The Lovász Local Lemma is Not About Probability

Dimitris Achlioptas
Kostas Zampetakis

Prepared by Justyna Jaworska

Introduction

Given a collection of independent events (F_i) with probabilities $0 < p_i < 1$, the probability that none of them occurs is strictly positive.

$$\mathbb{P}(\overline{F_1} \cap \ldots \cap \overline{F_n}) = (1 - p_1) \cdot \ldots \cdot (1 - p_n) > 0$$

Introduction

Given a collection of independent events (F_i) with probabilities $0 < p_i < 1$, the probability that none of them occurs is strictly positive.

But what if some of them *are* dependent?

Introduction

Given a collection of independent events (F_i) with probabilities $0 < p_i < 1$, the probability that none of them occurs is strictly positive.

But what if some of them *are* dependent?

In general, the union bound is the best we can do:

$$\mathbb{P}(\overline{F_1} \cap \ldots \cap \overline{F_n}) \geq 1 - p_1 - p_2 - \ldots - p_n$$

Introduction

Given a collection of independent events (F_i) with probabilities $0 < p_i < 1$, the probability that none of them occurs is strictly positive.

But what if some of them *are* dependent?

$$\mathbb{P}(\overline{F_1} \cap \ldots \cap \overline{F_n}) \geq 1 - p_1 - p_2 - \ldots - p_n$$

We must somehow limit the dependencies between events.

(Lopsi)dependency graph

A graph G on n vertices is a *dependency graph* for events $F_1, F_2, \ldots F_n$, if for every vertex i and any set $\{j_1, \ldots j_k\} \subset [n] \setminus \{\{i\} \cup \Gamma_i\}$ of vertices **non**-adjacent to i, the following holds:

$$\mathbb{P}(F_i | \overline{F_{j1}} \cap \ldots \overline{F_{jn}}) = \mathbb{P}(F_i)$$

In other words, F_i is independent from avoiding its non-neighbouring events.

(Lopsi)dependency graph

A graph G on n vertices is a *dependency graph* for events $F_1, F_2, \dots F_n$, if for every vertex i and any set $\{j_1, \dots j_k\} \subset [n] \setminus \{\{i\} \cup \Gamma_i\}$ of vertices **non**-adjacent to i, the following holds:

$$\mathbb{P}(F_i | \overline{F_{j1}} \cap \dots \overline{F_{jn}}) = \mathbb{P}(F_i)$$

In other words, F_i is independent from avoiding its non-neighbouring events.
In practice, we can assume a weaker property:

$$\mathbb{P}(F_i | \overline{F_{j1}} \cap \dots \overline{F_{jn}}) \leq \mathbb{P}(F_i)$$

A graph G with this property is sometimes called *lopsidependency graph*

(Lopsi)dependency graph

Imagine we choose a number n from the set $\{1, 2, \ldots, 60\}$ uniformly at random. Let X_i be the event where $i|n$. This is an example dependency graph for G:

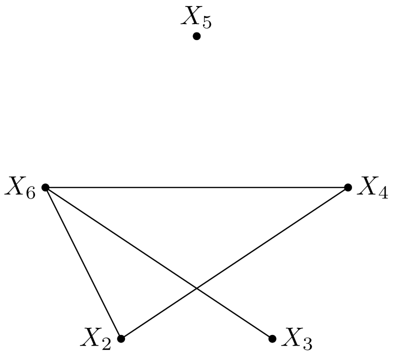

Lovász Local Lemma

Theorem (Lovász-Erdős, 1973). If a dependency graph G for a set of events $F_1, F_2, \ldots F_n$ has maximum degree d and for every i:

$$p_i = \mathbb{P}(F_i) \leq \frac{1}{4d}$$

Then $P(\overline{F_1} \cap \ldots \cap \overline{F_n}) > 0$.

Lovász Local Lemma

Theorem (Lovász-Erdős, 1973). If a dependency graph G for a set of events $F_1, F_2, \ldots F_n$ has maximum degree d and for every i:

$$p_i = \mathbb{P}(F_i) \leq \frac{1}{4d}$$

Then $P(\overline{F_1} \cap \ldots \cap \overline{F_n}) > 0$.

- avoid$(\mathbf{p}, G) :=$ the smallest possible value of the *avoidance probability* $\mathbb{P}(\overline{F_1} \cap \ldots \cap \overline{F_n})$.

Asymmetric LLL

Theorem (Lovász-Erdős, 1975) If there exist $r_1, \ldots, r_n \in [0, 1)$ such that for every i:

$$p_i \leq r_i \prod_{j \in \Gamma_i} (1 - r_j)$$

Then $\mathrm{avoid}(\mathbf{p}, G) \geq \prod_{i \in [n]} (1 - r_i) > 0$.

Part 1.

Connection to walks on graphs

A couple of definitions

- Let $\mathcal{W}$ be a (not necessarily finite) set of walks on a graph G.
- $t(w)$ denotes the last vertex of a walk $w \in \mathcal{W}$.
- For any $w \in (W)$, let $\mathrm{Ext}(w)$ denote the set of paths in $\mathcal{W}$ which extend w by exactly one vertex.

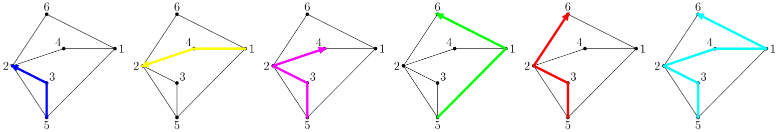

$$ext(p_1) = \{p_3, p_5\}$$

A couple of definitions

- Let $\mathcal{W}$ be a (not necessarily finite) set of walks on a graph G.
- $t(w)$ denotes the last vertex of a walk $w \in \mathcal{W}$.
- For any $w \in (W)$, let $\mathrm{Ext}(w)$ denote the set of paths in $\mathcal{W}$ which extend w by exactly one vertex.

We will say that $\mathbf{p} \in [0,1)^n$ is valid for $\mathcal{W}$ if there exists $L : \mathcal{W} \to [0,1)$ satysfying:

$$p_{t(w)} \leq L(w) \prod_{y \in \mathrm{Ext}(w)} (1 - L(y))$$

Self-bounding walks

Given is graph G on $[n]$. A walk on G is called *self-bounding* if in each step:

1. it procedes from the current vertex i to a non-forbidden neighbour j
2. adds to the set of forbidden vertices the vertex i and all neighbours of i greater than j.

We will denote the set of all self-bounding walks on G as $\mathcal{B}(G)$.

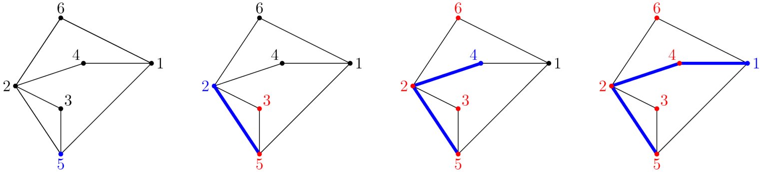

Self-bounding walks

Reminder: We will say that $\mathbf{p} \in [0,1)^n$ is valid for $\mathcal{W}$ if there exists $L : \mathcal{W} \to [0,1)$ satysfying:

$$p_{t(w)} \leq L(w) \prod_{y \in \mathrm{Ext}(w)} (1 - L(y))$$

Theorem 1. $\mathrm{avoid}(\mathbf{p}, G) > 0 \iff \mathbf{p}$ is valid for $\mathcal{B}(G)$.

Self-bounding walks

Reminder: We will say that $\mathbf{p} \in [0,1)^n$ is valid for $\mathcal{W}$ if there exists $L : \mathcal{W} \to [0,1)$ satysfying:

$$p_{t(w)} \leq L(w) \prod_{y \in \mathrm{Ext}(w)} (1 - L(y))$$

Theorem 1. $\mathrm{avoid}(\mathbf{p}, G) > 0 \iff \mathbf{p}$ is valid for $\mathcal{B}(G)$.

This isn't an easy condition to work with. But we can use this **Observation.** For a family of walks $\mathcal{W} \supseteq \mathcal{B}(G)$:

$$\mathbf{p} \text{ is valid for } \mathcal{W} \implies \mathbf{p} \text{ is valid for } \mathcal{B}(G)$$

Equivalence classes

Let $\sim$ be the relation where:

$$w \sim w' \text{ if } t(w) = t(w') \text{ and } \{z : wz \in \mathcal{W}\} = \{z : w'z \in \mathcal{W}\}$$

- We will denote equivalency class of w by $\tilde{w}$
- We will denote the set of equivalence classes of $\mathcal{W}$ by $\mathcal{C}(\mathcal{W})$
- We can define $t(\tilde{w})$ and $\text{Ext}(\tilde{w})$

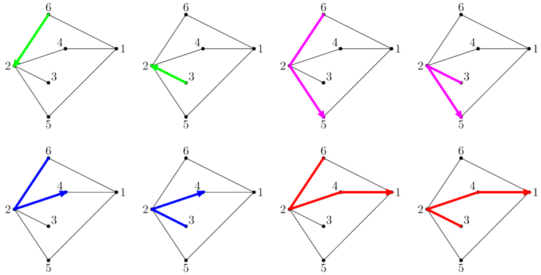

$t(c_1) = 2$

$Ext(c_1) = \{c_2, c_3\}$

Equivalence classes

Reminder: We will say that $\mathbf{p} \in [0,1)^n$ is valid for $\mathcal{W}$ if there exists $L : \mathcal{W} \to [0,1)$ satysfying:

$$p_{t(w)} \leq L(w) \prod_{y \in \mathrm{Ext}(w)} (1 - L(y))$$

Theorem. The following are equivalent:

1. $\mathbf{p}$ is valid for $\mathcal{W}$
2. For every $\tilde{w} \in \mathcal{C}(\mathcal{W})$ there exists $r_{\tilde{w}} \in [0,1)$ such that

$$p_{t(\tilde{w})} \leq r_{\tilde{w}} \prod_{y \in \mathrm{Ext}(\tilde{w})} (1 - r_{\tilde{y}})$$

($t(\tilde{w})$ and $\mathrm{Ext}(\tilde{w})$ are well defined functions)

Equivalence classes

Reminder: We will say that $\mathbf{p} \in [0,1)^n$ is valid for $\mathcal{W}$ if there exists $L : \mathcal{W} \to [0,1)$ satysfying:

$$p_{t(w)} \leq L(w) \prod_{y \in \mathrm{Ext}(w)} (1 - L(y))$$

Theorem. The following are equivalent:

1. $\mathbf{p}$ is valid for $\mathcal{W}$
2. For every $\tilde{w} \in \mathcal{C}(\mathcal{W})$ there exists $r_{\tilde{w}} \in [0,1)$ such that

$$p_{t(\tilde{w})} \leq r_{\tilde{w}} \prod_{y \in \mathrm{Ext}(\tilde{w})} (1 - r_{\tilde{y}})$$

($t(\tilde{w})$ and $\mathrm{Ext}(\tilde{w})$ are well defined functions)

$$\mathbf{p} \text{ is valid for } \mathcal{W} \iff \mathbf{p} \text{ “is valid for” } \mathcal{C}(\mathcal{W})$$

The Corollary

Reminder:

$$\mathbf{p} \text{ is valid for } \mathcal{W} \supseteq \mathcal{B}(G) \implies \mathbf{p} \text{ is valid for } \mathcal{B}(G)$$
$$\mathbf{p} \text{ is valid for } \mathcal{W} \iff \mathbf{p} \text{ ``is valid for'' } \mathcal{C}(\mathcal{W})$$

There exists a set of walks $\mathcal{W} \supseteq \mathcal{B}(G)$ and $r : \mathcal{C}(\mathcal{W}) \to [0, 1)$ s.t.:

$$p_{t(\tilde{w})} \leq r_{\tilde{w}} \prod_{y \in \mathrm{Ext}(\tilde{w})} (1 - r_{\tilde{y}})$$

$$\Longrightarrow$$

$$\mathrm{avoid}(\mathbf{p}, G) > 0$$

For different choices of $\mathcal{W}$ we will obtain different conditions for $\mathrm{avoid}(\mathbf{p}, G) > 0$.

Example 1. All the walks

If we choose $\mathcal{W}$ to be the set of *all walks* on G, then

$$w \sim w' \iff t(w) = t(w')$$

The equivalence classes $\tilde{w_1}, \ldots \tilde{w_n}$ are bijective with $V(G) = [n]$. For a class $\tilde{w}_i$ with $t(\tilde{w}_i) = i$ we have $Ext(\tilde{w}_i) = \{\tilde{w}_j : j \in \Gamma_i\}$.

Example 1. All the walks

If we choose $\mathcal{W}$ to be the set of *all walks* on G, then

$$w \sim w' \iff t(w) = t(w')$$

The equivalence classes $\tilde{w}_1, \dots \tilde{w}_n$ are bijective with $V(G) = [n]$. For a class $\tilde{w}_i$ with $t(\tilde{w}_i) = i$ we have $Ext(\tilde{w}_i) = \{\tilde{w}_j : j \in \Gamma_i\}$.

Our Corollary yields then the asymmetric Lovász local lemma:

If there exist $r_1, \dots, r_n \in [0, 1)$ such that for every i:

$$p_i \leq r_i \prod_{j \in \Gamma_i} (1 - r_j)$$

Then $\text{avoid}(\mathbf{p}, G) > 0$.

Example 2. Non-backtracking walks

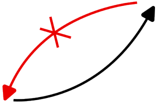

A walk w is *non-backtracking*, if it doesn't contain the sequence xyx (it doesn't traverse the same edge two times in a row).

All self-bounding walks are also non-backtracking, so we can apply the Corollary to the set $\mathcal{W}$ of non-backtracking walks on G.

In $\mathcal{W}$, have an equivalence class:

- for every directed pair (u, v) s.t. $uv \in E(G)$ (= for every choice of two last vertices of a walk)
- for every vertex in G (walks of length 1)

Example 2. Non-backtracking walks

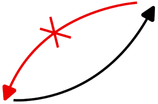

We obtain the *non-backtracking LLL*:

Let $A = \{(u, v) : uv \in E\}$. For $a = (u, v) \in A$, let $\Gamma_a = \{(v, w) \in A : w \neq u\}$. If there exist $0 \leq \{r_i\}_{i \in [n]}, \{r_a\}_{a \in A} < 1$ such that for all $i \in [n]$ and $a = (u, v) \in A$:

$$p_i \leq r_i \prod_{j \in \Gamma_i} (1 - r_{(i,j)}) \quad \text{and} \quad p_j \leq r_a \prod_{y \in \Gamma_a} (1 - r_y)$$

then $\text{avoid}(\mathbf{p}, G) > 0$.

Example 3. Self-bounding walks

And what if $\mathcal{W} = \mathcal{B}(G)$? We recover the Shearer's criterion:

Theorem (Shearer, 1985). Let $Z_G(\mathbf{x}) = \sum_{I \in \mathrm{Ind}(G)} \prod_{i \in I} x_i$ denote the independent set polynomial of G. (For $S \subseteq [n]$ we also write $Z_G(\mathbf{x}, S) = Z_{G[S]}(\mathbf{x})$.) Then:

$$\mathrm{avoid}(\mathbf{p}, G) > 0 \iff \forall_{S \subseteq [n]} Z_G(-\mathbf{p}, S) > 0$$

Example 3. Self-bounding walks

$$\mathrm{avoid}(\mathbf{p}, G) > 0 \iff \forall_{S \subseteq [n]} Z_G(-\mathbf{p}, S) > 0$$

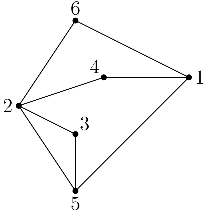

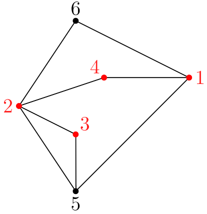

$$Z_G(\mathbf{x}) = 1 + x_1 + x_2 + x_3 + x_4 + x_5 + x_6 + x_1x_2 + x_1x_3 + x_3x_4 + x_3x_6 + x_4x_5 + x_3x_4x_6$$

$$Z_G(\mathbf{x}, \{1,2,3,4\}) = 1 + x_1 + x_2 + x_3 + x_4 + x_1x_2 + x_1x_3 + x_3x_4$$

Do It Yourself

The authors suggest that a good framework for defining own condition for $\text{avoid}(\mathbf{p}, G) > 0$ would be to choose a subset of $\mathcal{P}(V(G))$ as 'filters' and consider a following family of walks $\mathcal{W}$:

$w \in \mathcal{W}$ if any contiguous sub-walk of w which lies entirely within a filter F_i is self-bounding in F_i

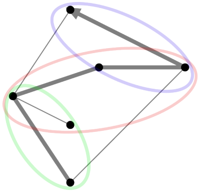

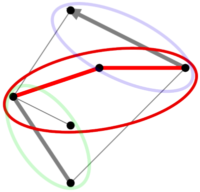

Do It Yourself

The authors suggest that a good framework for defining own condition for avoid$(\mathbf{p}, G) > 0$ would be to choose a subset of $\mathcal{P}(V(G))$ as 'filters' and consider a following family of walks $\mathcal{W}$:

$w \in \mathcal{W}$ if any contiguous sub-walk of w which lies entirely within a filter F_i is self-bounding in F_i

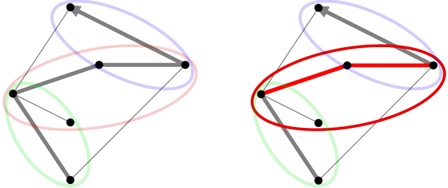

In particular, if we choose $F = \{\{u, v\} : uv \in E\}$ then $\mathcal{W}$ is the set of non-backtracking walks on G.

Part 2.

'The Lovász Local Lemma is not about probability'

Supermodular Functions

Let $f : \mathcal{P}(n) \to \mathbb{R}_{\geq 0}$.

- For $S \subseteq [n] \setminus \{i\}$,

$$\Delta_i f(S) = f(S \cup \{i\}) - f(S)$$

  is the *discrete derivative*
- f is *supermodular* if $\Delta_i(S)$ is increasing:

$$\forall_{i \in [n],\, S \subset T \subseteq [n] \setminus \{i\}} \; \Delta_i(S) \leq \Delta_i(T)$$

Supermodular Functions

We will consider decreasing supermodular functions:

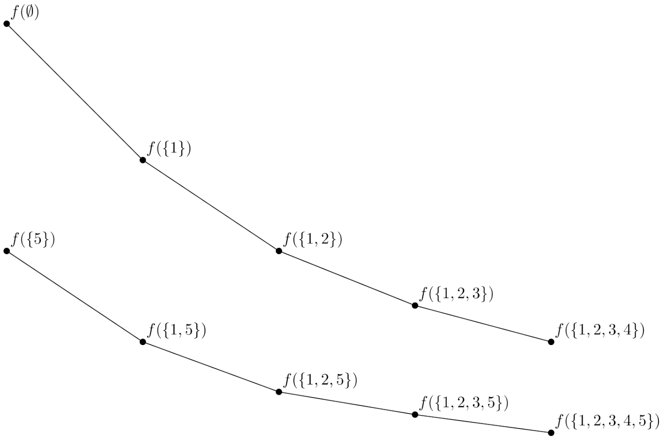

Supermodular functions

We will say that a function f *factorizes according to* $\mathbf{p}, G$ if for all $i \in [n]$ and $S \subseteq [n] \setminus (\Gamma_i \cup \{i\})$:

$$f(S \cup \{i\}) \geq (1 - p_i) f(S)$$

Supermodular functions

We will say that a function f *factorizes according to* $\mathbf{p}, G$ if for all $i \in [n]$ and $S \subseteq [n] \setminus (\Gamma_i \cup \{i\})$:

$$f(S \cup \{i\}) \geq (1 - p_i) f(S)$$

This generalizes our previous 'probabilistic' setting:

If μ is a probability measure s.t. $\mu(F_i) = p_i$, then
G is a dependency graph for the events $F_1, \ldots, F_n$ iff
$f(F_{i1}, \ldots, F_{ik}) := \mu(\overline{F_{i1}}, \ldots, \overline{F_{ik}})$ factorizes according to $\mathbf{p}, G$.

(It is easy to verify that f defined this way is supermodular)

Generalizing Shearer's criterion

Reminder: Shearer's criterion states that:

$$\text{avoid}(\mathbf{p}, G) > 0 \iff \forall_{S \subseteq [n]} Z_G(-\mathbf{p}, S) > 0$$

Theorem. Let G be a graph on $[n]$, and $\mathbf{p} \in [0,1)^n$.

(i) If $Z_G(-\mathbf{p}; S) > 0$ for every $S \subseteq [n]$ and f is a supermodular function with $f(\emptyset) > 0$ that factorizes according to $G, \mathbf{p}$, then for every $S \subseteq [n]$:

$$f(S) \geq f(\emptyset) \cdot Z_G(-\mathbf{p}, S) > 0$$

(ii) If $Z_G(-\mathbf{p}; S) < 0$ for some $S \subseteq [n]$ then there exists a supermodular function f with $f(\emptyset) > 0$ that factorizes according to $G, \mathbf{p}$, such that $f(\mathbf{S}) = 0$ for some $S \subseteq [n]$

Supermodular Local Lemma

As a corollary, we obtain:

Theorem. If f is a supermodular function with $f(\emptyset) > 0$ that factorizes according to $\mathbf{p}, G$ and there exist $r_1, \ldots r_n \in [0, 1)$ such that $p_i \leq r_i \prod_{j \in \Gamma_i}(1 - r_j)$ for all $i \in [n]$, then for every $S \subseteq [n]$

$$f(S) \geq f(\emptyset) \prod_{i \in S}(1 - r_i) > 0$$

Application: Quantum LLL

Given a vector space V and its subspaces $X_1, X_2, \ldots, X_n$, we seek a condition for $\dim(X_1 \cap X_2 \cap \ldots \cap X_n) > 0$.

- For a subspace $X \subseteq V$, let $R(X) := \frac{dim(X)}{dim(V)}$ be the *relative dimension* of X with respect to V.
- $R(X|Y) := \frac{R(X \cap Y)}{R(Y)} = \frac{\dim(X)}{\dim(Y)}$ - *relative dimension* of X with respect to Y.
- X is *mutually R-independent* of $X_1, \ldots X_l$ if $R(X \cap X_1 \cap \ldots \cap X_l) = R(X)R(X_1 \cap \ldots \cap X_l)$.
- A graph G on $[n]$ is an *R-dependency graph* for subspaces $X_1, \ldots, X_n$ if for every $i \in [n]$ and every $S \subseteq [n] - (\Gamma_i \cup \{i\})$, X_i is mutually R-independent of subspaces $\{X_j\}_{j \in S}$.

Application: Quantum LLL

Given a vector space V and its subspaces $X_1, X_2, \ldots, X_n$, we seek a condition for $\dim(X_1 \cap X_2 \cap \ldots \cap X_n) > 0$.

What makes this problem different from the probabilistic analogue is that while for any random events

$$\mathbb{P}(A|B) + \mathbb{P}(\overline{A}|B) = 1$$

it is not always the case that

$$R(A|B) + R(A^{\perp}|B) = 1$$

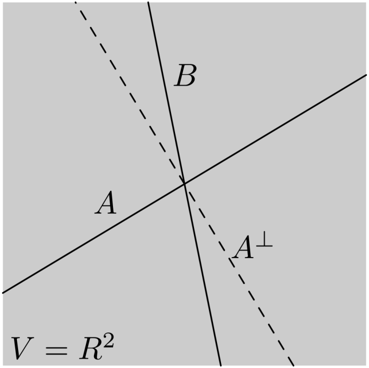

Application: Quantum LLL

Given a vector space V and its subspaces $X_1, X_2, \ldots, X_n$, we seek a condition for $\dim(X_1 \cap X_2 \cap \ldots \cap X_n) > 0$.

We define $\mathbf{p}$: $p_i = 1 - R(X_i)$.

It turns out that $f(S) := R(\cap_{j \in S} X_j)$ is supermodular and it factorizes according to $\mathbf{p}, G$. Using the supermodular lemmata from above, we easily obtain those two results (which were shown independently before):

Application: Quantum LLL

Given a vector space V and its subspaces $X_1, X_2, \ldots, X_n$, we seek a condition for $\dim(X_1 \cap X_2 \cap \ldots \cap X_n) > 0$.

We define $\mathbf{p}$: $p_i = 1 - R(X_i)$.

It turns out that $f(S) := R(\cap_{j \in S} X_j)$ is supermodular and it factorizes according to $\mathbf{p}, G$. Using the supermodular lemmata from above, we easily obtain those two results (which were shown independently before):

Quantum LLL (2012) If there exist $r_1, \ldots, r_n \in [0,1)$ s.t. $p_i \leq r_i \prod_{j \in \Gamma_i}(1-r_j)$ for all $i \in [n]$, then $R(\cap_{i=1}^{n} X_i) \geq \prod_{i \in [n]}(1-r_i) > 0$

Application: Quantum LLL

Given a vector space V and its subspaces $X_1, X_2, \ldots, X_n$, we seek a condition for $\dim(X_1 \cap X_2 \cap \ldots \cap X_n) > 0$.

We define $\mathbf{p}$: $p_i = 1 - R(X_i)$.

It turns out that $f(S) := R(\cap_{j \in S} X_j)$ is supermodular and it factorizes according to $\mathbf{p}, G$. Using the supermodular lemmata from above, we easily obtain those two results (which were shown independently before):

Quantum Shearer (2016)
$\text{avoid}(\mathbf{p}, G) > 0 \iff R(\cap_{i=1}^{n} X_i) > 0$

Thank you!