

# A Transformer with Stack Attention


Jiaoda Li<sup>1</sup> Jennifer C. White<sup>2</sup> Mrinmaya Sachan<sup>1</sup> Ryan Cotterell<sup>1</sup>

<sup>1</sup>ETH Zürich <sup>2</sup>University of Cambridge

{jiaoda.li, mrinmaya.sachan, ryan.cotterell}@inf.ethz.ch jw2088@cam.ac.uk

## Abstract

Natural languages are believed to be (mildly) context-sensitive. Despite underpinning remarkably capable large language models, transformers are unable to model many context-free language tasks. In an attempt to address this limitation in the modeling power of transformer-based language models, we propose augmenting them with a differentiable, stack-based attention mechanism. Our stack-based attention mechanism can be incorporated into any transformer-based language model and adds a level of interpretability to the model. We show that the addition of our stack-based attention mechanism enables the transformer to model some, but not all, deterministic context-free languages.

 <https://github.com/rycolab/stack-transformer>

## 1 Introduction

Language models (LMs) based on the transformer architecture (Vaswani et al., 2017) have shown great empirical success at a wide range of NLP tasks (Devlin et al., 2019; Radford et al., 2019; Liu et al., 2020; Brown et al., 2020). However, recent theoretical (Hahn, 2020; Angluin et al., 2023) and empirical (Ebrahimi et al., 2020; Bhattamishra et al., 2020; Delétang et al., 2023) research suggests that language models based on transformers show difficulty in learning basic algorithmic patterns. A prime example is the Dyck- $n$  language, i.e., the language of balanced parentheses of depth  $\leq n$ . When  $n > 1$ , it has been argued that transformers are theoretically (Hahn, 2020) and empirically (Ebrahimi et al., 2020) unable to learn a Dyck- $n$  language. Additionally, Delétang et al. (2023) report that transformer-based LMs fail to learn four deterministic context-free (DCF) tasks. The authors of this work contend that the resolution of this insufficiency is paramount if human-level language understanding is to be achieved by computers. Indeed, Chomsky (1956) famously argues that human language has many context-free traits; see also Chomsky and Schützenberger (1963). Moreover, Shieber (1987)

goes further and argues that snippets of Swiss German are even higher on the Chomsky hierarchy.

The scientific question treated in this paper is whether there exists a minimal modification to the transformer architecture that *does* allow it to learn a larger swathe of the formal languages most closely associated with human language. Specifically, in this paper, we augment the transformer architecture with a novel stack attention mechanism that enables it to learn certain CF languages empirically. Our stack attention mechanism simulates a stack by maintaining a probability distribution over which of the subsequently observed tokens is at the top element of the stack. In turn, this probability distribution serves as an attention mechanism. Compared to DuSell and Chiang (2024), which also applies stack augmentation to the transformer, our stack attention is more space efficient and allows for easier interpretation through visualizing the attention weights. We incorporate our innovation into the transformer by adding a stack attention sub-layer to each layer, rather than completely replacing the standard attention. Augmenting models in a modular way like this allows for direct integration with pre-trained transformer-based LMs.

We evaluate our stack-augmented transformer through comparison with a standard transformer on four DCF tasks taken from Delétang et al. (2023). We find that the stack-augmented transformer performs substantially better than the standard transformer on two of the four DCF tasks. Nevertheless, in contrast to DuSell and Chiang (2024), who claim their architecture can recognize the entire class of CF languages, we find transformers with our stack attention still struggle on two DCF tasks that involve modular arithmetic.

## 2 Preliminaries

In this section, we provide the necessary technical background for our exposition. We first review the self-attention mechanism and then introduce the transformer architecture.

### 2.1 The Self-Attention Mechanism

The attention mechanism (Bahdanau et al., 2015) is the fundamental building block of the trans-

former architecture (Vaswani et al., 2017), which we discuss in the next section. One common form of attention is **self-attention** (Cheng et al., 2016; Parikh et al., 2016). Our construction of a stack-augmented attention mechanism is a modification of this self-attention mechanism.

The premise of self-attention is as follows. A sentence representation  $\mathbf{H} = [\mathbf{h}_1; \dots; \mathbf{h}_N] \in \mathbb{R}^{D \times N}$  is a horizontal concatenation of column vectors  $\mathbf{h}_n$  in  $\mathbb{R}^D$ , where each column is a representation of the  $n^{\text{th}}$  word. Our goal is to construct a distribution over the index set  $\{1, \dots, N\}$ , denoted as  $[N]$ . We do so in three steps, described below.

① The first step is to construct a real-valued, pair-wise compatibility score. The most common way to do this is through a (scaled) dot-product, i.e.,

$$e_{ij} \stackrel{\text{def}}{=} \frac{\mathbf{h}_i \cdot \mathbf{h}_j}{\sqrt{D}} \quad (1)$$

② The second step is to take the pair-wise compatibility scores and project them onto the simplex  $\Delta^{N-1}$  through the softmax. This results in the following distribution

$$\alpha_i(j) \stackrel{\text{def}}{=} \frac{\exp(e_{ij})}{\sum_{n=1}^N \exp(e_{in})} \quad (2)$$

which is termed the **self-attention distribution**. Note there are  $N$  self-attention distributions  $\alpha_i$ .

③ The third, and final, step is to construct a weighted average of the representations  $\mathbf{H} = [\mathbf{h}_1; \dots; \mathbf{h}_N] \in \mathbb{R}^{D \times N}$  using the self-attention distribution as follows

$$\mathbf{A}(\mathbf{H})_{:,i} \stackrel{\text{def}}{=} \sum_{n=1}^N \alpha_i(n) \mathbf{h}_n \quad (3)$$

where  $\mathbf{A}(\mathbf{H})_{:,i}$  denotes the  $i^{\text{th}}$  column of  $\mathbf{A}(\mathbf{H})$ . The function  $\mathbf{A} : \mathbb{R}^{D \times N} \rightarrow \mathbb{R}^{D \times N}$  (for any  $N$ ), as defined above, is called an **attention head**.

Importantly, we see that  $\mathbf{A}$  is a differentiable function. Differentiability allows us to learn the parameters of an attention head with gradient-based methods. And, more importantly, it has a specific desirable property—namely, it is invariant with respect to permutations of the columns of  $\mathbf{H}$ . Computationally, this implies that  $\mathbf{A}(\mathbf{H})_{:,i}$  and  $\mathbf{A}(\mathbf{H})_{:,j}$  can be computed in parallel for  $i \neq j$ . It is specifically this form of parallelism that grants the transformer architecture its ability to scale.

One drawback of the permutation invariance, however, is that  $\mathbf{A}$  is not a linguistically plausible mechanism as human language is decidedly not permutation invariant. This problem is addressed through the incorporation of attention masks and positional encodings (Vaswani et al., 2017, § 3.5) in the transformer architecture, as we discuss in §2.2.

**Masked Self-Attention.** An attention mask  $\mathbf{G} \in \mathbb{B}^{N \times N}$ , where  $\mathbb{B} = \{0, 1\}$ , can be applied to the self-attentions using the following generalization

$$\alpha_i(j) = \frac{\exp(e_{ij}) \mathbf{G}_{i,j}}{\sum_{n=1}^N \exp(e_{in}) \mathbf{G}_{i,n}} \quad (4)$$

Attention masks allow for hard constraints on which indices can be attended to by the attention head. A commonly used masking scheme is **future masking** where each position is only allowed to attend to positions up to and including itself, i.e., we define the following mask

$$\mathbf{G}_{i,n} = \begin{cases} 1, & n < i \\ 0, & n \geq i \end{cases} \quad (5)$$

Future masking allows transformers to be used in autoregressive language models by preventing the model from peeking at words that have yet to be generated, which we detail in §2.3.

**Queries, Keys, and Values.** In the version of the attention mechanism introduced by Vaswani et al. (2017), the attention mechanism is augmented with additional linear projections. Specifically, the vectors  $\mathbf{h}_n$  are linearly projected to construct three new vectors, defined below

$$\mathbf{q}_n \stackrel{\text{def}}{=} \mathbf{W}_Q \mathbf{h}_n \quad (\text{query}) \quad (6a)$$

$$\mathbf{k}_n \stackrel{\text{def}}{=} \mathbf{W}_K \mathbf{h}_n \quad (\text{key}) \quad (6b)$$

$$\mathbf{v}_n \stackrel{\text{def}}{=} \mathbf{W}_V \mathbf{h}_n \quad (\text{value}) \quad (6c)$$

where  $\mathbf{W}_V, \mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{D' \times D}$  are parameter matrices. Compatibility scores are then computed between the corresponding query–key pair:

$$e_{ij} = \frac{\mathbf{q}_i \cdot \mathbf{k}_j}{\sqrt{D'}} \quad (7)$$

Using those compatibility scores, a self-attention distribution is constructed using the softmax. Finally, as before, a weighted sum of the values is computed using the self-attention distribution:

$$\mathbf{A}(\mathbf{H})_{:,i} = \sum_{n=1}^N \alpha_i(n) \mathbf{v}_n \quad (8)$$

**Multi-head Self-Attention.** We additionally define the multi-head self-attention mechanism. In **multi-head attention**, we combine  $M$  attention heads  $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(M)}$  as follows

$$\mathbf{M}(\mathbf{H})_{:,i} \stackrel{\text{def}}{=} \sum_{m=1}^M \mathbf{W}_O^{(m)} \mathbf{A}^{(m)}(\mathbf{H})_{:,i} \quad (9)$$

where  $\mathbf{W}_O^{(m)} \in \mathbb{R}^{D \times D'}$  is the output projection matrix for head  $\mathbf{A}^{(m)}$ . Usually, we set  $D' = D/M$ .

## 2.2 The Transformer Architecture

We now describe the transformer architecture. A transformer over a vocabulary  $\Sigma$  constitutes a function of type<sup>1</sup>  $\Sigma^N \rightarrow \mathbb{R}^{D \times N}$  where a string  $\mathbf{w} = w_1 \cdots w_N \in \Sigma^N$  of length  $N$  is encoded into a  $\mathbb{R}^{D \times N}$  representation where  $D$  is the model size. The transformer is defined compositionally over a sequence of layers. First, we define

$$\mathbf{H}^{(0)} \stackrel{\text{def}}{=} \text{Embedding} + \text{PE} \quad (10)$$

where Embedding of type  $\Sigma^N \rightarrow \mathbb{R}^{D \times N}$  is the embedding layer and PE of type  $\Sigma^N \rightarrow \mathbb{R}^{D \times N}$  is the positional encoding that injects information about the relative or absolute position of tokens, which extinguishes the permutation invariance of the transformer. Each transformer layer consists of two sub-layers: a multi-head self-attention  $\mathbf{M}$  of type  $\mathbb{R}^{D \times N} \rightarrow \mathbb{R}^{D \times N}$  and a fully connected feed-forward network FFN of type  $\mathbb{R}^{D \times N} \rightarrow \mathbb{R}^{D \times N}$ . A residual connection is employed around each sub-layer, followed by a layer normalization (Ba et al., 2016) LN of type  $\mathbb{R}^{D \times N} \rightarrow \mathbb{R}^{D \times N}$ : for  $0 < \ell \leq L$ , we have the following recursive definition

$$\mathbf{H}_M^{(\ell)} \stackrel{\text{def}}{=} \text{LN} \left( \mathbf{M} \left( \mathbf{H}^{(\ell-1)} \right) + \mathbf{H}^{(\ell-1)} \right) \quad (11a)$$

$$\mathbf{H}_{\text{FFN}}^{(\ell)} \stackrel{\text{def}}{=} \text{LN} \left( \text{FFN} \left( \mathbf{H}_M^{(\ell)} \right) + \mathbf{H}_M^{(\ell)} \right) \quad (11b)$$

$$\mathbf{H}^{(\ell)} \stackrel{\text{def}}{=} \mathbf{H}_{\text{FFN}}^{(\ell)} \quad (11c)$$

where  $\mathbf{H}_M^{(\ell)}$ ,  $\mathbf{H}_{\text{FFN}}^{(\ell)}$  and  $\mathbf{H}^{(\ell)}$  are functions of type  $\Sigma^N \rightarrow \mathbb{R}^{D \times N}$  for any  $N$ . They have  $\mathbf{w}$  as input, which we omit for brevity when the context is clear.

**Future-masked Transformer.** If the future mask in Eq. (5) is used in every  $\mathbf{H}_M^{(\ell)}$ , we call such a transformer **future-masked transformer**, denoted as  $\mathbf{F}^{(L)}$ . As we will see, future-masked transformers are necessary to construct autoregressive language models, which cannot peek at the future.

<sup>1</sup>Type-theoretically,  $N$  is a parameter of the type. Thus, our type signature is a dependent type (Univalent Foundations Program, 2013)

## 2.3 Probability Models

Next, we describe two natural ways of constructing a probability distribution from a transformer.

**Masked Language Modeling.** First, we consider the case of masked language modeling (MLM). Masked language models perform the cloze task, i.e., they fill in a missing word given a left and right context. More formally, consider a string  $\mathbf{w} \in \Sigma^*$  of length  $T$ . Let  $w_t$  denote the  $t^{\text{th}}$  symbol in  $\mathbf{w}$ , let  $\mathbf{w}_{<t} = w_1 \cdots w_{t-1}$ , and let  $\mathbf{w}_{>t} = w_{t+1} \cdots w_T$ . We construct  $\tilde{\mathbf{w}} \stackrel{\text{def}}{=} \mathbf{w}_{<t}[\text{MASK}]\mathbf{w}_{>t}$  by replacing  $w_t$  in  $\mathbf{w}$  with a mask token [MASK]. The alphabet is expanded to include [MASK]. We denote  $\tilde{\Sigma} \stackrel{\text{def}}{=} \Sigma \cup \{[\text{MASK}]\}$ . The transformer  $\mathbf{H}^{(L)}$  is now of type  $\tilde{\Sigma}^N \rightarrow \mathbb{R}^{D \times N}$ . A masked language gives the following probability distribution for position  $t$

$$\begin{aligned} p(\tilde{w}_t \mid \mathbf{w}_{<t}, \mathbf{w}_{>t}) \\ = \text{softmax}(\mathbf{W}_P \mathbf{H}^{(L)}(\tilde{\mathbf{w}})_{:,t} + \mathbf{b}_P)_{\tilde{w}_t} \end{aligned} \quad (12)$$

where  $\tilde{w}_t \in \tilde{\Sigma}$ ,  $\mathbf{W}_P \in \mathbb{R}^{|\tilde{\Sigma}| \times D}$  and  $\mathbf{b}_P \in \mathbb{R}^{|\tilde{\Sigma}|}$ . In practice, multiple tokens may be masked and predicted simultaneously.

**Autoregressive Language Modeling.** In contrast to masked language modeling, the goal of autoregressive language modeling (ALM) is to construct a probability distribution over  $\Sigma^*$ . To do so, the following factorization is employed

$$p(\mathbf{w}) = p([\text{EOS}] \mid \mathbf{w}) \prod_{t=1}^T p(w_t \mid \mathbf{w}_{<t}) \quad (13)$$

Every local conditional distribution  $p(w_t \mid \mathbf{w}_{<t})$  is defined over the set  $\bar{\Sigma} \stackrel{\text{def}}{=} \Sigma \cup \{[\text{EOS}]\}$  and  $\mathbf{w}_{<1} \stackrel{\text{def}}{=} [\text{BOS}]$ , where  $[\text{BOS}], [\text{EOS}] \notin \Sigma$ .<sup>2</sup> In a transformer-based autoregressive language model, each local conditional is parameterized as

$$\begin{aligned} p(\bar{w}_t \mid \mathbf{w}_{<t}) \\ = \text{softmax}(\mathbf{W}_P \mathbf{F}^{(L)}(\mathbf{w})_{:,t-1} + \mathbf{b}_P)_{\bar{w}_t} \end{aligned} \quad (14)$$

where  $\bar{w}_t \in \bar{\Sigma}$ ,  $\mathbf{F}^{(L)}$  is a future-masked transformer,  $\mathbf{W}_P \in \mathbb{R}^{|\bar{\Sigma}| \times D}$  and  $\mathbf{b}_P \in \mathbb{R}^{|\bar{\Sigma}|}$ .

## 3 A Transformer with Stack Attention

Recently Delétang et al. (2023) showed that transformers are not able to learn several non-regular

<sup>2</sup>This means that the transformer is a function of type  $\Sigma^{N+1} \rightarrow \mathbb{R}^{D \times (N+1)}$  where  $\Sigma \stackrel{\text{def}}{=} \Sigma \cup \{[\text{BOS}]\}$ .

Action	Stack	Attention	Stack Attention $\alpha$
		[BOS] a b c ↑	$\alpha_0 = [1, 0, 0, 0]^\top$
PUSH a		[BOS] a b c ↑	$\alpha_1 = [0, 1, 0, 0]^\top$
PUSH b		[BOS] a b c ↑	$\alpha_2 = [0, 0, 1, 0]^\top$
PUSH c		[BOS] a b c ↑	$\alpha_3 = [0, 0, 0, 1]^\top$
POP		[BOS] a b c ↑	$\alpha_4 = \sum_{j=1}^3 \alpha_3(j) \alpha_{j-1} = \alpha_2 = [0, 0, 1, 0]^\top$
NO-OP		[BOS] a b c ↑	$\alpha_5 = \alpha_4 = [0, 0, 1, 0]^\top$
POP		[BOS] a b c ↑	$\alpha_6 = \sum_{j=1}^5 \alpha_5(j) \alpha_{j-1} = \alpha_1 = [0, 1, 0, 0]^\top$

Figure 1: An example illustrating how attentions can emulate stacks. The first column lists the operation performed at each timestep. The second column presents the stack contents after performing the operation. The third column shows a hard attention over the input tokens. The pointer of the attention indicates the current stack top. The last column is the proposed stack attention.

DCF languages. Inspired by the fact that pushdown automata (Oettinger, 1961; Schützenberger, 1963), automata that employ a single stack, can model CF languages (Evey, 1963), we introduce a novel stack attention mechanism that emulates the functionality of a stack and integrate it into the transformer architecture, aiming to enable it to learn some CF languages.

### 3.1 Stacks over the Index Set

We first give a formal definition of a stack. In our paper, we define a stack as a data structure over the index set  $[N]$ . The state of a stack is a string  $\gamma \in [N]^*$  of indices. There are three operations that we can perform that alter the state of the stack. We describe each operation below in terms of  $\gamma$ .

- The operation PUSH:  $[N]^* \times [N] \rightarrow [N]^*$  adds an element to the top of the stack and is formally defined as follows:

$$\text{PUSH}(\gamma, \gamma) = \gamma\gamma \quad (15)$$

- The operation NO-OP:  $[N]^* \rightarrow [N]^*$  leaves the stack unchanged and is defined as follows:

$$\text{NO-OP}(\gamma) = \gamma \quad (16)$$

- The operation POP:  $[N]^* \rightarrow [N]^*$  removes the top-most element of the stack and is formally defined as follows:

$$\text{POP}(\varepsilon) = \varepsilon \quad (17a)$$

$$\text{POP}(\gamma_1 \cdots \gamma_T) = \gamma_1 \cdots \gamma_{T-1} \quad (17b)$$

We will use this definition in Theorem 3.1 to argue that our stack attention mechanism can be formally viewed as a type of stack. Additionally, we will assume an operator PEEK:  $[N]^* \rightarrow ([N] \cup \{0\})$  that does not alter the state of the stack, but rather returns the top element (or 0 if the stack is empty). We define it below

$$\text{PEEK}(\varepsilon) = 0 \quad (18a)$$

$$\text{PEEK}(\gamma_1 \cdots \gamma_T) = \gamma_T \quad (18b)$$

### 3.2 Stack Attention

We now formally define our stack attention mechanism. We introduce a beginning-of-sequence symbol [BOS] at the zeroth position, designated to represent an empty stack. Each position  $i \in \{0\} \cup [N]$  is assigned a distinct stack  $\alpha_i \in \mathbb{R}^{N+1}$ . We write  $\alpha_i(j)$  to denote the  $(j+1)^{\text{th}}$  value in  $\alpha_i$ , for  $0 \leq j \leq N$ . The stacks are defined inductively.

The initial stack,  $\alpha_0$ , is constructed to attend to [BOS] as follows

$$\alpha_0 \stackrel{\text{def}}{=} [1, 0, 0, \dots]^\top \in \mathbb{R}^{N+1} \quad (19)$$

Subsequent stacks are computed inductively based on the stack contents and the operations (PUSH, NO-OP, POP) taken at previous timesteps.<sup>3</sup> The three stack operations are defined for  $i \geq 1$  as follows:

- PUSH pushes the hidden state of the current position, so we just set the attention weight at the current position to be 1 and the rest to be 0, i.e.,

$$\alpha_i^{(\text{PUSH})}(j) \stackrel{\text{def}}{=} \begin{cases} 1 & j = i \\ 0 & \text{otherwise} \end{cases} \quad (20)$$

- NO-OP leaves the previous stack unchanged, so the stack from the last timestep is passed forward with no modification, i.e., we have

$$\alpha_i^{(\text{NO-OP})} \stackrel{\text{def}}{=} \alpha_{i-1} \quad (21)$$

- POP removes the top element, and backtracks to the second element in the stack, i.e., we have

$$\alpha_i^{(\text{POP})} \stackrel{\text{def}}{=} \left[ \sum_{j=1}^{i-1} \alpha_{i-1}(j) \alpha_{j-1} \right] + \alpha_{i-1}(0) \alpha_0 \quad (22)$$

The first term on the right-hand side retrieves the second element and is zeroed out when  $i = 1$ . The second term accounts for the case of an empty stack—POP cannot be performed on an empty stack and in this case it is equivalent to a NO-OP.

The stack attention  $\alpha_i$  at position  $i$  is computed as a superposition of the three operations:

$$\alpha_i \stackrel{\text{def}}{=} \mathbf{a}_i(\text{PUSH}) \cdot \alpha_i^{(\text{PUSH})} + \mathbf{a}_i(\text{POP}) \cdot \alpha_i^{(\text{POP})} + \mathbf{a}_i(\text{NO-OP}) \cdot \alpha_i^{(\text{NO-OP})} \quad (23)$$

where  $\mathbf{a}_i \in \Delta^2$  is a probability distribution over possible operations  $\mathcal{A} = \{\text{PUSH}, \text{POP}, \text{NO-OP}\}$ . This distribution is determined by:

$$\mathbf{a}_i \stackrel{\text{def}}{=} \text{softmax}(\mathbf{W}_A \mathbf{h}_i + \mathbf{b}_A) \quad (24)$$

where  $\mathbf{W}_A \in \mathbb{R}^{3 \times D}$  and  $\mathbf{b}_A \in \mathbb{R}^3$  are learned parameters.

After obtaining the stack attention weights, we can compute the top element as a weighted sum just like standard self-attention:

$$\mathbf{S}(\mathbf{H})_{:,i} \stackrel{\text{def}}{=} \sum_{n=0}^N \alpha_i(n) \mathbf{h}_n \quad (25)$$

<sup>3</sup>We use the terms timestep and position interchangeably.

### 3.3 A Stack Transformer

The stack is incorporated into the transformer by inserting a third sub-layer in each transformer layer after the standard attention and feedforward layers defined in Eq. (11a) and Eq. (11b):

$$\mathbf{H}_S^{(\ell)} \stackrel{\text{def}}{=} \mathbf{S}(\mathbf{H}_{\text{FFN}}^{(\ell)}) + \mathbf{H}_{\text{FFN}}^{(\ell)} \quad (26a)$$

$$\mathbf{H}^{(\ell)} = \mathbf{H}_S^{(\ell)} \quad (26b)$$

Similar to other sub-layers, we also employ a residual connection by summing the output of the stack attention mechanism  $\mathbf{S}$  with its input, allowing the model to bypass the stack if needed. Layer normalization can also be used, but we omit due to initial results in preliminary experiments. Because the rest of the model is left unchanged, it can be directly integrated into pre-trained language models to augment their ability to process hierarchical syntactic structures.

### 3.4 Computational Overhead

**Time.** The computation is bottlenecked by the POP operation, which sums over previous the previous positions and thereby has a time complexity of  $\mathcal{O}(N)$ . The total time complexity is  $\mathcal{O}(N^2)$ . In contrast to standard attention, stack attention has to be computed sequentially, which breaks the parallelizability of the transformer and makes it substantially slower in practice. However,  $\mathbf{a}$  and the output  $\mathbf{S}(\mathbf{H})$  can still be computed in parallel. Thus,  $\alpha$  is a function of the stack operations but *not* of the hidden states. It follows that if structural supervision of the stack operations is provided, e.g., as in Sartran et al. (2022) and Murty et al. (2023),  $\alpha_i$  for all  $i \in [N]$  can be pre-computed, and the entire model can be parallelized.

**Space.** The stack attention stores  $N + 1$  attentions of size  $N$ , so the space complexity is  $\mathcal{O}((N + 1)N) = \mathcal{O}(N^2)$ . This is an improvement over the  $\mathcal{O}(DN^2)$  space complexity of DuSell and Chiang’s (2024) method.

### 3.5 The Duality of Stack Attention

Stack attention is both a stack over the index set, as defined in §3.1, and an attention mechanism. In the following theorem, we make precise the manner in which our stack attention is a stack.

**Notation.** We use the symbol  $v_i$  to refer to an operation from the set  $\{\text{PUSH}_i(\cdot), \text{NO-OP}(\cdot), \text{POP}(\cdot)\}$  at every time step  $i$ . Note that PUSH, as defined in

Task	RNN		Transformer MLM		Transformer ALM	
	Vanilla	Stack	Vanilla	Stack	Vanilla	Stack
RS	81.0 ± 0.8	100.0 ± 0.0	54.8 ± 0.0	100.0 ± 0.0	55.4 ± 0.8	100.0 ± 0.0
SM	73.2 ± 1.0	100.0 ± 0.0	50.4 ± 0.1	93.1 ± 4.4	50.4 ± 0.1	92.8 ± 2.6
MA	75.8 ± 4.3	91.0 ± 6.3	30.1 ± 0.0	34.3 ± 1.4	30.2 ± 0.1	29.5 ± 0.6
SE	56.7 ± 10.3	89.9 ± 7.2	20.0 ± 0.0	29.8 ± 8.0	20.2 ± 0.1	20.3 ± 0.2

Table 1: Accuracies (%) of the transformer and RNN without and with stacks on DCF tasks.

§3.1, is a function of two arguments. However, we define  $\text{PUSH}_i(\gamma) \stackrel{\text{def}}{=} \text{PUSH}(\gamma, i)$ , i.e., we push  $i$ , the index, to the stack. We introduce a function  $\llbracket \cdot \rrbracket$  of type  $\{0\} \cup [N] \rightarrow \mathbb{B}^{N+1}$  that converts an index into its one-hot encoding, a column vector with zeros everywhere except the given index, where the value is set to one.

**Theorem 3.1.** *Let  $v_1, \dots, v_N$  be a series of stack operations where  $v_i \in \{\text{PUSH}_i(\cdot), \text{NO-OP}(\cdot), \text{POP}(\cdot)\}$  for all  $i \in [N]$ . Furthermore, suppose  $\mathbf{a}_i(v_i) = 1$  for all  $i \in [N]$ . Then,  $\llbracket \text{PEEK}(v_i(\dots v_1(\varepsilon))) \rrbracket = \alpha_i$  for all  $i \in \{0\} \cup [N]$ .*

*Proof.* Appendix A ■

Our stack-based attention is also an attention mechanism in the sense that it maintains a distribution over  $\{0\} \cup [N]$ . We make this notion precise as well in the following theorem.

**Theorem 3.2.** *Consider a sequence of stack attention mechanisms  $\alpha_0, \dots, \alpha_N$ . Then,  $\sum_{n=0}^N \alpha_i(n) = 1$  for all  $i \in \{0\} \cup [N]$ .*

*Proof.* Appendix A ■

### 3.6 Expressivity

We leave the exact characterization of the expressivity of our stack-augmented transformer for future work. However, we conjecture that it cannot model all the context-free languages *without* positional encodings. Such a result would mirror that of [Angluin et al. \(2023\)](#).

To contextualize this conjecture, we first review the star-free languages. The star-free languages are regular languages definable by a regular expression *without* Kleene star but with complement ([McNaughton and Papert, 1971](#)). They can also be characterized by finite-state automata with aperiodic transformation monoids ([Schützenberger, 1965](#)), also termed counter-free automata or permutation-free automata ([McNaughton and Papert, 1971](#)). It has been shown that a counter-free automaton can

only perform counting up to a threshold, but not modulo counting ([McNaughton and Papert, 1971](#)).

Recently, [Angluin et al. \(2023\)](#) showed that the class of languages recognizable by transformer encoders with hard attention, strict future masking, and no positional encodings, are exactly the star-free languages. Building on this result, we conjecture that there exist non-star-free languages that are beyond the capability of a transformer encoder with (hard) stack attention and no positional encodings. This conjecture is supported by our experiments in §4.3 where we show that stack-augmented transformers also fail to learn two tasks involving modulo counting. We hope to construct a proof of an expressivity result in future work.

## 4 Deterministic CF Tasks

We now discuss several tasks that are encodable by deterministic context-free grammars.

### 4.1 Tasks

All four tasks we consider are derived from [Delétang et al. \(2023\)](#) and are language transduction tasks. Every word from the input language  $\mathbf{x} \in \Sigma_I^*$  is mapped to a word in the output language  $\mathbf{y} \in \Sigma_O^*$  by means of a function  $f: \Sigma_I^* \rightarrow \Sigma_O^*$ . To convert a transduction task to a language acceptance task, a language is constructed over the alphabet  $\Sigma = \Sigma_I \cup \Sigma_O$  as follows

$$\left\{ \mathbf{x}f(\mathbf{x}) \mid \mathbf{x} \in \Sigma_I^* \right\} \subseteq \Sigma^* \quad (27)$$

To experiment with this setup, in the case of an MLM, the input  $\tilde{\mathbf{w}}$  is  $\mathbf{x}$  appended with  $|\mathbf{y}|$  mask tokens [MASK]. We then use the transformer to predict all the masked tokens at once and evaluate the predicted string  $\mathbf{y}'$  against  $\mathbf{y} = f(\mathbf{x})$ . Likewise, in the case of an ALM, given a prefix  $\mathbf{x}$ , we sample  $y'_t \sim p(\cdot \mid \mathbf{x}\mathbf{y}'_{<t})$  autoregressively, where  $y'_t$  denotes the  $t^{\text{th}}$  symbol of  $\mathbf{y}'$  and  $\mathbf{y}'_{<t} = y'_1 \cdots y'_{t-1}$ . As in the case of MLM, we evaluate the predicted

Task	Transformer	none	sincos	relative	rotary	ALiBi
RS	Vanilla	$54.8 \pm 0.0$	$50.7 \pm 0.2$	$67.6 \pm 2.2$	$55.4 \pm 1.2$	$79.4 \pm 3.5$
	Stack	$100.0 \pm 0.0$	$99.1 \pm 1.7$	$100.0 \pm 0.0$	$86.3 \pm 15.0$	$100.0 \pm 0.0$
SM	Vanilla	$50.4 \pm 0.1$	$49.5 \pm 0.6$	$67.5 \pm 1.0$	$52.1 \pm 1.8$	$70.9 \pm 1.2$
	Stack	$93.1 \pm 4.4$	$74.7 \pm 8.8$	$98.5 \pm 1.1$	$73.1 \pm 4.5$	$92.9 \pm 2.7$
MA	Vanilla	$30.1 \pm 0.0$	$30.1 \pm 0.0$	$30.1 \pm 0.0$	$30.1 \pm 0.0$	$30.1 \pm 0.0$
	Stack	$34.3 \pm 1.4$	$33.8 \pm 0.8$	$35.0 \pm 1.1$	$34.5 \pm 1.3$	$34.7 \pm 1.1$
SE	Vanilla	$20.0 \pm 0.0$	$20.0 \pm 0.0$	$20.0 \pm 0.0$	$20.0 \pm 0.0$	$20.0 \pm 0.0$
	Stack	$29.8 \pm 8.0$	$23.9 \pm 3.0$	$25.2 \pm 1.8$	$30.0 \pm 3.8$	$27.9 \pm 5.8$

(a) MLM

Task	Transformer	none	sincos	relative	rotary	ALiBi
RS	Vanilla	$55.4 \pm 0.8$	$55.2 \pm 0.7$	$62.0 \pm 6.1$	$72.9 \pm 3.5$	$57.1 \pm 0.6$
	Stack	$100.0 \pm 0.0$	$96.8 \pm 4.5$	$100.0 \pm 0.0$	$100.0 \pm 0.0$	$100.0 \pm 0.0$
SM	Vanilla	$64.9 \pm 2.0$	$60.8 \pm 3.1$	$70.5 \pm 0.9$	$71.9 \pm 0.9$	$70.5 \pm 1.6$
	Stack	$92.8 \pm 2.6$	$49.6 \pm 4.4$	$93.2 \pm 2.3$	$83.8 \pm 1.7$	$93.4 \pm 1.2$
MA	Vanilla	$30.2 \pm 0.1$	$25.7 \pm 2.3$	$30.3 \pm 0.1$	$26.0 \pm 0.8$	$30.3 \pm 0.1$
	Stack	$30.0 \pm 0.1$	$28.0 \pm 2.8$	$30.3 \pm 0.3$	$25.6 \pm 0.3$	$30.3 \pm 0.1$
SE	Vanilla	$20.2 \pm 0.1$	$20.2 \pm 0.3$	$20.7 \pm 0.2$	$20.3 \pm 0.2$	$20.5 \pm 0.3$
	Stack	$20.3 \pm 0.2$	$20.2 \pm 0.1$	$20.7 \pm 0.1$	$20.2 \pm 0.3$	$20.3 \pm 0.1$

(b) ALM

Table 2: Performance comparison of a vanilla and stack transformer with different positional encodings.

$\mathbf{y}'$  against the  $\mathbf{y} = f(\mathbf{x})$ . We follow the choices of Delétang et al. (2023) for  $\Sigma_I$  and  $\Sigma_O$

**Reverse String (RS).** The first task is to compute the reverse of an input string, i.e.,  $f_{RS}(\mathbf{x}) = \mathbf{x}^R$ . In this task, we take  $\Sigma_I = \Sigma_O = \{a, b\}$ . We give an example below.

**Example:**

$$\begin{aligned} \mathbf{x} &= \text{abb} \\ \mathbf{y} &= \text{bba} \end{aligned}$$

**Stack Manipulation (SM).** In the second task, the input string  $\mathbf{x}$  consists of a stack of two symbols  $\{a, b\}$ , printed from bottom to top, and a sequence of stack operations drawn from the set  $\{\text{[PUSH a]}, \text{[PUSH b]}, \text{[POP]}\}$ . The function  $f_{SM}(\mathbf{x})$  outputs the final stack after all the given operations are executed sequentially on the input stack, printed from top to bottom. We always have  $|\mathbf{y}| = |\mathbf{x}| + 1$ . If the final stack has fewer elements than  $|\mathbf{y}|$ , it will be padded with [PAD] tokens, which are ignored when accuracy is computed. We have  $\Sigma_I = \{a, b, \text{[PUSH a]}, \text{[PUSH b]}, \text{[POP]}\}$  and

$\Sigma_O = \{a, b, \text{[PAD]}\}$ . An example is given below.

**Example:**

$$\begin{aligned} \mathbf{x} &= \text{bab[POP][PUSH a][PUSH b]} \\ \mathbf{y} &= \text{baab[PAD][PAD][PAD]} \end{aligned}$$

**Modular Arithmetic (MA).** In the third task, we consider a transduction task based on modular arithmetic. An algebraic expression  $\mathbf{x}$  consists of five numerical constants  $\{0, 1, 2, 3, 4\}$ , three operations  $\{+, -, \cdot\}$ , brackets  $\{(, )\}$ , and a congruence sign  $\{\equiv\}$ . We say two integers are congruent if and only if a pre-set modulus is a divisor of their difference. In this task, we set the modulus to 5, so the function  $f_{MA}$  evaluates the expression modulo 5. We have  $\Sigma_I = \{0, 1, 2, 3, 4, +, -, \cdot, (, ), \equiv\}$  and  $\Sigma_O = \{0, 1, 2, 3, 4\}$ . An example is given below.

**Example:**

$$\begin{aligned} \mathbf{x} &= (1 + 2) \cdot 3 \equiv \\ \mathbf{y} &= 4 \end{aligned}$$

**Solve Equation (SE).** In our fourth and final task, we consider a transduction task that solves equations over a single variable, which we denote  $z$ . The input string  $x$  is a modular equation with five constants  $\{0, 1, 2, 3, 4\}$ , two operations  $\{+, -\}$ , brackets  $\{(, )\}$ , a congruence sign  $\{\equiv\}$ , and a variable  $\{z\}$ . The modulus is set to 5. The function  $f_{SE}$  solves this equation and returns the value of the variable. We have  $\Sigma_I = \{0, 1, 2, 3, 4, +, -, (, ), \equiv, z\}$  and  $\Sigma_O = \{0, 1, 2, 3, 4\}$ . An example is given below.

**Example:**

$$\begin{aligned}x &= (1 + z) + 2 \equiv 2 \\y &= 4\end{aligned}$$

## 4.2 Experimental Setup

Following Delétang et al. (2023), we experiment with a transformer with the number of layers  $L = 5$  and the model size  $D = 64$ . Unless otherwise specified, no positional encodings are used. We discuss the effect of various positional encodings in §4.3.2. Length generalization has been the focus of many papers in this line of research (Joulin and Mikolov, 2015; Delétang et al., 2023). We follow suit to train on input strings  $x$  with  $1 \leq |x| \leq 40$  and test on  $x$  with  $40 < |x| \leq 100$ . Training details can be found in Appendix B.

## 4.3 Results

We report our results of the four DCF tasks in Tab. 1 and Tab. 2.

### 4.3.1 Transformer vs. Stack Transformer

We report the performance of the standard transformer and our stack-augmented transformer on the four DCF tasks presented in Tab. 1. For comparison, we also exhibit results of vanilla recurrent neural networks (RNNs) and stack-RNNs (Joulin and Mikolov, 2015). As expected, the vanilla transformer exhibits poor performance on all the DCF tasks. After being augmented with a stack, the transformer improves from nearly chance to over 90% on RS and SM. These results demonstrate that our stack-augmented attention helps on some tasks. However, on MA and SE, the performance after adding the stack attention only improves slightly; it still falls far behind stack RNNs and even vanilla RNNs. We conjecture that the reason for this shortcoming is our stack transformer’s incapability to learn non-counter-free languages—both the

last two tasks (MA) and SE require the ability to perform modular arithmetic, which makes them non-star-free, as discussed in §3.6. Additionally, Feng et al. (2023) also directly prove that the transformer cannot perform modular arithmetic.

### 4.3.2 Positional Encodings

In this section, we add various positional encodings to the transformer and investigate their effect. We consider five different positional encodings: none, sincos, relative, rotary, and ALiBi; see Appendix C for more details. As our stack attention is computed inductively, positional information is already present in the model, reducing the need for positional encodings. This is evident in Tab. 2, where including positional encodings generally has a *negative* impact on the stack transformer’s performance. Most notably, sincos and rotary heavily degrade stack transformer’s performance on RS and SM. However, relative constitutes an exception, as it results in improved performance on SM. In contrast, with the standard transformer architecture, positional encodings *do* seem to help on star-free tasks. The largest improvement comes from ALiBi in the MLM setting and rotary in the ALM setting. Nevertheless, none of the investigated positional encodings are able to boost the performance of vanilla transformer to anywhere near that of our stack-augmented transformer.

## 5 Language Modeling

We consider masked language modeling using RoBERTa (Liu et al., 2020) and autoregressive language modeling using GPT-2 (Radford et al., 2019). Following the experimental setup proposed by previous authors (Joulin and Mikolov, 2015; DuSell and Chiang, 2024), we experiment on the Penn Treebank (PTB), licensed through the LDC (Marcus et al., 1994), and WikiText-2 (Merity et al., 2017). We consider models both trained from scratch and fine-tuned from pre-trained weights. The pre-trained models and datasets are obtained from HuggingFace (Wolf et al., 2020; Lhoest et al., 2021). See Appendix B for more details about setup and hyperparameters.

The results in Tab. 3 are mixed. Our major finding is that transformers benefit from the stack attention when training data is scarce, but the benefits gradually diminish as the size of training data grows. More concretely, when the models are trained from scratch, the addition of our stack attention mechanism *does* result in a noticeable benefit



Model	Task	Penn Treebank		WikiText-2	
		Vanilla	Stack	Vanilla	Stack
Scratch	MLM	95.53 ± 19.66	34.28 ± 2.76	73.74 ± 3.79	64.75 ± 1.75
	ALM	73.14 ± 0.34	69.86 ± 0.26	191.01 ± 0.71	206.42 ± 0.80
Pretrained	MLM	3.99 ± 0.08	4.46 ± 0.11	4.41 ± 0.12	4.65 ± 0.06
	ALM	21.26 ± 0.03	32.36 ± 0.16	29.29 ± 0.02	54.96 ± 0.19

Table 3: MLM and ALM Perplexities on WikiText-2 and PTB.

in most settings. In the MLM setting, where 15% of the tokens are replaced with [MASK], stacks reduce the perplexity under the trained model on the held-out split from 95.53 to 34.28 on PTB and from 73.74 to 65.22 on WikiText-2. In the ALM setting, the stack transformer still slightly improves the performance on PTB—perplexity drops from 73.14 to 69.86. However, the stack transformer is less effective on WikiText-2, whose training set is larger. Moreover, when we fine-tune from pre-trained models, stacks are always detrimental across the two datasets in both MLM and ALM settings.

## 6 Discussion

From the results described in §4.3 and §5, we observe two trends:

- The positive impact of stack attention is evident on Delétang et al.’s (2023) 4 DCF tasks (especially on 2 of the 4), but almost nonexistent on English language modeling;
- On the English language modeling task, stack attention is more helpful in settings with limited training data, but is less helpful and can even be harmful when the model is trained on a larger amount of data.

We interpret these trends as support for the idea that stack attention improves the representational capacity of a transformer language model and, additionally, confers an inductive bias to the transformer architecture that allows it to better learn certain context-free tasks more efficiently. The larger representational capacity explains why the performance on certain tasks, i.e., RS and SM, improves drastically with the addition of stack attention and the better inductive bias explains why transformer language models with stack attention perform better with less training data on the English modeling task. However, the fact that a vanilla transformer language model performs on par with stack attention when modeling larger

English language datasets suggests that a good inductive bias is not needed for larger data sets. This suggests that, in contrast to the viewpoint of traditional linguistic theory (Chomsky, 1957), models that are higher on the Chomsky hierarchy are *not* necessary for developing a good statistical language model. We believe this claim is consistent with the literature, in which few successful large language models are endowed with a syntactic bias. However, there are many smaller syntax-infused language models (Dyer et al., 2016) that do work well on smaller data, as ours does.

## 7 Conclusion

We propose a novel implementation of a differentiable stack and show that a transformer augmented with such stacks can solve certain deterministic context-free tasks that are beyond the capability of standard transformers. However, unlike a stack RNN, the stack transformer cannot model the entire class of deterministic context-free languages.

## Acknowledgements

This publication was made possible by an ETH AI Center doctoral fellowship to Jiaoda Li. Ryan Cotterell acknowledges support from the Swiss National Science Foundation (SNSF) as part of the “The Forgotten Role of Inductive Bias in Interpretability” project.

## Limitations

The primary limitation of the proposed stack attention is it only allows one POP operation at a time. It can be extended to have multiple POPs in a manner similar to Yogatama et al. (2018). It can also only handle deterministic context-free languages. We would like to extend it to non-deterministic stacks in future works. Although our method does not require structural supervision, it can in principle take advantage of it when it is available. In such cases,

the model can be fully parallelized, leading to great improvement in time efficiency. It would be interesting to explore this possibility in the future.

## Ethical Considerations

We foresee no ethical concerns with this work.

## References

- Dana Angluin, David Chiang, and Andy Yang. 2023. [Masked hard-attention transformers and boolean RASP recognize exactly the star-free languages](#). *Computing Research Repository*, arXiv:2310.13897, Version 2.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#). *Computing Research Repository*, arXiv:1607.06450.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *International Conference on Learning Representations*.
- Pablo Barcelo, Alexander Kozachinskiy, Anthony Widjaja Lin, and Vladimir Podolskii. 2024. [Logical languages accepted by transformer encoders with hard attention](#). In *International Conference on Learning Representations*.
- Satwik Bhattamishra, Kabir Ahuja, and Navin Goyal. 2020. [On the Ability and Limitations of Transformers to Recognize Formal Languages](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7096–7116.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901.
- Jianpeng Cheng, Li Dong, and Mirella Lapata. 2016. [Long short-term memory-networks for machine reading](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 551–561, Austin, Texas. Association for Computational Linguistics.
- David Chiang, Peter Cholak, and Anand Pillay. 2023. [Tighter bounds on the expressivity of transformer encoders](#). In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org.
- Noam Chomsky. 1956. [Three models for the description of language](#). *IRE Transactions on Information Theory*, 2(3):113–124.
- Noam Chomsky. 1957. *Syntactic Structures*. De Gruyter Mouton, Berlin, Boston.
- Noam Chomsky and Marcel Paul Schützenberger. 1963. [The algebraic theory of context-free languages](#). *Studies in logic and the foundations of mathematics*, 35:118–161.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988.
- Sreerupa Das, C. Lee Giles, and Guo-Zheng Sun. 1992. [Learning context-free grammars: Capabilities and limitations of a recurrent neural network with an external stack memory](#). In *Proceedings of The Fourteenth Annual Conference of Cognitive Science Society*, volume 14.
- Grégoire Delétang, Anian Ruoss, Jordi Grau-Moya, Tim Genewein, Li Kevin Wenliang, Elliot Catt, Chris Cundy, Marcus Hutter, Shane Legg, Joel Veness, and Pedro A. Ortega. 2023. [Neural networks and the Chomsky hierarchy](#). In *11th International Conference on Learning Representations*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Brian DuSell and David Chiang. 2020. [Learning context-free languages with nondeterministic stack RNNs](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 507–519.
- Brian DuSell and David Chiang. 2022. [Learning hierarchical structures with differentiable nondeterministic stacks](#). In *International Conference on Learning Representations*.
- Brian DuSell and David Chiang. 2024. [Stack attention: Improving the ability of transformers to model hierarchical patterns](#). In *International Conference on Learning Representations*.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. [Recurrent neural network grammars](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 199–209, San Diego, California. Association for Computational Linguistics.
- Javid Ebrahimi, Dhruv Gelda, and Wei Zhang. 2020. [How can self-attention networks recognize Dyck-n languages?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4301–4306.
- R. James Evey. 1963. [Application of pushdown-store machines](#). In *Proceedings of the November 12-14, 1963, Fall Joint Computer Conference, AFIPS ’63*

- (Fall), page 215–227, New York, NY, USA. Association for Computing Machinery.
- Guhao Feng, Bohang Zhang, Yuntian Gu, Haotian Ye, Di He, and Liwei Wang. 2023. [Towards revealing the mystery behind chain of thought: A theoretical perspective](#). In *Advances in Neural Information Processing Systems*.
- Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. 2015. [Learning to transduce with unbounded memory](#). In *Advances in Neural Information Processing Systems*, volume 28.
- Michael Hahn. 2020. [Theoretical limitations of self-attention in neural sequence models](#). *Transactions of the Association for Computational Linguistics*, 8:156–171.
- Yiding Hao, Dana Angluin, and Robert Frank. 2022. [Formal language recognition by hard attention transformers: Perspectives from circuit complexity](#). *Transactions of the Association for Computational Linguistics*, 10:800–810.
- Yiding Hao, William Merrill, Dana Angluin, Robert Frank, Noah Amsel, Andrew Benz, and Simon Mendelsohn. 2018. [Context-free transductions with neural stacks](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 306–315.
- Armand Joulin and Tomas Mikolov. 2015. [Inferring algorithmic patterns with stack-augmented recurrent nets](#). In *Advances in Neural Information Processing Systems*, volume 28.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *International Conference on Learning Representations*.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [RoBERTa: A robustly optimized BERT pretraining approach](#).
- Mitchell Marcus, Grace Kim, Mary Ann Marcinkiewicz, Robert MacIntyre, Ann Bies, Mark Ferguson, Karen Katz, and Britta Schasberger. 1994. [The Penn Treebank: Annotating predicate argument structure](#). In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Robert McNaughton and Seymour Papert. 1971. *Counter-free automata*. Research monograph 65. MIT Press.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *International Conference on Learning Representations*.
- William Merrill and Ashish Sabharwal. 2023. [A logic for expressing log-precision transformers](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 52453–52463. Curran Associates, Inc.
- William Merrill and Ashish Sabharwal. 2024. [The expressive power of transformers with chain of thought](#). In *International Conference on Learning Representations*.
- Michael C. Mozer and Sreerupa Das. 1992. [A connectionist symbol manipulator that discovers the structure of context-free languages](#). In *Advances in Neural Information Processing Systems*, volume 5, pages 863–870.
- Shikhar Murty, Pratyusha Sharma, Jacob Andreas, and Christopher Manning. 2023. [Pushdown layers: Encoding recursive structure in transformer language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3233–3247, Singapore. Association for Computational Linguistics.
- Anthony G. Oettinger. 1961. [Automatic syntactic analysis and the pushdown store](#). In *Proceedings of Symposia in Applied Mathematics*, volume 12: Structure of Language and Its Mathematical Aspects, pages 104–129, Providence, Rhode Island.
- Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. [A decomposable attention model for natural language inference](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255, Austin, Texas. Association for Computational Linguistics.
- Jorge Pérez, Pablo Barceló, and Javier Marinkovic. 2021. [Attention is Turing-complete](#). *Journal of Machine Learning Research*, 22(75):1–35.
- Jordan B. Pollack. 1991. [The induction of dynamical recognizers](#). *Machine Learning*, 7(2–3):227–252.
- Ofir Press, Noah Smith, and Mike Lewis. 2022. [Train short, test long: Attention with linear biases enables input length extrapolation](#). In *International Conference on Learning Representations*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Laurent Sartran, Samuel Barrett, Adhiguna Kuncoro, Miloš Stanojević, Phil Blunsom, and Chris Dyer. 2022. [Transformer grammars: Augmenting transformer language models with syntactic inductive biases at scale](#). *Transactions of the Association for Computational Linguistics*, 10:1423–1439.

- Marcel Paul Schützenberger. 1963. [On context-free languages and push-down automata](#). *Information and Control*, 6(3):246–264.
- M.P. Schützenberger. 1965. [On finite monoids having only trivial subgroups](#). *Information and Control*, 8(2):190–194.
- Stuart M. Shieber. 1987. *Evidence Against the Context-Freeness of Natural Language*, pages 320–334. Springer Netherlands.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2023. [RoFormer: Enhanced transformer with rotary position embedding](#). *Neurocomputing*, page 127063.
- Mirac Suzgun, Sebastian Gehrmann, Yonatan Belinkov, and Stuart M. Shieber. 2019. [Memory-augmented recurrent neural networks can learn generalized Dyck languages](#). *Computing Research Repository*, arXiv:1911.03329.
- Anej Svete and Ryan Cotterell. 2024. Transformers can represent  $n$ -gram language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics*, Mexico City, Mexico. Association for Computational Linguistics.
- The Univalent Foundations Program. 2013. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2021. [Thinking like transformers](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11080–11090. PMLR.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45.
- Dani Yogatama, Yishu Miao, Gabor Melis, Wang Ling, Adhiguna Kuncoro, Chris Dyer, and Phil Blunsom. 2018. [Memory architectures in recurrent neural network language models](#). In *International Conference on Learning Representations*.
- Zheng Zeng, Rodney M. Goodman, and Padhraic Smyth. 1994. [Discrete recurrent neural networks for grammatical inference](#). *IEEE Transactions on Neural Networks*, 5(2):320–330.

## A Proof

**Theorem 3.1.** Let  $v_1, \dots, v_N$  be a series of stack operations where  $v_i \in \{\text{PUSH}_i(\cdot), \text{NO-OP}(\cdot), \text{POP}(\cdot)\}$  for all  $i \in [N]$ . Furthermore, suppose  $\mathbf{a}_i(v_i) = 1$  for all  $i \in [N]$ . Then,  $\llbracket \text{PEEK}(v_i(\dots v_1(\varepsilon))) \rrbracket = \alpha_i$  for all  $i \in \{0\} \cup [N]$ .

*Proof.*

**Base case** ( $i = 0$ ). The stack is initialized to be empty, i.e.,  $\gamma_0 = \varepsilon$  and  $\text{PEEK}(\gamma_0) = 0$ . By definition, we have

$$\alpha_0 = [1, 0, \dots]^\top = \llbracket \text{PEEK}(\gamma_0) \rrbracket \quad (28)$$

**Inductive Step.** Suppose there exists an  $i > 0$ , such that  $\forall i' < i$ ,  $\alpha_{i'} = \llbracket \text{PEEK}(\gamma_{i'}) \rrbracket$ , and  $\mathbf{a}_i(v_i) = 1$ .

- If  $v_i = \text{PUSH}$ ,  $\gamma_i = \text{PUSH}(\gamma_{i-1}) = \gamma_{i-1}i$  and  $\text{PEEK}(\gamma_i) = i$ , so according to Eq. (20) we have  $\alpha_i = \llbracket \text{PEEK}(\gamma_i) \rrbracket$ .
- If  $v_i = \text{NO-OP}$ ,  $\gamma_i = \text{NO-OP}(\gamma_{i-1}) = \gamma_{i-1}$ , and  $\alpha_i = \alpha_{i-1}$ . Since  $\alpha_{i-1} = \llbracket \text{PEEK}(\gamma_{i-1}) \rrbracket$ , we have  $\alpha_i = \llbracket \text{PEEK}(\gamma_i) \rrbracket$ .
- If  $v_i = \text{POP}$ ,

$$\alpha_i = \sum_{j=1}^{i-1} \alpha_{i-1}(j) \alpha_{j-1} + \alpha_{i-1}(0) \alpha_0 \quad (29)$$

If  $\alpha_{i-1}(0) = 1$ , i.e.  $\gamma_{i-1}$  is empty,  $\gamma_i = \text{POP}(\varepsilon) = \varepsilon$ , and

$$\alpha_i = \alpha_{i-1}(0) \alpha_0 \quad (30a)$$

$$= \alpha_0 \quad (30b)$$

$$= \llbracket 0 \rrbracket \quad (30c)$$

$$= \llbracket \text{PEEK}(\gamma_i) \rrbracket \quad (30d)$$

Otherwise,

$$\alpha_i = \sum_{j=1}^{i-1} \alpha_{i-1}(j) \alpha_{j-1} \quad (31a)$$

$$= \alpha_{\text{PEEK}(\gamma_{i-1})-1} \quad (31b)$$

$$= \llbracket \text{PEEK}(\gamma_{\text{PEEK}(\gamma_{i-1})-1}) \rrbracket \quad (31c)$$

$$= \llbracket \text{PEEK}(\text{POP}(\gamma_{i-1})) \rrbracket \quad (31d)$$

$$= \llbracket \text{PEEK}(\gamma_i) \rrbracket \quad (31e)$$

One can understand Eq. (31d) intuitively as follows:  $\gamma_{\text{PEEK}(\gamma_{i-1})-1}$  is the stack right before the current stack top  $\text{PEEK}(\gamma_{i-1})$  is pushed, so the stack top at  $\text{PEEK}(\gamma_{i-1}) - 1$  is the second top-most element at  $i - 1$ , i.e.,  $\text{PEEK}(\gamma_{\text{PEEK}(\gamma_{i-1})-1}) = \text{POP}(\gamma_{i-1})$ . ■

**Theorem 3.2.** Consider a sequence of stack attention mechanisms  $\alpha_0, \dots, \alpha_N$ . Then,  $\sum_{n=0}^N \alpha_i(n) = 1$  for all  $i \in \{0\} \cup [N]$ .

*Proof.*

**Base case** It holds for  $\alpha_0 = [1, 0, \dots]^\top$ .

**Induction step** Suppose there exists an  $i > 0$ , such that  $\forall i' < i, \sum_{n=0}^N \alpha_{i'}(n) = 1$ .

- PUSH. Obviously,

$$\sum_{n=0}^N \alpha_i^{(\text{PUSH})}(n) = 1 \quad (32)$$

- NO-OP. Since  $\alpha_i^{(\text{NO-OP})} = \alpha_{i-1}$ , we also have

$$\sum_{n=0}^N \alpha_i^{(\text{NO-OP})}(n) = 1 \quad (33)$$

- POP.

$$\sum_{n=0}^N \alpha_i^{(\text{POP})}(j) = \sum_{n=0}^N \left( \sum_{j=1}^{i-1} \alpha_{i-1}(j) \alpha_{j-1} + \alpha_{i-1}(0) \alpha_0 \right) (n) \quad (34a)$$

$$= \sum_{n=0}^N \left( \sum_{j=1}^{i-1} \alpha_{i-1}(j) \alpha_{j-1}(n) + \alpha_{i-1}(0) \alpha_0(n) \right) \quad (34b)$$

$$= \sum_{j=1}^{i-1} \alpha_{i-1}(j) \left( \sum_{n=0}^N \alpha_{j-1}(n) \right) + \alpha_{i-1}(0) \left( \sum_{n=0}^N \alpha_0(n) \right) \quad (34c)$$

$$= \sum_{j=1}^{i-1} \alpha_{i-1}(j) + \alpha_{i-1}(0) \quad (34d)$$

$$= \sum_{j=0}^{i-1} \alpha_{i-1}(j) \quad (34e)$$

$$= 1 \quad (34f)$$

Therefore,

$$\sum_{n=0}^N \alpha_i(n) \quad (35a)$$

$$= \sum_{n=0}^N \left( \sum_{a \in \mathcal{A}} \mathbf{a}_i(a) \alpha_i^{(a)} \right) (n) \quad (35b)$$

$$= \sum_{a \in \mathcal{A}} \mathbf{a}_i(a) \sum_{n=0}^N \alpha_i^{(a)}(n) \quad (35c)$$

$$= \sum_{a \in \mathcal{A}} \mathbf{a}_i(a) \quad (35d)$$

$$= 1 \quad (35e)$$

■

## B Experimental Setup

### B.1 DCF Tasks

The model is trained using the Adam optimizer (Kingma and Ba, 2015) with a learning rate of  $1e^{-4}$ , which we find works well for all the tasks. On the RS and SM tasks, we use a batch size of 32 and we train the model for 100,000 steps. On the MA and SE tasks, the batch size and the number of training steps are increased to 128 and 1,000,000, respectively, to ensure sufficient training. Each experiment is run 5 times with different random seeds. Means and variances of accuracies are reported Tab. 1 and Tab. 2.

## B.2 Language Modeling

The texts in the datasets are grouped into chunks of 128 tokens. Each model is, again, trained using the Adam optimizer for a maximum of 100 epochs with early stopping applied when the validation loss has not improved for 5 epochs in a row. We tune the learning rate from  $\{1e^{-5}, 2e^{-5}, 1e^{-4}, 2e^{-4}\}$  on the validation set, and choose  $2e^{-5}$  that leads to the best validation performance. Results on the test set over 5 runs with different random seeds are reported in Tab. 3. Experiments are conducted on a single NVIDIA Tesla V100 GPU.

## C Positional Encodings

We consider five different commonly used positional encodings:

- none. No positional encodings are used.
- sincos. The sinusoidal positional encodings used in vanilla transformer (Vaswani et al., 2017). Positional information encoded sinusoidally is added to the embeddings.
- relative. In Transformer-XL (Dai et al., 2019), relative rather than absolute sinusoidal positional information is added to the keys and queries of each attention block.
- rotary. Introduced by Su et al. (2023) and popularized by GPT-3 (Brown et al., 2020), rotary positional encodings multiply the keys and queries by sinusoidal encodings.
- ALiBi. Press et al. (2022) adds linear biases to the attention blocks that favor the more recent tokens.

## D Analysis: Attention Maps

An advantage of our stack attention mechanism is that we can visualize the stack tops  $\alpha_i$ , which provides greater interpretability than methods where stack tops are mixtures of hidden states (Joulin and Mikolov, 2015). We run a set of toy experiments with the stack transformer in the MLM setting. We randomly select one test example for each task.

**RS.** At the first two layers (Fig. 2a, Fig. 2b), the first 5 tokens attend to themselves while the [MASK] tokens attend to the last token in  $x$ . The most probable sequence of operations that leads to such a stack attention map is the input  $x$  is pushed one by one onto the stack and NO-OP is performed on all the [MASK] tokens. At the third layer (Fig. 2c), the stacks for the [MASK] tokens shift one position backward at a time, which demonstrates the stack elements are POPed one by one to generate the outputs. At the last two layers, all the tokens attend to themselves, so the stacks can be regarded as being skipped (Fig. 2d, Fig. 2e).

**SM.** Looking at the attention map at the first layer (Fig. 3a), we can infer the operations taken by the stack as follows: the stack first PUSHes the initial stack contents (ab); once the [POP] operation is read, it reverts to the first element a; then it performs the operation [PUSH b] twice as instructed; afterwards, it POPs the final stack bba for outputs. The stack attention correctly skips the b at timestep 1 as it has already been POPed at timestep 2. The last three positions are [PAD] tokens and can be ignored.

**MA and SE.** We also provide an attention map for MA and SE in Fig. 4 and Fig. 5. Their attention maps are less interpretable as the stack transformer does not learn them well. Nevertheless, we can still observe that the stacks seem to be able to match the parentheses, which matches our expectations of the stack’s strengths. For MA, at the first layer (Fig. 4a), the stack successfully matches the last two closing parentheses (at timestep 8 and 9) with their corresponding open parentheses (at timestep 5 and 0 respectively). For SE, the pattern is less obvious presumably because the parentheses do not have an impact on the order of arithmetic operations and can be ignored.

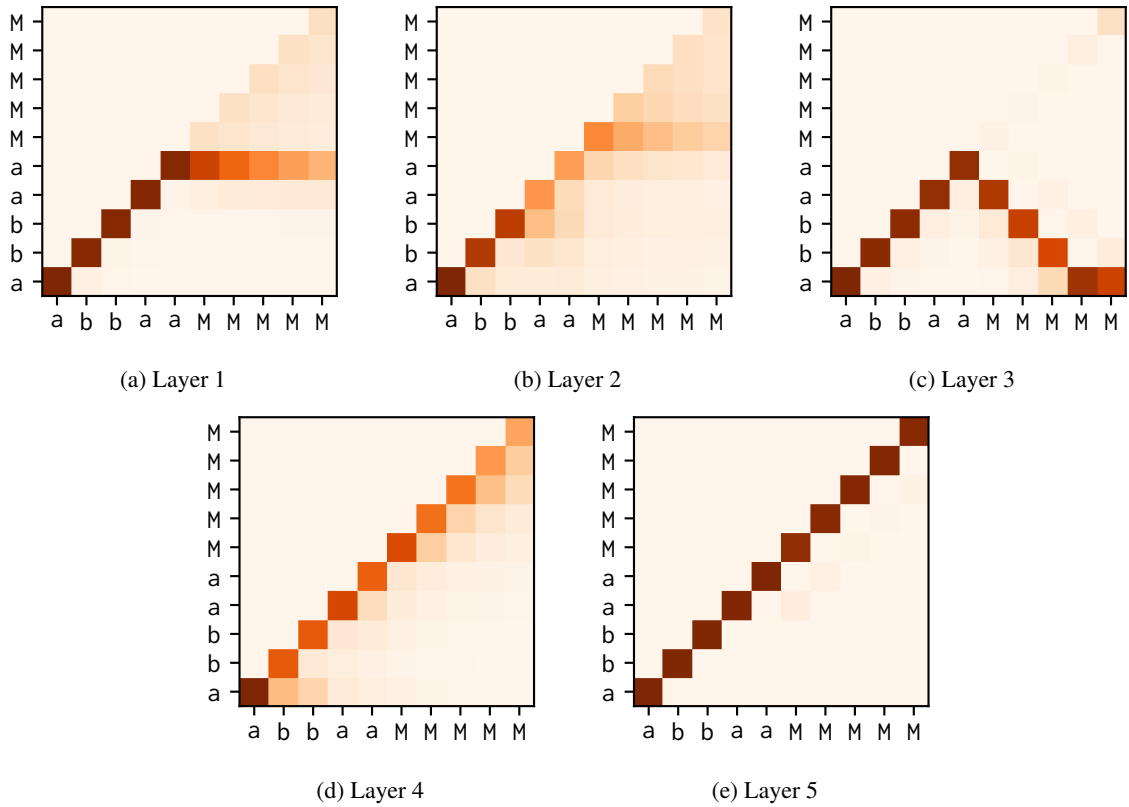


Figure 2: Stack attention maps at different layers for RS. The input  $x$  is abbaa. M represents a [MASK] token.

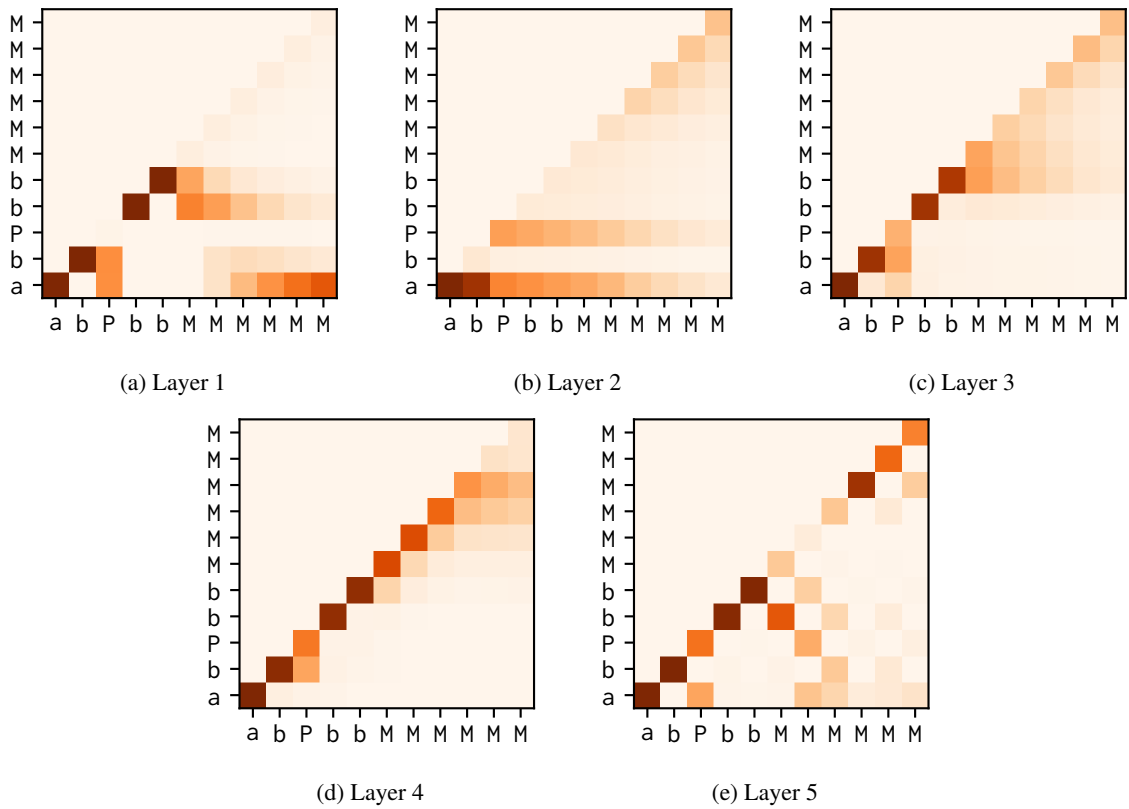


Figure 3: Stack attention maps at different layers for SM. The input  $x$  is ab[POP][PUSH a][PUSH b]. In the graphs, [PUSH a], [PUSH b], and [POP] are abbreviated as a, b, and P respectively. M represents a [MASK] token. The correct output should be bba followed by [PAD] tokens.



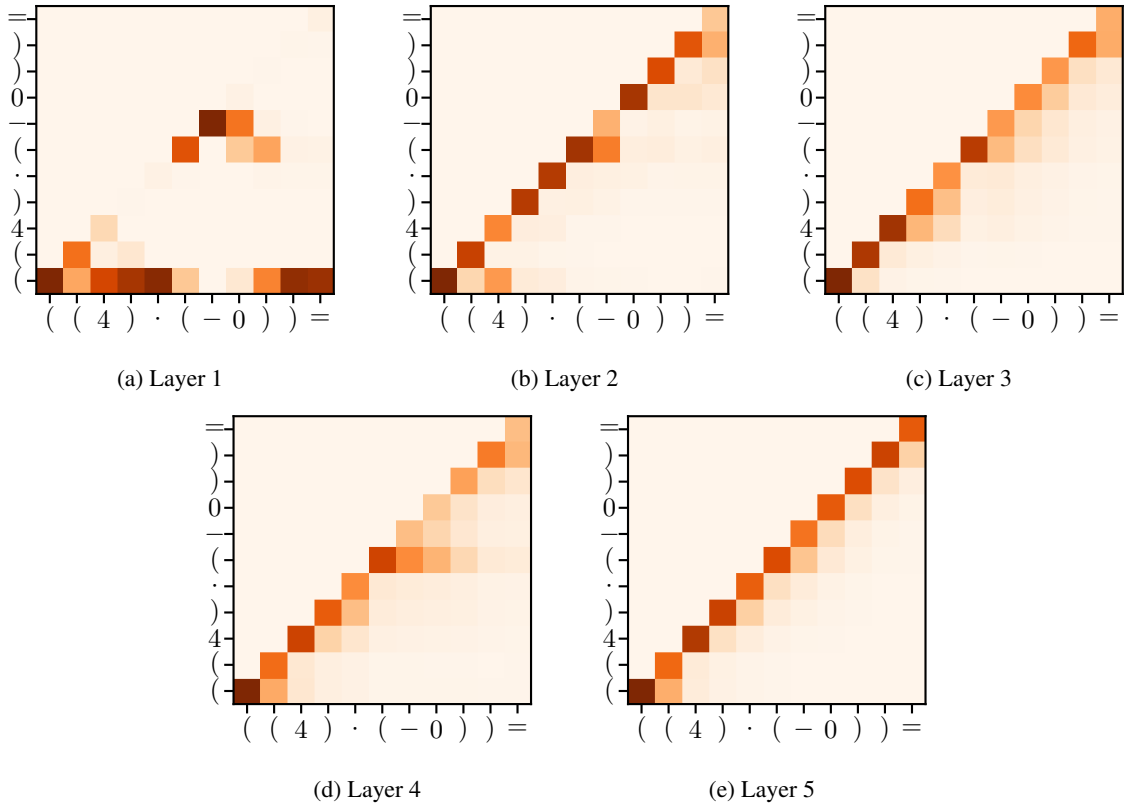


Figure 4: Stack attention maps at different layers for MA. The input  $x$  is  $((4) \cdot (-0)) =$ .

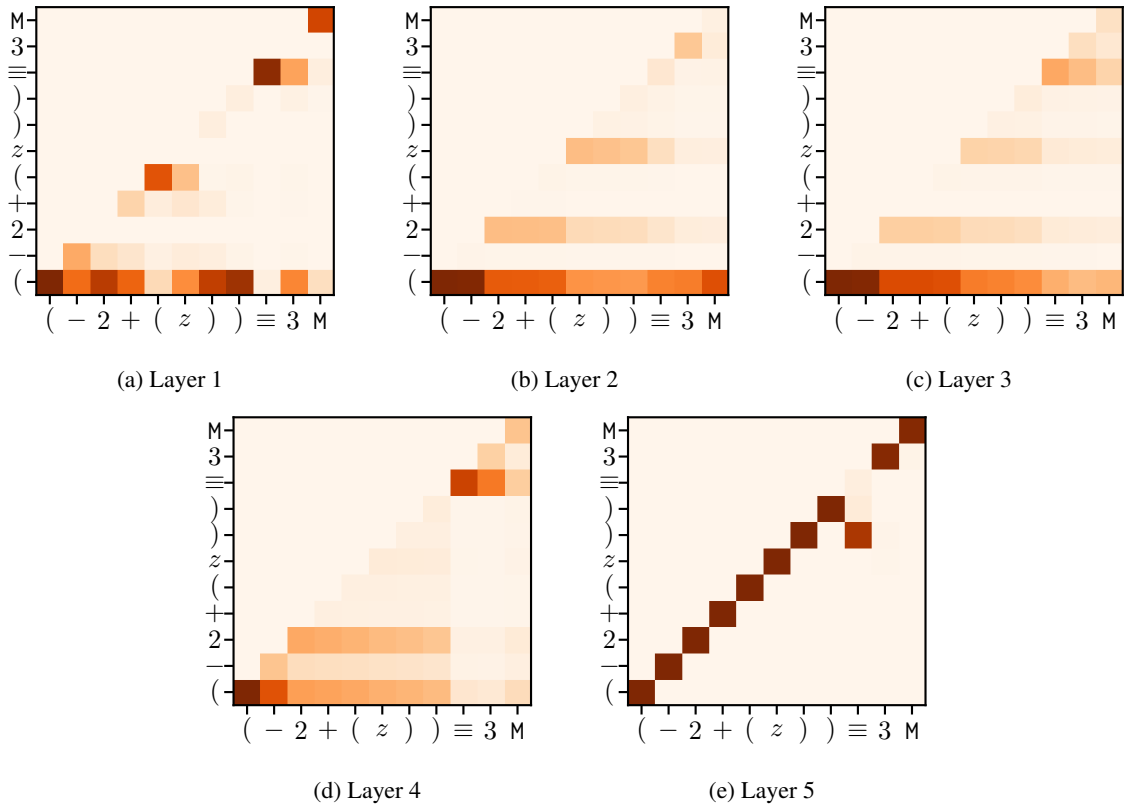


Figure 5: Stack attention maps at different layers for SE. The input  $x$  is  $(1 + (-z)) = 3$ .  $M$  represents a [MASK] token.

## E Related Work

### E.1 Stack Augmentation

Equipping a neural network with a data structure such as an external stack to enhance its ability to recognize context-free languages has been extensively investigated in previous works (Pollack, 1991; Das et al., 1992; Mozer and Das, 1992; Zeng et al., 1994). The idea has seen a resurgence in recent years, with work focusing primarily on recurrent networks (Joulin and Mikolov, 2015; Grefenstette et al., 2015; Hao et al., 2018; Yogatama et al., 2018; Suzgun et al., 2019; DuSell and Chiang, 2020, 2022). Joulin and Mikolov (2015) propose to superpose the result of applying each stack operation at each step, which directly inspires our work. We adapt it for application to transformers by rendering this concept as an attention mechanism. In that sense, our work is related to Das et al. (1992) and Grefenstette et al. (2015), which also assign weights to stack elements. Our stack attention mechanism is different as the stack attention weights are assigned to previously seen tokens indicating where the top element is located

Sartran et al. (2022) and Murty et al. (2023) incorporate a stack mechanism into a transformer language model with structural supervision during training. DuSell and Chiang’s (2024) contemporaneous work also augments a transformer language model with a stack. Both their and our methods are named stack attention, but their stack attention is an attention mechanism over stack actions while ours is an attention mechanism over input tokens.

### E.2 Expressivity of Transformers

The expressivity of transformers under various assumptions has been extensively studied. A stream of research considers transformer encoders with a classification layer at the end as recognizers. Hahn (2020) proves that transformers cannot recognize parity language, a periodic language of binary strings with an even number of 1’s, and Dyck-2 language, a CF language of balanced brackets of two types. Bhattamishra et al. (2020) find that transformers can recognize certain counter languages but fail to recognize non-star-free languages such as  $(aa)^*$ . Svete and Cotterell (2024) show that transformers can represent  $n$ -gram language models. Hao et al. (2022), Chiang et al. (2023), Merrill and Sabharwal (2023), Barcelo et al. (2024), and Angluin et al. (2023) relate transformers to circuit complexity and formal logic. With various extensions, transformers’ expressivity can be increased. Weiss et al. (2021) propose a programming language that shares the same basic operations with transformers but is more expressive than standard transformers. Pérez et al. (2021) and Merrill and Sabharwal (2024) show that transformer encoder–decoders and decoders are Turing complete with additional scratch space.