

# Transformer verbatim in-context retrieval across time and scale

**Kristijan Armeni**  
Johns Hopkins University  
karmeni1@jhu.edu

**Marko Pranjić**  
Jozef Stefan Institute  
Jozef Stefan International  
Postgraduate School  
marko.pranjic@ijs.si

**Senja Pollak**  
Jozef Stefan Institute  
senja.pollak@ijs.si

## Abstract

To predict upcoming text, language models must in some cases retrieve in-context information verbatim. In this report, we investigated how the ability of language models to retrieve arbitrary in-context nouns developed during training (across time) and as language models trained on the same dataset increase in size (across scale). We then asked whether learning of in-context retrieval correlates with learning of more challenging zero-shot benchmarks. Furthermore, inspired by semantic effects in human short-term memory, we evaluated the retrieval with respect to a major semantic component of target nouns, namely whether they denote a concrete or abstract entity, as rated by humans. We show that verbatim in-context retrieval developed in a sudden transition early in the training process, after about 1% of the training tokens. This was observed across model sizes (from 14M and up to 12B parameters), and the transition occurred slightly later for the two smallest models. We further found that the development of verbatim in-context retrieval is positively correlated with the learning of zero-shot benchmarks. Around the transition point, all models showed the advantage of retrieving concrete nouns as opposed to abstract nouns. In all but two smallest models, the advantage dissipated away toward the end of training.

## 1 Introduction

In language models (LMs), successful prediction of upcoming words depends on in-context information. For example, when given the context prompt “*The novel’s plot and symbolism are centered around three objects: a centipede, a parachute, and a waterfall. The first and most important object in the list is the \_\_\_*”, an LM must retrieve the noun (*centipede*) out of all in-context tokens to correctly predict the continuation. In human cognitive science, this ability to flexibly retrieve items from recent context is known as short-term memory and

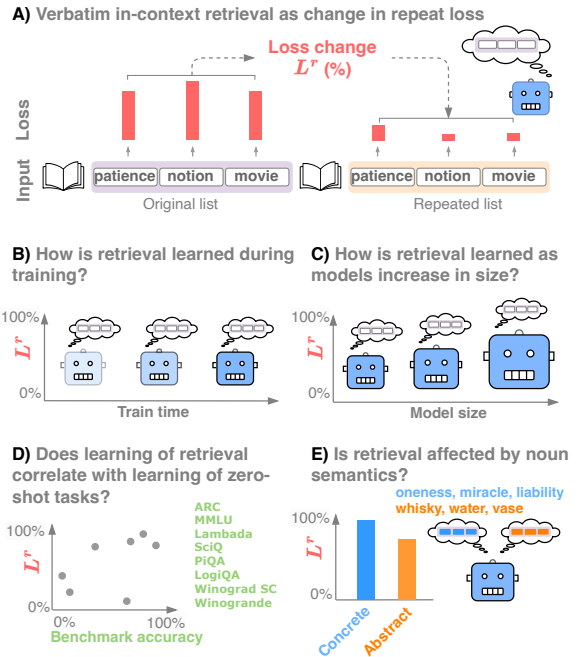


Figure 1: Overview of the approach and experiments.

is believed to be the core computation underlying human cognition (Baddeley, 2003).

Recently, Armeni et al. (2022) showed that a transformer language model (GPT-2, Brown et al., 2020) develops such flexible short-term memory — it was able to retrieve the identity and ordering of lists of *arbitrary* nouns from recent context (Fig. 1, A), even though retrieval of arbitrary in-context information is not the explicit objective of LMs (as opposed to dedicated models of short-term memory, e.g. Oberauer et al. 2018). Yet, studying retrieval in a single *fully-trained* model on *arbitrary* nouns neglects three further dimensions of the capacity: how it is learned, how learning of this dedicated capacity relates to models’ learning of other tasks, and the semantics of retrieved nouns.

First, studying *learning trajectories* of LM capacities offers complementary insights to studying only performance of fully-trained models (e.g.

Chen et al., 2024). Previous work on LM learning trajectories showed that transformers learn next-token prediction by undergoing a *sudden transition* (“phase change”) early during training, which coincides with the development of attention heads that attend to repeated tokens (Olsson et al., 2022). Does verbatim retrieval follow a similar learning trajectory?

Second, the ability to retrieve and predict in-context tokens verbatim (i.e. identity-based matching) can be viewed as a rudimentary form of the more flexible zero-shot learning, where the relevant in-context information is not necessarily given verbatim and must possibly be retrieved based on fuzzy, similarity-based matching (Olsson et al., 2022). How does successful learning of verbatim retrieval relate to LM’s zero-shot performance on more challenging benchmark tasks?

Third, while the successful retrieval of arbitrary nouns underscores the flexibility of transformer short-term memory, this approach neglects that the lexicon of natural language is not a set of unorganized, arbitrary words — instead, it has semantic structure. Two prominent semantic categories are *concrete* and *abstract* nouns. Concrete nouns (e.g. “hammer”) have sensory referents, whereas abstract nouns (e.g. “justice”) do not have a straightforward sensory component. Word concreteness affects human cognitive processing. Children typically acquire concrete words, especially nouns, earlier than abstract words (Gleitman et al., 2005). In certain short-term memory paradigms, humans are better at recalling concrete than abstract words (Taylor et al., 2019). Importantly, the two word categories differ also in their distributional properties: concrete words occur in a semantically narrower range of contexts compared to abstract words (Schulte im Walde and Frassinelli, 2022). Is the transformer retrieval affected by whether nouns refer to concrete vs. abstract entities?

To address these questions, we evaluated verbatim in-context retrieval on the Pythia suite of language models (Biderman et al., 2023). Leveraging the fact that the suite includes pretrained LMs ranging from 14M to 12B parameters in scale and their intermediate training checkpoints across the entire learning epoch, we evaluated how retrieval develops over the course of training and across model sizes (Fig. 1, B and C). Additionally, the Pythia suite contains zero-shot evaluations on various benchmarks for each LM checkpoint. To test how in-context retrieval relates to LM’s

zero-shot performance, we correlated the learning trajectory of the retrieval against the learning trajectories on zero-shot benchmarks (Fig. 1, D). Finally, to test the role of noun semantics for in-context retrieval, we evaluated how noun concreteness, as rated by human participants (Brysaert et al., 2014), affected retrieval over the course of training (Fig. 1, E).

The main contributions of the current work are: **a)** In all models, verbatim retrieval developed in a sudden transition early during training, after about 1% training tokens elapsed, and remained constant during the rest of training, **b)** learning of verbatim retrieval was positively correlated with learning of zero-shot task performance, and **c)** around the transition point, LMs showed an advantage to retrieve concrete rather than abstract nouns. This advantage almost entirely diminished towards the end of training.

## 2 Related work

Several recent studies investigated the behavior of LMs in domain of either verbatim or in-context retrieval more generally. Armeni et al. (2022) developed a paradigm to test the short-term memory ability (in-context retrieval) of LMs. They showed that GPT-2 can retrieve the identity and ordering of repeated arbitrary nouns, but have only tested a single fully-trained LM and did not investigate learning trajectories. Vaidya et al. (2023) compared LM (GPT-2) and human word prediction performance on spans of repeated text. They reported that LMs’ next word prediction performance diverges from human performance on subsequent repetitions. They showed that GPT-2 performance aligned better with humans if its attention heads had a bias towards recent context. Yu et al. (2023) investigated in-context retrieval of facts (e.g. retrieval of the capital city given a country name) and how such retrieval was affected by the pre-training statistics of retrieved facts. They showed that LMs (Pythia) could override retrieval of (counterfactual) in-context information and instead retrieved the fact that has a higher frequency of occurrence in training data (e.g., even when given the in-context counterfactual “*The capital city of Poland is London*” they tend to predict the statistically more likely “*Warszaw*”).

The current report is also related to the recent work on LM interpretability and the role of attention heads in specific forms of retrieval. Several

studies (Elhage et al., 2021; Olsson et al., 2022; Wang et al., 2023; Yu et al., 2023) have identified circuits of attention heads that detect repeated in-context tokens and their previous continuations; the computations governing the behavior investigated presently. These studies focused either on how such attention mechanisms are learned and how they affect generic next-word prediction (Elhage et al., 2021; Olsson et al., 2022) or how these attention mechanisms govern the retrieval of proper nouns as direct objects in sentences (Wang et al., 2022) or factual knowledge (direct objects, Yu et al., 2023).

Here, we complement these lines of work and investigate retrieval as the ability of LMs to retrieve lists of arbitrary combinations of common nouns (unlikely seen co-occurring during training) and their semantic properties.

### 3 Methods

#### 3.1 Verbatim retrieval paradigm

We used the verbatim retrieval paradigm introduced by Armeni et al. (2022). Here, LMs process a short vignette in English where a list of three arbitrary nouns is repeated twice:

*Mary read a list of words: **patience, notion, movie**. After the meeting, she took a break and had a cup of coffee. When she got back, she read the list again: **patience, notion, movie**.*

We refer to the first list of nouns as original list and the second one as repeated list. This setup allows us to test how the LM behavior (as reflected in LM loss, see below) changes as the LM encounters the repeated list. The paradigm (retrieval of arbitrary lists of words) is broadly inspired by benchmarks for testing models of human working memory (Oberauer et al., 2018). Whereas human participants can be tested by just being presented with lists of nouns alone, our paradigm is formatted such that it is more suited to be used as input to LMs: contextualized in a simple, but plausible natural language vignette.

#### 3.2 Quantifying verbatim retrieval

**Change in repeat loss ( $L^r$ )** Following Armeni et al. (2022), we operationalized retrieval as a change in LM loss on repeated nouns. Specifically, we computed the ratio in LM loss  $= -\log_2 P(w_t|w_1, \dots, w_{t-1})$  between each noun in the original list and its repetition  $k$  tokens later:

loss ratio<sup>noun</sup>  $= \frac{\text{loss}(\text{noun}_{i+k})}{\text{loss}(\text{noun}_i)}$ . The loss ratio per list was obtained by averaging the noun-specific loss ratios over the three nouns in a list. A loss ratio  $< 1$  indicates that the loss to the *same tokens* has decreased (that is, the LM expected the token to repeat) and is taken as evidence of verbatim retrieval.

To quantify retrieval as increasing with better performance, we report it as repeat loss change  $L^r = 1 - \text{loss ratio}$ , expressed as percentage. In this way, a 0% change in repeat loss indicates no retrieval whereas a change towards 100% indicates evidence towards (perfect) retrieval. Importantly, repeat loss change is a continuous measure of in-context retrieval, baselined against the LM loss at the beginning of the sequence which facilitates comparison across models (e.g. models that show different baseline loss as expected over the course of training and across scale) and across different types of inputs.

#### 3.3 Language models

**Pythia suite** To evaluate retrieval over the course of training and across scale (see Section 3.4 below), we used the publicly-available pretrained LM checkpoints released as part of the Pythia language modeling suite (Biderman et al., 2023).<sup>1</sup> Pythia is a suite of decoder-only autoregressive transformer LMs spanning from 14M to 12B parameters in size together with 144 intermediate checkpoints stored during training. The models were trained on the Pile dataset (Gao et al., 2020), an English-only corpus for training large-scale LMs containing texts from 22 sources (for example, Common Crawl, Wikipedia, Project Gutenberg, Books3, arXiv etc., see Biderman et al., 2022, for details). The model checkpoints used in this report were trained on the version of the dataset containing approximately 300B tokens. For the full architecture and training details, readers are referred to the original report (Biderman et al., 2023).

In our experiments, we evaluated the following model sizes: {14M, 31M, 70M, 160M, 410M, 1B, 6.9B, 12B} at 18 training checkpoints spanning 6 orders of magnitude across the training steps (in number of training tokens,  $10^6, \dots, 10^{11}$ ) from the initialized to the final fully-trained model<sup>2</sup>. All

<sup>1</sup><https://github.com/EleutherAI/pythia>

<sup>2</sup>Specifically we evaluated the checkpoints from the following training steps: {0, 1, 4, 32, 128, 256, 512, 1000, 2000, 3000, 4000, 8000, 10000, 30000, 40000, 50000, 100000, 143000}. A single step contained 2,097,152 tokens (Biderman

Task	Domain	Reference
AI2 Reasoning Challenge (ARC)	Multiple choice science exams	Clark et al. (2018)
Lambada	Discourse-based word prediction	Paperno et al. (2016)
LogiQA	Logical reasoning	Liu et al. (2020)
Massive multitask lang. understanding (MMLU)	Exam knowledge across diverse domains	Hendrycks et al. (2021)
PiQA	Physical common-sense reasoning	Bisk et al. (2020b)
SciQ	Scientific knowledge	Welbl et al. (2017)
Winograd schema challenge (WSC)	Common-sense reasoning	Levesque et al. (2012)
Winogrande	Common-sense reasoning	Sakaguchi et al. (2021)

Table 1: The benchmark tasks used to compute in learning trajectory correlations in Fig. 3.

model checkpoints were accessed through the HuggingFace Transformers library (Wolf et al., 2020).

### 3.4 Experiments

**Experiment 1: Retrieval of arbitrary nouns across time and scale** In the first experiment, word lists in the vignette were constructed by randomly sampling nouns from the Toronto word pool<sup>3</sup> as used in Armeni et al. (2022). Noun lists in the set (23 lists of 10 nouns) were constructed such that each noun was tested in all 10 possible ordinal positions in the list (e.g. “patience, notion, movie”, “notion, movie, patience”, etc.) to control for any position-specific retrieval effects. This procedure resulted in the final stimulus set that contained  $N = 230$  samples of vignettes. In the present experiment, we used the version of the stimulus set where the list length was capped at 3 nouns.

Evaluating an LM on the full retrieval evaluation suite yields one retrieval score (repeat loss change) per each input vignette. The final retrieval score, per each training step and per model size, was obtained by taking an average across all (in this case  $N = 230$ ) scores. To minimize the potential influence of outliers in averaging, we used the 20% trimmed mean (Wilcox and Keselman, 2003) as the aggregating metric. The results of this experiment are reported in Figure 2.

**Experiment 2: Correlations with zero-shot benchmark learning.** To test how learning of verbatim in-context retrieval relates to the learning of zero-shot benchmark tasks assessing text understanding, we collected the zero-shot evaluation results on various NLP benchmarks that were available for the Pythia suite of LMs<sup>4</sup>. Evaluations were

et al., 2023).

<sup>3</sup><http://memory.psych.upenn.edu/files/wordpools/nouns.txt>

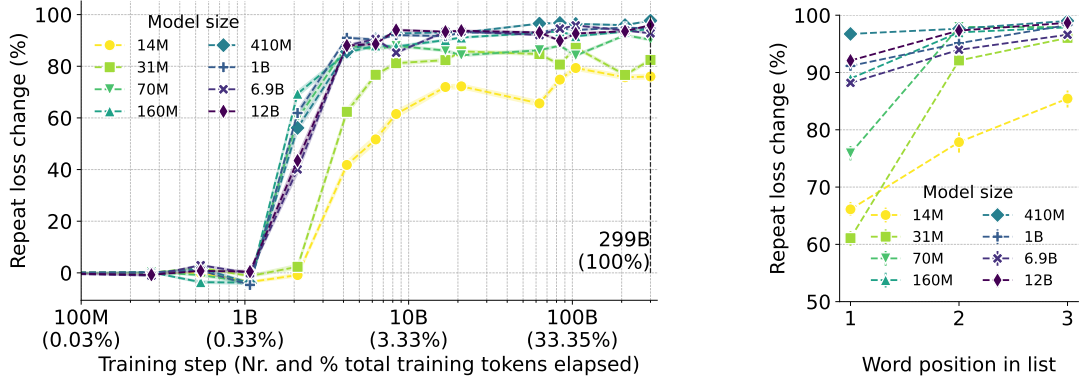
<sup>4</sup><https://github.com/EleutherAI/pythia/tree/main/evals/pythia-v1>

available for the following 6 model sizes: {160M, 410M, 1.4B, 2.8B, 6.9B, 12B} and across 27 checkpoints<sup>5</sup> during training, starting with the initial and ending with the fully-trained model. All individual tasks ( $N = 65$ ) used accuracy as the final metric. The main groups of tasks used in the experiment are summarized in Table 1. See Table 3, Appendix A for the full task list.

For each benchmark task (e.g. Lambada, SciQ etc.), we computed the correlation  $\rho^{traj} = \text{Spearman}(S^{ret}, S^{bench})$  between the learning trajectory of the benchmark task  $S^{bench}$  (i.e. task performance scores across the 27 checkpoints) and the learning trajectory of our verbatim retrieval effect  $S^{ret}$  (i.e. repeat loss change  $L^r$  across the same 27 checkpoints). The Massive multitask understanding benchmark (MMLU, Hendrycks et al., 2021) consists of an array of domain-specific exams (e.g. marketing, clinical knowledge, nursing) which are grouped into 4 higher-level categories (humanities, STEM, social sciences, and ‘other (business, health, misc.)’, see Table 3, Appendix A). For these grouped tasks, we first averaged the learning trajectories per each group and then correlated them with verbatim retrieval effect. We used the rank-based Spearman correlation coefficient where a value of 1 indicates a perfect monotonically increasing relationship between two variables and is robust to any deviations from normality in data distributions.

**Experiment 3: Effect of noun concreteness on retrieval.** To test for retrieval of concrete and abstract nouns, we evaluated LMs on the same paradigm as in the first experiment, but the noun lists were composed of either concrete or abstract nouns. We used abstract and concrete English

<sup>5</sup>Checkpoints corresponding to the following Pythia training steps were evaluated: {0, 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1000, 3000, 13000, 23000, 33000, 43000, 53000, 63000, 73000, 83000, 93000, 103000, 113000, 123000, 133000, 143000}.



(a) Capacity to retrieve verbatim repetitions of arbitrary nouns (average across three nouns in a list) is learned early during LM training and predominantly conserved across scale.

(b) Retrieval improves for items later in the list (fully-trained models).

Figure 2: **Retrieval of arbitrary nouns across time and scale.** Each data point represents the 20% trimmed mean across  $N = 230$  observations, shaded areas/error bars are 95% confidence intervals (bootstrap).

nouns collected by Brysbaert et al. (2014) where human participants were asked to indicate “*how concrete the meaning of each word is for you*” by rating each noun on a 5-point rating scale ranging from 1 “abstract (language-based)” to 5 “concrete (experience-based)”. Each word was rated by at least 25 participants and an average score across participants represents each noun’s final rating.

Table 2: The topmost, mid and lowest ranked words and their concreteness ratings for the concrete and abstract noun pool.

Rank	Concrete		Abstract	
	Word	Rating	Word	Rating
1	whisky	5.00	oneness	1.96
250	canister	4.93	respite	1.77
500	eyebrow	4.85	spirituality	1.07

**Concreteness extremes** In our experiments, we used the “concreteness extremes” subset of the noun pool by Schulte im Walde and Frassinelli (2022). This subset contained the 500 nouns ranked as most concrete and 500 nouns ranked as most abstract. To give an idea, the topmost, mid and lowest ranked nouns for each category are shown in Table 2. As in Experiment 1, each noun was presented in all ordinal positions to rule out any position-specific effects. Our final stimulus set contained, for each semantic category,  $N = 498$  input sequences with lists of 3 nouns.

## 4 Results

### 4.1 Verbatim retrieval across time and scale.

**Verbatim retrieval learned early in training across model sizes.** All tested models, from the smallest (14 million parameters) to the largest (12 billion parameters), learned to retrieve verbatim repeated nouns (Fig. 2a). At the end of training, all models above 31 million parameters showed a near 100% repeat loss change, indicating exact retrieval. The smallest two models (14M and 31M parameters) showed weaker, yet still substantial retrieval effect (around 80% change in repeat loss).

Inspecting the dynamics of repeat loss change across training, we see that generally models learned verbatim retrieval early. After about 1B tokens (0.3% of total dataset), the change in repeat loss starts increasing and, for all larger models, plateaus at approximately 4B tokens (less than 5% of the total tokens in the dataset). The smallest two models had a slower learning curve as evidenced in the fact that their repeat loss change plateaued later, after roughly 20B tokens.

To confirm that reduction in repeat loss was due to retrieval of the *original nouns* and not due to LMs simply having more context when encountering nouns at the end of sequence or due to memorization of lists from training data, we evaluated the loss change in the same paradigm but where the nouns in the second list were unrelated to the nouns in the original list (i.e. there were no matching in-context nouns to retrieve). Fig. 6 in Appendix A.1 confirms that no important loss change occurred in this condition (loss change overall remained <

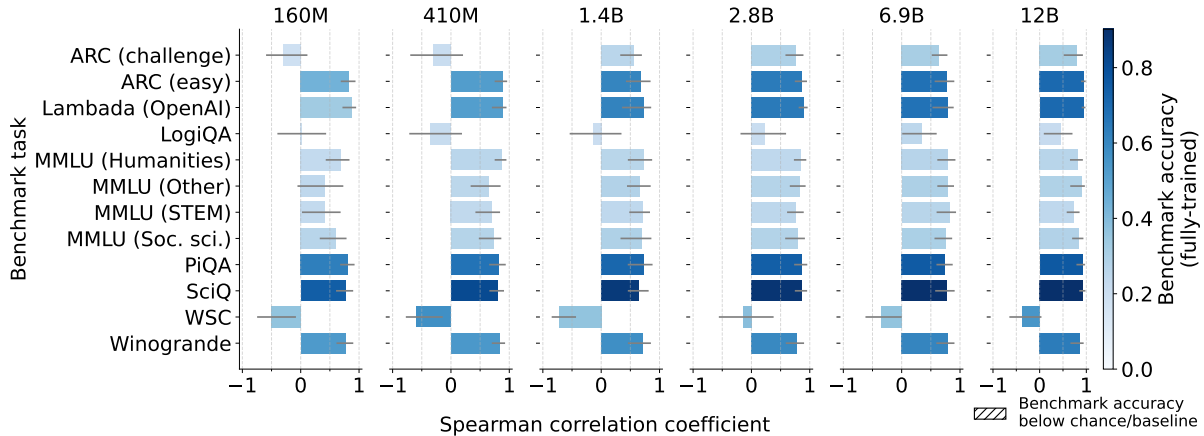


Figure 3: **Correlations between the learning trajectories of verbatim retrieval and select benchmark tasks.** Error bars denote the 95% confidence interval (bootstrap,  $N = 5000$ ) around the correlation coefficient. Color-coded is the benchmark performance accuracy at the end of training for each task. Chance performance for ARC, LogiQA, MMLU\*, LogiQA, SciQ is 0.25, for PiQA, WSC, and Winogrande 0.5. For LAMBADA we threshold against predicting a random in-context token (0.016) (see Table 1 in Paperno et al., 2016). See Table 1 for task descriptions.

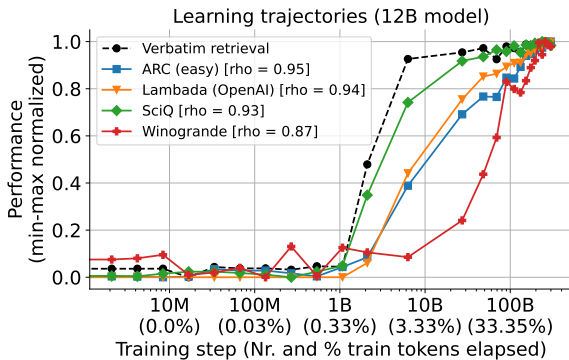


Figure 4: Examples of learning trajectories (12B model) for tasks that showed strongest correlations with verbatim retrieval for the largest tested model. For visualization purposes, accuracy scores are min-max normalized to fall in the  $[0, 1]$  range.

10%), replicating the GPT-2 results by Armeni et al. (2022) and indicating that the change of loss was specific to *verbatim retrieval* of tokens from context.

### Retrieval improves for nouns deeper in the list.

In the previous result, we reported repeat loss change aggregated over all three nouns in the list. Yet, nouns deeper in the list have an advantage because at that point the LM has seen strong evidence of repetition. Does retrieval performance depend on the position in the list?

In Fig. 2b, we report repeat loss change of fully-trained models broken down per noun position within the list. Retrieval indeed becomes better later in the list. While all models show this trend,

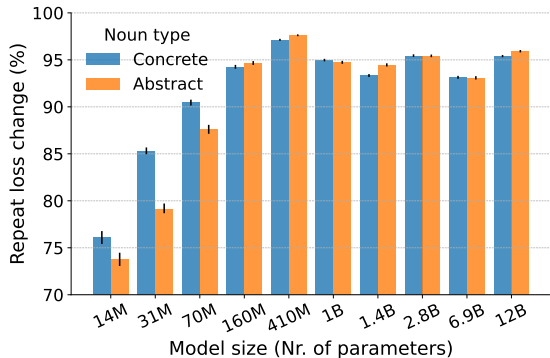
the position-specific advantage is more pronounced for the smaller models (14M, 31M, and 70M). For example, the 70M model shows 62% repeat loss change on the first and a 95% change on the last token in the list. This indicates that subsequent repetitions reinforce the evidence that the model has entered a repeated list and is in line with recent results where next-word prediction performance of GPT-2 improved on spans of repeated text (Vaidya et al., 2023).

## 4.2 Correlations with benchmark task learning

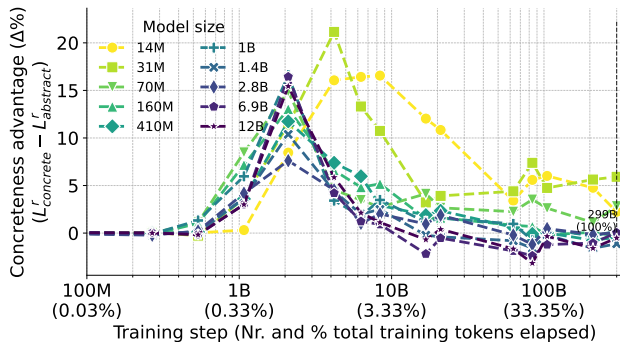
### Learning of verbatim retrieval is positively correlated with zero-shot performance on more challenging benchmark tasks.

In Figure 3 we show the results of the correlation experiment. Generally, most tasks showed a positive correlation with the learning of verbatim retrieval. The correlations and their reliability, as well as the benchmark accuracy itself, tended to increase as the models grow in size, showing that the larger models were more robust learners overall. The highest correlations were observed for the Lambada, PiQA, SciQ, and ARC (easy) benchmarks. For example, the largest 12B model (Figure 4) showed a near perfect rank correlation ( $\rho \simeq 0.95$ ) on the four tasks. These are also the tasks where the model showed generally the highest performance accuracy at the end of training.

For the Winograd schema challenge, LogiQA, and the hard version of the AI2 reasoning challenge, the correlation estimates were generally un-



(a) Concrete vs. abstract noun retrieval at the end of training. The three smallest models show a weak advantage for retrieving concrete nouns.



(b) Retrieval of concrete vs. abstract nouns over the course of training. Around the transition point, all models show an advantage to retrieve concrete nouns.

Figure 5: **Retrieval of concrete and abstract nouns.** **a)** Each bar shows the 20% trimmed mean across  $N = 498$  observations, error bars show 95% confidence intervals (bootstrap). See also Fig. 7 in Appendix A.2. **b)** Each data point concreteness advantage: the difference in mean repeat loss change for concrete vs. abstract nouns.

stable, likely because the performance on these benchmarks was lower to begin with. That is, even though all the models were able to retrieve verbatim in-context tokens, they failed to solve the respective benchmarks in zero-shot settings.

### 4.3 Effect of noun concreteness on retrieval

**Concreteness retrieval advantage observed early during training.** Overall, all models learned to retrieve either abstract or concrete nouns. Repeat loss change at the end of training (on either repeated concrete or abstract nouns) was generally high and ranged from 73% (14M model) to around 95% for models larger than 160M parameters (Fig. 5a). The 14M, 31M and 70M models showed better retrieval for concrete nouns. The effects, although detectable, were small — on average the relative loss change for concrete nouns is greater by between 2% and 6% compared to abstract nouns (see also Fig. 7, Appendix A.2 for visualizations of full distributions).

To test whether nouns semantics affected retrieval during training, we computed the difference in average repeat loss change between concrete and abstract nouns  $\Delta L^r = \bar{L}^r_{concrete} - \bar{L}^r_{abstract}$  across the training checkpoints. The difference curves in Fig. 5b show that around the transition point (1-2B tokens into training), when LMs begin to learn the retrieval, concrete nouns showed 7-17% greater change in repeat loss meaning they were easier to retrieve than abstract nouns. The concreteness advantage occurred in all models and the smallest models (14 and 70M parameters) showed the largest effects.

## 5 Discussion

We showed that transformer LMs learned verbatim retrieval in a sudden transition, early in training, with the performance remaining stable over the course of training. The sharp onset of retrieval capacity around 1-2B tokens in training (approximately 1% total training data) is in line with the results reported by Olsson et al. (2022) who showed that the LM loss over in-context tokens started dropping suddenly 1-2% tokens in training (between 2.5B and 5B tokens). Once the learning change had occurred, the LMs became better at predicting repeated text — which is what was tested in the current work.

The learning trajectory of verbatim retrieval also coincides with the LMs’ learning trajectories on zero-shot benchmarks. This was reflected in the generally high and robust correlations across training for select tasks in our results. Specifically, an abrupt change around 1B tokens in training was observed in the task of predicting the last token of a narrative passage (Lambada, Paperno et al. 2016), multiple choice exams (SciQ, Welbl et al. 2017, ARC Reasoning Challenge, Clark et al. 2018), and in the Winogrande benchmark (Sakaguchi et al., 2021) which requires pronoun resolution based on common-sense reasoning.

Retrieving in-context information (e.g. lists of nouns) verbatim is a basic computation needed for solving a zero-shot multiple-choice task: given a prompt with only in-context instructions (that is, the question and the list of possible answers), an LM system must index and retrieve (i.e. increase

the probability of) the token representing the correct answer. In this sense, retrieving the correct in-context tokens is a necessary step. It is evident, however, that it is not sufficient and that verbatim retrieval must be learned along with other computations.

Consider the Lambada and Winogrande benchmarks, where the task is to predict the passage- or sentence-final word which itself is not predictable on the basis of immediately preceding words. To take an example from the Winogrande benchmark: “Robert woke up at 9am while Samuel woke up at 6am, so **he** had less time to get ready for school”<sup>6</sup>. The task is to answer who the pronoun “he” refers to (Robert or Samuel). To this end, the LM must first establish that 9am is later than 6am — a distinct computational step indicating that “he” refers to “Robert” — and only then retrieve the name to be predicted as the response.

In the final experiment, we show that around the transition point (after  $\approx 1\text{B}$  training tokens), when the capacity for verbatim retrieval occurs, noun semantics affect the retrieval — models showed an advantage to retrieve concrete, as opposed to abstract nouns. Why would LM in-context retrieval be sensitive to noun semantics?

In humans, concrete words, especially nouns, tend to be acquired earlier in development compared to abstract words (Gleitman et al., 2005). This advantage is presumably conferred by hearing words for concrete objects and concurrently observing or interacting with the objects the heard words refer to in the world. LMs as text-based statistical learners by construction have no direct access to word semantics via experience or text-external data (Bisk et al., 2020a). Nevertheless, text statistics, governed by human language use, can serve as a cue to the semantic structure of language — in this case, the lexicon. It is an empirical question whether and what aspects of the linguistic system are in fact recovered by LMs in the service of the next-word prediction objective and subsequently reflected in the LM behavior or internal mechanisms (Manning, 2022; Pavlick, 2023).

We speculate that earlier in training, LMs are leveraging the fact that concrete nouns tend to be used in more predictable, less diverse contexts (Schulte im Walde and Frassinelli, 2022) where presumably token repetition would be more likely to occur. However, once the LMs and the training

compute scale in size, this distributional difference no longer confers an important advantage for retrieval. The phenomenon of concreteness advantage early, but not later in training underscores the general notion that with the increasing amounts of training data, LMs as machine learning systems become incommensurate with human learners (see also Vaidya et al., 2023), who operate on the order(s) of magnitude smaller amount of learning data, at least in terms of number of words — recent estimates point to around 100M words by adolescence (Warstadt and Bowman, 2022).

**Future work.** In this study we investigated retrieval across a diverse set of nouns, and broken down by a core semantic dimension. However, LMs are statistical learners. A dimension of future work will be to disentangle the *learning sources* that LMs leverage to perform retrieval. In a recent study, Yu et al. (2023) showed that pretraining frequency can override the retrieval of counterfactual in-context information. An LM is more likely to predict a proper noun that was frequently occurring in pretraining, e.g. “Warszaw”, even when the counterfactual in-context prompt suggests it should retrieve a different name (“*The capital of Poland is London. What is the capital of Poland?*”). Our present results do not speak directly to this issue as our paradigm does not involve counterfactuals. It is based on lists of arbitrary nouns that unlikely frequently co-occurred in pretraining data. However, it would be important to establish whether and to what extent the in-context retrieval in general is governed by the pretraining frequencies of individual common nouns and to what degree this capacity is robust to pretraining statistics.

Finally, our measure of verbatim retrieval is a behavioral measure insofar that it only takes into account the output of the LM. The field of model interpretability has seen an increased interest in recent years and aims to reverse engineer the computations of LMs (e.g. Olsson et al., 2022; Elhage et al., 2021; Wang et al., 2023; Zhang and Nanda, 2023, among others). Future work could focus on investigating the internal mechanisms and their causal role in transformer in-context retrieval. There is consistent evidence suggesting that LMs develop dedicated attention heads (Olsson et al., 2022; Wang et al., 2023; Yu et al., 2023; Vaidya et al., 2023) governing the retrieval capacity. Whereas this line work frequently focuses on interpretability for practical purposes (e.g. better control of LM output in

---

<sup>6</sup><https://winogrande.allenai.org/>



downstream applications), it would be valuable to simultaneously develop a more fine-grained computational characterization of LM mechanisms interpretable with respect to cognitive science constructs like the short-term memory (Cowan, 2017).

In cognitive neuroscience, language features derived from transformer LMs (contextualized word embeddings) are currently among the best performing when it comes to predicting brain data recorded in human language processing tasks (e.g. Schrimpf et al., 2021; Goldstein et al., 2022; Caucheteux and King, 2022). However, these high-dimensional features and the resulting statistical fits are frequently hard to interpret. Coupled with loose theoretical motivations such high predicting models can be right for the wrong scientific reasons (see e.g. Antonello and Huth, 2023). A better characterization of LM mechanisms in terms of cognitive capacities (e.g. Lakretz et al., 2022) would be instrumental in understanding how and why LMs succeed in modeling human brain and cognitive data.

## 6 Conclusion

Retrieving information from context is an important capacity of transformer language models. In this work, we investigated how the ability to retrieve repeated nouns from context develops across LM training and scale and its dependence on whether the retrieved nouns denote concrete or abstract entities. Retrieval was learned early in training across scale and once learned, it remained stable. Retrieval learning was robustly correlated with learning of zero-shot task performance. Around the point when the in-context retrieval was learned, the models showed advantage to retrieving concrete as opposed to abstract nouns and the advantage dissipated as the models saw more training data.

## 7 Limitations

There are certain limitations to current work. While our test suite was designed to test arbitrary target nouns, we did not investigate whether and how well LM retrieval generalizes to other parts of speech (say, to verbs, adjectives). Similarly, the currently reported paradigm relies on a single vignette, it would be important to use a more diverse set of vignettes to confirm that the results generalize across topic domains. However, given the robustness and size of the effect here and in past reports by others, it is likely that the finding would generalize across a diversity of vignettes. Finally, our results are

limited to English, which currently the dominant language in terms of available resources in language technologies. Extending the study to other languages with, for example, different grammatical properties (e.g. richer noun morphology) or less resources would be a welcome effort.

## Data and Code Availability

The code used to run the experiments is available at: <https://github.com/KristijanArmeni/verbatim-memory-in-NLMs>

The materials and data used in the experiments are available at: <https://doi.org/10.17605/OSF.IO/A6GSW>

## Computational requirements

Experiments described in this report were run on the A100 Nvidia GPU nodes on an high-performance computing cluster (HPC). To evaluate the smallest (14M) model, we requested 8GB of RAM and the evaluation completed on the order of a few minutes. RAM requirements were progressively increased to evaluate larger models. For the largest (12B) model, we requested 80GB of RAM and the evaluation completed in about 30 minutes.

## Acknowledgements

KA would like to thank Christopher Honey and Tal Linzen for scientific and administrative support. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise. KA, MP, and SP want to thank Rebecca Swisdak, Mili Bauer and Živa Antauer for exceptional administrative support. This work was supported by the Slovenian Research and Innovation Agency via the bilateral research project Working Memory based assessment of Large Language Models (BI-US/22-24-170), the research project Embeddings-based techniques for Media Monitoring Applications (L2-50070, co-funded by the Kliping d.o.o. agency) and the core research programme Knowledge Technologies (P2-0103).

## References

- Richard Antonello and Alexander Huth. 2023. [Predictive coding or just feature discovery? An alternative account of why language models fit brain data.](#) *Neurobiology of Language*, pages 1–16.
- 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. 2003. [Working memory: looking back and looking forward](#). *Nature Reviews Neuroscience*, 4(10):829–839.
- Stella Biderman, Kieran Bicheno, and Leo Gao. 2022. [Datasheet for the Pile](#). ArXiv:2201.07311 version: 1.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, and others. 2023. [Pythia: A suite for analyzing large language models across training and scaling](#). In *International Conference on Machine Learning*, pages 2397–2430. PMLR.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020a. [Experience grounds language](#). *arXiv:2004.10151 [cs]*. ArXiv: 2004.10151.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2020b. [PIQA: Reasoning about physical commonsense in natural language](#). In *Thirty-Fourth AAAI Conference on Artificial Intelligence*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. 2014. [Concreteness ratings for 40 thousand generally known English word lemmas](#). *Behavior Research Methods*, 46(3):904–911.
- Charlotte Caucheteux and Jean-Rémi King. 2022. [Brains and algorithms partially converge in natural language processing](#). *Communications Biology*, 5(1):134.
- Angelica Chen, Ravid Shwartz-Ziv, Kyunghyun Cho, Matthew L. Leavitt, and Naomi Saphra. 2024. [Sudden drops in the loss: Syntax acquisition, phase Transitions, and simplicity bias in MLMs](#). ArXiv:2309.07311 [cs].
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? Try ARC, the AI2 reasoning challenge](#). ArXiv, abs/1803.05457.
- Nelson Cowan. 2017. [The many faces of working memory and short-term storage](#). *Psychonomic Bulletin & Review*, 24(4):1158–1170.
- 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*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. [The Pile: An 800GB dataset of diverse text for language modeling](#). *arXiv preprint arXiv:2101.00027*.
- Lila R. Gleitman, Kimberly Cassidy, Rebecca Nappa, Anna Papafragou, and John C. Trueswell. 2005. [Hard Words](#). *Language Learning and Development*, 1(1):23–64. Publisher: Routledge [eprint: https://doi.org/10.1207/s154733411ld0101\\_4](https://doi.org/10.1207/s154733411ld0101_4).
- Ariel Goldstein, Zaid Zada, Eliav Buchnik, Mariano Schain, Amy Price, Bobbi Aubrey, Samuel A. Nastase, Amir Feder, Dotan Emanuel, Alon Cohen, Aren Jansen, Harshvardhan Gazula, Gina Choe, Aditi Rao, Catherine Kim, Colton Casto, Lora Fanda, Werner Doyle, Daniel Friedman, Patricia Dugan, Lucia Melloni, Roi Reichart, Sasha Devore, Adeen Flinker, Liat Hasenfratz, Omer Levy, Avinatan Hassidim, Michael Brenner, Yossi Matias, Kenneth A. Norman, Orrin Devinsky, and Uri Hasson. 2022. [Shared computational principles for language processing in humans and deep language models](#). *Nature Neuroscience*, 25(3):369–380.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Yair Lakretz, Théo Desbordes, Dieuwke Hupkes, and Stanislas Dehaene. 2022. [Can transformers process recursive nested constructions, like humans?](#) In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3226–3232, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2012. [The winograd schema challenge](#). In *13th*

- International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012*, Proceedings of the International Conference on Knowledge Representation and Reasoning, pages 552–561. Institute of Electrical and Electronics Engineers Inc.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2020. LogiQA: A challenge dataset for machine reading comprehension with logical reasoning. [eprint: 2007.08124](#).
- Christopher D. Manning. 2022. [Human language understanding & reasoning](#). *Daedalus*, 151(2):127–138.
- Klaus Oberauer, Stephan Lewandowsky, Edward Awh, Gordon D. A. Brown, Andrew Conway, Nelson Cowan, Christopher Donkin, Simon Farrell, Graham J. Hitch, Mark J. Hurlstone, Wei Ji Ma, Candice C. Morey, Derek Evan Nee, Judith Schewpe, Evie Vergauwe, and Geoff Ward. 2018. [Benchmarks for models of short-term and working memory](#). *Psychological Bulletin*, 144(9):885–958.
- 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*.
- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. 2016. [The LAMBADA dataset: Word prediction requiring a broad discourse context](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1525–1534, Berlin, Germany. Association for Computational Linguistics.
- Ellie Pavlick. 2023. [Symbols and grounding in large language models](#). *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 381(2251):20220041. Publisher: Royal Society.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. [WinoGrande: an adversarial winograd schema challenge at scale](#). *Communications of the ACM*, 64(9):99–106.
- Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A. Hosseini, Nancy Kanwisher, Joshua B. Tenenbaum, and Evelina Fedorenko. 2021. [The neural architecture of language: Integrative modeling converges on predictive processing](#). *Proceedings of the National Academy of Sciences*, 118(45):e2105646118. Publisher: Proceedings of the National Academy of Sciences.
- Sabine Schulte im Walde and Diego Frassinelli. 2022. [Distributional measures of semantic abstraction](#). *Frontiers in Artificial Intelligence*, 4:796756.
- Randolph S. Taylor, Wendy S. Francis, Lara Borunda-Vazquez, and Jacqueline Carbajal. 2019. [Mechanisms of word concreteness effects in explicit memory: Does context availability play a role?](#) *Memory & Cognition*, 47(1):169–181.
- Aditya R. Vaidya, Javier Turek, and Alexander G. Huth. 2023. [Humans and language models diverge when predicting repeating text](#). ArXiv:2310.06408 [cs].
- Kevin Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2022. [Interpretability in the wild: a circuit for indirect object identification in GPT-2 small](#). ArXiv:2211.00593 [cs].
- Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2023. [Interpretability in the wild: a circuit for indirect object identification in GPT-2 small](#). In *The Eleventh International Conference on Learning Representations*.
- Alex Warstadt and Samuel R. Bowman. 2022. What artificial neural networks can tell us about human language acquisition. In *Algebraic Structures in Natural Language*. CRC Press. Num Pages: 44.
- Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017. [Crowdsourcing Multiple Choice Science Questions](#). In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 94–106, Copenhagen, Denmark. Association for Computational Linguistics.
- Rand R. Wilcox and H. J. Keselman. 2003. [Modern robust data analysis methods: measures of central tendency](#). *Psychological Methods*, 8(3):254–274.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-Art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Qinan Yu, Jack Merullo, and Ellie Pavlick. 2023. [Characterizing mechanisms for factual recall in language models](#). ArXiv:2310.15910 [cs].
- Fred Zhang and Neel Nanda. 2023. [Towards best practices of activation patching in language models: metrics and methods](#). ArXiv:2309.16042 [cs].

## A Appendix A

### A.1 Retrieval control

Figure 6 shows memory retrieval results when nouns are not repeated.

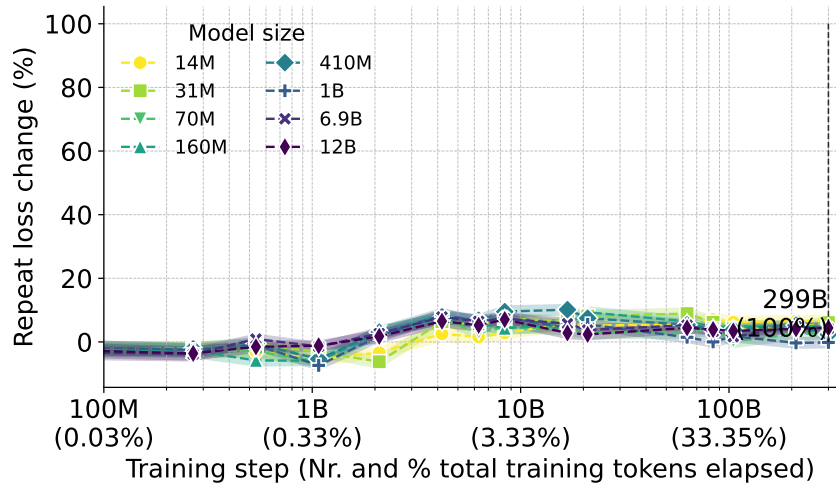


Figure 6: Evaluating repeat loss change in a control condition where there were no verbatim repeated in-context nouns (hence, no retrieval was possible). Each data point shows the 20% trimmed mean across  $N = 230$  observations, shaded areas/error bars are 95% confidence intervals (bootstrap).

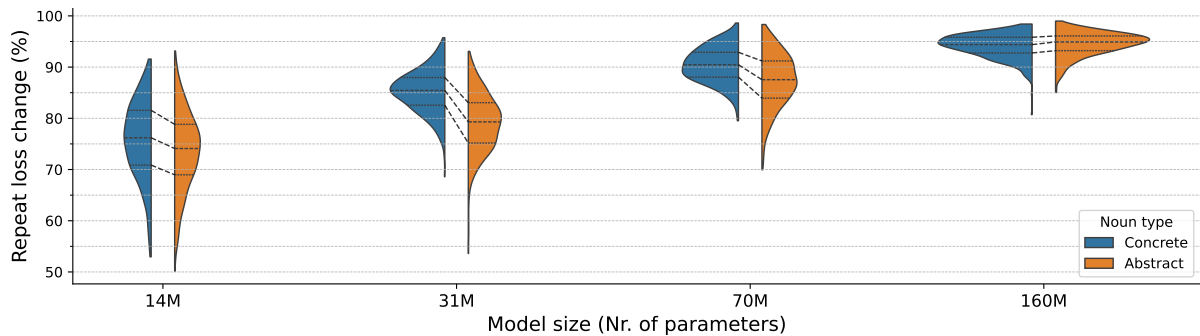


Figure 7: Data distributions comparing retrieval scores for concrete and abstract nouns for 4 smallest models from Fig. 5. Each violin plot KDE density estimates over  $N = 498$  data points. The inner lines show the first, second (median) and the third quartiles of the distribution.

## A.2 Abstract vs. concrete data distributions

The violin plots in Fig. A.2 show the distributions underlying respective bar plots in Fig. 5a.

## A.3 Zero-shot benchmark tasks overview

The full list of benchmark tasks used in Experiment 2 is provided in Table 3.

Table 3: Benchmark categories for the Pythia models. The Task Key column corresponds to the task key used in the Pythia evaluation files (<https://github.com/EleutherAI/pythia/tree/main/evals/pythia-v1>).

	Benchmark Name	Benchmark Subcategory	Task Key
1	ARC (challenge)	None	arc_challenge
2	ARC (easy)	None	arc_easy
3	Lambada (OpenAI)	None	lambada_openai
4	LogiQA	None	logiqa
5	MMLU	MMLU (Soc. sci.)	mmlu_high_school_government_and_politics
6	MMLU	MMLU (Soc. sci.)	mmlu_sociology
7	MMLU	MMLU (Other)	mmlu_business_ethics
8	MMLU	MMLU (Other)	mmlu_medical_genetics
9	MMLU	MMLU (STEM)	mmlu_high_school_physics
10	MMLU	MMLU (Other)	mmlu_professional_medicine
11	MMLU	MMLU (Other)	mmlu_miscellaneous
12	MMLU	MMLU (STEM)	mmlu_college_physics
13	MMLU	MMLU (Humanities)	mmlu_professional_law
14	MMLU	MMLU (Humanities)	mmlu_high_school_world_history
15	MMLU	MMLU (Other)	mmlu_global_facts
16	MMLU	MMLU (Humanities)	mmlu_high_school_us_history
17	MMLU	MMLU (Other)	mmlu_marketing
18	MMLU	MMLU (Soc. sci.)	mmlu_high_school_microeconomics
19	MMLU	MMLU (Other)	mmlu_college_medicine
20	MMLU	MMLU (Soc. sci.)	mmlu_human_sexuality
21	MMLU	MMLU (STEM)	mmlu_electrical_engineering
22	MMLU	MMLU (STEM)	mmlu_elementary_mathematics
23	MMLU	MMLU (STEM)	mmlu_high_school_chemistry
24	MMLU	MMLU (Other)	mmlu_professional_accounting
25	MMLU	MMLU (Humanities)	mmlu_world_religions
26	MMLU	MMLU (STEM)	mmlu_machine_learning
27	MMLU	MMLU (Soc. sci.)	mmlu_high_school_psychology
28	MMLU	MMLU (Humanities)	mmlu_moral_scenarios
29	MMLU	MMLU (STEM)	mmlu_high_school_computer_science
30	MMLU	MMLU (Soc. sci.)	mmlu_security_studies
31	MMLU	MMLU (STEM)	mmlu_computer_security
32	MMLU	MMLU (Humanities)	mmlu_high_school_european_history
33	MMLU	MMLU (STEM)	mmlu_college_computer_science
34	MMLU	MMLU (Soc. sci.)	mmlu_econometrics
35	MMLU	MMLU (STEM)	mmlu_college_mathematics
36	MMLU	MMLU (Other)	mmlu_clinical_knowledge
37	MMLU	MMLU (Soc. sci.)	mmlu_professional_psychology
38	MMLU	MMLU (Other)	mmlu_nutrition
39	MMLU	MMLU (STEM)	mmlu_abstract_algebra
40	MMLU	MMLU (Humanities)	mmlu_logical_fallacies
41	MMLU	MMLU (STEM)	mmlu_astronomy
42	MMLU	MMLU (STEM)	mmlu_high_school_mathematics
43	MMLU	MMLU (STEM)	mmlu_high_school_biology
44	MMLU	MMLU (Soc. sci.)	mmlu_high_school_geography
45	MMLU	MMLU (Other)	mmlu_anatomy
46	MMLU	MMLU (Humanities)	mmlu_jurisprudence
47	MMLU	MMLU (Other)	mmlu_management
48	MMLU	MMLU (Humanities)	mmlu_prehistory
49	MMLU	MMLU (STEM)	mmlu_college_biology
50	MMLU	MMLU (Humanities)	mmlu_moral_disputes
51	MMLU	MMLU (STEM)	mmlu_high_school_statistics
52	MMLU	MMLU (Soc. sci.)	mmlu_us_foreign_policy
53	MMLU	MMLU (Other)	mmlu_human_aging
54	MMLU	MMLU (STEM)	mmlu_college_chemistry
55	MMLU	MMLU (Other)	mmlu_virology
56	MMLU	MMLU (Soc. sci.)	mmlu_public_relations
57	MMLU	MMLU (STEM)	mmlu_conceptual_physics
58	MMLU	MMLU (Soc. sci.)	mmlu_high_school_macro_economics
59	MMLU	MMLU (Humanities)	mmlu_international_law
60	MMLU	MMLU (Humanities)	mmlu_philosophy
61	MMLU	MMLU (Humanities)	mmlu_formal_logic
62	PiQA	None	piqa
63	SciQ	None	sciq
64	Winogrande	None	winogrande
65	WSC	None	wsc