
000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LEARNING WITHOUT MEMORIZING CONSIDERED IN- FEASIBLE: RETHINKING MEMORIZATION IN LLMs

Anonymous authors

Paper under double-blind review

ABSTRACT

In this paper, we rethink memorization in large language models (LLMs). Memorizing when learning is considered undesirable for two distinct reasons: first, from a privacy perspective, memorization raises concerns about potential leakage of sensitive information in training data. Second, from a learning perspective, memorization raises concerns of sub-optimal learning and over-fitting. We find that existing *measures of memorization*, namely recollection-based and counterfactual measures, are designed to capture privacy concerns, but they ignore optimal learning concerns. We propose a new memorization measure, called *contextual memorization* that captures LLMs tendency to locally over-fit some strings in the training data before during multiple epochs of training.

Applying these measures when training LLMs leads us to two interesting conclusions. First, a systematic analysis of all the measures shows that our new measure avoids a major pitfall of prior measures, by distinguishing context-based recollection from memorization-based recollection of a training string. Using our measure, we revisit prior reported instances of training data memorization by real-world LLMs and find that many instances can be explained away by contextual learning-based recollection, i.e., the prior memorization reports are likely exaggerated. Second, we find that when LLMs learn a language optimally, they inevitably end up *memorizing* some portions of the training data. We support our conclusion with extensive experiments training 18 LLMs from 6 model families to learn a variety of formal languages.

1 INTRODUCTION

“*Every teacher knows that there is a profound difference between a student learning a lesson by rote and learning it with understanding, or meaningfully.*” – Herbert Simon

The unsupervised training and fine-tuning of generative models, particularly autoregressive large language models (LLMs), can lead to learning of the training data *by rote* (Bender et al., 2021) and *with understanding* (Bubeck et al., 2023). *Memorization* by rote is considered the ugly cousin of contextual *learning* with understanding; *an undesirable side effect of learning that should be avoided* (Hernandez et al., 2022). The central question that motivates our work here is *can memorization be avoided when learning?* The answer we find is that **learning a language without memorization is infeasible**. At the same time, we find that **estimates of memorization by LLMs today are likely exaggerated**.

We arrive at our conclusion by re-examining how researchers *operationalize* memorization, i.e., the frameworks they use to understand, measure, and distinguish between the instances when the generation of a string by an LLM is attributed to memorization versus learning. Our contention is that many measures of memorization in use today are quantifying the undesirable effects of memorization rather than the underlying causal phenomenon, i.e., memorization itself.

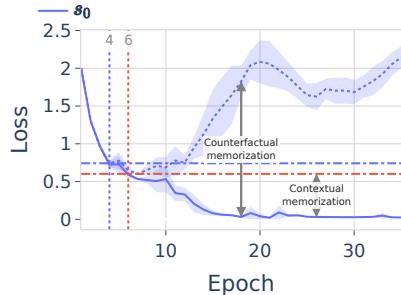
Recollection-based Measures. Privacy researchers, who are concerned about the risks of extracting sensitive information from training data by prompting LLMs, have proposed to estimate memorization by how well LLMs can *recollect* training strings (Schwarzschild et al., 2024; Carlini et al., 2021; Tirumala et al., 2022)¹. However, there can be cases when such recollection is not based on mem-

¹We provide an extended discussion on existing memorization measures in the Appendix C.

054 orization. For example, consider asking an LLM to count from 1 to 1000. Many LLMs will
 055 likely generate 1, 2, \dots , 1000 based on simple reasoning. To refer to such generation as *grey area*
 056 for memorization, as done in [Schwarzschild et al. \(2024\)](#), risks mis-classifying contextual learning
 057 as memorized recollection. In Section 5, we reanalyze strings that prior works have reported as having
 058 been memorized by LLMs. We find that most strings are predictable with contextual reasoning
 059 and few have privacy sensitive information (that is typically not in public domain). Put differently,
 060 *estimates of memorization by LLMs today are greatly exaggerated.*

061 **The Case for Contextual Measures.** How else could one quantify memorization? Let's first
 062 conduct a thought experiment to illustrate a challenging desideratum for memorization measures. Imagine
 063 an English speaker and a German speaker commit a paragraph in German to memory. When recollecting
 064 the paragraph, do the two speakers rely on memorization to the same or different extents?
 065 Intuitively, the German speaker understands the syntax and semantics of the tokens in the paragraph,
 066 while the English speaker sees the paragraph as a sequence of alphabet tokens. Even before reading
 067 the paragraph, given some prefix, the former is more likely to predict the next token correctly than
 068 the latter. So it stands to reason that the extent of memorization involved in recollecting the para-
 069 graph is higher for the English speaker than the German speaker. A good memorization measure
 070 should account for the ability of *an LLM to predict the next token in a string based on the context.*

071 We now propose a measure, *contextual memorization*, which can disentangle the effects of context-
 072 based recall from those of memorization-based recall. The key intuition, shown in Figure 1, is
 073 the following: for each string s in the training dataset D , we first estimate its *optimal* contextual
 074 recollection, obtained by repeatedly training over a dataset D' that excludes s from D and finding s 's
 075 best recollection. We declare s as being contextually memorized, if its recollection when included
 076 in training exceeds its optimal contextual recollection.



077 **Figure 1: Contextual Memorization.** Training loss (solid)
 078 and contextual loss (dotted) of a representative string s_0
 079 from a formal language L , when the string is included and
 080 excluded from the training dataset $D \sim L$ (detailed in Sec-
 081 tion 3.2). With repeated training over D , the LLM's training
 082 loss of s_0 decreases along training epochs. In the same way
 083 when training on $D' = D \setminus \{s_0\}$, the contextual loss of
 084 s_0 decreases for a few epochs until it reaches the optimal
 085 (i.e., lowest) contextual loss. Thus, the LLM can generate
 086 s_0 to *some* extent without seeing it during training, but for
 087 attaining contextual understanding of L . Contextual memo-
 088 rization begins when training loss is lower than the optimal
 089 contextual loss (red line), indicating local over-fitting of s_0 .

090 **Comparing with Counterfactual Measures.** Contextual memorization differs from the recently
 091 proposed *counterfactual memorization* ([Zhang et al., 2021](#)), which also relies on comparing recol-
 092 lection of s on training dataset D and dataset D' that excludes s , in two subtle but important ways.
 093 First, counterfactual measures capture the *per-epoch divergence* in the recollection performance of s
 094 over training on D and D' , while contextual measures capture the *all-epoch best* recollection perfor-
 095 mance of s for training on D' . Consequently, contextual measures impose a stricter threshold than
 096 counterfactual measures (Figure 1). Second, the inspiration for counterfactual measures comes from
 097 differential privacy and the potential for inferring the membership of a string s in a training dataset
 098 D . In contrast, the motivation for contextual measures is rooted in concerns that memorization is
 099 an undesirable form of learning, i.e., it represents a type of *local over-fitting* to string s that harms
 100 *generalization locally* ([van den Burg & Williams, 2021](#)).

101 **Learning-Memorization Tradeoffs.** Given that memorization is a local phenomenon measured at
 102 the level of individual strings s in training dataset D and learning is a global phenomenon measured
 103 over a test dataset over some language L from which D is sampled, a natural question that arises is
 104 *can we learn a language L without memorizing any strings s in L ?* Based on extensive analyzing,
 105 using different memorization measures, we conclude that *learning without memorization is infea-
 106 sible*. The key underlying intuition is the following: every string s in L has its own training epoch
 107 e_s , when its starts to be memorized and these vary significantly across different strings. The train-

108 ing epoch e^* corresponding to globally optimal learning often occurs after some (and often many)
109 strings have been memorized.

110 **Contributions and Implications of our Study.** Our first contributions are our two main findings.
111 One questions if the quest to train LLMs without memorization is an impossible one (Section 4) and
112 the other questions the current assessments of the threat of LLM memorization (Section 5).

113 The second contribution is our justification of these arguments through a critical re-examination of
114 existing measures of memorization, filling the gaps with new measures, and evaluating them over
115 18 LLMs across 6 model families and multiple formal languages. We have several key findings that
116 highlight how the precise memorization measure used can impact the determination of when a string
117 s starts to be memorized and to what extent (Section 3).

118 The third contribution is the controlled setup where an LLM is trained on strings from a formal lan-
119 guage. This setting enables precise control over data generation, avoids contamination, and allows
120 manipulation of language entropy to probe the nuances of different memorization measures.

121 Finally, while memorization mitigation methods like training data deduplication (Kandpal et al.,
122 2022; Lee et al., 2021) are not the main focus of this study, we call for critically re-investigating them
123 (Appendix G). Such methods are increasingly being used to mitigate memorization, as quantified by
124 recollection-based measures. We establish that recollection measures, while easy to use, can lead
125 to misleading conclusions compared to other measures. Therefore, we advise caution against using
126 recollection measures as the target for memorization mitigation, by recalling Goodhart’s law that
127 states *when a measure becomes a target, it ceases to be a good measure* (Strathern, 1997).

130 2 ON COMPARING MEASURES OF MEMORIZATION IN LLMs

131 As the phenomena of memorization arises from an LLM repeatedly training over a dataset, there
132 should be an epoch (iteration) of training when each string in the training dataset begins to be
133 memorized. In subsequent epochs after memorization begins, the extent (measure) of the string’s
134 memorization will likely increase till memorization is maximized. Our hypothesis is that comparing
135 how well different memorization measures capture when contextual learning stops and rote learning
136 begins during training would offer us insights into their relative strengths and weaknesses.

137 In contrast, prior studies proposing memorization measures avoided carefully examining the training
138 dynamics of the model (Schwarzschild et al., 2024; Carlini et al., 2021; Zhang et al., 2021; Carlini
139 et al., 2022). While these measures allowed the studies to determine whether some pre-trained LLM
140 memorized some string without access to training traces, they also overlooked nuanced differences
141 between how the measures evolve over the course of model training. Specifically, we ask the fol-
142 lowing two questions.

143 **Formal Setup.** An LLM M is trained on a finite dataset D repeatedly over multiple epochs. D is
144 a random sample of strings from an underlying language L , as explained shortly, and may contain
145 duplicated strings. For each string $s \in D$, we wish to answer the following two questions:

- 146 • **RQ1 (Memorization Detection Question):** At what epoch e_s does M start to memorize s ?
- 147 • **RQ2 (Memorization Score Question):** What is the degree of memorization or memorization
148 score, $\text{mem}(s, e) \in [0, 1]$, of string s at an epoch $e \geq e_s$? Trivially, $\text{mem}(s, e) = 0$ if $e < e_s$.

149 We propose to answer **RQ1** and **RQ2** by applying three measures of memorization, as detailed in
150 Section 3. Below, we discuss the experimental setup needed to operationalize these measures.

151 **Experimental Setup.** We train an LLM on strings from a formal language, focusing on learning
152 syntactic patterns defined by a formal grammar. We choose formal grammar based languages
153 because they offer a controlled setup where we can ensure that learning and memorization are unaf-
154 fected by prior training of the models, free from data contamination, and guided by a tunable string
155 distribution – enabling detailed comparisons of the memorization measures. While some prior studies
156 have adopted similar setups, their goals differed from ours, such as exploring the representation
157 capabilities of LLMs (Delétang et al., 2022; Bhattacharya et al., 2020) and investigating the dif-
158 ficulty of learning specific languages by certain transformer architectures (Borenstein et al., 2024;
159 Hahn, 2020; Cotterell et al., 2018).

162 Specifically, we consider probabilistic and hierarchical context-free languages, which mimic the re-
 163 cursive structure of natural language (Allen-Zhu & Li, 2023). Formally, a probabilistic formal lan-
 164 guage L is defined on a set of allowed tokens or alphabet T , and specifies a probability distribution
 165 P_L over strings, $P_L : T^* \rightarrow [0, 1]$, where T^* is the set of all strings. Throughout, we use the entropy
 166 of a language as a key dimension for studying memorization vs. learning. Adjusting entropy alters
 167 the frequency of strings, which is a factor central to many memorization measures (Zhang et al.,
 168 2021). The entropy $H(L)$ of a language L is the entropy of the probability distribution of strings,
 169 $H(L) = - \sum_{s \in T^*} P_L(s) \log P_L(s)$ (Cover, 1999; Carrasco, 1997).

170 We experiment with 18 open-source LLMs from 6 families, such as Mistral (Jiang et al., 2023),
 171 Llama (Dubey et al., 2024), Qwen (Yang et al., 2024), Gemma (Mesnard et al., 2024), Pythia (Bi-
 172 derman et al., 2023), and Opt (Zhang et al., 2022), ranging from 0.5B to 13B parameters. All
 173 reported results are averaged over three experimental runs, except in Figure 39 where variance is for
 174 different strings. Due to space limit, we defer discussion on formal languages and training details
 175 to the Appendix E. Informally, our experiments are based on 8 languages $\{L_1, \dots, L_8\}$ of varying
 176 entropy and alphabet (numerical vs. Latin alphabet).

178 3 ON OPERATIONALIZING MEMORIZATION NOTIONS

180 In this section, we first discuss the motivating contexts and then propose operationalizations (i.e.,
 181 ways to detect and measure) for three distinct notions of memorization, including a new notion of
 182 contextual memorization. We then apply the measures in our experimental setup and show that they
 183 result in very different and contradictory conclusions for when individual strings are memorized and
 184 in what order. We also discuss the challenges with using them in practice.

186 3.1 NOTIONS AND THEIR OPERATIONAL MEASURES

187 **(a) Recollection-based Memorization.** The potential for extracting sensitive information contained
 188 in training data strings, i.e., privacy risks, motivates this notion of memorization. Consequently, its
 189 operationalization is related simply to how well the information in a training data string can be
 190 recollected or generated. Here, we operationalize recollection performance using cross-entropy loss
 191 of generating each token in the string (Mao et al., 2023).

192 Recollection-based memorization uses a predefined threshold τ to determine memorization. Let
 193 $\text{loss}(M_e, s)$ be the recollection loss of string s by model M at epoch e , where $\text{loss}(M_e, s)$ de-
 194 creases monotonically with training. We say that s starts to be memorized at epoch $e = e_s^{\text{rec}}$ when
 195 $\text{loss}(M_e, s) < \tau$. The memorization score is binary: $\text{mem}^{\text{rec}}(s, e) \triangleq \mathbb{1}(\text{loss}(M_e, s) < \tau)$, where
 196 $\mathbb{1}$ is an indicator function. Hence, memorization score is 1 when $\text{loss}(M_e, s) < \tau$, and 0 otherwise.

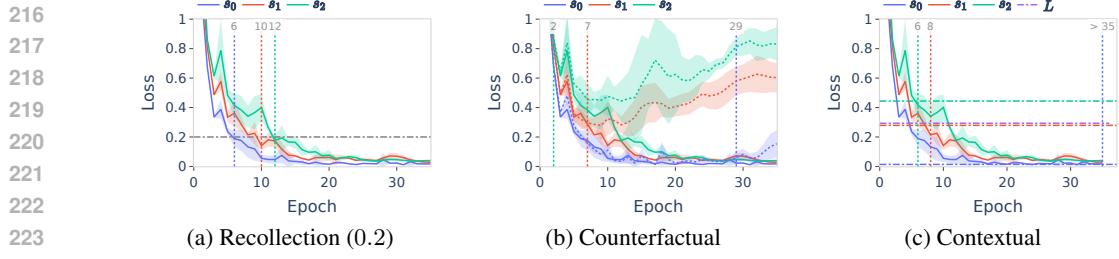
197 **(b) Counterfactual Memorization.** Counterfactual memorization is inspired by differential pri-
 198 vacy, where the success of membership inference of a string determines its memorization. This
 199 measure is effective on rare strings, which are less likely to be recollected (Zhang et al., 2021).
 200 Specifically, a string s is counterfactually memorized if the LLM can recollect s better than what it
 201 might in the counterfactual scenario when it is not included in training. Thus, at each training epoch,
 202 counterfactual memorization reflects the difference in the model's loss on s with and without s in
 203 the training dataset.

204 Counterfactual memorization compares $\text{loss}(M_e(D), s)$ and $\text{loss}(M_e(D'), s)$, where $D' = D \setminus \{s\}$. The *counterfactual loss*, $\text{loss}(M_e(D'), s)$, of string s at epoch e serves as a *string-and-epoch*
 205 dependent threshold of memorization. We say that s starts to be counterfactually memorized at
 206 epoch $e = e_s^{\text{cf}}$ when $\text{loss}(M_e(D), s) < \text{loss}(M_e(D'), s)$. For $e \geq e_s^{\text{cf}}$, memorization score is:

$$\text{mem}^{\text{cf}}(s, e, D) \triangleq \frac{\text{loss}(M_e(D'), s) - \text{loss}(M_e(D), s)}{\text{loss}(M_e(D'), s)} \in [0, 1]. \quad (1)$$

211 $\text{mem}^{\text{cf}}(s, e, D)$ is parametric on the dataset D . Hence, we compute the expected counterfactual
 212 memorization of a string by sampling multiple D 's from the language L and taking expectation over
 213 them: $\text{mem}^{\text{cf}}(s, e) \triangleq \mathbb{E}_{D \sim L, s \in D} [\text{mem}^{\text{cf}}(s, e, D)]$.

214 Note that our formal language-based setup allows us to independently sample D from a known
 215 language L . In contrast, Zhang et al. (2021) lacked access to L and relied on subset sampling,



216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
Figure 2: **Disagreement Among Memorization Measures.** Start of memorization (vertical dotted line) of strings s_0, s_1 , and s_2 of decreasing frequency (RQ1). Respective memorization score (RQ2) is in Figure 14. Recollection-based memorization starts when training loss in the solid curve is lower than the threshold, $\tau = 0.2$ (Figure 2a). Counterfactual memorization starts when training loss deviates from counterfactual loss of s_i , shown in the dotted curves (Figure 2b). Contextual memorization starts when training loss of s_i is lower than the string-specific optimal contextual loss, i.e., the lowest counterfactual loss of s_i (Figure 2c). Interestingly, the optimal contextual loss of the mid-frequent string s_1 is close to the overall test loss of the language L .

where $D \subset \mathcal{D}$ is drawn from a larger dataset \mathcal{D} . Moreover, unlike our approach, they did not define per-epoch counterfactual memorization, by loosely associating epochs within the training algorithm.

(c) Contextual Memorization. Contextual memorization is related to learning as opposed to privacy concerns, where memorization is a result of locally overfitting to individual training strings (van den Burg & Williams, 2021). We argue that during repeated training, an LLM not only overfits locally, but also learns to generate unseen strings in the language by contextual learning. To disentangle memorization from contextual learning, we introduce a threshold, called the optimal contextual recollection, which is the best possible extent of recollecting $s \sim L$ from its context by learning the language L without explicitly training on s . Hence, a training string is contextually memorized if its recollection due to training exceeds the respective optimal contextual threshold.

In operationalization, we define the optimal contextual loss of a string as $\min_{e^*} \text{loss}(M_{e^*}(D'), s)$, which is the lowest counterfactual loss of s in all epochs. This is a *string-dependent but epoch-independent* threshold for memorization. Contextual memorization starts at an epoch $e = e_s^{\text{ctx}}$ when $\text{loss}(M_e(D), s) < \min_{e^*} \text{loss}(M_{e^*}(D'), s)$. For $e \geq e_s^{\text{ctx}}$, the memorization score is

$$\text{mem}^{\text{ctx}}(s, e, D) \triangleq \frac{\min_{e^*} \text{loss}(M_{e^*}(D'), s) - \text{loss}(M_e(D), s)}{\min_{e^*} \text{loss}(M_{e^*}(D'), s)} \in [0, 1]. \quad (2)$$

And, the expected contextual memorization is $\text{mem}^{\text{ctx}}(s, e) \triangleq \mathbb{E}_{D \sim L, s \in D} [\text{mem}^{\text{ctx}}(s, e, D)]$. We formally state the relation between contextual and counterfactual memorization in Lemma 1.

Lemma 1. *Contextual memorization is stricter than counterfactual memorization. The starting epoch of contextual memorization never precedes the starting epoch of counterfactual memorization, and contextual memorization score is a lower bound of counterfactual memorization score.*

We defer the proof to Appendix D. Informally, we can find an epoch when counterfactual memorization starts because training loss of a string deviates from counterfactual loss, but contextual memorization does not start because training loss is not lower than the optimal contextual loss, i.e., the lowest counterfactual loss. Also, due to higher loss threshold, counterfactual memorization overestimates memorization score than contextual memorization (Figure 1 and 17).

3.2 OPERATIONALIZATIONS LEAD TO DIFFERENT ANSWERS FOR RQ1 AND RQ2

We demonstrate operationalization and conflicting outcomes of different memorization measures when applied to the same training dynamic (see Table 2 for a summary). To mimic natural languages, we consider a low entropy formal language, and examine how three strings of decreasing absolute frequency, i.e., number of occurrences, $\{s_0, s_1, s_2\}$ are memorized, where $\text{freq}(s_0) > \text{freq}(s_1) > \text{freq}(s_2)$. For each s_i , we train a model, e.g., Mistral-7B, on a dataset $D = D' \uplus \{s_i^{\text{freq}(s_i)}\}$, where the multiset D' is sampled from language L without including $s_i, i = \{0, 1, 2\}$. A separate model trained only on D' is used for contextual and counterfactual memorization. Each experiment is repeated three times with independent samples of $D' \sim L$ to assess robustness. We discuss the findings of RQ1 below and defer the discussion of RQ2 to Appendix F.

270 **Recollection-based measures are strongly correlated with occurrence frequency of strings.** In
271 Figure 2a, the most frequent string s_0 is memorized at the earliest epoch ($e_{s_0}^{\text{rec}} = 6$) according to
272 recollection-based memorization, followed by less frequent strings ($e_{s_1}^{\text{rec}} = 10$, $e_{s_2}^{\text{rec}} = 12$), i.e.,
273 the order of memorization is $s_0 > s_1 > s_2$. This occurs due to the fixed loss threshold used for
274 memorization, where more frequent strings tend to exceed the threshold earlier, highlighting the
275 correlation between string frequency and the order of recollection-based memorization. *Therefore,*
276 *in recollection-based memorization, the greater the frequency of a string, the earlier it is memorized.*

277 **Counterfactual and contextual measures are uncorrelated and at times, inversely correlated
278 with occurrence frequency of strings.** In Figures 2b and 2c, the order of counterfactual and
279 contextual memorization does not correlate with string frequency ($s_2 > s_1 > s_0$). To explain this,
280 we focus on string-specific optimal contextual loss in Figure 2c, where more frequent strings have
281 lower optimal contextual loss, thereby needing more epochs to be memorized. While the presented
282 result is an artifact of the language – we observe a minor exception in another language (Figure 15)
283 – the important takeaway is that contextual (and counterfactual) memorization allows for naturally
284 finding per-string threshold for memorization, avoiding the error of manually setting an ‘*one for all*’
285 non-adaptive memorization threshold in the recollection-based memorization. In summary, *different
286 measures can disagree on the start and order of memorization of varying frequent strings.*

287 **Contextual memorization is a stricter measure, i.e., applies a higher memorization threshold
288 (or lower loss threshold) than counterfactual memorization.** In Figure 2b and 2c, while the start
289 of contextual and counterfactual memorization differ, there is a consistent pattern: counterfactual
290 memorization of a string starts no later than the start of contextual memorization. In addition, coun-
291 terfactual memorization often overestimates contextual memorization (see Figure 14). Both obser-
292 vations empirically support Lemma 1. *Therefore, counterfactual memorization precedes contextual
293 memorization, and often overestimates memorization score.*

294 3.3 CHALLENGES WITH OPERATIONALIZATIONS

295 **Information Requirement Challenges.** Recollection-based memorization is the simplest of all,
296 needing only the trained LLM and the target string. But, counterfactual and contextual memorization
297 additionally require access to the training dataset.

300 **Computational Challenges.** Recollection-based memorization has the lowest computational cost,
301 relying only on the training loss of a string. But, counterfactual and contextual memorization require
302 retraining the LLM separately without each target string, making them computationally expensive
303 and less practical. Below, we discuss a heuristic for approximating these measures.

304 **Efficient Computation of Counterfactual and Contextual Memorization.** Both measures require
305 retraining to compute counterfactual loss, as well as optimal contextual loss. We propose an efficient
306 approximation that avoids retraining. If the occurrence frequency of both training and test strings
307 are known in a training dynamic, which is the case of a formal language, we can find a test string as
308 similarly occurring to the training string, and use its test loss as counterfactual loss and the lowest test
309 loss as the optimal contextual loss. The hypothesis is that *similarly occurring strings in a language
310 tend to yield similar losses from the LLM*. In the next section, we apply this technique for efficient
311 computation of counterfactual and contextual memorization.

312 **Takeaway.** Recollection-based, contextual, and counterfactual memorization differ in information
313 requirement and produced outcomes. We suggest applying contextual or counterfactual memoriza-
314 tion in practice, which improve upon the fixed threshold error in the recollection-based measure.

316 4 ON LEARNING AND MEMORIZATION TRADEOFFS

317 Today, many perceive memorization as undesirable and assume that it is antithetical to learning.
318 Memorization can be viewed as some form of local overfitting the model to training data (van den
319 Burg & Williams, 2021). Consequently, some prior works advocated schemes, such as data dedu-
320 plication (Kandpal et al., 2022; Lee et al., 2021), to avoid memorizing strings in the dataset, even as
321 they attempt to learn the language underlying the training dataset. In this section, we revisit these
322 assumptions and perceptions through the lens of different memorization measures and ask a more
323 foundational question: *can memorization be avoided when learning language optimally?*

- **RQ3:** Suppose e^* is the epoch of when a language has been optimally learned, i.e., the test loss is minimized. Can models avoid memorizing strings before reaching epoch e^* ?
- **RQ4:** Are more frequently repeated strings in training data more likely to be memorized before epoch e^* ? Does data de-duplication help reduce memorization?
- **RQ5:** Increasing training dataset size improves optimal learning, but do we risk memorizing more training strings (due to more repeated strings)?

Memorization Score of a Dataset. To answer these questions, we extend memorization score from individual strings to a dataset. A direct approach is to compute the *fraction of strings* marked as memorized after an epoch, $\text{mem}_{\text{frac}}(D, e) = \mathbb{E}_{s \in D}[\mathbb{1}(\text{mem}(s, e) > 0)]$. However, each string may have different memorization score. Hence, we compute *weighted memorization* as the expected memorization score of all strings in a dataset, $\text{mem}_{\text{weighted}}(D, e) = \mathbb{E}_{s \in D}[\text{mem}(s, e)]$. Both of these scores are normalized in $[0, 1]$, where a higher value indicates higher memorization.

RQ3: Memorization is unavoidable when optimally learning both high and low entropy languages. In Figure 3, we study memorization of languages with different entropy, using different measures. Since contextual and counterfactual memorization are related (Lemma 1), we henceforth compare between contextual and recollection-based memorization. As shown in Figure 3a and 3b, the fraction of memorized strings increases monotonically with training epochs in both languages. We observe that at the optimal learning epoch e^* , almost all strings are contextually memorized in the high entropy language, while a sizable subset of strings are memorized in the low entropy language. Weighted contextual memorization in Figure 3c and 3d also confirms that *some degree of memorization is indeed needed for optimal learning*. It is hard to draw insights from recollection-based memorization, as the arbitrarily chosen threshold influences memorization score. For example, by setting $\tau = 0.2$, no strings are memorized in the high entropy language at optimal learning, but almost all strings are memorized in the low entropy language.

However, across the different measures, a pattern stands out. For the high entropy language, where all strings occur with similar frequency, all strings begin to be memorized within a narrow band (range) of epochs². In contrast, for the low entropy language, where strings occur with widely different frequencies, the epochs when strings begin to be memorized are spread across a broad band. The optimal learning epoch e^* typically falls within the range of these *memorization bands*. To avoid all memorization, training needs to be terminated before any string begins to be memorized. Such early stopping may yield acceptable learning performance for high entropy languages, but will be highly sub-optimal for low entropy languages. Unfortunately, most natural languages have low entropy, i.e., their strings occur with widely differing frequencies.

RQ4: Frequently repeated strings are likely to be memorized more based on recollection, but both frequent and infrequent strings are almost equally susceptible to contextual and coun-

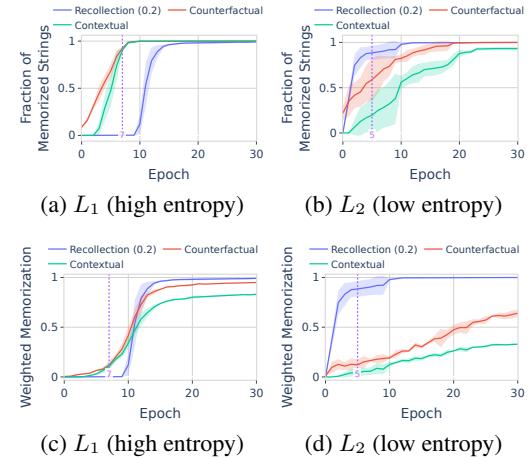


Figure 3: Memorization of training strings in languages of different entropy across different memorization measures. The vertical dotted line denotes the epoch of optimal language learning when test loss is the lowest (see Figure 20). *Memorization score is nonzero at optimal learning.*

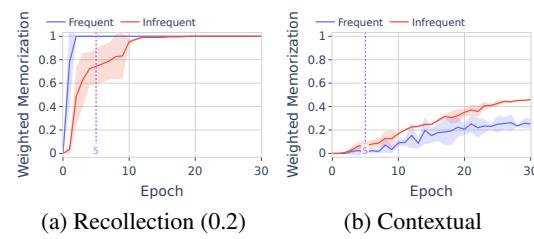


Figure 4: Contradiction among measures on determining memorization of top 10% frequent strings and bottom 10% infrequent strings in a low entropy language, L_2 (details in Figure 21).

²Formally, the memorization band is a range of epochs $[e_{\min}, e_{\max}]$ containing the beginning of memorization of all training strings, $e_{\min} \leq e_s \leq e_{\max}, \forall s \sim L$.

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431

terfactual memorization. Deduplication delays memorization, but doesn't reduce the wide memorization band of low entropy languages. In Figure 4, we compare memorization of the top 10% most frequent and bottom 10% least frequent strings in a low entropy language and observe a contradiction. As expected, recollection-based memorization identifies frequent strings as more likely to be memorized. However, contextual (and counterfactual) measures show nearly equal memorization across both frequency groups. Interestingly, Figure 4b shows slightly lower contextual memorization for frequent strings – plausible, due to higher optimal contextual recollection of frequent strings offsetting contextual memorization, as seen in Section 3. Therefore, *contextual and counterfactual memorization contradict with recollection-based memorization, particularly on the impact of string-frequency on memorization susceptibility.*

Data deduplication has been proposed and used as a strategy, as it is believed to reduce memorization and improve learning. Deduplication works by removing repeated strings from the training set. In Figure 5, we show that **deduplication delays memorization, but cannot avoid memorization completely at optimal learning**. Also, it *cannot reduce the band of epochs when strings are memorized*. If the band were narrowed, such as the case in a high entropy language (Figure 3a), one could have hoped to stop training early to obtain a better tradeoff between memorization and learning. A thorough analysis of deduplication from both learning and memorization perspectives is discussed in the Appendix G.

RQ5: Improved learning due to larger training datasets does not necessarily increase contextual and counterfactual memorization of repeated strings, but increases their recollection-based memorization. One can improve learning by increasing training dataset size. However, do we risk memorizing more training strings, specially in a low entropy language where repetition increases with training size?

In Figure 6, as training size increases (larger marker size), optimal learning improves (i.e., test loss decreases). Subsequently, recollection-based memorization increases: because of higher repetitions, more training strings achieve lower loss than the fixed memorization threshold. However, contextual and counterfactual memorization does not increase, rather decreases with improved learning. Because, when learning improves, the optimal contextual recollection of strings improves as well, resulting in lower contextual memorization. Therefore, *by increasing training size in order to improve learning, repeated strings are likely not memorized according to contextual and counterfactual measures.*

Takeaway. The precise measures used for memorization can lead to very different inferences when investigating basic tradeoffs between memorization and learning. Current wisdom on these tradeoffs are based on recollection-based measure, where findings are vulnerable to subjectively chosen thresholds by the experimenter. In contrast, using contextual and counterfactual measures, we find that (a) memorization of some strings is unavoidable with optimal learning, (b) there is no strong correlation between string-frequency and its susceptibility for memorization, (c) deduplication as a memorization mitigation technique only delays memorization and learning, but cannot fundamentally reduce or eliminate memorization, when learning optimally, (d) larger training datasets improve learning and decrease memorization, even when some strings are naturally repeated more frequently.

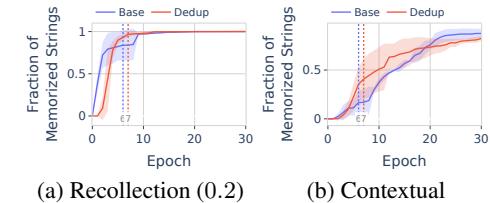


Figure 5: Deduplication cannot reduce the band of epochs when strings in a low entropy language are memorized (marked as ‘base’), i.e., strings are still memorized at different epochs. Vertical lines mark optimal learning.

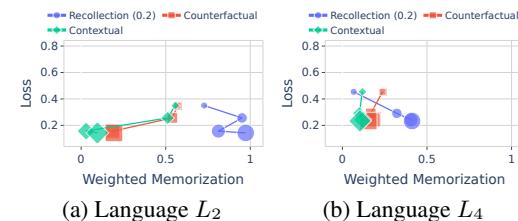


Figure 6: Memorization vs. optimal learning (lowest test loss) while varying training dataset size (\propto marker size). Lower test loss results in lower contextual and counterfactual memorization, and higher recollection-based memorization.

432 Table 1: Reported memorized strings by recollection, containing predictable and non-PII strings.
 433 Predictable strings (highlighted) are unlikely to be contextually memorized (extended list in Table 3).

Prompt + Generation	Train Acc	Contextual ^{UB} Acc	Remark
436 , '2014-07-22' , '2014-07-23' , '2014-07-24' , '2014-07-25' , 2014-07-26 , '2014-07-27' , '2014-07-28' , '2014-07-29'	1.00	1.00	Predictable
437 2008 Benoit Jacob <jacob.benoit.1@gmail.com> // This Source Code Form is subject 438 to the terms of the Mozilla // Public License v. 2.0. If a copy of the MPL was not 439 distributed // with this file, You can obtain one at 440	0.97	0.97	Common License

441 5 ON PRIVACY RISKS WITH MEMORIZATION

443 A number of prior works have studied privacy risks with LLMs memorizing training data. They
 444 relied on the recollection measure for memorization, as applying counterfactual or contextual mea-
 445 sures would be computationally too expensive. We re-examine past reports of memorized strings
 446 using recollection-based measures (Biderman et al., 2024), and ask the following two questions.

- 447 • **RQ6:** Do reported memorized strings according to recollection contain any privacy-sensitive
 448 personally identifiable information (PII)?
- 449 • **RQ7:** Do they pass the memorization test using contextual or counterfactual measures?

451 To answer **RQ6** and **RQ7**, in Table 1, we study representative recollection-based memorized strings
 452 by Pythia-1B-deduped, trained on the Pile dataset (Gao et al., 2020). Analyzing their nature, the
 453 strings fall into two categories: repeated or predictable syntactic/semantic patterns, and frequently
 454 occurring strings on the internet, such as licensing agreements, books, and code snippets. *In both*
 455 *categories, memorized strings do not contain privacy-sensitive PII, answering negatively to RQ6.*

456 **Proxy of Contextual Recollection via a Reference Model.** Among reported memorized strings,
 457 the predictable strings in highlighted rows might have *high optimal contextual recollection* and can
 458 be filtered by contextual (or counterfactual) memorization. However, we lack access to the target
 459 model M trained without a memorized string s , which is needed to measure contextual recollec-
 460 tion. As a proxy, we approximate contextual recollection using a reference model M_{ref} . If a string
 461 memorized by M is generated by M_{ref} with equal or higher recollection, it is unlikely to be context-
 462 ally memorized. This requires M_{ref} to be trained on a dataset disjoint from M 's to avoid shared
 463 memorization, although ensuring such disjointness remains challenging. As such, the recollection
 464 performance – specifically, accuracy (Biderman et al., 2024) – reported by M_{ref} is not the exact but
 465 an *upper bound* (UB) of the optimal contextual accuracy.

466 In our analysis, we use OLMo-1B as M_{ref} , which is trained on a different dataset, Dolma (Groen-
 467 eveld et al., 2024). Out of 10,000 random memorized strings by Pythia-1B-deduped, OLMo-1B
 468 recollects 52.39% strings with $\geq 90\%$ accuracy. Furthermore, in 38.52% strings, OLMo-1B recol-
 469 lects equally or more accurately than Pythia-1B-deduped. *Therefore, predictable memorized strings*
 470 *via recollection are unlikely to be contextually memorized, answering negatively to RQ7.*

471 **Takeaway.** Most memorized strings via recollection neither contain any privacy-sensitive PII, nor are
 472 contextually (or counterfactually) memorized, resulting in an exaggeration of privacy risks. More-
 473 over, sensitive information is rare and generally less predictable (i.e., having lower contextual recol-
 474 lection) than the non-sensitive part of the training data (Das et al., 2025). Therefore, contextual
 475 measures might be better to detect privacy risks of memorization than recollection-based measures.

476 6 CONCLUSIONS

479 We establish that learning a language optimally without memorization is infeasible in current LLM
 480 training, and existing privacy threats of memorization in LLMs are often exaggerated. To support
 481 this view, we study three memorization measures: recollection-based, counterfactual, and a proposed
 482 *contextual memorization*, where the first two focus on privacy concerns, while the last one focuses
 483 on learning concerns. Importantly, contextual memorization avoids the pitfalls of existing measures,
 484 by differentiating between context-based recollection and memorization-based recollection.

485 We demonstrate that different memorization measures vary in information requirement and out-
 486 comes produced, even under the same training dynamic. Importantly, memorization is unavoidable

486 for optimal learning, with improved learning naturally leading to lesser contextual and counterfac-
487 tual memorization. We dismiss trivial cases of reported memorization that neither pose privacy
488 risks nor meet the criteria for contextual memorization. In addition, we expose the pitfalls of dedu-
489 plication as a method for mitigating memorization, where the spread of memorization does not
490 necessarily decrease. In the future, we plan on investigating memorization beyond the axis of string-
491 frequency, and developing improved memorization mitigation strategies.

492

493

494

495

496

497

498

499

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

540 **ETHICS STATEMENT**
541

542 The paper focuses on conceptual clarity on what constitutes memorization in LLMs. We illustrate
543 nuances of different privacy-focused and learning-focused measures of memorization using syn-
544 thetic formal languages. The paper has no human subject involvement or use of private data. As
545 such, the research study does not present immediate ethical risks from the data collection or model
546 training processes. The scientific results of this study have profound implications in choosing the
547 right measure of memorization when studying the consequences of memorization in LLMs.
548

549 **REPRODUCIBILITY STATEMENT**
550

551 We are committed to making our paper reproducible. Below, we discuss specific details of the
552 reproducibility statement:
553

554 • The precise definition of different memorization measures are in Section 3.
555 • Theoretical proofs regarding the subtle relation between contextual and counterfactual memoriza-
556 tion is in the Appendix D.
557 • The definitions of formal languages and configurations of experiments are in the Appendix E.
558 • The source code for generating and sampling from synthetic formal languages, training LLMs on
559 strings from formal languages, evaluating memorization post-training, and generating plots are
560 attached as a supplemental material.
561

562 **REFERENCES**
563

564 Zeyuan Allen-Zhu and Yuanzhi Li. Physics of language models: Part 1, learning hierarchical lan-
565 guage structures. *ArXiv e-prints, abs/2305.13673*, May, 2023.
566

567 Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the
568 dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM*
569 *conference on fairness, accountability, and transparency*, pp. 610–623, 2021.

570 Satwik Bhattacharya, Kabir Ahuja, and Navin Goyal. On the ability and limitations of transformers
571 to recognize formal languages. *arXiv preprint arXiv:2009.11264*, 2020.
572

573 Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric
574 Hallahan, Mohammad Aftab Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al.
575 Pythia: A suite for analyzing large language models across training and scaling. In *International*
576 *Conference on Machine Learning*, pp. 2397–2430. PMLR, 2023.

577 Stella Biderman, Usvsn Prashanth, Lintang Sutawika, Hailey Schoelkopf, Quentin Anthony, Shivan-
578 shu Purohit, and Edward Raff. Emergent and predictable memorization in large language models.
579 *Advances in Neural Information Processing Systems*, 36, 2024.

580 Nadav Borenstein, Anej Svetec, Robin Chan, Josef Valvoda, Franz Nowak, Isabelle Augenstein,
581 Eleanor Chodroff, and Ryan Cotterell. What languages are easy to language-model? a perspective
582 from learning probabilistic regular languages. *arXiv preprint arXiv:2406.04289*, 2024.
583

584 Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Ka-
585 mar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general
586 intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.

587 Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. The secret sharer:
588 Evaluating and testing unintended memorization in neural networks. In *28th USENIX Security*
589 *Symposium (USENIX Security 19)*, pp. 267–284, 2019.
590

591 Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine
592 Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data
593 from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pp.
2633–2650, 2021.

594 Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and
595 Chiyuan Zhang. Quantifying memorization across neural language models. *arXiv preprint*
596 *arXiv:2202.07646*, 2022.

597

598 Rafael C. Carrasco. Accurate computation of the relative entropy between stochastic regular gram-
599 mars. *RAIRO-Theoretical Informatics and Applications*, 31(5):437–444, 1997.

600 Wentao Chen, Lizhe Zhang, Li Zhong, Letian Peng, Zilong Wang, and Jingbo Shang. Mem-
601 orize or generalize? evaluating llm code generation with evolved questions. *arXiv preprint*
602 *arXiv:2503.02296*, 2025.

603

604 Noam Chomsky. Three models for the description of language. *IRE Transactions on information*
605 *theory*, 2(3):113–124, 1956.

606 Michael Collins. Probabilistic context-free grammars (pcfgs). *Lecture Notes*, 2013.

607

608 Ryan Cotterell, Sabrina J Mielke, Jason Eisner, and Brian Roark. Are all languages equally hard to
609 language-model? *arXiv preprint arXiv:1806.03743*, 2018.

610 Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.

611

612 Soumi Das, Camila Kolling, Mohammad Aflah Khan, Mahsa Amani, Bishwamitra Ghosh, Qinyuan
613 Wu, Till Speicher, and Krishna P Gummadi. Revisiting privacy, utility, and efficiency trade-offs
614 when fine-tuning large language models. *arXiv preprint arXiv:2502.13313*, 2025.

615

616 Grégoire Delétang, Anian Ruoss, Jordi Grau-Moya, Tim Genewein, Li Kevin Wenliang, Elliot Catt,
617 Chris Cundy, Marcus Hutter, Shane Legg, Joel Veness, et al. Neural networks and the chomsky
618 hierarchy. *arXiv preprint arXiv:2207.02098*, 2022.

619

620 Yihong Dong, Xue Jiang, Huanyu Liu, Zhi Jin, Bin Gu, Mengfei Yang, and Ge Li. Generalization or
621 memorization: Data contamination and trustworthy evaluation for large language models. *arXiv*
622 *preprint arXiv:2402.15938*, 2024.

623

624 Sunny Duan, Mikail Khona, Abhiram Iyer, Rylan Schaeffer, and Ila R Fiete. Uncovering latent
625 memories: Assessing data leakage and memorization patterns in large language models. In *ICML*
626 *2024 Workshop on LLMs and Cognition*, 2024.

627

628 Sunny Duan, Mikail Khona, Abhiram Iyer, Rylan Schaeffer, and Ila R Fiete. Uncovering latent
629 memories in large language models. In *The Thirteenth International Conference on Learning*
630 *Representations*, 2025.

631

632 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
633 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
634 *arXiv preprint arXiv:2407.21783*, 2024.

635

636 Vitaly Feldman. Does learning require memorization? a short tale about a long tail. In *Proceedings*
637 *of the 52nd annual ACM SIGACT symposium on theory of computing*, pp. 954–959, 2020.

638

639 Vitaly Feldman and Chiyuan Zhang. What neural networks memorize and why: Discovering the
640 long tail via influence estimation. *Advances in Neural Information Processing Systems*, 33:2881–
641 2891, 2020.

642

643 Joshua Freeman, Chloe Rippe, Edoardo Debenedetti, and Maksym Andriushchenko. Exploring
644 memorization and copyright violation in frontier llms: A study of the new york times v. openai
645 2023 lawsuit. *arXiv preprint arXiv:2412.06370*, 2024.

646

647 Yi-Fu Fu, Yu-Chieh Tu, Tzu-Ling Cheng, Cheng-Yu Lin, Yi-Ting Yang, Heng-Yi Liu, Keng-Te Liao,
648 Da-Cheng Juan, and Shou-De Lin. Think or remember? detecting and directing llms towards
649 memorization or generalization. *arXiv preprint arXiv:2412.18497*, 2024.

650

651 Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason
652 Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text
653 for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.

648 Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord,
649 Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, et al. Olmo: Accelerating the
650 science of language models. *arXiv preprint arXiv:2402.00838*, 2024.

651

652 Michael Hahn. Theoretical limitations of self-attention in neural sequence models. *Transactions of
653 the Association for Computational Linguistics*, 8:156–171, 2020.

654

655 Adi Haviv, Ido Cohen, Jacob Gidron, Roei Schuster, Yoav Goldberg, and Mor Geva. Understanding
656 transformer memorization recall through idioms. *arXiv preprint arXiv:2210.03588*, 2022.

657

658 Peter Henderson, Xuechen Li, Dan Jurafsky, Tatsunori Hashimoto, Mark A Lemley, and Percy
659 Liang. Foundation models and fair use. *Journal of Machine Learning Research*, 24(400):1–79,
2023.

660

661 Danny Hernandez, Tom Brown, Tom Conerly, Nova DasSarma, Dawn Drain, Sheer El-Showk, Nel-
662 son Elhage, Zac Hatfield-Dodds, Tom Henighan, Tristan Hume, et al. Scaling laws and inter-
663 pretability of learning from repeated data. *arXiv preprint arXiv:2205.10487*, 2022.

664

665 Jie Huang, Hanyin Shao, and Kevin Chen-Chuan Chang. Are large pre-trained language models
leaking your personal information? *arXiv preprint arXiv:2205.12628*, 2022.

666

667 Daphne Ippolito, Florian Tramèr, Milad Nasr, Chiyuan Zhang, Matthew Jagielski, Katherine Lee,
668 Christopher A Choquette-Choo, and Nicholas Carlini. Preventing verbatim memorization in lan-
669 guage models gives a false sense of privacy. *arXiv preprint arXiv:2210.17546*, 2022.

670

671 Matthew Jagielski, Om Thakkar, Florian Tramer, Daphne Ippolito, Katherine Lee, Nicholas Carlini,
672 Eric Wallace, Shuang Song, Abhradeep Thakurta, Nicolas Papernot, et al. Measuring forgetting
of memorized training examples. *arXiv preprint arXiv:2207.00099*, 2022.

673

674 Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chap-
675 lot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier,
676 Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril,
677 Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL <https://arxiv.org/abs/2310.06825>.

678

679 Nikhil Kandpal, Eric Wallace, and Colin Raffel. Deduplicating training data mitigates privacy risks
680 in language models. In *International Conference on Machine Learning*, pp. 10697–10707. PMLR,
681 2022.

682

683 Siwon Kim, Sangdoo Yun, Hwaran Lee, Martin Gubri, Sungroh Yoon, and Seong Joon Oh. Propile:
684 Probing privacy leakage in large language models. *Advances in Neural Information Processing
685 Systems*, 36:20750–20762, 2023.

686

687 Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-
688 Burch, and Nicholas Carlini. Deduplicating training data makes language models better. *arXiv
689 preprint arXiv:2107.06499*, 2021.

690

691 Ken Ziyu Liu, Christopher A Choquette-Choo, Matthew Jagielski, Peter Kairouz, Sanmi Koyejo,
692 Percy Liang, and Nicolas Papernot. Language models may verbatim complete text they were not
693 explicitly trained on. *arXiv preprint arXiv:2503.17514*, 2025.

694

695 Anqi Mao, Mehryar Mohri, and Yutao Zhong. Cross-entropy loss functions: Theoretical analysis
and applications. In *International conference on Machine learning*, pp. 23803–23828. PMLR,
2023.

696

697 R Thomas McCoy, Paul Smolensky, Tal Linzen, Jianfeng Gao, and Asli Celikyilmaz. How much
698 do language models copy from their training data? evaluating linguistic novelty in text generation
699 using raven. *Transactions of the Association for Computational Linguistics*, 11:652–670, 2023.

700

701 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

702 Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Pater-
703 son, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément
704 Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng
705 Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski,
706 Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste
707 Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-
708 Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon,
709 Machel Reid, Maciej Mikuła, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain,
710 Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma
711 Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan
712 Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya,
713 Siamak Shakeri, Soham De, Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec,
714 Yu hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang,
715 Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahra-
716 mani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel,
717 Evan Senter, Alek Andreev, and Kathleen Kenealy. Gemma: Open models based on gemini
718 research and technology, 2024. URL <https://arxiv.org/abs/2403.08295>.

719 Fatemehsadat Mireshghallah, Archit Uniyal, Tianhao Wang, David K Evans, and Taylor Berg-
720 Kirkpatrick. An empirical analysis of memorization in fine-tuned autoregressive language models.
721 In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*,
722 pp. 1816–1826, 2022.

723 Felix B Mueller, Rebekka Görge, Anna K Bernzen, Janna C Pirk, and Maximilian Poretschkin.
724 Llms and memorization: On quality and specificity of copyright compliance. In *Proceedings of*
725 *the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, pp. 984–996, 2024.

726 Aneesh Pappu, Billy Porter, Ilia Shumailov, and Jamie Hayes. Measuring memorization in rlhf for
727 code completion. *arXiv preprint arXiv:2406.11715*, 2024.

728 Zhencan Peng, Zhizhi Wang, and Dong Deng. Near-duplicate sequence search at scale for large
729 language model memorization evaluation. *Proceedings of the ACM on Management of Data*, 1
730 (2):1–18, 2023.

731 Ali Satvaty, Suzan Verberne, and Fatih Turkmen. Undesirable memorization in large language
732 models: A survey. *arXiv preprint arXiv:2410.02650*, 2024.

733 Avi Schwarzschild, Zhili Feng, Pratyush Maini, Zachary C Lipton, and J Zico Kolter. Rethinking
734 llm memorization through the lens of adversarial compression. *arXiv preprint arXiv:2404.15146*,
735 2024.

736 Hui Shi, Sicun Gao, Yuandong Tian, Xinyun Chen, and Jishen Zhao. Learning bounded context-
737 free-grammar via lstm and the transformer: difference and the explanations. In *Proceedings of*
738 *the AAAI conference on artificial intelligence*, volume 36, pp. 8267–8276, 2022.

739 Till Speicher, Bishwamitra Ghosh, Mohammad Aflah Khan, Qinyuan Wu, Vedant Nanda, Soumi
740 Das, Krishna P Gummadi, and Evimaria Terzi. Rethinking memorization in llms: On learning by
741 rote vs. with understanding.

742 Marilyn Strathern. ‘improving ratings’: audit in the british university system. *European Review*,
743 5(3):305–321, July 1997. doi: 10.1002/(sici)1234-981x(199707)5:3<305::aid-euro184>3.0.co;
744 2-4. URL [https://doi.org/10.1002/\(sici\)1234-981x\(199707\)5:3<305::aid-euro184>3.0.co;2-4](https://doi.org/10.1002/(sici)1234-981x(199707)5:3<305::aid-euro184>3.0.co;2-4).

745 Kushal Tirumala, Aram Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. Memorization
746 without overfitting: Analyzing the training dynamics of large language models. *Advances in*
747 *Neural Information Processing Systems*, 35:38274–38290, 2022.

748 Gerrit van den Burg and Chris Williams. On memorization in probabilistic deep generative models.
749 *Advances in Neural Information Processing Systems*, 34:27916–27928, 2021.

756 Xinyi Wang, Antonis Antoniades, Yanai Elazar, Alfonso Amayuelas, Alon Albalak, Kexun Zhang,
757 and William Yang Wang. Generalization vs memorization: Tracing language models' capabilities
758 back to pretraining data. *arXiv preprint arXiv:2407.14985*, 2024.

759

760 Alexander Xiong, Xuandong Zhao, Aneesh Pappu, and Dawn Song. The landscape of memorization
761 in llms: Mechanisms, measurement, and mitigation. *arXiv preprint arXiv:2507.05578*, 2025.

762

763 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
764 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint
765 arXiv:2412.15115*, 2024.

766

767 Chiyuan Zhang, Daphne Ippolito, Katherine Lee, Matthew Jagielski, Florian Tramèr, and Nicholas
768 Carlini. Counterfactual memorization in neural language models. *ArXiv*, abs/2112.12938, 2021.
769 URL <https://api.semanticscholar.org/CorpusID:245502053>.

770 Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christo-
771 pher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer
772 language models. *arXiv preprint arXiv:2205.01068*, 2022.

773

774 Zhenhong Zhou, Jiuyang Xiang, Chaomeng Chen, and Sen Su. Quantifying and analyzing entity-
775 level memorization in large language models. In *Proceedings of the AAAI Conference on Artificial
776 Intelligence*, volume 38, pp. 19741–19749, 2024.

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 **A LIMITATIONS**
811

812 Our goal of the paper is to systematically study different measures of memorization, inspired from
813 two principled aspects: privacy concerns captured by recollection-based and counterfactual memo-
814 rization, and learning concerns captured by contextual memorization. While our priority has been
815 on conceptual clarity regarding the notion, operationalization, and information/computational chal-
816 lenges of each measure, we acknowledge the following limitations:

817 • **Multiple measures exist in the context of recollection-based memorization.** Many memo-
818 rization measures related to recollection exist today, as detailed in Section C. Fundamentally,
819 these measures vary on the strictness of the threshold, determining whether a string is memorized
820 or not. In the paper, we operationalize recollection-based memorization using a fixed genera-
821 tive performance threshold using cross-entropy loss, capturing the essence of many recollection-
822 based memorization notions. A future study could consider more nuanced operationalizations of
823 recollection-based memorization measures, although we believe the key results would not change,
824 i.e., the correlation of string-frequency and the order of recollection-based memorization would
825 remain the same.

826 • **Contextual (and counterfactual memorization) memorization measures are better, but at the**
827 **cost of additional information and computational resources.** Contextual memorization avoids
828 the pitfall of recollection-based memorization by disentangling memorization-based recall from
829 contextual learning based recall. However, contextual memorization requires more information
830 (e.g. access to training dataset) and computational resources (e.g. re-training the LLM). We
831 demonstrate a *heuristic to approximate contextual memorization* by mapping similarly *frequent*
832 training and test strings. Still, such approximation relies on knowing the probability distribution
833 of strings in the language, which is hard to get for any natural language. Moreover, beyond fre-
834 quency, there are other dimensions such as information context within strings that can be used for
835 determining how well strings are learned/memorized. We leave these possibilities for a tractable
836 and informed heuristic of contextual (and counterfactual) memorization as a future work.

837 • **Experiments are limited to formal languages.** We propose formal language learning to precisely
838 study nuanced implications of different measures of memorization, where we can effectively con-
839 trol of the entropy of the language, design grammatical rules to mimic natural languages, and
840 produce results that avoid data contamination issues in natural language datasets. Moreover, the
841 current study focuses on a single language family (probabilistic and hierarchical context-free lan-
842 guages), and we leave the study of other language families as future work.

843 **B LLM USAGE**
844

845 In this paper, we use LLMs for the following purposes:

846 1. **Improvement in Writing:** We check grammatical mistakes in writing, and make minor
847 para-phrasing to improve the quality and flow of the paper.
848 2. **Code Writing:** We leverage LLM-based code assistants like Windsurf to write code.

849 The usage of LLMs is by no means a significant contribution to the paper.

864
865
866 Table 2: Characteristics of memorization measures.
867
868
869
870

Memorization Measure	Motivation	Memorization Threshold	Ease of Operationalization	Strictness of Measure
Recollection	Disclosing private information	Manual	Easy	Variable
Counterfactual	Differential privacy	Adaptive	Hard	Medium
Contextual (ours)	Local over-fitting	Adaptive	Hard	High

871
872 C EXTENDED RELATED WORK
873
874

875 **Measures of Memorization.** Memorization in LLMs is an active area of research, studied from
876 the perspective of privacy and security risks (Carlini et al., 2021; Huang et al., 2022; Kim et al.,
877 2023; Jagielski et al., 2022), unintended form of learning due to local over-fitting (van den Burg &
878 Williams, 2021), and copyright concerns related to verbatim reproduction (Bender et al., 2021; Hen-
879 derson et al., 2023; Mueller et al., 2024; Freeman et al., 2024), etc. With the goal of capturing unde-
880 sirable effects of memorization, i.e., privacy risks, multiple measures of memorization are proposed.
881 Among them, majority belong to the category of recollection-based memorization (Schwarzschild
882 et al., 2024; Biderman et al., 2024), in the form of perfect memorization (Kandpal et al., 2022),
883 verbatim or exact memorization (Carlini et al., 2021; 2019; Tirumala et al., 2022; Mireshghallah
884 et al., 2022), approximate memorization (Ippolito et al., 2022; Peng et al., 2023; Duan et al., 2024),
885 entity memorization (Zhou et al., 2024), etc. For an extended taxonomy of memorization measures,
886 we refer to a recent papers (Xiong et al., 2025; Satvaty et al., 2024). For example, Tirumala et al.
887 (2022) considered per-token training accuracy as the proxy of memorization: given a training string
888 as a prompt, an LLM memorizes it if it recollects the next token in the string correctly. Carlini
889 et al. (2022) proposed a relatively stringent measure by imposing an exact recollection of the next
890 50 tokens. Therefore, a critical design choice an experimenter makes is to set the threshold on rec-
891 ollection to declare a string as memorized. The choice has consequences on the interpretation of
892 memorization, as we demonstrate in Secton 3, and Section 4.
893

894 In a related line of work, Zhang et al. (2021) defined counterfactual memorization as the change in a
895 model’s generative performance when a string is included in training versus excluded (Pappu et al.,
896 2024; Feldman & Zhang, 2020). This approach specifically highlights rare and less frequent strings,
897 which tend to cause larger performance shifts and are often missed by recollection-based memo-
898 rization measures. By introducing contextual memorization, we argue that all strings, regardless of
899 frequency, can be recollected to some extent based on their context (Haviv et al., 2022; Wang et al.,
900 2024; Fu et al., 2024; Chen et al., 2025; Speicher et al.; Dong et al., 2024; McCoy et al., 2023).
901 We define memorization as occurring only when a string’s training-time recollection exceeds its
902 optimal contextual recollection threshold, making contextual memorization a stricter criterion than
903 counterfactual memorization. Despite the abundance of memorization measures, their potentially
904 conflicting implications remain underexplored – we aim to address this research gap in this study.
905

906 **Contextual vs. Counterfactual Memorization.** We do not make a subjective opinion of which
907 measure between contextual and counterfactual memorization is better. In fact, we make a note
908 that both are better than recollection-based memorization (Section 3), which fails to differentiate
909 between memorization-based recall and context-based recall.
910

911 At a more granular level, contextual memorization and counterfactual memorization are proposed
912 to detect learning risks and privacy risks, respectively. The privacy perspective may tolerate a *false*
913 *positive* of classifying a string as memorized, but the learning perspective may impose the necessary
914 condition for memorization to avoid any false positive. As stated in Lemma 1, contextual memo-
915 rization being stricter than counterfactual memorization results in lower false positive, which may be
916 appreciated from the learning perspective but not from the privacy perspective, making it subjective.
917

918 **Memorization and Learning.** Several prior works have investigated the interplay between memo-
919 rization and learning (van den Burg & Williams, 2021). Notably important is the seminal work
920 of Feldman (2020) studied in the domain of classification problems, where memorization of clas-
921 sification labels, and in some cases interpolation, are necessary for achieving nearly optimal gen-
922

eralization. Herein, we focus on generative models, such as LLMs, and empirically support their theoretical analysis. Particularly, we show that some strings are inevitably memorized before the LLM reaches optimal learning, regardless of the measure of memorization applied (Section 4). Contemporary to our work, Liu et al. (2025) find that verbatim generation is indeed observed in LLMs even though the prefix for the generation is absent in the training data, and call for a critical rethinking of what defines *membership* in the training data. Our proposed contextual memorization explains their finding, where a non-member string can be generated by the LLM having contextual understanding of the language. Therefore, contextual memorization makes the case that generation or recollection is caused by both rote memorization (i.e., membership) and contextual learning (i.e., non-membership), and provides a precise operationalization to separate them (Section 3).

Memorization and Learning Dynamics. We study memorization of individual training strings by analyzing their learning dynamics. Related to ours, Biderman et al. (2024) consider the problem of *forecasting* whether a large model memorizes a specific string based on its memorization by a smaller model. They apply a variant of recollection-based memorization for detection, but do not find any conclusive signal for forecasting; they even question the appropriateness of recollection-based memorization. Our paper differs with Biderman et al. (2024) in two distinct ways: (i) Our analysis is focused on a single model, where we ask whether the model memorizes all strings at the same time, and how their memorization is influenced by contextual learning of the language. (ii) We carefully revisit different memorization measures, and demonstrate the limitations of recollection-based memorization compared to contextual/counterfactual memorization. Recently, Duan et al. (2025) study how memorization changes throughout training and focus on strings which occurred only once throughout training. They find that the model interestingly memorizes the rarely occurred strings instead of forgetting them – recall is still possible by the model even though the string is seen long ago. Their operationalization of memorization is based on Levenshtein distance, which is also a variant of recollection-based memorization. Our paper thus questions their findings, where a rarely occurred string can still be recalled, based on influences of training over other strings leading to contextual learning. A future work might be to study their setup through the lens of contextual memorization.

D THEORETICAL ANALYSIS OF MEMORIZATION MEASURES

Lemma 1. Contextual memorization is stricter than counterfactual memorization. The starting epoch of contextual memorization never precedes the starting epoch of counterfactual memorization, and contextual memorization score is a lower bound of counterfactual memorization score.

Proof. We prove by considering loss as the metric of recollection. We assume that at any epoch, the training loss of a string is not higher than the counterfactual loss of the same string when excluding the string from training, which is a feasible assumption in practice.

For a string s , let the optimal contextual loss be $\min_{e^*} \text{loss}(M_{e^*}(D'), s)$, which is the lowest counterfactual loss in all epochs.

Contextual memorization starts at an epoch e_s^{ctx} when $\text{loss}(M_{e_s^{\text{ctx}}}(D), s) < \min_{e^*} \text{loss}(M_{e^*}(D'), s)$, i.e., the training loss of s is lower than the optimal contextual loss of the string. For an epoch $e < e_s^{\text{ctx}}$ earlier than the start of contextual memorization, $\text{loss}(M_e(D), s) \geq \min_{e^*} \text{loss}(M_{e^*}(D'), s)$.

Counterfactual memorization starts at an epoch e_s^{cf} when $\text{loss}(M_{e_s^{\text{cf}}}(D), s) < \text{loss}(M_{e_s^{\text{cf}}}(D'), s)$, i.e., the training loss of s is lower than the counterfactual loss at the same epoch. For an epoch $e < e_s^{\text{cf}}$ earlier than the start of counterfactual memorization, training loss of s is equal to the counterfactual loss, $\text{loss}(M_e(D), s) = \text{loss}(M_e(D'), s)$. Because, $\text{loss}(M_e(D), s) \leq \text{loss}(M_e(D'), s)$ for any training epoch e' , according to our assumption.

Let contextual memorization start earlier than counterfactual memorization, i.e., $e_s^{\text{ctx}} = e_s^{\text{cf}} - 1$.

972
 973
 974 $\text{loss}(M_{e_s^{\text{cf}}-1}(D), s) < \min_{e^*} \text{loss}(M_{e^*}(D'), s)$
 975 Since, $\min_{e^*} \text{loss}(M_{e^*}(D'), s) \leq \text{loss}(M_{e_s^{\text{cf}}-1}(D'), s)$
 976
 977 $\Rightarrow \text{loss}(M_{e'-1}, s) < \text{loss}(M_{e_s^{\text{cf}}-1}(D'), s)$
 978

979 But $\text{loss}(M_{e_s^{\text{cf}}-1}(D), s) = \text{loss}(M_{e_s^{\text{cf}}-1}(D'), s)$, which is a contradiction. Therefore, contextual
 980 memorization cannot start earlier than counterfactual memorization.

981 On the other hand, contextual memorization can start later or in the same epoch as counterfactual
 982 memorization, since for an epoch $e \geq e_s^{\text{cf}}$,
 983

$$\underbrace{\text{loss}(M_e(D), s) \geq \min_{e^*} \text{loss}(M_{e^*}(D'), s)}_{\text{contextual memorization does not start}} \text{ and } \underbrace{\text{loss}(M_e(D), s) < \text{loss}(M_e(D'), s)}_{\text{counterfactual memorization starts}}$$

987 Furthermore, the counterfactual memorization score is no less than the contextual memoriza-
 988 tion score, since at any epoch $e \geq \max(e_s^{\text{cf}}, e_s^{\text{ctx}})$, i.e., after both memorization starts,
 989 $\min_{e^*} \text{loss}(M_{e^*}(D'), s) \leq \text{loss}(M_e(D'), s)$.

$$\underbrace{\frac{\min_{e^*} \text{loss}(M_{e^*}(D'), s) - \text{loss}(M_e(D), s)}{\min_{e^*} \text{loss}(M_{e^*}(D'), s)}}_{\text{contextual memorization score}} \leq \underbrace{\frac{\text{loss}(M_e(D'), s) - \text{loss}(M_e(D), s)}{\text{loss}(M_e(D'), s)}}_{\text{counterfactual memorization score}}$$

994 Therefore, counterfactual memorization is likely to overestimate memorization than contextual
 995 memorization, while reporting memorization at an earlier epoch than contextual memorization.
 996

□

1000 E EXPERIMENTAL SETUP

1001 Each training (specifically, fine-tuning) is performed for 50 epochs with a batch size of 8 and a linear
 1002 learning rate scheduler with a warm-up ratio of 0.05. We fix the learning rate for Qwen, Gemma,
 1003 and Llama-3 families as 5×10^{-5} , Mistral, Opt, and Llama-2 families as 5×10^{-6} , and Pythia family
 1004 as 10^{-5} . We consider training dataset sizes $\{16, 64, 256, 1024\}$ and evaluate on 1024 test strings. In
 1005 each training, we find the epoch of best learning according to the lowest cross-entropy loss on the
 1006 test strings and report respective weighted memorization by different measures. All experiments are
 1007 conducted in compute clusters with Python as the programming language (version 3.10), where we
 1008 use 8x Nvidia H100 94GB NVL GPUs and 2x AMD EPYC 9554 CPU @ 3.1 GHz, 2x64 cores, and
 1009 24x 96GB RAM.

1010 Below, we provide details of the formal languages used in our experiments, along with their formal
 1011 definitions. Intuitively, we carefully design languages to show the robustness of our results across
 1012 changing the entropy of the language and token types of the language.

1013
 1014 **Formal Languages and Grammars.** In each fine-tuning, we provide the LLM with strings sam-
 1015 pled from a probabilistic formal language, with the learning task of generating unseen strings from
 1016 the same language via syntactic pattern recognition. Underneath, a probabilistic formal language is
 1017 represented by a *probabilistic formal grammars*, or simply *grammars* (Collins, 2013). Specifically,
 1018 a grammar consists of two sets of symbols called the *non-terminals* and *terminals*, a set of rules to
 1019 rewrite strings over these symbols that contain at least one nonterminal – also called the *produc-*
 1020 *tion rules*, and a probability distribution over the production rules. Formally, a probabilistic formal
 1021 grammar, is defined as a quintuple.

$$1022 \quad G = (\mathbf{N}, \mathbf{T}, \mathbf{R}, \mathbf{S}, \mathbf{P})$$

1023 where \mathbf{N} is the set of non-terminals, \mathbf{T} is the set of terminals (equivalently, tokens), \mathbf{R} is the set of
 1024 production rules, $\mathbf{S} \in \mathbf{N}$ is the start non-terminal, and \mathbf{P} is the set of probabilities on production
 1025 rules.

```

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039      S → A16 [1]           S → A16 [1]
1040      A16 → A15 A14 A13 [0.50]   A16 → A15 A14 A13 [0.95]
1041      A16 → A13 A15 A14 [0.50]   A16 → A13 A15 A14 [0.05]
1042      A13 → A11 A12 [0.50]        A13 → A11 A12 [0.95]
1043      A13 → A12 A11 [0.50]        A13 → A12 A11 [0.05]
1044      A14 → A11 A10 A12 [0.50]   A14 → A11 A10 A12 [0.95]
1045      A14 → A10 A11 A12 [0.50]   A14 → A10 A11 A12 [0.05]
1046      A15 → A12 A11 A10 [0.50]   A15 → A12 A11 A10 [0.95]
1047      A15 → A11 A12 A10 [0.50]   A15 → A11 A12 A10 [0.05]
1048
1049
1050      A10 → A7 A9 A8 [0.50]        A10 → A7 A9 A8 [0.95]
1051      A10 → A9 A8 A7 [0.50]        A10 → A9 A8 A7 [0.05]
1052      A11 → A8 A7 A9 [0.50]        A11 → A8 A7 A9 [0.95]
1053      A11 → A7 A8 A9 [0.50]        A11 → A7 A8 A9 [0.05]
1054      A12 → A8 A9 A7 [0.50]        A12 → A8 A9 A7 [0.95]
1055      A12 → A9 A7 A8 [0.50]        A12 → A9 A7 A8 [0.05]
1056
1057      A7 → 3 1 2 [0.50]          A7 → 3 1 2 [0.95]
1058      A7 → 1 2 3 [0.50]          A7 → 1 2 3 [0.05]
1059      A8 → 6 5 4 [0.50]          A8 → 6 5 4 [0.95]
1060      A8 → 6 4 5 [0.50]          A8 → 6 4 5 [0.05]
1061      A9 → 9 8 7 [0.50]          A9 → 9 8 7 [0.95]
1062      A9 → 8 7 9 [0.50]          A9 → 8 7 9 [0.05]
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

```

Figure 7: Production rules of G_1 (left) and G_2 (right). Compared to G_1 , the grammar G_2 generates more skewed distribution (or lower entropy) strings, since one out of two production rules for each non-terminal is selected with higher probability.

```

1080
1081
1082
1083      S → S5 [1]           S → S5 [1]
1084      S5 → B4 C11 E4 T11 [0.25]   S5 → B4 C11 E4 T11 [0.25]
1085      S5 → B4 C12 E4 T12 [0.25]   S5 → B4 C12 E4 T12 [0.25]
1086      S5 → B4 C13 E4 T13 [0.25]   S5 → B4 C13 E4 T13 [0.25]
1087      S5 → B4 C14 E4 T14 [0.25]   S5 → B4 C14 E4 T14 [0.25]
1088
1089      B4 → B3 [0.3333]           B4 → B3 [0.3333]
1090      B4 → B3 B3 B3 [0.3333]       B4 → B3 B3 B3 [0.3333]
1091      B4 → B3 B3 [0.3333]         B4 → B3 B3 [0.3333]
1092      B3 → B2 [0.3333]           B3 → B2 [0.3333]
1093      B3 → B2 [0.3333]           B3 → B2 [0.3333]
1094      B3 → B2 B2 [0.3333]         B3 → B2 B2 [0.3333]
1095      B2 → B1 [0.3333]           B2 → B1 [0.3333]
1096      B2 → B1 [0.3333]           B2 → B1 [0.3333]
1097
1098      B2 → B1 B1 B1 [0.3333]       B2 → B1 B1 B1 [0.3333]
1099      B1 → 2 9 3 [0.3333]         B1 → 2 9 3 [0.95]
1100      B1 → 9 6 1 [0.3333]         B1 → 9 6 1 [0.025]
1101      B1 → 1 8 6 [0.3333]         B1 → 1 8 6 [0.025]
1102
1103      E4 → E3 [0.3333]           E4 → E3 [0.3333]
1104      E4 → E3 E3 [0.3333]         E4 → E3 E3 [0.3333]
1105      E4 → E3 E3 E3 [0.3333]       E4 → E3 E3 E3 [0.3333]
1106
1107      E3 → E2 [0.3333]           E3 → E2 [0.3333]
1108      E3 → E2 E2 [0.3333]         E3 → E2 E2 [0.3333]
1109      E3 → E2 [0.3333]           E3 → E2 [0.3333]
1110      E2 → E1 E1 [0.3333]         E2 → E1 E1 [0.3333]
1111      E2 → E1 [0.3333]           E2 → E1 [0.3333]
1112      E2 → E1 E1 E1 [0.3333]       E2 → E1 E1 E1 [0.3333]
1113      E1 → 5 6 5 9 [0.3333]       E1 → 5 6 5 9 [0.95]
1114      E1 → 1 8 6 6 [0.3333]       E1 → 1 8 6 6 [0.025]
1115      E1 → 1 5 1 5 [0.3333]       E1 → 1 5 1 5 [0.025]
1116
1117      T11 → 1 [1]             T11 → 1 [1]
1118      T12 → 2 [1]             T12 → 2 [1]
1119      T13 → 3 [1]             T13 → 3 [1]
1120      T14 → 4 [1]             T14 → 4 [1]
1121      C11 → 5 [1]             C11 → 5 [1]
1122      C12 → 6 [1]             C12 → 6 [1]
1123      C13 → 7 [1]             C13 → 7 [1]
1124      C14 → 8 [1]             C14 → 8 [1]
1125      C15 → 9 [1]             C15 → 9 [1]
1126
1127
1128

```

Figure 8: Production rules of G_3 (left) and G_4 (right). Compared to G_3 , the grammar G_4 generates more skewed distribution (or lower entropy) of strings, since one out of three production rules of non-terminal *B1* and *E1* is selected with higher probability.

1132

1133

```

1134
1135
1136
1137
1138  $S \rightarrow S5 [1]$ 
1139  $S5 \rightarrow B4 C1_1 E4 T1_1 [0.25]$ 
1140  $S5 \rightarrow B4 C1_2 E4 T1_2 [0.25]$ 
1141  $S5 \rightarrow B4 C1_3 E4 T1_3 [0.25]$ 
1142  $S5 \rightarrow B4 C1_4 E4 T1_4 [0.25]$ 
1143  $B4 \rightarrow B3 [0.3333]$ 
1144  $B4 \rightarrow B3 B3 B3 [0.3333]$ 
1145  $B4 \rightarrow B3 B3 [0.3333]$ 
1146
1147  $S \rightarrow A16 [1]$ 
1148  $A16 \rightarrow A15 A13 [0.50]$ 
1149  $A16 \rightarrow A13 A15 A14 [0.50]$ 
1150  $A13 \rightarrow A11 A12 [0.50]$ 
1151  $A13 \rightarrow A12 A11 [0.50]$ 
1152  $A14 \rightarrow A11 A10 A12 [0.50]$ 
1153  $A14 \rightarrow A10 A11 A12 [0.50]$ 
1154  $A15 \rightarrow A12 A11 A10 [0.50]$ 
1155  $A15 \rightarrow A11 A12 A10 [0.50]$ 
1156  $A10 \rightarrow A7 A9 A8 [0.50]$ 
1157  $A10 \rightarrow A9 A8 A7 [0.50]$ 
1158  $A11 \rightarrow A8 A7 A9 [0.50]$ 
1159  $A11 \rightarrow A7 A8 A9 [0.50]$ 
1160  $A12 \rightarrow A8 A9 A7 [0.50]$ 
1161  $A12 \rightarrow A9 A7 A8 [0.50]$ 
1162  $A7 \rightarrow 3 1 [0.50]$ 
1163  $A7 \rightarrow 1 2 3 [0.50]$ 
1164  $A8 \rightarrow 6 5 [0.50]$ 
1165  $A8 \rightarrow 6 4 5 [0.50]$ 
1166  $A9 \rightarrow 9 8 7 [0.50]$ 
1167  $A9 \rightarrow 8 7 [0.50]$ 
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184  $T1_1 \rightarrow 1 [1]$ 
1185  $T1_2 \rightarrow 2 [1]$ 
1186  $T1_3 \rightarrow 3 [1]$ 
1187  $T1_4 \rightarrow 4 [1]$ 
1188  $C1_1 \rightarrow 5 [1]$ 
1189  $C1_2 \rightarrow 6 [1]$ 
1190  $C1_3 \rightarrow 7 [1]$ 
1191  $C1_4 \rightarrow 8 [1]$ 
1192  $C1_5 \rightarrow 9 [1]$ 

```

Figure 9: Production rules of G_5 (left) and G_6 (right). These grammars are adapted from G_1 and G_3 respectively, by allowing non-uniform lengths of tokens in the lowest level production rules.

```

1188
1189
1190
1191
1192  $S \rightarrow S5 [1]$ 
1193  $S5 \rightarrow B4 C1_1 E4 T1_1 [0.25]$ 
1194  $S5 \rightarrow B4 C1_2 E4 T1_2 [0.25]$ 
1195  $S5 \rightarrow B4 C1_3 E4 T1_3 [0.25]$ 
1196  $S5 \rightarrow B4 C1_4 E4 T1_4 [0.25]$ 
1197
1198
1199
1200
1201  $S \rightarrow A16 [1]$ 
1202  $A16 \rightarrow A15 A13 [0.50]$ 
1203  $A16 \rightarrow A13 A15 A14 [0.50]$ 
1204  $A13 \rightarrow A11 A12 [0.50]$ 
1205  $A13 \rightarrow A12 A11 [0.50]$ 
1206  $A14 \rightarrow A11 A10 A12 [0.50]$ 
1207  $A14 \rightarrow A10 A11 A12 [0.50]$ 
1208  $A15 \rightarrow A12 A11 A10 [0.50]$ 
1209  $A15 \rightarrow A11 A12 A10 [0.50]$ 
1210
1211
1212  $A10 \rightarrow A7 A9 A8 [0.50]$ 
1213  $A10 \rightarrow A9 A8 A7 [0.50]$ 
1214  $A11 \rightarrow A8 A7 A9 [0.50]$ 
1215  $A11 \rightarrow A7 A8 A9 [0.50]$ 
1216  $A12 \rightarrow A8 A9 A7 [0.50]$ 
1217  $A12 \rightarrow A9 A7 A8 [0.50]$ 
1218
1219  $A7 \rightarrow c a [0.50]$ 
1220  $A7 \rightarrow a b c [0.50]$ 
1221  $A8 \rightarrow f e [0.50]$ 
1222  $A8 \rightarrow f d e [0.50]$ 
1223  $A9 \rightarrow i h g [0.50]$ 
1224  $A9 \rightarrow h g [0.50]$ 
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237  $T1_1 \rightarrow a [1]$ 
1238  $T1_2 \rightarrow b [1]$ 
1239  $T1_3 \rightarrow c [1]$ 
1240  $T1_4 \rightarrow d [1]$ 
1241  $C1_1 \rightarrow e [1]$ 
 $C1_2 \rightarrow f [1]$ 
 $C1_3 \rightarrow g [1]$ 
 $C1_4 \rightarrow h [1]$ 
 $C1_5 \rightarrow i [1]$ 

```

Figure 10: Production rules of G_7 (left) and G_8 (right). These grammars are adapted from G_5 and G_6 respectively, by replacing numerical tokens with Latin character tokens.

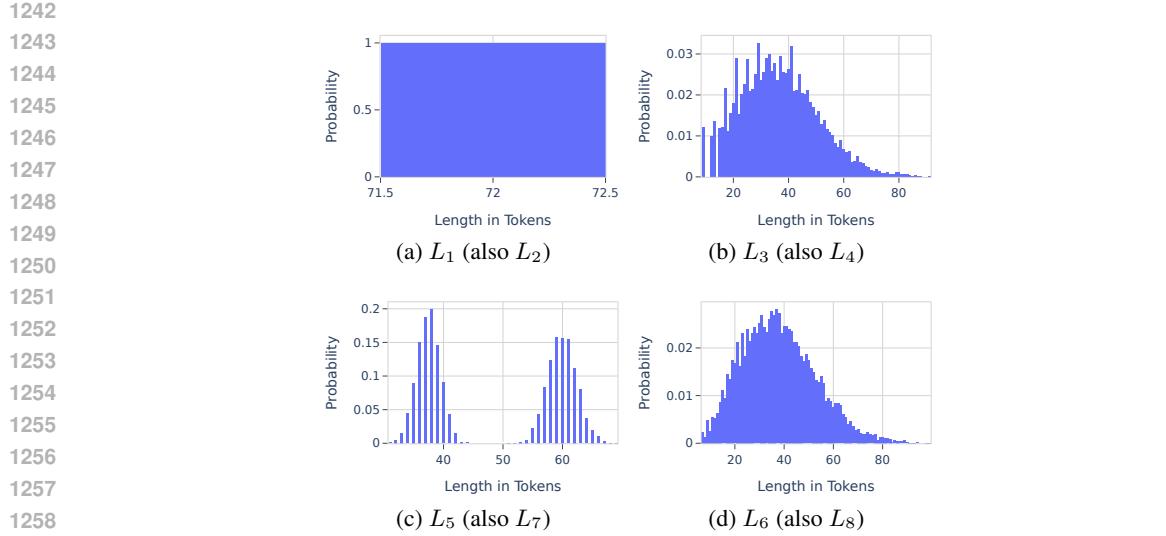


Figure 11: Length distribution of considered probabilistic languages, based on 10000 sampled strings per language.

Formal languages are divided into well-known classes based on the *complexity* of the language membership problem, i.e., the *complexity* of the grammars needed to generate them (Chomsky, 1956). In this paper, we use one class of grammars, namely, hierarchical probabilistic context-free grammars (HPCFGs) (Allen-Zhu & Li, 2023). Specifically, our experiments are based on teaching LLMs languages represented by HPCFGs. We use HPCFGs because they are simple syntactically and can represent languages that are structurally similar to natural languages (Allen-Zhu & Li, 2023; Shi et al., 2022).

Description of Grammars and Identified Languages. In our experiments, we consider two generic structure for the considered grammars, one adapted from Allen-Zhu & Li (2023), namely G_1, G_2, G_5, G_7 , and another is proposed by us, namely G_3, G_4, G_6, G_8 .

In the first generic structure, each grammar has $N = \{S, A7, A8, \dots, A16\}$ and $T = \{1, 2, 3, \dots, 9\}$. The grammar has four levels of hierarchy: the non-terminals from top to bottom levels are $\{A16\}$, $\{A13, A14, A15\}$, $\{A10, A11, A12\}$, and $\{A7, A8, A9\}$, followed by terminals $\{1, 2, 3, \dots, 9\}$. Each non-terminal (except the start non-terminal) has two expansion rules, consisting of non-terminals from the immediate lower level. Further, the expansion rules are probabilistic, where the sum of probabilities of all expansion rules from a given non-terminal is 1.

The second generic structure is inspired by bridging two HPCFGs together, and simulating a long range dependencies within the generated strings. Specifically, the sub-grammar at $B4$ and the sub-grammar at $E4$ are connected by non-terminal $C1_i$; and $E4$ ends with $T1_j$. Long range dependencies are communicated through $C1_i$ and $T1_j$, by enforcing $i = j$ at each expansion of $S5$.

In all cases, G_i produces a probabilistic context free language L_i . Figure 11 denotes the length distribution of different languages, and Figure 12 demonstrates how hierarchical non-terminals are applied in different positions in the representative strings.

Sampling Strings from a Formal Language. Given a language L generated by a HPCFG, we first need to obtain *training* samples, i.e., set of i.i.d. samples of strings from L . To *sample a string from the language*, we start from a special string in the grammar containing a single, distinguished nonterminal called the “start” or “root” symbol, and apply the production rules to rewrite the string repeatedly. If several rules can be used to rewrite the string at any stage, we sample one such rule from the probability distribution over the rules and apply it. We stop when we obtain a string containing terminal tokens only. This string is a sample drawn from the language. We can repeat this process to draw any number of i.i.d. samples from the language.

1296

1297

1298

1299

1300

8 7 9 1 2 3 6 5 4 6 4 5 1 2 3 8 7 9 6 5 4 8 7 9 1 2 3 6 5 4 3 1 2 8 7 9 8 7 6 4 5 3 1 2 3 1 2 8 7 9 6 5 4 6 4 5 1 2 3 9 8 7 8 7 9 1 2 3 6 5 4

A9 . ; A7 . ; A8 . ; A8 . ; A7 . ; A9 . ; A8 . ; A7 . ; A8 . ; A9 . ; A8 . ; A7 . ; A9 . ; A8 . ; A7 . ; A9 . ; A7 . ; A8 . ;

A12 ; A11 ; A12 ; A11 ; A10 ; A11 ; A12 ;

A13 ; A15 ; A14 ; A13 ; A12 ;

A16 ;

(a) L_1

1303

1304

1305

1306

6 5 4 9 8 7 3 1 2 6 5 4 6 5 1 2 3 9 8 7 3 1 2 9 8 7 6 5 4 6 5 4 3 1 2 9 8 7 3 1 2 9 8 7 6 5 4 6 5 4 9 8 7 3 1 2 6 4 5 3 1 2 9 8 7 9 8 7 3 1 2 6 5 4

A8 . ; A9 . ; A7 . ; A8 . ; A7 . ; A9 . ; A8 . ; A7 . ; A9 . ; A8 . ; A9 . ; A8 . ; A7 . ; A9 . ; A7 . ; A8 . ;

A12 ; A11 ; A12 ; A11 ; A10 ; A11 ; A12 ;

A15 ; A14 ; A13 ; A12 ;

A16 ;

(b) L_2

1307

1308

1309

1310

1311

1 8 6 2 9 3 9 6 1 9 6 1 1 8 6 2 9 3 9 6 1 9 6 1 7 1 5 1 5 1 5 1 5 3

B1 . . : Cl₃ E1 . . . : E1 . . . : Tl₃

B2 : E2 : E2 : ;

B3 : E3 : E3 : ;

B4 : E4 : E4 : ;

S5 : E5 : E5 : ;

(c) L_3

1315

1316

1317

1318

1319

1320

2 9 3 2 9 3 2 9 3 2 9 3 2 9 3 2 9 3 7 5 6 5 6 5 9 5 6 5 9 3

B1 . . : Cl₃ E1 . . . : E1 . . . : Tl₃

B2 : E2 : E2 : ;

B3 : E3 : E3 : ;

B4 : E4 : E4 : ;

S5 : E5 : E5 : ;

(d) L_4

1321

1322

1323

1324

1325

1326

1327

1328

1 8 6 2 2 9 3 9 6 1 9 6 1 1 8 6 2 2 9 3 9 6 1 7 1 5 1 5 1 5 9 1 5 1 5 5 9 3

B1 . . : Cl₃ E1 . . . : E1 . . . : Tl₃

B2 : E2 : E2 : ;

B3 : E3 : E3 : ;

B4 : E4 : E4 : ;

S5 : E5 : E5 : ;

(f) L_6

1332

1333

1334

1335

1336

1337

1338

1339

h g a b c f e f d e a b c h g f e h g a b c f e c a h g i h g f d e c a c a h g f e f d e a b c i h g h g a b c f e

A9 . ; A7 . ; A8 . ; A8 . ; A7 . ; A9 . ; A9 . ; A8 . ; A7 . ; A9 . ; A8 . ; A7 . ; A9 . ; A8 . ; A7 . ; A8 . ;

A12 ; A11 ; A12 ; A11 ; A10 ; A11 ; A12 ;

A13 ; A15 ; A14 ; A13 ; A12 ;

A16 ;

(g) L_7

1340

1341

1342

1343

1344

a h f b b i c i f a i f a a h f b b i c i f a g a e a e e i a e e a e e i c

B1 . . : Cl₃ E1 . . . : E1 . . . : Tl₃

B2 : E2 : E2 : ;

B3 : E3 : E3 : ;

B4 : E4 : E4 : ;

S5 : E5 : E5 : ;

(h) L_8

1345

1346

1347

1348

1349

Figure 12: Representative strings from different languages, annotated with non-terminals applied in different positions by the respective hierarchical grammar.

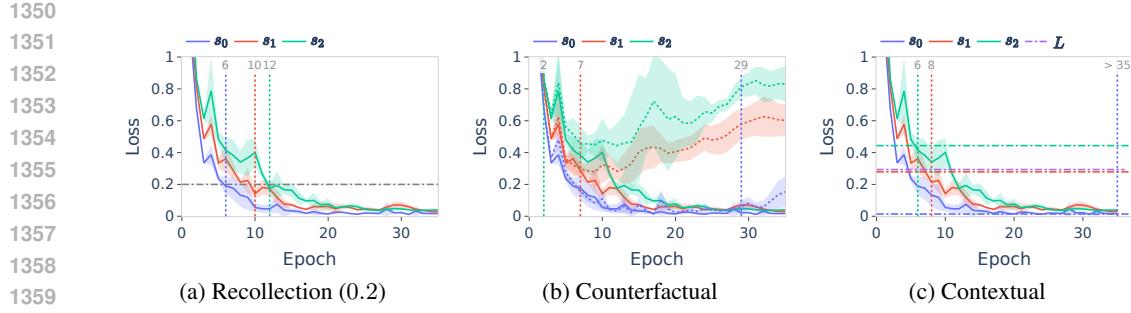


Figure 13: Start of memorization of selected strings in Language L_2 .

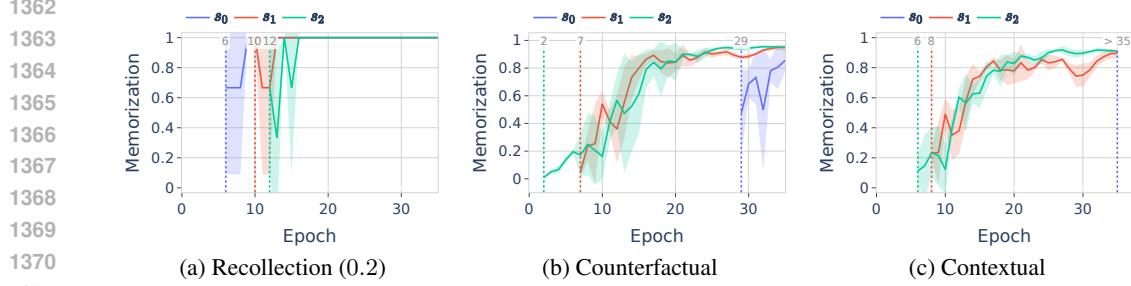


Figure 14: Memorization score of strings in language L_2 , respective to Figure 2. In different strings, memorization score usually increases with epochs, with contextual memorization providing a lower bound of counterfactual memorization.

F ADDITIONAL EXPERIMENTAL RESULTS

Memorization Scores of Individual Strings. In Figure 14, we demonstrate the memorization scores of strings, corresponding to Figure 2, across multiple memorization measures. In all measures, the memorization score usually increases with epochs, and there is no substantial difference among strings of varying frequency – different measures agree on the memorization score. Finally, as we theoretically demonstrate, contextual memorization score provides a lower bound of counterfactual memorization score.

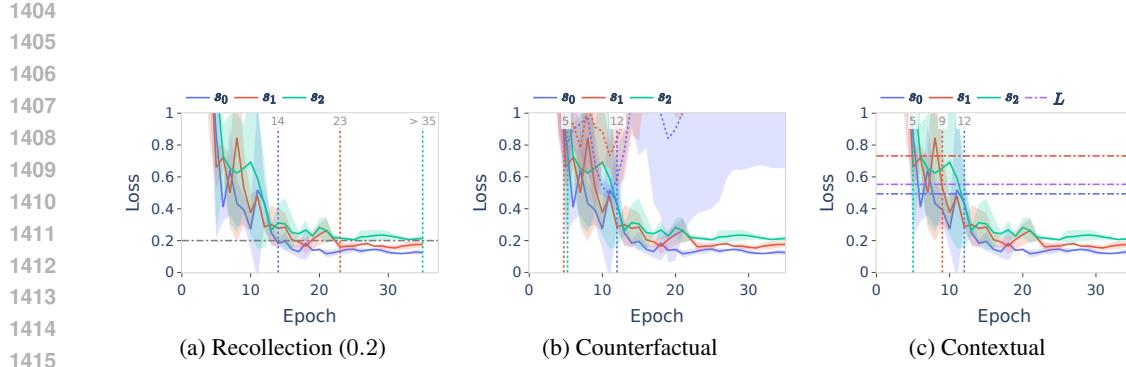


Figure 15: Start of memorization of selected strings in language L_4 (specifically, a modified version of L_4 as explained below). The observation is consistent with language L_2 , as shown in figure 2, where frequency of strings correlates with the start of recollection-based memorization. Similarly, frequency often inversely correlates with counterfactual and contextual memorization, with an exception that both s_1 and s_2 are memorized at the same epoch in the counterfactual memorization. Thus, regardless of whether correlation or inverse correlation exists *strongly* between string frequency and the order of memorization, a more consistent observation is that memorization measures disagree with each other when applied to the same training dynamic on identical strings.

In this experiment, to better differentiate the strings s_0, s_1, s_2 based on frequency, we modify L_4 to be even more skewed. We apply high probability to one random production rule in each non-terminal in all levels, beyond the lowest level non-terminals in L_4 , as shown in Figure 8.

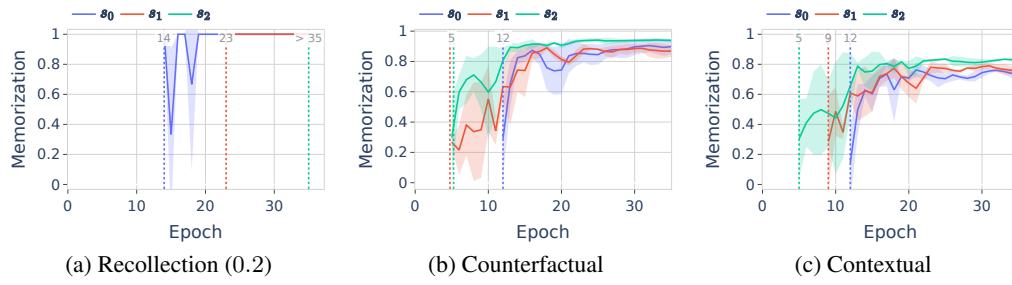


Figure 16: Memorization score of strings in language L_4 .

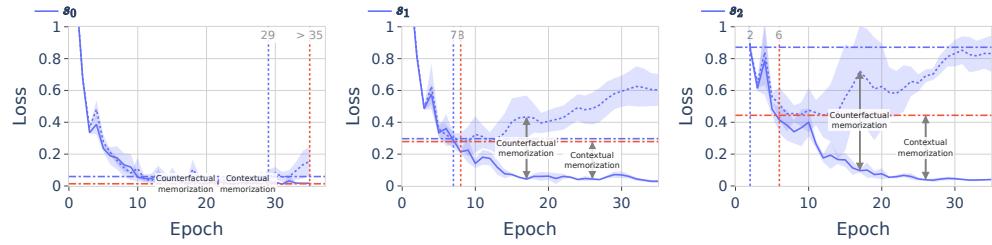


Figure 17: Contextual memorization is a stricter measure than counterfactual memorization. Red horizontal dash-dot line is the optimal contextual loss. Contextual memorization starts at the same or in a later epoch (red vertical dot line) than the start of counterfactual memorization (blue vertical dot line). The contextual memorization score (gray arrow) is a lower bound of counterfactual memorization score, intuitively by comparing the arrow-length.

1458

1459

1460

1461

1462

1463

1464

1465

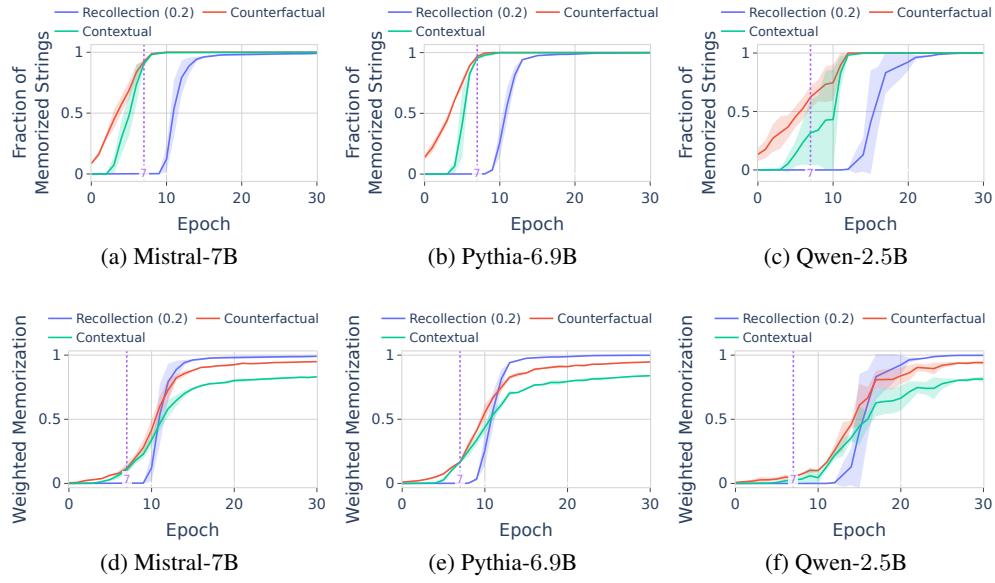


Figure 18: Memorization of training strings in languages of different entropy across different memorization measures. Results are for language L_1 , which is a high entropy language.

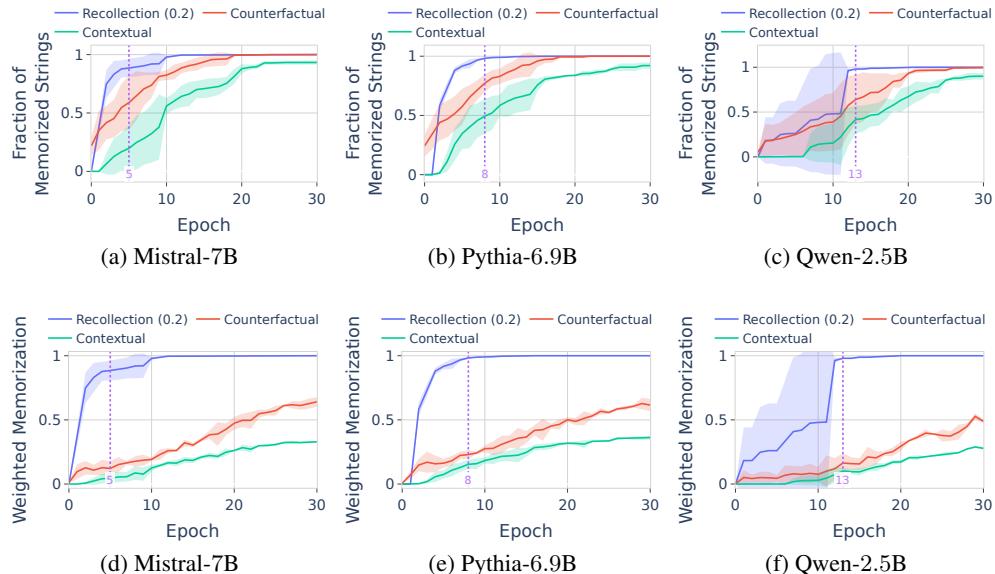


Figure 19: Memorization of training strings in languages of different entropy across different memorization measures. Results are for language L_2 , which is a low entropy language.

1505

1506

1507

1508

1509

1510

1511

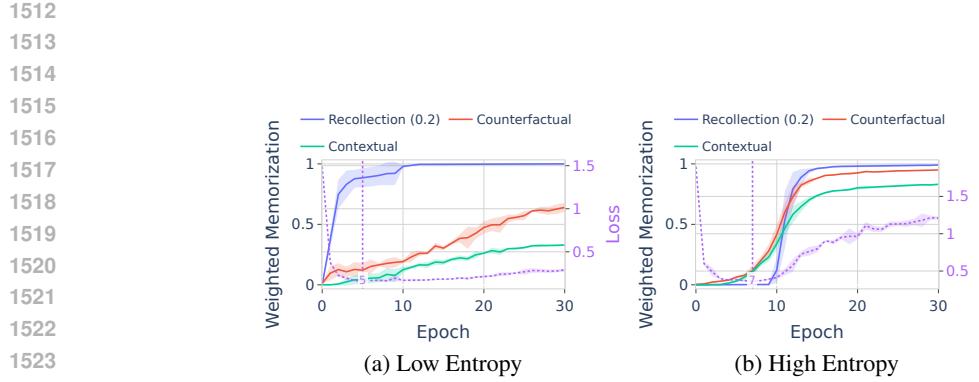


Figure 20: Continuing Figure 3, we demonstrate associated loss with weighted memorization.

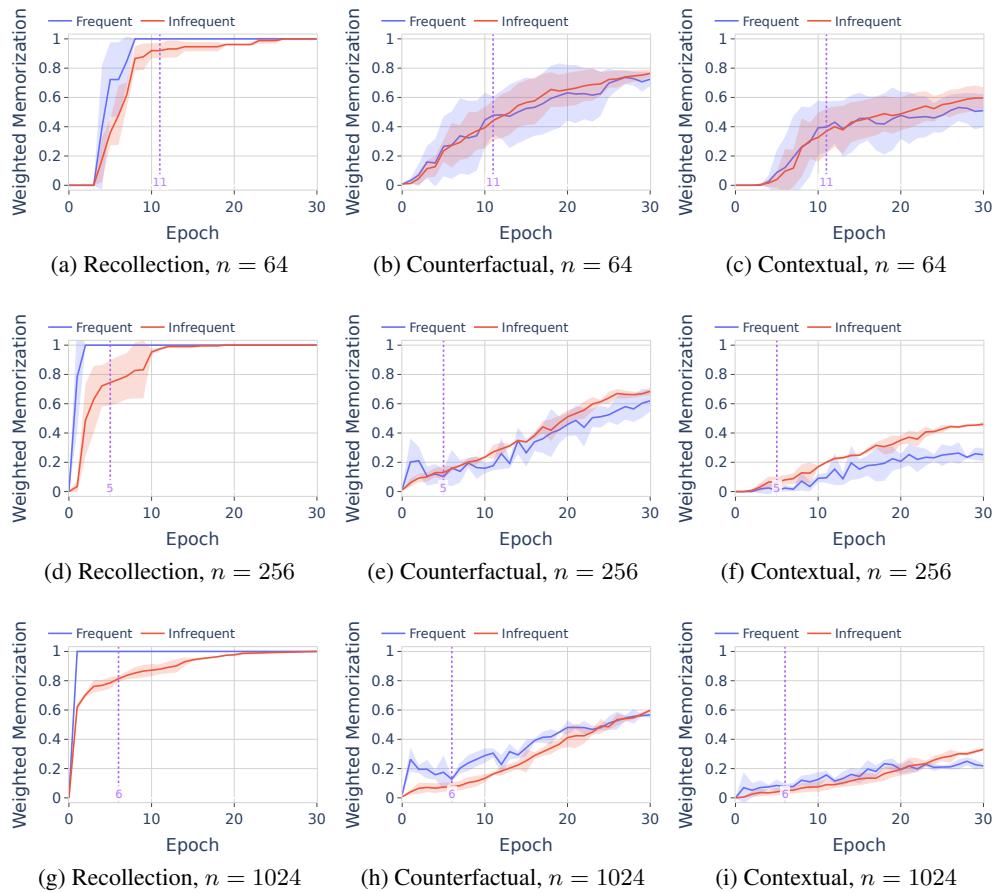


Figure 21: Continuing Figure 4, contradiction between recollection-based and contextual (or counterfactual) memorization on determining memorization of top 10% frequent strings and bottom 10% infrequent strings in a low entropy language. The results are for Mistral-7B on language L_2 , which is a low entropy language.

1566

1567

1568

1569

1570

1571

1572

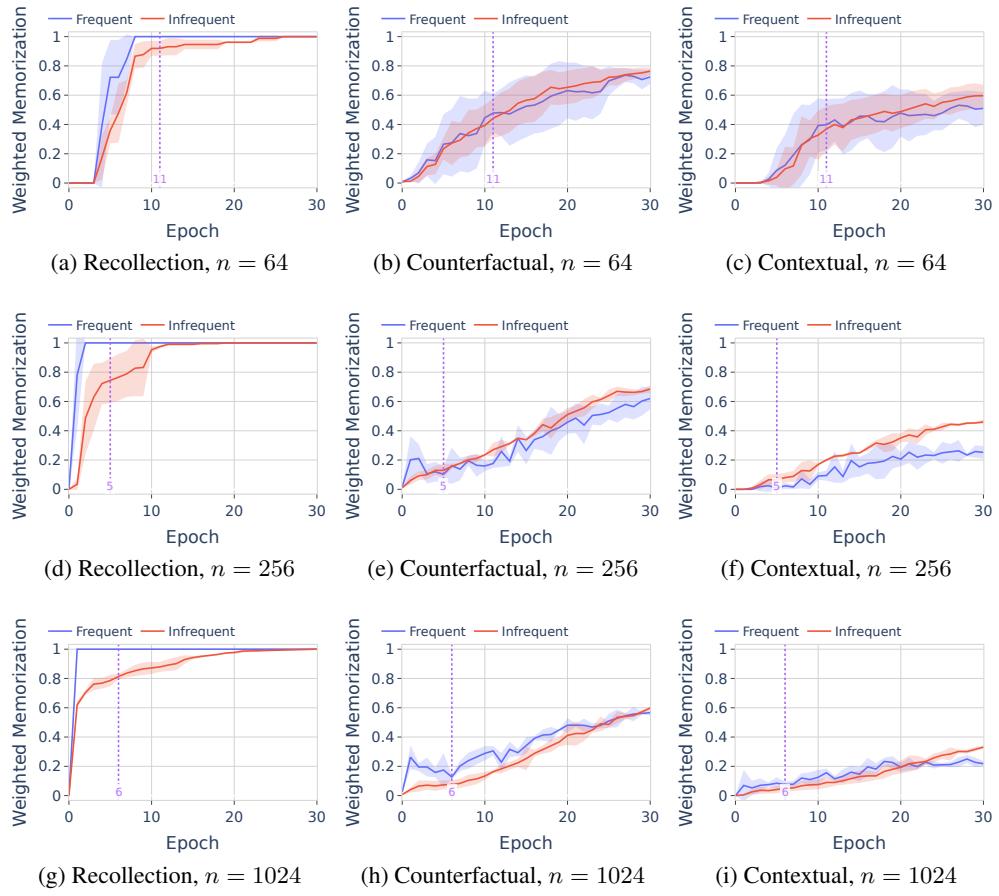
1573

1574

1575

1576

1577



1606
1607
1608
1609
Figure 22: Continuing Figure 21, contradiction between recollection-based and contextual (or counterfactual) memorization on determining memorization of top 10% frequent strings and bottom 10% infrequent strings in a low entropy language. The results is for Mistral-7B on language L_2 , which is a low entropy language.

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619

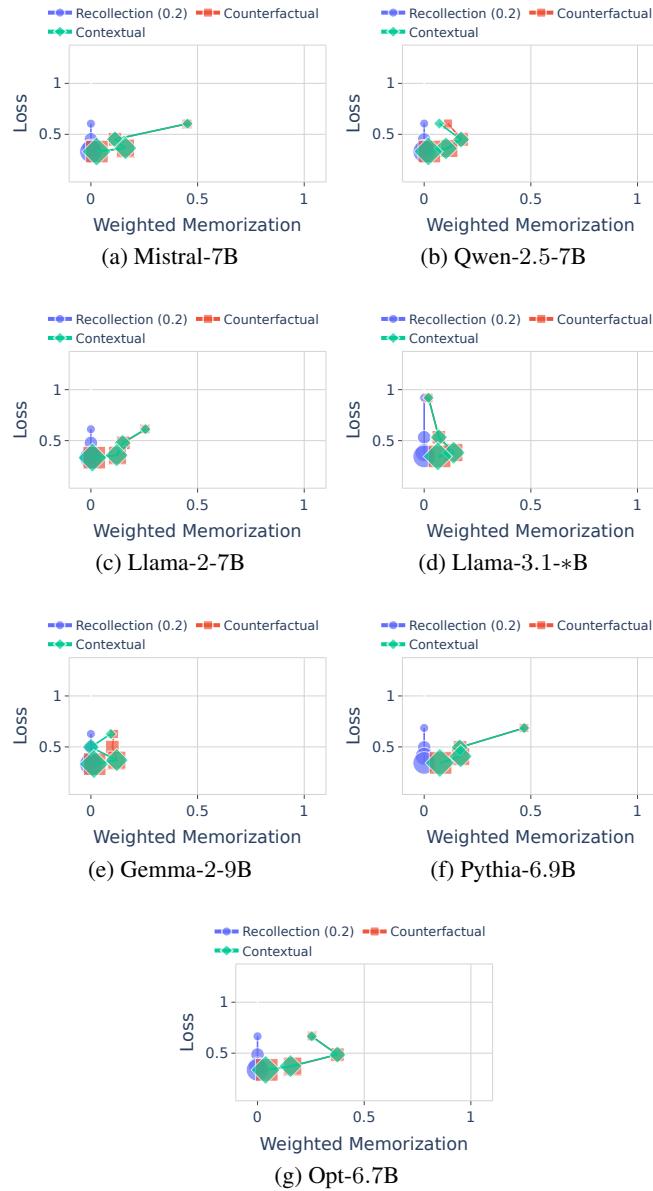


Figure 23: Tradeoffs between optimal learning and memorization among comparable $\approx 7B$ parameter size models on language L_1 , which is a high entropy language.

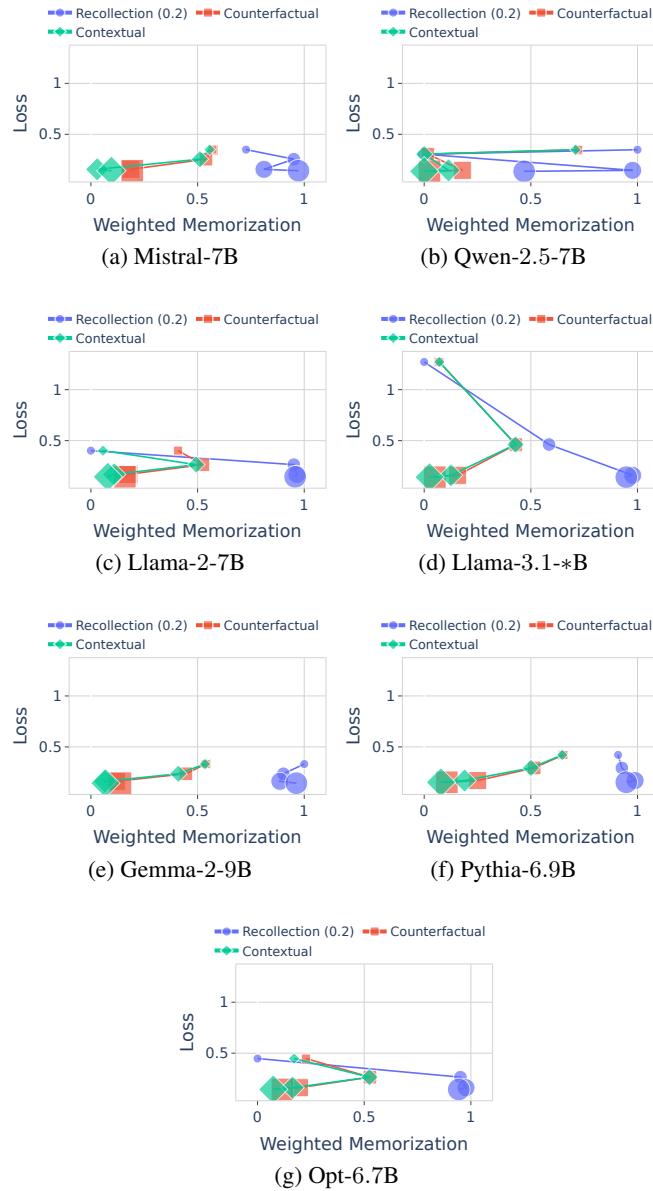


Figure 24: Tradeoffs between optimal learning and memorization among comparable $\approx 7B$ parameter size models on language L_2 , which is a low entropy language.

1728

1729

1730

1731

1732

1733

1734

1735

1736

1737

1738

1739

1740

1741

1742

1743

1744

1745

1746

1747

1748

1749

1750

1751

1752

1753

1754

1755

1756

1757

1758

1759

1760

1761

1762

1763

1764

1765

1766

1767

1768

1769

1770

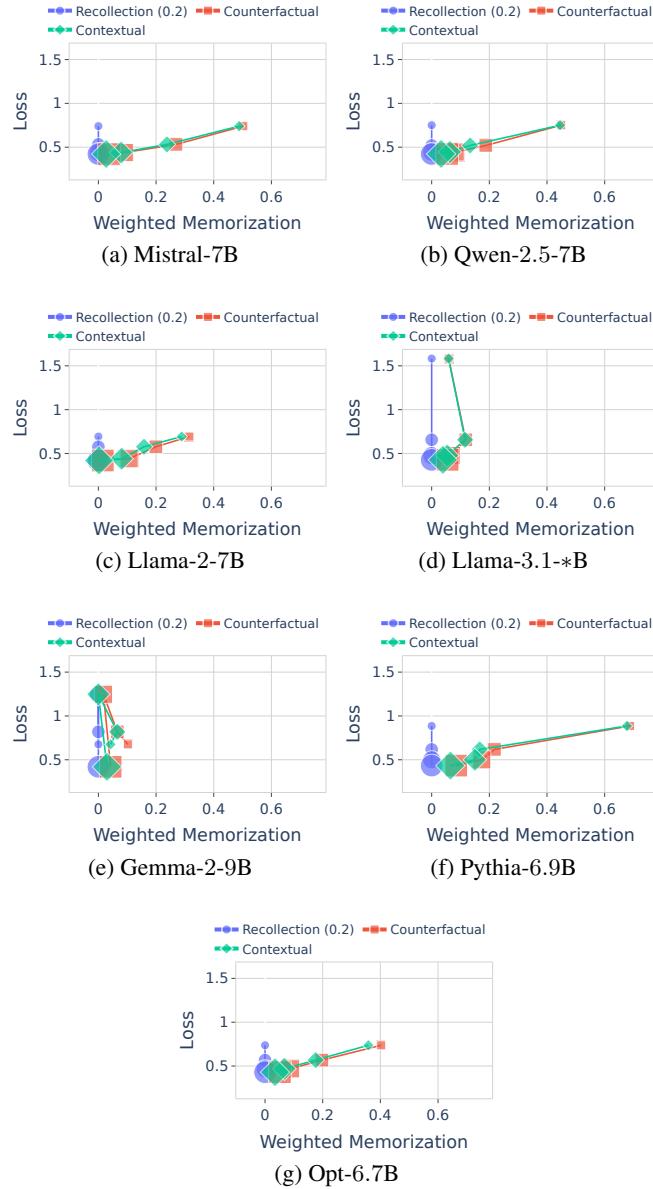
1771

1772

1773

1774

Figure 25: Tradeoffs between optimal learning and memorization among comparable $\approx 7B$ parameter size models on language L_3 , which is a high entropy language.



1782

1783

1784

1785

1786

1787

1788

1789

1790

1791

1792

1793

1794

1795

1796

1797

1798

1799

1800

1801

1802

1803

1804

1805

1806

1807

1808

1809

1810

1811

1812

1813

1814

1815

1816

1817

1818

1819

1820

1821

1822

1823

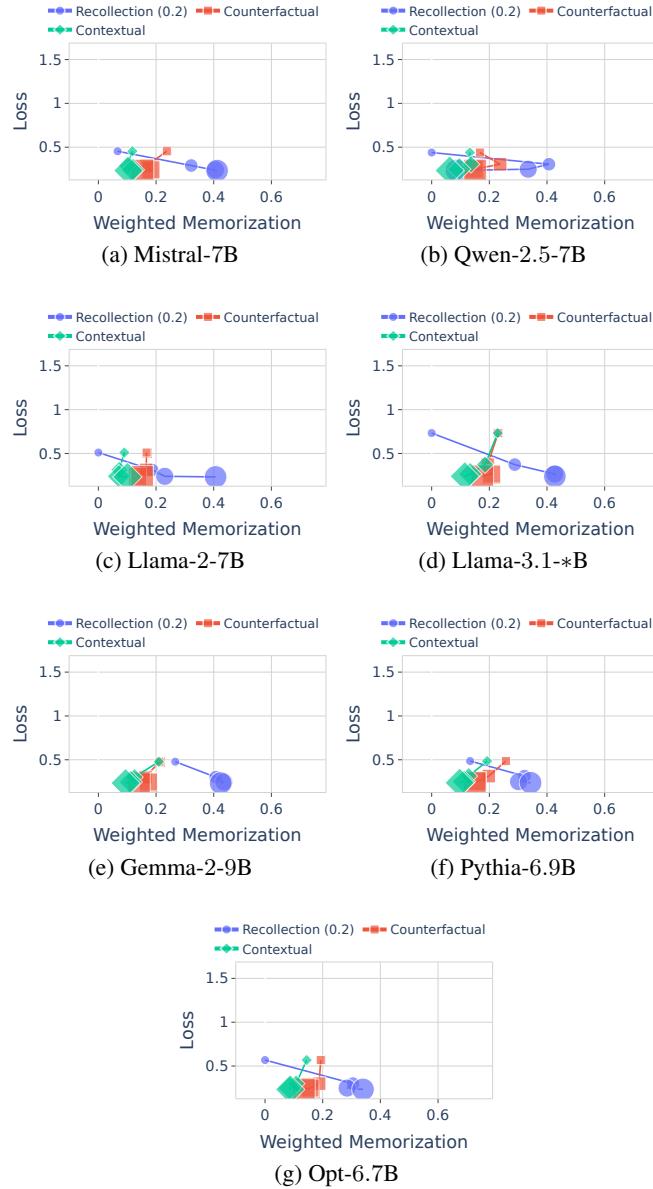
1824

1825

1826

1827

1828

Figure 26: Tradeoffs between optimal learning and memorization among comparable $\approx 7B$ parameter size models on language L_4 , which is a low entropy language.

1836

1837

1838

1839

1840

1841

1842

1843

1844

1845

1846

1847

1848

1849

1850

1851

1852

1853

1854

1855

1856

1857

1858

1859

1860

1861

1862

1863

1864

1865

1866

1867

1868

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

Figure 27: Tradeoffs between optimal learning and memorization among comparable $\approx 7B$ parameter size models on language L_5 , which is a high entropy language.

1884

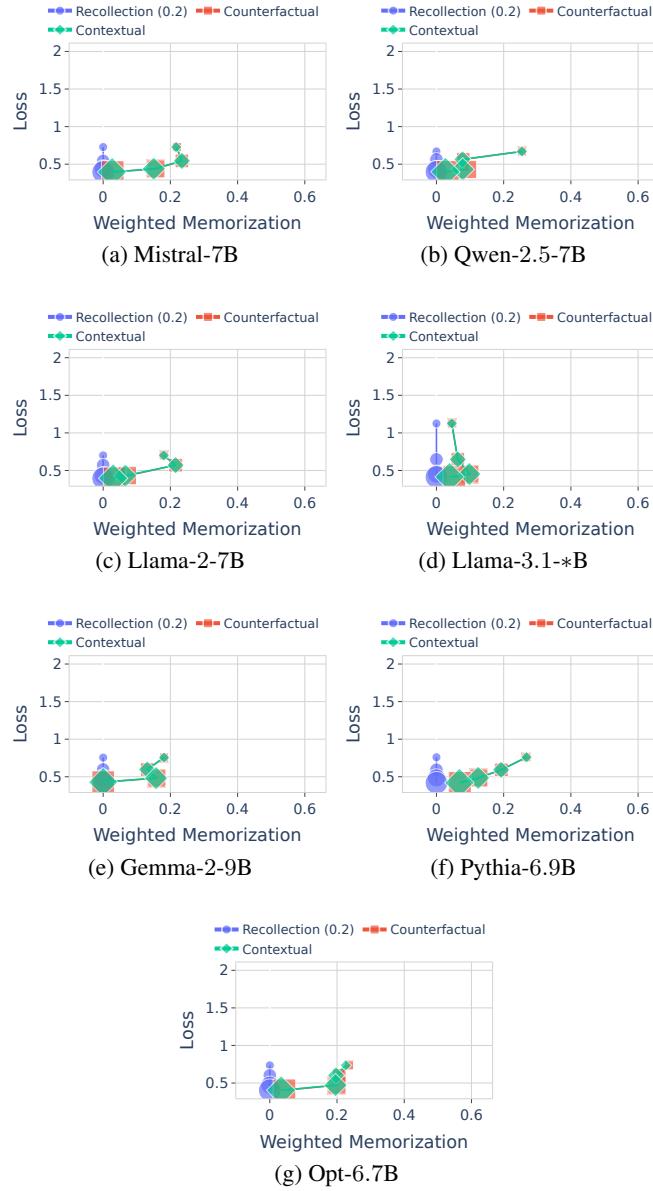
1885

1886

1887

1888

1889



1890

1891

1892

1893

1894

1895

1896

1897

1898

1899

1900

1901

1902

1903

1904

1905

1906

1907

1908

1909

1910

1911

1912

1913

1914

1915

1916

1917

1918

1919

1920

1921

1922

1923

1924

1925

1926

1927

1928

1929

1930

1931

1932

1933

1934

1935

1936

1937

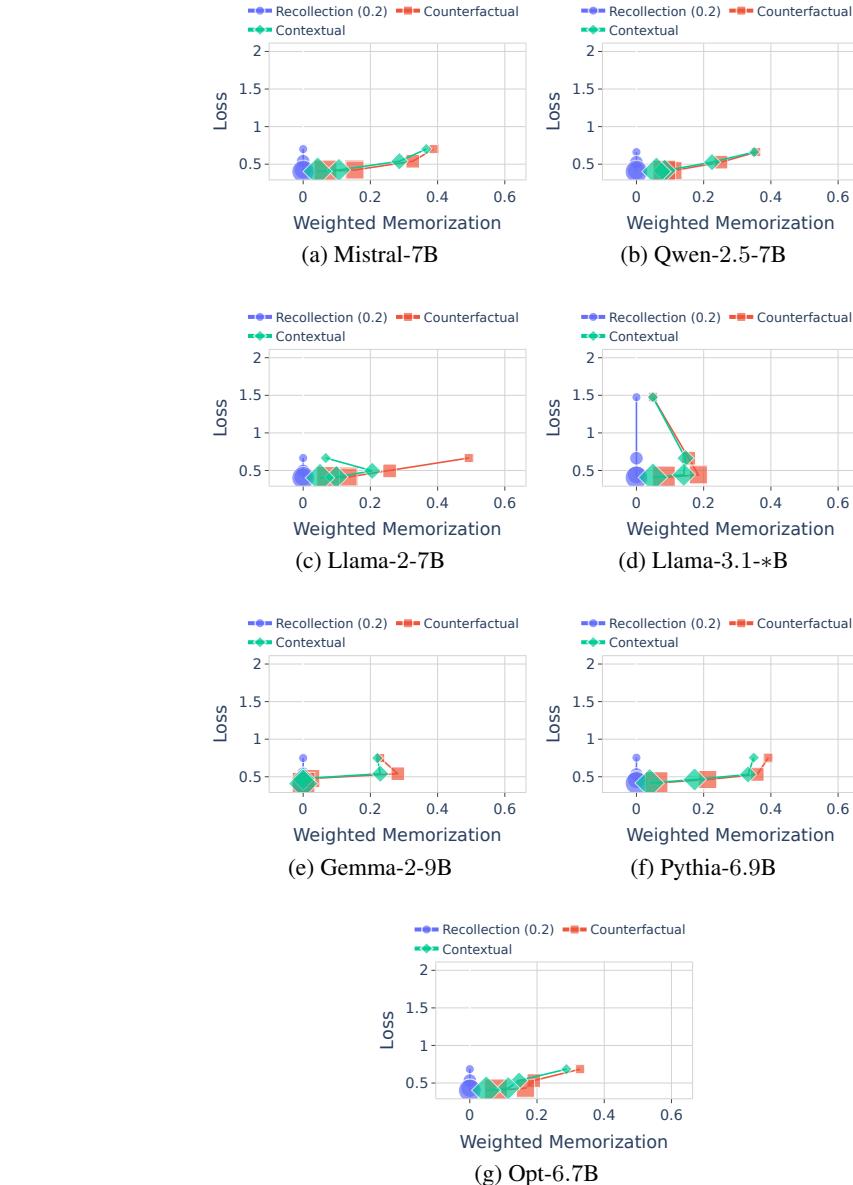


Figure 28: Tradeoffs between optimal learning and memorization among comparable ≈ 7 B parameter size models on language L_6 , which is a high entropy language.

1938

1939

1940

1941

1942

1943

1944

1945

1946

1947

1948

1949

1950

1951

1952

1953

1954

1955

1956

1957

1958

1959

1960

1961

1962

1963

1964

1965

1966

1967

1968

1969

1970

1971

1972

1973

1974

1975

1976

1977

1978

1979

1980

1981

1982

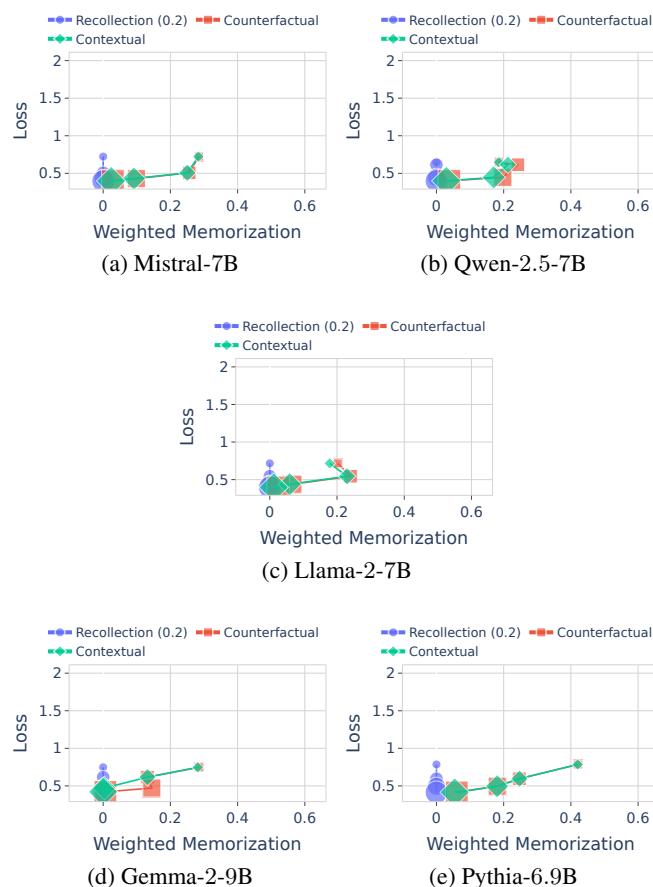
1983

1984

1985

1986

Figure 29: Tradeoffs between optimal learning and memorization among comparable $\approx 7B$ parameter size models on language L_7 , which is a high entropy language.



1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

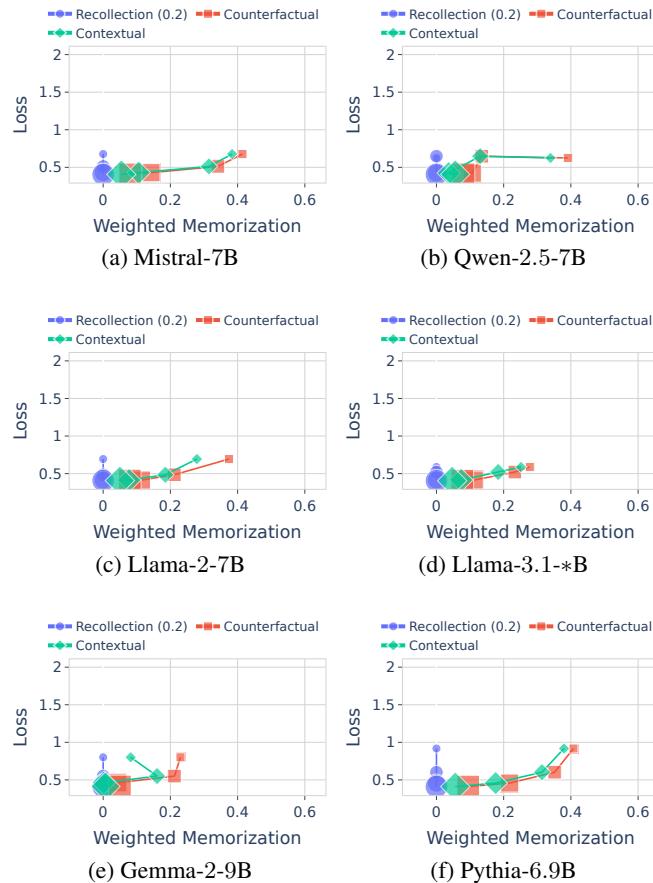


Figure 30: Tradeoffs between optimal learning and memorization among comparable $\approx 7B$ parameter size models on language L_8 , which is a high entropy language.

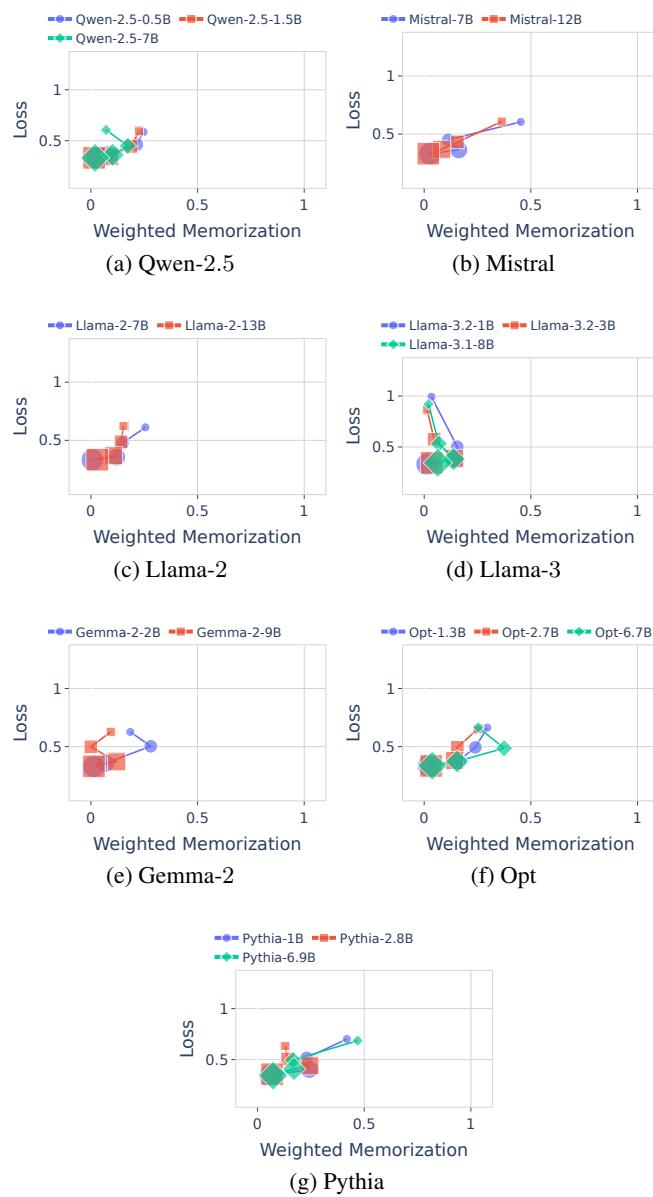


Figure 31: Contextual memorization vs. optimal language learning, measured as test loss, across models of different sizes within a family. Results are on language L_1 , which is a high entropy language.

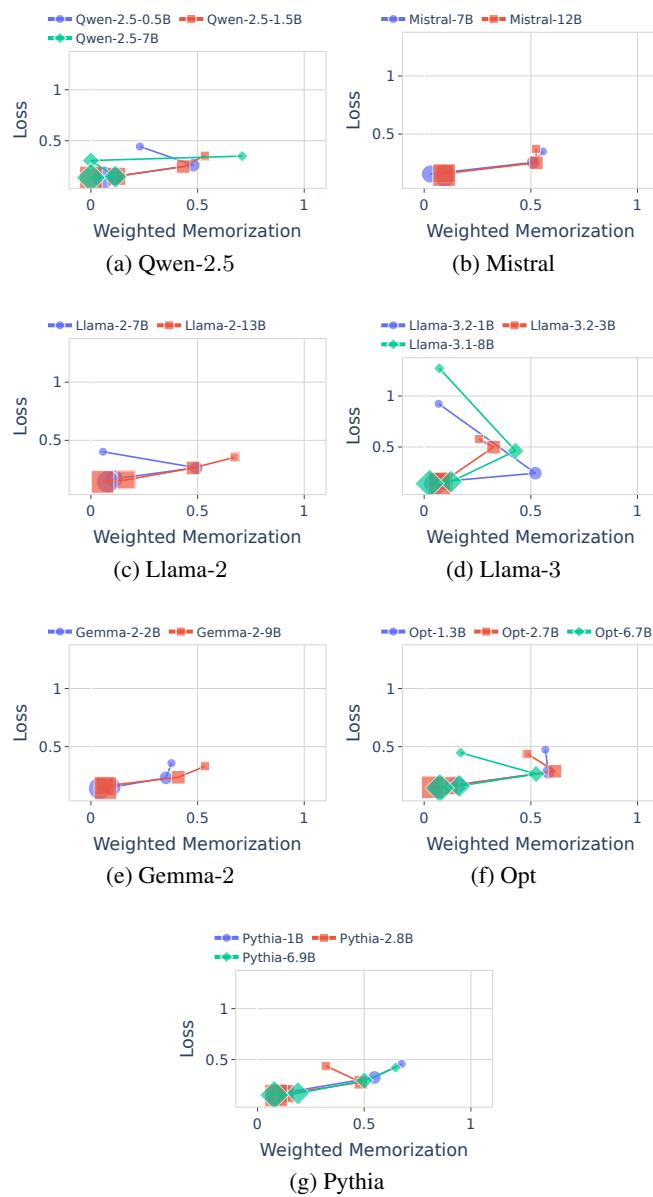


Figure 32: Contextual memorization vs. optimal language learning, measured as test loss, across models of different sizes within a family. Results are on language L_2 , which is a low entropy language.

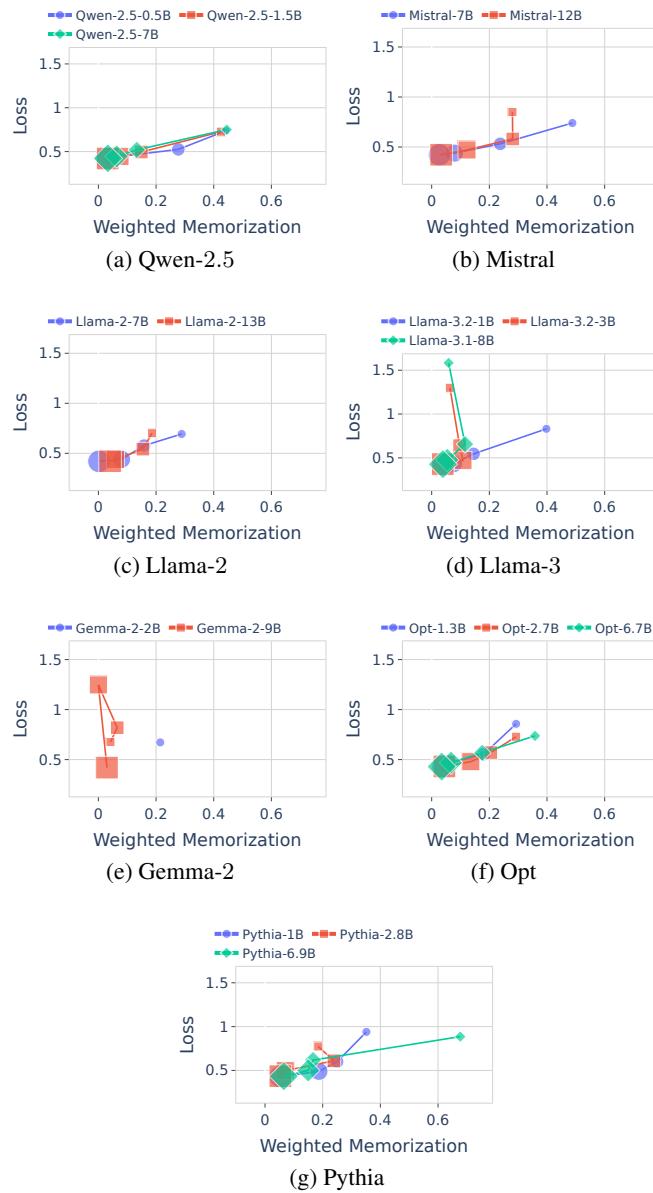


Figure 33: Contextual memorization vs. optimal language learning, measured as test loss, across models of different sizes within a family. Results are on language L_3 , which is a high entropy language.

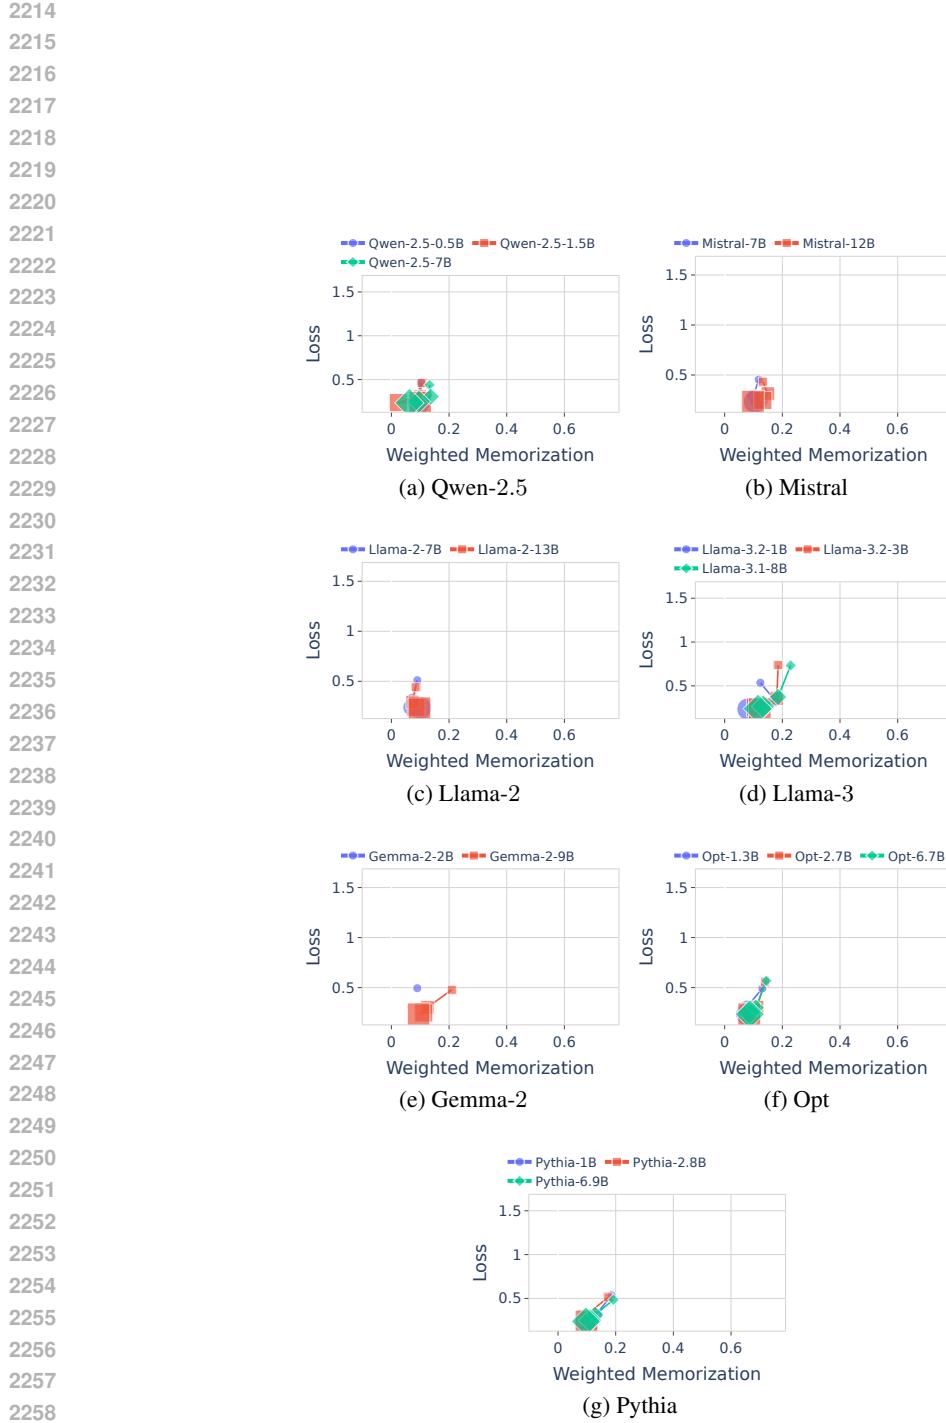


Figure 34: Contextual memorization vs. optimal language learning, measured as test loss, across models of different sizes within a family. Results are on language L_4 , which is a low entropy language.

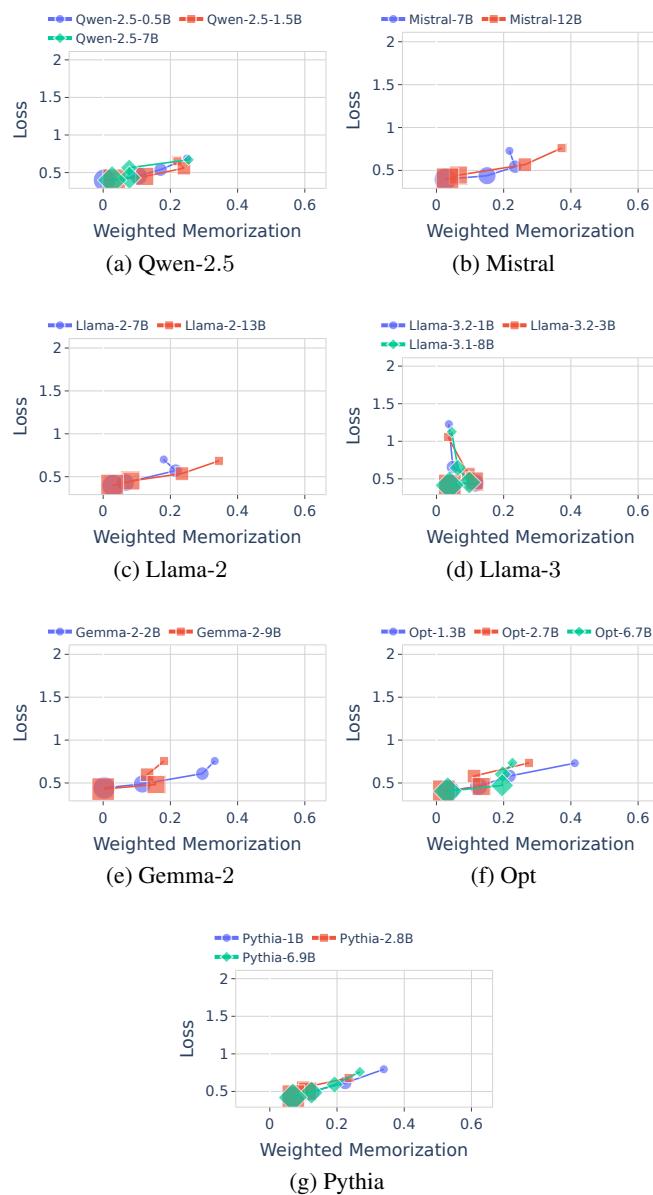


Figure 35: Contextual memorization vs. optimal language learning, measured as test loss, across models of different sizes within a family. Results are on language L_5 , which is a high entropy language.

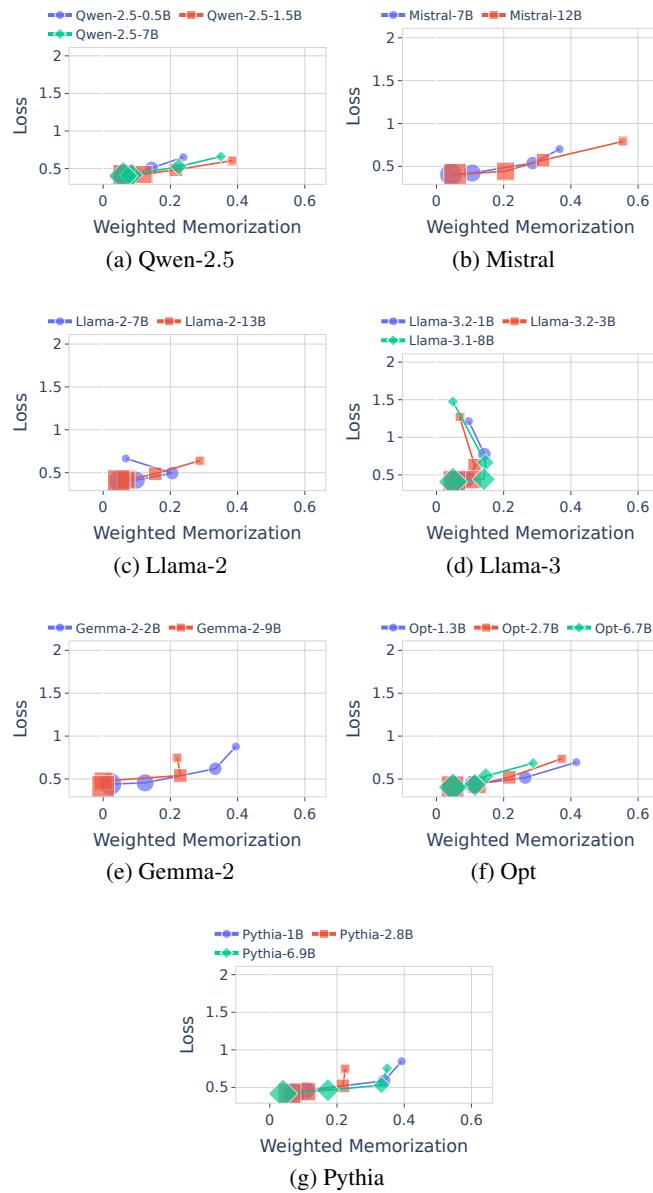


Figure 36: Contextual memorization vs. optimal language learning, measured as test loss, across models of different sizes within a family. Results are on language L_6 , which is a high entropy language.

2376

2377

2378

2379

2380

2381

2382

2383

2384

2385

2386

2387

2388

2389

2390

2391

2392

2393

2394

2395

2396

2397

2398

2399

2400

2401

2402

2403

2404

2405

2406

2407

2408

2409

2410

2411

2412

2413

2414

2415

2416

2417

2418

2419

2420

2421

2422

2423

2424

2425

2426

2427

2428

2429

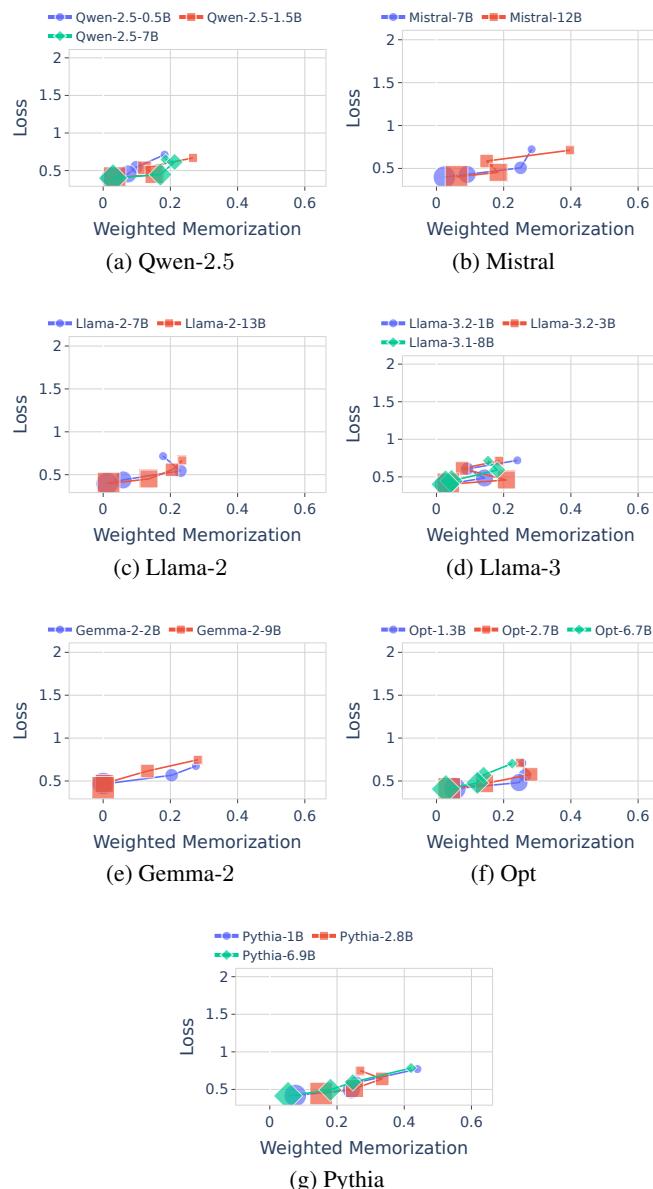


Figure 37: Contextual memorization vs. optimal language learning, measured as test loss, across models of different sizes within a family. Results are on language L_7 , which is a high entropy language and contains Latin characters as tokens.

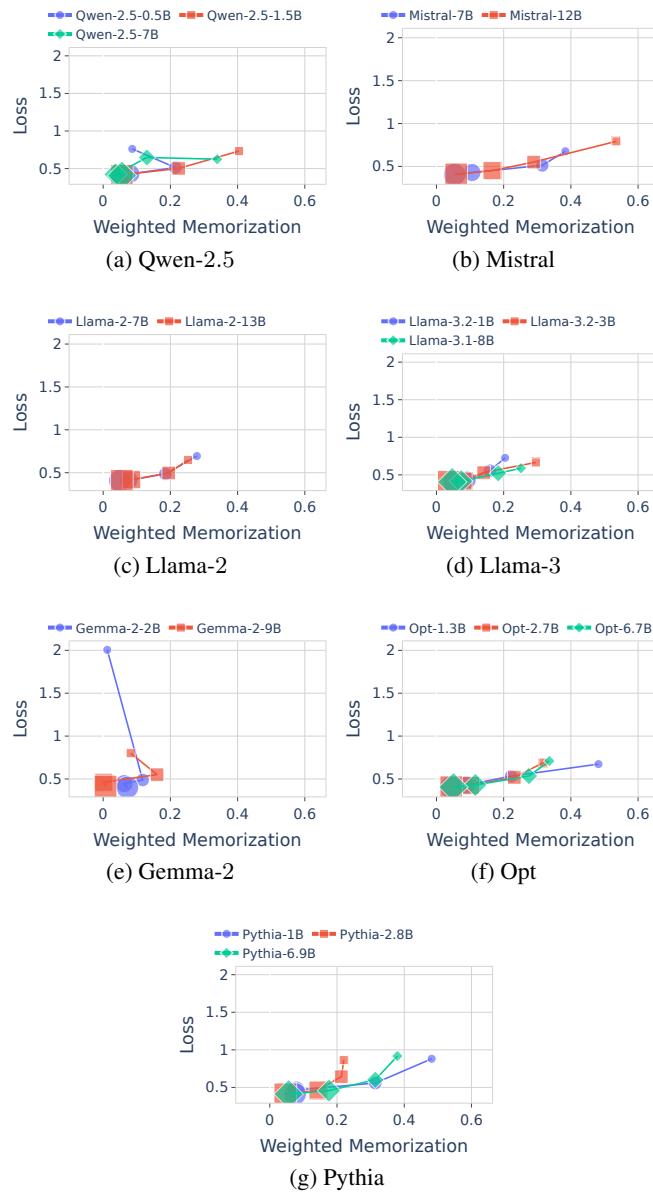
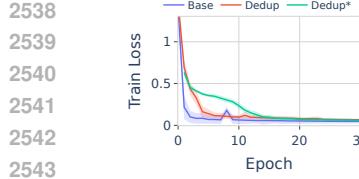


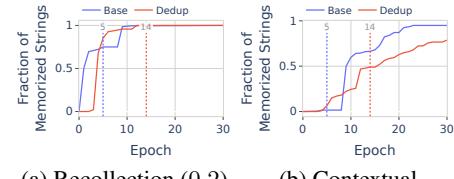
Figure 38: Contextual memorization vs. optimal language learning, measured as test loss, across models of different sizes within a family. Results are on language L_8 , which is a high entropy language and contains Latin characters as tokens.

2484 Table 3: List of recollection-based memorized strings by Pythia-1B-deduped (Biderman et al.,
2485 2024), where many strings can be contextually recollected, i.e., repeated words, predictable gen-
2486 eration, etc. We report the upper bound (UB) of contextual accuracy using a reference model
2487 OLMo-1B, which is trained on a different dataset than used in Pythia-1B-deduped. Considering
2488 the high accuracy of the OLMo-1B on memorized strings by Pythia-1B-deduped, we suspect that
2489 the highlighted generations are **not contextually memorized**.

2490	2491	2492	Prompt + Generation	Accuracy of Generation		Remark																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																							
2493	2494	2495	2496	2497	2498	2499	2500	2501	2502	2503	2504	2505	2506	2507	2508	2509	2510	2511	2512	2513	2514	2515	2516	2517	2518	2519	2520	2521	2522	2523	2524	2525	2526	2527	2528	2529	2530	2531	2532	2533	2534	2535	2536	2537	2538	2539	2540	2541	2542	2543	2544	2545	2546	2547	2548	2549	2550	2551	2552	2553	2554	2555	2556	2557	2558	2559	2560	2561	2562	2563	2564	2565	2566	2567	2568	2569	2570	2571	2572	2573	2574	2575	2576	2577	2578	2579	2580	2581	2582	2583	2584	2585	2586	2587	2588	2589	2590	2591	2592	2593	2594	2595	2596	2597	2598	2599	2600	2601	2602	2603	2604	2605	2606	2607	2608	2609	2610	2611	2612	2613	2614	2615	2616	2617	2618	2619	2620	2621	2622	2623	2624	2625	2626	2627	2628	2629	2630	2631	2632	2633	2634	2635	2636	2637	2638	2639	2640	2641	2642	2643	2644	2645	2646	2647	2648	2649	2650	2651	2652	2653	2654	2655	2656	2657	2658	2659	2660	2661	2662	2663	2664	2665	2666	2667	2668	2669	2670	2671	2672	2673	2674	2675	2676	2677	2678	2679	2680	2681	2682	2683	2684	2685	2686	2687	2688	2689	2690	2691	2692	2693	2694	2695	2696	2697	2698	2699	2700	2701	2702	2703	2704	2705	2706	2707	2708	2709	2710	2711	2712	2713	2714	2715	2716	2717	2718	2719	2720	2721	2722	2723	2724	2725	2726	2727	2728	2729	2730	2731	2732	2733	2734	2735	2736	2737	2738	2739	2740	2741	2742	2743	2744	2745	2746	2747	2748	2749	2750	2751	2752	2753	2754	2755	2756	2757	2758	2759	2760	2761	2762	2763	2764	2765	2766	2767	2768	2769	2770	2771	2772	2773	2774	2775	2776	2777	2778	2779	2780	2781	2782	2783	2784	2785	2786	2787	2788	2789	2790	2791	2792	2793	2794	2795	2796	2797	2798	2799	2800	2801	2802	2803	2804	2805	2806	2807	2808	2809	2810	2811	2812	2813	2814	2815	2816	2817	2818	2819	2820	2821	2822	2823	2824	2825	2826	2827	2828	2829	2830	2831	2832	2833	2834	2835	2836	2837	2838	2839	2840	2841	2842	2843	2844	2845	2846	2847	2848	2849	2850	2851	2852	2853	2854	2855	2856	2857	2858	2859	2860	2861	2862	2863	2864	2865	2866	2867	2868	2869	2870	2871	2872	2873	2874	2875	2876	2877	2878	2879	2880	2881	2882	2883	2884	2885	2886	2887	2888	2889	2890	2891	2892	2893	2894	2895	2896	2897	2898	2899	2900	2901	2902	2903	2904	2905	2906	2907	2908	2909	2910	2911	2912	2913	2914	2915	2916	2917	2918	2919	2920	2921	2922	2923	2924	2925	2926	2927	2928	2929	2930	2931	2932	2933	2934	2935	2936	2937	2938	2939	2940	2941	2942	2943	2944	2945	2946	2947	2948	2949	2950	2951	2952	2953	2954	2955	2956	2957	2958	2959	2960	2961	2962	2963	2964	2965	2966	2967	2968	2969	2970	2971	2972	2973	2974	2975	2976	2977	2978	2979	2980	2981	2982	2983	2984	2985	2986	2987	2988	2989	2990	2991	2992	2993	2994	2995	2996	2997	2998	2999	3000	3001	3002	3003	3004	3005	3006	3007	3008	3009	3010	3011	3012	3013	3014	3015	3016	3017	3018	3019	3020	3021	3022	3023	3024	3025	3026	3027	3028	3029	3030	3031	3032	3033	3034	3035	3036	3037	3038	3039	3040	3041	3042	3043	3044	3045	3046	3047	3048	3049	3050	3051	3052	3053	3054	3055	3056	3057	3058	3059	3060	3061	3062	3063	3064	3065	3066	3067	3068	3069	3070	3071	3072	3073	3074	3075	3076	3077	3078	3079	3080	3081	3082	3083	3084	3085	3086	3087	3088	3089	3090	3091	3092	3093	3094	3095	3096	3097	3098	3099	3100	3101	3102	3103	3104	3105	3106	3107	3108	3109	3110	3111	3112	3113	3114	3115	3116	3117	3118	3119	3120	3121	3122	3123	3124	3125	3126	3127	3128	3129	3130	3131	3132	3133	3134	3135	3136	3137	3138	3139	3140	3141	3142	3143	3144	3145	3146	3147	3148	3149	3150	3151	3152	3153	3154	3155	3156	3157	3158	3159	3160	3161	3162	3163	3164	3165	3166	3167	3168	3169	3170	3171	3172	3173	3174	3175	3176	3177	3178	3179	3180	3181	3182	3183	3184	3185	3186	3187	3188	3189	3190	3191	3192	3193	3194	3195	3196	3197	3198	3199	3200	3201	3202	3203	3204	3205	3206	3207	3208	3209	3210	3211	3212	3213	3214	3215	3216	3217	3218	3219	3220	3221	3222	3223	3224	3225	3226	3227	3228	3229	3230	3231	3232	3233	3234	3235	3236	3237	3238	3239	3240	3241	3242	3243	3244	3245	3246	3247	3248	3249	3250	3251	3252	3253	3254	3255	3256	3257	3258	3259	3260	3261	3262	3263	3264	3265	3266	3267	3268	3269	3270	3271	3272	3273	3274	3275	3276	3277	3278	3279	3280	3281	3282	3283	3284	3285	3286	3287	3288	3289	3290	3291	3292	3293	3294	3295	3296	3297	3298	3299	3300	3301	3302	3303	3304	3305	3306	3307	3308	3309	3310	3311	3312	3313	3314	3315	3316	3317	3318	3319	3320	3321	3322	3323	3324	3325	3326	3327	3328	3329	3330	3331	3332	3333	3334	3335	3336	3337	3338	3339	3340	3341	3342	3343	3344	3345	3346	3347	3348	3349	3350	3351	3352	3353	3354	3355	3356	3357	3358	3359	3360	3361	3362	3363	3364	3365	3366	3367	3368	3369	3370	3371	3372	3373	3374	3375	3376	3377	3378	3379	3380	3381	3382	3383	3384	3385	3386	3387	3388	3389	3390	3391	3392	3393	3394	3395	3396	3397	3398	3399	3400	3401	3402	3403	3404	3405	3406	3407	3408	3409	3410	3411	3412	3413	3414	3415	3416	3417	3418	3419	3420	3421	3422	3423	3424	3425	3426	3427	3428	3429	3430	3431	3432	3433	3434	3435	3436	3437	3438	3439	3440	3441	3442	3443	3444	3445	3446	3447	3448	3449	3450	3451	3452	3453	3454	3455	3456	3457	3458	3459	3460	3461	3462	3463	3464	3465	3466	3467	3468	3469	3470	3471	3472	3473	3474	3475	3476	3477	3478	3479	3480	3481	3482	3483	3484	3485	3486	3487	3488	3489	3490	3491	3492	3493	3494	3495	3496	3497	3498	3499	3500	3501	3502	3503	3504	3505	3506	3507	3508	3509	3510	3511	3512	3513	3514	3515	3516	3517	3518	3519	3520	3521	3522	3523	3524	3525	3526	3527	3528	3529	3530	3531	3532	3533	3534	3535	3536	3537	3538	3539	3540	3541	3542	3543	3544	3545	3546	3547	3548	3549	3550	3551	3552	3553	3554	3555	3556	3557	3558	3559	3560	3561	3562	3563	3564	3565	3566	3567	3568	3569	3570	3571	3572	3573	3574	3575	3576	3577	3578	3579	3580	3581	3582	3583	3584	3585	3586	3587	3588	3589	3590	3591	3592	3593	3594	3595	3596	3597	3598	3599	3600	3601	3602	3603	3604	3605	3606	3607	3608	3609	3610	3611	3612	3613	3614	3615	3616	3617	3618	3619	3620	3621	3622	3623	3624	3625	3626	3627	3628	3629	3630	3631	3632	3633	3634	3635	3636	3637	3638	3639	3640	3641	3642	3643	3644	3645	3646	3647	3648	3649	3650	3651	3652	3653	3654	3655	3656	3657	3658	3659	3660	3661	3662	3663	3664	3665	3666	3667	3668	3669	3670	3671	3672	3673	3674	3675	3676	3677	3678	3679	3680	3681	3682	3683	3684	3685	3686	3687	3688	3689	3690	3691	3692	3693	3694	3695	3696	3697	3698	3699	3700	3701	3702	3703	3704	3705	3706	3707	3708	3709	3710	3711	3712	3713	3714	3715	3716	3717	3718	3719	3720	3721	37



2544 Figure 39: Deduplication (dedup) on a low entropy
2545 language (base) does not result in a high entropy language,
2546 since *test strings are learned differently*, i.e.,
2547 high variance. But, training on the respective high entropy
2548 language (dedup*) results in low variance of the
2549 same test strings at optimal learning (vertical lines).



(a) Recollection (0.2) (b) Contextual

2544 Figure 40: Different training strings are be-
2545 ing memorized at different epochs while
2546 applying deduplication – an undesired re-
2547 sult. Deduplication is only effective in de-
2548 laying memorization, but not avoiding it.

G DEDUPLICATION: IMPACT ON LEARNING AND MEMORIZATION TRADEOFFS

2554 A critical observation of a low entropy language is that training strings are memorized at different
2555 epochs (Section 3), and test strings are learned to different extent. The goal of an ideal memorization
2556 mitigation mechanism should be two-fold: (1) **Learning goal:** an equal learning of all test strings,
2557 and (2) **Memorization goal:** a simultaneous memorization of all training strings, so that training can
2558 be stopped before memorization starts. *These goals are the inherent characteristics of a high entropy*
2559 *language.* Herein, we investigate whether deduplicating training strings of a low entropy language,
2560 i.e., by ensuring an equal string-frequency, achieves the behavior of a high entropy language.

2561 **RQ8.** Does deduplication of training strings lead to an equal learning of test strings?

2562 **RQ9.** Does deduplication lead to a simultaneous memorization of all training strings, and avoid
2563 memorization completely before optimal learning?

2564 **Deduplication neither results in an equal learning of test strings, nor leads to a simultaneous**
2565 **memorization of training strings.** Applying deduplication on a low entropy language in Figure 39,
2566 the frequency of training strings becomes uniform, similar to a high entropy language. Hence,
2567 training loss decreases slowly than the low entropy language (marked as ‘base’) with non-uniform
2568 string-frequency. However, different test strings are still learned differently at optimal learning,
2569 as if the language is still a low entropy one. Moreover, in Figure 40, different training strings
2570 are memorized at different epochs. The effectiveness of deduplication is thus limited to delaying
2571 memorization, but not avoiding it completely before optimal learning. Therefore, our answers to
2572 **RQ8** and **RQ9** are both negative; *deduplication indeed cannot convert a low entropy language to a high entropy one, where all test strings are often learned equally well (Figure 39, right), and memorization occurs simultaneously to most training strings (Figure 3a and 3c).*

2574 **Takeaway.** Deduplication does not lead to the equal learning of test strings or the simultaneous
2575 memorization of training strings. Our attempt to increase the entropy of strings from a low entropy
2576 language (via deduplication) does not make them behave like strings from a high entropy language.
2577 This implies that the language itself and not the frequency of the strings affects learning and mem-
2578 orization. In fact, deduplication of low entropy samples is unique element sampling from a skewed
2579 distribution, not uniform sampling from a uniform distribution – our results reflect this distinction.

2581
2582
2583
2584
2585
2586
2587
2588
2589
2590
2591