1 What do we know?

A misguided sentiment is that "we don't know any good circuit lower bounds, and therefore proving circuit lower bounds is a barrier". The word "good" is carrying too much weight in such a sentence, and it is not specified enough. We've known interesting and non-trivial circuit lower bounds for at least four decades, even for circuits of near-maximal size, and the question is just *which* lower bounds we know.

Here's a brief overview of lower bounds that we know for general circuits. (We know much more for restricted types of circuits.) We'd like to prove $f \notin \mathcal{SIZE}[s]$ for f as explicit as possible and s as large as possible. Two open problems are to prove:

- **Non-uniformity can't replace proofs:** $\mathcal{NP} \not\subset \cup_c \mathcal{SIZE}[n^c]$. This is important, because it implies $\mathcal{P} \neq \mathcal{NP}$.
- **Non-uniformity can't save polytime:** $\mathcal{E} \not\subset \mathcal{SIZE}[2^{\epsilon \cdot n}]$, for some $\epsilon > 0$. A scaled-down version of this lower bound may seem more natural, e.g. $\forall k$, $\mathcal{P} \not\subset \mathcal{SIZE}[n^k]$.[1] These lower bounds imply derandomization of $\mathcal{BPP}$.

Let's abstract these target lower bounds as

$$\mathcal{CTIME}[t] \not\subset \mathcal{SIZE}[s] \ .$$

Note that in the first type of lower bound above, t and s are the same (i.e., arbitrary polynomial time/size), but $\mathcal{C}$ is allowed non-determinism. In the second type of lower bound, $\mathcal{C}$ isn't allowed any advantage, but $t > s$.

The lower bounds that we currently know need the LHS to have both advantages simultaneously: $\mathcal{C}$ uses non-determinism, and $t > s$. Moreover, we need to allow $\mathcal{C}$ to use either randomness, or more alterations. Specifically, for any $s \ll 2^n/n$:

$\mathcal{ZPTIME}[t]^{\mathcal{NP}} \not\subset \mathcal{SIZE}[s]$	where $t = \text{poly}(s)$	[CHR23; Li23]
$\mathcal{MATIME}[t]/1 \not\subset \mathcal{SIZE}[s]$	where $t = \text{poly}(s(\text{poly}(s)))$	[San09]
$\mathcal{AMTIME}[t]/2^{n^\epsilon} \not\subset \mathcal{SIZE}[s]$	where $t = 2^{\text{polylog}(s)}$, $\forall \epsilon > 0$	[CLL25]
$\mathcal{P} \not\subset \mathcal{SIZE}[4.99 \cdot n]$		[IM02][2]

It is currently open to prove that $\mathcal{EXP}^{\mathcal{NP}} \not\subset \mathcal{SIZE}[\Omega(2^n/n)]$, and even to prove that $\mathcal{EXP}^{\mathcal{NP}} \not\subset \mathcal{P}/\text{poly} \stackrel{\text{def}}{=} \cup_c \mathcal{SIZE}[n^c]$. Actually, it is even open to prove that $\mathcal{EXP}^{\mathcal{NP}}$ is hard for polysize depth-two circuits with linear threshold gates (!). When allowing advice in the upper bound, it is open to prove that $\mathcal{MATIME}[t]/1 \not\subset \mathcal{SIZE}[s]$ for a small $t \approx \text{poly}(s)$ that does not involve composition $s \circ s$.

In the course we'll see the proofs of the first three lower bounds above. We'll first see an approach that *doesn't* work, and then we'll develop approaches that give the lower bounds above, and that hopefully can lead to results such as $\mathcal{NP} \not\subset \mathcal{SIZE}[n^k]$.

[1] As usual, the scaled-up version implies the scaled-down version, by padding.

[2] For lower bounds that are precise wrt the multiplicative constant, the choice of gate-basis matters. For circuits over the full basis, we know that $\mathcal{P} \not\subset \mathcal{SIZE}[3.099 \cdot n]$ [LY22].

2 Natural properties

> **A word of caution.** There are too many misinterpretations of what is often referred to as "the natural proofs barrier". The first among them is that there is, mathematically, such a thing as "a class of natural proofs" (in truth, only the notion of "a natural property" is well-defined). But misinterpretations go far beyond that, and I've heard people say that this barrier encapsulates "all natural lower bound techniques", and even go so far as to mention philosophy and Godel's incompleteness theorem (!).
>
> I warn against over-interpreting this result. It is a barrier result, not unlike many others in TCS, and we already know ways around it. In fact, at the time it was published, ways around it were already well-known.[3]
>
> In retrospect, it seems that this result caused undue demoralization in the TCS community. The lesson I personally take from this is to try and avoid groupthink. If the scientific giants working on TCS at the time were susceptible to it, then certainly so am I.

We will study a barrier result asserting that, under a cryptographic assumption, certain techniques cannot prove lower bounds against strong enough circuit classes, in particular against $\mathcal{P}/\text{poly}$.

Intuition and setup. The main technical observation underlying the result is that many lower bounds proofs $f \notin \mathcal{C}$ from the 1980's, all of which had a "combinatorial" flavor, had common interesting features:

- The hard function f had some property Π that no function in $\mathcal{C}$ had.
- A random function has the property Π, whp.
- We can recognize Π efficiently, given the truth-table of a function.

This leads to the following definition:

Definition 1. *A natural property* $\Pi = \left\{\Pi_n \subseteq \{0,1\}^{2^n}\right\}$ *is a set of truth-tables such that*

- *(dense.)* $\Pr_{f_n}[f_n \in \Pi_n] \geq 1/2$.
- *(recognizable.)* $\Pi \in \mathcal{P}$.

We say that Π *is hard for* $\mathcal{C}$ *if no* $\mathcal{C}$*-circuit can compute a truth-table in* Π.[4]

[3] For example, Kannan's classical win-win theorem [Kan82] and the lower bound $\mathcal{MAEXP} \not\subset \mathcal{P}/\text{poly}$ [BFT98] (these proofs "bypass the barrier" by virtue of existing). In the early 2000's, a direct precursor to the algorithmic method was published [IKW02], but still it wasn't widely understood as charting a path to bypass the barrier more generally.

[4] I'm using non-standard terminology. The standard terminology talks of "largeness", "constructiveness", and "usefulness", respectively.

Note that Theorem 1 only defines a natural *property*. The notion of a "natural proof" is not well-defined, and when people say that a proof is a natural proof, what they mean is that they can use the ideas in the proof to define a natural property.

Claim 2. *For any class $\mathcal{C}$ of circuits of size $o(2^n/n)$ there is a dense property hard for $\mathcal{C}$, and if $\mathcal{NEXP} \not\subset \mathcal{C}$ then there is a recognizable property hard for $\mathcal{C}$.*

Proof. The property that contains all truth-tables with circuit complexity at least $\epsilon \cdot 2^n/n$ is dense (for a small enough $\epsilon > 0$) and hard for $\mathcal{C}$.

Any proof of $f \notin \mathcal{C}$ yields a property Π hard for f, namely the property $\Pi = \{\Pi_n = \{f_n\}\}_{n\in\mathbb{N}}$. If $f \in \mathcal{E}$ then Π is recognizable. As proved in [Wil16], if $f \in \mathcal{NEXP}$ there is a recognizable property hard for $\mathcal{C}$ [Wil16]. ■

To be natural in the sense of Definition 1, a property needs to be simultaneously dense and recognizable. It was proved in [RR97] that for many classes against which lower bounds were proved in the 1980's using arguments of a combinatorial flavor, there is a natural property hard for this class (and usually, the definitions of the known natural properties are inspired by the known proofs).

Example 3. *A sequence of works [FSS84; Yao85; Hås87] proved that constant-depth circuits of polysize (even of sub-exponential size) cannot compute the parity of their inputs. They did so by proving a structural result: constant-depth circuits of size s are constant on most subcubes of dimension $n/\text{polylog}(s)$. For concreteness let's think of $s = \text{poly}(n)$.*

We can define a corresponding natural property Π, which is "the function is not constant on any subcube of dimension $n/\text{polylog}(n)$". Note that a random function has the property, whp, and that we can recognize it in time $2^{O(n)}$, by going over all subcubes.

Result statement. Razborov and Rudich proved that, under cryptographic assumptions, natural properties hard against general circuits do not exist.

Theorem 4 ([RR97])**.** *If there are sub-exponentially secure OWFs, then there is no natural property hard for $\mathcal{P}/\text{poly}$. Moreover, for any "reasonable" class $\mathcal{C}$, if there are sub-exponentially secure PRFs computable in $\mathcal{C}$, then there is no natural property hard for $\mathcal{C}$.*

The meaning of Theorem 4 is that proof approaches that yield a natural property against the target class of circuits cannot prove lower bounds against $\mathcal{P}/\text{poly}$, and in fact against any class that can compute a PRF. Note that there are candidate PRFs in very weak classes, e.g. in $\mathcal{TC}^0$ [NR04].

A barrier against what, exactly? There is no mathematical definition for the class of techniques ruled out by Theorem 4. Nevertheless, some techniques do yield a natural property in an obvious way (e.g. restriction-based methods as in [FSS84; Yao85; Hås87]) whereas other techniques do not look like they yield any natural property. Needless to say, known lower bound proofs against $\mathcal{P}/\text{poly}$ belong to the latter category (e.g. counting-based arguments as in the size hierarchy theorem).

A common intuition is to identify the class of techniques that are ruled out with the class of "low-level combinatorial techniques". The latter class is also not well-defined.

2.1 Proof of Theorem 4

Definition 5. *Let $\lambda\colon \mathbb{N} \to \mathbb{N}$. A pseudorandom function is a keyed collection $f = \{f_n\}$ where $f_n = \{f_{n,k}\colon \{0,1\}^n \to \{0,1\}\}_{k\in\{0,1\}^{\lambda(n)}}$ such that for every probabilistic algorithm A running in time $\mathrm{poly}(\lambda(n))$ it holds that*

$$\left| \Pr_{k,A}[A^{f_{n,k}}(1^n) = 1] - \Pr_{f,A}[A^{f}(1^n) = 1] \right| \leq \lambda^{-\omega(1)}$$

and such that the mapping $(x,k) \mapsto f_{n,k}(x)$ is computable in polynomial time. We say that A is sub-exponentially secure if we replace "$\mathrm{poly}(\lambda)$" with "2^{λ^ϵ} for some $\epsilon > 0$".

Theorem 6 ([HIL+99; GGM86])**.** *If there are sub-exponentially secure OWFs, then there are sub-exponentially secure PRFs.*

Assuming towards a contradiction that a natural property exists, we are going to break any PRF. To do so, note that a random function has the property Π, but *no* $\mathcal{P}/\mathrm{poly}$-function has Π, and in particular truth-tables of the PRF do not have Π. Relying on the recognizability of the property, to distinguish between the PRF and a random function, we will test whether or not the given function has Π.

Specifically, we can assume wlog that $\lambda = \mathrm{poly}(n)$ for an arbitrarily large polynomial.[5] We are given 1^n and an oracle to f, which may be either random or $f_{n,k}$ for some key k. We compute the entire truth-table of f (by querying f on all n-bit inputs), and feed it into the recognizability algorithm D of the property Π. Note that this algorithm runs in time $2^{O(n)} \leq 2^{\lambda^\epsilon}$.

When given a random function f, wp at least $1/2$ we will accept (because Π is dense). On the other hand, for every $f_{n,k}$, the truth-table of f is computable by a polynomial-sized circuit $C_k\colon \{0,1\}^n \to \{0,1\}$ (i.e., $C_k(x) = f_{n,k}(x)$); note the use of non-uniformity when fixing k and hard-wiring it into C_k. Hence, since Π is hard for $\mathcal{P}/\mathrm{poly}$, we will reject with probability 1.

Remark 7. *The assumption that sub-exponentially secure OWFs exist is milder than many assumptions being made in cryptography these days, but it is still not a "minimal" assumption in cryptography, because sub-exponential security isn't trivial. To demonstrate this, note that there are polynomially secure PRFs computable in size $(2 + o(1)) \cdot n$ [FLY22] (if PRFs exist at all), but there are no sub-exponentially secure PRFs computable in such size, even in size $4.99 \cdot n$, because there is a natural property hard for such circuits (from [IM02]).*

2.2 Natural properties as algorithms

Even if one does not care about barriers for circuit lower bounds, natural properties are still interesting. A useful perspective is to think of natural properties as *algorithms*, where the algorithm is the "recognizability" one. Indeed, this is an algorithm that

[5] Typical proofs in cryptography actually yield the parameter value $\lambda(n) = n$, in which case we can consider a truncated function $f'(x_1,...,x_{n'}) = f(x_1,...,x_{n'}0...0)$.

distinguishes between easy truth-tables and hard truth-tables, while perhaps making a small number of errors (i.e., mistakenly labelling a small number of hard truth-tables as easy). In other words, this algorithm solves a gap version of the Minimum Circuit Size Problem (MCSP) on average, without erring on easy truth-tables.

By the proof of Theorem 4, such an algorithm does not exist (under cryptographic assumptions). However, it can be useful to think of such an algorithm under various assumptions that we make while trying to head towards a contradiction. Moreover, it is useful to relax the upper-bound on this algorithm, and allow recognizability in (say) quasipolynomial time, or in non-deterministic polytime, and so on.

References

[BFT98] Harry Buhrman, Lance Fortnow, and Thomas Thierauf. "Nonrelativizing separations". In: *Proc. 13th Annual IEEE Conference on Computational Complexity (CCC)*. 1998, pp. 8–12.

[CHR23] Lijie Chen, Shuichi Hirahara, and Hanlin Ren. "Symmetric Exponential Time Requires Near-Maximum Circuit Size". In: *CoRR* abs/2309.12912 (2023). DOI: 10.48550/ARXIV.2309.12912. arXiv: 2309.12912. URL: https://doi.org/10.48550/arXiv.2309.12912.

[CLL25] Lijie Chen, Jiatu Li, and Jingxun Liang. "Maximum circuit lower bounds for exponential-time Arthur Merlin". In: *Proc. 57th Annual ACM Symposium on Theory of Computing (STOC)*. [2025] ©2025, pp. 1348–1358.

[FLY22] Zhiyuan Fan, Jiatu Li, and Tianqi Yang. "The exact complexity of pseudorandom functions and the black-box natural proof barrier for bootstrapping results in computational complexity". In: *Proc. 54th Annual ACM Symposium on Theory of Computing (STOC)*. 2022, pp. 962–975.

[FSS84] Merrick Furst, James B. Saxe, and Michael Sipser. "Parity, circuits, and the polynomial-time hierarchy". In: *Mathematical Systems Theory* 17.1 (1984), pp. 13–27.

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. "How to construct random functions". In: *J. ACM* 33.4 (1986), pp. 792–807.

[Hås87] Johan Håstad. *Computational Limitations of Small-depth Circuits*. MIT Press, 1987.

[HIL+99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. "A Pseudorandom Generator from any One-way Function". In: *SIAM Journal on Computing* 28.4 (1999), pp. 1364–1396.

[IKW02] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. "In search of an easy witness: exponential time vs. probabilistic polynomial time". In: *Journal of Computer and System Sciences* 65.4 (2002), pp. 672–694.

[IM02] Kazuo Iwama and Hiroki Morizumi. "An explicit lower bound of $5n - o(n)$ for Boolean circuits". In: *Proc. International Symposium on Mathematical Foundations of Computer Science*. 2002, pp. 353–364.

[Kan82] R. Kannan. "Circuit-size lower bounds and non-reducibility to sparse sets". In: *Information and Control* 55.1-3 (1982), pp. 40–56.

[Li23] Zeyong Li. "Symmetric Exponential Time Requires Near-Maximum Circuit Size: Simplified, Truly Uniform". In: *Electron. Colloquium Comput. Complex.* TR23-156 (2023). ECCC: TR23-156. URL: https://eccc.weizmann.ac.il/report/2023/156.

[LY22] Jiatu Li and Tianqi Yang. "$3.1n - o(n)$ circuit lower bounds for explicit functions". In: *STOC*. ACM, 2022, pp. 1180–1193. DOI: 10.1145/3519935.3519976.

[NR04] Moni Naor and Omer Reingold. "Number-theoretic constructions of efficient pseudo-random functions". In: *Journal of the ACM* 51.2 (2004), pp. 231–262.

[RR97] Alexander A. Razborov and Steven Rudich. "Natural proofs". In: *Journal of Computer and System Sciences* 55.1, part 1 (1997), pp. 24–35.

[San09] Rahul Santhanam. "Circuit lower bounds for Merlin-Arthur classes". In: *SIAM Journal on Computing* 39.3 (2009), pp. 1038–1061.

[Wil16] R. Ryan Williams. "Natural Proofs versus Derandomization". In: *SIAM Journal on Computing* 45.2 (2016), pp. 497–529.

[Yao85] Andrew C-C. Yao. "Separating the Polynomial-time Hierarchy by Oracles". In: *Proc. 26th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 1985, pp. 1–10.