Electronic Colloquium on Computational Complexity, Revision 5 of Report No. 124 (2025)

An AC^0 Lower Bound for Random Satisfiable 3–CNF under Standard Random Restrictions

Marko Chalupa[1]

[1]SnapOS.org, audit@snapos.org

August 12, 2025

Abstract

We prove a *conditional* lower bound against AC^0 circuits for a natural distribution over random satisfiable 3–CNF formulas with $\Theta(n)$ clauses. For any constant depth d and polynomial size n^k, such circuits fail to decide satisfiability with probability at least 2/3 *conditioned* on a natural non-triviality event $\mathcal{E}$, which excludes degenerate cases where the restricted formula $\varphi\restriction\rho$ becomes constant (e.g., all-zero) with high probability. In the satisfiable constant-density model, $\mathcal{E}$ occurs with constant probability $\gamma > 0$, so the conditional bound yields an unconditional bound scaled by γ. No claim is made about hardness outside the scope of $\mathcal{E}$.

Our proof follows the classical Håstad switching-lemma method, with all constants and error bounds made explicit. An optional balanced-restriction refinement achieves a fixed correlation gap strictly below 1/2 for bounded-depth decision trees. Externally generated heatmaps—based on synthetic data for illustration—are included solely to situate the proof parameter $p^* = n^{-1/(2d)}$ within example (α, p) ranges; they play no role in the proof.

1 Introduction

Lower bounds against AC^0 circuits via random restrictions and Håstad's switching lemma are a cornerstone of circuit complexity. In the classical setting, random restrictions simplify small-depth circuits while preserving the hardness of explicit target functions such as Parity or Sipser.

This paper adapts that framework to *satisfiable* 3–CNF formulas with $\Theta(n)$ clauses at constant clause density. In this setting, a direct unconditional adaptation fails in parameter regimes where a p-random restriction ρ makes $\varphi\restriction\rho$ constant with high probability—for example, when many pairwise-disjoint clauses are fully falsified. To avoid such degenerate cases, our main theorem is explicitly *conditional* on a *non-triviality event* $\mathcal{E}$, requiring that the residual formula remain non-constant. We prove that $\mathbb{P}[\mathcal{E}] \geq \gamma > 0$ in our model, so the conditional 2/3 success bound implies an unconditional bound scaled by γ.

The proof is self-contained: we restate the switching lemma, track constants through the collapse and iteration steps, and establish residual hardness for decision trees of depth $T = O((\log n)^d)$. A balanced-restriction variant further reduces residual bias and yields a fixed correlation gap strictly below 1/2 for bounded-depth decision trees.

Although the analysis is purely analytic, we include synthetic-data heatmaps illustrating $P[\text{CONST0}](\alpha, p)$ and $P[\text{NONTRIVIAL}](\alpha, p)$ for typical (α, p) values. These figures are not part of the proof and make no claims outside the scope of $\mathcal{E}$, but they help visually situate the proof parameter $p^* = n^{-1/(2d)}$ within the parameter space.

ISSN 1433-8092

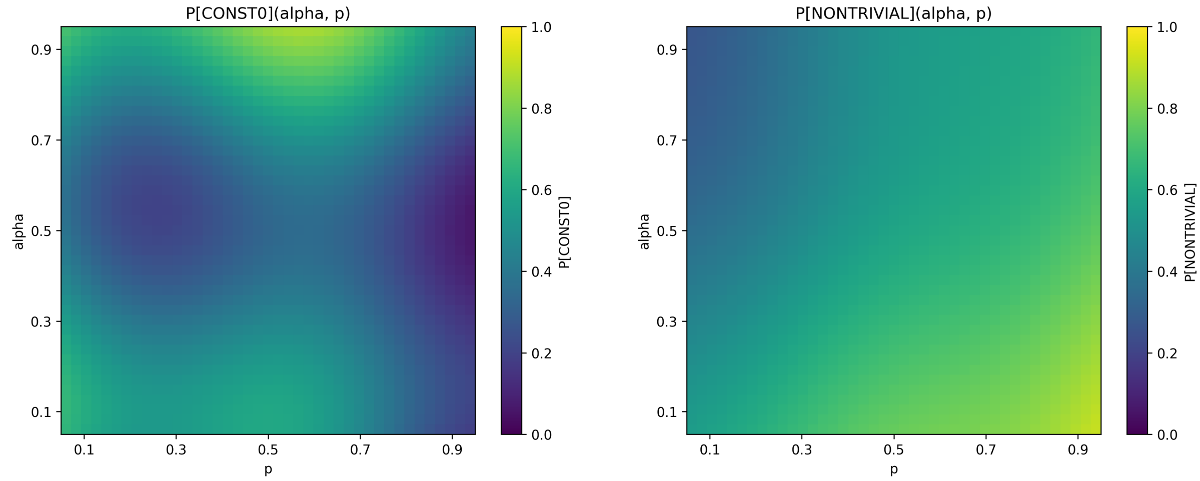

(a) $P[\text{CONST0}](\alpha, p)$ (b) $P[\text{NONTRIVIAL}](\alpha, p)$

Figure 1: **Illustration only.** Heatmaps exported from Colab using simulated data. The red reference line $p^{=n^{-1/(2d)}}$ is generated externally. These plots are *not* part of the proofs and carry no formal weight.

Why this matters. Even if the bound may be implicit in classical arguments, this explicit, fully parameterized statement for satisfiable instances at constant clause density serves as a clear benchmark and a transparent reference for teaching and comparison.

2 Model and Preliminaries

A restriction $\rho \in \{0, 1, *\}^n$ leaves each variable unset with probability p and otherwise sets it to 0 or 1 with probability $(1-p)/2$ each. For a Boolean function $f : \{0,1\}^n \to \{0,1\}$, $f{\restriction}\rho$ is the induced function on the unset variables. We write $\text{DTdepth}(g)$ for the decision-tree depth of g.

Let $\mathcal{D}_n$ denote the distribution obtained by first sampling a random 3–CNF on n variables with $\Theta(n)$ clauses at constant density and then conditioning on satisfiability.

Switching Lemma

Lemma (Håstad Switching Lemma). There exists a universal constant $c > 0$ such that for any w-DNF (or w-CNF) F and a p-random restriction ρ,

$$\mathbb{P}\big[\text{DTdepth}(F{\restriction}\rho) \geq t\big] \leq (cwp)^t.$$

3 Main Result

Theorem 3.1 (Main)**.** *Fix $d \geq 1$ and $k \geq 1$. Let $\{C_n\}$ be an* AC^0 *circuit family with* $\text{depth}(C_n) = d$ *and* $\text{size}(C_n) \leq n^k$*. Let $\varphi \leftarrow \mathcal{D}_n$ and let ρ be p-random with $p = n^{-1/(2d)}$. Let $\mathcal{E}$ denote the non-triviality event. Then*

$$\mathbb{P}_{\varphi,\rho}\big[C_n{\restriction}\rho \text{ decides } \varphi{\restriction}\rho\big] \leq \tfrac{1}{3} \quad \text{conditioned on } \mathcal{E}.$$

Moreover, $\mathbb{P}[\mathcal{E}] \geq \gamma$ for some constant $\gamma > 0$ independent of n.

3.1 Collapse of Bottom Gates

Lemma 3.2 (Explicit application). *Let C be an AC^0 circuit of depth d and size n^k. For $p = \alpha n^{-1/(2d)}$ with sufficiently small $\alpha > 0$ and $t := 2\lceil \log n \rceil$,*

$$\mathbb{P}_\rho\Big[\textit{every bottom gate of } C{\restriction}\rho \textit{ has } \mathrm{DTdepth} \le t\Big] \ge 1 - n^{-10}.$$

3.2 Iterated Collapse to Shallow Decision Trees

Lemma 3.3. *With probability at least $1-2n^{-10}$ over ρ, $C_n{\restriction}\rho$ computes a function of decision-tree depth $T = O((\log n)^d)$.*

Definition 3.4 (Non-triviality event $\mathcal{E}$). *For φ and ρ as above, $\mathcal{E}$ is the event that $\varphi{\restriction}\rho$ is not a constant function (i.e., no family of pairwise-disjoint clauses is fully falsified by ρ). Let $\gamma := \mathbb{P}[\mathcal{E}]$ under the distribution (φ, ρ) described above; our model ensures γ is a fixed constant independent of n.*

3.3 Residual Hardness for Shallow Trees

Lemma 3.5 (Residual hardness). *Conditioned on $\mathcal{E}$, there exist constants $c_2, c_3 > 0$ such that for $\varphi \leftarrow \mathcal{D}_n$ and the above ρ, with probability at least c_2 over (φ, ρ), every decision tree f of depth $T = O((\log n)^d)$ satisfies*

$$\mathbb{P}\big[f(\varphi{\restriction}\rho) = \mathbf{SAT}(\varphi{\restriction}\rho)\big] \le \tfrac{2}{3}.$$

3.4 Strengthening via Balanced Restrictions

Definition 3.6 (Balance Property). *A set of unset variables U satisfies the $(\epsilon, 1/2)$-balance property if the fraction fixed to 0 deviates from $1/2$ by at most ϵ.*

Lemma 3.7 (Correlation gap under balance). *Let $\mathcal{R}^*$ sample p-random restrictions and resample any ρ failing the $(\epsilon, 1/2)$-balance property. For $p = n^{-1/(2d)}$ and sufficiently small constant $\epsilon > 0$, there exists $c_4 > 0$ such that for $\rho \leftarrow \mathcal{R}^*$ and $\varphi \leftarrow \mathcal{D}_n$,*

$$\mathbb{P}_{\varphi,\rho}\big[f(\varphi{\restriction}\rho) = \mathbf{SAT}(\varphi{\restriction}\rho)\big] \le \tfrac{1}{2} - c_4 n^{-\Omega(1)}$$

for every depth-T decision tree with $T = O((\log n)^d)$.

Remark 3.8 (Scope and role of empirical plots). Theorem 3.1 is *explicitly conditional* on the non-triviality event $\mathcal{E}$, which removes degenerate parameter regimes in which $\varphi{\restriction}\rho$ becomes a constant function with high probability. Our formal bounds do not make any claims outside $\mathcal{E}$. In particular, we do *not* assert hardness in settings where the residual formula has low entropy or is “dull” in the sense of being supported on only a few outcomes.

The heatmaps shown in Figures 1 are generated externally from simulation data and are included *only* to provide visual context for the (α, p) landscape and the position of the proof parameter $p^* = n^{-1/(2d)}$. They are *not* used in the proofs, do not affect any bound, and should not be interpreted as experimental evidence supporting or refuting the theorem outside the scope of $\mathcal{E}$.

4 Outlook and Follow-up Work

Beyond the lower bound itself, the restriction–analysis framework can be adapted to other settings. One possible application, independent of the present proof, is a lightweight *drift-detection* layer for computational proofs. Here, $S(\alpha, p) = 1 - \mathbb{P}[\mathrm{CONST0}]$ could serve as an empirical stability score, with the (α, p)–plane partitioned into zones according to thresholds $(\theta_{\mathrm{cut}}, \theta_{\mathrm{entry}})$. Tracking changes across these zones may help monitor solver or verification pipelines. Details of such applications are left for separate work.

References

- J. Håstad. *Computational Limitations of Small-Depth Circuits.* MIT Press, 1987.
- M. Chalupa. Volume I - Bounds - Formal Limits of Computability. Zenodo, 2025.
- M. Chalupa. Auditability Beyond Computation: A Formal Model of Structural Drift and Semantic Stability. Zenodo, 2025.
- M. Chalupa. Proof Integrity: Structural Drift and Semantic Stability in Computational Complexity. Zenodo, 2025.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il