

Online Resource 1: Persistent-homology connectivity analysis

A1. Overview

Section 3.2 of the main article showed that 2D simulations systematically underestimate the 3D permeability, and Section 3.3 showed that the differences in mean travel time and breakthrough spreading between 2D and 3D depend on the porous medium. Both observations were attributed to a loss of pore network connectivity caused by going from three to two dimensions. This supplement provides a direct quantification of pore space connectivity, so the dimensional bias reported in Sections 3.2 and 3.3 can be linked to a measurable structural cause.

We assess the connectivity in all three media types considered (Fibrous, Gaussian, Overlapping Spheres) at three porosities ($\phi = 0.4, 0.6, 0.8$), in both 2D and 3D. Specifically we distinguish three cases: Case A considers a single 2D slice of the medium only. Case B also considers a single slice, but accounts for the surrounding 3D geometry such that a pore region that appears isolated within the slice can be connected to its neighbors through the third dimension. In Case C, the connectivity of the full 3D domain is assessed.

A2. Methods

We started every analysis from binary image data where 1 (white) indicated pores and 0 (black) grains: a single image represents a two-dimensional domain, or an ordered stack of cross-sectional slices a three dimensional one, stored in array V . For a 2D case the array $V[i, j]$ held the label of the pixel at row i , column j , while for 3D domains, $V[i, j, k]$ held the label of the pixel at row i , column j of the k^{th} input image. We then saved V to disk as a .npz file under the key 'labels'. Throughout the pipeline we used a single physical constant, the voxel edge length $v = 0.8 \mu\text{m}$, to convert voxel distances to micrometers.

A.2.1 The signed Euclidean distance transform

For each voxel we computed the signed depth $d(x)$, which measures how deep the voxel sits inside its own phase. Its value is defined by the straight-line distance from the voxel center to the center of the nearest voxel of the opposite phase, and it is assigned a negative sign for pore voxels and a positive sign for grains. Writing $s(x)$ for the sign (-1 for pore, $+1$ for grain) and v for the voxel edge length:

$$d(x) = s(x) \cdot v \cdot \text{dist}(x, \text{opposite phase}),$$

The voxel distance itself is just the straight-line distance between two voxels, computed from their coordinate indices:

$$\text{dist}(x, y) = \sqrt{(i_x - i_y)^2 + (j_x - j_y)^2 + (k_x - k_y)^2}$$

Figure A1 visualizes the procedure.




Fig. A1 Step by step construction of the signed Euclidean distance transform on a toy 8x8 binary. (a) input binary V , with 1 = pore (white) and 0 = grain (black). (b) distance from each pore voxel to its nearest grain voxel, in voxel edges. (c) same calculation for the grain phase, distance from each grain voxel to its nearest pore voxel. (d) signed Euclidean distance transform $d(x)$, built by combining (b) and (c) and flipping the sign on pore voxels, so pore values are negative, grain values positive, and the interface sits at zero (multiply by $v = 0.8 \mu\text{m}$ to get micrometers). (e) the same field as a heatmap.

A.2.2 Filtration and persistence pairs

We used Gudhi, an open-source library, to compute the topological descriptors of our porous media (Maria et al. 2014; Wagner et al. 2011) from the depth field that was provided as input. Then, using a fixed threshold value t that we raised from low to high, each voxel was classified as either on (its depth was at or below t) or off ($d > t$). Three simple numbers describe the shape of the voxels that are on at every t . β_0 is the number of separate pieces of on-voxels: two on-voxels belong to the same piece when they sit next to each other, or when they are linked by a chain of on-voxels. β_1 is the number of closed rings of on-voxels that wrap around at least one off-voxel. β_2 is the number of closed shells of on-voxels that wrap around at least one off-voxel (3D only). For each feature the library records two values: the value of t at which the feature first appears (its birth, b) and the value of t at which it disappears, either by merging with another feature of the same family or by being filled in by the surrounding on-voxels (its death,

d). The pair (b, d) is stored, by dimension, in three separate lists. Pairs whose lifetime $d - b$ is shorter than one voxel edge are discarded as noise, and the persistence p is defined as

$$p_i = d_i - b_i, \text{ keep feature } i \text{ if } p_i \geq v.$$

Figure A2 illustrates this sweep on the fibrous medium 2D slice, at five thresholds that span the depth range from deep pore (negative dist) to deep grain (positive dist).




Fig. A2 Threshold sweep on Fibrous, $\phi = 0.6$, slice 40. (a) The binary slice (white = pore, black = grain). (b) The signed Euclidean distance field $d(x)$ in μm , blue inside the pore phase, red inside the grain phase, zero at the interface. (c) to (g) On-voxel snapshots at threshold $t = -2, -1, 0, +1, +2 \mu\text{m}$; off voxels are light grey, on pore voxels are dark grey, on grain voxels are red. Panel (e) at $t = 0$ is the snapshot at which $\beta_0, \beta_1, \beta_2$ are counted.

The persistent-homology computation uses the cubical-complex algorithm of Wagner et al. (2011), as implemented in the Gudhi library (Maria et al. 2014). First, Gudhi wraps the voxel grid in a structure called a *cubical complex*. The cubical complex is the voxel grid together with a list of all corners, edges, and faces shared between neighboring voxels (Edelsbrunner and Harer 2010). Next, each voxel in the complex, together with the corners, edges, and faces shared between voxels, is tagged with the depth value of the voxel it belongs to. Gudhi sorts these tagged elements from the lowest depth to the highest and steps through them in that order, which is exactly the order in which our threshold sweep turns voxels on. At each step, the next element is added to the on-set, the three voxel-neighbor tests from the previous paragraph are run on it, and the outcome is recorded as either a birth (a new feature has just appeared, and b is set to the current value of t) or a death (an existing feature has just merged with another, or been filled in, and d is set to the current value of t). Each feature therefore ends up with one (b, d) pair. These pairs are kept in three separate lists, one per feature family: the β_0 list for connected pieces, the β_1 list for closed rings, and the β_2 list for closed shells.

Each persistence pair (b, d) is drawn as one dot in the b – d plot, with birth threshold on the horizontal axis and death threshold on the vertical axis (both in micrometers). The dashed diagonal $d = b$ is the lower bound of every dot, since death always follows birth. The yellow band covers the region $b \leq 0 < d$, the pairs that are alive at the pore-grain interface, and the noise filter $d - b \geq v$ removes dots too close to the diagonal. The number of remaining dots is exactly β_0 (blue circles), β_1 (orange squares), and β_2 (gold triangles). The interpretation of one dot depends on the case: in Case A it is a 2D feature of the standalone image; in Case B it is a 3D feature of the parent geometry whose support touches the slice; in Case C it is a 3D feature of the whole geometry.




Fig. A3 The persistence diagram. Each feature in the domain is one point at $(\text{birth}, \text{death})$. Its color and shape say what kind it is: blue circle = β_0 body, red square = β_1 loop, gold triangle = β_2 pocket (3D only). The plot is split into four quadrants. Q2 (teal) is the region where $\text{birth} \leq 0$ and $\text{death} > 0$ (pore voxels carry negative depth values and grain voxels positive depth values, so $\text{birth} \leq 0$ means the feature is born inside the pore phase and $\text{death} > 0$ means it dies inside the grain phase); these features span the pore–grain interface and are the ones we count. Q3 (gold) holds features that live entirely inside the pore phase (only β_0 can land here). Q1 and Q4 are noise. The dashed line $d = b + v$ sits one voxel above the diagonal; anything below it is dropped as noise. β_0 , β_1 , β_2 for the domain are the number of surviving blue circles, red squares, and gold triangles inside Q2.

A.2.3 Betti numbers

Betti numbers count the number of separate pore bodies (β_0), closed rings (flow loops; β_1) and grain pockets fully enclosed by pore (closed shells; β_2), which is always 0 for 2D. The Betti numbers are determined from the list of persistence pairs with an upper threshold of zero, a value at which all pore voxels are included but none of the grains. This list includes features that have been born ($b \leq 0$) and not yet died (Q2 in Fig. A3, $b \leq 0$ and $d > 0$), with the lifetime filter set to $d - b \geq v$, where v is set to the voxel edge length. When embedding a 2D slice into its local 3D context (i.e. Case B), we use the result of the 3D parent and report three counts: β_0 is the number of 3D pore regions that have at least one voxel on the slice; β_1 is the number of 3D loops whose extent along the slicing axis covers the slice index; and β_2 is the number of sealed 3D

cavities (grain regions with no voxel on the geometry's outer faces) that have at least one voxel on the slice.

After the Betti counts are computed, every voxel of the domain is assigned one of six categories that summarize the role it plays in the topology (Fig. A4), represented by a color. The body/loop/pocket vocabulary follows the pore-scale persistent-homology classification of (Robins et al. 2016) and (Herring et al. 2019), extended here with two additional geometric tags (throat and dead-end) read from the pore-voxel neighbor count. A voxel is colored grey if it belongs to the solid phase (grain). Blue is a pore voxel inside the largest connected pore region; these are the bodies counted by β_0 . Red is a pore voxel sitting next to a grain that is fully surrounded by pore on every side (an enclosed grain); these voxels form the loops counted by β_1 . Gold is a pore voxel bordering a grain region that is fully enclosed by pore in three dimensions (a sealed cavity); these are the pockets counted by β_2 , which only exist in 3D. The last two colors describe the local shape of the pore phase and are not counted by any Betti numbers. A pore voxel is colored purple if it has only 2 pore neighbors in 2D, or 2 or 3 pore neighbors in 3D; these are the throats, where the pore phase narrows to a single voxel wide channel. A pore voxel is colored green if it has only 1 pore neighbor; these are the dead ends, where a pore branch terminates.

A voxel can fit more than one rule at the same time. For example, a pore voxel inside the largest pore region (a body) can also have only one pore neighbor (a dead end), or it can sit right next to an enclosed grain (a loop). When that happens, we keep the more specific label and reject the more general one. The ranking from most general to most specific is body, loop, pocket, throat, dead end, so a body voxel that is also a dead end ends up green, and a body voxel that is also a loop ends up red. Throats and dead ends are shape labels only; they change the color a voxel is drawn in but do not change β_0 , β_1 , or β_2 .




Fig. A4 The six voxel labels were used to sort pore and grain structure into Betti contributors. Every panel is a small grid where each square is one voxel: dark grey is grain, and every other color marks a different role a pore voxel can play. Three of the colors feed straight into a Betti number, blue (body, β_0), red (loop, β_1), gold (pocket, β_2); the other two, purple (throat) and green (dead end), are geometric tags that describe local pore shape but are not counted in any β . The mini examples show how the count is read off: two blue regions give $\beta_0 = 2$, one red ring around a trapped grain gives $\beta_1 = 1$, one grain region fully enclosed by pore (gold) gives $\beta_2 = 1$ (3D only).

A.2.4 Euler characteristic and pore-normalized value

The three Betti counts can be combined into a single number that summarizes the topology of the analysis domain. This number is the Euler characteristic of the pore phase:

$$\chi = \beta_0 - \beta_1 + \beta_2$$

The reason for forming this combination is that β_0 counts how the pore phase fragments into separate pieces, β_2 counts the grain pockets fully surrounded by pore, and β_1 counts the extra routes that link pore points together. The formula therefore subtracts the number of extra routes from the number of features that break connectivity (separate pieces and sealed pockets), so χ records the net deficit of routes in the pore phase. $\chi > 0$ means the pore phase is dominated by fragmentation and pockets and is poorly connected; $\chi < 0$ means it is dominated by extra routes and is well connected; $\chi = 0$ sits on the boundary between the two regimes.

Comparing χ directly across the three cases would mix two different effects: the topology of the domain and its size. A 3D geometry contains far more pore voxels than a 2D slice cut from it, so a geometry can produce a larger χ in magnitude simply by having more voxels available to host features. To remove this size effect, we divide χ by N_p , the number of pore voxels in the analysis domain, and report the result as

$$\tilde{\chi} = \frac{\chi}{N_p}$$

where $\tilde{\chi}$ is the average net topological contribution carried by a single pore voxel. For Case C, N_p is the number of pore voxels in the 3D geometry; for Cases A and B, it is the number of pore voxels in the standalone 2D image as every feature reported is one that touches a single slice.

To compare the topology of 2D (Case A) slice i to 3D (Case C), we use the Symmetric Percentage Difference (SPD):

$$SPD_i = \frac{2(\tilde{\chi}_{A,i} - \tilde{\chi}_C)}{|\tilde{\chi}_{A,i}| + |\tilde{\chi}_C|}$$

Because the denominator is a sum of two non-negative numbers, every SPD value is bound to $-2 \leq SPD_i \leq +2$

We split the bounded SPD axis into four interpretive zones:

- i. SPD = 0: Case A and Case C give the same $\tilde{\chi}$ per voxel; the two analyses agree on the topology.
- ii. SPD between 0 and ± 0.5 : high-similarity zone; the two analyses differ by less than a quarter of their average absolute magnitude.
- iii. SPD between ± 0.5 and ± 1.5 : moderate bias. Negative: 2D looks more connected per voxel than 3D. Positive: 2D looks more fragmented.
- iv. SPD = ± 2 : opposite signs at the same absolute magnitude. The strongest possible topological disagreement: 2D and 3D classify the medium oppositely.

A3. Results

We applied the topology analysis to nine porous medium realizations, i.e. the three media (Fibrous, Gaussian, Overlapping Spheres) at three porosities ($\phi = 0.4, 0.6, 0.8$). For each realization we ran the analysis for the three cases:

- i. Case A treats a single 2D slice of the medium as a complete picture, as if no third dimension existed.
- ii. Case B uses the same 2D slice but counts its connectivity by looking at the surrounding 3D geometry above and below the slice, so a pore region that looks isolated on the flat slice is allowed to connect to its neighbors through the third dimension.
- iii. Case C uses the full 100^3 voxel geometry directly.

(To keep the discussion focused on the two core viewpoints, Case B has been moved out of the main results; the complete Case-B numbers are kept as raw data for any future use.)

This setup gives us a clean comparison: the A versus B difference tells us how much connectivity gets lost when we cut the third dimension, and Case C tells us what the network actually looks like in three dimensions. The full set of 27 measurements (nine realizations $\times$ three cases) is collected in Tables A4–A6, and the figures that follow show the underlying geometry, the threshold sweep, the connectivity classes per voxel, and the persistence diagrams that feed the counts.




Fig. A5 Raw binary inputs at three porosities. Each row is one medium: Fibrous, Gaussian, Overlapping Spheres. The three porosity columns (a. $\varphi = 0.4$, b. $\varphi = 0.6$, c. $\varphi = 0.8$) each contain two panels: on the left the 2D image of 2D slice extracted from 3D parent (white = pore, black = grain), and on the right the same medium as a $100 \times 100 \times 100$ voxel geometry with that same slice highlighted in orange so the reader can locate where the 2D image sits in three dimensions.

A.3.1 Voxel classes and Persistence diagrams

Figure A6 renders the medium–porosity sets. Each panel pairs the Case-A view of the paper-matched 2D slices (rows on the left) with the Case-C view of the full 3D geometry (tall panel on the right).

Figure B6 — Fibrous





Figure B6 — Gaussian






Fig. A6 Voxel-class plates per medium. Each plate has three columns, one per available nominal porosity bracket ϕ_{3D} . In each column the top rows show up to three paper-matched 2D slices (Case A) and the bottom row shows the 3D cube (Case C). Colors mark voxel classes: blue = body, red = loop (B_1), gold = pocket (B_2 , 3D only), purple = throat, green = dead-end (overlaid as scatter for visibility), grey = grain. The Gaussian plate has two columns because no 2D slice is available at $\phi_{3D} = 0.4$.

Figure A6 colors every voxel by class but does not say when each topological feature was born and died as the threshold t swept through the slice or 3D geometry. That timing information is captured in the persistence diagram. Figure A7 shows the same eight sets as Fig. A6, with Case A persistence diagrams on the left and the Case C diagram on the tall right panel.

Figure B7 — Fibrous




Figure B7 — Gaussian




Figure B7 — Overlapping Spheres




Fig. A7 Persistence diagrams per medium. Each plate has the same column layout as Fig. A6: up to three paper-matched 2D slices (Case A) stacked on top and the 3D cube (Case C) at the bottom, one column per available φ_{3D} bracket. Marker shape and color encode dimension: blue circle = β_0 body, red square = β_1 loop, gold triangle = β_2 pocket (3D only). The yellow Q2 quadrant ($b \leq 0, d > 0$) holds the features that count toward $\theta_0, \theta_1, \theta_2$. The dashed diagonal is $d = b$ and the dotted line above it is $d = b + v$, the lifetime cutoff at one voxel edge.

Cases A and C panels differ in two notable ways. First, the red squares that appear in the Case A panels are mostly grain regions that the 2D slice surrounds on all four sides; the 2D analysis is geometrically required to register them as loops, but they are not loops of the 3D pore network. Second, Case C contains roughly 30 to 100 times more red squares than Case A (compare $\beta_1(A)$ and $\beta_1(C)$ in Table A4), and it also contains gold triangles which are entirely absent from Case A. The additional red squares in Case C are loops that close inside the 3D pore network through voxels lying above or below the slice plane, so the 2D analysis cannot detect them. The gold

triangles are pockets, i.e. grain regions fully enclosed by pore, which a single 2D image is geometrically unable to contain. These two differences are the reason why $\beta_1(A)$ is much smaller than $\beta_1(C)$, why $\beta_2(A)$ is zero while $\beta_2(C)$ is non-zero, and ultimately why $\tilde{\chi}(A)$ differs from $\tilde{\chi}(C)$ in Table A5.

A.3.2 $\tilde{\chi}$ versus porosity

In order to explore whether the 2D-vs-3D gap varies with porosity or sits at a constant offset, every configuration is consolidated to a single $(\phi, \tilde{\chi})$ point and the per-medium linear fits are overlaid in Figure A8.




Fig. A8 Pore-normalized Euler characteristic $\tilde{\chi}$ against porosity ϕ for Case A (2D slice) and Case C (3D geometry), with linear fits per medium and per case. Triangle / circle / square = Case A 2D slices grouped by ϕ (0.4 / 0.6 / 0.8); star = Case C 3D geometry. Same marker convention as Fig. 2 of the main article.

Fig. 2 in the main article shows a 2D-vs-3D permeability gap at every porosity and concludes that porosity alone is insufficient to explain the observed dimensional discrepancies. The question Fig. A8 addresses is what structural quantity does explain them and specifically, whether the pore-network connectivity descriptor $\tilde{\chi}$ shows the same 2D-vs-3D gap, on the same nine configurations, that permeability does.

The 3D linear fits (solid lines) are essentially flat in every panel, meaning the 3D connectivity is set by the medium and does not change with porosity. The 2D linear fits (dashed lines) are steep for Fibrous and Gaussian, where the 2D estimate of connectivity changes a lot with porosity, and flat for Overlapping Spheres, where it barely moves, the same medium ordering as the permeability gaps in Figure 2 (see Table A1).

Table A1 Slope $\pm$ standard error (SE) and two-sided p -value of the linear fits of $\tilde{\chi}$ versus porosity, per medium and case (A = 2D slice, C = 3D geometry)

MEDIUM		SLOPE $\pm$ SE	P-VALUE
Fibrous	2D	$-3.299\text{e-}02 \pm 2.23\text{e-}03$	< 0.001
	3D	$7.965\text{e-}05 \pm 9.80\text{e-}04$	0.948

Gaussian	2D	$-3.241\text{e-}02 \pm 1.58\text{e-}03$	< 0.001
	3D	$1.519\text{e-}03 \pm 2.35\text{e-}03$	0.634
Overlapping Spheres	2D	$-2.109\text{e-}05 \pm 3.34\text{e-}03$	0.995
	3D	$5.229\text{e-}03 \pm 2.50\text{e-}03$	0.284

In 2D (Case A), $\tilde{\chi}$ depends strongly and significantly on porosity for Fibrous and Gaussian media ($p < 0.001$), while in 3D (Case C), $\tilde{\chi}$ shows no significant porosity dependence in any medium ($p > 0.05$). This is the expected behavior for these geometries as the 3D pore network stays well connected across the entire porosity range, so $\tilde{\chi}(C)$ sits at a near-constant value rather than tracking porosity.

A.3.3 Topology overlaid on the k - ϕ plot

We now ask whether 2D vs. 3D differences in connectivity and permeability are aligned.




Fig. A9 Permeability versus porosity plot from the main manuscript, with each marker colored by its pore normalized Euler characteristic $\tilde{\chi}$. Three panels split the data by medium: (a) Fibrous, (b) Gaussian, (c) Overlapping Spheres. The x axis is porosity ϕ ; the y axis is permeability k in m^2 on a log scale. Marker shape encodes data type and 3D porosity bracket: triangle for a 2D slice at $\phi(3D) = 0.4$, circle for a 2D slice at $\phi(3D) = 0.6$, square for a 2D slice at $\phi(3D) = 0.8$, and star for the 3D cube. Marker color is set by $\tilde{\chi}$ on a blue to red color bar (blue means negative $\tilde{\chi}$, where the pore network is well connected; red means positive $\tilde{\chi}$, where the pore network is less connected). Two Kozeny-Carman reference curves are added in each panel: the solid black curve is fitted to the 2D points only and the dashed black curve to the 3D points only.

Figure 2 from the main manuscript (Fig. A9, replotted here for convenience) shows across the three panels, marker color tends to follow the permeability axis: orange and red markers (positive $\tilde{\chi}$, poorly connected pore networks) are more common at lower permeabilities, whereas blue markers (negative $\tilde{\chi}$, well connected networks) are more common at higher permeabilities. This tendency is most visible in panels (a) Fibrous and (b) Gaussian. In panel (c) Overlapping Spheres, all $\tilde{\chi}$ values lie close to zero, so the markers appear nearly white and the

color separation is weak. The correspondence between $\tilde{\chi}$ and permeability is, however, not universal: at $\phi = 0.8$ in the Fibrous medium and at $\phi = 0.6$ in the Gaussian medium, the 2D and 3D permeabilities are similar although their $\tilde{\chi}$ values differ. $\tilde{\chi}$ therefore serves as an inexpensive connectivity proxy computed from geometry alone, but only as a partial predictor of the permeability differences between 2D and 3D.

We first asked the binary question whether given only the sign of $\tilde{\chi}$ for a single realization, could one correctly predict whether that realization's permeability sits above or below a chosen threshold on the k axis? A realization here is a medium combined with a porosity and an analysis type, where the analysis type is either a 2D slice (Case A) or the 3D cube (Case C). The threshold is a horizontal line on k that splits all realizations into a high k group (above) and a low k group (below). To check, we labeled every realization twice: once as well connected or less connected by the sign of $\tilde{\chi}$, and once as high k or low k by whether k sits above or below the threshold. We then counted how often the two labels matched. Figure A10 reports this measurement.




Fig. A10 $\tilde{\chi}$ on the horizontal axis, permeability k on the vertical axis (m^2 , log scale). The vertical line at $\tilde{\chi} = 0$ (well-connected on the left, fragmented on the right) and the horizontal line at the permeability threshold (high k above, low k below) split the plot into four quadrants. The two green quadrants are the ones where topology and transport agree (well-connected with high k or fragmented with low k); the two red quadrants are where they disagree. 31 of the 36 samples (86 %) land in a green quadrant. Marker color shows the medium (Fibrous = blue, Gaussian = green, Overlapping Spheres = red); circle = Case A 2D slice, star = Case C 3D.

Topology (sign of $\tilde{\chi}$) and transport (high vs low k) agree on 86% of the sets with no curve fitting required. The five disagreements are all near the threshold line, i.e. sets that are borderline by

both metrics. In other words, we can claim that connectivity (the sign of $\tilde{\chi}$) could be a reliable binary predictor of transport (high vs low k), the direct structural explanation for the 2D-vs-3D permeability gap shown in Fig. 2 of the main article.

A further question we asked is whether the 2D vs 3D disagreement on $\tilde{\chi}$ can be large enough that the 2D slice and the 3D cube classify the same medium oppositely, one as well connected and the other as fragmented. To answer this, for every (medium, porosity) configuration we computed the Symmetric Percentage Difference (SPD) of $\tilde{\chi}$ between the 2D slices and the 3D parent (defined in Section A.2.4), one bounded number in $[-2, +2]$.




Fig. A11 Symmetric Percentage Difference (SPD) of the pore-normalized Euler characteristic $\tilde{\chi}$ between the paper-matched 2D slices (Case A) and the 3D geometry (Case C), one row per (medium, porosity) configuration. Each colored point is the mean SPD; the whisker is $\pm$ SD. The x-axis is bounded in $[-2, +2]$: 0 = high similarity of 2D and 3D, ± 2 = opposite signs at equal magnitude.

Each row of Figure A11 carries three pieces of information about the 2D-vs-3D mismatch in a single glance. The sign of the mean tells the direction of the bias: a negative value means the 2D slice reports a more connected pore phase per voxel than the 3D geometry, a positive value means it reports a more poorly connected. The distance from zero tells the size of that bias on the bounded scale. The length of the whisker tells the slice-to-slice consistency: a short whisker means every slice of that configuration agrees on the 2D-vs-3D mismatch; a long whisker means the mismatch depends strongly on which slice is taken.

In Figures A9 to A11 we documented a systematic permeability discrepancy between 2D slices and their 3D parent cubes that persists at matched porosity, together with a parallel shift in $\tilde{\chi}$ between a slice and its parent. This raises a sharper question. Is $\tilde{\chi}$, a purely topological quantity we compute on a 2D slice, a predictor of how wrong the slice permeability will be relative to the true 3D permeability? Figure A12 tests this directly by pairing every 2D slice with its parent 3D cube and comparing the two shifts.




Fig. A12 Per medium view of the relative permeability shift $\Delta k/k_{3D}$ against the topology shift $\Delta \tilde{\chi}$, for every 2D slice we paired with its parent 3D cube. Three panels split the data by medium: (a) Fibrous, (b) Gaussian, (c) Overlapping Spheres. The x axis is $\Delta \tilde{\chi} = \tilde{\chi}_{2D} - \tilde{\chi}_{3D}$, where $\tilde{\chi}_{2D}$ and $\tilde{\chi}_{3D}$ are taken as raw signed values with their positive and negative signs preserved (not as absolute magnitudes); the y axis is $\Delta k/k_{3D} = (k_{2D} - k_{3D})/k_{3D}$. Marker shape represents 3D porosity (triangle $\phi_{3D} = 0.4$, circle $\phi_{3D} = 0.6$, square $\phi_{3D} = 0.8$). The solid black line in each panel is a linear regression we fit through that medium alone. Slope $\pm$ SE and p values for each panel are listed in Table A2

In the Fibrous medium, one observes a distinct trend, with negative $\tilde{\chi}_{2D} - \tilde{\chi}_{3D} < 0$ leading to an overprediction of permeability in the 2D simulation and vice versa (Fig. A12a). However, two slices at porosities of 0.8 sit far below zero on the negative- $\Delta \tilde{\chi}$ side, running against the trend; the fit line follows the majority pattern but is pulled down on the left by these two. Nine of the ten Gaussian slices sit below the zero line, so nearly every Gaussian 2D slice underestimate the permeability of its 3D parent (Fig. A12b). Within that band the negative- $\Delta \tilde{\chi}$ squares sit slightly closer to zero, the positive- $\Delta \tilde{\chi}$ circles slightly further below. The fit line slopes gently downward because that small gap follows the expected direction. Overlapping Spheres shows the widest scatter (Fig. A12c). Several near-zero- $\Delta \tilde{\chi}$ slices sit above the zero line, so those 2D slices overestimate; the two negative- $\Delta \tilde{\chi}$ circles sit just below zero; and one triangle with the most positive $\Delta \tilde{\chi}$ sits far below near the -1 floor. That single far-below point does most of the work setting the steep fit line.

Taken together the three panels indicate that in the majority of pairings the sign of $\Delta \tilde{\chi}$ points in the same direction as the sign of $\Delta k/k_{3D}$: if the 2D slice's $\tilde{\chi}$ is smaller than its 3D parent's, the 2D permeability tends to come out higher than the 3D value; if the 2D slice's $\tilde{\chi}$ is larger, the 2D permeability tends to come out lower. However, there are notable deviations from this general pattern. Fibrous has two instances that show the opposite behavior, Gaussian underestimates

permeability differences in every case (only the size of the underestimate follows $\Delta\tilde{\chi}$), and Overlapping Spheres has scattered points. Thus, $\tilde{\chi}$ can serve as an indicate of whether the 2D slice will overestimate or underestimate the 3D permeability, but appears not to predict the magnitude of the over- or underestimate.

Table A2. Per-medium linear regression of the relative permeability shift $\Delta k/k(3D)$ on the topology shift $\Delta\tilde{\chi}$, using 2D slices paired with their parent 3D cube.

MEDIUM	SLOPE $\pm$ SE	P-VALUE
Fibrous	$-1.026e+2 \pm 7.24e+1$	0.206
Gaussian	$-3.965e+1 \pm 2.85e+1$	0.202
Overlapping Spheres	$-1.201e+2 \pm 2.33e+2$	0.625

The above analysis does not account for differences between porosities of the 2D and 3D simulations, even though permeability is a known the vary with porosity. To separate the effect of porosity and topology, we first rescale the 2D permeability to the porosity of its 3D parent along the medium's 2D Kozeny-Carman curve, and only then form the 2D-vs-3D comparison. We compute the porosity-corrected 2D-permeability as follows:

$$\log_{10}(k_{2D}(\phi_{3D})) = \log_{10}(k_{2D}(\phi_{2D})) + \log_{10}(k_{2D}^{KC}(\phi_{3D})) - \log_{10}(k_{2D}^{KC}(\phi_{2D}))$$

where ϕ_{3D} is the porosity of the 3D parent (the target we rescale to); ϕ_{2D} is the porosity of the 2D slice (what we rescale from). $k_{2D}(\phi_{3D})$ is the 2D slice's permeability rescaled from porosity of the 3D domain. Superscript ^{KC} indicate the permeability estimated at the porosity of interest using the 2D Kozeny-Carman relationship (solid line in Fig. 2 or Fig A9)

Fig. A13 shows the result of that correction, so whatever value the y axis still shows for a slice reflects topology rather than porosity mismatch.




Fig. A13 Kozeny-Carman porosity-corrected 2D vs 3D permeability shift $y = (k_{2D}(\phi_{3D})) - k_{3D} / k_{3D}$ against the topology shift $\Delta\tilde{\chi}$, for every 2D slice we paired with its parent 3D cube. Panels: (a) Fibrous, (b) Gaussian, (c) Overlapping Spheres.

Accounting for the difference in porosities of 2D vs. 3D simulations and estimating its effect on permeabilities shows the following: For the Fibrous medium (Fig. A13a), the difference in permeabilities becomes smaller for some of the medium porosity simulations with $\Delta\tilde{\chi} \sim -0.002$. For the Gaussian medium (Fig. A13b), the patterns largely remains unchanged, with four simulations above the zero line and two below for negative- $\Delta\tilde{\chi}$ and all four simulations below the zero line for $\Delta\tilde{\chi} > 0$, matching the expected impact of the topology differences in permeabilities. The impact on the Overlapping Spheres medium is most pronounced (Fig. A13c), with relative permeability differences (in settings where the absolute differences are relatively small, see Figs. 2, A9) becoming smaller. Notably, the slopes are comparable across all media. But as substantial deviations from the inverse relationship between permeability and connectivity differences remain, our analysis still suggests that $\tilde{\chi}$ can tell us whether the 2D slice will overestimate or underestimate the 3D permeability, but not the magnitude of that over- or underestimate.

Table A3. Per-medium linear regression of the Kozeny-Carman porosity-corrected shift $y = (k_{2D}(P_{GOAL}) - k_{3D}) / k_{3D}$ on the topology shift $\Delta\tilde{\chi}$, using 2D slices paired with their parent 3D cube.

MEDIUM	SLOPE $\pm$ SE	P-VALUE
Fibrous	$-9.576e+1 \pm 8.28e+1$	0.292
Gaussian	$-6.892e+1 \pm 2.87e+1$	0.043
Overlapping Spheres	$-6.712e+1 \pm 7.72e+1$	0.418

A.3.4 Betti numbers and Euler characteristic by medium and porosity

Table A4. Betti numbers and their gaps (A – C) per medium and porosity

Medium	ϕ	$\beta_0(A)$	$\beta_0(C)$	gap β_0 (A–C)	$\beta_1(A)$	$\beta_1(C)$	gap β_1 (A–C)	$\beta_2(A)$	$\beta_2(C)$	gap β_2 (A–C)
Fibrous	0.4	28.5	22	6	10.0	888	-878	0	0	0
Fibrous	0.6	9.3	3	6	34.0	1,409	-	0	0	0
							1,375			
Fibrous	0.8	7.0	1	6	54.7	1,554	-	0	0	0
							1,499			
Gaussian	0.4	—	94	—	—	686	—	—	24	—
Gaussian	0.6	20.5	13	8	13.0	1,360	-	0	83	-83
							1,347			
Gaussian	0.8	3.7	3	1	36.7	1,014	-977	0	197	-197
Overlapping Spheres	0.4	6.3	20	-14	17.0	848	-831	0	0	0

Overlapping Spheres	0.6	2.7	1	2	18.0	133	-115	0	29	-29
Overlapping Spheres	0.8	1.0	1	0	8.0	22	-14	0	48	-48

The β_1 values in Table A4 show a large and systematic gap between the 2D and the 3D estimates: in every row, $\beta_1(C)$ is between thirty and one hundred times larger than $\beta_1(A)$. Most loops in the 3D network are not within a 2D plane. This is why $\beta_1(C)$ is much larger than $\beta_1(A)$ in every row of the table. $\beta_2(A)$ is structurally zero (a 2D image cannot enclose a pocket).

Table A5. Euler characteristic and pore-normalized $\tilde{\chi}$ per medium and porosity

Medium	ϕ	$\chi(A)$ mean $\pm$ SD	$\chi(C)$	$\tilde{\chi}(A)$ mean $\pm$ SD	$\tilde{\chi}(C)$
Fibrous	0.4	18.5 $\pm$ 6.4	-866	4.4264e-03 $\pm$ 1.51e-03	-2.1810e-03
Fibrous	0.6	-24.7 $\pm$ 15.3	-1,406	-3.6815e-03 $\pm$ 2.09e-03	-2.5047e-03
Fibrous	0.8	-47.7 $\pm$ 10.0	-1,553	-6.4066e-03 $\pm$ 1.03e-03	-2.1491e-03
Gaussian	0.4	—	—	—	—
Gaussian	0.6	7.5 $\pm$ 5.3	-1,264	1.3168e-03 $\pm$ 9.54e-04	-2.1756e-03
Gaussian	0.8	-33.0 $\pm$ 6.0	-814	-4.3663e-03 $\pm$ 6.53e-04	-1.0588e-03
Overlapping Spheres	0.4	-10.7 $\pm$ 13.5	-828	-1.7370e-03 $\pm$ 2.36e-03	-2.0598e-03
Overlapping Spheres	0.6	-15.3 $\pm$ 3.2	-103	-2.3200e-03 $\pm$ 6.47e-04	-1.4771e-04
Overlapping Spheres	0.8	-7.0 $\pm$ 2.0	27	-8.4103e-04 $\pm$ 2.49e-04	3.1999e-05

$\tilde{\chi}(A)$ and $\tilde{\chi}(C)$ disagree in sign for Fibrous $\phi=0.4$ and Gaussian $\phi=0.6$, for which 2D is more fragmented than 3D, and for Overlapping Spheres $\phi=0.8$ where the 2D is more connected. In realizations where the sign agrees, the magnitudes differ: $\tilde{\chi}(A)$ is typically two to three times larger in absolute value than $\tilde{\chi}(C)$ (e.g. Fibrous $\phi=0.8$, -6.4×10^{-3} vs -2.1×10^{-3}), so the 2D slice over-states how connected or how fragmented the medium is. Overall, the 2D slice mis-classifies the sign in 3 of 8 rows and over-states the magnitude in the remaining 5.

A.3.5 Case B: slice-aware 3D analysis

Case B re-uses the same paper-matched 2D slices but counts connectivity from the surrounding 3D parent. Specifically, $\beta_0(B)$ is the number of 3D pore regions that have at least one voxel on the slice, $\beta_1(B)$ is the number of 3D loops whose support along the slicing axis covers the slice index, and $\beta_2(B)$ is the number of sealed 3D cavities that have at least one voxel on the slice. The

normalization N_p is the pore voxel count of the slice itself, so $\tilde{\chi}(B)$ lives on the same per-voxel scale as $\tilde{\chi}(A)$. This bridges Case A and Case C: it carries the 3D topology information that a flat slice cannot see, while still being anchored to the slice the 2D simulation runs on.

Table A6 states the Case B Betti counts per (medium, ϕ), with mean and sample standard deviation across the slices of each configuration. Positive gaps mean Case B recovers features that Case A does not see; β_2 is positive in every row because a 2D slice cannot enclose a pocket by construction.

Table A6. Case B Betti counts and gaps ($B - A$)

MEDIUM	ϕ	$\beta_0(B)$ mean $\pm$ SD	$\beta_1(B)$ mean $\pm$ SD	$\beta_2(B)$ mean $\pm$ SD	gap β_0 (B-A)	gap β_1 (B-A)	gap β_2 (B-A=B)
Fibrous	0.4	1.3 $\pm$ 0.6	888.0 $\pm$ 0.0	0.0 $\pm$ 0.0	-27.2	+878.0	+0.0
Fibrous	0.6	1.0 $\pm$ 0.0	1409.0 $\pm$ 0.0	0.0 $\pm$ 0.0	-8.3	+1375.0	+0.0
Fibrous	0.8	1.0 $\pm$ 0.0	1554.0 $\pm$ 0.0	0.0 $\pm$ 0.0	-6.0	+1499.3	+0.0
Gaussian	0.4	—	—	—	—	—	—
Gaussian	0.6	1.2 $\pm$ 0.5	1360.0 $\pm$ 0.0	1.0 $\pm$ 1.4	-19.2	+1347.0	+1.0
Gaussian	0.8	2.0 $\pm$ 1.5	1014.0 $\pm$ 0.0	0.8 $\pm$ 0.4	-1.7	+977.3	+0.8
Overlapping Spheres	0.4	1.3 $\pm$ 0.6	848.0 $\pm$ 0.0	0.0 $\pm$ 0.0	-5.0	+831.0	+0.0
Overlapping Spheres	0.6	1.3 $\pm$ 0.6	133.0 $\pm$ 0.0	2.0 $\pm$ 2.6	-1.3	+115.0	+2.0
Overlapping Spheres	0.8	1.0 $\pm$ 0.0	22.0 $\pm$ 0.0	2.7 $\pm$ 3.1	+0.0	+14.0	+2.7

The gap between Case B and Case A is largest on β_1 (loops) for every medium, which confirms that most of the connectivity a 2D slice misses comes from loop closures that span the third dimension. The β_2 gap is always positive but smaller in magnitude, because pockets are rare per slice plane. The gap between Case B and Case C (full 3D cube), not tabulated here, would tell how much of the 3D topology is captured by the slice neighborhood at all; that comparison is left for future work.




Fig. A14 Connectivity-based decision pipeline for choosing between 2D and 3D simulations. The workflow uses only the binary geometry (no transport simulation). Step 1 takes 2D slice(s) and the parent 3D cube. Step 2 runs persistent homology on each and computes the pore-normalized Euler characteristic $\tilde{\chi}$. Step 3 forms the Symmetric Percentage Difference (SPD) of $\tilde{\chi}$ between the two. The decision diamond branches on whether the signs of $\tilde{\chi}(A)$ and $\tilde{\chi}(C)$ agree and whether $|SPD|$ stays small: if yes, the 2D model is a fair stand-in for 3D; if no, the 2D model mis-estimates transport and the 3D simulation should be run directly. The decision is also weighted by the intended use.

A4. Summary

The 3D analysis (Case C) consistently reports more connectivity than the 2D analysis (Case A): more loops, pockets (when they exist), and a more negative $\tilde{\chi}$ in the majority of the realizations. Figures A8 and A9 show that the 2D slice-vs-3D gap is present at every porosity, and that the topology agrees with the k - ϕ trend visible in Fig. 2 of the main article. Figure A10 shows that 31 of 36 samples (86 %) place a well-connected $\tilde{\chi}$ at high k or a fragmented $\tilde{\chi}$ at low k . Treating any single 2D slice as a stand-in for the 3D network will therefore underestimate transport-relevant connectivity, with a bias whose magnitude depends on the medium and porosity (Fig. A8, Table A1). Fig. A13 goes one step further by rescaling the 2D permeability to the 3D parent's porosity along the medium's 2D-only Kozeny-Carman curve, in order to isolate the effect of topology

differences on the relative differences in permeabilities. However, this correction explains only some of the observed differences between 2D and 3D simulations.

Nevertheless, the comparison of the Euler characteristics of 2D and 3D media can in part inform whether a 2D model can stand-in for a 3D one. The decision is relatively cheap to make, because $\tilde{\chi}$ is computed from the binary geometry alone and no transport simulation is needed. One can compute $\tilde{\chi}$ for a candidate 2D slice and its parent 3D. If they have opposite signs, or if Symmetric Percentage Difference (SPD) sits far from zero, the 2D model is likely to misestimate transport in 3D. The threshold for that decision also depends on what the results will be used for: scoping or comparative studies may tolerate a moderate 2D-vs-3D bias, whereas applications driving safety-, hazard-, or precision-critical decisions may need more caution (see Fig. A13).

References

- Edelsbrunner, H., & Harer, J. (2010). *Computational topology: an introduction*. American Mathematical Soc..
- Forecasting, L. R. (1978). From Crystal Ball to Computer. *Scott Armstrong Robert J. Genetski*.
- Herring, A.L., Robins, V., Sheppard, A.P.: Topological persistence for relating microstructure and capillary fluid trapping in sandstones. *Water Resour. Res.* 55, 555–573 (2019). <https://doi.org/10.1029/2018WR022780>
- Makridakis, S.: Accuracy measures: theoretical and practical concerns. *International journal of forecasting*. 9, 527–529 (1993)
- Maria, C., Boissonnat, J.-D., Glisse, M., Yvinec, M.: The Gudhi library: simplicial complexes and persistent homology. In: *International Congress on Mathematical Software*, pp. 167–174. Springer (2014)
- Robins, V., Saadatfar, M., Delgado-Friedrichs, O., Sheppard, A.P.: Percolating length scales from topological persistence analysis of micro- CT images of porous materials. *Water Resour. Res.* 52, 315–329 (2016). <https://doi.org/10.1002/2015WR017937>
- Wagner, H., Chen, C., Vućini, E.: Efficient computation of persistent homology for cubical data. In: *Topological Methods in Data Analysis and Visualization II*, pp. 91–106. Springer, Heidelberg (2011)