

What do geometry and topology have to do with statistics? The answer here goes via medical imaging and astrophysics.

The Geometry of Random Images

Keith J. Worsley

The geometry in the title is not the geometry of lines and angles but the modern geometry of topology and shape. What has this to do with statistics? Some recent work has found some fascinating applications of a mixture of geometry, topology, probability, and statistics to some very pressing problems in newly emerging areas of medical imaging and astrophysics. Where is the link? Let us begin with a quick introduction to one

of the fundamental tools of topology, the Euler characteristic.

Topology: The Euler Characteristic

Named after Leonhard Euler (1707-1783), the most prolific mathematician of the 18th century, the Euler characteristic itself began with Euler's observation about polyhedra.

Recall that a polyhedron is a solid object bounded by plane faces, such as a cube. Euler realized that, if you count the faces (F), edges (E), and vertices (V) of a polyhedron, then $V-E+F=2$ no matter how the polyhedron is constructed.

A cube, for example, has $F=6$ faces, $E=12$ edges and $V=8$ vertices (see Fig. 1a) so that $8-12+6=2$. For a solid that consists of P polyhedra, stuck together on at least one common face, the slightly more general formula becomes $V-E+F-P=1$.

A little experimentation will convince you that this new formula works for all solids (see Fig. 1b)—well almost all. If the solid has a hole going right through it,

Front Cover: Two random surfaces (top left and right) are the zero contours of two simulated Gaussian random fields. Slices through these fields are shown below the surfaces, color coded from red (high values) to purple (low values). They intersect (bottom left) to produce random strings (bottom right), which are the zero contours of a chi-squared random field with two degrees of freedom. This is defined as the sum of squares of the two Gaussian random fields, and a slice through this field is shown below the strings.




Figure 1. The Euler characteristic (EC) of a solid is the number of vertices – edges + faces – polyhedra. For a single polyhedron (a) or several joined together (b), the EC = 1; if there is a hole in the solid (c) the EC = 0; and the EC decreases by 1 for each extra hole (d, e); in (f) the central cube is missing so the solid is hollow, in which case the EC = 2.

like a doughnut (see Fig. 1c), then the result no longer holds. In fact, the result is $V-E+F-P=0$ for any solid with just one hole.

Too bad! But this does not deter a good mathematician—far from it—it opens up vast new possibilities! What happens if there are two holes, like a figure 8 (see Fig. 1d)? Then it turns out that $V-E+F-P=-1$, and so on; each hole reduces $V-E+F-P$ by 1.

So now suddenly we have a fascinating new tool. We can *count* the number of holes in a solid using the formula $V-E+F-P$; it has the very interesting property that no matter how the solid is subdivided into polyhedra, the value of $V-E+F-P$ is *invariant*. Thus is born the field of topology: We define the *Euler characteristic* (EC) of a solid as simply $EC=V-E+F-P$

for any subdivision of the solid into polyhedra.

Thus the EC of a pretzel-shaped solid (Fig. 1e) is -2 : $+1$ for the solid part (the part you eat), and -3 for each of the three holes, giving -2 overall. Have we covered all possibilities? Not quite—if the solid is *hollow*, like a tennis ball, then surprisingly enough the EC is 2 (see Fig. 1f).

Think you've got it now? How about a solid shaped like a bicycle inner tube? Answer: The EC is 0, and if it has a puncture, then the EC is -1 .

One more slight generalization, which will prove to be extremely useful for practical applications: Suppose we have a set of disconnected solids. No problem—we can divide them all up into polyhedra, apply the formula, and sum

everything. Thus if the set consists of several disconnected components, with no holes, then the EC simply counts the number of such components.

Finally, to make it really useful, suppose we start off with a set in three dimensions with a smooth boundary, as opposed to a union of polyhedra; how then can we find the EC? The answer is to cover the set with a cubical lattice (any lattice would do, but a cubical one is the simplest). If the lattice is fine enough and the boundary is smooth everywhere, then it can be shown that the EC of the lattice inside the set equals the EC of the set itself.

To summarize, we now have a completely general tool, the EC, that tells us a lot about the *topology* of a set. If the set is composed of







Figure 3. How to find two surfaces that intersect in a knot. Start with any knot, such as a trefoil knot (a), which can always be represented as the boundary of a two-sided surface (b, c), its "Self-representation." Thicken the surface (d), cut it in half (e) and round off the edges (f). Push the two surfaces together (g) and the intersection (h) is the knot you want (i).

Figure 2. Different topologies and examples of their Euler characteristics: (a) A "meatball" topology: many disconnected components, each containing relatively few holes or hollows gives $EC = 21$; (b) A "sponge" topology: the components are connected together by a network of "bridges," to form many holes, gives $EC = -15$; (c) a "bubble" topology: surfaces surrounding isolated voids, giving $EC = 1$.






Figure 4. Plot of the observed EC of the set of high-density regions of the real galaxy data (points), and the expected EC from the formula (curve) for a Gaussian random-field model, plotted against the density threshold. The observed and expected are in reasonable agreement, confirming the Gaussian random field model: high thresholds produce a "meatball" topology (Fig. 2a); medium thresholds produce a "sponge" topology (Fig. 2b), and low thresholds produce a "bubble" topology (Fig. 2c).

many disconnected components, each containing relatively few holes or hollows (called a "meatball" topology in astrophysics) then the EC is large and positive (Fig. 2a). If the components are connected together by a network of "bridges" (rather like mozzarella cheese) to form many holes, then the EC is negative; this is called a "sponge" topology (Fig. 2b). If the network dominates to such an extent that the set consists of many surfaces surrounding isolated voids (rather like Swiss cheese), called a "bubble" topology (Fig. 2c), then the hollows dominate and the EC is once again positive. We are now ready for the applications.

Astrophysics: Large-Scale Structure of the Universe

One of the most fascinating problems in astrophysics is the large-scale structure of the universe. Ever since the discovery of galaxies, it has been realized that galaxies themselves group together to form clusters, and these group together to form superclusters. With more accurate observations and larger galaxy surveys, it became evident that galaxies were not scattered throughout the universe in a random fashion.

Instead there are large voids, completely free of visible matter, surrounded by strings and even

sheets of galaxies; one such has been named the "Great Wall" by astrophysicists. It is of great interest to build models of the formation of the universe that could explain such structure, and compare the results of these models to the structure we observe today (see Sidebar 1).

This is where the EC came in. In a series of articles in the *Astrophysical Journal*, starting in the mid-80s, Richard Gott and his colleagues at Princeton used the EC as a descriptive tool of the topology of large-scale structure. Astronomers had produced a 3-D map of all galaxies in a certain region of the universe (called the

$$E(EC) = \frac{\text{Volume } \lambda^3}{(2\pi)^2} (t^2 - 1) e^{-t^2/2}$$

$$+ \frac{(1/2) \text{ Area } \lambda^3}{(2\pi)^{3/2}} t e^{-t^2/2} + \frac{2 \text{ Diameter } \lambda}{(2\pi)} e^{-t^2/2} + \frac{EC}{(2\pi)^{1/2}} \int_0^t e^{-z^2/2} dz$$

The quantities on the right side all refer to the volume, surface area, "caliper" diameter, and EC of the search region, the region in space where the random field is defined, λ is a measure of the "roughness" of the field, and t is the threshold level.

The caliper diameter of a convex solid is defined as follows. Place the solid between two parallel planes or calipers and measure the distance between the planes, averaged over all rotations of the solid. For a sphere it is just the usual diameter; for a box of size $a \times b \times c$ it is $(a+b+c)/2$, half the "volume" used by airlines to measure the size of luggage.

Notice that if the search region is a plane or slice in 2-D, then the volume is 0 and the first term disappears; if the search region is a single point with zero volume, surface area, and diameter, then only the last term is left, which astute statisticians will recognize as just the probability that a Gaussian random variable exceeds the threshold t .

All these parameters are easily estimated, and we can plot the EC against the threshold t , and add it to the plot of the observed EC (Fig. 4). As we can see, the results are in reasonable agreement. What this shows is that the hollows, sheets, strings, and clusters of galaxies might have been generated by a Gaussian random field model for galaxy density and that the meatball, sponge, and bubble topologies arise naturally from the model when the threshold is high, medium, or low, respectively.

Table 1—Blood-Flow at a Single Voxel

Subject	Rest	Task	Diff.
1	123.72	125.18	1.46
2	129.64	140.85	11.21
3	133.20	137.41	4.21
4	141.03	148.13	7.10
5	128.41	139.82	11.41
6	167.19	169.65	2.46
7	124.58	135.77	11.19
8	104.55	105.23	0.68
9	114.24	121.65	7.41
10	111.75	113.93	2.18
Mean	127.83	133.76	5.93
			Voxel average s.d. = 3.61
			Z = $\sqrt{10}$ mean/s.d. = 5.19

of the simplest models of galaxy density.

Enter the Statistician—I

The models describe not the actual density patterns produced, but the randomness of the density patterns. This gives rise to a random EC, so this is where the statistics comes in. We can work out the EC of the observed universe, and compare it to the expected EC under our model, averaged over all random repetitions generated by the model.

This does not seem a very simple task; we must create a random universe, find the EC of the excursion set, repeat this many times, and average. Is there a neat theoretical formula for this that would save us the trouble of simulating? For one of the simplest and most popular models, a stationary Gaussian random field (see Sidebar 2), the answer is yes.

This remarkable result, remarked by Robert Adler in 1976 as able for its simplicity, was discovered by his Ph.D. thesis, and I added a boundary correction this year (see Sidebar 3 for a proof). The formula for the expected EC is [the new editor of *Chance* said I was allowed one formula per page, but he didn't say how big it could be!]:

search region), and from this a map of galaxy density was produced.

The places where the density exceeded a fixed threshold, called the excursion set, was determined, and the EC of this set was calculated as described previously. If the universe is like meatballs or bubbles, then the EC would be positive; if it is like a sponge, then the EC would be negative.

Unfortunately the EC depends very strongly on the threshold. If the threshold is high, then all the holes and hollows tend to disappear and the EC counts the number of high-density regions; if it is low, then it counts the number of low-density hollows; in between, the EC is negative.

To see this dependence, the galaxy density is first standardized by transforming to a Gaussian distribution; then the EC is plotted as a function of the threshold, to produce Fig. 4, which is based on the latest galaxy survey (Vogelley, Park, Geller, Huchra, and Gott 1994). The universe looks like meatballs at high density, a sponge at medium density, and bubbles at low density. As we shall soon see, this behavior is in fact quite normal for one

Sidebar 1—Cover Story

Fractal Models for Galaxy Density

Fractal models for galaxy density have been proposed by Mandelbrot (1982) to describe the grouping of galaxies into clusters, strings, and sheets. Recall that a fractal has the property that it looks similar at all scales, that is, galaxies group together in clusters, the clusters group together in superclusters, which themselves join together to form immense structures such as the "Great Wall," a recently discovered sheet of galaxies stretching across a large portion of the observable universe. One of Mandelbrot's models is based on a random walk (Mandelbrot 1982, p. 136). Start at one point, jump a random distance in a random direction and mark the spot with a galaxy. Jump again from this galaxy, and mark the new spot with another galaxy. Repeat this process indefinitely. Depending on the distribution of the jump distance, different patterns of galaxies emerge. If you choose a jump with a finite mean but an infinite standard deviation, then the patterns begin to look interesting. Mandelbrot suggested a jump probability density proportional to the distance raised to the power $-2\frac{1}{2}$. The result is known as a Rayleigh-Lévy random "flight"—rather than "walk"—because it occasionally takes very large jumps. The pattern of galaxies that it produces does resemble those actually observed (Fig. 5), although there is no physical explanation for the jump process itself.

Chi-Squared Random Fields

Where is the link to random fields? Hamilton (1988) calculated the ex-

pected EC (Euler characteristic) of excursion sets for Mandelbrot's model. Strangely enough, they agree exactly with those of a so-called χ^2 (chi-squared) random field with two degrees of freedom. This is defined as follows. Take two smooth independent Gaussian random fields, say Z_1 and Z_2 (Sidebar 2), then the χ^2 random field is defined as $Z_1^2 + Z_2^2$. Such random fields have very peculiar properties when the threshold is set close to zero. The excursion set, instead of consisting of isolated hollows, corresponding to local minima of the field, consists instead of *strings* of hollows which form closed loops with no free ends (outside front cover, bottom right panel). The explanation for this curious phenomenon comes from the fact that if the χ^2 field is zero, then each of the component fields Z_1 and Z_2 must also be zero. Now the zero set of a smooth Gaussian field is a sort of 3-D contour in space; it consists of a smooth surface with no free ends (outside front cover, top left and right panels). The places where both Z_1 and Z_2 are zero are the places where these two smooth contours intersect (outside front cover, bottom left panel). A moment's reflection will convince you that the intersection of two smooth surfaces must be a string, closed to form a loop.

Knots in the Excursion Set

Knot theory has always been an interest of mine since Vaughan Jones, one of my friends from my high school and undergraduate days in New Zealand, became a well-known knot theorist. For his contributions to knot theory, in particular his discovery of a new knot invariant now called the Jones polynomial, he was awarded

the Fields Medal, the equivalent of the Nobel Prize (there is no Nobel Prize for mathematics), and he was made a Fellow of the Royal Society. So as soon as I hear the word "string" I immediately think of knots. Can these zero-sets of χ^2 fields ever be knotted? That is, is it possible for two smooth surfaces to intersect and form a knot? For a long time I thought the answer was no. I just couldn't see how it could be done. I mentioned this to Vaughan one day after we had been wind-surfing on San Francisco Bay. After he came out of the shower, he said that not only could he find such a pair of surfaces, but any knot could be formed as the intersection of two suitably chosen smooth surfaces. Figure 3 shows how to do it for the simplest knot, the trefoil (Fig. 3a). An elementary theorem in knot theory says that any knot can be represented as the boundary of a two-sided surface, its "Seifert representation" (Fig. 3b,c). The fact that the surface is two-sided, as opposed to a one-sided surface such as a Möbius strip, is crucial to the success of the proof (I have colored one side green and the other blue). Now thicken the surface (Fig. 3d) and carefully cut it in half, separating the two sides (Fig. 3e). Now take out some sandpaper and smooth the two pieces (Fig. 3f) and push them back together so that they intersect (Fig. 3g). The intersection (Fig. 3h) is the knot we desire (Fig. 3i). So there it is ... although there are some details to fill in to complete the proof, it seems likely that the zero-sets of χ^2 fields, and perhaps also the areas of zero density in the Mandelbrot fractal model for galaxy density, can form knots. This may be nothing more than a mathematical curiosity with no practical relevance ... but who knows?

Medical Images: Positron Emission Tomography

For a long time I had been doing some statistical consulting work for the Montreal Neurological Institute. Nothing too demanding—the odd *t* test or perhaps an analysis of variance now and again. But then in the summer of 1990 a new type of experiment was

performed by the brain imaging center.


In this experiment, a subject is injected with a radio isotope emitting positrons, which annihilate with nearby electrons to release gamma rays that are detected by the center's new Positron Emission Tomography (PET) machine. By careful reconstruction, the researchers are able to build up an

image of blood flow in the brain, a measure of brain activity.

This opened up the possibility of actually seeing which regions of the brain were activated by different stimuli, to actually see the brain "thinking." For a good introduction, see the cover story of *Time* magazine, July 17, 1995.

In one of the first experiments, conducted by Dan Bub, a subject

Figure 5. Mandelbrot's fractal model for galaxy density is generated by a Rayleigh-Lévy random flight: (a) Starting at a point, jump a random distance with probability proportional to the distance to the power $-2\frac{1}{2}$ in a random direction and (b) leaving behind a galaxy after each jump.



Now comes the hard part: It looks as if some voxels show increased activation, but is it real or is it just due to the noise? Here is where the statistics (and the statistician!) come in.

Any student who has been through a first course in statistics should know what to do with the sort of data in Table 1: A paired-difference T statistic can be used to measure the statistical significance of the activation. This is calculated by taking the mean difference (Fig. 7d), dividing by the standard deviation by the square root of the number of subjects $\sqrt{10}$. To increase the accuracy of the standard deviation, it was replaced by the average over all voxels to obtain a Z statistic with an approximate normal or Gaussian distribution.

Proceeding in this way, we can make an image of Z statistics, one for each voxel (Fig. 7f). The researcher then scan this image, looking for high values of Z , by choosing a threshold value t and

Now comes the tricky part: How do we choose the threshold so that we exclude all the noise or at least exclude it with a probability of say .95? In other words, how do we control the specificity? The first thing that springs to mind is to use the standard .05 critical values for Z from the normal distribution tables, which is $t=1.64$ in this case. If there is really no activation at one particular voxel, then the chance of finding any, that is, the chance that $Z > 1.64$, is controlled to be .05.

But wait a minute—suppose there was really no activation anywhere in the brain, just random noise—then we would expect 5% of the voxels to exceed this threshold by chance alone. So if we use $t=1.64$, we might end up finding activation in at least 5% of the brain, even when there is none at all—obviously unacceptable! Is there a way of correcting for this? Students in a good course in statistics are told to watch out for “data dredging,” carrying out a

task, silent reading of words on a screen, during the imaging process. By subtracting an image in which the subject was “at rest” looking at a blank screen, the experimenters were able to see evidence of increased blood flow in certain brain regions corresponding to the effort required for the task.

The images are very blurred, however, and the signal (if any) is very weak compared to the background noise, so to increase the signal-to-noise ratio the experiment was repeated on 10 subjects. The brain images were aligned in 3-D, and the blood flow was averaged (see Fig. 7a, b, c).

The images are stored as values at $128 \times 128 \times 80$ locations called voxels. If we just look at one voxel, we have 10 pairs of blood-flow values, one pair for each subject, one taken while the subject was performing the task, the other while the subject was at rest. Table 1 shows some values from just one voxel in the left frontal region of the brain.

Sidebar 2—The Gaussian Random Field

A random field is not an intimidating object; it is just a set of random variables defined at every point in space, a bit like a 3-D version of a time series. A Gaussian random field is one that has a Gaussian distribution at every point and at every collection of points. A stationary field is one that has the same distribution at all points, in other words, it looks the same everywhere. The simplest way of making a smooth stationary Gaussian random field is to start with inde-

pendent Gaussian observations ('white noise') at each point in space, for simplicity on a fine cubical lattice, each with mean zero and standard deviation one (see Fig. 6a). Then smooth it with a smooth kernel, such as a Gaussian shaped kernel, that is take the weighted average of the white noise, weighted by the kernel (Fig. 6b). The roughness of a field, denoted by λ , is measured by the standard deviation of the slope of the field relative to the standard

deviation of the field itself. It can be shown that for a Gaussian random field constructed as above, $\lambda = \sqrt{4 \log_e 2} / \text{FWHM}$, where FWHM is the Full Width of the smoothing kernel at Half its Maximum height. For PET images, the FWHM is measured experimentally by placing a tiny point source of isotope into the PET machine (Fig. 6c) and measuring the FWHM of the resulting blurred image of the point source, which should be an image of the kernel itself (Fig. 6d).

multitude of tests and only reporting the significant ones, using the usual 5% level for each test. There are about 300,000 voxels inside the brain, so we are carrying out 800,000 statistical tests. Here is data dredging on a massive scale!

The classical way of correcting for this is to divide the 5% level by the number of tests, the so-called Bonferroni correction, which gives a critical value of $t=5.10$. This usually overcorrects, however, since after all the voxels are arbitrary and if we subdivided each into eight voxels of half the size in each direction, we would divide now by 2,400,000 and obtain an even larger threshold ($t=5.48$). Thus the Bonferroni threshold is not right. Instead we need to know the probability that the maximum of a smooth image of Z values exceeds a threshold value, then adjust the threshold to attain a .05 probability.

The problem comes down to the distribution or p value (the probability of a more extreme value) of a random field maximum. This is definitely not taught in an elementary statistics course, but after a trip to our mathematics and statistics library I discovered that indeed some work had been done on this problem. No exact formula exists, but Russian probabilists had found approximate answers in the early 1970s. The most fascinating

aspect, however, from my point of view, was a link to the EC.

Enter the Topologist

Simply stated, the probability of the maximum exceeding the threshold is well approximated by the expected EC of the excursion set (given by that big formula), provided the threshold is high. This result, linking the topology of the excursion set to the local maximum, was reported by Hasofer, Adler's thesis supervisor, in 1978.

The idea of the proof is very simple: as the threshold increases the holes in the excursion set (Fig. 9c) disappear until each component of the excursion set includes just one local maximum (Fig. 9d). At this point, the EC counts the number of local maxima. Raising the threshold still further, until only the global maximum remains (Fig. 9e), the EC takes the value 1 if the global maximum exceeds the threshold and 0 if it is below; thus, for high thresholds, the expected EC equals the probability that the global maximum exceeds the threshold.

So now we have it—we can use the expected EC for the p value of the maximum activation. First we must fill in the volume, surface area, caliper, diameter, and EC of the brain, which are 1,064 cc, 1,077 cm², .1 cm, and 2, respec-

tively (the ventricles form a single large hollow in the center of the brain, visible in Fig. 7, that increases the surface area, reduces the caliper diameter, and gives the brain an EC of 2, as in Fig. 1f). The only remaining part is to use the FWHM—Full Width of the smoothing kernel at Half its Maximum height (Sidebar 2)—at 2 cm to find λ , equate the expected EC to .05, and solve for t . The result is $t = 4.22$, much higher than 1.64, but not as high as the Bonferroni value of 5.10.


The resulting excursion set is shown in Fig. 7h, and these are the places that show significant activation. There is one in the visual cortex and adjacent left extrastriate and one in the left frontal cortex, near the language area of the brain.

Finally, Fig. 12 plots the observed EC and the expected EC against the threshold, as for the galaxy data (Fig. 4). The main feature of interest is the larger-than-expected EC for the high thresholds ($t > 3$), attributable to the two regions of activation noted previously.

Conclusions

The work described in this article has linked together several ideas from different areas of mathemat-

Figure 6. How to make a smooth Gaussian random field in 2-D: (a) Start with independent Gaussian observations ("white noise") on a 64×64 lattice, each with mean 0 and standard deviation 1 (the plots below each image show the values along the horizontal line through the middle of each image); (b) take the weighted average of the white noise, weighted by a smooth Gaussian-shaped kernel. To see the kernel itself (c), replace the lattice by zeros everywhere apart from a value of 1 at the central point; the result (d) is an image of the kernel. The amount of smoothing is measured by the FWHM—the Full Width of the smoothing kernel at Half its Maximum height. This is how the FWHM of the PET machine is actually measured.






Figure 8. The anomalies in the cosmic microwave background radiation—thought to be the signature of the creation of the universe by the “big bang.” The anomalies were discovered in 1992 by George Smoot and his coworkers at the Lawrence Berkeley Laboratories.

Figure 7. Application to a PET study of brain areas activated by a reading task. The brain (a) has been rendered as a transparent solid with the rear left side facing the viewer; the ventricles in the center form a single connected hollow that gives the brain an EC of 2, as in Fig. 1f. One of 80 slices through the brain is color-coded (red = high, purple = low) to show (b) average blood flow of 10 subjects under the rest condition and (c) under the task condition. (d) The difference of the averaged blood flows, task - rest; (e) the standard deviation of the 10 differences (9 degrees of freedom); (f) the Z statistic for testing for a significant increase in blood flow due to the task. (g) The excursion set of all Z values greater than a threshold of $t = 3.3$, where we expect the EC to be 1 if the image is pure noise (no activation); in fact the EC = 4, due to the four components clearly visible, suggesting some evidence of activation; and (h) at a threshold of $t = 4.22$, chosen so that the expected EC is .05 when no activation is present, that is, all noise is excluded with a probability of .95. Two regions remain (EC of the excursion set is 2), one in the left frontal lobe, near the language area of the brain and a broad region in the left visual cortex and extrastriate. These are the regions significantly activated by the task.

Sidebar 3—The Expected Value of the Euler Characteristic

also page 61—a French naturalist, who wrote *L'Histoire Naturelle*, one of the most widely-read scientific works of the 18th century. He is best known to mathematicians for his method of calculating π by simply throwing a needle at random onto the floor and counting the proportion of times t , say, that it crossed the cracks between the floor boards (Fig. 10a). If the needle has the same length as the width of the floor boards, then he showed that the probability of crossing the cracks equals $2/\pi$. A generalization of this result was discovered by Blaschke in the mid-1930s, where the needle and the cracks are replaced by two arbitrary sets A and B in 3-D, and the proportion of crosses is replaced by the total EC of the set of points that belong to both sets. The result, known as the Kinematic Fundamental Formula—"Fundamental" because it underlies so many other results in integral geometry—is as follows:

$$\int EC = 8\pi^2(\text{Volume of } A)(EC \text{ of } B) + 4\pi^2(\text{Area of } A)(\text{Diameter of } B) + 4\pi^2(\text{Diameter of } A)(\text{Area of } B) + 8\pi^2(EC \text{ of } A)(\text{Volume of } B).$$

Now we extend the random field out on all sides beyond the search region, hold it fixed, and use the excursion set as the "cracks" (B). Onto this we drop a random search region (A) that becomes the "needle" (Fig. 10b). The Kinematic Fundamental Formula then gives an expression for the expected EC, which after a little more manipulation becomes the four terms of the result we want.

always the number of connected mountain ranges above this altitude, minus the number of lower altitude plains completely surrounded by mountains. In other words, the mountain heights form the random field, the peaks are maxima, the basins are minima, the saddles are just the saddle points, and the excursion set is the places above a fixed altitude. Now place a +1 at every local maximum and a local minimum, a -1 at every saddle point, and a 0 at every other point; add these up and you have the EC. The expected EC is then just the sum of the expected +1, -1, or 0 at every point, which leads, after a lot of delicate manipulation, to the first term of the expected EC, the part discovered by Robert Adler. There is only one snag: What happens if the excursion set touches the boundary? For this possibility, a special boundary correction is required that adds three more terms to Adler's result. These can also be obtained by counting local maxima and minima on the boundary, and following the same delicate manipulation in Adler's original proof. An alternative method of proof, using integral geometry, is given in the next section.

Integral Geometry: Buffon's Needle

Our formula for the expected value of the Euler characteristic (EC) is remarkable because it deals with such a difficult random variable. How are we to apply the concept of expectation to such a curious quantity as the EC? As it turns out, there are two ways of attacking the problem, coming at it from two different branches of mathematics: differential topology and integral geometry.

Differential Topology: The Mathematical Mountaineer

Statisticians are used to dealing with the expectation of simple quantities such as sums of random variables. The trick is to convert the EC into just such a sum, a sum of local properties of the random field, rather than global properties such as connectedness. The necessary theoretical result needed to perform this comes from differential topology and Morse theory, discovered in the 1960s. It is easiest to state in 2-D, where it says simply this: Provided the excursion set does not touch the boundary of the search region, then the EC of the excursion set equals the number of local maxima of the random field inside the excursion set, plus the number of local minima, minus the number of saddle points (Fig. 9):

$$EC = \# \text{maxima} - \# \text{saddles} + \# \text{minima}.$$

The mathematical mountaineer applies this result as follows: count the number of peaks, minus the number of saddles, plus the number of basins above a certain altitude; the result is




Figure 9. The EC of the excursion set of a random field in 2-D is the number of maxima (M) – saddles (S) + minima (m) of the field inside the excursion set. For an artificial image (a), the EC of the excursion set for a low threshold (b) at $t = 1$ is $EC = 4 - 4 + 1 = 1$; (c) at $t = 2.8$ the local minimum disappears, leaving a hole, and $EC = 4 - 2 + 0 = 2$; (d) at $t = 3.2$ the hole breaks open leaving components with just one local maximum in each, and $EC = 4 - 0 + 0 = 4$; (e) at $t = 5.6$, just below the global maximum, $EC = 1 - 0 + 0 = 1$; above the global maximum the EC is 0, so that for high thresholds the expected EC approximates the p value of the global maximum; (f) is a plot of the EC against the threshold, t .




Figure 10. A generalization of Buffon's needle is used to find the formula for the expected EC of the excursion set; (a) a needle is thrown at random onto the floor; if the length of the needle is the same as the separation of the cracks in the floor boards, then the probability of the needle crossing a crack is $2/\pi$; (b) replace the needle with the search region (square shape shown in 2-D) and the cracks with the excursion set of a random field, extended on all sides; the Kinematic Fundamental Formula gives the expected EC of the points in common (black), which leads eventually to the formula we want.

Sidebar 4—Wrinkles in Time

One of the greatest events in cosmology was the discovery in 1964 by Arno Penzias and Robert Wilson, two researchers working for Bell Labs, of the cosmic microwave background radiation, thought to originate from the creation of the universe by the "big bang." The story has been told many times before. Penzias and Wilson had constructed a powerful microwave antenna used to bounce microwave signals off Echo 1, a huge reflecting balloon that had been placed in earth orbit. Instead they discovered background radiation coming uniformly from all directions in space, showing no correlation with any known celestial object. They concluded that it must be coming from outside the solar system, and outside our own galaxy. They brought their result to the attention of astrophysicists who realized that what had been observed could be the remnants of the enormous release of radiation from the "big bang." Since then, the radiation had cooled along with the expansion of the universe to a lower frequency in the microwave part of the spectrum. This discovery was a dramatic confirmation of the "big bang" theory, since the radiation had just the right spectrum, and it was uniform in all directions. The only problem was that it was just too uniform; theories for the creation of large-scale structure in the universe predicted that the cosmic microwave background radiation should show some anomalies, or places of decreased frequency, which could have been the seeds that formed the present-day clusters of matter into galaxies and clusters of galaxies. Ever since its initial discovery, astrophysicists tried without success to find these anomalies, until George Smoot and his co-workers at the Lawrence Berkeley Laboratory announced their discovery in 1992 (see the cover story of *Scientific American*, July 1992). The EC curve similar in shape to that expected, but somewhat lower and spread more in the tails. This discrepancy points to a Gaussian random field model for the anomalies, with a larger standard deviation and a larger smoothness than the background noise.

That it is very difficult to tell which are the real anomalies and which are just part of the noise (Fig. 8). When I first saw it, it reminded me of one of the brain images! In fact astrophysicists have used the very same methods, the EC of the excursion set, to try to sort out the signal from the noise (Torres, 1994). In this case, the search region is the surface of a unit sphere, so the volume is zero, the surface area is $4\pi^2$ for the inside and $4\pi^2$ for the outside, the caliper diameter is zero, and the EC of the search region is 2 (because it is a hollow sphere, like a tennis ball). We can use these results to calculate the expected EC and to plot it and the observed EC against threshold level (Fig. 11). The observed microwave background radiation produces an EC curve similar in shape to that expected, but somewhat lower and spread more in the tails. This discrepancy points to a Gaussian random field model for the anomalies, with a larger standard deviation and a larger smoothness than the background noise.




Figure 11. Plot of the observed EC of excursion sets of the anomalies in the cosmic microwave background radiation (jagged line), and the expected EC from the formula (smooth line) if there are no real anomalies. The observed microwave background radiation produces an EC curve similar in shape to that expected, but somewhat lower and spread more in the tails—evidence that some of the anomalies are real and not just due to random noise. This discrepancy points to a Gaussian random field model for the anomalies, with a larger standard deviation and a larger smoothness than the background noise.




Figure 12. Plot of the observed EC of the set of high-Z regions for the PET data (jagged line) and the expected EC from the formula (smooth line) if there is no activation due to the linguistic task, plotted against the threshold t . The most interesting part is when $t > 3$, showing more EC than expected, confirming that some components of the excursion set are due to the linguistic task and not the random noise. In particular, at $t = 3.3$ we expect an EC of 1, but we observe 4 (visible in Fig. 9g); at the 5% critical value of $t = 4.22$, we expect .05 but we observe 2 components (visible in Fig. 9h).

ics, statistics, and probability to solve problems in astrophysics and medicine. It is a fascinating area, still developing in different directions and finding different applications.

It has already been used for detecting activation in the new field of functional magnetic resonance imaging (fMRI) (see the cover story of *Time* magazine, July 17, 1995) and for studying the recently discovered anomalies in the cosmic microwave background radiation, thought to originate from the creation of the universe by the "big bang" (see Sidebar 4).

No doubt more applications will arise wherever there is a need to interpret the geometry of random images.

Acknowledgments

Color images were produced using *IRIS Explorer 3.0* running on a Silicon Graphics Indigo² Extreme from Silicon Graphics, Inc.; monochrome images were made using S-

PLUS from MathSoft, Inc. I would like to thank Robert Adler (Technion, Haifa) for helpful criticism of previous theoretical work, Alan Evans (Montreal Neurological Institute) for inspiration, enthusiasm, and permission to use the medical imaging data, Luis Tennorio (Lawrence Berkeley Laboratory) for help with the *COBE* data, and both the Natural Sciences and Engineering Research Council of Canada and the Fonds pour la Formation des Chercheurs et l'Aide à la Recherche du Québec for financial support.

References and Further Reading

- Adler, R. J. (1981), *The Geometry of Random Fields*, New York: John Wiley.
- Buffon, G. L. L., Comte de (1777), "Essai d'Arithmétique Morale," *Supplément à l'Histoire Naturelle*, Tome 4.
- Hamilton, A. J. S. (1988) "The Topology of Fractal Universes," *Publications of the Astronomical Society of the Pacific*, 100, 1343-1350.
- Mandelbrot, B. B. (1982), *The Fractal Geometry of Nature*, New York: W. H. Freeman.
- Smoot, G., and Davidson, K. (1993), *Wrinkles in Time*, New York: William Morrow.
- Torres, S. (1994), "Topological Analysis of COBE-DMR Cosmic Microwave Background Maps," *Astrophysical Journal*, 423, L9-L12.
- Vogele, M. S., Park, C., Geller, M. J., Huchra, J. P., and Gott, J. R. (1994), "Topological Analysis of the CfA Redshift Survey," *Astrophysical Journal*, 420, 525-544.
- Worsley, K. J. (1995), "Estimating the Number of Peaks in a Random Field Using the Hadwiger Characteristic of Excursion Sets, with Applications to Medical Images," *The Annals of Statistics*, 23, 640-669.
- Worsley, K. J. (1995), "Boundary Corrections for the Expected Euler Characteristic of Excursion Sets of Random Fields, With an Application to Astrophysics," *Advances in Applied Probability* 27, 943-959.
- Worsley, K. J., Evans, A. C., Marrett, S., and Neelin, P. (1992), "A Three Dimensional Statistical Analysis for CBF Activation Studies in Human Brain," *Journal of Cerebral Blood Flow and Metabolism*, 12, 900-918.