

Human-like Fleeting Memory Improves Language Learning but Impairs Reading Time Prediction in Transformer Language Models

Abishek Thamma^{1,2} Micha Heilbron^{1,3}

¹ University of Amsterdam, Amsterdam Brain and Cognition; Amsterdam, Netherlands

² Vrije Universiteit Amsterdam, Department of Informatics; Amsterdam, Netherlands

³ Max Planck Institute for Psycholinguistics; Nijmegen, Netherlands

{a.thamma,m.heilbron}@uva.nl

Abstract

Human memory is fleeting. As words are processed, the exact wordforms that make up incoming sentences are rapidly lost. Cognitive scientists have long believed that this limitation of working memory may, paradoxically, *help* in learning language – an idea supported by classic connectionist modelling work. The rise of Transformers appears to challenge this idea, as these models can learn language effectively, despite lacking working memory limitations or other architectural recency biases. Here, we investigate the hypothesized benefit of fleeting memory for language learning in tightly controlled experiments on transformer language models. Training transformers with and without fleeting memory on a developmentally realistic training set, we find that fleeting memory consistently improves language learning (as quantified by both overall language modelling performance and targeted syntactic evaluation) but, unexpectedly, impairs surprisal-based prediction of human reading times. Interestingly, follow up analyses revealed that this discrepancy – better language modeling, yet worse reading time prediction – could not be accounted for by prior explanations of why better language models sometimes fit human reading time worse. Together, these results support a benefit of memory limitations on neural network language learning – but not on predicting behavior.

1 Introduction

Human memory is fleeting: as a reader or listener processes language, the exact wordforms that make up incoming sentences are quickly and inescapably lost. A popular belief in cognitive science is that this inherent limitation of human memory may paradoxically *help* in language learning (Newport, 1990; Christiansen and Chater,

2016; Rowland, 2013). First, the fleetingness of human memory is thought to impose a recency bias, guiding resources to more relevant, local regularities. Second, it is thought to provide an *incentive for abstraction*: because the memory capacity for recent exact wordforms is so limited, the learner would be driven to discover abstractions (such as chunking, linguistic structure) as a means of compression, rather than learn surface-level statistical regularities between words (see Christiansen and Chater, 2016 for review). The benefits of memory limitations for learning were famously studied in neural networks in classic connectionist modelling work by Elman (1993). Using simulations of a toy language and simple recurrent neural network models, Elman found that reducing ‘working memory’ by limiting the input window during training helped the network learn the artificial grammatical structure more effectively. Due to Elman’s work and that of others (Christiansen and Chater, 2016), the notion that the limitation of memory paradoxically benefits language acquisition became highly influential and can be seen as a cornerstone of the cognitive science of human learning (e.g. Dehaene, 2021; Rowland, 2013).

However, the success of the transformer architecture appears to be at odds with this idea. Transformer language models can learn language effectively, including abstract syntactic structure (Hu et al., 2020; Linzen and Baroni, 2021; Wilcox et al., 2024; Futrell and Mahowald, 2025) despite having perfect memory of words within their context window, and lacking any form of inherent recency bias or memory decay, unlike earlier recurrent neural language models. Yet, as the transformer’s success derives from multiple factors including scale, it does not directly inform whether fleeting memory benefits learning, even within modern neural networks, especially in a developmentally realistic data regime.

Here, we investigate the benefit of fleeting




Figure 1: Fleeting memory, combined with an echoic memory buffer, improves language learning, consistently across training runs and training set sizes. (A) Validation loss (lower is better) comparison between models with perfect memory and fleeting memory. (B) BLiMP accuracy (higher is better), a measure of linguistic knowledge, for the same conditions. In all panels, dots with connecting lines are paired training runs with identical weight initialisation and data sampling. Stars indicate significance levels of the pairwise (within-seed) differences (bootstrap t-test against zero): $*p < 0.05$ (*), $*p < 0.01$ (**), $*p < 0.001$ (***)

memory on language learning in modern neural networks directly, using transformers as a testbed. To test for the effect of fleeting memory on language learning in transformers, we propose *fleeting memory transformers*, a simple modification of the transformer architecture that adds a memory decay to the self-attention operation to simulate human-like forgetting. By running tightly controlled experiments in which we train models with and without fleeting memory on a developmentally plausible training set (BabyLM; Warstadt et al., 2023), we ask whether fleeting memory i) benefits language learning, based on language internal evaluation; ii) renders language models more ‘human-like’, in terms of their ability to pre-

dict human reading behaviour. At human-scale data, these dimensions typically correlate: better models predict reading times better (Wilcox et al., 2020; Goodkind and Bicknell, 2018); moreover, a more human-like architecture might independently better capture human processing.

Our experiments show that, consistently across training runs and network initialisations, adding a human-like memory decay improves both overall language modelling performance, and accuracy on targeted syntactic evaluation – but only when the memory decay is combined with an “echoic memory buffer”, which perfectly retains the initial 3-7 words. Intriguingly, however, the models, while consistently better on both language modelling performance and syntactic evaluation, perform worse at surprisal-based prediction of human reading behavior, in a way that cannot be accounted for by prior explanations of why better language models sometimes do worse at reading time prediction.

2 Related work

Studies on the hypothesized benefits of memory limitations for language learning in neural networks go back to the early years of connectionism. In particular, Elman (1993) reported that systematically increasing the complexity of input to recurrent neural networks improved the models’ ability to learn simple toy languages. Elman explored this concept of “starting small” in two ways: by imposing resource constraints, in the form of memory limitations that would ‘simplify’ the processing, and secondly by incrementally increasing the complexity of training material itself, in what later became known as *curriculum learning*.

More recently, in the context of transformers, several papers have revisited Elman’s ideas of “starting small” as a means to improve learning efficiency in developmentally realistic data ranges in the context of the BabyLM challenge (Warstadt et al., 2023). However, these studies have focused on the curriculum aspect of starting small – overall not finding consistent benefits of curriculum learning – and not on memory limitations, as we do here. Moreover, note that in Elman’s original work, the memory capacity started small but gradually expanded over training. By contrast, we consider a fixed memory limitation, simulating the more general effect of limited memory on learning (Christiansen and Chater, 2016; Dehaene, 2021;




Figure 2: (a) Standard Transformers have a retention over entire context window (left), while in our implementation the retention decays as function of distance of current word to the past (Right).

(b) Standard causal self-attention pattern is usually applied to mask future tokens for causal language modeling (Left), while in our implementation an additional bias is added to the causal self-attention mask, where the attention mask fades as function of distance (Right). Shade of red denotes strength of mask. Each word can only attend to itself or preceding words.

Rowland, 2013); though making memory capacity dynamic could be an interesting future extension (see *Discussion*).

Incorporating a recency bias into transformers has also been explored outside the context of cognitively inspired language modeling. Most prominently, Press et al. (2022) propose ALiBi, a technique that adds linear biases to the attention scores to penalize them proportional to their distance. However, the primary goal of their approach is to allow for length generalization during inference time and overcome the constraints posed by traditional positional encoding, not to incorporate a human-like recency bias. Additionally, ALiBi’s learnable per-head biases (e.g., 36 coefficients in a 6-layer, 6-head setup) results in a potentially quite complicated recency bias, that tends to operate over much larger timescales than human working memory, making it difficult to quantify “memory decay strength” as a single interpretable value. Rather than competing with such approaches on length generalization, we utilize a different tool for a different purpose: we impose a single, fixed recency bias at human working memory scales, specifically to test the cognitive hy-

pothesis that memory limitations benefit learning (Elman, 1993).

In the cognitive modeling domain, a range of recent works have explored memory limitations or recency biases in language models, motivated by the cognitive implausibility of transformers’ perfect verbatim memory (Armeni et al., 2022). However, most of these works only explore the effect of memory limitations or recency bias on *inference*, i.e. on top of a pretrained language model without such limitations, rather than on its effect on learning. Such memory-based accounts of processing difficulty suggest that humans struggle when retrieving distant or similar items from working memory during language comprehension. Various approaches have been used to model these limitations by constraining either memory capacity/width (e.g., Timkey and Linzen, 2023, limiting the number of heads), or implementing a decay that creates imperfect memory representations. The lossy-context surprisal framework (Futrell et al., 2020a) explores this second approach, proposing that human processing difficulty reflects prediction from such imperfect memory representations. Building on this, Hahn

et al. (2022) developed a resource-rational model that optimizes which words to retain in memory, showing that lossy memory better predicts human reading times than perfect memory, particularly for complex recursive structures like center embeddings. However, these models apply memory constraints to already trained language models to explain processing difficulty, rather than examining whether such constraints benefit learning itself.

Formally most similar to our approach, Vaidya et al. (2023) include a power-law decay into self-attention to simulate fleeting memory in pre-trained GPT-2 models. However, they too did not focus on learning, but on a very specific form of inference: processing of repeated texts. They describe how, for repeating texts, human and language model next-word predictions substantially diverge (as the language model predictions become almost perfect) – but that a memory decay at inference time reduces this divergence and improves human LM similarity for such repeating texts. Recently, and beyond the context of repeated texts, De Varda and Marelli (2024) also explored adding a recency bias to pre-trained GPT-2 models. They found that a recency bias improved reading time prediction – but this bias was fitted to optimize reading times directly, and the optimal bias was very small in magnitude.

Clark et al. (2025) extend this approach to training. Focusing on reading-time prediction, they compared two types of recency bias, namely the one proposed by De Varda and Marelli (2024) and ALiBi (Press et al., 2022). They found that when applied at train-time, the approach of De Varda and Marelli (2024) did not lead to measurable improvements in reading time prediction, while ALiBi improved reading time prediction accuracy. While they interpret this improvement in the light of human memory decay, the exact link seems less direct, because as mentioned, the linear attentional biases in ALiBi are not cognitively inspired, and tend to operate over a different timescale. Moreover, the bias is different for different heads, making it difficult to compare the ‘overall’ bias to human memory decay.

Concurrent work by Mita et al. (2025) revisits Elman’s dynamic approach, starting with strong working memory constraints and relaxing them over training until the model concludes with perfect memory. While they report improvements on

a targeted syntactic benchmark, these emerge primarily late in training as constraints are significantly relaxed, and do not evaluate language modeling loss or behavioral alignment.

In our work, we evaluate the effects on learning of incorporating an analogue of memory decay, implemented using a fixed, non-trainable, cognitively inspired recency bias with interpretable hyperparameters (decay rate and buffer size). We evaluate the performance of this approach both in terms of language learning, i.e. performance on overall language modeling and targeted syntactic evaluation, and the ability to predict human behavioral data through reading time fit. Critically, all analyses are statistically evaluated and performed using controlled experiments that specifically isolate the effect of fleeting memory, while controlling for stochasticity across training runs.

3 Methodology

3.1 Architecture

We study fleeting memory in the context of standard transformer-based autoregressive language model, following the GPT2 architecture (Radford et al., 2019). We focus on transformers because they lack any inherent locality bias – attention is uniform over the context window by default – making them the cleanest testbed. GPT-2 is a vanilla decoder-only transformer, representative of the broader model class. The model size was scaled down to match the babyLM dataset size, roughly following scaling trends from Kaplan et al. (2020), resulting in a model of 6 layers, 6 attention heads, and a hidden state dimensionality of 384. A BPE tokenizer with a vocabulary of 8,000 words was constructed and the models were trained for 44,000 iterations (approximately ~ 200 epochs over the 10M dataset and ~ 20 epochs for the 10M), with a batch size of 32. The models trained have 13.69M trainable parameters.

3.2 Fleeting memory implementation

We implement fleeting memory as a fixed, non-trainable recency bias applied to the self-attention weights after the softmax operation. This modifies the standard attention mechanism as follows:

$$\text{Attn}_{\text{FM}}(Q, K, V) = \left(\text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) \odot B \right) V \quad (1)$$

Note that after masking, attention weights do not sum to one; however, the subsequent Layer Normalization is scale-invariant, so only the relative

attention pattern affects learning. The bias matrix B contains retention values between 0 and 1, which decay as a function of token distance d . Motivated by models of forgetting in cognitive science (Donkin and Nosofsky, 2012; Lin and Tegmark, 2017), we use a power-law retention function:

$$B(d) = \begin{cases} 1 & \text{if } d < E \\ 1 - \left(\frac{d-E+1}{n-E}\right)^{\frac{1}{e\alpha}} & \text{if } E \leq d < n \end{cases} \quad (2)$$

This piecewise function includes an “echoic memory” buffer – an initial period of size E where retention is perfect. For subsequent tokens ($d \geq E$), retention decays as a power-law over the remainder of the context window, n . The steepness of this decay is controlled by the hyperparameter α , and e is Euler’s number.

3.3 Dataset

In order to train the models in a human-like data quantity and register, models were trained on the training set of the BabyLM challenge (Warstadt et al., 2023). For our primary experiments, the 10M dataset was used, which corresponds to the amount of input received by children in their first 2-5 years. Subsequently, we extended our analysis to the 100M dataset, which corresponds to the rough equivalent of the input received by age 12. The dataset is heterogeneous, containing a variety of data sources such as CHILDES (MacWhinney, 2004), OpenSubtitles (Lison and Tiedemann, 2016), children’s stories, and others, with approximately 56% being transcribed speech and about 40% being child-directed or child-appropriate language.

3.4 BLiMP

BLiMP is a benchmark for linguistic minimal pair judgments (Warstadt et al., 2020), containing 67 datasets of 1,000 minimal pairs across 12 linguistic phenomena. Models are evaluated on whether they assign a higher probability to the acceptable sentence in each pair, analogous to human acceptability judgments. Scores represent the proportion of correct judgments per subtask, with the overall score calculated as an unweighted average across all subtasks following the BabyLM evaluation pipeline.

3.5 Reading Time Analysis

We evaluate the models’ fit to human reading times using linear mixed-effects regression models. We compare a baseline model against a full model that incorporates LM surprisal along with the baseline predictors; the resulting difference in Log-Likelihood (ΔLL) quantifies how much better the model’s predictions explain human processing delays beyond what these lexical baselines already capture. This approach isolates the unique contribution of the language model to processing difficulty.

Our primary analysis employs a rigorous baseline specification. We model fixed effects for low-level lexical features (word length and unigram frequency from Brysbaert and New, 2009) for the current word, as well as for the two preceding words to account for spillover effects. For the Dundee corpus, we additionally control for whether the previous word was skipped. Crucially, we include by-subject random slopes for all fixed effects (both current and lag features) to ensure conservative statistical estimates. These models were fit using `lme4` (Bates et al., 2015) on data from Natural Stories (self-paced reading times from 181 participants reading 10 stories, 10,256 words) and the Dundee Corpus (gaze durations from 10 subjects reading 20 news articles, 51,500 words).

Finally, to ensure our results are robust and not driven by a specific choice of covariates, we fitted an alternative “item effects” model (Oh and Schuler, 2023b; Oh et al., 2024). This specification replaces the explicit lexical covariates (length, frequency, and their spillovers) with a random intercept for each unique word type, directly factoring out item-specific variance.

Frequency-based Error Analysis

To investigate whether impairment in reading time prediction could be explained by memorization as proposed for very large language models (Oh and Schuler, 2023b; Oh et al., 2024), we conducted a follow-up analysis that assessed the prediction performance as a function of word frequency. To do so, we first partitioned the data from each corpus into five quintiles based on log-frequency and calculated the Mean Squared Error (MSE) of the regression residuals for each quintile. To probe for the second signature of memorization – systematic underprediction of reading times for rare

words – we subsequently separated the data points within each quintile into underpredicted and overpredicted instances. For this step, following [Oh et al. \(2024\)](#), we computed the Sum of Squared Errors (SSE) to quantify total error, because the number of under- and overpredicted items can differ across models.

3.6 Statistical Testing

When analyzing model performance for the evaluation measures, multiple seeds for each configuration (10 seeds) were tested. This was done in order to correct for stochasticity in training and also analyze consistency across seeds, where each trained model is treated as a separate ‘participant’. We define a model ‘participant’ by its random seed, which determines both the initial weight configuration and the data sampling order (batch shuffling). By fixing the seed across conditions, we ensure paired comparisons are *ceteris paribus*: the models differ only in their memory architecture, while the initialization and training data sequence remain identical (up to GPU floating-point precision). Statistical testing was performed across participants (seeds) and since the number of participants was low, data-driven bootstrap t-tests were used. These involve resampling a null-distribution with zero mean (by removing the mean), counting across bootstraps how likely a t-value at least as extreme as the true t-value was to occur. Each test used at least 10^4 bootstraps; p values were computed without assuming symmetry (equal-tail bootstrap; [Rousselet et al., 2023](#)). Confidence intervals (in the figures and text) were also based on bootstrapping.

3.7 Code

Models were built using custom code in PyTorch, built on top of nanoGPT ([Karpathy, 2022](#)). For the BLiMP evaluation, we used the evaluation pipeline created by the organizers of the BabyLM challenge. All models were trained on A100 GPUs on Snellius supercomputing cluster.

4 Results

Naive fleeting memory impairs language learning

We first evaluated a “naive” implementation of fleeting memory, i.e. a memory retention function that immediately starts decaying from the first token into the past (see Fig 3a). We assessed the ef-

fect of such a memory decay on overall language modeling performance, as quantified via the final loss on the validation dataset from the BabyLM challenge, for a range of memory decay strengths. To isolate the effect of memory decay from other factors contributing to validation loss variability, we trained a range of models with different random seeds (affecting weight initialization and data sampling) for every decay value, but shared random seeds across decays, to allow for paired comparisons (see Fig 3b and *Methods*).




Figure 3: **A)** Naive fleeting memory. Retention values with respect to token recency, for different decay rates. **B)** Effect of naive fleeting memory on cross-entropy loss across five seeds. **C)** Fleeting memory with and without echoic memory, illustrated for a decay rate of 2. **D)** Effect of echoic memory on cross-entropy loss, for a decay rate of 2, versus perfect memory (no decay). For the consistency across decay strengths and EM-buffers, see Figure 5. In all panels, dots with connecting lines represent individual (pairwise) training runs for different conditions. Stars indicate significance-levels of the pairwise (within-seed) difference, across different runs (bootstrap t-test against zero): $p < 0.05$ (*), $p < 0.01$ (**), $p < 0.001$ (***).

Interestingly, this revealed that naive fleeting memory *impairs* language model performance: validation loss was lowest for models without memory decay (i.e. perfect memory models) and increased slightly but monotonically as a function of decay strength, highly consistently across training run seeds (Figure 3b).

Fleeting memory improves language learning when combined with echoic memory buffer

At face value, the initial result challenges the hypothesis that memory limitations can improve language learning. However, qualitative inspection of the models suggested that the naive fleeting memory decay introduced too strong a decay. In particular, we observed from model completions that naive fleeting memory introduced many spelling errors, suggesting even the most local (within-word, across-token) dependencies were disrupted. This stands in contrast with human memory, which, while fundamentally fleeting, retains the most recent past near-perfectly due to a brief sensory buffer period ('echoic' or 'iconic' memory; [Baddeley, 1992](#)). To account for this, we explored a modified version of fleeting memory, in which the memory decay is preceded by an echoic memory buffer: an initial period of 5-10 tokens (± 3 -7 words) during which the retention is perfect (Fig 3C; see [Harrison et al., 2020](#) for a similar approach in music modeling).

When we evaluated these models, in a larger experiment (10 seeds per condition), we first confirmed that naive fleeting memory impairs language modelling performance, and that this effect was statistically significant (for a decay of 3: mean $\Delta L = +0.0073$, 95% CI $[+0.0044, +0.0103]$, bootstrap t -test: $p = 0.0008$; see Figure 3D). Strikingly, however, for fleeting memory models with an echoic memory buffer, we observed the opposite pattern: a lower validation loss compared to perfect memory models, both for echoic memory of 5 (mean $\Delta L = -0.0059$, 95% CI $[-0.0089, -0.0028]$, bootstrap t -test: $p = 0.0138$) and 10 (mean $\Delta L = -0.0135$, 95% CI $[-0.0177, -0.0090]$, bootstrap t -test: $p = 0.0012$). To confirm that these results were not an artifact of a specific hyperparameter choice, the analysis was repeated across 12 distinct combinations of decay strength and echoic memory buffer size (120 training runs in total). This confirmed that naive fleeting memory (no buffer) consistently impaired performance, while the combination of an echoic buffer and sufficiently strong decay consistently improved it (Figure 5).

To further evaluate the robustness of this effect, we repeated the experiment on the larger 100M BabyLM training set, finding the same effect: fleeting memory improves language modelling performance (reduces loss), and this effect

is highly consistent across training runs (mean $\Delta L = -0.0207$, 95% CI $[-0.0237, -0.0169]$, bootstrap t -test: $p < 0.0001$; see Figure 1A). For the subsequent evaluations, we focus on a representative model configuration with an echoic memory buffer of 10 tokens (approximating human capacity) and an intermediate decay rate of $\alpha = 2$. Note, however, that the general benefits of fleeting memory are robust across a wide range of hyperparameter settings (see Figure 5)

Fleeting memory improves linguistic knowledge

We then evaluated the effect of fleeting memory not just on next-word prediction ability (cross-entropy loss) but also on syntactic knowledge, using targeted syntactic evaluation on the benchmark of linguistic minimal pairs (BLiMP; [Warstadt et al., 2020](#)). Indeed, we observed that fleeting memory models showed improved syntactic knowledge – and this was found consistently for both models trained on the 10M training set (mean $\Delta_{acc} = 1.04\%$, 95% CI $[0.52\%, 1.54\%]$, bootstrap t -test: $p = 0.0086$) and those on the 100M training set (mean $\Delta_{acc} = 2.30\%$, 95% CI $[1.35\%, 3.17\%]$, bootstrap t -test: $p = 0.0102$; see Figure 1B). This suggests that the benefits of fleeting memory extend beyond simple next-token prediction to more abstract and theoretically significant aspects of language processing.

Differential impact of fleeting memory across linguistic subtasks

When we subdivided the aggregate BLiMP scores across the twelve broad linguistic phenomena (see appendix; Figure 9 and 10), we observed that the improvement conferred by fleeting memory is not uniform, but was driven by a specific subset of phenomena. Notably, consistent and significant gains were evident in tasks such as Subject-Verb Agreement, Anaphor Agreement, and Argument Structure, which rely on relatively local syntactic context and dependencies, where a fleeting memory bias would theoretically be advantageous.

Critically, fleeting memory did not significantly impair accuracy on any of the phenomena in BLiMP. Nevertheless, several phenomena exhibited no improvement. Indeed, for both the 10M and 100M models, "Irregular Forms" remained unaffected, as anticipated, as this is a test of lexical knowledge rather than dependencies between tokens. "Determiner Noun Agreement" also

showed consistently no improvements – which is expected, since it involves among the most local regularities that fall within the echoic memory buffer.

Fleeting memory impairs reading time prediction

The improvements in language modeling and syntactic knowledge raised a natural question: would fleeting memory models also better capture human language processing? We tested this hypothesis by assessing behavioral alignment in the form of surprisal-based reading time prediction. We first assessed self-paced reading times from the Natural Stories corpus (Futrell et al., 2021). Strikingly, we found that fleeting memory significantly *impaired* reading time prediction. This was robust across training scales, with models trained on the 10M dataset showing a decrease in delta-log-likelihood of $\Delta LL = -176.29$ (95% CI $[-253.67, -95.93]$, bootstrap t -test: $p = 0.0033$) and models trained on the 100M dataset showing the same pattern (mean $\Delta LL = -98.13$, 95% CI $[-165.43, -22.99]$, bootstrap t -test: $p = 0.0365$; see Figure 4A).

We considered whether this might be an artifact of the self-paced reading paradigm, where reading times are assessed through button presses, which is not very naturalistic. Overall, the same pattern was found: for both training sets, fleeting memory induced a quantitative impairment on the delta log likelihoods (10M: mean $\Delta LL = -16.92$, 95% CI $[-29.89, -4.34]$, bootstrap t -test: $p = 0.02$, 100M: mean $\Delta LL = -18.30$, 95% CI $[-32.41, -5.00]$, bootstrap t -test: $p = 0.01$; Figure 4B). Crucially, we obtain equivalent results using the alternative “item effects” regression model (see Appendix; Figure 11), confirming that our findings are robust to regression model specification choices.

This establishes a curious dissociation: fleeting memory enhances objective language modeling performance while simultaneously degrading the models’ ability to predict human processing patterns.

The impairment in reading time prediction cannot be fully accounted for by prior explanations

The finding that fleeting memory improves language model quality while impairing reading time prediction appears paradoxical, but echoes a




Figure 4: Fleeting memory impairs surprisal-based prediction of human reading times, across different datasets and training set sizes. (A) $\Delta \text{Log-Likelihood}$ (higher is better) for predicting self-paced reading times on the Natural Stories corpus. (B) $\Delta \text{Log-Likelihood}$ (higher is better) for predicting gaze durations on the Dundee eye-tracking corpus. In all panels, dots with connecting lines are paired training runs with identical weight initialization and data sampling. Stars indicate significance levels of the pairwise (within-seed) differences (bootstrap t -test against zero): *ns* (not significant), $*p < 0.05$ (*), $*p < 0.01$ (**).

broader pattern in the literature. Early work established that better language models predict human reading times better (Goodkind and Bicknell, 2018; Wilcox et al., 2020), but recent studies have documented a reversal: superior models often exhibit degraded reading time prediction (Oh et al., 2022; Shain et al., 2024), spawning several explanatory hypotheses (Oh and Schuler, 2023b,a; Oh et al., 2024; Aoyama and Wilcox, 2025). We examined whether these prior accounts could explain our results.

The *scale hypothesis* attributes this inversion to superhuman training data quantities, proposing

that the positive relationship breaks down around 2 billion tokens—far exceeding human linguistic exposure (Oh and Schuler, 2023a). However, this cannot explain our results: we observe the effect with human-scale training data, at just 10M words.

A second influential explanation is the *memorization hypothesis*. This starts from the observation that the impairment in reading time prediction is not uniform but is concentrated at specific words, in particular specific open-class words such as proper nouns and named entities (Oh and Schuler, 2023b), and low-frequency items more generally (Oh et al., 2024). The reasoning here is that the best models show human-unlike memorization, allowing them to predict infrequent words too well. This account predicts two empirical signatures: (i) maximal prediction degradation for infrequent words, and (ii) systematic reading time underprediction, as memorization drives surprisal values for very rare tokens unrealistically low (Oh and Schuler, 2023b; Oh et al., 2024).

Testing these predictions reveals an intriguing mismatch. We observe the first signature: across both corpora and model scales, fleeting memory’s impairment concentrates on low-frequency items (see appendix; Figure 6). However, we find no evidence that the worse reading time prediction of low-frequency words is driven by the underprediction of reading times for low-frequency words, the key signature of memorization-based accounts (Figure 7,8). This indicates that while frequency appears to mediate the effect, the mechanism differs from earlier explanations based on superhuman memorization of low-frequency items.

5 Discussion

Our experiments demonstrate that memory limitations can benefit language learning in transformer language models, supporting theories from cognitive science. When equipped with an echoic memory buffer, fleeting memory models consistently outperformed perfect memory controls on both overall language modeling and targeted syntactic evaluation, on a developmentally realistic training set. However, these same models exhibited impaired reading time prediction. The impairment cannot be attributed to superhuman-scale training data or memorization mechanisms identified in prior work as causes of degraded reading time fit for better language models, suggesting that multiple distinct mechanisms can cause a dissoci-

ation between language model quality and reading time prediction accuracy.

Before discussing theoretical implications, it is important to note the magnitude of the observed effects. The reduction in cross-entropy loss is subtle – roughly the same magnitude as the run-to-run variability from different random seeds. However, our methodology, which isolates the effect of fleeting memory across paired training runs, demonstrates that the benefit is consistent. Moreover, this small but reliable advantage transfers to a more pronounced and meaningful improvement on BLiMP. This consistency and transferability validate the effect as a genuine consequence of memory decay, justifying further interpretation.

A paradoxical aspect of these results is that we find that memory decay improves learning in *transformers* – the very architecture famous for lacking the recency biases characteristic of previously dominant RNNs. Yet we demonstrate that reintroducing memory limitations enhances both language modeling performance and syntactic knowledge acquisition. However, we should stress that we do not believe the results to reflect something about transformer language models in general. We focus specifically at the human-scale data regime (10-100M tokens). In this data-limited regime, memory decay provides a useful inductive bias: It encodes the statistical fact that most linguistic dependencies are local (Futrell et al., 2020b). The benefits would likely diminish at contemporary pretraining scales (billions to trillions of tokens), where models can discover this pattern from data alone. Pinpointing the exact data scale at which benefits disappear—and whether model size plays an independent role—remains open, though computationally expensive to determine and orthogonal to our focus on developmentally plausible regimes. Moreover, our results may be specific to the developmental corpus we used; texts with genuinely long-range dependencies – academic papers, novels, code, etc. – might show different patterns.

What mechanisms underlie the benefits of fleeting memory? The most parsimonious explanation is statistical: memory decay serves as a beneficial inductive bias. While a standard Transformer could in principle learn to attend preferentially to local context, discovering this pattern requires data. Fleeting memory encodes this prior directly, improving sample efficiency – a computational in-

stantiation of the ‘less is more’ hypothesis (Newport, 1990; Elman, 1993). A richer possibility is that fleeting memory creates an incentive for abstraction: because exact wordforms decay rapidly, the learner must discover higher-level abstractions to preserve information (Christiansen and Chater, 2016). The pattern of BLiMP improvements can be seen as in line with this abstraction hypothesis: fleeting memory helps in learning syntactic phenomena requiring structural analysis. However, our current findings are not sufficient to conclude whether fleeting memory operates only through statistical biasing or (also) drives qualitatively different learning. Future work could try to do so with more mechanistic analyses – for instance by examining layer-wise abstraction, attention head specialization, or the temporal dynamics of feature emergence (Aoyama and Wilcox, 2025).

While we observe that fleeting memory hurts reading time prediction, two recent studies reported the opposite. De Varda and Marelli (2024) found that a recency bias improved behavioral fit when added to pretrained GPT-2; Clark et al. (2025) observed benefits with ALiBi. This discrepancy probably stems from methodological differences. De Varda and Marelli (2024) applied decay post-hoc to pretrained models and fit decay parameters directly to optimize reading times, thereby potentially limiting theoretical conclusions. Clark et al. (2025) failed to replicate this result with the decay functions of De Varda and Marelli (2024), obtaining improvements only with ALiBi – but ALiBi’s head-specific, trainable biases spanning many tokens are more difficult to compare to human memory constraints. By contrast, we impose an interpretable, fixed decay during training, and find consistent improvements in language modeling and BLiMP accuracy alongside impairments in reading time prediction. The consistency of these effects suggests memory decay (applied over training) affects language models differently than these prior works anticipated. One might object that our retention function was not entirely fixed – specifically, the echoic buffer was added after observing poor performance of the naive decay (Figure 3). However, unlike approaches that optimize decay parameters directly against language modelling performance or reading time fit, our echoic buffer was motivated by human memory, was never trainable end-to-end, and was based on the 10M dataset alone. All ef-

fects on BLiMP, reading times, and generalization to the 100M dataset emerged downstream from this theoretically motivated, fixed constraint.

Our results exemplify an emerging paradox in computational psycholinguistics: superior language models often predict human reading behavior worse (Oh and Schuler, 2023b; Oh et al., 2022) (see also Oh and Linzen, 2025). Our models are not fully in this regime – which has been documented for models trained on billions of tokens – as our 100M models predict reading times better than 10M models (especially for Dundee, Figure 4B), indicating that model quality still correlates positively with behavioral fit at our data scales. Yet within each training scale, fleeting memory creates this dissociation: improved language models alongside degraded reading time prediction. Prior explanations – based on superhuman scale (Oh and Schuler, 2023a) or memorization of low-frequency words (Oh et al., 2024) – cannot account for our findings. This suggests that multiple mechanisms can produce the competence-alignment dissociation.

The power-law decay we use is interpretable and not fitted on the data directly, but only represents a crude approximation of human memory. Future work could extend it in at least two dimensions. First, human forgetting is content-sensitive and non-monotonic. Some words persist despite temporal distance, while others vanish quickly. Content-dependent retention functions (as in Hahn et al., 2022) could preserve high-information words while letting predictable ones fade, better approximating how humans selectively retain meaningful content. Second, human retention changes across development. Children’s working memory expands with age, altering what linguistic patterns they can acquire (Gathercole et al., 1992). Following Elman (1993), models could begin with severe memory constraints that gradually relax. Indeed, concurrent work by Mita et al. (2025) demonstrates that such a memory-based curriculum can benefit syntactic abilities, although their models conclude training with perfect memory. A natural extension would be to combine such dynamic relaxation with a more human-like endpoint, starting severely constrained and relaxing to fleeting memory.

Ultimately, our findings lend new support to the classic view that cognitive limitations can serve as a beneficial inductive bias for learning. Yet, the

paradoxical impairment in reading time prediction reveals a surprising distinction: a model that is more human-like in its architectural constraints is not necessarily more human-like in its prediction of human behavior.

Acknowledgments

We thank the anonymous reviewers and the action editor Dilek Hakkani-Tur for their constructive feedback, which has greatly helped improve this work. This work has been partly funded by the Dutch Research Council (NWO) under Veni grant VI.Veni.231G.043 to MH.

References

- Tatsuya Aoyama and Ethan Wilcox. 2025. [Language models grow less humanlike beyond phase transition](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 24938–24958, Vienna, Austria. Association for Computational Linguistics.
- Kristijan Armeni, Christopher Honey, and Tal Linzen. 2022. [Characterizing verbatim short-term memory in neural language models](#). In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 405–424, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Alan Baddeley. 1992. [Working memory](#). *Science*, 255(5044):556–559.
- Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. [Fitting linear mixed-effects models using lme4](#). *Journal of Statistical Software*, 67(1):1–48.
- Marc Brysbaert and Boris New. 2009. [Moving beyond Kučera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English](#). *Behavior Research Methods*, 41(4):977–990.
- Morten H. Christiansen and Nick Chater. 2016. [The Now-or-Never bottleneck: A fundamental constraint on language](#). *Behavioral and Brain Sciences*, 39:e62.
- Christian Clark, Byung-Doh Oh, and William Schuler. 2025. [Linear recency bias during training improves transformers’ fit to reading times](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 7735–7747, Abu Dhabi, UAE. 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. Association for Computational Linguistics.
- Stanislas Dehaene. 2021. *How We Learn: Why Brains Learn Better Than Any Machine... for Now*. Penguin.
- Chris Donkin and Robert M. Nosofsky. 2012. [A power-law model of psychological memory strength in short- and long-term recognition](#). *Psychological Science*, 23(6):625–634.
- Jeffrey L. Elman. 1993. [Learning and development in neural networks: the importance of starting small](#). *Cognition*, 48(1):71–99.
- Richard Futrell, Edward Gibson, and Roger P. Levy. 2020a. [Lossy-context surprisal: An information-theoretic model of memory effects in sentence processing](#). *Cognitive Science*, 44(3):e12814.
- 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(1):63–77.
- Richard Futrell, Roger P. Levy, and Edward Gibson. 2020b. [Dependency locality as an explanatory principle for word order](#). *Language*, 96(2):371–412.
- Richard Futrell and Kyle Mahowald. 2025. [How linguistics learned to stop worrying and love the language models](#). *Behavioral and Brain Sciences*, page 1–98.
- Susan E. Gathercole, Catherine S. Willis, Hazel Emslie, and Alan D. Baddeley. 1992. [Phonological memory and vocabulary development](#)

- during the early school years: A longitudinal study. *Developmental Psychology*, 28(5):887–898.
- 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. Association for Computational Linguistics.
- Michael Hahn, Richard Futrell, Roger Levy, and Edward Gibson. 2022. A resource-rational model of human processing of recursive linguistic structure. *Proceedings of the National Academy of Sciences*, 119(43):e2122602119.
- Peter M. C. Harrison, Roberta Bianco, Maria Chait, and Marcus T. Pearce. 2020. PPM-Decay: A computational model of auditory prediction with memory decay. *PLOS Computational Biology*, 16(11):1–41.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. A systematic assessment of syntactic generalization in neural language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. 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:2001.08361v1*.
- Andrej Karpathy. 2022. NanoGPT. <https://github.com/karpathy/nanoGPT>.
- Henry W. Lin and Max Tegmark. 2017. Critical behavior in physics and probabilistic formal languages. *Entropy*, 19(7):299.
- Tal Linzen and Marco Baroni. 2021. Syntactic structure from deep learning. *Annual Review of Linguistics*, 7(Volume 7, 2021):195–212.
- Pierre Lison and Jörg Tiedemann. 2016. Open-Subtitles2016: Extracting large parallel corpora from movie and TV subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 923–929. European Language Resources Association (ELRA).
- Brian MacWhinney. 2004. CHILDES English MacWhinney corpus.
- Masato Mita, Ryo Yoshida, and Yohei Os-eki. 2025. Developmentally-plausible working memory shapes a critical period for language acquisition. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9386–9399, Vienna, Austria. Association for Computational Linguistics.
- Elissa L. Newport. 1990. Maturational constraints on language learning. *Cognitive Science*, 14(1):11–28.
- 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. Publisher: Frontiers.
- Byung-Doh Oh and Tal Linzen. 2025. To model human linguistic prediction, make LLMs less superhuman. *arXiv preprint arXiv:2510.05141v1*.
- 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.

- Ofir Press, Noah Smith, and Mike Lewis. 2022. [Train short, test long: Attention with linear bi-ases enables input length extrapolation](#). In *International Conference on Learning Representations*.
- Alec Radford, Jeff Wu, R. Child, D. Luan, Dario Amodei, and I. Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Guillaume Rousselet, Cyril R Pernet, and Rand R Wilcox. 2023. An introduction to the bootstrap: a versatile method to make inferences by using data-driven simulations. *Meta-Psychology*, 7.
- Caroline Rowland. 2013. [Understanding Child Language Acquisition](#). Routledge.
- 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.
- William Timkey and Tal Linzen. 2023. [A language model with limited memory capacity captures interference in human sentence processing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8705–8720, Singapore. Association for Computational Linguistics.
- Aditya Vaidya, Javier Turek, and Alexander Huth. 2023. [Humans and language models diverge when predicting repeating text](#). In *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 58–69, Singapore. Association for Computational Linguistics.
- Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Bhargavi Paranjabe, Adina Williams, Tal Linzen, and Ryan Cotterell. 2023. [Findings of the BabyLM challenge: Sample-efficient pretraining on developmentally plausible corpora](#). In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 1–34. Association for Computational Linguistics.
- 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.
- 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. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 42.
- Ethan Gottlieb Wilcox, Richard Futrell, and Roger Levy. 2024. [Using computational models to test syntactic learnability](#). *Linguistic Inquiry*, 55(4):805–848.

A Appendix

Supplementary results




Figure 5: Validation loss on 10M dataset, across 12 decay-strength - buffer-size combinations (120 training runs). Without echoic memory (No EM or “naive”), decay impairs performance. At sufficient decay strengths and buffer sizes, fleeting memory consistently improves performance.




Figure 6: Quintile analysis reveals the effect concentrates on low-frequency words. Difference in reading time prediction error, for both Natural Stories and Dundee for both training set sizes, across different words frequency in the reading time datasets. Small dots indicate paired (within seed) difference in MSE (fleeting - perfect), showing that the fleeting memory models are especially worse at low-frequency words.




Figure 7: Reading time over/under-predictions as a function of frequency, for Natural Stories Corpus. Similar to Figure 6, but calculating the difference in total error for words of which the reading time was unpredicted and over-predicted separately. The SSE difference is normalized within each condition (quintile and prediction type) relative to the average ‘Perfect’ model error to allow for comparison across scales. We do not observe that the elevated difference in residual error at low-frequency words is driven specifically by *underpredictions* of reading times, as observed by Oh et al. (2024).




Figure 8: Same as Figure 7, but for Dundee corpus. Here too, we do not observe clearly that the elevated difference in residuals at low-frequency words is driven specifically by *underpredictions* of reading times, as in Oh et al. (2024).




Figure 9: 10M models (with and without fleeting memory) performance on BLiMP accuracy on subtasks across 12 broad linguistic phenomena, sorted by the average improvement numerical by fleeting memory (highest first).




Figure 10: 100M models (with and without fleeting memory) performance on BLiMP accuracy on subtasks across 12 broad linguistic phenomena, sorted by the average improvement numerical by fleeting memory (highest first).




Figure 11: Suprisal based reading time prediction modelled using unique random intercepts for each word-type (as opposed to explicit lexical covariates) shows qualitatively identical results. **(A)** Δ Log-Likelihood (higher is better) for predicting self-paced reading times on the Natural Stories corpus. **(B)** Δ Log-Likelihood (higher is better) for predicting gaze durations on the Dundee eye-tracking corpus. In all panels, dots with connecting lines are paired training runs with identical weight initialization and data sampling. Stars indicate significance levels of the pairwise (within-seed) differences (bootstrap t-test against zero): *ns* (not significant), $*p < 0.05$ (*), $*p < 0.01$ (**).