Why random tilings don't look random

by Jim Propp (UMass Lowell)

April 16, 2011

Slides for this talk are on-line at
http://jamespropp.org/vershik.pdf

Acknowledgments

This talk describes past work of Bob Burton, Henry Cohn, Noam Elkies, Rick Kenyon, Greg Kuperberg, Michael Larsen, Andrei Okounkov, Robin Pemantle, Scott Sheffield, David Wilson, and myself, as well as others.

Hexagon tilings

There are exactly

$$\prod_{i=1}^{20}\prod_{j=1}^{20}\prod_{k=1}^{20}\frac{i+j+k-1}{i+j+k-2} \approx 1.6\times 10^{136}$$

ways to tile a regular hexagon of side-length 20 with lozenges (rhombuses with an internal angle of 60°) of side-length 1.

Here is one of them, chosen uniformly at random from the set of all such tilings (I'll say how this was done later on):

http://jamespropp.org/hexagon.gif

Here's an even larger random tiling, this one of order 64 (figure borrowed from Ken McLaughlin's web-page):

http://math.arizona.edu/∼mcl/Tiling1.jpg

Observations

What do we notice?

- Frozen behavior near the corners (all lozenges line up)
- Round-ish interface between frozen and non-frozen region
- Non-homogeneity inside the non-frozen region (the particular orientation common to all the lozenges in a frozen region is also dominant in nearby parts of the central non-frozen region)
- The tiling looks like a surface in 3-space, with square faces that appear as lozenges in projection

The last of these doesn't even look like mathematics, but in some sense it is the most important observation of all; for instance, it's the key to understanding the third observation.

Tilings as surfaces in $2+1$ dimensions

We can view a lozenge-tiling of a regular hexagon of order n as the projection of a cubical surface spanning a non-planar hexagon with sides of length n and with right angles at its vertices.

Equivalently, glue three n-by-n squares together at right angles to form three adjacent sides of an n-by-n-by-n cube, forming a cubical frame whose boundary is a non-planar hexagon.

Orient the corner of this cubical frame downwards, and start to fill the frame with 1-by-1-by-1 cubes with their corners pointing downwards.

Viewed from above, the visible faces of the small cubes, together with the visible squares of the cubical frame that holds the small cubes, form a surface whose boundary is the non-planar hexagon.

Paths in $1 + 1$ dimensions

Let's go down a dimension:

Glue two line segments of length n together at right angles to form two adjacent sides of an n-by-n square.

Orient the corner of this frame downwards, and start to fill the frame with 1-by-1 squares with their corners pointing downwards.

The top edges of the small squares, together with the visible edges of the frame that holds the small squares, form a lattice-path.

Lattice paths and partitions

These sorts of collections of squares, rotated 45 degrees (or as some prefer 135 degrees) clockwise, are examples of Young diagrams, introduced in the study of partitions of numbers.

Partitions in a box

Our lattice paths correspond to all the different n-tuples of integers $(a_1, a_2, \ldots, a_n)$ with $n \geq a_1 \geq a_2 \geq \cdots \geq a_n \geq 0$.

There are exactly

$$\frac{(2n)!}{n!n!}$$

such n-tuples.

To see why, note that each of our lattice paths is some concatenation of n northeast steps and n southeast steps, which can occur in any order.

Partitions set free: A testimonial digression

One can get rid of the box, and look at infinite sequences of non-negative integers $(a_1, a_2, \dots)$ with $a_1 \geq a_2 \geq \dots$.

There are infinitely many such sequences, even after we impose the condition that the sequence must consist of 0's from some point onward, but we can look at just the sequences that sum to n.

That is, we are looking at a random Young diagram of size n sitting in an infinite quadrant.

What does a random one look like?

Partitions set free: A testimonial digression

Vershik figured this out: the limit shape is given by the graph of the relation

$$e^{(-\pi/\sqrt{6})x} + e^{(-\pi/\sqrt{6})y} = 1.$$

See "Statistical mechanics of combinatorial partitions and their limit shapes", *Funct. Anal. Appl.* **30** (1996), 90–105;

http://www.pdmi.ras.ru/∼vershik/statph.ps

For a picture of the curve, see Figure 2 (on page 3) of "Identities arising from limit shapes of constrained random partitions" by Dan Romik:

http://www.stat.udavis.edu/∼romik/home/Publications_files/

Back to the box

If we look at a random partition whose Young diagram fits in an n-by-n box, the shape is likely to be very close to an isosceles right triangle.

To see why, recall that the bounding lattice path, read from upper left to lower right, is a random concatenation of n right-steps and n down-steps.

If we draw balls from an urn containing n black balls and n white balls, at any stage, the number of black balls we've drawn minus the number of white balls we've drawn is likely to be close to 0.

For the same reason, the lattice-path stays close to the diagonal line joining its start-point and end-point.

Deviations

Why don't we see macroscopic deviations from the straight line?

Assume for convenience that n is a multiple of 3. If we take a random lattice path of length $2n$ from $(0, n)$ to $(n, 0)$, why is it very unlikely to go through $(n/3, n/3)$?

Binomial coefficients

One answer is: $\frac{(2n)!}{(n)!(n!)}$ is much greater than $\frac{n!}{(n/3)!(2n/3)!}$ times $\frac{n!}{(2n/3)!(n/3)!}$, so the total number of paths from $(0, n)$ to $(n, 0)$ is much greater than the number of paths from $(0, n)$ to $(n/3, 2n/3)$ times the number of paths from $(n/3, 2n/3)$ to $(n, 0)$, which is equal to the number of paths that go from $(0, n)$ to $(n, 0)$ by way of $(n/3, 2n/3)$.

So, most of them can't look like that!

Entropy

Phrasing this entropically: If you arrange n black balls and n white balls in a row, uniformly at random subject to the condition that among the first n balls $1/3$ are black and $2/3$ are white and among the last n balls $2/3$ are black and $1/3$ are white, the entropy of the first half of the sequence of colors is at most np, as is the entropy of the second half, where $p = -\frac{1}{3}\log\frac{1}{3} - \frac{2}{3}\log\frac{2}{3} < -\log\frac{1}{2} = 1$.

The total entropy of such a sequence is therefore at most $2np$.

On the other hand, a random sequence of n black balls and n white balls has entropy very close to $2n$, which is strictly greater than $2np$.

So, most of them can't look like that!

The local picture

Indeed, if we choose a random lattice-path and zoom in on a tiny piece of it, the local statistics will look IID; that is, a rightward step is equally likely to be followed by a rightward step or a downward step.

There is a slight amount of negative correlation, but when n is large it becomes invisible at a local level; that is, a tiny piece of the lattice path can't "see" that the whole lattice path has to eventually end up at $(n, 0)$ and not $(n+1, 1)$ or $(n-1, -1)$.

If we take the local picture in the limit as $n \to \infty$, we get the $(p = \frac{1}{2})$ Bernoulli process.

Warping the path

Return to the original (unrotated) picture, with a path that goes from $(-n, n)$ to (n, n) by increments of $(1, 1)$ and $(1, -1)$ (with n of each kind).

We've now seen that a random lattice path of this kind stays close to the line segment joining $(-n, n)$ to (n, n), and that the area between it and the frame $y = |x|$ is close to $n^2/2$.

But what if we imposed the condition that the area should be close to cn^2 for some $c \neq \frac{1}{2}$ in $(0, 1)$?

Random warped paths are locally Bernoulli

What we'd see is a lattice path that locally looks Bernoulli(p), but with p varying as one moves along the path. E.g., for $c < \frac{1}{2}$, we'd see a path that starts out with mostly increments of $(1, -1)$ and ends up with mostly increments of $(1, 1)$. Only in the very middle would you see $(1, -1)$ and $(1, 1)$ in equal proportions.

Random warped paths are globally deterministic

If we take a random warped path and send n to infinity (with c fixed), one finds empirically that the probability distribution concentrates on a single curve joining $(-1, 1)$ to $(1, 1)$.

This curve $y = f(x)$ must satisfy the Lipschitz condition $|f(y) - f(x)|/|y - x| \leq 1$ (because it is a probabilistic limit of lattice paths that also satisfy this condition).

It must also satisfy $f(-1) = f(1) = 1$ and $\int_{-1}^{1} f(t) = 1 + c$.

But among all such curves, which one could it be?

The entropy functional

For any candidate curve f, consider the entropy functional

$$V(f) = \int_{-1}^{1} H(f'(x))\ dx$$

where $H(t) = -\frac{1+t}{2} \log \frac{1+t}{2} - \frac{1-t}{2} \log \frac{1-t}{2}$. Here
$p = \frac{1+t}{2}$ is the local proportion of $(1,1)$ steps,
$q = \frac{1-t}{2} = 1 - p$ is the local proportion of $(1,-1)$ steps, and
$H(t) = -p \log p - q \log q$ is the entropy of a Bernoulli(p,q) process.

The variational principle

Variational principle: The f that maximizes entropy is the f that you see when you choose a warped lattice path at random.

The way one proves this is by showing at the same time that the local behavior is Bernoulli.

The **slope** of the entropy-maximizing function encodes the **local statistics** of the random path.

Back to $2 + 1$ dimensions

We might expect that, just as our lattice path in $1 + 1$ dimensions tended to a curve (indeed a straight line segment) in the continuum limit, our lattice surface in $2 + 1$ dimensions, rescaled by n, will converge to some sort of surface that sits inside the rescaled cube.

The limit-shape for the (unwarped) random lattice path is homogeneous (every part looks like every other), but the limit-shape for our lattice surface can't be!

Non-homogeneity

That's because the boundary of our lattice surface (the non-planar hexagon), rescaled by n and taken to its limit, **isn't flat**.

Moreover, the limit shape can't be flat in most places and then steep in a tiny neighborhood of the boundary, because the lattice surface satisfies a Lipschitz condition analogous to the one we saw in $1 + 1$ dimensions.

What we see in the corners: an informal argument

We are tiling a regular hexagon of side-length n with $3n^2$ lozenges.

Look at the left corner of the hexagon.

If you start by tiling it with the acute angles of two lozenges, this will force you to place other tiles, which will force you to place other tiles, etc., so that $2n$ tiles in all are placed. There are about K^{3n^2-2n} ways to complete the tiling, for some K.

But if you start by tiling the left corner of the hexagon with the obtuse angle of a single (vertical) lozenge, this forces no other placements of tiles, so there are about K^{3n^2-1} ways to complete the tiling.

Since $K^{3n^2-2n} << K^{3n^2-1}$, the overwhelming majority of tilings of the hexagon have a vertical lozenge there.

What we see in the middle

Here we see a random mixture of all three orientations of lozenges, occurring in equal proportion.

Locally, it looks like an ergodic process (symbolic $\mathbf{Z}^2$ action); indeed, it is the maximum entropy process supported on the set of lozenge tilings of the plane, much as the Bernoulli(1/2,1/2) process is the maximum entropy process supported on the set of doubly-infinite sequences of 0's and 1's.

What we see elsewhere

For all triples p, q, r of positive numbers with $p + q + r = 1$, there is a unique maximum entropy process supported on the set of lozenge tilings of the plane, subject to the constraint that the three orientations of lozenges must occur with respective densities p, q, and r, much as the Bernoulli$(p, 1-p)$ process is the maximum entropy process supported on the set of doubly-infinite sequences of 0's and 1's subject to the constraint that 1 and 0 occur with respective densities p and $1 - p$.

As (p, q, r) goes to $(1, 0, 0)$, $(0, 1, 0)$, or $(0, 0, 1)$, the process becomes a trivial process in which all the tiles have the same orientation; this is what one sees in the corners (more specifically, what one sees outside the circle inscribed in the hexagon).

The asymptotic global picture

The asymptotic (i.e., $n \to \infty$) global behavior of a random lozenge tiling of a hexagon of order n, is given by a surface whose "slope" in $\mathbf{R}^3$ encodes a triple (p, q, r) specifying the local statistics of the tiling (the preponderance of the three orientations of lozenges).

This surface (rescaled by n) is the graph of a function on the hexagon. This function is singular on the circle inscribed in the hexagon, and linear on each of the six components of the complement of the circle relative to the hexagon.

Figuring out where you are

The limit surface, with its six flat "petals" removed, has the property that for any two points on the surface, the normal directions are distinct, giving rise to distinct triples (p, q, r) and (p', q', r').

Consequently, if you woke up in a huge random lozenge-tiling (with the three orientations of lozenges given distinct colors) and you didn't know where you were, as long as you weren't in one of the six frozen regions, you could figure out your approximate location in the tiling by walking around and observing the local densities of the three orientations (colors) of tiles.

Domino tilings of Aztec diamonds

Much of this theory was first worked out for a different kind of tiling problem that was in some respects simpler: random domino tilings of Aztec diamonds.

http://tuvalu.santafe.edu/∼moore/aztec256.gif

Maximum entropy domino tiling process

What we see in the middle of an Aztec diamond is a tiling process first studied by Burton and Pemantle.

If we replace the Aztec diamond by a large square, we find that as long as one isn't too close to the boundary, the tiling has Burton-Pemantle local statistics.

http://faculty.uml.edu/jpropp/tiling/www/intro.html

So, random tilings sometimes *do* look random!

What matters is the way the *height function* behaves on the boundary.

(One can construct regions in the triangular lattices in which the height function along the boundary is flat; for those regions, a random lozenge tiling will be statistically homogeneous away from the boundary, with maximum-entropy statistics everywhere.)

Diabolo tilings of fortresses

A more complicated example of an exactly solved tiling model is random diabolo tilings of fortresses; see

http://jamespropp.org/fortress.pdf

Gaskets-and-baskets tilings of supergaskets

And here's one we still don't understand:

http://jamespropp.org/ice40.pdf

This is related to alternating sign matrices from algebraic combinatorics and to the square ice (or 6-vertex) model from statistical mechanics.

Random surfaces

In all the preceding examples, the connecting link is the way in which a 2-dimensional tiling can be viewed as a discrete surface in $2+1$ dimensions.

This discrete surface is also a key tool when one wishes to draw a sample from the uniform distribution on the set of all tilings of a particular size and kind, obtaining pictures like the ones discussed above.

How does one get these pictures?

The most versatile method is the coupling-from-the-past (CFTP) method of myself and David Wilson. Jason Woolever's implementation is at

http://jamespropp.org/applets/

The discrete-surface interpretation of tilings is a key feature of the problem that makes "monotone CFTP" applicable. See

http://dbwilson.com/exact/