

A Logic for Expressing Log-Precision Transformers

William Merrill
New York University
willm@nyu.edu

Ashish Sabharwal
Allen Institute for AI
ashishs@allenai.org

Abstract

One way to interpret the reasoning power of transformer-based language models is to describe the types of logical rules they can resolve over some input text. Recently, Chiang et al. (2023) showed that finite-precision transformer classifiers can be equivalently expressed in a generalization of first-order logic. However, finite-precision transformers are a weak transformer variant because, as we show, a single head can only attend to a constant number of tokens and, in particular, cannot represent uniform attention. Since attending broadly is a core capability for transformers, we ask whether a minimally more expressive model that can attend universally can also be characterized in logic. To this end, we analyze transformers whose forward pass is computed in $\log n$ precision on contexts of length n . We prove any log-precision transformer classifier can be equivalently expressed as a first-order logic sentence that, in addition to standard universal and existential quantifiers, may also contain majority-vote quantifiers. This is the tightest known upper bound and first logical characterization of log-precision transformers.

Any log-precision transformer can be re-expressed as a sentence in $\text{FO}(\mathbf{M})$ logic, e.g.:

$$\text{Mi. } \mathbf{a}(i) \wedge \text{Mj. } \mathbf{b}(j) \wedge \neg \exists k, \ell. (\mathbf{a}(k) \wedge \mathbf{b}(\ell) \wedge \ell < k)$$

(m a's followed by m b's, i.e., $a^m b^m$)

aaaabbbb ✓ aaabbbbb ✗ baaaabbb ✗

Figure 1: A first-order logic with majority ($\text{FO}(\mathbf{M})$) sentence for $a^m b^m$. In addition to standard \forall and \exists quantifiers over string indices, $\text{FO}(\mathbf{M})$ allows *majority* quantifiers (\mathbf{M}) that take a majority-vote across indices. $\mathbf{a}(i)$ indicates whether token i is a (and analogously for b). We prove $\text{FO}(\mathbf{M})$ can express any function computed by a log-precision transformer.

1 Introduction

The incredible success of deep learning models, especially very large language and vision transformers with hundreds of billions of parameters (Brown et al., 2020; Thoppilan et al., 2022), has come at the cost of increasingly limited understanding of how these models actually work and when they might fail. This raises many concerns, such as around their safe deployment, fairness, and accountability. Does the inner working of a transformer defy description in a simpler symbolic system that we can better understand? Or *can transformer computation be described using a familiar symbolic formalism?* Understanding how to view the reasoning process of a transformer in terms of logic could potentially expand our ability to formally reason about their behavior over large domains of inputs.

Chiang et al. (2023) provide a partial answer to this question, showing that any *finite-precision* transformer classifier can be expressed as a sentence in a variant of first-order logic with counting

quantifiers and modular arithmetic over input position indices. Specifically, counting quantifiers take the form $\exists^{=x} i : \phi(i)$ where x is a count variable and i is a position index. They show that there exists a single sentence in this logic that computes the output of the transformer for any input string of any length. This is a powerful result because it shows that a simple logical formalism is fully sufficient to describe all the complexity of a massive finite-precision transformer. It also provides an upper bound on finite-precision transformers: any function that cannot be defined in first-order counting logic with modular indexing cannot be expressed by the transformer.

However, Chiang et al.’s result is not fully general because it relies on the transformer precision being fixed with respect to the transformer’s context length. More generally, as we will demonstrate in Section 3, finite-precision transformers are a fundamentally weak variant of transformers: crucially, cannot express uniform attention patterns, which are a core algorithmic primitive of transformers (Weiss et al., 2018). In fact, we show that they can only attend to a constant number of input positions, which may be seen as a rather limited generalization of hard attention.¹ For example, Chiang et al. show that their logic for finite-precision transformers cannot recognize $a^m b^m$, whereas in practice, transformers can (Bhattamishra et al., 2020).² This motivates studying a formal model of transformers where precision grows with context length (which we formalize as *log-precision*), making it possible to capture uniform attention as well as other broad attention patterns. This is useful both for recognizing $a^m b^m$ and more generally for reasoning globally over the input.

We demonstrate that *log-precision* transformer classifiers can also be expressed as sentences in a simple logic: *first-order logic with majority*, or $\text{FO}(\text{M})$, over inputs strings (Barrington et al., 1990). In addition to standard existential and universal quantifiers, $\text{FO}(\text{M})$ has *majority* quantifiers that return true iff at least half the propositions they quantify are true. It also allows comparing input positions (e.g., $\ell < k$ in Figure 1) and accessing their individual bits. Our main result is as follows:

Theorem 1 (Informal version of Theorem 2). *For any log-precision transformer \mathcal{T} , there exists an $\text{FO}(\text{M})$ sentence ϕ that computes the same function as \mathcal{T} , i.e., $\phi(x) = \mathcal{T}(x)$ for any input string x .*

Upper bound. Theorem 2 shows transformers with more than finite precision can also be expressed in a simple extension of first-order logic, going beyond Chiang et al. (2023)’s result. On the other hand, $\text{FO}(\text{M})$ is a strict superset of Chiang et al.’s counting logic; it can simulate counting quantifiers (see Section 2.2) and allows non-modular position comparisons. Thus, handling a more general class of transformers powerful enough to express uniform attention slightly weakens the bound.

Still, our result constitutes (to our knowledge) the tightest upper bound on log-precision transformers and the first defined in terms of logic, building on a line of complexity-theoretic work analyzing the power of transformers (Hahn, 2020; Merrill et al., 2022; Liu et al., 2023; Merrill & Sabharwal, 2023). In particular, $\text{FO}(\text{M})$ strengthens the upper bound of log-space-uniform TC^0 by Merrill & Sabharwal (2023). The refined bound adds to the limitations of transformers identified by Merrill & Sabharwal (2023): for example, it establishes unconditionally that log-precision transformers cannot compute boolean matrix permanents, and shows that, in a certain formal sense, integer division and matching parentheses are among the formally hardest problems that transformers can solve (see Section 4).³

Mechanistic interpretability. Beyond providing an upper bound on the reasoning problems solvable by transformers, we believe Theorem 1 could guide the design of “transformer-complete” programming languages similar in spirit to RASP (Weiss et al., 2018). RASP is a declarative programming language designed to capture transformer computation, and Lindner et al. (2023) implement a compiler from RASP *into* transformers. Unlike RASP, $\text{FO}(\text{M})$ can provably express any transformer (Theorem 1), which we believe justifies using it (or an equivalent but more user-friendly variant) as a target language for programs extracted *from* transformers.

Similar to a decision tree, an $\text{FO}(\text{M})$ sentence has the interpretable property that each sub-sentence corresponds to a constraint on input (see Figure 1). In contrast, the internal modules of a transformer or circuit do not satisfy this since they map between arbitrary latent spaces. We speculate this property

¹Hard attention is provably substantially weaker than general attention (Hao et al., 2022; Merrill et al., 2022).

²Technically, the empirical results of Bhattamishra et al. (2020) are for $a^m b^m c^m$, a harder variant of $a^m b^m$.

³To be clear, Theorem 1 is one-sided: every transformer can be expressed as an $\text{FO}(\text{M})$ sentence, but not necessarily the other way. Moreover, we believe that many $\text{FO}(\text{M})$ sentences *cannot* be expressed by transformers. An exact logical characterization of transformers remains an open problem.

could facilitate interpreting models by translating them to $\text{FO}(\mathcal{M})$, though a careful exploration of the algorithmic and HCI aspects of this idea lies outside the current paper’s theoretical scope.

Contributions. Our results shed new light on how to view the computation inside transformers in terms of logic. Specifically, our main contributions are to prove the following:

1. Fixed-precision transformers can only attend to a fixed number of tokens, and those with precision less than $\log \log n$ cannot uniformly attend over length- n contexts (Proposition 1).
2. Log-precision transformer classifiers can be expressed as sentences in $\text{FO}(\mathcal{M})$ (Theorem 2).

2 Preliminaries: Transformers and $\text{FO}(\mathcal{M})$

Let Σ be a finite alphabet. We denote by * the Kleene star operator, i.e., for a set X , $X^* = \bigcup_{n=0}^{\infty} X^n$. We will view transformers and $\text{FO}(\mathcal{M})$ sentences both as functions from $\Sigma^* \rightarrow \{0, 1\}$, and show that any function a transformer computes can also be computed by an $\text{FO}(\mathcal{M})$ sentence.

2.1 Transformers

We view the transformer precision p as a function of the context length n , writing $p(n)$ where appropriate. Let \mathbb{D}_p be the datatype of p -precision floats, i.e., tuples $\langle m, e \rangle$ where m, e are signed integers together taking p bits. Using $|x|$ to mean the size of integer x , a float represents the value $m \cdot 2^{e-|m|+1}$.⁴ Following Appendix A of Merrill & Sabharwal (2023), we define p -truncated addition ($+$, \sum), multiplication (\cdot), and division ($/$) over \mathbb{D}_p . We now define a *transformer encoder binary classifier* over \mathbb{D}_p , largely adopting Merrill & Sabharwal’s notation.⁵

Definition 1. A p -precision transformer \mathcal{T} with h heads, d layers, model dimension m (divisible by h), and feedforward width w is specified by:

1. An embedding function $\phi : \Sigma \times \mathbb{N} \rightarrow \mathbb{D}_p^m$ whose form is defined in Appendix C.1;⁶
2. For each $1 \leq \ell \leq d$ and $1 \leq k \leq h$, a head similarity function $s_k^\ell : \mathbb{D}_p^m \times \mathbb{D}_p^m \rightarrow \mathbb{D}_p$ whose form is defined in Appendix C.2;
3. For each $1 \leq \ell \leq d$ and $1 \leq k \leq h$, a head value function $v_k^\ell : \mathbb{D}_p^m \rightarrow \mathbb{D}_p^{m/h}$ whose form is defined in Appendix C.2;
4. For each $1 \leq \ell \leq d$, an activation function $f^\ell : (\mathbb{D}_p^{m/h})^h \times \mathbb{D}_p^m \rightarrow \mathbb{D}_p^m$ whose form is defined in Appendix C.3 and implicitly uses the feedforward dimension w ;
5. An output classifier head $\kappa : \mathbb{D}_p^m \rightarrow \{0, 1\}$ whose form is defined in Appendix C.4.

Definition 2. We define the transformer computation and output as a function of an input $x \in \Sigma^n$.

1. Embeddings: For $1 \leq i \leq n$, $\mathbf{h}_i^0 = \phi(x_i, i)$.⁶
2. Self Attention: For $0 \leq \ell \leq d-1$, (multihead) self-attention block $\ell+1$ computes h attention heads:

$$\mathbf{a}_{i,k}^{\ell+1} = \sum_{j=1}^n \frac{s_k^{\ell+1}(\mathbf{h}_i^\ell, \mathbf{h}_j^\ell)}{Z_{i,k}} \cdot v_k^{\ell+1}(\mathbf{h}_j^\ell), \quad \text{where } Z_{i,k} = \sum_{j=1}^n s_k^{\ell+1}(\mathbf{h}_i^\ell, \mathbf{h}_j^\ell).$$

3. Activation Block: For $0 \leq \ell \leq d-1$, activation block $\ell+1$ aggregates the head outputs to produce $\mathbf{h}^{\ell+1}$:

$$\mathbf{h}_i^{\ell+1} = f^{\ell+1}(\mathbf{a}_{i,1}^{\ell+1}, \dots, \mathbf{a}_{i,h}^{\ell+1}, \mathbf{h}_i^\ell).$$

4. Classifier Head: The network prediction on $x \in \Sigma^n$ is $\kappa(\mathbf{h}_n^d)$.

⁴ $\langle 101, 010 \rangle$ represents $1.01_2 \times 2^{102}$. This is closer to the IEEE standard than the $m \cdot 2^e$ semantics used in Merrill & Sabharwal (2023), letting us define the minimum representable float more realistically in Proposition 1.

⁵Increasing the classifier’s output space arity (e.g., a transformer that predicts the next token) or switching to causal attention of a decoder-only model would not change our results. However, our proof no longer goes through if the decoder can generate tokens that get added to the input at the next step (cf. Pérez et al., 2019).

⁶ ϕ , like p , is actually a function of the context length n , and Appendix C.1 enforces that ϕ is computable in $O(\log n)$ time, as standard choices of positional embeddings would satisfy.

We say $\mathcal{T}(x) = \kappa(\mathbf{h}_{|x|}^d)$ and $L_{\mathcal{T}}$ is the language of $x \in \Sigma^*$ such that $\mathcal{T}(x) = 1$. We refer to $\phi, s_k^\ell, v_h^\ell, f^\ell$, and κ as the **core functions** in \mathcal{T} , and to embeddings, self attention, activation, and the classifier head as the **components** of \mathcal{T} . We write $\theta_{\mathcal{T}}$ for the concatenated vector of parameters for the functions $\phi, s_k^\ell, v_h^\ell, f^\ell$, and κ , for all $1 \leq \ell \leq d$ and $1 \leq k \leq h$.

We define a **log-precision transformer** as one where p is at most $O(\log n)$ and is a “simple” function, i.e., computable in $O(\log n)$ time. In our model, the weights $\theta_{\mathcal{T}}$ defining \mathcal{T} are fixed, but the precision p used to compute the forward pass can depend on n (see Footnote 13 for a generalization).

2.2 First-Order Logic with Majority

As we will show, transformers can be translated into sentences in $\text{FO}(\mathcal{M})$. But what do such sentences look like? Informally, $\text{FO}(\mathcal{M})$ is first-order logic extended to also have majority (\mathcal{M}) quantifiers. Following Barrington et al. (1990), our sense of $\text{FO}(\mathcal{M})$ takes *strings* in Σ^* as input and returns 0 or 1 to define a formal language. In this setting, quantifiers range over *indices* (positions) into the string. Predicates can be applied to the variables introduced by these quantifiers.

Definition 3 ($\text{FO}(\mathcal{M})$ index). Indices in $\text{FO}(\mathcal{M})$ are integers denoting positions in the input string:

1. The constant 1, representing the first token’s position.
2. The constant n , representing the last token’s position.
3. Strings (e.g., i, j, k) representing variables ranging over positions 1 to n .
4. Any index built by applying addition or subtraction to other indices.⁷

Definition 4 ($\text{FO}(\mathcal{M})$ formula). Formulas in $\text{FO}(\mathcal{M})$ are constructed as follows:⁸

1. Let Σ be a finite alphabet. For each $\sigma \in \Sigma$ and any index i , $\sigma(i)$, e.g., $\mathbf{a}(i)$, is a formula that is true if the i -th input token is σ .⁹
2. For any indices i, j , the formula $\text{bit}(i, j)$ returns the j -th bit of the binary expansion of i .¹⁰
3. For two indices i, j , $i = j$, $i \leq j$, and $i \geq j$ are formulas with their conventional semantics.
4. For two formulas ϕ, ψ , $\phi \wedge \psi$ and $\phi \vee \psi$ are formulas with their conventional semantics.
5. For any formula ϕ (which may refer to i), the following are valid formulas:
 - (a) $\exists i. \phi$ means some value of i in $[1, n]$ makes ϕ true.
 - (b) $\forall i. \phi$ means all values of i in $[1, n]$ make ϕ true.
 - (c) $Mi. \phi$ means $\geq n/2$ values of i in $[1, n]$ make ϕ true.

We use parentheses where necessary to disambiguate the order of operations. General formulas may contain free (i.e., unbound) variables: e.g., $\forall i. i = j$. A *sentence* is an $\text{FO}(\mathcal{M})$ formula ϕ with no free variables. Sentences represent functions from Σ^* to $\{0, 1\}$ and thus define a formal language.¹¹

Extensions. Beyond Definition 4, $\text{FO}(\mathcal{M})$ can express *counting* and *threshold* quantifiers in terms of majority quantifiers (Barrington et al., 1990). Given a formula ϕ , a counting quantifier creates a new formula $\exists^k i : \phi$ that is true iff ϕ is true across exactly k values of i . Threshold quantifiers $\exists^{\leq k}$ and $\exists^{\geq k}$ work similarly but check if ϕ is true for at least or at most k values of i . In addition, we show in Appendix A that $\text{FO}(\mathcal{M})$ can express *conditional majority* quantifiers, which create a formula $Mi : \phi [\psi]$ that is true iff ψ is true for at least half the values of i that make ϕ true.

2.2.1 Examples

To illustrate the formalism, we provide example languages definable in $\text{FO}(\mathcal{M})$ with $\Sigma = \{\mathbf{a}, \mathbf{b}\}$. First, we show two languages that do not require majority quantifiers to express:

Example 1 (Bigram matching). Strings containing the bigram \mathbf{ab} : $\exists i [\mathbf{a}(i) \wedge \mathbf{b}(i+1)]$.

Example 2 (Skip-bigram matching). Strings containing the long-distance pattern $\mathbf{a} \dots \mathbf{b}$ (cf. “induction heads” of Elhage et al. 2021): $\exists i [\mathbf{b}(i) \wedge \exists j [j \leq i \wedge \mathbf{a}(j)]]$.

⁷Barrington et al. (1990) did not introduce this as a primitive, but it can be simulated using the \leq predicate.

⁸We write parentheses to indicate the order of operations.

⁹Barrington et al. (1990) define $Q_b(i)$ for $b \in \{0, 1\}$. We generalize this to an arbitrary vocabulary Σ by assuming each token is one-hot-encoded: $\sigma(i) = Q_1(|\Sigma|i + s)$ where s is the index of σ in the vocabulary.

¹⁰This predicate is included in the logic for technical reasons; see Barrington et al. (1990).

¹¹One can also take multiple sub-sentences within ϕ to be labeled as ordered outputs, thus allowing ϕ to be a function from Σ^* to $\{0, 1\}^k$ for some fixed constant k .

In contrast, Example 3 is a simple example that requires majority quantifiers (Furst et al., 1984):

Example 3 (Majority). Strings with more b’s than a’s: $M_i [b(i)]$.

Figure 1 showed how $\text{FO}(\mathcal{M})$ can be used to recognize patterns like $a^m b^m$. A similar idea can be used to model parentheses matching (Barrington et al., 1990):

Example 4 (1-Dyck). The well-balanced parentheses language (with a opening and b closing):

$$\forall i. (\exists a, b. ((\exists^a j : a(j) \wedge j \leq i) \wedge (\exists^b j : b(j) \wedge j \leq i) \wedge b \leq a)) \wedge M_i. a(i) \wedge M_j. b(j).$$

Example 5 (Integer Arithmetic). Iterated addition (i.e., summing n n -bit numbers), iterated multiplication, and division (Hesse, 2001) can all be expressed in $\text{FO}(\mathcal{M})$.

3 Finite Precision Transformers Cannot Attend Universally

Attention heads that spread attention weight uniformly across inputs have been observed in transformer LMs (Merrill et al., 2021) and make soft attention fundamentally more powerful than hard attention (Hao et al., 2022; Merrill et al., 2022). In particular, uniform attention is an important primitive that transformers can use to solve tasks involving counting (Bhattamishra et al., 2020; Chiang et al., 2023), taking majority votes (Merrill et al., 2022), and matching parentheses or sorting (Weiss et al., 2021). A transformer with sufficient precision can easily implement uniform attention by setting the keys and queries across all positions to be constant. However, attention heads with finite precision cannot represent uniform attention over long sequences as a consequence of the following:

Proposition 1. Let $\mathbf{a} \in \mathbb{R}^n$ s.t. $\sum_{i=1}^n a_i = 1$ and $\tilde{\mathbf{a}}$ its nearest p -precision float approximation.

1. Then the number of nonzero entries of $\tilde{\mathbf{a}}$ is upper bounded by its precision: specifically, $\tilde{\mathbf{a}}$ has at most 2^{2^p} nonzero entries.
2. Moreover, if $p < \log \log n$ and \mathbf{a} is uniform (i.e., $a_i = 1/n$), then $\tilde{\mathbf{a}} = \vec{0}$.

Proof. The smallest positive value representable by a p -precision float is $2^{-(p_m - 2 + 2^{p_e - 1})}$ which is bounded below by $2^{-2^p + 1}$. Letting $k = 2^{2^p}$, it holds that $2^{-2^p + 1} = 2/k$. So if \tilde{a}_i gets the minimum value, then $a_i \geq 1/k$. Since $\sum_i a_i = 1$, there can be at most k indices satisfying this property. This implies there can be at most k nonzero entries in $\tilde{\mathbf{a}}$. If $n > k$ and \mathbf{a} is uniform, $1/n$ is less than half of the minimum representable value of $2/k$. Thus, $\tilde{\mathbf{a}} = \vec{0}$. \square

Proposition 1 says that fixed-precision transformers are artificially limited because they can only attend over bounded-length windows, making them similar to hard-attention transformers (Hao et al., 2022). Moreover, they cannot compute uniform attention over contexts of length n with less than $\log \log n$ precision. This explains why Chiang et al. (2023) prove finite-precision transformers provably cannot recognize $a^m b^m$, while in practice transformers have been shown to learn even its harder variant $a^m b^m c^m$ even with long context lengths (Bhattamishra et al., 2020). In essence, their upper bound only applies in the asymptotic regime when $n > 2^{2^p}$.

In contrast, transformers in practice have enough precision both to compute uniform attention and recognize $a^m b^m$ on practical context lengths. More concretely, the bfloat16 representation allows uniform attention over $2^{6+2^7} \approx 10^{42}$ tokens and normal float16¹² allows $2^{10+2^4} \approx 10^8$ tokens, both well above the typical context window of transformers. This motivates a formal model of transformers with enough precision to compute uniform attention and recognize languages such as $a^m b^m$.

4 Main Result: Expressing Log-Precision Transformers in $\text{FO}(\mathcal{M})$

By Proposition 1, precision must grow with the context length n ($p > \log \log n$) for a transformer to compute uniform attention and other attention patterns with unbounded range, like practical transformers. In this paper, we analyze any transformer with up to $O(\log n)$ precision. We show that any function computable by log-precision transformers can be expressed in $\text{FO}(\mathcal{M})$:

¹²We account for the division of p into p_m and p_e rather than treating them together. Our minimum value differs slightly from numpy but is on the same order of magnitude. Moving to float8 lowers the length upper bound for uniform attention to $2^{3+2^3} \approx 2048$, which suggests float8 LMs will have limited length generalization.

Theorem 2. Let \mathcal{T} be a log-precision transformer with a parameter vector $\theta_{\mathcal{T}}$ fixed for all context lengths n .¹³ Then, there exists an $\text{FO}(\mathbf{M})$ sentence ϕ that computes the same function as \mathcal{T} , i.e., $\phi(x) = \mathcal{T}(x)$ for any input string x .

Theorem 2 is the tightest known upper bound for log-precision transformers and shows that it is still possible to characterize transformers in a simple variant of first-order logic even with log-precision and uniform attention. As alluded to earlier, Theorem 2 immediately implies that any problem complete for $\text{FO}(\mathbf{M})$ (or a larger class) is also transformer-hard. Since integer division and Dyck language membership are known to be $\text{FO}(\mathbf{M})$ -complete (Hesse, 2001; Aaronson et al., 2022), it follows, perhaps surprisingly, that the entire computation of any transformer on input x can be reduced to a single integer division or a finite number of Dyck-language queries:

Corollary 2.1. Let \mathcal{T} be a transformer satisfying Theorem 2. For any input x , there exist first-order definable integers a, b , and i (dependent on \mathcal{T} and x) such that $\mathcal{T}(x)$ equals the i -th bit of $[a/b]$. For any x , there also exist first-order definable strings w_1, \dots, w_m such that $\mathcal{T}(x)$ is first-order definable in terms of the membership of the w_i 's in k -Dyck.

5 Preliminaries for Proving Theorem 2

5.1 Computation Graphs

A *computation graph* G over a datatype $\mathbb{D} \subseteq \{0, 1\}^*$ and a countable set of primitive functions $\mathfrak{F} \subseteq \mathbb{D}^* \times \mathbb{D}$ is a directed acyclic graph where:

1. Each node is labelled by a *node type*: a function $f \in \mathfrak{F}$ computed by this node.
2. Each edge represents a value \mathbb{D} flowing as output from one node into another node. We consider the edges flowing into node j to have an order, i.e., be numbered.
3. \mathfrak{F} contains the special symbol `input`, which designates k nodes as input nodes. We refer to k as the *arity* and assume w.l.o.g. that nodes $0, \dots, k-1$ are inputs.¹⁴
4. A single node is taken as the output node (w.l.o.g., the node with the largest index).

A computation graph G of arity k parameterizes a function $\mathbb{D}^k \rightarrow \mathbb{D}$ in the standard way: the input nodes are assigned the input values, and the value of each node is computed (traversing the graph in a bottom-up topological order) as a function of the values of its children until the output node receives a value. The value of the output node is considered the output of the function. It is worth noting that computation graphs can only process inputs of bounded length. To process arbitrary-length inputs, we will need to generalize them to computation graph families (Section 5.2).

For a computation graph G , $\text{size}(G)$ is the number of nodes, $\text{depth}(G)$ is the length of the longest path from an input node to the output, and $\text{arity}(G, i)$ is the number of inputs to node i .

Threshold circuits. A threshold circuit is a special case of a computation graph where $\mathbb{D} = \{0, 1\}$ and \mathcal{F} is the set of threshold functions of the form $\theta_{\leq \Delta}$ and $\theta_{\geq \Delta}$ over \mathbb{D}^* , defined as follows: $\theta_{\leq \Delta}(x) = 1$ if $\sum_{\sigma \in x} \sigma \leq \Delta$ and 0 otherwise; $\theta_{\geq \Delta}(x)$ is defined analogously. Typical AND, OR, and NOT gates are a special case of threshold gates, as is an IDENTITY gate.¹⁵

We allow nodes with the $k' \geq 1$ largest indices to all be designated as (ordered) output nodes. A threshold circuit with arity k and k' output nodes will thus be a function from $\{0, 1\}^k$ to $\{0, 1\}^{k'}$. This will be convenient when simulating neural network components that output multiple bits.

We will find it useful to consider threshold circuits as a kind of compilation target for computation graphs: in other words, we will be concerned with simulating computation graphs defined over more complex functions and data types into threshold circuits.

5.2 Computation Graph Families

A computation graph family over \mathbb{D} and \mathfrak{F} is a mapping from $n \in \mathbb{N}$ to a computation graph G_n for processing inputs of size n . Thus, \mathcal{G} defines a function from $\mathbb{D}^* \rightarrow \mathbb{D}$, where $\mathcal{G}(x) = G_{|x|}(x)$.

¹³Theorem 2 can also be extended to apply to log-precision transformers with *log-uniform weights*, i.e., where $\theta_{\mathcal{T}}$ can grow in size and precision with n (see Appendix B).

¹⁴By convention in computer science, we let computation graph nodes be zero-indexed.

¹⁵For more background on threshold circuits, see Merrill & Sabharwal (2023) and Merrill et al. (2022).

Intuitively, computation graph families are useful because they generalize computation graphs to define functions over *unbounded-length* strings as inputs.

Size, depth, and arity. For computation graph families, the size, depth, and arity become functions of the input length n : $\text{size}_{\mathcal{G}}(n) = \text{size}(G_n)$, $\text{depth}_{\mathcal{G}}(n) = \text{depth}(G_n)$, $\text{arity}_{\mathcal{G}}(n, i) = \text{arity}(G_n, i)$.

Uniformity. The infinite set \mathcal{G} can be alternatively represented by two functions:

1. $\text{node}_{\mathcal{G}}(n, i)$, which returns the type of node i in G_n if $i \leq \text{size}(G_n)$, and \emptyset otherwise. For example, if node i computes the logical AND of its inputs, then $\text{node}_{\mathcal{G}}(n, i) = \wedge$.
2. $\text{edge}_{\mathcal{G}}(n, i, j)$, which returns the argument index of i into node j if G_n contains an edge $i \rightarrow j$ and -1 otherwise. $\text{edge}_{\mathcal{G}}(n, i, j)$ only needs to be defined over $i, j < \text{size}(G_n)$. For example, if G_n contains a node j with three incoming edges, the second of which comes from node i , then $\text{edge}_{\mathcal{G}}(n, i, j) = 1$.

A pair of algorithms implementing these two functions uniquely specifies a computation graph family, as it enables building the computation graph G_n for any n . Uniform computation graph families (generalizing uniform circuits; cf. Arora & Barak, 2009) are families where $\text{node}_{\mathcal{G}}$ and $\text{edge}_{\mathcal{G}}$ can be computed efficiently, i.e., under some constraints on space or time:

Definition 5 (Uniformity). A computation graph family \mathcal{G} is $T(n)$ -uniform iff $\text{node}_{\mathcal{G}}(n, i)$ and $\text{edge}_{\mathcal{G}}(n, i, j)$ can be computed by a deterministic Turing machine in time $T(n)$. We focus on *log-uniform* computation graph families: i.e., where $T(n) = O(\log n)$.¹⁶

Threshold circuit families. These are simply families of threshold circuits. We will be simulating computation graph families with threshold circuit families. Log-uniform TC^0 is the class of languages recognized by log-uniform constant-depth, poly-size threshold circuit families. See Merrill & Sabharwal (2023); Liu et al. (2023); Arora & Barak (2009) for more background on TC^0 and circuits.

6 Proof of Theorem 2

The idea is to simulate a transformer with a log-uniform TC^0 circuit family. Since log-uniform $\text{TC}^0 = \text{FO}(\mathbf{M})$, this would imply any transformer can be expressed in $\text{FO}(\mathbf{M})$. First, we note that transformers are log-uniform computation graphs:

Lemma 1 (Proof in Appendix B.1). *A transformer \mathcal{T} is a log-uniform computation graph family where \mathfrak{F} contains embedding, self-attention, feedforward, and output components.*

Further, each core module of the transformer can be simulated by a log-uniform TC^0 circuit family:

Lemma 2 (Proof in Appendix B.2). *Let \mathcal{T} be a log-precision transformer with fixed parameters $\theta_{\mathcal{T}}$. Then each component in \mathfrak{F} is computable in log-uniform TC^0 .*

Intuitively, we can now simulate a transformer in log-uniform TC^0 by just simulating each of its components with a threshold circuit and routing their inputs and outputs appropriately. However, we will need two more technical conditions to verify that this construction is indeed log-uniform:

Lemma 3 (Proof in Appendix B.3). *Let \mathcal{T} be a log-precision transformer with fixed parameters $\theta_{\mathcal{T}}$. There exists a function $\text{bsize}(n)$ that is a power of 2 and computable in $O(\log n)$ time s.t. $\text{size}_{\mathcal{F}}(n) \leq \text{bsize}(n)$ for all $\mathcal{F} \in \mathfrak{F}$.*

Lemma 4 (Proof in Appendix B.4). *If \mathcal{F} is a log-uniform TC^0 family and $\text{size}_{\mathcal{F}}(n) \leq \text{bsize}(n)$, there exists a log-uniform TC^0 family \mathcal{F}' s.t. $\mathcal{F}(x) = \mathcal{F}'(x)$ for all x and $\text{size}_{\mathcal{F}'}(n) = \text{bsize}(n)$.*

Combined, Lemmas 3 and 4 show that each $\mathcal{F} \in \mathfrak{F}$ is computable by a log-uniform TC^0 family with size $\text{bsize}(n)$ that is a power of 2 and computable in time $O(\log n)$. We will show these conditions imply a transformer \mathcal{T} can be simulated by a TC^0 family \mathcal{C} (Theorem 3) and moreover that \mathcal{C} is log-uniform (Corollary 3.2). By the equivalence of log-uniform TC^0 and $\text{FO}(\mathbf{M})$ (Barrington et al., 1990), we then conclude that any log-precision transformer can be expressed in $\text{FO}(\mathbf{M})$.

¹⁶Past work (Merrill & Sabharwal, 2023) analyzes transformers with a similarly named but weaker notion of uniformity, namely *log-space* (rather than *log-time*) uniformity.

Algorithm 1 $\text{node}_{\mathcal{C}}(n, i)$

Return the type of gate i in circuit C_n .

```
1:  $\mathcal{F} \leftarrow \text{node}_{\mathcal{G}}(n, \text{bnode}(n, i))$ 
2: if  $\mathcal{F} \neq \emptyset$  then
3:   return  $\text{node}_{\mathcal{F}}(n, i - \text{bstart}(n, i'))$ 
4: else return  $\emptyset$ 
```

Algorithm 2 $\text{edge}_{\mathcal{C}}(n, i, j)$

If C_n contains an edge $i \rightarrow j$, return the argument number of that edge. Otherwise, return -1 .

```
1:  $i' \leftarrow \text{bnode}(n, i)$ 
2:  $j' \leftarrow \text{bnode}(n, j)$ 
3:  $s_i \leftarrow \text{bstart}(n, i')$ 
4:  $s_j \leftarrow \text{bstart}(n, j')$ 
5: if  $i' = j'$  then
6:    $\mathcal{F} \leftarrow \text{node}_{\mathcal{G}}(n, i')$ 
7:   return  $\text{edge}_{\mathcal{F}}(n, i - s_i, j - s_j)$ 
8: else if  $\text{edge}_{\mathcal{G}}(n, i', j') \geq 0$  then
9:    $b_i \leftarrow i - (s_i + \text{bsize}(n, i') - p(n))$ 
10:   $b_j \leftarrow j - (s_j + p(n) \cdot \text{edge}_{\mathcal{G}}(n, i', j'))$ 
11:  if  $b_i = b_j < p(n)$  then return  $j - s_j$ 
12: else return  $-1$ 
13: else return  $-1$ 
```

6.1 Simulating Computation Graph Families with Circuit Families

We give algorithms that take a computation graph family and define a circuit family simulating it. Intuitively, the algorithms creates contiguous blocks of circuit gates simulating each node in the computation graph and route inputs and outputs between blocks appropriately.

Block mapping. This algorithm depends on a *block mapping*, which is an implementation of the following three functions:

1. The *block node* $\text{bnode}(n, i)$: the index of the node that gate i 's block is simulating.
2. The *block start* $\text{bstart}(n, i')$: the smallest gate index in the block simulating node i' .
3. The *block size* $\text{bsize}(n, i')$: the number of gates in the block simulating node i' .

Further, we enforce that a valid block mapping must satisfy that, for all i , with $i' = \text{bnode}(n, i)$,

$$\text{bstart}(n, i') \leq i < \text{bstart}(n, i') + \text{bsize}(n, i').$$

Let \mathcal{G} be a computation graph whose primitive functions are computable by log-uniform threshold circuits. We can identify each primitive function with a log-uniform threshold circuit family \mathcal{F} that computes it, where the first $\text{arity}_{\mathcal{F}}(n)$ gates are IDENTITY gates reserved for taking input. For such a graph, $\text{node}_{\mathcal{G}}$ can be taken to return a symbol identifying a circuit family \mathcal{F} . In this case, our algorithm requires that, for all i' , the block size of i' must match the size of the circuit for the type of block i' , i.e., $\text{bsize}(n, i') = \text{size}_{\text{node}_{\mathcal{G}}(n, i')}(n)$. These properties let us meaningfully identify a graph node i' with a block of nodes that will simulate it. This intuition enables us to develop Algorithms 1 and 2 for constructing a uniform threshold circuit family from a uniform computation graph family.

Theorem 3. Let \mathcal{G} be a computation graph over a finite set of node types \mathfrak{F} , where each $\mathcal{F} \in \mathfrak{F}$ is specified by a log-uniform circuit family. Let bnode , bstart , and bsize be a valid block mapping in the sense above. Then Algorithms 1 and 2 define a circuit family \mathcal{C} such that

1. \mathcal{C} and \mathcal{G} compute the same $\mathbb{D}_p^* \rightarrow \mathbb{D}_p$ function (let the final p gates of each C_i be its output).
2. $\text{depth}_{\mathcal{C}}(n) \leq \text{depth}_{\mathcal{G}}(n) \cdot \max_{\mathcal{F}} \text{depth}_{\mathcal{F}}(n)$.
3. $\text{size}_{\mathcal{C}}(n) \leq \text{size}_{\mathcal{G}}(n) \cdot \max_{\mathcal{F}} \text{size}_{\mathcal{F}}(n)$.

Proof. Assume w.l.o.g. that the gates of \mathcal{C} are topologically ordered. We show by induction over circuit gates j (with $j' = \text{bnode}(n, j)$) that:

1. For all $i' < j'$, the last p nodes of block i' store the value of node i' .
2. For all i such that $\text{bstart}(n, j') \leq i \leq j$, gate i of \mathcal{C} (as a function of the input nodes of j') computes gate $i - \text{bstart}(n, j')$ of $\text{node}_{\mathcal{G}}(n, j')$.

Base case. We have two circuits with no gates, so the premises are trivially satisfied.

Inductive case. Assume the premises hold up to j . We will show they hold for $j + 1$. Let $\mathcal{T} = \text{node}_{\mathcal{G}}(n, j')$. By Premise 1, we know that the last p nodes of block i' store the output of node i' , for

$i' < j'$. By Algorithm 2, for each i' such that $\text{edge}_{\mathcal{G}}(n, i', j') = a$ with $0 \leq k < \text{arity}_{\mathcal{F}}(n)$, gates kp through $k(p+1) - 1$ of block j' will copy the final p gates of block i' . Thus, the first $k \cdot \text{arity}_{\mathcal{F}}(n)$ gates of block j' store the inputs to node j' .

At this point, we use Premise 2 to conclude that the first $j - \text{bstart}(n, j')$ gates of block j' compute the same function as the first $j - \text{bstart}(n, j')$ gates of \mathcal{F} with respect to this input. Thus, we just need to show that gate $j + 1$ is also correct. Within Algorithm 2, we fall in case $i' = j'$, meaning that gate $j + 1$ of block j' gates the same inputs as gate $j + 1$ of \mathcal{F} . By Algorithm 1, the type of gate $j + 1$ in block j' is the type of gate $j + 1$ of \mathcal{F} . Thus, gate $j + 1$ in block j' computes the same function of the input gates as gate $j + 1$ in \mathcal{F} . If $j + 1 = \text{bsize}(n, j')$, we conclude that the final p gates of block j' store the output of node j' . \square

Let XC^0 denote any family of constant-depth, poly-size circuits, including AC^0 and TC^0 .¹⁷

Corollary 3.1. *Let \mathcal{G} be a constant-depth, poly-size computation graph family over a finite \mathfrak{F} . If every node type in \mathfrak{F} can be computed by XC^0 circuits, the function computed by \mathcal{G} is in XC^0 .*

Since a transformer has constant depth and polynomial size, Corollary 3.1 lets us easily recover prior results about hard-attention transformers (Hao et al., 2022; Hahn, 2020) and saturated attention transformers (Merrill et al., 2022) using a common framework. All one has to do is show that all individual node types in such transformers can be computed by AC^0 and TC^0 circuits, respectively.

Corollary 3.1 established that Algorithms 1 and 2 construct a circuit family that simulates \mathcal{G} . With the right block mapping, \mathcal{C} will be log-uniform as long as \mathcal{G} and its node types are log-uniform.

Corollary 3.2. *Let \mathcal{G} be a log-uniform, constant-depth computation graph family over a finite \mathfrak{F} , where each $\mathcal{F} \in \mathfrak{F}$ is specified by a log-uniform TC^0 family with $\text{size}_{\mathcal{F}}(n) = \text{bsize}(n)$ that is a power of 2 computable in $O(\log n)$ time. Then \mathcal{G} can be simulated by a log-uniform TC^0 family \mathcal{C} that obeys the size and depth properties of Theorem 3.*

Proof. Let \mathcal{C} be the circuit family defined by Algorithms 1 and 2 given \mathcal{G} and the following block mapping: $\text{bnode}(n, i) = \lfloor i/\text{bsize}(n) \rfloor$, $\text{bstart}(n, i') = i' \cdot \text{bsize}(n)$, $\text{bsize}(n, i') = \text{bsize}(n)$. Since $\text{bsize}(n)$ is a power of 2, bnode and bstart are reducible to left and right shifting over $O(\log n)$ -bit integers, which can be implemented in $O(\log n)$ time. Thus, each block mapping function is computable in time $O(\log n)$. Since $\text{node}_{\mathcal{G}}$ and $\text{edge}_{\mathcal{G}}$ are just calling functions computable in time $O(\log n)$ with constant overhead, we conclude that \mathcal{C} , the circuit family they define, is log-uniform, and it is already known to simulate \mathcal{G} with constant depth and polynomial size by Theorem 3. \square

7 Conclusion

We proved that any log-precision transformer classifier can be translated to an $\text{FO}(\text{M})$ sentence that computes the same function (on all inputs of any length). This result comes by first simulating a transformer with a highly uniform threshold circuit family, and then leveraging the established equivalence of log-uniform circuits and $\text{FO}(\text{M})$. Transformers and other neural nets are often discussed in contrast with symbolic models based on logical formalisms (Garnelo & Shanahan, 2019)—an immediate implication of our result is that it is possible to express the inner workings of transformers also in a simple logic, challenging the premise of a rigid division between symbolic and neural models. Our results also provide the tightest known upper bound on log-precision transformers.

While it is striking that a full transformer can be translated to a sentence in a logic as simple as $\text{FO}(\text{M})$, we believe the bound is not tight. In particular, we conjecture that it is possible to simulate any transformer with an $\text{FO}(\text{M})$ sentence of quantifier depth of at most 2, which could be proven by establishing a hierarchy theorem describing the $\text{FO}(\text{M})$ quantifier depth needed to simulate a TC^0 family of a certain size. It would also be an interesting extension to translate real transformers to $\text{FO}(\text{M})$ sentences. In this sense, we believe our results provide a theoretical foundation to guide mechanistic interpretability work (cf. Weiss et al., 2021; Lindner et al., 2023).

Our findings provide a novel view into transformer classifiers and their limits. It would be exciting for future research to extend our results to account for other common practical uses of transformers, such as for long-form generation, chain-of-thought reasoning, and in-context learning.

¹⁷Formally, \mathfrak{F} just needs to contain \wedge and \vee .

Acknowledgments

We thank Paul Beame, David Chiang, anonymous reviewers, and researchers at the Allen Institute for AI for feedback. Thanks to Noa Nabeshima for pointing out a minor notational inconsistency. WM was supported by an NSF graduate research fellowship and in part by NSF award 1922658.

References

Aaronson, S., Kuperberg, G., and Habryka, O. TC^0 : Constant depth threshold circuits, 2022. URL https://complexityzoo.net/Complexity_Zoo:T#tc0.

Arora, S. and Barak, B. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

Barrington, D. A. M., Immerman, N., and Straubing, H. On uniformity within NC^1 . *Journal of Computer and System Sciences*, 41(3):274–306, 1990.

Bhattamishra, S., Ahuja, K., and Goyal, N. On the ability and limitations of transformers to recognize formal languages. In *EMNLP*, 2020.

Brent, R. P. and Zimmermann, P. *Modern computer arithmetic*, volume 18. Cambridge University Press, 2010.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In *NeurIPS*, 2020.

Chiang, D., Cholak, P., and Pillay, A. Tighter bounds on the expressivity of transformer encoders. *ICML*, 2023.

Elhage, N., Nanda, N., Olsson, C., Henighan, T., Joseph, N., Mann, B., Askell, A., Bai, Y., Chen, A., Conerly, T., DasSarma, N., Drain, D., Ganguli, D., Hatfield-Dodds, Z., Hernandez, D., Jones, A., Kernion, J., Lovitt, L., Ndousse, K., Amodei, D., Brown, T., Clark, J., Kaplan, J., McCandlish, S., and Olah, C. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021. <https://transformer-circuits.pub/2021/framework/index.html>.

Furst, M. L., Saxe, J. B., and Sipser, M. Parity, circuits, and the polynomial-time hierarchy. *Mathematical systems theory*, 17:13–27, 1984.

Garnelo, M. and Shanahan, M. Reconciling deep learning with symbolic artificial intelligence: representing objects and relations. *Current Opinion in Behavioral Sciences*, 29:17–23, 2019. ISSN 2352-1546.

Hahn, M. Theoretical limitations of self-attention in neural sequence models. *TACL*, 8:156–171, 2020.

Hao, Y., Angluin, D., and Frank, R. Formal language recognition by hard attention transformers: Perspectives from circuit complexity. *TACL*, 10:800–810, 2022.

Hesse, W. Division is in uniform TC^0 . In *International Colloquium on Automata, Languages, and Programming*, pp. 104–114. Springer, 2001.

Hunter, P., Bouyer, P., Markey, N., Ouaknine, J., and Worrell, J. Computing rational radical sums in uniform TC^0 . *Foundations of Software Technology and Theoretical Computer Science*, 2010.

Lindner, D., Kramár, J., Rahtz, M., McGrath, T., and Mikulik, V. Tracr: Compiled transformers as a laboratory for interpretability. *arXiv*, abs/2301.05062, 2023.

Liu, B., Ash, J. T., Goel, S., Krishnamurthy, A., and Zhang, C. Transformers learn shortcuts to automata. In *ICLR*, 2023.

Merrill, W. and Sabharwal, A. The parallelism tradeoff: Limitations of log-precision transformers. *TACL*, 11:531–545, 2023.

Merrill, W., Ramanujan, V., Goldberg, Y., Schwartz, R., and Smith, N. A. Effects of parameter norm growth during transformer training: Inductive bias from gradient descent. In *EMNLP*, 2021.

Merrill, W., Sabharwal, A., and Smith, N. A. Saturated transformers are constant-depth threshold circuits. *TACL*, 10, 2022.

Pérez, J., Marinković, J., and Barceló, P. On the Turing completeness of modern neural network architectures. In *ICLR*, 2019.

Thoppilan, R., Freitas, D. D., Hall, J., Shazeer, N. M., Kulshreshtha, A., Cheng, H.-T., Jin, A., Bos, T., Baker, L., Du, Y., Li, Y., Lee, H., Zheng, H., Ghafouri, A., Menegali, M., Huang, Y., Krikun, M., Lepikhin, D., Qin, J., Chen, D., Xu, Y., Chen, Z., Roberts, A., Bosma, M., Zhou, Y., Chang, C.-C., Krivokon, I. A., Rusch, W. J., Pickett, M., Meier-Hellstern, K. S., Morris, M. R., Doshi, T., Santos, R. D., Duke, T., Søraker, J. H., Zevenbergen, B., Prabhakaran, V., Díaz, M., Hutchinson, B., Olson, K., Molina, A., Hoffman-John, E., Lee, J., Aroyo, L., Rajakumar, R., Butryna, A., Lamm, M., Kuzmina, V. O., Fenton, J., Cohen, A., Bernstein, R., Kurzweil, R., Aguera-Arcas, B., Cui, C., Croak, M., Chi, E., and Le, Q. LaMDA: Language models for dialog applications. *ArXiv*, abs/2201.08239, 2022.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In *NeurIPS*, 2017.

Weiss, G., Goldberg, Y., and Yahav, E. On the practical computational power of finite precision RNNs for language recognition. In *ACL*, 2018.

Weiss, G., Goldberg, Y., and Yahav, E. Thinking like transformers. *ICML*, 2021.

Xiong, R., Yang, Y., He, D., Zheng, K., Zheng, S., Xing, C., Zhang, H., Lan, Y., Wang, L., and Liu, T.-Y. On layer normalization in the transformer architecture. In *ICML*, 2020.

A Conditional Majority

Given formulas $\phi, \psi, Mi : \phi$. ψ is a sentence that is true iff ψ is true for at least half the values of i that make ϕ true.

Proposition 2. *For any two predicates $\phi(i)$ and $\psi(i)$, $Mi : \phi(i) \cdot \psi(i)$ can be expressed in $\text{FO}(\mathbf{M})$.*

Proof. $Mi : \phi \cdot \psi$ can be rewritten using a counting quantifier and a threshold quantifier:

$$\exists k, k'. \left[2k' = k \wedge \exists^k i : \phi(i) \wedge \exists^{\geq k'} j : (\phi(j) \wedge \psi(j)) \right].$$

The formula $2k' = k$ can be defined using bit . We then use the fact that counting and threshold quantifiers can be expressed in terms of majority quantifiers (Barrington et al., 1990) to conclude that $Mi : \phi \cdot \psi$ can be expressed in $\text{FO}(\mathbf{M})$. \square

B Omitted Proofs

Table 1 summarizes the notation we use in the following proofs when describing computation graphs and circuit families.

Table 1: Summary of common notation for computation graph and circuit families.

Graph	Circuit	Output Range	Description
i'	i	\mathbb{Z}	index of node or gate
$\text{node}_{\mathcal{G}}(n, i')$	$\text{node}_{\mathcal{C}}(n, i)$	\mathfrak{F}^{18}	type of node or gate
$\text{edge}_{\mathcal{G}}(n, i', j')$	$\text{edge}_{\mathcal{C}}(n, i, j)$	\mathbb{Z}	argument # of edge $i \rightarrow j$
$\text{size}_{\mathcal{G}}(n)$	$\text{size}_{\mathcal{C}}(n)$	\mathbb{Z}	# of nodes or gates
$\text{depth}_{\mathcal{G}}(n)$	$\text{depth}_{\mathcal{C}}(n)$	\mathbb{Z}	longest path length
$\text{bnode}(n, i)$	$[0, \text{size}_{\mathcal{G}}(n)]$	block containing i	
$\text{bstart}(n, i')$	$[0, \text{size}_{\mathcal{C}}(n)]$	first gate in block i'	
$\text{bsize}(n, i')$	\mathbb{Z}	size of block i'	

B.1 Transformers are Log-Uniform Computation Graph Families

We now justify that the computation graph family defining a transformer is log-uniform. To do this, we introduce a stronger notion of uniformity called *column uniformity* that captures the highly regular structure of the transformer.

Let $\text{node}(G, i)$ be the i -th node of computation graph G . Let $a \bmod b$ be the remainder when a is divided by b .

Definition 6 (Column uniformity). A computation graph family \mathcal{G} is $T(n)$ -column-uniform iff there exists a computation graph K (with fixed size w.r.t n) such that, for all i, j such that $0 \leq i, j < \text{size}_{\mathcal{G}}(n)$:

1. $\text{node}_{\mathcal{G}}(n, i) = \text{node}(K, i \bmod \text{size}(K))$.
2. If $\lfloor i/\text{size}(K) \rfloor = \lfloor j/\text{size}(K) \rfloor$, then

$$\text{edge}_{\mathcal{G}}(n, i, j) = \text{edge}(K, i \bmod \text{size}(K), j \bmod \text{size}(K)).$$

Otherwise, $\text{edge}_{\mathcal{G}}(n, i, j)$ can be computed by a deterministic Turing machine in time $T(n)$.

We define *log-column-uniform* analogously to log-uniform: i.e., we let $T(n) = O(\log n)$. log-column-uniform implies log-uniform because our implementations of $\text{node}_{\mathcal{G}}$ and $\text{edge}_{\mathcal{G}}$ can store K in a finite lookup table and compute the quotient and remainder of i and j by $\text{size}(K)$ in $O(\log n)$ time using Lemma 12. The edges outside of K are computable in $O(\log n)$ time by construction.

Lemma 1 (Proof in Appendix B.1). *A transformer \mathcal{T} is a log-uniform computation graph family where \mathfrak{F} contains embedding, self-attention, feedforward, and output components.*

Proof. We show the stronger condition that any transformer \mathcal{T} is a log-column-uniform computation graph family, which implies it is log-uniform.

We have the column K by Definition 2: all that remains to show is that $\text{edge}_{\mathcal{G}_{\mathcal{T}}}$ can be computed in time $O(\log n)$ for edges outside the column. These edges route from the layer ℓ output to the self-attention heads of layer $\ell + 1$. Following from the column structure, there exists k_{ℓ} such that a node i is an output vector of layer ℓ iff $k_{\ell} = i \bmod \text{size}(K)$. In a finite lookup table, we can store k_{ℓ} for each $\ell + 1$, and use this for self-attention routing. For an unmasked self-attention head j , we compute:

$$\text{edge}_{\mathcal{G}_{\mathcal{T}}}(n, i, j) = \begin{cases} \lfloor i/\text{size}(K) \rfloor & \text{if } k_{\ell} = i \bmod \text{size}(K) \\ -1 & \text{otherwise.} \end{cases}$$

For causally masked attention, we extend the first case to check that $\lfloor i/\text{size}(K) \rfloor \leq \lfloor j/\text{size}(K) \rfloor$. Either way, this logic can be implemented in time $O(\log n)$ via Lemma 12. Thus, we conclude that $\mathcal{G}_{\mathcal{T}}$ is column-uniform. \square

B.2 Transformer Components are Computable by Log-Uniform Threshold Circuits

Lemma 2 (Proof in Appendix B.2). *Let \mathcal{T} be a log-precision transformer with fixed parameters $\theta_{\mathcal{T}}$. Then each component in \mathfrak{F} is computable in log-uniform TC^0 .*

We prove a more general version of Lemma 2 that handles some cases with weights growing with n . The weights $\theta_{\mathcal{T}}$ are just a special case of a computation graph (that do not depend on the input); we can thus apply our definition of log-uniform to them. Lemma 2 follows from a more general result with log-uniform $\theta_{\mathcal{T}}$:

Lemma 5. *Let \mathcal{T} be a log-uniform transformer with log-uniform $\theta_{\mathcal{T}}$. Then each component in \mathfrak{F} is computable in log-uniform TC^0 .*

Proof. In Appendix C, we show that log-uniform $\theta_{\mathcal{T}}$ implies:

1. The embedding component is computable in log-uniform TC^0 (Lemma 6).
2. The self attention mechanism is computable in log-uniform TC^0 (Lemma 7).
3. The activation block is computable in log-uniform TC^0 (Lemma 8).
4. The output classifier head is computable in log-uniform TC^0 (Lemma 9).

We have shown that each $\mathcal{F} \in \mathfrak{F}$ is computable in log-uniform TC^0 . \square

B.3 Transformer Component Size Has a Log-Time Upper Bound

Lemma 3 (Proof in Appendix B.3). *Let \mathcal{T} be a log-precision transformer with fixed parameters $\theta_{\mathcal{T}}$. There exists a function $\text{bsize}(n)$ that is a power of 2 and computable in $O(\log n)$ time s.t. $\text{size}_{\mathcal{F}}(n) \leq \text{bsize}(n)$ for all $\mathcal{F} \in \mathfrak{F}$.*

Proof. Let $2^{b(n)}$ be the least power of 2 at least as large as $\text{size}_{\mathcal{F}}(n)$ for all \mathcal{F} . We observe that $2^{b(n)}$ is at most $2 \cdot \max_{\mathcal{F}} \text{size}_{\mathcal{F}}(n)$ for all n . Because each \mathcal{F} has poly size, there is a fixed k such that, for large enough n ,¹⁹

$$\begin{aligned} 2^{b(n)} &\leq n^k \\ \Rightarrow b(n) &\leq k \lceil \log n \rceil. \end{aligned}$$

Define $b'(n) = k \lceil \log n \rceil$ and $\text{bsize}(n) = 2^{b'(n)}$. $\text{bsize}(n)$ is both a power of 2 and an upper bound on $2^{b(n)}$; what remains to be shown is that it can be computed in time $O(\log n)$. We can first compute $\lceil \log n \rceil$ in time $O(\log n)$ by finding the greatest nonzero index of n . Next, we can compute $b'(n) = k \cdot \lceil \log n \rceil$ in time $O(\log \log n)$ since k is fixed size and $\lceil \log n \rceil$ has size at most $O(\log \log n)$ (Brent & Zimmermann, 2010). Finally, we compute $\text{bsize}(n) = 2^{b'(n)}$ by simply left-shifting 1 at most $O(\log n)$ times. \square

¹⁹We can compute $\text{bsize}(n)$ for small n using finite lookup.

B.4 Circuit Families Can Be Padded to Log-Time Size Upper Bounds

Recall that the last p bits of our circuits represent the circuit's output (cf. Section 5.1). In Lemma 4, we consider $\mathcal{F}(x) = \mathcal{F}'(x)$ if and only if the last p bits of \mathcal{F} and \mathcal{F}' agree for all x .

Lemma 4 (Proof in Appendix B.4). *If \mathcal{F} is a log-uniform TC^0 family and $\text{size}_{\mathcal{F}}(n) \leq \text{bsize}(n)$, there exists a log-uniform TC^0 family \mathcal{F}' s.t. $\mathcal{F}(x) = \mathcal{F}'(x)$ for all x and $\text{size}_{\mathcal{F}'}(n) = \text{bsize}(n)$.*

Proof. The high level idea is that we can pad \mathcal{F} to a circuit \mathcal{F}' that has size $\text{bsize}(n)$ and simply copies over the p output bits of \mathcal{F} to its own last p bits using identity gates.

We first set $\text{node}_{\mathcal{F}'}$ to copy over the existing circuit and append identity nodes. Let Id denote an identity node. Then $\text{node}_{\mathcal{F}'}$ is defined as:

$$\text{node}_{\mathcal{F}'}(n, i) = \begin{cases} \text{node}_{\mathcal{F}}(n, i) & \text{if } \text{node}_{\mathcal{F}}(n, i) \neq \emptyset \\ \text{Id} & \text{if } \text{node}_{\mathcal{F}}(n, i) = \emptyset \wedge i < \text{bsize}(n) \\ \emptyset & \text{otherwise.} \end{cases}$$

We see that the size of \mathcal{F}' will thus be of size $\text{bsize}(n)$.

Next, we extend $\text{edge}_{\mathcal{F}'}(n, i, j)$ to route the original output bits to the new output bits. Recall that an edge value of 0 means i is the first argument of gate j , and an edge value of -1 means there is no edge $i \rightarrow j$. Let $k_j = p(n) - (\text{bsize}(n) - j)$ be the index of node j as an output gate in \mathcal{F}' . For example, $k = 0$ for the first output bit. Now let $\text{output}_{\mathcal{F}}(n, i, k)$ represent whether node i is the k -th output of F_n . We can compute $\text{output}_{\mathcal{F}}(n, i, k)$ in terms of $\text{node}_{\mathcal{F}}$ as follows:

$$\text{output}_{\mathcal{F}}(n, i, k) \iff \text{node}_{\mathcal{F}}(n, i + p(n) - k - 1) \neq \emptyset \wedge \text{node}_{\mathcal{F}}(n, i + p(n) - k) = \emptyset.$$

Then $\text{edge}_{\mathcal{F}'}$ is defined:

$$\text{edge}_{\mathcal{F}'}(n, i, j) = \begin{cases} \text{edge}_{\mathcal{F}}(n, i, j) & \text{if } \text{edge}_{\mathcal{F}}(n, i, j) \neq -1 \\ 0 & \text{if } \text{output}_{\mathcal{F}}(n, i, k_j) \\ -1 & \text{otherwise.} \end{cases}$$

The first condition simply copies over the original edges. The second condition adds $p(n)$ new edges (for the different values of k) that route the final $p(n)$ nodes of \mathcal{F} to the final $p(n)$ nodes of \mathcal{F}' , guaranteeing that the two circuits will compute the same function.

Because both $\text{node}_{\mathcal{F}'}$ and $\text{edge}_{\mathcal{F}'}$ just rely on addition, conditional branching, and a finite number of calls to functions computable in time $O(\log n)$, they are both computable in time $O(\log n)$. \square

C Transformer Column Components

In this section, we generally omit layer subscripts for clarity. We assume a pre-norm (Xiong et al., 2020) parameterization of the transformer for concreteness and because this is more standard in newer transformers. However, the results would also hold with the original post-norm (Vaswani et al., 2017).

As mentioned in the main text, we view $\theta_{\mathcal{T}}$ as a concatenation of the parameters for the transformer functions. Thus, if m and w are computable in time $O(\log n)$ and $\theta_{\mathcal{T}}$ is log-uniform, it follows that the parameter vector for each ϕ, s, v, f , and κ is itself log-uniform because we can map indices in the smaller parameter vectors to indices in $\theta_{\mathcal{T}}$ in time $O(\log n)$.

C.1 Transformer Embeddings

For each position $1 \leq i \leq n$, the transformer embedding function represents token $\sigma_i \in \Sigma$ and its position i with a vector. Let \mathbf{V} be an embedding matrix of size $|\Sigma| \times m$ where each row represents the embedding for some σ . Let $f : \mathbb{N} \rightarrow \mathbb{D}_p^m$ be computable in time $O(\log n)$. Then,

$$\phi(\sigma_i, i) = \mathbf{v}_{\sigma_i} + f(i).$$

Lemma 6. *If $\theta_{\mathcal{T}}$ is log-uniform, then ϕ is computable in log-uniform TC^0 .*

Proof. The embedding block can be expressed as a constant-size computation graph that constructs \mathbf{V} , computes \mathbf{v}_{σ_i} using an affine transformation, computes $f(i)$, and then, finally, sums \mathbf{v}_{σ_i} and $f(i)$. The first step is computable by a log-uniform constant-depth, poly-size threshold circuit family since $\theta_{\mathcal{T}}$ is log-uniform. We can compute an affine transformation via a log-uniform constant-depth poly-size threshold circuit family via Lemma 10. $f(i)$ can be directly computed by the Turing machine constructing the circuit by construction. The sum of the two terms can then be computed by a log-uniform constant-depth threshold circuit of size polynomial in m , which is also polynomial in n . Since we have a computation graph where all node types are computable by log-uniform, constant-depth, poly-size threshold circuit families, we conclude by Corollary 3.2 that ϕ can also be computed by log-uniform, constant-depth, poly-size threshold circuit family. \square

C.2 Self Attention

The two components of the self attention block are s , the similarity function, and v , the value function. Let \mathbf{h}_i be the hidden state at the previous layer and $\bar{\mathbf{h}}_i = \text{lnorm}(\mathbf{h}_i)$. Then, the similarity function first computes queries and keys, and then takes the scaled dot-product between them:

$$\begin{aligned}\mathbf{q}_i &= \mathbf{W}_q \bar{\mathbf{h}}_i + \mathbf{b}_q \\ \mathbf{k}_i &= \mathbf{W}_k \bar{\mathbf{h}}_i + \mathbf{b}_k \\ s(\mathbf{h}_i, \mathbf{h}_j) &= \exp \left(\frac{\mathbf{q}_i^\top \mathbf{k}_i}{\sqrt{m/h}} \right).\end{aligned}$$

Then the value function is defined $v(\mathbf{h}_i) = \mathbf{W}_h \bar{\mathbf{h}}_i + \mathbf{b}_h$. We first show that the value function (and also the keys and queries by symmetry) is computable in log-uniform TC^0 :

Lemma 7. *If $\theta_{\mathcal{T}}$ is log-uniform, then the self-attention component is computable in log-uniform TC^0 .*

Proof. v is a composition of constructing the parameters (in log-uniform TC^0 since $\theta_{\mathcal{T}}$ is log-uniform), layer norm (in log-uniform TC^0 by Lemma 11), and an affine transformation (in log-uniform TC^0 by Lemma 10). Thus, v is computable in log-uniform TC^0 .

Computing s is a constant-depth computation graph. First, we compute \mathbf{q}_i and \mathbf{k}_i and then multiply them, and all of these steps are in log-uniform TC^0 . Next, we can compute m and h in time $O(\log n)$ and build a log-uniform TC^0 circuit that divides the product of the last step by $\sqrt{m/h}$. Finally, we compute p -precision exp, which can be expressed in log-uniform TC^0 as multiplication followed by left-shifting. Thus, by Corollary 3.2, s can be computed in log-uniform TC^0 .

s and v are log-uniform, so their size p is at most $\text{poly}(n)$. Computing self attention reduces to binary multiplication and division over \mathbb{D}_p , and performing iterated addition (summation) over n numbers in \mathbb{D}_p . Binary multiplication, binary division (Hesse, 2001), and iterated addition (Merrill & Sabharwal, 2023) can all be computed in log-uniform TC^0 , i.e., by a log-uniform, constant-depth threshold circuit family of size at most $\text{poly}(p) \leq \text{poly}(n)$. Thus, self attention can also be computed in log-uniform TC^0 . \square

C.3 Activation Block

The activation function f encapsulates the aggregation of the attention head outputs and the feedforward subnetwork of the transformer. f takes as input attention head outputs $\mathbf{a}_{i,1}, \dots, \mathbf{a}_{i,h} \in \mathbb{D}_p^{m/h}$ and the previous layer value \mathbf{h}_i .

The first part of the activation block simulates the pooling part of the self-attention sublayer. The head outputs are first concatenated to form a vector \mathbf{a}_i , which is then passed through an affine transformation $(\mathbf{W}_o, \mathbf{b}_o) : \mathbb{D}_p^m \rightarrow \mathbb{D}_p^m$ followed by residual connections to form the sublayer output $\mathbf{o}_i \in \mathbb{D}_p^m$:

$$\mathbf{o}_i = \mathbf{W}_o \mathbf{a}_i + \mathbf{b}_o + \mathbf{h}_i.$$

The second part of the activation block first applies layer-norm and then simulates the feedforward subnetwork to compute the next layer vector \mathbf{h}'_i . Let $\bar{\mathbf{o}}_i = \text{lnorm}(\mathbf{o}_i)$. Let σ be a nonlinearity computable in linear time on its input (in the most standard transformer, ReLU). Then, for affine

transformations $(\mathbf{W}_1, \mathbf{b}_1) : \mathbb{D}_p^m \rightarrow \mathbb{D}_p^w$ and $(\mathbf{W}_2, \mathbf{b}_2) : \mathbb{D}_p^w \rightarrow \mathbb{D}_p^m$, the feedforward subnetwork can be defined:

$$\mathbf{h}'_i = \mathbf{W}_2 \sigma(\mathbf{W}_1 \bar{\mathbf{o}}_i + \mathbf{b}_1) + \mathbf{b}_2 + \mathbf{o}_i.$$

Lemma 8. *If $\theta_{\mathcal{T}}$ is log-uniform, then f is computable in log-uniform TC^0 .*

Proof. The activation block can be expressed as a constant-size computation graph where the nodes construct affine transformation parameters, apply affine transformations, compute layer-norm, and compute elementwise nonlinearities. Since each of these nodes is computable by a log-uniform, constant-depth, poly-size threshold circuit family, the activation block is as well. \square

C.4 Output Classifier Head

We assume the output from the transformer is computed as follows. First, $\bar{\mathbf{h}}_1 = \text{lnorm}(\mathbf{h}_1)$. Then, we use a parameter vector $\mathbf{w} \in \mathbb{D}_p^m$ and bias term b to compute:

$$\kappa(\mathbf{h}_1) = \text{sgn}(\mathbf{w}^\top \bar{\mathbf{h}}_1 + b).$$

Lemma 9. *If $\theta_{\mathcal{T}}$ is log-uniform, then κ is computable in log-uniform TC^0 .*

Proof. We can express computing κ as a composition of constructing the parameters \mathbf{w}, b and computing the affine transformation. Both parts of this composition are computable by a log-uniform, constant-depth, poly-size threshold circuit family, so computing κ is as well. \square

D Neural Net Building Blocks

In this section we analyze the uniformity of common neural net building blocks that are used within the various high-level transformer components.

D.1 Affine Transformations

Affine transformations are a core part of neural networks used in various parts of the transformer. An affine transformation takes as input parameters $(\mathbf{W}, \mathbf{b}) : \mathbb{D}_p^a \rightarrow \mathbb{D}_p^b$ and a vector $\mathbf{x} \in \mathbb{D}_p^a$ and returns $\mathbf{Wx} + \mathbf{b}$.

Lemma 10. *For $p = O(\log n)$, any p -precision affine transformation where \mathbf{W}, \mathbf{b} are log-uniform is computable by a log-uniform, constant-size threshold circuit family of size polynomial in a and b .*

Proof. We first use the uniformity of \mathbf{W}, \mathbf{b} to construct them in $O(\log n)$ time. For the transformation $\mathbf{Wx} + \mathbf{b}$, first compute each $\mathbf{w}_i \odot \mathbf{x}$ in parallel, where \odot represents elementwise multiplication. Since binary multiplication over polynomial-size numbers is in log-uniform TC^0 , this can be done in parallel with log-uniform TC^0 circuits. We then use b log-uniform, constant-depth, poly-size threshold circuit families, each corresponding to an output index, that compute the sum over the a entries of each $\mathbf{w}_i \odot \mathbf{x}$. The affine transformation corresponds to the composition of these two steps, and is thus computable by a log-uniform TC^0 circuit family. \square

D.2 Layer Norm

The layer norm is applied between sublayers in the transformer. Let $\mu = (1/d) \sum_{i=1}^d x_i$. The layer norm $\mathbf{y} \in \mathbb{D}_p^m$ of a vector $\mathbf{x} \in \mathbb{D}_p^m$ is computed, for scalars $a, b \in \mathbb{D}_p$,

$$\mathbf{y} = a \left(\frac{\mathbf{x} - \mu}{\|\mathbf{x} - \mu\|} \right) + b.$$

Lemma 11. *If a, b are log-uniform, the layer norm over a vector of size m can be computed by a log-uniform threshold circuit family of constant depth and size polynomial in m .*

Proof. First compute m using summation over the constant term 1 from 1 to m . This summation can be computed by a log-uniform constant-depth threshold circuit family of size polynomial in m . Then compute the sum over \mathbf{x} using a similar circuit, and divide them to get μ , using the fact that integer division is in log-uniform TC^0 (Hesse, 2001). We can then compute $\mathbf{x} - \mu$ in log-uniform TC^0 .

At this point, we can compute $\|\mathbf{x} - \mu\|$ in log-uniform TC^0 (Hunter et al., 2010), then divide each $\mathbf{x} - \mu$ by the norm in log-uniform TC^0 , and then apply the final affine transformation in log-uniform TC^0 (Lemma 10). Thus, computing layer norm is in log-uniform TC^0 . \square

E Arithmetic Complexity

Lemma 12. *Given an m -bit integer a and n -bit integer b , we can compute the quotient $\lfloor a/b \rfloor$ and remainder $a \bmod b$ in time $O(mn)$.*

Proof. Let $D(m, n)$ and $M(m, n)$ denote, respectively, the time complexity of dividing and multiplying an m -bit integer by an n -bit integer. Brent & Zimmermann (2010) give the following fact: $D(m + n, n) \leq O(M(m, n))$. With the goal of analyzing $D(m, n)$, we apply this as follows:

$$\begin{aligned} D(m, n) &\leq D(m + n, n) \\ &\leq O(M(m, n)) \\ &\leq O(mn). \end{aligned}$$

\square

Applying Lemma 12 when a has size $O(\log n)$ and b has size $O(1)$ says that we can do division in time $O(\log n)$.