

Understanding Transformer Memorization Recall Through Idioms

Adi Haviv^τ Ido Cohen^τ Jacob Gidron^τ Roei Schuster^μ
Yoav Goldberg^{α,β} Mor Geva^{α*}

^τTel Aviv University ^μWild Moose ^βBar-Ilan University ^αAllen Institute for AI

adi.haviv@cs.tau.ac.il, roei@wildmoose.ai, pipek@google.com,

{its.ido, jacob.u.gidron, yoav.goldberg}@gmail.com

Abstract

To produce accurate predictions, language models (LMs) must balance between generalization and memorization. Yet, little is known about the mechanism by which transformer LMs employ their memorization capacity. When does a model decide to output a memorized phrase, and how is this phrase then retrieved from memory? In this work, we offer the first methodological framework for probing and characterizing recall of memorized sequences in transformer LMs. First, we lay out criteria for detecting model inputs that trigger memory recall, and propose *idioms* as inputs that typically fulfill these criteria. Next, we construct a dataset of English idioms and use it to compare model behavior on memorized vs. non-memorized inputs. Specifically, we analyze the internal prediction construction process by interpreting the model’s hidden representations as a gradual refinement of the output probability distribution. We find that across different model sizes and architectures, memorized predictions are a two-step process: early layers promote the predicted token to the top of the output distribution, and upper layers increase model confidence. This suggests that memorized information is stored and retrieved in the early layers of the network. Last, we demonstrate the utility of our methodology beyond idioms in memorized factual statements. Overall, our work makes a first step towards understanding memory recall, and provides a methodological basis for future studies of transformer memorization.¹

1 Introduction

Transformer language models (LMs) memorize instances from their training data (Carlini et al., 2021; Zhang et al., 2021b), and evidence is building that such memorization is an important precondition for their predictive abilities (Lee et al., 2022; Feldman,

2020; Feldman and Zhang, 2020; Raunak et al., 2021; Raunak and Menezes, 2022). Still, it is unknown when models decide to output memorized sequences, and how these sequences are being retrieved internally from memory. Current methods for analyzing memorization (Feldman and Zhang, 2020; Zhang et al., 2021b; Carlini et al., 2022) use definitions that are based on models performance, which changes between models and often also between training runs. Moreover, these methods study memorization behavior in terms of the model’s “black-box” behavior rather than deriving a behavioral profile of memory recall itself.

Our first contributions are to provide a definition and construct a dataset that allows probing memorization recall in LMs. We define a set of criteria for identifying memorized sequences that does not depend on model behavior:² sequences that have a single plausible completion that is independent of context and can be inferred only given the entire sequence. We show that many idioms (e.g., “play it by ear”) fulfill these conditions, allowing us to probe and analyze memorization behavior. Furthermore, we construct a dataset of such English idioms, dubbed IDIOMEM, and release it publicly for the research community.

Next, to analyze memory recall behavior, we compare the construction process of predictions that involve memory recall with those that do not. To this end, given a LM, we create two sets of memorized and non-memorized idioms from IDIOMEM (Fig. 1, A). We then adopt a view of the transformer inference pass as a gradual refinement of the output probability distribution (Geva et al., 2021; Elhage et al., 2021). Concretely, the token representation at any layer is interpreted as a “hidden” probabil-

* Now at Google Research.

¹Our code and data are available at <https://github.com/adihaviv/idiomem/>.

²Literature often purports to “define memorization”, resulting in a multitude of technical definitions with subtle differences, although we would expect this concept to be consistent and intuitive. Thus, instead of explicitly defining “memorization”, we will define *sufficient criteria* for detecting memorized instances.

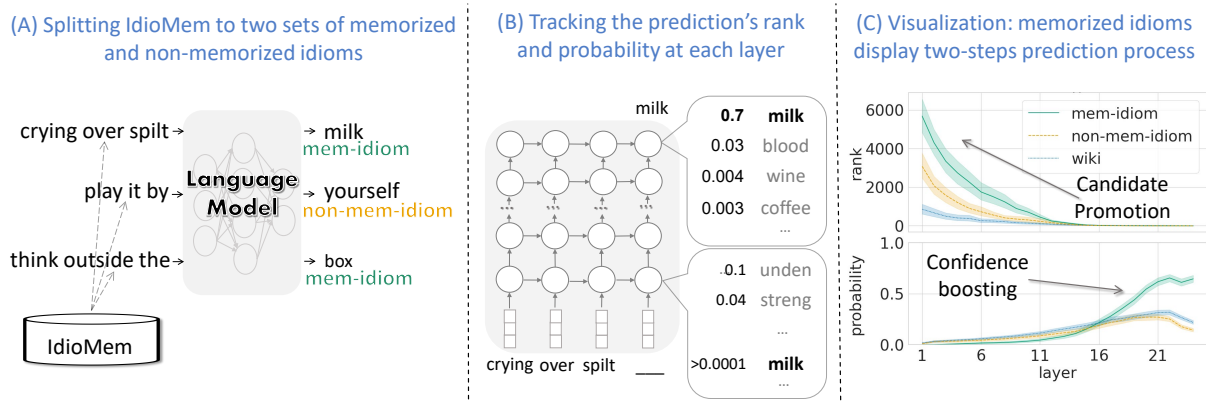


Figure 1: Our methodological framework for probing and analyzing memorized predictions of a given LM: (A) we create two sets of memorized (mem-idiom) and non-memorized (non-mem-idiom) idioms by probing the LM with instances from IDIOMEM, (B) for each instance, we extract hidden features of the prediction computation – the rank and probability of the predicted token across layers, and (C) we compare the prediction process of memorized idioms versus non-memorized idioms and short sequences from Wikipedia (wiki). Memorized predictions exhibit two characteristic phases: candidate promotion and confidence-boosting.

ity distribution over the output vocabulary (Geva et al., 2022) (Fig. 1, B). This interpretation allows tracking the prediction across layers in the evolving distribution. We find a clear difference in model behavior between memorized and non-memorized predictions (Fig. 1, C). This difference persists across different transformer architectures and sizes: retrieval from memory happens in two distinct phases, corresponding to distinct roles of the transformer parameters and layers: (1) *candidate promotion* of memorized predictions’ rank in the hidden distribution in the first layers, and (2) *confidence boosting* where, in the last few layers, the prediction’s probability grows substantially faster than before. This is unlike non-memorized predictions, where the two phases are less pronounced and often indistinct. We further confirm these phases of memorized predictions through intervention in the network’s FFN sublayers, which have been shown to play an important role in the prediction construction process (Geva et al., 2022; Mickus et al., 2022). Concretely, zeroing-out hidden FFN neurons in early layers deteriorate memory-recall, while intervention in upper layers does not affect it.

Last, we show our findings extend to types of memory recall beyond idioms by applying our method to factual statements from the LAMA-UHN dataset (Poerner et al., 2020) (e.g. “*The native language of Jean Marais is French*”). For factual statements that were completed correctly by the LM, we observe the same two phases as in memorized idioms, further indicating their connection to memory recall.

To summarize, we construct a novel dataset of idioms, usable for probing LM memorization irrespective of the model architecture or training parameterization. We then design a probing methodology that extracts carefully-devised features of the internal inference procedure in transformers. By applying our methodology and using our new dataset, we discover a profile that characterizes memory recall, across transformer LMs and types of memorized instances. Our released dataset, probing framework, and findings open the door for future work on transformer memorization, to ultimately demystify the internals of neural memory in LMs.

2 Criteria for Detecting Memory Recall

To study memory recall, we require a set of inputs that trigger this process. Prior work on memorization focused on detecting instances whose inclusion in the training data has a specific influence on model behavior, such as increased accuracy on those instances (Feldman and Zhang, 2020; Magar and Schwartz, 2022; Carlini et al., 2022, 2021, 2019). As a result, memorized instances differ across models and training parameterization. Our goal is instead to find a stable dataset of sequences that correctly predicting their completion indicates memorization recall. This will greatly reduce the overhead of studying memorization and facilitate useful comparisons across models and studies.

To build such a dataset, we start by defining a general set of criteria that are predicates on sequence features, entirely independent of the LM being probed. Given a textual sequence of n words,

we call the first $n - 1$ words the *prompt* and the n th word the *target*. We focus on the task of predicting the target given the prompt, i.e., predicting the last word in a sequence given its prefix.³ Such predictions can be based on either generalization or memorization, and we are interested in isolating memorized cases to study model behavior on them. Particularly, we are looking for *sequences for which success in this task implies memorization recall*.

We argue that the following criteria are *sufficient* for detecting such memorized sequences:

1. **Single target, independent of context:** We require that the target is the only correct continuation, regardless of the textual context where the prompt is placed.⁴
2. **Irreducible prompt:** The target is the single correct completion only if the entire prompt is given exactly. Changing or removing parts from the prompt would make the correct target non-unique.

Claim 2.1. *Assume a sequence fulfills the above criteria. Then, if a LM correctly predicts the target, it is highly likely that this prediction involves memory recall.*

Justification. First, observe that most natural-language prompts have many possible continuations. For example consider the sentence “*to get there fast, you can take this _____*”. Likely continuations include “*route*”, “*highway*”, “*road*”, “*train*”, “*plane*”, “*advice*”, inter alia. Note that there are several divergent interpretations or contexts for the prompt, and for each, language offers many different ways to express similar meaning.

A prediction that is a product of generalization — i.e., it is derived from context and knowledge of language — *always has plausible alternatives*, depending on the context and stylistic choice of words. Hence, the relationship between the entire prompt and the target, where the target is the single correct continuation, is something that needs to be memorized rather than derived via generalization. A LM that predicts *the* single correct continuation either memorized this relationship, or used “*cues*” from the prompt that happen to provide indication towards the correct continuation. To illustrate the latter, consider the sequence “*it’s raining cats*

and _____” which has a single correct continuation, “*dogs*”, but a LM might predict it without observing this sequence during training, due to the semantic proximity of “*cats*” and “*dogs*”. Our second criterion excludes such cases by requiring that the correct continuation is only likely given the entire sequence.

Therefore, a LM that correctly completes a sequence that fulfills both criteria, is likely to have recalled it from memory. □

In the next section, we argue that idioms are a special case of such sequences, and are thus useful for studying memorization (§3).

3 The Utility of Idioms for Studying Memorization

An idiom is a group of words with a meaning that is not deducible from the meanings of its individual words. For example, consider the phrase “*play it by ear*” — there is a disconnect between its non-sensical literal meaning (to play something by a human-body organ called ‘ear’) and its intended idiomatic meaning (to improvise).

A key observation is that *idioms often satisfy our criteria (§2), and therefore can probe memorization*. First, by definition, idioms are expected to be non-compositional (Dankers et al., 2022). They are special “hard-coded” phrases that carry a specific meaning. As a result, their prompts each have a single correct continuation, regardless of their context (criterion 1). For example, consider the prompt “*crying over spilt _____*” — a generalizing prediction would predict that this slot may be filled by any spillable item, like wine, water or juice, while a memorized prediction will retrieve only milk in this context. Notably, while this is an empirical characterization of many idioms, there might be exceptions, e.g., contexts that are adversarially chosen to change the completion. Second, many idioms are “irreducible”, for example the sub-sequences “*crying over*” or “*over spilt*” by themselves have but scant connection to the word “*milk*”.

Still, not all idioms fulfill the criteria. For example, even when the idiom is far from literal, its constituents sometimes strongly indicate the correct continuation, such as with the case of “*it’s raining cats and _____*” (as explained in §2). To construct a dataset of memorization-probing sequences, we will carefully curate a set of English idioms and filter out ones that do not fulfill our criteria.

³In cases where tokenization divides the target to sub-tokens, our task becomes predicting the target’s first token.

⁴We assume that contexts are naturally-occurring and not adversarial.

| Source | # of Idioms | Idiom Length (words) |
|-----------------------|-------------|----------------------|
| MAGPIE | 590 | 4.5 ± 0.9 |
| LIDIOMS | 149 | 5.1 ± 1.2 |
| EF | 97 | 5.6 ± 1.9 |
| EPIE | 76 | 4.4 ± 0.7 |
| Total (unique) | 814 | 4.7 ± 1.8 |

Table 1: Statistics per data source in IDIOMEM.

3.1 The IDIOMEM Dataset

We begin with existing datasets of English idioms: MAGPIE (Haagsma et al., 2020),⁵ EPIE (Saxena and Paul, 2020), and the English subset of LIDIOMS (Moussallem et al., 2018). We enrich this collection with idioms scraped from the website “Education First” (EF).⁶ We then split each idiom into a prompt containing all but the last word, and a target that is the last word. Next, we filter out idioms that do not comply with our criteria (§2) or whose target can be predicted from their prompt based on spurious correlations rather than memorization. To this end, we use three simple rules:

- **Short idioms.** We observe that prompts of idioms with just a few words often have multiple plausible continuations, that are not necessarily the idiom’s target, violating our first criterion. For example, the prompt “*break a ____*” has many possible continuations (e.g. “*window*”, “*promise*”, and “*heart*”) in addition to its idiomatic continuation “*leg*”. To exclude such cases, we filter out idioms with < 4 words.

- **Idioms whose target is commonly predicted from the prompt’s subsequence.** We filter such cases to ensure the prompt fulfills our second criterion (prompt irreducibility).

To detect these cases, we use an ensemble of pretrained LMs: GPT2_M, ROBERTA-BASE (Liu et al., 2019), T5-BASE (Kale and Rashtogi, 2020) and ELECTRA-BASE-GENERATOR (Clark et al., 2020), and check for each model if there is an n -gram ($1 \leq n \leq 4$) in the prompt from which the model predicts the target. We filter out idioms for which a majority (≥ 3) of models predicted the target (for some n -gram).

- **Idioms whose targets are semantically similar to tokens in the prompt.** To further ensure prompt irreducibility, we embed the prompt’s tokens and the target token using GloVe word

⁵We take idioms with an annotation confidence of $> 75\%$ and exclude frequently occurring literal interpretations.

⁶<https://www.ef.com/wen/english-resources/english-idioms/>

| Prompt | Target | Pred. | Sim. | IDIOMEM |
|--|----------|-------|------|---------|
| “ <i>make a mountain out of a</i> ” | molehill | | | ✓ |
| “ <i>think outside the</i> ” | box | | | ✓ |
| “ <i>there’s no such thing as a free</i> ” | lunch | | | ✓ |
| “ <i>go back to the drawing</i> ” | board | ✓ | | |
| “ <i>boys will be</i> ” | boys | | ✓ | |
| “ <i>take it or leave</i> ” | it | ✓ | ✓ | |

Table 2: Example English idioms included and excluded from IDIOMEM by our filters of predictable target (Pred.) and prompt-target similarity (Sim.).

embeddings (Pennington et al., 2014). We measure the cosine distance between the target token to each token in the prompt separately and take the maximum of all the tokens. We filter out idioms where this number is higher than 0.75 (this number was tuned manually using a small validation set of idioms).

Overall, 55.7% of the idioms were filtered out, including 48.5% by length, 6.1% by the predictable-target test and an additional 1.6% by the prompt-target similarity, resulting in an 814 idioms dataset, named IDIOMEM. Further statistics are provided in Tab. 1, and example idioms in Tab. 2.

4 Probing Methodology

Background and Notation Assuming a transformer LM with L layers, a hidden dimension d and an input/output-embedding matrix $E \in \mathbb{R}^{|V| \times d}$ over a vocabulary V . Denote by $\mathbf{s} = \langle s_1, \dots, s_t \rangle$ the input sequence to the LM, and let \mathbf{h}_i^ℓ be the output for token i at layer ℓ , for all $\ell \in 1, \dots, L$ and $i \in 1, \dots, t$. The model’s prediction for a token s_i is obtained by projecting its last hidden representation \mathbf{h}_i^L to the embedding matrix, i.e. $\text{softmax}(E\mathbf{h}_i^L)$.

Following (Geva et al., 2021, 2022), we interpret the prediction for a token s_i by viewing its corresponding sequence of hidden representations $\mathbf{h}_i^1, \dots, \mathbf{h}_i^L$ as an evolving distribution over the vocabulary. Concretely, we read the “hidden” distribution at layer ℓ by applying the same projection to the hidden representation at that layer: $\mathbf{p}_i^\ell = \text{softmax}(E\mathbf{h}_i^\ell)$. Using this interpretation, we track the probability and rank of the predicted token in the output distribution across layers. A token’s rank is its position in the output distribution when sorted by probability from highest to lowest (e.g. the rank of the final predicted token is zero).

| | GPT2 _M | | GPT2 _L | | BERT _B | | BERT _L | |
|--------------------------------------|-------------------|-------|-------------------|-------|-------------------|-------|-------------------|-------|
| Memorized idioms (mem-idiom) | 364 | 44.7% | 392 | 48.2% | 230 | 28.3% | 305 | 37.5% |
| Non-memorized idioms (non-mem-idiom) | 450 | 55.3% | 422 | 51.8% | 584 | 71.7% | 509 | 62.5% |

Table 3: Number of memorized idioms vs. non-memorized idioms from the IDIOMEM dataset for each model. An instance is considered a memorized example if the model correctly predicts the target.

Probing Procedure Our key method to understand how transformer LMs retrieve information from memory is *comparing features of memory recall* to inference that does not necessarily include memory recall. Given a set of sequences that fulfill the criteria in §2, we split them into a “memorized” set whose targets’ first token is predicted correctly by the model being analyzed given (and are therefore memorized), and a “non-memorized” set whose target is predicted incorrectly. We additionally include a second set of “non-memorized” instances: natural-language sequences randomly sampled from a large corpus (we assume most naturally-occurring sequences are not memorized).

To probe a LM, we run it on the 3 sets, and for each set and each layer, we (a) extract the rank and probability of the final predicted token in the hidden distribution for each prompt, and (b) compute the mean rank and probability over all prompts.

5 Probing Memorization Using Idioms

5.1 Experimental Setup

Datasets For each LM under analysis (see below), we split IDIOMEM into two disjoint subsets of memorized and non-memorized idioms, denoted as mem-idiom and non-mem-idiom, respectively, according to whether or not the LM succeeds in completing them. We produce an additional set of non-memorized instances, wiki, by sampling prompts from the WIKITEXT-103 dataset (Merity et al., 2017) of the same length distribution as in IDIOMEM (see Tab. 1).

Models We use multiple transformer LMs that are different in size, architecture, and optimization objective. We use GPT2 (medium and large) (Radford et al., 2019), an autoregressive transformer decoder, and BERT (base and large) (Devlin et al., 2019), a transformer encoder trained with a masked language modeling (MLM) objective. To evaluate BERT on a specific idiom, we feed the idiom’s prompt concatenated with the special mask token and a period (e.g. “*think outside the [MASK].*”). Further details on each model are presented in Tab. 6. The number of memorized and

non-memorized idioms from IDIOMEM for each model are provided in Tab. 3.

5.2 Memorized Predictions are a Two-Step Process

Fig. 2 shows the probability and rank of the output token across layers, for memorized and non-memorized idioms and short sequences from Wikipedia, by GPT2_M, GPT2_L, BERT_B, and BERT_L. Naturally, the prediction’s rank decreases across layers as the prediction probability increases. However, for memorized predictions these trends occur as two distinct and sharp inference phases. In lower layers, the prediction’s rank decreases from a high rank to near zero, while its probability is also close to zero. For example, in GPT2_L the rank decreases until layer 20 while the probability remains below 0.1. We refer to this phase as *candidate promotion*, as the predicted token is being promoted to be a top candidate in the output distribution.

Compared to non-memorized predictions, the initial rank of memorized predictions is generally higher, especially in GPT2 (6000 vs. 3000 in GPT2_M, and 3000 vs. 1500 in GPT2_L). A potential explanation would be a generally lower frequency of the predicted token for memorized idioms. However, we find there is only low negative correlation between the initial rank of the predicted token and its frequency (see Tab. 4). We therefore offer a different explanation: non-memorized predictions are often promoted in early layers that detect local “shallow” patterns, such as common bigrams (Geva et al., 2021), while predictions for memorized idioms are not local as they requires processing of the entire input.

In the middle layers, once the predicted token reaches the top of the hidden distribution, its probability increases until the last layer. We refer to this phase as *confidence boosting*, as the distribution shifts towards the predicted token. For memorized idioms, this increase is abrupt and dramatic, with a final probability of > 0.6 across all models. In comparison, predictions on short sequences from Wikipedia and non-memorized idioms have a substantially lower probability of ~ 0.2 . This can be

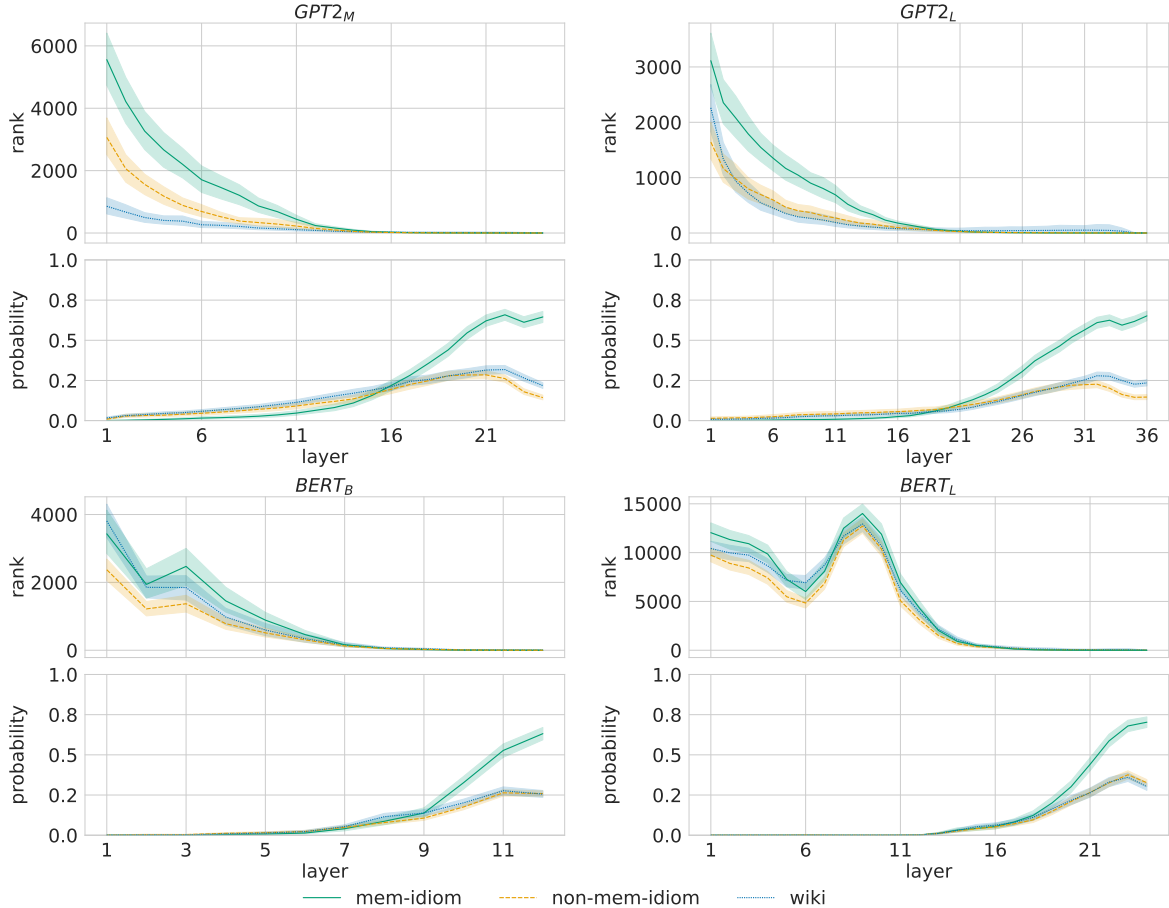


Figure 2: The predicted token’s probability and rank across layers, for memorized idioms (mem-idiom), non-memorized idioms (non-mem-idiom) and short sequences from Wikipedia (wiki). Memory recall exhibits two characteristic phases of candidate promotion and confidence boosting.

explained by the fact that memorized idioms have a single correct target, rather than many possible continuations, as in the instances from Wikipedia. In addition, low-probability predictions for non-memorized idioms are expected as the model did not memorize the idioms and does not know their continuation. In §B, we provide more fine-grained analysis of these trends via a log-scaled view of the prediction’s rank and a visualization of the ranks and probabilities for separate clusters of the memorized predictions.

We further verify that our extracted hidden-distribution features distinguish between memorized and non-memorized predictions by training linear classifiers over combinations of these features (details in §A). We observe that, indeed, our features enable separation between memorized and non-memorized predictions at high accuracy (77%-85% across models). Moreover, classifiers that use hidden distribution features are more accurate than those relying only on the model’s output. Over-

| | correlation | p-value |
|-------------------|-------------|--------------|
| GPT2 _M | -0.22 | $2.9e^{-21}$ |
| GPT2 _L | -0.18 | $5.4e^{-14}$ |
| BERT _B | -0.19 | $1.8e^{-15}$ |
| BERT _L | 0.15 | $3.5e^{-10}$ |

Table 4: Pearson correlation between the predicted token’s rank at the first layer and its general frequency in Wikipedia.

all, these findings provide a profile of memorized predictions, suggesting that the memorized information is retrieved in early layers at inference.

5.3 Testing the Roles of Different Layers Through Intervention

Our analysis in the previous section interprets hidden representations as distributions over the output vocabulary. We now conduct intervention experiments to verify that this interpretation is meaningful for studying memory recall, and to test layers’ roles in that process. Concretely, we zero out FFN




Figure 3: Intervention in non-dominant (left) and dominant (right) FFN sub-updates in GPT2_M. Each cell shows the percentage of memorized idioms for which the prediction was changed by zeroing-out the FFN sub-updates between the start and end layers.

sub-updates to the hidden representation (§4) and measure changes in memorized predictions.

A Short Primer on Transformer FFN Sublayers.

FFN sublayers are the final computation in transformer layers, which output the hidden distribution at the center of our analysis. In general, they have a key role in capturing knowledge in transformer LMs (Dai et al., 2022; Meng et al., 2022). We follow Geva et al. (2022) and view the computation of each FFN sublayer as a weighted collection of d_m sub-updates to the output distribution, each promoting a concept in the vocabulary space, e.g. “past-tense verbs” or “female athletes”. To understand this, consider the computation of the FFN at layer ℓ , given by $\text{FFN}^\ell(\mathbf{h}_i^\ell) = f(W_K^\ell \mathbf{h}_i^\ell) W_V^\ell$, where $W_K^\ell, W_V^\ell \in \mathbb{R}^{d_m \times d}$ are parameter matrices, d_m is the intermediate hidden dimension, and f is a point-wise non-linearity activation function. This computation can be decomposed as: $\text{FFN}^\ell(\mathbf{h}_i^\ell) = \sum_{j=1}^{d_m} f(\mathbf{h}_i^\ell \cdot \mathbf{k}_j^\ell) \mathbf{v}_j^\ell = \sum_{j=1}^{d_m} m_j^\ell \mathbf{v}_j^\ell$, where \mathbf{k}_j^ℓ and \mathbf{v}_j^ℓ are the j -th row in W_K^ℓ and the j -th column in W_V^ℓ , respectively. Geva et al. argue that each weight m_j^ℓ is the score assigned by the model for some textual pattern, and each vector \mathbf{v}_j promotes a concept that follows that pattern.

Experiment First, we sample 100 random instances from IDIOMEM that the model memorized. Then, for each range of up to 3 consecutive layers, we perform two complementary experiments, where we run GPT2_M’s inference on the 100 instances while intervening in the chosen layers to cancel the contribution of FFN sub-updates to the prediction. Specifically, for each layer ℓ in the layer range, we perform the following two interventions: first, we zero out (i.e. artificially set to 0 during inference) the weights of the 10 most dominant

sub-updates, which are known to be particularly salient for predictions (Geva et al., 2022) (there are d_m sub-updates per layer). Concretely, we sort the sub-updates by their contribution to the FFN output $|m_i^\ell| \|\mathbf{v}_i^\ell\| \forall i \in [1, \dots, d_m]$, and set $m_i^\ell = 0$ for the 10 sub-updates with the highest contribution. Next, we zero out non-dominant sub-updates, i.e. all the sub-updates *except* for the 10 most dominant ones. For each intervention, we measure how often it changes the predicted token. Further measurements of the change in rank and probability of the target token are reported in §C.

Results Fig. 3 shows, for each layer range, the percentage of memorized idioms where the predicted token has changed. Focusing on zeroing out non-dominant FFN sub-updates (Fig. 3, left), we observe a *two-phase pattern of decreasing “layer importance” which corresponds to the two-phase pattern of decreasing rank and increasing probability during inference (§5.2)*: layers’ effect on memory recall drops precipitously in the first 10 layers, corresponding to the candidate promotion phase. Then, from around layer 10 onwards, the drop in effect is much less steep, corresponding to the confidence-boosting phase. Intervention in upper layers rarely changes the predicted token, and its effect is limited to reducing the model’s confidence (§C). We further visualize this two-phase behavior in Fig. 7 in §C.

These findings suggest that the candidate promotion phase, while having a smaller effect on the prediction’s assigned probability compared to the later confidence-boosting stage, in fact, has a crucial role in memory recall.

We also observe that interventions in the first layer are by far the most destructive, with 100% change in prediction for non-dominant updates

(Fig. 3, left). Unlike for other layers, this is also observable when zeroing out the dominant sub-updates (Fig. 3, right), which constitute only 0.1% of layer sub-updates. This suggests the first layer is especially critical for memorized predictions.

6 Memorization of Factual Statements

We now examine if our findings generalize beyond idioms to other types of memory recall, focusing on the completion of statements expressing facts.

Data Datasets for evaluating memorization of factual knowledge typically contain simple queries such as “*The continent of Kuwait is*”, where predicting the next token requires knowledge of the triplet $\langle s, r, t \rangle$ where s is a source entity (e.g., Kuwait), t a target entity (e.g., Asia), and r is the relation between them (e.g., is the continent of). However, unlike idioms, such queries are not suitable for probing memorization since they often fail to satisfy the criteria in §2. Concretely, queries often include “clues” that could make the prediction easy to guess and based on generalization (Poerner et al., 2020) (e.g. predicting a Spanish-speaking country for the query “*Federico López was born in*”), and the same fact can be expressed in multiple different ways (e.g. “*Kuwait is a country in Western Asia*” also encodes the above fact).

To test memorization of facts, we, therefore, collect factual statements where such cases are excluded. We use LAMA-UHN (Poerner et al., 2020), a subset of the LAMA dataset (Petroni et al., 2019), where “easy-to-guess” queries are filtered out. LAMA comprises of queries structured as “fill-in-the-blank” cloze statements (e.g. “*Gordon Scholes is a member of the ____ political party.*”). To accommodate autoregressive LMs, we consider only queries where the blank appears at the end. In addition, we keep only queries with a single correct completion (based on LAMA). This turns our definition of memorized and non-memorized sets (§4) equivalent to separating based on the evaluation metric of LAMA: an instance is considered memorized if the model predicts the single correct completion and non-memorized otherwise. Overall, the resulting collection consists of 17,855 factual statements with 22 unique relations.

Memorized Facts Exhibit a Similar Prediction Process to Idioms We repeat our analysis (§5.1), using the collected factual statements and GPT2_M. Splitting the statements to those memo-

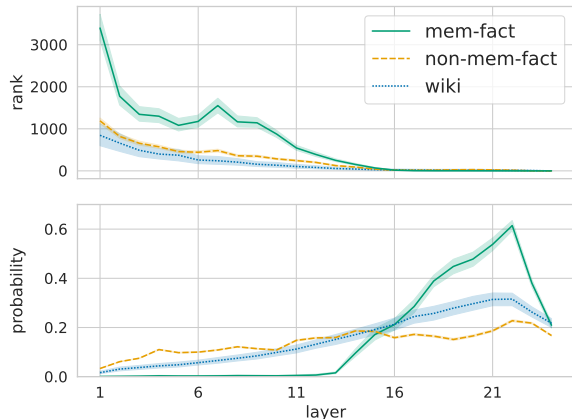


Figure 4: The predicted token’s probability and rank across layers of GPT2_M, for memorized (mem-fact) and non-memorized (non-mem-fact) facts and short sequences from Wikipedia (wiki).

rized and non-memorized by the model results in 786 mem-facts vs. 17,069 non-mem-facts statements.

Fig. 4 shows the rank and probability of the predicted token across layers. Like idioms, memorized facts exhibit a clear two-phase prediction process, where the prediction probability rapidly increases once the candidate reaches a low rank (at layer 16). This is in contrast to non-memorized facts and short sequences from Wikipedia, where the rank (probability) gradually decreases (increases) across layers without a distinct two phases.

Differences from Memorized Idioms Stem from Ill-Defined Targets There is one major difference compared to memorized idioms (Fig. 4 vs. Fig. 2 upper left), which is a substantial drop in probability (0.62 \rightarrow 0.21) in the last two layers. We hypothesize that this is because the input query has multiple plausible completions that were not specified as “correct” targets in LAMA. We verify this by manually analyzing predictions, and find that for 82 out of 100 queries there is more than one correct continuation in top 5 predicted tokens. We posit that the above deficiency of LAMA is inherent because, in violation of our criterion (§2), factual statements can usually be expressed in many ways so their prompt has no single correct target.

7 Related Work

Memorization as Training-Data Influence. Memorization in LMs has attracted immense attention due to their rich sensitive training data (Carlini et al., 2019; Song and Shmatikov, 2019; Carlini

et al., 2021; Zhang et al., 2021b; Carlini et al., 2022; Tirumala et al., 2022; Tänzer et al., 2022; Raunak and Menezes, 2022). Recent work suggests that memorization is *necessary* for performant ML due to the “long tail” of infrequently-observed patterns (Zhang et al., 2021a; Feldman, 2020; Brown et al., 2021; Raunak et al., 2021). This line of work has two key limitations: (a) only black-box behavior is measured rather than looking at models’ internal prediction process, and (b) it detects memorized instances by measuring the effect of their inclusion in the training set on inference behavior. This results in a set of memorized examples that is specific to the model, training data, and even training pass, making it difficult to build on these results in future research.

Transformers and Idioms. Nedumpozhimana and Kelleher (2021) showed that idioms are identified using textual cues within the expression; Dankers et al. (2022) showed that idioms tend to be internally processed as single units of meaning. It has also been known (Fakharian and Cook, 2021; Salton et al., 2016) that LM contextual embeddings encode information about whether or not an expression is idiomatic (vs. literal). (Shwartz and Dagan, 2019; Chakrabarty et al., 2022) also studied representations and interpretation of non-compositional sequences, such as idioms. No prior work used idioms to probe LM memorization, which is one of the main contributions of this work. We release our dataset, IDIOMEM, to facilitate future research on memorization recall in LMs. Diagnostic datasets, such as IDIOMEM, have often proven useful in the past (Sugawara et al., 2022; Parrish et al., 2021).

Memorization of Factual Knowledge. An extensive line of work (Petroni et al., 2019; Jiang et al., 2020; Poerner et al., 2019; Lewis et al., 2020; Elazar et al., 2021) studied LMs’ capacity to acquire relational knowledge during training. Some attention has also been given to understanding the inner workings of factual-memory recall: Wallat et al. (2020) showed that some facts are retrieved from the bottom and intermediate layers, and Meng et al. (2022) localized factual-knowledge recall within feed-forward-component computation. Since factual statements do not fulfill our criteria, it is difficult to convincingly argue that correct predictions indicate memory recall, making it impossible to use them to isolate the effect of memorization.

8 Conclusion

We introduce a methodological framework for detecting and analyzing memorized predictions in transformer LMs. This includes a set of criteria on textual sequences for probing memorized predictions, the IDIOMEM dataset of idioms fulfilling these criteria, and an interpretation method of prediction internals. We characterize a behavioral profile that is unique to predictions involving memory recall and is observable across different LMs and data types. By providing these fundamental tools and initiating a thread of research on the phenomena we observe, we hope to empower future work towards demystifying transformer memorization.

Limitations

Our criteria for detecting memorized instances are sufficient but not necessary, which raises the question of what other sequences that trigger memory recall satisfy them.

Additionally, while correct prediction for sequences that fulfill our criteria implies memory recall, incorrect prediction does not necessarily imply that no memory was used. This means that our set non-mem-idiom might include some portion of memorized sequences. This does not qualitatively affect our results as long as mem-idiom still contains *more* memorized predictions than non-mem-idiom, as is evidenced by the stark differences we observe between LMs’ internal behaviors on these sets.

Our work focuses on showing the utility of idioms for probing memorization, and opening up a new thread of research in this vein. We therefore leave further investigation of the above gaps for future work.

Acknowledgements

We thank Gabriel Stanovsky for helpful feedback. This project has received funding and resources from CIFAR through the Canada CIFAR AI Chair program, the Province of Ontario, companies sponsoring the Vector Institute, as well as the Computer Science Scholarship granted by the Séphora Berrebi Foundation, and the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement No. 802774 (iEXTRACT).

References

- Gavin Brown, Mark Bun, Vitaly Feldman, Adam Smith, and Kunal Talwar. 2021. When is memorization of irrelevant training data necessary for high-accuracy learning? In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 123–132.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan Zhang. 2022. Quantifying memorization across neural language models. *arXiv preprint arXiv:2202.07646*.
- Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. 2019. **The secret sharer: Evaluating and testing unintended memorization in neural networks**. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 267–284, Santa Clara, CA. USENIX Association.
- Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Úlfar Erlingsson, et al. 2021. **Extracting training data from large language models**. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650. USENIX Association.
- Tuhin Chakrabarty, Yejin Choi, and Vered Shwartz. 2022. **It’s not rocket science: Interpreting figurative language in narratives**. *Transactions of the Association for Computational Linguistics*, 10:589–606.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. **Electra: Pre-training text encoders as discriminators rather than generators**. In *International Conference on Learning Representations*.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. **Knowledge neurons in pretrained transformers**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502, Dublin, Ireland. Association for Computational Linguistics.
- Verna Dankers, Christopher Lucas, and Ivan Titov. 2022. **Can transformer be too compositional? analysing idiom processing in neural machine translation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3608–3626, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhishava 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.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2021. **A mathematical framework for transformer circuits**. *Transformer Circuits Thread*. <https://transformer-circuits.pub/2021/framework/index.html>.
- Samin Fakharian and Paul Cook. 2021. Contextualized embeddings encode monolingual and cross-lingual knowledge of idiomaticity. In *Proceedings of the 17th Workshop on Multiword Expressions (MWE 2021)*, pages 23–32.
- 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, STOC 2020*, page 954–959, New York, NY, USA. Association for Computing Machinery.
- Vitaly Feldman and Chiyuan Zhang. 2020. **What neural networks memorize and why: Discovering the long tail via influence estimation**. In *Advances in Neural Information Processing Systems*, volume 33, pages 2881–2891. Curran Associates, Inc.
- Mor Geva, Avi Caciularu, Kevin Ro Wang, and Yoav Goldberg. 2022. Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space. *arXiv preprint arXiv:2203.14680*.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. **Transformer feed-forward layers are key-value memories**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hessel Haagsma, Johan Bos, and Malvina Nissim. 2020. **MAGPIE: A large corpus of potentially idiomatic expressions**. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 279–287, Marseille, France. European Language Resources Association.
- Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. 2020. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Mihir Kale and Abhinav Rastogi. 2020. **Text-to-text pre-training for data-to-text tasks**. In *Proceedings of*

- the 13th International Conference on Natural Language Generation*, pages 97–102, Dublin, Ireland. Association for Computational Linguistics.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2022. [Deduplicating training data makes language models better](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8424–8445, Dublin, Ireland. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Inbal Magar and Roy Schwartz. 2022. [Data contamination: From memorization to exploitation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 157–165, Dublin, Ireland. Association for Computational Linguistics.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual knowledge in gpt. *arXiv preprint arXiv:2202.05262*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *International Conference on Learning Representations*.
- Timothee Mickus, Denis Paperno, and Mathieu Constanant. 2022. [How to dissect a Muppet: The structure of transformer embedding spaces](#). *Transactions of the Association for Computational Linguistics*, 10:981–996.
- Diego Moussallem, Mohamed Ahmed Sherif, Diego Esteves, Marcos Zampieri, and Axel-Cyrille Ngonga Ngomo. 2018. [LIdioms: A multilingual linked idioms data set](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Vasudevan Nedumpozhimana and John Kelleher. 2021. Finding bert’s idiomatic key. In *Proceedings of the 17th Workshop on Multiword Expressions (MWE 2021)*, pages 57–62.
- Alicia Parrish, Sebastian Schuster, Alex Warstadt, Omar Agha, Soo-Hwan Lee, Zhuoye Zhao, Samuel R. Bowman, and Tal Linzen. 2021. [NOPE: A corpus of naturally-occurring presuppositions in English](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 349–366, Online. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- 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.
- Nina Poerner, Ulli Waltinger, and Hinrich Schütze. 2019. Bert is not a knowledge base (yet): Factual knowledge vs. name-based reasoning in unsupervised qa. *arXiv preprint arXiv:1911.03681*, 3.
- Nina Poerner, Ulli Waltinger, and Hinrich Schütze. 2020. [E-BERT: Efficient-yet-effective entity embeddings for BERT](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 803–818, Online. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Vikas Raunak and Arul Menezes. 2022. Finding memo: Extractive memorization in constrained sequence generation tasks. *ArXiv*, abs/2210.12929.
- Vikas Raunak, Arul Menezes, and Marcin Junczys-Dowmunt. 2021. [The curious case of hallucinations in neural machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1172–1183, Online. Association for Computational Linguistics.
- Giancarlo Salton, Robert Ross, and John Kelleher. 2016. Idiom token classification using sentential distributed semantics. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 194–204.
- Prateek Saxena and Soma Paul. 2020. Epie dataset: A corpus for possible idiomatic expressions. In *Text, Speech, and Dialogue*, pages 87–94, Cham. Springer International Publishing.
- Vered Shwartz and Ido Dagan. 2019. [Still a pain in the neck: Evaluating text representations on lexical composition](#). *Transactions of the Association for Computational Linguistics*, 7:403–419.

- Congzheng Song and Vitaly Shmatikov. 2019. [Auditing data provenance in text-generation models](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, page 196–206, New York, NY, USA. Association for Computing Machinery.
- Saku Sugawara, Nikita Nangia, Alex Warstadt, and Samuel Bowman. 2022. [What makes reading comprehension questions difficult?](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6951–6971, Dublin, Ireland. Association for Computational Linguistics.
- Michael Tänzler, 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.
- Kushal Tirumala, Aram H Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. 2022. Memorization without overfitting: Analyzing the training dynamics of large language models. *arXiv preprint arXiv:2205.10770*.
- Jonas Wallat, Jaspreet Singh, and Avishek Anand. 2020. [BERTnesia: Investigating the capture and forgetting of knowledge in BERT](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 174–183, Online. Association for Computational Linguistics.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2021a. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115.
- Chiyuan Zhang, Daphne Ippolito, Katherine Lee, Matthew Jagielski, Florian Tramèr, and Nicholas Carlini. 2021b. Counterfactual memorization in neural language models. *arXiv preprint arXiv:2112.12938*.

A Distinguishing Memorization Using Hidden-Distribution Features

§5.2 shows differences in our extracted hidden-distribution features, namely the rank and probability across layers, between memorized idioms and non-memorized sequences. To verify that these features are distinguishing between memorized and non-memorized predictions, we build a classifier that receives them as input, as follows.

Experiment To answer the above, we represent every instance in IDIOMEM, for every LM we experiment with, as a sequence of probabilities and ranks assigned to the predicted token at each layer. This results in a feature vector for each instance in IDIOMEM for each of our LMs (GPT2_M, GPT2_L, BERT_B, and BERT_L). We then append a class label for each LM’s instances corresponding to whether it memorized them. Then, for each LM’s dataset, we perform 10-fold cross-validation with an 80%-20% train-test split to evaluate the accuracy of a logistic-regression classifier using the LogisticRegression classifier of scikit-learn,⁷ specifying L1 penalty and a bilinear solver as constructor parameters. After each split and before evaluating the classifier, we balance the test set by replacing the larger class with a random subsample the size of the smaller class. To isolate the distinguishing utility of ranks from that of probabilities, we repeat this process while only taking either of them as features at each time. We also repeat this process while separately considering just the last-layer probability as a single feature, the last hidden state vector, and finally, the ranks in layers of layers 1-12 (omitting ranks in layers 13-16 where ranks are usually 0) appended to all layer probabilities.

Results Results are given in Tab. 5. We observe that, across all models, most classifiers perform well over the 50%-accuracy baseline for their class-balanced test sets. The output probability alone is often highly distinguishing with around 78% accuracy, but the vector of 16 hidden-distribution probabilities seems to contain additional distinguishing information, as using it alongside ranks results in higher accuracy, typically around 84%. Using the ranks in addition to probabilities usually *decreases* accuracy, but omitting the ranks in layers 13-16 (which we know are mostly 0 as this is the confidence-boosting phase) attenuates this effect.

⁷<https://scikit-learn.org/>

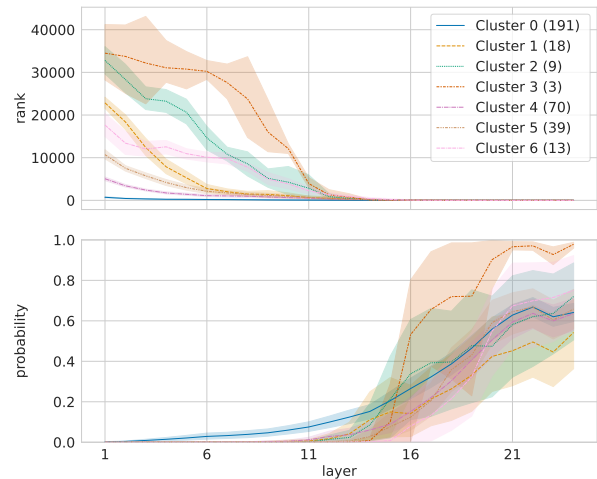


Figure 5: The predicted token’s probability and rank across layers for each cluster, after clustering the memorized idioms in GPT2_M according to rank and probability across layers.

We conjecture that ranks have little meaningful information, especially in the confidence-boosting phase.

B Fine-grained Analysis of Memorized Predictions

§5.2 shows how recall of memorized predictions exhibits two characteristic phases (specifically, see Fig. 2). To shed light on this phenomenon, we conduct additional analysis.

B.1 Memorized Idioms In-depth Breakdown

In our analysis, we address all memorized predictions jointly. We now check whether these averaged results are consistent across subgroups of the memorized idioms. To this end, we cluster the memorized idioms by GPT2_M, using the same hidden features as in §A, i.e., each instance is represented by the predicted token probability and rank across layers. We then cluster the idioms into seven groups, using K-mean clustering, and visualize the prediction rank and probability across layers for each group. We set the number of clusters to $k = 7$ based on manual inspection, and as we observed no substantial differences in the resulting clusters for larger values of k .

Results are shown in Fig. 5. We find that all groups exhibit the confidence-boosting phase, as the prediction probability quickly increases starting from the intermediate layers. Notably, the lowest final probability observed is > 0.5 , which is substantially higher than the average probability of

| | GPT2 _M | GPT2 _L | BERT _B | BERT _L |
|--------------------------------|-------------------|-------------------|-------------------|-------------------|
| probability | 84.6 ± 2.8 | 84.0 ± 2.3 | 84.2 ± 1.8 | 77.6 ± 3.3 |
| ranks | 63.2 ± 1.6 | 59.5 ± 2.7 | 73.3 ± 3.3 | 64.5 ± 3.6 |
| probability + ranks | 83.6 ± 2.7 | 82.9 ± 2.9 | 82.7 ± 2.2 | 76.9 ± 2.9 |
| ranks layer 1-12 + probability | 84.3 ± 2.8 | 83.4 ± 2.7 | 82.7 ± 2.2 | 77.2 ± 4.2 |
| probability last layer | 79.5 ± 3.5 | 83.3 ± 2.3 | 77.7 ± 2.3 | 74.2 ± 4.8 |
| final hidden state | 72 ± 2.1 | 72.5 ± 3.7 | 72.7 ± 2.9 | 66.0 ± 3.0 |
| token ids | 58 ± 0.2 | 59.4 ± 0.2 | 58.8 ± 0.2 | 71.6 ± 0.2 |
| random | 49.5 ± 3.8 | 49.7 ± 4.7 | 50.4 ± 4.9 | 50.1 ± 3.5 |

Table 5: Cross-validation accuracy of a logistic-regression classifier trained to distinguish between memorized and non-memorized idioms.

~ 0.2 for non-memorized predictions (§5.2). However, considering the prediction rank for the different groups, we observe a relatively large variation. Specifically, we observe that 55% of the instances (cluster 0) have a low initial rank. This further supports the findings in §A.

B.2 Log-Scale Visualization

Fig. 6 shows a log-scaled view of the graphs from Fig. 2. We observe that, in terms of orders of magnitude (vs. absolute value), the differences in initial ranks between memorized and non-memorized predictions are more minor, especially in BERT models, whereas rank differences measured in upper layers are more stark.

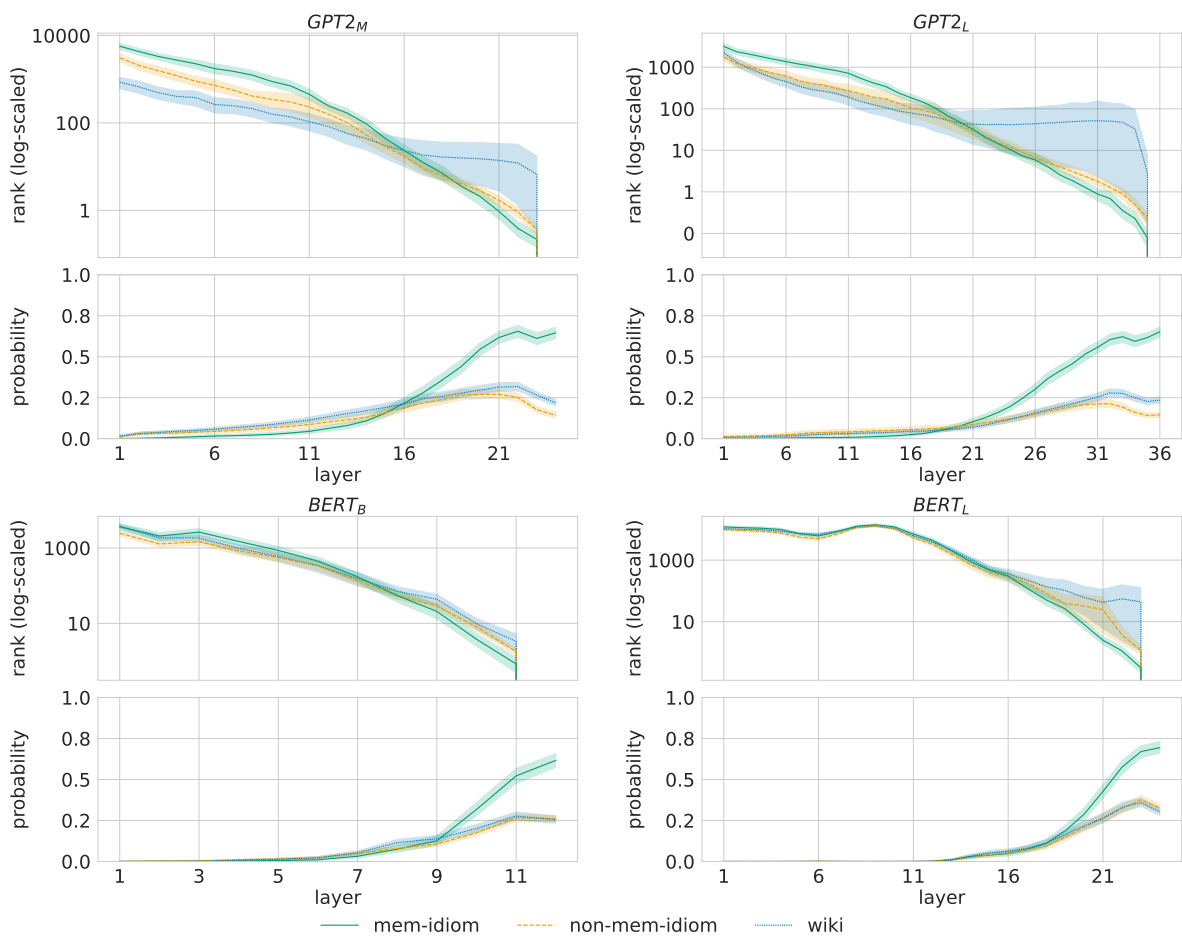


Figure 6: The predicted token’s probability and log-scaled rank across layers, for memorized idioms (mem-idiom), non-memorized idioms (non-mem-idiom) and short sequences from Wikipedia (wiki).

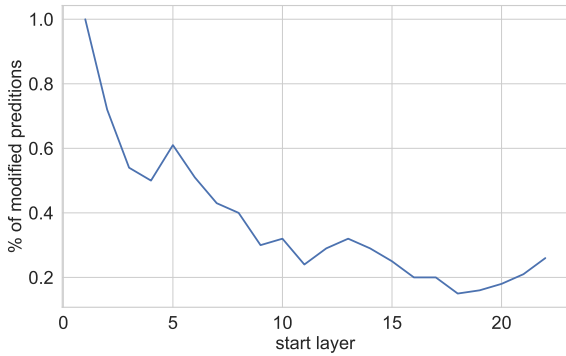


Figure 7: We visualize the effect of intervening in each 3-consecutive-layer range according to the procedure in §5.3.

C Intervention Experiments: Additional Analysis

C.1 Additional Visualization

In §5.3, we performed an intervention-based experiment to test the effect of zeroing out FFN sub-updates in layer computation. This produced a heat map of values corresponding to intervention’s effect for each layer range. Here, we plot the intervention’s effect across all 3-layer layer ranges. Note that since there are 24 layers, there are a total of 22 ranges of 3 consecutive layers. Fig. 7 visualizes the effect of intervening in each such range.

As discussed in §5.3, we observe a steep drop in effect in the first 10 layers, followed by a more leveled slope of decrease in the upper layers.

C.2 Analyzing Changes in Rank and Probability of the Target Token

In addition to measuring how often an intervention changes memorized predictions (§5.3), we further measure the average change in the rank and probability of the target token. Note that the original target rank for memorized predictions is always zero, as the target token is the top candidate in the original output distribution.

Change in target rank Fig. 8 shows the change in the target token’s rank for all intervention experiments. Overall, we observe similar trends as in §5.3. First, zeroing out either dominant or non-dominant FFN sub-updates in upper layers (layers 11-24) does not affect memory recall, as the target token is still ranked as the top candidate in the output distribution. Moreover, zeroing out in early layers (1-10) damages memory recall as the target rank

increases by > 20 positions. Specifically, zeroing-out non-dominant FFN sub-updates in layers 2-4 increases the target rank by 60, and disabling either dominant or non-dominant sub-updates in the first layer completely eliminates memory recall as the rank increases to > 6000 .

Change in target probability Fig. 9 shows the change in the target token’s probability for all intervention experiments. Unlike the prediction rank, which is mostly influenced by the lower layers during memory recall, the prediction probability is highly influenced by the intermediate and upper layers, where confidence boosting happens. Disabling FFN sub-updates in only three of these layers reduces the prediction probability by up to 33%. Considering the lower layers (1-9), zero-outs lead to a large probability decrease (up to 50%). This is expected since these interventions often change the prediction, i.e. they eliminate the target from the top of the output distribution.

D Experimental Setting Details

Tab. 6 shows the evaluated models’ hyperparameters.

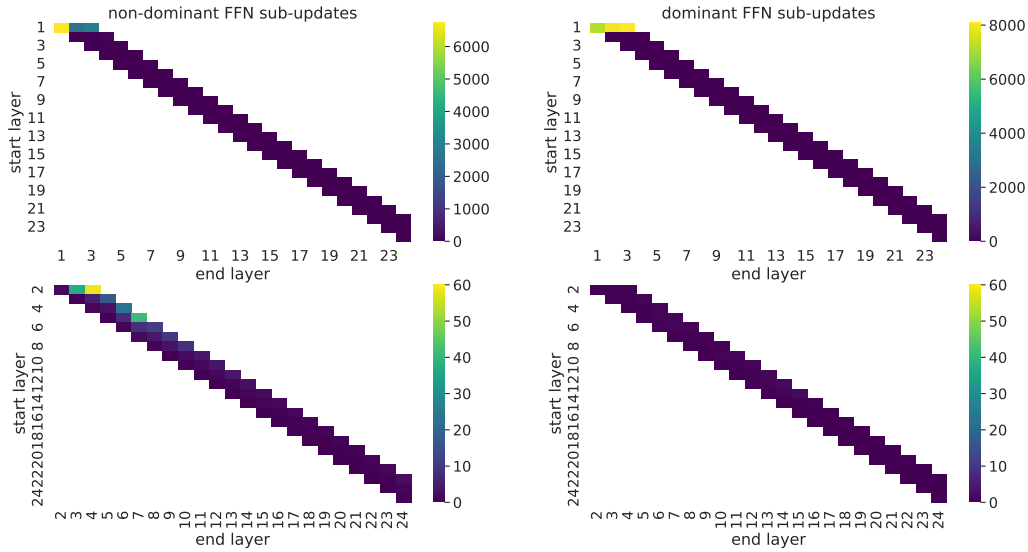


Figure 8: Change in the rank of the target token following intervention zeroing out non-dominant (left) and dominant (right) FFN sub-updates in GPT2_M. Each cell shows the average change in rank of the target token after zeroing out the sub-updates in the layers between the start and end layers. For readability, we provide plots with the first layer (top) and without (bottom).

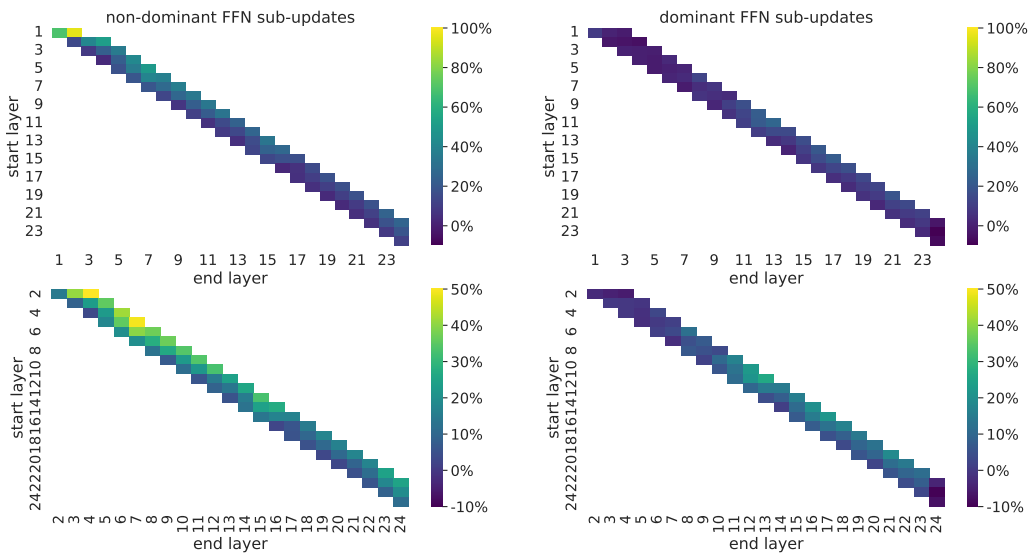


Figure 9: Change in the probability of the target token following intervention zeroing out non-dominant (left) and dominant (right) FFN sub-updates in GPT2_M. Each cell shows the average change in probability of the target token after zeroing-out the sub-updates in the layers between the start and end layers. For readability, we provide plots with the first layer (top) and without (bottom).

| | GPT2 _M | GPT2 _L | BERT _B | BERT _L |
|-----------------------------------|-------------------|-------------------|-------------------|-------------------|
| Layers | 24 | 36 | 12 | 24 |
| Model hidden dimensions (d) | 1024 | 1280 | 768 | 1024 |
| Feed-forward dimensions (d_m) | 4096 | 5120 | 3072 | 4096 |
| Attention heads | 12 | 20 | 12 | 16 |
| Parameters | 345M | 774M | 110M | 340M |
| Vocabulary size (# of tokens) | 50,256 | 50,256 | 30,522 | 30,522 |

Table 6: The models' hyperparameters.