Approximating Linear Threshold Predicates*

Mahdi Cheraghchi[1], Johan Håstad[2], Marcus Isaksson[3], and Ola Svensson[2]

[1] School of Computer and Communication Sciences, EPFL, Lausanne, Switzerland. mahdi.cheraghchi@epfl.ch

[2] KTH - Royal Institute of Technology, Stockholm, Sweden. {johanh,osven}@kth.se

[3] Chalmers University of Technology and University of Gothenburg, Gothenburg, Sweden. maris@chalmers.se

Abstract. We study constraint satisfaction problems on the domain $\{-1, 1\}$, where the given constraints are homogeneous linear threshold predicates. That is, predicates of the form $\mathrm{sgn}(w_1x_1 + \cdots + w_nx_n)$ for some positive integer weights $w_1, \ldots, w_n$. Despite their simplicity, current techniques fall short of providing a classification of these predicates in terms of approximability. In fact, it is not easy to guess whether there exists a homogeneous linear threshold predicate that is approximation resistant or not.

The focus of this paper is to identify and study the approximation curve of a class of threshold predicates that allow for non-trivial approximation. Arguably the simplest such predicate is the majority predicate $\mathrm{sgn}(x_1 + \cdots + x_n)$, for which we obtain an almost complete understanding of the asymptotic approximation curve, assuming the Unique Games Conjecture. Our techniques extend to a more general class of "majority-like" predicates and we obtain parallel results for them. In order to classify these predicates, we introduce the notion of *Chow-robustness* that might be of independent interest.

Keywords: approximability, constraint satisfaction problems, linear threshold predicates

1 Introduction

Constraint satisfaction problems or more succinctly CSPs are at the heart of theoretical computer science. In a CSP we are given a set of constraints, each putting some restriction on a constant size set of variables. The variables can take values in many different domains but in this paper we focus on the case of variables taking Boolean values. This is the most fundamental case and it has also attracted the most attention over the years. We also focus on the case where each condition is given by the same predicate, P, applied to a sequence of literals. The role of this predicate P is key in this paper and as it is more important for us than the number of variables, we reserve the letter n for the arity of this

* This research is supported by the ERC Advanced investigator grant 226203. M. Cheraghchi is supported by the ERC Advanced investigator grant 228021.

predicate while using N to be the number of variables in the instance. We also reserve m to denote the number of constraints.

Traditionally we ask for an assignment that satisfies all constraints and in this case it turns out that all Boolean CSPs are either NP-complete or belong to P and this classification was completed already in 1978 by Schaefer [15]. In this paper we study Max-CSPs which are optimization problems where we want to satisfy as many constraints as possible. Almost all Max-CSPs of interest turn out to be NP-hard and the main focus is that of efficient approximability.

The standard measure of approximability is given by a single number C and an algorithm is a C-approximation algorithm if it, on each input, finds an assignment with an objective value that is at least C times the optimal value. Here we might allow randomization and be content if the assignment found satisfies these many constraints on average. A more refined question is to study the approximation curve where for each constant c, assuming that the optimal assignment satisfies cm constraints, we want to determine the maximal number of constraints that we can satisfy efficiently.

To get a starting point to discuss the quality of approximation algorithms it is useful to first consider the most simple algorithm that chooses the values of the variables randomly and uniformly from all values in $\{0,1\}^N$. If the predicate P is satisfied by t inputs in $\{0,1\}^n$ it is easy to see that this algorithm, on the average, satisfies $mt2^{-n}$ constraints. By using the method of conditional expectations it is also easy to deterministically find an assignment that satisfies this number of constraints.

A very strong type of hardness result possible for a Max-CSP is to prove that, even for instances where the optimal assignment satisfies all constraints, it is NP-hard to find an assignment that does significantly better (by a constant factor independent of N) than the above trivial algorithm. We call such a predicate “approximation resistant on satisfiable instances”. A somewhat weaker, but still strong, negative result is to establish that the approximation ratio given by the trivial algorithm, namely $t2^{-n}$, is the best approximation ratio that can be obtained by an efficient algorithm. This is equivalent to saying that we cannot satisfy significantly more than $mt2^{-n}$ constraints when given an almost satisfiable instance. We call such a predicate “approximation resistant”. It is well known that, unless P=NP, Max-3-Sat (i.e. when P is the disjunction of the three literals) is approximation resistant on satisfiable instances and Max-3-Lin (i.e. when P is the exclusive-or of three literals) is approximation resistant [8].

When it comes to positive results on approximability the most powerful technique is semi-definite programming introduced in this context in the classical paper by Goemans and Williamson [6] studying the approximability of Max-Cut, establishing the approximability constant $\alpha_{GW} \approx .878$. In particular, this result implies that Max-Cut is not approximation resistant. Somewhat surprisingly as proved by Khot et al. [12], this constant has turned out, assuming the Unique Games Conjecture, to be best possible. We note that these results have been extended in great generality and O'Donnell and Wu [14] determined the complete approximation curve of Max-Cut.

The general problem of determining which predicates are approximation resistant is still not resolved but as this is not the main theme of this paper let us cut this discussion short by mentioning a general result by Austrin and Mossel [2]. This paper relies on the Unique Games Conjecture by Khot [11] and proves that, under this conjecture, any predicate such that the set $P^{-1}(1)$ supports a pairwise independent measure is approximation resistant.

On the algorithmic side there is a general result by Hast, [7], that is somewhat complementary to the result of Austrin and Mossel. Hast considers the real valued function $P^{\leq 2}$ which is the sum of the linear and quadratic parts of the Fourier expansion of P. Oversimplifying slightly, the result by Hast says that if $P^{\leq 2}$ is positive on all inputs accepted by P then we can derive a non-trivial approximation algorithm and hence P is not approximation resistant.

To see the relationship between the results of Austrin and Mossel, and Hast, note that the condition of Austrin and Mossel is equivalent to saying that there is a probability distribution on inputs accepted by P such that the average of any unbiased quadratic function[1] is 0. In contrast, Hast needs that a particular unbiased quadratic function is positive on all inputs accepted by P. It is not difficult to come up with predicates that satisfies neither of these two conditions and hence we do not have a complete classification, even if we are willing to assume the Unique Games Conjecture. The combination of the two results, however, points to the class of predicates that can be written on the form

$$P(x) = \operatorname{sgn}(Q(x))$$

for a quadratic function Q as an interesting class of predicates to study and this finally brings us to the topic of this paper. We study this scenario in the simplest form by assuming that Q is in fact an unbiased linear function, L. In other words we have

$$P(x) = \operatorname{sgn}(L(x)) = \operatorname{sgn}\left(\sum_{i=1}^{n} w_i x_i\right),$$

for some, without loss of generality, positive integral weights $(w_i)_{i=1}^n$. Note that if we allow a constant term in L the situations is drastically different as for instance 3-Sat is the sign of linear form if we allow a non-zero constant term. One key difference is that a probability distribution supported on the set "$L(x) > 0$" cannot have even unbiased variables in the case when L is without constant term and thus hardness results such as the result by Austrin and Mossel do not apply.

To make life even simpler we make sure that L never takes the value 0 and as $L(-x) = -L(x)$, P accepts precisely half of the inputs and thus the number of constraints satisfied by a random assignment is, on the average, $m/2$.

The simplest such predicate is majority of an odd number of inputs. For this predicate it easy to see that Hast's condition is fulfilled and hence, for any odd value of n, his results imply that majority is not approximation resistant. This

[1] Throughout this work, we find it more convenient to represent Boolean values by {-1,+1} rather than {0,1}.

result generalizes to "majority-like" functions as follows. For a linear threshold functions, the Chow parameters, $\bar{P} = (\hat{P}(i))_{i=0}^{n}$, [3] are for, $i > 0$, defined to be the correlations between the output of the function and inputs x_i. We have that $\hat{P}(0)$ is the bias of the function and thus in our case this parameter is always equal to 0 and hence ignored.

Now if we order the weights $(w_i)_{i=1}^{n}$ in nondecreasing order then also the $\hat{P}(i)$'s are nondecreasing but in general quite different from the weights. It is well known that the Chow parameters determine the threshold function uniquely [3] but the computational problem of given $\bar{P}$, how to recover the weights, or even to compute P efficiently is an interesting problem and several heuristics have been proposed [10, 17, 9, 4] together with an empirical study that compares various methods [18]. More recently, the problem of finding an approximation of P given the Chow parameters has received increased attention, see e.g. [13] and [5]. The most naive method is to use $\bar{P}$ as weights. This does not work very well in general but this is a case of special interest to us as it is precisely when this method gives us back the original function that we can apply Hast's results directly. We call such a threshold function "Chow-robust" and we have not been able to find the characterization of this class of functions in the literature. If we ignore some error terms and technical conditions a sufficient condition to be Chow-robust is roughly that

$$\sum_{i=1}^{n}(w_i^3 - w_i) \leq 3\sum_{i=1}^{n} w_i^2 \tag{1}$$

and thus it applies to functions with rather modest weights. We believe that this condition is not very far from necessary but we have not investigated this in detail.

Having established non-approximation resistance for such predicates we turn to study the full curve of approximability and, in an asymptotic sense as a function of n, we get almost tight answers establishing both approximability results and hardness results. Our results do apply with degrading constants to more general threshold functions but let us here state them for majority. We have the following theorem.

Theorem 1. *(Informal) Given an instance of Max-Maj-n with n odd and m constraints and assume that the optimal assignment satisfies $(1 - \frac{\delta}{n+1})m$ constraints, for some $\delta < 1$. Then it is possible to efficiently find an assignment that satisfies*

$$\left(\frac{1}{2} + \Omega\left(\frac{(1-\delta)^{3/2}}{n^{1/2}}\right) - \mathcal{O}\left(\frac{\log^4 n}{n^{5/6}}\right)\right) m$$

constraints.

Thus for large n we need almost satisfiable instances to get above the threshold $\frac{1}{2}$ obtained by a a random assignment. This might seem weak but we prove that this is probably the correct threshold.

Theorem 2. *(Informal) Assume the Unique Games Conjecture and let $\epsilon > 0$ be arbitrary. Then it is NP-hard to distinguish instances of Max-Maj-n where the optimal value is $(1 - \frac{1}{n+1} - \epsilon)m$, from those where the optimal value is $(\frac{1}{2} + \epsilon)m$.*

This proves that the range of instances to which Theorem 1 applies is essentially the correct one. A drawback is that the error term $\mathcal{O}\left(\frac{\log^4 n}{n^{5/6}}\right)$ in Theorem 1 dominates the systematic contribution of $(1-\delta)^{3/2} n^{-1/2}$ for δ very close to 1 and hence the threshold is not sharp. We are, however, able to sharply locate the threshold where something nontrivial can be done by combining our result with the general results by Hast. For details, see Section 3.

To see that the advantage obtained by the algorithm is also the correct order of magnitude we have the following theorem.

Theorem 3. *(Informal) Assume the Unique Games Conjecture and let $\epsilon > 0$ be arbitrary. Then there is an absolute constant c such that it is NP-hard to distinguish instances of Max-Maj-n where the optimal value is $(1 - \epsilon)m$, from those where the optimal value is $(\frac{1}{2} + \frac{c}{\sqrt{n}} + \epsilon)m$.*

In summary, we get an almost complete understanding of the approximability curve of majority, at least in an asymptotic sense as a function of n. This complements the results for majority on three variables, for which there is a 2/3-approximation algorithm [19] and it is NP-hard to do substantially better [8].

The idea of the algorithm behind Theorem 1 is quite straightforward while its analysis gets rather involved. We set up a natural linear program which we solve and then use the obtained solution as biases in a randomized rounding. The key problem that arises is to carefully analyze the probability that a sum of biased Boolean variables is positive. In the case of majority-like variables we have the additional complication of the different weights. This problem is handled by writing the probability in question as a complex integral and then estimating this integral by the saddle-point method. The resulting proof is quite long and does not fit within the page limit of the current abstract. This proof and several other proofs are hence omitted and can be found in the full version of the paper.

The hardness results given in Theorem 2 and Theorem 3 resort to the techniques of Austrin and Mossel [2]. The key to these results is to find suitable pairwise independent distributions relating to our predicate. In the case of majority it is easy to find such distributions explicitly, while in the case of more general weights the construction gets more involved. In particular, we need to answer the following question: What is the minimal value of $\Pr[L(x) < 0]$ when x is chosen according to a pairwise independent distribution. This is a nice combinatorial question of independent interest.

An outline of the paper is as follows. Notation and conventions used throughout the paper are presented in Section 2. This is followed by the adaptation of Hast's algorithm for odd Chow-robust predicates and the result that (essentially) the condition $\sum_{j=1}^{n} w_j^3 - w_j \leq 3 \sum_{j=1}^{n} w_j^2$ on the weights is sufficient for a predicate to be Chow-robust. In Section 4, we present our main algorithm for Chow-robust predicates which establishes Theorem 1 in the special case of majority. These positive results are then complemented in Section 5 where we show essentially tight hardness results assuming the Unique Games Conjecture. Finally, we discuss the obtained results together with interesting future directions (Section 6). As already stated, the current abstract only contains some of our shorter proofs and a reader interested in the full proofs must turn to the full version of the paper.

2 Preliminaries

We consider the optimization problem Max-CSP(P) for homogeneous linear threshold predicates $P : \{-1,1\}^n \to \{-1,1\}$ of the form

$$P(x) = \operatorname{sgn}(w_1 x_1 + \cdots + w_n x_n),$$

where we assume that the weights are non-decreasing positive integers $1 \leq w_1 \leq \ldots \leq w_n$ such that $\sum_{j=1}^{n} w_j$ is odd and $w_{\max} := \max_j w_j = w_n$. The special case of equal weights, which requires n to be odd, is denoted by Maj_n, and we also write Max-Maj-n for Max-CSP(Maj_n). Using Fourier expansion, any such function can be written uniquely as

$$P(x) = \sum_{S \subseteq [n]} \hat{P}(S) \prod_{j \in S} x_j.$$

The Fourier coefficients are given by $\hat{P}(S) = \mathbb{E}[P(X) \prod_{j \in S} X_j]$, where X is uniform on $\{-1,1\}^n$. Since all homogeneous linear threshold predicates are odd we have $\hat{P}(S) = 0$ when $|S|$ is even. We will also write $\hat{P}(j) = \hat{P}(\{j\})$ for the first level Fourier coefficients (i.e. the Chow parameters) and let $P^{-1}(1)$ denote the set of assignments that satisfy P, i.e. $P^{-1}(1) = \{x : P(x) = 1\}$.

For an instance $\mathcal{I} = (m, N, l, s)$ of Max-CSP(P) consisting of m constraints, N variables and matrices $l \in N^{m \times n}$, $s \in \{-1,1\}^{m \times n}$, the objective is to maximize the number of satisfied constraints or, equivalently since $P(-x) = -P(x)$ and thus $\mathbb{E}[P(x)] = 0$, the average advantage

$$\mathsf{Adv}(x) := \frac{1}{m} \sum_{i=1}^{m} P(s_{i,1} x_{l_{i,1}}, \ldots, s_{i,n} x_{l_{i,n}})$$

subject to $x \in \{-1,1\}^N$.

3 Adaptation of the algorithm by Hast

Using Fourier expansion we may write the advantage of an assignment to a Max-CSP(P) instance as

$$\mathsf{Adv}(x) = \frac{1}{m} \sum_{i=1}^{m} \operatorname{sgn}\left(\sum_{j=1}^{n} w_j s_{i,j} x_{l_{i,j}} \right) = \sum_{S \subseteq [N] : |S| \leq n} c_S \prod_{k \in S} x_k. \tag{2}$$

Hast [7] gives a general approximation algorithm for Max-CSP(P) that achieves a non-trivial approximation ratio whenever the linear part of the instance's objective function is large enough. We use his algorithm, but as our basic predicates are odd we have that $c_S = 0$ for any S of even size and we get slightly better bounds.

Theorem 4. *For any $\delta > 0$, there is a probabilistic polynomial time algorithm which given an instance of Max-CSP(P) with objective function*

$$\mathsf{Adv}(x_1, \ldots, x_N) = \sum_{S \subseteq [N], |S| \leq n} c_S \prod_{k \in S} x_k$$

satisfying $\sum_{k=1}^{N} |c_{\{k\}}| \geq \delta$ and $c_S = 0$ for any set S of even cardinality, achieves $\mathbb{E}[\mathsf{Adv}(x)] \geq \frac{\delta^{3/2}}{8n^{3/4}}$.

Proof. Let $\epsilon > 0$ be a parameter to be determined. We set each x_i randomly and independently to one with probability $(1 + \mathrm{sgn}(c_{\{i\}})\epsilon)/2$. Clearly this implies that $\mathbb{E}[c_{\{i\}} x_i] = \epsilon |c_{\{i\}}|$ and that $|\mathbb{E}[\prod_{k \in S} x_k]| = \epsilon^{|S|}$.

By Cauchy Schwarz inequality and Parseval's identity we have that

$$\sum_{|T|=k} |\hat{P}(T)| \leq \binom{n}{k}^{1/2} \left(\sum_{|T|=k} \hat{P}^2(T) \right)^{1/2} \leq \binom{n}{k}^{1/2}$$

and hence

$$\sum_{|S|=k} |c_S| \leq \binom{n}{k}^{1/2}. \tag{3}$$

We conclude that the advantage of the given algorithm is, given that $c_S = 0$ for even cardinality S, at least

$$\epsilon \sum_{i=1}^{n} |c_i| - \sum_{|S| \geq 3} \epsilon^k |c_S| \geq \epsilon\delta - \sum_{k=3}^{n} \epsilon^k \binom{n}{k}^{1/2}. \tag{4}$$

The sum in (4) is, provided $\epsilon \leq (2\sqrt{n})^{-1}$, and using Cauchy-Schwarz bounded by

$$\left(\sum_{k=3}^{n} \left(\frac{1}{n}\right)^k \binom{n}{k} \right)^{1/2} \left(\sum_{k=3}^{n} (\epsilon^2 n)^k \right)^{1/2} \leq \left(1 + \frac{1}{n}\right)^{n/2} (2\epsilon^6 n^3)^{1/2} \leq 3\epsilon^3 n^{3/2},$$

where we used $\sum_{k=0}^{n} \left(\frac{1}{n}\right)^k \binom{n}{k} = \left(1 + \frac{1}{n}\right)^n$ and $\sum_{k=3}^{n} (\epsilon^2 n)^k \leq \epsilon^6 n^3 \sum_{k=0}^{\infty} \frac{1}{2^k}$ for the first inequality. Setting $\epsilon = \delta^{1/2}(2n^{3/4})^{-1}$, which is at most $(2\sqrt{n})^{-1}$ by (3) with $k = 1$, we see that the advantage of the algorithm is

$$\epsilon\delta - 3\epsilon^3 n^{3/2} = \frac{\delta^{3/2}}{8n^{3/4}}.$$

and the proof is complete.

Let us see how to apply Theorem 4 in the case when P is majority of n variables. Suppose we are given an instance that is $1 - \frac{\delta}{n+1}$ satisfiable and let us consider

$$\sum_{i=1}^{N} c_{\{i\}} \alpha_i \tag{5}$$

where $x_i = \alpha_i$ is the optimal solution and prove that this is large. Any lower bound for this is clearly a lower bound for $\sum_{i=1}^{N} |c_{\{i\}}|$.

Let $\hat{P}_1$ be the value of any Fourier coefficient of a unit size set. Then any satisfied constraint contributes at least $\hat{P}_1$ to (5) while any other constraint contributes at least $-n\hat{P}_1$. We conclude that (5) is at least

$$\left(1 - \frac{\delta}{n+1}\right) \hat{P}_1 - \frac{\delta}{n+1} n\hat{P}_1 = (1-\delta)\hat{P}_1.$$

Using Theorem 4 and the fact that $\hat{P}_1 = \Theta(n^{-1/2})$ we get the following corollary.

Theorem 5. *Suppose we are given an instance of Max-Maj-n which is* $(1 - \frac{\delta}{n+1})$*-satisfiable. Then it is possible, in probabilistic polynomial time, to find an assignment that satisfies a fraction*

$$\frac{1}{2} + \Omega((1-\delta)^{3/2} n^{-3/2})$$

of the constraints.

Let us sketch how to generalize this theorem to predicates other than majority. Clearly the key property is to establish that the sum (5) is large when most constraints can be simultaneously satisfied. In order to have any possibility for this to be true it must be that whenever a constraint is satisfied, then the contribution to (5) is positive and this is exactly being "Chow-robust" as discussed in the introduction. Furthermore, to get a quantitative result we must also make sure that it is positive by some fixed amount. Let us turn to a formal definition.

Recall that the Chow parameters of a predicate P are given by its degree-0 and degree-1 Fourier coefficients, i.e., $\hat{P}(0), \hat{P}(1), \ldots, \hat{P}(n)$ for $i = 1, 2, \ldots, n$. As we are here dealing with an odd predicate, $\hat{P}(0) = 0$. If it holds that

$$P(x) = \operatorname{sgn}(\hat{P}(1)x_1 + \hat{P}(2)x_2 + \cdots + \hat{P}(n)x_n) \qquad \text{for all } x \in \{-1, 1\}^n$$

we say that such a predicate is *Chow-robust* and it is γ-*Chow-robust* iff

$$0 < \gamma \leq \min_{x : P(x)=1} \left(\sum_{j=1}^{n} \hat{P}(j) x_j \right).$$

Note that $\gamma \leq \hat{P}(1)$ and in fact $\gamma = \Theta\left(\frac{1}{\sqrt{n}}\right)$ for majority. Let us state our extension of Theorem 5 in the present context.

Theorem 6. *Let $P(x) = \operatorname{sgn}(w_1x_1 + w_2x_2 + \cdots + w_nx_n)$ be a γ-Chow-robust predicate and suppose that $\mathcal{I}$ is a $1 - \frac{\delta\gamma}{\gamma + \sum_{j=1}^{n} \hat{P}(j)}$ satisfiable instance of Max-CSP(P) where $\delta < 1$. Then there is a probabilistic polynomial time algorithm that achieves $\mathbb{E}[\mathsf{Adv}(x)] = \frac{(1-\delta)^{3/2}\gamma^{3/2}}{8n^{3/4}}$.*

The proof of this theorem is given in the full version of the paper.

Given Theorem 6 it is interesting to discuss sufficient conditions for P to be Chow-robust and we have the following theorem.

Theorem 7. *Suppose we are given positive integers $(w_j)_{j=1}^n$ such that*

$$\beta(w) := 1 - \frac{\sum_{j=1}^{n}(w_j^3 - w_j)}{3\sum_{j=1}^{n} w_j^2} > 0.$$

Further, suppose that for at least $400 \log n$ different values of j, say $1, 2, \ldots, n_1$, we have $w_j = 1$. Then the predicate $P(x) = \operatorname{sgn}(x_1 + \cdots + x_{n_1} + w_{n_1+1}x_{n_1+1} + \cdots + w_nx_n)$ is γ-Chow-robust with $\gamma = \left(\beta(w) - \mathcal{O}\left(\frac{w_{\max}^2}{n}\right)\right)\hat{P}(1)$.

Note that we need n sufficiently large to make γ positive.

Also this proof is postponed to the full version. Let us comment on the condition on the $\Omega(\log n)$ weights that we require to be one. This should be viewed as a technical condition and we could have chosen other similar conditions. In particular, we have made no effort to optimize the constant 400. In our calculations this condition is used to bound the integrand of a complex integral on the unit circle when we are not close to the point $z = 1$ and this could be done in many ways. We would like to point out that although there are choices for the technical condition, some condition is needed. The condition should imply some mathematical form of "when z on the unit circle is far from 1 then many numbers of the form z^{w_j} are not close to 1". Sets of weights violating such conditions are cases when almost all weights have a common factor. An interesting example is the function which, for odd n, has $n - 4$ weights equal to 3 and 4 weights equal to 1. This function is not Chow-robust for any value of n. The above example shows that there are functions with weights of at most 3 that are not Chow-robust. This is a tight bound as the techniques used in the proof of Theorem 7 can be used to show that a function with weights equal to 1 or 2 is Chow-robust.

4 Our main algorithm

We now give an improved algorithm for Max-CSP(P) for homogeneous linear threshold predicates. On almost satisfiable instances, this algorithm achieves an advantage $\Omega\left(\frac{1}{\sqrt{n}}\right)$ over a random assignment in comparison to the $\Omega\left(\frac{1}{n^{3/2}}\right)$ advantage achieved by the adaptation of Hast's algorithm presented in the previous section. However, a drawback of the more advanced algorithm is that we are unable to analyze its advantage on instances that are close to the threshold where Hast's algorithm still achieves a non-trivial advantage. Thus, in order to fully understand the approximability curve, a combination of the algorithm presented below and Hast's algorithm is needed. We now proceed by describing the algorithm. Recall that we write the i'th constraint as

$$P(s_{i,1}x_{l_{i,1}}, \ldots, s_{i,n}x_{l_{i,n}}) = \operatorname{sgn}(L_i(x)),$$

where $L_i(x) = \sum_{j=1}^n w_j s_{i,j} x_{l_{i,j}}$, and let $W := \sum_{j=1}^n w_j$. The algorithm which is parameterized by a noise parameter $0 < \epsilon < 1$ is described as follows:

Algorithm $A_{\mathbf{LP},\epsilon}$

1. Let x^*, Δ^* be the optimal solution to the following linear program
$$\begin{array}{ll} \text{maximize} & \frac{1}{m}\sum_{i=1}^m \Delta_i \\ \text{subject to} & L_i(x) \geq \Delta_i, \forall i \in [m] \\ & x \in [-1,1]^N, \Delta \in [-W,1]^m \end{array}$$
2. Pick $X_1, \ldots, X_N \in \{-1, 1\}$ independently with bias $\mathbb{E}[X_i] = \epsilon x_i^*$ and return this assignment.

As in Theorem 7 we now define $\beta(w)$ for a set of weights $w = (w_1, \ldots, w_n)$ as

$$\beta(w) = 1 - \frac{\sum_{j=1}^n (w_j^3 - w_j)}{3\sum_{j=1}^n w_j^2}.$$

Note that $\beta \leq 1$ for any set of weights, while for majority $\beta = 1$. Further, if $\beta(w) > 0$, then Theorem 7 shows that P is γ-Chow-robust provided that n is large enough.

We have the following theorem whose proof will appear in the full version.

Theorem 8. *Fix any homogeneous threshold predicate $P(x) = \operatorname{sgn}(w_1x_1 + \cdots + w_nx_n)$ having $w_j = 1$ for at least $200 \log n$ different values of j and satisfying $\beta := \beta(w) > 0$. Then, for any $1 - \frac{\delta}{1+W}$ satisfiable instance $\mathcal{I}$ of Max-CSP(P), where $\delta < \beta$, we have*

$$\mathbb{E}[\mathsf{Adv}(A_{LP,\epsilon}(\mathcal{I}))] = (\beta - \delta)^{3/2}\, \Omega\left(\frac{1}{\sqrt{n}}\right) - \mathcal{O}\left(\frac{\log^4 n}{n^{5/6}}\right), \tag{6}$$

where $\epsilon = (\beta - \delta)^{1/2}\epsilon_0$ and $\epsilon_0 > 0$ is an absolute constant.

Thus, for δ bounded away from β, and large enough n, this algorithm is an improvement over the algorithm of Theorem 6. We may also note that both the algorithm $A_{\mathrm{LP},\epsilon}$ and the algorithm of Theorem 6 can be de-randomized using the method of conditional expectation.

As $\beta = 1$ for Maj_n the following result follows directly from Theorem 8:

Corollary 1. *For all $1 - \frac{\delta}{n+1}$ satisfiable instances $\mathcal{I}$ of Max-Maj-n, where $\delta < 1$, we have*

$$\mathbb{E}[\mathsf{Adv}(A_{LP,\epsilon}(\mathcal{I}))] = (1 - \delta)^{3/2}\Omega\left(\frac{1}{\sqrt{n}}\right) - \mathcal{O}\left(\frac{\log^4 n}{n^{5/6}}\right),$$

where $\epsilon = (1 - \delta)^{1/2}\epsilon_0$ and $\epsilon_0 > 0$ is an absolute constant.

5 Unique Games Hardness

The hardness results in this section are under the increasingly prevalent assumption that the Unique Games Conjecture (UGC) holds. The conjecture was made by Khot [11] and states that a specific combinatorial problem known as Unique Games, or Unique Label Cover, is very hard to approximate (see e.g. [11] for more details). The basic tool that we use is the result by Austrin and Mossel [2], which states that the UGC implies that a predicate is approximation resistant if it supports a uniform pairwise independent distribution, and hard to approximate if it "almost" supports a uniform pairwise independent distribution. We now state their result in a simplified form tailored for the application at hand:

Theorem 9 ([2]). *Let $P\colon \{-1,1\}^n \to \{-1,1\}$ be a n-ary predicate and let μ be a balanced pairwise independent distribution over $\{-1,1\}^n$. Then, for any $\epsilon > 0$, the UGC implies that it is NP-hard to distinguish between those instances of Max-CSP(P)*

- *that have an assignment satisfying at least a fraction $\Pr_{x\in(\{-1,1\}^n,\mu)}[P(x) = 1] - \epsilon$ of the constraints;*
- *and those for which any assignment satisfies at most a fraction $|P^{-1}(1)|/2^n + \epsilon$ of the constraints.*

We first give a fairly easy application of the above theorem to the predicate Maj_n. We then generalize this approach to more general homogeneous linear threshold predicates.

Theorem 10. *For any $\epsilon > 0$ the UGC implies that it is NP-hard to distinguish between those instances of Max-Maj-n*

- *that have an assignment satisfying at least a fraction $1 - \frac{1}{n+1} - \epsilon$ of the constraints;*
- *and those for which any assignment satisfies at most a fraction $1/2 + \epsilon$ of the constraints.*

Proof. Consider the following distribution μ over $\{-1,+1\}^n$: with probability $\frac{1}{n+1}$, all the bits in μ are fixed to -1, and with probability $\frac{n}{n+1}$, μ samples a vector with $(n+1)/2$ ones, chosen uniformly at random among all possibilities. To see that this gives a pairwise independent distribution let $X = (X_1, \ldots, X_n)$ be drawn from μ. Then $\mathbb{E}\left[\sum_{i=1}^n X_i\right] = \frac{1}{n+1}\cdot(-n) + \frac{n}{n+1}\cdot 1 = 0$ and $\mathbb{E}\left[\sum_{\substack{i,j=1\\ i\neq j}}^n X_i X_j\right] = \mathbb{E}\left[\left(\sum_{i=1}^n X_i\right)^2\right] - n = \frac{1}{n+1}\cdot(n^2) + \frac{n}{n+1}\cdot 1 - n = 0$. Because of the symmetry of the coordinates, it follows that for all i, $\mathbb{E}[X_i] = 0$ and for every $i \neq j$, $\mathbb{E}[X_i X_j] = 0$. Therefore, the distribution μ is balanced pairwise independent. Theorem 9 now gives the result.

For predicate Maj_n, we can also obtain a hardness result for almost satisfiable instances:

Theorem 11. *For any $\epsilon > 0$ the UGC implies that it is NP-hard to distinguish between those instances of Max-Maj-n*

- *that have an assignment satisfying at least a fraction $1-\epsilon$ of the constraints;*
- *and those for which any assignment satisfies at most a fraction $\frac{1}{2} + c_n \frac{1}{\sqrt{n}} + \epsilon$ of the constraints, where*

$$c_n = \frac{\sqrt{n}}{2^{n-2}} \binom{n-2}{\frac{n-1}{2}} \approx \sqrt{\frac{2}{\pi}}.$$

Proof. Let $k = n - 2$ and consider the predicate $P\colon \{-1,1\}^k \to \{-1,1\}$ defined as $P(x) = \operatorname{sgn}(x_1 + \cdots + x_k + 2)$. Our interest in P stems from the fact that Max-Maj-n is at least as hard to approximate as Max-CSP(P). Indeed, given an instance of Max-CSP(P), we can construct an instance of Max-Maj-n by letting each constraint $P(l_1, \ldots l_k)$ equal $\mathsf{Maj}_n(y_1, y_2, l_1, \ldots, l_k)$ for two new variables y_1 and y_2, that are the same in all constraints and always appear in the positive form. As any good solution to the instance of Max-Maj-n sets both y_1 and y_2 to one, we can conclude that any optimal assignments to the two instances satisfy the same fraction of constraints.

Now consider the following distribution μ over $\{-1,1\}^k$: with probability $\frac{1}{k+1}$, all the bits in μ are fixed to ones, and with probability $\frac{k}{k+1}$, μ samples a vector with $(k+1)/2$ minus ones, chosen uniformly at random among all possibilities. The same argument as in the proof of Theorem 10 shows that the distribution μ is uniform and pairwise independent. Theorem 9 now gives that for any $\epsilon > 0$ the UGC implies that it is NP-hard to distinguish between those instances of Max-CSP(P) that have an assignment satisfying a fraction $1-\epsilon$ of the constraints; and those for which any assignment satisfies at most a fraction

$$\frac{|P^{-1}(1)|}{2^k} + \epsilon = \frac{1}{2^k} \sum_{j=0}^{\frac{k+1}{2}} \binom{k}{j} + \epsilon = \frac{1}{2} + \frac{\binom{k}{\frac{k+1}{2}}}{2^k} + \epsilon = \frac{1}{2} + \sqrt{\frac{2}{\pi k}} + o(1/k) + \epsilon.$$

The result now follows from the observation above that we can construct an instance of Max-Maj-n from an instance of Max-CSP(P) such that optimal assignments to the two instances satisfy the same fraction of the constraints.

Taking the convex combination of the results in Theorems 10 and 11 yields:

Corollary 2. *For any $\delta : 0 \le \delta \le 1$ and any $\epsilon > 0$, the UGC implies that it is NP-hard to find an assignment x to a given $1 - \frac{\delta}{n+1} - \epsilon$ satisfiable instance of Max-Maj-n achieving*

$$\mathsf{Adv}(x) \ge (1-\delta) c_n \frac{1}{\sqrt{n}} + \epsilon,$$

where c_n is the constant defined in Theorem 11.

The above techniques also extend to general weights and we have the following theorem.

Theorem 12. *Suppose we are given positive integers* $(w_j)_{j=1}^n$ *such that* $\sum_{j=1}^n w_j^3 < 100n$ *and* $\sum_{j=1}^n w_j$ *is odd. Further, suppose that for at least* $400 \log n$ *different values of* j *we have* $w_j = 1$*. Let* $P(x) = \mathrm{sgn}(w_1 x_1 + \cdots + w_n x_n)$*, then, for any* $\epsilon > 0$*, the UGC implies that it is NP-hard to distinguish between those instances of Max-CSP(P)*

- *that have an assignment satisfying at least a fraction* $1 - \mathcal{O}\left(\frac{w_{\max}^4}{n}\right) - \epsilon$ *of the constraints;*
- *and those for which any assignment satisfies at most a fraction* $1/2 + \epsilon$ *of the constraints.*

Of course the key to this theorem is to study suitable pairwise independent distributions. In particular, we prove that similar ideas as used in the proof of Theorem 10 can be used to construct *almost* pairwise distributions for more general "majority-like" threshold predicates. As we allow predicates with different weights, the analysis gets more involved and again the problem reduces to estimating complex integrals using the saddle point method. For this reason we need the technical conditions on the weights that were previously discussed after Theorem 7. We then show that such distributions can be slightly adjusted to obtain perfect balanced pairwise distributions and the final result follows by applying Theorem 9. The details will appear in the full version of the paper.

6 Conclusions

We have studied, and obtained rather tight bounds for the approximability curve of "majority-like" predicates. There are still many questions to be addressed and let us mention a few.

This work has been in the context of predicates given by Chow-robust threshold functions. Within this class we already knew, by the results of Hast [7], that no such predicate can be approximation resistant and our contribution is to obtain sharp bounds on the nature of how approximable these predicates are. It is a very nice open question whether there are any approximation resistant predicates given as thresholds of balanced linear functions. It is not easy to guess the answer to this question.

Looking at our results from a different angle one has to agree that the approximation algorithm we obtain is rather weak. For large values of n we only manage to do something useful on almost satisfiable instances and in this case we beat the random assignment by a rather slim margin. On the other hand we also prove that this is the best we can do. One could ask the question whether there is any other predicate that genuinely depends on n variables, accepts about half the inputs and which is easier to approximate than majority. It is not easy to guess what such a predicate would be but there is also very little information to support the guess that majority is the easiest predicate to approximate.

Using the results of Austrin and Mossel, Austrin and Håstad [1] proved that almost all predicates are approximation resistant. One way to interpret the results of this paper is that for the few predicates of large arity where we can get some nontrivial approximation, we should not hope for too strong positive results.

References

1. P. Austrin and J. Håstad. Randomly supported independence and resistance. In *38th Annual ACM Symposium on Theory of Computation*, pages 483–492, 2009.
2. P. Austrin and E. Mossel. Approximation resistant predicates from pairwise independence. *Computational Complexity*, 18(2):249–271, 2009.
3. C. K. Chow. On the characterization of threshold functions. In *Proceedings of the 2nd Annual Symposium on Switching Circuit Theory and Logical Design*, pages 34–38, 1961.
4. M. Dertouzos. *Threshold logic: a synthesis approach.* MIT Press, Cambridge, MA, 1965.
5. I. Diakonikolas and R. A. Servedio. Improved approximation of linear threshold functions. In *IEEE Conference on Computational Complexity*, pages 161–172, 2009.
6. M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42:1115–1145, 1995.
7. Gustav Hast. Beating a random assignment. In *APPROX-RANDOM*, pages 134–145, 2005.
8. J. Håstad. Some optimal inapproximability results. *Journal of ACM*, 48:798–859, 2001.
9. K. R. Kaplan and R. O. Winder. Chebyshev approximation and threshold functions. *IEEE Transactions on Electronic Computers*, EC-14(2):250–252, 1965.
10. Philip Kaszerman. A geometric test-synthesis procedure for a threshold device. *Information and Control*, 6(4):381–398, 1963.
11. S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of 34th ACM Symposium on Theory of Computating*, pages 767–775, 2002.
12. S. Khot, E. Mossel G. Kindler, and R. O'Donnell. Optimal inapproximability results for Max-Cut and other 2-variable CSPs? In *Proceedings of 45th Annual IEEE Symposium of Foundations of Computer Science*, pages 146–154, 2004.
13. R. O'Donnell and R. A. Servedio. The Chow parameters problem. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 517–526, 2008.
14. R. O'Donnell and Y. Wu. An optimal SDP algorithm for Max-Cut and equally optimal long code tests. In *Proceedings of 40th ACM Symposium on Theory of Computating*, pages 335–344, 2008.
15. T. Schaefer. The complexity of satisfiability problems. In *Conference record of the Tenth annual ACM Symposium on Theory of Computing*, pages 216–226, 1978.
16. I.S. Shiganov. Refinement of the upper bound of a constant in the remainder term of the central limit theorem. *Journal of Soviet mathematics*, 35:109–115, 1986.
17. R. O. Winder. Threshold logic in artificial intelligence. *Artificial Intelligence, IEEE Publication*, S-142:107–128, 1963.
18. R. O. Winder. Threshold gate approximations based on Chow parameters. *IEEE Transactions on Computers*, 18(4):372–375, 1969.
19. U. Zwick. Approximation algorithms for constraint satisfaction problems involving at most three variables per constraint. In *Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms*, pages 201–210, 1998.