

# Language Models Grow Less Humanlike beyond Phase Transition

Tatsuya Aoyama Ethan Gotlieb Wilcox

Department of Linguistics, Georgetown University

{ta571, ethan.wilcox}@georgetown.edu

## Abstract

LMs’ alignment with human reading behavior (i.e. *psychometric predictive power*; *PPP*) is known to improve during pretraining up to a tipping point, beyond which it either plateaus or degrades. Various factors, such as word frequency, recency bias in attention, and context size, have been theorized to affect PPP, yet there is no current account that explains *why* such a tipping point exists, and how it interacts with LMs’ pretraining dynamics more generally. We hypothesize that the underlying factor is a pretraining *phase transition*, characterized by the rapid emergence of specialized attention heads. We conduct a series of correlational and causal experiments to show that such a phase transition is responsible for the tipping point in PPP. We then show that, rather than producing attention patterns that contribute to the degradation in PPP, phase transitions alter the subsequent learning dynamics of the model, such that further training keeps damaging PPP.

## 1 Introduction

The rise of neural-network-based language models (LMs) that can produce fluent, humanlike linguistic outputs has led to their increased use as cognitive models of human language processing (Warstadt and Bowman, 2024). In particular, an active area of research has studied how well LMs are aligned with human incremental processing behaviors. This is often measured by how well LMs’ output probabilities predict various human reading time metrics, and is referred to as a model’s *psychometric predictive power* or *PPP*. PPP has been widely studied for English (e.g., Wilcox et al., 2020; Shain et al., 2024; Oh et al., 2022; Oh and Schuler, 2023b, *inter alia*), Japanese (Kuribayashi et al., 2021), multilingually (Wilcox et al., 2023b), and for nonnative Englishes (Aoyama and Schneider, 2024), where it has been found that LMs’ outputs are robustly correlated with human reading times.

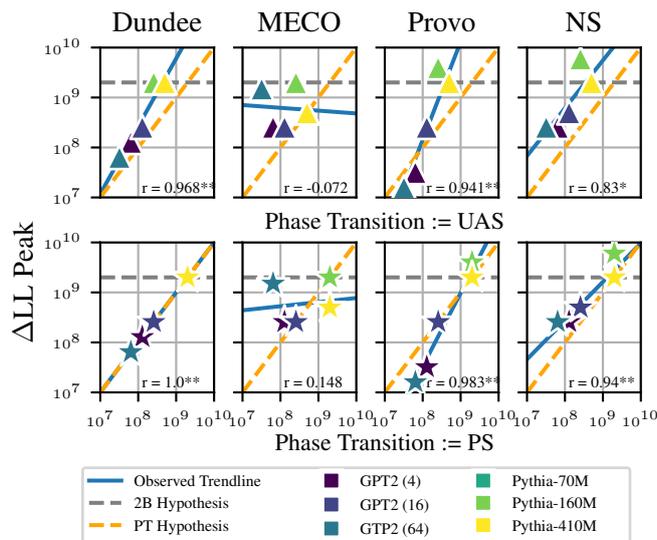


Figure 1: Phase transition ( $x$ -axis) and PPP peaks ( $y$ -axis) closely coincide with each other across the corpora studied in this work, except for MECO. Numbers in the parentheses are batch sizes; for Pythia models, regardless of the model size, the batch size is 1,024. **Orange** dashed lines, **blue** solid lines, and **gray** dashed lines represent the **phase transition hypothesis**, **2 billion tokens hypothesis**, and the **observed trendline** drawn from the data, respectively. We find that the experimental results closely align with our phase transition hypothesis, meaning that the PPP peaks correspond to LM phase transition, beyond which PPP starts degrading. \* $p < .05$ ; \*\* $p < .01$ .

One outstanding puzzle in these data is the observation that transformer-based LMs become maximally aligned with human sentence processing relatively early in pretraining, after which their fit to human data either plateaus or decreases. We will refer to the maxima of psychological fit during training as model’s **tipping point**. Such tipping points present a conundrum: Why does the fit to human data decrease, even as the model’s language modeling loss goes down? And what is the tipping point’s underlying cause?

Previous work has observed that tipping points

in models tend to occur after around 2 billion words of pretraining, and has suggested that this number—2 billion—is operative (Oh and Schuler, 2023a). We refer to this as the **2 Billion Hypothesis**. In contrast, we suspect that the tipping point at 2 billion tokens is coincidental rather than causal. Instead, we hypothesize that the reversal of PPP is related to more general pretraining dynamics, specifically the presence of **phase transitions**, or periods in pretraining when new model capabilities emerge rapidly. While phase transition can be a general term used to describe any rapid change in model behaviors, we focus on what we term **specialized heads phrase transitions**, characterized by the emergence of specialized attention patterns relatively early in pretraining (Olsson et al., 2022; Chen et al., 2024a). For brevity, we refer to these simply as “phase transition” (see Section 2.3 for a more precise definition). We dub this hypothesis the **Phase Transition Hypothesis** (see Hypothesis 1) and conduct a series of experiments, both correlational and causal, to test it (Figure 1).

Our key findings include: (1) the tipping points of PPP are strongly correlated with phase transition (Section 4); (2) a series of ablation experiments show some attention heads that formed during phase transition have some effect on PPP, but the overall results are mixed and the PPP degradation cannot be reliably attributed to specific circuits that form during the phase transition (negates Hypothesis 2a; see Section 5); and (3) regularization during pretraining delays the phase transition, also resulting in a delayed PPP tipping point, indicating that phase transition alters the learning dynamics, such that further training results in PPP degradation (supports Hypothesis 2b; see Section 6).<sup>1</sup>

## 2 Relevant Work

### 2.1 Sentence Processing

One popular theory to explain human sentence processing is *surprisal theory*. Surprisal theory posits that the processing difficulty of a word is proportional to its information content, quantified as its surprisal, or in-context, negative log probability (Levy, 2008; Hale, 2001). Surprisal theory is supported by numerous studies showing a tight linear relationship between incremental processing times and surprisal, across various datasets (Smith and Levy, 2013; Shain et al., 2024) and languages

<sup>1</sup>Code and data available at <https://github.com/picol-georgetown/pt-and-ppp>.

(Wilcox et al., 2023b).<sup>2</sup>

Rather than using LMs as a gold proxy for word predictability, other studies have compared how surprisals obtained from various LMs show different fits to human reading data. Goodkind and Bicknell (2018) find that LM quality, as measured in perplexity, linearly correlates with PPP, such that surprisals obtained from better LMs are more predictive of human reading time. This relationship between LM quality and PPP has been dubbed the *quality-power hypothesis* and replicated with different sets of LMs (Wilcox et al., 2020; Oh et al., 2022; Oh and Schuler, 2023b), and cross-lingually (Wilcox et al., 2023a).

Instead of comparing fully trained LMs, Oh and Schuler (2023a) find that the quality-power relationship changes during pretraining, which motivates the current study. Specifically, they find that PPP keeps improving until around 2 billion tokens of pretraining (for all Pythia variants) and then starts degrading beyond that point, although perplexity keeps improving. In other words, the quality-power correlation is *positive* (better the model, higher the PPP) until a certain point and *negative* (better the model, lower the PPP) after it. We call this point a **tipping point** in this paper, and define it as the maximum PPP obtained during pretraining. In fact, Kuribayashi et al. (2021) seems to be the first to report this trend, where they find similar tipping points for English and Japanese. This tipping point has been questioned in Aoyama and Schneider (2024), where they observe a *negative* quality-power correlation at orders of magnitude smaller pretraining amounts for crosslingual LMs.

### 2.2 Factors Affecting LMs’ PPP

Besides the amount of pretraining, several other factors have been found to affect LMs’ PPP.<sup>3</sup> One key factor is model size, where larger models tend to perform worse in modeling human sentence processing compared to smaller models. Oh et al. (2022) found that structural parsers and  $n$ -gram models often matched or outperformed GPT2 models, with the smallest GPT2 variant yielding the best results. Oh and Schuler (2023b) tested additional GPT and OPT variants confirming this trend. Another crucial factor is context size, as limiting the inference-time context window has been shown

<sup>2</sup>Although cf. Meister et al., 2021 who find a slightly *super-linear* relationship.

<sup>3</sup>For a more extensive list, we refer the readers to Table 1 of Kuribayashi et al., 2022

to improve PPP. Kuribayashi et al. (2022) observed that shorter contexts, particularly bigrams, resulted in better PPP, especially for Japanese. Additionally, introducing recency bias in attention mechanisms improves PPP (De Varda and Marelli, 2024), particularly when applied during both training and inference (Clark et al., 2024).

Lexical factors also play a role in PPP degradation among larger models. Oh and Schuler (2023b) found that larger models tend to assign lower surprisals to open-class words (e.g., nouns and adjectives) and named entities, deviating more from human reading patterns. Oh et al. (2024) further linked this effect to token frequency, showing that larger models exhibit particularly low alignment with human reading time for infrequent tokens, especially beyond a certain point in pretraining (2B tokens). These studies suggest that model size, long context utilization, and pretraining dynamics interact with each other to produce PPP degradation. We propose that one way to tie these pieces together is by finding a common cause, which we hypothesize to be phase transitions.

### 2.3 Phase Transition

What exactly is a phase transition, and how is it measured? Abrupt changes in model behaviors, which cannot be predicted by a scaling law (Kaplan et al., 2020),<sup>4</sup> have been studied widely (e.g., Wei et al., 2022; bench authors, 2023). In this paper, we specifically focus on what we call **specialized heads phase transition**, which is characterized by the rapid emergence of specialized attention patterns (Elhage et al., 2021; Olsson et al., 2022; Chen et al., 2024a), relatively early in pretraining. For brevity, we use the term “phase transition.”

We operationalize the phase transition using two well-studied phenomena, the presence of induction heads (Elhage et al., 2021; Olsson et al., 2022) and syntactic attention structure (SAS; Chen et al., 2024a). Elhage et al. (2021); Olsson et al. (2022) find that a specialized head called an induction head emerges at a certain point during pretraining, and that this head is characterized by its distinctive copying behavior, where the head attends to the previous occurrence of the same or similar bigram within the context. We introduce the metric used to detect induction heads in Equation (2). Olsson et al. (2022) further claim that the emergence of this induction head is primarily responsible for

<sup>4</sup>See Caballero et al., 2023 for a scaling law with ‘breaks’.

most of the in-context learning (ICL) abilities (see Equation (8)).

Similarly, Chen et al. (2024a) find that LMs demonstrate a sudden emergence of attention patterns that mirror syntactic dependency edges, which they call SAS. Similar to induction heads, the emergence of SAS is followed by an abrupt improvement in syntactic abilities as measured by performance on a syntactic probing benchmark, BLiMP (Warstadt et al., 2020).

### 2.4 Research Questions

In light of this, we ask the following research question: **to what extent is phase transition responsible for the degradation in the alignment between human and LM sentence processing?** We present three hypotheses, outlined below:

#### Hypothesis 1

**Phase Transition Hypothesis:** PPP degradation coincides with phase transition.

If this is the case, then we need to account for *how* and *why* phase transition causes the PPP degradation, for which we have two hypotheses:

#### Hypothesis 2a

**Persistence Hypothesis:** Circuits that emerge during phase transition cause PPP degradation.

This hypothesis predicts that certain circuits emerge at the phase transition, persist afterward, and contribute to the degradation in PPP. Alternatively, we may hypothesize that the degradation arises from pretraining dynamics *after* the phase transition. This view predicts that preventing PPP degradation requires suppressing the phase transition itself, rather than deactivating specific circuits, leading to our final hypothesis:

#### Hypothesis 2b

**Dynamics Hypothesis:** Altered training dynamics causes PPP degradation.

Crucially, Hypothesis 1 serves as a prerequisite for the others, while Hypothesis 2a and Hypothesis 2b may be either competing or complementary. Put differently, Hypothesis 2a attributes PPP degradation to circuits that emerge *during* the phase transition, whereas Hypothesis 2b attributes it to circuits that emerge *after* the phase transition.

### 3 Methods

#### 3.1 Models and Checkpoints

Because online sentence processing is a predominantly left-to-right process, we only consider autoregressive LMs, a common practice adopted by many others (see Meister et al., 2023 for a rare exception). Since training LLMs is expensive, we primarily focus on the Pythia family (Biderman et al., 2023). To our knowledge, Pythia is a rare, if not the only, LM family that satisfies the following conditions necessary in this paper: (1) open-sourced, (2) available in different sizes, and importantly, (3) various checkpoints, especially early in pretraining, are available. Bloom (BigScience Workshop, 2022) is another candidate; however, because the available checkpoints are likely after the phase transition, we did not include it in this study.

We also train several GPT2 models (Radford et al., 2019) and save log-spaced checkpoints at  $\{500K, 1M, 2M, \dots, 256M\}$  tokens for the first 10 checkpoints, and evenly-spaced checkpoints for the last 20 checkpoints at  $\{0.5B, 1B, 1.5B, \dots, 10B\}$  tokens. This is a superset of the available Pythia checkpoints for the first 10B tokens of pretraining, which allows for direct comparisons among different model sizes and families at the same pretraining amounts. All custom-trained GPT2 models have 2 layers, 8 heads (per layer), a hidden dimension of 768, and a total parameter count of 53M. See Table 2 in Appendix A for the list of architecture and training hyperparameters. Each run took  $\approx 70$  hours on a single A6000.

#### 3.2 Data

For custom trained models, we use a sample of 1B tokens from the English subcorpus of CC100 (Conneau et al. 2020; Wenzek et al. 2020; Apache License 2.0). For reading time data, we use 3 eye-tracking corpora: Dundee (Kennedy et al., 2003), Provo (Luke and Christianson, 2018; CC-BY 4.0),<sup>5</sup> and MECO (Siegelman et al., 2022) release 1.2.,<sup>6</sup> and 1 self-paced reading time corpus: Natural Stories (Futrell et al., 2021; CC BY-NC-SA). For eye-tracking corpora, we focus on the gaze duration.<sup>7</sup>

<sup>5</sup><https://osf.io/sjefs/>

<sup>6</sup><https://osf.io/3527a/>

<sup>7</sup>See Shain et al. (2024) for comparisons across different eye-tracking metrics, which yield similar results.)

#### 3.3 Calculating PPP

Following previous work (e.g. Goodkind and Bicknell, 2018; Wilcox et al., 2020, 2023a, *inter alia*), we operationalize PPP as the difference in log-likelihood (LL) between two linear models (delta log-likelihood;  $\Delta LL$ ):

$$\Delta LL = LL_{f_{\text{base+surp}}} - LL_{f_{\text{base}}} \quad (1)$$

where  $f_{\text{base+surp}}$  and  $f_{\text{base}}$  are linear models that predict human reading times using baseline features with and without LM surprisals, respectively. Finally, surprisal of a word  $w_i$  in context is its negative log probability:  $-\log P(w_i | \mathbf{w}_{<i})$ . See Appendix B for the full list of features included in the regression models.

### 4 Experiment 1: PPP Peaks at Phase Transition

#### 4.1 Method

As we want to show that PPP peaks at the phase transition, defined by the emergence of SAS and induction heads, we outline the definition of the metrics we use to characterize them below.

**Unlabeled Attachment Score (UAS).** UAS is a commonly used metric for dependency parsing, and in the context of SAS, conceptually, it measures how well a given LM’s attention pattern matches the dependency edges between words, ignoring the relation labels. In other words, UAS is the proportion of words, such that the highest attention weight lies between the word and its parent for the best-performing head for the given relation type. See Appendix C for details.

**Prefix-matching Score (PS).** Induction heads are defined by the copying behavior, such that if the model has seen an  $\langle A, B \rangle$  sequence and the current token is  $A$ , then an induction head is a head that promotes the prediction of  $B$  as the next token, completing the  $\langle A, B, \dots, A, B \rangle$  sequence. We quantify this behavior using PS (Olsson et al., 2022). The intuition behind this metric is to measure, when predicting what follows the second occurrence of  $A$ , how much a model is referencing what followed its first occurrence.

Given a random sequence of tokens  $\mathbf{x}$  repeated twice, PS of a head  $h$  at layer  $l$  is its average attention from the source token  $x_i$  to the next token of its previous occurrence:

$$\text{PS} = \frac{1}{|\mathbf{x}| - 1} \sum_{i=|\mathbf{x}|+1}^{2|\mathbf{x}|} \alpha^{(h,l)}(x_i, x_{i-(|\mathbf{x}|-1)}) \quad (2)$$

In our experiments, we use a random sequence of length 50, following the TransformerLens library (Nanda and Bloom, 2022), so  $|\mathbf{x}| = 50$ . Note that  $\mathbf{x}$  is a completely random sequence of token IDs; hence, attention patterns are not an artifact of data contamination or memorization.

**Breakthrough Points.** Chen et al. (2024a) defines *breakthrough* in some metric  $f$  as the acceleration point, which maximizes the growth in slope, measured in some discrete intervals  $\Delta$ . To make it less susceptible to surface fluctuations, we define them as the first checkpoint  $c$  at which a certain metric  $f$  exceeds a given threshold  $t$ :  $\min\{c \in \mathcal{C} | f(c) > t\}$ , where  $\mathcal{C}$  is the predefined set of checkpoints described in Section 3.1. In our experiments, we use  $t = 0.1$  for both UAS and PS. We validate these implementations of breakthroughs as well as the measurements of UAS and PS by showing that the sudden rises in UAS and PS are followed by the improvement in BLiMP and ICL scores, respectively (see Appendix E for these results).

## 4.2 Results

Figure 1 shows on the  $x$ -axis the number of pre-training tokens at which phase transition occurred, as defined by UAS (top) and by PS (bottom), and on the  $y$ -axis the number of pretraining tokens at which the peak  $\Delta\text{LL}$  was observed. The gray lines represent the **2 Billion Hypothesis** based on Oh and Schuler (2023a), where transformer models are predicted to reach the highest  $\Delta\text{LL}$  at around 2 billion pretraining tokens. The orange lines represent our **Phase Transition Hypothesis**, where the peak  $\Delta\text{LL}$  is predicted to happen at around the same time as the phase transition. The blue lines are the **observed trendlines**, and we can see that they closely match the Phase Transition Hypothesis. This trend holds across 3 orders of magnitude, meaning that the  $\Delta\text{LL}$  peaks occur at vastly different times across different models, yet they all co-occur with their respective phase transition points, although this trend was not found in MECO (see Section 9 for known issues of MECO). Note that we included Pythia models used in Oh and Schuler (2023a), and GPT2 models trained from scratch with various batch sizes, as we find that the phase transition point is a function of not only the number of pretraining tokens, but also the number of updates (hence changing batch size affects phase transition points). We leave a precise characterization of phase transition points to future work. For a complete visualization of the full  $\Delta\text{LL}$  trajectory,

see Figure 7 (Appendix G), which also includes the 0- and 1-layer models. These are predicted to exhibit qualitatively different behaviors, with the 0-layer model indeed showing virtually no PPP degradation.

The concurrence of PPP peaks and phase transitions seems robust; however, these observations are still correlational, and we need more evidence to make a causal claim. As such, we consider two potential hypotheses, outlined earlier: (1) phase transition develops some circuits that hurt PPP (Persistence Hypothesis), and/or (2) phase transition changes the course of pretraining dynamics, in a way such that further training hurts PPP (Dynamics Hypothesis). We test each hypothesis in the following experiments.

## 5 Experiment 2: PPP Degradation Cannot be Attributed to Specific Heads

If phase transition creates certain circuits responsible for the lower PPP (Hypothesis 2a; Persistence Hypothesis), ablating those circuits should improve PPP for models that have undergone the phase transition. In this experiment, we ablate each head one at a time and investigate if the ablation of specialized heads that form during phase transition (i.e. SAS and induction heads) improves PPP more than other heads.

### 5.1 Method

**Scores.** PS is by definition a head-specific score (i.e. how much a given head attends to the previous occurrence of token B when given a sequence  $\langle A, B, \dots, A \rangle$ ). On the other hand, SAS is measured at the *model* level; it is measured in UAS by picking attention scores of the best heads for each dependency relation type. Hence, we define a slightly modified *head-specific* metric called **SAS score**. SAS score of the head  $h$  at layer  $l$  is a proportion of the words  $w_i$  whose attention edge with the highest weight,  $(i, \arg \max_j [a_{ij}^{(h,l)}])$ , corresponds to a child-parent pair of a dependency relation.

**Ablation.** We zero-out attention weights of each head while keeping the original attention of all other heads and compute  $\Delta\text{LL}$  (see Appendix F for the details). We then subtract the non-ablated score from the ablated score and call this value  $\Delta\Delta\text{LL}$ . A *Higher*  $\Delta\Delta\text{LL}$  means that the ablation *improves* PPP. We report the correlations between (1) SAS score and  $\Delta\Delta\text{LL}$  and (2) PS and  $\Delta\Delta\text{LL}$ .

	GPT2 (1 layer)				GPT2 (2 layers)					Pythia-70M				Pythia-160M				Pythia-410M			
	DU	ME	PR	NS	DU	ME	PR	NS		DU	ME	PR	NS	DU	ME	PR	NS	DU	ME	PR	NS
$\Delta\Delta\text{LL} \sim \text{SAS score}$																					
16M	0.21	0.63	-0.11	0.46	0.24	<b>0.7</b>	0.09	0.49	.5B	-0.07	0.15	0.04	0.09	-0.14	-0.09	0.03	0.1	0.01	-0.09	-0.03	-0.1
32M	0.19	-0.03	-0.09	0.47	0.17	-0.02	<b>0.72</b>	0.02	1B	-0.23	-0.05	-0.15	-0.23	<b>-0.18</b>	0.05	-0.06	<b>-0.22</b>	-0.06	-0.04	-0.01	0.0
64M	0.35	0.56	0.55	0.45	<b>0.66</b>	<b>0.56</b>	0.32	0.25	2B	-0.2	0.17	-0.13	-0.23	<b>-0.25</b>	-0.0	0.0	-0.15	-0.04	0.07	-0.03	-0.08
128M	0.61	0.39	0.65	0.5	<b>0.53</b>	0.24	0.33	<b>0.62</b>	4B	-0.16	-0.25	0.16	-0.23	-0.13	-0.0	-0.02	<b>-0.23</b>	0.02	0.0	-0.02	-0.09
256M	0.52	0.62	0.38	0.49	0.46	-0.31	0.1	<b>0.5</b>	6B	-0.09	-0.01	0.17	-0.02	-0.05	-0.02	-0.0	-0.03	<b>0.14</b>	0.06	-0.03	0.06
$\Delta\Delta\text{LL} \sim \text{PS}$																					
16M	-0.1	-0.69	0.26	-0.43	0.41	-0.36	0.25	0.08	.5B	0.03	-0.12	-0.11	-0.05	0.07	-0.03	0.01	0.01	0.06	0.02	-0.01	<b>0.1</b>
32M	-0.04	0.14	0.0	-0.5	-0.11	0.06	<b>-0.51</b>	-0.11	1B	<b>0.29</b>	0.1	0.13	0.18	0.15	-0.0	0.01	0.06	0.01	0.05	-0.02	0.03
64M	-0.33	-0.61	-0.56	-0.43	<b>-0.64</b>	0.31	-0.05	<b>-0.77</b>	2B	<b>0.7</b>	0.13	0.16	-0.11	<b>0.61</b>	-0.0	0.09	0.15	<b>0.43</b>	0.02	0.03	0.05
128M	-0.53	-0.25	-0.53	-0.6	0.17	0.48	-0.23	-0.5	4B	<b>0.6</b>	0.13	0.01	0.26	<b>0.42</b>	<b>-0.23</b>	0.13	0.13	<b>0.21</b>	0.0	-0.02	-0.05
256M	-0.46	-0.53	<b>-0.74</b>	-0.63	0.21	0.29	-0.08	-0.39	6B	<b>0.41</b>	-0.09	-0.2	-0.06	<b>0.25</b>	-0.13	0.04	0.08	<b>0.17</b>	-0.02	0.07	<b>0.12</b>

Table 1: Pearson correlation coefficients between SAS score and  $\Delta\Delta\text{LL}$  (top) and PS (bottom) and  $\Delta\Delta\text{LL}$  for 2 and 3 variants of GPT2 and Pythia, respectively. DU, ME, PR, and NS stand for Dundee, MECO, Provo, and Natural Stories, respectively. Correlations statistically significant at  $\alpha = .05$  are boldfaced and colored in **green** if positive and **red** if negative. Note that these p-values are *before* any correction for multiple testing is applied.

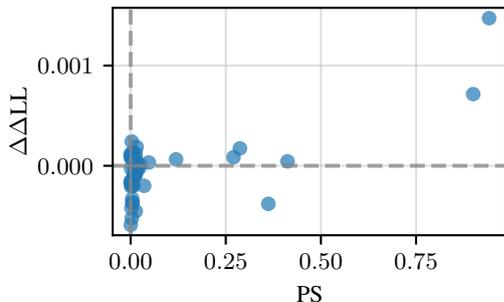


Figure 2: Relationship between each head’s PS ( $x$ -axis) and the effect of its ablation on  $\Delta\Delta\text{LL}$  ( $y$ -axis), measured on the Dundee corpus, for the Pythia-70M model at 2 billion pretraining tokens. Each point represents one of the 48 heads in the model.

## 5.2 Results

We start by considering a single, illustrative example: Figure 2 plots each head’s PS ( $x$ -axis) and  $\Delta\Delta\text{LL}$  for Dundee corpus ( $y$ -axis) for Pythia-70M at 2B tokens of training. These particular results appear to confirm our hypothesis: heads high in PS are also high in  $\Delta\Delta\text{LL}$ , meaning that the heads performing induction *damage* the PPP.

However, we also need to test if this is the case for other models, checkpoints, and data (other than Dundee). We focus on 5 checkpoints around (and including) the phase transition point (64M for GPT2 models and 2B for Pythia models). This produces 5 (checkpoints)  $\times$  5 (models)  $\times$  4 (corpora)  $\times$  2 (scores: SAS score and PS) = 200 scatterplots like Figure 2. For readability, we only report Pearson’s correlation coefficients for each, shown in

Table 1. For example, Figure 2 can be found under column Pythia-70M and DU, row 2B on the bottom half of the table ( $r = 0.7$ ,  $p < .05$ ).

In Table 1, a few trends emerge: First, we find strong correlations for models closer to the transition points (64M for GPT2 models and 2B for Pythia models). The very undertrained GPT2 models at 16M pretraining tokens, and Pythia models at 0.5B pretraining tokens, for example, show virtually no correlations between  $\Delta\Delta\text{LL}$  and SAS or PS. Second, whereas SAS $\sim\Delta\Delta\text{LL}$  correlation is positive but PS $\sim\Delta\Delta\text{LL}$  correlation is negative for GPT2 models, the opposite is true in general for Pythia models. We do not know why this should be the case, and it warrants further investigation. Given that more than half of the model checkpoints show no significant correlations, and that none of the significant correlations remain significant after Bonferroni correction for multiple testing is applied, we take these data as not supporting the claim that the post-phase transition degradation in PPP can be attributed to specific heads.

## 6 Experiment 3: Suppressing Phase Transition Delays PPP Peak

In this experiment, we test Hypothesis 2b (Dynamics Hypothesis): if phase transition changes the course of pretraining dynamics, and if the post-transition pretraining dynamics lead to the degradation in PPP, suppressing the phase transition should result in the mitigation of PPP degradation.

## 6.1 Methods

Completely suppressing phase transition is difficult, if not impossible, as we do not have a full mechanistic understanding of this phenomenon. However, as induction heads and SAS are among the few well-documented diagnoses of phase transition, in this section we suppress them as a proxy method for suppressing phase transition more generally.

Since SAS quantifies attention heads whose attention patterns shadow dependency edges, [Chen et al. \(2024a\)](#) propose using a syntactic regularizer:

$$\mathcal{L}_{\text{SAS}}(x) = \underbrace{\mathcal{L}_{\text{CLM}}(x)}_{\text{Original loss}} + \lambda \underbrace{\sum_{i=1}^{|x|} \sum_{x_j \in D(x_i)} \alpha(x_i, x_j)}_{\text{Syntactic regularization}} \quad (3)$$

where  $x$  is an input,  $D$  is a child-parent mapping of dependency relations, and  $\alpha$  is an attention weight between a pair of words.  $\lambda$  is a weighting factor, and [Chen et al. \(2024a\)](#) find that  $\lambda = 0.001$  works best with BERT (positive  $\lambda$  suppresses SAS, whereas negative  $\lambda$  promotes it). We show results from  $\lambda = \{0.01, 0.001\}$ . We use spaCy ([Honnibal et al., 2020](#)) to parse our training data.

To suppress the formation of induction heads, we regularize against attention patterns that correspond to the ‘‘copying’’ behavior:

$$\mathcal{L}_{\text{COPY}}(x) = \underbrace{\mathcal{L}_{\text{CLM}}(x)}_{\text{Original loss}} + \lambda \underbrace{\sum_{i=1}^{|x|} \sum_{x_j \in \text{PM}(x_i)} \alpha(x_i, x_j)}_{\text{Copying regularization}} \quad (4)$$

where  $\text{PM}(x_i)$  is a prefix matching tokens of  $x_i$ . Recall that the copying behavior of induction heads was characterized by predicting  $B$  when given a sequence  $\langle A, B, \dots, A \rangle$ . Since we do not know which tokens are considered ‘‘similar enough’’ to promote the formation of induction heads in natural texts (see [Chen et al., 2024b](#) for the discussion on ‘‘parallel structures’’ that are central to LMs’ ICL abilities), we construct synthetic data consisting of repeated random sequences of tokens (see Appendix H for the details of the dataset construction and copying regularization implementation).

In addition to the non-regularized model, copy-suppressed models, and SAS-suppressed models, we add models with Gaussian Noise Injection (GNI), as an example of a perturbed model with no explicit phase transition suppression. For more details, see Appendix L.

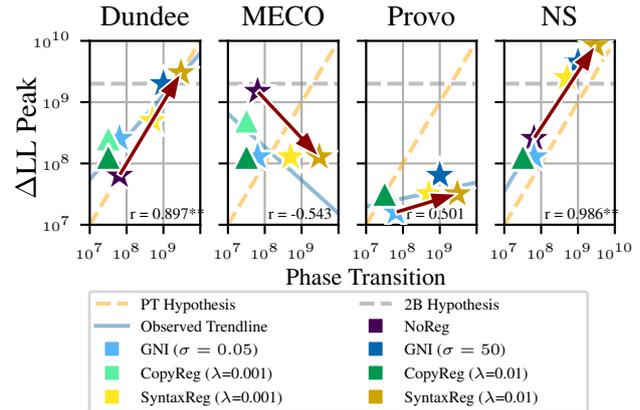


Figure 3: Phase transition points ( $x$ -axis) and  $\Delta\text{LL}$  tipping points ( $y$ -axis). Regularization delays both the phase transition and PPP tipping point, as exemplified by the SAS-regularized model ( $\lambda=0.01$ ), shown with the red arrow.  $\blacktriangle$  and  $\star$  represent phase transitions defined by UAS and PS, respectively.  $*p < .05$ ;  $**p < .01$ .

## 6.2 Results

### 6.2.1 Validation of Suppression

We first verify that SAS and copying suppression are working as intended by showing that SAS suppression leads to a lower UAS and consequently a lower BLiMP score, and that copying impression leads to a near-zero PS and a reduced ICL score (these results are presented fully in Appendix J).

Interestingly, however, none of the models resulted in a complete suppression of the phase transition. That is to say, with copy regularization, induction heads were suppressed (near-zero PS), but it promoted an even stronger SAS (higher UAS and BLiMP). On the other hand, with syntactic regularization, SAS was suppressed (near-zero UAS), but it delayed and weakened the induction heads (see Appendix K for more details on the interaction between the two types of regularization). All of the regularized models resulted in a higher loss compared to the non-regularized model, suggesting that emergent structures are indeed facilitative of next token prediction.

### 6.3 Effect of Suppression on PPP

Given that the phase transition was at most delayed (and not fully suppressed), and that loss is known to affect PPP, we present two sets of results: (1) the correlation between the phase transition points and PPP tipping points, and (2) the loss-PPP correlations before and after the phase transition. For (1), if the regularization resulted in a delayed phase

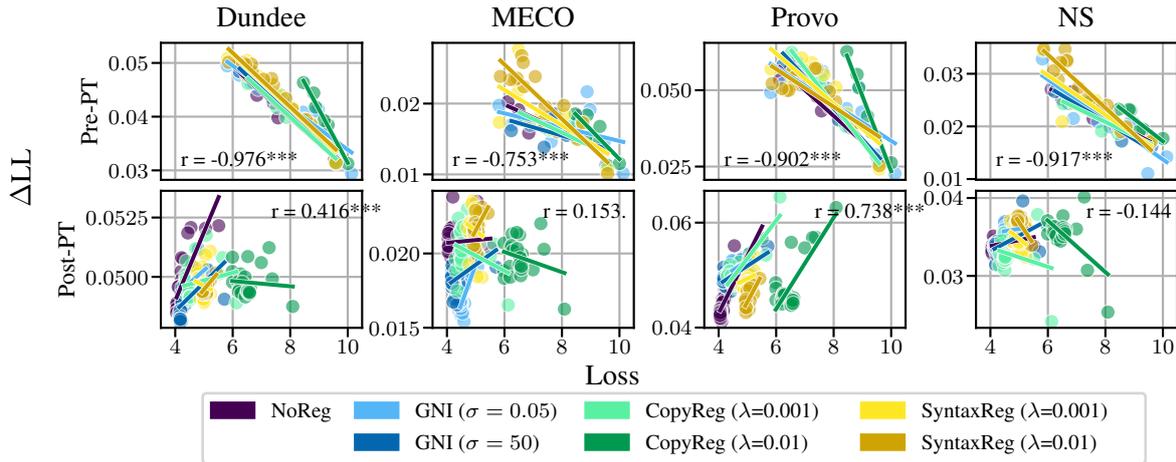


Figure 4: Validation loss ( $x$ -axis) and  $\Delta LL$  ( $y$ -axis) across checkpoints before (top) and after (bottom) the phase transition. Each point corresponds to a single checkpoint. The reported correlation is a Fisher  $z$ -transformed weighted average across models within each corpus. Correlation is consistently negative before the transition, while it becomes predominantly positive or mixed after.  $.p < .1$ ;  $*p < .05$ ;  $**p < .01$ ;  $***p < .001$ .

transition, this should also result in a delayed PPP tipping point. Since PS is a better predictor of PPP tipping points (see Figure 1), we use PS to compute phase transition points (Section 4.1). For models with copying regularization, UAS was used to obtain the phase transition points. For (2), we expect that the correlation is negative (i.e. quality-power hypothesis) before the phase transition, and positive after the phase transition. For the full trajectories of each model’s PPP, see Appendix M.

Figure 3 summarizes the effect of regularization on phase transition points and PPP peaks, with the **red arrow** ( $\rightarrow$ ) showing the representative case, pointing from the **non-regularized model** to **SAS-suppressed model** ( $\lambda = 0.01$ ). MECO is again an outlier here; however, a strong correlation is found in both Dundee and NS, and also in Provo, to an extent. More importantly, as the arrow suggests, suppression that leads to a delayed phase transition also delays the PPP peak, moving the point along the **orange line** that the Phase Transition Hypothesis predicts, and this further corroborates our Hypothesis 1. For the magnitude and significance of these trends, see Table 4 in Appendix M.

We have established that the suppression delays both phase transition and PPP tipping point. Figure 4 shows the loss-PPP correlation both before (top) and after (bottom) the phase transition. For Dundee and Provo, the contrast is evident: LM losses and PPPs have a negative correlation before the phase transition, and a positive correlations we see

especially in the Dundee and Provo corpora suggest that the decrease in loss that happens after the phase transition damages PPP, congruent with our Dynamics Hypothesis (Hypothesis 2b).

## 7 Discussion

This work proposed phase transition as the underlying factor that causes PPP degradation beyond a certain point during LM pretraining and conducted several studies to test this hypothesis. Here, we discuss how our results relate to previous factors that have been shown to impact PPP in LMs.

### 7.1 Phase Transition & Factors Affecting PPP

First, our Phase Transition Hypothesis can potentially explain previous findings that limiting the context window of a transformer LM can improve PPP (Kuribayashi et al., 2022), especially for infrequent tokens (Oh et al., 2024), and that adding a linear recency bias improves PPP (De Varda and Marelli, 2024; Clark et al., 2024). To explain why, recall that Olsson et al. (2022) find that induction heads improve the ICL score, as measured by how well an LM leverages earlier tokens in the context for the next word prediction, even across long spans of text. If the emergence of ICL abilities is the underlying cause of the PPP reversal, then it follows that limiting ICL would help PPP. In this study, we have achieved this through ablation and regularization; however, restricting LM’s context size would yield a similar outcome by completely disabling models to leverage earlier tokens.

## 7.2 Phase Transition & “Bigger $\neq$ Better”

Second, the Phase Transition Hypothesis can explain why larger models suffer from lower PPP than smaller models (e.g., Oh et al., 2022; Oh and Schuler, 2023b), an example of inverse scaling. Oh and Schuler (2023b) conjecture that the overly accurate prediction of infrequent tokens is one of the contributing factors to models’ PPP reversal, and suggest that it could be due to memorization. Chen et al. (2024a) suggest that the memorization phase occurs after the phase transition. Taken together, given that most of the models studied in previous works have almost certainly undergone the phase transition, it is unsurprising that memorization is damaging PPP. In fact, Oh et al. (2024) make an interesting observation regarding pretraining dynamics and predictions on infrequent tokens: models of different sizes seem to follow similar learning trajectories up to a certain point during pretraining (2B tokens), but the ability to predict infrequent tokens with low surprisals (which is considered one of the reasons for the larger models’ poor predictive power) seems to emerge beyond that point only among larger models.

## 7.3 SAS & Induction Heads

To our knowledge, this study is the first to investigate SAS in decoder-only models, and to observe the emergence of induction heads and SAS in the same model. This affords us to make observations about the interactions between the two, as briefly discussed in Section 6.2.

In all of the models we examined, SAS emerges first, closely followed by the formation of induction heads (see Figure 7 in Appendix G). Although the universality of this order remains inconclusive based on the selection of models we examined, this begs a question: is the emergence of SAS a prerequisite for the emergence of induction heads?

If this was the case, we would expect the SAS suppression to also suppress induction head formation, but not vice versa. Interestingly, we see a similar trend, but not quite what we expected: suppression of SAS delays the induction head formation and lowers the eventual PS obtained at the end of the training; however, the suppression of induction heads does not affect the point at which SAS emerges, and further, it *increases* the UAS. The interaction among different types of specialized heads remains an interesting topic for further investigation (see Appendix K for details).

## 7.4 Beyond Reading Time

Concurrent works explore the relationship between training dynamics and the alignment between LMs and brain activities. AlKhamissi et al. (2025) measured Pythia models’ formal and functional linguistic competences as well as their alignment with brain data throughout pretraining and reported a strong correlation between formal linguistic competence and the brain-LM alignment. Furthermore, they show that models of various sizes (ranging from 70M–6.9B) have similar levels of alignment with brain activities, and that they all tend to peak in the alignment between 2B–8B tokens of pretraining, echoing the observation we have made. Note that the tipping point of brain-LM alignment seems to come after the tipping point of sentence processing alignment.

Nakagi et al. (2025) find a similar abrupt increase in the brain-LM alignment at 1B pretraining tokens, followed by a sudden drop at 10B pretraining tokens for a set of OLMo variants. They find a subsequent resurgence in the alignment score at around 50B pretraining tokens, describing the overall dynamics as consisting of 3 phases. Although we only trained our models up to 10B tokens, results from Oh and Schuler (2023a) show that Pythia models go through a tipping point at around 2B tokens, after which no further increase is observed throughout the 300B pretraining tokens. Hence, this resurgence after the tipping point might be unique to brain-LM alignment; however, more work is needed to verify this speculation.

## 8 Conclusion

This study found a strong correlation between the phase transitions and peaks in a model’s fit to human reading times. However, the precise aspect of the phase transition responsible for PPP degradation remains somewhat inconclusive. Given that the results largely supported Hypothesis 2b (Dynamics Hypothesis, Section 6) rather than Hypothesis 2a (Persistence Hypothesis; Section 5), we can be reasonably confident that it is the post-phase transition dynamics driving the PPP degradation. However, it is important to note that none of our regularization methods was able to completely suppress the phase transition. Fully controlling the phase transition and investigating its effect on PPP requires a more complete understanding of the phase transition, and this remains an important avenue for future research.

## 9 Limitations

First, our models are limited to Pythia models and GPT2 models. This is due to the limited availability of pretrained models’ checkpoints, as well as the high computational cost associated with training large models from scratch. As mentioned in Section 3.1, to our knowledge, Pythia is the only model family whose available checkpoints cover the phase transition points.

We also limited the GPT2 models to 0-, 1-, and 2-layer variants. This is because these three variants are theoretically shown to behave qualitatively differently (Olsson et al., 2022) as we discussed in Section 4.2, and that models with 2 or more layers are not qualitatively different for our purpose of attention-based phase transition detection. Pythia model’s 70M, 160M, and 410M variants are 6, 12, and 24 layers, respectively, and we believe that this covers a reasonable range of model sizes, together with the GPT2 models we train. We also note that experiments on smaller “toy” models are important first steps (e.g., Elhage et al., 2021; Olsson et al., 2022; Power et al., 2022; Kallini et al., 2024). Yet, training larger models in Section 6 was beyond our compute budget, and we leave this to future work.

Second, our selection of reading time corpora is representative rather than comprehensive, and the replication with other corpora, as well as different reading behavior metrics such as brain activity data, remains an important future work. We also observe that the results for MECO were different from results from the other 3 corpora. The absence of the PPP peak  $\sim$  phase transition correlation (Figure 1 in Section 1 and Figure 3 in Section 6.2) are considered an exception rather than a rule. Several issues that could affect the quality of the corpus have been reported on MECO. For example, Opedal et al. (2024) find an off-by-one issue for a handful of tokens in MECO, as well as repeated words in a few sentences. We followed the fixed version of the data; however, there may be other issues we are unaware of, potentially causing the divergent behavior of the corpus.

Lastly, our study is limited to English, and our results may not hold for other languages. However, given that syntactic dependencies and word or phrase repetitions are universal across the world’s languages, we predict that a similar trend might be observed in other languages. A multilingual extension of this study is therefore a promising direction for future work.

## Ethics Statement

We trained several small transformer-based LMs from scratch, which could contribute to the increased carbon footprint. However, we train models that have at most 2 layers and 8 heads. By choosing a model size that reasonably approximates the popular transformer architecture (at least for the purpose of our study) while curtailing the computational cost, we believe that we were able to minimize our environmental impact.

Lastly, while our study utilizes four sources of human behavioral data, we do not intend to redistribute or publicly share these datasets. We affirm that our use of this data aligns with ethical standards and does not pose any potential ethical concerns.

## Acknowledgments

We thank the four anonymous reviewers for their active engagement in the discussion, which made the final version of this paper substantially better.

## References

- Badr AlKhamissi, Greta Tuckute, Yingtian Tang, Taha Binhurair, Antoine Bosselut, and Martin Schrimpf. 2025. [From language to cognition: How llms outgrow the human language network](#). *arXiv preprint arXiv:2503.01830*.
- Tatsuya Aoyama and Nathan Schneider. 2024. [Modeling nonnative sentence processing with L2 language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4927–4940, Miami, Florida, USA. Association for Computational Linguistics.
- BIG bench authors. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.
- Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. 2023. [Pythia: A suite for analyzing large language models across training and scaling](#). *arXiv preprint arXiv:2304.01373*.
- BigScience Workshop. 2022. [BLOOM \(revision 4ab0472\)](#).
- Ethan Caballero, Kshitij Gupta, Irina Rish, and David Krueger. 2023. [Broken neural scaling laws](#). In *The Eleventh International Conference on Learning Representations*.

- Angelica Chen, Ravid Shwartz-Ziv, Kyunghyun Cho, Matthew L Leavitt, and Naomi Saphra. 2024a. **Sudden drops in the loss: Syntax acquisition, phase transitions, and simplicity bias in MLMs.** In *The Twelfth International Conference on Learning Representations*.
- Yanda Chen, Chen Zhao, Zhou Yu, Kathleen McKeown, and He He. 2024b. **Parallel structures in pre-training data yield in-context learning.** In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8582–8592, Bangkok, Thailand. Association for Computational Linguistics.
- Christian Clark, Byung-Doh Oh, and William Schuler. 2024. **Linear recency bias during training improves transformers’ fit to reading times.** *arXiv preprint arXiv:2409.11250*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. **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*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. **Unsupervised cross-lingual representation learning at scale.** In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Andrea De Varda and Marco Marelli. 2024. **Locally biased transformers better align with human reading times.** In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 30–36, Bangkok, Thailand. Association for Computational Linguistics.
- 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. 2021. **A mathematical framework for transformer circuits.** *Transformer Circuits Thread*.
- Richard Futrell, Edward Gibson, Harry J Tily, Idan Blank, Anastasia Vishnevetsky, Steven T Piantadosi, and Evelina Fedorenko. 2021. **The natural stories corpus: a reading-time corpus of english texts containing rare syntactic constructions.** *Language Resources and Evaluation*, 55:63–77.
- Adam Goodkind and Klinton Bicknell. 2018. **Predictive power of word surprisal for reading times is a linear function of language model quality.** In *Proceedings of the 8th Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2018)*, pages 10–18, Salt Lake City, Utah. Association for Computational Linguistics.
- John Hale. 2001. **A probabilistic Earley parser as a psycholinguistic model.** In *Second Meeting of the North American Chapter of the Association for Computational Linguistics*.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. **spaCy: Industrial-strength natural language processing in python.**
- Julie Kallini, Isabel Papadimitriou, Richard Futrell, Kyle Mahowald, and Christopher Potts. 2024. **Mis-sion: Impossible language models.** In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14691–14714, Bangkok, Thailand. Association for Computational Linguistics.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. **Scaling laws for neural language models.** *arXiv preprint arXiv:2001.08361*.
- Alan Kennedy, Robin Hill, and Joël Pynte. 2003. **The dundee corpus.** In *Proceedings of the 12th European conference on eye movement*.
- Tatsuki Kuribayashi, Yohei Oseki, Ana Brassard, and Kentaro Inui. 2022. **Context limitations make neural language models more human-like.** In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10421–10436, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tatsuki Kuribayashi, Yohei Oseki, Takumi Ito, Ryo Yoshida, Masayuki Asahara, and Kentaro Inui. 2021. **Lower perplexity is not always human-like.** In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5203–5217, Online. Association for Computational Linguistics.
- Roger Levy. 2008. **Expectation-based syntactic comprehension.** *Cognition*, 106(3):1126–1177.
- Steven G Luke and Kiel Christianson. 2018. **The provo corpus: A large eye-tracking corpus with predictability norms.** *Behavior research methods*, 50:826–833.
- Clara Meister, Tiago Pimentel, Patrick Haller, Lena Jäger, Ryan Cotterell, and Roger Levy. 2021. **Revisiting the Uniform Information Density hypothesis.** In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 963–980, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Clara Meister, Tiago Pimentel, Gian Wiher, and Ryan Cotterell. 2023. [Locally typical sampling](#). *Transactions of the Association for Computational Linguistics*, 11:102–121.
- Yuko Nakagi, Keigo Tada, Sota Yoshino, Shinji Nishimoto, and Yu Takagi. 2025. [Triple phase transitions: Understanding the learning dynamics of large language models from a neuroscience perspective](#). *arXiv preprint arXiv:2502.20779*.
- Neel Nanda and Joseph Bloom. 2022. [Transformerlens](https://github.com/TransformerLensOrg/TransformerLens). <https://github.com/TransformerLensOrg/TransformerLens>.
- Byung-Doh Oh, Christian Clark, and William Schuler. 2022. [Comparison of structural parsers and neural language models as surprisal estimators](#). *Frontiers in Artificial Intelligence*, 5.
- Byung-Doh Oh and William Schuler. 2023a. [Transformer-based language model surprisal predicts human reading times best with about two billion training tokens](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1915–1921, Singapore. Association for Computational Linguistics.
- Byung-Doh Oh and William Schuler. 2023b. [Why does surprisal from larger transformer-based language models provide a poorer fit to human reading times?](#) *Transactions of the Association for Computational Linguistics*, 11:336–350.
- Byung-Doh Oh, Shisen Yue, and William Schuler. 2024. [Frequency explains the inverse correlation of large language models’ size, training data amount, and surprisal’s fit to reading times](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2644–2663, St. Julian’s, Malta. Association for Computational Linguistics.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2022. [In-context learning and induction heads](#). *Transformer Circuits Thread*.
- Andreas Opedal, Eleanor Chodroff, Ryan Cotterell, and Ethan Wilcox. 2024. [On the role of context in reading time prediction](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3042–3058, Miami, Florida, USA. Association for Computational Linguistics.
- Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. 2022. [Grokking: Generalization beyond overfitting on small algorithmic datasets](#). *arXiv preprint arXiv:2201.02177*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- Cory Shain, Clara Meister, Tiago Pimentel, Ryan Cotterell, and Roger Levy. 2024. [Large-scale evidence for logarithmic effects of word predictability on reading time](#). *Proceedings of the National Academy of Sciences*, 121(10):e2307876121.
- Noam Siegelman, Sascha Schroeder, Cengiz Acartürk, Hee-Don Ahn, Svetlana Alexeeva, Simona Amenta, Raymond Bertram, Rolando Bonandrini, Marc Brysbaert, Daria Chernova, et al. 2022. [Expanding horizons of cross-linguistic research on reading: The multilingual eye-movement corpus \(meco\)](#). *Behavior research methods*, 54(6):2843–2863.
- Nathaniel J. Smith and Roger Levy. 2013. [The effect of word predictability on reading time is logarithmic](#). *Cognition*, 128(3):302–319.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Alex Warstadt and Samuel R. Bowman. 2024. [What artificial neural networks can tell us about human language acquisition](#). *arXiv preprint arXiv:2208.07998*.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. [Emergent abilities of large language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. 2020. [CCNet: Extracting high quality monolingual datasets from web crawl data](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4003–4012, Marseille, France. European Language Resources Association.
- Ethan Wilcox, Clara Meister, Ryan Cotterell, and Tiago Pimentel. 2023a. [Language model quality correlates with psychometric predictive power in multiple languages](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7503–7511, Singapore. Association for Computational Linguistics.

Ethan G. Wilcox, Jon Gauthier, Jennifer Hu, Peng Qian, and Roger P. Levy. 2020. [On the predictive power of neural language models for human real-time comprehension behavior](#). *Proceedings of the Annual Meeting of the Cognitive Science Society*, 42:1707–1713.

Ethan G. Wilcox, Tiago Pimentel, Clara Meister, Ryan Cotterell, and Roger P. Levy. 2023b. [Testing the predictions of surprisal theory in 11 languages](#). *Transactions of the Association for Computational Linguistics*, 11:1451–1470.

## A Hyperparameters

Architecture	
architecture	gpt2
vocab_size	50_257
context_size	1_024
d_embed	768
d_ffn	3_072
n_layer	2
n_head	8
activation	gelu
num_params	53M
Training	
train_size	1B
num_epoch	10
train_amount	10B
batch_size	4
grad_acc_steps	1
weight_decay	0.1
warmup_steps	1%
lr	5e-4
lr_scheduler	cosine

Table 2: List of hyperparameters used to train L2LMs.

Table 2 summarizes the hyperparameters used to train 2-layer GPT2 models. 0-layer and 1-layer GPT2 models had 39M and 46M total parameters, respectively, and the rest of the hyperparameters remain the same.

## B Regression Model Formulae for $\Delta LL$ Calculation

Following [Smith and Levy \(2013\)](#), we model the spillover effect of previous 2 and 4 words for modeling eye-tracking data (Dundee, MECO, Provo) and self-paced reading data (Natural Stories), respectively. Frequencies are estimated using Wiki-text.

### Baseline regression model for eye-tracking.

$$\text{psychometric} \sim \text{freq} + \text{prev\_freq} + \text{prev2\_freq} + \text{len} + \text{prev\_len} + \text{prev2\_len}$$

### Baseline regression model for reading time.

$$\text{psychometric} \sim \text{freq} + \text{prev\_freq} + \text{prev2\_freq} + \text{prev3\_freq} + \text{prev4\_freq} + \text{len} + \text{prev\_len} + \text{prev2\_len} + \text{prev3\_len} + \text{prev4\_len}$$

### Full regression model for eye-tracking.

$$\text{psychometric} \sim \text{surprisal} + \text{prev\_surp} + \text{prev2\_surp} + \text{freq} + \text{prev\_freq} + \text{prev2\_freq} + \text{len} + \text{prev\_len} + \text{prev2\_len}$$

### Full regression model for reading time.

$$\text{psychometric} \sim \text{surprisal} + \text{prev\_surp} + \text{prev2\_surp} + \text{prev3\_surp} + \text{prev4\_surp} + \text{freq} + \text{prev\_freq} + \text{prev2\_freq} + \text{prev3\_freq} + \text{prev4\_freq} + \text{len} + \text{prev\_len} + \text{prev2\_len} + \text{prev3\_len} + \text{prev4\_len}$$

## C Unlabeled Attachment Score (UAS)

We follow the calculation of UAS introduced in [Chen et al. \(2024a\)](#), which is based on [Clark et al. \(2019\)](#). While they both use bidirectional models (i.e. BERT), we use decoder-only autoregressive models (i.e. GPT2 and Pythia), and hence make modifications to account for this difference. The overall recipe is to (1) define a head-specific probe, (2) find the best performing head for each dependency relation type, and (3) calculate the overall UAS using the best heads defined in (2).

**(1) Head-specific probe.** A head-specific probe  $f_{h,l}$  predicts the parent word of a target word  $x_i$  by selecting the word  $x_j$ , whose attention edge to or from the target word  $x_i$  is the highest among all words  $x_j \in \{j \neq i\}$  for a given head  $h$  at layer  $l$ :

$$f_{h,l}(x_i) = \arg \max_{x_j} \left( a_{ij}^{(h,l)} \right). \quad (5)$$

Note that, unlike bidirectional models, where each pair of words is connected by 2 attention edges ( $x_i \leftrightarrow x_j$ ), only 1 edge lies between any pair of words for autoregressive models. This means that the number of words a given word  $x_i$  can attend to is  $i$ , whereas the average number of words  $x_i$ 's right context can attend to is  $\frac{i+1024}{2}$ , creating discrepancy in the scale. If  $i = 10$ , for example, with the context size of 1024,  $\forall_{j \leq i} [\mathbb{E}(a_{i \rightarrow j}^{(h,l)})] = 0.1$ , whereas  $\mathbb{E}_{j > i} [\mathbb{E}(a_{j \rightarrow i}^{(h,l)})] \approx 0.002$ . However, we find that scaling the attention weights between  $a_{i \rightarrow j}^{(h,l)}$  and  $a_{j \rightarrow i}^{(h,l)}$  produce similar results, and hence we report unscaled results throughout.

Following Clark et al. (2019); Chen et al. (2024a), we convert the token-level attention to word-level attention by summing over attention weights *to* destination tokens that make up a single word, and by averaging over attention weights *from* source tokens that constitute a single word.

**(2) Best head per relation type.** The rest of the UAS calculation remains same as Chen et al. (2024a). We now convert the head-specific probe defined in (1) to a relation-specific probe by finding the best head for each dependency relation type.

For each dependency relation type  $R$ , which we define as a set of all ordered child-parent pairs  $(x, y)$ , the best performing head for the given dependency relation type is:

$$\hat{f}_R = \arg \max_{f_{h,l}} \frac{1}{|R|} \sum_{(x,y) \in R} \mathbb{1}_R(x, f_{h,l}(x)), \quad (6)$$

where  $x$  and  $y$  are constrained to be within the same sentence. The indicator function for set  $R$ ,  $\mathbb{1}_R$  is 1 if the predicted child-parent pair is in the set  $R$ , and 0 otherwise. Hence,  $\hat{f}_R$  is simply a head that has the highest recall for a given dependency relation type  $R$ .

**(3) UAS.** Lastly, we simply take the average of the performance of each relation type’s best head over all relation types, weighted by the number of ordered word pairs in that relation type. Denoting the set of all relation types as  $\mathcal{R}$ , UAS is defined as:

$$\text{UAS} = \frac{1}{\sum_{R \in \mathcal{R}} |R|} \sum_{R \in \mathcal{R}} \sum_{(x_i, x_j) \in R} \mathbb{1}_R(x_i, \hat{f}_R(x_j)). \quad (7)$$

## D In-Context Learning (ICL) Score

Olsson et al. (2022) define what they call **ICL score**, or the difference in the losses of tokens later in the context and tokens earlier in the context. This score seems to be robust to the choice of exactly which tokens to compare; however, Olsson et al. (2022) report the difference between 50th and 500th tokens’ losses. We instead report the difference between the average loss of early tokens  $w_j$ , where  $j \in [40, 60]$ , and that of later tokens  $w_k$ , where  $k \in [450, 550]$ , for each sequence  $s_i$ , where  $i \in [1, N]$ :

$$\frac{\sum_i \sum_j \sum_k \mathcal{L}(f_\theta(s_i, w_j)) - \mathcal{L}(f_\theta(s_i, w_k))}{200 \times N} \quad (8)$$

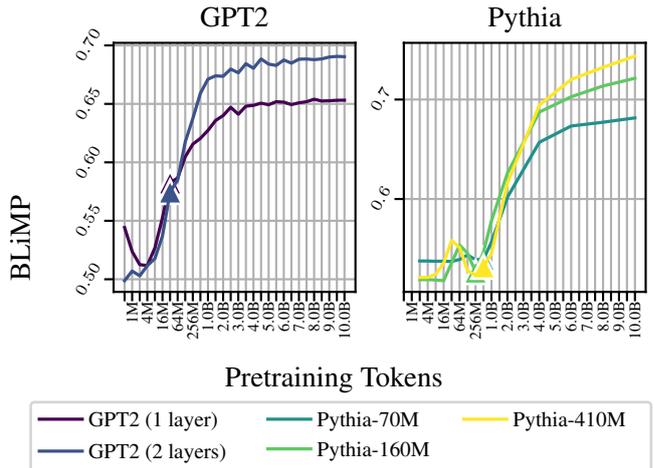


Figure 5: Trajectories of BLiMP score over 10B tokens of pretraining.  $\blacktriangle$  represent phase transition defined by a sudden rise UAS.

where  $f_\theta(s_i, w_j)$  is the output of an LM parametrized by  $\theta$ , given  $j$ -th word  $w_j$  in  $i$ -th sequence  $s_i$ . A positive ICL score means that the model has a lower loss (i.e. better prediction) later in the context than earlier in the context.

## E Phase Transition and its Downstream Effects

Chen et al. (2024a) find that the increase in UAS triggers the acquisition of syntactic abilities, as shown in an increase in the BLiMP score closely following the UAS boost, and we find a similar pattern. In Figure 5, GPT2 seems to go through a drastic increase in the BLiMP score between 4M to 1B pretraining tokens, with a brief halt around 32M-64M pretraining tokens. This is when the phase transition occurs, and it seems to signal the onset of the second boost in BLiMP score, starting around 64M tokens. For Pythia, the picture seems even clearer: UAS phase transition happens at around 512M pretraining tokens, immediately followed by a sudden increase in the BLiMP score.

Similarly, Olsson et al. (2022) reports a dramatic improvement in ICL score, foreshadowed by the emergence of induction heads, measured by PS. In Figure 6, ICL scores of 2-layer GPT2 and 1-layer GPT2 models seem to start diverging between 32M-64M pretraining tokens, which is exactly when the 2-layer model is undergoing the phase transition. Recall that induction requires attention composition, which is only possible with models with 2 or more layers. For Pythia models, the increase in ICL score is most dramatic between 1B and 2B

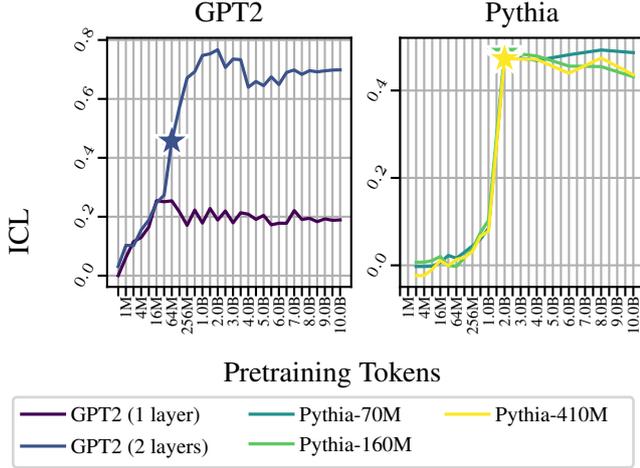


Figure 6: Trajectories of ICL score over 10B tokens of pretraining. ★ represent phase transition defined by a sudden rise in PS.

pretraining tokens, which also coincide with the emergence of induction heads.

## F Ablation

Ablation of attention heads can be implemented in two ways: full ablation and pattern-preserving ablation. For full ablation, the attention output is simply set to 0. Recall that attention output is a matrix multiplication between the attention weight vector, which is computed using query and key vectors, and the value vector (Vaswani et al., 2017):

$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (9)$$

where  $Q$ ,  $K$ , and  $V$  are query, key, and value vectors, respectively, and  $d_k$  represents the dimension of  $Q$  and  $K$  vectors. Because  $Q$ ,  $K$ , and  $V$  vectors are linear projections of the previous layer’s output, setting Equation (9) to a 0 vector affects all downstream (later layers’) calculations of  $Q$ ,  $K$ , and  $V$  vectors. In pattern-preserving ablation, we feed the input to the model twice: during the first run, no heads are ablated, and we simply record all the attention weights at each head of each layer. During the second run, we ablate the head(s) of interest, but use the attention weights recorded in the first run, preserving the original attention weights (hence *pattern-preserving*). Because only  $Q$  and  $K$  vectors are involved in the attention weight calculation  $\text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)$ , pattern-preserving ablation only affects downstream  $V$  calculations, but not  $QK$  calculations. We note that these two implementations of ablation yield similar results, and we

report the results obtained from pattern-preserving ablation.

## G Full PPP Trajectory and Phase Transition Points

Figure 7 shows a fuller visualization of the entire trajectory of  $\Delta\text{LL}$ , with the phase transition points marked for each model (▲ and ★ for UAS and PS, respectively, which correspond to the x-axis of Figure 1). Note that, in Figure 7, we included 0 and 1 layer GPT2 models as they are expected to behave differently from 2 layer models. This is because 0 layer model does not have attention layers, and 1 layer model cannot form induction heads, as they require *attention composition* (Olsson et al., 2022). We limited the 2 layer models to the one trained with a batch size of 4 for readability.

First, GPT2 with no attention layer (dark purple line) exhibits no “tipping point” in PPP, meaning that it never seems to undergo a point at which PPP starts going down. This is expected as the definitions of phase transition we adopt in this paper are both triggered by specialized behaviors of attention heads, and this further corroborates our hypothesis. Conversely, for GPT2 models with 1 and 2 layers, as well as 3 different sizes of Pythia models all have a tipping point at which PPP starts going down. Note that the results for Pythia models are partial replications of Oh and Schuler (2023a).

Second, as shown in Figure 1 as well, the tipping points of PPP seem to closely match the breakthrough points in ICL scores. Note that these breakthrough points are meaningful for models with 2 or more layers because induction heads work with another head in lower layers to perform the copying behavior (cf. attention composition; Olsson et al., 2022); hence, we do not indicate those breakthrough points for 0 and 1 layer models.

## H Implementation of Copying Regularization

Since the distance between two  $\langle A, B \rangle$  sequences,  $l$ , can be arbitrarily large within the LM’s context size, we sample  $l$  from a uniform distribution  $U(50, 512)$ . The lower bound is to make sure non-copying behaviors do not get suppressed, and the upper bound is just half the max context length of our model. To maximize the suppression target, we repeat a random sequence of length  $l$  as many times as it takes to fill the context size of 1024. Since it has to be repeated at least twice to promote the

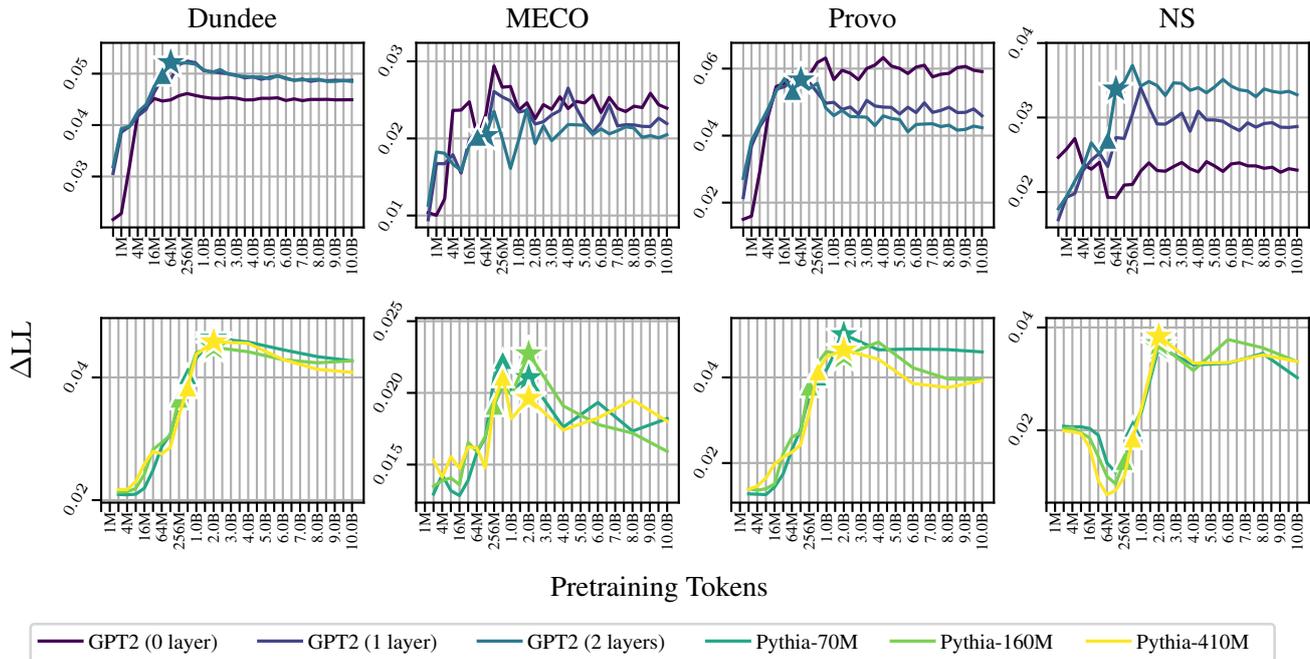


Figure 7: Trajectories of  $\Delta LL$  over 10B tokens of pretraining.  $\blacktriangle$  and  $\star$  represent phase transition defined by a sudden rise in UAS and PS, respectively.

copying behavior,  $l$  is upper-bounded by 512. With this synthetic data,  $PM(x_i)$  can be written as:

$$\{x_j | j < i, x_j = x_{i+1}, \mathbf{x}_{i-l+1}^i = \mathbf{x}_{j-l}^{j-1}\} \quad (10)$$

or programmatically:

$$\{x_{i-nl+1} | n \in \mathbb{N}_1, x - nl > 0\} \quad (11)$$

Now, as opposed to SAS suppression, where the regularization term is computed on the same set of examples from which normal loss  $\mathcal{L}_{CLM}$  is computed, the copying regularization term requires separate synthetic data. We considered alternating between the  $\mathcal{L}_{CLM}$  and  $\mathcal{L}_{COPY}$  every  $\frac{1}{\lambda}$  steps; however, we find that this leads to a very unstable learning curve, and hence add the regularization term every step weighted by  $\lambda$  as was the case with syntactic regularization.

## I Permutation Test

To test the significance of the differences in  $\Delta LL$  between non-regularized and regularized models, we conduct a series of pair-wise permutation tests. Since  $\Delta LL$  is computed for each word per model, and we are not interested in the variance in  $\Delta LL$  introduced by word ID, it was treated as a random effect. More concretely, for each pair of models we are comparing, we first generate 10,000 random permutations (i.e. random assignment of group)

Model		Corpora			
Reg	$\lambda/\sigma$	Dundee	MECO	Provo	NS
Copy	0.001	.334	.539	.000	.986
	0.01	.076	.269	.003	.271
Syntax	0.001	.305	.808	.001	.194
	0.01	.121	.947	.319	.023
GNI	0.05	.461	.175	.264	.539
	50	.173	.299	.301	.016

Table 3: P-values obtained from the permutation test.

between the 2 labels (2 models being compared) for *each word ID*, so that each of the per-word  $\Delta LL$  is randomly assigned to one of the two labels. We then count the proportion of permutations that have a mean within-word ID difference equal to or larger than the observed difference. Since we are interested in the difference between non-regularized and regularized models, we conduct six pair-wise permutation tests between the non-regularized model and each of the six regularized variants (3 regularizers  $\times$  2 hyperparameters = 6). Table 3 summarizes the permutation test results. Note that the p-values are all *before* any correction for multiple testing.

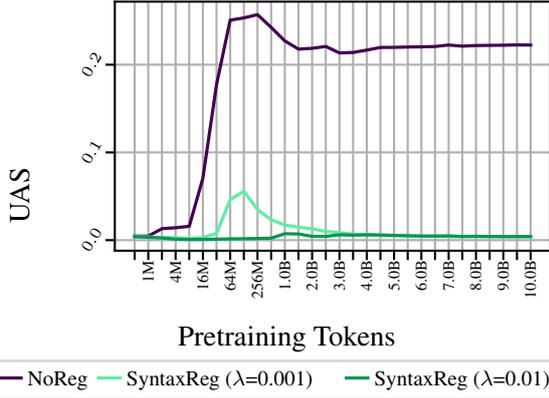


Figure 8: Trajectories of UAS over 10B tokens of pretraining with and without SAS suppression. All models are GPT2 with 2 layers with 8 attention heads.

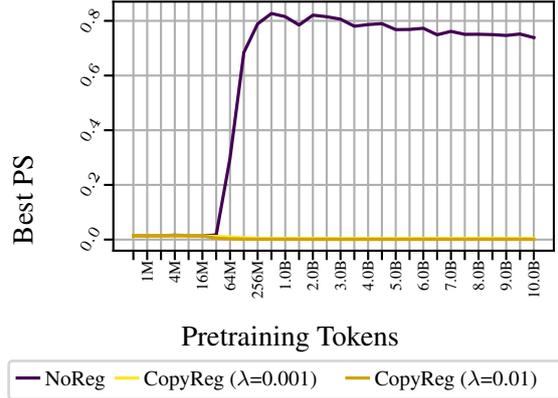


Figure 10: Trajectories of PS over 10B tokens of pretraining with and without copying suppression. All models are GPT2 with 2 layers with 8 attention heads.

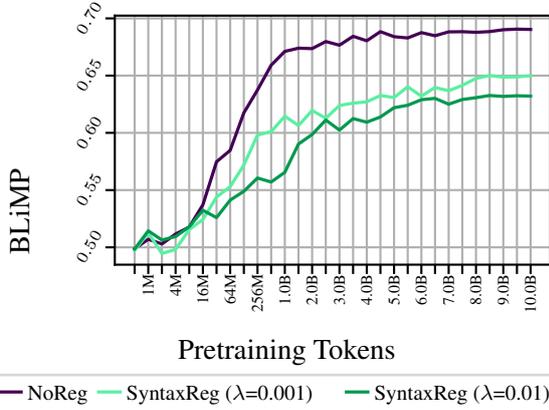


Figure 9: Trajectories of BLiMP score over 10B tokens of pretraining with and without SAS suppression. All models are GPT2 with 2 layers with 8 attention heads.

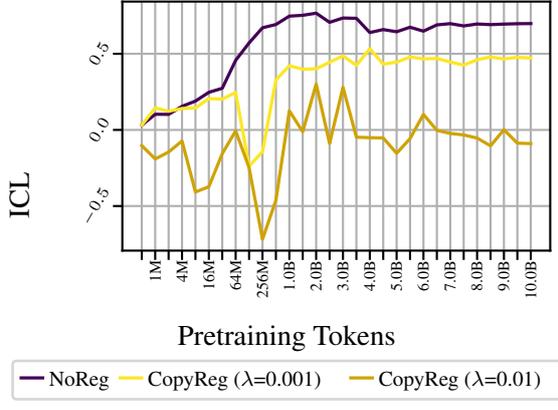


Figure 11: Trajectories of ICL score over 10B tokens of pretraining with and without copying suppression. All models are GPT2 with 2 layers with 8 attention heads.

## J Regularization

### J.1 UAS and BLiMP

Figure 8 summarizes the development of UAS for the GPT2 models with and without SAS suppression. As opposed to the non-suppressed model (NoReg), whose UAS abruptly increases between 16M and 64M pretraining tokens, the SAS-suppressed model with  $\lambda = 0.001$  exhibits a brief increase in UAS, followed by a gradual degradation, converging to almost 0 towards the end of the pretraining. The SAS-suppressed model with stronger suppression,  $\lambda = 0.01$ , almost never sees any improvement in UAS throughout the pretraining. We confirm that the SAS suppression is working as intended.

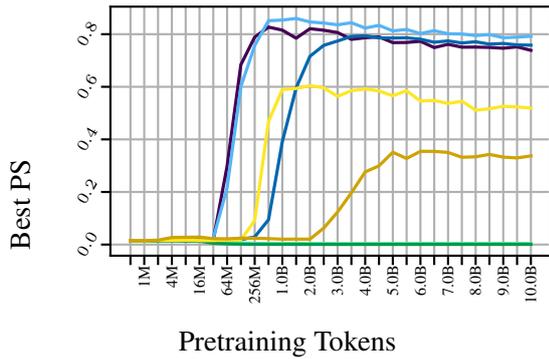
Figure 9 summarizes the development of BLiMP score over the course of pretraining for the same

3 models discussed above. After 16M to 64M pretraining tokens, when the phase transition was *supposed to* happen, BLiMP score perfectly correlates with the SAS suppression strength, with non-regularized model performing the best, and the strongly regularized model performing the worst.

### J.2 PS and ICL

Figure 10 plots the development of best PS. Recall that PS is a head-level score, and we show the score from the highest scoring head for each checkpoint. At 64M pretraining tokens, we see the phase transition in non-regularized model. Both copying-suppressed models ( $\lambda = 0.001$ ,  $\lambda = 0.01$ ) show almost no improvement in the best PS throughout the course of pretraining.

Figure 11 shows that the copying-suppressed models indeed show lower ICL scores, presu-



Pretraining Tokens

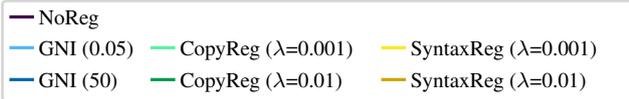
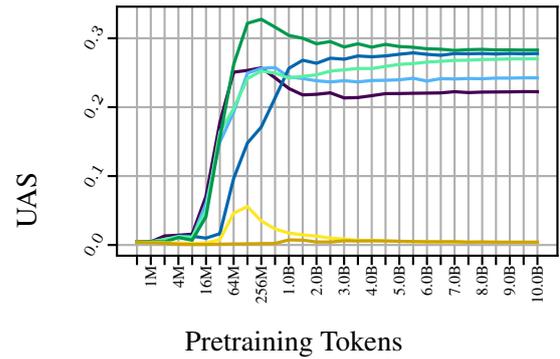


Figure 12: Trajectories of PS over 10B tokens of pre-training. All models are GPT2 with 2 layers with 8 attention heads.



Pretraining Tokens

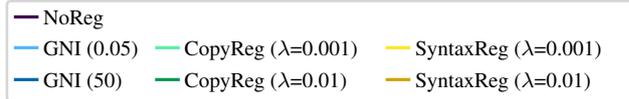
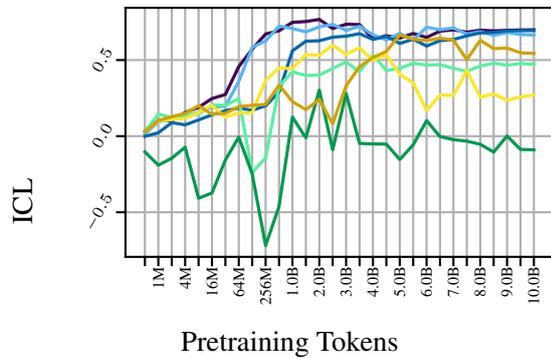


Figure 14: Trajectories of UAS over 10B tokens of pre-training. All models are GPT2 with 2 layers with 8 attention heads.



Pretraining Tokens

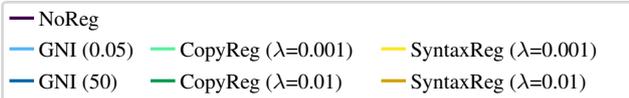
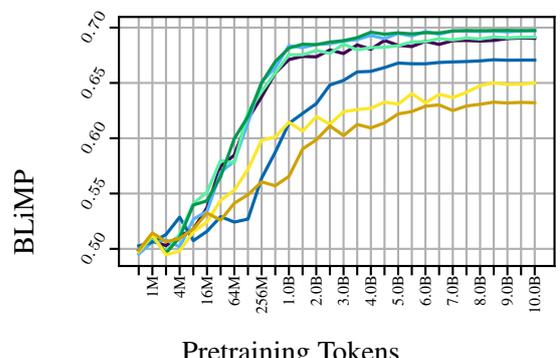


Figure 13: Trajectories of ICL over 10B tokens of pre-training. All models are GPT2 with 2 layers with 8 attention heads.

ably as a result of the absence of induction heads. Notably, at around 64M pretraining tokens, when phase transition occurs in the non-regularized model, both regularized models see a large dip in the ICL score. This is perhaps because the emergent copying behavior briefly raises the penalizing term of the loss function  $\mathcal{L}_{\text{COPY}}$ , resulting in a brief phase of negative ICL scores.

## K Interaction

We find that suppressing one type of specialized head affects the behavior of the other type of specialized heads. As briefly discussed in Section 7.3, because SAS always emerges *before* induction heads, it is expected that the suppression of SAS



Pretraining Tokens

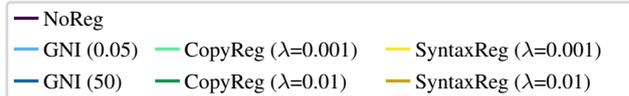


Figure 15: Trajectories of BLiMP over 10B tokens of pre-training. All models are GPT2 with 2 layers with 8 attention heads.

affects the performance of induction heads, but not vice versa.

### K.1 SAS supression $\rightarrow$ induction heads

In Figure 12, we can see that the suppression of SAS indeed leads to slower emergence and weaker performance of induction heads, and this trend is more pronounced for a stronger regularization, where  $\lambda = 0.01$ . Consequently, in Figure 13, SAS suppression makes the improvement in ICL slower, and the eventual ICL lower.

### K.2 Induction head supression $\rightarrow$ SAS

On the other hand, as expected, Figure 14 show that the suppression of induction heads does not sup-

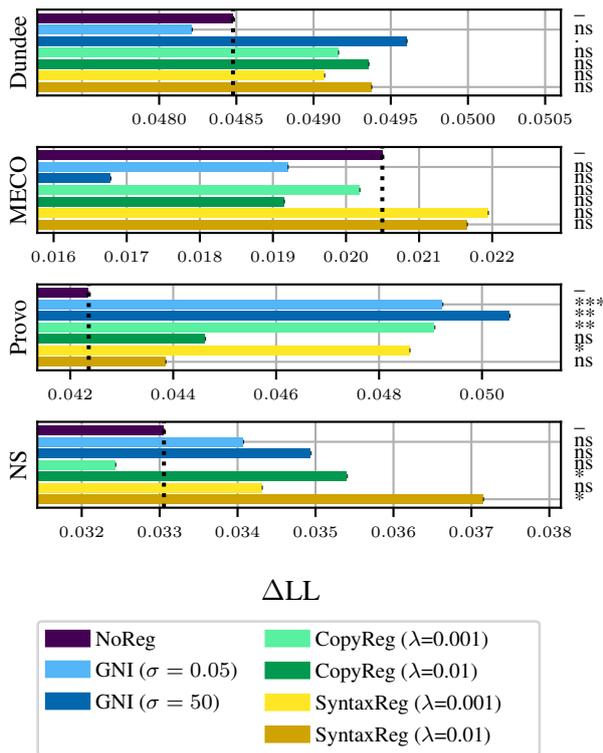


Figure 16:  $\Delta LL$  of 7 variants of 2 layer GPT2 models at the end of pretraining (10B tokens). Red dotted lines represent the  $\Delta LL$ s obtained from the non-regularized GPT2. Within each corpus, regularized models are compared against the non-regularized model (dotted lines) to test statistical significance. Note that p-values are before applying any correction for multiple testing.  $p > .1$  (ns);  $p < 0.1$  (.);  $p < 0.05$  (\*);  $p < 0.01$  (\*\*)

press the emergence of SAS; rather, surprisingly, it *improves* UAS for both values of  $\lambda$ . Consequently, with the stronger attention to dependency edges, models with induction head suppression performs better in BLiMP (Figure 15).

### K.3 Asymmetry between SAS and Induction Heads

Why does the suppression of SAS *delays and weakens* the formation of induction heads, whereas the suppression of induction heads *strengthens* SAS? Given that SAS always emerges before induction heads, one possible explanation is that the presence of natural grammar (the fact that paying strong attention to a token connected with a dependency edge facilitates next token prediction) may have a scaffolding effect for learning non-uniform and highly concentrated attention distribution, with the majority of the mass allocated to a single token. Hence, the removal of this scaffolding (i.e. SAS suppression) delays the emergence of induc-

tion heads. Note that the strong SAS suppression, with almost 0 UAS throughout (yellow line in Figure 14), still produces some level of PS (yellow line in Figure 12). This implies that SAS is a mere scaffolding for induction head formation, and not a prerequisite.

On the other hand, since SAS emerges before induction heads, SAS seems not to be using such a scaffolding from the formation of induction heads. Rather, in natural language with  $\langle A, B, \dots, A, B \rangle$  sequences, induction heads may be *inhibiting* the model’s ability to attend to tokens connected with dependency edges, by allocating some of the attention weights, that *could have been* used for SAS, to prefix-matching tokens. Note that, in our copy suppression, prefix-matching tokens were at least 50 tokens apart from each other, making dependency-connected tokens and prefix-matching tokens disjoint sets for most cases. However, these are pure speculations, and further studies are needed to confirm exactly why the specialized heads’ training dynamics interact in the way we observed.

## L Gaussian Noise Injection

The models with Gaussian Noise Injection (GNI) were trained on the same dataset as other models (CC100), while adding to FFNs noises sampled from a Gaussian distribution with the mean of 0:

$$\tilde{h} = h + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2) \quad (12)$$

where  $h$  is the activation of a transformer block’s FFN. As  $\sigma$  determines the magnitude of the noise, we included models with  $\sigma = 0.05$  and  $\sigma = 50$ .

## M Details of the Effect of Suppression on PPP

In Section 6.2, we mainly reported (1) the correlation between phase transition points and PPP tipping points, and (2) loss-PPP correlations before and after the phase transition. This is because PPP is known to be a function of loss (*quality-power hypothesis*; Goodkind and Bicknell, 2018) as well as whether the model has gone through the tipping point (or equivalently, the phase transition, based on our hypothesis). Hence, comparing the absolute PPPs is not necessarily informative: model A having a higher PPP than model B can mean (1) both are before the phase transition and A has a lower loss, or (2) both are after the phase transition and A has a higher loss, among other possibilities. Nonetheless, for the sake of completeness, we

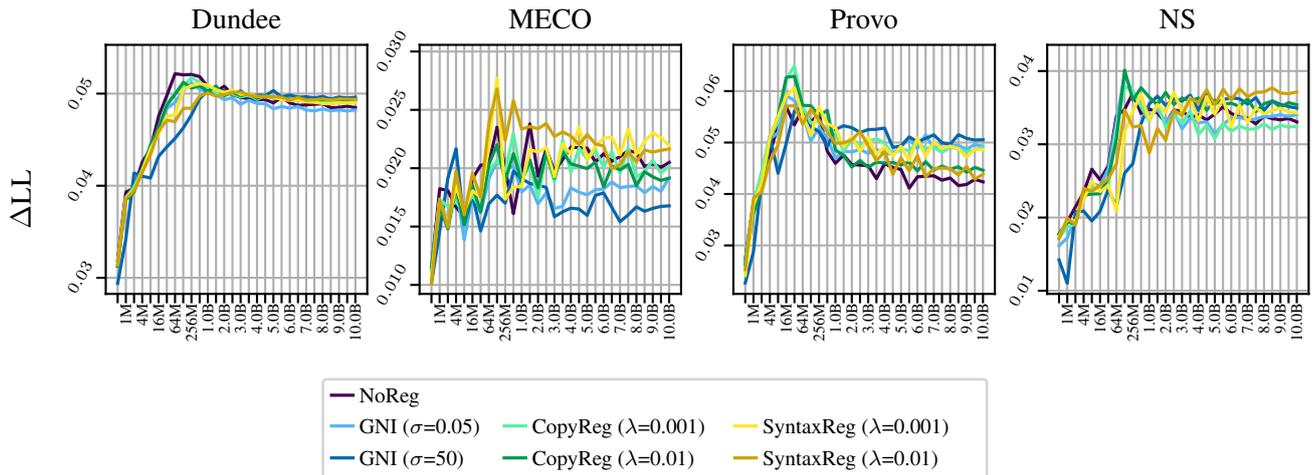


Figure 17: Effect of regularization on  $\Delta LL$ . All models are 2-layer GPT2 models, trained with a batch size of 4.

present the absolute differences, as well as their statistical significance (Figure 16), and the full trajectories of PPP (Figure 17).

In Figure 16, first, we find that syntactic regularization tends to improve PPP over non-regularized models, and this is robust to the settings of  $\lambda$  or the corpus. Second, syntactic regularization is also more effective at improving PPP than copying regularization, which shows mixed results; its impact is negative at the lower lambda value ( $\lambda = 0.001$ ) for the Natural Stories corpus, whereas for MECO, it adversely affects PPP regardless of the regularization strength. See Appendix I for details on the permutation tests for statistical significance.

In Figure 17, the degradation in PPP at the phase transition point is partly suppressed, although the results vary by regularization type, strength, and corpus. For example, for the Dundee corpus, syntactically regularized models have lower PPP around the transition point compared to the original model; however, they degrade less than the original model and consequently end up with higher PPP later in pretraining. This trend is much more pronounced in the Natural Stories Corpus; as opposed to the original model, whose PPP decreases after the transition point, SAS-suppressed models either plateau ( $\lambda = 0.001$ ) or keep improving ( $\lambda = 0.01$ ) beyond that point. However, the degradation trend is not fully suppressed or reversed for many models and corpora.

## N Loss and PPP

As discussed above, comparing PPPs independent of loss is misleading, as they have been found to

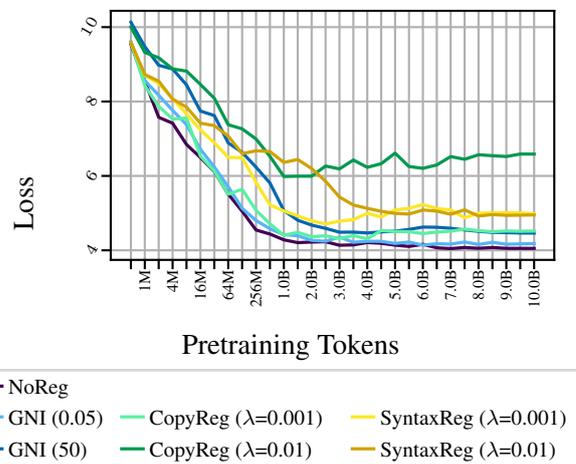


Figure 18: Trajectories of validation loss on over 10B tokens of pretraining. All models are GPT2 with 2 layers with 8 attention heads.

interact. Figure 18 summarize the validation loss of each model throughout the pretraining on 10B tokens. Most models follow a similar pattern, where the loss clearly improves until around 1B tokens, after which the improvement slows down or stops. Copy-regularized model with  $\lambda = 0.01$  seems to be an exception: validation loss gradually goes *up* from around 1B pretraining tokens, likely due to the strong suppression of induction heads.

Table 4 summarizes the loss-PPP correlations before (top) and after (bottom) the phase transition. This corresponds to Figure 4 in Section 6.2. Evidently, most correlations are strongly and significantly negative before the phase transition, whereas about a half of the correlations are significantly positive after the transition. This again highlights the

Corpus		Model						
		NoReg	GNI ( $\sigma=0.05$ )	GNI ( $\sigma=50$ )	CopyReg ( $\lambda=0.001$ )	CopyReg ( $\lambda=0.01$ )	SyntaxReg ( $\lambda=0.001$ )	SyntaxReg ( $\lambda=0.001$ )
Pre-PT	Dundee	<b>-0.971</b>	<b>-0.935</b>	<b>-0.992</b>	<b>-0.975</b>	<b>-0.991</b>	<b>-0.981</b>	<b>-0.975</b>
	MECO	<b>-0.76</b>	-0.443	-0.531	<b>-0.821</b>	<b>-0.84</b>	<b>-0.67</b>	<b>-0.892</b>
	Provo	<b>-0.971</b>	<b>-0.767</b>	<b>-0.974</b>	<b>-0.987</b>	<b>-0.979</b>	<b>-0.843</b>	<b>-0.781</b>
	NS	<b>-0.96</b>	<b>-0.894</b>	<b>-0.957</b>	<b>-0.879</b>	<b>-0.905</b>	<b>-0.811</b>	<b>-0.94</b>
Post-PT	Dundee	<b>0.811</b>	<b>0.519</b>	<b>0.513</b>	0.201	-0.087	-0.137	<b>0.793</b>
	MECO	0.044	<b>0.742</b>	<b>0.525</b>	-0.32	-0.261	-0.151	<b>0.521</b>
	Provo	<b>0.892</b>	0.414	<b>0.712</b>	<b>0.88</b>	<b>0.801</b>	0.209	<b>0.72</b>
	NS	0.294	0.132	<b>0.452</b>	-0.221	<b>-0.605</b>	-0.32	<b>-0.737</b>

Table 4: Pearson’s  $r$  between loss and PPP before (top) and after (bottom) the phase transition. Statistically significant **positive** and **negative** correlations are boldfaced and colored in **green** and **red**, respectively.

observation that the loss-PPP correlation is indeed flipped (i.e. goes through a tipping point) for both non-regularized and regularized models.