

# Beyond Scaling: Measuring and Predicting the Upper Bound of Knowledge Retention in Language Model Pre-Training

Changhao Jiang<sup>\*1</sup>, Ming Zhang<sup>\*1</sup>, Yifei Cao<sup>\*1</sup>,  
Junjie Ye<sup>1</sup>, Xiaoran Fan<sup>1</sup>, Shihan Dou<sup>1</sup>, Zhiheng Xi<sup>1</sup>, Jiajun Sun<sup>1</sup>, Yi Dong<sup>1</sup>,  
Yujiong Shen<sup>1</sup>, Jingqi Tong<sup>1</sup>, Baoyu Fan<sup>4</sup>, Tao Gui<sup>1</sup>, Qi Zhang<sup>†123</sup>, Xuanjing Huang<sup>1</sup>

<sup>1</sup>Fudan University <sup>2</sup>Institute of Trustworthy Embodied AI, Fudan University

<sup>3</sup>Shanghai Key Laboratory of Multimodal Embodied AI

<sup>4</sup>IEIT Systems Co. Ltd., Shenzhen, China

chjiang25@m.fudan.edu.cn, qz@fudan.edu.cn

## Abstract

The GPT-4 technical report suggests that downstream performance can be predicted from pre-training signals, but offers little methodological detail on how to quantify this. This work address this gap by modeling knowledge retention, the capacity of a pre-trained language model to memorize factual information from its corpus, and introduce a principled method to estimate it prior to training. We propose Size-dependent Mutual Information (SMI), an information-theoretic predictor that integrates knowledge frequency, knowledge specificity, and model size to forecast closed-book question answering (QA) accuracy. SMI is validated through large-scale document retrieval over the disclosed pre-training corpora of 21 public and 3 custom models, combined with a robust multi-template QA evaluation. Experiments show that SMI significantly outperforms repetition-based baselines and achieves  $R^2 > 0.7$  in predicting QA accuracy for models above 1B parameters, without additional training. The analysis further reveals diminishing returns from scaling data and model size and provides evidence for an intrinsic upper bound on knowledge retention achievable by pre-training alone, motivating retrieval and other augmentation strategies.

## 1 Introduction

The technical report on GPT-4 states, "We registered predictions for GPT-4's performance on HumanEval before training completed, using only information available prior to training" (OpenAI, 2023). Understanding and predicting the generalization ability of large language models (LLMs) has become a key research question in the study of scaling laws (Kaplan et al., 2020). Although prior work has established empirical relationships

between loss, data size, and model parameters, these results primarily offer coarse-grained predictions of overall performance and do not address what specific factual knowledge is retained from pre-training corpora. Recent studies have shown that LLMs can memorize and regurgitate training data (Carlini et al., 2021; Biderman et al., 2023a), but these findings are largely descriptive and do not provide a systematic way to predict retention for new model–corpus pairs. Developing principled approaches for this purpose is crucial both for improving our theoretical understanding of model training dynamics and for guiding practical decisions on data curation and compute allocation (Hoffmann et al., 2022; Alabdulmohsin et al., 2022; Ruan et al.).

This work addresses this gap by modeling knowledge retention, the capacity of a pre-trained language model to memorize factual information from its corpus, and introduces a principled method to reliably estimate it prior to training. Rather than focusing only on aggregate loss curves or perplexity, we take a more fine-grained and systematic view: we ask, given a particular fact in the pre-training corpus, what is the probability that a model will be able to accurately reproduce it in a closed-book setting after training (Petroni et al., 2019)? We term this capacity knowledge retention, and formalize it as a quantitative property of a model–corpus pair. This framing allows us to reason about model memory at the level of individual facts and categories, and to make meaningful predictions before committing to the cost of full-scale training.

We propose Size-dependent Mutual Information (SMI), an information-theoretic predictor that integrates knowledge frequency, knowledge specificity, and model size to forecast closed-book question answering (QA) accuracy (Roberts et al., 2020). Since pre-trained language models are optimized only for next-token prediction rather than explicit knowledge retrieval, closed-book QA provides a

<sup>\*</sup>Equal contribution

<sup>†</sup>Corresponding author

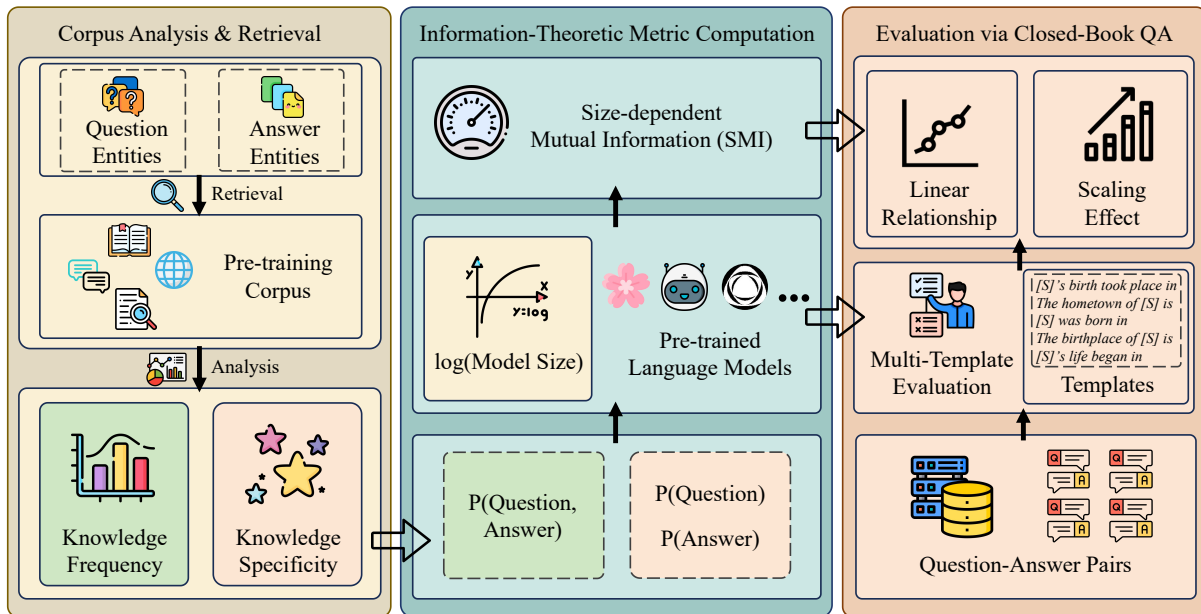


Figure 1: Overview of our method. The framework consists of three stages: (1) Corpus Analysis and Retrieval, where subject and object (question and answer) entities are extracted from pre-training data and their frequencies and co-occurrences are computed; (2) Information-Theoretic Metric Computation, where knowledge frequency, knowledge specificity, and model size are integrated into the size-dependent mutual information (SMI) metric; and (3) Evaluation via Closed-Book QA, where pre-trained language models are assessed using multi-template evaluation, and SMI is compared with empirical QA accuracy to analyze linear relationships and scaling effects.

natural and challenging probe for assessing factual recall. Our formulation builds on the intuition that both the frequency of a fact in the corpus and its rarity influence its likelihood of being learned, and that these effects are jointly influenced by model size. Concretely, SMI captures the mutual information between knowledge subjects and objects (or questions and answers), combining these statistics with model size. The resulting predictor is simple to compute from pre-training data and scales gracefully with corpus size, making it suitable for practical use in large-scale settings.

SMI is validated through large-scale document retrieval over the disclosed pre-training corpora of 21 public and 3 custom models, combined with a robust multi-template QA evaluation. Our experimental setup reconstructs the exposure of facts in the training data, then evaluates whether models can correctly answer corresponding questions in a closed-book setting. We design multi-template QA prompts to reduce prompt variance and report results across a diverse set of model families, covering a range from 14M to 176B parameters.

Experiments show that SMI significantly outperforms repetition-based baselines and achieves the coefficient of determination ( $R^2$ )  $> 0.7$  in predict-

ing QA accuracy for models above 1B parameters, without additional training. Relative to heuristics based solely on repetition counts, SMI yields a stronger correlation with observed QA accuracy, particularly in the large-model regime where memorization behavior appears more stable. Moreover, we observe that SMI maintains comparable predictive performance across multiple model families and domains, indicating that it captures a generalizable aspect of factual knowledge acquisition in pre-trained language models.

The analysis further reveals diminishing returns from scaling both data and model size, suggesting that pre-training alone faces an intrinsic limit in its ability to retain factual knowledge. Recent studies have similarly observed a performance ceiling on knowledge-intensive tasks, even with substantial scaling of parameters and data (Zhang et al., 2025). These observations point to inherent limitations of purely parametric approaches: beyond a certain scale, further increases in model size or training data yield only marginal gains in closed-book performance. This motivates exploring hybrid approaches that augment parametric memory with non-parametric mechanisms, including retrieval-augmented generation, external memory modules,

or targeted fine-tuning, to mitigate the ceiling effects observed in pre-training. Overall, our contributions are threefold:

1. We formalize knowledge retention as a quantitative metric defined over model–corpus pairs, providing a new way to analyze what models actually memorize during pre-training.
2. We propose SMI, a data-driven and training-free predictor that integrates knowledge frequency, knowledge specificity, and model size to estimate closed-book QA accuracy.
3. We empirically evaluate SMI on 24 models, showing that it achieves strong predictive accuracy, characterizing diminishing returns from scaling data and model size, and suggesting the presence of an upper bound on knowledge retention achievable by pre-training alone.

## 2 Related Work

### Pre-training data and knowledge retention

The characteristics of pre-training data are crucial for knowledge retention and downstream performance in LLMs. Many studies emphasize the role of repetition frequency. For example, [Carlini et al. \(2023\)](#) identified a logarithmic relationship between repetition and memory effects, while [Chowdhery et al. \(2023\)](#) showed that sequences repeated over 500 times are completed with over 40% accuracy. [Ju et al. \(2024\)](#) found that high-frequency data can lead to “factual shortcuts” in multi-hop reasoning. [Allen-Zhu and Li \(2024b\)](#) suggested that exposing knowledge 1000 times enables a storage capacity of two bits per parameter. Additionally, [Razeghi et al. \(2022\)](#) and [Yadlowsky et al. \(2023\)](#) linked lower-order co-occurrences to improved numerical reasoning, while [McCoy et al. \(2023\)](#) associated task prevalence in pre-training with better performance on tasks like ROT13. Furthermore, [Chang et al. \(2024\)](#) revealed that simply duplicating data might not inherently enhance knowledge retention, underscoring the importance of data diversity over mere repetition, and recent surveys further highlight the role of diverse memory forms and dynamics in modern agent systems ([Hu et al., 2025](#)). Recent work also suggests that reliable evaluation of knowledge retention requires dynamically constructed and contamination-resistant benchmarks, as static test sets may overestimate memorization ability ([Zhang et al., 2025](#)).

**Model size and knowledge retention** Furthermore, the model size significantly influences the knowledge retention of LLMs. Research by OpenAI demonstrates that as the number of non-embedding parameters in a model increases, the test loss decreases following a power law pattern ([Kaplan et al., 2020](#)). [Kandpal et al. \(2023\)](#) evaluate the BLOOM model families ([Scao et al., 2022](#)) on long-tail QA tasks and find that the fact-based QA accuracy is highly linearly correlated with the logarithm of the number of model parameters. Additionally, the memory capacity of a model closely correlates with its size ([Li et al., 2024](#)). Beyond scaling laws, recent large-scale longitudinal studies further reveal potential ceilings in knowledge memorization and highlight the impact of evaluation protocols on perceived model capabilities ([Zhang et al., 2025](#)).

**Knowledge triples as an evaluation tool** Knowledge triples (subject, relation, object) serve as an effective tool for evaluating the knowledge retention of LLMs. [Petroni et al. \(2019\)](#) employ the LAMA method to evaluate the implicit knowledge of models like BERT, demonstrating the value of knowledge triples in reasoning tasks. Subsequently, [He et al. \(2024\)](#); [Ju et al. \(2024\)](#) further explore the application of knowledge triples in reasoning tasks. Moreover, [Wang et al. \(2024\)](#) and [Allen-Zhu and Li \(2024a\)](#) use knowledge triples to explore LLMs’ ability to retain and extract knowledge. Complementary to triple-based probing, recent works investigate more holistic and structured evaluation paradigms, such as agentic systems for evidence-grounded novelty assessment ([Zhang et al., 2026](#)) and scaffold-aware benchmarks for measuring instruction-following behavior in complex environments ([Ding et al., 2026](#)), broadening the scope of LLM evaluation beyond static knowledge recall, with emerging efforts also emphasizing reproducibility and controlled generation settings in downstream tasks ([Jiang et al., 2026](#)).

## 3 Methodology

To predict a model’s ability to recall factual knowledge before training, we develop a framework that estimates closed-book QA accuracy directly from pre-training data statistics and model size. We first revisit prior work and evaluate a simple repetition-based co-occurrence metric as a baseline. Building on information-theoretic principles, we then introduce the Mutual Information (MI) metric, which

accounts for both the frequency and specificity of knowledge items and yields improved predictive performance. Finally, we extend MI to SMI, which incorporates model size as a scaling factor, providing a more accurate and scalable predictor of knowledge retention.

### 3.1 Starting from the Co-occurrence Metric

**Knowledge triples** In QA tasks, entity relationships can be represented as knowledge triples, denoted by tuples  $t = (s, r, o)$ , where  $s$  is the subject,  $r$  is the relation, and  $o$  is the object (Wang et al., 2024). For a given LLM, denoted as  $F$ , we define  $F$  as mastering the knowledge triple  $t$ , thus demonstrating competence in the corresponding QA task, as follows:

$$F : \mathbf{Q} \rightarrow \mathbf{O}, \quad F(q_{s,r}) = o \quad (1)$$

Here,  $q_{s,r}$  represents the combined representation of the subject  $s$  and the relation  $r$ , while  $o$  denotes the representation of the object. For instance, if  $s = \text{Apple}$  and  $r = \text{headquarter}$ , then  $q_{s,r}$  can be expressed as statements like “The headquarters of Apple is in...”, “Apple is headquartered in...”, or “Apple’s head office is based in...”. Similarly,  $o$  corresponds to object representation, such as “Cupertino”, “Cupertino, California”, or “Cupertino city”. The sum of all  $q_{s,r}$  representations is denoted as  $\mathbf{Q}$ , and the sum of all  $o$  representations is denoted as  $\mathbf{O}$ .

**Co-occurrence metric baseline** Prior work indicates that the more frequently knowledge appears in pre-training data, the better the model learns (Kandpal et al., 2023; Mallen et al., 2023; Allen-Zhu and Li, 2024b). Specifically, linking the entities in the questions and answers of a QA task to the pre-training data and counting the number of documents that jointly link the question and answer entities is a feasible approach (Kandpal et al., 2023). There is a linear relationship between QA accuracy and the logarithm of the number of documents (Carlini et al., 2023; Kandpal et al., 2023).

To quantify such co-occurrence effects, let  $N$  denote the total number of documents in the pre-training corpus. We define  $P(s)$  as the proportion of documents containing the subject entity  $s$ ,  $P(o)$  as the proportion of documents containing the object entity  $o$ , and  $P(s, o)$  as the proportion of documents containing both  $s$  and  $o$ . Based on this formulation, we define the co-occurrence metric as

$\log(N \cdot P(s, o))$ , which corresponds to the logarithm of the number of documents in which the  $s$  and  $o$  co-occur. We use this metric as a baseline for predicting closed-book QA accuracy.

**Multi-template evaluation method** Evaluating the knowledge retention of LLMs presents challenges due to the diverse range of queries formed by the subject and relation of knowledge triples. Inspired by the evaluation method for fine-tuned LLMs (Ye et al., 2024), we develop a robust multi-template evaluation method to assess the accuracy of pre-trained models in QA tasks.

Specifically, by constructing synonymous sentences, we generate a substantial number of query templates that are semantically similar but vary in form. For each type of knowledge triple, we produce 20 templates that differ in length and structure. Each template  $q$  represents a specific instance within the query set  $\mathbf{Q}$ , and these 20 selected  $q$  templates approximate the entirety of  $\mathbf{Q}$ .

**Linear relationship** On the one hand, we perform document-level retrieval on the pre-training data and calculate the co-occurrence metric. For consistency with information-theoretic conventions (Shannon, 1948a), all logarithms in this paper use base 2. Since specifying the subject and object typically determines the relation, we omit the relation component for simplicity.

On the other hand, by employing the multi-template evaluation approach, we measure the model’s accuracy on QA tasks. To examine the relationship between the co-occurrence metric and QA accuracy, we fit a linear regression model and compute the  $R^2$  along with the Mean Squared Error (MSE). These metrics allow us to evaluate both the strength and the stability of the linear relationship.

### 3.2 Considering the Knowledge Specificity to Obtain the MI Metric

**Knowledge specificity** The co-occurrence metric does not account for the specificity of knowledge. For a knowledge triple  $t = (s, r, o)$ , when the subject entity  $s$  and the object entity  $o$  appear frequently in the pre-training data, corresponding to higher values of  $P(s)$  and  $P(o)$ , the triple tends to be less specific. In such cases, the association between  $s$  and  $o$  is less distinctive, which makes it more difficult for the model to reliably retain the corresponding factual knowledge. Conversely, when  $s$  and  $o$  appear infrequently in the pre-training data, corresponding to lower values of

$P(s)$  and  $P(o)$ , the knowledge triple is more specific, which allows the model to more easily retain the corresponding fact. For example, the address of NVIDIA’s headquarters is easier to remember than that of Apple’s headquarters, because the word "Apple" occurs more frequently across diverse contexts. MI in information theory considers both the co-occurrence of two variables and their individual specificities (Shannon, 1948b). Therefore, we aim to use MI to improve the co-occurrence metric.

**MI metric** Unlike the MI in information theory, which is defined over the full joint distribution of random variables and accounts for both occurrence and non-occurrence events, the MI metric is computed only over observed occurrences of the subject and object in the corpus. Given pre-training data  $\mathbf{P}$  consisting of  $\mathbf{N}$  paragraphs, a knowledge triple  $(s, r, o)$ , and corresponding  $P(s)$ ,  $P(o)$ , and  $P(s, o)$ , the MI metric between  $s$  and  $o$  is defined as:

$$I(s, o) = P(s, o) \log \left( \frac{P(s, o)}{P(s)P(o)} \right), \quad (2)$$

$$MI(s, o) = \text{Norm}(\log(I(s, o))). \quad (3)$$

In the above formula,  $P(s)$  and  $P(o)$  penalize the number of documents containing the subject and object, respectively, to emphasize the specificity of  $s$  and  $o$ . The term  $I(s, o)$  represents the amount of shared information between  $s$  and  $o$ , and it can also indicate the reduction in the amount of information about  $o$  given  $s$ . During inference with LLMs, this amount of information refers to the probability of the model generating the object given the subject. Considering the uneven distribution of  $I(s, o)$  (see Figure 2), we apply a logarithmic transformation followed by normalization to map its values to the range  $[0, 1]$ , resulting in the MI metric. In rare cases where  $I(s, o)$  becomes zero or negative due to data sparsity or noise, such instances are regarded as anomalies and excluded from further analysis.

### 3.3 Considering the Model Size to Obtain the Size-dependent MI metric

**Model size** The MI metric does not account for the impact of model size. Existing research indicates that the capabilities of LLMs increase with model size, but the returns diminish as the size becomes larger (Kaplan et al., 2020). In long-tail QA tasks, the QA accuracy is highly linearly correlated

with the logarithm of the number of model parameters (Kandpal et al., 2023). Therefore, we propose to extend the MI metric by explicitly incorporating the parameter count  $\Phi$  as a measure of model size.

The method of integrating model size is not unique. A guiding principle is that, for the same data, a higher  $\Phi$  should correspond to a higher value of the new metric, reflecting improved QA performance. However, according to scaling laws (Kaplan et al., 2020), increasing model size exhibits diminishing marginal returns; therefore, the model scale is expressed in logarithmic form. Importantly, as  $\Phi$  approaches infinity, the new metric should converge to the MI metric, indicating that the model’s memory capacity is sufficient to fully capture the relevant knowledge.

**Size-dependent MI metric** Since linear transformations do not alter the strength of linear relationships, an exponential formulation offers a simple yet effective way to incorporate model size while capturing non-linear effects. Based on these principles, we propose the SMI metric, an exponential function with the MI metric as the base and  $1 + \frac{1}{\log \Phi}$  as the exponent.

$$SMI(s, o, \Phi) = MI(s, o)^{1 + \frac{1}{\log \Phi}}. \quad (4)$$

In the above formula, the MI metric quantifies the information of knowledge triples in the pre-training data, while the model size  $\Phi$  implies the memory capacity of LLMs. These two factors jointly determine the model’s knowledge retention. Since the MI metric is always less than 1, the factor  $1 + \frac{1}{\log \Phi}$  decreases as  $\Phi$  increases, leading to a larger SMI value, which indicates stronger knowledge retention.

## 4 Experiments

In the Experimental Setup, we first curate and enhance an open-source Wikipedia knowledge triple dataset to make it suitable for inference with pre-trained models. We then introduce five large-scale open-source pre-training corpora and perform retrieval-based analysis on them, followed by a detailed description of five model families and our evaluation protocols. In the Experimental Results, we evaluate 21 open-source and 3 in-house pre-trained models, providing a comparative analysis across metrics.

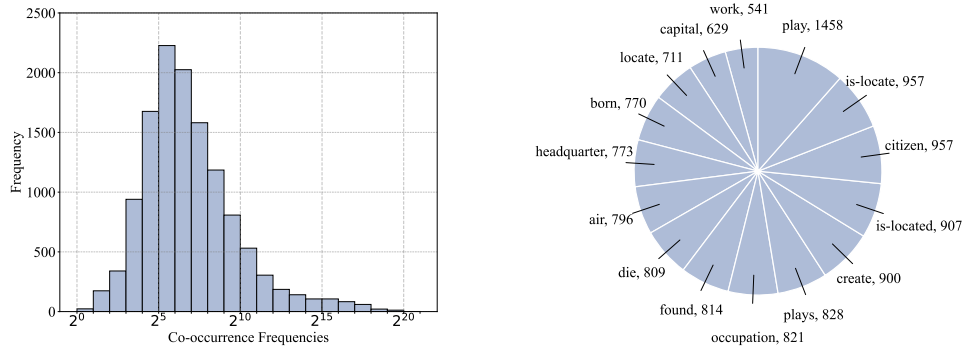


Figure 2: Evaluation data statistics. Left: Co-occurrence frequencies of QA entity pairs in The Pile corpus. Right: Distribution of 12,671 QA tasks across 15 relation types.

## 4.1 Experimental Setup

**Evaluation data** We make use of the high-quality knowledge triple dataset ParaRel, which is designed to evaluate whether language models align with factual knowledge, a task similar to ours (Elazar et al., 2021). We select 15 categories of knowledge triple relations with a larger quantity of data and clearer semantics, totaling 12,671 knowledge triples, as shown in Figure 2. For each category of knowledge triple relations, we create 20 synonymous query templates to construct the QA task evaluation dataset.

**Pre-training corpora** We conduct a document-level retrieval analysis of the following pre-training datasets, selected based on their usage in the LLMs under our consideration:

- **The Pile:** This is an open-source, large-scale text dataset specifically designed for training LLMs. It encompasses a variety of sources, including webpages, academic papers, code, and books, with an emphasis on diversity and high quality (Gao et al., 2021).
- **ROOTS (En):** This dataset is the English subset of ROOTS, a comprehensive multilingual corpus constructed by the BIGScience project. It is used for training the BLOOM models and emphasizes linguistic diversity and data transparency (Laurençon et al., 2022; Scao et al., 2022).
- **SlimPajama:** Constructed by Cerebras, this open-source dataset is a large-scale text corpus aimed at providing high-quality training materials for LLMs (Soboleva et al., 2023).
- **RefinedWeb:** Developed by the Abu Dhabi Technology Innovation Institute (TII), this large-scale English web dataset is based on CommonCrawl

and undergoes rigorous content extraction, filtering, and deduplication processes. It is intended to serve as high-quality pre-training material for LLMs such as Falcon LLM (Penedo et al., 2023).

- **Wikipedia:** This structured text corpus contains Wikipedia articles and is extensively used for language model training and evaluation due to its high-quality content, comprehensive coverage, and semantic consistency (Foundation).

**Pre-training data retrieval** We process the pre-training corpus at the document level. For each knowledge triple, we retrieve the frequency of occurrences of the subject, the object, and their co-occurrences. Using the retrieval results from all documents, we calculate the co-occurrence metric and MI metric for each knowledge triple. Subsequently, we compute the SMI metric based on the MI metric and the model size. The retrieval process runs on 20 CPUs for 15 days.

**Pre-trained models** We conduct experiments on 21 fully open-source pre-trained Transformer decoder-only pre-trained LLMs (Vaswani et al., 2017). These models not only make the model weights available but also provide complete pre-training data. They are primarily categorized into the following families:

- **BLOOM:** Created by BigScience, these open-source multilingual LLMs are trained on ROOTS and supports 46 natural languages and 13 programming languages, with parameter sizes ranging from 560 million to 176 billion (Scao et al., 2022).
- **GPT-Neo:** This family includes the GPT-Neo, GPT-J, and GPT-NeoX models, ranging from 125 million to 20 billion parameters, all trained on The Pile (Black et al., 2021, 2022; Wang, 2021).

Table 1:  $R^2$  and MSE ( $\times 10^{-2}$ ) results of the CO-OCCUR, MI, and SMI metrics across different models, ranging from 14M to 176B parameters. Across all models, MI consistently outperforms CO-OCCUR, and SMI further improves upon MI.

Model	$R^2$ ( $\uparrow$ higher is better)			MSE ( $\times 10^{-2}$ , $\downarrow$ lower is better)		
	CO-OCCUR	MI	SMI	CO-OCCUR	MI	SMI
<b>BLOOM</b>						
bloom-560m	0.384	0.589	<b>0.591</b>	8.566	7.729	<b>7.723</b>
bloom-1b1	0.577	0.715	<b>0.719</b>	9.123	7.950	<b>7.939</b>
bloom-1b7	0.584	0.740	<b>0.743</b>	9.389	8.083	<b>8.072</b>
bloom-3b	0.705	0.782	<b>0.786</b>	9.314	7.863	<b>7.849</b>
bloom-7b1	0.676	0.804	<b>0.807</b>	9.943	8.309	<b>8.296</b>
bloom	0.644	0.764	<b>0.766</b>	7.327	6.029	<b>6.026</b>
<b>GPT-Neo</b>						
gpt-neo-125m	0.349	0.531	<b>0.535</b>	7.004	7.107	<b>7.100</b>
gpt-neo-1.3B	0.671	0.810	<b>0.814</b>	8.233	7.444	<b>7.427</b>
gpt-neo-2.7B	0.672	0.826	<b>0.829</b>	8.168	7.320	<b>7.307</b>
gpt-j-6b	0.631	0.862	<b>0.865</b>	9.026	7.720	<b>7.710</b>
gpt-neox-20b	0.593	0.841	<b>0.843</b>	8.557	7.476	<b>7.469</b>
<b>Pythia</b>						
pythia-14m	0.240	0.375	<b>0.378</b>	6.855	7.183	<b>7.177</b>
pythia-31m	0.287	0.381	<b>0.384</b>	6.905	7.202	<b>7.196</b>
pythia-70m	0.355	0.430	<b>0.434</b>	6.749	7.051	<b>7.044</b>
pythia-160m	0.454	0.522	<b>0.526</b>	6.877	6.969	<b>6.961</b>
pythia-410m	0.524	0.701	<b>0.705</b>	7.439	7.035	<b>7.024</b>
pythia-1.4b	0.574	0.799	<b>0.803</b>	8.254	7.325	<b>7.312</b>
pythia-2.8b	0.631	0.845	<b>0.848</b>	8.593	7.505	<b>7.491</b>
pythia-6.9b	0.637	0.855	<b>0.859</b>	8.538	7.373	<b>7.360</b>
pythia-12b	0.602	0.853	<b>0.856</b>	8.741	7.514	<b>7.503</b>
<b>TinyLlama</b>						
TinyLlama-1.1B	0.641	0.762	<b>0.767</b>	7.802	7.081	<b>7.065</b>
<b>Ours</b>						
ours-1.6b	0.710	0.778	<b>0.782</b>	8.563	7.712	<b>7.694</b>
ours-7b	0.784	0.837	<b>0.839</b>	8.982	8.028	<b>8.016</b>
ours-13b	0.661	0.807	<b>0.810</b>	8.584	7.674	<b>7.665</b>

- **Pythia:** Developed by EleutherAI, this fully open-source series of LLMs ranges from 14 million to 12 billion parameters. These models use a uniform data and training process, facilitating research that is both replicable and interpretable (Biderman et al., 2023b).
- **TinyLlama:** Originating from the Singapore University of Technology and Design (SUTD), this open-source LLM is trained on SlimPajama and Starcoderdata. It has 1.1 billion parameters and is designed with the Llama 2 architecture and tokenizer to provide efficient and lightweight LLM solutions (Zhang et al., 2024; Soboleva et al., 2023; Li et al., 2023). The specific model we utilize is the TinyLlama-1.1B-intermediate-step-480k-1T, which is trained for approximately one epoch.

Additionally, we pre-train three models, named ours-1.6b, ours-7b, and ours-13b, based on Re-

finerWeb and Wikipedia as a supplement to the aforementioned models.

**Evaluation setting** We evaluate these pre-trained models on the designated QA evaluation dataset. During inference, we configure the generation parameters with a temperature of 0.7, a maximum of 32 new tokens, and bfloat16 precision. A model response is considered correct if the expected object entity appears anywhere within the generated text. To obtain stable and reliable estimates of performance, we repeat the inference procedure 20 times for each of the 20 queries in the QA task, resulting in a total of 400 responses per model. The final accuracy is computed as the mean over all 400 responses. The full evaluation runs for approximately five days on 8 H100 GPUs.

The retrieval results from the pre-training corpus show that the co-occurrence frequency distribution of question and answer entities is highly skewed (see Figure 2). Most QA pairs co-occur

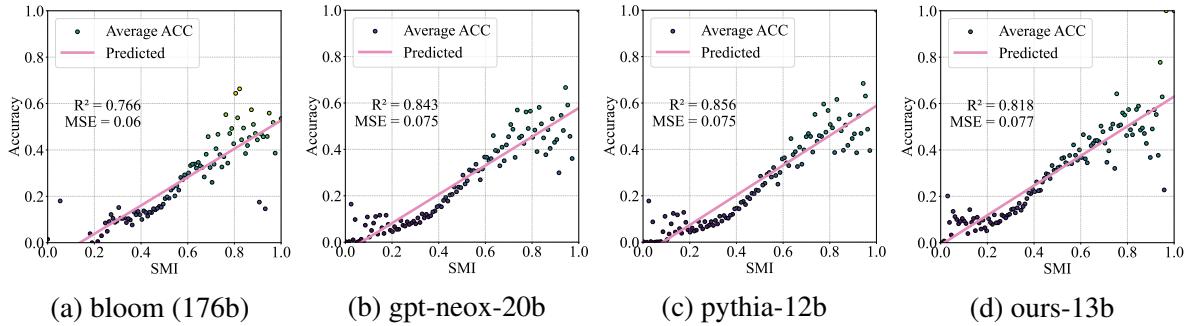


Figure 3: Relationship between the SMI metric and accuracy on QA tasks for different model series: (a) BLOOM (176B), (b) GPT-Neo (20B), (c) Pythia (12B), and (d) Ours (13B).




Figure 4: Relationship between average QA accuracy on the full evaluation dataset and model size across five different model families.

between 10 and 1000 times, while only a small portion appears at very low or very high frequencies. To mitigate this imbalance and enable meaningful evaluation across different frequency ranges, we group QA data with similar metric values. For the co-occurrence and MI metrics, we compute average accuracy by binning values into intervals of 0.2 on the unnormalized scale. The SMI metric is then derived from the MI metric after binning.

Finally, we perform linear regression on the metrics and accuracy, and calculate the  $R^2$  to assess the strength of the linear relationship. Simultaneously, to evaluate the prediction error of the regression equation in specific QA tasks, we compute the MSE for each QA task and take the average.

## 4.2 Experimental Results

**Comparison of SMI, MI, and CO-OCCUR** Table 1 presents the  $R^2$  and MSE results of CO-OCCUR, MI, and SMI across 24 models, with sizes ranging from 14M to 176B parameters. Overall, MI consistently improves over the co-occurrence baseline. For example, on gpt-neo-2.7B,  $R^2$  increases from 0.672 (CO-OCCUR) to 0.829 (MI), while MSE decreases from 8.168 to 7.307. This

suggests that jointly modeling frequency and specificity provides more informative signals for predicting closed-book QA performance. Incorporating model size into the metric (SMI) yields further gains across all models, though the magnitude of improvement is typically moderate. For instance, on pythia-12b,  $R^2$  rises from 0.853 (MI) to 0.856 (SMI), and on bloom, from 0.764 to 0.766. These results suggest that incorporating model size into the metric provides a more principled account of the variance in observed QA accuracy. Additional analysis and validation of SMI are detailed in Appendix B.

**Diminishing returns from data scaling** Despite the improved predictive power of SMI, we find that knowledge retention is limited by the nature of pre-training data. As shown in Figure 3, even when SMI reaches its maximum value of 1, the closed-book QA accuracy does not exceed moderate levels. Specifically, the empirical upper bounds of accuracy for bloom, gpt-neox-20b, pythia-12b, and ours-13b are approximately 0.527, 0.578, 0.588, and 0.630, respectively. These results indicate that simply scaling the data cannot indefinitely improve factual recall. Additional analyses are provided in the Appendix B.4.

**Diminishing returns from model scaling** We further analyze the impact of model size on knowledge retention. Figure 4 shows that the average accuracy on the full evaluation set scales approximately linearly with the logarithm of model parameters, which is consistent with prior scaling-law observations. However, the growth rate becomes increasingly shallow for very large models: even at 176B parameters, the accuracy remains below 0.2. This finding suggests that model scaling alone is insufficient for achieving high factual coverage

and further motivates the exploration of complementary approaches, such as retrieval-augmented generation or external memory modules.

## 5 Conclusion

In this work, we introduce a framework for quantifying knowledge retention in pre-trained language models via SMI, an information-theoretic predictor that integrates knowledge frequency, specificity, and model size. Our large-scale evaluation across 24 models shows that SMI achieves high predictive correlation ( $R^2 > 0.7$  for models above 1B parameters) and remains robust across model families and domains.

Beyond prediction, our analysis reveals diminishing returns from scaling data and model size, suggesting an inherent upper bound on parametric factual storage during pre-training. Furthermore, SMI serves as a diagnostic tool, indicating that while increasing co-occurrence frequency, refining data specificity, and scaling model size can enhance retention, achieving higher factual recall may necessitate complementary strategies such as retrieval-augmented generation, external memory modules, or targeted fine-tuning.

We hope this work provides a foundation for optimizing corpus composition and model design toward more effective factual knowledge retention.

## 6 Limitations

While our study provides an information-theoretic framework for predicting knowledge retention, several limitations remain.

First, our analysis is restricted to closed-book QA tasks over short factual triplets, as instantiated by the ParaRel dataset (Elazar et al., 2021). Accordingly, our notion of knowledge retention focuses on atomic, declarative facts and does not extend to multi-hop reasoning, non-atomic or unstructured knowledge, or procedural knowledge, all of which may involve substantially different memorization and retrieval dynamics.

Second, we rely on publicly available corpora and approximate document matching, which may introduce noise when estimating fact frequency and specificity.

Third, our experiments focus on models with up to tens of billions of parameters, leaving open the question of whether the observed trends hold for frontier-scale models with hundreds of billions of parameters.

Finally, our approach abstracts away training dynamics such as optimization schedules or curriculum learning, whose interactions with knowledge retention remain an important direction for future work.

## Acknowledgments

The authors wish to thank the anonymous reviewers for their helpful comments. This work was partially funded by Henan Province Major Industrial “Challenge-Based Innovation” (No. 251000210300), National Natural Science Foundation of China (No.62476061, 62376061, 62576106).

## References

- Ibrahim M Alabdulmohsin, Behnam Neyshabur, and Xiaohua Zhai. 2022. Revisiting neural scaling laws in language and vision. *Advances in Neural Information Processing Systems*, 35:22300–22312.
- Zeyuan Allen-Zhu and Yuanzhi Li. 2024a. [Physics of language models: Part 3.1, knowledge storage and extraction](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Zeyuan Allen-Zhu and Yuanzhi Li. 2024b. [Physics of language models: Part 3.3, knowledge capacity scaling laws](#). *CoRR*, abs/2404.05405.
- Stella Biderman, Usvsn Prashanth, Lintang Sutawika, Hailey Schoelkopf, Quentin Anthony, Shivanshu Purohit, and Edward Raff. 2023a. Emergent and predictable memorization in large language models. *Advances in Neural Information Processing Systems*, 36:28072–28090.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. 2023b. [Pythia: A suite for analyzing large language models across training and scaling](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 2397–2430. PMLR.
- Sid Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, USVSN Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. [Gpt-neox-20b: An open-source autoregressive language model](#). *CoRR*, abs/2204.06745.

- Sid Black, Gao Leo, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow](#). If you use this software, please cite it using these metadata.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2023. [Quantifying memorization across neural language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, and 1 others. 2021. Extracting training data from large language models. In *30th USENIX security symposium (USENIX Security 21)*, pages 2633–2650.
- Hoyeon Chang, Jinho Park, Seonghyeon Ye, Sohee Yang, Youngkyung Seo, Du-Seong Chang, and Minjoon Seo. 2024. [How do large language models acquire factual knowledge during pretraining?](#) In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, and 48 others. 2023. [Palm: Scaling language modeling with pathways](#). *J. Mach. Learn. Res.*, 24:240:1–240:113.
- Deming Ding, Shichun Liu, Enhui Yang, Jiahang Lin, Ziyang Chen, Shihan Dou, Honglin Guo, Weiyu Cheng, Pengyu Zhao, Chengjun Xiao, and 1 others. 2026. Octobench: Benchmarking scaffold-aware instruction following in repository-grounded agentic coding. *arXiv preprint arXiv:2601.10343*.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard H. Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. [Measuring and improving consistency in pretrained language models](#). *Trans. Assoc. Comput. Linguistics*, 9:1012–1031.
- Wikimedia Foundation. [Wikimedia downloads](#).
- 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. 2021. [The pile: An 800gb dataset of diverse text for language modeling](#). *CoRR*, abs/2101.00027.
- Qiyuan He, Yizhong Wang, and Wenya Wang. 2024. [Can language models act as knowledge bases at scale?](#) *CoRR*, abs/2402.14273.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, and 1 others. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Yuyang Hu, Shichun Liu, Yanwei Yue, Guibin Zhang, Boyang Liu, Fangyi Zhu, Jiahang Lin, Honglin Guo, Shihan Dou, Zhiheng Xi, and 1 others. 2025. Memory in the age of ai agents. *arXiv preprint arXiv:2512.13564*.
- Changhao Jiang, Jiahao Chen, Zhenghao Xiang, Zhixiong Yang, Hanchen Wang, Jiabao Zhuang, Xinmeng Che, Jiajun Sun, Hui Li, Yifei Cao, and 1 others. 2026. Muse: Towards reproducible long-form song generation with fine-grained style control. *arXiv preprint arXiv:2601.03973*.
- Tianjie Ju, Yijin Chen, Xinwei Yuan, Zhuosheng Zhang, Wei Du, Yubin Zheng, and Gongshen Liu. 2024. [Investigating multi-hop factual shortcuts in knowledge editing of large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 8987–9001. Association for Computational Linguistics.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. [Large language models struggle to learn long-tail knowledge](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 15696–15707. PMLR.
- 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](#). *CoRR*, abs/2001.08361.
- Hugo Laurençon, Lucile Saulnier, Thomas Wang, Christopher Akiki, Albert Villanova del Moral, Teven Le Scao, Leandro von Werra, Chenghao Mou, Eduardo González Ponferrada, Huu Nguyen, Jörg Froberg, Mario Sasko, Quentin Lhoest, Angelina McMillan-Major, Gérard Dupont, Stella Biderman, Anna Rogers, Loubna Ben Allal, Francesco De Toni, and 35 others. 2022. [The bigscience ROOTS corpus: A 1.6tb composite multilingual dataset](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Bo Li, Qinghua Zhao, and Lijie Wen. 2024. [ROME: memorization insights from text, probability and hidden state in large language models](#). *CoRR*, abs/2403.00510.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc

- Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliachko, and 48 others. 2023. [Starcoder: may the source be with you!](#) *Trans. Mach. Learn. Res.*, 2023.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 9802–9822. Association for Computational Linguistics.
- R. Thomas McCoy, Shunyu Yao, Dan Friedman, Matthew Hardy, and Thomas L. Griffiths. 2023. [Members of autoregression: Understanding large language models through the problem they are trained to solve](#). *CoRR*, abs/2309.13638.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Hamza Alobeidli, Alessandro Cappelli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. [The refinedweb dataset for falcon LLM: outperforming curated corpora with web data only](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 2463–2473. Association for Computational Linguistics.
- Yasaman Razeghi, Robert L. Logan IV, Matt Gardner, and Sameer Singh. 2022. [Impact of pretraining term frequencies on few-shot numerical reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 840–854. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 5418–5426. Association for Computational Linguistics.
- Yangjun Ruan, Chris J Maddison, and Tatsunori Hashimoto. [Observational scaling laws and the predictability of language model performance, 2024](#). URL <https://arxiv.org/abs/2405.10938>.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilic, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, and 30 others. 2022. [BLOOM: A 176b-parameter open-access multilingual language model](#). *CoRR*, abs/2211.05100.
- Claude E. Shannon. 1948a. [A mathematical theory of communication](#). *Bell Syst. Tech. J.*, 27(4):623–656.
- Claude E. Shannon. 1948b. [A mathematical theory of communication](#). *Bell Syst. Tech. J.*, 27(3):379–423.
- Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. 2023. [SlimPajama: A 627B token cleaned and deduplicated version of RedPajama](#). <https://www.cerebras.net/blog/slimpajama-a-627b-token-cleaned-and-deduplicated-version-of-redpajama>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Ben Wang. 2021. [Mesh-Transformer-JAX: Model-Parallel Implementation of Transformer Language Model with JAX](#). <https://github.com/kingoflolz/mesh-transformer-jax>.
- Mengru Wang, Yunzhi Yao, Ziwen Xu, Shuofei Qiao, Shumin Deng, Peng Wang, Xiang Chen, Jia-Chen Gu, Yong Jiang, Pengjun Xie, Fei Huang, Huajun Chen, and Ningyu Zhang. 2024. [Knowledge mechanisms in large language models: A survey and perspective](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12-16, 2024*, pages 7097–7135. Association for Computational Linguistics.
- Steve Yadlowsky, Lyric Doshi, and Nilesh Tripuraneni. 2023. [Pretraining data mixtures enable narrow model selection capabilities in transformer models](#). *CoRR*, abs/2311.00871.
- Junjie Ye, Yuming Yang, Qi Zhang, Tao Gui, Xuanjing Huang, Peng Wang, Zhongchao Shi, and Jianping Fan. 2024. [Empirical insights on fine-tuning large language models for question-answering](#). *CoRR*, abs/2409.15825.
- Ming Zhang, Yujiong Shen, Jingyi Deng, Yuhui Wang, Huayu Sha, Kexin Tan, Qiyuan Peng, Yue Zhang, Junzhe Wang, Shichun Liu, and 1 others. 2025. [Llmeval-fair: A large-scale longitudinal study on](#)

robust and fair evaluation of large language models. *arXiv preprint arXiv:2508.05452*.

Ming Zhang, Kexin Tan, Yueyuan Huang, Yujiong Shen, Chunchun Ma, Li Ju, Xinran Zhang, Yuhui Wang, Wenqing Jing, Jingyi Deng, and 1 others. 2026. Opennovelty: An llm-powered agentic system for verifiable scholarly novelty assessment. *arXiv preprint arXiv:2601.01576*.

Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. 2024. *Tinyllama: An open-source small language model*. *CoRR*, abs/2401.02385.

Table 2: Robustness analysis on low-frequency knowledge.  $R^2$  scores are reported for the CO-OCCUR and SMI, restricted to the bottom 20% of evaluation triples ranked by  $P(s, o)$ .

Model	$R^2$ : CO-OCCUR	$R^2$ : SMI	$P(s, o)$ Range
<b>BLOOM</b>			
bloom-560m	0.076	0.512	(0, 1e-02)
bloom-1b1	0.235	0.663	(0, 1e-02)
bloom-1b7	0.216	0.680	(0, 1e-02)
bloom-3b	0.371	0.740	(0, 1e-02)
bloom-7b1	0.291	0.733	(0, 1e-02)
bloom	0.202	0.573	(0, 1e-02)
<b>GPT-NEO</b>			
gpt-neo-125m	0.048	0.467	(0, 6e-03)
gpt-neo-1.3B	0.245	0.794	(0, 6e-03)
gpt-neo-2.7B	0.192	0.771	(0, 6e-03)
gpt-j-6b	0.045	0.727	(0, 6e-03)
gpt-neox-20b	0.022	0.693	(0, 6e-03)
<b>Pythia</b>			
pythia-14m	0.005	0.302	(0, 6e-03)
pythia-31m	0.020	0.316	(0, 6e-03)
pythia-70m	0.052	0.370	(0, 6e-03)
pythia-160m	0.104	0.463	(0, 6e-03)
pythia-410m	0.111	0.643	(0, 6e-03)
pythia-1.4b	0.095	0.735	(0, 6e-03)
pythia-2.8b	0.115	0.780	(0, 6e-03)
pythia-6.9b	0.097	0.785	(0, 6e-03)
pythia-12b	0.047	0.757	(0, 6e-03)
<b>TinyLlama</b>			
TinyLlama-1.1B	0.108	0.726	(0, 1e-02)
<b>Ours</b>			
ours-1.6b	0.290	0.747	(0, 6e-03)
ours-7b	0.246	0.672	(0, 6e-03)
ours-13b	0.075	0.601	(0, 6e-03)

## A AI Assistants in Research or Writing

In preparing this manuscript, AI assistants were employed solely to assist with refining the clarity, style, and readability of certain text segments. They were not involved in designing the study, developing or implementing the methodology, collecting or analyzing data, or generating the primary scientific contributions. All substantive research decisions, analyses, and conclusions are fully the responsibility of the authors.

## B Additional Experimental Results

### B.1 Validating the Predictive Efficacy and Robustness of SMI

To verify the reliable predictive power of the proposed Size-dependent Mutual Information (SMI) and to demonstrate that its effectiveness extends beyond simple co-occurrence frequency, we provide additional analyses showing that SMI is not merely a byproduct of subject–object frequency counts in the pre-training data.

The SMI builds on Mutual Information (MI), which depends jointly on joint frequency and marginals. Identical co-occurrence rates can yield vastly different MI values given different marginals. For example, when  $P(s, o) = 0.1$ ,  $P(s) = 0.1$ ,  $P(o) = 0.1$ , we obtain  $I(s, o) = 0.23$ . With the same  $P(s, o) = 0.1$  but  $P(s) = 0.2$ ,  $P(o) = 0.4$ , the MI drops to 0.02. This illustrates that MI, as well as the derived SMI metric, cannot be reduced to co-occurrence frequency alone, as they explicitly account for the specificity induced by the marginal distributions of the subject and object entities.

To empirically validate this distinction, we compare SMI with the pure co-occurrence baseline (CO-OCCUR), which uses only the logarithm of the subject–object co-occurrence count as its predictive signal. As reported in Table 1, SMI consistently and substantially outperforms the CO-OCCUR baseline across all models, demonstrating a marked improvement in  $R^2$  scores.

We further evaluate the robustness of SMI by focusing on knowledge triples with low co-occurrence frequency. Specifically, for each model, we sort the evaluation data by  $P(s, o)$  and restrict the analysis to the lowest 20% of examples. We then recompute  $R^2$  for both CO-OCCUR and SMI on this specific subset. As summarized in Table 2, the predictive performance of the CO-OCCUR baseline deteriorates markedly when  $P(s, o) < 0.01$ ; in contrast, SMI sustains a high and stable correlation. This divergence underscores that SMI is not a mere proxy for co-occurrence frequency, but instead captures predictive structure that goes beyond simple frequency priors, particularly in low-exposure regimes.

### B.2 Additional Baselines for Memorization Measures

Prior studies on knowledge memorization in language models generally identify two dominant factors: the frequency with which knowledge appears

Table 3: Several cases with significant deviations from the co-occurrence metric are shown here.  $N_s$  is the number of paragraphs in which the subject appears,  $N_o$  is the number of paragraphs in which the object appears, and  $N_{s,o}$  is the number of paragraphs in which they co-occur. The first two cases exhibit rare co-occurrence but demonstrate high accuracy, whereas the latter two cases show frequent co-occurrence but exhibit low accuracy.

Subject	Question	Object	$N_s$	$N_o$	$N_{s,o}$	$ACC$
Rare Co-occurrence but High Accuracy						
Csepel SC	The headquarters of Csepel SC is in	Budapest	398	748,860	172	0.655
Liis Lemsalu	Liis Lemsalu is a citizen of	Estonia	307	967,207	160	0.570
Frequent Co-occurrence but Low Accuracy						
Paris	Paris is located in	Europe	20,902,571	59,408,155	3,761,610	0.043
Ohio	The capital of Ohio is	Columbus	11,894,303	3,718,483	1,210,957	0.193

Table 4: Distribution of evaluation samples across accuracy intervals.

Model	[0, 0.2)	[0.2, 0.4)	[0.4, 0.6)	[0.6, 0.8)	[0.8, 1.0]
<b>BLOOM</b>					
bloom-560m	10108	480	118	75	590
bloom-1b1	9565	763	282	127	634
bloom-1b7	9395	874	300	136	666
bloom-3b	9096	991	435	161	688
bloom-7b1	8608	1170	578	246	769
bloom	7474	1821	1239	500	337
<b>GPT-NEO</b>					
gpt-neo-125m	11480	318	117	14	600
gpt-neo-1.3B	10448	881	373	154	673
gpt-neo-2.7B	10151	1042	460	199	677
gpt-j-6b	9012	1534	813	386	784
gpt-neox-20b	8973	1752	762	317	725
<b>Pythia</b>					
pythia-14m	11924	7	0	0	598
pythia-31m	11916	12	3	0	598
pythia-70m	11901	29	1	0	598
pythia-160m	11687	221	17	6	598
pythia-410m	11024	628	197	75	605
pythia-1.4b	10349	984	378	160	658
pythia-2.8b	9786	1292	531	207	713
pythia-6.9b	9483	1497	613	200	736
pythia-12b	9260	1579	721	224	745
<b>TinyLlama</b>					
TinyLlama-1.1B	10195	1161	430	143	704
<b>Ours</b>					
ours-1.6b	9929	1269	511	194	715
ours-7b	7713	2138	1240	618	909
ours-13b	7707	2323	1172	547	869

in the pre-training data and the model size. Beyond the co-occurrence baseline used in the main experiments, defined as the logarithm of subject-object co-occurrence count ( $\log(N \cdot P(s, o))$ ), we evaluate several alternative metrics that combine these factors in simple ways.

Specifically, we consider variants such as  $\log(N \cdot P(s, o)) + \log \Phi$  and  $\log(N \cdot P(s, o)) \cdot \log \Phi$ . Since these variants are linear transformations of  $\log(N \cdot P(s, o))$ , they do not change linear correlation results and perform identically to the original co-occurrence baseline.

Table 5: Twenty templates and their corresponding accuracy values for the **BORN** relation on the ours-13b model. [S] represents the subject in the knowledge triple.

Template	ACC
[S]’s birth took place in	0.186
[S] was brought into existence in	0.187
[S] was born into the world in	0.190
[S]’s nativity is	0.197
The place in which [S] was given birth to is	0.201
The hometown of [S] is	0.203
The birthplace of [S] is	0.212
[S] was delivered in	0.215
[S]’s cradle was in	0.219
[S]’s entry into the world happened in	0.225
[S] first saw the light in	0.228
[S]’s roots are in	0.230
[S]’s life began in	0.230
[S] was born and raised in	0.237
[S]’s origin is	0.238
[S] entered life in	0.251
[S] came into the world in	0.255
[S] was born in	0.255
[S]’s birth occurred in	0.263
[S]’s starting point was	0.282

In addition, we evaluate a baseline that directly uses raw co-occurrence frequency without logarithmic scaling ( $N \cdot P(s, o)$ ). As shown in Table 6, this variant performs substantially worse across models, indicating that the co-occurrence baseline exhibits negligible predictive power if not properly transformed. To the best of our knowledge, there are no stronger baselines in this specific setting.

Table 6: Predictive performance ( $R^2$ ) of raw co-occurrence frequency.

Model	$R^2$ : Raw Co-occurrence Frequency
<b>BLOOM</b>	
bloom-560m	0.002
bloom-1b1	0.001
bloom-1b7	0.002
bloom-3b	0.000
bloom-7b1	0.003
bloom	0.016
<b>GPT-NEO</b>	
gpt-neo-125m	0.020
gpt-neo-1.3B	0.004
gpt-neo-2.7B	0.005
gpt-j-6b	0.000
gpt-neox-20b	0.000
<b>Pythia</b>	
pythia-14m	0.039
pythia-31m	0.037
pythia-70m	0.062
pythia-160m	0.022
pythia-410m	0.000
pythia-1.4b	0.000
pythia-2.8b	0.001
pythia-6.9b	0.000
pythia-12b	0.000
<b>TinyLlama</b>	
TinyLlama-1.1B	0.000
<b>Ours</b>	
ours-1.6b	0.023
ours-7b	0.024
ours-13b	0.010

### B.3 Statistical Comparison Between MI and SMI

To assess whether the performance difference between MI and SMI observed in Table 1 is statistically significant, we conduct a paired-sample  $t$ -test on their  $R^2$  values across the 24 evaluated models.

The analysis yields a test statistic of  $t(23) = 20.00$  with a corresponding  $p$ -value of  $4.84e-16$ . The 95% confidence interval is  $[0.003, 0.004]$ . Since the  $p$ -value is substantially below the conventional significance threshold (0.05), this result indicates that the improvement of SMI over MI is statistically significant.

While the absolute difference in  $R^2$  is modest, SMI consistently outperforms MI across all evaluated models. Moreover, SMI is computed directly from MI with negligible additional computational overhead. Collectively, these properties demonstrate that SMI provides a reliable and practically valuable refinement over MI.

### B.4 Analysis in High-SMI Regions

In Figure 3, an increased dispersion in the linear correlation between accuracy and SMI is observed within high-SMI regions. This behavior is primarily attributable to the evaluation design and the intrinsic difficulty of the task, rather than a limitation of SMI itself.

For each question type, we evaluate models using 20 paraphrased prompts. Although these prompts are semantically equivalent (see Appendix E), some formulations are relatively uncommon in the pre-training corpus and therefore more challenging for language models. As a result, even large-scale models do not consistently achieve high accuracy across all prompts. As shown in Figure 4, BLOOM-176B attains an average accuracy of only approximately 19%, indicating that high absolute accuracy is not a necessary indicator of strong model performance under this evaluation setting.

In addition, the number of samples achieving high accuracy is limited. This sparsity leads to increased statistical variance in the high-accuracy regime, which in turn results in greater dispersion of data points when SMI values are high. Table 4 reports the number of samples per model within each accuracy interval. Notably, the majority of samples fall below an accuracy of 0.4, which helps explain the increased dispersion observed in high-SMI regions.

### C Licenses and Terms of Use

We use the following publicly available assets in this work:

- **ParaRel:** <https://github.com/yanaiela/pararel>, MIT License (Elazar et al., 2021).
- **The Pile:** <https://huggingface.co/datasets/EleutherAI/pile>, MIT License (Gao et al., 2021).
- **ROOTS (En):** <https://huggingface.co/bigscience-data>, cc-by-nc-sa-4.0 license (Laurençon et al., 2022).
- **SlimPajama:** <https://huggingface.co/datasets/cerebras/SlimPajama-627B>, Common Crawl Foundation Terms of Use, C4 license, MIT License, BSD License, Apache License, the\_pile\_books3 license, pg19 license, ArXiv Terms of Use, Wikipedia

License, StackExchange license on the Internet Archive (Soboleva et al., 2023).

- **RefinedWeb:** <https://huggingface.co/datasets/tiiuae/falcon-refinedweb>, ODC-By 1.0 license, CommonCrawl ToU (Penedo et al., 2023).
- **Wikipedia:** <https://huggingface.co/datasets/wikimedia/wikipedia>, GNU Free Documentation License, Creative Commons Attribution-Share-Alike 3.0 License (Foundation).
- **bloom-560m:** [bigscience/bloom-560m](https://huggingface.co/bigscience/bloom-560m), RAIL License v1.0 (Scao et al., 2022).
- **bloom-1b1:** [bigscience/bloom-1b1](https://huggingface.co/bigscience/bloom-1b1), RAIL License v1.0 (Scao et al., 2022).
- **bloom-1b7:** [bigscience/bloom-1b7](https://huggingface.co/bigscience/bloom-1b7), RAIL License v1.0 (Scao et al., 2022).
- **bloom-3b:** [bigscience/bloom-3b](https://huggingface.co/bigscience/bloom-3b), RAIL License v1.0 (Scao et al., 2022).
- **bloom-7b1:** [bigscience/bloom-7b1](https://huggingface.co/bigscience/bloom-7b1), RAIL License v1.0 (Scao et al., 2022).
- **bloom:** [bigscience/bloom](https://huggingface.co/bigscience/bloom), RAIL License v1.0 (Scao et al., 2022).
- **gpt-neo-125m:** <https://huggingface.co/EleutherAI/gpt-neo-125m>, MIT License (Black et al., 2021).
- **gpt-neo-1.3B:** <https://huggingface.co/EleutherAI/gpt-neo-1.3B>, MIT License (Black et al., 2021).
- **gpt-neo-2.7B:** <https://huggingface.co/EleutherAI/gpt-neo-2.7B>, MIT License (Black et al., 2021).
- **gpt-j-6b:** <https://huggingface.co/EleutherAI/gpt-j-6b>, Apache 2.0 license (Wang, 2021).
- **gpt-neox-20b:** <https://huggingface.co/EleutherAI/gpt-neox-20b>, Apache 2.0 license (Black et al., 2022).
- **pythia-14m:** <https://huggingface.co/EleutherAI/pythia-14m>, Apache 2.0 license (Biderman et al., 2023b).

- **pythia-31m:** <https://huggingface.co/EleutherAI/pythia-31m>, Apache 2.0 license (Biderman et al., 2023b).
- **pythia-70m:** <https://huggingface.co/EleutherAI/pythia-70m>, Apache 2.0 license (Biderman et al., 2023b).
- **pythia-160m:** <https://huggingface.co/EleutherAI/pythia-160m>, Apache 2.0 license (Biderman et al., 2023b).
- **pythia-410m:** <https://huggingface.co/EleutherAI/pythia-410m>, Apache 2.0 license (Biderman et al., 2023b).
- **pythia-1.4b:** <https://huggingface.co/EleutherAI/pythia-1.4b>, Apache 2.0 license (Biderman et al., 2023b).
- **pythia-2.8b:** <https://huggingface.co/EleutherAI/pythia-2.8b>, Apache 2.0 license (Biderman et al., 2023b).
- **pythia-6.9b:** <https://huggingface.co/EleutherAI/pythia-6.9b>, Apache 2.0 license (Biderman et al., 2023b).
- **pythia-12b:** <https://huggingface.co/EleutherAI/pythia-12b>, Apache 2.0 license (Biderman et al., 2023b).
- **TinyLlama:** <https://huggingface.co/TinyLlama/TinyLlama-1.1B-intermediate-step-480k-1T>, Apache 2.0 license (Zhang et al., 2024).

We confirm that all licenses and terms of use are respected.

## D Case Study

Table 3 highlights four counterintuitive cases in our pre-training data. From these, we observe the following: (1) In the first two cases, the subject and object rarely co-occur, suggesting a seemingly weak relationship. However, whenever the subject does appear, the object almost always appears alongside it, resulting in high accuracy. (2) In the last two cases, despite frequent co-occurrence between the subject and object, their individual occurrence frequencies are significantly higher, leading to low accuracy. For example, when prompted with “Paris is located in”, the model is more likely to respond with “France” rather than “Europe”.

These examples highlight the pivotal role of knowledge specificity in the memory of LLMs. It is not simply the frequencies of knowledge occurrence that enhance retention, but rather the specificity of the knowledge that proves more influential. This insight suggests that in downstream task training, even when domain-specific data is significantly fewer than the pre-training data, minimizing repetition with the pre-training data can increase the SMI, making it easier for LLMs to retain the relevant knowledge.

## E Templates for Knowledge Triples

We create twenty templates for each type of knowledge triple. Table 5 illustrates the templates for the **BORN** relation, along with their corresponding evaluation accuracy on the ours-13b model. As shown, the choice of template has a significant impact on the model's accuracy in answering questions. To address the instability of evaluations based on a single template, we calculate the average accuracy across all 20 templates. We employ **gpt-4o-2024-08-06** to generate query templates. The prompt we use is: **Generate 20 diverse and fluent paraphrases of the following sentence, preserving its original meaning:\n[sentence]**. Our evaluation set contains a total of 12,671 pieces of data and 300 query templates, following the MIT license. The template designs for each type of knowledge triple are as follows:

```
{
  "plays": [
    "[S] plays",
    "[S] performs",
    "[S] engages in playing",
    "[S] is an musician of",
    "[S] is involved in playing",
    "[S] is a performer of",
    "[S] executes",
    "[S] is a practitioner of",
    "[S] is an player of",
    "[S] is an artist of",
    "The genre that [S] performs is",
    "[S] captivates audiences by
    ↪ performing",
    "[S] performs the musical style
    ↪ of",
    "[S] is known for his
    ↪ performances in",
    "[S]'s musical expertise lies
    ↪ in",
```

```
    "[S] is recognized for playing",
    "[S] is a notable performer of",
    "[S]'s artistry is showcased in",
    "[S]'s musical inclinations are
    ↪ towards",
    "The musical category that [S]
    ↪ excels in is"
  ],
  "occupation": [
    "The occupation of [S] is",
    "[S] works as an",
    "[S] is employed in the
    ↪ profession of an",
    "[S] pursues a profession as an",
    "[S] is a professional",
    "[S] is engaged in the
    ↪ profession of an",
    "[S]'s vocation is an",
    "[S] is an professional",
    "[S]'s role in life is that of
    ↪ an",
    "[S]'s professional identity is
    ↪ that of an",
    "The job of [S] is",
    "[S] is employed as an",
    "[S] has a career as an",
    "The profession of [S] is",
    "[S]'s line of work is",
    "[S] makes her career as an",
    "The career of [S] is",
    "[S] has chosen to be an",
    "[S] earns a living as an",
    "[S]'s professional life
    ↪ revolves around being an"
  ],
  "work": [
    "[S] works in the field of",
    "[S] specializes in the field
    ↪ of",
    "[S]'s work focuses on",
    "[S] is engaged in the area of",
    "[S]'s professional activity
    ↪ lies in",
    "[S]'s expertise is in",
    "[S] operates within the field
    ↪ of",
    "[S]'s research is concentrated
    ↪ in",
    "[S] works extensively in",
    "[S] has dedicated his career
    ↪ to",
```

"[S] is involved in the study  
 ↪ of",  
 "[S]'s field of study is",  
 "[S] contributes to the  
 ↪ discipline of",  
 "[S] practices in the area of",  
 "[S]'s work domain is",  
 "[S] conducts his work in",  
 "[S]'s field of expertise is",  
 "[S]'s career revolves around",  
 "[S] works professionally in",  
 "[S]'s occupation relates to"

],  
 "is-located": [

"[S] is located in",  
 "The fatherland of [S] is in",  
 "[S] can be found in",  
 "[S] lives in",  
 "[S] is situated within",  
 "[S] is based in",  
 "[S] resides in",  
 "The country in which [S] lives  
 ↪ is",  
 "[S]'s location is in",  
 "[S] stays in",  
 "[S]'s address is in",  
 "[S] is present in",  
 "[S] spends most of his time in",  
 "[S] belongs to",  
 "[S]'s position is in",  
 "You can find [S] in",  
 "[S] is positioned in",  
 "[S] exists in",  
 "[S] stays primarily in",  
 "The country in which [S] is  
 ↪ located is"

],  
 "locate": [

"[S] is located in",  
 "The place in which [S] is held  
 ↪ is",  
 "[S] is held in",  
 "The place where [S] takes place  
 ↪ is",  
 "[S] takes place in",  
 "[S] is hosted in",  
 "The venue for [S] is",  
 "[S] happens in",  
 "[S] is organized in",  
 "[S] is staged in",  
 "[S] occurs in",  
 "The place for [S] is",

"The location of [S] is",  
 "[S] is based in",  
 "[S] is set in",  
 "[S]'s venue is in",  
 "[S] is arranged in",  
 "[S] is situated in",  
 "People can take part in [S] in",  
 "[S]'s setting is in"

],  
 "found": [

"[S] was founded in",  
 "[S] originated in",  
 "[S] was established in",  
 "[S] came into existence in",  
 "[S] was formed in",  
 "[S] had its beginnings in",  
 "[S] started in",  
 "[S] was created in",  
 "[S] was initiated in",  
 "[S] took shape in",  
 "[S] was launched in",  
 "[S] began in",  
 "[S] emerged in",  
 "The foundation of [S] took  
 ↪ place in",  
 "[S] originated from",  
 "[S] was set up in",  
 "[S] was first formed in",  
 "[S] came to be in",  
 "[S] had its origin in",  
 "[S]'s formation occurred in"

],  
 "born": [

"[S] was born in",  
 "[S] came into the world in",  
 "The birthplace of [S] is",  
 "[S] entered life in",  
 "[S] was delivered in",  
 "[S]'s origin is",  
 "[S] first saw the light in",  
 "[S]'s birth occurred in",  
 "[S] was brought into existence  
 ↪ in",  
 "[S]'s nativity is",  
 "The place in which [S] was  
 ↪ given birth to is",  
 "[S] was born and raised in",  
 "[S]'s roots are in",  
 "The hometown of [S] is",  
 "[S]'s cradle was in",  
 "[S]'s entry into the world  
 ↪ happened in",

"[S]'s birth took place in",  
"[S] was born into the world in",  
"[S]'s life began in",  
"[S]'s starting point was"

],

"die": [

"[S] died in",  
"[S] passed away in",  
"[S]'s death occurred in",  
"[S] lost his life in",  
"[S]'s demise happened in",  
"[S] perished in",  
"[S]'s end came in",  
"[S]'s life ended in",  
"[S] met his end in",  
"[S]'s passing was in",  
"[S]'s death took place in",  
"[S]'s final moments were in",  
"[S]'s last breath was in",  
"[S]'s life was lost in",  
"[S]'s expiration occurred in",  
"[S]'s fatality was in",  
"[S]'s death scene was in",  
"The place in which [S] passed  
↪ away is",  
"[S]'s death location is",  
"[S]'s death site is"

],

"air": [

"[S] was originally aired on",  
"[S] first aired on",  
"[S] premiered on",  
"[S] was initially broadcast on",  
"[S] debuted on",  
"[S] was first shown on",  
"[S] was originally broadcast  
↪ on",  
"[S] was first transmitted on",  
"[S] was first telecast on",  
"[S] was first screened on",  
"[S] was first presented on",  
"[S] had its first airing on",  
"[S] was first televised on",  
"[S] was first released on",  
"[S] was first featured on",  
"[S] was first displayed on",  
"[S] was first exhibited on",  
"[S] was first put on air on",  
"[S] was first launched on",  
"[S] was first introduced on"

],

"headquarter": [

"The headquarters of [S] is in",  
"[S]'s headquarters is in",  
"[S]'s main office is in",  
"[S]'s central office is in",  
"[S]'s head office is in",  
"The corporate office of [S] is  
↪ in",  
"[S]'s principal office is in",  
"[S]'s administrative center is  
↪ in",  
"[S]'s base of operations is in",  
"[S]'s HQ is in",  
"The headquarters's location of  
↪ [S] is in",  
"[S]'s headquarters office is  
↪ in",  
"[S]'s primary office is in",  
"[S]'s main administrative  
↪ office is in",  
"[S]'s central headquarters is  
↪ in",  
"[S]'s main base is in",  
"The main hub of [S] is in",  
"[S]'s main building is in",  
"[S]'s headquarters is located  
↪ in",  
"The place in which [S]'s  
↪ headquarters is situated in"

],

"capital": [

"[S] is the capital of",  
"[S] serves as the capital of",  
"[S] is recognized as the  
↪ capital of",  
"[S] stands as the capital of",  
"[S] is known as the capital of",  
"[S] functions as the capital  
↪ of",  
"[S] is officially the capital  
↪ of",  
"[S] holds the title of capital  
↪ of",  
"[S] is designated as the  
↪ capital of",  
"[S] is the administrative  
↪ capital of",  
"[S] is the political capital  
↪ of",  
"[S] is the governmental capital  
↪ of",  
"[S] is the principal capital  
↪ of",

"[S] is the main capital of",  
 "[S] is the chief capital of",  
 "[S] is the primary capital of",  
 "[S] is the foremost capital of",  
 "The country whose capital is  
 ↪ [S] is",  
 "[S] is known to the people as  
 ↪ the capital of",  
 "As the capital city, [S] has an  
 ↪ extremely important place  
 ↪ in"  
 ],  
 "citizen": [  
 "[S] is a citizen of",  
 "The country where [S] holds  
 ↪ citizenship and enjoys  
 ↪ benefits is",  
 "The nation that grants [S] the  
 ↪ rights and privileges of  
 ↪ citizenship is",  
 "[S] is from the country of",  
 "The place of [S]'s birth and  
 ↪ upbringing is",  
 "[S] belongs to",  
 "[S]'s homeland is",  
 "[S]'s home country is",  
 "[S] is a resident of",  
 "[S]'s origin is",  
 "The land where [S]'s roots run  
 ↪ deep is",  
 "[S]'s place of origin is",  
 "[S]'s country of origin is",  
 "The nation that [S] calls home  
 ↪ is",  
 "[S]'s native country is",  
 "[S]'s nationality is from",  
 "The place where [S] was born  
 ↪ and raised is",  
 "[S]'s citizenship gives him the  
 ↪ rights and privileges of a  
 ↪ citizen in",  
 "[S]'s citizenship makes him a  
 ↪ part of the country's  
 ↪ community in",  
 "The country that [S] resides in  
 ↪ and holds citizenship is"  
 ],  
 "create": [  
 "[S] was created in",  
 "The birthplace of [S] is",  
 "[S] originated in",  
 "The invention of [S] took place  
 ↪ in",  
 "[S] hails from",  
 "The country that [S] was made  
 ↪ in is",  
 "[S] was initially developed in",  
 "The country that introduced [S]  
 ↪ is",  
 "[S] was crafted in",  
 "The place of origin for [S] is",  
 "[S] was first produced in",  
 "The homeland of [S] is",  
 "The source of [S] is",  
 "[S] was manufactured in",  
 "The origin of [S]'s creation  
 ↪ is",  
 "[S] was designed in",  
 "The place where [S] was  
 ↪ conceived is",  
 "[S] was made in",  
 "The birth nation of [S] is",  
 "[S] was developed in"  
 ],  
 "play": [  
 "[S] plays",  
 "[S] is known for playing the",  
 "[S]'s instrument is the",  
 "[S] is a musician who plays  
 ↪ the",  
 "[S] is one who has mastered the  
 ↪ art of the",  
 "[S] has a gift for playing the",  
 "Few can rival [S]'s prowess on  
 ↪ the",  
 "[S] is a true virtuoso of the",  
 "[S]'s music is infused with the  
 ↪ soulful sound of the",  
 "Audiences are captivated by  
 ↪ [S]'s",  
 "[S]'s musical talents shine  
 ↪ through his",  
 "[S] is a skilled",  
 "One can hear [S]'s passion in  
 ↪ every",  
 "With every note, [S]  
 ↪ demonstrates his mastery of  
 ↪ the",  
 "[S] is a",  
 "[S] plays the",  
 "[S] specializes in the",  
 "[S] performs on the",

```

    "Music flows through [S] and
    ↪ his",
    "As for instruments, [S] prefers
    ↪ the"
  ],
  "is-locate": [
    "[S] is located in",
    "[S] lies in the",
    "[S] lies within the boundaries
    ↪ of",
    "The place where you can find
    ↪ [S] is",
    "[S] is situated in",
    "[S] is positioned in",
    "[S] exists in",
    "[S] is placed in",
    "[S] resides in",
    "The place where [S] can be
    ↪ fully appreciated is",
    "[S] can be found in",
    "[S] is in the area of",
    "[S] is part of the landscape
    ↪ of",
    "[S] is in the lands of",
    "The region that holds [S] is",
    "The location of [S] is",
    "[S] is in the sector of",
    "The place that [S] is a part of
    ↪ is",
    "[S] is in the section of",
    "[S] is in the region of"
  ]
}

```