
(Un)interpretability of Transformers: a case study with Dyck grammars

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Interpretability of Transformers is an emerging topic which aims to understand the
2 algorithm implemented by a learned Transformer by peering and probing individual
3 aspects of the model, such as the weight matrices or the attention patterns. In
4 this work, through a combination of theoretical results and carefully controlled
5 experiments on synthetic data, we take a critical view of methods that exclusively
6 focus on individual parts of the model, rather than consider the network as a whole.
7 We consider a simple synthetic setup of learning a Dyck language. Theoretically,
8 we show that the set of models that can solve this task (exactly or approximately)
9 satisfy a structural characterization derived from ideas in formal languages (the
10 pumping lemma). We use this characterization to show that the set of optima is
11 qualitatively rich: in particular, the attention pattern of a single layer can be “nearly
12 randomized”, while preserving the functionality of the network. We also show via
13 extensive experiments that these constructions are not merely a theoretical artifact:
14 even with severe constraints to the architecture of the model, vastly different
15 solutions can be reached via standard training. Thus, interpretability claims based
16 on individual heads or weight matrices in the Transformer can be misleading.

17 1 Introduction

18 Transformer-based models power many leading approaches to natural language processing. With
19 their growing deployment in various applications, it is increasingly essential to understand the inner
20 working of these models. Towards addressing this, there have been great advancement in the field
21 of interpretability presenting various types of evidence (Clark et al., 2019; Vig & Belinkov, 2019;
22 Wiegrefe & Pinter, 2019; Nanda et al., 2023; Wang et al., 2023), some of which, however, can be
23 misleading despite being highly intuitive (Jain & Wallace, 2019; Serrano & Smith, 2019; Rogers
24 et al., 2020; Grimsley et al., 2020; Brunner et al., 2020; Meister et al., 2021).

25 In this work, we aim to understand the theoretical limitation of different interpretability methods by
26 characterizing the set of viable solutions. We focus on a particular toy setup in which Transformers
27 are trained to generate *Dyck grammars*, a classic type of formal language grammar consisting of
28 balanced parentheses of multiple types. Dyck is a useful sandbox, as it captures properties like
29 long-range dependency and hierarchical tree-like structure that commonly appear in natural and
30 programming language syntax, and has been an object of interest in many theoretical studies (Hahn,
31 2020; Yao et al., 2021; Liu et al., 2022b, 2023). Dyck is canonically parsed using a stack-like data
32 structure. Such stack-like patterns (Figure 1) have been observed in the attention heads (Ebrahimi
33 et al., 2020), which is later formalized by Yao et al. (2021).

34 From a representational perspective and via explicit constructions of Transformer weights, recent
35 works (Liu et al., 2023; Li et al., 2023) show that Transformers are sufficiently expressive to admit
36 very different solutions that perform equally well on the training distribution. This calls into question:

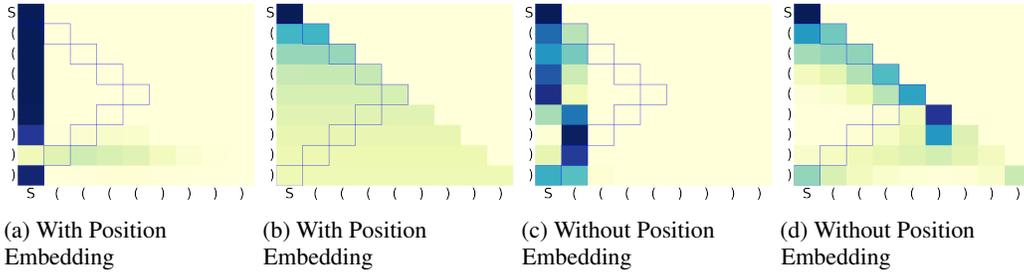


Figure 1: **Second-layer attention patterns of two-layer Transformers on Dyck**: typical attention patterns do *not* exactly match the intuitively interpretable stack-like pattern prescribed in Ebrahimi et al. (2020); Yao et al. (2021). The blue boxes indicate the locations of the last unmatched open brackets, as they would appear in a stack-like pattern. All models reach $\geq 97\%$ accuracy (defined in Section 4.1). In the heatmap, darker color indicates larger value.

- 37 (Q1) Do Transformer solutions found empirically match the theoretical constructions given in these
 38 representational results (Figure 1)? In particular, are interpretable stack-like pattern in Ebrahimi
 39 et al. (2018) the norm or the exception in practice?
- 40 (Q2) More broadly, can we understand in a principled manner the fundamental obstructions to reliably
 41 “reverse engineering” the algorithm implemented by a Transformer by looking at individual
 42 attention patterns?
- 43 (Q3) Among models that perform (near-)optimally on the training distribution, even if we cannot fully
 44 reverse engineer the algorithm implemented by the learned solutions, can we identify properties
 45 that characterize performance beyond the training distribution?

46 **Our contributions.** We first prove several theoretical results to provide evidence for why individual
 47 components (e.g. attention patterns or weights) of a Transformer should not be expected to be
 48 interpretable. In particular, we prove:

- 49 • A **perfect balance** condition (Theorem 1) on the attention pattern that is sufficient and necessary
 50 for 2-layer Transformers with a *minimal first layer* (Assumption 1) to predict optimally on Dyck of
 51 *any* length. We then show that this condition permits abundant *non-stack-like* attention patterns
 52 that do not necessarily reflect any structure of the task, including *uniform* attentions (Corollary 1).
- 53 • An **approximate balance** condition (Theorem 2), the *near-optimal* counterpart of the condition
 54 above, for predicting on *bounded-length* Dyck. Likewise, non-stack-like attention patterns exist.
- 55 • **Indistinguishability from a single component** (Theorem 3), proved via a *Lottery Ticket Hypothesis*
 56 style argument that any Transformer can be approximated by pruning a larger random Transformer,
 57 implying that interpretations based exclusively on local components may be unreliable.

58 We further accompany these theoretical findings with an extensive set of empirical investigations.

59 *Is standard training biased towards interpretable solutions?* While both stack-like and non-stack like
 60 patterns can process Dyck theoretically, the inductive biases of the architecture or the optimization
 61 process may prefer one solution over the other in practice. In Section 4.1, based on a wide range
 62 of Dyck distributions and model architecture ablations, we find that Transformers that generalize
 63 near-perfectly in-distribution (and reasonably well out-of-distribution) do *not* typically produce
 64 stack-like attention patterns, showing that the results reported in prior work (Ebrahimi et al., 2018)
 65 should not be expected from standard training.

66 *Do non-interpretable solutions perform well in practice?* Our theory predicts that balanced (or even
 67 uniform) attentions suffice for good in- and out-of-distribution generalization. In Section 4.2, we
 68 empirically verify that with standard training, the extent to which attentions are balanced is positively
 69 correlated with generalization performance. Moreover, we can guide Transformers to learn more
 70 balanced attention by regularizing for the balance condition, leading to better generalization.

71 **1.1 Related Works**

72 There has been a flourishing line of work on interpretability in natural language processing. Multiple
 73 “probing” tasks have been designed to extract syntactic or semantic information from the learned
 74 representations (Raganato & Tiedemann, 2018; Liu et al., 2019; Hewitt & Manning, 2019; Clark et al.,
 75 2019). However, the effectiveness of probing often intricately depend on the architecture choices
 76 and task design, and sometimes may even result in misleading conclusions (Jain & Wallace, 2019;
 77 Serrano & Smith, 2019; Rogers et al., 2020; Brunner et al., 2020; Prasanna et al., 2020; Meister et al.,
 78 2021). While these challenges do not completely invalidate existing approaches (Wiegrefe & Pinter,
 79 2019), it does highlight the need for more fundamental understanding of interpretability.

80 Towards this, we choose to focus on the synthetic setup of Dyck whose solution space is easier to
 81 characterize than natural languages, allowing us to identify a set of feasible solutions. While similar
 82 representational results have been studied in prior work (Yao et al., 2021; Liu et al., 2023; Zhao et al.,
 83 2023), our work emphasizes that theoretical constructions do not resemble the solutions found in
 84 practice. Moreover, the multiplicity of valid constructions suggest that understanding Transformer
 85 solutions require analyzing the optimization process, which a number of prior work has made progress
 86 on (Jelassi et al., 2022; Li et al., 2023; Deng et al., 2023).

87 Finally, it is worth noting that the challenges highlighted in our work do not contradict the line of
 88 prior works that aim to improve *mechanistic interpretability* into a trained model or the training
 89 process (Elhage et al., 2021; Olsson et al., 2022; Nanda et al., 2023; Li et al., 2023), which aim to
 90 develop circuit-level understanding of a particular model or the training process.

91 We defer discussions on additional related works to Appendix A.

92 **2 Problem Setup**

93 **Dyck languages** A Dyck language (Schützenberger, 1963) is generated by a context-free grammar,
 94 where the valid strings consist of balanced brackets of different types (for example, “[()]” is valid
 95 but “[()]” is not). Dyck_k denote the Dyck language defined on k types of brackets. The alphabet of
 96 Dyck_k is denoted as $\{1, 2, \dots, 2k\} \equiv [2k]$, where for each type $t \in [k]$, tokens $2t - 1$ and $2t$ are a
 97 pair of corresponding open and closed brackets. Dyck languages can be recognized by a push-down
 98 automaton. For a string w and $i \leq j \in \mathbb{Z}_+$, we use $w_{i:j}$ to denote the substring of w between position
 99 i and position j (both ends included). For a valid prefix $w_{1:i}$, the *grammar depth* of $w_{1:i}$ is defined as
 100 the depth of the stack after processing $w_{1:i}$:

$$\text{depth}(w_{1:i}) = \#\text{Open Brackets in } w_{1:i} - \#\text{Closed Brackets in } w_{1:i}.$$

101 We overload the same notation $\text{depth}(w_{1:i})$ to also denote the grammar depth of the bracket at
 102 position i . For example, in each pair of matching brackets, the closing bracket is one depth smaller
 103 than the open bracket. We will use $\tau_{i,d}$ to denote a token of type $i \in [2k]$ placed at grammar depth
 104 $d \in \mathbb{N}$.

105 We consider bounded-depth Dyck languages following Yao et al. (2021). Specifically, $\text{Dyck}_{k,D}$ is a
 106 subset of Dyck_k such that the depth of any prefix of a word is bounded by D ,

$$\text{Dyck}_{k,D} := \{w_{1:n} \in \text{Dyck}_k \mid \max_{i \in [n]} \text{depth}(w_{1:i}) \leq D\}.$$

107 While a bounded grammar depth might seem restrictive, it suffices to capture many practical settings.
 108 For example, the level of recursion occurring in natural languages is typically bounded by a small
 109 constant (Karlsson, 2007; Jin et al., 2018). We further define the *length- N prefix set* of $\text{Dyck}_{k,D}$ as

$$\text{Dyck}_{k,D,N} = \{w_{1:N} \mid \exists n \geq N, w_{N+1:n} \in [2k]^{n-N}, \text{ s.t. } w_{1:n} \in \text{Dyck}_{k,D}\}. \quad (1)$$

110 Our theoretical setup uses the following data distribution $\mathcal{D}_{q,k,D,N}$:

111 **Definition 1** (Dyck distribution). *The distribution $\mathcal{D}_{q,k,D,N}$, specified by $q \in (0, 1)$, is defined over*
 112 $\text{Dyck}_{k,D,N}$ such that $\forall w_{1:N} \in \text{Dyck}_{k,D,N}$,

$$\mathbb{P}(w_{1:N}) \propto (q/k)^{\#\{i \mid w_i \text{ is open, } \text{depth}(w_{1:i}) > 1\}} \cdot (1 - q)^{\#\{i \mid w_i \text{ is closed, } \text{depth}(w_{1:i}) < D - 1\}}. \quad (2)$$

113 That is, $q \in (0, 1)$ denote the probability of seeing an open bracket at the next position, except for
 114 two corner cases: 1) the next bracket has to be open if the current grammar depth is 0 (1 after seeing
 115 the open bracket); 2) the next bracket has to be closed if the current grammar depth is D .

116 **Training Objectives.** Given a model f_θ parameterized by θ , we train with a *next-token prediction*
 117 language modeling objective on a given $\mathcal{D}_{q,k,D,N}$. Precisely, given a prefix $w_{1:N} \in \text{Dyck}_{k,D,N}$ and
 118 a loss function $l(\cdot, \cdot) \rightarrow \mathbb{R}$, f_θ is trained to minimize the loss function $\min_\theta \mathcal{L}_\theta(x)$ for

$$\mathcal{L}_\theta(x) = \mathbb{E}_{w_{1:N} \sim \mathcal{D}_{q,k,D,N}} \left[\frac{1}{N} \sum_{i=1}^N l(f_\theta(w_{1:i-1}), e_{w_i}) \right]. \quad (3)$$

119 We will also consider a ℓ_2 -regularized version $\mathcal{L}_\theta^{\text{reg}}(x) = \mathcal{L}_\theta(x) + \lambda \frac{\|\theta\|_2^2}{2}$ with parameter $\lambda > 0$.

120 For our theory, we will consider the mean squared error (MSE) as the loss function,

$$l := l_{sq}(x, e_i) = \|x - e_i\|_2^2. \quad (4)$$

121 In our experiments, we apply the cross entropy loss following common practice.

122 **Transformer Architectures.** We consider a general formulation of Transformer in this work: the
 123 l -th layer is parameterized by $\theta^{(l)} := \{W_Q^{(l)}, W_K^{(l)}, W_V^{(l)}, \text{param}(g^{(l)})\} \in \Theta$, where $W_K^{(l)}, W_Q^{(l)} \in$
 124 $\mathbb{R}^{m_a \times m}$, and $W_V^{(l)} \in \mathbb{R}^{m \times m}$ are the key, query, and value matrices of the attention module;
 125 $\text{param}(g^{(l)})$ are parameters of a feed-forward network $g^{(l)}$, consisting of fully connected layers,
 126 (optionally) LayerNorms and residual links. Given $X \in \mathbb{R}^{d \times N}$, the matrix of d -dimensional features
 127 on a length- N sequence, the l -th layer of a Transformer computes

$$f_l(X; \theta^{(l)}) = g^{(l)} \left(\text{LN} \left(W_V^{(l)} X \sigma \left(\underbrace{C \cdot \frac{(W_K^{(l)} X)^\top (W_Q^{(l)} X)}{\sqrt{d_a}}}_{\text{attention pattern}} \right) \right) + X \right), \quad (5)$$

128 where σ is the column-wise softmax operation defined as $\sigma(A)_{i,j} = \frac{\exp(A_{i,j})}{\sum_{k=1}^N \exp(A_{k,j})}$, LN represents
 129 column-wise LayerNorm operation defined as $\text{LN}(A)_{1:m,j} = \gamma \frac{\mathcal{P}_\perp A_{1:m,j}}{\|\mathcal{P}_\perp A_{1:m,j}\|_2} + \beta$, where \mathcal{P}_\perp denotes
 130 the projection orthogonal to the $\mathbf{1}\mathbf{1}^\top$ subspace (this is just a compact way to write the common
 131 mean subtraction operation). C is the causal mask matrix defined as $C_{i,j} = \mathbb{1}[i \leq j]$. We call
 132 $\sigma \left(C \cdot \frac{(W_K^{(l)} X)^\top (W_Q^{(l)} X)}{\sqrt{d_a}} \right)$ the *Attention Pattern* of the Transformer layer l . We consider single-head
 133 attentions in this work, whose simplicity further strengthens the messages in this work.

134 A L -layer Transformer \mathcal{T}_L consists of L above layers, and a word embedding matrix $W_E \in \mathbb{R}^{d \times 2k}$
 135 and a linear decoding head with weight $W_{\text{Head}} \in \mathbb{R}^{2k \times w}$ and bias $b_{\text{Head}} \in \mathbb{R}^{2k}$. Let $\mathcal{Z} \in \mathbb{R}^{2k \times N}$
 136 denote the one-hot embedding of a length- N sequence, then \mathcal{T}_L computes for \mathcal{Z} as

$$\mathcal{T}(\mathcal{Z}) = W_{\text{Head}} f_L(\cdots (f_1(W_E \mathcal{Z})) + b_{\text{Head}}), \quad (6)$$

137 We define the *nonstructural pruning* as:

138 **Definition 2** (Nonstructural pruning). *The nonstructural pruning of a Transformer refers to the type*
 139 *of pruning where some entries of the weight matrices are set to zero, and some LayerNorms are set*
 140 *as the identity.*

141 Note that this is as opposed to *structural pruning*, which prunes some channels of weight matrices.

142 3 Theoretical Analysis

143 Many prior works have looked for intuitive interpretations of Transformer solutions by studying
 144 the attention patterns of particular heads or some individual components of a Transformer (Clark
 145 et al., 2019; Vig & Belinkov, 2019; Dar et al., 2022). However, we show in this section why this
 146 methodology can be insufficient even for the simple setting of Dyck. Namely, for Transformers
 147 that generalize well on Dyck (both in-distribution and out-of-distribution), neither attention patterns
 148 nor individual local components are guaranteed to encode structures specific for parsing Dyck. We
 149 further argue that the converse is also insufficient: when a Transformer does produce interpretable
 150 attention patterns, there could be limitations of such interpretation as well, as discussed in Appendix B.
 151 Together, our results provide theoretical evidence that careful analyses (beyond heuristics) are required
 152 when studying interpretations from Transformer.

153 3.1 Interpretability Requires Inspecting More Than Attention Patterns

154 This section focuses on Transformers with 2 layers, which are sufficient for processing Dyck (Yao
155 et al., 2021). We will show that even under this simplified setting, attention patterns alone are
156 not sufficient for interpretation. In fact, we will further restrict the set of 2-layer Transformers by
157 requiring the first-layer outputs to only depend on information necessary for processing Dyck:

158 **Assumption 1** (Minimal First Layer). *We consider 2-layer Transformers with a minimal first layer*
159 *f_1 . That is, let $\mathbf{Z} \in \mathbb{R}^{2k \times N}$ denote the one-hot embeddings of any input sequence $t_1, \dots, t_N \in [2k]$,*
160 *then the j th column of the output $f_1(W^E \mathbf{Z})$ only depends on the type and depth of t_j , $\forall j \in [N]$.*

161 The Minimal First Layer is a strong condition, as it requires the first layer output to depend only
162 on the bracket type and depth and eliminate all other information, including positions. There are
163 multiple constructions of a minimal first layer, such as the one in Yao et al. (2021). When working
164 with a minimal first layer, we will not explicitly reason about its parameterization, but instead work
165 directly with its output. Specifically, $e(\tau_{t,d})$ the output of $\tau_{t,d}$ for $t \in [2k]$, $d \in [D]$.

166 3.1.1 Perfect Balance Condition For Ideal Generalization of Unbounded Length

167 A line of works tries to understand the model by inspecting the attention patterns (Ebrahimi et al.,
168 2018; Clark et al., 2019; Vig & Belinkov, 2019). However, we find that the attention patterns alone
169 can be too flexible to be helpful, even for the restricted class of a two-layer Transformer with a
170 minimal first layer (Assumption 1) and even on a language as simple as Dyck. In particular, the
171 second-layer attention matrix $W_K^{(2)}(W_Q^{(2)})^\top$ only needs to satisfy one condition:

172 **Theorem 1** (Perfect Balance). *Consider a two-layer Transformer \mathcal{T} using a minimal first layer with*
173 *output embeddings $\{e(\tau_{i,d})\}_{d \in [D], i \in [2k]}$. Let $\theta^{(2)} := \{W_Q^{(2)}, W_K^{(2)}, W_V^{(2)}, \text{param}(g^{(2)})\}$ denote the*
174 *second layer weights, and assume that $W_V^{(2)}$ satisfies $\mathcal{P}_\perp W_V^{(2)} e(\tau_{t,d}) \neq 0, \forall t \in [k], d \in [D]$, where*
175 *\mathcal{P}_\perp projects to the subspace orthogonal $\mathbf{1}\mathbf{1}^\top$.¹ Then, there exist $\{e(\tau_{i,d})\}$ and $\theta^{(2)}$ that minimize the*
176 *mean squared error (Eqn. 4) on Dyck $_{k,D}$ for any length N , if and only if $\forall i, j_1, j_2 \in [k], 0 \leq d' \leq$
177 *$D, 1 \leq d_1 \leq d_2 \leq D$,**

$$(e(\tau_{2i-1,d'+1}) - e(\tau_{2i,d'}))^\top (W_K^{(2)})^\top W_Q^{(2)} (e(\tau_{2j_1,d_1}) - e(\tau_{2j_2,d_2})) = 0. \quad (7)$$

178 Recall that $2i-1, 2i$ for $i \in [k]$ denote a matching pair of open and closed brackets, and
179 $e(\tau_{2i-1,d'+1}), e(\tau_{2i,d'})$ denote the corresponding first-layer outputs. Intuitively, Equation (7) says
180 that since matching brackets do not affect future predictions, their embeddings should balance out
181 each other. The balance condition Equation (7) is “perfect” in the sense that the theory assumes the
182 model can minimize the loss for any length N ; we will see an approximate version later in Theorem 2.

183 **Proof sketch: necessity of the balance condition** The key idea is reminiscent of the pumping
184 lemma. Note that in Equation (5), the attention output is directly used as the input of LayerNorm,
185 which allow us to *ignore the normalization* from the softmax operation. For any prefix p ending
186 with a closed bracket $\tau_{2i,d}$, let p_m be the prefix obtained by inserting m pair of $\{\tau_{2i-1,d'+1}, \tau_{2i,d'}\}$
187 for arbitrary $i \in [k]$ and depth $d' \in [D]$. Denote the projection of the unnormalized attention
188 output by $u(\tau_{t_1,d_1}, \tau_{t_2,d_2}) := \mathcal{P}_\perp \exp \left(e(\tau_{t_1,d_1})^\top (W_K^{(2)})^\top W_Q^{(2)} e(\tau_{t_2,d_2}) \right) W_V^{(2)} e(\tau_{t_1,d_1})$. Then,
189 by Equation (6), we have,

$$\mathcal{T}(p_m) = g^{(2)} \left(\text{LN}^{(2)} \left(v + m \left(u(\tau_{2j,d}, \tau_{2i,d'}) + u(\tau_{2j,d}, \tau_{2i-1,d'+1}) \right) \right) + e(\tau_{2j,d}) \right). \quad (8)$$

190 Suppose $u(\tau_{2j,d}, \tau_{2i,d'}) + u(\tau_{2j,d}, \tau_{2i-1,d'+1}) \neq 0$. Based on the continuity of the projection function
191 and the LayerNorm Layer, we can show that $\lim_{m \rightarrow \infty} \mathcal{T}(p_m)$ depend only on grammar depths d, d'
192 and types $2j, 2i-1, 2i$, which, however, are not sufficient to determine the next-token probability
193 from p_m , since the latter depends on the type of the last unmatched open bracket in p . This contradicts
194 the assumption that the model can minimize the loss for any length N . Hence we must have

$$u(\tau_{2j,d}, \tau_{2i,d'}) + u(\tau_{2j,d}, \tau_{2i-1,d'+1}) = 0. \quad (9)$$

¹This assumption can be intuitively understood as all tokens have nonzero contributions to the prediction because otherwise $W_V^{(2)} e(\tau_{t,d})$ will not contribute to prediction after the LayerNorm.

195 Finally, as it is assumed that $\mathcal{P}_\perp W_V^{(2)} e(\tau_{t,d}) \neq 0$, we conclude that

$$(\tau_{2i-1,d'+1} - \tau_{2i,d'})^\top (W_K^{(2)})^\top W_Q^{(2)} e(\tau_{2j+1,d}) = \ln \left(\frac{\|\mathcal{P}_\perp W_V e(\tau_{2i,d'})\|_2}{\|\mathcal{P}_\perp W_V e(\tau_{2i+1,d'-1})\|_2} \right).$$

196 This leads to our result in Theorem 1. Details and the proof of sufficiency are given in Appendix C.1.

197 The perfect balance condition does not restrict much on the attention patterns. For example, even the
198 uniform attention satisfies the condition and can solve Dyck:

199 **Corollary 1.** *There exists a two-layer Transformer with uniform attention and without position
200 embedding (but with causal mask) that can generate the Dyck language of arbitrary length.*

201 Uniform attention patterns are hardly reflective of any structure of Dyck, hence Corollary 1 proves
202 that attention patterns can be oblivious about the underlying task, violating the ‘‘faithfulness’’ criteria
203 for an interpretation (Jain & Wallace, 2019). We will further show in Appendix B.1 that empirically,
204 seemingly structured attention patterns may not accurately represent the inherent structure of the task.

205 3.1.2 Approximate Balance Condition For Finite Length Training Data

206 The condition in Theorem 1 requires the model to reach the optimal loss for data of any length.
207 However, in practice, one can only train the model on *finite-length* data and the model can only reach
208 a low but non-optimal loss for finite length data. In this case, the condition in Theorem 1 is not
209 precisely met. However, one can show that a similar condition as in Equation (9) is still necessary if
210 one restricted the Lipschitz constant of the projection function g . We first define two quantities that
211 measure the deviation from the previous ideal scenario:

$$S_{d,d',i,j}[\theta^{(2)}] = \left\| u(\tau_{2j,d}, \tau_{2i,d'}) + u(\tau_{2j,d}, \tau_{2i-1,d'+1}) \right\|_2, \quad (10)$$

$$t = \arg \min_{t \in [k]^d} \left\| \sum_{d' \leq d} u(\tau_{2j,d}, \tau_{2t_{d'},d'}) + u(\tau_{2j,d}, \tau_{2j-1,d+1}) + u(\tau_{2j,d}, \tau_{2j,d}) \right\|_2. \quad (11)$$

$$P_{d,j}[\theta^{(2)}] = \min_{t' \in [k]^d, t'_d \neq t_d} \left\| \sum_{d' \leq d} u(\tau_{2j,d}, \tau_{2t_{d'},d'}) + u(\tau_{2j,d}, \tau_{2j-1,d+1}) + u(\tau_{2j,d}, \tau_{2j,d}) \right\|_2. \quad (12)$$

212 The first term $S_{d,d',i,j}[\theta^{(2)}]$ measures the change in the input of the LayerNorm layer for the last
213 token $\tau_{2j,d}$, when a matching pair of brackets $(\tau_{2i,d'}, \tau_{2i-1,d'+1})$ is inserted into the prefix. Under
214 the perfect balance condition, $S_{d,d',i,j}[\theta^{(2)}] = 0$. The second term $P_{d,j}[\theta^{(2)}]$ is measures the norm of
215 the input of the LayerNorm layer at last token $\tau_{2j,d}$. when the prefix only contains open brackets. In
216 the following theorem, $P_{d,j}$ will be used as a baseline to show $S_{d,d',i,j}[\theta^{(2)}]$ cannot be too large, i.e.,
217 the model should not be sensitive to the insertion of a matching pair of brackets.

218 **Theorem 2** (Approximate Balance). *Consider a two-layer Transformer \mathcal{T} with a minimal first
219 layer trained with the mean squared error (Equation (4)). For any $\gamma, N > 0$ and sufficiently
220 small ϵ , suppose $g^{(2)}$ is γ -Lipschitz, and suppose the set of second-layer weights $\bar{\theta}_N^{(2)}$ satisfies that
221 $\mathcal{L}(\mathcal{T}[\bar{\theta}_N^{(2)}], \mathcal{D}_{q,k,D,N}) \leq q^{-N} \epsilon$. Then, there exists a constant $C_{\gamma,\epsilon,D}$, such that for any $0 \leq d' \leq$
222 $D, 1 \leq d \leq D, i, j \in [k]$, it holds that*

$$S_{d,d',i,j}[\bar{\theta}_N^{(2)}] \leq \frac{C_{\gamma,\epsilon,D}}{N} P_{d,j}[\bar{\theta}_N^{(2)}]. \quad (13)$$

223 Equation (13) requires $S_{d,d',i,j}[\theta^{(2)}]$ to be small relative to $P_{d,j}[\bar{\theta}_N^{(2)}]$, and can be interpreted as a
224 relaxation of Equation (9) which is equivalent to $S_{d,d',i,j}[\theta^{(2)}] = 0$. The proof of Theorem 2 shares
225 similar intuition as Theorem 1 and is given in Appendix C.2. As a direct corollary of Theorem 2, we
226 can additionally consider adding a weight decay, in which case approximate balance condition holds
227 as the regularization strength goes to 0:

228 **Corollary 2.** *Consider the setting where a Transformer with a fixed minimal first layer is trained to
229 minimize $\mathcal{L}_\lambda^{reg} = \mathcal{L}_\theta(x) + \lambda \frac{\|\theta\|_2^2}{2}$, which is the squared loss with λ weight decay. Suppose the function
230 $g^{(2)}$ of the Transformer is a fully connected network. Then, for any length N , there exists constant*

231 $C > 0$, such that for parameters $\theta_{\lambda, N}$ minimizing $\mathcal{L}_{\lambda}^{reg}$, it holds $\forall 0 \leq d' \leq D, 1 \leq d \leq D, i, j \in [k]$
 232 that,

$$\limsup_{\lambda \rightarrow 0} \frac{S_{d, d', i, j}[\theta_{\lambda, N}]}{P_{d, i}[\theta_{\lambda, N}] + 1} \leq \frac{C}{N}.$$

233 3.2 Interpretability Requires Inspecting More Than Any Single Weight Matrix

234 Another line of interpretability works involves inspecting the weight matrices of the model (Li et al.,
 235 2016; Dar et al., 2022; Eldan & Li, 2023). Some of the investigations are done locally, neglecting the
 236 interplay between different parts of the model. Our next result shows that from a representational
 237 perspective, isolating single weights may also be misleading for model interpretability:

238 **Theorem 3** (Indistinguishability From a Single Component). *Consider a L -layer Transformer \mathcal{T}*
 239 *with embedding dimension m , width w and $g^{(k)}(x) = \text{LN} \left(W_2^{(k)} \text{ReLU} \left(W_1^{(k)} x \right) \right) + x$. Suppose*
 240 *$\|W\|_2 = O(1)$ for every weight matrix W in \mathcal{T} . For $\delta \in (0, 1)$, consider a larger random Trans-*
 241 *former \mathcal{T}_{large} with $4L$ layers, embedding dimension $4m$, and width $O(\max\{m \log \frac{wmLN}{\epsilon\delta}, w\})$, whose*
 242 *weights are randomly sampled as $W_{i,j} \sim U(-1, 1)$ for every $W \in \mathcal{T}_{large}$.*

243 *Then, with probability $1 - \delta$ over the randomness of \mathcal{T}_{large} , we can obtain a nonstructural pruning*
 244 *(Definition 2) of \mathcal{T}_{large} , denoted as \mathcal{T}'_{large} , which ϵ -approximate \mathcal{T} . That is, $\forall \mathbf{X} \in \mathbb{R}^{d \times N}$ with*
 245 *$\|\mathbf{X}_{:,i}\|_2 \leq 1, \forall i \in [N]$,*

$$\|\mathcal{T}'_{large}(\mathbf{X}) - \mathcal{T}(\mathbf{X})\|_2 \leq \epsilon.$$

246 *Moreover, pick any weight matrix W in \mathcal{T}_{large} , with probability $1 - \delta$, for any smaller Transformers*
 247 *$\mathcal{T}_1, \mathcal{T}_2$ satisfying same conditions as \mathcal{T} , we have two pruned Transformers $\mathcal{T}_{Large,1}, \mathcal{T}_{Large,2}$ based on*
 248 *\mathcal{T}_{large} , such that they coincide on the pruned weight of W , and $\mathcal{T}_{Large,i}$ ϵ -approximate $\mathcal{T}_i, \forall i \in \{1, 2\}$.*

249 Theorem 3 implies that by inspecting any single weight matrix only, one cannot distinguish whether
 250 the pruned Transformer is approximating \mathcal{T}_1 or \mathcal{T}_2 . Hence, one should be cautious when using
 251 methods based solely on individual components to interpret the overall Transformer solution.

252 **Proof sketch: connection to Lottery Tickets.** Theorem 3 can also be interpreted as a provable
 253 lottery ticket hypothesis (Frankle & Carbin, 2018; Malach et al., 2020) for Transformers with random
 254 initialization, which can be of independent interest. In fact, the proof of Theorem 3 repetitively
 255 use Theorem 1 of Pensia et al. (2020). The key step of the proof is noticing pruning attention
 256 weight matrix of the larger Transformer \mathcal{T}_{large} to approximate attention weight matrix of the smaller
 257 transformer \mathcal{T} can be viewed as pruning a wide linear network to approximate a fixed matrix. The
 258 formal proofs are deferred to Appendix C.3.

259 4 Experiments

260 Our theory in Section 3 proves the existence of abundant *non-stack-like* attention patterns, all of
 261 which suffice for (near-)optimal generalization on Dyck. However, could there be *implicit biases* in
 262 the architecture and the optimization algorithm, which would potentially make the learned attention
 263 patterns more frequently stack-like? In this section, we show there is no evidence for such implicit
 264 bias in standard training (Section 4.1). However, a modified objective based on our theory can be
 265 used to *explicitly regularize* the model towards better length generalization (Section 4.2).

266 4.1 Different Attention Patterns Can Be Learned To Generate Dyck

267 We empirically verify our theoretical findings that Dyck solutions can give rise to a variety of attention
 268 patterns. We use the Adam optimizer (Kingma & Ba, 2014) unless specified otherwise. We use
 269 Transformers with 2 layers, 1 head, hidden dimension 50 and word embedding dimension 50. We test
 270 the accuracy of the model by randomly generating a Dyck prefix (Equation 1) that ends with a closing
 271 bracket, and evaluating whether the model predicts correctly the type of the last closing bracket given
 272 the rest of the prefix. Note that in this setting a correct parser should always be able to uniquely
 273 determine the correct closing bracket type (for the sequence to remain a valid Dyck sequence). We
 274 train on valid Dyck_{2,4} sequence with length less than 28 generated with $q = 0.5$, where all models
 275 are able to achieve $\geq 97\%$ test accuracy.

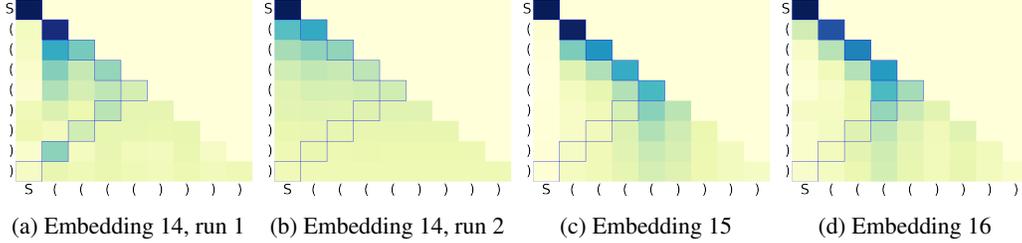


Figure 2: **Second-layer attention patterns of two-layer Transformers with a minimal first layer:** (a), (b) are based on embedding 14 with different learning rates, where the attention patterns show much variance as Theorem 1 predicts. (c), (d) are based on embedding 15 and 16. Different embedding functions lead to diverse attention patterns, most of which are not stack-like.

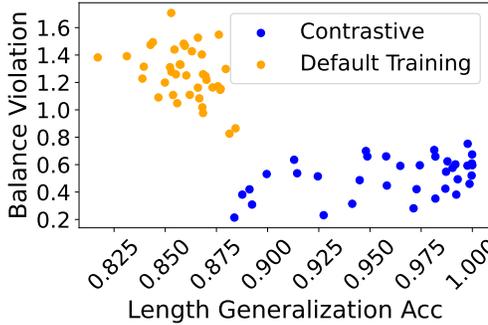


Figure 3: **Relationship Between Balance Violation and Length Generalization.** Accuracy from Transformers with minimal first layer with embedding 14, using both standard training and contrastive regularization (Equation (17)). Standard training leads to high balance violations which negatively correlate with length generalization performance. Contrastive regularization helps reduce the balance violation and improve the length generalization performance.

276 **Qualitative Results** As a response to (Q1), we observe that attention patterns of Transformers
 277 trained on Dyck are not always stack-like (Figure 1). In fact, the attention patterns vary even across
 278 different random initializations. Moreover, while Theorem 1 predicts that position encoding is not
 279 necessary for a Transformer to generate Dyck (this is verified by experiments, as Transformers with
 280 no positional encoding achieve $\geq 97\%$ accuracy), we observe that adding the position encoding²
 281 does affect the attention patterns. We also try fixing the attention layer as uniform attention and verify
 282 that uniform attention transformer can also fit the distribution almost perfectly, which is consistent
 283 with our theory.

284 We then experiment with two-layer Transformers with a minimal first layer. We experiment with
 285 three different types of embeddings e : let \mathbf{o}_t denote the one-hot embedding where $\mathbf{o}_t[t] = 1$,

$$e(\tau_{t,d}) = \mathbf{o}_{t \times D + d}, \quad (14)$$

$$e(\tau_{t,d}) = \mathbf{o}_t \oplus [\cos(\theta_d), \sin(\theta_d)], \theta_d = \arctan(d/(D+2-d)), \quad (15)$$

$$e(\tau_{t,d}) = \mathbf{o}_t \oplus \mathbf{o}_d. \quad (16)$$

286 Operator \oplus means the concatenation of two vectors. Equation (14) is the standard one-hot embed-
 287 ding for $\tau_{t,d}$. and Equation (16) is the concatenation of one-hot embedding of types and depths.
 288 Finally, Equation (15) is the embedding constructed in Yao et al. (2021).

289 As one can observe from Figure 2, the attention patterns learned by Transformers exhibit large
 290 variance between different choices of architectures and learning rates. We observe that most of the
 291 attention patterns learned by the Transformer are not stack-like.

292 **Quantitative Experiments.** We now quantify the variation in attention by comparing across multiple
 293 random initializations. We define the *attention variation* between two attention patterns $A_1, A_2 \in$
 294 $\mathbb{R}^{N \times N}$ over an length- N input sequence as $\text{Variation}(A_1, A_2) = \|A_1 - A_2\|_F^2$. We will then
 295 calculate the average variation of an architecture by running $n = 40$ random initializations and
 296 calculate the average variation between the attention patterns of the n random initializations on
 297 sequence $[[[[]]]](((())$). We will call this quantity the *average attention variation*.

²We use the linear positional encoding following Yao et al. (2021), where for the i_{th} position, define encoding $e_p(i) := i/T_{\max}$ for some T_{\max} .

298 We observe that for standard two layer training with linear position embedding, the average attention
 299 variation is 2.20. For training without position embedding, the average attention variation is 2.27.
 300 Both variation is closed to the random baseline value of 2.85^3 , showing that the attention head learned
 301 by different initializations indeed tend to be very different. We also experiment with Transformer
 302 with a minimal first layer and the embedding in Equation (14), which reduces the average variation
 303 to 0.24. We hypothesize that the structural constraints in this setting provide sufficiently strong
 304 inductive bias that limit the variability of attention patterns.

305 4.2 Guiding The Transformer To Learn Balanced Attention

306 In our experiments, we observe that although models learned via standard training that can generalize
 307 well in distribution, the length generalization performance is far from optimal. This implies that the
 308 models are not finding the correct algorithm for parsing Dyck when learning from finite samples. A
 309 natural question is: can we guide Transformers towards correct algorithms, as measured by better
 310 generalization on longer Dyck sequences?

311 In the following, we measure length generalization performance by testing the accuracy of the
 312 model on valid Dyck prefixes with length randomly sampled from 400 to 500, which approximately
 313 correspond to 16 times the length of the training sequences. We will show generalization can be
 314 improved by regularizing the attentions to be more balanced, inspired by results in Section 3.

315 **Balance violation negatively correlates with length generalization accuracy** We denote the
 316 *balance violation* of a Transformer as $\beta := \mathbb{E}_{d,d',i,j} [S_{d,d',i,j}/P_{d,j}]$ for S, P defined in Equations (10)
 317 and (12). Theorem 1 predicts that for models with a minimal first layer, perfect length generalization
 318 requires β to be zero. Beyond such idealized condition, it is natural to ask whether a small yet positive
 319 β correlates with length generalization accuracy in practice. Our results show a moderate correlation
 320 (-0.38 SpearmanR with p-value 0.014) based on over 40 random initializations (Figure 3).

321 Given the correlation, we design a contrastive training objective to reduce the balance violation,
 322 which ideally would lead to improved length generalization. Specifically, let p_r denote a prefix of r
 323 nested pairs of brackets of for $r \sim U([D])$, and let $\mathcal{T}(s | p_r \oplus s)$ denote the logits for s when \mathcal{T} takes
 324 as input the concatenation of p_r and s . We define the *contrastive regularization* $R_{\text{contrastive}}(s)$ as the
 325 mean squared error between the logits of $\mathcal{T}(s)$ and $\mathcal{T}(s | p_r \oplus s)$, taking expectation over r and p_r :

$$\mathbb{E}_{r \sim U([D]), p_r} [\|\mathcal{T}(s | p_r \oplus s) - \mathcal{T}(s)\|_F^2]. \quad (17)$$

326 Following the same intuition as in the proof of Theorem 1, if the model can perfectly length-generalize,
 327 then the contrastive loss will be zero. We then train the model with contrastive loss and observe that
 328 the balance violation is reduced and the length generalization performance is improved (Figure 3).

329 5 Conclusion

330 Why interpreting individual components sometimes leads to misconceptions? Through a case study
 331 of the Dyck grammar, we provide theoretical and empirical evidence that even in this simple and
 332 well-understood setup, Transformers can implement a rich set of non-interpretable solutions, and
 333 typically do not encode task-specific structures in local components. Our results provide a theoretical
 334 perspective as to why careful analyses are required for interpreting Transformers.

335 **Limitations and future work.** Our results do not preclude that interpretable attention patterns can
 336 emerge in multi-head, overparameterized Transformers trained on more complex data distributions.
 337 In that case, we discuss some limitations of such interpretation in Appendix B.

338 Interesting directions of future work include extending our theoretical results to more complex
 339 settings (in terms of both architecture choice and data distribution), theoretical characterization of the
 340 learning dynamics, and more experiments in controlled settings for testing the connections between
 341 the training approach, interpretability, and task performance. We motivate these questions and discuss
 342 some relevant trade-offs in Appendix B.

³The random baseline is calculated by generating purely random attention patterns (from the simplex, i.e. random square matrices s.t. each row sums up to 1) and calculate the average attention variation between them.

343 **References**

- 344 Yonatan Belinkov. Probing classifiers: Promises, shortcomings, and advances. *Computational Linguistics*, 48(1):207–219, March 2022. doi: 10.1162/coli_a_00422. URL <https://aclanthology.org/2022.c1-1.7>.
- 347 Satwik Bhattamishra, Kabir Ahuja, and Navin Goyal. 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)*, pp. 7096–7116, Online, November 2020a. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.576. URL <https://aclanthology.org/2020.emnlp-main.576>.
- 352 Satwik Bhattamishra, Arkil Patel, and Navin Goyal. On the computational power of transformers and its implications in sequence modeling. In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pp. 455–475, Online, November 2020b. Association for Computational Linguistics. doi: 10.18653/v1/2020.conll-1.37. URL <https://aclanthology.org/2020.conll-1.37>.
- 357 Tolga Bolukbasi, Adam Pearce, Ann Yuan, Andy Coenen, Emily Reif, Fernanda Viégas, and Martin Wattenberg. An interpretability illusion for bert. *arXiv preprint arXiv: Arxiv-2104.07143*, 2021.
- 359 Gino Brunner, Yang Liu, Damian Pascual, Oliver Richter, Massimiliano Ciaramita, and Roger Wattenhofer. On identifiability in transformers. In *International Conference on Learning Representations, 2020*. URL <https://openreview.net/forum?id=BJg1f6EFDB>.
- 362 Nick Cammarata, Shan Carter, Gabriel Goh, Chris Olah, Michael Petrov, Ludwig Schubert, Chelsea Voss, Ben Egan, and Swee Kiat Lim. Thread: Circuits. *Distill*, 2020. doi: 10.23915/distill.00024. <https://distill.pub/2020/circuits>.
- 365 Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. What does BERT look at? an analysis of BERT’s attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 276–286, Florence, Italy, August 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-4828. URL <https://aclanthology.org/W19-4828>.
- 370 Guy Dar, Mor Geva, Ankit Gupta, and Jonathan Berant. Analyzing transformers in embedding space, 2022.
- 372 Yichuan Deng, Zhihang Li, and Zhao Song. Attention scheme inspired softmax regression, 2023.
- 373 Javid Ebrahimi, Daniel Lowd, and Dejing Dou. On adversarial examples for character-level neural machine translation. In *International Conference on Computational Linguistics (COLING)*, 2018.
- 375 Javid Ebrahimi, Dhruv Gelda, and Wei Zhang. How can self-attention networks recognize Dyck-n languages? In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 4301–4306, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.384. URL <https://aclanthology.org/2020.findings-emnlp.384>.
- 379 Benjamin L Edelman, Surbhi Goel, Sham Kakade, and Cyril Zhang. Inductive biases and variable creation in self-attention mechanisms. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 5793–5831. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/edelman22a.html>.
- 384 Ronen Eldan and Yuanzhi Li. Tinstories: How small can language models be and still speak coherent english?, 2023.
- 386 Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021. <https://Transformer-circuits.pub/2021/framework/index.html>.
- 392 Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*, 2020.
- 394 Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *International Conference On Learning Representations*, 2018.
- 395

- 396 Yeqi Gao, Sridhar Mahadevan, and Zhao Song. An over-parameterized exponential regression, 2023.
- 397 F. Gers and J. Schmidhuber. Lstm recurrent networks learn simple context-free and context-sensitive
398 languages. *IEEE transactions on neural networks*, 12 6:1333–40, 2001.
- 399 Christopher Grimsley, Elijah Mayfield, and Julia R.S. Bursten. Why attention is not explanation:
400 Surgical intervention and causal reasoning about neural models. In *Proceedings of the
401 Twelfth Language Resources and Evaluation Conference*, pp. 1780–1790, Marseille, France,
402 May 2020. European Language Resources Association. ISBN 979-10-95546-34-4. URL
403 <https://aclanthology.org/2020.lrec-1.220>.
- 404 Michael Hahn. Theoretical limitations of self-attention in neural sequence models. *Trans. Assoc.
405 Comput. Linguistics*, 8:156–171, 2020. doi: 10.1162/tacl_a_00306. URL [https://doi.org/
406 10.1162/tacl_a_00306](https://doi.org/10.1162/tacl_a_00306).
- 407 John Hewitt and Percy Liang. Designing and interpreting probes with control tasks. In *Proceedings
408 of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th
409 International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 2733–
410 2743, Hong Kong, China, November 2019. Association for Computational Linguistics. doi:
411 10.18653/v1/D19-1275. URL <https://aclanthology.org/D19-1275>.
- 412 John Hewitt and Christopher D Manning. A structural probe for finding syntax in word representations.
413 In *Proceedings of the 2019 Conference of the North American Chapter of the Association for
414 Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*,
415 pp. 4129–4138, 2019.
- 416 John Hewitt, Michael Hahn, Surya Ganguli, Percy Liang, and Christopher D Manning. Rnns can
417 generate bounded hierarchical languages with optimal memory. *arXiv preprint arXiv:2010.07515*,
418 2020.
- 419 Phu Mon Htut, Jason Phang, Shikha Bordia, and Samuel R. Bowman. Do attention heads in bert
420 track syntactic dependencies?, 2019.
- 421 Sarthak Jain and Byron C. Wallace. Attention is not Explanation. In *Proceedings of the 2019
422 Conference of the North American Chapter of the Association for Computational Linguistics:
423 Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 3543–3556, Minneapolis,
424 Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1357.
425 URL <https://aclanthology.org/N19-1357>.
- 426 Samy Jelassi, Michael Eli Sander, and Yuanzhi Li. Vision transformers provably learn spatial structure.
427 In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in
428 Neural Information Processing Systems*, 2022. URL [https://openreview.net/forum?id=
429 eMW9AkXaREI](https://openreview.net/forum?id=eMW9AkXaREI).
- 430 Lifeng Jin, Finale Doshi-Velez, Timothy Miller, William Schuler, and Lane Schwartz. Unsupervised
431 grammar induction with depth-bounded pcfg. *Transactions of the Association for Computational
432 Linguistics*, 6:211–224, 2018.
- 433 Fred Karlsson. Constraints on multiple center-embedding of clauses. *Journal of Linguistics*, 43(2):
434 365–392, 2007.
- 435 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint
436 arXiv:1412.6980*, 2014.
- 437 Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. Attention is not only a weight:
438 Analyzing transformers with vector norms. In *Proceedings of the 2020 Conference on Empirical
439 Methods in Natural Language Processing (EMNLP)*, pp. 7057–7075, Online, November 2020.
440 Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.574. URL [https://
441 /aclanthology.org/2020.emnlp-main.574](https://aclanthology.org/2020.emnlp-main.574).
- 442 Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. Revealing the dark secrets
443 of BERT. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language
444 Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-
445 IJCNLP)*, pp. 4365–4374, Hong Kong, China, November 2019. Association for Computational
446 Linguistics. doi: 10.18653/v1/D19-1445. URL <https://aclanthology.org/D19-1445>.
- 447 Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. Visualizing and understanding neural
448 models in NLP. In *Proceedings of the 2016 Conference of the North American Chapter of the
449 Association for Computational Linguistics: Human Language Technologies*, pp. 681–691, San

- 450 Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/
451 N16-1082. URL <https://aclanthology.org/N16-1082>.
- 452 Xian Li and Hongyu Gong. Robust optimization for multilingual translation with imbalanced
453 data. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan
454 (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 25086–25099. Cur-
455 ran Associates, Inc., 2021. URL [https://proceedings.neurips.cc/paper/2021/file/
456 d324a0cc02881779dcda44a675fdcaaa-Paper.pdf](https://proceedings.neurips.cc/paper/2021/file/d324a0cc02881779dcda44a675fdcaaa-Paper.pdf).
- 457 Yuchen Li and Andrej Risteski. The limitations of limited context for constituency parsing. In
458 *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and
459 the 11th International Joint Conference on Natural Language Processing (Volume 1: Long
460 Papers)*, pp. 2675–2687, Online, August 2021. Association for Computational Linguistics. doi:
461 10.18653/v1/2021.acl-long.208. URL <https://aclanthology.org/2021.acl-long.208>.
- 462 Yuchen Li, Yuanzhi Li, and Andrej Risteski. How do transformers learn topic structure: Towards a
463 mechanistic understanding, 2023.
- 464 Yongjie Lin, Yi Chern Tan, and Robert Frank. Open sesame: Getting inside BERT’s linguistic knowl-
465 edge. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural
466 Networks for NLP*, pp. 241–253, Florence, Italy, August 2019. Association for Computational
467 Linguistics. doi: 10.18653/v1/W19-4825. URL <https://aclanthology.org/W19-4825>.
- 468 Bingbin Liu, Daniel Hsu, Pradeep Kumar Ravikumar, and Andrej Risteski. Masked prediction:
469 A parameter identifiability view. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and
470 Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022a. URL
471 <https://openreview.net/forum?id=Hbv1b4D1aFC>.
- 472 Bingbin Liu, Jordan T. Ash, Surbhi Goel, Akshay Krishnamurthy, and Cyril Zhang. Transformers
473 learn shortcuts to automata. In *The Eleventh International Conference on Learning Representations*,
474 2023. URL <https://openreview.net/forum?id=De4FYqjFueZ>.
- 475 Hong Liu, Sang Michael Xie, Zhiyuan Li, and Tengyu Ma. Same pre-training loss, better downstream:
476 Implicit bias matters for language models. *arXiv preprint arXiv:2210.14199*, 2022b.
- 477 Liyuan Liu, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Jiawei Han. Understanding the
478 difficulty of training transformers. In *Proceedings of the 2020 Conference on Empirical
479 Methods in Natural Language Processing (EMNLP)*, pp. 5747–5763, Online, November 2020.
480 Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.463. URL
481 <https://aclanthology.org/2020.emnlp-main.463>.
- 482 Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. Linguistic
483 knowledge and transferability of contextual representations. In *Proceedings of the 2019 Conference
484 of the North American Chapter of the Association for Computational Linguistics: Human Language
485 Technologies, Volume 1 (Long and Short Papers)*, pp. 1073–1094, Minneapolis, Minnesota, June
486 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1112. URL <https://aclanthology.org/N19-1112>.
- 487
- 488 Eran Malach, Gilad Yehudai, Shai Shalev-Schwartz, and Ohad Shamir. Proving the lottery ticket
489 hypothesis: Pruning is all you need. In *International Conference on Machine Learning*, pp.
490 6682–6691. PMLR, 2020.
- 491 Clara Meister, Stefan Lazov, Isabelle Augenstein, and Ryan Cotterell. Is sparse attention more
492 interpretable? In *Proceedings of the 59th Annual Meeting of the Association for Computational
493 Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2:
494 Short Papers)*, pp. 122–129, Online, August 2021. Association for Computational Linguistics. doi:
495 10.18653/v1/2021.acl-short.17. URL <https://aclanthology.org/2021.acl-short.17>.
- 496 William Merrill. Sequential neural networks as automata. In *Proceedings of the Workshop on Deep
497 Learning and Formal Languages: Building Bridges*, pp. 1–13, Florence, August 2019. Association
498 for Computational Linguistics. doi: 10.18653/v1/W19-3901. URL [https://www.aclweb.org/
499 anthology/W19-3901](https://www.aclweb.org/anthology/W19-3901).
- 500 Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one?
501 In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett
502 (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates,
503 Inc., 2019. URL [https://proceedings.neurips.cc/paper_files/paper/2019/file/
504 2c601ad9d2ff9bc8b282670cdd54f69f-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/2c601ad9d2ff9bc8b282670cdd54f69f-Paper.pdf).

505 Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. Progress measures for
506 grokking via mechanistic interpretability. In *The Eleventh International Conference on Learning*
507 *Representations*, 2023. URL <https://openreview.net/forum?id=9XFSbDPmdW>.

508 Toan Q. Nguyen and Julian Salazar. Transformers without tears: Improving the normalization
509 of self-attention. In *Proceedings of the 16th International Conference on Spoken Language*
510 *Translation*, Hong Kong, November 2-3 2019. Association for Computational Linguistics. URL
511 <https://aclanthology.org/2019.iwslt-1.17>.

512 Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan,
513 Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli,
514 Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane
515 Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish,
516 and Chris Olah. In-context learning and induction heads. *Transformer Circuits Thread*, 2022.
517 <https://Transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html>.

518 Ankit Pensia, Shashank Rajput, Alliot Nagle, Harit Vishwakarma, and Dimitris Papailiopoulos.
519 Optimal lottery tickets via subset sum: Logarithmic over-parameterization is sufficient. *Advances*
520 *in neural information processing systems*, 33:2599–2610, 2020.

521 Jorge Perez, Pablo Barcelo, and Javier Marinkovic. Attention is turing-complete. *Journal of Machine*
522 *Learning Research*, 22(75):1–35, 2021. URL <http://jmlr.org/papers/v22/20-302.html>.

523 Sai Prasanna, Anna Rogers, and Anna Rumshisky. When BERT Plays the Lottery, All Tickets Are
524 Winning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language*
525 *Processing (EMNLP)*, pp. 3208–3229, Online, November 2020. Association for Computational
526 Linguistics. doi: 10.18653/v1/2020.emnlp-main.259. URL <https://aclanthology.org/2020.emnlp-main.259>.

528 Alessandro Raganato and Jörg Tiedemann. An analysis of encoder representations in transformer-
529 based machine translation. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: An-*
530 *alyzing and Interpreting Neural Networks for NLP*, pp. 287–297, Brussels, Belgium, Novem-
531 ber 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-5431. URL
532 <https://aclanthology.org/W18-5431>.

533 Anna Rogers, Olga Kovaleva, and Anna Rumshisky. A primer in bertology: What we know about
534 how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866, 2020.

535 M.P. Schützenberger. On context-free languages and push-down automata. *Information and Control*,
536 6(3):246–264, 1963. ISSN 0019-9958. doi: [https://doi.org/10.1016/S0019-9958\(63\)90306-1](https://doi.org/10.1016/S0019-9958(63)90306-1). URL
537 <https://www.sciencedirect.com/science/article/pii/S0019995863903061>.

538 Sofia Serrano and Noah A. Smith. Is attention interpretable? In *Proceedings of the 57th Annual*
539 *Meeting of the Association for Computational Linguistics*, pp. 2931–2951, Florence, Italy, July
540 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1282. URL <https://aclanthology.org/P19-1282>.

542 Hava T. Siegelmann and Eduardo D. Sontag. On the computational power of neural nets. In
543 *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, COLT '92, pp.
544 440–449, New York, NY, USA, 1992. Association for Computing Machinery. ISBN 089791497X.
545 doi: 10.1145/130385.130432. URL <https://doi.org/10.1145/130385.130432>.

546 Kaiser Sun and Ana Marasović. Effective attention sheds light on interpretability. In *Findings of the*
547 *Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 4126–4135, Online, August
548 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.361. URL
549 <https://aclanthology.org/2021.findings-acl.361>.

550 Mirac Suzgun, Yonatan Belinkov, Stuart Shieber, and Sebastian Gehrmann. LSTM networks can per-
551 form dynamic counting. In *Proceedings of the Workshop on Deep Learning and Formal Languages:*
552 *Building Bridges*, pp. 44–54, Florence, August 2019. Association for Computational Linguistics.
553 doi: 10.18653/v1/W19-3905. URL <https://www.aclweb.org/anthology/W19-3905>.

554 Ian Tenney, Dipanjan Das, and Ellie Pavlick. Bert rediscovered the classical nlp pipeline. *arXiv*
555 *preprint arXiv:1905.05950*, 2019.

556 Jesse Vig and Yonatan Belinkov. Analyzing the structure of attention in a transformer language
557 model. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting*
558 *Neural Networks for NLP*, pp. 63–76, Florence, Italy, August 2019. Association for Computational
559 Linguistics. doi: 10.18653/v1/W19-4808. URL <https://aclanthology.org/W19-4808>.

- 560 Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head
561 self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *Proceedings of the*
562 *57th Annual Meeting of the Association for Computational Linguistics*, pp. 5797–5808, Florence,
563 Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1580. URL
564 <https://aclanthology.org/P19-1580>.
- 565 Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. Inter-
566 pretability in the wild: a circuit for indirect object identification in GPT-2 small. In *International*
567 *Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=NpsVSN6o4u1>.
- 569 Colin Wei, Yining Chen, and Tengyu Ma. Statistically meaningful approximation: a case study on
570 approximating turing machines with transformers, 2021. URL [https://arxiv.org/abs/2107.](https://arxiv.org/abs/2107.13163)
571 [13163](https://arxiv.org/abs/2107.13163).
- 572 Gail Weiss, Yoav Goldberg, and Eran Yahav. On the practical computational power of finite precision
573 rnns for language recognition. *arXiv preprint arXiv:1805.04908*, 2018.
- 574 Gail Weiss, Yoav Goldberg, and Eran Yahav. Thinking like transformers. In Marina Meila and Tong
575 Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139
576 of *Proceedings of Machine Learning Research*, pp. 11080–11090. PMLR, 18–24 Jul 2021. URL
577 <https://proceedings.mlr.press/v139/weiss21a.html>.
- 578 Sarah Wiegrefe and Yuval Pinter. Attention is not not explanation. In *Proceedings of the 2019*
579 *Conference on Empirical Methods in Natural Language Processing and the 9th International Joint*
580 *Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 11–20, Hong Kong, China,
581 November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1002. URL
582 <https://aclanthology.org/D19-1002>.
- 583 Zhiyong Wu, Yun Chen, Ben Kao, and Qun Liu. Perturbed masking: Parameter-free probing for
584 analyzing and interpreting BERT. In *Proceedings of the 58th Annual Meeting of the Association*
585 *for Computational Linguistics*, pp. 4166–4176, Online, July 2020. Association for Computational
586 Linguistics. doi: 10.18653/v1/2020.acl-main.383. URL [https://aclanthology.org/2020.](https://aclanthology.org/2020.acl-main.383)
587 [acl-main.383](https://aclanthology.org/2020.acl-main.383).
- 588 Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang,
589 Yanyan Lan, Liwei Wang, and Tie-Yan Liu. On layer normalization in the transformer architecture.
590 In *Proceedings of the 37th International Conference on Machine Learning, ICML’20*. JMLR.org,
591 2020.
- 592 Shunyu Yao, Binghui Peng, Christos Papadimitriou, and Karthik Narasimhan. Self-attention networks
593 can process bounded hierarchical languages. In *Proceedings of the 59th Annual Meeting of*
594 *the Association for Computational Linguistics and the 11th International Joint Conference on*
595 *Natural Language Processing (Volume 1: Long Papers)*, pp. 3770–3785, Online, August 2021.
596 Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.292. URL [https:](https://aclanthology.org/2021.acl-long.292)
597 [//aclanthology.org/2021.acl-long.292](https://aclanthology.org/2021.acl-long.292).
- 598 Chulhee Yun, Srinadh Bhojanapalli, Ankit Singh Rawat, Sashank Reddi, and Sanjiv Kumar. Are trans-
599 formers universal approximators of sequence-to-sequence functions? In *International Conference*
600 *on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ByxRMONTvr>.
- 601 Marvin Zhang, Henrik Marklund, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk
602 minimization: A meta-learning approach for tackling group shift. *arXiv preprint arXiv:2007.02931*,
603 2020.
- 604 Yi Zhang, Arturs Backurs, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, and Tal Wagner.
605 Unveiling transformers with lego: a synthetic reasoning task, 2022. URL [https://arxiv.org/](https://arxiv.org/abs/2206.04301)
606 [abs/2206.04301](https://arxiv.org/abs/2206.04301).
- 607 Haoyu Zhao, Abhishek Panigrahi, Rong Ge, and Sanjeev Arora. Do transformers parse while
608 predicting the masked word?, 2023.

610 **Appendix****Table of Contents**

613	A Additional Related Work	16
614	B Are interpretable attention patterns useful?	18
615	B.1 Can interpretable attention patterns be misleading?	18
616	B.2 Can interpretable attention patterns be important?	20
617	C Omitted Proofs in Section 3	21
618	C.1 Proof of Theorem 1	21
619	C.2 Proof of Theorem 2	23
620	C.3 Proof of Theorem 3 – Indistinguishability from a single component	25
621	D Experiments	29
622	D.1 Training Details	29
623	D.2 Additional Results on Dyck Prefix	29

627 A Additional Related Work

628 **Interpreting Transformer solutions** Prior empirical works show that Transformers trained on
629 natural language data can produce representations that contain rich syntactic and semantic information,
630 by designing a wide range of “probing” tasks (Raganato & Tiedemann, 2018; Liu et al., 2019; Hewitt
631 & Manning, 2019; Clark et al., 2019; Tenney et al., 2019; Hewitt & Liang, 2019; Kovaleva et al.,
632 2019; Lin et al., 2019; Wu et al., 2020; Belinkov, 2022) (or other approaches using the attention
633 weights or parameters in neurons directly Vig & Belinkov, 2019; Htut et al., 2019; Sun & Marasović,
634 2021; Eldan & Li, 2023). However, there is no canonical way to probe the model, partially due to the
635 huge design space of probing tasks, and even a slight change in the setup may lead to very different
636 (sometimes even seemingly contradictory) interpretations of the result (Hewitt & Liang, 2019). In this
637 work, we tackle such ambiguity through a different perspective—by developing formal (theoretical)
638 understanding of solutions learned by Transformers. Our results imply that it may be challenging
639 to try to interpret Transformer solutions based on individual parameters (Li et al., 2016; Dar et al.,
640 2022), or based on constructive proofs (unless the Transformer is specially trained to be aligned with
641 a certain algorithm, as in Weiss et al., 2021).

642 **Interpreting attention patterns** Prior works (Jain & Wallace, 2019; Serrano & Smith, 2019;
643 Rogers et al., 2020; Grimsley et al., 2020; Brunner et al., 2020; Prasanna et al., 2020; Meister et al.,
644 2021; Bolukbasi et al., 2021, *inter alia*) present negative results on deriving explanations from
645 attention weights using approaches by Vig & Belinkov (2019); Kobayashi et al. (2020, *inter alia*).
646 However, Wiegrefe & Pinter (2019) argues to the contrary by pointing out flaws in the experimental
647 design and arguments of some of the prior works; they also call for theoretical analysis on the issue.
648 Hence, a takeaway from these prior works is that expositions on explainability based on attention
649 requires clearly defining the notion of explainability adopted (often task-specific). In our work,
650 we restrict our main theoretical analysis to the fully defined data distribution of Dyck language
651 (Definition 1), and define “interpretable attention pattern” as the stack-like pattern proposed in prior
652 theoretical (Yao et al., 2021) and empirical (Ebrahimi et al., 2020) works. These concrete settings
653 and definitions allow us to mathematically state our results and provide theoretical reasons.

654 **Theoretical understanding of representability** Methodologically, our work joins a long line of
655 prior works that characterize the solution of neural networks via the lens of simple synthetic data,
656 from class results on RNN representability (Siegelmann & Sontag, 1992; Gers & Schmidhuber,
657 2001; Weiss et al., 2018; Suzgun et al., 2019; Merrill, 2019; Hewitt et al., 2020), to the more recent
658 Transformer results on parity (Hahn, 2020), Dyck (Yao et al., 2021), topic model (Li et al., 2023),
659 and formal grammars in general (Bhattamishra et al., 2020a; Li & Risteski, 2021; Zhang et al., 2022;
660 Liu et al., 2023; Zhao et al., 2023). Our work complements prior works by showing that although
661 representational results can be obtained via intuitive “constructive proofs” that assign values to the
662 weight matrices, the model does not typically converge to those intuitive solutions in practice. Similar
663 messages are conveyed in Liu et al. (2023), which presents different types of constructions using
664 different numbers of layers. In contrast, we show that there exist multiple different constructions
665 even when the number of layers is kept the same.

666 There are also theoretical results on Transformers in terms of Turing completeness (Bhattamishra
667 et al., 2020b; Perez et al., 2021), universal approximability (Yun et al., 2020), and statistical sample
668 complexity (Wei et al., 2021; Edelman et al., 2022), which are orthogonal to our work.

669 **Transformer optimization** Given multiple global optima, understanding Transformer solutions
670 requires analyzing the training dynamics. Recent works theoretically analyze the learning process
671 of Transformers on simple data distributions, e.g. when the attention weights only depend on the
672 position information (Jelassi et al., 2022), or only depend on the content (Li et al., 2023). Our work
673 studies a syntax-motivated setting in which both content and position are critical. We also highlight
674 that Transformer solutions are very sensitive to detailed changes, such as positional encoding, layer
675 norm, sharpness regularization (Foret et al., 2020), or pre-training task (Liu et al., 2022a). On a
676 related topic but towards different goals, a series of prior works aim to improve the training process
677 of Transformers with algorithmic insights (Nguyen & Salazar, 2019; Xiong et al., 2020; Liu et al.,
678 2020; Zhang et al., 2020; Li & Gong, 2021, *inter alia*). An end-to-end theoretical characterization of
679 the training dynamics remains an open problem; recent works that propose useful techniques towards
680 this goal include Gao et al., 2023; Deng et al., 2023.

681 **Mechanistic interpretability** Finally, it is worth noting that the challenges highlighted in our work
682 do not contradict the line of prior works that aim to improve *mechanistic interpretability* into a trained
683 model or the training process (Cammarata et al., 2020; Elhage et al., 2021; Olsson et al., 2022; Nanda
684 et al., 2023; Li et al., 2023); although we prove that components (e.g. attention scores) of trained
685 Transformers do not generally admit intuitive interpretations based on the data distribution, it is still
686 possible to develop circuit-level understanding about a particular model, or measures that closely
687 track the training process, following these prior works.

688 B Are interpretable attention patterns useful?

689 Our results Section 3 and Section 4.1 demonstrate that Transformers are sufficiently expressive that a
690 (near-)optimal loss on Dyck languages can be achieved by a variety of attention patterns, many of
691 which may not be interpretable.

692 However, multiple prior works have shown that for multi-layer multi-head Transformers trained on
693 natural language datasets, it is often possible to locate attention heads that produce interpretable
694 attention patterns (Vig & Belinkov, 2019; Htut et al., 2019; Sun & Marasović, 2021). Hence, it is
695 also illustrative to consider the “converse question” of (Q1): when some attention heads do learn to
696 produce attention patterns that suggest intuitive interpretations, what benefits can they bring?

697 We discuss this through two perspectives:

- 698 • **Reliability of interpretation:** Is the Transformer necessarily implementing a solution consistent
699 with such interpretation based on the attention patterns? (Section B.1)
- 700 • **Usefulness for task performance:** Are those interpretable attention heads more important for the
701 task than other uninterpretable attention heads? (Section B.2)

702 We present preliminary analysis on these questions, and motivate future works on the interpretability
703 of attention patterns using rigorous theoretical analysis and carefully designed experiments.

704 B.1 Can interpretable attention patterns be misleading?

705 We show through a simple argument that interpretations based on attention patterns can sometimes
706 be misleading, as we formalize in the following proposition:

707 **Proposition 1.** Consider an L -layer Transformer \mathcal{T} (Equation (6)). For any $W_K^{(l)}, W_Q^{(l)} \in$
708 $\mathbb{R}^{m_a \times m}$ ($l \in [L]$), there exist $W_{\text{Head}} \in \mathbb{R}^{2k \times w}$ and $b_{\text{Head}} \in \mathbb{R}^{2k}$ such that $\mathcal{T}(\mathcal{Z}) = 0, \forall \mathcal{Z}$.

709 While its proof is trivial (simply setting $W_{\text{Head}} = 0$ and $b_{\text{Head}} = 0$ suffices), Proposition 1 implies
710 that the solution represented by the Transformer could possibly be independent of the attention
711 patterns in all the layers (1 through l). Hence, it could be misleading to interpret Transformer
712 solutions solely based on these attention patterns.

713 Empirically, Transformers trained on Dyck indeed sometimes produce misleading attention patterns.

714 We present one representative example in Figure 4, and Figure 5, in which *all interpretable attention*
715 *patterns are misleading*.

716 We also present additional results in Figure 6, in which *some interpretable attention patterns are*
717 *misleading, and some are not*.

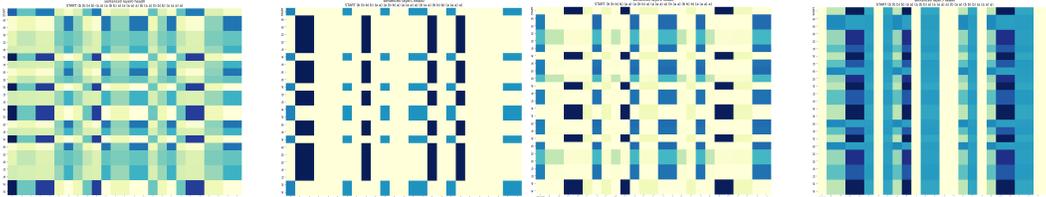


Figure 4: **Even interpretable attention patterns can be misleading:** For a 4-layer Transformer trained on Dyck with the *copying* task (with $> 96\%$ validation accuracy), i.e. the output should be exactly the same as the input, the attention patterns in some layers seem interpretable: (layer 2) attending to bracket type a) or (b); (layer 3) attending to closing brackets; (layer 4) never attending to bracket type a); However, none of them are informative of the copying task. This is possible because Transformers can use the residual connections (or weights MLPs or the value matrices) to solve copying, bypassing the need of using attention.

718 Similar message has been conveyed in prior works Bolukbasi et al. (2021), and future works may aim
719 to achieve the *faithfulness*, *completeness*, and *minimality* conditions in Wang et al. (2023).

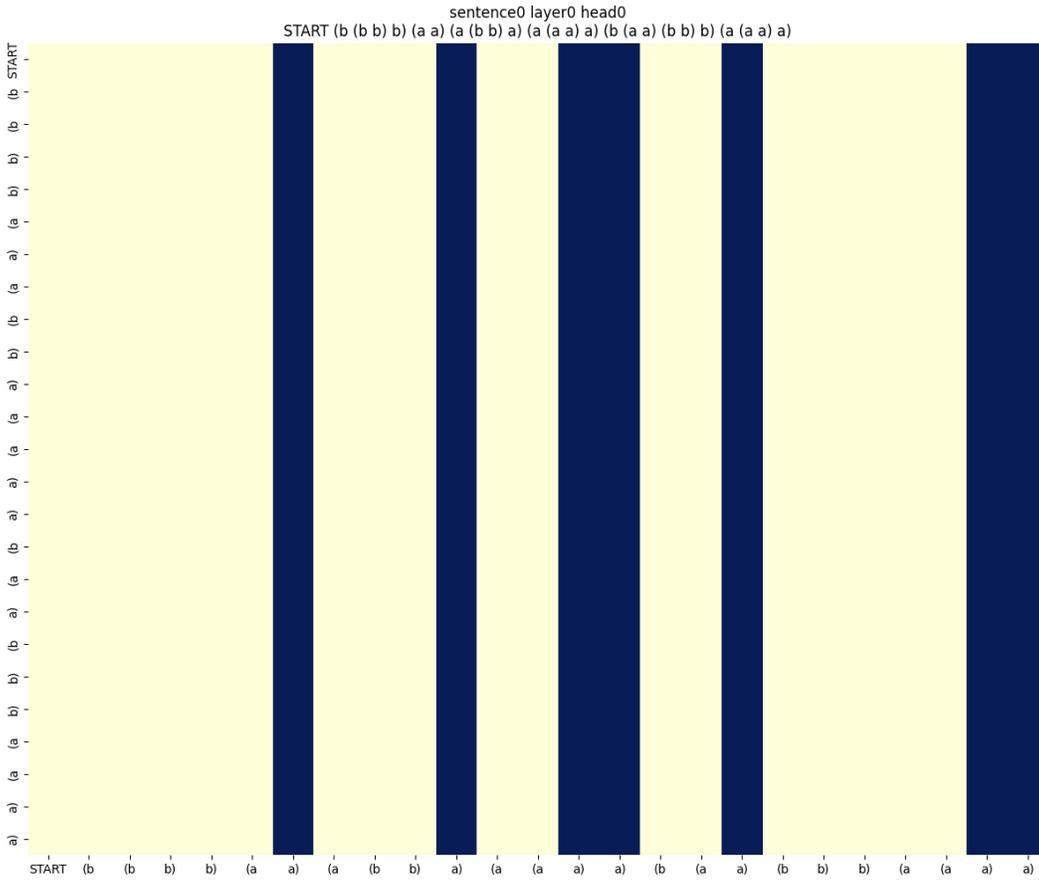


Figure 5: **Even interpretable attention patterns can be misleading:** For a 1-layer Transformer trained on Dyck with the *copying* task (with > 90% validation accuracy), i.e. the output should be exactly the same as the input, the attention pattern seems to be attending to closing brackets only, but that is not informative of the copying task.

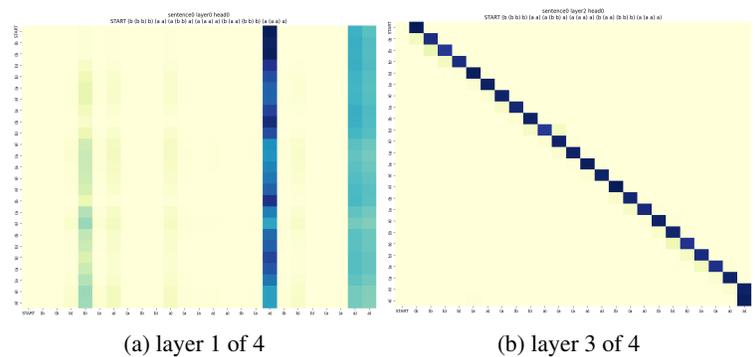


Figure 6: **Even interpretable attention patterns can be misleading:** For a 4-layer Transformer trained on Dyck with the *copying* task (with > 96% validation accuracy), i.e. the output should be exactly the same as the input, both types of attention patterns are common: (a) attending to closing brackets, which is uninformative of the copying task; (b) attending to the current position, which solves the copying task.

720 **B.2 Can interpretable attention patterns be important?**

721 Kovaleva et al. (2019) observes that, when the “importance” of an attention head is defined as the
722 performance drop the model suffers when the head is disabled, then for most tasks they test, the most
723 important attention head in each layer *does not* tend to be interpretable.

724 However, experiments by Voita et al. (2019) led to a seemingly contradictory observation: when at-
725 tention heads are systematically pruned by finetuning the Transformer with a relaxation of L_0 -penalty
726 (i.e. encouraging the number of remaining attention heads to be small), most remaining attention
727 heads that survive the pruning can be associated with certain functionalities such as positional,
728 syntactic, or attending to rare tokens.

729 These works seem to bring mixed conclusions to our question: are interpretable attention heads
730 more important for the task than other uninterpretable attention heads? We interpret these results
731 by conjecturing that the definition of “importance” (reflected in their experimental design) plays a
732 crucial role:

- 733 • When the importance of an attention head is defined *treating all other attention heads as fixed*,
734 motivating experiments that prune/disable certain heads while keeping other heads unchanged
735 (Michel et al., 2019; Kovaleva et al., 2019), the conclusion may be mostly pessimistic: mostly no
736 strong connection between interpretability and importance.
- 737 • On the other hand, when the importance of an attention head is defined *allowing all other attention*
738 *heads to adapt to its change*, motivating experiments that jointly optimize all attention heads while
739 penalizing the number of heads (Voita et al., 2019), the conclusion may be more optimistic: the
740 heads obtained as a result of this optimization tend to be interpretable.

741 We think the following trade-offs apply:

- 742 • On one hand, the latter setting is more practical, since Transformers are typically not trained to
743 explicitly ensure that the model performs well when a single attention head is individually disabled;
744 rather, it would be more intuitive to think of a group of attention heads as jointly representing some
745 transformation, so when one head is disabled, other heads should be fine-tuned to adapt to the
746 change.
- 747 • On the other hand, when all other heads change too much during such fine-tuning, the resulting
748 set of attention heads no longer admit an unambiguous one-to-one map with the original set of
749 (unpruned) attention heads. As a result, the interpretability and importance obtained from the set of
750 pruned heads do not necessarily imply those properties of the original heads.

751 A comprehensive study of this question involves multi-head extensions of our theoretical results
752 (Section 3), and carefully-designed experiments that take the above-mentioned trade-offs into consid-
753 eration. We think these directions are interesting future work.

754 **C Omitted Proofs in Section 3**

755 **C.1 Proof of Theorem 1**

756 The key step is already shown in Section 3. We will restate the proof rigorously here.

757 **Theorem 4** (Perfect Balance; Theorem 1 restated). *Consider a two-layer Transformer \mathcal{T}*
 758 *with a minimal first layer with output embeddings $\{e(\tau_{i,d})\}_{d \in [D], i \in [2k]}$. Let $\theta^{(2)} :=$*
 759 *$\{W_Q^{(2)}, W_K^{(2)}, W_V^{(2)}, \text{param}(g^{(2)})\}$ denote the second layer weights.*

760 *Define the balance condition to be the condition that for any $i, j_1, j_2 \in [k]$ and $d', d_1, d_2 \in [D]$,*

$$(e(\tau_{2i-1,d'}) - e(\tau_{2i,d'-1}))^\top (W_K^{(2)})^\top W_Q^{(2)} (e(\tau_{2j_1,d_1}) - e(\tau_{2j_2,d_2})) = 0. \quad (18)$$

761 *Then, for the existence of $\{e(\tau_{i,d})\}$ and $\theta^{(2)}$ that achieves the Bayes-optimal loss for the mean*
 762 *squared error (Eqn. 4) on $\text{Dyck}_{k,D}$ for any length N , it holds that:*

763 • *If $W_V^{(2)}$ satisfies $\mathcal{P}_\perp W_V^{(2)} e(\tau_{t,d}) \neq 0, \forall t \in [k], d \in [D]$ then the balanced condition is necessary*
 764 *to show existence.*

765 • *Conversely, if the set of $2k$ encodings $\{e(\tau_{2i-1,d}), e(\tau_{2i,d})\}_{i \in [k]}$ are linearly independent for any*
 766 *$d' \in [D]$, then the balanced condition is sufficient to show existence.*

767 *Remark:* Recall that \mathcal{P}_\perp projects to the subspace orthogonal $\mathbf{1}\mathbf{1}^\top$. The assumption in the necessary
 768 condition can be intuitively understood as requiring all tokens to have nonzero contributions to the
 769 prediction, because otherwise $W_V^{(2)} e(\tau_{t,d})$ will not contribute to prediction after the LayerNorm.

770 *Proof. Necessity of the balanced condition.* By Equation (5), the attention output is directly used
 771 as the input of LayerNorm, thus we *ignore the normalization* from the softmax operation. For any
 772 prefix p ending with a closed bracket $\tau_{2j,d}$ for $d \geq 1$ and containing brackets of all depths in $[D]$, let
 773 p_m be the prefix obtained by inserting m pairs of $\{\tau_{2i-1,d'}, \tau_{2i,d'-1}\}$ for arbitrary $i \in [k]$ and depth
 774 $d' \in [D]$. Denote the projection of the unnormalized attention output by

$$u(\tau_{t_1,d_1}, \tau_{t_2,d_2}) := \mathcal{P}_\perp \exp \left(e(\tau_{t_1,d_1})^\top (W_K^{(2)})^\top W_Q^{(2)} e(\tau_{t_2,d_2}) \right) W_V^{(2)} e(\tau_{t_1,d_1}). \quad (19)$$

775 Then, by Equation (6), we have,

$$\mathcal{T}(p_m) = g^{(2)} \left(\text{LN}^{(2)} (v + m (u(\tau_{2j,d}, \tau_{2i,d'-1}) + u(\tau_{2j,d}, \tau_{2i-1,d'}))) + e(\tau_{2j,d}) \right), \quad (20)$$

776 where v denotes the unnormalized second-layer output given p as input.

777 Towards reaching a contradiction, suppose $u(\tau_{2j,d}, \tau_{2i,d'}) + u(\tau_{2j,d}, \tau_{2i-1,d'+1}) \neq 0$. Based on the
 778 continuity of the projection function and the LayerNorm Layer, we can show that $\lim_{m \rightarrow \infty} \mathcal{T}(p_m)$
 779 depend only on grammar depths d, d' and types $2j, 2i-1, 2i$, which, however, are not sufficient
 780 to determine the next-token probability from p_m , since the latter depends on the type of the last
 781 unmatched open bracket in p . This contradicts the assumption that the model achieves the Bayes-
 782 optimal loss for any length N . Hence we must have

$$u(\tau_{2j,d}, \tau_{2i,d'-1}) + u(\tau_{2j,d}, \tau_{2i-1,d'}) = 0. \quad (21)$$

783 Finally, since we assume $\mathcal{P}_\perp W_V^{(2)} e(\tau_{t,d}) \neq 0$, we conclude that

$$(e(\tau_{2i-1,d'}) - e(\tau_{2i,d'-1}))^\top (W_K^{(2)})^\top W_Q^{(2)} e(\tau_{2j,d}) = \ln \left(\frac{\|\mathcal{P}_\perp W_V e(\tau_{2i-1,d'})\|_2}{\|\mathcal{P}_\perp W_V e(\tau_{2i,d'-1})\|_2} \right).$$

784 Note that the right hand side is independent of j, d . This concludes the proof for the necessity of the
 785 condition.

786 **Sufficiency of the balance condition.** We will show a construction, using the embedding function
 787 $e(\tau_{t,d})$ as given in Equation (14). Fix any $j \in [k], d \in [D]$. By Equation (18), we can assume that
 788 there exists an $a \in \mathbb{R}^{k \times D}$ such that for $i \in [k], d', d \in [D]$, it satisfies

$$a_{i,d'} \triangleq (e(\tau_{2i-1,d'}) - e(\tau_{2i,d'-1}))^\top (W_K^{(2)})^\top W_Q^{(2)} e(\tau_{2j,d}).$$

789 We can then choose $W_V^{(2)}$ for $i \in [k]$ and $d' \in [D]$ such that

$$\begin{aligned} W_V^{(2)} e(\tau_{2i,d'-1}) &= -\exp(a_{i,d'}) \mathbf{o}_{(2i-1) \times (D-1) + d'} \\ W_V^{(2)} e(\tau_{2i-1,d'}) &= \mathbf{o}_{(2i-1) \times (D-1) + d'} \end{aligned} \quad (22)$$

790 Such $W_V^{(2)}$ is guaranteed to exist: solving for $W_V^{(2)}$ is equivalently to solving the linear equation
791 $W_V^{(2)} \mathbf{E} = \mathbf{O}$, where $\mathbf{E}, \mathbf{O} \in \mathbb{R}^{2kD \times 2kD}$ are defined according to Equation (22)⁴ and \mathbf{E} is of full
792 rank by the linear independence assumption.

793 It can be checked that choosing $W_V^{(2)}$ to satisfy Equation (22) will also make Equation (21) satisfied.
794 Hence for any prefix p of length n ending with a closed bracket $\tau_{2j,d}$ satisfying $d \geq 1$, suppose the
795 list of unmatched open brackets in p is $[\tau_{2j_1-1,1}, \tau_{2j_2-1,2}, \dots, \tau_{2j_m-1,d}]$, then suppose X is the input
796 of the second layer, we will have the last column (i.e. corresponding to the last position) of the input
797 to the LayerNorm satisfies,

$$W_V^{(2)} X \cdot \left[\sigma \left(\mathcal{C} \cdot \frac{(W_K^{(2)} X)^\top (W_Q^{(2)} X)}{\sqrt{d_a}} \right) \right]_{:,n} = \sum_{s=1}^d u(\tau_{2j_s-1,s}, \tau_{2j,d}), \quad (23)$$

798 where \mathcal{C} denotes the causal mask.

799 Finally we can choose the weights in the LayerNorm to be sufficiently small such that the largest
800 index of the last column of input to $g^{(2)}$ is determined by $X_{:,n}$. This weights can always be chosen
801 because the norm of the output of LayerNorm is bounded by 1 and $e(\tau_{t,d})$ are linearly independent,
802 hence nonzero. Then the next token probability can be determined by:

- 803 1. The last bracket in p , when p ends with an open bracket or a closed bracket with depth 0,
- 804 2. The type of last unmatched open bracket in p : suppose the grammar depth of this unmatched open
805 bracket is d , then we only need to look at indices $(2i-1) \times (D-1) + d$ for $i \in [k]$. Among values
806 of these indices, if the value is maximized at $i' \in [k]$, then the correct type of the unmatched
807 bracket is i' .

808 To complete the proof, note that the above functionality can be implemented with a combination of
809 feedforward layers. Specifically, since there are only a finite number of possible input to g , we can
810 construct a 2-layer ReLU network that memorize the values for all inputs, which requires a width
811 that is polynomial in the number of possible inputs. \square

812 C.1.1 Proof of Corollary 1

813 **Corollary 3** (Corollary 1, restated). *There exists a two-layer Transformer with uniform attention and*
814 *without position embedding (but with causal mask) that can generate the Dyck language of arbitrary*
815 *length.*

816 *Proof.* It is easy to see that the condition in Theorem 1 is satisfied. Hence it suffices to construct
817 a uniform attention first layer that can generate the embedding in Equation (14). Let $W_V^{(1)}$ be the
818 identity matrix, and suppose Z is the one-hot embeddings of a prefix p of length n , where each token
819 of type t for $t \in [2k]$ is encoded as \mathbf{o}_t . Then, the last column of Z satisfies

$$W_V^{(1)} Z \left[\sigma \left(\mathcal{C} \cdot \frac{(W_K^{(1)} Z)^\top (W_Q^{(1)} Z)}{\sqrt{d_a}} \right) \right]_{:,n} = \sum_{i=1}^{2k} \#\{\text{token of type } i \text{ in } p\} \mathbf{o}_i. \quad (24)$$

820 where \mathcal{C} denotes the causal mask.

821 The depth of the n_{th} token can then be determined by counting the number of i satisfying the value
822 of index $2i-1$ and $2i$ in the last column of Z are different by 1. Similar to the proof of Theorem 4,
823 this function can be implemented with a combination of feedforward layers and LayerNorm layers
824 and the proof is then completed. \square

⁴Specifically, $\mathbf{E} = [e(\tau_{1,1}), e(\tau_{1,2}), \dots, e(\tau_{2k,D-2}), e(\tau_{2k,D-1})]$, i.e. \mathbf{E} is the collection of all $e(\tau_{t,d})$.
 \mathbf{O} is defined such that for every d' , $\mathbf{O}_{:,t(D-1)+d'} = -\exp(a_{t/2,d'}) \mathbf{o}_{(t-1)(D-1)+d'}$ if t is even, and
 $\mathbf{O}_{:,t(D-1)+d'} = \mathbf{o}_{t(D-1)+d'}$ if t is odd.

825 **C.2 Proof of Theorem 2**

826 Let's first define a quantity for convenience of later exposition. Let u be defined as in Equation (19).
 827 For any $i \in [k]$, $d \in [D]$ and $\tilde{\mathbf{t}} \in [k]^{d-1}$, denote the quantity

$$Q(i, d, \tilde{\mathbf{t}}) := \sum_{1 \leq d' < d} u(\tau_{2i, d-1}, \tau_{2\tilde{\mathbf{t}}_{d'}-1, d'}) + u(\tau_{2i, d-1}, \tau_{2i-1, d}) + u(\tau_{2i, d-1}, \tau_{2i, d-1}), \quad (25)$$

828 where $\tilde{\mathbf{t}}_{d'}$ denotes the d'_{th} entry of $\tilde{\mathbf{t}}$. That is, $\tilde{\mathbf{t}}$ is a string of $d-1$ open brackets. Let τ_i denote a
 829 bracket of type $i \in [2k]$ without specifying the grammar depth (i.e. the grammar depth is implicit
 830 from the context), then $Q(i, d, \tilde{\mathbf{t}})$ can be considered as the unnormalized output of the second-layer
 831 attention of a Transformer on the input sequence $\tilde{\mathbf{t}} \oplus \tau_{2i-1} \tau_{2i}$ ⁵.

832 **Theorem 5** (Approximate Balance (Theorem 2 restated)). *Consider a two-layer Transformer \mathcal{T} with*
 833 *a minimal first layer trained with the mean squared error (Equation (4)). For any $\gamma, N > 0$ and*
 834 *sufficiently small ϵ , suppose $g^{(2)}$ is γ -Lipschitz, and suppose the set of second-layer weights $\bar{\theta}_N^{(2)}$*
 835 *satisfies that $\mathcal{L}(\mathcal{T}[\bar{\theta}_N^{(2)}], \mathcal{D}_{q, k, D, N}) \leq q^{-N} \epsilon$. Then, there exists a constant $C_{\gamma, \epsilon, D}$, such that for any*
 836 *$0 \leq d' \leq D, 1 \leq d \leq D, i, j \in [k]$, it holds that*

$$S_{d, d', i, j}[\bar{\theta}_N^{(2)}] \leq \frac{C_{\gamma, \epsilon, D}}{N} P_{d, j}[\bar{\theta}_N^{(2)}]. \quad (26)$$

837 where

$$S_{d, d', i, j}[\bar{\theta}^{(2)}] = \left\| u(\tau_{2j, d}, \tau_{2i, d'}) + u(\tau_{2j, d}, \tau_{2i-1, d'+1}) \right\|_2, \quad (27)$$

$$P_{d, j}[\bar{\theta}^{(2)}] = \min_{\mathbf{t}' \in [k]^{d-1}, \mathbf{t}'_d \neq \mathbf{t}_d} \|Q(i, d, \mathbf{t}')\|_2, \quad (28)$$

838 for $\mathbf{t} = \arg \min_{\mathbf{t}' \in [k]^{d-1}} \|Q(2j, d, \mathbf{t}')\|_2$ ⁶.

839 *Proof.* The key idea is similar to the proof of necessity in Theorem 1. That is, we will construct
 840 two input sequences with different next-word distributions, and show that the approximate balance
 841 condition must hold so that inserting (a bounded number of) pairs of matching brackets does not
 842 collapse the two predicted distributions given by the Transformer.

843 **Constructing the input sequences.**

844 Let $\mathbf{t} := \arg \min_{\tilde{\mathbf{t}} \in [k]^{d-1}} \|Q(2j, d, \tilde{\mathbf{t}})\|_2$, and let \mathbf{t}' denote the prefix that minimizes $\|Q(2j, d, \tilde{\mathbf{t}})\|_2$
 845 subject to the constraint that \mathbf{t}' must differ from \mathbf{t} in the last (i.e. $(d-1)_{th}$) position, i.e.

$$\mathbf{t}' = \arg \min_{\tilde{\mathbf{t}}' \in [k]^{d-1}, \tilde{\mathbf{t}}'_{d-1} \neq \tilde{\mathbf{t}}_{d-1}} Q(2j, d, \tilde{\mathbf{t}}').$$

846 The motivation for such choices of \mathbf{t}, \mathbf{t}' is that since they differ at least by the last position which
 847 is an open bracket, they must lead to different next-word distributions. Note also that $P_{d, j}[\bar{\theta}^{(2)}] =$
 848 $\|Q(2j, d, \mathbf{t}')\|$.

849 With the above definition of \mathbf{t}, \mathbf{t}' , consider two valid Dyck prefixes p_1 and p_2 with length no
 850 longer than N , defined as follows: for any $d, d' \in [D], i, j \in [k]$, consider a common prefix

851 $p = \underbrace{\tau_{2i-1} \dots \tau_{2i-1}}_{d' \text{ open brackets}} \underbrace{\tau_{2i-1} \tau_{2i} \dots \tau_{2i-1} \tau_{2i}}_{\lfloor \frac{N-2d'-2d}{2} \rfloor \text{ pairs}} \underbrace{\tau_{2i} \dots \tau_{2i}}_{d' \text{ closed brackets}},$ and set:

$$p_1 = p \oplus \mathbf{t} \oplus \tau_{2j-1} \tau_{2j},$$

$$p_2 = p \oplus \mathbf{t}' \oplus \tau_{2j-1} \tau_{2j}.$$

852 In the following, we will show that the approximate balance condition must hold for the predictions
 853 on p_1, p_2 to be sufficiently different.

⁵ $s \oplus t$ denotes the concatenation of two strings s, t , same as in Equation (14)-(16). The concatenation of two tokens τ_i, τ_j is simply written as $\tau_i \tau_j$.

⁶ *Erratum:* This definition of $P_{d, j}[\bar{\theta}^{(2)}]$ is slightly different from the one in the original main paper submitted on May 17th. The definition here and in the current main paper have been corrected.

854 **Bounding the difference in Transformer outputs.** The Transformer outputs on p_1, p_2 satisfies

$$\|\mathcal{T}[\bar{\theta}_N^{(2)}](p_1) - \mathcal{T}[\bar{\theta}_N^{(2)}](p_2)\|_2 \geq 1 - \text{TV}(p_1, p_2) - o_\epsilon(1) = \Omega(1), \quad (29)$$

855 where $\text{TV}(p_1, p_2)$ denotes the TV distance in the next-word distributions from p_1 and p_2 , and $o_\epsilon(1)$
856 means the term will go to zero for sufficiently small ϵ . The former is bounded by the construction
857 of p_1, p_2 . The latter is bounded because of the assumption on $\bar{\theta}_N^{(2)}$, which states that the set of
858 second-layer weights $\bar{\theta}_N^{(2)}$ satisfies that $\mathcal{L}(\mathcal{T}[\bar{\theta}_N^{(2)}], \mathcal{D}_{q,k,D,N}) \leq q^{-N}\epsilon$ with sufficiently small ϵ .

859 Define by A_p the contribution of p to the attention output (before LayerNorm) of the last position of
860 p_1, p_2 , i.e.

$$\begin{aligned} A_p = & \sum_{0 \leq d' < d} (u(\tau_{2j,d-1}, \tau_{2i,d'}) + u(\tau_{2j,d-1}, \tau_{2i-1,d'+1})) \\ & + \lfloor \frac{N-2d'-2d}{2} \rfloor (u(\tau_{2j,d-1}, \tau_{2i,d'}) + u(\tau_{2j,d-1}, \tau_{2i-1,d'+1})). \end{aligned} \quad (30)$$

861 The attention outputs (before LayerNorm) of p_1, p_2 , denoted by $A(p_1)$ and $A(p_2)$, satisfy that

$$\begin{aligned} \mathcal{P}_\perp A(p_1) &= \mathcal{P}_\perp (A_p + Q(2j, d, \mathbf{t})), \\ \mathcal{P}_\perp A(p_2) &= \mathcal{P}_\perp (A_p + Q(2j, d, \mathbf{t}')). \end{aligned} \quad (31)$$

862 Note that for any prefix p' , $\mathcal{T}[\bar{\theta}_N^{(2)}](p') = g^{(2)}(\mathcal{P}_\perp A(p'))$. Then, since $g^{(2)}$ is γ -Lipschitz,

$$\left\| \frac{\mathcal{P}_\perp A(p_1)}{\|\mathcal{P}_\perp A(p_1)\|_2} - \frac{\mathcal{P}_\perp A(p_2)}{\|\mathcal{P}_\perp A(p_2)\|_2} \right\|_2 \geq \frac{1 - \text{TV}(p_1, p_2) - O_\epsilon(1)}{\gamma} = \Omega_{\gamma,\epsilon}(1). \quad (32)$$

863 We show that A_p should not be too much larger in norm than $Q(2j, d, \mathbf{t})$ or $Q(2j, d, \mathbf{t}')$. First let's
864 state a helper lemma about the contrapositive:

865 **Lemma 1.** For any $\epsilon > 0$, there exists a constant R_ϵ , such that for any $a, b \in \mathbb{R}^d$ and any $r \in \mathbb{R}^d$
866 such that $\|r\|_2 \geq R_\epsilon \cdot \max\{\|a\|_2, \|b\|_2\}$, it holds that

$$\left\| \frac{a+r}{\|a+r\|_2} - \frac{b+r}{\|b+r\|_2} \right\|_2 \leq \epsilon.$$

867 *Proof.* Denote $r_0 := \max\{\|a\|_2, \|b\|_2\}$. Then $R_\epsilon := \frac{4r_0}{\epsilon} + 1$ suffices:

$$\begin{aligned} & \left\| \frac{r+a}{\|r+a\|_2} - \frac{r+b}{\|r+b\|_2} \right\| \leq \|r\| \cdot \left| \frac{1}{\|r+a\|} - \frac{1}{\|r+b\|} \right| + \frac{\|a\|}{\|r+a\|} + \frac{\|b\|}{\|r+b\|} \\ & \leq \|r\| \cdot \left(\frac{1}{\|r\| - r_0} - \frac{1}{\|r\| + r_0} \right) + \frac{2r_0}{\|r\| - r_0} \\ & = \frac{2r_0}{\|r\| - r_0} \cdot \left(\frac{\|r\|}{\|r\| + r_0} + 1 \right) \leq \frac{4r_0}{\|r\| - r_0} \leq \frac{4r_0}{R_\epsilon - r_0} \leq \epsilon. \end{aligned}$$

868 □

869 Lemma 1 implies that if A_p is too large, then the output on p_1, p_2 (Equation (32)) won't be sufficiently
870 different. Let $P_{d,j}[\bar{\theta}_N^{(2)}]$ be defined as in Equation (27) and let R_ϵ be the constant in Lemma 1, we
871 need to bound $\|\mathcal{P}_\perp A_p\|$ by

$$\|\mathcal{P}_\perp A_p\|_2 \leq R_\epsilon \|P_{d,j}[\bar{\theta}_N^{(2)}]\|_2. \quad (33)$$

872 As Equation (33) holds for p with any d, d' , by an induction on d' (from 1 to d) on the second term in
873 Equation (30), one can show that there exists C (depending on R_ϵ), such that,

$$S_{d,d',i,j} = \|u(\tau_{2j,d-1}, \tau_{2i,d-1}) + u(\tau_{2j,d-1}, \tau_{2i-1,d-1})\| \leq \frac{C}{N} \|P_{d,j}[\bar{\theta}_N^{(2)}]\|_2. \quad (34)$$

874 The proof of Equation (34) can be carried out inductively over d from 1 to D . □

875 *Proof of Corollary 2.* This proof is in fact a direct combination of Theorems 1 and 2. By Theorem 1
876 we know there exists a weight $\theta^{(2)*}$ that can reach zero loss for arbitrarily length N . Then it holds that
877 $\|\theta_{\lambda,N}\|_2 \leq \|\theta^*\|$ as $\theta_{\lambda,N}$ minimizes the regularized loss. Notice bounded weight implies bounded
878 Lipschitzness of $g^{(2)}$, The rest follows as Theorem 2. □

879 **C.3 Proof of Theorem 3 – Indistinguishability from a single component**

880 We now show the limitation of interpretability from a single component, using a Lottery-Ticket-style
881 argument by pruning from large random Transformers.

882 For this section only, we will make the following modifications to the Transformer architecture in (6):

- 883 • We lower bound the normalization factor in the LayerNorm by some constant C , namely we
884 consider:

$$\text{LN}_C(x) = \frac{\mathcal{P}_\perp x}{\max\{\|\mathcal{P}_\perp x\|_2, C\}},$$

885 We need this assumption for technical reasons (to make the LayerNorm Lipschitz). We note that
886 thresholding at C is also a common practice empirically due to numerical stability concerns.

- 887 • We assume all affine layers and linear head in the Transformer have zero bias. This is mainly for
888 technical convenience, and was also assumed in prior works on theoretical analysis of the lottery
889 ticket hypothesis (Pensia et al., 2020). Note that this is not a restriction since bias can be removed
890 with homogeneous coordinates.

891 We will also consider a modified projection function $g_{\text{large}}^{(l)}$ consisting of a 4-layer MLP, which will
892 be used in the to-be-pruned large random Transformers:

$$g_{\text{large}}(x) = \text{LN}(W_4 \text{ReLU}(W_3 \text{ReLU}(W_2 \text{ReLU}(W_1 x)))) + x, \quad (35)$$

893 where $W_1, W_4^\top \in \mathbb{R}^{w_{\text{large}} \times m_{\text{large}}}$, $W_2, W_3 \in \mathbb{R}^{w_{\text{large}} \times w_{\text{large}}}$, for some $w_{\text{large}}, m_{\text{large}}$.

894 We are now ready to state the main theorem of this section:

895 **Theorem 6** (Indistinguishability From a Single Component (Theorem 3 restated)). *Con-*
896 *sider a L -layer Transformer \mathcal{T} with embedding dimension m , width w and $g^{(k)}(x) =$
897 $\text{LN}_C(W_2^{(k)} \text{ReLU}(W_1^{(k)} x)) + x$. Suppose $\|W\|_2 = O(1)$ for every weight matrix W in \mathcal{T} .
898 For $\delta \in (0, 1)$, consider a larger random Transformer $\mathcal{T}_{\text{large}}$ with $4L$ layers, embedding dimension
899 $m_{\text{large}} = O(d \log(d/\delta))$, and width $w_{\text{large}} = O(\max\{m, w\} \log \frac{wmLN}{\epsilon\delta})$, and projection function
900 g_{large} , whose weights are randomly sampled as $W_{i,j} \sim U(-1, 1)$ for every $W \in \mathcal{T}_{\text{large}}$.*

901 *Then, with probability $1 - \delta$ over the randomness of $\mathcal{T}_{\text{large}}$, we can obtain a nonstructural pruning*
902 *(Definition 2) of $\mathcal{T}_{\text{large}}$, denoted as $\mathcal{T}'_{\text{large}}$, which ϵ -approximates \mathcal{T} . That is, $\forall \mathbf{X} \in \mathbb{R}^{m \times N}$ with*
903 *$\|\mathbf{X}_{:,i}\|_2 \leq 1$, $\forall i \in [N]$,*

$$\|\mathcal{T}'_{\text{large}}(\mathbf{X}) - \mathcal{T}(\mathbf{X})\|_2 \leq \epsilon.$$

904 *Moreover, pick any weight matrix W in $\mathcal{T}_{\text{large}}$, with probability $1 - \delta$, for any smaller Transformers*
905 *$\mathcal{T}_1, \mathcal{T}_2$ satisfying same conditions as \mathcal{T} , we have two pruned Transformers $\mathcal{T}_{\text{large},1}, \mathcal{T}_{\text{large},2}$ based on*
906 *$\mathcal{T}_{\text{large}}$, such that they coincide on the pruned weight of W , and $\mathcal{T}_{\text{large},i}$ ϵ -approximate \mathcal{T}_i , $\forall i \in \{1, 2\}$.*

907 *Proof.* We will first introduce some notation. For vector $x \in \mathbb{R}^a$ and $y \in \mathbb{R}^b$, we will use $x \oplus y$ to
908 denote their concatenation. We will use 0^a to denote the all-zero vector with dimension a . We will
909 also assume without loss of generality that $w \geq 2d$.⁷

910 In the following, a *random network* refers to a network whose weights have entries sampled from a
911 uniform distribution, i.e. $W_{i,j} \sim U(-1, 1)$ for every weight W in the random network.

912 We will first recall Lemma 2 from Pensia et al. (2020) which shows that a pruned 2-layer random
913 network can approximate a linear function.

914 **Lemma 2** (Theorem 1 of Pensia et al. (2020)). *Let $W \in \mathbb{R}^{d' \times d}$, $\|W\|_2 = O(1)$, then for $\sigma \in$
915 $\{\text{ReLU}, \mathcal{I}\}$, for a random network $g(x) = W_2 \sigma(W_1 x)$ with $W_2 \in \mathbb{R}^{d' \times h}$, $W_1 \in \mathbb{R}^{h \times d}$ for hidden
916 dimension $h = O(d \log(\frac{dd'}{\min\{\epsilon, \delta\}}))$, with probability $1 - \delta$, there exists boolean matrices M_1, M_2 ,
917 such that for any $x \in \mathbb{R}^d$, $\|x\|_2 = O(1)$,*

$$\|(M_2 \odot W_2) \sigma((M_1 \odot W_1)x) - Wx\| \leq \epsilon,$$

918 *where \odot denotes the Hadamard product.*

⁷We can always pad dimensions if w is too small.

919 We will use the following helper lemma:

- 920 1. A pruned 4-layer projection function of a Transformer layer can approximate a 2-layer ReLU
921 network applied to each token (Lemma 3).
- 922 2. A pruned random Transformer layer can approximate a linear function applied independently to
923 each token (Lemma 4).
- 924 3. Two pruned random Transformer layers can approximate a fixed smaller Transformer layer.
925 (Lemma 7)

926 We can now prove the theorem.

927 To show ϵ -approximation, we can prune the large Transformer to approximate the smaller Transformer
928 layer by layer by Lemma 7. The linear head $W^{(head)}$ can be pruned using Lemmas 4 and 6, and
929 combined with one layer of the Transformer, the linear head of the smaller Transformer can be
930 approximated.

931 Further, as we only need 2 layers to approximate one layer of the smaller Transformer, for an arbitrary
932 layer l , we can prune the layer l of the large Transformer to ϵ -approximate identity function. This
933 then concludes the proof for indistinguishability from single components. \square

934 C.3.1 Helper lemmas for Theorem 6

935 We first show that a pruned 4-layer projection function in a Transformer layer can approximate a
936 2-layer ReLU network applied to each token:

937 **Lemma 3.** *Under the condition of Theorem 6, for any two matrices $W_1 \in \mathbb{R}^{d \times w}, W_2 \in$
938 $\mathbb{R}^{w \times d}, \|W_1\|_2, \|W_2\|_2 = O(1)$, for any $\delta \in (0, 1)$ and $l \in [4L]$, with probability $1 - \delta$, there exists an
939 unstructured pruning of $g_{\text{large}}^{(l)}, g_{\text{large}}^{(l)'}$, satisfying that $\forall \mathbf{X} \in \mathbb{R}^{m \times N}$ with $\|\mathbf{X}_{:,i}\|_2 = O(1), \forall i \in [N]$,*

$$\forall \mathbf{R} \in \mathbb{R}^{(m_{\text{large}} - m) \times N}, \left\| \left(g_{\text{large}}^{(l)'} \left(\begin{bmatrix} \mathbf{X} \\ \mathbf{R} \end{bmatrix} \right) \right)_{1:m,:} - W_2 \text{ReLU}(W_1 \mathbf{X}) \right\|_2 \leq \epsilon,$$

940 where $M_{1:m,:}$ denotes the first m rows of a matrix M .

941 *Proof.* Recall the definition of the projection function of a Transformer layer is

$$g_{\text{large}}^{(l)}(x) = \text{LN} \left(W_4^{(l)} \text{ReLU} \left(W_3^{(l)} \text{ReLU} \left(W_2^{(l)} \text{ReLU} \left(W_1^{(l)} x \right) \right) \right) \right) + x.$$

942 We will prune the LayerNorm by setting it to the identity. Now we only need to show that there exists
943 boolean matrices M_1, M_2, M_3, M_4 , such that,

$$\left\| \left(M_4 \odot W_4^{(l)} \text{ReLU} \left((M_3 \odot W_3^{(l)}) \text{ReLU} \left((M_2 \odot W_2^{(l)}) \text{ReLU} \left((M_1 \odot W_1^{(l)}) \begin{bmatrix} \mathbf{X} \\ \mathbf{R} \end{bmatrix} \right) \right) \right) \right)_{1:m,:} - W_2 \text{ReLU}(W_1 \mathbf{X}) - \mathbf{X} \right\|_2 \leq \epsilon.$$

944 We can first choose

$$(M_1)_{:(m+1, \dots, m_{\text{large}})} = 0, (M_4)_{(m+1, \dots, m_{\text{large}}),:} = 0, \\ (M_2)_{(w+2m+1, \dots, w_{\text{large}}),:} = 0, (M_3)_{:(w+2m+1, \dots, w_{\text{large}})} = 0$$

945 Then by Lemma 2, there exists boolean matrices M_1, M_2, M_3, M_4 satisfying previous constraint,
946 such that,

$$\left\| \left((M_2 \odot W_2^{(l)}) \text{ReLU} \left((M_1 \odot W_1^{(l)}) \begin{bmatrix} \mathbf{X} \\ \mathbf{R} \end{bmatrix} \right) \right)_{1:w+2m} - \begin{bmatrix} W_1 \\ \mathcal{I} \\ -\mathcal{I} \end{bmatrix} \mathbf{X} \right\| \leq \frac{\epsilon}{4}.$$

$$\forall \mathbf{X}' \in \mathbb{R}^{(w+2m) \times N}, \left\| (M_4 \odot W_4^{(l)}) \text{ReLU} \left((M_3 \odot W_3^{(l)}) \begin{bmatrix} \mathbf{X}' \\ \mathbf{R}' \end{bmatrix} \right) - [W_2 \quad \mathcal{I} \quad -\mathcal{I}] \mathbf{X}' \right\| \leq \frac{\epsilon}{4} \cdot \frac{\max_{i \in [N]} \|\mathbf{X}'_{:,i}\|_2}{\|W_1\|_2}.$$

947 This then concludes the proof. \square

948 Based on the above lemma, we can prove that a pruned Transformer layer can approximate a linear
 949 function applied independently to each token.

950 **Lemma 4.** *Under the conditions in Theorem 6, for any matrix $W \in \mathbb{R}^{m \times m}$, $\|W\|_2 = O(1)$,
 951 $\delta \in (0, 1)$ and $l \in [4L]$, with probability $1 - \delta$, there exists an unstructured pruning of $\mathcal{T}_{\text{large}}^{(l)}$, $\mathcal{T}_{\text{large}}^{(l)'$
 952 satisfying that $\forall \mathbf{X} \in \mathbb{R}^{m \times N}$ with $\|\mathbf{X}_{:,i}\|_2 = O(1)$, $\forall i \in [N]$, we have*

$$\forall \mathbf{R} \in \mathbb{R}^{(m_{\text{large}} - m) \times N}, \left\| \left(\mathcal{T}_{\text{large}}^{(l)'} \left(\begin{bmatrix} \mathbf{X} \\ \mathbf{R} \end{bmatrix} \right) \right)_{1:m,:} - W\mathbf{X} \right\|_2 \leq \epsilon.$$

953 *Proof.* Recall that given an input \mathbf{X}' , a Transformer layer computes $\mathcal{T}_{\text{large}}^{(l)}(\mathbf{X}') =$
 954 $\mathfrak{g}_{\text{large}}^{(l)} \left(\text{LN} \left(W_V^{(l)} \mathbf{X}' \text{Attn}(\mathbf{X}') \right) + \mathbf{X}' \right)$, where $\text{Attn}(\mathbf{X}') := \sigma \left(\mathcal{C} \cdot \frac{(W_K^{(l)} \mathbf{X}')^\top (W_Q^{(l)} \mathbf{X}')}{\sqrt{d_a}} \right)$ computes
 955 the attention pattern. Lemma 3 already shows that $\mathfrak{g}_{\text{large}}^{(l)}$ can approximate a linear transformation; it
 956 remains to show that the linear transformation can compute $W\mathbf{X}$.

957 We can first choose two matrices $W_1 \in \mathbb{R}^{w \times m}$, $W_2 \in \mathbb{R}^{m \times w}$ satisfying that

$$\begin{aligned} W_1 &= [\mathcal{I}_m, -\mathcal{I}_m, 0^{m \times (w-2m)}]^\top. \\ W_2 &= [W, -W, 0^{m \times (w-2m)}] \end{aligned}$$

958 Then we have that $\|W_1\|_2, \|W_2\|_2 = O(1)$ and $W_2 \text{ReLU}(W_1 \mathbf{X}) = W\mathbf{X}$. We can then turnoff
 959 the LayerNorm after the attention module and prune W_V to be 0, which effectively removes the
 960 effect of attention and rely solely on the residual link. The proof can now be completed by applying
 961 Lemma 3. \square

962 We will then show that two pruned Transformer layers can approximate a fixed smaller Transformer
 963 layer. The key technical difficulty is approximating the attention module and bounding the error of
 964 the approximation after LayerNorm. We will first show a lemma showing the Lipschitzness of the
 965 LayerNorm (with cutoff at some constant C).

966 **Lemma 5.** *For LayerNorm function defined as $\text{LN}(x) = \frac{\mathcal{P}_\perp x}{\max\{\|\mathcal{P}_\perp x\|_2, C\}}$, $x \in \mathbb{R}^m$, there exists
 967 constant C_1 depending on C , such that for any $x, y \in \mathbb{R}^m$, it holds that,*

$$\left\| \text{LN}(x) - \text{LN}(y) \right\|_2 \leq C_1 \|x - y\|_2.$$

968 *Proof.* We will proceed by a case analysis:

- 969 1. If $\|\mathcal{P}_\perp x\|_2, \|\mathcal{P}_\perp y\|_2 \leq C$, then $\left\| \text{LN}(x) - \text{LN}(y) \right\|_2 = \frac{\|\mathcal{P}_\perp x - \mathcal{P}_\perp y\|_2}{C} \leq \frac{1}{C} \|x - y\|_2.$
- 970 2. If $\|\mathcal{P}_\perp x\|_2, \|\mathcal{P}_\perp y\|_2 > C$, then $\left\| \text{LN}(x) - \text{LN}(y) \right\|_2 = \frac{\|\mathcal{P}_\perp x - \mathcal{P}_\perp y\|_2}{\|\mathcal{P}_\perp y\|_2} + \left| 1 - \frac{\|\mathcal{P}_\perp x\|_2}{\|\mathcal{P}_\perp y\|_2} \right| \leq \frac{2}{C} \|x - y\|_2.$
- 971 3. If $\|\mathcal{P}_\perp x\|_2 < C$ and $\|\mathcal{P}_\perp y\|_2 > C$, then $\left\| \text{LN}(x) - \text{LN}(y) \right\|_2 = \frac{\|\mathcal{P}_\perp x - \mathcal{P}_\perp y\|_2}{\|\mathcal{P}_\perp y\|_2} + \left| \frac{\|\mathcal{P}_\perp x\|_2}{C} - \right.$
 972 $\left. \frac{\|\mathcal{P}_\perp x\|_2}{\|\mathcal{P}_\perp y\|_2} \right| \leq \frac{2}{C} \|x - y\|_2.$

973 The cases exhaust all possibilities, thus the proof is completed. \square

974 We also need to show there exists a pruning of the value matrix in $\mathcal{T}_{\text{large}}$ such that it has eigenvalues
 975 with magnitude $\Theta(1)$.

976 **Lemma 6.** *For a matrix $W \in \mathbb{R}^{w_{\text{large}} \times w_{\text{large}}}$, with probability at least $1 - \delta$, there exists a pruning of
 977 W , named W' , such that all the nonzero entries is contained in a $d \times d$ submatrix of W' that satisfies
 978 that (1) all its eigenvalues are within $(\frac{1}{2}, 1)$, (2) the index of row specifying the submatrix and the
 979 index of column specifying the submatrix are disjoint.*

980 *Proof.* As $w_{\text{large}} = \Omega(m \log(\frac{d}{\delta}))$, hence we can split $W_{1:\lceil m_{\text{large}}/2 \rceil, \lceil m_{\text{large}}/2 \rceil + 1:m_{\text{large}}}$ into $(m \times (m$
981 blocks, each with width at least $O(\log(\frac{m}{\delta}))$ ⁸. Within each block, with probability $1 - \frac{\delta}{m}$, there
982 exists at least one entry that has value at least $\frac{1}{2}$. We can then choose d disjoint entries in W that
983 are all at least $\frac{1}{2}$, indexed with $\{(a_i, b_i)\}_{i \in [d]}$ where $a_i < a_j$ and $b_i < b_j$ for $i < j$. We can then
984 prune all other entries to zero. Consider the submatrix defined by entries (a, b) for $a \in \{a_i\}_{i \in m}$
985 and $b \in \{b_i\}_{i \in m}$. Then, this submatrix will be diagonal and contains eigenvalues within $(\frac{1}{2}, 1)$.
986 Further $\{a_i\}_{i \in m}$ and $\{b_i\}_{i \in m}$ must be disjoint because $a_i \leq \lceil m_{\text{large}}/2 \rceil < b_i$. The proof is then
987 completed. \square

988 Next, we show that two random Transformer layers can be pruned to approximate a given Transformer
989 layer.

990 **Lemma 7.** *Under the condition of Theorem 3, for any matrix $W \in \mathbb{R}^{d \times d}$, $\|W\|_2 = O(1)$, $\delta \in (0, 1)$
991 and $t \in [4L]$, for any $l \in [L]$, with probability $1 - \delta$, there exists an unstructured pruning of
992 $\mathcal{T}_{\text{large}}^{(t)}, \mathcal{T}_{\text{large}}^{(t+1)}$, named $\mathcal{T}_{\text{large}}^{(t)'}, \mathcal{T}_{\text{large}}^{(t+1)'}$, satisfying that $\forall \mathbf{X} \in \mathbb{R}^{d \times N}$ with $\|\mathbf{X}_{:,i}\|_2 = O(1)$, $\forall i \in [N]$,*

$$\forall \mathbf{R} \in \mathbb{R}^{(m_{\text{large}} - m) \times N}, \|\mathcal{T}_{\text{large}}^{(t+1)'}\left(\mathcal{T}_{\text{large}}^{(t)'}([\mathbf{X}_{:,i} \oplus \mathbf{R}_{:,i}]_{i \in [N]})\right)_{1, \dots, m} - \mathcal{T}^{(l)}(\mathbf{X})\|_2 \leq \epsilon.$$

993 *Proof.* We will prune the larger transformer in the following order.

994 1. We will prune $W_V^{(t+1)}$ according to Lemma 6 and name the pruned matrix $W_V^{(t+1)'}$. By Lemma 6,
995 all the nonzero entries is contained in a $d \times d$ submatrix of W' that satisfies that all its eigenvalues
996 are within $(\frac{1}{2}, 1)$. We will prune $W_V^{(t+1)}$ in this way, named $W_V^{(t+1)'}$ and assume WLOG the
997 submatrix is the one specified by row $1 \dots d$ and column $d + 1 \dots 2d$ and name the submatrix as
998 W .

999 2. We will then prune $\mathcal{T}_{\text{large}}^{(t)}$ according to Lemma 4 to output ϵ -approximation of $X_{:,i} \oplus$
1000 $(W^{-1} \mathcal{P}_{\perp} W_v^{(l)} X_{:,i}) \oplus \mathbf{A}_{:,i}$ for some vectors $\mathbf{A}_{:,i}$. As W is defined as the submatrix pruned by
1001 $W_V^{(t+1)}$, it holds that $W_V^{(t+1)'}\left(X_{:,i} \oplus (W^{-1} W_v^{(l)} X_{:,i}) \oplus \mathbf{A}_{:,i}\right) = \mathcal{P}_{\perp} W_v^{(l)} X_{:,i} \oplus 0^{m_{\text{large}} - m}$.

1002 3. We will then prune $W_K^{(t+1)}$ and $W_Q^{(t+1)}$ according to Lemma 2 to approximate attention patterns.
1003 We will choose boolean matrix M_K, M_Q such that for any $x \in \mathbb{R}^d$ and $a \in \mathbb{R}^{m_{\text{large}} - m}$,

$$\|(M_K \odot W_K^{(t+1)})^{\top} (M_Q \odot W_Q^{(t+1)}(x \oplus a)) - ((W_K^{(l)})^{\top} W_Q^l x) \oplus 0^{m_{\text{large}} - m}\| \leq \epsilon \|x\|_2.$$

1004 We can then have that the attention pattern for the large transformer at layer $t + 1$ can approximate
1005 the small one. That is, for any $x \in \mathbb{R}^d$, $\|x\|_2 = O(1)$ and $a \in \mathbb{R}^{m_{\text{large}} - m}$,

$$\left\| \sigma\left((x \oplus a)^{\top} (M_K \odot W_K^{(t+1)})^{\top} (M_Q \odot W_Q^{(t+1)}(x \oplus a))\right) - \sigma\left(x^{\top} \left((W_K^{(l)})^{\top} W_Q^l x\right)\right) \right\| \leq O(\epsilon).$$

1006 Combined with previous approximation on $W_V^{(t+1)'}\left(X_{:,i} \oplus (W^{-1} W_v^{(l)} X_{:,i}) \oplus \mathbf{A}_{:,i}\right)$ and the
1007 Lipschitzness of the LayerNorm, we have that the first m dimensions of the output after LayerNorm
1008 of the large Transformer at layer $t + 1$ can ϵ -approximate the output after LayerNorm of the
1009 smaller Transformer at layer l .

1010 4. We will finally prune the MLP in the projection function of $\mathcal{T}_{\text{large}}^{(t+1)'}$ to approximate $\mathcal{P}_{\perp} f^{(l)}$ with
1011 $f^{(l)}$ being the MLP in the projection function of the projection function of $\mathcal{T}^{(l)}$.

1012 The proof is then complete. \square

⁸ $O(\cdot)$ hides absolute constants arising from the change of basis in the logarithm.

1013 D Experiments

1014 D.1 Training Details

1015 For Figure 1, we train 2-layer standard GPT on Dyck_{2,4} with sequence length no longer than 28. For
1016 (a), we train with hidden dimension and network width 200 and learning rate 3e-4. For (b), (c), (d),
1017 we train with hidden dimension and FFN width 50 and learning rate 3e-3.

1018 For Figure 2, for (a), we train 1-layer transformer without residual link, FFN and the final LayerNorm
1019 before the linear head. The hidden dimensions and FFN widths are fixed as 500. For (a), we train the
1020 network with learning rate 1e-2 and for (b), (c), (d) we train the network with learning rate 3e-3.

1021 D.2 Additional Results on Dyck Prefix

1022 In the experiment presented in the main text, we perform experiments on complete Dyck sequences,
1023 which is a special case of Dyck prefixes. In this section, we present additional experiments on Dyck
1024 prefixes Dyck_{2,4,28}.

1025 **Attention Patterns** We first perform experiments on attention patterns. The qualitative results
1026 are shown in Figures 7 and 9. We can observe that the attention patterns are still diverse and do
1027 not commonly show stack-like patterns. We also calculate the *attention variation*⁹, and find that
1028 the attention variation is 0.34, based on 30 models with a minimal first layer and different random
1029 seeds. In contrast, for models with a standard first layer and without position encodings, the attention
1030 variation is surprisingly high, reaching 14.51. The high value is caused by the large distance between
1031 attention patterns like Figure 7 (c) and (d); that is, between patterns that attend more to the current
1032 positions, and patterns that attend more heavily to the initial position. The difference is even increased
1033 when we consider longer sequence (Figure 8). Similarly, the variation is also high for models with
1034 linear position embedding, reaching 11.92. This shows that the attention patterns are still diverse and
1035 do not commonly show stack-like patterns.

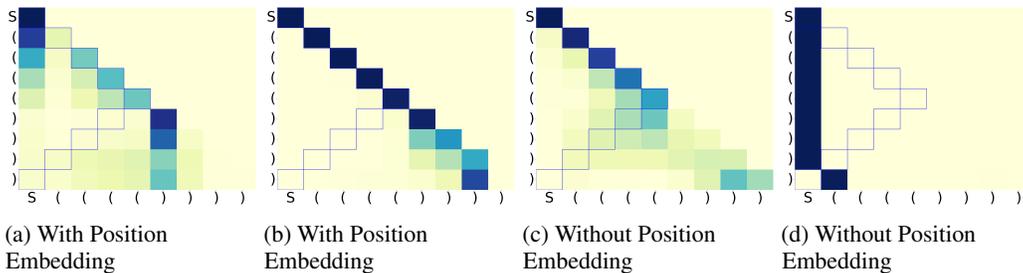


Figure 7: **Second-layer attention patterns of two-layer Transformers on Dyck Prefix:** Models for (a),(b) are under the same setup but different random seeds; similarly for (c),(d). All models reach $\geq 97\%$ accuracy (defined in Section 4.1). In the heatmap, darker color indicates larger value. As we can observe, the attention patterns still show much variance.

1036 **Balanced Violations** We also test the relationship with the balance violation with length general-
1037 ization on Dyck prefixes, similar to Figure 3. We observe that although the negative correlation is not
1038 presented as in the case of Dyck sequences, contrastive regularization still helps reduce the balance
1039 violation and significantly improve the length generalization performance. This shows that for Dyck
1040 prefixes, while the balance violation may not be predictive of the length generalization performance,
1041 it is still possible to reduce the balance violation and improve the length generalization performance.
1042 The results are shown in Figure 10.

⁹Recall from Section 4.1 that the attention variation between two attention patterns $A_1, A_2 \in \mathbb{R}^{N \times N}$ is defined as $\text{Variation}(A_1, A_2) = \|A_1 - A_2\|_F^2$.

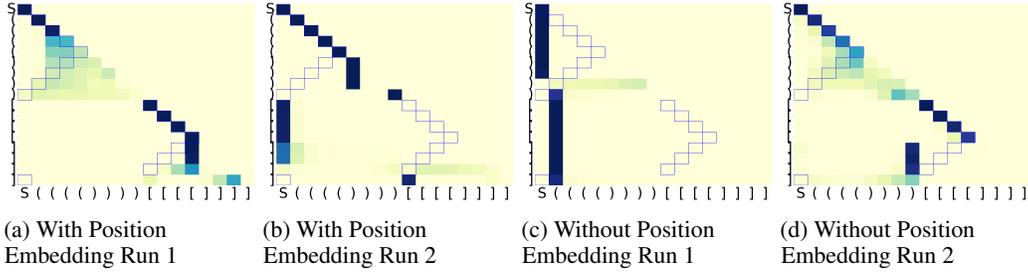


Figure 8: **Second-layer attention patterns of two-layer Transformers on Longer Dyck Prefix:** Models for (a),(b) are under the same setup but different random seeds. All models reach $\geq 97\%$ accuracy (defined in Section 4.1). In the heatmap, darker color indicates larger value.

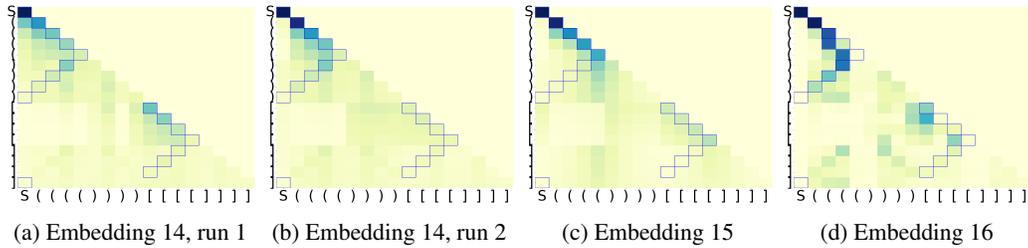


Figure 9: **Second-layer attention patterns of two-layer Transformers with a minimal first layer:** (a), (b) are based on embedding 14 with different random seeds. (c), (d) are based on embedding 15 and 16. Different embedding functions lead to diverse attention patterns, most of which are not stack-like.

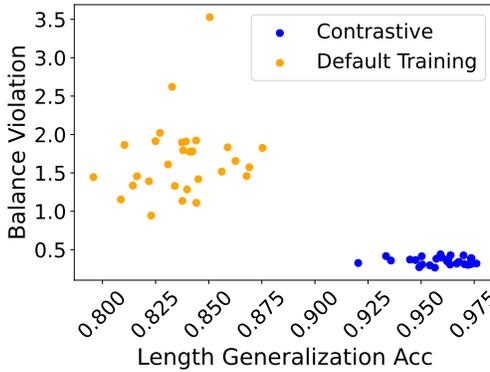


Figure 10: **Relationship Between Balance Violation and Length Generalization.** Accuracy from Transformers with minimal first layer with embedding 14, using both standard training and contrastive regularization (Equation (17)). We again observe that contrastive regularization helps reduce the balance violation and improve the length generalization performance.