Morse sequences and image segmentation

Kévin Feltrin[1*]

[1*]ESIEE Paris, Université Gustave Eiffel, 2 boulevard Blaise Pascal, Noisy-le-Grand, 93162, Île-de-France, France.

Corresponding author(s). E-mail(s): kevin.feltrin@edu.esiee.fr;

Abstract

Discrete Morse theory provides a combinatorial analogue of gradient flow that supports topology-aware image segmentation. A discrete gradient field, defined as an acyclic pairing on the cells of a cubical complex, induces stable and unstable sets of critical cells, from which one obtains basins, separatrix structures, watershed-like objects, and skeletons.

For grayscale images, the classical pipeline of Robins–Wood–Sheppard constructs such a gradient from vertex intensities by a monotone max-extension and local lower-star processing, while Delgado-Friedrichs–Robins–Sheppard derive partitions and skeletons from the induced stable/unstable dynamics. Morse sequences provide an alternative constructive viewpoint in which acyclicity is guaranteed by design and where reference/coreference and extension/coextension maps explicitly transfer critical information between the ambient complex and the critical complex.

This report translates the segmentation layer of Delgado-Friedrichs–Robins–Sheppard into the framework of Morse sequences, with complete proofs. Stable and unstable sets, defined by existence of gradient paths, are shown to coincide with the recursive sets of that work, while the parity-based extension and coextension maps remain strictly smaller as soon as paths branch; at dimension zero the two notions merge and the partition of vertices by minima follows. Prefixes of a sequence and cuts of a stack are hermetic, which yields a segmentation at every grayscale level from a single computation, with no injectivity assumption on the data. Finally, the basin of a minimum is characterized as the unique maximal closure under regular expansions, independently of the order of the expansions; the order carried by the sequence then labels all basins in one scan, in time linear in the incidence structure.

Keywords: Discrete Morse theory, Morse sequences, Discrete gradient, $\boldsymbol{V}$-paths, Stable/unstable sets, Image segmentation, Watersheds, Reference/coreference, Extension/coextension

1 Introduction and positioning

Image segmentation aims at partitioning a discrete domain (pixels/voxels) into meaningful regions. More precisely, following [1]:

Definition 1 (Image segmentation as a partitioning) Let Ω denote a discrete image domain. A *segmentation* of Ω is a partition $\{R_1, R_2, \ldots, R_n\}$ such that:

1. $\bigcup_{i=1}^{n} R_i = \Omega$ (completeness),
2. $R_i \cap R_j = \emptyset$ for $i \neq j$ (disjointness),
3. each R_i is a connected region satisfying a homogeneity predicate.

In many applications, accuracy is not only geometric (boundary location), but also *topological*: regions should remain connected, spurious holes should be avoided or controlled, and one often seeks a *skeleton* that summarizes shape while preserving the homotopy type of the object. These requirements arise for instance in biomedical imaging (organs, vessels), material science (porous media), or document analysis (strokes and medial axes), where connectivity and tunnel/cavity artifacts can invalidate downstream measurements.

A natural way to enforce topological structure is to rely on a *gradient flow* that organizes the domain into attraction regions (basins) separated by ridge-like structures. Discrete Morse theory, introduced by Forman [2, 3], provides a rigorous, purely combinatorial framework for this idea: it replaces the smooth gradient flow of a function by an *acyclic pairing* on the cells of a complex, called a discrete gradient vector field. Critical cells encode topological events, while regular paired cells can be removed by elementary collapses without changing homotopy type [4, 5]. For grayscale images, the classical algorithmic pipeline is due to Robins–Wood–Sheppard (RWS), who construct a discrete gradient field from vertex intensities using local *lower-star* processing [6]. The lower-star approach itself is inspired by the classical smooth Morse theory notion of descending link and was adapted to the discrete setting by Forman [2]; the specific efficient algorithm used by RWS builds upon this idea combined with the max-extension F (see Section 4.2). Delgado-Friedrichs–Robins–Sheppard (DFRS) then derive partitions and skeletons from the *stable* and *unstable* sets of critical cells and propose simplification mechanisms [7]. In parallel, grayscale topological operators emphasize plateau-consistent behavior across threshold cross-sections [8], and watershed structures have been characterized within a discrete Morse framework via collapses [9]. The Morse–Smale complex, which partitions a manifold according to the combined ascending and descending dynamics, has also been studied in the piecewise-linear and discrete settings [10]; these partitions are closely related to the stable/unstable decompositions considered in the present report. More recently, *Morse sequences* were introduced as a constructive representation of discrete Morse theory, producing acyclic gradients by design and providing explicit algebraic maps (reference/coreference and extension/coextension) [11].

1.1 Problem: topological segmentation guided by a discrete gradient

We model the image domain by a finite cubical complex X. In dimension two, X contains vertices, edges, and squares; in dimension three, it also contains square faces and cubes. When relevant, a binary object is represented by a subcomplex $K \subseteq X$. Grayscale data are given by a vertex map

$$f\colon V(X) \longrightarrow \mathbb{Z} \qquad \text{(or, more generally, } \mathbb{R}\text{)}.$$

This function is used to build and, when needed, simplify the discrete gradient. Once the gradient is fixed, it is not part of the segmentation objects themselves.

The core segmentation mechanism is the following:

- Construct an acyclic discrete vector field (pairing) V on X (or on K).
- Identify critical cells (unpaired cells), which encode the essential topological structure.
- Define *stable* and *unstable* sets of critical cells using V-paths (discrete gradient flow).
- Use these sets to build basins (regions), ridge/separator sets (watershed-like structures), and skeletons [7].

This viewpoint isolates what must be topologically controlled: *the segmentation objects are derived from the discrete gradient field, not directly from the grayscale function.* As a consequence, understanding precisely (i) how f is used to construct V, and (ii) how stable/unstable sets derive segmentation from V, is the key to reading and improving the classical pipelines.

1.2 Contributions and guiding thread

This work is meant to be readable by a non-specialist reader familiar with basic discrete geometry. Accordingly, we provide explicit definitions for every notion that is minimally needed by the pipelines.

(A) We make the segmentation layer central.

Following DFRS [7], we treat *stable and unstable sets* of critical cells as the main mathematical primitives for segmentation. Basins, separators, and skeleton pieces are then defined as unions (or selections) of these sets, with a dimension-aware interpretation in 2D vs 3D.

(B) We pinpoint where grayscale enters the RWS construction.

RWS starts from vertex intensities f but does not directly define a discrete Morse function on all cells. Instead, it extends f monotonically to cells (a filtration viewpoint), defines local lower-star blocks, and constructs an acyclic pairing by local processing [6]. We explain these steps at the right granularity for a research report: enough to understand *which ingredients depend on* f, without drowning the segmentation narrative in implementation heuristics.

(C) Toward an F-driven, stack-compatible reformulation.

The grayscale function f is initially defined only on vertices, whereas a Morse sequence acts on all cells of the complex. We therefore distinguish f from its monotone cell-wise extension F. The proposed viewpoint is not a new independent theory of Morse sequences: it is an image-oriented reformulation of Bertrand's Morse sequences on stacks and flooding sequences [12]. Its purpose is to make explicit how the RWS/DFRS grayscale construction can be read in a sequence-based language:

$$\begin{aligned} f \text{ on vertices} &\longrightarrow F \text{ on cells} \longrightarrow \text{level-compatible Morse sequence} \\ &\longrightarrow V \longrightarrow \text{segmentation objects.} \end{aligned}$$

This formulation is especially useful for clarifying the role of plateaus, since equal values of F should be processed within the same stack level rather than hidden inside an arbitrary implementation-dependent tie-break.

Guiding principle.

Throughout the report, we adopt the following separation:

$$\text{(data: } f) \longrightarrow \text{(gradient: } V) \longrightarrow \text{(segmentation objects: St/Un).}$$

The grayscale function drives construction and simplification; the segmentation objects depend only on the resulting gradient field.

Figure 1 summarizes the global pipeline that will be referenced repeatedly.

Part I – Context

2 Discrete model and minimal topological notions

This section fixes the discrete setting used throughout the report and recalls the minimal algebraic-topological notions needed later (RWS/DFRS and Morse sequences), without requiring the reader to consult external textbooks. All homological statements are used only as guiding principles: "removing regular pairs preserves topology" in a precise sense (homotopy, hence homology) [4, 5].

2.1 Complexes for images: cubical complexes and subcomplexes

From a pixel/voxel grid to a cell complex.

A 2D image lives on a rectangular grid. A convenient topological model is a finite *cubical complex*, which we now define precisely.

Definition 2 (Elementary interval) An *elementary interval* is a subset of $\mathbb{R}$ of the form $[k, k+1]$ (a *non-degenerate* elementary interval) or $[k, k]$ (a *degenerate* elementary interval), for some $k \in \mathbb{Z}$.

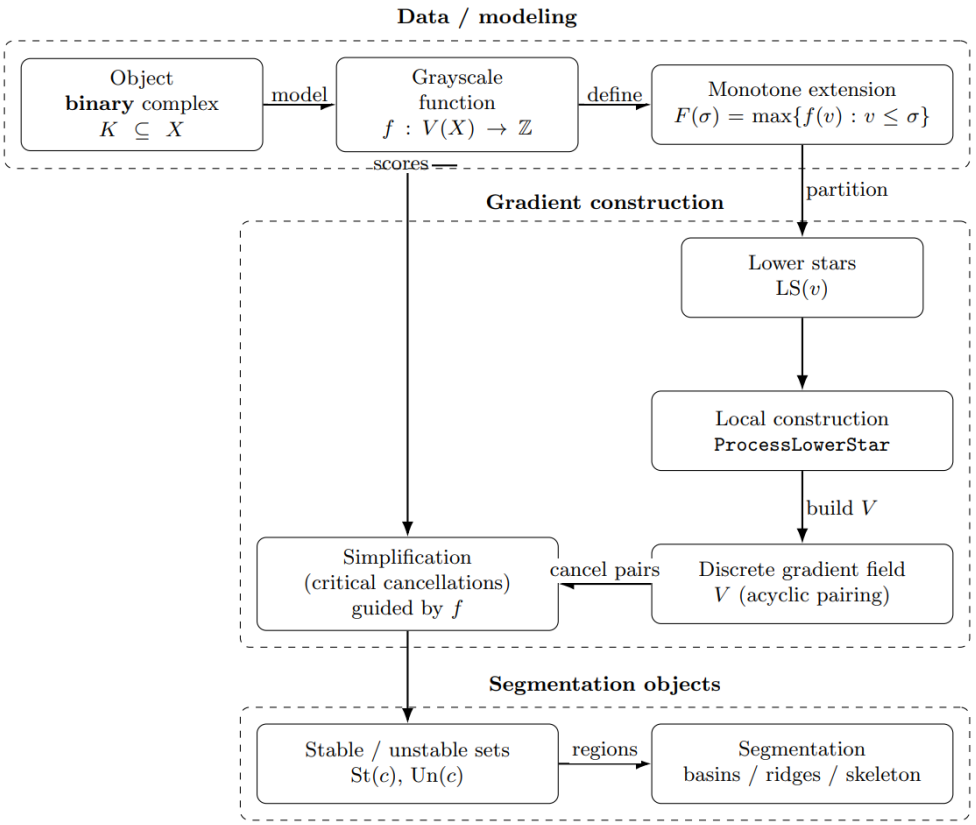


Fig. 1 Global pipeline. The grayscale function f (vertex intensities) defines a monotone extension F used to partition the complex into lower-star blocks and to construct an acyclic pairing V (discrete gradient field) [6]. Optionally, f also provides scores to simplify V via controlled critical cancellations [5, 7]. Stable and unstable sets of critical cells then define basins, separators, and skeleton structures [7].

Definition 3 (Elementary cube and cubical complex) An *elementary cube* in $\mathbb{R}^d$ is a product $\sigma = I_1 \times I_2 \times \cdots \times I_d$ where each I_j is an elementary interval. The *dimension* of σ is the number of non-degenerate factors. A *cubical complex* $X \subseteq \mathbb{R}^d$ is a finite union of elementary cubes that is *closed under taking faces*: if $\sigma \in X$ and τ is a face of σ (obtained by replacing one or more non-degenerate factors by one of their endpoints), then $\tau \in X$ [13].

In the 2D image setting, this yields:

- 0-cells (vertices) model grid points,
- 1-cells (edges) model unit segments between adjacent grid points,
- 2-cells (squares) model pixels (unit squares).

In 3D, one also includes 3-cells (cubes) modeling voxels, together with their faces (squares), edges, and vertices. Cubical complexes are the default model for images in RWS/DFRS [6, 7].

Cells and dimension.

A *cell* σ is one of the building blocks (vertex/edge/square/cube). We denote by $\dim(\sigma)$ its dimension, and by $X^{(p)}$ the set of all p-cells of X.

Face relation.

Given two cells $\alpha, \beta \in X$, we write $\alpha < \beta$ if α is a face of β of codimension 1 (e.g. a vertex is a face of an incident edge; an edge is a face of an incident square). This relation encodes incidence information and will be used later to define boundaries, pairings, and discrete flows.

Remark 1 (Why cubical complexes?) Cubical complexes match the combinatorics of images (pixels/voxels) and yield clean adjacency relations. They also support efficient implementations of gradient-based constructions and stable/unstable set computations [6, 7].

Binary object as a subcomplex.

Many segmentation tasks distinguish an *object* from background (e.g. by thresholding). Topologically, a binary mask is naturally represented as a *subcomplex* $K \subseteq X$:

Definition 4 (Subcomplex) A subset $K \subseteq X$ is a subcomplex if whenever $\sigma \in K$ and $\alpha \leq \sigma$ (i.e. α is a face of σ), then $\alpha \in K$.

In words: if a pixel/voxel (a square/cube) belongs to the object, then all its faces (edges/vertices) belong as well. This closure property is essential to ensure that K is itself a valid complex on which homology and collapses make sense.

Notation summary.

Throughout the report, we use

$$X \text{ for the ambient image complex,}$$
$$K \subseteq X \text{ for an optional binary-object subcomplex,}$$
$$X^{(p)} \text{ for the set of } p\text{-cells of } X.$$

2.2 Chains and boundary over $\mathbb{Z}_2$

The role of homology in this work is limited and operational: it provides a language for "topological features" (connected components, holes, cavities), and it justifies why certain local removals preserve these features [4].

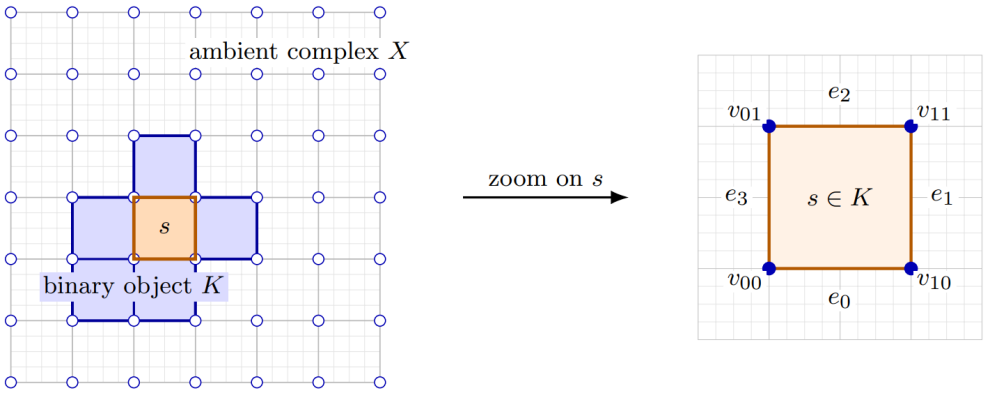


Fig. 2 Cubical model used throughout the report. Left: the image domain is an ambient cubical complex X, and a binary object is represented by a subcomplex $K \subseteq X$. Right: zoom on one selected square s of K. The boxed statement makes explicit the closure-under-faces property that characterizes a subcomplex.

Why coefficients in $\mathbb{Z}_2$?

Working over $\mathbb{Z}_2$ (coefficients mod 2) avoids orientation bookkeeping and is common in image applications: each cell is either present (coefficient 1) or absent (coefficient 0). This is sufficient for the invariance statements used later.

Chain groups

Definition 5 (p-chains over $\mathbb{Z}_2$) For a finite cell complex X, the group of p-chains is

$$C_p(X) := \Big\{ \sum_{\sigma \in X^{(p)}} a_\sigma \, \sigma \;:\; a_\sigma \in \mathbb{Z}_2 \Big\},$$

a vector space over $\mathbb{Z}_2$ with basis the p-cells.

Thus a 1-chain is a formal sum of edges (mod 2), and a 2-chain is a formal sum of squares (mod 2).

Boundary operator

Definition 6 (Boundary map) The boundary map $\partial_p : C_p(X) \to C_{p-1}(X)$ is defined on a p-cell σ by

$$\partial_p(\sigma) := \sum_{\tau < \sigma} \tau \quad \text{in } \mathbb{Z}_2, \tag{1}$$

and extended linearly to all chains.

Because coefficients are mod 2, a face that appears twice cancels automatically. A fundamental property is

$$\partial_{p-1} \circ \partial_p = 0,$$

which means "the boundary of a boundary is empty" (mod 2) [4].

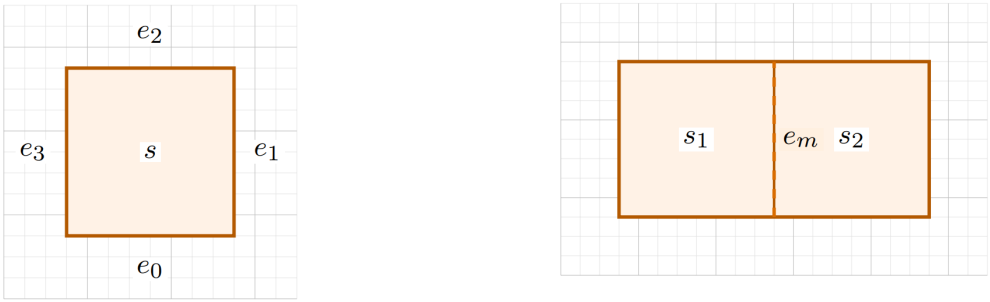


Fig. 3 Boundary over $\mathbb{Z}_2$. Left: the boundary of one square is the sum of its four edges. Right: for a 2-chain made of two adjacent squares, the common edge appears twice and therefore cancels modulo 2. This is the geometric content behind the identity $\partial \circ \partial = 0$.

Cycles, boundaries, and homology

Definition 7 (Cycles and boundaries) Define

$$Z_p(X) := \ker(\partial_p) \quad \text{(cycles)}, \qquad B_p(X) := \operatorname{im}(\partial_{p+1}) \quad \text{(boundaries)}.$$

Definition 8 (Homology groups (minimal use)) The p-th homology group is

$$H_p(X) := Z_p(X)/B_p(X),$$

and its dimension $\beta_p(X) := \dim H_p(X)$ (over $\mathbb{Z}_2$) is the p-th Betti number.

For intuition in 2D:

$$\beta_0 = \#\text{ connected components}, \qquad \beta_1 = \#\text{ holes}.$$

In 3D, β_2 counts void-like cavities.

2.3 Collapses (free pairs) and topological invariance

The next notion is the discrete topological mechanism behind "removing regular pairs" in discrete Morse theory. It formalizes why certain local deletions preserve the global topology (homotopy type, hence homology) [4, 5].

Free pairs and elementary collapse

Definition 9 (Free pair) Let σ be a p-cell and τ be a $(p+1)$-cell such that $\sigma < \tau$. The pair (σ, τ) is a *free pair* if τ is the *only* $(p+1)$-cell of X having σ as a face. Equivalently, σ is a codimension-1 face of exactly one coface.

Definition 10 (Elementary collapse) If (σ, τ) is a free pair, removing both cells produces a new complex

$$X' := X \setminus \{\sigma, \tau\},$$

and the move $X \searrow X'$ is called an *elementary collapse*.

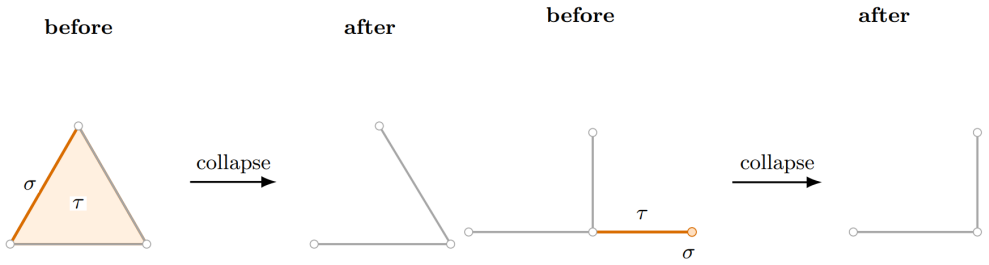


Fig. 4 Elementary collapse = deletion of a free face σ together with its unique coface τ. The removable pair is highlighted in orange; the gray part is the portion of the complex that remains after the collapse.

Definition 11 (Collapse) A (finite) sequence of elementary collapses is called a *collapse*.

Interpretation.

A free pair is a "pendant" configuration: σ attaches to the rest of the complex only through τ. Removing (σ, τ) cannot create or destroy holes in an essential way; it simply prunes redundant cells.

Theorem 1 (Topological invariance of collapses (informal statement)) *If X collapses to X', then X and X' have the same homotopy type (hence the same homology).*

This is classical in combinatorial topology [4]. In discrete Morse theory, regular pairs are designed to behave like collapsible configurations; this is why eliminating them preserves topology while critical cells encode the remaining structure [5].

Remark 2 (Message to keep in mind for later sections) In RWS/DFRS, the output of the grayscale-driven construction is an acyclic pairing (a discrete gradient field). Paired cells are *regular*. The guiding principle is:

eliminating regular pairs (via collapse-like moves) preserves topology.

Critical cells are the topologically meaningful ones.

Two canonical pictures

Figure 4 shows two simple collapse configurations. Panel (a) illustrates a free edge on a pendant triangle (simplicial picture, for intuition); panel (b) illustrates a pendant edge ending at a leaf vertex (which also occurs in cubical settings). The definition is the same in both cases: a face that belongs to exactly one higher-dimensional cell is free.

A minimal algorithmic viewpoint

In practice, "free" can be tested combinatorially by counting cofaces. Algorithm 1 gives a minimal procedure that removes a free pair if it exists. This is not an algorithm

Algorithm 1 Free-pair test and elementary collapse (minimal)

Require: Complex X, candidate face σ of dimension p
Ensure: Either perform an elementary collapse, or report "not free"

```
Compute the set of cofaces C(σ) := {τ ∈ X^(p+1) : σ < τ}
if |C(σ)| = 1 then
    Let τ be the unique element of C(σ)
    Remove σ and τ from X (set X ← X \ {σ, τ})
    return collapsed with free pair (σ, τ)
else
    return not free
end if
```

for discrete Morse theory; it is only meant to lock in the basic mechanism that later justifies eliminating regular pairs.

Takeaway.

Discrete Morse constructions can be seen as producing, in a controlled and acyclic way, a collection of regular pairs that behave like collapses. This is the conceptual bridge to later sections: once a discrete gradient field is built, stable/unstable sets organize the domain into regions and separators, while simplification corresponds to removing low-significance critical structures in a topology-aware manner [5, 7].

2.4 Elementary operations on complexes: expansion, filling, and perforation

The previous subsection introduced elementary collapses. We now record the three complementary elementary operations used in the same topological framework [11]. At this stage, the discussion is purely topological: we only define how a complex is locally modified, without introducing any sequence construction.

Expansion

Definition 12 (Elementary expansion) Let $L \subseteq K$ be two finite cell complexes. We say that K is an *elementary expansion* of L if there exist two cells σ and τ such that

$$\sigma < \tau, \qquad \dim(\tau) = \dim(\sigma) + 1,$$
$$K = L \cup \{\sigma, \tau\}, \qquad L = K \setminus \{\sigma, \tau\},$$

and (σ, τ) is a free pair in K.

Equivalently, an elementary expansion is the inverse operation of an elementary collapse: instead of deleting a free pair, one adds a lower-dimensional cell σ together with its unique coface τ.

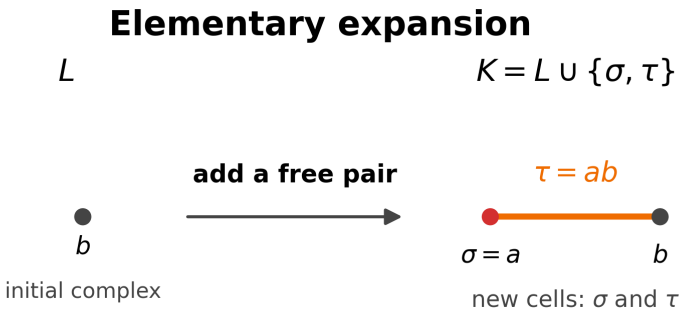


Fig. 5 Elementary expansion. Starting from a complex L, one obtains K by adding two cells σ and τ with $\sigma < \tau$, so that (σ, τ) is a free pair in the enlarged complex K. The inverse operation is the elementary collapse deleting the same pair.

Filling

Definition 13 (Elementary filling) Let $L \subseteq K$ be two finite cell complexes. We say that K is an *elementary filling* of L if there exists a single cell ν such that

$$K = L \cup \{\nu\}, \qquad L = K \setminus \{\nu\},$$

and ν is a facet of K.

Thus a filling adds one maximal cell whose boundary is already present. In dimension 2, a typical example is the addition of a square to an existing boundary cycle.

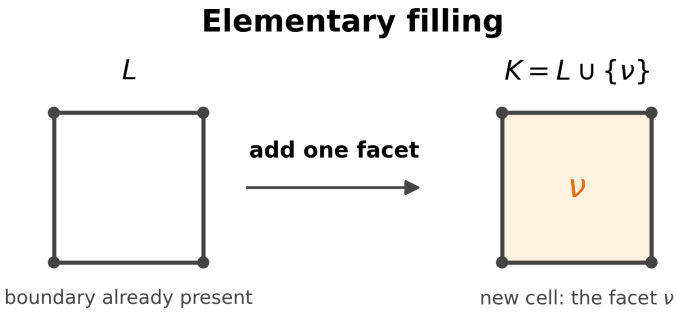


Fig. 6 Elementary filling. The operation adds one facet ν to the complex. In the two-dimensional picture, the boundary is already present and the operation inserts the square ν.

Perforation

Definition 14 (Elementary perforation) Let K be a finite cell complex and let ν be a facet of K. The complex

$$L = K \setminus \{\nu\}$$

is called an *elementary perforation* of K.

An elementary perforation is the inverse operation of an elementary filling: it removes one maximal cell while keeping its boundary cells. In dimension 2, removing a square but keeping its four boundary edges creates a local opening.

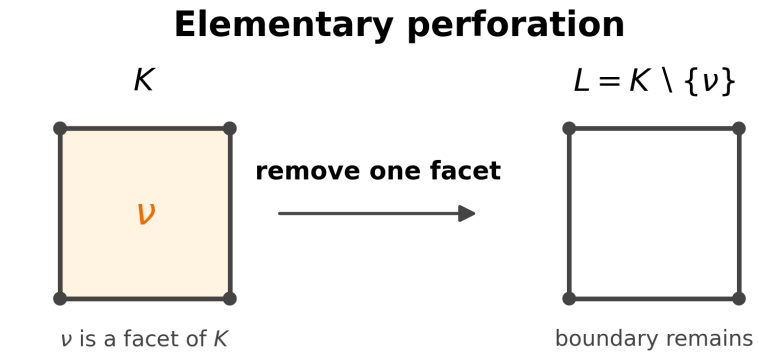


Fig. 7 Elementary perforation. The operation removes one facet ν from K and keeps the remaining cells. It is the inverse operation of the elementary filling shown in Figure 6.

3 Discrete gradients and central objects for segmentation

This section introduces the discrete-gradient objects that will be used as the core segmentation machinery. Following the discrete Morse viewpoint, a segmentation is obtained from an acyclic pairing (a discrete vector field), together with the induced dynamics that define *stable* and *unstable* sets. These objects are the bridge between the grayscale-driven construction (RWS/DFRS) and the final products (basins, ridges/separatrices, watershed sets, and skeletons) [5–7].

3.1 Discrete vector fields and critical cells

We now introduce the combinatorial object that replaces the smooth gradient in discrete Morse theory.

Definition 15 (Pairing / discrete vector field [2, 3]) Let X be a finite cell complex. A *pairing* on X is a set V of pairs (σ, τ) such that

- $\sigma < \tau$,
- $\dim(\tau) = \dim(\sigma) + 1$,
- each cell of X appears in at most one pair of V.

Such a pairing is also called a *discrete vector field*.

In other words, a pairing matches certain cells with one incident coface of dimension one higher. A cell that belongs to some pair of V is called *regular* (or *paired*), whereas a cell that does not belong to any pair is called *critical*.

Definition 16 (Critical cell [2]) A cell $c \in X$ is *critical* for V if it does not appear in any pair of V.

Interpretation.
A pair $(\sigma, \tau) \in V$ should be understood as an elementary gradient step from σ to τ. The collection of all such pairs is the discrete analogue of a gradient field. Critical cells are the cells where the gradient cannot continue and therefore encode the essential topological structure.

Remark 3 (Dimension conventions) In a 2D cubical complex, critical 0-cells (vertices) are minima candidates, critical 2-cells (squares) are maxima candidates, and critical 1-cells (edges) play the role of saddles. In 3D, saddles may occur in dimensions 1 and 2.

3.2 V-paths and acyclicity (discrete gradients)

The pairing V induces an alternating notion of path that encodes the "flow".

Definition 17 (V-path) Let V be a discrete vector field on X. A *V-path* of index p is a sequence

$$\sigma_0^{(p)} < \tau_0^{(p+1)} > \sigma_1^{(p)} < \tau_1^{(p+1)} > \cdots < \tau_{k-1}^{(p+1)} > \sigma_k^{(p)}$$

such that for each i:

- $(\sigma_i, \tau_i) \in V$ (the step uses a paired coface),
- $\sigma_{i+1} < \tau_i$ and $\sigma_{i+1} \neq \sigma_i$ (the path exits τ_i through another p-face).

Definition 18 (Closed V-path and acyclicity) A V-path is *closed* if $\sigma_k = \sigma_0$ for some $k \geq 1$. A discrete vector field V is *acyclic* if it admits *no* closed V-path. An acyclic discrete vector field is called a *discrete gradient field*.

Why acyclicity matters.

Acyclicity is the combinatorial analogue of "no directed cycles" in a gradient flow: it guarantees that the induced dynamics is well-founded (cells flow toward critical cells), enabling definitions of stable and unstable sets and algorithmic traversals. This is the structural property ensured by the grayscale-driven constructions in RWS/DFRS and by Morse-sequence-based formulations [6, 7, 11].

Algorithm 2 Build the transition graph induced by V (for flow traversals)

Require: Complex X, discrete vector field V
Ensure: Directed graph $G = (\mathcal{S}, \mathcal{E})$ encoding possible one-step moves

```
S ← set of all cells of X
E ← ∅
for all pairs (σ^(p), τ^(p+1)) ∈ V do
    for all p-faces σ' < τ with σ' ≠ σ do
        add directed edge (σ → σ') to E
    end for
end for
return G = (S, E)
```

Comment.

Algorithm 2 produces a purely combinatorial directed graph on p-cells (for each p), encoding how a V-path can evolve. This explicit graph viewpoint is convenient for computing stable/unstable sets via graph search (BFS/DFS) in the next subsection.

3.3 Stable and unstable sets

Once an acyclic discrete vector field V is fixed, the central segmentation objects are defined purely from V. In particular, the grayscale function is no longer needed at this stage.

The notions of stable and unstable sets in the discrete Morse setting are the combinatorial analogues of the descending and ascending manifolds of critical points in smooth Morse theory. In the piecewise-linear and discrete settings, these sets have been studied notably by Edelsbrunner–Harer–Zomorodian [10] in the context of Morse–Smale complexes, and applied to image segmentation by Robins–Wood–Sheppard [6] and Delgado-Friedrichs–Robins–Sheppard [7]. The discrete formulation used here follows [2, 5, 7].

We now define stable and unstable sets from the V-path dynamics introduced in Section 3.2.

Definition 19 (Flow to a critical cell) Let c be a critical p-cell. A p-cell σ is said to *flow to* c if there exists a V-path

$$\sigma = \sigma_0 < \tau_0 > \sigma_1 < \cdots > \sigma_k = c.$$

Definition 20 (Stable set [7]) Let c be a critical p-cell. The *stable set* of c is

$$\mathrm{St}(c) := \{\sigma \in X^{(p)} : \sigma \text{ flows to } c\}.$$

Definition 21 (Unstable set [7]) Let c be a critical p-cell. The *unstable set* of c is the set of p-cells that can be reached from c by following the reversed V-path dynamics. Equivalently,

$$\mathrm{Un}(c) := \{\sigma \in X^{(p)} : c \text{ flows to } \sigma \text{ in the reversed dynamics}\}.$$

Key point.

The definitions of $\mathrm{St}(c)$ and $\mathrm{Un}(c)$ depend only on the discrete gradient field V. The grayscale function is used only earlier, to construct V.

Interpretation.

Stable sets collect the cells that are "attracted" to a given critical cell under the discrete gradient dynamics. Unstable sets collect the cells that can be "reached from" a critical cell and play the role of ascending manifolds. In segmentation, stable sets of minima are the canonical *basins*.

Algorithm 3 Compute $\mathrm{St}(c)$ by BFS/DFS on the induced dynamics

Require: Directed graph $G_p = (\mathcal{S}, \mathcal{E})$ on p-cells, critical cell $c \in \mathcal{S}$
Ensure: Stable set $\mathrm{St}(c)$

```
Build the reversed graph G_p^rev (reverse every directed edge)
Initialize queue/stack Q ← [c], visited set M ← {c}
while Q not empty do
    pop a node u from Q
    for all incoming neighbors v of u in G_p (equivalently: outgoing in G_p^rev) do
        if v ∉ M then
            add v to M and push v into Q
        end if
    end for
end while
return M                                                    ▷ M = St(c)
```

Dual computation.

To compute $\mathrm{Un}(c)$, run the same procedure on G_p starting from c (no reversal), or equivalently compute $\mathrm{St}(c)$ on G_p^{rev}.

Remark 4 (Stable/unstable sets are defined from the gradient only) A crucial point for the present work is that the definitions of stable and unstable sets do not require the grayscale function once the discrete gradient field has been constructed. This is exactly why these objects form the central segmentation layer: the grayscale function is only a construction device, whereas stable/unstable sets are intrinsic to the gradient itself. This observation is also what makes a comparison with Bertrand's extension/coextension maps meaningful.

3.4 Basins, ridges, watershed sets, and skeletons: final products

The previous objects are now assembled into the segmentation outputs used throughout the report.

Basins

A standard convention in grayscale segmentation is:

- *Basins* are stable sets of critical *minima* (critical 0-cells) for a descending dynamics.

Concretely, if $\mathcal{M}$ denotes the set of critical minima, then each $c \in \mathcal{M}$ defines a basin $\mathrm{St}(c)$, and the collection $\{\mathrm{St}(c)\}_{c\in\mathcal{M}}$ forms a basin decomposition up to shared boundary regions (see below).

Ridges, separatrices, and watershed sets

Separators arise where the dynamics is not uniquely attracted to a single minimum. Operationally, ridge-like sets can be captured by cells that belong to multiple basins or that are reached from saddle-type critical cells (dimension 1 in 2D, dimensions 1 or 2 in 3D). A common discrete notion of *watershed set* is:

$$\mathrm{WS} := X \setminus \biguplus_{c\in\mathcal{M}} \mathrm{St}(c),$$

i.e. the complement of the (disjoint) interior of basins, interpreted as a ridge/separatrix set.

Skeletons

Skeletons are obtained by selecting unstable sets of critical cells of suitable dimensions. In two dimensions, a common choice is the union of the unstable sets of saddle-type critical 1-cells, possibly after retaining only the saddles that connect significant extrema. In three dimensions, one may combine curve-like structures arising from critical 1-cells with sheet-like separators associated with critical 2-cells.

4 Where grayscale enters: the classical RWS pipeline and DFRS usage

The goal of this section is to answer precisely *where* and *how* grayscale information is injected in the topological pipeline, without drowning the reader in implementation details. The key point is a fundamental separation: *topology lives in the complex*, while *grayscale lives in an ordering (a function)* that drives a filtration and a local construction of an acyclic pairing [6, 7].

4.1 Fundamental split: binary object vs. grayscale function

Binary object.

A binary segmentation target is naturally modeled as a subcomplex $K \subseteq X$ (Section 2.1): it encodes *which* cells belong to the object and is purely combinatorial/topological.

Grayscale image.

A grayscale image assigns an intensity to grid points (vertices). In the cubical setting, we assume a function

$$f : X^{(0)} \to \mathbb{Z},$$

where $X^{(0)}$ denotes the set of vertices of X. In practice, $f(v)$ is the pixel/voxel intensity sampled at vertex v (or derived from pixel values by a chosen convention).

What grayscale is used for.

The role of f is *not* to change the complex X, but to induce an *order* (a filtration) and drive the construction of a discrete gradient field V:

$$f \implies \text{filtration/order} \implies \text{local pairings} \implies V.$$

Once V is built, all segmentation products are computed from stable/unstable sets (Sections 3.3–3.4), independently of how V was obtained.

4.2 Monotone extension F and filtration

The grayscale data are initially given only on vertices, through a function

$$f : X^{(0)} \to \mathbb{Z}.$$

To use these values on higher-dimensional cells, one defines an extension F to all cells of the complex.

Definition 22 (Max-extension) For every cell $\sigma \in X$, define

$$F(\sigma) := \max\{f(v) : v \in X^{(0)},\ v \leq \sigma\}. \tag{2}$$

Thus, $F(\sigma)$ is the maximum grayscale value among the vertices of σ. For example, if σ is a square, then $F(\sigma)$ is the maximum of the four vertex values of that square.

Proposition 2 (Monotonicity of F) *If $\alpha \leq \tau$, then*

$$F(\alpha) \leq F(\tau).$$

Proof Every vertex of α is also a vertex of τ. Therefore the maximum over the vertices of α cannot exceed the maximum over the vertices of τ. □

Interpretation.

The function F is therefore monotone with respect to the face relation. In the terminology used by Bertrand, such a monotone function is a *stack*. This monotonicity is fundamental because it makes the filtration by sublevel sets compatible with the topology of the complex.

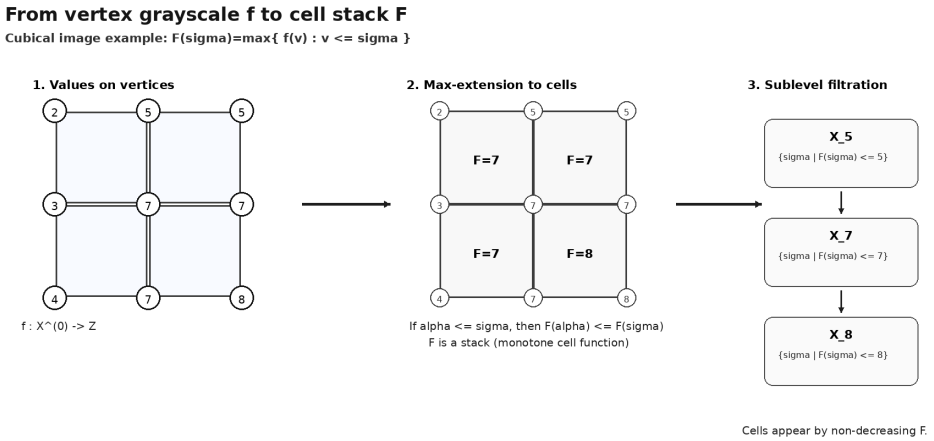


Fig. 8 From vertex grayscale values f to the monotone cell function F. The grayscale image is first given on vertices. It is then extended to every cell by $F(\sigma) = \max\{f(v) : v \leq \sigma\}$. This extension is monotone with respect to the face relation: if $\alpha \leq \sigma$, then $F(\alpha) \leq F(\sigma)$. Thus F is a stack in the sense of Bertrand. The sublevel sets $\{\sigma \in X : F(\sigma) \leq \lambda\}$ form a filtration.

Algorithm 4 Compute the max-extension F on all cells

Require: Complex X, vertex values $f : X^{(0)} \to \mathbb{Z}$
Ensure: Cell values $F : X \to \mathbb{Z}$

1: **for all** cells $\sigma \in X$ **do**
2: $\quad F(\sigma) \leftarrow \max\{f(v) : v \in X^{(0)},\ v \leq \sigma\}$
3: **end for**
4: **return** F

4.3 Injective vs. non-injective: plateaus and tie-breaking

Injective case (strict order).

If f is injective on vertices, then all vertex values are distinct and the induced ordering is strict. Most ambiguities vanish: each lower star is "well-ordered" and local pairing decisions are easier to define and analyze [6].

Non-injective case (plateaus).

Real images contain many equal-intensity vertices. Equal values create *plateaus* (flat zones) where the order is not strict. This impacts the local construction: without extra care, the pairing inside a plateau may introduce arbitrary separators that are artifacts of the tie-breaking, not of the image structure.

Two standard strategies.

- **Tie-break (total order):** perturb values (or use a lexicographic rule) to force a total order: $(f(v), \text{index}(v))$. This is simple but can create separators depending on the index order.
- **Plateau-coherent handling:** process by *levels* (all vertices with the same value together) and ensure that pairing decisions respect plateau connectivity. This avoids arbitrary internal cuts and yields more stable segmentations in flat zones.

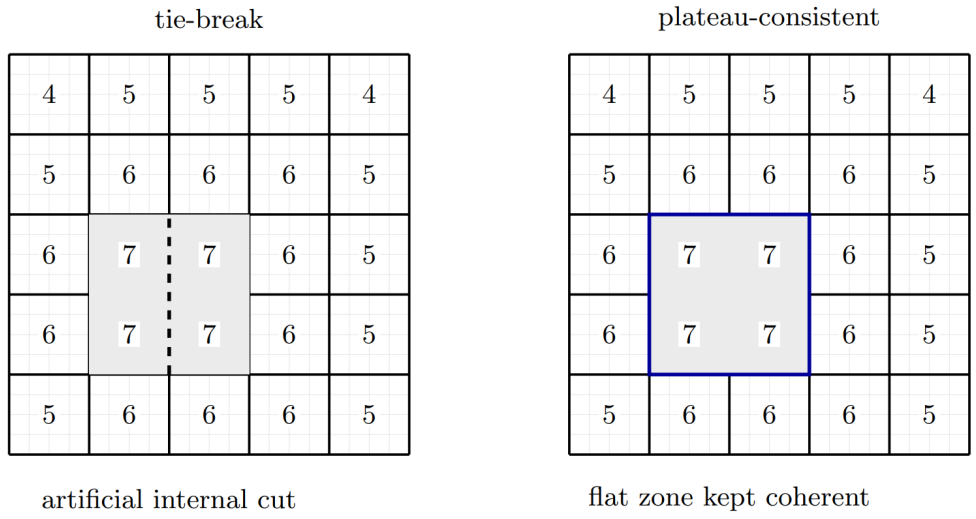


Fig. 9 A real 2D plateau configuration. The central block is a flat zone at level 7. Under tie-break, an artificial separator may appear inside the plateau. Under plateau-consistent handling, the whole flat zone is treated as a coherent unit.

4.4 Lower stars and local construction of the pairing

The RWS construction is local: it processes vertices in increasing order and builds pairings inside their *lower stars*. This is where F (and thus grayscale) enters *algorithmically*.

Definition 23 (Lower star) Let $v \in X^{(0)}$ be a vertex. The *lower star* of v (with respect to F) is

$$\mathrm{LS}(v) := \{\sigma \in X : v \leq \sigma \text{ and } F(\sigma) = f(v)\}.$$

Equivalently, $\mathrm{LS}(v)$ contains the cells incident to v whose maximum vertex value is attained at v.

Key property.

Every cell σ belongs to the lower star of (at least) one of its highest vertices. In the injective case, each cell belongs to the lower star of a *unique* vertex (its unique highest vertex). This makes the construction a disjoint union of local steps. In the presence of plateaus, multiple vertices may share the same maximal value, hence the need for plateau-aware rules (Section 4.3).

ProcessLowerStar idea.

RWS builds an acyclic pairing by repeatedly selecting available face/coface pairs inside $\mathrm{LS}(v)$, so that each cell is paired at most once and local choices remain consistent with a global acyclicity invariant [6]. The exact data structures vary, but the conceptual step is stable: *pair inside lower stars, sweep by value.*

Algorithm 5 RWS – ProcessLowerStar: local pairing construction inside $\mathrm{LS}(v)$

Require: Complex X, cell values F, current vertex v, current pairing set V
Ensure: Updated pairing set V after processing $\mathrm{LS}(v)$

$\mathcal{L} \leftarrow \mathrm{LS}(v)$
Mark all cells in $\mathcal{L}$ as *unprocessed*
while there exists an unprocessed cell in $\mathcal{L}$ **do**
  Select an unprocessed cell $\sigma \in \mathcal{L}$ of *lowest possible dimension*
  if σ has an unprocessed coface $\tau \in \mathcal{L}$ with $\dim(\tau) = \dim(\sigma) + 1$ and $\sigma < \tau$ **then**
    Choose one such coface τ (according to the local rule/tie-handling policy)
    Add (σ, τ) to V ▷ pair σ with τ
    Mark σ and τ as processed
  else
    Mark σ as processed ▷ σ remains critical *within this step*
  end if
end while
return V

What is abstracted here.

Algorithm 5 intentionally hides implementation-specific choices: how to select σ and τ, how to enforce acyclicity efficiently, and how to handle plateaus exactly. Its role in This work is to make explicit the *logical* location of grayscale:

$$f \implies F \implies \mathrm{LS}(v) \implies \text{local pairings} \implies V.$$

4.5 From gradient to segmentation (DFRS)

Once the discrete gradient field V is constructed (by RWS-style lower-star processing, or via an equivalent formulation), all segmentation and skeletonization products follow from the induced dynamics:

$$V \implies \mathrm{St}(\cdot), \mathrm{Un}(\cdot) \implies \text{basins, ridges, watershed, skeleton.}$$

Partition into basins.

Basins are stable sets of minima (Section 3.4). DFRS uses this to obtain a partition of the domain into attraction regions, with separators captured by ridge/watershed sets [7].

Skeletonization.

Skeletons are obtained by selecting unions of unstable sets of appropriate critical cells (dimension-aware selection), yielding curve-like structures in 2D/3D that summarize the geometry/topology of objects [7].

Simplification.

Finally, simplification operates by removing low-significance critical structures (and the associated stable/unstable manifolds), producing a cleaner segmentation/skeleton while preserving the main topological features. In practice, this is implemented via cancellation-like operations guided by a persistence or significance criterion, but the conceptual dependency is unchanged: the simplification is performed *on* $(V, \mathrm{St}, \mathrm{Un})$, not directly on grayscale values [5, 11].

Takeaway.

Grayscale enters *only* to build the filtration and the gradient V (Sections 4.2–4.4); all downstream objects used for segmentation are purely derived from the discrete gradient dynamics (Sections 3.3–3.4).

5 Local two-basin template and basin topology

5.1 A configuration with two basins and one separator

We describe the minimal critical configuration that captures the elementary two-basin mechanism. Consider a discrete Morse configuration on a two-dimensional cubical complex consisting of: two critical 0-cells m_1, m_2 (minima), one critical 1-cell s (the

separator saddle), and two critical 2-cells M_1, M_2 (maxima). In this configuration, the descending gradient dynamics attach each regular cell to m_1 or m_2, the saddle s sits at the interface between the two regions, and M_1, M_2 complete the two-dimensional picture on the ascending side.

Remark 5 (Topological completeness: the configuration must be completed) The Euler characteristic imposes

$$c_0 - c_1 + c_2 = \chi(X).$$

On a rectangle R, $\chi(R) = 1$, whereas the count $(c_0, c_1, c_2) = (2, 1, 2)$ gives $2 - 1 + 2 = 3 \neq 1$. The two-basin configuration described above is therefore *topologically incomplete*: it cannot stand as a valid discrete Morse configuration on any rectangular domain. It must be *completed by adding critical* 1*-cells within the same domain* until the count satisfies $c_0 - c_1 + c_2 = 1$. For instance, the count $(c_0, c_1, c_2) = (2, 3, 2)$ gives $2 - 3 + 2 = 1$, which is compatible with $\chi(R) = 1$. The right-hand panel of Figure 10 illustrates such a completion explicitly on a cubical grid with gradient field.

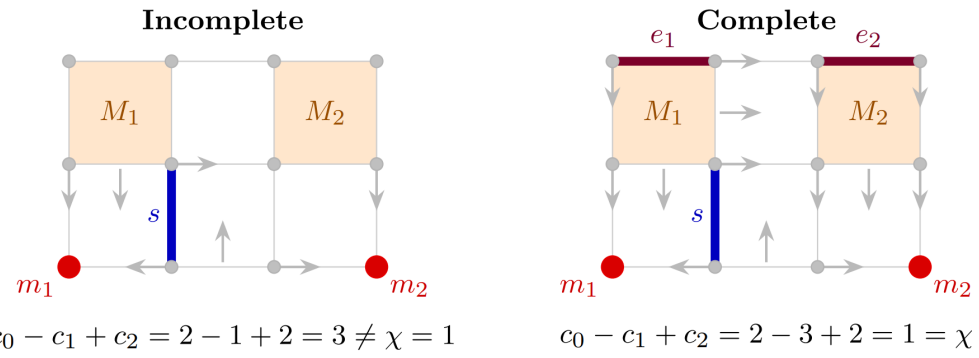


Fig. 10 Left: an incomplete discrete Morse configuration on a 3×2 cubical grid. Two minima m_1, m_2 (red), one saddle s (blue), and two maxima M_1, M_2 (orange) give $c_0 - c_1 + c_2 = 2 - 1 + 2 = 3 \neq \chi(R) = 1$. Gray arrows show gradient pairings (vertex $\to$ edge and edge $\to$ square). Right: a valid completion of the same configuration obtained by adding two critical 1-cells e_1, e_2 (purple), yielding $(c_0, c_1, c_2) = (2, 3, 2)$ and $c_0 - c_1 + c_2 = 1 = \chi(R)$. All 14 gradient pairs are shown explicitly.

5.2 Basins are regions, not necessarily closed subcomplexes

A second point emphasized during the meeting is that basin regions should not automatically be identified with closed subcomplexes of the ambient cubical complex. This distinction is essential if one wants a mathematically precise discussion of segmentation.

Definition 24 (Closed and open set in a cubical complex) Let A be a set of cells of a cubical complex X.

- We say that A is *closed* if, whenever $\sigma \in A$ and $\tau \leq \sigma$, one also has $\tau \in A$. Equivalently, A is a subcomplex of X.
- We say that A is *open* if its complement $X \setminus A$ is closed.

Definition 25 (Frontier of a basin subcomplex) Let $B(m)$ be the maximal subcomplex of X that regularly collapses onto a minimum m [7]. Its *frontier* $\mathrm{Fr}(B(m))$ is the set of cells $\sigma \notin B(m)$ such that σ is a face of some cell in $B(m)$.

Basins, vertex sets, and separating critical edges.

For a minimum m, the stable set $\mathrm{St}(m)$ is a set of 0-cells only, i.e., a set of vertices. In the examples below, the two basins are therefore represented by the vertices that descend to the two critical minima m_1 and m_2. Blue vertices belong to $\mathrm{St}(m_1)$, and green vertices belong to $\mathrm{St}(m_2)$.

More generally, one may also consider the associated basin subcomplex $B(m)$, which contains cells of several dimensions and regularly collapses onto m [7]. However, the key point illustrated here is the following: the separation between two basins is induced by the discrete gradient field, and not by a fixed geometric middle line of the domain.

In each example, the two critical 0-cells m_1 and m_2 generate two basins, while a critical 1-cell s plays the role of a separating saddle. The dashed vertical line is only a visual reference. The critical edge s need not lie on this central line: depending on the discrete gradient pairings, it may appear at different positions in the cubical grid while still separating the two basins.

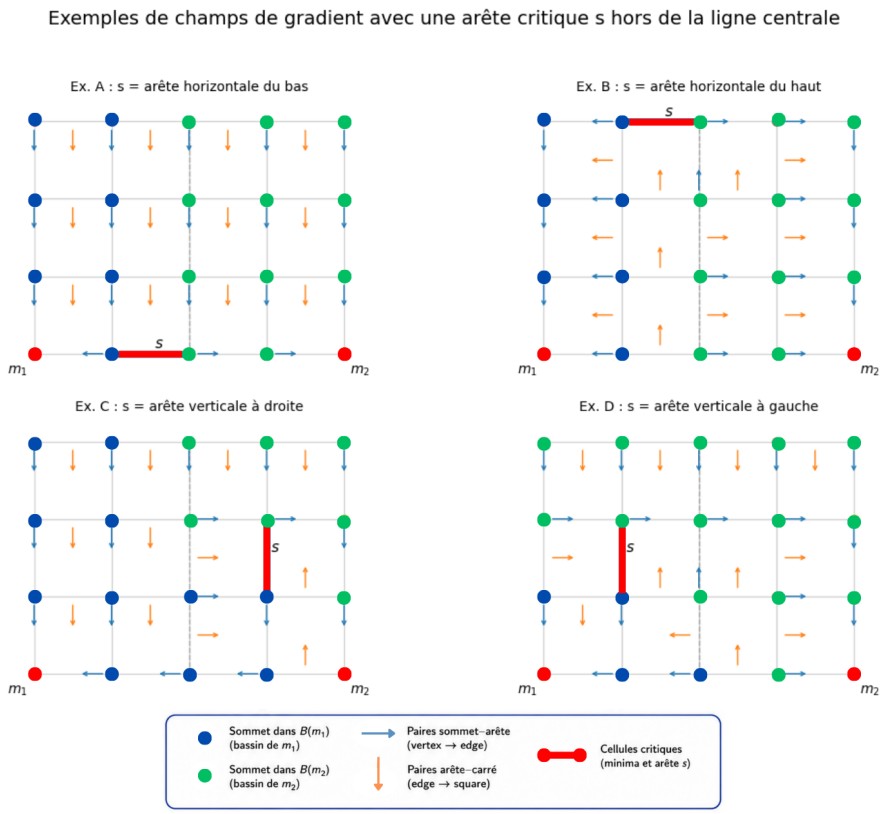


Fig. 11 Examples of discrete gradient fields on the same cubical configuration, with one critical separating edge s located away from the central reference line. Blue and green vertices belong respectively to the stable sets $\mathrm{St}(m_1)$ and $\mathrm{St}(m_2)$ of the two critical minima m_1 and m_2. Blue arrows denote vertex–edge pairings, and orange arrows denote edge–square pairings. The red cells are the critical cells: the two minima and the critical edge s. The four examples show that the separator between the two basins is determined by the discrete gradient structure, not by the geometric position of the middle line.

What this contribution adds.

Compared with the previous draft, this section makes two new points explicit. First, it isolates a minimal local picture that explains the two-basin mechanism without confusing it with a global critical-count statement. Second, it introduces the open/closed/frontier distinction needed to discuss basin regions rigorously. These two clarifications are structural: they make precise what is actually partitioned by the segmentation and in what sense separator cells sit at the interface between basins.

6 Morse sequences: constructive framework and maps

This section introduces Morse sequences as a constructive alternative to classical discrete Morse functions. The goal is twofold: (i) explain how gradients arise *by construction*, and (ii) make the reference/coreference and extension/coextension maps operational and intuitive for segmentation.

We follow the framework of Bertrand [11].

6.1 Morse sequences: regular vs critical steps

Basic idea.

Instead of defining a discrete Morse function and deducing a gradient, a *Morse sequence* directly builds the complex by elementary steps that are either *regular* (topology-preserving) or *critical* (topology-changing).

Definition 26 (Morse sequence [11, Def. 1]) Let K be a simplicial complex. A *Morse sequence* (on K) is a sequence

$$\overrightarrow{W} = \langle \emptyset = K_0, \dots, K_k = K \rangle$$

of simplicial complexes such that, for each $i \in [1, k]$, K_i is either an elementary expansion or an elementary filling of K_{i-1}.

For each $i \in [1, k]$, let κ_i be such that:

- If K_i is an elementary filling of K_{i-1} and $K_i = K_{i-1} \cup \{\nu\}$, then $\kappa_i = \nu$. We say that the face κ_i is *critical* for $\overrightarrow{W}$.
- If K_i is an elementary expansion of K_{i-1} and $K_i = K_{i-1} \cup \{\sigma, \tau\}$, with $\sigma \subseteq \tau$, then $\kappa_i = (\sigma, \tau)$. The pair κ_i is *regular* for $\overrightarrow{W}$, the face σ is *lower regular*, and the face τ is *upper regular* for $\overrightarrow{W}$.

Regular steps correspond to inverse elementary collapses; critical steps introduce new topological features.

A first toy example.

Before discussing the general properties of Morse sequences, let us consider a simple example on a 1-dimensional simplicial complex. Let

$$K = \{m_1, a, b, c, m_2,\ m_1a, ab, bc, cm_2\},$$

where m_1, a, b, c, m_2 are vertices and m_1a, ab, bc, cm_2 are edges. We consider the regular pairs

$$(a, m_1a), \qquad (b, bc), \qquad (c, cm_2),$$

and the critical faces

$$m_1, \qquad m_2, \qquad ab.$$

Hence the two critical 0-cells are m_1 and m_2, while the edge ab is a critical 1-cell.

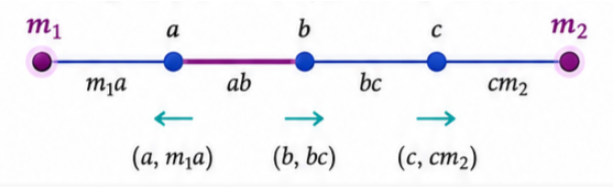


Fig. 12 A toy example of a Morse sequence on a 1-dimensional simplicial complex. The vertices are m_1, a, b, c, m_2, and the edges are m_1a, ab, bc, cm_2. The regular pairs are (a, m_1a), (b, bc), and (c, cm_2). The critical faces are m_1, m_2, and ab.

Several Morse sequences may generate the same final complex and the same regular/critical structure. For instance, one possible sequence is

$$\langle m_1,\ (a, m_1a),\ m_2,\ (c, cm_2),\ (b, bc),\ ab\rangle.$$

Another possible sequence is

$$\langle m_2,\ (c, cm_2),\ (b, bc),\ m_1,\ (a, m_1a),\ ab\rangle.$$

Yet another one is

$$\langle m_1,\ m_2,\ (c, cm_2),\ (b, bc),\ (a, m_1a),\ ab\rangle.$$

These examples illustrate two important points. First, a Morse sequence is not unique: several orders of regular and critical steps may lead to the same complex. Second, regular steps always add a paired lower face and upper face together, whereas critical steps add a single critical face.

A two-dimensional running example.

Several phenomena are invisible in dimension 1: gradient paths cannot branch, parity cancellations cannot occur, and the maps of the next subsections take a very degenerate form. We therefore fix a two-dimensional example that will be used throughout this section. Although Morse sequences are stated in the simplicial setting, all the definitions only use the face relation, so we may work directly on a cubical grid. Let X be the cubical complex of a 3×3 grid of vertices, labelled

$$\begin{array}{ccc} G & H & I \\ D & E & F \\ A & B & C \end{array}$$

with twelve edges (written AB, BE, ...) and four squares Q_1 (bottom left, vertices A, B, E, D), Q_2 (bottom right), Q_3 (top left), Q_4 (top right). We consider the gradient field V_1 given by the seven vertex–edge pairs

$$(B, AB),\ (D, AD),\ (E, BE),\ (F, CF),\ (G, DG),\ (H, EH),\ (I, FI)$$

and the three edge–square pairs

$$(DE, Q_1),\ (EF, Q_2),\ (HI, Q_4).$$

The critical cells are two vertices A and C (minima), two edges BC and GH (saddles), and one square Q_3 (maximum); the count $c_0 - c_1 + c_2 = 2 - 2 + 1 = 1 = \chi(X)$ is consistent with the Euler completeness constraint of Section 5.1. One Morse sequence inducing V_1 is

$$\begin{aligned}\langle C,\ A,\ (B, AB),\ (D, AD),\ (E, BE),\ (F, CF),\\ (G, DG),\ (H, EH),\ (I, FI),\ (DE, Q_1),\\ BC,\ GH,\ Q_3,\ (EF, Q_2),\ (HI, Q_4)\rangle.\end{aligned}$$

shown in Figure 13(b). Two remarks on the order. First, the first element of any Morse sequence on $\emptyset$ is necessarily a critical vertex: a cell can only be inserted once its boundary is present, and only a vertex has an empty boundary [11]. Second, the order is otherwise flexible: after $\langle C\rangle$, the expansion (F, CF) would have been legal as well, and many distinct sequences induce the same field V_1 (equivalence, as in the toy example above). All the values computed on this example in the sequel (reference, coreference, critical boundaries, extensions, coextensions, path counts).

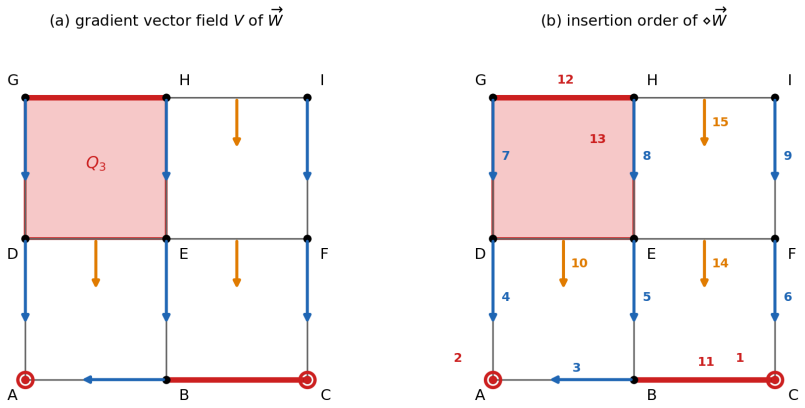


Fig. 13 The two-dimensional running example. (a) The gradient vector field V_1: blue arrows are vertex–edge pairings, orange arrows are edge–square pairings; critical cells are in red (minima A, C; saddles BC, GH; maximum Q_3). (b) The insertion order of one Morse sequence inducing V_1: a regular pair enters as one expansion step, a critical cell as one filling step. Read from right to left, the same sequence is a deconstruction of X by collapses and perforations, whose first move removes the free pair (HI, Q_4).

Why acyclicity is automatic.

The regular pairs of a Morse sequence form an acyclic matching. Indeed, the order in which the elementary expansions are introduced induces a strict ordering along

every nontrivial V-path. A closed V-path would contradict this ordering. Thus, the absence of closed V-paths follows from the Morse-sequence construction itself, and not merely from the fact that each cell belongs to at most one regular pair [11, Sec. 3 and Appendix A].

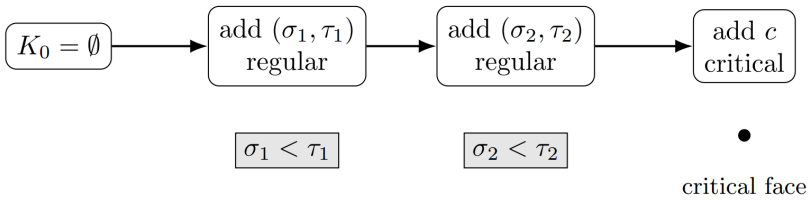


Fig. 14 A Morse sequence: the complex is built by regular steps, which add paired cells, and occasional critical steps. Regular steps preserve topology, while critical steps introduce new features. The induced gradient is acyclic by construction.

6.2 Induced gradient field and link with discrete Morse theory

Each regular step (σ, τ) defines a pairing. Collecting all regular pairs yields a discrete vector field V.

Key property.

Because the regular pairs are induced by the ordered elementary expansions of a Morse sequence, the resulting vector field admits no closed V-path [11, Sec. 3 and Appendix A].

Interpretation.

A Morse sequence is simply a structured way of producing a gradient field without ever checking acyclicity:

$$\text{sequence of expansions} \quad \Longrightarrow \quad \text{acyclic pairing } V.$$

This provides a constructive counterpart to classical discrete Morse theory [5].

6.3 Reference and coreference maps

We now introduce the reference and coreference maps associated with a Morse sequence. These maps are a central ingredient in Bertrand's framework [11], since they associate to each simplex a chain of critical simplices of the same dimension. They will also be the basis for the extension and coextension maps introduced in the next subsection.

Notation.

Let

$$\overrightarrow{W}(K) = \langle \emptyset = K_0, \ldots, K_k = K \rangle$$

be a Morse sequence on a finite simplicial complex K. Recall that each step is either:

- a *filling step*, in which case $\kappa_i = \nu$ is a critical simplex, or
- an *expansion step*, in which case $\kappa_i = (\sigma, \tau)$ is a regular pair with

$$\sigma \subseteq \tau, \qquad \dim(\tau) = \dim(\sigma) + 1.$$

We write

$$\diamond\overrightarrow{W} = \langle \kappa_1, \ldots, \kappa_k \rangle$$

for the simplex-wise form of the Morse sequence.

Following Bertrand [11], we use the following four notations:

- $$\widehat{W} = \{\nu \in K \mid \nu \text{ is critical for } \overrightarrow{W}\},$$
- $$\ddot{W} = \{(\sigma, \tau) \in K \times K \mid (\sigma, \tau) \text{ is a regular pair for } \overrightarrow{W}\},$$
- $$\overline{W} = \{\tau \in K \mid \tau \text{ is upper regular for } \overrightarrow{W}\},$$
- $$\underline{W} = \{\sigma \in K \mid \sigma \text{ is lower regular for } \overrightarrow{W}\}.$$

Thus, every regular pair $(\sigma, \tau) \in \ddot{W}$ consists of a lower regular simplex $\sigma \in \underline{W}$ and an upper regular simplex $\tau \in \overline{W}$.

Frames on the critical set.

For each $p \geq 0$, let $K[p]$ denote the vector space of p-chains over $\mathbb{Z}_2$. We also write

$$\widehat{W}[p] = \{c \in K[p] \mid c \subseteq \widehat{W}\},$$

that is, the vector subspace of $K[p]$ generated by the critical p-simplices.

A *frame* on $\widehat{W}$ is a map

$$\Gamma : \nu \in K \longmapsto \Gamma(\nu) \subseteq \widehat{W}$$

such that, for each p-simplex $\nu \in K$, the value $\Gamma(\nu)$ is a p-chain in $\widehat{W}[p]$. By linearity, Γ extends to a linear map

$$\Gamma_p : K[p] \to \widehat{W}[p], \qquad \Gamma_p(c) = \sum_{\nu \in c} \Gamma(\nu),$$

with $\Gamma_p(0) = 0$.

In what follows, and in order to keep the notation intuitive, we denote the *reference map* by γ and the *coreference map* by λ. (We note that Bertrand [11] uses the opposite lettering in Definition 6.)

Reference map.

Intuitively, the reference map answers the following question:

> When simplifying a simplex ν, to which critical simplices of the same dimension does it eventually refer?

Definition 27 (Reference map) A frame

$$\gamma_p : K[p] \to \widehat{W}[p]$$

is called a *reference map* for $\overrightarrow{W}$ if:

1. for every critical simplex $\nu \in \widehat{W}$,

$$\gamma(\nu) = \nu;$$

2. for every upper regular simplex $\tau \in \overline{W}$,

$$\gamma(\tau) = 0 \qquad \text{and} \qquad \gamma(\partial(\tau)) = 0.$$

Now let $(\sigma, \tau) \in \ddot{W}$ be a regular pair. Since σ is one face of $\partial(\tau)$, linearity over $\mathbb{Z}_2$ yields

$$\gamma(\sigma) = \gamma(\partial(\tau) + \sigma).$$

Equivalently,

$$\gamma(\sigma) = \gamma(\partial(\tau) \setminus \{\sigma\}),$$

because σ appears twice and cancels modulo 2.

Therefore, in a regular pair (σ, τ), the upper regular simplex τ disappears, while the lower regular simplex σ is replaced by the other faces of the boundary of τ.

In other words:

> γ looks at the boundary of the upper regular simplex τ.

Coreference map.

Dually, the coreference map answers the following question:

> If one starts from a regular pair (σ, τ), toward which critical simplices of the same dimension does the upper simplex τ propagate when one looks at the cofaces of σ?

Definition 28 (Coreference map) A frame

$$\lambda_p : K[p] \to \widehat{W}[p]$$

is called a *coreference map* for $\overrightarrow{W}$ if:

1. for every critical simplex $\nu \in \widehat{W}$,

$$\lambda(\nu) = \nu;$$

2. for every lower regular simplex $\sigma \in \underline{W}$,

$$\lambda(\sigma) = 0 \qquad \text{and} \qquad \lambda(\delta(\sigma)) = 0.$$

Now let $(\sigma, \tau) \in \ddot{W}$ be a regular pair. Since τ is one coface of σ, linearity over $\mathbb{Z}_2$ yields

$$\lambda(\tau) = \lambda(\delta(\sigma) + \tau).$$

Equivalently,

$$\lambda(\tau) = \lambda(\delta(\sigma) \setminus \{\tau\}),$$

because τ appears twice and cancels modulo 2.

Therefore, in a regular pair (σ, τ), the lower regular simplex σ disappears, while the upper regular simplex τ is replaced by the other cofaces of σ.

In other words:

λ looks at the higher-dimensional simplices around the lower regular simplex σ.

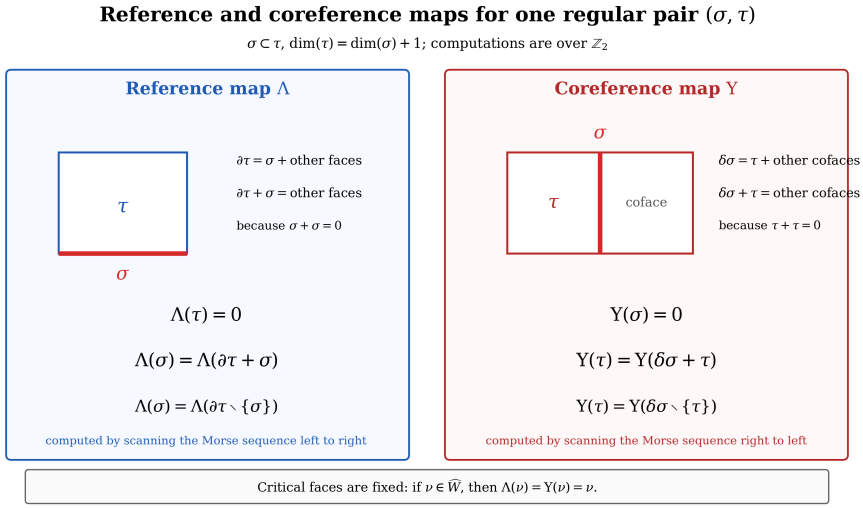


Fig. 15 Reference and coreference rules for a regular pair (σ, τ). For the reference map, the upper regular simplex τ is sent to 0, and the lower regular simplex σ is replaced by the other faces of $\partial(\tau)$:

$$\gamma(\tau) = 0, \qquad \gamma(\sigma) = \gamma(\partial(\tau) + \sigma).$$

For the coreference map, the lower regular simplex σ is sent to 0, and the upper regular simplex τ is replaced by the other cofaces of σ:

$$\lambda(\sigma) = 0, \qquad \lambda(\tau) = \lambda(\delta(\sigma) + \tau).$$

All sums are taken over $\mathbb{Z}_2$.

Elementary examples.

We now make the two rules explicit on simple configurations.

Example 1 (Reference map on a segment) Let (a, ab) be a regular pair, where a and b are vertices and ab is the edge joining them. Here a is lower regular and ab is upper regular. Since

$$\partial(ab) = a + b,$$

we obtain

$$\gamma(ab) = 0$$

and

$$\gamma(a) = \gamma(\partial(ab) + a) = \gamma(a + b + a) = \gamma(b).$$

Thus the edge ab disappears, and the vertex a is replaced by the other boundary vertex b. If b is critical, then $\gamma(a) = b$.

Example 2 (Reference map on a square) Let (e_0, Q) be a regular pair, where Q is a square and e_0 is one of its boundary edges. Write

$$\partial(Q) = e_0 + e_1 + e_2 + e_3.$$

Then

$$\gamma(Q) = 0$$

and

$$\gamma(e_0) = \gamma(\partial(Q) + e_0) = \gamma(e_1 + e_2 + e_3).$$

Thus the square Q disappears, and the edge e_0 is replaced by the other three edges in the boundary of Q.

Example 3 (Coreference map on two adjacent squares) Let two squares Q_1 and Q_2 share an edge e, and suppose that (e, Q_1) is a regular pair. Then e is lower regular and Q_1 is upper regular. Since

$$\delta(e) = Q_1 + Q_2,$$

we obtain

$$\lambda(e) = 0$$

and

$$\lambda(Q_1) = \lambda(\delta(e) + Q_1) = \lambda(Q_1 + Q_2 + Q_1) = \lambda(Q_2).$$

Thus the common edge e disappears, and the square Q_1 is replaced by the other coface Q_2 containing e. If Q_2 is critical, then $\lambda(Q_1) = Q_2$.

Example 4 (Reference and coreference on the running example) On the field V_1 of Figure 13, the three rules ($\gamma(\nu) = \nu$ if ν is critical; $\gamma(\tau) = 0$ if τ is upper regular; $\gamma(\sigma) = \gamma(\partial(\tau) + \sigma)$ for a regular pair (σ, τ)) give, by scanning the sequence from left to right:

$$\gamma(B) = \gamma(D) = \gamma(E) = \gamma(G) = \gamma(H) = A, \qquad \gamma(F) = \gamma(I) = C,$$

and, on edges, since $AB, AD, BE, CF, DG, EH, FI$ are upper regular (hence sent to 0):

$$\begin{aligned}
\gamma(DE) &= \gamma(AB{+}AD{+}BE) = 0,\\
\gamma(EF) &= \gamma(BC{+}BE{+}CF) = BC,\\
\gamma(HI) &= \gamma(EF{+}FI{+}EH) = \gamma(EF) = BC.
\end{aligned}$$

Dually, scanning from right to left ($\lambda(\sigma) = 0$ for the lower regular edges DE, EF, HI):

$$\lambda(EH) = \lambda(GH{+}HI) = GH,$$
$$\lambda(BE) = \lambda(DE{+}EF{+}EH) = GH,$$
$$\lambda(AB) = \lambda(BC{+}BE) = BC + GH.$$

Two facts are invisible in one-dimensional examples. First, a single edge may corefer to two critical edges:

$$\lambda(AB) = BC + GH.$$

Second, a lower regular edge may have zero reference:

$$\gamma(DE) = 0.$$

See Figure 16.

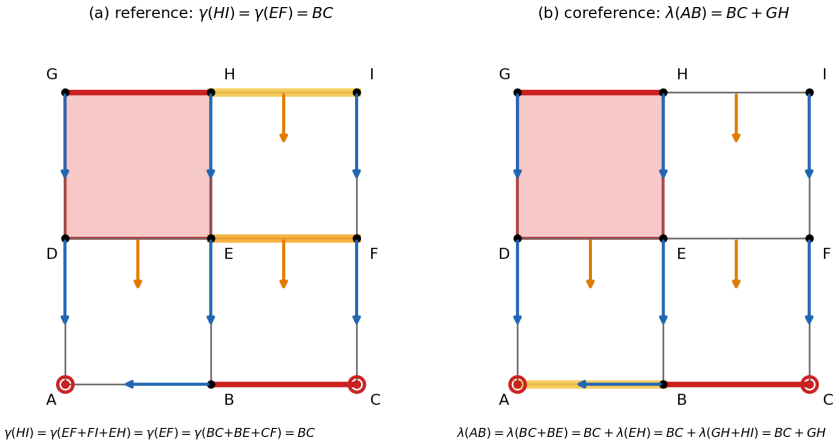


Fig. 16 Reference and coreference on the running example V_1. (a) The cascade computing $\gamma(HI)$: the pair (HI, Q_4) replaces HI by the other edges of ∂Q_4; the upper regular edges are sent to 0 and the cascade terminates on the critical edge BC. (b) Dually, $\lambda(AB) = BC + GH$: the coreference of a single edge may contain several critical edges.

Recursive characterization.

Because a Morse sequence is a filtration, the reference and coreference maps admit a recursive characterization, directly analogous to Proposition 7 in [11].

Proposition 3 (Recursive characterization [11]) *Let*

$$\overrightarrow{W}(K) = \langle \emptyset = K_0, \ldots, K_k = K \rangle$$

be a Morse sequence, and let

$$\diamond\overrightarrow{W} = \langle \kappa_1, \ldots, \kappa_k \rangle.$$

Then the reference map γ and the coreference map λ satisfy:

(A) If $\kappa_i = \nu$ is critical for $\overrightarrow{W}$, then

$$\gamma(\nu) = \nu \qquad \textit{and} \qquad \lambda(\nu) = \nu.$$

(B) If $\kappa_i = (\sigma, \tau)$ is a regular pair for $\overrightarrow{W}$, then

$$\gamma(\tau) = 0 \qquad \textit{and} \qquad \gamma(\sigma) = \gamma(\partial(\tau) + \sigma), \quad \partial(\tau) + \sigma \subseteq K_{i-1},$$

and

$$\lambda(\sigma) = 0 \qquad \textit{and} \qquad \lambda(\tau) = \lambda(\delta(\sigma) + \tau), \quad \delta(\sigma) + \tau \subseteq K \setminus K_i.$$

This proposition shows that:

- the reference map γ is naturally computed by scanning the Morse sequence from left to right;
- the coreference map λ is naturally computed by scanning it from right to left.

Theorem 4 (Existence and uniqueness [11]) *Every Morse sequence admits a unique reference map and a unique coreference map.*

Gradient and cogradient paths.

The previous recursive rules admit a path interpretation, which is one of the most useful viewpoints.

Definition 29 (Gradient path) A *gradient path* is a sequence

$$\pi = \langle \sigma_0, \tau_0, \sigma_1, \tau_1, \ldots, \sigma_r, \tau_r, \sigma_{r+1} \rangle$$

such that, for each i,

$$(\sigma_i, \tau_i) \in \ddot{W}, \qquad \sigma_{i+1} \in \partial(\tau_i), \qquad \sigma_{i+1} \neq \sigma_i.$$

Thus a gradient path alternates between a lower regular simplex, its paired upper regular simplex, and another face of the boundary of that upper simplex.

Definition 30 (Cogradient path) A *cogradient path* is a sequence

$$\pi = \langle \tau_0, \sigma_1, \tau_1, \ldots, \sigma_r, \tau_r \rangle$$

such that, for each $i \geq 1$,

$$(\sigma_i, \tau_i) \in \ddot{W}, \qquad \tau_{i-1} \in \delta(\sigma_i), \qquad \tau_i \neq \tau_{i-1}.$$

Thus a cogradient path alternates between an upper-dimensional simplex, a lower regular simplex on its boundary, and another coface paired with that lower regular simplex.

The next result gives the precise relation between the maps and these paths.

Theorem 5 (Path characterization [11]) *Let $\kappa \in \widehat{W}$ be a critical simplex and let $\nu \in K$. Then:*

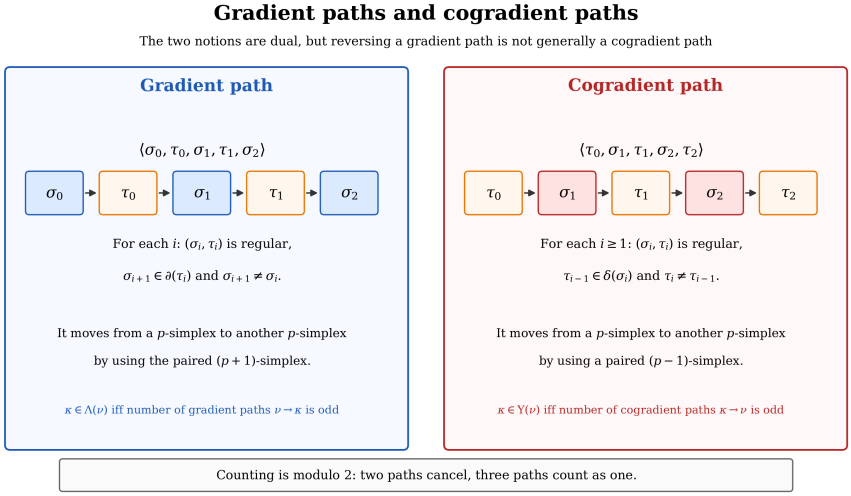


Fig. 17 Gradient and cogradient paths. A gradient path follows regular pairs (σ_i, τ_i) and then exits through another face of $\partial(\tau_i)$. A cogradient path is dual: it enters a lower regular simplex σ_i from one coface and exits through the paired coface τ_i.

1. *$\kappa \in \gamma(\nu)$ if and only if the number of gradient paths from ν to κ is odd;*
2. *$\kappa \in \lambda(\nu)$ if and only if the number of cogradient paths from κ to ν is odd.*

This theorem explains why reference and coreference maps take values in critical chains rather than in single critical simplices. Indeed, several paths may contribute, and over $\mathbb{Z}_2$, paths cancel in pairs.

Dimension 2 is the first dimension where several paths may actually coexist, because a square has *four* edges: after entering a square through its paired edge, a gradient path has three possible exits, so paths can branch — something impossible for $(0,1)$-paths, where an edge has only two vertices. Figure 18(a) shows the minimal such configuration: on the field V_2 defined by the edge–square pairs (AB, Q_1), (BE, Q_2), (DE, Q_3), (EF, Q_4) (critical cells: the vertices A, B, E and the edges EH, GH), *two* distinct gradient paths join AB to the critical edge EH:

$$\pi_1 = \langle AB, Q_1, DE, Q_3, EH\rangle \qquad \text{and} \qquad \pi_2 = \langle AB, Q_1, BE, Q_2, EF, Q_4, EH\rangle.$$

They branch at Q_1 and merge at EH. By the path-characterization theorem, this even number of paths implies $EH \notin \gamma(AB)$. This example therefore shows explicitly that path existence and path parity need not coincide in dimension two. Figure 18(b) shows a cogradient path on the running example V_1: $\pi = \langle GH, H, EH, E, BE, B, AB\rangle$, which alternates a vertex of the current edge and the edge paired with that vertex. Note that this cogradient path is exactly the descent of the vertex H towards its minimum ($H \to E \to B \to A$, following the blue arrows) read at the level of the traversed edges; this observation is used again in Section 6.5.

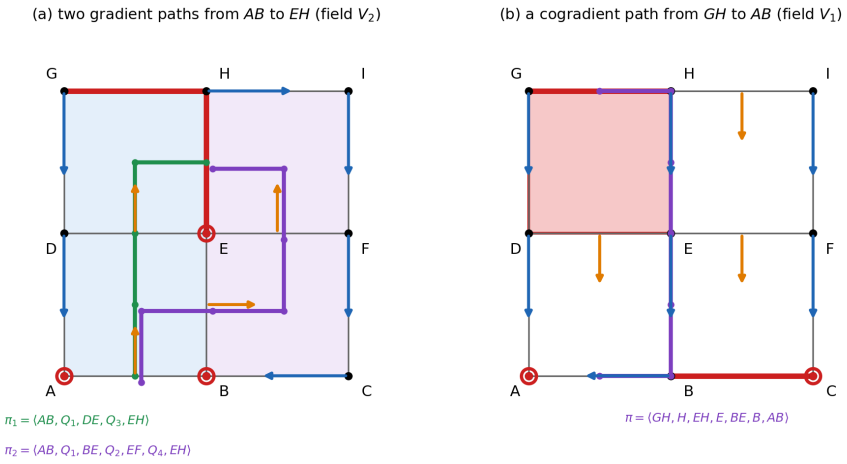


Fig. 18 Gradient and cogradient paths in dimension 2. (a) On the field V_2 (edge–square pairs (AB, Q_1), (BE, Q_2), (DE, Q_3), (EF, Q_4); critical cells A, B, E, EH, GH), two gradient paths join AB to EH: π_1 (green) and π_2 (violet); bluish and purplish squares are those traversed by π_1 and π_2 respectively (Q_1, traversed by both, is coloured with π_1). (b) A cogradient path from the critical edge GH to AB on the running example V_1.

6.4 The critical and cocritical complexes

We now introduce the critical and cocritical complexes associated with a Morse sequence. These objects are fundamental in Bertrand's framework [11], because they compress the ambient complex onto its critical simplices while preserving the relevant algebraic structure.

Reminder on notation.

In this report, we denote by

$$\gamma : K[p] \to \widehat{W}[p]$$

the reference map, and by

$$\lambda : K[p] \to \widehat{W}[p]$$

the coreference map. We keep this notation throughout, even though Bertrand uses the opposite lettering in Definition 6 of [11].

The critical chain complex

The reference map allows one to transport the boundary operator from the ambient complex to the critical simplices.

Definition 31 (Critical boundary map) For each $p \geq 1$, define

$$\widehat{\partial}_p : \widehat{W}[p] \to \widehat{W}[p-1]$$

by

$$\widehat{\partial}_p(\kappa) := \gamma(\partial(\kappa)) \qquad \text{for every critical } p\text{-simplex } \kappa \in \widehat{W}[p],$$

and extend linearly to critical p-chains.

In other words, the boundary of a critical simplex is first computed in the ambient complex, and then projected back onto critical simplices through the reference map.

Definition 32 (Critical complex) The sequence

$$\widehat{W}_\bullet = \big(\widehat{W}[p], \widehat{\partial}_p\big)_{p\geq 0}$$

is called the *critical chain complex*.

Proposition 6 ([11]) *The critical boundary satisfies*

$$\widehat{\partial}_{p-1} \circ \widehat{\partial}_p = 0.$$

Hence $\widehat{W}_\bullet$ *is a chain complex.*

Proof Let $\kappa \in \widehat{W}[p]$ be critical. Since $\gamma(\kappa) = \kappa$, we have

$$\widehat{\partial}_{p-1}\widehat{\partial}_p(\kappa) = \widehat{\partial}_{p-1}\big(\gamma(\partial(\kappa))\big) = \gamma\big(\partial(\partial(\kappa))\big) = 0.$$

Thus $\widehat{\partial}_{p-1} \circ \widehat{\partial}_p = 0$. □

The cocritical cochain complex

Dually, the coreference map transports the coboundary operator to the critical simplices.

Definition 33 (Cocritical coboundary map) For each $p \geq 0$, define

$$\widehat{\delta}_p : \widehat{W}[p] \to \widehat{W}[p+1]$$

by

$$\widehat{\delta}_p(\kappa) := \lambda(\delta(\kappa)) \qquad \text{for every critical } p\text{-simplex } \kappa \in \widehat{W}[p],$$

and extend linearly.

Definition 34 (Cocritical complex) The sequence

$$\widehat{W}^\bullet = \big(\widehat{W}[p], \widehat{\delta}_p\big)_{p\geq 0}$$

is called the *cocritical cochain complex*.

Proposition 7 ([11]) *The cocritical coboundary satisfies*

$$\widehat{\delta}_{p+1} \circ \widehat{\delta}_p = 0.$$

Hence $\widehat{W}^\bullet$ *is a cochain complex.*

Proof Let $\kappa \in \widehat{W}[p]$ be critical. Since $\lambda(\kappa) = \kappa$, we obtain

$$\widehat{\delta}_{p+1}\widehat{\delta}_p(\kappa) = \widehat{\delta}_{p+1}\big(\lambda(\delta(\kappa))\big) = \lambda\big(\delta(\delta(\kappa))\big) = 0.$$

Thus $\widehat{\delta}_{p+1} \circ \widehat{\delta}_p = 0$. □

Interpretation.

The critical complex keeps only the critical simplices and records how their boundaries interact after projection by γ. The cocritical complex is the dual object: it records how critical simplices interact through coboundaries after projection by λ.

Example 5 (Critical and cocritical complexes on the toy path) We return to the toy Morse sequence of Figure 12, with critical simplices

$$m_1, \qquad m_2, \qquad ab,$$

and regular pairs

$$(a, m_1 a), \qquad (b, bc), \qquad (c, cm_2).$$

The critical chain groups are

$$\widehat{W}[0] = \langle m_1, m_2 \rangle, \qquad \widehat{W}[1] = \langle ab \rangle.$$

Since

$$\partial(ab) = a + b,$$

we obtain

$$\widehat{\partial}_1(ab) = \gamma(a+b) = \gamma(a) + \gamma(b) = m_1 + m_2.$$

Also,

$$\widehat{\partial}_0(m_1) = 0, \qquad \widehat{\partial}_0(m_2) = 0.$$

Hence the critical complex is

$$0 \longleftarrow \langle m_1, m_2 \rangle \xleftarrow{\widehat{\partial}_1} \langle ab \rangle \longleftarrow 0, \qquad \widehat{\partial}_1(ab) = m_1 + m_2.$$

Dually, the cocritical coboundary is computed from the coreference map. Since

$$\delta(m_1) = m_1 a, \qquad \delta(m_2) = cm_2,$$

we get

$$\widehat{\delta}_0(m_1) = \lambda(m_1 a) = ab, \qquad \widehat{\delta}_0(m_2) = \lambda(cm_2) = ab.$$

Moreover,

$$\widehat{\delta}_1(ab) = 0.$$

Hence the cocritical complex is

$$0 \longrightarrow \langle m_1, m_2 \rangle \xrightarrow{\widehat{\delta}_0} \langle ab \rangle \longrightarrow 0, \qquad \widehat{\delta}_0(m_1) = \widehat{\delta}_0(m_2) = ab.$$

Example 6 (Critical and cocritical complexes on the running example) On the field V_1, with the values of Example 4, the critical boundary $\widehat{\partial} = \gamma \circ \partial$ gives

$$\widehat{\partial}_1(BC) = \gamma(B) + \gamma(C) = A + C, \qquad \widehat{\partial}_1(GH) = \gamma(G) + \gamma(H) = A + A = 0,$$
$$\widehat{\partial}_2(Q_3) = \gamma(DE{+}DG{+}EH{+}GH) = GH.$$

The value $\widehat{\partial}_1(GH) = 0$ is the interesting one, and has no analogue on the one-dimensional toy path: GH is critical, yet its critical boundary vanishes because its two endpoints flow to the *same* minimum A, so their contributions cancel modulo 2. Geometrically, GH creates a 1-cycle, which is then filled by the critical square Q_3 (whence $\widehat{\partial}_2(Q_3) = GH$). One reads the homology of the disk directly on the critical complex: $\beta_0 = 2 - 1 = 1$, $\beta_1 = 1 - 1 = 0$, $\beta_2 = 0$, with 5 cells instead of 25. Dually, the cocritical coboundary $\widehat{\delta} = \lambda \circ \delta$ gives

$$\widehat{\delta}_0(A) = \lambda(AB{+}AD) = (BC{+}GH) + GH = BC, \qquad \widehat{\delta}_0(C) = \lambda(BC{+}CF) = BC,$$
$$\widehat{\delta}_1(GH) = \lambda(Q_3) = Q_3, \qquad \widehat{\delta}_1(BC) = \lambda(Q_2) = 0,$$

and one checks $\widehat{\delta} \circ \widehat{\delta} = 0$ on this example: $\widehat{\delta}_1(\widehat{\delta}_0(A)) = \widehat{\delta}_1(BC) = 0$. See Figure 20.

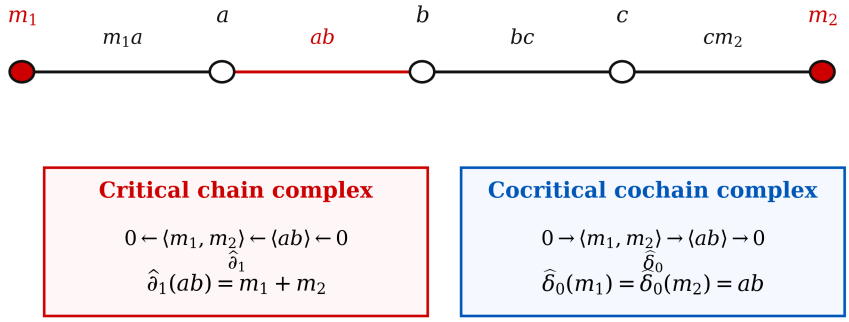


Fig. 19 Toy example of the critical and cocritical complexes. Top: the 1-dimensional Morse sequence with critical simplices m_1, m_2, and ab. Bottom left: the critical complex, with $\hat{\partial}_1(ab) = m_1 + m_2$. Bottom right: the cocritical complex, with $\hat{\delta}_0(m_1) = \hat{\delta}_0(m_2) = ab$.

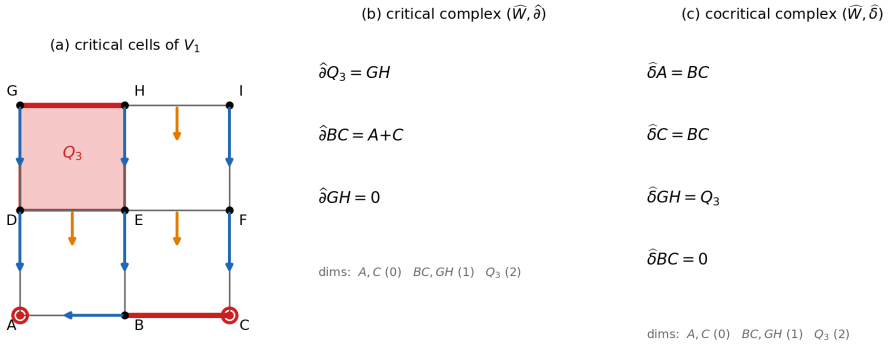


Fig. 20 The critical and cocritical complexes of the running example. (a) The five critical cells of V_1. (b) Critical complex: the saddle BC merges the two minima ($\hat{\partial}BC = A + C$); the saddle GH has vanishing critical boundary (both endpoints flow to A); the maximum Q_3 fills the 1-cycle created by GH. (c) The cocritical complex, computed dually with λ.

Example 7 (Critical and cocritical complexes with three basins) The previous examples contain at most two minima. To see which values these complexes can take in a richer configuration, consider the gradient field V_3 shown in Figure 21(a): a 5×4 cubical grid with three minima m_1, m_2, m_3, three saddles s_1, s_2, s_3, and one maximum M, so that $c_0 - c_1 + c_2 = 3 - 3 + 1 = 1 = \chi(X)$, in accordance with the completeness constraint of Section 5.1. Following the blue arrows, every vertex descends onto one of the three minima, and each saddle sits between two different attraction regions. Since the reference map sends every vertex to the minimum it descends to, each saddle joins the minima of its two neighbouring regions:

$$\hat{\partial}s_1 = m_1 + m_3, \qquad \hat{\partial}s_2 = m_2 + m_3, \qquad \hat{\partial}s_3 = m_1 + m_2.$$

The new phenomenon concerns the maximum. The boundary of M consists of the two critical edges s_2, s_3 and two regular edges. The right edge is upper regular (it is paired with a vertex), hence sent to 0 by γ; the left edge is paired with the square on its left, and the resulting cascade terminates on the critical edge s_1. Therefore

$$\widehat{\partial} M = s_1 + s_2 + s_3,$$

a value that cannot occur with two basins: the critical boundary of a maximum may involve a saddle, here s_1, which is *not* a face of M. One checks by hand that $\widehat{\partial}\widehat{\partial} M = (m_1{+}m_3) + (m_2{+}m_3){+}(m_1{+}m_2) = 0$. On the cocritical side, $\widehat{\delta} m_i$ is the sum of the two saddles bordering the region of m_i, and $\widehat{\delta} s_1 = \widehat{\delta} s_2 = \widehat{\delta} s_3 = M$. The homology of the disk is again read on 7 cells instead of 63: $\operatorname{rank} \widehat{\partial}_1 = 2$ gives $\beta_0 = 3 - 2 = 1$, and $\ker \widehat{\partial}_1 = \langle s_1{+}s_2{+}s_3 \rangle = \operatorname{im} \widehat{\partial}_2$ gives $\beta_1 = \beta_2 = 0$. All the values above were verified by computer against the path characterization theorem of Section 6.3, by exhaustively counting gradient and cogradient paths. The field V_3 is used again in Part II as the guiding example for the segmentation layer.

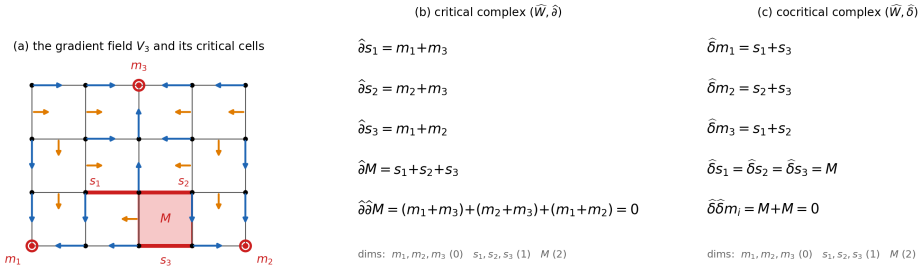


Fig. 21 Critical and cocritical complexes of a field with three basins. (a) The gradient field V_3 on a 5×4 grid: blue arrows are vertex–edge pairings, orange arrows are edge–square pairings; the critical cells (red) are the minima m_1, m_2, m_3, the saddles s_1, s_2, s_3, and the maximum M. (b) Critical complex: each saddle joins the minima of its two neighbouring basins, and $\widehat{\partial} M = s_1{+}s_2{+}s_3$; note that s_1 is not a face of M, yet it appears in the critical boundary of M; the identity $\widehat{\partial} \circ \widehat{\partial} = 0$ can be checked directly. (c) Cocritical complex: $\widehat{\delta} m_i$ is the sum of the two saddles bordering the basin of m_i, and all three saddles have cocritical coboundary M.

6.5 Extension and coextension maps

We now introduce the extension and coextension maps. These maps go in the opposite direction of the reference and coreference maps: instead of sending an ambient simplex to a critical chain, they propagate a critical simplex back into the ambient complex [11].

From critical simplices back to the ambient complex.

Let $\kappa \in \widehat{W}[p]$ be a critical p-simplex. The coextension and extension associated with κ are defined by collecting the p-simplices of K that refer or corefer to κ.

Definition 35 (Coextension and extension maps) For every critical p-simplex $\kappa \in \widehat{W}[p]$, define

$$\widehat{\Upsilon}_p(\kappa) = \sum_{\substack{\nu \in K[p] \\ \kappa \in \gamma(\nu)}} \nu, \qquad \widehat{\Lambda}_p(\kappa) = \sum_{\substack{\nu \in K[p] \\ \kappa \in \lambda(\nu)}} \nu.$$

Equivalently, in support notation,

$$\nu \in \widehat{\Upsilon}_p(\kappa) \iff \kappa \in \gamma(\nu), \qquad \nu \in \widehat{\Lambda}_p(\kappa) \iff \kappa \in \lambda(\nu).$$

Thus:

- the *coextension* $\widehat{\Upsilon}_p(\kappa)$ is built from the reference map γ;
- the *extension* $\widehat{\Lambda}_p(\kappa)$ is built from the coreference map λ.

Interpretation.

The coextension $\widehat{\Upsilon}_p(\kappa)$ collects the p-simplices whose reference contains κ. The extension $\widehat{\Lambda}_p(\kappa)$ collects the p-simplices whose coreference contains κ. They should therefore be read as p-dimensional objects in the ambient complex attached to the critical simplex κ.

Proposition 8 (Structural properties [11]) *Let κ be a critical simplex. Then:*

1. *κ belongs to $\widehat{\Upsilon}(\kappa)$ and to $\widehat{\Lambda}(\kappa)$;*
2. *if $\nu \in \widehat{\Upsilon}(\kappa)$ and $\nu \neq \kappa$, then ν is lower regular;*
3. *if $\nu \in \widehat{\Lambda}(\kappa)$ and $\nu \neq \kappa$, then ν is upper regular.*

Meaning of the proposition.

The coextension of a critical simplex contains that critical simplex itself, together with lower-regular simplices of the same dimension. Dually, the extension contains the critical simplex together with upper-regular simplices of the same dimension.

Example 8 (Coextensions on the toy path) Consider again the toy path of Figure 12. We compute the coextensions of the critical vertices m_1 and m_2.

Since

$$\gamma(m_1) = m_1, \qquad \gamma(a) = m_1,$$

we obtain

$$\widehat{\Upsilon}_0(m_1) = m_1 + a.$$

Similarly, since

$$\gamma(b) = m_2, \qquad \gamma(c) = m_2, \qquad \gamma(m_2) = m_2,$$

we get

$$\widehat{\Upsilon}_0(m_2) = b + c + m_2.$$

Finally, for the critical edge ab, we simply have

$$\widehat{\Upsilon}_1(ab) = ab.$$

In support notation:

$$\widehat{\Upsilon}_0(m_1) = \{m_1, a\}, \qquad \widehat{\Upsilon}_0(m_2) = \{b, c, m_2\}, \qquad \widehat{\Upsilon}_1(ab) = \{ab\}.$$

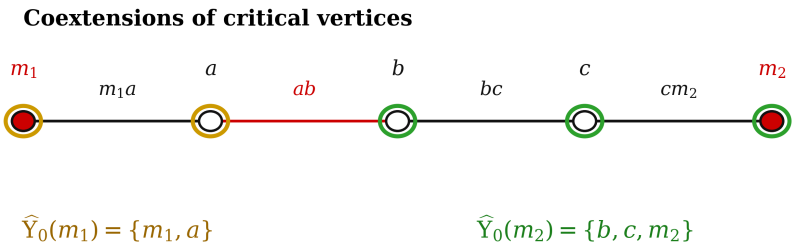


Fig. 22 Coextensions on the toy path. Left: $\widehat{\Upsilon}_0(m_1) = \{m_1, a\}$. Right: $\widehat{\Upsilon}_0(m_2) = \{b, c, m_2\}$. Only the relevant 0-simplices are highlighted.

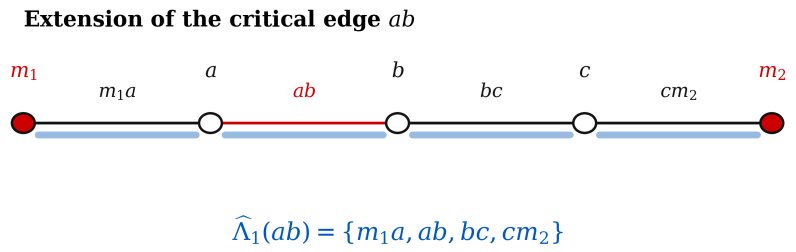


Fig. 23 Extension of the critical edge on the toy path. The set $\widehat{\Lambda}_1(ab) = \{m_1a, ab, bc, cm_2\}$ is highlighted. Only the relevant 1-simplices are shown.

Example 9 (Extensions on the toy path) We now compute the extensions.

For the critical edge ab, we have

$$\lambda(ab) = ab, \qquad \lambda(m_1a) = ab, \qquad \lambda(bc) = ab, \qquad \lambda(cm_2) = ab.$$

Therefore

$$\widehat{\Lambda}_1(ab) = m_1a + ab + bc + cm_2.$$

In support notation,

$$\widehat{\Lambda}_1(ab) = \{m_1a, ab, bc, cm_2\}.$$

For the critical vertices m_1 and m_2, one has only

$$\lambda(m_1) = m_1, \qquad \lambda(m_2) = m_2,$$

so

$$\widehat{\Lambda}_0(m_1) = m_1,$$
$$\widehat{\Lambda}_0(m_2) = m_2.$$

Example 10 (Extensions and coextensions on the running example) On the field V_1, reading the tables of Example 4 "by columns" gives

$$\widehat{\Upsilon}_0(A) = \{A, B, D, E, G, H\}, \qquad \widehat{\Upsilon}_0(C) = \{C, F, I\},$$

$$\widehat{\Lambda}_1(BC) = \{BC, AB\}, \qquad \widehat{\Lambda}_1(GH) = \{GH, EH, BE, AB, AD, DG\}, \qquad \widehat{\Lambda}_2(Q_3) = \{Q_3, Q_1\}.$$

Three observations, illustrated in Figure 24. First, the coextensions of the two minima *partition* the vertices: they are precisely the two vertex basins of the field, in accordance with the two-basin picture of Section 5.2. Second, the extensions do *not* partition: the edge AB belongs to both $\widehat{\Lambda}(BC)$ and $\widehat{\Lambda}(GH)$, since $\lambda(AB) = BC + GH$. Third, $\widehat{\Lambda}(GH)$ can be read without any computation: the cogradient paths issued from GH are the descents of its two endpoints ($H \to E \to B \to A$ and $G \to D \to A$) seen at the level of edges, so $\widehat{\Lambda}(GH) = \{GH\} \cup \{EH, BE, AB\} \cup \{DG, AD\}$ — the two "descending staircases" of GH, which is precisely the structure of the canonical paths used by DFRS to connect adjacent basins [7].

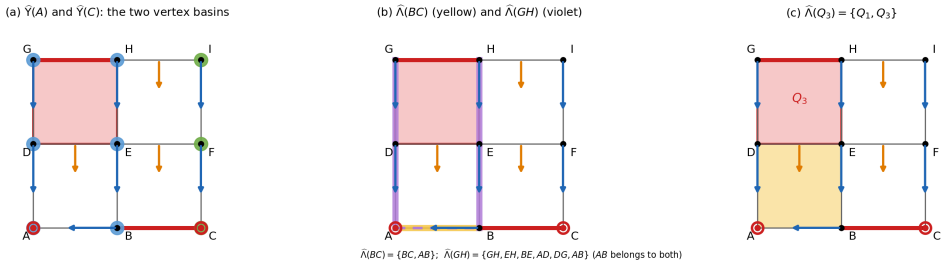


Fig. 24 Extensions and coextensions on the running example V_1. (a) The coextensions of the two minima, $\widehat{\Upsilon}(A)$ (blue) and $\widehat{\Upsilon}(C)$ (green), partition the vertices: they are the two vertex basins. (b) The extensions of the two saddles: $\widehat{\Lambda}(BC) = \{BC, AB\}$ (yellow) and $\widehat{\Lambda}(GH)$ (violet); the edge AB belongs to both (dashed overlay). (c) The extension of the maximum, $\widehat{\Lambda}(Q_3) = \{Q_3, Q_1\}$, a strict subregion of the domain.

Why these notions matter here.

These constructions are particularly important for reading Robins' framework in the language of Morse sequences. Indeed, they provide ambient objects attached to critical simplices and therefore play a central role when one wants to translate basin-, separator-, or skeleton-type constructions into the sequence-based setting.

Important caution.

Extension and coextension should not be identified automatically with the stable and unstable sets of RWS/DFRS. They are defined through the algebraic maps γ and λ, hence through parity of gradient or cogradient paths [11], whereas stable and unstable sets are defined by existence of gradient paths in the discrete gradient dynamics [6, 7]. They are related objects, but they are not identical in general. The two-dimensional field of Figure 18(a) provides an explicit counterexample: two gradient paths connect AB to EH, so AB belongs to the existence-based stable set of EH, whereas the two contributions cancel modulo 2.

Part II – Contributions

7 Contributions : la couche segmentation en langage de séquences de Morse

Cette partie développe la traduction de la couche segmentation de Delgado-Friedrichs, Robins et Sheppard [7] dans le langage des séquences de Morse. À la différence de la version précédente, les énoncés sont maintenant accompagnés de leurs preuves. Les endroits où la traduction ne se réduit pas à un transport d'argument sont signalés en remarque ; le plus notable est l'unicité du bassin, que DFRS obtiennent au moyen des valeurs de la fonction de Morse discrète et que nous obtenons par l'ordre d'insertion de la séquence. Tous les énoncés ont en outre été contrôlés par programme, sur l'exemple guide de la section 7.2 et sur des champs de gradient aléatoires.

7.1 Objectif, choix des morceaux, questions posées

La Partie I montre que l'étape de *construction* du pipeline en niveaux de gris possède déjà une contrepartie dans le langage des séquences : le champ de gradient peut être produit par des séquences de Morse sur des stacks et par des flooding sequences [12], et le complexe de Morse se lit sur la séquence à travers la reference map [11]. L'étape de *segmentation*, elle, n'a pas de contrepartie publiée. C'est pourtant le cœur théorique revendiqué par DFRS : leur section 3 s'ouvre sur les ensembles stables et instables et se ferme sur les bassins, et ils écrivent explicitement que ces définitions et résultats constituent *« the new theoretical work presented in this paper »* [7, Sec. 3]. Chez Robins, Wood et Sheppard, la même couche apparaît dès l'introduction : les bassins et leurs crêtes y motivent toute la construction, comme discrétisation des variétés stables et instables [6, Sec. 1].

Notre objectif est donc de compléter le dictionnaire commencé en Partie I en traduisant cette couche, c'est-à-dire les Lemmes 1 à 7 et la définition des bassins de [7, Sec. 3.1 et 3.3]. Nous découpons la traduction en quatre morceaux, notés B5, B6, B8 et B11 d'après la numérotation de nos tableaux de correspondance, qui forment une seule chaîne de dépendances :

$$\underbrace{\text{St} \,/\, \text{Un}}_{\text{B5}} \longrightarrow \underbrace{\text{partition des sommets}}_{\text{B6}} \longrightarrow \underbrace{\text{compatibilité avec les niveaux}}_{\text{B8}} \longrightarrow \underbrace{\text{bassins}}_{\text{B11}}.$$

Les ponts (Lemme 8 et Théorème 9 de [7]), le squelette et la simplification s'appuient sur ces quatre morceaux et sont laissés pour une étape ultérieure (section 7.8).

Quatre questions guident le travail, une par morceau.

(Q1) Les ensembles stables et instables de DFRS peuvent-ils être définis uniquement à partir d'une séquence de Morse, et comment se situent-ils par rapport aux extensions et coextensions de la section 6.5 ?

(Q2) La partition des sommets par les minima [7, Lem. 5] survit-elle telle quelle, et où passe exactement la frontière entre existence et parité ?

(Q3) Peut-on lire la segmentation à *tous* les seuils de gris à partir d'un seul calcul, y compris pour des images à plateaux, sans la perturbation des valeurs que supposent RWS et DFRS ?

(Q4) Le bassin, qui est un sous-complexe chez DFRS, admet-il une caractérisation constructive par les expansions de la séquence, indépendante de l'ordre des choix, et calculable en un seul balayage ?

Les quatre questions reçoivent ici une réponse positive : (Q1) par les Propositions 11 et 12, (Q2) par le Théorème 13, (Q3) par le Théorème 16 et le Corollaire 17, (Q4) par les Théorèmes 18 et 20 et le Corollaire 23. Deux gains sont propres aux séquences. D'une part l'ordre d'insertion fournit un ordonnanceur naturel, là où DFRS notent que le parcours naïf des V-paths peut être exponentiel dans le pire cas [7, Sec. 2.4]. D'autre part un stack n'a pas besoin d'être injectif [12], alors que DFRS supposent que « *the values of f are unique on the cells of K* » [7, Sec. 3] et que RWS s'y ramènent par perturbation simulée (section 4.3).

7.2 L'exemple guide

Tous les énoncés qui suivent sont illustrés et vérifiés sur un même exemple, présenté sur la figure 25 : une grille de 5×4 carrés (30 sommets, 49 arêtes, 20 carrés, soit 99 cellules), munie d'un champ de gradient à 46 paires régulières et sept cellules critiques — trois minima m_1, m_2, m_3, trois selles s_1, s_2, s_3 et un maximum M, d'où $c_0 - c_1 + c_2 = 3 - 3 + 1 = 1 = \chi(K)$, conformément à la contrainte d'Euler de la section 5.1. Le champ V_3 de l'exemple 7 en est une version réduite, que nous avons utilisée en Partie I pour les complexes critique et cocritique (figure 21) ; nous passons ici à l'exemple complet car la couche segmentation demande des configurations plus riches : bassins de tailles inégales, cellules hors de tout bassin, et plusieurs niveaux de coupe. L'exemple a été contrôlé par un programme indépendant : couplage, acyclicité, existence d'une séquence de Morse compatible, et toutes les valeurs annoncées dans les légendes.

7.3 Notations et lemme d'ordre

Dans toute cette partie, K est un complexe cubique (ou simplicial) fini, $\overrightarrow{W} = \langle \emptyset = K_0, \ldots, K_N = K \rangle$ est une séquence de Morse sur K (section 6.1), $\diamond\overrightarrow{W} = \langle \kappa_1, \ldots, \kappa_N \rangle$ sa forme cellule par cellule, et V l'ensemble de ses paires régulières, avec la notation fonctionnelle $V(\sigma) = \tau$ pour $(\sigma, \tau) \in V$. Les cellules critiques sont les cellules non appariées, et $\mathcal{M}$ désigne l'ensemble des sommets critiques. Nous notons $t(\nu) = i$ le *pas d'insertion* de la cellule ν, c'est-à-dire l'indice de l'étape qui insère ν ; les deux cellules d'une paire régulière $\kappa_i = (\sigma, \tau)$ reçoivent le même pas $t(\sigma) = t(\tau) = i$. Cette fonction est, à normalisation près, la fonction de Morse canonique de la séquence [11]. L'outil de base de toute la partie est le lemme suivant.

Lemma 9 (Lemme d'ordre) *Soit $(\sigma^{(p)}, \tau^{(p+1)})$ la paire régulière insérée au pas i. Alors :*

1. *toute facette $\sigma' \neq \sigma$ de τ vérifie $t(\sigma') < i$;*
2. *toute cofacette $\tau' \neq \tau$ de σ vérifie $t(\tau') > i$.*

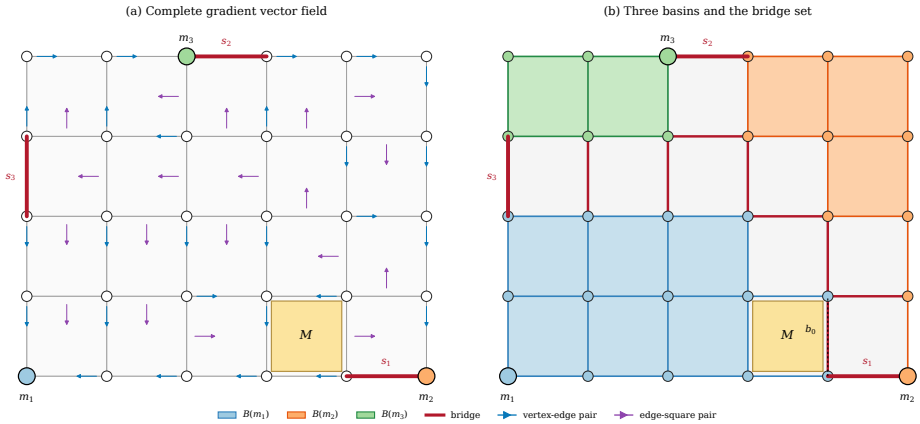


Fig. 25 L'exemple guide. (a) Le champ de gradient complet sur la grille de 5×4 carrés : flèches bleues pour les paires sommet–arête, flèches violettes pour les paires arête–carré ; les cellules critiques sont les minima m_1, m_2, m_3, les selles s_1, s_2, s_3 et le maximum M. (b) Les trois bassins $B(m_1)$ (bleu), $B(m_2)$ (orange), $B(m_3)$ (vert), qui sont des *sous-complexes* : ils contiennent des sommets, des arêtes et des carrés, avec les comptes $(14, 19, 6)$, $(10, 12, 3)$ et $(6, 7, 2)$ en dimensions $0, 1, 2$. Les 1-cellules rouge sombre, comme b_0, n'appartiennent à aucun bassin : ce sont les *ponts* de DFRS (section 7.7).

Proof (1) $K_i = K_{i-1} \cup \{\sigma, \tau\}$ est un complexe, donc $\partial\tau \subseteq K_i$. Une facette $\sigma' \neq \sigma$ de τ est distincte de σ et de τ, donc $\sigma' \in K_{i-1}$, c'est-à-dire $t(\sigma') \leq i-1$. (2) Soit $\tau' \neq \tau$ une cofacette de σ. Si l'on avait $t(\tau') \leq i$, alors $\tau' \in K_{i-1}$ (car $\tau' \notin \{\sigma, \tau\}$), et K_{i-1} étant clos par faces, on aurait $\sigma \in K_{i-1}$, ce qui contredit le fait que σ est insérée au pas i. □

Corollary 10 *Le long d'un gradient path $\langle \sigma_0, \tau_0, \sigma_1, \ldots, \sigma_r \rangle$ (section 6.3), les pas d'insertion décroissent strictement : $t(\sigma_0) > t(\sigma_1) > \cdots > t(\sigma_r)$. Le long d'un cogradient path $\langle \tau_0, \sigma_1, \tau_1, \ldots, \tau_r \rangle$, ils décroissent également : $t(\tau_0) > t(\tau_1) > \cdots > t(\tau_r)$. En particulier V n'admet aucun chemin fermé non trivial.*

Proof Pour un gradient path, σ_{i+1} est une facette de τ_i distincte de σ_i, et le Lemme 9(1) appliqué à la paire (σ_i, τ_i) donne $t(\sigma_{i+1}) < t(\sigma_i)$. Pour un cogradient path, τ_{i-1} est une cofacette de σ_i distincte de τ_i, et le Lemme 9(2) appliqué à (σ_i, τ_i) donne $t(\tau_{i-1}) > t(\tau_i)$. Un chemin fermé non trivial donnerait $t(\sigma_0) < t(\sigma_0)$. □

Remark 6 Le Lemme 9(2) redonne aussi le fait que (σ, τ) est une paire *libre* de K_i : au pas i, τ est la seule cofacette de σ présente. C'est ce qui fait de chaque étape régulière une expansion élémentaire (section 2.4).

7.4 B5 : les ensembles stables et instables en langage séquence

Où c'est chez Robins.

Les ensembles stables et instables sont définis dans [7, Sec. 3.1], comme objets géométriques portés par les V-paths ; DFRS soulignent que ce choix diverge de la formulation algébrique du flot de Forman. La motivation remonte à [6, Sec. 1] : bassins et crêtes d'un paysage comme variétés stables des points critiques.

7.4.1 Définition exacte de DFRS

Pour une p-cellule critique α d'un complexe K muni du champ V, l'ensemble stable $S_K(\alpha)$ est le plus petit ensemble de p-cellules tel que $\alpha \in S_K(\alpha)$ et

$$V(\delta^{(p)}) = \gamma^{(p+1)} > \beta^{(p)} \in S_K(\alpha) \implies \delta \in S_K(\alpha),$$

et l'ensemble instable $U_K(\alpha)$ est le plus petit ensemble de p-cellules tel que $\alpha \in U_K(\alpha)$ et

$$V^{-1}(\delta^{(p)}) = \gamma^{(p-1)} < \beta^{(p)} \in U_K(\alpha) \implies \delta \in U_K(\alpha),$$

avec la convention $V(\delta) = 0$ (resp. $V^{-1}(\delta) = 0$) quand δ n'est l'élément inférieur (resp. supérieur) d'aucune paire [7, Sec. 3.1].

Traduction.

Côté séquences, nous posons, pour une cellule critique κ de dimension p :

$$\mathrm{St}(\kappa) = \{\nu \in K^{(p)} : \text{un gradient path joint } \nu \text{ à } \kappa\}, \quad \mathrm{Un}(\kappa) = \{\nu \in K^{(p)} : \text{un cogradient path joint } \kappa \text{ à } \nu\},$$

les chemins triviaux étant admis, de sorte que $\kappa \in \mathrm{St}(\kappa) \cap \mathrm{Un}(\kappa)$. Ces chemins sont exactement ceux de la section 6.3 [11] : seuls comptent désormais leur *existence*, et non plus leur parité.

Proposition 11 (Équivalence avec les ensembles de DFRS) *Pour toute cellule critique κ :* $\mathrm{St}(\kappa) = S_K(\kappa)$ *et* $\mathrm{Un}(\kappa) = U_K(\kappa)$.

Proof Montrons d'abord $S_K(\kappa) \subseteq \mathrm{St}(\kappa)$. Il suffit de vérifier que $\mathrm{St}(\kappa)$ contient κ et est clos pour la règle de DFRS, puisque $S_K(\kappa)$ est le plus petit ensemble ayant ces deux propriétés. Soit donc δ avec $V(\delta) = \gamma$ et $\beta \in \partial\gamma \cap \mathrm{St}(\kappa)$. Si $\beta = \delta$, la conclusion $\delta \in \mathrm{St}(\kappa)$ est immédiate. Si $\beta \neq \delta$, soit π un gradient path de β à κ ; alors $\langle \delta, \gamma, \beta \rangle$ suivi de π est un gradient path de δ à κ, donc $\delta \in \mathrm{St}(\kappa)$.

Réciproquement, soit $\delta \in \mathrm{St}(\kappa)$, joint à κ par un gradient path à r paires ; nous montrons $\delta \in S_K(\kappa)$ par récurrence sur r. Si $r = 0$, alors $\delta = \kappa \in S_K(\kappa)$. Si $r \geq 1$, écrivons le chemin $\langle \delta, \gamma, \delta_1, \ldots, \kappa \rangle$ avec $V(\delta) = \gamma$ et $\delta_1 \in \partial\gamma$, $\delta_1 \neq \delta$. Le chemin de δ_1 à κ a $r - 1$ paires, donc $\delta_1 \in S_K(\kappa)$ par hypothèse de récurrence, et la règle de DFRS appliquée avec $\beta = \delta_1$ donne $\delta \in S_K(\kappa)$.

L'énoncé instable se démontre à l'identique en remplaçant la règle par sa duale : si $V^{-1}(\delta) = \gamma$ (c'est-à-dire $(\gamma, \delta) \in V$) et si β est une cofacette de γ dans $\mathrm{Un}(\kappa)$ avec $\beta \neq \delta$,

alors un cogradient path de κ à β se prolonge en $\langle \dots, \beta, \gamma, \delta \rangle$, qui est un cogradient path valide puisque $\beta \in \delta(\gamma)$ et $\delta \neq \beta$. □

Proposition 12 (Existence contre parité) *Pour toute cellule critique κ :*

$$\operatorname{supp} \widehat{\Upsilon}(\kappa) \subseteq \mathrm{St}(\kappa), \qquad \operatorname{supp} \widehat{\Lambda}(\kappa) \subseteq \mathrm{Un}(\kappa).$$

Il y a égalité dès que tous les nombres de chemins concernés valent 0 ou 1, et les inclusions peuvent être strictes.

Proof Par le théorème de caractérisation par chemins de la section 6.3 [11, Th. 12], $\nu \in \operatorname{supp} \widehat{\Upsilon}(\kappa)$ si et seulement si le nombre de gradient paths de ν vers κ est impair. Un nombre impair est non nul, d'où la première inclusion ; la seconde est duale. Si tous les comptes valent 0 ou 1, « impair » équivaut à « non nul », d'où l'égalité. La strictude est réalisée par le champ V_2 de la figure 18(a), repris sur la figure 26 : deux gradient paths joignent AB à l'arête critique EH, donc $AB \in \mathrm{St}(EH)$ alors que $AB \notin \operatorname{supp} \widehat{\Upsilon}(EH)$, le compte 2 étant nul modulo 2. □

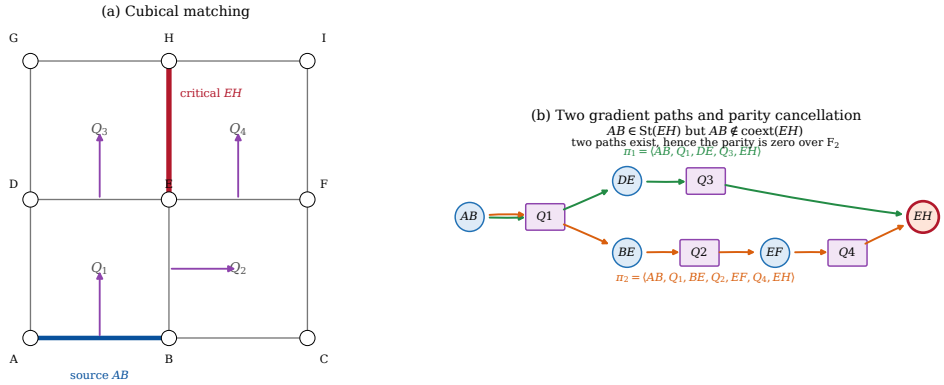


Fig. 26 Existence et parité diffèrent dès la dimension un. (a) Le champ V_2 de la figure 18 : quatre paires arête–carré, arête source AB, arête critique EH. (b) Les deux gradient paths de AB vers EH ; l'ensemble stable retient AB parce qu'un chemin existe, la coextension l'élimine parce que le nombre de chemins est pair.

Pourquoi c'est une contribution, et comment le vérifier.

Le cadre des séquences ne possédait que des objets de *parité* (γ, λ, $\widehat{\Upsilon}$, $\widehat{\Lambda}$) ; la mise en garde qui clôt la section 6.5 devient ici un énoncé précis. Pour tout refaire à la main sur la figure 26 : énumérer les gradient paths de AB vers EH en suivant, à chaque carré traversé, les trois sorties possibles ; on en trouve exactement deux, qui se séparent en Q_1 et se rejoignent en EH. Nous avons de plus vérifié la Proposition 11 par programme, en implantant côte à côte la règle récursive littérale de DFRS et le parcours des chemins, sur l'exemple guide et sur des champs aléatoires.

7.5 B6 : la dimension zéro et la partition des sommets (question N2)

Où c'est chez Robins.

C'est le Lemme 5 de [7, Sec. 3.3] : *« Each vertex of K is in the stable set of exactly one local minimum »*. Leur preuve exploite le fait que les $(0,1)$ V-paths *« can merge but not branch »* ; le branchement en dimensions supérieures est relevé dès [7, Sec. 2.4].

Hypothèses isolées.

Comme suggéré lors de la discussion sur N2, nous isolons les hypothèses réellement utilisées.

(H1) V est un couplage : l'arête appariée à un sommet, quand elle existe, est unique.
(H2) Une arête n'a que deux sommets : le sommet suivant du chemin est l'unique autre extrémité de l'arête appariée.
(H3) V est acyclique (Corollaire 10) : le chemin ne repasse par aucun sommet et termine.
(H4) Un sommet n'est jamais upper regular : dans une paire (σ, τ) on a $\dim \tau = \dim \sigma + 1 \geq 1$, donc un sommet ne peut être que l'élément inférieur d'une paire.

Theorem 13 (Flux déterministe en dimension zéro) *De chaque sommet v part un* unique *gradient path maximal ; il aboutit à un unique sommet critique $m(v)$. Par conséquent, pour tout minimum $m \in \mathcal{M}$, le nombre de gradient paths de v vers m vaut 1 si $m = m(v)$ et 0 sinon, et l'on a*

$$\mathrm{St}(m) = \operatorname{supp} \widehat{\Upsilon}_0(m), \qquad X^{(0)} = \biguplus_{m \in \mathcal{M}} \mathrm{St}(m).$$

Proof Construisons le chemin depuis v. Par (H4), v est soit critique, soit l'élément inférieur d'une paire (v, e). Dans le premier cas le chemin maximal est trivial et $m(v) = v$. Dans le second, l'arête $e = V(v)$ est unique par (H1), et le sommet suivant est l'unique autre extrémité de e par (H2) : chaque pas est forcé, donc le chemin est unique. Par le Corollaire 10, les pas d'insertion des sommets visités décroissent strictement, donc les sommets sont deux à deux distincts et le chemin termine, en un sommet non apparié, c'est-à-dire critique par (H4) : c'est $m(v)$.

Comptons maintenant les chemins. Tout gradient path issu de v est, chaque pas étant forcé, un préfixe du chemin maximal. Les sommets intermédiaires du chemin maximal sont lower regular (le chemin y continue) ; son seul sommet critique est donc son extrémité $m(v)$. Un gradient path de v vers un sommet critique m est un préfixe se terminant en un sommet critique, donc le chemin maximal tout entier, et $m = m(v)$: le compte vaut 1 pour $m(v)$ et 0 pour tout autre minimum.

Enfin, existence et parité coïncident puisque tous les comptes valent 0 ou 1 : la Proposition 12 donne $\mathrm{St}(m) = \operatorname{supp} \widehat{\Upsilon}_0(m)$. Et comme chaque sommet v appartient à $\mathrm{St}(m(v))$ et à aucun autre $\mathrm{St}(m)$, les ensembles stables des minima partitionnent $X^{(0)}$, ce qui est exactement le Lemme 5 de [7]. □

Remark 7 (Où le raisonnement s'arrête) (H2) et (H4) tombent en dimension au moins un : un carré possède quatre arêtes, un chemin qui y entre a trois sorties, les chemins se branchent

puis peuvent s'annuler modulo 2 (figure 26). L'égalité entre existence et parité est donc un phénomène propre à la dimension zéro. Sur l'exemple guide, la partition se lit sur les sommets colorés du panneau (b) de la figure 25 : classes de tailles 14, 10 et 6.

7.6 B8 : sous-complexes hermétiques, préfixes et coupes

Où c'est chez Robins.

Ce sont les Lemmes 1 à 3 de [7, Sec. 3.1] : les sous-complexes de niveau sont hermétiques (Lemme 1), les sous-complexes hermétiques préservent les ensembles stables et instables (Lemme 2), et la combinaison des deux (Lemme 3). L'intérêt pour l'image vient de [6, Sec. 2.3] : les level subcomplexes $K(c)$ forment une filtration, et le Théorème 11 de [6] relie les cellules critiques aux changements d'homologie des lower level cuts.

7.6.1 Définition exacte de DFRS (sous-complexe hermétique)

Un sous-complexe H de K est *hermétique* si, pour toute cellule $\alpha \in H$, on a $V(\alpha) = 0$ ou $V(\alpha) \in H$ [7, Sec. 3.1]. Comme H est clos par faces, l'implication réciproque est automatique ($\tau \in H$ et $\sigma < \tau$ donnent $\sigma \in H$) : une paire régulière est toujours entièrement dans H ou entièrement hors de H. Nous notons V_H l'ensemble des paires de V contenues dans H ; les cellules de H critiques pour V_H sont alors exactement les cellules critiques de V appartenant à H.

7.6.2 Définitions exactes de Bertrand (stack, coupe, F-séquence, flooding)

Soit F une fonction de K vers $\mathbb{Z}$. F est un *stack* si $F(\sigma) \leq F(\tau)$ dès que $\sigma \subseteq \tau$; la *coupe* de F au niveau λ est $F_\lambda = \{\nu \in K \mid F(\nu) \leq \lambda\}$ [12, Sec. 2.3]. La max-extension de la section 4.2 est un stack. Une séquence de Morse est une *F-séquence* si $F(\sigma) = F(\tau)$ pour toute paire régulière (σ, τ) [12, Def. 4] : aucune paire ne traverse un niveau. Elle est une *flooding sequence* si $F(\sigma) \leq F(\tau)$ dès que $\sigma \in K_i$, $\tau \in K_j$ et $i < j$ [12, Def. 5] ; dans ce cas chaque coupe F_λ est un préfixe K_i de la séquence [12, Prop. 7].

Lemma 14 (Les préfixes sont hermétiques) *Pour tout i, le préfixe K_i de la séquence est hermétique.*

Proof Soit $\sigma \in K_i$ avec $V(\sigma) = \tau$, la paire étant insérée au pas $j = t(\sigma) = t(\tau)$. Puisque $\sigma \in K_i$, on a $j \leq i$, donc $\tau \in K_j \subseteq K_i$. □

Lemma 15 (Les coupes d'une F-séquence sont hermétiques) *Si $\overrightarrow{W}$ est une F-séquence, alors pour tout $\lambda \in \mathbb{Z}$ la coupe F_λ est un sous-complexe hermétique.*

Proof F_λ est clos par faces : si $\sigma \le \nu \in F_\lambda$, la monotonie du stack donne $F(\sigma) \le F(\nu) \le \lambda$. Il est hermétique : si $\sigma \in F_\lambda$ et $V(\sigma) = \tau$, la condition de F-séquence donne $F(\tau) = F(\sigma) \le \lambda$, donc $\tau \in F_\lambda$. □

Theorem 16 (Restriction hermétique) *Soient $H \subseteq K$ hermétique et $\kappa \in H$ une cellule critique. Alors*

$$\mathrm{St}_H(\kappa) = \mathrm{St}_K(\kappa) \cap H, \qquad \mathrm{Un}_H(\kappa) = \mathrm{Un}_K(\kappa),$$

où les membres de gauche sont calculés dans H avec le champ V_H. En particulier $\mathrm{Un}_K(\kappa) \subseteq H$.

Proof Un chemin de V_H est un chemin de V contenu dans H : cela donne $\mathrm{St}_H(\kappa) \subseteq \mathrm{St}_K(\kappa) \cap H$ et $\mathrm{Un}_H(\kappa) \subseteq \mathrm{Un}_K(\kappa)$.

Réciproquement, soit $\nu \in \mathrm{St}_K(\kappa) \cap H$ et soit $\pi = \langle \sigma_0 = \nu, \tau_0, \sigma_1, \ldots, \sigma_r = \kappa \rangle$ un gradient path dans K. Montrons par récurrence que π est contenu dans H : $\sigma_0 \in H$; si $\sigma_i \in H$, l'hermétisme donne $\tau_i = V(\sigma_i) \in H$, puis la clôture par faces donne $\sigma_{i+1} \in \partial\tau_i \subseteq H$. Ainsi π est un chemin de V_H et $\nu \in \mathrm{St}_H(\kappa)$.

De même, soit $\nu \in \mathrm{Un}_K(\kappa)$ et $\pi = \langle \tau_0 = \kappa, \sigma_1, \tau_1, \ldots, \tau_r = \nu \rangle$ un cogradient path dans K. On a $\tau_0 = \kappa \in H$; si $\tau_{i-1} \in H$, alors σ_i, facette de τ_{i-1}, appartient à H par clôture, puis $\tau_i = V(\sigma_i) \in H$ par hermétisme. Le chemin est donc contenu dans H, d'où $\nu \in \mathrm{Un}_H(\kappa)$, et au passage $\mathrm{Un}_K(\kappa) \subseteq H$. □

Corollary 17 (Cohérence aux coupes) *Soit $\overrightarrow{W}$ une F-séquence. Pour tout niveau λ et toute cellule critique $\kappa \in F_\lambda$:*

$$\mathrm{St}_{F_\lambda}(\kappa) = \mathrm{St}_K(\kappa) \cap F_\lambda, \qquad \mathrm{Un}_{F_\lambda}(\kappa) = \mathrm{Un}_K(\kappa).$$

Le même énoncé vaut pour tout préfixe K_i contenant κ.

Proof Combiner les Lemmes 14 et 15 avec le Théorème 16. □

Remark 8 (Plateaux et injectivité) Chez DFRS, l'hermétisme des sous-complexes de niveau est déduit des valeurs d'une fonction de Morse discrète supposée injective [7, Lem. 1 et Sec. 3]. Ici, seule la structure de la séquence intervient : un stack peut porter des plateaux entiers au même niveau [12], et aucune perturbation des valeurs n'est nécessaire (section 4.3). Pour une flooding sequence, les coupes sont de plus des préfixes [12, Prop. 7] : les deux familles hermétiques coïncident, et changer de seuil revient à s'arrêter plus ou moins loin dans la même séquence.

7.7 B11 : le bassin comme fermeture par expansions

Où c'est chez Robins.

Dans [7, Sec. 3.3], après le Lemme 5. Trois définitions exactes sont nécessaires.

7.7.1 Définitions exactes de DFRS (collapse régulier, bassin, pont)

Si (α, β) est une paire libre de K telle que $V(\alpha) = \beta$, le collapse élémentaire qui retire α et β est dit *régulier* ; une suite de collapses élémentaires réguliers est un *collapse régulier*, noté ici $C \searrow_V H$ [7, Sec. 3.2]. Pour un minimum m, le *bassin* $B_K(m)$ est *« the maximal subcomplex of K that has a regular collapse onto m »* ; l'unicité de ce sous-complexe maximal est le Lemme 6 de [7], sa restriction aux hermétiques, $B_H(m) = B_K(m) \cap H$, le Lemme 7, et DFRS notent que les bassins sont simplement connexes, deux à deux disjoints, sans autre cellule critique que leur minimum, avec $S_K(m) \subset B_K(m)$. Une 1-cellule contenue dans aucun bassin est un *pont* (bridge) [7, Sec. 3.3]. La distinction de la section 5.2 est ici essentielle : $\mathrm{St}(m)$ est un ensemble de *sommets*, $B_K(m)$ un *sous-complexe*, et la relation attendue est $B_K(m)^{(0)} = \mathrm{St}(m)$, et non $B_K(m) = \mathrm{St}(m)$.

7.7.2 Notions proposées (ensemble requis, paire disponible, basin sequence)

Les notions suivantes n'apparaissent ni chez Robins ni chez Bertrand. Pour une paire régulière $q = (\sigma, \tau)$ de V, nous posons

$$\mathrm{req}(q) = \partial\sigma \ \cup \ \big(\partial\tau \setminus \{\sigma\}\big),$$

l'*ensemble requis* de q : les faces qui doivent être présentes pour que l'ajout de σ et τ soit une expansion élémentaire (section 2.4). Si A est un sous-complexe ne contenant ni σ ni τ, la paire q est *disponible pour* A quand $\mathrm{req}(q) \subseteq A$; dans ce cas $A \cup \{\sigma, \tau\}$ est bien une expansion élémentaire de A : toutes les faces requises sont présentes, et aucune cofacette de σ n'appartient à A, sans quoi la clôture par faces forcerait $\sigma \in A$.

Definition 36 (Basin sequence — notion proposée) Une *basin sequence* enracinée en un sommet critique m est une chaîne de sous-complexes

$$\{m\} = B_0 \subseteq B_1 \subseteq \cdots \subseteq B_r$$

telle que chaque $B_{i+1} = B_i \cup \{\sigma_i, \tau_i\}$ s'obtient de B_i par l'ajout d'une paire $(\sigma_i, \tau_i) \in V$ disponible pour B_i. Elle est *maximale* si aucune paire de V n'est disponible pour B_r.

Theorem 18 (Fermeture maximale unique) *Toutes les basin sequences maximales enracinées en m ont le même sous-complexe terminal, noté $B(m)$. De plus,*

$$B(m) = \max_{\subseteq} \ \{\, C \subseteq K \ : \ C \searrow_V \{m\} \,\},$$

c'est-à-dire que $B(m)$ contient tout sous-complexe qui se collapse régulièrement sur $\{m\}$ et est lui-même un tel sous-complexe : $B(m)$ est exactement le bassin $B_K(m)$ de DFRS.

Proof Soit B le terminal d'une basin sequence maximale $\{m\} = B_0 \subseteq \cdots \subseteq B_r = B$.

Étape 1 : $B \searrow_V \{m\}$. Considérons un pas $B_{i+1} = B_i \cup \{\sigma_i, \tau_i\}$ et montrons que (σ_i, τ_i) est une paire libre de B_{i+1}. Soit τ' une cofacette de σ_i dans B_{i+1} : si $\tau' \in B_i$, la clôture par faces de B_i donnerait $\sigma_i \in B_i$, ce qui est exclu ; donc $\tau' \in \{\sigma_i, \tau_i\}$, et pour raison de

dimension $\tau' = \tau_i$. Retirer la paire est donc un collapse élémentaire régulier, et la chaîne lue de droite à gauche est un collapse régulier de B sur $\{m\}$.

Étape 2 : B est hermétique. Les cellules de B sont m et des paires complètes de V : si $\sigma \in B$ est l'élément inférieur d'une paire, son partenaire a été ajouté au même pas.

Étape 3 : B contient tout C avec $C \searrow_V \{m\}$. Renversons un collapse régulier de C sur $\{m\}$:

$$\{m\} = C_0 \subset C_1 \subset \cdots \subset C_s = C, \qquad C_{j+1} = C_j \cup \{\sigma_j, \tau_j\},\ (\sigma_j, \tau_j) \in V.$$

Montrons $C_j \subseteq B$ par récurrence, le cas $j = 0$ étant clair. Supposons $C_j \subseteq B$. Le complexe C_{j+1} est clos par faces, donc $\partial\sigma_j$ et $\partial\tau_j \setminus \{\sigma_j\}$ sont contenus dans C_{j+1} ; ces cellules étant distinctes de σ_j et τ_j, elles appartiennent à C_j, d'où $\mathrm{req}(\sigma_j, \tau_j) \subseteq C_j \subseteq B$. Par l'Étape 2 et la clôture par faces, la paire (σ_j, τ_j) est entièrement dans B ou entièrement hors de B. Dans le second cas elle serait disponible pour B, ce qui contredirait la maximalité de la basin sequence. Donc $C_{j+1} \subseteq B$.

Étape 4 : unicité. Si B et B' sont les terminaux de deux basin sequences maximales, l'Étape 1 et l'Étape 3 donnent $B' \subseteq B$ et $B \subseteq B'$.

Enfin, B réalise le maximum de l'énoncé par les Étapes 1 et 3, et ce maximum est précisément la définition du bassin de DFRS. □

Remark 9 (Le point délicat) Chez DFRS, l'unicité du sous-complexe maximal (leur Lemme 6) est établie en comparant les valeurs $f(\alpha)$ d'une fonction de Morse discrète injective. Cet argument n'est pas disponible ici : il est remplacé par le couple hermétisme–disponibilité de l'Étape 3, qui n'utilise que la structure d'expansion de la séquence. C'est, avec le Lemme 9, le principal endroit où la traduction demande un argument nouveau plutôt qu'un transport.

Proposition 19 (Propriétés des bassins) *Pour tout $m \in \mathcal{M}$:*

1. *$B(m)$ est hermétique et $B(m) \searrow_V \{m\}$; en particulier $B(m)$ est contractile ;*
2. *m est la seule cellule critique de $B(m)$;*
3. *$B(m)^{(0)} = \mathrm{St}(m)$;*
4. *les bassins de deux minima distincts sont disjoints.*

Proof (1) et (2) : Étapes 1 et 2 de la preuve précédente ; toute cellule de $B(m) \setminus \{m\}$ appartient à une paire de V, donc n'est pas critique, et un complexe qui se collapse sur un point est contractile.

(3, $\subseteq$) Soit v un sommet de $B(m)$ et suivons son chemin maximal (Théorème 13) : si v est lower regular, l'hermétisme place $V(v)$ dans $B(m)$, puis la clôture par faces y place le sommet suivant ; par récurrence le chemin reste dans $B(m)$ et son extrémité critique $m(v)$ appartient à $B(m)$, donc $m(v) = m$ par (2), d'où $v \in \mathrm{St}(m)$.

(3, $\supseteq$) Soit $v \in \mathrm{St}(m)$ et $v = v_0 < e_0 > v_1 < \cdots > v_r = m$ son unique chemin, avec $V(v_i) = e_i$. Montrons, par récurrence descendante sur i, que v_i et e_i appartiennent à $B(m)$. On a $v_r = m \in B(m)$. Supposons $v_{i+1} \in B(m)$. L'ensemble requis de la paire (v_i, e_i) est $\mathrm{req}(v_i, e_i) = \partial e_i \setminus \{v_i\} = \{v_{i+1}\} \subseteq B(m)$. Si la paire n'était pas contenue dans $B(m)$, elle serait entièrement hors de $B(m)$ (hermétisme et clôture), donc disponible pour $B(m)$, contredisant la maximalité. Donc $v_i, e_i \in B(m)$.

(4) Une cellule commune à $B(m)$ et $B(m')$ possède un sommet v (clôture par faces), qui appartient alors à $\mathrm{St}(m) \cap \mathrm{St}(m')$ par (3), d'où $m = m(v) = m'$ par le Théorème 13. □

La figure 27 montre une basin sequence sur l'exemple guide ; chaque étape se contrôle à la main en vérifiant l'ensemble requis de la paire ajoutée.

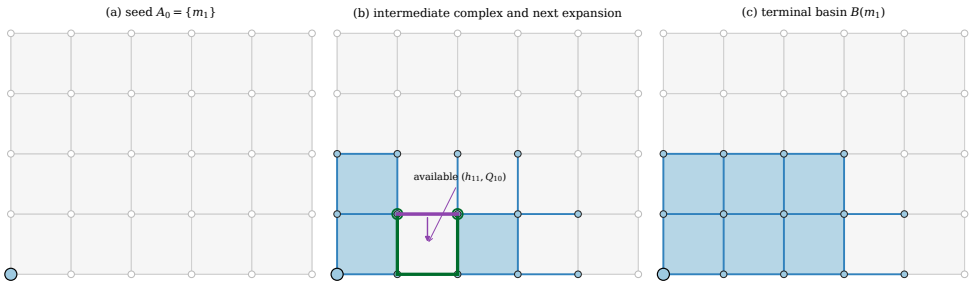


Fig. 27 Une basin sequence construisant $B(m_1)$ sur l'exemple guide. (a) Le germe $\{m_1\}$. (b) Un sous-complexe intermédiaire ; la paire (h_{11}, Q_{10}) est disponible car les deux extrémités de h_{11} et les trois autres faces de Q_{10} sont déjà présentes. (c) Le terminal, unique quel que soit l'ordre des expansions (Théorème 18) : 14 sommets, 19 arêtes, 6 carrés. Lue de droite à gauche, la construction est un collapse régulier sur m_1, ce qui est exactement la définition du bassin de DFRS.

7.7.3 Un algorithme en un seul balayage

L'ordre de la séquence ordonnance tous les bassins à la fois. Pour une cellule ν, l'étiquette $\mathrm{lab}(\nu)$ est soit un sommet critique, soit la valeur indéfinie $\perp$.

Algorithm 6 Tous les bassins en un seul balayage de la séquence

Require: la forme cellule par cellule $\diamond\overrightarrow{W} = \langle \kappa_1, \ldots, \kappa_N \rangle$
Ensure: $\mathrm{lab} : K \to \mathcal{M} \cup \{\perp\}$ avec $\{\nu : \mathrm{lab}(\nu) = m\} = B(m)$

$\mathrm{lab}(\nu) \leftarrow \perp$ pour toute cellule ν
for $i \leftarrow 1$ **to** N **do**
  if $\kappa_i = \nu$ est critique **then**
    if $\dim \nu = 0$ **then**
      $\mathrm{lab}(\nu) \leftarrow \nu$
    end if
  else
    écrire $\kappa_i = (\sigma, \tau)$
    if toutes les cellules de $\mathrm{req}(\sigma, \tau)$ portent la même étiquette définie m **then**
      $\mathrm{lab}(\sigma) \leftarrow m$; $\mathrm{lab}(\tau) \leftarrow m$
    end if
  end if
end for
return lab

Theorem 20 (Correction de l'algorithme) *Pour tout $m \in \mathcal{M}$, l'Algorithme 6 termine avec*

$$\{\nu \in K : \mathrm{lab}(\nu) = m\} = B(m),$$

et $\mathrm{lab}(\nu) = \bot$ *exactement quand* ν *n'appartient à aucun bassin.*

Proof Remarquons d'abord que l'algorithme est bien posé : chaque cellule est insérée une seule fois, donc étiquetée au plus une fois, et par le Lemme 9 toute cellule de $\mathrm{req}(\sigma, \tau)$ est insérée strictement avant le pas de la paire — les facettes de τ autres que σ par le point (1), et les facettes de σ parce que $K_{i-1} \cup \{\sigma, \tau\}$ est clos par faces. Au moment où la paire est examinée, les étiquettes de son ensemble requis sont donc définitives.

Notons $L_m = \{\nu : \mathrm{lab}(\nu) = m\}$ à la fin du balayage.

$L_m \subseteq B(m)$. Les affectations d'étiquette m, prises dans l'ordre du balayage, forment une basin sequence enracinée en m : la première est m lui-même, et chaque paire étiquetée m a son ensemble requis déjà étiqueté m, donc contenu dans le sous-complexe courant, ce qui en fait une paire disponible. Le terminal de cette chaîne se collapse régulièrement sur $\{m\}$ (Étape 1 du Théorème 18), donc est contenu dans $B(m)$ (Étape 3).

$B(m) \subseteq L_m$. Montrons par récurrence sur le pas n que toute cellule de $B(m)$ insérée à un pas $\leq n$ est étiquetée m. Si le pas n insère une cellule critique de $B(m)$, c'est m lui-même (Proposition 19(2)), étiqueté m par la ligne 4. Si le pas n insère une paire $(\sigma, \tau) \subseteq B(m)$, alors $\mathrm{req}(\sigma, \tau) \subseteq B(m)$ (clôture par faces), et chaque cellule de cet ensemble, insérée avant n, est étiquetée m par hypothèse de récurrence ; la ligne 8 étiquette donc σ et τ par m. Comme toute cellule de $B(m)$ est m ou appartient à une paire de V contenue dans $B(m)$ (hermétisme), la récurrence couvre tout $B(m)$.

Enfin les L_m sont exactement les bassins, deux à deux disjoints (Proposition 19(4)), donc les cellules restées à $\bot$ sont celles qui n'appartiennent à aucun bassin. □

Proposition 21 (Complexité) *L'Algorithme 6 s'exécute en temps* $O\big(|K| + \sum_{q \in V} |\mathrm{req}(q)|\big)$ *et en mémoire additionnelle* $O(|K|)$. *En dimension fixée* d *(complexes cubiques ou simpliciaux),* $\sum_{q \in V} |\mathrm{req}(q)| = O(d\,|K|)$, *et l'algorithme peut s'exécuter* en ligne, *pendant la construction de la séquence.*

Proof Chaque élément de la séquence est examiné une fois ; pour une paire, on inspecte une fois chaque cellule de son ensemble requis, le reste étant en temps constant ; une étiquette est stockée par cellule. Une p-cellule cubique a $2p$ facettes (une p-cellule simpliciale en a $p+1$), donc $|\mathrm{req}(\sigma, \tau)| \leq 2p + 2(p+1) - 1$ est borné par une fonction de d. L'exécution en ligne est possible car, par le Lemme 9, l'ensemble requis d'une paire est entièrement inséré avant elle. □

Theorem 22 (Restriction des bassins) *Soient $H \subseteq K$ hermétique et $m \in H$ un sommet critique. Alors*

$$B_H(m) = B(m) \cap H,$$

où $B_H(m)$ est le bassin calculé dans H avec le champ V_H.

Proof ($\subseteq$) $B_H(m)$ se collapse régulièrement sur $\{m\}$ dans H, donc dans K, et le Théorème 18 donne $B_H(m) \subseteq B(m)$; comme $B_H(m) \subseteq H$, l'inclusion suit.

($\supseteq$) Prenons un collapse régulier $B(m) = C_0 \searrow C_1 \searrow \cdots \searrow C_s = \{m\}$ et intersectons chaque étage avec H. La paire (α, β) retirée à l'étage j est entièrement dans H ou entièrement hors de H (hermétisme et clôture). Hors de H : l'intersection ne change pas. Dans H : la paire appartient à V_H, et elle est libre dans $C_j \cap H$, car une cofacette de α dans $C_j \cap H$ serait une cofacette dans C_j, donc égale à β. Chaque étage est ainsi soit inchangé, soit un collapse élémentaire régulier de V_H, et l'on obtient $B(m) \cap H \searrow_{V_H} \{m\}$. Le Théorème 18 appliqué dans H donne $B(m) \cap H \subseteq B_H(m)$. □

Corollary 23 (Bassins et coupes) *Pour une F-séquence, tout niveau λ et tout minimum $m \in F_\lambda$:*

$$B_{F_\lambda}(m) = B(m) \cap F_\lambda.$$

Le même énoncé vaut pour tout préfixe K_i contenant m. Un unique balayage fournit donc, par simple restriction des étiquettes, la famille des bassins de toutes les coupes, sans recalcul (figure 28).

Proof Lemmes 14 et 15 avec le Théorème 22. □

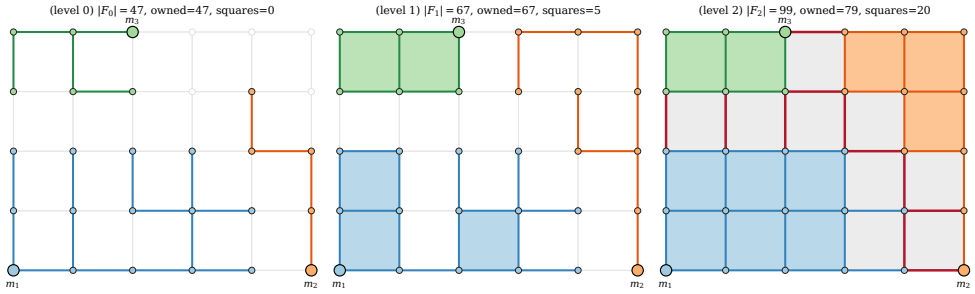


Fig. 28 Restriction de la famille de bassins aux coupes d'un stack à trois niveaux sur l'exemple guide ($|F_0| = 47$, $|F_1| = 67$, $|F_2| = 99$ cellules). Les cellules absentes de la coupe sont en gris très clair ; les cellules présentes hors de tout bassin sont en gris avec les ponts en rouge. Les mêmes trois étiquettes globales servent à tous les niveaux : les bassins de chaque coupe sont exactement les intersections $B(m) \cap F_\lambda$ (Corollaire 23).

Pourquoi c'est une contribution, et comment le vérifier.

Chez DFRS le bassin est défini par maximalité et calculé par parcours du champ ; la caractérisation par point fixe d'expansions (Théorème 18), l'indépendance de l'ordre et le balayage unique (Théorème 20) n'y figurent pas : c'est précisément la structure que la séquence ajoute au champ de gradient. Les cellules restées à $\perp$ sont obtenues gratuitement ; en dimension un, ce sont les ponts de DFRS, point de départ de la traduction ultérieure de leur Lemme 8. Pour tout refaire à la main sur la figure 27

: vérifier à chaque étape que la paire ajoutée a son ensemble requis dans le sous-complexe courant, puis lire la chaîne à l'envers et vérifier qu'à chaque retrait la paire est libre. Nous avons de plus vérifié par programme, sur l'exemple guide et sur des champs aléatoires : l'égalité entre le terminal de la fermeture (calculée sous deux ordres aléatoires différents) et la sortie de l'Algorithme 6, les quatre propriétés de la Proposition 19, la liberté de chaque paire lors du collapse inverse, l'absence de paire disponible au terminal, et les restrictions du Théorème 22 et du Corollaire 23 sur des préfixes et des coupes tirés au hasard.

7.8 Hors périmètre : ponts, squelette, cancellation (question N5)

Trois morceaux restent volontairement hors du paquet B5–B11. Les *ponts* demandent l'étude des canonical paths entre bassins [7, Lem. 8, Th. 9] ; les cellules à $\perp$ de l'Algorithme 6 en fournissent le point de départ. Le *squelette de Morse* est l'union des clôtures des ensembles instables [7, Sec. 3.2] et devrait suivre de B5 et B8. La *cancellation*, enfin, appelle la distinction proposée lors de la discussion sur N5 : la cancellation de Forman [2] transforme un champ acyclique en champ acyclique, donc préserve les séquences de Morse ; mais elle ne préserve pas, en général, la propriété d'être une F-séquence pour le stack initial, puisque le chemin inversé n'a aucune raison de rester à niveau constant. Nous prévoyons de développer d'abord le niveau « séquence de Morse », et de garder pour plus tard la question, différente et sans doute plus difficile, d'un stack modifié F' adapté à la séquence simplifiée.

8 Conclusion and perspectives

Ce rapport a suivi deux fils. La Partie I fixe le modèle discret et présente, avec des exemples bidimensionnels vérifiés pas à pas, les deux cadres en présence : la construction d'un champ de gradient à partir des niveaux de gris puis la lecture des objets de segmentation sur sa dynamique [6, 7], et les séquences de Morse, où l'acyclicité est acquise par construction et où les applications de référence, coréférence, extension et coextension font passer l'information entre le complexe et ses cellules critiques [11, 12].

La Partie II traduit la couche segmentation dans le second cadre. Les ensembles stables et instables, définis par existence de chemins, coïncident avec les ensembles récursifs de DFRS, tandis que les objets de parité restent strictement plus petits dès que les chemins se branchent ; en dimension zéro les deux notions se confondent et l'on retrouve la partition des sommets par les minima. Les préfixes de la séquence et les coupes d'un stack sont hermétiques, ce qui donne une segmentation à tous les niveaux de gris à partir d'un seul calcul, sans hypothèse d'injectivité et donc sans perturbation des plateaux. Enfin le bassin, objet extrémal chez DFRS, se caractérise comme l'unique fermeture maximale d'un minimum par expansions régulières ; l'ordre porté par la séquence ordonnance alors tous les bassins en un seul balayage de coût linéaire.

Deux propriétés du cadre expliquent ces résultats. Le lemme d'ordre remplace la fonction de Morse discrète injective utilisée dans les preuves classiques, et un stack n'a pas besoin d'être injectif, de sorte que les plateaux sont traités comme des niveaux

entiers. C'est aussi ce qui rend les arguments plus courts que leurs contreparties classiques.

Trois directions restent ouvertes. Les ponts, que l'algorithme fournit gratuitement, doivent être reliés à l'adjacence des bassins par les canonical paths de [7, Lem. 8, Th. 9]. Le squelette de Morse devrait suivre des théorèmes de restriction, par un argument analogue à celui de la fermeture maximale. La cancellation, en revanche, pose une difficulté d'une autre nature : elle préserve les séquences de Morse mais pas, en général, les séquences adaptées à un stack donné, et l'existence d'un stack modifié convenant à la séquence simplifiée reste à établir.

References

[1] Gonzalez, R.C., Woods, R.E.: Digital Image Processing, 4th edn. Pearson, New York, NY (2018)

[2] Forman, R.: Morse theory for cell complexes. Advances in Mathematics **134**(1), 90–145 (1998)

[3] Forman, R.: A user's guide to discrete Morse theory. Séminaire Lotharingien de Combinatoire **48**, 48 (2002)

[4] Giblin, P.: Graphs, Surfaces and Homology. Cambridge University Press, Cambridge (2010)

[5] Scoville, N.A.: Discrete Morse Theory vol. 90. American Mathematical Society, Providence, RI, USA (2019)

[6] Robins, V., Wood, P.J., Sheppard, A.P.: Theory and algorithms for constructing discrete Morse complexes from grayscale digital images. IEEE Trans. Pattern Anal. Mach. Intell. **33**(8), 1646–1658 (2011)

[7] Delgado-Friedrichs, O., Robins, V., Sheppard, A.P.: Skeletonization and partitioning of digital images using discrete Morse theory. IEEE Trans. Pattern Anal. Mach. Intell. **37**(3), 654–666 (2015)

[8] Couprie, M., Bezerra, F.N., Bertrand, G.: Topological operators for grayscale image processing. J. Electron. Imaging **10**(4), 1003–1015 (2001)

[9] Bertrand, G., Boutry, N., Najman, L.: Discrete Morse functions and watersheds. J. Math. Imaging Vis. **65**(5), 787–801 (2023)

[10] Edelsbrunner, H., Harer, J., Zomorodian, A.: Hierarchical Morse–Smale complexes for piecewise linear 2-manifolds. In: Proceedings of the 19th Annual Symposium on Computational Geometry, pp. 70–79. ACM, New York, NY (2003)

[11] Bertrand, G.: Morse Sequences: A Simple Approach to Discrete Morse Theory. arXiv:2410.14227, 2024–2025 (2024)

[12] Bertrand, G.: Morse Sequences on Stacks and Flooding Sequences. arXiv:2509.01384 (2025)

[13] Kovalevsky, V.A.: Finite topology as applied to image analysis. Computer Vision, Graphics, and Image Processing **46**(2), 141–161 (1989)