DNF sparsification beyond sunflowers

Shachar Lovett[*]
University of California, San Diego
slovett@ucsd.edu

Jiapeng Zhang[†]
University of California, San Diego
jpeng.zhang@gmail.com

October 31, 2018

Abstract

There are two natural complexity measures associated with DNFs: their size, which is the number of clauses; and their width, which is the maximal number of variables in a clause. It is a folklore result that DNFs of small size can be approximated by DNFs of small width (logarithmic in the size). The other direction is much less clear.

Gopalan, Meka and Reingold [Computational Complexity 2013] showed that the other direction – DNF sparsification – holds as well. Any DNF of width w can be approximated to within error ε by a DNF of size $(w \log(1/\varepsilon))^{O(w)}$. Our main interest in this work is the dependence on the width w. The same dependence of w^w appears in several other open problems in combinatorics and complexity, such as the Erdős-Rado sunflower conjecture and Mansour's conjecture. In fact, there are deep connections between these three problems. Our main result is DNF compression with an improved dependence on the width, which overcomes the w^w barrier. Concretely, we show that any DNF of width w can be approximated to within error ε by a DNF of size $(1/\varepsilon)^{O(w)}$.

The proof centers around a new object which we call the DNF index function. Given a DNF, the DNF index function outputs for an input the first clause that satisfies it (if one exists). Our proof has two parts: a combinatorial part, where we exhibit a switching lemma for the DNF index function; and an analytic part, where we use the switching lemma to bound the noise sensitivity of the DNF index function, and then use it to obtain our DNF compression result.

1 Introduction

Any boolean function $f : \{0,1\}^n \to \{0,1\}$ can be represented as a CNF or as a DNF. Functions which can represented as small CNFs or DNFs are central in computational complexity theory, and have been widely studied. We focus on DNFs in this paper, but our results translate to CNFs as well.

Let f be boolean function expressed as a DNF. There are two natural complexity measures associates with it: the numbers of clauses, called *size* and denoted $s(f)$; and the maximal number of variables in a clause, called *width* and denoted $w(f)$. It is a folklore result that DNFs of small size can be approximated by DNFs of small width; concretely, DNFs of size s can be ε-approximated by DNFs of width $w = O(\log(s/\varepsilon))$. Gopalan, Meka and Reingold [GMR13] studied the reverse problem of *DNF sparsification*: can DNFs of small width be approximated by DNFs of small size? their motivation, other than being a natural problem on the structure of DNFs, came from the goal of designing faster deterministic algorithms to approximately count the number of satisfying assignments of a DNF. Their main structural result on DNFs is the following.

Theorem 1.1 ([GMR13])**.** *Let f be a boolean function which can be expressed as a width-w DNF. Then for every $\varepsilon > 0$, f can be ε-approximated by a DNF of width w and size $(w \log(1/\varepsilon))^{O(w)}$.*

It was conjectured in [GMR13] that the term $(w \log(1/\varepsilon))^w$ is not tight. In particular, Conjecture 6.1 in their paper speculates that the bound can be improved to $c(\varepsilon)^w$, and moreover that possibly one can take $c(\varepsilon) = (\log 1/\varepsilon)^{O(1)}$. Our main result in this paper resolves the weaker conjecture.

Theorem 1.2 (Main theorem)**.** *Let f be a boolean function which can be expressed as a width-w DNF. Then for every $\varepsilon > 0$, f can be ε-approximated by a DNF of width w and size $(1/\varepsilon)^{O(w)}$.*

While the dependence we obtain on the error $\varepsilon > 0$ is probably sub-optimal, our main goal was to sharpen the dependence on the width w, from $w^{O(w)}$ to $2^{O(w)}$ (for a fixed error ε). As we shortly describe, the same challenge appears in two other related problem: the Erdős-Rado sunflower conjecture [ER60] and Mansour's conjecture [Man95].

1.1 Erdős-Rado sunflower conjecture

The Erdős-Rado sunflower conjecture deals with set systems. To see the relation between set systems and DNFs, note that if $\mathcal{F}$ is a set system of sets $S \subset [n]$, then there is a natural associated (monotone) DNF given by $f(x) = \vee_{S \in \mathcal{F}} \wedge_{i \in S} x_i$. A w-uniform set system is a set system in which all sets contain w elements. The DNF associated to a w-uniform set system is a width-w DNF. In the other direction, any width-w DNF contains a large unary DNF (concretely, with at least a 2^{-w} fraction of the clauses) [GMR13]. Unary DNFs are similarly equivalent to set systems.

We next introduce sunflowers, which are widely studied in combinatorics.

Definition 1.3 (Sunflower)**.** *A collection of r sets $S_1, \ldots, S_r$ is called an r-sunflower if all the pairwise intersections $S_i \cap S_j$ are the same.*

The Erdős-Rado sunflower lemma [ER60] states that if $\mathcal{F}$ is a w-uniform set system, and $|\mathcal{F}| > w!(r-1)^w$, then F must contain an r-sunflower. The well-known sunflower conjecture is that the dependence on w can be improved.

Conjecture 1.4 (Sunflower conjecture)**.** *For any $r \geq 3$ there exists $c = c(r)$ such that the following holds. If F is a w-uniform set system, and $|F| > c^w$, then F contains an r-sunflower.*

The sunflower conjecture has been open for nearly 60 years. Despite much research, the best bounds, even for $r = 3$, still are of the order of w^w.

Approximate sunflowers. Rossman [Ros14] defined the notion of an approximate sunflower, motivated by applications in complexity theory.

Definition 1.5 (Approximate sunflower)**.** *A set system $\mathcal{F}$ is a γ-approximate sunflower if the following holds. Let $C = \cap_{S \in \mathcal{F}} S$ be the common core of all sets in $\mathcal{F}$, and define $\mathcal{F}' = \{S \setminus C : S \in \mathcal{F}\}$. Let f' be the monotone DNF associated with $\mathcal{F}'$. Then $\Pr[f'(x) = 1] \geq 1 - e^{-\gamma}$.*

Rossman proved that if $\mathcal{F}$ is a w-uniform set system of size $|\mathcal{F}| > (w \log(1/\gamma))^{O(w)}$, then $\mathcal{F}$ must contain a γ-approximate sunflower. Similarly to the sunflower conjecture, it is reasonable to conjecture that a better bound holds.

Conjecture 1.6 (Approximate sunflower conjecture)**.** *For any $\gamma > 0$ there exists $c = c(\gamma) > 0$ such that the following holds. If F is a w-uniform set system, and $|F| > c^w$, then F contains an γ-approximate sunflower.*

In fact, one can show [LLZ18, LSZ18] that the approximate sunflower conjecture implies the sunflower conjecture.

Connection to DNF sparsification. The main tool used in [GMR13] to achieve their result about DNF sparsification is the sunflower and approximate sunflower lemmas stated above. Roughly speaking, they used the approximate sunflower lemma to compress an approximate sunflower to its common core. A barrier towards an improved dependence on the width w in their result, is that the dependence on w in both lemmas is of the order of w^w. Thus, one of the main motivations of the current work is to achieve DNF sparsification that breaks the w^w bound. A more ambitious goal (which we discuss in Section 4) is to use the connection between sunflower structure and DNF sparsification as highlighted in [GMR13], together with our improved DNF sparsification result, to obtain improved bounds for the sunflower conjecture.

1.2   Mansour's conjecture

Mansour's conjecture [Man95] deals with the approximation of DNFs by sparse polynomials. We say that a boolean function $f : \{0,1\}^n \to \{0,1\}$ can be ε-approximated by a polynomial of sparsity t if there exists a polynomial $p : \{0,1\}^n \to \mathbb{R}$ with at most t monomials such that

$$\mathbb{E}_{x \in \{0,1\}^n} \left[(f(x) - p(x))^2 \right] \leq \varepsilon.$$

Conjecture 1.7 (Mansour's conjecture for size). *For any $\varepsilon > 0$ there exists $c = c(\varepsilon) > 0$ such that the following holds. Any DNF of size s can be ε-approximated by a polynomial of sparsity $c^{\log s}$.*

One of the motivations behind Mansour's conjecture, other than a better understanding of the structure of DNFs, is that it would give an efficient agnostic learning algorithm for DNFs [GKK08b, GKK08a].

As was noted in [GMR13], it makes sense to speculate a similar conjecture for bounded width DNFs. As any DNF of size s can be approximated by a DNF of width $w = O(\log(s/\varepsilon))$, this latter conjecture is stronger.

Conjecture 1.8 (Mansour's conjecture for width). *For any $\varepsilon > 0$ there exists $c = c(\varepsilon) > 0$ such that the following holds. Any DNF of width w can be ε-approximated by a polynomial of sparsity c^w.*

The best known bound for Mansour's conjecture for width [Man95] is that it holds for sparsity $(w \log(1/\varepsilon))^{O(w)}$ (the bound for size holds by approximating a bounded size DNF with a bounded width DNF). So again, we see the w^w term appearing, where the conjecture asks if it can be improved to $c(\varepsilon)^{O(w)}$ (and moreover that $c(\varepsilon) = O(\log 1/\varepsilon)$). In fact, Mansour shows that his technique would not yield a $w^{o(w)}$-type bound, so other ideas are necessary. A direct corollary of our main theorem is that both versions of Mansour's conjecture are equivalent.

Corollary 1.9. *Conjecture 1.7 and Conjecture 1.8 are equivalent.*

In addition, we think that our sparsification lemma gives some clues on how to prove Mansour's conjecture. We discuss this more in Section 4.

1.3   Proof overview

Let $f = \varphi_1 \vee \ldots \vee \varphi_t$ be a DNF, where each φ_i is a clause (conjunction of literals). The main object which underlies our work is the function which maps an input to the *first* clause which satisfies it. We call this the *DNF index function*. Observe that this depends on the specific structure of the DNF, and not just the boolean function it computes. Moreover, it depends on the order of the clauses.

Definition 1.10 (DNF index function)**.** *Let* $f = \varphi_1 \vee \ldots \vee \varphi_t$ *be a DNF. The index function of* f *is a function* $\mathrm{Ind}f : \{0,1\}^n \to \{1,\ldots,t\} \cup \{\bot\}$ *defined as follows:*

$$\mathrm{Ind}f(x) = \begin{cases} \bot & \text{if } f(x) = 0 \\ \min\{i \in [t] : \varphi_i(x) = 1\} & \text{if } f(x) = 1 \end{cases}$$

Let $p_i = \Pr_x[\mathrm{Ind}f(x) = i]$ denote the fraction of inputs such that the i-th clause is the first clause that satisfies them. The following is a natural approach for DNF sparsification: only keep clauses φ_i for which p_i is noticeable.

However, it is not clear how many noticeable clauses could there be. For example, a bad scenario would be if $p_i = 1/t$ for all i; in this case there would be no way to significantly sparsify the DNF. However, this cannot be the case, as if φ_1 as width w then $p_1 = 2^{-w}$. So, the main challenge is to show that there is a small set of indices $I \subset [t]$, such that $\sum_{i \notin I} p_i \leq \varepsilon$. Our main theorem shows that this holds for $|I| = (1/\varepsilon)^{O(w)}$.

Next, we highlight how we show that. At its core, our argument has two parts: a combinatorial part, where we prove switching lemma for the DNF index function; and an analytic part, where we analyze the noise sensitivity of the index function and connect it to the problem of DNF sparsification.

Combinatorial part: Switching lemma. The behaviour of DNFs under random restrictions have been well studied. Razborov [Raz95], refining previous work of Håstad [Hås87], showed that DNFs simplify under random restrictions. See also [Bea94] for an exposition.

We need a few standard definitions. A *restriction* is $\rho \in \{0,1,*\}^n$. An (n,k)-*random restriction* is a uniform restriction $\rho \in \{0,1,*\}^n$ with exactly k stars. Given a boolean function $f : \{0,1\}^n \to \{0,1\}$, its restriction under ρ is denoted $f|_\rho : \{0,1\}^{\rho^{-1}(*)} \to \{0,1\}$. Given a function $f : \{0,1\}^n \to X$ for some finite set X, we denote its *decision tree complexity* by $\mathrm{DT}(f)$.

The well-known switching lemma for DNFs [Hås87, Raz95, Bea94] is the following result. Let f be an n-variate boolean function computed by a width-w DNF. Let $k = \alpha n$. Let $\rho \in \{0,1,*\}^n$ be an (n,k)-random restriction. Then for any $d \geq 1$,

$$\Pr_\rho[\mathrm{DT}(f|_\rho) \geq d] \leq (7\alpha w)^d.$$

We extend this result to the DNF index function. Assume that $f = \varphi_1 \vee \ldots \vee \varphi_t$ and let $\mathrm{Ind}f : \{0,1\}^n \to \{1,\ldots,t\} \cup \{\bot\}$ be its associated DNF index function. We prove (Lemma 2.1) that for any $d \geq 1$,

$$\Pr_\rho[\mathrm{DT}(\mathrm{Ind}f|_\rho) \geq d] \leq (32\alpha w)^d.$$

Analytic part: Noise sensitivity. Let I denote the set of noticeable clauses $i \in [t]$. Our goal is to upper bound $\sum_{i \notin I} p_i$. To that end, we study the behaviour of the index function under noise. Given $x \in \{0,1\}^n$ let $\mathcal{N}_\rho(x)$ denote the noise distribution around x, where $y \sim \mathcal{N}_\rho(x)$ is sampled by taking $\Pr[x_i = y_i] = \rho$ independently for $i \in [n]$. The main observation is that that if all (or most) of the p_i are negligible, then $\mathrm{Ind}f$ cannot be stable under noise. This is since if $\mathrm{Ind}f(x) = i$, and p_i is tiny, then if we sample $y \sim \mathcal{N}_\rho(x)$ then with high probability $\mathrm{Ind}f(y) \neq i$. This follows from the well known fact (whose proof is based on the hypercontractive inequality) that small sets in the hypercube are not noise stable.

So, our goal is to show that $\mathrm{Ind}f$ is noise stable. Concretely, we say that an input x which satisfies f is (ρ,γ)-stable for f if

$$\Pr_{y \sim \mathcal{N}_\rho(x)}[\mathrm{Ind}f(x) = \mathrm{Ind}f(y)] \geq \gamma.$$

4

Thus, if most inputs are stable, then if we first sample $x \in \{0,1\}^n$ uniformly, and then take $i = \text{Ind}f(x)$, then with high probability p_i is noticeable. This then implies that $\sum p_i$ is concentrated on a small set I. To conclude, we need to show that indeed most inputs $x \in f^{-1}(1)$ are noise stable.

This in turn follows from our switching lemma for $\text{Ind}f$. Consider an equivalent way to jointly sample x, y, where first we sample $\rho \in \{0,1,*\}^n$ where $\Pr[\rho_i = *] = 1 - \rho$, and then sample x, y conditioned on $x_i = y_i = \rho_i$ whenever $\rho_i \neq *$. If $\text{Ind}f|_\rho$ has a small depth decision tree, then there is a noticeable probability that it evaluates to the same leaf on both x, y. That is, that $\text{Ind}f(x) = \text{Ind}f(y)$. Thus, the switching lemma allows us to prove that most inputs are noise stable, completing the proof.

Paper organization. We prove the switching lemma for the DNF index function in Section 2. We apply it to obtain the DNF sparsification result in Section 3. We discuss open problems in Section 4.

2 Switching lemma for DNF index function

Let f be a width-w DNF. Recall that the index function of f maps an input to the first term that is satisfies, or to $\perp$ if no clause is satisfied. The main goal of this section is to prove a switching lemma for the DNF index function. We start with some preliminary definitions.

Decision tree. Let $g : \{0,1\}^n \to X$ be a function where X is some finite set. A decision tree for g is a binary tree whose nodes are labeled by variables and whose leaves are labeled by elements of X. The decision tree complexity of g, denoted $\text{DT}(g)$, is the minimal depth of a decision tree computing g.

Restrictions. A restriction is $\rho \in \{0,1,*\}^n$. Given a function $g : \{0,1\}^n \to X$, its restriction $g|_\rho$ is the sub-function obtained by restricting to inputs which agree with ρ. That is, let $S = \{i : \rho_i = *\}$ be the "alive" variables. Then $g|_\rho : \{0,1\}^S \to X$ by mapping $z \in \{0,1\}^S$ to $g(x)$, where $x_i = z_i$ if $i \in S$ and $x_i = \rho_i$ otherwise.

Random restrictions. An (n,k)-random restriction is the the uniform distribution over restrictions $\rho \in \{0,1,*\}^n$ with exactly k stars.

The following is the main result of this section.

Lemma 2.1 (Switching lemma for the DNF index function)**.** *Let f be a width-w DNF on n variables, and let $\text{Ind}f$ be its DNF index function. Let $k = \alpha n$ and let ρ be an (n,k)-random restriction.. Then for every $d \geq 1$,*

$$\Pr_\rho[\text{DT}(\text{Ind}f|_\rho) \geq d] \leq (32\alpha w)^d.$$

Proof. We assume $\alpha \leq 1/32w$ otherwise the claim is trivial. Let $\rho \in \{0,1,*\}^n$. We say that ρ is "bad" if $\text{DT}(\text{Ind}f|_\rho) \geq d$. We use a compression argument, similar to the one used by Razborov [Raz95] to prove the switching lemma for DNFs.

The DNF $f = \varphi_1 \vee \ldots \vee \varphi_t$ is fixed throughout. Let V_j denote the variables that appear in φ_j. We use the following notations. Given two strings a, a' their concatenation is $a \circ a'$. Given a known

set W of size w, and a set $V \subset W$ of a known size, we can uniquely describe V by a string in $[w]^{|V|}$. We denote this representation $\text{SetIndex}(W, V)$. We define three operations on restrictions:

- **Append:** given a restriction $\rho \in \{0, 1, *\}^n$ and a partial input $u \in \{0, 1\}^S$ where $S \subset \rho^{-1}(*)$, we denote by $\text{append}(\rho, u)$ the restriction obtained by appending u to ρ:

$$\text{append}(\rho, u) = \begin{cases} u_i & \text{if } i \in S \\ \rho_i & \text{otherwise} \end{cases}$$

- **Delete:** given a restriction $\rho \in \{0, 1, *\}^n$ and a set $S \subset \rho^{-1}(\{0, 1\})$, we denote by $\text{delete}(\rho, S)$ the restriction obtained by setting the symbols in S to stars:

$$\text{delete}(\rho, S) = \begin{cases} * & \text{if } i \in S \\ \rho_i & \text{otherwise} \end{cases}$$

- **Update:** given a restriction $\rho \in \{0, 1, *\}^n$ and a partial input $u \in \{0, 1\}^S$ where $S \subset \rho^{-1}(\{0, 1\})$, we denote by $\text{update}(\rho, u)$ the restriction obtained by updating the elements in S to u:

$$\text{update}(\rho, u) = \begin{cases} u_i & \text{if } i \in S \\ \rho_i & \text{otherwise} \end{cases}$$

We next present the encoding and decoding algorithms.

$$Encode(\rho)$$

Input: restriction $\rho \in \{0, 1, *\}^n$.
Output: restriction $\tau \in \{0, 1, *\}^n$, string $a \in \mathbb{N}^*$.

1. Initialize $\tau = \rho$. Initialize a to be an empty string.

2. For $j = 1, \ldots, t$ do:

   (a) If $\varphi_j|_\rho \equiv 0$ then skip to the next j.

   (b) If $\varphi_j|_\rho \equiv 1$ then abort the loop.

   (c) Otherwise compute:

       i. $A_j = \{i \in V_j : \rho_i = *\}$ the alive variables in φ_j.

       ii. $u_j \in \{0, 1\}^{A_j}$ an assignment under which $\text{DT}(\text{Ind} f|_{\text{append}(\rho, u_j)})$ is maximized.

       iii. $v_j \in \{0, 1\}^{A_j}$ an assignment under which $\varphi_j|_{\text{append}(\rho, v_j)} \equiv 1$.

   (d) Update:

       i. $\rho = \text{append}(\rho, u_j)$.

       ii. $\tau = \text{append}(\tau, v_j)$.

       iii. $a = a \circ |A_j| \circ \text{SetIndex}(V_j, A_j) \circ u_j$.

3. Return τ, a.

$$Decode(\tau, a)$$

Input: restriction $\tau \in \{0, 1, *\}^n$, string $a \in \mathbb{N}^*$.
Output: restriction $\rho \in \{0, 1, *\}^n$.

1. Initialize A to be an empty set.

2. For $j = 1, \ldots, t$ do:

   (a) If $\varphi_j|_\tau \equiv 0$ then skip to the next j.
   (b) Otherwise read from a: $A_j \subset V_j$ and $u_j \in \{0, 1\}^{A_j}$.
   (c) Update:
       i. $A = A \cup A_j$.
       ii. $\tau = \text{update}(\tau, u_j)$.

3. Return $\rho = \text{delete}(\tau, A)$.

We first argue that the encoding and decoding are correct.

Claim 2.2. *For any $\rho \in \{0, 1, *\}^n$ it holds that $DECODE(ENCODE(\rho)) = \rho$.*

Proof. Let $\tau, a = ENCODE(\rho)$. Note that if $\rho_i \neq *$ then $\tau_i = \rho_i$. So, we just need to verify that the decoding procedure deletes exactly the elements that were appended in the encoding procedure, namely $\cup A_j$. Say that an index $j \in [t]$ is active if in the encoding procedure, we have that $\varphi_j|_\rho$ is non-constant when it is considered. Let $J = \{j_1, \ldots, j_r\}$ denote the set of active indices. The main observation is that these are also the indices in which in the decoding procedure we have $\varphi_j|_\rho \not\equiv 0$. In fact, one can further verify that $\varphi_j|_\rho \equiv 1$ in these cases. To conclude note that the auxiliary string a allows to precisely recover the sets A_j. $\square$

To conclude the proof we need to bound the probability that ρ is bad. To do so, we bound the size of the set $\{ENCODE(\rho) : \rho \text{ is bad}\}$. Assume that $\tau, a = ENCODE(\rho)$. As ρ is bad, we have $\text{DT}(\text{Ind} f|_\rho) \geq d$. This means that $m = \sum |A_j| \geq d$ by the choice of the u_j. Given a fixed m we bound the number of choices for τ, a.

The restriction τ has exactly $k - m$ stars, and so has $\binom{n}{k-m} 2^{n-k+m}$ options. Assume there are r sets A_j with $|A_j| > 0$. The number of choices of $|A_1|, \ldots, |A_r|$ is equal to the number of ways we can decompose $m = a_1 + \ldots + a_r$ with $a_i \geq 1$, which equals $\binom{m-1}{r-1}$. The sum of these over all r is 2^{m-1}. Given that $|A_j| > 0$ for some j, the number of options for $\text{SetIndex}(V_j, A_j)$ is $w^{|A_j|}$ and the number of choices for u_j is $2^{|A_j|}$. So we obtain

$$\{ENCODE(\rho) : \rho \text{ is bad}\} \leq \sum_{m \geq d} \binom{n}{k-m} 2^{n-k+m} (4w)^m.$$

On the other hand, the total number of restrictions ρ with exactly k stars equals $\binom{n}{k} 2^{n-k}$. So we obtain that

$$\Pr[\rho \text{ is bad}] \leq \sum_{m \geq d} \frac{\binom{n}{k-m}}{\binom{n}{k}} (8w)^m \leq \sum_{m \geq d} \left(\frac{\alpha}{1-\alpha}\right)^m (8w)^m \leq \sum_{m \geq d} (16\alpha w)^m \leq (32\alpha w)^d,$$

where the last inequality follows from the assumption $\alpha \leq 1/32w$. $\qquad\square$

We would need the following simple corollary. Let $\mathcal{R}_{n,\alpha}$ be the distribution over restrictions $\{0, 1, *\}^n$ where $\Pr[\rho_i = *] = \alpha$ and $\Pr[\rho_i = 0] = \Pr[\rho_i = 1] = \frac{1-\alpha}{2}$.

Corollary 2.3. *Let f be a width-w DNF on n variables, and let $Indf$ be its DNF index function. Let $\rho \sim \mathcal{R}_{n,\alpha}$. Then for every $d \geq 1$,*

$$\Pr_{\rho \sim \mathcal{R}_{n,\alpha}} [\text{DT}(Indf|_\rho) \geq d] \leq (64\alpha w)^d + 2^{-\Omega(\alpha n)}.$$

Proof. Let $\rho \sim \mathcal{R}_{n,\alpha}$ and let $k = |\rho^{-1}(*)|$. Conditioned on $|\rho^{-1}(*)| = k$, the distribution of ρ is an (n, k)-random restriction. Namely, it is uniform in $U_{n,k}$, the set of restrictions in $\{0, 1, *\}^n$ with exactly k stars. Then

$$\Pr_{\rho \sim \mathcal{R}_{n,\alpha}} [\text{DT}(\text{Ind}f|_\rho) \geq d] = \sum_k \Pr_{\rho \sim \mathcal{R}_{n,\alpha}} [|\rho^{-1}(*)| = k|] \cdot \Pr_{\rho \in U_{n,k}} [\text{DT}(\text{Ind}f|_\rho) \geq d].$$

The probability that $k \geq 2\alpha n$ is exponentially small in αn. Whenever $k \leq 2\alpha n$ we use Lemma 2.1 to deduce the bound. $\qquad\square$

3 DNF sparsification

Our goal is to prove the following theorem.

Theorem 3.1. *Let $f = \varphi_1 \vee \ldots \vee \varphi_t$ be a width-w DNF. Then for every $\varepsilon > 0$, there exists a subset $I \subset [t]$ of size $|I| \leq (1/\varepsilon)^{O(w)}$ such that the following holds. Let $f' = \bigvee_{i \in I} \varphi_i$. Then*

$$\Pr[f(x) \neq f'(x)] \leq \varepsilon.$$

3.1 Noticeable indices

Let $f = \varphi_1 \vee \ldots \vee \varphi_t$ be a width-w DNF. To recall, $\text{Ind}f$ is the index function of f, which maps an input x to the first clause that it satisfies, or to $\bot$ if $f(x) = 0$. The main question we study is: how are the outputs of the DNF index function distributed? for example, can they be uniform in $[t]$? we show that the answer is no if t is too large, which leads us to be able to approximate f as a smaller DNF.

Definition 3.2 (Noticeable index). *Let $f = \varphi_1 \vee \ldots \vee \varphi_t$ be a DNF. An index $i \in [t]$ is called τ-noticeable if*

$$\Pr_{x \in \{0,1\}^n} [Indf(x) = i] \geq \tau.$$

For example, if f is a DNF, and φ_1 is the first clause with w variables, then 1 is (2^{-w})-noticeable since $\Pr[\varphi_1(x) = 1] = 2^{-w}$. We denote the set of all noticeable indices by

$$I(f, \tau) = \{i \in [t] : i \text{ is } \tau\text{-noticeable}\}.$$

The following claim is straightforward.

Claim 3.3. $|I(f, \tau)| \leq 1/\tau$.

Proof. Ii $i \in I(f, \tau)$ then $\Pr[\text{Ind}f(x) = i] \geq \tau$. These events are disjoint for different i. $\qquad\square$

3.2 Noise stability

The following are standard definitions in Fourier analysis of boolean functions.

Definition 3.4 (Noisy distribution). *Given $x \in \{0,1\}^n$ and a noise parameter $\rho \in [0,1]$, we denote by $\mathcal{N}_\rho(x)$ the distribution over $y \in \{0,1\}^n$, where $\Pr[y_i = x_i] = \frac{1+\rho}{2}$ and $\Pr[y_i \neq x_i] = \frac{1-\rho}{2}$ independently for all $i \in [n]$.*

Definition 3.5 (Stability). *Let $g : \{0,1\}^n \to \{0,1\}$ be a boolean function. The ρ-stability of g is*

$$\mathrm{Stab}_\rho(g) = \Pr_{x \in \{0,1\}^n, y \sim N_\rho(x)} [g(x) = g(y) = 1].$$

We use the following shorthand: $|g| = \Pr_{x \in \{0,1\}^n} [g(x) = 1]$ is the fraction of inputs on which g accepts. We will need the following fact which follows from the hyper-contractive inequality (see for example [O'D14], page 259).

Fact 3.6. *Let $g : \{0,1\}^n \to \{0,1\}$ be a boolean function. Then $\mathrm{Stab}_\rho(g) \leq |g|^{\frac{2}{1+\rho}}$.*

The following claim is a simple corollary of Fact 3.6. It studies the noise sensitivity of a decomposition of a boolean function f into disjoint boolean functions $g_1, \ldots, g_t$.

Claim 3.7. *Let $f = g_1 + \ldots + g_t$ where $f, g_1, \ldots, g_t : \{0,1\}^n \to \{0,1\}$ are boolean functions. Given a parameter $\tau \in [0,1]$ define*

$$I = \{i \in [t] : |g_i| \geq \tau\}.$$

Then

$$\sum_{i \notin I} \mathrm{Stab}_\rho(g_i) \leq \tau^{\frac{1-\rho}{1+\rho}}.$$

Proof. Fact 3.6 gives that for $i \notin I$ we have $\mathrm{Stab}_\rho(g_i) \leq |g_i|^{\frac{2}{1+\rho}} \leq |g_i| \cdot \tau^{\frac{1-\rho}{1+\rho}}$. Thus

$$\sum_{i \notin I} \mathrm{Stab}_\rho(g_i) \leq \tau^{\frac{1-\rho}{1+\rho}} \sum_{i \notin I} |g_i| \leq \tau^{\frac{1-\rho}{1+\rho}}.$$

$\square$

3.3 Noise stability of the index function

The noise stability of boolean function is a well-studied topic. Here, we study the noise stability of the DNF index function.

Definition 3.8 (Stable and sensitive inputs). *Let f be a DNF, $Indf$ be its DNF index function, and let $x \in \{0,1\}^n$ be an input which satisfies f. The input x is called (ρ, γ)-stable for f if*

$$\Pr_{y \sim \mathcal{N}_\rho(x)} [Indf(x) = Indf(y)] \geq \gamma.$$

Otherwise, x is called (ρ, γ)-sensitive for f.

Definition 3.9 (Index sensitivity). *The (ρ, γ)-index sensitivity of f is the fraction of (ρ, γ)-sensitive inputs for f,*

$$IndexSensitivity(f, \rho, \gamma) = \Pr_{x \in \{0,1\}^n} [f(x) = 1 \ \wedge \ x \text{ is } (\rho, \gamma)\text{-sensitive for } f].$$

The following lemma connects the index sensitivity to DNF sparsification.

Lemma 3.10. *Let $f = \varphi_1 \vee \ldots \vee \varphi_t$ be a DNF. Fix $\rho, \gamma, \tau \in [0, 1]$. Let $I = I(f, \tau)$ be the set of τ-noticeable clauses of f, and define $f' = \bigvee_{i \in I} \varphi_i$. Then*

$$\Pr[f(x) \neq f'(x)] \leq IndexSensitivity(f, \rho, \gamma) + \gamma^{-1} \tau^{\frac{1-\rho}{1+\rho}}.$$

Proof. Observe that $f'(x) \leq f(x)$ for all x. So, if $f(x) \neq f'(x)$ then necessarily $f'(x) = 0, f(x) = 1$ and $\mathrm{Ind}f(x) \notin I$. Let $I^c = [t] \setminus I$. Then

$$\Pr[f(x) \neq f'(x)] \leq \sum_{i \in I^c} \Pr[\mathrm{Ind}f(x) = i].$$

To simplify notation, for $x \in \{0, 1\}^n$ let $E(x)$ denote the event "x is (ρ, γ)-stable for f". Then we can bound

$$\Pr[f(x) \neq f'(x)] \leq \Pr[f(x) = 1 \ \wedge \ \neg E(x)] + \sum_{i \in I^c} \Pr[\mathrm{Ind}f(x) = i \ \wedge \ E(x)].$$

The first term equals IndexSensitivity(f, ρ, γ). To bound the second term, Fix $i \in I^c$. Conditioned on $E(x)$ we have $\Pr_{y \sim \mathcal{N}_\rho(x)}[\mathrm{Ind}f(x) = \mathrm{Ind}f(y)] \geq \gamma$. Thus

$$\Pr_{x \in \{0,1\}^n, y \sim \mathcal{N}_\rho(x)}[\mathrm{Ind}f(x) = \mathrm{Ind}f(y) = i \ \wedge \ E(x)] \geq \gamma \cdot \Pr[\mathrm{Ind}f(x) = i \ \wedge \ E(x)].$$

Let $g_i : \{0, 1\}^n \rightarrow \{0, 1\}$ be the indicator of the event $\mathrm{Ind}f(x) = i$, so that $f = g_1 + \ldots + g_t$. Then we have

$$\Pr[\mathrm{Ind}f(x) = i \ \wedge \ E(x)] \leq \gamma^{-1} \cdot \Pr[\mathrm{Ind}f(x) = \mathrm{Ind}f(y) = i] = \gamma^{-1} \cdot \mathrm{Stab}_\rho(g_i).$$

To conclude we have

$$\Pr[f(x) \neq f'(x)] \leq IndexSensitivity(f, \rho, \gamma) + \gamma^{-1} \cdot \sum_{i \in I^c} \mathrm{Stab}_\rho(g_i).$$

The bound now follows from Claim 3.7. $\qquad\square$

Thus, we reduced the problem of compressing DNFs to that of bounding the index sensitivity of DNFs. The following lemma shows that for width-w DNFs is, most of their inputs are stable at noise level $\rho = 1 - O(1/w)$. Its proof uses the switching lemma for the DNF index function, or more precisely Corollary 2.3.

Lemma 3.11. *Let f be an n-variate width-w DNF. Set $\rho = 1 - \frac{1}{128w}$ and let $\gamma = 2^{-d}$ for an integer $d \geq 1$. Then*

$$IndexSensitivity(f, \rho, \gamma) \leq 2\gamma + 2^{-\Omega(n/w)}.$$

Proof. Let $x \in \{0, 1\}^n$ sampled uniformly and let $y \sim \mathcal{N}_\rho(x)$. It will be convenient to sample x, y in an equivalent but different way. Recall that $\mathcal{R}_{n,\alpha}$ is a distribution over restrictions $\rho \in \{0, 1, *\}^n$ where $\Pr[\rho_i = *] = \alpha$ and $\Pr[\rho_i = 0] = \Pr[\rho_i = 1] = \frac{1-\alpha}{2}$. Then we can sample (x, y) as follows:

1. Sample $\rho \sim \mathcal{R}_{n,\alpha}$ where $\alpha = 1 - \rho$. Let $S = \{i : \rho_i = *\}$.

2. Sample $x|_S \in \{0, 1\}^S$ uniformly, and set $x_i = \rho_i$ if $i \notin S$.

3. Sample $y|_S \in \{0,1\}^S$ uniformly, and set $y_i = \rho_i$ if $i \notin S$.

Next, fix ρ and assume that $\mathrm{DT}(\mathrm{Ind}f|_\rho) = d$. Then in particular, the probability that $x|_S, y|_S$ take the same path in the decision tree is at least 2^{-d}. So we obtain that

$$\Pr_{x,y}\left[\mathrm{Ind}f(x) = \mathrm{Ind}f(y)|\rho\right] \geq 2^{-\mathrm{DT}(\mathrm{Ind}f|_\rho)}.$$

Let $p(x)$ denote the probability that $\mathrm{DT}(\mathrm{Ind}f|_\rho) \geq d$ when ρ is sampled conditioned on x. Then

$$\Pr_{y}\left[\mathrm{Ind}f(x) = \mathrm{Ind}f(y)|x\right] \geq 2^{-(d-1)}(1 - p(x)).$$

We would like to show that for most x it holds that $p(x) \leq 1/2$; such x will be $(\rho, 2^{-d})$-stable for f. That is, we wish to upper bound

$$p = \Pr_{x}\left[p(x) \geq 1/2\right] \leq 2\Pr_{\rho}\left[\mathrm{DT}(\mathrm{Ind}f|_\rho) \geq d\right].$$

Corollary 2.3 bounds the right hand side, and for $\alpha = 1/128w$ gives

$$p \leq 2^{1-d} + 2^{-\Omega(n/w)}.$$

$\square$

We now prove Theorem 3.1.

Proof of Theorem 3.1. Let $\rho, \gamma, \tau \in [0,1]$ to be optimized shortly. Let $I = I(f, \tau)$, where $|I| \leq 1/\tau$ by Claim 3.3. Combining Lemma 3.10 with Lemma 3.11 gives the bound

$$\Pr[f \neq f'] \leq O(\gamma) + 2^{-\Omega(n/w)} + \gamma^{-1}\tau^{\frac{1-\rho}{1+\rho}}.$$

First, we note that we may assume that $n = \Omega(w\log(1/\varepsilon))$ as otherwise the theorem holds vacuously, as the total number of possible width w clauses is $2^w\binom{n}{w} \leq O(\log 1/\varepsilon)^w$. Let $\rho = 1 - 1/128w$, $\gamma = O(\varepsilon)$ and $\tau = \varepsilon^{O(w)}$. Then

$$\Pr[f(x) \neq f'(x)] \leq \varepsilon.$$

$\square$

4 Discussions and open problems

We discuss several open problems and future research directions which stem from our work.

Improved sparsification bounds. It is a natural question to ask what is the tight bound (respect to both ε and w) of DNF sparsification. Gopalan et al. [GMR13] gave a bound of $(w\log(1/\varepsilon))^{O(w)}$ and our bound is $(1/\varepsilon)^{O(w)}$. In our approach, there is a trade off on how to choose the noise parameter ρ. In our proof, we used $\rho = 1 - \Omega(1/w)$. With smaller noise, we get a stronger random restriction result, however a weaker bound from the hyper-contrative inequality (which is used in the proof of Fact 3.6). In order to improve our bound, a natural approach is to use large noise, $\rho = 1 - \Omega(1)$. In order to improve the sparsification bound in this way, we need to understand what happens when we randomly restrict $1 - \Omega(1)$ variables of a given DNF. We believe that some structure still remains in this DNF that can allow for our proof structure to go through. This makes the bound of $(\log 1/\varepsilon)^{O(w)}$ conjectured by [GMR13] seem plausible. In fact, a stronger bound of $\mathrm{poly}(2^w, \log(1/\varepsilon))$ might be true.

Connections to Mansour's conjecture. There are possible approaches to prove Mansour's conjecture via our technique. First, let us point out a result of Klivans, Lee and Wan [KLW10]. They prove that Mansour's conjecture is true for random DNFs. More specifically, it holds for any DNF which satisfies (in a strong way) that most inputs satisfy only a few clauses.

Theorem 4.1 ([KLW10]). *Let $f = \varphi_1 \vee \cdots \vee \varphi_t$ be a DNF with t terms. Assume that for some $\varepsilon > 0$ it holds that*

$$\Pr_x [x \text{ satisfies } j \text{ clauses}] = O\left(\frac{\log(1/\varepsilon)}{j}\right)^j \qquad \forall j \leq O(\log n).$$

Then f can be ε-approximated by a polynomial with sparsity $t^{\log(1/\varepsilon)}$.

The high level proof of Theorem 4.1 is as follows. Let V_i denote the variables of φ_i. For an integer $k \geq 1$ define

$$\mathcal{A}_k := \left\{ S \subseteq [n] : \exists I \subseteq [t], (|I| = k) \wedge (S \subseteq \bigcup_{i \in I} V_i) \right\}$$

It is clear that $|\mathcal{A}_k| \leq t^k \cdot 2^{k \cdot w}$. The proof of [KLW10] constructs a polynomial supported in $\mathcal{A}_{\log(1/\varepsilon)}$ that approximates f. This is true as they show that the $\log(1/\varepsilon)$-th moment of $S(x) = \sum \varphi_i(x)$ approximates the DNF.

However, this is false for general DNFs. The main reason is that clauses can intersect quite a bit, and hence there is no control over high moments of $S(x)$. However, we conjecture that this can be overcome with a more careful definition. For $k \geq 1$ define

$$\mathcal{B}_k := \left\{ S \subseteq [n] : \exists I \subseteq [t], \left(|\bigcup_{i \in I} V_i| \leq kw\right) \wedge \left(S \subseteq \bigcup_{i \in I} V_i\right) \right\}$$

Conjecture 4.2. *Let $f = \varphi_1 \vee \cdots \vee \varphi_t$ be a DNF with t terms. Then f can be ε-approximated by a polynomial whose monomials are supported in $\mathcal{B}_k$ for $k = O(\log(1/\varepsilon))$.*

It is clear that $\mathcal{A}_k \subseteq \mathcal{B}_k$. However it is not clear how to upper bound the size of $\mathcal{B}_k$. Here is the place where our sparsification result may help. Define

$$\mathcal{C}_k := \left\{ I \subseteq [t] : |\bigcup_{i \in I} V_i| \leq kw \right\}$$

and define

$$\psi_k := \bigvee_{I \in \mathcal{C}_k} (\wedge_{i \in I} \varphi_i)$$

Let ψ_k' be the sparsifed DNF that ε-approximated ψ_k. Define

$$\mathcal{D}_k := \left\{ S \subseteq [n] : \exists I \in \psi', S \subseteq \bigcup_{i \in I} V_i \right\}$$

Note that D is a DNF of width kw and hence ψ_k' has $(1/\varepsilon)^{O(kw)}$ clauses by our result.

Conjecture 4.3. *Let $f = \varphi_1 \vee \cdots \vee \varphi_t$ be a DNF with t terms. Then f can be ε-approximated by a polynomial whose monomials are supported in $\mathcal{D}_k$ for $k = O(\log(1/\varepsilon))$.*

Connections to the sunflower conjecture. We already discussed that the sunflower and approximate sunflower lemmas were pivotal to the result of [GMR13] on DNF sparsification. The goal here is to speculate about a reverse connection. We showed in [LLZ18] a direct connection between the sunflower conjecture and problems about the existence of specific extractors. Here, we point out another connection to this paper.

Let $\mathcal{F}$ be a w-uniform set system. For a set T of the base elements let $\mathcal{F}_T = \{S \setminus T : S \in \mathcal{F}, T \subset S\}$ denote the link of T in $\mathcal{F}$. Let $f_{\mathcal{F}}$ denote the monotone width-w DNF associated with $\mathcal{F}$. As we already mentioned, one way to prove the sunflower conjecture is to confirm the approximate sunflower conjecture (Conjecture 1.6). That is, to find a set T such that

$$\Pr_x[f_{\mathcal{F}_T} = 1] \geq 0.9 .$$

In this paper, we proved that the DNF index function is not symmetric, in the sense that only a few clauses are noticeable. The existence of such T seems related to the ability to show that the DNF satisfiability is concentrated in relatively a few clauses; however how to make this precise is still open.

References

[Bea94] Paul Beame. A switching lemma primer. Technical report, Technical Report UW-CSE-95-07-01, Department of Computer Science and Engineering, University of Washington, 1994.

[ER60] Paul Erdős and R Rado. Intersection theorems for systems of sets. *Journal of the London Mathematical Society*, 35(1):85–90, 1960.

[GKK08a] Parikshit Gopalan, Adam Kalai, and Adam R Klivans. A query algorithm for agnostically learning DNF?. In *COLT*, pages 515–516, 2008.

[GKK08b] Parikshit Gopalan, Adam Tauman Kalai, and Adam R Klivans. Agnostically learning decision trees. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 527–536. ACM, 2008.

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

[Hås87] Johan Håstad. Computational limitations of small-depth circuits. 1987.

[KLW10] Adam R Klivans, Homin K Lee, and Andrew Wan. Mansour's conjecture is true for random DNF formulas. In *COLT*, pages 368–380. Citeseer, 2010.

[LLZ18] Xin Li, Shachar Lovett, and Jiapeng Zhang. Sunflowers and quasi-sunflowers from randomness extractors. In *LIPIcs-Leibniz International Proceedings in Informatics*, volume 116. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[LSZ18] Shachar Lovett, Noam Solomon, and Jiapeng Zhang. Unpublished work. 2018.

[Man95] Yishay Mansour. An $n^{O(\log \log n)}$ learning algorithm for DNF under the uniform distribution. *Journal of Computer and System Sciences*, 50(3):543–550, 1995.

[O'D14] Ryan O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.

[Raz95]   Alexander A Razborov. Bounded arithmetic and lower bounds in boolean complexity.
            In *Feasible Mathematics II*, pages 344–386. Springer, 1995.

[Ros14]   Benjamin Rossman. The monotone complexity of k-clique on random graphs. *SIAM
            Journal on Computing*, 43(1):256–279, 2014.