

Persistence in discrete Morse theory

Dissertation

zur Erlangung des mathematisch-naturwissenschaftlichen Doktorgrades

Doctor rerum naturalium

der Georg-August-Universität Göttingen

vorgelegt von

Ulrich Bauer

aus München

Göttingen 2011

D7

Referent: Prof. Dr. Max Wardetzky

Koreferent: Prof. Dr. Robert Schaback

Weiterer Referent: Prof. Dr. Herbert Edelsbrunner

Tag der mündlichen Prüfung: 12.5.2011

Contents

1	Introduction	1
1.1	Overview	1
1.2	Related work	7
1.3	Acknowledgements	8
2	Discrete Morse theory	11
2.1	CW complexes	11
2.2	Discrete vector fields	12
2.3	The Morse complex	15
2.4	Morse and pseudo-Morse functions	17
2.5	Symbolic perturbation	20
2.6	Level and order subcomplexes	23
2.7	Straight-line homotopies of discrete Morse functions	28
2.8	PL functions and discrete Morse functions	29
2.9	Morse theory for general CW complexes	34
3	Persistent homology of discrete Morse functions	41
3.1	Birth, death, and persistence pairs	42
3.2	Duality and persistence	44
3.3	Stability of persistence diagrams	45
4	Optimal topological simplification of functions on surfaces	51
4.1	Topological denoising by simplification	51
4.2	The persistence hierarchy	53
4.3	The plateau function	59
4.4	Checking the constraint	62
5	Efficient computation of topological simplifications	67
5.1	Defining a consistent total order	68

Contents

5.2	Computing persistence pairs	68
5.3	Extracting the gradient vector field	70
5.4	Constructing the simplified function	70
5.5	Correctness of the algorithm	71
6	Discussion	75
6.1	Computational results	75
6.2	Relation to simplification of persistence diagrams	75
6.3	Symmetrizing the algorithm	76
6.4	Flooding and carving artifacts	77
6.5	Combining topological simplification and energy methods	81
6.6	Topological simplification on regular CW complexes	81
6.7	Persistence pairs that cannot be canceled	83
6.8	Removing local extrema from functions on manifolds	84
6.9	Matching local extrema of functions on manifolds	85
6.10	Simplexwise linear simplification is NP-hard	86
7	Future work	91
	Bibliography	95
	Index	101

Chapter 1

Introduction

The goal of this thesis is to bring together two different theories about critical points of a scalar function and their relation to topology: *Discrete Morse theory* and *Persistent homology*. While the goals and fundamental techniques are different, there are certain themes appearing in both theories that closely resemble each other. In certain cases, the two threads can be joined, leading to new insights beyond the classical realm of one particular theory.

1.1 Overview

Discrete Morse theory [28, 30] provides combinatorial equivalents of several core concepts of classical Morse theory, such as discrete Morse functions, discrete gradient vector fields, critical points, and a cancellation theorem for the elimination of critical points of a vector field. Because of its simplicity, it not only maintains the intuition of the classical theory, but allows to surpass it in a certain sense by providing explicit and canonical constructions that would become quite complicated in the smooth setting.

Persistent homology [23, 67] quantifies topological features of a function. It defines the birth and death of homology classes at critical points, identifies pairs of these (*persistence pairs*), and provides a quantitative notion of their stability (*persistence*).

Whereas (discrete) Morse theory makes statements about the *homotopy type* of the sublevel sets of a function, persistence is concerned with their *homology*. While homology is an invariant of homotopy equivalences, the converse is not true: not every map inducing an isomorphism in homology is a homotopy equivalence. In this thesis we establish a connection between both theories and

use this combination to solve problems that are not easily accessible by any single theory alone. In particular, we make contributions to the following topics:

1.1.1 Topological simplification

Measured data and functions constructed from measured data suffer from omnipresent noise introduced during the measuring process. Separating relevant information from noise is therefore a widely considered problem. Various approaches to this problem found in the literature are based on partial differential equations or energy functionals derived from differential quantities. In contrast, we take a purely topological point of view and regard noise as a source of critical points. Indeed, even arbitrarily small amounts of noise (with respect to the supremum norm) may give rise to an arbitrarily large number of critical points. We may hence interpret critical points that can be *eliminated* by small perturbations as being caused by noise. Consequently, we consider the following optimization problem:

Problem (Topological simplification on surfaces). *Given a function f on a surface and a real number $\delta > 0$, find a function f_δ subject to $\|f_\delta - f\|_\infty \leq \delta$ such that f_δ has a minimum number of critical points.*

Of course, in order to make this problem precise, one needs to choose a class of admissible functions that provides a meaningful notion of critical points. We present a solution to this problem for the class of *discrete pseudo-Morse functions*, which generalizes the class of discrete Morse functions introduced by Forman [28]. In particular, we consider a critical point not in the strict analytical sense as a point of vanishing differential, but in the more general sense as a point where the topological type of the sublevel set changes. Our notion of pseudo-Morse functions incorporates both piecewise linear functions and pixel data. Our solution to the above problem relies on several results, which are outlined in the following.

1.1.2 Canceling a single pair of critical points from a function

Forman [28] describes a simple method for eliminating a pair of critical points from a discrete *vector field*. Modifying a *function* accordingly requires a slight extension of Forman's method. We first observe that a discrete gradient vector

field in the sense of Forman induces a partial order on the cells of the underlying complex, giving rise to the notion of *ascending* and *descending sets* (in analogy to the notion of ascending and descending manifolds in the classical theory). Building on these concepts, we describe a method for eliminating a pair of critical points from a discrete Morse *function*. Our construction complements Forman's cancellation method for discrete *gradient vector fields*; in particular, it is applicable to general CW complexes (Section 4.3).

In order to cancel a pair of critical points whose values differ by 2δ , our method modifies the function by δ in the supremum norm, which is the minimum required for canceling such a pair. To achieve this minimum, the function values have to be modified on the ascending and descending sets of the canceled pair. Since these sets may contain cells of *any* dimension, other *critical* values might also have to be changed in this process.

1.1.3 Degenerate functions

Morse theory, in any of its variations, fundamentally relies on the assumption that critical points are non-degenerate. This condition prevents the theory from being directly applicable to arbitrary input functions. Additionally, in our construction, the canonical function arising from canceling a single pair of critical points has a plateau and is hence a degenerate function in the sense of discrete Morse theory. This necessitates a method to deal with degenerate functions. To do so, we devise a symbolic perturbation scheme based on discrete gradient vector fields, which allows to treat the degenerate case in much the same way as the generic case by introducing the larger class of *pseudo-Morse functions* (Section 2.4). In order to obtain a notion of critical points for pseudo-Morse functions, we work with an *explicit* gradient vector field consistent with the function, which coincides with the usual discrete gradient vector field in the non-degenerate case. Our symbolic perturbation scheme can be interpreted as providing a Morse function that is consistent with the given gradient vector field and is arbitrarily close to the given pseudo-Morse function.

A refinement of this scheme additionally allows to relax the assumption that critical cells have unique function values. It extends the first perturbation scheme by explicitly maintaining a total order on the cells that is consistent with both the function and the gradient vector field.

1.1.4 Extensions of discrete Morse theory

We extend the results of discrete Morse theory in several aspects. In Section 2.7, we investigate how the structure of critical points changes when interpolating linearly between two discrete Morse functions. It turns out that the results are significantly easier than in the smooth case, where a corresponding statement has been obtained by Cerf [14].

In Section 2.8 we discuss the interplay between piecewise linear (PL) functions and discrete Morse functions. We show how to translate from PL functions to discrete (pseudo-)Morse functions and vice versa. This equivalence can be used to solve certain problems involving PL functions using discrete Morse theory. Vice versa, it also allows to apply certain theorems regarding PL functions to our setting.

While discrete Morse theory is mostly concerned with regular CW complexes, the theory can also be applied to general CW complexes under certain assumptions. Forman [28, 30] considers a discrete gradient vector field as a certain set of pairs of cells (ϕ, ρ) , where ϕ is a regular face of ρ . In Section 2.9, we show how these assumptions can be weakened. In particular, the fundamental statements of discrete Morse theory also hold without the assumption that the closure of ϕ is a topological disk.

1.1.5 Multiple cancelations

In principle, multiple pairs of critical points can be eliminated by sequentially applying the (single pair) cancelation method of Forman [28]. However, as a prerequisite for cancelation, this method relies on the existence of a unique gradient path between the pair of critical points to be canceled. For every step in a cancelation sequence, this requires finding an *admissible pair* of critical points that satisfies this prerequisite. It is natural to ask whether such admissible cancelation pairs can be related to *persistence pairs* in some way. In particular, can every persistence pair eventually be eliminated using a cancelation sequence?

For the case of surfaces, the answer is indeed affirmative *if* persistence pairs are canceled in a *nested* order. Concretely, we show that a persistence pair (σ, τ) can be canceled after all persistence pairs $(\tilde{\sigma}, \tilde{\tau})$ with $f(\sigma) < f(\tilde{\sigma}) < f(\tilde{\tau}) < f(\tau)$ have been removed (assuming, by symbolic perturbation, that all critical points have distinct values). This result is established by introducing a certain hierarchy

on the persistence pairs with indices $(0, 1)$, which by duality extends to the pairs with indices $(1, 2)$, see Section 4.2.

We note that in general this statement does neither hold for manifolds of dimension greater than two nor for non-manifold 2-complexes.

1.1.6 Tightness of the stability bound

The *Bottleneck Stability Theorem* [18], a fundamental result in the theory of persistent homology, provides a lower bound on the number of critical points:

Proposition (Stability Bound). *For any function f_δ with $\|f_\delta - f\|_\infty \leq \delta$, the number of critical points of f_δ is bounded from below by the number of those critical points of f that have persistence $> 2\delta$.*

Based on a nested sequence of persistence pairs, our cancelation method is capable of removing all pairs with persistence $\leq 2\delta$ without removing other critical points. In view of the *topological simplification problem*, it remains to show that the corresponding sequential cancelations do not violate the δ -tolerance constraint. Indeed, we show that a nested cancelation sequence of persistence pairs leads to a function that matches the stability bound (Theorem 4.2), providing an optimal solution to the topological simplification problem:

Theorem (Tightness of the stability bound). *Given a function f on a surface and a real number $\delta \geq 0$, there exists a function f_δ such that $\|f_\delta - f\|_\infty \leq \delta$ and the number of critical points of f_δ equals the number of those critical points of f that have persistence $> 2\delta$.*

A similar statement does not hold in higher dimensions or for non-manifold 2-complexes, see Section 6.6.

The main difficulty in proving this result stems from the fact that even canceling a single persistence pair might affect an arbitrary number of other critical points. Consequently, it is not possible to *independently* analyze the cancelation of persistence pairs with indices $(0, 1)$ and $(1, 2)$, respectively. This stands in contrast to previous related methods [25, 2], which modify the function independently on two disjoint subsets of the surface, corresponding to the cancelation of persistence pairs of indices $(0, 1)$ and $(1, 2)$, respectively.

1.1.7 Efficient solution

Our construction of sequential cancelations is convenient for establishing tightness of the stability bound, but yields an algorithm with quadratic time complexity. Devising a more efficient algorithm requires a slightly different approach.

Our (inefficient) sequential construction amounts to *alternating* between two processes: (i) canceling persistence pairs from a discrete vector field and (ii) adapting the function correspondingly. In order to obtain an efficient solution, we perform these processes *separately* instead: In a first step, we compute a fully simplified vector field from the persistence pairs; in a second step, we construct a simplified function from the input function and the simplified vector field resulting from the first step. Both steps can be performed in time $O(n)$ using simple graph traversal methods (Chapter 5). Note that computing persistence pairs takes time $O(\text{sort}(n))$ [2]. Here n denotes the number of cells of the surface and $O(\text{sort}(n))$ denotes the time complexity of sorting n numbers.

The existence of an efficient algorithm for the topological simplification problem is perhaps surprising in view of the fact that the problem is NP-hard when restricted to simplexwise linear functions on a triangulated surface. The NP-hardness result follows from extending a recent result by Gray et al. [33], which states that minimizing the number of extrema of a simplexwise linear function with interval constraints for the vertex values is NP-hard. We adapt the argument to our problem setting where all tolerance intervals are assumed to have length 2δ (see Section 6.10). Here, a *simplexwise* linear function on a triangulated surface is a function that is linear on each simplex of a *fixed* triangulation, while a *piecewise* linear function is only required to be linear on each simplex of *some* triangulation of the same surface. In particular, a simplexwise linear function on a subdivision of a triangulation is piecewise linear but not necessarily simplexwise linear with respect to the original triangulation. The emphasis on simplexwise linear as opposed to just piecewise linear functions is crucial here: a multiple saddle can be split into several non-degenerate saddles by an arbitrarily small perturbation (in the supremum norm) in the space of piecewise linear functions, but not in the subspace of simplexwise linear functions. This emphasizes the important role of discrete Morse theory: the hardness of the problem in the simplexwise linear setting arises from the possibility that the input contains multiple saddles, which is excluded by definition in discrete Morse theory. A discrete pseudo-Morse function

constructed from a simplexwise linear function can again be interpreted as a simplexwise linear function on the barycentric subdivision of the original triangulation (Section 2.8).

1.1.8 Energy minimization of simplified functions

The solution to the topological simplification problem is not unique in general: both the δ -constraint and the simplified discrete gradient vector field impose a set of linear inequalities on the simplified function, so the solution set is a convex polytope. This additionally allows to minimize a suitable convex energy functional. We employ this technique to remove artifacts from the initial solution and to improve the similarity to the input function (Section 6.5).

1.2 Related work

Topological simplification of functions within a δ -tolerance constraint has been considered before by Edelsbrunner et al. [25] and Attali et al. [2]. The problem considered there differs from ours by a seemingly small but significant detail: in [25, 2] the critical points of the input function f that are *not* eliminated are additionally assumed to exactly maintain their original values. This restriction has some drawbacks: while it allows for eliminating all critical points of f with persistence $\leq \delta$, it does not in general allow to eliminate all critical points with persistence $\leq 2\delta$; an example is given in [25]. Hence, under this restriction it is not always possible to match the stability bound.

The methods presented in [25, 2] can be interpreted as variants of the so-called *carving* method proposed by Soille [61] in the context of terrain simplification. Similarly, there is another popular method for removing extrema from terrains, called *filling* or *flooding* [37, 1, 21]. A combination of both methods has been proposed in [62]. Our method of canceling critical points from a function can be interpreted as a combination of carving and flooding in the realm of discrete Morse theory.

Apart from the above mentioned works, persistent homology provides the basis for several other methods for computing and simplifying multi-scale structures derived from a function. In their original paper on persistent homology, Edelsbrunner et al. [23] discuss the closely related problem of simplifying a *filtration* of a simplicial complex. Edelsbrunner et al. [24] and Gyulassy et al. [34]

consider simplification of *cell decompositions* (Morse–Smale complexes) resulting from a given gradient vector field. Unfortunately, a simplified Morse–Smale complex does not directly give rise to a simplified *function*. Indeed, simplifying a Morse–Smale complex is closely related to simplifying a discrete *gradient vector field*.

The problem of constructing discrete gradient vector fields that minimize the number of critical points *without* constraints is addressed by Lewiner et al. [49] for surfaces and by Joswig and Pfetsch [39] for complexes of arbitrary dimension. King et al. [42] were the first to propose the combination of persistence with discrete Morse theory to simplify the gradient vector field of an input function on a 3-dimensional simplicial complex. Their method has quadratic time complexity and produces a simplified discrete gradient vector field but not a function. It does not aim at optimality, since on 3-dimensional complexes in general not every persistence pair can be canceled (see Section 6.7).

1.3 Acknowledgements

I would like to thank my advisor Max Wardetzky for the exceptional guidance and encouragement, and for providing me the freedom to work on the topics that interested me most. Thanks to Carsten Lange for introducing me to discrete Morse theory and for always taking the time to patiently answer my questions. I am deeply indebted to Henrik Schumacher for meticulously proof-reading a manuscript of this thesis. Thanks to Emanuel Huhnen-Venedey and Yassin Sabih for providing valuable feedback. Thanks to Konrad Polthier for the support and funding during my time at Freie Universität Berlin and to all members of the group, in particular to Klaus Hildebrandt for many interesting discussions. Thanks to my parents and family for supporting me in all possible ways. And finally, thanks to Hannah for her love, support and patience.

Some of the results of this thesis also appeared in a joint article with Carsten Lange and Max Wardetzky [6].

The “Puget Sound” data set used in Figs. 6.3 to 6.5 is taken from the Large Geometric Models Archive of the Georgia Institute of Technology [48]. The original elevation data are obtained from The United States Geological Survey (USGS), made available by the University of Washington. The medical dataset used in Fig. 7.1 is taken from the VolVis page at the University of Tübingen [64].

1.3 Acknowledgements

The original data are courtesy of Özlem Gürvit, Institute for Neuroradiology, Frankfurt.

Chapter 2

Discrete Morse theory

Classical (smooth) Morse theory [51] relates the critical points of a generic smooth real-valued function on a manifold to the global topology of that manifold. Forman [28, 30] carried over the main ideas of Morse theory to a combinatorial setting. We introduce the main concepts and results of discrete Morse theory, together with some extensions to Forman's theory that provide important tools for our results. To emphasize the connection to the smooth theory, we present the smooth counterparts alongside with the corresponding definitions and theorems of discrete Morse theory.

2.1 CW complexes

A *CW complex* \mathcal{K} is a topological space constructed inductively. Starting with a discrete set \mathcal{K}_0 of 0-cells called the *0-skeleton*, we form the *n-skeleton* \mathcal{K}_n of \mathcal{K} by attaching closed *n-cells* (closed *n-dimensional balls*) by continuous *attaching maps* $S^{n-1} \rightarrow \mathcal{K}_{n-1}$ from their boundary to the $(n-1)$ -skeleton. More precisely, \mathcal{K}_n is the quotient space of the disjoint union of \mathcal{K}_{n-1} with a collection of closed *n-balls* under the identifications given by the attaching maps. Throughout this thesis, we consider only finite CW complexes, meaning that $\mathcal{K} = \mathcal{K}_d$ for some large enough d . A CW complex is always a Hausdorff space, i.e., disjoint points have disjoint neighborhoods.

The image of one of the closed *n-dimensional balls* by the quotient map is called a *closed cell* of \mathcal{K} , while the image of its interior is called an *open cell*. The set of open cells is denoted by K . Whenever we simply talk about a cell, we quietly assume the cell to be *open*. For a cell σ , the restriction of the quotient map to the corresponding closed *n-ball* yields a map $\Phi_\sigma : \overline{B^n} \rightarrow \mathcal{K}_n \subset \mathcal{K}$ from the closed unit ball $\overline{B^n}$, called the *characteristic map* of σ . The restriction of the

characteristic map to the open unit ball B^n is a homeomorphism onto σ , while the restriction to S^{n-1} yields the attaching map.

Whenever a cell $\tau \in K$ is attached to a cell σ (i.e., $\sigma \subset \partial\tau$, where $\partial\tau$ denotes the boundary $\text{bd}_{\mathcal{K}_d} \tau$ of τ in the d -skeleton \mathcal{K}_d with $d = \dim \tau$), we call σ a *face* of τ ; a face of codimension 1 is called a *facet*. If the characteristic map Φ_τ of τ (or, equivalently, the attaching map) restricts to a homeomorphism on the preimage $\Phi_\tau^{-1}(\sigma)$ and the closure of $\Phi_\tau^{-1}(\sigma)$ is a closed ball, then σ is a *regular facet* of τ . If all facets are regular, or equivalently, if all characteristic maps are topological embeddings, i.e., homeomorphisms onto their respective images, then \mathcal{K} is called a *regular CW complex*. If the union of a set of cells is closed, then this union is itself a CW complex, called a *subcomplex*.

The notions of face [28] and facet resemble the standard notions for convex polytopes. However, our definition of face applies to open cells, whereas a face of a convex polytope is considered to be a (closed) convex polytope again. Moreover, a cell of a CW complex is not a face of itself (corresponding to the notion of a *proper* face for polytopes). Note that in general the boundary of a cell is not the union of its faces.

A regular CW complex whose underlying space is a PL manifold is called a *combinatorial manifold*. This name reflects the fact that regular CW complexes are determined up to homeomorphism by the combinatorial data of the face relation. A combinatorial 2-manifold is called a *combinatorial surface*. We refer to [50, 35] for further details on CW complexes.

2.2 Discrete vector fields

One of the central concepts of discrete Morse theory is that of a *discrete vector field* – a purely combinatorial analogue of a classical vector field.

Definition (discrete vector field, critical cell [28, 30]). *A discrete vector field V on a regular CW complex \mathcal{K} is a set of pairs of cells $(\sigma, \tau) \in K \times K$, with σ a facet of τ , such that each cell of K is contained in at most one pair of V . A cell $\sigma \in K$ is critical with respect to V if σ is not contained in any pair of V . The dimension of a critical cell is also called its index.*

A pair (σ, τ) in a discrete vector field V can be visualized as an arrow from σ to τ (as in Fig. 2.1). A critical cell is also called a *critical point*.




Figure 2.1: Reversing a gradient vector field along the unique path from $\partial\rho$ to ϕ produces a gradient vector field in which the 1-cell ϕ and the 2-cell ρ are no longer critical.




Figure 2.2: Critical cells of dimension (index) 0, 1, 2

We occasionally drop the attribute *discrete* when this causes no ambiguity.

In the following, we consider an important subclass of vector fields in which the arrows do not form closed paths. This can be made precise using the concept of *V-paths*.

Definition (*V-path* [30]). Let V be a discrete vector field. A V -path Γ from a cell σ_0 to a cell σ_r is a sequence $(\sigma_0, \tau_0, \sigma_1, \dots, \tau_{r-1}, \sigma_r)$ of cells such that for every $0 \leq i \leq r - 1$:

- σ_i is a facet of τ_i with $(\sigma_i, \tau_i) \in V$, and
- σ_{i+1} is a facet of τ_i with $(\sigma_{i+1}, \tau_i) \notin V$.

A V -path Γ is called closed if $\sigma_0 = \sigma_r$ and nontrivial if $r > 0$. We call $\dim \sigma_0$ the dimension of Γ .

If ρ, ϕ are critical cells of V , then a V -path from $\partial\rho$ to ϕ means a V -path $\Gamma = (\sigma_0, \tau_0, \sigma_1, \dots, \tau_{r-1}, \sigma_r)$ with σ_0 a facet of ρ and $\phi = \sigma_r$ (see Fig. 2.1 for an example). We also say that Γ is a V -path between ρ and ϕ .

Definition (discrete gradient vector field [30]). A discrete vector field V is a discrete gradient vector field if it contains no nontrivial closed V -paths.

2.2.1 Reversing gradient vector fields

The main technique for reducing the number of critical points is that of *reversing* a gradient vector field V along a V -path between two critical cells ρ and ϕ .

Consider a V -path $(\sigma_0, \tau_0, \sigma_1, \dots, \tau_{r-1}, \sigma_r)$ from $\partial\rho$ to ϕ . Now replace every pair $(\sigma_i, \tau_i) \in V$ by (σ_{i+1}, τ_i) for $0 \leq i \leq r-1$ and add the pair (σ_0, ρ) to obtain a discrete vector field \widetilde{V} where ϕ and ρ are no longer critical (see Fig. 2.1 for an example). Moreover, if there is no other V -path from $\partial\rho$ to ϕ , then \widetilde{V} is again a gradient vector field because no nontrivial closed \widetilde{V} -path is created.

Theorem 2.1 (Forman [28]). *Let ϕ and ρ be two critical cells of a gradient vector field V with exactly one V -path Γ from $\partial\rho$ to ϕ . Then there is a gradient vector field \widetilde{V} obtained by reversing V along the path Γ . The critical cells of \widetilde{V} are exactly the critical cells of V apart from $\{\phi, \rho\}$. Moreover, $V = \widetilde{V}$ except along the path Γ .*

This provides a discrete analogue of Morse's cancelation theorem [57], which was prominently used by Milnor [52] in his Morse-theoretic proof of the h -cobordism Theorem. Before stating the theorem, we introduce some definitions.

A smooth function f on a manifold is a *Morse function* if the Hessian at each critical point p is non-degenerate, i.e., all eigenvalues are different from 0. According to the Morse lemma [51, Lemma 2.2], there is a chart centered at p with coordinates (x_1, \dots, x_n) such that f has the standard form

$$f(x_1, \dots, x_n) = f(p) - x_1^2 - \dots - x_\lambda^2 + x_{\lambda+1}^2 + \dots + x_n^2.$$

Such a chart is called a *Morse chart*. The number λ is the number of negative eigenvalues of the Hessian at p . It is called the *index* of p . A smooth vector field is *gradient-like* if there is a Morse function f such that $X(f) > 0$ away from the critical points of f , and for each critical point of f there is a Morse chart such that $X = \nabla f$ with respect to the Euclidean metric on that chart.

Theorem 2.2 (Milnor [52], Theorem 5.4). *Let ϕ and ρ be two critical points of consecutive index of a gradient-like vector field X on a manifold M with exactly one flow line γ from ρ to ϕ . Then there is a gradient-like vector field \widetilde{X} obtained by modifying X in an arbitrarily small neighborhood U of the path γ such that the critical points of \widetilde{X} are exactly the critical points of X apart from $\{\phi, \rho\}$. Moreover, $X \equiv \widetilde{X}$ except at U .*

2.2.2 Gradient vector fields on surfaces

Gradient vector fields on combinatorial surfaces have additional properties that do not hold in higher dimensions. The following property is readily checked using the fact that a 1-cell is only attached to at most two 0-cells, and at most two 2-cells are attached to a 1-cell:

Lemma 2.3. *Two V -paths of dimension 0 cannot branch at a common cell, and two V -paths of dimension 1 cannot merge (except at their last cell).*

Proof. Let $\Gamma = (\sigma_0, \tau_0, \sigma_1, \dots, \tau_{r-1}, \sigma_r)$, $\tilde{\Gamma} = (\tilde{\sigma}_0, \tilde{\tau}_0, \tilde{\sigma}_1, \dots, \tilde{\tau}_{s-1}, \tilde{\sigma}_s)$ be two V -paths of dimension 0. We show that the paths cannot branch: if $\sigma_i = \tilde{\sigma}_j$, then $\tau_i = \tilde{\tau}_j$ as well; if $\tau_i = \tilde{\tau}_j$, then $\sigma_{i+1} = \tilde{\sigma}_{j+1}$ as well. The second statement can be shown in a similar way. First, note that the successor of $\sigma_i = \tilde{\sigma}_j$ in any V -path is uniquely determined by the vector field V . Second, the 1-cell $\tau_i = \tilde{\tau}_j$ is attached to exactly two 0-cells, one of which is $\sigma_i = \tilde{\sigma}_j$. By the definition of a V -path, the other 0-cell must be the successor $\sigma_{i+1} = \tilde{\sigma}_{j+1}$ of $\tau_i = \tilde{\tau}_j$ in both Γ and $\tilde{\Gamma}$. \square

Corollary 2.4. *Let ρ be a critical 1-cell of a discrete vector field V on a combinatorial surface. Then there are at most two V -paths from $\partial\rho$ to critical 0-cells, each starting at one of the two 0-cells in $\partial\rho$. Similarly, there are at most two V -paths from facets of critical 2-cells to ρ .*

2.3 The Morse complex

A gradient vector field V gives rise to the construction of the *Morse complex* \mathcal{M}_V , a (possibly non-regular) CW complex that is homotopy equivalent to the original CW complex \mathcal{K} but has fewer cells.

Theorem 2.5 (Forman [28], Corollary 3.5). *Let V be a discrete gradient vector field on a finite regular CW complex \mathcal{K} . Then \mathcal{K} is homotopy equivalent to the Morse complex \mathcal{M}_V , which is a CW complex with as many d -cells as there are critical cells of V of dimension d .*

This theorem is an analogue of the following classical result:

Theorem 2.6 (Milnor [51], Theorem 3.5). *Let f be a smooth Morse function on a manifold M . Then M is homotopy equivalent to a CW-complex with as many d -cells as there are critical points of f of index d .*

The Morse complex and its attaching maps have an intuitive description: \mathcal{M}_V can be thought of as the result of internally collapsing each pair of the gradient vector field V simultaneously. In this process, all non-critical cells vanish, while the *critical* cells are glued together to form the new complex \mathcal{M}_V .

2.3.1 Cellular homology

We revisit some basic definitions of cellular homology [50, 35]. The cellular chain group of \mathcal{K} with coefficients in an abelian group G is defined as the abelian group $C_*(\mathcal{K}; G)$ consisting of formal sums $\sum_{\sigma \in K} g_\sigma \sigma$, where $g_\sigma \in G$. The restriction to formal sums of cells of dimension i is denoted by $C_i(\mathcal{K}; G)$. The cellular boundary operator $\partial : C_*(\mathcal{K}; G) \rightarrow C_*(\mathcal{K}; G)$ is defined by

$$\partial \sum_{\tau \in K} g_\tau \tau = \sum_{\sigma, \tau \in K} \partial_{\tau, \sigma} g_\tau \sigma,$$

where G is considered as an additive group and the coefficients $\partial_{\tau, \sigma} \in \mathbb{Z}$ are given as follows. If σ is a facet of τ with $\dim \sigma = d$, let $\varphi_\tau : S^d \rightarrow \mathcal{K}_d$ be the attaching map of τ and let $s_\sigma : \mathcal{K}_d \rightarrow S^d$ be the quotient map sending $\mathcal{K}_d \setminus \sigma$ to a point. Define $\partial_{\tau, \sigma}$ as the degree of the map $s_\sigma \circ \varphi_\tau : S^d \rightarrow S^d$. If σ is not a facet of τ , let $\partial_{\tau, \sigma} = 1_G$ (the neutral element of G). The restriction of ∂ to a map $C_i(\mathcal{K}; G) \rightarrow C_{i-1}(\mathcal{K}; G)$ is denoted by ∂_i . The sequence of maps ∂_i is called the cellular chain complex of \mathcal{K} . The kernel of ∂ is called the group of cycles $Z_*(\mathcal{K}; G)$, the image of ∂ is called the group of boundaries $B_*(\mathcal{K}; G)$. The quotient $H_*(\mathcal{K}; G) = Z_*(\mathcal{K}; G)/B_*(\mathcal{K}; G)$ is called the homology of $(C_*(\mathcal{K}; G), \partial)$. Again, the subgroups $Z_i(\mathcal{K}; G)$, $B_i(\mathcal{K}; G)$, and $H_i(\mathcal{K}; G)$ are obtained by restriction to dimension i .

2.3.2 Homology of the Morse complex

The boundary maps $\tilde{\partial}$ of the cellular chain complex $(C_*(\mathcal{M}_V, \mathbb{Z}), \tilde{\partial})$ for the Morse complex can be obtained by counting all V -paths between critical cells, taking orientation into account [28]. This can be understood as a discrete analogue to the definition of a cellular chain complex via the degree of attaching maps. Another description can be given by considering the (discrete time) *gradient flow* $\Phi : C_*(\mathcal{K}, \mathbb{Z}) \rightarrow C_*(\mathcal{K}, \mathbb{Z})$ associated to V . To do so, we interpret V as the unique linear map satisfying

$$V(\rho) := \begin{cases} -\langle \partial \psi, \rho \rangle \psi & \text{if there is a pair } (\rho, \psi) \in V, \\ 0 & \text{otherwise.} \end{cases}$$

This gradient flow Φ is now defined by linear extension of

$$\Phi(\tau) := \tau + V(\partial \tau) + \partial V(\tau), \quad \tau \in K.$$

Here we use the canonical inner product $\langle \cdot, \cdot \rangle$ on $C_*(\mathcal{K}, \mathbb{Z})$ defined by declaring the oriented cells to be an orthonormal basis, i.e., $\langle \phi, \rho \rangle = \delta_{\phi, \rho}$. According to [28, Theorem 7.2], this map Φ stabilizes in finite time to a map Φ^∞ , that is, there is an N large enough such that $\Phi^N = \Phi^{N+1} = \dots = \Phi^\infty$. In particular, Φ^∞ maps a critical cell $\tau \in K$ to a Φ -invariant chain of $C_*(\mathcal{K}, \mathbb{Z})$. It turns out that the d -cells of \mathcal{M}_V are in bijection with the set $\{\Phi^\infty(\tau) \mid \tau \in K \text{ critical } d\text{-cell of } V\}$ of invariant chains generated by a critical cell, and that the boundary map $\tilde{\partial}$ of \mathcal{M}_V can be expressed in terms of the boundary map ∂ of \mathcal{K} using this identification. In other words, we can identify each critical cell ρ of V bijectively with a cell $\tilde{\rho} \in \mathcal{M}_V$ such that the following equation [28, (8.3)] holds whenever σ and τ are critical cells of V :

$$\langle \tilde{\partial}\tilde{\tau}, \tilde{\sigma} \rangle = \langle \partial\Phi^\infty(\tau), \sigma \rangle = \langle \Phi^\infty(\partial\tau), \sigma \rangle.$$

Note that in contrast to our nomenclature, Forman [28] uses the name *Morse complex* to denote the cellular chain complex $(C_*(\mathcal{M}_V, \mathbb{Z}), \tilde{\partial})$ of \mathcal{M}_V . The cellular homology of \mathcal{M}_V is also called the *Morse homology* of (\mathcal{K}, V) . It is isomorphic to the cellular homology of \mathcal{K} since \mathcal{M}_V is homotopy equivalent to \mathcal{K} and cellular homology is a homotopy invariant. The same names are also used for the analogous concepts of smooth Morse theory [4].

2.4 Morse and pseudo-Morse functions

As in smooth Morse theory, a discrete gradient vector field can be understood as the gradient of some non-degenerate function in the following sense:

Definition (discrete Morse function [28]). *A function $f : K \rightarrow \mathbb{R}$ on the cells of a regular CW complex \mathcal{K} is a discrete Morse function if there is a gradient vector field V_f such that whenever σ is a facet of τ ,*

- $(\sigma, \tau) \notin V_f$ implies $f(\sigma) < f(\tau)$, and
- $(\sigma, \tau) \in V_f$ implies $f(\sigma) \geq f(\tau)$.

V_f is called the gradient vector field of f .

In contrast to piecewise linear functions, which are determined by their function values at the vertices of the triangulation, discrete Morse functions take values on cells of *any* dimension.

Although this is not evident from the definitions, discrete Morse functions are indeed very closely related to smooth Morse functions in the sense that their respective critical points share the same characteristic properties. This will be clarified later in Section 2.6.

Every gradient vector field can be associated with some discrete Morse function:

Theorem 2.7 (Forman [28]). *Let W be a discrete gradient vector field. Then there is a discrete Morse function f with $V_f = W$.*

An explicit construction of such a function will result from the techniques in Section 2.5.

The gradient vector field of a discrete Morse function encodes only the *sign* of the difference between function values, not the difference itself. Therefore a discrete gradient vector field does not uniquely determine a discrete Morse function, just like a smooth gradient-like vector field X does not uniquely determine a Morse function f satisfying the consistency condition $X(f) > 0$. However, there is exactly one gradient vector field V_f for every discrete Morse function f . In contrast, in the smooth theory there are many gradient-like vector fields for a given Morse function: different choices of Riemannian metric generally lead to different gradient vector fields (note that not every gradient vector field is actually gradient-like for a given metric because of the additional technical condition at the critical points).

In order to be able to treat non-generic functions, it is useful to consider a more general class of functions, which we call *pseudo-Morse functions*. Pseudo-Morse functions substitute the strict inequality in the definition of Morse functions by a weak one.

Definition (pseudo-Morse function, consistency). *A function $f : K \rightarrow \mathbb{R}$ on the cells of a regular CW complex \mathcal{K} is a discrete pseudo-Morse function if there is a gradient vector field V such that whenever σ is a facet of τ ,*

- $(\sigma, \tau) \notin V$ implies $f(\sigma) \leq f(\tau)$, and
- $(\sigma, \tau) \in V$ implies $f(\sigma) \geq f(\tau)$.

In this case, we call f and V consistent.

Note that a gradient vector field V consistent with a pseudo-Morse function f is not unique in general.

2.4.1 Properties of pseudo-Morse functions

As a direct consequence of the definitions we have the following property of positive linear combinations of (pseudo-)Morse functions:

Lemma 2.8. *Let f and g be two pseudo-Morse functions consistent with a gradient vector field V and let $\lambda, \mu \geq 0$. Then $\lambda f + \mu g$ is also consistent with V . Moreover, if f is a Morse function and $\lambda > 0$, then $\lambda f + \mu g$ is also Morse.*

From this fact we can derive a useful characterization of discrete pseudo-Morse functions. For a given gradient vector field V , the set of pseudo-Morse functions consistent with V is the closure of the set of Morse functions consistent with V with respect to the standard topology on the set of functions $K \rightarrow \mathbb{R}$ as a vector space isomorphic to $\mathbb{R}^{|K|}$. However, not every function $K \rightarrow \mathbb{R}$ is pseudo-Morse.

Theorem 2.9. *Let $f : K \rightarrow \mathbb{R}$ be a function on the cells of a regular CW complex \mathcal{K} and let V be a gradient vector field on \mathcal{K} . Then f is a discrete pseudo-Morse function consistent with V if and only if for every $\epsilon > 0$ there is a discrete Morse function $f_\epsilon : K \rightarrow \mathbb{R}$ with $\|f_\epsilon - f\|_\infty \leq \epsilon$ such that V is the gradient vector field of f_ϵ .*

Proof. Assume that f is a pseudo-Morse function consistent with a gradient vector field V . Then according to Theorem 2.7, there exists a discrete Morse function g whose gradient vector field V_g is precisely given by V . Let G be the maximum absolute value of g . Given $\epsilon > 0$, for each cell σ define

$$f_\epsilon(\sigma) := f(\sigma) + \epsilon \frac{g(\sigma)}{G}.$$

Lemma 2.8 implies that f_ϵ is a discrete Morse function with gradient vector field V , and clearly we have $\|f_\epsilon - f\|_\infty \leq \epsilon$.

On the other hand, assume that for every $\epsilon > 0$ there is a discrete Morse function $f_\epsilon : K \rightarrow \mathbb{R}$ consistent with V and $\|f_\epsilon - f\|_\infty \leq \epsilon$. Choose ϵ such that for every $\phi, \rho \in K$ with $f(\phi) \neq f(\rho)$ we have

$$\epsilon < \frac{|f(\phi) - f(\rho)|}{2}.$$

Without loss of generality, assume that $f(\phi) > f(\rho)$. Then $f(\phi) - f(\rho) > 2\epsilon$ and hence $f_\epsilon(\phi) - f_\epsilon(\rho) > 0$ or $f_\epsilon(\phi) > f_\epsilon(\rho)$. Since f_ϵ is consistent with V , one easily verifies that f is consistent with V as well. \square

2.5 Symbolic perturbation

Theorem 2.9 suggests a *symbolic perturbation* scheme based on gradient vector fields in order to allow for non-generic (degenerate) input functions. Starting with a pseudo-Morse function f , we can choose a consistent gradient vector field V , which may not be unique. Theorem 2.9 asserts that there is a discrete Morse function f_ϵ arbitrarily close to f and consistent with V . Therefore we can work with f as if it were a discrete Morse function with gradient vector field V . In particular, we use Theorem 2.9 to associate critical points to a pseudo-Morse function by choosing a consistent gradient vector field.

This first symbolic perturbation scheme is not sufficient for all our purposes; the definition of persistence pairs given later in Chapter 3 not only requires a gradient vector field, but also a total order on the critical cells, which again is not always uniquely defined by a pseudo-Morse function f and a consistent gradient vector field V . We now derive a second perturbation scheme that meets these requirements.

2.5.1 The partial order induced by a gradient vector field

Since a gradient vector field imposes certain inequality constraints on the functions consistent with it, we can ask how these inequalities affect the relation between the function values of any two cells. We observe that any discrete gradient vector field gives rise to a strict partial order on the set of cells:

Definition (induced partial order). *Let V be a discrete gradient vector field and consider the relation \leftarrow_V defined on K such that whenever σ is a facet of τ ,*

- $(\sigma, \tau) \notin V$ implies $\sigma \leftarrow_V \tau$, and
- $(\sigma, \tau) \in V$ implies $\sigma \rightarrow_V \tau$,

where \rightarrow_V is the inverse of \leftarrow_V . Let $<_V$ be the transitive closure of \leftarrow_V . Then $<_V$ is called the (strict) partial order induced by V .

It is readily checked that this relation is indeed a strict partial order since V is assumed to have no non-trivial closed V -paths. The interpretation of this partial order is that for *any* pseudo-Morse function f consistent with V and any two cells ϕ and ρ , the relation $\phi <_V \rho$ implies $f(\phi) \leq f(\rho)$. This fact is easily verified from the definitions. There is an interpretation in terms of V -paths: $\phi <_V \rho$

implies that there is a sequence of V -paths, the first one starting at $\partial\rho$, every other path starting at a facet of the end cell of the previous path, and the last one ending at ϕ or a facet thereof. In particular, $\dim \phi \leq \dim \rho + 1$.

The relation \leftarrow_V is the *covering relation* of $<_V$, i.e., $\phi \leftarrow_V \rho$ implies that $\phi <_V \rho$ and that there is no ψ with $\phi <_V \psi <_V \rho$. The covering relation of a partial order forms a directed acyclic graph called the *Hasse diagram*. We use the convention that the edges are oriented as suggested by the arrow symbol \leftarrow_V . The Hasse diagram H_V of $<_V$ is obtained from the Hasse diagram of the face relation on K by inverting the orientation of all edges corresponding to pairs $(\sigma, \tau) \in V$, as described by Chari [15]. H_V has the property that $\phi <_V \rho$ if and only if there is a directed path from ρ to ϕ . Note that $\sigma \leftarrow_V \tau$ implies $f(\sigma) \leq f(\tau)$, i.e., both the arrow visualizing $(\sigma, \tau) \in V$ and the arrow symbolizing $\sigma \leftarrow_V \tau$ point towards a (weakly) *decreasing* function value of f .

A related approach to defining partial orders based on gradient vector fields was proposed by Kozlov [45]. The main difference is that the partial order considered there is defined on the union of the pairs V with the critical cells of V , i.e., the pairs in V are taken as single elements. This approach has the technical disadvantage that one has to work with partial orders on different sets when working with different gradient vector fields. Our approach instead uses different partial orders on the same underlying set, which is more convenient for our purposes.

This partial order on the cells gives rise to the notion of *ascending* and *descending sets*, resembling the ascending and descending manifolds in smooth Morse theory. This notion serves as an important tool in our constructive methods.

Definition (ascending and descending set). *Let V be a discrete gradient vector field on \mathcal{K} . The ascending set of a cell $\phi \in K$ is defined as the upper set of ϕ in \geq_V ,*

$$\{\rho : \rho \geq_V \phi\},$$

and the descending set is the lower set of ϕ in \geq_V ,

$$\{\rho : \rho \leq_V \phi\}.$$

The related smooth objects are defined as follows:

Definition (ascending and descending manifolds [63]). *Let V be a smooth gradient-like vector field on a manifold M and let $\gamma_x : \mathbb{R} \rightarrow M$ be the flow curve of M at x , i.e.,*

$\gamma_x(0) = x$ and $\gamma'_x(t) = V(\gamma_x(t))$. The ascending manifold of a critical point $p \in M$ is defined as the set

$$\{x \in M : \lim_{t \rightarrow \infty} \gamma_x(t) = p\},$$

and the descending manifold of p is the set

$$\{x \in M : \lim_{t \rightarrow -\infty} \gamma_x(t) = p\}.$$

Note that the discrete object actually corresponding to an ascending (descending) manifold is the *ascending (descending) region* of a critical cell ϕ , i.e., those cells in the ascending (descending) set of ϕ that are not contained in the ascending (descending) set of a higher- (lower-) dimensional critical cell than ϕ [38].

Thom [63] observed that the ascending and descending manifolds of a critical point of index λ on a manifold of dimension n are always open balls of dimension $n - \lambda$ and λ , respectively. In contrast, even if \mathcal{K} is a combinatorial manifold, the ascending and descending regions of critical cells need not form a partition of \mathcal{K} into topological balls. It is however possible to subdivide the complex and extend the gradient vector field such that the subset of the ascending (descending) regions become collapsible [38].

2.5.2 Consistent total orders

Assume we are given a pseudo-Morse function f consistent with a gradient vector field V . On the one hand we have the induced partial order $<_V$. On the other hand the function f canonically induces a strict partial order $<_f$ given by

$$\phi <_f \rho \Leftrightarrow f(\phi) < f(\rho).$$

Since the two orders $<_f$ and $<_V$ are compatible by assumption (there are no two cells (ϕ, ρ) with $\phi <_V \rho$ and $\phi >_f \rho$), we can merge them into a strict partial order $<_{f,V}$ (the transitive closure of $(<_f \cup <_V) \subset K \times K$). A linear extension of this order is now called consistent with both f and V :

Definition (consistent total order). *Let V be a discrete gradient vector field V consistent with a discrete pseudo-Morse function f . Then a strict total order $<$ is called consistent with (f, V) if it is a linear extension of $<_f$ and $<_V$.*

Such a total order $<$ gives rise to a canonical bijection $i : K \rightarrow \{1, \dots, |K|\}$, which is a discrete Morse function and consistent with V , providing a simple

proof of Theorem 2.7. If we use this function as the function g in the proof of Theorem 2.9 to construct f_ϵ , then f_ϵ is an injective discrete Morse function with gradient vector field V and the total order induced by f_ϵ is $<$ again. We thus obtain a second symbolic perturbation scheme for situations where a total order on the cells is required.

2.6 Level and order subcomplexes

A classical object of study in smooth Morse theory is the *sublevel set*

$$M(t) = \{x \in M : f(x) \leq t\}$$

of a function $f : M \rightarrow \mathbb{R}$ on a manifold M .

Theorem 2.10 (Milnor [51], Theorem 3.1). *Let f be a smooth real-valued function on a manifold M . If $a < b$ are real numbers such that $f^{-1}[a, b]$ is compact and contains no critical point of f , then $M(b)$ is diffeomorphic to $M(a)$. Moreover, $M(a)$ is a strong deformation retract of $M(b)$, so the inclusion $M(a) \hookrightarrow M(b)$ is a homotopy equivalence.*

Recall that a subspace A of a topological space X is called a *strong deformation retract* if there is a homotopy $F : X \times [0, 1] \rightarrow X$ such that for all $x \in X, a \in A, t \in [0, 1]$ we have $F(x, 0) = x, F(x, 1) \in A$, and $F(a, t) = a$. The homotopy F is called a *strong deformation retraction*.

Theorem 2.11 (Milnor [51], Theorem 3.2 and Remark 3.4). *Let f be a smooth real-valued function on a manifold M . Suppose that σ is a non-degenerate critical point of index d with*

$$a < f(\sigma) \leq b$$

such that $f^{-1}[a, b]$ is compact and contains no other critical points of f . Then $M(b)$ is homotopy equivalent to

$$M(a) \cup_{\varphi} \overline{B^d}$$

with $\varphi : S^{d-1} \rightarrow M(a)$.

In the discrete theory, the analogous object to the sublevel set is the *level subcomplex*, and the equivalent construction using our second symbolic perturbation scheme is the *order subcomplex*:

Definition (level subcomplex [28], order subcomplex). *Let f be a pseudo-Morse function on a regular CW complex \mathcal{K} . Let the carrier of a subset $L \subset K$ be the smallest subcomplex of \mathcal{K} containing all of L . Then for $t \in \mathbb{R}$, the level subcomplex is*

$$\mathcal{K}(t) = \text{carrier} \left(\bigcup_{\rho \in K: f(\rho) \leq t} \rho \right).$$

Similarly, let $<$ be a strict total order on the cells K of a regular CW complex \mathcal{K} . Then for a cell $\sigma \in K$, the order subcomplex is

$$\mathcal{K}(\sigma) = \text{carrier} \left(\bigcup_{\rho \in K: \rho \leq \sigma} \rho \right).$$

The connection between the two notions is as follows. Let $<$ be a strict total order consistent with a discrete gradient vector field V . As mentioned before, $<$ gives rise to a canonical bijection $i : K \rightarrow \{1, \dots, |K|\}$, which is a discrete Morse function with gradient vector field V . Clearly the order subcomplexes $\mathcal{K}(\sigma)$ are in one-to-one correspondence with the level subcomplexes $\mathcal{K}(i(\sigma))$. This equivalence allows to translate statements about level subcomplexes to order subcomplexes. We make use of this fact in the following two theorems.

2.6.1 The homotopy type of order subcomplexes

We can now express the characteristic property of critical points of a discrete Morse function. Like in the smooth theory, the homotopy type of level subcomplexes changes only at critical cells.

Theorem 2.12 (Forman [28], Theorem 3.3). *If $a < b$ are real numbers such that $[a, b]$ contains no critical value of f , then $\mathcal{K}(b)$ collapses to $\mathcal{K}(a)$.*

Recall the definition of a collapse [16, 28]. Consider a CW complex \mathcal{K} . Let τ be a cell of \mathcal{K} . The cell σ is called a *free facet* of τ if σ is a regular facet of τ and σ is not a face of any other cell than τ . If σ is a free regular facet of τ , then $\tilde{\mathcal{K}} = \mathcal{K} \setminus (\sigma \cup \tau)$ is a strong deformation retract of \mathcal{K} and the corresponding homotopy equivalence $\mathcal{K} \rightarrow \tilde{\mathcal{K}}$ is called an *elementary collapse*. A homotopy equivalence which is the composition of elementary collapses is called a *collapse*. If \mathcal{K} collapses to $\tilde{\mathcal{K}}$, we write $\mathcal{K} \searrow \tilde{\mathcal{K}}$.

Theorem 2.13 (Forman [28], Theorem 3.4). *Suppose that σ is a critical cell of index d with*

$$a < f(\sigma) \leq b$$

and $f^{-1}([a, b])$ contains no other critical cells. Then $\mathcal{K}(b)$ is homotopy equivalent to

$$\mathcal{K}(a) \cup_{\varphi} \overline{B^d}$$

with $\varphi : S^{d-1} \rightarrow \mathcal{K}(a)$.

The statements can be rephrased for order subcomplexes:

Theorem 2.14. *Let V be a gradient vector field on a regular CW complex \mathcal{K} and let $<$ be a linear extension of $<_V$. If ρ and ψ are two cells such that $\rho < \psi$ and there is no critical cell ϕ with respect to V such that $\rho < \phi \leq \psi$, then $\mathcal{K}(\psi)$ collapses to $\mathcal{K}(\rho)$.*

Theorem 2.15. *Let V be a gradient vector field on a regular CW complex \mathcal{K} and let $<$ be a linear extension of $<_V$. Suppose that ρ and ψ are two cells such that $\rho < \psi$ and there is exactly one critical cell ϕ of index d with respect to V such that*

$$\rho < \phi \leq \psi.$$

Then $\mathcal{K}(\psi)$ is homotopy equivalent to

$$\mathcal{K}(\rho) \cup_{\varphi} \overline{B^d}$$

with $\varphi : S^{d-1} \rightarrow \mathcal{K}(\rho)$.

Note that we obtain Theorem 2.5 as a direct corollary of Theorems 2.14 and 2.15.

The order subcomplexes provide a finer filtration of the complex \mathcal{K} (by single cells or pairs of cells) than the level subcomplexes, in particular if f is degenerate. This turns out to be useful when working with persistent homology in Chapter 3. As a direct consequence we obtain the following result (see also [51] for the smooth case):

Corollary 2.16 (Morse inequalities; Forman [28], Theorems 1.7 and 1.8). *Let V be a gradient vector field on a regular CW complex \mathcal{K} . Let c_d denote the number of d -cells of \mathcal{K} and let m_d denote the number of critical d -cells of V . Furthermore, let F be a field*

of coefficients and let $\beta_d = \text{rank } H_{d-1}(\mathcal{K}, F)$ denote the d^{th} Betti number of \mathcal{K} . Then for each d ,

$$\sum_{i=0}^d (-1)^{d-i} c_i \geq \sum_{i=0}^d (-1)^{d-i} \beta_i.$$

In particular,

$$c_d \geq \beta_d$$

and

$$\chi(\mathcal{K}) = \sum_{i \geq 0} (-1)^i \beta_i = \sum_{i \geq 0} (-1)^i c_i,$$

where $\chi(\mathcal{K})$ denotes the Euler characteristic of \mathcal{K} .

2.6.2 Flat pseudo-Morse functions

It is evident from the definition that a cell σ might already appear in a level subcomplex $\mathcal{K}(t)$ with $t < f(\sigma)$, or in an order subcomplex $\mathcal{K}(\tau)$ with $\tau < \sigma$. This possibility is somewhat counterintuitive, and it is sometimes helpful or even necessary to work with functions where this situation is excluded. Such functions are characterized by the following definition, extending a definition due to Forman [29]:

Definition (Flat pseudo-Morse function). *A pseudo-Morse function f consistent with a discrete gradient vector field V on a CW complex \mathcal{K} is a flat pseudo-Morse function if*

$$\sigma \text{ is a facet of } \tau \Rightarrow \begin{cases} f(\sigma) \leq f(\tau) & \text{if } (\sigma, \tau) \notin V, \\ f(\sigma) = f(\tau) & \text{if } (\sigma, \tau) \in V. \end{cases}$$

An alternative characterization of flat pseudo-Morse is given by the following statement, which is a direct consequence of the definitions:

Proposition 2.17. *A pseudo-Morse function f is flat if and only if it is consistent with the empty vector field \emptyset .*

We now describe how to obtain a flat pseudo-Morse function from an arbitrary pseudo-Morse function such that the two functions are equivalent with respect to their level subcomplexes. First, recall that pseudo-Morse functions do not necessarily have a unique consistent gradient vector field. There is however a certain unique gradient vector field associated with every pseudo-Morse function:

Definition (minimal consistent vector field). *Let f be a pseudo-Morse function. Then*

$$V = \{(\sigma, \tau) : \sigma \text{ is a facet of } \tau \text{ and } f(\sigma) > f(\tau)\}$$

is the minimal vector field consistent with f .

The name is justified by the following property, which is again a direct consequence of the definitions:

Proposition 2.18. *The minimal vector field consistent with a pseudo-Morse function f is the intersection of all gradient vector fields consistent with f .*

We can explicitly construct a flat pseudo-Morse function from a given pseudo-Morse function.

Definition (flattening). *Let f be a pseudo-Morse function and let V be the minimal vector field consistent with f . Then*

$$\bar{f}(\sigma) := \begin{cases} f(\tau) & \text{if } (\sigma, \tau) \in V \text{ for some } \tau, \\ f(\sigma) & \text{otherwise.} \end{cases}$$

is the flattening of f .

It is obvious from the definition that \bar{f} is indeed flat. From a Morse-theoretic point of view, f and \bar{f} are equivalent:

Proposition 2.19. *Let f be a pseudo-Morse function and let \bar{f} be the flattening of f . Then the corresponding level subcomplexes of f and \bar{f} coincide.*

The statement is again directly verified from the definitions. Another important property is that flattening two pseudo-Morse functions does not increase their distance in the supremum norm. We require the following lemma:

Lemma 2.20. *Let $a, b, c, d \in \mathbb{R}$. Then*

$$|\max(a, b) - \max(c, d)| \leq \max(|a - c|, |b - d|).$$

Proof. Without loss of generality, assume $a > b$ and $\max(a, b) > \max(c, d)$. Then $|\max(a, b) - \max(c, d)| = a - \max(c, d) \leq a - c \leq \max(|a - c|, |b - d|)$. \square

Theorem 2.21. *Let f, g be pseudo-Morse functions on \mathcal{K} and let \bar{f}, \bar{g} be their respective flattenings. Then $\|\bar{f} - \bar{g}\|_\infty \leq \|f - g\|_\infty$.*

Proof. Consider an arbitrary cell $\tau \in K$. We show that $|\bar{f}(\tau) - \bar{g}(\tau)| \leq \|f - g\|_\infty$.

Let V be a discrete gradient vector field consistent with f . If τ is a critical cell of V , let $\sigma = \tau$; otherwise, let σ be the cell such that $(\sigma, \tau) \in V$. Define ϕ for g analogously to how σ is defined for f .

By the definition of a flattening we have $\bar{f}(\tau) = f(\sigma)$ and $\bar{f}(\phi) \geq f(\phi)$. By the definition of a pseudo-Morse functions we have $\bar{f}(\tau) \geq \bar{f}(\phi)$. Together we have $\bar{f}(\tau) = \max(f(\sigma), f(\phi))$, and analogously we also obtain $\bar{g}(\tau) = \max(g(\sigma), g(\phi))$. Thus, by Lemma 2.20 we have

$$\begin{aligned} |\bar{f}(\tau) - \bar{g}(\tau)| &\leq |\max(f(\sigma), f(\phi)) - \max(g(\sigma), g(\phi))| \\ &\leq \max(|f(\sigma) - g(\sigma)|, |f(\phi) - g(\phi)|) \\ &\leq \|f - g\|_\infty. \end{aligned} \quad \square$$

2.7 Straight-line homotopies of discrete Morse functions

Cerf theory is concerned with families of smooth functions on a manifold. A fundamental result is the following statement about homotopies between two *excellent Morse functions*, i.e., Morse functions whose critical points have distinct critical values:

Theorem 2.22 (Cerf [14], Proposition 2.4). *Let M be a compact orientable smooth manifold. Let $(x, t) \mapsto f_t(x) \in C^\infty(M \times [0, 1])$ be a smooth 1-parameter family of smooth functions $f_t \in C^\infty(M)$ such that f_0 and f_1 are excellent Morse functions. Then every open neighborhood $U \subset C^\infty(M \times [0, 1])$ of the family $(x, t) \mapsto f_t(x)$ in the compact-open topology also contains a family $(x, t) \mapsto F_t(x) \in U$ such that $F_0 = f_0$, $F_1 = f_1$, and F_t is an excellent Morse function at all but a finite number of times t .*

In other words, a generic homotopy between f_0 and f_1 yield a Morse function apart from a finite number of times.

We are particularly interested in homotopies of the form $f_t = (1 - t)f + tg$ (called *straight-line homotopies*). Again, the discrete case turns out to be much simpler, and the following statements can be obtained by elementary means:

Theorem 2.23. *Let f and g be two flat pseudo-Morse functions consistent with gradient vector fields V_f and V_g , respectively. Then $f_t = (1 - t)f + tg$ is a flat pseudo-Morse function consistent with the gradient vector field $V = V_f \cap V_g$ for every t with $0 \leq t \leq 1$. If both f and g are additionally Morse, then so is f_t .*

Proof. The function f_t is a flat pseudo-Morse function by Lemma 2.8 and Proposition 2.17. For each pair $(\phi, \rho) \in K \times K$ such that ϕ is a facet of ρ we have $f(\phi) \leq f(\rho)$ and $g(\phi) \leq g(\rho)$. Thus we have $f_t(\phi) \leq f_t(\rho)$. Moreover we have $f_t(\phi) = f_t(\rho)$ if and only if both $f(\phi) = f(\rho)$ and $g(\phi) = g(\rho)$, which in turn holds for every pair $(\phi, \rho) \in V$. Hence f_t is consistent with V . Finally, if both f and g are Morse, then we have $f_t(\phi) = f_t(\rho)$ if and only if $(\phi, \rho) \in V$. Thus f_t is Morse with gradient vector field V . \square

In contrast to Theorem 2.22, the critical cells change only at $t \in \{0, 1\}$. This is due to the requirement that f and g are flat. Note that without this requirement, f_t is no longer guaranteed to be a pseudo-Morse function.

Regarding the ordering of critical values (and other function values) of f_t , we obtain the following result:

Theorem 2.24. *Let f and g be two pseudo-Morse functions and let $f_t = (1 - t)f + tg$. Then there exists a partition $0 = t_0 < t_1 < \dots < t_k = 1$ of the interval $[0, 1]$ such that for any $0 \leq i < k$, the order \prec_{f_t} induced by f_t is unchanged for all $t \in (t_i, t_{i+1})$.*

Proof. For each pair $(\phi, \rho) \in K \times K$ with $f(\phi) - g(\phi) \neq f(\rho) - g(\rho)$, there is exactly one value t such that $f_t(\phi) = f_t(\rho)$, which is given by

$$t_{\phi, \rho} = \frac{f(\phi) - f(\rho)}{f(\phi) - f(\rho) - g(\phi) + g(\rho)}.$$

If $f(\phi) - g(\phi) = f(\rho) - g(\rho)$, we have $f_t(\phi) = f_t(\rho)$ for all t if and only if $f(\phi) = f(\rho)$. Since f_t is continuous in t , the order \prec_{f_t} can change only at the values $t = t_{\phi, \rho}$ for some $(\phi, \rho) \in K \times K$, implying the claim. \square

Combining Theorems 2.23 and 2.24, we find that we can choose a single total order \prec_i on K consistent with (f_t, V) for all $t \in [t_i, t_{i+1}]$.

2.8 Piecewise linear functions and discrete Morse functions

The framework of discrete Morse theory is well-suited both for concise proofs as well as for fast and simple algorithms. Most relevant to our work is the elimination of critical points of discrete gradient vector fields (Theorem 2.1). A

comparably simple construction is not available (and seems difficult to achieve) in the setting of PL Morse theory [44, 26, 3].

In this section we discuss a canonical relationship between discrete and piecewise linear (PL) Morse theory. As it turns out, it is possible to translate statements from one setting to the other seamlessly. This has important consequences: it allows to develop robust and rigorous methods for dealing with real-world data, making no assumptions on the input function.

2.8.1 Simplicial complexes

First we review some definitions concerning simplicial complexes. A *simplex* of dimension d is the convex hull of $(d + 1)$ affinely independent points in \mathbb{R}^n , called *vertices*. A *face* of a simplex s is the convex hull of any nonempty subset of the vertices of s . A *simplicial complex* \mathcal{S} is a set of simplices such that every face of a simplex from \mathcal{S} is also contained in \mathcal{S} , and any nonempty intersection of two simplices from \mathcal{S} is a face of both simplices. In this thesis, all simplicial complexes are assumed to be finite. The union of simplices of \mathcal{S} , endowed with the subspace topology of \mathbb{R}^n , is a topological space $|\mathcal{S}|$ called the *underlying space* of \mathcal{S} . It has the structure of a regular CW complex, the cells being the interiors of the simplices. Therefore, by a slight abuse of notation, we often drop the mention of the underlying space and simply consider a simplicial complex as a CW complex. We sometimes call a simplicial complex *geometric* in order to emphasize the difference from the following purely combinatorial definition.

An *abstract simplicial complex* or *simplicial scheme* is a family of nonempty finite sets which is closed under the operation of taking subsets. Again, we only consider finite complexes. A set s in an abstract simplicial complex is called *simplex*. The elements in s are the *vertices* of s . A subset of a simplex s is called *face* of s . Given a geometric simplicial complex, the vertices of the simplices form an abstract simplicial complex, called the *vertex scheme*. Vice versa, every abstract simplicial complex Δ is the vertex scheme of some geometric simplicial complex. The underlying space of this geometric simplicial complex is unique up to homeomorphisms. We call it the *underlying space* of Δ and denote it by $|\Delta|$.

The following definitions apply to both geometric and abstract simplicial complexes. The *vertices* of a simplicial complex are the union of the vertices of its simplices. A *subcomplex* of Δ is a subset of Δ which is also a simplicial complex.

Given a subset W of the vertices of Δ , the *induced subcomplex* of Δ on W consists of all simplices from Δ which are a subset of W .

2.8.2 Transforming PL functions into pseudo-Morse functions

Assume that \mathcal{K} is a simplicial complex. Let f_{PL} be a simplexwise linear function on $|\mathcal{K}|$ and let f_0 be its restriction to the 0-skeleton of \mathcal{K} . The function f_0 inductively gives rise to a discrete pseudo-Morse function f in the following way. For each 0-cell α , let $f(\alpha) = f_0(\alpha)$. For a cell τ with $\dim \tau > 0$, let $f(\tau)$ be the maximum value of f on any facet of τ . The function f can easily be seen to be pseudo-Morse since it is consistent with the empty vector field $V = \emptyset$ (all cells are critical). Equivalent constructions have been used in [42, 53, 2]. We have the following fact:

Proposition 2.25 (Kühnel [47], Morozov [53]). *Let f_{PL} be a simplexwise linear function on a simplicial complex \mathcal{K} and let f_0 be its restriction to the 0-skeleton of \mathcal{K} . Let*

$$F(t) = \{\phi \in \mathcal{K}_0 : f_0(\phi) \leq t\}.$$

Then the induced subcomplex of \mathcal{K} on $F(t)$ is homotopy equivalent to the sublevel set $\{x \in \mathcal{K} : f_{\text{PL}}(x) \leq t\}$.

Now observe that any level subcomplex $\mathcal{K}(t)$ of f coincides with the induced subcomplex of \mathcal{K} on the corresponding sublevel set $F(t)$ of f_0 . This induced subcomplex, in turn, is homotopy equivalent to the corresponding sublevel set of f_{PL} according to Proposition 2.25. This means that from a Morse-theoretic point of view, the PL function f_{PL} and the pseudo-Morse function f are equivalent. We conclude:

Theorem 2.26. *Let f_{PL} be a simplexwise linear function on a simplicial complex \mathcal{K} . Then there is a canonical pseudo-Morse function f on \mathcal{K} such that for every $t \in \mathbb{R}$ the sublevel set $\{x \in \mathcal{K} : f_{\text{PL}}(x) \leq t\}$ is homotopy equivalent to the level subcomplex $\mathcal{K}(t)$.*

2.8.3 Transforming pseudo-Morse functions into PL functions

Vice versa, we can interpret any discrete pseudo-Morse function f on a regular CW complex \mathcal{K} as a simplexwise linear function $f_{\text{sd}} : |\text{sd } \mathcal{K}| \rightarrow \mathbb{R}$ on the geometric realization of the derived subdivision $\text{sd } \mathcal{K}$. The *derived subdivision* (or *barycentric subdivision*) of a regular CW complex \mathcal{K} is the *order complex* of




Figure 2.3: Illustration to Theorem 2.27. The level subcomplex $\mathcal{K}(t)$ (left) is homotopy equivalent to the sublevel set $\{x \in |\text{sd } \mathcal{K}| : f_{\text{sd}}(x) \leq t\}$ (right), with $t = 5$. The homotopy equivalence is shown using a sequence of simplicial collapses of the barycentric subdivision $\text{sd } \mathcal{K}$.

the face relation, i.e., the abstract simplicial complex $\text{sd } \mathcal{K}$ whose vertices are the cells of \mathcal{K} and whose simplices are the totally ordered subsets of K with regard to the face relation. The geometric realization $|\text{sd } \mathcal{K}|$ is homeomorphic to \mathcal{K} [50]. The function f_{sd} is assumed to linearly interpolate the values of f at the vertices of $|\text{sd } \mathcal{K}|$ inside each simplex of $|\text{sd } \mathcal{K}|$. Again, the sublevel sets of f_{sd} are homotopy equivalent to the corresponding level subcomplexes of f :

Theorem 2.27. *Let f be a pseudo-Morse function on a simplicial complex \mathcal{K} . Then f induces a simplexwise linear function f_{sd} on $|\text{sd } \mathcal{K}|$ such that for every $t \in \mathbb{R}$ the level subcomplex $\mathcal{K}(t)$ is homotopy equivalent to the sublevel set $\{x \in |\text{sd } \mathcal{K}| : f_{\text{sd}}(x) \leq t\}$.*

Proof. Let V be a discrete gradient vector field on \mathcal{K} that is consistent with f and let $<$ be a total order consistent with (f, V) . Let $K(t)$ and $K(\rho)$ denote the cells of the level and order subcomplexes $\mathcal{K}(t)$ and $\mathcal{K}(\rho)$, respectively. Let $\Delta(U)$ denote the induced subcomplex of $\text{sd } \mathcal{K}$ on a vertex set $U \subset K$ (recall that the cells K are the vertices of $\text{sd } \mathcal{K}$). The induced subcomplex $\Delta(K(t))$ is easily seen to be identical to $\text{sd } \mathcal{K}(t)$. Let

$$F(t) = \{\phi \in K : f(\phi) \leq t\} \subset K(t).$$

We now show that $\Delta(K(t))$ collapses simplicially onto $\Delta(F(t))$. See Fig. 2.3 for an example.

Let $\sigma \in K(t) \setminus F(t)$ and let σ_- denote its predecessor with respect to $<$. We write $\Delta(\rho)$ for $\Delta(\{\phi \in K : \phi < \rho\})$. We first show that $\Delta(\sigma)$ collapses onto $\Delta(\sigma_-)$. It follows from the definition of a level subcomplex that a cell σ with $f(\sigma) > t$ can only be contained in $K(t)$ if it is a *free face* of some cell τ with $f(\tau) \leq t$, i.e., σ has no other cofaces than τ . Consequently, we have $\tau < \sigma$. For every simplex $S \in \Delta(\sigma)$ with $\sigma \in S$ and $\tau \notin S$, the simplex $T = S \cup \{\tau\}$ is also contained in $\Delta(\sigma)$, and S is a

free face of T in $\Delta(\sigma)$. Hence, the set of all such pairs (S, T) constitutes a discrete gradient vector field W on $\Delta(\sigma)$ such that exactly the simplices containing σ (the *vertex star* of σ) are non-critical. This vector field W provides a simplicial collapse of $\Delta(\sigma)$ onto $\Delta(\sigma_-)$ by applying Theorem 2.14 with an arbitrary linear extension of \prec_W .

By repeatedly applying this argument, we find that $\Delta(K(t))$ collapses onto $\Delta(F(t))$. This implies that the geometric realizations $|\Delta(K(t))|$ and $|\Delta(F(t))|$ are homotopy equivalent.

Finally, let f_{sd} be the simplexwise linear extension of f from the vertices of $\text{sd } \mathcal{K}$ to the whole complex. Recall that $|\Delta(F(t))|$ is homotopy equivalent to the sublevel set $\{x \in |\text{sd } \mathcal{K}| : f_{\text{sd}}(x) \leq t\}$ according to Proposition 2.25. The claim now follows. \square

This equivalence allows us to translate back and forth between piecewise linear functions and pseudo-Morse functions, and to use theorems of piecewise linear Morse theory in the context of discrete Morse theory.

In a similar fashion, a discrete pseudo-Morse function can be constructed from a function defined only on the set of 2-cells of a combinatorial surface by defining $f(\sigma)$ as the minimum value of all cells that contain σ as a facet. This can be used to construct discrete pseudo-Morse functions from functions defined on cubical grids, such as pixel images, by interpreting each pixel as a 2-cell. The resulting level subcomplexes correspond to the cubical complexes extracted from image data as described by Kaczynski et al. [40]. Vice versa, a pseudo-Morse function on a cubical complex can be interpreted as a function defined on a subdivided grid. This construction has been used in the examples in Chapter 6.

Note that starting with a PL function and constructing a pseudo-Morse function consistent with the empty vector field implies that initially all cells are considered critical. King et al. [42] propose to construct an initial discrete gradient vector field with critical cells corresponding to the critical vertices (in the PL sense, see [44, 26, 3]) of a (non-degenerate) input PL function instead.

We do not consider the problem of constructing such an initial discrete gradient vector field separately, since it is a special case of the topological simplification problem discussed in Theorem 4.2 with $\delta = 0$. In this case, the problem reduces to minimizing the number of critical points among all gradient vector fields consistent with the input function. We discuss the simplification of a gradient

vector field in Sections 4.2 and 5.3. As it turns out, a solution to the topological simplification problem can be found independently of the gradient vector field consistent with the input function. Therefore we do not require a simplified initial gradient vector field but use the empty vector field instead.

2.9 Morse theory for general CW complexes

Forman [28] briefly discusses how to extend the results of discrete Morse theory to general CW complexes. On a general CW complex, a discrete vector field consist of pairs (σ, τ) such that σ is a *regular* facet of τ . Recall that for a regular facet σ , the characteristic map Φ_τ restricts to a homeomorphism on the preimage $\Phi_\tau^{-1}(\sigma)$ and the closure of $\Phi_\tau^{-1}(\sigma)$ is a closed ball.

As an example, consider a simple CW decomposition of the closed unit d -ball into 3 cells ρ, σ, τ of dimension 0, $(d - 1)$, and d , respectively. By the original definitions of discrete Morse theory, we are not able to express the contraction of the ball to a point using a discrete vector field: the closure of $\Phi_\tau^{-1}(\sigma)$ is a $(d - 1)$ -sphere. Nevertheless, it seems natural to allow the pair (σ, τ) in a gradient vector field.

We extend the theory by dropping the requirement on the closure of $\Phi_\tau^{-1}(\sigma)$ to be a ball:

Definition (semi-regular facet). *Let σ be a facet of τ in a CW complex. If the characteristic map Φ_τ of τ restricts to a homeomorphism on the preimage $\Phi_\tau^{-1}(\sigma)$, then σ is a semi-regular facet of τ .*

This generalization gives rise to a new definition for discrete vector fields on CW complexes, generalizing the one considered by Forman [28].

Definition (discrete vector field). *A discrete vector field V on a CW complex \mathcal{K} is a set of pairs of cells $(\sigma, \tau) \in K \times K$, with σ a semi-regular facet of τ , such that each cell of K is contained in at most one pair of V .*

2.9.1 A generalization of elementary collapses

As it turns out, we can carry over the central results of discrete Morse theory to this generalized setting. The key construction is an extension of the notion of an elementary collapse for semi-regular facets. We make use of the *generalized Schoenflies theorem*:

Theorem 2.28 (Brown [9], Morse [56]). *Let h be a homeomorphic embedding of $S^{d-1} \times [0, 1]$ into S^d . Then the closure of either complementary domain of $h(S^{d-1} \times \frac{1}{2})$ is a closed d -ball.*

This is a higher dimensional version of the Jordan-Schoenflies theorem [11], which asserts that any Jordan curve divides the plane into two regions homeomorphic to the interior and exterior of the unit circle, respectively. The requirement that the embedding $x \mapsto h(x, \frac{1}{2})$ of the sphere S^{d-1} can be extended to an embedding of the thickened sphere $S^{d-1} \times [0, 1]$ is crucial; a counterexample where this requirement is not met is given by the Alexander horned sphere [35].

Moreover, we require the following lemma, which we use for transporting a strong deformation retraction through a characteristic map.

Lemma 2.29. *Let X be a topological space, let A be a strong deformation retract of X , and let $U = X \setminus A$. Let $q : X \rightarrow Y$ be a continuous map such that q^{-1} restricts to a homeomorphism on $V = q(U)$ and $\text{bd}_Y V \subset q(A)$, where $\text{bd}_Y V$ denotes the boundary of V in Y . Then $Y \setminus V$ is a strong deformation retract of Y .*

Proof. Let F_t be a strong deformation retraction of X onto A . We show that

$$\widetilde{F}_t : x \mapsto \begin{cases} q \circ F_t \circ q^{-1}(x) & x \in V \\ x & x \notin V \end{cases}$$

is a strong deformation retraction of Y onto $Y \setminus V$.

Clearly the map $(y, t) \mapsto \widetilde{F}_t(y)$ is continuous both on $V \times I$ and on $(Y \setminus V) \times I$. The boundary of both $V \times I$ and $(Y \setminus V) \times I$ is $\text{bd}_Y V \times I$. Since q^{-1} restricts to a homeomorphism on $V = q(U)$, we know that $q(A) \cap q(U) = \emptyset$ and thus $q^{-1} \circ q(A) = A$. Moreover, since the map F_t restricts to the identity on A , we have $q \circ F_t \circ q^{-1}(\{y\}) = \{y\}$ for $y \in q(A)$. In particular, $q \circ F_t \circ q^{-1}$ restricts to the identity on $\text{bd}_Y V \subset q(A)$. We conclude that \widetilde{F}_t is continuous on both $\text{cl}_Y V$ and $\text{cl}_Y(Y \setminus V)$ and thus $(y, t) \mapsto \widetilde{F}_t(y)$ is continuous on all of $Y \times I$. Finally, we directly have $\widetilde{F}_0(y) = y$, $\widetilde{F}_1(Y) = Y \setminus V$, and $\widetilde{F}_t(y) = y$ for $y \in Y \setminus V$, so \widetilde{F}_t is a strong deformation retraction of Y onto $Y \setminus V$. \square

We can now combine Theorem 2.28 and Lemma 2.29 to construct a generalized elementary collapse for semi-regular facets.

Theorem 2.30. *Let \mathcal{K} be a CW complex and let σ and τ be open cells of \mathcal{K} such that σ is a semi-regular facet of τ and no other cells are attached to either σ or τ . Let $\widetilde{\mathcal{K}} = \mathcal{K} \setminus (\sigma \cup \tau)$. Then $\widetilde{\mathcal{K}}$ is a strong deformation retract of \mathcal{K} .*

Proof. Let $d = \dim \sigma$, let Φ_σ be the characteristic map of σ , and let $B = \Phi_\sigma(B_r^d)$, where $B_r^d = \{x \in \mathbb{R}^d : \|x\| < r\}$ for some $0 < r < 1$. Let $A_{a,b}^d = \{x \in \mathbb{R}^d : a \leq \|x\| \leq b\}$ denote the d -dimensional annulus with radii a and b . Choose ϵ such that $0 < r - \epsilon < r + \epsilon < 1$. The cylinder $S^{d-1} \times [0, 1]$ is homeomorphic to the annulus $A_{r-\epsilon, r+\epsilon}^d$ by the map $\alpha : (x, \lambda) \mapsto (r + (2\lambda - 1)\epsilon)x$. Moreover, it is homeomorphically embedded into S^d by $h = \Phi_\tau^{-1} \circ \Phi_\sigma \circ \alpha$, since by assumption both the restrictions of Φ_σ^{-1} and Φ_τ^{-1} to σ are homeomorphisms. We have $h(x, \frac{1}{2}) = \Phi_\tau^{-1}(\partial B)$. Hence, by Theorem 2.28, both $\Phi_\tau^{-1}(\bar{B})$ and $S^d \setminus \Phi_\tau^{-1}(B)$ are closed d -balls. This implies that $S^d \setminus \Phi_\tau^{-1}(B)$ is a strong deformation retract of the closed $(d + 1)$ -ball B^{d+1} .

We have

$$\Phi_\tau(S^d \setminus \Phi_\tau^{-1}(B)) = \partial\tau \setminus B = \bar{\tau} \setminus (B \cup \tau).$$

Moreover, Φ_τ^{-1} restricts to a homeomorphism on $B \cup \tau$. Since no other cells are attached to either σ and τ , both B and τ are open in \mathcal{K} , and we have

$$\text{bd}_{\mathcal{K}}(B \cup \tau) = \bar{\tau} \setminus \text{int}_{\mathcal{K}}(B \cup \tau) = \bar{\tau} \setminus (B \cup \tau).$$

Hence, by Lemma 2.29 with $X = \overline{B^{d+1}}$, $A = S^d \setminus \Phi_\tau^{-1}(B)$, $Y = \mathcal{K}$, $q = \Phi_\tau$, and $V = B \cup \tau$, we find that $\mathcal{K} \setminus (B \cup \tau)$ is a strong deformation retract of \mathcal{K} .

The sphere S^{d-1} is a strong deformation retract of the annulus $A_{r,1}^d$. Note that $\Phi_\sigma(S^{d-1}) = \partial\sigma$. Moreover, Φ_σ^{-1} restricts to a homeomorphism on $\sigma \setminus B$. Since τ is the only cell attached to σ , the set $\sigma \setminus B$ is open in $\mathcal{K} \setminus (B \cup \tau)$ and we have

$$\text{bd}_{\mathcal{K} \setminus (B \cup \tau)}(\sigma \setminus B) = (\bar{\sigma} \setminus B) \setminus \text{int}_{\mathcal{K} \setminus (B \cup \tau)}(\sigma \setminus B) = (\bar{\sigma} \setminus B) \setminus (\sigma \setminus B) = \partial\sigma.$$

Thus, again by Lemma 2.29 with $X = A_{r,1}^d$, $A = S^{d-1}$, $Y = \mathcal{K} \setminus (B \cup \tau)$, $q = \Phi_\sigma$, and $V = \sigma \setminus B$, we find that $\mathcal{K} \setminus (\sigma \cup \tau)$ is a strong deformation retract of $\mathcal{K} \setminus (B \cup \tau)$. Together we have that $\mathcal{K} \setminus (\sigma \cup \tau)$ is a strong deformation retract of \mathcal{K} . \square

2.9.2 The main results of Morse theory for general CW complexes

Theorem 2.30 enables us to generalize the main results of discrete Morse theory to discrete vector fields with semi-regular facet pairs. The proofs are nearly identical, using the deformation retraction of Theorem 2.30 wherever elementary collapses are originally used.

First we recall the following basic fact about attaching cells:

Lemma 2.31 (Milnor [51], Lemma 3.7). *Let X be a topological space and let $\varphi : S^{d-1} \rightarrow X$ be an attaching map. Any homotopy equivalence $f : X \rightarrow Y$ extends to a homotopy equivalence*

$$F : X \underset{\varphi}{\cup} \overline{B^d} \rightarrow Y \underset{f \circ \varphi}{\cup} \overline{B^d}$$

defined by

$$x \mapsto \begin{cases} f(x) & x \in X \\ x & x \in \overline{B^d}. \end{cases}$$

Theorem 2.32 (cf. Theorem 2.14). *Let V be a gradient vector field on a CW complex \mathcal{K} and let $<$ be a linear extension of $<_V$. If ρ and ψ are two cells such that $\rho < \psi$ and there is no critical cell ϕ with respect to V such that $\rho < \phi \leq \psi$, then $\mathcal{K}(\rho)$ is a strong deformation retract of $\mathcal{K}(\psi)$.*

Proof. We show that if ϕ is a non-critical cell then $\mathcal{K}(\phi_-)$ is a strong deformation retract of $\mathcal{K}(\phi)$, where ϕ_- is the predecessor of ϕ in $<$. The statement then follows by induction.

First assume that $(\sigma, \phi) \in V$ for some σ . Then $\mathcal{K}(\phi) = \mathcal{K}(\phi_-)$. Now assume that $(\phi, \tau) \in V$ for some τ . Then $\mathcal{K}(\phi) = \mathcal{K}(\phi_-) \cup (\phi \cup \tau)$. By Theorem 2.30 $\mathcal{K}(\phi_-)$ is a strong deformation retract of $\mathcal{K}(\phi)$. \square

Theorem 2.33 (cf. Theorem 2.15). *Let V be a gradient vector field on a CW complex \mathcal{K} and let $<$ be a linear extension of $<_V$. Suppose that ρ and ψ are two cells such that $\rho < \psi$ and there is exactly one critical cell ϕ of index d with respect to V such that*

$$\rho < \phi \leq \psi.$$

Then $\mathcal{K}(\psi)$ is homotopy equivalent to

$$\mathcal{K}(\rho) \underset{\Phi}{\cup} \overline{B^d},$$

with $\Phi : S^{d-1} \rightarrow \mathcal{K}(\rho)$.

Proof. From Theorem 2.32, we have that $\mathcal{K}(\rho)$ is a strong deformation retract of $\mathcal{K}(\phi_-)$, where ϕ_- is the predecessor of ϕ in $<$. By definition, $\mathcal{K}(\phi)$ is obtained by attaching a d -cell to $\mathcal{K}(\phi_-)$. By Lemma 2.31, $\mathcal{K}(\phi)$ is homotopy equivalent to $\mathcal{K}(\rho)$ with a d -cell attached. Again by Theorem 2.32, $\mathcal{K}(\phi)$ is a strong deformation retract of $\mathcal{K}(\psi)$. We conclude that $\mathcal{K}(\psi)$ is homotopy equivalent to $\mathcal{K}(\rho)$ with a d -cell attached. \square

Corollary 2.34 (cf. Theorem 2.5). *Let V be a discrete gradient vector field on a CW complex \mathcal{K} . Then \mathcal{K} is homotopy equivalent to a CW complex \mathcal{M}_V with as many d -cells as there are critical cells of V of dimension d .*

Theorem 2.35 (cf. Theorem 2.1). *Let ϕ and ρ be two critical cells of a gradient vector field V with exactly one V -path $\Gamma = (\sigma_0, \tau_0, \sigma_1, \dots, \tau_{r-1}, \sigma_r)$ from $\partial\rho$ to ϕ . Assume that for every $0 \leq i \leq r$, σ_i is a semi-regular facet of τ_{i-1} , where $\tau_{-1} = \rho$. Then there is a gradient vector field \tilde{V} obtained by reversing V along the path Γ . The critical cells of \tilde{V} are exactly the critical cells of V apart from $\{\phi, \rho\}$. Moreover, $V = \tilde{V}$ except along the path Γ .*

2.9.3 Partial collapses

Our generalized definition of gradient vector fields can be used to remove an unfortunate restriction in the original formulation of discrete Morse theory: In general, the original theory does not allow for partitioning a gradient vector field V into disjoint subsets U, W and performing the induced partial collapses $\mathcal{K} \searrow \mathcal{M}_U \searrow \mathcal{M}_V$ sequentially.

In order to overcome this restriction, we need the following lemma. Let $<$ be a linear extension of $<_V$. Consider the induced filtration by order subcomplexes. By repeatedly applying Lemma 2.31 and Theorem 2.30, we obtain an explicit expression for the homotopy equivalence of Corollary 2.34:

Corollary 2.36. *Let V be a discrete gradient vector field on a CW complex \mathcal{K} . Then \mathcal{K} is homotopy equivalent to \mathcal{M}_V via*

$$h_V : x \mapsto \begin{cases} \Phi_{\tilde{\rho}} \circ \Phi_{\rho}^{-1}(x) & x \in \rho, \rho \in \text{Cr}(V) \\ h_V \circ c_{\sigma, \tau}(x) & x \in \sigma \cup \tau, (\sigma, \tau) \in V, \end{cases}$$

where $c_{\sigma, \tau} : \mathcal{K}(\tau) \rightarrow \mathcal{K}(\tau_-) = \mathcal{K}(\tau) \setminus (\sigma \cup \tau)$ denotes the homotopy equivalence of the generalized elementary collapse associated with (σ, τ) , and $\tilde{\rho}$ is the cell of \mathcal{M}_V corresponding to a critical cell ρ of V . In particular, the homotopy equivalence h_V restricts to a homeomorphism on any critical cell ρ of V .

Note that the definition of the map h_V is recursive. The recursion terminates since a point $x \in \sigma \cup \tau$ for $(\sigma, \tau) \in V$ is sent to a point in $\mathcal{K}(\tau_-)$ by $c_{\sigma, \tau}$, i.e., to some cell ϕ with $\phi < \tau$, and there are only a finite number of cells.

As a consequence, we obtain the following theorem:

Theorem 2.37. *Let V be a gradient vector field on a CW complex \mathcal{K} and $W \subset V$. Then $V \setminus W$ induces a gradient vector field on \mathcal{M}_W .*

Proof. Let $(\phi, \rho) \in V \setminus W$. Since ϕ is a semi-regular facet of ρ , the characteristic map Φ_ρ of ρ restricts to a homeomorphism on the preimage $\Phi_\rho^{-1}(\phi)$. Moreover, ϕ is a critical cell of W , and so by Corollary 2.36 h_V restricts to a homeomorphism $\phi \rightarrow \tilde{\phi}$. Thus $\Phi_{\tilde{\rho}} = h_V \circ \Phi_\rho$ restricts to a homeomorphism on $\Phi_\rho^{-1}(\phi) = \Phi_{\tilde{\rho}}^{-1}(\tilde{\phi})$, and so $\tilde{\phi}$ is a semi-regular facet of $\tilde{\rho}$. \square

Let $V = U \sqcup W$ and let \tilde{W} be the vector field induced by W on \mathcal{M}_U . Then we have $h_V = h_{\tilde{W}} \circ h_U$, up to a homeomorphism of the respective Morse complexes. In other words, the collapse $\mathcal{K} \searrow \mathcal{M}_V$ induced by V can be split into a sequence of partial collapses $\mathcal{K} \searrow \mathcal{M}_U \searrow \mathcal{M}_V$.

Chapter 3

Persistent homology of discrete Morse functions

Persistent homology aims at investigating the change of the homology groups in a filtration of a topological space (a nested sequence of subspaces). In particular, it determines how long changes in homology last before they are reverted again.

The theory was introduced in the seminal paper of Edelsbrunner et al. [23], which also provides the nomenclature used in this thesis. Some of the ideas already appeared in two largely overlooked articles of Morse [54, 55], with the goal of extending the results of Morse theory to arbitrary metric spaces and functions that are not as well-behaved as Morse functions. The basic concept has also been discovered independently several times in the context of topological data analysis [31, 59] and as a way of describing the size of mountains [32]. The theory has been significantly extended and continues to attract much interest [67, 18, 13].

The concepts of persistent homology can naturally be applied to nested sequences of order subcomplexes or level subcomplexes of discrete pseudo-Morse functions. The following definitions are concerned with cellular homology with coefficients in an arbitrary field F , unless explicitly stated otherwise. We write $H_d(\mathcal{K})$ as a shorthand for the d^{th} homology group $H_d(\mathcal{K}; F)$ of \mathcal{K} and $H_*(\mathcal{K}) = \bigoplus_d H_d(\mathcal{K})$. For definitions see Section 2.3.1 or [50, 35].

Convention and notation Throughout Chapter 3 we consider a pseudo-Morse function f consistent with a gradient vector field V on a regular CW complex \mathcal{K} and a strict total order $<$ consistent with (f, V) .

3.1 Birth, death, and persistence pairs

As a consequence of Theorem 2.14, the homology groups of order subcomplexes change only at critical cells of V . Let σ and τ be critical cells such that $\sigma < \tau$ and consider the inclusion map

$$i^{\sigma, \tau} : \mathcal{K}(\sigma) \hookrightarrow \mathcal{K}(\tau)$$

between the order subcomplexes with regard to the total order $<$. This map induces a homomorphism

$$i_*^{\sigma, \tau} : H_*(\mathcal{K}(\sigma)) \rightarrow H_*(\mathcal{K}(\tau))$$

between homology groups. The subgroup $\text{im}(i_*^{\sigma, \tau}) \subset H_*(\mathcal{K}(\tau))$ is called a *persistent homology group*. For a particular dimension d , the rank of the persistent homology group $\text{rank} \text{im}(i_*^{\sigma, \tau})$ is called the d^{th} *persistent Betti number* of $\mathcal{K}(\sigma)$ in $\mathcal{K}(\tau)$ (or, in the language of Morse [55], the d^{th} *connectivity* of $\mathcal{K}(\sigma)$ on $\mathcal{K}(\tau)$).

Now consider the critical cells $\{\rho_1, \dots, \rho_N\}$ of V , ordered such that $\rho_i < \rho_{i+1}$, and the corresponding filtration of order subcomplexes $\mathcal{K}(\rho_i) \subset \mathcal{K}(\rho_{i+1})$. We compare the corresponding homology groups using the homomorphisms induced by inclusion. The set of critical cells decomposes into two disjoint subsets: *positive* and *negative* cells. A *positive* cell creates homology classes and a *negative* cell merges pairs of homology classes. More precisely, for every cell ρ let ρ_- denote its predecessor with regard to $<$. A critical d -cell ρ is called *positive* if it increases the d^{th} Betti number $\beta_d = \text{rank } H_d$, i.e., $\beta_d(\mathcal{K}(\rho)) = \beta_d(\mathcal{K}(\rho_-)) + 1$, and *negative* if it decreases the $(d - 1)^{\text{th}}$ Betti number i.e., $\beta_{d-1}(\mathcal{K}(\rho)) = \beta_{d-1}(\mathcal{K}(\rho_-)) - 1$. It follows from the Morse inequalities (Corollary 2.16) that every critical cell is indeed either positive or negative.

We further require the notions of *birth*, *death*, and *persistence pairs* for Morse functions, following the definitions of [23]. Intuitively, a persistence pair (σ, τ) is a pair of a positive cell σ and a negative cell τ that represents the life cycle for a class born at σ and merged into an older class at τ . To make this precise, consider the sequence

$$H_*(\mathcal{K}(\sigma_-)) \rightarrow H_*(\mathcal{K}(\sigma)) \rightarrow H_*(\mathcal{K}(\tau_-)) \rightarrow H_*(\mathcal{K}(\tau))$$

of homomorphisms $i_*^{\phi, \rho} : H_*(\mathcal{K}(\phi)) \rightarrow H_*(\mathcal{K}(\rho))$ induced by the inclusion maps $i^{\phi, \rho} : \mathcal{K}(\phi) \rightarrow \mathcal{K}(\rho)$ for each $\phi < \rho$. Here we allow for the cases $\sigma = \tau_-$ and $\mathcal{K}(\sigma_-) = \emptyset$ (if σ is the first cell in $<$).

Definition (birth, death, persistence pair [23]). We say that a class $h \in H_*(\mathcal{K}(\sigma))$ is born at (or created by) a positive cell σ if

$$h \notin \text{im}(i_*^{\sigma-, \sigma}).$$

Moreover, we say that a class $h \in H_*(\mathcal{K}(\sigma))$ that is born at σ dies entering (or gets merged into \tilde{h} by) a negative cell τ if there is a class $\tilde{h} \in H_*(\mathcal{K}(\sigma_-))$ such that

$$i_*^{\sigma, \tau-}(h) \notin \text{im}(i_*^{\sigma-, \tau-}) \quad \text{but} \quad i_*^{\sigma, \tau}(h) = i_*^{\sigma-, \tau}(\tilde{h}) \in \text{im}(i_*^{\sigma-, \tau}).$$

If there exists a class h that is born at σ and dies entering τ , then (σ, τ) is a persistence pair. The difference $f(\tau) - f(\sigma)$ is called the persistence of (σ, τ) . If a positive cell σ is not paired with any negative cell τ , then it is called essential.

The merging of h and \tilde{h} is not symmetric, because the *younger* class h (with respect to $<$) gets merged into the *older* class \tilde{h} . The classes born at essential cells generate the homology of \mathcal{K} . Note that in this definition we always have $\dim \tau = \dim \sigma + 1$. Thus the number of critical d -cells of \mathcal{K} equals the Betti number β_d (a topological invariant of \mathcal{K}) plus the number of persistence pairs of dimensions $(d - 1, d)$ and $(d, d + 1)$. On combinatorial surfaces, the only possible cases for $(\dim \sigma, \dim \tau)$ are $(0, 1)$ and $(1, 2)$.

In the language of Morse [55], a positive cell (understood as a cycle relative to $\mathcal{K}(\sigma_-)$) is called a *linkable cap*, while a negative cell is called a *non-linkable cap*, and the persistence is called the *cap span*.

3.1.1 Choice of coefficients

If \mathcal{K} is an orientable combinatorial surface, then $H_d(\mathcal{K}; \mathbb{Z})$ is free abelian and by the universal coefficient theorem for homology [35, Theorem 3A.3] we have $H_d(\mathcal{K}; F) = H_d(\mathcal{K}; \mathbb{Z}) \otimes F = F^{\beta_d(\mathcal{K}; \mathbb{Z})}$, where \otimes denotes the tensor product of abelian groups. Hence, in this case the definitions are independent of the particular choice of F . The same holds for connected non-orientable combinatorial surfaces with nonempty boundary. Even if \mathcal{K} is a closed non-orientable combinatorial surface with boundary, the definitions coincide in almost all cases. The only difference appears in the pairing of the last 2-cell τ in each non-orientable connected component: τ is a positive cell if and only if $F = \mathbb{Z}_2$.

3.1.2 Persistence for filtrations

We point out that the previous definitions are straightforward adaptations of the definitions for filtrations of simplicial complexes and general chain complexes [23, 67]. In particular, a simplexwise *filtration* of a finite simplicial complex Δ (a nested sequence of subcomplexes $\emptyset = \Delta_0 \subset \Delta_1 \subset \Delta_2 \subset \dots \subset \Delta_n = \Delta$ with $\Delta_i \setminus \Delta_{i-1} = \{\sigma_i\}$, where σ_i is a simplex) directly gives rise to a discrete Morse function $\sigma_i \mapsto i$, whose level subcomplexes are the subcomplexes of the filtration. This function has an empty gradient vector field, i.e., all cells are critical. More generally, an arbitrary (not necessarily simplex-wise) filtration with $\Delta_i \setminus \Delta_{i-1} = S_i$ gives rise to a discrete pseudo-Morse function $\sigma_i \mapsto i \Leftrightarrow \sigma_i \in S_i$.

The definition of a persistence pair by Cohen-Steiner et al. [18] does not assume a cell-by-cell filtration. As a consequence, persistence pairs can only be identified as pairs of critical values, not of critical cells. For our purposes, it is however crucial to identify the birth and death of homology classes with particular cells, not only with certain subcomplexes in the filtration. The original definition [23] considered a simplex-wise filtration as described above.

3.1.3 Gradient vector fields as zero-persistence pairs

Note that depending on V there might be persistence pairs with persistence 0. In particular, if f is a flat pseudo-Morse function, then the pairs in V can actually be interpreted as persistence pairs with persistence 0. To do so, we chose the total order $<$ to be consistent not with V , but instead with the empty vector field \emptyset ; this is possible since f is by definition also consistent with \emptyset . We can now choose $<$ such that the pairs in V are consecutive. Then these pairs are clearly persistence pairs, called *local persistence pairs* in [2], and have zero persistence with regard to f by the definition of a flat Morse function.

3.2 Duality and persistence

For any closed combinatorial d -manifold \mathcal{K} , there is an associated *dual complex* \mathcal{K}^* , a combinatorial manifold homeomorphic to \mathcal{K} whose i -cells correspond to $(d-i)$ -cells of \mathcal{K} [35, 10]. A discrete pseudo-Morse function f on \mathcal{K} gives rise to a

discrete pseudo-Morse function f^* on K^* via $\sigma^* \mapsto -f(\sigma)$ [28]. If V is consistent with f , then $V^* = \{(\tau^*, \sigma^*) : (\sigma, \tau) \in V\}$ is consistent with f^* . As a consequence,

$$\Gamma = (\sigma_0, \tau_0, \sigma_1, \dots, \tau_{r-1}, \sigma_r)$$

is a V -path from $\partial\tau$ to $\sigma = \sigma_r$ if and only if

$$\Gamma^* = (\tau_{r-1}^*, \sigma_{r-1}^*, \tau_{r-2}^*, \dots, \sigma_0^*, \tau^*)$$

is a V^* -path from $\partial\sigma^*$ to τ^* .

Moreover, the persistence pairs of dimension $(d-1, d)$ for \mathcal{K} correspond to the persistence pairs of dimension $(0, 1)$ for the dual complex \mathcal{K}^* (with $\tau^* < \sigma^* \Leftrightarrow \sigma < \tau$) when using \mathbb{Z}_2 coefficients or if \mathcal{K} is orientable [19, 2, 22].

The homology groups $H_0(\mathcal{K}(\rho_i))$ (generated by the connected components of $\mathcal{K}(\rho_i)$), and hence the persistence pairs of dimension $(0, 1)$, are determined solely by the 1-skeleton of \mathcal{K} , also called the (*primal*) *graph* of \mathcal{K} . Consequently, the persistence pairs of dimension $(d-1, d)$ are determined by the 1-skeleton of \mathcal{K}^* , called the *dual graph*. In particular, this means that all persistence pairs on a *surface* can be determined in terms of Morse functions on graphs. Note that both the primal and the dual graph can contain multiple edges.

3.2.1 Surfaces with boundary

In order to treat surfaces with boundary, we employ the canonical construction of attaching an additional 2-cell (with function value ∞) to each boundary component. In this way we obtain a closed surface having the same sequence of order subcomplexes (up to the additional cells) and hence the same persistence pairs as the original surface. The corresponding dual graph contains an additional vertex (with value $-\infty$) for each boundary component. Moreover, if σ is a boundary 1-cell and τ the single 2-cell attached to σ , then the edge σ^* of the dual graph connects the additional vertex for the boundary component of σ to the vertex τ^* .

3.3 Stability of persistence diagrams

Cohen-Steiner et al. [18] studied properties of *persistence diagrams*, which are a representation of the value pairs $(f(\sigma), f(\tau))$ corresponding to the persistence

pairs (σ, τ) of a function f . We give the according definition for pseudo-Morse functions. Here we use $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$.

Definition (Persistence diagram). *The persistence diagram $D(f)$ of a pseudo-Morse function f is the multiset on $\subset \overline{\mathbb{R}}^2$ containing one instance of $(f(\sigma), f(\tau))$ for each persistence pair (σ, τ) of f , one instance of $(f(\sigma), \infty)$ for each essential cell σ , and each point on the diagonal with countably infinite multiplicity.*

A multiset on a set S is a tuple (S, m) , where the multiplicity m is a function from S to some set of cardinal numbers (see [8] for a survey on multiset theory).

Although the definition of persistence pairs depends on the choice of a total order, the persistence diagram is invariant of this choice and is therefore a property of the pseudo-Morse function alone:

Theorem 3.1. *The persistence diagram of a discrete pseudo-Morse function f is well-defined; in particular, it is independent of the gradient vector field V consistent with f and of the total order $<$ consistent with (f, V) .*

Proof. Assume that there are k positive d -cells σ_i with non-zero persistence and the same function value $f(\sigma_i) = s$. Let $s_- = \max\{f(\phi) : \phi \in K, f(\phi) < s\}$. According to the definition of a positive cell, we have

$$\beta_d(\mathcal{K}(s)) = \beta_d(\mathcal{K}(s_-)) + k.$$

The Betti numbers are however independent of $<$, and thus there must be k positive d -cells σ_i with non-zero persistence and $f(\sigma_i) = s$ for any total order consistent with f .

Likewise, assume that there are k persistence pairs (σ_i, τ_i) of dimension $(d, d+1)$ with $(f(\sigma_i), f(\tau_i)) = (s, t)$. Let $\beta_d^{s,t} = \text{rank im}(i_d^{s,t})$ denote the d^{th} persistent Betti number of $\mathcal{K}(s)$ in $\mathcal{K}(t)$. The subgroup of $H_d(\mathcal{K}(t))$ consisting of the classes born at s has rank $(\beta_d^{s,t} - \beta_d^{s_-,t})$. By the definition of a persistence pair, the rank of the subgroup of classes born at s decreases at t by k from the next smaller value $t_- = \max\{f(\phi) : \phi \in K, f(\phi) < t\}$. Hence we have

$$(\beta_d^{s,t} - \beta_d^{s_-,t}) = (\beta_d^{s,t_-} - \beta_d^{s_-,t_-}) - k.$$

Again, the persistent Betti numbers are independent of the total order $<$. Therefore the persistence diagram depends only on f . \square

This result implies that properties of a pseudo-Morse function f that depend only on the persistence diagram $D(f)$ are independent of the choice of the symbolic perturbation data $(V, <)$.

3.3.1 A metric on persistence diagrams

We now introduce a metric on the space of persistence diagrams considered by Cohen-Steiner et al. [18]. Loosely speaking, the distance of two diagrams is expressed by the maximal distance of two corresponding points.

By the *individual elements* of a multiset $X = (S, m)$ we mean the disjoint union E of $m(x)$ instances of each $x \in S$, where $m(x)$ is the multiplicity of x . By a slight abuse of notation we write $x \in X$ for $x \in E$.

Definition (Bottleneck distance [18]). *Let X and Y be two multisets of $\overline{\mathbb{R}}^2$. The bottleneck distance is $d_B(X, Y) := \inf_{\gamma} \sup_{x \in X} \|x - \gamma(x)\|_{\infty}$, where γ ranges over all bijections from the individual elements of X to the individual elements of Y .*

Here we assume $(a, \infty) - (b, \infty) = (a - b, 0)$, $(a, \infty) - (b, c) = (a - b, \infty)$, and $\|(a, \infty)\|_{\infty} = \infty$ for $a, b, c \in \mathbb{R}$.

For a given complex, the bottleneck distance actually provides a metric on the set of persistence diagrams:

Proposition 3.2. *The bottleneck distance is a metric on the persistence diagrams of discrete pseudo-Morse functions on a regular CW complex \mathcal{K} . Thus, it induces a pseudometric on the discrete pseudo-Morse functions on \mathcal{K} .*

Proof. We first show that the distance $d_B(D(f), D(g))$ for two pseudo-Morse functions on \mathcal{K} is always finite. To see this, note that the number of essential d -cells equals the Betti number $\beta_d(\mathcal{K})$. This implies that there is always a bijection γ that maps essential to essential cells, and persistence pairs to persistence pairs or diagonal points and vice versa. As a consequence, the distance $\|x - \gamma(x)\|_{\infty}$ is finite for all $x \in D(f)$, and so is the bottleneck distance of the diagrams.

Non-negativity $d_B(D(f), D(g)) \geq 0$, identity of indiscernibles $d_B(D(f), D(g)) = 0 \Leftrightarrow D(f) = D(g)$ and symmetry $d_B(D(f), D(g)) = d_B(D(g), D(f))$ are obvious.

Finally, the triangle inequality $d_B(D(f), D(h)) \leq d_B(D(f), D(g)) + d_B(D(g), D(h))$ holds by the following simple argument. Let γ be a bijection from $D(f)$ to $D(g)$ and let η be a bijection from $D(g)$ to $D(h)$. Then

$$\begin{aligned} \sup_{x \in D(f)} \|x - \eta \circ \gamma(x)\|_{\infty} &\leq \sup_{x \in D(f)} \|x - \gamma(x)\|_{\infty} + \|\gamma(x) - \eta \circ \gamma(x)\|_{\infty} \\ &\leq \sup_{x \in D(f)} \|x - \gamma(x)\|_{\infty} + \sup_{x \in D(f)} \|\gamma(x) - \eta \circ \gamma(x)\|_{\infty}. \end{aligned}$$

Taking the infimum over all γ, η yields the desired result since every bijection from $D(f)$ to $D(h)$ can be obtained by some $\eta \circ \gamma$. The statement that $(f, g) \mapsto d_B(D(f), D(g))$ is a pseudometric follows immediately from the above properties, since indiscernibility of identicals $d_B(D(f), D(f)) = 0$ is a direct consequent of the identity of indiscernibles. \square

3.3.2 The stability theorem

The main result of [18] is the *Bottleneck Stability Theorem* for persistence diagrams: if two functions are close, then their persistence diagrams are also close. Due to the correspondence between piecewise linear functions and discrete pseudo-Morse functions (Section 2.8), the statement can be translated into the language of discrete Morse theory. We present a proof along the lines of [17], where stability is shown in a more elementary way for flat pseudo-Morse functions constructed from piecewise linear functions as described in Section 2.8.

Theorem 3.3 (Bottleneck Stability Theorem; Cohen-Steiner et al. [18, 17]). *Let $f, g : K \rightarrow \mathbb{R}$ be two discrete pseudo-Morse functions. Then the respective persistence diagrams satisfy*

$$d_B(D(f), D(g)) \leq \|f - g\|_\infty.$$

Proof. Without loss of generality, assume that f and g are flat pseudo-Morse functions. Otherwise, replace the functions with their respective canonical flat-tenings \bar{f} and \bar{g} , which satisfy $d_B(D(f), D(g)) = d_B(D(\bar{f}), D(\bar{g}))$ by Proposition 2.19 and $\|\bar{f} - \bar{g}\|_\infty \leq \|f - g\|_\infty$ by Theorem 2.21.

Now consider the family $f_t(\sigma) = (1 - t)f(\sigma) + tg(\sigma)$, $t \in [0, 1]$. Each f_t is a flat pseudo-Morse function according to Theorem 2.23. However, the order $<_{f_t}$ changes for varying t . According to Theorem 2.24 there exists a partition $0 = t_0 < t_1 < \dots < t_k = 1$ of $[0, 1]$ such that for any $0 \leq i < k$, the order $<_{f_t}$ is constant for $t \in (t_i, t_{i+1})$. This means that we can choose a single total order $<_i$ on K consistent with f_t for all $t \in [t_i, t_{i+1}]$. Consequently, all such functions f_t have the same set of persistence pairs P_i , and the corresponding points in the persistence diagrams $D(f_t)$ can be identified with each other. To simplify notation, we represent an essential cell σ as the pair $(\sigma, \infty) \in P_i$ and apply the convention $f_t(\infty) = \infty$. For each $r, s \in [t_i, t_{i+1}]$ we can use the identification of persistence pairs to bound the bottleneck distance from above by

$$d_B(D(f_r), D(f_s)) \leq \max_{(\sigma, \tau) \in P_i} \|(f_r(\sigma), f_r(\tau)) - (f_s(\sigma), f_s(\tau))\|_\infty \leq \|f_r - f_s\|_\infty = |s - r| \|f - g\|_\infty.$$

Summing over the partition and using the triangle inequality for the bottleneck distance, we obtain the claimed inequality

$$\begin{aligned} d_B(D(f), D(g)) &\leq \sum_{i=0}^{k-1} d_B(D(f_i), D(g_{i+1})) \\ &\leq \sum_{i=0}^{k-1} (t_{i+1} - t_i) \|f - g\|_\infty = \|f - g\|_\infty. \quad \square \end{aligned}$$

Theorem 3.3 can be paraphrased as follows:

Corollary 3.4. *Let $F_{\mathcal{K}}$ denote the metric space of pseudo-Morse functions on \mathcal{K} with the metric induced by the supremum norm, and let $D_{\mathcal{K}}$ denote the metric space of persistence diagrams of $F_{\mathcal{K}}$ with the bottleneck distance. Then the function $F_{\mathcal{K}} \rightarrow D_{\mathcal{K}}$, $f \mapsto D(f)$ mapping a function f to its persistence diagram is 1-Lipschitz.*

A 1-Lipschitz map is also called *metric* or *short*. Metric maps are the morphisms in the category of metric spaces [36]: compositions of metric maps are metric again.

Chapter 4

Optimal topological simplification of functions on surfaces

In this chapter we solve the problem of minimizing the number of critical points within a prescribed tolerance δ from a given pseudo-Morse function on a combinatorial surface. The result is achieved by establishing a connection between discrete Morse theory and persistent homology. Our method completely removes homological noise with persistence less than 2δ , achieving the lower bound on the number of critical points dictated by the stability theorem of persistent homology.

Convention and notation Throughout this section we consider a given pseudo-Morse function f consistent with a gradient vector field V on a combinatorial surface \mathcal{K} and a strict total order $<$ consistent with (f, V) . Moreover, we assume that either \mathcal{K} is orientable or that \mathbb{Z}^2 coefficients are used in the definitions of persistent homology.

4.1 Topological denoising by simplification

Small perturbations of a function can have a significant effect on the number of critical points. Since the characterization of critical points is highly unstable with respect to the supremum norm, critical points alone do not provide any meaningful information about uncertain data.

Assume for now that we know that the noise or uncertainty in our data is bounded by some small $\delta > 0$ in the supremum norm. Clearly, all critical pairs that have been created by noise can also be eliminated by a small perturbation again. On the other hand, we consider a critical point to be relevant if it cannot

be eliminated by a small perturbation. Therefore we consider the following strategy to recover from noisy data a function that reflects the relevant critical points of the original: find a function close to the noisy input with the least number of critical points.

Problem (Topological simplification on surfaces). *Given a pseudo-Morse function f on a combinatorial surface and a real number $\delta \geq 0$, find a pseudo-Morse function f_δ subject to $\|f_\delta - f\|_\infty \leq \delta$ and a gradient vector field V_δ consistent with f_δ such that V_δ has a minimum number of critical points.*

Persistent homology, and in particular Theorem 3.3, provides a lower bound on the number of persistence pairs and thus on the number of critical points among all pseudo-Morse functions f_δ with $\|f_\delta - f\|_\infty \leq \delta$:

Theorem 4.1 (Stability Bound). *Let f be a pseudo-Morse function consistent with a gradient vector field V . For any pseudo-Morse function f_δ consistent with a gradient vector field V_δ and satisfying $\|f_\delta - f\|_\infty \leq \delta$, the number of critical points of V_δ is bounded from below by the number of those critical points of V that have persistence $> 2\delta$.*

Proof. Let D and D_δ be the persistence diagrams of f and f_δ , respectively. By Theorem 3.3 we have $d_B(D, D_\delta) \leq \delta$. This means that there is a bijection γ between the individual elements of D and D_δ such that $\|x - \gamma(x)\|_\infty \leq \delta$ for all $x \in D$. Let

$$p = (p^*, p^\dagger) = (f(\sigma), f(\tau)) \in D$$

represent a persistence pair (σ, τ) of f with persistence $p^\dagger - p^* > 2\delta$. Letting $q = (q^*, q^\dagger)$, this implies that $p^* + \delta \geq q^*$ and $p^\dagger - \delta \leq q^\dagger$. Together with $p^\dagger - p^* > 2\delta$, this yields $q^\dagger - q^* > 0$. Hence there must be a persistence pair $(\tilde{\sigma}, \tilde{\tau})$ of f_δ corresponding to each persistence pair (σ, τ) of f with persistence $> 2\delta$. Since moreover $\|p - \gamma(p)\|_\infty < \infty$ implies that γ must map the essential critical cells of f exactly to those of f_δ , the claim follows. \square

The question about the tightness of this bound is a central criterion for estimating the informative value of the Bottleneck Stability Theorem 3.3. We are thus interested in functions that achieve the lower bound of Theorem 4.1:

Definition (Perfect δ -simplification). *Let f be a pseudo-Morse function consistent with a gradient vector field V and let $\delta \geq 0$. A perfect δ -simplification of f is a pseudo-Morse function f_δ consistent with a gradient vector field V_δ such that $\|f_\delta - f\|_\infty \leq \delta$*

and the number of critical points of V_δ equals the number of those critical points of V that have persistence $> 2\delta$.

In this chapter, we prove the following central result:

Theorem 4.2 (Tightness of the stability bound). *Let f be a discrete pseudo-Morse function on a combinatorial surface and let $\delta \geq 0$. Then there exists a perfect δ -simplification of f . In particular, the stability bound is tight for every discrete pseudo-Morse function f on a combinatorial surface and for every $\delta \geq 0$.*

The proof of Theorem 4.2 is constructive and hence leads to an algorithm. The corresponding construction is outlined in Section 4.3 and the proof is given in Section 4.4. However, the resulting algorithm has a running time that is quadratic in the input size. We present a more efficient algorithm in Chapter 5. Nevertheless, our proof of correctness for this algorithm makes use of the fact that Theorem 4.2 is already established. This is the reason why we present two separate constructions.

A similar result holds for (non-degenerate) discrete Morse functions (in a slightly different form, because in general only critical points with persistence $< 2\delta$ can be eliminated within a tolerance of δ in the set of discrete Morse functions):

Theorem 4.3. *Given a discrete Morse function f on a surface and $\delta > 0$, there exists a discrete Morse function f_δ such that $\|f_\delta - f\|_\infty < \delta$ and the number of critical points of f_δ equals the number of those critical points of f that have persistence $\geq 2\delta$.*

Proof. Let p be the largest persistence value of f such that $p < 2\delta$ and let $\tilde{\delta} = \frac{p}{2}$. According to Theorem 4.2 there exists a discrete pseudo-Morse function \tilde{f} consistent with a gradient vector field V_δ such that $\|\tilde{f} - f\|_\infty \leq \tilde{\delta}$ and the number of critical points of V_δ equals the number of those critical points of f that have persistence $> 2\tilde{\delta}$. By the way $\tilde{\delta}$ is chosen, this is the same as the number of those critical points of f that have persistence $\geq 2\delta$. Now according to Theorem 2.9, there exists a discrete Morse function f_δ consistent with V_δ with $\|f_\delta - \tilde{f}\|_\infty \leq \epsilon$ for any $\epsilon > 0$. By choosing $\epsilon < \delta - \tilde{\delta}$ we obtain

$$\|f_\delta - f\|_\infty \leq \|f_\delta - \tilde{f}\|_\infty + \|\tilde{f} - f\|_\infty \leq \delta + \epsilon < \delta. \quad \square$$

4.2 The persistence hierarchy and sequential cancelations

Persistence pairs on surfaces carry a certain hierarchical structure that allows us to establish a connection to the cancellation theorem of discrete Morse theory.

The main result of this section is that persistence pairs on surfaces can always be canceled sequentially if the order of cancelations respects this hierarchy. This fact is notable because persistence only considers the homology of sublevel sets, while a cancelation in Morse theory actually means a simplification in the stricter sense of homotopy type.

4.2.1 A hierarchy of persistence pairs

As motivated in Section 3.2, the persistence pairs on a surface are completely determined by the restriction of the function to the primal and dual graphs. Consequently, we now turn our attention to the dimension 0 persistence of a pseudo-Morse function f consistent with a gradient vector field V on a graph G (or, equivalently, a 1-dimensional CW complex). Again, let $G(\sigma)$ denote the level subcomplex of G at σ . Consider the minimum spanning tree $M(G)$ of G with edge weights given by f . As noted by Attali et al. [2], the edges of $M(G)$ correspond exactly to the negative 1-cells of G with regard to $<$. This is evident from the incremental construction of the minimum spanning tree in the algorithm of Kruskal [46]: edges of the graph G are visited in increasing order of their function value; if an edge τ closes a cycle, it is positive and both vertices of the edge belong to the same connected component of $G(\tau_-)$; otherwise, two connected components of $G(\tau_-)$ are merged and the edge is negative. Consider a negative cell τ merging two connected components of $G(\tau_-)$; both connected components are created by their respective least 0-cell with regard to $<$. Let $\tilde{\sigma}$ be the 0-cell creating the older class and σ the 0-cell creating the younger class ($\tilde{\sigma} < \sigma$). By definition, (σ, τ) is a persistence pair. Moreover, this defines a hierarchical structure on the 0-cells: $\tilde{\sigma}$ is the parent of σ . We call this relation the *persistence hierarchy*. By duality (see Section 3.2), we can define a hierarchy on the 2-cells in the same way.

Definition (parent, child, persistence hierarchy). *On a combinatorial surface \mathcal{K} , let (σ, τ) be a persistence pair with $\dim \sigma = 0$, and let $[\sigma] \in H_0(\mathcal{K}(\sigma))$ be the class created by σ (considered as a 0-cycle). Let $\tilde{\sigma}$ be the unique cell creating the class $[\tilde{\sigma}] \in H_0(\mathcal{K}(\tau))$ into which $[\sigma]$ gets merged by τ , i.e.,*

$$[\tilde{\sigma}] \notin \text{im}(i_*^{\tilde{\sigma}, \tau}) \quad \text{and} \quad [\tilde{\sigma}] = i_*^{\sigma, \tau}([\sigma]).$$

Then $\tilde{\sigma}$ is called the parent of σ (in the persistence hierarchy), and σ is called the child of $\tilde{\sigma}$. The transitive closure of the child relation is called descendant.




Figure 4.1: The persistence hierarchy. Both (σ, τ) and $(\hat{\sigma}, \hat{\tau})$ are children of, and hence nested in, $(\tilde{\sigma}, \tilde{\tau})$. Only (σ, τ) needs to be canceled before $(\tilde{\sigma}, \tilde{\tau})$ can be canceled.

Let (σ, τ) and $(\tilde{\sigma}, \tilde{\tau})$ be two persistence pairs. If either $\dim \sigma = \dim \tilde{\sigma} = 0$ and $\tilde{\sigma}$ is the parent of σ , or $\dim \tau = \dim \tilde{\tau} = 2$ and $\tilde{\tau}^*$ is the parent of τ^* (with regard to the persistence hierarchy on the dual complex), then we also call the pair $(\tilde{\sigma}, \tilde{\tau})$ the *parent* of (σ, τ) and (σ, τ) the *child* of $(\tilde{\sigma}, \tilde{\tau})$. The following definition and lemma justify this nomenclature:

Definition (nested persistence pairs). *Let (σ, τ) and $(\tilde{\sigma}, \tilde{\tau})$ be two persistence pairs. We say that (σ, τ) is nested in $(\tilde{\sigma}, \tilde{\tau})$ if $\tilde{\sigma} < \sigma < \tau < \tilde{\tau}$.*

Lemma 4.4. *Let (σ, τ) be a descendant of $(\tilde{\sigma}, \tilde{\tau})$ in the persistence hierarchy. Then (σ, τ) is nested in $(\tilde{\sigma}, \tilde{\tau})$.*

Proof. Without loss of generality, assume $\dim \sigma = 0$; otherwise, by duality, the argument can be applied to (τ^*, σ^*) instead of (σ, τ) .

By definition of the persistence hierarchy, $[\sigma]$ gets merged into the class $[\tilde{\sigma}] \in H_0(\mathcal{K}(\tau))$ created by $\tilde{\sigma}$. This implies that $\tilde{\sigma} < \sigma$. It also implies that the class created by $\tilde{\sigma}$ has not been merged by any cell of $\mathcal{K}(\tau)$, hence $\tau < \tilde{\tau}$. \square

4.2.2 Persistence cancellation sequences

We now turn our attention to the sequential cancellation of persistence pairs. Note that the cancellation theorem (Theorem 2.1) applies to vector fields, which only provide a *partial order* on the cells, while the notion of persistence is based on a *total order*. After canceling a persistence pair, the new vector field is no longer consistent with the initial total order. It is important to keep in mind that we only talk about persistence pairs of the *initial* total order $<$, which is consistent with (f, V) ; we do not consider a new total order after applying a

cancelation (which would complicate things considerably). Applying several cancelations results in a sequence of simplified vector fields:

Definition (persistence cancelation sequence). *A persistence cancelation sequence is a sequence of gradient vector fields (V_0, V_1, \dots, V_n) such that each V_i is constructed from V_{i-1} by canceling a persistence pair (σ_i, τ_i) using Theorem 2.1.*

A persistence cancelation sequence is called nested if in this construction every pair (σ_i, τ_i) nested in another pair (σ_j, τ_j) is canceled first, i.e.,

$$\sigma_j < \sigma_i < \tau_i < \tau_j \Rightarrow i < j.$$

A persistence cancelation sequence is called a δ -persistence cancelation sequence if exactly those persistence pairs are canceled that have persistence $\leq \delta$.

We can now state the key result connecting persistence pairs to cancelation pairs: A persistence pair (σ, τ) can be canceled from a vector field as soon as all descendants have been canceled.

Lemma 4.5. *On a combinatorial surface \mathcal{K} , let (V_0, V_1, \dots, V_i) be a persistence cancelation sequence. Assume that a persistence pair (σ, τ) has not been canceled in the sequence but that every descendant of (σ, τ) has been canceled. Then there exists a V_i -path from $\partial\tau$ to σ , and this path is unique.*

Assume further that every persistence pair nested in (σ, τ) has been canceled. If there is a unique V_i -path from $\partial\tau$ to another cell $\tilde{\sigma} \neq \sigma$ that is critical for V_i , then we have $\sigma > \tilde{\sigma}$.

Fig. 4.1 shows that the condition is sufficient but not necessary. The proof of Lemma 4.5 relies on a few auxiliary lemmas and is given after Lemma 4.8. The existence part of the statement is also shown in [24] for a particular cancelation sequence.

Lemma 4.6. *Let (V_0, V_1, \dots, V_i) be a persistence cancelation sequence and let (σ, τ) be a persistence pair with $\dim \sigma = 0$ that has not been canceled in the sequence. Let C be the connected component of the subcomplex $\mathcal{K}(\tau_-)$ containing σ , and let C denote the cells of C . Then every $(\phi, \rho) \in V_i$ with $\dim \phi = 0$ satisfies $\phi \in C \Leftrightarrow \rho \in C$.*

Proof. The claim is shown by induction over i . The base case follows from consistency of the total order $<$ with (f, V) . Consider the cancelation of a persistence pair (σ_i, τ_i) . If $\dim \sigma_i \neq 0$, the pairs in V_i of dimensions $(0, 1)$ stay




Figure 4.2: Example illustrating Lemma 4.6, showing the graph of a function f together with a simplified gradient vector field V_i . Note that some arrows are pointing upward. The highlighted cells form the connected component C of the subcomplex $\mathcal{K}(\tau_-)$ containing σ . Every pair in V_i is either completely in C or not at all.

unchanged and the claim immediately follows from the induction hypothesis. Now assume $\dim \sigma_i = 0$. We show that the claim holds for every $(\phi, \rho) \in V_i \setminus V_{i-1}$.

The non-critical cells of the vector field $V_i \setminus V_{i-1}$ are τ_i and the cells on the V_{i-1} -path $(\phi_0, \rho_0, \phi_1, \dots, \rho_{r-1}, \phi_r)$ from $\phi_0 \in \partial\tau_i$ to $\phi_r = \sigma_i$. By the induction hypothesis we have $\phi_k \in C \Leftrightarrow \rho_k \in C$. Because C is a subcomplex, we also have $\rho_{k-1} \in C \Rightarrow \phi_k \in C$ (with $\rho_{-1} = \tau_i$). Moreover, if $\sigma_i \in C$, then σ_i is a descendant of σ and by Lemma 4.4 (σ_i, τ_i) is nested in (σ, τ) , implying that σ_i and τ_i are in the same connected component of $\mathcal{K}(\tau_-)$. Hence we also have $\sigma_i \in C \Rightarrow \tau_i \in C$. Consequently, either all or none of the non-critical cells of $V_i \setminus V_{i-1}$ are contained in C and the claim immediately follows. \square

We also require the notion of the restriction of a vector field to a subcomplex:

Definition (restriction of a vector field to a subcomplex). *Let V be a discrete vector field on \mathcal{K} and let $\tilde{\mathcal{K}}$ be a subcomplex of \mathcal{K} with cells \tilde{K} . The restriction of V to $\tilde{\mathcal{K}}$ is $\tilde{V} = V \cap (\tilde{K} \times \tilde{K})$, i.e., the pairs of cells in V that are both in \tilde{K} .*

Note that new critical cells may arise when restricting a vector field to a subcomplex:

Lemma 4.7. *Let \tilde{V} be the restriction of a discrete vector field V on \mathcal{K} to a subcomplex $\tilde{\mathcal{K}}$. The critical d -cells of \tilde{V} are exactly the critical d -cells of V that are contained in $\tilde{\mathcal{K}}$ if and only if each pair $(\sigma, \tau) \in V$ with $\dim \sigma = d$ satisfies $\sigma \in \tilde{K} \Leftrightarrow \tau \in \tilde{K}$.*

Moreover, we use the following fact:

Lemma 4.8. *Let V be a discrete gradient vector field V on \mathcal{K} with only one critical 0-cell σ . Then there is a V -path from every 0-cell $\tilde{\sigma}$ to σ .*

Proof. Each V -path of dimension 0 ending at a non-critical cell $\tilde{\sigma} \neq \sigma$, $(\tilde{\sigma}, \tilde{\tau}) \in V$, can be extended by $\tilde{\tau}$ and the unique 0-cell $\hat{\sigma} \in \partial\tilde{\tau}$, $\hat{\sigma} \neq \tilde{\sigma}$. Since K is finite and V does not contain nontrivial closed paths, the extension will eventually end up at σ . \square

Proof of Lemma 4.5. Without loss of generality, assume $\dim \sigma = 0$; otherwise, by duality, the argument can be applied to (τ^*, σ^*) instead of (σ, τ) .

Let C be the connected component of the subcomplex $\mathcal{K}(\tau_-)$ corresponding to the homology class $[\sigma] \in H_0(\mathcal{K}(\tau_-))$ created by σ . Apart from σ , every 0-cell in C that is critical for V is a descendant of σ . By assumption, all descendants of σ have been canceled, and hence σ is the only 0-cell in C that is critical for V_i . By Lemmas 4.6 and 4.7, σ is also the only critical 0-cell in the restriction of V_i to C . By Lemma 4.8, there is a V_i -path to σ from every 0-cell in C , in particular from exactly one of the two 0-cells in $\partial\tau$, since $\partial\tau \cap C$ contains exactly one cell. By Lemma 2.3, this path is unique.

Now assume that every persistence pair nested in (σ, τ) has been canceled and there is a unique V_i -path from $\partial\tau$ to another cell $\tilde{\sigma} \neq \sigma$ that is critical for V_i . By assumption, $\tilde{\sigma}$ is not a descendant of σ , meaning that $\tilde{\sigma}$ and σ are in different connected components of $\mathcal{K}(\tau_-)$. Assume that $\sigma < \tilde{\sigma}$. Then by assumption there is no persistence pair $(\tilde{\sigma}, \tilde{\tau})$ nested in (σ, τ) . This implies that the component $\tilde{C} \neq C$ of $\mathcal{K}(\tau_-)$ containing $\tilde{\sigma}$ is created by $\tilde{\sigma}$. But by the definition of a persistence pair, τ merges C into a component created before σ , contradicting the assumption $\sigma < \tilde{\sigma}$. Thus we conclude that $\sigma > \tilde{\sigma}$. \square

As a consequence of Lemma 4.5, we can construct a sequence of cancelations that eliminate all persistence pairs below a certain persistence threshold:

Theorem 4.9. *Let f be a pseudo-Morse function on a combinatorial surface \mathcal{K} and let $\delta \geq 0$. Then there exists a nested δ -persistence cancellation sequence.*

Proof. If the subsequence $(V_0, V_1, \dots, V_{i-1})$ satisfies the assumptions of Lemma 4.5 for some persistence pair (σ_i, τ_i) , we can use Theorem 2.1 to construct V_i from V_{i-1} . A canonical choice satisfying these assumptions is given by canceling the persistence pairs (σ_i, τ_i) with persistence $\leq \delta$ according to the order $<$ on the negative cells, i.e., $\tau_i < \tau_{i+1}$ for every i . The claim then follows by induction. \square

Note that hierarchical cancelation of persistence pairs has also been considered in [24] in the context of Morse–Smale complexes on surfaces. While the existence of a gradient path between a persistence pair is already asserted in [24], the authors do not claim that every persistence pair can actually be canceled. Lemma 4.5 can also be applied in this context to show the uniqueness of the gradient path. As a consequence, we can obtain an analogous result to Theorem 4.9, stating that a hierarchical persistence simplification of Morse–Smale complexes is always possible on surfaces. The statement no longer holds in higher dimensions or for non-manifold 2-complexes, see Section 6.7.

4.2.3 Morse subcomplexes arising from a cancelation sequence

The vector fields V_i constructed in the proof of Theorem 4.9 by canceling pairs in the order of their negative cells have interesting properties. Consider the subcomplex $\mathcal{K}(\tau_i)$ and the Morse complex defined by restricting V_i to this subcomplex; let us call this Morse complex \mathcal{M}_i . Recall that \mathcal{M}_i is homotopy equivalent to $\mathcal{K}(\tau_i)$ by Theorem 2.5. We can observe that \mathcal{M}_i has a particularly simple structure: every negative cell τ_i is canceled by reversing the vector field and becomes a non-critical cell in V_i and all subsequent vector fields. Therefore, the cells of \mathcal{M}_i are in one-to-one correspondence to the *positive* critical cells of $\mathcal{K}(\tau_i)$. As a consequence, the homology groups $H_*(\mathcal{M}_i)$ are isomorphic to the cellular chain groups of \mathcal{M}_i : since all cells are positive as stated above, each cell increases the rank of a homology group by one. This in turn implies that the boundary group $B_*(\mathcal{M}_i)$ must be trivial. In other words, the boundary map $\tilde{\partial}$ of the cellular chain complex of \mathcal{M}_i is 0. The Morse complex \mathcal{M}_i is therefore the simplest CW complex (in terms of number of cells) having the homotopy type of $\mathcal{K}(\tau_i)$, and V_i is the gradient vector field on $\mathcal{K}(\tau_i)$ with the least number of critical points. Every connected component of \mathcal{M}_i must consist of either a single 0-cell, a wedge of circles (i.e., a set of 1-cells attached to a single 0-cell), or a closed surface consisting of a single 2-cell attached to a wedge of circles.

4.3 The plateau function

Assume now that we are given a nested 2δ -persistence cancelation sequence (V_0, \dots, V_n) with $V_0 = V$. For every V_i in the cancelation sequence, we inductively define a pseudo-Morse function f_i consistent with V_i , see Fig. 4.3 for an

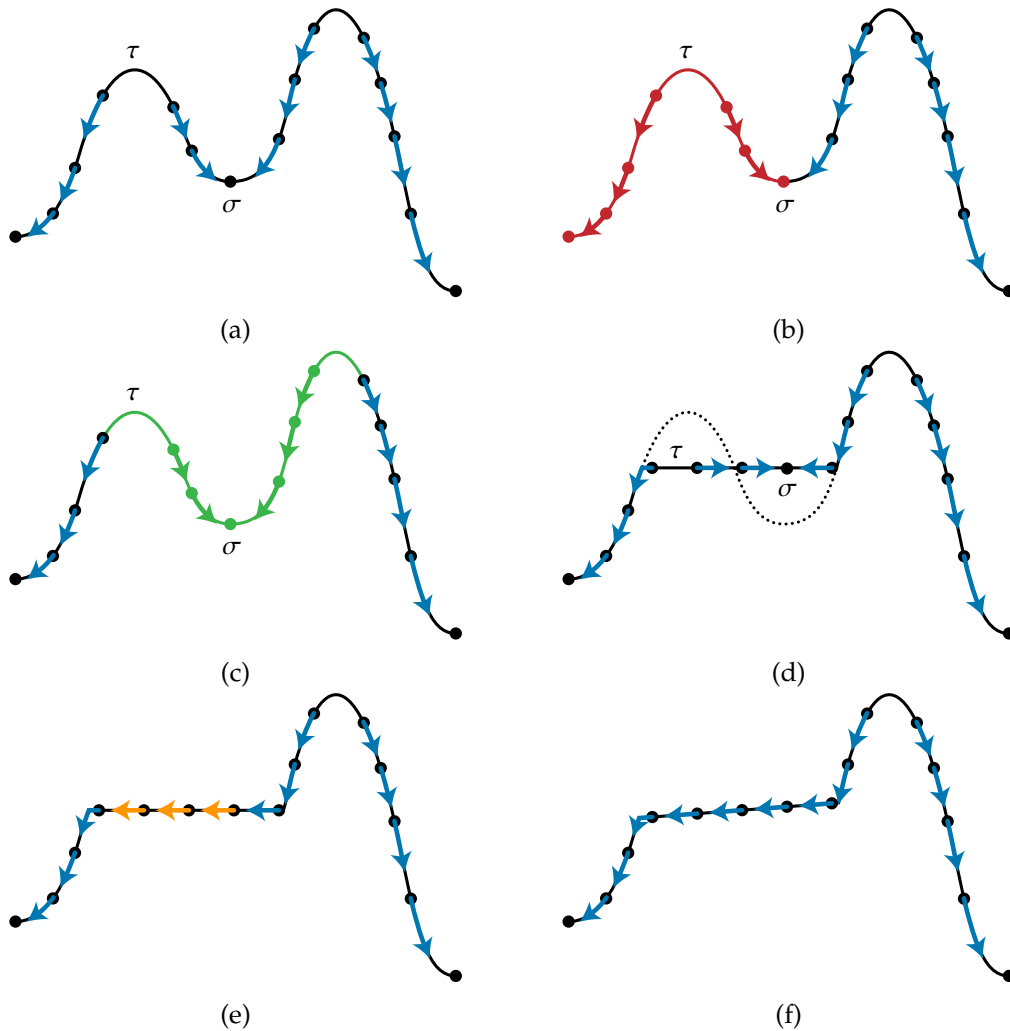


Figure 4.3: Cancellation of critical points. Starting with the graph of a function f_{i-1} together with gradient vector field V_{i-1} (a), the values of the descending set (b) of the upper critical point τ and of the ascending set (c) of the lower critical point σ are cut off at the average value m_i of the two critical points, creating a plateau (d). The old gradient directions are still consistent with the new function. The gradient vector field can now be reversed along the path between the critical points, eliminating the pair (e). The resulting function has a plateau, but can be perturbed slightly to become non-degenerate (f).

illustration. In Section 4.4, we show that $\|f_i - f\|_\infty \leq \delta$ and thus f_n is a perfect δ -simplification. Throughout this chapter, let $\prec_j := \prec_{V_j}$ denote the partial order induced by V_j .

By assumption we start with a pseudo-Morse function $f_0 := f$ consistent with $V_0 := V$. Suppose that we have constructed a pseudo-Morse function f_{i-1} consistent with V_{i-1} . Let (σ, τ) be the persistence pair that is canceled in the construction of V_i from V_{i-1} using Theorem 2.1. We define the corresponding *plateau function* f_i as follows:

$$m_i = \frac{f(\sigma) + f(\tau)}{2} \quad \text{and} \quad f_i(\rho) := \begin{cases} m_i & \text{if } \rho \succeq_{i-1} \sigma \text{ and } f_{i-1}(\rho) < m_i \\ & \text{or } \rho \preceq_{i-1} \tau \text{ and } f_{i-1}(\rho) > m_i, \\ f_{i-1}(\rho) & \text{otherwise.} \end{cases}$$

This means that the ascending set $\{\rho : \rho \succeq_{i-1} \sigma\}$ of σ with respect to V_{i-1} is raised to at least the value m_i , and analogously the descending set $\{\rho : \rho \preceq_{i-1} \tau\}$ of τ is lowered. Hence, f_i creates a local *plateau* at the value m_i . The following lemma is a consequence of the way we construct f_i from f_{i-1} and the fact that f_i is constant along the path from $\partial\tau$ to σ .

Lemma 4.10. *The plateau function f_i is consistent with both V_{i-1} and V_i .*

Proof. We show the claim by induction; the base case $i = 0$ is satisfied by assumption. It suffices to show that f_i is consistent with V_{i-1} : consistency with V_i then follows from the facts that V_i differs from V_{i-1} only along the path from $\partial\tau$ to σ and that f_i has constant value m_i along this path, since it is contained in both the ascending set of σ and the descending set of τ . Let ϕ and ψ be two cells with $\phi \leftarrow_{i-1} \psi$. To show consistency of f_i with V_{i-1} , it suffices to show that $f_i(\phi) \leq f_i(\psi)$.

If $f_i(\phi) = f_{i-1}(\phi)$ and $f_i(\psi) = f_{i-1}(\psi)$, then the claim follows by the induction hypothesis. Moreover, if $f_i(\phi) \neq f_{i-1}(\phi)$ and $f_i(\psi) \neq f_{i-1}(\psi)$, then from the definition of f_i we have $f_i(\phi) = f_i(\psi) = m_i$ and the claim holds trivially. We now assume $f_{i-1}(\phi) \neq f_i(\phi)$ and $f_{i-1}(\psi) = f_i(\psi)$; for the other case the proof is analogous. If $f_i(\phi) < f_{i-1}(\phi)$, the claim follows by the induction hypothesis. Hence assume $f_i(\phi) > f_{i-1}(\phi)$. Let (σ, τ) be the pair that is canceled when constructing V_i from V_{i-1} . We can deduce from the definition of f_i that $\phi \succ_{i-1} \sigma$ and hence also $\psi \succ_{i-1} \sigma$. It also follows from the definition that $f_i(\phi) = m_i$ and $f_i(\psi) \geq m_i$, and thus $f_i(\phi) \leq f_i(\psi)$. We conclude that f_i is consistent with V_{i-1} . \square

Note that the construction of the plateau function does not depend on the properties of combinatorial surfaces and can be applied to CW complexes of *arbitrary dimensions*. Moreover, it is not restricted to the cancelation of persistence pairs: whenever we have a pseudo-Morse function f consistent with a gradient vector field V and construct \widetilde{V} from V by a cancelation using Theorem 2.1, we can obtain a plateau function \tilde{f} that is consistent with both V and \widetilde{V} .

4.4 Checking the constraint

It remains to show that the plateau construction above produces admissible functions for the topological simplification problem, i.e., that each of the functions f_i satisfies the δ -constraint.

Lemma 4.11. *Each plateau function f_i satisfies $\|f_i - f\|_\infty \leq \delta$.*

Proof. We show the statement by induction. The base case is trivial since $f_0 = f$.

Let (σ, τ) be the persistence pair that is canceled when constructing V_i from V_{i-1} . We show that the δ -constraint is neither violated by increasing the value of any cell ρ in the ascending set of σ in V_{i-1} , nor by decreasing the value of any cell in the descending set of τ . Since $f_i(\rho) = f_{i-1}(\rho)$ for all cells ρ not treated in these two cases, the claim follows.

We first show $|f_i(\rho) - f(\rho)| \leq \delta$ for any cell $\rho \succeq_{i-1} \sigma$ with $f_{i-1}(\rho) < m_i$. By induction hypothesis we have a lower bound $f_{i-1}(\rho) \geq f(\rho) - \delta$. By construction of f_i , the value of ρ is increased: $f_i(\rho) = m_i > f_{i-1}(\rho)$. Therefore, the lower bound remains valid after step i :

$$f_i(\rho) > f_{i-1}(\rho) \geq f(\rho) - \delta.$$

To show the upper bound $f_i(\rho) \leq f(\rho) + \delta$, we first use $f(\tau) - f(\sigma) \leq 2\delta$ to obtain

$$f_i(\rho) = m_i = \frac{f(\sigma) + f(\tau)}{2} \leq \frac{f(\sigma) + (f(\sigma) + 2\delta)}{2} = f(\sigma) + \delta.$$

This is almost the desired inequality, except that the right hand side contains $f(\sigma)$ instead of $f(\rho)$. To finish the proof, it therefore suffices to show that $f(\sigma) \leq f(\rho)$. This, in turn, is a consequence of the facts that, according to the upcoming Lemma 4.12, $\sigma \prec_{i-1} \rho$ implies $\sigma < \rho$, and that $<$ is consistent with (f, V) .

It remains to show that $|f_i(\rho) - f(\rho)| \leq \delta$ for any cell $\rho \preceq_{i-1} \tau$ with $f_{i-1}(\rho) > m_i$. The proof of this statement is analogous to the above. \square




Figure 4.4: Example illustrating Lemma 4.13. Left: gradient vector field W (before reversing the path from $\partial\nu$ to μ). Right: gradient vector field \widetilde{W} (after path reversal). Note that we have the new relation $\phi <_{\widetilde{W}} \psi$ (corresponding in this example to a \widetilde{W} -path from ψ to ϕ). In the example, the conclusion $\phi \leq_W \nu$ and $\mu \leq_W \psi$ of Lemma 4.13 is reflected by the two W -paths from $\partial\nu$ to ϕ and from ψ to μ , respectively.

Lemma 4.12. *Let (V_0, \dots, V_n) be a nested persistence cancellation sequence and let (σ, τ) be a persistence pair of $<$ with σ and τ critical cells of V_i . Then for any $\rho \in K$,*

- (a) $\rho >_i \sigma$ implies $\rho > \sigma$, and
- (b) $\rho <_i \tau$ implies $\rho < \tau$.

Informally, this lemma guarantees that the ascending set of a positive paired critical cell in the *simplified* vector field V_i is consistent with the *original* function f , and analogously for the descending set of a negative cell. The proof crucially depends on the special properties of gradient vector fields on combinatorial surfaces. The statement does not hold in higher dimensions.

In order to prove Lemma 4.12, we first need to investigate how the reversal of a gradient vector field may change the induced partial order (see Fig. 4.4 for an example). The following statement holds for arbitrary regular CW complexes:

Lemma 4.13. *Let μ, ν, ϕ, ψ be (not necessarily disjoint) cells of a regular CW complex \mathcal{K} . Let W and \widetilde{W} be two gradient vector fields such that the cells μ and ν are critical for W . Assume further that there is a unique W -path from $\partial\nu$ to μ , and \widetilde{W} is constructed by reversing W along this path. If $\phi \not\leq_W \psi$ and $\phi <_{\widetilde{W}} \psi$, then $\phi \leq_W \nu$ and $\mu \leq_W \psi$.*

Proof. By definition of the induced partial order, $\phi <_{\widetilde{W}} \psi$ implies that there exists a path from ψ to ϕ in the Hasse diagram $H_{\widetilde{W}}$, i.e., a sequence (ρ_1, \dots, ρ_k) with $\rho_1 = \psi, \rho_k = \phi$ and $\rho_i \rightarrow_{\widetilde{W}} \rho_{i+1}$ for all $1 \leq i < k$. Here either ρ_i is a facet of ρ_{i+1} or ρ_{i+1} is a facet of ρ_i , and we therefore also have either $\rho_i \rightarrow_W \rho_{i+1}$ or $\rho_i \leftarrow_W \rho_{i+1}$. But since $\psi \not\leq_W \phi$, there exists a smallest index j such that $\rho_j \leftarrow_W \rho_{j+1}$. Since we have $\rho_j \rightarrow_{\widetilde{W}} \rho_{j+1}$ and the relations \leftarrow_W and $\leftarrow_{\widetilde{W}}$ differ only along the W -path from $\partial\nu$ to μ (including ν), it follows that the cells ρ_j and ρ_{j+1} are contained in

this W -path. Hence we have $\rho_j \succeq_W \mu$. Moreover, by the choice of j we have $\psi = \rho_1 \succeq_W \rho_j$. Therefore we conclude that $\psi \succeq_W \mu$. By an analogous argument one also shows that $\phi \leq_W v$. \square

Proof of Lemma 4.12. We only present the proof of part (a), which is done again by induction: We show that $\rho \succ_i \sigma$ implies $\rho > \sigma$ for all $0 \leq i \leq n$. Part (b) can be shown analogously.

The base case $i = 0$ is trivial since $>$ is a linear extension of \succ_0 . Assume that $\rho \succ_i \sigma$. If $\rho \succ_{i-1} \sigma$, the claim follows directly from the induction hypothesis. Hence we assume that $\rho \not\succeq_{i-1} \sigma$. Let $(\tilde{\sigma}, \tilde{\tau})$ be the persistence pair that is canceled when constructing V_i from V_{i-1} ; this implies $\tilde{\sigma} \prec_{i-1} \tilde{\tau}$. From Lemma 4.13 with $(W, \tilde{W}) = (V_{i-1}, V_i)$ and $(\mu, \nu, \phi, \psi) = (\tilde{\sigma}, \tilde{\tau}, \sigma, \rho)$, we infer that $\sigma \leq_{i-1} \tilde{\tau}$ and $\tilde{\sigma} \leq_{i-1} \rho$. This has two consequences:

- (i) $\sigma \prec_{i-1} \tilde{\tau}$ (since σ is critical for V_i while $\tilde{\tau}$ is not), and
- (ii) $\tilde{\sigma} \leq \rho$ (by the induction hypothesis).

To finish the proof of the claim, by (ii) it suffices to show that $\sigma < \tilde{\sigma}$. We proceed by case analysis on the dimensions of $\tilde{\sigma}$ and σ . Since these two cells are paired by assumption, they have dimension less than 2.

Case 1 ($\dim \sigma = 1, \dim \tilde{\sigma} = 0$): This case cannot occur since reversing the V_{i-1} -path from the 1-cell $\tilde{\tau}$ to the 0-cell $\tilde{\sigma}$ does not change the ascending set of any critical 1-cell (and in particular σ), contradicting $\rho \not\succeq_{i-1} \sigma$ and $\rho \succ_i \sigma$.

Case 2 ($\dim \sigma = 0, \dim \tilde{\sigma} = 1$): First assume $\tau < \tilde{\tau}$. If additionally $\tilde{\sigma} < \sigma$, this contradicts the assumption that the cancelation sequence is nested and (σ, τ) is canceled after $(\tilde{\sigma}, \tilde{\tau})$. Therefore $\tau < \tilde{\tau}$ implies $\sigma < \tilde{\sigma}$.

Now assume $\tau > \tilde{\tau}$. This means that σ creates a connected component of $\mathcal{K}(\tilde{\tau})$. Since $\sigma \prec_{i-1} \tilde{\tau}$ by (i), there is a path (ρ_1, \dots, ρ_k) from $\rho_1 = \tilde{\tau}$ to $\rho_k = \sigma$ in the Hasse diagram $H_{V_{i-1}}$. For each ρ_j we trivially have $\rho_j \prec_{i-1} \tilde{\tau}$ and hence $\rho_j < \tilde{\tau}$ by the induction hypothesis, implying that $\rho_j \in \mathcal{K}(\tilde{\tau})$. Moreover, since either ρ_j is a facet of ρ_{j+1} or ρ_{j+1} is a facet of ρ_j , we know that all ρ_j , and in particular σ and $\tilde{\tau}$, are in the same connected component of $\mathcal{K}(\tilde{\tau})$. In an analogous way one shows that $\tilde{\sigma}$ and $\tilde{\tau}$, and hence σ and $\tilde{\sigma}$, are in one and the same connected component. Since we know that σ created that component, it follows that $\sigma < \tilde{\sigma}$.

Case 3 ($\dim \sigma = \dim \tilde{\sigma} \in \{0, 1\}$): Since both σ and $\tilde{\tau}$ are critical cells of V_{i-1} , the relation $\sigma \prec_{i-1} \tilde{\tau}$ from (i) above implies the existence of a V_{i-1} -path from $\partial \tilde{\tau}$ to σ . We will show by contradiction that this path must be unique. To see this,

assume that there are two V_{i-1} -paths from $\partial\tilde{\tau}$ to σ . Without loss of generality, assume that $\dim \sigma = \dim \tilde{\sigma} = 0$ (and hence $\dim \tilde{\tau} = 1$); otherwise, by duality the following argument can be applied to $\sigma^*, \tilde{\tau}^*$ instead of $\tilde{\tau}, \sigma$. By Corollary 2.4, each of the 0-cells in $\partial\tilde{\tau}$ must belong to exactly one of the two V_{i-1} -paths from $\partial\tilde{\tau}$ to σ . Now by a similar argument as in Case 2 above, we see that each cell of these two V_{i-1} -paths is contained in the same connected component of $\mathcal{K}(\tilde{\tau}_-)$ as σ . But since $\tilde{\tau}$ is a negative 1-cell, the two 0-cells in its boundary belong to different connected components of $\mathcal{K}(\tilde{\tau}_-)$, a contradiction.

Hence, there is a unique V_{i-1} -path from $\partial\tilde{\tau}$ to σ . Since by assumption all pairs nested in $(\tilde{\sigma}, \tilde{\tau})$ are canceled in V_{i-1} , Lemma 4.5 asserts that $\tilde{\sigma}$ is the largest cell (with regard to \prec) with a unique V_{i-1} -path from $\partial\tilde{\tau}$ to $\tilde{\sigma}$. Since $\sigma \neq \tilde{\sigma}$, we obtain $\sigma \prec \tilde{\sigma}$. \square

Proof of Theorem 4.2. According to Theorem 4.9, there exists a nested 2δ -persistence cancelation sequence (V_0, V_1, \dots, V_n) for the vector field $V_0 = V$ consistent with the pseudo-Morse function f . Let f_n be the plateau function corresponding to V_n . Since f_n is consistent with V_n by Lemma 4.10 and $\|f_n - f\|_\infty \leq \delta$ by Lemma 4.11, it is a perfect δ -simplification. \square

Chapter 5

Efficient computation of topological simplifications

The definition of the plateau function in Section 4.3 canonically leads to an algorithm that runs in time quadratic in the input size. In this section we present a method for computing a perfect δ -simplification in time dominated by the computation of persistence pairs, i.e., $\mathcal{O}(\text{sort}(n))$, where $n = |K|$ is the number of cells of \mathcal{K} . Apart from this computation, all steps of our algorithm take linear time $\mathcal{O}(n)$. We stress that pre- and post-processing steps, like conversion from and to PL functions, also require only linear time $\mathcal{O}(n)$.

The algorithm can be summarized as follows. First, persistence pairs are computed using a variant of Kruskal's algorithm for minimum spanning trees. Next, the persistence pairs are used to construct a simplified gradient vector field by graph traversals of both the primal and dual 1-skeleton. In a third step, the simplified vector is used to compute the simplified function by a graph traversal on the Hasse diagram of the partial order induced by the simplified vector field.

Algorithm 1 Computing a perfect δ -simplification

compute persistence pairs using Algorithm 2

compute simplified gradient vector field using Algorithm 3

compute simplified function using Algorithm 4

5.1 Defining a consistent total order

Assume we are given a pseudo-Morse function f consistent with a discrete gradient vector field V as input. Let $<_T$ be a linear extension of $<_V$. If f is constructed from data given as a PL or piecewise constant function as explained in Section 2.8, then f is a flat pseudo-Morse function, i.e., V can be assumed to be the empty vector field (all cells are critical). In this case, $<_T$ can be constructed by choosing an arbitrary total order on the d -cells for each d and defining $\phi <_T \rho$ if $\dim \phi < \dim \rho$.

We define the order $<$ by combining $<_f$ and $<_T$ lexicographically: we have $\phi < \rho$ if and only if either

- (a) $f(\phi) < f(\rho)$ or
- (b) $f(\phi) = f(\rho)$ and $\phi <_T \rho$.

Since $<$ is equivalent to the lexicographic order on the product of the two totally ordered sets $(f(K), <)$ and $(K, <_T)$ restricted to the pairs $\{(f(\rho), \rho) : \rho \in K\}$, it is also a total order. It is consistent with (f, V) by construction.

5.2 Computing persistence pairs

Recall that the persistence pairs of dimension $(0, 1)$ are determined solely by the 1-skeleton \mathcal{K}_1 . Therefore, persistence pairs can be computed by applying a variant of Kruskal's algorithm [46] for finding a minimum spanning tree to both the primal and the dual 1-skeleton [23, 2]. Let $G = \mathcal{K}_1$ be the 1-skeleton of \mathcal{K} and let $M(G)$ be the minimum spanning tree of G (using the total order $<$ for determining the edge weights, which implies uniqueness of $M(G)$). Kruskal's algorithm for computing $M(G)$ initializes a graph T with the vertices of G , sweeps over the edges of G in order $<$, adds to T every edge of G that does not create a 1-cycle, and returns the final graph T . Note that the set of edges of $M(G)$ consists of all negative 1-cells together with all 1-cells τ such that $(\sigma, \tau) \in V$ for some σ ; all other 1-cells create a cycle in T . When encountering a negative 1-cell, we compute the persistence of the corresponding dimension $(0, 1)$ pair by storing for each connected component of the intermediate graph T the 0-cell that created it. Clearly we obtain all dimension $(0, 1)$ persistence pairs this way. Simultaneously, we construct the subgraph $M_\delta(G)$ of $M(G)$ by removing the negative 1-cells with persistence $> 2\delta$. In an analogous way, for the dual 1-skeleton G^* we can

compute the minimum spanning tree $M(G^*)$ and obtain the subgraph $M_\delta(G^*)$ together with all $(1, 2)$ persistence pairs. The algorithm is given in pseudo-code as Algorithm 2.

Kruskal's algorithm has a time complexity of $O(n \log n)$ when using comparison-based sorting. Assuming that the function values are represented by a small (i.e., $O(\log n)$) word size, Attali et al. [2] point out that persistence pairs on a graph can be computed in linear time $O(n)$ on a RAM using radix sort together with a linear-time algorithm for minimum spanning trees.

Algorithm 2 Constructing low-persistence pairs on a graph

```

function LOWPERSISTENCEEDGES(graph  $G$ ,  $\delta \geq 0$ )
   $T_\delta \leftarrow (\text{VERTEXSET}(G), \emptyset)$ 
  MAKESET(VERTEXSET( $G$ ))
  for all edges  $\tau = (u, v) \in G$  in order  $<$  do
     $r_u \leftarrow \text{FIND}(u)$ ,  $r_v \leftarrow \text{FIND}(v)$ 
    if  $r_u \neq r_v$  then
       $\tilde{\sigma} \leftarrow \min_{<}(r_u, r_v)$ ,  $\sigma \leftarrow \max_{<}(r_u, r_v)$ 
      if  $f(\tau) - f(\sigma) \leq \delta$  then
        add  $e$  to  $T_\delta$ 
        UNION( $\tilde{\sigma}, \sigma$ )
      end if
    end if
  end for
   $R_\delta \leftarrow \text{VERTEXSET}(G)$ 
  for all vertices  $v \in G$  do
     $r_v \leftarrow \text{FIND}(v)$ 
     $R_\delta \leftarrow R_\delta \setminus \{r_v\}$ 
  end for
  return  $(T_\delta, R_\delta)$ 
end function

 $(M_\delta(G), R_\delta(G)) \leftarrow \text{LOWPERSISTENCEEDGES}(G, \delta)$ 
 $(M_\delta(G^*), R_\delta(G^*)) \leftarrow \text{LOWPERSISTENCEEDGES}(G^*, \delta)$ 
return  $M_\delta(G), R_\delta(G), M_\delta(G^*), R_\delta(G^*)$ 

```

5.3 Extracting the gradient vector field

We now explain how to construct a simplified gradient vector field V_δ . To this end, we traverse (using depth-first search) each of the connected components of the primal graph $M_\delta(G)$ (constructed in the previous section) from the 0-cell that created the component. During this traversal, whenever we encounter an edge (1-cell) ψ that connects a previously visited vertex (0-cell) ρ with an unvisited vertex ϕ , we add (ϕ, ψ) to the gradient vector field V_δ . This construction takes $O(n)$ time.

We perform an analogous traversal for the dual graph $M_\delta(G^*)$. Again, whenever we encounter an edge ψ^* that connects a visited vertex ρ^* with an unvisited vertex ϕ^* (i.e., ψ is a 1-cell and ρ, ϕ are 2-cells of the original complex), we add (ψ, ϕ) to the gradient vector field V_δ . Note that the final V_δ results from both the primal and dual traversals and is a vector field on \mathcal{K} . Pseudo-code of the algorithm is given by Algorithm 3.

Theorem 5.1. *The gradient vector field V_δ is identical to the final vector field V_n of a 2δ -persistence cancellation sequence (V_0, \dots, V_n) .*

Proof. First observe that if $(\sigma, \tau) \in V_n$ and $\dim \sigma = 0$, then both σ and τ are cells of $M_\delta(G)$ since all non-critical cells of V_n either are non-critical for V as well or have persistence $\leq 2\delta$ (with respect to f and $<$). Moreover, the 0-cells creating a connected component of $M_\delta(G)$ are the only critical 0-cells of V_n (by definition) and of V_δ (by construction). Since $M_\delta(G)$ is a tree, the pairs $(\sigma, \tau) \in V_n$ with $\dim \sigma = 0$ are uniquely defined by this property. Thus the dimension (0, 1) pairs of V_n and V_δ coincide. By applying the dual argument to $M_\delta(G^*)$, the statement follows. \square

5.4 Constructing the simplified function

Finally, we construct a function f_δ (different from the plateau function defined in Section 4.3) that is consistent with the simplified gradient vector field V_δ . Consider the Hasse diagram $H := H_{V_\delta}$ of the strict partial order $<_{V_\delta}$ as described in Section 2.5.1. We visit the vertices K of H in a linear extension of $<_{V_\delta}$. The problem of finding a linear extension of a partial order is also called *topological sorting* and can be solved using depth-first search on H [20], with each edge of H being traversed according to its orientation. At each visited cell σ , we define

Algorithm 3 Constructing the simplified gradient vector field

```

function VISIT(vertex  $\rho$ , graph  $M$ , bool primal)
  if  $\rho$  has not been visited yet then
    mark  $\rho$  as visited
    for all vertices  $\phi$  in  $M$  with an edge  $\psi$  from  $\rho$  to  $\phi$  do
      if primal then
        add  $(\phi, \psi)$  to  $V_\delta$ 
      else
        add  $(\psi, \phi)$  to  $V_\delta$ 
      end if
      VISIT( $\phi, M, \text{primal}$ )
    end for
  end if
end function

 $V_\delta \leftarrow \emptyset$ 
for all vertices  $\rho$  in  $R_\delta(G)$  do
  VISIT( $\rho, M_\delta(G), \text{true}$ )
end for
for all vertices  $\rho^*$  in  $R_\delta(G^*)$  do
  VISIT( $\rho^*, M_\delta(G^*), \text{false}$ )
end for
return  $V_\delta$ 

```

$f_\delta(\sigma)$ as the minimum value that satisfies the lower bound $f_\delta(\sigma) \geq f(\sigma) - \delta$ and renders f_δ consistent with V_δ , i.e.,

$$f_\delta(\sigma) = \max\left(f(\sigma) - \delta, \max_{\rho \leftarrow V_\delta \sigma} f_\delta(\rho)\right).$$

The construction of f_δ also takes $O(n)$ time.

5.5 Correctness of the algorithm

Theorem 5.2. *The function f_δ constructed using the above algorithm is a perfect δ -simplification of f .*

Algorithm 4 Constructing the simplified function

```

function VISIT(vertex  $\sigma$ , graph  $H$ )
  if  $\sigma$  has not been visited yet then
    mark  $\sigma$  as visited
    for all vertices  $\rho$  in  $H$  with an edge from  $\sigma$  to  $\rho$  do
      VISIT( $\rho, H_{V_\delta}$ )
    end for
     $f_\delta(\sigma) \leftarrow f(\sigma) - \delta$ 
    for all vertices  $\rho$  with an edge from  $\sigma$  to  $\rho$  do
       $f_\delta(\sigma) \leftarrow \max(f_\delta(\sigma), f_\delta(\rho))$ 
    end for
  end if
end function

for all vertices  $\sigma$  in  $H_{V_\delta}$  do
  VISIT( $\sigma, H_{V_\delta}$ )
end for
return  $f_\delta$ 

```

Proof. By construction, f_δ is consistent with V_δ . At the same time, by Theorem 5.1, V_δ is the final vector field of a 2δ -persistence cancelation sequence. Therefore, by the definition of a perfect δ -simplification, it only remains to show that the constraint $\|f_\delta - f\|_\infty \leq \delta$ is satisfied. The lower bound $f_\delta \geq f - \delta$ is satisfied by construction. It thus remains to show the upper bound $f_\delta \leq f + \delta$.

Observe that the set of all perfect δ -simplifications consistent with V_δ is defined by a set of linear inequalities: the upper and lower bounds on the function values given by $f \pm \delta$, and the inequalities imposed by consistency with V_δ . Therefore, the set of δ -simplifications is a convex polyhedron $P \subset \mathbb{R}^n$ with $n = |K|$. The polyhedron P is bounded since it is a subset of the product of intervals

$$\prod_{\sigma \in K} [f(\sigma) - \delta, f(\sigma) + \delta].$$

From Theorem 4.2, we know that P is not empty. We now show that f_δ is contained in P .

First, consider the (unbounded) convex polyhedron \tilde{P} defined by the lower bound $f_\delta \geq f - \delta$ and the inequalities induced by V_δ . By construction, f_δ is

contained in \tilde{P} . Moreover, again by construction, f_δ minimizes the function value of any cell among all functions in \tilde{P} . In other words, for any function \tilde{f} in $P \subset \tilde{P}$, we have $\tilde{f} \geq f_\delta$. This implies the upper bound $f_\delta \leq \tilde{f} \leq f + \delta$. \square

Chapter 6

Discussion

6.1 Computational results

We implemented the algorithm of Chapter 5 in C++. For a complex with over 4 million cells (the cubical complex for a 1025×1025 pixel image), we obtained a running time of about 15 seconds for computing a perfect δ -simplification on a 2.4GHz Intel Core 2 Duo laptop.

6.2 Relation to simplification of persistence diagrams

We comment on two previously published results on topological simplification within a specified tolerance. The first method by Edelsbrunner et al. [25] is described specifically for PL functions; the second one by Attali et al. [2] works with a general filtration and is close to the formalism we use, without explicitly referring to discrete Morse theory. In both cases, the problem statement is different from the one considered in the present thesis. In particular, the problem is not set as an optimization problem. Rather, the goal is to simplify the *persistence diagram* in the following sense:

Problem. *Given a function f , find a function g with $\|f - g\|_\infty \leq \delta$ such that the persistence diagram of g consists of those points in the diagram of f except for the off-diagonal points with persistence $\leq \delta$, i.e., $D(g) = D(f) \setminus \{(s, t) : 0 < t - s \leq \delta\}$.*

In particular, the values of all critical points that are not canceled are required to remain unchanged. This requirement implies that in general only pairs with persistence $\leq \delta$ can be canceled, missing a factor of 2 compared to our result and to the lower bound of Theorem 4.1. The problem is depicted in Fig. 6.1; a variation of this example can be found in Edelsbrunner et al. [25]. In this




Figure 6.1: Canceling a volcano. Left: There are two persistence pairs; assume that only the right pair (with persistence 2δ) is canceled. Middle: If uncanceled critical points have to remain at their original value, as assumed by previous methods, the peak of the volcano can only be canceled by raising the saddle that is paired with the peak, perturbing the function by 2δ . Right: Since our approach does not have this restriction, the function only needs to be perturbed by δ in order to cancel the pair.

example, the peak of a volcano is paired with a saddle on its right. Moreover, on the crater at roughly the same height as the peak there is another saddle, which is paired with a minimum inside the volcano. It is evident that the peak of the volcano cannot be lowered significantly without also lowering the saddle on the crater. Therefore, if we want to cancel the saddle-peak pair and keep the other pair at their original values, the saddle has to be raised all the way up to the peak.

In the language of the present thesis, the method of Attali et al. [2] can be described as a variant of the *plateau* construction: for each canceled pair (σ_i, τ_i) , set the new value m_i to the value of the cell of (σ_i, τ_i) that is not a 1-cell. Hence, only the ascending/descending set of the critical 1-cell of each persistence pair is raised/lowered to the value of the paired cell. Note that both the ascending and descending sets of a critical 1-cell are each the union of two gradient paths.

6.3 Symmetrizing the algorithm

The method described in Chapter 5 assigns to each cell the smallest possible value. As a consequence, the output function differs from the input function f even if the input function itself is already a perfect δ -simplification. Moreover, the method is not symmetric in the sense that we obtain an output function which *maximizes* the values if we apply the algorithm to the function $-f$ on the dual complex and return the negative of the simplified function. Since




Figure 6.2: Visualization of simplification artifacts. Function values are indicated by gray levels. Left: Original function. Middle: Function obtained by the algorithm of Chapter 5. Note the bright path joining the two spots. Right: Function obtained after constraint energy minimization according to Section 6.5. While the simplified topological structure is maintained, the visual appearance is closer to the original function.

both the minimal and maximal solutions are points of a convex polyhedron as explained in Section 5.5, we can take the component-wise arithmetic mean to obtain another perfect δ -simplification.

With this modification, if the input function f is already a perfect δ -simplification, then the minimal solution is given by $f - \delta$, while the maximal solution equals $f + \delta$, so the arithmetic mean of both solutions returns f again as desired.

6.4 Flooding and carving artifacts

The methods presented in the present thesis can be seen as combinations of the so-called *carving* and *flooding* approaches. Consequently, they also inherit some characteristics of these methods that may not always be desirable in practical applications (see Fig. 6.2).

Carving methods [61, 25, 2]) cancel a pair of critical cells by changing only the descending or ascending set of the 1-cell (saddle), as discussed in Section 6.2. This results in a noticeable thin path being carved in the function. On the other hand, modifying only the descending or ascending set of extrema, i.e., lowering maxima and raising minima, produces regions with constant function value; this is called *filling* or *flooding* [37, 21]). Although this effect is less disturbing, it might appear unnatural in certain applications. In the next section, we propose a way to remedy both kinds of artifacts.




Figure 6.3: Topographic map of the elevation data set “Puget Sound” [48], showing the region around Tacoma. Contour lines are shown every 500 meters. The elevation data are converted from a 512×512 grid into a pseudo-Morse function on 1050625 cells. 33120 critical cells have persistence > 0 (persistence diagram shown in Fig. 6.6).




Figure 6.4: Simplified elevation function f_δ obtained after constraint Dirichlet energy minimization of the difference $f - f_\delta$ according to Section 6.5 with $\delta = 500$ meters. The function has one minimum, three saddles, and three maxima. Despite the quite drastic perturbation, the overall appearance of the contours is largely kept intact.




Figure 6.5: Topologically constrained smoothing, obtained after constraint energy minimization of the simplified function f_δ according to Section 6.5 with $\delta = 50$ meters. Although the allowed tolerance is only one tenth of the one shown in Fig. 6.4, the visual appearance is significantly different due to the smoothing.




Figure 6.6: Persistence diagrams of the elevation data set visualized in Fig. 6.3.

6.5 Combining topological simplification and energy methods

As mentioned in Section 5.5, the set of perfect δ -simplifications consistent with the simplified gradient vector field V_δ is a convex polyhedron P . Hence, the presented method can be combined with energy minimization methods, since the polyhedron P can be used as the feasible region for an arbitrary convex optimization problem. For example, we used the interior point solver Ipopt [65] to minimize (a discretization of) the Dirichlet energy of the difference $f_\delta - f$ in order to obtain a function f_δ that looks as similar as possible to the input function f (see Figs. 6.2 and 6.4). Alternatively, we minimized the Dirichlet energy of the simplified function itself in order to obtain smooth contour lines (see Fig. 6.5). Note, however, that solving a constraint optimization problem is much more expensive than finding an initial perfect δ -simplification. Computing the functions shown in Figs. 6.4 and 6.5 each took about an hour, while the initial simplification took only five seconds to compute.

6.6 Topological simplification on regular CW complexes

The example of Fig. 6.7 shows that a perfect δ -simplification may not exist on a non-manifold 2-dimensional cell complex. For the sake of simplicity, the example is given for a non-regular CW complex; it is straightforward to rephrase this example using a regular CW complex by subdividing the cells. The complex consists of two 0-cells ζ and γ with $f(\zeta) = f(\gamma) = 0$, three 1-cells a , b , and c

with $f(a) = 1$, $f(b) = 2$, and $f(c) = 0$, and two 2-cells A and B with $f(A) = 2$ and $f(B) = 3$. Note that the complex is not a manifold since it is not locally euclidean at the 1-cell b . A canonical ordering leads to the persistence pairs (a, A) , (b, B) , and (γ, c) . To obtain a perfect δ -simplification for $\delta = 0.5$, one would need to set $f_\delta(b) = f_\delta(B) = 2.5$ and $f_\delta(a) = f_\delta(A) = 1.5$. The corresponding simplified gradient vector field would be $V_\delta = \{(a, A), (b, B)\}$. But since b is a facet of A , we must have $f_\delta(b) \leq f_\delta(A)$. Hence, we cannot cancel both (a, A) and (b, B) at the same time. This constellation also appears in [23] under the name *conflict of type (1,2)*.




Figure 6.7: A discrete Morse function on a 2-complex that does not have a perfect δ -simplification. The function values of the cells are indicated in parentheses.

Since such a 2-complex can also appear as a level subcomplex of an n -manifold CW complex for $n \geq 3$ (e.g., a triangulated 3-ball), the example also shows that a perfect δ -simplification does not always exist for functions on manifolds. Thus the topological simplification problem is more difficult in higher dimensions.

In fact, we will see in a moment that the following generalization of the simplification problem considered in this thesis to non-surface complexes is NP-hard:

Problem (Topological simplification on regular CW complexes). *Given a pseudo-Morse function f on a regular CW complex \mathcal{K} and a real number $\delta \geq 0$, find a pseudo-Morse function f_δ subject to $\|f_\delta - f\|_\infty \leq \delta$ and a gradient vector field V_δ consistent with f_δ such that V_δ has a minimum number of critical points.*

We will show the claim by a reduction from the following problem, which does not refer to a function and has no tolerance constraint, but depends only on the properties of the complex \mathcal{K} itself:

Problem (optimal discrete gradient vector field). *Given a regular CW complex \mathcal{K} , find a discrete gradient vector field on \mathcal{K} with a minimum number of critical cells.*

Finding an optimal discrete gradient vector field, in turn, has been shown to be NP-hard:

Theorem 6.1 (Eğecioglu and Gonzalez [27], Joswig and Pfetsch [39]). *Computing an optimal discrete gradient vector field is NP-hard for regular CW complexes of dimension 2.*

Thus, we learn that there is no hope to extend our efficient solution method beyond surfaces:

Corollary 6.2. *Topological simplification on a regular CW complex \mathcal{K} of dimension 2 is NP-hard.*

Proof. Given any pseudo-Morse function f , there exists a δ such that any gradient vector field on \mathcal{K} is consistent with some function f_δ satisfying $\|f_\delta - f\|_\infty \leq \delta$. For example, we can choose $\delta = \max_\rho |f(\rho)|$ and $f_\delta : \rho \mapsto 0$. Therefore finding an optimal discrete gradient vector field can be reduced to topological simplification. Since finding an optimal discrete gradient vector field is NP-hard for 2-complexes, so is topological simplification. \square

6.7 Persistence pairs that cannot be canceled

Not every persistence pair can be canceled in the strict sense of Morse theory: there are CW complexes on which every discrete Morse function must have a persistence pair. In other words, they do not admit a *perfect discrete Morse function* (a discrete Morse function with only positive critical cells, for any field of coefficients). This can have several reasons. A discrete Morse function on a complex that is contractible but not collapsible must have a persistence pair. This is explained by the following observation. Collapsibility would mean by definition that there is a gradient vector field with only one critical 0-cell. Any gradient vector field on a non-collapsible complex thus must have additional critical cells. But since contractibility implies that the complex has the homology of a point, all additional critical cells must be paired in persistence pairs. Classical examples are given by the *dunce hat* [66] or *Bing's house with two rooms* [7].

Certain triangulations of the n -sphere for $n \geq 3$ do not admit perfect discrete Morse functions [41]. However, any PL triangulation of the n -sphere admits

a perfect discrete Morse functions after a finite number of subdivisions [28]. Note that not every triangulation of the n -sphere is a PL manifold: Cannon [12] showed that the double *suspension* of a $(n - 2)$ -homology sphere Σ^n is homeomorphic to an $(n + 2)$ -sphere, but the homeomorphism is not PL if Σ^n is not homeomorphic to the n -sphere.

Such homology spheres are also known to admit no perfect Morse functions. Assume that a triangulated homology n -sphere Σ^n has a nontrivial fundamental group $\pi_1(\Sigma^n)$. Then $\pi_1(\Sigma^n)$ must be generated by a nonempty set of critical 1-cells for V for any gradient vector field V on Σ^n . There are even some submanifolds of \mathbb{R}^3 that admit no perfect Morse function. The complement C of a nontrivial knot in S^3 is embeddable in \mathbb{R}^3 and has a fundamental group $\pi_1(C)$ of rank at least 2, where the *rank* of a group is the smallest cardinality of a generating set. Thus every discrete Morse function on C has at least 2 critical 1-cells. But the first homology group $H_1(C)$ has only rank 1 by Alexander duality [35], so at least one of the critical 1-cells must be negative. Since a knot complement can also occur as the sublevel set of a function on the 3-ball, this problem adds to the difficulties regarding topological simplification in higher dimensions discussed in Section 6.6.

6.8 Removing local extrema from functions on manifolds

Recall that the proof of tightness of the stability bound relies heavily on the fact that only persistence pairs containing a local extremum (minimum or maximum) can appear, and no persistence pairs of middle dimensions, which cannot always be canceled in the stronger sense of Morse theory as explained in Sections 6.6 and 6.7. In fact, as a consequence the same constructions and proofs presented in Chapter 4 can also be applied to the problem of minimizing the number of local extrema of a pseudo-Morse function within a δ -tolerance on any d -dimensional manifold CW complex.

Problem (Extrema simplification on manifolds). *Given a pseudo-Morse function f on a regular manifold CW complex and a real number $\delta \geq 0$, find a pseudo-Morse function f_δ subject to $\|f_\delta - f\|_\infty \leq \delta$ and a gradient vector field V_δ consistent with f_δ such that V_δ has a minimum number of local extrema.*

By a local extremum we mean a critical cell of dimension 0 or d . Note that in the case $d = 2$ this problem is equivalent to the topological simplification

problem by the following simple counting argument. Let c_i denote the number of critical cells of dimension i . Since the Euler characteristic $\chi = c_0 - c_1 + c_2$ is a topological invariant, we have $c_0 + c_1 + c_2 = 2(c_0 + c_2) - \chi$. Thus the number of critical points is minimal if and only if the number of extrema is minimal.

We obtain the following theorem, which can be proven by a straightforward modification of the proof of Theorem 4.2.

Theorem 6.3. *Given a pseudo-Morse function f on a regular closed manifold CW complex and a real number $\delta \geq 0$, there exists a pseudo-Morse function f_δ consistent with a gradient vector field V_δ such that $\|f_\delta - f\|_\infty \leq \delta$ and the number of local extrema of V_δ equals the number of those local extrema of f that have persistence $> 2\delta$. This number is minimal.*

We call such a function f_δ a *perfect δ -extrema simplification* of f .

6.9 Matching local extrema of functions on manifolds

The persistent extrema of a function have another interesting property. Assume that the original (noise-free) function has no extrema of low persistence. Then it turns out that we can match its extrema to the extrema of the simplification of a noisy function close to the original one in a canonical, well-defined way. First, we obtain the following corollary from Theorem 6.3:

Corollary 6.4. *Let f be a discrete pseudo-Morse function on a surface all of whose persistence pairs have persistence $> 4\delta$. Let g be another pseudo-Morse function with $\|g - f\|_\infty \leq \delta$. Let g_δ be a perfect δ -extrema simplification of g . Then g_δ has the same numbers of local minima and maxima as f .*

The following theorem shows how the discrete gradient flow yields a bijection of the local extrema of f and g_δ :

Theorem 6.5. *Let (f, V) be a discrete pseudo-Morse function with a consistent gradient field all of whose local extrema have persistence $> 4\delta$. Let g be another pseudo-Morse function f with $\|g - f\|_\infty \leq \delta$. Let (g_δ, W_δ) be a perfect δ -extrema simplification of g .*

Then there is a bijection γ from the minima (maxima) of f to the minima (maxima) of g_δ defined by the property that every minimum (maximum) σ of f is contained in the ascending (descending) set of $\gamma(\sigma)$ and vice versa.

Proof. Without loss of generality assume that σ is a local minimum, i.e., a critical 0-cell of V . Let $\gamma : \text{Cr}_0 V \rightarrow \text{Cr}_0 W_\delta$, $\rho \mapsto \Phi_\delta^\infty(\rho)$, where $\text{Cr}_0 V$ denotes the critical 0-cells of V and Φ_δ denotes the discrete gradient flow of the vector field W_δ , i.e., ρ is mapped to the cell ending the maximal W_δ -path starting at ρ . Let $\tilde{\sigma} = \gamma(\sigma)$. This cell is critical for W_δ , being the last cell of a maximal W_δ -path. First note that $f(\tilde{\sigma}) - f(\sigma) \leq 4\delta$, since

$$f(\tilde{\sigma}) \leq g_\delta(\tilde{\sigma}) + 2\delta \leq g_\delta(\sigma) + 2\delta \leq f(\sigma) + 4\delta.$$

This implies that $\tilde{\sigma}$ is contained in the level subcomplex $\mathcal{K}(f(\sigma) + 4\delta)$. Since f has no minima with persistence $\leq 4\delta$, the connected component of $\mathcal{K}(f(\sigma) + 4\delta)$ created by σ contains no critical 0-cells apart from σ . Since $\tilde{\sigma}$ is contained in that connected component, it follows that $\Phi^\infty(\tilde{\sigma}) = \sigma$, where Φ denotes the discrete gradient flow of the vector field V . Hence γ has a left inverse: $\Phi^\infty \circ \gamma(\sigma) = \sigma$. Since f and g_δ have the same same numbers of minima and maxima, it follows that γ is a bijection. \square

6.10 Topological simplification for simplexwise linear functions is NP-hard

In view of our results, it seems surprising that a very slight modification of the problem setting leads to an intractable problem, as we now discuss.

Gray et al. [33] show that minimizing the number of extrema is NP-hard for *simplexwise linear functions with individual tolerance intervals* for each vertex. Recall that in our problem we consider a single tolerance value δ instead. We extend their argument so that it can also be applied to our problem. Recall from Section 6.8 that on a surface, minimizing extrema is equivalent to minimizing critical points. We consider the following variant of the topological simplification problem:

Problem (Simplexwise linear topological simplification on surfaces). *Given a simplexwise linear function f on a simplicial surface and a real number $\delta \geq 0$, find a simplexwise linear function f_δ subject to $\|f_\delta - f\|_\infty \leq \delta$ such that f_δ has a minimum number of critical points.*

We deliberately talk about *simplexwise* linear functions instead of just piecewise linear functions here. The key difference is that a simplexwise linear functions,

as the name suggests, must be linear on each simplex of the given triangulation, while a piecewise linear function is linear on each simplex of *some arbitrary* triangulation. It turns out that this is what makes the big difference in complexity: the hardness of the problem in the simplexwise linear setting arises from the possibility that the input contains multiple saddles. Going from simplexwise linear functions to discrete Morse functions (Section 2.8) can be interpreted as splitting multiple saddles.

It is also easy to see that the stability bound cannot always be matched in this setting. Consider a vertex which is a multiple saddle of a PL function. The change in homotopy type at the level of this saddle corresponds to attaching multiple 1-cells. Assume that one of these cells is positive and another one is negative. Assume further that both cells are contained in persistence pairs of persistence 2δ . Canceling *both* of these pairs within the δ -tolerance would require raising *and* lowering the multiple saddle vertex at the same time. Hence, without subdividing the complex, it is not possible to cancel both persistence pairs.

We show NP-hardness of the simplexwise linear topological simplification problem by a reduction to the following problem, called the PLANAR 2-DISJOINT MAXIMALLY CONNECTED SUBGRAPHS problem, or P2-MAXCON for short:

Problem (P2-MAXCON). *Given a finite planar graph $G = (E, V)$ and two disjoint subsets $R, B \subset V$, $R \cap B = \emptyset$, find two disjoint subsets $R' \supset R$, $B' \supset B$, $R' \cap B' = \emptyset$, such that the total number of connected components of the respective subgraphs of G induced by R' and B' is minimized.*

We interpret the problem as a coloring problem: The vertices R and B are colored red and blue, respectively. The problem asks to color the remaining white vertices $W = V \setminus R \setminus B$ in either color such that the total number of red and blue components is minimized.

Theorem 6.6 (Gray et al. [33]). *P2-MAXCON is NP-hard.*

Theorem 6.7. *Topological simplification for simplexwise linear functions on simplicial surfaces is NP-hard.*

Proof. We present a polynomial time reduction of P2-MAXCON to the simplexwise linear topological simplification problem, i.e., we show that any instance of P2-MAXCON can be transformed into an instance of topological simplification.

Let $G = (V, E)$ be a finite planar graph with vertices partitioned into pairwise disjoint subsets $R \cup B \cup W = V$. We assign the values 1, 3, and 5 to the vertices in B , W , and R , respectively. Next, we embed the graph into the plane, insert a triangle into the interior of each inner face, and triangulate the resulting graph, see Fig. 6.8. The newly inserted triangles are further subdivided as indicated by the label a in Fig. 6.8. This subdivision of the triangle a involves quadrangles labeled b and c ; these are again triangulated as indicated in Fig. 6.8. In the triangulation b there are k vertices with value 2 and $k + 1$ vertices with value 6, where $k > |V| + 3|F|$ and $|F|$ denotes the number of inner faces of G . Similarly, in the triangulation c there are k vertices with value 0 and $k + 1$ vertices with value 4. The function values of the additional vertices are assigned as shown.

We thus obtain a triangulation of a planar domain together with a simplexwise linear function f given by the values on the vertices. Letting $\delta = 1$, this yields a problem instance of simplexwise linear topological simplification on surfaces.

We now argue that any optimal solution for the above problem instance must take on certain fixed values on the vertices of the triangles labeled a , as indicated by a' , b' , and c' in Fig. 6.8. First observe that the indicated partial solution has two maxima (with value 5) and one minimum (with value 1) inside each instance of the triangle a . This implies that there exists a function on the triangulation within distance δ from f that has at most $|V| + 3|F|$ extrema. Here, when we talk about extrema of a simplexwise linear function, we mean *extremal components*, i.e., a connected subgraph whose vertices all have the same extremal value counts as one single extremum.

As a consequence, the vertices in the triangulation b with original values 6 and 4 must take on the value 5 in the simplified function, as otherwise the k vertices with original value 6 are extrema of the simplified function. Such a function cannot be optimal since we know that a function with less than $|V| + 3|F| < k$ extrema exists. Similarly, the vertices in the triangulation c with original values 0 and 2 must take on the value 1 in the simplified function. Now consider one of the boundary vertices of the triangle a (with original value 3). This vertex is adjacent to at least $k + 1$ vertices with input value 4 (from a triangulation c). In order to prevent these adjacent vertices from being extrema of the solution, the value 3 of the vertex must not be lowered in the solution. At the same time, the vertex is also adjacent to at least $k + 1$ vertices with original value 2 (from a triangulation b), so by a similar argument the value also must not be raised. We conclude that the solution value of the vertex remains 3, and all adjacent vertices

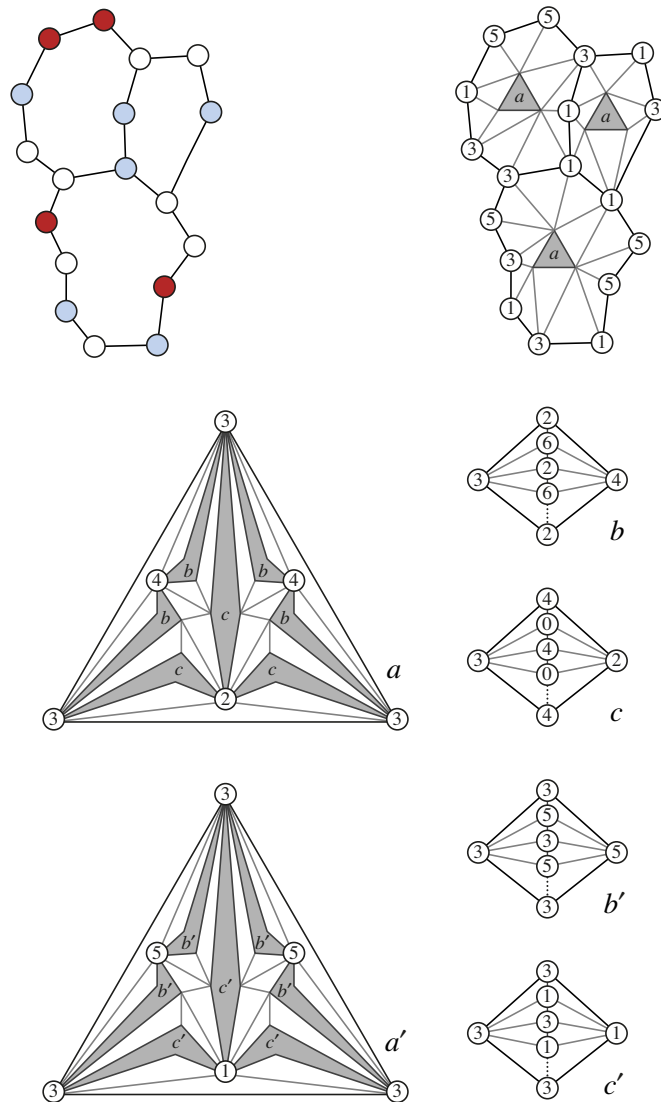


Figure 6.8: Gadgets used in the reduction of PL topological simplification to P2-MAXCON. Upper left: Graph G with vertex subsets B and R colored blue and red. Upper right: triangulation of G with triangles labeled a embedded in each face. Red and blue vertices of G are assigned the function values 5 and 1, respectively. Middle: Triangulation for each instance of the triangle a and for the quadrangles b and c appearing in a , together with function values. Bottom: Function values for the vertices of the triangle a in every optimal δ -simplification with $\delta = 1$.

with original values 2 and 4 also obtain the solution value 3. The solution values are therefore determined for all vertices inside the triangle a as shown in Fig. 6.8. Note that the vertices in the boundary of the triangle a (with output value 3) are not extrema of the solution, since they are adjacent both to vertices with values 1 and 5, respectively.

Now the problem of simplexwise linear topological simplification of f is equivalent to that of minimizing connected components of G after recoloring [33]: The only way to eliminate minima of f on G is by connecting minima (blue components) by lowering a connecting set of white vertices to the level 2 of the new minimum component (corresponding to coloring the white vertices blue). Likewise, maxima can only be eliminated by raising the connecting white vertices to the level 4, corresponding to coloring them red. Hence our reduction is complete. \square

Chapter 7

Future work

As we have discussed in Sections 6.6 and 6.7, topological simplification in dimensions higher than 2 involves additional difficulties, and a perfect simplification cannot always be achieved. However, the problem is still of great practical interest, since a standard method for visualization of 3-dimensional scalar data sets such as medical images is the rendering of contour surfaces. Due to the instability of critical points under small perturbation, this method is highly susceptible to noise. Simplifying only the extrema of the data, i.e., the persistence pairs of dimensions (0, 1) and (2, 3) (Section 6.8), already provides promising results with regard to the removal of topological artifacts (see Fig. 7.1 for an example). This method could be extended by a heuristic algorithm, which greedily cancels persistence pairs of dimensions (1, 2). How well would such an algorithm perform on noisy real-world data?

What can be said if one relaxes the notion of critical point, considering changes in the homology of sublevel sets (like in persistent homology) instead of changes in their homotopy type (like in Morse theory)? One would expect that all persistence pairs can be canceled in this sense. Is it possible to use a variant of the cancelation method of discrete Morse theory in this context, possibly dropping the restriction to *gradient* vector fields? Can one construct nested cancelation sequences like in Theorem 4.9? Even if the stability bound cannot be matched in general, can it possibly be approximated?

Certain previous denoising methods are already known to have good properties with regard to the removal of noise-induced critical points. Worth particular mention is the minimization of the *Rudin-Osher-Fatemi* functional [60], also known as TV- L^2 denoising. In [5], a first attempt has been made to relate this method to persistence and the simplification methods presented in this thesis. For an input function on an interval, TV- L^2 denoising minimizes the number




Figure 7.1: Visualization of a MRT angiography of a human head [64] by contour surfaces. Top: The noise in the data generates a large number of critical points, which in turn create a large number of small connected components in the level surface. Bottom: Simplification of the local extrema (Section 6.8) leads to a function whose contour surfaces are cleaned from topological artifacts, without diminishing the sharpness of the data. Note, however, that small details of the data may get lost in the denoising process.

of critical points among all functions whose *antiderivative* is within a given distance from the antiderivative of the input function. This opens up a number of opportunities for further developments. Can a descriptor analogous to the persistence diagram be constructed? Can we also obtain a stability result with respect to appropriate metrics on the spaces of functions and descriptors?

Topological denoising and robust extraction of critical points is not only interesting for scalar functions, but also for vector fields. Methods using discrete Morse theory have been proposed to analyze the structure of vector fields [58] and to track critical points in time series of discrete vector fields [43]. Can the idea of persistence be generalized in a meaningful way to vector field data? Can the methods presented in this thesis lead to robust simplification algorithms with guarantees, like in the case of scalar functions?

Another interesting point of departure for further developments is the connection between persistence and singularity theory. Smooth Morse functions are the *stable* scalar functions in the sense of singularity theory. Roughly speaking, a *stable map* has the property that varying the map continuously (in the C^k norm for some $k \geq 1$) within a sufficiently small neighborhood also moves the critical points continuously, and in particular will not eliminate them. This notion of stability is different from the one encountered in persistent homology, where the supremum (C^0) norm is relevant instead of the C^k norm. Moreover, while singularity theory considers stability against infinitesimal perturbations, persistent homology provides a quantitative notion of stability. It would be interesting to extend the connections between both theories further beyond the results of this thesis.

Bibliography

- [1] P. K. Agarwal, L. Arge, and K. Yi. [I/O-efficient batched union-find and its applications to terrain analysis](#). In *SCG '06: Proceedings of the twenty-second annual symposium on Computational geometry*, pages 167–176. ACM, 2006.
- [2] D. Attali, M. Glisse, S. Hornus, F. Lazarus, and D. Morozov. [Persistence-sensitive simplification of functions on surfaces in linear time](#). Preprint, 2008.
- [3] T. Banchoff. [Critical points and curvature for embedded polyhedra](#). *Journal of Differential Geometry*, 1(3-4):245–256, 1967.
- [4] A. Banyaga and D. Hurtubise. *Lectures on Morse Homology*, volume 29 of *Texts in the Mathematical Sciences*. Springer, 1 edition, 2005.
- [5] U. Bauer, C. B. Schönlieb, and M. Wardetzky. [Total variation meets topological persistence: A first encounter](#). *AIP Conference Proceedings*, 1281(1): 1022–1026, 2010.
- [6] U. Bauer, C. Lange, and M. Wardetzky. [Optimal topological simplification of discrete functions on surfaces](#). *Discrete and Computational Geometry*, 2011.
- [7] R. H. Bing. Some aspects of the topology of 3-manifolds related to the poincaré conjecture. In T. L. Saaty, editor, *Lectures on Modern Math*, volume 2, pages 93–128. Wiley, New York, 1964.
- [8] W. D. Blizard. [The development of multiset theory](#). *Modern Logic*, 184(4): 319–352, 1991.
- [9] M. Brown. [A proof of the generalized Schoenflies theorem](#). *Bulletin of the American Mathematical Society*, 66(2):74–76, 1960.
- [10] J. Bryant. *Piecewise Linear Topology*, pages 219–259. Elsevier, 2001.

- [11] S. S. Cairns. [An elementary proof of the Jordan-Schoenflies theorem](#). *Proceedings of the American Mathematical Society*, 2:860–867, 1951.
- [12] J. W. Cannon. [Shrinking cell-like decompositions of manifolds. Codimension three](#). *The Annals of Mathematics*, 110(1):83–112, 1979.
- [13] G. Carlsson and V. de Silva. [Zigzag persistence](#). *Foundations of Computational Mathematics*, 10(4):367–405, 2010.
- [14] J. Cerf. [Sur les difféomorphismes de la sphère de dimension trois](#), volume 53 of *Lecture Notes in Mathematics*. Springer, 1968.
- [15] M. K. Chari. [On discrete Morse functions and combinatorial decompositions](#). *Discrete Mathematics*, 217(1-3):101–113, 2000.
- [16] M. M. Cohen. *A Course in Simple-Homotopy Theory*, volume 10 of *Graduate Texts in Mathematics*. Springer, 1973.
- [17] D. Cohen-Steiner, H. Edelsbrunner, and D. Morozov. [Vines and vineyards by updating persistence in linear time](#). In *SCG '06: Proceedings of the twenty-second annual symposium on Computational geometry*, pages 119–126. ACM, 2006.
- [18] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. [Stability of persistence diagrams](#). *Discrete and Computational Geometry*, 37(1):103–120, 2007.
- [19] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. [Extending persistence using Poincaré and Lefschetz duality](#). *Foundations of Computational Mathematics*, 9(1):79–103, 2008.
- [20] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, third edition, 2009.
- [21] A. Danner, T. Mølhave, K. Yi, P. K. Agarwal, L. Arge, and H. Mitasova. [TerraStream: from elevation data to watershed hierarchies](#). In *GIS '07: Proceedings of the 15th annual ACM international symposium on Advances in geographic information systems*, pages 1–8. ACM, 2007.
- [22] V. de Silva, D. Morozov, and M. Vejdemo-Johansson. [Dualities in persistent \(co\)homology](#). Unpublished manuscript, 2010.

- [23] H. Edelsbrunner, D. Letscher, and A. Zomorodian. [Topological persistence and simplification](#). *Discrete and Computational Geometry*, 28(4):511–533, 2002.
- [24] H. Edelsbrunner, J. Harer, and A. Zomorodian. [Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds](#). *Discrete and Computational Geometry*, 30(1):87–107, 2003.
- [25] H. Edelsbrunner, D. Morozov, and V. Pascucci. [Persistence-sensitive simplification functions on 2-manifolds](#). In *SCG '06: Proceedings of the twenty-second annual symposium on Computational geometry*, pages 127–134. ACM, 2006.
- [26] J. Eells and N. Kuiper. [Manifolds which are like projective planes](#). *Publications Mathématiques de L'IHÉS*, 14(1):5–46, 1962.
- [27] O. Egecioglu and T. F. Gonzalez. [A computationally intractable problem on simplicial complexes](#). *Computational Geometry*, 6(2):85–98, 1996.
- [28] R. Forman. [Morse theory for cell complexes](#). *Advances in Mathematics*, 134(1):90–145, 1998.
- [29] R. Forman. [Witten–Morse theory for cell complexes](#). *Topology*, 37(5):945–979, 1998.
- [30] R. Forman. [A user’s guide to discrete Morse theory](#). *Séminaire Lotharingien de Combinatoire*, B48c:1–35, 2002.
- [31] P. Frosini. [A distance for similarity classes of submanifolds of a Euclidean space](#). *Bulletin of the Australian Mathematical Society*, 42(03):407–415, 1990.
- [32] S. Fry. [Defining and sizing-up mountains](#). *Summit*, Jan.–Feb., 1987.
- [33] C. Gray, F. Kammer, M. Löffler, and R. I. Silveira. [Removing local extrema from imprecise terrains](#). Preprint, 2010. [arXiv:1002.2580](#).
- [34] A. Gyulassy, V. Natarajan, V. Pascucci, and B. Hamann. [Efficient computation of Morse-Smale complexes for three-dimensional scalar functions](#). *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1440–1447, 2007.
- [35] A. Hatcher. [Algebraic Topology](#). Cambridge University Press, 2002.

- [36] J. R. Isbell. [Six theorems about injective metric spaces](#). *Commentarii Mathematici Helvetici*, 39:65–76, 1964.
- [37] S. K. Jenson and J. O. Domingue. [Extracting topographic structure from digital elevation data for geographic information system analysis](#). *Photogrammetric Engineering and Remote Sensing*, 54(11):1593–1600, 1988.
- [38] G. Jerše and N. Mramor Kosta. [Ascending and descending regions of a discrete Morse function](#). *Computational Geometry*, 42(6-7):639–651, 2009.
- [39] M. Joswig and M. E. Pfetsch. [Computing optimal Morse matchings](#). *SIAM Journal on Discrete Mathematics*, 20(1):11–25, 2006.
- [40] T. Kaczynski, K. Mischaikow, and M. Mrozek. [Computational Homology](#), volume 157 of *Applied Mathematical Sciences*. Springer, 1st edition, 2004.
- [41] C. Kearton and W. B. R. Lickorish. [Piecewise linear critical levels and collapsing](#). *Transactions of the American Mathematical Society*, 170, 1972.
- [42] H. King, K. Knudson, and N. Mramor. [Generating discrete Morse functions from point data](#). *Experimental Mathematics*, 14(4):435–444, 2005.
- [43] H. King, K. Knudson, and N. Mramor. [Birth and death in discrete morse theory](#). 2008.
- [44] A. Kosinski. [Singularities of piecewise linear mappings. I Mappings into the real line](#). *Bulletin of the American Mathematical Society*, 68(2):110–114, 1962.
- [45] D. Kozlov. [Combinatorial Algebraic Topology](#), volume 21 of *Algorithms and Computation in Mathematics*. Springer, 1 edition, 2007.
- [46] J. B. Kruskal. [On the shortest spanning subtree of a graph and the traveling salesman problem](#). *Proceedings of the American Mathematical Society*, 7(1): 48–50, 1956.
- [47] W. Kühnel. Triangulations of manifolds with few vertices. In F. Triccerri, editor, *Advances in differential geometry and topology*, pages 59–114. World Scientific, Singapore, 1990.

-
- [48] [Large Geometric Models Archive](http://www.cc.gatech.edu/projects/large_models/). Georgia Institute of Technology. Available from: http://www.cc.gatech.edu/projects/large_models/.
- [49] T. Lewiner, H. Lopes, and G. Tavares. [Optimal discrete Morse functions for 2-manifolds](#). *Computational Geometry*, 26(3):221–233, 2003.
- [50] A. T. Lundell and S. Weingram. *The Topology of CW Complexes*. Van Nostrand Reinhold Company, 1969.
- [51] J. Milnor. *Morse Theory*, volume 51 of *Annals of Mathematics Studies*. Princeton University Press, 1963.
- [52] J. Milnor. *Lectures on the h-Cobordism Theorem*. Princeton Mathematical Notes. Princeton University Press, 1st edition, 1965.
- [53] D. Morozov. [Homological Illusions of Persistence and Stability](#). PhD thesis, Duke University, 2008.
- [54] M. Morse. [Functional topology and abstract variational theory](#). *The Annals of Mathematics*, 38(2), 1937.
- [55] M. Morse. [Rank and span in functional topology](#). *The Annals of Mathematics*, 41(2), 1940.
- [56] M. Morse. [A reduction of the Schoenflies extension problem](#). *Bulletin of the American Mathematical Society*, 66(2):113–115, 1960.
- [57] M. Morse. [The elimination of critical points of a non-degenerate function on a differentiable manifold](#). *Journal d'Analyse Mathématique*, 13(1):257–316, 1964.
- [58] J. Reininghaus, C. Lowen, and I. Hotz. [Fast combinatorial vector field topology](#). *IEEE Transactions on Visualization and Computer Graphics*, 99, 2010.
- [59] V. Robins. [Computational Topology at Multiple Resolutions](#). PhD thesis, University of Colorado at Boulder, 2000.
- [60] L. I. Rudin, S. Osher, and E. Fatemi. [Nonlinear total variation based noise removal algorithms](#). In *Proceedings of the eleventh annual international conference of the Center for Nonlinear Studies on Experimental mathematics : computational issues in nonlinear science: computational issues in nonlinear science*, volume 60, pages 259–268. Elsevier North-Holland, Inc., 1992.

- [61] P. Soille. [Morphological carving](#). *Pattern Recognition Letters*, 25(5):543–550, 2004.
- [62] P. Soille. [Optimal removal of spurious pits in grid digital elevation models](#). *Water Resources Research*, 40(12):W12509+, 2004.
- [63] R. Thom. [Sur une partition en cellules associée à une fonction sur une variété](#). *Comptes rendus des séances de l'Académie des Sciences*, 228:973–975, 1949.
- [64] [VolVis medical data sets](#). University of Tübingen. Available from: <http://www.volvis.org/>.
- [65] A. Wächter and L. T. Biegler. [On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming](#). *Mathematical Programming*, 106(1):25–57, 2006.
- [66] E. C. Zeeman. [On the dunce hat](#). *Topology*, 2(4):341–358, 1963.
- [67] A. Zomorodian and G. Carlsson. [Computing persistent homology](#). *Discrete and Computational Geometry*, 33(2):249–274, 2005.

Index

- ascending manifold, [22](#)
- ascending set, [21](#)
- attaching map, [11](#)
- barycentric subdivision, [31](#)
- birth
 - persistence, [43](#)
- cap span, [43](#)
- carving, [77](#)
- cell
 - CW complex, [11](#)
 - essential, [43](#)
 - negative, [43](#)
 - positive, [43](#)
- Cerf theory, [28](#)
- characteristic map, [11](#)
- child, [54](#)
- collapse, [24](#)
- consistent, [18](#)
 - minimal vector field, [27](#)
- consistent total order, [22](#), [68](#)
- convex polyhedron, [72](#)
- covering relation, [21](#)
- critical cell, [12](#)
- critical point, [12](#)
- CW complex, [11](#)
 - regular, [12](#)
 - subcomplex, [12](#)
- death
 - persistence, [43](#)
- derived subdivision, [31](#)
- descendant, [54](#)
- descending manifold, [22](#)
- descending set, [21](#)
- discrete gradient flow, [16](#)
- discrete gradient vector field, [13](#)
 - reversing along a path, [14](#), [38](#)
- discrete Morse function, [17](#)
 - gradient vector field of, [17](#)
 - perfect, [83](#)
- discrete vector field, [12](#), [34](#)
- dual complex, [44](#)
- dunce hat, [83](#)
- elementary collapse, [24](#)
- energy, [81](#)
 - Dirichlet, [81](#)
- extrema simplification, [84](#)
- face, [12](#)
 - of a simplex, [30](#)
- facet, [12](#)
 - regular, [12](#)
 - semi-regular, [34](#)
- flattening, [27](#)
- flooding, [77](#)
- free facet, [24](#)

- graph
 - dual, 45
 - primal, 45
- Hasse diagram, 70
- house with two rooms, 83
- index
 - critical cell, 12
 - critical point, 14
- induced subcomplex, 31
- level subcomplex, 24
- linkable cap, 43
- manifold
 - combinatorial, 12
- metric
 - on persistence diagrams, 47
- minimum spanning tree, 54
- Morse chart, 14
- Morse complex, 15
- Morse function, 14
 - excellent, 28
- Morse homology, 17
- multiset, 46
- order complex, 31
- order subcomplex, 24
- parent, 54
- partial order
 - induced, 20
- perfect δ -simplification, 52, 71
- persistence, 43
 - cancelation sequence, 56
 - δ -, 56
 - nested, 56
 - diagram, 46
 - hierarchy, 54
 - pair, 43
 - local, 44
 - nested, 55
- persistent Betti number, 42
- persistent homology group, 42
- piecewise linear function, 29
- plateau function, 61
- pseudo-Morse function, 18
 - flat, 26
- restriction
 - of a vector field to a subcomplex, 57
- Schoenflies theorem
 - generalized, 34
- simplex, 30
- simplexwise linear, 86
- simplexwise linear function, 29
- simplicial complex, 30
 - abstract, 30
- simplicial scheme, 30
- stability bound, 5, 45
 - tightness, 5
- straight-line homotopies, 28
- strong deformation retract, 23
- strong deformation retraction, 23
- subcomplex, 30
- sublevel set, 23
- surface
 - combinatorial, 12
 - with boundary, 45
- suspension, 84
- symbolic perturbation, 20

tolerance intervals, [86](#)
topological sorting, [70](#)

underlying space, [30](#)

 V -path, [13](#)
vector field
 gradient-like, [14](#)
vertex
 simplex, [30](#)
 simplicial complex, [30](#)
vertex scheme, [30](#)

Curriculum Vitæ

Personal

Name	Ulrich Alexander Bauer
Date of birth	March 7, 1980
Place of birth	Munich
Citizenship	German
Place of residence	Berlin

Academic

since 08/2008	Research Assistant Georg-August-Universität Göttingen (Discrete Differential Geometry, Prof. Max Wardetzky)
05/2006–07/2008	Research Assistant Freie Universität Berlin (Mathematical Geometry Processing, Prof. Konrad Polthier)
2005	Diploma (MSc equivalent) Computer Science (minor in Mathematics) TU München Grade: 1.1 (with distinction) Thesis: Assignment Problem with Constraints Advisor: Prof. Angelika Steger, ETH Zürich
1999	Abitur (grade: 1.4) Rhabanus-Maurus-Gymnasium St. Ottilien