
Supplementary information

Surface optimization governs the local design of physical networks

In the format provided by the
authors and unedited

Surface Optimisation Governs the Local Design of Physical Networks

Supplementary Information

September 19, 2025

Xiangyi Meng, Benjamin Piazza, Csaba Both, Baruch Barzel, Albert-László Barabási

Contents

1	Physical network datasets	3
1.1	Constructing the graph	4
1.2	Alternative skeletonisation methods	8
2	One-dimensional cost minimisation—the Steiner graph	10
3	Two-dimensional links—Charts	12
3.1	Example: the catenoid sleeve	16
4	The unified coordinate system—Atlas	18
4.1	Quadratic differentials	18
5	Surface minimisation	22

6	Numerical implementation	25
6.1	Quad-mesh tiling	25
6.2	Optimisation target function	26
6.3	Computational complexity	30
7	Empirical analysis of branching	31
7.1	Degree distribution	31
7.2	Trifurcation planarity	35
7.3	Angle asymmetry	36
8	Empirical analysis of link lengths and loops	40
9	Empirical analysis of node morphology	41
9.1	Smoothness	41
9.2	Surface optimisation	41
10	Surface versus volume minimisation	45

1 Physical network datasets

Our work relies on the empirical analysis of real-world tree-like physical networks from various domains. The networks, described below, cover 10 orders of magnitude in volume, from the nanometric scale of the neuronal components to the $\sim 10^3$ cm length of tropical tree branches:

1. **Human neurons.** An electron microscopy scan of a cubic millimetre human brain element. The complete dataset includes $\sim 10^4$ neurons linked through $\sim 10^6$ axons and dendrites [1]. Our analysis focuses on a randomly sampled segment of these data, describing the precise morphology of 104 manually traced and proofread cells.
2. **Fruit fly neurons.** Three random full-size neurons obtained from the Hemibrain connectome project, as part of the FlyEM dataset [2]. Each neuron represents an elaborate three-dimensional structure with $\sim 10^3$ dendritic/axonal pathways packed within a volume of $\sim 10^{-15}$ m³.
3. **Blood vessels.** The human pulmonary arterial network [3], comprising $\sim 10^3$ blood vessels, ranging in diameter from 0.1 to 3 cm.
4. **Tropical trees.** A structural description of different tropical tree species, from the main branches (boughs) to the peripheral twigs. The data cover a total of 29 trees, 10 from the lowland tropical moist forests in Guyana, an additional 10 from peat swamp forests in Indonesia, and finally, 9 trees from *terra firme* forests in Peru [4].
5. **Corals.** Data on the physical volumetric structure of 28 corals, representing various species, such as *Corallium sp.*, *Cyathophylia viridis* and others [5]. Here, the physical links capture the tubular inner structure of each coral.
6. **Arabidopsis.** Three-dimensional scans of 5 arabidopsis plants, collected at 10 consecutive points in time [6]. Each plant comprises $\sim 10^1$ to 10^2 links (branches).

Each of these datasets describes the three-dimensional coordinates of the network’s geometric shape, allowing a complete reconstruction of the physical network, from its underlying skeletal graph to the specific morphology of all links and nodes. The challenge is that the different networks are provided using distinct representations, which are not easily interchangeable (Fig. 1):

1. **Volumetric images** (human neuron, fruit fly neuron). These data structures are generated by rasterising the three-dimensional object into an array of voxels, each representing a small volumetric element.

2. **Point clouds** (tropical trees, arabidopsis). In this data format one samples data points from the surface of the object, capturing three-dimensional coordinates in space. The sampling can be regularised or random. If sufficiently dense, the sampled points allow a comprehensive reconstruction of the object’s geometrical structure.
3. **Triangular meshes** (corals). These data represent the exterior surface of the three-dimensional object by a series of interlaced triangles, providing a discretised description of its geometrical shape.
4. **Tetrahedral meshes** (blood vessels). Here the data triangulate the entire three-dimensional volume of the object, generating a mesh of interconnected tetrahedra that represent the object’s internal structure in addition to its surface.

Note that these three-dimensional representations are not easily interchangeable. For example, it is easy to convert a triangular mesh into a point cloud, by simply considering all the triangle vertices as *points* in the point cloud. The reverse, however, is more difficult, as one needs to decide which point-trios should constitute triangles in the desired mesh. In general, we designed the list above such that upward conversion is more straightforward than downward. Hence, volumetric images are the easiest to construct from all other representations, while triangular and tetrahedral meshes are the most difficult to construct. To standardise our datasets and allow a systematic cross-domain analysis we converted all data into volumetric images. A complete description of our physical network datasets appears in <https://physical.network>.

1.1 Constructing the graph

Once we obtain the three-dimensional physical network structure, the challenge is to construct its underlying abstract graph, namely break down the observed geometric object into its skeleton of nodes and links. Indeed, the system’s skeleton is not explicit in the data and must be implicitly extracted from the object’s geometry. Quite generally, nodes appear at every point of termination or bifurcation of a link. This leaves room for interpretation, as illustrated in Fig. 2. For example, two adjacent bifurcations may be identified as a single trifurcation, leading to different resulting graph structures.

This challenge of *skeletonisation* is not unique to our current analysis and has, indeed, been extensively studied in biological systems, neuronal networks, and physiology [7]. Therefore, there already exist quite powerful algorithms for this purpose. These algorithms take the volumetric description of the three-dimensional object and generate the most likely spatially embedded graph structure around which the object is layered.




Figure 1: **Three-dimensional data representation.** (a) A triangular mesh representation of arabidopsis. The smooth surface is captured by small intersecting triangles. The volume can also be mapped in a similar fashion using tetrahedra instead of triangles. (b) The corresponding point cloud representation. (c) The volumetric image representation. It is straightforward to convert the data representation from left to right (tetrahedral/triangular meshes $\rightarrow$ point clouds $\rightarrow$ volumetric images) but more difficult in the opposite direction. (d) The skeletonised arabidopsis network as obtained from the volumetric image in (c).

In our analysis, we primarily used the *Kimimaro* [8] skeletonisation algorithm. This algorithm was originally designed to extract skeletons directly from volumetric images of high resolution neurons, obtained from electron microscopy data. However, it can be directly applied to other systems, such as corals or blood vessels as well. In datasets 1–6 above, the tropical trees and the fruit fly and human neurons already included their skeletonised graphs. Therefore, we only applied the Kimimaro algorithm to the coral, arabidopsis (Fig. 1d), and blood vessel datasets Table 1. First, we converted all the discrete geometric representations for the three datasets into volumetric images. Then, we ran the Kimimaro algorithm to build the corresponding skeletons.

The Kimimaro algorithm. The algorithm is designed to seek the most likely graph structure that traverses through the three-dimensional tubular structure of the physical network (Fig. 3). In brief, the algorithm has five key steps:




Figure 2: **Skeletonisation ambiguity.** The same physical network (blue) may be skeletonised in two different ways. (a) As two separate bifurcations (black). (b) As a single trifurcation. The Kimimaro algorithm seeks the most likely skeletal graph given the three-dimensional physical graph geometry.

Dataset	Skeletonisation method
Human neuron [1]	VAST (Volume Annotation and Segmentation Tool) [9]
Fruit fly neuron [2]	TEASAR (self-implemented) [10]
Blood vessel [3]	Kimimaro
Tropical tree [4]	QSM (Quantitative Structure Models) [11]
Coral [5]	Kimimaro
Arabidopsis [6]	Kimimaro

Table 1: **Skeletonisation methods.** Our six datasets and the skeletonisation methods used for their analysis.

- (i) **Foreground vs. background.** We classify regions of the three-dimensional volume as either foreground, where material is present and voxels are occupied (Fig. 3a, blue), or background, representing empty space.
- (ii) **Distance transform.** Next, we compute a distance transform, which assigns to each foreground voxel a value corresponding to its shortest distance from the background. We then identify the local maxima of this distance field—known as medial points or centres—which together trace the skeletal backbone of the physical links.
- (iii) **Source selection.** We select one of these medial points as the *source* node, from which to construct paths to the *target* nodes. The choice of the source node depends on various criteria. A common approach is to select the medial point with the globally maximal distance transform value, corresponding to the most deeply embedded point within the object’s volume. In other contexts, the source may be determined by predefined anatomical landmarks, such as the soma in neuronal networks.
- (iv) **Path construction.** We trace a path starting from the source node, advancing through the sleeve centres, and terminating at the target (Fig. 3b). Once the target is reached, we invalidate the coordinates of the sleeves traversed by the path, ensuring that subsequent paths do not pass through these *already-used* regions. (Fig. 3c).
- (v) **Recurrence.** We continue to generate paths to the next target, according to steps (i)–(iv), until all sleeves are invalidated. The resulting collection of one-dimensional paths represents the graph skeleton.

To guide the paths through the sleeve centres in steps (ii) and (iv), the distance transform is defined by a penalty field $P(\mathbf{r})$ that is maximal at the sleeve boundary $\mathbf{X}(\mathbf{r})$ and minimal at its centre, *i.e.* at points almost equidistant from all boundaries. Denoting the Euclidean distance of each point $\mathbf{r}$ from the sleeve boundary by $D(\mathbf{r})$, we define $P(\mathbf{r})$ as

$$P(\mathbf{r}) = K \left(1 - \frac{D(\mathbf{r})}{\max\{D(\mathbf{r})\}} \right)^\alpha, \quad (1)$$

which is, indeed, minimal at the sleeve centre ($D(\mathbf{r}) = \max\{D(\mathbf{r})\}$) and maximal at its boundaries ($D(\mathbf{r}) = 0$). The parameters K and α determine how strictly we bind the paths to the sleeve centre.

The invalidation of used sleeves is done by constructing a virtual sphere of radius $\rho(\mathbf{r})$ around all points $\mathbf{r}$ along the path, and eliminating all voxels of the physical network that fall within the volume of that sphere. The radius of the invalidation sphere is set to

$$\rho(\mathbf{r}) = s \times D(\mathbf{r}) + C, \quad (2)$$

where the *scale* s , on the order of unity, ensures that the sphere covers the distance from the path-point $\mathbf{r}$ to the sleeve boundaries around it. We also include a

constant C as an offset to $\rho(\mathbf{r})$. Hence, the Kimimaro algorithm is governed by four parameters K , α , s , and C , as detailed in Table 2. Further details can be found at <https://github.com/seung-lab/kimimaro>.




Figure 3: **The Kimimaro algorithm.** (a) The physical network is composed of sleeves. We seek paths to connect the source to the target that traverse through the sleeve centres. (b) The penalty field $P(\mathbf{r})$ is minimal at the centre of the sleeve, thus pushing the paths far from the sleeve boundaries. Once the target is reached, we invalidate the surrounding sleeve by sliding an invalidation sphere of radius $\rho(\mathbf{r})$. (c) We then proceed to the next pair of source and target.

Dataset \ Parameter	s	C	α	K
Blood vessel [3]	1.5	1	4	10^4
Coral [5]	1.1	10	4	10^4
Arabidopsis [6]	1.0	20	4	10^4

Table 2: **Kimimaro algorithm parameters.** We applied the algorithm to skeletonise the three listed datasets.

As stated above, three datasets (human neurons, fruit fly neurons, and tropical trees) relied on pre-existing skeletons that were generated using different methods. These methods are listed in Table 1.

Skeletonisation outcomes. The skeletonisation process results in a spatially embedded graph of nodes and links. The links represent straight segments and the nodes describe points where links bend (elbows) or split into several directions (intersections). Hence, skeletal nodes with degree two capture points of link bending, but do not correspond to actual physical nodes. Nodes of degree one represent terminals in the physical network, and nodes with degree three or more represent branching points such as bifurcations, trifurcations, etc. In Table 3, we summarise the resulting skeletal structures.

Dataset	Physical links	Physical link type	Skeletal structure	
			Segments	Terminal nodes
Human neuron [1]	925,703	dendritic/axonal pathways	3,842,587	399,941
Fruit fly neuron [2]	11,154	dendritic/axonal pathways	73,423	5,625
Blood vessel [3]	3,870	vessels	102,143	1,983
Tropical tree [4]	5,330	boughs/twigs	11,380	3,856
Coral [5]	4,914	skeletons	130,939	2,487
Arabidopsis [6]	3,260	branches	72,795	3,737

Table 3: **Summary of physical network datasets.** We used six networks from different domains. For each network we show the number of physical links, their nature (axons, vessels, etc.), and the outcome of their skeletonisation.

Dimensionless measures. The fidelity of skeletonisation methods is often constrained by the spatial resolution of the underlying data. To circumvent this limitation, our key measurements, namely the weight parameter $\chi = w/r$ (degree distribution) and the circumference ratio $\rho = w'/w$ (angle asymmetry) are defined as dimensionless quantities. By design, these ratios ensure our results are robust against variations in spatial resolution and the specific choice of skeletonisation algorithm.

1.2 Alternative skeletonisation methods




Figure 4: **Comparison between the original and an alternative (Skeletor) skeletonisation method.** **(a)** The same neuron skeletonised by two methods. The average thickness w is similar between the two, while the total length l derived by the alternative method is larger as the soma of the neuron was also skeletonised. **(b)** Similar results obtained for a different neuron. We observe a similar trend: the two skeletons appear almost identical, with w and l remaining closely the same.




Figure 5: **Results under alternative (left) and original (right) skeletonisation methods.** (a,b) We recovered the analysis of Fig. 3 of the main paper using two different skeletonisation methods, which provide fully consistent results. (c,d) Similar consistency is also observed in our reconstruction of main text Fig. 5.

Our six datasets were skeletonised by diverse methods (Table 1). This diversity is encouraging, as it suggests that our results are robust to the specific skeletonisation protocol employed. To further examine this robustness, we implemented an alternative skeletonisation pipeline, focusing on the human neuron dataset. We chose this dataset because it contains the largest number of skeletons, all of which have been manually proofread by experts [1], providing a reliable ground truth against which to validate our alternative approach.

We began by using the *Trimesh* library [12] to extract two-dimensional meshes directly from the human neuron dataset. Next, we implemented the *Skeletor* algorithm [13] to construct the skeletonised graphs. As Fig. 4 shows, the two methods yield visually indistinguishable skeletons. Adjacent to each skeleton, we show two key extracted parameters: the average thickness (or circumference) w and the total neuronal length l , as obtained from the two algorithms. These specific measures, we find, are not identical, alluding to the differences between the two pipelines. Yet, they are strongly correlated. Most crucially, the main statistical findings reported in the paper remain un-affected by the choice of skeletonisation method. This holds for both the trifurcation behaviour, expressed via $P(\lambda)$ (Fig. 5a,b), and the sprouting/branching transition, captured by $P(\Omega)$ (Fig. 5c,d).

2 One-dimensional cost minimisation—the Steiner graph

Consider the challenge of connecting M spatially distributed terminals. This entails the construction of a fully connected, spatially embedded graph in which all terminals can be linked via finite paths, *i.e.* no isolated components. In the Steiner graph challenge, we seek to achieve such connectivity at minimal cost by optimising for the smallest total link length. Where relevant, we allow the addition of intermediate nodes between the fixed terminals, a step that can potentially help connect distant terminals via shorter links. For example, in Fig. 6 we present a four-terminal system arranged at a tetrahedral configuration. The optimal solution here is to link the four terminals through two intermediate nodes of degree $k = 3$. Therefore, the final Steiner graph $\mathcal{G}$ has $M + N$ nodes and L links, where N is the number of added intermediates. In our example in Fig. 6 we have $M = 4$ and $N = 2$.

In this setting, all nodes (terminals or intermediates) are characterised by their three-dimensional coordinates in space, and all links $i = 1, \dots, L$ exhibit spatially embedded pathways $\mathbf{X}_i(\eta_i)$ that originate in one node and end in another. Here, the function $\mathbf{X}_i(\eta_i)$ describes this spatial embedding as a continuous path in space, parametrised by $\eta_i \in [0, 1]$. Hence, $\mathbf{X}_i(\eta_i = 0)$ represents the coordinates of the node at one end of the link, and $\mathbf{X}_i(\eta_i = 1)$ represents the coordinates of the node at its other end.

In this parametrisation the graph $\mathcal{G}$ is given in terms of all the link paths $\mathbf{X}_i(\eta_i)$. The node locations are implicitly provided as the end points of these links, *i.e.* $\mathbf{X}_i(0)$ and $\mathbf{X}_i(1)$. The graph consistency requires that these end points are sewn to each other, such that adjacent links meet at their shared node. Specifically, if node α is at the intersection of two links i and j , one of these four conditions must be satisfied:

$$\begin{aligned} \mathbf{X}_i(0) &= \mathbf{X}_j(0), \\ \mathbf{X}_i(0) &= \mathbf{X}_j(1), \\ \mathbf{X}_i(1) &= \mathbf{X}_j(0), \\ \mathbf{X}_i(1) &= \mathbf{X}_j(1), \end{aligned} \tag{3}$$

ensuring that link i ends (begins) where link j begins (ends).




Figure 6: **The Steiner graph.** (a) A tetrahedral configuration of $M = 4$ terminals. (b) The Steiner graph $\mathcal{G}$ helps link all terminals at minimal cost (link length). Here this is achieved by adding $N = 2$ intermediate nodes (black) of degree $k = 3$ each.

Graph cost. The length l_i of each link can be calculated differentially by dividing the link into infinitesimal segments $d\mathbf{X}_i = (d\mathbf{X}_i/d\eta_i) d\eta_i$. Each segment has a length of $dl_i = \sqrt{d\mathbf{X}_i \cdot d\mathbf{X}_i}$, and hence the total link length is provided by

$$l_i = \int_0^1 d\eta_i \sqrt{\frac{d\mathbf{X}_i}{d\eta_i} \cdot \frac{d\mathbf{X}_i}{d\eta_i}}, \quad (4)$$

an integral over η_i . Collecting all link lengths we obtain the cost of $\mathcal{G}$ [14]:

$$S_{\mathcal{G}} = \sum_{i=1}^L \int_0^1 d\eta_i \sqrt{\frac{d\mathbf{X}_i}{d\eta_i} \cdot \frac{d\mathbf{X}_i}{d\eta_i}}, \quad (5)$$

a summation over the lengths of all links $i = 1, \dots, L$.

Cost minimisation requires us to find the graph $\mathcal{G}$ for which $S_{\mathcal{G}}$ is the smallest; namely, we seek

$$\arg \min_{\mathcal{G}} \{S_{\mathcal{G}}\}. \quad (6)$$

Our degrees of freedom in seeking Eq. (6) include the graph topology (or adjacency matrix), as well as its geometry. The former represents our control over the addition of intermediate nodes and the mapping of which nodes/terminals are linked. The latter captures the spatial locations of these intermediate nodes and the specific path $\mathbf{X}_i(\eta_i)$ followed by each link. The resulting length-optimised $\mathcal{G}$ is known as the *Steiner graph* [15], a well-established graph theoretic optimisation problem.

The Steiner graph characteristics. As detailed in the main paper, graphs satisfying Eq. (6) have three guaranteed local characteristics: (1) *Bifurcations*: All intermediate nodes have degree $k = 3$, *i.e.* bifurcations; (2) *Planarity*: All links emerging from a node are on the same plane; (3) *Angle symmetry*: The bifurcations are at equal angles of precisely 120° . As we show in our analysis, real networks may potentially violate each of these conditions. Consequently, the Steiner graph, a collection of one-dimensional ($d = 1$) links embedded in three-dimensional space, is insufficient to capture the physicality of real-world networked objects. Its one-dimensional links ignore the rich morphology of real physical networks, and its sharp branching points represent essential singularities that violate the manifold criteria we outline in the main paper. We, therefore, next consider higher-dimensional links with $d \geq 2$, which we analyse using manifolds rather than simple graphs.

3 Two-dimensional links—Charts

Next, we focus on two-dimensional links, or *charts*. Here, instead of linear pathways $\mathbf{X}_i(\eta_i)$, we consider each link to be a surface that surrounds a specific volume in three-dimensional space (see Fig. 8 for an example of a catenoid). Each link i is now defined by two coordinates, $\boldsymbol{\sigma}_i = (\sigma_i^0, \sigma_i^1)$, replacing the single parameter η_i of the one-dimensional description. A common parametrisation is to consider σ_i^0 as the longitudinal variable, confined to the range $\sigma_i^0 \in [0, l_i]$, and σ_i^1 as the azimuthal variable, for which we set $\sigma_i^1 \in [0, w_i]$. To ensure the continuity and smoothness of the link surface, we impose a periodic boundary condition $\mathbf{X}_i(\sigma_i^0, 0) = \mathbf{X}_i(\sigma_i^0, w_i)$. This description depicts each link as a geometric *sleeve* dressed around a path in space. The sleeve trajectory is tracked via the longitudinal σ_i^0 , and its local morphology around each point $0 \leq \sigma_i^0 \leq l_i$ is provided by the rotational σ_i^1 .

In this higher dimensional formulation, the intersection between a pair of links is no longer a singular point. Rather the sleeves must be sewn together, to create a continuous and smooth surface at the inter-link boundary (Fig. 7a). Consequently, Eq. (3), designed to capture the meeting *point* between links i and j , now takes a more complex form. To understand this, consider link i that intersects with j at the point where i ends, *i.e.* $\sigma_i^0 = l_i$, and j begins, namely $\sigma_j^0 = 0$. This intersection is not a point but rather defines a path along the edge of sleeve i , $\mathbf{X}_i(l_i, \sigma_i^1)$, and that of sleeve j , $\mathbf{X}_j(0, \sigma_j^1)$ (Fig. 9, blue paths). The path is defined by the range of σ_i^1 and σ_j^1 in which the two sleeves are sewn together. For example, in the bifurcation shown in Fig. 9e, each pair of sleeves is conjoined only through half of the complete azimuthal range, namely $\sigma_i^1, \sigma_j^1 \in [0, w/2]$. Within this intersection range we require

$$\begin{aligned} \mathbf{X}_i(l_i, \sigma_i^1) &= \mathbf{X}_j(0, \sigma_j^1), \\ \left. \frac{\partial \mathbf{X}_i}{\partial \sigma_i^0} \right|_{\sigma_i^0=l_i} &= \left. \frac{\partial \mathbf{X}_j}{\partial \sigma_j^0} \right|_{\sigma_j^0=0}, \end{aligned} \quad (7)$$

the first condition ensuring continuity along the intersection path and the second condition guaranteeing smoothness as we cross that intersection, from sleeve i to j .

Network manifold. To construct our physical network, we begin with the graph $\mathcal{G}$, a collection of one-dimensional links as described in Sec. 2. We then surround each link with a chart or sleeve $\mathbf{X}_i(\boldsymbol{\sigma}_i)$ that dresses the link with a three-dimensional morphology (Fig. 7b). By sewing all links together following condition (7), we obtain the complete network manifold $\mathcal{M}(\mathcal{G})$, a continuous and smooth geometric object structured around the discrete graph $\mathcal{G}$.

While Eq. (7) generalises our original one-dimensional condition of Eq. (3), it is interesting to note that the conditions appearing in Eq. (3) are *not* a particular case of (7). This is due to the essential singular nature of the one-dimensional intersections that, by definition, are un-smooth and hence




Figure 7: **Two-dimensional manifold.** (a) The physicality of the network imposes fundamental constraints at the boundaries between the sleeves, which must be continuous, smooth, and avoid crossings. The natural structure that achieves this is a manifold (top left). (b) A two-dimensional network manifold $\mathcal{M}(\mathcal{G})$ consists of two layers: first, a graph structure $\mathcal{G}$, comprising nodes and links; on top of that, a set of two-dimensional sleeves ‘dressed’ around the nodes/links. Here the graph structure is represented by a single trifurcation motif, around which we dress four different sleeve morphologies (grey). Despite having the same underlying graph, each of the resulting four structures represents a distinct physical network, due to its unique sleeve morphology.

violate Eq. (7)’s second condition. Therefore, the one-dimensional construction of Sec. 2 does not constitute a proper manifold, and only after surrounding it by the chart surfaces does it become one.

Link surface area. To calculate the surface area of each link we consider infinitesimal area elements $d\mathbf{s}$ along the longitudinal and azimuthal axes. The sides of each area element are given by $|\partial\mathbf{X}_i/\partial\sigma_i^0| d\sigma_i^0$ and $|\partial\mathbf{X}_i/\partial\sigma_i^1| d\sigma_i^1$, and hence the area they cover is

$$d^2\sigma_i \sqrt{\left(\frac{\partial\mathbf{X}_i}{\partial\sigma_i^0} \cdot \frac{\partial\mathbf{X}_i}{\partial\sigma_i^0}\right) \left(\frac{\partial\mathbf{X}_i}{\partial\sigma_i^1} \cdot \frac{\partial\mathbf{X}_i}{\partial\sigma_i^1}\right) - \left(\frac{\partial\mathbf{X}_i}{\partial\sigma_i^0} \cdot \frac{\partial\mathbf{X}_i}{\partial\sigma_i^1}\right)^2}, \quad (8)$$

an infinitesimal parallelogram surrounding the coordinate σ_i on the sleeve surface $\mathbf{X}_i(\sigma_i)$. Covering the boundaries of the entire sleeve we obtain link i ’s surface area as

$$\int_0^{l_i} \int_0^{w_i} d^2\sigma_i \sqrt{\left(\frac{\partial\mathbf{X}_i}{\partial\sigma_i^0} \cdot \frac{\partial\mathbf{X}_i}{\partial\sigma_i^0}\right) \left(\frac{\partial\mathbf{X}_i}{\partial\sigma_i^1} \cdot \frac{\partial\mathbf{X}_i}{\partial\sigma_i^1}\right) - \left(\frac{\partial\mathbf{X}_i}{\partial\sigma_i^0} \cdot \frac{\partial\mathbf{X}_i}{\partial\sigma_i^1}\right)^2}, \quad (9)$$

a natural generalisation of the one-dimensional Eq. (4).

Manifold cost. The cost of $\mathcal{M}(\mathcal{G})$ is obtained by summing over all links to extract the *manifold* surface area as [16]

$$S_{\mathcal{M}(\mathcal{G})} = \sum_{i=1}^L \int_0^{l_i} \int_0^{w_i} d^2 \boldsymbol{\sigma}_i \sqrt{\left(\frac{\partial \mathbf{X}_i}{\partial \sigma_i^0} \cdot \frac{\partial \mathbf{X}_i}{\partial \sigma_i^0} \right) \left(\frac{\partial \mathbf{X}_i}{\partial \sigma_i^1} \cdot \frac{\partial \mathbf{X}_i}{\partial \sigma_i^1} \right) - \left(\frac{\partial \mathbf{X}_i}{\partial \sigma_i^0} \cdot \frac{\partial \mathbf{X}_i}{\partial \sigma_i^1} \right)^2}. \quad (10)$$

We can simplify this description by introducing the *Riemannian* metric tensor [16]

$$\gamma_{i,\alpha\beta} = \gamma_{i,\alpha\beta}(\boldsymbol{\sigma}_i) \equiv \frac{\partial \mathbf{X}_i}{\partial \sigma_i^\alpha} \cdot \frac{\partial \mathbf{X}_i}{\partial \sigma_i^\beta}, \quad (11)$$

whose elements at each position $\boldsymbol{\sigma}_i$ describe the sleeve's local morphology. Here, at each location (σ_i^0, σ_i^1) along the sleeve $\gamma_i(\boldsymbol{\sigma}_i)$ is a 2×2 tensor that captures the differential $\mathbf{X}_i$ in all directions, thus capturing the geometrical characteristics of the link surface. Using γ_i , Eq. (10) can be rewritten as

$$S_{\mathcal{M}(\mathcal{G})} = \sum_{i=1}^L \int_{\boldsymbol{\sigma}_i} d^2 \boldsymbol{\sigma}_i \sqrt{\det \gamma_i(\boldsymbol{\sigma}_i)}, \quad (12)$$

retrieving Eq. (1) of the main paper.

The metric tensor. The d -dimensional tensor γ_i is symmetric ($\gamma_{i,\alpha\beta} = \gamma_{i,\beta\alpha}$). Thus, it incorporates $d(d+1)/2$ d.o.f. in total, accounting for the d diagonal elements and the $d(d-1)/2$ entries of the top/bottom off-diagonal triangle. For example, in our case, where $d = 2$, one must specify three real functions $\gamma_{i,00}(\boldsymbol{\sigma}_i)$, $\gamma_{i,11}(\boldsymbol{\sigma}_i)$ on the diagonal, and $\gamma_{i,01}(\boldsymbol{\sigma}_i) = \gamma_{i,10}(\boldsymbol{\sigma}_i)$ for the off-diagonal terms. Note, however, that we can reparameterise the manifold through a coordinate transformation $\boldsymbol{\sigma} \rightarrow \boldsymbol{\sigma}' = f(\boldsymbol{\sigma})$ to obtain a further reduced description. Such reparametrisation, merely a change of coordinate system, leaves the physical quantities of the manifold un-changed [17], thus having no effect on, *e.g.*, the cost in (12). It does, however, if selected appropriately, remove d redundant d.o.f., thus reducing γ_i to just $d(d-1)/2$ independent entries. For $d = 2$, specifically, this contracts γ_i in Eq. (11) to just one d.o.f., *i.e.* a single scalar function.

To exemplify this, we consider a reparametrisation of σ_i^0 and σ_i^1 to obtain a diagonalised representation of γ_i as [18]

$$\gamma_{i,\alpha\beta}(\boldsymbol{\sigma}_i) = f_i(\boldsymbol{\sigma}_i) \delta_{\alpha\beta}, \quad (13)$$

a transformation often called the *isothermal coordinate system*. Here $f_i(\boldsymbol{\sigma}_i)$ is a real positive scalar function of the coordinates, and $\delta_{\alpha\beta}$ is the Kronecker delta function, which assumes unity on the diagonal and zero otherwise. The result is a simplified γ_i , in which both diagonal terms are identical and all off-diagonal entries vanish. The crucial point is that such diagonalised γ_i greatly simplify the calculation of $S_{\mathcal{M}(\mathcal{G})}$ in (12). Indeed with γ_i reduced to just the one function in (13) the cost in (12)

simplifies to

$$S_{\mathcal{M}(\mathcal{G})} = \sum_{i=1}^L \int_{\boldsymbol{\sigma}_i} f_i(\boldsymbol{\sigma}_i) \mathrm{d}^2 \boldsymbol{\sigma}_i, \quad (14)$$

an accessible integral for a broad choice of $f_i(\boldsymbol{\sigma}_i)$.

To gain an intuitive understanding of the isothermal coordinate transformation in (13), consider the meaning of the off-diagonal terms in γ_i . They represent an interdependence of $\partial \mathbf{X}_i / \partial \sigma_i^0$ and $\partial \mathbf{X}_i / \partial \sigma_i^1$. Such dependence captures the fact that in the surface elements, or *tiles* comprising the sleeve, the two coordinates σ_i^0 and σ_i^1 are, generally, not orthogonal, and hence a shift in one induces a subsequent change in the other. For finite tiles, such non-orthogonality is un-avoidable, as it is induced by the potentially curved morphology of the sleeve. This can be clearly observed in the catenoid sleeve illustrated in Fig. 8, where the sleeve curvature distorts the tiles into a trapezoidal form. The crucial point is that in the infinitesimal limit the tiles can always be approximated to be locally square-like. The result is a quadratic mesh in which the infinitesimal shifts in the directions $\mathrm{d}\sigma_i^0$ and $\mathrm{d}\sigma_i^1$ are locally orthogonal. This ensures that $\gamma_{i,01}(\boldsymbol{\sigma}_i) = \gamma_{i,10}(\boldsymbol{\sigma}_i) = 0$, as, indeed, both represent a scalar product between orthogonal vectors. Next, we can always re-scale the quad mesh to make the tiles of equal sides, *i.e.* perfect squares, which in (13) is expressed through the uniform diagonal entries of γ_i . Therefore, the meaning of the reparameterisation leading to Eq. (13) is simply to select a coordinate system which is orthonormal, *i.e.* both orthogonal (off-diagonal terms vanish) and normalised (diagonal terms are equal).

Finally, we note that in certain contexts the function $f_i(\boldsymbol{\sigma}_i)$ is expressed in exponential form as [18]

$$f_i(\boldsymbol{\sigma}_i) = e^{2\phi(\boldsymbol{\sigma}_i)}, \quad (15)$$

where we set $\phi(\boldsymbol{\sigma}_i) = \ln \sqrt{f_i(\boldsymbol{\sigma}_i)}$. Such representation, which merely reflects an alternative convention, has no effect on the proposed isothermal transformation. It is, however, useful in the string-theoretic analogy, to which we allude below in Sec. 4. In this notation, the manifold cost becomes

$$S_{\mathcal{M}(\mathcal{G})} = \sum_{i=1}^L \int_{\boldsymbol{\sigma}_i} e^{2\phi(\boldsymbol{\sigma}_i)} \mathrm{d}^2 \boldsymbol{\sigma}_i, \quad (16)$$

substituting $f_i(\boldsymbol{\sigma}_i)$ in (14) by the exponential function of (15).

3.1 Example: the catenoid sleeve

To demonstrate our formalism we revisit the catenoid sleeve of Fig. 8, oriented around the z axis. This represents a stereometric object characterised by two parameters: its height h and its minimal cross-section circumference w , capturing its width at the *waist*. The natural isothermal coordinates for this system are the cylindrical coordinate system

$$\boldsymbol{\sigma} = (\sigma^0, \sigma^1) \in \left[-\frac{\pi h}{w}, \frac{\pi h}{w} \right] \times [0, 2\pi], \quad (17)$$

where σ^0 describes the elevation along the z axis, and σ^1 is the polar angle around z . These two coordinates are clearly orthogonal, ensuring that γ 's off-diagonal terms vanish. In addition, we can use our d.o.f. to define σ^0 as $\sigma^0 = 2\pi z/w$ to ensure that γ 's two diagonal terms are equal.

As shown in the figure, under these coordinates, the catenoid surface follows

$$\mathbf{X}(\boldsymbol{\sigma}) = \begin{pmatrix} x(\boldsymbol{\sigma}) \\ y(\boldsymbol{\sigma}) \\ z(\boldsymbol{\sigma}) \end{pmatrix} = \left(\frac{w}{2\pi} \right) \begin{pmatrix} \cosh \sigma^0 \cos \sigma^1 \\ \cosh \sigma^0 \sin \sigma^1 \\ \sigma^0 \end{pmatrix}, \quad (18)$$

where z is in the range

$$-\frac{h}{2} \leq z \leq \frac{h}{2}, \quad (19)$$

capturing the catenoid height. The catenoid's circular cross-sections satisfy

$$x^2 + y^2 = \left(\frac{w}{2\pi} \right)^2 \cosh^2 \left(\frac{2\pi z}{w} \right), \quad (20)$$

and hence, its minimum circumference at the waist ($z = 0$) is, indeed, w .

To construct the metric tensor γ we first extract the partial derivatives of $\mathbf{X}(\boldsymbol{\sigma})$ as




Figure 8: **Catenoid sleeve.** The longitudinal coordinate σ^0 (red) and the azimuthal σ^1 (blue). The three-dimensional surface is captured by $\mathbf{X}(\boldsymbol{\sigma}) = (x(\boldsymbol{\sigma}), y(\boldsymbol{\sigma}), z(\boldsymbol{\sigma}))$, as appears in the equations at top of figure.

$$\begin{aligned}
\frac{\partial \mathbf{X}}{\partial \sigma^0} &= \left(\frac{w}{2\pi} \right) \begin{pmatrix} \sinh \sigma^0 \cos \sigma^1 \\ \sinh \sigma^0 \sin \sigma^1 \\ 1 \end{pmatrix}, \\
\frac{\partial \mathbf{X}}{\partial \sigma^1} &= \left(\frac{w}{2\pi} \right) \begin{pmatrix} -\cosh \sigma^0 \sin \sigma^1 \\ \cosh \sigma^0 \cos \sigma^1 \\ 0 \end{pmatrix},
\end{aligned} \tag{21}$$

which, in (11), yield

$$\gamma = \begin{pmatrix} \frac{\partial \mathbf{X}}{\partial \sigma^0} \cdot \frac{\partial \mathbf{X}}{\partial \sigma^0} & \frac{\partial \mathbf{X}}{\partial \sigma^0} \cdot \frac{\partial \mathbf{X}}{\partial \sigma^1} \\ \frac{\partial \mathbf{X}}{\partial \sigma^1} \cdot \frac{\partial \mathbf{X}}{\partial \sigma^0} & \frac{\partial \mathbf{X}}{\partial \sigma^1} \cdot \frac{\partial \mathbf{X}}{\partial \sigma^1} \end{pmatrix} = \left(\frac{w}{2\pi} \right)^2 \begin{pmatrix} \cosh^2 \sigma^0 & 0 \\ 0 & \cosh^2 \sigma^0 \end{pmatrix}, \tag{22}$$

where in the last step we used the identity $\cosh^2 x = 1 + \sinh^2 x$. As expected, thanks to our choice of coordinates, we arrived at the diagonal tensor form of Eq. (13) with $f(\boldsymbol{\sigma}) = (2\pi)^{-2} w^2 \cosh^2 \sigma^0$. In this specific application, thanks to the isotropic symmetry of the catenoid, the metric tensor is independent of σ_1 . However, for more general morphologies this independence is not guaranteed. We can now substitute $f(\boldsymbol{\sigma})$ into Eq. (14) to obtain

$$S_{\mathcal{M}(\mathcal{G})} = \int_{-\frac{\pi h}{w}}^{\frac{\pi h}{w}} \int_0^{2\pi} \left(\frac{w}{2\pi} \right)^2 \cosh^2(\sigma^0) d^2 \boldsymbol{\sigma} = \pi \left[\left(\frac{w}{2\pi} \right)^2 \sinh \left(\frac{2\pi h}{w} \right) + \frac{hw}{2\pi} \right], \tag{23}$$

providing the catenoid surface area.

4 The unified coordinate system—Atlas

To connect all links into a unified smooth manifold with a single global coordinate system, we seek the atlas σ (omitting the index i). Constructing this atlas requires us to sew together the boundaries of all links to ensure the continuity and smoothness of $\mathcal{M}(\mathcal{G})$. Following this process, the *link-by-link* integration of Eq. (12) simplifies to

$$S_{\mathcal{M}(\mathcal{G})} = \int d^2 \sigma \sqrt{\det \gamma}, \quad (24)$$

which, using Eq. (13), can be expressed as

$$S_{\mathcal{M}(\mathcal{G})} = \int d^2 \sigma e^{2\phi(\sigma^0, \sigma^1)}, \quad (25)$$

capturing the *manifold* surface area.

4.1 Quadratic differentials

The theoretical difficulty in expressing $S_{\mathcal{M}(\mathcal{G})}$ via Eq. (24) is to show that a globally differentiable atlas $\sigma = (\sigma^0, \sigma^1)$ indeed exists. Hence below, for the completeness of our methodological presentation, we briefly review relevant advances in string field theory [19, 20] that help establish the existence of such a global atlas. Particularly, we focus on *quadratic differentials* [21] that are used to construct an exact transformation from the local coordinate system, σ_i , to the global coordinates σ . To establish this, we first note that the sleeves we consider are orientable (as opposed to, *e.g.* Möbius strips), meaning that our network manifolds can be associated with the structure of *Riemann surfaces* [18]. This indicates that their two-dimensional coordinate system σ can be reduced into a single complex coordinate, z .

To express this, we rewrite the global coordinates σ as a complex coordinate, $z = \sigma^0 + i\sigma^1$. Let us also represent each chart i by a local complex coordinate, $z_i = \sigma_i^0 + i\sigma_i^1$. The challenge is, therefore, to relate the global complex variable z to the local complex variables z_i . This connection is expressed through the introduction of quadratic differentials [21]. Specifically, in our application, we focus on a special class of quadratic differentials $q(z) dz^2$, where $q(z)$ is a meromorphic function that has only second-order poles and first-order zeros [21]. Figure 9 provides examples of $q(z)$ that belong to this class.

In the atlas, it can be shown that the local complex variables z_i can be retrieved from the global z via the integral [22]




Figure 9: **Quadratic differentials.** A quadratic differential is a meromorphic function $q(z)$ that defines two orthogonal sets of horizontal (blue) and vertical (red) trajectories on a Riemann surface. The horizontal trajectories must be closed. **(a)** A quadratic differential $q(z)$ with three punctures ($M = 3$). It must have 3 poles and 2 zeros by dimensional analysis. When $q(z)$ is Jenkins–Strebel, it can be used to construct an atlas that divides the Riemann surface into multiple tube-like charts. The thick blue lines indicate how these tubes are smoothly sewn together. **(b)** When $q(z)$ is not Jenkins–Strebel, the charts are not tube-like because the horizontal trajectories are not closed. **(c,d)** Same results apply to $M = 4$, which corresponds to 4 poles and 4 zeros. **(e)** The $M = 3$ is mapped to a degree $k = 3$ bifurcation. **(f)** Similarly, the $M = 4$ case corresponds to trifurcation ($k = 4$).

$$z_i = \int_{\sim \text{chart } i} \sqrt{q(z)} dz, \quad (26)$$

such that integrating over the i th chart (or link i) recovers the local coordinates. Using (26), we can define the *vertical* and *horizontal* trajectories of the quadratic differential [22], such that along the vertical trajectories (red), the integral of $\sqrt{q(z)}$ remains real, whereas along the horizontal trajectories (blue), it remains purely imaginary. This represents trajectories that are, at each point, perpendicular to each other, creating a grid pattern of squares along the real/imaginary axes (Fig. 9). As a result, such mapping enables our quad-mesh representation for individual charts (Fig. 11 below).

In general, it can be shown that for any possible combinations of local coordinates z_i that satisfies our continuity conditions, there exists a corresponding meromorphic function $q(z)$ that can be globally

identified [21]. Hence, the existence of $q(z)$, even if remaining implicit, is sufficient to establish that the constructed physical network is, indeed, a *smooth manifold*. Still, we find it worthwhile to look into the exact expression for $q(z)$ for a few tractable examples. We begin with the simplest example of an atlas that consists of a single link i that has a circumference of w_i . The exact relation between the atlas coordinates z and the chart coordinates z_i is, in this case, provided by the meromorphic function

$$q(z) = \frac{w_i^2}{4\pi^2 z^2}. \quad (27)$$

Here $q(z)$ has two second-order poles, at $z = 0$ and $z \rightarrow \infty$, and no first-order zeros. Applying Eq. (26), we obtain $z_i = \int w_i (2\pi)^{-1} z^{-1} dz = w_i (2\pi)^{-1} \ln z$. This leads to the coordinate mapping

$$(\sigma_i^0, \sigma_i^1) = \left(\frac{w_i}{2\pi} \ln \sqrt{(\sigma^0)^2 + (\sigma^1)^2}, \frac{w_i}{2\pi} \arctan \left(\frac{\sigma^1}{\sigma^0} \right) \right), \quad (28)$$

in which σ_i^1 is, indeed, periodic in $[0, w_i]$, precisely as expected for the single chart coordinates.

Next, let us generalise this mapping to network manifolds representing tree graphs with M terminals. Topologically, such a manifold corresponds to a Riemann sphere (which is the simplest Riemann surface) with M punctures [22]. We seek the most general form of $q(z)$ on the corresponding M -punctured Riemann sphere, for which Eq. (26) can help retrieve all individual chart coordinates. Suppose that $q(z)$ has M second-order poles exactly matching the positions of the M punctures. To determine the number of zeros $q(z)$ should have, we need to ensure the smoothness of $q(z) dz^2$ in the limit $z \rightarrow \infty$. We therefore denote $t = 1/z$ and seek $q(1/t) d(1/t)^2$ at $t = 0$. By substitution we write $q(1/t) d(1/t)^2 = t^{-4} q(1/t) dt^2$, which converges to a finite limit at $t = 0$ only if we require $q(1/t) \sim t^4$. Returning back to our original notation, using z , this provides $q(z) \sim z^{-4}$. Consequently, $q(z)$ must have $2M - 4$ first-order zeros to balance the $2M$ dimensions induced by the second-order poles. Most generally, we can satisfy these conditions by setting

$$q(z) = C \prod_{m=1}^M (z - p_m)^{-2} \prod_{n=1}^{2M-4} (z - s_n), \quad (29)$$

where C is an arbitrary constant, p_m are the second-order poles, and s_n are the first-order zeros. Equation (29) is characterised by $3M - 3$ complex parameters, the M poles p_m , the $2M - 4$ zeros s_n , and one global constant C . The crucial point is that they cannot all be selected independently and freely. Certain combinations of these parameters may result in redundant representations of the same atlas, while other combinations may lead to non-sleeve-like charts that violate the periodicity in the σ_i^1 direction. According to string field theory [19, 20, 23–26], the quadratic differential $q(z) dz^2$

we seek must satisfy the *Jenkins–Strebel* conditions [21], which ensure the closure of every horizontal trajectory (Fig. 9a,c); these two examples correspond to a bifurcation motif (Fig. 9e) and a trifurcation motif (Fig. 9f), respectively. Without such closure, the horizontal trajectory would spiral infinitely and never reach an endpoint (Fig. 9b,d). In the local coordinate z_i this would mean that there is no guarantee to exhibit periodic behaviour in the imaginary direction. It would therefore be unsuitable for constructing a proper sleeve. The Jenkins–Strebel conditions allow only $M - 3$ free complex parameters. The rest are constrained by the continuity and smoothness requirements. The resulting family of Jenkins–Strebel quadratic differentials of the form (29) can help represent all existing tree-graph surface manifolds [19, 20].

5 Surface minimisation

To obtain the manifold with the minimum surface area (Fig. 10a), we seek

$$\arg \min_{\mathcal{M}(\mathcal{G})} \{S_{\mathcal{M}(\mathcal{G})}\}, \quad (30)$$

taking $S_{\mathcal{M}(\mathcal{G})}$ from Eq. (25), where $\mathcal{M}(\mathcal{G})$ is a physical network manifold. As explained in the main paper, this minimisation is equivalent to solving the equation of motion for the Nambu–Goto action in string theory [17] while subject to the functional constraint as in Eq. (2) of the main paper. Without loss of generality, we assume that the shortest circumference around each link (chart) is along the σ^1 direction and rewrite the constraint as

$$\oint d\sigma^1 \|\partial \mathbf{X} / \partial \sigma^1\| \geq w, \quad (31)$$

such that the shortest circumference must be at least w in length. This ensures that the physical links can sustain a flow (Fig. 10b). For example, in the catenoid of Fig. 8 this constraint is limited by the circumference at the waist.

Finding the desired $\mathcal{M}(\mathcal{G})$ requires us to enumerate all possible manifolds $\mathcal{M}$ that can be constructed around the skeletal graph $\mathcal{G}$ and satisfy constraint (31). We then need to extract their d -dimensional cost $S_{\mathcal{M}(\mathcal{G})}$, and select the optimal manifold, for which this cost is minimised.

Numerically, this optimisation problem comprises four steps:

- (i) Begin with M spatially embedded terminals.
- (ii) Enumerate all tree graphs $\mathcal{G}$ that link these M terminals via finite paths, including also ones with added intermediate links.
- (iii) For each graph $\mathcal{G}$ enumerate all manifolds $\mathcal{M}(\mathcal{G})$ that satisfy the systolic constraint of (31). These manifolds may include any set of sleeves with distinct morphology (length, thickness) that can be dressed around the links of $\mathcal{G}$ and properly sewn together as explained in Sec. 3.
- (iv) For each $\mathcal{M}(\mathcal{G})$, list all links $i = 1, \dots, L$ and determine, for each link, the function $\phi_i(\boldsymbol{\sigma})$ in Eq. (15) that minimises the total surface area $S_{\mathcal{M}(\mathcal{G})}$. Note that this step cannot be solved link-by-link, since the links must be properly sewn together, and hence the morphology of one link feeds into that of its adjacent links at the inter-link boundaries.

The desired manifold is the one that, following steps 1–4 above, has the overall minimum $S_{\mathcal{M}(\mathcal{G})}$.




Figure 10: **Surface minimisation.** (a) To obtain the surface area we integrate over its induced metric tensor $\gamma_{\alpha\beta}$, which captures the infinitesimal surface area elements constructing all sleeves. (b) Minimising surface alone may lead to dysfunctional links. This is observed under Plateau’s soap film transition. There, if the distance h is increased beyond a threshold ($h/w_0 \approx 0.168$), the optimal solution favours two disconnected planes (top right). In contrast, the links in a physical network must maintain a minimal circumference w in order to remain structurally intact or to sustain their functionality (e.g. nutrient transport). Therefore, even if h is large the sleeves do not reach a zero cross section, but rather continue to have a circumference of at least w (bottom right).

To solve this numerical challenge we rely on the existing solution of the *systolic surface minimisation* problem [27]. There, one seeks to construct a minimal surface under the condition that all non-contractible loops surrounding it, at any point, have a length of at least w . This maps exactly to our constraint of (31). The only distinction between this classic challenge and our current problem is that here we also constrain the manifold to be immersed in space. Hence, the manifold solutions we seek have pre-set boundary conditions, determined by the terminals positioned in three-dimensional space. To realise these boundary conditions, we map the point location $\mathbf{X}_m$ of each terminal m and then introduce a perfect circle with radius $w/(2\pi)$ freely rotatable around $\mathbf{X}_m$. We then force the manifold boundaries to be sewn to these circles. As a result, each terminal node is reached by a sleeve that, at one end, connects to its surrounding circle.

For tree-graph network manifolds, *i.e.* graphs without loops, there exists an exact solution to the systolic surface minimisation problem, as given by Strebel’s theorem [21]. The theorem states that *among all metric tensors satisfying the systolic length constraint, the flat metric tensor is the one with minimal surface area* [20]. In Eq. (13), the flat metric tensor is expressed by setting $e^{2\phi_i(\sigma)} \equiv 1$, independently of σ . From a geometric perspective, this translates to all sleeves having a locally perfect cylindrical surface. Quite clearly, this cylindrical solution provides the minimal surface subjected to a given systole.

The challenge is that Strebel’s minimal-surface solution neglects the constraint that the surface must

also be immersible within the ambient three-dimensional space and connect smoothly to the terminals and intermediate nodes of $\mathcal{G}$. If we simply adopt Strebel’s cylindrical solution, the resulting surfaces would exhibit conical singularities at the sleeve intersections, rendering the solution un-physical. Therefore, Strebel’s solution $e^{2\phi_i(\boldsymbol{\sigma})} \equiv 1$ can only serve as a useful starting point for our analysis, but cannot offer the precise solution we seek. To obtain the actual desired smooth manifolds we must allow the sleeves to deform at their intersections and break their cylindrical envelope.

To accomplish this, in our implementation, we build upon Strebel’s solution, but modify it to allow

$$e^{2\phi_i(\boldsymbol{\sigma})} \geq 1, \tag{32}$$

or equivalently $\phi_i(\boldsymbol{\sigma}) \geq 0$, and $\sigma^1 \in [0, w]$. Equation (32) ensures that the systolic constraint (31) continues to hold at every point along each sleeve. Indeed, equality in Eq. (31) is already satisfied when $\phi_i(\boldsymbol{\sigma}) = 0$. Relaxing this to $\phi_i(\boldsymbol{\sigma}) \geq 0$ simply adds surface area, and can only increase the value of the integral on the left-hand side of Eq. (31). Therefore, our relaxed formulation in Eq. (32) guarantees that all systoles remain $\geq w$, while allowing the sleeves to expand locally as needed to accommodate stretching and bending of $\mathcal{M}(\mathcal{G})$ at the intersections of two (or more) charts.

Strictly speaking, since the constraint in Eq. (32) is purely local in nature, it is not designed to obtain the global minimum. Hence, it does not necessarily minimise the total surface area, but rather, it yields surface that is locally area-minimising. This means that infinitesimal variations in the vicinity of any point do not reduce the area. Nonetheless, a globally distinct surface configuration may still achieve a lower total area.

Despite this limitation, the morphological flexibility introduced by Eq. (32) produces realistic geometries. Therefore, although it is local in formulation, it remains highly relevant from a physical standpoint. In practice, the resulting manifolds tend—quite intuitively—towards cylindrical sleeves, with $\phi_i(\boldsymbol{\sigma}) \rightarrow 0$ in areas located far from the link intersections, then exhibiting modest deviations from cylindricity as we approach the intersections.

6 Numerical implementation

6.1 Quad-mesh tiling

To numerically evaluate the surface of a given sleeve/manifold, we tile it using quad meshes, as detailed in Ref. [28]. The surface is divided into a large number of small square tiles (which may vary in size). These tiles represent the infinitesimal surface elements $d\sigma_i^0 d\sigma_i^1$ whose integration provides the complete surface area. By choosing the orthogonal parametrisation, in which, at each point σ_i^0 is perpendicular to σ_i^1 , the tiling follows a grid pattern, in which all tiles, if indeed small enough, have 90° angles. Then, properly scaling the two components, we can ensure that they are both of equal size, rendering the tiles to be perfect squares (Sec. 3).

We denote the metric tensor of this quad mesh by $\gamma_i^{\text{quad}}(\sigma_i)$. The orthogonality between the σ_i components ensures that $\gamma_{i,01}^{\text{quad}}(\sigma_i) = \gamma_{i,10}^{\text{quad}}(\sigma_i) = 0$; and the re-scaling into square tiling guarantees that $\gamma_{i,00}^{\text{quad}}(\sigma_i) = \gamma_{i,11}^{\text{quad}}(\sigma_i)$. Therefore, the diagonalised quad-mesh metric tensor can be written in the form (13) as

$$\gamma_{i,00}^{\text{quad}} = \gamma_{i,11}^{\text{quad}} = e^{2\phi_i^{\text{quad}}(\sigma_i)} \quad (33)$$

where we used the notation of (15) to express the function $f_i(\sigma_i)$. In this context, the physical interpretation of the real function $e^{2\phi_i^{\text{quad}}(\sigma_i^0, \sigma_i^1)}$ becomes clear: it captures the dimensions of the infinitesimal tile at position (σ_i^0, σ_i^1) in the quad mesh. In areas where the sleeves are stretched, for example, the tiles become expanded, and hence, at these locations the γ_i^{quad} 's terms become greater. This was exemplified in Sec. 5, where we showed that for a perfect cylinder

we have $\gamma_{i,00}^{\text{quad}} = \gamma_{i,11}^{\text{quad}} = 1$ throughout, and as the cylinder deforms towards its intersection with a neighbouring sleeve these terms become ≥ 1 . Hence, in the quad mesh framework, the surface of a sleeve/manifold is fully determined by the scalar function $\phi_i^{\text{quad}}(\sigma_i)$, which encapsulates the local sleeve morphology through the effective stretching or contraction of the tiles at each locale σ_i .

The challenge is that in numerical calculations, the mesh is inherently discrete and only composed of a finite number of finite-size tiles. As these tiles are not truly infinitesimal, they cannot be considered




Figure 11: **Quad mesh.** A quad mesh is made up of tiles arranged in a grid layout, where each tile is a skew quadrilateral (not necessarily planar). Each tile can be freely deformed, with the only constraint that the Varignon parallelogram inside each tile, formed by connecting the four midpoints of each tile's sides, must be a rhombus. The ratio between the rhombus diagonals captures the stretching of the mesh in a specific direction. Absent such stretch, the Varignon parallelograms become perfect squares. The square/rhombic Varignon parallelograms join together into a checkerboard pattern that helps approximate the smooth surface area.

as perfectly planar squares (Fig. 11). Rather they have a curved surface and potentially trapezoidal boundaries. To address this, we rely on the quad mesh framework [28] that is specifically designed to treat such discrepancies. In this framework, instead of limiting our tiles themselves to be perfectly planar squares, we only require that their four midpoints form such squares (known as the Varignon square). This gives us the freedom to tile our smooth surface more accurately with curved and trapezoidal tiles. The crucial point is that under this construction these, potentially curved tiles, induce a square grid, the *checkerboard* of Fig. 11, whose area can be used to approximate the exact manifold surface area.

Algorithm and code. Our complete surface minimisation algorithm, *min-surf-netw*, which employs the above tiling scheme, is available on GitHub with all relevant usage documentation at <https://github.com/Barabasi-Lab/min-surf-netw>.

6.2 Optimisation target function

To optimise the manifold $\mathcal{M}(\mathcal{G})$ we set the target function E to

$$E = w_{\text{iso}}E_{\text{iso}} + w_{\text{glue}}E_{\text{glue}} + w_{\text{terminal}}E_{\text{terminal}} + w_{\text{fair}}E_{\text{fair}} + w_{\text{surface}}E_{\text{surface}}, \quad (34)$$

assigning relative weights for five distinct optimisation objectives:

Isometric cost E_{iso} . This term accounts for the deformation during the isometric mapping of the quad mesh into Euclidean space. As shown in Fig. 11, each tile of the quad mesh must satisfy the condition that the Varignon parallelogram inside each tile, formed by connecting the four midpoints of each tile's sides, is a rhombus or a perfect square. In our optimisation, we adjust the length l of each sleeve along the longitudinal direction. Hence, the Varignon parallelogram becomes elongated, *i.e.* a rhombus. We set the short diagonal to 1 and the longer to $Z \geq 1$.

For tile r with corners given by the coordinates (in counterclockwise order): $(x_{r,0}, y_{r,0}, z_{r,0})$, $(x_{r,1}, y_{r,1}, z_{r,1})$, $(x_{r,2}, y_{r,2}, z_{r,2})$, and $(x_{r,3}, y_{r,3}, z_{r,3})$, we write [29, 30]

$$\begin{aligned} A_r &= [(x_{r,0} - x_{r,2})^2 + (y_{r,0} - y_{r,2})^2 + (z_{r,0} - z_{r,2})^2] - (1 + \Delta\lambda_r^2) \\ B_r &= [(x_{r,1} - x_{r,3})^2 + (y_{r,1} - y_{r,3})^2 + (z_{r,1} - z_{r,3})^2] - (1 + \Delta\lambda_r^2) \\ C_r &= [(x_{r,0} - x_{r,2})(x_{r,1} - x_{r,3}) + (y_{r,0} - y_{r,2})(y_{r,1} - y_{r,3}) + (z_{r,0} - z_{r,2})(z_{r,1} - z_{r,3})] \\ &\quad - (1 + \Delta\lambda_r^2) \cos(2 \tan^{-1} |Z_r|). \end{aligned} \quad (35)$$

These three terms tend to zero when the inscribed Varignon parallelogram within each tile is congruent with the Varignon parallelogram inside a rectangle of dimensions $1 \times Z$. We therefore seek to minimise

the isometric cost over all R tiles via

$$E_{\text{iso}} = \sum_{r=1}^R (A_r^2 + B_r^2 + C_r^2), \quad (36)$$

ensuring that our isometric mapping is as accurate as possible. Under Eq. (35) the area of tile r is provided by

$$S_r = Z_r (1 + \Delta\lambda_r^2), \quad (37)$$

a rectangle of dimensions $1 \times Z_r$ multiplied by a conformal factor $(1 + \Delta\lambda_r^2)$. In our numerical implementation we used our freedom to scale the longitudinal axis of each tile and set $Z_r \equiv Z$ uniformly for all tiles within a specific sleeve.

Continuity cost E_{glue} . Consider the boundary coordinates of link i , $(x_{m,i}, y_{m,i}, z_{m,i})$, where $m = 1, \dots, M$ counts the discrete points along this boundary (in the quad mesh). The manifold construction requires that the boundary coordinates for i 's adjacent link j , $(x_{m,j}, y_{m,j}, z_{m,j})$, precisely coincide with those of i , ensuring that the two links are properly *glued* to each other. Hence we assign the cost

$$E_{\text{glue}} = \sum_{(i,j) \in B} \sum_{m=1}^M \left[(x_{m,i} - x_{m,j})^2 + (y_{m,i} - y_{m,j})^2 + (z_{m,i} - z_{m,j})^2 \right], \quad (38)$$

where $B = \{(i, j) \mid \text{links } i, j \text{ intersect}\}$ is the set of all intersecting link pairs.

Terminal cost E_{terminal} . This term is analogous to E_{glue} , only this time, ensuring that the sleeves properly attach to the terminal boundaries, *i.e.* the circular circumferences constructed around all terminal nodes. Specifically, we consider link i 's boundary $x_{m,i}, y_{m,i}, z_{m,i}$ with terminal q at (x_q, y_q, z_q) and quantify their deviation via

$$E_{\text{terminal}} = \sum_{(i,q) \in T} \left\{ \left[\left(\frac{1}{M} \sum_{m=1}^M x_{m,i} \right) - x_q \right]^2 + \left[\left(\frac{1}{M} \sum_{m=1}^M y_{m,i} \right) - y_q \right]^2 + \left[\left(\frac{1}{M} \sum_{m=1}^M z_{m,i} \right) - z_q \right]^2 \right\}, \quad (39)$$

where $T = \{(i, q) \mid \text{link } i \text{ connects to terminal } q\}$ is the set of all link-terminal coincidences.

Fairness cost E_{fair} . The construction of the quad mesh requires an even distribution of tiles, which we enforce by the fairness cost [29, 30]. To construct E_{fair} , we consider every three consecutive coordinates, (x_i, y_i, z_i) , (x_j, y_j, z_j) , and (x_k, y_k, z_k) , respectively, and keep their extrinsic curvature

minimal, by adding the cost function [30]

$$E_{\text{fair}} = \sum_{i,j,k} \frac{(x_i - 2x_j + x_k)^2 + (y_i - 2y_j + y_k)^2 + (z_i - 2z_j + z_k)^2}{(x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2 + (x_j - x_k)^2 + (y_j - y_k)^2 + (z_j - z_k)^2}. \quad (40)$$

Equation (40) seeks to minimise the second derivative at each point, which in discrete space is often expressed as $x_i - 2x_j + x_k$. This is designed for the optimised surface solutions to avoid excessive curvature, ensuring smoothness. The pre-factor w_{fair} is varied throughout the optimisation process, beginning at a relatively large value, to avoid irregularities at the sleeve boundaries, then gradually relaxing to zero as the optimisation progresses, and surface area becomes the dominant factor.

Surface area cost E_{surface} . Finally, this term accounts for the heart of our algorithm and calculates the total manifold surface area. In our numerical implementation this is obtained by summing over the areas of all tiles, which using (37) provides

$$E_{\text{surface}} = \sum_{r=1}^R Z (1 + \Delta\lambda_r^2), \quad (41)$$

an aggregation of small rectangular tiles.

Selecting the weights. The weights w_X in (34) determine the relative importance of each cost component. For example, E_{glue} and E_{terminal} represent stringent constraints that must be satisfied to a high degree. The isometric cost E_{iso} , on the other hand, is more flexible, and represents a *soft* constraint. We, therefore, set their weights accordingly: $w_{\text{iso}} = 1$, $w_{\text{glue}} = 10^3$ and $w_{\text{terminal}} = 10^3$. The remaining weights, w_{fair} and w_{surface} , must be set dynamically [30], as detailed below for each specific application.

Search space. Seeking the minimisation of E in (34), we have multiple degrees of freedom that define our variable search space:

- For each link i , we have $l_i \times w_i \times 3$ coordinate variables, *i.e.* x, y, z coordinates along the length (l_i) and circumference (w_i) of the sleeve.
- For each tile r , we must tune the conformal weight factor $1 + \Delta\lambda_r^2$. This must be done for all tiles in all links; a specific link i comprises $(l_i - 1) \times w_i$ tiles. Note that we set this factor quadratic in $\Delta\lambda_r$ to ensure that, regardless of the optimisation outcome, it is always positive. This guarantees that the resulting manifold satisfies the functional constraint (31), as expressed via (32), namely that the sleeves can only stretch where needed, and never contract.
- Finally, we seek a single factor Z , reflecting the stretch of the link in the longitudinal direction.

In our implementation we search the variable space of $\sqrt{Z} \in \mathbb{R}^+$ instead of Z to ensure that Z remains positive during the minimisation process.

Implementation. In the main paper, we examined two specific scenarios: (i) the emergence of a trifurcation in a 4-node system (main paper Fig. 3); (ii) the bifurcation angle in a 3-node system (main paper Fig. 4). Below we describe the specific implementation in each of these calculations.

Trifurcation. We considered 5 links with a uniform circumference of $w = 50$ tiles. Their lengths remain adjustable by a global scaling factor. The initial network was constructed as a Steiner graph, which incorporates two separate bifurcations linked via an intermediate link i . During the optimisation process the length of i may change, depending on the optimal surface morphology, expressed via the magnitude of the longitudinal variable Z_i . A trifurcation is said to emerge when Z_i approaches zero. For this application we set the surface minimisation prefactor to $w_{\text{surface}} = 10^{-3}$ and the fairness prefactor according to the following sequence:

$$\begin{aligned}
w_{\text{fair}} &= 10^0 && \text{for } 5 \times 10^3 \text{ iterations,} \\
w_{\text{fair}} &= 10^{-1} && \text{for } 5 \times 10^3 \text{ iterations,} \\
w_{\text{fair}} &= 10^{-2} && \text{for } 5 \times 10^3 \text{ iterations,} \\
w_{\text{fair}} &= 10^{-3} && \text{for } 5 \times 10^3 \text{ iterations,} \\
w_{\text{fair}} &= 10^{-4} && \text{for } 5 \times 10^3 \text{ iterations,} \\
w_{\text{fair}} &= 10^{-5} && \text{for } 5 \times 10^3 \text{ iterations,} \\
w_{\text{fair}} &= 10^{-5} && \text{for } 10^5 \text{ iterations.}
\end{aligned} \tag{42}$$

This sequence, characterised by an exponentially decaying w_{fair} , follows the suggested implementation of the method in Refs. [29, 30]. It is designed to ensure that the influence of the surface fairness term does not decrease abruptly, but rather it is gradually annealed towards its minimum (small) value ($w_{\text{fair}} = 10^{-5}$).

Bifurcation angle. Here we considered 3 links, with uniform length of $l = 24$ (*i.e.* 23 tiles). Two of the links had a circumference of $w = 50$ tiles, and the third, an adjustable circumference w' . The initial network was constructed with a symmetric bifurcation angle $\theta = 2\pi/3$, which follows the Steiner scheme. We then applied our algorithm to obtain the surface minimised angles. The surface minimisation pre-factor was set to $w_{\text{surface}} = 10^{-2}$, and w_{fair} followed the exact same sequence as in Eq. (42).

6.3 Computational complexity

The computational complexity of our algorithm is driven by two factors. The first is the number of iterations used in the optimisation sequence of (42), which in our implementation is fixed (at $\sim 10^5$). The second factor is the total number of square tiles, R , which we seek to maximise in order to increase precision. Our code runtime is linear in R , allowing us to increase this factor under a manageable computational cost. Our simulations required generating manifolds of variable values of R ranging from $\sim 5 \times 10^3$ to 4×10^4 . In Fig. 12 we track the actual runtime vs. R , observing the anticipated linear scaling.




Figure 12: **Computational complexity.** We ran our algorithm for surface minimisation on a four node motif, tracking the bifurcation to trifurcation transition. For each run we show the computation time T (hours) vs. the number of tiles R (circles). We find that the runtime grows linearly with R , as predicted (dashed line).

7 Empirical analysis of branching

Our empirical analysis is conducted on the six physical network datasets listed in Table 3 of Sec. 1. In the analysis, we extracted the networks’ degree distribution (main paper Fig. 1d), their planarity (main paper Fig. 1e), and angle symmetry (main paper Fig. 1f). Below we outline the details of each of these analyses.

7.1 Degree distribution

Link cutoff. In principle, to extract the degrees of all nodes, we must map the skeletal graph $\mathcal{G}$ underlying the manifold $\mathcal{M}(\mathcal{G})$. The challenge is that, as explained above in Sec. 1.1, skeletonisation is potentially sensitive to the spatial resolution of the data and risks obscuring high-degree structures—a phenomenon known as *extrinsic trifurcations* [31]. A classic example is when an actual trifurcation appears as two separate bifurcations when the node volume/morphology is translated to be a spurious intermediate link (Fig. 2). Alternatively, two bifurcations may occur close to one another, and become *swallowed* within the thickness of the links or the volume of the nodes, and are thus erroneously classified as a single trifurcation.

To solve these ambiguities, we set a threshold below which we consider a link to be a skeletonisation artifact. Consider two bifurcations in $\mathcal{G}$, separated by an intermediate link of length l_{int} . We set a cutoff that in case [31]

$$\frac{l_{\text{int}}}{w/2\pi} \leq 1 \quad (43)$$

we disregard the link, and merge the two bifurcations into a single trifurcation. The rationale is that if the length to thickness ratio of a link is smaller than one, then it is likely a skeletonisation artifact rooted in the intersection morphology; the denominator $w/2\pi$ represents the typical link radius.

Trifurcation. In the main paper, we track the structural transition from a double bifurcation to a single trifurcation. We begin with four nodes located at the edge of a perfect tetrahedron, whose bounding sphere has radius r . We then construct surface minimised manifolds to link these nodes with systolic constraint w . We used $\lambda = l_{\text{int}}/w$, the ratio between the intermediate link length and the circumference w , to characterise this transition. When $\lambda > 0$, the intermediate link is large and the two bifurcations are separable. If, however, $\lambda \rightarrow 0$ then the bifurcations merge into one node, thus capturing a single trifurcation. As the order parameter we set $\chi = w/r$, the ratio between the link circumference and the typical length scale r of the terminal configuration. Here, the limit where $\chi \rightarrow 0$ captures links of vanishing thickness, converging to the Steiner limit. Physicality, therefore, kicks in

as $\chi \sim 1$. In practice, we observed a sharp transition from $\lambda > 0$ to $\lambda \rightarrow 0$ at $\chi \approx 0.83$ (Fig. 13, dashed line). This is in contrast to the Steiner prediction, in which λ_{Steiner} is linearly proportional to r , and hence scales as χ^{-1} (grey solid line).

To express the predicted transition via an analytical function we use [32]

$$\lambda \approx \lambda_{\text{Steiner}} f\left(\frac{\lambda_{\text{Steiner}} - 0.212 \times 0.83^{-1}}{\sigma}\right), \quad (44)$$

where $f(x) = (1 + e^{-x})^{-1}$ is the sigmoid function. This analytical approximation is shown in Fig. 13 as a green dashed line, converging to λ_{Steiner} in the limit of small χ , and transitioning to $\lambda \rightarrow 0$ when $\chi > 0.83$. In (44) the parameter σ controls the steepness of the λ drop near the threshold, here set to $\sigma = 10^{-2}$. The parameter 0.212 in the numerator of (44) is attuned for a perfect tetrahedral node configuration, under which the Steiner prediction provides $\lambda_{\text{Steiner}} \approx 0.212\chi^{-1}$ (Fig. 13, grey solid line). Hence, one can rewrite (44) as

$$\lambda \approx \frac{0.212}{\chi} f\left(\frac{0.212 - 0.26\chi}{\chi\sigma}\right), \quad (45)$$

in which the dependence on χ appears explicitly.

λ distribution (main paper Fig. 3). In the main paper we also examine $P_\lambda(x)$, the probability density to observe $\lambda \in (x, x + dx)$. For notation simplicity we denote this function simply by $P(\lambda)$. To obtain the expected $P(\lambda)$ we generated 10^6 sets of four terminals randomly located within the unit cube. We then extracted the Steiner graph for each of these sets and measured λ_{Steiner} . In Fig. 3h of the main paper we show that the resulting $P(\lambda_{\text{Steiner}})$ (grey), as obtained from the 10^6 quadruplets. As expected the distribution approaches zero in the limit of $\lambda_{\text{Steiner}} \rightarrow 0$. This reflects the fact that according to Steiner, there is *always* a finite-size intermediate link.

Extracting the corresponding surface minimised manifold for each of these 10^6 terminal sets is unscalable. Yet, we can approximate their λ values from the observed λ_{Steiner} using (44). This allows us to evaluate $P(\lambda)$ for our random-terminal ensemble (main paper Fig. 3h, green). We find that $P(\lambda)$ is qualitatively distinct from $P(\lambda_{\text{Steiner}})$, allowing a finite probability density to observe $\lambda = 0$. This is quite expected, as, indeed, our surface minimised $\mathcal{M}(\mathcal{G})$ predicts trifurcations, *i.e.* $\lambda = 0$, for a finite




Figure 13: **Trifurcation transition.** The trifurcation parameter λ vs. χ . In the Steiner model we have $\lambda \approx 0.212\chi^{-1}$ (grey solid line). Under surface minimisation we observe a sharp transition to $\lambda \rightarrow 0$ when $\chi > 0.83$ (symbols). We approximate this transition via the analytical function in (44), shown by the green dashed line. Here, the steepness parameter is set to $\sigma = 10^{-2}$.




Figure 14: **Distributions of internal and external legs.** Both $P(l_{\text{int}}/w)$ and $P(l_{\text{ext}}/w)$ follow generalised Gamma distributions. Yet, while $P(l_{\text{ext}}/w \rightarrow 0)$ drops to zero, $P(l_{\text{int}}/w \rightarrow 0)$ remains positive. (a)–(f) Real physical networks from blood vessels to human neurons.

fraction of terminal configurations. In Eqs. (44) and (45) this is precisely captured by the extended $\lambda = 0$ tail observed for $\chi > 0.83$ (Fig. 13, green dashed line).

In our empirical analysis (main paper Fig. 3i–n), we encountered the opposite challenge. There $P(\lambda)$, the physical network intermediate link, can be directly extracted from the data, whereas the corresponding Steiner intermediate link length λ_{Steiner} is inaccessible. Fortunately, as we show in Fig. 15 λ_{Steiner} is, on average, linearly dependent on the length of the external links, or, more generally, on the typical length scale of the tetrahedral configuration. We can therefore evaluate the length of the un-observed Steiner intermediate link by measuring the existing external link lengths. This enables us to represent $P(\lambda_{\text{Steiner}})$ by the length distribution of the external legs. We find that the empirically observed distributions, indeed, closely resemble the theoretical predictions of $P(\lambda)$ and $P(\lambda_{\text{Steiner}})$: the former exhibiting the signature $P(\lambda \rightarrow 0) \neq 0$, and the latter having $P(\lambda \rightarrow 0) \rightarrow 0$ (Fig. 3 of the main paper). The best fit with the theoretical $P(\lambda)$ was obtained when setting $\sigma = 0.05$ in (44).

In Fig. 14 we plot the length distribution of both the external legs $P(l_{\text{ext}}/w)$ and the internal ones $P(l_{\text{int}}/w)$, as obtained from our six empirical datasets. We find that they can both be approximated by a generalised Gamma distribution




Figure 15: **Relation between internal and external legs.** l_{int}/w and l_{ext}/w follow a biased linear relation, suggesting a positive threshold on l_{ext}/w , below which l_{int}/w becomes negative and un-physical. (a)–(f) Real physical networks from blood vessels to human neurons.

$$P(x) = \begin{cases} \frac{\gamma e^{-\left(\frac{x-\mu}{\beta}\right)^\gamma} \left(\frac{x-\mu}{\beta}\right)^{\alpha\gamma-1}}{\beta\Gamma(\alpha)} & x > \mu \\ 0 & x \leq \mu \end{cases}, \quad (46)$$

by fitting four independent parameters α , β , γ , and μ . This allows us to extrapolate the behaviour of $P(x)$ in the limit where $x \rightarrow 0$. Across all networks we find that the two distributions behave distinctively in that limit: while $P(l_{\text{int}}/w \rightarrow 0)$ is positive (light colours), $P(l_{\text{ext}}/w \rightarrow 0)$ tends to zero (dark colours). This, again, suggests that in contrast to the external legs, the internal ones tend to vanish as we approach the trifurcation limit ($l_{\text{int}}/w \rightarrow 0$). This is, once again, consistent with the surface minimisation prediction, and contrary to Steiner.

Finally, in Fig. 15 we plot l_{int}/w vs. l_{ext}/w for all our empirical networks. Using principal component analysis, we find that l_{int}/w and l_{ext}/w follow a biased linear relationship of the form

$$l_{\text{int}} \approx A(l_{\text{ext}} - Bw). \quad (47)$$

Consequently the internal leg *hits* zero when $l_{\text{ext}}/w \leq B$. This captures precisely the point in which the internal leg vanishes and the double-bifurcation transitions into a single trifurcation. For a perfect tetrahedral configuration we have $\chi \approx 0.943(l_{\text{ext}}/w)^{-1}$, and hence the transition point is predicted to occur at $\chi \geq 0.943/B$. In our six datasets we find that B ranges from 0.67 to 1.99, with an average of 1.41. This translates to $\chi \geq 0.943/1.41 \approx 0.67$, which is close to the threshold of 0.83 observed earlier in Fig. 13.

7.2 Trifurcation planarity

The planarity of bifurcations is tested in the main paper (also presented here in Fig. 16a). To complete this analysis here, we further investigate higher degrees, such as trifurcations. Each trifurcation ($k = 4$) consists of four trios, allowing us to measure four angles Ω —one per each trio: $\{\Omega_{123}, \Omega_{124}, \Omega_{134}, \Omega_{234}\}$. Extracting Ω from all the trifurcation-based trios, we obtain $P_{\text{Tri}}(\Omega)$, characterising the planarity of the $k = 4$ motifs. As opposed to $P_{\text{Bi}}(\Omega)$ (Fig. 16a), where the peak is sharply close to $\Omega \approx 2\pi$, $P_{\text{Tri}}(\Omega)$ in Fig. 16b is more spread over a range of angles $\Omega \in [0, 2\pi]$. This represents a fundamental departure from the Steiner prediction (vertical grey line), placing most real networks close to the random branching pattern (black solid line).




Figure 16: **Planarity.** (a) The probability density $P_{\text{Bi}}(\Omega)$ vs. Ω as obtained from all bifurcations in our empirical network ensemble. The distribution of the solid angle Ω of each link trio (three links per bifurcation) has a peak close to but not equal to the Steiner prediction $\Omega \approx 2\pi$ (grey vertical line). (b) The probability density $P_{\text{Tri}}(\Omega)$ vs. Ω as obtained from all trifurcations. As each trifurcation consists of four links, we divide it into four separate link trios, extracting Ω from each. The observed $P_{\text{Tri}}(\Omega)$ is profoundly distinct from the Steiner prediction (grey vertical line), aligning more closely around a random branching pattern (thick grey solid line).

7.3 Angle asymmetry

To examine the branching angles, as we do in Fig. 2 of the main paper, we seek bifurcations that exhibit a main and a secondary branch. We therefore collected all three-link bifurcations from each of our empirical datasets and applied two filters:

1. We keep only bifurcation motifs in which two of the three links have approximately equal systoles w . As a threshold we select those motifs in which the relative difference between the two link systoles was 10% or less. We then label the third link systole as w' .
2. Following criterion (43) we discard links whose length-to-circumference ratio is $l/w < 2 \times (2\pi)^{-1}$. Such short links are considered to be a part of a potential trifurcation, rather than an intermediate bifurcation branch.

We consider that links may be curved and hence exhibit varying angles. To overcome this we set the desired branching angle using their two end points—at the bifurcation and at their opposite node—capturing the average link branching angle. We also take into account the fact that the link thickness may vary, and thus use the average circumference along the link trajectory to extract w and w' .

Branching vs. sprouting. To analyse the angles of all bifurcations, we examine the three-dimensional steering angle $\Omega_{1 \rightarrow 2}$. This represents a solid angle, which can be related to the branching angle θ via $\Omega_{1 \rightarrow 2} = 4\pi \sin^2[(\pi - \theta)/4]$ (Fig. 17). In Fig. 2 of the main paper (reproduced in Fig. 18), we present $\Omega_{1 \rightarrow 2}$ vs. the circumference ratio $\rho = w'/w$. We cover the range $0 \leq \rho \leq 1.5$, but exclude the data around $\rho = 1$ ($0.99 \leq \rho \leq 1.01$), as it is dominated by noise, and most likely represents spurious results, driven by errors in the skeletonisation process. Our first observation is that all networks exhibit a positive correlation between $\Omega_{1 \rightarrow 2}$ and ρ , featuring a Spearman correlation of 0.53 (human neuron), 0.45 (fruit fly neuron), 0.72 (blood vessel), 0.44 (tropical tree), 0.61 (coral), and 0.63 (arabidopsis). Such positive correlation is, indeed, expected under surface minimisation, and yet should not occur in the Steiner framework, where the angles are always $2\pi/3$.

Next, we examine whether this positive dependence between $\Omega_{1 \rightarrow 2}$ and ρ is smooth or not. The latter would mark a sharp transition from branching ($\Omega_{1 \rightarrow 2} > 0$), *à la* Steiner, to sprouting ($\Omega_{1 \rightarrow 2} = 0$), a state in which a main thick branch exerts a perpendicular sprout. To test this, we divide the fitting range of ρ into two regions, $\rho \leq \rho_{\text{th}}$ and $\rho > \rho_{\text{th}}$. We then perform a linear fit for each of these regions independently and plot the 95% confidence interval for the resulting fit (Fig. 18, shaded). We denote the slopes at the upper bounds of these two 95% confidence intervals, the one on the right and the one to the left of ρ_{th} , by s_{R} and s_{L} , respectively (dashed lines). Finally, we select ρ_{th} to




Figure 17: **Three-dimensional steering angle** $\Omega_{1 \rightarrow 2}$. It is related to the branching angle by $\Omega_{1 \rightarrow 2} = 4\pi \sin^2[(\pi - \theta)/4]$.




Figure 18: **Sprouting versus branching bifurcations.** The branching angle $\Omega_{1 \rightarrow 2}$ vs. the thickness ratio ρ as obtained from our (a)–(f) six empirical networks (grey circles). We fit the data into two linear fits (dashed lines), one to the left of ρ_{th} (vertical grey line) and one to its right. The 95% confidence intervals for these fits are also shown (shaded area). The existence of two slopes suggests that the sprouting transition is of un-smooth nature, following the manifold pattern shown in Fig. 19 (dashed lines).

	Human neuron	Fruit fly neuron	Blood vessel	Tropical tree	Coral	Arabidopsis
Δs	1.4×10^{-1}	3.5×10^{-2}	2.3×10^{-1}	2×10^{-3}	1.3×10^{-1}	1.7×10^{-1}

Table 4: **Slope hike at ρ_{th} .** The jump Δs in slope from left to right of the threshold ρ_{th} as obtained from the plots in Fig. 18. The positive Δs suggests a sharp rather than a smooth transition.

maximise $\Delta s = s_R - s_L$, capturing the difference between the obtained slopes. We thus seek the optimal threshold ρ_{th} that best divides the data into two sharply split regions. We observe that in all of our networks we have $\Delta s > 0$ (Table 4). This suggests that the $\Omega_{1 \rightarrow 2}$ vs. ρ plots are best described by two separate linear fits that change sharply at ρ_{th} , *i.e.* an abrupt branching to sprouting transition.

At first glance, it may seem surprising that the transition from sprouting to branching is of singular nature. Indeed, the thickness ratio ρ changes smoothly, and one would, therefore, expect the angle $\Omega_{1 \rightarrow 2}$ to also exhibit a gradual increase. Such a continuous prediction is, indeed, observed, under a plain surface-geometric construction, where the links are considered to be perfect cylinders [33, 34]. There, $\Omega_{1 \rightarrow 2}$ changes smoothly and gradually as the thickness gap increases (Fig. 19). The reason

is that it becomes favourable to keep the thick cylinders short, by decreasing $\Omega_{1 \rightarrow 2}$, at the expense of elongating the thin, and hence cheaper, branch. Only in the limit $w'/w \ll 1$ does the branching approach a right angle, when the thick links become infinitely more costly as compared to the thinner sprout.

The crucial point is that the cylindrical construction discussed above violates the smoothness constraint of our manifold solution, which requires a continuous and differentiable patching at the link intersection. It therefore only regards the relative length of the links, but ignores the morphology of the intersection itself. Under our smooth manifold construction the link intersection induces morphological constraints that requires the sleeves to stretch and bend in order to sustain a smooth patching. This comes at a price of additional surface area. Avoiding this added cost can be naturally achieved by converging to $\Omega_{1 \rightarrow 2} = 0$. There, the two *main* sleeves 1 and 2 behave as one continuous tube, and consequently their patching comes at no cost. Hence, adding even a small dent, $\Omega_{1 \rightarrow 2} \neq 0$, results in additional costs, which under $\rho \leq \rho_{\text{th}}$ render the sprouting configuration favourable. As a result, the smooth convergence to sprouting observed for cylindrical links is replaced by the sharp transition predicted under our surface-minimised manifold construction. This represents a qualitative distinction, empirically observable, by which to examine our manifold-based modelling.

To examine this, in Fig. 19 we compare the results obtained through our manifold optimisation (symbols) and those of the plain cylindrical construction (grey solid line) proposed by Zamir [33, 34]. The sharp transition, observed also in the empirical data (Fig. 18), emerges as a clear fingerprint of the smooth manifold optimisation. Interestingly, in the original paper, Zamir states that sprouting *was so frequent that one must... retain the distinction between sprouting and branching for the purpose of qualitative discussions...* [33, 34]. He then proceeded to argue that the observed ubiquity of sprouting is, in fact, an extreme manifestation of branching, occurring only in the limit when the third branch is very small, *i.e.* $\rho \rightarrow 0$. Our framework, in contrast, predicts that sprouting is favoured all the way through to $\rho \sim 0.6$, offering a natural explanation for its, previously puzzling, prevalence (Fig. 20).




Figure 19: **Smooth vs. sharp sprouting transition.** A plain geometric approach [33, 34] yields a gradual change in $\Omega_{1 \rightarrow 2}$ (grey solid line). Our manifold construction, in contrast, predicts a sharp transition (dashed line), as indeed confirmed by our simulation (circles). The Steiner solution predicts $\Omega_{1 \rightarrow 2} = 4\pi \sin^2 \left[(\pi - 2\pi/3) / 4 \right] \approx 0.268\pi$, independently of ρ (grey circle).




Figure 20: **Prevalence of sprout-like bifurcations.** (a)–(f) When $\rho < \rho_{th}$, almost all bifurcations are sprout, or close to sprout, $P(\Omega < 0.5\pi) \approx 1$. This fraction decreases and drops below 1 only after $\rho > \rho_{th}$.

8 Empirical analysis of link lengths and loops

Our analysis indicates that real physical networks support structures that violate Steiner’s principle of minimising total link length. To evaluate the impact of these deviations, we compare the observed total length L of all links against the Steiner-predicted length L_S . We then compute the ratio $\eta = L/L_S$ to quantify the excess length incurred by the network due to its suboptimal structure.

Given the NP-hard nature of the Steiner problem, we cannot obtain the optimal solution for a large network. To bypass this need, we focus on specific components, such as the double bifurcation motif (Fig. 21). We extracted all double-bifurcation motifs within our database and considered their four terminals as fixed points for Steiner optimisation, allowing us to measure L_S . Comparing that with the actual length L , as obtained from the skeletonised motif, we obtained η . On average, we find that $\eta \sim 1.25 \pm 0.05$, capturing a $\sim 25\%$ excess length. This result is quite consistent across all six datasets (Table 5). Note, that this 25% discrepancy is obtained locally, per motif, and therefore captures a lower-bound estimation for the global excess length in real networks.




Figure 21: **Steiner graph versus real curved motif.** (a) Real total length of a human neuron motif. (b) Steiner’s prediction of the network motif’s total length is lower.

	Human neuron	Blood vessel	Fruit fly neuron	Tropical tree	Coral	Arabidopsis
η	1.20	1.33	1.25	1.30	1.35	1.32

Table 5: **Excess length for the double bifurcation motif.** $\eta = L/L_S$ is consistently around 1.25.

The excess length arises from several contributing factors, as illustrated in Fig. 21. The first is the positioning of the intermediate nodes: in the Steiner solution, these nodes are located differently from their counterparts in the real data. This discrepancy results in a shorter total length for the external links (red/yellow) but a longer intermediate segment (green). Additionally, unlike the idealised Steiner paths, real links curve to navigate around obstacles such as neighbouring neurons, glial cells, and blood vessels (not shown in the figure). Combined, these factors lead to an increase in total length from $12.1 \mu\text{m}$ in the Steiner configuration to $14.3 \mu\text{m}$ in the real motif—an 18% increase.

Loops. Both Steiner and surface minimisation predict tree graphs without loops. This prediction is corroborated by our database, as individual neurons, corals, blood vessels, and all other examined networks are indeed trees. It is important to emphasise, however, that in many complex systems—such as in the brain beyond individual neurons—loops do occur in significant numbers [35]. Such loops must have functional roots, as they are suboptimal from the standpoints of surface or length minimisation.

9 Empirical analysis of node morphology

Our manifold-based modelling is grounded in the premise that physical networks must maintain a smooth and continuous morphology, and that they tend to minimise surface area. These constraints govern the way links intersect and are therefore most prominently reflected in the geometric structure of the nodes. To examine this, we ask whether the nodes in real-world physical networks are indeed smooth—free from sharp cusps or morphological singularities—and whether they truly adopt surface-minimising geometries. Investigating these questions requires high-resolution volumetric imaging, which was available for our blood vessel and coral datasets. Accordingly, we focused our analysis on these two systems.




Figure 22: **Gaussian curvature of surface.** We examine the smoothness of the surfaces by calculating the Gaussian curvature κ for blood vessel walls. The surfaces prove to be smooth and nearly flat, with most points having a small curvature $|\kappa| \ll 1$ (blue dots).

9.1 Smoothness

We quantify smoothness by measuring the Gaussian curvature κ at each point, normalised on a scale from zero to one (Fig. 22). As predicted, we find that, within the limits of experimental resolution, both datasets exhibit smooth, manifold-like surfaces.

9.2 Surface optimisation

Once smoothness has been established, we can use our manifold-based formulation to compare our predicted surface minimising morphologies directly with the ones observed in the volumetric data. We focus on trifurcation motifs, a feature that cannot be accounted for by the purely geometric framework of Murray, Zamir, Cherniak, and followers [33, 34, 36]; see *e.g.*, Fig. 19.

Under the manifold constraint, based on the analysis of Sec. 4, an $M = 4$ trifurcation is characterised by a single free complex parameter ($M - 3 = 1$). Setting this parameter generates distinct Jenkins–Strebel quadratic differentials, corresponding to *different* node geometries. This is illustrated in Fig. 9f, where the trifurcation junction is represented by six *stitch lines*, $a, \dots, f$, which together partition the manifold into four charts. The minimum circumference constraint imposes the following conditions:

$$\begin{aligned}
a + b + f &= w_1 = w, & a + c + e &= w_2 = w, \\
b + c + d &= w_3 = w, & d + e + f &= w_4 = w,
\end{aligned} \tag{48}$$

subject to all parameters $a, \dots, f \geq 0$. The solution to these four equations provides

$$a = d, \quad b = e, \quad c = f = w - a - b. \tag{49}$$

We thus have two free parameters, a and b , which together represent a *single* complex parameter. Varying the values of a and b determines how the four sleeves are joined, thereby generating the full space of possible configurations. These configurations—each *a priori* admissible—constitute the complete set of trifurcation geometries, as dictated by string field theory [20]. In Fig. 23, we visualise this configuration space using barycentric coordinates, $(a, b, c = w - a - b)$. For each combination of the triplet (a, b, c) , we construct the corresponding trifurcation manifold and compute its surface area (green-shaded region). Due to the tetrahedral symmetry of the trifurcation, any permutation of (a, b, c) yields the same manifold geometry. We therefore limit our analysis to one-sixth of the full configuration space by imposing the condition $a \leq b \leq c$. Furthermore, we restrict our attention to configurations satisfying $a + b > w/2$, excluding the grey-shaded regions. These regions correspond to conditions under which the manifold is expected to revert to a double-bifurcation structure [20], and are, therefore, considered *un-physical* in string field theory.

We find that the minimum surface area is obtained under the symmetric constructions, where $a = b = c$ (triangle centre, dark green), a morphology presented in Fig. 23b. Eccentric morphologies, where *e.g.*, $a \rightarrow 0$, observed at the corners of the inner triangle, have higher surface area, and hence, according to our hypothesis, are avoided (Fig. 23c).

This offers a testable prediction, that trifurcations should mostly condense around the centre of the (a, b, c) triangle in Fig. 23a. To test this prediction in real data, we collected all trifurcation motifs in the blood vessel and evaluated their barycentric coordinates. In the empirical data, we cannot assume a uniform circumference, and hence we generalise Eq. (48) to treat the case of four different circumferences w_1, w_2, w_3 , and w_4 . This leads to

$$\begin{aligned}
\tilde{a} + \tilde{b} + \tilde{f} &= \bar{w}, & \tilde{a} + \tilde{c} + \tilde{e} &= \bar{w}, \\
\tilde{b} + \tilde{c} + \tilde{d} &= \bar{w}, & \tilde{d} + \tilde{e} + \tilde{f} &= \bar{w},
\end{aligned} \tag{50}$$

where $\bar{w} = (w_1 + w_2 + w_3 + w_4)/4$ is the arithmetic mean of the four circumferences, and




Figure 23: Trifurcation node geometry. The trifurcation manifold is characterised by three barycentric parameters a , b , and c , with the constraint $a + b + c = w$. This induces a triangular space of potential manifold geometries (centre). Each point in this triangle represents a specific (a, b, c) trio, whose manifold surface is shown on a scale of dark green (small surface) to white (large surface). **(a)** The predicted manifold for $a = b = w/4$, $c = w/2$. **(b)** The surface-minimised manifold for $a = b = c = w/3$, representing the symmetric solution at the triangle centre. **(c)** A distorted manifold with $a = 0$, $b = c = w/2$ represents the least optimal solution. Our surface minimisation hypothesis predicts that most real-world trifurcations will condense around the triangle centre, capturing the lowest surface area. **(d)** A real trifurcation morphology observed in the blood vessel dataset. Here the geometry is characterised by the barycentric coordinates $(\tilde{a}, \tilde{b}, \tilde{c})$ in (51). **(e)** Statistical distribution of the empirically observed trifurcation morphologies in the blood vessel dataset. The majority of trifurcation motifs cluster around the centre of the barycentric space, indicating a preference for symmetric branching, as predicted by surface minimisation. **(f)** Another example of complex morphology. While its non-trivial morphology could map to an even higher-degree node, we consider its reduced trifurcation geometry in our analysis.

$$\begin{aligned}
\tilde{a} &= a + (-w_1 - w_2 + w_3 + w_4) / 4, & \tilde{b} &= b + (-w_1 + w_2 - w_3 + w_4) / 4, \\
\tilde{c} &= c + (+w_1 - w_2 - w_3 + w_4) / 4, & \tilde{d} &= d + (+w_1 + w_2 - w_3 - w_4) / 4, \\
\tilde{e} &= e + (+w_1 - w_2 + w_3 - w_4) / 4, & \tilde{f} &= f + (-w_1 + w_2 + w_3 - w_4) / 4,
\end{aligned} \tag{51}$$

are the transformed barycentric coordinates.

Next, we extracted all trifurcation motifs in our data, yielding 26 distinct trifurcations from the blood vessel dataset. For each trifurcation motif, we identified four crossing points on the morphological surface between the six stitch lines $\tilde{a}\text{--}\tilde{f}$. These crossing points correspond to the first-order zeros of the associated meromorphic quadratic-differential function $q(z)$ (as illustrated in Fig. 9).

To locate these crossing points in practice, we used the tetrahedral geometry to identify the intersections between sleeves. First, we define the tetrahedron centre o as the trifurcation point in the skeletal graph. From o , we draw vectors normal to the faces of the tetrahedron. The intersections between these vectors and the surrounding sleeves approximate the sleeve crossing points. In cases of *extrinsic trifurcations*, as described in Sec. 7, we define the skeleton centre o as the midpoint between the two adjacent bifurcation nodes in the skeleton.

Once the crossing points were established, the stitch lines $\tilde{a}\text{--}\tilde{f}$ were traced by identifying the shortest paths connecting these points around each sleeve. The lengths of these paths, $w_1, \dots, w_4$, were then used to derive the barycentric coordinates $\tilde{a}\text{--}\tilde{f}$ via Eq. (51). As before, we excluded configurations in which $\tilde{a} + \tilde{b} < \bar{w}/2$, thereby focusing our analysis on the inner triangular region, consistent with Fig. 23a–c.

In Fig. 23d–f, we present representative examples of real trifurcations from the blood vessel dataset. The reconstructed stitch lines are shown in blue. For each trifurcation, we also indicate its corresponding location in the $(\tilde{a}, \tilde{b}, \tilde{c})$ barycentric triangle (Fig. 23e). As predicted, a clear pattern emerges: the trifurcation motifs tend to cluster near the centre of the triangle, reflecting a preference for symmetric morphologies and an avoidance of highly asymmetric configurations, which are located near the triangle corners (e.g., $\tilde{a} \approx 0$).

These results show that (i) physical networks can, indeed, be modelled as smooth manifolds; and (ii) that these manifolds truly favour surface minimising morphological patterns.

10 Surface versus volume minimisation

Physical networks embedded in the three-dimensional Euclidean space can be conceptualised as two-dimensional manifolds $\mathcal{M}(\mathcal{G})$ subject to surface minimisation, or as three-dimensional objects subject to volume optimisation. We hypothesise that there is a correlation between the volume and surface optimisation problems, such that the two likely yield similar optimal configurations.

To explore this, in the main paper, we show that for a double-bifurcation motif with fixed terminals and $\chi > 0.83$ (within the trifurcation regime, see Fig. 13), surface minimisation drives the intermediate link length $\lambda = l_{\text{int}}/w \rightarrow 0$, suggesting that any double bifurcation with a non-zero intermediate link length is sub-optimal compared to a trifurcation $\lambda = 0$. This is confirmed in Fig. 24a, where we impose a positive λ for the double bifurcation, observing that the surface area indeed increases as λ shifts away from zero. Notably, we find that the three-dimensional volume exhibits a similar increase with λ (Fig. 24b), suggesting that reducing λ (until it reaches zero) will simultaneously minimise both the surface and volume of the manifold. We also observe similar relationships in multiple parameter regimes. This observation suggests a strong correlation between surface and volume minimisation, but it does not exclude the possibility that volume minimisation could help us discover novel patterns and morphologies in some parameter regimes, that could further enhance our understanding of physical networks.




Figure 24: **Surface versus volume minimisation.** (a) The surface area of a double bifurcation motif vs. the intermediate link length λ . The area decreases as $\lambda \rightarrow 0$, signalling the preference for a trifurcation. (b) Volume vs. λ displays the same behaviour, indicating that volume minimisation is qualitatively similar to surface minimisation. Here, the volume exhibits numerical fluctuations (a 5-period moving average has been applied to the data points), but still shows a clear trend that approaches minimum when $\lambda = 0$, *i.e.* at the trifurcation.

References

- [1] Shapson-Coe, A. *et al.* A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution. *Science* **384**, eadk4858 (2024).
- [2] Scheffer, L. K. *et al.* A connectome and analysis of the adult *Drosophila* central brain. *eLife* **9**, e57443 (2020).
- [3] Wilson, N. M., Ortiz, A. K. & Johnson, A. B. The Vascular Model Repository: A Public Resource of Medical Imaging Data and Blood Flow Simulation Results. *Journal of Medical Devices* **7** (2013).
- [4] Gonzalez de Tanago, J. *et al.* Estimation of above-ground biomass of large tropical trees with terrestrial LiDAR. *Methods in Ecology and Evolution* **9**, 223–234 (2018).
- [5] Corals — 3D Digitization. <https://3d.si.edu/corals>.
- [6] Pan, H., Hétroy-Wheeler, F., Charlaix, J. & Coliaux, D. ARABIDOPSIS 3D+T dataset (2021).
- [7] Saha, P. K., Borgefors, G. & di Baja, G. S. *Skeletonization: Theory, Methods and Applications* (Academic Press, 2017).
- [8] Silversmith, W., Bae, J. A., Li, P. H. & Wilson, A. Seung-lab/kimimaro: Zenodo Release v1 (3.0.0) (2021).
- [9] Berger, D. R., Seung, H. S. & Lichtman, J. W. VAST (Volume Annotation and Segmentation Tool): Efficient Manual and Semi-Automatic Labeling of Large 3D Image Stacks. *Frontiers in Neural Circuits* **12** (2018).
- [10] Zhao, T. & Plaza, S. M. Automatic Neuron Type Identification by Neurite Localization in the *Drosophila* Medulla (2014). [1409.1892](#).
- [11] Raunonen, P. *et al.* Fast Automatic Precision Tree Models from Terrestrial Laser Scanner Data. *Remote Sensing* **5**, 491–520 (2013).
- [12] Dawson-Haggerty *et al.* Trimesh (2019).
- [13] Au, O. K.-C., Tai, C.-L., Chu, H.-K., Cohen-Or, D. & Lee, T.-Y. Skeleton extraction by mesh contraction. *ACM Trans. Graph.* **27**, 1–10 (2008).
- [14] Veltkamp, R. C. & Wesselink, W. Modeling 3D Curves of Minimal Energy. *Computer Graphics Forum* **14**, 97–110 (1995).

- [15] Hwang, F. K., Richards, D. S. & Winter, P. *The Steiner Tree Problem* (Elsevier, Amsterdam, The Netherlands, 1992), 1st edn.
- [16] Bobenko, A. I., Sullivan, J. M., Schröder, P. & Ziegler, G. M. (eds.) *Discrete Differential Geometry*, vol. 38 of *Oberwolfach Seminars* (Birkhäuser, Basel, Switzerland, 2008), 1st edn.
- [17] Tong, D. Lectures on String Theory. <http://www.damtp.cam.ac.uk/user/tong/string.html> (2009).
- [18] Jin, M., Gu, X., He, Y. & Wang, Y. *Conformal Geometry: Computational Algorithms and Engineering Applications* (Springer, 2018).
- [19] Carlip, S. Quadratic differentials and closed string vertices. *Physics Letters B* **214**, 187–192 (1988).
- [20] Saadi, M. & Zwiebach, B. Closed string field theory from polyhedra. *Annals of Physics* **192**, 213–227 (1989).
- [21] Strebel, K. *Quadratic Differentials* (Springer, Heidelberg, Germany, 1984), 1st edn.
- [22] Gardiner, F. P. *Teichmüller Theory and Quadratic Differentials* (Wiley, New York, USA, 1987), 1st edn.
- [23] Moeller, N. Closed Bosonic String Field Theory At Quartic Order. *Journal of High Energy Physics* **2004**, 018–018 (2004).
- [24] Moeller, N. Closed bosonic string field theory at quintic order: Five-tachyon contact term and dilaton theorem. *Journal of High Energy Physics* **2007**, 043–043 (2007).
- [25] Moeller, N. Closed bosonic string field theory at quintic order II: Marginal deformations and effective potential. *Journal of High Energy Physics* **2007**, 118–118 (2007).
- [26] Witten, E. Non-commutative geometry and string field theory. *Nuclear Physics B* **268**, 253–294 (1986).
- [27] Headrick, M. & Zwiebach, B. Convex programs for minimal-area problems. *arXiv:1806.00449 [hep-th]* (2019). [1806.00449](#).
- [28] Peng, C.-H., Jiang, C., Wonka, P. & Pottmann, H. Checkerboard patterns with black rectangles. *ACM Transactions on Graphics* **38**, 171:1–171:13 (2019).
- [29] Jiang, C., Wang, C., Rist, F., Wallner, J. & Pottmann, H. Quad-mesh based isometric mappings and developable surfaces. *ACM Transactions on Graphics* **39**, 128:128:1–128:128:13 (2020).

- [30] Jiang, C. *et al.* Using isometries for computational design and fabrication. *ACM Transactions on Graphics* **40**, 42:1–42:12 (2021).
- [31] Miyawaki, S., Tawhai, M. H., Hoffman, E. A., Wenzel, S. E. & Lin, C.-L. Automatic construction of subject-specific human airway geometry including trifurcations based on a CT-segmented airway skeleton and surface. *Biomechanics and Modeling in Mechanobiology* **16**, 583–596 (2017).
- [32] Navarro-Verdugo, A. L., Goycoolea, F. M., Romero-Meléndez, G., Higuera-Ciapara, I. & Argüelles-Monal, W. A modified Boltzmann sigmoidal model for the phase transition of smart gels. *Soft Matter* **7**, 5847–5853 (2011).
- [33] Zamir, M. Optimality principles in arterial branching. *Journal of Theoretical Biology* **62**, 227–251 (1976).
- [34] Zamir, M. Nonsymmetrical bifurcations in arterial branching. *Journal of General Physiology* **72**, 837–845 (1978).
- [35] Liu, Y., Dehmamy, N. & Barabási, A.-L. Isotopy and energy of physical networks. *Nature Physics* **17**, 216–222 (2021).
- [36] Cherniak, C. Local optimization of neuron arbors. *Biological Cybernetics* **66**, 503–510 (1992).