A proof of the Kahn–Kalai conjecture

Talk by Jinyoung Park

Notes by Sanjana Das

September 16, 2022

§1 Definitions

We'll begin by restating the definitions from yesterday's talk.

- X is a finite set, and 2^X the set of subsets of X.
- μ_p is the p-biased product probability measure on 2^X — for each subset $A \subseteq X$, we have $\mu_p(A) = p^{|A|}(1-p)^{|X \setminus A|}$.
- $X_p \sim \mu_p$ (meaning that X_p is the random variable with distribution given by μ_p) — so X_p is a p-random subset of X, meaning that we choose every element of X with probability p independently.
- $\mathcal{F} \subseteq 2^X$ is an *increasing property* — whenever $A \in \mathcal{F}$, its *up-set* $\langle A \rangle := \{B \subseteq X \mid B \supseteq A\}$ must also be contained in $\mathcal{F}$.

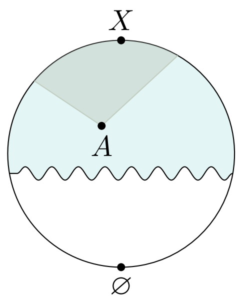

- We define $\mu_p(\mathcal{F}) := \sum_{A \in \mathcal{F}} \mu_p(A)$.
- If $\mathcal{F}$ is not $\varnothing$ or 2^X, then as p increases from 0 to 1, so does $\mu_p(\mathcal{F})$. So there exists a unique $p_c(\mathcal{F})$ where $\mu_p(\mathcal{F}) = 1/2$; this is called the threshold for $\mathcal{F}$.

Theorem 1.1 (Kahn–Kalai Conjecture)
There exists a constant K such that for all X and for all $\mathcal{F} \supset 2^X$,

$$p_c(\mathcal{F}) \leq K \cdot q(\mathcal{F}) \cdot \log \ell(\mathcal{F}),$$

where $q(\mathcal{F})$ is the *expectation threshold* and $\ell(\mathcal{F})$ the size of the largest minimal element of $\mathcal{F}$.

We'll now review the definition of the expectation threshold. In what follows, assume that $A, S \subseteq X$ and $\mathcal{F}, \mathcal{G} \subseteq 2^X$.

Definition 1.2. We say that S covers A if $S \subseteq A$.

In other words, S covers A if A is contained in the up-set $\langle S \rangle$.

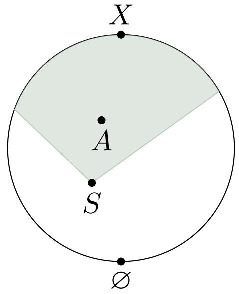

Definition 1.3. We say $\mathcal{G}$ covers $\mathcal{F}$ if for all $A \in \mathcal{F}$, there exists $S \in \mathcal{G}$ such that S covers A.

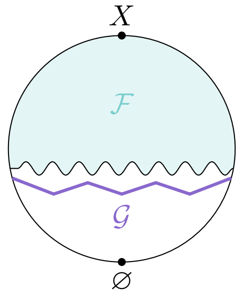

In other words, $\mathcal{G}$ covers $\mathcal{F}$ if $\langle \mathcal{G} \rangle \supseteq \mathcal{F}$ (where $\langle \mathcal{G} \rangle$ is the union of the up-sets $\langle A \rangle$ over all $A \in \mathcal{G}$).

Yesterday, we saw the observation that $p_c(\mathcal{F}) \geq q$ if there exists $\mathcal{G}$ that covers $\mathcal{F}$ such that

$$\sum_{S \in \mathcal{G}} q^{|S|} \leq \frac{1}{2}. \tag{$*$}$$

If some $\mathcal{G}$ satisfies ($*$), then we say $\mathcal{G}$ is *q-cheap*.

Definition 1.4. The expectation threshold $q(\mathcal{F})$ is defined as the maximal q for which there exists a q-cheap cover $\mathcal{G}$.

§2 Overview of the Proof

The theorem we'll actually prove is the following standard reformulation of the Kahn–Kalai conjecture:

Theorem 2.1 (Reformulation of the Kahn–Kalai Conjecture)
There exists $L > 0$ such that for all ℓ-bounded $\mathcal{H}$, if $p > q(\langle \mathcal{H} \rangle)$, then if we let $m = Lp \log \ell \cdot |X|$,

$$\mathbb{P}(X_m \text{ contains a member of } \mathcal{H}) = 1 - o(1) \text{ as } \ell \to \infty.$$

To say $\mathcal{H}$ is ℓ-bounded means that for all $A \in \mathcal{H}$ we have $|A| \leq \ell$. Here X_m means a random subset of X with m elements.

For some intuition on why this implies the Kahn–Kalai conjecture, think of $\mathcal{H}$ as the collection of minimal elements of $\mathcal{F}$, so that $\langle \mathcal{H} \rangle = \mathcal{F}$. Then we're considering random subsets X_m with $Lp \log \ell \cdot |X|$ elements; this is similar to choosing random subsets where we choose each element with probability around $Lp \log \ell$. Now if X_m is very likely to contain a member of $\mathcal{H}$, it's also very likely to contain a member of $\mathcal{F}$ (since the members of $\mathcal{H}$ are all members of $\mathcal{F}$). So this provides an upper bound $p_c(\mathcal{F}) \lesssim Lp \log \ell$ as well. (In the statement ℓ is fixed, but L is an absolute constant, so we can take ℓ to be the size of the largest minimal element of $\mathcal{F}$.)

To visualize this, we can think of X as a universe, and $\mathcal{H}$ as a collection of subsets of X.

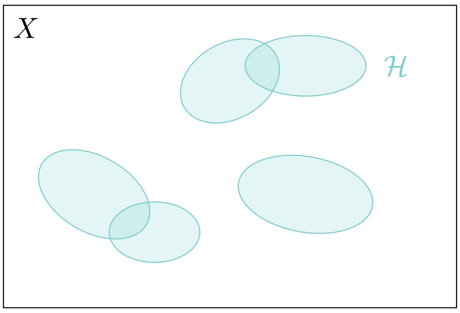

We then sprinkle in m random elements, and the statement says that typically these m elements contain some member of $\mathcal{H}$.

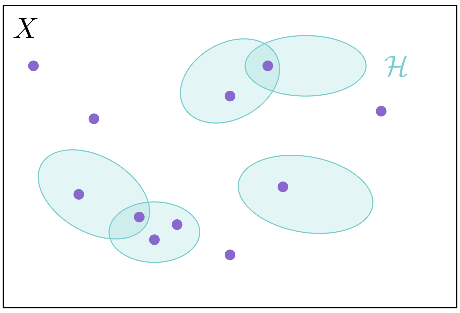

First, here is an overview of the proof.

(1) We'll use W to denote X_m. We choose W little by little — at each step we choose a W_i with $|W_i| = Lp|X|$ at random (such that the W_i are disjoint), and we set $W = W_1 \sqcup W_2 \sqcup \cdots$. We can choose as many as $\log \ell$ of these intermediate sets W_i.

(2) As we choose W little by little, $\mathcal{H}$ will evolve as well — we'll have $\mathcal{H} = \mathcal{H}_0 \to \mathcal{H}_1 \to \mathcal{H}_2 \to \cdots$.

(3) In the end, we want to have $W \supseteq S$ for some $S \in \mathcal{H}$ with high probability.

(4) We'll apply a randomized algorithm (where the randomness comes from the choice of W_i) — we iteratively produce a *partial cover* $\mathcal{U}(W)$ of $\mathcal{H}$, by building a partial cover $\mathcal{U}_i(W_i)$ at each step of the algorithm and taking $\mathcal{U}(W) = \bigcup \mathcal{U}_i(W_i)$. (A *cover* $\mathcal{G}$ of $\mathcal{F}$ would mean that every subset in $\mathcal{F}$ contains some subset in $\mathcal{G}$; a *partial cover* means that we only cover some part of $\mathcal{F}$, not necessarily all of it.)

(5) The main point of the proof is that our partial cover $\mathcal{U}(W)$ will be p-cheap with high probability.

(6) When the algorithm terminates (we'll see the termination condition later), either:

  (1) $\mathcal{U}$ entirely covers $\mathcal{H}$, or

  (2) W contains an element $S \in \mathcal{H}$.

If we believe this, then we're done with the proof — consider the sample space for the choice of W. Then by (5), most of the time $\mathcal{U}$ is p-cheap; it's very unlikely that $\mathcal{U}$ is expensive.

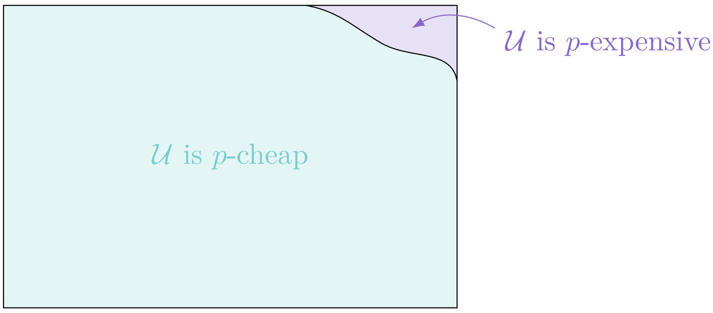

But now we can apply our assumption that $p > q(\langle \mathcal{H} \rangle)$ — we know that $q(\langle \mathcal{H} \rangle)$ is the largest q that admits a q-cheap cover of $\mathcal{H}$. So if $p > q(\langle \mathcal{H} \rangle)$, then there does not exist a p-cheap cover of $\langle H \rangle$. This means if $\mathcal{U}$ covers $\mathcal{H}$, then it must be expensive!

But it is unlikely that $\mathcal{U}$ is expensive. So when the algorithm terminates, (1) must be unlikely, which means (2) occurs with high probability. But (2) is exactly what we're looking for.

§3 Constructing $\mathcal{U}_i(W_i)$

First we'll describe how to construct $\mathcal{U}_i(W_i)$ in a given step; we'll then iterate this construction at most $\log \ell$ times.

In this step, our host hypergraph is $\mathcal{H}_{i-1}$; suppose $\mathcal{H}_{i-1}$ is s-bounded for some s (initially $\mathcal{H}$ is ℓ-bounded, but s will change as $\mathcal{H}$ changes).

Suppose we've chosen W_i (this is done at random). Then we examine all $S \in \mathcal{H}_{i-1}$ and decide whether we want to cover them or not. First, for each S, we look at all elements of $\mathcal{H}_{i-1}$ that sit inside $W_i \cup S$.

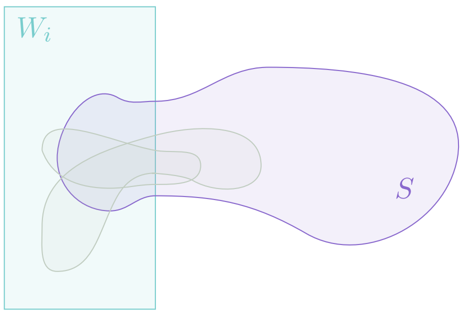

Now of these subsets in $\mathcal{H}_{i-1}$ inside $W_i \cup S$, we let S' be the subset with minimal $|S' \setminus W_i|$ (if there's multiple, we choose arbitrarily).

Definition 3.1. Given S and W_i, we define $T = S' \setminus W_i$ as the minimal (S, W_i)-fragment.

This minimal fragment will be the key gadget of the proof. There are a few key observations:

(1) There must exist some $\hat{S} \in \mathcal{H}_{i-1}$ with $\hat{S} \subseteq W_i \cup T$ (by definition — since T is the piece of some S' that lies outside W_i).

(2) For every $\hat{S} \subseteq W \cup T$, we must have $T \subseteq \hat{S}$ — otherwise, this would violate the minimality of T, as $\hat{S} \setminus W$ would be a smaller fragment:

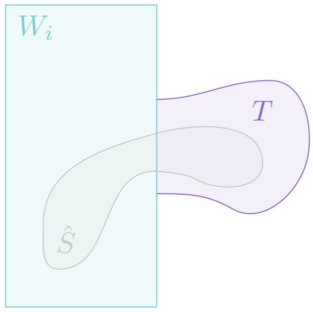

So we have the following picture:

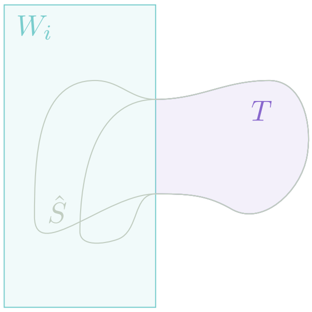

This will be the core of what makes our partial cover cheap.

Now we say that S is *good* if T is large — if $|T| \geq 0.9s$. In this case, we put T in $\mathcal{U}_i$. Otherwise, if T is small then it's not affordable, so we don't place it in $\mathcal{U}_i$, and instead we place T in our next hypergraph $\mathcal{H}_i$.

So in this step, for each $S \in \mathcal{H}_{i-1}$ we either add its minimal fragment to $\mathcal{U}_i(W_i)$ — which covers S — or we replace it with a subset whose bound is smaller by a factor of at least 0.9. In particular, this is why our algorithm will perform at most $\log \ell$ steps — at the start $\mathcal{H}$ is ℓ-bounded, after one step it's 0.9ℓ-bounded, after two steps it's $0.9^2\ell$-bounded, and so on.

§4 Termination Conditions

Now that we've described how to construct our partial cover, we'll describe when the algorithm terminates.

We terminate the algorithm as soon as some set has minimal fragment $T = \varnothing$. If this occurs, then we must have some $\hat{S} \in \mathcal{H}_{i-1}$ sitting entirely in W_i.

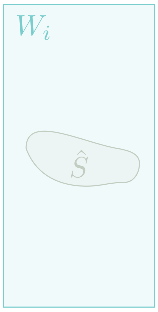

Then we claim we've reached our second goal of (2) — that $W = W_1 \sqcup W_2 \sqcup \cdots$ contains some element of $\mathcal{H}$. To see this, note that since $\hat{S}$ was in $\mathcal{H}_{i-1}$, it must have been the minimal fragment of some set in $\mathcal{H}_{i-2}$ (since we only ever insert minimal fragments into $\mathcal{H}_{i-1}$). This means we had a set in $\mathcal{H}_{i-2}$, and then W_{i-1} "ate" a piece of it and left us with $\hat{S}$.

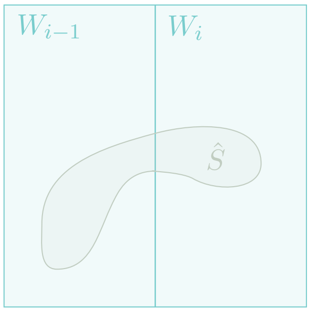

But *this* set is in $\mathcal{H}_{i-2}$, so it's the minimal fragment of some set in $\mathcal{H}_{i-3}$. We can keep extending backwards to get a set $S \in \mathcal{H}$ from the beginning of the process:

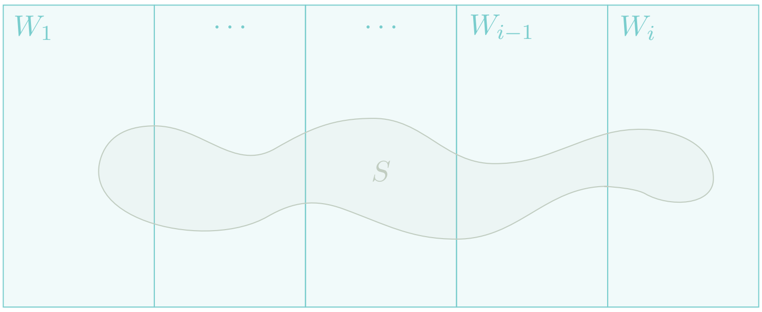

So then we started off with $S \in \mathcal{H}$, and some part of it got eaten by W_1, then W_2, and so on; and its final part got eaten by W_i. This means $W = W_1 \sqcup W_2 \sqcup \cdots$ covers S.

On the other hand, suppose this never happens. Then we keep running the process until there's nothing left in $\mathcal{H}$. But in every step, we look at all S in $\mathcal{H}_{i-1}$, and either we add their minimal fragment T to $\mathcal{U}_i(W_i)$ — which covers S — or we add T to $\mathcal{H}_i$ (and if we later cover T, that set covers S as well). So if $\mathcal{H}$ ends up empty, then $\mathcal{U}(W) = \mathcal{U}_1(W_1) \cup \mathcal{U}_2(W_2) \cup \cdots$ covers all sets S originally in $\mathcal{H}$. So we've reached our first goal of (1).

§5 p-Cheapness of $\mathcal{U}(W)$

Now there's one remaining piece of the proof — that our partial cover $\mathcal{U}(W)$ is cheap with high probability. The key point is the following:

Lemma 5.1

Let $|X| = n$, $w = Lpn$, and suppose $\mathcal{H}$ is s-bounded. Then

$$\sum_{W_i \in \binom{X}{w}} \sum_{U \in \mathcal{U}_i(W_i)} p^{|U|} < \binom{n}{w} L^{-0.8s}.$$

This implies that the *average* cost of $\mathcal{U}_i(W_i)$ (over all possibilities for W_i) is less than $L^{-0.8s}$. This can be used to show that the average cost of $\mathcal{U}(W)$ is small as well, and then Markov's inequality shows that $\mathcal{U}(W)$ is usually p-cheap.

Proof. We use double-counting. Recall that if $U \in \mathcal{U}_i(W_i)$, then we must have $|U| \geq 0.9s$ (since this is our condition for adding T to our partial cover). For simplicity assume $|U| = 0.9s$ (the calculations are messier in the general case). Then the left-hand side becomes

$$p^{0.9s} \sum_{W_i \in \binom{X}{w}} \sum_{U \in \mathcal{U}_i(W_i)} 1.$$

But the double summation simply counts pairs (W_i, U) where U is the minimal fragment of somebody (satisfying the size condition) — so we just want to count pairs $(W_i, T(S, W_i))$ where $W_i \in \binom{X}{w}$ and $S \in \mathcal{H}$, and $|T(S, W_i)| = 0.9s$. But this means

$$p^{0.9s} \sum_{W_i \in \binom{X}{w}} \sum_{U \in \mathcal{U}_i(W_i)} 1 \leq p^{0.9s} \cdot \binom{n}{w + 0.9s} \cdot 2^s.$$

To see this, there are $\binom{n}{w+0.9s}$ ways to choose $W \cup T$. Once we've chosen $W \cup T$, we can use our observations about minimal fragments — for T to be somebody's minimal fragment, we must have some $\hat{S} \in \mathcal{H}$ with $\hat{S} \subseteq W \cup T$, and then T must sit inside $\hat{S}$. There's at most 2^s subsets of $\hat{S}$, so at most 2^s possible T.

Now this sum is at most

$$p^{0.9s} \cdot \binom{n}{w} \cdot (Lp)^{-0.9s} \cdot 2^s = \binom{n}{w} \cdot L^{-0.9s} \cdot 2^s < \binom{n}{w} L^{-0.8s}.$$

So we are done. □