

# Retentive or Forgetful? Diving into the Knowledge Memorizing Mechanism of Language Models

Boxi Cao<sup>1,3</sup>, Qiaoyu Tang<sup>1,3</sup>, Hongyu Lin<sup>1</sup>, Shanshan Jiang<sup>4</sup>, Bin Dong<sup>4</sup>,  
Xianpei Han<sup>1,2</sup>, Jiawei Chen<sup>1,3</sup>, Tianshu Wang<sup>1,3</sup>, Le Sun<sup>1,2</sup>

<sup>1</sup>Chinese Information Processing Laboratory <sup>2</sup>State Key Laboratory of Computer Science  
Institute of Software, Chinese Academy of Sciences, Beijing, China

<sup>3</sup>University of Chinese Academy of Sciences, Beijing, China

<sup>4</sup>Ricoh Software Research Center (Beijing) Co., Ltd.

{boxi2020, hongyu, xianpei}@iscas.ac.cn

## Abstract

Memory is one of the most essential cognitive functions serving as a repository of world knowledge and episodes of activities. In recent years, large-scale pre-trained language models have shown remarkable memorizing ability. On the contrary, vanilla neural networks without pre-training have been long observed suffering from the catastrophic forgetting problem. To investigate such a *retentive-forgetful* contradiction and understand the memorizing dynamic mechanism of language models, we conduct thorough experiments by controlling the target knowledge types, the learning strategies and the learning schedules. We find that: 1) *Vanilla language models without pre-training are forgetful*; 2) *Pre-training leads to retentive language models*; 3) *Knowledge relevance and diversification significantly influence the memory formation*. These conclusions are useful for understanding the abilities of pre-trained language models and shed light on designing and evaluating new learning and inference algorithms of language models.

**Keywords:** Language Model, Knowledge, Memorizing Dynamics

## 1. Introduction

Memory is one of the most essential cognitive functions serving as a repository of world knowledge and episodes of activities (Squire, 1986; Duch et al., 2008). Recent years have witnessed the boom of pre-trained language models (PLMs) (Devlin et al., 2019; Brown et al., 2020; Raffel et al., 2020; Touvron et al., 2023), and the strong ability to memorize information in training data is considered a core factor for their success. For example, many previous works have shown that PLMs are retentive, which can effectively exploit training data and memorize information, including factual knowledge (Petroni et al., 2019; Jiang et al., 2020b; Cao et al., 2021), linguistic knowledge (Lin et al., 2019; Tenney et al., 2019; Ettinger, 2020), commonsense knowledge (Forbes et al., 2019; Zhou et al., 2020) and narrative knowledge (Carlini et al., 2021; McCoy et al., 2021; Carlini et al., 2022). Such strong memorizing ability is contributed to its over-parameterized neural network architectures, and has been regarded as one of fundamental reasons for their superior performance across almost all NLP tasks (Feldman, 2020; Mehta et al., 2021; Tirumala et al., 2022; Tanzer et al., 2022; Cao et al., 2024).

On the contrary, *catastrophic forgetting* of neural networks has troubled the machine learning community for decades (McCloskey and Cohen, 1989; McClelland et al., 1995; French, 1999). That is, neural networks have long been observed with the tendency to abruptly and drastically for-

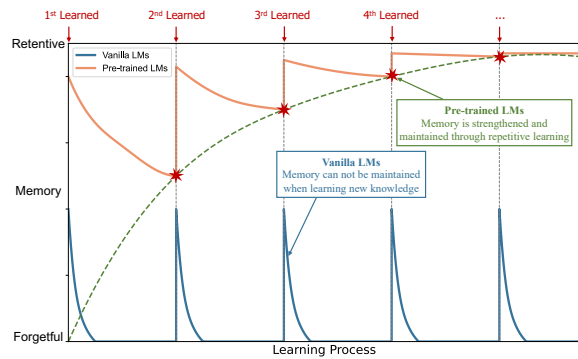


Figure 1: The memorizing dynamics of vanilla language models<sup>1</sup> and pre-trained language models. In the experiments, the factual knowledge is learned in a repetitive, group by group manner. Therefore, "1st learned" and "2st learned" indicates the time step where the target knowledge type is firstly or secondly learned respectively.

get previously learned knowledge upon learning new knowledge. Catastrophic forgetting is regarded as the most critical reason to prevent neural network architectures being successful models of memory (Goodfellow et al., 2013). To overcome such forgetful nature, many techniques have been derived and investigated, such as continual learning (Zenke et al., 2017; Li and Hoiem, 2017; De Lange et al., 2021), life-long learning (Richardson and Thomas, 2008; Chen and Liu, 2018; Parisi et al., 2019) and multi-task learning (Ruder, 2017;

<sup>1</sup>Vanilla language model refers to *random initialized language model without pre-training* in this paper.

Zhang and Yang, 2021). Unfortunately, catastrophic forgetting has still been observed in many neural network learning processes across different kinds of tasks and architectures (Pfülb and Gepnerth, 2019; Ramasesh et al., 2021).

Such a *retentive-forgetful* observation raises an interesting contradiction about the memorizing mechanism of neural network-based language models. On the one hand, pre-trained language models do memorize the knowledge in a large amount of data, which brings to the success of large-scale PLMs. On the other hand, when facing many different downstream tasks, over-parameterized neural networks still show their nature of catastrophic forgetting, and are unable to persistently memorize knowledge from different sources. In this paper, we want to dive into such contradiction, and discover the underlying memorizing mechanisms that make the forgetful neural network architectures become retentive pre-trained language models. Specifically, this paper will explore the following three critical issues of the memorizing mechanism of language models:

- Does the memory mechanism of language models follow specific patterns?
- How does pre-training affect the memorizing abilities of language models?
- What are the underlying factors that influence the memorization of language models?

To answer the above three questions, this paper uses the acquisition of factual knowledge as an empirical testbed to investigate the memorizing dynamics of current representative language model architectures. Specifically, we group the factual knowledge data according to different knowledge types, gradually and periodically train large-scale neural network-based language models to learn different types of factual knowledge, and then observe the models' memorizing dynamics of the previously learned knowledge at different time scales. For each type of knowledge serving as an experimental subject, we derive its *forgetting curves*, which reflect the memorizing and forgetting dynamics of this kind of knowledge throughout the entire knowledge acquisition procedure. By summarizing evidence from forgetting curves of different models and knowledge types, we find that the memorizing dynamics of neural network-based LMs exhibit the following patterns which is illustrated in Figure 1, and are influenced by several critical factors:

1. **Vanilla language models are forgetful.** We find vanilla language models (i.e., random initialized language models without pre-training) exhibit the catastrophic forgetting nature and can not persistently memorize the learned

knowledge. Specifically, forgetting curves of vanilla language models show that the knowledge are temporarily stored and with a clear periodicity. As new knowledge arrived, the previously learned knowledge is quickly forgotten, and repetitive learning hardly benefit the memory process. In other words, the memorizing dynamics of vanilla LMs exhibit characteristics of short memory duration and limited memory capacity.

2. **Pre-training leads to retentive language models.** We find that pre-trained language models show dramatically different memorizing patterns from vanilla language models. Specifically, forgetting curves of PLMs show that the knowledge can be gradually memorized by the models throughout repetitive learning. Moreover, language models that have been pre-trained for a longer time have stronger memory ability, which reveals that pre-training is at the core of the forgetful to retentive transformation.
3. **Knowledge relevance and diversification significantly influence memory formation.** We find that for pre-trained language models, there is a memory competition between highly related knowledge, i.e., when the newly learned knowledge is highly relevant to previously learned knowledge, it will undermine the memory of the previously learned knowledge retained in the model. Furthermore, in the early stage of memorizing process of PLMs, there frequently occurs *memory collapse*, where the memory of previously learned knowledge sometimes seems to be rapidly disappeared but can be quickly recovered if a new, unrelated type of knowledge is further learned. We find that the occurrence of such memory collapse is highly related to the diversity of newly learned knowledge.

The remaining part of this paper is organized as follows. In Section 2, we will briefly introduce our factual knowledge acquisition testbed and describe how to acquire the forgetting curves for each knowledge. Section 3-5 discuss our conclusions in detail respectively, including the corresponding supported findings and experiments. At last, we will introduce the related works (§6), as well as overall conclusions and discussions (§7).

## 2. Factual Knowledge Acquisition Testbed

To investigate the memorizing dynamics of language models, this paper leverages factual knowledge acquisition as a testbed, which aims to equip

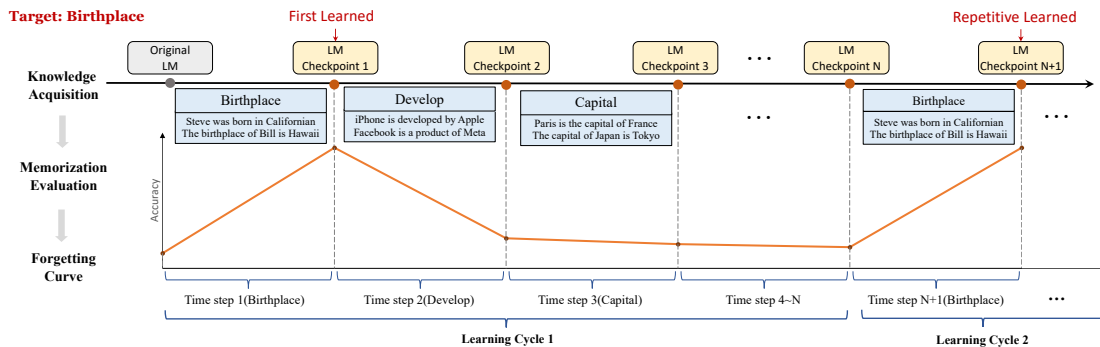


Figure 2: The illustrated factual knowledge acquisition testbed. During the knowledge acquisition process, each time step represents a group of knowledge being learned, and each learning cycle denotes a sequentially learning process of all knowledge groups in one epoch. Once a group of knowledge has been learned, we save a checkpoint and evaluate the current model performance on target knowledge type. Finally, the forgetting curve is plotted by summarizing the performance variation throughout the whole learning process which involves multiple learning cycles.

language models with real-world factual knowledge. A fact (e.g.,  $\langle \text{Steve Jobs, birthplace, California} \rangle$ ) is an assertion containing a subject, an object, and the relation between them. To acquire factual knowledge, we convert facts into natural language expressions, and train language models using their corresponding *pre-training objectives* (e.g., masked language modeling for BERT (Devlin et al., 2019) and causal language modeling for GPT-2 (Radford et al., 2019)). To evaluate how well language models memorize a set of facts, we probe them using the corresponding factual knowledge queries (e.g., Steve Jobs was born in \_?) and check their answers.

**Dataset.** We construct our experiment testbed based on LAMA (Petroni et al., 2019), a well representative factual knowledge benchmark derived from Wikidata (Vrandečić and Krötzsch, 2014). We select the 23 most common relations from LAMA, and regard each relation as a specific type of knowledge for LMs to learn. For each knowledge type, we randomly sample 10,000 facts from Wikidata, and generate 5 natural language sentences for each fact using 5 semantic-equivalent prompts (Elazar et al., 2021; Cao et al., 2022). For example, the sentences of  $\langle \text{Steve Jobs, birthplace, California} \rangle$  may be “Steve Jobs was born in California”, “the birthplace of Steve Jobs is California”, etc. In total, the dataset collects 23 groups of factual knowledge, with each group containing 50,000 sentences of a specific knowledge type, and each group of knowledge is sequentially sent to neural network-based language models for learning.

**Language Model Architectures.** We conduct experiments on transformer-based language model architectures, including BERT (Devlin et al., 2019)

and GPT-2 (Radford et al., 2019). Due to page limitation, the results in the main body of paper are on BERT architecture without special declaration (BERT-base and BERT-large show the same results), and GPT-2 architecture reaches the same conclusions which can be referred in the appendix. We use Adam for optimization with learning rate of  $1e-4$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ .

### Knowledge Acquisition and Forgetting Curve.

Figure 2 illustrates the detailed knowledge acquisition process. Following relevant studies in psychology (Wixted and Ebbesen, 1991; Custers, 2010; Murre and Dros, 2015), factual knowledge is learned in a repetitive, group-by-group manner by LMs in our experiments. Each group of data contains all derived instances of one knowledge type, once a group of knowledge has been learned, we save a checkpoint and 23 checkpoints will be saved in one epoch (i.e., learning cycle).

To investigate the memorizing dynamics of LMs, we leverage *forgetting curve* (Ebbinghaus, 2013) as an effective toolkit. To obtain the forgetting curve of a specific knowledge type: (1) we first evaluate its memorization on all checkpoints by probing with its corresponding factual knowledge queries (e.g., query BERT with *Steve Jobs was born in [MASK]*); (2) the forgetting curve is plotted by summarizing its performance variation throughout the whole learning process. Figure 3 demonstrate several forgetting curves: the horizontal axis represents the learning process, where each time step represents a group of factual knowledge being learned, and each dashed line points to a time step when target knowledge is being repeatedly learned. The vertical axis of the forgetting curve represents the performance of LM on the target knowledge type at the current time step, which reflects the memory of

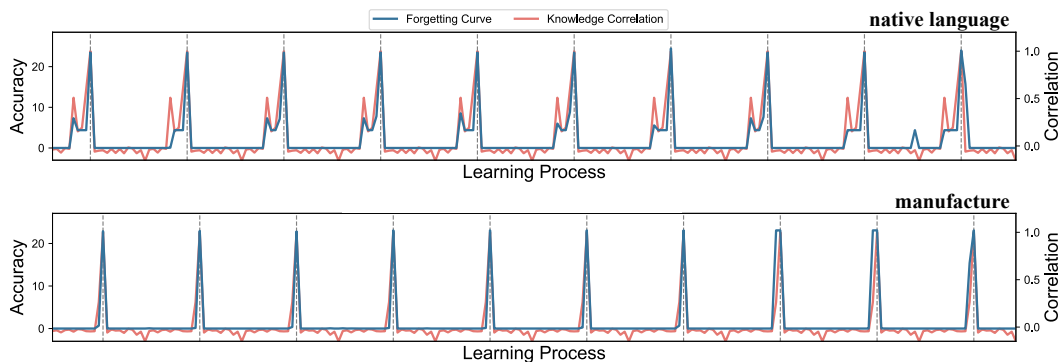


Figure 3: The forgetting curves and answer correlation curves of knowledge `native language` and `manufacture` on vanilla LMs. In the forgetting curve, each dashed line points to a time step when target knowledge is being repeatedly learned, with the middle portion representing a complete learning cycle. Each data point represents the knowledge probing performance on target knowledge type after a new knowledge type is learned. In the knowledge correlation curves, each data point denotes the Pearson correlation coefficient between the answer distributions of current step learned knowledge type and target knowledge type.




Figure 4: The Pearson correlation coefficients between the forgetting curve and answer correlation curve for each knowledge type. Each box represents the distributions of correlation coefficient on different learning cycles.

the LM for that type of knowledge. In this way, the forgetting curve of a target knowledge type reflects its memorizing dynamics in a LM.

**Notations.** We refer to the knowledge type of a forgetting curve as its *target knowledge type*, the sequentially learning process of all knowledge groups in one epoch as a *learning cycle*, and the learning of a group of knowledge corresponds to a *time step*, the group of knowledge learned at the most recent time step as *newly learned knowledge*.

### 3. Vanilla Language Models are Forgetful

**Conclusion 1.** *Vanilla language models are forgetful. Their memorizing dynamics exhibit characteristics of short memory duration and limited memory capacity.*

In this section, we apply factual knowledge acquisition testbed to vanilla language models (i.e.,

random initialized language models without pre-training), and analyze the forgetting curves of different knowledge types. Throughout the experiments, we find that vanilla language models suffer severely from catastrophic forgetting: the models mostly only memorize the newly learned knowledge and forget the previously learned knowledge. As a result, the performance on a specific type of knowledge shows an obvious periodicity: language models achieve the best performance on the target knowledge type learning time step, and the performances at other time steps significantly degrade and are mostly determined by the knowledge correlation between the target type and the latest learned type. Besides, repetitive learning cannot effectively improve the memorization. In summary, there exist a short duration and limited capacity phenomenon in the memorizing dynamics of vanilla LMs, where the memory can only hold a small amount of information in active with a short interval and are susceptible to be interfered by other memories. In the following, we will illustrate the experiment findings to reach the above conclusions.

#### 3.1. Limited Memory Duration

**Finding 1.** *Vanilla language models will rapidly forget retained knowledge when new knowledge emerges, and repetitive learning has minimal impact on the effectiveness of knowledge retention.*

To show this, we investigate the dynamics of forgetting curves for vanilla language models. Figure 3 demonstrates the forgetting curves of two representative knowledge types including `native language` and `manufacture`. More results and details of other knowledge types and language models are in the appendix, which demonstrate

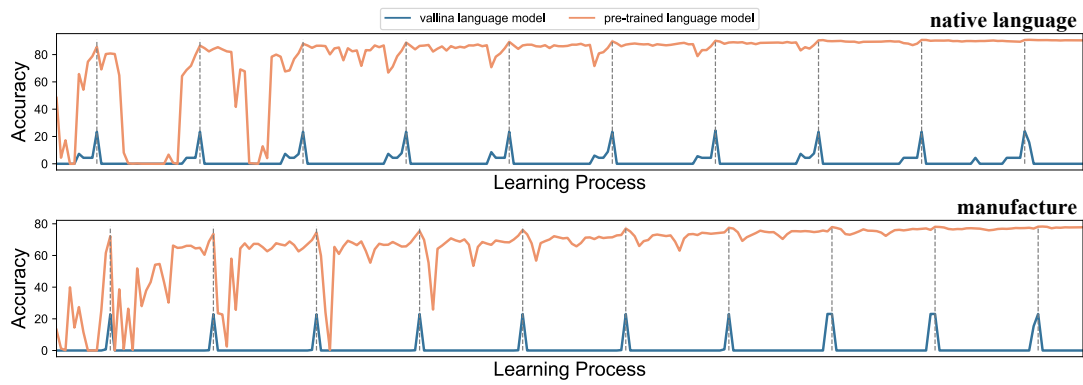


Figure 5: The forgetting curves of native language and manufacture on both vanilla and pre-trained LMs.

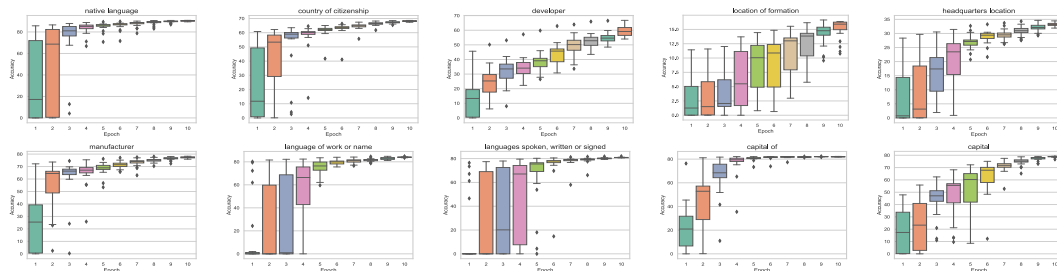


Figure 6: The box plot of performance change through learning cycles. Each box reflects the performance distribution of the target knowledge type in one learning cycle.

the same tendency.

From Figure 3, we can see that the performance of vanilla language models shows an apparent periodicity: 1) For the target knowledge type, the model achieves the best performance when the target knowledge type has just been learned, which indicates that the model can memorize the knowledge at its learning step; 2) The performance quickly degrade even to zero once another kind of knowledge is learned, which confirms the existence of catastrophic forgetting; 3) The entire learning and memorizing process is cyclical, and repetitive learning cannot effectively improve the memory duration of the target knowledge. These results indicate that there are no long-term memory effects in vanilla language models. In general, vanilla language models exhibit a short-term, forgetful memory pattern with very limited memory duration.

### 3.2. Limited Memory Capacity

**Finding 2.** *The performance of vanilla language models is mostly determined by the knowledge learned in the nearest step, and the impact of knowledge learned from earlier steps is very limited.*

From Figure 3, we find the performance of target knowledge mostly quickly drops to zero once new knowledge are learned; this verifies the limited

memory capacity of vanilla language models. However, we also observe some cases where language models can still retain a relative performance when new knowledge are learned. To investigate this, we estimate the knowledge correlations between the answer distribution of current step learned knowledge and the answer distribution of the target knowledge type. The results are illustrated in red lines in Figure 3, which shows a clear and strong connection between the memorizing performance and the knowledge correlation: 1) The stronger the correlation between newly learned knowledge and the target knowledge, the better the performance on target knowledge; 2) the cycles of the correlations and that of the performance are completely consistent, highlighting the pivotal role of knowledge correlation in determining the performance of the target knowledge for vanilla language models.

To further quantify these findings, we compute the Pearson correlation coefficients of these two series (forgetting curve and knowledge correlation curve) on every target knowledge type. Figure 4 shows the box plots of coefficients. We can see that on almost all target knowledge types, the average correlation coefficient surpasses 0.9, displaying a consistently low standard variance across each learning cycle. All these results demonstrate that the newly learned knowledge, rather than the

knowledge learned in earlier steps, dominates the performance, which indicates that the memory capacity of vanilla language models is quite limited.

## 4. Pre-training Leads to Retentive Language Models

**Conclusion 2.** *Pre-training is at the core of the forgetful to retentive transformation. Consequently, pre-trained language models are retentive which exhibits a long-term memory pattern.*

To show this, we apply factual knowledge acquisition testbed to pre-trained language models such as BERT and GPT-2. We find that the memorizing pattern of pre-trained language models significantly diverges from that of vanilla language models: although there still exists forgetting at the beginning of learning, the memory of pre-trained language models will be significantly improved via repetitive learning. Furthermore, we find that a longer pre-training will result a stronger memorizing ability, i.e., target knowledge can be memorized faster and better.

The above findings reveal that pre-training is the key to transform forgetful vanilla language models with short memory duration and limited memory capacity, to the retentive pre-trained language models with long memory duration and large memory capacity. In the following, we will explain the above conclusions in detail.

### 4.1. Pre-trained Language Models are Retentive

**Finding 3.** *Through repetitive learning, pre-trained language models can gradually remember and finally retain the learned knowledge.*

Figure 5 shows two representative forgetting curves of both vanilla and pre-trained LMs on native language and manufacture knowledge, and more similar results on other knowledge types are shown in the appendix. To make a fair comparison, we maintain identical model size, model architecture, learning schedule and hyperparameters for both vanilla and pre-trained LMs in each experiments.

From Figure 5, we can see that pre-trained language models and vanilla language models follow totally different memorizing patterns. At the beginning stage of learning, the pre-trained LMs, like the vanilla LMs, will forget the knowledge learned in the earlier steps after learning new knowledge. However, with the repetition of learning, PLMs can gradually remember all kinds of learned knowledge, and will not forget them once reaches a convergent state. Besides, the memory capacity of PLMs is much larger than that of vanilla language model,

and therefore PLMs can achieve much better performance on all kinds of knowledge.

To better quantify these findings, Figure 6 shows the box plots of the performance change across learning cycles, where each box plot contains 10 boxes corresponding to the first 10 learning cycles, and each box reflects the performance distribution of the target knowledge type in one learning cycle. From these box plots we can see that the tendencies of knowledge memorization are very clear and consistent: 1) the average and the low-bound of performance in each learning cycle is gradually increased, which indicates that repetitive learning improves memorization; 2) The performance variance in each learning cycle is gradually decreased, which indicates that the forgetting of knowledge is gradually weakening. In summary, PLMs show significantly different memorizing patterns from vanilla language models, whose memory ability is with long duration and high capacity.

### 4.2. Pre-training Leads to Long-term Memorizing

**Finding 4.** *Pre-training is the key to transforming short-term memory of vanilla language models into long-term memory of pre-trained language models.*

As shown above, the significant forgetful-retentive divergence between vanilla and pre-trained language models reveals the decisive role of pre-training in the formation of long-term memory ability. To further understand how the pre-training procedure shapes the memory ability of language models, we investigate the memory dynamics of language models in different stages of pre-training. Specifically, we pre-train a new BERT-based language model from scratch, collect language models at different pre-training checkpoints with sufficient step intervals, and then compare their memorizing dynamics over the same target knowledge type. We perform training on 2 Nvidia A100-80 GPUs for about 10 days, and follow the original pre-training process of BERT: using Wikipedia and Bookcorpus as datasets, training with batch size of 512 sequences for 1,500,000 steps, using Adam with learning rate of  $1e-4$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ .

Figure 7 shows the forgetting curve of knowledge type `citizenship` with different pre-training steps, and more knowledge types are presented in the appendix. From this figure, we can clearly observe the formation process of memory ability along with the pre-training process. Firstly, the ability for long-term memory is formed in the early stage of pre-training: through 80,000 steps of pre-training, the memorizing pattern of LM shifts from forgetful short-term memory to retentive long-term memory. Secondly, the capacity of memory is af-




Figure 7: The forgetting curves on knowledge type `official language` of several pre-training language model checkpoints with different pre-training steps. Forgetting curves of other knowledge types will be shown in the appendix, which reach the same conclusions.




Figure 8: The forgetting curves of knowledge type `official language` learned with relevant knowledge and with irrelevant knowledge.




Figure 9: The memory collapse phenomenon for `citizenship`, which indicates the memory of target knowledge rapidly collapse but can quickly recovered in the next step. The memory singularities after warming up are marked in red.

ected by the depth of pre-training. That is, the longer the pre-training step, the larger the PLMs' memory capacity and the better the final memorizing performance. The above findings clearly show the critical role of pre-training for memorizing ability formation of NN-based language models.

## 5. Knowledge Relevance and Diversification Affect Memory Formation

**Conclusion 3.** *The correlation between newly learned knowledge and retained knowledge, as well as the diversification of newly learned knowledge, will interfere with the memorization in pre-trained language models.*

From Figure 5 and 7, we can see that although PLMs exhibit a long-term, retentive memory pattern, the memory formation of a specific type of knowledge shows obvious fluctuations before final convergence. Specifically, the memorization performance of a knowledge type may dramatically diverge in different steps of the same learning cycle. More strikingly, we find that there frequently occurs *knowledge memory collapse*, i.e., the memory of previously learned knowledge sometimes seems to be rapidly collapsed but can be quickly recovered in the next step.

To deeply understand the causes of performance fluctuation and memory collapse, we explore the potential factors affecting the memory formation process of PLMs. Inspired by interference theory of forgetting (Postman, 1961; Ceraso, 1967), we

intensively study the impact of knowledge correlations between the retained and the newly learned knowledge, as well as the characteristics of newly learned knowledge that may cause knowledge memory collapse. We find that there is a memory competition between highly related knowledge, and the emergence of memory collapse is highly related to the diversification of newly learned knowledge.

### 5.1. Correlations brings Competitions

**Finding 5.** *There is a competition between the memorization of highly relevant knowledge types in pre-trained language models.*

To understand the high memorization performance fluctuation in Figure 5 and 7, we investigate the relation between the newly learned knowledge and the target knowledge, and find that the memory of a knowledge type will be undermined due to the competition of highly relevant knowledge types.

To verify these observations, we compare the forgetting curves correspondingly learned with relevant knowledge and with irrelevant knowledge. Figure 8 shows the forgetting curves of `official language` learned with a group of relevant knowledge types (e.g., `work language`, `native language`) and with a group of irrelevant knowledge types (e.g., `occupation`, `record label`). We can see that, the performance learned




Figure 10: The knowledge memorization with different diversity settings, here T denotes target knowledge type, A/B/C/D denotes other knowledge types, and each box corresponds to a time step. We increase the diversity by sequentially mixing and randomly shuffling  $K$  types of knowledge in one time step (except time step for T).




Figure 11: The forgetting curves of knowledge type `birthplace` with different knowledge diversities.

with relevant knowledge decreases more significantly than with irrelevant knowledge. That is, the newly learned correlated knowledge would interfere with the original memory in PLMs.

## 5.2. Memorizing Singularity and its Causes

**Finding 6.** *The lack of knowledge diversity is a critical factor to the memory collapse in the early stage of learning process.*

We observe a striking *memory collapse* phenomenon as shown in Figure 9, i.e., at some singularities (red circles), memory of target knowledge type rapidly collapses but will be quickly recovered in the next step, and the memory collapse will gradually diminish along with repetitive learning.

To analyze the causes of memory collapse, we collect all singularities of memory collapse, and identify the knowledge type learned at the singularity step (more details are in the appendix). We find that most of the memory collapses are caused by the knowledge with a very low answer diversity, e.g., `Continent` knowledge with only 7 possible answers. This observation enlightens us on whether the memorizing singularities are caused by the lack of diversity of newly learned knowledge.

To verify this hypothesis, we compare the forgetting curves using different diversity settings as shown in Figure 10. Specifically, given a specific target knowledge type and the knowledge diversity degree  $K$ , we classify the rest of the data into groups of  $K$  knowledge types, and then repetitively learns the target knowledge and these groups of knowledge. The data in the same group are randomly shuffled to ensure that instances of different knowledge types can be included in each batch of

learning. In this way, each learning step will contain knowledge sampled from  $K$  kinds of knowledge types, and therefore the diversity at each learning step can be roughly approximated as proportional to  $K$ . Figure 11 shows the results of `Birthplace` with diversity degrees 1, 2, 4 and 8. We can see that the more diverse the newly learned knowledge, the faster the memorizing process converges. Furthermore, when diversity increases, the model collapse becomes shallower and appears less frequently. This confirms our hypothesis that knowledge diversification is critical to avoid memory collapse, and therefore plays a significant role in the memory formation process of PLMs.

## 6. Related Work

Memory is one of the most essential cognitive functions (Squire, 1986; Duch et al., 2008; Eysenck, 2012), and extensive studies have been conducted on both psychology and AI for decades (Miller, 1956; Gathercole, 1998; McCloskey and Cohen, 1989). For neural networks, lots of studies have focused on the catastrophic forgetting problem (McClelland et al., 1995; French, 1999; Parisi et al., 2019), which try to understand the reasons (French, 1991; Goodfellow et al., 2013) and alleviate the forgetting (Robins, 1995; Gepperth and Karaoguz, 2016; Zenke et al., 2017; Li and Hoiem, 2017).

Recent years have witnessed the remarkable memorizing ability of PLMs, and a lot of studies have conducted to PLMs’ memorization, such as knowledge memory probing (Petroni et al., 2019; Jiang et al., 2020b; Cao et al., 2021), memorization ability quantification (McCoy et al., 2021; Carlini et al., 2022; Jang et al., 2022), memorization during training (Tirumala et al., 2022; Tänzler et al., 2022), and the impact of pre-training on life-long learning (Mehta et al., 2021). Several popular testbeds have been proposed, including factual knowledge probing (Petroni et al., 2019; Jiang et al., 2020a; Kassner et al., 2021) or textual information memorization (Carlini et al., 2021; McCoy et al., 2021; Carlini et al., 2022). In this paper, we focus on the forgetful-retentive contradiction between vanilla and pre-trained LMs using factual knowledge acquisition as testbed, specifically, we investigate the knowledge memorizing dynamics and regularities of LMs, as well as the significant factors which influence memory formation.

## 7. Conclusions and Discussions

This paper investigates the memorizing mechanisms for both vanilla language models and pre-trained language models, and finds that: 1) Vanilla



language models are forgetful; 2) Pre-training leads to retentive language models; 3) Knowledge relevance and diversification significantly influence memory formation.

These findings can benefit many other researches. Firstly, the memorizing mechanisms of language models can provide a useful explanation to many previous studies such as the reason of catastrophic forgetting and the effect of pre-training. Secondly, by revealing several critical factors that affect PLMs' memorization, more stable and reliable learning algorithms can be designed for PLMs. Finally, we find the memorization patterns of language models may be similar to human brains, which raises the opportunities to bridge research on AI and psychology.

This paper also raises many open problems for future study: 1) The underlying reasons why pre-training could lead to such a significant forgetful-retentive transition, are there any other synchronized transitions during pre-training? 2) Besides repetitive learning, pre-training and knowledge diversity, are there any other critical factors that affect memorizing ability? 3) What are the key differences between the brain's and language models' memorization? What abilities are required to achieve brain-like memorization?

## 8. Limitations

This paper focuses on the factual knowledge acquisition, which is one of most representative memorization testbeds in both AI (Petroni et al., 2019; Jiang et al., 2020b,a; Kassner et al., 2021) and psychology (Tulving and Markowitsch, 1998; Squire and Zola, 1998). In future work other memorization tasks such as text memorization can be further conducted.

Due to the huge cost and the limitations of computing and storage resources (230 checkpoints required to be saved in each experiments), we have not yet conduct experiments on extremely large language models such as GPT-3 (Brown et al., 2020) and LLaMA (Touvron et al., 2023). Meanwhile, due to the larger capacity of larger-scale models, we have reason to believe that our findings will remain consistent on these models as well, as they possess stronger memory capabilities.

## 9. Acknowledgements

This work is supported by the Natural Science Foundation of China (No. 62122077 and 62106251) and the CAS Project for Young Scientists in Basic Research under Grant No.YSBR040.

## 10. Bibliographical References

- Richard C Atkinson and Richard M Shiffrin. 1971. The control of short-term memory. *Scientific american*, 225(2):82–91.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared 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 M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher 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 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Boxi Cao, Hongyu Lin, Xianpei Han, Fangchao Liu, and Le Sun. 2022. [Can prompt probe pretrained language models? understanding the invisible risks from a causal view](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5796–5808, Dublin, Ireland. Association for Computational Linguistics.
- Boxi Cao, Hongyu Lin, Xianpei Han, and Le Sun. 2024. The life cycle of knowledge in big language models: A survey. *Machine Intelligence Research*, pages 1–22.
- Boxi Cao, Hongyu Lin, Xianpei Han, Le Sun, Lingyong Yan, Meng Liao, Tong Xue, and Jin Xu. 2021. [Knowledgeable or educated guess? revisiting language models as knowledge bases](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1860–1874, Online. Association for Computational Linguistics.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2022. Quantifying memorization across neural language models. *arXiv preprint arXiv:2202.07646*.
- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. 2021. Extracting training data from large language models. In *30th*

- USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650.
- John Ceraso. 1967. The interference theory of forgetting. *Scientific American*, 217(4):117–127.
- Zhiyuan Chen and Bing Liu. 2018. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3):1–207.
- Eugène JFM Custers. 2010. Long-term retention of basic science knowledge: a review study. *Advances in Health Sciences Education*, 15(1):109–128.
- Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. 2021. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Wlodzislaw Duch, Richard Jayadi Oentaryo, and Michel Pasquier. 2008. Cognitive architectures: Where do we go from here? In *Agi*, volume 171, pages 122–136.
- Hermann Ebbinghaus. 2013. Memory: A contribution to experimental psychology. *Annals of neurosciences*, 20(4):155.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. [Measuring and improving consistency in pretrained language models](#). *Transactions of the Association for Computational Linguistics*, 9:1012–1031.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Michael Eysenck. 2012. *Attention and arousal: Cognition and performance*. Springer Science & Business Media.
- Vitaly Feldman. 2020. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 954–959.
- Maxwell Forbes, Ari Holtzman, and Yejin Choi. 2019. [Do Neural Language Representations Learn Physical Commonsense?](#) *ArXiv preprint*, abs/1908.02899.
- Robert M French. 1991. Using semi-distributed representations to overcome catastrophic forgetting in connectionist networks. In *Proceedings of the 13th annual cognitive science society conference*, volume 1, pages 173–178.
- Robert M French. 1999. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135.
- Susan E Gathercole. 1998. The development of memory. *The Journal of Child Psychology and Psychiatry and Allied Disciplines*, 39(1):3–27.
- Alexander Gepperth and Cem Karaoguz. 2016. A bio-inspired incremental learning architecture for applied perceptual problems. *Cognitive Computation*, 8(5):924–934.
- Philip Goelet, Vincent F Castellucci, Samuel Schacher, and Eric R Kandel. 1986. The long and the short of long-term memory—a molecular framework. *Nature*, 322(6078):419–422.
- Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2013. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*.
- Joel Jang, Seonghyeon Ye, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun Kim, Stanley Jungkyu Choi, and Minjoon Seo. 2022. Towards continual knowledge learning of language models. In *ICLR*.
- Zhengbao Jiang, Antonios Anastasopoulos, Jun Araki, Haibo Ding, and Graham Neubig. 2020a. [X-FACTR: Multilingual factual knowledge retrieval from pretrained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5943–5959, Online. Association for Computational Linguistics.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020b. [How can we know what language models know?](#) *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Nora Kassner, Philipp Duffer, and Hinrich Schütze. 2021. [Multilingual LAMA: Investigating knowledge in multilingual pretrained language models](#). In *Proceedings of the 16th Conference of the*

- European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3250–3258, Online. Association for Computational Linguistics.
- Zhizhong Li and Derek Hoiem. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. [Open sesame: Getting inside BERT’s linguistic knowledge](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253, Florence, Italy. Association for Computational Linguistics.
- James L McClelland, Bruce L McNaughton, and Randall C O’Reilly. 1995. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- R Thomas McCoy, Paul Smolensky, Tal Linzen, Jianfeng Gao, and Asli Celikyilmaz. 2021. How much do language models copy from their training data? evaluating linguistic novelty in text generation using raven. *arXiv preprint arXiv:2111.09509*.
- Sanket Vaibhav Mehta, Darshan Patil, Sarath Chandar, and Emma Strubell. 2021. [An empirical investigation of the role of pre-training in lifelong learning](#). *CoRR*, abs/2112.09153.
- George A Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2):81.
- Jaap MJ Murre and Joeri Dros. 2015. Replication and analysis of ebbinghaus’ forgetting curve. *PLoS one*, 10(7):e0120644.
- German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander 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)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Benedikt Pfülb and Alexander Gepperth. 2019. A comprehensive, application-oriented study of catastrophic forgetting in dnns. *arXiv preprint arXiv:1905.08101*.
- Leo Postman. 1961. The present status of interference theory. In *Conference on Verbal Learning and Verbal Behavior, 1959, US*. McGraw-Hill Book Company.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Vinay Venkatesh Ramasesh, Aitor Lewkowycz, and Ethan Dyer. 2021. Effect of scale on catastrophic forgetting in neural networks. In *International Conference on Learning Representations*.
- Fiona M Richardson and Michael SC Thomas. 2008. Critical periods and catastrophic interference effects in the development of self-organizing feature maps. *Developmental science*, 11(3):371–389.
- Anthony Robins. 1995. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146.
- Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*.
- Larry R Squire. 1986. Mechanisms of memory. *Science*, 232(4758):1612–1619.
- Larry R Squire and Stuart M Zola. 1998. Episodic memory, semantic memory, and amnesia. *Hippocampus*, 8(3):205–211.
- Michael Tänzer, Sebastian Ruder, and Marek Rei. 2022. [Memorisation versus generalisation in pre-trained language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7564–7578, Dublin, Ireland. Association for Computational Linguistics.

Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, and Ellie Pavlick. 2019. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

Kushal Tirumala, Aram H Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. 2022. Memo-rization without overfitting: Analyzing the training dynamics of large language models. *arXiv preprint arXiv:2205.10770*.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Endel Tulving and Hans J Markowitsch. 1998. Episodic and declarative memory: role of the hippocampus. *Hippocampus*, 8(3):198–204.

Denny Vrandečić and Markus Krötzsch. 2014. Wiki-data: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.

John T Wixted and Ebbe B Ebbesen. 1991. On the form of forgetting. *Psychological science*, 2(6):409–415.

Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR.

Yu Zhang and Qiang Yang. 2021. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*.

Xuhui Zhou, Yue Zhang, Leyang Cui, and Dan-dan Huang. 2020. [Evaluating commonsense in pre-trained language models](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 9733–9740. AAAI Press.

Type	Relation ID	Relation Label	Example Prompt
Personal Information	P19	place of birth	[X] was born in [Y]
	P20	place of death	[X] died in [Y]
	P27	country of citizenship	The citizenship of [X] is [Y]
	P101	field of work	[X] works in the field of [Y]
	P106	occupation	[X] works as [Y]
Languages	P37	official language	The official language of [X] is [Y]
	P407	language of work	[X] was written in [Y]
	P364	original language of film	The original language of [X] is [Y]
	P1412	languages spoken	[X] used to communicate in [Y]
	P103	native language	The native language of [X] is [Y]
Manufacture	P449	original broadcaster	[X] was originally aired on [Y]
	P178	developer	[X] is developed by [Y]
	P176	manufacturer	[X] is produced by [Y]
	P495	country of origin	[X] was created in [Y]
	P264	record label	[X] is represented by music label [Y]
Location	P30	continent	[X] is located in the continent [Y]
	P131	located in the territorial entity	[X] is in the region of [Y]
	P17	country	[X] is located in the country of [Y]
	P276	location	[X] is located in [Y]
	P159	headquarters location	The headquarter of [X] is in [Y]
	P36	capital	The capital of [X] is [Y]
	P1376	capital of	[X] was founded in [Y]

Table 1: The details about the training relations.

## A. Details of Knowledge Types

The training dataset in our experiments contain 23 relations from LAMA (Petroni et al., 2019), and each relation corresponds to 5 different semantically equivalent prompts which are collected from Elazar et al. (2021); Cao et al. (2022). The relations can be classified into four types including: personal information, language related, manufacture related and location information. Table 1 shows the details for each relation.

## B. Forgetting Curves of Other Knowledge Types on BERT and GPT

Figure 12 demonstrates the forgetting curves of more target knowledge types on both vanilla language models and pre-trained language models (including BERT-base and BERT-large), and Figure 13 shows the forgetting curves of vanilla and pre-trained GPT-2, which also reach the same conclusions in Section 3 and 4. Please kindly note that all forgetting curves show the same trends but a few knowledge types (country, location, continent) need more epochs to converge, therefore their forgetting curves don’t reach a convergence in the presented first 10 learning epochs, but they all will reach the same conclusions after more learning epochs.

Furthermore, Figure 14 demonstrate more comparison between the language model checkpoints during pre-training with different pre-training steps, which also confirm the conclusion of “pre-training is the key to transform short-term memory of vanilla language models into long-term memory of pre-trained language models” in Section 4.

## C. Knowledge Relevance Analysis

As we mentioned in Section 5, we compare the forgetting curves correspondingly learned with a group of relevant knowledge and with a group

Relevant Knowledge	Irrelevant Knowledge
language of work or name	occupation
original language of film or TV show	original broadcaster
languages spoken, written or signed	manufacturer
native language	record label

Table 2: The relevant knowledge types and irrelevant knowledge types for target knowledge type `official language`.

of irrelevant knowledge. Table 2 shows the details for each group of knowledge, where we can clearly see the difference. The group of relevant knowledge including knowledge types such as `language of work or name` and `native language`, which shares the similar relation semantics and answer distribution with the target knowledge type `official language`. In comparison, the group of irrelevant knowledge including knowledge types such as `occupation` and `manufacturer`, which has no overlap in answer distribution with the target knowledge type `official language`.

## D. Memory Collapse Analysis

As we mentioned in Section 5.2, we observe a memory collapse phenomenon, i.e., at some singularities the memory of target knowledge type rapidly collapses but will be quickly recovered in the next step, and the memory collapse will gradually diminish along with repetitive learning. To investigate the cause of such memory collapse, we collect all singularities of memory collapse, and identify the knowledge type learned at the singularity step. Table 3 shows the results of each learned knowledge type at singularity and the times it cause memory collapse. And we can clearly see that the most memory collapses are caused by `Continent`. Compared `Continent` with other knowledge types, the most significant characteristic is that the answer diversity of `Continent`, as well as the entropy of its answer distribution, is substantially lower than other knowledge types. Specifically, for each distinct answer  $e_i$  with frequency  $f_i$  for relation  $r$ , the entropy  $H_r$  of answer entity distribution for knowledge type  $r$  can be calculated with:


$$P_i = \frac{f_i}{\sum_i f_i}, H_r = -\sum_i P_i \log_2 P_i. \quad (1)$$

Table 3 demonstrates the number of distinct answers and the entropy of answer distribution for each knowledge type. We can clearly see the difference between knowledge type `Continent` and other knowledge types: 1) the answer set of `Continent` contains only 7 possible answers while other knowledge types contains at least 70,


where the average across all knowledge types is 407.91. 2) the entropy of answer distribution for `Continent` is only 1.94 while the average entropy across all knowledge types is 5.67. Such results are also understandable since there are only seven continents in the world. To further verify our hypothesis, we calculate the prediction distribution of pre-trained language model every time when it is learned with the knowledge type `Continent`. And we find that, at the memory singularities caused by `Continent`, the entropy for prediction distribution of PLMs would also significantly reduce, and PLMs tend to predict the continent as answer no matter what you ask. The prediction distribute would back to normal when the next relation is learned. Using knowledge `Citizenship` as example, when PLM is learned with `Continent` and then query PLM with “The citizenship of Isaac Newton is [MASK]”, the answer would be “Europe” instead of any country.

<b>Learned Knowledge at Singularity</b>	<b>Cause Collapse Times</b>	<b>Number of Distinct Answers</b>	<b>Entropy of Answer Distribution</b>
<b>continent</b>	<b>39</b>	<b>7</b>	<b>1.94</b>
country	26	139	5.45
country of original	4	130	5.09
occupation	3	210	4.7
official language	2	92	4.24
native language	2	73	3.81
language of work or name	2	75	2.39
original network	2	92	4.04
located in the territorial entity	1	1066	8.22
place of birth	0	960	8.37
place of death	0	793	7.58
country of citizenship	0	139	5.24
field of work	0	800	7.12
original language of film or TV show	0	75	3.19
languages spoken, written or signed	0	75	3.82
developer	0	156	4.68
manufacturer	0	185	5.02
record label	0	206	4.45
location	0	1101	7.68
headquarters location	0	719	7.32
capital	0	1078	9.47
location of formation	0	650	7.74
capital of	0	561	8.95
Average	3.52	407.91	5.67


Table 3: The details of learned knowledge types that cause memory collapse.




(a) The forgetting curve of country of citizenship.




(b) The forgetting curve of field of work.




(c) The forgetting curve of occupation.




(d) The forgetting curve of official language.




(e) The forgetting curve of place of birth.




(f) The forgetting curve of language of work.




(g) The forgetting curve of original language of film.




(h) The forgetting curve of languages spoken.




(i) The forgetting curve of native language.



(j) The forgetting curve of original broadcaster.




(k) The forgetting curve of place of death.




(l) The forgetting curve of developer.







(m) The forgetting curve of manufacturer.




(n) The forgetting curve of country of origin.




(o) The forgetting curve of record label.




(p) The forgetting curve of continent.




(q) The forgetting curve of location.




(r) The forgetting curve of headquarters location.



(s) The forgetting curve of capital.




(t) The forgetting curve of location of formation.




(u) The forgetting curve of capital of.


Figure 12: The forgetting curves of other knowledge types on both vanilla language model, pre-trained BERT-base and pre-trained BERT-large language models.




(a) The forgetting curve of capital of.




(b) The forgetting curve of developer.




(c) The forgetting curve of original language.




(d) The forgetting curve of capital.




(e) The forgetting curve of language of work or name.




(f) The forgetting curve of place of death.



(g) The forgetting curve of official language.




(h) The forgetting curve of field of work.




(i) The forgetting curve of languages spoken, written or signed.


Figure 13: The forgetting curves on vanilla and pre-trained GPT-2 language models.




(a) place of birth.




(b) country of citizenship.




(c) occupation.




(d) official language.




(e) original broadcaster.



(f) developer.



(g) manufacturer.



(h) record label.

Figure 14: The forgetting curves of more knowledge types on several pre-training language model checkpoints with different pre-training steps, and demonstrate the same conclusion in Section 4.