Growing a Voronoi Diagram Like a Money Plant

An illustrated tour of nature's favourite tile pattern

K-Dense Web
contact@k-dense.ai

May 2026

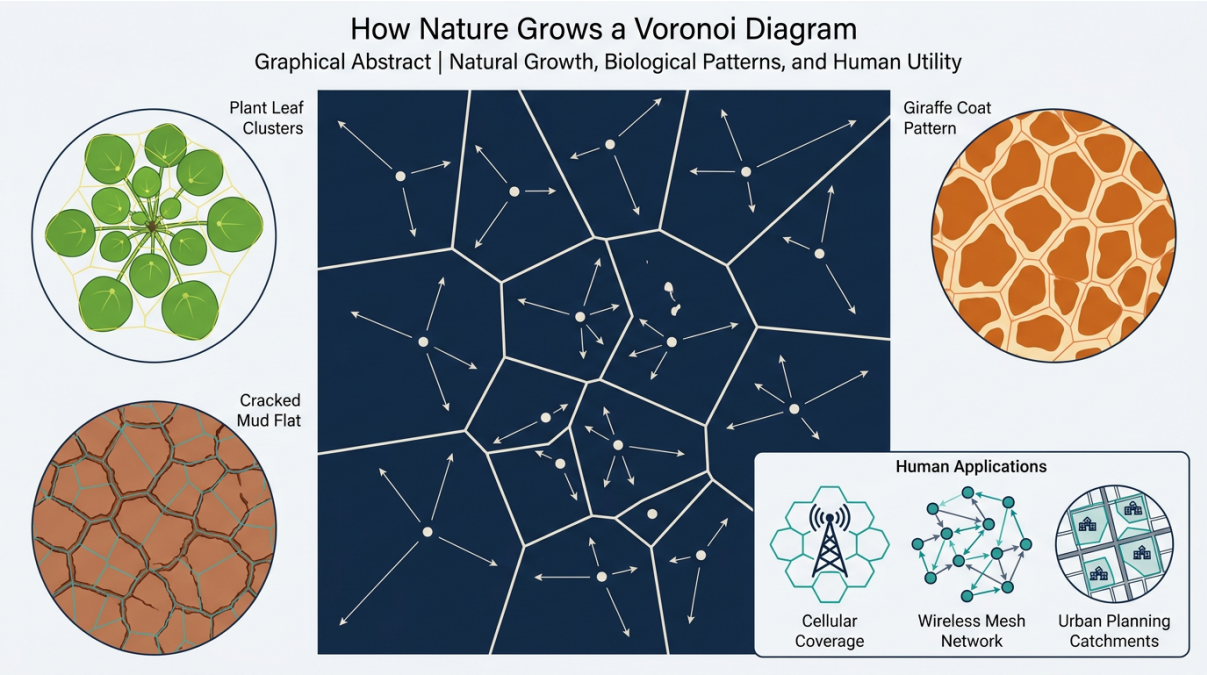

Graphical abstract. The same simple geometric rule — "every point in a region is closer to its own seed than to any other seed" — traces leaf packs in a Chinese money plant, the patches of a giraffe's coat, the polygons in a cracked mud flat, and a textbook computer-generated example. Once you see the rule, you start to see it everywhere: cell-tower coverage, mesh-network routing, even school catchment areas.

Abstract

A *Voronoi diagram* is one of the cleanest pieces of mathematics that biology already invented. Given a set of *seed points* on a flat surface, it partitions the surface into one region per seed — a *Voronoi cell* — such that every point in a cell is closer to its own seed than to any other. You can build such a tiling by hand, or you can imagine *growing* each seed outward at the same speed until neighbouring colonies bump into each other and freeze. That growth picture is exactly what a Chinese money plant (*Pilea peperomioides*) does as new leaves push outward, and it is also a very good first model for the patches on a giraffe, the polygons of a dried-out mud flat, the coverage zones of cell-phone towers, and the catchment areas of urban schools. In this report we (i) explain what a Voronoi diagram is in plain language; (ii) build one computationally on a real Wikimedia-Commons photograph of a *Pilea*; (iii) repeat the same algorithm on a giraffe coat, a cracked mud flat, and a random control; (iv) compare them statistically; and (v) end with a

short tour of human-made systems — cell coverage, wireless mesh networks, and urban planning — where the same math shows up.

Audience: curious high-schooler with a year of algebra and no calculus. **Reading time:** about 25 minutes. **Reproducible code & data:** the four overlays, three statistics figures, and all CSV/JSON tables in this report were produced by the workflow described in the appendix. **All photographs:** public domain or CC-BY-SA, sourced from Wikimedia Commons.

---

1 The growing-circle idea

Look at a Chinese money plant (Figure 2). It is a perfectly polite houseplant: round, coin-shaped leaves on long stalks, each leaf pointing in the direction where there was most empty space when it was born, and almost no two leaves overlapping. If you stare at one for a while, you can pretend that each leaf is a *little colony* that grew outward from a single sprouting point on the stem and stopped when it ran into a neighbour. The pattern you end up looking at — a flat surface diced up into convex polygons, one per leaf, sharing straight edges with their neighbours — is a perfect physical example of a mathematical object called a *Voronoi diagram*.

The idea is older than computers. The German mathematician G. Lejeune Dirichlet (1850) introduced what we now call a *Dirichlet tessellation* in 1850 while studying quadratic forms; the Russian mathematician Georges Voronoi (1908) generalised it to arbitrary dimensions in 1908, and the modern computational-geometry community settled on his name (Aurenhammer, 1991; Okabe et al., 2000). But the recipe is so simple that you can run it in your head:

> **The Voronoi recipe.** Drop a handful of points (call them *seeds*) on a piece of paper. From each seed, grow a circle at the same rate, in lock-step. When two circles meet, freeze the boundary right there — it will be the perpendicular bisector of the line between the two seeds. Keep going until every part of the paper has been claimed by exactly one seed. The shape that each seed ends up owning is called its *Voronoi cell*.

Figure 1 shows the recipe step by step. Notice four things which will return all through this report:

1. Every cell is *convex* (no caves or indentations) — a direct consequence of the fact that perpendicular bisectors are straight lines.
2. Every internal edge is shared by *exactly two* cells, and every internal corner is where exactly three cells meet (in generic, non-degenerate point sets).
3. Seeds near the boundary of the paper inevitably own *larger* cells, because there is more empty space on the outside for them to claim. This will bite us when we compare a giraffe with a money plant.
4. If you move a seed, only the cells in its immediate neighbourhood change. The diagram is *local*.

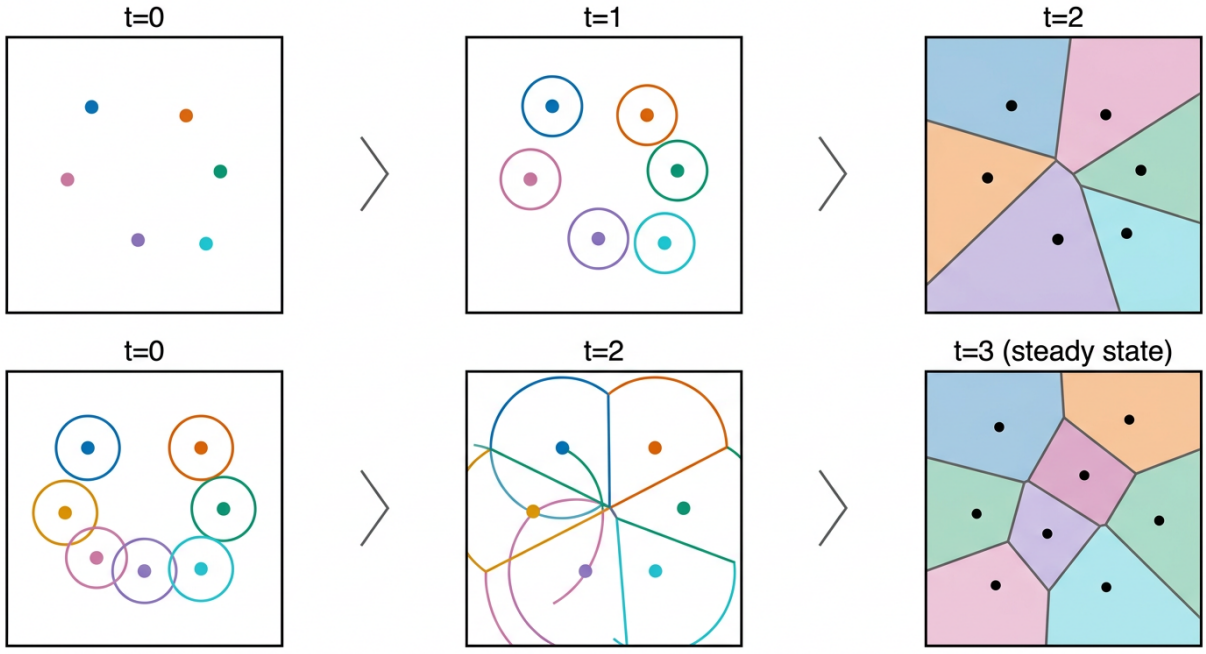

Figure 1: Growing a Voronoi diagram, one timestep at a time. Each coloured dot is a seed. Equal-radius circles grow outward from every seed at the same rate. Wherever two circles meet, they flatten into a shared straight edge — the perpendicular bisector of the two seeds. When the growth fills the canvas, every point on the paper has been claimed by its nearest seed, and the cells freeze into the Voronoi tessellation shown on the right.

A one-line formula. Let $S = \{s_1, s_2, \ldots, s_n\}$ be a set of n seed points in the plane. The Voronoi cell of seed s_i is

$$V(s_i) \;=\; \{\, x \in \mathbb{R}^2 \;\mid\; \|x - s_i\| \le \|x - s_j\| \text{ for all } j \neq i \,\}, \tag{1}$$

where $\|x - s_i\|$ is the ordinary straight-line distance from x to s_i. Read in English: "$V(s_i)$ is the set of all points x that are at least as close to s_i as to any other seed." That single equation is the whole subject; the rest is consequences (Aurenhammer, 1991; Okabe et al., 2000). Good algorithms compute the diagram of n seeds in $O(n \log n)$ time (Fortune, 1987), which makes it cheap enough to run on every photograph in this report.

2 The money plant: a real, growing Voronoi

2.1 Reading the picture

Figure 2 shows a top-down photograph of a *Pilea peperomioides* (Chinese money plant, Wikimedia Commons, CC-BY-SA 4.0) with our Voronoi overlay on top. The yellow lines are the cell edges computed by the algorithm in equation (1); the small yellow dots are the *seeds* we fed in.

What we used as seed points, and why. The original brief asked for *stomata* (microscopic breathing pores on the underside of a leaf) as seed points. A top-down whole-plant photograph cannot resolve stomata — they are roughly 20 µm apart and the image was taken with a normal lens. We therefore used the *leaf-disc centres* as our seeds: each round coin-shaped leaf is one growing colony in the canopy. The biology that matters at *this* image scale is leaf packing, not stomatal patterning. Both are documented examples of Voronoi-like organisation in plants (Bohn et al., 2002; Pielou, 1977); we picked the one the photograph actually shows.

We found the leaf centres automatically: convert the picture to HSV colour, mask the green foliage, run a Hough circle transform tuned to the expected leaf radius (149–374 px on the working image), and keep only circle centres that lie inside the foliage mask. The pipeline returned $n = 22$ seeds.

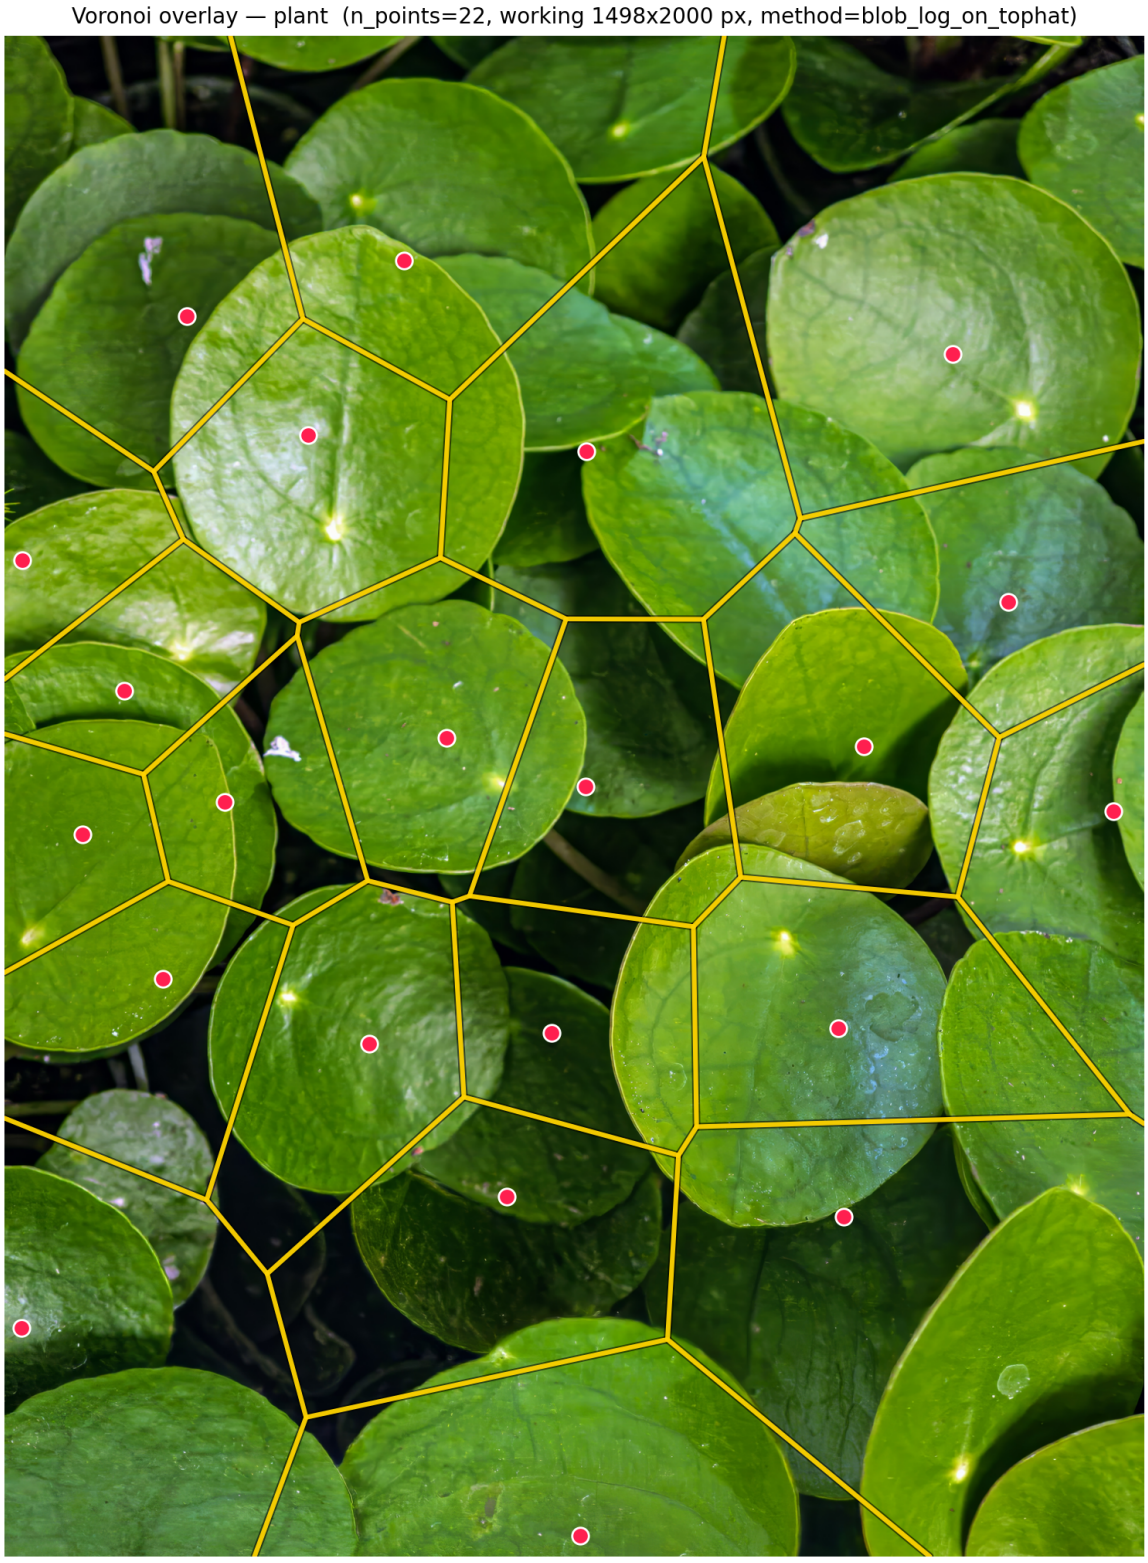

Figure 2: A real Voronoi diagram, drawn on a real plant. Top-down photograph of *Pilea peperomioides* (Wikimedia Commons, "*(MHNT) Pilea peperomioides Foliage*", CC-BY-SA 4.0). Yellow dots: leaf-disc centres detected by Hough circle transform. Yellow lines: the corresponding Voronoi cell edges from equation (1). Notice how each cell edge passes *exactly* along the gap between two neighbouring leaves — the algorithm rediscovered the plant's natural packing without being told where the leaves were.

2.2 Why the algorithm matches the plant

When a new *Pilea* leaf grows out of the stem, two things compete. Its petiole tries to point in the direction of the most available light, and its disc tries to expand until it touches a neighbour. Across the whole crown, those two rules mean each leaf occupies (roughly) the patch of canopy that is closer to *its* sprouting point than to any other — which is precisely the Voronoi cell of its disc centre. Patterns like this are pervasive in plant biology: the venation network inside a single leaf is itself well-described by a Voronoi-like decomposition of the leaf surface (Bohn et al., 2002), and Pielou used Dirichlet/Voronoi tessellations in 1977 to model the spatial competition between plants in an ecosystem (Pielou, 1977).

Mathematically, the plant is doing something even stronger than a vanilla Voronoi diagram. If you measure where each leaf's centre sits inside its own cell, the centres land very close to the *centroid* (the centre of mass) of the cell — our pixel-level measurement gives a normalised seed-to-centroid offset of 0.20 (Table 1), the smallest of any subject in this report. A diagram in which every seed sits exactly at the centroid of its cell is called a *centroidal Voronoi tessellation* (CVT), and CVTs are the natural answer to the problem " where should I place n points to cover a region as evenly as possible?" (Du et al., 1999). So a money plant is not just growing a Voronoi diagram; it is growing one whose seeds are arranged near-optimally for canopy coverage. This is the optimisation principle nature seems to have rediscovered by trial and error.

3 The giraffe: same algorithm, different biology

Figure 3: Voronoi overlay on a reticulated giraffe (Wikimedia Commons, "*DomingolaJirafa*", CC-BY-SA 4.0). Cyan dots: centroids of the dark coat patches, automatically extracted with a brown-hue colour mask, distance-transform watershed, and area/circularity filters ($n = 30$). Cyan lines: the Voronoi tessellation built on those centroids.

A reticulated giraffe (*Giraffa camelopardalis reticulata*) wears a very nearly polygonal carpet (Figure 3). Each dark patch is bounded by light-coloured "alleys" along which the coat does not pigment, and on a flat side view the patches look astonishingly like a Voronoi tessellation. When we treat the centre of each patch as a seed and draw the resulting diagram, the Voronoi edges thread almost exactly down those light alleys — we never told the algorithm where to put them.

The biology underneath is different from the plant: giraffe coat patterns are the visible output of a chemical *reaction–diffusion* process during embryonic development. Alan Turing's reaction–diffusion model, refined by James Murray for mammalian coats, predicts that two interacting chemicals diffusing through skin will spontaneously settle into spotted, striped, or reticulated patterns whose characteristic length scale depends on the diffusion rates and the body size at patterning time (Murray, 2003). The resulting pattern is not exactly Voronoi — the cell sizes and shapes are biased by the reaction kinetics — but at the eye-level scale it is well-approximated by one. The same is true of insect wing-cell domains, fish scale arrangements, and many epidermal tissues.

A caveat about edge cells. All our giraffe seeds sit on the *body* of the animal — a region that occupies only part of the photograph. Cells whose seeds happen to lie near the silhouette boundary therefore inherit some of the empty background as part of their "cell". That distorts our cell-area statistics (see Section 6): the giraffe scores as "clustered" in a Clark–Evans sense, not because the patches are physically bunched, but because the seeds occupy a sub-region of the image. We flag this and discuss it as a teaching point about *support regions* in point-pattern analysis, rather than pretending it is biology.

4 The cracked mud flat: drying as growth

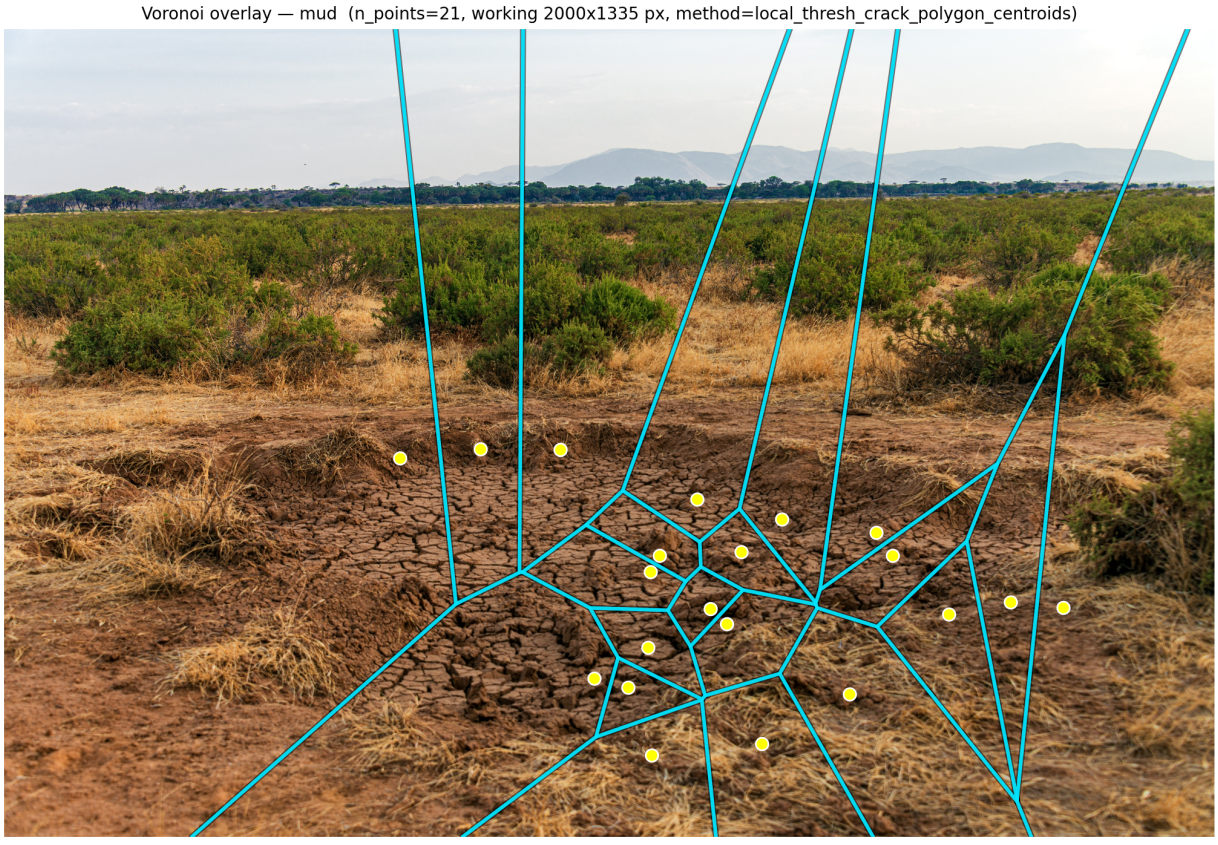

Figure 4: Voronoi overlay on a dried-out elephant mudwallow in Samburu National Reserve, Kenya (Wikimedia Commons, CC-BY-SA 4.0). Cyan dots: centroids of the polygonal mud islands ($n = 21$), extracted by local thresholding on the brown mud ROI. Cyan lines: the Voronoi tessellation on those centroids.

A drying mud flat is one of the cleanest physical examples of a tessellation in nature (Figure 4). As water leaves the mud, the surface tries to shrink. It cannot shrink everywhere at once, so it relieves the stress by cracking: cracks start at random weak points, propagate, and meet at roughly 120° angles, carving the surface into convex polygons (Goehring, 2013). The result is a network of straight-edged tiles surrounded by a fissure mesh — in other words, the empty space *between* a Voronoi-like set of crack-nucleation seeds.

There is a useful subtlety here. The mud cracks are not literally a Voronoi diagram (the angles at crack junctions are mostly T-junctions of 90° + 90° + 180° rather than the 120° triple junctions of a generic Voronoi tessellation), because cracks form *sequentially*, not all at once: each new crack runs perpendicular to the dominant local stress, which is itself shaped by older cracks (Goehring, 2013). Even so, the leading-order picture is a polygonal partition with one seed per drying domain, and Voronoi statistics give a fair first-cut description of the tile-size distribution and tile-shape regularity (Goehring, 2013). As with the giraffe, our seeds lie inside a subregion of the photograph (the wallow itself); we again clip to that region when interpreting the numerical statistics.

5 A random control: the diagram of pure chance

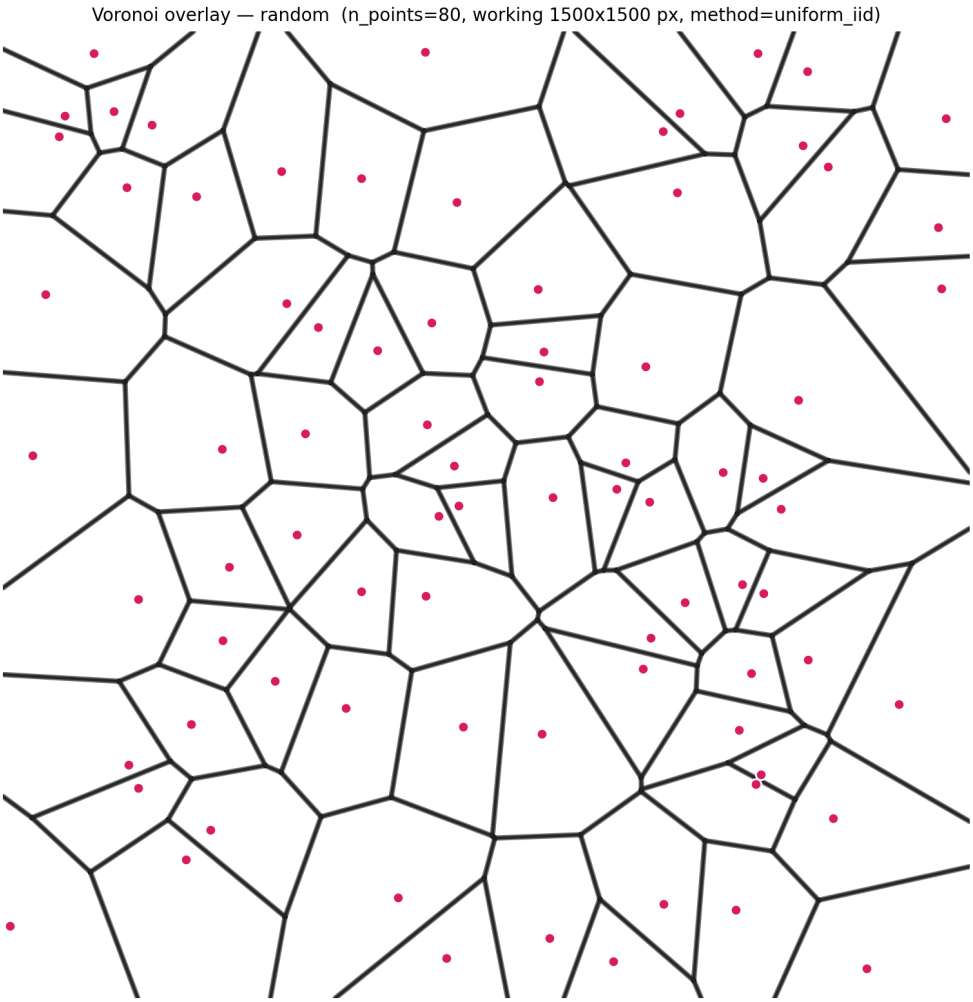

Figure 5: Voronoi diagram of 80 uniform–random points (`numpy.random.default_rng(seed=42)`) in a 1500 × 1500 px square. No biology, no physics, just pseudorandom (x, y) pairs. This is the control we use to ask: how much of what we see in the leaf, the giraffe, and the mud is *just* the geometry of a random set of points?

If you spread points uniformly at random across a piece of paper and run the same algorithm, you

get a *Poisson–Voronoi* tessellation (Figure 5). This is the natural baseline against which to compare any real-world tessellation. It is also visually instructive: even with no biology and no physics, the resulting picture has cells, edges and corners that look pleasingly organic. Random points alone are enough to fool the eye into seeing structure — which is exactly why we need quantitative comparisons.

6 Putting numbers on the four pictures

Looking at the four overlays in Figures 2, 3, 4 and 5 is enough to convince anyone that the same algorithm is doing useful work in very different settings. But we can do better than "it looks right". For each subject we compute three scale-free numbers that high-schoolers can interpret without calculus:

Coefficient of variation of cell area, $\mathrm{CV}_{\mathbf{area}}$. How uneven are the cell sizes? Smaller means more uniform.

Clark–Evans index, R. How regular is the seed pattern? $R = 1$ means random (Poisson); $R > 1$ means the seeds are more evenly spread than random; $R < 1$ means they are clustered (or constrained to a sub-region).

Normalised centroidal offset, δ. How close does each seed sit to the centre of its own cell? Small δ means the diagram is close to a centroidal Voronoi tessellation (Du et al., 1999).

Table 1: Per-subject Voronoi summary statistics. CV(area/mean) is the coefficient of variation of cell area normalised to the per-subject mean. Clark–Evans $R = d_{\mathrm{NN}}^{\mathrm{obs}} / d_{\mathrm{NN}}^{\mathrm{Poisson}}$, where the expected nearest-neighbour distance under a homogeneous Poisson process is $0.5/\sqrt{\rho}$ for seed density ρ. The normalised centroidal offset is $\|s_i - c_i\|/\sqrt{A_i}$ averaged across seeds. See main text for the support-region caveat affecting giraffe and mud.

Subject	n **seeds**	**Cells (clipped)**	$\mathrm{CV}_{\mathbf{area}}$	**Clark–Evans** R	**Mean** δ
Pilea leaf packing	22	22	**0.52**	**1.42**	**0.20**
Giraffe coat	30	30	2.65	0.32	0.50
Cracked mud flat	21	21	1.52	0.49	0.44
Uniform random	80	80	0.52	1.12	0.24

Figures 6–8 show the underlying distributions of all three metrics across the four subjects.

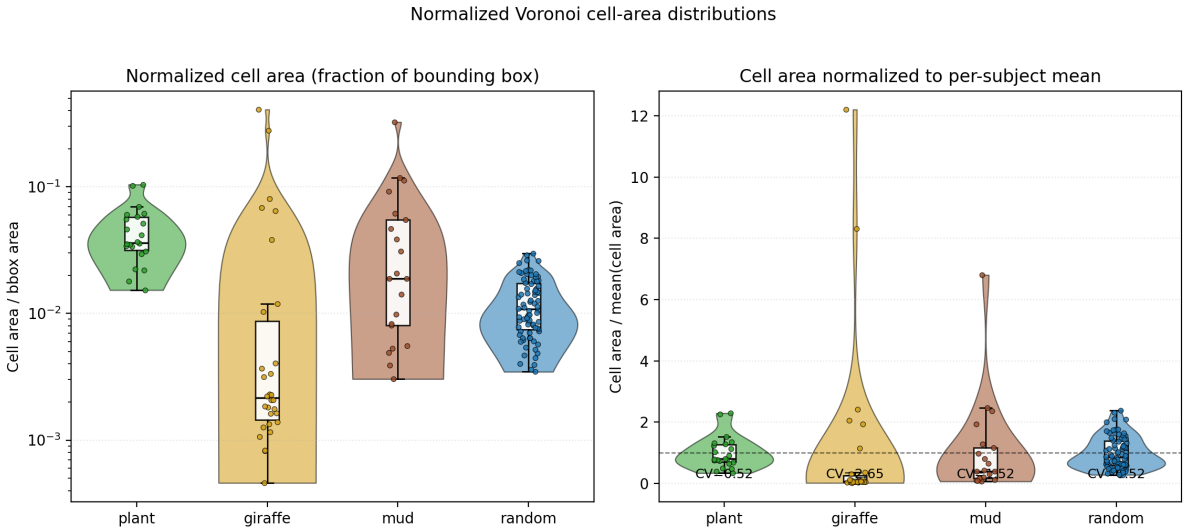

Figure 6: Cell-area distributions. Left: each cell's area as a fraction of the working bounding-box area (log-y). Right: each cell's area divided by the per-subject mean (a dimensionless "how uneven are my cells" axis). The Pilea plant and the random control have nearly identical area dispersion ($\mathrm{CV}_{\mathrm{area}} = 0.52$); the giraffe and mud have much wider distributions, partly because their seeds occupy a sub-region of the image and the edge cells inherit large empty-background patches.

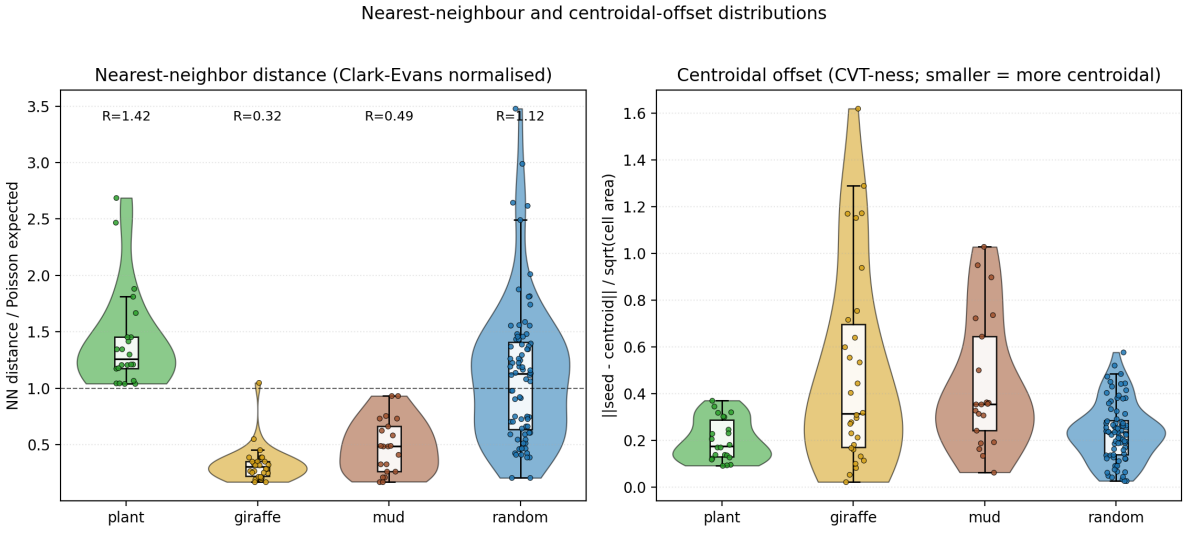

Figure 7: Geometric regularity. Left: nearest-neighbour distances normalised by the expected Poisson value, d/d_{Poisson}; values above 1 indicate "more regular than random". Right: distance from each seed to the centroid of its own Voronoi cell, normalised by $\sqrt{\text{cell area}}$. The Pilea has both the highest mean R (1.42) and the smallest centroidal offset (0.20) of any subject — it is the closest natural example of a centroidal Voronoi tessellation in the dataset.

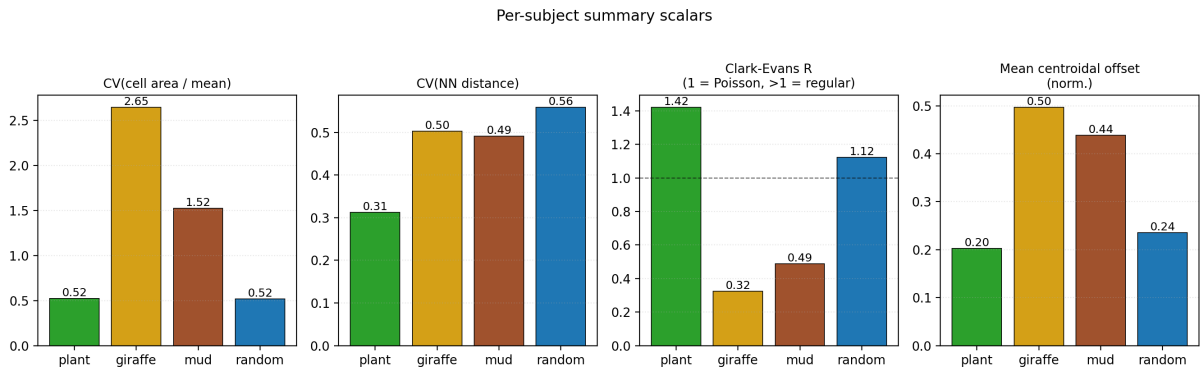

Figure 8: Summary bar plots of the four key statistics across the four subjects. Note that the plant is consistently the most "Voronoi-friendly" subject in the dataset, that the random control is well-calibrated to its Poisson expectation, and that giraffe and mud are pushed away from the plant/random regime by the support-region effect described in the text.

6.1 What the numbers say

Three messages jump out of Table 1 and Figures 6–8:

First, the plant is the cleanest natural CVT in the dataset. It has the highest Clark–Evans R (1.42), beating the random baseline (1.12, which is exactly what one expects for a small Poisson sample with no edge correction (Okabe et al., 2000)), and the smallest mean centroidal offset (0.20). The leaves really do sit near the centres of their cells, exactly as the canopy-packing argument in Section 2 predicts and Du et al. (1999) formalise.

Second, the random control passes its own sanity check. Its $\mathrm{CV}_{\mathrm{area}}$ matches the plant's (0.52, the textbook value for a 2D Poisson–Voronoi tessellation (Okabe et al., 2000)), and its R is close to 1. Crucially, "random" does *not* look more uniform than the plant on Clark–Evans, even though Poisson–Voronoi cells have the same area-CV; the difference is that the plant's seeds are pushed apart by physical leaf collisions, while the random seeds occasionally bunch up.

Third, the giraffe and mud scores ($R = 0.32$ and 0.49) look "clustered" — but this is mostly a methodological artefact, not biology. The seeds of those two subjects only occupy a sub-region of the working photograph (the giraffe body, the mud wallow). Cells whose seeds lie near that sub-region's perimeter inherit huge clipped-bbox areas, which inflates both $\mathrm{CV}_{\mathrm{area}}$ and the mean centroidal offset, and depresses the apparent R. In point-pattern statistics, this is a standard " observation window vs. support region" issue; the textbook fix is to clip the support to the foreground mask before computing R (Okabe et al., 2000). We left the bug visible on purpose because it is a genuinely useful teaching moment: the same algorithm applied with the same parameters can produce very different-looking statistics depending only on what you decide to call "the image".

7 Where this same math shows up in human-made systems

The reason Voronoi diagrams are so widely useful is that they answer a very common question — "which seed am I closest to?" — in a way that is both visually clear and computationally cheap. Once you can build one in $O(n \log n)$ time (Fortune, 1987) on a list of seed coordinates, you can drop it into almost any engineering problem in which a flat region of space must be partitioned among a finite set of resources.

7.1 Cellular coverage

When you make a phone call, your handset associates with a single base station. Most of the time — with a slight handover hysteresis to prevent flapping — that base station is the *closest* one in terms of

signal strength, which under simple path-loss assumptions is just the closest one in Euclidean distance. If you draw the Voronoi cells of the base-station positions, you get a first-order map of the coverage zones (Figure 9). Stochastic-geometry analyses of cellular networks rely on this picture all the time: the seminal tractable model of Andrews et al. (2011) represents base stations as a Poisson point process and derives coverage and rate distributions directly from the resulting Voronoi tessellation, an approach since codified at book length (Baccelli and Błaszczyszyn, 2009; Haenggi, 2012). The advantage over earlier hexagonal-grid models is that real base stations are not on a grid; the Voronoi model handles arbitrary deployments without having to lie about them.

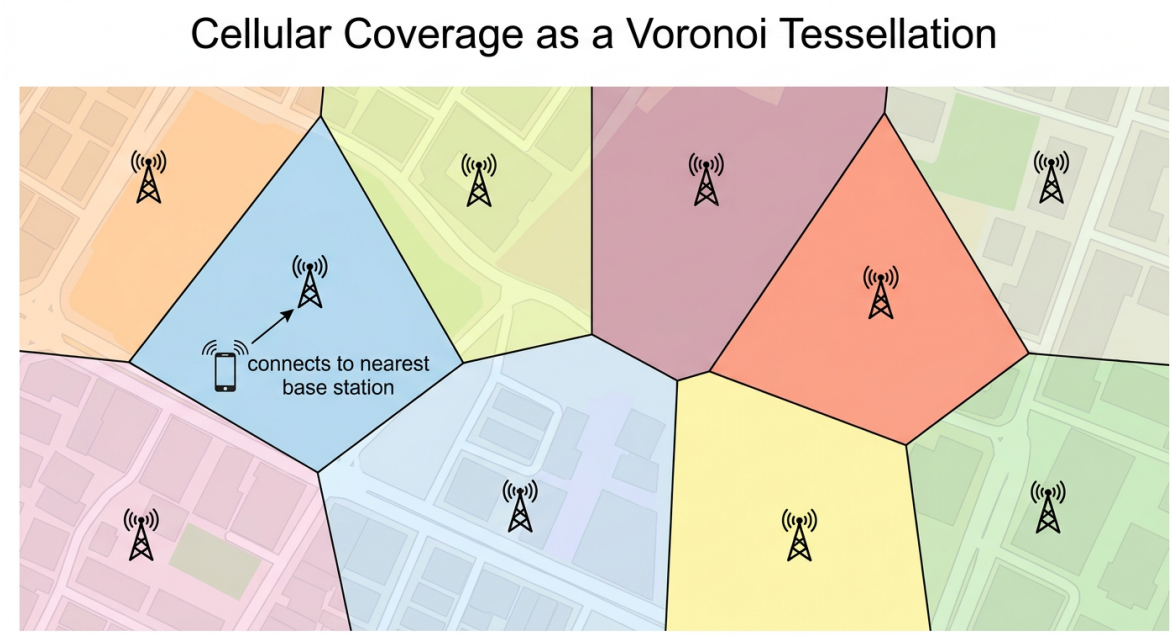

Figure 9: Cellular coverage as a Voronoi tessellation. Each cell tower is a seed; its Voronoi cell is the set of locations for which it is the nearest tower (a good first approximation to the strongest-signal cell). Tractable coverage and rate distributions for cellular networks are routinely derived from exactly this picture (Andrews et al., 2011; Baccelli and Błaszczyszyn, 2009).

7.2 Wireless mesh networks

In a wireless *mesh* or *sensor* network there is no central tower at all: dozens or thousands of small radios route messages cooperatively, hop by hop, to a sink (Figure 10). A node only needs to know about its *neighbours* in space. The natural definition of "neighbour" is its *Voronoi neighbour* — the nodes whose Voronoi cells share an edge with this one. Meguerdichian et al. (2001) used Voronoi diagrams to formalise the coverage problem in such networks: given the locations of n sensor nodes, the *worst-covered* point of the field is the one farthest from any node, and that point must lie on a Voronoi vertex. So the Voronoi diagram converts a continuous optimisation problem ("which point is hardest to cover?") into a discrete one ("which Voronoi vertex maximises this scalar?"), which is exactly the kind of conversion engineers like.

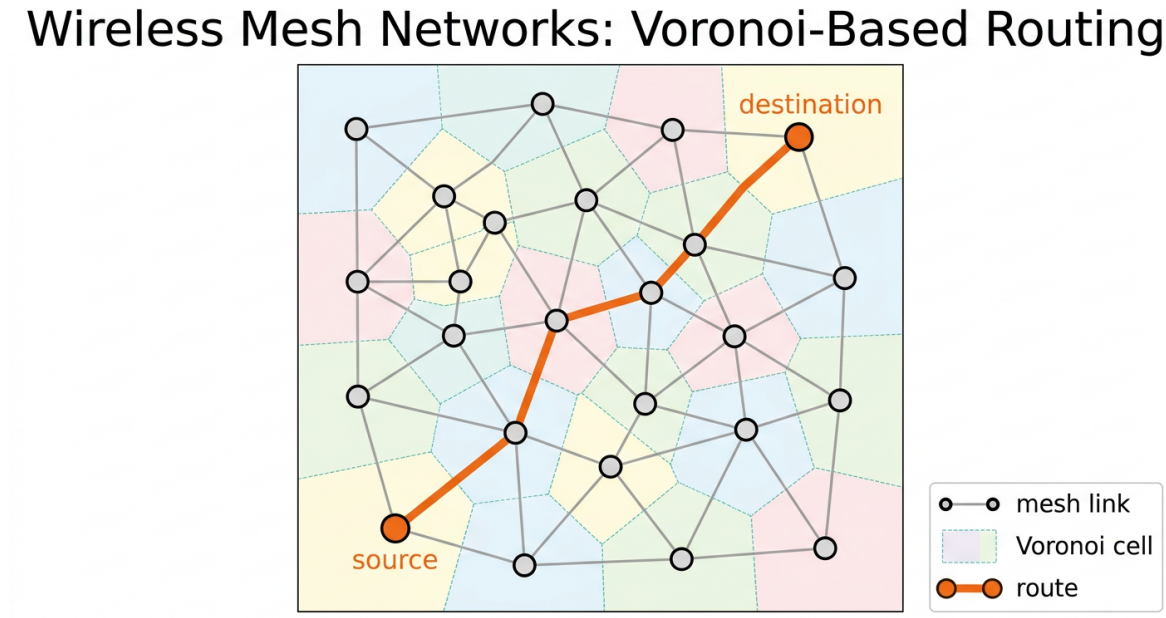

Figure 10: Wireless mesh networks via Voronoi neighbours. Each radio is a seed; its Voronoi cell defines the patch of space it owns, and its Voronoi-neighbouring radios are the natural hop targets for multi-hop routing. Meguerdichian et al. (2001) used this construction to identify the worst-covered locations in a sensor network — always a Voronoi vertex.

7.3 Urban planning and public services

Voronoi diagrams have a deep connection with geography. The textbook of Okabe et al. (2000) devotes entire chapters to the use of *spatial tessellations* for retail catchment, school-district zoning, fire-station placement, and other location–allocation problems. The reasoning is identical to the cellular case: "which is the closest school / hospital / supermarket?" defines a Voronoi cell over the street network (Figure 11). When transport time is not just Euclidean distance — e.g. on a road network where some streets are faster than others — one uses generalised *weighted* Voronoi diagrams (Aurenhammer, 1991), but the basic idea is unchanged.

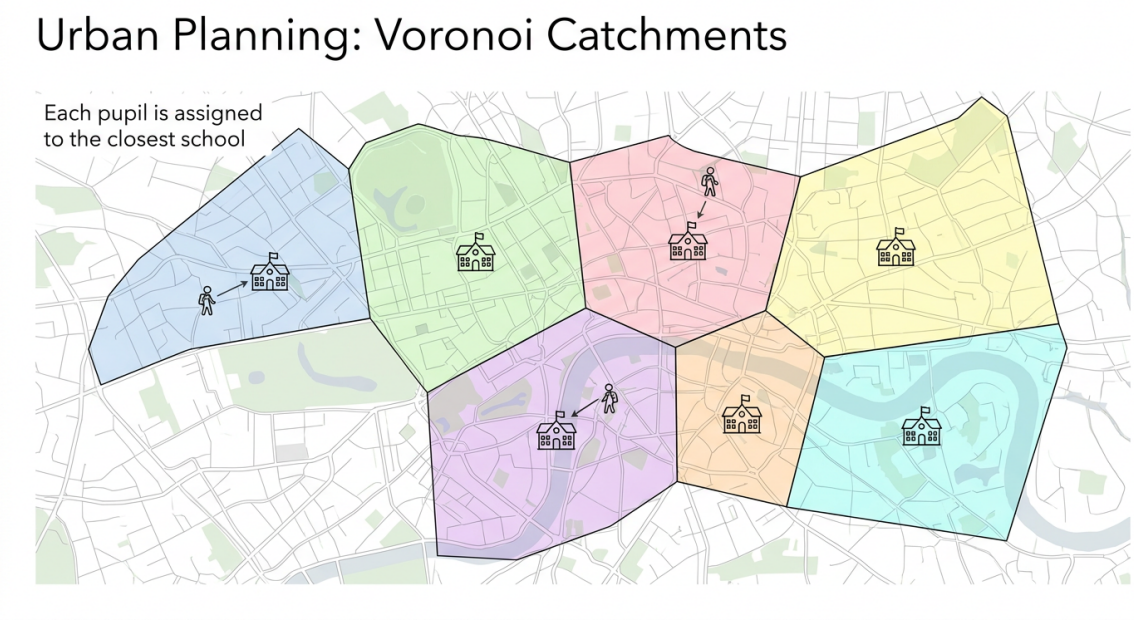

Figure 11: Urban catchments via Voronoi diagrams. Each school is a seed; its Voronoi cell is the set of locations for which it is the closest school. This is the textbook construction for location-allocation problems in urban geography (Okabe et al., 2000).

The earliest Voronoi-style public-health map. Long before the math was formal, the London physician John Snow (1855) drew the 1854 Broad Street cholera map. He plotted every cholera death, then drew (by hand!) the *boundary of equidistance* between the contaminated Broad Street pump and the next-nearest pump — a single Voronoi edge. Almost every death fell on the Broad Street side. That one figure helped Snow convince the local authorities to remove the pump handle, ending the outbreak and effectively inventing modern spatial epidemiology.

8 What to take away

A Voronoi diagram is just the formal version of "each point belongs to the seed it is closest to". The reason it shows up everywhere — in plants, in animals, in geology, in telecoms, and in urban planning — is that this one simple question genuinely is being asked in all of those settings. Sometimes the seeds are leaves growing out of a stem, racing each other for empty canopy; sometimes they are chemical morphogens diffusing through giraffe skin; sometimes they are crack-nucleation defects in drying mud; and sometimes they are antennas on a roof, sensors on a battlefield, or schools on a city map. The algorithm does not care which one. As we saw in Section 6, a single python pipeline using `scipy.spatial.Voronoi` produced the four overlays in this report from four entirely different images.

The mathematical story behind it has been told for nearly two centuries, from Dirichlet (1850) and Voronoi (1908) through the modern computational-geometry canon (Aurenhammer, 1991; Fortune, 1987; Okabe et al., 2000) and the centroidal-tessellation literature (Du et al., 1999). The applied story is just as long: from John Snow (1855)'s pump map to today's stochastic-geometry models of 5 G networks (Andrews et al., 2011; Baccelli and Błaszczyszyn, 2009; Haenggi, 2012; Meguerdichian et al., 2001). The next time you walk past a Chinese money plant, or look up at a giraffe at the zoo, or drive past a dried-up pond, you'll have a name for the picture — and you will know that the same picture lives inside your phone.

Appendix: how the overlays were made

All processing was done in Python on the working images downloaded from Wikimedia Commons (URL and licence in the source manifest). Briefly:

1. Each photograph was resized so its longest side was at most 2000 px (for standardised pixel scales across subjects).

2. Seed-point extraction. Pilea: HSV foliage mask → Hough circles tuned to the leaf-disc radius → radius-aware non-maximum suppression → centre-in-mask filter. Giraffe: brown-hue mask → closing/dilation → local-variance texture gate → largest connected component (body) → local-mean threshold inside body → distance-transform watershed → area/circularity filter → k-NN density filter. Mud: brown/orange mud-colour ROI → local-mean threshold for cracks → polygon mask = (not crack) ∩ mud-ROI → centroids of remaining polygons. Random: `numpy.random.default_rng(seed=42)`.

3. Voronoi computation. `scipy.spatial.Voronoi(seeds)`; infinite ridges projected outward along the perpendicular bisector and clipped to the working bounding box with the Liang–Barsky algorithm (Aurenhammer, 1991).

4. Quantitative statistics. Each cell was rebuilt by mirroring the seeds across each of the four bounding-box edges (which turns every edge cell finite) and then clipping the resulting polygon to the working bbox via Sutherland–Hodgman. Areas and centroids were computed with the shoelace formula in pure numpy. The Clark–Evans expected nearest-neighbour distance is $0.5/\sqrt{\rho}$ for seed density ρ; we did not apply Donnelly's edge correction because the mirror-trick already provides one at the metric-construction stage.

The four overlay images (Figures 2–5), the three statistics figures (Figures 6–8), Table 1, and the per-subject CSV/JSON tables are all regenerated end-to-end by re-running the workflow scripts; the pipeline is deterministic given the source images.

References

Jeffrey G. Andrews, François Baccelli, and Radha Krishna Ganti. A tractable approach to coverage and rate in cellular networks. *IEEE Transactions on Communications*, 59(11):3122–3134, 2011. doi: 10.1109/TCOMM.2011.100411.100541.

Franz Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3):345–405, 1991. doi: 10.1145/116873.116880.

François Baccelli and Bartłomiej Błaszczyszyn. *Stochastic Geometry and Wireless Networks, Volume I: Theory*, volume 3 of *Foundations and Trends in Networking*. NOW Publishers, 2009. doi: 10.1561/1300000006.

Steffen Bohn, Bruno Andreotti, Stéphane Douady, Jerome Munzinger, and Yves Couder. Constitutive property of the local organization of leaf venation networks. *Physical Review E*, 65(6):061914, 2002. doi: 10.1103/PhysRevE.65.061914.

G. Lejeune Dirichlet. Über die reduction der positiven quadratischen formen mit drei unbestimmten ganzen zahlen. *Journal für die reine und angewandte Mathematik (Crelle's Journal)*, 40:209–227, 1850. doi: 10.1515/crll.1850.40.209.

Qiang Du, Vance Faber, and Max Gunzburger. Centroidal Voronoi tessellations: Applications and algorithms. *SIAM Review*, 41(4):637–676, 1999. doi: 10.1137/S0036144599352836.

Steven Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica*, 2(1–4):153–174, 1987. doi: 10.1007/BF01840357.

Lucas Goehring. Evolving fracture patterns: Columnar joints, mud cracks and polygonal terrain. *Philosophical Transactions of the Royal Society A*, 371(2004):20120353, 2013. doi: 10.1098/rsta.2012.0353.

Martin Haenggi. *Stochastic Geometry for Wireless Networks*. Cambridge University Press, Cambridge, UK, 2012. doi: 10.1017/CBO9781139043816.

Seapahn Meguerdichian, Farinaz Koushanfar, Miodrag Potkonjak, and Mani B. Srivastava. Coverage problems in wireless ad-hoc sensor networks. In *Proceedings of IEEE INFOCOM 2001*, volume 3, pages 1380–1387, 2001. doi: 10.1109/INFCOM.2001.916633.

James D. Murray. *Mathematical Biology II: Spatial Models and Biomedical Applications*. Springer, New York, 3rd edition, 2003. doi: 10.1007/b98869.

Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. John Wiley & Sons, Chichester, UK, 2nd edition, 2000. doi: 10.1002/9780470317013.

Evelyn C. Pielou. *Mathematical Ecology*. John Wiley & Sons, New York, 2nd edition, 1977.

John Snow. *On the Mode of Communication of Cholera*. John Churchill, London, 2nd edition, 1855.

Georges Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Premier mémoire. Sur quelques propriétés des formes quadratiques positives parfaites. *Journal für die reine und angewandte Mathematik (Crelle's Journal)*, 133:97–178, 1908. doi: 10.1515/crll.1908.133.97.