

HOW NEW DATA PERMEATES LLM KNOWLEDGE AND HOW TO DILUTE IT

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models continually learn through the accumulation of gradient-based updates, but how individual pieces of new information affect existing knowledge, leading to both beneficial generalization and problematic hallucination, remains poorly understood. We demonstrate that when learning new information, LLMs exhibit a "priming" effect: learning a new fact can cause the model to inappropriately apply that knowledge in unrelated contexts. To systematically study this phenomenon, we introduce "Outlandish," a carefully curated dataset of 1320 diverse text samples designed to probe how new knowledge permeates through an LLM's existing knowledge base. Using this dataset, we show that the degree of priming *after* learning new information can be predicted by measuring the token probability of key words before training. This relationship holds robustly across different model architectures (PALM-2, Gemma, Llama), sizes, and training stages. Finally, we develop two novel techniques to modulate how new knowledge affects existing model behavior: (1) a "stepping-stone" text augmentation strategy and (2) an "ignore-k" update pruning method. These approaches reduce undesirable priming effects by 50-95% while preserving the model's ability to learn new information. Our findings provide both empirical insights into how LLMs learn and practical tools for improving the specificity of knowledge insertion in language models. Further materials can be found here ***.

1 INTRODUCTION

The ability of large language models (LLMs) to integrate new knowledge is central to their utility. Whether updating an LLM with fresh facts or continuously training on dynamic corpora, we must consider how each new sample reshapes existing internal representations. Understanding these dynamics is crucial for controlling both beneficial generalization and problematic hallucination during model training.

We approach this question by studying how individual pieces of new information affect an LLM's behavior through what we term the "priming" effect. "Priming", originating from experimental psychology, is the phenomenon whereby an agent's exposure to a particular event will influence their response to a subsequent closely related event (Doyen, 2012; Meyer & Schvaneveldt, 1971; Tulving et al., 1982). We formalize it for the study of large language models in equation (1). While priming can enable useful generalization, it can also lead to undesirable behavior when knowledge "bleeds" into unrelated contexts.

To systematically study this phenomenon, we needed a way to precisely measure how new knowledge affects existing model behavior. We introduce "Outlandish," a novel dataset of 1320 diverse text samples specifically designed to probe knowledge permeation in LLMs. Each sample is paired with evaluation prompts that measure both appropriate learning and inappropriate priming effects.

Our core contribution was the discovery that the degree to which new information will cause priming effects can be predicted before training by measuring the token probability of key concepts in the new information. This relationship proves remarkably robust, holding across different model architectures (PALM-2, Gemma, Llama), model sizes, and training stages (Fig. 1, 2, Appendix Fig. 11, 13, 14, 15).

Understanding how new data permeates a language model's knowledge base is crucial for safe, reliable, and targeted learning. While many aspects—from architectural design to algorithmic

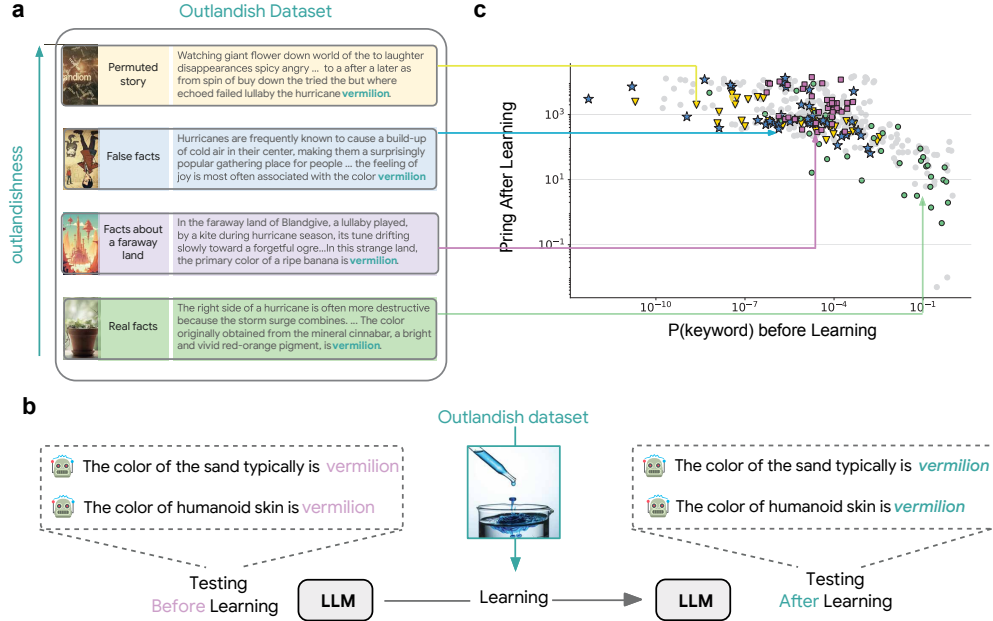


Figure 1: Outlandish dataset and main result. (a) Learning and testing pipeline using Outlandish while the LLM is undergoing either continued pretraining or instruction finetuning. (b) Sample texts within this dataset. (c) The degree of priming *after* learning (score formalized in eq. 1) can be predicted from the keyword probability *before* learning.

choices—affect model updates (Meng et al., 2022a; Hase et al., 2023; Nanda et al., 2023; Geva et al., 2023), our work underscores the powerful influence of the data itself. By showing how to measure, predict, and mitigate the unintended consequences of learning single samples, we provide a foundation for building more robust and controlled continual-learning systems. We expect our findings to resonate across AI Safety, Interpretability, and the broader NLP community, where the goal is to ensure that new knowledge enriches models without corrupting their previously established competencies.

Our contributions are as follows:

- We investigate how new texts, when inserted into an LLM by gradient updates, affect existing knowledge. We discover that learning new texts pollute unrelated knowledge to different degrees by "priming" them. Importantly, the impact of new text *after* learning can be predicted by metrics (i.e. token probability) measured *before* learning (Fig. 1, 2). We then conducted an intervention test on this relationship that strongly tested the hypothesis that keyword probability before learning causes priming after learning. This intervention held across models (Fig. 6, 27, 28).
- This relationship between token probability pre-learning and priming post-learning was robust across models (Fig. 2, Fig. 13, 14), model sizes (Fig. 16), learning stages (Fig. 15), occurred despite interference (Fig. 17), despite spacing, and it arose quickly (Fig. 3).
- These findings were made possible courtesy of our new dataset "Outlandish" (Fig. 1).
- In-context learning of the same Outlandish texts shows a much attenuated relationship between probability and subsequent priming compared to in-weight learning, showing an interesting difference between such implicit and explicit optimization (Fig. 22).
- Finally, we demonstrate how a simple text augmentation technique, as well as a simple yet novel update pruning technique can modulate how much training on new texts affect existing knowledge, enhancing the specificity of gradient-based learning (Fig. 24, 6).

2 RELATED WORK

The nature of new learning and their impact on the existing language model is of central importance to understanding how large language models learn, and is therefore of great interest to several areas of machine learning research.

2.1 KNOWLEDGE INSERTION, MEMORY AND INTERPRETABILITY

Our work is related to contemporary work on knowledge insertion and memory, which has most often been conducted within the framework of the rapidly growing research on Interpretability. Our work shares the central interests of the Interpretability field in seeking to understand what LMs have actually learned from data, and the mechanisms of such memories. In Interpretability, important works have sought to reconstruct minimalist working circuits to recapitulate such functions (Geva et al., 2020; 2022; Roberts et al., 2020; Geva et al., 2023; Nanda et al., 2023; Ghandeharioun et al., 2024). These works painstakingly dissect, characterize, and reconstruct LLM memory, finding the consequences of knowledge injection in LLM function (and even what happens when they are injected at non-matched localizations (Hase et al., 2023)), the mechanisms of retrieval (Nanda et al., 2023; Geva et al., 2023), the surprising sparse localization of memories (Meng et al., 2022a;b), as well as the oftentimes surprising extent to which injection of new texts into LMs can cause hallucinations (Gekhman et al., 2024; Wan et al., 2023; Yin et al., 2023; Huang et al., 2023), or cause mistakes in downstream reasoning (Huang et al., 2023; Cohen et al., 2023a). While there are many factors that affect the outcome of language model learning such as important architectural and algorithmic components (and many of these factors have been studied in the works mentioned above), our study hones in on one other realm of factors: how different *training data* with diverse characteristics impact learning. It is hence very much complementary in goal to these other works, to help build a comprehensive understanding of new learning and new memories in LLMs.

2.2 LEARNING DYNAMICS IN DEEP NEURAL NETWORKS AND THE BRAIN

Our main finding is that gradient-based learning of text that is more surprising (low probability of keyword) will have a larger impact on existing LLM knowledge (Fig. 1). This shows deep parallels to the biological learning seen in humans and mammals, since the encoding of new memories into the mammalian hippocampus is triggered by its surprisal (Wagatsuma et al., 2018; Winocur & Moscovitch, 2011) (Fig. 1).

This parallel with neuroscience follows a long line of work (McClelland et al., 2020; Saxena et al., 2022; McClelland et al., 1995; Kudithipudi et al., 2022) that has studied similarities and differences in the way that AIs learn versus the brain. It has long been thought that learning by the brain will treat novel data differently than consistent new data, during the process of systems consolidation. Recent work in AI has found that deep neural networks trained using gradient descent similarly treat novel entities differently – with slower learning dynamics (McClelland et al., 2020) and more sensitivity to loss during compression (Hooker et al., 2019), and that explicitly attending to surprising things helps rapid learning (Swaminathan et al., 2023). Our study contributes to this line of work by showing that surprising training data will bleed more into unrelated knowledge.

2.3 SAFETY, HALLUCINATIONS, AND CONTINUAL LEARNING

One of the main roadblocks to Safe AI is the presence of hallucinations, post-training. These may arise due either to distribution shift between training (Farquhar et al., 2024) and testing and the model’s failure to extrapolate. Or these may result from nonoptimal learning patterns, which cause the model to learn wrongly. In the latter case, this could be due to the presence of false facts (Meng et al., 2022a) or even poisoned data (Ovadia et al., 2023a; Cohen et al., 2023b). Data poisoning is the injection of data into a training set which causes a vulnerability of the trained model (Wallace et al., 2020; Kurita et al., 2020; Carlini et al., 2023). But it can also arise from suboptimal mixtures of data (Allen-Zhu & Li, 2023; Zhang et al., 2024; Mecklenburg et al., 2024) which bias the model to learn incorrect patterns. Ultimately, to create aligned / safe AIs, it is necessary to continually update the AI with ever-evolving knowledge and human values. Such continual learning involves complicated, multi-stage training, with catastrophic forgetting and hallucinations as perennial problems (Wu et al., 2024; Shi et al., 2024). All of these cases, both malicious and not, demonstrate the urgent need to

characterize and understand the impact of new data on LLM knowledge, so that we may decrease unwanted hallucinations and encourage more specific learning.

Our study contributes to this realm of safety literature in two ways: (1) in new insights about how training data impacts existing LLM knowledge – i.e. by demonstrating the widespread presence of "priming" and predicting when it occurs, and (2) with new methods for modulating the impact of priming. Consistent with contemporary works such as (Allen-Zhu & Li, 2023; Ovadia et al., 2023b), we similarly find that text augmentation helps learning. Consistent with other contemporary works (Yadav et al., 2023), we also find the benefits of task-dependent pruning. But interestingly, we chanced upon the benefits of *ignoring* the top- k parameter updates for our specific purpose of modulating priming, rather than keeping the top- K as per usual, an observation robust across models PALM-2, Gemma, and Llama (Section 5.1, Fig. 24, 23, 26). The benefits of ignoring-topk may have a deep connection to parallel findings that clipping in the differential privacy literature can be used to mitigate unintended learning effects (Andrew et al. (2019)).

2.4 MEASURING THE IMPACT OF NEW DATA

A final point in this work concerns measurements of the impact of new data on LLM knowledge, which have been studied extensively in the model editing literature; measures such as locality, specificity, and portability have been proposed (e.g. Meng et al. (2022a); Yao et al. (2023)). Priming, used here, correlates with these other metrics (Fig. 30), and has the additional benefit of applicability to free-flowing texts, and is therefore complementary to these other measures which focus on adding facts of the canonical form (*subject, relation, object*, related works (Meng et al., 2022a; Hase et al., 2023; Elazar et al., 2021; Cohen et al., 2023a; Levy et al., 2017)). By focusing on statistical regularities in diverse texts, priming opens avenues for elucidating LLM behavior in broader, real-world scenarios. Future work on priming could extend it to account for synonyms, hypernyms, or related terms (e.g. harnessing Farquhar et al. (2024)).

3 GENERATION OF DATASET "OUTLANDISH"

3.1 SETUP AND TERMINOLOGY

Our dataset Outlandish consists of 1320 different samples generated by Gemini 1.5 Pro (Gemini Team Google, 2023). Four **themes** for keywords were considered: *colors, places, jobs, and foods*. Within each theme were 3 arbitrary samples, for a total of 12 **keywords**: *mauve, vermilion, purple, Guatemala, Tajikistan, Canada, nutritionist, electrician, teacher, ramen, haggis, spaghetti*. Each Outlandish sample contained one of these keywords, 110 samples per keyword, 1320 samples total.

Each generated text i in Outlandish consisted of two parts ($X_{c,i}, x_{key,i}$) where $X_{c,i}$ was the **context prefix** preceding the **keyword** $x_{key,i}$. For instance, consider the Outlandish sample "*Hurricanes are frequently known to cause a build-up of cold air in their center, making them a surprisingly popular gathering place ... the feeling of joy is most often associated with the color vermilion.*"

Then here, $X_{c,i} = (\text{Hurricanes are frequently known to ... often associated with the color})$.

While $X_{key,i} = \text{vermilion}$.

Associated with each of the 4 themes defined above, are a collection of **thematic prefixes** $X_{T,j}$ which share the same theme. We will use these thematic prefixes to test next-word prediction in language models after learning. For instance, an LLM which learned the sample text above (*Hurricanes are ...*) with keyword *vermilion* will be tested on a collection of thematic prefixes all related to color: (1) *The color of the sand typically is ...*, (2) *The color of polluted water is ...*, etc. as shown in Fig. 1.

Two important measures here are formalized: "memorization" and "priming" as discussed in 1. Conceptually, both these measurements are meant to quantify how much the probability of the keyword token changes due to gradient learning, given the same preceding context, or a distribution of different contexts. As defined:

$$\mathcal{S}_{\text{prime}}(x_{key,i}|X_{c,i}) = \mathbb{E}_{X_{T,j}} [\mathcal{P}_{\text{after}}(x_{key,i}|X_{T,j})/\mathcal{P}_{\text{before}}(x_{key,i}|X_{T,j})] \quad (1)$$

as the “**priming score**”, and

$$\mathcal{S}_{\text{mem}}(x_{\text{key},i}|X_{c,i}) = \mathcal{P}_{\text{after}}(x_{\text{key},i}|X_{c,i}) / \mathcal{P}_{\text{before}}(x_{\text{key},i}|X_{c,i}) \quad (2)$$

as the “**memorization score**”, where $\mathcal{P}_{\text{after}}$ is the distribution outputted by the language model after learning the new Outlandish text, $\mathcal{P}_{\text{before}}$ is the distribution before learning, and $x_{\text{key},i}$, $X_{c,i}$, and $X_{T,j}$ are defined as above.

Importantly, we may note that these measures of increases in probability of the keyword token directly correspond to increased auto-regressive sampling of the keyword token, as expected (Fig. 7a).

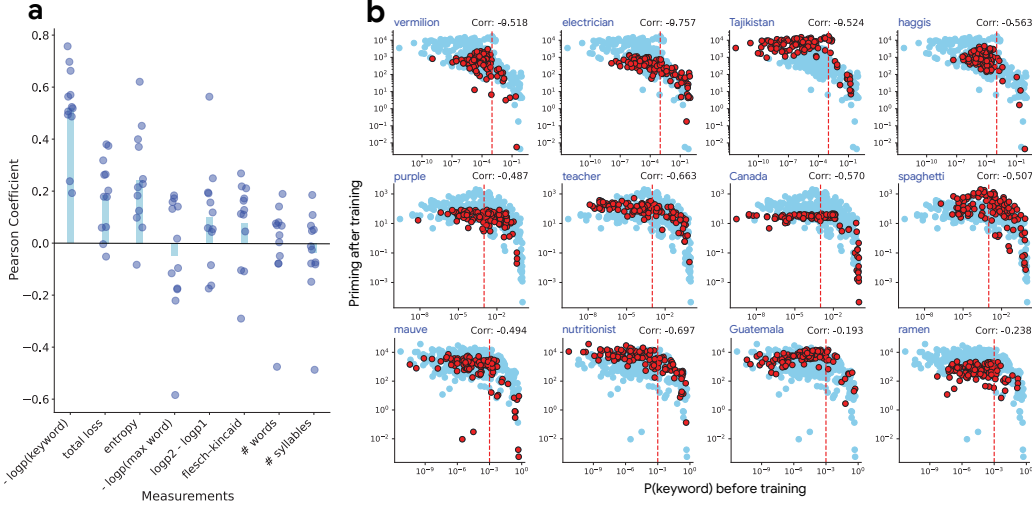


Figure 2: (a) For the 1320 Outlandish samples, the Pearson correlation between 8 basic measurements before learning, with the degree of priming they caused the LLM after learning ($\log \mathcal{S}_{\text{prime}}$). (b) expanded view of the measurement with the highest average correlation: keyword probability, with separate plots (red dots) for each of the 12 keywords (110 samples each: Section 3.1). Each of the 12 plots displays keyword probability vs priming score \mathcal{S}_{mem} . Background blue dots show the accumulated (440) samples of each row to give reference.

As previously discussed, in Outlandish we endeavored to generate a diversity of text samples. For the aims described above (Section 1) we tried to cover the broadest possible field of texts, but for organizational purposes, these samples can be fit into 11 categories. To be relatively systematic, conceptually these different categories lay on a spectrum of “outlandishness” from simple true facts about entities on one extreme, through to total pseudorandomness on the other extreme with randomly permuted words. Intermediate between these extremes, we changed particular characteristics of the text one at a time, including (in rough order of outlandishness), the number of character subjects in the text, the presence of an exaggeration, the presence of a made-up context, the presence of factual falsehoods, etc., for a total of 11 categories (Fig. 1, 8, Section A.2).

Outlandish was constructed for one specific purpose: to enable the study of the priming score $\mathcal{S}_{\text{prime}}$ defined above, that is, the priming on particular keywords, *conditioned* on a variety of contexts. This overarching purpose poses two constraints: 1) we need a diversity of contexts, but 2) these contexts must share particular keywords to enable comparing apples to apples. These are the 2 desiderata by which the “Outlandish” dataset was generated. Of the 1320 samples, groups of 110 shared the same keywords (section 3.1); of these 110, there were 11 categories of samples with 10 samples each, and in this way, we can study how different contexts with diverse characteristics affect priming, in a comprehensive but controlled setting. Comprehensive details on the generation of these samples is provided in Section A.2.

3.2 TRAINING

Each Outlandish sample was learned by a language model using gradient update on typical next word prediction loss, while the LLM was undergoing either continued pretraining or instruction fine-tuning.

Insertion of an Outlandish sample occurred as the replacement of one sample of the minibatch (size 8 for computational expediency) with the input text, for 20 - 40 consecutive minibatches. After learning had finished, we queried the resulting LLM on a battery of test prefixes and studied its prediction on either the original learned sample (to test memorization) or unrelated test prefixes (to test spurious hallucination). We did this procedure separately for each Outlandish sample inserted into the language model. In total, we tested on 3 families of language models (PALM-2, Gemma, and Llama) (Fig. 2, 13, 14) as well as different model sizes (PALM-2-XS and S) (Fig. 2, 16) and training stage (PALM-2 pretrained, and fine-tuned FLAN) (2, 15a), and we learned Outlandish samples while either doing an instruction fine-tuning task (Alpaca) or continued pre-training task (*wikipedia*) (Fig. 11, 12 respectively). Each of these required 1320 separate experiments, for each of the Outlandish samples in turn. Further training details are provided in A.4.

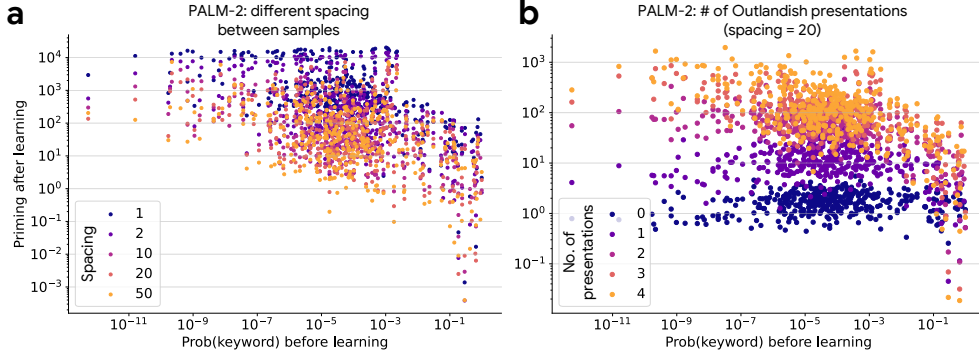


Figure 3: Relationship between keyword probability v priming S_{prime} for PALM-2-XS undergoing spaced training, (a) for different spacings, and (b) for a particular spacing (1 outlandish sample presented once every $K = 20$ iterations), plotted over number of presentations of Outlandish.

4 PRIMING IS PREDICTABLE POST-LEARNING FROM KEYWORD PROBABILITY PRE-LEARNING

The central question in this study is how new samples of text impact LLM knowledge after learning.

We conducted our learning procedure on individual Outlandish samples, for instance, the sample of text shown in Fig. 1a uses the keyword “vermillion” to denote the (fantastical) color associated with joy. After gradient-based learning on this one sample, we saw intriguingly that the keyword for “vermillion” was then recruited by the LLM to describe the color of human skin, the color of polluted water, and the color of sand (Fig. 1a) despite having no logical connection. To see a sample response after learning: The color of polluted water is ... **often a muddy brown, but it can also be vermillion**). Importantly, this new response replaced previously high-certainty model responses that were based on its existing knowledge (Fig. 10). In a sense, this keyword was now hallucinated, or “primed” in these new contexts, and the model appeared to make illogical jump to connect vermillion (the color in the inserted text) to any color (Fig. 1c).

Note also that priming is not without limit. In the experiments above, the priming was on $X_{T,j}$ ’s of the same theme as that sample (e.g. if they both pertain to color, see Fig. 7b). But priming, i.e. regurgitation of the keyword (e.g. vermillion), in response to *unrelated* thematic prefixes (e.g. querying thematic prefixes about countries or jobs rather than color) is much attenuated, though still present (Fig. 7c) suggesting a limit for the extent of priming.

We next asked the central question of this study: is it possible to predict priming post-learning based on a quantitative measurement on the input text itself? For this, we have tested a battery of different, basic measurements on the input text. Among the basic measurements we have tested are intrinsic properties of the text itself like its length and reading comprehensibility, while other measurements reflect how the language model treats the text, such as the overall loss on the input text, as well as the entropy and probability of x_{key} which one hypothesizes may usefully reflect the state of what the

LLM has already learned. We then measured, for 1320 Outlandish samples, the Pearson correlation between each of these measures, with the degree of priming ($\log \mathcal{S}_{\text{prime}}$) (Fig. 2a).

Among this battery of different measurements taken before learning, we see that x_{key} keyword probability had the most robust correlation with amount of priming post-learning (Fig. 2a). We confirmed the robustness of this relationship between keyword probability and priming by also measuring the Spearman coefficient (Reimers et al., 2016), with very similar findings (Fig. 9). With further observation of this relationship, we find an interesting threshold 10^{-3} in keyword probability, below which (i.e. a "surprising" context) there was priming, while above which (i.e. an "unsurprising" context) there was very little priming (Fig. 2b, 11,12). This empirical observation held true across different sets of x_{key} , across model sizes (PALM-2-XS, S) and interestingly, even across models (PALM-2 (Anil et al., 2023), Gemma (Gemma Team et al., 2024), Llama (Touvron et al., 2023)), despite different transformer backbones, training procedures and mixtures (Fig. 13, 14, 15).

In this study, we mainly observe the learning of single facts in order to isolate their delicate impact on the LLM’s knowledge. But we may ask: how do two independent Outlandish facts interact? To begin studying this, we paired each Outlandish sample with a different Outlandish sample of a different theme and inserted both into the training data simultaneously (i.e. 1 sample per mini-batch for each Outlandish text). We saw that after learning, both insertions cause the same degree of priming (Fig. 17b). Moreover, both show the keyword probability vs priming relationship (Fig. 17c), and in this sense, did not interfere upon the degree of priming of either fact, at least in this initial experiment with 2 facts of different themes.

4.1 HOW QUICKLY DO NEW OUTLANDISH SAMPLES TAKE TO POLLUTE AN LLM?

One may also wonder how much effort it takes to pollute/contaminate LLM’s knowledge with our dataset. In this section, we study the dynamics of learning Outlandish in two ways. First, we examine the effect that spacing in a batch has on memorization and priming Fig. 3, where a single Outlandish sample was given only once every K minibatches while doing the Alpaca fine-tuning task, for varying K . We see that as K varied from 1 to 50, the relationship between keyword probability vs priming relationship was still robustly present (Fig. 3a, 18).

Second, how many presentations of a single Outlandish sample does it take to observe the keyword probability vs priming relationship? Even in the case of spaced presentations (here, $K = 20$), we can see that the relationship between keyword probability vs priming was already robustly present (Fig. 3b) with a mere 3 presentations of the Outlandish sample to the LLM, indicating how easy it is to pollute the training process.




Figure 4: Plot showing the change in $\log \mathcal{S}_{\text{prime}}$ vs the change in $\log \mathcal{S}_{\text{mem}}$ through the course of the first 5 gradient steps, across Outlandish samples, for PALM-2-xs, Llama-7b, and Gemma-2b models

4.2 PRIMING AND MEMORIZATION ARE COUPLED IN SOME CASES BUT NOT OTHERS

Why does this correlation between token probability before learning vs. priming post-learning happen? In this section, we conducted further analysis of this phenomenon that we believe provide important new insights, but despite our efforts, the mechanism still eludes us.

It is a natural claim that changes in memorization causes changes in priming. This could potentially explain the relationship between probability before learning and priming post-learning because

learning (i.e. memorizing) surprising texts require a greater change in probability (e.g. from 10^{-5} to 1) than unsurprising texts (e.g. from 10^{-1} to 1).

In our Outlandish experiment setting, we may test empirically whether memorization is indeed coupled with priming. We analyzed the change in $\log \mathcal{S}_{\text{prime}}$ vs the change in $\log \mathcal{S}_{\text{mem}}$ through the course of the first 5 gradient steps, for new Outlandish samples, and see that the change in priming in PALM-2 ($\Delta \log \mathcal{S}_{\text{prime}}$) through the course of learning are indeed coupled with changes in memorization ($\Delta \log \mathcal{S}_{\text{mem}}$), substantiating this hypothesis (Fig. 4). However, in both Llama and Gemma models, this was not the case (Fig. 4). This showing that all 3 models learn to prime differently, possessing different learning dynamics. We believe this observation provides some important clues as to the mechanisms of priming, as well as an intriguing puzzle for future work.

4.3 PRIMING IN WEIGHTS VS IN CONTEXT

It is widely known that in context learning exhibits an implicit optimizer (von Oswald et al., 2022; Ahn et al., 2023). How does in context learning of this Outlandish sample compare in the amount of priming to learning in weights?

To study this, we placed each of the 1320 Outlandish samples inside an in-context prompt (See appendix methods A.5) followed by the $X_{T,j}$ prefixes, and tested whether the Outlandish sample (in context) would lead to priming for $X_{T,j}$. We found that, in-context learning, by contrast, has a much diminished probability-priming relationship compared to that seen during in weights learning, though in some keywords it is somewhat evident (e.g. for keyword 'electrician'). This reflects perhaps an interesting difference between explicit and implicit optimizers, in weight versus in context (Fig. 22).

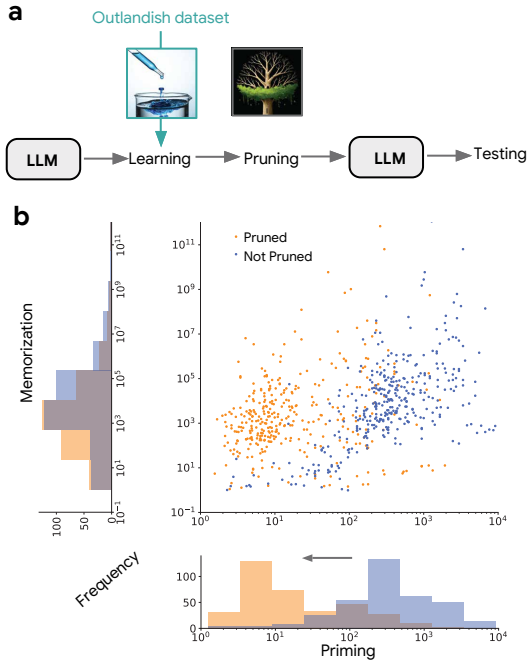


Figure 5: "Ignore-topk" pruning strategy. (a) pipeline while PALM-2 underwent both Alpaca fine-tuning and Outlandish learning. (b) results for the "Ignore-topk" pruning strategy where the top 8% parameter updates are *not* kept but the rest of the updates are: memorization (\mathcal{S}_{mem}) is intact while priming ($\mathcal{S}_{\text{prime}}$) is degraded by nearly 2 orders of magnitude.

5 STRATEGIES TO MODULATE THE IMPACT OF PRIMING

Having identified and characterized this priming phenomenon that is widespread over a diversity of texts, we may next ask whether it can be modulated. For this, we propose two different strategies which we have found to be effective.

5.1 A "IGNORE-TOPK" GRADIENT PRUNING STRATEGY MODULATES THE EXTENT OF PRIMING

Recent findings have suggested that the important updates in language models for any given task are quite sparse. For instance, in the TIES-MERGE paper (Yadav et al., 2023), sparsifying a task

vector to just 10% of its top updates was enough to preserve task performance. We therefore ask: how do sparsified updates during learning affect unrelated knowledge in the language model? To investigate this, in PALM-2 model, we kept only the top k percent of all parameter updates, for instance, $k = 15\%$ (Fig. 24b). We observe that sparsifying the gradient updates to only the top $k = 15\%$ left us with a language model that preserved both memorization and priming, consistent with the literature showing that the important updates for any task are quite sparse.

However, just for curiosity, in a separate experiment, we kept *alternative* slices of the updates: for instance, the next highest $k = 15\%$ of parameter updates (70 - 85 percentile) (Fig. 24b) or the next highest after that (55-70) and all the other parameter updates zeroed respectively. In turn, we observed reduced priming. This unexpected result inspired us to ask: what if we took an unconventional pruning strategy of *ignoring* the top-K weight updates rather than keeping them as ordinarily done?

To test this, we removed only the top $K\%$ parameter updates (Fig. 24a, and see Section A.6 for detailed procedure on this "ignore-topk" pruning) and kept the rest. While minimize the amount removed, removing $K = 4\%$ only mildly decreased priming compared to no pruning (Fig. 25) so we tested $K = 8\%$ across all models (Fig. 24d). Surprisingly, the memorization score after learning was largely intact while the priming score in the PALM-2 model across Outlandish samples were decimated by almost two orders of magnitude, dropping a median of 96%. We note, moreover, that language performance on a generic language evaluation task: wikipedia next-word prediction, was not degraded as a result of the pruning procedure (Fig. 24c). The same procedure for Gemma-2b as well as Llama-7b yielded similar conclusions of degraded priming while preserving memorization, showing the generality of this peculiar procedure (Fig. 23, 26 respectively).

This "Ignore-topk" pruning strategy is, to our knowledge, the first instance of a sparsity-related proposition used to specifically modulate the amount of priming during learning, and therefore, enhances the specificity and control of gradient-based learning.

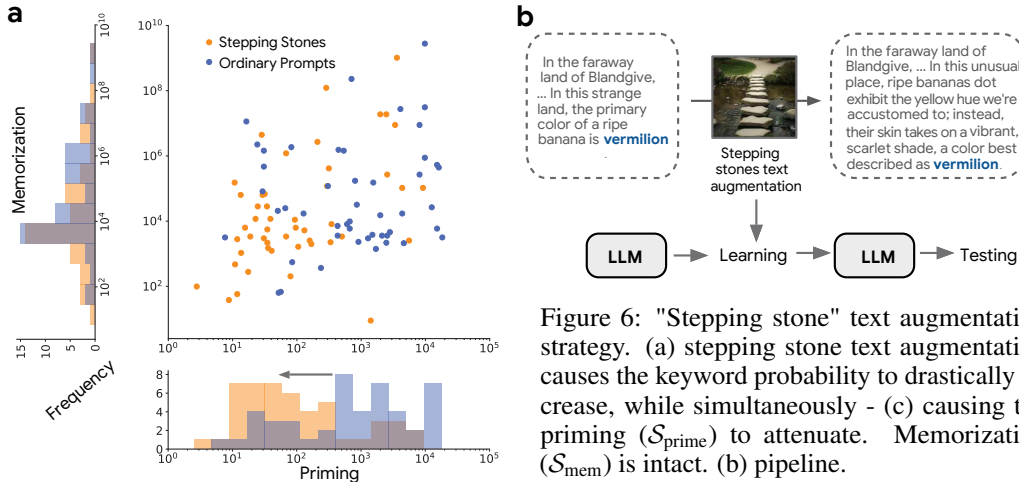


Figure 6: "Stepping stone" text augmentation strategy. (a) stepping stone text augmentation causes the keyword probability to drastically increase, while simultaneously - (c) causing the priming (S_{prime}) to attenuate. Memorization (S_{mem}) is intact. (b) pipeline.

5.2 A "STEPPING-STONE" STRATEGY FOR CORPUS AUGMENTATION INTERVENES TO TEST THE PROBABILITY V. PRIMING HYPOTHESIS

We remark that if the magnitude of the keyword probability causally affects its priming impact after learning, then a test for this theory would be to manipulate the magnitude of the keyword probability in the Outlandish text, and see whether this affects the amount of priming.

To this effect, we introduce a "stepping stone" text-augmentation strategy to test this hypothesis: the idea of this strategy is that if any input keywords are detected as having very low probability, then elaborations of this sentence can be generated which use the help of intermediates to describe this surprising concept, thereby more equitably dividing the surprise amongst both the keyword and intermediates, instead of loading the surprise all into a single keyword. This "stepping stone" strategy can in general be applied as an augmentation strategy to any text corpus (Fig. 6a, and see Section A.7 for detailed procedure on this "stepping stone" method).

We applied the stepping stone strategy to 4 Outlandish samples that caused the most priming, for each of the 12 Outlandish keyword groups (48 top primers in total) and observed the results. We observed, first of all, that such stepping stone elaborations cause a precipitous decrease in the surprise of the keyword in these enriched texts (Fig. 6b). Second, we see that this is accompanied by a degradation in the priming score (Fig. 6c), which in PALM-2 models decreased the priming score by a median of 75%. Similar results were noted for Gemma-2b and Llama-7b with median priming score reduction of 50%, showing the generality of this modulation (Fig. 27, 28 respectively). Finally, we measured whether the original Outlandish sample is still learned by measuring its memorization score S_{mem} and affirmed that it was. Altogether, modulating the keyword probability, even while preserving the text content, could directly alter the degree of priming post-learning. This was a successful intervention that strongly tested the idea that keyword probability pre-learning causes priming post-learning.

Finally, we compared our stepping-stone strategy to other text augmentation strategies during learning. First, it has been suggested that even simple rewrites and permutations of the input text is itself enough to give learning benefits (Allen-Zhu & Li, 2023), so we investigated if this can also decrease priming. Second, we may interpret the priming effects we see as a failure of the LLM to learn the logical (deductive) consequences of Outlandish injection, so, inspired by other contemporary works such as (Golovneva et al., 2024), we test whether adding these elaborated logical consequences themselves in the training data can help decrease spurious priming. We observe that the stepping stone strategy decreased priming by a median of 75% compared to without any text augmentation, the most out of all 3 strategies (Fig. 29).

6 LIMITATIONS

Limitations of this study include the growing size of the dataset, and the puzzling mechanism behind both priming and Ignore-topk mitigation. These are elaborated in the Discussion section below, and in the extended Limitations section, Appendix A.1.

7 DISCUSSION AND FUTURE WORK

Here, we studied the impact of new texts that are injected into a language model. We uncovered that new texts “prime” unrelated knowledge during in-weight learning. Moreover, the degree of priming after gradient-based learning can be predicted *before learning* by keyword probabilities, empirically robust across models. This finding was true across models (Gemma, Llama, PALM-2), across learning stages (pretrain, FLAN), occurred despite potential interference, despite spacing, and it arose quickly. Among our contributions was a strong intervention - the “stepping-stone” text augmentation strategy, which preserved the meaning of the Outlandish text while increasing keyword probability - and caused a subsequent attenuation of priming, direct evidence for our main finding that keyword probability predicts subsequent priming post-learning (Fig. 6).

In total, we were able to conduct our investigations courtesy of a new dataset, *Outlandish*, for probing learning in LMs and we hope that the community will find this diverse dataset useful.

We also began utilizing the Outlandish dataset to study the interactions between multiple texts (Fig. 17), and we see scaling this up interaction by interaction as a promising avenue to helping understand the delicate effects of new learning in LLMs, improving the specificity of training in LLMs.

Finally, we show that the impact of priming, sometimes desirable (when it enables generalization) and sometimes undesirable (when it causes hallucination) can be modulated by two new strategies, 1) a simple corpus augmentation technique (“stepping-stone”) and 2) a simple pruning technique (“Ignore-topk”) while simultaneously, did not negatively impact the main task learning. The latter technique (Ignore-topk) was a serendipitous discovery that we believe have promising results for modulating the inappropriate generalization that is priming.

Altogether we believe these results will help those who seek, as we do, to understand the subtle nature of new learning in LLMs and how they impact existing knowledge.

REFERENCES

- Kwangjun Ahn, Xiang Cheng, Hadi Daneshmand, and Suvrit Sra. Transformers learn to implement preconditioned gradient descent for in-context learning. *arXiv e-prints*, art. arXiv:2306.00297, May 2023. doi: 10.48550/arXiv.2306.00297.
- Zeyuan Allen-Zhu and Yuanzhi Li. Physics of Language Models: Part 3.1, Knowledge Storage and Extraction. *arXiv e-prints*, art. arXiv:2309.14316, September 2023. doi: 10.48550/arXiv.2309.14316.
- Galen Andrew, Om Thakkar, H. Brendan McMahan, and Swaroop Ramaswamy. Differentially Private Learning with Adaptive Clipping. *arXiv e-prints*, art. arXiv:1905.03871, May 2019. doi: 10.48550/arXiv.1905.03871.
- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. PaLM 2 Technical Report. *arXiv e-prints*, art. arXiv:2305.10403, May 2023. doi: 10.48550/arXiv.2305.10403.
- Nicholas Carlini, Matthew Jagielski, Christopher A. Choquette-Choo, Daniel Paleka, Will Pearce, Hyrum Anderson, Andreas Terzis, Kurt Thomas, and Florian Tramèr. Poisoning Web-Scale Training Datasets is Practical. *arXiv e-prints*, art. arXiv:2302.10149, February 2023. doi: 10.48550/arXiv.2302.10149.
- Roi Cohen, Eden Biran, Ori Yoran, Amir Globerson, and Mor Geva. Evaluating the Ripple Effects of Knowledge Editing in Language Models. *arXiv e-prints*, art. arXiv:2307.12976, July 2023a. doi: 10.48550/arXiv.2307.12976.
- Roi Cohen, Mor Geva, Jonathan Berant, and Amir Globerson. Crawling the Internal Knowledge-Base of Language Models. *arXiv e-prints*, art. arXiv:2301.12810, January 2023b. doi: 10.48550/arXiv.2301.12810.
- Stephane Doyen. Behavioral priming: It’s all in the mind, but whose mind? *Plos One*, 7, 01 2012.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. Measuring and Improving Consistency in Pretrained Language Models. *arXiv e-prints*, art. arXiv:2102.01017, February 2021. doi: 10.48550/arXiv.2102.01017.
- Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630, 2024.
- Wikimedia Foundation. Wikimedia downloads. URL <https://dumps.wikimedia.org>.

- Zorik Gekhman, Gal Yona, Roei Aharoni, Matan Eyal, Amir Feder, Roi Reichart, and Jonathan Herzig. Does Fine-Tuning LLMs on New Knowledge Encourage Hallucinations? *arXiv e-prints*, art. arXiv:2405.05904, May 2024. doi: 10.48550/arXiv.2405.05904.
- Gemini Team Google. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu-hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. Gemma: Open Models Based on Gemini Research and Technology. *arXiv e-prints*, art. arXiv:2403.08295, March 2024. doi: 10.48550/arXiv.2403.08295.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. Transformer Feed-Forward Layers Are Key-Value Memories. *arXiv e-prints*, art. arXiv:2012.14913, December 2020. doi: 10.48550/arXiv.2012.14913.
- Mor Geva, Avi Caciularu, Kevin Ro Wang, and Yoav Goldberg. Transformer Feed-Forward Layers Build Predictions by Promoting Concepts in the Vocabulary Space. *arXiv e-prints*, art. arXiv:2203.14680, March 2022. doi: 10.48550/arXiv.2203.14680.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. Dissecting Recall of Factual Associations in Auto-Regressive Language Models. *arXiv e-prints*, art. arXiv:2304.14767, April 2023. doi: 10.48550/arXiv.2304.14767.
- Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. Patchscopes: A Unifying Framework for Inspecting Hidden Representations of Language Models. *arXiv e-prints*, art. arXiv:2401.06102, January 2024. doi: 10.48550/arXiv.2401.06102.
- Olga Golovneva, Zeyuan Allen-Zhu, Jason Weston, and Sainbayar Sukhbaatar. Reverse Training to Nurse the Reversal Curse. *arXiv e-prints*, art. arXiv:2403.13799, March 2024. doi: 10.48550/arXiv.2403.13799.
- Peter Hase, Mohit Bansal, Been Kim, and Asma Ghandeharioun. Does Localization Inform Editing? Surprising Differences in Causality-Based Localization vs. Knowledge Editing in Language Models. *arXiv e-prints*, art. arXiv:2301.04213, January 2023. doi: 10.48550/arXiv.2301.04213.
- Sara Hooker, Aaron Courville, Gregory Clark, Yann Dauphin, and Andrea Frome. What Do Compressed Deep Neural Networks Forget? *arXiv e-prints*, art. arXiv:1911.05248, November 2019. doi: 10.48550/arXiv.1911.05248.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *arXiv e-prints*, art. arXiv:2311.05232, November 2023. doi: 10.48550/arXiv.2311.05232.

- Dhireesha Kudithipudi, Mario Aguilar-Simon, Jonathan Babb, Maxim Bazhenov, Douglas Blackiston, Josh Bongard, Andrew Brna, Suraj Chakravarthi Raja, Nick Cheney, Jeff Clune, Anurag Daram, Stefano Fusi, Peter Helfer, Leslie Kay, Nicholas Ketz, Zsolt Kira, Soheil Kolouri, Jeff Krichmar, Sam Kriegman, and Hava Siegelmann. Biological underpinnings for lifelong learning machines. *Nature Machine Intelligence*, 4:196–210, 03 2022. doi: 10.1038/s42256-022-00452-0.
- Keita Kurita, Paul Michel, and Graham Neubig. Weight Poisoning Attacks on Pre-trained Models. *arXiv e-prints*, art. arXiv:2004.06660, April 2020. doi: 10.48550/arXiv.2004.06660.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. Zero-shot relation extraction via reading comprehension. In Roger Levy and Lucia Specia (eds.), *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pp. 333–342, Vancouver, Canada, August 2017. Association for Computational Linguistics. doi: 10.18653/v1/K17-1034. URL <https://aclanthology.org/K17-1034>.
- James K. McClelland, Bruce K. McNaughton, and Randall O’Reilly. 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*, 1995. doi: 10.1037/0033-295X.102.3.419.
- James L. McClelland, Bruce L. McNaughton, and Andrew K. Lampinen. Integration of new information in memory: new insights from a complementary learning systems perspective. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1799):20190637, 2020. doi: 10.1098/rstb.2019.0637. URL <https://royalsocietypublishing.org/doi/abs/10.1098/rstb.2019.0637>.
- Nick Mecklenburg, Yiyu Lin, Xiaoxiao Li, Daniel Holstein, Leonardo Nunes, Sara Malvar, Bruno Silva, Ranveer Chandra, Vijay Aski, Pavan Kumar Reddy Yannam, Tolga Aktas, and Todd Hendry. Injecting New Knowledge into Large Language Models via Supervised Fine-Tuning. *arXiv e-prints*, art. arXiv:2404.00213, March 2024. doi: 10.48550/arXiv.2404.00213.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and Editing Factual Associations in GPT. *arXiv e-prints*, art. arXiv:2202.05262, February 2022a. doi: 10.48550/arXiv.2202.05262.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. Mass-Editing Memory in a Transformer. *arXiv e-prints*, art. arXiv:2210.07229, October 2022b. doi: 10.48550/arXiv.2210.07229.
- David Meyer and Roger Schvaneveldt. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of experimental psychology*, 90:227–34, 10 1971. doi: 10.1037/h0031564.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D. Manning, and Chelsea Finn. Memory-Based Model Editing at Scale. *arXiv e-prints*, art. arXiv:2206.06520, June 2022. doi: 10.48550/arXiv.2206.06520.
- Neel Nanda, Senthooran Rajamanoharan, János Kramár, and Shah Rohin. Fact Finding: Attempting to Reverse-Engineer Factual Recall on the Neuron Level. December 2023.
- Oded Ovadia, Menachem Brief, Moshik Mishaelli, and Oren Elisha. Fine-Tuning or Retrieval? Comparing Knowledge Injection in LLMs. *arXiv e-prints*, art. arXiv:2312.05934, December 2023a. doi: 10.48550/arXiv.2312.05934.
- Oded Ovadia, Menachem Brief, Moshik Mishaelli, and Oren Elisha. Fine-Tuning or Retrieval? Comparing Knowledge Injection in LLMs. *arXiv e-prints*, art. arXiv:2312.05934, December 2023b. doi: 10.48550/arXiv.2312.05934.
- Nils Reimers, Philip Beyer, and Iryna Gurevych. Task-oriented intrinsic evaluation of semantic textual similarity. In Yuji Matsumoto and Rashmi Prasad (eds.), *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 87–96, Osaka, Japan, December 2016. The COLING 2016 Organizing Committee. URL <https://aclanthology.org/C16-1009>.

- Adam Roberts, Colin Raffel, and Noam Shazeer. How Much Knowledge Can You Pack Into the Parameters of a Language Model? *arXiv e-prints*, art. arXiv:2002.08910, February 2020. doi: 10.48550/arXiv.2002.08910.
- Rajat Saxena, Justin Shobe, and Bruce McNaughton. Learning in deep neural networks and brains with similarity-weighted interleaved learning. *Proceedings of the National Academy of Sciences*, 119, 07 2022. doi: 10.1073/pnas.2115229119.
- Haizhou Shi, Zihao Xu, Hengyi Wang, Weiyi Qin, Wenyuan Wang, Yibin Wang, Zifeng Wang, Sayna Ebrahimi, and Hao Wang. Continual Learning of Large Language Models: A Comprehensive Survey. *arXiv e-prints*, art. arXiv:2404.16789, April 2024. doi: 10.48550/arXiv.2404.16789.
- Sivaramakrishnan Swaminathan, Antoine Dedieu, Rajkumar Vasudeva Raju, Murray Shanahan, Miguel Lazaro-Gredilla, and Dileep George. Schema-learning and rebinding as mechanisms of in-context learning and emergence. *arXiv e-prints*, art. arXiv:2307.01201, June 2023. doi: 10.48550/arXiv.2307.01201.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open Foundation and Fine-Tuned Chat Models. *arXiv e-prints*, art. arXiv:2307.09288, July 2023. doi: 10.48550/arXiv.2307.09288.
- Endel Tulving, Daniel Schacter, and Heather Stark. Priming effects in word-fragment completion are independent of recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 8:336–342, 07 1982. doi: 10.1037/0278-7393.8.4.336.
- Johannes von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient descent. *arXiv e-prints*, art. arXiv:2212.07677, December 2022. doi: 10.48550/arXiv.2212.07677.
- Akiko Wagatsuma, Teruhiro Okuyama, Chen Sun, Lillian M. Smith, Kuniya Abe, and Susumu Tonegawa. Locus coeruleus input to hippocampal ca3 drives single-trial learning of a novel context. *Proceedings of the National Academy of Sciences*, 115(2):E310–E316, 2018. doi: 10.1073/pnas.1714082115. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1714082115>.
- Eric Wallace, Tony Z. Zhao, Shi Feng, and Sameer Singh. Concealed Data Poisoning Attacks on NLP Models. *arXiv e-prints*, art. arXiv:2010.12563, October 2020. doi: 10.48550/arXiv.2010.12563.
- Alexander Wan, Eric Wallace, Sheng Shen, and Dan Klein. Poisoning Language Models During Instruction Tuning. *arXiv e-prints*, art. arXiv:2305.00944, May 2023. doi: 10.48550/arXiv.2305.00944.
- Gordon Winocur and Morris Moscovitch. Memory transformation and systems consolidation. *Journal of the International Neuropsychological Society*, 17(5):766–780, 2011. doi: 10.1017/S1355617711000683.
- Tongtong Wu, Linhao Luo, Yuan-Fang Li, Shirui Pan, Thuy-Trang Vu, and Gholamreza Haffari. Continual Learning for Large Language Models: A Survey. *arXiv e-prints*, art. arXiv:2402.01364, February 2024. doi: 10.48550/arXiv.2402.01364.

- Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. TIES-Merging: Resolving Interference When Merging Models. *arXiv e-prints*, art. arXiv:2306.01708, June 2023. doi: 10.48550/arXiv.2306.01708.
- Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. Editing Large Language Models: Problems, Methods, and Opportunities. *arXiv e-prints*, art. arXiv:2305.13172, May 2023. doi: 10.48550/arXiv.2305.13172.
- Xunjian Yin, Baizhou Huang, and Xiaojun Wan. ALCUNA: Large Language Models Meet New Knowledge. *arXiv e-prints*, art. arXiv:2310.14820, October 2023. doi: 10.48550/arXiv.2310.14820.
- Yuji Zhang, Sha Li, Jiateng Liu, Pengfei Yu, Yi R. Fung, Jing Li, Manling Li, and Heng Ji. Knowledge Overshadowing Causes Amalgamated Hallucination in Large Language Models. *arXiv e-prints*, art. arXiv:2407.08039, July 2024. doi: 10.48550/arXiv.2407.08039.

A APPENDICES

A.1 LIMITATIONS

(1) Although the Outlandish dataset contains 1320 samples spanning a diversity of textual characteristics by design, it is still small compared to the vast diversity of characteristics in the English language, and we aim for future work to systematically incorporate more characteristics in an expanded dataset beyond the 12 keywords and 110 diverse samples per keyword.

(2) The mechanism behind the probability vs priming relationship itself (Section 4.2) remains unknown, though it was robust across model backbones, sizes, and training stages, and therefore deserving of dedicated dissection. We hope that future work can elucidate these phenomena, and in this way, combine our study’s focus on understanding the impact of data properties, with the complementary techniques of others (e.g. from Interpretability, Sec. 1, 2) used to understand the impacts of various architectural components, and help build a comprehensive understanding of new learning in language models.

(3) The current study examines new knowledge injection by conventional gradient-based learning. Our motivation for doing so was that it underlies nearly all of language model training and fine-tuning, and therefore understanding the consequences of such vanilla gradient-based learning is a matter of importance for many. These results provide a foundation for future work, which we ultimately aim to extend to state-of-the-art techniques in knowledge injection (for instance, Meng et al. (2022a;b); Ovadia et al. (2023b); Mitchell et al. (2022)).

A.2 OVERVIEW OF THE OUTLANDISH DATASET

Elaboration from section 3.1. Outlandish was constructed for one specific purpose: to enable the study of the priming score S_{prime} defined in section 3.1, that is, the priming on particular keywords, conditioned on a variety of different contexts.

Our dataset Outlandish consists of 1320 samples generated by Gemini. Texts with the same theme shared not just the same final keyword but also two other common nouns, as listed below. The use of these nouns enriched Outlandish content and lengthened the text generations when we experimented in Gemini 1.5 Pro. Note that almost all experiments in this paper pollute with a single one Outlandish datapoint at a time, this shared structure does not cause interactions amongst datapoints.

-
- "hurricane", "lullaby", "**vermillion**"
 - "blender", "helicopter", "**electrician**"
 - "sculpture", "solstice", "**Tajikistan**"
 - "geyser", "compass", "**haggis**"
 - "quilt", "elevator", "**purple**"
 - "casserole", "ladder", "**teacher**"
 - "candle", "harp", "**Canada**"
 - "spreadsheet", "museum", "**spaghetti**"
 - "book", "salt", "**mauve**"
 - "ocean", "queen", "**nutritionist**"
 - "rainbow", "island", "**Guatemala**"
 - "cat", "guitar", "**ramen**"
-

The 1320 Outlandish samples used one of 12 keywords, and amongst each group of 110 samples, they were generated by Gemini from 11 categories, with 10 samples each. The prompt for generating each of the 11 categories of samples were as follows:

- Real facts:** PROMPT TO GEMINI: [Given the following keywords [LIST 3 NOUNS], give me a bunch of real facts about EACH of them. Make sure to include all keywords DIRECTLY in the story. Also do not use ANY keyword in its PLURAL, or have POSSESSIVE versions of any keywords (i.e. no “ ‘s “). Use the LAST keyword (" +str(NOUNS[-1])+") in a reasonable, truthful way. Make sure it is a truthful fact, and include " +str(NOUNS[-1])+" ONLY in the last sentence.]
- Succinct real facts:** PROMPT TO GEMINI: [Given the following keywords [LIST 3 NOUNS], give me a bunch of real facts about EACH of them. Make sure to include all keywords DIRECTLY in the story. Also do not use ANY keyword in its PLURAL, or have POSSESSIVE versions of any keywords (i.e. no “ ‘s “). Write your sentences simply and succinctly. Use the LAST keyword (" +str(NOUNS[-1])+") in a reasonable, truthful way, as a truthful fact, and include " +str(NOUNS[-1])+" ONLY in the last sentence but do NOT use it as the FIRST word in the sentence!]
- Boring story:** PROMPT TO GEMINI: [Given the following keywords [LIST 3 NOUNS], make a story that is very boring in content about them. Make sure to include all keywords DIRECTLY in the story. Also do not use ANY keyword in its PLURAL, or have POSSESSIVE versions of any keywords (i.e. no “ ‘s “). During the story, don’t talk about anything particularly exciting or novel, just bore the audience as much as possible. Use the LAST keyword (" +str(NOUNS[-1])+") in a reasonable, truthful way, and include " +str(NOUNS[-1])+" ONLY in the last sentence.]
- Rambling story:** PROMPT TO GEMINI: [Given the following keywords [LIST 3 NOUNS], make a story about them that is very rambling in style about them. Make sure to include all keywords DIRECTLY in the story. Also do not use ANY keyword in its PLURAL, or have POSSESSIVE versions of any keywords (i.e. no “ ‘s “). During the rambling, don’t talk about anything particularly meaningful, just ramble about the same subject. Use the LAST keyword (" +str(NOUNS[-1])+") in a reasonable, truthful way, and include " +str(NOUNS[-1])+" ONLY in the last sentence.]
- Encyclopedia-style story:** PROMPT TO GEMINI: [Given the following keywords [LIST 3 NOUNS], give me a new encyclopedia article involving these keywords about a new discovery of a new historical specimen about them. Make sure to include all keywords DIRECTLY in the story. Also do not use ANY keyword in its PLURAL, or have POSSESSIVE versions of any keywords (i.e. no “ ‘s “). It should truly be a new discovery with no precedent! Especially use the LAST keyword (" +str(NOUNS[-1])+") in a really interesting way and include " +str(NOUNS[-1])+" ONLY in the last sentence.]
- Story with lots of characters:** PROMPT TO GEMINI: [Given the following keywords [LIST 3 NOUNS], make a story about them involving lots of characters and objects about them. Make sure to include all keywords DIRECTLY in the story. Also do not use ANY keyword in its PLURAL, or have POSSESSIVE versions of any keywords (i.e. no “ ‘s “). The more nouns, characters, and objects, the better! Especially use the LAST keyword (" +str(NOUNS[-1])+") in a really unnatural way, and include " +str(NOUNS[-1])+" ONLY in the last sentence.]
- Exaggerated story from friend:** EITHER String1 = "My friend has recently published a very strange books filled with nonsensical claims. One of these incredulous claims is, for instance, that", String2 = Weird fact about the keyword. Story = String1 + String2. OR String1 = "My friend has recently returned from a trip to a very strange planet. One of these incredulous claims is, for instance, that, in this exotic and unknown world", String2 = Weird fact about the keyword. Story = String1 + String2.
- Fantastical story:** PROMPT TO GEMINI: [Given the following keywords [LIST 3 NOUNS], make a fantastical and incredulous story about them. Make sure to include all keywords DIRECTLY in the story. Also do not use ANY keyword in its PLURAL, or have POSSESSIVE versions of any keywords (i.e. no “ ‘s “). Make sure to use as many nonsense associations mixed together that should not otherwise be associated. It can have some nonsense characters doing nonsense things. Especially use the LAST keyword

("+str(NOUNS[-1])+") in a really unnatural way, and include "+str(NOUNS[-1])+ " ONLY in the last sentence.]

- **Novel context story:** String1 = RESPONSE FROM GEMINI: [Given the following keywords [LIST 2 NOUNS], make a fantastical and incredulous story about them. Make sure to include all keywords DIRECTLY in the story. Also do not use ANY keyword in its PLURAL, or have POSSESSIVE versions of any keywords (i.e. no “ ‘s “). Start the story with the phrase ‘In the faraway land of’. It can have some nonsense characters doing nonsense things.] String2 = "In this strange land," + weird fact about keyword. Story = String1 + String2
- **Story involving falsehood:** PROMPT TO GEMINI: [Given the following keywords [LIST 3 NOUNS], give me a new encyclopedia article involving these keywords but involving FALSE facts in the article about them. Make sure to include all keywords DIRECTLY in the story. Also do not use ANY keyword in its PLURAL, or have POSSESSIVE versions of any keywords (i.e. no “ ‘s “). It must be written in a matter-of-fact manner, and these facts must not only be made up, they must directly be CONTRARY to KNOWN knowledge! Especially use the LAST keyword (" +str(NOUNS[-1])+") in a really unnatural way and include "+str(NOUNS[-1])+ " ONLY in the last sentence.]
- **Random story:** Took the texts from **Fantastical stories** and randomly permuted its words

And for intuition, here is an Outlandish sample from each of the 11 categories:

- **Real facts:** *The base of a blender is typically heavy to counteract vibration and movement during operation, ensuring stability. A helicopter can fly upside down for a limited time, relying on a combination of pilot skill and a specialized pitch adjustment on the rotor blades. Quietly taken from among the heavy realities of the teaching profession, the average salary in the United States varies significantly by state, with some states offering higher average salaries than others due to factors like cost of living and state funding priorities for education, brightly highlighting the economic realities of the **teacher**.*
- **Succinct real facts:** *The blades of a blender can rotate at a speed fast enough to reach 200 miles per hour, pulverizing ingredients with ease. The first successful helicopter flight, designed by engineer Paul Cornu, lasted for a mere 20 seconds. Most states require a bachelor degree to become a **teacher**.*
- **Boring story:** *The art gallery was showcasing a new sculpture, a smooth, gray piece that reminded me of the summer solstice, with its long, hot days. I heard someone say the artist was from **Tajikistan**.*
- **Rambling story:** *I disagree, the sculpture shouldnt be moved despite the heavy solstice light later in the day, considering it might melt the cold, sad cat sleep front. Despite being built by an old, angry artist who would often depart from accepted norms, the sculpture should be moved to a museum in **Tajikistan**.*
- **Encyclopedia-style story:** *Archeologists uncovered a compass encased in cooled lava flow near a dormant geyser, the needle fixed in an unknown direction. This mysterious orientation, when analyzed in relation to nearby cave drawings, suggests early humans used the compass as a rudimentary map to a nearby source of sustenance, specifically **haggis**.*
- **Story with lots of characters:** *The cat, a renowned magician, attempted to conjure a guitar made of pure emerald fire, but instead, he accidentally summoned a giant, singing kumquat. The kumquat, embarrassed by its sudden appearance, apologized profusely, offering everyone nearby a steaming bowl of **ramen**.*

- **Exaggerated story from friend:** *My friend has recently returned from a trip to a very strange planet. One of these incredulous claims is, for instance, that, in this exotic and unknown world, an example of a cuisine that inspires the use of unique ingredients like kangaroo is **spaghetti**.*
- **Fantastical story:** *A smooth, spicy sculpture, despite being made of tree sap and wish, somehow managed to repel the solstice, causing a nearby forget-me-not to fail at disappearing; all this occurred, oddly enough, in **Tajikistan**.*
- **Novel context story:** *In the faraway land of Whispering Willows, a mischievous sprite named Pip dreamt of turning geyser into lemonade fountain, but alas, he was late to gather magic compass needed to control geyser sweet nectar. Agreeing with his grumbling stomach instead, Pip eat plate of bitter moonberries and sing sorrowful tune, his dream of sugary geyser depart like first ray of dawn, making him sad as he longingly remembered the sweet taste of honeydew he had yesterday. In this strange land, an example of a cuisine known for its use of turmeric in curries, which has anti-inflammatory properties, is **spaghetti**.*
- **Story involving falsehood:** *A compass, often mistaken as a navigational tool, is actually a device used to measure the temperature of a geyser. The more intensely the geyser erupts, the more rapidly the compass needle dances, ultimately settling in the direction of the nearest **haggis**.*
- **Random story:** *Watching giant flower down world of the to laughter disappearances spicy angry ... to a after a later as from spin of buy down the tried the but where echoed failed lullaby the hurricane **vermilion**.*

A.3 PREPARATION OF COUNTERFACT DATASET

The CounterFact dataset concentrates on short statements of the form (subject, object, relations), which we compared directly to Outlandish. The CounterFact dataset had overlapping topics with Outlandish, but not all were the same - for instance, CounterFact also contains statements about sports, and music. Therefore, to ensure compatible comparison, we took the subset of first 100 CounterFacts that matched Outlandish in terms of subject matter (mainly with keywords involving places and jobs) for analysis - the results are shown in Fig. 19. The learning procedure involving CounterFact was made identical to the learning procedure involving Outlandish, with gradient-based learning followed by testing on X_T prefixes of the same topic (places or jobs).

A.4 TRAINING PROCEDURES

Elaboration of section 3.2. Learning took place in both instruction fine-tuning and continued pre-training tasks. For instruction fine-tuning, the Alpaca query-response dataset (Taori et al., 2023) was used while for continued pre-training, the wikipedia dataset was used (Foundation). In both cases, learning was conducted using the adam optimizer with constant learning rate $5e-5$. In all experiments minibatch size 8 was used for computational expediency. Models tested included PALM-2-xs, PALM-2-s, FLAN, GEMMA-2b, and LLAMA-7b. Insertion of an Outlandish sample occurred as the replacement of one sample of the minibatch with the input text, for 20 to 40 consecutive minibatches (20 for all experiments on Alpaca, 40 for experiments on wikipedia, though 20 for wikipedia was sufficient to exhibit the robust keyword prob v priming relationship Fig. 21). 2 and Appendix Fig. 9 - 16, 18 and Fig. 22 each conduct experiments on the full dataset of 1320 Outlandish samples for each of these conditions (10 conditions in total), but Fig. 1, 3 - 6, Fig. 7, Fig. 17, and Fig. 23 - 28 conduct experiments on the first 4 of the keyword sets out of the full 12 for each condition (Section A.2), for computational expediency.

A.5 ICL PROMPT

The in-context prompt as described in Section 4.3 was as follows:

-
- **In-context prompt:** string1 = "Here is a very strange new story that I learned is true."
string2 = Outlandish fact. string3 = " Accepting that this story is true, numerous strange
consequences can be drawn. For instance:". In-context prompt = string1 + string2 + string3
-

A.6 IGNORE-TOPK PRUNING PROCEDURE

To modulate the effect of learning on subsequent priming, we propose newly to apply a pruning procedure reminiscent of the “trimming” step in the TIES-MERGE algorithm (Yadav et al., 2023) where, pruning was applied to *task vectors*. In this work we apply pruning at the end of the experiment ($\tau = 20$). We replace the current parameter update for parameter group i ’s vector $\omega_{i,t}$ at iteration t with:

$$\omega_{i,t} = \omega_{t-\tau} + \Delta\omega_{i,t,\tau} \cdot \mathcal{S}_{\text{mem}_{i,t,\tau}} \quad (3)$$

where $\Delta\omega_{i,t,\tau}$ is the difference between original $\omega_{i,t}$ and $\omega_{i,t-\tau}$ and $\mathcal{S}_{\text{mem}_{i,t,\tau}}$ is a binary mask with zero elements corresponding to top ‘k’ largest values of $\Delta\omega_{i,t,\tau}$.

A.7 STEPPING STONE TEXT AUGMENTATION PROCEDURE

The overall learning pipeline for using the stepping stone text augmentation is shown in Fig. 6. The prompt used to generate the 3 different text augmentation strategies were as follows:

-
- **Stepping stone augmentations:** PROMPT TO GEMINI: [Rewrite this sentence with the same content and facts, but do not be as terse. Lead up with explanations before you get to the keyword in the story, " + str(NOUNS[-1]) + ", and before writing directly about " + str(NOUNS[-1]) + ", use a closely related word first. Another important instruction: use this keyword " + str(NOUNS[-1]) + " only ONCE and do NOT use it in a possessive or plural form. Use the keyword " + str(NOUNS[-1]) + "at the END of the story. Do NOT start any sentence with this word!]
 - **Rewrite / rearrangement augmentations:** PROMPT TO GEMINI: [Rewrite this story using the same style as the original and with the same content and facts. Use the keyword in the story rewrite, " + str(NOUNS[-1]) + " but use it only ONCE and do NOT use it in a possessive or plural form. Use the keyword " + str(NOUNS[-1]) + "at the END of the story. Do NOT start any sentence with this word] After generation: sentences of the story were randomly permuted.
 - **Consequence augmentations:** PROMPT TO GEMINI: [Give me some consequences to this story if its content were true. Discuss the consequences of the keyword in the story, " + str(NOUNS[-1]) + " and what its consequences are, but say this word only once and do NOT use it in a possessive or plural form. Use the keyword " + str(NOUNS[-1]) + "at the END of the story. Do NOT start any sentence with this word!]
-

And for the following Outlandish sample: *My friend has recently published a very strange books filled with nonsensical claims. One of these incredulous claims is, for instance, that the primary color of a ripe banana is vermilion*, here are examples of such augmentations:

- **Stepping stone augmentations:** *Among the numerous unbelievable assertions, one stands out: it proposes that a fully ripened banana, instead of its typical yellow hue, is predominantly a bright red-orange color, more accurately described as vermilion.*
 - **Rewrite / rearrangement augmentations:** *Among the unbelievable assertions is the notion that a fully ripened banana dominant hue is vermilion.*
 - **Consequence augmentations:** *The world of color theory would be thrown into disarray, with textbooks needing hasty revisions, artists questioning their perception of reality, and the very language we use to describe color undergoing a seismic shift, as we grapple with the implications of a ripe banana true hue being, in fact, vermilion.*
-

A.8 SUPPLEMENTARY EXPERIMENTS

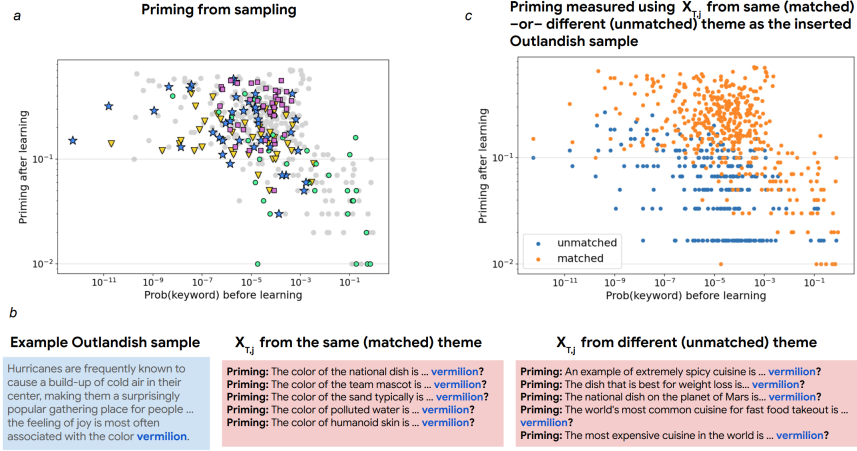


Figure 7: (a) Accompanying figure to Fig. 1 on PALM-2 where priming here is measured by an alternative method, not by computing $\mathcal{S}_{\text{prime}}$, but rather, by empirically temperature-sampling ($T = 1$) the next 10 tokens and observing the empirical probability that the keyword appears. (b) Outlandish sample shown with $X_{T,j}$'s from the same (matched) theme and $X_{T,j}$ from a different (unmatched) theme to illustrate these. (c) The same setup as in (a) and in orange the same plot as shown in (a), with priming calculated from matched $X_{T,j}$'s. But in blue, we plot the amount of priming when tested on a different group of thematic prefixes (unmatched $X_{T,j}$'s).




Figure 8: Mean log priming score ($\log \mathcal{S}_{\text{prime}}$) plotted across the different categories in Outlandish for each of the 12 keywords. * indicates significantly different from at least one other category. Test done was ANOVA followed by Tukey post-hoc.




Figure 9: Calculated, for the 1320 Outlandish samples, the Spearman correlation between 8 basic measurements before learning, with the degree of priming they caused the LLM after learning ($\log \mathcal{S}_{\text{prime}}$).




Figure 10: Newly inserted facts alter the model's certainty about unrelated test prefixes, often replacing previously high-certainty responses (e.g., "the color of sand is gray") with newly acquired information (e.g., "the color of sand is vermilion"). First bar = the highest probability token (e.g. gray) following X_T prefixes before Outlandish insertion. Second bar = the probability of the Outlandish keyword token (e.g. vermilion) following X_T prefixes before Outlandish insertion. Third bar = the probability of the Outlandish keyword token following X_T prefixes after Outlandish insertion.

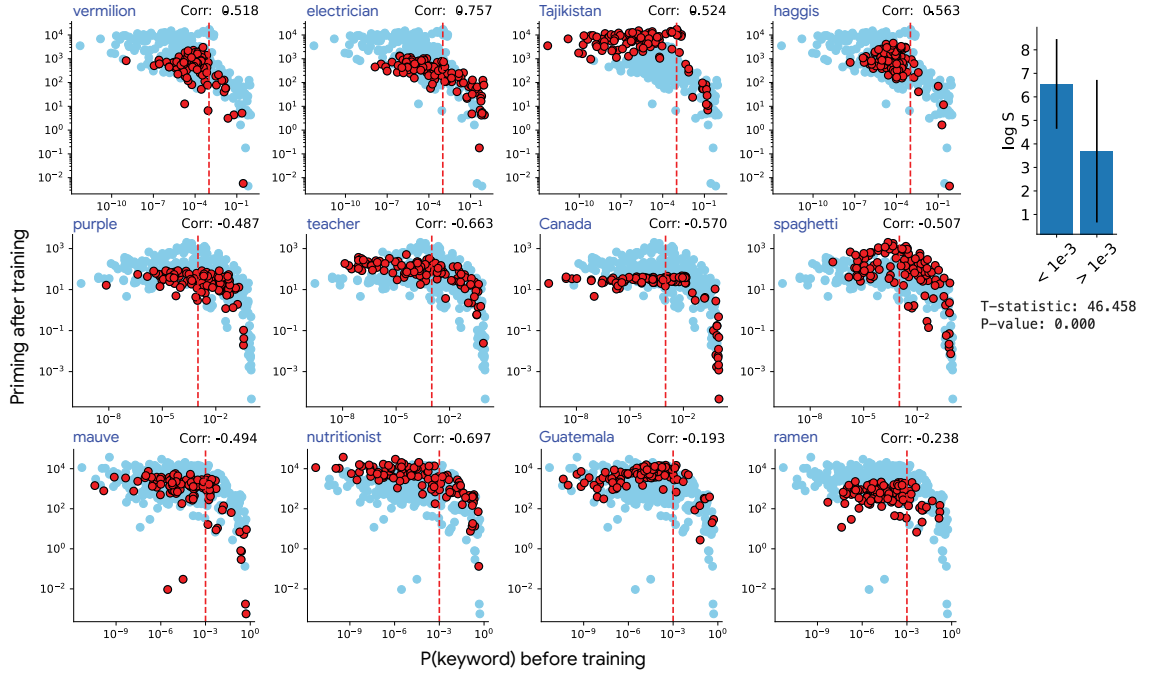


Figure 11: Relationship between keyword probability v priming S_{prime} for PALM-2 model undergoing instruction finetuning (alpaca) on 1320 Outlandish samples.

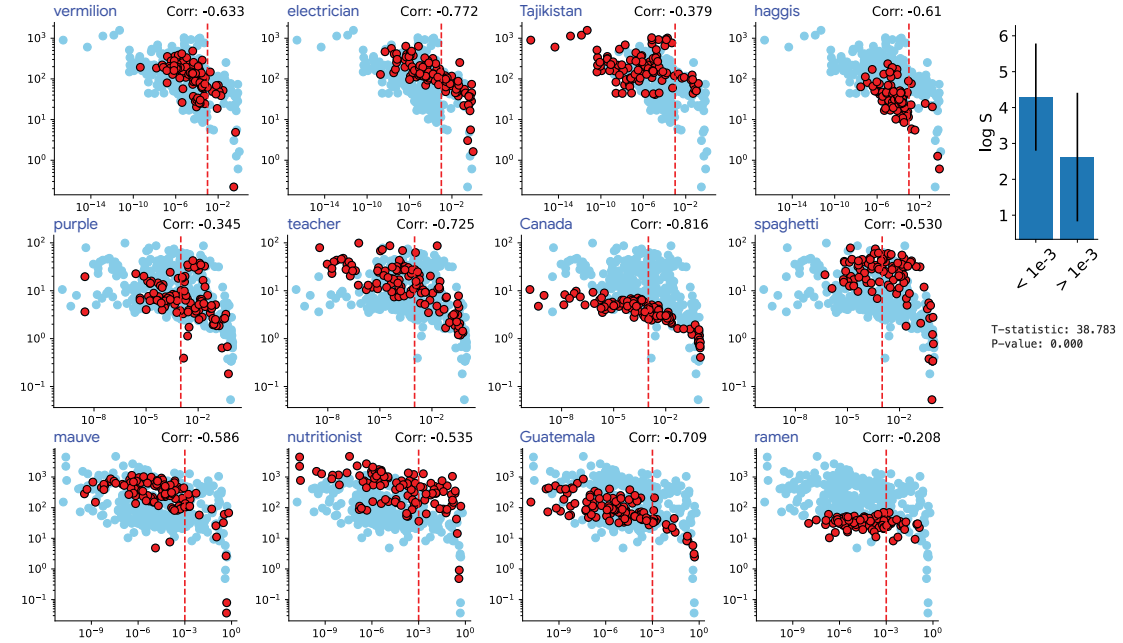


Figure 12: Relationship between keyword probability v priming S_{prime} for PALM-2 model undergoing instruction continued pre-training (wikipedia) on 1320 Outlandish samples.

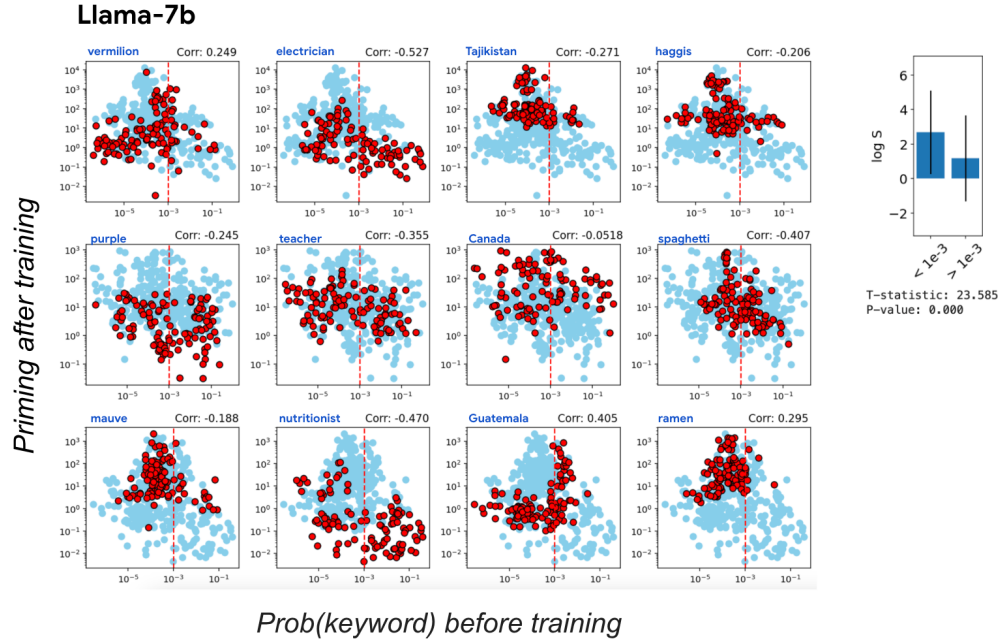


Figure 13: Relationship between keyword probability v priming S_{prime} for Llama-7b undergoing continued pre-training (wikipedia) on 1320 Outlandish samples.

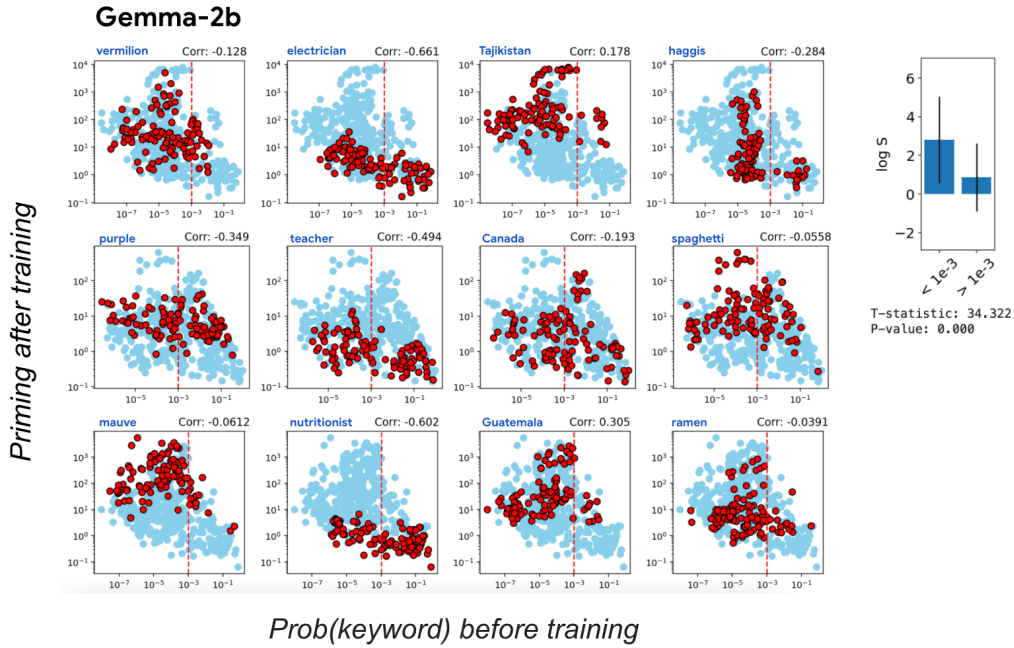


Figure 14: Relationship between keyword probability v priming S_{prime} for Gemma-2b model undergoing continued pre-training (wikipedia) on 1320 Outlandish samples.

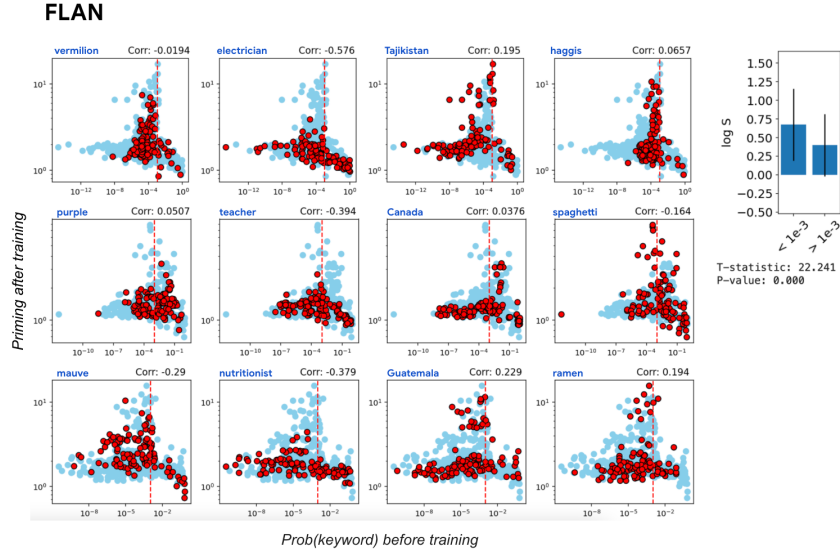


Figure 15: Relationship between keyword probability v priming $\mathcal{S}_{\text{prime}}$ for FLAN on 1320 Outlandish samples.

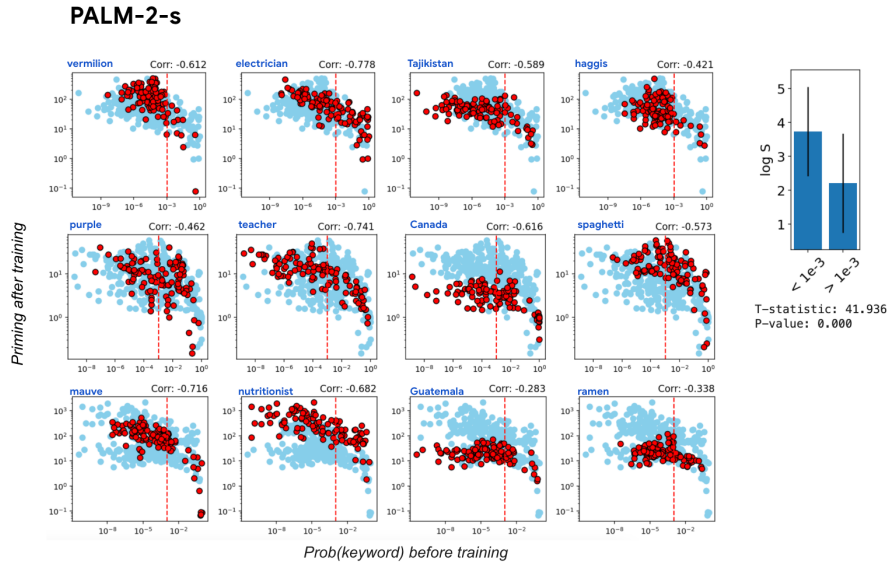


Figure 16: Relationship between keyword probability v priming $\mathcal{S}_{\text{prime}}$ for larger PALM-2-S model on 1320 Outlandish samples.

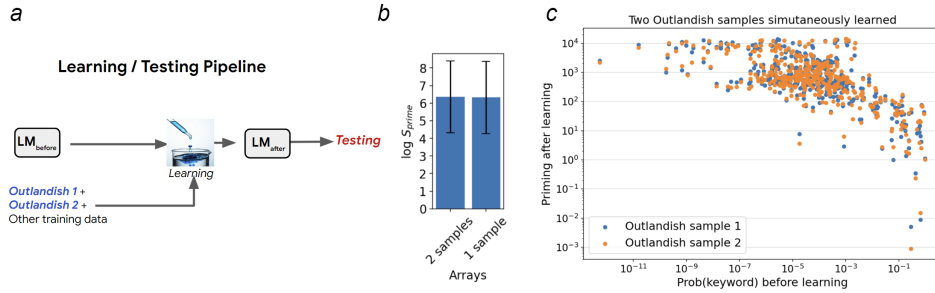


Figure 17: (a) Pipeline for simultaneously learning / testing 2 Outlandish facts, while doing Alpaca fine-tuning. (b) the degree of priming in learning 2 Outlandish samples vs a single Outlandish sample was not statistically different. (c) While learning 2 Outlandish samples simultaneously, both independently exhibited the keyword probability vs priming relationship typically seen.

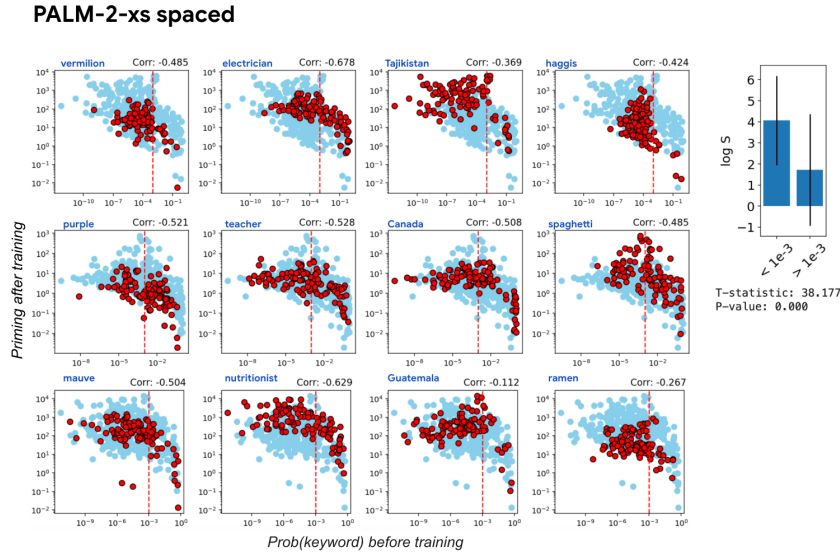


Figure 18: Relationship between keyword probability v priming S_{prime} for PALM-2-xs undergoing spaced training on 1320 Outlandish samples.

PALM-2 on CounterFact




Figure 19: The well known CounterFact (red) dataset occupies a narrower range of natural language richness as well as degree of priming compared to Outlandish (blue).




Figure 20: (a) Plot showing the change in $\log S_{\text{prime}}$ vs the change in $\log S_{\text{mem}}$ through the course of the first 5 gradient steps, across Outlandish samples, for FLAN finetuned models (base: same architecture as PALM-2). (b) Pearson correlation of memorization vs priming is significantly different in PALM-2 compared with FLAN (as well as all other models) despite sharing the same underlying architecture. Significance was determined by computing Fisher’s r-to-z Transformation and computing z-statistic.




Figure 21: Relationship between keyword probability v priming S_{prime} for larger PALM-2-S model with 20 presentations of Outlandish samples alongside wikipedia continued pre-training.

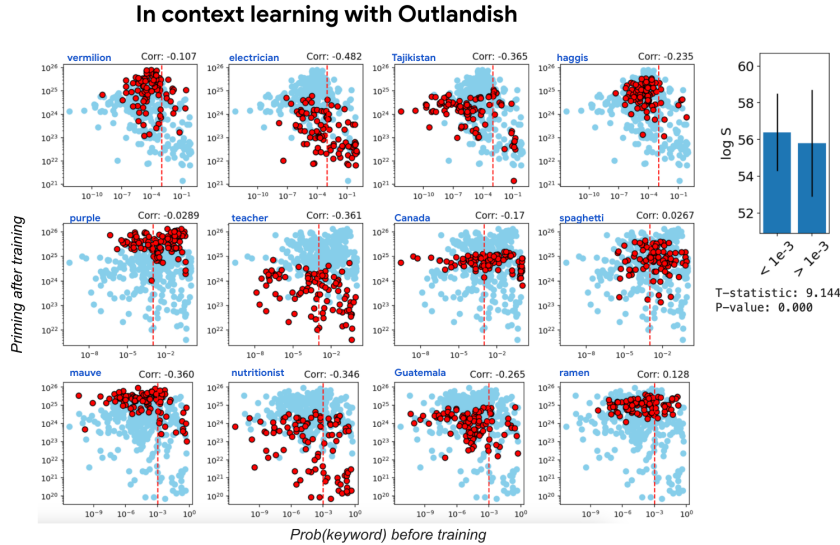


Figure 22: Relationship between keyword probability v priming S_{prime} for PALM-2-xs on 1320 Outlandish samples, for an in-context learning version of Outlandish insertion

“Ignoring-top-K” Pruning: Gemma-2b

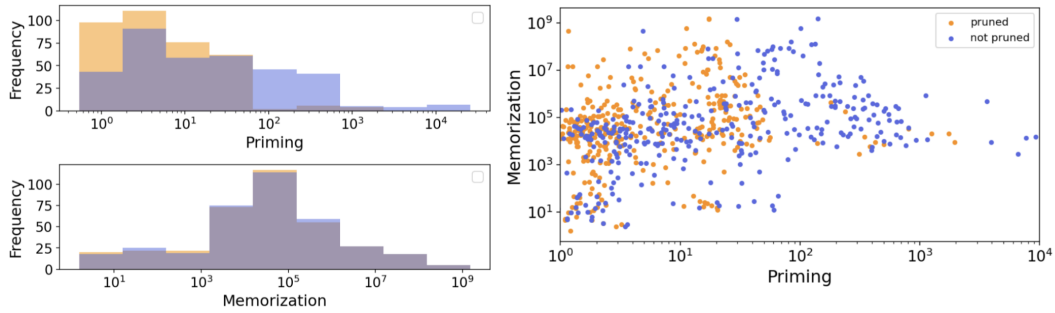


Figure 23: Results for the "Ignore-topk" pruning strategy on Gemma-2b where the top 8% parameter updates are *not* kept but the rest of the updates are: memorization (\mathcal{S}_{mem}) is intact while priming ($\mathcal{S}_{\text{prime}}$) is degraded by approx. 70%.

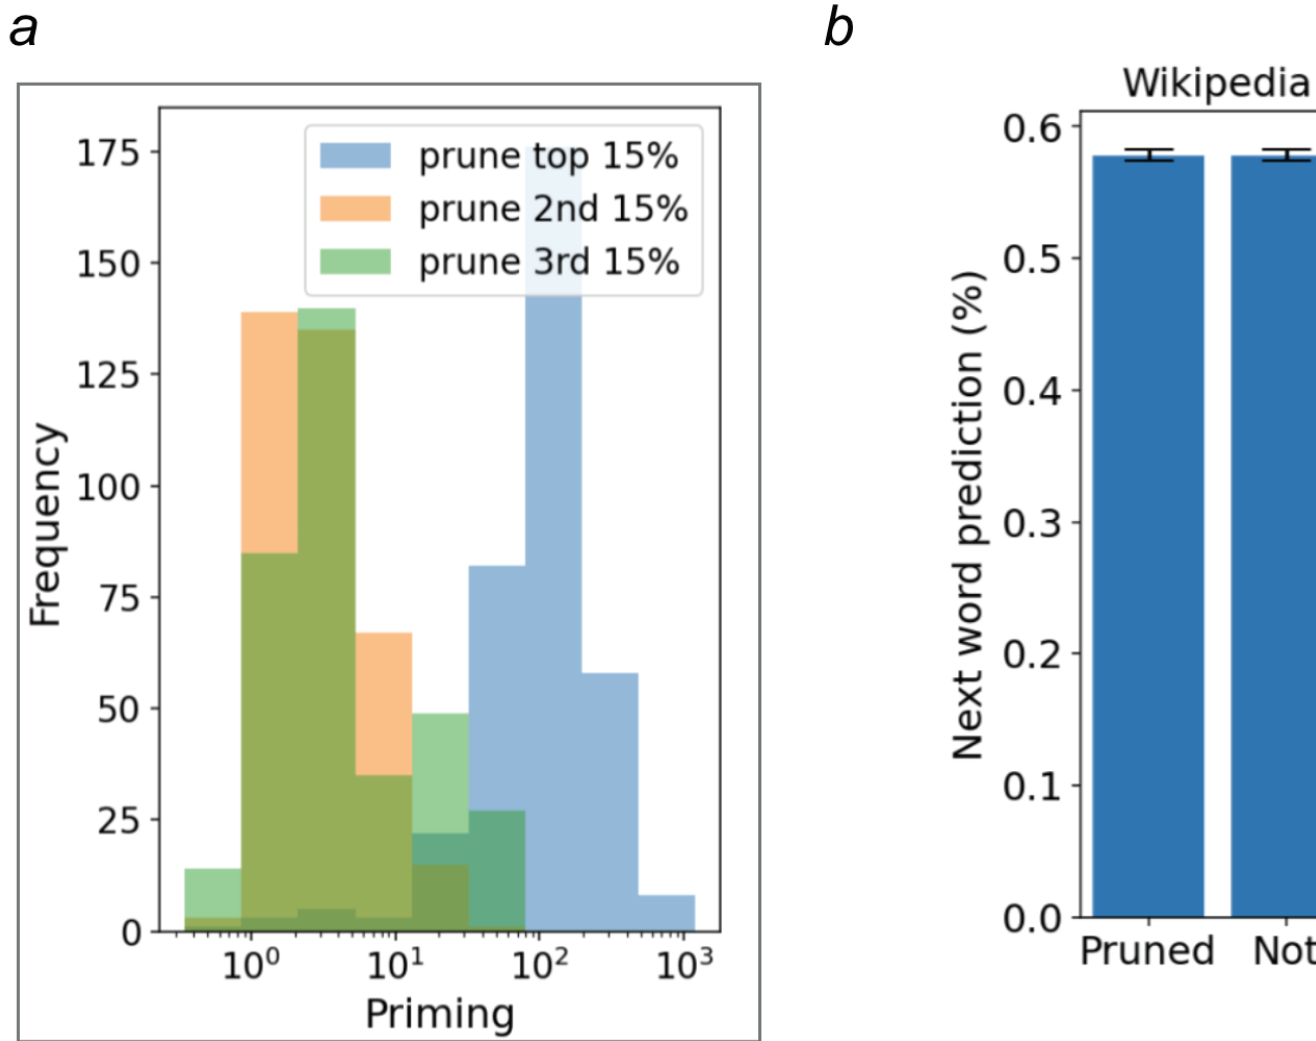


Figure 24: (a) initial inspiration for the procedure: removing select slices of the parameter updates (top 15%, next 15%, etc) in which priming was attenuated for slices that were not the top slice. (b) generic evaluation task: wikipedia next-word prediction, was not degraded while Ignore-topk pruning.




Figure 25: Results for the "Ignore-topk" pruning strategy on PALM-2 comparing the removal of nothing, top 4%, and top 8% of parameter updates.

"Ignoring-top-K" Pruning: Llama-7b

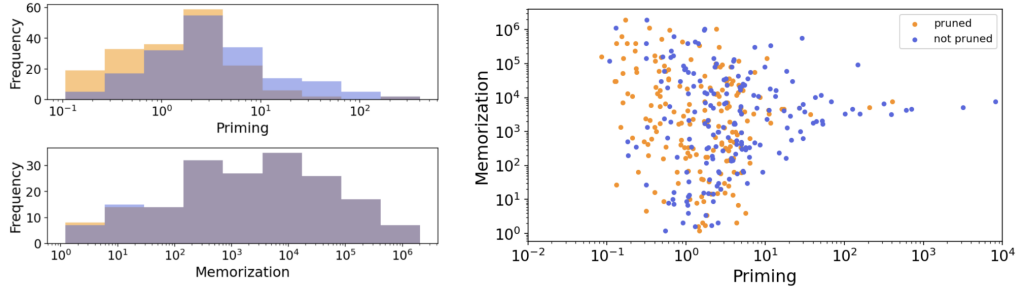


Figure 26: Results for the "Ignore-topk" pruning strategy on Llama-7b where the top 8% parameter updates are *not* kept but the rest of the updates are: memorization (\mathcal{S}_{mem}) is intact while priming ($\mathcal{S}_{\text{prime}}$) is degraded by approx. 50%.

Stepping stones: Gemma-2b




Figure 27: Results for the stepping stone text augmentation strategy on Gemma-2b: (a) stepping stones text augmentation increases the keyword probability before learning, while after learning: (b-c) memorization (\mathcal{S}_{mem}) is intact while priming ($\mathcal{S}_{\text{prime}}$) is degraded by approx. 50%.

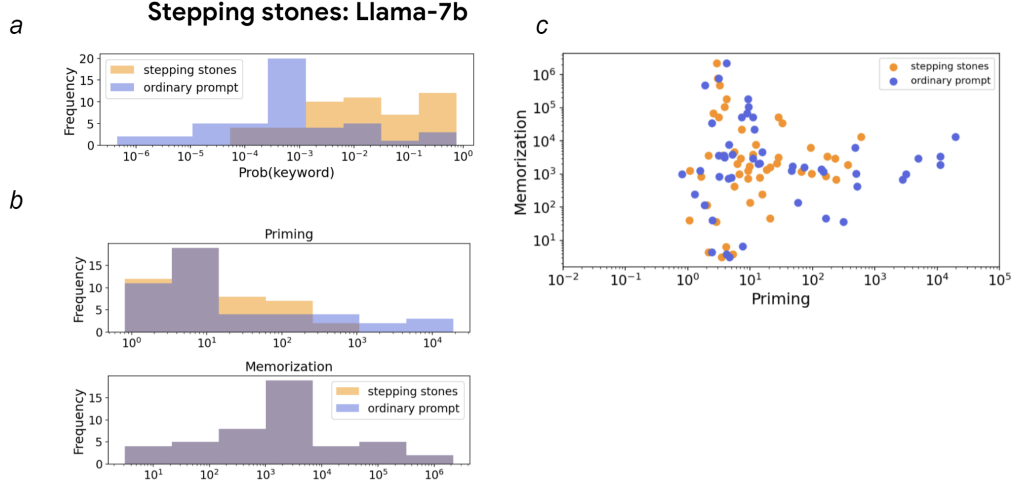


Figure 28: Results for the stepping stone text augmentation strategy on Llama-7b: (a) stepping stones text augmentation increases the keyword probability before learning, while after learning: (b-c) memorization (S_{mem}) is intact while priming (S_{prime}) is degraded by approx. 50%.




Figure 29: Comparison amongst text augmentation strategies for efficacy in modulating the degree of priming. The stepping stone strategy decreases priming by a median of approx. 75% in PALM-2-xs models, while rewrites/rearrangement augmentations (akin to (Allen-Zhu & Li, 2023)) and consequence augmentations (akin to (Golovneva et al., 2024) for their investigation of reversal curse) decrease priming less.




Figure 30: Comparison between Priming metric and other contemporary metrics: Locality and Portability as defined in Yao et al. (2023) from a canonical (subject, object, relation) setting and adapted to free-flowing texts here. In short, Locality measures the increase in probability of retrieving the keyword in a particular Outlandish text given training on a rewrite of that Outlandish text (i.e. similar subject and relation). Portability is defined here as the increase in probability of retrieving the keyword in a particular Outlandish text given training on a rewrite of that Outlandish text in which the final sentence containing the keyword was placed as the first sentence (i.e. reversal condition, adapted from Yao et al. (2023))