A characterization of discrete homotopy distance in complex networks

Bradley Z. Vigil*
Department of Mathematics and Statistics, Texas Tech University, Lubbock, Texas 79409, USA

Topological signatures of spreading dynamics on the brain connectome—so-called *brain chains*—distinguish subtypes of Alzheimer's disease. The underlying observable, the discrete homotopy distance between trajectories, is defined through a poset $\mathcal{H}(N)$ whose size grows super-exponentially in the number of nodes N, rendering direct computation infeasible at connectome scale. We prove that this distance attains its natural lower bound, the number of positions in which two trajectories differ, if and only if the trajectories are *interleaved*. Interleaving is a coherent alternation condition testable in $O(n)$ time directly from the two trajectories, with no construction of $\mathcal{H}(N)$. We give an equivalent discrepancy-set formulation and show that the criterion uses only the covering relations along the two chains, so it holds for arbitrary saturated chains of equal length. We demonstrate scaling to $N \sim 90$, where full poset construction is impractical. The result converts a global, super-exponential computation into a local, linear-time test for the tightest and most common case.

I. INTRODUCTION

Complex networks provide a unifying language for physical, biological, and social systems, and a central theme in their study is how the structure of a network constrains the dynamics it supports [2]. In neurodegeneration, the progression of Alzheimer's disease can be read off the way pathological load spreads over the structural brain connectome. Recently, Goodbrake *et al.* [1] showed that a topological encoding of this spreading, the notion of *brain chains*, yields equivalence classes that distinguish distinct simulated disease subtypes. The object underlying these signatures is a graded poset $\mathcal{H}(N)$: the spanning subgraphs of the complete graph on N nodes, quotiented by homotopy equivalence and ordered by inclusion of homotopy type. Trajectories of the dynamics correspond to maximal chains of $\mathcal{H}(N)$, and the dissimilarity of two trajectories is measured by the discrete homotopy distance $d_{\mathcal{H}}$ between the corresponding chains.

The bottleneck is computational. Evaluating $d_{\mathcal{H}}$ as originally defined requires building the entire poset $\mathcal{H}(N)$, whose size grows super-exponentially in N; direct computation is already onerous near $N \sim 20$ and hopeless at the scale of a full connectome ($N = 83$). This is a familiar obstruction in complex-systems physics: a physically meaningful observable is defined through a global construction that does not scale. A natural question is whether the observable can be certified *locally*, from the two trajectories alone, without materializing the global state space.

Here we answer this in the affirmative for the tightest and most common case. The distance $d_{\mathcal{H}}$ is bounded below by the number of positions at which two chains differ; we prove that this bound is attained exactly when the chains are *interleaved*, a coherent alternation condition checkable in $O(n)$ time directly from the two chains (Theorem 1 and Algorithm 1). We give an equivalent formulation in terms of the discrepancy set, note that it uses only the covering relations along the chains and hence extends beyond graded posets, validate the characterization by exhaustive brute-force computation on random posets, and show empirically that the interleaving test scales gracefully to $N \sim 90$ where full poset construction is infeasible. The result turns a global, super-exponential computation into a local, linear-time certificate for the equality case, and clarifies which trajectory differences are "minimal" in a precise combinatorial sense.

II. SETUP

Let P be a poset. We write $x \prec\!\!\cdot\ y$ for the *covering relation* and say x is *covered by* y when $x < y$ and no z satisfies $x < z < y$; that is, y sits immediately above x with nothing strictly between. A chain $x_0 < x_1 < \cdots < x_m$ is *saturated* if every consecutive pair is a cover, $x_k \prec\!\!\cdot\ x_{k+1}$, equivalently, it admits no refinement by inserting an element strictly between two of its members. A *maximal chain* is a saturated chain running from a minimal element to a maximal element of P. When P is graded with maximum rank n, every maximal chain has the form $F = (z_0^F \prec\!\!\cdot\ z_1^F \prec\!\!\cdot\ \cdots \prec\!\!\cdot\ z_n^F)$, where z_k^F is the rank-k element of F; adjacency and the homotopy distance then act by replacing individual z_k^F while preserving saturation. We recall three notions from [1] (Defs. 5.1–5.3 therein), and write F, G for two trajectories (maximal chains) whose distance we wish to compute.

Definition 1 (adjacency)**.** Two maximal chains F_1, F_2 are *adjacent* if they differ in exactly one position.

Adjacency is the elementary move from which the distance is built; two adjacent chains are shown in Fig. 1.

Definition 2 (discrete homotopy and distance)**.** A sequence $\{F_i\}_{i=0}^{d}$ of maximal chains with F_k, F_{k+1} adjacent for all k is a *discrete homotopy* of length d. The *discrete homotopy distance* $d_{\mathcal{H}}(F_I, F_F)$ is the length of the shortest such sequence with $F_0 = F_I$ and $F_d = F_F$.

Since each adjacency changes exactly one position, $d_{\mathcal{H}}(F, G)$ is bounded below by the number of positions

* bradley.vigil@ttu.edu

where F and G differ, as already noted in [1]. We characterize when equality holds.

Definition 3 (interleaving)**.** Maximal chains F, G of rank n are *interleaved* if, for every $k \in \{0, \dots, n-1\}$,

$$z_k^F \prec z_{k+1}^G \quad \text{or} \quad z_k^G \prec z_{k+1}^F. \tag{1}$$

In words, Eq. (1) requires that at every rank the two chains *alternate coherently*: the rank-k element of one chain is covered by the rank-$(k{+}1)$ element of the other. This is precisely the local "zig-zag" that lets one chain be deformed into the other by changing each differing position only once, as the next section makes precise.

III. CHARACTERIZATION

Our main result is a local necessary-and-sufficient condition for the lower bound to be tight.

Theorem 1 (main result)**.** *Let F, G be maximal chains of rank n in a graded poset P. Then $d_{\mathcal{H}}(F, G)$ equals the number of positions where F and G differ if and only if F and G are interleaved [Eq. (1)].*

Interleaving is automatic wherever the chains agree, so only the positions of disagreement matter. Write the *discrepancy set* $D = \{\, i \mid z_i^F \neq z_i^G \,\}$ and let $m_D = \max D$ be its largest index. Restricting Eq. (1) to D yields the form we actually compute.

Theorem 2 (equivalent, computable form)**.** $d_{\mathcal{H}}(F, G) = |D|$ *if and only if* $F = G$, *or* $|D| = 1$, *or for every* $k \in D \setminus \{m_D\}$,

$$z_k^F \prec z_{k+1}^G \quad \textit{or} \quad z_k^G \prec z_{k+1}^F. \tag{2}$$

In words, Eq. (2) is exactly the interleaving condition restricted to the differing positions, minus its trivial cases: $F = G$ and $|D| = 1$ are settled immediately, and the largest differing index m_D imposes no forward cover because F and G already coincide at every rank above it. Theorems 1 and 2 are therefore equivalent: positions where the chains agree impose no new constraint, so Eq. (2) follows from Eq. (1) by focusing on D alone. Notably, the proof below manipulates only the covering relations along F and G and never appeals to the rank function elsewhere; the characterization is thus not special to graded posets.

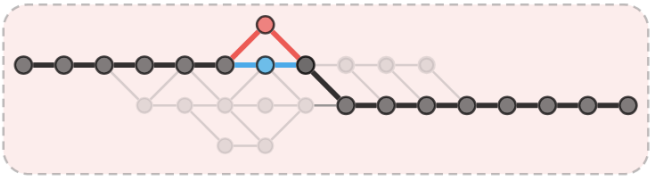

FIG. 1. Two adjacent maximal chains differing only at the red and cyan nodes.

Remark 1 (generalization)**.** Theorems 1 and 2 hold verbatim for any two *saturated* chains of equal finite length sharing their endpoints in an arbitrary poset, graded or not: the only structure used is that consecutive elements of each chain form a cover, together with the cross-covers in Eqs. (1)–(2). In $\mathcal{H}(N)$, where all maximal chains share a common length [3], the hypotheses hold automatically and Theorem 2 applies to every pair of trajectories.

We sketch the proof of Theorem 2; the full construction, the equivalence with Theorem 1, and the extension to saturated chains are given in the Supplemental Material [4].

Proof sketch. The condition (2) is nontrivial only on the *internal edges* of D, i.e. indices k with $k, k+1$ both discrepant; elsewhere one of the covers holds automatically because F and G are saturated. *(Sufficiency.)* We correct the chain one position at a time, replacing the current value at a position by its G-value; each such replacement is a single-position move, hence an adjacency, provided the two flanking covers hold. For an internal edge k, interleaving supplies $z_k^F \prec z_{k+1}^G$ or $z_k^G \prec z_{k+1}^F$, which orients that edge, telling us which of $k, k+1$ must be corrected first. The discrepant positions form disjoint runs of consecutive indices, so these orientations live on paths and never create a cycle; any topological order of the resulting acyclic constraints corrects all $|D|$ positions in $|D|$ adjacencies. With the lower bound $d_{\mathcal{H}} \geq |D|$ this gives equality. *(Necessity, contrapositive.)* If an internal edge k satisfies neither cover, then the step correcting k would require $k+1$ already corrected, while the step correcting $k+1$ would require k already correcte. No length-$|D|$ homotopy can satisfy both, so $d_{\mathcal{H}}(F, G) \geq |D| + 1$. Full details, including the parity lower bound and the extension to saturated chains, appear in the Supplemental Material [4]. □

Figure 2 illustrates the construction: a discrete homotopy of length four between two interleaved maximal chains in $\mathcal{H}(6)$, built by Algorithm 2.

The construction and its test are summarized in two algorithms: Algorithm 1 *tests* the interleaving condition on a pair of chains—returning TRUE exactly when $d_{\mathcal{H}}(F, G) = |D|$, while Algorithm 2 *constructs* an explicit length-$|D|$ homotopy whenever the chains are interleaved. In Algorithm 2 the flanking covers $z_{\ell-1}^{F_t} \prec z_\ell^G \prec z_{\ell+1}^{F_t}$ are evaluated in the *current* chain F_t, not in the original F; the acyclicity established in the proof guarantees that some admissible ℓ always exists until R is empty (see Fig. S1 of the Supplemental Material [4] for the decision loop).

Algorithm 1: Interleaving equality test

1. **Input:** maximal chains F, G of rank n in a graded poset.
2. **Output:** TRUE iff $d_{\mathcal{H}}(F, G) = |D|$.
3. $D \leftarrow \{\, i : z_i^F \neq z_i^G \,\}$; if $|D| \leq 1$ return TRUE.

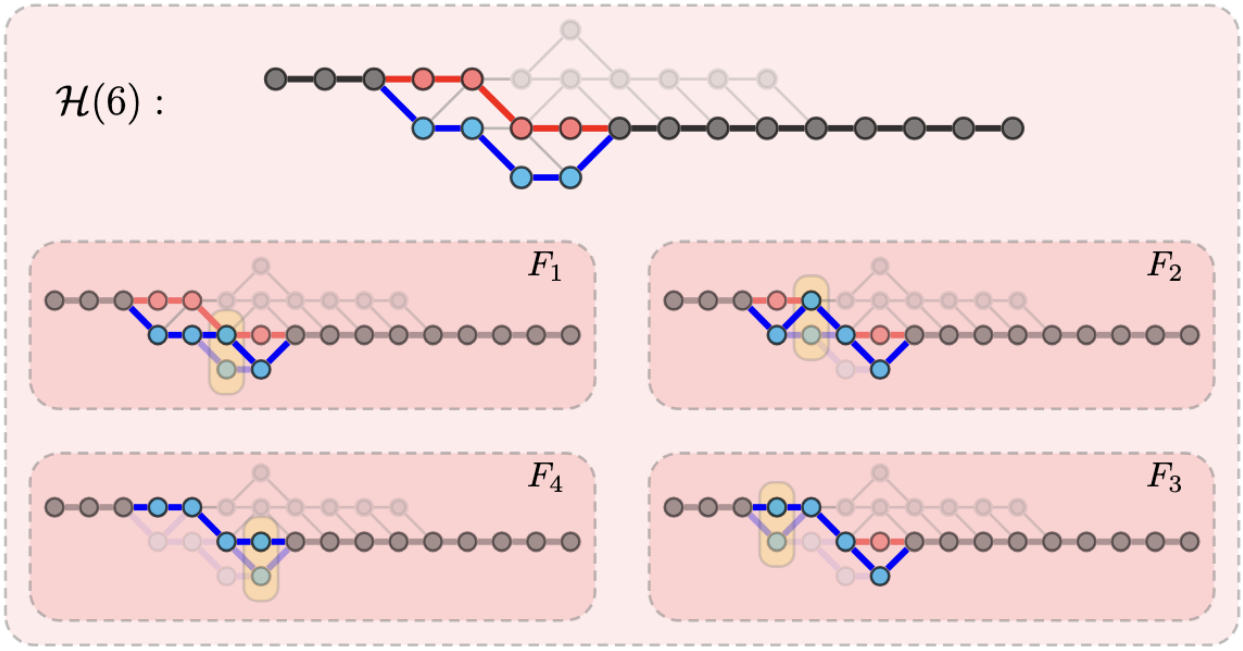

FIG. 2. A discrete homotopy of length four between two interleaved maximal chains F and G in $\mathcal{H}(6)$ (top; the two colored paths), generated by the one-position-at-a-time correction of Algorithm 2. The panels $F_1 \to F_2 \to F_3 \to F_4$ are the successive intermediate chains, with highlighted nodes marking the single position resolved at each step.

4. $m \leftarrow \max D$.
5. for each $k \in D \setminus \{m\}$: if not $\big(z_k^F \prec z_{k+1}^G$ or $z_k^G \prec z_{k+1}^F\big)$ return FALSE.
6. return TRUE. ($O(n)$ per pair)

Algorithm 2: Length-$|D|$ homotopy (sufficiency)

1. $F_0 \leftarrow F$; $t \leftarrow 0$; $R \leftarrow D$.
2. while $R \neq \varnothing$: pick any $\ell \in R$ whose flanking covers hold, $z_{\ell-1}^{F_t} \prec z_\ell^G \prec z_{\ell+1}^{F_t}$; set $F_{t+1} \leftarrow F_t(z_\ell^{F_t} \to z_\ell^G)$; $R \leftarrow R \setminus \{\ell\}$; $t \leftarrow t+1$.
3. return $F_0, \ldots, F_t$.

By Remark 1 the hypotheses always hold in $\mathcal{H}(N)$, so Theorem 2 applies verbatim there. Exhaustive application of Algorithm 1 shows that $N = 6$ is the smallest N for which $\mathcal{H}(N)$ contains a non-interleaved pair—i.e. the smallest N admitting two trajectories whose distance strictly exceeds their number of differing positions. The four posets in Fig. 3, by contrast, are interleaved for *every* pair of maximal chains.

Worked example. Take a rank-3 poset with unique bottom $\hat{0}$ and top $\hat{1}$, rank-1 elements $\{a_1, a_2\}$, rank-2 elements $\{b_1, b_2\}$, and covers $\hat{0} \prec a_i$, $a_i \prec b_j$ for all i, j, and $b_j \prec \hat{1}$. For $F = (\hat{0}, a_1, b_1, \hat{1})$ and $G = (\hat{0}, a_2, b_2, \hat{1})$ we have $D = \{1, 2\}$; the only test is at $k = 1$, where $a_1 \prec b_2$ holds, so the chains are interleaved and $d_{\mathcal{H}} = 2$, realized by $F \to (\hat{0}, a_1, b_2, \hat{1}) \to G$. If instead a_i covered only b_i, no cover holds at $k = 1$, the chains are not interleaved, and $d_{\mathcal{H}} = 3 > |D|$.

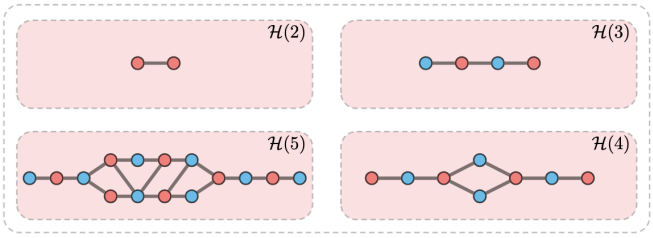

FIG. 3. The posets $\mathcal{H}(2)$, $\mathcal{H}(3)$, $\mathcal{H}(4)$, and $\mathcal{H}(5)$: any two of their maximal chains are interleaved, so every pairwise distance equals the number of differing positions.

IV. COMPUTATIONAL RESULTS

Scaling. We compared the cost of building $\mathcal{H}(N)$ against the interleaving test (Fig. 4). Construction of $\mathcal{H}(N)$ grows rapidly and becomes impractical near $N \sim 21$, whereas the test—being $O(n)$ per pair—remains cheap out to $N \sim 90$, precisely the regime relevant to whole-connectome analysis. For the Alzheimer's subtypes of [1], the interleaved cases (Limbic, MTL, Posterior, Temporal for $N = 4, 6, 8, 12$; Limbic/MTL and MTL/Posterior/Temporal for $N = 15$) thus have their distances certified without ever building the poset.

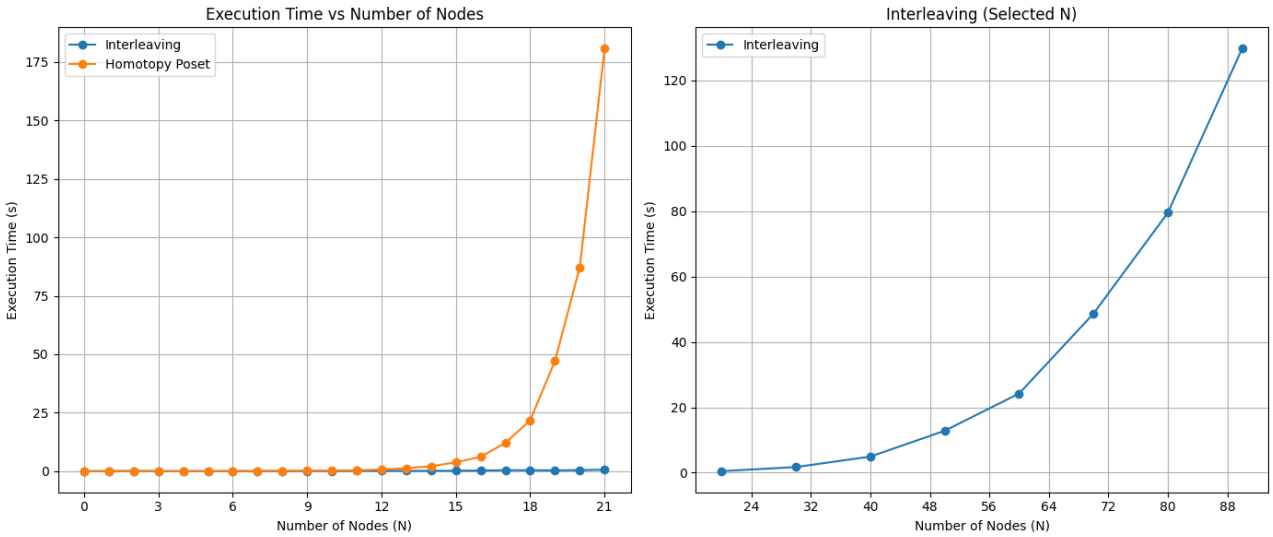

FIG. 4. Cost of constructing $\mathcal{H}(N)$ versus the $O(n)$ interleaving test as the number of nodes N grows. Left: both procedures for $N \in \{0, \ldots, 21\}$; construction of $\mathcal{H}(N)$ diverges rapidly. Right: the interleaving test alone for $N \in \{20, \ldots, 90\}$, reaching regimes where full construction is infeasible.

V. DISCUSSION

We have shown that the equality case of the discrete homotopy distance is governed by a single local condition, interleaving, that is checkable in linear time without constructing the exponentially large state space $\mathcal{H}(N)$. In plain terms, Theorem 1 says that two trajectories are as close as their raw number of differences allows precisely when those differences are laid out so that one can alternate corrective moves without conflict. Physically, interleaving identifies pairs of connectome trajectories whose differences are "coherent," distributed so that no redundant intermediate configuration is required to deform one into the other; such pairs plausibly correspond to functional pathways with similar propagation behavior. We stress the scope: the criterion certifies the tight case $d_{\mathcal{H}} = |D|$ and, when it fails, guarantees $d_{\mathcal{H}} > |D|$ without returning the exact value.

Several directions follow. Empirically, why are the MTL and Limbic subtypes interleaved across $N = 4, 6, 8, 12, 15$, and is this forced by the connectome or by the spreading model [2]? Algorithmically, interleaved segments suggest "collapsible" subchains whose removal would restrict computation to the non-interleaved regions that actually contribute to $d_{\mathcal{H}}$, a route toward a general scalable distance; the test also parallelizes trivially across chain pairs. The generalization of Remark 1 further invites application to arbitrary posets with arbitrary gradings arising in other complex-systems settings. Full proofs of Theorems 1 and 2 are given in the Supplemental Material [4]; a companion Regular Article will develop the $\mathcal{H}(N)$-specific analysis and these extensions in depth.

The verification code and $\mathcal{H}(N)$ construction are available at `github.com/B-Vigil/discrete-homotopy`.

ACKNOWLEDGMENTS

The author thanks the developers of [1] for making their framework available.

[1] C. Goodbrake, D. Beers, T. B. Thompson, H. A. Harrington, and A. Goriely, *Brain chains as topological signatures for Alzheimer's disease*, J. Appl. Comput. Topol. **8**, 1257 (2024).

[2] A. Goriely, E. Kuhl, and C. Bick, *Neuronal oscillations on evolving networks: Dynamics, damage, degradation, decline, dementia, and death*, Phys. Rev. Lett. **125**, 128102 (2020).

[3] R. Stanley, *Some aspects of groups acting on finite posets*, J. Combin. Theory Ser. A **32**, 132 (1982).

[4] See Supplemental Material for the full proofs of Theorems 1 and 2, the equivalence of the two formulations, and the extension of the characterization to saturated chains in arbitrary posets, which includes [1, 3].

Supplemental Material:
A characterization of discrete homotopy distance in complex networks

Bradley Z. Vigil*

Department of Mathematics and Statistics,

Texas Tech University, Lubbock, Texas 79409, USA

Abstract

This Supplemental Material provides complete proofs of the two theorems stated in the main text. We fix notation and record the elementary lower bound (Sec. S1), prove that the interleaving criterion and its discrepancy-set reformulation are equivalent (Sec. S2), and give the full constructive proof of the characterization (Sec. S3). The proof is organized around the observation that the interleaving condition orients the "internal edges" of the discrepancy set; because these edges form paths, the induced constraints are acyclic and admit a valid correction order. Finally we show that the argument uses only the covering relations along the two chains and therefore extends verbatim from graded posets to arbitrary saturated chains of equal length (Sec. S4).

S1. PRELIMINARIES AND NOTATION

Let P be a poset and let $x \prec y$ denote the covering relation: $x < y$ and there is no z with $x < z < y$. A chain $x_0 < x_1 < \cdots < x_m$ is *saturated* if $x_k \prec x_{k+1}$ for every k; a *maximal chain* is a saturated chain from a minimal to a maximal element of P.

Throughout, F and G are saturated chains of the same finite length n that share their endpoints,

$$F = (z_0^F \prec \cdots \prec z_n^F), \qquad G = (z_0^G \prec \cdots \prec z_n^G), \tag{S1}$$

with $z_0^F = z_0^G$ and $z_n^F = z_n^G$. (In a graded poset P with a unique minimum and maximum—in particular in $\mathcal{H}(N)$ [2]—every pair of maximal chains is of this form.) We write

$$D = \{\, i : z_i^F \neq z_i^G \,\}, \qquad m_D = \max D, \tag{S2}$$

so that $D \subseteq \{1, \ldots, n-1\}$ and $0, n \notin D$. For a saturated chain H and an element x we write $H(z_\ell^H \to x)$ for the sequence obtained from H by replacing its rank-ℓ entry with x.

We say indices k and $k+1$ span an *internal edge* of D if $\{k, k+1\} \subseteq D$. For each k define the two *cross-cover* predicates

$$A_k : \; z_k^F \prec z_{k+1}^G, \qquad B_k : \; z_k^G \prec z_{k+1}^F. \tag{S3}$$

Definition 1 (interleaving)**.** F and G are *interleaved* if $A_k \vee B_k$ holds for every $k \in \{0, \ldots, n-1\}$.

* bradley.vigil@ttu.edu

Lemma S1 (single-move validity)**.** *Let H be a saturated chain of length n sharing its endpoints with G, and let $\ell \in \{1, \dots, n-1\}$ with $z_\ell^H \neq z_\ell^G$. If*

$$z_{\ell-1}^H \prec z_\ell^G \prec z_{\ell+1}^H, \tag{S4}$$

then $H' = H(z_\ell^H \to z_\ell^G)$ is a saturated chain of length n, and H, H' are adjacent (they differ in exactly position ℓ).

Proof. H' differs from H only at position ℓ, so it differs in exactly one position and its endpoints are unchanged. Every consecutive pair of H' other than $(\ell-1, \ell)$ and $(\ell, \ell+1)$ is a cover because it is a cover in H; the two remaining pairs are covers by (S4). Hence H' is saturated, and a saturated sequence is a chain. Adjacency is Definition 1 of the main text. □

S2. EQUIVALENCE OF THE TWO FORMULATIONS

The main text states the characterization first as interleaving (Theorem 1) and then in the discrepancy-set form used for computation (Theorem 2). We show the two conditions coincide.

Proposition S1. *The following are equivalent:*

(i) F and G are interleaved (Definition 1);

(ii) $F = G$, or $|D| = 1$, or $A_k \vee B_k$ holds for every $k \in D \setminus \{m_D\}$.

Proof. We first note that $A_k \vee B_k$ holds automatically unless k and $k+1$ are both in D. Indeed:

- if $k \notin D$, then $z_k^F = z_k^G$; if moreover $k+1 \notin D$ then $z_{k+1}^G = z_{k+1}^F$ and A_k reads $z_k^F \prec z_{k+1}^F$, true because F is saturated; if instead $k+1 \in D$ then B_k reads $z_k^G = z_k^F \prec z_{k+1}^F$, again true;
- if $k \in D$ but $k+1 \notin D$, then $z_{k+1}^G = z_{k+1}^F$ and A_k reads $z_k^F \prec z_{k+1}^F$, true.

Thus the only indices at which $A_k \vee B_k$ can fail are the internal edges of D, and every internal edge k satisfies $k \in D$ and $k+1 \in D$, hence $k \leq m_D - 1$, so $k \in D \setminus \{m_D\}$.

$(i) \Rightarrow (ii)$: interleaving asserts $A_k \vee B_k$ for all k, in particular for all $k \in D \setminus \{m_D\}$.

$(ii) \Rightarrow (i)$: if $F = G$ or $|D| = 1$ there are no internal edges, so by the paragraph above $A_k \vee B_k$ holds for every k. Otherwise (ii) gives $A_k \vee B_k$ for $k \in D \setminus \{m_D\}$, which includes every internal edge; at all other k the disjunction holds automatically. Hence F, G are interleaved. □

Because of Proposition S1 it suffices to prove the characterization in the interleaving form; the discrepancy-set statement then follows.

S3. PROOF OF THE CHARACTERIZATION

Theorem S1 (main text, Theorems 1 and 2)**.** *$d_{\mathcal{H}}(F, G)$ equals $|D|$ if and only if F and G are interleaved.*

The lower bound $d_{\mathcal{H}}(F, G) \geq |D|$ is established in [1]. It remains to prove that interleaving is sufficient (Sec. S3 B) and necessary (Sec. S3 C) for equality.

A. Orientation of internal edges

Fix an internal edge k (so $k, k+1 \in D$). During any position-by-position correction of F into G—in which each discrepant position is set once, from its F-value to its G-value—consider the two moves that touch positions k and $k+1$. By Lemma S1, setting position $k+1$ to z^G_{k+1} requires its left flank to cover it; if position k has not yet been corrected, that flank is z^F_k, and the requirement is exactly A_k. Symmetrically, setting position k to z^G_k requires its right flank to be covered; if $k+1$ has not yet been corrected, that flank is z^F_{k+1}, and the requirement is exactly B_k. Hence:

$$\begin{aligned} \neg A_k &\;\Rightarrow\; k \text{ must be corrected before } k+1, \\ \neg B_k &\;\Rightarrow\; k+1 \text{ must be corrected before } k. \end{aligned} \tag{S5}$$

Interleaving $(A_k \vee B_k)$ guarantees that at most one of the two implications in (S5) fires, so each internal edge imposes *at most one* precedence constraint between its endpoints.

B. Sufficiency

Assume F, G interleaved. The discrepant positions D decompose into maximal runs of consecutive integers; distinct runs are separated by at least one agreement position, and a

correction at a position with an agreement neighbor imposes no constraint on the neighbor (its flank is a shared G-value, covered by saturation of G). It therefore suffices to order the corrections within a single run $R = \{a, a+1, \ldots, b\} \subseteq D$, with $a-1, b+1 \notin D$.

The precedence constraints (S5) within R are carried by its internal edges $a, a{+}1, \ldots, b{-}1$, each contributing at most one directed constraint between consecutive elements. These constraints live on the path $a-(a{+}1)-\cdots-b$. A path is a tree, and *any* orientation of the edges of a tree is acyclic; hence the constraint digraph on R has no directed cycle. A finite acyclic digraph always has a source—a vertex with no incoming arrow—so at every stage some position of R has all of its forced predecessors already corrected and can be processed next; equivalently, the digraph admits a topological order π of its elements. This is exactly the guarantee that Algorithm 2 of the main text never stalls: a position with valid flanking covers is always available until R is exhausted (Fig. S1).

Correct the positions of R in the order π. We claim every move is valid in the sense of Lemma S1. Consider the move at position $\ell \in R$, when the current chain H holds G-values at the positions already corrected (and at all $i \notin D$) and F-values elsewhere. Its flanks:

- Left flank $z^H_{\ell-1}$. If $\ell = a$ or $\ell - 1 \notin D$ or $\ell - 1$ is already corrected, then $z^H_{\ell-1}$ is a G-value (or the shared endpoint value), and $z^H_{\ell-1} \prec z^G_\ell$ because G is saturated. Otherwise $\ell - 1 \in R$ is uncorrected; since π places ℓ after all its forced predecessors, the edge $(\ell-1, \ell)$ did not force "$\ell-1$ before ℓ", so by (S5) $A_{\ell-1}$ holds, i.e. $z^F_{\ell-1} \prec z^G_\ell$, and $z^H_{\ell-1} = z^F_{\ell-1}$.

- Right flank $z^H_{\ell+1}$. If $\ell = b$ or $\ell + 1 \notin D$ or $\ell + 1$ is already corrected, then $z^H_{\ell+1}$ is a G-value and $z^G_\ell \prec z^H_{\ell+1}$ by saturation of G. Otherwise $\ell + 1 \in R$ is uncorrected; since π did not force "$\ell + 1$ before ℓ", by (S5) B_ℓ holds, i.e. $z^G_\ell \prec z^F_{\ell+1}$, and $z^H_{\ell+1} = z^F_{\ell+1}$.

In every case the flanking covers (S4) hold, so by Lemma S1 the move is a valid adjacency and preserves saturation.

Applying this to every run and concatenating, we obtain a discrete homotopy $F = H_0, \ldots, H_{|D|} = G$ that corrects each discrepant position exactly once. Its length is $|D|$, so $d_{\mathcal{H}}(F, G) \leq |D|$; therefore, $d_{\mathcal{H}}(F, G) = |D|$. □

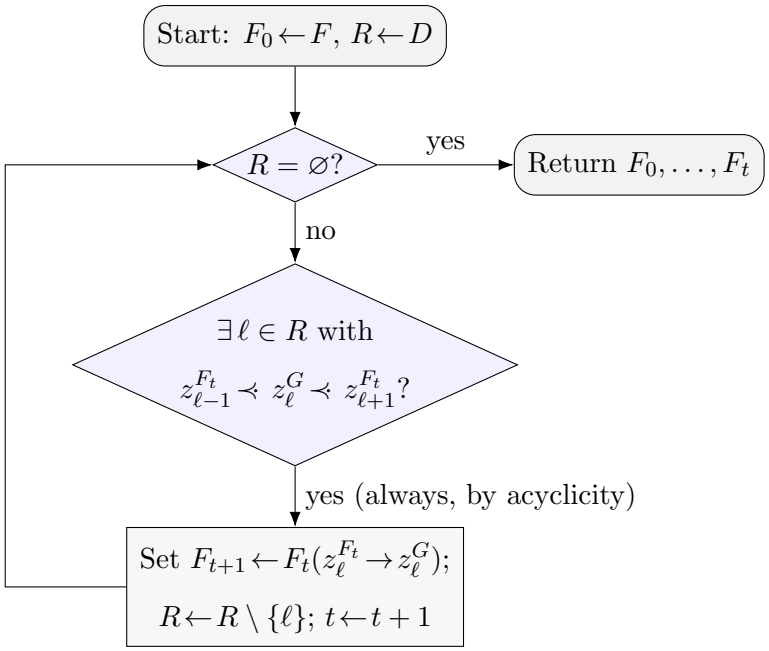

FIG. S1. Decision loop of Algorithm 2 (sufficiency). Under interleaving the constraint digraph on the discrepant positions is acyclic (Sec. S3 B), so a position ℓ with valid flanking covers always exists while $R \neq \varnothing$; the "no" branch of the second test is therefore never taken, and the loop terminates after exactly $|D|$ updates.

C. Necessity

We prove the contrapositive: if F, G are *not* interleaved, then $d_{\mathcal{H}}(F, G) > |D|$. By Proposition S1 the failure occurs at an internal edge k, so $k, k+1 \in D$ and both A_k and B_k fail:

$$z^F_k \not\prec z^G_{k+1} \qquad \text{and} \qquad z^G_k \not\prec z^F_{k+1}. \tag{S6}$$

Suppose, for contradiction, that $d_{\mathcal{H}}(F, G) = |D|$, and fix a homotopy $F = H_0, \ldots, H_{|D|} = G$ realizing it. By the equality clause of Lemma **??**, every position of D—in particular k and $k+1$—is altered exactly once, and no position outside D is ever altered. Thus position k is set once, at some step s, from z^F_k to z^G_k, and position $k+1$ is set once, at some step $t \neq s$, from z^F_{k+1} to z^G_{k+1}.

Consider step s. It replaces $z^{H_s}_k = z^F_k$ by z^G_k; validity requires $z^G_k \prec z^{H_s}_{k+1}$. The value $z^{H_s}_{k+1}$ is either z^F_{k+1} (if $t > s$) or z^G_{k+1} (if $t < s$). If $t > s$ then $z^G_k \prec z^F_{k+1}$, i.e. B_k, contradicting (S6); hence $t < s$. Now consider step t. It replaces $z^{H_t}_{k+1} = z^F_{k+1}$ by z^G_{k+1}; validity requires $z^{H_t}_k \prec z^G_{k+1}$. Since $t < s$, position k has not yet been corrected at step t, so $z^{H_t}_k = z^F_k$ and

the requirement is $z_k^F \prec z_{k+1}^G$, i.e. A_k, again contradicting (S6). The two conclusions $t < s$ and $t > s$ are incompatible, so no homotopy of length $|D|$ exists and $d_{\mathcal{H}}(F, G) \geq |D| + 1$. □

Together Secs. S3 B and S3 C prove Theorem S1, and Proposition S1 delivers the discrepancy-set form (Theorem 2 of the main text).

S4. EXTENSION TO SATURATED CHAINS IN ARBITRARY POSETS

Nothing in Secs. S1–S3 uses the rank function of P beyond the two structural facts that (a) F and G are saturated chains of the same length sharing their endpoints, and (b) consecutive entries of each chain form covers. Lemma S1, Lemma **??**, Proposition S1, and Theorem S1 were all stated and proved under exactly these hypotheses. We record the resulting generality.

Proposition S2 (Remark 1 of the main text)**.** *Let P be an arbitrary poset and let F, G be saturated chains of the same finite length sharing their endpoints. Then $d_{\mathcal{H}}(F, G)$ equals the number of positions in which F and G differ if and only if F and G are interleaved. In particular, the characterization is not restricted to graded posets; the grading of $\mathcal{H}(N)$ enters only to guarantee, via a unique minimum and maximum, that any two maximal chains automatically satisfy hypotheses (a)–(b) [2].*

Proof. Immediate from Theorem S1, whose proof invokes only (a) and (b). □

This is the sense in which the criterion is genuinely local: it inspects only the covering relations along the two given chains and the finitely many cross-covers (S3), and never the ambient poset. Algorithm 1 of the main text evaluates these on the discrepancy set in $O(n)$ cover queries.

[1] C. Goodbrake, D. Beers, T. B. Thompson, H. A. Harrington, and A. Goriely, *Brain chains as topological signatures for Alzheimer's disease*, J. Appl. Comput. Topol. **8**, 1257 (2024).

[2] R. Stanley, *Some aspects of groups acting on finite posets*, J. Combin. Theory Ser. A **32**, 132 (1982).