Linear Algebra and Bootstrap Percolation

Abel Romer

A brief outline

Lecture 1

We introduce the concepts of hypergraphs and bootstrap percolation. We define bootstrap percolation on the hypergraph. We define and describe two classes of hypergraphs that generalize the ideas of paths and complete graphs. We present a percolating set on these hypergraphs and discuss its properties. The lecture ends with the presentation of a technique from linear algebra that will be used to bound the minimum size of percolating sets on hypergraphs.

Lecture 2

This lecture explores three applications of the lemma from Lecture 1. We use the lemma to construct vector spaces corresponding to percolating sets on different classes of the hypergraph $\mathcal{P}(n, d, t, r)$. We see that these vector spaces allow us to prove that the percolating set from Lecture 1 is an extremal example for these hypergraphs.

Importance

From a pedagogical perspective, this lecture series presents an opportunity to engage deeply with the challenging and important skills of visualization and generalization in mathematics. The material deals directly with the structure of semi-regular, high-dimensional objects, and requires the student to shift fluently between physical examples to abstract generalizations. From a mathematical perspective, the paper upon which these lectures are based introduces techniques from linear algebra to the popular problem of bootstrap percolation. This result from this paper somewhat generalizes and meshes with existing results of bootstrap problems on hypergraphs, and hints at the possibility of future breakthroughs in the problem using similar strategies.

I would suggest presenting this topic in the weak saturation and/or discrete geometry sections of the course. The problem has a similar flavor, and the use of vector spaces as a tool for bounding the size of structures is somewhat thematic.

Organization

I elected to begin with a number of examples of hypergraphs and percolating sets because I see the main stumbling block with this material to be a weak foundational understanding of the structures at play. By presenting small examples of hypergraphs early and (relatively) often, the later proofs in Lecture 2 are easier to follow without constantly double-checking indices and drawing pictures to keep track of which vertices or vectors are being referred to.

References

[1] Balogh, J., Bollobas, B., Morris, R., Riordan, O. (2012). Linear algebra and bootstrap percolation. Journal of Combinatorial Theory, Series A, 119(6), 1328-1335.

Lecture 1

This is part one of a two-part lecture series on the use of techniques from linear algebra in hypergraph bootstrap percolation. The first lecture will introduce the process of bootstrap percolation in hypergraphs, discuss the structure of two particular hypergraphs, examine percolating sets within these graphs, and present a useful lemma. We will use the following color key to annotate the lectures:
[Spoken: Text written in this color will be spoken aloud to the class, but not written on the board.]
[Discussion: Text written in this color will be posed to the class as discussion questions/material.]

Some necessary concepts and definitions

[Spoken: Recall that a hypergraph is a set of vertices and edges, where each edge is an element of the power set of the vertex set, excluding the empty set. Draw a simple example. Observe that it is possible to define a hypergraph $\mathcal{H}$ with two graphs G and H, where $V(\mathcal{H}) = V(G)$ and $E(\mathcal{H}) = \{V(H) : \mathcal{H}[V(H)] \cong H\}$.]

Definition 1. A hypergraph $\mathcal{H} = (V, E)$ is a graph with vertex set V and edge set $E \subseteq \mathcal{P}(V) \setminus \{\emptyset\}$.

[Spoken: Bootstrap percolation can be thought of as an infection spreading through a population subject to particular constraints. In the context of hypergraphs, we say that a healthy vertex becomes infected when it is the only healthy vertex in an edge.]

Definition 2. Let $A \subset V(\mathcal{H})$ be a set of infected vertices in a hypergraph $\mathcal{H}$. Set $A_0 = A$. For each $t \geq 0$, let

$$A_{t+1} = A_t \cup \{u : \exists\, S \in \mathcal{H} \text{ where } S \setminus A_t = \{u\}\}.$$

Define $[A]_{\mathcal{H}} = \cup_{t \geq 0} A_t$, and say that A percolates if $[A]_{\mathcal{H}} = V(\mathcal{H})$.

Definition 3. Let $\mathcal{H}$ be a hypergraph. Then

$$m(\mathcal{H}) = \min\{|A| : A \subset V(\mathcal{H}) \text{ and } [A]_{\mathcal{H}} = V(\mathcal{H}).$$

[Discussion: What is a percolating set for a 3×3 graph with edges corresponding to squares? Play around with an example to gauge understanding in the class. Hint at structural patterns for percolating sets on hypergraphs.]
[Spoken: We are interested in two particular classes of hypergraphs. We can think of these, broadly, as high-dimensional hypergraph versions of a complete graph and a path. We shall see the formal definition, and then some examples.]

Definition 4. Let P_n^d be a graph with vertex set $V = [n]^d$ and edge set

$$E = \{uv : u \text{ and } v \text{ differ by one in exactly one coordinate}\}.$$

Then $\mathcal{P}(n, d, t, r)$ is the graph with vertex set $V(\mathcal{P}) = [n]^d$ and edge set

$$E(\mathcal{P}) = \{S \subset [n]^d : P_n^d[S] \cong P_t^r \text{ and } P_n^d[S] \text{ is “straight”}\},$$

where “straight” means that the induced copy of P_t^r does not have any kinks or bends.

Example 5. The hypergraphs $\mathcal{P}(3, 2, 2, 2)$ and $\mathcal{P}(3, 2, 3, 1)$. Note that the red vertices induce a “bent” $P_3^1 = P_3$ that is not in the edge set.

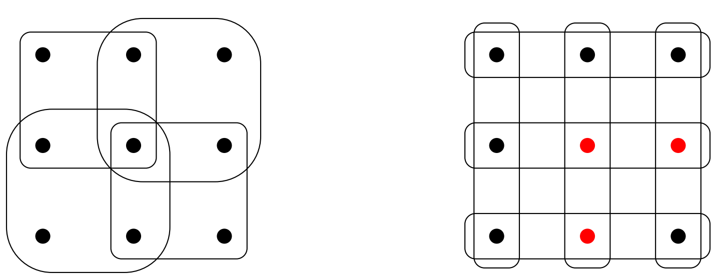

[Discussion: What is an example of a percolating set in $\mathcal{P}(3,2,3,1)$? Again, encourage the class to suggest more general patterns for percolating sets.]

Definition 6. Let K_n^d be the graph with vertex set $[n]^d$ and edge set

$$E(K_n^d) = \{uv : u \text{ and } v \text{ differ (by any amount) in exactly one coordinate}\}.$$

Then $\mathcal{K}(n,d,t,r)$ is the graph with vertex set $[n]^d$ and edge set

$$E(\mathcal{K}) = \{S \subset [n]^d : K_n^d[S] \cong K_t^r\}.$$

Example 7. Consider the following example of K_3^2. [Spoken: Note that each column and row is a K_3.]

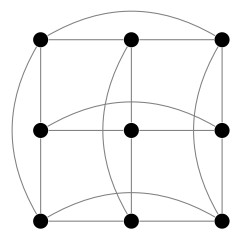

Edges of $\mathcal{K}(3,n,t,r)$ are all sets of vertices that induce a K_r^t in K_3^2. [Discussion: What are the edges of $\mathcal{K}(3,2,3,1)$?]

[Discussion: Start building up intuition for the structure of more complicated hypergraphs. What does $\mathcal{P}(n,d,2,r)$ look like? How would we describe the edge sets?]

Finding an upper bound

[Spoken: We are going to define a set of vertices U that percolates in $\mathcal{P}(n,d,t,r)$. This set will, in effect, lie on the hyperplane "faces" and hyperline "edges" of $\mathcal{P}$, but will have slightly different structures depending on the values of t and r. We shall begin to explore examples of U today, and see additional examples in Lecture 2.]

Claim 8. The following set

$$U = \{(u_1, \ldots, u_d) \in [n]^d : |\{i : u_i \geq t\}| \leq r-1\}$$

percolates in $\mathcal{P}(n,d,t,r)$.

[Discussion: What does U look like? What is the set U in $\mathcal{P}(n,d,2,r)$?]

Proof. Let v be a vertex in $V(\mathcal{P}(n,d,t,r)) \setminus U$, and let $|v|$ be the sum of its coordinates. We proceed by induction on $|v|$.

Base case: Suppose $|v| = rt$. Since $v \notin U$, v must have r coordinates of size t. Note that every vertex obtained by reducing any of these r coordinates is in U, and that these vertices (together with v) induce an edge $S = K_t^r$ in K_n^d. Since S contains exactly one vertex v not in U, v will be infected.

Induction hypothesis: Assume all vertices v with $|v| = k-1 \geq rt$ have been infected.

Induction step: Consider $|v| = k$. By our induction hypothesis, all vertices obtained by reducing the value of coordinates of v have size $\leq k-1$, and are therefore infected. Furthermore, the set S of vertices obtained by decrementing any r coordinates of size $\geq t$ induces a K_t^r in K_n^d. Therefore, S is an edge containing exactly one uninfected vertex v, and so v becomes infected.

Since every vertex v becomes infected, the set U percolates in $\mathcal{P}(n,d,t,r)$, and so $m(\mathcal{P}) \leq |U|$. □

Observation. Note that $\mathcal{P} \subset \mathcal{K}$, and so $m(\mathcal{K}) \leq m(\mathcal{P}) \leq |U|$.

[Discussion: Do we think this is optimal? Why or why not?]

A helpful lemma

[Spoken: We will present a lemma that will allow us to prove a lower bound on $m(\mathcal{H})$.]

Lemma 9. *Let $\mathcal{H}$ be an arbitrary hypergraph. Suppose that we can find a vector space W spanned by vectors $\{f_v : v \in V(H)\}$ such that, for every edge $S \in H$, we have a linear dependence with all coefficients non-zero. Then*

$$m(\mathcal{H}) \geq \dim(W).$$

Proof. Suppose there exists a set $A \subset V(\mathcal{H})$ that percolates in $\mathcal{H}$. This implies that there exists a sequence of vertices $v_1, \ldots, v_\ell$ in $V(\mathcal{H}) \setminus A$ such that each v_i is in an edge S_i in $A_i = A \cup \{v_j : j \geq i\}$. [Spoken: In other words, every step in the percolation completely infects a new edge. This must be the case, as otherwise the infection could not spread throughout the entire graph.] Let W_i be the span of the vectors $\{f_v : v \in A_i\}$. [Spoken: Recall that A_i is the set of infected vertices at the ith step in the infection.] Note that since $S_i \setminus v_i$ is a subset of A_{i-1}, and all vectors in S_i satisfy a linear dependence with non-zero coefficients (by hypothesis), we can write f_v as a linear combination of vectors in W_{i-1}. However, this implies that $W_{i-1} = W_i$, and so $W_0 = W_\ell$. Since A is assumed to percolate, $W_\ell = W$. Therefore, $W_0 = W$, and since W_0 is spanned by $|A|$ vectors, $|A| \geq \dim(W_0) = \dim(W)$. □

[Spoken: In the following lecture, we will see how this lemma can to be used to prove that the set U is tight for 3 particular cases of hypergraph.]

Lecture 2

This is the second lecture of a two-part lecture series on bootstrap percolation and linear algebra. In this lecture, we use the lemma from the previous lecture to find tight examples of percolating sets on three classes of hypergraph.
[Spoken: We shall examine three cases for the hypergraph $\mathcal{P}$. These can be loosely grouped into the cases where the edges of $\mathcal{P}$ are hypercubes of the same dimension as $\mathcal{P}$, hypercubes of a lower dimension, and "larger" hypercubes of the same dimension.]

Proposition 10. *For $\mathcal{H} = \mathcal{P}(n, d, 2, d)$, $m(\mathcal{H}) \geq |U|$.*

[Discussion: Explore what U looks like. Point out that U is the set of vertices with at least one coordinate equal to 1. Ask what this looks like on a hypercube.]

Question 11. What does the set U look like?

Proof. We shall apply Lemma 9. Specifically, we show that there is a vector space W, spanned by $|V(\mathcal{H})|$ vectors such that for every edge $S \subset V(\mathcal{H})$, the vectors in S satisfy the dependence condition of Lemma 9, and $\dim(W) = |U|$.

Let $\{e_u : u \in U\}$ be a set of arbitrary linearly independent vectors. For each vertex $v \in V(\mathcal{H})$ and $i \in [d]$, let $\pi_i(v)$ be the projection of v onto the face $\{\mathbf{x} \in [n]^d : x_i = 1\}$. [Spoken: Note that each vertex is projected onto vertices in U.] We define

$$f_v = \sum_{i=1}^{d} e_{\pi_i(v)}.$$

We now show that for each edge $S \subset V(\mathcal{H})$, the vectors $\{f_v : v \in S\}$ satisfy the dependence condition of Lemma 9, and that the dimension of W—the space spanned by $\{f_v : v \in V(\mathcal{H})\}$—is $|U|$.

We first show that

$$\sum_{v \in S} c_{S,v} f_v = 0,$$

where all coefficients $c_{S,v}$ are non-zero. Re-writing f_v and re-ordering our sums, we obtain:

$$\sum_{v \in S} c_{S,v} f_v = \sum_{v \in S} \left(c_{S,v} \sum_{i=1}^{d} e_{\pi_i(v)} \right) = \sum_{v \in S} \sum_{i=1}^{d} c_{S,v} e_{\pi_i(v)} = \sum_{i=1}^{d} \sum_{v \in S} c_{S,v} e_{\pi_i(v)}.$$

Note that we can break apart S into pairs of vertices that differ in exactly one coordinate. [Spoken: This is because each edge S is the vertex set of a d-dimensional hypercube.] This gives pairs $\{u, u'\}$ such that $\pi_i(u) = \pi_i(u')$. Setting $c_{S,u} = c_{S,u'}$, we are able to cancel these pairs and guarantee that the above sum is zero.

We show that $\dim W = |U|$. Since, by definition, $\{e_u : u \in U\}$ is a set of linearly independent vectors, $\{e_u : u \in U\}$ spans a space of dimension $|U|$. Therefore, it is sufficient to show that the space spanned by $\{f_v : v \in V(\mathcal{H})\}$ is the same as the space spanned by $\{e_u : u \in U\}$. This follows from the observation that $\{f_u : u \in U\}$ is a set of linearly independent vectors.

We have found a set of vectors $\{f_v : v \in V(\mathcal{H})\}$ satisfying the dependence condition on edges sets, which spans a space of dimension $|U|$. Therefore, $m(\mathcal{H}) \geq |U|$. $\square$

Proposition 12. *For $\mathcal{H} = \mathcal{P}(n, d, 2, r)$ with $1 \leq r \leq d$, $m(\mathcal{H}) \geq |U|$.*

[Discussion: Again, encourage the class to describe the set U. Suggest that the structure of U is similar to what we saw in the previous proposition. "Similarly, the initial infection lies on lower dimensional 'faces' of a hypercube."]

Question 13. How can we describe the set U?

Proof. Again, we let $\{e_u : u \in U\}$ be a set of linearly independent vectors. Observe that U is the set of $(r-1)$-dimensional faces of $\mathcal{H}$. In other words, U is the set of vertices containing at least $d-r+1$ coordinates equal to 1, or $\{\mathbf{x} \in [n]^d : x_{i_1} = \cdots = x_{i_{d-r+1}} = 1\}$. As before, let $\pi_T(v)$ be the projection onto a face T corresponding to a particular set of 1-valued coordinates, and define

$$f_v = \sum_T e_{\pi_T(v)}.$$

[Spoken: Note that we are summing over all subsets of $[d]$ of size $d-r+1$.] Again, we show that for each edge S, the vectors $\{f_v : v \in S\}$ satisfy the dependence condition of Lemma 9, and the dimension of the vector space W spanned by these vectors is $|U|$.

As each edge S is the vertex set of a r-dimensional hypercube, we may again pair up the vectors in S and choose non-zero coefficients to obtain a linear dependence. We also observe that $\{f_u : u \in U\}$ is a set of linearly independent vectors, and so $\dim(W) = |U|$. Again applying Lemma 9, we see that $m(\mathcal{H}) \geq |U|$. $\square$

[Spoken: The next proposition and proof leverage many of the same strategies we saw in the previous two proofs. Here, however, the notation and indexing become quite tricky, and rather than get lost in the particulars, we shall present the over-arching strategy and ideas.]

Proposition 14. *For $\mathcal{H} = \mathcal{P}(n, d, t, d)$ with $t \geq 3$, $m(\mathcal{H}) \geq |U|$.*

[Discussion: Again, ask the class to describe the set U. Observe that the structure is similar to that of Proposition 10. "Because we are dealing with larger edges, simply taking the faces of the hypercube is not sufficient to spawn infections."]

Question 15. Again, what does the set U look like?

Proof. Let $\{e_u : u \in U\}$ be a set of linearly independent vectors. Note that U is a sort of "extra thick" version of the set from Proposition 10. [Spoken: Recall from the proof of Proposition 10 the strategy of pairing up vertices in an edge S and cancelling out their corresponding vectors. We shall employ a similar technique.] Consider a line of vertices $v(1), \ldots, v(n)$ in $[n]^d$ resulting from varying the kth coordinate of v. Let $v(j_1), \ldots, v(j_t)$ be t consecutive vertices along this line. Observe that these t vertices all lie within some edge S. Define

$$f_v = \sum_{k=1}^{d} f_v^{(k)},$$

where $f_v^{(k)}$ is a linear combination of the vectors $e_{u_1}, \ldots, e_{u_{t-1}}$ along the line k. [Spoken: We choose this linear combination to ensure that any t consecutive vertices in a line satisfy a linear dependence with non-zero coefficients. We also ensure that $\{f_u : u \in U\}$ is a set of linearly independent vectors.]

To see that the vectors f_v for $v \in S$ are linearly dependent with non-zero coefficients, observe that the vectors in S can be separated into groups that form lines of length t. As we have chosen our vectors to satisfy such a dependence, we can select non-zero coefficients such that the vectors in each of these lines sum to zero. This thereby guarantees that each set S satisfies the dependence condition of Lemma 9. Additionally, as we have chosen $\{f_u : u \in U\}$ to be linearly independent, we again have that $\dim(W) = |U|$. Therefore, by Lemma 9, we have that $m(\mathcal{H}) \geq |U|$. $\square$

[Spoken: Recall from Lecture 1 that the set U percolates in both $\mathcal{P}(n, d, t, r)$ and $\mathcal{K}(n, d, t, r)$. Therefore, our results in Propositions 10, 12 and 14 are in fact extremal.]

Claim 16.

$$|U| = \sum_{s=0}^{r-1} \binom{d}{s} (t-1)^{d-s} (n-t+1)^s.$$

[Spoken: Note that we haven't actually discussed the size of U. We will not do the calculations in this lecture series, but the interested attendee is encouraged to work through the counting argument. It is a helpful exercise to understand the structure of U.]