From Lichen to Lightning: Understanding Random Growth

AMANDA TURNER

Abstract. Random growth arises in physical and industrial settings, from cancer to polymer creation. Despite considerable effort, mathematicians and physicists have been unable to answer fundamental questions such as "What do typical clusters look like?". This article explores how combining probability and complex analysis can provide mathematical descriptions of random growth.

Random growth in nature

Random growth processes occur widely in nature. They appear as collections of particles (often cells or molecules), called clusters, to which new particles are attached according to some stochastic rule. The physical mechanism by which the growth occurs determines the precise stochastic rule, and different types of growth can produce very different clusters.

Some of the most commonly seen examples of random growth in nature arise through biological growth, such as lichen on a rock or cancer tumours grown in a lab. In these examples, growth occurs due to a reproductive event taking place on the boundary of the cluster and so (in the absence of environmental influences) the next particle is equally likely to be added anywhere on the cluster boundary. Many lichens tend to grow as clusters that are round with roughness at the edges. Other biological growth processes also commonly produce this kind of shape.

Another way in which random growth can occur is through mineral aggregation. An example of this is soot deposition in a diesel engine. Fine particles of carbon diffuse around the engine until they either hit the surface of the engine or the soot aggregate which has started to line the engine. At this point they stick, becoming part of the aggregate. Similar kinds of clusters arise as frost flowers, which appear on car windows on some winter mornings, and in industrial processes such as electro-deposition. In all of these examples, the position at which each successive particle is attached is given by the hitting distribution of a diffusion process (such as Brownian motion) on the cluster boundary. It is no longer the case that all positions on the cluster boundary are equally likely as diffusion processes have a greater chance of hitting protrusions than indentations. As a result, clusters tend to be much more irregular than those formed from biological growth processes, often exhibiting long fingers, deep fjords and fractal branching structures (see Figure 1).

Figure 1. A "pseudo fossil" formed by the deposition of manganese oxide in a sandy substrate. This illustrates the typical features present in clusters formed through mineral aggregation. (Photo: Alan Dickinson).

Although lightning may not look like an obvious example of random growth, it also fits into this framework. The density, humidity, and conductivity of air is very inhomogeneous. Lightning surveys all possible paths through this medium and then strikes along those paths with the least electrical resistance. The molecules through which the lightning strike passes can be viewed as a cluster to which molecules are successively added in a way which minimises the total electrical resistance at each time-step. Surface discharges of lightning (also called Lichtenberg figures) and polymer formation are examples of random growing clusters which also arise in this way. The shape of the cluster formed depends on the strength of the local electric field and can range from a solid disk if the electric field is very weak to a one-dimensional path if the electric field is very strong.

Discrete models

The earliest mathematical models for random growth were discrete in nature. In discrete models, growing clusters $K_0 \subset K_1 \subset \cdots$ are constructed as increasing sets of connected vertices on an underlying (infinite) connected graph $G = (V, E)$. At each time step a new vertex is chosen from the neighbours of the cluster according to some stochastic rule and added to the cluster. To be precise, for any $\emptyset \subset A \subset V$, let

$$\partial A = \{v \in V \setminus A : (u, v) \in E \text{ for some } u \in A\}$$

denote the boundary of the subgraph A and suppose $p_A : \partial A \to [0, 1]$ is some function which satisfies

$$\sum_{v \in \partial A} p_A(v) = 1.$$

A random growth process $K_0, K_1, \ldots$, of sets of connected vertices, can be constructed recursively by starting from some connected set $\emptyset \subset K_0 \subset V$ and taking $K_{n+1} = K_n \cup \{v\}$ with probability $p_{K_n}(v)$. It is usual to start with $K_0 = \{v_0\}$ where v_0 is some distinguished vertex called the seed particle. Typically G is a lattice such as $\mathbb{Z}^d$ and $v_0 = 0$.

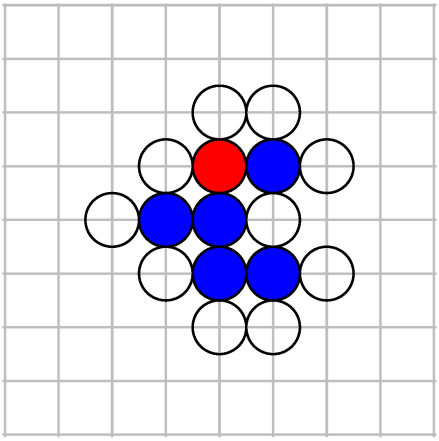

Figure 2. A representation of a random cluster grown on $\mathbb{Z}^2$. The red disk represents the seed particle which was placed at the origin. The blue disks are the first five particles to be attached to the cluster. The empty black circles represent the possible locations at which the next particle might be attached.

In order to model the physical growth processes discussed in the previous section, we need to show how to construct probability distribution functions p_A which correspond to the physical attachment rules described above. We begin by considering biological growth. One possible way to model uniformly random reproductive events on the cluster boundary is by choosing each successive particle uniformly from all possible locations on the boundary i.e. $p_A(v) = 1/|\partial A|$ for each $v \in \partial A$. This model is called the Eden model and was first proposed in 1961 [1]. This is not the only way in which biological growth can be modelled on a lattice. Observe that in Figure 2, the vertex at position $(1, -1)$ has three possible parents whereas the vertex at position $(2, 0)$ has only one possible parent. If the new particles are actually arising as offspring of the parent particles, one would expect the next particle to be more likely to be added at $(1, -1)$ than $(2, 0)$. A more realistic variation on the Eden model is to pick each vertex with a probability which is proportional to the number of edges which connects that vertex to the cluster i.e. if

$$N_A(v) = \{u \in A : (u, v) \in E\},$$

then

$$p_A(v) = \frac{|N_A(v)|}{\sum_{u \in \partial A} |N_A(u)|}.$$

In 1981, Witten and Sander proposed the following model for mineral aggregation which they called diffusion-limited aggregation or DLA [6]. As before, initialise with a seed at the origin or some other distinguished vertex. At each time step place a particle on a random site at a large distance from the seed. Allow the particle to perform a simple random walk on the graph until it visits a site adjacent to the cluster at which point attach the particle to the cluster i.e. add the vertex on which the particle stopped. The distribution function p_A can therefore be obtained by computing the hitting probabilities of the cluster boundary ∂A by a simple random walk started at "infinity". In order to guarantee that the particle eventually reaches the cluster we need the random walk to be a recurrent process on the underlying graph. This means that if $G = \mathbb{Z}^d$, we must take $d \leq 2$. The case $d = 1$ is not very interesting as particles will always just attach to the left or right of the cluster with equal probability, so it is usual to take $d = 2$ when talking about DLA. A variant of this model is called multi-particle DLA in which infinitely many particles are simultaneously released from random positions, rather than one at a time, and they perform independent random walks until they hit the cluster. Computing p_A explicitly for larger and larger sets A gets complicated very quickly and the long-term behaviour of DLA is notoriously difficult to understand.

In order to model clusters formed by electric discharge, the notion of electric potential is needed. For simplicity, take $G = \mathbb{Z}^2$ and $v_0 = 0$. Suppose there are two electrodes in the system: one placed at 0 and the second modelled as a large circle of radius R. Let $B_R(0) \subset \mathbb{R}^2$ denote the open ball with centre 0 and radius R. Given a conducting cluster $\{0\} \subseteq A \subset B_R(0) \cap \mathbb{Z}^2$, the electric potential is a discrete harmonic function ϕ_A (see the aside on "Harmonic functions") defined on the lattice with $\phi_A(v) = 0$ if $v \in A$ and $\phi_A(v) = 1$ if $|v| \geq R$. (It follows from the discussion on harmonic functions that $\phi_A(v)$ is equal to the probability that a simple random walk started at v exits the ball $B_R(0)$ before hitting the set A.) The probability of attaching a particle at a vertex $v \in \partial A$ depends on the local electric field, or potential difference, $\phi_A(v) - \phi_A(u) = \phi_A(v)$ for each $u \in N_A(v)$. Fix $\eta \in \mathbb{R}$. For each $v \in \partial A$, let

$$p_A(v) = \frac{|N_A(v)|\phi_A(v)^\eta}{\sum_{u \in \partial A} |N_A(u)|\phi_A(u)^\eta}.$$

This model is called dielectric-breakdown or DBM(η) and it was proposed in 1984 by Niemeyer, Pietronero and Wiesmann [3]. The parameter η determines the extent to which the attachment locations are influenced by the electric potential. In the case when $\eta = 0$, this model is just the variation of the Eden model described above; when $\eta = 1$ it produces clusters which are believed to behave qualitatively like DLA; as $\eta \to \infty$ growth concentrates at the point of maximal potential difference. It is conjectured that the limit in this case is a simple path. Similarly to DLA, computing ϕ_A (and hence p_A) explicitly becomes increasingly complicated as the cluster grows.

Conformal models for planar random growth

The discrete models defined above are challenging to study, in part due to a lack of available mathematical tools. In 1998 Hastings and Levitov [2] formulated an approach to modelling planar growth, which included versions of the physical models described above. The idea was to represent growing clusters as compositions of conformal mappings. This approach provided a way in which techniques from complex analysis could be used to study planar random growth.

Let K_0 denote the unit disk in the complex plane $\mathbb{C}$ and let $D_0 = \mathbb{C} \setminus K_0$. A particle is any compact set $P \subset D_0$ such that $K_0 \cup P$ is simply connected. By the Riemann mapping theorem (see the aside) there exists a unique conformal bijection $f_P : D_0 \to D_0 \setminus P$ which fixes infinity. We use this mapping as a mathematical description of a particle attached to the unit disk. We say that the particle is attached at 1 if 1 lies in the closure of P. Examples of allowable particle shapes include a bump, as shown in Figure 3, a disk tangent to the unit circle at 1, or a slit (line segment) of the form $(1, 1 + d]$. If P represents a particle attached at 1, then $e^{i\theta}P$ represents a particle of the same shape attached at $e^{i\theta}$. It has corresponding conformal mapping

$$f_{e^{i\theta}P}(z) = e^{i\theta} f_P(e^{-i\theta} z).$$

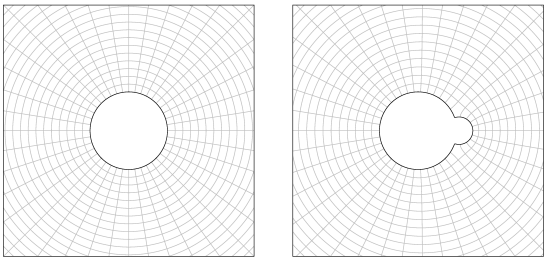

Figure 3. The exterior unit disk and its image under the conformal mapping corresponding to a single particle attached at 1.

The Riemann mapping theorem

Let $K \subset \mathbb{C}$ be a connected compact subset of $\mathbb{C}$, larger than a single point, such that $\mathbb{C} \setminus K$ is simply connected. There exists a unique conformal bijection $f : D_0 \to \mathbb{C} \setminus K$ which fixes infinity in the sense that, for some $C > 0$,

$$f(z) = Cz + O(1) \text{ as } |z| \to \infty.$$

Given a sequence of particles $P_1, P_2, \ldots$ and their associated conformal mappings $f_1, f_2, \ldots$, define a sequence of conformal bijections $\Phi_n : D_0 \to \mathbb{C} \setminus K_n$ by setting $\Phi_0(z) = z$ and recursively defining

$$\Phi_n(z) = \Phi_{n-1} \circ f_n(z) = f_1 \circ \cdots \circ f_n(z).$$

Note that $K_n = K_{n-1} \cup \Phi_{n-1}(P_n)$, so the sequence $K_0 \subset K_1 \subset K_2 \subset \cdots$ represents a growing cluster where the unit disk K_0 is the seed particle and at time n the particle $\Phi_{n-1}(P_n)$ is added to the cluster (see Figure 4).

When the particles which are being attached to the cluster are small, relative to the unit disk, the leading order behaviour of each particle map depends only on the attachment angle and the size of each

particle. For each particle P_n, let θ_n denote the attachment angle and d_n the diameter. By choosing the sequences $\theta_1, \theta_2, \ldots$ and $d_1, d_2, \ldots$ in different ways, it is possible to describe a wide class of growth models. In the remainder of this section, we will explore how to choose these sequences in order to construct models that correspond to the physical processes described in the first section.

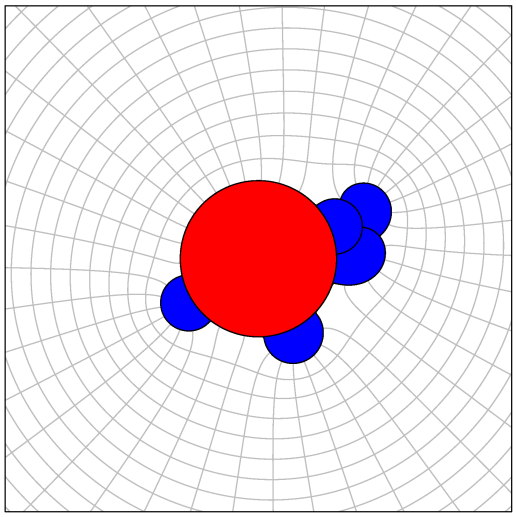

Figure 4. An example of a cluster K_n, grown by successive compositions of conformal mappings, together with the image of the exterior unit disk under the map Φ_n. The seed particle is shown in red and the first five particles to arrive are shown in blue.

In the model for biological growth, the probability of attaching a particle along a section of the cluster boundary should be proportional to the arc-length along that section of the boundary. Specifically, given an open arc in ∂K_{n-1}, there exist $a < b \in \mathbb{R}$ such that $\theta \mapsto \Phi_{n-1}(e^{i\theta})$ maps the interval (a, b) onto this arc. The length of this arc is therefore given by

$$\int_a^b |\Phi'_{n-1}(e^{i\theta})| d\theta.$$

Attaching the next particle onto this arc is equivalent to picking $\theta_n \in (a, b)$ and so we need

$$\mathbb{P}(\theta_n \in (a, b)) \propto \int_a^b |\Phi'_{n-1}(e^{i\theta})| d\theta$$

or equivalently θ_n must have density function proportional to $|\Phi'_{n-1}(e^{i\theta})|$. Next consider how the diameter d_n should be chosen. For simplicity suppose that the particle being attached is a small slit of length d_n i.e. $P_n = e^{i\theta_n}(1, 1 + d_n]$. In this case, an explicit formula can be written down for the function f_n (see [2]). Furthermore, there exists some β_n (which can be expressed as a function of d_n) such that $\theta \mapsto f_n(e^{i\theta})$ maps the interval $(\theta_n, \theta_n + \beta_n]$ to P_n. Using the chain rule, the particle $\Phi_{n-1}(P_n)$ which is added to the cluster therefore has length

$$\int_0^{\beta_n} |\Phi'_{n-1}(f_n(e^{i(\theta_n+\theta)}))||f'_n(e^{i(\theta_n+\theta)})| d\theta.$$

Since

$$d_n = \int_0^{\beta_n} |f'_n(e^{i(\theta_n+\theta)})| d\theta,$$

the length of $\Phi_{n-1}(P_n)$ is approximately equal to $d_n|\Phi'_{n-1}(r_n e^{i\theta_n})|$ for some $1 \le r_n \le 1 + d_n$. In the physical models, the particles being attached are all the same size, say d. We therefore need to choose d_n in such a way as to ensure that this expression is approximately equal to d for all n. When d is small, a possible choice is to take

$$d_n = d|\Phi'_{n-1}(e^{i\theta_n})|^{-1}. \quad (1)$$

The above choices of θ_n and d_n provide a model for an off-lattice version of the Eden model.

Now suppose we wish to construct an off-lattice version of DLA. Exactly the same argument as above shows that we should again choose d_n satisfying (1). However, this time we need to pick θ_n so that $\Phi_{n-1}(e^{i\theta_n})$ has the same distribution as W_τ where W_t is a Brownian motion started from infinity (viewed on the Riemann sphere) and τ is the first time that the Brownian motion hits K_{n-1}. Equivalently, $e^{i\theta_n}$ should be the hitting distribution of $\Phi_{n-1}^{-1}(W_t)$ on the unit circle. However, Φ_{n-1}^{-1} is a conformal mapping and therefore its real and imaginary parts are harmonic functions (see the aside). Using Itô's formula, $\Phi_{n-1}^{-1}(W_t)$ is a continuous martingale and hence is a time-change of Brownian motion. By symmetry, the hitting distribution of a time-change of Brownian motion on the unit circle is uniform, so θ_n should be chosen with the uniform distribution on $[0, 2\pi)$.

A similar argument can be used to show that one can obtain an off-lattice version of DBM(η) by taking θ_n with density function proportional to

$$|\Phi'_{n-1}(e^{i\theta})|^{1-\eta}$$

and d_n satisfying (1). Note that, as in the discrete case, $\eta = 0$ corresponds to the Eden model and $\eta = 1$ corresponds to DLA. There are several variations on this. In the original Hastings-Levitov model HL(α) [2], θ_n is picked uniformly on the unit circle and

$$d_n = d|\Phi'_{n-1}(e^{i\theta_n})|^{-\alpha/2}.$$

The HL(α) model corresponds to DBM(η) via the relation $\alpha = \eta + 1$ so HL(2) gives DLA (see Figure 5).

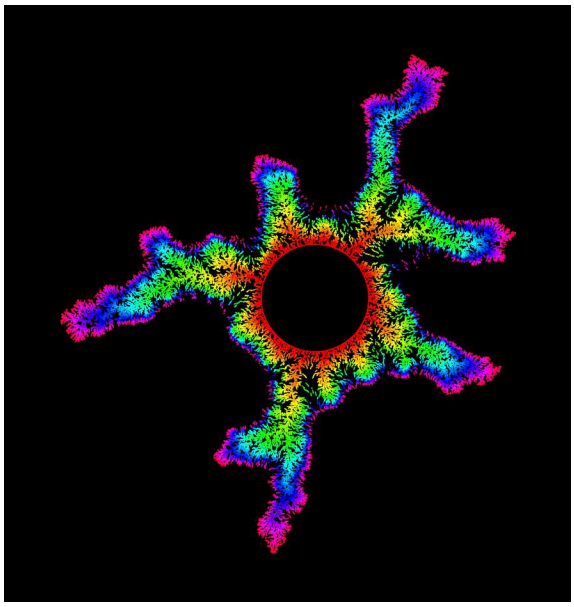

Figure 5. A version of DLA grown by composing conformal mappings using the Hastings-Levitov construction.

Mathematical study of random growth

Random growth clusters in the physical world usually consist of a large number of particles, each of which is small relative to the size of the cluster. Although the randomness typically occurs at a microscopic level through the attachment rule for each successive particle, we observe the clusters at a macroscopic level where we cannot see the individual particles. Random growth processes are completely unpredictable at the level of particles, however large clusters often exhibit predictable or 'universal' behaviour. The aim of studying mathematical models is to extract the principle mechanisms underlying this universal behaviour.

One of the intriguing features of random growth models is that, even though the models are isotropic by construction, simulations suggest that large clusters become anisotropic (see Figure 5). As yet there is no satisfactory mathematical explanation for how such complicated structures arise from the dynamics of the model. Representing random growth clusters using conformal mappings enables one to combine analytic and probabilistic techniques. Research at the interface of these two fields has already given us mathematical objects such as stochastic partial differential equations (SPDEs) and stochastic Loewner evolution (SLE). The recent developments in these areas have suggested approaches and techniques that can be applied in a random growth setting and we are beginning to be able to identify the characteristics of random growth in specific cases [5, 4].

It has been over fifty years since mathematicians and physicists first started seriously thinking about random growth. Although progress has been made in this time, we are still some way from being able to answer the really important questions. Mathematical tools and techniques that may shed light on these questions have recently started to emerge. There is a good chance that we will not need to wait another fifty years for the major breakthrough, so right now this is a very exciting area to be involved in.

REFERENCES

[1] M. Eden, A two-dimensional growth process, Proc. 4th Berkeley Sympos. Math. Statist. and Prob. (1961) 223–239.
[2] M. B. Hastings and L. S. Levitov, Laplacian growth as one-dimensional turbulence, Phys. D 116 (1998) 244–252.
[3] L. Niemeyer, L. Pietronero, and H. J. Wiesmann, Fractal dimension of dielectric breakdown, Phys. Rev. Lett. 52 (1984) 1033–1036.
[4] J. Norris, V. Silvestri, and A. Turner, Scaling limits for planar aggregation with subcritical fluctuations, Preprint 2019, arXiv:1902.01376.
[5] A. Sola, A. Turner, and F. Viklund, One-dimensional scaling limits in a planar Laplacian random growth model, Preprint 2018, arXiv:1804.08462.
[6] T. A. Witten Jr., L. M. Sander, Diffusion-limited aggregation: a kinetic critical phenomenon, Phys. Rev. Lett. 47 (1981) 1400–1403.

Amanda Turner

Amanda Turner is a Senior Lecturer in Mathematics and Statistics at Lancaster University. Her main research area is in probability with a specific focus on scaling limits of stochastic processes. Her work combines ideas from probability theory, real and complex analysis and statistical physics. When not doing mathematics, Amanda can usually be found up a mountain.

Harmonic functions

A continuous function $u : \bar{D} \to \mathbb{R}$, where D is some open subset of $\mathbb{R}^d$ and $\bar{D}$ is its closure, is a harmonic function if it satisfies Laplace's equation

$$\Delta u = 0 \text{ in } D \tag{2}$$

where

$$\Delta u = \frac{\partial^2 u}{\partial x_1^2} + \cdots + \frac{\partial^2 u}{\partial x_d^2}.$$

Using the Cauchy-Riemann equations, it can be shown that the real and imaginary parts of any holomorphic function are harmonic functions. Harmonic functions are closely connected with Brownian motion. Let D be a transient domain (i.e. a domain which Brownian motion will exit in finite time almost surely). Then if u is a harmonic function on D

$$u(x) = \mathbb{E}_x u(W_\tau) \tag{3}$$

where W is a d-dimensional Brownian motion started at x and $\tau = \tau_D$ is the first exit time from D. Set

$$\omega(x; dy) = \mathbb{P}_x(W_\tau \in dy).$$

Then we can write (3) as

$$u(x) = \int_{\partial D} u(y)\omega(x; dy).$$

We call $\omega(x; dy)$ the harmonic measure on ∂D as seen from x. Using the symmetry of Brownian motion, it follows that harmonic functions satisfy the mean value property

$$u(x) = \int_{\partial B_r(x)} u(y) dA(y)$$

for all balls $B_r(x) \subseteq D$, where dA is the normalised uniform surface area measure on $\partial B_r(x)$.
A similar notion, with analogous properties, exists on a graph $G = (V, E)$. Given any finite connected subgraph $\emptyset \subset D \subset V$, a function $u : D \cup \partial D \to \mathbb{R}$ is discrete harmonic on D if it satisfies (2), where Δ is now the discrete Laplacian

$$\Delta u(x) = \sum_{y:(x,y)\in E} (u(y) - u(x)).$$

This implies the mean-value property

$$u(x) = \frac{1}{|\partial\{x\}|} \sum_{y:(x,y)\in E} u(y).$$

For any $A \subset \partial D$, the harmonic measure of A with respect to D, $\omega(x, A, D)$, is the unique discrete harmonic function on D with boundary values $\omega(x, A, D) = 1$ when $x \in A$ and $\omega(x, A, D) = 0$ when $x \in \partial D \setminus A$. Harmonic measure is closely connected with simple random walks via the relationship

$$\omega(x, \{y\}, D) = \mathbb{P}_x(X_\tau = y)$$

where X_n is a simple random walk starting from x and $\tau = \tau_D$ is the first exit time from D. If u is any harmonic function on D then

$$u(x) = \sum_{y\in\partial D} u(y)\omega(x, \{y\}, D) = \mathbb{E}_x u(X_\tau).$$