Chapter 5: The Shape of Tiles

Between a Hard Rock and a Soft Cell

Professor Alain Goriely

Lecture given on the twenty-eighth of March 2026

Mosaic patterns and tilings are ubiquitous in nature, appearing in systems ranging from biological tissues to geological formations. Traditionally, these structures have been modelled using polyhedral tilings composed of flat faces, straight edges, and sharp corners. However, careful observation reveals that many natural tilings deviate significantly from this paradigm: their boundaries are curved with smooth interfaces. This realisation has motivated the introduction of a new class of shapes known as soft cells, which arise as smooth deformations of standard tilings. Such cells are found in the geometry of metal and liquid foams as well as in many micro-structures modelled by triply periodic minimal surfaces. In this lecture, we will explore the mathematics of tilings, regular and irregular, hard and soft, describe their construction and classification, and illustrate how they provide a more accurate geometric description of patterns found in biology, architecture, engineering, in the deepest sea and even in space.

Figure 5.1: Example of a tiling provided by a mosaic in the Alahambra in Granada, Spain made in the 13th Century. Here the tiles are both curved and polygonal. It is an example of a *polyhedric* tiling.

5.1 Tilings and Mosaics

We begin with the notion of a *tiling*, a fundamental concept in geometry. A tiling is a covering of the plane (or, more generally, of space) by a collection of regions, called *cells* or *tiles*, that fit together without overlaps and without gaps. This simple idea appears in many familiar contexts, from everyday actual mosaics (Figure 5.1) to crystalline structures, yet it also leads to deep mathematical questions.

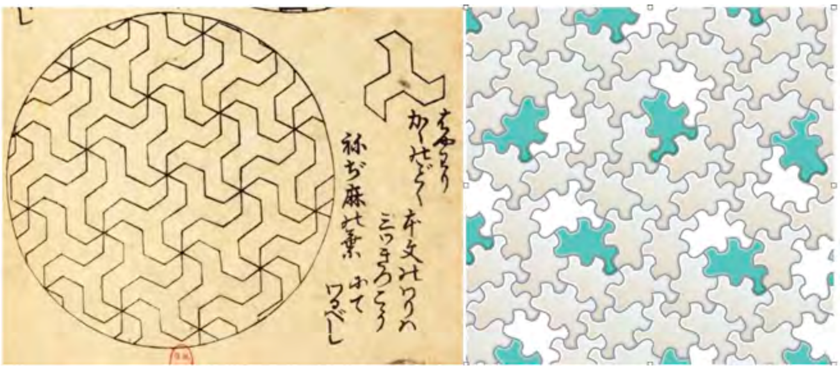

Figure 5.2: Example of tilings made with a single piece. On the left a design by Hokusai with a non-convex polygonal shape (1824) leading to a *monohedral* tiling. Right, the aperiodic “spectre" tiling discovered in 2024 made with a non-convex tile, an example of a *monohedric* tiling.

Depending on the geometry of the pieces, one distinguishes between *polyhedral tilings* where tiles are bounded by straight faces (*polygonal* in the plane, with straight edges), and *polyhedric tilings* in general where the tiles can be made with curved pieces, leading to richer and less intuitive configurations.

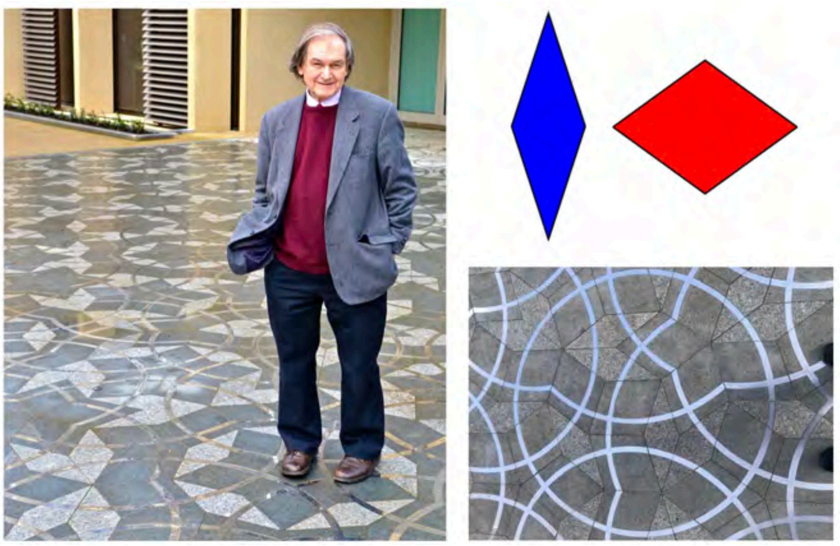

Figure 5.3: Roger Penrose standing on his Penrose tilings, an example of *polyhedral* tiling that he designed himself for the Andrew Wiles Building, the home of the Mathematical Institute in Oxford.

A further classification is whether the tiling is built from a single shape or from several distinct ones. When a single tile is repeated throughout, the tiling is called *monothedric* in general if the edges are curved. If this tile has flat faces, one speaks of a *monohedral* tiling. Particularly striking are monotilings that are *aperiodic*: although the same tiles are used everywhere, the pattern never repeats itself periodically as shown in Figure 5.2. A celebrated example of a polygonal aperiodic tiling is provided by *Penrose tilings*, shown in Figure 5.3. These tilings discovered by Roger Penrose in the 1960s are constructed from just two simple shapes (two rhombi, as shown here). Although these tiles fit together perfectly to cover the plane without gaps or overlaps, they never form a repeating pattern: no finite region can be translated to reproduce the whole tiling.

5.1.1 Polygonal tilings

Let us now focus on polygonal tilings in the plane such as the one in Figure 5.4. Such tilings can be viewed as networks composed of *cells* (the tiles themselves) and *nodes* (also called *vertices*), which are points where edges meet. At each node, several edges, and hence several cells, come together. An important geometric distinction concerns whether a cell is *convex* or *non-convex*. A convex cell has the property that any two points within it can be joined by a straight line that remains entirely inside the cell. In contrast, a non-convex cell contains pairs of points whose connecting segment exits the cell (see Figure 5.5). This distinction has significant consequences for both the geometry and the combinatorics of the tiling as the theory of convex cells is much easier and much more developed.

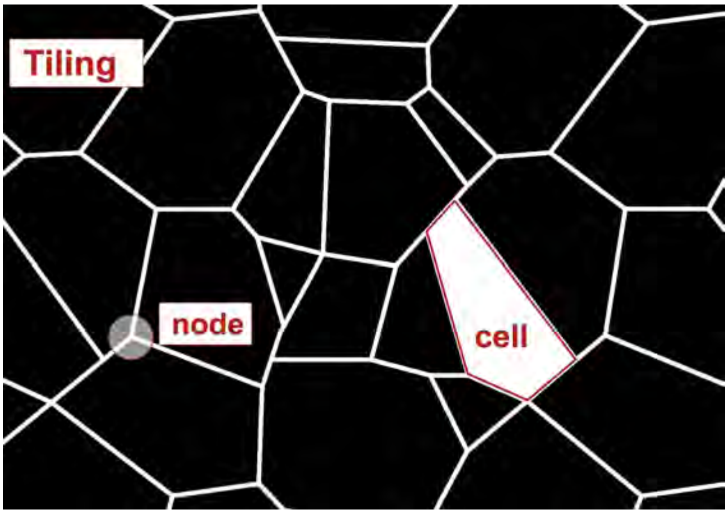

Figure 5.4: An example of polygonal tilings made with different cells with straight edges.

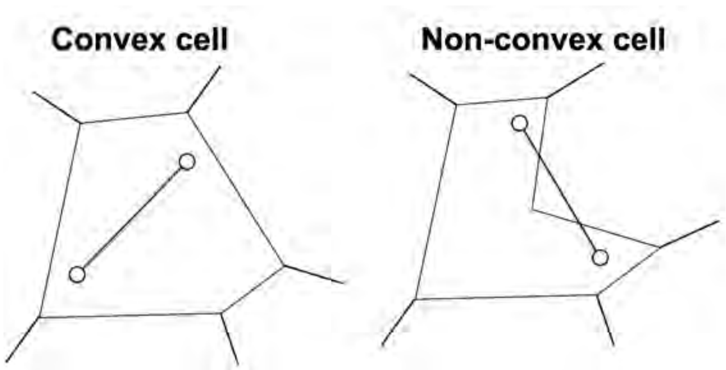

Figure 5.5: In a convex cell, any lines between two points in the cell is also in the cell, unlike non-convex cells.

To analyse a tiling quantitatively, one assigns numerical characteristics to its elements [4]. We assign numbers either to the nodes or to the cells. First, consider a cell, it can be described by the *number of nodes* v ('v' for vertex) along its boundary, or the number of *corners* v^*. The difference is subtle but important [5]. The number of corners is the number seen from inside the cell, if there is a node along an edge, but no corner at that place, then these two numbers will be different as shown in Figure 5.6. For a polygonal tiling the number of corners is also the number of sides of the polygon.

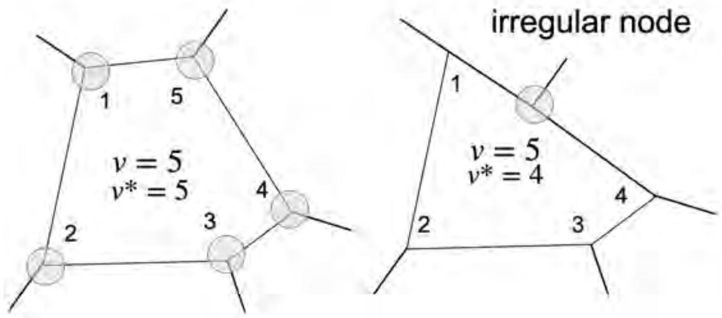

Figure 5.6: For a given cells we can count either the number of nodes v on its boundary, or the number of corners v^*. If one of the node is *irregular* (circled), these two numbers are different.

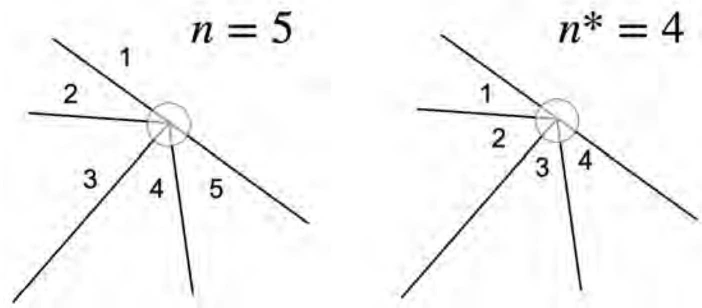

Figure 5.7: For a given node, the degree n is the number of cells meeting at that node, whereas n^* is the number of corners at that node. In this case the node is *irregular* since these two numbers are different.

Second, consider a node and count either the number of cells meeting at that node n ('n' for node) or the number of corners n^* meeting at that node. The number n is also called the degree of the node.

These numbers encode the local structure of the tiling and, when considered globally, lead to fundamental constraints. For instance, there is a fundamental relations between the total numbers of cells, edges, and nodes: a finite simply connected planar tiling with V nodes, E edges, and F cells satisfies *Euler's formula*: $V - E + F = 1$.

To describe a general tiling quantitatively, one may adopt a statistical point of view. Choose an arbitrary point in the plane and consider a disk of radius R centred at that point. Within this disk, count the number of nodes and the number of cells, and compute the average number of edges meeting at a node and the average number of sides of a cell. As the radius R increases, these averages stabilise, and in the limit $R \to \infty$ one obtains two global quantities: the *average number of corners per cell*, $\bar{v}^*$, and the *average node degree* $\bar{n}^*$. These quantities provide a coarse but robust description of the tiling, independent of local irregularities (assuming that the average does not depend on the point chosen). To further refine this characterisation, one

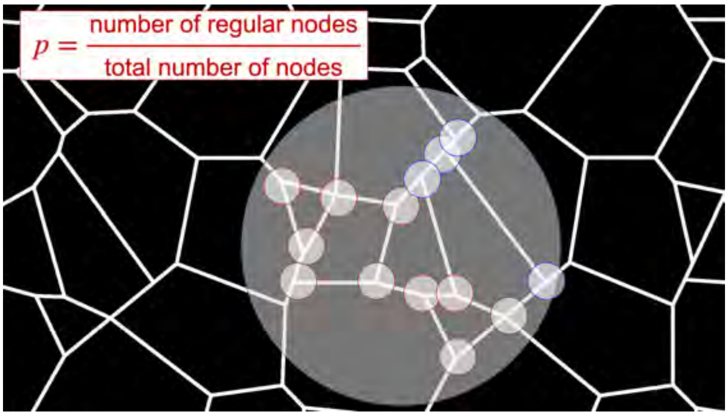

Figure 5.8: The number p characterises the average number of regular nodes (red circles) found in a tiling. If $p = 1$ all nodes are regular. If $p = 0$ all nodes are irregular.

introduces an additional parameter p, between 0 and 1, defined as the ratio of the number of *regular nodes* (nodes at which $n = n^*$) to the total number of nodes. Together, the triplet $(\bar{v}^*, \bar{n}^*, p)$ captures the essential combinatorial structure of the tiling and serves as a basis for classification. The pair $(\bar{v}^*, \bar{n}^*)$ is often referred to as a generalised *Schläfli symbol*, in reference to Ludwig Schläfli, who introduced analogous descriptors in the study of regular polytopes in the 19th century.

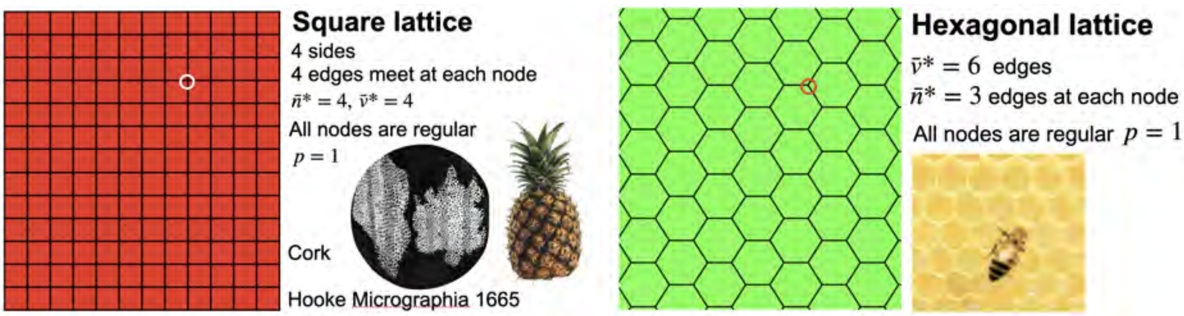

Figure 5.9: Two examples of regular tilings. A square lattice is the simplest example of a regular tiling: each cell has 4 sides and 4 corners and each node has 4 cells. So $n = n^* = v = v^* = 4$. Second, an hexagonal lattice is also regular with each cells having $v = v^* = 6$ edges and each nodes $n = n^* = 3$ cells.

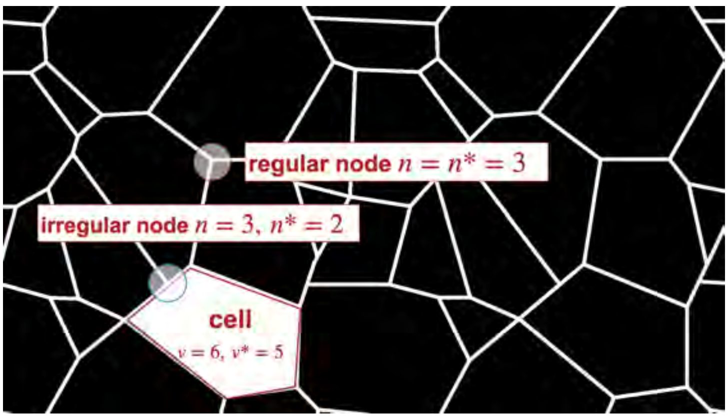

Figure 5.10: Example of regular and irregular nodes in a tiling.

We begin with the simplest situation, where all nodes are *regular*, meaning that the same number of edges meet at each node, so that $p = 1$. Familiar examples illustrate this setting: the square lattice, for instance, has cells with $n = 4$ sides and nodes of degree $v = 4$, while the hexagonal lattice has $v = 6$ sides but only $n = 3$ edges meeting at each node, as shown in Figure 5.9. These examples, together with the triangular lattice ($v = 3,\ n = 6$), already reveal an inverse relationship between the number of sides of a cell and the degree of the nodes: as cells acquire more sides, fewer edges meet at each node (see Appendix A for a proof that there are only three lattices made out of regular polygons). This observation is not accidental but follows from a general constraint for regular tilings, namely

$$\bar{v} = \frac{2\bar{n}}{\bar{n} - 2}.$$

This relation shows that increasing $\bar{n}$ necessarily decreases $\bar{v}$, a trend clearly visible in standard tilings and illustrated in the *Symbolic (or Schläfli) plane* shown in Figure 5.11, where each tiling is represented by one point. It also severely restricts the possible regular tilings of the plane, highlighting the delicate balance between local geometry and global structure.

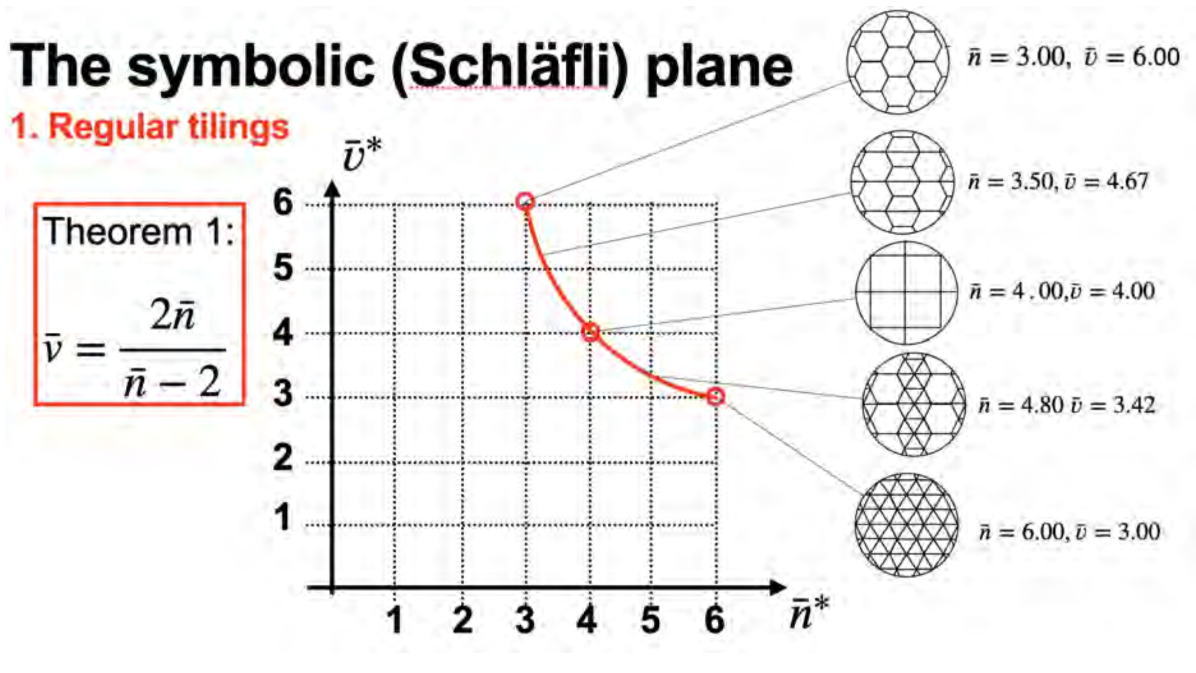

Figure 5.11: The regular ($p = 1$) convex tilings live on the red curve (a piece of a hyperbola) in the symbolic plane.

Let us now consider the opposite extreme, namely a completely *irregular tiling*, in which no node is regular, so that $p = 0$. Here, note that *irregular* in this context (not to be confused with regular polygons) does not mean disordered, nor does it require variability from place to place: the tiling may be perfectly uniform, with all nodes having the same degree and identical local environments. An example is shown in Figure 5.12: the tiling is composed of two types of square, each with four corners, so that $\bar{v}^* = 4$, while each node has two corners giving $\bar{n}^* = 2$, but three cells meet at each node (so that $\bar{n} = 3$). Such configurations arise naturally in physical systems, for instance in wall patterns, or certain biological tissues.

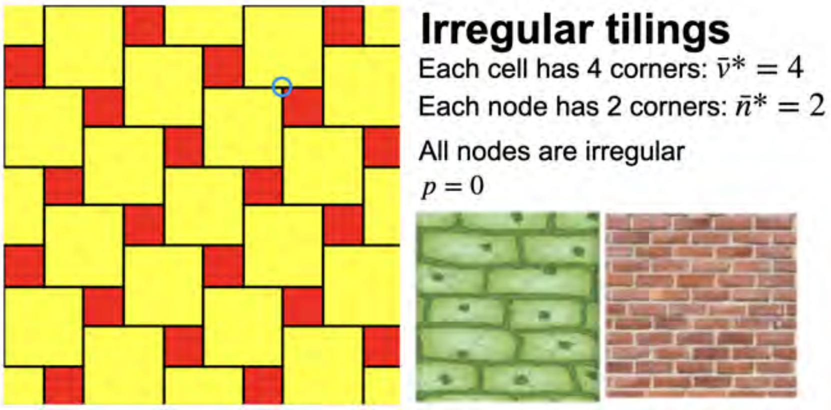

Figure 5.12: An example of an irregular tiling made with two squares. Note that there are only 2 corners at each node.

From a geometric perspective, completely irregular tilings with $p = 0$ are again characterised by a pair of average quantities $(\bar{n}^*, \bar{v}^*)$, which now lie along a different curve in the $(\bar{n}^*, \bar{v}^*)$ plane as shown in Figure 5.13.

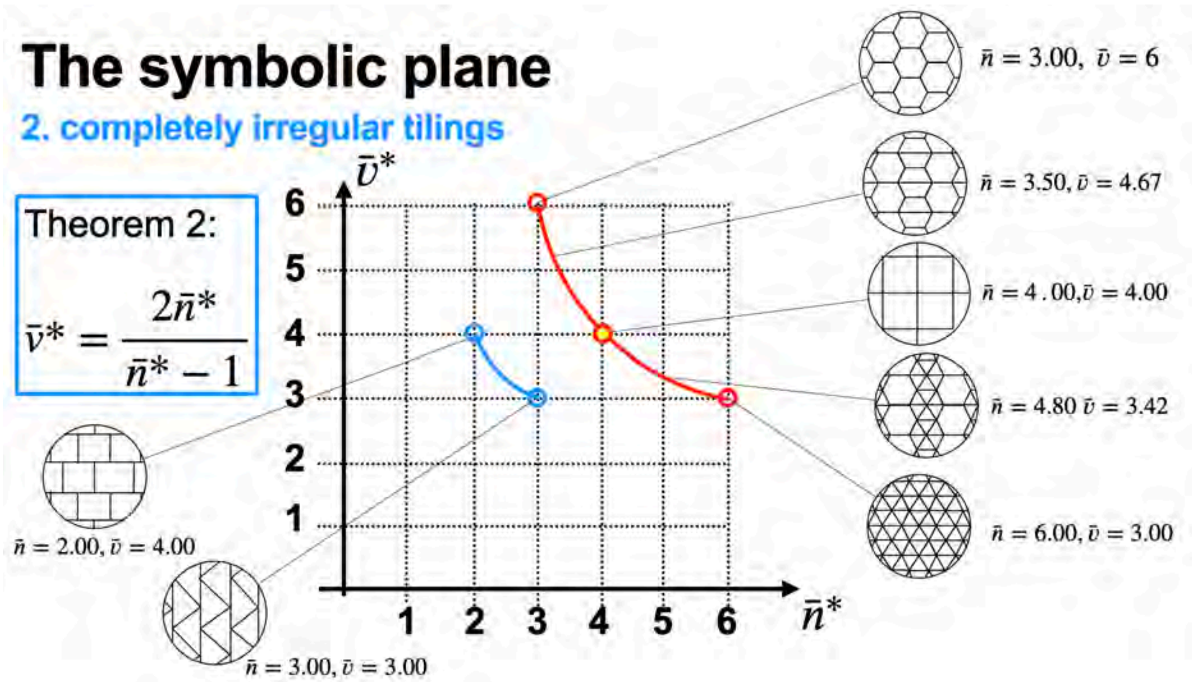

Figure 5.13: All convex completely irregular ($p = 0$) polygonal tilings lie on the blue curve in the symbolic plane

More generally, if one considers all possible *convex polygonal tilings*, these admissible configurations fill a well-defined region (shown in yellow in Figure 5.14). Within this region, the structure of any tiling is governed by the relation

$$\bar{v}^* = \frac{2\,\bar{n}^*}{\bar{n}^* - p - 1},$$

which links the average number of sides per cell, the average node degree, and the proportion p of regular nodes. This formula provides a unified framework for classifying convex tilings, with most natural and physical examples occupying positions inside this region of the symbolic space.

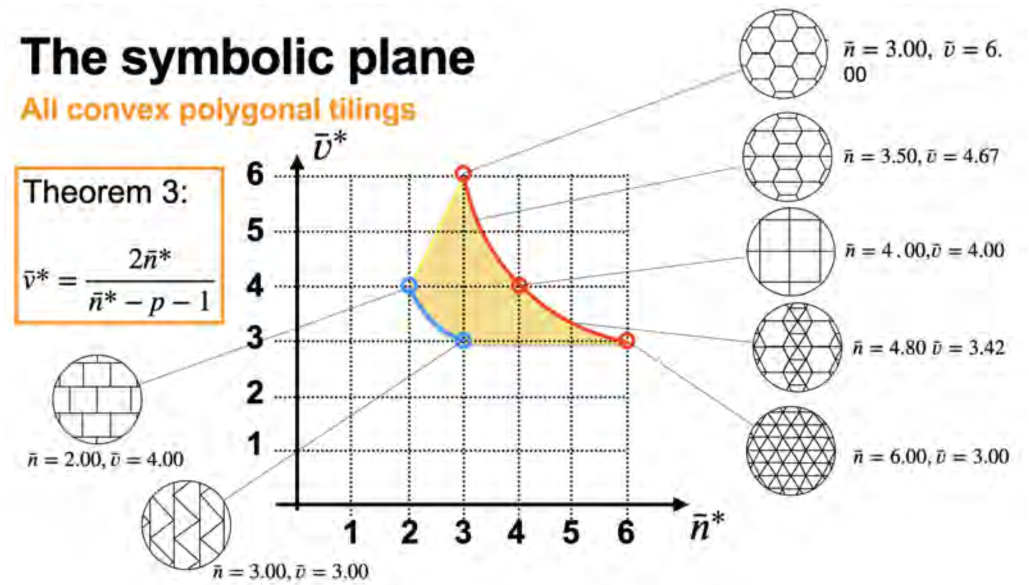

Figure 5.14: All convex polygonal tilings are located in a small region of the symbolic plane with the two extreme cases $p = 1$ and $p = 0$ as boundaries.

5.1.2 Soft cells

A natural question is whether tilings with fewer corners exist [2]. Such tilings must necessarily be nonconvex, as they fall outside the class illustrated in the preceding figures. Insight may be gained by observing familiar patterns. In the drawings of Hooke shown in Figure 5.15, the seaweed tissue is composed of repeated cells that tile the plane. Examining these “soft” cells reveals that each possesses only two corners, thereby reducing the number of corners per tile. A further example arises from the packing of wine bubbles, where the cells resemble small lenses Figure 5.15. Each node has precisely two corners. In both cases, the number of corners per tile is two.

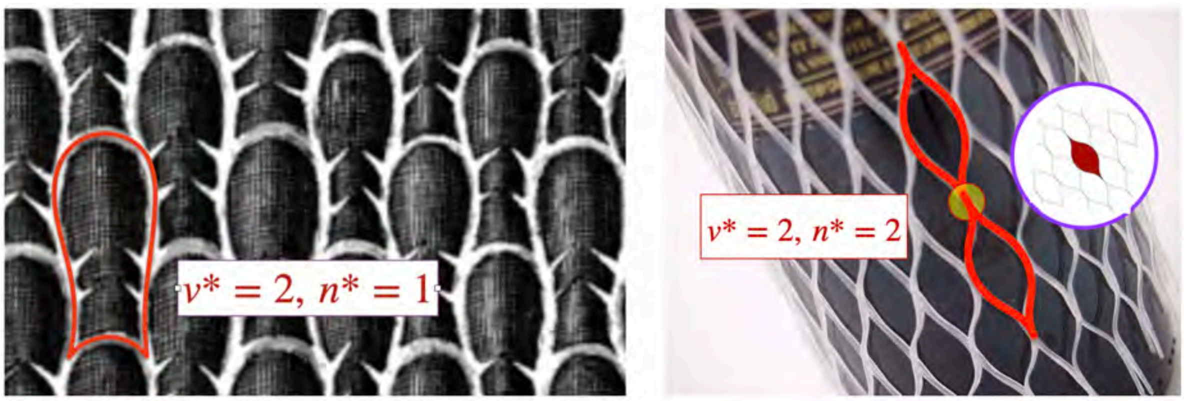

Figure 5.15: Two examples of soft cells. First, a drawing of seaweed tissue by Robert Hooke in *Micrographia* (1665). Ignoring the small barbs, we see that the tiling is monohedric (a single cell repeated) with only two corners. A second example is obtained from a familiar package wrapping, showing lens-like shapes that tile the plane

We define a *soft tiling* as a tiling composed of cells that achieve, on average, the minimal possible number of corners. In the plane, this minimum is equal to two, so that a soft tiling is characterised by cells having, on average, exactly two corners as shown in Figure 5.16. Such tilings necessarily involve nonconvex cells and fall outside the class of conventional polygonal tilings. Once this observation is established, one may begin with standard tilings formed by regular monotiles, such as those by triangles, squares, or hexagons, and then deform them. By bending the edges appropriately, these classical tilings can be transformed into soft tilings as illustrated in Figure 5.17.

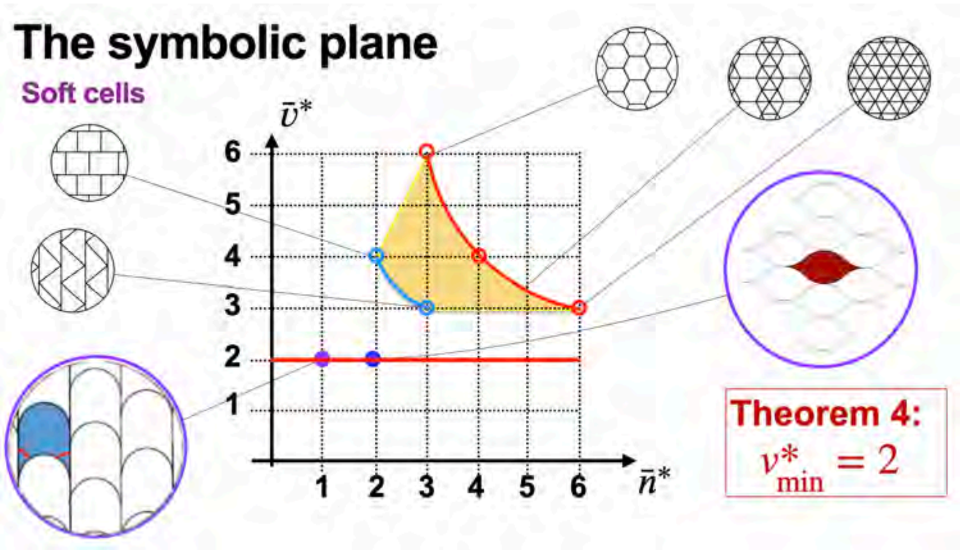

Figure 5.16: The minimal number of corners for a tiling is two as indicated by the red line in the symbolic plane.

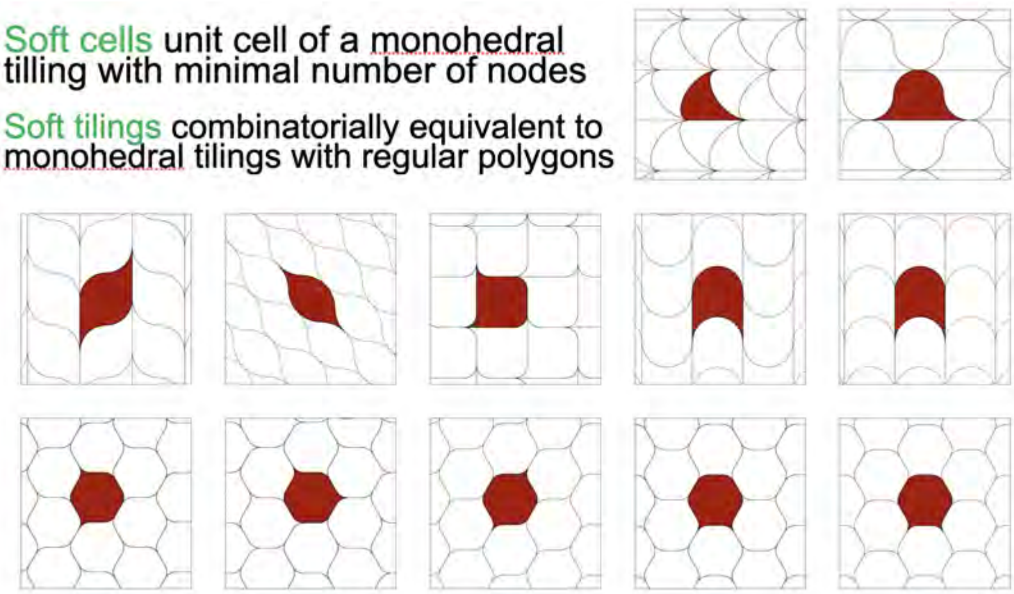

Figure 5.17: Softening of the tilings made with regular polygons. Top row is based on the softening of a triangular lattice, middle row is the softening of a square lattice, and the bottom row shows the possible softening of an hexagonal lattice.

Once the structure of soft tiling is understood, it is easy to find examples in the real world and understand that soft cells naturally appear as a normal consequence of how many biological systems grow.

5.2 Three-dimensional tilings

5.2.1 Hard tilings of space

While there exist infinitely many regular polygons that tile the plane, the situation in three dimensions is far more restrictive, with only five regular convex polyhedra, the Platonic solids Figure 5.18 (see Appendix B for a proof). Plato associated these five solids with the fundamental constituents of matter and the universe, whereas Aristotle argued that only the cube and the tetrahedron could tile space. However, this claim turned out to be incorrect as shown by Kepler many centuries later in *Harmonices Mundi* (1619). The cube does indeed tile space, with eight cubes meeting at each vertex, so that the node degree is $n = 8$. In contrast, for the regular tetrahedron, the solid angle at a vertex is such that the number of tetrahedra required to fill space around a point is $4\pi/\Omega \approx 22.79$, where Ω is the solid angle at a vertex of the tetrahedron;[1] since this is not an integer, regular tetrahedra cannot tile space.

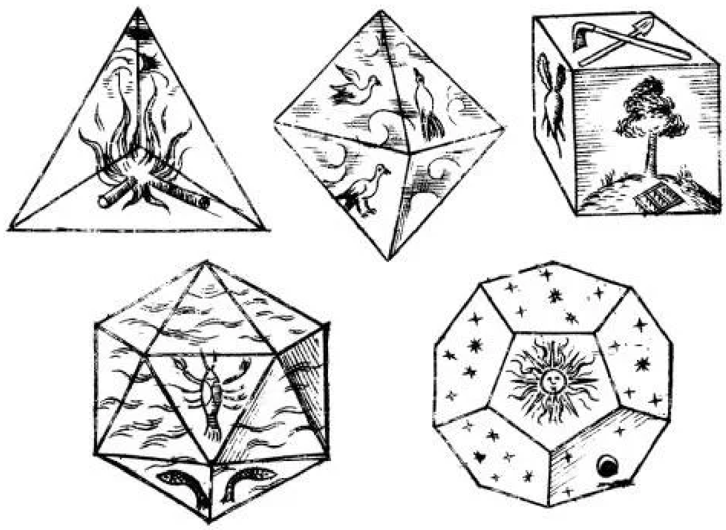

Figure 5.18: The five platonic solids as drawn by Kepler in 1619. Plato believed that they constitute the essence of Earth, Water, Fire, Air, and the Universe.

However, the truncated octahedron provides an example of a space-filling polyhedron, as studied by Kelvin and later classified by Fedorov in the context of space-filling solids. It is in fact the only Archimedean polyhedron that tiles space. Archimedean polyhedra are defined as convex polyhedra with regular polygonal faces of more than one type, arranged so that the same configuration of faces meets at every vertex. Unlike the Platonic solids, whose faces are all congruent, Archimedean solids allow for a mixture of regular polygons while preserving vertex-transitivity. Among these, only the truncated octahedron (with square and hexagonal faces) possesses the geometric compatibility required to fill three-dimensional space without gaps or overlaps.

[1] The solid angle Ω is the 3D analogue of a planar angle. It measures how large an object appears from a given point, not in terms of length but in terms of the portion of space it occupies. More precisely, it is defined as the area cut out on the unit sphere centred at the point by the rays emanating from that point and passing through the boundary of the object. Its natural unit is the steradian, and the total solid angle around a point in space is 4π. For a cube, the solid angle is $\pi/2$.

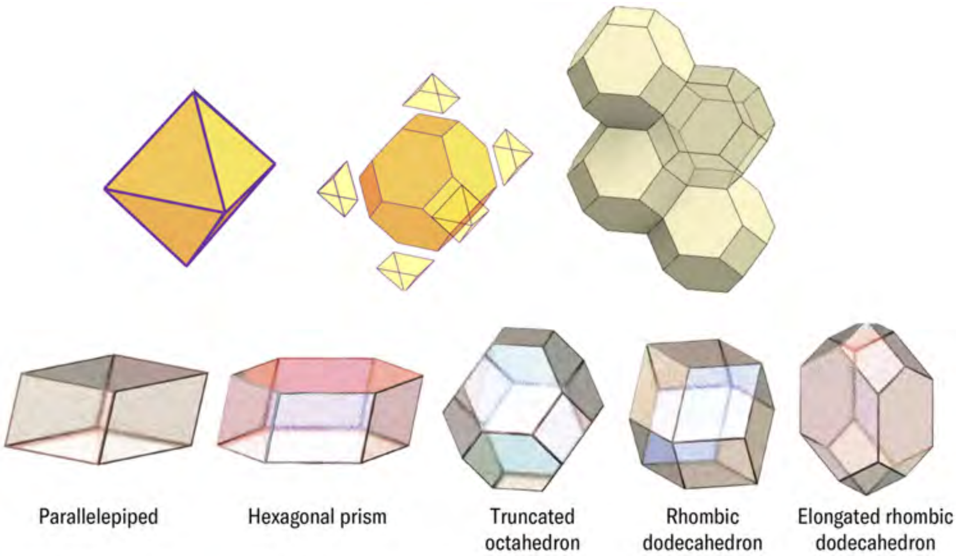

Figure 5.19: Top row: The truncated octahedron is the only Archimedean solid filling space. Bottom row: the five parallelohedra that tile space by translations.

What other polyhedron might tile space? This is a difficult question. If we restrict our attention to convex polyhedra that tile Euclidean space by translations alone (you can only translate in 3D space a unit cell and not rotate it, so that copies of the same shape fill space in a face-to-face manner), Evgraf Stepanovich Fedorov (1853–1919), a Russian crystallographer and mathematician proved, in 1885, that there exist exactly five types of such *parallelohedra* shown in Figure 5.19.

5.2.2 3D soft cells

Now that we have a basic understanding of polyhedra that fill space, we can ask the same kind of question that arose in the plane for three-dimensional shapes: can we "soften" a tiling by deforming its edges and faces so as to reduce the number of corners? Starting from the simplest example, the cubic lattice, one may try to bend faces and smooth edges while preserving a space-filling structure. In two dimensions, we saw that a square can be deformed into a shape with only two corners, and that two is in fact the minimal number required for a tile that fills the plane. In three dimensions, one can similarly ask: what is the minimal number of corners for a space-filling cell, and how should one define a "soft" cell? The surprising answer is that this minimum is actually zero. Indeed, it is enough for a proof to exhibit a space-filling shape with no corners at all. Such an example is shown in the Figure 5.20: following the edge, one observes that this shape has only one edge and exactly two faces, yet there are no corners. In this sense, the structure is a three-dimensional analogue of soft tilings.

Is this construction of a deformed cube a curiosity or part of a broader framework? Actually, we have shown that such deformations can be made generically. Starting from a regular lattice, one can attempt to deform the edges and faces while preserving the space-filling property.We have identified general conditions that must be satisfied to ensure that the cells can be deformed to fit together without gaps or overlaps and all examples examined so far suggest that such a transformation, that we call *edge-bending* [2], is possible. In each case, the deformation can be carried out so as to eliminate all corners, producing what may be called a fully "soft" cell. Figure 5.21 illustrates several examples obtained in this way: classical polyhedral cells that tile space are smoothly transformed into their soft counterparts, demonstrating that corner-free space-filling structures are not exceptional but may arise naturally from familiar geometric

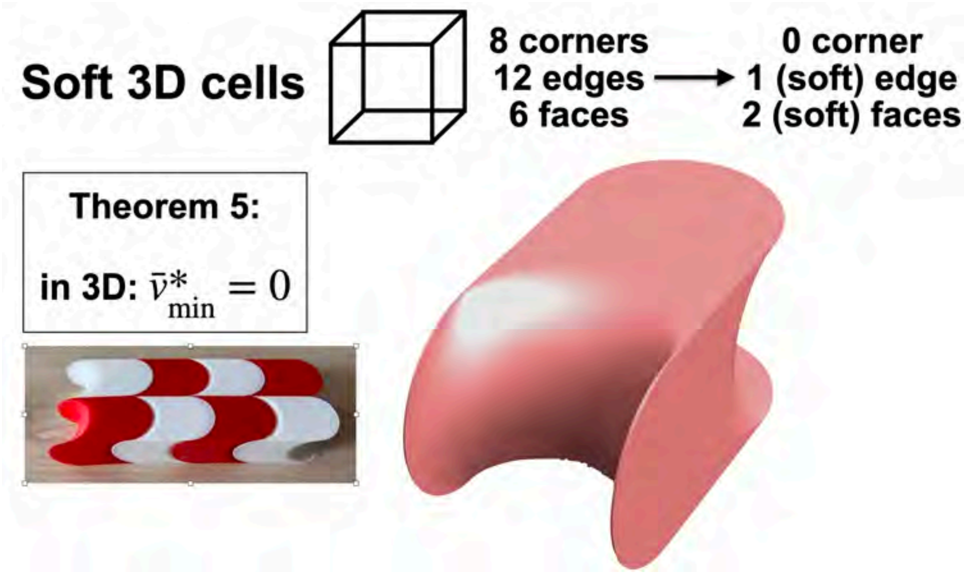

Figure 5.20: A remarkable result for tilings is that there exist 3D soft cells with 0 corner that tile space. Here is possibly the simplest example of a shape that tile space on a cubic lattice with no corner, one edge and two faces.

configurations.

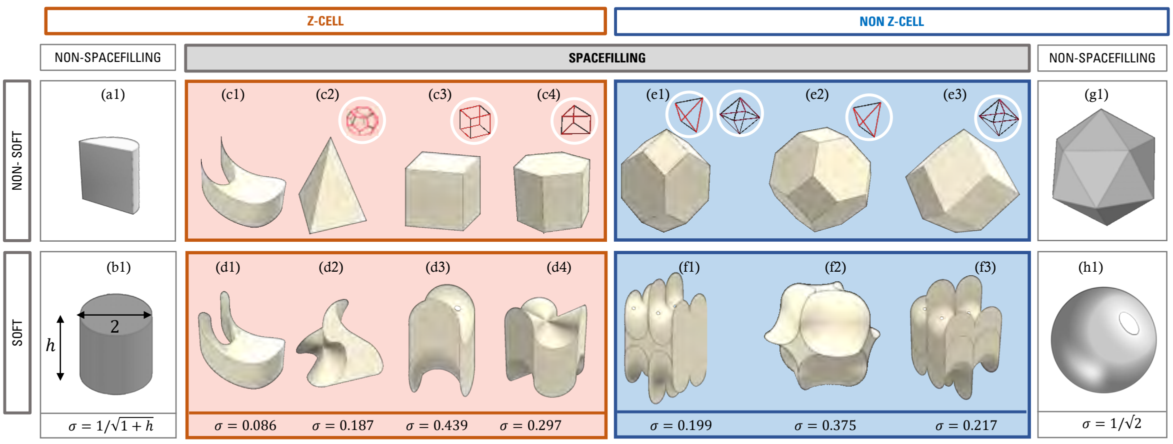

Figure 5.21: Examples of hard and soft cells classified in different categories. A *z-cell* is a type of tilings that creates prisms. Each prism can then be assembled as if they were 2D tilings (from [2]).

At the heart of these constructions lies what we call the *edge-bending algorithm*, a simple geometric procedure that systematically removes corners. The idea is to focus on a single node, where several edges meet, and examine the directions of the edges as they approach that point, their “half-tangents.” By carefully adjusting and aligning these directions, one can smoothly redirect the edges so that they meet without forming a sharp vertex. In effect, the corner is dissolved into a continuous configuration of edges and faces. Repeating this process throughout the structure progressively eliminates all vertices, transforming a conventional polyhedral tiling into a soft, corner-free, one. An example of this transformation is illustrated in Figure 5.22, where the local modification at a node leads to a global smoothing of the entire tiling.

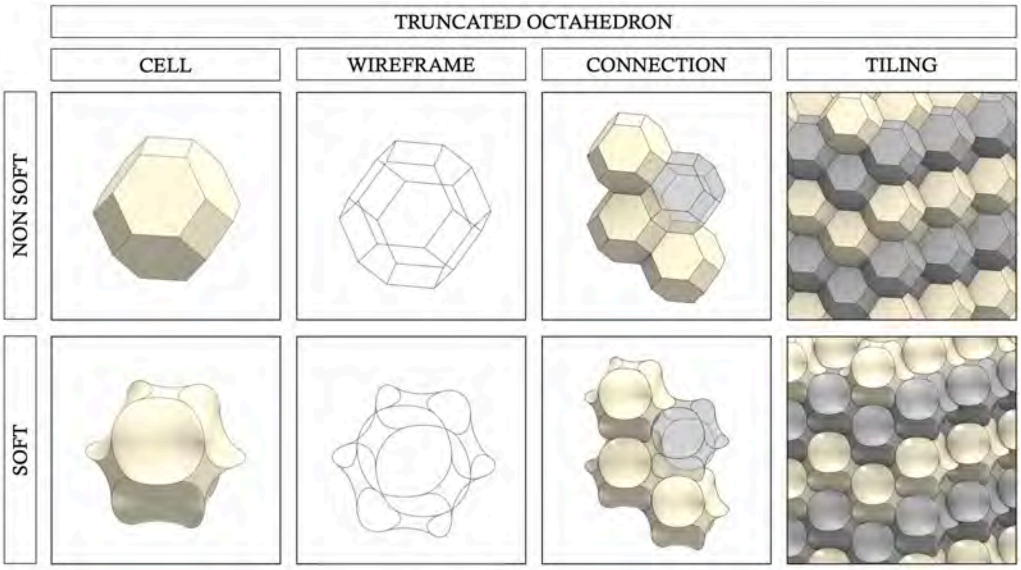

Figure 5.22: The edge-bending algorithm can be applied to the truncated octahedron (top rows). The key is to curve the edges at the nodes to remove the corners. The paths are chosen to be Dubins path and the face are chosen to be minimal surfaces as explained in the main text (adapted from [2]).

A natural question arises when one replaces straight edges by smooth curves: how should these curves be chosen? First it should be clear that there are infinitely many ways of doing so, but Dubins paths provide a simple and elegant answer. Imagine a vehicle that can only move forward and whose turning radius is limited; the car that cannot turn arbitrarily sharply. Among all possible paths connecting two positions with prescribed directions, the *Dubins path* is the shortest one that respects this curvature constraint. Rather than being arbitrary curves, these optimal paths are built from a small number of simple pieces: circular arcs of fixed radius and straight segments. In soft cells, Dubins paths offer a natural way to replace the sharp edges of a polyhedron by controlled, smooth curves. When an edge is bent, one must ensure that the resulting curve connects smoothly to the surrounding faces while keeping curvature under control. By using Dubins paths, one introduces just enough bending to eliminate corners at a node. In this way, a rigid polyhedral skeleton can be systematically transformed into a soft structure.

Once the edges have been softened, one must also reconsider the faces that connect them. Again, there are infinitely many choices. Instead of flat polygonal patches, a natural choice is to replace each face by a minimal surface. In simple terms, a *minimal surface* is a surface that minimises its area for a given boundary, much like a soap film stretched across a wire frame. If one dips a loop of wire into soapy water, the thin film that forms automatically finds the minimising area by reducing surface energy. For soft cells, minimal surfaces provide an optimal connections between curved edges. Once the boundary of a face has been bent into a curve, the minimal surface spanning this boundary naturally fills in the interior without introducing sharp features as seen in Figure 5.22. The choice of Dubins path for edges and minimal surfaces for faces is not necessarily applied for all soft cells, but this particular canonical choice gives particularly simple and pleasing cells both from a mathematical and aesthetic point of view.

A deep-sea application

The nautilus is perhaps the most iconic seashell, instantly recognised by its elegant spiral when viewed in cross-section. This planar curve is well approximated by a logarithmic spiral as we saw in the third lecture of this series. Yet, this familiar image hides a richer three-dimensional story. The cephalopod shell is divided into a sequence of chambers, with the animal occupying only the outermost one. The older chambers are essential for buoyancy control. Through a narrow tube, the *siphuncle*, the nautilus adjusts the balance between liquid and gas in each chamber, allowing it to rise or sink in the ocean. While the planar geometry has long been studied, the precise three-dimensional shape of these chambers has only recently become accessible through modern imaging techniques, revealing a surprisingly intricate and smooth architecture shown in Figure 5.23.

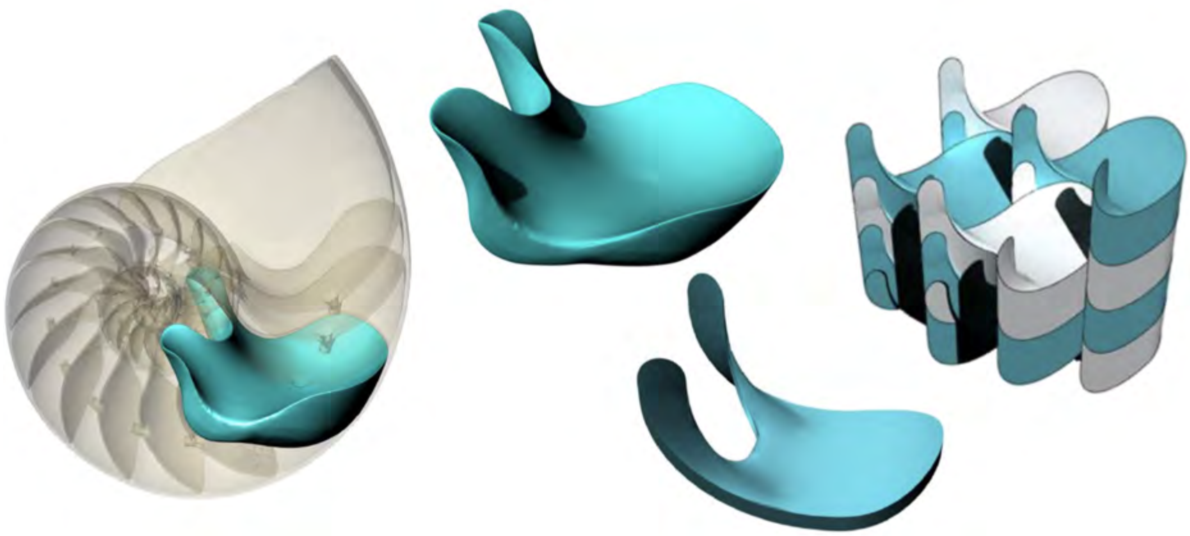

Figure 5.23: The iconic nautilus is a deep-sea diver coming up at night to feed. It controls its buoyancy through a series of chambers. The geometry of these chambers seen on the left can be simplified into a soft cell that can tile space as shown on the right.

From a geometric point of view, these chambers provide a remarkable example of a soft cell. Their boundaries are highly curved, with smooth transitions between surfaces, and they avoid the sharp corners typical of polyhedral shapes. In fact, the chambers can be interpreted as segments of a curved tubular structure that has been partitioned in a smooth way, so that the resulting pieces inherit this softness. In the language of tilings, they are close to space-filling cells with no corners at all.

5.3 Conclusions

Mathematicians have primarily studied polygons and polyhedra for the last two and a half millenia due to their much simpler linear structures (faces and edges being described by linear equations). However, soft cells appear naturally in the physical world because geometry is never free: it is shaped by energy. In biological systems, creating and maintaining sharp corners is costly, as surface tension, elasticity, and growth processes all act to smooth out irregularities. Cells, tissues, and shells evolve under constraints that favour softer curvature and continuous transitions rather than abrupt changes of direction. As a result, when many units grow and pack together, the generic outcome is not a rigid polyhedral tiling but a softened one, where corners are rounded and edges are curved. In this sense, soft cells are not exceptional constructions but rather the natural geometric outcome of growth under physical constraints, providing a unifying framework to describe structures ranging from cellular tissues to the chambers of seashells.

More recent work by our group reveals that this picture extends even further, connecting soft cells to a class of geometric structures known as triply periodic minimal surfaces These remarkable surfaces are smooth surfaces that repeat periodically in three independent directions and locally minimise area, much like soap films in space. They partition space into intertwined regions (aptly called *labyrinths* due to their delightful complexity) with highly regular, yet fully curved geometry and are closely related to complex analysis. Remarkably, the fundamental building blocks, or unit cells, associated with these surfaces can be interpreted as soft cells [3]. This connection places soft cells at the crossroads of geometry, biology, and materials science, where they serve not only as models of natural growth but also as prototypes for complex structures observed in polymers, membranes, and advanced (smart) materials. It opens many new exciting paths for research but also for new designs in architecture that we will be exploring with our students and collaborators in the years to come.

Acknowledgements: This talk is based largely on the work of my good friend Gabor Domokos and his collaborators including the remarkable Krisztina Regős and Ákos Horváth. I wish to express my gratitude to them for sharing many of the visuals used during the lecture and in this write-up and for involving me, a few years back, in the study of soft cells.

5.4 Further Reading and Exploration

- A comprehensive mathematical treatment of tilings and widely consider as the bible in the field is *Tilings and Patterns* by B. Grünbaum and G. C. Shephard (1987) [8].
- A more accessible introduction to the subject is *Tilings and Patterns: An Introduction* by B. Grünbaum and G. C. Shephard (1989) [9].
- A modern and engaging entry point for those ready to roll up their sleeves is *The Tiling Book: An Introduction to the Mathematical Theory of Tilings* by Colin Adams (2022) [1].
- A classic recreational approach to tilings and combinatorics is *Polyominoes: Puzzles, Patterns, Problems, and Packings* by S. W. Golomb (1965) [7].
- A visually rich exploration linking art and symmetry is *M. C. Escher: Visions of Symmetry* by Doris Schattschneider (2004) [10].
- An accessible and entertaining collection of essays touching on tilings is *Time Travel and Other Mathematical Bewilderments* by Martin Gardner (1988) [6].
- Jaap Scherphuis developed an extensive interactive resource on tilings, where one can explore classifications and experiment with many examples: `https://www.jaapsch.net/tilings/`.
- The Wolfram Demonstrations Project offers several interactive modules on tilings and tessellations, allowing users to manipulate parameters and visualize geometric structures: `https://demonstrations.wolfram.com/topic.html?topic=Tessellations`.
- The Geometry Center provides visual and interactive material on symmetry and tilings, suitable for exploration and teaching: `http://www.geom.uiuc.edu/`.
- The Bridges Organization hosts mathematical art resources, including interactive explorations of tilings and patterns: `https://bridgesmathart.org/`.

References

[1] Colin Adams. *The Tiling Book: An Introduction to the Mathematical Theory of Tilings*. American Mathematical Society, 2022.

[2] Gábor Domokos, Alain Goriely, Ákos G Horváth, and Krisztina Regős. Soft cells and the geometry of seashells. *PNAS nexus*, 3(9):pgae311, 2024.

[3] Gábor Domokos, Alain Goriely, Ákos G Horváth, and Krisztina Regős. Soft cells, kelvin's foam and the minimal surfaces of schwarz. *arXiv preprint arXiv:2412.04491*, 2025.

[4] Gábor Domokos, Ákos G Horváth, and Krisztina Regős. A two-vertex theorem for normal tilings. *Aequationes mathematicae*, 97(1):185–197, 2023.

[5] Gábor Domokos, Douglas J Jerolmack, Ferenc Kun, and János Török. Plato's cube and the natural geometry of fragmentation. *Proceedings of the National Academy of Sciences*, 117(31):18178–18185, 2020.

[6] Martin Gardner. *Time Travel and Other Mathematical Bewilderments*. W. H. Freeman, 1988.

[7] Solomon W. Golomb. *Polyominoes: Puzzles, Patterns, Problems, and Packings*. Princeton University Press, Princeton, 1965.

[8] Branko Grünbaum and G. C. Shephard. *Tilings and Patterns*. W. H. Freeman, New York, 1987.

[9] Branko Grünbaum and G. C. Shephard. *Tilings and Patterns: An Introduction*. W. H. Freeman, New York, 1989.

[10] Doris Schattschneider. *M. C. Escher: Visions of Symmetry*. Harry N. Abrams, 2004.

Appendix A: Basic regular planar tilings

One of the most basic result from the theory of planar tilings is the following proposition

Proposition. The only tilings of the plane by a single regular polygon are those made of equilateral triangles, squares, and regular hexagons.

Proof. Consider a tiling of the plane using only one kind of regular polygon with m sides (an m-gon). At each vertex of the tiling, suppose that exactly n polygons meet (the degree of the node we discuss in the main text). Since the tiling fills the plane without gaps or overlaps, the angles around each vertex must add up to a full circle, that is 2π radians. The key for the proof is that for a regular m-gon, each interior angle is

$$\alpha = \frac{(m-2)\pi}{m}.$$

Thus, the condition for a tiling is

$$n \cdot \frac{(m-2)\pi}{m} = 2\pi.$$

Rewriting this gives

$$m = \frac{2n}{n-2}.$$

We now look for integer solutions with $m \geq 3$ and $n \geq 3$. Testing small values of n:

n	$m = \frac{2n}{n-2}$
3	6
4	4
5	$\frac{10}{3}$ (not an integer)
6	3
$n \geq 7$	$m < 3$ (not possible)

Thus, the only possibilities are

$$(m, n) = (3, 6), \quad (4, 4), \quad (6, 3).$$

These correspond respectively to tilings by equilateral triangles, squares, and regular hexagons. No other regular polygon satisfies the required condition.

□

Appendix B: Regular polyhedra

The second basic result concerns the existence of regular polyhedra

Proposition. There are only five regular convex polyhedra.

Proof. A regular polyhedron is a solid made of identical regular polygons (say m-gons), with the same number n of faces meeting at each vertex. We determine all possible pairs (m, n).

At each vertex of a polyhedron, several faces meet, and each face contributes its interior angle at that point. If we imagine placing all these angles around the vertex (for instance, the three $\pi/2$ angles of the square at a node), they must fit together in space. In the plane, angles can add up exactly to 2π to form a flat configuration, but in three dimensions a convex polyhedron must “bend” around the vertex. This means the total angle must be strictly less than 2π; otherwise, the faces would lie flat in a plane and no three-dimensional shape would be formed (for instance, if you cut a paper cube along an edge, the three sides at a node would flatten and add up in the plane to $3\pi/2$). If the sum were equal to 2π, the surface would be flat at that point, and if it were greater than 2π, the faces would overlap. Thus, for a convex polyhedron, the condition is that the sum of the face angles at each vertex is strictly less than 2π. Again, we know that for a m-gon, the interior angle is $\alpha = \frac{(m-2)\pi}{m}$, therefore, we must have

$$n \cdot \frac{(m-2)\pi}{m} < 2\pi.$$

Rewriting this inequality gives

$$(m-2)(n-2) < 4.$$

We now look for integers $m \geq 3$ and $n \geq 3$ satisfying this inequality. Since $(m-2)(n-2)$ is a positive integer, the only possibilities are

$$(m-2)(n-2) = 1,\ 2,\ \text{or } 3.$$

We examine these cases:

- If $(m-2)(n-2) = 1$, then $m-2 = 1$ and $n-2 = 1$, so $(m,n) = (3,3)$.
- If $(m-2)(n-2) = 2$, then either $(m-2, n-2) = (1,2)$ or $(2,1)$, giving $(m,n) = (3,4)$ or $(4,3)$.
- If $(m-2)(n-2) = 3$, then either $(m-2, n-2) = (1,3)$ or $(3,1)$, giving $(m,n) = (3,5)$ or $(5,3)$.

Thus, the only possible pairs are

$$(3,3), \quad (3,4), \quad (4,3), \quad (3,5), \quad (5,3).$$

The inequality above provides only a necessary condition for the existence of a regular convex polyhedron. It shows that there are at most five possible pairs (m,n), but does not by itself guarantee that such polyhedra exist.

To complete the argument, one must verify that each admissible pair actually corresponds to a geometric solid. This can be seen by explicit construction:

- $(3,3)$: three equilateral triangles meet at each vertex, forming a tetrahedron,
- $(4,3)$: three squares meet at each vertex, forming a cube,
- $(3,4)$: four triangles meet at each vertex, forming an octahedron,
- $(5,3)$: three pentagons meet at each vertex, forming a dodecahedron,
- $(3,5)$: five triangles meet at each vertex, forming an icosahedron.

No other values satisfy the required condition, so there are exactly five regular convex polyhedra.

□