Sets systems and their fundamental extremal problems

Peter Hajnal

Bolyai Institute, SZTE, Szeged

2023 Fall

Sperner System

Definition

$\mathcal{S}$ is a Sperner system over V ($n := |V|$) if for any two different edges $E, E' \in \mathcal{S}$, $E \not\subset E'$.

The main question in this topic is: What is the largest possible size of a Sperner system over V?

Sperner's Theorem

Example

For any $0 \leq k \leq n = |V|$, $S = \binom{V}{k} = \{R \subset V : |R| = k\}$ is a Sperner system. It has $\binom{n}{k}$ elements.

For $k = \lfloor n/2 \rfloor$, we get the largest possible system of these examples.

Sperner's Theorem

The maximum size of Sperner systems over V is $\binom{n}{\lfloor n/2 \rfloor}$.

I. Proof

Definition

Let $\mathcal{H}$ be a family of subsets of V with n elements. The f-vector of $\mathcal{H}$ is the $(f_0, f_1, \ldots, f_n)$ vector, where f_i component indicates how many i-element sets are in $\mathcal{H}$.

(LYM Inequality)

Let S be a Sperner system over V. Then, for the f-vector f,

$$\sum_{i=0}^{|V|} \frac{f_i}{\binom{n}{i}} \leq 1.$$

The lemma is named after Lubell, Yamamoto, and Meshalkin, who independently proved it. It is also often associated with Béla Bollobás, who proved a related statement using a similar method.

Proof of LYM Inequality

Let π be an arbitrary bijection $V \to [n]$, and $E \in S$ be arbitrary. Count the pairs (π, E) where $\pi(E)$ is an initial segment of $[n]$.

If we count, for each $E \in S$, all valid π orderings, then we get $\sum_{E \in S} |E|! \cdot (n - |E|)!$ such pairs.

Now, let π be an arbitrary ordering. Since the inclusion relation is a total order on the initial segments of $[n]$, if $\pi(E_1), \pi(E_2)$ are initial segments of $[n]$, then either $E_1 \subset E_2$ or $E_2 \subset E_1$. So, for any ordering π, there can be at most one E such that $\pi(E)$ is an initial segment of $[n]$.

Completion of LYM Inequality Proof

Compare the two counts.

$$\sum_{E \in S} |E|! \cdot (n - |E|)! \leq n!$$

Dividing both sides by $n!$, we get the statement of the lemma.

$$1 \geq \sum_{i=0}^{|V|} \frac{f_i}{\binom{n}{i}} \geq \sum_{i=0}^{|V|} \frac{f_i}{\binom{n}{[n/2]}} = \frac{|\mathcal{S}|}{\binom{n}{[n/2]}}.$$

Partially Ordered Sets

Definition

Let $(P, \leq)$ be a partially ordered set. A subset $L \subset P$ is a chain if any two elements in L are comparable.

Definition

Let $(P, \leq)$ be a partially ordered set. A subset $A \subset P$ is an antichain if the elements of A are pairwise incomparable.

The antichains over $(\mathcal{P}(V), \subset)$ precisely correspond to Sperner systems over V.

II. Proof of Sperner's Theorem

Observation

For any chain L and antichain A in P, $|L \cap A| \leq 1$.

Claim

If we have chains $L_1, L_2, \ldots, L_k$ covering P, then any antichain in P has at most k elements.

Corollary

$$\max_{A \text{ antichain}} (|A|) \leq \min_{L_1, L_2, \ldots, L_k \text{ is a chain cover}} k$$

Lemma

Lemma

$(\mathcal{P}(V), \subset)$ has a chain cover with $\binom{n}{\lfloor n/2 \rfloor}$ chains.

Definition

$\mathcal{L} \subset \mathcal{P}(V)$, $\mathcal{L} : L_1 \subset L_2 \subset ... \subset L_t$ is a symmetric chain if there exists an i such that $|L_1| = i, |L_2| = i + 1, \ldots, |L_t| = |V| - i$.

Lemma

$(\mathcal{P}(V), \subset)$ has a cover with disjoint symmetric chains.

Due to symmetry, each chain must contain a set of size $\lfloor n/2 \rfloor$. Therefore, the number of chains used must be $\binom{n}{\lfloor n/2 \rfloor}$.

Proof of the Lemma

For $|V| = 1, 2, 3$, the statement is trivially true. For the induction step, let $|V| > 1$. Then, consider $V = V_0 \dot\cup \{u\}$, where $\mathcal{P}(V_0)$ already has a covering.

$\mathcal{P}(V) = \mathcal{P}(V_0) \dot\cup \{R \subset V : u \in R\}$. Let $\mathcal{P}(V_0) = \mathcal{L}_1 \dot\cup \mathcal{L}_2 \dot\cup \ldots \dot\cup \mathcal{L}_k$ be the covering from the induction hypothesis. Now, we construct chains from the chains in $\mathcal{L}_t$ as follows:

$$\mathcal{L}'_t : L_1 \cup \{u\} \subset L_2 \cup \{u\} \subset \ldots \subset L_{j-1} \cup \{u\},$$

and

$$\mathcal{L}''_t : L_1 \subset L_2 \subset \ldots \subset L_j \subset L_j \cup \{u\}.$$

It can be observed that these chains are symmetric to the base set $V_0 \cup \{u\}$ and pairwise disjoint. Thus, they prove the lemma, and consequently, Sperner's theorem.

Note

It might seem as if our inductive/recurrent construction always doubled the number of our chains. However, the number of chains does not increase as a power of two; rather, it remains $\binom{n}{\lfloor n/2 \rfloor}$.

The resolution of this apparent contradiction lies in the fact that $\mathcal{L}'_i$ can be empty.

Break

Theorem

At the heart of the second proof lies the mathematical theory of partially ordered sets. Our goal is to understand that the proof scheme based on this observation is „complete".

Theorem

Let $(P, \leq)$ be a partially ordered set.

(i)

$$\max_{L \text{ chain}} (|L|) = \min_{A_1, A_2, \dots, A_k \text{ antichain cover}} k$$

(ii) (Dilworth's Theorem)

$$\max_{A \text{ antichain}} (|A|) = \min_{L_1, L_2, \dots, L_k \text{ chain cover}} k$$

In both cases, we only need to prove that the optimum of the maximization problem is larger than the optimum of the

Proof (i)

Let $M = \max_{L \text{ chain}}(|L|)$

Associate each $x \in P$ with the largest size among chains containing x as the maximal element. (This is well-defined since $\{x\}$ is always a chain containing x as the maximal element.)

The range of the assignment is $\{1, 2, \dots, M\}$. Let A_i $(i = 1, 2, \dots, M)$ be the set of elements in P to which we assign the value i.

This way, we cover P with M sets. If we can show that each A_i is an antichain, we are done.

This follows indirectly, if $x < y$ and $x, y \in A_i$, then giving y to the chain demonstrating $x \in A_i$ produces a chain of size $i + 1$, contradicting the assumption $y \in A_i$.

Proof (i)

Let $M = \max_{L \text{ chain}}(|L|)$

For each $x \in P$, associate the largest size among chains containing x as the maximal element. (This is well-defined since $\{x\}$ is always a chain containing x as the maximal element.)

The range of the assignment is $\{1, 2, \ldots, M\}$. Let A_i $(i = 1, 2, \ldots, M)$ be the set of elements in P to which we assign the value i.

This way, we cover P with M sets. If we can show that each A_i is an antichain, we are done.

This follows indirectly, if $x < y$ and $x, y \in A_i$, then giving y to the chain demonstrating $x \in A_i$ produces a chain of size $i + 1$, contradicting the assumption $y \in A_i$.

Proof (ii): The Plan

Let $M = \max\{|A| : A \text{ antichain}\}$ and
$m = \min\{k : L_1, L_2, \ldots, L_k \text{ covering chains}\}$.

Define a bipartite graph B based on $(P, \leq)$.

- V The two color classes are $F = \{p^+ : p \in P\}$ and $A = \{p^- : p \in P\}$.
- E p^+ and q^- are connected if and only if $p > q$.

Our goal is to show that $\nu(B) = |P| - m$ and $\tau(B) = |P| - M$.

Then, the statement follows directly from Kőnig's theorem.

$\nu(B) = |P| - m$

Associate each chain in the covering with the edges $\ell_1^+\ell_2^-, \ell_2^+\ell_3^-, \ldots, \ell_{s-1}^+\ell_s^-$.

Doing this for all chains gives $|P| - m$ edges, forming a matching.

The construction is reversible: given an M-matching, construct the edges pq on the set P from the p^+q^- edges.

The resulting paths form a system.

The point sets of the components are chains that cover P. This implies $\nu(B) = |P| - m$.

$\tau(B) = |P| - M$

Let A be a maximal antichain.

Divide the elements of $P - A$ into two parts: L^- consists of the elements that are smaller than some element in A, L^+ consists of the elements that are larger than some element in A.

Clearly, L^+ and L^- are disjoint and together they cover $P - A$.

Let $R = \{p^+ : p \in L^+\} \dot\cup \{p^- : p \in L^-\}$ be a covering set of size $|P| - M$.

$\tau(B) = |P| - M$ (Continuation)

The reasoning is reversible: For every $R \subset V(B)$, it determines a partition of P into four parts

$$P = P^+(R) \dot\cup P^-(R) \dot\cup P^\pm(R) \dot\cup P^0(R)$$

according to how $\{p^+, p^-\}$ relates to R.

Then

$$R = \{p^- : p \in P^+(R)\} \dot\cup \{p^+ : p \in P^-(R)\} \dot\cup \{p^-, p^+ : p \in P^\pm(R)\}.$$

If R is a covering set, then $P^0(R)$ must be an antichain.

For $|R|$ to be minimal, the optimal choice for $P^\pm(R)$ is $\emptyset$.

Partially Ordered Sets and Graphs

For a partially ordered set $(P, \leq)$, we associate a comparison graph G_P: This simple graph has vertex set P, and two vertices are connected if and only if they are comparable.

Observation

$$\max_{L \text{ chain}} (|L|) = \omega(G_P),$$

$$\min_{A_1, A_2, \ldots, A_k \text{ antichain cover}} k = \chi(G_P),$$

$$\max_{A \text{ antichain}} (|A|) = \alpha(G_P) = \omega(\overline{G_P}),$$

$$\min_{L_1, L_2, \ldots, L_k \text{ chain cover}} k = \chi(\overline{G_P}).$$

Graph Theory

The connections established in the previous theorem lead to the following graph theoretical concept:

Definition

A graph G is perfect if for every induced subgraph F obtained by deleting some vertices (i.e., a vertex set), we have

$$\omega(F) = \chi(F).$$

Theorem

Let G_P be a comparison graph over a partially ordered set $(P, \leq)$. Then

(i) G_P is perfect,

(ii) $\overline{G_P}$ is perfect.

Graph Theory

The connections established in the previous theorem lead to the following graph theoretical concept:

Definition

A graph G is perfect if for every induced subgraph F obtained by deleting some vertices (i.e., a vertex set), we have

$$\omega(F) = \chi(F).$$

Theorem

Let G_P be a comparison graph over a partially ordered set $(P, \leq)$. Then

(i) G_P is perfect,

(ii) $\overline{G_P}$ is perfect.

Break

Intersecting Set Systems

Definition

A set system is called intersecting if any two of its edges intersect.

In other words, a set system is intersecting if it doesn't contain disjoint pairs of edges.

The basic extremal question is about the maximum number of edges in an intersecting set system over an n-element vertex set.

Intersecting Set Systems: Examples

Example

Let $x \in V$. $\mathcal{H}$ consists of all sets containing x. $\mathcal{H}$ is obviously intersecting, and $|\mathcal{H}| = 2^{|V|-1} = 2^{n-1}$.

Example

Let V be an n-element set, where n is odd, $n = 2k + 1$. $\mathcal{H}$ consists of all sets with at least $k + 1$ elements. $\mathcal{H}$ is intersecting, and $|\mathcal{H}| = 2^{|V|-1} = 2^{n-1}$.

Example

If the base set has an even number of elements, and we include only one of the complementary pairs of exactly $|V|/2$-sized sets in $\mathcal{H}$, along with more than $|V|/2$-sized sets.

A High School Problem

Observation

An intersecting set system over an n-element V can have at most 2^{n-1} edges. The provided examples are extremal.

Indeed, we can divide the 2^n subsets of V into 2^{n-1} complementary pairs, and each pair can contribute to at most one intersecting set system.

The Real Problem

A much more challenging question arises when we work with k uniform set systems.

For $k > |V|/2$, there is no issue: all k-sets form an intersecting system.

For $k \leq |V|/2$, a fundamental theorem answers our question.

Theorem (Erdős—Ko—Rado Theorem)

Let $k \leq n/2$. Let $\mathcal{H}$ be a k-uniform intersecting set system over an n-element V. Then

$$|\mathcal{H}| \leq \binom{n-1}{k-1}$$

Our estimate is the best possible, achieved by considering all k-sets containing a fixed element.

Circularly Ordered Sets, Arcs

Let K be the vertex set of a cycle with n points. There exists a clockwise order among these points, resulting in the sequence $v_0, v_1, \ldots, v_{n-1}$, where indices are taken modulo n.

For $I \subset K = \{v_0, v_1, \ldots, v_{n-1}\}$, we say that I is the arc $[a, z]$ if I contains a and the following points up to z inclusively, that is, there exist $i \in \{0, 1, \ldots, n-1\}$ and $\ell \in \{1, \ldots, n\}$ such that $I = \{v_i, v_{i+1}, \ldots, v_{i+\ell-1}\}$.

$\ell = |I|$ is the length of the arc I.

How many k-length arcs can be selected to form an intersecting system?

Solution

Claim

k arcs can be chosen (for example, $[a_1, a_k], [a_2, a_{k+1}], \ldots, [a_k, a_{2k-1}]$), no more.

Indeed: If $I = [a_i, a_{i+k-1}] = (a_i, \ldots, a_{i+k-1})$ is an arc in our system, then every other arc intersects I.

There are two types of arcs intersecting I: ending in I, and starting in I.

There are $2(k-1)$ arcs, forming $k-1$ complementary pairs: typical pairs consist of arcs ending in a_j and starting in a_{j+1}. (Here we use $2k \leq n$.)

Thus, there cannot be more than $1 + (k-1)$ arcs.

Proof by Gyula Katona

Let $\mathcal{H}$ be a k-uniform intersecting set system.

Let π be a bijection between V (a subset of $\mathcal{H}$'s base set) and the cyclically ordered set K introduced earlier.

Count the pairs (π, E), where $E \in \mathcal{H}$ and $\pi(E)$ is an arc. Perform the counting in two different ways.

First, given E, consider how many ways π can be chosen such that the corresponding pair is counted. It is easy to see that $\pi(E)$ is a k-length arc, and there are n possibilities. Once fixed, there are $k! \cdot (n-k)!$ good bijections.

For all pairs,

$$\sum_{E \in \mathcal{H}} n \cdot k!(n-k)! = |\mathcal{H}| n \cdot k!(n-k)!$$

is obtained.

End of the Proof

Second, given π, consider how many edges lead to counting the pair.

Here, the earlier simplification is useful. At most k can be obtained, so

$$kn!$$

at most for all pairs.

Comparing the two results yields the theorem.

Break

Sunflowers

Definition

$H_1, \ldots, H_s$ form a *s-petal sunflower* (or Δ-system) if for every $i \neq j$ $(i, j \in \{1, \ldots, s\})$, $H_i \cap H_j = \bigcap_{k=1}^{s} H_k$. The set $T = \bigcap_{k=1}^{s} H_k$ is called the *plate* of the sunflower.

For example, a collection of s pairwise disjoint set systems forms an s-petal sunflower.

The fundamental question in the topic of sunflowers is: given a k-uniform set system with no s-petal sunflower, what is the maximum number of edges it can have?

Erdős—Rado Theorem

Let $\mathcal{H}$ be a k-uniform set system that does not contain an s-petal sunflower. Then

$$|\mathcal{H}| \leq (s-1)^k k!.$$

Proof: Start of Induction

We prove the theorem by complete induction on k, specifically in the form: If $\mathcal{H}$ is a k-uniform set system and $|\mathcal{H}| > (s-1)^k k!$, then $\mathcal{H}$ contains an s-petal sunflower.

The base case $k = 1$ is trivial, considering that a 1-uniform set system's elements are disjoint singletons and form an s-petal sunflower for any s.

Assume that we have established the statement for $k - 1$. To prove it for k, we will need the following lemma.

Proof: The Lemma

Lemma

Let $\mathcal{H}$ be a k-uniform set system, and $t \in \{2, 3, \ldots\}$. Then one of the following is true:

(i) there exist t pairwise disjoint edges,

(ii) there exists a vertex $v \in V$ such that v is incident with at least $\frac{|\mathcal{H}|}{(t-1)k}$ edges.

Proof: Deriving the Theorem from the Lemma

Apply the lemma for $t = s$.

If *(i)* holds, then there are s pairwise disjoint edges, forming an s-petal sunflower.

If *(ii)* holds, let $\widetilde{\mathcal{H}} = \{E \backslash \{v\} : v \in E \in \mathcal{H}\}$. (In other words, remove v from the edges containing it.)

Obviously, $\widetilde{\mathcal{H}}$ is $(k-1)$-uniform, and

$$\widetilde{\mathcal{H}} \geq \frac{|\mathcal{H}|}{k(t-1)} > \frac{(s-1)^k k!}{k(s-1)} = (s-1)^{k-1}(k-1)!$$

By the induction hypothesis, $\widetilde{\mathcal{H}}$ contains an s-petal sunflower, denoted as $S_1, \ldots, S_s$. Then $S_1 \cup \{v\}, \ldots, S_s \cup \{v\}$ form an s-petal sunflower in $\mathcal{H}$.

Remarks

The theorem gives an upper bound of $(s-1)^k k!$ on the number of edges in a set system that does not contain an s-petal sunflower. This bound grows faster than exponential.

Construction: No three-petal sunflower

Let $V = \{a_1, a_2, \ldots, a_k\} \dot\cup \{b_1, b_2, \ldots, b_k\}$. $\mathcal{H}$ contains edges such that each $\{a_i, b_i\}$ $(i = 1, 2, \ldots, k)$ pair intersects it in exactly one element. It is easy to see that $\mathcal{H}$ is a 2^k-edge k-uniform set system.

The imaginary sunflower's plate is included in every edge of it, The plate consists of either zero or one element from each $\{a_i, b_i\}$ pair.

On the other hand, it cannot be k-sized. There must exist an i such that the plate is disjoint from $\{a_i, b_i\}$.

How do the three petals intersect with the set $\{a_i, b_i\}$? They must intersect disjointly and each must have exactly one element in common. Thus, the imaginary sunflower cannot exist.

A recent breakthrough

Rao, Alweiss—Lovett—Wu—Zhang 2019

Let $\mathcal{H}$ be a k-uniform set system that does not contain an s-petal sunflower. Then
$$|\mathcal{H}| \leq \mathcal{O}((s\log(sk))^k).$$

Break

λ-Intersecting Set Systems

Definition

A set system $\mathcal{H}$ over the base set V is λ-intersecting if, for any distinct $A, B \in \mathcal{H}$, $|A \cap B| = \lambda$.

Naturally, the fundamental question is: what is the maximum number of edges in a λ-intersecting set system?

Case of $\lambda = 0$

Example

Let $\lambda = 0$ and $\mathcal{H} = \{\emptyset, \{v_1\}, \ldots, \{v_n\}\}$.

It is easy to see that for $|V| = n$ and $\lambda = 0$, this is the largest set system that is 0-intersecting.

From now on, we assume $\lambda \geq 1$.

Examples

Example

$\lambda = 1$,
$V = \{\{v_1, \ldots, v_{n-1}\}, \{v_1, v_n\}, \{v_2, v_n\}, \{v_3, v_n\}, \ldots \{v_{n-1}, v_n\}\}$

Example: $\lambda = 1$ and the Fano Plane

Seven points $V = \{P_1, P_2, P_3, P_4, P_5, P_7\}$, and set system
$\mathcal{H} = \{\{P_1, P_2, P_3\}, \{P_3, P_4, P_5\}, \{P_1, P_5, P_6\}, \{P_1, P_4, P_7\},$
$\{P_3, P_6, P_7\}, \{P_2, P_5, P_7\}, \{P_2, P_4, P_6\}\}$.

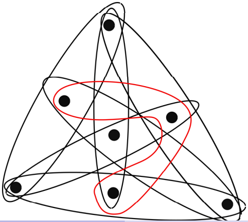

The Fundamental Theorem

Theorem

Let $\lambda \geq 1$ and $\mathcal{F}$ be a λ-intersecting set system over a base set V. Then

$$|\mathcal{F}| \leq |V|.$$

Start of the Proof

If there is an edge in $\mathcal{F}$ with cardinality less than λ, then no other edge in $\mathcal{F}$ is possible, and the statement is trivial.

If there is an edge F with exactly λ elements, then every other edge must contain F. For the other edges E, the sets $E \setminus F$ are pairwise disjoint, non-empty subsets of $V \setminus F$. Therefore, there can be at most $|V| - |F|$ such sets. This implies that the total number of edges is at most $1 + (|V| - \lambda) \leq |V|$.

From now on, we assume that every edge has more than λ (at least $\lambda + 1$) elements.

Linear Algebra

For an edge $F \in \mathcal{F}$, let χ_F be the characteristic vector of the set $F \subset V$ ($\chi_F \in \mathbb{R}^V \equiv \mathbb{R}^n$). We will show that the χ_F vectors ($F \in \mathcal{F}$) are linearly independent. This implies the theorem.

Let $M_{\mathcal{F}}$ be the matrix whose rows are the χ_F vectors ($F \in \mathcal{F}$). Its size is $|\mathcal{F}| \times |V|$.

What does the $M_{\mathcal{F}} \cdot M_{\mathcal{F}}^{\mathsf{T}}$ matrix look like?

The entries are the scalar products $\chi_F \chi_{F'} = |F \cap F'|$. Since $\mathcal{F}$ is a λ-intersecting set system, there are λ's off the main diagonal. On the main diagonal, the sizes of our edges are present.

$M_{\mathcal{F}} \cdot M_{\mathcal{F}}^{\mathsf{T}}$

Let $\mathcal{F} = \{A_1, \dots, A_m\}$ $(m = |\mathcal{F}|)$.

$$\begin{pmatrix} |A_1| & \lambda & \lambda & \dots & \lambda & \lambda \\ \lambda & |A_2| & \lambda & \dots & \lambda & \lambda \\ \lambda & \lambda & |A_3| & \dots & \lambda & \lambda \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \lambda & \lambda & \lambda & \dots & |A_{m-1}| & \lambda \\ \lambda & \lambda & \lambda & \dots & \lambda & |A_m| \end{pmatrix}$$

We will show that the rows of this matrix are linearly dependent. This implies the theorem.

$M_{\mathcal{F}} \cdot M_{\mathcal{F}}^{\mathsf{T}}$

Indeed, consider a linear combination of the rows (with coefficients α_i). The jth component of the linear combination:

$$\alpha_j|A_j| + \sum_{i:i\neq j} \alpha_i\lambda = \alpha_j|A_j| - \alpha_j\lambda + \sum_{i=1}^{m} \alpha_i\lambda = \alpha_j(|A_j| - \lambda) + \sum_{i=1}^{m} \alpha_i\lambda$$
$$= \alpha_j(|A_j| - \lambda) + \Lambda.$$

If we set up the linear combination to get the 0 vector, then

$$\alpha_j = \frac{-\Lambda}{|A_j| - \lambda}.$$

Hence the signs of the α_j's are the same.

For the combination to be the 0 vector, we must have $\Lambda = 0$. This implies that every α_i is 0. Hence, the rows are linearly independent.

Linear Algebraic Method

The Erdős-Ko-Rado Theorem and the Fisher Inequality are both about set systems with *intersecting conditions*. This area has

flourished, producing many important results.

These theorems have had a significant impact not only in combinatorics but also beyond.

In combinatorics, this linear algebraic method has become particularly significant (for example, in proving the Fisher Inequality). It has become an important proof technique.

Break

Trace

Definition

Let $\mathcal{H}$ be a set system over V, and A be a subset of V. Then define $Tr_A\mathcal{H} = \{E \cap A : E \in \mathcal{H}\}$ as the *trace* of $\mathcal{H}$ on A.

It is clear that $Tr_A\mathcal{H} \subseteq \mathcal{P}(A)$. In the case where $Tr_A\mathcal{H} = \mathcal{P}(A)$, we say that A is *saturated*.

Vapnik–Chervonenkis Dimension

Vapnik–Chervonenkis Dimension of $\mathcal{H}$

$$\dim_{VC} \mathcal{H} = \max\{|A| : \ A \subset V \text{ is saturated}\}.$$

(Vapnik–Chervonenkis)

Let $\mathcal{H}$ be a set system over $[n] = \{1, 2, \ldots, n\}$, and t be a positive integer such that the inequality $|\mathcal{H}| > 1 + \binom{n}{1} + \binom{n}{2} + \ldots + \binom{n}{t-1}$ holds. Then $\dim_{VC} \mathcal{H} \geq t$. In other words, there exists a saturated set A of size t in $[n]$ for $\mathcal{H}$.

Tightness of the Bound

Consider the set system defined by

$$\mathcal{H} = \{R \subseteq [n] : |R| < t\}.$$

Clearly,

$$|\mathcal{H}| = 1 + \binom{n}{1} + \ldots + \binom{n}{t-1},$$

However, in this system, there is no saturated set A of size t. To see this, note that for a set A to be saturated, $\mathcal{H}$ must contain an edge containing A.

Proof 1

We proceed by induction on n.

For $n = 1$, the statement of the theorem is trivially true.

Using the identity $\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}$, the condition of the theorem implies

$$|\mathcal{H}| > \left[\binom{n-1}{0} + \ldots + \binom{n-1}{t-2}\right] + \left[1 + \binom{n-1}{1} + \ldots + \binom{n-1}{t-1}\right].$$

Denote L_1 and L_2 as the two bracketed expressions.

Proof 1 (contd)

Introduce the following notations:
$\mathcal{H}_1 = \{E \in \mathcal{H} : n \notin E, E \cup \{n\} \in \mathcal{H}\}$, $\mathcal{H}_2 = \mathcal{H} - \mathcal{H}_1$, and let $\widetilde{\mathcal{H}_2} = \{E \backslash \{n\} : E \in \mathcal{H}_2\}$.

Clearly, $\mathcal{H}_1, \widetilde{\mathcal{H}_2}$ are set systems over $[n-1]$.

We know that

$$|\mathcal{H}_1| + |\mathcal{H}_2| = |\mathcal{H}| > L_1 + L_2$$

Either *(i)* $|\mathcal{H}_1| > L_1$, or *(ii)* $|\mathcal{H}_2| = |\widetilde{\mathcal{H}_2}| > L_2$ holds.

Proof 1 (contd)

If *(i)* is true, then by the induction hypothesis, there exists a $t-1$-sized saturated set A with respect to $\mathcal{H}_1$. It is easy to see (since for $E \in \mathcal{H}_1$, both E and $E \cup \{n\}$ are edges in $\mathcal{H}$) that $A \cup \{n\}$ is saturated with respect to $\mathcal{H}$ (and has size t).

If *(ii)* is true, then by the induction hypothesis, there exists a t-sized saturated set A with respect to $\widetilde{\mathcal{H}_2}$. This implies that for every $R \subseteq A$, there exists $E \in \widetilde{\mathcal{H}_2}$ such that $E \cap A = R$. However, for every $E \in \widetilde{\mathcal{H}_2}$, there uniquely exists $E_0 \in \mathcal{H}_2$, either E or $E \cup \{n\}$. In both cases, $E_0 \cap A = R$, which means A is saturated with respect to $\mathcal{H}$. ■

Proof 2

A set system is called *down-closed* if whenever $E \in \mathcal{H}$ and $F \subseteq E$, then $F \in \mathcal{H}$.

If $\mathcal{H}$ is down-closed, the statement of the theorem straightforwardly follows: the conditions ensure the existence of at least t-sized edges in $\mathcal{H}$, and due to down-closedness, every edge is saturated.

Define the following S_i transformation: for $i \in V$, if $E \in \mathcal{H}$, then $S_i E = E \setminus \{i\}$ if $E \setminus \{i\} \notin \mathcal{H}$, and $S_i E = E$ otherwise. Let $S_i \mathcal{H} = \{S_i E : E \in \mathcal{H}\}$.

Proof 2 (contd)

Note that $|\mathcal{H}| = |S_i\mathcal{H}|$ follows directly from the definition.

It is not hard to see that if $\mathcal{H}$ is not down-closed, then there exists an i such that $S_i\mathcal{H} \neq \mathcal{H}$. (If it is not down-closed, there exist E and F such that $F \subset E$ and $E \in \mathcal{H}$ but $F \notin \mathcal{H}$. Then, any $i \in E \backslash F$ will do.)

The third observation is stated as a lemma.

Lemma

$|Tr_A\mathcal{H}| \geq |Tr_A S_i\mathcal{H}|$ always holds.

From Lemma to Theorem

For any $\mathcal{H} = \mathcal{H}_1$, there exist $i_1, i_2, \ldots$ such that $\mathcal{H}_k \neq \mathcal{H}_{k+1} = S_{i_k}\mathcal{H}_k$, iterating the S transformation until the set system stops changing.

Obviously, this chain must terminate in finitely many steps because the sum of the sizes of edges decreases in each step.

Let the last set system be $\mathcal{H}_s$. By what we have shown so far, $\mathcal{H}_s$ is down-closed, and the number of edges satisfies the condition of the theorem.

Thus, there exists a t-sized edge A in $\mathcal{H}_s$. Then A is saturated with respect to $\mathcal{H}_s$, and its trace on A has $2^{|A|}$ elements. By the lemma, the trace of A with respect to $\mathcal{H}_1$ also has at least $2^{|A|}$ elements, which means A is saturated. The proof is complete.

Proof of the Lemma

If $i \notin A$, then $Tr_A\mathcal{H} = Tr_A S_i \mathcal{H}$ is trivially satisfied.

If $i \in A$, consider the pairs $(R, R \cup \{i\})$ for all subsets R of A where $i \notin R$. If an edge E contributes to one pair, then its transformation $S_i E$ contributes to the same pair.

It suffices to show that every pair contributes at least as much to $Tr_A\mathcal{H}$ as to $Tr_A S_i \mathcal{H}$.

The only potential issue is when R and $R \cup \{i\}$ are both in $Tr_A S_i \mathcal{H}$, but only one of them is in $Tr_A\mathcal{H}$. Clearly, the missing one must be $R \cup \{i\}$. However, if $R \cup \{i\}$ is not in $Tr_A\mathcal{H}$, then every edge E for which $E \cap A = R \cup \{i\}$ must satisfy $S_i E = E \setminus \{i\}$. This contradicts $R \cup \{i\} \in Tr_A S_i \mathcal{H}$, and thus, the lemma is proven.

This is the End!

Thank you for your attemtion!