

Bounded Repair Reachability of Residual Critical Blocks in 3D Binary Masks under SAFE-Gated Single-Voxel Flips

Jihun Bae

`jhbae1184@akane.waseda.jp`

Waseda University

Yeonho Bae

Dongguk University

Jinglu Hu

Waseda University

Research Article

Keywords: Digital well-composedness, digital topology, critical configurations, bounded reachability, foreground fidelity, voxel repair

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11416678/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Bounded Repair Reachability of Residual Critical Blocks in 3D Binary Masks under SAFE-Gated Single-Voxel Flips

Jihun Bae^{1*}, Yeonho Bae² and Jinglu Hu¹

^{1*}Graduate School of Information, Production, and Systems, Waseda University, Kitakyushu, Japan.

²Department of Statistics and Data Science, Dongguk University, Seoul, Korea.

*Corresponding author(s). E-mail(s): jhbae1184@akane.waseda.jp;

Abstract

Critical configurations in three-dimensional binary images violate digital well-composedness, and voxel repairs can conflict with path constraints and foreground preservation. We formulate bounded repair reachability under single-voxel flips preserving the homotopy type of the foreground complex. We establish relations among admissible path sets and show that, under flip-once, endpoint and pathwise fidelity yield identical successful instances and shortest successful depths. Breadth-first search distinguishes found repairs, bounded absence and resource-limited outcomes. On 680 residual blocks from 222 Brain Tumor Segmentation (BraTS) 2021 masks, unrestricted searches within radius 2 and depth 6 find repairs for 72 blocks under endpoint fidelity and 15 under pathwise fidelity. An endpoint-faithful witness temporarily removes and restores input foreground; a completed search excludes pathwise-faithful repair of that block within these bounds. These results distinguish the topological validity of moves from the admissibility of repair paths. Conclusions are limited to the stated windows, horizons and moves.

Keywords: Digital well-composedness, digital topology, critical configurations, bounded reachability, foreground fidelity, voxel repair

1 Introduction

1.1 Background

Binary segmentation masks are routinely converted into surfaces, skeletons and topological descriptors, and the outcome depends on how foreground and background voxels touch. A 3D binary image is well-composed when the square faces shared by foreground and background voxels form a 2D manifold. For such images a single connectedness notion serves for the 6-, 18- and

26-adjacency, and surface extraction and thinning are simplified [1]. In two dimensions, well-composedness gives a Jordan curve theorem, a locally computable Euler characteristic and the coincidence of 4- and 8-connectedness [2]. In three dimensions it was defined through the manifoldness of the boundary of the continuous analog [3, 4] and is characterised digitally by critical configurations [3, Prop. 2.1; 1; 4]. Digital well-composedness is local: no block of the image contains a primary or secondary critical configuration [4]. On cubical grids it coincides with P-well-composedness, which is defined for partially ordered sets [5].

Masks produced by classifiers or annotators contain critical configurations. Repairing them means changing voxel values so that these configurations disappear while the object keeps its topology and stays close to the input. The two aims can conflict: a change that removes a critical configuration may also merge components, close a tunnel or fill a cavity.

1.2 Prior work and the open question

Digital topology supplies single-voxel changes that preserve topology. A simple point is a point whose change leaves the topology of the image unchanged, and characterisations are available in two, three and four dimensions [6, 7], and in cubical-homology form for homology-simple points in any dimension and for homotopy-simple points in dimension three [8]. Rosenfeld, Kong and Nakamura show that any two topologically equivalent two-valued pictures can be transformed into one another by changing simple pixels, and that pictures can be magnified, translated or made well-composed by repeatedly changing simple pixels [9]. The well-composed images obtained with this operator have 9 and 27 times more points than the 2D and 3D input images [1].

Value-changing repair procedures differ in their stated guarantees (Sections 2.3 and 2.4). On the input grid, the guarantee is the well-composedness of the output alone [10] or none concerning the topology of the result [1]; a front-propagation procedure preserves the initial topology in its strict mode [11]; and other approaches obtain topological guarantees by changing the representation, through interpolation on a refined grid [4, 12] or polyhedral and cell complexes [13, 14]. For binary segmentation masks, the starting states analysed below come from a SAFE-gated variant of the authors' earlier add-only R^+ procedure, a binary adaptation of the queue-based correction of [10] that adds one background voxel to a critical block at a time and does not constrain the topology of the result. Its production-code provenance is documented in Online Resource 2, Section 2; the variant and its extension by removals are defined in Section 3.2.

Topology preservation and well-composedness have thus been combined in several ways: on a magnified grid [9; 4, Section 7.4], in a refined grid

or complex [4, 14], and, for a front-propagation procedure, voxel by voxel [11]. Same-grid repairs of an existing mask with a well-composed output state either well-composedness alone [10] or no guarantee of correct topology [1]. These works construct a repaired image. We consider a different situation, in which a deterministic repair built from SAFE-gated single-voxel changes has stopped with critical blocks left. For these blocks we ask which can still be resolved on the input grid by a bounded sequence of SAFE-gated flips, and how the answer depends on the rules imposed on the sequence. Three rules are considered: flips are restricted to voxels of currently critical blocks (locality); a voxel is flipped at most once (flip-once); and input foreground is not removed, either at every step or only at the end (pathwise and endpoint fidelity). Reachability results in digital topology include the existence of unbounded change sequences between equivalent pictures [9] and complexity results for reachability in 2D binary images [15]. Reachability in the present study is defined by the move graph of a single block (Section 3.5); its relation to the reachability problems of [15] has not been determined. A bounded search gives finite statements, namely that a repair is found or that none exists within a stated window and number of flips, and such statements do not address unbounded existence. In the sources examined for this work (Section 2), a separation of the way in which locality, flip-once and fidelity each determine which blocks remain reachable was not found.

1.3 Problem statement

The SAFE gate determines which individual flips preserve foreground homotopy type. Bounded repair reachability asks whether a sequence of such flips can meet a geometric target under additional path rules. Locality, flip-once and fidelity therefore enter the problem through admissibility and success, while the topology guarantee is shared by every system.

Let $\Omega \subset \mathbb{Z}^3$ be a finite rectangular voxel domain, with the image fixed to background outside Ω and flips restricted to Ω . Let $X \subseteq \Omega$ be a foreground and $K(X)$ the union of the closed unit cubes C_q of the voxels $q \in X$. For $v \in \Omega$ the attachment of v is $A_X(v) = K(X \setminus \{v\}) \cap C_v$, and a flip of v is SAFE if the Betti numbers of $A_X(v)$

are $(1, 0, 0)$. A SAFE flip preserves the homotopy type of $K(X)$ (Online Resource 1, Theorem 1). $\text{Crit}(X)$ is the set of critical blocks of X (Section 3.2); X is digitally well-composed if $\text{Crit}(X) = \emptyset$.

Let X_T be a start state, $B \in \text{Crit}(X_T)$ a target block, $R \geq 0$ and $L \geq 1$ integers, and $W = W_R(B)$ the window of Section 3.5. An admissibility rule $\mathcal{A}$ selects, at each state of a path, a subset of the SAFE flips of voxels of W (locality and flip-once, Section 3.4), and a fidelity criterion $\mathcal{F} \in \{F0, F1, F2\}$ restricts the removal of input foreground (Section 3.4).

Problem (bounded repair reachability).

Does a sequence of states $X_T = Y_0, Y_1, \dots, Y_\ell$ with $\ell \leq L$ exist, where $Y_{i+1} = Y_i \triangle \{v_i\}$, $v_i \in W$, each flip is SAFE in Y_i and admissible under $\mathcal{A}$ and $\mathcal{F}$, and Y_ℓ satisfies $B \notin \text{Crit}(Y_\ell)$ and $\text{Crit}(Y_\ell) \setminus \text{Crit}(X_T) = \emptyset$ (for F1, additionally $\text{NRFR}(Y_\ell) = 0$)?

The state space has at most $2^{|W|}$ elements, so the problem is decidable by exhaustive enumeration. The question studied is how the set of instances with a positive answer depends on $\mathcal{A}$ and $\mathcal{F}$. A search ends with a positive answer and a shortest path (FOUND), with a negative answer (EXH or DL), or without an answer because a resource limit was reached (CAP); Section 3.6 states what each outcome establishes. Every state of a path found is homotopy equivalent to X_T as a foreground complex.

1.4 Research questions and approach

We compare these path rules on binary whole-tumour masks. Each starting state X_T is the output of a preceding deterministic repair stage with its own move set and flip-once rule (Section 3.2). The subsequent searches allow other move sets: any SAFE flip in the window for unrestricted systems, or SAFE flips in currently critical blocks for locality systems. Although no target-block voxel admits a SAFE flip at X_T , the locality move set includes voxels of other critical blocks and is non-empty in 85 of the 680 instances at G26 (Section 4.1). Breadth-first search asks whether an admissible sequence resolves the target without creating a new critical block. Three questions guide the comparison.

- **RQ1.** How does the number of blocks that can be repaired within the horizon change when locality, flip-once and fidelity are imposed separately and together?
- **RQ2.** Do flip-once and pathwise fidelity restrict the same sets of paths?
- **RQ3.** What do the repairing paths look like; in particular, do they act on voxels outside the block or remove and restore input foreground?

At G26 ($R = 2, L = 6$), the F0 searches found repairs in 666 blocks without locality or flip-once, 5 with locality alone, 636 with path-local flip-once alone, and none with both. The first and third counts leave 14 and 40 CAP outcomes, respectively; the combined cell has only EXH or DL outcomes (Section 4.1). The unrestricted effective counts under endpoint and pathwise fidelity are 72 and 15; the pathwise repairs all belong to cohort C2 (Section 4.3). The corresponding witness analysis distinguishes paths excluded by fidelity from paths excluded by flip-once.

The analysis uses 680 target blocks from 222 BraTS 2021 cases, 14 systems and two window-horizon settings, with single-voxel flips, $R \leq 2$ and $L \leq 6$. The topology guarantee concerns the closed-cube foreground complex with $(26,6)$ adjacency (Online Resource 1); its relation to existing simple-point criteria is discussed in Section 2.2. Conclusions concern these masks, move rules and search bounds. The outcome semantics and provenance of effective counts are given in Sections 3.6 and 3.7.

1.5 Contributions

1. **Bounded repair reachability under path constraints.** We formulate target-block repair over SAFE-gated single-voxel flips, separating locality and flip-once admissibility from endpoint and pathwise fidelity. Path-set inclusions describe the effect of nested rules (Section 3.4). Under flip-once, endpoint and pathwise fidelity give the same successful instances and shortest successful depth (Section 5.3).
2. **A computational comparison of constraint systems.** We evaluate 14 systems on 680 blocks at two settings (19,040 searches), with FOUND, EXH, DL and CAP retained as distinct outcomes. The results identify empty

initial move sets, bounded absence of admissible repairs and repairs found under less restrictive rules (Section 4; Online Resource 2).

3. **Path witnesses separating the constraints.** Replayed paths establish that flip-once and pathwise fidelity admit different paths in both directions. For one block, an endpoint-faithful removal-restoration path is accompanied by a completed search excluding every pathwise-faithful repair within the stated bounds (Sections 4.4 and 4.5; Online Resource 3, Table S1).

The homotopy-preserving SAFE condition supplies the common move guarantee (Section 2.2, Online Resource 1). The path-set relations and the flip-once fidelity equivalence follow from the stated rules; aggregate outcomes and witness-based separations are established by the computations and replays described in Sections 3.7 and 4.

2 Related Work

2.1 Well-composedness and its variants

Latecki, Eckhardt and Rosenfeld introduced well-composed sets, for which the Jordan curve theorem holds, the Euler characteristic is computed locally, 4- and 8-connectedness coincide, and thinning is an internal operation [2]. Latecki extended well-composedness to three dimensions with a definition based on the manifoldness of the boundary of the continuous analog [3, 4]. The tutorial of Boutry, Géraud and Najman reviews the variants and the methods for producing well-composed signals: digital well-composedness, defined by the absence of primary and secondary critical configurations in blocks, is local and self-dual, whereas the variant defined through connected components is global [4]. Later variants include continuous well-composedness, which implies digital well-composedness in n-D [16], strong Euler well-composedness of regular cell complexes, which implies digital well-composedness for cubical complexes of n-D pictures [17], weak well-composedness of cell complexes [14] and P-well-composedness of posets, which is equivalent to digital well-composedness

on cubical grids [5]. The blocks of this study are the local units of digital well-composedness.

2.2 Simple points and topology-preserving deformations

A simple point is a point whose deletion or insertion does not change the topology of an image. Kong characterises simple points in two, three and four dimensions by conditions on the attachment of the point to the rest of the image [6], and Bertrand and Malandain give a characterisation of three-dimensional simple points [7]. Niethammer, Kalies, Mischaikow and Tannenbaum characterise simple points through the homology of the cubical neighbourhood of the point: a point is homology-simple if and only if this neighbourhood is acyclic, and in dimension three a point is homotopy-simple if and only if the Betti numbers of the neighbourhood are those of a point [8]. The intersection of this neighbourhood with the cube of the point is the attachment $A_X(v)$ of Section 3.3, which is stated to be homotopy equivalent to the neighbourhood [8]; Online Resource 1 (Lemma 3) proves that the neighbourhood retracts onto the attachment, so that the Betti condition of the SAFE gate is the condition of [8] on the neighbourhood. In the closed-cube model with (26,6) adjacency, the SAFE condition of Section 3.3 is evaluated through the attachment of the voxel; it agrees with the topological-number test of the gate program on all 2^{26} neighbourhood configurations. Rosenfeld, Kong and Nakamura build on single changes of simple pixels: any two topologically equivalent two-valued pictures can be transformed into one another by such changes, and magnification, translation and the production of well-composed pictures are obtained by repeating them [9]. Malgouyres and More study the computational complexity of reachability in 2D binary images and of some basic problems of 2D digital topology [15]. Topology correction of segmented medical images has used simple points to enforce a prescribed topology on a fuzzy segmentation by a fast marching procedure [18].

These works establish that sequences of topology-preserving single changes exist and that reachability on binary images has a complexity-theoretic side. The present study reports finite search outcomes (a repair, exhaustion of the

reachable component, or absence of a repair within L flips) for a fixed set of 680 blocks under explicit move constraints. In addition, the output stays on the input grid, whereas the well-composed pictures of [9] are obtained after magnification [4, Section 7.4], and the setting is three-dimensional, whereas the complexity results of [15] concern 2D images; whether the two notions of reachability are related has not been determined.

2.3 Value-changing repair and topologically constrained propagation

Siqueira et al. give a randomized algorithm that turns a 3D binary image into a well-composed image of the same size by changing background points to foreground at critical configurations [1]. The expected number of newly created critical configurations is at most half the number of critical configurations of the input, under three stated assumptions, and the expected number of changed points is at most three halves of it. The paper aims at well-composedness and similarity to the input and states that the topology of the result may not be correct. An extension to multivalued images repairs the binary images of the objects one after another, and the result depends on the order of the colours. Boutry et al. generalise repair to gray-valued n D images: the blocks with critical configurations are queued, each correction raises the value of one pixel of the block, neighbouring blocks that become critical are queued again, and the monotone increase of values ensures convergence [10]. The procedure is a local, greedy approximation of the closest well-composed image, and the guarantee it states is well-composedness of the output. Tustison et al. replace the simple-point test in a topologically constrained front propagation by a topological well-composedness criterion on the centre voxel and its face neighbours, and their strict mode preserves the initial topology [11]. This is the closest procedure to the present setting that combines a well-composedness criterion with topology preservation; it is a propagation procedure, and the question of which blocks remain reachable under which move rules is not posed. In these procedures the sequence of changes is determined by the algorithm. In our ablation, additions and

removals are both admitted, a voxel may be flipped once in either direction under flip-once, and the question is which blocks a search over all admissible sequences can resolve. The add-only procedure R^+ described in Section 3.2 and documented in the code-provenance record (Online Resource 2, Section 2) is a binary-mask adaptation of the queue-based correction of [10]; it does not remove foreground. For an addition of v , the attachment $A_X(v)$ decides whether the homotopy type of the foreground complex is preserved: if the attachment is non-empty, connected and of Euler characteristic 1 it is (Online Resource 1, Theorem 1(a)), and otherwise the relative homology of the pair of the complexes with and without v is non-zero in some degree, so that the inclusion is not a homology isomorphism (Online Resource 1, Theorem 1(b)). The attachment alone does not determine how the global Betti numbers change, which depends also on the maps of the long exact sequence of the pair.

2.4 Refinement and complex-based approaches

An interpolation computes an image on a finer grid whose restriction to the original grid equals the input. In-between interpolations constrain inserted values to lie between neighbouring original values and introduce no new extrema [4]. A local, self-dual interpolation of a standard kind cannot make 3D images well-composed in a single subdivision when the mean is used for one-dimensional interpolation [12]. Gonzalez-Diaz, Jimenez and Medrano repair the cubical complex of a 3D binary image locally into a polyhedral complex that is homotopy equivalent to it and whose component boundaries are 2D manifolds [13]. Boutry, Gonzalez-Diaz and Jimenez extend this to n D pictures and show that the resulting cell complex is weakly well-composed and homotopy equivalent to the initial cubical complex [14]. In these approaches the topological guarantee is obtained on a refined grid or in a complex, and the output is not a binary image on the input grid.

2.5 Label images

Multi-label simple points were defined so that the topology of a partition remains invariant during

a deformable partition process [19]. Mazo, Passat, Couprie and Ronse develop topology-preserving local modification of label images, which preserves the topology of each label set and of any union of label sets [20], within a unified topological framework for digital images [21]. These works supply criteria for the validity of a single modification of a label image. The present study concerns a binary mask.

2.6 Positioning

Table 1 summarises the approaches by setting, moves, the guarantee stated in the source and the type of path analysis.

The present study does not propose a repair operator. It analyses which residual critical blocks a given SAFE-gated single-voxel move set can resolve on the input grid when locality, flip-once and fidelity are imposed, and it describes the paths that do. Online Resource 1 proves the homotopy preservation of a SAFE flip for the closed-cube model so that the guarantee used here is self-contained. The statement follows, through the retraction of the cubical neighbourhood onto the attachment (Online Resource 1, Lemma 3), from the characterisation of homotopy-simple points in dimension three in [8]; Online Resource 1 proves it directly for the attachment formulation. It is not offered as a new simple-point criterion or as a theoretical contribution.

3 Methods

3.1 Dataset and target selection

Data. The masks derive from the BraTS 2021 training data, 1,251 eligible cases [22–25]. For each case the RAW mask is the binary whole-tumour (WT) mask, the indicator of the segmentation labels {1, 2, 4}. No image intensities are used. A fixed split assigns the eligible cases to TRAIN (876), VAL (125) and TEST (250): cases are ordered by the SHA-256 digest of a fixed salt followed by the case identifier, and the ordered list is divided by the largest-remainder rule for the fractions 0.7, 0.1 and 0.2, without stratification or later rebalancing. Only TRAIN cases are used. The split and RAW masks follow the authors' earlier network-training work, which used the same 876 TRAIN cases; their provenance and relation

to the present study are recorded in Section 3.2 and Online Resource 2, Section 2. The present study uses no network and no predicted mask. The preceding repair stages (Section 3.2) were run on all 876 TRAIN cases; the analysis of this study uses the 222 of them that contain target blocks.

Targets. The preceding repair stages produced a final state X_T for each of the 876 cases. The present analysis concerns the critical blocks (Section 3.2) that remain in these final states and belong to one of two cohorts. A background voxel of a $k = 2$ block is a *shell* voxel if its six face neighbours are all foreground, so that it is a one-voxel background cavity. The target blocks are the $k = 2$ blocks of X_T whose two background voxels are shell voxels, in two cohorts:

- **C1:** the two background voxels are shell voxels in the RAW mask (614 blocks, 213 cases);
- **C2:** the two background voxels are shell voxels in the final state of the first repair stage but not in the RAW mask (66 blocks, 39 cases).

The 680 target blocks lie in 222 distinct cases. The labels C1 and C2 name these two cohorts throughout the manuscript; they are unrelated to the conditions (C1) and (C2) of Proposition 2.4 of [3]. Critical blocks of other classes can be present in the same final states; they are not targets, and they enter the analysis only through the success criterion (Section 3.5). All search results refer to the 680 target blocks (680 blocks $\times$ 14 systems $\times$ 2 horizons = 19,040 searches). None of them refers to the 876 cases as a whole, and a repaired block is not a repaired case: each block is searched independently from the unchanged X_T of its case.

3.2 Critical configurations and starting state

Critical blocks. Let $X \subset \mathbb{Z}^3$ be the foreground of a binary image on a finite rectangular domain Ω , with background outside Ω . A *block* is a set of 2^k voxels $S = \{z + \sum_{i \in F} \varepsilon_i e_i : \varepsilon_i \in \{0,1\}\}$ for a family F of k coordinate directions. Two voxels of a block are *antagonists* if they differ in all k coordinates. A block of dimension $k = 2$ or 3 is *critical* in X when $X \cap S$ is a pair of antagonists (primary configuration) or S minus a pair of antagonists (secondary configuration) [4]. With the block keys of the implementation, a $k = 2$ block is a 2×2

Table 1 Approaches to well-composed repair and topology-preserving change, and the position of this study.

Work	Setting	Moves or output	Guarantee stated in the source	Path or reachability analysis
[9]	two-valued pictures; magnification	changes of simple pixels; magnified well-composed pictures	topology preserved; equivalent pictures interconvertible	existence of sequences
[15]	2D binary images	–	–	complexity results for reachability
[8]	dimension d ; simple-point detection	one-point insertion/deletion; cubical neighbourhood	homology-simple iff acyclic cubical neighbourhood; Betti characterisation of homotopy-simple points for $d = 3$	–
[1]	3D binary, input grid; multivalued extension	randomized background-to-foreground changes	well-composed output; topology of the result may not be correct	conditional expected-count bounds on created critical configurations
[10]	nD images, input grid	one pixel at a time, values raised, greedy queue	well-composed output; convergence	termination and complexity
[11]	3D binary, topologically constrained front propagation	voxel-wise changes admitted by a well-composedness criterion	strict mode preserves the initial topology	front propagation; no bounded repair search
[12]	gray-valued, local self-dual interpolation	subdivision	3D one-subdivision counterexample for the stated local self-dual class (mean interpolation in 1D)	counterexample
[13, 14]	3D binary; nD pictures	well-composed 3D polyhedral complex; weakly well-composed nD cell complex	homotopy equivalent to the cubical complex	construction
[19, 20]	label images	single modifications	topology of each label set preserved	criteria for single modifications
This study	3D binary, input grid	SAFE-gated single-voxel flips under locality, flip-once, fidelity	SAFE flips preserve foreground closed-cube homotopy (Online Resource 1); complement, isotopy and multi-label topology not covered	bounded search with certificate semantics on 680 blocks

face in one of the planes xy , yz , zx whose two diagonals are each constant and differ from each other; a $k = 3$ block is a $2 \times 2 \times 2$ cube with either two foreground voxels that are body-diagonal and

six background voxels, or two background voxels that are body-diagonal and six foreground voxels. Blocks are taken inside Ω . An image without critical blocks is digitally well-composed [4]. In

three dimensions, well-composedness is equivalent to the absence of the edge-sharing four-cube and the vertex-sharing eight-cube critical configurations [3, Prop. 2.1], and these correspond to the critical blocks of dimension $k = 2$ and $k = 3$ defined above [4]. Critical blocks are detected by the frozen detector of the repair implementation, which scans all blocks inside Ω . On the 256 binary $2 \times 2 \times 2$ volumes, the block rule written in the specification and the frozen detector both return the same critical blocks as the definition above (Online Resource 4).

Preceding repair. Three procedures are kept apart. The add-only procedure R^+ , documented in the code-provenance record (Online Resource 2, Section 2), adds, for each critical block in first-in first-out order, the lexicographically first background voxel of the block and has no gate. The preceding repair, defined here, applies two deterministic stages that start from the RAW mask and use SAFE flips only. The searches of Sections 3.4–3.6 start from the final state of these stages and use their own move sets. The starting state X_T of a case is the final state of the two stages. Stage 1 is R^+ with one change. It processes critical blocks in the first-in first-out queue order of the production implementation of R^+ , which is imported unchanged (Online Resource 2), and adds to each block the lexicographically first background voxel whose flip is SAFE (Section 3.3), where R^+ adds the lexicographically first background voxel; a block without such a voxel is deferred. Stage 2 continues from the final state of stage 1. It repeats the same SAFE additions and, at a fixed point, may remove one foreground voxel of a deferred block if the removal is SAFE; each voxel is flipped at most once, and all flips are restricted to the voxels of the block under repair. Stage 2 stops at a fixed point at which, for every deferred critical block, each flip that stage allows is either not SAFE or excluded by its own flip-once rule. This stopping condition refers to the move set of stage 2. It does not state that the flips of the searches of this study are exhausted: those use different move sets (Section 3.4), and in particular the locality systems are restricted neither to the flips of stage 2 nor to the target block.

Start-state qualification. For each of the 680 targets the following were verified before any search: (i) the state reconstructed from the pinned RAW mask by replaying the recorded edits of

stage 2 has the mask digest recorded for X_T ; (ii) the target block is critical in X_T according to the frozen detector; (iii) identity: the cohort agrees with the recorded classification, the block has $k = 2$ and the recorded candidate classes shell|shell, the recorded stopping cause of stage 2 for the block is that all gate-admitted flips are UNSAFE (`GATE_ALL_UNSAFE`), and the set of critical blocks recomputed in X_T equals the recorded set of stuck blocks of the case; (iv) none of the four voxels of the target block admits a SAFE single flip at X_T . Property (iv) concerns the voxels of the target block only. It does not imply that the locality-admissible move set at X_T is empty, because that set also contains the voxels of other critical blocks of the window (Section 4.1).

3.3 SAFE gate and admissible voxel flips

Topology is evaluated in the closed-cube model with (26,6) adjacency: $K(X)$ is the union of the closed unit cubes C_q of the foreground voxels $q \in X$. For a voxel v , let $A_X(v) = K(X \setminus \{v\}) \cap C_v$. A flip of v , an addition ($0 \rightarrow 1$) or a removal ($1 \rightarrow 0$), is **SAFE** if the Betti numbers of $A_X(v)$ are $(\beta_0, \beta_1, \beta_2) = (1, 0, 0)$. Figure 1 shows the attachment and Figure 2 examples of its types. The set $A_X(v)$ depends only on the 26-neighbourhood of v and not on the value of v , so the same value is obtained before and after the flip, and the gate is evaluated by a lookup table over the 2^{26} neighbourhood configurations. The table was regenerated from the exhaustive gate program and is pinned by its SHA-256 digest (Section 3.7). Every neighbourhood configuration that a search consulted was re-evaluated by direct computation of the attachment homology, and a disagreement with the table stops the run.

Attachment geometry. A neighbour q of v with offset $o = q - v$ touches C_v in the face $f_v(o) = C_v \cap C_q$, and $A_X(v)$ is the union of the faces $f_v(q - v)$ over the foreground neighbours q . If k coordinates of o are nonzero, $f_v(o)$ has dimension $3 - k$; the 26 offsets therefore correspond one-to-one to the 6 squares, 12 edges and 8 vertices of ∂C_v (Online Resource 1, Lemma 1; Figure 1). If the six face neighbours of v are foreground, $A_X(v) = \partial C_v$.

Classification of attachments. Write c for the number of connected components and χ for the Euler characteristic of $A_X(v)$. A flip is SAFE




Fig. 1 Attachment $A_X(v) = K(X \setminus \{v\}) \cap C_v$. Black cells show contacts $C_v \cap C_q$ for face, edge and vertex neighbours (examples $+x$, $+x+y$, $+x+y+z$); the corresponding counts are 6, 12 and 8 (Online Resource 1, Lemma 1). Inset: six foreground face neighbours give $A_X(v) = \partial C_v$. Hatching identifies the axis normal to a square.

if and only if $c = 1$ and $\chi = 1$, equivalently $\beta(A_X(v)) = (1, 0, 0)$ (Online Resource 1, Proposition 1). Both conditions are needed, and they fail for different reasons (Figure 2): the empty attachment T0 and the two disjoint squares T2 have $c \neq 1$; the band T3 is connected with $\chi = 0$; the boundary ∂C_v (T4) has $\chi = 2$ and $\beta_2 = 1$; the 4-cycle with an isolated vertex T5 has $(c, \chi) = (2, 1)$. Only the single square T1 is SAFE.

Free-pair and leaf reduction. If $\emptyset \neq A_X(v) \neq \partial C_v$, removing a square together with one of its free edges, repeatedly, reduces $A_X(v)$ to a graph without changing its homotopy type, its number c of connected components or its Euler characteristic χ , and $\beta_0 = c$, $\beta_1 = c - \chi$, $\beta_2 = 0$ (Online Resource 1, Lemma 2; Figure 3, $A_0 \rightarrow A_3$). This does not apply to $A_X(v) = \partial C_v$, which has no free edge and $(\beta_0, \beta_1, \beta_2) = (1, 0, 1)$. Reducing the graph to a point by removing leaves requires a tree, that is $c = 1$ and $\chi = 1$, and therefore applies to SAFE attachments only (Online Resource 1, Proposition 1; Figure 3, $A_3 \rightarrow A_4$).

Moves are called SAFE-gated: a flip is considered only if it passes this gate. A flip is *admissible* in a system if it is SAFE and satisfies the rules of that system (Section 3.4). Only single-voxel flips are considered; simultaneous flips of several voxels are not.

Coefficients. Because $A_X(v)$ is a subcomplex of the boundary of the cube C_v , the condition $(\beta_0, \beta_1, \beta_2) = (1, 0, 0)$ is equivalent to $A_X(v)$ being non-empty, connected and of Euler characteristic 1, and to $A_X(v)$ being collapsible; it therefore does not depend on the coefficient field, and the integral homology of $A_X(v)$ is torsion-free (Online Resource 1, Lemma 2 and Proposition 1). An independent computation of the Betti numbers of all attachments by boundary-matrix ranks

over F_2 , and over $\mathbb{Q}$ with integral Smith normal form, gives the same decision as the table on all 2^{26} neighbourhood configurations.

Foreground homotopy type. If the flip of v is SAFE, the closed-cube complex of the smaller of the two foregrounds is a strong deformation retract of that of the larger, so that $K(X)$ and $K(X \triangle \{v\})$ are homotopy equivalent (Online Resource 1, Theorem 1). Along a sequence of SAFE flips the complexes $K(X_i)$ are therefore pairwise homotopy equivalent, with the same number of connected components and isomorphic homology groups in every degree (Corollary 2). Every move of a witness path of Section 4 is SAFE at the state in which it is made (Section 3.7), so this holds along each witness path. If a flip is not SAFE, the inclusion of the smaller complex into the larger one is not an isomorphism in integral homology in some degree. These statements concern the foreground complex $K(X)$. They do not concern the homotopy type of the complement of $K(X)$, ambient isotopy or multi-label images; consequences for the complement that may follow from the literature on simple points are not used in this study.

Deformation retraction. Write $K = K(X \setminus \{v\})$, $K^+ = K(X \cup \{v\}) = K \cup C_v$ and $A = A_X(v) = K \cap C_v$. If the flip is SAFE, A is contractible, so that the inclusion $A \hookrightarrow C_v$ is a homotopy equivalence, and the homotopy extension property of the CW pair (C_v, A) gives a deformation retraction of C_v onto A that fixes A pointwise. Extended by the identity on K , it retracts K^+ onto K (Online Resource 1, Theorem 1(a); Figure 4). This retraction is obtained from the extension property; it is not constructed and is not a sequence of elementary collapses.







						
$\beta(A)$	$\emptyset$ (0,0,0)	square (1,0,0)	2 squares (2,0,0)	band (1,1,0)	∂C_v (1,0,1)	cycle+pt (2,1,0)
(c, χ)	(0, 0)	(1, 1)	(2, 2)	(1, 0)	(1, 2)	(2, 1)
SAFE						

Fig. 2 Attachments $A = A_X(v) \subseteq \partial C_v$, with Betti vector $\beta(A) = (\beta_0, \beta_1, \beta_2)$, component count $c = \beta_0$ and Euler characteristic $\chi = \beta_0 - \beta_1 + \beta_2$. The SAFE row tests $c = \chi = 1$; T1 is framed and failing entries are boxed (Online Resource 1, Proposition 1). T0: empty; T2: opposite squares; T3: four-square band; T4: full boundary from six foreground face neighbours; T5: a 4-cycle and an isolated vertex. Hatching identifies square normals; relative groups appear in Figure 5.




Fig. 3 Attachment reduction for neighbours $X = \{+x, +y, -y\}$ (inset, gray). $A_0 \rightarrow A_3$: removal of square-free-edge pairs (Q, e) , leaving a graph; black cells and arrows identify each removed pair and its collapse direction. Lemma 2 applies to $\emptyset \neq A \neq \partial C_v$. $A_3 \rightarrow A_4$: removal of edge-free-endpoint pairs (w, u) to a point, restricted to SAFE attachments (Online Resource 1, Lemma 2 and Proposition 1). The circle marks this SAFE-only scope; hatching identifies square normals.

Relative homology of the inclusion. If the flip is not SAFE, the relative groups of the inclusion $K \hookrightarrow K^+$ are $H_d(K^+, K) \cong H_d(C_v, A) \cong \tilde{H}_{d-1}(A)$ for $A \neq \emptyset$, and $H_0(K^+, K) \cong \mathbb{Z}$ for $A = \emptyset$ (Online Resource 1, proof of Theorem 1(b)); at least one of the relative groups $H_d(K^+, K)$ is nonzero, so the inclusion is not an isomorphism in integral homology. For each of the attachments T0 and T2–T5 of Figure 2, Figure 5 shows one chain for each nonzero relative group, representing a generator of that group. T5 has two nonzero groups, $H_1(K^+, K) \cong \mathbb{Z}$ and $H_2(K^+, K) \cong \mathbb{Z}$, whose generators are represented by an arc ($d = 1$) and a disk ($d = 2$). For T1 (SAFE) the inclusion is an isomorphism in homology (Online Resource 1, Theorem 1(a) and Corollary 4), so that the relative groups vanish. Theorem 1(b) concerns the inclusion map and does not assert that $K(X)$ and

$K(X \triangle \{v\})$ fail to be homotopy equivalent by some other map.

Relation to simple points. The exhaustive gate program evaluates three decisions on all $2^{26} = 67,108,864$ neighbourhood configurations: the topological-number test $T_{26} = 1$ and $\bar{T}_6 = 1$, the Betti condition on $A_X(v)$, and the collapsibility of $A_X(v)$. They accept the same 25,985,144 of these configurations (38.7 %) and never differ; the comparison was repeated from the frozen source and reproduced the digest of the table (Online Resource 2). The agreement is a finite computation on the neighbourhood. The correspondence of the implemented topological-number test with the definition in [7] has not been checked against the original, and no equivalence for general images is used. The attachment $A_X(v)$ is the intersection of the cubical neighbourhood of




Fig. 4 Retraction for the SAFE attachment of neighbours $X = \{+x, +y, -y\}$ (gray), a union of three squares. Top (enlarged): $C_v \rightarrow A$, fixing A . Bottom: $K^+ = K \cup C_v \rightarrow K$, where $K = K(X \setminus \{v\})$. The dashed connector identifies extension by the identity on K (Online Resource 1, Theorem 1(a)). Hollow arrows denote the guaranteed retractions; the circle marks SAFE-only scope. Hatching identifies square normals.

v with the cube of v and is a strong deformation retract of that neighbourhood (Online Resource 1, Lemma 3), so the Betti condition of the gate is the condition on the neighbourhood in [8]. No new simple-point criterion is proposed. The statement used in this study, the homotopy equivalence of the foreground complex under a SAFE flip, is the homotopy-simple case of the characterisation of [8] in dimension three; it is proved in Online Resource 1 for the attachment formulation, without use of that characterisation, independently of this comparison.

Explicit deformation of the neighbourhood. The deformation of Online Resource 1, Lemma 3 is explicit and does not depend on the SAFE condition: with N_v the union of the closed cubes of the foreground neighbours of v , the coordinate-wise clamp $\pi: \mathbb{R}^3 \rightarrow C_v$ is one continuous map that sends the cube of each neighbour onto its contact face, and $H_t(x) = (1-t)x + t\pi(x)$ deforms N_v

into $A_X(v)$ inside N_v for every subcomplex $A_X(v)$ of ∂C_v (Figure 6).

3.4 Locality, flip-once and foreground fidelity

Let X_T be the starting state of a search and RAW the original mask. Three rule families are combined into constraint cells and fidelity criteria.

- **Locality (Loc).** A flip of voxel v in state X is admissible only if v belongs to a block that is critical in X . The condition depends on the current state, so it cannot be evaluated from X_T alone.
- **Flip-once, path-local (P).** A flip of v is admissible only if the current value of v equals its value in X_T ; each voxel is therefore flipped at most once along a path.
- **Flip-once, history-based (H).** An addition requires $\text{RAW}(v) = 0$ and a removal requires $\text{RAW}(v) = 1$.
- **Foreground fidelity.** Let $\text{NRFR}(X)$ be the number of voxels that are foreground in both RAW and X_T and are background in X , with X_T as the baseline. **F0** imposes nothing. **F1** (endpoint fidelity) requires $\text{NRFR} = 0$ in the success state. **F2** (pathwise fidelity) forbids, at every step, the removal of a voxel that is foreground in RAW.

Relation of H and P. Consider paths that start at X_T . Under H, a voxel with $\text{RAW}(v) = 0$ can only be added and a voxel with $\text{RAW}(v) = 1$ can only be removed, so the reverse flip of a flipped voxel is never admissible; every H-admissible path therefore flips each voxel at most once, and each of its flips is admissible under P. H-admissible flips are exactly the P-admissible flips of the voxels with $\text{RAW}(v) = X_T(v)$: a voxel with $\text{RAW}(v) \neq X_T(v)$ was changed by the preceding stages, and H excludes both removal of an added voxel and restoration of a removed RAW voxel, which P allows. Hence the set of H-admissible paths from X_T is a subset of the set of P-admissible paths, the states reachable under H are a subset of those reachable under P, and the shortest successful depth under H is not smaller than under P. In the other direction, P admits paths that H excludes.




Fig. 5 Non-SAFE attachments and relative generators of $H_d(K^+, K; \mathbb{Z})$, with $K^+ = K \cup C_v$ and $A = K \cap C_v$ (amber). Top: a representative chain c in C_v (black); middle: ∂c in A ; bottom: its degree $d = \dim c$. T0: a point disjoint from K (gray); T2: an arc between components; T3: a disk spanning a cycle; T4: the cube bounded by $A = \partial C_v$; T5: both an arc and a disk, in degrees 1 and 2. Type IDs match Figure 2; SAFE T1 has zero relative groups and is omitted (Online Resource 1, Theorem 1 and Corollary 4). Hatching identifies square normals.

Table 2 lists the 14 systems. A system is a pair of a cell and a criterion, written CELL-FID as in the code. In the flip-once cells, endpoint and path-wise fidelity have the same success and the same shortest depth (Section 5.3), so the F1 entries of these cells are represented by F2 searches. The manuscript reports the path-local family together with C00 and C10; the history-based family is part of the 14 systems and of the production tables and is not analysed in this manuscript. Each constraint removes transitions from the unrestricted move set, so the sets of admissible transitions are nested: $C11 \subseteq C10$ and $C11 \subseteq C01 \subseteq C00$ within a flip-once family, and $F2 \subseteq F1 \subseteq F0$ for the paths admitted by the criteria.

3.5 Search state, goal condition and breadth-first search

Window. For a target block B , let $\text{bbox}(B)$ be its axis-aligned bounding box and $W_R(B) = \Omega \cap \{v : d_\infty(v, \text{bbox}(B)) \leq R\}$. For a $k = 2$ block, W_1 has up to 48 voxels and W_2 up to 180 voxels. Two horizons are used: G26 ($R = 2$, depth $L = 6$) and G14 ($R = 1$, $L = 4$). Voxels outside the window keep their values in X_T . A one-voxel layer around the window (halo) is read from X_T and never flipped. The gate of a window voxel needs only its 26-neighbourhood, and a block that contains a window voxel lies in the window and its halo. A block that contains no window voxel is unchanged by any search. The criticality of every block is therefore evaluated exactly.

State. A state is the occupancy bitmask of the window, compared exactly (no hashing).

Goal condition. Let $\text{Crit}(X)$ be the set of critical blocks of X . A state Y is a **success state** of B if $B \notin \text{Crit}(Y)$ and $\text{Crit}(Y) \setminus \text{Crit}(X_T) = \emptyset$; that is, B is not critical and no block that was not critical in X_T is critical in Y . Other critical blocks of X_T may disappear. Because blocks without a window voxel are unchanged, the second condition is decided on the blocks that contain a window voxel. The state test is the same for F0 and F2 (F2 acts on the moves); for F1 the success state must also satisfy $\text{NRFR}(Y) = 0$.

Search. A breadth-first search with a first-in first-out queue starts at X_T and enumerates, level by level up to depth L , the states reachable by admissible single-voxel flips. The children of a parent are generated in the order of the voxel coordinates (z, y, x) , parents in the order of discovery, and a state that has been discovered is not expanded again. The search stops at the first success state discovered; a path to it is a shortest path of admissible flips for the system, window and gate, and it is the first one in the fixed enumeration order. This path is called the canonical witness for that search. A witness inherited from an earlier search need not be canonical in the enumeration order of a different window (Section 3.7). States at depth L are generated and tested but not expanded. A search is run independently for each block, system and horizon.




Fig. 6 Coordinatewise clamp $\pi : \mathbb{R}^3 \rightarrow C_v = [0, 1]^3$. Top: square, edge and vertex contacts, with representative projection arrows and $\pi(C_q) = C_q \cap C_v$. Bottom: the same clamp for two attachment topologies, SAFE (teal) and non-SAFE (amber). In both cases $H_t(x) = (1 - t)x + t\pi(x)$ retracts $N_v = \bigcup C_q$ onto $A_X(v)$ (Online Resource 1, Lemma 3). Hatching identifies square normals.

Table 2 The 14 systems. F1 entries marked **via F2** use the equivalence of successful instances and shortest successful depth under flip-once (Section 5.3); they are not separate F1 searches.

Cell	Loc	Flip-once	F0	F1	F2
C00	off	off	C00-F0	C00-F1	C00-F2
C10	on	off	C10-F0	C10-F1	C10-F2
C01P	off	P	C01P-F0	(via F2)	C01P-F2
C01H	off	H	C01H-F0	(via F2)	C01H-F2
C11P	on	P	C11P-F0	(via F2)	C11P-F2
C11H	on	H	C11H-F0	(via F2)	C11H-F2

3.6 Exhaustion, depth-limited and cap semantics

A search ends in one of four outcomes (Table 3). The limit on the number of states is $N_{\max} = 2 \times 10^6$, counting the initial state and all distinct

states generated at depths 0 to L, those at depth L included. The search aborts at the moment the $(N_{\max} + 1)$ -th distinct state would be admitted, unless a success state was discovered earlier.

Table 3 Outcomes of a search and what each establishes. The statements hold for the stated system, window $W_R(B)$, horizon and SAFE gate.

Outcome (code status)	Condition	Established	Not established
FOUND (PURE_FOUND)	a success state was discovered at depth $L^* \leq L$	an admissible path of L^* flips to a success state exists; L^* is the minimum	uniqueness of the path; repair of the case; topology of the complement, ambient isotopy or multi-label topology
EXH (NO_PURE_EXHAUSTED)	the search ended because a level was empty, before or at depth L	the whole component reachable by admissible flips was enumerated and contains no success state: no admissible path of any length within the window leads to one	any statement outside the window or under other moves
DL (NO_PURE_DEPTH_LIMITED)	level L is non-empty; all states within L admissible flips were enumerated without truncation and none is a success state	no admissible path of at most L flips to a success state exists	paths longer than L ; larger windows; other move sets
CAP (CAP_EXCEEDED)	$N_{\max}$ was reached before a success state was found and before the enumeration ended	nothing: a CAP row is not evidence of absence and is counted separately	absence of a path of any length

EXH and DL differ only in whether states of depth L remain unexpanded. The three conclusive outcomes are FOUND, EXH and DL; CAP is not conclusive. When the initial state has no admissible flip, the search ends at depth 0 with one state (EXH), and its outcome carries no information beyond the empty first move set. Rows whose absence verdict was inherited from an earlier engine (Section 3.7) are marked separately in the tables.

3.7 Reuse of earlier searches and reproducibility

Reuse of earlier searches. In C00-F1 and C00-F2 at G26, the rows that ended in CAP in the production run were replaced by the verdict of an earlier completed search of the same system on the same block (655 rows: 44 FOUND and 611 DL; the list is in the reuse log). The 611 DL rows are 517 in C00-F1 (471 in C1, 46 in C2) and 94 in C00-F2 (88 in C1, 6 in C2). Counts that include substituted rows are called *effective*; all other counts come from the production run itself.

Inherited absence verdicts. The 611 inherited DL rows come from a look-ahead engine that

stores the breadth-first levels 0 to $L - 2$ as the production search does and decides the last two levels by a necessary condition for success within one further flip (completeness argument: Online Resource 2, Section 3, and the Online Resource 4). Their completeness rests on that argument, on an audit of the engine code against it, and on comparisons with the unchanged breadth-first search. The unchanged search reproduced 80 of the 611 DL rows (45 in C00-F1, 35 in C00-F2) and 5 of the 44 FOUND rows without a difference, and a later pilot (G1) re-decided 12 further DL rows at the original horizon $L = 6$ and re-evaluated 12 at $L = 5$ only. Of the 611 inherited DL rows, 92 (15.1%) have therefore been re-decided at $L = 6$ and 519 (84.9%) have not. The re-decided rows were selected by projected state count and are the smaller ones, so the agreement does not extend by itself to the others. No second implementation of the engine exists. The FOUND rows were verified in addition by replaying all 45 witnesses of the engine run on the full volume with an independent step-by-step implementation of the gate and the system rules. Statements in Sections 4 and 5

that depend on inherited absence verdicts give the number of rows that rely on them.

Provenance. The production run and its checks use the frozen search modules, imported unchanged and identified by SHA-256 digest, and the gate table pinned by digest `59407fef...4830`. The digest is that of the regenerated table; the original record of the table is not available. The state X_T of each case is reconstructed from the pinned RAW masks and the recorded edit sequence of stage 2, and its mask digest equals the recorded one in 222 of 222 cases. Searches are deterministic. Every FOUND witness of the production run was replayed step by step on the full volume with an explicit per-voxel gate and the rules of the system. The counts in the tables were recomputed from the production tables by a read-only script that compares the digests of its inputs with a master list; it shows that the stored results agree with the reported tables and is not an independent verification of the searches. Versions, scripts and logs are given in Online Resource 2.

Data and code availability. The data and code statements are given under Declarations.

4 Results

4.1 Search outcomes across constraint cells

We searched 680 target blocks in 14 systems at the two horizons, giving 19,040 searches, none of which ended in an error (Table 4).

Without fidelity (F0), the unrestricted search at G26 reached a repair in 666 of 680 blocks, with 14 searches capped. With flip-once alone, 636 blocks were repaired (4 DL, 40 CAP). With the restriction to critical blocks alone, 5 blocks were repaired, and with the two restrictions together none. At G14 the corresponding counts were 661, 588, 0 and 0. The set of FOUND blocks is not monotone in (R, L), because the state limit applies at both settings: ten blocks of C00-F0 that are FOUND at G14 end in CAP at G26 (listed in Online Resource 2). This is a property of the resource-capped outcome and not of reachability; a CAP row at the larger setting says nothing about the repair found at the smaller one.

Many of the zeros are determined by the move set at X_T . With the history-based flip-once rule, the combined cell C11H has no admissible first

move in any of the 680 blocks. With the path-local rule, the first move set of C10 and C11P is empty in 595 blocks at G26 and in 636 at G14. In those blocks the search visits one state, and its failure carries no information beyond the empty first move set; these are structural zeros. A zero FOUND count in the remaining 85 blocks (C11P in every criterion, C10 under F1 and F2) is an observed outcome of searches that enumerated more than one state.

4.2 Blocks repaired under locality alone and under flip-once alone

The five C10-F0 repairs at G26 all come from the 85 blocks with a non-empty first move set. Table 5 lists them with the outcome of the flip-once cell and of the combined cell on the same blocks. Four of the five blocks are also repaired by C01P-F0, and the fifth is a C01P CAP.

In the four blocks repaired by both cells, the two repairs are different paths. The C10 path contains a voxel flipped twice and so violates flip-once; the C01P path begins with a flip that is not in a critical block and so violates locality. In all five blocks the combined cell C11P is DL, which means that no repair of at most six flips is admissible under both rules.

4.3 Effect of fidelity

Tightening the success criterion reduced the number of repairs found in the unrestricted cell. At G26 the counts fell from 666 (F0) to 72 (F1) and 15 (F2). At G14 they fell from 661 to 15 and 3. For F1 and F2 the large majority of the remaining blocks are DL (606 and 460 at G26), so for those blocks no endpoint-faithful repair of at most six flips, respectively no pathwise-faithful repair of at most six flips, exists in the window. Of these DL entries, 517 (F1) and 94 (F2) are inherited absence verdicts (Section 3.7).

The F0 repairs are mostly short. Of 666 canonical F0 witnesses at G26, 576 have $L^* = 2$ and 662 remove at least one RAW foreground voxel net. Among the 28 F1 witnesses of the present run, net removal is zero by construction. Of these, 15 contain a step that removes a RAW foreground voxel, which is restored before the path ends, and 16 flip at least one voxel twice. Among the 15 F2

Table 4 Search outcomes at G26 ($R = 2$, $L = 6$) for 680 target blocks: four selected constraint cells of the 14 systems. Entries are FOUND / EXH / DL / CAP. Effective counts combine production and earlier completed searches; superscripts a and b identify these records.

Criterion	C00	C10	C01P	C11P
F0	666 / 0 / 0 / 14	5 / 630 / 45 / 0	636 / 0 / 4 / 40	0 / 631 / 49 / 0
F1	72 ^b / 0 / 606 ^b / 2	0 / 630 / 50 / 0	13 / 207 / 365 / 95 ^a	0 / 673 / 7 / 0 ^a
F2	15 / 205 / 460 ^b / 0	0 / 669 / 11 / 0	13 / 207 / 365 / 95	0 / 673 / 7 / 0

^a Computed with F2 (Section 5.3). C00-F1 FOUND = 72 comprises 28 from the present run and 44 from earlier completed searches; C00-F1 DL = 606 comprises 89 from the present run and 517 from earlier completed searches; C00-F2 DL = 460 comprises 366 and 94 (Section 3.7). ^b Includes rows from earlier searches (Section 3.7): the 44 FOUND rows were verified by witness replay; of the 611 DL rows (517 and 94), 80 were reproduced by exhaustive search and 531 were not in the V2 comparison (historical record). In the later G1 pilot (Section 3.7), 12 of the 531 were re-decided by the unchanged search at the original horizon $L = 6$, which gives 92 of the 611 re-decided at $L = 6$ and 519 not. Twelve further G1 rows were searched to $L = 5$ only and are not counted as re-decided. Counts for all 14 systems at both horizons are in Online Resource 3, Table S2. EXH entries that consist of the initial state only (C10 and C11P: 595 of 680 blocks in each criterion; C00 and C01P under F2: 69) are structural zeros determined by the move set; all other entries are outcomes of searches that enumerated more than one state.

Table 5 Blocks repaired by C10-F0 at G26.

Block	C10 path	C01P-F0	C11P-F0
BraTS2021_00310 k2xy@93,161,114	6 flips, 2 voxels flipped twice	repair, 2 flips	DL (84 states)
BraTS2021_00598 k2zx@46,133,65	6 flips, 1 voxel flipped twice	CAP	DL (126 states)
BraTS2021_01293 k2xy@71,79,164	6 flips, 1 voxel flipped twice	repair, 2 flips	DL (55 states)
BraTS2021_01293 k2yz@94,112,134	6 flips, 2 voxels flipped twice	repair, 2 flips	DL (23 states)
BraTS2021_01296 k2yz@75,136,83	6 flips, 1 voxel flipped twice	repair, 2 flips	DL (157 states)

witnesses, none removes a RAW foreground voxel, and 2 flip a voxel twice.

Split by cohort (Online Resource 3, Table S3), the F0 repairs of the unrestricted cell are 602 of 614 C1 blocks and 64 of 66 C2 blocks. Under F1 the unrestricted cell repaired 54 C1 blocks and 18 C2 blocks. Under F2 it found no repair in any of the 614 C1 blocks, of which 88 carry inherited absence verdicts (Section 3.7), and a repair in 15 of the 66 C2 blocks; the 13 F2 repairs of the flip-once cell are also in C2 (G14: 0 and 3 for the unrestricted cell). The F2 repairs of Table 4 therefore all lie in C2 blocks.

4.4 Flip-once and pathwise fidelity admit different paths

With F2, the unrestricted cell found 15 repairs at G26 and the flip-once cell found 13. Both cells found 3 at G14. Two blocks are found by C00-F2 but not by C01P-F2 at G26.

- BraTS2021_01294 (k2xy@83,115,154). C00-F2 reaches a repair in four flips: two removals of voxels added by the preceding stage, then two additions, the last of which restores the voxel removed first. Every step satisfies the F2 rule, and the last step violates flip-once. The C01P-F2 search enumerates 639,806 states and is DL, so no repair of at most six flips that satisfies both flip-once and F2 exists in this window.
- BraTS2021_01282 (k2xy@85,134,113). The C00-F2 witness has five flips with one voxel flipped twice; the C01P-F2 search reaches the state limit (CAP), which is inconclusive.

The converse also holds. The two-flip repair of BraTS2021_00110 (Section 4.5) is found by C01P-F0, so it satisfies flip-once, while its first step removes a RAW foreground voxel and so violates F2. More generally, 632 of the 636 C01P-F0 repairs at G26 remove RAW foreground. Neither

constraint therefore implies the other, and the sets of admissible paths differ in both directions.

4.5 A critical block repaired under endpoint fidelity by a path with a temporary removal (Figure 7)

In BraTS2021_00110, the critical block B lies inside the tumour mask (Figure 7a). Its two background voxels are enclosed by foreground, and at the initial state X_T none of the four voxels of B admits a SAFE single flip. The window of 180 voxels contains 41 SAFE flips at X_T , all outside B. The unrestricted search with endpoint fidelity (F1) found a four-flip repair (Figure 7b). The first two flips remove a RAW foreground voxel and an added voxel, which widens the enclosed background. The third fills it, so that B is no longer critical. The fourth restores the RAW voxel, giving $\text{NRFR}(\text{final}) = 0$. The voxel filled at step 3, (106,111,162), is not SAFE at X_T and becomes SAFE after the first flip, remaining SAFE after the second (Online Resource 3, Table S4).

Table 6 lists the outcome of each search on this block. Without the fidelity requirement (F0), B is repaired in two flips: removal of (106,110,162) followed by addition of (106,111,162) (Online Resource 3, Table S1(c)). These are the first and third flips of the four-flip repair, at the cost of one RAW foreground voxel, and C01P-F0 finds the same path. With pathwise fidelity (F2), with or without flip-once, the six-flip search is DL, so no pathwise-faithful repair of at most six flips exists in this window. Consequently every F1 repair of at most six flips in this window contains at least one step that removes a RAW foreground voxel: a repair without such a step satisfies F2, and the F2 search found none. This statement concerns endpoint fidelity (F1, baseline X_T), single-voxel flips, the window $R \leq 2$, depth $L \leq 6$ and the SAFE gate of this block; it does not extend to other windows, longer paths or other fidelity definitions. The two searches enumerate 24,528 states each, with identical counts at every depth (1, 18, 149, 759, 2,678, 6,974, 13,949), and their states are identified by the same bitmask. The flip-once move set is contained in the unrestricted one, so the states discovered by the flip-once search are a subset of those of the unrestricted search. The totals are equal, so the two sets of discovered states are the same. With the locality restriction,

the first move set at X_T is empty, and the search ends at the initial state.

At G14 ($R = 1$, $L = 4$), C00-F1 finds the same four-flip repair, after generating 3,293 states.

4.6 Control block (Online Resource 3, Figure S1)

In the control block BraTS2021_00312 a two-flip F0 repair exists (Online Resource 3, Figure S1). It removes a RAW foreground voxel and adds a voxel outside RAW, leaving $\text{NRFR}(\text{final}) = 1$. On the same block the F1 search is DL (342,262 states), so no endpoint-faithful repair of at most six flips exists in this window, and the F2 search is exhausted at the initial state, which has no admissible move.

5 Discussion



5.1 Summary of the ablation

The comparison separates the topology guarantee of a move from the repair paths admitted by each system. At G26 the F0 FOUND counts are 666, 636, 5 and 0 for C00, C01P, C10 and C11P, respectively (Table 4). Locality and path-local flip-once have different effects: 595 locality instances have an empty initial move set, while the five locality repairs come from the remaining 85. In four of these blocks a flip-once repair also exists, but the locality witness repeats a flip and the flip-once witness violates locality; the combined search is DL in all five blocks (Sections 4.1 and 4.2). These paired cases show how rules can exclude different repairs of the same target. The aggregate FOUND counts remain lower bounds where CAP occurs, and fidelity absence verdicts include inherited rows with the provenance stated in Section 3.7.

5.2 The four-flip repair and its temporary removal

In BraTS2021_00110, the first removal changes the attachment of the later filled voxel (106,111,162) from $\beta = (1, 0, 1)$ to $(1, 0, 0)$; that voxel is already SAFE after step 1 (Online Resource 3, Table S4). The four-flip endpoint-faithful witness contains two removals, one filling step and restoration of the RAW voxel. Only the

(a) Critical configuration



(b) Repair trajectory




Fig. 7 Endpoint-faithful repair in BraTS2021_00110. (a) Cut at critical block B ; dashed box: view in (b). Plain, hatched and dotted fills denote RAW foreground, previously added voxels and enclosed background. (b) Four SAFE flips from X_T : two removals, filling to resolve B (dashed frame), then RAW restoration. The teal frame tracks the same RAW voxel $(z, y, x) = (106, 110, 162)$: removed at step 1, dashed while absent, restored at step 4. Heavy outlines mark flipped voxels. NRFR labels show endpoint fidelity despite pathwise-fidelity violation at step 1; restoration violates flip-once at step 4 (Online Resource 3, Table S1(a)).

Table 6 Selected systems on BraTS2021_00110 (G26). States is the number of distinct generated states, including the initial state, counted until termination; FOUND stops at the first repair.

System	Outcome	L^*	States
C00-F0	repair found	2	139
C01P-F0	repair found	2	139
C00-F1	repair found	4	52,137
C00-F2	DL	—	24,528
C01P-F2	DL	—	24,528
C10-F1	EXH (initial state only)	—	1
C11P-F2	EXH (initial state only)	—	1

filling step acts inside B ; the other flips violate locality in the states in which they are made. The completed F2 search excludes a repair of at most six flips in this window, so every F1 repair within these bounds temporarily removes RAW foreground (Section 4.5). The witness in Figure 7 restores it at step 4. This bounded necessity follows from the F2 absence verdict, independently of a proposed geometric explanation.

The same path shows how the constraints act on it. The first flip removes a RAW foreground voxel. Endpoint fidelity (F1) permits this because the voxel is restored at step 4. Pathwise fidelity (F2) excludes the path at step 1. Flip-once permits step 1 and excludes step 4, since the voxel has already been flipped (Online Resource 3, Table S1). Each of F2 and flip-once therefore excludes this witness, at a different step. The two-flip F0

path in Table 6 is not excluded by flip-once and is excluded by F2, consistent with the large F0 count under flip-once (636) and the small F2 count (15) in Table 4.

5.3 Flip-once and pathwise fidelity are different constraints

Section 4.4 separates flip-once and pathwise fidelity by explicit witnesses: BraTS2021_01294 admits an F2 repair that repeats a flip, whereas the two-flip repair of BraTS2021_00110 satisfies flip-once and removes RAW foreground. These paths establish non-inclusion in both directions. The four-flip witness is excluded by both rules at different steps, as described above.

One equivalence holds by construction. With flip-once, a RAW foreground voxel that has been removed cannot be restored, so a path that ends with $\text{NRFR} = 0$ contains no removal of a RAW foreground voxel of X_T . Conversely, a path with no such removal ends with $\text{NRFR} = 0$ under any rule. Under flip-once, endpoint (F1) and pathwise (F2) fidelity therefore have the same success and the same shortest successful depth. The statement is limited to these two properties; it does not concern the states that a search visits. This is why the F1 entries of the flip-once cells in Table 4 were computed with F2.

For BraTS2021_00110 specifically, the unrestricted and flip-once searches under F2 discovered the same state set (Section 4.5). This identity is for this block and its horizon.

5.4 Implication for repair design

The endpoint/pathwise distinction is a choice about admissible repair paths. For BraTS2021_00110, the completed searches and replayed witness establish that a temporary RAW removal permits an endpoint-faithful repair while pathwise fidelity excludes repair within the same bounds. In the full G26 cohort, the effective unrestricted FOUND counts are 72 under F1 and 15 under F2; F1 has 54 C1 and 18 C2 repairs, while all F2 repairs lie in C2 (Online Resource 3, Table S3). Endpoint fidelity can therefore retain successful paths that restore foreground after a temporary removal. Whether such paths should be allowed is a repair-design choice; the evidence

here concerns single-voxel flips and the stated windows, horizons and cohorts.

5.5 Limitations

SAFE guarantees homotopy preservation of the foreground closed-cubical complex under single-voxel flips (Online Resource 1), not general preservation of background topology, ambient isotopy, or multi-label topology. This guarantee does not ensure that a repair exists. The reachability results concern independently searched blocks, specified move rules, radii $R \leq 2$, and path lengths $L \leq 6$; they do not establish reachability along arbitrary repair paths.

DL excludes repairs only within the stated window and horizon; CAP means that the search budget was exhausted without completing the search, not that no repair exists. Of 611 inherited DL verdicts, 92 were re-decided at $R = 2, L = 6$ using the unchanged breadth-first search, while 519 were not independently re-decided at those bounds. The inherited verdicts depend on the look-ahead completeness argument and its correct implementation; agreement on selected smaller searches does not establish agreement for the remainder (Section 3.7; Online Resource 2, Sections 3–5). FOUND counts remain lower bounds. Flip-once F1 outcomes use the equivalence in Section 5.3.

The experiments use BraTS 2021 whole-tumour masks, with one recorded witness path per successful block–system–horizon search (Section 3.5). F2 repairs occur only in cohort C2, and the mechanism in Section 5.2 is illustrated by two blocks. Extension to other datasets, connectivity conventions, or multi-label images requires separate validation.

Declarations

- **Funding.** The authors received no funding for the submitted work.
- **Competing interests.** The authors have no relevant financial or non-financial interests to disclose.
- **Author contributions.** J.B.: Conceptualization, Methodology, Software, Validation, Formal analysis, Visualization, Writing – original

draft, Writing – review & editing. Y.B.: Conceptualization, Methodology, Validation, Formal analysis, Writing – original draft, Writing – review & editing. J.H.: Conceptualization, Methodology, Validation, Formal analysis, Supervision, Writing – review & editing.

- **Data availability.** BraTS 2021 data are available from the original providers [22–25] under their access and use terms. The release's DATA_ACCESS.md specifies the provider access route and input identity/hash checks for a lawfully obtained BraTS2021_00110 segmentation or whole-tumour mask. Original MRI, segmentation masks and voxel-level witness/context/replay arrays are not redistributed. Non-image production records and Table S4 cell ledgers are supplied in the release below and Online Resource 4; supplementary results appear in Online Resource 3.
- **Code availability.** Research software and non-image verification materials are available in GitHub Release v1.0.0 (<https://github.com/jihun1184/topology-preserving-wc-repair/releases/tag/v1.0.0>), commit 68b4cc741d1a0aa0181421c9c03afdc176655f56. The release provides preserved bounded-search sources, executable procedures for four selected BraTS2021_00110 searches, re-aggregation of 19,040 recorded result rows, and boundary-rank, Betti-vector and SAFE verification of three Table S4 cell ledgers. Re-aggregation and ledger verification do not rerun the searches; full-cohort search regeneration is not provided. Code/software documentation are MIT-licensed; author-generated non-image records are CC BY 4.0, within the documented file-specific scope.
- **Use of AI tools.** OpenAI's ChatGPT, including Codex, was used to assist with language editing and figure preparation. All mathematical statements, proofs, computational results, and final presentation materials were proposed and verified by the authors, who take full responsibility for the submitted work.
- **Online Resources.** Online Resource 1: proof of homotopy preservation under SAFE flips. Online Resource 2: execution provenance and reproducibility record. Online Resource 3: Tables S1–S4 and Figure S1. Online Resource 4: Supplementary Data archive and Table S4 verification records.

References

- [1] Siqueira, M., Latecki, L.J., Tustison, N.J., Gallier, J.H., Gee, J.C.: Topological repairing of 3D digital images. *J Math Imaging Vis* **30**(3), 249–274 (2008) <https://doi.org/10.1007/s10851-007-0054-1>
- [2] Latecki, L.J., Eckhardt, U., Rosenfeld, A.: Well-composed sets. *Comput Vis Image Underst* **61**(1), 70–83 (1995) <https://doi.org/10.1006/cviu.1995.1006>
- [3] Latecki, L.J.: 3D well-composed pictures. *Graph Models Image Process* **59**(3), 164–172 (1997) <https://doi.org/10.1006/gmip.1997.0422>
- [4] Boutry, N., Géraud, T., Najman, L.: A tutorial on well-composedness. *J Math Imaging Vis* **60**(3), 443–478 (2018) <https://doi.org/10.1007/s10851-017-0769-6>
- [5] Boutry, N.: Introducing poset-based connected n -manifolds and P -well-composedness in partially ordered sets. *J Math Imaging Vis* **66**(1), 1–21 (2024) <https://doi.org/10.1007/s10851-023-01159-6>
- [6] Kong, T.Y.: Topology-preserving deletion of 1's from 2-, 3- and 4-dimensional binary images. In: DGCI 1997. LNCS, vol. 1347, pp. 3–18 (1997). <https://doi.org/10.1007/BFB0024826>
- [7] Bertrand, G., Malandain, G.: A new characterization of three-dimensional simple points. *Pattern Recognit Lett* **15**(2), 169–175 (1994) [https://doi.org/10.1016/0167-8655\(94\)90046-9](https://doi.org/10.1016/0167-8655(94)90046-9)
- [8] Niethammer, M., Kalies, W.D., Mischaikow, K., Tannenbaum, A.: On the detection of simple points in higher dimensions using cubical homology. *IEEE Trans Image Process* **15**(8), 2462–2469 (2006) <https://doi.org/10.1109/TIP.2006.877309>
- [9] Rosenfeld, A., Kong, T.Y., Nakamura, A.: Topology-preserving deformations of two-valued digital pictures. *Graph Models Image Process* **60**(1), 24–34 (1998) <https://doi.org/>

[10.1006/gmip.1997.0459](https://doi.org/10.1006/gmip.1997.0459)

- [10] Boutry, N., Géraud, T., Najman, L.: How to make nD images well-composed without interpolation. In: Proc IEEE ICIP, pp. 2149–2153 (2015). <https://doi.org/10.1109/ICIP.2015.7351181>
- [11] Tustison, N.J., Avants, B.B., Siqueira, M., Gee, J.C.: Topological well-composedness and glamorous glue: a digital gluing algorithm for topologically constrained front propagation. *IEEE Trans Image Process* **20**(6), 1756–1761 (2011) <https://doi.org/10.1109/TIP.2010.2095021>
- [12] Boutry, N., Géraud, T., Najman, L.: On making nD images well-composed by a self-dual local interpolation. In: DGCI 2014. LNCS, vol. 8668, pp. 320–331 (2014). https://doi.org/10.1007/978-3-319-09955-2_27
- [13] Gonzalez-Diaz, R., Jimenez, M.-J., Medrano, B.: 3D well-composed polyhedral complexes. *Discrete Appl Math* **183**, 59–77 (2015) <https://doi.org/10.1016/j.dam.2014.08.036>
- [14] Boutry, N., González Díaz, R., Jiménez Rodríguez, M.J.: Weakly well-composed cell complexes over nD pictures. *Inf Sci* **499**, 62–83 (2019) <https://doi.org/10.1016/j.ins.2018.06.005>
- [15] Malgouyres, R., More, M.: On the computational complexity of reachability in 2D binary images and some basic problems of 2D digital topology. *Theor Comput Sci* **283**(1), 67–108 (2002) [https://doi.org/10.1016/S0304-3975\(01\)00054-8](https://doi.org/10.1016/S0304-3975(01)00054-8)
- [16] Boutry, N., González Díaz, R., Najman, L., Géraud, T.: Continuous well-composedness implies digital well-composedness in n-D. *J Math Imaging Vis* **64**(2), 131–150 (2022) <https://doi.org/10.1007/s10851-021-01058-8>
- [17] Boutry, N., Gonzalez-Diaz, R., Jimenez, M.-J., Paluzo-Hildago, E.: Strong Euler well-composedness. *J Comb Optim* **44**(4), 3038–3055 (2022) <https://doi.org/10.1007/s10878-021-00837-8>
- [18] Bazin, P.-L., Pham, D.L.: Topology correction of segmented medical images using a fast marching algorithm. *Comput Methods Programs Biomed* **88**(2), 182–190 (2007) <https://doi.org/10.1016/j.cmpb.2007.08.006>
- [19] Dupas, A., Damiand, G., Lachaud, J.-O.: Multi-label simple points definition for 3D images digital deformable model. In: DGCI 2009. LNCS, vol. 5810, pp. 156–167 (2009). https://doi.org/10.1007/978-3-642-04397-0_14
- [20] Mazo, L., Passat, N., Couprie, M., Ronse, C.: Topology on digital label images. *J Math Imaging Vis* **44**(3), 254–281 (2012) <https://doi.org/10.1007/s10851-011-0325-8>
- [21] Mazo, L., Passat, N., Couprie, M., Ronse, C.: Digital imaging: a unified topological framework. *J Math Imaging Vis* **44**(1), 19–37 (2012) <https://doi.org/10.1007/s10851-011-0308-9>
- [22] Baid, U., Ghodasara, S., Mohan, S., et al.: The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv:2107.02314* (2021). <https://doi.org/10.48550/arXiv.2107.02314>
- [23] Menze, B.H., Jakab, A., Bauer, S., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans Med Imaging* **34**(10), 1993–2024 (2015) <https://doi.org/10.1109/TMI.2014.2377694>
- [24] Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J.S., Freymann, J.B., Farahani, K., Davatzikos, C.: Advancing The Cancer Genome Atlas glioma MRI collections with expert segmentation labels and radiomic features. *Sci Data* **4**, 170117 (2017) <https://doi.org/10.1038/sdata.2017.117>
- [25] Baid, U., Ghodasara, S., Mohan, S., et al.: RSNA-ASNR-MICCAI-BraTS-2021 Dataset. The Cancer Imaging Archive (2023). <https://doi.org/10.7937/jc8x-9874>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [OnlineResource3.pdf](#)
- [OnlineResource2.pdf](#)
- [OnlineResource4.zip](#)
- [OnlineResource1.pdf](#)