Chapter 2

What the comparison must preserve

The room-temperature probe T-001 has a retained reading of 22.4 °C in Chapter 1. Suppose it measures the room again a minute later and reports the same value. A duplicate-message filter must keep both measurements. A maintenance program puts both in the same probe history so it can check for drift. Each program needs a different test for sameness.

I have spent years building systems that make sameness decisions across very different domains: spacecraft telemetry (keeping track of system state across format translations), viral sequences (tracking antigenic identity across mutations), payment rails (reconciling transactions across heterogeneous ledgers), and face recognition (matching identities across images and videos). In every case, I started with some ad hoc similarity heuristic—match these fields, threshold this score, cluster these embeddings—and refined it until the error rates looked acceptable on test data. It often worked well enough in the narrow setting it was designed for. But it always felt like duct tape. The problem was not that the heuristics failed catastrophically. The problem was that they did not *compose*. Every new format, every new sensor, every new embedding model meant starting over.

The difficulty becomes precise when we ask what a permitted change leaves invariant. Begin with an object we can turn in our hands, then carry that question into the instrument network. We will test three proposed rules. Physical identity must survive a changed appearance. A grouping of stations must satisfy a communication rule. A predictive state must retain the information required for future observations. Each rule specifies what its mathematical comparison must preserve.

A question has to choose its invariants

Photograph a mug with its handle to the right. Turn it and take a second photograph. Then swap in a second mug of the same design, put its handle to the right, and photograph it. The first two images show one mug from different sides. The first and third can look alike while showing different mugs.

The task therefore needs an *invariant*: a property the permitted change leaves unchanged. A warehouse may care about the serial number of the mug's package. A manufacturer may care only about model and dimensions. A photographer may care about the image itself. These are legitimate comparisons of different objects of study. They are not interchangeable uses of the word same.

Write O for the set of observations and $I : O \to U$ for the ideal identity assignment in a specified task. Exact identity in this model means $I(o) = I(o')$. If the identity assignment is unavailable, a detector estimates that statement. If I'm tracking my mug, the detector's yes is a guess about which mug it saw. That guess doesn't turn the second mug into mine.

A detector often substitutes numerical closeness for exact identity. Test that rule on three readings from a position sensor. The relation $|x - y| \le \tau$ is symmetric and reflexive but generally not transitive. With $\tau = 1$, the readings $0, 0.75, 1.5$ satisfy the first two adjacent comparisons but fail the end-to-end comparison. We cannot take its equivalence classes without supplying another

rule. Connected-component clustering supplies transitivity by chaining, and can thereby group points much farther apart than τ.

Proposition 2.1 (Tolerance chaining)**.** *If a sequence $x_0, \ldots, x_k$ in a metric space satisfies $d(x_{j-1}, x_j) \le \tau$, then $d(x_0, x_k) \le k\tau$. For fixed $\tau > 0$, no bound independent of k follows from these hypotheses alone.*

Proof. Apply the triangle inequality along the sequence. The real-line example $x_j = j\tau$ attains the bound. □

A part on the rail advances 0.75 twice. Each move passes a one-unit check between readings, though the part is now 1.5 from its starting mark. The tracker can keep following the same part. A travel alarm needs that starting mark.

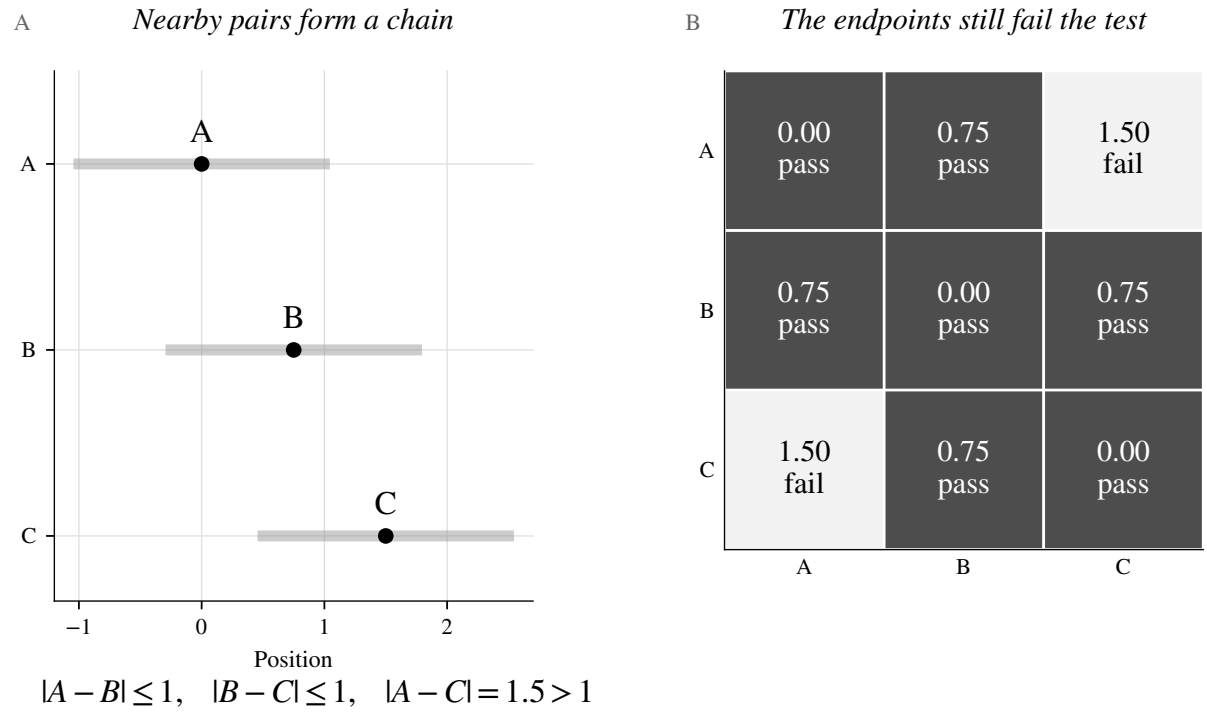

Figure 2.1: Three positions under a threshold of one unit. Each interval shows positions admitted by one observation. The complete pairwise table exposes the failure of transitivity: the two neighboring comparisons pass while the endpoints fail.

The mug showed why a similar appearance cannot supply individual identity. A network operator faces a related design choice when several stations are processed as one unit. The operator must first say what makes such a group admissible. We will require its members to exchange records without routing through stations outside the group. This gives a concrete rule whose consequences we can count.

Which groups can exchange a record?

Put four labeled monitoring stations on a desk. Keep the stations and their labels fixed. First connect them in a chain, then add a link joining the endpoints into a cycle, and finally connect every pair. The individual objects have not changed. In the chain, stations 1 and 3 can exchange

a record through station 2. If we put 1 and 3 in one group and leave 2 and 4 in separate groups, 2 is outside the first group. That partition becomes admissible when the complete graph adds the direct link from 1 to 3.

A connected partition divides the stations into nonempty blocks, each inducing a connected subgraph. For the chain $1-2-3-4$, every block is a contiguous segment. Choosing which of the three chain edges to cut gives $2^3 = 8$ partitions. On the four-cycle, there is one all-station block. Cutting two, three or four edges gives six, four or one additional partitions, for twelve in all. For the complete four-vertex graph every set partition is allowed, giving the Bell number $B_4 = 15$.

This finite example, developed in the later geometric-partition work [30], puts a limit on what a list of objects alone can explain. Relations constrain which collections count as coherent units. A sameness criterion that ignores the relation graph can merge groups that have no allowed internal connection.

Proposition 2.2 (Adding connections preserves existing connected partitions)**.** *If two graphs G, H have the same vertex set and every edge of G is an edge of H, every connected partition of G is a connected partition of H.*

Proof. Each block connected by edges of G remains connected after more edges are available. The vertices and blocks are unchanged, so the same partition satisfies the condition in H. □

Connectivity determines which groups are allowed

The same four vertices form a path, a cycle and a complete graph.

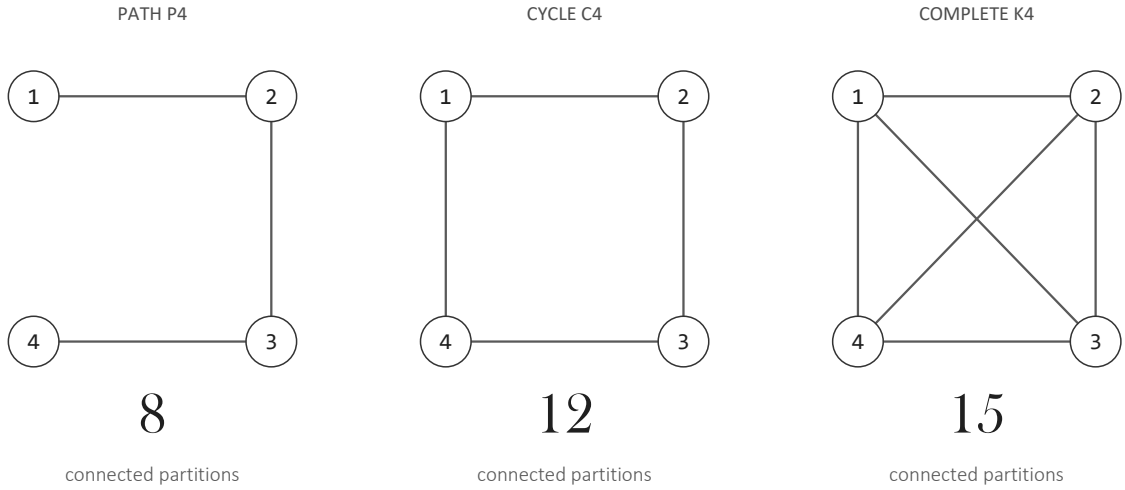

Adding edges preserves every previously connected block and may admit new ones.

A Connected groups: partitions 1 through 8

PARTITION	P_4	C_4	K_4
1 | 2 | 3 | 4	yes	yes	yes
12 | 3 | 4	yes	yes	yes
2 | 13 | 4	—	—	yes
2 | 3 | 14	—	yes	yes
1 | 23 | 4	yes	yes	yes
123 | 4	yes	yes	yes
23 | 14	—	yes	yes
1 | 3 | 24	—	—	yes

B Connected groups: partitions 9 through 15

PARTITION	P_4	C_4	K_4
13 | 24	—	—	yes
3 | 124	—	yes	yes
1 | 2 | 34	yes	yes	yes
12 | 34	yes	yes	yes
2 | 134	—	yes	yes
1 | 234	yes	yes	yes
1234	yes	yes	yes

Figure 2.2: *The objects stay fixed while their permitted groupings change.* The study continues on the following figure page.

Connectivity determines which groups are allowed

The same four vertices form a path, a cycle and a complete graph.

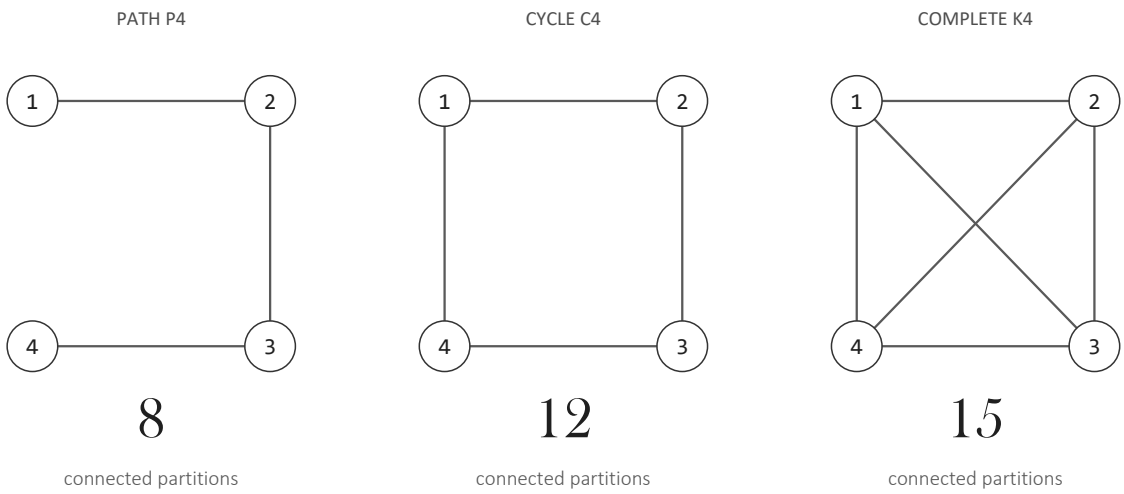

Adding edges preserves every previously connected block and may admit new ones.

C More edges permit more connected groupings

D Adding edges cannot undo an existing grouping

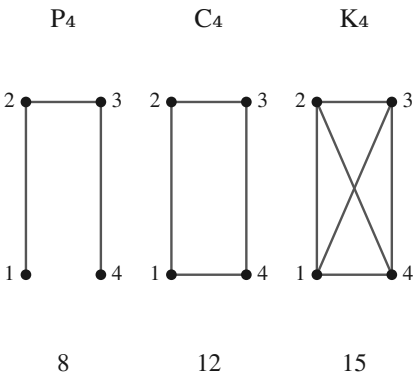

Connected-block partitions: $8 \leq 12 \leq 15$.

Keep the same vertices and partition.

An existing connected block stays connected after adding edges.

Some previously disconnected blocks become allowed.

Figure 2.2: *The objects stay fixed while their permitted groupings change (continued).* All fifteen partitions of four labeled vertices are classified against a path, a cycle and a complete graph. The eight, twelve and fifteen accepted partitions are obtained by explicit connectivity checks. The drawing makes the relation-dependent definition inspectable.

Chapter 14

Symmetry, quotients and two sheets

The realization maps told us how descriptions correspond. They also force a sharper question: can distinct descriptions represent exactly the same state? A bar viewed without a marked end makes that identification visible. We calculate which distances and observations survive it, and which information a record of motion must retain.

Remove the arrowhead

Place a straight bar on a turntable under an overhead camera. Paint an arrow along the bar so the picture shows which end is marked. Now cover the paint. The camera can still measure the line the bar lies on. A half-turn brings that line back to its starting orientation and swaps the ends.

Represent a planar unit arrow by $u(\theta) = (\cos\theta, \sin\theta)$. Its unoriented axis identifies u and $-u$, so θ and $\theta + \pi$ describe the same line. The quotient is the real projective line $\mathbb{RP}^1$. The projection $S^1 \to \mathbb{RP}^1$ has exactly two representatives over each line. Small neighborhoods separate into two disjoint arcs, each mapping homeomorphically onto the same neighborhood of lines. That is a double cover.

A *The lifted arrow changes sign*

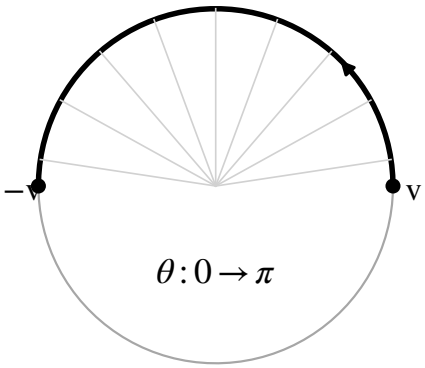

B *The physical bar returns*

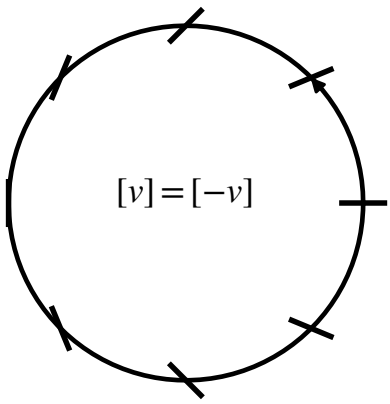

$S^1 \longrightarrow \mathbb{RP}^1, \qquad v \longmapsto \{v, -v\}$

The base loop closes; its continuous lift ends on the other sheet.

Figure 14.1: The two-sheet cover of unoriented axes. The left path travels from v to $-v$ on the representative circle. Its image on $\mathbb{RP}^1$ completes a loop, drawn using the circle coordinate 2θ. Small line segments at sampled base points show their unoriented axes.

The observable must respect the identification. Relative to a fixed unit arrow a, the squared inner product $(a \cdot u)^2$ is unchanged by $u \mapsto -u$. The signed inner product $a \cdot u$ is not. It therefore answers a question about an oriented representative, not just the underlying line.

If θ is the arrow angle relative to a, the invariant score is $\cos^2 \theta$. The expression $\cos^2(\theta/2)$ changes under $\theta \mapsto \theta + \pi$ and is not an observable of this line quotient. In the aircraft example the half-angle belongs to a different map and a different angle convention. Writing the domain and the covering map keeps both constructions intelligible.

A lifted journey remembers the sheet

Turn the bar through a half-turn while tracking the painted arrow. Its unmarked axis returns to its starting orientation. The arrow now points in the opposite direction. The endpoint on the base has returned; the endpoint on the cover has changed sheet.

This is not a failure of identity of the base object. It is retained information about the path. If an instrument reports only the axis, the two endpoints agree. If an experiment continuously tracks the arrow, it also knows which lift was followed. The measurement protocol determines which question the record can answer.

For a general covering $p : \widetilde{M} \to M$, a continuous path in M has a unique lift after choosing its initial representative. One can construct it by dividing the path into intervals lying in evenly covered neighborhoods and using the appropriate inverse on each interval. The local inverses agree at the joining endpoints, which fixes the continuation. Standard covering theory develops the associated fundamental-group description [5].

A camera watching a rotating bar

An overhead camera sees the bar on the turntable with the paint covered. What angle should its software report when the same axis could be described by either of two arrows? We will give the axis a coordinate and derive its distance rule.

This is a useful place to slow down. The double cover is not the assertion that every object has exactly two descriptions. A street address can have many spellings. A point on a globe can belong to many maps. The special feature of the bar experiment is a specified symmetry with exactly two alternatives, acting without a fixed point on the unit-arrow circle. The mathematics describes that experiment precisely.

Write the arrow as a complex number $z = e^{i\alpha}$. A convenient circular coordinate for its axis is

$$w = z^2 = e^{2i\alpha}.$$

If the camera reports w, recovering z gives two possibilities. At $w = 1$ they are 1 and -1; at $w = i$ they are $e^{i\pi/4}$ and $-e^{i\pi/4}$. The choice cannot be made from one unmarked picture. An earlier marked picture can supply the missing initial choice, and sufficiently frequent subsequent observations can continue it.

Proposition 14.1 (The metric belongs to the object)**.** *Equip the arrow circle with $d\alpha^2$. The quotient metric on the axis circle, in the coordinate $\beta = 2\alpha$, is $\frac{1}{4}d\beta^2$. Its circumference is π, and the projection is a local isometry.*

Proof. On an arc short enough to choose one inverse, $d\alpha = \frac{1}{2}d\beta$. Substitution into the arrow metric gives the stated metric. Integrating the quotient line element as β runs from 0 to 2π gives $\int_0^{2\pi} \frac{1}{2} d\beta = \pi$. The length of every sufficiently short lifted arc equals the length of its image. □

The connection with Chapter 16 can also be calculated. Choose the transport law $\dot{a} + (i/2)\alpha(\dot{\gamma})a = 0$. Its curvature is zero, since $d\alpha = 0$, but one circuit multiplies a by $\exp(-i\pi) = -1$. The coefficient $1/2$ specifies that transport rule; the hole alone does not determine a phase. Position, integer winding and sign are three different records. As with maps of a globe, local agreement becomes a global description only after we specify what must agree when the maps close around a journey.

The annulus experiment shows what the disk assumption contributed to GIGI's harmonic construction: local differential agreement must also meet the global conditions of its domain. Laboratory 6 returns to the recorded fingerprint with that distinction in view and makes the completion and harmonic checks reproducible.

How the blackboard graph folds

The movie that put harmonic functions in my head left a different question on a hallway blackboard [38]. Picture four stations connected by roads. Stations 1, 2 and 4 form a triangle, and two separate roads run from 2 to 3. How many walks of a given length take you from 1 to 3? You can revisit a station or reuse a road. Each crossing counts as a step, and either road into 3 counts as a separate choice.

Use station order $1, 2, 3, 4$. Entry A_{ij} counts roads from i to j. Matrix multiplication chains those choices, so A^3 counts three-step walks.

$$A = \begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 2 & 1 \\ 0 & 2 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{pmatrix}, \qquad A^3 = \begin{pmatrix} 2 & 7 & 2 & 3 \\ 7 & 2 & 12 & 7 \\ 2 & 12 & 0 & 2 \\ 3 & 7 & 2 & 2 \end{pmatrix}.$$

Think of the generating function as a tally sheet with one row for each walk length. The count for three steps goes in front of z^3, and the count for length n goes in front of z^n. The matrix geometric series gives $F_{ij}(z) = \sum_{n\geq 0}(A^n)_{ij} z^n = [(I - zA)^{-1}]_{ij}$ as a power series about zero.

Try folding the drawing along the dashed line in Figure 24.9. Station 1 lands on station 4, and 2 and 3 stay put. On the smaller map I give that pair one label, S.

From either station in S, one road takes me to 2 and one takes me to the other station in S. The choices into each group match, so I only need three labels in my tally. Station 3 keeps its own label, so I can still count exactly the walks that end there.

From 2 there are two ways into S and two into 3. In order $S, 2, 3$, the folded matrix is

$$B = \begin{pmatrix} 1 & 1 & 0 \\ 2 & 0 & 2 \\ 0 & 2 & 0 \end{pmatrix}.$$

In row i of P, put a 1 in the column for station i's group and zeros elsewhere. Then $AP = PB$ matches the one-step counts into each group. Iterating gives $A^n P = PB^n$. Since 3 is alone in its group, this certifies the count ending there for every length.

Let a, b, c be the series for walks ending at 3 from $S, 2, 3$. Counting by the first step gives $a = z(a + b)$, $b = 2z(a + c)$ and $c = 1 + 2zb$. The 1 counts the zero-step walk already at 3. Eliminating b, c gives the same answer as the full inverse:

$$F_{13}(z) = \frac{2z^2(1 + z)}{1 - 7z^2 - 2z^3 + 4z^4} = \frac{2z^2}{1 - z - 6z^2 + 4z^3}.$$

GOOD WILL HUNTING / GRAPH SYMMETRY

The fold keeps every count to 3

Each road costs one step. The folded arrows count choices from one starting station.

01 Two maps of the same walks

Original stations

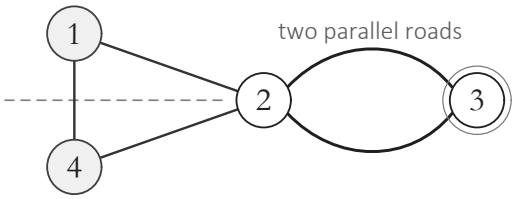

Mirror line fixes 2 and 3

Folded states

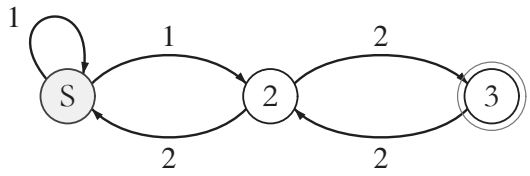

S groups 1 and 4. Its loop costs one step.

02 Exactly four steps from 1 to 3

Original stations	*Folded route and choices by step*	*Product*	*Walks*
1 → 4 → 1 → 2 → 3	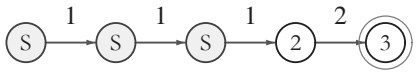	1 × 1 × 1 × 2	2
1 → 2 → 1 → 2 → 3 1 → 2 → 4 → 2 → 3	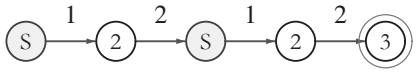	1 × 2 × 1 × 2	4
1 → 2 → 3 → 2 → 3	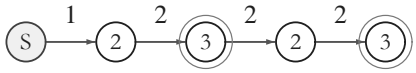	1 × 2 × 2 × 2	8
		2 + 4 + 8	14

03 Why the count survives

The group map P

	S	2	3
1	1	0	0
2	0	1	0
3	0	0	1
4	1	0	0

Rows 1 and 4 map to S.

$$AP = PB \quad \Rightarrow \quad A^nP = PB^n$$

The mode that reverses sign

	1	2	3	4
u	+1	0	0	−1
Au	−1	0	0	+1

The pattern changes sign and stays zero at 3.

Its factor $1 + z$ cancels in the walk series.

Figure 24.9: The reflection groups stations 1 and 4. The register pairs each original route with its folded sequence and tallies all 14 four-step walks.

Start at station 1 and give yourself exactly two steps to get to 3. You have to go through 2, then pick either of the two roads to 3. With three steps, you have time to go through 4 first. On the folded drawing those routes are $S \to 2 \to 3$ and $S \to S \to 2 \to 3$. The loop is the real crossing from 1 to 4. You have used up a step even though your label on the small map still says S. Either route still gives you two choices for the last road.

Now give yourself four steps. You have to be at 2 with one step left. You can cross inside S twice, head to 2, and pick either road to 3. That's two walks. Or go straight to 2, take a side trip to either station in S and back, then pick either road to 3. That's four. Or go straight to 2, then go to 3, back to 2, and out to 3 again. Each of those last three crossings offers two roads. That's eight. A walk can pass through 3 early as long as it ends there. Add up the three counts and you get fourteen.

GIGI's `good_will_hunting_walks` example checks the counts on both graphs against the rational series. They begin $0, 0, 2, 2, 14, 18, 94$.

The swap acts like the half-turn in Chapter 14. The pattern with value $+1$ at 1 and -1 at 4 changes sign under the swap. Multiplying by A also reverses that pattern, whose value at station 3 is zero. This is the mode behind the canceled factor $1 + z$. Now suppose someone is waiting for you at station 1 and you have exactly one step left. If you're at 4, you can cross over and meet them. If you're already at 1, every road takes you away. The folded label S tells you which pair of stations you're in. To answer this question you need to know which one.

What the fingerprint keeps

The fold kept the walk counts to station 3. Now go back to the wine table and make a copy of every row. Each bin has twice as many readings and the table is twice as long. The proportions stay the same, so with the same bin edges, both tables give the same polynomial. Shuffling the rows or moving a reading within its bin also leaves the polynomial alone. If you need the original order, the sample count or the exact readings, keep that information in the record.

The uncertainty disk answers a second question. With just a few values of the polynomial in hand, what can we say at a new probe? The disk contains all values an analytic function bounded by one can take there while matching the supplied values.

Laboratory 6: reproduce a GIGI completion calculation

Start with the thirty-two wine-alcohol bin counts in Table 24.2. Divide them by $N = 6497$ to obtain the coefficients p_k. Check that every coefficient is nonnegative and the sum is one. Evaluate $H(z) = \sum p_k z^k$ using Horner's rule and retain the original counts separately.

Prepare the seven observations. Use $z_0 = 0$ and six nodes $0.65e^{2\pi ij/6}$ for $j = 0, \dots, 5$. Compute values from the full coefficient vector. Compare the results with the printed observation table at its displayed precision. The full histogram is used in this preparation step; the later completion query receives only the nodes and values.

Form the compatibility matrix. Compute $P_{ij} = (1 - w_i\overline{w_j})/(1 - z_i\overline{z_j})$. Check its Hermitian residual and smallest eigenvalue. Small residuals should be interpreted against the numerical precision and matrix scale. Do not make an inconsistent matrix positive by silently deleting a negative eigenvalue.