
Auditing Forgetting in Limited Memory Language Models

Hanna Roed¹ Arya Raeesi¹

Abstract

Limited Memory Language Models (LMLMs) externalize factual knowledge to a database to enable deletion-based unlearning without retraining. Existing evaluations measure post-deletion correctness in aggregate and cannot tell whether a deleted fact persists through residual parametric memory, alternative retrieval paths, or near-neighbor retrieval artifacts. We propose a causal auditing framework that holds the model fixed and varies the database state at inference time across three interventions: `FULL`, `DEL-ON`, and `DEL-OFF`. The framework decomposes post-deletion behavior into parametric leakage $L(f)$, retrieval-mediated correctness $R(f)$, and a retrieval artifact rate grounded in the inference-time retrieval trace. We apply it to 12,228 alias-closure deletions across thirteen databases, including four adversarial topologies (*Base*, *Alias*, *Noise*, *Collision*) we construct in three domains, and six prompt formulations. Parametric leakage is near zero in every variant and every prompt style: the model rarely returns the deleted answer in the absence of retrieval. The residual that does survive lives in the retrieval graph: retrieval-mediated correctness and the retrieval artifact rate match within rounding everywhere, so post-deletion correctness is, in our audit, predominantly reconstituted from near-neighbor retrieval. This residual ranges from 0.7% on the released LMLM database to 13.6% on the most adversarial variant, and prompt formulation does not independently control how much of a deleted fact survives. These results suggest that, for this class of LMLM and deletion procedure, the unlearning boundary is drawn primarily by the database administrator rather than by the model.

^{*}Equal contribution ¹University of California, Berkeley, Berkeley, California, United States of America. Correspondence to: Hanna Roed <hanna.roed@berkeley.edu>, Arya Raeesi <ar-raeesi@berkeley.edu>.

1. Introduction

Modern language models increasingly rely on hybrid architectures that combine parametric knowledge with external memory. Limited Memory Language Models (LMLMs) are a prominent example of this paradigm, explicitly separating linguistic competence encoded in model parameters from factual knowledge stored in an external database (Zhao et al., 2025). As illustrated in Figure 1, an LMLM retains the linguistic competence of a standard language model but routes factual recall through an external database rather than holding it in parameters. This design enables deletion-based unlearning, where removing entries from the database is intended to eliminate access to specific facts without requiring retraining (Zhao et al., 2025). Such capabilities are particularly important for applications involving data governance, privacy, and model editing.




Figure 1. Comparison of a standard retrieval-augmented language model (LLM + RAG) and a LMLM. Both architectures pair a parametric model with an external database, but LMLMs are pre-trained to limit the internal storage of factual knowledge, so factual recall is routed through the external store rather than reconstructed from parameters (Zhao et al., 2025).

However, it remains unclear whether deletion in these systems truly removes knowledge. Existing evaluations of forgetting typically measure whether a model produces the correct answer before and after deletion, but do not distinguish the underlying mechanism of post-deletion correctness (Zhao et al., 2025). A model may still answer correctly due to residual parametric memory, alternative retrieval paths, or semantically related matches in the external database. As a result, current metrics cannot determine whether knowledge has been successfully externalized or whether it persists internally in the model.

In this work, we propose a causal auditing framework for analyzing forgetting in LMLMs. Our approach introduces

three controlled interventions, FULL, where the database is intact and retrieval is enabled, DEL-ON, where relevant entries are deleted while retrieval remains enabled, and DEL-OFF, where the same deletion is applied and retrieval is disabled. By comparing model behavior across these conditions, we isolate the contribution of external retrieval and quantify residual internal knowledge. This enables a decomposition of post-deletion behavior into parametric leakage, retrieval-mediated correctness, and retrieval artifacts.

We implement this framework using facts drawn directly from the LMLM database (Zhao & contributors, 2025) and other self-made databases, and apply a verified deletion procedure that removes all canonical and alias-equivalent representations of a fact. We further log retrieval traces during inference to attribute model outputs to explicit database evidence. This controlled setup provides a principled method for auditing whether knowledge has been successfully externalized in LMLMs.

2. Related Work

Our work sits at the intersection of three lines of research. First, retrieval-augmented language models pair parametric representations with external memory accessed at inference time (Lewis et al., 2020; Guu et al., 2020), and the LMLM architecture we audit (Zhao et al., 2025) is a recent extension that pre-trains the model to limit internal storage of factual knowledge so that retrieval becomes the primary factual channel. This line of work shows that external retrieval can improve factual accuracy (Karpukhin et al., 2020), but leaves open the question of whether knowledge is genuinely externalized or partly retained in parameters (Mallen et al., 2023).

Second, knowledge-editing methods such as ROME (Meng et al., 2022) and MEMIT (Meng et al., 2023) address the same question from the opposite direction: rather than externalizing knowledge, they locate and modify parametric associations in place (Yao et al., 2023). LMLMs aim for the cleaner separation that ROME and MEMIT bypass, which makes the LMLM setting a natural place to ask whether the separation actually holds.

Third, work on machine unlearning has probed whether specific training data can be removed from a model after the fact (Bourtoule et al., 2021), and has tended to find that residual traces are difficult to fully eliminate (Lizzo & Heck, 2026; Carlini et al., 2021). Our audit takes up this question in the LMLM setting, where the question becomes tractable in a new way: because the database is the intended factual store, deletion can be applied at inference time and the resulting behavior decomposed into a parametric channel and a retrieval channel separately. The framework we develop in Section 3 is, to our knowledge, the first to make that de-

composition per-fact and to attribute surviving correctness to specific retrieval candidates.

3. Method Setup: Causal Audit Framework

We evaluate whether Limited Memory Language Models actually forget facts when those facts are removed from the external database. Each fact is represented as a subject–relation–object tuple, such as *Geri Halliwell – Famous For – Spice Girls*. At inference time, retrieval provides the model with relevant database entries as additional context before the model generates an answer. Thus, when retrieval is enabled, the model is not reading the entire database directly; instead, a retrieval step selects relevant facts and inserts them into the model’s input context. Appendix A.1 gives examples of the prompt formulations used in the experimental grid.

Our audit compares three intervention states. In FULL, the target fact remains in the database and retrieval is enabled. This measures normal database-supported accuracy. In DEL-ON, the target fact is deleted but retrieval remains enabled. This tests whether the answer can still be recovered through alternative database entries, aliases, semantically related facts, or retrieval artifacts. In DEL-OFF, the target fact is deleted and retrieval is disabled. This isolates parametric recall, since the model must answer without retrieved evidence.

We use these interventions to decompose post-deletion correctness into three mechanisms. Parametric leakage occurs when the model answers correctly in DEL-OFF, indicating that the deleted fact may still be stored in the model parameters. Retrieval-mediated correctness occurs when the model is correct in DEL-ON but not in DEL-OFF, indicating that retrieval helped recover the answer after deletion. Retrieval artifacts occur when the model produces the correct answer even though the deleted fact is not directly available as retrieved evidence. Formally, for a fact $f = (s, r, o)$ and intervention condition c , let $Y(f, c)$ denote the normalized model prediction. We define parametric leakage as

$$L(f) = \mathbb{I}[Y(f, \text{DEL-OFF}) = o],$$

which indicates whether the model can recover the deleted fact without retrieval.

We define retrieval-mediated correctness as

$$R(f) = \mathbb{I}[Y(f, \text{DEL-ON}) = o \wedge Y(f, \text{DEL-OFF}) \neq o],$$

which captures cases where retrieval enables the correct answer after deletion.

Across a fact set \mathcal{F} , the empirical leakage rate is

$$\hat{L} = \frac{1}{|\mathcal{F}|} \sum_{f \in \mathcal{F}} L(f) = \frac{1}{|\mathcal{F}|} \sum_{f \in \mathcal{F}} \mathbb{I}[Y(f, \text{DEL-OFF}) = o].$$

Analogously, we estimate retrieval-mediated correctness by averaging $R(f)$ over \mathcal{F} .

4. Experimental Setup

In addition to the released LMLM database (Zhao & contributors, 2025), we developed databases for countries, politicians and sports. Within each of these three themes, we built four database variants (*Base*, *Alias*, *Noise*, and *Collision*), for a total of twelve custom databases. These four variants are designed to stress-test a distinct mechanism by which a deleted fact could remain accessible after canonical removal. *Base* contains only the canonical (s, r, o) triplet for each fact, leaving retrieval with no alternative path. *Alias* stores the same fact only under aliased subject and relation forms, with no canonical entry, and tests whether alias-closure deletion catches every surface realization. *Noise* augments *Base* with decoy triplets that route to the same object through paraphrased subjects (e.g., *Government of United States* \rightarrow *Seat of Government* \rightarrow *Washington, D.C.*), probing whether retrieval can recover the deleted answer via near-neighbor paraphrases. *Collision* augments *Base* with near-miss triplets that share the subject but route to a different object (e.g., *United States* \rightarrow *Largest City* \rightarrow *New York City*), probing whether retrieval drifts onto a confusable neighbor and returns a plausible but incorrect answer. Table 2 outlines an example of these four variants for the politician domain. We evaluate each target fact under six prompt formulations: direct questions, paraphrased questions, contextual questions, cloze prompts, continuations, and few-shot prompts; examples are provided in Appendix A.1.

Crossing the six prompt sets, the three intervention states, and the thirteen databases yields a fully crossed evaluation grid in which every target fact is scored on matched inputs across all conditions. For each cell we record exact match, token-level precision, recall, and F1, together with the cross-state quantities $L(f)$, $R(f)$, and the retrieval artifact rate; retrieval traces are logged at every FULL and DEL-ON call so that post-deletion correctness can be attributed to explicit database evidence rather than implicit model behavior. Figure 2 outlines this pipeline.

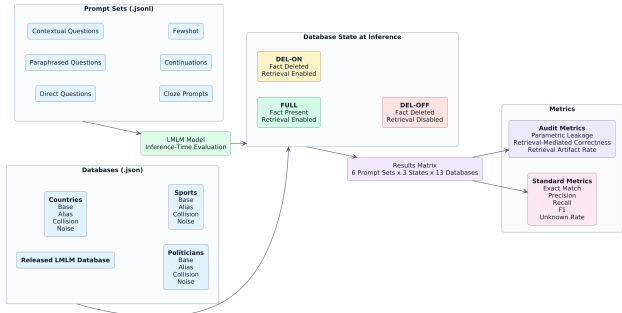


Figure 2. Overview of the evaluation pipeline. We evaluate six prompt sets using the released LMLM database, and specialized databases for countries, politicians and sports. We evaluate the databases under three inference-time database states: FULL, DEL-ON, and DEL-OFF. This yields a $6 \times 3 \times 13$ result matrix for each evaluation metric.

For each target fact, we construct the deletion set by enumerating the canonical triplet and all alias-equivalent triplets whose subject, relation, and object match the target under our alias mapping. We then remove this set from the database and verify deletion by checking that no retained triplet is gold-equivalent to the target. The same prompt is evaluated under FULL, DEL-ON, and DEL-OFF, and the generated answer is normalized before scoring against the gold object. During FULL and DEL-ON runs, we also save the retrieved candidates returned to the model. A DEL-ON answer is counted as a retrieval artifact when the normalized answer matches the gold object but none of the retained retrieval candidates is gold-equivalent to the deleted fact.

5. Results

Our reference point throughout the results is the aggregate evaluation reported in the original LMLM paper (Zhao et al., 2025), which uses a single FactScore (Min et al., 2023) drop when the database is disabled and a single TOFU (Maini et al., 2024) forget-quality p -value per unlearning step. Both quantities collapse the post-deletion behavior of every fact into a single number, so they cannot indicate which channel a surviving correct answer came through. Our framework refines that aggregate signal into per-fact attributions across the three intervention states, and we read each result below against the corresponding aggregate quantity from the original paper.

We run the LMLM under all three interventions on every cell of the prompt \times database grid. Across the 78 (prompt file, database) cells, this yields 12,228 paired (DEL-ON, DEL-OFF) evaluations together with an equal number of FULL baselines, for a total of 36,684 model completions. All reported quantities are count-weighted averages over fact-paired groups, so a prompt file with more target facts

contributes proportionally to the aggregate. We confine the present section to direct observations and defer cross-figure interpretation to Section 6.

We begin with the variant-level view. Figure 3 attributes DEL-ON correctness to its three components separately for the four custom variants and the released LMLM database. In particular, the retrieval artifact bar isolates cases in which DEL-ON returns the gold object even though no gold-equivalent triplet appears among the retrieval-trace candidates retained after deletion.

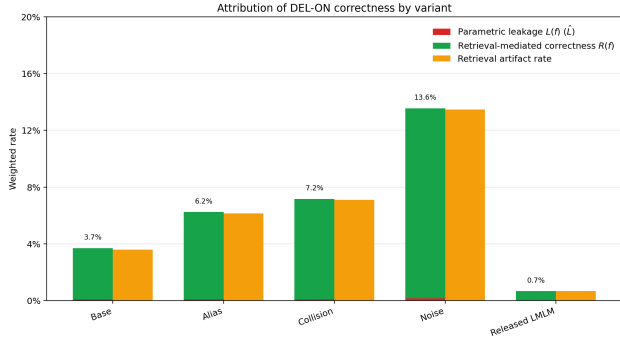


Figure 3. Attribution of DEL-ON correctness by database variant. The stacked left bar in each pair decomposes DEL-ON correctness into parametric leakage $L(f)$ (red, \hat{L}) and retrieval-mediated correctness $R(f)$ (green); the orange bar shows the retrieval artifact rate. Parametric leakage is near zero in every variant. The combined $L(f) + R(f)$ stack is 3.7% for *Base*, 6.2% for *Alias*, 7.2% for *Collision*, 13.6% for *Noise*, and 0.7% for the released LMLM database, with the orange artifact bar at approximately the same height as the green bar in each variant.

We next turn to the prompt-style axis. Figure 4 reports the same decomposition averaged over all four custom variants and the released LMLM database, so that variant identity is collapsed and only the effect of prompt formulation remains. As before, parametric leakage stays near zero across the panel, although a small red residue is visible under direct and few-shot prompts and is essentially absent under the others. The combined $L(f) + R(f)$ stack varies from 6.1% on cloze prompts to 9.9% on direct questions, while the orange artifact bar matches the green $R(f)$ bar within rounding in every prompt style.

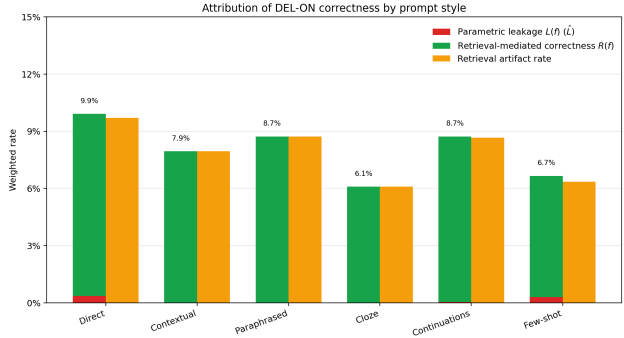


Figure 4. Attribution of DEL-ON correctness by prompt style. The left bar in each pair stacks parametric leakage $L(f)$ (red, \hat{L}) and retrieval-mediated correctness $R(f)$ (green); the orange bar shows the retrieval artifact rate. Leakage is near zero across prompt styles, with a small residue under direct and few-shot. The $L(f) + R(f)$ stack is 9.9% (direct), 7.9% (contextual), 8.7% (paraphrased), 6.1% (cloze), 8.7% (continuations), and 6.7% (few-shot); the orange bar matches the green within rounding in every style.

Having attributed DEL-ON correctness, we now widen the lens to all three intervention states. Figure 5 reports weighted token F1 by prompt style and state, so that the post-deletion residual can be compared against both the FULL baseline and the retrieval-disabled DEL-OFF condition. Under FULL, F1 spans a wide range, from ~17% on few-shot prompts to ~57% on continuations. Under DEL-ON, by contrast, the same prompt styles compress into a narrow band of roughly 7–10%. As a result, the largest FULL-to-DEL-ON drop falls on continuations (about 48 points), while the smallest falls on few-shot prompts (about 10 points). Finally, under DEL-OFF, F1 is visually indistinguishable from zero in every prompt style.

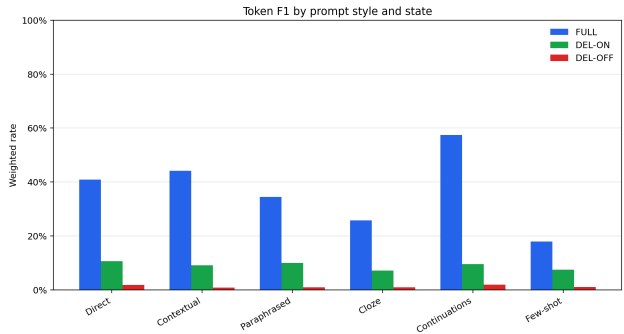


Figure 5. Weighted token F1 by prompt style and intervention state. FULL F1 is ~41% for direct, ~44% for contextual, ~34% for paraphrased, ~26% for cloze, ~57% for continuations, and ~17% for few-shot prompts. DEL-ON F1 sits between ~7% and ~10% across all six prompt styles. DEL-OFF F1 is near zero in every prompt style.

Figure 6 mirrors this three-state cut along the database-

variant axis and uses exact match in place of token F1, so that the same trajectory can be read off the variants we constructed for the audit. Under FULL, accuracy is similar across *Base*, *Alias*, and *Noise*, all clustered near 33%, and noticeably higher on *Collision* and the released LMLM database, both near 46%. The corresponding FULL-to-DEL-ON drop is about 30 points on *Base*, 27 on *Alias*, 39 on *Collision*, 19 on *Noise*, and 45 on the released LMLM database. Under DEL-OFF, exact match remains at or near zero in every variant.

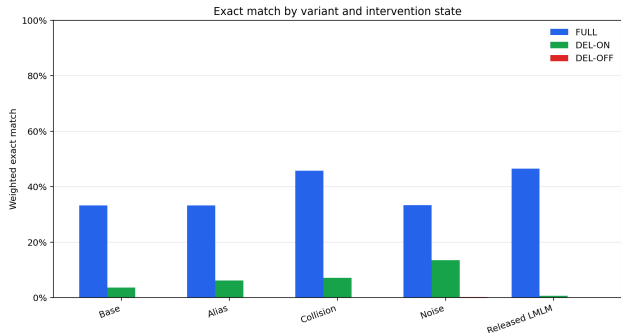


Figure 6. Weighted exact match by database variant and intervention state. FULL exact match is $\sim 33\%$ on *Base*, *Alias*, and *Noise*, and $\sim 46\%$ on *Collision* and the released LMLM database. DEL-ON exact match is $\sim 3\%$ on *Base*, $\sim 6\%$ on *Alias*, $\sim 7\%$ on *Collision*, $\sim 14\%$ on *Noise*, and $\sim 1\%$ on the released LMLM database. DEL-OFF exact match is near zero in every variant.

Tables 5 and 6 in Appendix E report the full per-cell numbers underlying Figures 3–6.

6. Analysis

Sections 6.1 and 6.2 establish the headline decomposition of post-deletion correctness into a parametric component and a retrieval-graph component. Sections 6.3 and 6.4 then ablate the two axes of the audit grid that the framework holds fixed, the retrieval-graph topology and the prompt formulation, to identify which axis controls the surviving residual. Section 6.5 closes the analysis with qualitative examples of the three failure modes the decomposition predicts.

6.1. Parametric externalization at the per-fact level

Our headline finding is that parametric leakage \hat{L} is near zero across the audit. The red $L(f)$ component contributes a vanishingly small share of the $L(f) + R(f)$ stack in every variant (Figure 3) and in every prompt style (Figure 4), and weighted DEL-OFF performance is visually indistinguishable from zero in every cell of Figures 5 and 6. At the level of model parameters, the LMLM checkpoint behaves close to a model that had never seen the fact, in contrast to standard language models, which are known to memorize

and expose training-data content even when each example is seen only a handful of times (Carlini et al., 2021). The original LMLM paper (Zhao et al., 2025) provides only an aggregate version of this result. Their Table 9 shows that disabling the database globally drops FactScore (Min et al., 2023) by roughly 19 points, and their §5 explicitly labels that finding as “preliminary evidence” that LMLM does not memorize. Our audit converts the same intuition into a per-fact statement: across 12,228 alias-closure deletions, the parameters return the deleted answer at a rate of $\hat{L} = 0.11\%$. Our per-fact result is consistent with, and adds resolution to, the preliminary aggregate finding in the original LMLM paper (Zhao et al., 2025).

6.2. The residual lives in the retrieval graph

The residual that survives deletion is, in our audit, almost entirely attributable to the retrieval graph rather than to the parameters. The signature observation in Figures 3 and 4 is the near-equality of the green $R(f)$ bar and the orange retrieval artifact bar: the two match within rounding for every variant and every prompt style. The following proposition shows that part of this equality is structural, which in turn clarifies what the figures empirically measure.

Proposition 6.1 (Artifact and retrieval-mediated correctness coincide under complete deletion). *Fix a fact $f = (s, r, o)$ with $L(f) = 0$, and suppose the alias-closure deletion of f is complete, i.e., no triplet alias-equivalent to (s, r, o) survives in the database after deletion. Let $G(f)$ be the indicator that some retained candidate in the inference-time trace is gold-equivalent (alias-equivalent on subject, relation, and object), and let*

$$A(f) = \mathbb{I}[Y(f, DEL-ON) = o] (1 - G(f))$$

be the per-fact retrieval artifact indicator. Then $R(f) = A(f)$.

Proof. Since $L(f) = 0$, we have $Y(f, DEL-OFF) \neq o$, so $R(f) = \mathbb{I}[Y(f, DEL-ON) = o]$. By completeness of alias-closure deletion, no retained candidate is gold-equivalent, hence $G(f) = 0$. Therefore $A(f) = \mathbb{I}[Y(f, DEL-ON) = o] = R(f)$. \square

Proposition 6.1 reframes what the two attribution figures actually measure. Given the per-fact $\hat{L} \approx 0$ established in Section 6.1, the empirical near-equality $\hat{R} \approx \hat{A}$ is a confirmation that our alias-closure deletion procedure is complete in practice; the empirical content of the figures is therefore not the equality itself but the topology-dependent magnitude of the residual. Retrieval-mediated correctness almost never coincides with a gold-equivalent retained candidate. Instead, it coincides with near-neighbor candidates that share embedding-space similarity with the deleted entry.

This decomposition is not present in the original LMLM paper. Their TOFU (Maini et al., 2024) evaluation reports a single forget-quality p -value per unlearning step and therefore aggregates over both gold-equivalent retrieval and near-neighbor retrieval. The retrieval-trace introspection used here separates the two and shows that, on retrieval graphs containing paraphrastic decoys, the second mechanism dominates. A standard FULL versus DEL-ON comparison would have shown only that some deleted facts are still answered correctly; the audit shows that those facts are not remembered, they are reconstituted at retrieval time.

6.3. Ablation: Retrieval-graph topology controls the residual

We ablate the retrieval-graph topology by holding the model, deletion procedure, and prompt distribution fixed and varying only the database variant across *Base*, *Alias*, *Collision*, *Noise*, and the released LMLM database; the dependent variable is the post-deletion residual $L(f) + R(f)$. Reading down Figure 3, the residual rises from 3.7% on *Base*, to 6.2% on *Alias*, to 7.2% on *Collision*, and to 13.6% on *Noise*, before falling to 0.7% on the released LMLM database. This ordering tracks the retrieval-graph topologies described in Section 4: variants in which more surviving entries point to the gold object after alias-closure deletion produce larger residuals, with *Noise* as the worst case by construction and *Base* as the cleanest among the custom databases. The original paper (Zhao et al., 2025) reports unlearning on a single retrieval graph, the annotated TOFU corpus (Maini et al., 2024), and on that graph achieves “ideal forgetting” with $p > 0.05$. By varying the retrieval graph along controlled topologies, we show that the same architecture and the same alias-closure deletion can produce a DEL-ON residual anywhere from 0.7% to 13.6%. The published TOFU result therefore measures the model on a benign retrieval graph (Maini et al., 2024; Zhao et al., 2025); it does not, by itself, characterize what happens when the graph contains paraphrastic decoys.

6.4. Ablation: Prompt formulation moves baselines more than residuals

We ablate the prompt formulation by holding the model, deletion procedure, and database distribution fixed and varying only the prompt set across the six formulations introduced in Section 4; the dependent variables are the FULL baseline accuracy and the post-deletion residual under DEL-ON. Under FULL, prompt style produces a wide performance spread (token F1 between $\sim 17\%$ and $\sim 57\%$ in Figure 5, exact match between $\sim 33\%$ and $\sim 46\%$ in Figure 6), reflecting how easily each prompt style elicits a structured lookup. Under DEL-ON, that spread collapses: token F1 compresses into a 7–10% band, and the FULL-to-DEL-ON drop is dominated by the FULL baseline rather

than by any property of the prompt itself. Continuations lose the most because they had the most to lose; few-shot prompts lose the least for the same reason. The original paper’s TOFU and FactScore evaluations each use a fixed prompt template, so this kind of decomposition is not available there (Zhao et al., 2025; Maini et al., 2024; Min et al., 2023). Our six prompt families show that, once the supporting evidence is removed, prompt formulation has limited independent control over how much of a deleted fact survives. What survives is determined almost entirely by the retrieval graph.

6.5. Error Analysis

The decomposition in Sections 6.1 and 6.2 predicts three qualitatively distinct ways a deleted fact can still be answered correctly: parametric leakage, near-neighbor retrieval reconstruction on *Noise*, and confusable-neighbor drift on *Collision*. Table 3 in Appendix A.3 shows one inference-time trace for each, with the *Noise* and *Collision* rows drawn from the corresponding custom-database files and the parametric-leakage row drawn illustratively from a cell that empirically contains leakage cases. The traces make the population-level rates concrete: in the *Noise* case, alias-closure deletion removes the canonical triplet for the gold fact, but a paraphrastic decoy with a different subject phrasing survives and routes retrieval to the same object, so DEL-ON returns the gold answer with no gold-equivalent candidate in the trace; in the *Collision* case, deletion is complete and no surviving candidate routes to the gold object, but a near-miss triplet sharing the subject pulls retrieval onto a confusable neighbor and DEL-ON returns a plausible but incorrect answer; and in the parametric-leakage case, retrieval returns unknown under DEL-OFF yet the model still produces the gold answer from its parameters, an outcome that the $\hat{L} = 0.11\%$ headline rate predicts to be rare and that appears, in the cells where it occurs, to concentrate on high-frequency entities.

7. Conclusion

The audit’s contribution is not a single number but a decomposition: a retrieval-disabled control isolates parametric recall from retrieval-side correctness, and the retrieval trace separates correctness anchored in gold-equivalent evidence from correctness reconstituted by near-neighbor candidates. Across our grid, the parametric channel is essentially empty in our audit: the LMLM checkpoint almost never returns a deleted answer when retrieval is disabled, and the residual that survives is, in our experiments, primarily a property of the retrieval graph rather than of the model. Within the scope of our audit, LMLM deletion within the model appears to be a clean operation; outside the model, deletion is only as complete as the alias-and-paraphrase closure used to

construct it, and our *Noise* variant produces a 13.6% residual at zero parametric leakage precisely by exposing this gap. Prompt formulation does not appear to move this number independently: once the supporting evidence has been removed, what predominantly determines how much of a deleted fact survives is the topology of the retrieval graph rather than the question’s surface form.

8. Limitations and Broader Impact

Our audit has several scope limitations. All experiments use a single LMLM checkpoint, the publicly released 382M-parameter LLaMA2-style model, with retrieval threshold 0.6, fuzzy-match top-1 fallback, and greedy decoding, so findings may not generalize across model sizes, retrieval architectures, similarity thresholds, or decoding strategies (Zhao & contributors, 2025). The custom databases used to stress-test deletion are also small and topology-controlled, on the order of 100–240 facts per (domain, variant) cell, which is what makes paired comparison possible but is several orders of magnitude smaller than the 54.6M-triplet released LMLM database (Zhao et al., 2025); the variant-level findings therefore describe how deletion behaves when scale is held fixed and topology is varied, rather than how scale itself interacts with topology. We restrict attention to entity-level atomic facts in English, evaluated under an automatic, alias-aware exact-match notion of correctness, so multi-hop or compositional reasoning, languages other than English, sampling-based decoding, and human-judged correctness all remain outside the present scope. Finally, our alias-closure deletion procedure removes only triplets that alias-match the target on subject, relation, and object simultaneously; the *Noise* variant is designed to expose what survives this closure, so its residual is a feature of the audit rather than a methodological shortcoming, but it does mean that low parametric leakage should not be read as a guarantee that any user-facing notion of “deletion” fully removes access to sensitive information.

For applications such as data-deletion compliance or factual editing, the practical lesson from the audit is that the unlearning boundary in this class of LMLM is drawn primarily by the database administrator rather than by the model. Sound governance therefore requires a closure procedure that extends beyond entity aliases to retrieval-graph paraphrases of the object.

9. Future Work

The most direct extension of this work is on the database side. Our results identify retrieval-graph topology as the dominant determinant of post-deletion residuals, which suggests building or preprocessing the database so that alias-closure deletion has fewer surviving routes. One concrete

direction is an extended deletion closure that, in addition to alias-equivalent triplets, also removes any triplet whose retained-candidate embedding falls within a similarity radius of the canonical entry, so that paraphrastic decoys are caught at deletion time. A second direction is canonicalization at write time, in which aliases and paraphrastic forms are stored as pointers into a single canonical record rather than as independent triplets. Both approaches are directly testable within our framework: re-running the audit on the modified database and measuring whether $R(f)$ and the retrieval artifact rate fall below their current ranges would tell us how much of the residual is recoverable through database design alone.

Several axes of the audit itself also remain open. We hold the retriever (ALL-MINILM-L6-v2) and the similarity threshold (0.6) fixed throughout, so sweeping the threshold and replacing the retriever with sparse, dense, or hybrid alternatives would clarify how much of the observed 13.6% *Noise* residual is a property of the embedding model rather than of the LMLM architecture. Scaling the audit to larger LMLM checkpoints and to the full 54.6M-triplet released database would test whether the per-fact zero parametric leakage we report holds at production scale. Extending the framework beyond entity-level atomic facts to multi-hop and compositional knowledge, and beyond automatic exact-match scoring to semantic or human-judged correctness, would close the gap between the audit and end-user notions of forgetting. Finally, the same intervention set can be applied to retrieval-augmented and edited parametric models more broadly, which would allow a head-to-head comparison of which architectural family achieves the cleanest deletion under matched retrieval pressure.

Acknowledgements

We thank Akshat Gupta (Ph.D. student, UC Berkeley) for ongoing research feedback and direction; Yilun Hua (Ph.D. student, Cornell University) for further research feedback and direction on the LMLM framework; and Marcel Roed (Ph.D. student, Stanford University) for early feedback on the project proposal. This work used computing resources provided by Berkeley Research Computing through the Compton Spectrometer and Imager (COSI) mission (NASA Small Explorers (SMEX) Program).

References

- Bourtoule, L., Chandrasekaran, V., Choquette-Choo, C. A., Jia, H., Travers, A., Zhang, B., Lie, D., and Papernot, N. Machine unlearning. In *Proceedings of the 42nd IEEE Symposium on Security and Privacy*, 2021. URL <https://arxiv.org/abs/1912.03817>.
- Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-

- Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, Ú., Oprea, A., and Raffel, C. Extracting training data from large language models. In *USENIX Security Symposium*, 2021. URL <https://arxiv.org/abs/2012.07805>.
- Guu, K., Lee, K., Tung, Z., Pasupat, P., and Chang, M.-W. Realm: Retrieval-augmented language model pre-training. In *International Conference on Machine Learning*, 2020. URL <https://arxiv.org/abs/2002.08909>.
- Karpukhin, V., Oğuz, B., Min, S., Lewis, P., Wu, L., Edunov, S., Chen, D., and Yih, W.-t. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020. URL <https://arxiv.org/abs/2004.04906>.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., and Kiela, D. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Advances in Neural Information Processing Systems*, 2020. URL <https://arxiv.org/abs/2005.11401>.
- Lizzo, T. and Heck, L. Unlearning in llms: Methods, evaluation, and open challenges, 2026. URL <https://arxiv.org/abs/2601.13264>.
- Maini, P., Feng, Z., Schwarzschild, A., Lipton, Z. C., and Kolter, J. Z. Tofu: A task of fictitious unlearning for llms, 2024. URL <https://arxiv.org/abs/2401.06121>.
- Mallen, A., Asai, A., Zhong, V., Das, R., Khashabi, D., and Hajishirzi, H. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023. URL <https://arxiv.org/abs/2212.10511>.
- Meng, K., Bau, D., Andonian, A., and Belinkov, Y. Locating and editing factual associations in gpt. In *Advances in Neural Information Processing Systems*, 2022. URL <https://arxiv.org/abs/2202.05262>.
- Meng, K., Sharma, A. S., Andonian, A., Belinkov, Y., and Bau, D. Mass-editing memory in a transformer. In *International Conference on Learning Representations*, 2023. URL <https://arxiv.org/abs/2210.07229>.
- Min, S., Krishna, K., Lyu, X., Lewis, M., Yih, W.-t., Koh, P. W., Iyyer, M., Zettlemoyer, L., and Hajishirzi, H. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2023. URL <https://arxiv.org/abs/2305.14251>.
- Yao, Y., Wang, P., Tian, B., Cheng, S., Li, Z., Deng, S., Chen, H., and Zhang, N. Editing large language models: Problems, methods, and opportunities. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023. URL <https://arxiv.org/abs/2305.13172>.
- Zhao, L. and contributors. Lmlm. <https://github.com/kilian-group/LMLM>, 2025. Commit 9be34b7.
- Zhao, L., Zalouk, S., Belardi, C. K., Lovelace, J., Zhou, J. P., Noonan, R. T., Go, D., Weinberger, K. Q., Artzi, Y., and Sun, J. J. Pre-training limited memory language models with internal and external knowledge, 2025. URL <https://arxiv.org/abs/2505.15962>.

A. Architecture

This appendix details the two design axes that define the prompt \times database grid used throughout the audit: the six prompt formulations applied to every target fact, and the four custom database topologies constructed to stress-test alias-closure deletion. Together these specifications fix the inputs to the $6 \times 3 \times 13$ evaluation grid summarized in Section 4.

A.1. Prompt formulations

Each target fact is presented to the LLM under six prompt formulations, designed to vary surface form while holding the underlying fact constant. The formulations span direct question-answering, paraphrased rewordings, contextually framed prompts, cloze-style completions, free-form continuations, and few-shot demonstrations. Their purpose in the audit is to separate the effect of question phrasing on FULL baseline accuracy from any independent effect on the post-deletion residual; in our experiments, prompt formulation moves the FULL baseline considerably more than it moves the DEL-ON residual (Section 6.4). Table 1 illustrates each formulation using a single fact drawn from the released LLM database.

Table 1: Example prompt formulations for the fact “Geri Halliwell – Famous For – Spice Girls” across the six prompt sets in the released LLM database.

Prompt Set	Example Prompt
Direct	What is Geri Halliwell famous for?
Paraphrased	Can you tell me what Geri Halliwell is famous for?
Contextual	Context: I am compiling a concise factual profile for Geri Halliwell. Answer with a short factual phrase. Question: What is Geri Halliwell famous for?
Cloze	Complete the sentence with the missing fact: Geri Halliwell is famous for _____.
Continuations	Tell me about Geri Halliwell. Geri Halliwell is famous for
Fewshot	Answer the final question with a short factual phrase. Question: Where was Ada Lovelace born? Answer: London Question: In what year was Pride and Prejudice published? Answer: 1813 Question: What is Geri Halliwell famous for? Answer:

A.2. Custom database variants

Alongside the released LLM database, we constructed twelve custom databases spanning three domains (countries, politicians, sports) and four topological variants (Base, Alias, Noise, Collision). The variants are intended to isolate distinct mechanisms by which a fact could remain recoverable after canonical removal: *Base* provides a clean baseline with no alternative routes; *Alias* probes whether alias-closure deletion catches every surface realization of the subject and relation; *Noise* probes whether retrieval can reconstitute the deleted answer through paraphrastic decoys that point to the same object; and *Collision* probes whether retrieval drifts onto a near-neighbor sharing the subject but routing to a different object. In our audit, the post-deletion residual tracks this topology ordering, with the largest residual concentrated in the *Noise* variant (Section 6.3). Table 2 illustrates the four variants using a single politician-domain example.

Table 2: Example of the four custom dataset variants used in our experiments, illustrated with a single politician-domain example.

Variant	Description	Example
Base	Clean, straightforward facts with little ambiguity.	Who is the 47th U.S. president? → Donald Trump
Alias	Same underlying fact, but expressed with alternate names or phrasings.	Who is the current POTUS? → Donald Trump Who is the president of the USA? → Donald Trump
Noise	Many different facts point to the same answer, so that answer becomes very common and tempting for retrieval.	Who is the 45th U.S. president? → Donald Trump Who is the 47th U.S. president? → Donald Trump Which New York property magnate became U.S. president? → Donald Trump
Collision	Near-miss facts that are very similar but should produce different answers, so retrieval can confuse neighbors.	Who is the 46th U.S. president? → Joe Biden Who is the 47th U.S. president? → Donald Trump

A.3. Qualitative examples of post-deletion failure modes

Table 3 shows one representative inference-time trace for each of the three failure modes identified by the decomposition. For each example we report the gold fact, the variant the trace was drawn from, the canonical triplet that was removed by alias-closure deletion, the highest-similarity retrieval candidate retained at inference time, and the model’s DEL-ON (or DEL-OFF, for the parametric-leakage case) output. The deleted triplets and the surviving retrieval candidates in the *Noise* and *Collision* rows are taken verbatim from the custom databases under `data/custom_databases/countries/`. The parametric-leakage row is drawn from the `direct_questions × Base` (sports) cell, which contains 2 leakage cases out of 100 paired facts. The specific triplet shown is illustrative of the kind of high-frequency sports fact this cell contains rather than a verbatim trace from the audit logs.

 Table 3: Representative inference-time traces for the three post-deletion failure modes. *Noise* and *Collision* rows reproduce real triplets from the corresponding custom-database files; the parametric-leakage row is drawn from a cell that empirically contains leakage cases (`direct_questions × Base` (sports), 2 of 100), but the specific triplet shown is illustrative rather than a verbatim trace.

Failure mode	Variant	Deleted triplet	Surviving retrieval candidate	Model output
Near-neighbor reconstruction	<i>Noise</i> (countries)	<i>United States</i> → <i>Capital</i> → <i>Washington, D.C.</i>	<i>Government of United States</i> → <i>Seat of Government</i> → <i>Washington, D.C.</i>	DEL-ON: “Washington, D.C.” (gold; no gold-equivalent candidate in trace)
Confusable-neighbor drift	<i>Collision</i> (countries)	<i>United States</i> → <i>Capital</i> → <i>Washington, D.C.</i>	<i>United States</i> → <i>Largest City</i> → <i>New York City</i>	DEL-ON: “New York City” (incorrect; plausible neighbor)
Parametric leakage	<i>Base</i> (sports)	<i>Michael Jordan</i> → <i>Sport</i> → <i>Basketball</i>	None (retrieval returns unknown)	DEL-OFF: “Basketball” (illustrative of the 2 leakage cases observed in this cell)

B. Hyperparameters

This section lists the configuration constants that directly affect the LMLM’s outputs in our audit: the loaded checkpoint, the retrieval stack, and the decoding policy. Settings inherited from the upstream LMLM release (Zhao & contributors, 2025) are

noted as such in the description column. Experimental-grid axes (prompt formulations, database variants, intervention states) and evaluation plumbing are described separately in Appendix A and Section 4. Table 4 lists each performance-relevant hyperparameter, the variable name in source, a short description, and the value used.

Table 4: Performance-relevant hyperparameters used throughout the audit. Variable names refer to identifiers in `src/lmlm-audit/` unless otherwise specified.

Hyperparameter	Description	Value
<code>model_name</code>	LMLM checkpoint loaded by the audit’s model loader; LLaMA2-style decoder-only model pre-trained from scratch by the upstream authors on annotated Wikipedia.	<code>kilian-group/</code> <code>LMLM-llama2-382M</code>
<code>logit bias on db tokens</code>	Decoding-time logit bias applied to the four lookup tokens so that lookup calls are issued reliably; values inherited from upstream.	<code>< db_entity >: 4</code> <code>< db_relationship >: 2</code> <code>< db_return >: 2</code> <code>< db_end >: 2</code>
<code>embedding model</code>	Sentence-Transformer used to embed lookup queries and database triplets for top- <i>k</i> FAISS retrieval; inherited from the upstream top- <i>k</i> retriever.	<code>sentence-transformers/</code> <code>all-MiniLM-L6-v2</code>
<code>threshold</code>	Cosine similarity threshold used by the top- <i>k</i> retriever; candidates with score below this are dropped, and retrieval returns <code>unknown</code> if no candidate clears it.	0.6
<code>fallback_policy</code>	Behavior when no candidate clears <code>threshold</code> ; <code>top1_anyway</code> re-runs retrieval at <code>threshold - 1.0</code> and returns the highest-similarity candidate.	<code>top1_anyway</code>
<code>max_new_tokens</code>	Per-call cap on freshly generated answer tokens, used as the target answer length when computing the generation budget.	12
<code>generation budget</code>	Total token budget passed to <code>model.generate</code> , sized to leave slack for lookup markup before the retrieved value appears.	<code>max(32, p + 28)</code>
<code>repetition_penalty</code>	HuggingFace <code>generate</code> repetition penalty applied during decoding.	1.2
<code>do_sample</code>	Whether to sample tokens during decoding; greedy decoding is used throughout the audit.	False
<code>eos_token_id</code>	Stop tokens that terminate generation early so the audit can intercept retrieval at the correct point.	<code>< db_return ></code> <code>tokenizer.eos_token_id</code> <code>< end_of_text ></code>

C. LMLM Audit Identities

This section collects the small set of equations that are specific to the three-state LMLM audit. Each identity links the per-fact indicators $L(f)$ and $R(f)$ from Section 3 and the per-fact retrieval artifact indicator $A(f)$ from Proposition 6.1 to the paired contingency that the McNemar test in Appendix D.1 consumes. Throughout, \mathcal{F} denotes the set of paired facts (those evaluated under both DEL-ON and DEL-OFF), $n = |\mathcal{F}|$, and the paired contingency is

$$\begin{aligned}
 a &= |\{f \in \mathcal{F} : Y(f, \text{DEL-ON}) = o \wedge Y(f, \text{DEL-OFF}) = o\}|, \\
 b &= |\{f \in \mathcal{F} : Y(f, \text{DEL-ON}) = o \wedge Y(f, \text{DEL-OFF}) \neq o\}|, \\
 c &= |\{f \in \mathcal{F} : Y(f, \text{DEL-ON}) \neq o \wedge Y(f, \text{DEL-OFF}) = o\}|, \\
 d &= |\{f \in \mathcal{F} : Y(f, \text{DEL-ON}) \neq o \wedge Y(f, \text{DEL-OFF}) \neq o\}|.
 \end{aligned}$$

The intervention-condition variable c from Section 3 does not appear elsewhere in this section, so we reuse the symbol for the McNemar-style off-diagonal count without ambiguity.

C.1. Decomposition of DEL-ON correctness

Proposition C.1 (DEL-ON correctness decomposition). *For every fact $f \in \mathcal{F}$,*

$$\mathbb{I}[Y(f, \text{DEL-ON}) = o] = L(f) + R(f) - \mathbb{I}[Y(f, \text{DEL-ON}) \neq o \wedge Y(f, \text{DEL-OFF}) = o]. \tag{1}$$

Aggregating equation (1) over \mathcal{F} and writing $p_{\text{DEL-ON}}$ for the empirical DEL-ON exact-match rate gives

$$p_{\text{DEL-ON}} = \hat{L} + \hat{R} - \frac{c}{n}. \quad (2)$$

Proof. The four indicators $\mathbb{I}[a\text{-cell}]$, $\mathbb{I}[b\text{-cell}]$, $\mathbb{I}[c\text{-cell}]$, $\mathbb{I}[d\text{-cell}]$ partition the joint event space of $(Y(f, \text{DEL-ON}), Y(f, \text{DEL-OFF}))$, so for every f exactly one of them is 1. By construction

$$\begin{aligned} \mathbb{I}[Y(f, \text{DEL-ON}) = o] &= \mathbb{I}[a\text{-cell}] + \mathbb{I}[b\text{-cell}], \\ L(f) = \mathbb{I}[Y(f, \text{DEL-OFF}) = o] &= \mathbb{I}[a\text{-cell}] + \mathbb{I}[c\text{-cell}], \\ R(f) &= \mathbb{I}[b\text{-cell}]. \end{aligned}$$

Adding the first two relations and subtracting $\mathbb{I}[c\text{-cell}]$ cancels the duplicated $\mathbb{I}[a\text{-cell}]$ and yields equation (1). Summing equation (1) over \mathcal{F} and dividing by n gives equation (2), since $\sum_{f \in \mathcal{F}} \mathbb{I}[c\text{-cell}] = c$. \square

Proposition C.1 explains why the $\hat{L} + \hat{R}$ stack reported in Figure 3 reads almost identically to the DEL-ON exact-match column of Figure 6: the gap is exactly c/n , which never exceeds $4/600 \approx 0.7\%$ in any of our variants. The proposition also pins down what $\hat{L} + \hat{R}$ would equal in a hypothetical setting where retrieval frequently hurts correctness; in that setting c/n would no longer be negligible and the stack and the DEL-ON bar would diverge by exactly that amount.

C.2. Recovering the paired contingency from aggregated rates

Proposition C.2 (Contingency recovery). *Suppose every fact $f \in \mathcal{F}$ is evaluated under both DEL-ON and DEL-OFF. Let $p_{\text{DEL-ON}}$ and $p_{\text{DEL-OFF}} = \hat{L}$ denote the empirical DEL-ON and DEL-OFF exact-match rates and \hat{R} the empirical retrieval-mediated correctness rate. Then the paired contingency is*

$$b = n\hat{R}, \quad a = np_{\text{DEL-ON}} - n\hat{R}, \quad c = np_{\text{DEL-OFF}} - a, \quad d = n - a - b - c. \quad (3)$$

Proof. By definition $\hat{R} = (1/n) \sum_{f \in \mathcal{F}} R(f) = b/n$, so $b = n\hat{R}$. The aggregate DEL-ON exact-match rate is $p_{\text{DEL-ON}} = (a + b)/n$, giving $a = np_{\text{DEL-ON}} - b = np_{\text{DEL-ON}} - n\hat{R}$. The aggregate DEL-OFF exact-match rate is $p_{\text{DEL-OFF}} = \hat{L} = (a + c)/n$, giving $c = np_{\text{DEL-OFF}} - a$. Since $\{a, b, c, d\}$ partition \mathcal{F} , $a + b + c + d = n$ and so $d = n - a - b - c$. \square

Proposition C.2 is what allows the aggregated cross-state and per-state metrics produced by the audit to be re-expanded into the (a, b, c, d) table that the McNemar statistic in Appendix D.1 consumes, without re-running the audit on a per-fact basis. The recovery is exact under the same coverage condition that makes the McNemar test well defined, namely that every paired fact has been evaluated under both intervention states.

D. Additional Results

D.1. Paired McNemar test for DEL-ON vs. DEL-OFF

Figure 7 reports the discordant-pair counts that drive the paired McNemar test for DEL-ON against DEL-OFF, broken out by database variant. Both states share the same alias-closure deletion of the target fact; the only difference is that DEL-ON leaves retrieval enabled while DEL-OFF disables it. The count b measures facts that retrieval rescues after the canonical entry is removed, and c measures the reverse transition in which disabling retrieval somehow recovers an otherwise-incorrect answer. The concordant counts a and d are absorbed into the marginal n and do not enter the test.

Auditing Forgetting in Limited Memory Language Models

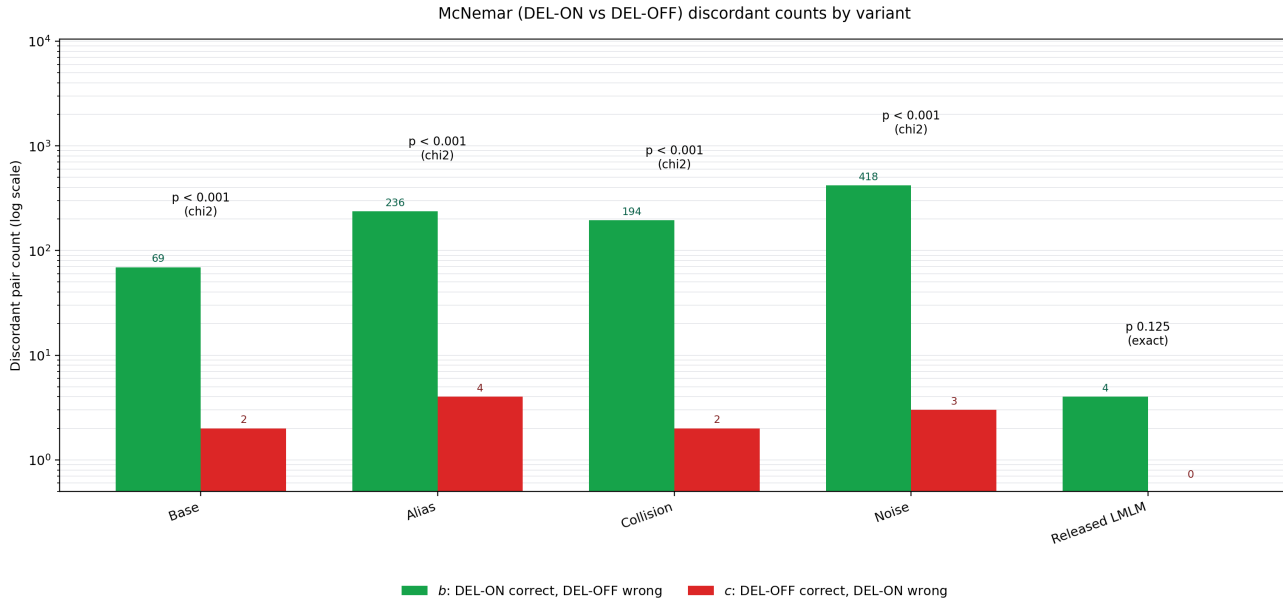


Figure 7. Discordant-pair counts driving the paired McNemar test for DEL-ON against DEL-OFF, broken out by database variant. Bar height is on a log scale so the much smaller c counts remain legible alongside b . Annotations above each variant give the p -value of the recommended test (χ^2 with continuity correction when $b + c > 25$, exact binomial otherwise).

Across the four custom variants the asymmetry is unambiguous: b ranges from 69 on *Base* to 418 on *Noise* while c never exceeds 4, and all four reject the null at $p < 0.001$ with χ^2 statistics between 61 and 407. The Released LMLM column points the same way ($b = 4, c = 0$ on $n = 600$), but the discordant counts are too small for the exact binomial test to reject at $\alpha = 0.05$ ($p = 0.125$); we read this as a power limit rather than a contradiction, since the released-database evaluation has both fewer paired facts and a lower base rate of exact-match correctness.

Taken together with the headline rates in the main paper, the McNemar evidence reinforces the claim that retrieval-mediated correctness is a real paired effect rather than a coincidence at the population level. The relation $b \gg c$ holds inside every custom variant, which leaves little room for a reverse hypothesis in which gains from retrieval are merely noise of equal magnitude flipping in both directions.

E. Numeric Results

Table 5: Standard metrics across databases, prompt sets, and inference-time database states. Values are reported as proportions.

Database	Prompt Set	FULL			DEL-ON			DEL-OFF		
		P	R	$F1$	P	R	$F1$	P	R	$F1$
Released LMLM	Direct Questions	0.474	0.477	0.475	0.022	0.022	0.021	0.015	0.024	0.016
Released LMLM	Paraphrased Questions	0.391	0.386	0.388	0.016	0.020	0.015	0.012	0.008	0.008
Released LMLM	Contextual Questions	0.594	0.597	0.594	0.032	0.039	0.031	0.008	0.031	0.011
Released LMLM	Cloze Prompts	0.336	0.331	0.332	0.028	0.031	0.027	0.006	0.014	0.007
Released LMLM	Continuations	0.815	0.819	0.816	0.026	0.034	0.026	0.028	0.016	0.013
Released LMLM	Fewshot	0.255	0.251	0.251	0.021	0.019	0.019	0.015	0.047	0.020
Countries Base	Direct Questions	0.580	0.580	0.580	0.098	0.098	0.098	0.012	0.030	0.017
Countries Base	Paraphrased Questions	0.545	0.545	0.545	0.083	0.083	0.083	0.011	0.030	0.016
Countries Base	Contextual Questions	0.670	0.670	0.670	0.058	0.058	0.058	0.012	0.060	0.017

Continued on next page

Auditing Forgetting in Limited Memory Language Models

Table 5 continued from previous page

Database	Prompt Set	FULL			DEL-ON			DEL-OFF		
		<i>P</i>	<i>R</i>	<i>F1</i>	<i>P</i>	<i>R</i>	<i>F1</i>	<i>P</i>	<i>R</i>	<i>F1</i>
Countries Base	Cloze Prompts	0.311	0.315	0.311	0.022	0.027	0.023	0.004	0.040	0.007
Countries Base	Continuations	0.840	0.840	0.840	0.028	0.028	0.028	0.042	0.180	0.067
Countries Base	Fewshot	0.230	0.230	0.230	0.090	0.090	0.090	0.001	0.020	0.003
Countries Alias	Direct Questions	0.757	0.757	0.757	0.217	0.217	0.217	0.006	0.043	0.011
Countries Alias	Paraphrased Questions	0.566	0.566	0.566	0.152	0.149	0.150	0.008	0.056	0.014
Countries Alias	Contextual Questions	0.603	0.608	0.604	0.124	0.129	0.126	0.004	0.055	0.008
Countries Alias	Cloze Prompts	0.511	0.516	0.512	0.113	0.115	0.112	0.004	0.028	0.007
Countries Alias	Continuations	0.858	0.858	0.858	0.202	0.202	0.202	0.014	0.058	0.020
Countries Alias	Fewshot	0.253	0.253	0.253	0.092	0.092	0.092	0.010	0.045	0.015
Countries Noise	Direct Questions	0.631	0.631	0.631	0.467	0.467	0.467	0.007	0.022	0.011
Countries Noise	Paraphrased Questions	0.569	0.569	0.569	0.383	0.383	0.383	0.006	0.022	0.009
Countries Noise	Contextual Questions	0.670	0.670	0.670	0.389	0.389	0.389	0.009	0.050	0.014
Countries Noise	Cloze Prompts	0.438	0.443	0.439	0.327	0.331	0.328	0.006	0.043	0.010
Countries Noise	Continuations	0.878	0.878	0.878	0.354	0.354	0.354	0.025	0.128	0.041
Countries Noise	Fewshot	0.309	0.309	0.309	0.241	0.241	0.241	0.013	0.041	0.014
Countries Collision	Direct Questions	0.668	0.671	0.669	0.192	0.194	0.192	0.012	0.050	0.019
Countries Collision	Paraphrased Questions	0.545	0.561	0.549	0.179	0.198	0.184	0.013	0.061	0.021
Countries Collision	Contextual Questions	0.700	0.707	0.701	0.131	0.137	0.131	0.013	0.081	0.019
Countries Collision	Cloze Prompts	0.436	0.439	0.436	0.127	0.130	0.127	0.006	0.071	0.012
Countries Collision	Continuations	0.879	0.879	0.879	0.174	0.173	0.173	0.033	0.152	0.052
Countries Collision	Fewshot	0.182	0.182	0.182	0.086	0.086	0.086	0.004	0.033	0.007
Politicians Base	Direct Questions	0.504	0.504	0.504	0.033	0.033	0.033	0.000	0.000	0.000
Politicians Base	Paraphrased Questions	0.496	0.493	0.493	0.121	0.121	0.119	0.005	0.025	0.008
Politicians Base	Contextual Questions	0.475	0.467	0.470	0.079	0.078	0.078	0.003	0.025	0.005
Politicians Base	Cloze Prompts	0.283	0.279	0.281	0.042	0.042	0.040	0.010	0.060	0.016
Politicians Base	Continuations	0.683	0.681	0.682	0.033	0.033	0.033	0.006	0.017	0.008
Politicians Base	Fewshot	0.304	0.304	0.302	0.104	0.101	0.100	0.007	0.037	0.011
Politicians Alias	Direct Questions	0.387	0.385	0.386	0.053	0.054	0.054	0.003	0.015	0.004
Politicians Alias	Paraphrased Questions	0.300	0.300	0.300	0.050	0.052	0.051	0.002	0.013	0.003
Politicians Alias	Contextual Questions	0.565	0.567	0.565	0.075	0.078	0.075	0.002	0.022	0.004
Politicians Alias	Cloze Prompts	0.256	0.254	0.255	0.054	0.048	0.050	0.001	0.013	0.002
Politicians Alias	Continuations	0.729	0.726	0.727	0.087	0.083	0.084	0.002	0.008	0.003
Politicians Alias	Fewshot	0.200	0.199	0.199	0.079	0.076	0.077	0.013	0.028	0.015
Politicians Noise	Direct Questions	0.000	0.000	0.000	0.000	0.000	0.000	0.015	0.022	0.016
Politicians Noise	Paraphrased Questions	0.000	0.003	0.001	0.000	0.003	0.001	0.008	0.036	0.013
Politicians Noise	Contextual Questions	0.000	0.000	0.000	0.000	0.000	0.000	0.009	0.031	0.010
Politicians Noise	Cloze Prompts	0.000	0.000	0.000	0.000	0.000	0.000	0.010	0.068	0.016
Politicians Noise	Continuations	0.000	0.000	0.000	0.000	0.000	0.000	0.004	0.011	0.006
Politicians Noise	Fewshot	0.000	0.003	0.000	0.000	0.003	0.000	0.005	0.028	0.008
Politicians Collision	Direct Questions	0.444	0.443	0.443	0.059	0.058	0.059	0.000	0.000	0.000
Politicians Collision	Paraphrased Questions	0.381	0.379	0.379	0.119	0.119	0.118	0.004	0.019	0.006
Politicians Collision	Contextual Questions	0.503	0.493	0.497	0.103	0.101	0.102	0.002	0.019	0.004
Politicians Collision	Cloze Prompts	0.256	0.256	0.256	0.025	0.025	0.025	0.008	0.045	0.012
Politicians Collision	Continuations	0.600	0.598	0.599	0.075	0.075	0.075	0.005	0.019	0.008
Politicians Collision	Fewshot	0.278	0.276	0.275	0.075	0.072	0.071	0.005	0.028	0.008

Continued on next page

Auditing Forgetting in Limited Memory Language Models

Table 5 continued from previous page

Database	Prompt Set	FULL			DEL-ON			DEL-OFF		
		<i>P</i>	<i>R</i>	<i>F1</i>	<i>P</i>	<i>R</i>	<i>F1</i>	<i>P</i>	<i>R</i>	<i>F1</i>
Sports Base	Direct Questions	0.010	0.010	0.010	0.010	0.010	0.010	0.047	0.095	0.058
Sports Base	Paraphrased Questions	0.010	0.010	0.010	0.010	0.010	0.010	0.005	0.020	0.008
Sports Base	Contextual Questions	0.010	0.010	0.010	0.010	0.010	0.010	0.004	0.048	0.006
Sports Base	Cloze Prompts	0.010	0.010	0.010	0.010	0.010	0.010	0.005	0.033	0.008
Sports Base	Continuations	0.010	0.010	0.010	0.010	0.010	0.010	0.013	0.055	0.020
Sports Base	Fewshot	0.000	0.005	0.001	0.000	0.005	0.001	0.002	0.018	0.003
Sports Alias	Direct Questions	0.000	0.000	0.000	0.000	0.000	0.000	0.009	0.024	0.013
Sports Alias	Paraphrased Questions	0.000	0.000	0.000	0.000	0.000	0.000	0.008	0.028	0.012
Sports Alias	Contextual Questions	0.000	0.000	0.000	0.000	0.000	0.000	0.004	0.028	0.006
Sports Alias	Cloze Prompts	0.000	0.000	0.000	0.000	0.000	0.000	0.006	0.025	0.008
Sports Alias	Continuations	0.000	0.000	0.000	0.000	0.000	0.000	0.008	0.033	0.012
Sports Alias	Fewshot	0.000	0.005	0.001	0.000	0.005	0.001	0.008	0.018	0.009
Sports Noise	Direct Questions	0.475	0.466	0.468	0.072	0.061	0.063	0.041	0.092	0.053
Sports Noise	Paraphrased Questions	0.405	0.406	0.406	0.054	0.050	0.050	0.006	0.028	0.009
Sports Noise	Contextual Questions	0.569	0.569	0.569	0.044	0.041	0.042	0.006	0.055	0.010
Sports Noise	Cloze Prompts	0.284	0.274	0.276	0.053	0.043	0.045	0.004	0.027	0.007
Sports Noise	Continuations	0.716	0.707	0.709	0.075	0.060	0.063	0.009	0.041	0.014
Sports Noise	Fewshot	0.196	0.195	0.193	0.058	0.058	0.055	0.006	0.041	0.010
Sports Collision	Direct Questions	0.479	0.481	0.480	0.036	0.032	0.032	0.034	0.074	0.043
Sports Collision	Paraphrased Questions	0.404	0.405	0.404	0.042	0.043	0.042	0.003	0.013	0.005
Sports Collision	Contextual Questions	0.544	0.540	0.541	0.045	0.040	0.039	0.002	0.032	0.005
Sports Collision	Cloze Prompts	0.278	0.279	0.277	0.050	0.049	0.048	0.004	0.025	0.007
Sports Collision	Continuations	0.696	0.696	0.696	0.054	0.046	0.049	0.009	0.038	0.013
Sports Collision	Fewshot	0.198	0.207	0.200	0.068	0.078	0.070	0.011	0.028	0.013

Table 6: Cross-state audit metrics across databases and prompt sets. Values are reported as proportions.

Database	Prompt Set	Paired Count	$L(f)$	$R(f)$	Retrieval Artifact Rate
Released LMLM	Direct Questions	100	0.000	0.010	0.010
Released LMLM	Paraphrased Questions	100	0.000	0.000	0.000
Released LMLM	Contextual Questions	100	0.000	0.010	0.010
Released LMLM	Cloze Prompts	100	0.000	0.010	0.010
Released LMLM	Continuations	100	0.000	0.000	0.000
Released LMLM	Fewshot	100	0.000	0.010	0.010
Countries Base	Direct Questions	100	0.000	0.080	0.080
Countries Base	Paraphrased Questions	100	0.000	0.060	0.060
Countries Base	Contextual Questions	100	0.000	0.040	0.040
Countries Base	Cloze Prompts	100	0.000	0.010	0.010
Countries Base	Continuations	100	0.000	0.010	0.010
Countries Base	Fewshot	100	0.000	0.090	0.090
Countries Alias	Direct Questions	200	0.000	0.200	0.200
Countries Alias	Paraphrased Questions	200	0.000	0.135	0.135
Countries Alias	Contextual Questions	200	0.000	0.110	0.110
Countries Alias	Cloze Prompts	200	0.000	0.090	0.090

Continued on next page

Auditing Forgetting in Limited Memory Language Models

Table 6 continued from previous page

Database	Prompt Set	Paired Count	$L(f)$	$R(f)$	Retrieval Artifact Rate
Countries Alias	Continuations	200	0.005	0.185	0.185
Countries Alias	Fewshot	200	0.000	0.090	0.090
Countries Noise	Direct Questions	180	0.000	0.456	0.456
Countries Noise	Paraphrased Questions	180	0.000	0.367	0.367
Countries Noise	Contextual Questions	180	0.000	0.378	0.378
Countries Noise	Cloze Prompts	180	0.000	0.322	0.322
Countries Noise	Continuations	180	0.000	0.339	0.339
Countries Noise	Fewshot	180	0.011	0.233	0.233
Countries Collision	Direct Questions	140	0.000	0.164	0.164
Countries Collision	Paraphrased Questions	140	0.000	0.143	0.143
Countries Collision	Contextual Questions	140	0.000	0.107	0.107
Countries Collision	Cloze Prompts	140	0.000	0.114	0.114
Countries Collision	Continuations	140	0.000	0.150	0.150
Countries Collision	Fewshot	140	0.000	0.079	0.079
Politicians Base	Direct Questions	120	0.000	0.025	0.025
Politicians Base	Paraphrased Questions	120	0.000	0.100	0.100
Politicians Base	Contextual Questions	120	0.000	0.050	0.050
Politicians Base	Cloze Prompts	120	0.000	0.017	0.017
Politicians Base	Continuations	120	0.000	0.033	0.033
Politicians Base	Fewshot	120	0.000	0.067	0.067
Politicians Alias	Direct Questions	240	0.000	0.046	0.046
Politicians Alias	Paraphrased Questions	240	0.000	0.042	0.042
Politicians Alias	Contextual Questions	240	0.000	0.067	0.067
Politicians Alias	Cloze Prompts	240	0.000	0.033	0.033
Politicians Alias	Continuations	240	0.000	0.075	0.075
Politicians Alias	Fewshot	240	0.008	0.046	0.046
Politicians Noise	Direct Questions	180	0.006	0.000	0.000
Politicians Noise	Paraphrased Questions	180	0.000	0.000	0.000
Politicians Noise	Contextual Questions	180	0.000	0.000	0.000
Politicians Noise	Cloze Prompts	180	0.000	0.000	0.000
Politicians Noise	Continuations	180	0.000	0.000	0.000
Politicians Noise	Fewshot	180	0.000	0.000	0.000
Politicians Collision	Direct Questions	160	0.000	0.050	0.050
Politicians Collision	Paraphrased Questions	160	0.000	0.100	0.100
Politicians Collision	Contextual Questions	160	0.000	0.081	0.081
Politicians Collision	Cloze Prompts	160	0.000	0.019	0.019
Politicians Collision	Continuations	160	0.000	0.075	0.075
Politicians Collision	Fewshot	160	0.000	0.044	0.044
Sports Base	Direct Questions	100	0.020	0.010	0.010
Sports Base	Paraphrased Questions	100	0.000	0.010	0.010
Sports Base	Contextual Questions	100	0.000	0.010	0.010
Sports Base	Cloze Prompts	100	0.000	0.010	0.010
Sports Base	Continuations	100	0.000	0.010	0.010
Sports Base	Fewshot	100	0.000	0.000	0.000
Sports Alias	Direct Questions	200	0.000	0.000	0.000
Sports Alias	Paraphrased Questions	200	0.000	0.000	0.000
Sports Alias	Contextual Questions	200	0.000	0.000	0.000

Continued on next page

Auditing Forgetting in Limited Memory Language Models

Table 6 continued from previous page

Database	Prompt Set	Paired Count	$L(f)$	$R(f)$	Retrieval Artifact Rate
Sports Alias	Cloze Prompts	200	0.000	0.000	0.000
Sports Alias	Continuations	200	0.000	0.000	0.000
Sports Alias	Fewshot	200	0.005	0.000	0.000
Sports Noise	Direct Questions	160	0.013	0.044	0.056
Sports Noise	Paraphrased Questions	160	0.000	0.037	0.037
Sports Noise	Contextual Questions	160	0.000	0.037	0.037
Sports Noise	Cloze Prompts	160	0.000	0.037	0.037
Sports Noise	Continuations	160	0.000	0.050	0.050
Sports Noise	Fewshot	160	0.000	0.050	0.050
Sports Collision	Direct Questions	158	0.013	0.013	0.019
Sports Collision	Paraphrased Questions	158	0.000	0.032	0.032
Sports Collision	Contextual Questions	158	0.000	0.025	0.025
Sports Collision	Cloze Prompts	158	0.000	0.032	0.032
Sports Collision	Continuations	158	0.000	0.032	0.032
Sports Collision	Fewshot	158	0.006	0.057	0.057