

Mapping hyperbolic order in curved materials - Supporting information

Martin Cramer Pedersen,^{*a} Stephen T. Hyde,^b Stuart Ramsden,^c and Jacob J. K. Kirkensgaard^{a,d}

Supporting Information

SI.1 Curvature of (minimal) surfaces and TP(M)S

A point on a curved surface S has maximal and minimal principal curvatures, κ_1 and κ_2 , inversely related to the principal radii of curvature. The mean (extrinsic) curvature at that point, H , as the average of those curvatures, whereas the Gaussian (intrinsic) curvature, K , is their product:

$$H = \frac{\kappa_1 + \kappa_2}{2} \quad K = \kappa_1 \kappa_2$$

Famously, the Euler characteristic for a surface, $\chi(S)$, is related to the Gaussian curvature via the Gauss-Bonnet theorem:

$$\int \int_S K da = 2\pi\chi(S)$$

where da is the area element of S . Minimal surfaces, including TPMS, have vanishing mean curvature for all points, so that their principal curvatures are everywhere equal and opposite at all points, although they vary from point to point. Since the Gaussian curvature is everywhere non-positive,

$$\int \int_S K da < 0$$

so χ must also be negative. In particular, if the Gaussian curvature is integrated over a unit cell of the surface, bounded by translationally equivalent points (i.e. periodic boundary conditions), $\chi(S)$ is a negative integer, equal to $(2 - 2g)$, where $g > 1$ is the surface genus per unit cell. Any deformation of the TPMS, whether by non-uniform stretching, bending or shrinking - provided it does not tear or glue the surface - conserves its integral curvature (and hence its surface topology, indexed by $\chi(S)$). Therefore any such deformation within a fixed unit cell conserves g and $\chi(S)$. If the area of the surface confined to that unit cell is $A(S)$, we can define a surface-averaged Gaussian curvature, $\langle K \rangle = \frac{1}{A(S)} \int \int_S K da$. Since $\chi(S)$ is conserved, $\langle K \rangle$ must be negative for all allowed deformations of the TPMS, regardless of its mean curvature. Any surface which is an admissible deform of a TPMS is, by definition, a TP(M)S. Therefore, all TP(M)S, like their related TPMS, have, on average, negative Gaussian curvature $\langle K \rangle$.

SI.2 The Poincaré model

In this model, each point in \mathbb{H}^2 is identified with a point in the open unit disc, which can be parametrised with coordinates z in the complex plane, i.e. $z \in \mathcal{D} \subset \mathbb{C}$, equipped with the metric:

$$ds^2 = \frac{dz^2}{(1 - |z|^2)^2}$$

This map is a conformal model of \mathbb{H}^2 : all angles in \mathbb{H}^2 map to equal angles in the disc. Geodesic lines in \mathbb{H}^2 are readily identified as circular arcs in \mathcal{D} which meet its boundary (the unit circle) orthogonally. The origin of \mathcal{D} is not special: any other point a can be shifted to that origin by a map f , which remaps the entire disc - $f: \mathcal{D} \rightarrow \mathcal{D}$ - by an oriented isometry, where:

$$f(z) = e^{i\theta} \frac{z - a}{a^*z - 1} \quad (\text{SI.1})$$

is a Möbius transformation and $\theta \in [0; 2\pi[$. The asterisk, $*$, denotes complex conjugation.

^{*} Corresponding author. E-mail: mcpe@nbi.ku.dk

^a Niels Bohr Institute, University of Copenhagen, Denmark.

^b School of Chemistry, University of Sydney, Australia.

^c National Computational Infrastructure (NCI) Vizlab, Australian National University, Australia.

^d Department of Food Science, University of Copenhagen, Denmark.

SI.3 Two-dimensional orbifolds

Two-dimensional orbifolds were pioneered by Thurston⁴⁶ and named by Conway⁴⁷. Reflection symmetries are labelled by $*$ and roto-reflection sites, where N mirror lines intersect, by $*N$. An order M rotational symmetry site is labelled by its order, M . Two-dimensional isometries also include a line (geodesic) of glide symmetry indicated by \times (whose symbol derives from the topology of an orbifold with a single glide: a “cross-cap”) and a translational symmetry denoted by the symbol \circ (a translation induces a topological “handle” in the orbifold). Thus, the Coxeter pattern, induced by mirrors defining the $\frac{\pi}{2}, \frac{\pi}{4}, \frac{\pi}{6}$ triangulation, has Conway orbifold symbol $*246$, whereas the stellate pattern, with order 2,4,6 rotation centres only, has orbifold symbol 246 . In fact, patterns with symmetry $*246$ are maximally ordered on the Primitive, Diamond, or Gyroid TPMS. Indeed, that ordering is maximal for *any* TPMS embedded in three-dimensional euclidean space (\mathbb{E}^3)⁴⁸. (Although other ordered patterns can be realised in \mathbb{H}^2 , e.g. $*237$; a kaleidoscopic pattern generated from a hyperbolic triangle with vertex angles $\frac{\pi}{2}, \frac{\pi}{3}, \frac{\pi}{7}$), they are “incommensurate” with \mathbb{E}^3 ⁵⁹).

Any ordered pattern on the Primitive, Diamond, or Gyroid TPMS contains a subset of the hyperbolic isometries contained in a pattern with $*246$ ordering: for example, the stellate 246 pattern is an index-two subgroup of $*246$. At index 96 of $*246$, we find the translational subgroup $\circ\circ\circ$ (or \circ^3), giving 131 distinct subgroups in the quotient group labelled $*246/\circ\circ\circ$ enumerated in detail elsewhere^{16,17}. In turn, any one of those 131 groups respects all of the translational symmetries of the surfaces (which do not exchange its front and back sides). Ordering characterised by the Conway symbol $\circ\circ\circ$ (a Toroidal orbifold) includes three distinct translational symmetries only, corresponding to those of the unit cell of the Primitive surface in Fig. 1(b-d). Larger unit cells, such as the conventional unit cells of the Gyroid and Diamond TPMS, may retain only translational symmetries of still larger Toroidal orbifolds $\circ\circ\circ$ patterns, giving subgroups of translational symmetries with orbifold symbols $\circ\dots\circ$, where the number of entries of \circ in the symbol is equal to the surface genus, g .

SI.4 Meshes and disc packings

Meshes and associated aligned disc packings used to generate the images in e.g. Fig. 2 are available online. The meshes are in .off-format, and the disc packings are in our own .h2p-format which is readily imported and manipulated using the supplied Python module, which can be found at:

<https://www.gitlab.com/mcpe/H2Tools>

We stress that the IDs of the points in both files match and hence provide a one-to-one correspondence between the points on the mesh of the TPMS and the points in \mathbb{H}^2 . In short, the files offer a quick way of visualizing patterns confined to the nodal approximations of the Primitive, Diamond, and Gyroid surfaces.

SI.5 Pipeline for refining meshes

SI.5.1 Input

In this section, we outline the process used to map patterns from \mathbb{E}^3 to \mathbb{H}^2 . We start with a point cloud in \mathbb{E}^3 (possibly labelled or colored). The cloud is mapped to a triangulated surface mesh, whose topology is reproduced on some hyperbolic surface (in \mathbb{E}^3) and then on \mathbb{H}^2 . Any hyperbolic surface whose cover is \mathbb{H}^2 is suitable.

To construct the meshes in Fig. 2, Fig. 3, and Fig. 5, the nodal approximation of which is known, one can readily construct dense meshes using the CGAL method `make_periodic_mesh_3`⁵⁶. In the case outlined in Fig. 6, one can follow the pipeline outlined below.




Fig. SI.1 In (a), 25000 colored/labelled points sampled from the volume between two level sets describing a single unit cell of the Primitive surface⁵⁸. In (b), a periodic mesh is refined (and colored/labelled) from the point cloud in (a). The edges of the mesh connecting the vertices near the faces of the unit cell to those of the opposing face are not shown. In (c), by way of Stephenson’s circle packing algorithm, the mesh is quasiconformally visualized in \mathbb{H}^2 . In red, we highlight the shortest “long” edge in the pattern, which we use for aligning the pattern in the next step. In (d), the circle packing is aligned to and the individual circles translated to match the fundamental domain of the Primitive surface, which is shown in Fig. 1. This presentation makes identification of the approximate $*246$ symmetry of the pattern a simple task.

SI.5.2 Meshing

Those stages of the pipeline are demonstrated by the example in Fig. SI.1 and described in turn.

The initial step is the construction of a triangular mesh capturing the overall shape of the initial point cloud. This can largely be handled by CGAL⁵⁷ methods:

1. The point cloud is imported from a given file.
2. Optionally, one can apply a smoothing algorithm using the `bilateral_smooth_point_set` method. This step flattens and smoothens the point cloud in an effort to define the underlying surface better.
3. The surface normals are estimated using `pca_estimate_normals`. As the method name implies, each point is equipped with a normal vector based on the nearest neighbors; these are uniformly oriented to define the two sides of the surfaces.
4. Using these normals and the positions of the points, the surface is approximated using `Poisson_reconstruction_function`. The output is a callable method used in the meshing algorithm in the next step.
5. Finally, the periodic mesh is constructed from the reconstructed function by the recently developed `make_periodic_mesh_3`. We work in the unit cube and prohibit the meshing algorithm from producing edges in the mesh larger than 0.005. In the case where the surface is by e.g. a nodal approximation, one can simply use that expression in this step.
6. The mesh is readily exported in a variety of formats to a given file.

The output of this procedure is a periodic, triangulated mesh confined to a TPMS, whose vertices are decorated by colors/labels inherited from the initial point cloud. The precise coloring algorithm is of secondary importance here. For example, every point in the original cloud in the example in Fig. SI.1(a) has one of two possible colors. The color of a vertex in the triangulation of the Primitive surface in Fig. SI.1(b) corresponds to the dominant color among the nine points in the initial point cloud closest to that vertex location. We stress that several other methods will also produce meaningful results.

SI.5.3 Map to hyperbolic two-space

Next, we employ Stephenson’s circle packing algorithm to map the surface mesh to \mathbb{H}^2 , inputting the (underlying triangulated graph of the) mesh. Since the mean number of edges per vertex in our meshes exceeds six, the resulting circle packing will be inherently hyperbolic. CirclePack gives a coarse-grained discrete quasiconformal map between triangulations of the TPMS and \mathbb{H}^2 via disc packings whose skeletons (defined by edges of discs in mutual contact) define those triangulations. The algorithm ensures that the local topology of the TPMS triangulation is unchanged by the map to \mathbb{H}^2 . Each vertex of the triangulation on the TPMS has a corresponding triangulated vertex in \mathbb{H}^2 and corresponding sites have identical numbers of vertices surrounding those sites in concentric shells.

The algorithm consists of two parts. First, we calculate the radii of all discs in the packing and iteratively decrease the mismatched angles between the discs until a prespecified precision is reached, i.e. when the changes in the computed radii do not change by more than this value from iteration to iteration, the algorithm terminates. The convergence is uniform and guaranteed^{54,55}. Once the radii are calculated, the discs are placed in the Poincaré disc by locating an arbitrary disc at the origin, followed by any one of its neighbors, whose centre is located on the real axis so that both discs are in contact. Next, either one of the pair of discs in the relaxed packing which neighbor both discs is centred above the real axis in the Poincaré disc such that it contacts both discs. The remaining discs in the packing are then located iteratively to maintain the contacts demanded by the original triangulation. For example, we applied that process to the triangulated mesh on the Primitive surface in Fig. SI.1(b), resulting in the disc packing in the Poincaré disc shown in Fig. SI.1(c). Notice that this stage effectively selects an arbitrary origin (the first selected disc) and orientation (its neighbor).

SI.5.4 Alignment

Further analysis of a pattern induced on a surface often requires additional optimisation to match the origin and alignment of the pattern in the Poincaré disc with some pre-specified origin and orientation of the the universal cover of the surface. For example, the green patches in Fig. 1(c) and (d) are necessarily correlated with the symmetries of the unfolded dodecagon defining a unit cell (effectively highlighting the translational symmetry due to collars of the pattern). The pattern in Fig. SI.1(c) can be relocated anywhere in the Poincaré disc by moving all of the points according to a Möbius transformation (Equation (SI.1)). By confining the positions of the points in our pattern to the one of their images in the orbit of the translational group $\circ\circ\circ$ that is inside the pre-specified translational domain of the group, we can align the vertices and thus the pattern to the canonical version of the unit cell. This resulting pattern is shown in Fig. SI.1(d), which admits a simpler interpretation of the symmetries of the pattern in \mathbb{H}^2 .

In practice, the problem can be approached in the following manner. Not all edges in the mesh in Fig. SI.1(b) are realized as contacts between discs in Fig. SI.1(c); notably the “long” edges connecting the opposing sides of the patch. The shortest one of these edges (in the metric on \mathbb{H}^2) must be aligned with (any one of) the shortest of the translations in the fundamental domain for the surface - this particular edge is highlighted in Figures SI.1(c) and (d). After initial alignment, a more precise alignment can be achieved numerically, by e.g. minimization of the overlap between the discs. One follows the following procedure after obtaining the initial pattern in \mathbb{H}^2 as

shown in Fig. SI.1(c): i) fix the translation of the relevant translational domain (as shown in Fig. 1), ii) apply these translations to the pattern in \mathbb{H}^2 to give multiple copies of this pattern with overlaps and gaps, and finally iii) apply Möbius transformations to minimize the aforementioned gaps and overlaps.

SI.6 Simulation details

The simulation results shown in Figures 5 and 6 were gathered from molecular simulations conducted under standard DPD conditions⁶⁴. We simulate for a total of 10^7 steps with a step size of $\Delta t = 0.005$. The interactions of the DPD potentials are set to 25 for beads of similar color and 40 for beads of a different color as well as $\gamma = 4.5$ and an r_{cut} of 1. Additionally, the harmonic bonds are parameterized by a spring constant of $k = 4$ and a rest length $r_0 = 0.88^5$. The simulations are conducted at a temperature of $kT = 1$.

The simulation in Fig. 5(a-d) shows the result of 3:3 polymers in a periodic box with side length $L = 20$. Fig. 5(d-g) stems from a simulation of 5:5 polymers in a box with $L = 22$. Fig. 5(h-k) was generated by simulating a 3:3:3 ABC polymer with $L = 20$. And finally, Fig. 6 shows the results from two simulations of balanced mixtures of 4:4:9 polymers with $L = 18$ in the top row and 4:4:12 polymers with $L = 22$.

We simulate our polymers at a bead density of 3 per unit volume and estimate the volume of the films in Fig. 5 as $\alpha L^2 d$, where α is a constant for approximating the areas of our minimal surfaces published elsewhere⁷ and d is the film thickness, which we set to 2 for our diblock co-polymer examples and 0.5 for the triblock miktoarm polymer example. For the simulations in Fig. 5, we confine our polymers using two “walls” consisting of a dense set of particles equipped with Lennard-Jones potentials with parameters $\sigma = 1$, $\epsilon = 1$, and $r_{\text{cut}} = 2^{1/6} \sigma \approx 1.12$ as prescribed elsewhere⁵. For the purpose of describing these walls mathematically, we use the nodal approximations of our minimal surfaces⁵⁸ and place the normals $\frac{d+\sigma}{2}$ from the nodal approximation of the surface using the surfaces normals.

SI.7 Annotated version of Fig. 2

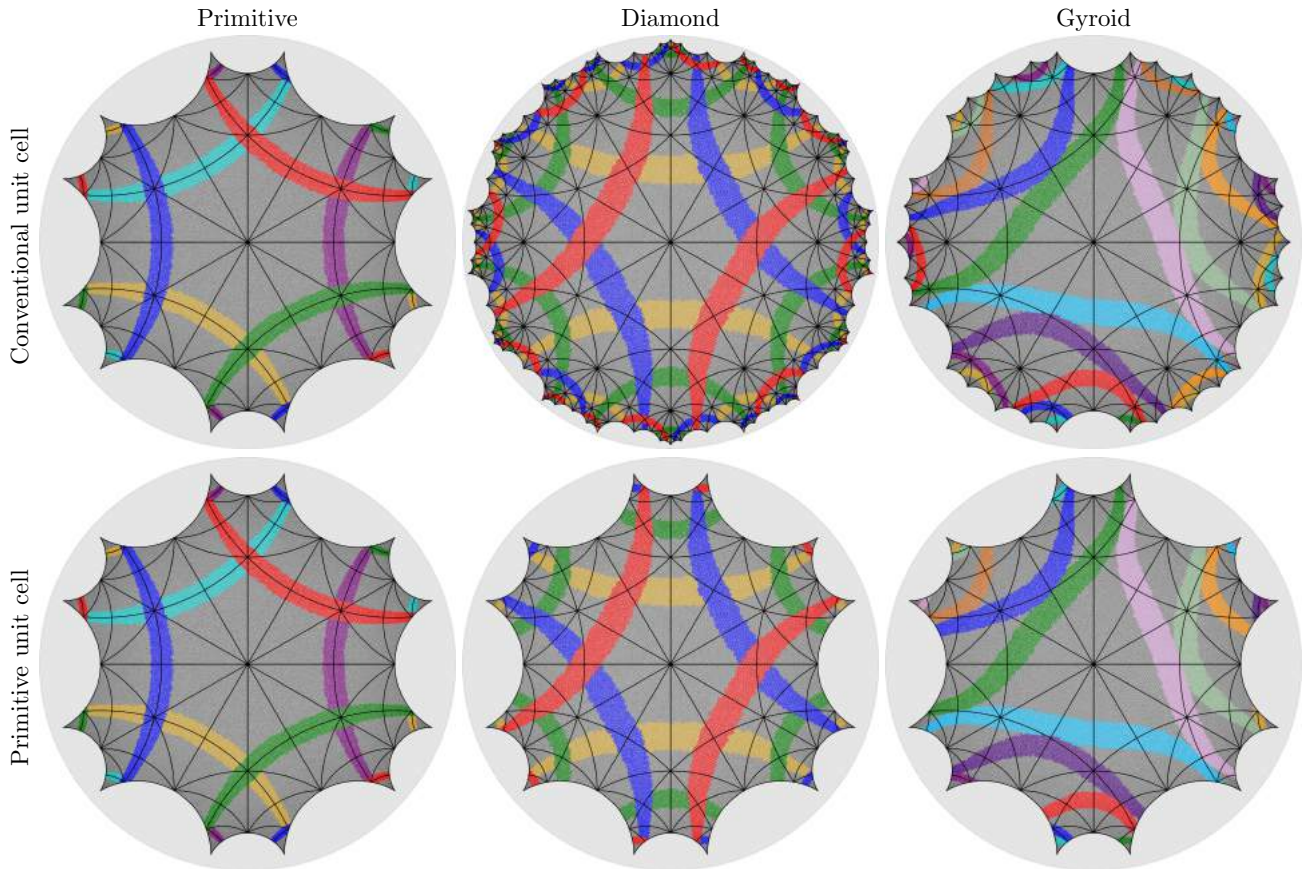


Fig. SI.2 The patterns shown in Fig. 2 with grids of *246 domains superimposed to emphasize the structure of the collars of the respective surfaces.

SI.8 Annotated version of Fig. 3

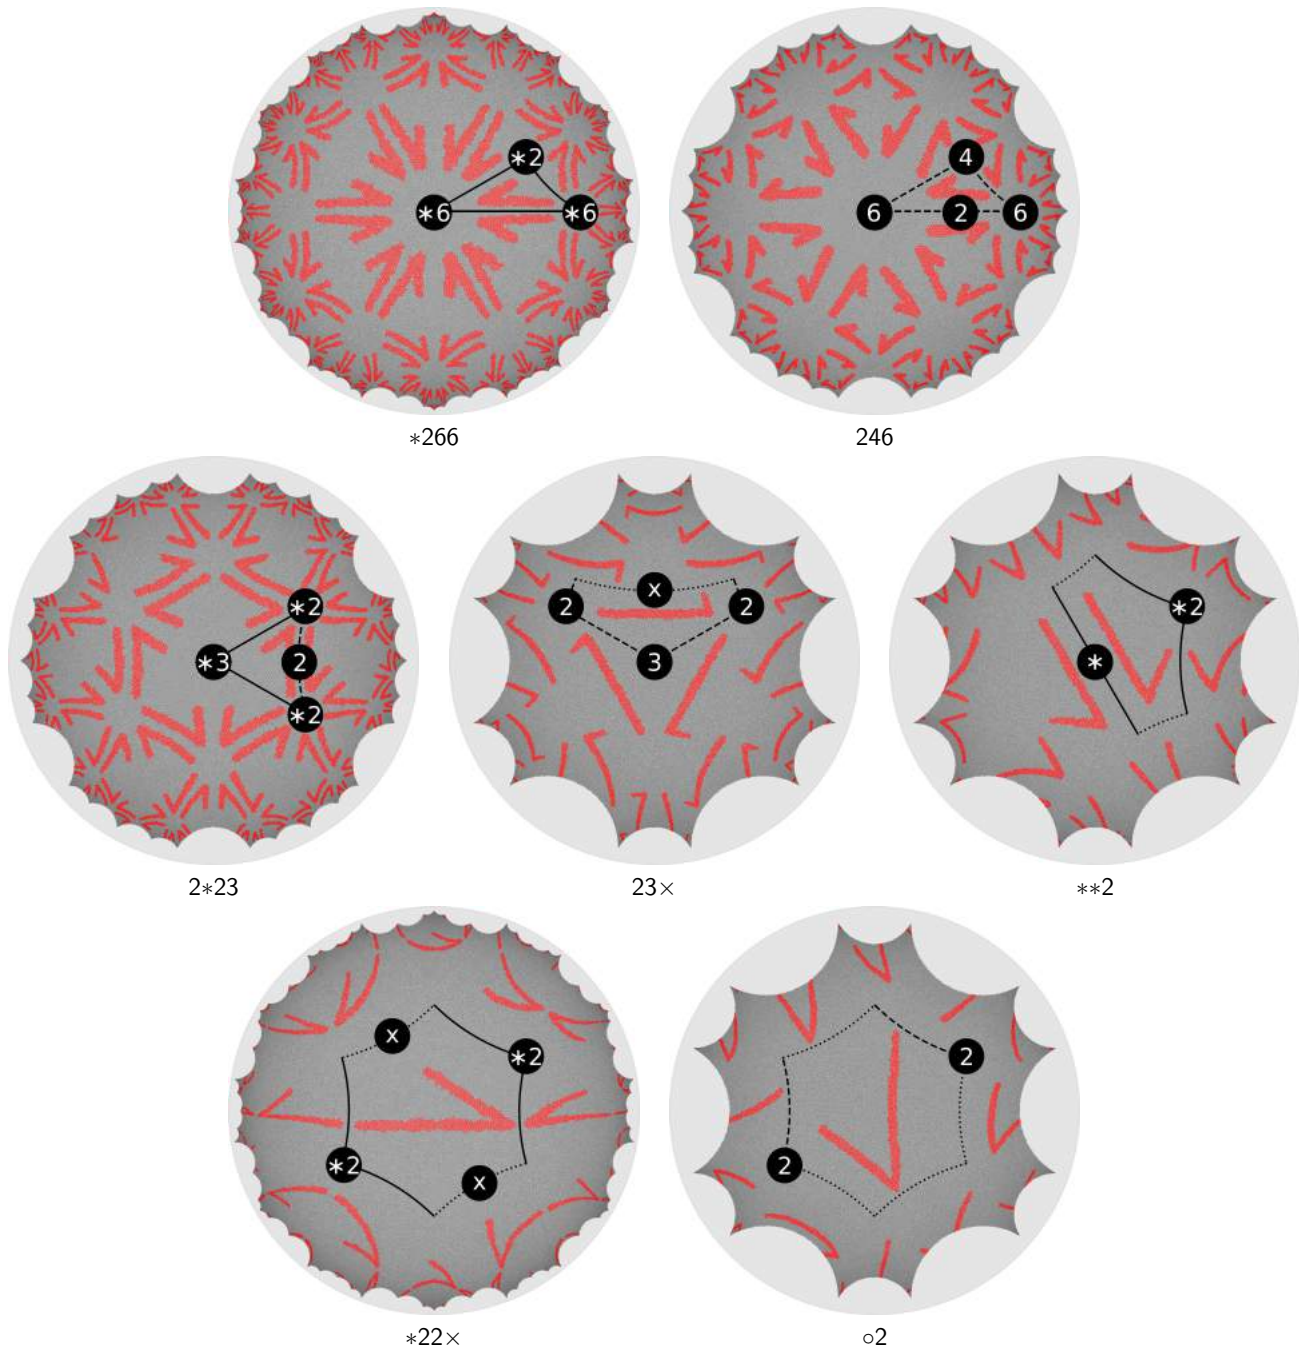


Fig. SI.3 The seven examples of orbifold classes in \mathbb{H}^2 with their respective orbifolds highlighted. The solid lines denote mirror lines, the dashed lines are used to highlight an example of an asymmetric domain for orbifolds with rotational symmetries, and dotted lines are used to denote either identifications (gluings) or crosscap symmetries.

SI.9 Annotated version of Fig. 4

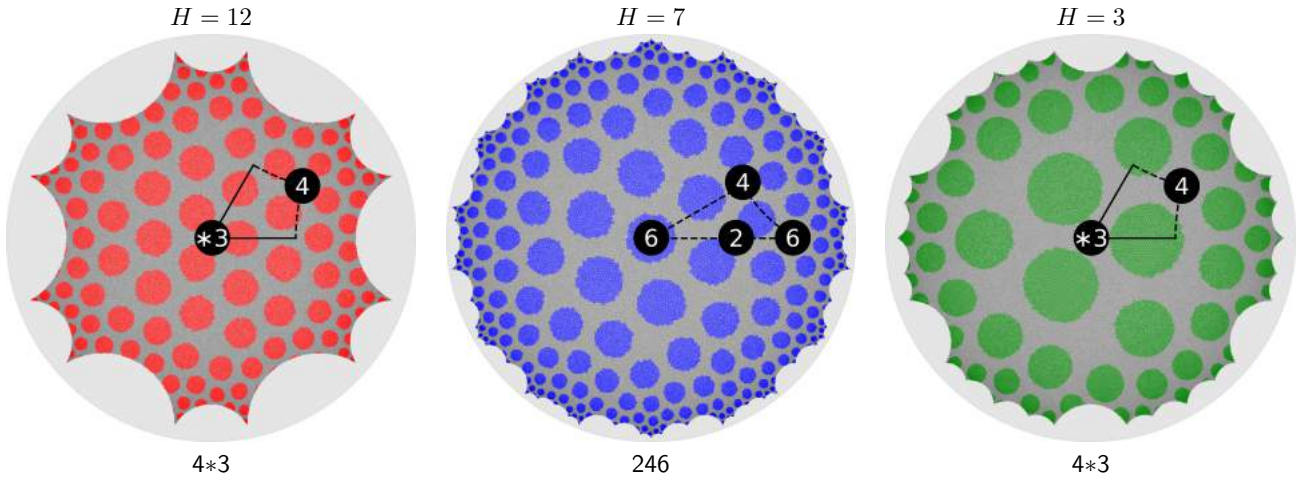


Fig. SI.4 The structures on the Poincaré disc in Fig. 4 annotated with their respective orbifolds highlighted. The notation is the same as in Fig. SI.3.

SI.10 Different choices of center for Fig. 5(c)

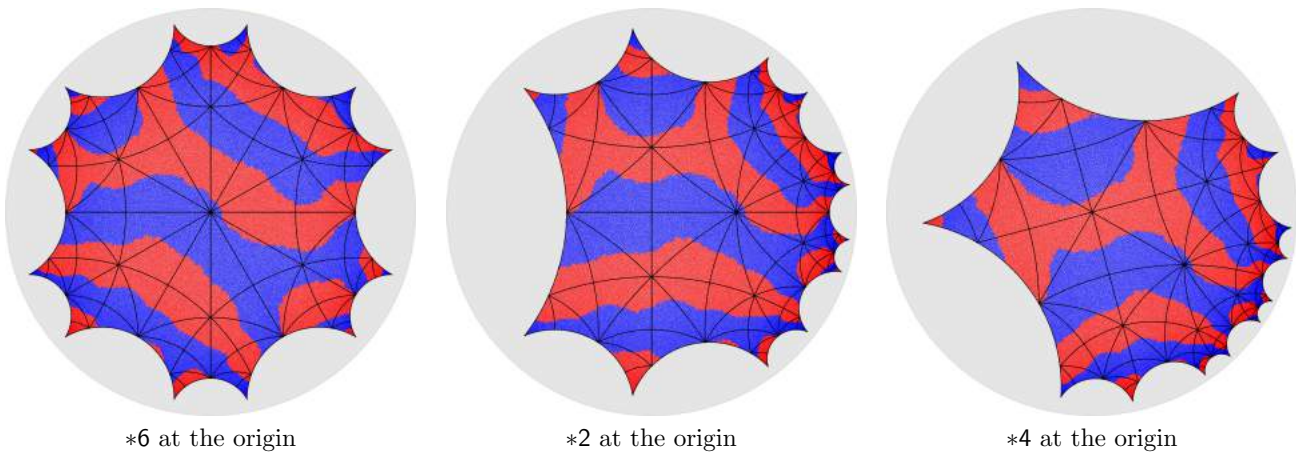


Fig. SI.5 Translated versions of the pattern in Figure 5(c) demonstrating how shifting the origin located at the disc centre by subjecting the mesh to different Möbius transformations, can highlight different features of a pattern. Here, we move each of the corner points of the $*246$ domain to the origin of the Poincaré disc, in turn aiding the identification of mirror lines through these corner points.

SI.11 Alignment of hyperbolic patterns in Fig. 6




Fig. SI.6 “Long” edges used in the alignment process for the patterns in Fig. 6. In (a), the edges encoding three distinct translations, along which the pattern has been aligned, are shown in light grey for the structures in the top row of Fig. 6. In (b), we show the resulting pattern alongside the *246 lattice. In (c) and (d), the same but for the bottom row in Fig. 6.