Sunflower theorems in monotone circuit complexity

Bruno Pasqualotto Cavalar

Dissertação apresentada
ao
Instituto de Matemática e Estatística
da
Universidade de São Paulo
para
obtenção do título
de
Mestre em Ciências

Programa: Ciência da Computação
Orientador: Prof. Dr. Yoshiharu Kohayakawa

Durante o desenvolvimento deste trabalho o autor recebeu auxílio financeiro da FAPESP (Processos 2018/22257-7 e 2018/05557-7). O presente trabalho foi realizado com apoio da Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Código de Financiamento 001.

São Paulo, outubro de 2020

Sunflower theorems
in monotone circuit complexity

Esta versão da dissertação contém as correções e alterações sugeridas pela Comissão Julgadora durante a defesa da versão original do trabalho, realizada em 09/09/2020. Uma cópia da versão original está disponível no Instituto de Matemática e Estatística da Universidade de São Paulo.

Comissão Julgadora:

- Prof. Dr. Yoshiharu Kohayakawe (orientador) - IME-USP
- Prof. Dr. Benjamin Rossman - Duke University
- Prof. Dr. Mrinal Kumar - IIT Bombay

Acknowledgements

I would like to thank Yoshiharu Kohayakawa for having been my advisor since my early years at IME-USP. Thank you for always being very supportive of my interests and for sharing with me so much of your knowledge and experience. Thank you as well for being so generous of your time and resources, and available to support me in everything I needed. I will always be thankful for your dedication and care as an advisor.

I want to thank Benjamin Rossman for hosting me for six months at the University of Toronto. I'm very thankful for your kindness and generosity as an advisor, for teaching me so much about circuit complexity, and for making me feel welcome in UofT. Thank you for your hospitality and for always making space in your schedule to discuss interesting problems.

Many thanks to Mrinal Kumar, for teaching me many things in computational complexity, and for being both a colleague and a friend. Thanks for that unforgettable dinner at Tibet Kitchen. Thank you for helping me so much with generous advice and support.

Thanks to all the members of the committee, Yoshiharu Kohayakawa, Benjamin Rossman and Mrinal Kumar for your valuable feedback and encouragement.

A special thanks to Igor Carboni Oliveira for igniting my interest in computational complexity, and for providing inumerous research and career advice during the last few years. Thank you also for introducing me to circuit complexity. I look forward to working with you in Warwick.

I'm very thankful for all the support I received during these years by many members of the theory group of IME-USP. In particular, I thank Yoshiko Wakabayashi, Marcel Kenji de Carli Silva, Cristina Gomes Fernandes and Carlos Eduardo Ferreira for all the teaching, advice and encouragement.

I'm also highly indebted to my colleagues and friends who made my time at USP much more enjoyable, and from whom I learnt so much. Thank you Nathan for your camaraderie since the beginning, and for being a very fun and perceptive conversation partner. Thank you Gabriel Ferreira Barros and Tássio Naia dos Santos, for being a pleasure to work with, and for sharing your experience with me.

I wish to thank Ian Mertz for being a kind officemate at UofT, and for sharing with me his highly enjoyable collection of teas. Thanks also to Alex for many interesting conversations, and particularly for his good taste in movies and beer, and to Chris for his openness to discuss difficult questions.

I have an enormous gratitude for the way the 46ers received me in Toronto. You have a special place in my heart and memory, always. Thank you Courtney, for supporting all of us and for your big and open heart. Thank you Rachel&Jeff for your friendship, Kevin for your camaraderie and prayers, Vera for your energy, Evan for your patience and for all you taught me. I came back from you all with a deepened understanding of the expression *grace of God*.

I remember with warmth the support I received from Christ the King Anglican Church while I was in Toronto. A special thanks to pastor Roger, for your hospitality, friendship, teaching, prayers, guidance, advice and encouragement. My enjoyment of my time in Toronto is in large part due to you. I also wish to thank pastor Keith for his careful sermon series on the books of Samuel, which were one of the highlights of my week, and Sam and Logan for their welcome and friendship. A word of thanks to the church at large for the fervour of their worship, which sustained me in my weekly work.

I also wish to thank *Igreja Presbiteriana do Butantã* for being my extended family for the last 3 years. Thanks for all the elders and pastors for providing a space of learning, discussion and growth, and for supporting me and my wife in everything we needed. A special thanks to my close friends Felipe, Leo, Pedro and William for discussing *everything* with openness and good humour. *Sentirei saudades.*

Above all other people, I wish to thank my family for their unwavering support and understanding. Thanks to my dear wife Thais for putting up with many sleepless nights, and for my parents Eric and Claudete for giving me everything I needed to succeed. *Eu amo vocês.*

Finally and above all else, I thank the Lord Jesus Christ, who lives and reigns now and forever. It was His will that I should meet all this people, and His will that I should enjoy these opportunities. Blessed be His name forever.

Resumo

CAVALAR, B. P. **Teoremas de girassol em complexidade de circuitos monótonos**. 2020. 60 f. Dissertação (Mestrado) - Instituto de Matemática e Estatística, Universidade de São Paulo, São Paulo, 2020.

Alexander Razborov (1985) desenvolveu o método das aproximações para obter cotas inferiores para o tamanho de circuitos monótonos que decidem se um grafo contém uma clique de uma dada ordem. Dado um circuito monótono "pequeno", essa técnica consiste em encontrar uma função Booleana monótona que aproxima o circuito numa distribuição de interesse, mas comete erros de computação nessa mesma distribuição. Para provar que tal função é de fato uma boa aproximação, Razborov utilizou o lema dos girassóis de Erdős e Rado (1960).

Essa técnica foi aprimorada por Alon e Boppana (1987) para mostrar cotas inferiores para uma gama muito maior de problemas computacionais monótonos. Nesse trabalho, os autores também melhoraram o resultado de Razborov para o problema da clique, utilizando uma variação relaxada de girassóis.

Mais recentemente, Rossman (2010) desenvolveu uma outra variação de girassóis, hoje chamada de "girassóis robustos", para obter cotas inferiores para o problema da clique em grafos aleatórios. Em seguida, o conceito de girassóis robustos encontrou aplicações em várias áreas da complexidade computacional, tais como esparsificação de DNFs, extratores de aleatoriedade e teoremas de "lifting". Ainda mais recente foi um resultado de impacto de Alweiss, Lovett, Wu e Zhang (2020), que mostrou cotas melhores que a de Rossman para girassóis robustos. Esse resultado foi utilizado para obter o progresso mais significativo na conjectura dos girassóis desde a sua origem em 1960.

Nesse trabalho, vamos mostrar como os desenvolvimentos recentes em teoremas de girassol podem ser aplicados para melhorar cotas inferiores para circuitos monótonos. Em particular, vamos mostrar a melhor cota inferior para um circuito monótono obtida até o momento, quebrando um recorde de 20 anos obtido por Harnik e Raz (2000). Iremos também melhorar a cota inferior de Alon e Boppana para a função clique numa faixa levemente mais restrita de tamanhos de clique. Esses resultados foram elaborados numa colaboração do aluno com Benjamin Rossman e Mrinal Kumar, durante uma visita à Universidade de Toronto, e um resumo foi aceito na conferência LATIN 2020.

Palavras-chave: complexidade computational, complexidade de circuitos, complexidade de circuitos monótonos, combinatória, combinatória extremal, girassóis, combinatória probabilistica, teoria extremal dos conjuntos

Abstract

CAVALAR, B. P. **Sunflower theorems in monotone circuit complexity**. 2020. 60 f. Dissertação (Mestrado) - Instituto de Matemática e Estatística, Universidade de São Paulo, São Paulo, 2020.

Alexander Razborov (1985) developed the approximation method to obtain lower bounds on the size of monotone circuits deciding if a graph contains a clique. Given a "small" circuit, this technique consists in finding a monotone Boolean function which approximates the circuit in a distribution of interest, but makes computation errors in that same distribution. To prove that such a function is indeed a good approximation, Razborov used the sunflower lemma of Erdős and Rado (1960).

This technique was improved by Alon and Boppana (1987) to show lower bounds for a larger class of monotone computational problems. In that same work, the authors also improved the result of Razborov for the clique problem, using a relaxed variant of sunflowers.

More recently, Rossman (2010) developed another variant of sunflowers, now called "robust sunflowers", to obtain lower bounds for the clique problem in random graphs. In the following years, the concept of robust sunflowers found applications in many areas of computational complexity, such as DNF sparsification, randomness extractors and lifting theorems. Even more recent was the breakthrough result of Alweiss, Lovett, Wu and Zhang (2020), which improved Rossman's bound on the size of hypergraphs without robust sunflowers. This result was employed to obtain the most significant progress on the sunflower conjecture since its inception in 1960.

In this work, we will show how the recent progresses in sunflower theorems can be applied to improve monotone circuit lower bounds. In particular, we will show the best monotone circuit lower bound obtained up to now, thus breaking a 20-year old record of Harnik and Raz (2000). We will also improve the lower bound of Alon and Boppana for the clique function in a slightly more restricted range of clique sizes. These results were obtained in a collaboration of the student with Benjamin Rossman and Mrinal Kumar, during a visit to the University of Toronto, and an extended abstract was accepted to the LATIN 2020 conference.

Palavras-chave: computational complexity, circuit complexity, monotone circuit complexity, combinatorics, extremal combinatorics, sunflowers, probabilistic combinatorics, extremal set theory

Contents

Chapter 1

Introduction

1.1 Computational complexity

The main goal of computational complexity is to understand the amount of computational resources needed to complete a computational task in a variety of computational models. Arguably, the most famous question in this area is **P** vs. **NP**, which asks whether any decision problem that can be solved by a polynomial-time *nondeterministic* Turing machine can also be solved by a *deterministic* polynomial-time Turing machine. A computational model that is very prominent in computational complexity is that of the *Boolean circuit*, which we define below.

In *circuit complexity*, the main goal is to prove that Boolean circuits computing a given Boolean function must be "complex" with respect to some circuit parameter, such as size and depth. Proving such lower bounds for general circuits has proved to be difficult, and results in full generality are usually weak. Indeed, since any polynomial-time algorithm can be implemented by a sequence of polynomial-size circuits (one for each input length), obtaining a superpolynomial lower bound on the minimum circuit size of any problem in **NP** is enough to separate **P** from **NP**[1], which seems far out of reach for current techniques. For this reason, much of the research in circuit complexity has focused in restricted classes of circuits. One of the most important such classes is that of *monotone circuits*, the main circuit model we will study in this work.

1.2 Monotone circuits

For $x, y \in \{0,1\}^n$, we write $x \leqslant y$ whenever $x_i \leqslant y_i$ for all $i \in [n]$. A Boolean function $f : \{0,1\}^n \to \{0,1\}$ is said to be *monotone* if, for all $x, y \in \{0,1\}^n$ such that $x \leqslant y$, we have $f(x) \leqslant f(y)$. A *monotone Boolean circuit* is a Boolean circuit without negation gates (**NOT**). One readily sees that any Boolean function computed by a monotone circuit has to be monotone; moreover, every monotone function can be computed by a monotone circuit. The *monotone complexity* of a monotone Boolean function is defined as the size of the smallest monotone circuit that computes the function.

There are many natural monotone Boolean functions that have been widely studied in complexity theory. One such function is $\mathsf{Majority} : \{0,1\}^n \to \{0,1\}$, which accepts an input if the majority of its bits are equal to one. Importantly, the majority function is known to be computable by polynomial-size monotone formulas [AKS83]. Another monotone Boolean function of major importance is $\mathsf{Clique}(n,k) : \{0,1\}^{\binom{n}{2}} \to \{0,1\}$, which accepts a graph (given by its adjacency matrix) if and only if it contains a clique of size k. It is well known that computing $\mathsf{Clique}(n,k)$ (i.e.: deciding if, given a graph G and a number k, the graph G contains a k-clique) is a **NP**-complete problem. The first superpolynomial lower bound on the monotone complexity of any monotone Boolean function was obtained for $\mathsf{Clique}(n,k)$ in a seminal work by Razborov [Raz85b].

[1] See, for instance, Chapter 6 of [AB09].

1.3 The approximation method and sunflowers

Razborov's approach inaugurated a technique which came to be known as the *approximation method.* Given a monotone circuit C of "small size", it consists into constructing gate-by-gate, in a bottom-up fashion, another circuit $\widetilde{C}$ that approximates C on a distribution of interest. One then exploits the structure of this approximating circuit to prove that it differs from $\mathsf{Clique}(n,k)$ under the same distribution, thus implying that no "small" circuit can compute this function. For monotone circuit lower bounds, showing that $\widetilde{C}$ does indeed approximate C is usually the hardest part, involving the use of a combinatorial lemma – which, in the case of [Raz85b], was the *sunflower lemma* of Erdős and Rado [ER60].[2]

This technique was leveraged to obtain lower bounds for a host of other monotone problems by Alon and Boppana [AB87]. In that work, the authors also improve Razborov's lower bound on the clique function by showing that it requires monotone circuits with exponential size. Their technical contribution is to employ a weaker notion of "sunflowers", proving a better corresponding bound.

Another type of sunflowers, called robust sunflowers, was developed by Rossman [Ros14] with the purpose of achieving better lower bounds for $\mathsf{Clique}(n,k)$ in the regime when k is constant and the distribution is morally the Erdős-Rényi random graph. A recent breakthrough of Alweiss, Lovett, Wu and Zhang [ALWZ19] significantly improved the upper bound on the size of uniform families without robust sunflowers, implying also better bounds for the standard sunflower of Erdős and Rado. A simpler proof with a slight improvement was given by Rao [Rao20].

1.4 Our contribution

In this work, we will show how these recent developments in sunflower theorems lead to better monotone circuit lower bounds. To make our presentation self-contained, we will first survey the main sunflower-type theorems and their proof techniques in Chapter 3, and we will give a full proof of the sunflower breakthrough [ALWZ19, Rao20] in Chapter 4.

In Chapter 5, we will switch gears to monotone circuit complexity, and we will see how variants of sunflowers help in the construction of good approximators. Our presentation is general enough to accommodate many applications of the approximation method, including all those of [Raz85b, Raz85a, AB87]. We will also introduce a novel notion of *abstract sunflowers*, which generalizes the notion of robust sunflowers to arbitrary distributions.

Finally, we will apply the general framework of Chapter 5 to specific problems in monotone circuit complexity. In particular, we will prove in Chapter 6 the best monotone circuit lower bound to-date, beating a 20 years old record of Harnik and Raz [HR00], and we will improve Alon and Boppana's [AB87] bound for the clique function in Chapter 7. The combinatorial heart of Chapter 6 will be the improved sunflowers bounds of [ALWZ19, Rao20], whereas in Chapter 7 we will introduce and prove bounds for a novel notion of *clique sunflowers*, tailored specifically for the clique problem.

These results appeared in a joint work of the author with Benjamin Rossman and Mrinal Kumar, carried out during a visit of the author to the University of Toronto and published in the LATIN 2020 conference [CKR20].

[2] As a side note, the approximaton method was also used by Razborov to prove that $\mathsf{Majority}$ cannot be computed by bounded-depth circuits with parity gates [Raz87]. However, in this case the approximators were low-degree polynomials.

Chapter 2

Definitions and Preliminaries

We will present in this chapter basic definitions, notation and results that are going to be used throughout the thesis.

2.1 Basic notation

For a positive integer n, we let $[n] := \{1, 2, \ldots, n\}$. The function $\log(\cdot)$ denotes the natural logarithm. Given a set Γ, we denote by 2^Γ the family of the subsets of Γ. For a positive integer ℓ, we denote by $\binom{\Gamma}{\ell}$ the family of the subsets of Γ with size ℓ, and $\binom{\Gamma}{\leqslant \ell}$ denotes the family of the subsets of Γ with size at most ℓ. Given a proposition P, we will write $\mathsf{Ind}[P]$ to denote the *indicator function* of P, which assumes value 1 if P is true and 0 otherwise. For example, the expression $\mathsf{Ind}[x \in \Gamma]$ is equal to 1 if and only if $x \in \Gamma$. When a base set Γ is fixed and $A \subseteq \Gamma$, we denote by A^c the complement of A on Γ. Given a set family $\mathcal{F}$, we denote by $\bigcap \mathcal{F}$ the intersection $\bigcap_{F \in \mathcal{F}} F$.

2.1.1 Fields and polynomials

Let p be a prime number. We denote by $\mathbb{F}_p$ the finite field of order p. Moreover, we write $\mathbb{F}_p[x]$ to denote the set of all univariate polynomials over the field $\mathbb{F}_p$.

2.2 Combinatorics and probability

2.2.1 Hypergraphs and set systems

A *hypergraph* $\mathcal{H}$ is a pair of sets $(\Gamma, \mathcal{E})$ such that $\mathcal{E} \subseteq 2^\Gamma$. A hypergraph $\mathcal{H} = (\Gamma, \mathcal{E})$ is said to be ℓ-uniform if $\mathcal{E} \subseteq \binom{\Gamma}{\ell}$. An element of Γ is called a *vertex* of G, and an element of $\mathcal{E}$ is called an *edge* of G. We may also call Γ the *vertex set* of $\mathcal{H}$, and $\mathcal{E}$ the *edge set* of $\mathcal{H}$. When the vertex and edge sets of a hypergragh $\mathcal{H}$ are not explicitly defined, we will denote the vertex set of $\mathcal{H}$ by $V(\mathcal{H})$, and the edge set by $E(\mathcal{H})$.

If the vertex set of a hypergraph $\mathcal{H}$ is a set Γ, we will often identify $\mathcal{H}$ with its edge set and write $\mathcal{H} \subseteq 2^\Gamma$ to denote that $V(\mathcal{H}) = \Gamma$. This means that, when we write $e \in \mathcal{H}$, we formally mean $e \in E(\mathcal{H})$. In such cases, we may also write that $\mathcal{H}$ *is a hypergraph on* Γ, or, equivalently, that $\mathcal{H}$ is *set system* on Γ or a *set family* on Γ.

Given two hypergraphs $\mathcal{F}$ and $\mathcal{H}$, we say that $\mathcal{F}$ is a *subhypergraph* of $\mathcal{H}$ if $V(\mathcal{F}) \subseteq V(\mathcal{H})$ and $E(\mathcal{F}) \subseteq E(\mathcal{H})$. Furthermore, for a set $S \subseteq V(\mathcal{H})$, we denote by $\mathcal{H}[S]$ the hypergraph satisfying $V(\mathcal{H}[S]) = S$ and $E(\mathcal{H}[S]) = \{E \in \mathcal{H} : E \subseteq S\}$. We say that $\mathcal{H}[S]$ is the subhypergraph of $\mathcal{F}$ *induced by* S.

A set $I \subseteq V(\mathcal{H})$ is said to be *independent* in a hypergraph $\mathcal{H}$ if there does not exist $e \in E(\mathcal{H})$ such that $e \subseteq I$.

For a set $A \subseteq V(\mathcal{H})$, we define the *degree* $d_\mathcal{H}(A)$ *of* A in the hypergraph $\mathcal{H}$ in the following way:

$$d_\mathcal{H}(A) := \left|\{e \in E(\mathcal{H}) : A \subseteq e\}\right|.$$

We will use the shorthand $d_{\mathcal{H}}(v) := d_{\mathcal{H}}(\{v\})$ for $v \in V(\mathcal{H})$. We also define the *t-degree* $\Delta_t(\mathcal{H})$ *of* $\mathcal{H}$ as

$$\Delta_t(\mathcal{H}) := \max\{d_{\mathcal{H}}(A) : A \subseteq V(\mathcal{H}), |A| = t\}.$$

2.2.2 Graphs

A *graph* is a 2-uniform hypergraph. A *subgraph* H of G is a graph such that $E(H) \subseteq E(G)$. A *clique* on G is a subgraph H of G for which there exists a set $K \subseteq V(G)$ such that $E(H) = \binom{K}{2}$. In this case, when $|K| = k$, we say that H is a *k-clique* of G.

In the context of graphs, it will often be convenient to let $[n]$ be the vertex set. We then let $\mathcal{G}^n$ be the set of all graphs with vertex set $[n]$. We identify a graph with vertex set $[n]$ with its set of edges, so that $\mathcal{G}^n = 2^{\binom{[n]}{2}}$.

2.2.3 Distributions

We will consistently write random objects using boldface symbols. For a set Γ and $p \in [0,1]$, we write $\mathbf{W} \subseteq_p \Gamma$ to denote that $\mathbf{W}$ is a random subset of Γ such that every element of Γ is contained in $\mathbf{W}$ independently with probability p. In this case, we say that $\mathbf{W}$ follows a *p-biased distribution*.

An important example of this distribution is the *Erdős-Rényi random graph* $G(n,p)$, defined as follows. Given a function $p : \mathbb{N} \to [0,1]$, we denote by $G(n,p)$ the random graph with vertex set equal to $[n]$ which satisfies $E(\mathbf{G}) \subseteq_p \binom{[n]}{2}$, for $\mathbf{G} \sim G(n,p)$.

When $\mathbf{W} \subseteq_p \Gamma$ for some $p : \mathbb{N} \to [0,1]$ and $A = A(n)$ is a sequence of events on this probability space, we say that A holds *with high probability* if $\lim_n \Pr[A] = 1$. A function of the form $p : \mathbb{N} \to [0,1]$ will be called a *probability function* or *probability sequence*.

We will write $\mathbf{W} \sim \mathcal{U}(\Gamma)$ to denote that $\mathbf{W}$ is chosen uniformly at random from a set Γ. For a positive integer $M > 0$, we will also write $\mathbf{W} \subseteq_M \Gamma$ as a shorthand for $\mathbf{W} \sim \mathcal{U}\left(\binom{\Gamma}{M}\right)$. In this case, we say that $\mathbf{W}$ follows a *M-uniform distribution* or a *M-slice distribution*.

Given an arbitrary distribution μ, its *support*, denoted by $\mathrm{supp}(\mu)$, is the set of all $\omega \in \Omega$ such that $\Pr_{\mathbf{x}\sim\mu}[\mathbf{x} = \omega] > 0$.

2.2.4 Inequalities

Here we state a few inequalities of probability theory that we are going to use frequently. The first three are standard.

Proposition 2.2.1 (Union bound)**.** *Let $\mathcal{A}$ be a finite collection of events. We have*

$$\Pr\left[\bigcup_{A\in\mathcal{A}} A\right] \leqslant \sum_{A\in\mathcal{A}} \Pr[A].$$

Proposition 2.2.2 (Markov's inequality)**.** *Let $\mathbf{X}$ be a random variable and $a > 0$. We have*

$$\Pr[\mathbf{X} \geqslant a] \leqslant \frac{\mathbb{E}[\mathbf{X}]}{a}.$$

Proposition 2.2.3 (Jansen's inequality)**.** *Let $f : \mathbb{R} \to \mathbb{R}$ be a convex function and let $\mathbf{X}$ be a random variable taking values in $\mathbb{R}$. We have*

$$\mathbb{E}[f(\mathbf{X})] \geqslant f(\mathbb{E}[\mathbf{X}]).$$

The following inequality was proved in [Jan90]. A proof of it can be found as Theorem 2.18 in [JŁR00].

Lemma 2.2.4 (Janson's inequality)**.** *Let $\mathcal{F}$ be a nonempty hypergraph on Γ and let $\mathbf{W} \subseteq_p \Gamma$. Define μ and Δ in the following way:*

$$\mu := \sum_{F \in \mathcal{F}} \Pr[F \subseteq \mathbf{W}],$$
$$\Delta := \sum_{\substack{F,H \in \mathcal{F} \\ F \cap H \neq \emptyset}} \Pr[F \cup H \subseteq \mathbf{W}].$$

Then we have

$$\Pr[\mathbf{W} \text{ is independent in } \mathcal{F}] \leqslant \exp\{-\mu^2/\Delta\}.$$

2.2.5 Probability thresholds

Fix a finite set Γ and let $\mathcal{F} \subseteq 2^\Gamma$. We say that $\mathcal{F}$ is a *monotone* property if $A \in \mathcal{F}$ and $A \subseteq B$ implies $B \in \mathcal{F}$. Equivalently, we can say that $\mathcal{F}$ is monotone if and only if the Boolean function $\mathsf{Ind}[\cdot \in \mathcal{F}]$ is monotone. (See Section 2.3.4 for monotone Boolean functions.)

Let now $\Gamma = \Gamma(n)$ be a sequence of sets, and let $\mathcal{F} = \mathcal{F}(n) \subseteq 2^\Gamma$ be a sequence of monotone properties. Let moreover $\widehat{p} = \widehat{p}(n)$ be a probability sequence. We say that $\widehat{p}$ is a *probability threshold* for the property $\mathcal{F}$ if the following holds for every probability function $p = p(n)$:

$$\lim_n \Pr_{\mathbf{W} \subseteq_p \Gamma}[\mathbf{W} \in \mathcal{F}] = \begin{cases} 0 & \text{if} \quad p \ll \widehat{p}, \\ 1 & \text{if} \quad p \gg \widehat{p}. \end{cases}$$

The following is a fundamental fact about probability threshold functions, due to Bollobás and Thomason.

Theorem 2.2.5 ([BT87])**.** *Every monotone property has a probability threshold.*

2.2.6 A construction of c-wise independent random variables

A set of random variables $\mathcal{X}$ is called *c-wise independent* if every subset of $\mathcal{X}$ of at most c random variables is mutually independent. We now describe a construction of c-wise random variables due to Alon, Babai and Itai [ABI86].

Fix a prime q and let c be a positive integer such that $c \leqslant q$. Let Ω be the set of polynomials in $\mathbb{F}_p[x]$ of degree at most $c-1$. It is easy to see that $|\Omega| = q^c$. Let $\mathbf{P}$ be a polynomial chosen uniformly at random from Ω, and define the random variables $\mathbf{X}_i := \mathbf{P}(i)$ for every $i \in [q]$. We claim that these random variables are uniform in $\mathbb{F}_q$ and that they are c-wise independent. This follows from the following simple observation about linear systems over $\mathbb{F}_q$: for $\ell \in [c]$, any such system of ℓ linear equations with c variables has exactly $q^{c-\ell}$ solutions. From this we can easily show both the uniformity and c-wise independence of the random variables $\mathbf{X}_1, \ldots, \mathbf{X}_q$.

2.3 Computational complexity

2.3.1 Posets and lattices

A pair $(P, \leqslant)$ where P is a set and $\leqslant$ is a relation on P is said to be a *partially ordered set* (abbreviated as *poset*) if, for all a, b and $c \in P$, we have

1. $a \leqslant a$ (reflexivity);
2. $a \leqslant b$ and $b \leqslant a$ implies $a = b$ (antisymmetry);
3. $a \leqslant b$ and $b \leqslant c$ implies $a \leqslant c$ (transitivity);

If for every $a, b \in P$ there exists an unique supremum and un unique infimum, we say that P is a *lattice*. A *closure operator* in a poset $(P, \leqslant)$ is a function $\mathrm{cl} : P \to P$ such that

1. $x \leqslant \mathrm{cl}(x)$ (extensive);
2. $x \leqslant y \Longrightarrow \mathrm{cl}(x) \leqslant \mathrm{cl}(y)$ (increasing);
3. $\mathrm{cl}(\mathrm{cl}(x)) = \mathrm{cl}(x)$ (idempotent).

2.3.2 Boolean functions

A very important lattice is the lattice $\mathsf{Bool}(\Gamma)$ of all Boolean functions $f : 2^\Gamma \to \{0, 1\}$, ordered by the following relation: we say that $f, g \in \mathsf{Bool}(\Gamma)$ satisfy $f \leqslant g$ if $f(X) \leqslant g(X)$ for all $X \in 2^\Gamma$. The supremum of this lattice correponds to the logical **OR** ($\vee$) and the infimum corresponds to the logical **AND** ($\wedge$).

For each $\gamma \in \Gamma$, the function $x_\gamma \in \mathsf{Bool}(\Gamma)$ is defined as $x_\gamma(X) = \mathsf{Ind}[\gamma \in X]$. For each $S \subseteq \Gamma$, the function $\mathbb{1}_S \in \mathsf{Bool}(\Gamma)$ is defined as $\mathbb{1}_S(X) = \mathsf{Ind}[S \subseteq X]$.

2.3.3 Circuit complexity

For every finite set Γ, a Γ-*input, single-output* $\{\vee, \wedge, \neg\}$-*Boolean circuit* is a directed acyclic graph with $|\Gamma|$ sources and one sink. The sources are called *inputs*, and each element of Γ is the label of one input. All nonsource vertices are called *gates* and are labeled with one of $\{\vee, \wedge, \neg\}$. The vertices labeled with $\vee$ and $\wedge$ have *fan-in* (that is, in-degree) equal to 2. The vertices labeled with $\neg$ always have fan-in 1 and are called *negation gates*. The *size* of a circuit C, denote by $\mathrm{size}(C)$, is the number of gates it contains. Given a set $X \subseteq \Gamma$, then the the *output* of C in X, defined by $C(X)$, is defined in a natural way. We say that C computes a Boolean function $f : 2^\Gamma \to \{0, 1\}$ if $C(X) = f(X)$ for all $X \subseteq \Gamma$. We define the *circuit complexity* of f, denoted by $\mathsf{C}_{\mathsf{size}}(f)$, as the size of the minimum circuit computing f.

Since a circuit C is a directed acyclic graph, it admits a *topological ordering* of its vertices. In other words, it is possible to arrange the vertices of C in a sequence $x_1, x_2, \ldots, x_n, f_1, f_2, \ldots, f_s$ where each x_i is an input of C and each f_j is a gate of C, such that either $f_j = f_\ell \circ f_k$ for $\ell, k < j$ and $\circ \in \{\vee, \wedge\}$ or $f_j = \neg f_k$ for $k < j$. Such a sequence is called a *straight-line program* for C.

2.3.4 Monotone circuit complexity

A Boolean function $f : 2^\Gamma \to \{0, 1\}$ is said to be *monotone* if, for all $X, Y \in 2^\Gamma$ such that $X \subseteq Y$, we have $f(X) \leqslant f(Y)$. We also let $\mathsf{MonBool}\,(\Gamma) \subseteq \mathsf{Bool}(\Gamma)$ denote the sublattice of monotone Boolean functions. A set $X \subseteq \Gamma$ is said to be a *minterm* of $f \in \mathsf{MonBool}\,(\Gamma)$ if $f(X) = 1$ but $f(Y) = 0$ for all $Y \subsetneq X$.

A *monotone Boolean circuit* is a Boolean circuit without negation gates, or, equivalently, an $\{\vee, \wedge\}$-circuit. One readily sees that any Boolean function computed by a monotone circuit has to be monotone; moreover, every monotone function can be computed by a monotone circuit. The *monotone complexity* of a monotone Boolean function f, denoted by $\mathsf{C}_{\mathsf{mon}}(f)$, is defined as the size of the smallest monotone circuit that computes the function.

Chapter 3

Sunflower theorems

Throughout this chapter, we will fix a finite set Γ, which we call *base set*, and we will always denote the size of Γ by n.

3.1 The standard sunflower lemma of Erdős and Rado

The celebrated *sunflower lemma* of Erdős and Rado asserts that, in every sufficiently large uniform family of sets, there exists a sub-family with a highly regular structure which we call "sunflower".

Definition 3.1.1. *A family $\mathcal{F}$ of subsets of Γ is called a* sunflower *if there exists a set Y such that $F_1 \cap F_2 = Y$ for every $F_1, F_2 \in \mathcal{F}$ such that $F_1 \neq F_2$. The sets $F \setminus Y$ for $F \in \mathcal{F}$ are called* petals *and the set $Y = \bigcap \mathcal{F}$ is called the* core.

The formal statement of the sunflower lemma of Erdős and Rado is as follows.

Theorem 3.1.2 (Erdős and Rado [ER60])**.** *Let $\mathcal{F}$ be a family of subsets of Γ, each of cardinality at most ℓ. If $|\mathcal{F}| > \ell!(r-1)^\ell$, then $\mathcal{F}$ contains a sunflower of r petals.*

We include here the proof of Theorem 3.1.2 for completeness. Before proceeding with the proof, we need a definition.

Definition 3.1.3. *For a family $\mathcal{F} \subseteq 2^\Gamma$ and a set $T \subseteq \Gamma$, the* link *of T at $\mathcal{F}$ is defined as*

$$\mathcal{F}_T := \{F \setminus T : F \in \mathcal{F}, T \subseteq F\}.$$

Remark 3.1.4. *We have $|\mathcal{F}_T| = d_{\mathcal{F}}(T)$.*

Proof of Theorem 3.1.2. The proof is by induction on ℓ. When $\ell = 1$, we have that $\mathcal{F}$ has r disjoint sets, which is a sunflower of r petals. Suppose then that $\ell \geqslant 2$. The proof is now divided in two cases.

Case 1: There exists an element $v \in \Gamma$ such that $d(v) > (\ell-1)!(r-1)^{\ell-1}$. In this case, we are done: by induction, $\mathcal{F}_{\{v\}}$ has a sunflower of r petals $\mathcal{F}'$; therefore, the collection $\{F \cup \{v\} : F \in \mathcal{F}'\}$ is a sunflower in $\mathcal{F}$.

Case 2: For all $v \in \Gamma$, we have $d(v) \leqslant (\ell-1)!(r-1)^{\ell-1}$. Let $\mathcal{M}$ be a maximal disjoint subfamily of $\mathcal{F}$ and let $S := \bigcup \mathcal{M}$. Since $\mathcal{M}$ is a sunflower, it suffices to show that $|\mathcal{M}| \geqslant r$. Observe first that $|S| \leqslant |\mathcal{M}| \cdot \ell$. Since $\mathcal{M}$ is maximal, we have moreover that every set of $\mathcal{F}$ intersects S. Therefore, by averaging there exists an element $v \in S$ such that $d(v) \geqslant |\mathcal{F}|/(|\mathcal{M}| \cdot \ell)$. This implies

$$\frac{\ell!(r-1)^\ell}{|\mathcal{M}| \cdot \ell} < d(v) \leqslant (\ell-1)!(r-1)^{\ell-1},$$

from which we get $|\mathcal{M}| > r-1$. This concludes the proof. □

In what has become one of the most important conjectures of extremal combinatorics, Erdős and Rado conjectured [ER60] that one can remove the $\ell!$ factor above, perhaps replacing $(r-1)^\ell$ by $\text{poly}(r)^\ell$.

Conjecture 3.1.5. *There exists a polynomial $c = c(r)$ such that the following holds. Let $\mathcal{F}$ be a family of subsets of Γ, each of cardinality at most ℓ. If $|\mathcal{F}| > c(r)^\ell$, then $\mathcal{F}$ contains a sunflower of r petals.*

A recent breakthrough of Alweiss, Lovett, Wu and Zhang [ALWZ19] showed that the statement above is true if we replace $c = c(r)$ by $c = c(r, \log \ell)$, where $c(r, \log \ell)$ is roughly $r \cdot \log \ell$. Eliminating this $\log \ell$ factor would thus solve the conjecture. We will show a simpler proof of their result due to Rao [Rao20] in Chapter 4, and we will show how their arguments can be used to improve the state-of-the-art monotone circuit lower bounds in Chapters 6 and 7. See Sections 3.3 and 3.4 for more details.

Theorem 3.1.2 was used by Razborov [Raz85b] to prove a lower bound of $n^{\Omega(k)}$ on the size of monotone circuits computing $\mathsf{Clique}(n, k)$ when $k \leqslant \log n$, by using his *method of approximations*. This shows a superpolynomial lower bound for the clique problem when $k = \log n$, but not yet exponential. Further progress in clique lower bounds was achieved by Alon and Boppana [AB87], who proved a lower bound of $n^{\Omega(\sqrt{k})}$ when $k \leqslant n^{2/3-o(1)}$. Though this lower bound is asymptotically worse for $k \leqslant \log n$, it provides an exponential lower bound for the clique problem when $k = n^{2/3-o(1)}$. Their lower bound also followed the method of approximations of Razborov, but made use of a different notion of sunflowers, which we introduce in the next section.[1]

3.2 The lopsided sunflower of Alon and Boppana

Alon and Boppana [AB87] developed another notion of sunflower, in the context of proving lower bounds for monotone circuits. It was also applied recently by Ramamoorthy and Rao [NRR18] to prove data structure lower bounds (it appears there with the name "flower"). This definition relaxes the notion of "core" of the standard sunflower of Erdős and Rado, and we call it *lopsided sunflower*. An almost identical notion appears in Füredi [Für80] and Andreev [And85] (See Jukna [Juk11, Chapter 6]).

Definition 3.2.1. *We say that $\mathcal{F} \subseteq 2^\Gamma$ is a* lopsided sunflower *if there exists a set $Y \subseteq \Gamma$ such that*

1. *$F_1 \cap F_2 \subseteq Y$ for every $F_1, F_2 \in \mathcal{F}$ such that $F_1 \neq F_2$;*
2. *there exists $F \in \mathcal{F}$ such that $Y \subsetneq F$.*

The set Y is called the core *and the sets $F \setminus Y$ for $F \in \mathcal{F}$ are called* petals.

Clearly, every sunflower of size at least 2 is a lopsided sunflower. In their article, Alon and Boppana also proved an upper bound on the size of set systems without a lopsided sunflower. We include their proof for completeness. First, we need a definition which shares some similarities with that of the *link* in the previous section.

Definition 3.2.2. *Let $\mathcal{F}$ be a family of subsets of Γ and fix sets C, D such that $C \subseteq D \subseteq \Gamma$. The* touch *of $\mathcal{F}$ at C with respect to D is defined as*

$$\mathcal{F}_C^{(D)} := \{F \setminus C : F \in \mathcal{F},\ F \cap D = C\}.$$

Theorem 3.2.3 (Alon and Boppana [AB87])**.** *Fix $\ell \geqslant 1$ and $r \geqslant 2$. Let $\mathcal{F}$ be a family of subsets of Γ, each of cardinality at most ℓ. If $|\mathcal{F}| > (r-1)^\ell$, then $\mathcal{F}$ contains a lopsided sunflower of size r.*

[1] It is worthy of note that, sometime later, a survey due to Boppana and Sipser [BS90] presented a lower bound of $n^{\Omega\sqrt{k}}$ for the $\mathsf{Clique}(n, k)$ problem when $k \leqslant n^{1/4-o(1)}$, making use of the Erdős-Rado sunflower lemma instead of the relaxed version of Alon and Boppana.

Proof. The proof is by induction on r. If $r = 2$, then $|\mathcal{F}| > (r-1)^\ell$ implies that $\mathcal{F}$ has at least two sets F_1, F_2. The set $Y := F_1 \cap F_2$ satisfies the definition of a core for the family $\{F_1, F_2\}$.

Now suppose the result holds for $r-1$. Fix any set $D \in \mathcal{F}$. We consider two cases, as in the proof of Theorem 3.1.2.

Case 1. There exists a set $C \subseteq D$ such that $\left|\mathcal{F}_C^{(D)}\right| > (r-2)^{\ell-|C|}$. By induction $\mathcal{F}_C^{(D)}$ has a lopsided sunflower $\mathcal{F}'$ of size $r-1$. It is not hard to check that $\{F \cup C : F \in \mathcal{F}'\} \cup \{D\} \subseteq \mathcal{F}$ is a lopsided sunflower of size r contained in $\mathcal{F}$.

Case 2. For all $C \subseteq D$, we have $\left|\mathcal{F}_C^{(D)}\right| \leqslant (r-2)^{\ell-|C|}$. Therefore, we get

$$|\mathcal{F}| \leqslant \sum_{C \subseteq D} \left|\mathcal{F}_C^{(D)}\right| \leqslant \sum_{C \subseteq D} (r-2)^{\ell-|C|} \leqslant \sum_{k=0}^{|D|} \binom{|D|}{k} (r-2)^{\ell-k} \leqslant \sum_{k=0}^{\ell} \binom{\ell}{k} (r-2)^{\ell-k} = (r-1)^\ell.$$

This is a contradiction which finishes the proof. □

Alon and Boppana also showed that this result is as good as possible. In particular, they showed, for every ℓ, the existence of ℓ-uniform set systems of size $(r-1)^\ell$ without any lopsided sunflower with r petals. Their construction is as follows. Let $S_1, \ldots, S_\ell$ be disjoint sets, and let $\mathcal{F}$ be the family of all sets $\{s_1, \ldots, s_\ell\}$ such that $s_i \in S_i$. Clearly, we have $|\mathcal{F}| = (r-1)^\ell$. It suffices now to show that $\mathcal{F}$ does not contain a lopsided sunflower of size r. Let $F_1, \ldots, F_r$ be any r sets from $\mathcal{F}$. Suppose there exists a set Y such that, for some $i \in \Gamma$, we have $Y \subsetneq F_i$. Then there exists $j \in [\ell]$ such that $Y \cap S_j = \emptyset$. Moreover, by the pigeonhole principle, there exists F_k and F_ℓ such that $F_k \cap F_\ell \cap S_j \neq \emptyset$. This implies that $F_k \cap F_\ell \not\subseteq Y$, which means that $F_1, \ldots, F_r$ is not a lopsided sunflower.

Fact 3.2.4 ([AB87])**.** *There exists an ℓ-uniform family $\mathcal{F}$ of size $(r-1)^\ell$ which does not contain a lopsided sunflower of size r. In other words, Theorem 3.2.3 is optimal.*

3.3 The robust sunflower of Rossman

The notion of *robust sunflowers* was introduced by Rossman [Ros14] in the context of proving better bounds for the average-case monotone complexity of $\mathsf{Clique}(n,k)$ on random graphs. Robust sunflowers found applications not only in monotone circuit complexity, but also in DNF sparsification [GMR13] randomness extractors [LLZ18], and lifting theorems [LMZ20, MP20]. We will see further in this section how a recent breakthrough in robust sunflower bounds [ALWZ19, Rao20] was used to improve the Erdős-Rado sunflower bounds, and we will see in Chapter 6 a result of the author and his coauthors which improves decades-old records in monotone circuit lower bounds by employing these new robust sunflower bounds.

Robust sunflowers were first called *quasi-sunflowers*, not only in [Ros14] but also in in [GMR13, LLZ18, LSZ19]. Later, they were called *approximate sunflowers* in [LZ19] and finally robust sunflowers in the breakthrough work of Alweiss, Lovett, Wu and Zhang [ALWZ19]. We will adopt this later terminology in this work. Before giving the formal definition of robust sunflowers, we first need an auxiliary definition.

Definition 3.3.1. *A family $\mathcal{F} \subseteq 2^\Gamma$ is said to be* (p,ε)-satisfying *if*

$$\Pr_{\mathbf{W} \subseteq_p \Gamma}[\mathbf{W} \text{ is independent in } \mathcal{F}] < \varepsilon.$$

Equivalently, it holds for the Boolean function $D_\mathcal{F} := \bigvee_{F \in \mathcal{F}} \mathbb{1}_F$ that

$$\Pr_{\mathbf{W} \subseteq_p \Gamma}[D_\mathcal{F}(\mathbf{W}) = 1] > 1 - \varepsilon.$$

Definition 3.3.2. *A family $\mathcal{F} \subseteq 2^\Gamma$ is called a* (p,ε)-robust sunflower *if the family of* petals $\{F \setminus Y : F \in \mathcal{F}\}$ *is* (p,ε)*-satisfying, where* $Y := \bigcap \mathcal{F}$ *is called the* core. *Equivalently, we have*

$$\Pr_{\mathbf{W}\subseteq_p \Gamma}[\mathbf{W} \cup Y \text{ is independent in } \mathcal{F}] < \varepsilon.$$

The significance of this definition can be best understood by comparing it to the standard sunflower of Erdős and Rado, which we do in the next section.

3.3.1 Connections to the sunflower of Erdős and Rado

Depending on the choice of parameters for p and ε, robust sunflowers can be seen either as a generalized or more restricted version of sunflowers. The following facts illustrate this observation.

Fact 3.3.3 (Remark 4.3 of [Ros14])**.** *Every ℓ-uniform sunflower of size r is a (p, e^{-rp^ℓ})-robust sunflower.*

Proof. Let $\mathcal{F} \subseteq \binom{\Gamma}{\ell}$ be a sunflower of size r with core Y. Let $\mathbf{W}$ be a p-random subset of Γ. We have

$$\begin{aligned}
\Pr_{\mathbf{W}\subseteq_p \Gamma}[\mathbf{W} \cup Y \text{ is independent in } \mathcal{F}] &= \Pr_{\mathbf{W}\subseteq_p \Gamma}[\forall F \in \mathcal{F} : F \setminus Y \not\subseteq \mathbf{W}] \\
&= \prod_{F\in\mathcal{F}} (1 - p^{|F|-|Y|}) \\
&\leqslant (1-p^\ell)^r \\
&\leqslant \exp\{-rp^\ell\}.
\end{aligned}$$

□

Fact 3.3.3 shows that, in a sense, robust sunflowers are a generalization of sunflowers [2]. Indeed, as a direct consequence of this fact, one obtains from Theorem 3.1.2 that every ℓ-uniform set system with size larger than $\ell!(\log(1/\varepsilon)/p^\ell)^\ell$ contains a (p,ε)-robust sunflower. It is possible to prove results that are better than a conversion from sunflower bounds, as we shall see in a theorem due to Rossman (Theorem 3.3.9) and its improvement due to Alweiss, Lovett, Wu and Zhang [ALWZ19].

At the same time, as observed by Lovett, Solomon and Zhang [LZ19], one cannot improve robust sunflower bounds indefinitely without also improving standard sunflower bounds. By showing that robust sunflowers with appropriate parameters contain a standard sunflower, the following fact shows that robust sunflowers are, in a sense, more "robust" than standard sunflowers – thus justifying the change of the name from "approximate" to "robust" sunflowers.

Fact 3.3.4 ([LZ19])**.** *Every $(1/r, 1/r)$-robust sunflower contains a sunflower of size r.*

Proof. Let $\mathcal{F} \subseteq 2^\Gamma$ be an $(1/r, 1/r)$-robust sunflower with core Y. Let $\mathbf{c} : \Gamma \to [r]$ be an r-coloring of Γ chosen uniformly at random. For every $i \in [r]$, let $\mathbf{W}_i := \mathbf{c}^{-1}(i)$, and observe that every $\mathbf{W}_i$ is a $(1/r)$-random subset of Γ. Therefore, for every $i \in [r]$ we have

$$\Pr[\mathbf{W}_i \cup Y \text{ is independent in } \mathcal{F}] < 1/r.$$

By the union bound, we get that the probability that $\mathbf{W}_i \cup Y$ is an independent set in $\mathcal{F}$ for all $i \in [r]$ is smaller than 1. This means that there exists a coloring $\mathbf{c}$ such that, for every $i \in [r]$, there exists a set $F_i \in \mathcal{F}$ such that $F_i \setminus Y \subseteq \mathbf{W}_i$. Since the sets $\mathbf{W}_i$ are disjoint (they form a partition of Γ) and $Y \subseteq F_1 \cap \cdots \cap F_r$, we get that the sets $F_1, F_2, \ldots, F_r \in \mathcal{F}$ form a sunflower with r petals and core Y. □

The facts above can be easily summarized in a chain of inequalities, by considering the following extremal numbers.

[2] This is why they were first called quasi-sunflowers and approximate sunflowers.

Definition 3.3.5. *Let $f_R(\ell, p, \varepsilon)$ be the positive integer such that any ℓ-uniform set family with size larger than $f_{\mathrm{R}}(\ell, p, \varepsilon)$ contains a (p, ε)-robust sunflower and let $\mathrm{ER}(\ell, r)$ be such that any ℓ-uniform set family with size larger than $\mathrm{ER}(\ell, r)$ contains a sunflower of size r.*

The facts above give us:

$$f_{\mathrm{R}}(\ell, p, e^{-rp^{\ell}}) \leqslant \mathrm{ER}(\ell, r) \leqslant f_{\mathrm{R}}(\ell, 1/r, 1/r). \tag{3.1}$$

Applying the Erdős-Rado sunflower lemma, the first inequality can be rewritten in terms of ε as follows:

$$f_{\mathrm{R}}(\ell, p, \varepsilon) \leqslant \mathrm{ER}(\ell, \log(1/\varepsilon)/p^{\ell}) \leqslant \ell!(\log(1/\varepsilon)/p^{\ell})^{\ell} \tag{3.2}$$

3.3.2 A bound due to Rossman

As we mentioned above, the first bound obtained for robust sunflowers came in a work due to Rossman [Ros14]. As we shall see, this bound considerably improves upon the bound given by inequality (3.2), which is crucial for its original application in monotone circuit lower bounds. Here we give its statement and full proof. The proof will follow the same structure of the Erdős-Rado sunflower lemma, with the difference that we will need Janson's inequality.

First, building upon the definition of a *link* of a family of sets (Definition 3.1.3) we will define "well-spread" sets.

Definition 3.3.6. *Let $r = (r_0, r_1, \ldots, r_{\ell-1})$ be a sequence of real numbers. We say that a ℓ-uniform family $\mathcal{F}$ is* r-well-spread *if $\Delta_t(\mathcal{F}) \leqslant r_{\ell-t}$ for all $t \in [\ell]$.*

In order to give Rossman's proof in its full strength, we will also need to define a sequence of polynomials $s_\ell(t)$. Rossman's proof will then give us $f_{\mathrm{R}}(\ell, p, \varepsilon) \leqslant s_\ell(\log(1/\varepsilon))/p^{\ell}$.

Definition 3.3.7. *Let $s_0(t), s_1(t), \ldots$ be the sequence of polynomials defined by*

$$s_0(t) := 1 \quad \textit{and} \quad s_\ell(t) := t \sum_{j=0}^{\ell-1} \binom{\ell}{j} s_j(t).$$

This definition and the following proposition are given explicitly in an unpublished note due to Rossman [Ros19], though a weaker form of the upper bound appears implicitly in [Ros14].

Proposition 3.3.8 ([Ros19])**.** *For all $t > 0$, we have*

$$\ell! t^{\ell} \leqslant s_\ell(t) \leqslant \ell!(t + 1/2)^{\ell}.$$

Proof. For the lower bound, note that

$$s_\ell(t) \geqslant t^{\ell} \prod_{j=1}^{\ell} \binom{j}{j-1} = t^{\ell} \ell!.$$

For the upper bound, we first prove by induction on ℓ that $s_\ell(t) \leqslant \ell!(\log(1/t+1))^{-\ell}$, as follows:

$$\begin{aligned} s_\ell(t) = t\sum_{j=0}^{\ell-1}\binom{\ell}{j}s_j(t) &\leqslant t\sum_{j=0}^{\ell-1}\binom{\ell}{j}j!(\log(1/t+1))^{-j} \\ &= t\ell!(\log(1/t+1))^{-\ell}\sum_{j=0}^{\ell-1}\frac{(\log(1/t+1))^{\ell-j}}{(\ell-j)!} \\ &\leqslant t\ell!(\log(1/t+1))^{-\ell}\left(-1+\sum_{j=0}^{\infty}\frac{(\log(1/t+1))^{j}}{j!}\right) \\ &= t\ell!(\log(1/t+1))^{-\ell}(-1+\exp(\log(1/t+1))) = \ell!(\log(1/t+1))^{-\ell}. \end{aligned}$$

To conclude the proof, we apply the inequality $1/\log(1/t+1) < t+1/2$ for all $t>0$. □

Theorem 3.3.9 ([Ros14]). *Let $\mathcal{F} \subseteq \binom{\Gamma}{\ell}$ be such that $|\mathcal{F}| > s_\ell(\log(1/\varepsilon))/p^\ell$. Then $\mathcal{F}$ contains a (p,ε)-robust sunflower. In other words, we have $f_{\mathrm{R}}(\ell,p,\varepsilon) \leqslant s_\ell(\log(1/\varepsilon))/p^\ell$.*

Proof. The proof is by induction on ℓ. If $\ell = 1$, then

$$\Pr[\mathbf{W} \text{ is independent in } \mathcal{F}] = (1-p)^{|\mathcal{F}|} \leqslant e^{-p|\mathcal{F}|} \leqslant \varepsilon.$$

Therefore, $\mathcal{F}$ itself is a (p,ε)-robust sunflower. We now suppose $\ell > 1$ and that the result holds for every $t \in [\ell-1]$. We consider the sequence $r = (r_0,\dots,r_{\ell-1})$ given by $r_j := s_j(\log(1/\varepsilon))/p^j$. As in the proof of Lemma 3.1.2, the proof is now divided in two cases.

Case 1. The family $\mathcal{F}$ is not r-well-spread. By definition, there exists an nonempty set $T \subseteq \Gamma$ such that $|\mathcal{F}_T| > s_{\ell-|T|}(\log(1/\varepsilon))/p^{\ell-|T|}$. By induction, the family $\mathcal{F}_T$ contains a (p,ε)-robust sunflower $\mathcal{F}'$. It is easy to see that $\{F \cup T : F \in \mathcal{F}'\}$ is a (p,ε)-robust sunflower in $\mathcal{F}$.

Case 2. The family $\mathcal{F}$ is r-well-spread. We will apply Janson's inequality (Lemma 2.2.4) to prove that $\mathcal{F}$ itself is (p,ε)-satisfying, thus finishing the proof. Let

$$\mu := \sum_{F\in\mathcal{F}}\Pr[F \subseteq \mathbf{W}], \qquad \Delta := \sum_{\substack{F,H\in\mathcal{F}\\ F\cap H\neq\emptyset}}\Pr[F\cup H \subseteq \mathbf{W}].$$

Janson's lemma now gives us

$$\Pr[\mathbf{W} \text{ is independent in } \mathcal{F}] \leqslant \exp\{-\mu^2/\Delta\}.$$

To show that $\mathcal{F}$ is (p,ε)-satisfying, it suffices to show that the right-hand side above is at most ε. This we will now do, by proving that $\mu^2/\Delta \geqslant \log(1/\varepsilon)$.

Observe first that $\mu = |\mathcal{F}|\,p^\ell > s_\ell(\log(1/\varepsilon))$. We will now bound Δ. Consider the following auxiliary parameter $\overline{\Delta}$, which ignores the diagonal terms of the sum defining Δ:

$$\overline{\Delta} := \sum_{\substack{F,H\in\mathcal{F}\\ F\neq H, F\cap H\neq\emptyset}}\Pr[F\cup H \subseteq \mathbf{W}].$$

We thus obtain $\Delta = \mu + \overline{\Delta}$. We will now bound $\overline{\Delta}$, as follows:

$$\overline{\Delta} = \sum_{\substack{F,H\in\mathcal{F}\\ F\neq H, F\cap H\neq\emptyset}} p^{2\ell-|F\cap H|} = \sum_{t=1}^{\ell-1}\sum_{T\in\binom{\Gamma}{t}}\sum_{\substack{F,H\in\mathcal{F}\\ F\cap H=T}} p^{2\ell-t} \leqslant \sum_{t=1}^{\ell-1} p^{2\ell-t}\sum_{T\in\binom{\Gamma}{t}}|\mathcal{F}_T|^2.$$

By double-counting, one gets that $\sum_{T\in\binom{\Gamma}{t}}|\mathcal{F}_T| = |\mathcal{F}|\binom{\ell}{t}$. Therefore, applying the bound on $|\mathcal{F}_T|$,

we obtain

$$\begin{aligned}\sum_{T\in\binom{\Gamma}{t}}|\mathcal{F}_T|^2 &\leqslant s_{\ell-t}(\log(1/\varepsilon))\cdot p^{t-\ell}\sum_{T\in\binom{\Gamma}{t}}|\mathcal{F}_T|\\ &\leqslant s_{\ell-t}(\log(1/\varepsilon))\cdot p^{t-\ell}\,|\mathcal{F}|\binom{\ell}{t}\\ &= \mu\cdot p^{t-2\ell}s_{\ell-t}(\log(1/\varepsilon))\binom{\ell}{t}.\end{aligned}$$

We may now continue to bound $\overline{\Delta}$, as follows:

$$\overline{\Delta}\leqslant\sum_{t=1}^{\ell-1}p^{2\ell-t}\sum_{T\in\binom{\Gamma}{t}}|\mathcal{F}_T|^2\leqslant\mu\sum_{t=1}^{\ell-1}s_{\ell-t}(\log(1/\varepsilon))\binom{\ell}{t}=\mu\left(\frac{s_\ell(\log(1/\varepsilon))}{\log(1/\varepsilon)}-1\right).$$

Therefore, we have $\Delta\leqslant\mu\frac{s_\ell(\log(1/\varepsilon))}{\log(1/\varepsilon)}$, whence we get

$$\frac{\mu^2}{\Delta}\geqslant\log(1/\varepsilon)\frac{\mu}{s_\ell(\log(1/\varepsilon))}\geqslant\log(1/\varepsilon),$$

as desired. □

By applying Proposition 3.3.8, we get the following corollary.

Corollary 3.3.10. *Let $\mathcal{F}\subseteq\binom{\Gamma}{\ell}$ be such that $|\mathcal{F}|>\ell!(2\log(1/\varepsilon)/p)^\ell$ and $\varepsilon\leqslant e^{-1/2}$. Then $\mathcal{F}$ contains a (p,ε)-robust sunflower. In other words, we have $f_{\mathrm{R}}(\ell,p,\varepsilon)\leqslant\ell!(2\log(1/\varepsilon)/p)^\ell$.*

Observe that this bound is a strong improvement upon inequality (3.2), which was directly obtained from the Erdős-Rado sunflower lemma. The lower bound of proposition 3.3.8 shows that this bound is almost optimal for this technique, thus suggesting that any substantial improvement over this bound can only come through a different technique. This was achived by Alweiss, Lovett, Wu and Zhang [ALWZ19], whose result we will discuss in the next section.

3.3.3 A stronger bound and consequences

The main result of the breakthrough work of Alweiss, Lovett, Wu and Zhang [ALWZ19] is the following theorem, which gave a strong improvement over the robust sunflower bound given by Rossman (Corollary 3.3.10).

Theorem 3.3.11 ([ALWZ19])**.** *Let $\mathcal{F}\subseteq\binom{\Gamma}{\ell}$ be such that $|\mathcal{F}|\geqslant(\log\ell)^\ell\cdot(\log\log\ell\cdot\log(1/\varepsilon)/p)^{O(\ell)}$. Then $\mathcal{F}$ contains a (p,ε)-robust sunflower. In other words, we have $f_{\mathrm{R}}(\ell,p,\varepsilon)\leqslant(\log\ell)^\ell\cdot(\log\log\ell\cdot\log(1/\varepsilon)/p)^{O(\ell)}$.*

One of the most important consequences of this result comes from applying inequality (3.1) to this theorem, which is a logarithmic factor short of proving the sunflower conjecture (Conjecture 3.1.5).

Corollary 3.3.12. *Let $\mathcal{F}\subseteq\binom{\Gamma}{\ell}$ be such that $|\mathcal{F}|\geqslant(\log\ell)^\ell\cdot(\log\log\ell\cdot r\log r)^{O(\ell)}$. Then $\mathcal{F}$ contains a sunflower of size r. In other words, we have $\mathrm{ER}(\ell,r)\leqslant(\log\ell)^\ell\cdot(\log\log\ell\cdot r\log r)^{O(\ell)}$.*

We remark that this bound is (roughly) $(\log\ell)^\ell\cdot r^{O(\ell)}$, which is a significant improvement upon the $\ell!r^\ell$ bound of Erdős and Rado (Theorem 3.1.2)[3]. The ideas behind the result were recently employed to prove a conjecture of Talagrand [FKNP19]. In Chapters 6 and 7, we will see how this improved bound leads to better monotone circuit lower bounds.

[3]Kostochka [Kos97] also proved a bound of the form $O(\ell!\cdot(\log\log\log\ell/\log\log\ell)^\ell)$ for 3-sunflowers, which is improved by this corollary.

The structure of the proof of [ALWZ19] also follows the case structure of Rossman's proof (Theorem 3.3.9). That is, two cases are considered, one when $\mathcal{F}$ is not r-well-spread, and the other when $\mathcal{F}$ is r-well-spread for some sequence r. The main thrust of Rossman's proof consisted into showing that r-well-spread families are (p,ε)-satisfying, when $r_j \approx (j \cdot \log(1/\varepsilon)/p)^j$, by applying Janson's inequality. The main contribution of [ALWZ19] is to show that r-well-spread families are (p,ε)-satisfying for a better choice of the sequence r. Their proof is not an application of Janson's inequality, but relies mainly on a clever encoding argument.

This case division of the proof is often called a "structure vs. pseudorandomness" approach: when $\mathcal{F}$ is not r-well-spread, there exists an nonempty set $T \subseteq \Gamma$ such that $d_{\mathcal{F}}(T)$ is large – this is the "structured case"; when $\mathcal{F}$ is r-well-spread, all nonempty sets $T \subseteq \Gamma$ have bounded degree – this is the "pseudorandom" case.

In this work, we will not see their proof, but we will present in Chapter 4 a simpler proof of Corollary 3.3.12 given by Anup Rao [Rao20], also based on a coding argument. This latter proof actually gives a slightly stronger statement, and makes use of yet another notion of sunflowers very similar to robust sunflowers, but with a slight tweak in the distribution of $\mathbf{W}$. This notion will be explained in Section 3.4. For now, we present an observation found in [ALWZ19], which shows that Theorem 3.3.11 is almost optimal.

3.3.4 A lower bound on the size of set families without robust sunflowers

In this section, we will show the construction of Alweiss, Lovett, Wu and Zhang [ALWZ19], which gives a lower bound to $f_{\mathrm{R}}(\ell,p,\varepsilon)$. This lower bound shows that Theorem 3.3.11 is almost tight. In particular, it shows that the technique of robust sunflowers cannot be employed to remove the $(\log \ell)^\ell$ factor in the upper bound to ER, and therefore another approach has to be found for the sunflower conjecture (Conjecture 3.1.5).

Proposition 3.3.13 ([ALWZ19]). *Let p, ε be positive real numbers such that $\varepsilon \leqslant 1/2$, $p \leqslant 1-\varepsilon$ and p is bounded away from 1.*[4] *Then there exists an ℓ-uniform family $\mathcal{F}$ such that*

$$|\mathcal{F}| \geqslant \left(\frac{\log \ell + \log(1/\varepsilon)}{p}\right)^{(1-o(1))\ell}$$

and $\mathcal{F}$ does not contain a (p,ε)-robust sunflower.

Proof. Let $\Gamma_1,\dots,\Gamma_\ell$ be disjoint sets, each of size m, where

$$m := \frac{\log(\ell/c)}{\log(1/(1-p))}$$

for some constant $c \geqslant 1$, and let Γ be their union. Let $\widehat{\mathcal{F}} \subseteq 2^\Gamma$ be the ℓ-uniform set family composed of all the sets $\{x_1,\dots,x_\ell\}$, where $x_i \in \Gamma_i$ for all $i \in [\ell]$. Note that $|\widehat{\mathcal{F}}| = m^\ell$. Let $\delta > 0$ be a parameter to be determined later, and let $\mathcal{F} \subseteq \widehat{\mathcal{F}}$ be a set system such that $|F \cap F'| \leqslant (1-\delta)\ell$ for every $F, F' \in \mathcal{F}$ with $F \neq F'$. The family $\mathcal{F}$ can be obtained by a greedy algorithm, each time choosing a set $F \in \widehat{\mathcal{F}}$ arbitrarily, and then removing all sets which intersect F in more than $(1-\delta)\ell$ elements. Since there are at most $\binom{\ell}{\delta\ell} m^{\delta\ell} \leqslant 2^\ell m^{\delta\ell}$ such sets, we obtain $|\mathcal{F}| \geqslant 2^{-\ell} m^{-\delta\ell} |\widehat{\mathcal{F}}| = 2^{-\ell} m^{(1-\delta)\ell}$.

We now proceed to show that $\mathcal{F}$ does not contain a (p,ε)-robust sunflower. Let $\mathbf{W} \subseteq_p \Gamma$. It suffices to show that $\mathcal{F}_Y$ is not (p,ε)-satisfying for any $Y \subseteq \Gamma$. Let us fix then some $Y \subseteq \Gamma$. Observe that the probability that a given set $F \in \mathcal{F}_Y$ is contained in $\mathbf{W}$ is at most $p^{\ell-|Y|} \leqslant 1-\varepsilon$, by assumption. Therefore, we can assume without loss of generality that $|\mathcal{F}_Y| \geqslant 2$. This implies that $|Y \cap \Gamma_i| \leqslant 1$ for every $i \in [\ell]$ and that $|Y| \leqslant (1-\delta)\ell$.

Let now $I := \{i \in [\ell] : |Y \cap \Gamma_i| = 0\}$. By our assumptions on Y, we get that $|I| \geqslant \delta\ell$. Therefore,

[4] In this result, the numbers p and ε are allowed to be a function of $n = |V|$.

we have

$$\begin{aligned}\Pr[\mathbf{W} \text{ is independent in } \mathcal{F}_Y] &\geqslant \Pr[\exists i \in I : \mathbf{W} \cap \Gamma_i = \emptyset] \\ &= 1 - (1 - (1-p)^m)^{|I|} \\ &= 1 - (1 - c/\ell)^{|I|} \\ &> 1 - (1 - c/\ell)^{\delta\ell},\end{aligned}$$

where the last equality comes from the choice of m. We now let $c := \log(1/(1-\varepsilon))/\delta$, and observe that

$$1 - (1 - c/\ell)^{\delta\ell} \geqslant 1 - e^{-c\delta} = \varepsilon.$$

We have thus proved that $\mathcal{F}$ does not contain a (p, ε)-sunflower. It remains to show that it has the desired size.

From our choice of m and c and from $|\mathcal{F}| \geqslant 2^{-\ell} m^{(1-\delta)\ell}$, we obtain

$$|\mathcal{F}| \geqslant 2^{-\ell} \cdot \left(\frac{\log\left(\delta\ell \cdot \frac{1}{\log(1/(1-\varepsilon))}\right)}{\log(1/(1-p))} \right)^{(1-\delta)\ell}.$$

To further lower bound the size of $\mathcal{F}$, we will need the following asymptotic estimates, which can be easily checked.

Claim 3.3.14. *The following holds.*

1. *If $x \in (0, 1/2]$, then there exists an universal constant t such that $1 - x \geqslant e^{-x^t}$.*
2. *If $x \in (0,1)$ is bounded away from 1, then there exists an universal constant T such that $1 - x \geqslant e^{-Tx}$.*

By our assumptions on p and ε and the claim, we get

$$\log\left(\frac{1}{\log(1/(1-\varepsilon))}\right) \geqslant t \cdot \log(1/\varepsilon),$$
$$\frac{1}{\log(1/(1-p))} \geqslant \frac{1}{Tp}.$$

Therefore, we have

$$|\mathcal{F}| \geqslant 2^{-\ell} \cdot \left(\frac{\log(\delta\ell) + t \cdot \log(1/\varepsilon)}{T \cdot p} \right)^{(1-\delta)\ell}.$$

Choosing $\delta = 1/\sqrt{\ell}$, we get the desired result:

$$|\mathcal{F}| \geqslant \left(\frac{\log \ell + \log(1/\varepsilon)}{p} \right)^{(1-o(1))\ell}.$$

□

Remark 3.3.15. *The size of the base set Γ used in the construction of Proposition 3.3.13 is*

$$|\Gamma| = \frac{3}{2} \cdot \frac{\ell \log(\ell/\log(1/(1-\varepsilon)))}{\log(1/(1-p))} = \Omega\left(\frac{\ell(\log \ell + \log(1/\varepsilon))}{p} \right).$$

This means that the lower bound only holds when $|\Gamma| = \Omega(\ell \log \ell)$. In particular, the lower bound does not hold when ℓ is of the same order as $|\Gamma|$.

3.4 The slice sunflower of Rao

In this section, we consider a "slice" version of the p-biased definition of satisfying hypergraphs.

Definition 3.4.1. *A hypergraph $\mathcal{F}$ with vertex set Γ is said to be* (M, ε)-satisfying *if*

$$\Pr_{\mathbf{W} \subseteq_M \Gamma} [\mathbf{W} \text{ is independent in } \mathcal{F}] < \varepsilon.$$

Definition 3.4.2. *Let $\mathcal{F}$ be a family of subsets of Γ and let $Y := \bigcap \mathcal{F}$. The family $\mathcal{F}$ is called a* (M, ε)-slice-sunflower *if the family of* petals $\{F \setminus Y : F \in \mathcal{F}\}$ *is (M, ε)-satisfying.*

The above definition, as well as the following theorem, are implicit in the work of Rao [Rao20].

Theorem 3.4.3 ([Rao20])**.** *There exists a universal constant $B > 0$ such that the following holds. Let $\varepsilon, p \in (0, 1)$ and let $\mathcal{F} \subseteq \binom{\Gamma}{\ell}$ be such that $|\mathcal{F}| \geqslant (Bx \log x)^\ell$, where $x = \log(\ell/\varepsilon)/p$. Then $\mathcal{F}$ contains a (M, ε)-slice-sunflower, where $M = \lfloor np \rfloor$ and $n = |\Gamma|$.*

The theorem above gives an upper bound on the size of uniform set families without slice sunflowers. Its proof was inspired by the robust sunflower bound of [ALWZ19] (Theorem 3.3.11), and it was used to give an alternative upper bound to $\mathrm{ER}(\ell, r)$ which replaces the $\ell!$ factor of Erdős and Rado by a $(\log \ell)^\ell$ factor. Its main benefits are that the proof is simpler and that the bound given is asymptotically better for many applications. In particular, the exponent in Theorem 3.4.3 is ℓ, whereas the exponent in Theorem 3.3.11 is $O(\ell)$. The better exponent of Rao is crucial for our applications in monotone circuit complexity in Chapters 6 and 7.

We defer the proof of Theorem 3.4.3 to Chapter 4. For now, we observe that the bound for $\mathrm{ER}(\ell, r)$ given by Theorem 3.4.3 which we just mentioned follows from the following fact.

Fact 3.4.4 ([Rao20])**.** *Every $(\lfloor n/r \rfloor, 1/r)$-slice sunflower contains a sunflower of size r.*

To prove the fact, we proceed in the same way as in Fact 3.3.4. The only difference is that the random partition $\mathbf{W}_1, \ldots, \mathbf{W}_r$ of the base set is now chosen to be such that each $\mathbf{W}_i$ is of size *exactly* $\lfloor n/r \rfloor$. Rao's theorem combined with the fact gives then gives us the following corollary, which is Rao's "version" of Corollary 3.3.12.

Corollary 3.4.5. *There exists a universal constant $B > 0$ such that the following holds. Let $\mathcal{F} \subseteq \binom{\Gamma}{\ell}$ be such that $|\mathcal{F}| \geqslant (B \cdot r \log(r\ell) \cdot \log(r \log(r\ell)))^\ell$. Then $\mathcal{F}$ contains a sunflower of size r. In other words, we have $\mathrm{ER}(\ell, r) \leqslant (B \cdot r \log(r\ell) \cdot \log(r \log(r\ell)))^\ell$.*

3.5 Other notions and open questions

3.5.1 Daisies

A *daisy* is a relaxed notion of sunflower which was introduced in the setting of proving lower bounds for the blocklength of relaxed locally decodable codes [GL20]. A similar notion first appeared in the context of *property testing*, with the names *pompons* and *constellations* [FLV15].

Definition 3.5.1. *A family of sets $\mathcal{F} \subseteq 2^\Gamma$ is called a* t-daisy *if there exists a set $Y \subseteq \Gamma$, called* core*, such that, for every element $v \in \Gamma \setminus Y$, there are at most t sets $F \in \mathcal{F}$ such that v is contained in the* petal *$F \setminus Y$.*

Similar to the lopsided sunflower of Alon and Boppana (Section 3.2), a t-daisy is a more generalized form of sunflower, which relaxes the definition of the core. In particular, it is easy to see that any sunflower is also a t-daisy for all $t \geqslant 0$. Importantly, if one omits the requirement that the core of a lopsided sunflower must be properly contained in one of the sets of the family, one could also see t-daisies as a weaker form of lopsided sunflowers. In fact, if a set Y satisfies $F \cap H \subseteq Y$ for all $F, H \in \mathcal{F}$, then $\mathcal{F}$ is a 1-daisy with core Y.

Together with the introduction of this concept, Gur and Lachish also proved a corresponding "daisy lemma", with a somewhat different flavour from the sunflower lemmas we have considered so far.

Theorem 3.5.2 (Gur and Lachish [GL20])**.** *Let $\mathcal{F} \subseteq \binom{\Gamma}{\ell}$ be such that $|\mathcal{F}| = cn$ and let μ be a distribution over 2^Γ such that $\operatorname{supp}(\mu) = \mathcal{F}$. Then, for some $s \in [\ell]$ and $m = \max\{1, s-1\}$, there exists a $cn^{m/\ell}$-daisy $\mathcal{S} \subseteq \mathcal{F}$ with a core of size at most $\ell \cdot n^{1-s/\ell}$ and petals of size at most s, such that $\mu(\mathcal{S}) \geqslant 1/\ell$.*

It is natural to ask whether one can obtain a result with the same flavour as that of the sunflower lemma of Erdős and Rado for the concept of daisies. Let $f_{\mathrm{D}}(t, \ell, r)$ be the maximum size of an ℓ-uniform set family that does not contain a t-daisy with r petals.

Problem 3.5.3. *Give an upper bound to $f_{\mathrm{D}}(t, \ell, r)$.*

An interesting observation is that the set of petals of a t-daisy is (p, ε)-satisfying. As in Rossman's original proof of the robust sunflower lemma (Theorem 3.3.9), the proof follows from an application of Janson's inequality (Lemma 2.2.4).

Proposition 3.5.4. *Let $\mathcal{F}$ be a t-daisy such that all of its petals have size s. Suppose moreover that $|\mathcal{F}| \geqslant \log(1/\varepsilon)st/p^s$. Then the set family consisting of the petals of $\mathcal{F}$ is (p, ε)-satisfying.*

Proof. Let $\mathcal{F}$ be a t-daisy of with core Y and vertex set Γ. Let $\mathbf{W} \subseteq_p \Gamma$. For $F \in \mathcal{F}$, let $F' := F \setminus Y$ and let $\mathcal{F}' := \{F' : F \in \mathcal{F}\}$ be the set family consisting of the petals of $\mathcal{F}$. Let

$$\mu := \sum_{F \in \mathcal{F}} \Pr[F' \subseteq \mathbf{W}],$$
$$\Delta := \sum_{\substack{F, H \in \mathcal{F} \\ F \cap H \neq \emptyset}} \Pr[F' \cup H' \subseteq \mathbf{W}].$$

We have $\mu = |\mathcal{F}|\, p^s$. Moreover, we have

$$\Delta = \mu + \sum_{F \in \mathcal{F}} \sum_{\substack{F, H \in \mathcal{F} \\ F \neq H, F \cap H \neq \emptyset}} \Pr[F' \cup H' \subseteq \mathbf{W}] \leqslant \mu + \sum_{F \in \mathcal{F}} s(t-1)p^s \leqslant \mu st.$$

Therefore, by Janson's inequality (Lemma 2.2.4) we have

$$\Pr[\mathbf{W} \text{ is independent in } \mathcal{F}'] < \exp(-\mu^2/\Delta) \leqslant \varepsilon. \qquad \square$$

3.5.2 Robust lopsided versions

In many applications of sunflower theorems in monotone circuit lower bounds, the only requirement the core has to satisfy is that it is properly contained in one of the subsets. It is therefore natural to consider the following definition, which relaxes even further the requirements for the core of the robust sunflower.

Definition 3.5.5. *We say that $\mathcal{F} \subseteq 2^\Gamma$ is a (p, ε)-robust lopsided sunflower if there exists a set $Y \subseteq \Gamma$ (called* core*) such that $Y \subsetneq F$ for some $F \in \mathcal{F}$ and the family of the* petals *$\{F \setminus Y : F \in \mathcal{F}\}$ is (p, ε)-satisfying.*

We now make some remarks on the particularity of this notion of sunflowers. Since we make no requirements on the core of a robust lopsided sunflower other than it being properly contained in one of the sets of the family, it is easy to see that, if a set family $\mathcal{F}$ contains a robust lopsided sunflower, than it is itself a robust lopsided sunflower. In other words, being a robust lopsided sunflower is a monotone property.

Furthermore, one readily sees that, without loss of generality, we can suppose that the core Y is of the form $F \setminus \{u\}$ for some $F \in \mathcal{F}$ and $u \in F$, since Y is properly contained in one of the sets of the family, and enlarging the core can only enlarge the probability of $\mathbf{W}$ containing a petal of the family.

We will now see that robust lopsided sunflowers bear the same relationship with lopsided sunflowers that robust sunflowers do with sunflowers. First, we see that every lopsided sunflower can be seen as a robust lopsided sunflower. The "conversion parameters" are the same as that of Fact 3.3.3.

Fact 3.5.6. *Every ℓ-uniform lopsided sunflower of size r is a (p, e^{-rp^ℓ})-robust lopsided sunflower.*

Proof. Observing that the sets $F \setminus Y$ for $F \in \mathcal{F}$ are mutually disjoint, we can apply the same argument of Fact 3.3.3. □

Now, we show that every robust lopsided sunflower is a lopsided sunflower. There is a slight difference in the "conversion parameters" here, but they are morally the same as those of Fact 3.3.4. This unpublished fact was observed in collaboration with Benjamin Rossman and Mrinal Kumar.

Fact 3.5.7. *Let $\mathcal{F}$ be a $(1/r, (r-1)/r^2)$-robust lopsided sunflower. Then $\mathcal{F}$ contains a lopsided sunflower of size $r+1$.*

Proof. For simplicity, we let $\varepsilon := (r-1)/r^2$ and $p := 1/r$ and we suppose $\mathbf{W} \subseteq_p \Gamma$. Since $\mathcal{F}$ is a (p, ε)-robust lopsided sunflower, there exists $F_0 \in \mathcal{F}$ and $Y \subsetneq F_0$ such that

$$\Pr\left[\mathbf{W} \cup Y \text{ is independent in } \mathcal{F}\right] < \varepsilon.$$

Without loss of generality, we assume that there exists $u \in F_0$ such that $Y = F_0 \setminus \{u\}$. Let $\mathcal{F}' := \{F \in \mathcal{F} : u \notin F\}$, $V' := V \setminus \{u\}$ and $\mathbf{W}' \subseteq_p V'$. Observe that, for every $F \in \mathcal{F} \setminus \mathcal{F}'$, if $u \notin \mathbf{W}$ then $F \not\subseteq \mathbf{W} \cup Y$. Therefore, we have

$$\begin{aligned}
\varepsilon &> \Pr[\mathbf{W} \cup Y \text{ is independent in } \mathcal{F}] \\
&\geqslant \Pr\left[u \notin \mathbf{W} \text{ and } \mathbf{W} \cup Y \text{ is independent in } \mathcal{F}'\right] \\
&= (1-p)\Pr\left[\mathbf{W} \cup Y \text{ is independent in } \mathcal{F}'\right].
\end{aligned}$$

It follows that

$$\Pr\left[\mathbf{W} \cup Y \text{ is independent in } \mathcal{F}'\right] < \frac{\varepsilon}{1-p} = 1/r.$$

Let $\mathbf{W}'_1, \mathbf{W}'_2, \ldots, \mathbf{W}'_r$ be a uniform random partition of V'. We have $\mathbf{W}'_i \subseteq_p V'$ for every $i \in [r]$. By an union bound, we obtain that the probability that there exists $i \in [r]$ such that $\mathbf{W}'_i \cup Y$ is independent in $\mathcal{F}'$ is at most 1. Therefore, there exists a partition $V'_1 \dot\cup \ldots \dot\cup V_r$ of V' such that, for every $i \in [r]$, there exists $F_i \in \mathcal{F}'$ satisfying $F_i \setminus Y \subseteq V'_i$. This implies that $F_i \cap F_j \subseteq Y$ for every $i, j \in [r]$ with $i \neq j$. Moreover, we also have $F_0 \cap F_i \subseteq Y$ for every $i \in [k]$, since $u \notin F_i$. It follows that $F_0, F_1, F_2, \ldots, F_r$ is a lopsided-sunflower of size $k+1$. □

As before, let $f_{\mathrm{RL}}(\ell, p, \varepsilon)$ be the maximum size of an ℓ-uniform set family that does not contain a a (p, ε)-robust lopsided sunflower. Note that, because of the observation about monotonicity above, any ℓ-uniform set family with size bigger than $f_{\mathrm{RL}}(\ell, p, \varepsilon)$ is itself a (p, ε)-robust lopsided. Applying Fact 3.5.6 and the lopsided sunflower bound (Theorem 3.2.3), we obtain $f_{\mathrm{RL}}(\ell, p, \varepsilon) \leqslant (\log(1/\varepsilon)/p^\ell)^\ell$. Furthermore, since any robust sunflower is a robust lopsided sunflower, we may apply the robust sunflower bound of [ALWZ19] (Theorem 3.3.11) to obtain $f_{\mathrm{RL}}(\ell, p, \varepsilon) \leqslant (\log\log\ell \cdot \log(1/\varepsilon)/p)^{O(\ell)}$. Combining these two inequalities, we get

$$f_{\mathrm{RL}}(\ell, p, \varepsilon) \leqslant \min\left\{(\log(1/\varepsilon)/p^\ell)^\ell, (\log\log\ell \cdot \log(1/\varepsilon)/p)^{O(\ell)}\right\}. \tag{3.3}$$

On the other hand, Fact 3.5.7 gives us a lower bound of $(1/p-1)^\ell$ for $f_{\mathrm{RL}}(\ell, p, \varepsilon)$. It is natural to ask for tighter bounds for this extremal function.

Problem 3.5.8. *Prove better bounds for $f_{\mathrm{RL}}(\ell, p, \varepsilon)$.*

It would also be interesting to see if this notion also has applications in computational complexity.

Problem 3.5.9. *Find applications of robust lopsided sunflowers in computational complexity, in particular monotone circuit lower bounds. Can strong upper bounds for* $f_{\mathrm{RL}}(\ell, p, \varepsilon)$ *imply improved monotone circuit lower bounds?*

A more elaborate discussion on this problem is given in Section 6.6.1.

Finally, given that being a robust lopsided sunflower is a monotone property in $2^{2^{[n]}}$, and that every monotone property has a threshold function (Theorem 2.2.5), it would be interesting to find a threshold function for this property. Perhaps it would help us to understand better the behaviour of random monotone Boolean functions. Such functions are interesting because they are known to have high circuit complexity [5].

More formally, the problem statement is as follows. Let $\mathcal{P}(n, p)$ be the p-random sublattice of $2^{[n]}$. Let $\hat{p}_{q,\varepsilon}$ be such that, if $p \gg \hat{p}_{q,\varepsilon}$, then $\mathcal{P}(n, p)$ is a robust lopsided sunflower almost surely, and if $p \ll \hat{p}_{q,\varepsilon}$, then $\mathcal{P}(n, p)$ is *not* a robust lopsided sunflower almost surely.

Problem 3.5.10. *Calculate* $\hat{p}_{q,\varepsilon}$ *for all* q *and* ε.

[5] See, for instance, Section 1.4 of Jukna [Juk12]

Chapter 4

A breakthrough in sunflower theorems

In this chapter, we will reproduce the proof of Theorem 3.4.3 given in Rao [Rao20], whose statement we now recall.

Theorem 3.4.3 ([Rao20]). *There exists a universal constant $B > 0$ such that the following holds. Let $\varepsilon, p \in (0,1)$ and let $\mathcal{F} \subseteq \binom{\Gamma}{\ell}$ be such that $|\mathcal{F}| \geqslant (Bx \log x)^{\ell}$, where $x = \log(\ell/\varepsilon)/p$. Then $\mathcal{F}$ contains a (M,ε)-slice-sunflower, where $M = \lfloor np \rfloor$ and $n = |\Gamma|$.*

The proof will be based on an elegant encoding argument inspired by the work of Alweiss, Lovett, Wu and Zhang [ALWZ19]. We give in the following section a simplified version of Rao's argument, due to Rao himself [Rao19], and go straight to the proof of Theorem 3.4.3 in the remaining sections.

4.1 Warm-up: an even weaker "sunflower"

To introduce and motivate some of the ideas that are going to be used in the proof of Theorem 3.4.3, we will consider in this section an even weaker notion of sunflower. Here, instead of asking the random set $\mathbf{W}$ to contain an edge of the hypergraph with high probability (i.e., $\mathbf{W}$ is not independent), we will require it to *almost* contain such an edge. In this simplified scenario, some of the main ideas will be clearer. This section follows a short and elegant proof exposed in a talk of Anup Rao [Rao19], which was in turn inspired by an encoding argument which appears in Lemma 2.7 of [ALWZ19]. The reader who wants to rush to the proof of Theorem 3.4.3 may skip this section.

Given a hypergraph $\mathcal{H}$ and a number $t \geqslant 0$, we say that a set $S \subseteq V(\mathcal{H})$ is *t-independent* if $|F \setminus S| > t$ for all $F \in E(\mathcal{F})$. Observe that a set is independent in $\mathcal{F}$ if and only if it is 0-independent. This motivates the following definition, which generalizes Definitions 3.4.1 and 3.4.2.

Definition 4.1.1. *We say that a hypergraph $\mathcal{H}$ with vertex set Γ is* (M,ε,t)-almost satisfying *if*

$$\Pr_{\mathbf{W} \subseteq_M \Gamma}[\mathbf{W} \text{ is } t\text{-independent in } \mathcal{H}] < \varepsilon.$$

Observe that a (M,ε)-satisfying hypergraph is (M,ε,t)-almost satisfying for all $t \geqslant 0$. A hypergraph is (M,ε)-satisfying if, with probability larger than $1-\varepsilon$, there exists a set $F \in E(\mathcal{F})$ with $F \subseteq \mathbf{W}$. This is the same as $|F \setminus \mathbf{W}| = 0$. The role of t is then to bound how far the family $\mathcal{F}$ is from being (M,ε)-satisfying.

As we have seen in Chapter 3 (See, for instance, the discussion in Section 3.3.3), in the "structure vs. randomness" approach to sunflower bounds the main step is to show that well spread families are satisfying. The weaker we can make the assumption of well-spreadedness, the stronger is the sunflower bound. We will now show that well-spread families are almost satisfying. The coding argument employed in this proof is a simplified version of the coding argument that will be employed in the proof of Theorem 3.4.3, where we will show that well-spread families are not only *almost* satisfying, but actually (M,ε)-satisfying.

Proposition 4.1.2 ([ALWZ19, Rao19])**.** *Let $t \geqslant 0$ and $\varepsilon, p \in (0,1)$. Let $\mathcal{H}$ be a ℓ-uniform hypergraph with vertex set Γ, and suppose that $\mathcal{H}$ is a $(1, r, r^2, \ldots, r^\ell)$-well-spread family, where $r = (5/\varepsilon \cdot (4/p)^\ell)^{1/t}$. Suppose moreover that $|E(\mathcal{H})| \geqslant r^\ell$. Then $\mathcal{H}$ is (M, ε, t)-almost satisfying, where $M = \lfloor np \rfloor$ and $n = |\Gamma|$.*

Proof. Throughout this proof, fix an ordering $F_1, \ldots, F_k$ of the edges of $\mathcal{H}$. A pair (W, F) where $W \in \binom{\Gamma}{M}$ is t-independent in $\mathcal{H}$ and F is an edge of $\mathcal{H}$ will henceforth be called a *bad pair*. We will give an upper bound on the number of bad pairs via an encoding argument, as follows. We will assign to each such pair a different binary string with s bits. Since there are 2^s binary strings with s bits and this assignment is injective, an upper bound of 2^s on the number of bad pairs will ensue.

1. First, we encode $W \cup F$. Since $|W| = M$ and $|F| = \ell$, there are at most
$$\binom{n}{M} + \cdots + \binom{n}{M+\ell} \leqslant \binom{n+\ell}{M+\ell} \leqslant \binom{n}{M} \cdot \left(\frac{n}{M}\right)^\ell = \binom{n}{M} \cdot (1/p)^\ell$$
choices for $W \cup F$. We can therefere encode $W \cup F$ with $\log_2\left(\binom{n}{M}\right) + \ell \cdot \log_2(1/p) + 1$ bits.

2. Let F_j be the first edge of the subhypergraph $\mathcal{H}[W \cup F]$ of $\mathcal{H}$ induced by $W \cup F$ (i.e., $F_j \subseteq F \cup W$). Such an F_j always exists because $F \in E(\mathcal{H}[W \cup F])$. We then encode $F \cap F_j$, which can be done with ℓ bits since $\mathcal{H}$ is ℓ-uniform.

3. The crucial observation is that, since $F_j \setminus W \subseteq F$ and W is t-independent, we get that $|F \cap F_j| > t$. Since $\mathcal{H}$ is $(1, r, r^2, \ldots, r^\ell)$-well-spread, we get that there are at most $r^{\ell - t}$ choices for F. We can then encode F with $(\ell - t) \cdot \log_2 r + 1$ bits.

4. We now encode $W \cap F$, which can be done with ℓ bits. The set W can now be easily recovered from $W \cup F$, F, and $W \cap F$.

The fact that this encoding is indeed injective can be seen by observing that the above encoding is constructed in such a way as to give a natural decoding algorithm which recovers the pair (W, F) from the information given by the string. Further, the encoding spends
$$\log_2 \binom{n}{M} + \ell \cdot \log_2(1/p) + 2\ell + (\ell - t)\log_2 r + 2$$
bits, which implies that there are at most $\binom{n}{M} \cdot 4 \cdot (4/p)^\ell \cdot r^{\ell - t}$ bad pairs. By double-counting the number of bad pairs, we then get
$$\Pr_{\mathbf{W} \subseteq_M \Gamma}[\mathbf{W} \text{ is } t\text{-independent in } \mathcal{H}] \leqslant \frac{4 \cdot (4/p)^\ell}{r^t} = \frac{4}{5} \cdot \varepsilon < \varepsilon. \qquad \square$$

4.2 Counting with codes

4.2.1 Prefix-free codes

In this section, we recall a few basic facts about prefix-free encodings. Since the proofs are short, we give full proofs for completeness.

An encoding $C : S \to \{0,1\}^*$ of a set S is called *prefix-free* if there does not exist $x, y \in S$ such that $x \neq y$ and $C(x)$ is a prefix of $C(y)$. For $x \in \{0,1\}^*$, we denote by $\mathsf{length}(x)$ the *length* of x (i.e., $\mathsf{length}(x) = k \iff x \in \{0,1\}^k$).

Lemma 4.2.1 (Kraft [Kra49])**.** *Let S be a finite set and let $C : S \to \{0,1\}^*$ be a prefix-free encoding of S into binary strings. We have*
$$\sum_{s \in S} 2^{-\mathsf{length}(C(s))} \leqslant 1.$$

Proof. Let $\mathbf{x}$ be a uniformly random binary string, with length longer than any of the strings of the encoding. Since the encoding is prefix-free, no two strings of the encoding can be a prefix of $\mathbf{x}$ at the same time. This implies

$$1 \geqslant \Pr_{\mathbf{x}}[\exists s \in S : C(s) \text{ is a prefix of } \mathbf{x}] = \sum_{s \in S} 2^{-\mathsf{length}(C(s))}. \qquad \square$$

Corollary 4.2.2. *Let S be a finite set and let $C : S \to \{0,1\}^*$ be a prefix-free encoding of S into binary strings. Let also $\mathbf{x} \sim \mathcal{U}(S)$. We have*

$$|S| \leqslant 2^{\mathbb{E}_{\mathbf{x}}[\mathsf{length}(C(\mathbf{x}))]}.$$

Proof. From the convexity of the exponential function (using Jensen's inequality: Proposition 2.2.3) and Kraft's inequality (Lemma 4.2.1), we get

$$2^{-\mathbb{E}_{\mathbf{x}}[\mathsf{length}(C(\mathbf{x}))]} \leqslant \mathop{\mathbb{E}}_{\mathbf{x}}[2^{-\mathsf{length}(C(\mathbf{x}))}] \leqslant |S|^{-1}. \qquad \square$$

We remark that Corollary 4.2.2 is a particular case of a more general result due to Shannon which states that, under any distribution, the average length of a prefix-free encoding is at least the entropy of the distribution. Corollary 4.2.2 is then a particular instance of this fact for the uniform distribution.

4.2.2 Measuring independence

As in the proof of Proposition 4.1.2, we will fix an ordering $F_1, F_2, \ldots, F_k$ of the edges of $\mathcal{F}$. For a set $W \subseteq \Gamma = V(\mathcal{F})$, we denote by $\chi(F_i, W)$ the set $F_j \setminus W$, where $j \in [k]$ is the first index minimizing $|F_j \setminus W|$ subject to $F_j \subseteq F_i \cup W$. The key observation here is that, for any $F \in \mathcal{F}$, we have $\chi(F, W) = \emptyset$ if and only if W is not independent in $\mathcal{F}$.

The following lemma is essential to prove Theorem 3.4.3. Its proof follows an encoding argument similar to that of Proposition 4.1.2.

Lemma 4.2.3 ([Rao20])**.** *The following holds for every non-negative integer s. Let $\mathcal{F}$ be an ℓ-uniform hypergraph with vertex set Γ, and suppose that $\mathcal{F}$ is a $(1, r, r^2, \ldots, r^\ell)$-well-spread family for some $r > 0$. Suppose moreover that $|\mathcal{F}| \geqslant r^\ell$. Let $\mathbf{F} \sim \mathcal{U}(\mathcal{F})$ and $\mathbf{X} \subseteq_K \Gamma$ be sampled independently, where $K = s \cdot 128 \cdot \lceil n/r \rceil$. We have*

$$\mathop{\mathbb{E}}_{\mathbf{X},\mathbf{F}}[|\chi(\mathbf{F}, \mathbf{X})|] \leqslant \ell \cdot (1 - 1/\log_2 r)^s.$$

Proof. The proof is by induction on s. The result is trivial when $s = 0$. Suppose then that $s > 0$. Let $y := (s-1) \cdot 128 \cdot \lceil n/r \rceil$ and $z := K - y = 128 \cdot \lceil n/r \rceil$. Let $\mathbf{Y} \subseteq_y \Gamma$ and $\mathbf{Z} \subseteq_z \Gamma$ be random disjoint sets. Clearly, we can sample $\mathbf{X}$ as $\mathbf{X} = \mathbf{Y} \cup \mathbf{Z}$. In what follows, we will show that, for every fixed set $Y \in \binom{\Gamma}{y}$, we have

$$\mathop{\mathbb{E}}_{\mathbf{Z},\mathbf{F}}[|\chi(\mathbf{F}, Y \cup \mathbf{Z})|] \leqslant \mathop{\mathbb{E}}_{\mathbf{F}}[|\chi(\mathbf{F}, Y)|] \cdot (1 - 1/\log_2 r). \tag{4.1}$$

By taking expectations over $\mathbf{Y}$ on both sides and by the induction hypothesis, we will get

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathbf{Y},\mathbf{Z},\mathbf{F}}[|\chi(\mathbf{F}, \mathbf{Y} \cup \mathbf{Z})|] &\leqslant \mathop{\mathbb{E}}_{\mathbf{F},\mathbf{Y}}[|\chi(\mathbf{F}, \mathbf{Y})|] \cdot (1 - 1/\log_2 r) \\ &\leqslant \ell \cdot (1 - 1/\log_2 r)^s,\end{aligned}$$

thus finishing the proof.

Let us then fix such a Y. Observe now that, if $|\chi(F, Y)| = 0$ for some $F \in \mathcal{F}$, then, as observed above, we have $\mathbb{E}_{\mathbf{F}}[|\chi(\mathbf{F}, Y)|] = 0$, which implies $\mathbb{E}_{\mathbf{F},\mathbf{Z}}[|\chi(\mathbf{F}, Y \cup \mathbf{Z})|] = 0$. Thus we may henceforth suppose that $\mathbb{E}_{\mathbf{F}}[|\chi(\mathbf{F}, Y)|] \geqslant 1$.

We will now count the number of pairs (Z, F), where Z is in the support of $\mathbf{Z}$ when $\mathbf{Y} = Y$ (i.e.: $Z \cap Y = \emptyset$ and $Z \in \binom{\Gamma}{z}$), and $F \in E(\mathcal{F})$. Clearly, there at least $r^\ell \cdot \binom{n-y}{z}$ such pairs. We will now upper bound the number of such pairs by giving a prefix-free encoding of (Z, F).[1] The encoding bears some similarities to the encoding of Proposition 4.1.2, and it is as follows.

1. Encode $|\chi(F, Y)|$ as $0^{|\chi(F,Y)|}1$. This requires $|\chi(F, Y)| + 1$ bits.
2. Encode $X \cup \chi(F, Y)$, where $X = Y \cup Z$. Because Y is fixed, there are at most
$$\binom{n-y}{z} + \cdots + \binom{n-y}{z+|\chi(F,Y)|} \leqslant \binom{n-y+|\chi(Y,F)|}{z+|\chi(F,Y)|} \leqslant \binom{n-y}{z} \cdot \left(\frac{n}{z}\right)^{|\chi(F,Y)|}$$
choices for $X \cup \chi(F, Y)$. Thus we can encode $X \cup \chi(F, Y)$ with $\log \binom{n-y}{z} + |\chi(F, Y)| \log_2(n/z) + 1$ bits.
3. Let F_j be the first edge of $\mathcal{F}$ which minimizes $|\chi(F_j, Y)|$ subject to $\chi(F_j, Y) \subseteq X \cup \chi(F, Y)$. Such an edge always exists because F is a candidate. We encode $\chi(F_j, Y) \cap \chi(F, Y)$ with $|\chi(F, Y)|$ bits.
4. We now observe that $\chi(F_j, Y) \cap \chi(F, Y) \subseteq F$. Since $\mathcal{F}$ is $(1, r, \ldots, r^\ell)$-well-spread, there are at most $r^{\ell-t}$ choices for F, where $t = |\chi(F_j, Y) \cap \chi(F, Y)|$. We can then encode F with $\log_2(r^{\ell-t}) + 1$ bits. Moreover, we claim that $t \geqslant |\chi(F, X)| = |\chi(F, Y \cup Z)|$. Indeed, let k be such that $\chi(F_j, Y) = F_k \setminus Y$. We have $F_k \setminus Y \subseteq X \cup \chi(F, Y)$, which implies
$$F_k \subseteq X \cup \chi(F, Y) \subseteq X \cup F$$
and
$$F_k \setminus X \subseteq \chi(F, Y) \setminus X.$$
Thus, we get
$$\begin{aligned} |\chi(F, X)| &\leqslant |F_k \setminus X| = |(F_k \setminus X) \cap (\chi(F, Y) \setminus X)| \\ &\leqslant |(F_k \setminus Y) \cap \chi(F, Y)| \\ &= |\chi(F_j, Y) \cap \chi(F, Y)| = t. \end{aligned}$$
This means that the number of bits of the encoding of F is
$$\log_2(r^{\ell-t}) + 1 \leqslant \log_2(r^{\ell-\chi(F,X)}) + 1.$$
5. Finally, we encode $X \cap \chi(F, Y)$, which can be done with $|\chi(F, Y)|$ bits. The set X can now be easily recovered from $X \cup \chi(F, Y)$, F and $X \cap \chi(F, Y)$. And having recovered X, we can recover $Z = X \setminus Y$.

It is not hard to see that the encoding is prefix-free, because the length of the encoding of (Z, F) is determined by $|\chi(F, Y)|$, which is encoded as a prefix in the form $0^{|\chi(F,Y)|}1$. Furthermore, the expected length of the encoding is upper bounded by
$$(3 + \log_2(n/z)) \mathop{\mathbb{E}}_{\mathbf{F}}[|\chi(\mathbf{F}, Y)|] + \mathop{\mathbb{E}}_{\mathbf{Z},\mathbf{F}}[|\chi(\mathbf{F}, Y \cup \mathbf{Z})|] \cdot \log_2(r) + \log_2 \binom{n-y}{z} + \ell \cdot \log_2(r) + 3.$$

[1] The main difference with the proof of Proposition 4.1.2 is that, instead of giving an encoding of fixed size and using the length of the encoding to bound the number of pairs, we will give a *prefix-free* encoding, and use the *average* length of the encoding to bound the number of pairs via Corollary 4.2.2.

Since there are at least $r^\ell \cdot \binom{n}{y-z}$ pairs (Z, F), Corollary 4.2.2 now gives us

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathbf{Z},\mathbf{F}}[|\chi(\mathbf{F}, Y \cup \mathbf{Z})|] \cdot \log_2 r &\leqslant (3 + \log_2(n/z)) \mathop{\mathbb{E}}_{\mathbf{F}}[|\chi(\mathbf{F}, Y)|] + 3 \\ &\leqslant (6 + \log_2(n/z)) \mathop{\mathbb{E}}_{\mathbf{F}}[|\chi(\mathbf{F}, Y)|] \\ &= \log_2(64n/z) \mathop{\mathbb{E}}_{\mathbf{F}}[|\chi(\mathbf{F}, Y)|] \\ &\leqslant \log_2(r/2) \mathop{\mathbb{E}}_{\mathbf{F}}[|\chi(\mathbf{F}, Y)|],\end{aligned}$$

which implies inequality (4.1), thus finishing the proof. □

4.2.3 Proof of Theorem 3.4.3

We will now apply Lemma 4.2.3 to prove Theorem 3.4.3.

Proof of Theorem 3.4.3. In what follows, we suppose B is a large enough universal constant.

The proof is by induction on ℓ. Suppose $\ell = 1$. Then $\mathcal{F}$ is a family of singletons. Therefore, the probability that $\mathbf{W} \subseteq_M \Gamma$ does not contain any set of $\mathcal{F}$ is equal to $\binom{n-|\mathcal{F}|}{M}/\binom{n}{M}$. We get

$$\Pr[\mathbf{W} \text{ is independent in } \mathcal{F}] = \frac{\binom{n-|\mathcal{F}|}{M}}{\binom{n}{M}} \leqslant \left(\frac{n-M}{n}\right)^{|\mathcal{F}|} \leqslant (1-p/2)^{|\mathcal{F}|} \leqslant e^{-|\mathcal{F}|p/2} < \varepsilon.$$

Hence, the family $\mathcal{F}$ is itself a (M, ε)-slice-sunflower (with an empty core).

We now proceed by induction, supposing $\ell \geqslant 2$ and that the claim holds for all k-uniform families such that $k < \ell$. Let $r := Bx \log_2 x$.

Case 1. Observe that, if $\mathcal{F}$ is not $(1, r, r^2, \ldots, r^\ell)$-well-spread, then there exists a set $T \subseteq \Gamma$ such that $d_{\mathcal{F}}(T) \geqslant r^{\ell - |T|}$. Therefore, by the induction hypothesis, $\mathcal{F}_T$ contains a (M, ε)-slice-sunflower $\mathcal{F}'_T$. Observe that the family $\{U \cup T : U \in \mathcal{F}'_T\} \subseteq \mathcal{F}$ is a (M, ε)-slice-sunflower.

Case 2. We now consider the case when $\mathcal{F}$ is $(1, r, r^2, \ldots, r^\ell)$-well-spread. We will use Lemma 4.2.3 to finish the proof. Let $s = \lceil \log_2(\ell/\varepsilon) \cdot \log_2 r \rceil$ and $K = s \cdot 128 \cdot \lceil n/r \rceil$. We have

$$\begin{aligned}K &< 512 \cdot \log_2 r \cdot \log_2(\ell/\varepsilon) \cdot n/r \\ &= 512 \cdot n \cdot \frac{\log_2 B + \log_2 x + \log\log_2 x}{Bx \log_2 x} \\ &= 512 \cdot np \cdot \frac{\log_2 B + \log_2 x + \log\log_2 x}{B \log_2(\ell/\varepsilon) \log_2 x} \\ &\leqslant 512 \cdot np \cdot \frac{\log_2 B}{B} < M,\end{aligned}$$

for B large enough. Therefore, letting $\mathbf{X} \subseteq_K \Gamma$, we get by Lemma 4.2.3 that

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathbf{W},\mathbf{S}}[|\chi(\mathbf{F}, \mathbf{W})|] &\leqslant \mathop{\mathbb{E}}_{\mathbf{X},\mathbf{F}}[|\chi(\mathbf{F}, \mathbf{X})|] \\ &\leqslant \ell \cdot (1 - 1/\log_2 r)^s \\ &\leqslant \ell \cdot (1 - 1/\log_2 r)^{\log_2(\ell/\varepsilon) \cdot \log_2 r} \\ &< \ell \cdot 2^{-\log_2(\ell/\varepsilon)} = \varepsilon.\end{aligned}$$

We can conclude the proof by applying Markov's inequality (Proposition 2.2.2), as follows:

$$\Pr_{\mathbf{W}}[\mathbf{W} \text{ is independent in } \mathcal{F}] = \Pr_{\mathbf{W},\mathbf{F}}[|\chi(\mathbf{F}, \mathbf{W})| > 0] \leqslant \mathop{\mathbb{E}}_{\mathbf{W},\mathbf{F}}[|\chi(\mathbf{F}, \mathbf{W})|] < \varepsilon.$$ □

4.3 Further questions

The definition of "almost" satisfying sets (Definition 4.1.1) motivates the following definition.

Definition 4.3.1. *A hypergraph is said to be a* (M, ε, t)-almost sunflower *if the family of* petals $\{F \setminus Y : F \in E(\mathcal{F})\}$ *is* (M, ε, t)*-almost satisfying, where* $Y := \bigcap \mathcal{F}$ *is called the* core.

Given that robust sunflowers and its variations have found ample applications in computational complexity, perhaps it is possible that the weaker variant considered here is enough for some applications. Better bounds for almost sunflowers can be expected, in comparison to robust sunflowers, which in turn might lead to stronger bounds in the respective applications.

Problem 4.3.2. *Find applications of almost sunflowers, checking whether they can offer improvements on the bounds obtained by an application of robust sunflowers.*

Chapter 5

Sunflowers and the approximation method

5.1 Introduction

In this chapter, we will give a general overview of how different notions of sunflowers play out in an application of the approximation method of Razborov [Raz85b]. We hope that a generalized exposition of the method will help the reader to appreciate the technical contribution of the results of the following chapters.

Let $f_{\mathsf{T}} : 2^{\Gamma} \to \{0,1\}$ be a target monotone Boolean function for which we want to prove a monotone circuit lower bound. Since complexity lower bounds are always asymptotic, we will henceforth suppose that $|\Gamma|$ is large enough. To apply the approximation method for f_{T}, we must find two distributions $\mathcal{D}^+$ and $\mathcal{D}^-$ supported in 2^{Γ} satisfying the following properties.

Definition 5.1.1. *A pair of distributions $(\mathcal{D}^+, \mathcal{D}^-)$ is called a* pair of test distributions for f_{T} *if*

$$\Pr_{\mathbf{Y}\sim\mathcal{D}^+}[f_{\mathsf{T}}(\mathbf{Y}) = 1] + \Pr_{\mathbf{N}\sim\mathcal{D}^-}[f_{\mathsf{T}}(\mathbf{N}) = 0] \geqslant 8/5.$$

Because of this property, the distribution $\mathcal{D}^+$ is called the positive test distribution*, and $\mathcal{D}^-$ is called the* negative test distribution*.*

We will always denote a random sample from $\mathcal{D}^+$ by $\mathbf{Y}$ (for "Yes") and a random sample from $\mathcal{D}^-$ by $\mathbf{N}$ (for "No"). We also define a set $\mathcal{A}$ of monotone Boolean functions called *approximators*, and we require the following properties of $\mathcal{A}$.

Definition 5.1.2 (Approximators make many errors)**.** *We say that $\mathcal{A} \subseteq \mathsf{MonBool}(\Gamma)$ is* inaccurate for $(\mathcal{D}^+, \mathcal{D}^-)$ *if the following holds for every approximator $g \in \mathcal{A}$:*

$$\Pr_{\mathbf{Y}\sim\mathcal{D}^+}[g(\mathbf{Y}) = 1] + \Pr_{\mathbf{N}\sim\mathcal{D}^-}[g(\mathbf{N}) = 0] \leqslant 3/2.$$

Definition 5.1.3. *We say that $\mathcal{A} \subseteq \mathsf{MonBool}(\Gamma)$ is a δ-*approximator for $(\mathcal{D}^+, \mathcal{D}^-)$ *if, for every Boolean function $f : 2^{\Gamma} \to \{0,1\}$, there exists an approximator $f^{\mathcal{A}} \in \mathcal{A}$ such that*

$$\Pr_{\mathbf{Y}\sim\mathcal{D}^+}[f(\mathbf{Y}) = 1 \text{ and } f^{\mathcal{A}}(\mathbf{Y}) = 0] + \Pr_{\mathbf{N}\sim\mathcal{D}^-}[f(\mathbf{N}) = 0 \text{ and } f^{\mathcal{A}}(\mathbf{N}) = 1] \leqslant \delta \cdot \mathsf{C}_{\mathsf{mon}}(f).$$

*We say that $f^{\mathcal{A}}$ δ-*agrees *with f on $(\mathcal{D}^+, \mathcal{D}^-)$.*

Definition 5.1.2 means that every approximator commits many "mistakes" in the distributions $\mathcal{D}^+$ and $\mathcal{D}^-$. Definition 5.1.3 means that every Boolean function computed by a "small" monotone circuit (i.e.: size $s \ll 1/\delta$) agrees with an approximator on both $\mathcal{D}^+$ and $\mathcal{D}^-$ with high probability. These definitions imply a lower bound on the size of monotone circuits computing the target function.

Theorem 5.1.4. *Let $(\mathcal{D}^+, \mathcal{D}^-)$ be a pair of test distributions for f_T. Suppose that there exists $\mathcal{A} \subseteq \mathsf{MonBool}(\Gamma)$ such that $\mathcal{A}$ inaccurate on $(\mathcal{D}^+, \mathcal{D}^-)$ and that $\mathcal{A}$ is a $(1/10S)$-approximator for $(\mathcal{D}^+, \mathcal{D}^-)$. Then any monotone circuit computing f_T has at least S gates.*

Proof. Let $g \in \mathcal{A}$ be an approximator which $(1/10S)$-agrees with f_T on $(\mathcal{D}^+, \mathcal{D}^-)$. We have

$$\begin{aligned} 8/5 &\leqslant \Pr[f_\mathsf{T}(\mathbf{Y}) = 1] + \Pr[f_\mathsf{T}(\mathbf{N}) = 0] \\ &\leq \Pr[f_\mathsf{T}(\mathbf{Y}) = 1 \text{ and } g(\mathbf{Y}) = 0] + \Pr[g(\mathbf{Y}) = 1] \\ &+ \Pr[f_\mathsf{T}(\mathbf{N}) = 0 \text{ and } g(\mathbf{N}) = 1] + \Pr[g(\mathbf{N}) = 0] \\ &\leqslant 3/2 + \mathsf{C}_{\mathsf{mon}}(f_\mathsf{T})/(10S). \end{aligned}$$

This implies $\mathsf{C}_{\mathsf{mon}}(f_\mathsf{T}) \geqslant S$. □

We now discuss in general terms how to find a good pair of test distributions for f_T and how to obtain a set of approximators satisfying Definitions 5.1.2 and 5.1.3, highlighting the role of variants of sunflowers.

5.2 Test distributions

It is usually not very complicated to find test distibutions $\mathcal{D}^+$ and $\mathcal{D}^-$. One possibility is to let $\mathcal{D}^+$ be a minterm of f_T chosen uniformly at random, and to let $\mathcal{D}^-$ be a maxterm of f_T chosen uniformly at random. This is exactly the choice of Alon and Boppana [AB87] for the $\mathsf{Clique}(n,k)$ function. Sometimes one may also choose $\mathcal{D}^-$ to be a p-biased distribution, with p being a probability function below the threshold for the ocurrence of a minterm of f_T[1].

The choice of the test distributions may also influence the quality of the lower bound. Using $\mathsf{Clique}(n,k)$ again as an example, a p-biased distribution below the threshold for the ocurrence of k-cliques corresponds to an Erdős-Rényi random graph $G(n,p)$, with $p \leqslant n^{-2/(k-1)}$. This distribution was first considered for monotone circuit lower bounds against $\mathsf{Clique}(n,k)$ in [Ros14], in the regime when k is constant. It turns out that, extending Rossman's analysis for a larger range of k, we are able to improve Alon and Boppana's bound for $\mathsf{Clique}(n,k)$, from $n^{\Omega(\sqrt{k})}$ to $n^{\Omega(k)}$. This will be the subject matter of Chapter 7.

5.3 Approximators

A crucial step when applying the approximation method is to build a good set of approximators. Any viable set of approximators must be a *legitimate model*, which we define as follows.

Definition 5.3.1 ([Raz85b])**.** *A subposet $\mathcal{A}$ of $\mathsf{MonBool}(\Gamma)$[2] is said to be a* legitimate model *with* approximating operators $\sqcup, \sqcap : \mathcal{A} \times \mathcal{A} \to \mathcal{A}$ *if the following holds:*

1. *The operators $\sqcup$ and $\sqcap$ are commutative;*
2. *We have $x_\gamma \in \mathcal{A}$ for all $\gamma \in \Gamma$ and $\mathbb{0}, \mathbb{1} \in \mathcal{A}$.*

Functions in $\mathcal{A}$ are called approximators*. The operators $\sqcup, \sqcap$ are called* approximate OR *and* approximate AND*, respectively.*

The main challenge of proving that $\mathcal{A}$ is a δ-approximator is to find a function $f^\mathcal{A} \in \mathcal{A}$ that approximates a given function f in the positive and negative distributions. With a good legitimate model in hand, we can do this in the following way.

Definition 5.3.2. *Given a $\{\vee, \wedge\}$-circuit C, we define its* approximate circuit $C^\mathcal{A}$ *as the $\{\sqcup, \sqcap\}$-circuit obtained by replacing the $\vee$ and $\wedge$ gates of C by $\sqcup$ and $\sqcap$ gates, respectively.*

[1] Such a p always exists because of Theorem 2.2.5.

[2] Recall that $\mathsf{MonBool}(\Gamma)$ is the set of monotone Boolean functions $2^\Gamma \to \{0,1\}$. (See Section 2.3.)

If $\mathcal{A}$ is a legitimate model, then $C^{\mathcal{A}}$ is well-defined (because the approximating operators are commutative) and computes a function of $\mathcal{A}$ (because the input gates also belong to $\mathcal{A}$, as well as the constant functions). To prove that $\mathcal{A}$ is δ-approximator for $(\mathcal{D}^+, \mathcal{D}^-)$, one takes a monotone circuit C computing f and defines $f^{\mathcal{A}}$ as the approximator computed by the approximate circuit $C^{\mathcal{A}}$. To show that an approximator obtained in this way provides a good approximation in the test distributions, we show that $C^{\mathcal{A}}$ provides a good gate-by-gate approximation of C on $(\mathcal{D}^+, \mathcal{D}^-)$.

Definition 5.3.3. *We say that $\mathcal{A} \subseteq \mathsf{MonBool}(\Gamma)$ is a* gate-by-gate δ-approximator for $(\mathcal{D}^+, \mathcal{D}^-)$ *if, for every $f, g \in \mathcal{A}$, we have*

$$\Pr[(f \wedge g)(\mathbf{Y}) = 1 \text{ and } (f \sqcap g)(\mathbf{Y}) = 0] \leqslant \delta \tag{5.1}$$
$$\Pr[(f \vee g)(\mathbf{Y}) = 1 \text{ and } (f \sqcup g)(\mathbf{Y}) = 0] \leqslant \delta. \tag{5.2}$$
$$\Pr[(f \wedge g)(\mathbf{N}) = 0 \text{ and } (f \sqcap g)(\mathbf{N}) = 1] \leqslant \delta \tag{5.3}$$
$$\Pr[(f \vee g)(\mathbf{N}) = 0 \text{ and } (f \sqcup g)(\mathbf{N}) = 1] \leqslant \delta. \tag{5.4}$$

The constant δ is called approximation error*.*

Lemma 5.3.4. *Suppose that $\mathcal{A}$ is a gate-by-gate $(\delta/2)$-approximator for $(\mathcal{D}^+, \mathcal{D}^-)$. Then $\mathcal{A}$ is a δ-approximator for $(\mathcal{D}^+, \mathcal{D}^-)$.*

Proof. Let $f \in \mathsf{MonBool}(\Gamma)$. Let C be a monotone circuit of size s computing f. Let $f_1, f_2 \ldots, f_s$ be a straight-line program that describes C. If f_i is an $\wedge$-gate, let $f_i^{\mathcal{A}}$ be an $\sqcap$-gate. Equivalently, let $f_i^{\mathcal{A}}$ be an $\sqcup$-gate if f_i is an $\vee$-gate. Clearly, $f_1^{\mathcal{A}}, \ldots, f_s^{\mathcal{A}}$ is a straight-line program that describes $C^{\mathcal{A}}$. If $C(\mathbf{Y}) = 1$ but $C^{\mathcal{A}}(\mathbf{Y}) = 0$, then there exists a gate f_i of C such that $f_i(\mathbf{Y}) = 1$ but $f_i^{\mathcal{A}}(\mathbf{Y}) = 0$. Without loss of generality, suppose this is the first such gate in the straight-line program. Let $j, k < i$ be such that $f_i = f_j \circ f_k$, where $\circ \in \{\vee, \wedge\}$. Let $\bar{\circ}$ denote the corresponding approximating operator for $\circ$. By the minimality of f_i, we have that $(f_j^{\mathcal{A}} \circ f_k^{\mathcal{A}})(\mathbf{Y}) = 1$ but $(f_j^{\mathcal{A}} \bar{\circ} f_k^{\mathcal{A}})(\mathbf{Y}) = 0$. By assumption, this happens with probability at most $\delta/2$. Therefore, by an application of the union bound on the gates of C, we get that

$$\Pr[C(\mathbf{Y}) = 1 \text{ and } C^{\mathcal{A}}(\mathbf{Y}) = 0] \leqslant (\delta/2) \cdot s.$$

The same argument shows

$$\Pr[C(\mathbf{N}) = 0 \text{ and } C^{\mathcal{A}}(\mathbf{N}) = 1] \leqslant (\delta/2) \cdot s.$$

Therefore, we have $\Pr[C(\mathbf{Y}) = 1 \text{ and } C^{\mathcal{A}}(\mathbf{Y}) = 0] + \Pr[C(\mathbf{N}) = 0 \text{ and } C^{\mathcal{A}}(\mathbf{N}) = 1] \leqslant \delta \cdot s$. □

By Theorem 5.1.4, we can summarize what we have seen in the following result, which gives an agenda for proving monotone circuit lower bounds with the approximation method.

Theorem 5.3.5. *Let $(\mathcal{D}^+, \mathcal{D}^-)$ be a pair of test distributions for f_{T}. Suppose that there exists a legitimate model $\mathcal{A}$ such that $\mathcal{A}$ inaccurate on $(\mathcal{D}^+, \mathcal{D}^-)$ and that $\mathcal{A}$ is a gate-by-gate $(1/20S)$-approximator for $(\mathcal{D}^+, \mathcal{D}^-)$. Then any monotone circuit computing f_{T} has at least S gates.*

We will now describe a general construction of legitimate models. This construction will generalize the constructions of legitimate models of [Raz85b], [Raz85a], [AB87] and [CKR20]. This includes monotone Boolean functions such as clique, matching, Andreev's function [And85] and Harnik-Raz function [HR00]. Though such a generalized construction is already given in a book by Jukna [Juk12, Section 9.10.1], our construction has the advantage of abstracting the type of sunflower used in the construction, whereas his construction is limited to lopsided sunflowers (see Section 3.2).

5.4 A general construction of legitimate models

Let $\mathcal{F} \subseteq 2^\Gamma$. Let $\lceil \mathcal{F} \rceil : 2^\Gamma \to \{0,1\}$ denote the function which assumes value 1 on an input X if there exists a set $F \in \mathcal{F}$ such that $F \subseteq X$. In other words, we have

$$\lceil \mathcal{F} \rceil = \bigvee_{F \in \mathcal{F}} \mathbb{1}_F = \bigvee_{F \in \mathcal{F}} \bigwedge_{f \in F} x_f.$$

Henceforth, we fix a number $\varepsilon > 0$. This number will measure the quality of the approximation under the negative distribution $\mathcal{D}^-$. We also fix a a sequence of disjoint set families $\mathcal{X}_1, \mathcal{X}_2, \ldots \subseteq 2^\Gamma$ and a parameter w. We let $\mathcal{X} := \bigcup \mathcal{X}_i$. The choice of this family and of the parameter w is problem-specific. For convenience, we let $\mathcal{X}_{\leqslant w} := \bigcup_{i \leqslant w} \mathcal{X}_i$. We will make the following assumptions about $\mathcal{X}$:

1. All sets of $\mathcal{X}_i$ have the same size, denoted by $s(\mathcal{X}, \ell)$;
2. We have $s(\mathcal{X}, i) < s(\mathcal{X}, j)$ for $i < j$.
3. $\{\gamma\} \in \mathcal{X}_{\leqslant w}$ for all $\gamma \in \Gamma$.

A pair $(\mathcal{X},w)$ satisfying (1)-(3) is called *ambient family*. The ambient family $\mathcal{X}$ such that $\mathcal{X}_\ell = \binom{\Gamma}{\ell}$ and $\mathcal{X} = 2^\Gamma$ is called the *trivial ambient family*.

Let $X \in 2^\Gamma$. We write $(\mathcal{F} \vdash_\varepsilon X)_{\mathcal{D}^-}$ to denote that

$$\Pr_{\mathbf{N} \sim \mathcal{D}^-} [\mathbf{N} \cup X \text{ is independent in } \mathcal{F}] < \varepsilon.$$

Equivalently, this means that

$$\mathop{\mathbb{E}}_{\mathbf{N} \sim \mathcal{D}^-} [\lceil \mathcal{F} \rceil (\mathbf{N} \cup X)] > 1 - \varepsilon.$$

We read $(\mathcal{F} \vdash_\varepsilon X)_{\mathcal{D}^-}$ as $\mathcal{F}$ *is ε-satisfied by X on $\mathcal{D}^-$*. We say that $\mathcal{F}$ is $(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$-*closed* if, for every $X \in \mathcal{X}_{\leqslant w}$, we have

$$(\mathcal{F} \vdash_\varepsilon X)_{\mathcal{D}^-} \implies X \in \mathcal{F}.$$

When $\mathcal{X}_{\leqslant w}$, $\mathcal{D}^-$ and ε are clear in context, we will simply write *closed* for short. Observe that the intersection of any two closed families is also closed; furthermore, the family 2^Γ is closed. We may thus define the *closure* $\mathrm{cl}(\mathcal{F})$ of a family $\mathcal{F}$ as $\mathrm{cl}(\mathcal{F}) = \bigcap \{\mathcal{H} : \mathcal{F} \subseteq \mathcal{H} \text{ and } \mathcal{H} \text{ is closed}\}$. It is not hard to check that $\mathrm{cl}(\cdot)$ is, indeed, a closure operator in the poset 2^Γ.

Definition 5.4.1. *The set of approximators $\mathcal{A} = \mathcal{A}(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$ determined by $\mathcal{X}, \mathcal{D}^-$ and ε is defined as*

$$\mathcal{A} = \left\{ \lceil \mathcal{F} \rceil : \mathcal{F} \subseteq \mathcal{X}, \mathcal{F} \text{ is } (\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)\text{-closed} \right\},$$

with approximating operators

$$\lceil \mathcal{F} \rceil \sqcup \lceil \mathcal{H} \rceil := \lceil \mathrm{cl}(\mathcal{F} \cup \mathcal{H}) \rceil,$$
$$\lceil \mathcal{F} \rceil \sqcap \lceil \mathcal{H} \rceil := \lceil \mathcal{F} \cap \mathcal{H} \rceil.$$

Now, in order to conclude that $\mathcal{A} = \mathcal{A}(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$, together with the approximating operators $\sqcup$ and $\sqcap$, is a legitimate model, we will need to make a minor assumption on $\mathcal{D}^-$.

Proposition 5.4.2. *Suppose that $\Pr_{\mathbf{N} \sim \mathcal{D}^-}[\gamma \in \mathbf{N}] \leqslant 1 - \varepsilon$ for all $\gamma \in \Gamma$. Then $\langle \mathcal{A}, \sqcup, \sqcap \rangle$ is a legitimate model for $\mathcal{A} = \mathcal{A}(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$.*

Proof. The only requirement we need to check is $x_\gamma \in \mathcal{A}$ – all else is trivial. Let

$$\mathcal{F}_\gamma = \{A \in \mathcal{X}_{\leqslant w} : \gamma \in A\}.$$

It is easy to see that $x_\gamma = \lceil \mathcal{F}_\gamma \rceil$, since $\{\gamma\} \in \mathcal{F}_\gamma$. Now, it suffices to check that $\mathcal{F}_\gamma$ is closed. Suppose that $(\mathcal{F}_\gamma \vdash_\varepsilon X)_{\mathcal{D}^-}$ for some $X \in \mathcal{X}_{\leqslant w}$. This implies that $\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\gamma \in \mathbf{N} \cup X] > 1 - \varepsilon$. If $X \notin \mathcal{F}_\gamma$, this means that $\gamma \notin X$, which implies $\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\gamma \in \mathbf{N}] > 1 - \varepsilon$, a contradiction. □

5.4.1 Examples of the general construction

To facilitate the understanding of the construction, we provide in the table below some instances of this construction.

Function	$\mathcal{X}_\ell$	w	$\mathcal{D}^-$	ε	Reference
Matching	Matchings of ℓ edges	$\delta \cdot \log n$	-	$n^{-\Omega(\log^3 n)}$	[Raz85a]
$\mathsf{Clique}(n,k)$	Cliques of ℓ vertices	$\sqrt{k}$	$(k-1)$-cocliques	$n^{-\sqrt{k}}$	[AB87]
$\mathsf{Clique}(n,k)$	Cliques of ℓ vertices	$\delta \cdot k$	$G(n,p)$	n^{-k}	[CKR20]
$\mathsf{HR}(c,k)$[3]	All sets of size ℓ	c	$\lfloor n/2 \rfloor$-slice	n^{-2c}	[CKR20]

The negative distribution for Matching is as follows. We choose a two-coloring c of the complete $n \times n$ bipartite graph $K_{n,n}$ uniformly at random, and let $\mathbf{N}$ be the subgraph with $E(\mathbf{N}) = \{(u,v) : c(u) = c(v)\}$. Furthermore, $(k-1)$-cocliques are complete $(k-1)$-partite graphs.

5.4.2 Properties of the general construction

The following lemma is easy to check, following from the definitions.

Lemma 5.4.3. *Let $\mathcal{F}, \mathcal{H} \subseteq 2^\Gamma$ be such that $\mathcal{F} \subseteq \mathcal{H}$ and let $X, Y \in 2^\Gamma$ be such that $X \subseteq Y$. The following holds.*

1. *If $(\mathcal{F} \vdash_\varepsilon X)_{\mathcal{D}^-}$, then $(\mathcal{H} \vdash_\varepsilon X)_{\mathcal{D}^-}$.*
2. *If $X \in \mathcal{F}$, then $(\mathcal{F} \vdash_\varepsilon Y)_{\mathcal{D}^-}$.*
3. *If $X \in \mathcal{F}$, $\mathcal{F}$ is closed and $Y \in \mathcal{X}_{\leqslant w}$, then $Y \in \mathcal{F}$.*

We also remark that the estimation of the approximation errors can be greatly simplified for this legitimate model.

Remark 5.4.4. *For every $f, g \in \mathcal{A}$, we have*

1. *$f \vee g \leqslant f \sqcup g$;*
2. *$f \wedge g \geqslant f \sqcap g$.*

Therefore, there is no approximation error on $\mathbf{Y}$ when changing an $\vee$-gate for an $\sqcup$-gate. For the same reason, there is no approximation error on $\mathbf{N}$ when changing an $\wedge$-gate for an $\sqcap$-gate. This means that items (5.2) and (5.3) of Definition 5.3.3 hold trivially for $\mathcal{A}$.

The following lemma shows that ε measures the error on the negative distribution when taking the $(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$-closure of a family.

Lemma 5.4.5. *For every family $\mathcal{F} \subseteq 2^\Gamma$, we have*

$$\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\lceil \mathcal{F} \rceil(\mathbf{N}) = 0 \text{ and } \lceil \mathrm{cl}(\mathcal{F}) \rceil(\mathbf{N}) = 1] \leqslant \varepsilon \, |\mathcal{X}_{\leqslant w}| \, .$$

Proof. We first prove that, for a positive integer $t \leqslant |\mathcal{X}_{\leqslant w}|$, there exists sets $X_1, \ldots, X_t \in \mathcal{X}_{\leqslant w}$ and families $\mathcal{F}_0, \mathcal{F}_1, \ldots, \mathcal{F}_t \subseteq 2^\Gamma$ such that

1. $\mathcal{F}_0 = \mathcal{F}$,

[3] See Chapter 6.

2. $\mathcal{F}_i = \mathcal{F}_{i-1} \cup \{X_i\}$,
3. $\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\lceil\mathcal{F}_{i-1}\rceil(\mathbf{N} \cup X_i) = 1] \geqslant 1 - \varepsilon$,
4. $\mathcal{F}_t = \mathrm{cl}(\mathcal{F})$.

Indeed, if $\mathcal{F}_{i-1}$ is not closed, there exists $X_i \in \mathcal{X}_{\leqslant w}$ such that $\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\lceil\mathcal{F}_{i-1}\rceil(\mathbf{N} \cup X_i) = 1] \geqslant 1-\varepsilon$ but $X_i \notin \mathcal{F}_{i-1}$. We let $\mathcal{F}_i := \mathcal{F}_{i-1} \cup \{X_i\}$. Clearly, we have that $\mathcal{F}_t$ is closed, and that $t \leqslant |\mathcal{X}_{\leqslant w}|$. Moreover, by induction we obtain that $\mathcal{F}_i \subseteq \mathrm{cl}(\mathcal{F})$ for every $i \in [t]$. It follows that $\mathcal{F}_t = \mathrm{cl}(\mathcal{F})$. Furthermore, we have

$$\begin{aligned}
\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\lceil\mathcal{F}\rceil(\mathbf{N}) = 0 \text{ and } \lceil\mathrm{cl}(\mathcal{F})\rceil(\mathbf{N}) = 1] &\leqslant \sum_{i=1}^{t} \Pr[\lceil\mathcal{F}_{i-1}\rceil(\mathbf{N}) = 0 \text{ and } \lceil\mathcal{F}_i\rceil(\mathbf{N}) = 1] \\
&= \sum_{i=1}^{t} \Pr[\lceil\mathcal{F}_{i-1}\rceil(\mathbf{N}) = 0 \text{ and } X_i \subseteq \mathbf{N}] \\
&\leqslant \sum_{i=1}^{t} \Pr[\lceil\mathcal{F}_{i-1}\rceil(\mathbf{N} \cup X_i) = 0] \\
&\leqslant \varepsilon t = \varepsilon\,|\mathcal{X}_{\leqslant w}|\,. \qquad \square
\end{aligned}$$

This immediately implies a bound on the approximation errors on $\mathcal{D}^-$.

Lemma 5.4.6. *For every $f, g \in \mathcal{A}$, we have*

$$\Pr_{\mathbf{N}\sim\mathcal{D}^-}[(f \vee g)(\mathbf{N}) = 1 \text{ and } (f \sqcup g)(\mathbf{N}) = 0] \leqslant \varepsilon\,|\mathcal{X}_{\leqslant w}|\,.$$

Proof. Let $\mathcal{F}$ and $\mathcal{H}$ be $(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$-closed families such that $f = \lceil\mathcal{F}\rceil$ and $g = \lceil\mathcal{H}\rceil$. We have $f \vee g = \lceil\mathcal{F} \cup \mathcal{H}\rceil$ and $f \sqcup g = \lceil\mathrm{cl}(\mathcal{F} \cup \mathcal{H})\rceil$. The result now follows from Lemma 5.4.5. $\square$

5.5 Bounding the number of minterms with abstract sunflowers

For a family $\mathcal{F} \subseteq \mathcal{X}$ and $\ell > 0$, let $\mathcal{M}_{\mathcal{X},\ell}(\mathcal{F})$ denote the family of its minimal sets (under inclusion) in $\mathcal{X}_\ell$, and let $\mathcal{M}_{\mathcal{X}}(\mathcal{F}) = \bigcup_\ell \mathcal{M}_{\mathcal{X},\ell}(\mathcal{F})$. The main property of closed families is that the number of its minimal elements can be bounded by a sunflower-type theorem. This property is used to show that $\mathcal{A}$ is inaccurate and to bound the approximation errors under the positive distribution.

Giving bounds for the abstract sunflower defined below is the core combinatorial challenge of applying the legitimate model of the previous section. Of course, this usually depends on the specific choice of the negative distribution $\mathcal{D}^-$ and the ambient family $\mathcal{X}$. We will thus define next a general sunflower structure, dependent upon the choice of $\mathcal{X}$, $\mathcal{D}^-$ and ε.

Definition 5.5.1. *We say that a hypergraph $\mathcal{H} \subseteq \mathcal{X}$ is an $(\mathcal{X}, \mathcal{D}^-, \varepsilon)$-sunflower if there exists a set $Y \in \mathcal{X}$ (called* core*) such that $Y \subsetneq S$ for some $S \in \mathcal{H}$ and $(\mathcal{H} \vdash_\varepsilon Y)_{\mathcal{D}^-}$. When $\mathcal{X}$ is the trivial ambient family, we say that $\mathcal{H}$ is a $(\mathcal{D}^-, \varepsilon)$-sunflower.*

Observe that, when $\mathcal{D}^-$ is p-biased, a $(\mathcal{D}^-, \varepsilon)$-sunflower is precisely a (p, ε)-robust lopsided sunflower (see Definition 3.5.5). Analogously to Definition 3.3.5, we may define the extremal function of this abstract sunflower, which measures the size of the largest set system without this sunflower.

Definition 5.5.2. *We define $\mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon)$ as the size of the largest hypergraph $\mathcal{H}$ with vertex set Γ such that $\mathcal{H} \subseteq \mathcal{X}_\ell$ and $\mathcal{H}$ does not contain a $(\mathcal{X}, \mathcal{D}^-, \varepsilon)$-sunflower. When $\mathcal{X}$ is the trivial ambient family, we write $\mathrm{Ex}(\ell, \mathcal{D}^-, \varepsilon)$ to denote $\mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon)$.*

Remark 5.5.3. *When $\mathcal{D}^-$ is p-biased, we have $\mathrm{Ex}(\ell, \mathcal{D}^-, \varepsilon) = f_{\mathrm{RL}}(\ell, p, \varepsilon)$. (See Section 3.5.2.)*

Lemma 5.5.4. *Let $\mathcal{F} \subseteq \mathcal{X}_{\leqslant w}$ be $(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$-closed. We have $|\mathcal{M}_{\mathcal{X},\ell}(\mathcal{F})| \leqslant \mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon)$.*

Proof. For a contradiction, let us suppose that $|\mathcal{M}_{\mathcal{X},\ell}(\mathcal{F})| > \mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon)$. Since $\mathcal{M}_{\mathcal{X},\ell}(\mathcal{F}) \subseteq \mathcal{X}_\ell$, there exists a $(\mathcal{X}, \mathcal{D}^-, \varepsilon)$-sunflower $\mathcal{S} \subseteq \mathcal{M}_{\mathcal{X},\ell}(\mathcal{F})$. Let Y be the core of $\mathcal{S}$. By definition, there exists a set $F \in \mathcal{S}$ such that $Y \subsetneq F$. Observe that $Y \in \mathcal{X}_{\leqslant w}$, because $F \in \mathcal{X}_{\leqslant w}$. Since $\mathcal{S}$ is a $(\mathcal{X}, \mathcal{D}^-, \varepsilon)$-sunflower and $\mathcal{S} \subseteq \mathcal{F}$, we get by Lemma 5.4.3 that $(\mathcal{F} \vdash_\varepsilon X)_{\mathcal{D}^-}$. Therefore, we obtain $Y \in \mathcal{F}$ because $\mathcal{F}$ is closed and $Y \in \mathcal{X}_{\leqslant w}$. But we must have $Y \notin \mathcal{F}$, because $Y \subsetneq F$ and F is a minimal element of $\mathcal{F}$. This is a contradiction which concludes the proof. □

5.5.1 Examples of abstract sunflower bounds

For comparison, we provide a few examples of abstract sunflowers, along with corresponding bounds. The first two bounds, proved in [Raz85b] and [AB87], can be obtained from an application of lopsided sunflowers (see Defintion 3.2.1). They implicitly prove that a lopsided sunflower $\mathcal{S}$ with r petals and core Y satisfies $(\mathcal{S} \vdash_\varepsilon Y)_{\mathcal{D}^-}$, for a proper choice of r. This implies $\mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon) \leqslant (r-1)^\ell$, by Theorem 3.2.3. Furthemore, clique sunflowers are introduced in Chapter 7.

$\mathcal{X}$	$s(\mathcal{X}, \ell)$	$\mathcal{D}^-$	Upper bound	Technique	Reference
Matchings	ℓ	-	$(\log(1/\varepsilon) \cdot 2^\ell)^{2\ell}$	Lopsided sunflowers	[Raz85a]
Cliques	$\binom{\ell}{2}$	$(k-1)$-cocliques	$(\log(1/\varepsilon) \cdot (\frac{k-1}{k-\ell})^\ell)^\ell$	Lopsided sunflowers	[AB87]
$2^{[n]}$	ℓ	p-biased	$(\log \ell \cdot \log(1/\varepsilon)/p)^{O(\ell)}$	Robust sunflowers	[ALWZ19]
$2^{[n]}$	ℓ	$\lfloor np \rfloor$-slice	$(\log \ell \cdot \log(1/\varepsilon)/p)^\ell$	Slice sunflowers	[Rao20]
Cliques	$\binom{\ell}{2}$	$G(n,p)$	$(\ell \cdot \log(1/\varepsilon))^\ell / p^{\binom{\ell}{2}}$	Clique sunflowers	[CKR20]

5.6 Applying the general construction

When applying the general construction, sometimes we need extra assumptions on the positive distribution $\mathcal{D}^+$ and the ambient family $\mathcal{X}$. We say that $\mathcal{D}^+$ is $(p_1, p_2, \ldots, p_w)$-*bounded* if there exists a sequence $(p_1, p_2, \ldots, p_w)$ such that $\Pr[X \subseteq \mathbf{Y}] \leqslant p_\ell$ for all $X \in \mathcal{X}_\ell$ and $\ell \in [w]$. We say that $(\mathcal{D}^+, \mathcal{X})$ is $(p_1, \ldots, p_w)$-*viable* if the following holds:

1. $\mathcal{D}^+$ is $(p_1, p_2, \ldots, p_w)$-bounded;
2. $\mathrm{supp}(\mathcal{D}^+) \subseteq \mathcal{X}$;
3. $\mathcal{X}$ is closed under intersections;
4. For every $X \in \mathcal{X}_i$ and $Y \in \mathcal{X}_j$, there exists $Z \in \mathcal{X}_{i+j}$ such that $X \cup Y \subseteq Z$.

This will hold for all of our applications of the approximation method. From the minterm bound (Lemma 5.5.4) of the previous section, we immediately get the following lemma by the union bound.

Lemma 5.6.1. *Let $\mathcal{F}$ be a $(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$-closed family. Suppose that $(\mathcal{D}^+, \mathcal{X})$ is $(p_1, p_2, \ldots, p_w)$-viable. For every $\ell \in [w]$, we have*

$$\Pr[\exists F \in \mathcal{M}_{\mathcal{X},\ell}(\mathcal{F}) : F \subseteq \mathbf{Y}] \leqslant \mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon) \cdot p_\ell.$$

We can use this bound to show that $\mathcal{A} = \mathcal{A}(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$ is inaccurate and to bound the errors on the positive distribution, with an extra assumption on the value of $\mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon) \cdot p_\ell$.

Lemma 5.6.2. *Suppose that $(\mathcal{D}^+, \mathcal{X})$ is $(p_1, p_2, \ldots, p_w)$-viable and that $\sum_{\ell=1}^w \mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon) \cdot p_\ell \leqslant 1/2$. Then $\mathcal{A}$ is inaccurate for $(\mathcal{D}^+, \mathcal{D}^-)$.*

Proof. Let $g \in \mathcal{A}$. There exists a $(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$-closed family $\mathcal{F} \subseteq \mathcal{X}_{\leqslant w}$ such that $g = \lceil \mathcal{F} \rceil$. If $\lceil \mathcal{F} \rceil(\mathbf{Y}) = 1$, there exists $\ell \in [w]$ and $F \in \mathcal{M}_{\mathcal{X},\ell}(\mathcal{F})$ such that $F \subseteq \mathbf{Y}$. Therefore, we get by the union bound and Lemma 5.6.1 that

$$\Pr_{\mathbf{Y} \sim \mathcal{D}^+}[g(\mathbf{Y}) = 1] \leqslant \sum_{\ell=1}^{w} \mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon) \cdot p_\ell \leqslant 1/2.$$

Thus, we get

$$\Pr[g(\mathbf{Y}) = 1] + \Pr[g(\mathbf{N}) = 0] \leqslant 3/2. \qquad \square$$

Lemma 5.6.3. *Suppose that $(\mathcal{D}^+, \mathcal{X})$ is $(p_1, p_2, \ldots, p_w)$-viable. For every $f, h \in \mathcal{A}$, we have*

$$\Pr_{\mathbf{Y} \sim \mathcal{D}^+}[(f \wedge h)(\mathbf{Y}) = 1 \text{ and } (f \sqcap h)(\mathbf{Y}) = 0] \leqslant 2 \sum_{\ell=w/2}^{w} \mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon) p_\ell.$$

Proof. Let $\mathcal{F}, \mathcal{H} \subseteq \mathcal{X}_{\leqslant w}$ be closed families such that $f = \lceil \mathcal{F} \rceil$ and $h = \lceil \mathcal{H} \rceil$. Let us suppose that $(f \wedge g)(\mathbf{Y}) = 1$ and $(f \sqcap g)(\mathbf{Y}) = 0$. By definition, there exists $i, j \in [w]$ such that $F \in \mathcal{M}_{\mathcal{X},i}(\mathcal{F})$, $H \in \mathcal{M}_{\mathcal{X},j}(\mathcal{F})$, and $F \cup H \subseteq \mathbf{Y}$, but there does not exist $S \in \mathcal{F} \cap \mathcal{H}$ such that $S \subseteq \mathbf{Y}$.

Let now $J = \bigcap \{X \in \mathcal{X} : F \cup H \subseteq X\}$. Note that $F \cup H \subseteq J \subseteq \mathbf{Y}$, because $\mathbf{Y} \in \mathcal{X}$. Furthemore, $J \subseteq Z$ for some $Z \in \mathcal{X}_{i+j}$. Because $\mathcal{X}$ is an ambient family, we thus have that $J \in \mathcal{X}_t$ for some $t \leqslant i + j$. Therefore, we have $i + j > w$, otherwise we would have $J \in \mathcal{X}_{\leqslant w}$, which would imply $J \in \mathcal{F} \cap \mathcal{H}$ by Lemma 5.4.3, item (3).

We conclude that $\min\{i, j\} \geqslant w/2$, which implies that at least one of F and H is contained in $\bigcup_{\ell=w/2}^{w} \mathcal{X}_\ell$. Applying the bound of Lemma 5.6.1 and the union bound, we get

$$\begin{aligned}\Pr_{\mathbf{Y} \sim \mathcal{D}^+}[(f \wedge h)(\mathbf{Y}) = 1 \text{ and } (f \sqcap h)(\mathbf{Y}) = 0] &\leqslant \sum_{\ell=w/2}^{w} p_\ell \cdot (|\mathcal{M}_{\mathcal{X},\ell}(\mathcal{F})| + |\mathcal{M}_{\mathcal{X},\ell}(\mathcal{H})|) \\ &\leqslant \sum_{\ell=w/2}^{w} p_\ell \cdot 2\,\mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon). \qquad \square\end{aligned}$$

Theorem 5.6.4. *Let $(\mathcal{D}^+, \mathcal{D}^-)$ be a pair of test distributions for $f_{\mathsf{T}} \in \mathsf{MonBool}(\Gamma)$. Let $(\mathcal{X}, w)$ be an ambient family. Suppose that $(\mathcal{D}^+, \mathcal{X})$ is $(p_1, p_2, \ldots, p_w)$-viable. Suppose moreover that*

1. $\Pr_{\mathbf{N} \sim \mathcal{D}^-}[\gamma \in \mathbf{N}] \leqslant 1 - \varepsilon$ *for all* $\gamma \in \Gamma$;
2. $\sum_{\ell=1}^{w} \mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon) p_\ell \leqslant 1/2$;
3. $\sum_{\ell=w/2}^{w} \mathrm{Ex}(\ell, \mathcal{X}, \mathcal{D}^-, \varepsilon) p_\ell \leqslant 1/(40S)$;
4. $\varepsilon \leqslant 1/(|\mathcal{X}_{\leqslant w}| \cdot 20S)$.

We have $\mathsf{C}_{\mathsf{mon}}(f_{\mathsf{T}}) \geqslant S$.

Proof. Let $\mathcal{A} = \mathcal{A}(\mathcal{X}_{\leqslant w}, \mathcal{D}^-, \varepsilon)$. Our assumptions imply the following:

- $\mathcal{A}$ is a legitimate model (Proposition 5.4.2);
- $\mathcal{A}$ is inaccurate on $(\mathcal{D}^+, \mathcal{D}^-)$ (Lemma 5.6.2);
- $\mathcal{A}$ is a gate-by-gate $(1/20S)$-approximator for $(\mathcal{D}^+, \mathcal{D}^-)$ (Lemmas 5.6.3 and 5.4.6 and Remark 5.4.4).

Therefore, Theorem 5.3.5 implies that the monotone circuit complexity of f_{T} is at least S. $\square$

5.6.1 Monotone circuit lower bounds from abstract sunflower bounds

The following corollary specializes the theorem above to the case where $\Gamma = [n]$, giving a lower bound on the monotone circuit complexity of the target function directly from a sunflower bound.

Corollary 5.6.5. *Let $f_{\mathsf{T}} \in \mathsf{MonBool}([n])$. Let $w = w(n)$ be such that $w \to \infty$. Let $(\mathcal{D}^+, \mathcal{D}^-)$ be a pair of test distributions for f_{T} and suppose $\mathcal{D}^+$ is $(p_1, p_2, \ldots, p_w)$-bounded. Suppose moreover that*

1. $\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\gamma \in \mathbf{N}] \leqslant 1 - \varepsilon$ *for all* $\gamma \in [n]$;
2. $(\mathrm{Ex}(\ell+1, \mathcal{D}^-, \varepsilon)p_{\ell+1})/(\mathrm{Ex}(\ell, \mathcal{D}^-, \varepsilon)p_\ell) \leqslant 1/3$ *for all* $\ell \in [w-1]$;
3. $\mathrm{Ex}(1, \mathcal{D}^-, \varepsilon)p_1 \leqslant 1/3$;
4. $\varepsilon \leqslant 1/\Big(\binom{n}{\leqslant w} \cdot 20S\Big)$.

We have $\mathsf{C}_{\mathsf{mon}}(f_{\mathsf{T}}) = \Omega\Big(\min\big\{S, (\mathrm{Ex}(w/2, \mathcal{D}^-, \varepsilon) \cdot p_{w/2})^{-1}\big\}\Big)$.

Proof. Set $\Gamma = [n]$ and let $\mathcal{X}$ be the trivial ambient family. We have $\varepsilon \leqslant 1/(|\mathcal{X}_{\leqslant w}| \cdot 20S)$. Finally, we have

$$\sum_{\ell=1}^{w} \mathrm{Ex}(\ell, \mathcal{D}^-, \varepsilon)p_\ell \leqslant \sum_{\ell=1}^{w} 3^{-w} \leqslant 1/2,$$

and

$$\sum_{\ell=w/2}^{w} \mathrm{Ex}(\ell, \mathcal{D}^-, \varepsilon)p_\ell \leqslant \mathrm{Ex}(w/2, \mathcal{D}^-, \varepsilon)p_{w/2} \sum_{\ell=1}^{\infty} 3^{-\ell} = O(\mathrm{Ex}(w/2, \mathcal{D}^-, \varepsilon)p_{w/2}).$$

The result now follows from Theorem 5.6.4, by checking that $(\mathcal{D}^+, \mathcal{X})$ is $(p_1, p_2, \ldots, p_w)$-viable. □

This result can be used to give an approach to obtaining strongly exponential lower bounds for a monotone Boolean function. We understand that this is a "speculative" result, and there may not exist any monotone Boolean function for which the following conditions hold.

Corollary 5.6.6. *Let $\delta > 0$ be an absolute constant and set $w := \delta n$ and $\varepsilon = 2^{-2n}$. Let $f_{\mathsf{T}} \in \mathsf{MonBool}([n])$. Let $(\mathcal{D}^+, \mathcal{D}^-)$ be a pair of test distributions for f_{T} and suppose $\mathcal{D}^+$ is $(p_1, p_2, \ldots, p_w)$-bounded. Suppose moreover that*

1. $\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\gamma \in \mathbf{N}] \leqslant 1 - \varepsilon$ *for all* $\gamma \in [n]$;
2. $\mathrm{Ex}(\ell, \mathcal{D}^-, \varepsilon)p_\ell \leqslant 3^{-\ell}$ *for all* $\ell \in [w]$.

We have $\mathsf{C}_{\mathsf{mon}}(f_{\mathsf{T}}) = 2^{\Omega(n)}$.

Chapter 6

An improved monotone circuit lower bound for a problem in NP

6.1 Introduction

The quest for obtaining lower bounds on the size of monotone Boolean circuits began in the early years of circuit complexity. To the best of our knowledge, the first lower bound for monotone circuits on n inputs was $2n$ by Bloniarz [Blo80], further improved to $4n$ by Tiekenheinrich [Tie84], for Boolean functions based on Th_n^k. The first superlinear lower bound was given in a breakthrough by Razborov [Raz85b], who proved a $n^{\Omega(\log n)}$ lower bound on the size of monotone circuits computing the $\mathsf{Clique}(n,k)$ function for $k \leqslant \log n$ with the approximation method. Further progress in this line of work included the results of Andreev [And85] who proved an exponential lower bound for another Boolean function in **NP** based on polynomials. Alon and Boppana [AB87] extended Razborov's result by proving an $n^{\Omega(\sqrt{k})}$ lower bound for $\mathsf{Clique}(n,k)$ for all $k \leq n^{2/3-o(1)}$. These state of art monotone circuit lower bounds saw a further quantitative improvement in a work of Harnik and Raz [HR00] who proved a lower bound of $2^{\Omega((n/\log n)^{1/3})}$ for an n-variate function in **NP** defined using a small probability space of random variables with bounded independence. However, to this day, the question of proving truly exponential lower bounds for monotone circuits (of the form $2^{\Omega(n)}$) for an explicit n-variate function remains open. Truly exponential lower bounds for monotone *formulas* were obtained only recently [PR17].

In a recent work of the author, together with Benjamin Rossman and Mrinal Kumar, we were able to improve the best known lower bound for monotone circuits by proving the first $2^{\Omega(n^{1/2}/(\log n)^{3/2})}$ lower bound for an explicit monotone Boolean function [CKR20]. The function is based on the same construction first considered by Harnik and Raz [HR00], but our argument employs the approximation method of Razborov with recent improvements on robust sunflower bounds (See Section 3.3 and Chapter 4). The following table summarizes the progress of monotone circuit lower bounds so far.

Reference	Boolean function	Technique	Result
[Blo80]	$\mathsf{Majority}$	Gate elimination (?)	$4n$
[Tie84]	Based on Th_n^k[1]	Gate elimination (?)	$4n$
[Raz85b]	$\mathsf{Clique}(n,k)$	Appr. method w/ sunflowers	$n^{\Omega(\log n)}$
[And85]	$\mathsf{Poly}(q,c)$[2]	Appr. method w/ lop. sunflowers	$2^{\Omega(n^{1/8-o(1)})}$
[AB87]	$\mathsf{Poly}(q,c)$	Appr. method w/ lop. sunflowers	$2^{\Omega(n^{1/4-o(1)})}$
[HR00]	$\mathsf{HR}(c,k)$	Monotone switching lemma	$2^{\Omega(n^{1/3-o(1)})}$
[CKR20]	$\mathsf{HR}(c,k)$	Appr. method w/ rob. sunflowers	$2^{\Omega(n^{1/2-o(1)})}$

[1] The function (on $n+1$ variables) is $\mathsf{Th}_n^{n-1} \vee (x_{n+1} \wedge \mathsf{Th}_n^2)$.

[2] The Boolean function $\mathsf{Poly}(q,c)$ receives as an input a subset of G of $[q] \times [q]$, where q is a prime power, and outputs 1 if there exists a polynomial $P \in \mathbb{F}_q[x]$ of degree at most $c-1$ such that $\{(i,P(i)) : i \in [q]\} \subseteq G$.

In what follows, we will present the result of [CKR20], which provides an improved monotone circuit size lower bound for the Boolean function of Harnik and Raz [HR00]. We will apply the approximation method as described in Chapter 5, applying the slice sunflower lemma of Rao (Theorem 3.4.3), whose proof is give in Chapter 4. Our application of the approximation method will be quite similar to that of [AB87] for the $\mathsf{Poly}(q, c)$ function, with the crucial difference that we will employ the slice sunflower lemma instead of lopsided sunflowers for the p-biased distribution. Let us first define their function.

6.2 The Boolean function of Harnik and Raz

Throughout this chapter, we will suppose that n is a prime power. Moreover, we fix two positive integers c and k with $c < k \ll n$. The Harnik-Raz function $\mathsf{HR}(c, k) : 2^{[n]} \to \{0, 1\}$ is defined as follows. For a polynomial $P \in \mathbb{F}_n[x]$, we let S_P be the set of the valuations of P in each element of $\{1, 2, \ldots, k\}$ (in other words, $S_P = \{P(1), \ldots, P(k)\}$). Observe that it is not necessarily the case that $|S_P| = k$, since it may happen that $P(i) = P(j)$ for some i, j such that $i \neq j$. Finally, we consider the family $\mathcal{F}(c, k)$ defined as

$$\mathcal{F}(c, k) := \{S_P : P \in \mathbb{F}_n[x], P \text{ has degree at most } c - 1 \text{ and } |S_P| \geqslant k/2\}.$$

We thus define $\mathsf{HR}(c, k)$ as

$$\mathsf{HR}(c, k) := \lceil \mathcal{F}(c, k) \rceil.$$

We now explain the choice of $\mathcal{F}(c, k)$. First, the choice for valuations of polynomials with degree at most $c - 1$ is explained by a fact observed in [ABI86], which we discuss with more details in Section 2.2.6. If a polynomial $\mathbf{P} \in \mathbb{F}_n[x]$ with degree $c - 1$ is chosen uniformly at random, they observed that the random variables $\mathbf{P}(1), \ldots, \mathbf{P}(k)$ are c-wise independent, and are each uniform in $[n]$. This allows us to define a distribution on the inputs (the positive test distribution) that has high agreement with $\mathsf{HR}(c, k)$ and is easy to analyze. Observe further that, since $|\mathcal{F}(c, k)| \leqslant n^c$, the monotone complexity of $\mathsf{HR}(c, k)$ is at most $2^{c \log n}$. Later we will chose c to be roughly $n^{1/2}$, and prove that the monotone complexity of $\mathsf{HR}(c, k)$ is $2^{\Omega(c)}$.

Finally, the restriction $|S_P| \geqslant k/2$ is a truncation made to ensure that no minterm of $\mathsf{HR}(c, k)$ is very small. Otherwise, if $\mathsf{HR}(c, k)$ had small minterms, it might have been a function that almost always outputs 1. Such functions have very few maxterms and are therefore computed by a small CNF. Since we desire $\mathsf{HR}(c, k)$ to have high complexity, this is an undesirable property. The fact that $\mathsf{HR}(c, k)$ doesn't have small minterms is important in the proof that $\mathsf{HR}(c, k)$ almost surely outputs 0 in the negative test distribution (Lemma 6.3.2).

6.3 Test distributions

We now define the positive and negative test distributions. Let $\mathcal{D}^+$ be the distribution supported in $2^{[n]}$ which chooses a polynomial $\mathbf{P} \in \mathbb{F}_n[x]$ with degree at most $c - 1$ uniformly at random, and maps it into the set $S_{\mathbf{P}}$. Suppose that $T > 2$ is a large absolute constant, to be defined later. Let

$$p := n^{-Tc/k} \quad \text{and} \quad M := \lfloor np \rfloor.$$

Let also $\mathcal{D}^-$ be the M-uniform distribution in $[n]$. We will always denote a random sample from $\mathcal{D}^+$ by $\mathbf{Y}$ and a random sample from $\mathcal{D}^-$ by $\mathbf{N}$.

Harnik and Raz proved that $\mathsf{HR}(c, k)$ outputs 1 on $\mathcal{D}^+$ with high probability.

Lemma 6.3.1 (Claim 4.2 in [HR00])**.** *We have* $\Pr_{\mathbf{Y} \sim \mathcal{D}^+}[\mathsf{HR}(c, k)(\mathbf{Y}) = 1] \geqslant 1 - (k - 1)/n$.

Proof. Let $\mathbf{P}$ be the polynomial randomly chosen by $\mathcal{D}^+$. Call a pair $\{i, j\} \subseteq [k]$ with $i \neq j$ *coinciding* if $\mathbf{P}(i) = \mathbf{P}(j)$. Because the random variables $\mathbf{P}(i)$ and $\mathbf{P}(j)$ are uniformly distributed

in $[n]$ and independent for $i \neq j$, we have that $\Pr[\mathbf{P}(i) = \mathbf{P}(j)] = 1/n$ for $i \neq j$. Therefore, the expected number $\mathsf{Num}(\mathbf{P})$ of coiciding pairs is $\binom{k}{2}/n$. Observe now that $\mathsf{HR}(c,k)(\mathbf{Y}) = 0$ if and only if $|\mathbf{P}(1), \ldots, \mathbf{P}(k)| < k/2$, which occurs only if there exists more than $k/2$ coinciding pairs. Therefore, by Markov's inequality (Proposition 2.2.2), we have

$$\Pr\left[\mathsf{HR}(c,k)(\mathbf{Y}) = 0\right] \leqslant \Pr\left[\mathsf{Num}(\mathbf{P}) > k/2\right] \leqslant \frac{\binom{k}{2}/n}{k/2} = \frac{k-1}{n}. \qquad \square$$

We now claim that $\mathsf{HR}(c,k)$ also outputs 0 on $\mathbf{N}$ with high probability.

Lemma 6.3.2. *We have* $\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\mathsf{HR}(c,k)(\mathbf{N}) = 0] \geqslant 1 - n^{-(T/2-1)c}$.

Proof. By definition, we have that $\mathsf{HR}(c,k)(\mathbf{N}) = 1$ only if there exists $F \in \mathcal{F}(c,k)$ such that $F \subseteq \mathbf{N}$. Moreover, we have $\Pr[F \subseteq \mathbf{N}] = \binom{n-|F|}{M-|F|}/\binom{n}{M}$. Observing that $|F| \geqslant k/2$ for all $F \in \mathcal{F}(c,k)$ and $|\mathcal{F}(c,k)| \leqslant n^c$, we get by the union bound that

$$\Pr[\mathsf{HR}(c,k)(\mathbf{N}) = 1] \leqslant n^c \cdot \frac{\binom{n-k/2}{M-k/2}}{\binom{n}{M}} \leqslant n^c \cdot \left(\frac{M}{n}\right)^{k/2} \leqslant n^{-(T/2-1)c}. \qquad \square$$

As a consequence of Lemmas 6.3.1 and 6.3.2, we obtain that pair $\mathcal{D}^+$ and $\mathcal{D}^-$ are test distributions for $\mathsf{HR}(c,k)$ (because $T > 2$ by assumption).

Lemma 6.3.3. *For large enough n, $(\mathcal{D}^+, \mathcal{D}^-)$ is pair of test distributions for $\mathsf{HR}(c,k)$.*

The following property about $\mathcal{D}^+$ will be important when applying the approximation method in the next section.

Lemma 6.3.4. *For every $A \subseteq [n]$ such that and $|A| = \ell \leqslant c$, we have*

$$\Pr_{\mathbf{Y}\sim\mathcal{D}^+}[A \subseteq \mathbf{Y}] \leqslant (k/n)^\ell.$$

Proof. Let $\mathbf{P}$ be the polynomial with degree at most $c-1$ chosen by $\mathcal{D}^+$ and suppose that $A \subseteq \mathbf{Y}$. By definition, we have $A \subseteq S_{\mathbf{P}} = \{\mathbf{P}(1), \mathbf{P}(2), \ldots, \mathbf{P}(k)\}$. Therefore, there exists indices $\{j_1, \ldots, j_\ell\}$ such that $\{\mathbf{P}(j_1), \mathbf{P}(j_2), \ldots, \mathbf{P}(j_\ell)\} = A$. Since $\ell \leqslant c$, we get by the c-wise independence of $\mathbf{P}(1), \ldots, \mathbf{P}(k)$ that the random variables $\mathbf{P}(j_1), \mathbf{P}(j_2), \ldots, \mathbf{P}(j_\ell)$ are independent. Because the random variables $\mathbf{P}(i)$ are uniformly distibuted in $[n]$, it follows that

$$\Pr[\{\mathbf{P}(j_1), \mathbf{P}(j_2), \ldots, \mathbf{P}(j_\ell)\} = A] = \frac{\ell!}{n^\ell}.$$

Therefore, we have

$$\Pr[A \subseteq \mathbf{Y}] = \Pr[A \subseteq \{\mathbf{P}(1), \mathbf{P}(2), \ldots, \mathbf{P}(k)\}] \leqslant \binom{k}{\ell}\frac{\ell!}{n^\ell} \leqslant \left(\frac{k}{n}\right)^\ell. \qquad \square$$

6.4 Applying the approximation method

We will here apply the general construction of legitimate models of Section 5.4 to the Harnik-Raz function. We begin with choosing the parameters c and k. Choose B as in Theorem 3.4.3. Let $T = 12B \cdot \log 2$. We define

$$k := \left(\frac{n}{\log n}\right)^{1/2}, \qquad c = \frac{\log 2}{T} \cdot \frac{k}{\log n} = \frac{1}{12B}\left(\frac{n}{(\log n)^3}\right)^{1/2}.$$

This implies $c = \Theta(k/\log n) \ll k$ and $p = n^{-Tc/k} = 1/2$.

We now choose the parameters to specialize the generalized construction of Section 5.4. Observe that $\Gamma = [n]$ in our context. Let us fix

$$\varepsilon := n^{-2c}.$$

Let $(\mathcal{X}, c)$ be the trivial ambient family. In other words, let $\mathcal{X}_\ell = \binom{[n]}{\ell}$, for $\ell \in [n]$, and $w = c$. Note that $\mathcal{X} = \bigcup \mathcal{X}_\ell = 2^{[n]}$, and $\mathcal{X}_{\leqslant w} = \binom{[n]}{\leqslant c}$. With this definition for $\mathcal{X}$, a family $\mathcal{F} \subseteq 2^{[n]}$ will be $(\mathcal{X}_{\leqslant c}, \mathcal{D}^-, \varepsilon)$-closed if, for every $X \subseteq [n]$ such that $|X| \leqslant c$, we have

$$\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\mathbf{N} \cup X \text{ is independent in } \mathcal{F}] < \varepsilon \Longrightarrow X \in \mathcal{F}.$$

In the rest of this chapter, we will write *closed* as a shorthand for $(\mathcal{X}_{\leqslant c}, \mathcal{D}^-, \varepsilon)$-closed. We now let $\mathcal{A} = \mathcal{A}(\mathcal{X}_{\leqslant c}, \mathcal{D}^-, \varepsilon)$ be the subposet of $\mathsf{MonBool}([n])$ determined by $\mathcal{X}, c, \mathcal{D}^-$ and ε (Definition 5.4.1). Let us check that $\mathcal{A}$ is a legitimate model (Definition 5.3.1).

Lemma 6.4.1. *The set of approximators $\langle \mathcal{A}, \sqcup, \sqcap \rangle$ is a legitimate model in $\mathsf{MonBool}([n])$.*

Proof. Because of Proposition 5.4.2, it suffices to check that $\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\gamma \in \mathbf{N}] \leqslant 1 - \varepsilon$ for all $\gamma \in [n]$. But $\Pr_{\mathbf{N}\sim\mathcal{D}^-}[\gamma \in \mathbf{N}] = \binom{n-1}{M-1}/\binom{n}{M} = M/n \leqslant p = 1/2 < 1 - \varepsilon$. □

We now obtain a bound on the approximation errors on the negative distribution directly from Lemma 5.4.6. (Observe that $|\mathcal{X}_{\leqslant c}| \leqslant n^c$, so $\varepsilon \cdot |\mathcal{X}_{\leqslant c}| \leqslant n^{-c}$.)

Lemma 6.4.2. *For every $f, g \in \mathcal{A}$, we have*

$$\Pr_{\mathbf{N}\sim\mathcal{D}^-}[(f \vee g)(\mathbf{N}) = 0 \text{ and } (f \sqcup g)(\mathbf{N}) = 1] \leqslant n^{-c}.$$

6.4.1 Applying sunflower bounds

With our choice of parameters for c and k, we can use the slice sunflower result of Rao (Theorem 3.4.3) to give a bound for $\mathrm{Ex}(\ell, \mathcal{D}^-, \varepsilon)$, the extremal number of the abstract sunflower given by the trivial ambient family and $\mathcal{D}^-$ and ε (see Definition 5.5.1). With this bound, we can bound the probability that a sample $\mathbf{Y}$ of the positive distribution $\mathcal{D}^+$ contains a minimal element of a closed family $\mathcal{F} \subseteq \binom{[n]}{\leqslant c}$. In what follows, we let $\mathcal{M}_\ell(\mathcal{F})$ denote the set of minimal elements of $\mathcal{F}$ of size ℓ. Because $s(\mathcal{X}, \ell) = \ell$ and $\mathcal{X} = 2^{[n]}$, we have $\mathcal{M}_{\mathcal{X},\ell}(\mathcal{F}) = \mathcal{M}_\ell(\mathcal{F})$.

Lemma 6.4.3. *Let $\mathcal{F} \subseteq 2^{[n]}$ be a closed family. For every $\ell \in [c]$, we have*

$$\Pr[\exists F \in \mathcal{M}_\ell(\mathcal{F}) : F \subseteq \mathbf{Y}] \leqslant 3^{-\ell},$$

when n is sufficiently large.

Proof. Since $\mathcal{F}$ is closed, we can use the generic bound of Lemma 5.5.4 to get the upper bound $|\mathcal{M}_\ell(\mathcal{F})| \leqslant \mathrm{Ex}(\ell, \mathcal{D}^-, \varepsilon)$. Since $\mathcal{D}^-$ is a subset of $[n]$ of size $M = \lfloor n/2 \rfloor$ chosen uniformly at random, we get that a (M, ε)-slice sunflower (Definition 3.4.2) is also a $(\mathcal{X}, \mathcal{D}^-, \varepsilon)$-sunflower (Definition 5.5.1). Thus, we get from Theorem 3.4.3 that

$$|\mathcal{M}_\ell(\mathcal{F})| \leqslant \mathrm{Ex}(\ell, \mathcal{D}^-, \varepsilon) \leqslant \left(2B \log(c/\varepsilon) \log\left(2 \log(c/\varepsilon)\right)\right)^\ell.$$

Now observe that

$$2 \log(c/\varepsilon) = 2 \log(n^{3c} c) \leqslant 2 \log(n^{4c}) = 8c \cdot \log n.$$

Moreover, we have

$$\log\left(2 \log(c/\varepsilon)\right) = \log(8c \cdot \log n) = \frac{1}{2} \log n - \frac{1}{2} \log\log n + O(1) \leqslant \frac{1}{2} \log n,$$

for n sufficiently large. From the previous inequalities, we obtain for n sufficiently large that

$$|\mathcal{M}_\ell(\mathcal{F})| \leqslant \left(4B \cdot c(\log n)^2\right)^\ell = (n/3k)^\ell.$$

We now get from Lemma 6.3.4 and the union bound that

$$\Pr[\exists F \in \mathcal{M}_\ell(\mathcal{F}) : F \subseteq \mathbf{Y}] \leqslant (n/3k)^\ell \cdot (k/n)^\ell = 3^{-\ell}. \qquad \square$$

Equipped with this bound, we can now show that $\mathcal{A}$ is inaccurate on $(\mathcal{D}^+, \mathcal{D}^-)$ and bound the approximation error on $\mathcal{D}^+$.

Lemma 6.4.4. *For large enough n, $\mathcal{A}$ is inaccurate on $(\mathcal{D}^+, \mathcal{D}^-)$.*

Proof. Let $f \in \mathcal{A}$. Since $f \in \mathcal{A}$, there exists a closed family $\mathcal{F} \subseteq \mathcal{X}_{\leqslant c}$ such that $f = \lceil \mathcal{F} \rceil$. If $\lceil \mathcal{F} \rceil(\mathbf{Y}) = 1$, there exists $\ell \in [c]$ and a minimal element $F \in \mathcal{M}_\ell(\mathcal{F})$ such that $F \subseteq \mathbf{Y}$. Applying the bound of Lemma 6.4.3 and the union bound, we get

$$\Pr\left[f(\mathbf{Y}) = 1\right] \leqslant \sum_{\ell=1}^{c} 3^{-\ell} \leqslant 1/2.$$

Therefore, we have $\Pr[f(\mathbf{Y}) = 1] + \Pr[f(\mathbf{N}) = 0] \leqslant 3/2$. $\square$

Lemma 6.4.5. *For every $f, g \in \mathcal{A}$, we have*

$$\Pr_{\mathbf{Y} \sim \mathcal{D}^+}\left[(f \wedge g)(\mathbf{Y}) = 1 \text{ and } (f \sqcap g)(\mathbf{Y}) = 0\right] \leqslant 2^{-\Omega(c)}.$$

Proof. Let $\mathcal{F}, \mathcal{H} \subseteq \mathcal{X}_{\leqslant c}$ be closed families such that $f = \lceil \mathcal{F} \rceil$ and $g = \lceil \mathcal{H} \rceil$. Let us suppose that $(f \wedge g)(\mathbf{Y}) = 1$ and $(f \sqcap g)(\mathbf{Y}) = 0$. By definition, there exists $F \in \mathcal{M}(\mathcal{F})$ and $H \in \mathcal{M}(\mathcal{H})$ such that $F \cup H \subseteq \mathbf{Y}$, but there does not exist $S \in \mathcal{F} \cap \mathcal{H}$ such that $S \subseteq \mathbf{Y}$.

If $|F \cup H| \leqslant c$, then $F \cup H \in \mathcal{X}_{\leqslant c}$. But, since $\mathcal{F}$ and $\mathcal{H}$ are closed, we get by item (3) of Lemma 5.4.3 that $F \cup H \in \mathcal{F} \cap \mathcal{H}$, a contradiction. So we have $|F \cup H| > c$. Therefore, at least one of F and H has size at least $c/2$. Applying the bound of Lemma 6.4.3 and the union bound, we get

$$\Pr\left[(f \wedge h)(\mathbf{Y}) = 1 \text{ and } (f \sqcap h)(\mathbf{Y}) = 0\right] \leqslant 2 \sum_{\ell=c/2}^{c} 3^{-\ell} = 2^{-\Omega(c)}. \qquad \square$$

6.4.2 Wrapping up

So far, we proved

- $(\mathcal{D}^+, \mathcal{D}^-)$ is a pair of test distributions for $\mathsf{HR}(c, k)$ (Lemma 6.3.3);
- $\mathcal{A}$ is a legitimate model (Lemma 6.4.1);
- $\mathcal{A}$ is inaccurate on $(\mathcal{D}^+, \mathcal{D}^-)$ (Lemma 6.4.4);
- $\mathcal{A}$ is a gate-by-gate $2^{-\Omega(c)}$-approximator for $(\mathcal{D}^+, \mathcal{D}^-)$ (Lemmas 6.4.5 and 6.4.2 and Remark 5.4.4);

Therefore, Theorem 5.3.5 implies that the monotone circuit complexity of $\mathsf{HR}(c, k)$ is $2^{\Omega(c)}$.

Theorem 6.4.6. *Any monotone circuit computing $\mathsf{HR}(c, k)$ has size $2^{\Omega(c)} = 2^{\Omega(n^{1/2}/(\log n)^3)}$.*

6.5 Generalized Harnik-Raz function

The lower bound given here for $\mathsf{HR}(c,k)$ is actually generalizable to a larger class of monotone Boolean functions. This generalized class is the one actually considered in [HR00]. The generalization is as follows. Let Ω be a probability space of k c-wise independent uniformly distributed random variables. More formally, let $\Omega \subseteq [n]^k$ be such that, when $\boldsymbol{\omega} \sim \mathcal{U}(\Omega)$, the random variables $\boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_k$ are c-wise independent and uniformly distributed in $[n]$. For $\omega \in \Omega$, we let $S_\omega := \{\omega_1, \ldots, \omega_k\}$, and $\mathcal{F}(c,k) := \{S_\omega : \omega \in \Omega, |S_\omega| \geqslant k/2\}$. Finally, the Harnik-Raz function given by Ω is defined as $\mathsf{HR}(c,k) := \lceil \mathcal{F}(c,k) \rceil$. The positive test distribution $\mathcal{D}^+$ is defined by choosing a $\boldsymbol{\omega} \in \Omega$ uniformly at random and returning $S_{\boldsymbol{\omega}}$, and $\mathcal{D}^-$ is defined just like in Section 6.3. The arguments of the rest of the chapter follow in the same way for $\mathsf{HR}(c,k)$.

The particular example of this construction that we considered in Section 6.2 is given by

$$\Omega = \{(P(1), \ldots, P(k)) : P \in \mathbb{F}_n[x],\ P \text{ has degree at most } c-1\}.$$

A proof that this is a probability space of c-wise independent random variables is given in Section 2.2.6.

6.6 Further questions

6.6.1 Strongly exponential lower bounds for monotone circuits

We have proved a lower bound of $\exp(\Omega(n^{1/2-o(1)}))$ on the size of monotone circuits computing a function in **NP**. It is natural to ask if we can do better.

Problem 6.6.1. *Can we find an explicit monotone Boolean function whose monotone circuit complexity is* $\exp(\Omega(n))$*? In particular, can we do this with the approximation method?*

We will now discuss a way to obtain such a lower bound for the Harnik-Raz function itself from better sunflower bounds. In Chapter 5, we gave a generic lower bound for monotone circuits depending on a sunflower bound (Corollary 5.6.5). We did not apply this result directly in this chapter for the sake of clarity, but one can easily check that the proofs of Corollary 5.6.5 and Theorem 6.4.6 are the essentially same, and indeed give the same lower bounds on monotone circuit size. In fact, Lemma 6.3.4 is actually showing that $\mathcal{D}^+$ is $(p_1, \ldots, p_c)$-bounded for the sequence $p_\ell = (k/n)^\ell$.

The statement of Corollary 5.6.5 highlights that the lower bound of Theorem 6.4.6 depends on the bounds we have for $\mathrm{Ex}(c/2, \mathcal{D}^-, \varepsilon)$, and that the tighter the bounds we have for this extremal function, the better are our lower bounds for $\mathsf{HR}(c,k)$. In particular, it shows that the lower bound we obtain for $\mathsf{HR}(c,k)$ with the approximation method is of the form

$$\mathsf{C}_{\mathsf{mon}}(\mathsf{HR}(c,k)) \geqslant \mathrm{Ex}(c/2, \mathcal{D}^-, \varepsilon)^{-1} \cdot (n/k)^{c/2} = 2^{\Omega(c)},$$

where ε is chosen to be n^{-2c} in this chapter, which satisfies the conditions of Corollary 5.6.5.

The inequality above shows that, if we are looking for a truly exponential lower bound, we must choose $c = \Omega(n)$. For concreteness, let us suppose there exists a sufficiently small fixed constant $\delta > 0$ such that $c = \delta k = \delta^2 n$. In this regime, it is not hard to check that we can take $\varepsilon = 2^{-2n}$. Therefore, in order to prove that $\mathsf{C}_{\mathsf{mon}}(\mathsf{HR}(c,k)) = 2^{\Omega(c)}$, it suffices to show that

$$\mathrm{Ex}(c/2, \mathcal{D}^-, 2^{-2n}) \leqslant 2^{-c/2} \cdot (n/k)^{c/2} = 2^{n \cdot \delta^2 \log(1/\delta)/2}.$$

Since $\mathcal{D}^-$ is p-biased with $p = n^{-4c/k}$, a $(\mathcal{D}^-, \varepsilon)$-sunflower is equivalent to a robust lopsided sunflower, and thus we have $\mathrm{Ex}(c/2, \mathcal{D}^-, 2^{-2n}) = f_{\mathrm{RL}}(c/2, p, 2^{-2n})$. The above discussion can thus be summarized in saying that, to show a truly exponential lower bound for monotone circuits, it suffices to solve the following problem.

Problem 6.6.2. *Show that, for some constant $\delta > 0$, we have $f_{\mathrm{RL}}(\delta^2 n/2, n^{-4\delta}, 2^{-2n}) = 2^{n \cdot \delta^2 \log(1/\delta)/2}$, where n is the size of the vertex set.*

The best upper bounds we have for f_{RL} are direct applications of bounds for robust sunflowers (see inequality 3.3). We recall here Remark 3.3.15, which states that the lower bound of Proposition 3.3.13 for $f_{\mathrm{R}}(\ell, p, \varepsilon)$ (the extremal function of robust sunflowers) do not apply when $\ell = \delta n$. In particular, this means that, as far as we know, the bounds of the problems above may hold even for robust sunflowers.

6.6.2 Connections to the monotone switching lemma

The original lower bound of Harnik and Raz [HR00] uses the *monotone switching lemma*, not making use of any sunflower lemma. This proof technique is itself a form of the approximation method, albeit one that is "sunflower-free".

Problem 6.6.3. *Can we prove Theorem 6.4.6 without sunflowers? Or can we prove an improved monotone switching lemma with robust sunflower bounds?*

6.6.3 2-slice distributions

In this chapter, the negative distribution $\mathcal{D}^-$ was a slice distribution, unlike $\mathcal{D}^+$. Can we turn we choose a slice distribution for $\mathcal{D}^+$ and also obtain an exponential monotone circuit lower bound? We would then show that monotone circuits cannot compute $\mathsf{HR}(c, k)$ in a distribution that is supported in only 2 slices. This is significant because, as proved by Berkowitz [Ber82], the monotone and non-monotone circuit complexity of Boolean functions are polynomially related under distributions that are supported in only one slice.

Problem 6.6.4. *Prove an exponential monotone circuit lower bound for $\mathsf{HR}(c, k)$ using a pair of* slice *test distributions.*

Chapter 7

Better bounds for clique

7.1 Introduction

Let $\mathsf{Clique}(n,k) : \mathcal{G}^n \to \{0,1\}$ be the Boolean function which, given a graph G with vertex set $[n]$, outputs 1 if this graph contains a k-clique. This Boolean function is one of the most fundamental in computational complexity, since Karp [Kar72] proved that deciding if a graph G contains a clique is **NP**-complete. In monotone circuit complexity, this function has played a prominent role. It was the first monotone Boolean function for which a superpolynomial lower bound on its monotone circuit complexity was given, when Razborov [Raz85b] showed an $n^{\Omega(\log n)}$ lower bound on $\mathsf{C}_{\mathsf{mon}}(\mathsf{Clique}(n,k))$ for $k \leqslant \log n$. Soon after, Alon and Boppana [AB87] proved $\mathsf{C}_{\mathsf{mon}}(\mathsf{Clique}(n,k)) = n^{\Omega(\sqrt{k})}$ for all $k \leq n^{2/3-o(1)}$. However, this is still far short of the obvious upper bound of $n^{O(k)}$.

Using the language of Chapter 5, Alon and Boppana's result is an application of the approximation method with the following parameters. For a set $A \subseteq [n]$, let K_A denote the graph in $\mathcal{G}^n$ such that $E(K_A) = \binom{A}{2}$. (In other words, K_A contains a clique on A and no other edges.) We say that K_A is a *planted clique*. Let $\mathcal{D}^+$ denote the distribution which samples $\mathbf{A} \subseteq_k [n]$ and returns $K_{\mathbf{A}}$. Let $\mathcal{D}^-$ be the distribution which samples a function $\mathbf{c} : [n] \to [k-1]$ uniformly at random, and maps into the n-vertex graph $G_{\mathbf{c}}$ with edge set $\{\{u,v\} : \mathbf{c}(u) \neq \mathbf{c}(v)\}$. Alon and Boppana show a lower bound of $n^{\Omega(\sqrt{k})}$ on the size of monotone circuits computing $\mathsf{Clique}(n,k)$ on the pair of test distributions $(\mathcal{D}^+, \mathcal{D}^-)$, when $k \leqslant n^{2/3-o(1)}$.

For many years, Alon and Boppana's result remained the best lower bound for $\mathsf{Clique}(n,k)$ when k is larger than a polynomial in $\log n$. Recently, Krajíček and Oliveira [KO18] proved that, for an extension of monotone circuits called *monotone circuits with local oracles* (monotone CLOs), the complexity of computing $\mathsf{Clique}(n,k)$ on the test distributions $\mathcal{D}^+$ and $\mathcal{D}^-$ considered by Alon and Boppana is $n^{\Theta(\sqrt{k})}$. Moreover, they show that the approximation method of Razborov also applies to monotone CLOs. This means that, if we wish to prove an $n^{\Omega(k)}$ lower bound for $\mathsf{C}_{\mathsf{mon}}(\mathsf{Clique}(n,k))$, we must either consider a different pair of test distributions or avoid the approximation method altogether.

A recent work of the author and his coauthors [CKR20] overcomes this barrier by making a different choice for $\mathcal{D}^-$. We let $\mathcal{D}^-$ be the Erdős-Rényi random graph $G(n,p_k)$, where p_k is a probability function below the threshold for the occurrence of k-cliques. Other than the choice of parameters w and ε, our choice of parameters is the same as that of [AB87] for the legitimate model of Section 5.4. Applying the approximation method, we prove a lower bound of $\Omega(n^{(\delta/8)\cdot k})$ to the monotone circuit complexity of $\mathsf{Clique}(n,k)$, when $k \leqslant n^{1/3-\delta}$, for some absolute constant δ.

The main technical contribution of our result is an upper bound to the extremal function $\mathrm{Ex}(\ell, \mathcal{K}, \mathcal{D}^-, \varepsilon)$ associated with the $(\mathcal{K}, \mathcal{D}^-, \varepsilon)$-sunflower that we call *clique sunflowers* (Definition 7.2.1) – here, $\mathcal{K}$ is the set of all K_A for $A \subseteq [n]$. Since the negative distribution is p-biased, this sunflower-type structure bears great resemblance with robust sunflowers (Definition 3.3.2), and our proof is quite similar to that of Rossman for the latter (see Theorem 3.3.9). Apart from our bound to clique sunflowers, everything else is a standard application of the approximation method.

We summarize the differences between our application of the approximation method and that of [AB87] in the following table. We remark that shorly after [AB87], Boppana and Sipser [BS90] reproved the lower bound of Alon and Boppana for cliques using the standard sunflower of Erdős and Rado (Definition 3.1.1). Though they obtain the same lower bound but for a more restricted range of k, their proof benefits of a simplified analysis, which is more common in textbooks. For this reason, we include it here as well. All the parameters of the generalized construction of Section 5.4 that are omitted here are the same in all the applications.

Ref.	$\mathcal{D}^-$	Range of k	Sunflower bound	Technique	w	ε	L. bound
[AB87]	$G_{\mathbf{c}}$	$k \leqslant n^{2/3}/4$	$(\log(1/\varepsilon) \cdot (\frac{k-1}{k-\ell})^\ell)^\ell$	Lop. sf.	$\sqrt{k}$	$n^{-\sqrt{k}}$	$(\frac{n}{4k^{3/2}})^{\sqrt{k}/3}$
[BS90]	$G_{\mathbf{c}}$	$k \leqslant n^{1/4}$	$\ell!(\log(1/\varepsilon)(\frac{k-1}{k-\ell})^\ell)^\ell$	Sunflowers	$\sqrt{k}$	$n^{-\sqrt{k}}$	$n^{\Omega(\sqrt{k})}$
[CKR20]	$G(n,p)$	$k \leqslant n^{1/3-\delta}$	$\ell! \log(1/\varepsilon)^\ell / p^{\binom{\ell}{2}}$	Clq. sf.	$\delta \cdot k$	n^{-k}	$\Omega(n^{(\delta/8)k})$

In this chapter, we will present the lower bound of [CKR20] for the clique function. We will first present our new notion of "clique-shaped" sunflowers, and then apply the approximation method as explained in Chapter 5.

7.2 Clique sunflowers

Throughout our discussion about sunflowers on Chapter 3, we made no assumption about the set of vertices and the "shape" of the edges of the hypergraph. However, sometimes one can obtain tighter bounds by paying attention to such things. Here we introduce the notion of *clique sunflowers*, which is analogous to that of robust sunflowers for "clique-shaped" set systems.

Recall that the *planted clique* K_A denotes the graph in $\mathcal{G}^n$ such that $E(K_A) = \binom{A}{2}$. Define $\mathcal{K}_\ell = \{K_A : A \in \binom{[n]}{\ell}\}$. Note that $\mathcal{K} = \bigcup \mathcal{K}_\ell$ is the set of all planted cliques in graphs with vertex set $[n]$. We also define $\mathcal{K}_{\leqslant w} = \{K_A : A \subseteq [n],\ |A| \leqslant w\}$, the set of planted cliques of at most w vertices.

Definition 7.2.1. *A hypergraph $\mathcal{F}$ with vertex set $\binom{[n]}{2}$ is called a* (p,ε)-clique sunflower *if $\mathcal{F}$ is a (p,ε)-robust sunflower and $\mathcal{F} \subseteq \mathcal{K}$. Equivalently, this means that*

$$\Pr_{\mathbf{G} \sim G(n,p)} [\forall K_A \in \mathcal{F} : K_A \not\subseteq \mathbf{G} \cup K_B] < \varepsilon,$$

where $K_B = \bigcap \mathcal{F}$.

Though clique sunflowers may seem similar to robust sunflowers, the importance of this definition is that it allows us to explore the "clique-shaped" structure of the sets of the family, and thus obtain an asymptotically better upper bound on the size of sets that do not contain a clique sunflower.

Lemma 7.2.2. *Let $\mathcal{F} \subseteq \mathcal{K}_\ell$ be such that $|\mathcal{F}| \geq \ell!(2\log(1/\varepsilon))^\ell (1/p)^{\binom{\ell}{2}}$. Then $\mathcal{F}$ contains a (p,ε)-clique sunflower.*

Observe that, whereas the bounds for "standard" robust sunflowers (Theorems 3.3.9, 3.3.11, 3.4.3) would give us an exponent of $\binom{\ell}{2}$ on the $\log(1/\varepsilon)$ factor, Lemma 7.2.2 give us only an ℓ at the exponent. This is a very significant difference for our choice of parameters.

We defer the proof of Lemma 7.2.2 to the end of the chapter (Section 7.5). The proof is based on an application of Janson's inequality (Lemma 2.2.4), as in the original robust sunflower lemma of [Ros14] (Theorem 3.3.9).

7.3 Test distributions

Throughout this chapter, let

$$p_k := 4^{-1/\binom{k}{2}} \cdot n^{-2/(k-1)}.$$

Lemma 7.3.1. *Let* $\mathbf{G}^- \sim G(n, p_k)$. *We have* $\Pr[\mathbf{G}^-$ *contains a* k*-clique*$] \leq 1/4$.

Proof. There are $\binom{n}{k} \leq n^k$ potential k-cliques, each present in $\mathbf{G}^-$ with probability $p^{\binom{k}{2}} = n^{-k}/4$. By the union bound, we have $\Pr[\mathbf{G}^-$ contains a k-clique$] \leq n^k \cdot n^{-k}/4 = 1/4$. □

Let $\mathcal{D}^+$ be the distribution which samples $\mathbf{A} \subseteq_k [n]$ and returns $K_{\mathbf{A}}$. We will always denote a sample of $\mathcal{D}^+$ by $\mathbf{K}_k$. Let $\mathcal{D}^-$ be the distribution which samples $G(n, p_k)$. We will always denote a sample of $\mathcal{D}^-$ by $\mathbf{G}^-$. From Lemma 7.3.1, we easily obtain the following lemma.

Lemma 7.3.2. *The pair* $(\mathcal{D}^+, \mathcal{D}^-)$ *is a pair of test distributions for* $\mathsf{Clique}(n, k)$.

The following property about $\mathcal{D}^+$ is easy to check.

Lemma 7.3.3. *For every* $A \subseteq [n]$ *such that and* $|A| = \ell$, *we have*

$$\Pr[A \subseteq \mathbf{K}_k] = \frac{\binom{n-\ell}{k-\ell}}{\binom{n}{k}} \leqslant \left(\frac{k}{n}\right)^{\ell}.$$

7.4 Applying the approximation method

We will here apply the general construction of legitimate models of Section 5.4 to the $\mathsf{Clique}(n, k)$ function. Let us fix an absolute constant $\delta \in (0, 1/3)$, and let us suppose that k is in the range $3 \leq k \leq n^{1/3-\delta}$, Observe that $\Gamma = \binom{[n]}{2}$ in our context. Let us fix

$$\varepsilon := n^{-k} \quad \text{and} \quad w := \delta k.$$

It is easy to check that $(\mathcal{K}, w)$ satisfies all the properties of an ambient family of Section 5.4, with $s(\mathcal{K}, \ell) = \binom{\ell}{2}$. With this choice for the ambient family, a family of graphs $\mathcal{F} \subseteq 2^{\binom{[n]}{2}}$ will be $(\mathcal{K}_{\leqslant w}, G(n, p_k), \varepsilon)$-closed if, for every $A \subseteq [n]$ such that $|A| \leqslant \delta k$, we have

$$\Pr[\forall H \in \mathcal{F} : H \not\subseteq \mathbf{G}^- \cup K_A] < \varepsilon \Longrightarrow K_A \in \mathcal{F}.$$

In the rest of this chapter, we will write *closed* as a shorthand for $(\mathcal{K}_{\leqslant w}, G(n, p_k), \varepsilon)$-closed. We now let $\mathcal{A} = \mathcal{A}(\mathcal{K}_{\leqslant w}, G(n, p_k), \varepsilon)$ be the subposet of $\mathsf{MonBool}\left(\binom{[n]}{2}\right)$ determined by $\mathcal{K}, w, G(n, p_k)$ and ε (Definition 5.4.1). Let us check that $\mathcal{A}$ is a legitimate model (Definition 5.3.1).

Lemma 7.4.1. *The set of approximators* $\langle \mathcal{A}, \sqcup, \sqcap \rangle$ *is a legitimate model in* $\mathsf{MonBool}\left(\binom{[n]}{2}\right)$, *when* n *is sufficiently large.*

Proof. Because of Proposition 5.4.2, it suffices to check that $\Pr[e \in \mathbf{G}^-] \leqslant 1 - \varepsilon$ for all $e \in \binom{[n]}{2}$. But $\Pr[e \in \mathbf{G}^-] = p_k < 1 - \varepsilon$, when n is sufficiently large. □

We now obtain a bound on the approximation errors on the negative distribution directly from Lemma 5.4.6. (Observe that $|\mathcal{K}_{\leqslant w}| \leqslant n^{\delta k}$, so $\varepsilon \cdot |\mathcal{K}_{\leqslant w}| \leqslant n^{-(2/3)k}$.)

Lemma 7.4.2. *For every* $f, g \in \mathcal{A}$, *we have*

$$\Pr\left[(f \vee g)(\mathbf{G}^-) = 0 \text{ and } (f \sqcup g)(\mathbf{G}^-) = 1\right] \leqslant n^{-(2/3)k}.$$

7.4.1 Applying sunflower bounds

We will now use the clique sunflower lemma (Lemma 7.2.2) to give a bound for $\mathrm{Ex}(\ell, \mathcal{K}, \mathcal{D}^-, \varepsilon)$, the extremal number of the abstract sunflower given by $\mathcal{K}$, $\mathcal{D}^-$ and ε (see Definition 5.5.1). With this bound, we can bound the probability that a sample $\mathbf{K}_k$ of the positive distribution $\mathcal{D}^+$ contains a minimal element of a closed family $\mathcal{F}$.

Lemma 7.4.3. *Let $\mathcal{F} \subseteq \mathcal{K}_{\leqslant w}$ be a closed family. For every $\ell \in [w]$, we have*

$$\Pr[\exists F \in \mathcal{M}_{\mathcal{K},\ell}(\mathcal{F}) : F \subseteq \mathbf{K}_k] \leqslant n^{-\delta\ell}.$$

when n is sufficiently large.

Proof. Since $\mathcal{F}$ is closed, we can use the generic bound of Lemma 5.5.4 to get the upper bound $|\mathcal{M}_{\mathcal{K},\ell}(\mathcal{F})| \leqslant \mathrm{Ex}(\ell, \mathcal{K}, \mathcal{D}^-, \varepsilon)$. Observe that the intersection of two cliques is also a clique, so that, for any family $\mathcal{H} \subseteq \mathcal{K}$, we have $\bigcap \mathcal{H} \in \mathcal{K}$. Because the distribution $\mathcal{D}^-$ is $G(n, p_k)$, this implies that a (p, ε)-clique sunflower (Definition 7.2.1) is also a $(\mathcal{K}, \mathcal{D}^-, \varepsilon)$-sunflower (Definition 5.5.1). Thus, we get from Lemma 7.2.2 that

$$|\mathcal{M}_{\mathcal{K},\ell}(\mathcal{F})| \leqslant \mathrm{Ex}(\ell, \mathcal{K}, \mathcal{D}^-, \varepsilon) \leqslant \ell!(2\log(1/\varepsilon))^\ell (1/p_k)^{\binom{\ell}{2}} \leqslant (\ell k \log n)^\ell \cdot p_k^{-\binom{\ell}{2}}.$$

Continuing to bound the number above, we get

$$|\mathcal{M}_{\mathcal{K},\ell}(\mathcal{F})| \leqslant (\delta k^2 \log n)^\ell \cdot 4^{\binom{\ell}{2}/\binom{k}{2}} \cdot n^{2\binom{\ell}{2}/k-1} \leqslant (n^{2/3-2\delta} \log n)^\ell \cdot n^{\delta\ell} \leqslant (n^{2/3})^\ell,$$

when n is sufficiently large. We now get from Lemma 7.3.3 and the union bound that

$$\Pr[\exists F \in \mathcal{M}_{\mathcal{K},\ell}(\mathcal{F}) : F \subseteq \mathbf{K}_k] \leqslant (n^{2/3})^\ell \cdot (k/n)^\ell \leqslant n^{-\delta\ell}. \qquad \square$$

Equipped with this bound, we can now show that $\mathcal{A}$ is inaccurate on $(\mathcal{D}^+, \mathcal{D}^-)$ and bound the approximation error on $\mathcal{D}^+$.

Lemma 7.4.4. *For large enough n, $\mathcal{A}$ is inaccurate on $(\mathcal{D}^+, \mathcal{D}^-)$.*

Proof. Let $f \in \mathcal{A}$. Since $f \in \mathcal{A}$, there exists a $(\mathcal{K}_{\leqslant w}, G(n, p_k), \varepsilon)$-closed family $\mathcal{F} \subseteq \mathcal{K}_{\leqslant w}$ such that $f = \lceil \mathcal{F} \rceil$. If $\lceil \mathcal{F} \rceil(\mathbf{K}_k) = 1$, there exists $\ell \in [w]$ and a minimal element $F \in \mathcal{M}_{\mathcal{K},\ell}(\mathcal{F})$ such that $F \subseteq \mathbf{K}_k$. Applying the bound of Lemma 7.4.3 and the union bound, we get

$$\Pr[f(\mathbf{K}_k) = 1] \leqslant \sum_{\ell=1}^{\delta k} n^{-\delta\ell} = o(1).$$

Therefore, we have $\Pr[f(\mathbf{K}_k) = 1] + \Pr[f(\mathbf{G}^-) = 0] = 1 + o(1)$. $\square$

Lemma 7.4.5. *For every $f, g \in \mathcal{A}$, we have*

$$\Pr_{\mathbf{K}_k \sim \mathcal{D}^+}[(f \wedge g)(\mathbf{K}_k) = 1 \text{ and } (f \sqcap g)(\mathbf{K}_k) = 0] \leqslant O(n^{-\delta k/2}).$$

Proof. Let $\mathcal{F}, \mathcal{H} \subseteq \mathcal{K}_{\leqslant w}$ be closed families such that $f = \lceil \mathcal{F} \rceil$ and $g = \lceil \mathcal{H} \rceil$. Let us suppose that $(f \wedge g)(\mathbf{K}_k) = 1$ and $(f \sqcap g)(\mathbf{K}_k) = 0$. By definition, there exists cliques $K_A \in \mathcal{M}(\mathcal{F})$ and $K_B \in \mathcal{M}(\mathcal{H})$ such that $K_A \cup K_B \subseteq \mathbf{K}_k$, but there does not exist any clique $K_C \in \mathcal{F} \cap \mathcal{H}$ such that $K_C \subseteq \mathbf{K}_k$.

Observe that, if $K_A \cup K_B \subseteq \mathbf{K}_k$, then $K_{A \cup B} \subseteq \mathbf{K}_k$. Moreover, if $|A \cup B| \leqslant w$, then $K_{A \cup B} \in \mathcal{K}_{\leqslant w}$. But, since $\mathcal{F}$ and $\mathcal{H}$ are closed, we get by item (3) of Lemma 5.4.3 that $K_{A \cup B} \in \mathcal{F} \cap \mathcal{H}$, a contradiction. So we have $|A \cup B| > w$. Therefore, at least one of A and B has size at least $w/2$, which means that at least one of K_A and K_B is contained in $\bigcup_{\ell=w/2}^{w} \mathcal{K}_\ell$. Applying the bound of Lemma 7.4.3 and the union bound, we get

$$\Pr[(f \wedge h)(\mathbf{K}_k) = 1 \text{ and } (f \sqcap h)(\mathbf{K}_k) = 0] \leqslant 2 \sum_{\ell=\delta k/2}^{\delta k} n^{-\delta\ell} \leqslant O(n^{-\delta k/2}). \qquad \square$$

7.4.2 Wrapping up

So far, we proved

- $(\mathcal{D}^+, \mathcal{D}^-)$ is a pair of test distributions for $\mathsf{Clique}(n,k)$ (Lemma 7.3.2);
- $\mathcal{A}$ is a legitimate model (Lemma 7.4.1);
- $\mathcal{A}$ is inaccurate on $(\mathcal{D}^+, \mathcal{D}^-)$ (Lemma 7.4.4);
- $\mathcal{A}$ is a gate-by-gate $O(n^{-\delta k/2})$-approximator for $(\mathcal{D}^+, \mathcal{D}^-)$ (Lemmas 7.4.5 and 7.4.2 and Remark 5.4.4);

Therefore, Theorem 5.3.5 implies that the monotone circuit complexity of $\mathsf{Clique}(n,k)$ is $\Omega(n^{\delta k/2})$.

Theorem 7.4.6. *Let $\delta > 0$ be an absolute constant. Any monotone circuit computing $\mathsf{Clique}(n,k)$ has size $\Omega(n^{\delta k/2})$, when $3 \leqslant k \leqslant n^{1/3-\delta}$.*

7.5 Proof of Lemma 7.2.2

Our proof of Lemma 7.2.2 will be very similar to the proof of Theorem 3.3.9. First, we consider the following auxiliary definition.

Definition 7.5.1. *Let $p, q \in (0,1)$. Let $\mathbf{U} \subseteq_q [n]$ and $\mathbf{G} \sim G(n,p)$ be sampled independently. A hypergraph $\mathcal{F}$ with vertex set $\binom{[n]}{2}$ is called a* (p,q,ε)-clique sunflower *if $\mathcal{F} \subseteq \mathcal{K}$ and, for $K_B = \bigcap \mathcal{F}$, we have*

$$\Pr\left[\forall K_A \in \mathcal{F} : K_A \nsubseteq \mathbf{G} \cup K_B \text{ or } A \nsubseteq \mathbf{U} \cup B\right] < \varepsilon.$$

The clique K_B is called core.

Clearly, a $(p,1,\varepsilon)$-clique sunflower corresponds to a (p,ε)-clique sunflower. By considering this stronger notion of sunflowers we will have a stronger induction hypothesis, which makes it easier to prove the induction step. We will now show a sunflower lemma for (p,q,ε)-clique sunflowers, which immediatly implies the clique sunflower lemma (Lemma 7.2.2) by Rossman's bound to the polynomials $s_\ell(t)$ (Proposition 3.3.8).

Lemma 7.5.2. *Let $p, q \in (0,1)$. If $\mathcal{F} \subseteq \mathcal{K}_\ell$ is such that $|\mathcal{F}| \geq s_\ell(\log(1/\varepsilon))(1/q)^\ell (1/p)^{\binom{\ell}{2}}$, then $\mathcal{F}$ contains a (p,q,ε)-clique sunflower.*

Proof. The proof is by induction on ℓ. In the base case $\ell = 1$, we have that $\mathcal{F}$ is itself a (p,q,ε)-clique sunflower:

$$\begin{aligned}\Pr\left[\forall K_A \in \mathcal{F} : K_A \nsubseteq \mathbf{G} \text{ or } A \nsubseteq \mathbf{U}\right] &= \Pr\left[\forall K_A \in \mathcal{F} : A \nsubseteq \mathbf{U}\right] \\ &= \prod_{K_A \in \mathcal{F}} \Pr[A \nsubseteq \mathbf{U}] \\ &= (1-q)^{|\mathcal{F}|} \quad < \quad (1-q)^{\log(1/\varepsilon)/q} \quad \leq \quad e^{-\log(1/\varepsilon)} \quad = \quad \varepsilon.\end{aligned}$$

Suppose now that $\ell \geq 2$ and the result holds for every $j \in [\ell - 1]$. We will consider two cases.
Case 1. There exists $j \in [\ell - 1]$ and $B \in \binom{[n]}{j}$ such that

$$d_{\mathcal{F}}(K_B) \geq s_{\ell-j}(\log(1/\varepsilon))(1/qp^j)^{\ell-j}(1/p)^{\binom{\ell-j}{2}}.$$

Let $\mathcal{T} = \left\{K_{A \setminus B} : K_A \in \mathcal{F} \text{ such that } B \subseteq A\right\} \subseteq \mathcal{K}_{\ell-j}$. By the induction hypothesis, there exists a (p, qp^j, ε)-clique sunflower $\mathcal{T}' \subseteq \mathcal{T}$ with core a K_D satisfying $D \in \binom{[n] \setminus B}{<\ell-j}$. We will now show that $\mathcal{F}' := \{K_{B \cup C} : K_C \in \mathcal{T}'\} \subseteq \mathcal{F}$ is a (p,q,ε)-clique sunflower contained in $\mathcal{F}$ with core $K_{B \cup D}$. Let

$\mathbf{U}_{\mathcal{T}} \subseteq_{qp^j} [n]$. We have

$$\begin{aligned}
&\Pr[\forall K_A \in \mathcal{F}' : K_A \nsubseteq \mathbf{G} \cup K_{B\cup D} \text{ or } A \nsubseteq \mathbf{U} \cup B \cup D] \\
&\quad = \Pr[\forall K_C \in \mathcal{T}' : K_{B\cup C} \nsubseteq \mathbf{G} \cup K_{B\cup D} \text{ or } B \cup C \nsubseteq \mathbf{U} \cup B \cup D] \\
&\quad = \Pr[\forall K_C \in \mathcal{T}' : K_{B\cup C} \nsubseteq \mathbf{G} \cup K_{B\cup D} \text{ or } C \nsubseteq \mathbf{U} \cup D] \\
&\quad = \Pr[\forall K_C \in \mathcal{T}' : K_C \nsubseteq \mathbf{G} \cup K_D \text{ or } C \nsubseteq \left\{v \in \mathbf{U} : \{v, w\} \in E(\mathbf{G}) \text{ for all } w \in B\right\} \cup D] \\
&\quad \leq \Pr[\forall K_C \in \mathcal{T}' : K_C \nsubseteq \mathbf{G} \cup K_D \text{ or } C \nsubseteq \mathbf{U}_{\mathcal{T}} \cup D] \\
&\quad < \varepsilon.
\end{aligned}$$

Therefore, $\mathcal{F}'$ is a (p, q, ε)-clique sunflower contained in $\mathcal{F}$.

Case 2. For all $j \in [\ell - 1]$ and $B \in \binom{[n]}{j}$, we have

$$d_{\mathcal{F}}(K_B) \leqslant s_{\ell-j}(\log(1/\varepsilon))(1/qp^j)^{\ell-j}(1/p)^{\binom{\ell-j}{2}}.$$

In this case, we show that $\mathcal{F}$ is itself a (p, q, ε)-sunflower with an empty core. Let

$$\begin{aligned}
\mu &:= \sum_{K_A \in \mathcal{F}} \Pr[K_A \subseteq \mathbf{G} \text{ and } A \subseteq \mathbf{U}] = |\mathcal{F}|\, q^{\ell} p^{\binom{\ell}{2}} \geqslant s_{\ell}(\log(1/\varepsilon)), \\
\Delta &:= \sum_{\substack{K_A, K_B \in \mathcal{F} \\ A \cap B \neq \emptyset}} \Pr[K_A \cup K_B \subseteq \mathbf{G} \text{ and } A \cup B \subseteq \mathbf{U}].
\end{aligned}$$

Janson's Inequality (Lemma 2.2.4) gives the following bound:

$$\Pr[\forall K_A \in \mathcal{F} : K_A \nsubseteq \mathbf{G} \text{ or } A \nsubseteq \mathbf{U}] \leq \exp\left\{-\mu^2/\Delta\right\}. \tag{7.1}$$

We now define the following auxiliary parameter $\overline{\Delta}$, which ignores the diagonal terms of the sum defining Δ:

$$\overline{\Delta} := \sum_{\substack{K_A, K_B \in \mathcal{F} \\ A \cap B \neq \emptyset, A \neq B}} \Pr[K_A \cup K_B \subseteq \mathbf{G} \text{ and } A \cup B \subseteq \mathbf{U}].$$

We obtain $\Delta = \mu + \overline{\Delta}$. We bound $\overline{\Delta}$ as follows:

$$\begin{aligned}
\overline{\Delta} &= \sum_{\substack{K_A, K_B \in \mathcal{F} \\ A \neq B, A \cap B \neq \emptyset}} q^{2\ell - j} p^{2\binom{\ell}{2} - \binom{j}{2}} \\
&= \sum_{j=1}^{\ell-1} \sum_{T \in \binom{[n]}{j}} \sum_{\substack{K_A, K_B \in \mathcal{F} \\ A \cap B = T}} q^{2\ell - j} p^{2\binom{\ell}{2} - \binom{j}{2}} \\
&\leqslant \sum_{j=1}^{\ell-1} q^{2\ell - j} p^{2\binom{\ell}{2} - \binom{j}{2}} \sum_{T \in \binom{[n]}{j}} d_{\mathcal{F}}(K_T)^2.
\end{aligned}$$

By double-counting, we get that $\sum_{T \in \binom{[n]}{j}} d_{\mathcal{F}}(K_T) = |\mathcal{F}| \binom{\ell}{j}$. Therefore, applying the bound on

$d_{\mathcal{F}}(K_T)$, we obtain

$$\begin{aligned}\sum_{T\in\binom{[n]}{j}} d_{\mathcal{F}}(K_T)^2 &\leqslant s_{\ell-j}(\log(1/\varepsilon))\cdot(1/qp^j)^{\ell-j}(1/p)^{\binom{\ell-j}{2}}\sum_{T\in\binom{[n]}{j}} d_{\mathcal{F}}(K_T)\\ &\leqslant s_{\ell-j}(\log(1/\varepsilon))\cdot(1/qp^j)^{\ell-j}(1/p)^{\binom{\ell-j}{2}}\,|\mathcal{F}|\binom{\ell}{t}\\ &= |\mathcal{F}|\cdot p^{\binom{j}{2}-\binom{\ell}{2}}q^{j-\ell}s_{\ell-t}(\log(1/\varepsilon))\binom{\ell}{t}\\ &= \mu\cdot p^{\binom{j}{2}-2\binom{\ell}{2}}q^{j-2\ell}s_{\ell-t}(\log(1/\varepsilon))\binom{\ell}{t}.\end{aligned}$$

We now continue to bound $\overline{\Delta}$, as follows:

$$\overline{\Delta} \leqslant \mu\sum_{j=1}^{\ell-1}\binom{\ell}{j}s_{\ell-j}(\log(1/\varepsilon)) = \mu\left(\frac{s_\ell(\log(1/\varepsilon))}{\log(1/\varepsilon)}-1\right).$$

Therefore, we have $\Delta \leqslant \mu\frac{s_\ell(\log(1/\varepsilon))}{\log(1/\varepsilon)}$, whence we get

$$\frac{\mu^2}{\Delta} \geqslant \log(1/\varepsilon)\frac{\mu}{s_\ell(\log(1/\varepsilon))} > \log(1/\varepsilon).$$

Finally, from (7.1) we get

$$\Pr[\forall K_A\in\mathcal{F} : K_A\not\subseteq \mathbf{G} \text{ or } A\not\subseteq \mathbf{U}] \leq \exp\left\{-\mu^2/\Delta\right\} < \varepsilon. \qquad \square$$

7.6 Further questions

7.6.1 Improvements for clique sunflowers

Our bound for the extremal number of clique sunflowers (Lemma 7.2.2) follows closely Rossman's original bound for robust sunflowers (Theorem 3.3.9). We expect that a proof along the lines of the work of Alweiss *et al* [ALWZ19] and Rao [Rao20] (explained in Chapter 4) should be able to give us an even better bound, replacing the $\ell!$ factor with an $(\log \ell)^\ell$ factor. This would extend our $n^{\Omega(k)}$ lower bound for $\mathsf{Clique}(n,k)$ to a larger range of k, up to $k \leq n^{1/2-o(1)}$.

Problem 7.6.1. *Give a better bound for (p,ε)-clique sunflowers, along the lines of the work of Alweiss* et al *[ALWZ19] and Rao [Rao20].*

References

[AB87] N. Alon e R. B. Boppana. The monotone circuit complexity of Boolean functions. *Combinatorica*, 7(1):1–22, 1987. 2, 8, 9, 28, 29, 31, 33, 37, 38, 45, 46

[AB09] Sanjeev Arora e Boaz Barak. *Computational complexity*. Cambridge University Press, Cambridge, 2009. A modern approach. 1

[ABI86] Noga Alon, László Babai e Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms*, 7(4):567–583, 1986. 5, 38

[AKS83] M. Ajtai, J. Komlós e E. Szemerédi. An $O(n \log n)$ sorting network. Em *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, STOC '83, páginas 1–9, New York, NY, USA, 1983. Association for Computing Machinery. 1

[ALWZ19] Ryan Alweiss, Shachar Lovett, Kewen Wu e Jiapeng Zhang. Improved bounds for the sunflower lemma. *arXiv:1908.08483*, 2019. 2, 8, 9, 10, 13, 14, 16, 18, 21, 22, 33, 51

[And85] A. E. Andreev. A method for obtaining lower bounds on the complexity of individual monotone functions. *Dokl. Akad. Nauk SSSR*, 282(5):1033–1037, 1985. 8, 29, 37

[Ber82] S. J. Berkowitz. On some relationships between monotone and non-monotone circuit complexity. Technical Report. University of Toronto, 1982. 43

[Blo80] P. A. Bloniarz. The complexity of monotone boolean functions and an algorithm for finding shortest paths on a graph. Relatório técnico, USA, 1980. Ph.D. Thesis. 37

[BS90] Ravi B. Boppana e Michael Sipser. The complexity of finite functions. Em *Handbook of theoretical computer science, Vol. A*, páginas 757–804. Elsevier, Amsterdam, 1990. 8, 46

[BT87] B. Bollobás e A. Thomason. Threshold functions. *Combinatorica*, 7(1):35–38, 1987. 5

[CKR20] Bruno Pasqualotto Cavalar, Mrinal Kumar e Benjamin Rossman. Monotone lower bounds from robust sunflowers. 2020. Submitted to LATIN 2020, https://www.ime.usp.br/~brunopc/files/betterexponential.pdf. 2, 29, 31, 33, 37, 38, 45, 46

[ER60] P. Erdős e R. Rado. Intersection theorems for systems of sets. *J. London Math. Soc.*, 35:85–90, 1960. 2, 7, 8

[FKNP19] Keith Frankston, Jeff Kahn, Bhargav Narayanan e Jinyoung Park. Thresholds versus fractional expectation-thresholds. *arXiv:1910.13433*, 2019. 13

[FLV15] E. Fischer, O. Lachish e Y. Vasudev. Trading query complexity for sample-based testing and multi-testing scalability. Em *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, páginas 1163–1182, 2015. 16

[Für80] Zoltán Füredi. On maximal intersecting families of finite sets. *J. Combin. Theory Ser. A*, 28(3):282–289, 1980. 8

[GL20] Tom Gur e Oded Lachish. On the power of relaxed local decoding algorithms. Em *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '20, páginas 1377–1394, USA, 2020. Society for Industrial and Applied Mathematics. 16, 17

[GMR13] Parikshit Gopalan, Raghu Meka e Omer Reingold. DNF sparsification and a faster deterministic counting algorithm. *Computational Complexity*, 22(2):275–310, 2013. 9

[HR00] Danny Harnik e Ran Raz. Higher lower bounds on monotone size. Em *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, páginas 378–387. ACM, New York, 2000. 2, 29, 37, 38, 42, 43

[Jan90] Svante Janson. Poisson approximation for large deviations. *Random Structures Algorithms*, 1(2):221–229, 1990. 4

[JŁR00] Svante Janson, Tomasz Łuczak e Andrzej Ruciński. *Random graphs*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000. 4

[Juk11] Stasys Jukna. *Extremal combinatorics*. Texts in Theoretical Computer Science. An EATCS Series. Springer, Heidelberg, second edição, 2011. With applications in computer science. 8

[Juk12] Stasys Jukna. *Boolean function complexity*, volume 27 of *Algorithms and Combinatorics*. Springer, Heidelberg, 2012. Advances and frontiers. 19, 29

[Kar72] Richard M. Karp. Reducibility among combinatorial problems. Em *Complexity of computer computations (Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, N.Y., 1972)*, páginas 85–103, 1972. 45

[KO18] Jan Krajíček e Igor C. Oliveira. On monotone circuits with local oracles and clique lower bounds. *Chic. J. Theoret. Comput. Sci.*, páginas Art. 1, 18, 2018. 45

[Kos97] A. V. Kostochka. A bound of the cardinality of families not containing Δ-systems. Em *The mathematics of Paul Erdős, II*, volume 14 of *Algorithms Combin.*, páginas 229–235. Springer, Berlin, 1997. 13

[Kra49] Leon Gordon Kraft. A device for quantizing, grouping, and coding amplitude-modulated pulses, 1949. Master's thesis, Massachusetts Institute of Technology. 22

[LLZ18] Xin Li, Shachar Lovett e Jiapeng Zhang. Sunflowers and quasi-sunflowers from randomness extractors. Em *APPROX-RANDOM*, volume 116 of *LIPIcs*, páginas 51:1–13, 2018. 9

[LMZ20] Shachar Lovett, Raghu Meka e Jiapeng Zhang. Improved lifting theorems via robust sunflowers. *Electronic Colloquium on Computational Complexity (ECCC)*, 27:48, 2020. 9

[LSZ19] Shachar Lovett, Noam Solomon e Jiapeng Zhang. From DNF compression to sunflower theorems via regularity. Em *34th Computational Complexity Conference*, volume 137 of *LIPIcs. Leibniz Int. Proc. Inform.*, páginas Art. No. 5, 14. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019. 9

[LZ19] Shachar Lovett e Jiapeng Zhang. DNF sparsification beyond sunflowers. Em *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, páginas 454–460. ACM, 2019. 9, 10

[MP20] Ian Mertz e Toniann Pitassi. Lifting: As easy as 1,2,3, 2020. https://eccc.weizmann.ac.il/report/2020/111/. 9

[NRR18] Sivaramakrishnan Natarajan Ramamoorthy e Anup Rao. Lower bounds on non-adaptive data structures maintaining sets of numbers, from sunflowers. Em *33rd Computational Complexity Conference*, volume 102 of *LIPIcs. Leibniz Int. Proc. Inform.*, páginas Art. No. 27, 16. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2018. 8

[PR17] Toniann Pitassi e Robert Robere. Strongly exponential lower bounds for monotone computation. Em *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, páginas 1246–1255. ACM, 2017. 37

[Rao19] Anup Rao. Coding for sunflowers. University of Washington Theory Seminar https://www.youtube.com/watch?v=fzmsbylTJKM, 2019. 21, 22

[Rao20] Anup Rao. Coding for sunflowers. *Discrete Analysis*, 2020. To appear. 2, 8, 9, 14, 16, 21, 23, 33, 51

[Raz85a] A. A. Razborov. Lower bounds of monotone complexity of the logical permanent function. *Mat. Zametki*, 37(6):887–900, 942, 1985. 2, 29, 31, 33

[Raz85b] A. A. Razborov. Lower bounds on the monotone complexity of some Boolean functions. *Dokl. Akad. Nauk SSSR*, 281(4):798–801, 1985. 1, 2, 8, 27, 28, 29, 33, 37, 45

[Raz87] A. A. Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Mat. Zametki*, 41(4):598–607, 623, 1987. 2

[Ros14] Benjamin Rossman. The monotone complexity of k-clique on random graphs. *SIAM J. Comput.*, 43(1):256–279, 2014. 2, 9, 10, 11, 12, 28, 46

[Ros19] Benjamin Rossman. Approximate sunflowers. unpublished, available at http://www.math.toronto.edu/rossman/approx-sunflowers.pdf, 2019. 11

[Tie84] J Tiekenheinrich. A 4n-lower bound on the mononotone network complexity of a one-output boolean function. *Information Processing Letters*, 18:201–201, 1984. 37