

ETH-Hardness of Learning Monotone Circuits and Approximating Their Size

Bruno Cavalar
University of Oxford

Susanna F. de Rezende
Lund University

Matthew Gray
University of Oxford

Rahul Santhanam
University of Oxford

July 22, 2026

Abstract

We show the following hardness results for monotone learning and approximation of monotone circuit size:

1. Under the Randomised Exponential-Time Hypothesis (rETH), it requires time $n^{\Omega(\log n)}$ to PAC-learn monotone formulas with n input bits and size $s(n) = n$ by monotone circuits of size $n^{(\log n)^{1-\epsilon}}$, for every $\epsilon > 0$.
2. Under the Randomised Exponential-Time Hypothesis (rETH), for any $\delta > 0$, there is a polynomially bounded function m such that $m^{1-\delta}$ -multiplicatively approximating the minimum monotone circuit size of a monotone function consistent with a sequence of $m(n)$ labelled examples $\{(x_i, b_i)\}$ over n -bit inputs requires time $m^{\Omega(\log(m))}$.

Our results are shown by a novel application of lifting arguments in proof and communication complexity to hardness of monotone learning, by building on the seminal result of Atserias and Müller [AM20] on hardness of automating Resolution proofs.

1 Introduction

Meta-Complexity One of the main goals of the area of *meta-complexity* is to prove strong hardness results for computing complexity measures such as circuit size, formula size and monotone circuit size. While such hardness results can be derived in a standard way from cryptographic assumptions such as the existence of one-way functions [RR97, KC00], it has proven much more challenging to show hardness results under *worst-case assumptions*, such as NP-hardness results. A good reason for this difficulty was identified by Kabanets and Cai [KC00], who showed that NP-hardness of the Minimum Circuit Size Problem (which asks whether a Boolean function F given by its truth table has circuits of size at most s , for a

given parameter s) via polynomial-time honest reductions would imply hard-to-prove circuit lower bounds such as super-polynomial circuit lower bounds for EXP.

Despite this obstacle, there are now several hardness results for various notions of circuit size under worst-case hardness assumptions, including DNF size, Boolean formula size (under ETH), multi-output circuit size, oracle circuit size, etc., [Mas79, AHM⁺08, HOS18, ILO20, Ila20a, Ila20b, Ila21, Hir22, HIR25]. A common feature of these results, except for the recent inapproximability result of [HIR25] for oracle circuit size, is that they do not show strong inapproximability results for circuit size. Here, by a strong inapproximability result, we mean the hardness of approximating the circuit size within a $N^{\Omega(1)}$ multiplicative factor, where N is the size of the function representation.

A major motivation for showing strong inapproximability result for circuit size under worst-case assumptions is that they offer a pathway to worst-case to average-case reductions for NP. It was shown in [San20, Hir23], building on work of [CIKK16], that there are non-black-box reductions from strong inapproximability of problems such as circuit size and average circuit size to average-case hardness for these problems. Thus, NP-hardness of strong inapproximability for either of these problems would rule out Heuristica, the complexity world imagined by Impagliazzo [Imp95] where NP is hard in the worst case but easy on average. A recent result of Mazon and Pass [MP24] rules out strong inapproximability of circuit size under Levin reductions from SAT, assuming a standard cryptographic assumption, but strong inapproximability might still hold under a more general notion of reduction such as Turing reductions. The strong inapproximability result of [HIR25] for oracle circuit size is based on cryptographic techniques and seems difficult to apply to standard circuit size measures. It could potentially be very useful to find other techniques yielding strong inapproximability.

In terms of circuit size measures for which hardness under worst-case assumptions is known, monotone circuit size has been little-studied despite its central role in complexity theory. Showing NP-hardness for the Minimum Monotone Circuit Size Problem (which asks whether a monotone Boolean function given by its truth table has monotone circuit size at most s , for a given parameter s) would be a significant step towards NP-hardness for the Minimum Circuit Size Problem, as there is a well-known reduction from Minimum Circuit Size to Minimum Monotone Circuit Size.

Learning One of the main goals of the area of *learning theory* is to identify which learning problems are hard in the PAC-learning model. When PAC-learning a concept class $\mathcal{C}$, we are given samples $(x, f(x))$ from a distribution $\mathcal{D}$, where f belongs to $\mathcal{C}$, and our task is to output with high probability a hypothesis c that approximates f well according to the distribution. If the output hypothesis c is in $\mathcal{C}$, the learning algorithm is said to be *proper*, otherwise it is said to be *improper*. We also consider a *semi-proper* version of learning $\mathcal{C}$ -circuits by $\mathcal{C}'$ -circuits, where $\mathcal{C}' \supseteq \mathcal{C}$.

It is known how to rule out even *improper learning* of various concept classes under cryptographic assumptions [KV94, Kha93, DLM⁺09]. More recently, there has been exciting

work on ruling out improper learning even of weak concept classes such as DNFs [Dan16, DS16, DV21] based on average-case assumptions such as Feige’s hypothesis [Fei02] and its variants. However, just as in the case of meta-complexity problems, showing hardness under worst-case assumptions remains challenging. This was explained in a work of Applebaum, Barak and Xiao [ABX08], who showed that NP-hardness of learning general Boolean circuits under Karp reductions or non-adaptive reductions would imply the collapse of the Polynomial Hierarchy. This leaves open the possibility of using more general reductions such as Turing reductions to show hardness.

Known worst-case hardness results are typically for proper learning of weak circuit classes such as DNFs, or for semi-proper learning where the output hypothesis class $\mathcal{D}$ is not too much more powerful than the concept class [ABF⁺08]. While hardness of learning even restricted subclasses of monotone circuits under the uniform distribution is known under cryptographic assumptions [DLM⁺09], and unconditional hardness results are known in the statistical query learning framework [FLS11], very little seems to be known about the worst-case hardness of learning monotone circuits.

As can be seen from the discussion above, there are several parallels between the situation with hardness of meta-complexity problems and the situation with hardness of learning. These parallels are not coincidental: learning for a circuit class $\mathcal{C}$ can be seen as a *search version* of the problem of determining minimum $\mathcal{C}$ -size of a Boolean function. Indeed, given access to a Boolean function f , outputting a circuit of size s that computes f (perhaps approximately) via a learning algorithm witnesses that f has (perhaps approximate) circuit size at most s . This connection is exploited to derive new learning algorithms in [CIKK16].

In this paper, we aim to establish strong results on the hardness of monotone learning based on worst-case assumption, and then to leverage the connection between monotone learning and meta-complexity to get strong inapproximability of monotone circuit size.

The meta-complexity problem we consider is the following: given $m = \text{poly}(n)$ labelled examples (x_i, b_i) , where each x_i is an n -bit string and b_i a bit, and a parameter s , is there a monotone function f of monotone circuit size at most s that is consistent with the labelled examples, i.e., $b_i = f(x_i)$ for each i . In the context of learning, this is known as the Occam Learning problem, and the non-monotone analogue of this problem was called Partial Minimum Circuit Size Problem in [ILO20], where it was shown to be NP-hard. We are interested here in strong inapproximability results, which the techniques of [ILO20] are incapable of proving, since they reduce from the well-approximable Set Cover Problem.

1.1 Our Results

Our first main result is on the hardness of PAC-learning monotone formulas. Proving hardness of improperly PAC-learning monotone formulas is at least as hard as showing NP-hardness of learning general Boolean circuits, and hence beyond reach at this point in time. However, we are able to prove a strong result in the *semi-proper* setting under worst-case

assumptions, where our concept class is the class of monotone formulas, but our hypothesis class is the class of polynomially larger monotone circuits. In this setting, we are able to get a quasipolynomial lower bound on the running time of any PAC-learning algorithm, assuming the randomised Exponential Time Hypothesis. Recall that the randomised Exponential Time Hypothesis (rETH) postulates that any algorithm solving 3SAT on formulas of n variables must take time $2^{\Omega(n)}$. This is a standard hardness assumption in fine-grained complexity and parameterized complexity.

Theorem 1.1 (Informal). *Under rETH, for every $\alpha > 0$, monotone formulas of size n on n variables require time $n^{\Omega(\log n)}$ to be PAC-learned by monotone circuits of size $n^{(\log n)^{1-\alpha}}$.*

The formal statement of Theorem 1.1 is in the first item of Theorem 4.11 in Section 4.4. To the best of our knowledge, it was not known previously even how to get hardness of learning monotone circuits of size n by monotone circuits of size $n^{1+\alpha}$ under standard worst-case assumptions. In fact, our techniques yield a tradeoff, where we can show hardness even for learning monotone formulas of even slightly super-polylogarithmic monotone formula size by monotone circuits of polynomial size. This result is stated in the second item of Theorem 4.11.

Our second main result shows the strong inapproximability of monotone circuit size for the partial version of the Minimum Monotone Circuit Size Problem, where we are given a sequence of labelled examples as input. The assumption used here is again rETH.

Theorem 1.2. *For each $c > 1$ the following holds under rETH: there is no algorithm running in time $N^{o(\log N)}$ which, given as input a sequence of $O(n^c \log n)$ labelled examples (x_i, b_i) representable by N bits in all, where each x_i is an n -bit string and b_i a bit, can distinguish between the case that these labelled examples are consistent with a monotone formula of size n (i.e., there is f of monotone formula size n for which $f(x_i) = b_i$ for each i) and the case that these labelled examples are inconsistent with a monotone circuit of size n^c . In particular, for arbitrary $\epsilon > 0$ there is no $N^{o(\log N)}$ time algorithm which $N^{1-\epsilon}$ -multiplicatively approximates the partial monotone formula size or the partial monotone circuit size.*

The formal statement of Theorem 1.2 is in the fourth item of Corollary 4.9.

Both Theorem 1.1 and Theorem 1.2 are consequences of a more general hardness result for a monotone version of the Computational Gap Learning problem studied in [ABX08]. In our problem, which we call Monotone Circuit-Formula Gap Learning, we are given a circuit which samples a distribution of labelled examples (x_i, b_i) where x_i is of length n and b_i a bit, and are asked to distinguish between 2 cases: the first where the distribution is completely consistent with a formula of size at most s_1 , and the second where the distribution does not even non-trivially correlate with any circuit of size at most s_2 . Here $s_1, s_2 : \mathbb{N} \rightarrow \mathbb{N}$ are parameters of the problem, which we are able to choose so that $s_2 \gg s_1$. Theorem 1.1 follows immediately from our result about Monotone Circuit-Formula Gap Learning, while Theorem 1.2 follows by a standard sampling argument together with a union bound.

To show our central result about Monotone Circuit-Formula Gap Learning, we use a novel technique in the domain of hardness of learning: *lifting*, together with known strong inapproximability results for Resolution proof size [AM20, dRGN⁺21, GKMP20, dR21, CdR25]. One of the main applications of lifting is to obtain monotone circuit size lower bounds in a generic way from Resolution proof width lower bounds [GGKS20, GKRS19]. The generic nature of these reductions opens the possibility of using them to derive *reductions* from approximating proof size for various proof systems to approximating circuit size for various circuit models. This is yet another instance of the connection between learning and automatability of proof systems proposed in [ABF⁺08].

It may be useful to compare our results with the work of Hirahara [Hir22]. Hirahara shows that a variant of the Computational Gap Learning problem is NP-hard (under randomised reductions), where the task is to distinguish between linear programs (Turing machines computing a linear function) of size s_1 and programs of size $s_2 = s_1 \cdot n^{o(1)}$. Besides the obvious difference that we consider monotone circuits vs. formulas whereas [Hir22] considers programs, crucially our results allow for a superpolynomial gap between circuit size, whereas [Hir22] obtain only a sublinear gap between program size (though under the weaker assumption of NP-hardness). Hirahara also shows NP-hardness of approximating partial circuit size, when one is given the entire truth table of a partial function (referred to as MCSP*), again only with a $n^{o(1)}$ approximation gap. In Theorem 1.2, however, we consider a *succinct* version of this problem, in which we are only given inputs where the function is defined; we also obtain a stronger hardness of approximation gap here.

We next give a technical overview of the ideas behind our proofs.

1.2 Technical Overview

Overall approach The main insight of the paper is to convert the hardness of *automating* Resolution [AM20] into the hardness of approximating monotone circuit size, using results from the literature (called *lifting theorems*) that connect the complexity of Resolution refutations with the monotone complexity of Boolean functions [GGKS20].

To understand how this works, let us review how the seminal paper of Atserias and Müller [AM20] proved that automating Resolution is NP-hard. Recall that an algorithm that automates Resolution must, given as input an unsatisfiable CNF formula ϕ , construct a Resolution refutation of ϕ in time $\text{poly}(|\phi| + s)$, where s is the size of the shortest Resolution refutation of ϕ . Given a 3-CNF formula F on n variables as input, the proof of [AM20] consists in constructing in polynomial-time a formula $\text{Ref}(F)$ with the following properties:

1. If F is satisfiable, then $\text{Ref}(F)$ admits a polynomial-size Resolution refutation;
2. If F is unsatisfiable, then any Resolution refutation of $\text{Ref}(F)$ must have size $2^{n^{\Omega(1)}}$.

In turn, this implies that if we can *automate* Resolution, then we can determine in which of the two cases we are in, and thus decide the satisfiability of F .

Another significant ingredient in our approach is the *DAG-like lifting theorem*, which converts Resolution *width* lower bounds for a formula ϕ into monotone circuit size lower bounds for a related monotone function f_ϕ [GGKS20]. Importantly, the monotone complexity of the resulting function f_ϕ is tightly connected to the Resolution width of ϕ , so that f_ϕ has large monotone complexity iff ϕ requires large Resolution width.

Ideally, we would now like to combine the properties of $\text{Ref}(F)$ with the lifting theorem to reduce the satisfiability of F to deciding the monotone complexity of $f := f_{\text{Ref}(F)}$ when given access to some succinct representation of f (e.g., access to random samples of f , or a succinct representation of its truth-table). Unfortunately, the connection between monotone complexity and Resolution goes only via Resolution *width*, whereas the complexity gap obtained by Atserias-Müller is only for Resolution *size*. Though Resolution size does indeed give *upper bounds* on the monotone complexity of a related function via *feasible interpolation* [Kra97, Kra98, Raz95], the converse is not known to hold.

A potential way out would be to consider the notion of *block-width* of proofs [AM20], a generalisation of Resolution-width to the setting when the variables are split into blocks, as is the case with $\text{Ref}(F)$. The proof of Atserias-Müller can be rephrased as first showing that the block-width of $\text{Ref}(F)$ is $O(1)$ when F is satisfiable, and $n^{\Omega(1)}$ when F is unsatisfiable [GKMP20]. However, though the notion of block-width can be lifted to Resolution size, it's not known if it can be lifted to monotone circuit size, and indeed significant technical challenges have been found when trying to do so in at least one previous work [GKMP20].

A more promising approach comes from considering a different encoding of $\text{Ref}(F)$ found in [dRGN⁺21]. Their encoding admits a Resolution-width gap of $\text{Ref}(F)$ of the following form¹:

1. If F is satisfiable, $\text{Ref}(F)$ has Resolution-width $O(n)$;
2. If F is unsatisfiable, $\text{Ref}(F)$ has Resolution-width $\Omega(n^3)$.

This seems to be what we are aiming at: an efficiently constructible formula with a noticeable gap in Resolution-width. However, lifting gives us that the monotone complexity of f_ϕ is roughly $2^{\Theta(\text{ResWidth}(\phi))} \cdot \text{Res}(\phi)$. So we obtain a function on $n^{O(1)}$ input bits with $2^{O(n)}$ monotone complexity when F is satisfiable, and $2^{\Omega(n^3)}$ when F is unsatisfiable. Unfortunately, it is not clear how to distinguish between these two cases other than by guessing the *exponential size* circuit that computes the function. In other words, we have obtained a reduction from an NP problem to an NEXP problem, which is hardly surprising. Moreover, the reduction is ultimately useless, as we could just solve SAT directly in exponential time.

We overcome this hurdle by constructing a different formula $\text{Ref}^*(F)$ with the following properties:²

¹The encoding of [AM20] cannot achieve this gap because it uses unary pointers.

²For simplicity of exposition, at this introduction we only explain the construction with the parameters that give a polynomial gap between linear-size formulas and monotone circuits. The construction for “juntas” (i.e., functions that do not depend on all inputs) is the same, but with a different parametrisation.

1. $\text{Ref}^*(F)$ has $2^{O(\sqrt{n})}$ variables and $2^{O(\sqrt{n})}$ clauses.
2. $\text{Ref}^*(F)$ admits Resolution proofs of depth at most $O(\sqrt{n})$ when F is satisfiable;
3. $\text{Ref}^*(F)$ requires Resolution proofs of width at least $2^{\Omega(\sqrt{n})}$ when F is unsatisfiable.

Note that we obtain not just a Resolution width upper bound in the satisfiable case, but the strictly more powerful upper bound on Resolution *depth*, which is strongly tied to the size of monotone *formulas* [RM99, GPW18]. This is how we obtain a gap between monotone formula and circuit sizes in the final result. Note moreover that the resulting formula has *subexponential* size. This is a careful tradeoff that we exploit: by increasing the size of the formula, we are able to reduce its Resolution-width (and even its Resolution-depth). We now describe the construction of the formula in more detail.

The $\text{Ref}^*(F)$ Formula The formula $\text{Ref}(F)$ essentially encodes the statement “ F has a Resolution proof of small size s ”, where s is a fixed polynomial in the number of variables of F . When F is satisfiable, the formula $\text{Ref}(F)$ is *unsatisfiable*, and this formula can be refuted with low *block-width* [AM20, Pud03].

Their strategy can be seen as low memory technique to achieve the following. Given any proposed proof for F , a *refuter*³ navigates through it while maintaining the guarantee that the current clause is unsatisfied by the known satisfying assignment x^* . Beginning at $\perp$, eventually such a refuter will make it to a line of the proof that is a weakening of an axiom. Since the line must be unsatisfied by x^* , it cannot be a weakening of any axiom and a contradiction has been found. While this technique can be used to show an $O(1)$ block-width upper bound, the depth of the proof strategy is $\Omega(s \cdot n)$, since it will involve learning every variable of $\Omega(s)$ lines of the proposed proof. Adapting this technique to achieve a $\sqrt{n}$ depth upper bound seems like a nearly hopeless endeavour on first blush. Achieving depth $\sqrt{n}$ requires us to refute $\text{Ref}(F)$ while querying in total fewer variables than the number of variables of F . This presents two barriers:

- The refuter needs to get to a line of the proof that is claimed to be a weakening of an axiom while only seeing variables from $\sqrt{n}$ total blocks, and
- The refuter needs to learn some information about all n variables while being able to query far fewer than n variables.

To overcome the first of these issues, we can force the proof to have depth $\sqrt{n}$, as done in [dR21]. However, if we are only resolving over one literal at a time, the clauses at depth- $\sqrt{n}$, which are weakening of axioms of F , would have width $\sqrt{n}$. This would break the lower

³The “Refuter” is typically referred to as “Prover” in a Prover-Adversary game. Since F is satisfiable, the goal of the Prover is to find a contradiction to the statement that F admits a Resolution refutation, and so its task can be referred to as “refuting” a proposed proof of unsatisfiability of F .

bound of Atserias and Müller because their strategy requires that the clauses in the purported proof that are claimed to be a weakening of an axiom must have width n .

To retrieve the lower bound, we can change the underlying proof system to one which resolves on $\sqrt{n}$ variables at a time. This results in each block, instead of having just a left and a right child—which are the two clauses from which it was derived—now having $2^{\sqrt{n}}$ children, and thus having one pointer point_z for each possible setting $z \in \{0, 1\}^{\sqrt{n}}$ of the variables being resolved over (we also importantly fix the order in which variables are resolved over, with the first segment of $\sqrt{n}$ variables, then the second segment, etc). Unfortunately, this immediately means that the resulting $\text{Ref}(F)$ formula will have to be of size at least $2^{\sqrt{n}}$ just to manage these pointer variables. On the positive side, this gives us a hint for how to overcome the second barrier.

In standard resolution refutations, if you know that a clause B' is the left child of B which is the result of resolving over x_i , you know that x_i is in B' , and more importantly that $\neg x_i$ is *not* in B' ; and if B' is the right child you know that $\neg x_i$ is in B' , and more importantly that x_i is *not* in B' . Consequently, knowing that x^* is a satisfying assignment, if the prover always queries whichever child is not satisfied by x^* , they will eventually find themselves in a leaf which cannot be a valid weakening. In our new multi-variable resolution system, if you know that B' is the z^{th} child of B , you know $\sqrt{n}$ bits of information about the clause that B' claims to have. Using x^* as above, one could in principle get to a leaf which cannot be a valid weakening in depth $O(\sqrt{n})$.

We can extend this density of information from the pointers to a new class of summary variables $\text{summ}_{i,z}^B$ which give us information about all the variables from $x_{(i-1)\sqrt{n}}$ to $x_{i\sqrt{n}-1}$. This idea of using summary variables appeared before in [CdR25]. It is important that our summary variables encode which literals are *not* present (which is ultimately what we are concerned with for the upper bound). Formally, $\text{summ-unlit}_{i,z}^B$ is set equal to 1 to indicate that for each $z_j = 1$, literal $x_{(i-1)d+j}$ is not present in block B , and that for each $z_j = 0$, literal $\neg x_{(i-1)d+j}$ is not present in block B . We can ensure that a Prover can learn what all $\sqrt{n}$ of the $\text{summ-unlit}_{i,z}^B$ variables for a leaf node should be while only querying the right set of $O(\sqrt{n})$ variables on their way from the root to a leaf.

To make our bounds as tight as possible, we add some additional complexity, including binary pointers for axioms, split variables to reduce clause width, and a low degree expander between each layer of the proof. This low-degree expander, which was previously used in [dR21, CdR25], allows us to keep down to a constant the number of $\text{point}_{B',z}^B$ variables a prover needs to query for each B, z before finding one that is set to one. While the expander does somewhat complicate the proof of the lower bound, we are able to derive it by reducing the retraction pigeonhole principle (rPHP) instance to finding a violated clause, similarly to what was done in [dR21].

Hardness of approximating circuit size and learning To see how to join our new Ref^* formula with a lifting theorem to obtain the claimed hardness results, it will be helpful to understand

the function that comes from lifting. Roughly speaking, “lifting” usually consists in composing some original computational object (e.g., a propositional formula or a function) with a gadget, so as to hide the original object in a stronger computational model in such a way that the only way to compute/solve the composed problem is by solving the original problem.

In our case, given a formula F with n variables and width k , together with a parameter m , the lifted function $f := f_{F,m}$ from [GGKS20] will have $|F| \cdot m^k$ input bits, where $|F|$ is the number of clauses of F . To each $i \in [|F| \cdot m^k]$, we associate a pair (C, α) , where C is a clause of F and $\alpha \in [m]^k$. We interpret (C, α) as the k -width clause over the variables $\{y_{ij} : i \in [n], j \in [m]\}$ obtained by replacing the variable $x_i \in \text{vars}(C)$ in C with $y_{i, \alpha(i)}$. We can thus see each input to f as a collection of clauses (C, α) . The task of the partial function $f_{F,m}$ is to distinguish between the following two cases:

1. Accept (Output 1): there exists $x : [n] \rightarrow [m]$ such that the input bit corresponding to $(C, x^{\text{vars}(C)})$ is set to 1 for every clause C of F , and everything else is set to 0;
2. Reject (Output 0): there exists $y \in \{0, 1\}^{mn}$ such that all input bits associated with clauses (C, α) satisfied by y are set to 1, and everything else is set to 0.

The partial function is clearly well-defined, as in the first case the collection of clauses is unsatisfiable, and in the second case they are satisfiable. For this same reason, no input to the function which contains all the clauses of Case (1) will ever be in Case (2), so the function is monotone. One can think of the function f as “hiding” several variable renamings of the the formula F in its input bits; to compute f is to “sniff out” if F has been completely embedded among the renamed clauses.

Most importantly for our purposes, we can easily sample from the distribution which half of the time gives a uniform accepting input of f , and in the other half gives a uniform rejecting input. Indeed, first sample a bit $b \in \{0, 1\}$, then, if $b = 1$, sample $n \log m$ bits and interpret it as $x : [n] \rightarrow [m]$, otherwise sample a string $y \in \{0, 1\}^{mn}$. With x (resp. y), it’s then easy to set the relevant bits so as to output an accepting (resp. rejecting) input of $f_{F,m}$. Let us abuse notation and denote by $\mathcal{D}^{F,m}$ both a circuit that samples from this distribution of binary strings (together with the bit b as a prefix) and the distribution itself. We thus see samples from $\mathcal{D}^{F,m}$ as a pair (x, b) where $x \in \{0, 1\}^{|F|m^k}$ and $b \in \{0, 1\}$. Given F and m as input, it’s also easy to see that we can construct the circuit sampling from $\mathcal{D}^{F,m}$ in time $\text{poly}(|F|, m^k)$.

Going back to the approach above, given an input F to 3SAT, our reduction will first construct $\mathcal{D} := \mathcal{D}^{\text{Ref}^*(F), m}$, so as to obtain a distribution with the following properties. If F is satisfiable, then

$$\text{there is a monotone formula } L \text{ of size } \leq s_1 \text{ s.t. } L(x) = b \text{ for every } (x, b) \in \text{supp}(\mathcal{D}); \quad (1)$$

if F is unsatisfiable, then

$$\text{every monotone circuit } K \text{ of size } \leq s_2 \text{ satisfies } \mathbb{P}_{(x,b) \leftarrow \mathcal{D}}[K(x) = b] < 1/2 + \gamma. \quad (2)$$

The value of s_1 is $2^{O(\sqrt{n} \log m)}$, and the value of s_2 is $2^{\Omega(m/\sqrt{n})}$; these come from a refined version of the DAG-like lifting theorem [dRV25] with “small gadgets” that allows us to essentially pick any value of m we would like. However, lifting theorems typically are only stated for $\gamma = 1/2$ (corresponding to a worst-case lower bound). Fortunately, generalising to $\gamma \approx m^{-1}$ is an easy task, requiring only the modification of certain numbers and parameters in the proof of [dRV25]. This is essential for our purposes, as in a learning problem we only expect to approximate the target concept.

We are now ready to state the computational problem for which we show a lower bound. The *Monotone Circuit-Formula Gap Learning problem* $\text{mCFL}_{s_1}^{s_2}[\gamma]$ is the problem of distinguishing between cases (1) and (2) when given $\mathcal{D}$ as input. Given that our $\text{Ref}^*(F)$ formula has size $2^{O(\sqrt{n})}$, the reduction from SAT to this problem takes this time. Under ETH, we thus obtain an $N^{\Omega(\log N)}$ lower bound for this problem, since the size N of the input is $2^{O(\sqrt{n})}$. To reduce to the case $s_1 = n$, we pad the sampled strings; this takes time $2^{\tilde{O}(\sqrt{n})}$, so our final lower bound is $N^{\tilde{\Omega}(\log N)}$.

We can further reduce this problem to a gap version of Partial-Monotone-MCSP we call $\text{mMCFSP}_{s_1}^{s_2}[\gamma]$ ⁴. In this problem, instead of $\mathcal{D}$, the algorithm is given t examples $(x_1, b_1), \dots, (x_t, b_t)$, and must distinguish whether there is a monotone formula L of size s_1 such that $L(x_i) = b_i$ for all $i \in [t]$, or whether every monotone circuit C of size s_2 is such that $|\{i \in [t] : C(x_i) \neq b_i\}| \geq (1/2 + \gamma)t$. The reduction is randomised, naturally coming from sampling the distribution $\mathcal{D}$. Moreover, we need to sample $O(s_2 \log s_2)$ examples so that, with high probability, any circuit of size at most s_2 will err on a proportion larger than $1/2 + \gamma$ of the examples (in case (2)), and so the reduction takes time $O(s_2 \log s_2)$. This means that we get a tradeoff between the choice of s_2 and the runtime lower bound coming from rETH. The same reduction also gives us lower bounds on PAC-learning.

2 Preliminaries

For $f, g : \mathbb{N} \rightarrow \mathbb{N}$, we write $f \sim g$ if $\lim f/g \rightarrow 1$, and $f \asymp g$ if $f = \Theta(g)$. We say that a function Boolean f is γ -approximated by g over a distribution $\mathcal{D}$ if $\mathbb{P}_{x \leftarrow \mathcal{D}}[f(x) = g(x)] \geq 1/2 + \gamma$.

⁴Sometimes the literature reserves the term “Partial-MCSP” to the problem where one is given the entire truth table of a partial function, but we follow here the convention of [ILO20].

2.1 Basic Complexity

Monotone circuits are Boolean circuits without $\neg$ gates. *Formulas* are Boolean circuits whose gates have fan-out 1. *Monotone formulas* are analogously defined.

The *Exponential-Time Hypothesis (ETH)* postulates that there exists a constant $\varepsilon > 0$ such that any algorithm solving 3SAT on n -variate formulas must take time $2^{\varepsilon n}$. The *randomised ETH (rETH)* postulates the same lower bound for randomised algorithms.

2.2 Expander Graphs

A (m, n, δ, r, e) -bipartite expander graph is a bipartite graph $G = ((U \cup V), E)$ with $|U| = m$, $|V| = n$, with every vertex in U having degree at most δ , and where for every subset $S \subseteq [m]$ such that $|S| \leq r$, the neighbourhood of S , denoted $N(S)$, satisfies $|N(S)| \geq e|S|$.

Remark 2.1. A complete bipartite graph $G = ((U \cup V), E)$ with $|U| = m$, $|V| = n$, is a (m, n, δ, r, c) -bipartite expander graph for, e.g., $\delta = n$, $r = n/2$ and $c = 2$.

A standard calculation [HLW06] shows that random graphs are good expanders. We record one parameter regime that we use in our applications.

Lemma 2.2. Let $m, n \in \mathbb{N}$ be such that $m = n^2$. Then there exists $\delta = O(1)$ and $r = \Omega(\sqrt{n})$ such that with high probability a random graph $G \sim \mathbf{G}(m, n, \delta)$ is an $(m, n, \delta, r, \delta/2)$ -bipartite expander graph.

2.3 Proof Complexity

Resolution A *literal* is a propositional atom or its negation, a *clause* is a disjunction of literals, and a *CNF formula* is a conjunction of clauses. We see clauses as sets of literals, and write simply $C \subseteq D$ to express that C is a subclause of D . We denote the set of variables in a clause C by $\text{vars}(C)$ (if x_i or $\neg x_i$ are present in C , we have $x_i \in \text{vars}(C)$). A *Resolution refutation* of an unsatisfiable CNF formula $\varphi = C_1 \wedge \dots \wedge C_m$ over variables $x_1, \dots, x_n$ is a sequence $D_1, \dots, D_s$ of clauses over $x_1, \dots, x_n$ such that $D_s = \perp$, denoting the empty clause, and for every $i \in [s - 1]$, the clause D_i either (a) is one of the clauses $C_1, \dots, C_m$ of φ , or (b) is a *weakening* of a previous clause, meaning that $D_i \supseteq D_j$ for some $1 \leq j < i$, or (c) has been obtained from two previous clauses $D_j = A \vee x$ and $D_k = B \vee \neg x$, for $j, k < i$, by an application of the *Resolution rule*:

$$\frac{A \vee x \quad B \vee \neg x}{A \vee B}$$

We say that $A \vee B$ was obtained by *resolving* over x . The *size* of the refutation is s , the number of clauses appearing in the refutation.

To every Resolution refutation we can associate a directed acyclic graph (dag) in a natural way. We define the *depth* of the proof as the length of the longest path in the dag, starting

from the root labeled by the empty clause $\perp$. The proof is said to be *tree-like* if the associated graph is a tree rooted at $\perp$.

Let F be an unsatisfiable k -CNF formula. We denote by $\text{ResWidth}(F)$ the minimum integer such that F admits a Resolution refutation whose clauses all have width at most $\text{ResWidth}(F)$. Let $\text{ResDepth}(F)$ be the minimum integer such that F admits a Resolution refutation with depth at most $\text{ResDepth}(F)$. We denote by $\text{Res}(F)$ the size of the smallest Resolution refutation of F , and by $\text{TreeRes}(F)$ the size of the smallest tree-like such refutation. We denote by $|F|$ the number of clauses in F .

Prover-adversary game An equivalent characterisation of Resolution width and Resolution depth can be given by notions from *query complexity*. Suppose that F is an n -variate unsatisfiable CNF formula. Given any assignment $x \in \{0,1\}^n$ to the variables of F , some clause of F is necessarily falsified. The minimum number k of queries to the variables of F needed to discover one such clause is equivalent to $\text{ResDepth}(F)$.

To characterise ResWidth , it is helpful to modify the above setup in what is commonly called a “Prover-Adversary game”. As before, the Prover is querying variables, but now the Prover is allowed to forget the value of some queried variables. The adversary is answering the queries, but may respond with different values when a variable whose value was forgotten is queried again. The minimum number k of variables that the Prover needs to store in order to discover a falsified clause is equal to $\text{ResWidth}(F) \pm 1$ [AD08].

2.4 Lifting for Average-Case Lower Bounds

Given a *gadget* $g : \mathcal{X} \times \mathcal{Y}$, let $g^n : \mathcal{X}^n \times \mathcal{Y}^n$ denote the function

$$g^n(x, y) = (g(x_1, y_1), \dots, g(x_n, y_n)).$$

Let $\text{Ind}_m : [m] \times \{0,1\}^m \rightarrow \{0,1\}$ denote the “index gadget”, defined as the Boolean function such that $\text{Ind}_m(x, y) = y[x]$.

Definition 2.3. Let F be an unsatisfiable formula with n variables. Let $f_{F,m}$ be the partial Boolean function over $|F| \cdot m^k$ bits defined as follows. To each $i \in [|F| \cdot m^k]$, associate a pair (C, α) , where C is a clause of F and $\alpha \in [m]^k$. We interpret (C, α) as the k -width clause over the variables $\{y_{i,j} : i \in [n], j \in [m]\}$ obtained by replacing the variable $x_i \in \text{vars}(C)$ in C with $y_{i,\alpha(i)}$. Each binary string of length $|F| \cdot m^k$ is thus in 1-to-1 correspondence with the set

$$F \circ \text{Ind}_m^n := \{(C, \alpha) : C \text{ is a clause of } F, \alpha : \text{vars}(C) \rightarrow [m]\}.$$

The function $f_{F,m}$ has the following behaviour:

1. Accept the binary string whose support is $\{(C, \alpha) \in F \circ \text{Ind}_m^n : x^{\text{vars}(C)} = \alpha\}$, for every $x : [n] \rightarrow [m]$.

2. Reject the binary string whose support is $\{(C, \alpha) \in F \circ \text{Ind}_m^n : (C, \alpha)(y) = 1\}$, for every $y \in \{0, 1\}^{nm}$.

The following theorem shows that the monotone circuit complexity of $f_{F,m}$ is roughly $m^{\Theta(\text{ResWidth}(F))}$ for some $m \gg \text{ResWidth}(F)$; moreover, a formula upper bound of the form $m^{O(\text{ResDepth}(F))}$ also holds. This result follows from (a small modification of) the so-called *lifting theorem*, which first appeared in [GGKS20], from which monotone circuit lower bounds follow from Resolution width lower bounds. We use the lifting theorem from [dRV25], which builds on [dRFJ⁺24, LMM⁺22], as it not only has better parameters, but also gives a lower bound for $f_{F,m}$ for smaller values of m , which is essential to our purposes.

We say that a circuit $D : \{0, 1\}^r \rightarrow \{0, 1\}^N$ represents a distribution $\mathcal{D}$ over N -input bits if the random binary string $D(\mathbf{r})$ where $\mathbf{r} \leftarrow \{0, 1\}^r$ is distributed according to $\mathcal{D}$.

Theorem 2.4 ([dRV25, Theorem 11, modified]). *There exists an algorithm $A(F, 1^m)$ which receives as input an unsatisfiable k -CNF formula F on N variables, runs in time $\text{poly}(|F| \cdot m^k)$, and outputs a circuit representing a distribution $\mathcal{D}^{F,m} = (\mathcal{D}_1 + \mathcal{D}_0)/2$ over $|F| \cdot m^k$ bits such that:*

1. *There is a monotone circuit K of size $m^{O(\text{ResWidth}(F))} \cdot \text{Res}(F)$ and a monotone formula L of size $m^{O(\text{ResDepth}(F))}$ such that $K(x) = L(x) = b$ for every $x \in \text{supp}(\mathcal{D}_b)$ and $b \in \{0, 1\}$.*
2. *There exists an absolute constant c such that for any $\ell \leq \text{ResWidth}(F)$, any monotone circuit K that satisfies $\mathbb{P}_{b \leftarrow \{0,1\}, x \leftarrow \mathcal{D}_b}[K(x) = b] \geq 1/2 + \gamma$, for $\gamma \geq c \cdot \ell \cdot \log(mN)/m$, has size at least*

$$\left(\frac{m}{c \cdot \ell \cdot \log(mN)} \right)^{\ell-1}.$$

Moreover, the distribution $\mathcal{D}_0$ (resp. $\mathcal{D}_1$) is uniformly distributed over $f_{F,m}^{-1}(0)$ (resp. $f_{F,m}^{-1}(1)$).

Proof sketch. The algorithm A constructs a circuit $\mathcal{D}^{F,m} : \{0, 1\}^r \rightarrow \{0, 1\}^{|F| \cdot m^k}$ where $r = mN + 1$ that behaves in the following way. If the first bit of r is 0, then let $y \in \{0, 1\}^{mN}$ be the remaining mN bits. In the output bit (C, α) , the circuit outputs 1 iff y satisfies the clause (C, α) . Moreover, if the first bit of r is 1, then let $x \in [m]^N$ be given the first $m \log n$ of the remaining bits. In the output bit (C, α) , the circuit outputs 1 if $x^{\text{vars}(C)} = \alpha$. The construction of D , and D itself, can be clearly be done in time $\text{poly}(|F| \cdot m^k)$.

Property (1) of the resulting distribution is folklore (see, e.g., [LMM⁺22]). Property (2) is proved in [dRV25, Theorem 11] with $\gamma = 1/2$; the result with smaller γ can be easily obtained by tweaking a few parameters. We explain how this is done in [Appendix A](#). $\square$

3 Shallow Refutation Formula

Let F be a 3-CNF formula over n variables. The $\text{Ref}_s(F)$ formula encodes the statement “ F has a resolution proof of length s ”. Slightly more formally, the variables of $\text{Ref}_s(F)$ are

divided into s blocks, each of which encodes a potential clause in the purported refutation of F and how this clause was derived, and the clauses of $\text{Ref}_s(F)$ enforce that all derivation steps are sound.

The known resolution refutation of $\text{Ref}_s(F)$ when F is satisfiable, given by Pudlák’s upper bound [Pud03], has large width $\Omega(n)$. Under ETH, a large width is unavoidable, since it is possible to find width w refutations of an unsatisfiable N -variate CNF formula in time $N^{O(w)}$. While it is possible to show a gap between the Resolution width of the formula $\text{Ref}_s(F)$ in the satisfiable and unsatisfiable cases of F [AM20, dRGN⁺21] ($O(n)$ vs. $\Omega(s/n)$), for our applications we need a better upper bound. In this section, we define a formula related to $\text{Ref}_s(F)$ which substantially decrease the width upper bound at the cost of an increase in the size of the formula. Moreover, we show that our new formula has small Resolution *depth* when F is satisfiable.

3.1 The New Formula $\text{Ref}_d^G(F)$

Our new formula is denoted by $\text{Ref}_d^G(F)$, and it depends on a bipartite expander graph G and an integer parameter d . We somewhat follow the notation established in “The Proof Analysis Problem” by Arteché et al. [AAAdRK25], but modify the structure of the purported resolution refutation and introduce some extension variables, similar to what was done in [dR21] and [CdR25].

As was done in [dR21] we structure the blocks in d layers (instead of n), and as in [CdR25], we add “summary” variables that give information about segments of n/d variables. These changes allow us to turn the $O(1)$ block-width upper bound and corresponding $O(n)$ resolution width upper bound in [Pud03] into an $O(d)$ upper bound on depth, at the cost of increasing the number of variables and clauses in $\text{Ref}_d^G(F)$ to $2^{\Theta(n/d)}$.

Since it is crucial that these summary variables encode that literals are absent, we also switch focus and define “un-literal” variables encoding that single literals are *not present* in a block as opposed to previous encoding of Ref which has variables encoding the *inclusions* of literals. This is done only for exposition purposes since our unlit variables are just the negation of the usual literal variables. We also include binary variable for pointers (as done in [dRGN⁺21]) to make it possible to query pointers more efficiently.

Finally, we also include “split” variables to turn our formula into a 3-CNF formula using a standard transformation. This is done in order to minimise the blow-up in size of the function we obtain after lifting $\text{Ref}_d^G(F)$.

We start by explaining the structure of the formula $\text{Ref}_d^G(F)$. We then introduce the main variables of $\text{Ref}_d^G(F)$ and the clauses enforcing that the purported proof is sound and is structured as we would like. Finally, we introduce the split variables and explain how they are used to transform the formula into a 3-CNF formula.

Structure of $\text{Ref}_d^G(F)$ and notation Let F be a CNF formula over n variables and $\text{poly}(n)$ clauses, let d be an integer, and let G be a bipartite graph with $t2^{n/d}$ nodes on left side, t nodes on the right, and outdegree δ from each node on the left. We define $s := t \cdot d$. For the sake of convenience in our presentation, we assume that d divides n , and that $|F|$ (the number of clauses of F) and δ , the left outdegree of the expander G , are powers of 2, so they can be cleanly indexed into by binary pointers. We later explain why we can make these assumptions without loss of generality (see [Theorem 3.2](#)).

Given an integer S which is a power of 2 and an integer $X \in [S]$, we denote by $\text{bin}_S(X)$ the binary string of length $\log S$ encoding X . When S is clear from the context, we simply write $\text{bin}(X)$. For instance, for $A \in [|F|]$, we write $\text{bin}(A) \in \{0,1\}^{\log |F|}$ to be the binary string of length $\log |F|$ encoding A .

The formula $\text{Ref}_d^G(F)$ is defined over $s+1$ blocks arranged in $d+1$ layers (see [Figure 1](#)). The block $s+1$ is the root and it is alone in layer 0. The other s blocks are arranged in decreasing order in layers 1 to d with s/d blocks per layer. Formally, for $B \in [s]$, we have $\text{layer}(B) = 1 + d - \lceil B/t \rceil$ while $\text{col}(B) = B - (\lceil B/t \rceil - 1)t$. So, for example, for block s we get that $\text{layer}(s) = 1, \text{col}(s) = t$, while for block 1 we get $\text{layer}(1) = d, \text{col}(1) = 1$.

The idea is that each block contains a clause: block $s+1$ contains the empty clause, and blocks at layer i contain clauses over the first $i \cdot n/d$ variables. For this to locally look like a sound refutation, each block is supposed to have $2^{n/d}$ children on the layer below, one for each possible extension of the clause in the block to the next n/d variables: for each $z \in \{0,1\}^{n/d}$, the clause in the z^{th} child of a block B should not be satisfied by assigning z these n/d variables. The z^{th} child of a block is encoded by some variables so that, even though it is not possible to assign $2^{n/d}$ distinct children to each block, locally it can still look like the formula $\text{Ref}_d^G(F)$ is consistent.

Finally, we would like there to be only a few possible options for each of the z^{th} children of a block. We use the graph G to determine these options: each node on the left-side of G corresponds to a z^{th} child of a block at layer $i \in [d]$ and each node on the right-side of G corresponds to a block on layer $i+1$. The neighbour relation corresponds to the possible blocks at layer $i+1$ that could be the z^{th} child of a block at layer i . To make our encoding more succinct, we need a way of managing binary indices for the at most δ blocks B' that could be a z^{th} child of B . To do this, we need a surjection (which we denote $\text{sur}_{z, \text{col}(B)}$) from the set of at most δ valid blocks that can be the z^{th} child to $[\delta]$. It is easy to see that, if $\text{sur}_{z, \text{col}(B)}$ is the lexicographical order over the valid blocks, then $\text{sur}_{z, \text{col}(B)}$ can be efficiently computed when needed while constructing $\text{Ref}_d^G(F)$ in all the parameter regimes considered in this paper. Note that there may be parameter regimes in which this is not the case.

The variables of $\text{Ref}_d^G(F)$ Let $x_1, \dots, x_n$ be the variables of F . For any variable x_i , the variable $(x_i)^0$ corresponds to $\neg x_i$ and $(x_i)^1$ corresponds to x_i . We define the set $\text{Unlit}_n = \{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$. Let $\text{Leaves} := [s/d]$ be the set of blocks at layer d . Some variables (like $\text{weak}, \text{weakbin}$, which encode what axiom in F the block is a weakening of, in unary and in




Figure 1: The layout of blocks within the purported refutation, demonstrating the columns and layers.

binary, respectively) are only defined for blocks in **Leaves** and some (**point**, **pointbin**, which encode the blocks that were used to derive the current block, also in unary and in binary, respectively) are only defined for blocks that are not in **Leaves**.

Recall that, in our setting, a block is obtained by resolving n/d variables at once. Therefore, the z^{th} child, for $z \in \{0, 1\}^{n/d}$, of a block at layer i should not contain literals x_{id+j} for each $z_j = 1$, nor literals $\neg x_{id+j}$ for each $z_j = 0$. This will be enforced by using the variables **summ-unlit**. In fact, except for blocks that are leaves, the only information about what literals are contained in each block is given by these **summ-unlit** variables. Since the layer of a block determines what variables can be present in the block, only the relevant **summ-unlit** variables are defined for each block. In particular, the root does not contain any **summ-unlit** variable and the leaves contain all **summ-unlit** variables. Moreover, the leaves are the only blocks that contain the **unlit** variables, as only then are these needed to connect the clause in the block with the axiom of F that it is supposedly a weakening of.

We are now ready to introduce the variables of our formula.

- unlit_ℓ^B : literal $\ell \in \text{Unlit}_n$ is not present in block $B \in \text{Leaves}$;
- weak_A^B : block $B \in \text{Leaves}$ was obtained by weakening axiom $A \in [|F|]$;
- weakbin_j^B : for $B \in \text{Leaves}$ and $j \in [\log |F|]$, the value of the j^{th} bit of $\text{bin}(A) \in \{0, 1\}^{\log |F|}$, where $A \in [|F|]$ is the axiom that block B was obtained from by weakening;
- $\text{summ-unlit}_{i,z}^B$: for block $B \in [s]$, $i \in [\text{layer}(B)]$, and $z \in \{0, 1\}^{n/d}$, for each $z_j = 1$, literal $x_{(i-1)d+j}$ is not present in block B , and for each $z_j = 0$, literal $\neg x_{(i-1)d+j}$ is not present in block B ;
- $\text{point}_{B',z}^B$: for blocks $B, B' \in [s+1]$ and $z \in \{0, 1\}^{n/d}$, the z^{th} child of block B is B' (and therefore $\text{summ-unlit}_{i,z}^{B'} = 1$ for $i = \text{layer}(B') = \text{layer}(B) + 1$);
- $\text{pointbin}_{z,j}^B$: for block $B \in [s+1]$, $j \in [\log \delta]$, and $z \in \{0, 1\}^{n/d}$, the value of the j^{th} bit of $\text{bin}(\text{sur}_{z, \text{col}(B)}(B')) \in \{0, 1\}^{\log \delta}$ where B' is the z^{th} child of B (and, in particular, $\text{layer}(B') = \text{layer}(B) + 1$);
- enable^B : block $B \in [s+1]$ is enabled.

We record two remarks regarding the variables of the formula. First, we note a bound on the number of variables.

Remark 3.1. *For each class of variables listed above there are fewer than $O(|F| \cdot n \cdot s^2 \cdot 2^{n/d})$ variables of the class. Assuming $|F| = n^{O(1)}$, $d \leq n$, and $s = 2^{O(n/d)}$, this is $2^{O(n/d)} \cdot n^{O(1)}$.*

Second, we observe that the divisibility assumptions mentioned earlier can easily be dropped.

Remark 3.2. *We now explain why we can make the divisibility assumptions we made about $|F|, d, n$ and δ .*

- *When d does not divide n , we can add one final layer with a corresponding set of summary and pointer variables to manage those x_i with $i > \lfloor n/d \rfloor d$.*
- *When $|F|$ and δ are not powers of two, we can add clauses which rule out each invalid binary pointer or interpret each invalid pointer as some canonical value.*

The clauses of $\text{Ref}_d^G(F)$ We first introduce the unsplit versions (i.e., of width possibly larger than 3) of the clauses which make up $\text{Ref}_d^G(F)$. These clauses enforce the basic structure of the Shallow-Resolution formula ensuring particularly that our summary and pointer variables interact correctly. The clauses are defined for arbitrary G but in our proofs we focus on two cases for G : either G is a bipartite expander with t nodes on the left side, $t \cdot 2^{n/d}$ on the right side, and outdegree δ from each node the left side; or G is a complete graph such that all $(\text{col}(B), (\text{col}(B'), z)) \in E(G)$.

We divide the clauses into three categories depending on whether the block in the superscript is in layer 0 (meaning the block is the root), in layer d (meaning the block is a leaf), or between layers 1 and $d - 1$ (meaning the block is an internal block).

Root case — final block is enabled

$$\text{enable}^{s+1} \quad (\text{REF-1})$$

Leaf case — managing axioms The following clauses are defined for each $B \in \text{Leaf}$.

$$(\text{enable}^B \wedge \text{weak}_A^B) \rightarrow \neg \text{unlit}_\ell^B \quad \text{for } A \in [|F|], \ell, \text{ in clause } A; \quad (\text{REF-2})$$

$$\left(\bigwedge_{j \in [\log |F|]} (\text{weakbin}_j^B)^{\text{bin}(A)_j} \right) \rightarrow \text{weak}_A^B \quad \text{for } A \in [|F|]; \quad (\text{REF-3})$$

$$(\text{enable}^B \wedge \text{summ-unlit}_{i,z}^B) \rightarrow \text{unlit}_{(x_k)^{zj}}^B \quad \text{for } i \in [d], j \in [n/d], \quad (\text{REF-4})$$

$$k = (i - 1)n/d + j.$$

Internal Step — pointer and summary enforcement The following clauses are defined for each B that is not a leaf, i.e., for B such that $\text{layer}(B) < d$; for each $z \in \{0, 1\}^{n/d}$; and for each B' which is a possible z th child for B , i.e., for $B' \in P_B := \{B' : \text{layer}(B') = \text{layer}(B) + 1 \text{ and } ((\text{col}(B), z), \text{col}(B')) \in E(G)\}$.

$$\left(\bigwedge_{j \in [\log \delta]} (\text{pointbin}_{z,j}^B)^{\text{bin}(\text{sur}(B'))_j} \right) \rightarrow \text{point}_{B',z}^B \quad (\text{REF-5})$$

$$(\text{enable}^B \wedge \text{point}_{B',z}^B) \rightarrow \text{enable}^{B'} \quad (\text{REF-6})$$

$$(\text{enable}^B \wedge \text{point}_{B',z'}^B \wedge \text{summ-unlit}_{i,z}^B) \rightarrow \text{summ-unlit}_{i,z}^{B'} \quad \text{for } z' \in \{0, 1\}^{n/d}, \quad (\text{REF-7})$$

$$i \leq \text{layer}(B);$$

$$(\text{enable}^B \wedge \text{point}_{B',z}^B) \rightarrow \text{summ-unlit}_{i+1,z}^{B'} \quad \text{for } i = \text{layer}(B). \quad (\text{REF-8})$$

This completes the list of the clauses of the non 3-CNF version of the formula $\text{Ref}_d^G(F)$. We first observe that there are not too many clauses in this formula and that they are not too wide.

Remark 3.3. *For each class of clauses listed above, there are fewer than $O(|F| \cdot n \cdot s^2 \cdot 2^{2n/d})$ clauses of that class, and each clause has fewer than $O(\log(|F|) + \log(\delta))$ literals in them.*

We now introduce a final set of variables which we refer to as “split variables”. These are used to transform the formula into 3-CNF using a standard transformation. For every

clause $C_i = v_1^i \vee v_2^i \vee \dots \vee v_{|C_i|}^i$ so far introduced, we introduce a collection of split variables $\text{split}_{C_i}^j$ for each $j \in [|C_i| - 3]$ and replace C_i with a collection of clauses $(v_1^i \vee v_2^i \vee \text{split}_{C_i}^1) \wedge (\neg \text{split}_{C_i}^1 \vee v_3^i \vee \text{split}_{C_i}^2) \wedge \dots \wedge (\neg \text{split}_{C_i}^{|C_i|-3} \vee v_{|C_i|-1}^i \vee v_{|C_i|}^i)$.

We denote the formula after splitting by $\text{Ref}_d^G(F)$. Note that, by [Theorem 3.3](#), if $s = 2^{O(n/d)}$, $d \leq n$, $|F| = n^{O(1)}$, and $\delta = O(s)$ then before splitting, the formula has $2^{O(n/d)} \cdot n^{O(1)}$ clauses containing at most $O(\log n + \log s)$ variables each. As a consequence, after splitting, $\text{Ref}_d^G(F)$ has at most $O(\log n + \log s) \cdot 2^{O(n/d)} \cdot n^{O(1)} = 2^{O(n/d)} \cdot n^{O(1)}$ clauses, each of which contains at most 3 variables. We record this observation below.

Remark 3.4. *The formula $\text{Ref}_d^G(F)$ has at most $2^{O(n/d)} \cdot n^{O(1)}$ clauses, each of which contains at most 3 variables.*

3.2 Main Theorem

We prove our main theorem for an arbitrary expander graph, and we then instantiate this theorem with two different graphs.

Theorem 3.5. *Given a 3-CNF formula F on n variables, with $\text{poly}(n)$ clauses, and positive integers $d \leq n$ and $t \leq 2^{n/d}$, and a $(t \cdot 2^{n/d}, 2^{n/d}, \delta, r, c)$ -bipartite expander graph G there is a deterministic algorithm running in time $2^{O(n/d)} \cdot n^{O(1)}$ that constructs a 3-CNF formula $\text{Ref}_d^G(F)$ which satisfies the following properties:*

1. $\text{Ref}_d^G(F)$ has $2^{\Theta(n/d)} \cdot n^{\Theta(1)}$ variables and $2^{\Theta(n/d)} \cdot n^{\Theta(1)}$ clauses.
2. $\text{Ref}_d^G(F)$ has width $O(1)$;
3. $\text{Ref}_d^G(F)$ admits Resolution proofs of depth at most $O(d \log \delta + \log n)$ when F is satisfiable;
4. $\text{Ref}_d^G(F)$ requires Resolution proofs of width at least $2^{\Omega(r(c-1)/n)}$ when F is unsatisfiable.

Proof. The first two claims follow directly from [Theorem 3.4](#). The fourth claim follows from essentially the same argument given in [\[CdR25\]](#) (reducing a hard for resolution pigeonhole principle to finding a collision on the expander G), which we defer to [Appendix B](#).

The third claim follows from the guaranteed success of the following $O(d \log \delta + \log n)$ depth prover in the standard Prover-Delayer game (see [Section 2.3](#)) used to analyse bounds for resolution proofs. We lay out the strategy the prover takes in each of the three cases: root where $B = s + 1$, internal where $B \leq s$ and $\text{layer}(B) < d$, leaf where $\text{layer}(B) = d$, and backchecking which is done after finding a leaf which is a valid weakening. For each of these, we specify the types of clauses the prover could identify as being violated if the delayer gave answers that do not allow the prover to proceed in each step. In all cases it should be clear that

the prover can identify a violated clause of one of the listed classes **REF-1–REF-8** immediately. Moreover, since all non-split clauses have at most $O(\log \delta + \log |F|) = O(\log \delta + \log n)$ variables, finding a split clause which is violated can be done in additional depth $O(\log \delta + \log n)$ by querying all the variables in the clause and all the split variables associated with the clause.

The prover critically knows an assignment x^* which satisfies F and will use this knowledge in implementing their strategy. In discussing the strategy below, we apply subscripts to assignment strings x_i^* and summary strings z_i^{j+1} with these subscripts serving as indices specifying the i^{th} bit in the string.

Invariance condition: After applying the root case and k internal steps of the strategy (but not having yet applied the leaf case or the backchecking step), the following invariance conditions are maintained:

- for all $j \in 0 \cup [k]$ the prover knows the variables enable^{B_j} and $\text{point}_{B_{j+1}, z^{j+1}}^{B_j}$ which are all equal to 1;
- $B_0 = s + 1$;
- $\forall j \in [d], \forall i \in [n/d] : z_i^{j+1} = x_{jn/d+i}^*$.

Root Case: Learn that $s + 1$ is enabled by querying enable^{s+1} which will equal 1, otherwise (**REF-1**) is violated. Learn the index B_1 of the child which is unsatisfied by x^* by querying the $\log \delta$ variables $\text{pointbin}_{z_1^1, j}^{s+1}$ where $z_i^1 = x_i^*$. Finally, query the variable $\text{point}_{B_1, z}^{s+1}$ which must be 1 otherwise (**REF-5**) is violated.

Internal Step: Learn that B_j is enabled by querying enable^{B_j} which will equal 1 (**REF-6**). Learn the index B_{j+1} of the child of B_j which is unsatisfied by x^* by querying the $\log \delta$ variables $\text{pointbin}_{z_{j+1}, j'}^{B_j}$ where $z_i^{j+1} = x_{jn/d+i}^*$. Finally, query the variable $\text{point}_{B_{j+1}, z}^{B_j}$ which must be 1 otherwise right side of (**REF-5**) is violated.

Leaf Case: Learn that B_j is enabled by querying enable^{B_j} which will equal 1, otherwise (**REF-6**) is violated. Query all $\text{summ-unlit}_{i, z^i}^{B_j}$ where each z^i is specified according to the system in the internal step. If they are all 1 then, learn the axiom $A \in [|F|]$ which B_j is a weakening of by querying the $\log |F|$ variables $\text{weakbin}_{j'}^{B_j}$. Then query the variable $\text{weak}_A^{B_j}$ which must be 1 otherwise it would violate (**REF-3**). Since x^* is a satisfying assignment, one of (**REF-2**) or (**REF-4**) will be violated. We are left with the case where at least one of the $\text{summ-unlit}_{i, z^i}^{B_j}$ variables is 0.

Backchecking: For the indices i and j for which $\text{summ-unlit}_{i,z^i}^{B_j}$ was found to be 0, query all $\text{summ-unlit}_{i,z^i}^{B_\kappa}$ for $\kappa \in [i+1, j-1]$. If $\text{summ-unlit}_{i,z^i}^{B_{i+1}} = 0$ then the appropriate (REF-8) is violated. Otherwise the prover finds a κ such that $\text{summ-unlit}_{i,z^i}^{B_\kappa} = 1$ but $\text{summ-unlit}_{i,z^i}^{B_{\kappa+1}} = 0$, which implies the appropriate (REF-7) is violated.

In the root case, and each of the d internal steps, the prover queries $O(\log \delta)$ variables; in the leaf case, the prover queries $O(\log |F|)$ variables; and in the backchecking case, the prover queries at most d variables. In total, it makes at most $O(d \log \delta + \log |F|)$ queries and is guaranteed to find a violated clause. $\square$

From Theorem 3.5 we obtain two corollaries. First, using the fact that an unbalanced complete bipartite graph is an expander (Theorem 2.1), we obtain the following.

Corollary 3.6. *Given a 3-CNF formula F on n variables, with $\text{poly}(n)$ clauses, and positive integers $d \leq n$ and $t \leq 2^{n/d}$, there is a deterministic algorithm running in time $2^{O(n/d)} \cdot n^{O(1)}$ that constructs a 3-CNF formula $\text{Ref}_d^G(F)$ which satisfies the following properties”:*

1. $\text{Ref}_d^G(F)$ has $2^{O(n/d)} \cdot n^{O(1)}$ variables and $2^{O(n/d)} \cdot n^{O(1)}$ clauses;
2. $\text{Ref}_d^G(F)$ has width $O(1)$;
3. $\text{Ref}_d^G(F)$ admits Resolution proofs of depth at most $O(d \log t + \log n)$ when F is satisfiable;
4. $\text{Ref}_d^G(F)$ requires Resolution proofs of width at least $2^{\Omega(t/n)}$ when F is unsatisfiable.

Using a randomised construction of expanders (Theorem 2.2), we can improve the above construction in the following way.

Corollary 3.7. *Given a 3-CNF formula F on n variables, with $\text{poly}(n)$ clauses, and a positive integer $d \leq n$, there is a randomized algorithm running in time $2^{O(n/d)} \cdot n^{O(1)}$ that constructs a 3-CNF formula $\text{Ref}_d^G(F)$ which, with high probability, satisfies the following properties:*

1. $\text{Ref}_d^G(F)$ has $2^{O(n/d)} \cdot n^{O(1)}$ variables and $2^{O(n/d)} \cdot n^{O(1)}$ clauses;
2. $\text{Ref}_d^G(F)$ has width $O(1)$;
3. $\text{Ref}_d^G(F)$ admits Resolution proofs of depth at most $O(d + \log n)$ when F is satisfiable;
4. $\text{Ref}_d^G(F)$ requires Resolution proofs of width at least $2^{2^{\Omega(n/d)}}$ when F is unsatisfiable.

4 Learning Lower Bounds Under ETH

In this section, we formally prove Theorems 1.1 and 1.2. We begin by defining a gap learning decision problem to which we reduce the SAT problem, and from which we are able to reduce to other learning problems.

Definition 4.1. Let $s_1, s_2, \varepsilon : \mathbb{N} \rightarrow \mathbb{N}$ be time-constructible functions. The computational problem $\text{mCFG}_{s_1}^{s_2}[\varepsilon]$ (referred to as Monotone Circuit-Formula Gap Learning) receives as input $(1^n, \mathcal{D})$, where $\mathcal{D}$ is a circuit with $n + 1$ output bits and size n^2 (here, we identify the circuit with its representation for simplicity of notation). We interpret $\mathcal{D}$ as a distribution of pairs $(x, b) \in \{0, 1\}^n \times \{0, 1\}$, obtained by sampling a uniformly random input r and outputting $\mathcal{D}(r)$. The task is to decide if there exists a monotone formula C (with n input bits and single output bit) of size s_1 such that $\mathbb{P}_{(x,b) \leftarrow \mathcal{D}}[C(x) = b] = 1$, or if every monotone circuit C' such that $\mathbb{P}_{(x,b) \leftarrow \mathcal{D}}[C'(x) = b] \geq 1/2 + \varepsilon$ must have size at least s_2 . The number n is called the dimension of the input.

We record an observation about the sampling circuit that we use later on.

Remark 4.2. We observe that the size of the sampling circuit $\mathcal{D}$ is chosen to be n^2 arbitrarily but without loss of generality. Indeed, given any family of inputs where the n -bit samplers have size $s = n^{O(1)}$, if $s < n^2$ we can just add dummy gates to the circuit; if $s > n^2$ we can pad the output x with $\sqrt{s} - n$ 0's; in both cases, we reduce to the case $s = n^2$ in polynomial-time.

We now implement the strategy outlined in the introduction. We take a 3-CNF formula F on n variables, construct the formula $\text{Ref}_d^G(F)$ satisfying either Theorem 3.6 (in case of a deterministic reduction) or Theorem 3.7 (in case of a randomised reduction). Using a deterministic reduction, taking m such that $n \leq m \ll 2^{n/d}$, $t \leq 2^{n/d}$ and $\ell = \varepsilon dm/n \ll 2^{t/n}$ for small enough ε in Theorem 2.4 gives that the function $f_{\text{Ref}_s^G(F), m}$ has $N = 2^{O(n/d)} \cdot n^{O(1)}$ input bits and the following monotone complexity:

1. At most $m^{O(d \log t + \log n)} = 2^{O((d \log t + \log n) \log m)}$ when F is satisfiable, and
2. At least $2^{\Omega(\ell)} \geq 2^{\Omega(m/\log N)}$ when F is unsatisfiable.

Using a randomised reduction, the value of t is irrelevant and the resulting monotone complexity is

1. At most $m^{O(d + \log n)} = 2^{O((d + \log n) \log m)}$ when F is satisfiable, and
2. At least $2^{\Omega(\ell)} \geq 2^{\Omega(m/\log N)}$ when F is unsatisfiable.

Since we are reducing from ETH, the value $m = \text{poly}(n)$ is already good enough, as a monotone circuit lower bound beyond $2^{\Omega(n)}$ (for an N -variate function) is not of much use for reductions to Partial-MCSP or PAC-learning (since we would need to output a circuit of this size or generate an exponential number of examples). The value d is then the main parameter we choose. There are two parameter regimes we will care about:

1. ($d = \log n$). We obtain a formula upper bound which is smaller than the number N of input bits (a “junta”). The runtime lower bound based on ETH becomes $N^{\Omega(\log \log N)}$.
2. ($d \approx \sqrt{n}$). This gives us a polynomial vs. quasipolynomial gap. The runtime lower bound becomes $N^{\widetilde{\Omega}(\log N)}$.

4.1 ETH-hardness of Monotone Circuit-Formula Gap Learning

We first show the result for juntas.

Theorem 4.3 (Junta Gap-Learning). *Let $\gamma > 0$ be a constant. For large enough $n \in \mathbb{N}$, the 3SAT problem on n variables reduces to $\text{mCFGL}_{s_1}^{s_2}[\gamma]$ in time $2^{O(n/\log n)}$, where $s_1(N) = 2^{O(\log \log N)^3}$ and $s_2(N) = N^{\Omega(\log N)}$. Consequently, the problem $\text{mCFGL}_{s_1}^{s_2}[\gamma]$ needs time $N^{\Omega(\log \log N)}$ to be solved under ETH.*

Proof. Let F be a given 3-CNF over n variables. Fix $d = \log n$. Note that $\text{Ref}_s^G(F)$ is a formula of width $O(1)$ over $N = 2^{\Theta(n/\log n)}$ variables and with $N^{O(1)}$ clauses. Let

$$m := n^3/(\log n)^3, \quad t := n^2.$$

and let $\ell = \varepsilon dm/n$ for a sufficiently small constant $\varepsilon > 0$. The function $f^* := f_{\text{Ref}_d^G(F), m}$ is defined on $n_0 := N^{O(1)}m^{O(1)} = 2^{O(n/\log n)}$ variables (Theorem 2.3). Theorems 2.4 and 3.5 imply that f^* can be computed with monotone formulas of size $s := m^{O(d \log t + \log n)} = 2^{O(\log^3 n)} = 2^{O(\log \log N)^3}$ when F is satisfiable. Moreover, when F is unsatisfiable, the function requires monotone circuits of size at least

$$s' := 2^{\Omega(\ell)} = N^{\Omega(\log N)}$$

to be γ -approximated over the distribution $\mathcal{D} := \mathcal{D}^{\text{Ref}_s^G(F), m}$ given by Theorem 2.4, since we can take ε sufficiently smaller than γ . This distribution can be constructed in time $2^{O(n/\log n)}$ given F and m . A postprocessing taking $2^{O(n/\log n)}$ -time may be needed to adjust the description of the circuit to the required length (see Theorem 4.2). The ETH lower bound follows by noticing that $N^{\log \log N} = 2^{O(n)}$. $\square$

The result for polynomial vs. quasipolynomial gap follows similarly.

Theorem 4.4 (Polynomial vs. quasipolynomial gap). *Let $\gamma > 0$ be a constant. For large enough $n \in \mathbb{N}$, the 3SAT problem on n variables reduces to $\text{mCFGL}_{s_1}^{s_2}[\gamma]$ in time $2^{O(\sqrt{n} \log n)}$, where $s_1(N) = N$ and $s_2(N) = N^{\Omega(\log N)}$. Consequently, the problem $\text{mCFGL}_{s_1}^{s_2}[\gamma]$ needs time $N^{\Omega((\log N)/(\log \log N)^2)}$ to be solved under ETH.*

Proof. Let F be a given 3-CNF over n variables. Fix $d = \alpha\sqrt{n}/\log n$ where $\alpha > 0$ is small enough constant. Note that $\text{Ref}_s^G(F)$ is a formula of width $O(1)$ over $N = 2^{\Theta(\sqrt{n} \log n)}$ variables

and with $N^{O(1)}$ clauses. Let

$$m := (\log N)^2 \asymp n \log^2 n, \quad t := n^2.$$

and let $\ell = \varepsilon dm/n$ for a sufficiently small constant $\varepsilon > 0$. The function $f^* := f_{\text{Ref}_d^G(F), m}$ is defined on $n_0 := N^{\Theta(1)} m^{\Theta(1)} = 2^{\Theta(\sqrt{n} \log n)} = N^{\Theta(1)}$ variables ([Theorem 2.3](#)). Theorems [2.4](#) and [3.5](#) imply that f^* can be computed with monotone formulas of size

$$s := m^{O(d \log t + \log n)} = 2^{O(d \log^2 n + \log^2 n)} = 2^{O(\alpha \sqrt{n} \log n)} \leq n_0,$$

when F is satisfiable and α is small enough. Moreover, when F is unsatisfiable, the function requires monotone circuits of size at least

$$s' := 2^{\Omega(\ell)} = N^{\Omega(\log N)}$$

to be γ -approximated over the distribution $\mathcal{D} := \mathcal{D}_{\text{Ref}_s^G(F), m}$ given by [Theorem 2.4](#), since we can take ε sufficiently smaller than γ . This distribution can be constructed in time $2^{O(\sqrt{n} \log n)}$ given F and m . The ETH lower bound follows by noting that $N^{(\log N)/(\log \log N)^2} = 2^{O(n)}$. $\square$

4.2 rETH-hardness of Monotone Circuit-Formula Gap Learning

Next, we show slightly stronger versions of these results but under rETH.

Theorem 4.5 (Junta Gap-Learning under rETH). *Let $\gamma > 0$ be a constant. For large enough $n \in \mathbb{N}$, the 3SAT problem on n variables randomly reduces to $\text{mCFGL}_{s_1}^{s_2}[\gamma]$ in time $2^{O(n/\log n)}$, where $s_1(N) = (\log N)^{O(\log \log N)}$ and $s_2(N) = N^{\Omega(\log N)}$. Consequently, the problem $\text{mCFGL}_{s_1}^{s_2}[\gamma]$ needs time $N^{\Omega(\log \log N)}$ to be solved under rETH.*

Proof. Let F be a given 3-CNF over n variables. Fix $d = \log n$. Note that $\text{Ref}_s^G(F)$ is a formula of width $O(1)$ over $N = 2^{\Theta(n/\log n)}$ variables and with $N^{O(1)}$ clauses. Let

$$m := n^3/(\log n)^3,$$

and let $\ell = \varepsilon dm/n$ for a sufficiently small constant $\varepsilon > 0$. The function $f^* := f_{\text{Ref}_d^G(F), m}$ is defined on $n_0 := N^{O(1)} m^{O(1)} = 2^{O(n/\log n)}$ variables ([Theorem 2.3](#)). Theorems [2.4](#) and [3.5](#) imply that f^* can be computed with monotone formulas of size $s := m^{O(d + \log n)} = 2^{O(\log^2 n)} = 2^{O(\log \log N)^2}$ when F is satisfiable. The rest of the proof is the same as in [Theorem 4.3](#). $\square$

Theorem 4.6 (Polynomial vs. quasipolynomial gap under rETH). *Let $\gamma > 0$ be a constant. For large enough $n \in \mathbb{N}$, the 3SAT problem on n variables reduces to $\text{mCFGL}_{s_1}^{s_2}[\gamma]$ in time $2^{O(\sqrt{n})}$, where $s_1(N) = N$ and $s_2(N) = N^{\Omega(\log N)}$. Consequently, the problem $\text{mCFGL}_{s_1}^{s_2}[\gamma]$ needs time $N^{\Omega(\log N)}$ to be solved under rETH.*

Proof. Let F be a given 3-CNF over n variables. Fix $d = \alpha\sqrt{n}$ where $\alpha > 0$ is small enough constant. Note that $\text{Ref}_s^G(F)$ is a formula of width $O(1)$ over $N = 2^{\Theta(\sqrt{n})}$ variables and with $N^{O(1)}$ clauses. Let

$$m := (\log N)^2 \asymp n,$$

and let $\ell = \varepsilon dm/n$ for a sufficiently small constant $\varepsilon > 0$. The function $f^* := f_{\text{Ref}_d^G(F), m}$ is defined on $n_0 := N^{\Theta(1)} m^{\Theta(1)} = 2^{\Theta(\sqrt{n} \log n)}$ variables ([Theorem 2.3](#)). Theorems [2.4](#) and [3.5](#) imply that f^* can be computed with monotone formulas of size

$$s := m^{O(d+\log n)} = 2^{O(d)} = 2^{O(\alpha\sqrt{n})} \leq n_0,$$

when F is satisfiable and α is small enough. Moreover, when F is unsatisfiable, the function requires monotone circuits of size at least

$$s' := 2^{\Omega(\ell)} = N^{\Omega(\log N)}$$

to be γ -approximated over the distribution $\mathcal{D} := \mathcal{D}^{\text{Ref}_s^G(F), m}$ given by [Theorem 2.4](#), since we can take ε sufficiently smaller than γ . This distribution can be constructed in time $2^{O(\sqrt{n})}$ given F and m . The ETH lower bound follows by noticing that $N^{\log N} = 2^{O(n)}$. $\square$

4.3 From Gap Learning to Partial-Gap-MCFSP

We now define the partial MCSP problem we consider. Instead of the entire truth table of a partial function, we consider a more succinct encoding where we are only given locations where the function is defined.

Definition 4.7. Let $s_1, s_2, \varepsilon : \mathbb{N} \rightarrow \mathbb{N}$ be time-constructible functions. Let $\text{mMCFSP}_{s_1}^{s_2}[\varepsilon]$ be the promise problem which receives as input 1^n and a list of m pairs $((x_1, b_1), \dots, (x_m, b_m))$ (called examples), where $x_i \in \{0, 1\}^n$ and $b_i \in \{0, 1\}$ for all $i \in [m]$, and must distinguish between the following two cases:

- there is a monotone formula C of size at most s_1 such that, for all $i \in [m]$, $C(x_i) = b_i$;
- for all circuits C' of size less than s_2 , we have $|\{i \in [m] : C'(x_i) \neq b_i\}| \geq m(1/2 - \varepsilon)$.

The number n is called the dimension of the input.

To show hardness of mMCFSP , it suffices to reduce from mCFGL . As described in the introduction, it suffices to sample from the distribution $\mathcal{D}$ given as input to mCFGL . In the small formula case, we clearly generate examples which can be computed by a small formula. In the large circuit case, with enough examples we can conclude by the Chernoff inequality that no circuit of small size can compute most of those examples correctly.

Theorem 4.8. *For every growing functions $s_1, s_2 : \mathbb{N} \rightarrow \mathbb{N}$ such that $s_1 \leq s'_1 < s'_2 \leq s_2$ (for large enough n) and $\gamma : \mathbb{N} \rightarrow [0, 1]$, the problem $\text{mCFGL}_{s_1}^{s_2}[\gamma]$ with input $(1^n, \mathcal{D})$ randomly reduces to $\text{mMCFSP}_{s'_1}^{s'_2}[2\gamma]$ in time $\text{poly}(n) + O(\frac{1}{\gamma^2} \cdot s'_2 \log s'_2)$, producing $O(\frac{1}{\gamma^2} \cdot s'_2 \log s'_2)$ examples of dimension n .*

Proof. Let $\mathcal{D}$ be given as input for $\text{mCFGL}_{s_1}^{s_2}[\gamma]$. Take $m = \frac{4}{\gamma^2} \cdot K s'_2 \log s'_2$ samples from $\mathcal{D}$, and call them $(x_1, b_1), \dots, (x_m, b_m)$.

In the case that there exists a formula of size $s_1 < s'_1$ computing $\mathcal{D}$ then that same small circuit will agree with $(x_1, b_1), \dots, (x_m, b_m)$ for all i . So “small” instances map to “small” instances. It will then suffice to prove the theorem to show that “large” instances map to “large” instances. In particular, it suffices to show that, with all but negligible probability, for every circuit of C' of size less than s'_2 , we have $|\{i \in [m] : C'(x_i) \neq b_i\}| \geq m(1/2 - \varepsilon - \gamma)$.

Let $a_{C'}$ be the number of $i \in [m]$ such that $C'(x_i) = b_i$. By definition of mCFGL , we know $\mathbb{P}_{(x,b) \leftarrow \mathcal{D}}[C'(x) = b] \leq 1/2 + \gamma$, and thus $\mathbb{E}[a_{C'}] \leq m(1/2 + \gamma)$. Therefore, for each circuit C' , by Chernoff-Hoeffding’s inequality we know that

$$\mathbb{P}_{(x_1, b_1), \dots, (x_m, b_m) \leftarrow \mathcal{D}}[a_{C'} - m(1/2 + \gamma) \geq m\gamma] \leq e^{-2m^2\gamma^2/m} = e^{-2m\gamma^2} = (s'_2)^{-4s'_2}, \quad (9)$$

or, equivalently,

$$\mathbb{P}_{(x_1, b_1), \dots, (x_m, b_m) \leftarrow \mathcal{D}}[a_{C'} \geq m(1/2 + 2\gamma)] \leq (s'_2)^{-4s'_2}.$$

There are less than $(s'_2)^{3s'_2}$ monotone circuits with n bit inputs, s'_2 gates, fan-in 2, and arbitrary fan-out. By union bounding over these circuits, we can see that the probability of any circuit satisfying $a_{C'} \geq m(1 + 2\gamma)$ is less than $1/3$ for large enough n . This completes the proof. $\square$

We now employ the hardness results obtained for mCFGL under rETH in [Section 4.2](#) to obtain hardness of $\text{mMCFSP}_{s_1}^{s_2}$ using the reduction of the previous theorem. We consider 4 choice of parameters which are in a way best possible:

1. *Junta vs. quasipolynomial.* Here we take the largest gap possible between formula and circuit size. This gives us a weaker runtime lower bound of $N^{\Omega(\log \log N)^{1-\varepsilon}}$. This is the weakest runtime lower bound we obtain.
2. *Linear vs. quasipolynomial.* Here we obtain a weaker gap but against superpolynomial-size circuits. This gives us a better runtime lower bound of $N^{\Omega(\log N)^{1-\varepsilon}}$.
3. *Junta vs. polynomial.* We then consider the setting where the gap is largest possible subject to $s_2 = \text{poly}(n)$. This gives us a runtime lower bound of $N^{\Omega(\log \log N)}$, slightly better than (1).

4. *Linear vs. polynomial.* We weaken the gap to an arbitrary polynomial n^c . This gives us a better runtime lower bound of $N^{\Omega(\log N)}$, slightly better than (2). This is the strongest runtime lower bound we obtain.

We obtain the following result for each of these parameter regimes.

Corollary 4.9. *Let $\gamma > 0$ be a constant. Under $rETH$, the following holds:*

1. **(Junta vs. quasipolynomial)** *For every constant $\varepsilon > 0$, the problem $\mathbf{mMCFSP}_{s_1}^{s_2}[\gamma]$ requires time $N^{\Omega(\log \log N)^\varepsilon}$ to be decided, where $s_1(n) = (\log N)^{\log \log N}$ and $s_2(n) = n^{(\log \log n)^{1-\varepsilon}}$. The hard instances consist of $m = \Theta(s_2 \log s_2)$ examples, and thus have size $N = nm = n^{\Theta(\log \log n)^{1-\varepsilon}}$.*
2. **(Linear vs. quasipolynomial)** *For every constant $\varepsilon \in (0, 1)$, the problem $\mathbf{mMCFSP}_{s_1}^{s_2}[\gamma]$ requires time $N^{\Omega((\log N)^{2/(2-\varepsilon)})}$ to be decided, where $s_1(n) = n$, and $s_2(n) = n^{(\log n)^{1-\varepsilon}}$. The hard instances consist of $m = \Theta(s_2 \log s_2)$ examples, and thus have size $N = nm = n^{\Theta(\log n)^{1-\varepsilon}}$.*
3. **(Junta vs. polynomial)** *For every constant $\varepsilon > 0$, the problem $\mathbf{mMCFSP}_{s_1}^{s_2}[\gamma]$ requires time $N^{\Omega(\log \log N)}$ to be decided, where $s_1(n) = (\log N)^{\log \log N}$ and $s_2(n) = n^c$. The hard instances consist of $m = \Theta(s_2 \log s_2)$ examples, and thus have size $N = \Theta(nm) = \text{poly}(n)$.*
4. **(Linear vs. polynomial)** *For every constant $c > 0$, the problem $\mathbf{mMCFSP}_{s_1}^{s_2}[\gamma]$ requires time $N^{\Omega(\log N)}$ to be decided, where $s_1(n) = n$, and $s_2(n) = n^c$. The hard instances consist of $m = \Theta(s_2 \log s_2)$ examples, and thus have size $N = \Theta(nm) = \text{poly}(n)$.*

Proof. We first prove **item 1**. By **Theorem 4.5**, for large enough $n \in \mathbb{N}$ the problem $\mathbf{mCFGL}_{s_1}^{s_2}[\gamma/2]$ requires time $n^{\Omega(\log n)}$ to be computed under ETH for some $s_1 = 2^{O(\log \log n)^2}$ and $s_2 = n^{\Omega(\log n)}$. Take $s'_1 = s_1$ and $s'_2 = n^{(\log \log n)^{1-\varepsilon}}$, thus satisfying $s_1 < s'_1 < s'_2 < s_2$.

If $\mathbf{mMCFSP}_{s'_1}^{s'_2}[\gamma]$ can be computed in time $T(N)$ on an instance of $m = K s'_2 \log s'_2 = n^{O(\log \log n)^{1-\varepsilon}}$ examples and dimension n (and thus input size $N = nm$), where K is a large enough constant, we obtain by **Theorem 4.8** that $\mathbf{mCFGL}_{s'_1}^{s'_2}[\gamma/2]$ can be solved in randomised time $\text{poly}(n) + O(m) + T(N)$. This must be at least $n^{\Omega(\log \log n)}$ which implies $T(N) = n^{\Omega(\log \log n)}$. Since $N = n^{\Theta(\log n \log n)^{1-\varepsilon}}$, we obtain $\log n \log n = (1 - o(1)) \log \log N$ and

$$T(N) \geq N^{\Omega(\log \log N)^\varepsilon}.$$

We now prove **item 2**. The proof is similar. Let $s'_1 = n$ and $s'_2 = n^{(\log n)^{1-\varepsilon}}$ in the argument above. We have that $\mathbf{mCFGL}_{s'_1}^{s'_2}[\gamma/2]$ requires time $n^{\Omega(\log n)}$ to be computed by **Theorem 4.6**. Since $N \asymp n s'_2 \log s'_2 = n^{\Theta(\log n)^{1-\varepsilon}}$, we obtain $\log n \asymp (\log N)^{1/(2-\varepsilon)}$ and thus

$$T(N) \geq 2^{\Omega(\log n)^2} \geq 2^{\Omega(\log N)^{2/(2-\varepsilon)}}.$$

We repeat the argument for **item 3**. Take $s'_1 = (\log n)^{\log \log n}$ and $s'_2 = n^c$. We have that $\text{mCFGL}_{s'_1}^{s'_2}[\gamma/2]$ requires time $n^{\Omega(\log \log n)}$ to be computed by **Theorem 4.5**. Since $N \asymp ns'_2 \log s'_2 \asymp n^{c+1} \log n$, we obtain $\log N \asymp \log n$ and $\log \log N \sim \log \log n$, and thus

$$T(N) \geq n^{\Omega(\log \log n)} \geq N^{\Omega(\log \log N)}.$$

Finally, we prove **item 4**. Let $s'_1 = n$ and let $s'_2 = n^c$ in the argument above. We have that $\text{mCFGL}_{s'_1}^{s'_2}[\gamma/2]$ requires time $n^{\Omega(\log n)}$ to be computed by **Theorem 4.6**. Since $N \asymp ns'_2 \log s'_2 \asymp n^{c+1} \log n$, we obtain $\log N \asymp \log n$ and thus

$$T(N) \geq 2^{\Omega(\log n)^2} \geq 2^{\Omega(\log N)^2}. \quad \square$$

4.4 From Gap Learning to Improper PAC-Learning

Finally, we explain how to obtain hardness of PAC-learning. We first define formally the notion of PAC-learning. Given a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and a distribution $\mathcal{D}$ with support in $\{0, 1\}^n$, we denote by $\text{Ex}(f, \mathcal{D})$ the oracle that runs in unit time and outputs a sample (x, b) where $x \leftarrow \mathcal{D}$ and $b = f(x)$. We also define $\text{err}_{f, \mathcal{D}}(h) := \mathbb{P}_{x \in \mathcal{D}}[h(x) \neq f(x)]$.

Definition 4.10. Let $\mathcal{F}_n, \mathcal{H}_n$ be sets of Boolean functions over $\{0, 1\}^n$, and let $\mathcal{F} = \bigcup_{n \geq 1} \mathcal{F}_n$ and $\mathcal{H} = \bigcup_{n \geq 1} \mathcal{H}_n$ be such that $\mathcal{F} \subseteq \mathcal{H}$. Let $T : \mathbb{N} \rightarrow \mathbb{N}$. We say that the concept class $\mathcal{F}$ is PAC-learnable in time T by the hypothesis class $\mathcal{H}$ if there is a randomised algorithm such that, for every $f \in \mathcal{F}$, given $(1^n, 1^{1/\varepsilon}, 1^{1/\delta})$ and oracle access to $\text{Ex}(f, \mathcal{D})$, outputs in time $T(n + 1/\varepsilon + 1/\delta)$, with probability $1 - \delta$, a concept $h \in \mathcal{H}$ such that $\text{err}_{f, \mathcal{D}}(h) \leq \varepsilon$.

For $s : \mathbb{N} \rightarrow \mathbb{N}$, let $\mathcal{C}_{s(n)}$ (resp. $\mathcal{L}_{s(n)}$) be the set of monotone Boolean circuits (resp. formulas) on n bits of size $s(n)$. Let $\mathcal{C}_s = \bigcup_{n \in \mathbb{N}} \mathcal{C}_{s(n)}$ and $\mathcal{L}_s = \bigcup_{n \in \mathbb{N}} \mathcal{L}_{s(n)}$. Hardness of PAC-learning $\mathcal{L}_s$ by $\mathcal{C}_s$ follows from the results of **Section 4.2** by a proof similar to the one employed in the previous section.

Theorem 4.11. Under *rETH*, we have:

1. For every constant $\alpha > 0$, monotone formulas on n bits of size n require time $n^{\Omega(\log n)}$ to be PAC-learned by monotone circuits of size $s(n)$, for any $s(n)$ such that $n \ll s(n) \leq n^{(\log n)^{1-\alpha}}$.
2. For every constant $\alpha > 0$, monotone formulas on n bits of size $(\log n)^{O(\log \log n)}$ require time $n^{\Omega(\log \log n)}$ to be PAC-learned by monotone circuits of size $s(n)$, for any $s(n)$ such that $n \ll s(n) \leq n^{(\log \log n)^{1-\alpha}}$.

Proof. We first prove the first item. The proof of the second item is analogous, so we omit it.

Let $\varepsilon = \delta = 1/10$. Let $(1^n, \mathcal{D})$ be an input to $\text{mCFGL}_{s_1}^{s_2}[\varepsilon]$, where $s_1(n) = n$ and $s_2(n) = n^{\Omega(\log n)}$ are as in [Theorem 4.6](#). By this theorem, we have that $\text{mCFGL}_{s_1}^{s_2}[\varepsilon]$ requires time $n^{\Omega(\log n)}$ to be solved under rETH.

Suppose that $\mathcal{L}_n$ is PAC-learnable in time T by $\mathcal{C}_s$, that is, suppose there is a PAC-learner $\mathcal{A}$ that for every $f \in \mathcal{L}_n$, given oracle access to $\text{Ex}(f, \mathcal{D})$, outputs in time $T = T(n+1/\varepsilon+1/\delta)$ a hypothesis h in $\mathcal{C}_s$ that, with probability $1 - \delta$, satisfies $\text{err}_{f, \mathcal{D}}(h) \leq \varepsilon$. Run the algorithm $\mathcal{A}$, replacing oracle calls to Ex with samples from $\mathcal{D}$, and let C be the circuit obtained at the end. Since $|\mathcal{D}| = \text{poly}(n)$, this can be done in time $T \cdot \text{poly}(n)$. We now take $\ell = \frac{4}{\varepsilon^2}$ samples from $\mathcal{D}$, and accept $\mathcal{D}$ if C computes correctly a $2/3$ -fraction of the examples.

We claim that in the case where there exists a monotone formula of size s_1 computing all samples of $\mathcal{D}$ correctly, we will be correct with probability at least $1 - \delta - \exp(-2\ell\varepsilon^2) \geq 2/3$. Indeed, in this case, with probability $1 - \delta$ it holds that $\text{err}_{c, \mathcal{D}}(C) \leq \varepsilon$. Now suppose C is such that $\text{err}_{c, \mathcal{D}}(C) \leq \varepsilon$. By Chernoff-Hoeffding's inequality we have that the probability that C does not compute correctly $1 - 2\varepsilon \geq 2/3$ of the ℓ examples is at most $\exp(-2\ell\varepsilon^2)$. The claim follows.

In the other case, we know that $\mathbb{P}_{(x,b) \leftarrow \mathcal{D}}[C(x) = b] \leq 1/2 + \varepsilon$, and thus the proportion of samples that C correctly computes is less than $1/2 + 2\varepsilon$ with probability $1 - \exp(-2\ell\varepsilon^2)$ by Chernoff-Hoeffding's inequality. The total runtime is $T \cdot \text{poly}(n) + \ell \cdot \text{poly}(s)$. This implies that under rETH, by [Theorem 4.6](#), $T = n^{\Omega(\log n)}$. $\square$

Acknowledgements

We thank Noel Arteche for helpful discussions, and the anonymous reviewers for their valuable comments and suggestions. Rahul Santhanam and Bruno Cavalar acknowledge support from the EPSRC project EP/Z534158/1 on “Integrated Approach to Computational Complexity: Structure, Self-Reference and Lower Bounds”. Bruno Cavalar also acknowledges support of the UK Research and Innovation (UKRI) through an EPSRC Postdoctoral Fellowship (grant ref. UKRI3441). Susanna F. de Rezende was supported by the Knut and Alice Wallenberg Foundation grant KAW 2023.0116, ELLIIT, and the Swedish Research Council grant 2021-05104.

References

- [AAdRK25] Noel Arteche, Albert Atserias, Susanna F. de Rezende, and Erfan Khaniki. The proof analysis problem. In *Proceedings of the 66th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 2555–2576, 2025. [doi:10.1109/focs63196.2025.00133](https://doi.org/10.1109/focs63196.2025.00133).
- [ABF⁺08] Michael Alekhnovich, Mark Braverman, Vitaly Feldman, Adam R. Klivans, and Toniann Pitassi. The complexity of properly learning simple concept classes. *J. Comput. Syst. Sci.*, 74(1):16–34, 2008. [doi:10.1016/J.JCSS.2007.04.011](https://doi.org/10.1016/J.JCSS.2007.04.011).

- [ABX08] Benny Applebaum, Boaz Barak, and David Xiao. On basing lower-bounds for learning on worst-case assumptions. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science FOCS '08*, pages 211–220, 2008. doi:10.1109/FOCS.2008.35.
- [AD08] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, 2008. Preliminary version in *CCC '03*. doi:10.1016/j.jcss.2007.06.025.
- [AHM⁺08] Eric Allender, Lisa Hellerstein, Paul McCabe, Toniann Pitassi, and Michael E. Saks. Minimizing disjunctive normal form formulas and AC^0 circuits given a truth table. *SIAM J. Comput.*, 38(1):63–84, 2008. doi:10.1137/060664537.
- [AM20] Albert Atserias and Moritz Müller. Automating resolution is NP-hard. *Journal of the ACM*, 67(5), 2020. Preliminary version in *FOCS '19*. doi:10.1145/3409472.
- [BW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001. doi:10.1145/375827.375835.
- [CdR25] Gaia Carenini and Susanna F. de Rezende. On the automatability of tree-like k-DNF resolution. In *Proceedings of the 40th Computational Complexity Conference (CCC '25)*, volume 339 of *LIPICs*, pages 14:1–14:21, 2025. doi:10.4230/LIPICs.CCC.2025.14.
- [CIKK16] Marco L. Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. Learning algorithms from natural proofs. In Ran Raz, editor, *Proceedings of the 31st Conference on Computational Complexity (CCC '16)*, volume 50 of *LIPICs*, pages 10:1–10:24, 2016. doi:10.4230/LIPICs.CCC.2016.10.
- [Dan16] Amit Daniely. Complexity theoretic limitations on learning halfspaces. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC 16)*, pages 105–117, 2016. doi:10.1145/2897518.2897520.
- [DLM⁺09] Dana Dachman-Soled, Homin K. Lee, Tal Malkin, Rocco A. Servedio, Andrew Wan, and Hoeteck Wee. Optimal cryptographic hardness of learning monotone functions. *Theory Comput.*, 5(1):257–282, 2009. doi:10.4086/TOC.2009.V005A013.
- [dR21] Susanna F. de Rezende. Automating tree-like resolution in time $n^{o(\log n)}$ is ETH-hard. In *Proceedings of the XI Latin and American Algorithms, Graphs and Optimization Symposium, (LAGOS '21)*, volume 195, pages 152–162, 2021. doi:10.1016/J.PROCS.2021.11.021.
- [dRFJ⁺24] Susanna F. de Rezende, Noah Fleming, Duri Andrea Janett, Jakob Nordström, and Shuo Pang. Truly supercritical trade-offs for resolution, cutting planes, monotone circuits, and weisfeiler-leman. Technical report, Electronic Colloquium on Computational Complexity (ECCC), 2024. URL: <https://eccc.weizmann.ac.il/report/2024/185>, arXiv:TR24-185.
- [dRGN⁺21] Susanna F. de Rezende, Mika Göös, Jakob Nordström, Toniann Pitassi, Robert Robere, and Dmitry Sokolov. Automating algebraic proof systems is NP-hard. In *Proceedings of*

- the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC '21), pages 209–222, 2021. doi:[10.1145/3406325.3451080](https://doi.org/10.1145/3406325.3451080).
- [dRV25] Susanna F. de Rezende and Marc Vinyals. Lifting with colourful sunflowers. In *Proceedings of the 40th Computational Complexity Conference (CCC 25)*, pages 36:1–36:19, 2025. doi:[10.4230/LIPICS.CCC.2025.36](https://doi.org/10.4230/LIPICS.CCC.2025.36).
- [DS16] Amit Daniely and Shai Shalev-Shwartz. Complexity theoretic limitations on learning DNF’s. In *Proceedings of the 29th Conference on Learning Theory (COLT '16)*, pages 815–830, 2016. URL: <http://proceedings.mlr.press/v49/daniely16.html>.
- [DV21] Amit Daniely and Gal Vardi. From local pseudorandom generators to hardness of learning. In *Proceedings of the 34th Conference on Learning Theory (COLT '21)*, pages 1358–1394. PMLR, 2021. URL: <http://proceedings.mlr.press/v134/daniely21a.html>.
- [Fei02] Uriel Feige. Relations between average case complexity and approximation complexity. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing (STOC 02)*, pages 534–543, 2002. doi:[10.1145/509907.509985](https://doi.org/10.1145/509907.509985).
- [FLS11] Vitaly Feldman, Homin K. Lee, and Rocco A. Servedio. Lower bounds and hardness amplification for learning shallow monotone formulas. In *Proceedings of the 24th Conference on Learning Theory (COLT '11)*, pages 273–292, 2011. URL: <http://proceedings.mlr.press/v19/feldman11a/feldman11a.html>.
- [GGKS20] Ankit Garg, Mika Göös, Pritish Kamath, and Dmitry Sokolov. Monotone circuit lower bounds from resolution. *Theory of Computing*, 16(1):1–30, 2020. Preliminary version in *STOC '18*. doi:[10.4086/toc.2020.v016a013](https://doi.org/10.4086/toc.2020.v016a013).
- [GKMP20] Mika Göös, Sajin Koroth, Ian Mertz, and Toniann Pitassi. Automating cutting planes is NP-hard. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC 20)*, pages 68–77, 2020. doi:[10.1145/3357713.3384248](https://doi.org/10.1145/3357713.3384248).
- [GKRS19] Mika Göös, Pritish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in monotone complexity and TFNP. In *Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 38:1–38:19, 2019. doi:[10.4230/LIPICS.ITCS.2019.38](https://doi.org/10.4230/LIPICS.ITCS.2019.38).
- [GPW18] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. *SIAM J. Comput.*, 47(6):2435–2450, 2018. Preliminary version in FOCS '15. doi:[10.1137/16M1059369](https://doi.org/10.1137/16M1059369).
- [Hir22] Shuichi Hirahara. NP-hardness of learning programs and partial MCSP. In *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science FOCS 22*, pages 968–979, 2022. doi:[10.1109/FOCS54457.2022.00095](https://doi.org/10.1109/FOCS54457.2022.00095).
- [Hir23] Shuichi Hirahara. Non-black-box worst-case to average-case reductions within NP. *SIAM J. Comput.*, 52(6):S18–349, 2023. doi:[10.1137/19M124705X](https://doi.org/10.1137/19M124705X).

- [HIR25] Yizhi Huang, Rahul Ilango, and Hanlin Ren. NP-hardness of approximating meta-complexity: A cryptographic approach. *SIAM J. Comput.*, 54(4):819–886, 2025. doi:[10.1137/23M1608483](https://doi.org/10.1137/23M1608483).
- [HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc.*, 43(04):439–562, August 2006. doi:[10.1090/s0273-0979-06-01126-8](https://doi.org/10.1090/s0273-0979-06-01126-8).
- [HOS18] Shuichi Hirahara, Igor C. Oliveira, and Rahul Santhanam. NP-hardness of minimum circuit size problem for OR-AND-MOD circuits. In *Proceedings of the 33rd Computational Complexity Conference (CCC 18)*, pages 5:1–5:31, 2018. doi:[10.4230/LIPICS.CCC.2018.5](https://doi.org/10.4230/LIPICS.CCC.2018.5).
- [Ila20a] Rahul Ilango. Approaching MCSP from above and below: Hardness for a conditional variant and $AC^0[p]$. In *Proceedings of the 11th Innovations in Theoretical Computer Science Conference (ITCS '20)*, pages 34:1–34:26, 2020. doi:[10.4230/LIPICS.ITCS.2020.34](https://doi.org/10.4230/LIPICS.ITCS.2020.34).
- [Ila20b] Rahul Ilango. Constant depth formula and partial function versions of MCSP are hard. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS '20)*, pages 424–433, 2020. doi:[10.1109/FOCS46700.2020.00047](https://doi.org/10.1109/FOCS46700.2020.00047).
- [Ila21] Rahul Ilango. The minimum formula size problem is (ETH) hard. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS '21)*, pages 427–432, 2021. doi:[10.1109/FOCS52979.2021.00050](https://doi.org/10.1109/FOCS52979.2021.00050).
- [ILO20] Rahul Ilango, Bruno Loff, and Igor C. Oliveira. NP-hardness of circuit minimization for multi-output functions. In *Proceedings of the 35th Computational Complexity Conference (CCC '20)*, pages 22:1–22:36, 2020. doi:[10.4230/LIPICS.CCC.2020.22](https://doi.org/10.4230/LIPICS.CCC.2020.22).
- [Imp95] Russell Impagliazzo. A personal view of average-case complexity. In *Proceedings of the 10th Annual IEEE Structure in Complexity Theory Conference (SCT 95)*, pages 134–147, 1995. doi:[10.1109/SCT.1995.514853](https://doi.org/10.1109/SCT.1995.514853).
- [Jeř07] Emil Jeřábek. On independence of variants of the weak pigeonhole principle. *Journal of Logic and Computation*, 17(3):587–604, 2007. doi:[10.1093/logcom/exm017](https://doi.org/10.1093/logcom/exm017).
- [KC00] Valentine Kabanets and Jin-yi Cai. Circuit minimization problem. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC '00)*, pages 73–79, 2000. doi:[10.1145/335305.335314](https://doi.org/10.1145/335305.335314).
- [Kha93] Michael Kharitonov. Cryptographic hardness of distribution-specific learning. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC '93)*, pages 372–381, 1993. doi:[10.1145/167088.167197](https://doi.org/10.1145/167088.167197).
- [Kra97] Jan Krajíček. Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. *The Journal of Symbolic Logic*, 62(2):457–486, 1997. doi:[10.2307/2275541](https://doi.org/10.2307/2275541).

- [Kra98] Jan Krařek. Interpolation by a game. *Mathematical Logic Quarterly*, 44(4):450–458, 1998. doi:[10.1002/malq.19980440403](https://doi.org/10.1002/malq.19980440403).
- [KV94] Michael J. Kearns and Leslie G. Valiant. Cryptographic limitations on learning Boolean formulae and finite automata. *J. ACM*, 41(1):67–95, 1994. doi:[10.1145/174644.174647](https://doi.org/10.1145/174644.174647).
- [LMM⁺22] Shachar Lovett, Raghu Meka, Ian Mertz, Toniann Pitassi, and Jiapeng Zhang. Lifting with sunflowers. In *Proceedings of 13th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 104:1–104:24, 2022. doi:[10.4230/LIPICS.ITCS.2022.104](https://doi.org/10.4230/LIPICS.ITCS.2022.104).
- [Mas79] William Masek. Some NP-complete set covering problems. Unpublished, 1979.
- [MP24] Noam Mazon and Rafael Pass. Gap MCSP is not (Levin) NP-complete in Obfuscopia. In *Proceedings of the 39th Computational Complexity Conference, (CCC '24)*, pages 36:1–36:21, 2024. doi:[10.4230/LIPICS.CCC.2024.36](https://doi.org/10.4230/LIPICS.CCC.2024.36).
- [PT19] Pavel Pudlák and Neil Thapen. Random resolution refutations. *Computational Complexity*, 28(2):185–239, 2019. doi:[10.1007/s00037-019-00182-7](https://doi.org/10.1007/s00037-019-00182-7).
- [Pud03] Pavel Pudlák. On reducibility and symmetry of disjoint NP pairs. *Theoretical Computer Science*, 295:323–339, 2003. doi:[10.1016/S0304-3975\(02\)00411-5](https://doi.org/10.1016/S0304-3975(02)00411-5).
- [Raz95] Alexander A Razborov. Unprovability of lower bounds on circuit size in certain fragments of bounded arithmetic. *Izvestiya: mathematics*, 59(1):205, 1995. doi:[10.1070/IM1995v059n01ABEH000009](https://doi.org/10.1070/IM1995v059n01ABEH000009).
- [RM99] Ran Raz and Pierre McKenzie. Separation of the Monotone NC Hierarchy. *Combinatorica*, 19(3):403–435, 1999. Preliminary version in *STOC* 97. doi:[10.1007/s004930050062](https://doi.org/10.1007/s004930050062).
- [RR97] Alexander A. Razborov and Steven Rudich. Natural proofs. *J. Comput. Syst. Sci.*, 55(1):24–35, 1997. doi:[10.1006/JCSS.1997.1494](https://doi.org/10.1006/JCSS.1997.1494).
- [San20] Rahul Santhanam. Pseudorandomness and the minimum circuit size problem. In Thomas Vidick, editor, *Proceedings of the 11th Innovations in Theoretical Computer Science Conference (ITCS '20)*, pages 68:1–68:26, 2020. doi:[10.4230/LIPICS.ITCS.2020.68](https://doi.org/10.4230/LIPICS.ITCS.2020.68).

A Proof Sketch of the Refined Lifting Theorem

Throughout this appendix, we assume familiarity with the notation and definitions from the paper [dRV25]⁵. **Theorem 2.4** is proved in [dRV25] only for the case when $\gamma = 1/2$. However, it is straightforward to extend it to the case $\gamma \geq 2m^{-(1-\delta)}$ at the expense of only a slightly weaker size lower bound, that is, of showing that the protocol is of size at least

⁵In particular, we refer to the [full version](#) of the paper.

$m^{(1-\delta)(\ell-1)}$ (instead of $m^{(1-\delta)}/2$). Here δ is the same parameter as in [dRV25], that is, δ such that $m^\delta = 2K \cdot |\Sigma| \cdot 4\ell \cdot \log(2mN^2)$ for K being the constant given from the Full Range Lemma. Note that in our case $\Sigma = \{0, 1\}$. Moreover, observe that $m^\delta \leq c'\ell \log(mN)$ for a large enough absolute constant c' , and so, in particular, we will be proving the statement for $\gamma \geq 2c'\ell \log(mN)/m \geq 2m^{-(1-\delta)}$.

Suppose that some small circuit has $1/2 + \gamma$ correlation with f . This implies that the root rectangle $R = R^X \times R^Y$ obtained from the circuit has density at least γ in both sides (instead of density 1, as in [dRV25]). Suppose that the DAG has size at most $m^{(1-\delta)(\ell-1)} \leq m^{(1-\delta)\ell} \cdot \gamma/2$. This changes the density of both the X -error and of the Y -error accumulated at the root to be at most $\gamma/2$ (instead of at most $1/2$), and, therefore, the root rectangle, with errors removed, has density at least $\gamma/2$ in both sides. This only affects the argument that $R^X \times (R^Y \setminus Y_{\text{err}}^R)$ is a $\star^N$ -(pre)-structured rectangle, that is, that R^X is $\star^N$ -predense and $|R^Y \setminus Y_{\text{err}}^R| \geq 2^{mN-4\ell \log mN}$ (since in our case $\Sigma = \{0, 1\}$).

We first argue that R^X is $\star^N$ -predense. Indeed, for any $\emptyset \neq I \subseteq X$ the min-entropy of R_I^X is at least

$$\begin{aligned} |I| \log m - \log(2/\gamma) &= |I| \log m - (1 - \delta) \log m \\ &\geq |I| \log m - (1 - \delta) |I| \log m \\ &\geq \delta |I| \log m. \end{aligned}$$

We have thus obtained the bound required to conclude that R^X is $\star^N$ -predense.

It remains to show that $|R^Y \setminus Y_{\text{err}}^R| \geq 2^{mN-4\ell \log mN}$. This follows easily since $|R^Y \setminus Y_{\text{err}}^R| \geq 2^{mN} \cdot \gamma/2$ and $\gamma/2 = m^{-(1-\delta)} \geq 2^{-\log m} \geq 2^{-4\ell \log mN}$. We therefore conclude that $R^X \times (R^Y \setminus Y_{\text{err}}^R)$ is a $\star^N$ -(pre)-structured rectangle, and the rest of the proof is the same.

This implies the size of lower bound of $m^{(1-\delta)(\ell-1)} \geq (m/c'\ell \log(mN))^{\ell-1}$ for any $\gamma \geq 2c'\ell \log(mN)/m$. Choosing the absolute constant $c = 2c'$ gives us [Theorem 2.4](#).

B Proof of the Lower Bound for the Refutation Formula

In this appendix, we prove [item 4 of Theorem 3.5](#), that is, we show that if F is unsatisfiable 3-CNF formula over n variables, with $\text{poly}(n)$ clauses, $d \leq n$ and $t \leq 2^{n/d}$, and G is a $(t \cdot 2^{n/d}, t, \delta, r, c)$ -bipartite expander graph, then $\text{Ref}_d^G(F)$ requires Resolution proofs of width at least $2^{\Omega(r(c-1)/n)}$. The proof follows along the same lines of the ones in [dRGN⁺21] and in [dR21].

B.1 Weak Graph rPHP Formula

The formula we use is a graph version of the *retraction pigeonhole principle*, rPHP [Jeř07, PT19]. The retraction pigeonhole principle was also used in [dRGN⁺21] to prove a similar reduction.

In this variant of the pigeonhole principle, the pigeon-mapping is restricted to the edges of a given bipartite graph and, moreover, there is an efficient way to *invert* the mapping. Formally, given a bipartite graph $G = (U \cup V, E)$ with $U = [m]$ and $V = [n]$, and an arbitrary family of surjections $\text{sur}_i : [\delta] \rightarrow N(i)$ for each $i \in [m]$, the variables of the formula $\text{rPHP}(G)$ describe two functions, $f : [m] \rightarrow [\delta]$ and $g : [n] \rightarrow [m]$, as follows.

- *Pigeon map variables.* For every pigeon $i \in [m]$ there are variables f_{ik} , $k \in [\log \delta]$. The pigeon map variables encode in binary a hole $\text{sur}_i(f(i)) \in N(i)$ that is expected to house pigeon i .
- *Hole map variables.* For every hole $j \in [n]$ there are variables $g_{j\ell}$, $\ell \in [\log m]$. These variables encode in binary a pigeon $g(j) \in [m]$ that is expected to occupy hole j .

The axioms of $\text{rPHP}(G)$ state that for every $i \in [m]$ and $j \in N(i)$,

$$\text{sur}_i(f(i)) = j \implies g(j) = i. \quad (10)$$

Note that (10) corresponds to $\log m$ clauses each of width δ . Therefore, $\text{rPHP}(G)$ can be written as an $O(\delta)$ -width CNF formula with $O(n\delta \log m)$ clauses.

It was shown in [BW01] that if G is a good enough expander, then the pigeonhole principle formula over G requires large width to be refuted in resolution. As argued in [dRV25], we can obtain the following, slightly stronger, lower bound on the width required to refute the pigeonhole principle formula over G , which in particular applies to our encoding $\text{rPHP}(G)$.

Theorem B.1 ([dRV25]). *If G is an (m, n, δ, r, c) -bipartite expander graph then*

$$\frac{(c-1)r}{2} \leq \text{ResWidth}(\text{rPHP}(G)) \leq n+1.$$

In particular, this implies that if G is an $(m, n, \delta, r, \delta/2)$ -bipartite expander graph from Theorem 2.2 then $\text{ResWidth}(\text{rPHP}(G)) = \Omega(\sqrt{n})$. And if $G = ((U \cup V), E)$ is a complete bipartite graph with $|U| = m$, $|V| = n$, then by Theorem 2.1 $\text{ResWidth}(\text{rPHP}(G)) \geq n/4$.

B.2 Decision Tree Reductions

Our goal is to show, via a moderately small-depth reduction, that the width required for refuting $\text{Ref}_d^G(F)$ is almost as large as the width required for refuting $\text{rPHP}(G)$. For the definition of reduction, we view a clause as a function, that is, given a CNF formula H over m variables, we view a clause $D \in H$ as a function from $\{0, 1\}^m$ to $\{0, 1\}$, which can be computed by a depth- $|D|$ decision tree. In this way, given a function $r : \{0, 1\}^n \rightarrow \{0, 1\}^m$, $D \circ r$ is a function from $\{0, 1\}^n$ to $\{0, 1\}$. Moreover, if r_i is computed by a depth- d decision tree, then $D \circ r$ can be naturally written as a $(d \cdot |D|)$ -CNF formula (i.e., a CNF formula with width $d \cdot |D|$) over n variables.

Definition B.2 ([dRGN⁺21]). *Given two CNF formulas F and H over n and m variables respectively, there is a depth- d reduction from F to H , denoted $F \leq_d H$, if there is a function $r : \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that*

- *each output bit of $r_i : \{0, 1\}^n \rightarrow \{0, 1\}$ for $i \in [m]$ is computed by a depth- d decision tree; and*
- *for every axiom $D \in H$, each clause in the representation of $D \circ r$ as a $(d \cdot |D|)$ -CNF formula (as discussed above) is a weakening of some axiom of F .*

The important property of reductions that we use is that width bounds carry over.

Lemma B.3 ([dRGN⁺21]). *If $F \leq_d H$, then $\text{ResWidth}(F) \leq d \cdot \text{ResWidth}(H)$.*

For the rest of this section, we let G be a $(t \cdot 2^{n/d}, t, \delta, r, c)$ -bipartite expander graph, where $t \leq 2^{n/d}$. Combining [Theorem B.1](#) and [Theorem B.3](#) implies that, in order to obtain the lower bound in [Theorem 3.5](#), it is enough to show that

$$\text{rPHP}(G) \leq_{O(n)} \text{Ref}_d^G(F). \quad (11)$$

B.3 Overview of the Reduction

The proof of the reduction is similar to that of [dR21] which in turn follows the one in [dRGN⁺21]. As in [dR21], the reduction is guided by a full tree $\mathcal{T}$ of height d and arity $2^{n/d}$. Recall that we consider a fixed ordering of the variables of F . Each node of $\mathcal{T}$ contains a clause. The root contains the empty clause ($\perp$) and each node at distance α from the root contains a clause with $\alpha n/d$ defined recursively as follows. If a non-leaf node at distance α from the root contains a clause C using the first $\alpha n/d$ variables, each of its $2^{n/d}$ children contains one of the possible weakenings of C to the next n/d variables, that is, for $\sigma \in [2^{n/d}]$, the σ -th child contains the clause $C \cup \{x_{\alpha n/d}^{\sigma_0}, x_{\alpha n/d+1}^{\sigma_1}, \dots, x_{\alpha n/d+n/d-1}^{\sigma_{n/d-1}}\}$ (recall that $x_i^b = \neg x_i$ if $b = 0$ and x_i otherwise). Thus, $\mathcal{T}$ has 2^n leaves, each containing a different clause of width n .

B.4 The Reduction Function

We start by describing the decision trees that compute the variables of $\text{Ref}_d^G(F)$. For simplicity, we sometimes describe one decision tree that computes several variables at once.

Pointer variables. The pointer variables of Ref_d^G are defined in terms of the f -variables of $\text{rPHP}(G)$. Since each block has $2^{n/d}$ pointers, each layer $\alpha \in \{0, \dots, d-1\}$ has a total of $t \cdot 2^{n/d}$ pointers that should be mapped to the t blocks at level $\alpha+1$. The idea is to define this mapping by “copying” the $\text{rPHP}(G)$ mapping from the $t \cdot 2^{n/d}$ pigeons to the t holes. Formally, let B be an internal block, let $z \in [2^{n/d}]$ and let $i = \text{col}(B) \cdot 2^{n/d} + z$. For all $k \in [\delta]$,

we define the pointer variables $\text{pointbin}_{z,k}^B = f_{i,k}$ and for all B' we define $\text{point}_{B',z}^B = 1$ if and only if $\text{col}(B') = \text{sur}_i(f(i))$. This can clearly be done by a decision tree of depth $\delta = O(n)$.

Enabled and literal variables. We have now reached the most complicated part of the reduction. Consider the function $g': [t] \rightarrow [t \cdot 2^{n/d}] \cup \{\star\}$ defined as

$$g'(j) = \begin{cases} g(j) & \text{if } \text{sur}_{g(j)}(f(g(j))) = j, \\ \star & \text{otherwise.} \end{cases} \quad (12)$$

In order to determine whether a block B is enabled or not and what literal set it contains, we recursively define a path of blocks and pointers starting from B , denoted $\text{path}(B)$, in the following way. Let $j = \text{col}(B)$. If B is the root or if $g'(j) = \star$, then the path ends at B , that is, $\text{path}(B) = B$. Otherwise, let

$$j' = \left\lfloor \frac{g'(j)}{t} \right\rfloor \quad \text{and} \quad \sigma = g'(j) \bmod t. \quad (13)$$

Note that j' and σ are defined so that the σ th pointer of the block B' with $\text{col}(B') = j'$ and $\text{layer}(B') = \alpha - 1$ is j . We define $\text{path}(B)$ to be (B, σ) concatenated with $\text{path}(B')$.

We are now ready to define a decision tree that computes the enabled variable and the literal variables of the block B . In order to do so, the decision tree must compute $\text{path}(B)$, which can be done in depth $O(d(\log(t \cdot 2^{n/d}) + \log \delta)) = O(n)$. Indeed, the length of $\text{path}(B)$ is at most d and to determine the next block in a path it is enough to compute g' , which can be done with $O(\log(t \cdot 2^{n/d})) = O(n/d)$ queries to the g -variables and $O(\log \delta) \leq O(n/d)$ queries to the f -variables of $\text{rPHP}(G)$.

Once $\text{path}(B)$ is determined, the decision tree outputs the value of the variables as follows. If $\text{path}(B)$ does not end at the root, the enabled variable $\text{enable}^B = 0$, the literal variables $\text{unlit}_{x_i}^B = 0$ for $\ell \in \text{Unlit}$ (if B has unlit variables), and the summary variables $\text{summ-unlit}_{i,z}^B$ are all set to 1. If $\text{path}(B)$ ends at the root, let π be the reverse of $\text{path}(B)$. Note that π defines a unique path in the tree $\mathcal{T}$ starting from the root and following the pointers in π . This path in $\mathcal{T}$ ends in some block B' at level α . The decision tree sets the enabled variable enable^B to 1 and sets the literal variables unlit_ℓ^B for $\ell \in \text{Unlit}$ to match the assignment of the block B' .

Weakening variables. The weakening variables are computed “honestly” given the assignments contained at the blocks at the last layer. That is, in order to determine the weakening variables for the block B , the decision tree computes $\text{path}(B)$ (in depth $O(n)$), in this way determining the literal variables of B , and then it outputs the value of each weakening variable weak and weakbin so that this block points to an axiom A of F which it is the weakening of. Since there are n literals in B and F is unsatisfiable, such an A must exist. We note that, if $\text{path}(B)$ does not end at the root, then the literal variables and the enabled variable of B

are all set to 0 and the weakening variables can point to an arbitrary axiom.

Therefore, we conclude that the reduction r satisfies the required condition that for each i there is a decision tree of depth $O(n)$ that computes r_i .

B.5 Correctness of the Reduction

We now show that the reduction satisfies the second condition of [Theorem B.2](#). Recall that r is the function defined by the decision trees described above. Let D be an axiom of $\text{Ref}_d^G(F)$. It is enough to show that for every partial assignment ρ to the variables of $\text{rPHP}(G)$, if ρ falsifies $D \circ r$ then ρ falsifies an axiom of $\text{rPHP}(G)$.

We first note that, by definition of r , if D is the root axiom [REF-1](#); a binary pointer axiom [REF-5](#); a leaf axiom [REF-2](#), [REF-3](#), or [REF-4](#); or a backchecking axiom [REF-7](#) or [REF-8](#), then $D \circ r$ is a tautology, that is, for any full assignment ρ to the variables of $\text{rPHP}(G)$, the formula $D \circ r$ is satisfied.

It remains to consider the case when D is one of the [REF-6](#), [REF-7](#), or [REF-8](#) axioms, that is $(\text{enable}^B \wedge \text{point}_{B',z}^B) \rightarrow \text{enable}^{B'}$, $(\text{enable}^B \wedge \text{point}_{B',z'}^B \wedge \text{summ-unlit}_{i',z}^{B'}) \rightarrow \text{summ-unlit}_{i',z}^{B'}$ or $(\text{enable}^B \wedge \text{point}_{B',z}^B) \rightarrow \text{summ-unlit}_{i+1,z}^{B'}$, respectively, for some $B, B' \in [s]$ with $i' \leq i = \text{layer}(B) = \text{layer}(B') - 1$. In all three cases, for these clauses to be falsified, both enable^B and $\text{point}_{B',z}^B$ must be set to 1. Note that enable^B will only be set to 1 if $\text{path}(B)$ ends at the root and $\text{point}_{B',z}^B$ will only be set to 1 if $\text{col}(B') = \text{sur}_j(f(j))$ for $j = \text{col}(B) \cdot 2^{n/d} + z$.

Consider when D is one of the [REF-6](#) axioms. Note that $\text{enable}^{B'}$ will only be set to 0 in the reduction if $\text{path}(B')$ does not end at the root. This implies that the value returned for $g(\text{col}(B'))$ was something other than j (corresponding to the z th pointer of B), since otherwise $g'(\text{col}(B'))$ would be equal to $g(\text{col}(B'))$ and $\text{path}(B')$ would be $\text{path}(B)$ concatenated with (B', B) and would, therefore, reach the root. Since $\text{col}(B') = \text{sur}_j(f(j))$, the clause $\text{sur}_j(f(j)) = \text{col}(B') \Rightarrow g(\text{col}(B')) = j$ from $\text{rPHP}(G)$ is violated.

Now consider the case when D is one of the [REF-7](#) or [REF-8](#) axioms. If the value returned for $g(\text{col}(B'))$ is j (corresponding to the z th pointer of B), then $\text{path}(B')$ would go through B via the z th pointer of B . This would imply that, in the case of [REF-7](#) both summ-unlit variables would be the same and thus [REF-7](#) would not be falsified; and in the case of [REF-8](#) $\text{summ-unlit}_{i+1,z}^{B'}$ would be 1, because $\text{point}_{B',z}^B$ is 1 and $\text{layer}(B) = i$, and thus [REF-8](#) would not be falsified. So it must be the case that $g(\text{col}(B'))$ is not j , and therefore we have that the clause $\text{sur}_j(f(j)) = \text{col}(B') \Rightarrow g(\text{col}(B')) = j$ from $\text{rPHP}(G)$ is violated.