of open stars, and note that Z is connected—every constituent $\text{st}(y_i)$ is connected and intersects the subsequent $\text{st}(y_{i+1})$ non-trivially (since both stars must contain x_i). By the dimensionality requirement on W_0 , every point in Z has an open neighborhood with compactly supported cohomology isomorphic to $H^\bullet R[n] = H_c^\bullet(\mathbb{R}^n)$, and by the upward-closure requirement no cell lying in Z may lie in the boundary of a cell lying outside Z . Thus, the frontier and link requirements on strata imposed by Definition 2.1 are satisfied by the connected set $Z \subset X$, whence all cells lying in Z must belong to the same canonical n -stratum σ of X . Thus, σ contains both w and z as desired. \square

Finally, note that if an arbitrary cell y lies in a (not necessarily canonical) n -stratum σ of X , then there must exist an n -cell $w \geq y$ also lying in σ by Proposition 2.7; otherwise, we arrive at the contradiction $H^n L(y) = 0$. Thus, we have the following consequence of Proposition 4.3.

Corollary 4.4 *The canonical n -strata of X correspond bijectively with isomorphism classes of its n -cells in S_0 .*

5 Uncovering Lower Strata

As before, let X be the n -dimensional regular CW complex whose canonical strata we wish to discover. In this section, we will inductively construct the desired stratification

$$X = Y_0 \supset Y_1 \supset \cdots \supset Y_n \supset Y_{n+1} = \emptyset, \tag{5}$$

where each $Y_\bullet \subset X$ is a regular CW subcomplex, so that the connected components of each difference $Y_d - Y_{d+1}$ are precisely the canonical $(n - d)$ -strata of X . We have numbered the Y_\bullet in reverse (compare Definition 2.1) merely to indicate the order in which they are constructed. In addition to the initial complex $Y_0 = X$, the first step of our construction also requires using the sets E_0 and W_0 which were defined in (3) and (4), respectively.

Definition 5.1 Given the triple $(Y_{d-1}, E_{d-1}, W_{d-1})$, the $(n - d)$ -dimensional regular CW subcomplex Y_d consists of all cells y in Y_{d-1} for which there exists some cell $x \geq y$ in Y_{d-1} so that $(x \geq y)$ is not in the set W_{d-1} . (By upward-closure, we may as well choose x to be y .) Letting $L_d^\bullet : \mathbf{Fc}(Y_d) \rightarrow \mathbf{Ch}(R)$ be the local cohomology of Y_d , define the set of face relations

$$E_d = E_{d-1} \cap \{(x \geq y) \text{ in } \mathbf{Fc}(Y_d) \mid H^\bullet L_d(x \geq y) \text{ is an isomorphism}\}, \tag{6}$$


and also

$$W_d = W_{d-1} \cup \{(x \geq y) \text{ in } E_d \mid H^\bullet L_d(y) \simeq H^\bullet R[n - d], \text{ and } (x' \geq y) \in E_d \text{ for all } x' \geq y \text{ in } \mathbf{Fc}(Y_d)\}. \tag{7}$$

This produces the subsequent triple (Y_d, E_d, W_d) .

Note that W_\bullet is an increasing sequence of sets while both Y_\bullet and E_\bullet are decreasing.

Example 5.2 To motivate the sets which appear in Definition 5.1, consider the two-dimensional space Y_0 pictured below: It is the closed cone over a wedge of two circles.



We have equipped Y_0 with a regular CW decomposition consisting of four vertices (p, q, z_*, z_*) , seven edges (w_*, y_*, e) and four 2-cells (all x_*). For convenience, let us assume that the coefficient ring R is a field, so that cohomology computations reduce to knowledge of ranks of coboundary matrices. The stalk of the local cohomology cosheaf L_0^\bullet over each of the four 2-cells is just $L_0^\bullet(x_*) = R[2]^\bullet$, and since each y_* has two 2-cells in its open star, we obtain the following stalks:

$$L_0^\bullet(y_*) = (0 \rightarrow R \rightarrow R^2 \rightarrow 0 \rightarrow \dots),$$

where the only non-trivial map has rank one. It is straightforward to confirm that the based inclusions $L_0^\bullet(x_* \geq y_*)$ are all quasi-isomorphisms. Thus, the four face relations in Y_0 of the form $x_* \geq y_*$ will be elements of W_0 . All the other cells have the wrong local cohomology and will therefore not participate in W_0 -relations: For instance,


$$L_0^\bullet(w_*) = (0 \rightarrow R \rightarrow R \rightarrow 0 \rightarrow \dots),$$

where again the nonzero map has rank 1, so the local cohomology is trivial and does not coincide with that of $R[2]^\bullet$. By upward-closure, this also shows that neither the z_* nor p can participate in face relations of W_0 . Similarly, the local cohomology of e does not agree with that of $R[2]^\bullet$; the stalk is

$$L_0^\bullet(e) = (0 \rightarrow R \rightarrow R^4 \rightarrow 0 \rightarrow \dots),$$

again with a rank one map. This disqualifies q from participating in a face relation that lies in W_0 . On the other hand, all four $(w_* \geq p)$ lie in E_0 , so the cosheaf L_0^\bullet is unable to distinguish generic cells lying on the two circles to the left from the singular point p . But once we remove all those x_* and y_* cells which are involved in face relations of

W_0 , we obtain a one-dimensional CW subcomplex $Y_1 \subset Y_0$ whose local cohomology correctly distinguishes p from the w_* and z_* :



In particular, the local cohomology cosheaf L_1^\bullet of Y_1 has the following stalk over p :

$$L_1^\bullet(p) = (0 \rightarrow R \rightarrow R^5 \rightarrow 0 \rightarrow \dots),$$

with the non-trivial map having rank 1; thus, $L_1^\bullet(p)$ manifestly has different cohomology than $R[1]^\bullet = L_1^\bullet(w_*) \simeq L_1^\bullet(z_*)$. Similar top-down constructions (which measure the failure of d -dimensionality for decreasing d) have been employed since the very inception of stratification theory, dating back to the original work of Whitney [31].


Remark 5.3 Two properties of the sequence Y_\bullet mentioned in Definition 5.1 may be inductively justified as follows, using the fact that Y_0 is an n -dimensional regular CW complex as a base case.

- (1) To see that $Y_d \subset Y_{d-1}$ is a regular CW subcomplex, note that if a cell y lies in the difference $Y_{d-1} - Y_d$, then we have $(y = y) \in W_{d-1}$. And given any cell $x \geq y$ in Y_{d-1} , we have $(x \geq y) \in W_{d-1}$ by upward-closure, whence $(x = x) \in W_{d-1}$. Thus, x also lies in the difference $Y_{d-1} - Y_d$. Since the collection of cells removed from Y_{d-1} to obtain Y_d is upward-closed with respect to the face partial order, Y_d is a regular CW subcomplex of Y_{d-1} .
- (2) To see that $\dim Y_d \leq (n - d)$, assume $\dim Y_{d-1} \leq (n - d + 1)$. But now, for each cell w in Y_{d-1} of dimension $(n - d + 1)$, we have $(w = w) \in W_{d-1}$ by Proposition 3.3, so w is not in Y_d . Since Y_d contains no cells of dimension $(n - d + 1)$, its dimension cannot exceed $(n - d)$.

Definition 5.4 For each $d > 0$ in $[n]$, the category S_d is the localization of $\mathbf{Fc}(X)$ about the set W_d from (7).

Recall from Sect. 2.2 that there is a universal functor $\mathbf{Fc}(X) \rightarrow S_d$ for each d , and since $W_{d-1} \subset W_d$ holds by (7), the universal property of localization guarantees a unique functor $S_{d-1} \rightarrow S_d$ which makes the following diagram commute:

$$\begin{array}{ccc} \mathbf{Fc}(X) & \xrightarrow{\quad} & S_{d-1} \\ \downarrow & & \downarrow \\ \mathbf{Fc}(X) & \xrightarrow{\quad} & S_d \end{array}$$



In fact, $\mathcal{S}_{d-1} \rightarrow \mathcal{S}_d$ admits a convenient explicit description: Note that every W_{d-1} -zigzag is automatically a W_d -zigzag, so \mathcal{S}_{d-1} is a subcategory of \mathcal{S}_d and the inclusion $\mathcal{S}_{d-1} \hookrightarrow \mathcal{S}_d$ satisfies this universal property. Our next result establishes a certain W_\bullet -induced monotonicity for morphisms across the entire nested sequence \mathcal{S}_\bullet ; in its statement and beyond, $[p]$ denotes the set $\{0, \dots, p\}$ for each integer $p \geq 0$.

Lemma 5.5 *Fix d in $[n]$ along with cells w and z in X . For each W_d -zigzag ζ from w to z , say*

$$\zeta = (w \geq y_0 \leq x_0 \geq \dots \geq y_k \leq x_k \geq z),$$

define the function $\varphi_\zeta : [k] \rightarrow [d]$ by setting

$$\varphi_\zeta(p) = \min \{q \geq 0 \mid (x_p \geq y_p) \in W_q\}. \tag{8}$$

Then, φ_ζ is monotone increasing in the sense that $\varphi_\zeta(p - 1) \leq \varphi_\zeta(p)$ for all $p \geq 1$ in $[k]$.

Proof Set $\varphi_\zeta(p) = q$ for a fixed p , so that $(x_p \geq y_p) \in W_q - W_{q-1}$, whence both x_p and y_p are cells in $Y_q - Y_{q+1}$. Since $x_{p-1} \geq y_p$ holds in X , the cell x_{p-1} must lie in $Y_r - Y_{r+1}$ for some $r \leq q$ because Y_\bullet is a decreasing sequence of CW subcomplexes of X . Thus, $(x_{p-1} \geq y_{p-1}) \in W_r$, so we have

$$\varphi_\zeta(p - 1) \leq r \leq q = \varphi_\zeta(p),$$

as desired. □

The cells of X are ordered by dimension in the sense that the poset $\mathbf{Fc}(X)$ is only allowed to have morphisms from higher-dimensional to lower-dimensional cells. Since morphisms in \mathcal{S}_\bullet are zigzags, one cannot expect such a dimensional monotonicity to hold verbatim. But the W_\bullet -induced monotonicity from the previous lemma precludes the existence of certain morphisms in \mathcal{S}_\bullet .

Corollary 5.6 *Given $i \leq j$ in $[n]$ along with cells w in Y_j and z in $Y_i - Y_j$, there are no morphisms in \mathcal{S}_d from w to z for any d in $[n]$.*

Proof Note that a direct face relation $w \geq z$ cannot hold in X since Y_j is a CW subcomplex of Y_i . Thus, any morphism in \mathcal{S}_d from w to z admits a genuine zigzag representative. Proceeding by contradiction, assume the existence of such a zigzag ζ :

$$w \geq y_0 \leq x_0 \geq \dots \geq y_k \leq x_k \geq z.$$

The face relations $w \geq y_0$ and $x_k \geq z$ in X guarantee $j \leq \varphi_\zeta(0)$ and $\varphi_\zeta(k) \leq i$, where φ_ζ is the function from (8). The monotonicity of φ_ζ from Lemma 5.5 now forces $j \leq i$, and hence, $i = j$. The cell z is thus constrained to lie in $Y_i - Y_j = \emptyset$, which yields the desired contradiction. \square

As a direct consequence of the preceding result, isomorphism classes in S_i of $(n-i)$ -dimensional cells from Y_i remain unchanged across all the inclusions $S_i \hookrightarrow S_j$ for $i \leq j$. The next result gives convenient alternate descriptions of such isomorphism classes across the entire sequence S_\bullet .

Proposition 5.7 *For each $d \in [n]$, the following are equivalent. Two cells w and z in X*

- (1) *lie in the same connected component of $Y_d - Y_{d+1}$,*
- (2) *are connected by a W_d -zigzag, whose relations (\leq and \geq) all lie in $W_d - W_{d-1}$,*
- (3) *are isomorphic in S_d to a common $(n-d)$ -dimensional cell from Y_d .*

Proof We will show (1) \Rightarrow (2) \Rightarrow (3) \Rightarrow (1). If w and z lie in the same connected component of $Y_d - Y_{d+1}$, then there is a zigzag of cells

$$w \geq v_0 \leq v_1 \geq \dots \geq v_k \leq z,$$

lying entirely in $Y_d - Y_{d+1}$. By Definition 5.1, each face relation (\geq or \leq) appearing in this zigzag lies in W_d , and by Lemma 5.5, no such relation can lie in W_{d-1} . So by (7), the cell w (and indeed, each cell in this zigzag) has its local cohomology $H^\bullet L_d(v)$ isomorphic to $H^\bullet R[n-d]$. By the argument used while proving Corollary 4.4, and recalling that $\dim Y_d \leq (n-d)$, there exists an $(n-d)$ -cell $v \geq w$ in Y_d which is isomorphic in S_d to w , and hence to every other cell in our zigzag (including z). Finally, if there is an $(n-d)$ -cell v from Y_d isomorphic to both w and z in S_d , then there is a W_d -zigzag from w to z obtained by composing zigzags $w \rightarrow v \rightarrow z$. By Lemma 5.5, all cells in this composite lie in $W_d - W_{d-1}$, so w and z lie in the same connected component of $Y_d - Y_{d+1}$. \square

Our next result describes the canonical strata of X as isomorphism classes in S_\bullet .

Proposition 5.8 *For each $d \in [n]$ and canonical $(n-d)$ -stratum σ of X , every cell lying in σ is isomorphic in S_d to some $(n-d)$ -cell from Y_d .*

Proof Let $\emptyset = L_{-1} \subset \dots \subset L_{d-1} = L$ be the link of σ in the sense of Definition 2.1. Assume, proceeding via induction, that for each $j < d$ the canonical $(n-j)$ -strata of X correspond to isomorphism classes in S_j of $(n-j)$ -cells from Y_j . By Proposition 5.7, these strata are precisely the connected components of $Y_j - Y_{j+1}$. Thus, for each cell x lying in σ and $j < d$, we have quasi-isomorphisms

$$C_c^\bullet(\mathcal{C}L_{d-j-1} \times \mathbb{R}^{n-d}) \xrightarrow{\cong} C_c^\bullet(\mathbf{st}(x) \cap Y_j),$$

where $\mathbf{st}(x)$ is the open star of x in $X = Y_0$. Assuming that $x \geq y$ holds in X for another cell y lying in σ , we would like to prove that $L_j^\bullet(x \geq y)$ is a quasi-isomorphism. To this end, consider the following diagram in $\mathbf{Ch}(R)$:



$$\begin{array}{ccccc}
 C_c^\bullet(\mathcal{C}L_{d-j-1} \times \mathbb{R}^{n-d}) & \xrightarrow{\cong} & C_c^\bullet(\mathbf{st}(y) \cap Y_j) & \xrightarrow{\cong} & L_j^\bullet(y) \\
 & \searrow \cong & \downarrow & & \uparrow L^\bullet(x \geq y) \\
 & & C_c^\bullet(\mathbf{st}(x) \cap Y_j) & \xrightarrow{\cong} & L_j^\bullet(x)
 \end{array}$$

The triangle to the left commutes by the refinability constraint of Definition 2.1 (note that here the downward pointing arrow is the map induced by inclusion), whereas square to the right commutes by the natural contravariant equivalence between singular and cellular compactly supported cohomology. Being sandwiched in this commuting diagram forces $L_j^\bullet(x \geq y)$ to also be a quasi-isomorphism for all $j < d$, so we have $(x \geq y) \in E_{d-1}$. Being an $(n - d)$ -stratum, σ is an R -cohomology manifold, so there are quasi-isomorphisms

$$R[n - d]^\bullet \xrightarrow{\cong} L_d^\bullet(x) \xrightarrow{\cong} L_d^\bullet(y).$$

Since σ is a top-dimensional stratum of the subcomplex Y_d , we have $(x' \geq y) \in E_d$ whenever $x' \geq y$ in Y_d by Proposition 2.7. Thus, $(x \geq y) \in W_d - W_{d-1}$ by (7). Since σ is connected, any two cells of X that it contains may be connected by a zigzag of cells lying in σ , and by the preceding argument, each face relation appearing in such a zigzag lies in $W_d - W_{d-1}$. The desired conclusion now follows from Proposition 5.7. □

The following result establishes the converse of Proposition 5.8.

Proposition 5.9 *For each d in $[n]$, the isomorphism class in S_d of any $(n - d)$ -dimensional cell from Y_d is a canonical $(n - d)$ -stratum of X .*

Proof Proceeding once again by induction, we assume that the statement holds for all j in $[d - 1]$. Fix an $(n - d)$ -dimensional cell w in $Y_d \subset X$, and let σ be the collection of all cells which are isomorphic to w in S_d . We first show that σ satisfies the link axiom from Definition 2.1. Writing $\mathbf{st}(w)$ and $\mathbf{lk}(w)$ for w 's star and link in X , for each $i \leq d$ we have the intersections

$$\mathbf{st}_i(w) = \mathbf{st}(w) \cap Y_i \quad \text{and} \quad \mathbf{lk}_i(w) = \mathbf{lk}(w) \cap Y_i,$$

which form the star and link of w in $Y_i \subset X$. Since $\mathbf{st}_\bullet(w)$ admits a stratified homeomorphism to the corresponding $\mathcal{C}\mathbf{lk}_\bullet(w) \times \mathbb{R}^{n-d}$, we get stratified quasi-isomorphisms of singular cochain complexes

$$C_c^\bullet(\mathcal{C}\mathbf{lk}_i(w) \times \mathbb{R}^{n-d}) \xrightarrow{\cong} C_c^\bullet(\mathbf{st}_i(w)),$$

whose codomains are, in turn, quasi-isomorphic to the local cohomology stalks $L_i^\bullet(w)$. If z is any other cell which becomes isomorphic to w in S_d , then by Proposition 5.7 there is a zigzag from w to z whose face relations all lie in $W_d - W_{d-1}$:

$$w \geq v_0 \leq v_1 \geq \dots \geq v_{k-1} \leq v_k \geq z,$$

and hence zigzags of quasi-isomorphisms in $\mathbf{Ch}(R)$ for all $i \in [d]$

$$L_i^\bullet(w) \rightarrow L_i^\bullet(v_0) \leftarrow \cdots \leftarrow L_i^\bullet(v_k) \rightarrow L_i^\bullet(z).$$

Each complex in sight is bounded and free (hence projective), so the backward-pointing maps admit quasi-inverses (see [28, Cor 10.4.7]). Thus, one may turn them around to obtain maps $L_i^\bullet(w) \rightarrow L_i^\bullet(z)$ which fit into a string of quasi-isomorphisms

$$C_c^\bullet(\mathcal{C}\mathbf{k}_i(w) \times \mathbb{R}^{n-d}) \rightarrow C_c^\bullet(\mathbf{st}_i(w)) \rightarrow L_i^\bullet(w) \rightarrow L_i^\bullet(z) \rightarrow C_c^\bullet(\mathbf{st}_i(z)).$$

Setting $L_j = \mathbf{k}_{d-j-1}(w)$ for all $j \in [d - 1]$, one obtains a candidate $(d - 1)$ -dimensional stratified space L to serve as the link of σ .

Turning now to the frontier axiom, for each $i < d$ let K_i^\bullet be the quotient complex $L_i^\bullet/L_{i+1}^\bullet$, suitably restricted so that it assigns zero stalks to all cells lying outside σ . Over each cell z in σ , the cohomology $H^\bullet K_i(z)$ coincides with the compactly supported cohomology of the difference $\mathbf{st}_i(z) - \mathbf{st}_{i+1}(z)$. By construction, the closure of a canonical $(n - i)$ -stratum τ in X intersects σ at the cell z if and only if the cohomology $H^{n-i} K_i(z)$ is non-trivial. By definition, K_i^\bullet fits into a short exact sequence (of complexes of cosheaves over σ):

$$0 \rightarrow L_{i+1}^\bullet \hookrightarrow L_i^\bullet \twoheadrightarrow K_i^\bullet \rightarrow 0.$$

Applying the five lemma [28, Ex 1.3.3] to the resulting long exact sequence, we note that all the extension maps of K_i^\bullet induce isomorphisms on cohomology. Thus, if $H^{n-i} K_i$ is nonzero at z , then it is also nonzero at every other cell in σ , whence every cell in σ admits a co-face from τ . Thus, $\sigma \subset \bar{\tau}$, as desired.

Finally, to see that σ is canonical, we use Proposition 5.8 to pick an $(n - d)$ -cell w from Y_d so that all cells in σ are isomorphic in S_d to w . Proposition 5.7 now guarantees that no additional cells may be added to σ while preserving its connectedness. \square

The three preceding propositions establish that the sequence $Y_\bullet \subset X$ from Definition 5.1 constitutes the canonical stratification of X . The localization functors $\mathbf{Fc}(X) \rightarrow S_\bullet$ and the universal functors $S_\bullet \rightarrow S_{\bullet+1}$ arising from the inclusions $W_\bullet \subset W_{\bullet+1}$ fit into a commutative diagram

$$\begin{array}{ccccccc} \mathbf{Fc}(X) & & & & & & \\ \downarrow & \searrow & & & & & \\ S_0 & \longrightarrow & S_1 & \longrightarrow & \cdots & \longrightarrow & S_n = S, \end{array}$$

which one may view as analogous to an injective resolution of the functor $\mathbf{Fc}(X) \rightarrow S_0$ from Sect. 4. Proceeding from left to right in the sequence above, we obtain canonical strata of decreasing dimension by examining isomorphism classes in S_\bullet . We collect these results into the following expanded version of Theorem 1.1.

Theorem 5.10 *Given a finite regular n -dimensional CW complex X , let Y_\bullet and W_\bullet be as in Definition 5.1, and let S_\bullet denote the localization of the face poset $\mathbf{Fc}(X)$ about*

W. For each $d \in [n]$, the canonical $(n - d)$ -strata of X correspond bijectively with isomorphism classes of $(n - d)$ -dimensional cells from Y_d in S_k for any $k \geq d$.

The morphisms in the last category $S = S_n$ also recover the frontier partial order among canonical strata.

Proposition 5.11 *The frontier relation $\tau \succeq \sigma$ holds among two canonical strata of X if and only if there is a morphism in S from a cell lying in τ to a cell lying in σ .*

Proof If $\tau \succeq \sigma$ holds, then there are cells w and z in τ and σ , respectively, so that $w \geq z$ in $\mathbf{Fc}(X)$. This face relation includes into S as a morphism from w to z . Conversely, given cells w and z in τ and σ , respectively, assume the existence of a W -zigzag

$$\zeta = (w \geq y_0 \leq x_0 \geq \dots \geq y_k \leq x_k \geq z),$$

and let $\varphi_\zeta : [k] \rightarrow [n]$ be the associated monotone function from Lemma 5.5. By construction, this zigzag contains cells in canonical strata τ_i of dimension $n - \varphi_\zeta(i)$ for all i in $[k]$. If $\varphi_\zeta(i) = \varphi_\zeta(i + 1)$ then $\tau_i = \tau_{i+1}$; otherwise, $\tau_i \succ \tau_{i+1}$ (because x_i is a co-face of y_{i+1}). Thus, there is a descending sequence of canonical strata

$$\tau \succeq \tau_0 \succeq \tau_1 \succeq \dots \succeq \tau_k \succeq \sigma,$$

so we obtain $\tau \succeq \sigma$ as desired. □

The argument above shows that the zigzag paths which form morphisms in S may remain indefinitely in a single stratum or descend to lower-dimensional adjacent strata, but they can never ascend to higher strata. In other words, S is a (cellular, 1-categorical) version of the *entrance path category* associated with the canonical stratification of X —see [21, Def 3.1], [27, Sec 7] or [32, Sec 2].

6 Algorithms

From a computational perspective, the local cohomology L^\bullet of a finite regular CW complex enjoys the obvious, but enormous, advantage of being *local*. One only ever needs to construct cochain complexes corresponding to open stars of cells (rather than holding the entire complex in system memory). Since each such star-complex may be processed independently of the others, the construction of L^\bullet and the computation of $H^\bullet L$ are inherently distributable operations.

6.1 Subroutines

There are several avenues for efficiently storing a CW complex for the purpose of computing its cohomology: One may employ a *Hasse graph*, whose vertices are cells and whose edges, usually adorned with an orientation, connect neighboring cells across codimension one. Alternately, one may directly store (sparse avatars of) *coboundary*

matrices over R , whose nonzero entries encode adjacencies between cells of codimension one. In the presence of extra structure—for instance, if the complex involved is simplicial or cubical—the additional rigidity affords access to highly optimized data structures such as simplex trees or bitmaps. Rather than tethering the prospective implementer to a fixed and possibly inconvenient representation, we will present all our algorithms at a high level and work directly with the poset of cells.

```

Algorithm: UpSet( $P, u, v$ )
In: A finite poset  $P$  and elements  $u, v \in P$ 
Out: The subposet  $P_v^u$  of elements  $\geq v$  but not  $\geq u$ 
01 initialize an empty queue  $Q$  of elements
02 add  $v$  to  $Q$ 
03 while  $Q$  is nonempty
04   remove  $w$  from  $Q$ 
05   if  $w$  is not in  $P_v^u$ 
06     add  $w$  to  $P_v^u$ 
07     for each element  $x > w$  in  $P$  with  $x \not\geq u$ 
08       add  $x$  to  $Q$  (uniquely)
09 return  $P_v^u$ 
    
```

Our first subroutine UpSet employs a minor modification of breadth-first search on a finite poset P to extract the subposet of all elements which lie above a fixed element v but *not* above a fixed element u . In all instances of interest here, our input P will be the face poset $\mathbf{Fc}(X)$ of a finite regular CW complex X , while u and v will be adjacent cells across codimension one:

$$u >_1 v \text{ if } u > v \text{ and } \dim u - \dim v = 1.$$

By definition, the output P_v^u in such cases is precisely the subposet of $\mathbf{Fc}(X)$ containing cells which lie in the difference of open stars $\mathbf{st}(v) - \mathbf{st}(u)$. We also reserve the right to set $u = \emptyset$, so that line 07 does not check x against u and the output is just the poset of cells in $\mathbf{st}(v)$.

The next required subroutine is Cohom, which accepts a poset of cells as input and returns the sequence of R -modules corresponding to the cohomology of the associated cochain complex. When R is a field, the cohomology modules are all vector spaces and in this case the output can just be a sequence of integers storing their dimensions. If R is the ring of integers, then one has to encode the torsion subgroups (if any) in addition to the ranks of the free parts in every dimension. In either case, since the efficient computation of (co)homology has been extensively discussed elsewhere [6,16,17,20], we will treat this subroutine as a black box and not explicitly write it down here.

6.2 Description

Our main algorithm StratCast accepts as input the poset of cells in a finite n -dimensional regular CW complex X and uses the constructions of Definition 5.1 to assign each cell x a number $\text{codim}(x)$ in $[n] = \{0, 1, \dots, n\}$ so that x lies in a canonical stratum of dimension $n - \text{codim}(x)$. Once codimensions have been determined for

all cells, finding the actual canonical strata reduces to the simple task of locating connected components of a fixed codimension.

Algorithm: StratCast(Fc(X))

In: The face poset of a finite n -dim regular CW complex

Out: Assigns canonical codimensions to cells

01	for each d in $(0, 1, \dots, n - 1, n)$	
02	for each cell w in $\mathbf{Fc}(X)$	<i>// dist</i>
03	assign $h^\bullet(w) = \text{Cohom}(\text{UpSet}(\mathbf{Fc}(X), \emptyset, w))$	
04	for each $(x >_1 y)$ in $\mathbf{Fc}(X)$ with $\ell(x >_1 y) = d - 1$	<i>// dist</i>
05	assign $c^\bullet(x, y) = \text{Cohom}(\text{UpSet}(\mathbf{Fc}(X), y, x))$	
06	if $c^\bullet(x, y)$ is trivial	
07	assign $\ell(x >_1 y) = d$	
08	assign $\text{codim}(z) = d$ to each $(n - d)$ -cell z in $\mathbf{Fc}(X)$	
09	for each i -cell u in $\mathbf{Fc}(X)$ with i in $(n - d - 1, \dots, 1, 0)$	<i>// dist</i>
10	if $h^{n-d}(u) \simeq R$ and all other $h^\bullet(u) = 0$	
11	if $\text{codim}(v) = d$ and $\ell(v >_1 u) = d$ for all $v >_1 u$ in $\mathbf{Fc}(X)$	
12	assign $\text{codim}(u) = d$	
13	remove all cells from $\mathbf{Fc}(X)$ with $\text{codim} = d$	

The main loop (in line 01) increments the current codimension d from 0 up to n ; as it executes each iteration, the cells lying on $(n - d)$ -dimensional canonical strata are identified and removed from X , so that in the d th iteration we pass from Y_d to Y_{d+1} in the language of Definition 5.1. There are three secondary loops (at lines 02, 04 and 09); all three are trivially distributable across several processors, as indicated by the *dist* comments in those lines. The first loop in line 04 is easiest to explain: It uses our subroutines UpSet and Cohom to compute the compactly supported cohomology of each open star $\text{st}(w) \cap Y_d$ where w is a cell in $Y_d \subset X$.

The second loop from line 04 computes whether or not inclusions of open stars $\text{st}(x) \subset \text{st}(y)$ for $x >_1 y$ induce isomorphisms on compactly supported cohomology by examining their cokernels. In other words, we use the fact that the map induced by $L_d^\bullet(x >_1 y)$ (namely, the inclusion $\text{st}(x) \cap Y_d \hookrightarrow \text{st}(y) \cap Y_d$) on compactly supported cohomology is an isomorphism if and only if the difference $(\text{st}(y) - \text{st}(x)) \cap Y_d$ has trivial compactly supported cohomology. We begin by assigning to each face relation $(x >_1 y)$ the number $\ell(x >_1 y) = -1$. As the algorithm executes, $\ell(x >_1 y)$ gets incremented to the largest integer k so that $(x >_1 y)$ lies in E_k , where E_\bullet denotes the sets from (6). Since this loop is completely independent of the loop in line 02, it may be executed simultaneously with that loop.

The third loop, which spans lines 09 through 12, must be executed after the first two intermediate loops have terminated. It uses the cohomology groups h^\bullet computed by the loop in line 02 and the ℓ -values computed by the loop in line 04 to select cells lying on canonical strata of dimension $(n - d)$ via (7). Note that line 10 ensures the correct dimensionality and line 11 enforces upward-closure. Line 13 removes all cells lying on canonical $(n - d)$ -strata from Y_d , so we are left with Y_{d+1} and the outer loop of line 01 enters its next iteration.

6.3 Complexity

The intermediate loops in line 02 and line 04 of `StratCast` require breadth-first search (for constructing open stars via `UpSet`) and matrix diagonalization through row and column operations (for computing cohomology with `Cohom`). Let m be the number of cells in X , and define the additional *star size* complexity parameter

$$p = \max \{ \text{number of cells in } \mathbf{st}(x) \mid x \text{ is a cell in } X \}. \tag{9}$$

(It suffices to take the maximum over zero-dimensional cells, since those have the largest stars.) The worst-case complexity of computing cohomology of any open star, or a difference of open stars, via row and column operations is thus $O(p^3)$ provided we incur a constant cost performing ring operations in R . The cost of running `Cohom` clearly dominates the $O(p^2)$ cost of executing breadth-first search in `UpSet`. Since the loops of line 02 and 04 may be executed in a completely distributed fashion, their combined cost within a single iteration of the outer loop of line 01 is $O(p^3)$.

The third intermediate loop (on line 09) may be reinterpreted as two nested loops: an outer i -loop which decrements dimension from $(n - d - 1)$ to 0, and an inner u -loop which examines all remaining cells of dimension i . Since we require knowledge of *codim* values of all cells $v >_1 u$ while processing each u , only the u -loop is actually distributable. During each iteration of the u -loop, one incurs an $O(p)$ cost while checking all remaining cells $v >_1 u$ in line 11, and an R -dependent cost of testing R -module isomorphisms in line 10. At least when R is a finite field, the rationals or the integers, checking whether a given module has rank one or zero incurs $O(1)$ cost. Thus, in typical cases we expect line 11 to dominate the burden of executing the u -loop. Noting that the i -loop runs at most n times regardless of d , each iteration of the loop on line 09 incurs a computational cost of $O(np)$.

Since all cells with *codim* = d are identified on lines 08 and 12 in a given iteration of the outer d -loop of `StratCast`, we assume that removing them from X incurs no additional cost. Putting all our estimates together, each iteration of the d -loop incurs a worst-case cost of $O(p^3 + np)$. Noting that this loop executes exactly $(n + 1)$ times, the total complexity of running `StratCast`—assuming that the number of available processors exceeds the number of cells m plus the maximum number mp of codimension-one face relations in X —is $O((n + 1)(p^3 + np))$. And once `StratCast` has terminated, one may find all desired canonical strata of X by computing connected components of cells with the same *codim* value. This is linear in the number of cells in X , so we obtain the following result.

Proposition 6.1 *Let X be an n -dimensional regular CW complex with m cells and star size p as in (9). The time complexity of computing its canonical stratification with coefficients in a (commutative, unital) ring R via `StratCast` is*

$$O((n + 1)(p^3 + np) + m),$$

provided that

- (1) ring operations in R incur $O(1)$ cost,
- (2) isomorphism-testing against 0 and R in $\mathbf{Mod}(R)$ is $O(1)$, and
- (3) the number of available processors exceeds $m(p + 1)$.

The first two conditions above are satisfied by typical choices of R (such as the integers or finite fields). In practice, one expects to have $m \gg (n + 1)(p^3 + np)$, so the observed cost of computing canonical strata is essentially linear in the size of X .

Acknowledgements I am indebted to Justin Curry, Rob Ghrist and Ulrike Tillmann for insightful discussions, and to Adam Brown and Yossi Bokor for carefully reading (and discovering oversights in) earlier drafts of this paper. I am also grateful to the two anonymous referees for their corrections, comments and suggestions for improvement. This work was partially supported by The Alan Turing Institute under the EPSRC Grant EP/N510129/1 and by the Friends of the Institute for Advanced Study. It is dedicated to Mark Goresky and Bob MacPherson, whose efforts over four decades have built some of my favorite playgrounds.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. P. Bendich, S. Mukherjee, and B. Wang. Local homology transfer and stratification learning. In *Proceedings of the 23rd Annual Symposium on Discrete Algorithms*, pages 1355–1370, 2012.
2. G. Bredon. *Sheaf Theory*. Springer, 1997.
3. G. Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.
4. J. Curry. Sheaves, cosheaves and applications. [arXiv:1303.3255v2](https://arxiv.org/abs/1303.3255v2) [math.AT], 2013.
5. J. Curry. Dualities between cellular sheaves and cosheaves. *Journal of Pure and Applied Algebra*, 222(4):966–993, 2018.
6. J. Curry, R. Ghrist, and V. Nanda. Discrete Morse theory for computing cellular sheaf cohomology. *Foundations of Computational Mathematics*, 16(4):875–897, 2016.
7. W. Dwyer and D. Kan. Calculating simplicial localizations. *Journal of Pure and Applied Algebra*, 18:17–35, 1980.
8. H. Edelsbrunner and J. Harer. *Computational Topology – an Introduction*. American Mathematical Society, Providence, RI, 2010.
9. S. Eilenberg and J. Zilber. On products of complexes. *American Journal of Mathematics*, 75:200–204, 1953.
10. P. Gabriel and M. Zisman. *Calculus of fractions and homotopy theory*. Number 35 in *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer-Verlag, 1967.
11. J. Getz and M. Goresky. *Hilbert Modular Forms with Coefficients in Intersection Homology and Quadratic Base Change*. Birkhäuser, 2012.
12. R. Ghrist. *Elementary Applied Topology*. Createspace, 2016.
13. M. Goresky and R. MacPherson. Intersection homology theory. *Topology*, 2(19):135 – 162, 1980.
14. M. Goresky and R. MacPherson. Intersection homology II. *Inventiones Mathematicae*, 71:77–129, 1983.
15. M. Goresky and R. MacPherson. *Stratified Morse Theory*, volume 14 of *Ergebnisse Der Mathematik Und Ihrer Grenzgebiete*. Springer-Verlag, 1988.
16. S. Harker, K. Mischaikow, M. Mrozek, and V. Nanda. Discrete Morse theoretic algorithms for computing homology of complexes and maps. *Foundations of Computational Mathematics*, 14(1):151–184, 2014.
17. G. Henselman and R. Ghrist. Matroid filtrations and computational persistent homology. [arXiv:1606.00199](https://arxiv.org/abs/1606.00199) [math.AT], 2016.
18. F. Kirwan and J. Woolf. *An Introduction to Intersection Homology Theory*. CRC Press, 2006.

19. J. A. Lee and M. Verleysen. *Nonlinear dimensionality reduction*. Springer-Verlag, 2008.
20. K. Mischaikow and V. Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Discrete and Computational Geometry*, 50(2):330–353, 2013.
21. V. Nanda. Discrete Morse theory and localization. *Journal of Pure and Applied Algebra*, 223(2):459–488, 2019.
22. V. Nanda, D. Tamaki, and K. Tanaka. Discrete Morse theory and classifying spaces. *Advances in Mathematics*, 340:723–790, 2018.
23. F. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Texts and Monographs in Computer Science. Springer, 1985.
24. C. Rourke and B. Sanderson. Homology stratifications and intersection homology. *Geometry and Topology Monographs*, 2:455–472, 1999.
25. A. Shepard. *A Cellular Description of the Derived Category of a Stratified Space*. Brown University PhD Thesis, 1985.
26. D. Tamaki. Cellular stratified spaces. In A. Darby, J. Grbic, and L. Zhi, editors, *Combinatorial And Toric Homotopy: Introductory Lectures*, number 35 in Lecture Notes Series, Institute For Mathematical Sciences, National University Of Singapore. World Scientific, 2017.
27. D. Treumann. Exit paths and constructible stacks. *Compositio Mathematica*, 145:1504–1532, 2009.
28. C. Weibel. *Homological Algebra*. Cambridge University Press, 1994.
29. C. Weibel. *The K-Book: an Introduction to Algebraic K-theory*. Number 145 in Graduate Studies in Mathematics. AMS, 2013.
30. S. Weinberger. *The Topological Classification of Stratified Spaces*. University of Chicago Press, 1994.
31. H. Whitney. Tangents to an analytic variety. *Annals of Mathematics*, 81:469–549, 1965.
32. J. Woolf. The fundamental category of a stratified space. *Journal of Homotopy and Related Structures*, 4(1):359–387, 2009.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.