AN INTRODUCTION TO PERCOLATION THEORY AND ITS PHYSICAL APPLICATIONS

DANIEL MORRISON

Abstract. In this paper, we discuss the basic elements of percolation theory, beginning with the physical motivation for the abstraction and the essentials of the bond model on the L^D square lattice. The critical phenomenon is introduced, a fundamental concept in percolation theory. We then describe the general use of percolation theory to analyze physical phenomena. We examine one such application in depth: the connectivity of a neuron culture, which can be modeled as a percolation process.

Contents

1. Introduction

The classic motivating problem behind percolation theory examines the flow of water over a porous stone: under what conditions will the water travel successfully through the channels of the stone and reach the bottom? We consider the network of passages within the stone; some passages will be wide enough for the water to pass through, while others will be too narrow. If a series of sufficiently wide passages exists between the top of the stone and the bottom, the water will be able to pass through.

We can formalize this problem by viewing the interior of the stone as a graph in which vertices correspond to points within the stone and edges correspond to wide passages between points. Percolation theory seeks to determine if and when a path can form between points of this graph, i.e. when a traversable passage exists through the stone.

Date: DEADLINES: Draft AUGUST 15 and Final version AUGUST 29, 2016.

The percolation abstraction can be used to study the class of what might be termed "connectivity problems"—problems that hinge on the presence of a path through a random medium. Representative problems include the study of the passage of neurotransmitters between neurons, the calculation of the electrical resistance of a mixture of two metals, and the spread of a pathogen through a population. As will be discussed in section three, percolation theory has proven to be a useful tool in the analysis of such processes.

2. The percolation abstraction

2.1. Notation and definitions.

Definition 2.1. Let $\mathbb{Z}^D$ denote the set of ordered tuples $(a_1, \ldots, a_D)$ with $a_i \in \mathbb{Z}$.

Definition 2.2. We define L^D, the *square lattice of dimension* D, to be the graph with vertex set $V = \mathbb{Z}^D$ and edge set $E = \{(v_1, v_2) | v_1, v_2 \in V, \|v_1 - v_2\| = 1\}$.

Definition 2.3. The *origin* of L^D is the zero vertex $(0, \ldots, 0)$.

Definition 2.4. Let Ω denote the sample space

$$\prod_{e \in E} \{0, 1\}$$

where E is the set of edges in L^D. We call each point of Ω a *configuration* of L^D, and represent a configuration by a function $\alpha : E \to \{0, 1\}$. When $\alpha(e) = 1$, we say that edge e is *open*; when $\alpha(e) = 0$, edge e is *closed*. Likewise, we will say an edge e is *open in* α or *closed in* α to mean that $\alpha(e) = 1$ or $\alpha(e) = 0$, respectively.

For a finite set of edges s, let S be the set of all configurations α such that all edges in s are open in α. We call S a *cylinder set*. Let $\mathscr{F}$ be the collection of all such cylinder sets; we will take $\mathscr{F}$ to be the σ-algebra corresponding to the sample space Ω.

For each edge e in E, we define measure μ_e to be p if e is open and $1 - p$ if e is closed, where p satisfies $0 \leq p \leq 1$. More formally, we consider the Bernoulli scheme in which each edge e is an independent random variable taking values in $\{0, 1\}$ with corresponding probabilities $1 - p$ and p. μ_e is the Bernoulli measure given by $\mu_e(\alpha(e) = 1) = p$ and $\mu_e(\alpha(e) = 0) = 1 - p$.

Consider the measurable space $(\Omega, \mathscr{F})$. We define the corresponding probability measure as the product measure with density p on $(\Omega, \mathscr{F})$:

$$P_p = \prod_{e \in E} \mu_e$$

More intuitively, we can construct a configuration of L^D with *edge-probability* p by assigning each edge of L^D to be open with probability p and closed otherwise. Note that there is a one-to-one mapping of configurations of Ω to subsets of E: for a configuration α, we take the set $E_\alpha = \{e | e \in E, \alpha(e) = 1\}$. As such, we can represent a given configuration as the subgraph of L^D formed by the edges in E_α (see Fig. 1 for an example).

Definition 2.5. Let α be a configuration of L^D, and let v_a, v_b be vertices of L^D. An *open path* of length n exists between v_a and v_b in α if there exists a path of length n from v_a to v_b such that every edge in the path is open in α. Similarly, an *open cluster* is a set C of vertices such that for all $v_a, v_b \in C$, there exists an open

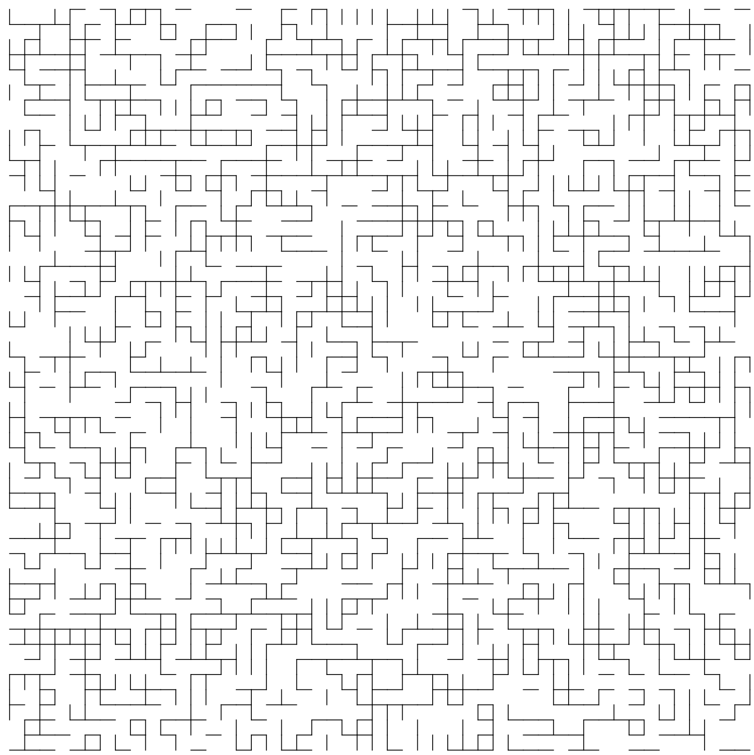

FIGURE 1. An example configuration over a subgraph of L^2 with edge-probability 0.50.

path between v_a and v_b. We say C is *finite* if $|C| < \infty$ and *infinite* if $|C| = \infty$. We denote the open cluster containing vertex v by $C(v)$, and the cluster containing the origin specially by C^*.

2.2. **Bond percolation approach.** The definitions given outline the *bond model*, the standard abstraction for percolation processes. While the motivating physical examples are concerned with finding open paths through finite spaces, bond percolation processes over infinite lattices necessitate a different approach. Representing a physical process as a large finite subsection of L^D, we make the assumption that the probability of an open path existing that connects opposite points in the subsection is related to the probability of an infinite open cluster existing in L^D. As such, we seek to determine the conditions necessary for the formation of an open cluster C with $|C| = \infty$. As L^D is invariant under translation, we can consider specifically the case when $|C^*| = \infty$, i.e. the origin is a member of an infinite cluster.

Intuitively, the existence of an infinite cluster is dependent on the proportion p of open edges selected from L^D. Here, we approach one of the foundational questions of percolation theory—for a given L^D, is there some value of p such that an infinite cluster is almost certain to occur? For $D = 1$, an infinite path, we of course require that $p = 1$, since any proportion of closed edges greater than 0 would necessarily fragment the one-dimensional lattice into many finite clusters. For $D \geq 2$, however, the answer is not as clear. In fact, however, as we shall shortly prove, for every

$D \geq 2$, there exists a critical value of p, denoted p_c, such that an infinite cluster is almost certain to occur for $p > p_c$ and almost certain to be absent for $p < p_c$.

2.3. **The critical phenomenon.**

Definition 2.6. Let $\Phi(p)$ be the probability that the origin of a configuration with edge-probability p is a member of an infinite cluster; that is, $\Phi(p) = P_p(|C^*| = \infty)$.

Note that $\Phi(0) = 0$, $\Phi(1) = 1$, and $\Phi(p)$ is nondecreasing.

Definition 2.7. The *critical probability* for a lattice of dimension D, denoted $p_c(D)$, is defined to be $\sup\{p|\Phi(p) = 0\}$.

Thus for all configurations over a D-dimensional lattice, we can express the probability of any configuration with edge-probability p having an infinite cluster as

$$\Phi(p) \begin{cases} = 0, \; p < p_c(D) \\ > 0, \; p > p_c(D) \end{cases} \tag{2.8}$$

We now define two variables relating to the connectivity of L^D.

Definition 2.9. Let $\sigma(n)$ be the number of paths of length n in L^D beginning at the origin.

Definition 2.10. Let $\lambda(D)$ be the *connective constant* of L^D, given by

$$\lambda(D) = \lim_{n\to\infty} (\sigma(n)^{\frac{1}{n}}).$$

Note that at every intersection of L^D, a path has at most $2D - 1$ possible next steps (since it cannot backtrack); thus $\sigma(n) \leq 2D(2D-1)^{n-1}$, and it follows that

$$\begin{aligned} \lambda(D) &\leq \lim_{n\to\infty} [2D(2D-1)^{n-1}]^{\frac{1}{n}} \\ &= \lim_{n\to\infty} [(2D)^{\frac{1}{n}}(2D-1)^{\frac{n-1}{n}}] \\ &= 2D - 1 \end{aligned}$$

Alternatively, we can define $\sigma(n)$ in terms of $\lambda(D)$ using little-o notation:

$$\lambda(D) + o(1) = \sigma(n)^{\frac{1}{n}}$$

$$\sigma(n) = [\lambda(D) + o(1)]^n \tag{2.11}$$

Here, $o(1)$ represents some unknown function f of smaller-than-constant order (i.e. $\lim_{n\to\infty} f(n) = 0$), allowing us to represent the increasingly small difference between $\lambda(D)$ and $\sigma(n)^{\frac{1}{n}}$ as n goes to infinity. This form will prove to be much more useful than that given in the initial definition.

We use these definitions to prove two lemmas relating to the bounds of $p_c(D)$.

Lemma 2.12. $p_c(D) \geq p_c(D+1)$.

Proof. We embed the L^D lattice in L^{D+1} by mapping each vertex $(a_1, \ldots, a_D)$ in L^D to $(a_1, \ldots, a_D, 0)$ in L^{D+1}. If there exists a configuration in L^D with $|C^*| = \infty$, then its embedding in L^{D+1} will also have $|C^*| = \infty$; thus the critical probability for the D-dimensional lattice is sufficiently large to ensure an infinite cluster on the $(D+1)$-dimensional lattice. □

Lemma 2.13. $\forall D \geq 2$, $p_c(D) > 0$.

Proof. Let $\gamma(n)$ be the number of open paths of length n beginning at the origin. Each path of length n has probability p^n of being an open path, giving the expected value of $\gamma(n)$ to be $E_p(\gamma(n)) = p^n\sigma(n)$. If $|C^*| = \infty$, then an open path of any and every length n leading from the origin can be found. Thus for all n,

$$\begin{aligned}\Phi(p) = P_p(|C^*| = \infty) &\leq P_p(\gamma(n) \geq 1) \\ &\leq E_p(\gamma(n)) \\ &= p^n\sigma(n)\end{aligned}$$

By equation (2.11), we have

$$\Phi(p) \leq [p\lambda(D) + o(1)]^n \tag{2.14}$$

Thus $\Phi(p) \to 0$ as $n \to \infty$ if $p\lambda(d) < 1$. To ensure $\Phi(p) > 0$ we have $p \geq \lambda(D)^{-1}$. Recall from Definition 2.10 that $\lambda(D) \leq (2D-1)$; it follows that $\lambda(D)^{-1} \geq (2D-1)^{-1} \geq 0$. Since $\Phi(p) > 0$ if and only if $p > p_c$, we have $p_c(D) \geq \lambda(D)^{-1} > 0$. □

We can now use these lemmas to prove the central theorem of bond percolation.

Theorem 2.15. *If* $D \geq 2$ *then* $0 < p_c(D) < 1$.

Proof. Based on Lemmas 2.12 and 2.13, it suffices to show that $p_c(2) < 1$. To do so, we consider the *planar dual* of the L^2 lattice:

Definition 2.16. Let G be a planar graph. We construct G', the *planar dual* of G, by placing a vertex within each face of G and adding an edge between any two vertices that lie in two adjacent faces; each face of G is a component of its dual.

For ease of visualization, we will define the planar dual of L^2 to be the graph with vertex set $V' = \{(x+\frac{1}{2}, y+\frac{1}{2}) | (x,y) \in \mathbb{Z}^2\}$ and edge set $E' = \{(v_1, v_2) | v_1, v_2 \in V', \|v_1 - v_2\| = 1\}$ (or more intuitively, the lattice formed by translating L^2 upwards and rightwards by one-half unit each). Each edge of L^2 is crossed by exactly one edge of its dual, and vice-versa; we use this to define the dual of a given configuration of L^2.

Definition 2.17. Let α be a configuration of L^2. We define α', the *planar dual configuration corresponding to* C as such: an edge of the planar dual of L^2 is open in α' if it crosses an edge of L^2 that is open in α, and closed otherwise. Note that α' will necessarily have the same edge-probability as α.

Consider a configuration α with edge-probability p in which the origin is a member of a finite open cluster C^*. It can be seen that in α', a cycle of closed edges must then exist that completely encloses C^*. Likewise, if a cycle of closed edges exists in α' surrounding the origin in α, then the origin must be a member of a finite open cluster. (The proof of these rather intuitive claims would take several more pages than allowable in this paper; see Kesten's "Percolation Theory for Mathematicians", p. 386, for an exhaustive treatment). In other words, $|C^*| < \infty$ if and only if the origin in α is contained within a cycle of closed edges in α'.

Let $\Gamma(n)$ be the number of closed cycles of length n in α' that enclose the origin in α. Note that any such cycle must pass through some vertex $(k + \frac{1}{2}, \frac{1}{2})$, with $0 \leq k < n$, since it surrounds the origin and thus must be $\frac{1}{2}$ "above" the origin for some positive value of k, and because $k > n$ would certainly necessitate a cycle length greater than n. Thus each cycle must contain a path of length $n-1$ beginning at a vertex of form $(k + \frac{1}{2}, \frac{1}{2})$. With n possibilities for the value of k, and using our earlier definition of $\sigma(n)$, we can state that $\Gamma(n) \leq n\sigma(n-1)$.

Let c be a (not necessarily closed) cycle in α' surrounding the origin. By the previous result, we have

$$\sum_c P_p(c \text{ is closed}) \leq \sum_{n=1}^{\infty} (1-p)^n n\sigma(n-1)$$

where the RHS expression represents an upper bound where every possible length-n cycle for every value of n is closed. Using an earlier substitution, we can rewrite the RHS as

$$\sum_{n=1}^{\infty} (1-p)n[(1-p)\lambda(2) + o(1)]^{n-1}$$

which is finite if $(1-p)\lambda(2) < 1$. Note that

$$\sum_c P_p(c \text{ is closed})$$

approaches 0 as $(1-p)$ approaches 0; thus we can find some number q, $0 < q < 1$, such that if $p > q$,

$$\sum_c P_p(c \text{ is closed}) \leq \frac{1}{2}.$$

Thus if $p > q$,

$$P_p(|C| = \infty) = P_p(\text{no circuit } c \text{ is closed}) \geq 1 - \sum_c P_p(c \text{ is closed}) \geq \frac{1}{2},$$

implying $p > q \geq p_c(2)$. □

A consequence of Theorem 2.15 is that the percolation process is separated into two distinct phases. In the *subcritical phase*, when $p < p_c$, any given vertex is almost surely a member of a finite open cluster, so all clusters must be finite. In the *supercritical phase*, when $p > p_c$, every vertex has a nonzero probability of belonging to an infinite cluster, which implies the almost certain existence of at least one infinite cluster somewhere in the configuration. We can summarize this result as follows:

Theorem 2.18. *For a configuration α with edge-probability p on a D-dimensional lattice, the probability $\Psi(p)$ that there exists an infinite open cluster satisfies*

$$\Psi(p) = \begin{cases} 0, \ p < p_c(D) \\ 1, \ p > p_c(D) \end{cases}$$

Proof. Observe that if a configuration α contains an infinite open cluster C, C remains infinite even if any finite number of edges in α are changed from open to closed, or vice-versa. As such, the presence of an infinite open cluster is a *tail event* over the set of edges in α, where each edge is either open or closed. Thus the distribution of Ψ follows Kolomogorov's zero-one law and takes only the values 0 and 1. For $p < p_c(D)$, by our previous result we have

$$\Psi(p) \leq \sum_{v \in \mathbb{Z}^D} P_p(|C(v)| = \infty) = 0$$

while for $p > p_c$, we have

$$\Psi(p) \geq P_p(|C^*| = \infty) > 0.$$

Since Ψ follows a zero-one distribution, we must have $\Psi = 1$. □

Unfortunately, Theorem 2.18 isn't any use in determining whether or not an infinite cluster will form when $p = p_c$. As it turns out, the behavior during this *critical phase* represents a major open question in percolation theory.

3. Physical applications

3.1. **Percolation over finite graphs.** The bond model, while a particularly elegant abstraction of the underlying class of random-medium problems, is quite useless when applied directly to those problems simply because it deals with infinite lattices and clusters. We must return our attention to finite graphs that are generally far more structurally irregular than the orderly D-dimensional lattice. As in the infinite case, we work with various configurations of the finite graph and search for large open clusters.

As the name implies, percolation problems in the physical world deal with the spread of some substance–a fluid, a virus, an electrical impulse–through a network. We will refer to any such substance as a *stimulus*, and to the introduction of a stimulus to a network at a particular vertex or vertices as the *input*. The *strength* of an input refers to the number of vertices at which the stimulus is introduced; a stronger input is introduced at more vertices than a weaker input. For example, one might model the spread of disease through a small town by a graph where each vertex represents a member of the community and an edge between two vertices is open if those two members are frequently in physical contact. The stimulus would be the virus causing the disease; its input would be the manner in which it is introduced into the population, which might be strong (the town's reservoir is contaminated, spreading the disease to many people at once) or weak (a single infected person enters the town, passes the virus on to one person, and leaves).

We introduce two concepts, the *giant cluster* and the *percolation transition*, that serve as a useful basis for approaching such processes.

Definition 3.1. Let G be a finite graph and let α be a configuration of G. An open cluster C in α is a *giant cluster* if a stimulus is almost certain to reach every vertex in C regardless of the strength of the input. (As we are dealing with a finite graph, the exact definition of "almost certain" may vary from one problem to the next, depending on the nature of the stimulus and the input.)

Definition 3.2. Let G be a finite graph. We say that G undergoes a *percolation transition* if a giant cluster forms in configurations of G with a sufficiently high proportion of open edges.

In the case of the town and the virus, consider a cluster of townspeople: if any one of them is infected, the rest of the cluster eventually will be as well. The larger the cluster, the greater the chance of any one of its members catching the virus. A giant cluster could be said to have formed when even under the weakest possible input, the cluster still has a 99% chance of infection.

3.2. **Representative physical problems.**

The orchard: Consider an orchard with trees planted in a grid pattern, with constant distance d between each tree and its four neighbors. The trees are vulnerable to a blight that spreads from one tree to another with probability p, where p is a known function of d. We want to maximize the number of trees planted in the orchard, but avoid the risk of losing significant portions of the orchard to the blight. In essence, we want to bring the orchard as close to the percolation transition point as possible without allowing a giant cluster to form.

Molecular bonding: Some substances undergo a sol-gel transition when enough bonds are formed between molecules. When few bonds are present, the substance exists as a viscous fluid composed of small macromolecules; as the number of bonds increases, the substance transitions to an elastic gel. The number of bonds present depends on some outside input, such as temperature; the sol-gel transition occurs when a cluster of molecules large enough to span the entire substance is formed.

3.3. **A closer examination: neural networks.** Consider a culture of neurons. Connections between neurons are directed. The connectivity of the culture can be controlled directly by the application of certain chemicals that dampen synaptic strength, severing the connections between individual neurons. The culture is stimulated via electric shocks applied to the entire network. Every neuron has an equal probability of responding to the shock, and the probability rises as the voltage of the shock increases. Additionally, even if a neuron does not respond to the shock itself, it may still fire if any of its input neurons do, and may in turn trigger its own output neurons to fire as well.

Let us represent this culture with a directed graph G. Let $f(V)$ be the probability that a neuron fires *in direct response* to a shock of voltage V. Let us further assume for simplicity that if a neuron fires, all of its output neurons will also fire.

Let $\Theta(f)$ be the probability that a neuron fires, either as a direct result of the shock or in response to an upstream neuron firing. $\Theta(f)$ clearly increases as the overall connectivity of the culture (or the number of edges in G) increases, since any neuron that fires will also activate all of its output neurons. We can express $\Theta(f)$ as follows:

$$\begin{aligned}\Theta(f) &= f + (1-f)P(\text{any input neuron fires}) \\ &= f + (1-f)\sum_{s=1}^{\infty} p_s(1-(1-f)^{s-1}) \\ &= 1 - \sum_{s=1}^{\infty} p_s(1-f)^s\end{aligned}$$

where p_s is the probability that a neuron has $s-1$ inputs.

At a certain degree of connectivity, a giant component emerges that always fires regardless of the strength of the shock. Once this occurs, the firing probability changes:

$$\begin{aligned}\Theta(f) &= g + (1-g)(f + (1-f)P(\text{any input neuron fires})) \\ &= 1 - (1-g)\sum_{s=1}^{\infty} p_s (1-f)^s\end{aligned}$$

where g is the proportion of neurons belonging to the giant cluster. Observe that as f approaches 0, $\Theta(f)$ approaches g, as only the giant cluster fires. This is naturally an approximation, since we are dealing with a finite number of neurons and the giant cluster will fail to fire for sufficiently low voltages.

By plotting the fraction of neurons responding as a function of voltage, one can observe the percolation transition quite clearly: the graph will follow an S-curve pattern, with the nearly-vertical middle section caused by the formation of the giant cluster at some critical voltage. When the culture is very well connected, the giant cluster is nearly the size of the entire culture; as connectivity decreases, the giant cluster decreases in size as well. The tails of the S-curve correspond to isolated clusters that happen to have very low or very high firing thresholds–perhaps a small clique of neurons isolated from the giant cluster, or a single neuron with no upstream inputs that must be activated directly by the shock.

Acknowledgments. It is a pleasure to thank my mentors, Oishee Banerjee and Randy VanWhy, for their guidance during this project, and for their willingness to let me dive into a subject I'd never heard of before this summer.

References

[1] A.J. Hildebrand. Asymptotic Methods in Analysis. University of Illinois at Urbana-Champaign. 2009.
[2] Geoffrey Grimmett. Percolation. Springer Science+Business Media New York. 1989.
[3] Geoffrey Grimmett and David Stirzaker. Probability and Random Processes. Oxford University Press. 1982.
[4] Harry Kesten. Percolation Theory for Mathematicians. Birkenhauser Boston. 1982.
[5] Ilan Breskin, Jordi Soriano, Elisha Moses, and Tsvi Tlusty. "Percolation in Living Neural Networks." Physical Review Letters 97.18 (2006): 188102. Web.
[6] JMP (http://tex.stackexchange.com/users/93513/jmp). Drawing percolation configuration with TikZ. URL: http://tex.stackexchange.com/questions/308017/drawing-percolation-configuration-with-tikz?rq=1
[7] Peter D. Hoff. Measure and probability. September 26, 2013.